HARVARD UNIVERSITY
Graduate School of Arts and Sciences

DISSERTATION ACCEPTANCE CERTIFICATE

The undersigned, appointed by the

Department of Physics

have examined a dissertation entitled

Progress towards an improved measurement of the electric dipole moment of the electron

presented by Daniel Gordon Ang

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

Signature _____

Typed name: Professor Gerald Gabrielse, Chair

Signature _____

Typed name: Professor Roxanne Guenette (Univ. of Manchester)

Signature _____

Typed name: Professor Isaac Silvera

Signature _____

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

Date: April 11, 2023

Progress towards an improved measurement of the electric dipole moment of the electron

A DISSERTATION PRESENTED
BY
DANIEL GORDON ANG
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
APRIL 2023

Progress towards an improved measurement of the electric dipole moment of the electron

Abstract

The Standard Model (SM) of particle physics is one of the most successful scientific theories, but is unable to explain fundamental features of the universe such as baryon asymmetry. The electron electric dipole moment (eEDM) is a powerful probe of physics beyond the SM. The ACME experiment seeks to measure the eEDM by performing spin precession on a beam of thorium monoxide molecules. In 2018, the ACME II experiment set an upper limit on the eEDM of $|d_{\mathrm{e}}| < 1.1 \times 10^{-29}\ e \cdot \mathrm{cm}$, which has stood as the most stringent published limit for the last five years. Since then, the next generation ACME III experiment has been developed with the goal of improving upon this measurement by another order of magnitude. Three major sensitivity upgrades (increased spin precession time, an electrostatic molecular lens, and improved photon detection) have been demonstrated, and all known sources of systematic uncertainty will be reduced to well below the projected statistical sensitivity. This work describes the ACME II measurement, the ACME III apparatus, a lifetime measurement of the EDM-sensitive H-state which allows an increase in spin precession time, upgrades to the photon detection and data acquisition systems, and improved control of magnetic fields in the experiment. With these advances, ACME III is well-poised to accomplish its goal of probing new physics at the scales of tens of TeV, comparable to the reach of the Large Hadron Collider.

Contents

Publications

Parts of this dissertation have been previously reported on and adapted from the following papers:

1. The ACME Collaboration: V. Andreev, **D.G. Ang**, D. DeMille, J.M. Doyle, J. Haefner, N.R. Hutzler, Z. Lasner, C. Meisenhelder, B.R. O'Leary, C.D. Panda, A.D. West, E.P. West, X. Wu, "Improved limit on the electric dipole moment of the electron," Nature **562**, 355-360 (2018)

2. C.D. Panda, C. Meisenhelder, M. Verma, **D.G. Ang**, J. Chow, Z. Lasner, X. Wu, D. DeMille, J.M. Doyle, G. Gabrielse, "Attaining the shot-noise-limit in the ACME measurement of the electron electric dipole moment", Journal of Physics B **52**, 235003 (2019).

3. X. Wu, Z. Han, J. Chow, **D.G. Ang**, C. Meisenhelder, C.D. Panda, E. West, G. Gabrielse, J.M. Doyle, 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).

4. T. Masuda, **D.G. Ang**, N. R. Hutzler, C. Meisenhelder, N. Sasao, S. Uetake, X. Wu, D. DeMille, G. Gabrielse, J.M. Doyle, K. Yoshimura, "Suppression of the optical crosstalk in a multi-channel silicon photomultiplier array", Optics Express **29**(11), 16914-16926 (2021).

5. X. Wu, P. Hu, Z. Han, **D.G. Ang**, C. Meisenhelder, G. Gabrielse, J.M. Doyle, D. DeMille, "Electrostatic focusing of cold and heavy molecules for the ACME electron EDM search", New Journal of Physics **24**, 073043 (2022).

6. **D.G. Ang**, C. Meisenhelder, C. Panda, X. Wu, D. DeMille, J. Doyle, G. Gabrielse, "Measurement of the $H^3\Delta_1$ Radiative Lifetime in ThO", Physical Review A **106**, 022808 (2022).

7. A. Hiramoto, T. Masuda, **D.G. Ang**, C. Meisenhelder, C. Panda, N. Sasao, S. Uetake, X. Wu, D. DeMille, J.M. Doyle, G. Gabrielse, K. Yoshimura, "SiPM module for the ACME III electron EDM search", Nuclear Instruments and Methods in Physics Research A **1045**, 167513 (2023).

8. 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, K. Yoshimura, "High-sensitivity low-noise photodetector using large-area silicon photomultiplier", Optics Express **31**(2), 1943-1957 (2023).

List of Figures

List of Tables

Acknowledgments

I would like to thank my advisor Professor Gerald Gabrielse for giving me the opportunity to work in his lab. His guidance, advice, and encouragement have been invaluable in completing the projects in this thesis and developing my scientific skills. His dedication towards difficult experiments that seek to uncover the most fundamental secrets of nature have also been a source of inspiration towards my own career. I truly seek to live by one of his personal utterances, "God decides, we measure."[1]

I am grateful to Professors John Doyle and David DeMille, the two other co-PIs of the ACME experiment whose technical advice and expertise have been invaluable throughout this project. I would also like to thank the other members of my committee, Professors Ronald Walsworth, Roxanne Guenette, and Isaac Silvera, for their helpful professional advice on this thesis and throughout my PhD career.

I would like to thank Cristian Panda, who patiently mentored me in my first few years as a graduate student. I am grateful to past and current members of the ACME experiment who have been a delight and privilege to work with and learn from, especially Cole Meisenhelder, Xing Wu, Brendon O'Leary, Zack Lasner, Vitaly Wirthl, Zhen Han, Peiran Hu, Maya Watts, Collin Diver, Takahiko Masuda, Ayami Hiramoto, and Xing Fan. I would also like to thank members of the Gabrielse lab during my time who contributed to a friendly atmosphere of scientific camaraderie: Tharon Morrison, Andra Ionescu, Nathaniel

[1] Gerald Gabrielse, plenary talk from the 2019 Annual Meeting of the American Scientific Affiliation at Wheaton College, 21 July 2019, Wheaton, Illinois.

McDonough, Mason Marshall, Daniel Zambrano, Elise Novitski, and Geev Nahal.

I am grateful to the Harvard and Northwestern physics departments for their administrative and logistical support – Laura Nevins, Pattee McGarry, Silke Exner, and Lisa Cacciabaudo, as well as Stan Cotreau for his teaching and technical assistance in making many important components of the experiment.

I would like to thank Professor Larry Hunter, my former undergraduate advisor at Amherst College, who introduced me to the art of precision measurement and has remained a dear friend, mentor, and source of inspiration.

This work has been supported by several important funding sources: the National Science Foundation, the Alfred P. Sloan Foundation, the Gordon and Betty Moore Foundation, Okayama University, and the Japan Society for the Promotion of Science. Additionally, I would like to acknowledge the generous contributions of Amherst College and the Mustard Seed Foundation in supporting my graduate studies.

I would like to give my highest thanks to my family: my parents Gideon and Yohana Ang and my brother Joseph, whose love, care, and companionship have formed and placed me in my current trajectory. I am blessed to have been able to be part of the Harvard Graduate Christian Fellowship and the Dudley World Music Ensemble, especially throughout the early years of my PhD. I am also grateful to my church family at Trinity Cambridge Church, thanks to whom I have undergone immense spiritual growth and nourishment. Finally, I would like to thank my wife Ziwei Zhang, whose support, love, and encouragement has been indispensable. *Soli Deo gloria.*

1

Background

The last century has seen immense progress in our understanding of the fundamental particles and forces. This was witnessed in the development of the Standard Model (SM) of particle physics in the latter half of the 20th century. The success of the SM has been exemplified in triumphs such as the discovery of the Higgs Boson at the Large Hadron Collider (LHC) [1, 2] and the experimental verification of its predicted value of the electron magnetic moment to a part in a trillion [3, 4]. However, fundamental puzzles remain. The SM is unable to explain matter-antimatter asymmetry [5], dark matter [6], or dark energy [7]. Without matter-antimatter asymmetry, there would be no stars, planets, or physicists, while dark matter and dark energy together comprise ∼95% of the matter-energy in the

universe. Thus, despite its successes, the SM effectively only accounts for a small slice of the material universe.

In response, new theories "beyond the SM" (BSM) have been proposed to explain these and other puzzles in particle physics such as the hierarchy problem, the strong CP problem, and the nature of quantum gravity [8, 9]. Many of these theories are being tested directly at large-scale experiments such as the LHC. However, many also predict the existence of permanent electric dipole moments (EDMs) of elementary particles, including the electron. This is the key motivation for our ACME experiment, which seeks to measure the electron EDM using a beam of thorium monoxide (ThO). The last decade has seen the immense potential of such smaller scale tabletop experiments to probe for new physics beyond the Standard Model.

1.1 The structure of this dissertation

This dissertation primarily reports my contributions to the ACME experiment from 2015-2022. It comprises of two parts: first, the ACME II experiment, which in 2018 improved the upper limit on the electron EDM by an order of magnitude, and second, the design and construction of the next generation ACME III apparatus, projected to improve upon the sensitivity of ACME II by another order of magnitude. Much of this work would not have been possible without collaboration with other members of the ACME collaboration. While most attention will be given to the projects in which I played a leading role, brief summaries of the progress accomplished on other components of the experiment will be given in order to provide a broader overview of the experimental effort.

In more detail:

- Chapter 1 provides the theoretical motivations for measuring the electron EDM and the basics of the measurement method used in the ACME experiment.
- Chapter 2 describes the ACME II measurement, including both the components of

the experimental apparatus and the systematic errors which were encountered. My main contributions to this measurement were in the areas of experimental control, data acquisition systems, data analysis, and maintenance of lasers.

- Chapter 3 is an overview of the new ACME III experiment now being commissioned.
- The next four chapters detail projects to design and improve components of the ACME III experiment in which I was deeply involved in. The measurement of the EDM-sensitive state lifetime is described in Chapter 4, photon detection in Chapter 5, data acquisition and experimental control in Chapter 6, and magnetic fields in Chapter 7.
- Chapter 8 summarizes the overall progress of the ACME experiment so far, as well as the future prospects for continuing to increase our sensitivity in measuring EDMs.

1.2 EDMs and and the search for new physics

1.2.1 Introduction to EDMs

In classical physics, *electric multipole moments* arise when the electric potential V of an arbitrary distribution of charge $\rho(\vec{r})$ is expressed as a sum of electric potentials of elementary combinations of single charges, such as monopoles, dipoles, quadrupoles, and so on, also known as the *multipole expansion* [10, ch. 4]. The dipole term is

$$V_{\text{dipole}} = \frac{\vec{d} \cdot \hat{r}}{r^2}, \tag{1.1}$$

where $\vec{d}$ is the *electric dipole moment* (EDM), defined as

$$\vec{d} = \int \rho(\vec{r}) \vec{r} d^3 r. \tag{1.2}$$

An EDM can thus be approximately understood as the degree to which V resembles the potential of a pair of positive and negative unit charges separated by a distance s, where we take the limit $s \to 0$. For the magnitude of $\vec{d}$ in Eq. 1.2 to be independent of coordinate system, the origin is commonly taken to be at the center of mass. $d \equiv |\vec{d}|$ can then be interpreted as the separation between the centers of charge and mass of the system. Its Hamiltonian contribution in the presence of an electric field $\vec{\mathcal{E}}$ is

$$H_{\text{EDM}} = -\vec{d} \cdot \vec{\mathcal{E}}, \tag{1.3}$$

where d is commonly given the units $e \cdot \text{cm}$, where e is the elementary charge.

It is possible for composite systems such as polar molecules and atoms to have an induced EDM in the presence of an electric field, which underlies the well-known linear Stark mixing of energy levels with opposite parity [11, p. 254]. Much more interesting is the possibility of *permanent* EDMs in fundamental particles and nucleons such as protons, electrons, neutrons, and muons, a search for which was first suggested by Purcell and Ramsey [12]. For the case of electrons, which are considered point-sized elementary particles within the SM, an electron EDM (eEDM) d_{e} would be an intrinsic property rather than a spatially extended distribution of charge. Because the electron has no known substructure, $\vec{d}_{\text{e}}$ must be parallel to the electron's spin $\vec{S}$, a consequence of the Wigner-Eckart theorem [11, p. 256]. Thus Eq. 1.3 becomes

$$H_{\text{EDM}} = -\vec{d}_{\text{e}} \cdot \vec{\mathcal{E}} = -d_{\text{e}} \vec{S} \cdot \vec{\mathcal{E}} \tag{1.4}$$

If such an EDM exists, it would be an instance of time reversal (T) symmetry violation, which makes the question of immense interest to particle physicists, as will be explained in the next section.

1.2.2 EDMs and discrete symmetries

A physical system is said to have a symmetry when its behavior remains the same after a certain transformation is applied [13, ch. 1]. Examples include spatial and time translation symmetries, both of which are *continuous* symmetries. A *discrete* symmetry is one which returns the same physical system after being applied twice. Three fundamental discrete transformations underlie the Standard Model:

1. **Charge conjugation (C)**: exchange of each particle for its antiparticle counterpart.
2. **Parity (P)**: inversion of spatial coordinates with respect to the origin, $\vec{x} \to -\vec{x}$.
3. **Time reversal (T)**: inversion of the direction of time, $t \to -t$.

In 1956, Lee and Yang famously suggested the possibility of P violation, which was later observed by Wu and coworkers in the decay of cobalt-60 atoms [14, 15]. In 1964, Cronin, Fitch, and coworkers observed the first evidence of CP violation in neutral kaon decay [16]. CP violation was later also discovered in B mesons [17] and D mesons [18].

Only CPT remains as an unbroken discrete symmetry within the SM. To this day, there is still no experimental evidence of a CPT violation. The preservation of CPT is a consequence of the CPT theorem, a rigorous proof of which assumes Lorentz symmetry [19, 20]. The CPT theorem is a cornerstone of quantum field theory whose violation would have far-reaching consequences [13, ch. 5]. It has been extensively tested in various experiments, such as those comparing the properties of matter and antimatter particles [21–24]. If we assume that the CPT theorem is true, then this means that a CP violation must imply a T violation, and vice versa.

Returning to the case of EDMs, suppose that a fundamental particle such as the electron has a non-zero EDM d_e (Eq. 1.4). If we apply a T transformation to the electron, this would reverse the direction of $\vec{S}$ but not $\vec{d}_\mathrm{e}$, resulting in a different electron than we started with (Fig. 1.2.1). Similarly, when we apply a P transformation, $\vec{d}_\mathrm{e}$ is reversed but not $\vec{S}$.

Thus, a non-zero electron EDM would violate P and T. Due to the CPT theorem, the T violation would also imply a CP violation. Because the existence of EDMs of fundamental particles would be a form of CP violation, they are interesting to theorists who seek to solve fundamental puzzles such as matter-antimatter asymmetry.

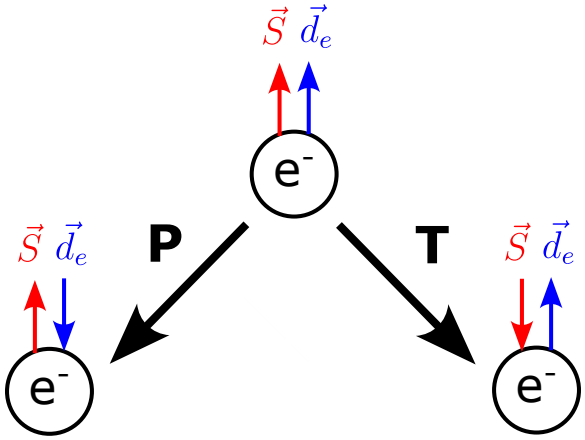

Figure 1.2.1: **EDMs and discrete symmetries.** If an electron EDM exists, then undergoing a P (left) or T (right) transformation will flip only the spin $\vec{S}$ or the EDM $\vec{d}_e$, transforming the particle. Thus, an electron EDM would violate P and T.

1.2.3 Matter-antimatter asymmetry

One of the major motivations for the search for more sources of CP violation is the problem of matter-antimatter asymmetry, also known as baryon asymmetry (see [5] for a review). As far as we can tell, essentially all matter around us is ordinary matter.[1] Searches for signs of antimatter particles in our solar system by cosmic ray detection have not come up with significant amounts [25], and other scenarios such as distant star systems or large regions of the universe being made of antimatter have been ruled out by the lack of observations of annihilation signatures [26]. This imbalance can be parameterized by the baryon

[1] Of course, this is putting aside the equally puzzling problems of the nature of dark matter and dark energy.

asymmetry of the universe (BAU), defined as the ratio [13, §11.3.2]

$$\eta \equiv \frac{\eta_B - \eta_{\vec{B}}}{\eta_\gamma}, \tag{1.5}$$

where η_B, $\eta_{\vec{B}}$, η_γ are the number densities of baryons, antibaryons, and photons respectively. The value of η has been experimentally determined to be $(6.143 \pm 0.190) \times 10^{-10}$, based on observations of the cosmic microwave background and Big Bang Nucleosynthesis (BBN) data [27]. What this means is that there was an excess of roughly one particle of matter for every 10^9 matter-antimatter particle pairs in the early universe, after which all the pairs annihilated, leaving only the remnant excess to constitute all the ordinary matter we encounter today.

To explain the origin of this asymmetry, theories of *baryogenesis* have been proposed. Scenarios in which the BAU originates from the initial conditions of the universe are incompatible with the strong evidence for cosmic inflation [28, 29]. Assuming initially equal amounts of baryons and antibaryons, Sakharov [30] identified three conditions that are required for baryogenesis, namely baryon number violation, C and CP violation, and departure from thermodynamic equilibrium (assuming that CPT is preserved). All three conditions are present in the Standard Model. In particular, CP violation is parameterized in the SM through a complex phase $\delta_{\text{CKM}} = (1.144 \pm 0.0027)$ in the Cabibbo–Kobayashi–Maskawa (CKM) matrix [31]. However, this is an insufficient amount of CP violation to explain the BAU [32–34].

Motivated by this puzzle, beyond-the-SM (BSM) theories which posit additional sources of CP violation have been proposed to explain baryogenesis [35, 36]. Most of these theories predict non-zero values of the EDMs of fundamental particles. In particular, many predict electron EDMs at around the current experimental limit of $10^{-29}\ e \cdot \text{cm}$ [37, 38]. The SM itself predicts a non-zero eEDM value of $d_\text{e} \approx 10^{-35}\ e \cdot \text{cm}$ [39], which is several orders of

magnitude smaller.[2] Thus, any non-zero result in a current eEDM experiment will be an immediate sign of new physics, in contrast to other probes such as the muon g-factor, where careful and technically complex calculations of the SM value are required to interpret any anomalies [41]. Conversely, a null result would be a strong constraint on these theories. As an example, the ACME II bound severely constrains the two-Higgs-doublet model of baryogenesis [36, 42].

1.2.4 Other theoretical motivations for electron EDM experiments

There are broader motivations for BSM theories besides baryogenesis. The existence of supersymmetric (SUSY) particles have been proposed to explain the hierarchy problem [43] and are also attractive dark matter (DM) candidates [44]. However, recent searches at the LHC have so far not found any signs of SUSY particles [45]. Electron EDM searches provide a complementary method to probe such theories. From dimensional analysis, one can make a general estimate of the scale Λ of new physics being probed by an electron EDM measurement [46–49]:

$$\Lambda^2 \sim m_{\mathrm{e}} \frac{e}{d_{\mathrm{e}}} \left(\frac{\alpha_{\mathrm{eff}}}{2\pi} \right)^n \kappa (\sin \phi_{\mathrm{CP}}) (\hbar c), \tag{1.6}$$

where m_{e} is the electron mass, e is the electron charge, α_{eff} encodes the coupling strength of the electron to new particles (with $\alpha_{\mathrm{eff}} = 4/137$ for electroweak interactions), n is the number of loops in the Feynman diagram from which the eEDM arises, ϕ_{CP} is the CP-violating phase (usually assumed to be of order 1), and $\kappa \approx 0.1 - 1$ is a model-dependent factor. Thus, the current best limit on the eEDM implies that $\Lambda \gtrsim 30$ TeV (one-loop processes) or $\Lambda \gtrsim 3$ TeV (two-loop). These figures exceed or are comparable to the scale of new physics being probed in the LHC, although they are obtained from much

[2] As explained in Ref. [39], this the "equivalent eEDM" from the SM contribution to the T-violating electron-nucleon coupling C_S (see §2.3) and is the "effective measurable eEDM" in experiments with paramagnetic systems such as ThO or YbF. The "pure" eEDM value has been calculated to be $d_{\mathrm{e}} \approx 10^{-40}\ e$ [40].

smaller scale tabletop experiments [49–51]. More detailed calculations find that eEDM limits provide strong constraints on various versions of supersymmetry [52–55]. Electron EDM experiments are also capable of probing alternative models of BSM physics beyond supersymmetry [46, 56, 57]. Finally, some EDM experiments can also be used to search for oscillating EDMs predicted by the existence of dark matter axionlike particles (ALPs) [58].

We have thus far focused on electron EDM experiments with paramagnetic molecules as this is the purview of the ACME experiment. Similar tabletop spin precession experiments with diamagnetic species such as TlF [59], ^{199}Hg [60], and ^{129}Xe [61] are sensitive to T-violating nucleon EDMs and nuclear Schiff moments. In addition, there are several large-scale efforts to measure the neutron EDM [62–64]. These efforts are complementary to electron experiments as they test BSM theories in the hadronic sector.

To summarize, there are broad theoretical motivations to look for the electron EDM and other forms of T-violations, and their existence may hold the key to solving some of the most fundamental puzzles in our current understanding of physics. State-of-the-art experiments measuring these quantities have function as background-free probes of new physics with extremely high energy reach. Such experiments will likely become even increasingly important in the next few decades, given the increasing costs and technical difficulties of continuing to scale up large-scale particle accelerators [65].

1.3 Overview of electron EDM measurements

1.3.1 Basic EDM measurement technique

The majority of EDM measurements use a spin precession technique which originated from the separated oscillating fields technique to measure nuclear magnetic moments, also known as Ramsey interferometry [66, 67]. The general scheme features a relatively long spin precession region sandwiched between two regions for state preparation and readout [68]. For simplicity, let us assume a spin-1/2 system, where the two states of the particle ($|\uparrow\rangle$

and $|\downarrow\rangle$) have opposite Zeeman and EDM shifts. Static electric ($\vec{\mathcal{E}}$) and magnetic ($\vec{\mathcal{B}}$) fields are applied along $\hat{z}$ throughout the entire apparatus. In the preparation region the spin of the particle(s) is prepared along $\hat{x}$ using RF pulses or lasers:

$$|\psi\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle + |\downarrow\rangle), \tag{1.7}$$

The particles proceed to the precession region where their magnetic moment $\vec{\mu}$ interacts with $\vec{\mathcal{B}}$ to result in precession in the xy plane. Due to the opposite Zeeman shifts, the two components will acquire phases with opposite sign, and similarly with the EDM term $\vec{d}\cdot\vec{\mathcal{E}}$ (Eq. 1.4). The particles reach the end of the spin precession after a certain time τ, at which point we have

$$|\psi\rangle = \frac{1}{\sqrt{2}}(e^{-i\phi_1}|\uparrow\rangle + e^{+i\phi_1}|\downarrow\rangle), \tag{1.8}$$

where $\phi_1 = -(\vec{\mu}\cdot\vec{\mathcal{B}}+\vec{d}\cdot\vec{\mathcal{E}})\tau/\hbar$. The accumulated phase is measured at the readout region by applying another pulse that projects ϕ_1 along $+\hat{y}$. The probability of finding the spin along $+\hat{y}$ is $P_+ = \sin^2\phi_1$. Typically, $|\vec{\mathcal{B}}|$ is set to maximize the change in P_+, i.e. $\phi_1 \approx \pi/4 + n\pi$ for integer n.

The experiment is then repeated with the relative directions of $\vec{\mathcal{B}}$ and $\vec{\mathcal{E}}$ reversed, such that $\phi_2 = -(\vec{\mu}\cdot\vec{\mathcal{B}}-\vec{d}\cdot\vec{\mathcal{E}})\tau/\hbar$. We can then extract the EDM term, $\phi^{\mathrm{EDM}} = (\phi_1-\phi_2)/2 = \vec{d}\cdot\vec{\mathcal{E}}$. As these are quantum measurements (i.e. the result being either 0 or 1), assuming there are no other sources of noise, the uncertainty in d will be [19, 69]

$$\delta d = \frac{\hbar}{2\tau\mathcal{E}\sqrt{N}}, \tag{1.9}$$

where N is the number of measurements. This is also known as the *shot-noise limit*.

This fundamental experimental scheme applies to neutron, atomic, and molecular EDM measurements, with only subtle variations [70, §2.2]. Experiments with trapped particles (such as the JILA ion trap eEDM experiment [51]) deliver temporally rather than spatially

separated preparation and readout pulses. To obtain the smallest possible uncertainty, Eq. 1.9 shows the precession time, electric field, and number of measurements must all be maximized. In practice, applied laboratory electric fields are limited to ∼100 kV/cm. τ is commonly limited by the radiative lifetime of the EDM-sensitive state (if it is not a ground state) and the velocity of the particles in a beam.

After Ramsey and Purcell first proposed the possibility of searching for EDMs of fundamental particles and nuclei in 1950 [12], early efforts focused on measuring the neutron EDM using the method discussed above on a beam of free neutrons. In 1957, the first experiment using this method successfully placed an upper bound of $|d_{\mathrm{n}}| < 5 \times 10^{-20}\ e \cdot \mathrm{cm}$ [71]. However, it is not possible to use this method on charged particles such as free electrons, where applying an electric field will result in the particles accelerating out of control. Thus, initial limits on the electron EDM were obtained using other techniques: measurements of the Lamb shift [72], atomic parity [73], electron g-factor [74, 75], and alpha particle scattering [76, 77]. These only managed to place an upper bound of $|d_{\mathrm{e}}| < 10^{-16}\ e \cdot \mathrm{cm}$.

1.3.2 Sandars and the evasion of Schiff's theorem

A significant advance occurred in 1965, when Sandars showed that it is possible to search for the electron EDM by measuring the EDM of a neutral atom [78]. Previously, it was thought that an electron EDM would not produce an atomic EDM, also known as Schiff's theorem [79]. Intuitively, this can be seen in the following argument [80]: In a non-relativistic atom with a point-sized nucleus, only electrostatic forces exist. When a homogeneous electric field is applied to a neutral atom, the atom does not undergo acceleration. Therefore, it must be the case that the atom is polarized in a way such that its internal electrostatic forces cancel out the externally applied electric field. Thus, electrons inside the atom on average experience zero electric field, and an electron EDM would have no effect.

However, Sandars showed that the above argument does not apply for heavy paramag-

netic atoms with a high atomic number Z, whose constituents move non-relativistically.[3] For an electron in an electric field $\boldsymbol{\mathcal{E}}$, the relativistic EDM Hamiltonian is [68, 80]

$$\begin{aligned} H_{\mathrm{EDM}} &= -d_{\mathrm{e}}\gamma^0\boldsymbol{\Sigma}\cdot\boldsymbol{\mathcal{E}} \\ &= -d_{\mathrm{e}}\boldsymbol{\Sigma}\cdot\boldsymbol{\mathcal{E}} - d_{\mathrm{e}}(\gamma^0-1)\boldsymbol{\Sigma}\cdot\boldsymbol{\mathcal{E}}, \end{aligned} \tag{1.10}$$

where γ^0 and $\boldsymbol{\Sigma} = \left(\begin{smallmatrix}\boldsymbol{\sigma} & 0\\ 0 & \boldsymbol{\sigma}\end{smallmatrix}\right)$ are standard Dirac matrices. When H_{EDM} is added as a perturbation to the Dirac Hamiltonian for a single-electron atom in an external electric field, the first term vanishes in the non-relativistic limit, which is Schiff's theorem [80]. The second term remains and contributes an EDM energy shift [81, 82]

$$\begin{aligned} \Delta E_{\mathrm{EDM}} &= \langle\psi|\, d_{\mathrm{e}}(\gamma^0-1)\boldsymbol{\Sigma}\cdot\boldsymbol{\mathcal{E}}\,|\psi\rangle \\ &= \left\langle \psi \middle| \begin{pmatrix} 0 & 0 \\ 0 & -2d_{\mathrm{e}}\boldsymbol{\sigma}\cdot\boldsymbol{\mathcal{E}} \end{pmatrix} \middle| \psi \right\rangle. \end{aligned} \tag{1.11}$$

As the operator above has odd parity, the shift vanishes if $|\psi\rangle$ is a parity eigenstate [83, §1.1.2]. Thus, a laboratory electric field must be applied to mix $|\psi\rangle$ with a state of opposite parity $|\psi'\rangle$ separated by. The polarization P expresses the degree of mixing,

$$P \approx -\mathcal{E}\frac{D}{\Delta E}, \tag{1.12}$$

where D is the transition dipole moment between the two parity doublets, ΔE is the energy difference, and $0 < P < 1$. In this way, the electron EDM value can be obtained by measuring the EDM of an atomic or molecular system. A more intuitive classical explanation is that the electron's EDM undergoes a length contraction when it moves at fast speeds near the nucleus, sidestepping the non-relativistic assumption of Schiff's

[3] Alternatively, the argument also does not apply to diamagnetic atoms with a large nucleus, which drives the search for hadronic EDMs and nuclear Schiff moments in atoms such as ^{199}Hg and TlF.

theorem [80]. This effect is more pronounced for certain heavy paramagnetic atoms, such that the average electric field experienced by the electron will be enhanced compared to the laboratory $\vec{\mathcal{E}}$. The EDM Hamiltonian becomes [68]

$$H_{\mathrm{EDM}} = \vec{d}_{\mathrm{e}} \cdot \vec{\mathcal{E}}_{\mathrm{eff}} = \vec{d}_{\mathrm{e}} \cdot \mathcal{Q} P \vec{\mathcal{E}} \tag{1.13}$$

where $\mathcal{Q}$ is an atom-dependent enhancement factor that roughly scales as Z^3 [84]. This equation favors experiments using atoms with high Z such as cesium ($R = \mathcal{Q}\mathcal{P} \approx 120$), thallium ($R \approx -600$), or francium ($R \approx 1150$) [85–87]. One difficulty of atomic EDM experiments, however, is that the parity doublets are usually electronic levels with $\Delta E \gtrsim 100$ THz, resulting in $P \sim 10^{-3}$. In contrast, molecular EDM experiments use rotational splittings or Ω-doubling splittings which are much smaller, allowing much greater polarization and thus larger $\mathcal{E}_{\mathrm{eff}}$.

1.4 Progress of electron EDM measurements

Figure 1.4.1 shows a plot of the progress in reducing the upper limit of the electron EDM in 70 years of searches.[4] Sandars' aforementioned discovery regarding atomic EDMs opened up a new era of rapid progress. Since then, an improvement of over eight orders of magnitude in precision has been seen. More recently, molecular EDM experiments initiated a remarkably fertile period of progress which has seen two orders of magnitude of improvement within a decade. No eEDM has yet been found. Today, the precision of state-of-the-art electron EDM experiments make them extremely powerful and sensitive probes of new physics. Here, we will briefly discuss some of these experiments including the challenges they encountered from systematic effects.

[4] Not all eEDM experimental results have been shown, but only those which were the best limit (or close to it) when they were published.

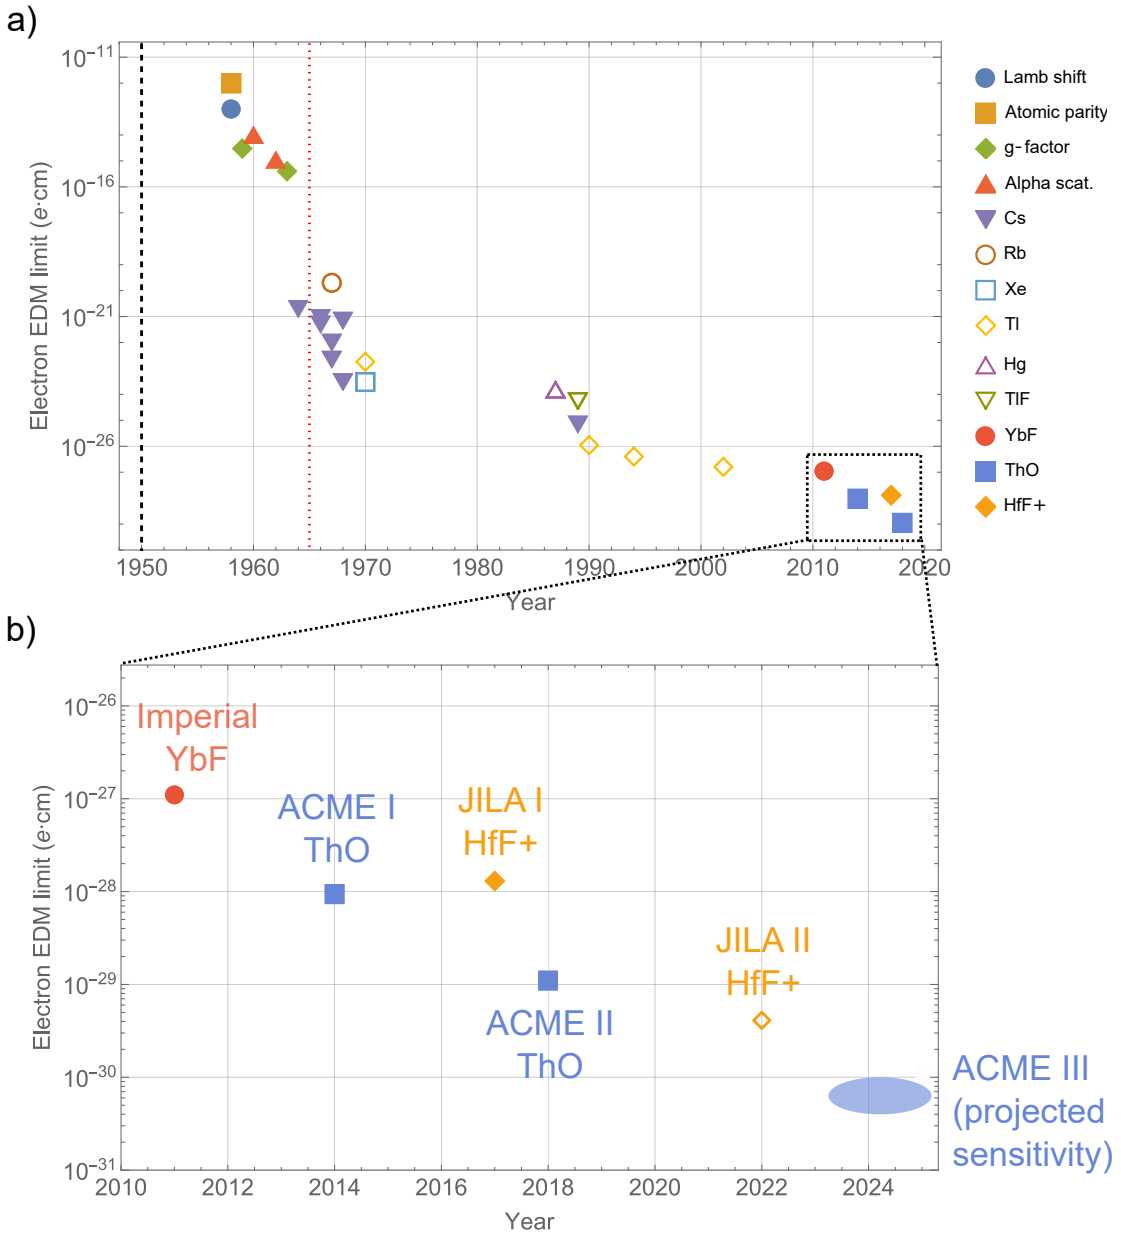

Figure 1.4.1: **Progress in electron EDM measurements**. a) After Purcell and Ramsey [12] initiated the search for EDMs for fundamental particles (vertical dashed black line), initial measurements used a variety of methods [73–77]. Sandars' seminal paper in 1965 [78] (vertical dotted red line) began a new era where atomic EDMs could be used to search for the eEDM, resulting in rapid progress in the 1960s with cesium experiments [88–92], and later in the 1990s with thallium [93–95]. b) The last decade has seen impressive improvement of nearly three orders of magnitude with the use of polar molecules, which have a much larger $\mathcal{E}_{\text{eff}}$ [49–51, 96]. Also shown are the limit in the recent preprint from the JILA II experiment [97] and the projected ACME III statistical sensitivity (Chapter 3).

1.4.1 Atomic electron EDM experiments

The first atomic eEDM measurement was performed by Sandars and coworkers with a beam of cesium atoms, reducing the previous bound by five orders of magnitude, $|d_e| < 2 \times 10^{-21}\ e \cdot \text{cm}$ [98]. The main systematic error came from a motional magnetic field $\vec{\mathcal{B}}_{\text{mot}} = \frac{1}{c}\vec{\mathcal{E}} \times \vec{v}$ arising from the interaction between the atom's magnetic dipole moment as it moves in the electric field $\vec{\mathcal{E}}$ [19, §3.5.3]. If $\vec{\mathcal{E}}$ is perfectly parallel to $\vec{\mathcal{B}}$, this would not be a problem as $\vec{\mathcal{B}}_{\text{mot}}$ is perpendicular to $\vec{\mathcal{B}}$. However, a small misalignment would cause an energy shift that masquerades as an EDM. This became a dominant systematic effect in many atomic beam EDM experiments. One method to suppress it is to build an apparatus with two atomic beams propagating in opposite directions, allowing reversal of the sign of the effect [99]. Another suppression method is to use atoms with high polarizability, such that the quadratic Stark shift is much larger than the Zeeman shift [68]. This second suppression method was first demonstrated in a xenon EDM experiment by Player and Sandars [100]. These and other improvements in Cs experiments through the 1960s led to the reduction of eEDM limit by a further two orders of magnitude [88–92, 99, 100].

Subsequent significant progress only came in the 1990s, starting with a highly sensitive Cs experiment by Hunter et al. [101] conducted in a cell which was not affected by motional magnetic field effects and improved upon the previous limit by over an order of magnitude. The leading systematic effect in the experiment came from leakage currents in the electric field plates. Eventually, the most precise atomic beam experiments were performed by Commins et al. with thallium [93–95]. With $R = -585$ and $\mathcal{E} = 122$ kV/cm, $\mathcal{E}_{\text{eff}} \approx 70$ MV/cm was achieved. In the last iteration, two pairs of counter-propagating atomic beams were built, the second pair being for sodium atoms as a co-magnetometer. Sodium has a different sensitivity to the eEDM compared to thallium, enabling more powerful differentiation of systematic from EDM shifts. This experiment set the best atomic limit on the electron EDM, $|d_e| < 1.6 \times 10^{-27}\ e \cdot \text{cm}$.

1.4.2 Molecular electron EDM experiments

The last decade has seen astonishing progress in the search for the eEDM through the use of molecular EDM measurements. Compared to atomic EDM experiments, heavy paramagnetic polar molecules are able to achieve much better sensitivity to the eEDM because of their much higher $\mathcal{E}_{\text{eff}} \approx 10-100$ GV/cm. This is possible because the molecules have parity doublets that are separated by much smaller splittings and are thus much more easily polarized. As an example, for ThO this splitting is several hundred kHz, $\sim 10^3 - 10^4$ smaller than in typical atoms [68]. However, molecules are often more difficult to produce in large numbers and state preparation and readout are usually more complex, requiring a large number of lasers. Diatomic molecule species which have been used in electron EDM experiments are shown in Table 1.4.1.

Molecular species	$\lvert\mathcal{E}_{\text{eff}}\rvert$ **(GV/cm)**	**eEDM bound (e · cm)**	**Ref.**
YbF	26	1.1×10^{-27}	[50, 102]
ThO	78	1.1×10^{-29}	[49, 103, 104]
HfF^+	23	4.1×10^{-30}	[97, 105]
ThF^+	35	-	[104, 106]
WC	36	-	[107]
PbO	25	1.7×10^{-26}	[108, 109]
BaF	6.5	-	[110, 111]

Table 1.4.1: Diatomic molecular species which have been used in electron EDM experiments, together with their calculated effective electric field values and the latest eEDM upper bound obtained (if applicable).

