



Advances in the Measurement of the Electron Electric Dipole Moment

Citation

Meisenhelder, Cole. 2023. Advances in the Measurement of the Electron Electric Dipole Moment. Doctoral dissertation, Harvard University Graduate School of Arts and Sciences.

Permanent link

<https://nrs.harvard.edu/URN-3:HUL.INSTREPOS:37375705>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

HARVARD UNIVERSITY
Graduate School of Arts and Sciences



DISSERTATION ACCEPTANCE CERTIFICATE

The undersigned, appointed by the
Department of Physics
have examined a dissertation entitled

Advances in the Measurement of the Electron Electric Dipole Moment

presented by Cole Mitchell Meisenhelder

candidate for the degree of Doctor of Philosophy and hereby
certify that it is worthy of acceptance.

Signature _____ *Gerald Gabrielse*

Typed name: Professor Gerald Gabrielse, Chair

Signature _____ *Masahiro Morii*

Typed name: Professor Masahiro Morii

Signature _____ *R. Walsworth*

Typed name: Professor Ronald Walsworth (Univ. of Maryland)

Date: May 3, 2023

Advances in the Measurement of the Electron Electric Dipole Moment

A DISSERTATION PRESENTED
BY
COLE MITCHELL MEISENHOLDER
TO
THE DEPARTMENT OF PHYSICS

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN THE SUBJECT OF
PHYSICS

HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS
MAY 2023

©2023 – COLE MITCHELL MEISENHELDER
ALL RIGHTS RESERVED.

Advances in the Measurement of the Electron Electric Dipole Moment

ABSTRACT

Searches for the electron electric dipole moment (eEDM) are powerful tests of physics beyond the standard model, as they probe for new evidence of time-reversal violating interactions. In 2018 our ACME collaboration reported a measurement of the eEDM which set a new limit of $|d_e| < 1.1 \times 10^{-29} \text{ e} \cdot \text{cm}$, which improved on the sensitivity of the previous best measurement, also set by the ACME collaboration, by an order of magnitude. With the development of a new ACME III apparatus, we intend to perform a new measurement with an order of magnitude improved sensitivity over the ACME II result. The new experiment has should both upgrade the statistical uncertainty of the measurement and upgrade the suppression of the effects of known sources of systematic error. To upgrade our statistical sensitivity we have implemented upgrades to our molecular beam flux, increased our precession time, and increased to our detection efficiency. To suppress known sources of systematic error and noise below our projected new sensitivity we have developed new magnetic shields, improved our experiment timing controls, and developed new methods to suppress laser polarization gradients in our system.

Contents

TITLE	i
COPYRIGHT	iii
ABSTRACT	iv
TABLE OF CONTENTS	viii
LIST OF FIGURES	xi
LIST OF TABLES	xii
I INTRODUCTION TO THE ELECTRON ELECTRIC DIPOLE MOMENT	1
1.1 EDMs, Discrete Symmetries and Physics Beyond the Standard Model	2
1.1.1 Baryon Asymmetry and the Sakharov Conditions	3
1.1.2 EDMs and T Violation	4
1.1.3 Energy Scales	5
1.2 Atomic and Molecular Measurements	7

1.2.1	Evading Schiff's Theorem	7
1.2.2	Experimental Uncertainty	8
1.2.3	Alternative Interpretation	9
1.3	A Brief History of electron EDM Measurements	9
1.4	This Thesis	12
ACKNOWLEDGEMENTS		1
2	THE ACME MEASUREMENT OVERVIEW	14
2.1	The ThO Molecule	15
2.1.1	Electric and Magnetic Field Properties of the H State	16
2.1.2	The Ω Doublet Structure	16
2.1.3	H State Lifetime	18
2.2	A Spin Precession measurement	18
2.3	Measurement Apparatus	20
2.3.1	The Buffer Gas Beam Source	23
2.3.2	Rotational Cooling	24
2.3.3	Electrostatic Focusing	25
2.3.4	State Preparation	28
2.3.5	Free Precession and Applied Fields	30
2.3.6	State Readout and Detection	31
2.4	ACME Switches	31
2.4.1	Block Level Switches and The Minimum EDM Measurement	32
2.4.2	Other Switch Timescales	35
3	THE ACME II MEASUREMENT	37

3.1	Upgrades over ACME I	38
3.1.1	STIRAP	40
3.1.2	Beam Geometry	42
3.1.3	Improved Collection Efficiency and Detection Wavelength	44
3.1.4	Electric Field Plates	45
3.1.5	Upgrades that Were not Implemented	48
3.2	Laser Systems	50
3.2.1	The Ablation Laser	50
3.2.2	ACME II Rotational Cooling Lasers	51
3.2.3	STIRAP Lasers	54
3.2.4	The Refinement and Readout Laser	56
3.3	Electric and Magnetic Field Control	58
3.3.1	Magnetometry	58
3.3.2	Magnetic Shielding	61
3.3.3	Magnetic Field Coils	64
3.3.4	Electric Field Measurements and \mathcal{E}^{nr}	65
3.4	Systematic Errors	66
3.4.1	A Short Aside About $\mathcal{N}\mathcal{E}$ -Correlated Laser Detunings	69
3.4.2	Magnetic Field Gradient Effects	70
3.4.3	STIRAP Laser Imperfections and Finite Refinement Laser Power	74
3.4.4	AC Stark Shift Effects and Polarization Gradients	76
3.4.5	Correlated Contrast Effects	78
3.4.6	Effects of an $\tilde{\mathcal{E}}$ Correlated Magnetic Field	79
3.5	The ACME II Result	80
3.5.1	Excess Noise from Timing Jitter	81

3.5.2	Excess Noise at Large Magnetic Fields	83
3.5.3	An Order of Magnitude Improved Limit on d_e	84
4	ACME III DEVELOPMENT	85
4.1	ACME III Statistical Upgrade Pathway	86
4.1.1	Increased Precession Time	88
4.1.2	The Electrostatic Lens and Improved Rotational Cooling	90
4.1.3	SiPM Detectors and Improved Collection Optics	98
4.1.4	Excess Noise Suppression	103
4.1.5	Rapid Beam Source Target Changing	116
4.1.6	Proposed Upgrades That Were Not Implemented	120
4.2	Systematic Error Suppression	122
4.2.1	Magnetic Field Systematics and Noise	123
4.2.2	Polarization Gradients	125
4.3	General Upgrades and Experiment Robustness	134
5	EXTENDING THE ACME III INTERACTION REGION	137
5.1	A New Measurement of the H State Lifetime	138
5.1.1	Methods and Apparatus	138
5.1.2	Systematic Uncertainties	145
5.2	Magnetic Shielding and Magnetometry	148
5.2.1	Shielding Theory	149
5.2.2	Shield Design Criteria	151
5.2.3	The ACME III Magnetic Shields	157
5.2.4	Degaussing	165
5.2.5	Magnetometry	167

5.3	The ACME III Interaction Region Vacuum Chamber	170
5.3.1	Design Considerations	170
5.3.2	Vacuum Levels	177
5.3.3	Mechanical Supports and Vibration Analysis	179
5.3.4	SiPM Mounting and Stray Light Blocking	181
6	CONCLUSION	186
APPENDIX A LASER POWER ESTIMATES FOR $X - A - H$ STIRAP		189
A.1	Previous measurements	190
A.2	Method 1: Direct Scaling of Rabi Frequencies	192
A.2.1	Power Results	193
A.3	Method 2: Comparisons to Optical Pumping	194
A.4	Conclusion	195
BIBLIOGRAPHY		214

Listing of figures

1.1	Discrete Symmetries of the Electron	6
1.2	History of Electron EDM Measurements	10
2.1	Relevant Electronic Energy Levels of ThO and Laser Transitions	15
2.2	H state level diagram	17
2.3	ACME II Experiment Schematic	21
2.4	ACME III Experiment Schematic.	22
2.5	ACME II Rotational Cooling Scheme.	25
2.6	X-C-Q STIRAP Level Diagram	26
2.7	ACME III Electrostatic Lens Schematic and Electric Field Strength.	27
2.8	ACME II State Preparation	29
2.9	ACME II Switching Structure.	33
3.1	ACME I State Preparation Scheme	41
3.2	The ACME II Magnetic Shields	62
3.3	A Model for the $\partial\mathcal{B}_z\partial z \times \partial\mathcal{E}^{\text{nr}}/\partial z$ Systematic	72

3.4	Statistics of the ACME II Measurement	82
4.1	EDM Sensitivity Gain as a Function of Precession Time	91
4.2	Simulations of the Electrostatic Lens Focusing	94
4.3	ACME III Rotational Cooling	97
4.4	Schematic of the SiPM Module	100
4.5	Comparison Between SiPMs and PMTs	102
4.6	Polarization Switching Timescales	105
4.7	Asymmetry Within a Polarization Cycle	106
4.8	Asymmetry Noise Dependence on Clock Syncing and $\Delta T_{\hat{X}-\hat{Y}}$	109
4.9	Asymmetry Noise Dependence on Sub-bin Selection	110
4.10	Timing Jitter Before Modifications	111
4.11	Fast Scale Timing Jitter After Clock Syncing	114
4.12	Timing Jitter After All Modifications	116
4.13	ACME II Signal Levels	117
4.14	ACME III Load Lock System for Rapid Target Changes	119
4.15	Large Supported Window	131
5.1	H State Lifetime Measurement Schematic	139
5.2	H State Lifetime Measurement Fluorescence Intensity Ratios versus Transit Time	144
5.3	H State Lifetime Measurement Pulse Shapes and Arrival Times	148
5.4	The ACME III Magnetic Shields	158
5.5	ACME III Magnetic Shield Panels	159
5.6	Magnetic Shield Frames	160
5.7	Magnetic Shield Design Simulation Results	162
5.8	The ACME Table	164

5.9	The Interaction Region Vacuum Chamber	172
5.10	Interaction Region Chamber Vibration Simulations	180
5.11	The Light Pipe Mounting System	183

List of Tables

3.1	ACME II Demonstrated Signal Gains	40
3.2	ACME II Combined Error Budget for $\omega^{\mathcal{N}\mathcal{E}}$	69
4.1	ACME III Demonstrated Statistical Gains	87
5.1	Lifetime Measurement Excitation Region Positions	142
5.2	Lifetime Measurement Combined Error Budget	145

Acknowledgments

IT SHOULD COME AS NO SURPRISE THAT THIS THESIS AND THE WORK THAT WENT INTO IT DID NOT HAPPEN IN A VACUUM. Every stage of this work was supported by the communities I am a part of, regardless of their relationships to the ACME experiment or to physics. Without these people, I would not have come this far.

I would first like to thank my advisor Jerry Gabrielse, for his guidance and enthusiasm for the work we've done over these last eight years. I would like to thank John Doyle and Dave DeMille for their advice and insights over the course of my time in the ACME collaboration. I would also like to thank my committee members Ron Walsworth and Masahiro Morii for their guidance, both scientific and professional.

I would like to thank all of the members of the ACME collaboration, past and present, as my colleagues and collaborators. Cristian Panda and Zack Lasner served as my mentors when I first joined ACME, and taught me many of the lab techniques I used in this work. Over the years I have been grateful to work with such wonderful colleagues as Daniel Ang, Xing Wu, Jonathan Haefner, Vitaly Wirthl, Zhen Han, Peiran Hu, Siyuan Liu, Bingjie Hao, Daniel Lascar, Takahiko Masuda,

Ayami Hiramoto, and Xing Fan. I have also been fortunate to get to work with the many members of the Gabrielse lab.

The work in this thesis was made possible by the incredible support we receive from staff members at both Harvard and Northwestern. Thank you to Laura Nevins, Lisa Cacciabaudo, Pattee McGarry, Jan Ragusa, Silke Exner, Stan Cotreau, Jim MacArthur, and Salomon Rodriguez.

I was fortunate to have the support from the broader community during my time at Harvard. Within the physics department I am grateful for the friendships I built within the incoming class from 2015, in particular my roommates Kristine Rezai, Laurel Anderson, and Sasha Brownsberger (and of course Holly and Campbell). I also want to thank the Harvard Graduate Students Union, which played an important role in the trajectory of my graduate career. I would especially like to thank the members of our first bargaining committee, who it was an honor to work with, for all of the highs and lows we shared, and the support you all showed me.

During my time in undergrad I found a lifelong group of friends who have shaped a great deal of who I am as a person, and as a physicist. You all have shown me such support and generosity over the last decade, and I will always be grateful. Thank you to Jeremy Boissevain, Gabriel Samach, Teddy Amdur, David Yan, Wei Li, Liz Deitz, Jack Hoover, Julia Cline, and Adrienne Favis.

Finally, I would like to thank my parents and my sister, for whom I cannot give enough thanks. Without your support, I never would have thought this was possible, and would not be where I am today.

But the Universe doesn't say what you want it to say

The Universe says, just what it says,

when it says what it says.

The Ernies, *Here and Now*

1

Introduction to the Electron Electric Dipole Moment

THE STANDARD MODEL OF PARTICLE PHYSICS IS ONE OF THE GREAT TRIUMPHS [1] OF MODERN PHYSICS. With the measurement of the Higgs boson in 2012 [2], all elementary particles predicted by the Standard Model (SM) have now been observed. Furthermore, the SM has been tested

at the the part per trillion level with precision measurements of the electron magnetic moment [3] in our research group and quantum field theory predictions [4]. However it is also true that the SM is unable to answer fundamental questions about the universe that remain. These questions, particularly why our universe is essentially entirely made up of matter (baryon asymmetry) [5], what the nature of dark matter is, and what the nature of dark energy is. This has motivated the exploration of theories of physics beyond the Standard Model (BSM) which may be able to answer these questions.

This thesis focuses on measurements of the electron electric dipole moment (EDM) made as part of the ACME collaboration, which serve as experimental tests of BSM physics. For nearly a decade the ACME I [6, 7] and ACME II [8] results, the latter described in this thesis, were the most precise measurements of the electron EDM. In section 1.1 we will consider the problem of baryon asymmetry, which motivates searches for particles that break certain discrete symmetries, and show that a non-zero electron EDM must inherently break these symmetries. As a result many BSM theories predict values of the electron EDM that are orders of magnitude above the SM prediction [9]. In section 1.2 we will describe how atoms and molecules can be used to measure the electron EDM, which have been for the last 60 years have been the most sensitive . Finally, in section 1.3 we will provide an overview of the historical progression of electron EDM measurements and look towards future measurements.

1.1 EDMs, DISCRETE SYMMETRIES AND PHYSICS BEYOND THE STANDARD MODEL

In order to understand the motivations behind the search for an electron EDM, we must first consider the three discrete symmetries, charge conjugation (C), parity reversal (P), and time reversal (T) which along with their combinations underpin the SM.

- **Charge Conjugation (C)** where we replace every particle with its antiparticle.

- **Parity Reversal (P)** where we mirror all spacial coordinates with the transformation

$$(x, y, z) \rightarrow (-x, -y, -z).$$

- **Time Reversal (T)** where we invert the direction of time with the transformation $t \rightarrow -t$.

While we describe these transformations as symmetries, they are not exact symmetries of the SM. In fact, violations of all three of these symmetries have been observed. The first evidence came in 1957 with the observation of P violation in the β -decay of ^{60}Co in the Wu experiment [10]. With the discovery of P violation, searches for symmetry violations extended to CP violation, where CP symmetry is the product of simultaneous C and P transformations. This symmetry was also found to be violated, first in measurements of the decay of the K_2^0 meson [11], and then in the B meson [12], and D meson [13]. Since then further evidence of CP violation [14] and T violation [15], which we will shortly show is equivalent to CP violation, have been observed.

Although neither C, P, or T is a symmetry of the SM, the combined symmetry CPT is an exact invariance of the SM. This CPT theorem requires Lorentz invariance, and CPT violation would have profound consequences [16]. While theories that include CPT violation exist [17], tests for CPT violation have returned no evidence of this violation across a wide range of sectors including Penning trap experiments comparing matter and antimatter in our own lab [18, 19]. So long as the CPT theorem holds, CP violations must also violate T in order to maintain CPT symmetry¹.

1.1.1 BARYON ASYMMETRY AND THE SAKHAROV CONDITIONS

For reasons we do not currently understand, our universe is made up essentially entirely of matter, with only small amounts of anti-matter. While we should be grateful that we can avoid annihilations, this observed baryon asymmetry is a deep question for which the SM does not have answers.

¹This is true for other combinations of C,P, and T, but for our purposes we are concerned with CP and T as we will see shortly.

This asymmetry has been well studied and can be characterized by the measured ratio of baryons to photons in the universe $\eta = 6.143 \pm 0.190 \times 10^{-10}$ from cosmological observations [5].

The most promising explanation for the observed baryon asymmetry was a solution proposed in 1967 by Sakharov[20]. This solution which is commonly known as the ‘‘Sakharov conditions’’, when assuming CPT symmetry relies on three conditions that must be met.

1. Baryon number must not be conserved.
2. Both C and CP symmetries must be violated
3. Interactions must occur out of thermal equilibrium, as occurred early in the Big Bang.

All of these conditions can be met within the SM, as the SM allows for a small amount of CP violation through the Cabibbo-Kobayashi-Maskawa (CKM) matrix [21, 22] in the quark sector. However, the amount of CP violation in the CKM matrix seems much too small to account for the observed baryon asymmetry [23, 24], so some other CP violating BSM mechanism is needed.

1.1.2 EDMs AND T VIOLATION

In 1950 Purcell and Ramsey [25] first proposed searches for T violating physics through the measurement of permanent electric dipole moments of nuclei and elementary particles as fundamentally T violating. Since then permanent EDMs have been measured for the neutron [26, 27, 28], the muon [29], the nuclei of atoms such as mercury [30, 31, 32] or radium [33], and the electron which is the subject of this work.

We now consider the electron EDM, which for a fundamental particle like the electron must be an inherent property of the particle, and not a spacial charge separation like in the classical case. However, the Hamiltonian of the electron EDM, as in the classical case, can be written as

$$\mathcal{H} = -\vec{d}_e \cdot \vec{\mathcal{E}}. \tag{1.1}$$

From the Wigner-Eckart theorem it must also be true that the vector \vec{d}_e must lie parallel to the electron spin vector \vec{S} . This can be seen intuitively as any misalignment between \vec{d}_e and \vec{S} would introduce an additional internal degree of freedom to the electron. This would suggest that another quantum number would be necessary to describe electrons, and from the Pauli exclusion principle more than two electrons would be able to occupy the same atomic orbital, which has not been observed.

With this picture of the electron EDM, we now consider how the electron behaves under both T and P transformations as shown in figure 1.1. Under a T transformation, the direction of \vec{d}_e does not change², while as an angular momentum the spin \vec{S} reverses to give $-\vec{S}$. Conversely, under a P transformation \vec{d}_e reverses direction, while the spin \vec{S} does not. As the relative sign between \vec{d}_e and \vec{S} change under these transformations, the particles before and after the transformation are distinguishable. This tells us that an electron EDM would violate both T and P. Assuming CPT invariance, an observed T violation would indicate a new source of CP violation.

1.1.3 ENERGY SCALES

In attempting to resolve unanswered questions from the SM, nearly all BSM theories introduce new interactions that violate T symmetry and that predict non-zero values of the electron EDM [34, 35, 36, 37]. These theories often predict values of the electron EDM near the current best limits $\sim 10^{-30}$ e · cm, which is many orders of magnitude higher than the SM prediction for the electron EDM [9]. The SM does predict a non-zero electron EDM, as described by the CKM matrix, the CP violation arises only in four-loop level calculations, with estimated values ranging from $\sim 10^{-35}$ – 10^{-38} e · cm [9, 36, 38]. This makes relatively small-size³ precision measurement experiments of the electron EDM powerful tools for testing BSM theories, as the current best measurements are

²This is intuitively seen in the classical picture of a spacial charge distribution which is time independent.

³As compared to typical particle physics experiments




Figure 1.1: Discrete Symmetries of the Electron. If an electron EDM does exist then it must be true that $\vec{d}_e \parallel \vec{S}$. If we apply a T transformation the EDM remains invariant, but the spin reverses. If instead we apply a P transformation the spin remains invariant but the EDM reverses. The original electron is distinguishable from the result of both the P and T transformations so a non-zero electron EDM must violate T and P symmetry.

completely free of the SM background. While the next generation of measurements are likely to remain free of the SM background, proposed new measurement techniques may approach the SM estimates in the near future [39].

We can also estimate the energy scales that current measurements of the electron EDM probe in the search for new physics. In a typical BSM theory we can estimate the rest mass Λ of new particles that produce an EDM in an n -loop interaction as [23, 16, 40]

$$\Lambda^2 \sim \kappa \left(\frac{\alpha}{2\pi} \right)^n \left(\frac{em_e c^2}{d_e} \right) \sin(\phi_T) (\hbar c), \quad (1.2)$$

where κ is a model dependent constant expected to be between $\sim 0.1 - 1$, α is the fine structure constant, m_e is the electron mass, and ϕ_T is a T violating phase. The latter is typically assumed to be of order $\sin(\phi_T) \sim 1$ because there is no compelling argument which requires ϕ_T to be small. For the ACME II result of $|d_e| < 1.1 \times 10^{-29} \text{ e} \cdot \text{cm}$, at the one-loop this probes new physics at the

level of $\Lambda_{1\text{-loop}} \approx 30$ TeV level, while at the two-loop level it probes physics at the $\Lambda_{2\text{-loop}} \approx 3$ TeV level [38, 41, 23]. These energies are higher than are probed at the Large Hadron Collider. More exact treatments of BSM models and the energy scales probed by electron EDM measurements can be found in [42, 34, 37].

1.2 ATOMIC AND MOLECULAR MEASUREMENTS

Attempts to measure the electron EDM with a free electron in a lab field would quickly find the electron accelerating away from the apparatus, a problem that is neatly solved with bound electrons. The original proposal by Purcell and Ramsey in 1950 [25], the first measurements in cesium atoms [43], and essentially all measurements of the electron EDM have thus been performed using electrons in polarized atoms or molecules. In these atoms and molecules it is possible to achieve effective electric fields such as the $\mathcal{E}_{\text{eff}} \approx 80$ GV/cm in the ThO molecule used for the ACME experiment [44, 45], which are many orders of magnitude higher than can be attained in a laboratory.

1.2.1 EVADING SCHIFF'S THEOREM

An initial consideration of this method would suggest that there can be no net electric field experienced by the bound electron, as such a field would simply rearrange the electron distribution. This principle, known as Schiff's theorem [46], only holds true in the non-relativistic limit. For heavy paramagnetic atoms and molecules with high Z this non-relativistic limit breaks down [47], and it is possible to measure an electron EDM. This is caused by a relativistic correction to the energy of the EDM in the applied field which can be understood as dipole moment of the electron (which has units of length) experiencing a length contraction near the nucleus [48]. We can express this correction as

$$U_{\text{EDM}} = - \left(\vec{d}_e \cdot \vec{\mathcal{E}} - \frac{\gamma}{1 + \gamma} \vec{\beta} \cdot \vec{d}_e \vec{\beta} \cdot \vec{\mathcal{E}} \right) \quad (1.3)$$

where $\vec{\beta} = \vec{v}/c$, $\gamma = 1/\sqrt{1 - \beta^2}$ is the Lorentz factor, and $\vec{\mathcal{E}}$ is the electric field. By Schiff's theorem the first $\vec{d}_e \cdot \vec{\mathcal{E}}$ term averages to zero, while the relativistic term does not. We can then group these relativistic factors into an effective electric field $\vec{\mathcal{E}}_{\text{eff}}$, which is dependent on the atom or molecule and its polarization, with larger values for higher Z and easily polarized molecules. This allows us to describe the EDM energy shift as

$$\langle U_{\text{EDM}} \rangle = -\vec{d}_e \cdot \vec{\mathcal{E}}_{\text{eff}}. \quad (1.4)$$

1.2.2 EXPERIMENTAL UNCERTAINTY

We can perform a measurement of the electron electric dipole moment in an atom or molecule by allowing the electron spin to precess in polarized atom or molecule in an external applied field. A more concrete picture of this method is given in section 2.2. We can understand the uncertainty of these measurements, where the spin has been allowed to precess for a time τ , through the uncertainty principle. This tells us that for a single measurement the uncertainty in the measured energy U_{EDM} must be

$$\delta U_{\text{EDM}} \approx \frac{\hbar}{2\tau}. \quad (1.5)$$

If we combine this with equation 1.4 we can see that the uncertainty in the measured value of d_e is given by

$$\delta d_e \approx \frac{\hbar}{2\tau \mathcal{E}_{\text{eff}}}. \quad (1.6)$$

If we then repeat this measurement N to improve our uncertainty, we should expect the uncertainty to obey Poisson statistics and scale as $1/\sqrt{N}$. This gives us the “shot-noise limit” for the uncertainty in an EDM measurement that is limited only by statistical noise, which we express as

$$\delta d_e = \frac{\hbar}{2\tau \mathcal{E}_{\text{eff}} \sqrt{N}} = \frac{\hbar}{2\tau \mathcal{E}_{\text{eff}} \sqrt{NT}}. \quad (1.7)$$

Where in the above equation \dot{N} is the measurement rate, limited by the photon detection rate in the ACME experiments, and T is the averaging time of the experiment. This uncertainty figure of merit is critical for EDM measurements and great effort is spent on minimizing this uncertainty, where τ and \mathcal{E}_{eff} , and \dot{N} are often limited by the atomic or molecular species and their properties.

1.2.3 ALTERNATIVE INTERPRETATION

It is worth noting that while we describe the quantity these atomic and molecular experiments measure as the electron EDM, there is an alternative interpretation. In general it is more accurate to state that the energy shift measured in these experiments is

$$U = -d_e \mathcal{E}_{\text{eff}} + W_S C_S \quad (1.8)$$

where C_S is a dimensionless electron-nucleon coupling parameter (that is also T violating) and W_S is a species-specific constant[45, 44, 9]. When we calculate the EDM value we are actually assuming that $C_S = 0$, but we could just as equally interpret the measured value with $d_e = 0$ instead.

If we were to do so for the ACME II measurement, this would result in $|C_S| < 7.3 \times 10^{-10}$ (90% confidence level)[8]. However, for many BSM models, the expected value of C_S is typically a smaller effect. If we make a similar energy scale argument as done earlier this result probes even higher energy scales, on the order of $\sim 10^5$ TeV [37].

1.3 A BRIEF HISTORY OF ELECTRON EDM MEASUREMENTS

If we look at the history of EDM measurements over the last 65 years incredible progress has been made. The limit on the electron EDM has been reduced by 14 orders of magnitude! On average this progress, which can be seen in figure 1.2, is equivalent to an order of magnitude every ≈ 4.5 years, which demonstrates an impressive sustained rate of advancement. While early measurements relied




Figure 1.2: History of Electron EDM Measurements. Bounds on the electron EDM placed by measurements over the last 65 years with the measurement system indicated by the legend. Early measurements were based on measurements of the Lamb shift [49], and measurements of the electron g factor [50, 51]. Following the first measurement in atomic cesium [43], measurements shifted to atomic systems such as Cs [52, 53], Rb [54], Xe [55], Hg [30], and Tl [56, 57, 58]. Measurements eventually moved to molecules and molecular ions including TlF [59], YbF [60, 61], PbO [62], HfF⁺ [63] and [64] (preprint), and ThO [6, 8]. This listing includes the best limits at the time and some novel systems.

on measurements of the Lamb shift [49] and the g factor of the electron [50, 51], once the shift to measurements in heavy paramagnetic atoms began in 1964 with the first cesium measurement [43] progress rapidly took off. Since then all competitive measurements of the electron EDM have been performed using this method in various atomic and molecular species.

With these experiments the primary pathway to improved sensitivity is often through improving the statistical sensitivity of the experiment, determined by equation 1.7, until the experiment becomes sensitive to systematic errors that can no longer be sufficiently suppressed. In the pursuit

of the best statistical sensitivity, experiments are designed to maximize \mathcal{E}_{eff} , τ , and \dot{N} ⁴. The values of \mathcal{E}_{eff} and τ are typically limited by properties of the atom or molecule, with τ often limited by the radiative lifetime of an excited state or the trapping time of a ground state. While N is often less directly limited by the species, this number often depends on properties of the experiment design as well as the properties of the atom or molecule. For these reasons, the selection of the atomic or molecular species plays an important role in determining the sensitivity of an experiment, and new possibilities in this space have led to significant leaps forward in the measurement of the electron EDM.

Cesium was not only the first atomic species to be used for a measurement of this kind, but it remained competitive with the best other measurements until the 1990s [43, 52, 53]. During this time other measurements used rubidium [54], xenon [55], mercury [30], and the first molecule to be used for an electron EDM measurement in TIF [59]. From 1990 until 2011, experiments with heavy thallium ($z = 81$), with a large effective electric field placed the tightest bounds on the electron EDM [56, 57, 58].

Since then the state of the art has shifted to using polar molecules, which due to their complex internal structure that makes them easily polarizable gives them an effective electric field that is typically $\sim 10^3 \times$ larger than was used in the thallium experiments. The first such molecule to set the best limit on the electron EDM was YbF in 2011, which improved on an earlier measurement with YbF to slightly improve upon the 2002 thallium measurement [60, 61]. The ACME I experiment, using ThO, which has the highest effective electric field used in an electron EDM measurement to date, improved upon this by an order of magnitude [6, 7]. In 2018 the ACME II measurement, which is included in this work, further improved upon this bound by nearly an order of magnitude, setting the current best published limit on the electron EDM of $|d_e| < 1.1 \cdot 10^{-29} \text{ e} \cdot \text{cm}$ [8].

⁴Maximizing T directly comes at the cost of the stability of an experiment and the patience of the operators.

Until the start of the EDM measurements at JILA using the molecular ion HfF^+ , all of the molecular measurements had been performed in molecular beams[63]. In these experiments the molecular ions are instead confined in an ion trap, which allows for much longer precession times than the molecular beam experiments, as a trade off with a much lower molecular density. This experiment measured a bound on the EDM in 2017 [63], which was close to the ACME II measurement, and have since reported a new measurement in a pre-print that reports what we expect to be the new best limit on the EDM of $|d_e| < 4.1. \times 10^{-30} \text{ e} \cdot \text{cm}$ [64].

Future measurements of the electron EDM are currently underway including the planned ACME III measurement, which is the subject of this work. With the ACME III measurement we hope to improve upon the ACME II limit by an order of magnitude once again. Similar next generations for the JILA experiment [65], using ThF^+ instead, and the YbF experiment are planned[66]. In the longer term new experiments with novel methods are planned using optically trapped polyatomic molecules [67], polar molecules embedded in rare-gas matrices [39], and super-heavy molecules with higher effective electric fields [68] have been proposed and are underway.

1.4 THIS THESIS

Over the course of this thesis we will describe the advances in the ACME measurement from the ACME II measurement through the development of the ACME III measurement. In chapter 2 we will provide a general picture of how the ACME experiment uses a beam of ThO molecules to perform a precession measurement and extract a measured value of d_e . In chapter 3 we will describe the specifics of the ACME II measurement, which concluded in 2018 and placed the current best published limit on the electron EDM of $|d_e| < 1.1. \times 10^{-29} \text{ e} \cdot \text{cm}$ [8]. We will especially focus on the systematic effects that we observed during the ACME II measurement as they will drive the discussion of systematic error suppression in the following two chapters. Chapters 4 and 5 will de-

scribe the development of the ACME III measurement which aims to measure the EDM with an order of magnitude improved sensitivity. with an overview of the ACME III statistical and systematic upgrades in chapter 4. Chapter 5 will focus on the development of the upgrades which increase the length of the interaction region to increase the precession time of our measurement, including the measurement which motivated these changes, and the development of new magnetic shields and interaction region systems.

The ordering of this thesis does not follow the typical ordering of steps for developing a new measurement, as I joined the ACME collaboration in the spring of 2017, when the ACME II experiment was already underway. I contributed to the ACME II measurement through upgrades to the laser systems, and collecting data for systematic studies and the final ACME II data set [8]. After the completion of ACME II, I worked to investigate and suppress the excess noise that limited the ACME II measurement [69]. Following that, my focus has been on the development of the ACME III measurement, which is described in the latter half of this thesis. In particular I worked primarily on the systems associated with the new interaction region, first on the new measurement of the radiative lifetime of the H state [70]. Following that work my focus went into the design and development of a new interaction region chamber to increase the precession time of our experiment, thereby increasing our statistical sensitivity, along with associated systems to increase our detection efficiency [71, 72, 73], and suppress polarization gradients in our windows that caused systematic effects in ACME II. In addition to this I made significant contributions to the design, construction, and testing of new magnetic shielding and magnetometry systems that will allow us to suppress both excess noise and known sources of systematic errors.

A universe in mourning, I don't know how to quit

It's gonna break my heart to see it blown to bits

Charly Bliss, *Blown to Bits*

2

The ACME Measurement Overview

THIS CHAPTER PROVIDES AN OVERVIEW OF THE GENERIC SCHEME OF THE ACME MEASUREMENT before the specifics of ACME II (completed) and ACME III (built) in later chapters. The first section describes the ThO molecule and the features that make the ACME measurement possible. The following sections describe the spin precession measurement that is performed, and how experimental state switches are used to extract the EDM value.




Figure 2.1: Relevant Electronic Energy Levels of ThO and Laser Transitions. Level diagram representing all of the electronic states of ThO that were used for ACME II, will be used for ACME III, or will be discussed in other sections of this work. The horizontal axis represents the magnitude of the quantum number Ω , while the vertical axis shows the hierarchy of states, not drawn to scale. The letter shows the label for each state, where X is the ground state, and H is the experiment state, while A , C , I , and Q are used in other auxiliary roles. Each straight arrow transition is color coded and matched to the label for the associated laser wavelength, while for the $I - X$ transition we only use spontaneous emission for detection.

2.1 THE ThO MOLECULE

We represent electronic states in ThO with the quantum numbers $|Y, J, M, \Omega\rangle$, where Y represents the electronic state, J represents the total angular momentum, M represents the projection of J onto the quantization axis \hat{z} , which in the ACME experiment is defined by applied lab fields, and Ω represents the projection of the electronic angular momentum onto the internuclear axis \hat{n} [7]. The relevant states for the ACME experiment are shown in figure 2.1. X is the ground state H is the EDM sensitive experiment state. The Q state is used for the electrostatic lens in ACME III. The C state is used for state preparation and rotational cooling. The I state is used for fluorescence read-out. The state that makes the ACME measurements possible is the $H^3\Delta_1$ excited electronic state of ThO. This unique state has properties, discussed in this section, that provide good sensitivity to the EDM while suppressing possible systematic effects.

2.1.1 ELECTRIC AND MAGNETIC FIELD PROPERTIES OF THE H STATE

The H state has two valence electrons in a $\sigma\delta$ configuration, where the σ orbital wavefunction amplitude has a large component centered around the heavy Th atom. This produces a relativistic effect as described in 1.2 that results in an internal effective electric field of $\mathcal{E}_{\text{eff}} \approx 78$ GV/cm [44, 45]. This large value of $\vec{\mathcal{E}}_{\text{eff}}$ provides a significant advantage to ThO as a candidate molecule over the current best eEDM competitor molecules and molecular ions in HfF+ with $\mathcal{E}_{\text{eff}} \approx 23$ GV/cm [64, 63], and YbF with $\mathcal{E}_{\text{eff}} = 14.5$ GV/cm [61, 60].

The H state also is remarkable decoupled from the magnetically field because its magnetic moment in the $|H, J = 1, \Omega = \pm 1, M\rangle$ states of $\mu_1 = g_1 \mu_B M = (-0.00440(5)) \mu_B M$, where μ_B is the Bohr magneton, and the g-factor is $g_1 = -.00440(5)$ [74, 75]. This tiny g-factor is a general feature of $^3\Delta_1$ states as the σ and δ orbitals nearly cancel each other's magnetic moments [76], and the primary source of the nonzero magnetic moment in the ThO H state is mixing with other electronic states [74]. This relative magnetic insensitivity significantly helps to suppress potential magnetic field related systematic errors.

2.1.2 THE Ω DOUBLET STRUCTURE

Nearly degenerate, metastable states with opposite parity are very rare in atoms. ThO and other molecules provide such states. The $|H, J = 1, \Omega = \pm 1\rangle$ manifolds form an Ω -doublet, where pairs of states of opposite parity are separated by a small splitting Δ_Ω , which is caused by the Coriolis effect in the molecule [77, 76]. While this is generically true for molecular states with $\Omega \neq 0$, it is of particular interest to the ACME experiment due to its behavior in the presence of a laboratory electric field $\vec{\mathcal{E}}$, which mixes the $M \neq 0$ states through the Stark shift $-\vec{D}_H \cdot \vec{\mathcal{E}}$, where \vec{D}_H is the electric dipole operator of the H state. By applying a laboratory electric field such that $|\vec{D}_H \cdot \vec{\mathcal{E}}| \gg \Delta_\Omega$ the molecule can be fully polarized in the lab frame. In the case of the $|H, J = 1, |\Omega| = 1\rangle$




Figure 2.2: H state level diagram. Level diagram showing the $|H, J = 1\rangle$ state in the presence of an applied electric field \mathcal{E} , which splits the $\tilde{\mathcal{N}}$ levels with the Stark shift of $2D_H\mathcal{E}$. For each of the four states with $\tilde{\mathcal{N}} = \pm 1$ the associated figure shows the orientation of the molecule and $\vec{\mathcal{E}}_{\text{eff}}$, as well as the orientation of the σ electron spin \vec{s} . The red arrows and dashed lines indicate how the Zeeman shift changes the energies of the $\tilde{\mathcal{N}} = \pm 1$ states, while the green arrows and dashed lines show how they shift in the presence of a non-zero electron EDM d_e . Figure adapted from [8].

state that we use for the ACME experiment[78], $\Delta_\Omega = 2\pi \times 360$ kHz [79], and $|D_H| = 2\pi \times 1$ MHz/V/cm [80], where these energies are given in terms of angular frequencies by setting $\hbar = 1$. This means that the molecule can be fully polarized ($\mathcal{P}_\mathcal{E} > 99\%$)[78] with laboratory fields as low as $|\vec{\mathcal{E}}| = 80$ V/cm, and we can easily apply a higher of $|\vec{\mathcal{E}}| = 140$ V/cm as an additional systematic check. Such small lab fields also make our experiment less sensitive to potential systematic effects related to leakage currents, motional fields, and geometric phases.

With the molecule fully polarized we define a new quantum number $\tilde{\mathcal{N}} \equiv \text{sgn}(\hat{n} \cdot \vec{\mathcal{E}}) = \pm 1$, where the internuclear axis \hat{n} is defined as pointing from the oxygen nucleus to the thorium nucleus. $\tilde{\mathcal{N}}$ determines the direction the molecule is oriented in the lab frame, and therefore determines the direction that $\vec{\mathcal{E}}_{\text{eff}}$ is oriented in the lab frame. As an applied lab field of $|\vec{\mathcal{E}}| = 80$ V/cm splits the $\tilde{\mathcal{N}} = \pm 1$ states by ~ 165 MHz, it is possible to spectroscopically select which $\tilde{\mathcal{N}}$ state is populated during state preparation, which along with the reversal of the laboratory field provides two independent ways to reverse the direction of $\vec{\mathcal{E}}_{\text{eff}}$, one with and one without changing the laboratory electric field. This ability to independently switch $\vec{\mathcal{E}}_{\text{eff}}$ is one of ACME's most powerful tools

for suppressing systematic effects as any potential systematic effect must correlate with not just one but both switches, as will be discussed further in section 2.4. [58, 81, 62]

2.1.3 H STATE LIFETIME

In the ACME experiment, particularly in ACME III, the H state radiative lifetime will limit the usable precession time of the experiment. The value we recently measured is $\tau_H = 4.2(5)$ ms [70]. The longer lifetime allows us to increase the precession time by a factor of 5 in ACME III versus ACME II. The details and implications of this measurement will be discussed in depth in section 5.1.

2.2 A SPIN PRECESSION MEASUREMENT

As discussed in section 1.2 a non-zero EDM in the presence of the effective electric field $\vec{\mathcal{E}}_{\text{eff}}$ will produce an energy shift given by $U = -\vec{d}_e \cdot \vec{\mathcal{E}}_{\text{eff}}$, where $\vec{d}_e = d_e \vec{s}/(\hbar/2)$, \vec{s} is the spin of the σ orbital electron, and \hbar is the reduced Planck constant. In the case of the ThO H state, \vec{s} is proportional to the molecule spin \vec{S} . This allows us to measure d_e as an energy splitting between the H states $|M = +1, \tilde{\mathcal{N}}\rangle$ and $|M = -1, \tilde{\mathcal{N}}\rangle$, which correspond to \vec{S} either aligned or anti-aligned with $\vec{\mathcal{E}}_{\text{eff}}$ as shown in figure 2.2. We measure this splitting through a precession measurement in the presence of applied electric and magnetic fields. First, we prepare an initial spin state in the xy plane with the applied fields $\vec{\mathcal{E}} = \mathcal{E} \hat{z}$ and $\vec{\mathcal{B}} = \mathcal{B}_z \hat{z}$ oriented along the \hat{z} axis. We can consider the specific case used in ACME where this initial state is given by

$$|\psi(t=0), \tilde{\mathcal{N}}\rangle = \frac{|M = +1, \tilde{\mathcal{N}}\rangle - |M = -1, \tilde{\mathcal{N}}\rangle}{\sqrt{2}}. \quad (2.1)$$

The fields $\vec{\mathcal{E}}_{\text{eff}}$ and $\mathcal{B}_z \hat{z}$ will exert a torque on the spin causing the spin to precess on average by an angle ϕ in the xy plane. After being allowed to precess for an amount of time τ this angle can be

represented as

$$\phi \approx \frac{-(\mu\tilde{\mathcal{B}}|\tilde{\mathcal{B}}_z| + \tilde{\mathcal{N}}\tilde{\mathcal{E}}d_e\mathcal{E}_{\text{eff}})\tau}{\hbar} \quad (2.2)$$

where as described in 2.1.1 $\mu = g_1\mu_B$, $\tilde{\mathcal{B}} \equiv \text{sgn}(\vec{\mathcal{B}} \cdot \hat{z})$, and $\tilde{\mathcal{E}} \equiv \text{sgn}(\vec{\mathcal{E}} \cdot \hat{z})$. It is worth noting at this point that $\tilde{\mathcal{N}}$, $\tilde{\mathcal{E}}$, $\tilde{\mathcal{B}}$ use the convention that will be used throughout this thesis to represent states of experimental switches as will be discussed in 2.4. After the initial prepared state shown in 2.1 has precessed by ϕ , we can represent the state as

$$|\psi(t = \tau), \tilde{\mathcal{N}}\rangle = \frac{e^{-i\phi} |M = +1, \tilde{\mathcal{N}}\rangle - e^{i\phi} |M = -1, \tilde{\mathcal{N}}\rangle}{\sqrt{2}} \quad (2.3)$$

We measure this phase ϕ by optically pumping from the H state to a higher lying state, and detecting the fluorescence as that state decays back to the ground state. In ACME I this was done by pumping to the *C* state[6], but both ACME II and ACME III use the *I* state[8], which decays back to the ground state with 512 nm light. The benefits of this change, which primarily resulted in more efficient photon detection, will be discussed in more detail in section 3.1. By switching between light that is linearly polarized in the basis \hat{X} and \hat{Y} , where \hat{X} is defined to be at an angle θ relative to the \hat{x} axis and \hat{Y} is orthogonal to \hat{X} , we are able to project the spin onto this basis. This switching is performed rapidly relative to a single molecular beam shot at 200 kHz in order to normalize out fluctuations between shot-to-shot variations in the number of molecules, which occur at the 20% level [82].

Once we have measured the fluorescence signal for each polarization, F_X and F_Y we compute an asymmetry [78]

$$\mathcal{A} \equiv \frac{F_X - F_Y}{F_X + F_Y} = C \cos(2(\phi - \theta)). \quad (2.4)$$

where \mathcal{C} is the measurement contrast, defined by

$$\mathcal{C} \equiv -\frac{1}{2} \frac{\partial \mathcal{A}}{\partial \theta} \approx \frac{1}{2} \frac{\partial \mathcal{A}}{\partial \phi} \quad (2.5)$$

In ACME II the contrast was $95\% \pm 2\%$ on average [8], which we measure by dithering the angle θ between the \hat{X} polarization and the \hat{x} axis. We set the values of $|\mathcal{B}_z|$ and θ such that $\phi - \theta \approx (\pi/4)(2n + 1)$ for an integer value of n in order to maximize the sensitivity of \mathcal{A} to changes in ϕ . We can then extract the measured phase, which we distinguish from the true phase ϕ by using the symbol Φ , which deviates from ϕ as a result of systematic effects. This phase can then be used to extract the EDM correlated phase and the size of the electron EDM as will be described in section 2.4.

2.3 MEASUREMENT APPARATUS

In this section, we consider the measurement apparatus for ACME II in figure 2.3 and ACME III in 2.4. While the most of the basic structure of the measurement is the same between ACME II and ACME III, there are significant differences, which will be discussed in more depth in chapters 4 and 5. For this section, we describe ACME as being composed of six stages 1) the buffer gas beam source; 2) rotational cooling; 3) electrostatic focusing (used only for ACME III); 4) state preparation; 5) free precession; and 6) state readout and detection.

At this time it is necessary to define the ACME coordinate system that will be used throughout this thesis. We define \hat{x} as the direction of travel of the molecular beam, \hat{y} points downwards, and \hat{z} is aligned with our applied electric field with it's sign determined by a right handed coordinate system. In ACME II, this meant that $+\hat{x}$ pointed North and $+\hat{z}$ pointed West. For ACME III, we moved the experiment from Harvard University to Northwestern University, where $+\hat{x}$ now points East, and $+\hat{z}$ now points North.




Figure 2.3: ACME II Experiment Schematic. A cold and slow ThO beam is first produced by ablation in the buffer gas beam source. We then use rotational cooling via optical pumping to enhance the population of the rotational ground state of X . Next the beam passes through fixed collimators to prevent molecules from reaching the ITO coated field plates. The molecular beam then enters a magnetically shielded region with five cylindrical μ -metal shields where fields \vec{E} and \vec{B} are applied by transparent field plates and field coils respectively. The molecules are coherently transferred from the ground state $|X, J = 0\rangle$ to the $|H, J = 1, \tilde{N}\rangle$ state using $X-C-H$ STIRAP with lasers sent through the experiment vertically. The refinement laser removes imperfections in the STIRAP prepared phase by optically pumping out the orthogonal spin state. The molecules are then allowed to precess freely in the applied fields for a distance $L = 22$ cm until they reach the readout laser. The detection laser projects out the spin state by rapidly switching between two polarization states \hat{X} , and \hat{Y} , where the resulting 512 nm fluorescence decay from $I \rightarrow X$ is collected by eight lens doublets. This light is focused into curved fused quartz light pipes where it travels to eight photomultiplier tubes for detection.




Figure 2.4: ACME III Experiment Schematic. A cold and slow ThO beam is first produced by ablation in the buffer gas beam source. We then use rotational cooling via optical pumping with a new overlapped method to enhance the population of the rotational ground state of X . We then use STIRAP to coherently transfer molecules to the $|Q, J = 2, M = 2, \Omega = -2\rangle$ state for lensing. Molecules in this state are then focused by an electrostatic lens consisting of cylindrical electrodes in a hexapole configuration that are charged to $\sim \pm 14$ kV. After the lens we reverse the STIRAP process to return the population back to the ground state. Next the beam passes through fixed collimators to prevent molecules from reaching the ITO coated field plates. The molecular beam then enters a magnetically shielded region with three layers of rectangular μ -metal shields where fields \vec{E} and \vec{B} are applied by transparent field plates and field coils respectively. The molecules are coherently transferred from the ground state $|X, J = 0\rangle$ to the $|H, J = 1, \tilde{N}\rangle$ state using $X - C - H$ STIRAP with lasers sent through the experiment vertically. The refinement laser removes imperfections in the STIRAP prepared phase by optically pumping out the orthogonal spin state. The molecules are then allowed to precess freely in the applied fields for an increased distance $L = 100$ cm (not shown to scale) until they reach the readout laser. The detection laser projects out the spin state by rapidly switching between two polarization states \hat{X} , and \hat{Y} , where the resulting 512 nm fluorescence decay from $I \rightarrow X$ is collected by eight lens doublets. This light is focused into straight fused quartz light pipes where it travels to eight silicon photomultipliers for detection. This figure was produced by Daniel Ang.

2.3.1 THE BUFFER GAS BEAM SOURCE

The ACME experiment begins with the production of a buffer gas cooled beam for which more details can be found in [83, 79]. A cylindrical pressed powder ceramic target made up of primarily ThO_2 , with diameter ~ 0.6 inches, is placed inside of a copper cell, which is cooled to ~ 16 K. In ACME II, the cell and the thermal shields were cooled using a single pulse tube refrigerator. In ACME III, a pair of pulse tube refrigerators will be used in a configuration which will provide more cooling power and that has been tested for multiple years on what was our test beam source, previously referred to locally as “Beam Box II”. This target is ablated by a pulse of focused light from a Nd:YAG laser¹ with pulse energy ~ 50 mJ, pulse duration of ~ 15 ns, and 50 Hz repetition rate. The ThO molecules produced in the ablation process are entrained in the Neon gas flow that runs through the cell at 40 sccm and which has thermalized with the cell. This neon flow thermalizes with the heavy ThO molecules, then exits the circular aperture of the cell creating a beam of neon, ThO, along with an unknown number of other neutral and ionic species which do not interact with the laser used in the experiment and can be neglected along with the neon.

To remove any ions we use an “ion sweeper” consisting of two charged plates that produce an electric field of 250 V/cm which deflects any ions from the beam. At this point the beam is extremely divergent with an angular full-width-half-maximum (FWHM) of $\sim 45^\circ$ [79], which is collimated by a “skimmer” aperture with a 6 mm opening. This beam divergence represents the largest loss in the ACME experiment, and is a significant motivation for both moving to the electrostatic lens that will be described in section 2.3.3 and shortening the distance between the cell and the lens as much as possible, as described in section 2.3.2. After the skimmer, the beam has a forward velocity of $v_{||} \approx 200$ m/s, however this value has been observed to vary at the 10% level from shot-to-shot. This process produces approximately 10×10^{10} molecules per pulse.

¹Litron Nano

2.3.2 ROTATIONAL COOLING

The molecules leaving the cell, with a temperature of ~ 4 K, have almost all of the population in the lowest four rotational levels of the electronic and vibrational ground state, $|X, J = 0, 1, 2, 3\rangle$. As the state preparation used in ACME II and III only transfers population from the $|X, J = 0\rangle$ state to the H state, we optically pump the population from $|X, J = 1\rangle$ and $|X, J = 2\rangle$ into $|X, J = 0\rangle$ in order to achieve a signal gain of approximately 2.5 [8]. While more signal could be achieved by also pumping out the $|X, J = 3\rangle$ state, in practice we choose not to pump out this state as it would require an additional laser for a significantly smaller gain.

In ACME II, this optical pumping used the transition from $X - C$ and was performed in two separate stages which are represented schematically in figure 2.5. At this point, we introduce the convention of representing states in the form of $|X, J = 0^+\rangle$, where the superscript indicates the parity of the state, which will be necessary for understanding the pumping scheme used. The first stage optical pumping drives the 690 nm transition $|X, J = 2^+\rangle \rightarrow |C, J = 1^-\rangle$, where the $|C, J = 1^-\rangle$ state can decay to either $|X, J = 2^+\rangle$ or $|X, J = 0^+\rangle$. To address all M sub-levels of the system and ensure that we repump molecules that decay back to $|X, J = 2^+\rangle$ we perform this step with 5-7 passes, where the polarization of the light on each pass is orthogonal to the last. Parity selection rules prevent us from being able to use the same scheme to drive the $|X, J = 1^-\rangle \rightarrow |C, J = 1^-\rangle$ transition, so instead an electric field of ~ 100 V/cm is used to weakly mix the opposite parity states $|C, J = 1^\pm\rangle$. In this applied field we use a similar multipass laser with alternating polarization to drive the transition $|X, J = 1^-\rangle \rightarrow |C, J = 1^{\text{Mixed}}\rangle$, which can then decay to the desired state $|X, J = 0^+\rangle$. For a more detailed discussion of this rotational cooling scheme see [84].

In ACME III we use a similar scheme as in ACME II, however the two stages are spatially overlapped in order to minimize the distance between the skimmer and the electrostatic lens aperture




Figure 2.5: ACME II Rotational Cooling Scheme. **a.** In the first step, with no applied field, we optically pump from $|X, J = 2^+\rangle \rightarrow |C, J = 1^-\rangle$, which decays into $|X, J = 0^+\rangle$. **b.** In the second step we apply an electric field to mix the parity state of the $|C, J = 1\rangle$ state so that we can drive the $|X, J = 1^-\rangle \rightarrow |C, J = 1^{Mixed}\rangle$ transition. From the $|C, J = 1^{Mixed}\rangle$ state population can decay to either $|X, J = 0^+\rangle$ or $|X, J = 2^+\rangle$, which limits the transfer efficiency in this step.

[85]. This is made possible by coupling both pumping lasers into a single fiber, and rapidly switching the applied electric field on and off. Further details of this new method will be discussed in section 4.1.

2.3.3 ELECTROSTATIC FOCUSING

In order to take better advantage of the highly divergent molecular beam, ACME III uses a new electrostatic focusing scheme. This work [85] (led by Xing Wu), and further details on this system and its benefits, can be found in section 4.1. In this scheme the large dipole moment of the $|Q, J = 2, M = 2, \Omega = -2\rangle$ state ($d_Q = 1.60(2) e a_0$) and the linear stark shift in a quadratic electric potential is used to focus a larger portion of the molecular beam into the detection region than would otherwise be possible [85].

The first step of this process is to use stimulated Raman adiabatic passage (STIRAP) [86] to coherently transfer population from the $|X, J = 0\rangle$ state into the $|Q, J = 2, M = 2, \Omega = -2\rangle$ state




Figure 2.6: X-C-Q STIRAP Level Diagram. Level diagram for the STIRAP processes used with the electrostatic lens to transfer population from $|X, J = 0, M = 0\rangle \rightarrow |Q, J = 2, M = 2\Omega = -2\rangle$ and back. The black lines show the energy levels after the Stark shift, while green lines show the Zeeman shifts. W_z is the Zeeman shift of the $|Q, J = 2, M = 2\Omega = -2\rangle$ state, ΔW_S is the Stark shift between the $|Q, J = 2, M = 2\Omega = -2\rangle$ and $|Q, J = 2, M = 1\Omega = -2\rangle$ states, and δ is the one-photon detuning of the STIRAP lasers. Figure adapted from [85].

[87] as shown in figure 2.6. This process, which is also used for state preparation in the H state, is able to transfer population with $\sim 90\%$ efficiency, far more than can be achieved through optical pumping, which is critical for processes intended to increase molecular flux. Using $|C, J = 1, M = 1\rangle$ as an intermediate state, the pump laser is tuned to the $X - C$ 690 nm transition, while the Stokes laser is tuned to the $C - Q$ 1196 nm transition. With no applied fields the $|Q, J = 2, M = 2, \Omega = \pm 2\rangle$ states are degenerate, so an applied electric field of 50.1 V/cm is applied along \hat{z} to split the Ω . In order to populate only the stretched state with $M = 2$, both lasers propagate along \hat{z} with σ polarization. As the $M = \pm 2$ states are degenerate without a magnetic field, we also apply a magnetic field of 5 G along \hat{z} in order to break the degeneracy so that imperfections in polarization do not also populate the $M = -2$ level.

Once in the stretched $M = 2$ Q state, the molecular beam enters a hexapole lens consisting of 6




Figure 2.7: ACME III Electrostatic Lens Schematic and Electric Field Strength. **a.** A schematic of the electrostatic lens showing the voltage distributions of the lens. The lens uses cylindrical electrodes with $d_0 \approx 2$ cm. The blue and red dashed lines are cut lines for which the electric field strength is plotted in **b.** **b.** The electric field strength produced inside the lens along the two cut lines. Figure adapted from [85].

equally spaced cylindrical electrodes that are 53 cm long, which can be charged up to ± 30 kV with alternating polarity on each electrode as shown in figure 2.7. This configuration produces a radially varying quadratic field strength, which for the linear Stark shift produces a quadratic potential and therefore a linear restoring force. This restoring force can then be tuned to focus the molecular beam in order to effectively produce an image at the detection region.

After this focusing the molecules must then be returned to the $|X, J = 0\rangle$ state before the experimental spin state can be prepared. This is done using $Q - C - X$ STIRAP in the exact reverse of the process used to populate the Q state. With a demonstrated efficiency of $\sim 90\%$ for each STIRAP direction, the gains from the focusing ($\approx 12\times$ higher signal) far outweigh the losses during STIRAP ($\approx 80\%$ efficiency).

2.3.4 STATE PREPARATION

After the population of the molecular beam has been transferred to the $|X, J = 0\rangle$ state, either by rotational cooling in ACME II or by $Q - C - X$ STIRAP in ACME III, the molecules enter a magnetically shielded region, which we refer to as the interaction region, where we first prepare the superposition state that will be allowed to precess. In this region we apply both electric and magnetic fields, that cause precession, as described in 2.2. Upon entering this region, the molecular beam is also collimated again with a set of square fixed collimators (25 mm \times 25 mm in ACME II) designed to prevent any of the beam from hitting the electric field plates.

The state preparation occurs in two stages, shown in figure 2.8, the first of which uses STIRAP to coherently transfer population from the $|X, J = 0\rangle$ state to the desired spin state $|\psi(t = 0), \tilde{\mathcal{N}}\rangle$ as defined in equation 2.1 [86]. This process uses the $|C, J = 1, M = 0\rangle$ state as the intermediate state, with the pump laser tuned to the 690 nm $X - C$ line and the Stokes laser tuned to the 1090 nm $C - H$ line. The lasers are sent vertically through the experiment with a slight overlap such that the passing molecular beam experiences the correct temporal overlap of the two lasers. The lasers must be sent through the experiment vertically because the initial and final states differ by only $\Delta M = \pm 1$, so the pump laser uses \hat{z} polarization ($\Delta M = 0$), while the Stokes laser uses \hat{x} polarization ($\Delta M = \pm 1$). For more details on this STIRAP process, see [88, 86], and for more information on the gains provided by STIRAP over ACME I, see section 3.1.

While STIRAP is very efficient at preparing the desired state $|\psi(t = 0), \tilde{\mathcal{N}}\rangle$, imperfections can produce an imperfect initial spin state, which can cause systematic errors as will be discussed in section 3.4. To suppress these effects we implement a refinement laser step, also referred to colloquially as the “cleanup laser”. This laser is tuned to the 703 nm transition between $|H, J = 1, \tilde{\mathcal{N}}\rangle$ and the $|I, J = 1, M = 0^+\rangle$ state, which has a short lifetime of only ~ 115 ns [89, 90]. By sending this laser through the field plates along the \hat{z} axis of the experiment, the polarization of the laser can be

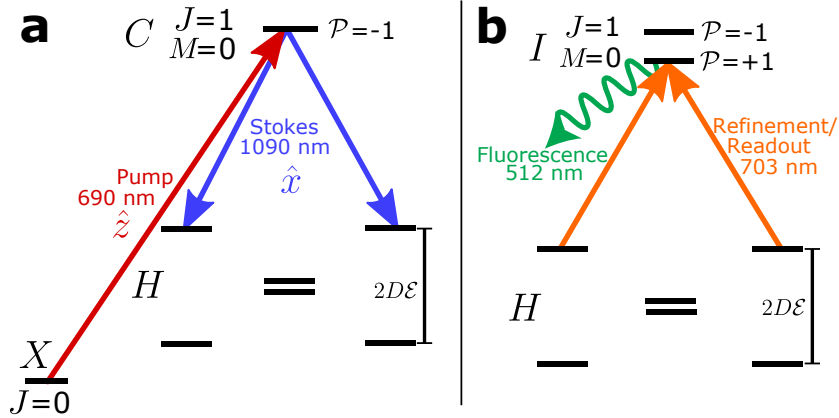


Figure 2.8: ACME II State Preparation. a. Level diagram showing how STIRAP was used to coherently transfer population from the $|X, J = 0, M = 0\rangle \rightarrow |H, J = 1, \tilde{N} = -1\rangle$ state in ACME II, which will also be used in ACME III. By tuning the resonance of the STIRAP lasers we can also address the other \tilde{N} . b. Level diagram showing the readout and refinement lasers which are used to either prepare a specific initial spin state, or project out the spin state by optically pumping from $|H, J = 1, \tilde{N}\rangle \rightarrow |I, J = 1, M = 0^+\rangle$. By tuning the 703 nm laser we are able to address either \tilde{N} , and we can also perform the \tilde{P} switch by addressing the $|I, J = 1, M = 0^-\rangle$ state. The 512 nm fluorescence light is used for detection when projecting out the spin state. This figure was adapted from [8].

aligned with the \hat{x} axis. This laser optically pumps out the bright state, which in this configuration is orthogonal to the desired state $|\psi(t = 0), \tilde{N}\rangle$, and the desired dark state is left behind. While this process could be performed with the C state as was done in ACME I [6, 7], in practice the $H - I$ transition dipole moment can be much more easily saturated than the $H - C$ transition, providing more powerful suppression of phase misalignment errors. We are also able to use the same laser system as the state readout described in section 2.3.6, reducing the complexity of our system and making the laser systems more robust.

It is important to note that for both the STIRAP lasers and the refinement laser, we are able to choose which \tilde{N} state is addressed, and perform rapid switching between states. This is critical for the switching described in sections 2.1.2 and 2.4. The methods for this switching will be discussed in detail in section 3.2.

2.3.5 FREE PRECESSION AND APPLIED FIELDS

Before the molecular beam reaches the state preparation lasers, it enters the interaction region, which runs from before the state preparation lasers until after the detection lasers. In this region, the molecules are exposed to uniform applied electric and magnetic fields $\vec{\mathcal{E}}$ and $\vec{\mathcal{B}}$ that define the \hat{z} axis. The electric field is produced by a pair of uniformly spaced glass field plates, separated by 4.5 cm for ACME II and 6 cm for ACME III, coated with a thin (20 nm in ACME II) layer of optically transparent conductive Indium Tin Oxide (ITO). These transparent field plates allow for both laser access along \hat{z} , and the detection of molecular fluorescence through them. The magnetic field is produced by a coil outside the vacuum chamber referred to as the cosine theta coil due to the angular dependence of the current density. This coil produces a highly uniform magnetic field in the region of interest, but we also have coils designed to apply auxiliary fields along \hat{x} and \hat{y} and the independent first order gradients ($\partial\mathcal{B}_z/\partial z$, $\partial\mathcal{B}_z/\partial y$, $\partial\mathcal{B}_z/\partial x$, $\partial\mathcal{B}_y/\partial y$, $\partial\mathcal{B}_y/\partial x$, and $\partial\mathcal{B}_x/\partial x$).

This region is also shielded from the Earth's magnetic field and stray magnetic fields in the room by a series of μ -metal magnetic shields. ACME II used a five layer shield, while ACME III will use a three layer shield, the development of which will be described in detail in section 5.2. In ACME II this led to a typical ambient field magnitude of $\sim 50 - 100 \mu\text{G}$, while in ACME III we intend to have an ambient field in the $\sim 1 - 10 \mu\text{G}$ range, which will be necessary to suppress excess noise associated with molecular beam velocity fluctuations.

Once the molecules have been transferred to the $|\psi(t=0), \tilde{\mathcal{N}}\rangle$ state by the refinement laser, they are allowed to precess freely in the applied fields for a fixed distance L between the refinement laser and the readout laser. In ACME I and II we used $L = 22 \text{ cm}$ [8], while for ACME III we will use $L = 100 \text{ cm}$ in order to make use of our better understanding of the H state radiative lifetime after our recent measurement [70]. This precession takes place for an amount of time τ determined by the forward velocity of the molecular beam $v_{||}$, which as previously noted can fluctuate from

shot-to-shot at the 10% level. After this time the molecular state can be represented as in equation

$$2.3 \text{ as } |\psi(t = \tau), \tilde{\mathcal{N}}\rangle = \frac{e^{-i\phi}|M=+1, \tilde{\mathcal{N}}\rangle - e^{i\phi}|M=-1, \tilde{\mathcal{N}}\rangle}{\sqrt{2}}.$$

2.3.6 STATE READOUT AND DETECTION

In the readout region, molecules are optically pumped from the $|H, J = 1, \tilde{\mathcal{N}}\rangle$ state to the $|I, J = 1, M = 0^+\rangle$ with a 703 nm laser similar to how the refinement laser prepared the state, shown in figure 2.8. The short lived I state decays to the ground state X by emitting a 512 nm photon. In order to project out the phase ϕ as described in section 2.2, we rapidly alternate the polarization between two orthogonal linear polarizations \hat{X} and \hat{Y} . This switching occurs at 200 kHz (i.e. $5 \mu\text{s}$) much faster than the molecular beam transit time ($\approx 10 \mu\text{s}$), so that each molecule is exposed to both polarizations, with enough dead time to prevent any overlap between the two polarizations. The signal produced by this switching scheme can be seen in figure 2.9.

The 512 nm fluorescence light then passes through the transparent field plates and is collected by eight pairs of lens doublets as shown in figures 2.3 and 2.4, which focus the light into fused quartz light pipes. These light pipes transfer the light with total internal reflection out of the vacuum chamber and the magnetic shields, where it exposes eight detectors. In ACME II these detectors were photomultiplier tubes (PMTs). To improve our collection efficiency, ACME III will use Silicon Photomultipliers (SiPMs), described in more detail in section 4.1.3.

2.4 ACME SWITCHES

In order to measure the EDM using the scheme described in the previous section, we must repeat the experiment many times in multiple configurations. We do this by performing a number of switches which typically either reverse the sign of applied field, or change a judiciously chosen calibration angle. These switches also allow us to study and significantly suppress systematic errors.

The structure of the ACME measurement timing and switches are shown in figure 2.9.

2.4.1 BLOCK LEVEL SWITCHES AND THE MINIMUM EDM MEASUREMENT

The basis of the EDM measurement is the “block” which consists of roughly 60 s of data taking, and consists of fast and independent switches of $\tilde{\mathcal{N}}$ (the relative sign of the internal electric field of the molecule and the applied electric field), $\tilde{\mathcal{E}}$ (the sign of the applied electric field), $\tilde{\theta}$ (the readout polarization basis dither), and $\tilde{\mathcal{B}}$ (the sign of the applied magnetic field). Within a single block, we measure Φ in each of the 2^4 different switch configurations. In practice our fastest switch is the $\tilde{\mathcal{N}}$ switch which we vary every 25 molecular beam shots which we average into a single “trace”, and takes about 0.6 s. We then vary the $\tilde{\mathcal{E}}$ state every four traces, corresponding to once every four $\tilde{\mathcal{N}}$ switches. We next vary $\tilde{\theta}$ every four $\tilde{\mathcal{E}}$ switches. Finally our slowest switch is the $\tilde{\mathcal{B}}$ switch which is limited by the long time (~ 5 s) required for degaussing the magnetic shields with each field reversal. We perform this switch once for every two $\tilde{\theta}$ switches and only once per block, roughly every 30 s. Additionally, to further suppress potential errors we randomly alternate the patterns of the switching, with $\tilde{\mathcal{N}}$ and $\tilde{\mathcal{B}}$ following either a $(+-)$ or $(-+)$ pattern, and $\tilde{\mathcal{E}}$ and $\tilde{\theta}$ following either a $(+ - - +)$ or $(- + + -)$ pattern.

With these measurements we can define “switch parity components”, which are odd under the specified switch [62]. We define a convention of using a superscript index u on a quantity to represent all of the switches which that quantity is odd under. Following this definition, the quantity is then even under any switches not represented by u . For any quantity that does not reverse with any switches, we use the superscript nr for non-reversing. We use this formalism to represent the measured phase Φ as a function of $\tilde{\mathcal{N}}$, $\tilde{\mathcal{E}}$, and $\tilde{\mathcal{B}}$ as

$$\Phi(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}) = \Phi^{nr} + \tilde{\mathcal{N}}\Phi^{\mathcal{N}} + \tilde{\mathcal{E}}\Phi^{\mathcal{E}} + \tilde{\mathcal{B}}\Phi^{\mathcal{B}} + \tilde{\mathcal{N}}\tilde{\mathcal{E}}\Phi^{\mathcal{N}\mathcal{E}} + \tilde{\mathcal{N}}\tilde{\mathcal{B}}\Phi^{\mathcal{N}\mathcal{B}} + \tilde{\mathcal{E}}\tilde{\mathcal{B}}\Phi^{\mathcal{E}\mathcal{B}} + \tilde{\mathcal{N}}\tilde{\mathcal{E}}\tilde{\mathcal{B}}\Phi^{\mathcal{N}\mathcal{E}\mathcal{B}}. \quad (2.6)$$

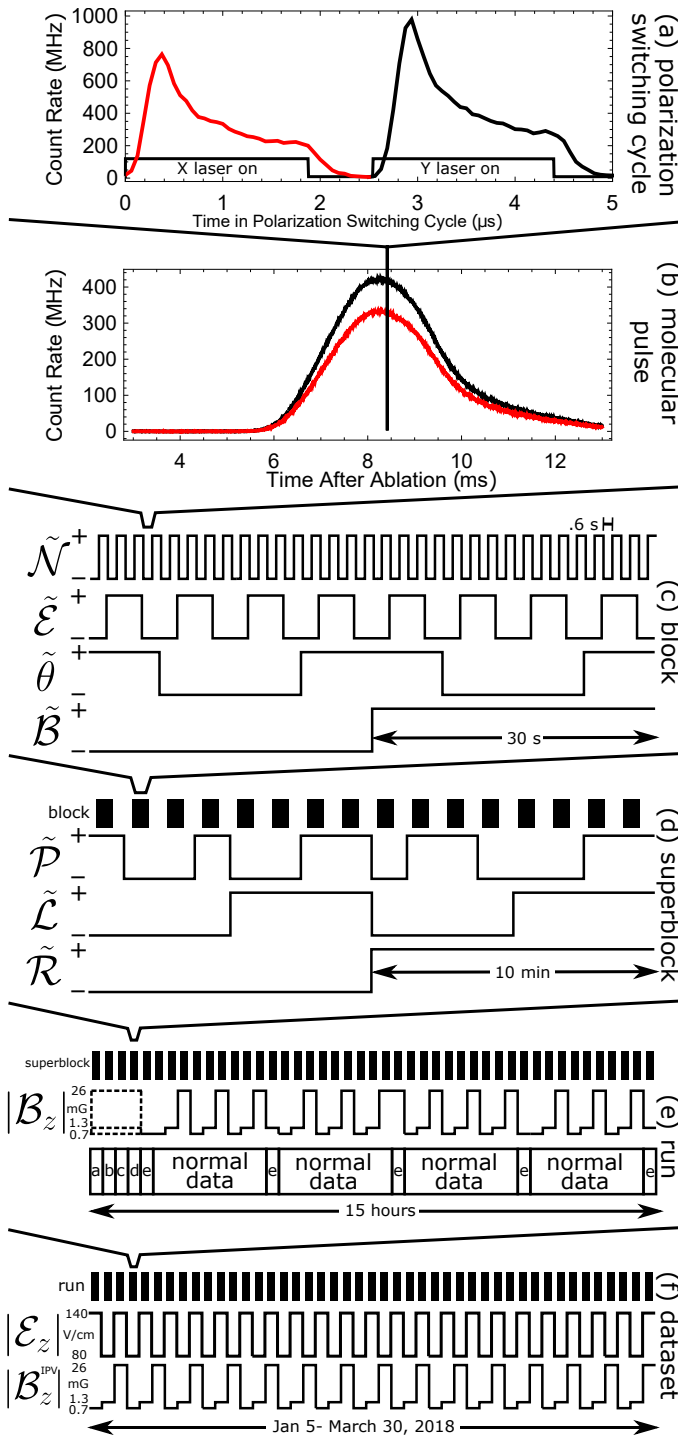


Figure 2.9: ACME II Switching Structure.

The full structure of switches and timing sequences used in the ACME II measurement. **a.** Within a molecular beam shot we switch between polarization states \hat{X} (signal shown in red) and \hat{Y} (signal shown in black) for the detection laser in order to project out the spin state. **b.** The signal from a single trace consisting of 25 averaged shots separated out into data from the two polarization states with \hat{X} signal shown in red, and \hat{Y} in black once again. **c.** The timing of a single block, with the four block level switches, \tilde{N} , \tilde{E} , $\tilde{\theta}$, and \tilde{B} , performed with a four-fold degeneracy. One block measurement takes roughly one minute of data taking. **d.** The timing of a single superblock measurement showing the relative timing of the three switches \tilde{P} , \tilde{L} , and \tilde{R} , which takes approximately twenty minutes of data taking. **e.** The structure of a single “run” of data taken within a day. On this time scale we switched between taking “normal data” at multiple values of $|\mathcal{B}_z|$, while also performing intentional parameter variations for parameters, labeled here as a-e, for a superblock. It should be noted that in addition to the $|\mathcal{B}_z|$ values shown in this figure, we also ran at $|\mathcal{B}_z| = 2.6$ mG. **f.** The structure for the entire data set used in the ACME II result. At this time scale we also varied the applied electric field magnitude $|\mathcal{E}|$ and the value of $|\mathcal{B}_z|$ associated with IPVs. This figure was adapted from [8].

We also define the transformation by which we can extract component in the parity basis as

$$\Phi^u(\mathcal{N}, \mathcal{E}, \mathcal{B}) = \frac{1}{2^8} \sum_{\tilde{\mathcal{N}}', \tilde{\mathcal{E}}', \tilde{\mathcal{B}}' = \pm 1} (\tilde{\mathcal{N}}')^{\frac{1-\mathcal{N}}{2}} (\tilde{\mathcal{E}}')^{\frac{1-\mathcal{E}}{2}} (\tilde{\mathcal{B}}')^{\frac{1-\mathcal{B}}{2}} \Phi(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}), \quad (2.7)$$

which can be extended to consider an arbitrary number of switches.

It is important now to consider what the individual switch components represent. Most importantly, $\Phi^{\mathcal{N}\mathcal{E}}$ is the EDM correlated phase component, used to ultimately extract d_e . This can be seen as it by definition reverses with both $\tilde{\mathcal{N}}$ and $\tilde{\mathcal{E}}$, which are the only switches we would expect to reverse the energy shift produced by the EDM given as $U = -\vec{d}_e \cdot \vec{\mathcal{E}}_{\text{eff}}$. $\Phi^{\mathcal{B}}$ is the Zeeman precession phase, which we use to extract the average precession time τ for each block using the equation $\Phi^{\mathcal{B}} = -\mu|\mathcal{B}_z|\tau/\hbar$. For ACME II, $\Phi^{\mathcal{B}} \approx \pi/4$ with an applied field of $|\mathcal{B}_z| = 26$ mG.

We can now combine what we have covered in this section and section 2.2 to calculate the EDM from a single block. We first compute the contrast as

$$\mathcal{C}(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}) = \frac{-1}{2\Delta\theta} [\mathcal{A}(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}, \tilde{\theta} = +1) - \mathcal{A}(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}, \tilde{\theta} = -1)] \quad (2.8)$$

where here $\Delta\theta$ is the dither angle of ~ 0.2 rad, and we average $\mathcal{A}(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}, \tilde{\theta})$ over the two $\tilde{\theta}$ states as

$$\mathcal{A}(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}) = \frac{\text{sgn}(\mathcal{C})}{2} [\mathcal{A}(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}, \tilde{\theta} = +1) + \mathcal{A}(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}, \tilde{\theta} = -1)]. \quad (2.9)$$

These definitions and equation 2.4 relate the EDM d_e to $\Phi^{\mathcal{N}\mathcal{E}}$ as

$$\Phi^{\mathcal{N}\mathcal{E}} = -d_e \mathcal{E}_{\text{eff}} \tau / \hbar, \quad (2.10)$$

from which d_e can be extracted most reliably.

2.4.2 OTHER SWITCH TIMESCALES

At slower timescales than the ≈ 60 s block we perform additional switches which are designed to suppress potential systematic errors. For a more in-depth discussion of these potential systematic errors and others see section 3.4. We first define the “superblock” structure which consists of three switches performed over the course of approximately 20 minutes. While not necessary for the minimal EDM measurement made with the block level switches, in practice we make use of all 2^7 block and superblock switch states when extracting $\Phi^{\mathcal{N}\mathcal{E}}$.

The first of these switches is the $\tilde{\mathcal{P}}$ switch, which is performed every two blocks (roughly every two minutes). This switch consists of changing which of the two parity states of the $|I, J = 1, M = 0^\pm\rangle$ state are used for the readout laser. This is done by tuning the readout laser frequency to select the correct resonance, where the two lines are split by $\Delta_\Omega = 90$ MHz [90]. The $\tilde{\mathcal{P}} = +1$ state addresses the $M = 0^+$ state as described in section 2.3.6, while $\tilde{\mathcal{P}} = -1$ corresponds to addressing the $M = 0^-$ state. Reversing $\tilde{\mathcal{P}}$ is equivalent to rotating the angle between the refinement and readout lasers, θ , by $\pi/2$. This serves to remove systematic effects that appear as a spurious asymmetry. Similar to the block switches, we randomize the order of this switch, while we do not randomize the order of the other two superblock switches.

The second superblock switch is the $\tilde{\mathcal{L}}$ switch, which is performed every four blocks, and consists of mechanically reversing the leads of the voltages that are applied to the field plates. In conjunction with switching the relays, we also switch the polarity of the supplies such that in an ideal system this switch should have no effect on the apparatus. This suppresses systematics that may be related to offsets in the voltage supplies.

The final superblock switch is the $\tilde{\mathcal{R}}$ switch, which is performed every eight blocks (≈ 10 min), and consists of physically rotating the polarization basis of the readout lasers by $\pi/2$. This is performed by mechanically rotating a wave plate in the path of both the \hat{X} and \hat{Y} readout lasers. This

switch has the same effect as the $\tilde{\mathcal{P}}$ switch, which further helps to suppress asymmetry related effects.

In addition to superblock switches we also take EDM data at multiple values of $|\mathcal{B}_z|$ and $|\mathcal{E}|$. In ACME II we used the values $|\mathcal{B}_z| = 0.7, 1.3, 2.6,$ and 26 mG, and $|\mathcal{E}| = 80$ and 140 V/cm. For ACME III, we are likely to use the same values of $|\mathcal{E}|$, but we will need to use new values of $|\mathcal{B}_z|$ to reduce a known excess noise source, discussed in sections 3.4 and 5.2.

In addition to the previously described switches which we consider to be “normal” data collection, we also interspersed our data collection with time spent under “intentional parameter variations” (IPVs). In these conditions one or more experiment parameters were intentionally set to deviate from their ideal parameters, often to an exaggerated degree relative to what could reasonably occur in the experiment. This allowed us to both measure the sensitivity of the EDM phase to known systematics and search for new sources of systematic error with high sensitivity. These IPVs and the related systematic effects will be discussed in detail in section 3.4.

But somewhere along the way

I started to smile again

I don't remember when

Somewhere along the way

Dawes, *Somewhere Along the Way*

3

The ACME II Measurement

IN 2018, ACME II REPORTED A NEW MEASUREMENT THAT SET A LIMIT ON THE ELECTRON EDM $|d_e| < 1.1 \times 10^{-29} e \text{ cm}$, which improved on the previous best limit set by ACME I by nearly an order of magnitude [8, 6]. This result is already very well documented in the theses of Cristian Panda [88], Zack Lasner [84], and Brendon O'Leary [91]. This chapter is my summary of the ACME II measurement I participated in. It will also provide context and a basis for under-

standing the development of the ACME III experiment, discussed in chapters 4 and 5. Section 3.1 discusses the way the experiment was upgraded to improve sensitivity after the ACME I result. Section 3.2 discusses the laser systems used for the ACME II measurement, many of which will be the same for ACME III. Section 3.3 discusses the electric magnetic field control systems of ACME II, which had to be redesigned after ACME II to accommodate the longer interaction region of ACME III. Section 3.4 discusses the systematic errors that contributed to the ACME II result, with a particular focus on those which factor heavily into the design of ACME III. Finally, section 3.5 discusses the statistics of the ACME II result.

3.1 UPGRADES OVER ACME I

The transition from ACME I to ACME II is marked by a significant number of changes that improve experiment sensitivity. In order to better understand the motivation behind each change, it is helpful to describe all changes as belonging to one of three categories. The first category is those changes which improve the statistics of the experiment by lowering the shot-noise limit. The second category is those changes which suppress known systematic effects that were either significant in the previous generation, or are expected to become significant in the new measurement. These first two categories will be discussed in this section. The final category of upgrades consists of those which either make the experiment more robust or address known difficulties with the previous generation, however the benefits of these upgrades are harder to quantify. In order to increase the robustness of the experiment and reduce time spent on experiment maintenance, improvements were made to the data acquisition system [91], the interaction region vacuum system, the laser systems as discussed in section 3.2, and the magnetic field monitoring as discussed in section 3.3.

In chapter 1 we showed that the sensitivity of a shot-noise limited measurement scales as

$$\delta d_e = \frac{\hbar}{2\mathcal{E}_{\text{eff}}\tau\sqrt{\dot{N}T}}, \quad (3.1)$$

which serves as guide for how we can improve each ACME experiment. While in an ideal case we would be able to target improvements at the two terms, \mathcal{E}_{eff} and τ , which the sensitivity scales as the inverse of, but in practice for ACME II this was not possible. In order to improve on \mathcal{E}_{eff} we would need to switch to another molecule, as we were already able to fully polarize ThO in ACME I [7, 6]. Although we now have a better understanding, during ACME I and II the H state lifetime was only known to be $\tau_H \gtrsim 1$ ms [78, 75, 92], and the precession time of $\tau \approx 1$ ms had been optimized based on this understanding. This leaves only the count rate \dot{N} , and the integration time T as candidates for improving the sensitivity of ACME II. The majority of the statistical sensitivity upgrades that were implemented in ACME II ultimately improved the count rate, as the integration time proved to be limited both by the amount of time spent collecting data during ACME I, and the duty cycle of the experiment. The methods for increasing the count rate in ACME II primarily focused on removing steps of the experiment with large losses in order to maximize the flux of usable molecules and improve the detection efficiency of the system. We can summarize these signal level gains in table 3.1, which is discussed in detail in the following sections. We also now clarify the distinction that will be made in this work between signal gain, which increases the fluorescence count rate of the experiment, and sensitivity gain, which directly changes δd_e , and scales as the square root of signal gain¹.

¹In order to make a better comparison between signal and sensitivity gain, we typically talk about improvements as gains to both values, where for sensitivity gain we are actually describing the gain for the quantity $1/\delta d_e$.

Upgrade	Demonstrated Signal Gain
<i>X-C-H</i> STIRAP for State Preparation	$\sim 12\times$
Beam Geometry	$\sim 7\times$
Collection Optics Efficiency and Detection Wavelength	$\sim 5\times$
Actual Combined Gain	$\sim 400\times$

Table 3.1: ACME II Demonstrated Signal Gains. Table showing the measured gain in the fluorescence count rate relative to ACME I for each of the count rate upgrades implemented in the ACME II measurement. Each upgrade was first demonstrated independently to determine the expected gain, then the combined gain was measured on the full ACME II apparatus, giving the actual combined gain.

3.1.1 STIRAP

The largest single signal gain in ACME II came from improving the efficiency of the state preparation method by implementing a STIRAP process. This work was led by Cristian Panda, and was reported both in his thesis [88], and in [86]. In ACME I, state preparation was performed as shown in figure 3.1 by optically pumping molecules from the $|X, J = 1^-, M = \pm 1\rangle$ state, which had been prepared with a different rotational cooling scheme from ACME II, into the $|A, J = 0^+, M = 0\rangle$ state [7, 6]. It should be noted that although this was performed in an applied electric field, $\Omega = 0$ for both the X and A states, so parity is still a good quantum number in this case. The excited $|A, J = 0, M = 0^+\rangle$ state can decay to the $|H, J = 1\rangle$ manifold that we use for the experiment, with a measured branching ratio of ~ 0.35 [93, 88]. However, this incoherent decay populates five out of the six states in the $|H, J = 1\rangle$ manifold, with $1/6$ of the population in each of the four $|H, J = 1, \tilde{N} = \pm 1, M = \pm 1\rangle$ states, and $1/3$ of the population in the parity allowed $M = 0$ state, $|H, J = 1, M = 0^-\rangle$. As such, only $1/6$ of the decay from the A state ended up in the desired initial state $|\psi(t = 0), \tilde{N}\rangle$ as defined in equation 2.1, while the remaining population in the desired \tilde{N} state had to be pumped out by the refinement laser (using optical pumping to the C state rather than the I state). Ultimately this resulted in a transfer efficiency of only 6% [7].

As STIRAP is a coherent process, it can be used to both prepare a desired spin aligned state in the

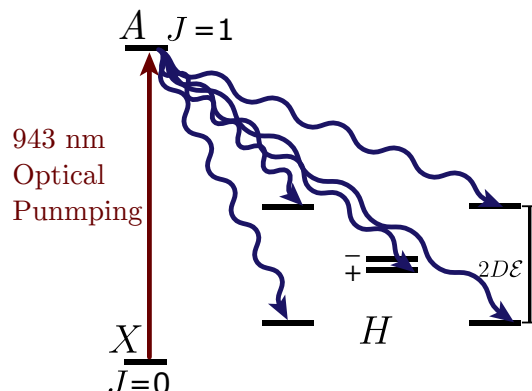


Figure 3.1: ACME I State Preparation Scheme. Molecules are excited to the $A, J = 1$ state via optical pumping with a 943 nm laser. Once in the A state the molecules spontaneously decay to the H state with a branching ratio of ~ 0.35 . This method only has a transfer efficiency of 6%. Once in the H state, a refinement laser prepares the spin aligned state.

final state, and select a particular N state state. These combined effects allowed us to demonstrate a $75\% \pm 5\%$ transfer efficiency from $|X, J = 0\rangle \rightarrow |\psi(t = 0)\tilde{\mathcal{N}}\rangle$ [86]. While this efficiency is $\sim 12\times$ higher than that of optical pumping, it is notable that this efficiency is less than has now been demonstrated for X - C - Q STIRAP, which has demonstrated 90% efficiency [87]. One possible explanation of this lower efficiency is due to the power requirements of saturating the 1090 nm C - H Stokes transition. Tests in which the power of the pump and Stokes lasers were independently varied indicated that the transfer efficiency was still increasing with Stokes laser power when the maximum available power of our fiber amplifier (10 W) was reached. This may be a possible source of further gain in the future, but this gain will be quite small ($\sim 20\%$).

One benefit that the ACME I scheme had over the ACME II state preparation method was that it was possible to switch the prepared state orientation, which was referred to as the $\tilde{\mathcal{G}}$ switch. This switch was performed as a global rotation by $\pi/2$ for both the refinement laser and the readout beams [6, 7]. However, for STIRAP we require the use of \hat{z} polarized light due to selection rules, so the STIRAP lasers are launched vertically through the experiment, and so \hat{y} polarized light is not

possible. The removal of this switch did not impact the systematic errors in ACME II, but has been considered as an option for ACME III.

3.1.2 BEAM GEOMETRY

The second largest signal gain in ACME II came from optimizing the geometry of the experiment in order to better take advantage of the large flux of molecules exiting the beam source that did not make it to the detection region. This upgrade, which ultimately resulted in a factor of $8\times$ larger beam flux, came from two modifications; decreasing the distance the beam traveled before the state preparation region, and increasing the separation of the field plates and the fixed collimators. Notably, these changes were made while still using the majority of the apparatus of ACME I, and the ultimate values chosen for these changed parameters reflect the limitations of this apparatus. As ACME III required a complete redesign of the interaction region, the designs of the new ACME III apparatus were not limited by the ACME I and ACME II apparatus, and the upgrades discussed in sections 4.1 and 5.3 reflect this fact.

The molecular beam source has a uniform angular distribution so we can easily see that the number of molecules that pass through a square fixed collimator a distance l from the beam source skimmer will scale as $\propto 1/l^2$. In ACME I, this distance between the two apertures was $l_I = 1.3$ m, which was largely set by the distance needed for rotational cooling, the length needed for formed bellows to join the interaction region chamber to the beam source, gate valves, space for an adjustment mechanism on the interaction region collimator, and a nipple to connect this “stem” region to the chamber inside the shields. ACME II was able to decrease this length to $l_{II} = 1.1$ m by converting the adjustable collimator to a low profile fixed separation collimator, reducing the length of the rotational cooling region, and changing to shorter length vacuum components including shorter gate valves and edge welded bellows. Ultimately this resulted in a gain factor of $\frac{l_I^2}{l_{II}^2} = 1.4$. However, this distance could not be further decreased due to the required distance for the ACME II rotational

cooling scheme, and the long nipple required to connect the interaction region chamber inside the five layer magnetic shields to the stem region.

Further signal improvements were made possible by reconsidering the aperture size of the interaction region collimator, which is designed to protect the field plates from the molecular flux. This was done so that ThO molecules and other beam source products could not build up on the conductive ITO field plate coating and create new patch potentials, which could cause new systematic errors. In order to do this the collimator was set so that no ballistic path from the beam source, modeled as a small disk, could pass through the collimator, with a 1 cm separation in ACME I, and come within 3 mm of the field plates, which were separated by 2.5 cm in ACME I². As such, increasing the collimator size also necessarily means increasing the separation of the field plates. While in an ideal case the beam flux would scale quadratically in the collimator size, in practice the efficiency of the collection optics also drops with field plate separation and with the larger molecular beam size. Cristian Panda performed simulations [88] in order to determine that with the ACME I collection optics the gain could be optimized by moving to a 2.5 cm collimator and a field plate separation of 4.5 cm, which combined with the gain of 1.4 listed above, gave a $\sim 7 - 8\times$ larger signal.

While changes were made to the collection optics in ACME II, it is helpful to separate out this gain as an independent factor as will be done in the next section. Although it did not impact the design of ACME II, it is worth considering that the increased beam size significantly increases the laser power required to saturate all of transitions used, both because the beam includes a larger range of transverse velocity and thus a larger Doppler width, and because the spatial area needed to be covered by the lasers is larger. However, this change may have also contributed to the slightly lower observed non-reversing electric field, \mathcal{E}^{nr} , due to the increased distance between any patch potentials

²Both Cristian Panda and Brendon O’Leary note in their theses that the original reason for this design choice has been “lost to history” [88, 91] so I can’t speculate as to the reason.

on the field plates and the molecular beam.

3.1.3 IMPROVED COLLECTION EFFICIENCY AND DETECTION WAVELENGTH

As described in the previous section, it became necessary to redesign the collection optics, which transferred the readout fluorescence to the detectors, in order to optimize for the larger field plate separation. The collection optics, which consist of eight lens doublets, which are placed on the outside of the electric field plates to ensure field homogeneity, were optimized using simulations performed by Cristian Panda [88] in LightTools. This resulted in tiling the lens doublets closer together and optimizing the angles of the lenses. In addition to this, the two 9 mm aperture bundled fibers used as light guides in ACME I were replaced with eight 16 mm fused quartz light pipes, each going to their own PMT. This change served two purposes, the first of which was that the light pipes do not suffer from the limited 65% fiber packing fraction of the fiber bundle. The second purpose was increasing the total detection area by a factor of 4 by adding an individual light pipe and PMT to each lens doublet, as opposed to combining four lens doublets as in ACME I.

In addition to these collection optics changes, ACME II changed from using the $|C, J = 1\rangle$ state for refinement and readout to using the $|I, J = 1\rangle$ state for both processes, as described in sections 2.3.4 and 2.3.6. The primary benefit of this change came from the new wavelength of the detected photons, 512 nm for the I state versus 690 nm for the C state, for which available PMTs have a much higher quantum efficiency³. In practice, this resulted in a $2.5\times$ gain in detector quantum efficiency over ACME II [91]. The branching rate from $|I\rangle \rightarrow |X\rangle$, measured to be 91% [89], was also significantly higher than the $|C\rangle \rightarrow |X\rangle$ branching ratio of only 75% as determined by Paul Hess [80]. While not a direct gain, this new detection wavelength did not overlap with the laser wavelengths used in any of the other processes in the experiment, which meant that any laser background

³Hamamatsu R7600U-300

could be removed with bandpass filters, as had also been the case in ACME I. Last, the I state radiative lifetime (~ 115 ns [89]) was much shorter than the C state radiative lifetime (~ 500 ns [80]). This meant that the polarization switching rate could be increased from 100 kHz to 200 kHz, and as a result the width of the readout laser along \hat{x} could be decreased, which helped reduce concerns about laser power introduced by the changes to the beam geometry. The new transition also required a change of the applied electric field compared to ACME I, which was chosen such that the X - I resonance was far from the X - C resonance used for STIRAP while maintaining a relative small applied field.

Unfortunately, these decisions were made after the field plates for ACME II had already been coated, and the coatings had not been optimized for 512 nm, resulting in a transfer efficiency through the field plates of only $\sim 70\%$. Despite this, the combined collection efficiency was measured to be $\sim 20\%$ [88], which combined with the quantum efficiency of the PMTs at 512 nm resulted in a combined detection efficiency of $\sim 4\%$. Ultimately this was roughly $5\times$ higher than in ACME I. However, this number still had a large margin for possible improvement, and was revisited in ACME III, as discussed in sections 4.1 and 4.1.3.

3.1.4 ELECTRIC FIELD PLATES

In ACME I, the largest systematic error was caused by a coupling between polarization ellipticity gradients in the refinement and readout lasers, which coupled with the AC Stark shift to prepare a portion of the phase in the spin state orthogonal to the desired initial spin aligned state. For a full description of this model see the ACME I result papers [6, 7]. This systematic, which was significantly suppressed due to the efforts described in this section, remained significant for ACME II and is discussed further in section 3.4. The ellipticity gradients in ACME I were caused by heating of the glass field plates by the refinement and readout lasers, which had a spatially varying intensity profile, with a $\sim 10\%/mm$ circular polarization component along \hat{x} [7, 6], which introduced stress

to the glass, causing a spatially varying birefringence. This particular effect has been observed in a number of different systems including Nd:YAG rods [94] and windows [95, 96]. Fortunately for ACME I, this systematic could be managed by rotating the polarization of the refinement laser such that the polarization aligned with the birefringence axis of the field plate, and lowering the incident laser power. This resulted in simply rotating the initial prepared state by an angle θ_{prep} relative to the \hat{x} axis. Unfortunately for ACME II, this was no longer possible as the STIRAP scheme required that the initial state be aligned with the \hat{x} axis due to the polarization requirements of the STIRAP lasers.

In order to understand the solution to this problem used in ACME II, it is helpful to consider the model for how this thermal stress-induced birefringence occurred. For this case we need only consider the spatial variation of the laser in the \hat{x} direction as the laser is sufficiently stretched along \hat{y} to engulf all molecules that the spatial variation along \hat{y} is small in the region we care about. If we also assume that there are no shear stresses, we can represent the difference between the indices of refraction in the \hat{x} and \hat{y} directions using the stress-optic law [97] as

$$\Delta n = K(\sigma_{xx} - \sigma_{yy}). \quad (3.2)$$

Where in this equation σ_{xx} and σ_{yy} are the principal stresses along the \hat{x} and \hat{y} axes, while K is a material property of the glass called the “stress-optic coefficient”. This results in a retardance for each laser given by $\Gamma = 2\pi\Delta n(t/\lambda)$, where λ is the laser wavelength and t is the thickness of the glass. We can now consider the stress caused by the heating rate $\dot{Q}(x)$ given by

$$\frac{\partial^2 \sigma_{yy}}{\partial x^2} = \frac{E\alpha_V}{\kappa} \dot{Q}(x), \quad (3.3)$$

where E is the Young’s modulus of the glass, α_V is the coefficient of thermal expansion of the glass,

and κ is the thermal conductivity of the glass [98]. With these quantities and the bulk absorbance of the glass, α , we can define a material property of the glass

$$\eta_{\text{thermal}} \equiv 2\pi K E \alpha_V \alpha / \kappa, \quad (3.4)$$

which is just a material property of the glass [7, 79], such that we can relate the intensity profile of the laser $I(x)$ to the stress induced birefringence as

$$\frac{\partial^2 \Gamma}{\partial x^2} = \eta_{\text{thermal}} \frac{t}{\lambda} I(x). \quad (3.5)$$

Accordingly it is possible to select a glass such that η is minimized.

In ACME I, the field plates were made from Borofloat glass [99], which was not chosen to optimize this parameter in any way, and has a much higher value of η_{thermal} than many other available glasses on the market, where $\eta_{\text{thermal, Borofloat}} \approx 20 - 30 \times 10^{-6} \text{ W}^{-1}$ [99]. Ultimately, ACME II found that the best alternative was Corning 7980 fused silica [100], as the roughly $\sim 6 \times$ smaller coefficient of thermal expansion, and $\sim 50 \times$ smaller absorbance result in $\eta_{\text{thermal, 7980}} \approx 6 \times 10^{-8} \text{ W}^{-1}$ [100]. This reduced the thermal stress induced birefringence due to the glass by about a factor of ~ 300 , well below the expected level of sensitivity for ACME II. This glass was also used for the relevant vacuum windows as these components also contribute to the ellipticity of the laser profile.

In this model it is also important to consider the absorbance of the ITO coating layer, which in ACME I had a thickness of 200 nm and an estimated absorbance of roughly $\sim 5 - 10\%$, with higher absorbance at 703 nm than 1090 nm [79]. In ACME I, this was comparable to the bulk absorbance of the borofloat glass, but this absorbance would dominate the absorbance of the Corning 7980 field plates. As such, in an effort to lower the absorbance of the ITO coating, the ACME II

field plates were coated with only 20 nm of ITO.


In order to verify that these changes would not introduce new issues, the resistivity of test pieces from the same coating run were measured to have a resistivity of 2000 – 4000 Ω /Square, which was sufficient for ACME II [88]. In addition to this, Vitaly Wirthl performed measurements with the polarimeter originally built by Paul Hess and Ben Spaun [101, 80, 93] in order to determine that the thermal stress induced birefringence was $\sim 0.1\%$ /mm at the maximum intended 703 nm laser power of 1 W[101]. This was also consistent with tests performed with up to 6 W of 1090 nm light on field plate samples. This successfully suppressed the thermal stress-induced birefringence below the ACME II sensitivity, however as discussed in section 3.4 stress-induced birefringence remained a significant source of systematic error in ACME II, and many of these considerations were revisited in the design of ACME III as described in section 4.2.2.

3.1.5 UPGRADES THAT WERE NOT IMPLEMENTED

In the development phase of ACME II multiple upgrade paths were considered, and significant effort was put into some systems which ultimately did not make it into the final apparatus. There were two such systems which are of particular interest, both for the amount of development put into the systems and for what we learned from them that we would eventually implement in ACME III.

THE THERMOCHEMICAL BEAM SOURCE

The first of these upgrade projects was the thermochemical beam source developed by Elizabeth Petrik West [102] and Jacob Baron. The idea behind this system was to create a new buffer gas beam source that used the high temperature reaction



where the parentheses indicate the phase of the reactants and product, in order to increase the beam source molecule flux. This system used pressed powder targets made up of 75% ThO₂ and 25% Th, which was heated by a power modulated fiber laser in a cell similar to the ablation cell.

While high beam fluxes were demonstrated this upgrade was not implemented for two primary reasons. The first was that the system required further study in order to either understand or control variations in the beam properties which could introduce noise to the experiment. The other reason, which is particularly relevant in this work is that the signal from a fresh thermochemical target decayed significantly (a factor of $\sim 3 - 5$) within the first 12 hours of use[102]. This would have required daily changes of the pressed powder targets which need to be kept at cryogenic temperatures while running in order to provide a significant gain over the ablation source. The ablation targets also show significant decay over time, and it was found that the gains of the thermochemical source would be smaller than hoped if compared to the ablation source with a similar target change frequency. While not implemented for ACME II, these considerations and what we learned about the ablation targets decay over time provided the motivation for the development of a load lock for rapid target changes which will be used for ACME III and is discussed in section 4.1.5.

THE GROUND STATE ELECTROSTATIC LENS

The second upgrade project was an electrostatic lens to focus molecules in the $|X, J = 1^-, M = 0\rangle$ or $|X, J = 2^+, M = 0\rangle$ state [88]. A lens was built and tested using quadrupole electrodes, which operated similarly to the lens we plan to use in ACME III, which is discussed in sections 2.3.3 and 4.1.2. This lens system required that the electrodes be charged to ± 30 kV in order to achieve a gain of roughly ~ 20 relative to ACME I. Unfortunately it turned out that leakage currents from the lens electrodes produced 60 keV bremsstrahlung x rays [103]. While significant effort was spent on preventing these x rays, particularly by Adam West, we were unsuccessful at preventing them with this lens system. Additionally, it was determined that a $\sim 2\times$

smaller gain could be achieved with the beam geometry changes described in section 3.1.2 without the addition of the lens. We revisited the concept of the electrostatic lens when developing ACME III, with the plan to perform the lensing in the recently studied Q state [87], where the lens which has been demonstrated requires roughly half of the applied potentials, and with improved manufacturing we have demonstrated charging to ± 30 kV without x ray production [85].