In 2011, the Imperial College YbF experiment was the first molecule eEDM experiment to exceed the precision of atomic eEDM experiments [50]. Due to the high polarizability of the YbF molecule, $\mathcal{E}_{\text{eff}} = 14.5$ GV/cm was achieved while only applying a laboratory $\mathcal{E} = 10$ kV/cm. This polarizability also suppressed the motional magnetic field effect far below the statistical sensitivity. Instead, the leading systematic effects were related to imperfect $\mathcal{E}$ reversal. The experiment slightly improved the upper bound on the eEDM by

a factor of 1.5 to $|d_e| < 1.05 \times 10^{-27}\ e \cdot \text{cm}$.

The ACME experiment uses the $H^3\Delta_1$ state of ThO to measure the eEDM, which can attain full polarization ($\mathcal{E}_{\text{eff}} = 80$ GV/cm) with a laboratory field of only 10 V/cm. An important property of ThO is the presence of two states where the EDM energy shifts have opposing signs, allowing them to be used as a pair of in-built co-magnetometer states for systematic error identification and rejection. With cryogenic buffer gas beam source technology [112], a large number of usable molecules can be produced in the beam. These and other advantageous features of ThO for an eEDM experiment are reviewed in §1.5.1. The ACME I experiment was completed in 2014 and successfully improved the previous eEDM limit by an order of magnitude [96, 113]. Like the YbF experiment, major systematic effects in ACME have arisen from imperfect $\mathcal{E}$-reversal. The second generation ACME II experiment was completed just four years later, improving on the upper limit by another order of magnitude to $|d_e| < 1.1 \times 10^{-29}\ e \cdot \text{cm}$, which at the time of writing still stands as the best published limit on the electron EDM. A more detailed account of the ACME II experiment is given in Chapter 2.

The JILA eEDM experiment was the first EDM experiment to use trapped molecular ions [51]. The measurement is performed in the metastable $^3\Delta_1$ state of HfF^+, which has several similar advantageous features as ThO, including a reasonably high $\mathcal{E}_{\text{eff}}$ (23 GV/cm), high polarizability (requiring only $\mathcal{E} = 24$ V/cm), and co-magnetometer states. Because the experiment is conducted in an ion trap, very long coherence times (∼1 s) are possible. Furthermore, the fact that the polarizing electric fields and ions are rotating relative to the laboratory is useful in screening out effects such as constant magnetic field offsets. In 2017, the first generation of this experiment trapped ∼1500 ions per shot with $\tau \approx 700$ ms, obtaining an eEDM limit within a factor of 1.4 of ACME I. The second generation JILA experiment has an increased coherence time of 1.5 s, $N \sim 2 \times 10^4$ per shot, and the ability to measure an EDM with two co-magnetometer states simultaneously [114]. Recently, the experiment revealed its final result, which improves upon the ACME II bound by a factor

of 2.4 [97, 115].

1.4.3 Future prospects

New generations of all three of the molecular EDM experiments described above are underway. Since their 2011 result, the YbF experiment has undergone significant upgrades, achieving a demonstrated sensitivity of $1.5 \times 10^{-28}\ e \cdot \mathrm{cm}/\sqrt{\mathrm{day}}$ [116]. Recent accomplishments in performing transverse laser coooling on the YbF beam may further improve their sensitivity by another two orders of magnitude [117]. The ACME III experiment, featured in this dissertation, has completed various demonstrated upgrades and is projected to improve upon the sensitivity of ACME II by at least an order of magnitude within the next few years (Chapter 3). The third generation of the JILA eEDM experiment is also under construction. It plans to use the ground state of the ThF^+ molecule, which has similar properties as HfF^+ while allowing even longer coherence times [106].

In the more distant future, a promising path forward for eEDM measurements lies in using molecules trapped in an optical lattice, which allows the combination of extremely long coherence times and large numbers. Such a method seems experimentally feasible, given recent experimental successes in creating magneto-optical traps for molecules, a required step to load the optical lattice [118–120], as well as the use of optical lattices for highly precise atomic clocks [121]. This technique has been proposed for YbF [122]. With projected $N = 10^6$ per shot, $\tau = 10$ s, and $\mathcal{E}_{\mathrm{eff}} = 17.5$ GV/cm, this would accomplish a statistical uncertainty of $2 \times 10^{-32}\ e \cdot \mathrm{cm}/\sqrt{\mathrm{day}}$. Certain species of polyatomic molecules have also been identified as suitable for this technique [123]. YbOH has a similar $\mathcal{E}_{\mathrm{eff}}$ as YbF but is much more easily polarizable and possesses internal co-magnetometer states.

A final method of measuring the eEDM which has been recently proposed is to use EDM-sensitive polar molecules embedded in a solid, transparent matrix of a rare gas such as argon [124, 125]. Such a matrix could be made using the recently developed ice-film nanocapacitors technique [126]. In the matrix, forces between ions in the molecule and

nearest neighbor inert gas atoms inhibit the rotation of the molecule and enable them to stay polarized even in the absence of a laboratory $\mathcal{E}$. The main advantage of this technique would be extremely large numbers, with $N \sim 10^{13}$ for a 1 cm^3 sample. Using a polar molecule such as BaF and $\tau \approx 100$ ms (limited by the density of molecules in the sample), a statistical uncertainty of $5 \times 10^{-34}\ e \cdot \mathrm{cm}/\sqrt{\mathrm{day}}$ could be obtained, which would be tantalizingly close to the SM prediction, $|d_\mathrm{e}| \sim 10^{-35}\ e \cdot \mathrm{cm}$ [39]. However, this technique is very new, and major experimental work remains to be done to fabricate such samples and study their suitability for an EDM experiment, especially for systematic effects which may be quite different from atomic or molecular EDM experiments.

1.5 Measuring the electron EDM with ThO

The ACME experiment measures the EDM by performing a spin precession measurement on a beam of thorium monoxide (ThO) molecules in the $|H, J = 1\rangle$ state. In this section, we will review the basic structure and properties of ThO which make it a compelling species for an EDM experiment.

1.5.1 The general structure of ThO

As a diatomic molecule, ThO possesses electronic, vibrational, and rotational degrees of freedom. It is best classified as a Hund's case (c) molecule where the orbital $\vec{L}$ and spin $\vec{S}$ angular momentums couple to the total electronic angular momentum $\vec{J}_e = \vec{L} + \vec{S}$ [127, §9.2]. However, J_e is not a good quantum number. Instead, a good quantum number is Ω, the projection of $\vec{J}_e$ onto the molecule internuclear axis $\hat{n}$:

$$\Omega = \vec{J}_e \cdot \hat{n}, \tag{1.14}$$

$$= \Lambda + \Sigma, \tag{1.15}$$

where Λ and Σ respectively refer to the projections of $\vec{L}$ and $\vec{S}$ onto the internuclear axis. A diatomic molecule can also rotate around its center of mass along its internuclear axis $\hat{n}$, resulting in the rotational angular momentum $\vec{R}$ that combines with Ω to give the total angular momentum $\vec{J} = \vec{R} + \Omega\hbar\hat{n}$. Hence J is also a good quantum number.

Each state of ThO is commonly described with the basis set $|Y, \nu, J, M, \Omega\rangle$. Y is a letter designating the electronic state (X, C, H, Q, etc.), where X indicates the ground state. Fig. 1.5.1 shows all electronic states used in the ACME II and III experiments and their associated commonly used transitions. ν denotes the vibrational quantum number. In this work virtually all states are in the lowest vibrational state $\nu = 0$, so this number will usually be omitted from the description of the state. J and M respectively denote the total angular momentum and its projection along the laboratory $\hat{z}$ axis. Ω has already been defined above. The ACME electron EDM measurement is conducted in the lowest ($J = 1$) rotational level of the H-state.

1.5.2 Advantageous features of the $H\ ^3\Delta_1$ state of ThO

Although ThO is a Hund's case (c) molecule, each electronic state can be expressed as a sum of Hund's case (a) states, where Λ and Σ are good quantum numbers. The H electronic state has a predominant composition of 98.4% $^3\Delta_1$ state ($S = 1$, $|\Lambda| = 2$, $|\Omega| = 1$),[5] which has two valence electrons in σ and δ orbitals. This results in several features that are favorable for an EDM measurement:

Large effective electric field. The σ electron spends a lot of time near the heavy Th nucleus, resulting in $\mathcal{E}_{\text{eff}} = 78(2)$ GV/cm,[6] the highest among any of the molecules being actively used for eEDM experiments.

Small magnetic moment. As $\Lambda + \Sigma = \Omega$ always holds (Eq. 1.15), in a $^3\Delta_1$ state,

[5] Hund's case (a) states are written with the spectroscopic notation $^{2S+1}|\Lambda|_{|\Omega|}$.

[6] This number is obtained by averaging the latest results from the two groups which have independently computed $\mathcal{E}_{\text{eff}}$: Skripnikov (2016), 79.9 GV/cm, 4% uncertainty [103] and Denis and Fleig (2016), 75.2 GV/cm, 3% uncertainty [128]. Previous calculations of $\mathcal{E}_{\text{eff}}$ can be found in [129–132].

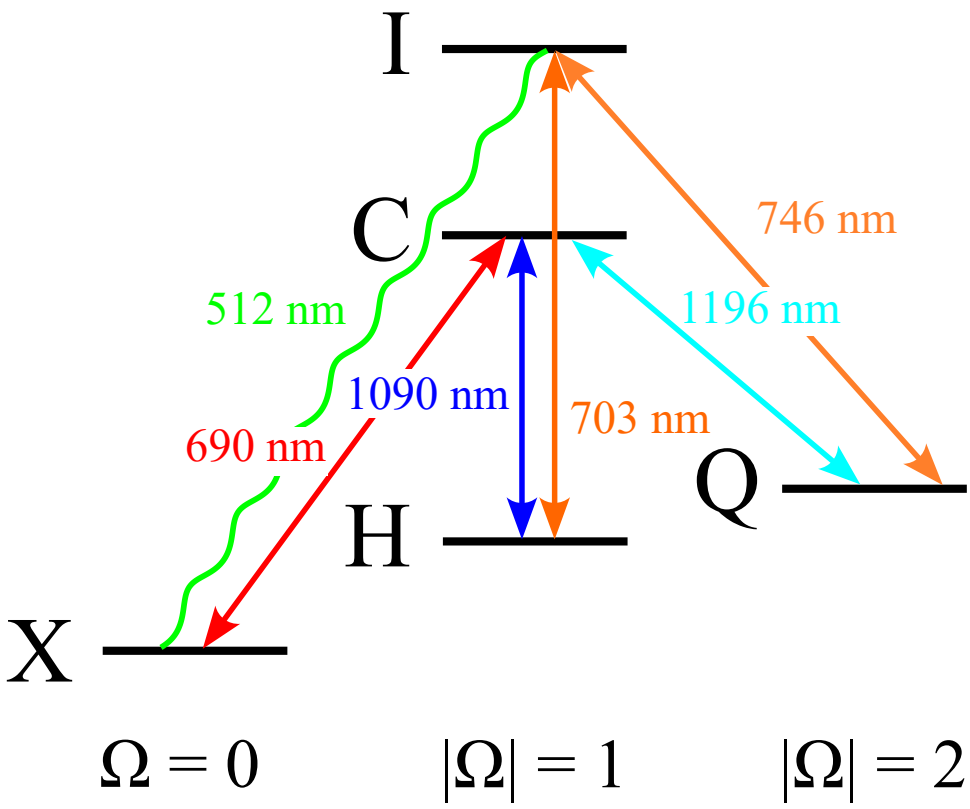

Figure 1.5.1: **Electronic levels of ThO**, including arrows indicating transitions used in the ACME II and III experiments discussed in this dissertation. Energy distance between states is not to scale. The Ω quantum number is the projection of the total electronic angular momentum $\vec{J}_e$ onto the molecule internuclear axis $\hat{n}$.

$\Lambda = -2\Sigma$. In other words, the projection of the spin angular momentum is always opposed to the projection of the orbital angular momentum. As the g-factor of the state $g_H \approx g_s\Sigma + g_L\Lambda = 0$, this means the H-state is magnetically very insensitive, which is very useful for suppressing magnetic field-related systematic effects. The value of g_H has been measured to be 0.0088(1) [133].

High polarizability. Ω-doubling is present in all states of ThO with $\Omega \neq 0$. In the $H\ ^3\Delta_1$ state, $\Omega = \pm 1$. These two values correspond to the clockwise and counterclockwise rotation of the angular momentum vector around the internuclear axis [134, p. 329]. In the absence of an electric field, these two states are degenerate, resulting in two parity eigenstates

$$|\psi\rangle = \frac{|\Omega = 1\rangle \pm |\Omega = -1\rangle}{\sqrt{2}}, \tag{1.16}$$

which have a very close spacing $\Delta_{|\Omega|,J=1} = 2\pi \times 362(2)$ kHz [83, §D.3]. In the ACME experiment, these states can be fully mixed with a very small laboratory electric field $\mathcal{E} = 10$ V/cm to achieve the maximum value of $\mathcal{E}_{\text{eff}}$ [135]. This is much less than the $\mathcal{E}$ required in several other experiments, such as the Imperial YbF experiment which requires $\mathcal{E}$ of several thousand kV/cm to polarize a much larger rotational splitting.

A low $\mathcal{E}$ requirement for polarization has several advantages. First, it strongly suppresses motional magnetic field and leakage current systematic effects, which have plagued many generations of EDM experiments. Due to this and the small magnetic moment of the H-state, these effects have never been a concern in the ACME experiment so far. Second, it gives flexibility in choosing $\mathcal{E}$ without sacrificing EDM sensitivity. In ACME II, data was taken at two different $\mathcal{E}$ (80 V/cm and 140 V/cm) as a systematic check. Lastly, it omits risks when working with high voltage field plates such as electric field breakdown [136].

Internal co-magnetometer states. Once the two parity eigenstates are fully mixed in an $\mathcal{E}$-field, two parity-mixed states emerge corresponding to whether $\vec{\mathcal{E}}_{\text{eff}}$ is aligned or anti-aligned with the internuclear axis $\hat{n}$.[7] These two states are distinguished by the quantum

[7] In the convention used in ACME literature, $\hat{n}$ is taken to point from the O atom towards the Th atom,

number $\tilde{\mathcal{N}} = \Omega M \text{sign}(\mathcal{E} \cdot \hat{n})$. The EDM energy shift is [83, §2.1.4]

$$\Delta E_{\text{EDM}} = -d_e \vec{S} \cdot \vec{\mathcal{E}}_{\text{eff}} = (\vec{S} \cdot \hat{n})(d_e \mathcal{E}_{\text{eff}}) = \Sigma(d_e \mathcal{E}_{\text{eff}}) = -\Omega(d_e \mathcal{E}_{\text{eff}}) = -M\tilde{\mathcal{N}}\tilde{\mathcal{E}} d_e \mathcal{E}_{\text{eff}}, \quad (1.17)$$

where we have defined $\tilde{\mathcal{E}} = \text{sign}(\mathcal{E} \cdot \hat{n})$ and used the fact that for the H-state, $\Omega = -\Sigma$. This means that by alternately performing the EDM measurement in the two $\tilde{\mathcal{N}} = \pm 1$ states, one can reverse the direction of the EDM shift without physically reversing the direction of the laboratory $\tilde{\mathcal{E}}$. Instead, switching $\tilde{\mathcal{N}}$ can be done by tuning the preparation lasers several tens of MHz, which is readily accomplished using an acousto-optic modulator (AOM). Hence, the two states function as internal co-magnetometers. This is an extremely useful feature for diagnosing and rejecting systematic errors, as only effects that are odd with respect to the reversal of $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ will cause a shift in the EDM channel $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$. An example would be a leakage current systematic effect, which to first order would reverse with $\mathcal{E}$ but not $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$.

In addition to the advantages afforded by the $^3\Delta_1$ structure of the H-state, ThO possesses several other advantages for an EDM measurement.

Adequate radiative lifetime. The H-state is a metastable state which can only radiatively decay to the ground state X (Fig. 1.5.1). Its lifetime has been measured to be $\tau_H = 4.2\ (5)$ ms (Chapter 4). While this is not as long as HfF^+ (3 s) or YbF and ThF^+ (ground state), it is long enough for an experiment using a slow molecular beam, unlike species with very short-lived lifetimes for their EDM state such as PbO [137]. Moreover, as explained in Chapter 3.1, there is much room for improvement in ACME III by extending the spin precession time to fully make use of the lifetime.

No nuclear spin. ThO does not have nuclear spin, which means that there is no need to deal with hyperfine structure, unlike YbF or HfF^+.

and a positive $\vec{\mathcal{E}}_{\text{eff}}$ is taken to be anti-aligned with $\hat{n}$, such that $\vec{\mathcal{E}}_{\text{eff}} = -\mathcal{E}_{\text{eff}}\hat{n}$. [113].

High-yield production method available. A cold, intense beam of ThO can be produced using cryogenic buffer gas beam technology [112, 138]. The beam source used in ACME I and II produced $\sim 3 \times 10^{11}$ ground state (X) molecules per steradian per pulse at a temperature of ~ 4 K, mean velocity of 180-200 m/s, and repetition rate of 50 Hz. This is a large number in comparison to experiments such as YbF [139]. Established atomic physics techniques such as rotational cooling and STIRAP [140] can be applied to efficiently transfer the molecules to the $|H, J = 1\rangle$ state. More recently, studies have revealed that it is possible to transfer the molecules to the Q-state, which is long-lived with a large linear Stark shift, allowing the building of a molecular lens to collimate the beam and further increase molecule yield [141, 142].

Well-known spectroscopy and accessible laser transitions. The structure of ThO has been studied since the 1960s (e.g. [143, 144]), removing the need to perform preliminary spectroscopy which is common in many new proposals for molecular EDM experiments. All the energy transitions used in the ACME experiments are accessible with commercially available diode lasers.

1.5.3 General EDM measurement scheme

ACME measures the EDM using the previously described general spin precession technique (§1.3.1). The precession occurs in the interaction region, where electric and magnetic fields are applied to molecules in the $|H, J = 1\rangle$ state of ThO. In the fully polarized (i.e. linear Stark) regime, this results two $\tilde{\mathcal{N}} = \pm 1$ co-magnetometer states of mixed parity and three energy shifts (Fig. 1.5.2).

Assuming fields $\vec{\mathcal{E}}$ and $\vec{\mathcal{B}}$ pointing along $\hat{z}$ with magnitudes $\mathcal{B} = |\vec{\mathcal{B}}|$, $\mathcal{E} = |\vec{\mathcal{E}}|$, the combined energy shift for the $M = \pm 1$ levels is [113, 145]

$$\Delta E(M, \tilde{\mathcal{N}}) = -D_1 \mathcal{E} \tilde{\mathcal{N}} - M \mu_1 \mathcal{B} \tilde{\mathcal{B}} - M \tilde{\mathcal{N}} \tilde{\mathcal{E}} d_e \mathcal{E}_{\text{eff}}, \tag{1.18}$$

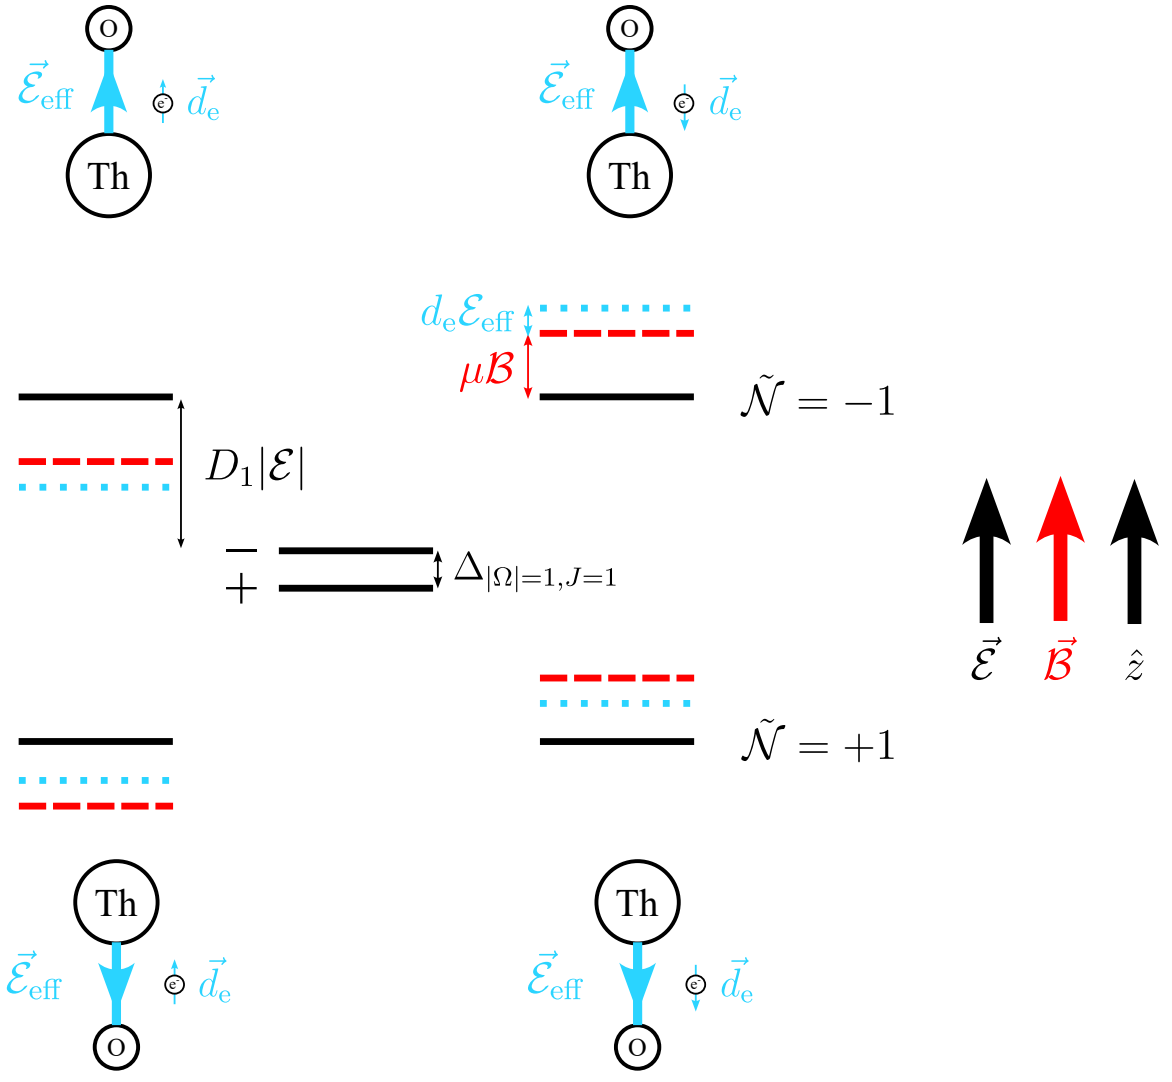

Figure 1.5.2: **Energy structure of the $H\ ^3\Delta_1$, $J = 1$ state in the presence of electric and magnetic fields**. In the absence of any fields, there are three degenerate M-states with an Ω-doublet splitting of $\Delta_{|\Omega=1|,J=1} = 362$ kHz. In the spin precession region of the ACME experiment, electric ($\mathcal{E}$) and magnetic ($\mathcal{B}$) fields are applied to the molecule along $\hat{z}$. The electric fields fully polarize the molecule, resulting in two $\tilde{\mathcal{N}} = \pm 1$ states of mixed polarity for the $M = \pm 1$ states and a Stark shift $D_1|\mathcal{E}|$. These states correspond to the direction of $\vec{d}_e$ being aligned or anti-aligned to the molecule's internal electric field $\vec{\mathcal{E}}_{\text{eff}}$, resulting in eEDM shifts of opposite sign. The magnetic field results in a Zeeman shift $\mu\mathcal{B}$ with opposite signs for $M = \pm 1$.

whose three terms correspond to the Stark, Zeeman, and EDM shifts respectively. The EDM shift term is the same as in Eq. 1.17. For the Stark shift, $D_1 = 2\pi \times 1$ MHz/(V/cm) is the electric dipole moment of $|H, J = 1\rangle$. (The $M = 0$ state is unaffected by the electric field.) For the Zeeman shift, $\tilde{\mathcal{B}} = \pm 1$ is the direction of the magnetic field (aligned or anti-aligned to $\hat{z}$) and μ_1 is the magnetic moment for $|H, J = 1\rangle$, which is slightly different for the two $\tilde{\mathcal{N}}$ states and also affected by mixing with other electronic and rotational states in the presence of an electric field. It can be expressed as

$$\mu_1(\mathcal{E}, \tilde{\mathcal{N}}) = (g_1 + \eta_1 \mathcal{E} \tilde{\mathcal{N}})\mu_B, \tag{1.19}$$

where $\bar{g}_1 \mu_B = -0.00440(5)\mu_B = -2\pi \times 6$ kHz/G and $\eta_1 = 0.79(1)$ nm/V [113, 145]. In the ACME experiment, $\mathcal{E} \sim 100$ V/cm and $\mathcal{B} \sim 1$ mG (for ACME II), resulting in Stark and Zeeman shifts of roughly hundreds of MHz and Hz respectively. In the spin precession region, the molecules are prepared in a superposition of $M = \pm 1$ in one of the $\tilde{\mathcal{N}}$ states:

$$|\psi(t = 0, \tilde{\mathcal{N}})\rangle = \frac{1}{\sqrt{2}}(|M = +1, \tilde{\mathcal{N}}\rangle + |M = -1, \tilde{\mathcal{N}}\rangle), \tag{1.20}$$

after which they precess within the xy-plane under the influence of $\mathcal{E}$ and $\mathcal{B}$ until $t = \tau$, acquiring a phase

$$|\psi(t = \tau, \tilde{\mathcal{N}})\rangle = \frac{1}{\sqrt{2}}(e^{-i\phi}\,|M = +1, \tilde{\mathcal{N}}\rangle + e^{+i\phi}\,|M = -1, \tilde{\mathcal{N}}\rangle), \tag{1.21}$$

which is measured by projecting the state by exciting the molecule to the short-lived I-state and detecting the resulting fluorescence from the decays. Only terms in Eq. 1.18 that depend on M contribute to this phase shift:

$$\phi \approx \frac{-(\mu_1 \mathcal{B}\tilde{\mathcal{B}} + \tilde{\mathcal{N}}\tilde{\mathcal{E}} d_e \mathcal{E}_{\text{eff}})\tau}{\hbar}. \tag{1.22}$$

By measuring ϕ multiple times while switching the values of $\tilde{\mathcal{N}}$ and $\tilde{\mathcal{E}}$, the EDM contribution to the phase $\phi^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ can be extracted. The precession time τ is typically measured by switching $\tilde{\mathcal{B}}$ and using the known values of μ_1 and $\mathcal{B}$. The value of $\mathcal{E}_{\text{eff}}$ is taken from molecular theory calculations [103, 128]. With these, it is possible to determine the electron EDM value d_{e}. Assuming the measurement is shot-noise limited, following Eq. 1.9 it will have an uncertainty of

$$\delta d_{\text{e}} = \frac{\hbar}{2\tau \mathcal{E}_{\text{eff}} \sqrt{\dot{n} T}}, \tag{1.23}$$

where $\dot{n}$ is the rate of detected molecules participating in the spin precession, and T is the experiment time. For ACME II, $\tau = 1$ ms, $\mathcal{E}_{\text{eff}} = 78$ GV/cm, $\dot{n} \sim 300,000$ molecules/pulse $= 1.5 \times 10^7$ Hz, $T \approx 350$ hours, giving $\delta d_{\text{e}} \sim 10^{-30}\ e \cdot \text{cm}$, which is within a factor of a few of the actual ACME II statistical uncertainty, $\delta d_{\text{e}} = 3.1 \times 10^{-30}\ e \cdot \text{cm}$.

2

The ACME II electron EDM measurement

The ACME II experiment in 2018 obtained a new upper limit on the electron EDM, about an order of magnitude smaller than the previous best limit also obtained from the ACME I experiment [49, 96]. Various efforts to upgrade the ACME I apparatus as part of the ACME II campaign have been described in the Refs. [146, 147]. The ACME II experiment and its results have been extensively discussed in Refs. [49, 70, 148]. This chapter provides sufficient background to understand the research efforts towards the next generation ACME III apparatus (Chapter 3). In the ACME II measurement, I made key contributions towards experimental control, data acquisition, and data analysis, in addition to assisting in commissioning of the apparatus, maintaining lasers and optics, and taking final run data.

2.1 Apparatus and Methods

A schematic of the whole ACME II beamline is shown in Fig. 2.1.1. The molecular beam travels along $+\hat{x}$, gravity points along $+\hat{y}$ and the electric and magnetic fields are applied along $\hat{z}$ [113]. Lasers are sent horizontally primarily along $+\hat{z}$ (East to West).

2.1.1 Molecular beam source

A beam of ThO molecules is produced via laser ablation in a cryogenic buffer gas beam source [112, 138]. The core of the beam source is a small copper chamber (commonly called the cell) held at 16-17 K using a pulse tube cooler[1] and a resistive heater. The cell contains two disc-shaped pressed thoria (ThO_2) ablation targets (0.6" diameter) and several openings for optical access, the buffer gas fill line, and the exit of the molecular beam. It is encased in a 4 K inner layer to cryopump the neon buffer gas and shield blackbody radiation, a 60-70 K middle layer to provide additional shielding, and an outer room-temperature vacuum chamber layer.

During operation, neon flows into the cell as a buffer gas at ≈ 40 SCCM.[2] Pulsed YAG laser beams[3] are sent into the cell at 50 Hz and focused with a lens near the surface of the ablation target. This creates pulses of ThO molecules alongside other byproducts. The molecules are cooled by elastic collisions with the neon and extracted through a 6 mm diameter aperture. Outside the cell, the gas rapidly expands and cools further until inter-molecular collisions are no longer significant, such that the molecules follow a ballistic trajectory throughout the rest of their journey through the beamline. This process produces $\sim 10^{11}$ molecules per pulse with an FWHM of 45 degrees. The longitudinal velocity of the molecules has a mean and standard deviation of ≈ 200 m/s and 13 m/s, although 10-20% faster speeds are common with newer ablation targets. The rotational temperature is ≈ 4

[1] Cryomech PT415.
[2] Standard cubic centimeters per minute.
[3] Litron Nano TRL 80-200, 1064 nm, 60-70 mJ pulses of ≈ 5 ns each.

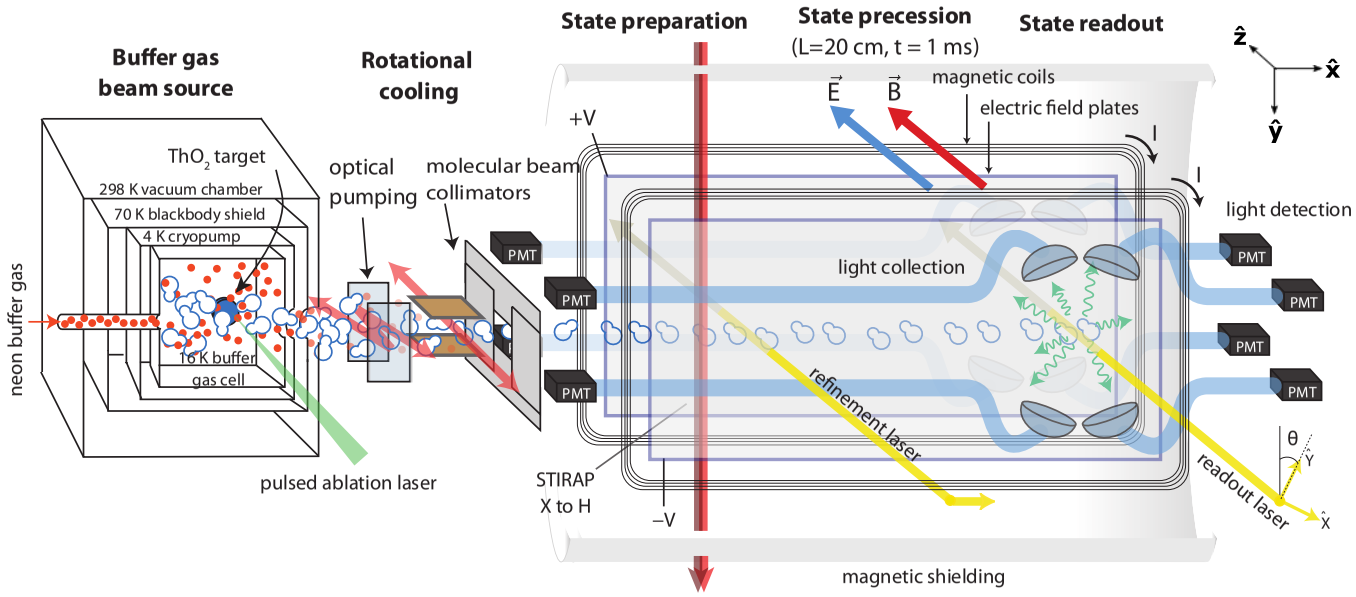

Figure 2.1.1: **Schematic of the ACME II experiment apparatus.** A cryogenic buffer gas beam source produces ThO molecules mostly in the ground state, which are then transferred to the EDM-sensitive $H^3\Delta_1$ state. The molecules enter an interaction region where they are prepared in an initial superposition of $|M = \pm 1\rangle$ states, undergo spin precession under the influence of electric and magnetic fields, and have their accumulated phase read out by laser-induced fluorescence. The 512 nm photons are detected by a set collection optics connected to PMTs.

K, corresponding to $\approx 90\%$ of molecules being in the $J = 0$–3 rotational states.

With a new ablation target, the beam source can typically produce molecules continuously for 10-20 minutes before the spot on the target is exhausted and the signal decreases below acceptable levels. At this point, the data taking is briefly paused while the ablation laser is redirected to a different spot on the target by adjusting a motorized mirror. After about 1-2 weeks of constant (18 hr/day) operation, the target can no longer be used and must be replaced.

The ACME beam source was originally developed for ACME I and its basic physics and features are essentially the same throughout ACME II and III, with the exception of some practical upgrades in the latter (§3.1.4, §3.3.3). More detailed accounts of the beam source (including its development) can be found in past ACME dissertations [83, 147].

2.1.2 Rotational cooling

Having exited the beam box, the molecules enter a module where optical pumping on three $X - C$ (690 nm) transitions is performed to transfer the populations in $J = 1,\ 2,\ 3$ to $|X, J = 0\rangle$, a process we term "rotational cooling". The $|C, \nu = 0\rangle \rightarrow |X, \nu = 0\rangle$ branching ratio has been measured to be 74% [141], making it suitable for rotational cooling. In the first stage of rotational cooling (Fig. 2.1.2a), pumping is performed on $|X, J = 2^+\rangle \rightarrow |C, J = 1^-\rangle$ and $|X, J = 3^-\rangle \rightarrow |C, J = 2^+\rangle$.[4] Due to parity selection rules, the molecules spontaneously decay to $|X, J = 0^+, 1^-\rangle$ respectively. However, this scheme does not work to transfer the population from $|X, J = 1^-\rangle$ to $|X, J = 0^+\rangle$, as these are states of opposite parity. Thus in the second stage an electric field $\mathcal{E} \approx 150$ V/cm is applied to fully mix the C-state parity doublets. Optical pumping is then performed on $|X, J = 1^+\rangle \rightarrow |C, J = 1, \mathcal{N} = +1\rangle$ (Fig. 2.1.2b), resulting in spontaneous decay to $|X, J = 0, 2\rangle$. Here and in the figure we use the $\mathcal{N} = \pm 1$ notation to denote the two sublevels of mixed parity, as

[4] The superscript denotes the parity of the state. For $\Omega = 0$ electronic states such as X, the parity is $\mathcal{P} = (-1)^J$, whereas in the absence of an electric field, electronic states with $|\Omega| \geq 1$ have sublevels of each parity due to Ω-doubling, as previously discussed in §1.5.2.

was previously done for the H-state (§1.5.2).

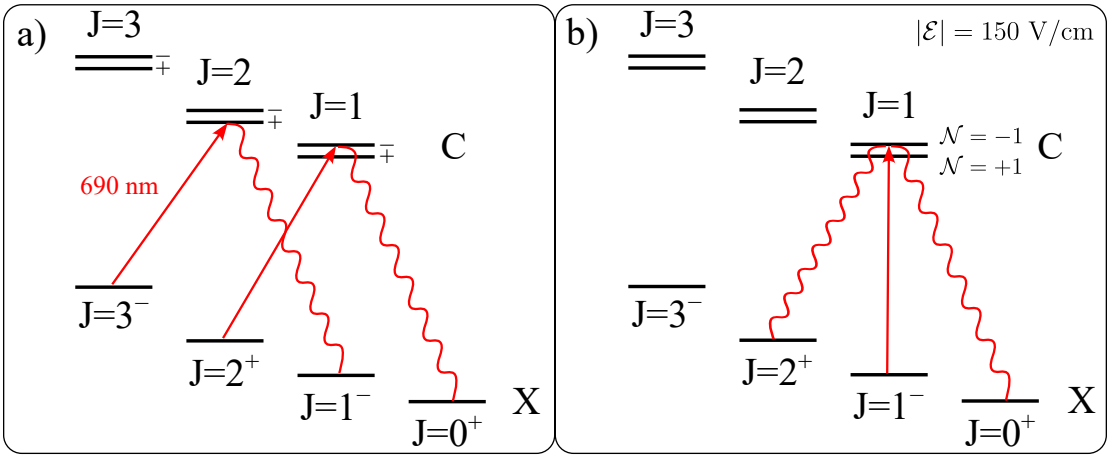

Figure 2.1.2: **ACME II rotational cooling.** Rotational cooling is applied using the 690 nm $X - C$ transition to transfer higher-lying rotational populations to $J = 0$. In the figure, only the $J \leq 3$ levels are displayed. (Because the C state has $|\Omega| = 1$, it has no $J = 0$ state.) a) In the first stage, optical pumping is performed on $|X, J = 2^+\rangle \rightarrow |C, J = 1^-\rangle$ and $|X, J = 3^-\rangle \rightarrow |C, J = 2^+\rangle$. The population spontaneously decays to $|X, J = 0^+, 1^-\rangle$. b) In the second stage, the $J = 1$ population is transferred to $J = 0$ via optical pumping on $|X, J = 1^-\rangle \rightarrow |C, J = 1\rangle$ in the presence of an electric field to fully mix the C-state parity doublet, whose sublevels of mixed parities are denoted by $\mathcal{N} = \pm 1$, analogous to the $\tilde{\mathcal{N}}$ sublevels in the H-state used for an EDM measurement (Fig. 1.5.2).

In each of the rotational cooling stages, the laser(s) are passed through the molecular beam 5-7 times with alternating $\hat{x}$ and $\hat{z}$ polarizations. This polarization switching is required to avoid dark states formed by superpositions of M-sublevels. Overall, rotational cooling increases the population in $|X, J = 0\rangle$ by a factor of ≈ 2.5. A more detailed account of the rotational cooling can be found in Lasner's thesis [70, §2.6].

2.1.3 Interaction region

The molecules then enter the interaction region where the spin precession occurs. The region is shielded with five layers of mu-metal magnetic shielding to achieve a background

$|B| \approx$ 100-200 μG. An electric field $\mathcal{E} = 80$ V/cm or 140 V/cm is applied with a pair of parallel glass field plates coated with indium tin oxide (ITO). A magnetic field $\mathcal{B}$ is applied along $\hat{z}$ with a set of rectangular cosine coils [113, §3.2.6]. $|\mathcal{B}|$ is regularly switched between 0.7 mG, 1.3 mG, and 26 mG.