3.2 LASER SYSTEMS

The ACME II experiment collected data for 16 hours per day, 5 days a week, for roughly 3 months during the final measurement run, and collected data at slightly slower rates for over a year. In order for this to be possible, it was necessary to ensure that all of the laser systems were as robust as possible so that the experiment did not lose unnecessary time to fixing unlocked lasers. Additionally, it was necessary to implement a large number of frequency switches for some of the lasers. In this section we will describe the laser systems used for ACME II, as well as efforts to ensure their robustness.

3.2.1 THE ABLATION LASER

From a maintenance perspective, the ablation laser was the simplest laser system we used in ACME II, as it did not require any frequency stabilization. We used a Nd:YAG pulsed laser⁴ with pulse energy ~ 50 mJ, pulse duration of ~ 15 ns, and repetition rate of 50 Hz to ablate the pressed powder targets. While this laser was doubled from 1064 nm to 532 nm the ratio of IR to green light has no significant effect on the ablation process [79], but the green light is used for alignment purposes. The laser is focused down to a small spot on the target surface by a lens approximately one focal distance away, however the exact spot size is not well known as the lens position is regularly manually

⁴Litron Nano TRL 80-200

adjusted to optimize molecular beam flux. In addition to this, the mirror which directs the beam into the cell is controlled by a pair of servos, that allow us to control the beam position on the target from the main experiment control computer. This was necessary as after selecting a new ablation spot on the target, the molecular flux would begin to decay very quickly, such that for a new target the signal level would drop below the threshold we set for good data within $\sim 20 - 30$ minutes for a brand new target, or ~ 5 minutes for a target near the end of its life. These servos allowed us to track where we had already ablated the target and depleted the signal, while also allowing us to easily see the signal response to the ablation spot location in real time.

We also discovered during the ACME II data run that the YAG power supply broadcasts voltage spikes every time the flash lamps fire (every 5 ms). This caused a dip in the fluorescence signal as the discharge was picked up by the data acquisition system. We were able to shield this effect by ensuring that the housing of the YAG power supply was properly closed during use and placing a Faraday cage around the laser head and umbilical. This should be kept in consideration for ACME III.

3.2.2 ACME II ROTATIONAL COOLING LASERS

The two rotational cooling lasers used in ACME II were external cavity diode lasers (ECDLs) built in the lab using the design described in Yulia Gurevich's thesis [104]. These lasers were frequency stabilized by locking to a scanning Fabry-Pérot transfer cavity, which was itself referenced to a stable laser. This stable laser was an Nd:YAG laser at 1064 nm that was frequency doubled and locked to a reference line in a molecular iodine cell, which was developed by Daniel Farkas [105]. With this method we are able to stabilize the frequency of each laser with a linewidth of roughly $\sim 1 - 2$ MHz. This reference source provided significantly more frequency stability than we could reasonably take advantage of during the ACME II experiment, but it was available to the ACME II experiment, and was used for both transfer cavity locks and calibrating our wave meters. Notably,

this was not transferred to Northwestern University when the experiment moved, and so new stable references have been implemented for ACME III. The two ECDLs, which optically pumped the $|X, J = 2^+\rangle \rightarrow |C, J = 1^-\rangle$ and $|X, J = 1^-\rangle \rightarrow |C, J = 1^{\text{Mixed}}\rangle$ transitions were built using ThorLabs HL6750MG diodes, which were rated for 50 mW, which after fiber coupling and locking was sufficient to provide the ~ 15 mW we needed for laser cooling.

The lasers are fiber coupled and sent to a table next to the stem region of the experiment, where they are expanded vertically by cylindrical lens telescopes to ~ 1.25 cm [84] in order to engulf the molecular beam at the stem. The light then reaches the molecular beam where it is reflected by a pair of mirrors so that the lasers pass through the molecular beam 7 times. Between each pass the lasers go through a $\lambda/4$ waveplate, then are reflected through that same waveplate so that the polarization alternates between \hat{y} and \hat{x} polarization with each pass, which is necessary to ensure that there is no dark state within the $|X, J = 2^+\rangle$ or $|X, J = 1^-\rangle$ manifolds.

WORK TO UPGRADE THE ROTATIONAL COOLING LASERS

During the ACME II measurement Daniel Ang and I spent significant time and effort on upgrading the rotational cooling lasers so that they would be more robust than the home built ECDLs we ultimately used. The idea behind this project was based on the work Cristian Panda had already done in developing the lasers used for the state preparation STIRAP in ACME II [86], which had demonstrated the ability to lock a laser for up to a week, as opposed to the typical ~ 1 day of locking we could achieve with the home built ECDLs and the transfer cavity method. This method relied on two upgrades, the first of which was using a commercial ECDL⁵ with a high quality AR coated laser diode diode⁶ with a much larger mode-hop free tuning range (in practice ~ 10 GHz at 690 nm) than the home built ECDLs.

⁵Toptica DL Pro

⁶LD-0695-0040-AR-1

The second improvement was referencing the laser frequency to ultra low expansion (ULE) glass cavity using the Pound-Drever-Hall (PDH) locking scheme[106]. The cavity consisted of a ULE glass spacer with one planar mirror and one concave mirror and had a finesse of $\sim 30,000$. The cavity was housed in a homemade vacuum chamber with the ability to temperature stabilize the cavity to minimize expansion or contraction of the ULE glass at the critical temperature of the glass. In this method an electro-optic modulator (EOM) is used to add 10 MHz sidebands to the laser carrier frequency, which is then aligned into the cavity. We detect the reflected light off of the cavity, and mix it with the original EOM modulation frequency to extract a DC component and a component oscillating at twice the modulation frequency, which we filter out. This provides an asymmetric error signal and we use a Toptica Digilock 110 module to provide proportional-integral-derivative (PID) feedback to the laser based on this signal, stabilizing both the laser current with a bandwidth of ≤ 5 MHz and the grating piezo with a slow bandwidth of ≤ 100 Hz. Tests performed by Cristian Panda demonstrated that this technique resulted in a laser linewidth below 150 Hz [88]. In order to control the laser frequency while locked to the ULE cavity, we put a double passed acousto-optic modulator (AOM) between the EOM and the cavity so that we can shift the laser output frequency relative to the cavity mode without disturbing the lock. Longer term drifts of the ULE cavity were corrected by an optical frequency comb⁷ referenced to a GPS stabilized signal.

In order to implement this system for rotational cooling, we needed to set up and lock two new lasers with this method, and ran into two challenges. The first challenge we found is that the available AR coated laser diodes could only provide up to 25 mW of power, which in an ideal system would have been sufficient to saturate our laser cooling transitions. However, unlike the lasers locked to the transfer cavities, in ACME II we were not able to lock to our ULE cavities in the room that housed our main apparatus, likely due to the large amounts of acoustic noise in the room, and

⁷Menlo Systems

instead housed our stable lasers in quieter room in a nearby building. Bundles of 100 m optical fibers were run between these two rooms by students working on the original ACME I experiment. This presented a significant issue, as the power lost over these long fibers meant that we did not have enough power at the experiment to saturate the laser cooling transitions. We initially considered using the ULE stabilized commercial lasers as a primary system to drive the home built lasers as a secondary system, but ultimately settled on amplifying the ULE stabilized lasers with a tapered amplifier providing up to 300 mW as we did for the STIRAP 690 nm laser.

The second challenge we encountered related to an unknown noise source in the PDH locking scheme which caused a drift in the background level of the error signal. While we could temporarily offset this drift when we initially locked the lasers, we found that over time this drift would unlock the laser on the timescale of a few hours to ~ 1 day, eliminating the primary benefit of this upgrade. One of the benefits of PDH locking is that it is generally insensitive to fluctuations in laser power, however it is sensitive to fluctuations in the laser phase, which can produce these changes in the error signal background level. While we did not successfully find the exact components that caused this issue, moving the lasers to a new table and realigning a number of the optics suppressed this drift. Ultimately we did not resolve this issue before the ACME II measurement began its final data run, so this upgrade was not implemented for ACME II, but will be implemented for ACME III.

3.2.3 STIRAP LASERS

The STIRAP laser system required two lasers, the pump laser at 690 nm requiring $\sim 15 - 20$ mW, and the Stokes laser at 1090 nm requiring < 10 W in order to saturate[86]. Both lasers were Topica DL Pro systems stabilized with the ULE cavity PDH locking scheme described in the previous section. In order to achieve the powers needed for STIRAP, both lasers needed to be amplified. The 690 nm laser used the same tapered amplifier described in the previous section to achieve a power

of 300 mW. The 1090 nm laser was amplified with a commercial fiber amplifier⁸ to 10 W, which as described in section 2.3.4 was not sufficient to fully saturate the STIRAP transition, but was the highest power easily available at the time. Both lasers are coupled into single mode fibers after amplification to clean up the spatial mode, and expanded before being focused into the state preparation region.

In order to select out a particular $\tilde{\mathcal{N}}$ during the STIRAP process we shift the two-photon resonance by tuning the 1090 nm Stokes laser. We also shift the Stokes laser frequency to shift the two-photon resonance to adjust from the different Stark shift when operating at $|\mathcal{E}| = 80$ V/cm and $|\mathcal{E}| = 140$ V/cm. This is done after the first fiber amplifier with a series of AOMs, where the laser carrier frequency is initially red detuned from both $\tilde{\mathcal{N}}$. The laser passes through four AOMs in series, where the first order deflected beam from each AOM is split off and recombined with a series of 50/50 beam splitters. This switching scheme allows us to select one out of four possible frequencies by turning on a single AOM and leaving the other three turned off. However, each AOM switch is only roughly $\sim 80\%$ efficient, and the two 50/50 beam splitters each beam passes through, followed by a new fiber coupling results in a total power loss of roughly $\sim 6 - 7\times$. In order to counteract this, the Stokes laser is amplified with a second 10 W fiber amplifier in the same room as the experiment to prevent further fiber coupling loss over long distances.

The two STIRAP lasers were launched vertically through the experiment apparatus and so a pair of structures were built. The first structure supported the laser breadboards, and was attached to the floor of the lab. The second structure, known as the “laser lounge” was a raised platform connected to the lab walls so that adjustments could be made to the STIRAP optics safely and comfortably. This platform also allowed vertical access to the experiment for magnetometry measurements. The two structures and the experiment were all kept separate mechanically to isolate vibrations.

⁸Nufern PSFA-1084-01-10W-1-3

However, as discussed in section 5.2, the new lab space for ACME III does not have an equivalent support structure to the laser lounge, so work is required to determine the best way to mount the ACME III *X-C-H* STIRAP lasers.

3.2.4 THE REFINEMENT AND READOUT LASER

Both the refinement laser and the readout laser for ACME II were provided by a single laser system. This system consisted of a a Toptica DL Pro ECDL at 703 nm locked to the same ULE cavity as the STIRAP lasers, and a high powered Ti:Sapph laser referenced to the ECDL. This was necessary as the ECDL only provided ~ 40 mW, while the refinement beam ultimately needed ~ 800 mW of power to suppress errors in the spin aligned state, and the readout beam needed ~ 100 mW to saturate the $X \rightarrow I$ transition. The Ti:Sapph laser⁹ was pumped by an 18 W pump laser¹⁰, and produced a typical output of $\sim 4 - 5$ W during the ACME II measurement¹¹.

The locking mechanism between the Ti:Sapph and the ECDL, originally implemented by Cristian Panda[88], and based on [107], relies on a beat note between the two lasers, created by overlapping the two beams on an avalanche photodiode. This signal is then amplified by an RF amplifier with a 2 GHz bandwidth¹², then mixed with the output of a custom made RF chirp generator¹³ to produce a signal that oscillates at the difference frequency between the lasers minus the frequency of the chirp generator, while the higher frequency beat component is filtered out. This lower frequency signal is then split evenly between two lines, one of which travels through a 1 m delay line

⁹M Squared SolsTiS

¹⁰Lighthouse Photonics Sprout G18

¹¹This power has decayed over time since the ACME II measurement and we have not reached 5 W since then. This decay is at least in part related to a surface within the proprietary laser head which needs to be cleaned every 1 - 2 years of use, however service visits have not been able to reach the 5 W we originally saw, and the laser is only specified for 3 W.

¹²ZFL-2000+

¹³This system was built by Jim McArthur at Harvard University.

before the signals are mixed back together. After filtering out the high frequency beat signal, we are left with a sinusoidal error signal. This error signal is sent to the Ti:Sapph control system and amplified before being used to provide PID feedback to a piezo in the Ti:Sapph cavity, with a bandwidth of up to 100 kHz[88]. By sweeping the frequency of the chirp generator, we can sweep the carrier frequency of the Ti:Sapph laser, with a demonstrated range of 400 MHz in 5 ms while maintaining the lock. This technique was used to switch the $\tilde{\mathcal{N}}$ addressed by both the cleanup and readout lasers, and as the frequency could be selected arbitrarily within a large enough range this system was also used for adjusting the lasers when the magnitude of $|\mathcal{E}|$ was changed.

Once the Ti:Sapph was locked it was split into two paths using a polarizing beam splitter (PBS), with roughly 2.5 W in the refinement path and roughly 1.5 W in the readout path. The readout path included a two state AOM switch like was used for the STIRAP Stokes laser, where the two states corresponded to the different $\tilde{\mathcal{P}}$ switch states and changed the addressed parity of the I state by shifting the laser frequency by roughly 90 Mhz. To provide independent tuning and compensate for the offset in the carrier frequency that this switching requires, the refinement laser also passed through an AOM before going to the experiment. After this the lasers were fiber coupled and sent to an optics table next to the interaction region windows.

Once on this table the refinement laser was sent through a PBS to clean up any polarization imperfections then sent through a $\lambda/2$ waveplate mounted in a digitally controlled rotation stage. The beam was then expanded and collimated vertically by a pair of cylindrical lenses before entering the interaction region. The readout laser, in contrast required additional AOM switches to create the polarization switching used to project out the measured phase. In order to do this, one of the two AOM deflected paths passed through a $\lambda/2$ waveplate before the two beams were recombined on a high extinction ratio PBS. The overlapped beams then passed through a $\lambda/2$ waveplate in a rotations stage before being expanded and sent into the detection region window. This waveplate allowed us to perform the $\tilde{\theta}$ switch and the $\tilde{\mathcal{R}}$ switch. It was important to ensure that the power in

each polarization state was identical to the other, as the AOM switches could produce asymmetries between the two beams, so this power was regularly monitored.

3.3 ELECTRIC AND MAGNETIC FIELD CONTROL

In order to both control the experiment state and measure and suppress systematic errors it is necessary to both control and monitor the electric and magnetic fields of the experiment. We will first discuss the magnetic field measurement methods and the control systems, which consist of a combination of passive shielding and three systems of magnetic field coils. We will then discuss the details of the electric field plates, and the method of measuring the electric field along the beam axis.

3.3.1 MAGNETOMETRY

The critical magnetometry in ACME II was performed with Bartington Mag-03 three-axis fluxgate magnetometers. These fluxgate sensors have small internal electronic offsets, so in order to accurately read out the magnetic field we mounted the sensors on a Zaber rotation stage and rotated the sensor about one axis by 180° or 90° to calibrate out these offsets with every measurement. In practice we mostly rotated by 90° as this allowed us to both measure these offsets and acts as a cross reference between the two axes which are effectively switched. It should be noted that this method only allows us to distinguish two axes of the true field from the offsets, and we must use a sensor which rotates about one of the other two axes in order to accurately measure all field components.

As an upgrade over ACME I, the ACME II interaction chamber included four “magnetometer pockets”, which consisted of welded tube flanges on the chamber walls that allowed a magnetometer in atmosphere to be inserted into the tube to a location just outside of the electric field plates in order to give us the best measurement sensitivity to the fields that matter without disrupting the experiment. We can insert the fluxgate sensors into these pockets attached to a drive shaft so that

they can be rotated from outside of the magnetic shielding. In this configuration the angle of the fluxgates is constrained by the fitting with the tube, which has a 2 mm clearance, which we estimate can result in a ~ 5 mrad, which is less than the specified orthogonality of the sensors of ~ 9 mrad. It is important to minimize these imperfections as they can cause significant offsets and mixing of the sensor axes can add noise to the measurements. For a full derivation of the ways in which these imperfections can contribute to our measurements see chapter 5 of Zack Lasner’s thesis[84]. Two of these pockets were located on the top surface of the chamber and allowed us to insert magnetometers in the $+\hat{y}$ direction, located at the positions $x = -5.5$ cm, and $z = \pm 8.9$ cm, where the origin is defined as the center of the molecular beam at the midpoint between the refinement and readout lasers. These pockets extended below the center of the molecular beam along \hat{y} , and the two sensors were able to be scanned for ± 7 cm along \hat{y} with a shared translation stage. The other two pockets allowed us to insert the magnetometers along $\pm z$, with one pocket on the two sides of the interaction chamber at the position $x = y = 0$ cm. These two pockets were designed so that the sensors could be placed as close as possible to the center of the interaction chamber at $z = \pm 20$ cm. Before the ACME II final data run began, Cristian Panda and Zack Lasner discovered that the fluxgates could get “stuck” in a configuration in which they broadcast extra noise at ~ 80 kHz that was picked up by the data acquisition system, so we opted to keep the fluxgates turned off during data collection, and measured the magnetic field at the start and end of each day during the measurement run. For these sensors we used the naming convention of FGYW, for FluxGate Y West, or FGXE, for FluxGate X East.

PROBULATION

In order to measure the magnetic field in the interaction region with as much accuracy as possible, we measured the magnetic field with a fluxgate sensor inserted along the beam axis $-\hat{x}$, which we call “probulation”. Since this process required us to insert the sensor between the electric field

plates, this process was performed before the ACME II measurement began, and once again when the ACME II data collection was complete, as we did not expose the field plates to atmosphere at any point between starting the measurement and the end of data collection due to concerns that this may cause the patch potentials on the field plates to change, shifting the non-reversing electric field. During this process we also insert the two fluxgate sensors into the two vertical pockets to calibrate the field along \hat{x} . The probulation fluxgate was housed in a rigid G10 tube to prevent angular deflections, which was aligned with the electric field plates and the beam axis using a theodolite¹⁴ mounted at the end of the beamline. The magnetometer could then be scanned ± 12 cm along the \hat{x} direction. With these scans we could measure the field along all coordinate axes, and enough of the gradients to extract those we could not measure directly using Maxwell's equations.

We perform this measurement by both rotating the magnetometers, and also by applying and reversing a magnetic field. This allows us to decompose the measurement into the same parity basis developed in section 2.4

$$\mathcal{B}_i = \mathcal{B}_i^{\text{nr}} + \tilde{\mathcal{F}}\mathcal{B}_i^{\mathcal{F}} + \tilde{\mathcal{B}}\mathcal{B}_i^{\mathcal{B}} + \tilde{\mathcal{F}}\tilde{\mathcal{B}}\mathcal{B}_i^{\mathcal{FB}}, \quad (3.7)$$

where here the subscript $i = \hat{x}, \hat{y}, \hat{z}$ represents the coordinate axis, $\tilde{\mathcal{F}} = \pm 1$ represents the rotation state of the magnetometers, and $\tilde{\mathcal{B}}$ represents the state of whatever magnetic field is applied, not just \mathcal{B}_z in this case. From this formulation we can extract both the ambient field, which dominates the term $\mathcal{B}_i^{\mathcal{F}}$, as it does not reverse with the applied field, and the applied fields, which dominate the channel $\mathcal{B}_i^{\mathcal{FB}}$. The other two parity channels are still of value, where $\mathcal{B}_i^{\text{nr}}$ is dominated by the electronic offsets of the sensors, and $\mathcal{B}_i^{\mathcal{B}}$ can provide information about imperfections with the rotation of the sensors.

¹⁴With ~ 10 arcseconds accuracy, this does not limit the alignment of the magnetometer at all.

3.3.2 MAGNETIC SHIELDING

In order to minimize the ambient field inside the interaction region we surrounded the interaction region chamber with five layers of cylindrical μ -metal¹⁵ shields with endcaps. As these shields were originally designed for the ACME I experiment, and only slightly modified for ACME II, see section 5.2 for a more in depth discussion of the design criteria for magnetic shields. Each layer of the shields consisted of a pair of 1.6 mm thick hemispherical μ -metal sheets, which were bolted to circular endcaps at the upstream and downstream ends of the interaction region and bolted to each other. These endcaps were mounted to an 8020 structure to fix the shields in place and the attachment between the two hemispheres minimized possible gaps in the μ -metal while providing additional structural support. For a diagram of the structure of the shield design and mounting see figure 3.2. The outermost shield layer had an outer diameter of 107 cm and was 130 cm long, while the innermost shield layer had an inner diameter of 76 cm and was ~ 88 cm long.

During the course of a single block measurement we must reverse the applied magnetic field \vec{B}_z within the shields, which we know caused the ambient field inside the shields to be magnetized up to the $\sim 100 \mu\text{G}$ level [84]. In order to reduce this we perform a “degaussing” procedure every time we switch the applied field. This process consisted of intentionally applying an oscillating magnetic field to the shields which is then ramped down to zero. This applied field forces the magnetization of the μ -metal to follow a hysteresis loop, and ultimately brings the magnetization near zero. In practice we performed this with a pair of 50 turn ribbon-cable coils wrapped around each of the 10 individual shield hemispheres. We typically applied a pulse lasting 1 s at 100 Hz, which was chosen in order to minimize the time needed for degaussing as this process and the settling time allowed for the magnetic field after the switching were the largest duty cycle loss in our switch structure. In hindsight, we determined that this frequency was not optimal for our μ -metal shields as the skin

¹⁵Amuneal Inc.




Figure 3.2: The ACME II Magnetic Shields. The ACME II magnetic shields consisted of five layers of paired half cylinders that were mounted on fixed endcaps. The endcaps, which were supported by an 8020 frame, were the only supported locations, with the endcaps supporting all of the weight of the cylinder. The two half cylinders were connected by overlapping their lips and fastening them with brass screws.

depth of the material, which can be given as $\delta = \sqrt{2\rho/(2\pi f\mu)}$, where ρ is the resistivity, and μ is the permeability, which for our μ -metal at 100 Hz is only roughly $\delta \approx 150 \mu\text{m}$ [108].

The ambient field within the magnetic shields was measured before the start of the ACME II measurement, and found to be at or below $|\mathcal{B}^{\mathcal{F}}| \approx 50 \mu\text{G}$ [84]. Additionally, based on measurements performed by Emil Kirilov when the shields were first built, we estimated the transverse shielding factor of the shields to be $\sim 10^5$ [75], which suggested that the contribution from the Earth's field to our ambient field was roughly an order of magnitude below this measured field. Additionally, we discovered that the ambient field increased significantly between the initial measurements and the beginning of the ACME II final data run, where the field had reached the level of $\sim 200 - 300 \mu\text{G}$. Upon discovering this increase, we spent a few weeks investigating the source of this change, and by removing individual shield layers determined that this change was most likely due to the magnetization of the innermost shield layer. Through this process we also determined that the outermost two layers of magnetic shielding had very little effect on the ambient field, however we continued to operate with all five layers in place for ACME II.

During this investigation we also discovered that some of the materials in the system were more magnetic than we had previously thought. In particular we discovered that some of the screws used to mount the shield layers had been mixed up with stainless steel screws instead of the brass screws we intended to use, and we replaced them. More critically, we found that the connectors used on the degaussing ribbon cabled had become magnetized over time, so we opted to individually solder the connections between each of the ribbon cable wires to remove any connectors inside the shields. We learned an important lesson through this, that while it is valuable to measure the magnetism of components we want to use, we also need to make sure we test them by intentionally magnetizing them first as well.

Ultimately, although we were able to determine that the increased ambient field was due to the innermost shield layer, we were not able to completely suppress this, even with degaussing at a much

lower frequency to account for the skin depth of the material. It is likely that the combination of stress from the mounting of the shields and the coarse handling methods required for assembly of the shields contributed to the magnetization over time. This seems particularly likely, as during the setup for probulation after the ACME II measurement, we observed that by bumping the shields support with the probulation table the ambient field jumped, as observed with the sensors in the magnetometer pockets. We expect that we could have returned to the prior ambient field levels by reannealing the μ -metal, however this would have required sending the shields out to an outside company, and we were able to operate with the larger ambient field for the duration of ACME II. This problem was a significant motivation behind the designs of the ACME III magnetic shields which are discussed in section 5.2.

3.3.3 MAGNETIC FIELD COILS

The primary applied magnetic field \mathcal{B}_z was applied by a combination of a cosine-theta coil and a pair of auxiliary “side” coils to increase the homogeneity of the applied field. This pair of coils was also used to apply the first order gradient $\partial\mathcal{B}_z/\partial x$. These coils, along with the coils for applying \mathcal{B}_x and the gradient $\partial\mathcal{B}_x/\partial x$, were wound on a cylindrical plastic structure between the interaction region chamber and the innermost magnetic shield. In addition to these coils we also used two pairs of coils, which were mounted to the top and bottom surfaces of the interaction region chamber in order to apply \mathcal{B}_y and the gradients $\partial\mathcal{B}_y/\partial x$, $\partial\mathcal{B}_y/\partial y$, and $\partial\mathcal{B}_y/\partial z$. With these coils all possible coordinate directions and all possible gradients could be applied in the interaction region. We used probulation to calibrate the applied currents to each of the coils.

During the initial probulation measurement we determined that the most significant gradient in the experiment was $\partial\mathcal{B}_z/\partial z \approx 4 \mu\text{G}/\text{cm}$ [84]. However, we applied compensating gradients to reduce all measured gradients to at most $\sim 1 \mu\text{G}$, which were kept constant throughout the ACME II measurement. As we will see in section 3.4, gradients in the magnetic field did contribute to our

systematic error budget, but at these levels they were not the limiting systematic errors [8].

An external system of “room coils” were also built as an undergraduate project by Zack Soule and Aaron Markowitz. These coils, which ran along the edges of the lab walls, were designed to minimize the effect of the Earth’s field in the interaction region. In practice, they were able to decrease the external field by roughly a factor of $\sim 5 \times$ [84]. These coils were ultimately not necessary for the ACME II measurement and most of ACME II occurred with these coils off, as the magnetic shields provided sufficient shielding and the magnetization of the shields dominated the ambient field in the interaction region.

3.3.4 ELECTRIC FIELD MEASUREMENTS AND \mathcal{E}^{nr}

During both ACME I and ACME II, multiple significant sources of systematic error coupled into the experiment through the non-reversing electric field \mathcal{E}^{nr} . Significant changes were made to the construction of the electric field plates between the two experiments in order to suppress these effects, which were discussed in section 3.1.4, but in order to quantify these effects it was necessary to measure the electric field. We measured the electric field in the interaction region with a modified version of the ACME measurement using a resonant microwave pulse, which was a method originally developed for ACME I [7]. We initially prepare the molecules in the H state with a spin aligned at $\pi/4$ relative to the \hat{x} axis. As the molecular beam passes through the precession region, a microwave pulse polarized along \hat{y} with frequency $2\pi \times 40$ GHz, which is near resonance for the zero field $|H, J = 1, M = \pm 1\rangle \rightarrow |H, J = 2, \tilde{\mathcal{P}}\rangle$ transition. By tuning the microwave pulse to act as a π pulse, for molecules on resonance we can empty out the polarization quadrature \hat{Y} , and use our standard detection method to compute the asymmetry. By scanning the detuning of the microwave resonance, we can then map the asymmetry into the Stark shift experienced by the molecules. We additionally are able to resolve the molecule arrival time, and with the assumption that slower moving molecules arrive at the detection region later, can map the arrival time onto the

position along \hat{x} where the molecules interacted with the microwave pulse. By combining these two factors we can extract a measurement of the electric field with field resolution better than 1 mV/cm [7], and with a spatial resolution on the order of $\sim 1 - 2$ cm.

With this method, we can perform the experiment while varying the switches $\tilde{\mathcal{N}}$, $\tilde{\mathcal{E}}$, and $\tilde{\mathcal{L}}$, and we can compute parity sums from these switches to better understand the behavior of the electric field. In addition to this, we were able to measure gradients in the electric field along \hat{y} by blocking half of the molecular beam either by blocking half of the STIRAP beams for $\partial\mathcal{E}/\partial z$, or by blocking half of the readout laser for $\partial\mathcal{E}/\partial y$. From this we extracted a measurement of \mathcal{E}^{nr} , which we determined to be $|\mathcal{E}^{\text{nr}}| \leq 5$ mV/cm, and that it was the largest at the center of the interaction region [84]. We also determined that the gradient $|\partial\mathcal{E}^{\text{nr}}/\partial y| \leq 1$ mV/cm², however we found that $|\partial\mathcal{E}^{\text{nr}}/\partial z| \leq 10$ mV/cm². The consequences of these non-reversing field values are discussed in the following section.

3.4 SYSTEMATIC ERRORS

As we expect the shift from a non-zero EDM to be proportional to $\vec{\mathcal{E}}_{\text{eff}}$, the true value of the EDM should reverse with both the $\tilde{\mathcal{N}}$ and $\tilde{\mathcal{E}}$ switches, and should reverse with no other switches. With this in mind, we measure the EDM-correlated frequency $\omega^{\mathcal{N}\mathcal{E}}$ while varying over 40 different experiment parameters [8] in order to search for any correlated shifts in the EDM value. In addition to this we also measure other parity channels for both ω and \mathcal{C} in order to look for unexpected shifts in those channels, which might leak into $\omega^{\mathcal{N}\mathcal{E}}$.

For a given parameter P , we vary P by a factor ΔP , typically $10\times$ larger than it's typical value, only using smaller variations when limited by our ability to vary the parameter in a few cases. We

then assume a linear dependence on P ¹⁶ [7, 109, 31] in order to calculate the slope

$$S_P = \partial\omega^{\mathcal{N}\mathcal{E}}/\partial P. \quad (3.8)$$

In cases where we observe a non-zero slope, or in cases where we expect to see a non-zero slope, we also measure the typical value of the parameter during normal data collection, \bar{P} , and use this to compute the shift to the EDM frequency

$$\omega_P^{\mathcal{N}\mathcal{E}} = \bar{P}S_P. \quad (3.9)$$

We can then also compute the uncertainty $\delta\omega_P^{\mathcal{N}\mathcal{E}}$ using standard error propagation techniques as

$$\delta\omega_P^{\mathcal{N}\mathcal{E}} = \sqrt{(\bar{P}\delta S_P)^2 + (\delta\bar{P}S_P)^2} \quad (3.10)$$

where $\delta\bar{P}$ and δS_P represent the uncertainty in their respective quantities. Non-zero shifts are then subtracted from our measured EDM frequency to give the reported EDM frequency. In cases where we did not observe a non-zero slope, we do not compute a shift, but for some cases we do compute the uncertainty by assuming $\delta\bar{P} \approx \bar{P}$. Similarly, for parameters which do not have an “ideal” value of \bar{P} , such as the refinement laser power, we do not compute a shift and only compute the uncertainty.

In order to determine which quantities to include our error budget, we use the same method as ACME [7, 6], which consists of three sets of criteria, referred to as class A, B, or C.

- **Class A** systematics are those for which a non-zero correlation of $\omega^{\mathcal{N}\mathcal{E}}$ with some parameter

¹⁶Where a non-zero slope is observed we measure the slope for multiple values of ΔP in order to confirm the linearity. We also have found that all of our models for these slopes exhibit linear dependence. However, it is still possible that some non-linear behavior is missed.

is found in our experiment. The measured shift $\omega_P^{\mathcal{N}\mathcal{E}}$ is subtracted and the uncertainty is added in quadrature with other errors. Additionally, we collected data with IPVs interleaved with normal data collection in order to monitor these shifts during the measurement. This data with IPVs was not averaged in with the final data set.

- **Class B** systematics are those for which we observed a signal in a channel other than $\omega^{\mathcal{N}\mathcal{E}}$ which we deemed important to understand, and that we did not understand, but which we did not observe to be correlated with $\omega^{\mathcal{N}\mathcal{E}}$. There were no such signals observed in ACME II, however this was not the case in ACME I [6], where an uncertainty was placed on $\omega^{\mathcal{N}\mathcal{E}}$ based on possible leakage into the EDM channel.
- **Class C** systematics were those where no correlation of $\omega^{\mathcal{N}\mathcal{E}}$ with a parameter was observed in our measurement, but such a correlation had been observed and not understood in a similar experiment. For these we include an uncertainty in our error budget as there may be related effects just below our sensitivity. For these parameters we consider the ACME I experiment [7, 6], as well as the YbF [109] and PbO [62] experiments.

Those parameters which made it into the ACME II error budget are listed along with their uncertainties and shifts in table 3.2. In this section we will provide an overview of each class A systematic included in the error budget, as well as how they were suppressed in ACME II. These systematics include magnetic field gradient effects ($\partial\mathcal{B}_z/\partial z$ and $\partial\mathcal{B}_z/\partial y$), STIRAP laser imperfections and finite refinement laser power ($\omega_{\text{ST}}^{\mathcal{N}\mathcal{E}}$ via $\theta_{\text{ST}}^{\text{H-C}}$, and $P_{\text{ref}}^{\mathcal{N}\mathcal{E}}$), AC Stark shift effects (\mathcal{E}^{nr}), correlated contrast effects ($|\mathcal{C}|^{\mathcal{N}\mathcal{E}}$ and $|\mathcal{C}|^{\mathcal{N}\mathcal{E}\mathcal{B}}$), and effects from an $\tilde{\mathcal{E}}$ correlated field ($\omega^{\mathcal{E}}$ via $\mathcal{B}_z^{\mathcal{E}}$). By definition, class C systematics included in our error bar did not have any observed non-zero correlation with the EDM channel. So absent a strong argument why they should start to show up in our measurements exactly at the sensitivity of ACME III, we will continue to study them, but do not yet expect them to be significant problems in ACME III.

Class	Parameter	Shift	Uncertainty
A	$\partial\mathcal{B}_z/\partial z$ and $\partial\mathcal{B}_z/\partial y$	7	59
A	$\omega_{\text{ST}}^{\mathcal{N}\mathcal{E}}$ (via $\theta_{\text{ST}}^{\text{H-C}}$)	0	1
A	$P_{\text{ref}}^{\mathcal{N}\mathcal{E}}$	–	109
A	\mathcal{E}^{nr}	–56	140
A	$ \mathcal{C} ^{\mathcal{N}\mathcal{E}}$ and $ \mathcal{C} ^{\mathcal{N}\mathcal{E}\mathcal{B}}$	77	125
A	$\omega^{\mathcal{E}}$ (via $\mathcal{B}_z^{\mathcal{E}}$)	1	1
C	Other magnetic field gradients (4)	–	134
C	Non-reversing magnetic field $\mathcal{B}_z^{\text{nr}}$	–	106
C	Transverse magnetic fields $\mathcal{B}_x^{\text{nr}}$ and $\mathcal{B}_y^{\text{nr}}$	–	92
C	Refinement and readout laser detunings	–	76
C	$\tilde{\mathcal{N}}$ -correlated laser detuning, $\Delta^{\tilde{\mathcal{N}}}$	–	48
Total systematic uncertainty		29	310
Statistical uncertainty			373
Total uncertainty			486

Table 3.2: ACME II Combined Error Budget for $\omega^{\mathcal{N}\mathcal{E}}$. Systematic and statistical error budget for $\omega^{\mathcal{N}\mathcal{E}}$ in the ACME II experiment [8]. All values are given for $\omega^{\mathcal{N}\mathcal{E}}$ in units of $\mu\text{rad/s}$.

3.4.1 A SHORT ASIDE ABOUT $\mathcal{N}\mathcal{E}$ -CORRELATED LASER DETUNINGS

Before describing any systematic error models in detail, it is helpful to first consider the method by which an $\mathcal{N}\mathcal{E}$ correlated laser detuning can occur as a result of the non-reversing electric field component \mathcal{E}^{nr} , as this effect will play a role in multiple systematic error models. We can represent the electric field as

$$\vec{\mathcal{E}} \cdot \hat{z} = \tilde{\mathcal{E}}\mathcal{E}^{\mathcal{E}} + \mathcal{E}^{\text{nr}}. \quad (3.11)$$

and the linear Stark shift, which depends on $\tilde{\mathcal{N}}$, as $\Delta_{\text{Stark}}^{\tilde{\mathcal{N}}} = -D_H|\mathcal{E}|\tilde{\mathcal{N}}$. If we substitute in our decomposed version of the electric field this becomes

$$\Delta_{\text{Stark}}^{\tilde{\mathcal{N}}} = -D_H|(\tilde{\mathcal{E}}\mathcal{E}^{\mathcal{E}} + \mathcal{E}^{\text{nr}})|\tilde{\mathcal{N}} = -D_H(|\mathcal{E}^{\mathcal{E}}| + \tilde{\mathcal{E}}\mathcal{E}^{\text{nr}})\tilde{\mathcal{N}} = -D_H|\mathcal{E}^{\mathcal{E}}|\tilde{\mathcal{N}} - D_H\mathcal{E}^{\text{nr}}\tilde{\mathcal{N}}\tilde{\mathcal{E}}. \quad (3.12)$$

In order to keep our lasers on resonance with the correct $\tilde{\mathcal{N}}$ we set the laser frequencies in our

system to take the form $f^{\mathcal{N}} = f_0 + \Delta f^{\mathcal{N}} \tilde{\mathcal{N}}$, where f_0 is the average laser frequency between the two $\tilde{\mathcal{N}}$ states, and $\Delta f^{\mathcal{N}}$ is the frequency applied to shift to the correct $\tilde{\mathcal{N}}$ state, which is typically set to closely match the Stark shift. In practice this results in a laser detuning, which we can group by dependence on $\tilde{\mathcal{E}}$ and $\tilde{\mathcal{N}}$,

$$\begin{aligned}
\Delta &= \Delta^{\text{nr}} + \Delta f^{\mathcal{N}} \tilde{\mathcal{N}} - \Delta_{\text{Stark}}^{\mathcal{N}} \\
&= \Delta^{\text{nr}} + (D_H |\mathcal{E}^{\mathcal{E}}| - \Delta f^{\mathcal{N}}) \tilde{\mathcal{N}} + D_H \mathcal{E}^{\text{nr}} \tilde{\mathcal{N}} \tilde{\mathcal{E}} \\
&= \Delta^{\text{nr}} + \Delta^{\mathcal{N}} \tilde{\mathcal{N}} + \Delta^{\mathcal{N}\mathcal{E}} \tilde{\mathcal{N}} \tilde{\mathcal{E}}.
\end{aligned} \tag{3.13}$$

From this we can now see that a non-reversing electric field component can cause an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ -correlated detuning in our lasers. Notably we have not described which laser systems are involved in this model yet because this effect occurs in all lasers which address transitions connected to one of the $\tilde{\mathcal{N}}$ states in our system. This includes the STIRAP lasers, the refinement laser, and the readout laser.

3.4.2 MAGNETIC FIELD GRADIENT EFFECTS

The first systematic effect that we will consider is related to non-reversing gradients in \mathcal{B}_z , specifically $(\partial \mathcal{B}_z / \partial z)^{\text{nr}}$ and $(\partial \mathcal{B}_z / \partial y)^{\text{nr}}$, for simplicity of notation in this section we will drop the superscript *nr* for the remainder of this chapter. We can model a systematic error in $\omega^{\mathcal{N}\mathcal{E}}$ due to these gradients as being caused by an $\mathcal{N}\mathcal{E}$ correlated shift in the center of mass of the molecular beam along the relevant axis for the gradient, $dz_{\text{cm}}^{\mathcal{N}\mathcal{E}}$ or $dy_{\text{cm}}^{\mathcal{N}\mathcal{E}}$. As the precession frequency in our experiment is dominated by Zeeman precession, such a shift would couple directly with a gradient to produce systematic shifts of the form

$$d\omega^{\mathcal{N}\mathcal{E}} = -\mu \partial \mathcal{B}_z / \partial z \cdot dz_{\text{cm}}^{\mathcal{N}\mathcal{E}} / \hbar \tag{3.14}$$

with an analogous shift for the gradient along \hat{y} . As the systematic models related to each of the gradients are analogous in this way, we will only describe the systematic effect related to a gradient along \hat{z} from here. We were able to develop models for two separate effects that could cause this shift in the center of mass of the molecular beam, one of which was related to a gradient $\partial\mathcal{E}^{\text{nr}}/\partial z$ ($\partial\mathcal{E}^{\text{nr}}/\partial y$ for $dy_{\text{cm}}^{\mathcal{N}\mathcal{E}}$), and one related to the non-reversing electric field \mathcal{E}^{nr} . Both of these effects occur in both the STIRAP lasers and the readout laser as a result of incomplete laser excitation.

Let us first consider the effect caused by the gradient $\partial\mathcal{E}^{\text{nr}}/\partial z$, which is represented schematically in figure 3.3. Once again for simplicity we will focus on describing the effect in the STIRAP lasers, chosen as this was the dominant source of this effect in ACME II [88]. When the STIRAP 2-photon detuning δ is non-zero, the state preparation efficiency η has a dependence on δ such that $\partial\eta/\partial\delta \neq 0 \propto \delta$. If there is also a gradient in the non-reversing electric field component $\partial\mathcal{E}^{\text{nr}}/\partial z$, this causes an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ -correlated two-photon detuning gradient along \hat{z} , $\partial\delta^{\mathcal{N}\mathcal{E}}/\partial z = D_H\partial\mathcal{E}^{\text{nr}}/\partial z$. When combined, these two effects result in a shift in the center of mass of the molecular beam, as STIRAP preferentially prepares molecules that experience a two-photon detuning closer to resonance, which we can represent as

$$dz_{\text{cm}}^{\mathcal{N}\mathcal{E}} \propto \left(\frac{\partial\eta}{\partial\delta}\right) \left(\frac{\delta^{\mathcal{N}\mathcal{E}}}{\partial z}\right) = \left(\frac{\partial\eta}{\partial\delta}\right) \left(D_H\frac{\partial\mathcal{E}^{\text{nr}}}{\partial z}\right). \quad (3.15)$$

Combining this with equation 3.14, we can see that

$$\omega_{\partial\mathcal{B}_z\partial z - \partial\mathcal{E}^{\text{nr}}/\partial z}^{\mathcal{N}\mathcal{E}} \propto \left(\frac{\partial\mathcal{B}_z}{\partial z}\right) \left(\frac{\partial\eta}{\partial\delta}\right) \left(\frac{\partial\mathcal{E}^{\text{nr}}}{\partial z}\right) \quad (3.16)$$

depends on three imperfections, the gradient $\partial\mathcal{B}_z/\partial z$, the two-photon detuning δ through $\partial\eta/\partial\delta$, and the gradient $\partial\mathcal{E}^{\text{nr}}/\partial z$. Analogous models exist for both gradients along \hat{z} and \hat{y} , for both the STIRAP lasers and the readout laser.




Figure 3.3: A Model for the $\partial B_z \partial z \times \partial \mathcal{E}^{\text{nr}} / \partial z$ Systematic. The gradient in $\partial \mathcal{E}^{\text{nr}}$ (blue arrows) results in a $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ -correlated shift of the STIRAP two-photon detuning. So long as $\delta \neq 0$ this produces a gradient in the STIRAP efficiency $\partial \eta^{\mathcal{N}\mathcal{E}} / \partial z$, which shifts the center of mass of the molecular beam by $d_z^{\mathcal{N}\mathcal{E}}$. This shift in the molecular beam couples to the gradient $\partial B_z / \partial z$ to produce an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ -correlated shift in the measured frequency. This figure was adapted from [8].

This model was verified by varying both an intentionally non-zero δ and deliberately applying a gradient $\partial B_z / \partial z$ and seeing agreement with our model. However, we did not have any means to apply an intentional $\partial \mathcal{E}^{\text{nr}} / \partial z$ in our system, so we compared the results of these tests with the value of this gradient measured as described in section 3.3.4. Additionally, this model was tested by comparing the detected fluorescence in the PMTs on the East and West sides of the experiment ($-\hat{z}$ and $+\hat{z}$). For a full accounting of these tests see Cristian Panda's thesis [88].

Let us now consider the effect caused by the non-reversing field component \mathcal{E}^{nr} , for which we will consider the case of the effect in the readout laser, where it was most prominent in ACME II[88]. Due to the diverging nature of our molecular beam, which we constrain with fixed collimators, there is a strong positive correlation between the position of a molecule along \hat{z} (and \hat{y}) relative to the center of the beam, and the transverse velocity of that molecule, which we can represent as $\partial v_z / \partial z$. This transverse velocity results in a Doppler shift for the laser resonance

$\Delta_z^{\text{Doppler}} = k_{\text{read}} v_z$. As a result, we have a positive correlation between the Doppler shift and the position of a given molecule, which we can represent as

$$\frac{\partial \Delta_z^{\text{Doppler}}}{\partial z} = \left(\frac{\partial \Delta_z^{\text{Doppler}}}{\partial v_z} \right) \left(\frac{\partial v_z}{\partial z} \right). \quad (3.17)$$

This results in a shift in the center of mass of the molecular beam that is excited by the readout laser if the readout laser detuning Δ changes by an amount $d\Delta$, which scales as

$$dz_{\text{cm}} = \frac{\partial z}{\partial \Delta_z^{\text{Doppler}}} d\Delta. \quad (3.18)$$

As we discussed in section 3.4.1, a non-zero \mathcal{E}^{nr} produces an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ -correlated detuning, which correlated this shift we just considered with the EDM channel. With this in mind we can represent this systematic shift as

$$\omega_{\partial \mathcal{B}_z / \partial z - \mathcal{E}^{\text{nr}}}^{\mathcal{N}\mathcal{E}} \propto \left(\frac{\partial \mathcal{B}_z}{\partial z} \right) \left(\frac{\partial z}{\partial v_z} \right) (\mathcal{E}^{\text{nr}}) \quad (3.19)$$

which depends on three quantities, two of which we can control.

This model was verified in a manner similar to the previous model, by varying the applied gradient $\partial \mathcal{B}_z / \partial z$ and the detuning of the readout laser Δ_{readout} . Additionally, unlike the previous model we were able to apply an intentional \mathcal{E}^{nr} , which further confirmed that the observed behavior matched our model. Once again, analogous models exist for both the magnetic field gradients along \hat{z} and \hat{y} for both the STIRAP and readout lasers. However, this effect was roughly $3 - 5 \times$ smaller than the effect related to gradients in \mathcal{E}^{nr} [88].

Once these systematic effects were understood, we were able to suppress them with a combination of tuning the magnetic field gradient and our laser detunings independently to minimize the slope of the systematic effect with variations in the other parameter. This allowed us to suppress the systematic to at least second order. Notably, as multiple otherwise independent effects depended on

the laser detunings δ and Δ_{readout} , we could not minimize these effects separately, and instead optimized our detunings to minimize the combined shift to the EDM channel. Additionally, we were able to suppress magnetic field gradients in the interaction region to below $\sim 1 \mu\text{G}/\text{cm}$ by applying offset gradients, and we continued to monitor these gradients on a daily basis as described in section 3.3.1. We also regularly measured \mathcal{E}^{nf} and its gradients in order to ensure that the systematic effects did not drift over time.

While these systematic effects were significantly smaller $\sim 2 - 3\times$ than the limiting systematic errors in ACME II, they will be of concern in ACME III if not managed appropriately. We are working to minimize \mathcal{E}^{nf} and its gradients for ACME III, but we do not have a method to significantly suppress these relative to their levels in ACME II. Our largest handle with which to control these effects will be our ability to control the magnetic field gradients in the experiment, which should be improved significantly with our new shield system and our new magnetic field coils, the details of which are discussed in section 5.2. Additionally, the electrostatic lens should improve the second effect we discussed in this section as the lens should decrease the correlation between molecule position and transverse velocity as it weakly focuses the molecular beam, and will slightly decrease the Doppler width of our molecular beam.

3.4.3 STIRAP LASER IMPERFECTIONS AND FINITE REFINEMENT LASER POWER

The next category of systematic effects we will consider includes two effects that depend on there being a finite laser power in the refinement beam. As described in section 2.3.4, after the STIRAP lasers prepare a spin aligned state along the vector \vec{S}_{ST} we use the refinement laser to reproject the spins along the vector \vec{e}_{ref} given by the polarization of the refinement laser. Without the refinement laser, an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ -correlated component of the STIRAP prepared phase, which we describe as $\theta_{\text{ST}}^{\mathcal{N}\mathcal{E}}$, would directly cause a systematic shift of the form $\omega_{\text{ST}}^{\mathcal{N}\mathcal{E}} = \theta_{\text{ST}}^{\mathcal{N}\mathcal{E}}/\tau$. We can then define an attenuation factor A_{ref} which describes the suppression of this $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ -correlated frequency due to the

STIRAP lasers as

$$A_{\text{ref}} \equiv \left(\frac{\partial \omega^{\mathcal{N}\mathcal{E}}}{\partial \omega_{\text{ST}}^{\mathcal{N}\mathcal{E}}} \right)^{-1}, \quad (3.20)$$

which will be necessary for describing these systematic effects. A_{ref} was experimentally verified to depend strongly on the laser power of the refinement laser, and to have a value $A_{\text{ref}} > 10^4$ for refinement laser powers above ~ 500 mW [88].

The first of the two systematic effects we will consider is caused by an ellipticity gradient in the STIRAP H - C Stokes laser, which is referred to in table 3.2 as $\omega_{\text{ST}}^{\mathcal{N}\mathcal{E}}$ (via $\theta_{\text{ST}}^{\text{H-C}}$). We control this ellipticity by tuning a $\lambda/2$ waveplate, which sets the angle between the initial polarization of the Stokes laser and the average birefringence axis of the combined optical system, which includes optics and the vacuum windows. Similar to the dominant systematic effect in ACME I, an ellipticity gradient introduced to a laser causes mixing between the intended bright and dark states, which couples with the AC Stark shift to cause phase evolution between the bright and dark states which may not be fully pumped out [6, 7]. In the presence of an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ -correlated two-photon detuning $\delta^{\mathcal{N}\mathcal{E}}$ (caused by \mathcal{E}^{nr}) this results in an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ -correlated frequency $\omega_{\text{ST}}^{\mathcal{N}\mathcal{E}}$, which is then suppressed by the refinement laser attenuation factor A_{ref} .

We measured this systematic regularly throughout the ACME II experiment by intentionally exaggerating the waveplate angle $\theta_{\text{ST}}^{\text{H-C}}$ and measuring the resulting EDM shift. This measurement was consistent with zero with the refinement laser turned on, as was the case during normal data collection, however this value and the associated uncertainty were included in the systematic error budget because this shift was observed to be non-zero with the refinement laser power set to zero. In ACME III we expect that this systematic effect will not become a new problem, as the available refinement laser power is already sufficient to suppress this below the level we are concerned about.

The second effect is caused by a non-zero angle $\theta_{\text{ST}}^{\text{ref}}$ between the STIRAP prepared phase \vec{S}_{ST} and the refinement laser polarization \vec{e}_{ref} . As the attenuation factor A_{ref} is dependent on the re-

finement laser power, if there is an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ -correlated power component $P_{\text{ref}}^{\mathcal{N}\mathcal{E}}$ then there will be an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ -correlated component of the attenuation factor $A_{\text{ref}}^{\mathcal{N}\mathcal{E}}$. It then follows that this produces an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ -correlated frequency shift $\omega_{P_{\text{ref}}^{\mathcal{N}\mathcal{E}}}$ so long as $\theta_{\text{ST}}^{\text{ref}} \neq 0$.

In ACME II we did not observe any evidence of a non-zero $P_{\text{ref}}^{\mathcal{N}\mathcal{E}}$ component, however in ACME I we observed evidence of an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ -correlated Rabi frequency component in the state preparation laser, which resulted in a signal in the $\omega^{\mathcal{N}\mathcal{E}\mathcal{B}}$ channel [7]. As a result of this we used the $\omega^{\mathcal{N}\mathcal{E}\mathcal{B}}$ channel and an intentionally applied $P_{\text{ref}}^{\mathcal{N}\mathcal{E}}$ to put an uncertainty limit on a possible $P_{\text{ref}}^{\mathcal{N}\mathcal{E}}$ in ACME II. Additionally, we were able to minimize this systematic effect by minimizing the angle $\theta_{\text{ST}}^{\text{ref}}$, and continued to monitor the systematic slope throughout the ACME II measurement. However, as we did not directly observe $P_{\text{ref}}^{\mathcal{N}\mathcal{E}}$, we only included an uncertainty in our error budget and not a shift. For ACME III it is likely that this systematic effect will still be significant, however we also expect our ability to place a limit on $P_{\text{ref}}^{\mathcal{N}\mathcal{E}}$ in the experiment to also decrease so long as a new source of $P_{\text{ref}}^{\mathcal{N}\mathcal{E}}$ does not appear in the experiment.

3.4.4 AC STARK SHIFT EFFECTS AND POLARIZATION GRADIENTS

As we have already discussed, in ACME I, the dominant systematic effect came from thermally induced ellipticity gradients in the preparation and readout laser beams [6, 7], which were significantly suppressed in ACME II by moving to field plates and windows made from higher quality glass (Corning 7980). These ellipticity gradients then coupled into the EDM channel through the same coupling with an AC Stark shift and the $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ -correlated detunings, $\Delta^{\mathcal{N}\mathcal{E}}$, which led to the $\omega_{\text{ST}}^{\mathcal{N}\mathcal{E}}$ systematic. Unfortunately, we discovered that the ACME II apparatus still had large enough gradients in the birefringence of the optics to produce a significant systematic effect, with the largest single uncertainty in our error budget, $140 \mu\text{rad/s}$.

Upon investigation we discovered that while the thermal stress induced gradients had been suppressed as expected, the remaining birefringence gradients were caused by stress in the glass of the

field plates and vacuum windows from their mounting. This was verified with the polarimeter developed by Vitaly Wirthl, which showed gradients at the $< 0.2\%$ /mm level for the combination of both field plates and both vacuum windows [101]. These measurements also determined that we had been fortunate enough that the birefringence axis of the glass was only $\sim 5^\circ$ offset from the \hat{x} axis, likely due to the rectangular symmetry of the glass. This systematic model was also verified by intentionally applying an \mathcal{E}^{nr} , and by repeatedly measuring \mathcal{E}^{nr} every two weeks during the ACME II final data run.

Notably, while the systematic shift was consistent with zero and had similar uncertainty to other systematics during normal operation, we found that the uncertainty grew significantly when we reversed the direction of propagation for the refinement and readout lasers, \vec{k} . We define the normal operation condition as $\vec{k} \cdot \hat{z} = +1$, and performed a reversal of $\vec{k} \cdot \hat{z}$ by creating a nominally identical optical setup on the opposite side of the experiment. However, we quickly discovered that the systematic slope $S_{\mathcal{E}^{\text{nr}}}$ was significantly larger in the $\vec{k} \cdot \hat{z} = -1$ configuration. This was most likely due to either a different alignment of the combined birefringence axis of the field plate and vacuum window glass or a larger overall birefringence gradient on that side of the experiment compared to the same pairing on the other side of the experiment. The readout and refinement lasers only pass through one window and one field plate before reaching the molecular beam, but our polarimetry system was only able to measure the combination of glass on both sides of the experiment so we could not verify this model exactly. As a result of this, only the $\vec{k} \cdot \hat{z} = +1$ data was included in the ACME II result and error budget.

As this systematic effect was the largest source of uncertainty in our error budget for ACME II, and depends only on \mathcal{E}^{nr} and the birefringence gradients of the glass, it is expected to pose a significant challenge for ACME III. Significant efforts have been undertaken by myself, Daniel Ang, and Peiran Hu to investigate alternative glasses to with lower stress-optic coefficients for ACME III. Work led by Peiran Hu has also developed new methods of mounting field plates to significantly

reduce residual stress on the glass. Work led by myself, Takahiko Masuda, Ayami Hiramoto, and Peiran Hu has also gone towards developing vacuum windows that will have lower birefringence gradients by using low stress-optic coefficient glass and reducing stress on the glass. The details of this work are discussed in more depth in sections 4.2.2 and 5.3.

3.4.5 CORRELATED CONTRAST EFFECTS

In ACME II we observed a significant correlation between $\omega^{\mathcal{N}\mathcal{E}}$ and the two contrast channels $|\mathcal{C}|^{\mathcal{N}\mathcal{E}}$ and $|\mathcal{C}|^{\mathcal{N}\mathcal{E}\mathcal{B}}$. We can understand the mechanism of this systematic effect by considering how we compute the measured frequency $\omega^{\mathcal{N}\mathcal{E}} = (\mathcal{A}/(2\mathcal{C}))^{\mathcal{N}\mathcal{E}}/\tau$. We can make the assumption that for a given channel u , $|\mathcal{C}|^u \ll |\mathcal{C}|^{\text{nr}}$, and Taylor expand $\omega^{\mathcal{N}\mathcal{E}}$ to find

$$\omega^{\mathcal{N}\mathcal{E}} = \frac{\mathcal{A}^{\mathcal{N}\mathcal{E}}}{2\tau|\mathcal{C}|^{\text{nr}}} - \omega^{\text{nr}} \frac{|\mathcal{C}|^{\mathcal{N}\mathcal{E}}}{|\mathcal{C}|^{\text{nr}}} - \omega^{\mathcal{B}} \frac{|\mathcal{C}|^{\mathcal{N}\mathcal{E}\mathcal{B}}}{|\mathcal{C}|^{\text{nr}}} - \dots, \quad (3.21)$$

where we have neglected to show any other components as their frequency components were much smaller than ω^{nr} and $\omega^{\mathcal{B}}$. The two frequency components included above are in fact quite large, as $\omega^{\mathcal{B}}$ is set by the Zeeman precession frequency, and ω^{nr} is set by the global polarization angle θ . The combination of these two values are intentionally chosen to maximize the sensitivity of our asymmetry. We determined that the average values of both $\langle |\mathcal{C}|^{\mathcal{N}\mathcal{E}} \rangle$ and $\langle |\mathcal{C}|^{\mathcal{N}\mathcal{E}\mathcal{B}} \rangle$ were consistent with zero in our experiment, but as we still saw large correlations with these values we included a shift and an uncertainty in the error budget.

In ACME III, while we expect the exact values of ω^{nr} and $\omega^{\mathcal{B}}$ to change, their combination is expected to shrink by a factor of $\sim 5\times$ as we will still select $\phi - \theta \approx \pi/4(2n + 1)$ in order to maximize the sensitivity of the asymmetry, but τ will be approximately $\sim 5\times$ longer, reducing ω^{nr} and $\omega^{\mathcal{B}}$ by a corresponding amount. In addition to this, as we have no current model to suggest that $|\mathcal{C}|^{\mathcal{N}\mathcal{E}}$ and $|\mathcal{C}|^{\mathcal{N}\mathcal{E}\mathcal{B}}$ should be non-zero in our experiment, unless we discover a new effect in ACME

III, we also expect the uncertainties we assign to $|\mathcal{C}|^{\mathcal{N}\mathcal{E}}$ and $|\mathcal{C}|^{\mathcal{N}\mathcal{E}\mathcal{B}}$ scale with our new sensitivity. As such, we do not expect this systematic effect to pose more of a challenge for ACME III than it did for ACME II, but it will still most likely be present.

3.4.6 EFFECTS OF AN $\tilde{\mathcal{E}}$ CORRELATED MAGNETIC FIELD

The final class A systematic effect included in our error budget came from $\omega^{\mathcal{E}}$, which can be caused by an $\tilde{\mathcal{E}}$ -correlated magnetic field, $\mathcal{B}_z^{\mathcal{E}}$. Similar effects have been seen in other EDM experiments, caused by leakage currents or geometric phase effects [58, 62, 16, 53]. However in our experiment these effects are significantly suppressed by the $\tilde{\mathcal{N}}$ switch. For large values of $\omega^{\mathcal{E}}$ though, there can still be leakage into the EDM channel due to a small electric field dependent difference between the g-factors of the two $\tilde{\mathcal{N}}$ states, which were measured by Nick Hutzler in ACME I [79]. If we decompose the g-factor as $g = g^{\text{nr}} + g^{\mathcal{N}}\tilde{\mathcal{N}}$, we can see that there will be an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ -correlated shift of the form

$$\omega_{\mathcal{B}_z^{\mathcal{E}}}^{\mathcal{N}\mathcal{E}} = -g^{\mathcal{N}}\mu_B\mathcal{B}_z^{\mathcal{E}} = (g^{\mathcal{N}}/g)\omega^{\mathcal{E}}. \quad (3.22)$$

While $\omega^{\mathcal{E}}$ was consistent with zero for the ACME II measurement, we still placed a limit on $\omega_{\mathcal{B}_z^{\mathcal{E}}}^{\mathcal{N}\mathcal{E}}$, as we were able to measure a non-zero $S_{\omega^{\mathcal{E}}}$, which is effectively a measure of the suppression due to the $\tilde{\mathcal{N}}$ switch. Additionally, we were able to place a limit on the possible leakage current in the experiment by treating the field plates as parallel plate capacitors and using the Stark shift in our molecules to determine how fast they discharge. For a full description of this method see Cristian Panda's thesis [88]. Ultimately we determined that this leakage current was ~ 25 pA, which was multiple orders of magnitude below the level that would impact our experiment. Based on this and the scale of the error associated with this effect in ACME II, we do not expect this effect to become significant in ACME III.

3.5 THE ACME II RESULT

The ACME II final data collection run took place from January 5 until March 30, 2018, and consisted of ~ 350 hours of data in the normal configuration. An additional ~ 150 hours of data collection was used for measuring and monitoring systematic errors during this time. This consisted of running the experiment roughly five days per week, 16 hours per day, with an 8 hour pause for warming up the beam box in order to remove neon ice buildup and subsequent cooling back to our operating temperature. This process required two shifts per day of monitoring the experiment, and is sufficiently long that we did not believe significant sensitivity gains could be reasonably realized by averaging longer.

Throughout the ACME II measurement the value of the measured EDM frequency was blinded by automatically adding an unknown but recorded offset to the value in all data analysis routines. In the final data analysis we used two independent data analysis routines, developed in separate software languages by Cristian Panda (in Mathematica) and Zack Lasner (in MATLAB)¹⁷. Once all of the data analysis was complete, and the two analysis routines were in agreement, we unblinded the data by removing the recorded offset. These processes were implemented in order to reduce any possible bias from the users during the analysis process and to ensure that we were not influenced by an unknown issue either with an analysis routine or a property of the language it was developed in.

The majority of the data fell within a Gaussian distribution, but we observed a significantly higher number of data points in the tails of the Gaussian distribution than we expected, which can be seen in figure 3.4. In order to properly account for these points, we investigated a number of different statistical methods. Ultimately we chose to perform an M -estimator analysis [110] on bootstrapped data [111, 112], for a full description of this analysis process see Cristian Panda and

¹⁷Additional analysis routines were developed by Daniel Ang (in Mathematica) and Jonathan Haefner (in C++), but these were not finalized in time for the final result and were not included, however they were consistent with the other two routines

Zack Lasner’s theses [88, 84]. This analysis needed to be performed independently for different values of $|\mathcal{B}_z|$ because, as we will discuss shortly, we observed different levels of noise at different magnetic field magnitudes. With this methodology we were able to construct 1σ confidence intervals corresponding to a 68% confidence level.

During the course of the ACME II measurement, we found that our measurement was limited by two sources of excess noise, which prevented us from reaching the shot-noise limit. In the following subsections we will discuss these noise sources, one of which was related to the magnetic field magnitude, while the other noise source was not well understood during the ACME II measurement.

3.5.1 EXCESS NOISE FROM TIMING JITTER

We determined that the scatter in the ACME II superbloc level data was significantly higher than the scatter in the data at the “group” level, which is constructed by combining polarization cycles within a single trace. We can characterize this excess noise by calculating the reduced chi-squared for the data, χ_r^2 . In the case of the excess noise which did not vary with the applied magnetic field magnitude, we found that $\chi_r^2 \approx 3$, but during the ACME II measurement could not track down the source of this noise. Shortly after the ACME II result was reported, work led by myself, Cristian Panda, and Mohit Verma determined that the noise was caused by a combination of timing noise in our DAQ and timing imperfections in the polarization switching scheme [69]. The exact details of this effect, which has now been well controlled are discussed in section 4.1.4, as the control of this effect effectively results in a signal gain of $\sim 3\times$ for ACME III.

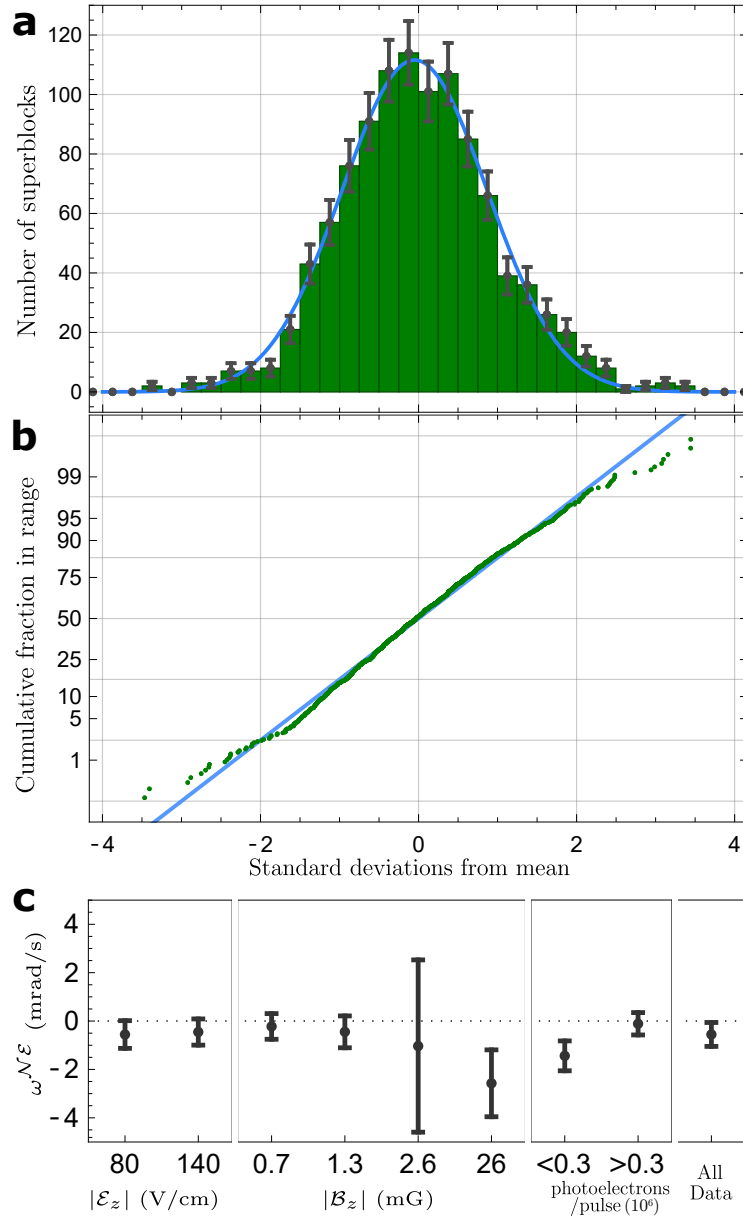


Figure 3.4: Statistics of the ACME II Measurement. **a)** Normalized histogram of the superblocks in the ACME II data set. The histogram has been centered and does not show the EDM value. Error bars represent the standard deviation in the bin expected for a Poisson distribution, while the blue curve is a Gaussian fit. **b)** Normal probability plot for the same data, where the blue line represents a normal distribution. Excess data in the tails of the distribution can clearly be seen as deviations from the normal distribution. **c)** Measured values of $\omega^{\mathcal{N}\mathcal{E}}$ for different experiment configurations. Error bars represent 1σ . This figure was adapted from [8].

3.5.2 EXCESS NOISE AT LARGE MAGNETIC FIELDS

We found that there was additional excess noise, characterized by $\chi_r^2 \approx 7$, which only occurred when we operated the experiment with $|\mathcal{B}_z| = 26$ mG. Through simulation and a direct measurement of the precession time τ ¹⁸ we determined that this was caused by fluctuations in the mean forward velocity of the molecular beam at the 0.05% level from shot-to-shot[8]. This effect can be understood by considering that the precession time τ is determined by a fixed distance L between refinement and readout lasers, so fluctuations in velocity will result in fluctuations in τ . We calculate τ by measuring the Zeeman precession phase, which we can write as $\phi^{\mathcal{B}} = -g\mu_B\mathcal{B}_z\tau = -g\mu_B\mathcal{B} \times \frac{L}{v}$. We can see how the standard deviation in the mean forward velocity σ_v impacts our uncertainty by writing

$$\sigma_{\phi^{\mathcal{B}}} \approx |g|\mu_B\mathcal{B}_z \frac{L}{v} \frac{\sigma_v}{v} = |\phi^{\mathcal{B}}| \frac{\sigma_v}{v}. \quad (3.23)$$

With this in mind, we can see that these velocity fluctuations will introduce excess noise into our experiment if this uncertainty is similar in scale to the statistical uncertainty in our measured asymmetry, as was the case in ACME II. However, we can control the scale of this uncertainty by changing the scale of the Zeeman precession, which we could easily do by operating the experiment at lower values of $|\mathcal{B}_z|$. Because of this, we chose to operate at $|\mathcal{B}_z| = 0.7, 1.3, 2.6$ mG for the remainder of the ACME II measurement, where this excess noise was found to not be a concern. This discovery served as the primary driver behind our projected magnetic field requirements in ACME III, as we believe we will need to operate at a lower magnetic field value of $|\mathcal{B}_z| \approx 100 - 200$ μ G in order to suppress this noise again. This requirement and its consequences are discussed further in section 5.2.

¹⁸This measurement was performed by creating a “notch” in the fluorescence signal by temporarily turning off the STIRAP lasers and measuring the time of flight. For more details on this see [88].

3.5.3 AN ORDER OF MAGNITUDE IMPROVED LIMIT ON d_e

The final result of the ACME II measurement was

$$\omega^{\mathcal{N}\mathcal{E}} = -510 \pm 373_{\text{stat}} \pm 310_{\text{sys}} \mu\text{rad/s}, \quad (3.24)$$

which we can use to compute the value of the electron EDM with the formula $d_e = -\hbar\omega^{\mathcal{N}\mathcal{E}}/\mathcal{E}_{\text{eff}}$ and the value $\mathcal{E}_{\text{eff}} = 78 \text{ GV/cm}$ to get

$$d_e = (4.3 \pm 3.1_{\text{stat}} \pm 2.6_{\text{sys}}) \times 10^{-30} e \text{ cm}. \quad (3.25)$$

This result has a combined statistical and systematic uncertainty ($\sigma_{d_e} = 4.0 \times 10^{-30} e \text{ cm}$) that is $12\times$ smaller than the ACME I result [6, 7], which was the previous best measurement of d_e , demonstrating more than an order of magnitude increased sensitivity to the electron EDM.

In order to put an upper bound on $|d_e|$, we use the Feldman-Cousins [113] prescription as was done in ACME I, which provides a rigorous method for determining whether to report a bound or a central value, and how to determine the confidence interval. This is particularly relevant as we can express the ACME II result as being $\langle |d_e| \rangle = 1.1 \times \sigma_{d_e}$. It should be noted that although our sensitivity is improved by more than an order of magnitude, the limit we set is also dependent on how far from zero our mean value of $|d_e|$ is, which can result in a larger bound than just the sensitivity would suggest. With a 90% confidence level we can then report

$$|d_e| < 1.1 \times 10^{-29} e \text{ cm}, \quad (3.26)$$

which is a factor of $8.6\times$ lower than the bound set by ACME I [6, 7].

*I'm a rattlesnake, babe, I'm like fuel on fire
So if you're gonna' get made don't be afraid of what
you've learned*

Blitzen trapper *Furr*

4

ACME III Development

SINCE THE CONCLUSION OF THE ACME II MEASUREMENT IN THE SPRING OF 2018, all efforts have been devoted to the development of a new generation of the ACME measurement, known as ACME III. Just as the ACME II measurement improved on the ACME I measurement sensitivity by an order of magnitude [8], we set out to improve on the ACME II sensitivity by an order of magnitude. This chapter provides an overview of all of the upgrades for ACME III, most of which have

been fully demonstrated at the time of this writing. I played a central role in the upgrades to increase the precession time of our measurement, to improve our detection systems, to reduce noise in our system, and to suppress systematic errors in our system.

Just as we described the upgrades over the ACME I experiment as belonging to three categories, so too will we describe the upgrades over the ACME II experiment. In section 4.1 we discuss all of the statistical upgrades to ACME III, however those related to the increased precession length of the experiment will also be discussed further in chapter 5. In section 4.2 we discuss upgrades intended to suppress known systematic error sources below the level required for the ACME III measurement. Finally, in section 4.3 we discuss general upgrades and changes to the experiment intended to improve the robustness of the experiment. It should be noted that the development of many of the upgrades discussed in this chapter were either a continuation or inspired by those upgrades for the ACME II experiment discussed in chapter 3, including those upgrades which were ultimately not implemented for that experiment.

4.1 ACME III STATISTICAL UPGRADE PATHWAY

In order to understand the statistical upgrades implemented for the ACME III measurement it is helpful to once again return to the figure of merit for the sensitivity of a shot-noise limited measurement

$$\delta d_e = \frac{\hbar}{2\mathcal{E}_{\text{eff}}\tau\sqrt{\dot{N}T}}. \quad (4.1)$$

Whereas in the case of ACME II, all statistical upgrades were based on increasing the photon count rate \dot{N} , in ACME III there are two notable upgrades that provide gains through other means, one that increases the precession time τ , and one that suppresses excess noise that in ACME II prevented us from actually reaching the shot-noise limit. These upgrades complement a series of upgrades to the photon count rate by increasing our molecular beam flux and improving our detection ef-

Upgrade	Demonstrated Count Rate Gain	Demonstrated δd_e Gain
Increased Precession Length	-	2.6
Electrostatic Lens	12	3.5
Improved Rotational Cooling	1.4	1.2
SiPM Detectors	2.3	1.5
Improved Collection Optics	1.7	1.3
Timing Jitter Noise Reduction	-	1.7
Ablation Target Load Lock	1.5	1.2
Projected Combined Gain		~ 43

Table 4.1: ACME III Demonstrated Statistical Gains. Table showing the gains associated with each individual statistical improvement in the ACME III experiment. Each gain factor has been individually verified, but actual combined signal level and EDM sensitivity in the ACME III apparatus have not yet been demonstrated at the time of this writing.

efficiency. Because of this it is helpful when considering the statistical gains in table 4.1 to wherever meaningful, break out both the gains to the photon count rate, and the corresponding δd_e improvement. However we do not show a combined count rate gain, as there is no longer a simple conversion to EDM sensitivity as there has been in the past.

Within this section we will individually describe each of the upgrades that were made in the development of the ACME III experiment. We will first discuss the increased precession length used in the experiment. We will then discuss the development of the electrostatic lens in the Q state, and the improved rotational cooling system, which was designed to make the lens more effective. We will then describe the new silicon photomultiplier (SiPM) detection system, and the re-optimized collection optics. Next we will describe the investigation into the excess noise found during the ACME II measurement and the methods used to suppress it. We will then describe the development of a new load-lock system to allow for the rapid changing of ablation targets. Finally, we will briefly discuss those statistical upgrades which were considered for the ACME III measurement, but which were not implemented.

4.1.1 INCREASED PRECESSION TIME

Both the ACME I and ACME II measurements used an $L = 22$ cm precession length, which with a typical molecular beam velocity of ~ 220 m/s corresponded to $\tau \approx 1$ ms [6, 7, 8]. This choice had been based on attempts to measure the lifetime made during the development of the original ACME experiment that suggested the radiative lifetime of the ThO H state was $\tau_H \gtrsim 1$ ms [78, 75, 92]. However, these measurements had all been performed inside of a buffer gas cell where collisions can introduce a significant alternative decay pathway, and multiple exponential decay constants were sometimes observed [92], without a conclusive explanation. The choice of $\tau \approx 1$ ms followed from this information, as there was little evidence that much sensitivity could be gained with a longer precession time, and a longer precession region presents a significantly more complicated engineering challenge.

After the completion of the ACME II measurement, in the hopes of finding an avenue to improve the ACME experiment, Cristian Panda, with assistance from Xing Wu and myself, performed a pair of test measurements using the ACME II apparatus and the setup on Beam Box II used for Q state measurements to see if we could perform a better measurement of the H state radiative lifetime in the beamline rather than a cell, where collisions would not be significant. This measurement suggested that the H state lifetime was closer to ~ 4 ms, significantly higher than previously thought, but was limited by the constraints of the existing apparatus, which had limited optical access, and with the apparatus available it was difficult to control systematic errors at the desired level. Based on these preliminary measurements, Daniel Ang and myself developed a specialized beam apparatus and performed a higher precision measurement that determined $\tau_H = 4.2 \pm 0.5$ ms [70], the details of which will be described in section 5.1.

Following this measurement it was clear that the ACME measurement could be extended to take advantage of our better understanding of the H state radiative lifetime. In order to do so we

needed to either significantly change the velocity of the molecular beam or build a new apparatus to accommodate the longer length. While methods exist in order to slow buffer gas cooled beams, such as the introduction of a two-stage cell rather than the single stage used in ACME, these methods do not provide sufficient slowing to fully take advantage of the radiative lifetime, and can come with more significant signal losses that would offset most of our gains [114]. As such, our best option was to build a new interaction region with a longer precession region, particularly as we already had other motivation to redesign our magnetic shields and our electric field plates to suppress magnetic field related noise and polarization gradient related systematic errors in ACME III. The design and development of this longer apparatus will be left to chapter 5.

In order to get the largest increase in EDM sensitivity it was important to consider three ways that the precession time contributes to δd_e . The first of these is quite clear from equation 4.1, that in converting from a measured phase to a frequency, the uncertainty in that frequency should scale inversely with τ . What is less clear initially, is that the photon count rate \dot{N} is also a function of τ , which contributes to \dot{N} in two ways. The first of these contributions comes from the fact that the population in the H state decays exponentially as a function of τ as $\dot{N} \propto e^{-\tau/\tau_H}$, simply due to the unavoidable radiative decay. If we combine these first two terms, which covers all of the effects which ignore engineering challenges, and effectively assume a perfectly collimated molecular beam we would find

$$\delta d_e(\tau) \propto \left(\frac{1}{\tau}\right) \left(\frac{1}{\sqrt{e^{-\tau/\tau_H}}}\right) = \frac{e^{\tau/(2\tau_H)}}{\tau}. \quad (4.2)$$

The second contribution to \dot{N} is harder to quantify analytically as a number of engineering decisions impact this, but is caused by the fact that the molecular beam is not perfectly collimated and is instead naturally diverging both over the distance leading to the precession region and in the precession region, and both the optics for ACME II and ACME III have a finite detection volume. If for now, we disregard the effects of the electrostatic lens, which weakly focuses the molecular beam

and also ignore the changes to the detection optics used in ACME III, this term would introduce a scaling term of $\dot{N} \propto 1/(l + \tau)^2$, where l is the distance from the beam source to the refinement laser, due to geometric expansion of the beam over a longer distance to give

$$\delta d_e(\tau) \propto \left(\frac{1}{\tau}\right) \left(\frac{1}{\sqrt{e^{-\tau/\tau_H}}}\right) (l + \tau). \quad (4.3)$$

The bound with perfect collimation shown in equation 4.2 and the bound with no lens or change to the detection optics in equation 4.3 act as limiting cases for the performance of the lens as shown in figure 4.1. In order to properly determine the gain in EDM sensitivity it was necessary to perform simulations of the electrostatic lens to account for both the effect of the weak focusing from the lens and the new detection optics. These simulations were performed by Xing Wu, and were also used to optimize the lens parameters once a precession length was chosen for ACME III. The results of these simulations are also shown in figure 4.1.

Due to the timelines involved with the development of the next generation of the experiment a precession time of $\tau \approx 5$ ms was chosen once preliminary results from the purpose-built lifetime measurement apparatus were available, and confirmed the results found by Cristian Panda. This value was chosen as a good value for τ both based on early simulations of the electrostatic lens and considerations for the engineering challenges associated building a longer interaction region in terms of both the systems themselves and the available lab space. From these considerations we chose a precession length of $\tau \approx 5$ ms, which corresponds to a precession length of $L = 100$ cm. With this choice of precession length, we expect to see an improvement in δd_e of approximately $\sim 2.6\times$.

4.1.2 THE ELECTROSTATIC LENS AND IMPROVED ROTATIONAL COOLING

While efforts to develop an electrostatic lens using the ground state X for ACME II ultimately were unable to overcome obstacles like the production of x-rays [103], new efforts led by Xing Wu to




Figure 4.1: EDM Sensitivity Gain as a Function of Precession Time. Projected EDM sensitivity relative to ACME II as a function of precession time for three different configurations based on the recent measurement of the H state radiative lifetime [70]. The red curve labeled “perfect collimation” is based on the estimate found using equation 4.2. The blue curve labeled “with lens collimation” corresponds to the projection found using the simulations of the electrostatic lens. The orange curve labeled “no lens” is based on the estimates found using the assumptions that the beam diverges as in ACME II with no electrostatic lens as represented in equation 4.3. The shaded regions of the curves represent the uncertainty from our uncertainty in the lifetime τ_H . The dashed lines indicate the precession time used for ACME II and the chosen precession time for ACME III. Figure adapted from [70].

study the ThO Q state made a new avenue possible. The Q state was found to be well suited to our purposes with the electrostatic lens for three reasons, which were all measured or verified in a recent study [87].

1. The first, which has already been discussed, is the large linear Stark shift of the $|Q, J = 2, M = 2, \Omega = -2\rangle$ state, with dipole moment $d_Q = 1.60(2) e a_0$. This Stark shift is nearly an order of magnitude larger than what can be found in the $|X, j = 2, M = 0, \Omega = 0\rangle$ state used in the previous electrostatic lens [85]. Based on this, larger trap depths can be achieved with lower applied fields, both improving the effectiveness of the lens and decreasing the likelihood of producing X-rays.
2. The next reason is that the radiative lifetime of the Q state is sufficiently long that there is no measurable decay out of the Q state during the transit through the 53 cm long lens. This recent measurement resulted in a lower bound on the radiative lifetime of $\tau_Q > 62$ ms [87].
3. The final reason is that the Q - C transition dipole moment is sufficiently large that we can reasonably perform STIRAP between X - C - Q in order to prepare molecules in the lens state with high efficiency, then return them to the ground state. This value has been measured to be $d_{Q-C} = 0.397(47) e a_0$, and both one and two stage STIRAP has been performed with a one (two) stage efficiency of 90% (80%) [87]. A detailed description of this STIRAP process can be found in section 2.3.3 and [87].

Notably, the Q state is also well suited to use in a magnetostatic lens using a Halbach array due to the large molecule-frame magnetic dipole moment $\mu_Q = 2.07(11)\mu_B$ [87, 84]. Both electrostatic and magnetostatic lenses were considered for ACME III, but an electrostatic method was chosen, primarily due to the tunability of the lens through the electrode voltage, which is not available in a magnetostatic system relying on permanent magnets.

Based on the measurements of the Q state, and significant simulation work, an electrostatic lens apparatus has been designed, constructed, and tested [85]. This work has been led by Xing Wu, with significant help from Peiran Hu and Zhen Han. As previously described in section 2.3.3 and shown in figure 2.7, the lens consists of six 53 cm long, 19 mm diameter cylindrical electrodes, spaced out at a radius of 19 mm. We apply a DC voltage $\pm V$ to the electrodes, alternating polarity for each electrode in order to produce the desired electric field. This configuration produces an electric field magnitude

$$\mathcal{E}(r) = \left(\frac{3V}{R^3}\right) r^2 \quad (4.4)$$

where r is the radial position, and R is the radial distance from the center of the lens to the surface of each electrode (19 mm). This results in a linear restoring force due to the linear Stark shift given by

$$\mathbf{F}(r) = -\frac{d}{dr} (-\mathcal{D}\mathcal{E}(r)) \hat{r} = \mathcal{D} \left(\frac{6V}{R^3}\right) r \hat{r}, \quad (4.5)$$

where $\mathcal{D} = d_Q \frac{\Omega M}{J(J+1)}$. Based on this formula it is clear that for states with $\mathcal{D} \geq 0$ this force is not restoring, however we have chosen to use the lensing state $|Q, J = 2, M = 2, \Omega = -2\rangle$ as for this state $\mathcal{D} = -\frac{4}{6}d_Q$ and the force is in fact restoring. For a long enough lens this force would cause sinusoidal motion, but for a carefully chosen length the lens acts analogously to a thick optical lens [85].

The lens design was optimized via numerical simulations in order to maximize the flux of molecules in the detection region. In practice this amounted to effectively focusing the molecules near the detection region as shown in figure 4.2, while also optimizing the following criteria. First, in order to maximize the number of molecules that enter the finite capture volume of the lens, the lens needed to be as close to the molecular beam source as was feasible. This was the primary motivation for the changes to rotational cooling implemented for ACME III, which will be discussed shortly. Second, the diameter of the lens and the lens voltage were optimized to accept the




Figure 4.2: Simulations of the Electrostatic Lens Focusing. Simulations performed by Xing Wu of different lens trajectories as viewed from above the apparatus. Green trajectories are those which reach the probe laser (in red). In a) with no applied voltage the beam is not focused. With either too little voltage (b) or too much voltage (d), fewer trajectories can be used. c) Optimal focusing is found with ± 14 kV. Figure adapted from [85].

largest possible range of transverse velocities while staying within the linear regime of the Stark shift ($\mathcal{E}_{max} \approx 30$ kV/cm [85]), and staying small enough to not require voltages beyond ± 20 kV. Once the lens was designed and assembled, it was tested using the ACME II apparatus to confirm that everything was operational and confirm the results of the simulations, which agreed well with the measurements. This optimization effectively provides a factor of $\sim 3.5\times$ improvement to our EDM sensitivity by increasing the molecular beam flux in the detection region by a factor of $\sim 12\times$.

In order to maximize the available improvements from the lens we have also opted to increase the separation of the interaction region electric field plates from 4.5 cm in ACME II, to 6 cm for ACME III. This also helps to prevent any molecules from colliding with the ITO field plates, which we are concerned could build up patch potentials and change \mathcal{E}^{nr} over time. In order to further prevent these collisions, we have included a fixed collimator between the lens and the field plates.

The lens is housed in its own purpose built vacuum chamber, replacing the majority of the stem

region from ACME II. This chamber has segments for performing STIRAP at either end, and on the upstream end attaches to the rotational cooling system, while on the downstream end it attaches to the nipple connected to the interaction region that passes through the magnetic shields. In order to minimize the distance between the beam source and the lens entrance, as well as the distance between the lens and the start of the precession region, this chamber has low profile gate valves on either end. We also mount multiple Geiger-Müller counters on the lens chamber in order to monitor for the production of any X-rays, however we have not observed X-ray production during normal operation with this lens.

IMPROVED ROTATIONAL COOLING

As molecular beam before the lens is highly divergent and the lens has a finite effective volume, any distance between the beam source and the lens opening comes at a significant signal loss. This space between the source and the lens is limited by our need for rotational cooling to concentrate the population of the molecular beam in the $|X, J = 0, M = 0, \Omega = 0\rangle$ state. As described in section 2.3.2, this was done in ACME II via a two step process, one optically pumping the $|X, J = 2^+\rangle \rightarrow |C, J = 1^-\rangle$ transition in the absence of an electric field, and one optically pumping the $|X, J = 1^-\rangle \rightarrow |C, J = 1^{\text{Mixed}}\rangle$ transition in the presence of an applied electric field to mix the parity of the C state [8]. However, in ACME II these two steps were performed sequentially, requiring separate spaces in the beamline, one with field plates and one without them, and separate optical setups. Ultimately this took up roughly ~ 20 cm in the ACME II beamline.

For ACME III we instead will perform rotational cooling in a single spatial stage, by rapidly switching the driving laser between the two transitions, which is synchronized with a switch between an applied electric field and no applied field. This scheme is shown in figure 4.3. To minimize the space required for the two lasers, they are coupled into a single shared fiber with orthogonal linear polarizations. We do this by first combining the two beams on a PBS. Once the two beams are

overlapped, we send them through a polarization EOM followed by another PBS. By switching the state of the polarization EOM, we can select out which of the two pump lasers is passed through the second PBS, with an extinction ratio of ~ 15 dB [85]. This switch state is typically changed every $8 \mu\text{s}$, in sync with an electric field being switched between 0 and ≈ 150 V/cm. After this PBS, we amplify the laser from ~ 10 mW to ~ 250 mW using a commercial tapered amplifier. As the lens can accept a large range of transverse velocities, we also want to make sure that we can fully saturate the 20 MHz full width at half maximum Doppler width of the molecular beam [85]. In order to reduce the laser power needed to do this, we use a phase EOM to introduce up to 90 sidebands to the beam with a spacing of 330 kHz. For the ACME II rotational cooling scheme we varied the polarization of the light between passes with a $\lambda/4$ waveplate to address all M sublevels. Here we instead use a Pockels cell to rapidly switch between orthogonal linear polarizations, which simplifies the optical setup at the beamline and reduces the required space. This polarization is switched between the two states every $2 \mu\text{s}$, so that for every electric field switch both polarization states each happen twice.

After the Pockels cell the lasers are coupled into the shared fiber and delivered to the beamline. At the beamline, the lasers, propagating along \hat{z} are multipassed 13 times through the molecular beam over a 2.5 cm region. This extremely compact setup is made possible by a pair of slightly offset right-angle prisms as opposed to the mirrors used in ACME II. These prisms were aligned and epoxied into place with an ultra high vacuum (UHV) compatible epoxy and are housed inside of the vacuum chamber for rotational cooling. A set of copper field plates inside the vacuum chamber produce the desired electric field along \hat{y} . This entire vacuum chamber is only ≈ 4 cm long, for a significant reduction compared to the ACME II rotational cooling. Overall this system has been demonstrated to produce an enhancement of ≈ 3.5 higher flux with rotational cooling turned on, which is a factor of $\approx 1.4 \times$ larger than was observed in ACME II [85, 8].

This method is also a significant improvement over the ACME II rotational cooling scheme as it




Figure 4.3: ACME III Rotational Cooling. Rotational Cooling Scheme for ACME III rotational cooling (a-c) and the demonstrated signal gain (d). **a)** Shows the energy levels and laser transitions involved. **b)** shows a schematic of the optical setup for combining the two lasers. **c)** is a timing diagram of the switches used for rotational cooling. Figure adapted from [85].

is more robust than the previous system. In particular both rotational cooling lasers are now commercial ECDLs¹ as opposed to the home built ECDLs used previously. Additionally, these lasers are locked to a ULE cavity with the PDH method that was also used for the STIRAP lasers and the 703 nm ECDL in ACME II, which we have demonstrated are able to remain locked for approximately ~ 1 week, as opposed to the ~ 1 day timescale associated with the transfer cavities used previously. Finally, the epoxied optics in vacuum are expected to be much more stable over time than the mirror system used in ACME II, which was external to the vacuum chamber and more sensitive to any possible unintended impacts.

¹Toptica DL Pro

4.1.3 SiPM DETECTORS AND IMPROVED COLLECTION OPTICS

The ACME II detectors were Hamamatsu R7600U-300 PMTs, which as described in section 3.1.3, were an upgrade over the PMTs used in ACME I, based on their higher quantum efficiency at the new 512 nm detection wavelength [84]. However, these detectors only had a 25% quantum efficiency at 512 nm [115], leaving potential room for improvement if a more efficient alternative could be found. In contrast to PMTs, silicon photomultipliers (SiPMs) can provide quantum efficiencies of nearly 50% at 512 nm, however, this higher quantum efficiency does not come for free, and SiPMs are known to have higher dark count rates than PMTs [116], as well as significant optical cross talk [117] and afterpulsing [118] that can result in excess noise [119]. During the development of ACME II some efforts to develop a SiPM based detector were led by Zack Lasner, but at that time we were unable to develop a robust system. Since the start of the ACME III development a group of collaborators from Okayama University have joined ACME and brought their experience with detector systems to create robust, high efficiency SiPM detectors. This work, which has primarily been led by Takahiko Masuda and Ayami Hiramoto, with assistance from Daniel Ang and myself, has been previously reported in [71, 73, 72]. Ten of these SiPM modules have now been built, eight of which will be used on the experiment with two spares.

SiPM detectors consist of a chip with an array of avalanche photodiode sensor pixels operated in Geiger mode, with the pixels connected in parallel. Our SiPM system is based on the commercially available Hamamatsu S13361-6075NE-04 [120], which is a 4×4 array of SiPMs in a 25 mm \times 25 mm package for a total of 16 sensors. Each individual sensor measures 4 mm \times 4 mm, and in order to take advantage of the total detector size, which is significantly larger than the 18 mm \times 18 mm effective area of the ACME II PMTs [115], we have increased the diameter of our fused quartz light pipes from 16 mm to 20 mm. These sensors have a quantum efficiency of $\approx 45\%$ at 512 nm. However, they also have an efficiency of $\sim 20\%$ at 703 nm which presents a problem as

the sensors will be susceptible to excess background from scattered 703 nm laser light. The dark count rate at room temperature is 2 Mcps/channel, which is orders of magnitude higher than the 6 kcps rate we saw in the ACME II detectors [84], and high enough to be a concern as we expect a count rate on the order of ~ 100 Mcps for each detector in ACME III. Additionally the optical cross talk probability without any intervention for this system has been measured to be 27.4% [71].

The dark count rate is the rate at which pulses are produced by a detector due to thermal processes rather than incident photons [121]. It is well known and stated by manufacturers that the dark count rate for a SiPM can be controlled by cooling the SiPM, as the dark count rate has a strong temperature dependence. For our system we have demonstrated that at -15°C the dark count is roughly only $\approx 3\%$ of the room temperature rate at ≈ 20 kcps/channel [71]. In practice we intend to operate at -20°C , where the dark count rate is further suppressed. This is made possible by cooling the SiPM and the aluminum circuit board it is attached to with a thermoelectric cooler, and cooling this with a closed loop chiller with 20°C water. To prevent condensation of atmospheric water onto the SiPM, the modules is housed inside of a vacuum chamber, however the vacuum requirements for this chamber are very straightforward in comparison to anything else on the experiment, as we require a vacuum below only 225 mtorr. Each of the eight SiPM modules on the experiment must therefore be connected to their own vacuum and water cooling lines, in addition to electronic control and signal lines. A schematic of the SiPM module can be found in figure 4.4

The optical cross talk in SiPMs is the effect that occurs when secondary photons are produced in the detection process of one detector pixel, which can then be picked up in a nearby pixel [117]. When this occurs, the secondary photons result in a higher than normal pulse associated with the original photon, and result in excess noise in the detection system. Similarly, after pulsing can occur when there is a time delay between the original photons detection and the detection of secondary photons produced by cross talk, which also contributes to the excess noise factor. We can quantify




Figure 4.4: Schematic of the SiPM Module. Cross section schematic of the SiPM module showing the locations of each of the three optical filters. The majority of the module volume is required for the cooling systems. Figure adapted from [73].

this excess noise factor or Fano factor [118] by measuring the output of the detector when exposed to a pulse of photons as

$$F = \frac{\sigma^2}{\mu} \quad (4.6)$$

where μ is the mean value of the observed number of photons, and σ is the standard deviation. The PMTs used in ACME II had an excess noise factor of $F_{\text{PMT}} = 1.2$, which is typically caused in PMTs by statistical nature of the gain process [119, 122]. With our system we have found that the optical cross talk can be suppressed by placing absorptive bandpass filters in the optical path to the SiPM. In particular, we place an absorptive green bandpass filter², which has a transmission peak near the detection wavelength at ~ 500 nm, directly on the SiPM and use index matching gel to couple the optical surfaces. This reduces the optical cross talk to 4.3% by absorbing the majority of

²Schott BG40

the secondary photons, which need not be at the same wavelength as the original photon, and are predominantly in the 600 – 1000 nm range [71]. In addition to this filter, we also use an absorptive bandpass filter for a vacuum window³, and an interference bandpass filter outside of the vacuum window before the light pipe⁴. These additional filters, which can be seen in figure 4.4 serve to reduce the background level caused by scattered 703 nm light from the readout and refinement lasers. These filter choices were based on measurements of the background count rate and signal count rate on the ACME II apparatus with the fluorescence source simulated with a delrin ball and a 512 nm laser alongside the refinement and readout lasers⁵. In total, the effects of these filters are that the system has now been measured to have an average excess noise factor of $F_{\text{SiPM}} \approx 1.07$, for a slight improvement over ACME II [73].

All ten SiPM modules have now been constructed, and have been thoroughly tested to verify their operation, including over a seven week long data run. In particular, they have been tested directly against the PMTs used in ACME II on the ACME II apparatus, including the smaller ACME II light pipes. The relative signal levels from this measurement can be seen in figure 4.5, which was more than a factor of 3 higher for the SiPM module. When accounting to the difference in effective area between the two sensors, the detection efficiency was $2.3 \times$ higher than the ACME II PMT, which is what we expect as the effective improvement for the ACME III apparatus with SiPMs. The system for mounting the SiPMs and lightpipes to the ACME III apparatus, as well as a discussion of methods to block stray light from the room are discussed in section 5.3.

³Schott BG39

⁴Semrock FF01-520/70

⁵the readout laser had to be turned off for tests with the delrin ball in place to prevent unrealistic scattering.




Figure 4.5: Comparison Between SiPMs and PMTs. Data taken on the ACME II apparatus to compare the detection efficiencies of the ACME II PMTs and the SiPM modules, normalized for single photon gains. Figure adapted from [73].

IMPROVED COLLECTION OPTICS

The decisions to replace the PMTs with a larger sized SiPM, along with the decision to increase the field plate separation to 6 cm for the electrostatic lens made it necessary to redesign the collection optics for ACME III. This work, which was led by Daniel Ang, was based on simulations in LightTools, just like how the ACME II collection optics were optimized, and the design process is described in detail in Daniel’s thesis [123]⁶. As before, the collection optics consist of eight pairs of lens doublets which are mounted just outside of the field plates, and are aimed at the detection volume where the readout laser intersects the molecular beam. Each doublet focuses the collected light into a light pipe which transmits the collected light out of the vacuum chamber and magnetic shields to the SiPM. Based on constraints related to the size of the interaction region chamber and the decision to use straight light pipes instead of the curved light pipes, the motivations for which

⁶At the time of this writing this has not been published.

are discussed in section 5.3, the angle of the optics were fixed at 45° . The most notable change made in this design process was to move away from commercial aspheric lenses, as they were limited in diameter to 75 mm, and instead order custom lenses with a 135 mm diameter for the larger lens, and a 106 mm diameter for the smaller lens in the doublet. The effect of this is clear as the optics are able to cover a larger solid angle than the off the shelf optics, and this is especially critical as the optics must now be further from the source due to the larger new field plate separation. Additionally, these simulations were used to optimize the diameter of the light pipes used in ACME III. Based on these simulations and experimental tests, we expect these new lenses and the larger light pipes to have a collection efficiency of $\sim 30\%$, and should collect a factor of $\approx 1.7\times$ higher signal than ACME II.

4.1.4 EXCESS NOISE SUPPRESSION

The ACME II measurement was unable to reach the shot-noise limit in practice due to a source of excess noise that was not well understood until after the completion of the measurement. This noise, which did not vary with any known experiment parameters, was characterized by the reduced chi-squared per degree of freedom of the EDM data set, which was $\chi_r^2 \approx 3$. The efforts to suppress this noise were led by myself, Cristian Panda, and Mohit Verma, and were ultimately successful in suppressing through three independent methods as measured on the ACME II apparatus, enabling us to reach the shot-noise limit in our tests as was reported in [69]. We believe that with these suppression methods this noise has been suppressed to well below the level of the expected ACME III sensitivity, which acts as an effective gain in EDM sensitivity of $1.7\times$. In this section we will describe in detail the model for how this noise enters our system, describe the three methods of suppressing this noise, and describe how this was verified.

In order to understand the source of the excess noise we must first consider the way we performed polarization switching in ACME II, as well as how we acquired our data and used the po-

larization switching to compute an asymmetry \mathcal{A} . As was described in section 3.2.4, in order to project out the phase of our molecules we rapidly switch the readout laser between the two polarizations \hat{X} and \hat{Y} using a pair of AOMs. These AOMs are controlled by TTL pulses generated by an SRS DG645 precision timing and delay generator that is synced to the timing of the molecular beam pulses, which control RF switches to turn the AOMs on and off. We switched the polarizations at a frequency of 200 kHz, with a total polarization cycle period $T = 5 \mu\text{s}$. We nominally had a $0.6 \mu\text{s}$ dead time between switches, which was chosen based on the lifetime of the I state [90] to ensure that there was no overlap between the \hat{X} and \hat{Y} polarization “bins”. We then detected the fluorescence associated with each polarization, S_X and S_Y , using our data acquisition (DAQ) digitizer⁷, which sampled at 16 MS/s, which corresponded to a spacing of $\Delta t = 62.5 \text{ ns}$ between each data point. This sample rate was chosen so that each full polarization period consisted of an even number of data points (80 total, with 40 assigned to each of \hat{X} and \hat{Y}). From this we define a polarization “sub-bin” which consisted of a series of consecutive data points within a single polarization bin used to integrate the measured fluorescence, and which were defined to be the same for the \hat{X} and \hat{Y} bins. This sub-bin timing can be seen in figure 4.6 along with the other relevant timing parameters for this discussion. We must also define the time $T_{\hat{X}-\hat{Y}}$, which is the time delay between the actual \hat{X} and \hat{Y} laser pulses, which ideally should be $T_{\hat{X}-\hat{Y}} = T/2$, but in practice has been observed to vary from this by as much as 200 ns. This discrepancy is caused by the way we use AOMs to control the laser pulses, as there is a propagation time associated with the acoustic wave in the AOM crystal [124, 125], which can vary with the alignment of the laser into the AOM crystal. The last term we define is $\Delta T_{\hat{X}-\hat{Y}} = T_{\hat{X}-\hat{Y}} - T/2$, which describes this discrepancy. In practice we worked to minimize $\Delta T_{\hat{X}-\hat{Y}}$ in ACME II by timing the laser pulses on a photodiode, with $\sim 40 \text{ ns}$ precision [69].