2.1.4 State preparation and refinement

In the interaction region, the molecules are first transferred from $|X, J = 0, M = 0\rangle$ to the $|H, J = 1, M = \pm 1, \tilde{\mathcal{N}}\rangle$ state by Stimulated Raman adiabatic passage (STIRAP, Fig. 2.1.3a). STIRAP is a well-known technique that allows nearly-complete transfer between two states (in this case $X \rightarrow H$) via an intermediary third state ($|C, J = 1, M = 0, \mathcal{P} = -1\rangle$) [149, 150]. A pair of co-propagating, partially overlapping laser beams are sent through the molecular beam: the Stokes beam ($H - C$, 1090 nm, 10 W, $\hat{x}$-polarized) followed by the Pump beam ($X - C$, 690 nm, 300 mW, $\hat{z}$-polarized). This counterintuitive pulse sequence causes the adiabatic evolution of the dark state in the system from X to H, resulting in highly efficient transfer. The two laser beams are near-resonant with detunings Δ_{Stokes} and Δ_{Pump} respectively as depicted in Fig. 2.1.3a. (Note that the STIRAP *two-photon detuning* is defined as $\delta \equiv (\Delta_{\text{Pump}} - \Delta_{\text{Stokes}})/2$. This will be an important quantity when discussing systematic errors in §2.2.)

Because of the high power requirement of the Stokes beam (due to the small transition dipole moment of the $H - C$ transition [151]), the lasers are sent in vertically along $\hat{y}$ instead of horizontally to avoid going through and heating up the field plates, which may cause systematic effects (§2.2.3). To achieve sufficient intensity, the beam waists must be compressed to ≈ 150 μm along $\hat{x}$. Because of the parity $\mathcal{P} = -1$ of the intermediate C-state and the polarizations of the Stokes and Pump beams, the STIRAP process results in the molecules being prepared in a coherent superposition of $|M = \pm 1\rangle$ states:

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

The frequency of the Stokes laser is constantly switched every 0.5 s using an AOM to alternately prepare the molecules in one of the $\tilde{\mathcal{N}} = \pm 1$ states.

In ACME II, this STIRAP process resulted in a transfer efficiency $\eta = 75\%$ of the molecules in $|X, J = 0\rangle$ to $|H, J = 1\rangle$, likely limited by the 1090 nm power and small misalignments due to the difficulty of sending in the beams vertically. This transfer efficiency is about 12 times better compared to the optical pumping transfer method in ACME I. More details of the implementation ACME II STIRAP can be found in Refs. [140, 148].

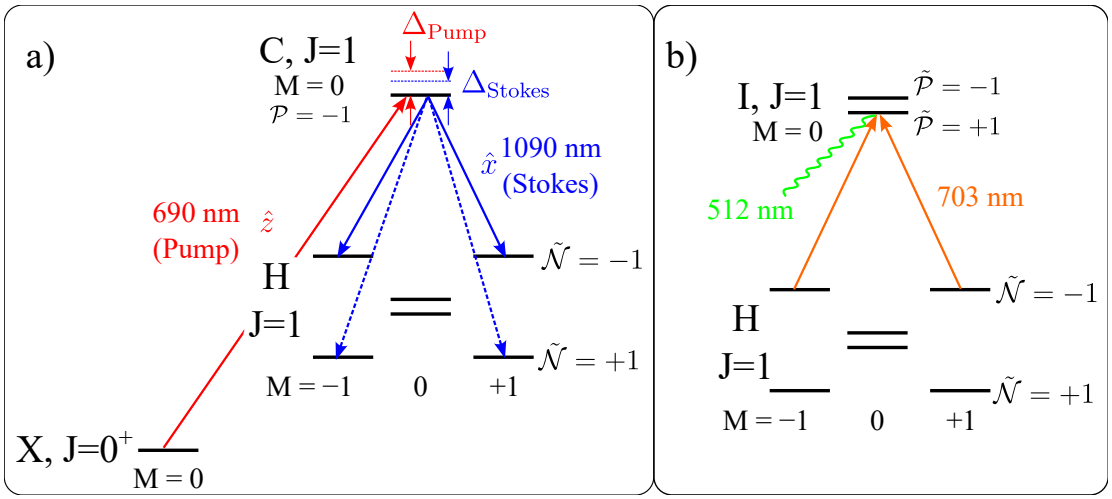

Figure 2.1.3: **ACME II state preparation.** a) Molecules are transferred with near-unity efficiency from X to H by STIRAP via the intermediate C state. The process results in the molecules being prepared in a coherent superposition of $M = \pm 1$ states. The frequency of the Stokes beam is constantly switched to prepare in either one of the $\tilde{\mathcal{N}} = \pm 1$ states. b) After STIRAP transfer, optical pumping on the $H - I$ transition is performed to refine the prepared state. The molecules then undergo spin precession and accumulate a phase ϕ, which is read out by laser-induced fluorescence on the same transition, but with with the polarization of the laser rapidly switched between $\hat{X}$ and $\hat{Y}$. The state refinement and readout lasers are switched between $|\tilde{\mathcal{N}} = \pm 1\rangle$ (together with the STIRAP Stokes laser). The readout laser is switched on longer time scales between $|I, \tilde{\mathcal{P}} = \pm 1\rangle$.

After STIRAP, the molecules are prepared in a coherent spin-state aligned along $\hat{x}$ which is in principle ready to undergo precession. However, imperfections in STIRAP

process (such as from laser misalignments) may introduce systematic errors. To suppress these, the state is reprojected with a 703 nm "refinement" laser[5] tuned to the $|H, J = 1, M = \pm 1, \tilde{\mathcal{N}}\rangle \rightarrow |I, J = 1, M = 0, \tilde{\mathcal{P}} = +1\rangle$ transition (Fig. 2.1.3b). The laser is sent in horizontally along $\hat{z}$ immediately downstream to the vertical STIRAP beams. This allows alignment of the polarization of the laser such that $\theta_{\text{ref}}^{ST} \approx 0$, where θ_{ref}^{ST} is the angle between the polarization and the STIRAP-prepared state (which is nominally along $\hat{x}$). The frequency of the laser is shifted according to which of the $\tilde{\mathcal{N}} = \pm 1$ states was previously prepared by STIRAP.

This process pumps out molecules in the state

$$|\psi\rangle = \frac{1}{\sqrt{2}}(|M = +1, \tilde{\mathcal{N}}\rangle - |M = -1, \tilde{\mathcal{N}}\rangle), \tag{2.2}$$

which are excited to the short-lived I-state and quickly decay to the ground state X, no longer participating in the spin precession. The remaining H-state molecules are those in the orthogonal spin state (Eq. 2.1).

2.1.5 State precession, readout, and detection

The molecules fly through the spin precession region (a distance of $L = 22$ cm, equivalent to a precession time $\tau \approx 1$ ms), precessing in the xy-plane under the influence of $\mathcal{E}$ and $\mathcal{B}$ which are aligned along $\hat{z}$. At $t = \tau$, the state accumulates a phase ϕ:

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

which is measured in the detection region by sending in a 703 nm "readout" laser beam along $\hat{z}$ to excite the molecules to the short-lived I-state. The laser is rapidly switched between two orthogonal polarizations (labeled $\hat{X}$ and $\hat{Y}$) at 200 kHz (1.9 μs per polarization and 0.6 μs dead time before switching, Fig. 2.1.4a). This effectively projects $|\psi(t = \tau)\rangle$ to

[5] M Squared SolsTiS TiSapph.

two orthogonal bases. The molecules then spontaneously decay back to the ground state, emitting 512 nm photons which are detected by a set of eight lens doublets (four each on the East and West sides). Each lens doublet is attached to a light pipe to transport the light outside the vacuum chamber. (The collection optics are described in more detail in §5.1.1.) The light pipes are connected to 8 photo-multiplier tube detectors (PMTs)[6] mounted just outside the interaction region. The PMTs generate photocurrents which are amplified and recorded by a digitizer data acquisition device. This produces fluorescence traces containing about 2-3 ms of fluorescence signal, as depicted in Fig. 2.1.4b. The total number of detected photons per molecular pulse is $\sim 300,000$.

2.1.6 Extraction of the phase

To extract the accumulated phase ϕ from the spin precession, we first calculate the asymmetry $\mathcal{A}$, which is defined as [113]

$$\mathcal{A} \equiv \frac{S_X - S_Y}{S_X + S_Y} = \mathcal{C}\cos\left[2(\phi - \theta)\right], \tag{2.4}$$

where S_X (S_Y) is the integrated fluorescence signal when the readout laser polarization is aligned along $\hat{X}$ ($\hat{Y}$), θ is the angle between the polarizations of the refinement and readout lasers, and $\mathcal{C}$ is the measurement contrast, defined as

$$\mathcal{C} \equiv \frac{1}{2}\frac{\partial \mathcal{A}}{\partial \theta} \approx \frac{1}{2}\frac{\partial \mathcal{A}}{\partial \phi}. \tag{2.5}$$

The definition of $\mathcal{A}$ in Eq. 2.4 makes it immune to periodic fluctuations in signal which occur due to the ThO production process. The values of $|\mathcal{B}|$ and θ are set such that

$$(\phi - \theta) \approx \frac{\pi}{4}(2n+1) \tag{2.6}$$

[6] Hamamatsu R7600U-300.

for the integer n. This maximizes $\mathcal{C}$, ensuring the phase is measured in the steepest and most sensitive region of the precession fringes. $\mathcal{C}$ is measured by dithering θ around the center of this region by ± 0.1 rad, with a typical value of 0.95.

The measured phase Φ is then computed by calculating

$$\Phi = \frac{\mathcal{A}}{2\mathcal{C}} + \frac{\pi}{4}(2n+1), \tag{2.7}$$

which is approximately equal to ϕ, unless there are excess noise sources or systematic effects. Φ contains contributions from magnetic and EDM precession (Eq. 1.22) and any other mechanisms Φ' that affect the phase:

$$\Phi \approx \frac{-(\mu_1 \mathcal{B}\tilde{\mathcal{B}} + \tilde{\mathcal{N}}\tilde{\mathcal{E}} d_e \mathcal{E}_{\text{eff}})\tau}{\hbar} + \Phi'(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}, ...). \tag{2.8}$$

2.1.7 Experimental switches, timescales, and data structures

Experiment switches

In order to extract the EDM contribution to Φ, which depends on the sign of $\tilde{\mathcal{N}}\tilde{\mathcal{E}} = \pm 1$, the entire measurement sequence above is repeated many times while varying $\tilde{\mathcal{N}}$, $\tilde{\mathcal{E}}$, and five other main binary switches, resulting in $2^7 = 128$ unique experimental configurations (Fig. 2.1.4c, d):

- $\tilde{\mathcal{N}}$ (every 25 molecular pulses, or ≈ 0.5 s): shifts the frequencies of the STIRAP, refinement, and readout lasers to perform the spin precession in either of the two co-magnetometer states (§1.5.2).
- $\tilde{\mathcal{E}}$ (every 4 $\tilde{\mathcal{N}}$ switches, or ≈ 2 s): reverses the direction of the applied laboratory electric field $\mathcal{E}$ along $\hat{z}$ by reversing the output of the voltage source connected to the field plates.

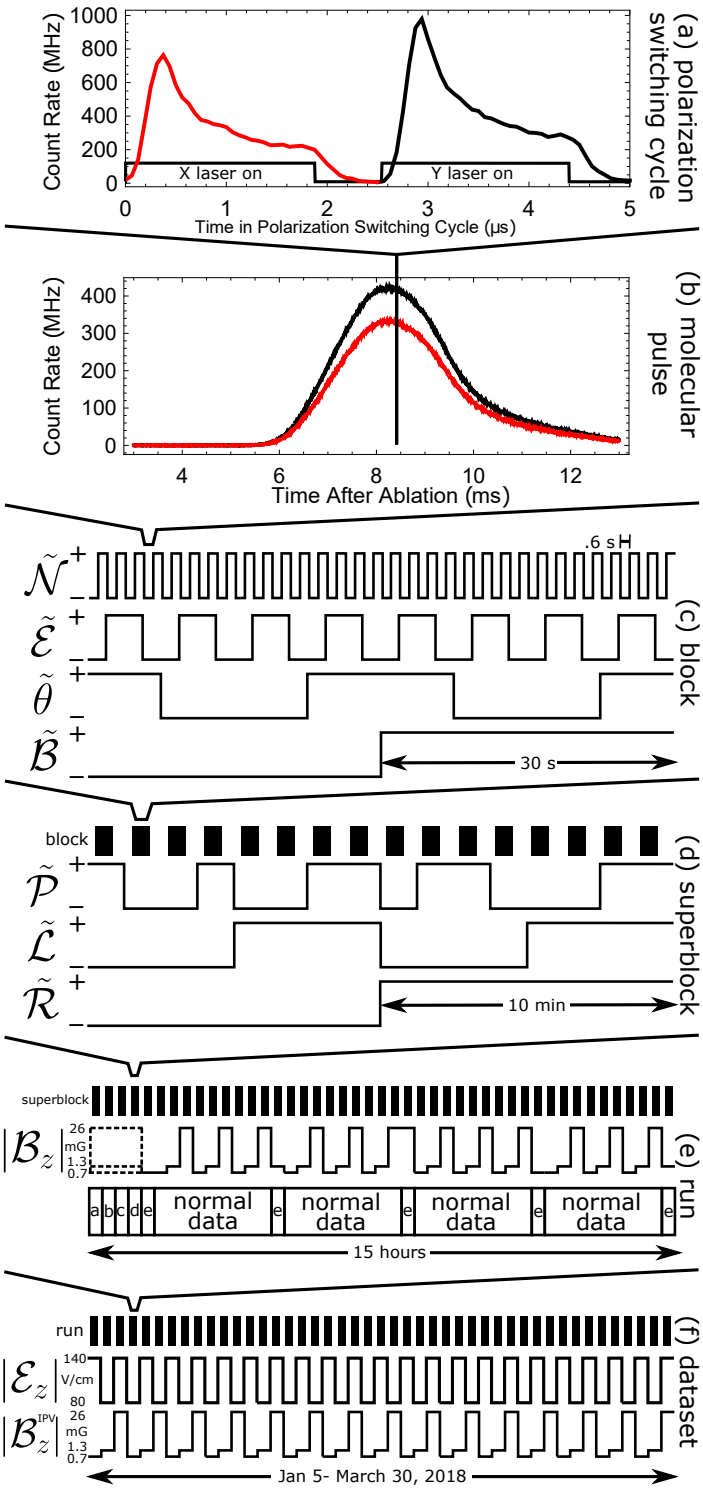

Figure 2.1.4: **ACME II experiment structure and timescales.** a) The readout laser is switched between $\hat{X}$ and $\hat{Y}$ polarizations at 200 kHz, resulting in distinct fluorescence bins (red and black). b) Fluorescence recorded versus time (25 averaged traces). c) Structure of a block, consisting of 64 traces taken in all 16 possible combinations of the $\tilde{\mathcal{N}}$, $\tilde{\mathcal{E}}$, $\tilde{\theta}$, $\tilde{\mathcal{B}}$ switches. d) Structure of a superblock, consisting of 32 blocks taken in all 8 possible combinations of the $\tilde{\mathcal{P}}$, $\tilde{\mathcal{L}}$, $\tilde{\mathcal{R}}$ switches. e) Structure of a run, consisting of 10-20 superblocks taken with $|\mathcal{B}_z|$ varied between three values, as well as some data taken while performing IPVs to monitor systematic slopes. f) Structure of the whole dataset, including variation of $|\mathcal{E}_z|$ and $|\mathcal{B}_z^{\mathrm{IPV}}|$. Figure adapted from [49].

- $\tilde{\theta}$ (every 4 $\tilde{\mathcal{E}}$ switches, or $\approx$ 10 s): rotates the angle of the polarization of the readout laser relative to the refinement beam. The typical step size is $\Delta\theta = 12°$. As noted above, this is performed to measure the contrast $\mathcal{C}$.

- $\tilde{\mathcal{B}}$ (every 2 $\tilde{\theta}$ switches, or $\approx$ 30 s): reverses the direction of the applied laboratory magnetic field $\mathcal{B}$ along $\hat{z}$ by reversing the output of the current source connected to the magnetic field coils.

- $\tilde{\mathcal{P}}$ (every 1-2 blocks, or $\approx$ 1-2 minutes): shifts the frequency of the readout laser to address different parity sub-levels of the I-state. This is physically equivalent to a $\pi/2$ rotation of the orthogonal basis used to measure Φ, and removes effects caused by differences between the $\hat{X}$- and $\hat{Y}$-polarized readout beams.

- $\tilde{\mathcal{L}}$ (every 4 blocks, or $\approx$ 4-6 minutes): switches the physical voltage sources connected to the two field plates while also reversing the output voltages to avoid reversing $\tilde{\mathcal{E}}$. This averages out any non-reversing voltage offsets in the power supplies.

- $\tilde{\mathcal{R}}$ (every 2 $\tilde{\mathcal{L}}$ switches, or $\approx$ 10 minutes: rotates $\lambda/2$ waveplates in the paths of the refinement and readout lasers to rotate their polarizations by $\pi/2$. This is equivalent to the $\tilde{\mathcal{P}}$ switch, but performed mechanically rather than spectroscopically.

The laser frequency-related switches ($\tilde{\mathcal{N}}$, $\tilde{\mathcal{P}}$) require shifts of only several hundred MHz at most and are thus implemented mainly with acousto-optical modulators (AOMs). The polarization rotation switches ($\tilde{\theta}$, $\tilde{\mathcal{R}}$) with $\lambda/2$ waveplates mounted on motorized rotation stages. The $\tilde{\mathcal{L}}$ switch is performed using a TTL-activated relay, and $\tilde{\mathcal{E}}$ and $\tilde{\mathcal{B}}$ are performed by configuring the respective voltage/current sources. All switches are programmatically controlled through a centralized software system (§2.1.9).

Data structures

The seven main switches and other more long-term switches define the data structures of the experiment:

- **Shot** (20 ms): every time a batch of ThO molecules produced is prepared, and detected. This is set by the 50 Hz repetition rate of the pulsed ablation laser. Within each shot, the polarization of the readout laser is switched at 200 kHz (Fig. 2.1.4a).
- **Trace** ($\approx$ 500 ms): 25 consecutive molecular pulses are averaged into one trace (Fig. 2.1.4b).
- **Block** ($\approx$ 1 minute): consists of 64 traces recorded with $2^4 = 16$ unique combinations of ($\tilde{\mathcal{N}}$, $\tilde{\mathcal{E}}$, $\tilde{\theta}$, $\tilde{\mathcal{B}}$), where each combination is repeated 4 times (Fig. 2.1.4c).
- **Superblock** ($\approx$ 15-20 minutes): consists of 16 blocks recorded with 8 unique combinations of ($\tilde{\mathcal{P}}$, $\tilde{\mathcal{L}}$, $\tilde{\mathcal{R}}$), where each combination is repeated twice (Fig. 2.1.4d).
- **Run** (1 day): all data taken during one day of running, typically 15-25 superblocks (Fig. 2.1.4e). There are two types of data taken:
    1. Normal EDM data, where the value of $|\mathcal{B}_z|$ is changed between three values (0.7 mG, 1.3 mG, 26 mG) every superblock.
    2. Systematic checks where known systematic effects arising from experimental imperfections are amplified to monitor their effect on the EDM value, also known as *intentional parameter variation* (IPV). These are taken at a single chosen magnetic field magnitude $|\mathcal{B}_z^{\mathrm{IPV}}|$. There are five regularly performed IPVs (labeled A-E in the figure): A) P_{ref}, B) $\mathcal{E}^{\mathrm{nr}}$, C) $P^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$, D) $\phi_{\mathrm{ST}}^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$, and E) $\partial\mathcal{B}_z/\partial z$. The systematic effects that motivate these IPVs will be explained in §2.2.
- **Dataset**: all data taken during the ACME II final data run (about 40 runs taken over 76 days, Fig. 2.1.4f). On different days, $|\mathcal{E}_z|$ is switched between 80 V/cm and

140 V/cm on consecutive days, and $|\mathcal{B}_z^{\mathrm{IPV}}|$ is switched between the same three values used when taking normal data within a single day (0.7 mG, 1.3 mG, 26 mG).

Transformation to parity basis

Having run the experiment many times with different configurations of switches, we obtain values of the measured phase for each configuration $\Phi(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}, ...)$. It is possible to write this phase in terms of components that are odd with respect to certain switches or combinations of them. Assuming that there are only three switches $\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}$, we have

$$\Phi(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}) = \Phi^{\mathrm{nr}} + \Phi^{\tilde{\mathcal{N}}}\tilde{\mathcal{N}} + \Phi^{\tilde{\mathcal{E}}}\tilde{\mathcal{E}} + \Phi^{\tilde{\mathcal{B}}}\tilde{\mathcal{B}} + \Phi^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}\tilde{\mathcal{N}}\tilde{\mathcal{E}} + \Phi^{\tilde{\mathcal{N}}\tilde{\mathcal{B}}}\tilde{\mathcal{N}}\tilde{\mathcal{B}} + \Phi^{\tilde{\mathcal{E}}\tilde{\mathcal{B}}}\tilde{\mathcal{E}}\tilde{\mathcal{B}} + \Phi^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}\tilde{\mathcal{B}}}\tilde{\mathcal{N}}\tilde{\mathcal{E}}\tilde{\mathcal{B}}, \quad (2.9)$$

where the "nr" subscript denotes the component of Φ which is non-reversing, i.e. unaffected by any switching. The EDM term would be the term which is odd under reversal of $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$, $\Phi^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$. The magnetic (Zeeman) precession phase would be $\Phi^{\tilde{\mathcal{B}}}$. Other phase parity components can also be analyzed to look at other effects affecting the phase which are suppressed by the switch(es); these are commonly termed different "switch-parity channels" of the experiment. To transform from the state basis to the parity basis we use [113]

$$\Phi^p(\mathcal{N}, \mathcal{E}, \mathcal{B}) = \frac{1}{2^3} \sum_{\tilde{\mathcal{N}}', \tilde{\mathcal{E}}', \tilde{\mathcal{B}}' = \pm 1} (\tilde{\mathcal{N}}')^{\frac{1-\mathcal{N}}{2}} (\tilde{\mathcal{E}}')^{\frac{1-\mathcal{E}}{2}} (\tilde{\mathcal{B}}')^{\frac{1-\mathcal{B}}{2}} \Phi(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}}), \quad (2.10)$$

where $\Phi(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}})$ and $\Phi^p(\mathcal{N}, \mathcal{E}, \mathcal{B})$ denote the phase in the state and parity bases respectively. For the latter, the three numbers in parentheses denote whether Φ^p is odd or even with respect to a switch, e.g. $(-1, -1, +1)$ for $\Phi^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$, $(+1, +1, -1)$ for $\Phi^{\tilde{\mathcal{B}}}$, and so on. Most discussion of data analysis and systematic effects in ACME literature expresses the measured phases in terms of this switch-parity basis. Of course, equations 2.9 and 2.10 can be generalized to include more binary switches such as $\tilde{\mathcal{P}}$, $\tilde{\mathcal{L}}$, $\tilde{\mathcal{R}}$, etc. and applied to quantities other than the phase.

2.1.8 Data acquisition

In the next two sections, I will describe the ACME data acquisition (DAQ) and experimental control systems. Because it will be revisited in the context of ACME III (Chapter 6), we will cover these topics in more detail than previous sections. Besides myself, Brendon O'Leary, Adam West, and Cristian Panda all made important contributions to the development of these systems for ACME II.

DAQ hardware and timing

The laser-induced fluorescence data from the molecules is recorded using an 8-channel NI PXIe-5171 FPGA digitizer, which has 14 bits of resolution and a 250 MHz maximum sampling rate. The digitizer is triggered at 50 Hz alongside the ablation laser (with a $\approx$ 6.5 ms offset to account for the flight time from the beam source to the detection region). At each trigger, 10 ms of data is recorded from all 8 channels, each sampling at 16×10^6 samples/s with a dynamic range of ± 2 V.[7] (The PMT current is amplified to match the dynamic range.) It results in 40 samples recorded for every polarization switching bin (X or Y), which is sufficient to time-resolve the decays (as seen in Fig. 2.1.4a).[8] One record consists of 2000 pairs of $\hat{X}$-$\hat{Y}$ polarization switching bins as depicted in the figure. All fast triggers in the experiment (including the timing of the 200 kHz polarization switching) are coordinated in hardware with a digital delay generator.[9] It was later found that without syncing the digitizer's internal clock to an external reference and setting the sampling rate to an integer divisor of 250 MHz, there will be some timing jitter in the data acquisition. This issue will be discussed more in §6.1.

Fluorescence records from 25 consecutive shots are acquired in 16-bit integer format,

[7] This sampling rate is about 3× that of ACME I, and chosen based on the 115 ns lifetime of the I-state [152], which is $\sim 1/3$ the lifetime of the C state used for detection in ACME I [153].

[8] The temporal shape of the fluorescence profile is discussed in detail in Paul Hess' dissertation [153, §4.2.2].

[9] SRS DG645.

then averaged into a single trace and stored in 32-bit integer format in a network attached storage (NAS) drive.[10] This results in a data streaming rate of

$$\begin{aligned} DSR_{\text{ACMEII}} &= \frac{1}{25} \times (16 \times 10^6 \text{ samples/s}) \times (10 \times 10^{-3} \text{ s/record}) \times (32 \text{ bits/sample}) \\ &\quad \times (50 \text{ records/s/channel}) \times 8 \text{ channels} \\ &= 10.2 \text{ MB/s}. \end{aligned} \tag{2.11}$$

Software and hardware limitations made it impossible to save individual records from each molecular pulse, which would have required an order of magnitude faster DSR. (This issue will be revisited in Chapter 6.) However, for diagnostic purposes, it was possible to save individual records for short periods of time to a local solid state drive (SSD), which can support DSR $\sim 100 - 200$ MB/s. Over the course of about 2 years (2017-19), the ACME II experiment generated about 25 TB of data, consisting of 8 TB of final run data and the rest being data from systematic searches and other trial runs.

DAQ software

The code to run the DAQ system is programmed in LabView 2014. To ensure sufficient processing speed, the data acquisition and storage was performed with a separate computer than the main experiment control computer. The digitizer is housed in a PXIe chassis[11] which is connected to the DAQ computer. The DAQ and main control computers are coordinated by using a lossless network stream [154]. The digitizer constantly acquires records at 50 Hz which are always read by the DAQ computer. However, the data is only saved to storage when the main control computer indicates through the network stream that the experiment is ready, having set all the required switches and other parameters. Finally, the network stream is also used to notify the DAQ computer of the number of each

[10] Synology DS1817+, with 8 hard drives of 6 TB each in a RAID6 configuration, giving about 30 TB of storage. Every night, the data is also copied to an identical second NAS for backup and analysis purposes.

[11] NI PXIe-1075.

trace which is generated by the main control computer.

The aforementioned scheme was adopted after problems were found with the initial communication scheme, which used digitally generated signal from the main control computer to control a gate for the triggers to the digitizer. With this older scheme, we discovered that occasional lags of several milliseconds in the computer led to a mismatch between the intended number of triggers and the actual number of traces being acquired and fetched by the computer. This resulted in a mislabeling of the acquired traces, which was confusing for data analysis. This episode demonstrated that ensuring seamless timing and communication for the DAQ system is crucial.

Besides acquiring and saving the data from the eight PMT channels, the DAQ software also performs some basic analysis functions. First, traces from the 8 PMTs are summed into a single combined trace which is simultaneously saved alongside the separate PMT data. The majority of post-experiment data analysis focused on these summed traces which is more manageable in terms of file size. Analysis of signal from different PMTs was only required in special situations such as to look for effects that cause an asymmetric distribution of fluorescence between the East and West collection optics. Second, the asymmetry $\mathcal{A}$ throughout the trace is computed by grouping pairs of adjacent $\hat{X}$-$\hat{Y}$ polarization switching bins and using the formula in Eq. 2.4. This is performed and displayed while the experiment is running, providing a rough measurement of the phase of each measurement in real-time, which is sometimes useful for diagnostic purposes.

Spare DAQ system

In addition to the main DAQ device, a second spare DAQ system was developed and tested. This system consisted of four NI PXI-5922 modules (2 channels each, 15×10^6 samples/s), some of which were previously used in ACME I. Initially, the system did not meet required performance standards for ACME II while using default LabView software drivers, leading to the purchase of the 8-channel FPGA device described above. After I took over the

project, subsequent development and optimization of the software (including incorporation of some faster C++ code for data processing) made it viable. The majority of ACME II data was taken with the FPGA device. However, this second DAQ system was useful as a comparison when diagnosing DAQ-related issues in the experiment, such as in the diagnosis of the timing issues described above.

2.1.9 Experimental control, monitoring, and logging

An overall schematic of the control and data acquisition systems is shown in figure 2.1.5. The basic structure of the system was developed by Brendon O'Leary based on the ACME I system. Around the start of final commissioning of the ACME II apparatus (summer 2016), I took over the maintenance and further refining of the system to full running specifications.

The main control (MC) computer coordinates the running of the entire experiment using a series of LabView programs (called virtual instruments or VIs). During a typical run, a VI allows the user to configure and select the experiment switch sequence, which for a normal EDM data run usually consists of several superblocks taken at all three values of $|\mathcal{B}_z|$. The generated switch sequence is fed to the "Master Run VI" (MRVI). The MRVI loops through each combination of switch states and executes the programmatic sequence depicted in Fig. 2.1.6. First, it grabs the current set of switch states within the switch sequence. Second, it implements the switch states by sending commands to various VIs controlling the relevant devices – AOMs, waveplates, and current and voltage sources generating $\mathcal{E}$ and $\mathcal{B}$. The duration of this step depends on the switch. For example, the magnetic shields are degaussed before flipping $\tilde{\mathcal{B}}$, which takes several seconds. When the parameters are all properly set and the experiment is ready, the MRVI generates the number of the trace to be taken. This number is then sent to the DAQ computer along with a notification to start saving traces, using the aforementioned lossless network stream. At the same time, the MRVI saves the associated *header* for each trace (which will be explained below). The

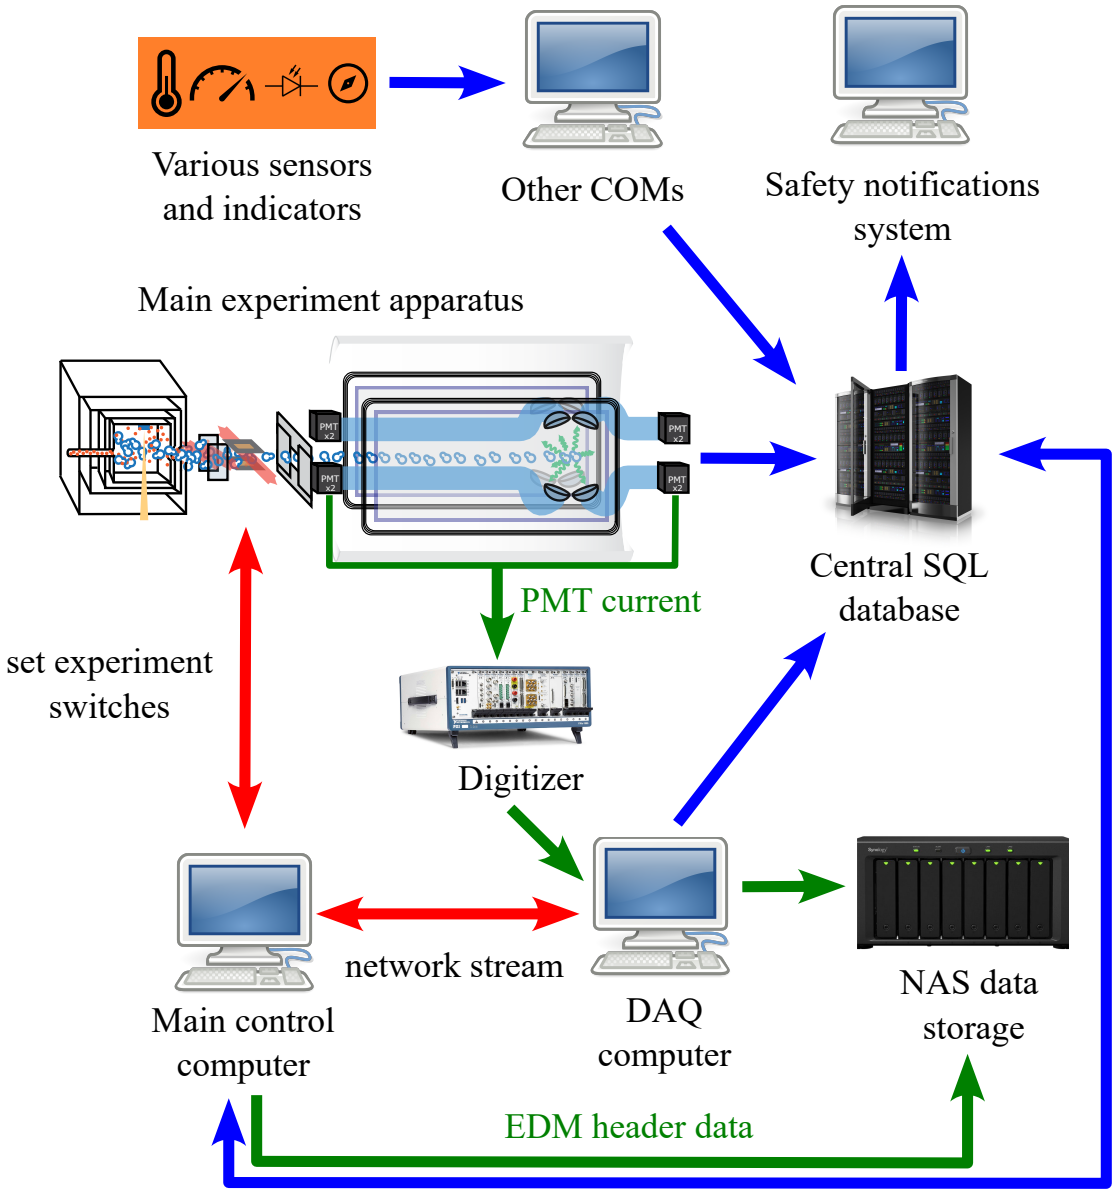

Figure 2.1.5: **ACME II experiment control and DAQ system.** Green arrows indicate flow of EDM data (including header data), blue arrows indicate the flow of auxiliary data (logged in the central database), and red arrows indicate connections for communication between devices.

DAQ computer then acquires data from 25 shots (500 ms), averages them into a single trace and saves it to the NAS. Finally, a notification to the MRVI on the ECC, after which the whole sequence can be repeated with the next set of switch states.

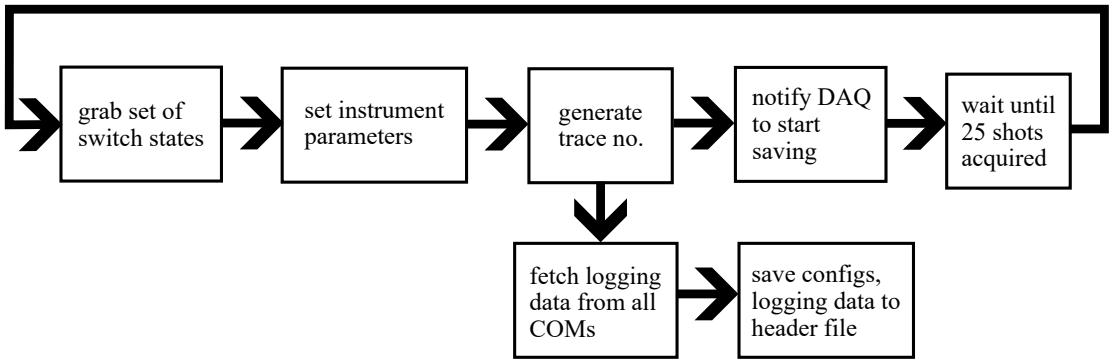

Figure 2.1.6: **Master Run VI program.** A simplified schematic of the procedure executed by the MRVI when running the experiment for a given switch state in an experiment switch sequence.

The header is an important part of the experiment control system, as it records crucial information for data analysis. A single header file is associated with each block, containing a single header for each of the 64 traces in the block. The header contains information about the switch states, DAQ parameters, instrument setpoints, and any auxiliary data associated with the experiment. This auxiliary data consists of a broad range of sensor, instrument, and environmental data collected by multiple computers in the experiment, such as beam source temperatures and pressures, magnetometer readings, laser locking statuses, laser power monitors, and so on – over 100 channels in total. When the experiment is not actively running (i.e. executing a switch sequence), this data is logged at slow intervals (> 5 s) to an SQL database hosted by a central server computer. When the experiment starts an active run, the logging of the data is sped up to $\lesssim 500$ ms intervals (the timescale

for a single averaged trace).[12] After it generates the trace number, the MRVI on the MC computer retrieves from the server all the auxiliary data gathered in the last 500 ms and inserts it into the header file. In this way, all auxiliary data that could be relevant to understanding the EDM data is gathered and consolidated in a single file stored in the same directory as the averaged fluorescence data. An example of a header can be found in Ref. [70, Appendix G].

For long-term monitoring, a set of safety VIs monitor certain critical parameters which must stay within a certain range (such as beam source temperatures and pressures) and sends email notifications if they stray outside this range. The SQL database also stores important configuration data for operating the instruments in the experiment and serves as a central hub for coordinating between VIs on different computers in the experiment.

2.1.10 Data analysis

Basic data analysis routine

An ACME data analysis routine comprises all of the steps needed to extract the EDM value d_e starting from the fluorescence traces acquired for each experiment switch configuration. Four analysis routines were written by different members of the ACME experiment. By design, each set of code contains slightly different choices on data cuts, binning, groupings, and so on, resulting in slightly different values of the extracted EDM, uncertainties, and systematic slopes. None of these differences affected the final result in a non-trivial way, giving confidence in the correctness of the analysis. Detailed descriptions of two analyses routines have already been given in the dissertations by Panda and Lasner [70, 148]. Here, I shall only sketch a summary of my own analysis routine, which has the same basic structure.

1. Start by taking a block of data where the averaged traces from all 8 PMTs have been summed into a single trace. Each block contains 64 traces of 160,000 samples each.

[12] For the data logged by multiple computers to be timestamped correctly, we ensure that the local computer clocks are synchronized to much better than 500 ms.

2. For each trace in the block, perform the following:

   (a) *Polarization switch binning*: partition each trace into $\hat{X}$ and $\hat{Y}$ polarization switching bin-pairs, which are 2×40 points each.

   (b) *Background subtraction*: take a portion of the trace before the signal (a certain number of $\hat{X} - \hat{Y}$ bin-pairs), average and then subtract from the entire trace. (It is assumed that the background is constant throughout the whole trace.)

   (c) *Signal integration*: integrate the signal within each polarization switching bin into a single number. Typically only the high-signal region within each bin (Fig. 2.1.4a) is integrated. This reduces the trace to ≈ 4000 points.

   (d) *Asymmetry calculation*: calculate the asymmetry $\mathcal{A}$ for each $\hat{X} - \hat{Y}$ bin-pair using Eq. 2.4. This reduces the trace to ≈ 2000 $\mathcal{A}$ values.

   (e) *Asymmetry grouping*: group a certain number of adjacent $\mathcal{A}$ values (usually about 20) together and take their mean and standard error. This error bar will then be propagated throughout the rest of the analysis. This results in the trace being reduced to ~ 100 averaged $\mathcal{A}$ values.

   (f) *Asymmetry filtering*: Filter out $\mathcal{A}$ values that are below a certain threshold. This throws away values in the low-signal regions of the trace (beginning and end), resulting in traces of $\sim$ 30-40 $\mathcal{A}$ values.

3. Group traces taken with the same switch-state combination and perform a weighted average for each group, resulting in 16 traces (due to the 4-fold degeneracy within a block).

4. Group traces according to $\tilde{\theta} = \pm 1$, resulting in two groups of 8 $(\tilde{\mathcal{N}}, \tilde{\mathcal{E}}, \tilde{\mathcal{B}})$ states. For each state, calculate the following (propagating uncertainties accordingly):

(a) The contrast:

$$\mathcal{C}(\tilde{\mathcal{N}},\tilde{\mathcal{E}},\tilde{\mathcal{B}}) = \frac{1}{2\Delta\theta}(\mathcal{A}(\tilde{\theta}=+1,\tilde{\mathcal{N}},\tilde{\mathcal{E}},\tilde{\mathcal{B}}) - \mathcal{A}(\tilde{\theta}=-1,\tilde{\mathcal{N}},\tilde{\mathcal{E}},\tilde{\mathcal{B}})), \tag{2.12}$$

where $\theta = 6^\circ$ is the full range of the readout laser polarization dither angle.

(b) The $\tilde{\theta}$-averaged, $\mathcal{C}$-corrected asymmetry:

$$\mathcal{A}(\tilde{\mathcal{N}},\tilde{\mathcal{E}},\tilde{\mathcal{B}}) = \frac{\mathrm{sign}(\mathcal{C})}{2}(\mathcal{A}(\tilde{\theta}=+1,\tilde{\mathcal{N}},\tilde{\mathcal{E}},\tilde{\mathcal{B}}) + \mathcal{A}(\tilde{\theta}=-1,\tilde{\mathcal{N}},\tilde{\mathcal{E}},\tilde{\mathcal{B}})) \tag{2.13}$$

(c) The phase:

$$\Phi(\tilde{\mathcal{N}},\tilde{\mathcal{E}},\tilde{\mathcal{B}}) = \frac{\mathcal{A}(\tilde{\mathcal{N}},\tilde{\mathcal{E}},\tilde{\mathcal{B}})}{2|\mathcal{C}(\tilde{\mathcal{N}},\tilde{\mathcal{E}},\tilde{\mathcal{B}})|} \tag{2.14}$$

5. Transform $\Phi(\tilde{\mathcal{N}},\tilde{\mathcal{E}},\tilde{\mathcal{B}})$ from the state basis to the parity basis using the formula in Eq. 2.10. This gives us 8 switch–parity channels: Φ^{nr}, $\Phi^{\tilde{\mathcal{N}}}$, $\Phi^{\tilde{\mathcal{B}}}$, $\Phi^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$, etc.

6. Calculate the precession time τ from the Zeeman precession phase $\Phi^{\tilde{\mathcal{B}}}$ by using the formula

$$\left(\Phi^{\tilde{\mathcal{B}}} + n\frac{\pi}{4}\right) = -g_1\mu_B\mathcal{B}_z\tau. \tag{2.15}$$

7. Calculate the precession frequencies $\omega = \Phi/\tau$ for each channel.

8. Consolidate each of the ω^p into single value by performing a weighted average of its values in time. This step can also be performed later if we wish to analyze the dependence of ω^p on molecule arrival time.

9. Add a blind offset to the EDM quantities $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$, $\Phi^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$. This blind is not revealed until the end of the experiment when all statistic and systematic error analyses have been completed.

10. Repeat steps 2-9 for each block in the superblock.

11. Perform a similar analysis at the superblock level, which has three switches ($\tilde{\mathcal{P}}$, $\tilde{\mathcal{L}}$, $\tilde{\mathcal{R}}$) and 2-fold degeneracy. Group and average together blocks taken with the same switch states, then convert to the parity basis using the same formula in Eq. 2.10.

After the last step, we have $2^6 = 64$ channels associated with each superblock for every combination of $\tilde{\mathcal{N}}$, $\tilde{\mathcal{E}}$, $\tilde{\mathcal{B}}$, $\tilde{\mathcal{P}}$, $\tilde{\mathcal{L}}$, $\tilde{\mathcal{R}}$. The $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ channel is the EDM value for that superblock, which we call one "EDM data point".

Extracting the EDM statistical mean and uncertainty

We perform the above analysis steps for each superblock in the dataset, including when searching for systematic errors. To extract the final average EDM value, all the EDM data points are grouped by the four $|\mathcal{B}_z|$ values. A mean and uncertainty for each set is computed using M-estimator analysis [155] performed on bootstrapped datasets [156]. We then perform a weighted mean of the four datasets, producing a single EDM mean value and statistical uncertainty.

In an experiment like ACME, the M-estimator method to calculate the mean is preferred instead of a simply assuming a Gaussian distribution because there could be non-Gaussian noise sources which may slightly change the distribution of the data especially in the tails. The M-estimator method prevents these outlier data points from inappropriately skewing the mean. This procedure is more exhaustively described in Ref. [148, §4.9.1].

Statistics of the dataset

The ACME II experiment was commissioned with full statistical sensitivity in fall 2016. After about a year of refining the system for consistency and robustness, searches for systematic errors began in fall 2017. The final data set was taken across 40 runs during January-March 2018, consisting of ∼1100 superblocks of normal EDM data with regularly interleaved systematic check data (see the experiment data structure described in §2.1.7).

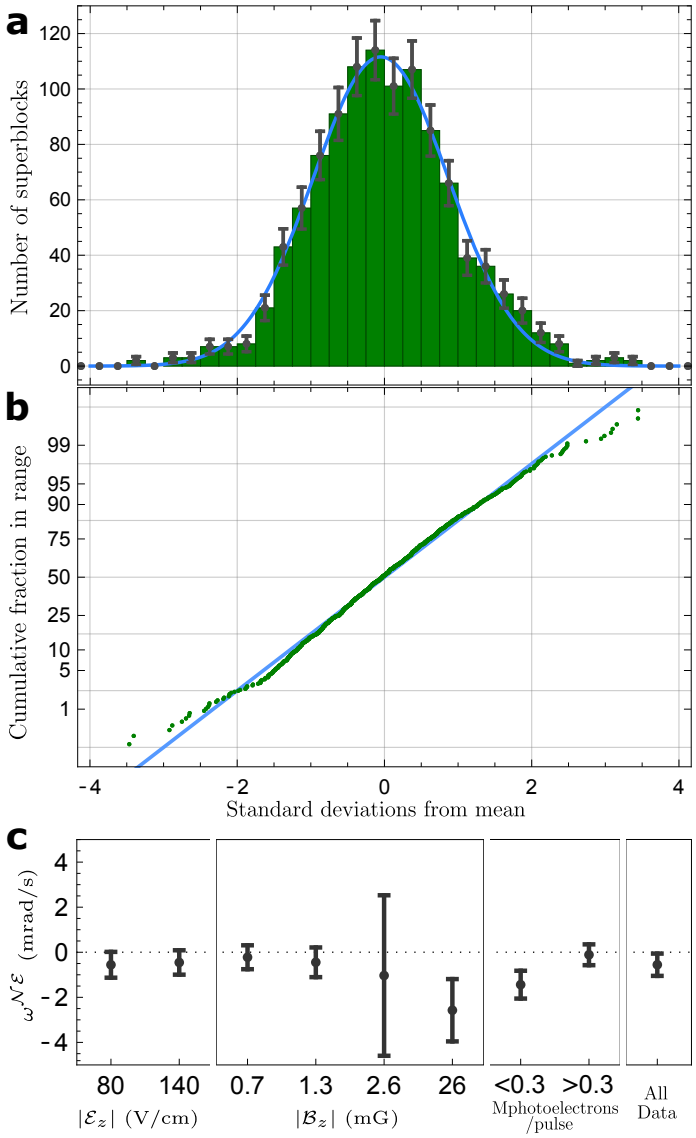

Figure 2.1.7: **Statistics of the ACME II dataset.** a) Histogram of measured EDM values $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ from each superblock. Data are normalized by uncertainty (scaled to reduced χ^2) and centered relative to the mean EDM value for the whole dataset. The blue line is a Gaussian fit to the histogram. b) Normal probability plot of the data (green points) compared with a normal distribution (blue line). Deviations from the normal distribution are observed in the tails. c) EDM values grouped by $|\mathcal{E}_z|$, $|\mathcal{B}_z|$, the block-averaged number of photoelectrons per pulse, and combined for all states, with error bars corresponding to 1σ. Figure is reproduced from [49].

This is equivalent to $T \sim 500$ hours of normal and systematic data which acquired over the course of 76 days, or a ∼20% duty cycle. This was a significant improvement over the $\sim 5\%$ duty cycle in ACME I.

The statistics of the dataset are summarized in the plots of Fig. 2.1.7. The measured EDM values follow a Gaussian distribution with some deviations in the tails. The majority of the data taken at $|\mathcal{B}_z| = 0.7,\ 1.3,\ 2.6$ mG have $\chi^2 \approx 3$, later discovered to be caused by timing asymmetry noise (§2.2.9). The small amount of data taken at $|\mathcal{B}_z| = 26$ mG was also affected by velocity noise (§2.2.8) and has $\chi^2 \approx 7$. During the entire data analysis, $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ channel was blinded by an offset which was not revealed until the main systematic data analysis was concluded.

2.1.11 Magnetic fields

Here we shall describe the magnetic fields in the apparatus, in order to provide useful background for the discussion of the same in Chapter 7. A fuller account of the features and characterization of the field can be found in Ref. [70, §5.2].

Magnetic shields

The entire interaction region is shielded by five concentric layers of cylindrical mu-metal magnetic shields. Each layer consists of two flat circular endcaps and two removable half-cylinders. Each half-cylinder is wrapped by a pair of degaussing coils. Degaussing is performed each time $\tilde{\mathcal{B}}$ is switched with a 100 Hz waveform which lasts for ∼1 s. When the ACME II apparatus was initially commissioned, ambient fields of ≈ 50 μG. However, after about a year of running, the ambient field was found to increase to ≈ 300 μG. After removing some slightly magnetic components and performing better degaussing, this was reduced to ≈ 150 μG. A possible explanation for this long-term degradation in shielding performance is unintended bumping of the shields.

Magnetic coils

The magnetic field coils are able to apply constant and gradient fields along all three axes in the interaction region. There are nominally three possible constant fields ($\mathcal{B}_x$, $\mathcal{B}_y$, $\mathcal{B}_z$) and nine first-order gradients ($\partial\mathcal{B}_x/\partial x$, $\partial\mathcal{B}_y/\partial y$, $\partial\mathcal{B}_x/\partial y$, etc.). However Maxwell's equations

$$\nabla \cdot \vec{\mathcal{B}} = 0 \implies \frac{\partial \mathcal{B}_x}{\partial x} + \frac{\partial \mathcal{B}_y}{\partial y} + \frac{\partial \mathcal{B}_z}{\partial z} = 0 \tag{2.16}$$

and

$$\nabla \times \vec{H} = \vec{J} = 0 \implies \begin{cases} \frac{\partial \mathcal{B}_x}{\partial y} = \frac{\partial \mathcal{B}_y}{\partial x} \\ \frac{\partial \mathcal{B}_x}{\partial z} = \frac{\partial \mathcal{B}_z}{\partial x} \\ \frac{\partial \mathcal{B}_y}{\partial z} = \frac{\partial \mathcal{B}_z}{\partial y} \end{cases} \tag{2.17}$$

imply that there are only five independent components to be applied; the rest can be inferred from the above equations. Thus, the five sets of coils depicted in figure 2.1.8 and described in Table 2.1.1 are able to apply magnetic gradients in all directions. The main $\mathcal{B}_z$ field is produced by a saddle-shaped cosine theta coil (orange) with shim coils (red) to improve the homogeneity of the field and allow the production of $\partial\mathcal{B}_z/\partial x$. These coils can also be used to produce $\partial\mathcal{B}_z/\partial z$. $\mathcal{B}_y$ and associated gradients are produced by four sets of coils above and below the interaction region (blue and green), and $\mathcal{B}_x$, $\partial\mathcal{B}_x/\partial x$ are produced by a pair of circular Helmholtz coils (yellow). The coils consist of resin-coated copper wires installed in grooves of an HDPE frame, with the exception of the $\mathcal{B}_x$ coils which use ribbon connectors. The main $\mathcal{B}_z$ coils are powered with a commercial power supply[13] and all other coils use a custom power supply constructed by Jim MacArthur at Harvard.

Field probulation using fluxgate magnetometers performed by Zack Lasner and Adam West found the applied $\mathcal{B}_z$ field homogeneity to be $\approx 1\%$ [70, §5.2.4]. While taking the final

[13] Krohn-Hite 521/522.

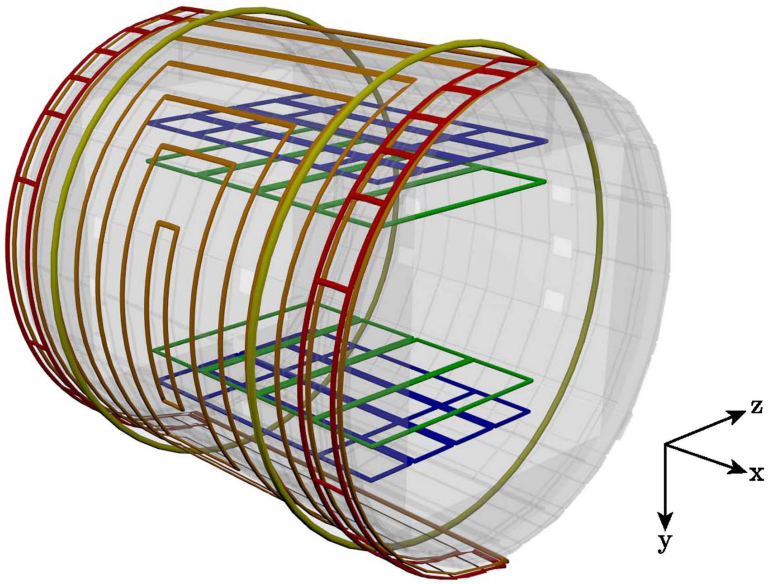

Figure 2.1.8: **ACME II magnetic coils.** Schematic of magnetic coils in ACME II apparatus used to produce constant and gradient magnetic fields along all three axes. Figure reproduced from Ref. [113], licensed under Creative Commons 3.0.

Coil colour	Fields produced	Field gradients produced
Orange	$\mathcal{B}_z$	$\partial\mathcal{B}_z/\partial z$
Red	$\mathcal{B}_z$	$\partial\mathcal{B}_z/\partial x$, $\partial\mathcal{B}_z/\partial z$
Yellow	$\mathcal{B}_x$	$\partial\mathcal{B}_x/\partial x$
Blue	$\mathcal{B}_y$	$\partial\mathcal{B}_y/\partial y$, $\partial\mathcal{B}_y/\partial z$
Green	$\mathcal{B}_y$	$\partial\mathcal{B}_y/\partial x$, $\partial\mathcal{B}_y/\partial y$

Table 2.1.1: Magnetic fields and gradients produced by individual coil sets in figure 2.1.8.

dataset, $|\mathcal{B}_z| = 0.7$, 1.3, 2.6, 26 mG were applied with the main and shim coils. During systematic checks, constant fields of up to 12.5 mG and field gradients up to 2 mG/cm are applied. During the final EDM dataset, the magnetic coils are used to cancel the systematic-error-contributing field gradients (§2.2.5) to < 10 μG/cm.

2.2 Systematic errors and excess noise

Assuming that the spin prcession measurement as detailed in the previous section is shot-noise limited, the uncertainty in the EDM will follow Eq. 1.23. This uncertainty can be increased by systematic errors (shifts in the EDM channel $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ caused by effects other than an electron EDM) and/or excess noise (effects that increase statistical uncertainty but do not cause a shift in $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$). The ACME II experiment was designed to be resilient against systematic errors. Many mechanisms causing phase shifts are filtered out due to the co-magnetometer setup of the $\tilde{\mathcal{N}}$-state – only mechanisms that are odd with respect to $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ can become a systematic error. Specific experimental imperfections known to cause systematic shifts in ACME I were redesigned or improved in ACME II. In this section, we will provide a summary of the major sources of systematic error and excess noise in ACME II, focusing especially on the largest contributors to the systematic error budget which were focal concerns in the development of ACME III. More comprehensive accounts of the systematic error search and characterization can be found in prior publications [49, 70, 148].

2.2.1 General procedure to search for systematics

A comprehensive systematic error search was undertaken before the final ACME II data run, which involved varying about 40 experimental parameters and seeing the effect on $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ and other channels. These include magnetic field gradients, non-reversing electric fields, laser detunings, powers, pointings, polarizations, and others (a full list is available in [148, Table 5.1.1]). Many of these parameters have an ideal value of zero (such as magnetic

field gradients).[14] For such a parameter X, we increase its magnitude by about an order of magnitude compared to its value X_{norm} during normal operating conditions. We then assume a linear relationship between $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ and X and fit the data to find the systematic slope $\partial\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}/\partial X$. The systematic error would then be $\omega_X^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}} = (\partial\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}/\partial X)X_{\text{norm}}$. Typically, we would take about one day of data for each X, and perform further investigations if $|(\partial\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}/\partial X)| > 0$ is observed with greater than 2σ significance. For such cases, we would attempt to understand the systematic mechanism and attempt to suppress the slope.

A table of all the systematic error contributions included in the ACME II systematic uncertainty is shown in Table 2.2.1. There are two main classes of systematic shifts. The most important category (Class A) are parameters which are known to cause a statistically significant shift in $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ if their value is non-zero. For these, the value of the slope is constantly monitored during the final run and the value of the shift is measured and subtracted from $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$. The second category (Class C) are parameters for which we did not observe a statistically significant shift, but have historically caused a systematic shift in other EDM experiments. For these, a systematic shift was not subtracted, but they were included in the systematic uncertainty budget.[15] Here and in the next chapter, we will focus our attention on Class A systematic effects, especially those which result in the largest systematic uncertainties. In the final analysis, all the shifts in Table 2.2.1 are subtracted from the mean EDM value (computed using the method described in §2.1.10) and their uncertainties added in quadrature to become the total systematic uncertainty.

[14] Other parameters (such as polarization switching frequency) have no ideal value, and no statistically significant shifts in $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ were observed when varying any of them. As such, we do not include them in the systematic error budget.

[15] There is another category of systematic errors (Class B) which was used in ACME I [113]. These occur if there is an unexplained signal in a channel which is not correlated to $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ but is deemed important to understand. No systematic effect of this class was observed in ACME II.

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

Table 2.2.1: **ACME II systematic shifts and their statistical uncertainties.** Units are in μrad/s. Table adapted from [49, 148]. Ordering of the parameters has been modified to reflect the order of discussion in this chapter.

2.2.2 The $\mathcal{E}^{\mathrm{nr}} \to \Delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ mechanism

Systematic effects in ACME are typically caused by a combination of mechanisms which influence the spin precession frequency ω. These mechanisms would not be a systematic unless they couple to a parameter that is $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-odd. One of the central mechanisms to produce such a parameter is the $\mathcal{E}^{\mathrm{nr}} \to \Delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ mechanism, where a non-reversing electric field $\mathcal{E}^{\mathrm{nr}}$ can give rise to an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-odd laser detuning. Such an $\mathcal{E}^{\mathrm{nr}}$ can arise due to effects such as patch potentials on the ITO coating of the field plates [157]. In the ACME II field plates, this was measured to be $\mathcal{E}^{\mathrm{nr}} = -2.6(1.6)$ mV/cm. Thus, the lab electric field $\mathcal{E}$ becomes

$$\mathcal{E} = \mathcal{E}_0\tilde{\mathcal{E}} + \mathcal{E}^{\mathrm{nr}}, \tag{2.18}$$

resulting in the level structure shown in Fig. 2.2.1, which shows the detuning from $|H, J = 1\rangle$ to an excited state f.

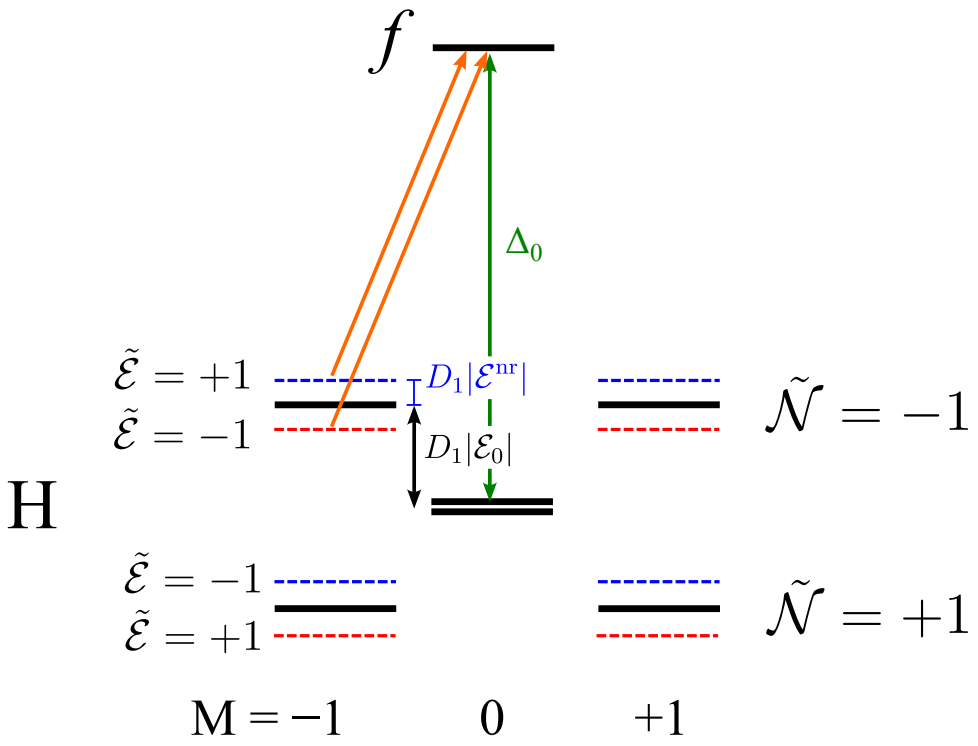

Figure 2.2.1: $\mathcal{E}^{\rm nr} \rightarrow \Delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ **systematic mechanism.** A non-reversing electric field $\mathcal{E}^{\rm nr}$ creates $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-odd shifts in the structure of the H-state, which results in the $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-odd parameter $\Delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ when molecules are excited with a laser to state f. This mechanism couples to other imperfections in the apparatus to create shifts in $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$.

The combined detuning between the two levels is

$$\Delta = \Delta_0 + \tilde{\mathcal{N}}(D_1|\mathcal{E}_0| - \Delta^{\tilde{\mathcal{N}}}_{\rm AOM}) + \tilde{\mathcal{N}}\tilde{\mathcal{E}}D_1|\mathcal{E}^{\rm nr}|, \tag{2.19}$$

where Δ_0 is the component of the detuning which does not change with any of the switches, and $\Delta^{\tilde{\mathcal{N}}}_{\rm AOM}$ is the shift in the AOM when the the $\tilde{\mathcal{N}}$ switch is flipped. The resulting $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-dependent detuning,

$$\Delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}} = \tilde{\mathcal{N}}\tilde{\mathcal{E}}D_1|\mathcal{E}^{\rm nr}| \tag{2.20}$$

is present in all instances in the interaction region where we drive a transition between H and an excited state: the STIRAP Pump and Stokes lasers, refinement laser, and readout

laser.

2.2.3 AC Stark shift systematic

The first systematic is an AC Stark shift effect arising from ellipticity gradients along $\hat{x}$ in the polarization of the refinement and readout lasers. The refinement laser beam has an approximately Gaussian intensity profile along $\hat{x}$. When molecules traverse the more intense portions, they are pumped out of the bright state, leaving a dark state that is intended to be the initial spin precession state $|\psi(t=0)\rangle$ (§2.1.4, Eq. 2.1). However, when they traverse the dimmer tails of the laser, they are not completely pumped out of the bright state. At the same time, the ellipticity gradient causes the light to interact with the dark state and cause an AC Stark shift. A similar effect occurs in the dimmer tail of the readout laser within the precession region. Altogether, this gives rise to a non-zero dependence of the precession frequency on the detuning, $\partial\omega/d\Delta$. Combined with the $\Delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ due to the $\mathcal{E}^{\mathrm{nr}} \rightarrow \Delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ mechanism (Eq. 2.19), it results a systematic shift in $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$. This was a known effect in ACME I and is described in more detail in previous publications [83, 113, 146].

The ellipticity gradient in the laser beam arises from birefringence gradients in the optical elements in the path of the laser before it reaches the molecule: a field plate, vacuum window, lenses, etc. These gradients can be precisely measured and monitored by using a polarimeter [158, 159]. In ACME I, a major contributor was thermal stress-induced birefringence in the field plates caused by absorption of high-powered laser beams [113]. To suppress this effect, the ACME II field plates were designed with a different glass[16] which has an order of magnitude lower thermal expansion [160]. This effect is also the reason why STIRAP is performed with laser beams sent vertically between the field plates instead of horizontally.

A second source of birefringence gradients are the side vacuum windows in the preparation and readout regions, which are most likely the result of mechanical (rather than

[16] Corning 7980 instead of Schott Borofloat.

thermal) stress birefringence. In ACME II, the total polarization gradient when sending in laser beams from the East side of the experiment (including both field plate and vacuum window) was $\lesssim 0.1\%/\text{mm}$.[17] Sending in the lasers from the West side of the experiment revealed a much larger ∼1-2%/mm gradient [148, §5.3.1]. This asymmetric result is best explained by the presence of mechanical stress birefringence on one side but not the other (in the field plates and/or the vacuum windows). Because of this asymmetry, the ACME II final data run was conducted by sending in the laser beams from the East side to suppress this systematic. The systematic was monitored daily during the final run by taking a few superblocks of data with an artificially large value of $\mathcal{E}^{\text{nr}}$. The regular $\mathcal{E}^{\text{nr}}$ in the experiment was also measured regularly every two weeks with microwave spectroscopy (described in more detail in [113, 148]). This led to the systematic shift of $-56(140)$ μrad/s in Table 2.2.1.

2.2.4 Other laser-related systematics

The vertical STIRAP laser beams are associated with two systematic effects.[18] The first effect arises from ellipticity gradients (similar to the effect the refinement and readout lasers described above), which is large in the STIRAP beams because they are focused to a very small waist [146, §4.3.4]. The second effect arises from "Stark interference" between E1 and M1 transitions [113, 146].[19] However, both of these systematics are suppressed by the refinement laser, resulting in the small $\omega_{\text{ST}}^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ shift of 0(1) μrad/s in Table 2.2.1.

Another AC Stark shift-related systematic is the $P_{\text{ref}}^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ shift. A misalignment $\theta_{\text{ref}}^{\text{ST}}$ between the orientation of the STIRAP-prepared state $\vec{S}_{\text{ST}}$ and the polarization of the refinement laser $\vec{\epsilon}_{\text{ref}}$ causes imperfect refinement of the initial spin state i.e. a non-zero slope $\partial\omega/\partial P_{\text{ref}}$. If the power of the refinement laser has an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-correlated component $P_{\text{ref}}^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}} > 0$, this slope

[17] The number denotes the percentage of circular polarization or S/I using the Stokes parameters [158].

[18] Because the STIRAP beams propagate along $\hat{y}$ and the spin precession occurs in the xy plane, unlike the case of the refinement and readout lasers, the effect of polarization gradients caused by the vacuum windows is insignificant. I thank Peiran Hu for our discussions to clarify this.

[19] A detailed explanation of Stark interference can be found in Ref. [146, §4.2].

leads to a systematic shift in $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$. A way to suppress this shift is to reduce $\partial\omega/\partial P_{\mathrm{ref}}$ by tuning $\vec{\epsilon}_{\mathrm{ref}}$ until $\theta_{\mathrm{ref}}^{\mathrm{ST}}$ is consistent with zero and increasing P_{ref} to saturate the optical pumping as much as possible. Another method is to reduce the $P_{\mathrm{ref}}^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ itself. However, in ACME II there was no direct evidence of such a component in P_{ref}. Still, ACME I had found that AC Stark shift effects could lead to a linear dependence of the $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}\tilde{\mathcal{B}}}$ channel on $P_{\mathrm{ref}}^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ [113]. The offset in this channel was used as a measure of $P_{\mathrm{ref}}^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ and used to set the 109 μrad/s in Table 2.2.1.

2.2.5 Magnetic field gradients

A second major source of systematics came from magnetic field gradients $\partial\mathcal{B}_z/\partial z$ and $\partial\mathcal{B}_z/\partial y$. For each type of gradient, this effect comprised of two different mechanisms, one which coupled to a constant $\mathcal{E}^{\mathrm{nr}}$ and another to the $\mathcal{E}^{\mathrm{nr}}$ gradient. The following systematic models were developed by Cris Panda and Zack Lasner in ACME II, and verified both experimentally and through numerical simulations.

First mechanism: coupling to constant $\mathcal{E}^{\mathrm{nr}}$

The first mechanism[20] results from the coupling of three physical effects. First, a $\partial\mathcal{B}_z/\partial z$ causes a precession frequency gradient $\partial\omega/\partial z$ (Fig. 2.2.2). Second, the molecules possess a strong position-transverse velocity correlation ($\partial v_z/\partial z$) which is a natural result of their ballistic trajectories from the beam source: faster molecules are more likely to be further away from $z = 0$. When the molecules traverse the readout laser beam (which propagates along $\hat{z}$), this results in a spatial *detuning* gradient $\partial\Delta_z/\partial z$: molecules further away from $z = 0$ have a larger Doppler shift along $\hat{z}$. Finally, the $\mathcal{E}^{\mathrm{nr}} \rightarrow \Delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ mechanism gives rise to an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-odd detuning $\Delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ as previously described (Eq. 2.19). These three effects couple

[20] In [148] this is called the second mechanism, but for this thesis I have decided to switch the numbering to provide a more natural flow of the explanation.

together to give

$$\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}_{\partial\mathcal{B}_z/\partial z,\ \mathcal{E}^{\mathrm{nr}}} = \left(g_1\mu_B\frac{\partial\mathcal{B}_z}{\partial z}\right)\left(\frac{1}{\frac{\partial v_z}{\partial z}\frac{\partial\Delta_z}{\partial v_z}}\right)\left(\frac{\partial\Delta_z^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}}{\partial\mathcal{E}^{\mathrm{nr}}}\mathcal{E}^{\mathrm{nr}}\right) = \frac{\partial\omega}{\partial z}\frac{\partial z}{\partial\Delta_z}\Delta_z^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}} \neq 0. \tag{2.21}$$

The same argument applies for $\partial\mathcal{B}_z/\partial y$ and produces a similar shift with the vertically propagating STIRAP laser beams.

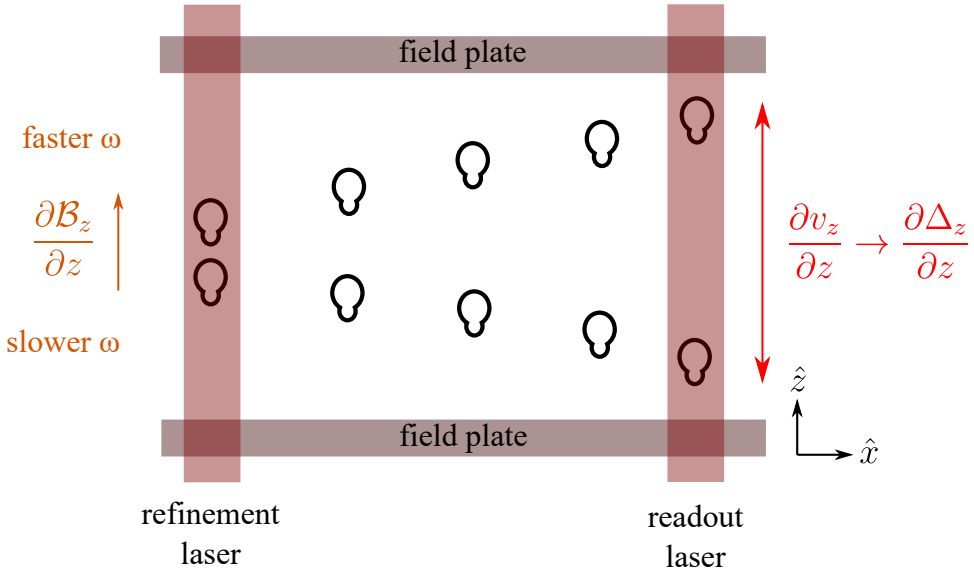

Figure 2.2.2: **First magnetic gradient systematic mechanism: coupling to $\mathcal{E}^{\mathrm{nr}}$.** A $\partial\mathcal{B}z/\partial z$ gradient leads to a precession frequency gradient $\partial\omega/\partial z$. A position-velocity gradient leads to a readout laser detuning gradient $\partial\Delta_z/\partial z$. Finally, the presence of an $\mathcal{E}^{\mathrm{nr}}$ leads to $\Delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$. Together, the three effects combine to give rise to a systematic shift in $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$. Besides the readout laser, this mechanism is also present for the vertically propagated STIRAP laser with the $\partial\mathcal{B}_z/\partial y$ gradient.

Second mechanism: coupling to $\mathcal{E}^{\mathrm{nr}}$ gradient

The second mechanism is also a product of three physical effects (Fig. 2.2.3). The first is the same as before: a $\partial\mathcal{B}_z/\partial z$ causes a precession frequency gradient $\partial\omega/\partial z$. The second effect is the dependence of STIRAP transfer efficiency η on the two-photon detuning δ

(§2.1.4), $\partial\eta/\partial\delta$, which is non-zero when δ is slightly off-resonant. The third is a analogous to the familiar $\mathcal{E}^{\mathrm{nr}} \rightarrow \Delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ mechanism, but applied to a gradient in $\mathcal{E}^{\mathrm{nr}}$. In this case, $\partial\mathcal{E}^{\mathrm{nr}}/\partial z$ gives rise to an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-correlated STIRAP detuning gradient $\partial\delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}/\partial z$. The second and third effects give rise to an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-correlated shift in the center of mass of the molecules that are prepared by STIRAP [49]:

$$dz_{\mathrm{CM}}^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}} = \frac{a^2}{3\eta_0}\frac{\partial\eta}{\partial\delta}\frac{\partial\delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}}{\partial z}, \tag{2.22}$$

where a is the half-width of the molecular beam and η_0 is the maximum STIRAP efficiency at $\delta = 0$. (Here a simple model where the molecular density is constant along $\hat{z}$ is assumed.) This population shift was confirmed by comparing the fluorescence from the East and West detectors when δ was increased.

Together with the first effect, this results in a systematic shift

$$\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}_{\partial\mathcal{B}_z/\partial z,\ \partial\mathcal{E}^{\mathrm{nr}}/\partial z} = \left(g_1\mu_B\frac{\partial\mathcal{B}_z}{\partial z}\right) dz_{\mathrm{CM}}^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}} = \frac{\partial\omega}{\partial z}\frac{a^2}{3\eta_0}\frac{\partial\eta}{\partial\delta}\frac{\partial\delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}}{\partial z} \neq 0. \tag{2.23}$$

Notice that unlike the first mechanism (magnetic gradients coupling to constant $\mathcal{E}^{\mathrm{nr}}$), this mechanism does not depend on a Doppler shift along the laser beam direction. Thus, it also occurs for $\partial\mathcal{B}_z/\partial y$ gradients coupling to $\partial\mathcal{E}^{\mathrm{nr}}/\partial y$. Furthermore, a similar mechanism exists for the readout laser beam for both $\partial\mathcal{B}_z/\partial z$ and $\partial\mathcal{B}_z/\partial y$, assuming it has a non-zero detuning Δ. Thus, there are four contributions to this second mechanism, although the STIRAP contribution is larger. The total shift introduced by the second mechanism is about several times larger than the first mechanism.

Suppression and monitoring

Both of the systematic mechanisms above depend on $\partial\mathcal{B}_z/\partial z$ and $\partial\mathcal{B}_z/\partial y$. They were suppressed by tuning the STIRAP and readout lasers such that the systematic slopes

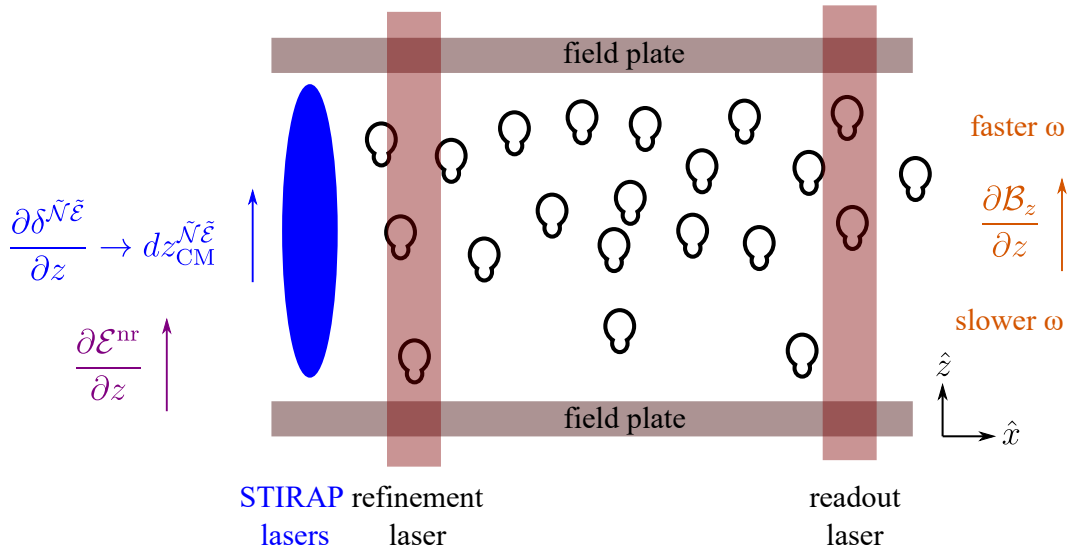

Figure 2.2.3: **Second magnetic gradient systematic mechanism: coupling to $\partial\mathcal{E}^{\mathrm{nr}}/\partial z$.** The presence of an $\partial\mathcal{E}^{\mathrm{nr}}/\partial z$ leads to a detuning gradient $\partial\Delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}/\partial z$, which together with a slightly off-resonant STIRAP 2-photon detuning $\delta \neq 0$ produces an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-odd shift in the center of mass of the prepared population $dz_{\mathrm{CM}}^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$, as seen in the greater number of molecules along $+\hat{z}$ in the figure. This combines with the precession frequency gradient $\partial\omega/\partial z$ caused by a $\partial\mathcal{B}_z/\partial z$, resulting in a systematic shift in $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$. An analogous mechanism applies for $\partial\mathcal{B}_z/\partial y$, as well as for the readout laser for both magnetic field gradients.

$\partial\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}/\partial(\partial\mathcal{B}_z/\partial z)$ and $\partial\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}/\partial(\partial\mathcal{B}_z/\partial y)$ are consistent with zero. As Lasner [70, p. 164] explains, this is effectively using the two mechanisms to cancel each other.

For the final data run, the magnetic gradients in the interaction region were minimized to (-1 ± 9) μG/cm by using the magnetic field coils and regularly monitored by inserting magnetometers in a pocket right next to the molecular beam. A more thorough characterization was performed before and after the final data run by inserting a magnetometer through the center of the spin precession region along $\hat{x}$ [70]. In the final data run, datasets taken with exaggerated magnetic field gradients were regularly interleaved between normal EDM datasets. This resulted in the 7(59) μrad/s shift in Table 2.2.1.