⁷NI PXI-5171R




Figure 4.6: Polarization Switching Timescales. (a) shows the timescales within a single polarization cycle. The gray shaded region of the two fluorescence signals indicates the sub-bin used for integration. A,B, and C indicate the three characteristic regions of the fluorescence signal, when excitation is first turned on, when the population becomes roughly evenly mixed between the H and I states, and when the excitation is turned off. The time $T_{\hat{X}-\hat{Y}}$ is shown relative to the timing of the two laser pulses. (b) shows the measured fluorescence signals vs time for each polarization as averaged over 25 shots to form a trace. Figure adapted from [69].

We can now see how a nonzero $\Delta T_{\hat{X}-\hat{Y}}$ can impact the asymmetry by looking at the time dependent instantaneous asymmetry

$$\mathcal{A}(t) = \frac{S_X(t) - S_Y(t)}{S_X(t) + S_Y(t)}, \quad (4.7)$$

where t is the time within the polarization bin, and which in practice we can make discrete by considering our digitization rate such that $t_i = i\Delta t$, where $i \in \{1, 2, 3, \dots, 40\}$ for ACME II. When $\Delta T_{\hat{X}-\hat{Y}} \neq 0$ $S_X(t_i)$ and $S_Y(t_i)$ do not correspond to the same portion of the fluorescence signal as the two data points have been collected at different times relative to the start of the actual laser pulses. This results in a time dependent shift of $\mathcal{A}(t)$ away from the mean of the asymmetry as shown in figure 4.7, which for $\Delta T_{\hat{X}-\hat{Y}} \ll T$ is proportional to $\Delta T_{\hat{X}-\hat{Y}}$ and the time dependent slope of the fluorescence signal. We can represent this shift as

$$\Delta \mathcal{A}(t) \approx \frac{1}{2S(t)} \frac{dS(t)}{dt} \Delta T_{\hat{X}-\hat{Y}}, \quad (4.8)$$

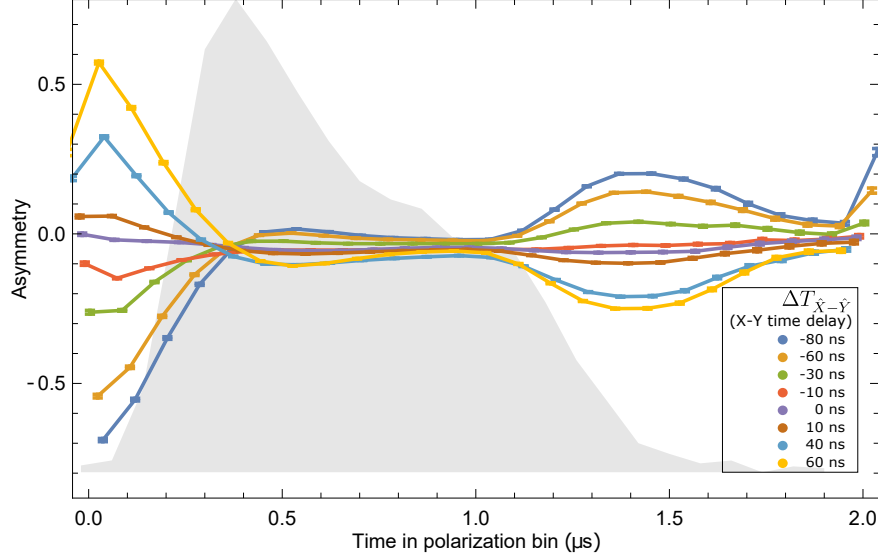


Figure 4.7: Asymmetry Within a Polarization Cycle. Asymmetry versus time within a polarization cycle, with the averaged signal $S(t)$ shown in gray. The different colored data sets correspond to different values of $\Delta T_{\hat{X}-\hat{Y}}$ as shown. The magnitude of the asymmetry is largest for large values of $\Delta T_{\hat{X}-\hat{Y}}$ and at times where the slope of the fluorescence signal is largest. Figure adapted from [69].

where $S(t) = (S_X(t) + S_Y(t))/2$ is the average of the \hat{X} and \hat{Y} fluorescence signals. On its own, this time dependent shift in the asymmetry would not be as much of a concern, as if we had a constant $\Delta T_{\hat{X}-\hat{Y}}$, then $\Delta \mathcal{A}(t)$ would produce a shift in the asymmetry that was not correlated with any experiment parameters, and would be subtracted out by our parity sums. Additionally, the sign of the asymmetry is reversed under the $\tilde{\mathcal{P}}$ and $\tilde{\mathcal{R}}$ switches, so any $\Delta \mathcal{A}(t)$ is further cancelled in our calculation of $\omega^{\mathcal{N}\mathcal{E}}$ as it is even under these switches.

However, in the presence of timing jitter between the laser pulses and the triggering of the DAQ, any nonzero $\Delta T_{\hat{X}-\hat{Y}}$ introduces noise in $\Delta \mathcal{A}(t)$. This noise can enter our measured phase in any channel which is odd with respect to an experimental switch which is changed at a frequency that is slower than the frequency that the noise occurs at. As our fastest switch $\tilde{\mathcal{N}}$ occurs every 0.6 s, any timing jitter faster than this time scale can introduce noise to the experiment. In consideration of this fact, one upgrade we considered for ACME III involved methods to switch $\tilde{\mathcal{N}}$ within a single

molecular beam pulse to suppress fast noise. We ultimately decided not to implement any such “fast $\tilde{\mathcal{N}}$ switch”, due to the complexity of the required changes to the experiment, and the fact that we did manage to control this timing jitter related noise.

From this model we expect that we should be able to control this noise source in two ways that are already clear, and one method that is dependent on our analysis method. The first method to control this noise is to ensure that there is no detectable timing jitter between the laser pulses and the trigger of the DAQ, as a constant $\Delta T_{\hat{X}-\hat{Y}}$ will subtract out. This effort, which will be described shortly took multiple months of intensive investigation to track down, as the noise was itself dependent on three flaws in our DAQ system, which all needed to be resolved. The second method is to minimize $\Delta T_{\hat{X}-\hat{Y}}$ with greater precision than was typically done in ACME II, which from equation 4.8 will reduce the magnitude of $\Delta\mathcal{A}(t)$ and the magnitude of the noise contribution. The effects of independently controlling the timing jitter and $\Delta T_{\hat{X}-\hat{Y}}$, based on measurements on the ACME II apparatus performed by Cristian Panda and myself can be seen in figure 4.8. Due to the requirements for suppressing the timing jitter, we changes the sampling rate of the DAQ from 16 MS/s to 12.5 MS/s in these measurements, which changed the time between digitization points to 80 ns. In order to keep an even number of data points within each polarization cycle, we changed the polarization switching frequency from 200 kHz to 250 kHz, and used a dead time between polarization bins of 0.8 μs , leaving 25 digitization points in each of the \hat{X} and \hat{Y} polarization bins. For the third method of controlling this noise, we must consider the fact that in our normal analysis we integrate $\mathcal{A}(t)$ to compute the mean asymmetry, which is shifted from the “true” mean asymmetry by the integral of $\Delta\mathcal{A}(t)$. This integral is dependent only on the value of $\Delta\mathcal{A}(t)$ at the endpoints of this integral t_{i0} and t_{if} , which define the sub-bin. As $\Delta\mathcal{A}(t)$ is proportional to $dS(t)/dt$, we can choose the endpoints of our sub-bin t_{i0} and t_{if} to minimize $dS(t)/dt$, which in turn minimizes the mean $\Delta\mathcal{A}$, reducing the noise in our measurement. The effect of this choice of sub-bin start and end points is shown in figure 4.9, conducted in the same configuration as the other

tests with 25 data points per bin. From these tests, we were able to determine that each of these three control methods individually suppressed this excess noise, and when combined suppressed this noise to the point where we were able to measure the reduced chi squared of our data set to be $\chi_r^2 = 0.87 \pm 0.40$, which was consistent with 1 as expected [69].

TIMING JITTER CONTROL

In order to describe the methods developed to suppress the timing jitter, we will first describe the ACME II DAQ system in more depth. The primary DAQ system in ACME II was a NI PXIe-5171R eight channel PXI Oscilloscope, which we typically refer to as the FPGA. This FPGA oscilloscope was housed in a PXIe-1075 18 slot chassis, which was connected to our DAQ computer by a PXIe-8375 MXI Express card paired with a PCIe-8375 for remote control. Each of the eight channels was normally connected to one of the eight experiment PMT signals, and was provided triggering information through the PFIO input, which is accessible through a NI breakout box that we connect to with a NI cable⁸. This triggering signal is produced by our DG645 and should have extremely low jitter of 10^{-8} [126], which we can confirm on a bench-top oscilloscope where we see no timing error. The LabView bit file for this FPGA was developed by NI engineers in contact with Adam West to allow us to use the PFIO trigger, but is otherwise the same as the example file available online.

We also have a backup system consisting of four PXI-5922 PXI oscilloscope cards. This system is simpler in some ways, but each card only uses two input channels, and each card has an option for a clock input and trigger. This clock input must be a square wave unlike the FPGA which can take a square or sine wave input. Additionally, this system has been configured from the start to use the

⁸SHH19-MH19-AUX. Note that this looks like an HDMI cable, and has the same ends, however it apparently has nonstandard wiring, which caused problems when Adam first tried to work with the system, so make sure that the Aux I/O cable is an NI cable.




Figure 4.8: Asymmetry Noise Dependence on Clock Syncing and $\Delta T_{\hat{X}-\hat{Y}}$. Measurements of excess noise characterized by χ_r^2 in the ACME II apparatus. (a) and (b) show the asymmetry noise as a function of time within a polarization cycle, where the colored data sets correspond to different values of $\Delta T_{\hat{X}-\hat{Y}}$ as shown. In (a) the DAQ system is not synced to an external clock, while in (b) the DAQ is synced to our Rb clock. The noise is clearly largest when the slope of the fluorescence signal is largest, for larger values of $\Delta T_{\hat{X}-\hat{Y}}$, and for the configuration without clock syncing. (c) shows the χ_r^2 averaged over an entire bin as a function of $\Delta T_{\hat{X}-\hat{Y}}$, with the blue data set using the clock syncing, and the red data set using no external clock. All values are calculated from 200 consecutive molecular pulses. Figure adapted from [69].





Figure 4.9: Asymmetry Noise Dependence on Sub-bin Selection. The asymmetry noise (χ_r^2) depends on the selection of the endpoints of the sub-bin. By choosing sub-bin endpoints for which $dS(t)/dt$ is smaller, we can minimize the excess asymmetry noise. The different colored data sets correspond to different values of $\Delta T_{\hat{X}-\hat{Y}}$ as shown. This data was collected with no clock syncing to separate the effects of different noise suppression methods. All values are calculated from 200 consecutive molecular pulses. Figure adapted from [69].

NI-SCOPE software driver.


While investigating the source of our excess noise, we discovered that our excess noise was correlated with the beginning and end of a molecular beam trace. Further investigation revealed that there was a significant source of timing noise in the DAQ system. This was very evident when the polarization switching control pulses from the DG645 were sent directly into the PXIe-5171R module and recorded in LabView. Specifically we sent in two signals which were 10 ms long pulse trains of 200 kHz square waves repeating at 50 Hz that are roughly 180° out of phase with each other. The trigger pulse was sent from the DG645 every 500 ms, as we did normally during the run of ACME II. We could see that the arrival time of these 50 Hz pulse trains varied over a range of about 120 ns. This 120 ns corresponded to shifting one full digitization point in either direction, and the traces could be seen to drift somewhat continuously over this range as shown in figure 4.10. This jitter also appeared to be linear and periodic within this range.

First Pulse of Consecutive Shots, PFI Trigger, No Clock, 16 MS/s



(a) The leading edge of consecutive 10 ms pulse trains overlapped, spaced by 50 Hz. No clock reference was used, and we sampled at 16 MS/s so each point corresponds to 62.5 ns.

Single Trigger Pulse Train, PFI Trigger, No Clock, 16 MS/s



(b) The rising edges of the 200kHz square wave within a single 10 ms pulse train overlapped, sampling at 16 MS/s. No clock reference was used, and the signal drifts at a rate of at least 6 ppm.

Figure 4.10: Timing Jitter Before Modifications The two time scales at which timing jitter was observed before implementing any of the three control methods described in this section. (a) shows the jitter between the leading edges of successive 50 Hz triggers. (b) shows the faster time scale where the leading edge of each polarization trigger within a single molecular beam shot are compared. In both cases, the drift is linear and periodic within the range. The vertical scale in both plots is in arbitrary units.

Once we had identified that there was a timing issue, we checked that it was not an issue with the signals themselves by connecting them to a standard lab scope and saw that the problem disappeared. This confirmed that what we were seeing was being caused by the DAQ and was not a true signal error. We were eventually able to fix the problem by implementing three changes.

Our first thought was that there must be some sort of clock syncing error, as we had not synced our SRS Rb clock⁹ to the DG645 and we had no reference clock synced to the FPGA itself. The 5171R has a timebase accuracy of only ± 25 ppm when using the onboard clock [127], which is enough to explain the behavior we saw. We first synced the DG645 to the clock using the clock input on the DG645, but saw no change in the behavior. We then determined that we could sync the clock to the FPGA through two methods. In either method we can use either a sine wave or a square wave. The first is to use the AUX I/O inputs to connect the clock, which is through pin 6 of the breakout box ¹⁰. This method requires that the input source of the clock be set in LabView to “VAL_CLK_IN”. The second is to use the onboard clock of the chassis, which has a BNC input for syncing with an external clock. This is probably preferable if we are syncing multiple cards to the clock, and requires the LabView source name “VAL_PXI_CLK”. For the FPGA we observed no difference between these methods, but with the PXI-5922 cards we saw one card show jitter when using this clock method while it was fine with the direct clock input. This needs to be investigated further before trusting though.

The PXI-5922 cards can be synced to a clock using either the same chassis clock sync method as the FPGA, or by using the front panel CLK IN¹¹ connection. There are some concerns with using the chassis clock as it may not cleanly distribute the clock signal to all boards as previously men-

⁹SRS FS725

¹⁰this can be found online in the getting started guide for the 5171

¹¹The LabView settings for the clock input with these cards are significantly easier to remember as LabView will create a dropdown menu of the available choices.

tioned. However, the front panel clock input required us to generate a square wave synced to our Rb clock as our clock source only outputs a sine wave. For this a SRS signal generator was sufficient. We also had to input a clock signal individually for all 4 boards.

When synced with a reference clock there was a clear improvement in the behavior as the drift rate between consecutive pulses clearly decreased, but the drift over three points remained. This was clear as the overlapping pulses no longer were able to drift over the full two point range in a single 500 ms triggering window, and the traces overlap significantly more.

While this did not fix the issue, it did improve the jitter that we were seeing. To try to better understand the source, we also looked at the overlap of the rising edges of the 200 kHz square wave within a single 10 ms pulse. We did this both with and without a clock and found that there was significant drift over a range between two digitization points that went away when synced to a clock. This drift without a clock corresponds to a drift rate of at least 6 ppm, which is only a lower bound, but is within the FPGA spec. These results can be seen in figures 4.10 and 4.11 where in the clock case all of the traces overlap so well as to almost look like a single trace. This is the behavior we want to see at this level, and indicated that the clock was doing its job, and most likely was not the problem anymore.

The next control method we discovered was related to an error in the way the FPGA attempted to sample at a non-integer division of its clock speed. The FPGA clock runs at 250 MHz, and so we were unsure of how the software actually handles decimation by non-integer values, like the 16 MS/s sample rate we had used in ACME II and in our tests up to this point. The manual indicates that the system has the option to decimate by n . The choice of 16 MS/s had been made for ACME II in spite of this in order to cleanly fit an even number of digitization points within the 200 kHz polarization switching cycle. In LabView we set a sample rate rather than the decimation, and the FPGA bit file directly takes this input and decides how to decimate. However, we were unable to determine directly from this file or from conversation with NI engineers how the program interprets a

Single Trigger Pulse Train, PFI Trigger, DG645 Clock, 16 MS/s




Figure 4.11: Fast Scale Timing Jitter After Clock Syncing The rising edges of the 200 kHz square wave within a single 10 ms long pulse train overlapped, sampling at 16 MS/s. The FPGA clock was synced to the clock output of the DG645, which in turn was synced to the Rb clock module. No drift is visible at this fast time scale once a stable clock is introduced. The vertical scale is in arbitrary units.

non-integer input for the decimation factor. The best understanding of how this must be occurring is that the FPGA somehow interpolates between the two nearest integer decimation rates to almost sample at the specified rate, however NI could not confirm how the system operates without integer decimation. We experimented with changing the sample rate and found that by switching to integer decimation the drift rate decreased significantly once again.

In trying to track down the cause of the remaining drift and periodicity, we repeated these tests with the analog trigger settings instead of the PFI trigger. For this we used the same trigger signal sent into channel 1 of the FPGA. We tried this test because the FPGA has a trigger jitter specification that is 8 ns for a digital trigger and one Sample Clock timebase period for an analog trigger. This difference initially did not make sense to us until discussing with an NI engineer who told us that the digital trigger uses a 125 MHz clock rather than the 250 MHz clock used for the analog trigger. The tests seemed to somewhat confirm this specification as when triggering off of the PFI

input we saw drift over three points at all sample rates, and when using an analog trigger we saw drift over just two points at all sample rates. However, this did not match the 8 ns spec for digital triggers because of how they scaled with the sample rate. NI R&D also could not confirm whether this behavior was expected or not.

The PXI-5922 has a similar specification for the analog triggers (the only option) so we repeated these tests with this system, and found that using integer decimation sampling rates and syncing to a clock we no longer had any discernible jitter. This confirmed that there must either be a problem with our FPGA code or the FPGA hardware itself that was out of spec.

The final change required to remove the noise was switching over from our old LabView code. We previously had used the ACQ session palette tools to control and read from the FPGA, but NI engineers suggested we try using the newer NI-SCOPE driver as the ACQ tools were no longer supported. Additionally, the tests with the PXI-5922 suggested that this might be a potential solution as it was one of the few differences between the code we used for the different systems. Switching over fixed the remaining issues, but neither we or NI can actually explain why this worked. It seems to be a bug in the old code that will never be fixed with unsupported drivers.

After implementing these three changes we were able to remove the noise to the level that we expect to be able to measure. In all future measurements with this system three things should always be implemented. First, always sync the scope to a clock source. Second, make sure that we are clearly choosing a sample rate that is an integer decimation of the sample clock. For ACME III this will most likely be 12.5 MS/s, as we don't want to increase our data taking rate for space reasons, and we will have to adjust our polarization switching rate to match the sample rate nicely. However, if we want the PXI-5922 modules to be completely interchangeable with the FPGA DAQ, we must choose a sample rate which is a common factor of both sample clocks, the closest of which is 10 MS/s. Third, never use the ACQ session palette tools to connect to the scopes as it is unsupported software with dangerous bugs we do not understand. The NI-SCOPE driver is currently




Figure 4.12: Timing Jitter After All Modifications The leading edge of consecutive 10 ms pulse trains overlapped, spaced at 50 Hz. The clock is synced to the Rb reference clock, using the NI-Scope drivers, using a digital trigger, and sampling at integer decimation with 25 MS/s. With all of these changes there is no remaining timing jitter that we can detect. Similar results are found with an analog trigger, which previously had less jitter than the digital trigger, and for the PXI-5922 module operating under the same conditions. The horizontal axis corresponds to the digitization point number, which in this case have a spacing of 40 ns. The vertical scale is in arbitrary units.

supported and shows no problems. For what it's worth, it is also easier to work with. The final result from all of these modifications can be seen in figure 4.12.

4.1.5 RAPID BEAM SOURCE TARGET CHANGING

The last upgrade we have implemented to increase our statistical sensitivity is a method for increasing the rate at which we change our ablation beam source targets. Anecdotally, when running the beam source it has long been clear to us that the signal from an ablation target decays steadily over the lifetime of the target, even with changes to the ablation spot, and the focal depth of the ablation spot on the target. This signal decay was always the motivation behind changing out the target every few weeks during intensive data collection, while the actual mass loss from the targets was small relative to their size. After the ACME II measurement, we performed an analysis of all of the signal levels in the ACME II final data set to study the signal decay, the result of which can be seen in fig-




Figure 4.13: ACME II Signal Levels. Signal level for each superblock over the course of the ACME II data set. Changing the ablation target causes the rapid jumps in signal level. Figure created by Zack Lasner.

ure 4.13. From this analysis it is clear that the signal level decays roughly linearly over time, and the signal typically decayed by a factor of $\sim 3 - 5\times$ before reaching the threshold for a target change. If we were to consider only the first run of each target (one day of data), compared to the average signal level over the lifetime of each target, we could expect a factor of $\approx 1.5\times$ higher signal. However, this would require us to be able to perform target changes every day. While replacing the ablation cell itself along with a target seemed to anecdotally increase the signal, the variance in initial signal levels with either new cells or new targets made it difficult for us to draw as strong a conclusion with our limited data.

During the ACME II final data collection, we typically performed target changes every 2 – 3 weeks. This was due to two factors, the first of which was that the ACME II cell housed two targets, one upstream and one downstream, with a slightly higher peak signal and longer lifetime for

the downstream target, which extended the time needed between target change operations. The target change operations were typically performed at the start of the week, as it took most of the weekend for the beam box to warm up to a high enough temperature to open up without causing condensation. The actual target change process, which typically took ~ 4 hours was then followed by a ~ 12 hour cooling to get back to our operating temperature. In order to make it possible to condense this process to the point where it could be performed every day without impacting our duty cycle, we needed to develop a load lock apparatus so that we do not need to warm up and cool down the beam box for each change. Additionally, by doing this we are able to reduce the amount of “radiation work” we need to do, where we don PPE to protect from exposed thorium targets, and exposed thorium dust. This is a significant safety and quality of life improvement for the experiment.

This system, which Zhen Han led the development of, can be seen in figure 4.14. The targets, which are mounted onto a copper block, can be inserted with a rod into a load lock chamber at the top of the beam box, where it can be screwed onto the end of a rotatable vertical translation rod. Once attached to this rod, the load lock chamber can be closed, and the vertical translation rod is lowered to the correct height for the target. Here the target block is transferred to a horizontal rod, which is connected to a linear bellows and a ferrofluid feedthrough to allow for rotation. This horizontal rod is threaded, but also has a smaller ball driver end on the tip instead of more threads. The threaded portion is used to hold the target block, which is then translated into position on the cell wall, where it slides onto mounting rods on the cell. The horizontal manipulation rod can then be disconnected from the target block, and the ball driver end of the tool can be used to tighten the four screws that provide good thermal contact between the target block and the cell itself and prevent gas flow leaks at the joint. In order to remove a target, the opposite procedure is used. Old targets can be sealed in plastic bags connected to the load lock chamber, which can be laminated closed to prevent any potential exposure to free thorium dust.




Figure 4.14: ACME III Load Lock System for Rapid Target Changes. Partial cross section of Beam Box II showing the components of the load lock system. The red x indicates the location where the ablation target is transferred from the vertical arm to the horizontal arm which allows access to the ablation cell.

With this system we believe that we will be able to perform target changes at the end of each day in roughly ~ 1 hour and expect to gain the full factor of $\sim 1.5 \times$ higher signal. This is convenient as it coincides with the deicing we already know is necessary after each day of running, so this will not negatively impact our experiment duty cycle. As we will not need to pause the experiment we may in fact increase our overall duty cycle slightly. However as we may still need to perform cell changes at an as yet unknown frequency, it remains to be seen how much this will impact our duty cycle.

4.1.6 PROPOSED UPGRADES THAT WERE NOT IMPLEMENTED

It is important to keep a record of any upgrade pathways we considered that ultimately were not successful, most of which considered improvements to the detection system in ACME III. The most significant of these proposals was to replace our detection method with an optical cycling transition to increase our detection efficiency by a factor of ~ 10 . These efforts, which mostly took place during the ACME II measurement, were led by Daniel Ang. The primary idea behind this was based on a recent measurement of the branching ratio from the I state to the X state [89], which reported a branching ratio of 91%. Based on this branching ratio measurement, we hoped to be able to drive the $X \rightarrow I$ 512 nm transition, to cycle and detect on average 10 photons from each molecule, which we hoped would increase our detection efficiency from $\sim 4\%$ in ACME II to nearly $\sim 100\%$ in conjunction with other upgrades. We developed a 512 nm laser system locked to a ULE cavity, and developed schemes to maintain the separation between the \hat{X} and \hat{Y} polarization bins during cycling. However, based on measurements in an apparatus built by Daniel Ang and myself, we found that we were unable to cycle as many as 10 photons like we had hoped. In addition to this, we found that there was a previously unappreciated issue with detecting by optical cycling due to the statistical nature of the decay pathways [128]. In essence the stochastic nature of how many transitions could occur before each molecule decayed to a dark state would add an excess noise factor which would limit the possible gain in our EDM sensitivity from optical cycling to a factor of

~ 2 , even if we could cycle an average of 10 photons. A full description of these measurements and the difficulties we found can be found in Daniel's thesis [123].

We also considered two other methods for upgrading the detection efficiency of the experiment, while these were discarded as infeasible before tests were performed that may not be true in the future. The first method considered was ionization of the ThO molecules, as the detection efficiency for ions on a multichannel plate is near unity. This detection of ions is used in other EDM experiments that already work with molecular ions [63, 64]. However, concerns of the availability of lasers and a lack of knowledge about the necessary transitions in ThO made this unattractive at the time. As part of both the considerations for ionization detection and for cycling detection, we also considered the development of a secondary detection chamber after the interaction region vacuum chamber where we would have the freedom to detect nearly the full solid angle of emitted photons. However, the costs of such a system, the difficulty of shelving molecules in a long lived state between the projection into the \hat{X} and \hat{Y} basis and detection, as well as space concerns in the lab made this infeasible at the time.

The last proposed upgrade we will describe comes from attempts to increase the flux of the molecular beam source. These efforts were implemented by Xing Wu, Jonathan Haefner, and myself. A number of tests were conducted on beam box II, similar to those tests performed in the past [79] to optimize the beam flux, such as changing the operating temperature, the YAG pulse rate, and the YAG pulse energy. However, these tests did not yield any noticeable improvements, as doubling the pulse frequency yielded approximately half the beam flux per pulse, and the increased YAG power did not increase the yield. We additionally attempted to implement a de Laval nozzle, which first converges then diverges, as the exit aperture from the cell, which previously was just a straight machined hole in the cell. This was based on promising results reported in [129], which suggested that we might be able to increase the collimation of the beam exiting the cell and thereby increase our flux. However we were not able to see any statistically significant improvement, in part due to

the large variability in beam flux whenever we modified the cell. The most promising result of any of these tests was the determination that we could cut our ceramic ThO₂ targets in half to reduce their thickness without impacting the beam flux. This suggested a method to allow us to extend our available target supply before having to press our own targets.

4.2 SYSTEMATIC ERROR SUPPRESSION

With the planned statistical improvements to the experiment it is also necessary to ensure that all known sources of systematic error in the ACME experiment are suppressed to at least the level of the new expected statistical sensitivity. If we consider the class A systematics that were included in the ACME II error budget, as shown in table 3.2, we can clearly see that the $\omega_{\text{ST}}^{\mathcal{N}\mathcal{E}}$ and $\omega^{\mathcal{E}}$ systematic errors can already be expected to remain far smaller than our statistical uncertainty. Additionally, we did not observe any source of $P_{\text{ref}}^{\mathcal{N}\mathcal{E}}$, and so we did not include a shift for this parameter. Based on this we do not have reason to believe that a new source of $P_{\text{ref}}^{\mathcal{N}\mathcal{E}}$ will occur in ACME III, and so we expect our ability to place a bound on the uncertainty associated with $P_{\text{ref}}^{\mathcal{N}\mathcal{E}}$ to scale with our overall statistical uncertainty. We can make a similar argument for the systematic contributions from $|\mathcal{C}|^{\mathcal{N}\mathcal{E}}$ and $|\mathcal{C}|^{\mathcal{N}\mathcal{E}\mathcal{B}}$ as here we also have no model for why these values should be non-zero, and our measurements were consistent with this. Additionally, these systematics are proportional to $\omega^{\mathcal{B}}$ and ω^{nr} , respectively, which are determined by the Zeeman phase and our global polarization offset θ , which are chosen to set $\phi - \theta \approx \pi/4(2n + 1)$. With our increased precession time, these frequencies should therefore both be scaled down by roughly a factor of 5×10^{12} . This leaves two categories of systematic errors which require consideration and modifications to the ACME III apparatus, which are discussed in the following two sections.

¹²This scaling is not exact as $\omega^{\mathcal{B}}$ will also be scaled with our choice of magnetic field magnitude, and ϕ^{nr} will compensate for the change in Zeeman precession to maintain our sensitivity to changes in the asymmetry.

4.2.1 MAGNETIC FIELD SYSTEMATICS AND NOISE

The first source of systematic errors we will consider is those related to magnetic fields, in particular the magnetic field gradients $\partial\mathcal{B}_z/\partial z$ and $\partial\mathcal{B}_z/\partial y$. We described the models for these systematics in section 3.4.2 there were in fact two models for sources that contributed to this component of the error budget, one of which relied on the non-reversing electric field \mathcal{E}^{nr} , and one of which relied on gradients in \mathcal{E}^{nr} and had a $\sim 3 - 5\times$ smaller effect than the other.

As shown in equations 3.16, the larger of these two systematic shifts scaled with $\partial\mathcal{E}^{\text{nr}}/\partial z$ (or $\partial\mathcal{E}^{\text{nr}}/\partial y$), a laser detuning¹³ through $\partial\eta/\partial\delta$, and the magnetic field gradient $\partial\mathcal{B}_z/\partial z$ (or $\partial\mathcal{B}_z/\partial y$). We do not currently know of a reliable method to decrease the scale of \mathcal{E}^{nr} or its gradients in our experiments, as we believe these are the source of patch potentials on the ITO surface of the field plates [130]. We can minimize the contribution from $\partial\eta/\partial\delta$ as we did in ACME II, by optimizing δ , however we do not expect the linewidths of the relevant transitions to change significantly for ACME III (although the Doppler width of the beam will decrease slightly thanks to the electrostatic lens). We therefore expect to still be able to use this as a tool to suppress this systematic in ACME III, but we still want a new suppression factor if possible. In contrast we do have the ability to apply greater control over the magnetic field gradients in ACME III. In ACME II, we were able to control the effect of this systematic effect in part by applying additional gradient fields to suppress the the gradients to below $\leq 1 \mu\text{G}/\text{cm}$. In ACME III we hope for the ambient gradients to be at or below this level based on our new magnetic shield system and our cosine theta coil system, for reasons we will describe shortly.

As equation 3.19 shows, the second source of magnetic field systematic shifts depends on \mathcal{E}^{nr} itself, the relationship between position and transverse velocity for our beam $\partial z/\partial v_z$, and the mag-

¹³In the STIRAP lasers this is the two-photon detuning δ , but this effect occurs in the readout laser to a lesser degree as well, where the detuning Δ is relevant instead.

netic field gradients. While the same arguments as before hold for \mathcal{E}^{nr} and $\partial\mathcal{B}_z/\partial z$ (and $\partial\mathcal{B}_z/\partial y$), we do expect to see an improvement in $\partial z/\partial v_z$ thanks to the electrostatic lens. In ACME II, there was a strong correlation between the position of a molecule and its transverse velocity due to the ballistic trajectories our molecules followed after exiting the source and the way we collimated the beam. However, in the presence of the electrostatic lens, this correlation will be much weaker due to the way the lens acts to focus the molecules into the detection region, which can be seen clearly from the trajectories shown in figure 4.2. Thanks to this, as well as the fact that this impact of this systematic pathway was smaller in ACME II, we expect the contribution from this effect to be sufficiently small for ACME III.

In order to fully describe our goals for the magnetic field gradients in ACME III and how well we expect to control them, it is important to also consider the excess noise we saw in ACME III that was correlated with the magnitude of the applied magnetic field, which we described in section 3.5.2. This noise, which was caused by fluctuations in the forward velocity of the molecular beam, was observed when we operated with $|\mathcal{B}_z| = 26$ mG, but we did not see any excess noise when operating at the values $|\mathcal{B}_z| = 0.7, 1.3, 2.6$ mG. Based on this evidence, we determined that for ACME III we expect to apply an order of magnitude smaller magnetic fields with $|\mathcal{B}_z| \sim 100 \mu\text{G}$ to have this noise scale with our order of magnitude smaller expected statistical sensitivity. We expect that even with the longer precession time (and therefore larger $|\phi^{\mathcal{B}}|$), we expect to lowest order that the longer precession time will not increase the the noise in our EDM measurement which from equation 3.23 should scale as $\sim \sigma_{\phi^{\mathcal{B}}}/|\phi^{\mathcal{B}}| \approx \sigma_v/v$. However, we are not yet certain of the exact scale at which this noise will occur in ACME III¹⁴ so we based our ACME III design goals on an assumed field magnitude of $\sim 100 \mu\text{G}$.

With this choice of $|\mathcal{B}_z| \sim 100 \mu\text{G}$ we set our design goal for the ambient magnetic field within

¹⁴We also will be using a new Litron YAG ablation laser, which seemed to also have an effect on this in ACME II when we swapped between the Quantel and Litron YAG lasers we had available.

the precession region to be $1 \mu\text{G}$ or less. However, we recognize the difficulty of achieving this goal, and expect that our experiment will be able to operate as expected so long as we actually achieve an ambient magnetic field $< 10 \mu\text{G}$. The details of how we plan to achieve this goal, through the development of a new system of magnetic shields, as well as the results at the time of this writing are discussed in section 5.2. It is worth noting that in ACME II, we were able to operate at $|\mathcal{B}_z| = 700 \mu\text{G}$ despite an ambient field on the order of $\sim 200 \mu\text{G}$, so if we have ambient field of $\leq 10 \mu\text{G}$, we should expect to be able to safely operate at fields as small as $\sim 30 - 50 \mu\text{G}$. At the time of this writing, without having assembled the complete apparatus, we have been able to achieve the goal of fields below $10 \mu\text{G}$, and we believe there may be further room for improvement with our apparatus, but this requires further study.

Based on these choices for the ambient magnetic field magnitude inside the precession region we also effectively place constraints on the gradients that contributed to systematic errors $\partial\mathcal{B}_z/\partial z$ and $\partial\mathcal{B}_z/\partial y$. As previously described, we set a design goal of $< 1 \mu\text{G}/\text{cm}$ for all magnetic field gradients in ACME III, based on the level to which we were able to control our gradients in ACME II. In practice, with our new apparatus, while achieving ambient fields of $< 10 \mu\text{G}$, we have observed that all gradients in the precession region appear to already be well below this goal. Additionally, with improved magnetometry relative to ACME II, discussed in section 5.2, as well as the new magnetic field coils, we expect to be able to further suppress magnetic field gradients if necessary. Based on this, we expect that with our new magnetic shields we should be able to sufficiently control these systematics in ACME III.

4.2.2 POLARIZATION GRADIENTS

The remaining systematic effect which we must consider for ACME III is the systematic effect caused by the combination of a non-reversing electric field \mathcal{E}^{nr} and polarization gradients in the refinement and readout lasers. This systematic effect was one of the most significant systematic un-

certainty contribution for both ACME I [6] and ACME II [8], and so is of particular concern for ACME III. As we have previously described in sections 3.1.4 and 3.4.4, the scale of this systematic is dependent on the size of \mathcal{E}^{nr} which causes an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ correlated detuning $\Delta^{\mathcal{N}\mathcal{E}}$, and the size of the polarization gradients in the refinement and readout lasers, which are caused by birefringence gradients in the optical path of the laser, including lenses, a vacuum window, and one field plate. Just as it is with other systematic errors, we do not have plans to reliably reduce \mathcal{E}^{nr} for ACME III, so all of our efforts have gone towards improving the birefringence gradients in the optical pathways.

In ACME I, these ellipticity gradients were primarily caused by thermal stress induced birefringence in the field plates, which were observed to be as large as 10 %/mm [7]. These large gradients were caused by heating of the field plates, which were made of Borofloat glass, from the high laser power required. We were able to compensate for these large gradients in ACME I by aligning the polarization of the refinement laser with the birefringence axis of the optical pathway. In the development of ACME II, the laser polarizations required for STIRAP made it impossible to vary the orientation of the prepared state to align with the birefringence axis. We were instead able to suppress this thermal stress induced birefringence by using a higher quality glass for our field plates (Corning 7980) for which this effect was inherently smaller as described in section 3.1.4. Unfortunately we did still see a systematic effect associated with polarization gradients during ACME II, with a much larger systematic slope in the $\vec{k} \cdot \hat{z} = -1$ configuration than observed for $\vec{k} \cdot \hat{z} = +1$. Based on measurements performed by Vitaly Wirthl [101], we determined that the birefringence gradients were primarily the result of mechanical stress in the field plates and vacuum windows. Unfortunately due to the constraints of our apparatus we were only able to measure the combined birefringence gradients from passing through both field plates and the vacuum windows on either side of the interaction region chamber, which showed that the combined gradients were < 0.2 %/mm. Estimates made by Peiran Hu, based on a linear scaling from the ACME I and II systematic uncertainties and the measured birefringence gradients in ACME I, suggest that the gradients in the $\vec{k} \cdot \hat{z} = +1$ optical

pathway for ACME II may have been roughly a factor of $\sim 5\times$ smaller than was found in this combined configuration. This is possible because there could have been some cancellation between the windows and field plates on the East and West sides of the experiment in addition to the difference in scale for the gradients on the two sides of the experiment.

For the ACME III experiment we have set a design goal that all of the birefringence gradients be ≤ 0.01 %/mm. This was chosen to be an order of magnitude smaller than the gradients measured in the combined two field plate and two window measurements. As the size of this systematic is believed to scale linearly with the size of the polarization gradients we expect this to be sufficient for an order of magnitude improved measurement. Additionally, it should be noted that there is another factor of $5\times$ suppression of this effect that we can expect thanks to the longer precession time for ACME III. This can be seen as this systematic effectively produces a small phase offset $d\theta_{\text{nr}}^{\mathcal{N}\mathcal{E}}$ that contributes to $\phi^{\mathcal{N}\mathcal{E}}$, so in converting to the EDM-correlated frequency this is scaled by τ , which will be $5\times$ larger for ACME III. This is true for both the refinement and readout lasers. While we do not need to consider this effect when comparing our design goal with the measured gradients [131], if the ACME II gradients were closer to the values estimated by Peiran Hu and if we consider this additional scaling, we still expect < 0.01 %/mm to be more than sufficient for ACME III.

NEW GLASS OPTIONS

Following the same chain of logic as we used in ACME II, we determined that one of the simplest paths to reducing these gradients would be to find an alternative glass. As we noted in equation 3.2 the stress-optic law determines the difference between the indices of refraction for \hat{x} and \hat{y} polarized light as

$$\Delta n = K(\sigma_{xx} - \sigma_{yy}), \quad (4.9)$$

where K is the stress-optic coefficient, which is a material property of the glass. So long as the stress in the glass is the same, if we can use glass with an order of magnitude smaller value of K then we could expect to meet our goals. We found two candidate glasses for this purpose, one of which has since been discontinued¹⁵, and Schott SF57HTUltra [132] which is a leaded glass that has a stress-optic coefficient that is specified to be $175\times$ smaller than the Corning 7980 used in ACME II [100]. However, we discovered that this glass is ill suited to our purposes. The first and most important issue is a practical one as the glass is not produced in blanks larger than $280 \times 160 \times 40$ mm. This is clearly an issue as ACME III requires field plates that are at least 1.2 m long (a 1 m long precession region with an additional length l on either end of the region to ensure field homogeneity that should be no shorter than half the height of the field plates). The limits of these dimensions have played a significant role in the designs of our vacuum windows used for STIRAP, the refinement laser, and the readout laser.

Similar to this, we found that in ACME II the beam shaping optics were not made from well controlled glass or mounted with care for the birefringence of the glass. based on work led by Peiran Hu, we intend to use high quality fused silica lenses for ACME III. Additional measurements with the polarimeter have also found improved ways to clamp these optics.

VACUUM WINDOWS

We plan to use SF57HTUltra for the vacuum windows in ACME II that are critical to our measurement, the STIRAP entrance window on the top of the chamber, the refinement laser windows on both sides of the chamber, and the readout laser windows on both sides of the chamber. For the other windows on the chamber, which are in intermediate positions along the beam line, we will use N-BK7 windows as we do not currently have plans for their use and they were included to

¹⁵Ohara PBH56

give us future flexibility. The first major change to these windows besides the choice of glass, is that the windows use circular glass blanks as opposed to the rectangular glass used for all windows in ACME I and II. The primary reason we chose this is that the symmetry of a circular window has been demonstrated to result in significantly lower birefringence than rectangular windows, even in cases where there is more stress in the circular window [133]. The second benefit of moving to a circular window is that the glass can be rotated if necessary.

In ACME II we believe we were lucky that the birefringence axis on the $\vec{k} \cdot \hat{z} = +1$ side of the apparatus was only $\sim 5^\circ$ offset from \hat{x} , which helped to keep our polarization gradients small. For ACME III we have made plans to be able to rotate the glass in our vacuum windows so that the combined birefringence of the window and field plate minimize the birefringence gradient seen by the refinement laser. This is effectively the inverse of what was done in ACME I, where we rotated the polarization of the lasers to match the birefringence axis. As the combined system of optics, one vacuum window, and one field plate can be modeled as a single optic with one birefringence axis, so long as the window birefringence gradients in the window are not significantly smaller than the other components, the window can be rotated to effectively minimize the birefringence gradients. We had originally planned to develop a vacuum window with a rotating seal that could be operated while under pressure. Commercial solutions for this exist, but no commercial options use only non-magnetic materials, so we could not use them in our experiment while maintaining our magnetic field targets. Instead we designed our own low profile rotary vacuum window based largely on the system described in [134], however the full development of this would have required more time than we had available. Instead we intend to rotate the vacuum window only while the chamber is vented, based on measurements with our polarimeter.


For the refinement and readout windows, we intend to use windows with a 15.8 cm diameter glass blank that is 12.7 mm thick (1/2 in), as this is the largest circular size that can be cut from the available glass blanks. The decision to make these windows as large as possible was based on the

size of the holes we placed in the magnetic shields, which were based on our intended laser beam height at the time, which was chosen to be 14.5 cm along \hat{y} (full width). While we now believe that our lasers will not need to be as tall as this, windows this large allow us to take full advantage of the fixed hole size in the magnetic shields. The thickness of the glass was chosen out of concern for the strength of the glass as a vacuum window, and we added additional supports from o-ring cord stock that run vertically along \hat{y} on either side of the small aperture of the window, which has dimensions of $\approx 1.5 \times 5.3$ in, which can be seen in figure 4.15. These supports were designed to use the same o-ring groove design as the actual sealing o-ring and to be under similar compression as that o-ring. The STIRAP entrance window¹⁶, as well as the window on the top of the chamber in the readout region, are designed in a similar way, with a smaller 5 in diameter. This smaller diameter was chosen to be the same as the auxiliary windows on the chamber, as our optical access for STIRAP is already limited to < 6 cm by the separation of the field plates and their guard rings. Tests performed under vacuum with a polarimeter at Okayama University have demonstrated that for the large glass blanks used in the refinement and readout region have birefringence gradients at or below the ~ 0.006 %/mm level. The design and testing of these windows was a collaborative effort between myself, Takahiko Masuda, Ayami Hiramoto, and Peiran Hu.


FIELD PLATES

As the SF57HT ultra is not available in large enough sizes for our field plates we initially investigated alternative options for constructing our field plates. I include these here as the records of what did not work are often as useful as what did work. We considered finding a company to remelt the glass

¹⁶In the design of the vacuum chamber we assumed that the STIRAP entrance window would be on top of the chamber as in ACME II, but it may be simpler in ACME III to have the entrance window on the bottom for ACME III as we do not have an equivalent to the laser lounge at Northwestern. If this is chosen, the window will not have the same supports, but these are less necessary for the smaller window diameter with the same thickness.



(a) Front view of the window aperture and support grooves.



(b) Cross section view of the supported window design.

Figure 4.15: Large Supported Window. Front view and cross section of the supported window design. The grooves parallel to the aperture house o-ring cord stock with the same diameter as the sealing o-ring.

to create a larger piece with the dimensions we wanted, but few companies work with this type of leaded glass and there was no way to ensure that the glass would still have the properties we desire, as they are dependent on the specific annealing process used [135]. We next considered using multiple rectangular panels along \hat{x} both in configurations where they were epoxied together or mounted independently. In both cases our largest concern was with our ability to maintain field homogeneity. We also considered circular insets of SF57HTUltra in a larger blank of soda-lime glass, however this still had issues related to the epoxy at the joint, which was difficult to make sufficiently flat to maintain field homogeneity. We also had concerns that were not fully addressed with regards to how well we could coat epoxies with ITO at the same time as the glass. At various times these investigations were led by myself, Daniel Ang, and Peiran Hu.

Eventually we found that our best option for ACME III was to still use Corning 7980 as in ACME II. To understand this choice we must make a distinction between the two causes of stress that are significant in our system (as thermal stress is not a concern anymore), which are residual stress in the glass and stress from clamping the glass. The residual stress in the glass is a function of the volume of the glass, and the annealing process used to manufacture the glass. As such we have very little control of this stress, and have found that in similar sized pieces of Corning 7980 and SF57HTUltra, the birefringence gradients are similar when the glass is only resting under gravity. This indicates that there must be far more residual stress in the SF57HTUltra than the Corning 7980. Measurements with the polarimeter found that in this configuration the Corning 7980 had birefringence gradients from residual stress at the 0.015 %/mm level, which is close enough to our design goal to work. By contrast the clamping stress is a function of our own designs so we investigated ways to minimize this. We found that almost all of the clamping stress was caused by force components that were tangential to the surface of the glass due to the way the guard rings held the glass in place for ACME II. By sandwiching spheres between the clamps and the glass, there can be no tangential forces applied to the glass, and so this stress is minimized. In tests with soda-lime glass,

the stress-induced birefringence gradients were reduced by an order of magnitude when using this clamping method. From this work, which has been led by Peiran Hu, we expect using this clamping method with Corning 7980 to be sufficient for ACME III.

OPTIONS WE DID NOT IMPLEMENT

We also considered two significant changes to our experimental method that could be used to suppress these effects, both of which may still be useful in future upgrades to the experiment. The first of these changes was to change the STIRAP transitions we used so that we could send the lasers horizontally through the molecular beam along \hat{z} . The primary constraint that prevented us from doing this with $X - C - H$ STIRAP is that the Stokes laser at 1090 nm is not fully saturated even with 10 W of laser power [86]. We were not comfortable putting that much laser power through our ITO coatings, in addition to concerns about the heating effects this would have on our \mathcal{E}^{nr} systematic. With an intermediate state that required less laser power, we could send the lasers in through the field plates without this concern, which would allow us to align the polarization of the lasers with the birefringence axis of the system. Additionally, in ACME II the STIRAP optics were ~ 1.5 m from the molecular beam due to constraints associated with the location of the laser lounge. This distance limited the quality of the beam shaping, and the beams were more sensitive to vibrations in the system. By using horizontal STIRAP we could place the optics right at the magnetic shields, reducing this distance.

We first considered the L or U state as the intermediate state, based on its large transition dipole moments [89], but eventually moved to considering the A state which had been used in ACME I. In particular the UV light required for the $X - L$ and $X - U$ transitions was a concern as we were concerned it could change the patch potentials in the ITO coating. By contrast the A state uses only IR and NIR transitions, and the $X - A$ transition has been well studied in the past [93]. We spent a significant amount of time searching for laser systems that we could use to generate the

estimated $\sim 1 - 2$ W at 1892 nm we would need to saturate the $A - H$ transition. Eventually we were able to find fiber lasers available at this wavelength from AdValue Photonics and Precilasers. However, before implementing this system we would still need to perform tests to determine how much laser power was needed at 943 nm. Due to the varying cost depending on how much 943 nm power we need, we were reluctant to proceed further with this. Our method for estimating the laser power required for $X - A - H$ STIRAP can be found in appendix A. Additional questions about the effect of the absorption of 1892 nm light in the ITO coating could not be answered without experimental testing.

The second option we considered was reintroducing the $\tilde{\mathcal{G}}$ switch from ACME I [6, 7], that rotates the initial prepared state and the readout polarization basis by $\pi/2$ to suppress effects related to polarization gradients. Due to the constraints of the $X - C - H$ STIRAP this would require either replacing this STIRAP with horizontal STIRAP, or introducing a second step of horizontal STIRAP after the $X - C - H$ STIRAP. While we believe this would be an effective method, at this time we expect to have sufficient control of the \mathcal{E}^{nr} systematics for ACME III, and the development of this STIRAP would require too much development at this time.

4.3 GENERAL UPGRADES AND EXPERIMENT ROBUSTNESS

In addition to those upgrades which improve our statistical or systematic uncertainty, we have also made a number of changes to the experiment that improve the robustness and long term stability of our experiment. Many of these are related to the interaction region vacuum chamber, our magnetic shields, and our magnetometry, as described in chapter 5. The remaining changes to our apparatus are largely related to our data acquisition system, our lasers, and related to our new lab space at Northwestern University.

The first of these changes is the work done by Daniel Ang to upgrade our data acquisition system

[123]. The most important change to this system is related to the rate at which we are able to save data. In ACME II we averaged together every 25 shots from our molecular beam into a single trace, and only saved this trace. We found that this averaging makes it difficult to investigate some properties of our data, which was the case for our excess noise from timing jitter. During the timing jitter investigation we temporarily saved every shot instead, and were able to better understand what was happening. For ACME III we now have the capability to always save every shot, or average shots together if desired. This required a significant overhaul of the entire system to allow for faster saving, more data storage, and faster communication with our data storage and between our computers.

The second category of changes in our system is related to the stability of our lasers. As we described earlier in section 4.1.2, we have now switched over to using lasers that are referenced to a ULE cavity for our rotational cooling. With this change, in addition to new considerations for referencing lasers to our frequency comb, we no longer will have any lasers in our experiment that are locked to transfer cavities. This extends the typical time that our lasers can stay locked from ~ 1 day to ~ 1 week, which is very important for our ability to operate continuously, while also improving the linewidth associated with the lock. In addition to this, we have replaced the pump laser for our Ti:Sapph laser in the hopes of reducing the amount of maintenance required to maintain both the output power of the laser and our lock.