2.2.6 Contrast correlations

The next systematic shift arises from observed correlations between $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ and components of the contrast magnitude $|\mathcal{C}|$.[21] Recall that in the data analysis procedure described in §2.1.10, we calculate

$$\omega = \frac{\mathcal{A}}{2\tau|C|} \tag{2.24}$$

from which we extract the EDM shift by taking the $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-odd component

$$\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}} = \left(\frac{\mathcal{A}}{2\tau|C|}\right)^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}. \tag{2.25}$$

However, there are higher-order $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-odd components of $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ which can be seen as follows. First, we expand Eq. 2.24 in terms of all its components in the parity basis $\mathcal{A}^p$, $|\mathcal{C}|^p$:

$$\omega = \frac{\sum_p \mathcal{A}^p}{2\tau \sum_q |\mathcal{C}|^q} = \frac{\sum_p \mathcal{A}^p}{2\tau|\mathcal{C}|^{\mathrm{nr}}} \left(\frac{1}{1+\sum_{q\neq \mathrm{nr}} \frac{|\mathcal{C}|^q}{|\mathcal{C}|^{\mathrm{nr}}}}\right), \tag{2.26}$$

then we use the fact that $|\mathcal{C}^{\mathrm{nr}}| \gg |\mathcal{C}^r|$, performing a Taylor expansion to obtain

$$\omega \approx \frac{\sum_p \mathcal{A}^p}{2\tau|\mathcal{C}|^{\mathrm{nr}}} \left(1 - \sum_q \frac{|\mathcal{C}|^q}{|\mathcal{C}|^{\mathrm{nr}}}\right). \tag{2.27}$$

Taking the $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-odd components only, we get

$$\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}} \approx \frac{\mathcal{A}^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}}{2\tau|C|^{\mathrm{nr}}} - \omega^{\mathcal{B}}\frac{|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}\tilde{\mathcal{B}}}}{|\mathcal{C}|^{\mathrm{nr}}} - \omega^{\mathrm{nr}}\frac{|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}}{|\mathcal{C}|^{\mathrm{nr}}} - \ldots \tag{2.28}$$

where apart from the main $\mathcal{A}^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ term, only the largest terms which arise from $\omega^{\mathcal{B}}$ and ω^{nr} are shown. These correspond to the Zeeman precession frequency and any global offset phases due to a misalignment between the initial prepared spin state and the readout polarization, i.e. $(\phi - \theta) - \frac{\pi}{4}(2n+1) \neq 0$ (see Eq. 2.6). In ACME II data, $|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}\tilde{\mathcal{B}}}$ and

[21] For this section, I am indebted to recent discussions with Peiran Hu.

$|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ were unsurprisingly the only two contrast channels for which correlations with $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ were observed, with $\partial\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}/\partial|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}\tilde{\mathcal{B}}}$ being ≈ 6 times larger than $\partial\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}/\partial|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ [70, p. 153]. No physical mechanisms are known to give rise to a non-zero $|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}\tilde{\mathcal{B}}}$ and $|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ and their average values across the dataset were consistent with zero. However, to be conservative, the decision was made to subtract their contributions and include the uncertainties in the total uncertainty budget, leading to the 77(125) μrad/sec shift in Table 2.2.1.

2.2.7 $\omega^{\tilde{\mathcal{E}}}$ shift

If there is a shift in the $\omega^{\tilde{\mathcal{E}}}$ channel, this could potentially lead to a a shift in $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ if there is a slope $\partial\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}/\partial\omega^{\tilde{\mathcal{E}}}$. This slope exists because of the small, $|\mathcal{E}|$-dependent difference in magnetic moments between the two $\tilde{\mathcal{N}} = 1$ states (Eq. 1.19). A mechanism to cause the $\omega^{\mathcal{E}}$ shift are leakage currents in the field plates, which were measured to be $I_{\text{leakage}} = 25$ pA [148, §5.8]. This value was used to determine the $\omega^{\mathcal{E}}$ shift of 1(1) μrad/s.

2.2.8 Velocity noise

Molecular beam velocity noise was a major issue in the ACME II experiment. The ablation process produces molecular beams with a mean velocity of ≈ 200 m/s which fluctuated $\approx 0.1\%$ from shot-to-shot fluctuations, leading to fluctuations in precession time τ and thus the measured phase Φ. These fluctuations do not average out because τ is measured by switching $\tilde{\mathcal{B}}$ to measure the Zeeman phase $\Phi^{\tilde{\mathcal{B}}}$, which only occurs once every block ($\approx$ 1 minute). The resulting noise scales with $|\mathcal{B}|$. When running at $|\mathcal{B}_z| = 26$ mG, the data has $\chi^2 \approx 7$. At smaller fields ($|\mathcal{B}_z| = 0.7,\ 1.3,\ 2.6$ mG) the data has $\chi^2 \approx 3$.

An attempt was made to investigate and possibly reduce this noise by measuring τ directly. A small "notch" in the molecules was induced by turning off the STIRAP Pump laser beam for a moment (100 μs) and measuring the position of the resulting dip in the detected fluorescence trace [148, §4.12]. However, this method was not able to measure the velocity in a trace to better than 0.1% because of the velocity dispersion within the notch.

This dispersion exists because molecules do not all leave the beam box at the same time, resulting in imperfect correlation between their velocity and precession times. Moreover, the notch method also resulted in loss of signal. In the end, this noise was suppressed by simply taking the majority of the final dataset at lower values of $|\mathcal{B}_z|$.

2.2.9 Timing noise

The last major source of excess noise is asymmetry noise arising from jitter in the timing of the data acquisition, due to a mismatch between the DAQ internal clock and sampling rate. This noise was not fully understood until after the conclusion of ACME II. As a result, a fuller account of the mechanism of this noise and its suppression is given in §6.1 as well as in a published paper [161]. After this setting was fixed, the experiment noise was reduced to the shot noise limit (Eq. 1.23).

2.3 Results and interpretation

After subtracting out all systematic shifts listed in Table 2.2.1 and performing the data analysis procedures described in 2.1.10, the final mean EDM precession frequency including statistical and systematic uncertainties was found to be

$$\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}} = (-510 \pm 373_{\text{stat}} \pm 310_{\text{syst}})\ \mu\text{rad/s}, \tag{2.29}$$

which we can convert into d_e by using $\mathcal{E}_{\text{eff}} = 78$ GV/cm (§1.5.2) and the relation $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}} = d_e \mathcal{E}_{\text{eff}}/\hbar$, giving

$$d_e = (4.3 \pm 3.1_{\text{stat}} \pm 2.6_{\text{syst}}) \times 10^{-30}\ e \cdot \text{cm} \tag{2.30}$$

$$= (4.3 \pm 4.0) \times 10^{-30}\ e \cdot \text{cm}, \tag{2.31}$$

where the combined statistical and systematic uncertainty is a factor of 12 smaller than the ACME I uncertainty [96, 113]. This result is consistent with a null value of d_e. It can be converted to into an upper limit by applying the Feldman-Cousins prescription [147, 162], yielding

$$|d_e| < 1.1 \times 10^{-29}\ e \cdot \text{cm} \tag{2.32}$$

at 90% confidence level. This limit smaller by a factor of 8.6 compared to the limit obtained in ACME I, which was the previous most stringent limit. At the time of writing, this result still stands as the most stringent published upper limit on the electron EDM. Using Eq. 1.6, the ACME II result probes physics at the 30 TeV (1-loop) or 1 TeV (2-loop) level, which is exceeds or is comparable to the scale probed at the LHC. A more detailed study by Cesarotti et al. [54] found that the ACME II result probed several SUSY scenarios at several TeV, including sleptons in 1-loop SUSY at 10 TeV.

Besides an electron EDM, $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ can also arise from a P- and T-violating scalar-pseudoscalar electron-nucleon interaction characterized by the dimensionless coupling parameter C_S [163]. Both effects could contribute to $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$, i.e.

$$\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}} = d_e \mathcal{E}_{\text{eff}} + W_S C_S, \tag{2.33}$$

where W_s is a molecule-specific constant whose value for ThO is $-2\pi\hbar \times 282$ kHz [103, 128].[22] The EDM limit in Eq. 2.32 assumes that $|C_S| = 0$. If instead we assume that $|d_e| = 0$, then this result would give us an upper bound[23]

$$|C_S| < 7.3 \times 10^{-10}. \tag{2.34}$$

[22] Note that the values of W_S reported in the cited papers must be multiplied by the ratio of $A/Z = 232/90$ for ThO. See the discussion in [113, §A.2].

[23] In Ref. [164], Flambaum et al. calculate various contributions to C_S, including from nucleon EDMs. This allows the extraction of implied limits on the proton and neutron EDMs from the limit on C_S set by ACME II. For the proton EDM, the limit is $\sim 10^{-23}\ e \cdot$cm with an uncertainty of 50%, which is remarkably only about a factor of ~50 from the most competitive limit [165].

Because W_S is molecule-specific, it is beneficial to perform eEDM experiments with different molecules in order to constrain the parameter space of possible contributions of C_s and d_e. The recent preprint from the JILA EDM experiment [97] which improves the eEDM upper limit by a factor of 2.4 is an important complement to the ACME II result, as it not only used a different molecular species but was also obtained with a very different technique (using trapped ions with long coherence times).

3

Overview of ACME III upgrades

While the ACME II measurement already probes new physics at the multi-TeV level, the next generation is pursuing statistical upgrades that will enable it to reach even higher scales, namely extending the spin precession time and increasing the number of detected molecules in the experiment. It also aims to suppress known systematic errors and sources of excess noise by a commensurate amount. This chapter is an overview of the ACME III apparatus currently under construction at Northwestern University. Subsequent chapters discuss in more detail features and upgrades which I occupied a leading role in developing: the H-state lifetime measurement which laid the foundation for the extended spin precession time (Chapter 4), detection efficiency improvements (Chapter 5), data acquisition and

experimental control (Chapter 6), and magnetic fields (Chapter 7). With all these upgrades, the next generation ACME III experiment is projected to measure the electron EDM with over an order of magnitude improvement in statistical sensitivity compared to ACME II and corresponding reductions in systematic errors and excess noise.

3.1 Statistical improvements

As discussed in Chapter 1, the statistical uncertainty of the electron EDM in the experiment (assuming that the shot-noise limit has been reached) is [113]

$$\delta d_{\mathrm{e}} = \frac{1}{2\tau \mathcal{E}_{\mathrm{eff}} \sqrt{\dot{n} T}} \tag{3.1}$$

where τ is the spin precession time, $\mathcal{E}_{\mathrm{eff}}$ is the effective electric field, $\dot{n}$ is the rate of detected molecules used for spin precession, and T is the experiment time. $\mathcal{E}_{\mathrm{eff}}$ is fixed by the choice of the molecule, which for ThO is 78 GV/cm. In ACME II, T could be extended in ACME III by taking data for a longer time, but realistically no more than a factor of ~ 2. The duty cycle of $\sim 20\%$ could be slightly improved by the addition of the ablation target load lock (see §3.1.4), which would reduce the need for weekly target changes. Dramatic improvements in duty cycle using other methods are more difficult. For example, previous experimental tests have shown that the 50 Hz rate of ablation is already at or close to the optimum amount production rate of molecules [148, p. 32].

Two avenues for significant improvement lie in τ and $\dot{n}$. Previously, τ was thought to be already optimized and limited by the lifetime of the H-state. However, a new measurement of the H-state lifetime [166] that it is longer than previously thought, opening the way towards an increase in τ. Substantial opportunities exist to increase $\dot{n}$ by performing molecular beam collimation to decrease geometrical losses [142] and increasing photon detection efficiency ($\sim 5\%$ in ACME II) [167, 168].

Thus, after a substantial period of development, three main statistical upgrades are

planned for ACME III:

1. An extended spin precession region that increases τ by ≈ 5 times which is essentially its optimum value given the newly measured radiative lifetime of the H-state.

2. An electrostatic molecular lens to improve the collimation of the molecular beam such that more molecules participate in the spin precession.

3. Upgrades to improve photon detection efficiency, consisting of enlarged collection optics and more efficient silicon photomultiplier (SiPM) detectors.

In addition, gains arise from the reducing excess data acquisition triggering noise (§6.1), improving rotational cooling, and implementing an ablation target load lock system (§3.1.4). A schematic of the newly upgraded apparatus is shown in Fig. 3.1.1.

3.1.1 Longer spin precession time

Some material in this section has been adapted from Ref. [166].

The spin precession time τ is primarily limited by τ_H, the radiative lifetime of the EDM-sensitive $H^3\Delta_1$ state. In ACME I and II, τ was set to ~ 1 ms based on earlier measurements of τ_H [69]. While adequate for practical purposes at the time, these measurements were all carried out in closed gas cell, resulting in complications from collisions with the buffer gas causing electronic quenching. No measurement had been done in a molecular beam, where collisions are not a major factor and is closest to the actual ACME experiment setup where spin precession occurs.

Measuring the $H^3\Delta_1$ radiative lifetime

Spurred by interesting results from some preliminary measurements of τ, we decided to build a dedicated experimental apparatus to perform a more conclusive measurement of τ_H in a molecular beam. The apparatus consisted of a long beamline with multiple locations to

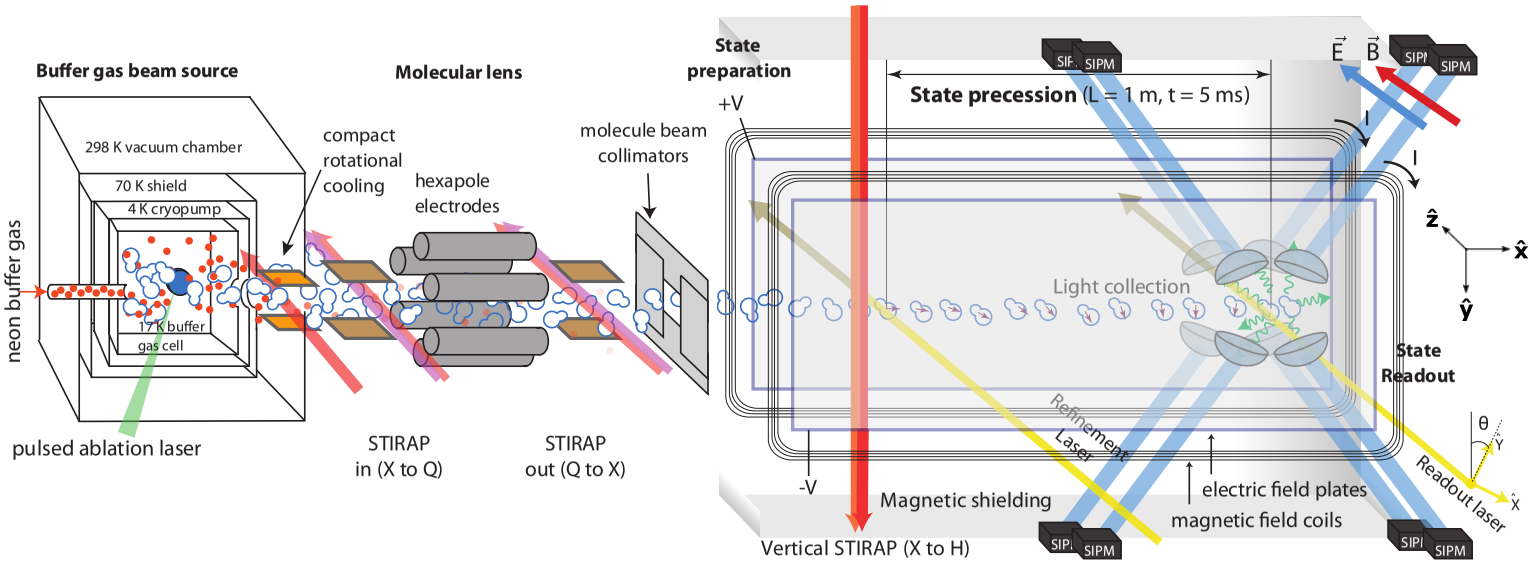

Figure 3.1.1: **Schematic of the new ACME III experiment apparatus.** Compared to the ACME II schematic (Fig. 2.1.1), there are three major statistical improvements: an extended spin precession region for $\tau = 5$ ms, an electrostatic molecular lens (including compactified rotational cooling), and newly improved photon detection (larger collection optics and SiPM detectors).

excite the molecules from to the H-state. After flying through the beamline, the remaining population was probed by laser-induced fluorescence at a detection region downstream. The signal from runs with different excitation regions was plotted and fit to extract the lifetime,

$$\tau_H = 4.2 \pm 0.5 \text{ ms}. \tag{3.2}$$

This measurement is described in Chapter 4 and also reported in a recent paper [166]. The value of τ_H is significantly longer than the spin precession time used in the ACME I and II experiments ($\tau \approx 1$ ms). Thus, a significant decrease in δd_e (Eq. 3.1) could result from increasing τ to something close to τ_H.

Determining the optimum spin precession time

The three curves in the plot of figure 3.1.2 show the EDM sensitivity gain (relative to ACME II where $\tau \approx 1$ ms) against τ for three different cases. The bands in the figure indicate uncertainty arising from the τ_H measurement. A second horizontal axis corresponding to the length of the interaction region L_{int} assuming a typical beam velocity of $v = 210$ m/s is shown.

In the limiting case where the molecular beam is perfectly collimated (red, upper curve in Fig. 3.1.2), the optimum τ for a given τ_H can be calculated by

$$\delta d_e \propto \frac{1}{\tau\sqrt{N}} \propto \frac{1}{\tau\sqrt{\exp\left(-\tau/\tau_H\right)}}. \tag{3.3}$$

The optimum value at $\tau = 2\tau_H$ is $\sim$20% smaller than when compared to $\tau = \tau_H$. Extending the coherence time from 1 ms in ACME II to about 5 ms (green vertical dashed line) will realize essentially all of this potential gain. This requires constructing an apparatus with a 5 times longer interaction region. Attaining this perfectly collimated limit with a molecular beam that spreads out, however, would also require a rather impractical scaling up of the radial size of the apparatus (i.e. perpendicular to the direction of the molecular beam).

The size of electric field plates and detection optics would need to increase, as must the power of the lasers needed to saturate the molecules over a larger volume.

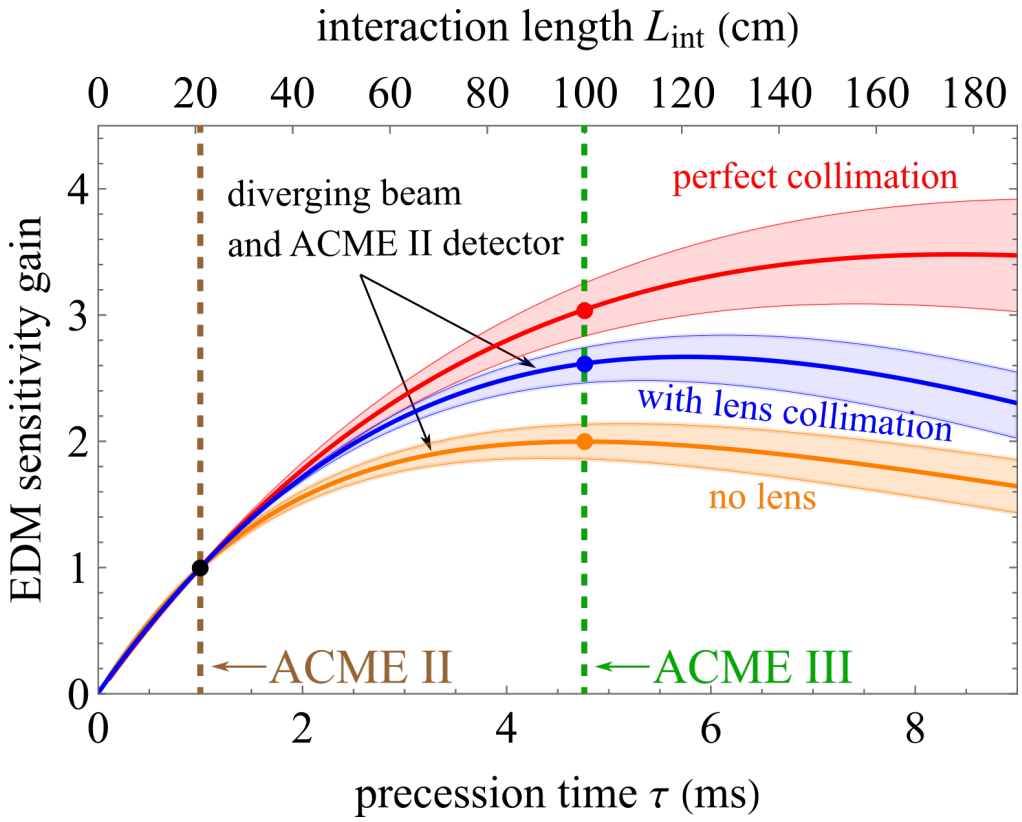

Figure 3.1.2: Projected EDM sensitivity gains over ACME II from extending the spin precession time given the measured lifetime τ_H for a perfectly collimated molecular beam (red, upper curve). The bands represent the effect of the uncertainty in τ_H, and the dashed lines indicate the coherence times for ACME II and the projected ACME III. For a diverging molecular beam in an apparatus that is made longer without increasing its radial dimensions, the longer τ_H increases the sensitivity by a factor of 2 (orange, lower curve). The sensitivity improves by up to 2.6 due to the effective collimation provided by the addition of an electrostatic lens for the molecules (blue, middle curve). (The additional sensitivity gain of 3.5 because the lens also captures more molecules is not included.) Figure adapted from Ref. [166].

Thus, it is much more practical to keep the radial size of the apparatus much the same as for ACME II. Extending the length of the spin precession region by itself would increase the eEDM sensitivity by about a factor of 2 (orange curve in Fig. 3.1.2). This factor can be estimated by using Eq. 3.3 while including the geometric factor $(L_0/(L_0 + v\tau))^2$ when

computing the number of molecules N, where $L_0 = 1.3$ m is the distance from the molecule beam source to the start of the precession region in ACME II.

Furthermore, numerical simulations [141, 142] suggest that the gain would increase from 2 to 2.6 if an electrostatic molecular lens is used to improve the collimation of the beam while keeping the radial size of the apparatus the same (blue middle curve in Fig. 3.1.2). This factor is achieved after optimizing the lens to maximize the number of molecules participating in the observed precession, because the lens cannot achieve perfect collimation. (Note that this only considers the collimating effect of the lens on the molecule trajectories in the spin precession region. An additional sensitivity gain of 3.5 because the lens captures more of the diverging molecules from the ablation source is not included in the figure. The molecular lens will be discussed in §3.1.2.)

Larger interaction region

Given this information, the ACME III apparatus was designed with a 100 cm long spin precession region, equivalent to $\tau \approx 4.7$ ms (green vertical dashed line in Fig. 3.1.2). This choice had momentous implications for the rest of the apparatus, as it required the complete redesign and rebuilding of all components of the interaction region to accommodate the longer precession, including the main vacuum chamber, magnetic shielding, magnetic coils, electric field plates, and collection optics.

Fig. 3.1.3 shows a comparison of the ACME II and III interaction vacuum chambers. The design and construction of the interaction region chamber, setup, and support table was led by Cole Meisenhelder and Daniel Lascar. The new vacuum chamber has dimensions 1568 mm $\times$ 591 mm $\times$ 533 mm. The dimensions were limited by the length of the spin precession region (along $\hat{x}$), the height of the field plates ($\hat{y}$), and the space required for the field plates and collection optics ($\hat{z}$). Both sides of the chamber ($\pm\hat{z}$) have two main side windows for the state preparation and probing.[1] The side windows are designed with the

[1] As seen in the figure, the holes for these windows were designed with a large diameter to accommodate

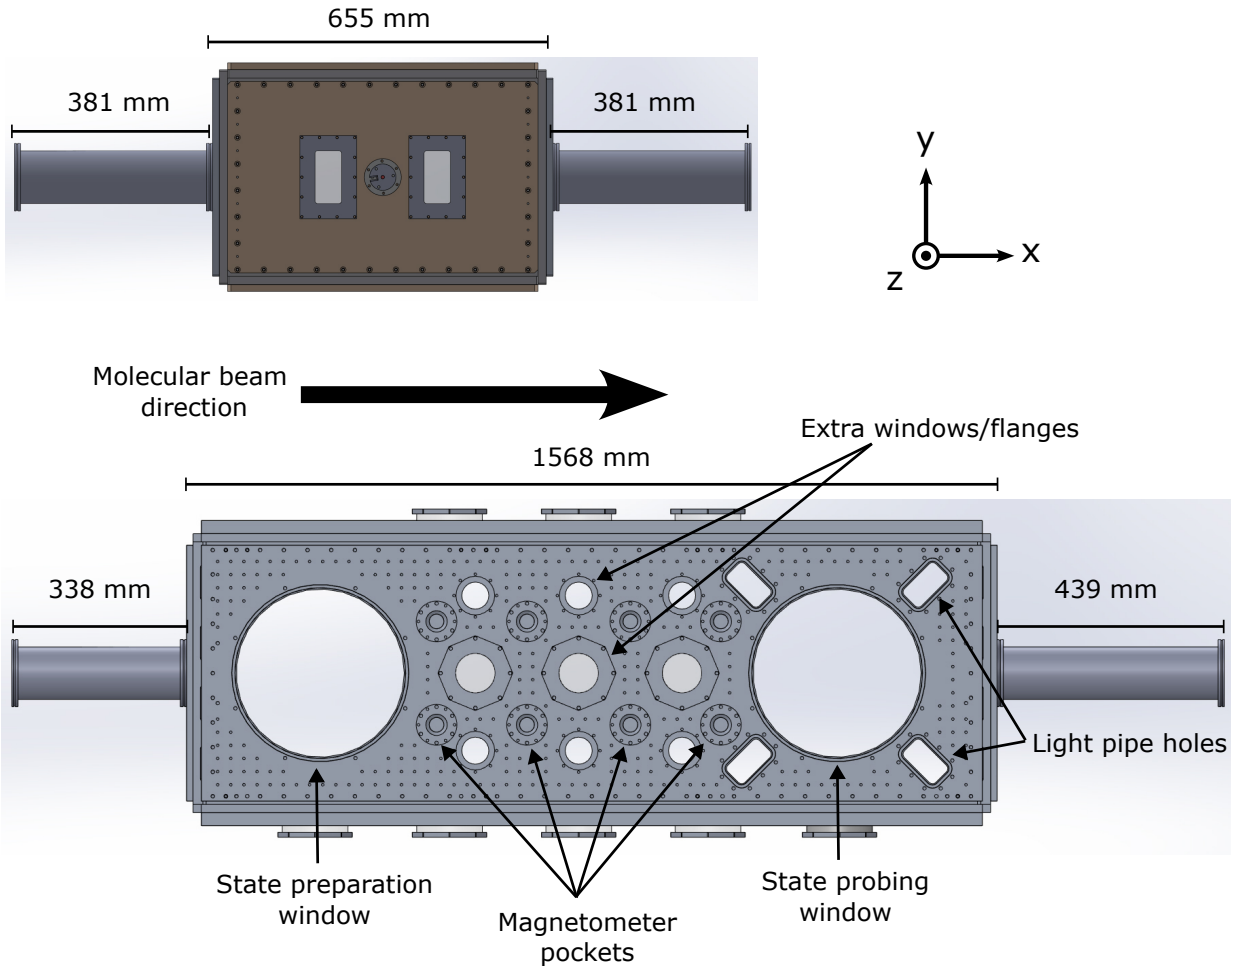

Figure 3.1.3: **ACME II and III vacuum scaled chamber comparison.** The new ACME III chamber is more than twice the previous length to accommodate the 5 times longer spin precession region. The center of the precession region is slightly offset along x in order to accommodate the downstream light pipe holes. Also shown are the upstream and downstream ISO-100 vacuum nipples which set the boundaries of the interaction region enclosed by the magnetic shields. These are connected to the electrostatic lens and beam dump chambers respectively.

goal of reducing stress birefringence which is known to cause systematic errors (see §3.2.1 for more discussion).

A major change is that the light pipes carrying photons from the collection optics in the detection region are no longer bent to emerge from the front and back ($\pm\hat{x}$) sides of the apparatus as in ACME II. Instead, the light pipes are straight and come out at a 45° angle. In order to conveniently position the light pipe holes, the center of the spin precession region is slightly offset along x (relative to the center of the chamber). More details on the mounting of the light pipes are described in §5.3. Many other openings were added to the chamber to give flexible optical access to the spin precession region. This included holes for inserting magnetometer pockets, as the apparatus was designed for improved monitoring of the magnetic fields (§7.2).

The main vacuum chamber is connected to the electrostatic lens chamber (upstream) and the beam dump chamber (downstream) using ISO-100 vacuum nipples (Fig. 3.1.3). The former nipple was designed to be shorter in order to minimize the distance between the lens and the detection region, which is beneficial for lensing. Finally, the entire interaction region is enclosed by three layers of rectangular mu-metal magnetic shields mounted on aluminum frames (not shown in the figure, discussed in §7.1).

3.1.2 Electrostatic molecular lens

In ACME II, due to the solid angle of the molecular beam, only about 0.007%[2] of ThO molecules produced from laser ablation were sufficiently collimated to successfully traverse the 1.2 m distance from the source to the interaction region without hitting the sides of the beamline. This provides much room for improvement. In addition, improved collimation would allow an increase of the spin precession time without having to significantly enlarge

a planned mechanism to enable vacuum *in situ* rotation of their birefringence axes. However, after some research, a different strategy was pursued to control and suppress the birefringence of the vacuum windows (§3.2.1).

[2]Based on the 1 in 20000 number for ACME I [113, §3.2.3] and the 1.4x enhancement in molecular signal due to a closer distance to the interaction region in ACME II [146, §5.4.1].

the radial size of the apparatus (§3.1.1). Since ThO does not have a set of suitable transitions for transverse laser cooling, an electrostatic lens was identified as the most promising path to accomplish this in ACME III.[3] The development of this lens was led by Xing Wu and the University of Chicago team. Technical details regarding the lens and the ThO Q-state used for lensing are described in two recent papers [141, 142]. Here we shall give a brief overview.

General principles of an electrostatic molecular lens

In an electrostatic molecular lens, a set of cylindrical electrodes are arranged around the axis parallel to the propagation direction of the molecular beam such that they create a harmonic electric potential [169]

$$U(r) \propto -r^{n/2}, \tag{3.4}$$

which implies an electric field with magnitude

$$|\mathcal{E}(r)| \propto r^{n/2-1}, \tag{3.5}$$

where r is the radial distance from the lens center axis and n depends on the number of electrodes (e.g. n=4 for a quadrupole lens and n=6 for a hexapole lens). As molecules traverse in the region between the electrodes, they experience a harmonic potential due to the Stark shift:

$$W(r) = -D\mathcal{E}(r) \propto -r^{n/2-1} \tag{3.6}$$

where D is the state-dependent dipole moment. This results in a radially directed restoring force

$$F(r) = dW(r)/dr \propto -r^{n/2-2}. \tag{3.7}$$

[3] A magnetostatic lens made using permanent magnets was also considered [70, 141], but an electrostatic lens was eventually chosen because of practical advantages, such as that it allows some amount of varying of the confining potential (by varying the voltage) and it does not suffer from the potential problem of slow vacuum pumping speeds in the narrow bore of the magnet.

In this way, the electrodes confine the molecules to the center of the lens, analogous to a thick optical lens focusing beams of light [170]. The effective focal length of the molecular lens depends on the molecular beam velocity, the electrode dimensions, and the dipole moment of the molecular lensing state. The choice of a quadrupole or hexapole lens depends on whether the Stark shift of the molecule is quadratic or linear. The Stark shift also determines the amount of voltage required to produce a sufficiently strong electric field.

Design of the ACME III lens

In ACME III, the $Q\ ^3\Delta_2$ state of ThO was chosen as the lensing state. Experimental investigations by Wu et al. [141] identified favorable properties for lensing, such as a strong linear Stark shift (particularly in the $|JM\Omega = 2, 2, -2\rangle$ state), a long lifetime (>62 ms), and a practical scheme for transfer in and out of the state from the ground X state using STIRAP through the C-state due to the large $Q-C$ transition dipole moment. Importantly, the strong Stark shift means that only $E \approx 30$ kV/cm is required to create a lens with a capture velocity of about 10 m/s.[4] (At even higher E, the Stark shift starts to become non-linear.) However, the goal of capturing the largest number of molecules (which would favor a high electrode voltage) must be balanced with the preference of operating at lower voltages to prevent technical challenges such as the production of X-rays [171].[5] The lens was also designed to maximize the number of molecules which are detected after participating in the spin precession, which meant setting the focus at a point between the state preparation and detection regions in the interaction region [142].

With these design goals in mind, a hexapole lens was designed with an internal bore radius $R = d_0 = 19$ mm (where R is the radius of the bore of the lens and d_0 is the electrode diameter, Fig. 3.1.4a), electrode length of 53 cm, and a typical electric field $E \lesssim$

[4]Previous attempts to construct a ThO molecular lens assumed the use of the X-state, which due to its quadratic Stark shift required a much higher voltage, leading to unintended X-ray production.

[5]To monitor and prevent X-ray production, multiple Geiger–Müller counters are installed around the lens chamber and linked to an interlock of the electrode voltage supplies.

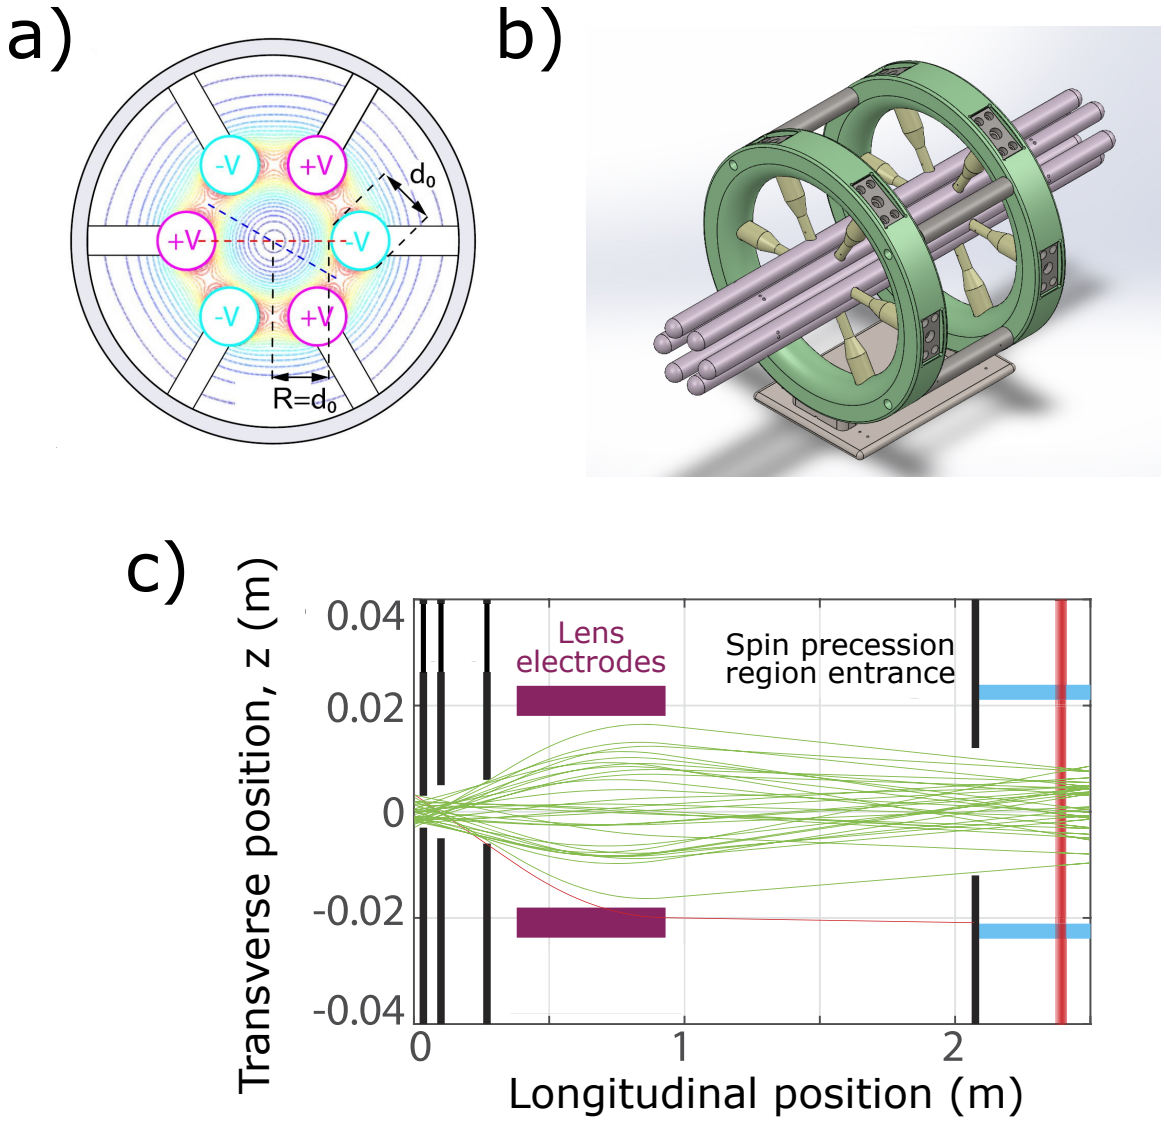

Figure 3.1.4: Design of the ACME III electrostatic molecular lens. a) Cross section of hexapole lens. The lens bore radius R is set to be the same as the lens electrode diameter d_0=19 mm. b) Design of the hexapole lens assembly. c) Simulated molecular trajectories from numerical simulation assuming $V = \pm 14$ kV, the setting found to be the optimum during experimental tests of the lens. Figure adapted from Ref. [142].

20 kV/cm. These parameters were chosen based on numerical simulations incorporating realistic velocity and spatial distribution of the molecules, beamline geometry, and the Stark shift of the $|JM\Omega = 2, 2, -2\rangle$ lensing state. During this design process it was also found that increasing the field plate separation to 6 cm (previously 4.5 cm) was beneficial to increase the molecular flux and prevent ThO molecules from being deposited on the field plates.

Transfer in and out of the lensing state

The lens setup includes modules along for efficiently transferring the molecules in and out of the lensing state. Initially, ThO molecules exit the beam source mostly in the ground X state with a $\sim$4 K Boltzmann distribution. Transfer into the lensing state occurs in two steps. First, rotational cooling is performed by shining 690 nm lasers tuned to the $X - C$, $P(2)$ and $Q(1)$ lines[6] in order to transfer the population in $J = 1, 2$ populations into the $J = 0$ state (Fig. 3.1.5a). This is virtually the same rotational cooling scheme as that deployed in ACME II,[7] but its implementation has been significantly enhanced in two ways.

First, in order to minimize the distance between the beam source and the electrostatic lens entrance (improving lensing), the rotational cooling module has been compactified. The two previously separated stages to pump the $P(2)$ and $Q(1)$ transitions are now combined into a single pumping module with a single optical path for the two lasers. This requires a number of optical components (Fig. 3.1.5b). The two lasers are prepared in orthogonal linear polarizations, combined on an input polarizing beam splitter (PBS), then passed through a polarization electro-optic modulator (EOM) and an output PBS. This allows for rapid (8 μs period) switching of the two lasers sent into the rotational cooling

[6] P(2) denotes the transition between $|X, J = 2\rangle \rightarrow |C, J = 1\rangle$ while Q(1) denotes $|X, J = 1\rangle \rightarrow |C, J = 1\rangle$.

[7] The third rotational cooling transition used in ACME II ($J = 3 \rightarrow J = 2$) is not implemented as it only added about 10% signal improvement at the cost of additional complexity.

chamber. As pumping on the $Q(1)$ line requires the mixing of the parity eigenstates of the C state (see §2.1.2), a set of electric field plates in the rotational cooling chamber applies an electric field of 150 V/cm along $\hat{y}$. This electric field only turns on when the $Q(1)$ laser is shining (Fig. 3.1.5c). The lasers are also switched rapidly (2 μs period) between two orthogonal polarizations using a Pockels cell. Thus, there are a total of 4 unique states of the system (two cooling transitions and two orthogonal polarizations).