The remaining category of changes to the experiment are those related to our move to Northwestern University and the differences between our new lab space and our space at Harvard University. The most consequential change here is that we need to develop new systems to block stray light from the room from entering the experiment and being detected by our SiPMs. Our new lab space is a shared space with other experiments, and so we do not have the ability to turn off all the lights in the room as we did in ACME II, which meant we only needed a small amount of black cloth over the detectors themselves. We have developed a system for blocking stray light access in the SiPMs, the light pipes and their mating to the vacuum chamber, as described further in section

5.3. This system has been demonstrated to have no difference with room lights on or off. However, further development is still ongoing to block stray light from entering through the windows we use for laser access. This system requires enclosing our laser tables, and coupling these enclosures to the vacuum chamber in a light tight way. One significant benefit of our new lab space is that the quiet laser lab is in an adjacent room to our main apparatus, rather than in a separate building as was the case for ACME I and ACME II. We can now couple lasers to our experiment using 30 m fibers instead of the 100 m fibers used previously, which will significantly reduce the power loss when transferring to the experiment. We have also purchased a new YAG laser for ablation, which has a smaller footprint, and does not require a connection to building water, which fit better in our available space at Northeastern, as we have no full sized laser tables in the same room as the main apparatus. Finally, we have also implemented two new commercial rack mounting systems¹⁷ which will house seven diode lasers and two tapered amplifiers, to account for the limited space available to us in our new laser lab.

¹⁷Toptica T-RACKs

900 cc's of raw, whining power

No outstanding warrants for my arrest

Hi diddle dee dee, goddamn

The pirate's life for me

The Mountain Goats, *Jenny*

5

Extending the ACME III Interaction Region

SINCE THE CONCLUSION OF ACME II, a significant focus of our lab has been on developing a new interaction region to increase our precession time as well as improve our magnetic field control for ACME III. In this chapter we will first describe the recent measurement of the H state radiative

lifetime that motivated this change to our precession time and the need to redesign our apparatus. We will then discuss the design and testing of our new magnetic shielding system, as well as our new magnetometry development. Finally, we will discuss the development of the interaction region vacuum chamber and the associated systems.

5.1 A NEW MEASUREMENT OF THE H STATE LIFETIME

After ACME II had been wrapped up, our group performed a pair of measurements of the H state lifetime, which were led by Cristian Panda, using the ACME II apparatus and the Q state measurement apparatus on Beam Box II. These preliminary measurements suggested the lifetime of the H state was most likely ~ 4 ms, and motivated us to develop a dedicated measurement apparatus on Beam Box II starting in the spring of 2019. This work was led by both Daniel Ang and myself, and was eventually reported in [70]. The general principle of this measurement, as was the case in the preliminary measurement, relied on exciting molecules into the H state at a fixed excitation region, then allowing the beam to travel to a detection region a known distance away, and measuring a fluorescence signal in this region. This fluorescence signal scales as e^{-t/τ_h} , where t is the travel time between the excitation and the detection. By performing this excitation in multiple regions, we can fit the fluorescence results to extract the radiative lifetime τ_H . Notably this measurement took place both before and after the 2020 COVID shutdowns of our labs, and so we collected two data sets under slightly different conditions which we will describe.

5.1.1 METHODS AND APPARATUS

For this measurement we used our test beam source Beam Box II, which had originally been used to test the thermochemical source developed by Elizabeth Petrik West [102]. This was converted into an ablation source in 2017 by Jonathan Haefner and myself, and later used for measurements




Figure 5.1: *H* State Lifetime Measurement Schematic. Diagram showing the outline of the lifetime measurement. Molecules are excited to the *H* state at one of the five excitation regions shown as E1-E5, which are separated by the distances shown. They are then excited to the *I* state by the probe laser and decay back to *X* emitting 512 nm photons that we detect. The coordinate system used for this measurement is shown to the right of the figure. Figure adapted from [70].

of the *Q* state by Xing Wu [87]. This beam source, which operates at 50 Hz uses the same design as the ACME II cell, and the beam flux and forward velocity were measured to be comparable to the beam from Beam Box I, however this source has more cooling power from two pulse tubes rather than the single pulse tube used on Beam Box I. We changed out the ablation target between the two data sets, which notably changed the average forward velocity of the beam from 210 m/s to 250 m/s. We used this beam source instead of Beam Box I and the ACME II beamline for these measurements because of the flexibility and space needed to create a new beamline vacuum chamber that had optical access at multiple positions over a much larger distance than had been used for the preliminary measurements. Additionally, we were able to build both this vacuum chamber and the electric field plates used for this measurement as extensions to the existing beamline chamber and field plates built for the *Q* state measurement [87]. The chamber was built from a combination of ISO100 crosses, as well as two modular vacuum chambers from Ideal Vacuum Inc that provided a large area window on the bottom surface for optical access that we used for detection. A schematic view of the measurement apparatus can be seen in figure 5.1.

Once the ThO molecules exit the buffer gas cell, they pass through a 690 nm *X* – *C* laser that

we use to collect an absorption signal to determine the time the molecular beam passed this point. This laser was a commercial Toptica DL Pro referenced to a ULE cavity as we have described previously in section 3.2.2. The molecules next pass through a pair of adjustable collimators made from razor blades mounted to vacuum translation stages, approximately 48 cm after the cell exit. These collimators, along with the aperture of the cell, are used to set the size of the molecular beam in our experiment, as well as set the range of transverse velocities and thereby control the Doppler width for our lasers. In an ideal measurement we would make our beam as small as possible to maximize our ability to saturate the transitions we use, and reduce the possible effects of systematic errors related to imperfect saturation, but this comes at a significant cost to our signal level. Based on these constraints, we typically operated with a 6 mm horizontal aperture and a 3 mm vertical aperture, which corresponded to a measured 1σ Doppler width of 6 ± 1 MHz for our 943 nm excitation laser.

After the molecules pass through the collimators they enter the electric field region, which is defined by a pair of aluminum electric field plates. The field plates, which are 1.32 m long, consisted of 6061 aluminum plates separated by 1.5 in. using PEEK spacers. These field plates were made up of four segments along \hat{x} that were mechanically and electrically connected, which were assembled separately in part because they were partially built to extend the two segments that already existed, and to make their assembly in the chamber easier. For this experiment we maintain some of the conventions of the ACME right handed coordinate system described in section 2.3 as we define $+\hat{x}$ as pointing along the direction the molecular beam propagates, and we define $+\hat{z}$ as the typical applied electric field direction, which here points vertically up, in contrast to the typical ACME coordinates. This defined \hat{y} as pointing horizontally, perpendicular to the beamline, and all of our lasers propagated along $\pm\hat{y}$, which determined the geometry of the field plates. The field plates had two viewports on the bottom plate, located 42 cm and 129 cm from the start of the field plates, which

were 1 in. square ITO coated windows with Thorlabs B AR coating¹. These viewports were epoxied into the field plates to be flush with the inner surface² using a conductive silver epoxy, and we verified the electrical connection between the glass and the aluminum. The viewport closer to the beam source was used occasionally for detection as a systematic check, while the farther viewport was used as the primary detection location. The electric field plates were typically charged to 38 V/cm using a Thorlabs 3-channel Piezo driver³, which was sufficient to fully polarize the $H, J = 1$ state.

Once in the electric field region, the molecules travel freely until they arrive at one of five excitation regions (labeled E1-E5) where a portion of the molecules are transferred to the H state using the ACME I state preparation method. At one of the excitation regions the molecules pass through a 943 nm laser that optically pumps the $|X, J = 0\rangle \rightarrow |A, J = 1\rangle$ transition. From the excited A state molecules spontaneously decay, with a branching ratio of ~ 0.35 to the $|H, J = 1\rangle$ state as we previously described in section 3.1.1. Of those molecules which decay into H , $1/6$ of the population decays into each of the four $|H, J = 1, \tilde{N} = \pm 1, M = \pm 1\rangle$ states, and the remaining $1/3$ decays to the $|H, J = 1, M = 0^-\rangle$. We use a single laser setup for all five excitation regions, which is based on a Toptica DL Pro at 943 nm that is locked to a transfer cavity referenced to our iodine clock using the same technique as the ACME II rotational cooling lasers described in section 3.2.2. This resulted in a typical linewidth of ~ 2 MHz for this laser. This laser was then coupled into a commercial tapered amplifier before going to our switching mechanism. The switching mechanism consisted of a series of AOMS that shared the same zeroth order beam, such that by turning on one AOM at a time we could send light from the first order beam of that AOM through a fiber to exactly one of the excitation regions and none of the others. With this system we had a laser power of 65 ± 7 mW, which we then quadruple-passed through the molecular beam with a pair of offset right

¹These were discontinued just as we built this system.

²The windows were thicker than the aluminum stock used for the field plates.

³Thorlabs MDT693B, with a max output of 150 V.

Excitation Region	Distance from Detection Region L_i (cm)
E1	18.0 ± 1.0
E2	43.5 ± 1.0
E3	65.0 ± 1.0
E4	87.0 ± 1.0
E5	121.5 ± 1.0

Table 5.1: Lifetime Measurement Excitation Region Positions. Table showing the positions of the lifetime measurement excitation regions relative to the detection region.

angle prisms. This increased the effective power to ~ 230 mW in order to improve our saturation of the $X - A$ transition. We used an independent optical setup at each excitation region to align the lasers, expand the beams to a $1/e^2$ height of ~ 1 cm, and quadruple-pass the beams. We measured the height of the molecular beam in these regions using razor blades clipping the excitation laser and determined that it was ≤ 0.8 cm. The excitation regions were spaced out by the distances shown in table 5.1

After passing through the excitation region, the molecules travel a distance L_i , where $i \in \{1, 2, 3, 4, 5\}$ refers to which excitation region laser they passed through, before arriving at the detection region. Over this distance the molecules in the H state undergo radiative decay back to X , leaving a population in the H state of

$$P_i = P_0 \eta_i e^{-t_i/\tau_H} = P_0 \eta_i e^{-L_i/(\tau_H v)}, \quad (5.1)$$

where P_0 is the initial population excited with efficiency η_i at excitation region E_i , t_i is the time of flight from the excitation region, which is set by the forward velocity v of the molecular beam.

In the detection region we use a fixed 703 nm probe laser that is linearly polarized along \hat{x} to excite the $|H, J = 1, \tilde{N} = +1\rangle \rightarrow |I, J = 1, M = 0^+\rangle$ transition. We used the same Ti:Sapph laser locked to a stable ECDL for this detection as was used for the refinement and readout laser used in

ACME II, which was described in section 3.2.4. During this measurement the readout laser had a typical operating power of 160 mW during the collection of the first data set, and 240 mW⁴ during the collection of the second data set. After excitation the molecules spontaneously decay back to the ground state, emitting 512 nm photons which we focus onto a PMT placed directly under the viewport in the detection region. We should note here that this measurement had a much lower signal to noise ratio than ACME II, most importantly because we were effectively using the ACME I optical pumping state preparation method rather than STIRAP, in addition to an overall lower collection efficiency due to the limitations of our collection optics geometry. By using a shared detection region we are less sensitive to a number of possible systematic effects, especially those related to saturation that we saw with the excitation regions. We then average this data over 64 shots in the first data set, and 32 shots in the second data set. We also use the average arrival time in the detection region along with the average arrival time in the $X - C$ absorption laser just outside the cell to compute v for each trace as the two positions are separated by a known distance of 178 ± 1 cm.

From the collected data we subtract off the background signal, which is averaged over the start and end of the trace, and integrate the non-background signal to compute the fluorescence intensity I_i . We use the fluorescence intensity from the highest signal excitation region, E1, to normalize the measurements from each of the four other excitation regions. This allows us to remove any effects caused by the decay of the signal associated with each ablation spot over a period of a few minutes, as well as any common mode population effects, as we effectively assume that $\eta_1 = \eta_i$ for all excitation regions. In order to do this normalization we switch between E1 and one of the other excitation regions every 7 – 9 s, and use the data from these two states to compute the ratio I_i/I_1 . We

⁴Our power here was limited by our ability to couple the laser into a single mode fiber, as the Ti:Sapph spatial mode degraded after the end of the ACME II measurement. Part of the change between the two data sets was that we had the laser serviced, which improved the spatial mode.




Figure 5.2: *H* State Lifetime Measurement Fluorescence Intensity Ratios vs Transit Time. Semilog plot of the fluorescence intensity ratios vs the transit time shown for E2-E5. The two data sets, which had different average beam velocities are shown in orange for the earlier data set, and blue for the later data set. The fit using equation 5.2 is shown with the black line. Figure adapted from [70].

then fit these ratios to

$$\frac{I_i}{I_1} = \exp\left(\frac{L_1 - L_i}{v\tau_H}\right). \quad (5.2)$$

From this we can extract the measured value of the radiative lifetime τ_H . The two data sets and their fit using this equation can be seen in figure 5.2. The result of this did not change when we performed a number of other analysis methods, which included analyzing the two data sets separately, analyzing the data while excluding any one of the excitation regions, and using the alternative fit function $Ae^{(L_1-L_i)/v\tau_H}$, which allows for different initial populations in the excitation regions.

5.1.2 SYSTEMATIC UNCERTAINTIES

In order to search for systematic errors in our experiment, we varied a large number of experimental parameters, and determined their contributions to our error bar as shown in table 5.2. As the majority of these parameter variations did not show any significant effect, we include only the four error sources from saturation effects in the excitation lasers, background uncertainty, the position of the excitation regions, and uncertainty from our velocity determination. The parameter variations which were not included in our error bar include changing the size of the adjustable collimators, shifting the center of the collimator aperture, varying the applied electric field, exciting different sub-levels of H , and saturation related effects in the readout laser.

Source of Uncertainty	Uncertainty (ms)
Excitation laser saturation	0.4
Background uncertainty	0.2
Excitation laser position	0.04
Velocity determination	0.05
Total systematic uncertainty	0.5
Statistical uncertainty from fitting	0.02
Total uncertainty	0.5

Table 5.2: Lifetime Measurement Combined Error Budget. Systematic and statistical error budget for the measurement of τ_H . Uncertainties are shown in units of ms.

The largest contribution to our error budget comes from errors relates to imperfect saturation of the excitation laser in different excitation regions. This is a particular concern in our apparatus as we have separate optical setups including beam shaping, alignment, and quadruple-passing optics for each individual excitation point. Additionally, the molecular beam is diverging, and so the size of the molecular beam varies between excitation regions. Finally, the dispersion of the beam forward velocity changes the shape of the pulse between each excitation region. Ultimately, this can result in variation between the excitation efficiencies η_i for each excitation region despite our assumption in equation 5.2 that $\eta_i = \eta_1$ for all excitation regions. For future measurements the most promising

upgrade path we considered would be to use a single optical setup which could be moved from one excitation to another mechanically so that saturation would be more uniform between excitation regions.

In order to study this effect we varied the collimator size, measured the size of the molecular beam, performed Doppler scans at each excitation region, and performed power scans at each region. The Doppler scans in particular showed that we were sensitive to the angle between the excitation lasers and the molecular beam as we could observe a double peak in the Doppler scan by misaligning the laser slightly. We additionally, used razor blade cuts in the excitation laser to ensure that the lasers were correctly aligned vertically relative to the molecular beam. From the power scans at each excitation region we were able to fit the fluorescence intensity as a function of laser power P in the form

$$I_i = I_{i,\max} \left(1 - e^{-P/P_S} \right) \quad (5.3)$$

where P_S is the saturation laser power and $I_{i,\max}$ is the expected maximum signal. From this we computed the excitation efficiency for each region as $\eta_i = I_i/I_{i,\max}$ assuming the typical laser power used in the measurement and we found that the excitation efficiencies varied by up to 7 %. Based on a simulation performed by Daniel Ang alongside this analysis, we found that excitation efficiencies with this level of variance could result in variation of the measured value with a standard deviance of 0.4 ms, which we assigned as the uncertainty for these effects.

The second largest contribution to our error budget came from uncertainty in the background level of our signal. A non-zero background could in theory be caused by the population of some unknown metastable reservoir state, possibly a higher vibrational state, that is populated during the ablation process, which then decays into the H state over time. This has been theorized as a potential systematic effect in previous lifetime measurements [92]. Such a decay would result in a constant offset in fluorescence intensities that is independent of excitation region, and could be seen

as a molecular pulse signal when the excitation lasers were turned off. In practice we did not directly observe any such signal during data taking, possibly limited by our signal to noise level. We placed a bound on this effect by considering data that was taken while detuning the excitation laser far from resonance, showing that this background effect was less than 3 % of I_1 , which is equivalent to a 0.2 ms uncertainty in the lifetime.

The third contribution to our error bar, which is quite small comes from our uncertainty in the position of each of the excitation regions. This uncertainty in the positions is limited by the fact that we quadruple-pass each excitation laser to improve saturation, which substantially increases the effective size of the regions. This distance is then used to extract the lifetime τ_H as shown in equation 5.2. Notably, this has a much larger effect on the data taken from excitation regions that are closer to the detection region, as the fractional uncertainty in the distance is larger, however the data from the closer regions is also weighted less than the farthest excitation regions in the overall fit of the lifetime. The result of this is only a small error component of 0.04 ms that does not significantly change our result.

The fourth and final contribution to our error bar comes from the uncertainty in the measurement of the transit time between the absorption laser and the detection region which we use to measure the forward velocity of the molecular beam v . This effect is shown in figure 5.3. This uncertainty is caused by the dispersion in v that causes the molecular beam pulse to spread out in time between the two locations from ~ 0.5 ms long in the absorption laser (shown in blue) to ~ 2 ms long in the detection region (shown in orange). When computing the arrival times for these pulses we found that there was typically a 0.1 ms difference between the results when computing the arrival time by considering the mean arrival time (solid lines) vs the arrival time of the peaks of the pulses (dashed lines). We ultimately calculated the arrival time from the mean arrival time, but found that this discrepancy could result in an uncertainty of 0.05 ms in the measured lifetime.

We combined these errors using standard uncorrelated error propagation alongside the 0.02 ms




Figure 5.3: H State Lifetime Measurement Pulse Shapes and Arrival Times. Figure showing averaged pulse shapes for the absorption laser (blue, left) and the fluorescence detection (orange, right). The arrival times for based on the peaks of the pulses are shown with dashed lines, while the arrival times based on the average of the pulse are shown with solid lines. The two different methods differ by 0.1 ms. The absorption trace has been scaled for comparison. Figure adapted from [70].

statistical uncertainty extracted from the fitted data. This resulted in a measurement with 1σ uncertainty of

$$\tau_H = 4.2 \pm 0.5 \text{ ms.} \quad (5.4)$$

5.2 MAGNETIC SHIELDING AND MAGNETOMETRY

We took the requirement for new, larger magnetic shields to enclose a longer precession length as an opportunity to completely redesign our magnetic shields. As we described in section 3.3.1, we discovered during ACME II that the magnetic field inside our interaction region had increased over time to the $\sim 100 - 200 \mu\text{G}$ level, which we believed was primarily caused by the remanent magnetization of the innermost shielding layer despite our best efforts to remove this field with degaussing. With this concern in mind, alongside our expectations for our conservative ambient magnetic field requirements of $\leq 1 \mu\text{G}$ (with $\leq 10 \mu\text{G}$ likely to be acceptable), and our goal that all magnetic field gradients be $< 1 \mu\text{G}/\text{cm}$, we set out to design, construct, and characterize a new three-layer mag-

netic shield system. This process, as well as the development of magnetometry methods to study this system, and the development of degaussing systems have been a multiyear project which at various stages has been led by different members of our research group, including myself.

5.2.1 SHIELDING THEORY

In order to achieve our ambient field goals in the interaction region we rely on passive magnetic shielding. While electric fields can be easily screened out and high enough frequency ac fields are shielded by the skin effect, low frequency and static magnetic fields require specialized materials and careful design. Static magnetic shielding in the style we use in the ACME experiment relies upon enclosing an area with multiple layers of a metal alloy⁵ that has an extremely high permeability. For ACME II we used Amuneal [108], while for ACME III we are using Co-NETIC AA [136] provided by Magnetic Shield Corp⁶, both of which are Nickel alloys. It should be noted that the permeability of these alloys is dependent both on the magnetization history of the material, and the strain that the material has experienced, as strains are well known to lower the permeability of the material [137]. The high permeability of the μ -metal effectively draws magnetic flux into the metal, where the flux is guided tangentially around the shielded volume with little leakage, before exiting the μ -metal on the other side of the shielded volume. In a perfectly enclosed region, the interior field would be governed entirely by leakage, however any realistic system will have gaps and intentional access holes that also provide a pathway for magnetic flux. This behavior can be seen as a consequence of Gauss' law and Ampère's law, which require that the normal components of \vec{B} and the tangential components of \vec{H} at the surfaces be continuous [137, 138, 139]. In practice, this can cause the magnetic field lines to change direction sharply at the boundary, entering from

⁵Typically referred to as μ -metal regardless of trade names.

⁶This was chosen at least in part due to the proximity of Magnetic Shield Corp to Northwestern University as we consulted with them in the early stages of our design.

the air nearly perpendicular to the surface, and flowing tangentially inside the thin metal. Perhaps more intuitively, this can also be viewed as analogous to an electrical circuit, where the high permeability metal has a very low magnetic reluctance, which is analogous to resistance, and magnetic flux is analogous to electrical current [140]. This concentration of flux inside a small volume of metal can result in saturation, which reduces the shielding of the system, however, for the scale of external field and the scale of our shielded volume we do not expect even our outermost shields to be near saturation.

We can characterize a shielding system by defining the shielding factor as

$$S = \frac{|\mathcal{B}_o|}{|\mathcal{B}_i|} \quad (5.5)$$

where \mathcal{B}_o is the magnetic field which would exist in the absence of the magnetic shields, which we typically estimate from the external field, and \mathcal{B}_i is the field inside the shielded region. In general, this shielding factor is a function of the frequency of the magnetic field, and depends on the geometry of the shielding, as well as the angle between the incident magnetic field and the shields if they are not symmetric. We neglect the frequency here, as we are primarily concerned with dc fields and quasi-dc changes to the external field⁷. However, the shielding factor for dc fields and quasi-dc fields will still differ, as the shielding factor depends on the permeability of the metal, which itself is a function of the applied field and our degaussing procedures [137].

Analytic solutions to the shielding factor exist for spherical shields, infinite cylinders, and infinite planar shields, and approximate solutions with analytical forms have been found for cylindrical shields with end caps [137, 138], but more general configurations require numerical estimation. For more general closed shapes with wall thickness t and a characteristic size for the enclosed volume D

⁷We are primarily concerned with these as our lab at NU houses a “fast solenoid” a few meters away which can ramp between $-6 \leftrightarrow 6$ T over the course of roughly half an hour to an hour.

we can estimate the shielding as

$$S \sim \eta \frac{\mu t}{D} \quad (5.6)$$

where μ is the relative permeability of the shield material and η is a geometric scaling factor.

To understand the interaction of multiple shield layers, as is typically the case, we can consider the specific case of two infinite cylindrical shields [137, 140]. For two thin, high permeability, concentric cylindrical shields with cross sectional areas A_1 and A_2 , and transverse shielding factors S_1^T and S_2^T , we can write the shielding factor as

$$S^T \approx 1 + S_1^T + S_2^T + S_1^T S_2^T \left(1 - \frac{A_1}{A_2}\right). \quad (5.7)$$

This formula can be extended generally to n shields, and a similar formula exists for the axial shielding of the cylinders. Two key takeaways from this formula are that the individual shielding factors do not simply add linearly, and in the limit of $A_1 \rightarrow A_2$, which would be the equivalent of making a single shield layer twice as thick, the total shielding factor is smallest. From this it is clear that multiple well spaced shields provide greater shielding than a single thick shield.

5.2.2 SHIELD DESIGN CRITERIA

When we set out to initially design our magnetic shield system we based our work on the following criteria, which are listed roughly in order of priority.

- We wanted to ensure that the ambient field within the interaction region volume, spanning 100 cm and a diameter of roughly ~ 5 cm, was below at worst $10 \mu\text{G}$, and ideally below $1 \mu\text{G}$. We additionally required that all gradients be below $1 \mu\text{G}/\text{cm}$. These technical limitations were chosen for the reasons discussed in section 4.2.1 related to our magnetic field noise and magnetic field gradient systematics from ACME II. As the ambient field inside

our lab is typically on the order of ~ 1 G or lower, this set our target shielding factor on the order of $\sim 10^5 - 10^6$.

- We wanted a system that would put the μ -metal under as little stress as possible, both during the assembly and disassembly of the system and in the assembled configuration. This was a lesson we learned from the ACME I and II shields, which were stressed both during assembly and in their normal configuration. This is crucial as any stress can lower the permeability of the material, and we believe this contributed to the increased ambient field we observed in ACME II. The previous shields consisted of five layers, where each layer was made up of two half cylinders and two fixed end caps as shown in figure 3.2. The only supported portion of the shields were the end caps, which the half cylinders were bolted to in addition to bolting to each other at their seams, which put the center of the cylinder under stress. Additionally, during the assembly and disassembly, the half cylinders were flexible enough that they could bend easily, even with their plastic support ribs to help maintain their shape, and in order to bolt them together they had to be flexed slightly away from their equilibrium shape, leaving them under stress. Finally, to hold the half cylinders in place while tightening them we relied upon a pair of turnbuckle straps, which applied localized stress during this process.
- We wanted a system that could be conveniently annealed again if we saw an increase in the ambient internal field over time which we could not remove with degaussing. The annealing process requires a hydrogen furnace, and while we initially considered designs that would allow us to use the hydrogen furnace in the lab at Northwestern, we ultimately decided to design based on the dimensions that can be annealed in a commercial furnace located nearby⁸. This also required that we design shields that can be reasonably packed in a way where they will not be stressed. This constraint and the previous constraint were our primary motiva-

⁸Exotic Metal Treating in Indianapolis, IN.

tions for designing a rectangular prism shield system rather than another cylindrical system.

- We wanted a system that would use only the amount of μ -metal we actually needed to achieve our technical goals. As μ -metal is both expensive and dense, this goal worked to reduce unnecessary costs and reduce the total mass of our system, which helps both lower the constraints on the support structures for our shields and to make the shields easier to manipulate. Notably, while investigating our increased ambient internal field during ACME II, we confirmed that there was no significant change associated with removing the outermost two layers of shields, indicating that we could hopefully use only a three layer system for ACME III. Additionally, in line with using less μ -metal, due to the limited lab space at Northwestern we wanted to ensure that the shields did not occupy a larger footprint than was necessary.
- We wanted to ensure that any μ -metal panels were joined in such a way that at every joint there was at least 3 cm of overlap, and that all seams were covered in such a way that no linear seams remained and the only gaps were small points. The effects of these overlaps have been well studied, as most large shield systems are made up of multiple pieces of either flat or cylindrical sheet metal that must be joined [141, 142, 143, 144]. We considered multiple methods for this, including using μ -metal foils to cover gaps, but ultimately decided to use bolts with cover panels, and bent edge pieces. We left the corners of the shields open as our prototyping found corner pieces had minimal effects as the holes without them were smaller than many of the other access holes in our shields, and the open corners made attaching degaussing coils more convenient. Additionally, we spaced out the bolts attaching the shields by no more than 10 in. and often 6 in. in order to keep the cover plates and edge covers as flush and well mated with the other plates as possible.
- We wanted a system that could be conveniently disassembled and reassembled safely and

quickly. More time and work required to assemble or disassemble the shields increases the barrier to making changes to any of the shields or the apparatus housed inside of them. This required that any parts requiring a crane be able to be lifted by the crane built above the interaction region, both in terms of weight and vertical clearance as we do not have high bay space like we did for ACME I and II. It also required that any parts that must be moved by hand be sufficiently light that they can be safely carried by ideally one or two people.

- We wanted to ensure that there was space built into the frame and μ -metal pieces so that we could easily run magnetic field coils for use in degaussing the system. In the rectangular geometry we chose this meant being able to have coils loop around every edge of the shield as well as two locations for loops in each direction around every face of the rectangular prism. This constraint was based in part on existing techniques for demagnetizing magnetically shielded rooms with similar geometry to ours [145, 144, 146].
- We wanted to ensure that we had sufficient access to the inside of the shields, which required us to place holes in the magnetic shield layers. These holes needed to be as small as possible as they serve as a way for external fields to leak into the shielded region, while still allowing sufficient access. The sizing of the largest holes was an iterative process between designing and simulating the effects of the holes. The list of holes we designed for included the following which are roughly ordered by size:
 - Holes for the refinement and readout lasers on the North and South sides of the shields.⁹ These four holes are the largest holes in our shields, and are rectangular with a height of 14.5 cm to allow us to expand our readout and refinement lasers vertically

⁹Up until now we have not needed to define the coordinate system of ACME III, which with respect to the apparatus is the same as ACME II as described in section 2.3. However as we are in a new lab space for ACME III, the coordinates are no longer the same with respect to the cardinal directions. $+\hat{x}$ now points East, and $+\hat{z}$ now points North.

to fully engulf the molecular beam. While we can clip the lasers closer to the molecular beam size, we initially planned not to do this based on observations during STIRAP development where clipping the beam caused Airy patterns to appear in the laser intensity profile [69]. This height was chosen based on early estimates of the molecular beam size after the electrostatic lens, but we now know that the beam size will be smaller than we initially thought.

- Holes on the upstream and downstream faces of the shields to allow ISO100 nipples for the beamline to enter and exit the shielded region. These are the second largest sets of holes, with a diameter of 10.8 cm. We later decided to use custom ISO100 nipples with a smaller diameter pipe to increase the clearance with these holes.
- Circular holes for the STIRAP windows on the top and bottom of the chamber near the readout region, as well as the 14 equally sized 3 in. diameter auxiliary windows on the chamber that provide optical access to the beamline if needed.
- Holes to allow the light pipes to connect the collection optics to the SiPMs just outside the shields. As we will discuss in section 5.3.4, we chose to use eight straight light pipes that each exit the chamber at a 45° angle. These holes are slots because of the angle the cylindrical light pipes pass through them¹⁰.
- Holes on the bottoms of the shields for the six support posts that the each of the shield layer frames rest on (for 1 in. diameter posts), as well as the six support posts for the magnetic field coil frame and the six support posts for the interaction region vacuum chamber (2 in. diameter posts).
- Pairs of holes near the readout and refinement windows as well as the STIRAP win-

¹⁰An oval is less complicated to machine and design than the actual elliptical projection of the cylinder onto the plane of the shield.

dow and the window above the readout window. These were designed to be used for inserting a tool to rotate our planned rotary windows and to readout the angular position of these windows. However, we currently do not plan to use an in situ rotary window so at this time these holes no longer have any function.

- Holes to insert magnetometers into the magnetometer pockets in the interaction region from the top surface and the North and South surfaces of the shields, as well as additional holes that would allow us to insert magnetometers between the shield layers.
 - Holes in the corners and in the edge covers that allow us to wind degaussing coils around the faces of the shields.
 - Holes in the μ -metal so that we could attach the plates to the frames with $5/16''$ –18 bolts.
 - Small holes that can be used to secure degaussing coils on the faces of the shields, these are the smallest holes in the shields with a diameter of 0.2 in..
- We wanted to specifically ensure that the bolts used to attach our μ -metal would not be a source of stress on the metal. In ACME II, the bolts were tapped into stainless steel parts welded into the μ -metal [147], which we did not want to replicate. Instead we ensured that the bolts all attached directly to the support frames, and we deliberately oversized the holes in the μ -metal so that the only contact with the μ -metal should be the flat aluminum surface of the frame and the washers on the bolts. After our initial test assembly, we determined that we needed to expand this clearance on the bolt holes, so almost all of these holes were expanded by hand at the same time as we added more bolt holes. One consequence of this hand drilling is that the position of many of the holes is not as well toleranced as the positions of the holes cut with the waterjet, and so care must be taken when assembling the shields to ensure that

the plates are positioned to minimize contact with the bolts.

- We wanted to minimize the distance between the exit of the electrostatic lens and the refinement laser, which was a paired design constraint with the interaction region vacuum chamber.
- We tried to maintain as much symmetry as possible in the design, both to limit the number of unique parts required, reduce possible gradients caused by asymmetries of the shields, and to eliminate as much future confusion as possible. We ultimately introduced two sources of asymmetry in our design, the first of which is related to the set of holes originally intended for turning the rotary windows. These holes break the reflection symmetry about the $\hat{x} - \hat{z}$ plane on the North and South faces, as well as the reflection symmetry about the $\hat{x} - \hat{y}$ plane on the top surface of the shields. The more significant asymmetry is that the center of the precession region is not centered in the shields along \hat{x} , and is ~ 15 cm closer to the upstream (West) surface of the outer layer than the downstream surface of the outer layer. This asymmetry was the result of two factors, the first of which was our desire to minimize the distance between the refinement laser and the lens, and the second of which was our decision to use straight light pipes. We needed to add length to the shields on the downstream end to prevent a conflict with the shield support frames and the light pipes, but we did not want to add this additional length on the upstream side of the shields.

5.2.3 THE ACME III MAGNETIC SHIELDS

The ACME III magnetic shields consist of three nested layers of rectangular prisms constructed out of high permeability Co-NETIC AA [136] sheets which are 1.6 mm thick, as shown in figure 5.4. This configuration was chosen for the ability to be constructed from flat sheet metal that can be mounted with minimal stress, and is similar to the work which has been demonstrated with magnet-




Figure 5.4: The ACME III Magnetic Shields. Vertical cross section of the ACME III magnetic shields. The shields are made from three nested rectangular prisms. Different types of μ -metal components are shown in false color to distinguish them.

ically shielded rooms [144, 146, 145]. Additionally, these flat sheet metal panels can be annealed in a commercial hydrogen furnace without the need for jigs to maintain curvature or control stress, and can easily be flat-packed for safe transport to and from the annealing facility. The plates were cut on a commercial waterjet, and the covers we use on the edges of the prisms were then bent to 90° .

The outermost layer has dimensions $2.26 \times 1.57 \times 1.57$ m, with the long axis aligned with the molecular beam along \hat{x} . The innermost layer has dimensions $1.86 \times 1.17 \times 1.17$ m, with a spacing of ≈ 10 cm between each layer. As can be seen in figure 5.5, each face of every layer consists of two large flat plates, which are mounted next to each other, while the seam between them is covered by one patch plate (two cover plates on the upstream and downstream surfaces due to the beamline nipples). Each edge of the prism is then covered by a single edge cover, which overlaps with the adjacent flat plates by roughly $\sim 3 - 6$ cm depending on the layer. Similar to the ACME II shields, the majority of the critical holes in the shields are placed in the slightly smaller patch plates, which are






Figure 5.5: ACME III Magnetic Shield Panels. Face view showing the North side of the magnetic shields. The side plates (beige) are first mounted to the frame, then the seam between the two plates is covered by the patch plate (green). Each of the edges is covered by a 90° bend edge cover.

easier to machine and anneal again if changes are necessary.

The μ -metal plates for each layer are mounted onto an aluminum frame, which is made up of an upper “U”-shaped frame and a flat base frame as shown in figure 5.6. The upper frame is made up of 60 series Minitec T-slotted aluminum extruded profiles (60 mm on a side), with solid aluminum runners at the bottom edge of the frame. We chose this particular extruded profile based on the availability of an “L”-shaped profile, which allows us most of the rigidity of a 60 mm member, while letting us use the 30 mm \times 30 mm negative space on one corner of the profile to create a pathway for our degaussing cables around every edge of the frame. The base frame is instead made up of solid aluminum bars in order to provide more rigidity to the frame, as this frame then rests on six 1 in. aluminum support posts that attach to the ACME table. Each of the frames for the three layers are mechanically independent of each other, except in that they are each separately mounted on the ACME table.



(a) The upper frame of the inner layer magnetic shield mounted to the ACME table.



(b) The upper frame of the inner layer magnetic shield separated from the base of the shield layers as designed to be removed with the crane.

Figure 5.6: Magnetic Shield frames. Diagrams showing how the upper frame of the magnetic shields is mounted on top of the base frame, and can be removed to access the shielded region.

We chose this split frame for each layer to simplify the process of accessing the shielded region, as there must be a split in the shields somehow so that we can access the shielded region to install or access the interaction region chamber, whose beam pipes extend out through the shields. This design allows us to keep the three base frames permanently installed on the ACME table, then our crane can lift off the upper “U”-shaped frame while requiring us to only remove the minimum number of μ -metal pieces from the shields. Our crane then allows us to move the upper frame to the space in our lab downstream of the beam line, and mount the upper frame on specially designed wheel blocks to be moved to storage. In the minimal version of this process it would only be necessary to disconnect the the four bottom edge covers for each layer, and remove panels from the bottom of the upstream and downstream surfaces, but the limits of our vertical clearance from our crane make it necessary to remove all panels on the upstream face for each layer. Additionally, we can nest the upper layers while they are on wheels, but this requires a tricky shuffling process as the outer layers

must be removed first. We can attach our crane using an aluminum frame that connects directly to the support frame for each layer, with no possible contact with the μ -metal.

We mount the μ -metal plates to the frame by sandwiching the sheet metal between the flat surfaces of the aluminum frame and a washer on the outside. In almost all locations we sandwich two overlapping pieces of μ -metal in this way to create smooth, well-bonded mating joints. These joints are held together by brass or titanium bolts¹¹, which attach to long brass or titanium T-nuts in the extruded frame or titanium threaded inserts in the solid bars of the base frame. We tighten all bolts to a torque of 34 in. · lbs¹² to prevent over-tightening that may cause unwanted stress. There are small gaps in some locations due to the slight warping of some of the μ -metal after the annealing process.

The design process we used to determine the dimensions of our shields, as well as the allowable sizes of holes in the shields was a collaboration between Daniel Lascar and myself performing the mechanical design work, and Siyuan Liu performing simulation work. These simulations were performed in COMSOL, and were magnetostatic simulations with permeability data provided by the manufacturer of the metal. These results found that with our hole configurations we had a shielding factor of $\sim 75000 - 125000$ depending on the orientation of the external field. For a 1 G external field, this resulted in internal fields of $\approx 8 - 14 \mu\text{G}$, with the largest internal field for an external field along \hat{x} as shown in figure 5.7. This met our goals as the actual field in our lab is both well below 1 G and the long axis of our shields is oriented along the East-West axis, so the Earth's field primarily aligns with directions for which we have better shielding.

The shield system was constructed over the course of roughly one year, which required a collective effort from all members of the ACME collaboration. The shields were annealed in the late fall

¹¹Brass is used for the middle and outer layers, while titanium is used for the inner layer as it is less magnetic.

¹²This value was chosen somewhat arbitrarily as “just past finger-tight”, but has so far worked well for us.

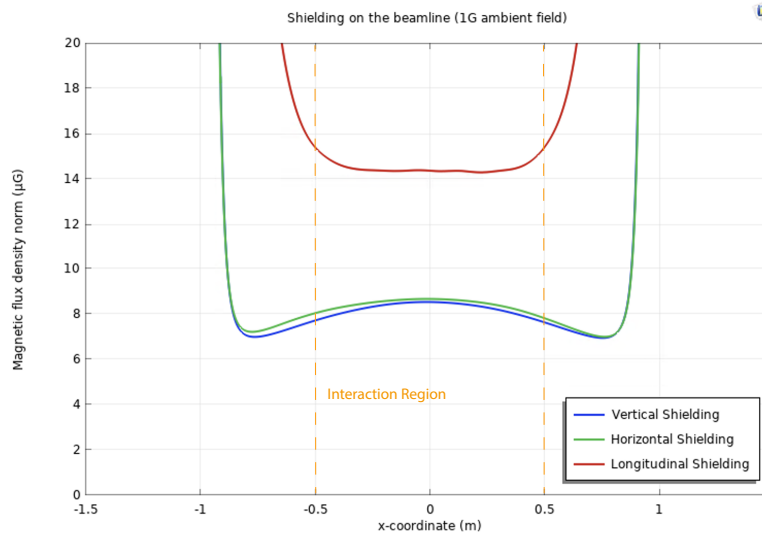


Figure 5.7: Magnetic Shield Design Simulation Results. Simulated magnetic fields inside our three layer shields, assuming a 1 G external field. The three curves correspond to the orientation of the external field used in the simulation. The results are shown along the line defined by $y = z = 0$ at the center of the molecular beam. The orange dashed lines indicate the interaction region as delineated by the readout and refinement lasers, which are separated by 100 cm along \hat{x} . Figure created by Siyuan Liu.

of 2021, with the shields fully assembled for the first time in the summer of 2022. Since that time we have been performing work to characterize our shields and our ability to degauss them. At the time of this writing, we have a preliminary degaussing method, with which we have been able to achieve fields along the beamline of $\approx 5 - 10 \mu\text{G}$ with the interaction region chamber installed, meeting our design goal of $< 10 \mu\text{G}$. Additionally, all of the magnetic field gradients we have measured in this low field configuration have been well below our requirement of $< 1 \mu\text{G}/\text{cm}$. However, this work is still ongoing and we hope to develop methods to consistently achieve lower fields inside the shielded region.

THE ACME TABLE

The three layers of magnetic shields, the magnetic field coils, and the interaction region chamber are all mounted separately on a support structure that we call the ACME table. The table is constructed out of four layers of stacked 3 in. \times 3 in. square 8020 aluminum extrusions, which are arranged so that the weight of the system (\sim 2400 kg for the fully assembled experiment) is transferred directly through the extruded profiles and not through any fasteners. We stack two layers of 0.5 in. aluminum plates atop the 8020 frame that have tapped holes with threaded inserts for mounting the support posts for the experiment components. These plates are split in half along \hat{z} and made from two layers so that the parts can be reasonably handled as these split plates still weigh \approx 60 kg. One final layer of 0.5 in. aluminum plate is placed on top of these plates with cutouts for each of the posts to serve as an alignment guide for the posts to compensate for the split in the attachment plates which could cause a misalignment. Due to the large mass of the components that would be mounted on the table, we performed static simulations to determine the rigidity of the table, and found that it should deflect by less than 0.25 mm, which we considered sufficient for the alignment of our apparatus, and the stresses involved were sufficiently low that we need not be concerned.

The table itself has the same footprint as the outer magnetic shield and stands 13.5 in. tall. With this table height, we decided to increase the height of the ACME beamline by 10 cm to 1.2 m from the floor. The four feet of the table were placed so that we have complete access to the underside of the table with a 9 in. clearance along both \hat{x} and \hat{z} for installing optics below the table for either STIRAP injection or dumping the STIRAP beams. For this we have five 4 in.¹³ diameter holes in the table aligned with each of the five windows on the underside of the chamber. We additionally ensured that the access under the table was large enough to insert pallet jacks to move the apparatus if necessary. A diagram of the table can be seen in figure 5.8. We mounted another 8020

¹³These are intentionally oversized relative to the windows and the openings in the shields.




Figure 5.8: The ACME Table. Figure showing the ACME table and the support structures for the magnetic shields, magnetic field coil, and interaction region chamber.

frame around the table, which we refer to as the shield guards. The primary purpose of this frame is to support removable, clear purple acrylic sheets that protect the μ -metal from any impacts. This frame is also used as the mounting structure of the SiPMs so that they can be mounted just outside of the shields.

ACCEPTABLE MATERIALS

In the design of the shield frames, as well as the interaction region chamber and the magnetic field coils, we opted to use exclusively what we believed to be non-magnetic materials. Our initial plans were to use primarily brass fasteners, with some titanium fasteners, alongside aluminum for the chamber and the structural components of the magnetic field coils and shields. However, in later testing with our QuSpin optically pumped magnetometers inside our small commercial three-layer test shields, we found that many of the components we had expected to be non-magnetic were in fact magnetized to a surprising degree. Most notably, we found that after exposure to a strong per-

manent magnet, brass fastener components often produced fields larger than $> 1 \mu\text{G}$ when placed ≈ 5.5 cm away from a sensor, with some brass components from our system producing fields as large as $> 100 \mu\text{G}$. Additionally, we found that even some aluminum pieces could produce significant fields at this distance. While almost all components will be significantly farther from the molecular beam than they were from the sensor in these tests, the large number of fasteners required in our experiment could result in a significant effect, particularly on the interaction chamber which is much closer to the beam than the shields. As such, we removed almost all brass¹⁴ from everything inside the inner shield, including the shield frame, replacing the brass with either titanium or aluminum, which all needed to be tested for magnetism.

5.2.4 DEGAUSSING

In the course of every block measurement of our experiment, we must reverse the direction of the applied magnetic field. As this applied magnetic field magnetizes the shields, we must demagnetize them each time before continuing with the measurement through a degaussing process. In this process we apply a sinusoidal driving field $\vec{\mathcal{H}}(t)$ using permanent coils to orient the magnetic domains of the shields as the magnetization of the shields follows the $\mathcal{B} - \mathcal{H}$ hysteresis curve. This field is slowly ramped down from an amplitude that saturates the material to $\mathcal{H} = 0$, forcing the magnetization of the shields to go towards zero [145]. As we first found during ACME II, this process depends on the frequency of the alternating drive, which must be low enough that the skin depth allows the driving field to fully penetrate the 1.6 mm thick metal, the rate at which the envelope function goes to zero [145], the amplitude of the driving field, which should initially saturate the material, and how closely the final applied field approaches $\mathcal{H} = 0$. We are still working to develop an optimized degaussing procedure with our apparatus, particularly as we face the added constraint

¹⁴The power lock fasteners used in the shield frame still contain brass.

that the time required to degauss after each field reversal is already the largest limit to our measurement duty cycle, and preliminary measurements indicate that an optimal degaussing frequency is likely to be 1 Hz or lower. We have also found during these tests that uncontrolled transient voltages on the coils will remagnetize the shields.

The degaussing field in our apparatus is supplied by a series of 108 10-turn coils constructed from ribbon cables. Each coil runs along the outer surface of the shields until reaching either a corner or an edge of the shield, then winds back along the inside surface of the μ -metal until encountering the next edge or corner, forming a closed loop around a single edge or face of the shields. Based in part on existing designs for degaussing magnetically shielded rooms [144, 146, 145] as well as our experience with the ACME II degaussing system, one coil runs along each of the 12 edges of the rectangular prism, and on each face of the prism two coils run in each of the two tangent directions. With this configuration we have 12 coils for each layer than apply a field along each of the three coordinate axes. This is intended to form a closed loop of flux in the shields that travels in each of the coordinate planes, and the number of coils is intended to keep the flux density homogeneous throughout the loop. This gives us 36 coils per shield layer, for a total of 108 coils.

THE SELF-SHIELDED COSINE THETA COIL

In the hopes of reducing the magnetization of the inner shield from our applied magnetic field B_z , the new cosine theta coil has been designed to be self-shielding. This self-shielding technology has long been used to reduce the external field produced by superconducting solenoids [148], and is achieved through the use of a second set of current paths on the outer surface of the coil in which the current flow is anti-aligned with the rest of the coil. This produces only a small field shift inside the intended field region, while allowing us to minimize the field at the inner surface of the magnetic shields. This coil was designed based on the work of multiple members of our research group and is currently under construction. Based on simulations performed by Daniel Ang we expect that

when applying $|\mathcal{B}_z| = 200 \mu\text{G}$, the field produced by the coil at all points on the inner surface of the inner shield should be $\leq 17 \mu\text{G}$ [123]. We anticipate that this should significantly reduce the magnetization of the inner shield during normal operation, and reduce the requirements of our degaussing system.

5.2.5 MAGNETOMETRY

With the goals of characterizing the magnetic shields and later monitoring the magnetic fields in our experiment we have developed three new magnetometry systems for ACME III. One of our primary goals with the development of these systems was to have more continuous monitoring and access to more monitoring locations than in ACME II so that we could more readily diagnose problems like the long term increase in the ambient field we saw.

OPTICALLY PUMPED MAGNETOMETERS

The primary sensor that we intend to use for ACME III is the the QuSpin QZFM Gen-3 optically pumped magnetometer (OPM) [149]. This three axis sensor is based on an optically pumped rubidium vapor cell, where the optical transmission is dependent on the magnetic field. This can be used with a modulated applied magnetic field oscillating at 923 Hz and a lock-in amplifier to produce an antisymmetric signal that scales with the magnetic field [150, 151]. These sensors were selected for their extremely high sensitivity of $< 0.2 \text{ nG}/\sqrt{\text{Hz}}$ in the band from 3 – 100 Hz within their linear operating range of $\pm 50 \mu\text{G}$, which can be extended to $\pm 500 \mu\text{G}$ by applying dc or quasi-dc offsets to the internal magnetic field coils, as the lock-in signal is only useful near a zero field configuration [149].

However, like the flux-gate sensors we used in ACME II, and the magnetoresistive sensors we also intend to use for ACME III, these sensors have unknown internal offsets. These offsets, which

we have measured to be on the order of $\sim 30 - 70 \mu\text{G}$ in our sensors, are believed to be caused by a combination of errors in the magnetic field caused by the lasers, light shifts, and slight magnetization of components in the head. In order to measure these offsets and subtract them from our measurements, we have built mechanical systems to precisely rotate the sensors both inside our three-layer test shields and the ACME III shields. As we described in section 3.3.1, by rotating the sensor by 180° (or 90°) it is possible to separate out the offsets and the ambient field. One important complication related to these sensors is that the rubidium cell that is the actual sensor is offset from the center of the sensor housing on all three axes [152], and so we must be careful to rotate around the actual sensor center and not the center of the housing. These measurements can be further improved by coupling these rotations with a reversal of an applied field, to also separate out the applied field and systematic effects related to our rotations, and this will likely occur once the magnetic field coils are installed in our apparatus. At the time of this writing, these offsets are the current limit on our ability to perform an absolute measurement of the ambient field inside our shields as we have only been able to measure them with an uncertainty of $\sim 1 - 2 \mu\text{G}$, and work is ongoing to improve this.

Based on the sensitivity and low field characteristics of the sensors they are well suited to our purposes for measuring the field inside our shielded region, however they cannot be used well outside our inner shield layer. These sensors have served as our primary tool for characterizing our magnetic shields, and we plan to use them in the future for both probulation, and for inserting into our magnetometer pockets. We had initially hoped to place a static array of these sensors on the surface of the interaction region vacuum chamber, but we determined that the internal offsets of these sensors were not likely to be stable enough for our purposes without the ability to rotate the magnetometers and measure the offsets.

MAGNETORESISTIVE SENSORS

In addition to the OPMs, we have also developed a system of lower sensitivity magnetoresistive sensors for characterizing and monitoring the ambient field in our lab and the field between layers of our magnetic shields. For this system we use the three-axis Twinleaf VMR Magnetoresistive Vector Magnetometer [153] which have a sensitivity of $3 \mu\text{G}$, but can operate up to $\pm 1 \text{ G}$, making it well suited to measuring the ambient lab field and the field between the magnetic shield layers, which is significantly higher than at the center of the shields. Allan deviation measurements we performed with these sensors indicate that inside the inner shields we can resolve fields $< 1 \mu\text{G}$ with 10 s of averaging, however this is only useful as an independent check of the more sensitive OPMs. Additionally, based on our measurements, these sensors have similar scale offsets to the OPMs.

We have created an array of 12 of these sensors inside our magnetic shields, with one sensor placed just inside roughly the center of each face of the middle and outer layer shields (6 sensors per shield layer), with an additional array of 6 ambient field sensors placed outside of our shields. This array is intended to allow us to monitor the field near our shields to detect and diagnose drifts of the field caused by magnetization of individual shield layers. We had originally intended to use an identical array of 6 sensors inside the inner layer, but found that these sensors were themselves not sufficiently non-magnetic for this purpose. Additional spare sensors can be used as a check for the QuSpin sensors and for inserting between the magnetic shield layers through holes designated for this purpose.

Q STATE CO-MAGNETOMETER

A third magnetometry method that relies on the large magnetic moment of the ThO Q state [87] has been developed by Xing Wu based on tests using the ACME II apparatus. This method, which could be used as an alternative to probed, can use the molecular beam itself to measure the

magnetic field through Zeeman precession of the $|Q, J = 2, M = \pm 1\rangle$ states, which can be read out using the $Q - I$ transition analogous to our normal phase readout. The tests of this method on the ACME II apparatus indicate that this method has a sensitivity $< 1 \mu\text{G}$, while requiring minimal modification of the ACME apparatus. At the time of this writing this method has not been tested on the ACME III beamline, as we have not yet commissioned the ACME III ThO beam in the interaction region, however we expect this method to be a powerful tool to measure the magnetic field along the actual beam axis during the ACME III data collection.

5.3 THE ACME III INTERACTION REGION VACUUM CHAMBER


In order to take advantage of what we learned from the H state lifetime measurement we needed to redesign the interaction region vacuum chamber in addition to the magnetic shields. This chamber needed to house the new field plates, which were not only roughly 1 m longer, but had a larger separation of 6 cm, and had been made taller to maintain the electric field homogeneity with the larger separation. We set out to develop a system that would meet the needs we knew we would have at the time, while also building in the flexibility to use the chamber in unforeseen ways. The development of this chamber and the associated systems has been led by myself with significant contributions from Daniel Ang, Ayami Hiramoto, Peiran Hu, and Takahiko Masuda.

5.3.1 DESIGN CONSIDERATIONS


We set out to design our interaction region with the design of the ACME I and II interaction region as a known starting point, which was a rectangular chamber constructed from a welded aluminum frame with large rectangular mating flanges on each of the six faces of the chamber. All of the smaller ports for windows, light pipes, and the beam nipples mate to these mating flanges. Additionally, the chamber has 0.63 in. thick lips that run along \hat{x} on the top and bottom of the chamber

to increase the rigidity of the chamber. With this in mind we designed the ACME III chamber with the following considerations, which are ordered roughly by priority. A diagram of the ACME III interaction region chamber and the associated systems can be seen in figure 5.9.


- The first order design constraint on the chamber was the dimensions of our new field plates, which were being developed in parallel by Peiran Hu. Based on simulations of the electric field homogeneity performed by Daniel Ang, we determined that the the field plates needed to be at least ≈ 130 cm long and ≈ 35 cm tall. We wanted to minimize the size of the chamber primarily because the size of the chamber determined the smallest possible dimensions for both the magnetic field coils and for the magnetic shields. From this consideration, as well as accommodations required for the paths of our straight light pipes, we set the dimensions of the chamber. The dimensions of our rectangular interaction region chamber are $157 \text{ cm} \times 58 \text{ cm} \times 59 \text{ cm}$.
- We chose to have the chamber constructed from welded aluminum, as we did for the previous chamber, due to the non-magnetic requirements of our apparatus. We also considered titanium as an alternative, but aluminum was more cost effective, and we have a good track record with the manufacturer of our chamber Atlas Technologies, who also manufactured the original ACME chamber. Notably, out of concern for the strength and longevity of holes tapped in the aluminum, every tapped hole associated with a sealing surface on the chamber includes titanium threaded inserts. However, these threaded inserts in the chamber were installed by the manufacturer and we have not found a sensitive and non-destructive method to test the magnetism of these inserts directly. However probulation performed with the chamber installed suggest that these are not a problem.
- We wanted to ensure that we could achieve similar vacuum levels as we did in ACME II, which typically reached pressures of $\sim 10^{-7}$ torr with the buffer gas flow off, and whose



(a) The interaction region vacuum chamber and associated systems.



(b) Side view of the chamber from the North.



(c) Top down view of the chamber

Figure 5.9: The Interaction Region Vacuum Chamber. Diagrams showing the new interaction region chamber. Components are shown in false color for easier identification.

pressure was primarily limited by water in the chamber [88]. Efforts to ensure that we can reach these levels are discussed in section 5.3.2

- We put in optical access windows for the known lasers that we will use for ACME III. This includes two pairs of windows on opposite sides of the chamber for the refinement and readout lasers, which are spaced by 100 cm along \hat{x} , and a pair of windows on the top and bottom of the chamber for the vertical $X-C-H$ lasers which are centered 1 cm upstream from the refinement laser windows. As we described in section 4.2.2, we originally intended to use a rotary window design for these windows, so we included large apertures in the chamber for the readout and refinement windows, with slightly smaller apertures on the top of the chamber for the STIRAP entrance window, as well as the window directly above the readout region. This additional window above the readout region provides us vertical access for lasers in the readout region, which we have considered for systems like optical cycling detection, but we do not have a planned use for this window. As we are no longer planning to use the rotary window design, the large apertures for these windows are covered by window adapter flanges as can be seen in figure 5.9.
- In addition to these windows we also included 14 auxiliary windows, which were included for flexibility, but do not have planned uses. All of these windows have the same 3 in. diameter aperture as the STIRAP windows, as a trade off between a large aperture and minimizing the size and number of holes in our magnetic shields. There are three of these windows between the refinement and readout regions on each side of the chamber, located at the center of the precession region ($x = 0$) and at $x = \pm 20$ cm. These windows were placed with the intention that they may be used as alternate refinement locations that could be used to either test for systematics related to the precession distance, or possibly for future measurements of the lifetime of various states of ThO. Additionally, we placed five evenly spaced windows

on the bottom surface of the chamber (at locations $x = 0, \pm 25, -51, 50 \text{ cm}^{15}$), and three corresponding auxiliary windows on the top surface between the readout and refinement regions. These windows can be seen in figure 5.9 in yellow.

- With our new design of the magnetic shields we needed a new way to extract the detected light so that it could be collected by our SiPMs. In ACME II this was done with eight static fused silica light pipes that were bent so that they exited the vacuum chamber at the upstream and downstream surfaces. This allowed them to remain in place when the shields were removed as they only passed through the shield end caps which were fixed in place. In contrast for our new shield design the only parts that are intended to be semi-permanent are the bottom surfaces of the shields. Rather than place the SiPMs under the shields as there is limited access below the shields, we decided to create a system of straight light pipes that exit the shields at a 45° angle, which can be disconnected from the chamber and removed before taking apart the magnetic shields. The details of this system are discussed in section 5.3.4. Notably, as was the case for our magnetic shields, the straight light pipes required that we introduce an asymmetry of $\approx 5 \text{ cm}$ to the chamber along \hat{x} to prevent conflicts between the light pipes and the welded frame of the chamber. This allowed us to keep the distance between the electrostatic lens and the refinement laser as small as possible to maximize the effect of the lens.
- We expanded the magnetometer pocket system that was introduced in ACME II to increase the number of locations where we can insert a magnetometer to probe the magnetic field near the molecular beam. We placed eight pockets on the top surface of the chamber, which allow a sensor to be inserted along \hat{y} , and eight pockets on each side of the chamber allowing sensors to be inserted along $\pm \hat{z}$. These magnetometer pockets follow roughly the

¹⁵Two of these windows correspond to the STIRAP lasers and the readout region.

same design as the ACME II pockets, with a larger inner diameter of $1 - 1/8$ in. to accommodate the larger diameter needed to rotate the OPMs about their sensor centers. On the North and West sides of the chamber the pockets are located in two rows of four pockets at $y = \pm 10$ cm and located along \hat{x} at $x = \pm 10, \pm 27.5$ cm. These two rows offset from the beamline, rather than a single row at $y = 0$ were chosen to allow for the large auxiliary windows along the beam axis. Similarly, for the magnetometer pockets on the top of the chamber we have two rows of four pockets located at $z = \pm 10$ cm (in the closer configuration) and located along \hat{x} at $x = \pm 5, \pm 34$ cm. These positions were chosen to allow measurements distributed across the precession region while placing sensors as close to the refinement and readout regions as we could. Additionally, the distribution of the pockets in two rows on each face allow us to better measure gradients across the beamline. These pockets are shown in figure 5.9 in shades of pink.

We designed these pockets so that the sensors could get as close as possible to the molecular beam without interfering with the field plates. As the field plate separation for ACME III had not been finalized when we ordered the chamber, the magnetometer pockets on the top of the chamber mate to a slot in the chamber flange not a circle, and the welded tube is offset from the center of the oval-shaped flange for the pocket, which allows us to change the \hat{z} position of these pockets by rotating the pocket 180° without additional machining.

- We chose the thickness of the aluminum plates used for the mating flanges as well as the welded frame to be 1 in. to minimize the deflection of the plates when under vacuum. We can calculate the maximum deflection of a flat plate under uniform stress that has all of its

edges fixed using the formula¹⁶

$$d_{\max} = \frac{\alpha qb^4}{Et^3} \quad (5.8)$$

where E is the Young's modulus, q is the pressure loading the plate, t is the thickness of the plate, b is the shorter dimension of the plate, and $\alpha = 0.0284$ is a dimensionless constant based on the aspect ratio of the plate [154]. With the largest opening in the welded frame being 17 in., the maximum deflection we should expect under vacuum is only $87 \mu\text{cm}$ (.003 in.). This is consistent with tests we have performed with the actual chamber while under vacuum.

- During ACME II we observed that vibrations of the field plates created time varying interference patterns in our lasers. Based on these observations we designed the mounting system of the chamber to suppress vibrations to decrease this effect. Details of these efforts are described in section 5.3.3.
- In order to mount the aluminum plate that holds the field plates and our collection optics, we included rows of tapped holes with threaded inserts along the inside surface of the bottom plate of the welded frame. Two runners attach to this part of the welded frame, and the plate with the field plates can be lowered into place with our crane once the top mating flange of the chamber is removed.
- We included 42 standard KF50 ports distributed across all of the chamber walls, which at the time of this writing are only intended to be used for electrical feedthroughs for the electric field plates. As these holes are not intended for optical access, there is no corresponding set of holes on the shields.

¹⁶This assumes the Poisson's ratio of the material is $\nu = 0.3$, which is close enough to the value of aluminum for our purposes.

- Due to the weight of the chamber, $\gtrsim 400$ kg, we included tapped holes for mounting shackles so that the chamber could be installed with our single axis crane.
- We included tapped holes with threaded inserts so that we could attach aluminum handles to each of the large rectangular side flanges of the chamber to ease their handling.
- In order to maximize our flexibility, on each of the six large rectangular mating flanges we included a breadboard pattern of tapped $1/4''-20$ blind screw holes spaced apart by 1 in. along both \hat{x} and \hat{y} . These patterns exist on both sides of each flange, with the tapped holes offset by 0.5 in. in each direction between the two sides to prevent conflicts. We included these so that we could mount sensors to the chamber interior or exterior, attach collimators to the interior of the chamber, and ideally handle any unforeseen future needs. Due to the large number of these holes, we did not include threaded inserts here as they would be both cost and labor prohibitive, and if necessary an adjacent hole could potentially be used instead.

The interaction region chamber was delivered in the fall of 2021, and since then has undergone testing to confirm the vacuum levels, check the compatibility of the chamber with the shields, and check the light blocking systems we have developed for the SiPMs.

5.3.2 VACUUM LEVELS

The pumping speed of the ACME II chamber was limited by the diameter of the ISO100 nipples used to connect the chamber to the vacuum pumps in the stem and dump region of the beamline. This was determined by a calculation of the conductance of each element in the beamline between the ACME II chamber and the turbo pumps based on the tables in [155], which found that the conductance of the ISO100 nipples were a factor of $2\times$ smaller than any other component in the system besides the ACME II rotational cooling chamber, which will not be present for ACME III.

However, the diameter of the ISO100 nipples are already among the largest holes placed in the magnetic shields, and we in fact opted to decrease the diameter of the nipples by 2 cm to allow more clearance with the shields. We considered additional pumping, but any such ports would either require placing magnetic materials inside the shields, or would be limited to significantly lower pumping speeds than the existing nipples. We were able to decrease the length of the upstream nipple by 1.5 in., but due to the asymmetry of the chamber and shields we increased the downstream nipple by a similar amount. As such we expect that we are still limited by the pumping speed through these nipples, as we have increased the conductance of the elements on both sides of the beamline. Additionally, to minimize outgassing, we use only Viton for our o-ring seals, and ensure that only vacuum rated materials are allowed inside the chamber, which mostly includes aluminum, glass, and PEEK.

After the delivery of the chamber we opted to test the vacuum of the chamber while blanked off. We waterjet cut aluminum blanks for each of the vacuum seals besides the commercial KF50 and ISO100 blanks. In order to ensure a good vacuum seal with these blanks we wet-sanded the sealing surface of each custom blank with four progressive grits of sandpaper up to 800 grit as this proved faster and more reliable than attempting to machine the surfaces with a fly cutter due to the size of the blanks. In this blanked off configuration we demonstrated that the chamber was helium leak tight to the 5×10^{-9} mbar L/s level. Additionally we baked the chamber in this configuration for more than one week at a temperature of $\sim 90 - 110^\circ$, until the pressure stalled at the 1×10^{-5} torr level. After this baking, the chamber was able to reach pressures of $< 4 \times 10^{-7}$ torr. Based on this, further baking will be necessary after the installation of the final components such as the magnetometer pockets, the field plates, and the windows, as these will introduce a large surface area to the chamber, which has not been baked. Additional tests with a residual gas analyzer will allow us to determine the likely sources of our background pressure.

5.3.3 MECHANICAL SUPPORTS AND VIBRATION ANALYSIS

As previously described, in ACME II we observed vibrations of the electric field plates, which coupled with reflections off the field plates produced time varying interference patterns in our readout and refinement lasers. We determined that the best way to support the chamber, given our magnetic shield design was using the same style of support posts that the shields used, which consists of six aluminum posts attached to the top of the ACME table with one post at each of the corners of the chamber and one at the center of each of the long edges of the chamber. However, as the support posts for the shield are necessarily much longer than for the shields, it does not take much of a stretch to see this system as a mass at the end of a spring. we investigated how different configurations for these posts impacted the vibrational spectrum of the chamber. In particular we considered configurations where the posts had either a diameter of 1 in. or 2 in., and configurations where we mechanically linked the tops of the posts supporting the chamber to posts of the same diameter that supported the magnetic field coils.

We studied these configurations by performing dynamic harmonic simulations in SolidWorks in conjunction with measurements of the actual vibration spectrum of the floor in our lab at Northwestern. For these simulations we used a simplified test mass in place of the chamber, as we expect the dominant effects to be from the posts and not flexing of the chamber. We determined that the 20 lowest order harmonic modes of the chamber (along both \hat{x} and \hat{z}) lie below 250 Hz, so we constrained our simulations to the range from 0.1 – 250 Hz. Initial simulations using a white noise drive spectrum showed that the 2 in. posts generally shifted the resonant frequencies of the spectrum higher and coupled higher frequency components of the table to the chamber, while the linking between the chamber and the magnetic field coils decreased the amplitude of the resonances. We also simulated the ACME II chamber structure as our best means to compare with what was observed during ACME II, however it should be noted that this is not an exact representation when




Figure 5.10: Interaction Region Chamber Vibration Simulations. Simulated vibration spectra for different configurations of the interaction region supports compared to the ACME II apparatus. The orange dashed curve is the scalar vibration spectrum in our lab at Northwestern. The blue curves show the simulated spectra of the ACME II apparatus driven by the measured spectrum. The configurations with coupling to the magnetic field coil are not shown.

using the measured drive spectrum from the Northwestern lab.

When using the vibration spectrum from our Northwestern lab as the driving spectrum in our simulations of ACME III configurations and the ACME II configuration, we found that the 2 in. posts provided the simplest method of keeping the displacement spectrum comparable to ACME II. The results of these simulations can be seen in figure 5.10, which shows the vibrational response spectra for the different diameter configurations along both \hat{x} and \hat{z} alongside the response from the ACME II simulations. Notably, while the simulations with the mechanical couplings suggested better responses with a white noise drive, with the measured noise spectrum these benefits largely disappeared. Based on these results the ACME III chamber is supported by 2 in. posts with no cou-

pling to the magnetic field coils.

5.3.4 SiPM MOUNTING AND STRAY LIGHT BLOCKING

One of the most critical systems connected to the interaction region chamber is the light pipe system which couples the detection optics to the SiPMs which must be mounted outside of the magnetic shields. As we previously described, we opted to develop a system where the light pipes were split into two components, one partially inside the vacuum chamber labeled the internal light pipe, and a longer light pipe which is mated to this light pipe with a coaxial end joint and connects to the SiPM outside the shields, labeled the external light pipe. This split system was also used in ACME II, but for our system we have developed methods to allow the external light pipe to be easily removed from the system whenever the shields need to be disassembled. This system can be used with no direct access to the joint itself, only the exposed end of the external light pipe. The system also uses only tested non-magnetic components and only minimally impacts the transmission efficiency of the light pipes. Additionally, one of the most crucial aspects of this system is that it must completely block out stray light which can increase the SiPM background, particularly as in our lab at Northwestern we cannot simply turn off all of the room lighting as we are operating in a space shared with other experiments. The development of this system, which interfaced between multiple projects on the experiment, was a collaboration between myself, Daniel Ang, Takahiko Masuda, and Ayami Hiramoto.

We now provide a description of this system and its operation, working from the collection optics outwards to the SiPM. This system is represented schematically in figure 5.11. The internal light pipe, which is a 7.5 in. fused silica rod with 20 mm diameter, is designed to terminate just inside the collection optics mounts, where it is loosely held radially by an o-ring in a PEEK clamp attached to the optics mounts. The internal light pipe is primarily held in place by the vacuum sealing o-ring,

which runs around the circumference of the light pipe, and is compressed by an angled clamp¹⁷, providing enough friction to fully constrain the light pipe. The end of this light pipe sticks out from the vacuum seal, which is angled at 45° by only ≈ 1 cm, and this portion is surrounded by a PEEK collar, whose end is flush with the end of the light pipe, which is connected to the light pipe via an o-ring groove inside the bore of the collar. This collar serves as added protection from the light pipe being pushed through the vacuum seal¹⁸, and also protects the end of the internal light pipe from collisions with the end of the external light pipe.

In order to mate the two light pipes, we use a coaxial end joint, where a similar PEEK collar is attached to the end of the external light pipe, this time using two o-rings rather than one¹⁹, which is also flush with the end of the external light pipe. In order to align the two light pipes, we mount a component we refer to as the external light pipe holder on the light pipe vacuum flange, with a Viton sheet gasket between the two to block any light from entering the system at this joint. The external light pipe holder is essentially an angled block of aluminum with a bore that is intended to be an engineering fit with the external light pipe collar, and constrains the external light pipe so that it is coaxial with the internal light pipe. The external light pipe slides into this bore until the collars on the two light pipes are in contact with one another²⁰.

To block any light entering the light pipe between the external holder and a clamp outside the shields we attach an aluminum tube to the external light pipe holder. This 1.25 in. outer diameter

¹⁷This design was based off of the mechanism of quick connect vacuum couplings, which were not available for our system in non-magnetic materials

¹⁸While these seals have been tested to $\sim 4\times$ atmospheric pressure, we intend to increase this safety factor by also epoxying the light pipe to the collar, however this requires further testing to confirm that this does not significantly impact the light pipe transmission efficiency.

¹⁹The choices of two o-rings and the longer collar were partly based on tests on how well the light pipe was constrained from angular deflection.

²⁰In practice we intentionally introduce a slight air gap between the two light pipes to prevent damage. We also considered soft silicone disks and index matching gel, but determined that these were not feasible without access to the joint.




Figure 5.11: The Light Pipe Mounting System. Schematic diagram showing the joint between the internal and external light pipes, the vacuum seal around the light pipes, and the light pipe alignment structures. The other end of the light pipe is held with the M2 clamp and couples to the SiPM as shown in figure 4.4.

tube, has a 1.5 in. outer diameter black Delrin sleeve epoxied to the end which is inserted into the shields, which has 1.5”–12 external threads²¹ on the end. This threads into a pair of mating threads in the external light pipe holder, forming a light-tight seal. This tube can be easily attached or detached with the shields in place, and also serves as a guide for the light pipe when it is inserted so that it does not contact the shields.

Once the external light pipe is inserted into the external holder and mated to the internal light pipe, the rotation and axial motion of the light pipe is fixed by a clamp we refer to as the M2 mount. This clamp, which is 3D printed, consists of two halves which clamp a pair of o-rings around the light pipe. This clamp also mates with the end of the aluminum tube, which end inside a recessed groove in the clamp, which can easily be wrapped in tape if necessary. The M2 mount is mounted to the same optics breadboard as the SiPM module using optics posts. The SiPM side of the mount has a lens tube which protects the sides of the light pipe as it couples into the SiPM module. The SiPM module is then mounted at the end of the light pipe on a breadboard which attaches to the shield guards frame. Where necessary additional light blocking can easily be applied to these accessible components using black tape and rubber tubes around the lens tubes.

This system has been mechanically tested, both on a test apparatus and on the ACME III apparatus with the shields in place, and has proven to be robust and repeatable. Additionally, we have tested the light blocking abilities of this system by measuring the output of the SiPMs and comparing it to the measured dark current of the same SiPM module. From these tests, with only a small amount of tape outside the shields, there is no detectable light reaching the SiPMs. It is worth noting that in these tests we found that the SiPMs were able to detect photons emitted by an ion gauge attached directly to the interaction region chamber. As a result care should be taken to ensure that

²¹We initially used a fine custom thread on the aluminum tubes, but found that these threads were difficult to have made with sufficient tolerances that they would not bind. The sleeve is softer than the aluminum threads it mates to, and can both be easily replaced or re-threaded with a die, while the mating threads can be re-threaded with a tap.

no gauges on the apparatus have a direct line of sight to the collection optics, which must be the case in the final configuration of the lens and interaction region chamber. Further tests and development are ongoing at this time to extend these types of light blocking methods to the windows on the interaction region chamber, which must be coupled directly to enclosures for the lasers on the optics tables next to the apparatus.

*In our hands is placed a power greater than their boarded
gold
Greater than the might of atoms, magnified a thousand
fold,
We can bring to birth a new world from the ashes of the
old,
For the union makes us strong.*

Pete Seeger, *Solidarity Forever*

6

Conclusion

THIS THESIS HAS DESCRIBED THE ACME II MEASUREMENT AND RESULT and the progress made towards the upcoming ACME III measurement. In chapter 1 we described how precision measurements of the electron EDM, like the ACME experiment, are incredibly powerful tools for probing new T violating physics beyond the Standard Model. Both the ACME I [6, 7] and ACME II [8] measurements set the most stringent limits on the electron EDM, improving upon the sensitivity of

the previous best measurement by an order of magnitude every ~ 4 years. With the ACME III experiment, which we describe in chapters 4 and 5, we aim to once again improve upon the sensitivity of the ACME II measurement by another order of magnitude.

The ACME II measurement, which placed a limit on the electron EDM of $|d_e| < 1.1 \times 10^{-29}$ e·cm remains the most stringent published limit on the electron EDM. This limit probed new physics at energy scales up to tens of TeV, well beyond what can be achieved at the Large Hadron Collider. This measurement, using techniques described in chapters 2 and 3, laid the foundation for the current ACME III measurement, both in developing our understanding of the systematic effects that we expect to see, and in influencing the upgrade pathways for ACME III.

A crucial preparation for ACME III was a precise measurement of the H state radiative lifetime in ThO [70]. The longer-than-realized lifetime of the state used for our measurement resulted in an increase in our experiment precession time. In order to implement this longer precession time in our experiment we designed a longer interaction region chamber, and new magnetic shields that will also allow us to suppress excess noise caused by velocity fluctuations and suppress systematic errors associated with magnetic field gradients. As described in chapter 5, these three projects and their integration with other systems, have made up a significant portion of my contributions to the development of ACME III over the last few years.

In chapter 4 we described the development of the upgraded ACME III experiment. Before developing the new apparatus we tracked down the source of the excess noise that prevented ACME II from reaching the shot-noise limit, and found methods to suppress this noise below our projected sensitivity [69]. With our new apparatus we have implemented an electrostatic lens to focus our molecular beam [85], an improved rotational cooling system [85], and a load-lock system to increase our ablation target changing rate in order to increase the usable flux of molecules. In addition to this we have implemented upgraded collection optics and new SiPM based detectors to increase our detection efficiency [71, 72, 73]. Alongside these upgrades, we have also developed new techniques to

suppress the dominant sources of systematic uncertainty in ACME II, which were associated with polarization gradients in our lasers and magnetic field gradients, to below the expected sensitivity of ACME III.

At this time the systems of the experiment have mostly been demonstrated individually and are now being brought together. We expect to soon commission the full ACME III ThO beam in the interaction region and begin collecting data for ACME III. Based on this we estimate that within approximately the next two years the ACME III experiment will be able to perform a new measurement of the electron EDM with more than an order of magnitude improved sensitivity compared to ACME II. When combined, all of the ACME III improvements are projected to give us a factor of $\sim 43\times$ greater sensitivity than ACME II, probing new physics up to the ~ 100 TeV level [37].



Laser Power Estimates for $X - A - H$

STIRAP

As described in section 4.2.2, we have investigated changing from STIRAP using the $X - C - H$ pathway to a new method using $X - A - H$. The A state was used for state preparation by optically pumping the $X - A$ 943 nm transition [6, 7]. However, we have never directly addressed the $A - H$ 1892 nm transition in our apparatus. Additionally, we need to perform tests in order to determine

how much laser power will be needed to saturate the two STIRAP transitions. This is of particular concern as we found that significantly more power was required to saturate STIRAP than had been observed for optical pumping on the same transitions. Before we purchase lasers to test the new STIRAP transition of $X - A - H$, we want to ensure that we will have enough power available with the lasers we might order. Here we'll go through two different empirical methods of estimating the required power based on previous measurements.

A.1 PREVIOUS MEASUREMENTS

For $X - C - H$ STIRAP in ACME II we were $> 90\%$ saturated with

$$P_{XC, STIRAP} = 15 \text{ mW},$$

$$P_{CH, STIRAP} = 10 \text{ W}.$$

It's worth noting that we were barely able to saturate STIRAP with the 10 W available to us with a Nufern fiber amplifier, but we were able to provide more power than required on the $X - C$ transition. This data can be found on pg. 60 of Cristian Panda's thesis [88].

In ACME I we used optical pumping on the $X - A$ transition (943 nm) for state preparation. With a double passed setup to fully excite from the $X, J = 1$ state, there was ~ 100 mW in one leg, and ~ 80 mW in the reflected leg. Additionally, based on minibeam measurements, Ben Spaun estimated that ACME I optical pumping could be saturated with as little as 50 mW of 943 nm light (pg. 64 of [93]). This lower bound of 50 mW is helpful, but we will primarily rely on the measured bound of 80 – 100 mW in ACME I for our calculations. Using the same saturation criteria, optical pumping on $H - C$ (1090 nm) in ACME I was used for the readout, with 2 W of power. One complication with this measurement however is that based on Ben's lablog post [156], the beam

sizes for the $X - A$ laser and the $H - C$ probe laser differed in size by a factor of 5. It's worth noting that the data shown here indicates that significantly less than 50 mW of 943 nm light is needed to saturate optical pumping in ACME I, as the fluorescence curve has essentially flattened out at that point. This also suggests that Ben's estimate for the $X - A$ transition dipole moment may also be an underestimation.

Additionally, Ben's lablog ([157, 158, 159]) describes measurements in the minibeam apparatus to measure saturation of the $X - A$ and $X - C$ transitions. Here Ben reports saturation intensities, and uses different beam sizes for the two lasers. So we will only refer to the saturation intensities here, and not the powers. For $X - C$, he measured a saturation intensity of 5 mW/cm², when he actually fit his curve with error bars, as opposed to the 10 mW/cm² stated in [157]. For $X - A$, he measured between 10 – 30 mW/cm² in [157, 158]. Here we really only need to take the upper measurement for our calculations, as that is the more conservative choice.

Finally, Ben extracted the transition dipole moments for $X - C$, $X - A$, and $H - C$ in his thesis (pg. 62 of [93]). These values are

$$d_{XC} = 0.34 \pm 0.13 ea_0,$$

$$d_{XA} = 0.071 \pm 0.038 ea_0,$$

$$d_{HC} = 0.022 \pm 0.009 ea_0.$$

However, more recent measurements by Xing Wu [87], produced a more accurate measurement of the $X - C$ transition dipole moment of

$$d_{XC} = 0.50 \pm 0.02 ea_0.$$

Additionally, Daniel Ang performed a measurement of the $X - A$ transition dipole moment [123] which found a higher value of

$$d_{XA} = 0.15 \pm 0.04 ea_0,$$

and from this also extracted the dipole moment of the $H - A$ transition

$$d_{HA} = 0.31 \pm 0.09 ea_0.$$

For this analysis we will use these more recent numbers, and note that there seems to be a systematic underestimation of the dipole moment in Ben's analysis, which will be important to consider when we use the value of d_{HC} .

A.2 METHOD 1: DIRECT SCALING OF RABI FREQUENCIES

The first empirical method we can use relies on directly comparing the Rabi frequency and power used in the ACME II STIRAP with the power required to achieve the same Rabi frequency in the new configuration. This can be done for both the $X - A$, and the $A - H$ transitions, so long as we have a way to reasonably estimate the transition dipole moment for $A - H$. This operation is helpful because it doesn't require us to consider any of the actual beam properties, as they are assumed to be

same for both STIRAP processes. We can therefore compare the required powers as follows:

$$\begin{aligned}
 \Omega_{XC} &= \Omega_{XA} \\
 d_{XC}E_{XC} &= d_{XA}E_{XA} \\
 d_{XC}\sqrt{I_{XC}} &= d_{XA}\sqrt{I_{XA}} \\
 d_{XC}^2I_{XC} &= d_{XA}^2I_{XA} \\
 d_{XC}^2P_{XC} &= d_{XA}^2P_{XA} \\
 P_{XA} &= \frac{d_{XC}^2}{d_{XA}^2}P_{XC}.
 \end{aligned} \tag{A.1}$$

An identical equation can be produced for $H - A$, using the dipole moment and power for the $H - C$ transition.

A.2.1 POWER RESULTS

Plugging these numbers into equation A.1 we get the following:

$$\begin{aligned}
 \boxed{P_{XA} \approx 170 \text{ mW}} \\
 \boxed{P_{HA} \approx 50 \text{ mW}}
 \end{aligned} \tag{A.2}$$

Notably, I don't display the standard deviation for this estimate, as the $> 50\%$ error on the dipole moments involving the A state does not propagate in a simple way. If we take the most extreme possible values of the dipole moments in question from Ben's measurements, we get maximum values of $P_{XA} = 3.7 \text{ W}$, and $P_{HA} = 2.7 \text{ W}$. These upper values are unlikely to be correct, and it is more likely that Ben's analysis underestimated the dipole moments in question, which would lower our power requirements.

A.3 METHOD 2: COMPARISONS TO OPTICAL PUMPING

Another way we can estimate the power requirements for the $X - A$ laser is to rely directly on how much power we required to optically pump each transition. This method is nice because it doesn't rely on the extracted transition dipole moments, which have very large uncertainties, and only uses the direct measurements of saturation. This method inherently assumes that the dependence on the transition dipole moment is fully captured by the saturation intensity for optical pumping. However, we do need to ensure that the correct Hönl-London factors are captured here. Fortunately, the Hönl-London factors for $|X, J = 1, M = 0\rangle \rightarrow |A, J = 0, M = 0\rangle$, which Ben used for optical pumping and the transition $|X, J = 0, M = 0\rangle \rightarrow |A, J = 1, M = 0\rangle$, which we plan to use for STIRAP are equivalent. So we only need to consider the relative Hönl-London factors of the transition we compare against. Notably, this method doesn't allow us to determine how much power we need for the A-H laser, as this transition has never been directly addressed in our experiment.

The cleanest comparison is from the measurements done in the minibeam setup [93].

Here we are considering only pump lasers for STIRAP, and these measurements were performed in the same place in the beam, and we can directly compare saturation intensities rather than powers. However, the Hönl-London factor for the optical pumping transition Ben used $|X, J = 1, M = \pm 1\rangle \rightarrow |C, J = 1, M = \pm 1\rangle$ is $\frac{1}{2}$, while the Hönl-London factor for the STIRAP transition $|X, J = 0, M = 0\rangle \rightarrow |C, J = 1, M = 0\rangle$ is $\frac{1}{\sqrt{3}}$. We can say that the STIRAP power required for $X - A$ then goes as:

$$P_{XA} = \frac{I_{Sat, XA}}{I_{Sat, XC}} P_{XC} \left(\frac{2}{\sqrt{3}} \right)^2.$$

Plugging in numbers, and taking the more conservative results from Ben's measurements:

$$\boxed{P_{XA} = 120 \text{ mW.}} \quad (\text{A.3})$$

We can attempt a slightly different comparison between the ACME I optical pumping measurements on $X - A$ and $H - C$. This calculation is complicated by the different beam sizes used in ACME I, with the $H - C$ transition having a $5\times$ larger beam size than the $X - A$ beam. However the Hönl-London factors for the optical pumping ($|H, J = 1, M = \pm 1\rangle \rightarrow |C, J = 1, M = 0\rangle$) and STIRAP ($|H, J = 1, M = \pm 1\rangle \rightarrow |C, J = 1, M = 0\rangle$) are the same and can be ignored here. So we find that:

$$P_{XA, STIRAP} = \frac{P_{OP, XA}}{(P_{OP, HC}/5)} P_{HC, STIRAP}.$$

If we plug in numbers we get:

$$\boxed{P_{XA} = 1.3 - 2.5 \text{ W.}} \quad (\text{A.4})$$

There is a very large discrepancy between these two sets of numbers, however I suspect that the real number is closer to the value calculated in section A.2. Comparing the STIRAP pump power to the STIRAP stokes power seems somewhat questionable to me, and the lower number we find here is a more direct comparison.

A.4 CONCLUSION

Based on these calculations, it is difficult to say with much certainty how much power we will require for $X - A - H$ STIRAP. However the most likely result is represented by the calculations

in section A.2, as the least assumptions are made. However, it is also likely that we will require less power than those calculations suggest, as more recent measurements indicate Ben may have underestimated the transition dipole moments.

Ultimately, the lower values suggested by the comparison with $X - C$ optical pumping make a good case for purchasing an $H - A$ laser and performing tests with the 943 nm TA¹ we already have, which is capable of producing 500 mW. This is probably also a good idea, as the upper bounds suggested by both the $H - C$ optical pumping comparison and the upper bound from the Rabi frequency method suggest there is a small possibility that we may not be able to perform STIRAP with a 3 W Ti:Sapph, which is the other option we have considered.

¹Toptica BoosTA

Bibliography

- [1] S. Guellati-Khelifa, “Searching for New Physics with the Electron’s Magnetic Moment,” *Physics* **16**, 22 (2023).
- [2] The ATLAS Collaboration, “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC,” *Physics Letters B* **716**, 1–29 (2012).
- [3] X. Fan, T. G. Myers, B. A. D. Sukra, and G. Gabrielse, “Measurement of the Electron Magnetic Moment,” *Physical Review Letters* **130**, 071801 (2023).
- [4] T. Aoyama, T. Kinoshita, and M. Nio, “Revised and improved value of the QED tenth-order electron anomalous magnetic moment,” *Physical Review D* **97**, 036001 (2018).
- [5] B. D. Fields, K. A. Olive, T.-H. Yeh, and C. Young, “Big-Bang Nucleosynthesis after Planck,” *Journal of Cosmology and Astroparticle Physics* **2020**, 010–010 (2020).
- [6] The ACME Collaboration, J. Baron, W. C. Campbell, D. DeMille, J. M. Doyle, G. Gabrielse, Y. V. Gurevich, P. W. Hess, N. R. Hutzler, E. Kirilov, I. Kozyryev, B. R. O’Leary, C. D. Panda, M. F. Parsons, E. S. Petrik, B. Spaun, A. C. Vutha, and A. D. West, “Order of Magnitude Smaller Limit on the Electric Dipole Moment of the Electron,” *Science* **343**, 269–272 (2014).
- [7] The ACME Collaboration, J. Baron, W. C. Campbell, D. DeMille, J. M. Doyle, G. Gabrielse, Y. V. Gurevich, P. W. Hess, N. R. Hutzler, E. Kirilov, I. Kozyryev, B. R. O’Leary, C. D. Panda, M. F. Parsons, B. Spaun, A. C. Vutha, A. D. West, and E. P. West, “Methods, Analysis, and the Treatment of Systematic Errors for the Electron Electric Dipole Moment Search in Thorium Monoxide,” *New Journal of Physics* **19**, 073029 (2017), [arxiv:1612.09318](https://arxiv.org/abs/1612.09318) [hep-ph, physics:physics].
- [8] The ACME Collaboration, V. Wirthl, 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, and X. Wu, “Improved limit on the electric dipole moment of the electron,” *Nature* **562**, 355–360 (2018).

- [9] Y. Ema, T. Gao, and M. Pospelov, “Standard Model prediction for paramagnetic EDMs,” *Physical Review Letters* **129**, 231801 (2022), arxiv:2202.10524 [hep-ex, physics:hep-ph, physics:nucl-th].
- [10] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, “Experimental Test of Parity Conservation in Beta Decay,” *Physical Review* **105**, 1413–1415 (1957).
- [11] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, “Evidence for the 2π Decay of the K^0 Meson,” *Physical Review Letters* **13**, 138–140 (1964).
- [12] K. Abe, K. Abe, R. Abe, I. Adachi, B. S. Ahn, H. Aihara, M. Akatsu, G. Alimonti, K. Asai, M. Asai, Y. Asano, T. Aso, V. Aulchenko, T. Aushev, A. M. Bakich, E. Banas, S. Behari, P. K. Behera, D. Beilina, A. Bondar, A. Bozek, T. E. Browder, B. C. K. Casey, P. Chang, Y. Chao, K.-F. Chen, B. G. Cheon, R. Chistov, S.-K. Choi, Y. Choi, L. Y. Dong, J. Dragic, A. Drutskoy, S. Eidelman, V. Eiges, Y. Enari, R. Enomoto, C. W. Everton, F. Fang, H. Fujii, C. Fukunaga, M. Fukushima, N. Gabyshev, A. Garmash, T. J. Gershon, A. Gordon, K. Gotow, H. Guler, R. Guo, J. Haba, H. Hamasaki, K. Hanagaki, F. Handa, K. Hara, T. Hara, N. C. Hastings, H. Hayashii, M. Hazumi, E. M. Heenan, Y. Higasino, I. Higuchi, T. Higuchi, T. Hirai, H. Hirano, T. Hojo, T. Hokuue, Y. Hoshi, K. Hoshina, S. R. Hou, W.-S. Hou, S.-C. Hsu, H.-C. Huang, Y. Igarashi, T. Iijima, H. Ikeda, K. Ikeda, K. Inami, A. Ishikawa, H. Ishino, R. Itoh, G. Iwai, H. Iwasaki, Y. Iwasaki, D. J. Jackson, P. Jalocho, H. K. Jang, M. Jones, R. Kagan, H. Kakuno, J. Kaneko, J. H. Kang, J. S. Kang, P. Kapusta, N. Katayama, H. Kawai, H. Kawai, Y. Kawakami, N. Kawamura, T. Kawasaki, H. Kichimi, D. W. Kim, H. Kim, H. J. Kim, H. Kim, S. K. Kim, T. H. Kim, K. Kinoshita, S. Kobayashi, S. Koishi, H. Konishi, K. Korotushenko, P. Krokovny, R. Kulasiri, S. Kumar, T. Kuniya, E. Kurihara, A. Kuzmin, Y.-J. Kwon, J. S. Lange, G. Leder, M. H. Lee, S. H. Lee, C. Leonidopoulos, Y.-S. Lin, D. Liventsev, R.-S. Lu, J. MacNaughton, D. Marlow, T. Matsumura, S. Matsui, S. Matsumoto, T. Matsumoto, Y. Mikami, K. Misono, K. Miyabayashi, H. Miyake, H. Miyata, L. C. Moffitt, G. R. Moloney, G. F. Moorhead, S. Mori, T. Mori, A. Murakami, T. Nagamine, Y. Nagasaka, Y. Nagashima, T. Nakadaira, T. Nakamura, E. Nakano, M. Nakao, H. Nakazawa, J. W. Nam, Z. Natkaniec, K. Neichi, S. Nishida, O. Ni-toh, S. Noguchi, T. Nozaki, S. Ogawa, T. Ohshima, Y. Ohshima, T. Okabe, T. Okazaki, S. Okuno, S. L. Olsen, H. Ozaki, P. Pakhlov, H. Palka, C. S. Park, C. W. Park, H. Park, L. S. Peak, M. Peters, L. E. Piilonen, E. Prebys, J. L. Rodriguez, N. Root, M. Rozanska, K. Rybicki, J. Ryuko, H. Sagawa, Y. Sakai, H. Sakamoto, M. Satpathy, A. Satpathy, S. Schrenk, S. Semenov, K. Senyo, Y. Settai, M. E. Sevier, H. Shibuya, B. Shwartz, A. Sidorov, S. Stanič, A. Sugi, A. Sugiyama, K. Sumisawa, T. Sumiyoshi, J.-I. Suzuki, K. Suzuki, S. Suzuki, S. Y. Suzuki, S. K. Swain, H. Tajima, T. Takahashi, F. Takasaki, M. Takita, K. Tamai, N. Tamura, J. Tanaka, M. Tanaka, G. N. Taylor, Y. Teramoto, M. Tomoto, T. Tomura, S. N. Tovey, K. Trabelsi, T. Tsuboyama, T. Tsukamoto, S. Uehara, K. Ueno, Y. Unno, S. Uno, Y. Ushiroda, S. E. Vahsen, K. E. Varvell, C. C. Wang, C. H. Wang, J. G. Wang, M.-Z. Wang, Y. Watanabe, E. Won, B. D. Yabsley, Y. Yamada, M. Yamaga, A. Yamaguchi, H. Yamamoto,

T. Yamanaka, Y. Yamashita, M. Yamauchi, S. Yanaka, J. Yashima, M. Yokoyama, K. Yoshida, Y. Yusa, H. Yuta, C. C. Zhang, J. Zhang, H. W. Zhao, Y. Zheng, V. Zhilich, and D. Žontar, “Observation of Large CP Violation in the Neutral B Meson System,” *Physical Review Letters* **87**, 091802 (2001).

- [13] R. Aaij, C. Abellán Beteta, B. Adeva, M. Adinolfi, C. A. Aidala, Z. Ajaltouni, S. Akar, P. Albicocco, J. Albrecht, F. Alessio, M. Alexander, A. Alfonso Albero, G. Alkhazov, P. Alvarez Cartelle, A. A. Alves, S. Amato, Y. Amhis, L. An, L. Anderlini, G. Andreassi, M. Andreotti, J. E. Andrews, F. Archilli, P. d’Argent, J. Arnau Romeu, A. Artamonov, M. Artuso, K. Arzymatov, E. Aslanides, M. Atzeni, B. Audurier, S. Bachmann, J. J. Back, S. Baker, V. Balagura, W. Baldini, A. Baranov, R. J. Barlow, S. Barsuk, W. Barter, M. Bartolini, F. Baryshnikov, V. Batozskaya, B. Batsukh, A. Battig, V. Battista, A. Bay, F. Bedeschi, I. Bediaga, A. Beiter, L. J. Bel, S. Belin, N. Beliy, V. Bellee, N. Belloli, K. Belous, I. Belyaev, E. Ben-Haim, G. Bencivenni, S. Benson, S. Beranek, A. Berezhnoy, R. Bernet, D. Berninghoff, E. Bertholet, A. Bertolin, C. Betancourt, F. Betti, M. O. Bettler, M. Van Beuzekom, I. Bezshyiko, S. Bhasin, J. Bhom, M. S. Bieker, S. Bifani, P. Billoir, A. Birnkraut, A. Bizzeti, M. Bjørn, M. P. Blago, T. Blake, F. Blanc, S. Blusk, D. Bobulska, V. Bocci, O. Boente Garcia, T. Boettcher, A. Bondar, N. Bondar, S. Borghi, M. Borisyak, M. Borsato, M. Boubdir, T. J. V. Bowcock, C. Bozzi, S. Braun, M. Brodski, J. Brodzicka, A. Brossa Gonzalo, D. Brundu, E. Buchanan, A. Buonauro, C. Burr, A. Bursche, J. Buytaert, W. Byczynski, S. Cadetdu, H. Cai, R. Calabrese, S. Cali, R. Calladine, M. Calvi, M. Calvo Gomez, A. Camboni, P. Campana, D. H. Campora Perez, L. Capriotti, A. Carbone, G. Carboni, R. Cardinale, A. Cardini, P. Carniti, K. Carvalho Akiba, G. Casse, M. Cattaneo, G. Cavallero, R. Cenci, M. G. Chapman, M. Charles, Ph. Charpentier, G. Chatzikonstantinidis, M. Chefdeville, V. Chekalina, C. Chen, S. Chen, S.-G. Chitic, V. Chobanova, M. Chrzaszcz, A. Chubykin, P. Ciambone, X. Cid Vidal, G. Ciezarek, F. Cindolo, P. E. L. Clarke, M. Clemencic, H. V. Cliff, J. Closier, V. Coco, J. A. B. Coelho, J. Cogan, E. Cogneras, L. Cojocariu, P. Collins, T. Colombo, A. Comerma-Montells, A. Contu, G. Coombs, S. Coquereau, G. Corti, C. M. Costa Sobral, B. Couturier, G. A. Cowan, D. C. Craik, A. Crocombe, M. Cruz Torres, R. Currie, C. D’Ambrosio, C. L. Da Silva, E. Dall’Occo, J. Dalseno, A. Danilina, A. Davis, O. De Aguiar Francisco, K. De Bruyn, S. De Capua, M. De Cian, J. M. De Miranda, L. De Paula, M. De Serio, P. De Simone, C. T. Dean, W. Dean, D. Decamp, L. Del Buono, B. Delaney, H.-P. Dembinski, M. Demmer, A. Dendek, D. Derkach, O. Deschamps, F. Dese, F. Dettori, B. Dey, A. Di Canto, P. Di Nezza, S. Didenko, H. Dijkstra, F. Dordei, M. Dorigo, A. Dosil Suárez, L. Douglas, A. Dovbnya, K. Dreimanis, L. Dufour, G. Dujany, P. Durante, J. M. Durham, D. Dutta, R. Dzhelyadin, M. Dziwiecki, A. Dziurda, A. Dzyuba, S. Easo, U. Egede, V. Egorychev, S. Eidelman, S. Eisenhardt, U. Eitschberger, R. Ekelhof, L. Eklund, S. Ely, A. Ene, S. Escher, S. Esen, T. Evans, A. Falabella, N. Farley, S. Farry, D. Fazzini, P. Fernandez Declara, A. Fernandez Prieto, F. Ferrari, L. Ferreira Lopes, F. Ferreira Rodrigues, S. Ferreres Sole, M. Ferro-Luzzi, S. Filippov, R. A. Fini, M. Fiorini, M. Firlej, C. Fitzpatrick, T. Fiutowski, F. Fleuret,

M. Fontana, F. Fontanelli, R. Forty, V. Franco Lima, M. Frank, C. Frei, J. Fu, W. Funk, C. Färber, M. Féo, E. Gabriel, A. Gallas Torreira, D. Galli, S. Gallorini, S. Gambetta, Y. Gan, M. Gandelman, P. Gandini, Y. Gao, L. M. Garcia Martin, B. Garcia Plana, J. García Pardiñas, J. Garra Tico, L. Garrido, D. Gascon, C. Gaspar, G. Gazzoni, D. Gerick, E. Gersabeck, M. Gersabeck, T. Gershon, D. Gerstel, Ph. Ghez, V. Gibson, O. G. Girard, P. Gironella Gironell, L. Giubega, K. Gizdov, V. V. Gligorov, D. Golubkov, A. Golutvin, A. Gomes, I. V. Gorelov, C. Gotti, E. Govorkova, J. P. Grabowski, R. Graciani Diaz, L. A. Granado Cardoso, E. Graugés, E. Graverini, G. Graziani, A. Grecu, R. Greim, P. Griffith, L. Grillo, L. Gruber, B. R. Gruberg Cazon, C. Gu, X. Guo, E. Gushchin, A. Guth, Yu. Guz, T. Gys, C. Göbel, T. Hadavizadeh, C. Hadjivasiliou, G. Haefeli, C. Haen, S. C. Haines, B. Hamilton, Q. Han, X. Han, T. H. Hancock, S. Hansmann-Menzemer, N. Harnew, T. Harrison, C. Hasse, M. Hatch, J. He, M. Hecker, K. Heinicke, A. Heister, K. Hennessey, L. Henry, E. Van Herwijnen, J. Heuel, M. Heß, A. Hicheur, R. Hidalgo Charman, D. Hill, M. Hilton, P. H. Hopchev, J. Hu, W. Hu, W. Huang, Z. C. Huard, W. Hulsbergen, T. Humair, M. Hushchyn, D. Hutchcroft, D. Hynds, P. Ibis, M. Idzik, P. Ilten, A. Inglese, A. Inyakin, K. Ivshin, R. Jacobsson, S. Jakobsen, J. Jalocha, E. Jans, B. K. Jashal, A. Jawahery, F. Jiang, M. John, D. Johnson, C. R. Jones, C. Joram, B. Jost, N. Jurik, S. Kandybei, M. Karacson, J. M. Kariuki, S. Karodia, N. Kazeev, M. Kecke, F. Keizer, M. Kelsey, M. Kenzie, T. Ketel, B. Khanji, A. Kharisova, C. Khurewathanakul, K. E. Kim, T. Kirn, V. S. Kirsobom, S. Klaver, K. Klimaszewski, S. Koliiev, M. Kolpin, R. Kopečna, P. Koppenburg, I. Kostiuik, S. Kotriakhova, M. Kozeiha, L. Kravchuk, M. Kreps, F. Kress, S. Kretzschmar, P. Krokovny, W. Krupa, W. Krzemien, W. Kucewicz, M. Kucharczyk, V. Kudryavtsev, G. J. Kunde, A. K. Kuonen, T. Kvaratskheliya, D. Lacarrere, G. Lafferty, A. Lai, D. Lancierini, G. Lanfranchi, C. Langenbruch, T. Latham, C. Lazzeroni, R. Le Gac, A. Leflat, R. Lefèvre, F. Lemaitre, O. Leroy, T. Lesiak, B. Leverington, H. Li, P.-R. Li, X. Li, Y. Li, Z. Li, X. Liang, T. Likhomanenko, R. Lindner, P. Ling, F. Lionetto, V. Lisovskyi, G. Liu, X. Liu, D. Loh, A. Loi, I. Longstaff, J. H. Lopes, G. Loustau, G. H. Lovell, D. Lucchesi, M. Lucio Martinez, Y. Luo, A. Lupato, E. Luppi, O. Lupton, A. Lusiani, X. Lyu, R. Ma, F. Machefer, F. Maciuc, V. Macko, P. Mackowiak, S. Maddrell-Mander, O. Maev, K. Maguire, D. Maisuzenko, M. W. Majewski, S. Malde, B. Malecki, A. Malinin, T. Maltsev, H. Malygina, G. Manca, G. Mancinelli, D. Marangotto, J. Maratas, J. F. Marchand, U. Marconi, C. Marin Benito, M. Marinangeli, P. Marino, J. Marks, P. J. Marshall, G. Martellotti, M. Martinelli, D. Martinez Santos, F. Martinez Vidal, A. Massafferri, M. Materok, R. Matev, A. Mathad, Z. Mathe, V. Matiunin, C. Matteuzzi, K. R. Mattioli, A. Mauri, E. Maurice, B. Maurin, M. McCann, A. McNab, R. McNulty, J. V. Mead, B. Meadows, C. Meaux, N. Meinert, D. Melnychuk, M. Merk, A. Merli, E. Michielin, D. A. Milanes, E. Millard, M.-N. Minard, L. Minzoni, D. S. Mitzel, A. Mogini, R. D. Moise, T. Mombächer, I. A. Monroy, S. Monteil, M. Morandin, G. Morello, M. J. Morello, J. Moron, A. B. Morris, R. Mountain, F. Muheim, M. Mukherjee, M. Mulder, C. H. Murphy, D. Murray, A. Mödden, D. Müller, J. Müller, K. Müller, V. Müller, P. Naik, T. Nakada, R. Nandakumar, A. Nandi, T. Nanut, I. Nasteva, M. Needham, N. Neri, S. Neubert, N. Neufeld, R. Newcombe, T. D.