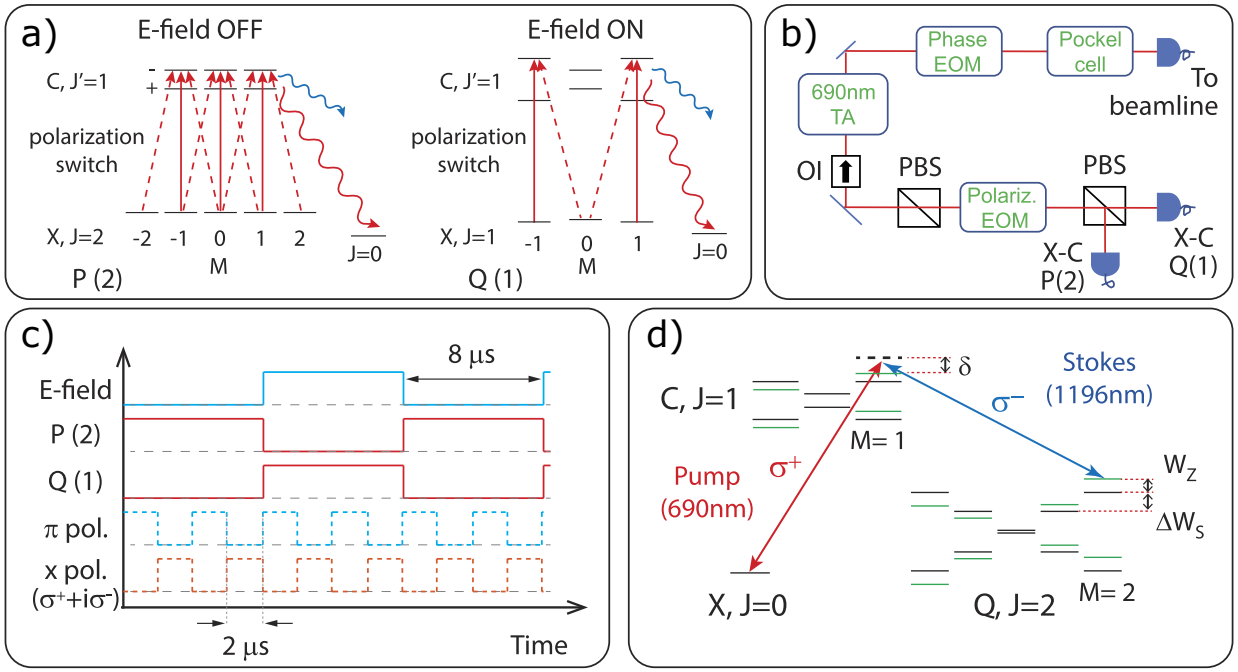

Figure 3.1.5: Schematics for rotational cooling and STIRAP for the electrostatic lens. a) Level diagrams depicting the two transitions (P(2) and Q(1)) used to increase the number of molecules in the $|X, J = 0\rangle$ state. b) Optical components for implementing enhanced rotational cooling. c) Modulation timing scheme for switching polarizations, electric fields, and rotational cooling line. d) Level diagram depicting the STIRAP process to efficiently transfer the molecular population from the X to Q and back for lensing. b) Figure adapted from Ref. [142].

Second, the saturation of the optical pumping has been improved. A tapered amplifier is used to increase the power of the 690 nm lasers from ∼15 mW (ACME II) to ∼100 mW. Sidebands are added to the lasers (the FWHM Doppler width of the molecules is 20 MHz).

Prisms are used to make 13 passes of the lasers through the molecular cloud.

As a result, the length of the rotational cooling region has been compressed from 20 cm to 4 cm, while experimental tests show that the rotational cooling enhances the population in $|X, J = 0\rangle$ by a factor of ∼3.5 at highest saturation, which is ∼1.4 times better than the rotational cooling enhancement in ACME II.

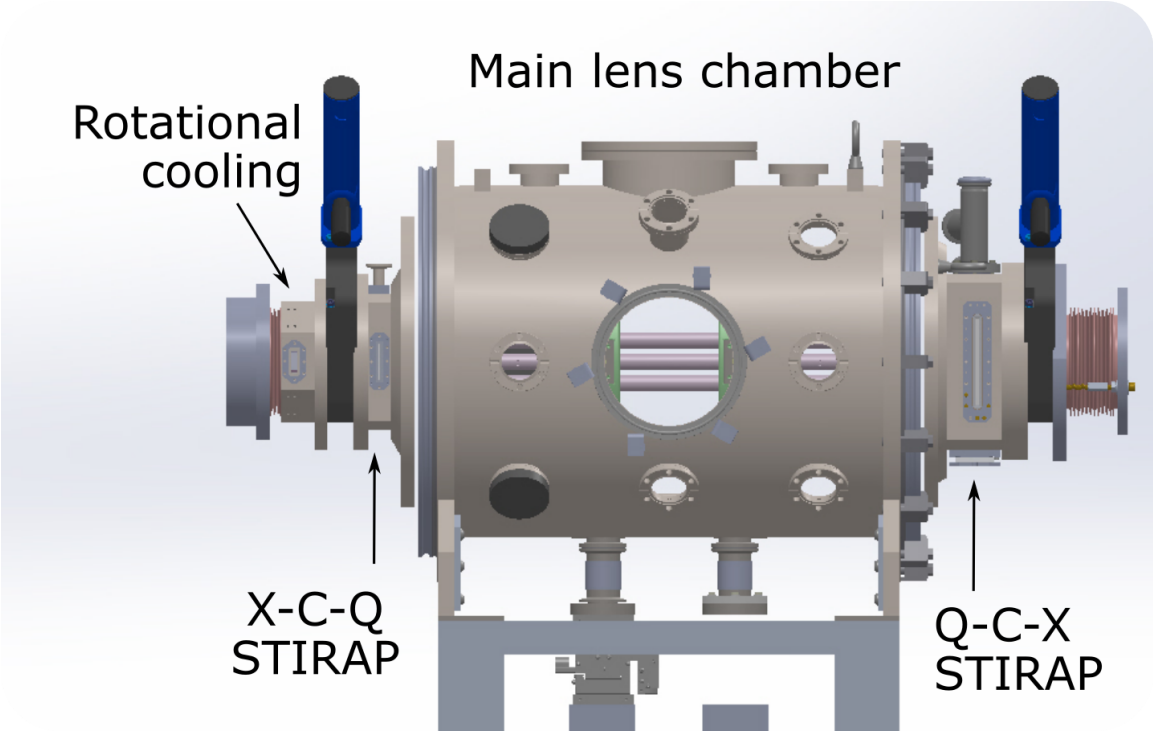

Figure 3.1.6: Design of the entire electrostatic lens system including the compactified rotational cooling chamber, the two STIRAP chambers, and the main lens chamber. The chambers also contain coils and field plates (not visible) to apply electric and magnetic fields.

In the second step, the molecules are transferred from $|X, J = 0\rangle$ to $|Q, JM\Omega = 2, 2, -2\rangle$ by using STIRAP via the $|C, JM\Omega = 1, 1, -1\rangle$ state (Fig. 3.1.5d). This is performed in a separate STIRAP region with magnetic coils and electric field plates which apply 5 G magnetic and 50 V/cm electric fields, in order to break the degeneracy of the $M = \pm 2$ and $\Omega = \pm 2$ states. Here, the two STIRAP laser beams (690 nm Pump and 1196 nm Stokes) are

sent in along the z-axis with σ polarization. After the molecules have undergone lensing, a second STIRAP using the same set of transitions is performed in reverse to transfer them back to the $|X, J = 0\rangle$ state. In order to prevent unwanted spin flips as the molecules travel between the STIRAP regions and the entrance/exit of the lens, the electric fields are designed to not flip sign in the transition regions. Previous tests have shown that ∼80% combined transfer efficiency for $X - C - Q$ and $Q - C - X$ STIRAP can be achieved with laser powers of ∼100 mW [141]. A schematic of the entire lens chamber (including rotational cooling, STIRAP, and the main lens chambers) is shown in Fig. 3.1.6.

Construction and testing of the lens

Development and testing of the electric lens was performed at Harvard using the ACME II beamline and apparatus. The lens (including the rotational cooling region) was installed in between the beam source and the ACME II interaction region. The second stage of the STIRAP back to $|X, J = 0\rangle$ was skipped, such that the molecules remained in the Q state as they enter the interaction region. At the end of interaction region a 746 nm probe laser excites the molecules to the I state, producing 512 nm photons from the $I - X$ decay (similar to ACME II). The distance between the exit of the lens to the probe laser is 148 cm, which is comparable to the planned ACME III distance (∼186 cm). During these tests, the electrode voltage was scanned from 0 to ±25 kV, with the optimal electrode voltage found to be $V = \pm 14$ kV. At this setting, the gain in detected molecules is 16.2(6) relative to the lens turned off, with a Doppler width of 2.3 MHz (1σ) at 746 nm, equivalent to 2.4 MHz at 703 nm (the frequency of the $H - I$ probe laser in ACME III). This is comparable to the 2.1 MHz Doppler width in ACME II [148]. Spatially, the molecular beam follows a Gaussian distribution with an FWHM of ∼1.3 cm. In contrast, the ACME II molecular beam was a flat top with 2.3 cm FWHM. Thus, while the Doppler width is slightly larger, the narrower Gaussian profile may actually improve the saturation with the molecular lens.

These experimental results are found to be consistent with the molecular trajectory sim-

ulations. After modifying these simulations slightly to reflect the ACME III beamline geometry, the lens is predicted to give a factor of 12 enhancement in molecular flux, equivalent to a gain in EDM sensitivity of 3.5 compared to ACME II. (This is on top of the collimation from the lens which allows an increase in precession time without increasing the radial size of the apparatus, as previously discussed in §3.1.1.) There is also the additional factor of 1.2 from the improved rotational cooling.

3.1.3 Improved photon detection efficiency

In ACME II, the collection optics covered ∼20% of the collection angle, and ∼25% collected were detected by the photomultiplier tubes (PMTs) due to limited quantum efficiency, resulting a total of only ∼5% of generated photons detected overall. Efforts were made to improve the photon detection efficiency. The first is a redesigned collection optics system utilizing larger diameter, lenses and light pipes which covered a larger solid angle. The second is the replacement of the PMTs with silicon photomultiplier (SiPM) detectors, which have about twice the quantum efficiency. Both of these improvements have been designed, constructed, and tested in an experimental mock-up of the ACME III detection system. These tests show a factor of 1.7 statistical gain in photon collection from the improved collection optics and a factor of 2.3 gain from the SIPMs, resulting in a combined gain of 2.0 in EDM statistical sensitivity. Details of both of these improvements are described in Chapter 5 as well as several recently published papers on the SIPM development [167, 168, 172].

Earlier efforts were also made to improve the detection efficiency by performing optical cycling on the $X - I$ transition, where the $I - X$ branching ratio has been measured to be 91% [152]. In principle, this would allow each molecule to produce more than one photon. However, after some experimental tests, the transition turned out to produce less photons than expected, and other practical and theoretical considerations led us to switch our experimental efforts to improving the detection efficiency through other means

as described in the previous paragraph. Appendix A provides a fuller description of the optical cycling project.

3.1.4 Other statistical improvements

In addition to the three major statistical improvements discussed above, there are two remaining improvements that are anticipated to further increase the statistical sensitivity of the ACME III apparatus.

Timing noise reduction

We have briefly mentioned the timing noise in the DAQ system used for the ACME II which resulted in the data having $\chi^2 \approx 3$ (§2.2.9). In 2019, this problem was solved by making several straightforward technical changes to the DAQ system: 1) changing the sampling rate to an integer multiple of the base clock rate of the DAQ oscilloscope, 2) locking the oscilloscope's clock to a stable external reference clock, and 3) updating the software drivers for the oscilloscope. These changes suppressed the excess noise to well below the ACME II shot noise limit, equivalent to a factor of 1.7 gain in EDM statistical sensitivity relative to ACME II. Details on this noise reduction are described in §6.1, as well as in a published paper [161].

Ablation target load lock system

The second additional statistical improvement comes from implementing an ablation target load lock system to enable quicker changing of targets during a run. During the ACME II final run, a freshly installed ThO_2 ablation target would have about 1-2 days of constant maximum signal followed by a long linear decay over the next few days (Fig. 3.1.7). The target would usually be replaced after about two weeks of constant use. Target replacement was a time-consuming process which involved warming up the beam box to room temperature (∼24 hours) and 2-3 personnel engaging in several hours of intense work while

wearing extensive personal protective equipment to protect from radioactive dust, after which it would take another ∼12 hours to fully cool down to operational temperatures. Thus, more frequent changing of targets was simply not practical.

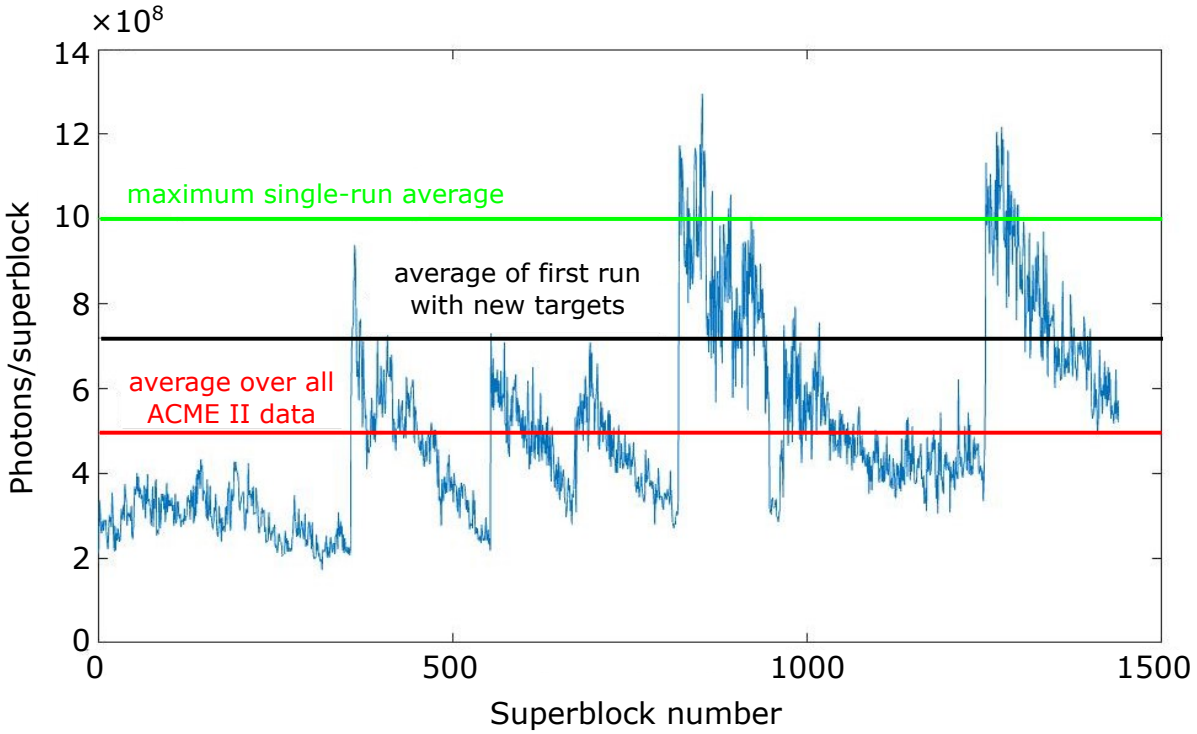

Figure 3.1.7: Plot of number of photons against superblock number in the ACME II final dataset. Spikes indicate ablation target changes. Data analyzed and plotted by Zack Lasner.

For ACME III, a mechanical system has been developed to allow frequent *in situ* changing of targets without having to warm up or open the beam box (Fig. 3.1.8). The development was led by Zhen Han (University of Chicago). Two mechanical arms (horizontal and vertical) are mounted on the beam box using vacuum-compatible feedthroughs. Using the arms, an old target can be detached from its mount in the ablation cell and extracted from the beam box via an airlock chamber at the top of the beam box. (The new target would be inserted and mounted in the same way.)

This system has been constructed and successfully tested. The entire target replacement

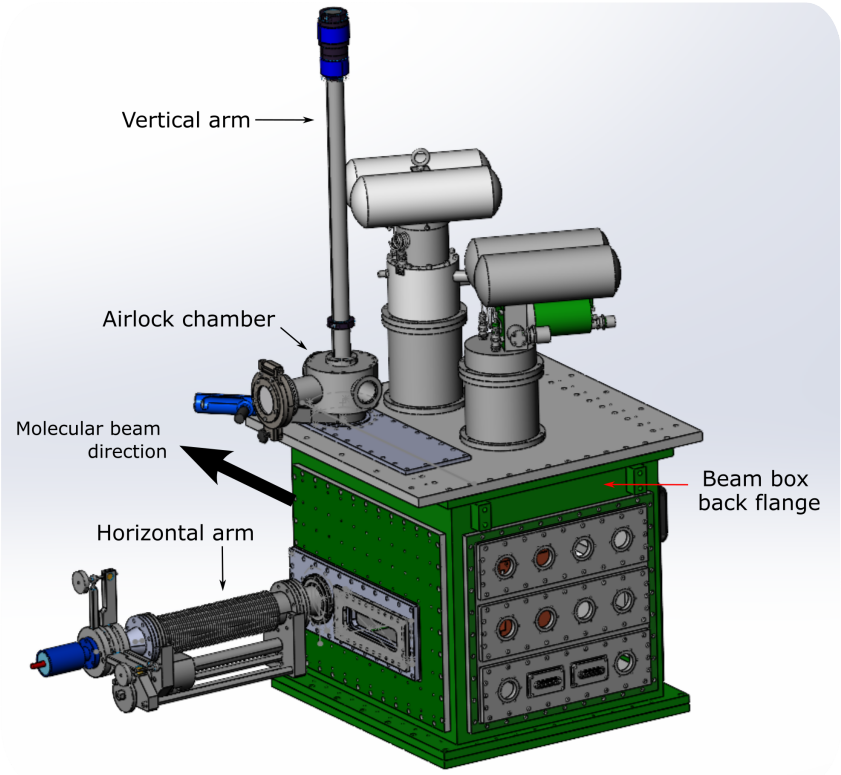

Figure 3.1.8: Design of ablation target load lock system (from model developed by Zhen Han) which enables *in situ* target replacement. Vertical and horizontal arms are mounted with vacuum feedthroughs on the beam box. These are used to screw or unscrew the ablation target from its mount. Targets can be inserted or extracted using the airlock chamber along the path of the vertical arm.

process takes ∼1 hour. During this time, only a small increase in the beam box temperature which stabilizes back to operating temperature within ∼15 minutes (due to the relatively warm temperature of the new target) has been observed. With this system, daily target changes will be possible, which will increase the absolute number of molecules used in the spin precession by a factor of ∼1.5, equivalent to a factor of 1.2 gain in EDM sensitivity.

The load lock system will also improve the general operating efficiency of the experiment. Previously, when operating with a target older than a few days, a human operator constantly made small adjustments to the ablation laser optics in order to find an ablation spot with good signal. This is no longer required with a constantly fresh target.

3.1.5 Summary

Putting together all the aforementioned statistical improvements, the ACME III experiment is projected to measure the electron EDM with a sensitivity larger by more than an order of magnitude compared to ACME II (Table 3.1.1). Assuming a similar amount of integration time, this would imply a statistical uncertainty of $\sim 10^{-31}$ e · cm. All of the improvements have been individually designed, constructed, and their statistical gains experimentally demonstrated. What remains is to assemble all the modules into a single integrated apparatus.

Improvement	**EDM statistical sensitivity gain**
Increased precession time	2.6
Electrostatic lens	3.5
Improved photon detection	2.0
Timing noise reduction	1.7
Ablation target load lock	1.2
Improved rotational cooling	1.2
Total	**45**

Table 3.1.1: Projected improvements in ACME III and their associated predicted gain in EDM sensitivity relative to ACME II.

3.2 Systematic error and excess noise suppression

In order to measure the electron EDM with an order of magnitude improved sensitivity, all systematic errors and excess noise must also be suppressed below the projected statistical uncertainty. In ACME II, the final statistical uncertainty was 373 μrad/s. Let us conservatively assume that the target uncertainty in ACME III is $\sim$40 μrad/s. Table 3.2.1 shows all class A systematic shifts in the ACME II uncertainty budget. These shifts come from known experimental imperfections that must be dealt with for the next generation. Out of the six shifts, four have uncertainties greater than the ACME III target. In addition to systematic shifts, there are known sources of excess noise that may prevent the experiment from reaching the shot noise limit.

Parameter	**Shift**	**Uncertainty**	**Section**
$\mathcal{E}^{\mathrm{nr}}$	-56	140	3.2.1
$\omega_{\mathrm{ST}}^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ (via $\theta_{\mathrm{ST}}^{H-C}$)	0	1	-
$P_{\mathrm{ref}}^{\mathcal{N}\mathcal{E}}$	-	109	3.2.1
$\partial\mathcal{B}_z/\partial z$ and $\partial\mathcal{B}_z/\partial y$	7	59	3.2.2
$\lvert\mathcal{C}\rvert^{\mathcal{N}\mathcal{E}}$ and $\lvert\mathcal{C}\rvert^{\mathcal{N}\mathcal{E}\mathcal{B}}$	77	125	3.2.3
$\omega^{\mathcal{E}}$(via $\mathcal{B}_z^{\mathcal{E}}$)	1	1	-

Table 3.2.1: **Class A systematic shifts and their statistical uncertainties in ACME II.** Units are in μrad/s. The total systematic uncertainty of ACME II was 310 μrad/s (including several class C systematic uncertainties). Also shown are the sections where we discuss strategies to reduce the larger systematic uncertainties to below the ACME III statistical sensitivity goal.

3.2.1 Polarization gradients and AC Stark shifts

As discussed in §2.2.3, several systematic errors have been known to arise from an AC Stark shift phase coupling to ellipticity gradients in the the refinement and readout laser polarizations to cause a non-zero slope $d\omega/d\Delta$. In the presence of $\mathcal{E}^{\mathrm{nr}}$, this couples to $\Delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ produced by the $\mathcal{E}^{\mathrm{nr}} \rightarrow \Delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ mechanism (§2.2.2) to result in a systematic shift in $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ (denoted by the "$\mathcal{E}^{\mathrm{nr}}$" entry in Table 3.2.1). Second, the $P^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ systematic uncertainty was

calculated based on the possibility of an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-odd component of the refinement laser power $P_{\mathrm{ref}}^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ arising from AC Stark effects. Better control of polarization gradients will reduce both of these systematic effects. For the vertical STIRAP lasers, a similar effect arises from ellipticity gradients (caused by the small waist of the beams) and Stark interference. This effect is washed out by the refinement laser, causing the small $\theta_{\mathrm{ST}}^{H-C}$ systematic shift which is well below the projected ACME III statistical sensitivity.

Typically, ellipticity gradients are caused by birefringence gradients in optical elements that the lasers traverse before reaching the molecules, such as the electric field plates and vacuum windows. A birefringence gradient can be induced by a thermal load (from absorption of a high powered laser beam) or mechanical stress. In ACME II the total ellipticity gradient was $\lesssim 0.1\%$/mm on the East side and $\sim$1-2%/mm on the West side, with the difference likely caused by differing mechanical stress birefringence in the vacuum windows.

Ideally, the total polarization gradient should be reduced by an order of magnitude to keep up with the improved statistical sensitivity of the experiment. However, the EDM phase ($\phi^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$) will be 5 times larger compared to ACME II (due to the increased precession time), which reduces the relative significance of a systematic phase shift caused by a local polarization gradient at the state preparation/readout regions by the same factor. This introduces a "slack factor" of 5 to the requirement.

Improved field plate mounting design to reduce mechanical stress birefringence

In ACME II, Corning 7980 fused silica was chosen as material for the field plates as it has more favorable properties for reducing thermal stress birefringence compared to the Schott Borofloat glass used in ACME I. The thickness of the ITO layer was also decreased from 200 nm to 20 nm to minimize absorption of laser power. As a result, the thermal stress induced birefringence per watt of absorbed laser power was reduced by a factor of

$\sim$7 compared to the ACME I field plates [83, Table A.1][158, §4.2.2]. The highest laser power sent through the field plates in ACME I was $\sim$4 W (1090 nm) [113], whereas ACME II used $\sim$1 W (703 nm).[8] Thus, the overall thermal birefringence was reduced by over an order of magnitude in ACME II. This is likely also adequate for ACME III, as there are also other methods to further decrease it, such as modulating (chopping) the laser power in between molecular shots.

However, as mentioned above, mechanical stress birefringence remained an issue even in ACME II, making it infeasible to send in lasers from the West side of the experiment. To improve this in ACME III, a different type of glass (Schott SF57HTUltra) was identified as favorable as it has stress birefrigence coefficient $K = 0.07$ GPa^{-1},[9] which is $\sim$200 times lower than Corning 7980 [160] while having similar thermal birefringence. The main disadvantage of the material is that it is not available in sizes larger than 16 cm $\times$ 16 cm. This is insufficiently large for the new field plates, which must be much larger along $\hat{x}$ and $\hat{y}$ to accommodate the longer spin precession region (100 cm) and further field plate separation (6 cm) respectively. Thus, in the early stages of development, several different composite field plate designs were explored, consisting of SF57HTUltra for the state preparation and readout sections and ordinary glass would be used for the other sections. These designs were found to have significant engineering challenges to ensure electric field homogeneity and smooth transitions between sections.

Further simulations and experimental tests by Peiran Hu (University of Chicago) led to the decision to use a single large piece of Corning 7980 as for the field plate (the same material as ACME II). In order to reduce the mechanical stress birefringence, a new field plate mounting system was designed where contact between the field plate and the mounting frame is mediated by small spheres similar to ball bearings. As a result, the clamping force is mostly normal to the surface of the plate instead of tangential, significantly reducing

[8] The absorption of the glass at the two wavelengths are about the same.

[9] Measured by Zhen Han using polarimetry at 703 nm.

the mechanical stress birefringence [173]. Experiments with a prototype mounting system designed based on this principle have shown that mechanical stress can be reduced to $\sim$0.01%/mm, which sufficiently fulfills ACME III requirements. The new field plates with improved mounting are currently under construction.

Improved vacuum windows

A second source of polarization gradients is mechanical stress birefringence from the side vacuum windows at the preparation and readout regions. The windows are designed to be circular instead of square, due to evidence in the literature that circular symmetry reduces stress birefringence [174, 175]. The relatively smaller size of the window (compared to the field plates) made it possible to use the SF57HTUltra glass which has a much lower stress birefringence coefficient. However, the lower fracture toughness (typical of glass with lead oxide content [176]) presents the danger of cracking from the differential vacuum pressure. Currently, experimental tests are being carried out at Okayama University to construct a mounting design which includes metal bar supports to relieve the vacuum pressure. With this design, it should be possible to reduce the mechanical stress birefringence from the vacuum windows by two orders of magnitude, fulfilling the ACME III requirement.

Suppressing polarization gradient-related systematic errors using the $\tilde{\mathcal{G}}$ switch

Another method to reduce systematic errors arising from polarization gradients is to rotate the polarization of the preparation and readout lasers to match the birefringence axis of the optical element. This was possible to do in ACME I as state preparation was done with laser beams only propagating horizontally along $\hat{z}$ (also known as the $\tilde{\mathcal{G}}$ or global polarization switch) [113], but not in ACME II where the $X-C-H$ STIRAP was performed with vertical laser beams, fixing the starting polarization of the molecules. To return this capability in ACME III, the possibility of switching to performing STIRAP with horizontal laser beams

was explored. This would require sending in the lasers through the field plates. It is not practical for the existing $X-C-H$ scheme as it requires a 10 W 1090 nm for the Stokes laser, which would likely cause an unacceptable amount of thermal stress-induced birefringence. Schemes using relatively strong transitions such as $X-L-H$ and $X-U-H$ [152] were considered, but these involved sending in an ultraviolet laser which may damage the ITO layer on the field plates. The most realistic possibility for performing horizontal STIRAP is using the $X-A-H$ scheme. To explore this possibility, the $X-A$ (943 nm) and $H-A$ (1892 nm) transition dipole moments were carefully measured by performing optical pumping on the $X-A$ transition (Appendix B), resulting in the following measured values:

$$d_{A-X} = (0.4 \pm 0.1)\ \text{D}, \tag{3.8}$$
$$d_{A-H} = (0.8 \pm 0.2)\ \text{D},$$

which implies a STIRAP power requirement of about 290 mW at 943 nm and 140 mW at 1892 nm (§B.4). Both of these are attainable with commercial lasers and sufficiently low for avoiding thermal stress-induced birefringence.

However, a general concern for STIRAP through the field plates is the possibility of reflection off the ITO coating by the 1892 nm laser, resulting in near-resonance optical pumping of the H-state. Besides degrading the amount of H-state population, this was feared to cause spatial or velocity-dependent systematic errors. This concern led to the theoretical exploration of other $\tilde{\mathcal{G}}$-switch-capable preparation schemes involving horizontal optical pumping (instead of STIRAP).

Overall, because of the cost for new lasers involved in any of these schemes (whether STIRAP or optical pumping), it is likely that implementation of the $\tilde{\mathcal{G}}$ switch will be seriously pursued only if the other measures taken to suppress the polarization gradient-related systematic errors described above are found to be insufficient.

3.2.2 Magnetic fields

In ACME II, magnetic fields were correlated with several systematic error and excess noise. Here we shall describe strategies for reducing these effects and their implementation in ACME III.

Magnetic field gradient systematics

A systematic shift in the $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ arises from two different mechanisms related to magnetic gradients and non-reversing electric fields (§2.2.5). In both cases, a magnetic field gradient such as $\partial\mathcal{B}_z/\partial z$ creates a spin precession phase gradient $\partial\omega/\partial z$:

- In the first mechanism, $\partial\omega/\partial z$ couples to a position-velocity gradient $\partial z/\partial v_z$ which results in a spatial gradient of detuning, $\partial\Delta/\partial z$. The $\mathcal{E}^{\mathrm{nr}} \to \Delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ mechanism (§2.2.2) then produces $\Delta^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ which couples with $\partial\Delta/\partial z$ to produce a systematic shift in $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$. This mechanism occurs in both the vertically propagated STIRAP (via $\partial B_z/\partial y$) and the horizontally propagated readout (via $\partial B_z/\partial z$) lasers.

- In the second mechanism, $\partial\omega/\partial z$ couples to an $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-correlated state preparation efficiency gradient produced by the STIRAP state preparation lasers in the presence of a nonzero $\mathcal{E}^{\mathrm{nr}}$ gradient. The readout laser contributes similar effects, although the STIRAP contribution is larger. An analogous mechanism exists for $\partial B_z/\partial y$ gradients with both lasers.

For the $\partial\mathcal{B}_z/\partial z$-related shifts, we can express the contributions of these two mechanisms in the two respective terms of the following equation:

$$\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}_{\partial\mathcal{B}_z/\partial z} = c_1 \frac{\partial\mathcal{B}_z}{\partial z}\frac{\partial z}{\partial v_z}\mathcal{E}^{\mathrm{nr}} + c_2 \frac{\partial\mathcal{B}_z}{\partial z}\frac{1}{\eta_0}\frac{\partial\eta}{\partial\delta}\frac{\partial\mathcal{E}^{\mathrm{nr}}}{\partial z}, \tag{3.9}$$

where δ is the STIRAP 2-photon detuning, η is the STIRAP efficiency (with η_0 being its value at $\delta = 0$), and c_1, c_2 are constants determined by largely fixed experiment parameters

such as the molecular beam size, transition dipole moments, and laser frequencies. The first term comes from the STIRAP laser and the second term from the readout laser. (For concision, we have omitted the readout laser contribution via the first mechanism which is dependent on its detuning Δ.) An analogous equation can be written for $\partial B_z/\partial y$.

In ACME II, cancellation of this shift was performed by tuning δ until the first term cancelled the second in Eq. 3.9. The same procedure can performed in ACME III. We can also examine the possibility of tuning other terms in the equation:

- Significant reduction of $\mathcal{E}^{\mathrm{nr}}$ and $\partial\mathcal{E}^{\mathrm{nr}}/\partial z$ is not practical since the exact mechanisms that lead to $\mathcal{E}^{\mathrm{nr}}$ in the field plates are not well-understood.[10]

- The use of a molecular lens (§3.1.2) will somewhat reduce position-velocity correlations present in the second term, as seen in the molecule trajectories in figure 3.1.4.

- Since both terms couple to magnetic field gradients $\partial\mathcal{B}_z/\partial z$ and $\partial\mathcal{B}_z/\partial y$, reducing or at least keeping the same level of gradients would be the most straightforward strategy to keep this systematic under control. The ACME II gradients were on the level of several μG/cm after active cancellation by gradient coils.

Molecular beam velocity noise

It is known that excess noise may enter the EDM channel from the applied $\mathcal{B}_z$ magnetic field coupling to the shot-to-shot molecular beam velocity fluctuations, resulting in noise in the measurement of precession time τ (§2.2.8). In ACME II the $\sim$0.1% fractional change in velocity between traces was enough to cause a noticeable effect. There are three possible methods to suppress this noise:

1. Reduce the shot-to-shot longitudinal velocity fluctuations in the molecular beam,

2. improve the precesion of the τ measurement, or

[10] As an example, in ACME I and II, despite the different ITO thicknesses of 200 nm and 20 nm, the magnitude of $\mathcal{E}^{\mathrm{nr}}$ was about the same, at the level of several mV/cm.

3. decrease the magnitude of the magnetic field, which decreases the size of the Zeeman phase $\phi^{\tilde{\mathcal{B}}}$ relative to the EDM phase $\phi^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$.

The most realistic way to accomplish method 1 would involve performing some sort of velocity selection on the molecules in a shot (such as by timing the preparation and/or detection lasers to only turn on at certain times relative to the ablation laser pulse, similar to a slotted disk velocity selector [177, Chapter 3]). However, simulations show that such a method would result in too much signal loss.

Method 2 was explored in ACME II by turning off the state preparation lasers for a short moment after each ablation laser pulse to create a "notch" (gap) in the cloud of molecules [148, §4.12]. The molecular beam velocity (and thus τ) could be measured in every shot by locating the center of the notch in the trace showing the arrival times of the molecules in the detection region. However, this method also causes some signal loss around the notch area and is limited by velocity dispersion. Still, given the significantly larger signal expected in ACME III, implementing this technique with multiple notches might be beneficial, if not entirely solving the problem.

The most realistic method to reduce systematic error is method 3. In ACME II, this was accomplished by taking most of the data at $|\mathcal{B}_z|$ = 1-2 mG. Accordingly, the ACME III apparatus must reduce $|\mathcal{B}_z|$ by an order of magnitude to 100-200 μG in order to suppress this noise below the target uncertainty.[11] To prevent large local field fluctuations causing systematic effects, this would also require reducing the background magnetic field from 100-300 μG in ACME II to $\sim$10 μG.

Magnetic field design in ACME III

Considering the aforementioned requirements regarding the magnetic fields, the target background field in ACME III was set to $\sim$1 μG, a conservative target exceeding the

[11] The 5 times increase in τ does not affect the requirement to decrease $|\mathcal{B}_z|$ by a factor of 10. Increasing τ affects $\phi^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ and $\phi^{\tilde{\mathcal{B}}}$ equally, but the excess noise primarily couples in only through $\phi^{\tilde{\mathcal{B}}}$. This is why in the absence of excess noise, changing $|\mathcal{B}_z|$ by itself does not affect the uncertainty in the EDM channel $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$.

anticipated minimum requirement of $\sim$10 μG. This required a complete redesign and rebuilding of the magnetic shielding to accommodate this goal and the longer interaction region, which was led by the Northwestern University team. Integrated with the magnetic shields is an expanded network of magnetometers to carefully measure and monitor the field. A new set of magnetic field coils was developed taking into account the much longer interaction region while aiming to be capable of applying magnetic fields (especially $\mathcal{B}_z$) with high homogeneity. As of the time of writing, the magnetic shields have been assembled and are currently being characterized for optimum degaussing, multiple methods of magnetometry have already been set up and tested, and the magnetic coils are being manufactured. Lastly, magnetometry with the molecular beam via the magnetically sensitive Q-state has also been experimentally tested, which will enable high-sensitivity measurement of the magnetic field in exactly the same region as the spin precession. All these magnetic field-related efforts are described in more detail in Chapter 7.

3.2.3 Contrast correlations

The final major systematic uncertainty comes from higher-order, $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$ correlated terms in the calculation of the contrast $\mathcal{C}$. The two significant terms arise from a non-zero $|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}\mathcal{B}}$ and $|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ coupling to $\omega^{\tilde{\mathcal{B}}}$ and ω^{nr} respectively (Eq. 2.28). In ACME II, the measured $\partial\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}/\partial|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}\mathcal{B}}$ was $\sim$6$\times$ larger than $\partial\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}/\partial|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$. There was no clear evidence of a physical mechanism to cause non-zero $|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}\mathcal{B}}$ and $|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$, and their average measured values across the whole dataset were consistent with zero. One possibility is that their true values are actually zero. In that case, in ACME III we would similarly obtain an average value of $|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}\mathcal{B}}$ and $|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$ which is consistent with zero with better statistical uncertainty. However, if they are not zero, then we would have to investigate and attempt to reduce this noise further.

A complementary strategy to guard against this possibility is to reduce $\partial\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}/\partial|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}\mathcal{B}}$ and $\partial\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}/\partial|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$. The planned reduction of $|\mathcal{B}_z|$ by an order of magnitude will simi-

larly reduce $\omega^{\tilde{\mathcal{B}}}$ and thus $\partial\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}/\partial|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}\mathcal{B}}$, which is the larger slope of the two. To reduce $\partial\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}/\partial|\mathcal{C}|^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$, we would have to reduce ω^{nr}, the non-reversing angle between the state preparation and readout bases. It may be possible to carefully monitor and tune this angle using a half-waveplate, which was not done in ACME II. Alternately, it would also be beneficial to investigate if it is possible to modify the standard data analysis routine (§2.1.10) to be immune to contrast correlations entirely.

3.2.4 Summary

The next generation ACME III apparatus has been designed to suppress the two leading sources of excess noise and systematic uncertainty: magnetic field gradients and polarization gradients. Some minor experimental tests of these features are still ongoing (such as with the polarization gradients of the vacuum windows), but clear and realistic plans exist to suppress all known sources of systematic uncertainty below the target statistical uncertainty. No systematic uncertainty is expected to be a roadblock towards achieving the projected order-of-magnitude improvement in EDM sensitivity.

3.3 General upgrades

In addition to the specific upgrades aimed to suppress statistical and systematic errors, several other upgrades were implemented in the ACME III apparatus to improve its general reliability and flexibility. In contrast to ACME I and II which were located at Harvard, the ACME III experiment will be assembled and operate at Northwestern University. The design and planning of the new apparatus had to take into account specific features of the infrastructure and conditions of the new lab. The new setting also brought the opportunity to rebuild and reoptimize all aspects of the apparatus based on previous experiences in ACME I and II.

3.3.1 Data acquisition and experimental control

In ACME II, traces recording detected photons from the molecules were acquired at 50 Hz (following the frequency of the ablation laser). However, only averages of groups of consecutive 25 traces were saved. Each averaged trace represented ∼500 ms of data, corresponding to the period of the fastest switch, $\tilde{\mathcal{N}}$. The DAQ system had software and hardware limitations which prevented long-term saving of all the raw traces. Thus, diagnosis of certain noise sources (such as the timing noise) was difficult. In ACME III, the data acquisition system has been improved to enable recording all 25 traces. This necessitated an upgrading of the hardware and software of the data acquisition system to deal with the much greater data throughput and storage. In addition, an online data flagging and analysis system has been developed to aid in sorting through the mass of data. The existing ACME experimental control software has also been integrated with the new system. Setup and testing of the new DAQ system have been completed at Northwestern University. These upgrades are described in Chapter 6.

3.3.2 Lasers

The ACME III experiment will require substantially more lasers than previous generations. Three STIRAP operations will be performed along the beamline ($X-C-Q$, $Q-C-X$, $X-C-H$), each requiring a pair of narrow linewidth lasers locked to high-finesse ultra low expansion (ULE) cavities. In addition, two 690 nm lasers are required for rotational cooling, one 703 nm laser for state refinement and readout, and a 746 nm ($Q-I$) laser for the Q-state magnetometry. Many of these lasers also include separate optical amplifier modules. To maintain a high experimental duty cycle, higher laser reliability is needed so that less time is spent adjusting and relocking lasers. In ACME III, all homemade external cavity diode lasers (ECDLs) have now been replaced with more reliable commercial ECDLs,[12]

[12]Mostly Toptica DL Pro or TA Pro systems. The 703 nm is an MSquared SolsTiS 4000.

virtually all of which are locked to ULE cavities using Toptica DigiLock modules. This ensures that these lasers have a linewidth no more than a few tens of Hz and decreases the need to perform frequent Doppler scans to locate the center of the transition. The YAG laser for ablation has also been replaced with a Litron Nano L 90-100 which has a much smaller footprint compared to its ACME II predecessor and does not require an open loop cooling system. The 10 W Nufern fiber amplifier for the 1090 nm STIRAP will be replaced by a newer 20 W Precilaser fiber amplifier, which may slightly improve saturation and increase the ∼75% STIRAP transfer efficiency in ACME II [140, Fig. 3c].

To lock the lasers, there are three ULE cavities housed in temperature-controlled vacuum chambers. One is a new system from Stable Laser Systems (dedicated to locking the electrostatic lens STIRAP lasers), whereas the other two are carried over from ACME II but have been rebuilt and rejuvenated after they were moved from Harvard to Northwestern. All the Toptica lasers are now housed in two Toptica T-racks. These enclosed laser rack systems improve thermal management, vibration isolation, and cable management while saving optical table space. Finally, a completely new comprehensive laser safety system (including many interlocks and shutters) has been built to accommodate the operating requirements at Northwestern. Besides myself, substantial contributions to efforts to move and upgrade the laser systems at NU have been made by Cole Meisenhelder, Xing Wu, and Collin Diver.

3.3.3 Beam source

In tandem with the development of the ablation target load lock system (§3.1.4), a new cryogenic buffer gas beam source has been constructed for ACME III. The basic design and dimensions remain the same, with a few improvements, such as 1) using oxygen-free high thermal conductivity (OFHC) copper instead of aluminum for the thermal radiation shields, improving their temperature distribution, 2) making the shields 50% thinner to reduce their thermal mass, and 3) gold-plating the outermost shield, reducing radiation

heat load from the outside. This project was conducted by Zhen Han and Zack Lasner. The new beam box was constructed and tested at Harvard with a ytterbium atomic beam before being recently moved to the final experiment site at Northwestern.

3.4 Conclusion

In conclusion, the ACME III campaign has so far demonstrated a realistic path towards an order of magnitude improved measurement of the electron compared to ACME II. Several statistical improvements have been experimentally demonstrated, with the principal improvements coming from increased spin precession time, electrostatic lens, and improved photon detection. The apparatus has been designed to suppress all known sources of excess noise and systematic uncertainties well below the projected statistical uncertainty, as well as with improvements to enhance reliability and experimental flexibility (such as the upgraded data acquisition system). Ongoing projects at the time of writing include the design and construction of the low-birefringence vacuum windows, characterization of the magnetic shield degaussing, and the construction of the field plates and magnetic coils. The next major phase of the experiment is to assemble all the components together and perform a spin precession measurement to demonstrate the projected statistical uncertainty including all the improvements (Table 3.1.1). At this point, the experiment would be ready for systematic error investigations before taking the final dataset.

4

Measurement of the $H\ ^3\Delta_1$ Radiative Lifetime in ThO

This chapter describes the $H\ ^3\Delta_1$ lifetime measurement which found its value to be several times longer than the ACME II precession time, $\tau = 1$ ms [166]. The measurement showed that we could obtain a significant gain in EDM statistical sensitivity by increasing τ in ACME III, as discussed in §3.1.1.

Some material in this section has been adapted from a prior publication [166].

4.1 Preliminary Measurements

Previous measurements of the H-state lifetime were performed in a cryogenic gas cell of ThO [69, 135, 178], where the buffer gas in the cell could quench the molecules to the H-state faster than the natural radiative lifetime, requiring the fitting of two exponential decays. Thus, the measurement in Ref. [69] was reported as a lower limit, $\tau_H \geq 1.8$ ms. This was adequate for ACME I and II, as the apparatus was designed for $\tau = 1$ ms. No measurements had been done in a molecular beam where collisions are not an issue and better approximates the actual ACME beamline.

In February 2019, two preliminary molecular beam measurements of τ_H were performed using existing beamlines. Cristian Panda, Xing Wu, and Cole Meisenhelder contributed to these measurements. In these experiments, molecules are excited to the H-state at one of two different regions along a beamline and the amount of remaining population was detected in a detection region downstream. From comparing the signal obtained when exciting with the closer vs. further excitation regions, one could extract the lifetime.

The first measurement was performed in the ACME II apparatus (Beam Box I) with two pairs of excitation regions spaced 3 cm and 99 cm apart, while the second was performed in the Beam Box II apparatus with a single pair spaced 33 cm apart. These measurements found $4.3 < \tau_H < 7.0$ ms. The large range of values reflected the fact that these setups were not originally designed for these measurements. However, they strongly suggested that τ_H is several times longer than 1 ms. If this were true, it would provide a path to significantly improving the statistical sensitivity of the next generation ACME apparatus by increasing the precession time by a similar factor. It would also have massive practical implications, as designing and building a five times longer spin precession region is much more difficult.

4.2 Main Measurement Method

4.2.1 Overview

Spurred by these interesting results, a dedicated apparatus for the H-state lifetime measurement was constructed (Fig. 4.2.1). I was the main researcher who led, carried out, and analyzed data in the experiments. In addition, Cole Meisenhelder also made significant contributions in constructing and maintaining the experimental setup. The apparatus consists of a ThO buffer gas beam source connected to a beamline with parallel electric field plates running across its entire length, allowing full polarization of the molecules throughout their entire journey. A laser is shone on the molecules at one of five regions along the beamline (termed excitation regions) to excite it to the H-state. After being transferred to the H-state, the molecules fly through the beamline until they reach a detection region at the end, where another laser is used to probe the remaining amount of molecules using laser-induced fluorescence (LIF). The measurement is repeated many times, each time switching the location along the beamline where the excitation laser is shone. The fluorescence signals obtained by exciting the molecules at different points along the beamline are then plotted and fit to an exponential function to determine the H-state radiative lifetime.

4.2.2 Buffer gas beam source

The ACME "Beam Box II" beam source was used as the cryogenic buffer gas beam source.[1] A Litron Nano L 90-100 laser is used to ablate the ThO_2 target, producing pulses of ThO molecules as in the main ACME experiment [138]. During the course of the measurement, a target change was performed, resulting in two datasets taken before and after the change.[2] The mean longitudinal velocity of the beam was 210 m/s and 250 m/s before and after the

[1] This beam source was converted from the former thermochemical beam source developed by Elizabeth West [147] into an ablation beam source used for ACME III development.

[2] The two datasets were also taken before and after the COVID pandemic-related shutdown in summer 2020, after which many components of the experimental setup had to be rebuilt.

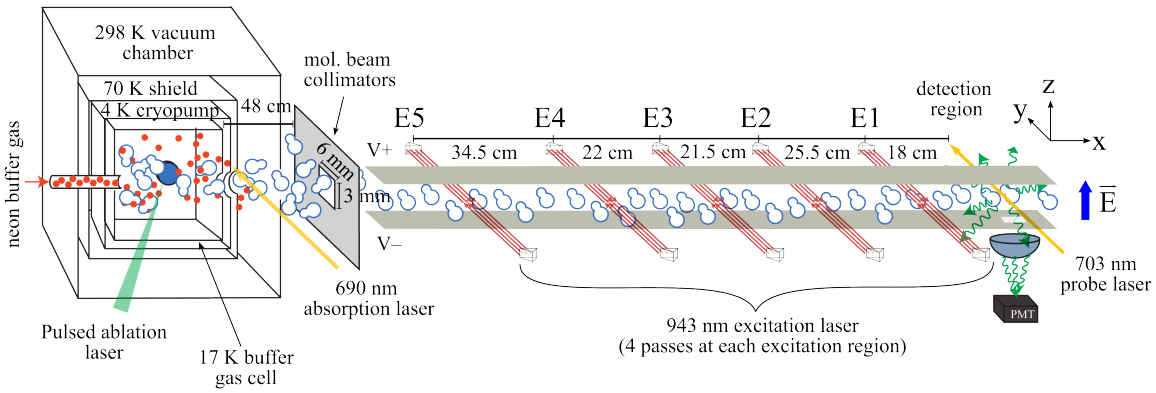

Figure 4.2.1: **Schematic of experimental setup used to probe the H state lifetime.** Molecules are excited to the H-state at one of the five excitation regions and the remaining population is probed at the detection region. Figure reproduced from Ref. [166].

target change. This change in velocity for a new target and an aging buffer gas cell was also observed during the ACME II experiment.

After being produced from laser ablation, the molecules are mostly in their ground electronic (X) and lowest vibrational states, with a rotational temperature of ∼4 K. A laser tuned to the 690 nm $X - C$ transition produces an optical absorption signal just after the buffer gas cell which is used as part of the determination of the molecular velocity. About 48 cm downstream from the cell aperture (which has a diameter of 5 mm), 6 mm horizontal and 3 mm vertical collimators control the size and distribution of transverse velocities in the molecular beam, giving a 1σ Doppler width of 6 ± 1 MHz along the y-axis for a laser wavelength of 943 nm. The molecular pulses are ∼0.5 ms in duration as they leave the source but expand to ≈2 ms by the time they arrive at a detection region that is 1.78 m away.

As is typical of the ACME beam source (§2.1.1), the molecular pulse intensity typically varies by 10-15% from pulse to pulse. This signal also slowly decays, dropping by as much as 50% after 5-15 minutes. We revive the signal by moving the ablation laser to a different location on the ThO_2 target [113]. To suppress the effect of these changes, normalization

was required (§4.2.5).

4.2.3 Beamline and excitation regions

The molecules then fly into the beamline, where a pair of ∼1.3 m long parallel plates produces a ∼38 V/cm electric field that is vertical along $\hat{z}$, perpendicular to the molecular beam direction $\hat{x}$ (Fig. 4.2.1).[3] Within this field, at one of five nominally-identical excitation regions (labeled E1-E5 in Fig. 4.2.1), an excitation laser transfers molecules into the metastable H state which is fully polarized within the applied electric field [69]. The five excitation regions are located at distances from the detection region L_i, measured from the center of the beam volume occupied by multiple passes of the excitation lasers (Table 4.2.1).

No.	**Distance to 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 4.2.1: Distances between the excitation regions (E#) and the detection region.

The 943 nm excitation laser is linearly polarized along $\hat{z}$ and excites the transition $|X, J = 0\rangle \rightarrow |A, J = 1\rangle$. About 30% of the molecules spontaneously decay to an incoherent mixture of M and $\mathcal{N}$ states within the $|H, J = 1\rangle$ manifold (Fig. 4.2.2a) [151]. (This state preparation method is identical to the method used in the ACME I experiment [96].)

The molecules freely propagate down the beamline while undergoing radiative decay from the metastable H state to the stable ground state X. The number of H state molecules reaching the detection region decreases exponentially as e^{-t/τ_H} with lifetime τ_H, where t is

[3] The mounting structure for the field plates were primarily designed and built by Cole Meisenhelder.

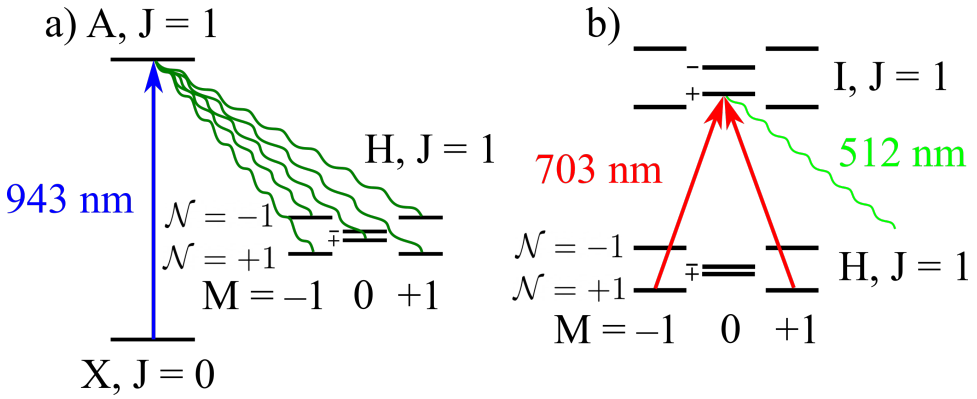

Figure 4.2.2: a) Optical pumping excitation of the H-state in an electric field. b) Detection of fluorescence after a second excitation. Figure reproduced from Ref. [166].

the time between excitation and detection. This remaining population is probed by optically pumping the $|H, J = 1\rangle \rightarrow |I, J = 1\rangle$ transition using a 703 nm probe laser, linearly polarized along $\hat{x}$ (Fig. 4.2.2b). The I-state is short-lived [152] and rapidly decays back to the ground state, producing 512 nm photons which are detected using a photomultiplier. The intensities are normalized to that observed for excitation at E1 and fit to an exponential decay curve to obtain τ_H.

4.2.4 Lasers

The excitation laser[4] is locked to an iodine-clock-stabilized laser using a slow scanning cavity transfer lock, resulting in a laser linewidth (FWHM) of ∼2 MHz. Software-controlled acousto-optic modulators (AOMs) coupled to optical fibers enable rapid switching of the optical path of the laser between different excitation regions. Each excitation region contains 65 mW (±10%) of laser power that is quadruple-passed through the molecular beam using a pair of prisms to improve saturation of the optical pumping, resulting in ∼230 mW of total circulating power (after accounting for transmission losses from the vacuum

[4]Toptica DL Pro with Roithner RLT0940-300GS diode followed by a Toptica BoosTA tapered amplifier. This laser was previously used as the state preparation laser in the ACME I experiment.

windows). For the second dataset, replacement of an optical isolator for the excitation laser and aging of the TA resulted in a loss of ∼20% laser power compared to the first dataset. The laser is linearly polarized along the z-axis and propagates along the y-axis (see Fig. 4.2.1). Each excitation region also contains an independent set of optics to expand the laser beam to a $1/e^2$ height and width of ∼1 cm and ∼0.1 cm respectively. This height was selected based on the measured height of the molecular beam at each excitation region of $\lesssim 0.8$ cm. Ensuring that we address all of the molecules reduces the complicating effects of molecular beam divergence.

The probe laser[5] is locked to a second, stabilized low-powered diode laser[6] using a delay line transfer lock scheme [148, 179]. The second laser is stabilized to a high-finesse optical cavity using a Pound-Drever-Hall locking scheme [180]. This results in the probe laser having a linewidth of ∼20 kHz. The first dataset was taken with a laser power of 160 mW, and the second dataset with 240 mW. At the detection region, the probe laser beam is expanded to the same size as the excitation laser and linearly polarized along the x-axis.

4.2.5 Data acquisition and analysis

Absorption and fluorescence data is acquired using an oscilloscope[7] sampling at 50,000 samples/s and 800 samples/record, resulting in records of length 16 ms. A digital pulse generator[8] is used for triggering data acquisition, with the trigger set such that the molecular pulse signal is centered within the trace.[9] In the first dataset, fluorescence from 64 consecutive molecular pulses was averaged into one trace. In the second dataset, 32 pulses were averaged to enable quicker change and optimization of the laser ablation spots on the ThO_2 target.[10] For each trace, we subtract a background by sampling the first 3 ms and

[5] M Squared SolsTiS TiSapph, the same probe laser as used in the ACME II experiment
[6] 703 nm Toptica DL Pro
[7] NI PXI-5922.
[8] SRS DG645.
[9] The trigger timing is adjusted between the first and second datasets to account for the different average molecular velocities.
[10] We verified that the number of records per trace did not affect the measured lifetime value.

last 4 ms and integrate the molecular pulse signal by sampling the region from 3 to 12 ms. This subtracts out background from electronic offsets and non-molecular photon scatter.

To suppress noise from signal fluctuations from the molecular beam source (as mentioned in §4.2.2), we use the large signal acquired from excitation at E1 for normalization. Other normalization schemes were tried (such as using the $X - C$ absorption signal), but they were less successful in suppressing noise. Thus, we switch between acquiring E1 data and one of E2-5 every 7-9 seconds, resulting in two groups of several traces, from which we calculate the relative intensity I_i/I_1. Figure 4.2.3 is a semilog plot of the observed relative intensities at the excitation regions $i = 2, ..., 5$. The ablation target used to produce the ThO molecules was changed between the first (orange points) and second (blue points) datasets.

The measured relative intensities are fit to

$$\frac{I_i}{I_1} = \exp\left(-\frac{L_i - L_1}{v\,\tau_H}\right) \tag{4.1}$$

to obtain the best fit straight line shown on the semilog plot. The measured L_i are in Table 4.2.1. The molecular beam velocity v is measured by subtracting the arrival times of the center of mass of the fluorescence trace at the detection region and the optical absorption trace from the 690 nm laser placed just after the ablation cell (as shown in Fig. 4.3.1). The statistical uncertainty in τ_H from the fit (0.02 ms) is computed using standard error estimation procedures from maximum likelihood estimation [181]. A computation using bootstrapped datasets [156] yields the same value. As is obvious from Fig. 4.2.3, and a reduced $\chi^2 = 4.4$, the statistical uncertainty is small compared to other sources of uncertainty that are discussed in the next section.

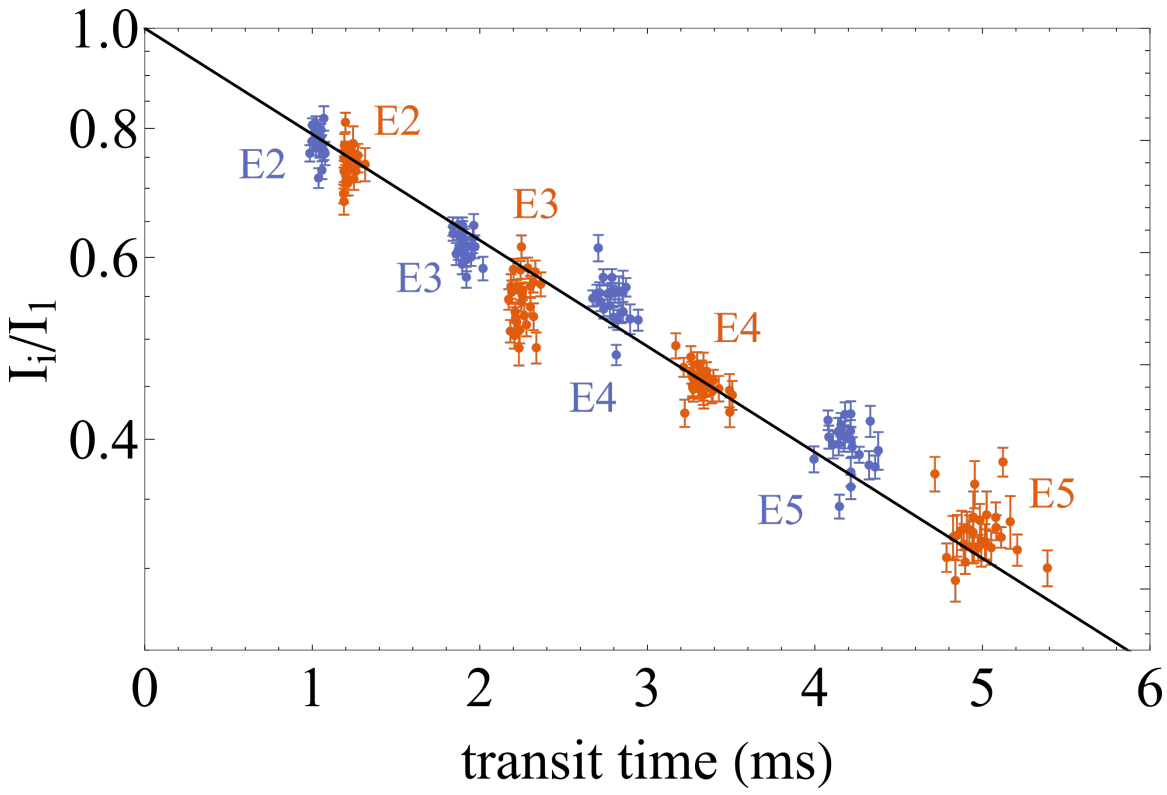

Figure 4.2.3: Semilog plot of the intensity ratios as a function of transit time between the excitation regions and fluorescence detector, with a fit to Eq. 4.1. The colors indicate two datasets taken at widely separated times, before (orange) and after (blue) an ablation target change that resulted in a change in ThO velocity. Uncertainty in the transit times arising from uncertainty in the positions of the excitation regions are addressed in §4.3. Note that adjacent data points have been consolidated in the plot for clarity. Figure reproduced from Ref. [166].

4.3 Uncertainty Analysis

4.3.1 Excitation Efficiencies

The largest of the systematic uncertainties comes from possible variations in excitation efficiencies at the excitation regions, given that Eq. 4.1 assumes an equal number of H-state molecules leave each excitation region. In practice, there could be variations between the excitation regions due to velocity dispersion, differences in laser illumination and saturation, or misalignments. Fluctuations from one excitation region to the next can be represented as a set of excitation efficiencies η_i that are slightly different for each excitation region. Such an effect could plausibly explain why datapoints from the same excitation region tend to cluster together above or below the fitted line in Fig. 4.2.3.

To estimate the spread of excitation efficiencies, the detected signal as a function of laser power P at excitation point i is fit to $S = S_{max}[1 - \exp(-P/P_s)]$ to extract the saturation power P_s and the extrapolated maximum signal S_{max}. The efficiency is $\eta_i = S/S_{max}$ for the laser power P used for the lifetime measurement. We obtain a spread in excitation efficiencies of up to 7% from the measured values and their uncertainties.

To determine the effect of this variation in excitation efficiency on the uncertainty of τ_H, the data is fit for each of a large set of efficiencies selected randomly from a Gaussian distribution with a 7% standard deviation. In effect, this simulates a range of scenarios including when the excitation efficiencies happen to be all skewed in one direction. The distribution of the resulting τ_H values has a standard deviation of 0.4 ms. This value is used as the estimated systematic uncertainty from differing excitation efficiencies at the excitation regions.

4.3.2 Background H-state population

A smaller contribution that only slightly increases the lifetime uncertainty comes from a possible background of H-state molecules in the signal region of the trace. This is distinct

from the background subtracted during data analysis (§4.2.5) which is sampled from the region before and after the signal within the trace. Such a background of molecules in the H-state could possibly come from population which is initially excited to short-lived, higher-lying electronic states by laser ablation and subsequently decay down to H. A similar effect was previously observed in the gas cell lifetime measurements [69]. Detuning the excitation laser far from resonance established that the background is below 3% of I_1, which corresponds to a 0.2 ms uncertainty. A more recent experimental study suggested that this background is actually smaller than 1%. However, since this study was conducted with a different beam source and beamline setup, we decided to use the conservative larger value of the background in the error budget.

4.3.3 Laser position

To get substantial (>90%) laser saturation, the laser beams cross the molecular beam four times at each excitation region. The resulting width of each interaction region makes it possible to localize the distance between the excitation regions and the detection region to about 1 cm. To estimate the resulting uncertainty in τ_H, the data is fit to fitting functions like Eq. 4.1 but with $L_i - L_1$ in each case offset by a value from a Gaussian distribution that is $\pm\sqrt{2}$ cm wide on average. From this we learn that the 1 cm uncertainty most likely contributes an uncertainty of ± 0.04 ms in uncertainty to τ_H, which can be neglected compared to the other uncertainties. The key here is that the 1 cm uncertainty is fractionally small compared to the longer $L_i - L_1$ values corresponding to excitation regions which are given more weight in the fit.

4.3.4 Velocity determination

The final source of uncertainty is from the determination of the molecular beam velocity v. This is deduced from the time that elapses between when an absorption signal is observed just after the buffer gas cell and the fluorescence signal is observed at a distance 178 ± 1 cm

away (Fig. 4.3.1). The uncertainty in τ_H from this uncertainty in the spacing is negligibly small for this large separation. A bigger contribution arises because dispersion in buffer gas cell exit times and velocities give rise to a spread of arrival times. The difference of the average arrival times (solid lines in Fig. 4.3.1) for each distribution is used to compute the velocity. The 0.1 ms difference of the peak (dashed lines in Fig. 4.3.1) and average arrival times of each distribution when applied to the dataset results in a 0.05 ms uncertainty in τ_H from the velocity determination.

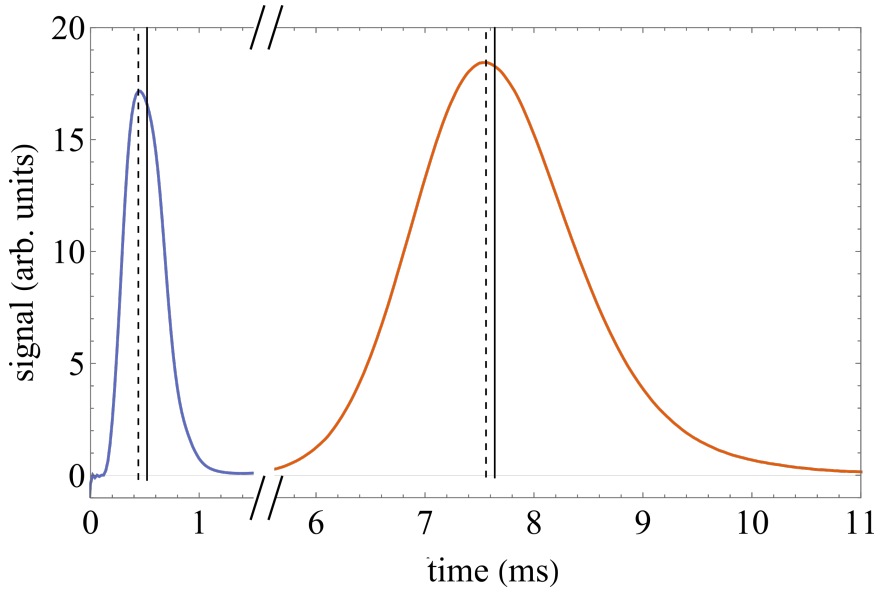

Figure 4.3.1: Average absorption near the molecule source (blue, left) and fluorescence in the detection region (orange, right) traces show the transit time and the longitudinal velocity dispersion of the molecular pulses. The peak time (dotted) and average time (solid) differ slightly. The displayed traces are averaged from the entire second dataset. The height of the absorption signal has been inverted and re-scaled so it can be compared to the detected fluorescence signal. Figure reproduced from Ref. [166].

4.3.5 Summary

The sources of uncertainties in the measurement are summarized in Table 4.3.1. Assuming these uncertainties are uncorrelated, they together produce the 0.5 ms uncertainty in

Eq. 4.2.

Source of uncertainty	**Uncertainty (ms)**
Fitting uncertainty	0.02
Excitation laser saturation	0.4
Background uncertainty	0.2
Excitation laser position	0.04
Velocity determination	0.05
Total uncertainty (1σ)	**0.5**

Table 4.3.1: Systematic and statistical uncertainties in the lifetime measurement.

4.4 Results and conclusion

The best fit result for the measured H state lifetime is

$$\tau_H = 4.2 \pm 0.5 \text{ ms}. \tag{4.2}$$

The uncertainty stated is a one standard deviation combined statistical and systematic uncertainty. As discussed in the last section, the systematic uncertainties dominate.

We performed three main tests to check the robustness of the result and its uncertainty. First, we analyzed the two datasets (orange and blue in Fig. 4.2.3) separately, given that they have average molecular velocities that differ by about 20%, as is clearly evident in Fig. 4.2.3. This gives $\tau_H = 4.1$ and 4.4 ms for the first and second datasets respectively. Second, we excluded data from one excitation region and refit the remaining data, performing this procedure for each of E2-5. Third, we refit the data while relaxing the requirement that $I_i/I_1 = 1$ at $(L_i - L_1)/v = 0$ (see Eq. 4.1), using the equation

$$\frac{I_i}{I_1} = A \exp\left(-\frac{L_i - L_1}{v\,\tau_H}\right), \tag{4.3}$$

where both A and τ_H are fit parameters. During all these tests, the individual uncertainties

in Tab. 4.3.1 did not vary by more than 0.1 ms, and the quadrature sum of the uncertainties remained unchanged. All the tests give results for τ_H that are consistent within the uncertainty in Eq. 4.2.

The result and uncertainty in this measurement of τ_H was enough to inform the design of the ACME III interaction region. However, if one were interested to further reduce the uncertainty, the best strategy would likely be to build a single movable optical setup for the excitation laser which is mounted on a rail running along the length of the beamline. This would significantly reduce the primary source of uncertainty, which arises from different excitation efficiencies caused by small differences in laser alignment and beam properties. Second, significantly increasing the intensity of the excitation laser would improve saturation and further reduce this effect. Finally, one could set up a second fluorescence detection region upstream to probe the magnitude of the signal (for example by using the $X - C$ transition) while taking data regularly, removing the need to use the E1 signal for normalization.

5

Detection improvements

In this chapter we discuss the improved photon detection for ACME III, consisting of improved collection optics and replacement of the conventional photomultiplier tube (PMT) detectors with silicon photomultiplier (SiPM) detectors. I was the lead researcher for the collection optics project. I also assisted the Okayama team towards the SiPM detector project and collaborated with Cole Meisenhelder, Takahiko Masuda, and Ayami Hiramoto on the mounting and integration of all components of the photon detection system.

5.1 Fluorescence collection

5.1.1 Collection optics in ACME II

In order to understand the development efforts in ACME III, we shall first describe the fluorescence collection optics previously used in the experiment. The ACME II collection optics consists of two sets of four lens doublets placed on two sides of the molecular beam ($\hat{z}$) right behind the field plates (Fig. 5.1.1). The readout laser beam is sent along $\hat{z}$ to the molecular beam through a diamond-shaped opening in the center.

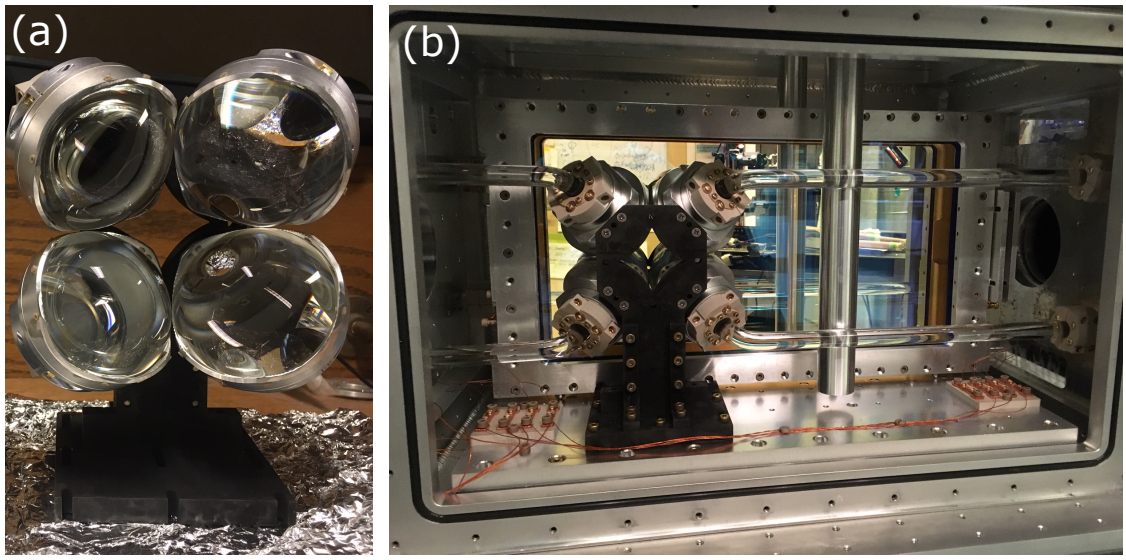

Figure 5.1.1: **ACME II collection optics.** a) Front view, showing the four lens doublets on each side of the experiment. b) View of the optics in the interaction region, showing the curved lightpipes which then exit the vacuum chamber through the front and back of the apparatus. Figure reproduced from Fig. 3.3.1 in Ref. [148].

The lens doublets are mounted with a diagonal orientation (41° from vertical) to point at the molecular light source. Each doublet consists of two plano-convex aspheric lenses with diameters 75 mm and 50 mm,[1] mounted using standard Thorlabs lens tubes and optics adapters (Fig. 5.1.5a). We shall term the two lenses the *large lens* and *small lens* respec-

[1] CVI LAG-75.0-50.0-C-SLMF-400-700 and LAG-50.0-35.0-C-SLMF-400-700.

tively. The lenses were originally chosen because they are among the largest commercially available aspheric lenses [83, §C.1].[2]

Light from each of the lens doublets is collected with a fused silica lightpipe (16 mm diameter) connected to a detector. The lightpipes are curved to align the straight end along $\hat{x}$, allowing them to be routed through magnetic shield endcaps on the front and back, where the detectors are mounted. The combined collection efficiency is about ∼20% [148, §3.3.1].

5.1.2 Design motivations and constraints

The design of the ACME III collection optics is influenced by several motivations and constraints. One of the primary motivations is to adapt to the increase in the field plate (FP) distance d_{FP} from 4.5 cm to 6 cm. This increase allows the electrostatic lens to focus more molecules without having any hitting the field plates (§3.1.2), but requires a commensurate increase in the solid angle subtended by the collection optics. The wider d_{FP} also means the molecules are further away from the collection optics ($\approx$ 8 cm in ACME III versus $\approx$ 5 cm in ACME II). These factors motivated the increased size of the collectors.

Second, the size of the vacuum chamber poses some spatial constraints. The width of the chamber has been fixed to 21" (533 mm) for various practical reasons, which restricts the maximum size of the collectors. Moreover, because of a lack of space to insert curved lightpipes as in ACME II, the decision was made to have straight lightpipes coming out diagonally from the chamber (Fig. 5.1.2). Appropriately placed holes have been cut on the magnetic shield layers following the diagonal path of the lightpipes.

Third, the straight lightpipes lead to some constraints on the angle of the lens mounts. In order to ensure that they all emerged on the sides (rather than the back) of the outermost magnetic shields, the angle is fixed to 45°, as seen in figure 5.1.2.

[2]The collection optics in ACME I were essentially the same except that it used less efficient fiber bundles instead of lightpipes.

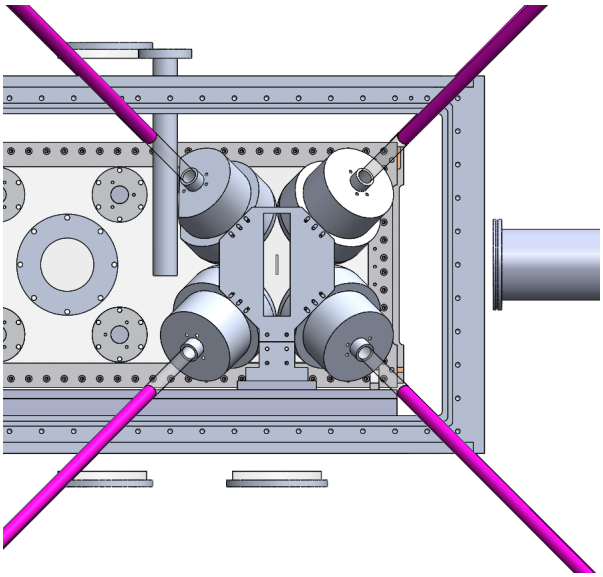

Figure 5.1.2: **Design of ACME III collection optics with lightpipes.** Shown are four lens doublets in custom mounts with straight lightpipes exiting the chamber diagonally. The lightpipes consist of a short (transparent) and long (magenta) segments as described in Section 5.3. (Part of one lens doublet mount has been partially cut without affecting the lenses inside to accommodate a vertical magnetometer pocket.)

Next, the readout laser beam needs to enter and exit the vacuum chamber at the center of the molecular source. It has typical 4σ dimensions of $\sim$3 mm (along $\hat{x}$) and $\sim$31 mm (along $\hat{y}$). This rules out any collection optic design that completely obstructs this optical access along $\pm\hat{z}$.

These constraints contributed towards deciding the basic design of the new collection optics. Alternative designs were considered, such as using elliptical reflectors, compound parabolic concentrators [83, §C.4], or adding collectors above and below the molecular light source. However, these designs were eventually found to be either unfeasible or add significant complexity without a commensurate increase in light collection efficiency. Thus, we settled on using the same basic optics design as ACME I and II, but with larger and more optimized choices of lenses, lightpipes, and their mounting. Since the previous design already used lenses which are among the largest available commercially, these new lenses must be custom manufactured.

5.1.3 Design process

The optics design process consisted of several stages. LightTools ray tracing software is used to calculate the predicted collection efficiency.[3] The simulations impose an anisotropic angular fluorescence distribution of the molecular source, reflecting the two possible distributions expected from the $I - X$ decay which depends on different parity states of I. For $\tilde{\mathcal{P}} = +1$ and $\tilde{\mathcal{P}} = -1$ these two distributions are linearly dependent on $\cos^2\theta$ and $\sin^2\theta$ respectively [133].

In the first stage, the ACME II collection optics setup is carefully reproduced in the simulation to serve as an accurate comparison. The aspheric lenses are modeled after the manufacture specifications [182]. A rotationally symmetric aspheric surface is typically specified by the expression [183]

[3] MATLAB scripts were integrated with LightTools to allow automated, rapid parameter scans.

$$Z = \frac{r^2}{R\left(1+\sqrt{1-(1+k)\frac{r^2}{R^2}}\right)} + Ar^4 + Br^6 + Cr^8 + Dr^{10}, \tag{5.1}$$

where Z is the surface sag, k is the conic constant, R is the radius of curvature, r is the radius of the surface measured from the optical axis, and A, B, C, D are higher-order coefficients. For a plano-convex lens, the distance from the edge of the convex side to the base of the plane side is called the thickness. The values of these parameters for the ACME II large and small lens can be found in Table 5.1.1. (Note that all the higher-order coefficients are zero, and $k = -1$ for a parabolic surface.)

The lightpipe is modeled as cylindrical with diameter of 16 mm. The material of all optics is assumed to be BK7 (index of refraction $n = 1.52$). The molecular light source is modeled as a uniform, cuboid, distribution of photons with dimensions 3 mm $\times$ 28 mm $\times$ 28 mm.

For the decays from $\tilde{\mathcal{P}} = \pm 1$, these simulations found light collection efficiencies ϵ of

$$\epsilon_{\cos^2\theta} = 0.29 \tag{5.2}$$

$$\epsilon_{\sin^2\theta} = 0.19, \tag{5.3}$$

which are reasonably close to the $\epsilon = 20\%$ measured in ACME II [148, §3.3.1]. The relative size of $\epsilon_{\cos^2\theta}$ to $\epsilon_{\sin^2\theta}$ remains the same for many different configurations, so we use only the smaller value of $\epsilon_{\sin^2\theta} = \epsilon$ for comparison throughout the rest of this section. We shall also denote the gain in ϵ relative to ACME II as G_ϵ.

We then model the expected molecular light source in ACME III. Based on numerical simulations with the electrostatic lens (§3.1.2), this results in a Gaussian (rather than uniform) distribution of molecules along $\hat{y}$ and $\hat{z}$ with $\sigma \approx 7$ mm. The width along $\hat{x}$ is determined by the readout laser width. Thus, we use a light source of the same dimensions (3 mm $\times$ 28 mm $\times$ 28 mm) imprinted with a $\sigma = 7$ mm Gaussian intensity distribution

(convolved with the aforementioned anisotropic distribution to reflect the angular distribution from the $I - X$ decays). The position and distances of the collection optics assume $d_{fp} = 6$ cm and an FP thickness of 0.5".[4]

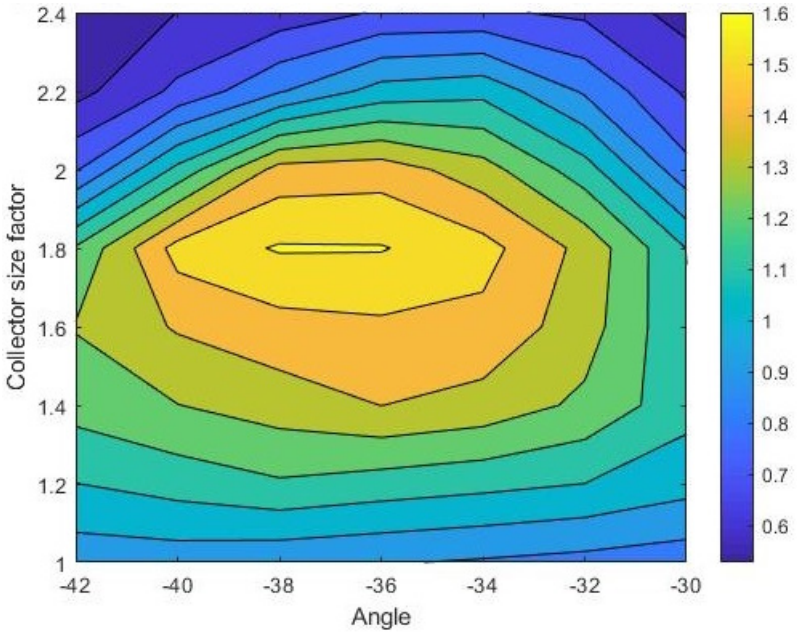

Figure 5.1.3: **Preliminary explorations of enlarging collection optics.** Plot of G_ϵ versus collector size enlargement factor and angle, obtained from LightTools simulations. The diameter and distances of the lenses are proportionally increased by a collector size factor while keeping the lightpipe diameter fixed at $d_{LP} = 24$ mm. The enlarged collectors are also proportionally moved away from the center to avoid conflicts.