Nguyen, C. Nguyen-Mau, S. Nieswand, R. Niet, N. Nikitin, N. S. Nolte, D. P. O'Hanlon, A. Oblakowska-Mucha, V. Obraztsov, S. Ogilvy, R. Oldeman, C. J. G. Onderwater, J. D. Osborn, A. Ossowska, J. M. Otalora Goicochea, T. Ovsiannikova, P. Owen, A. Oyanguren, P. R. Pais, T. Pajero, A. Palano, M. Palutan, G. Panshin, A. Papanestis, M. Pappagallo, L. L. Pappalardo, W. Parker, C. Parkes, G. Passaleva, A. Pastore, M. Patel, C. Patrignani, A. Pearce, A. Pellegrino, G. Penso, M. Pepe Altarelli, S. Perazzini, D. Pereima, P. Perret, L. Pescatore, K. Petridis, A. Petrolini, A. Petrov, S. Petrucci, M. Petruzzo, B. Pietrzyk, G. Pietrzyk, M. Pikies, M. Pili, D. Pinci, J. Pinzino, F. Pisani, A. Piucci, V. Placinta, S. Playfer, J. Plews, M. Plo Casasus, F. Polci, M. Poli Lener, M. Poliakov, A. Poluektov, N. Polukhina, I. Polyakov, E. Polycarpo, G. J. Pomery, S. Ponce, A. Popov, D. Popov, S. Poslavskii, E. Price, C. Prouve, V. Pugatch, A. Puig Navarro, H. Pullen, G. Punzi, W. Qian, J. Qin, R. Quagliani, B. Quintana, N. V. Raab, B. Rachwal, J. H. Rademacker, M. Rama, M. Ramos Pernas, M. S. Rangel, F. Ratnikov, G. Raven, M. Ravonel Salzgeber, M. Reboud, F. Redi, S. Reichert, A. C. Dos Reis, F. Reiss, C. Remon Alepuz, Z. Ren, V. Renaudin, S. Ricciardi, S. Richards, K. Rinnert, P. Robbe, A. Robert, A. B. Rodrigues, E. Rodrigues, J. A. Rodriguez Lopez, M. Roehrken, S. Roiser, A. Rollings, V. Romanovskiy, A. Romero Vidal, J. D. Roth, M. Rotondo, M. S. Rudolph, T. Ruf, J. Ruiz Vidal, J. J. Saborido Silva, N. Sagidova, B. Saitta, V. Salustino Guimaraes, C. Sanchez Gras, C. Sanchez Mayor-domo, B. Sanmartin Sedes, R. Santacesaria, C. Santamarina Rios, M. Santimaria, E. Santovetti, G. Sarpis, A. Sarti, C. Satriano, A. Satta, M. Saur, D. Savrina, S. Schael, M. Schellenberg, M. Schiller, H. Schindler, M. Schmelling, T. Schmelzer, B. Schmidt, O. Schneider, A. Schopper, H. F. Schreiner, M. Schubiger, S. Schulte, M. H. Schune, R. Schwemmer, B. Sciascia, A. Sciubba, A. Semennikov, E. S. Sepulveda, A. Sergi, N. Serra, J. Serrano, L. Sestini, A. Seuthe, P. Seyfert, M. Shapkin, T. Shears, L. Shekhtman, V. Shevchenko, E. Shmanin, B. G. Siddi, R. Silva Coutinho, L. Silva De Oliveira, G. Simi, S. Simone, I. Skiba, N. Skidmore, T. Skwarnicki, M. W. Slater, J. G. Smeaton, E. Smith, I. T. Smith, M. Smith, M. Soares, L. Soares Lavra, M. D. Sokoloff, F. J. P. Soler, B. Souza De Paula, B. Spaan, E. Spadaro Norella, P. Spradlin, F. Stagni, M. Stahl, S. Stahl, P. Stefko, S. Stefkova, O. Steinkamp, S. Stemmler, O. Stenyakin, M. Stepanova, H. Stevens, A. Stocchi, S. Stone, S. Stracka, M. E. Stramaglia, M. Straticiu, U. Straumann, S. Strokov, J. Sun, L. Sun, Y. Sun, K. Swientek, A. Szabelski, T. Szumlak, M. Szymanski, S. T'Jampens, Z. Tang, T. Tekampe, G. Tellarini, F. Teubert, E. Thomas, J. Van Tilburg, M. J. Tilley, V. Tisserand, M. Tobin, S. Tolk, L. Tomassetti, D. Tonelli, D. Y. Tou, R. Tourinho Jadallah Aoude, E. Tournefier, M. Traill, M. T. Tran, A. Trisovic, A. Tsaregorodtsev, G. Tuci, A. Tully, N. Tuning, A. Ukleja, A. Usachov, A. Ustyuzhanin, U. Uwer, A. Vagner, V. Vagnoni, A. Valassi, S. Valat, G. Valenti, H. Van Hecke, C. B. Van Hulse, R. Vazquez Gomez, P. Vazquez Regueiro, S. Vecchi, M. Van Veghel, J. J. Velthuis, M. Veltri, A. Venkateswaran, M. Vernet, M. Veronesi, M. Vesterinen, J. V. Viana Barbosa, D. Vieira, M. Vieites Diaz, H. Viemann, X. Vilasis-Cardona, A. Vitkovskiy, M. Vitti, V. Volkov, A. Vollhardt, D. Vom Bruch, B. Voneki, A. Vorobyev, V. Vorobyev, N. Voropaev, J. A. De Vries, C. Vázquez Sierra, R. Waldi, J. Walsh, J. Wang, M. Wang, Y. Wang, Z. Wang, D. R. Ward, H. M. Wark, N. K. Wat-

son, D. Websdale, A. Weiden, C. Weisser, M. Whitehead, G. Wilkinson, M. Wilkinson, I. Williams, M. R. J. Williams, M. Williams, T. Williams, F. F. Wilson, M. Winn, W. Wislicki, M. Witek, G. Wormser, S. A. Wotton, K. Wyllie, D. Xiao, Y. Xie, H. Xing, A. Xu, M. Xu, Q. Xu, Z. Xu, Z. Xu, Z. Yang, Z. Yang, Y. Yao, L. E. Yeomans, H. Yin, J. Yu, X. Yuan, O. Yushchenko, K. A. Zarebski, M. Zavertyaev, M. Zeng, D. Zhang, L. Zhang, W. C. Zhang, Y. Zhang, A. Zhelezov, Y. Zheng, X. Zhu, V. Zhukov, J. B. Zonneveld, S. Zucchelli, and LHCb Collaboration, “Observation of C P Violation in Charm Decays,” *Physical Review Letters* **122**, 211803 (2019).

- [14] M. Sozzi, *Discrete Symmetries and CP Violation: From Experiment to Theory*, Oxford Graduate Texts (Oxford University Press, Oxford, 2012).
- [15] J. P. Lees, V. Poireau, V. Tisserand, J. Garra Tico, E. Grauges, A. Palano, G. Eigen, B. Stugu, D. N. Brown, L. T. Kerth, Yu. G. Kolomensky, G. Lynch, H. Koch, T. Schroeder, D. J. Asgeirsson, C. Hearty, T. S. Mattison, J. A. McKenna, R. Y. So, A. Khan, V. E. Blinov, A. R. Buzykaev, V. P. Druzhinin, V. B. Golubev, E. A. Kravchenko, A. P. Onuchin, S. I. Serednyakov, Yu. I. Skovpen, E. P. Solodov, K. Yu. Todyshev, A. N. Yushkov, M. Bondioli, D. Kirkby, A. J. Lankford, M. Mandelkern, H. Atmacan, J. W. Gary, F. Liu, O. Long, G. M. Vitug, C. Campagnari, T. M. Hong, D. Kovalskyi, J. D. Richman, C. A. West, A. M. Eisner, J. Kroseberg, W. S. Lockman, A. J. Martinez, B. A. Schumm, A. Seiden, D. S. Chao, C. H. Cheng, B. Echenard, K. T. Flood, D. G. Hitlin, P. Ongmongkolkul, F. C. Porter, A. Y. Rakitin, R. Andreassen, Z. Huard, B. T. Meadows, M. D. Sokoloff, L. Sun, P. C. Bloom, W. T. Ford, A. Gaz, U. Nauenberg, J. G. Smith, S. R. Wagner, R. Ayad, W. H. Toki, B. Spaan, K. R. Schubert, R. Schwierz, D. Bernard, M. Verderi, P. J. Clark, S. Playfer, D. Bettoni, C. Bozzi, R. Calabrese, G. Cibinetto, E. Fioravanti, I. Garzia, E. Luppi, M. Munerato, L. Piemontese, V. Santoro, R. Baldini-Ferrolì, A. Calcaterra, R. de Sangro, G. Finocchiaro, P. Patteri, I. M. Peruzzi, M. Piccolo, M. Rama, A. Zallo, R. Contri, E. Guido, M. Lo Vetere, M. R. Monge, S. Passaggio, C. Patrignani, E. Robutti, B. Bhuyan, V. Prasad, C. L. Lee, M. Morii, A. J. Edwards, A. Adametz, U. Uwer, H. M. Lacker, T. Lueck, P. D. Dauncey, U. Mallik, C. Chen, J. Cochran, W. T. Meyer, S. Prell, A. E. Rubin, A. V. Gritsan, Z. J. Guo, N. Arnaud, M. Davier, D. Derkach, G. Grosdidier, F. Le Diberder, A. M. Lutz, B. Malaescu, P. Roudeau, M. H. Schune, A. Stocchi, G. Wormser, D. J. Lange, D. M. Wright, C. A. Chavez, J. P. Coleman, J. R. Fry, E. Gabathuler, D. E. Hutchcroft, D. J. Payne, C. Touramanis, A. J. Bevan, F. Di Lodovico, R. Sacco, M. Sigamani, G. Cowan, D. N. Brown, C. L. Davis, A. G. Denig, M. Fritsch, W. Gradl, K. Griessinger, A. Hafner, E. Prencipe, R. J. Barlow, G. Jackson, G. D. Lafferty, E. Behn, R. Cenci, B. Hamilton, A. Jawahery, D. A. Roberts, C. Dallapiccola, R. Cowan, D. Dujmic, G. Sciolla, R. Cheaib, D. Lindemann, P. M. Patel, S. H. Robertson, P. Biassoni, N. Neri, F. Palombo, S. Stracka, L. Cremaldi, R. Godang, R. Kroeger, P. Sonnek, D. J. Summers, X. Nguyen, M. Simard, P. Taras, G. De Nardo, D. Monorchio, G. Onorato, C. Sciacca, M. Martinelli, G. Raven, C. P. Jessop, J. M. LoSecco, W. F. Wang, K. Honscheid, R. Kass, J. Brau, R. Frey, N. B. Sinev, D. Strom, E. Torrence, E. Feltresi, N. Gagliardi, M. Margoni, M. Morandin, A. Pompili, M. Posocco,

M. Rotondo, G. Simi, F. Simonetto, R. Stroili, S. Akar, E. Ben-Haim, M. Bomben, G. R. Bonneaud, H. Briand, G. Calderini, J. Chauveau, O. Hamon, Ph. Leruste, G. Marchiori, J. Ocariz, S. Sitt, M. Biasini, E. Manoni, S. Pacetti, A. Rossi, C. Angelini, G. Batignani, S. Bettarini, M. Carpinelli, G. Casarosa, A. Cervelli, F. Forti, M. A. Giorgi, A. Lusiani, B. Oberhof, E. Paoloni, A. Perez, G. Rizzo, J. J. Walsh, D. Lopes Pegna, J. Olsen, A. J. S. Smith, A. V. Telnov, F. Anulli, R. Faccini, F. Ferrarotto, F. Ferroni, M. Gaspero, L. Li Gioi, M. A. Mazzoni, G. Piredda, C. Bünger, O. Grünberg, T. Hartmann, T. Leddig, H. Schröder, C. Voss, R. Waldi, T. Adye, E. O. Olaiya, F. F. Wilson, S. Emery, G. Hamel de Monchenault, G. Vasseur, Ch. Yèche, D. Aston, D. J. Bard, R. Bartoldus, J. F. Benitez, C. Cartaro, M. R. Convery, J. Dorfan, G. P. Dubois-Felsmann, W. Dunwoodie, M. Ebert, R. C. Field, M. Franco Sevilla, B. G. Fulsom, A. M. Gabareen, M. T. Graham, P. Grenier, C. Hast, W. R. Innes, M. H. Kelsey, P. Kim, M. L. Kocian, D. W. G. S. Leith, P. Lewis, B. Lindquist, S. Luitz, V. Luth, H. L. Lynch, D. B. MacFarlane, D. R. Muller, H. Neal, S. Nelson, M. Perl, T. Pulliam, B. N. Ratcliff, A. Roodman, A. A. Salnikov, R. H. Schindler, A. Snyder, D. Su, M. K. Sullivan, J. Va'vra, A. P. Wagner, W. J. Wisniewski, M. Wittgen, D. H. Wright, H. W. Wulsin, C. C. Young, V. Ziegler, W. Park, M. V. Purohit, R. M. White, J. R. Wilson, A. Randle-Conde, S. J. Sekula, M. Bellis, P. R. Burchat, T. S. Miyashita, E. M. T. Puccio, M. S. Alam, J. A. Ernst, R. Gorodeisky, N. Guttman, D. R. Peimer, A. Soffer, P. Lund, S. M. Spanier, J. L. Ritchie, A. M. Ruland, R. F. Schwitters, B. C. Wray, J. M. Izen, X. C. Lou, F. Bianchi, D. Gamba, S. Zambito, L. Lanceri, L. Vitale, J. Bernabeu, F. Martinez-Vidal, A. Oyanguren, P. Villanueva-Perez, H. Ahmed, J. Albert, Sw. Banerjee, F. U. Bernlochner, H. H. F. Choi, G. J. King, R. Kowalewski, M. J. Lewczuk, I. M. Nugent, J. M. Roney, R. J. Sobie, N. Tasneem, T. J. Gershon, P. F. Harrison, T. E. Latham, H. R. Band, S. Dasu, Y. Pan, R. Prepost, S. L. Wu, and The BABAR Collaboration, "Observation of Time-Reversal Violation in the B^0 Meson System," *Physical Review Letters* **109**, 211801 (2012).

- [16] I. B. Khriplovich, *Cp Violation without Strangeness: Electric Dipole Moments of Particles, Atoms, and Molecules* (Springer, Place of publication not identified, 2012).
- [17] V. A. Kostelecký and N. Russell, "Data tables for Lorentz and C P T violation," *Reviews of Modern Physics* **83**, 11–31 (2011).
- [18] G. Gabrielse, S. F. Hoogerheide, J. Dorr, and E. Novitski, "Precise Matter and Antimatter Tests of the Standard Model with e^- , e^+ , p , \bar{p} , and \bar{H} ," in *Fundamental Physics in Particle Traps*, Vol. 256, edited by W. Quint and M. Vogel (Springer Berlin Heidelberg, Berlin, Heidelberg, 2014) pp. 1–40.
- [19] J. DiSciaccia, M. Marshall, K. Marable, G. Gabrielse, S. Ettenauer, E. Tardiff, R. Kalra, D. W. Fitzakerley, M. C. George, E. A. Hessels, C. H. Storry, M. Weel, D. Grzonka, W. Oelert, T. Sefzick, and ATRAP Collaboration, "One-Particle Measurement of the Antiproton Magnetic Moment," *Physical Review Letters* **110**, 130801 (2013).

- [20] A. D. Sakharov, “Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe,” *JETP Letters* **5**, 24–27 (1967).
- [21] N. Cabibbo, “Unitary Symmetry and Leptonic Decays,” *Physical Review Letters* **10**, 531–533 (1963).
- [22] M. Kobayashi and T. Maskawa, “CP Violation in the Renormalizable Theory of Weak Interaction,” *Progress of Theoretical Physics* **49**, 652–657 (1973).
- [23] W. Bernreuther and M. Suzuki, “The electric dipole moment of the electron,” *Reviews of Modern Physics* **63**, 313–340 (1991).
- [24] M. B. Gavela, P. Hernandez, J. Orloff, and O. Pene, “Standard Model CP-violation and baryon asymmetry,” *Modern Physics Letters A* **09**, 795–809 (1994).
- [25] E. M. Purcell and N. F. Ramsey, “On the Possibility of Electric Dipole Moments for Elementary Particles and Nuclei,” *Physical Review* **78**, 807–807 (1950).
- [26] C. A. Baker, D. D. Doyle, P. Geltenbort, K. Green, M. G. D. van der Grinten, P. G. Harris, P. Iaydjiev, S. N. Ivanov, D. J. R. May, J. M. Pendlebury, J. D. Richardson, D. Shiers, and K. F. Smith, “Improved Experimental Limit on the Electric Dipole Moment of the Neutron,” *Physical Review Letters* **97**, 131801 (2006).
- [27] J. M. Pendlebury, S. Afach, N. J. Ayres, C. A. Baker, G. Ban, G. Bison, K. Bodek, M. Burghoff, P. Geltenbort, K. Green, W. C. Griffith, M. van der Grinten, Z. D. Grujić, P. G. Harris, V. Hélaine, P. Iaydjiev, S. N. Ivanov, M. Kasprzak, Y. Kermaidic, K. Kirch, H.-C. Koch, S. Komposch, A. Kozela, J. Krempel, B. Lauss, T. Lefort, Y. Lemièrre, D. J. R. May, M. Musgrave, O. Naviliat-Cuncic, F. M. Piegsa, G. Pignol, P. N. Prashanth, G. Quémener, M. Rawlik, D. Rebreyend, J. D. Richardson, D. Ries, S. Rocca, D. Rozpedzik, A. Schnabel, P. Schmidt-Wellenburg, N. Severijns, D. Shiers, J. A. Thorne, A. Weis, O. J. Winston, E. Wursten, J. Zejma, and G. Zsigmond, “Revised experimental upper limit on the electric dipole moment of the neutron,” *Physical Review D* **92**, 092003 (2015).
- [28] R. Schmid, *New Search for the Neutron Electric Dipole Moment Using Ultracold Neutrons at the Spallation Neutron Source*, Ph.D. thesis, California Institute of Technology (2014).
- [29] R. Chislett, “The muon EDM in the g-2 experiment at Fermilab,” *EPJ Web of Conferences* **118**, 01005 (2016).
- [30] S. K. Lamoreaux, J. P. Jacobs, B. R. Heckel, F. J. Raab, and N. Fortson, “New constraints on time-reversal asymmetry from a search for a permanent electric dipole moment of Hg 199,” *Physical Review Letters* **59**, 2275–2278 (1987).
- [31] W. C. Griffith, M. D. Swallows, T. H. Loftus, M. V. Romalis, B. R. Heckel, and E. N. Fortson, “Improved limit on the permanent electric dipole moment of ^{199}Hg ,” *Physical Review Letters* **102**, 101601 (2009), arxiv:0901.2328 [physics].

- [32] B. Graner, Y. Chen, E. G. Lindahl, and B. R. Heckel, “Reduced Limit on the Permanent Electric Dipole Moment of Hg 199,” *Physical Review Letters* **116**, 161601 (2016).
- [33] R. H. Parker, M. R. Dietrich, M. R. Kalita, N. D. Lemke, K. G. Bailey, M. Bishof, J. P. Greene, R. J. Holt, W. Korsch, Z.-T. Lu, P. Mueller, T. P. O’Connor, and J. T. Singh, “First Measurement of the Atomic Electric Dipole Moment of Ra 225,” *Physical Review Letters* **114**, 233002 (2015).
- [34] Y. Nakai and M. Reece, “Electric dipole moments in natural supersymmetry,” *Journal of High Energy Physics* **2017**, 31 (2017).
- [35] S. M. Barr and A. Zee, “Electric dipole moment of the electron and of the neutron,” *Physical Review Letters* **65**, 21–24 (1990).
- [36] M. Pospelov and A. Ritz, “CKM benchmarks for electron electric dipole moment experiments,” *Physical Review D* **89**, 056006 (2014).
- [37] C. Cesarotti, Q. Lu, Y. Nakai, A. Parikh, and M. Reece, “Interpreting the electron EDM constraint,” *Journal of High Energy Physics* **2019**, 59 (2019).
- [38] M. Pospelov and A. Ritz, “Electric dipole moments as probes of new physics,” *Annals of Physics* **318**, 119–169 (2005).
- [39] A. C. Vutha, M. Horbatsch, and E. A. Hessels, “Orientation-dependent hyperfine structure of polar molecules in a rare-gas matrix: A scheme for measuring the electron electric dipole moment,” *Physical Review A* **98**, 032513 (2018).
- [40] N. Fortson, P. Sandars, and S. Barr, “The Search for a Permanent Electric Dipole Moment,” *Physics Today* **56**, 33–39 (2003).
- [41] J. Engel, M. J. Ramsey-Musolf, and U. van Kolck, “Electric dipole moments of nucleons, nuclei, and atoms: The Standard Model and beyond,” *Progress in Particle and Nuclear Physics Fundamental Symmetries in the Era of the LHC*, **71**, 21–74 (2013).
- [42] J. L. Feng, “Naturalness and the Status of Supersymmetry,” *Annual Review of Nuclear and Particle Science* **63**, 351–382 (2013).
- [43] P. G. H. Sandars and E. Lipworth, “Electric Dipole Moment of the Cesium Atom. A New Upper Limit to the Electric Dipole Moment of the Free Electron,” *Physical Review Letters* **13**, 718–720 (1964).
- [44] L. V. Skripnikov, “Combined 4-component and relativistic pseudopotential study of ThO for the electron electric dipole moment search,” *The Journal of Chemical Physics* **145**, 214301 (2016).

- [45] M. Denis and T. Fleig, “In search of discrete symmetry violations beyond the standard model: Thorium monoxide reloaded,” *The Journal of Chemical Physics* **145**, 214307 (2016).
- [46] L. I. Schiff, “Measurability of Nuclear Electric Dipole Moments,” *Physical Review* **132**, 2194–2200 (1963).
- [47] P. Sandars, “The electric dipole moment of an atom,” *Physics Letters* **14**, 194–196 (1965).
- [48] E. D. Commins, J. D. Jackson, and D. P. DeMille, “The electric dipole moment of the electron: An intuitive explanation for the evasion of Schiff’s theorem,” *American Journal of Physics* **75**, 532–536 (2007).
- [49] G. Feinberg, “Effects of an Electric Dipole Moment of the Electron on the Hydrogen Energy Levels,” *Physical Review* **112**, 1637–1642 (1958).
- [50] D. F. Nelson, A. A. Schupp, R. W. Pidd, and H. R. Crane, “Search for an Electric Dipole Moment of the Electron,” *Physical Review Letters* **2**, 492–495 (1959).
- [51] D. T. Wilkinson and H. R. Crane, “Precision Measurement of the g Factor of the Free Electron,” *Physical Review* **130**, 852–863 (1963).
- [52] M. C. Weisskopf, J. P. Carrico, H. Gould, E. Lipworth, and T. S. Stein, “Electric Dipole Moment of the Cesium Atom. A New Upper Limit to the Electric Dipole Moment of the Electron,” *Physical Review Letters* **21**, 1645–1648 (1968).
- [53] S. A. Murthy, D. Krause, Z. L. Li, and L. R. Hunter, “New limits on the electron electric dipole moment from cesium,” *Physical Review Letters* **63**, 965–968 (1989).
- [54] E. S. Ensberg, “Experimental Upper Limit for the Permanent Electric Dipole Moment of Rb 85 by Optical-Pumping Techniques,” *Physical Review* **153**, 36–43 (1967).
- [55] M. A. Player and P. G. H. Sandars, “An experiment to search for an electric dipole moment in the 3P_2 metastable state of xenon,” *Journal of Physics B: Atomic and Molecular Physics* **3**, 1620–1635 (1970).
- [56] K. Abdullah, C. Carlberg, E. D. Commins, H. Gould, and S. B. Ross, “New experimental limit on the electron electric dipole moment,” *Physical Review Letters* **65**, 2347–2350 (1990).
- [57] E. D. Commins, S. B. Ross, D. DeMille, and B. C. Regan, “Improved experimental limit on the electric dipole moment of the electron,” *Physical Review A* **50**, 2960–2977 (1994).
- [58] B. C. Regan, E. D. Commins, C. J. Schmidt, and D. DeMille, “New Limit on the Electron Electric Dipole Moment,” *Physical Review Letters* **88**, 071805 (2002).

- [59] D. Cho, K. Sangster, and E. A. Hinds, “Tenfold improvement of limits on T violation in thallium fluoride,” *Physical Review Letters* **63**, 2559–2562 (1989).
- [60] J. J. Hudson, B. E. Sauer, M. R. Tarbutt, and E. A. Hinds, “Measurement of the Electron Electric Dipole Moment Using YbF Molecules,” *Physical Review Letters* **89**, 023003 (2002).
- [61] J. J. Hudson, D. M. Kara, I. J. Smallman, B. E. Sauer, M. R. Tarbutt, and E. A. Hinds, “Improved measurement of the shape of the electron,” *Nature* **473**, 493–496 (2011).
- [62] S. Eckel, P. Hamilton, E. Kirilov, H. W. Smith, and D. DeMille, “Search for the electron electric dipole moment using Ω -doublet levels in PbO,” *Physical Review A* **87**, 052130 (2013).
- [63] W. B. Cairncross, D. N. Gresh, M. Grau, K. C. Cossel, T. S. Roussy, Y. Ni, Y. Zhou, J. Ye, and E. A. Cornell, “Precision Measurement of the Electron’s Electric Dipole Moment Using Trapped Molecular Ions,” *Physical Review Letters* **119**, 153001 (2017).
- [64] T. S. Roussy, L. Caldwell, T. Wright, W. B. Cairncross, Y. Shagam, K. B. Ng, N. Schlossberger, S. Y. Park, A. Wang, J. Ye, and E. A. Cornell, “A new bound on the electron’s electric dipole moment,” *arXiv* (2022), 10.48550/ARXIV.2212.11841.
- [65] K. B. Ng, Y. Zhou, L. Cheng, N. Schlossberger, S. Y. Park, T. S. Roussy, L. Caldwell, Y. Shagam, A. J. Vigil, E. A. Cornell, and J. Ye, “Spectroscopy on the electron-electric-dipole-moment-sensitive states of ThF⁺,” *Physical Review A* **105**, 022823 (2022).
- [66] X. Alauze, J. Lim, M. A. Trigatzis, S. Swarbrick, F. J. Collings, N. J. Fitch, B. E. Sauer, and M. R. Tarbutt, “An ultracold molecular beam for testing fundamental physics,” *Quantum Science and Technology* **6**, 044005 (2021).
- [67] I. Kozyryev and N. R. Hutzler, “Precision Measurement of Time-Reversal Symmetry Violation with Laser-Cooled Polyatomic Molecules,” *Physical Review Letters* **119**, 133002 (2017).
- [68] R. Mitra, V. S. Prasanna, R. F. G. Ruiz, T. K. Sato, M. Abe, Y. Sakemi, B. P. Das, and B. K. Sahoo, “Towards C P -violation studies on superheavy molecules: Theoretical and experimental perspectives,” *Physical Review A* **104**, 062801 (2021).
- [69] C. D. Panda, C. Meisenhelder, M. Verma, D. G. Ang, J. Chow, Z. Lasner, X. Wu, D. DeMille, J. M. Doyle, and G. Gabrielse, “Attaining the shot-noise-limit in the ACME measurement of the electron electric dipole moment,” *Journal of Physics B: Atomic, Molecular and Optical Physics* **52**, 235003 (2019).
- [70] D. G. Ang, C. Meisenhelder, C. D. Panda, X. Wu, D. DeMille, J. M. Doyle, and G. Gabrielse, “Measurement of the H³ Δ_1 radiative lifetime in ThO,” *Physical Review A* **106**, 022808 (2022).

- [71] T. Masuda, D. G. Ang, N. R. Hutzler, C. Meisenhelder, N. Sasao, S. Uetake, X. Wu, D. DeMille, G. Gabrielse, J. M. Doyle, and K. Yoshimura, “Suppression of the optical crosstalk in a multi-channel silicon photomultiplier array,” *Optics Express* **29**, 16914 (2021).
- [72] T. Masuda, A. Hiramoto, D. G. Ang, C. Meisenhelder, C. D. Panda, N. Sasao, S. Uetake, X. Wu, D. P. DeMille, J. M. Doyle, G. Gabrielse, and K. Yoshimura, “High-sensitivity low-noise photodetector using a large-area silicon photomultiplier,” *Optics Express* **31**, 1943 (2023).
- [73] A. Hiramoto, T. Masuda, D. Ang, C. Meisenhelder, C. Panda, N. Sasao, S. Uetake, X. Wu, D. Demille, J. Doyle, G. Gabrielse, and K. Yoshimura, “SiPM module for the ACME III electron EDM search,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **1045**, 167513 (2023).
- [74] A. N. Petrov, L. V. Skripnikov, A. V. Titov, N. R. Hutzler, P. W. Hess, B. R. O’Leary, B. Spaun, D. DeMille, G. Gabrielse, and J. M. Doyle, “Zeeman interaction in $\text{ThO } H^3\Delta_1$ for the electron electric-dipole-moment search,” *Physical Review A* **89**, 062505 (2014).
- [75] A. C. Vutha, B. Spaun, Y. V. Gurevich, N. R. Hutzler, E. Kirilov, J. M. Doyle, G. Gabrielse, and D. DeMille, “Magnetic and electric dipole moments of the $H^3\Delta_1$ state in ThO ,” *Physical Review A* **84**, 034502 (2011).
- [76] G. Herzberg, *The Spectra and Structures of Simple Free Radicals: An Introduction to Molecular Spectroscopy* (Dover Publications, New York, 1988).
- [77] J. M. Brown and A. Carrington, *Rotational Spectroscopy of Diatomic Molecules*, Cambridge Molecular Science Series (Cambridge University Press, Cambridge ; New York, 2003).
- [78] A. C. Vutha, W. C. Campbell, Y. V. Gurevich, N. R. Hutzler, M. Parsons, D. Patterson, E. Petrik, B. Spaun, J. M. Doyle, G. Gabrielse, and D. DeMille, “Search for the electric dipole moment of the electron with thorium monoxide,” *Journal of Physics B: Atomic, Molecular and Optical Physics* **43**, 074007 (2010).
- [79] N. R. Hutzler, *A New Limit on the Electron Electric Dipole Moment: Beam Production, Data Interpretation, and Systematics*, Ph.D. thesis, Harvard University (2014).
- [80] P. W. Hess, *Improving the Limit on the Electron EDM: Data Acquisition and Systematics Studies in the ACME Experiment*, Ph.D. thesis, Harvard University (2014).
- [81] S. Bickman, P. Hamilton, Y. Jiang, and D. DeMille, “Preparation and detection of states with simultaneous spin alignment and selectable molecular orientation in PbO ,” *Physical Review A* **80**, 023418 (2009).

- [82] E. Kirilov, W. C. Campbell, J. M. Doyle, G. Gabrielse, Y. V. Gurevich, P. W. Hess, N. R. Hutzler, B. R. O’Leary, E. Petrik, B. Spaun, A. C. Vutha, and D. DeMille, “Shot-noise-limited spin measurements in a pulsed molecular beam,” *Physical Review A* **88**, 013844 (2013).
- [83] N. R. Hutzler, H.-I. Lu, and J. M. Doyle, “The Buffer Gas Beam: An Intense, Cold, and Slow Source for Atoms and Molecules,” *Chemical Reviews* **112**, 4803–4827 (2012).
- [84] Z. Lasner, *Order-of-Magnitude-Tighter Bound on the Electron Electric Dipole Moment*, Ph.D. thesis, Yale University (2019).
- [85] X. Wu, P. Hu, Z. Han, D. G. Ang, C. Meisenhelder, G. Gabrielse, J. M. Doyle, and D. DeMille, “Electrostatic focusing of cold and heavy molecules for the ACME electron EDM search,” *New Journal of Physics* **24**, 073043 (2022).
- [86] C. D. Panda, B. R. O’Leary, A. D. West, J. Baron, P. W. Hess, C. Hoffman, E. Kirilov, C. B. Overstreet, E. P. West, D. DeMille, J. M. Doyle, and G. Gabrielse, “Stimulated Raman adiabatic passage preparation of a coherent superposition of $\text{ThO } H^3\Delta_1$ states for an improved electron electric-dipole-moment measurement,” *Physical Review A* **93**, 052110 (2016).
- [87] X. Wu, Z. Han, J. Chow, D. G. Ang, C. Meisenhelder, C. D. Panda, E. P. West, G. Gabrielse, J. M. Doyle, and D. DeMille, “The metastable $Q^3\Delta_2$ state of ThO: A new resource for the ACME electron EDM search,” *New Journal of Physics* **22**, 023013 (2020).
- [88] C. D. Panda, *Order of Magnitude Improved Limit on the Electric Dipole Moment of the Electron*, Ph.D. thesis, Harvard University (2018).
- [89] D. L. Kokkin, T. C. Steimle, and D. DeMille, “Branching ratios and radiative lifetimes of the U, L, and I states of thorium oxide,” *Physical Review A* **90**, 062503 (2014).
- [90] D. L. Kokkin, T. C. Steimle, and D. DeMille, “Characterization of the I ($|\Omega|=1$)- $X^1\Sigma^+(0,0)$ band of thorium oxide,” *Physical Review A* **91**, 042508 (2015).
- [91] B. R. O’Leary, *In Search of the Electron’s Electric Dipole Moment in Thorium Monoxide: An Improved Upper Limit, Systematic Error Models, and Apparatus Upgrades*, Ph.D. thesis, Yale University (2016).
- [92] Y. Au, *Inelastic Collisions of Atomic Thorium and Molecular Thorium Monoxide with Cold Helium-3*, Ph.D. thesis, Harvard University (2013).
- [93] B. N. Spaun, *A Ten-Fold Improvement to the Limit of the Electron Electric Dipole Moment*, Ph.D. thesis, Harvard University (2014).
- [94] W. Koechner, “Thermal Lensing in a Nd:YAG Laser Rod,” *Applied Optics* **9**, 2548 (1970).

- [95] N. Solmeyer, K. Zhu, and D. S. Weiss, “Note: Mounting ultra-high vacuum windows with low stress-induced birefringence,” *Review of Scientific Instruments* **82**, 066105 (2011).
- [96] S. Eisenbach, H. Lotem, Z. Horvitz, G. Miron, M. Lando, and S. Gabay, “Thermally induced window birefringence in high-power copper vapor laser,” in *8th Meeting in Israel on Optical Engineering*, edited by M. Oron, I. Shladov, and Y. Weissman (Tel Aviv, Israel, 1993) pp. 13–26.
- [97] J. W. Dally and W. F. Riley, *Experimental Stress Analysis*, 3rd ed. (McGraw-Hill, New York, 1991).
- [98] J. R. Barber, *Elasticity*, 3rd ed., Solid Mechanics and Its Applications No. v. 172 (Springer, Dordrecht ; New York, 2010).
- [99] “SCHOTT BOROFLOAT® floated borosilicate glass,” <https://www.schott.com/en-us/products/borofloat-p1000314> (2022).
- [100] “HPFS_Product_Brochure_All_Grades_2015_07_21.pdf,” https://www.corning.com/media/worldwide/csm/documents/-HPFS_Product_Brochure_All_Grades_2015_07_21.pdf (2015).
- [101] V. Wirthl, *Polarimetry on the Advanced Cold Molecule Electron Electric Dipole Moment Experiment*, Master’s thesis, Technische Universität München (2016).
- [102] E. P. West, *A Thermochemical Cryogenic Buffer Gas Beam Source of TbO for Measuring the Electric Dipole Moment of the Electron*, Ph.D. thesis, Harvard Univeristy (2017).
- [103] A. D. West, Z. Lasner, D. DeMille, E. P. West, C. D. Panda, J. M. Doyle, G. Gabrielse, A. Kryskow, and C. Mitchell, “An Underappreciated Radiation Hazard from High Voltage Electrodes in Vacuum,” *Health Physics* **112**, 33–41 (2017).
- [104] Y. V. Gurevich, *Preliminary Measurements for an Electron EDM Experiment in TbO*, Ph.D. thesis, Harvard Univeristy (2012).
- [105] D. Farkas, *An Optical Reference and Frequency Comb for Improved Spectroscopy of Helium*, Ph.D. thesis, Harvard University (2006).
- [106] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, “Laser phase and frequency stabilization using an optical resonator,” *Applied Physics B Photophysics and Laser Chemistry* **31**, 97–105 (1983).
- [107] U. Schünemann, H. Engler, R. Grimm, M. Weidemüller, and M. Zielonkowski, “Simple scheme for tunable frequency offset locking of two lasers,” *Review of Scientific Instruments* **70**, 242–243 (1999).

- [108] “Magnetic Shielding Materials - Amuneal: Magnetic Shielding & Custom Fabrication,” <https://www.amuneal.com/magnetic-shielding/magnetic-shielding-materials-2> (2022).
- [109] D. M. Kara, I. J. Smallman, J. J. Hudson, B. E. Sauer, M. R. Tarbutt, and E. A. Hinds, “Measurement of the electron’s electric dipole moment using YbF molecules: Methods and data analysis,” *New Journal of Physics* **14**, 103051 (2012).
- [110] P. J. Huber, “Robust Estimation of a Location Parameter,” *The Annals of Mathematical Statistics* **35**, 73–101 (1964).
- [111] B. Efron, “Bootstrap Methods: Another Look at the Jackknife,” *The Annals of Statistics* **7** (1979), 10.1214/aos/1176344552.
- [112] B. Efron and R. Tibshirani, “[Bootstrap Methods for Standard Errors, Confidence Intervals, and Other Measures of Statistical Accuracy]: Rejoinder,” *Statistical Science* **1** (1986), 10.1214/ss/1177013817.
- [113] G. J. Feldman and R. D. Cousins, “Unified approach to the classical statistical analysis of small signals,” *Physical Review D* **57**, 3873–3889 (1998).
- [114] L. Anderegg, *Ultracold Molecules in Optical Arrays: From Laser Cooling to Molecular Collisions*, Ph.D. thesis, Harvard University (2019).
- [115] “R7600U_TPMH1317E.pdf,” https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/etd/R7600U_TPMH1317E.pdf (2019).
- [116] W.-S. Sul, C.-H. Lee, and G.-S. Cho, “Influence of Guard-Ring Structure on the Dark Count Rates of Silicon Photomultipliers,” *IEEE Electron Device Letters* **34**, 336–338 (2013).
- [117] M. Hampel, A. Fuster, C. Varela, M. Platino, A. Almela, A. Lucero, B. Wundheiler, and A. Etchegoyen, “Optical crosstalk in SiPMs,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **976**, 164262 (2020).
- [118] S. Vinogradov, T. Vinogradova, V. Shubin, D. Shushakov, and K. Sitarsky, “Probability distribution and noise factor of solid state photomultiplier signals with cross-talk and afterpulsing,” in *2009 IEEE Nuclear Science Symposium Conference Record (NSS/MIC)* (2009) pp. 1496–1500.
- [119] M. Teich, K. Matsuo, and B. Saleh, “Excess noise factors for conventional and superlattice avalanche photodiodes and photomultiplier tubes,” *IEEE Journal of Quantum Electronics* **22**, 1184–1193 (1986).

- [120] “S13361-3050_series_kapd1054e.pdf,” https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/ssd/s13361-3050_series_kapd1054e.pdf (2022).
- [121] “Mppc_kapd9008e.pdf,” https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/ssd/mppc_kapd9008e.pdf (2021).
- [122] F. Robben, “Noise in the Measurement of Light with Photomultipliers,” *Applied Optics* **10**, 776 (1971).
- [123] D. G. Ang, *Progress towards an Improved Measurement of the Electric Dipole Moment of the Electron*, Ph.D. thesis, Harvard University (2023).
- [124] C. Degenhardt, T. Nazarova, C. Lisdat, H. Stoehr, U. Sterr, and F. Riehle, “Influence of chirped excitation pulses in an optical clock with ultracold calcium atoms,” *IEEE Transactions on Instrumentation and Measurement* **54**, 771–775 (2005).
- [125] S. Falke, M. Misera, U. Sterr, and C. Lisdat, “Delivering pulsed and phase stable light to atoms of an optical clock,” *Applied Physics B* **107**, 301–311 (2012).
- [126] “Digital Delay/Pulse Generator - DG645,” <https://www.thinksrs.com/products/dg645.html> (2023).
- [127] “PXIe-5171,” <https://www.ni.com/en-us/support/model.pxie-5171.html> (2023).
- [128] Z. Lasner and D. DeMille, “Statistical sensitivity of phase measurements via laser-induced fluorescence with optical cycling detection,” *Physical Review A* **98**, 053823 (2018).
- [129] D. Xiao, D. M. Lancaster, C. H. Allen, M. J. Taylor, T. A. Lancaster, G. Shaw, N. R. Hutzler, and J. D. Weinstein, “Shaped nozzles for cryogenic buffer-gas beam sources,” *Physical Review A* **99**, 013603 (2019).
- [130] N. A. Robertson, J. R. Blackwood, S. Buchman, R. L. Byer, J. Camp, D. Gill, J. Hanson, S. Williams, and P. Zhou, “Kelvin probe measurements: Investigations of the patch effect with applications to ST-7 and LISA,” *Classical and Quantum Gravity* **23**, 2665–2680 (2006).
- [131] V. Wirthl, C. D. Panda, P. W. Hess, B. Spaun, and G. Gabrielse, “A self-calibrating polarimeter to measure Stokes parameters,” (2021), [arxiv:1703.00963](https://arxiv.org/abs/1703.00963) [physics].
- [132] “SF57HTULTRA | SCHOTT Advanced Optics,” <https://www.schott.com/shop/advanced-optics/en/Optical-Glass/SF57HTULTRA/c/glass-SF57HTULTRA> (2023).

- [133] J. Zhang, X. Zhang, D. Wu, X. Tian, M. Li, and H. Yu, "Investigating on stress-induced birefringence in vacuum windows of plasma-electrode Pockels cell," in *2009 Conference on Lasers & Electro Optics & The Pacific Rim Conference on Lasers and Electro-Optics* (2009) pp. 1–2.
- [134] R. W. Anderson, "A large, inexpensive, reliable rotary vacuum seal," *Review of Scientific Instruments* **45**, 122–122 (1974).
- [135] "Schott-tie-27-stress-in-optical-glass-eng.pdf," <https://www.schott.com/shop/medias/schott-tie-27-stress-in-optical-glass-eng.pdf> (2004).
- [136] "CP-1-brochure-web-spread.pdf," <https://store-w4rbnih33r.mybigcommerce.com/content/CP-1-brochure-web-spread.pdf> (2014).
- [137] T. J. Sumner, J. M. Pendlebury, and K. F. Smith, "Conventional magnetic shielding," *Journal of Physics D: Applied Physics* **20**, 1095–1101 (1987).
- [138] S. Celozzi, R. Araneo, and G. Lovat, *Electromagnetic Shielding*, Wiley Series in Microwave and Optical Engineering (Wiley-Interscience : IEEE Press, Hoboken, N.J, 2008).
- [139] T. Rikitake, *Magnetic and Electromagnetic Shielding* (Terra Scientific Pub. Co. ; D. Reidel Pub. Co. ; Sold and distributed in the U.S.A. and Canada by Kluwer Academic Publishers, Tokyo : Dordrecht ; Boston : Hingham, MA, U.S.A, 1987).
- [140] E. Paperno and I. Sasada, "Magnetic Circuit Approach to Magnetic Shielding," *Journal of the Magnetics Society of Japan* **24**, 40–44 (2000).
- [141] S. Dickerson, J. M. Hogan, D. M. S. Johnson, T. Kovachy, A. Sugarbaker, S.-w. Chiow, and M. A. Kasevich, "A high-performance magnetic shield with large length-to-diameter ratio," *Review of Scientific Instruments* **83**, 065108 (2012).
- [142] S. Malkowski, R. Y. Adhikari, J. Boissevain, C. Daurer, B. W. Filippone, B. Hona, B. Plaster, D. Woods, and H. Yan, "Overlap Technique for End-Cap Seals on Cylindrical Magnetic Shields," *IEEE Transactions on Magnetics* **49**, 651–653 (2013).
- [143] L. Hasselgren and J. Luomi, "Geometrical aspects of magnetic shielding at extremely low frequencies," *IEEE Transactions on Electromagnetic Compatibility* **37**, 409–420 (1995).
- [144] I. Altarev, E. Babcock, D. Beck, M. Burghoff, S. Chesnevskaya, T. Chupp, S. Degenkolb, I. Fan, P. Fierlinger, A. Frei, E. Gutmiedl, S. Knappe-Grüneberg, F. Kuchler, T. Lauer, P. Link, T. Lins, M. Marino, J. McAndrew, B. Niessen, S. Paul, G. Petzoldt, U. Schläpfer, A. Schnabel, S. Sharma, J. Singh, R. Stoepler, S. Stuibler, M. Sturm, B. Taubenheim, L. Trahms, J. Voigt, and T. Zechlau, "A magnetically shielded room with ultra low residual field and gradient," *Review of Scientific Instruments* **85**, 075106 (2014).

- [145] F. Thiel, A. Schnabel, S. Knappe-Grüneberg, D. Stollfuß, and M. Burghoff, “Demagnetization of magnetically shielded rooms,” *Review of Scientific Instruments* **78**, 035106 (2007).
- [146] I. Altarev, P. Fierlinger, T. Lins, M. G. Marino, B. Nießen, G. Petzoldt, M. Reisner, S. Stuber, M. Sturm, J. Taggart Singh, B. Taubenheim, H. K. Rohrer, and U. Schlöpfer, “Minimizing magnetic fields for precision experiments,” *Journal of Applied Physics* **117**, 233903 (2015).
- [147] A. C. Vutha, *A Search for the Electric Dipole Moment of the Electron Using Thorium Monoxide*, Ph.D. thesis, Yale University (2011).
- [148] G. Gabrielse and J. Tan, “Self-shielding superconducting solenoid systems,” *Journal of Applied Physics* **63**, 5143–5148 (1988).
- [149] “QZFM Gen-3 – QuSpin,” <https://quspin.com/products-qzfm/> (2023).
- [150] V. Shah, J. Osborne, J. Orton, and O. Alem, “Fully integrated, standalone zero field optically pumped magnetometer for biomagnetism,” in *Steep Dispersion Engineering and Opto-Atomic Precision Metrology XI*, edited by S. M. Shahriar and J. Scheuer (SPIE, San Francisco, United States, 2018) p. 51.
- [151] “Introduction to Zero-Field Magnetometer – QuSpin,” <http://quspin.com/products-qzfm-gen2-arxiv/zero-field-magnetometer-description/> (2023).
- [152] “QZFM Dimensions – QuSpin,” <http://quspin.com/products-qzfm/qzfm-dimensions/> (2023).
- [153] “Twinleaf — VMR,” <https://twinleaf.com/vector/VMR/> (2023).
- [154] R. J. Roark, W. C. Young, and R. G. Budynas, *Roark’s Formulas for Stress and Strain*, 7th ed. (McGraw-Hill, New York, 2002).
- [155] S. Dushman, *Scientific Foundations of Vacuum Technique*, 2nd ed. (Wiley, New York, 1962).
- [156] “Saturation in interation region | Graduate - 2008 - Ben Spaun,” (2009).
- [157] “June 4 spec run: Good optical pumping signal, 690 nm saturation measurements, and looking for E state | Graduate - 2008 - Ben Spaun,” (2009).
- [158] “Saturation of Optical pumping. | Graduate - 2008 - Ben Spaun,” (2009).
- [159] “I_sat measurements with error-bar (787 nm, 690 nm) | Graduate - 2008 - Ben Spaun,” (2009)

THIS THESIS WAS TYPESET using L^AT_EX, originally developed by Leslie Lamport and based on Donald Knuth's T_EX. The body text is set in 11 point Egenolff-Berner Garamond, a revival of Claude Garamont's humanist typeface. A template that can be used to format a PhD thesis with this look and feel has been released under the permissive MIT (X11) license, and can be found online at github.com/suchow/Dissertate or from its author, Jordan Suchow, at suchow@post.harvard.edu.