Next, we perform a preliminary exploration of the possible G_ϵ while increasing the overall size of the collectors. As the four lens doublets on each side are larger, their angle and position must be adjusted to avoid colliding with each other, the readout laser beam, or the field plates while still pointing at the molecular source. This results in plots such as in figure 5.1.3. We found that the optimum G_ϵ results from enlarging the collectors by no more than a factor of ∼1.8, equivalent to having a large lens diameter $d_{large} = 135$ mm.

[4]During this stage, we had not yet decided to set $d_{fp} = 6$ cm or 7.5 cm, so the simulations were done

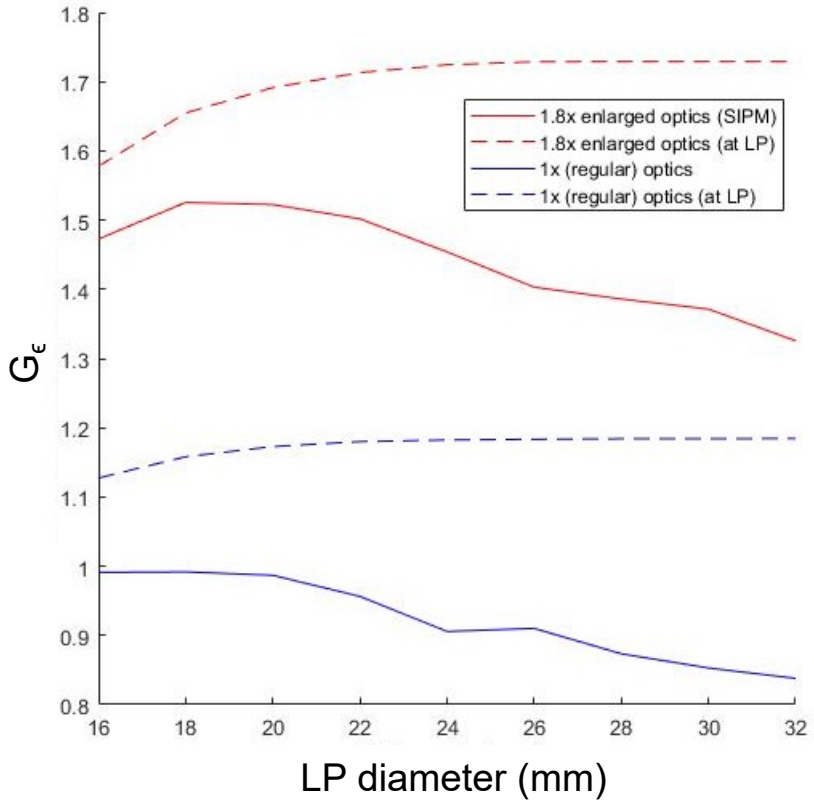

Figure 5.1.4: **Varying lightpipe diameter with a detector.** G_ϵ versus d_{LP} recorded at lightpipe (dashed lines) and SiPM placed 1 cm away (solid lines). Plots are shown for 1.8× enlarged optics (red upper pair of lines) and original ACME II optics (blue lower pair of lines). For simulations with SiPM, the gain decreases at larger LP diameter due to the wider range of collection angles emitted from the LP, some of which are not captured by the detector.

We subsequently perform more fine-tuned optimizations, including on the thickness, curvature, distances, and relative diameters of the two lenses, incorporating the aforementioned constraints (§5.1.2). We find that some additional gains could be made by optimizing the lens doublet mounts. We also incorporate the SiPM detectors (§5.2), which are modeled as a square 24 mm × 24 mm detector placed 1 cm away from the lightpipe. The distance arises from the thickness of the various filters placed before the SiPM. This constrained the diameter of the lightpipe d_{LP}: a larger diameter captures more rays from the lenses, but this is offset by the larger angular spread on the output (Fig. 5.1.4). Eventually, we set $d_{LP} = 20$ mm.

Finally, assuming $d_{large} = 135$ mm and $d_{LP} = 20$ mm, a multivariable algorithmic optimization is performed incorporating all variables: large lens curvature and thickness; small lens diameter, curvature, thickness, and position; and light pipe distance. This results in the final design shown in Fig. 5.1.5b, with the aspheric lens dimensions shown in Table 5.1.1. The figure also shows the custom lens mount designed for the lens doublet and the lightpipes. The simulation predicts this configuration to result in $G_\epsilon = 1.6$ for $d_{fp} = 6$ cm.[5]

	Large lens		**Small lens**	
Quantity	**ACME II**	**ACME III**	**ACME II**	**ACME III**
Diameter (mm)	75	135	50	106
Radius of curvature (mm)	26	54	18	32
Focal length (mm)	50	104	35	62
Thickness (mm)	29	47	21	45

Table 5.1.1: **ACME II and III aspheric plano-convex lens properties.** See Eq. 5.1 and explanation in text for reference. All units are in mm. Shown are the lens diameter $d = 2r$, radius of curvature R, focal length $f = R/(n-1)$, and thickness. The conic constant k is set to -1, and the higher-order coefficients are all zero.

with the aim of producing a configuration that works reasonably well for both.

[5]The design and configuration was actually optimized for $d_{fp} = 7.5$ cm (obtaining $G_\epsilon = 1.5$) while ensuring that it still had reasonable performance at 6 cm because at the time we had not definitively decided on the value of d_{FP}. It is possible that G_ϵ could be increased by another 10-20% if the design was reoptimized for 6 cm.

This gain can be compared to an estimate of the gain in the solid angle Ω subtended by the new collection optics. The new collectors are roughly 1.8× larger than before, but the molecular light source is ∼1.6× further away due to the increased d_{FP}. Since $\Omega = A/r^2$, the solid angle in ACME III is ∼1.3× larger. This is slightly lower than the simulation $G_\epsilon = 1.6$ likely because it does not account for the optimization in the relative sizes, curvatures, and positions of the lenses.

5.1.4 Manufacturing and testing

The lenses have been custom manufactured from BK-7 glass by Asphera, including AR coating optimized for 512 nm.[6] Simulations showed that a tolerance of 0.5 mm on the diameters was sufficient. The mounts for the lenses were custom made by the Harvard Scientific Instrument Shop, including custom threads and retaining rings for both lenses.

In October 2021, an experimental test was performed at Harvard comparing the new optics with the old ACME II optics (Fig. 5.1.6), in conjunction with Okayama University tests of the SiPM detectors. A fiber with light from a 512 nm laser is inserted into a Delrin ball to serve as a mock molecular light source. A detector is placed after the lightpipe. The test found $G_\epsilon = 1.7$, which is close to the predicted gain. Changing the distance of the lightpipe to the lens by several mm did not change this result by more than 10%, which means that small misalignments in inserting the lightpipes will likely not be a major issue. Using measurements of the ACME II collection efficiency as a baseline, ACME III has an expected collection efficiency of ∼30%. A more complete test is expected to be performed once all the optics are assembled in their final configuration in the experiment vacuum chamber.

[6] The total cost was about $\$20,000$ for 10 pairs of large and small lenses.

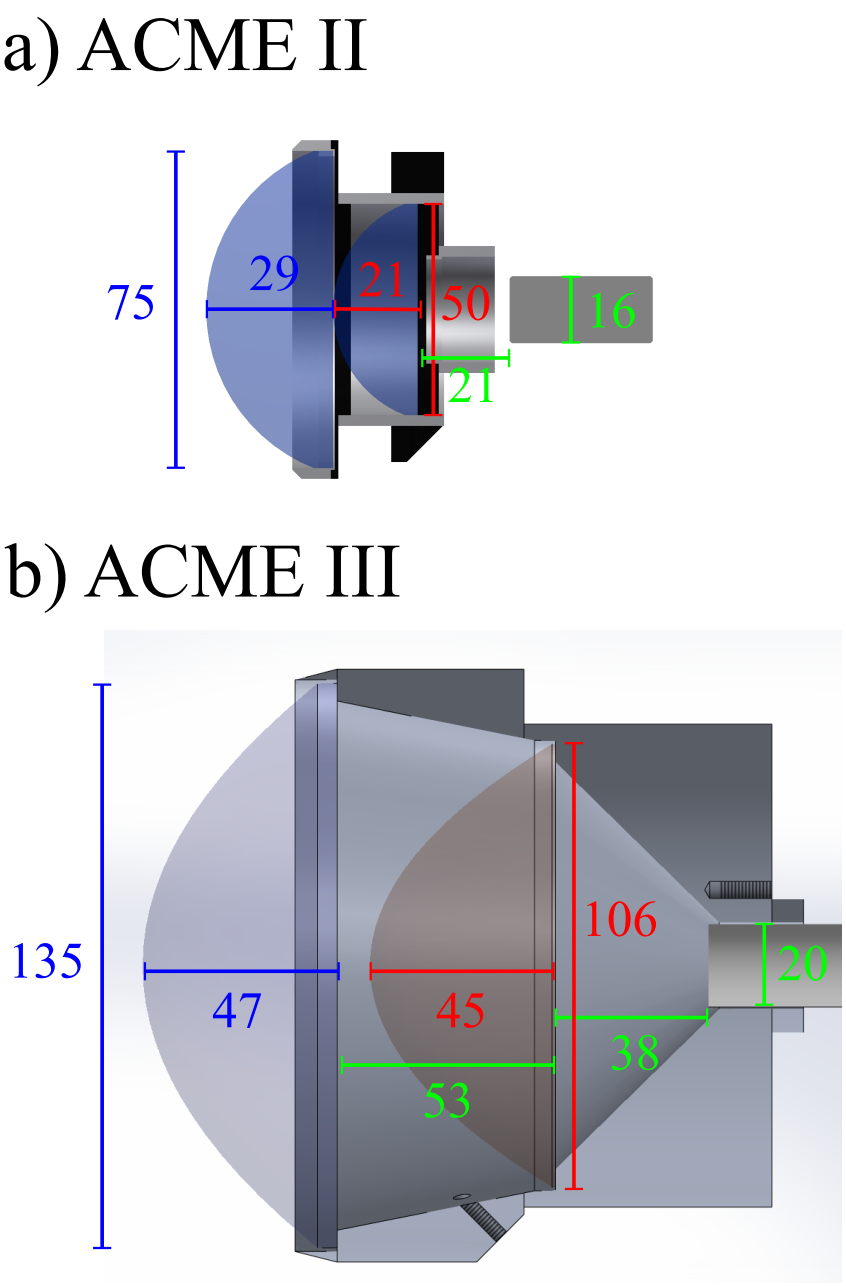

Figure 5.1.5: **Scaled comparison of collection optics used in ACME II and III.** All units are in mm. a) Lens doublet used in ACME I and II, consisting of two commercial aspheric lenses mounted in standard Thorlabs parts. ACME II used a lightpipe to collect the photons and transfer them to the PMTs outside the vacuum chamber [148]. Adapted from Fig. C.2 in Ref. [83]. b) ACME III larger lens doublet and lightpipe in a custom mount, obtained after the design process described in Section 5.1.3.

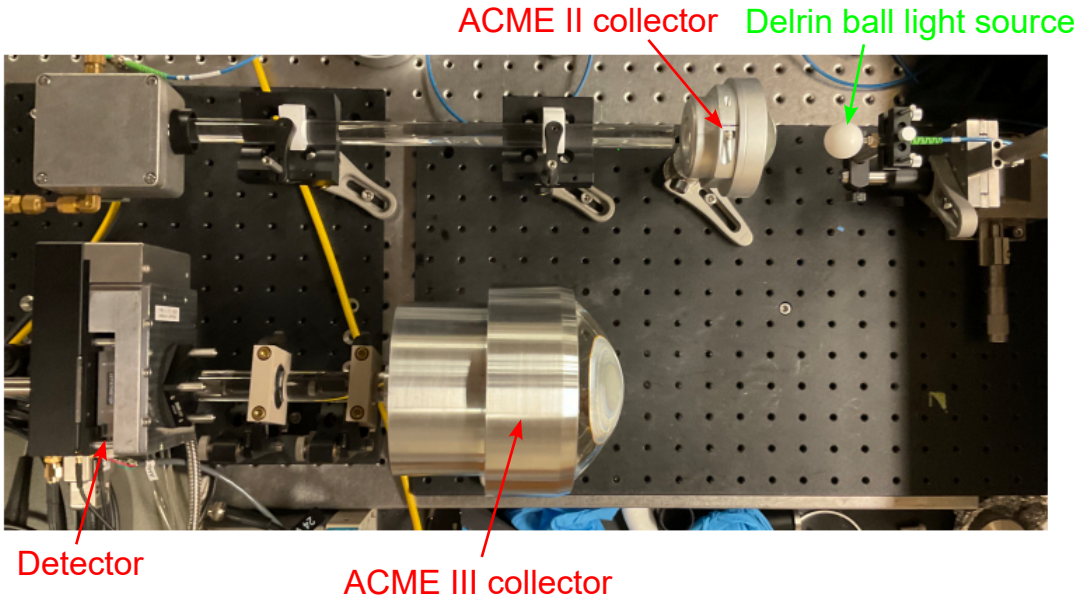

Figure 5.1.6: Setup used for experimental test of the new collection optics. Photo taken by Ayami Hiramoto.

5.2 Silicon photomultipliers

After the photons from the molecules are captured by the lens doublets and collected into the lightpipes, they are picked up by detectors placed just outside the interaction region. In ACME II, a conventional photomultiplier tube (PMT) detector (Hamamatsu R7600U-300) was used for each of the eight lightpipes and mounted on the endcaps of the experiment. These PMTs have a quantum efficiency of $\approx 25\%$ at 512 nm and an effective detection area of 18 mm $\times$ 18 mm [184].

Silicon photomultipliers (SiPMs) are a viable alternative to replace PMTs. They consist of an array of small silicon avalanche photo detectors (APDs) operating in Geiger mode, commonly called microcells or pixels [185]. While in operation, a bias voltage is applied to the SiPM. When a photon impinges on a pixel, it triggers an internal avalanche of photoelectrons that results in a large current. SiPMs have the advantage of having better quantum efficiency than PMTs. However, one major disadvantage is a high dark count rate (DCR), which is typically $\sim 3 \times 10^7$ cps per detector [186]. In contrast, the DCR

of a PMT is typically several orders of magnitude lower [184]. In ACME II, the signal per detector was $\sim 4 \times 10^7$ cps, which meant using SiPMs operating at room temperature would have been too noisy. In ACME III, the signal is projected to be larger by over an order of magnitude. Thus, the DCR requirement is less stringent.

The common method to suppress DCR is to operate SiPMs at subzero temperatures, as the DCR decreases by a factor of ∼2 every 10° C [185, 187]. This can be obtained using standard thermoelectric coolers (TECs). To avoid condensation and freezing, each SiPM must be placed in a vacuum chamber.

5.2.1 ACME III SiPM module design

There have been some past efforts at developing SiPM detectors for ACME [70, §6.3]. In 2019, the Okayama University team took over the project and has since developed a fully functioning SiPM module. The module characteristics and experimental results have been described in detail in Refs. [167, 168, 172]. Here we shall provide a summary.

A schematic of the module can be seen in figure 5.2.1. It consists of a Hamamatsu S13361-6075NE-04 SiPM, which has a 24 mm × 24 mm active area, quantum efficiency of ∼45% at 512 nm, and DCR $\sim 2 \times 10^6$ cps per channel, with 16 channels total [186]. A TEC cools the SiPM to −15° C while the module is under vacuum. This reduces DCR to 3% of the room temperature value, or $\sim 10^6$, which is about two orders of magnitude lower than the projected ACME III signal. Water cooling lines are connected to a thermoelectric chiller to stabilize the overall temperature of the module.

Between the lightpipe and the SiPM active area are three bandpass filters. The first is an interference filter (Semrock FF01-520/70) with a high extinction ratio to suppress scattered light in the interaction region. (This is the same filter used in ACME II with the PMTs.) The other two are absorptive filters (Schott BG39 and BG40) required to suppress optical cross-talk (OCT). OCT is a phenomenon where secondary photoelectrons produced

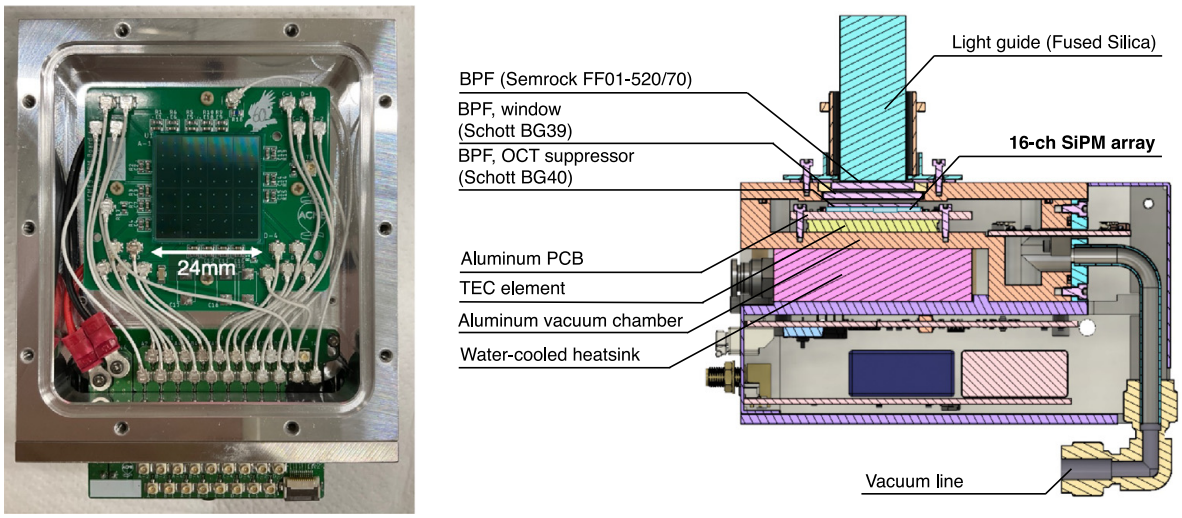

Figure 5.2.1: **ACME III SiPM module.** Photo (left) and cross-section of SiPM module developed for ACME III. Figure reproduced from Ref. [168].

by a SiPM pixel are reflected off an interference filter and redetected by a different pixel, resulting in excess noise [188]. By adding an absorptive filter in between, OCT probability is reduced from 24.7% to 4.3% [168, 172].

Each SiPM module has its own electronics package. At each channel, pole-zero cancellation technique [189] is used to stimulate a fast response. Outputs from the 16 channels are then summed into one channel using an operational amplifier and shaped by a Bessel filter. Thus, each detector has only one readout for a total of 8 detectors, the same number as in ACME II. This means the same DAQ device can be reused for ACME III (§6.2). At the same time, each module features a separate output with a multiplexer to allow reading of individual channels for diagnostic purposes. More details on the SiPM electronics can be found in Ref. [167].

The Fano or *excess noise factor* F of a detector is obtained when measuring signal from single photon pulses, defined as

$$F = \frac{\sigma^2}{\mu}, \tag{5.4}$$

where μ and σ are the mean and standard deviation of the number of observed photons,

with $F = 1$ for an ideal Poisson process. The average F measured for the SiPM modules is 1.07 [167], which is slightly better than the $F \sim 1.2$ measured for the ACME II PMTs [148, §3.3.6].

5.2.2 Experimental tests

Experimental tests of the SiPM module have been carried out at Okayama and Harvard. A test using a Delrin ball light source (§5.1.4) was carried out comparing the SiPM and ACME II PMT performance, finding a gain of ∼2.7. A more direct test was later carried out by mounting the SiPM module on the ACME II apparatus in place of the PMT [167]. The same setup as the molecular lens test (§3.1.2) was used, where molecules are transferred into the Q-state, then detected with laser-induced fluorescence on the 746 nm $Q - I$ transition, producing 512 nm photons from the $I - X$ decay. These tests found a gain of ∼2.3 in signal (Fig. 5.2.2). A total of ten SiPM modules have been built, and a long-term test found that their signals are stable to ∼1 %.

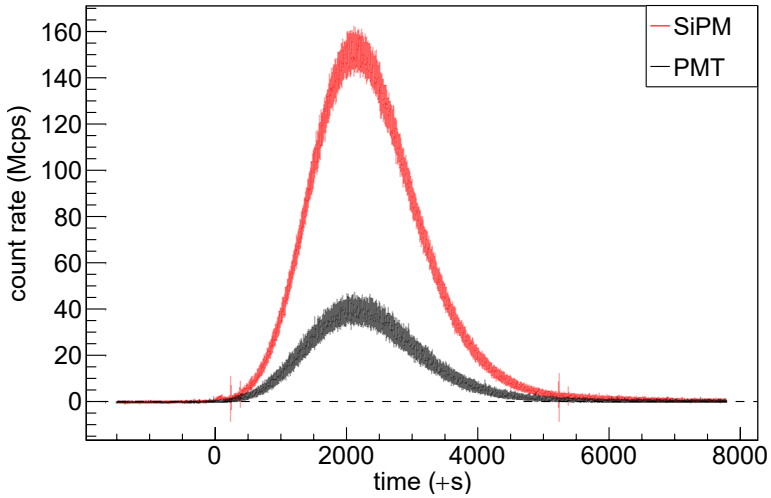

Figure 5.2.2: **Comparison of PMT and SiPM.** Signal detected from laser-induced fluorescence of ThO using PMT (black) and SiPM (red). Figure reproduced from Ref. [167].

5.3 Lightpipes and mounting of components

We shall now discuss how the SiPM, collection optics, and lightpipes (LPs) are integrated together and mounted onto the ACME III experimental chamber. The components in this mounting system were designed by Cole Meisenhelder and myself in collaboration with the Okayama team.

The lightpipes (LPs) are manufactured from fused silica[7] by Technical Glass Products, with a diameter of 20 mm. There are two LP segments (short and long, as seen in figure 5.1.2). The short segment ($L = 19.1$ cm) is inserted into the lens doublet mount to collect the light from the lenses. It exits the vacuum chamber through an angled flange attached to the wall of the chamber (Fig. 5.3.1b), which grips the other end of the LP with two O-rings, the inner one of which forms a vacuum seal.

The angled flange secures the point of connection with the long LP segment ($L = 85.1$ cm). This segment goes through the three magnetic shield layers to the SiPM module mounted just outside. A PEEK collar is securely wrapped around one end of the LP with two internal O-rings. The collar fits snugly into a slot in the flange. The flange is designed to allow the faces of the two LP segments to meet as close as possible with each other without touching.[8]

The other end of the long LP is supported by a custom plastic mount (commonly termed the "M2" mount), which secures the LP with a pair of O-rings (Fig. 5.3.2). This mount is connected to a standard optical post via a ball-and-socket mount[9] allowing for flexible rotation. The post rests on a small breadboard mounted on the frame of the magnetic

[7] The possibility of using acrylic lightpipes was explored as they would be lighter and less fragile, but a major issue was the lower bulk transmission of acrylic resulting in a loss of signal.

[8] Because the flange is normally physically inaccessible, it is not practical to use index-matching gel to interface between the two LP faces, which might require cleaning every time the LP is removed. Some tests were also performed to explore the possibility of using a solid, transparent silicone disk as an interface, but this was eventually not done because of the difficulty of mounting the disk securely and the predicted $\sim$10% loss in signal.

[9] Thorlabs TRB2.

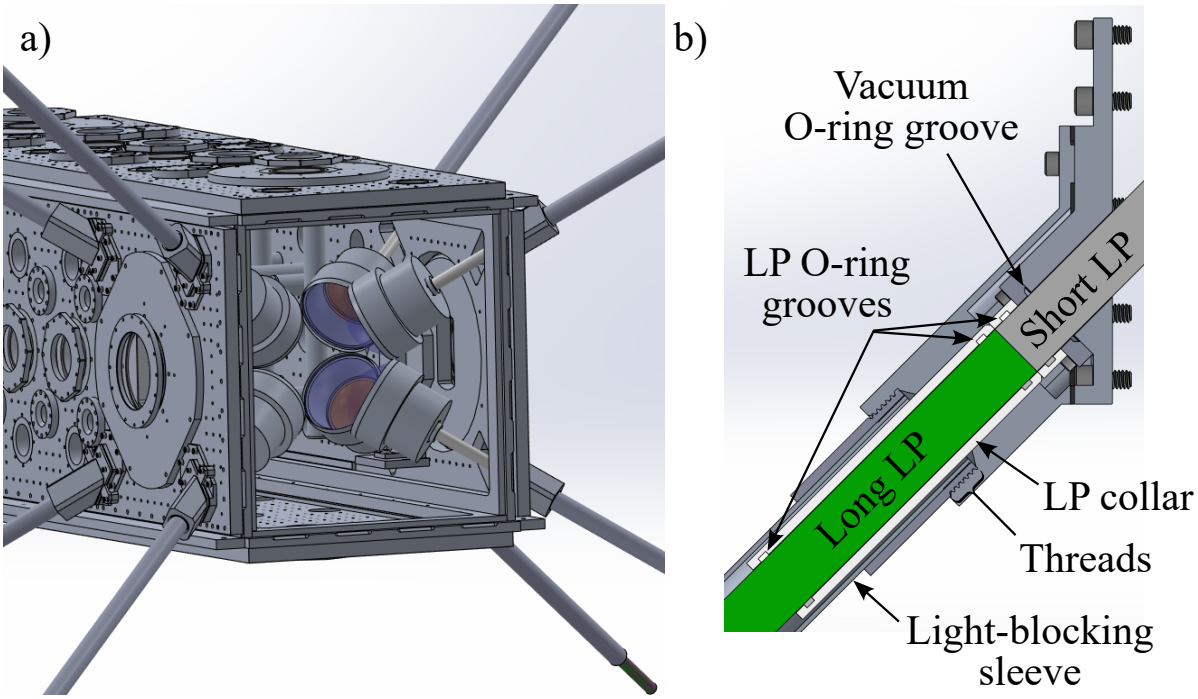

Figure 5.3.1: **Lightpipe angled flange.** a) View of vacuum chamber with collection optics lens doublets, LPs, and angled LP flange. The field plates and a vacuum flange have been removed for clarity. b) Cross-section of LP angled flange (designed by Cole Meisenhelder), where the short and long LPs connect. The long LP is held in place by friction from the LP collar. A sleeve with external threads is installed to the flange for light-blocking.

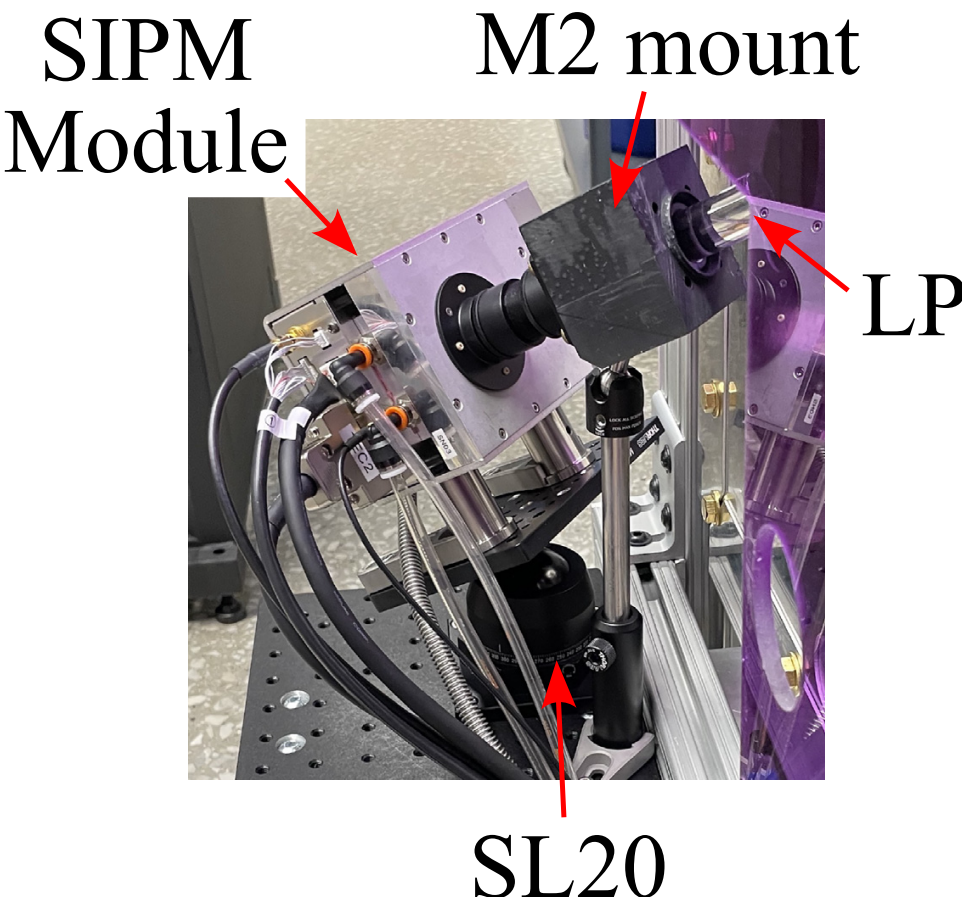

Figure 5.3.2: **SiPM and M2 mount.** Photograph of SiPM module mounted on an SL20 platform and aligned with the M2 mount, which also holds up one end of the long LP. (The other end is held by the angled LP flange shown in Fig. 5.3.1.) Light-blocking components (LP sleeve and rubber bellows for the SiPM lens tubes) are not shown. The SiPM and the M2 mount rest on a breadboard mounted on the magnetic shield guards. Figure adapted from [168].

shield guards. This breadboard also supports the Thorlabs SL20, a rotatable mounting platform for the SiPM module. Standard lens tubes are attached to the SiPM module to assist alignment with the M2 mount and the LP.

5.4 Stray light suppression

Unlike the previous lab at Harvard where ACME I and II were located, the new experimental space at Northwestern is located in a larger room shared with other experiments. Here it is impractical to turn off the room lights while operating. Thus, the new experiment must be designed with much better blocking of stray light to allow it to operate with all the lights turned on. As such, many components have been developed with light-blocking capabilities in mind. The LP angled flange (Fig. 5.3.1b) and M2 mount (Fig. 5.3.2) both contain a circular groove to allow an aluminum sleeve (threaded on one end) to encase the LP. This sleeve is to be installed before inserting the LP. Similarly, the connection from the M2 mount to the SiPM is encased with a rubber bellows (not shown in Fig. 5.3.2). All components surrounding the LP are painted black for maximum light absorption. Other light-blocking components (such as for the optical access holes in the magnetic shields) are still in development.

5.5 Conclusion

The improved detection system is projected to substantially improve photon detection efficiency in ACME III. The newly designed collection optics system is projected to give a 1.7 times gain in signal, while the SiPM detectors are projected to give a 2.3 times gain. Both of these components have been designed, built, tested, and integrated into a unified system. Together, they are projected to produce a 3.9 times improvement in overall detection efficiency to $\sim$ 20%, equivalent to a 2.0 times gain in EDM statistical sensitivity. Work remains to assemble all components in their final form on the apparatus

and implementing measures to suppress stray light.

6

Data acquisition and experimental control

The ACME III data acquisition system is designed towards two main goals. First is to eliminate the excess noise caused by the data acquisition timing, which initially prevented ACME II from reaching the shot-noise limit [148]. Second is to upgrade the system to be able to save traces of all molecular pulses (instead of only averages), and adequately deal with the resulting large volumes of data. We will discuss both of these goals in this chapter. Finally, we will discuss the experimental control system in ACME III and how it is integrated with the new DAQ system.

6.1 Timing asymmetry noise

As mentioned in §2.1.10 and §2.2.9, the ACME II final dataset had at best $\chi^2 \approx 3$, signifying the presence of excess noise which prevented it from reaching the shot-noise limit (Eq. 1.23). In 2019, investigations led by Cris Panda, Mohit Verma, and Cole Meisenhelder revealed that the noise was caused by timing jitter in the DAQ system coupling to a dependence of the asymmetry on the timing structure of the polarization switching, which has been reported before in Ref. [161]. Here we shall give a summary of the mechanism of this noise.

6.1.1 Asymmetry noise mechanism

In ACME II, polarization switching of the readout laser occurred at 200 kHz, or a period of $T = 5$ μs for each switching cycle (Fig. 6.1.1a). T is divided into $T_{\hat{X}}$ and $T_{\hat{Y}}$ polarization time bins. Within each bin, the $\hat{X}$ or $\hat{Y}$ laser is turned on for $T_{on} \approx 1.9$ μs and off for $T_{off} \approx 0.6$ μs. To implement the polarization switching, AOMs controlled by a hardware timing box[1] are used. Ideally, we aim to set $T_{\hat{X}} = T_{\hat{Y}} = T/2$. In practice, AOMs have a rise time of ~200 ns, causing an asymmetric delay between the pulses, $\Delta T_{\hat{X}-\hat{Y}} = T_{\hat{X}} - T/2$. Thus, we typically adjust $T_{\hat{X}}$ and $T_{\hat{Y}}$ such that $\Delta T_{\hat{X}-\hat{Y}} \approx 0$. However, this zeroing process is not perfect, which will have consequences described later.

The fluorescence records are taken at 15 M samples/s, or 80 samples every $T = 5$ μs. In the ACME data analysis procedure (§2.1.10), T is divided into two bins of 40 samples each which are assigned as belonging to S_X and S_Y respectively. The signal has time-dependent shape, which is determined by the quantum population dynamics dependent on laser and molecular beam properties [153, §4.2] (Fig. 6.1.1a). We then integrate the signal in a sub-region within the time bin (normally samples #3, 4, 5, ..., 34), from which the

[1] SRS DG645.

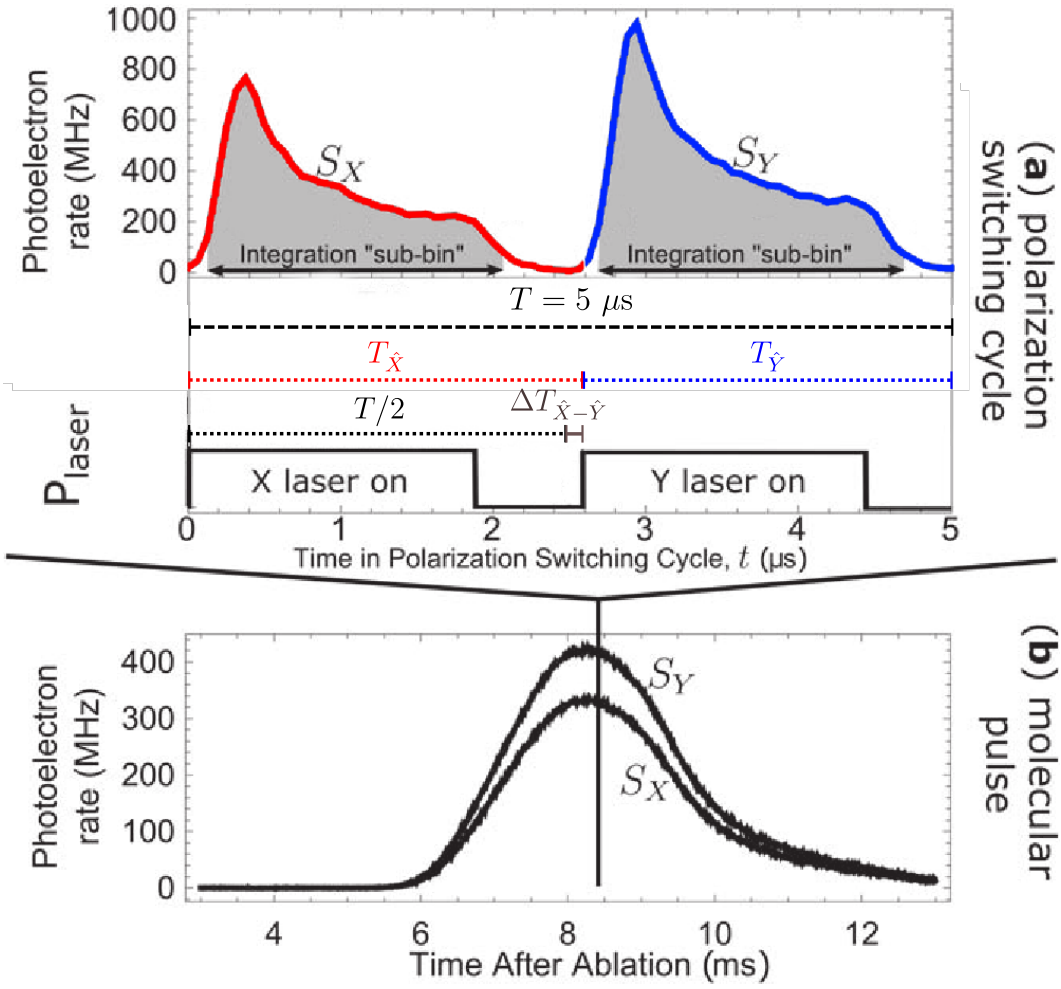

Figure 6.1.1: **Polarization switching timing structure.** a) Close-up of fluorescence data with polarization switching timing structure, taken from the 10 ms trace of the molecular pulse shown in b). Light from the $\hat{X}$ and $\hat{Y}$-polarized lasers are switched at a period of $T = 5\ \mu$s. Physical effects in the AOM cause slight timing mismatch $\Delta T_{\hat{X}-\hat{Y}}$ which may not be zeroed out perfectly. Figure adapted from Ref. [161].

asymmetry

$$\mathcal{A} \equiv \frac{S_X - S_Y}{S_X + S_Y} \tag{6.1}$$

is computed.

However, if $\Delta T_{\hat{X}-\hat{Y}}$ is not exactly zero, then there will be a mismatch between the time-dependent $S_X(t)$ and $S_Y(t)$, which causes the asymmetry to have a time dependence $\mathcal{A}(t)$. Assuming that $\Delta T_{\hat{X}-\hat{Y}}$ is small, then $S_Y(t) \approx S_X(t) + S'_X(t)\Delta T_{\hat{X}-\hat{Y}}$. Defining $S(t) \equiv (S_X(t) + S_Y(t))/2$ and using Eq. 6.1, we obtain

$$\mathcal{A}(t) \approx -\frac{S'(t)}{2S(t)}\Delta T_{\hat{X}-\hat{Y}}. \tag{6.2}$$

Thus, the value of integrated asymmetry $\mathcal{A}$ is dependent on $\Delta T_{\hat{X}-\hat{Y}}$ and the choice of integration sub-bin. This becomes an offset in the measured phase Φ which is $\tilde{\mathcal{N}}\tilde{\mathcal{E}}$-even and so does not affect the EDM phase $\Phi^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$. If it is constant, it will simply be subtracted out. However, investigations found that a non-zero $\Delta T_{\hat{X}-\hat{Y}}$ (of the order of tens of ns) combined with a global timing jitter in the DAQ system causes a noise which is too fast to be averaged out by even our fastest experiment switches, as seen in figure 6.1.2a. The magnitude of the noise depends on ΔT_{X_Y} and time within the polarization switching cycle, with the largest variance seen in the initial rising edge.

6.1.2 Origin and suppression of noise

The large timing variance in the DAQ was eventually found to be caused by two issues. First is the internal clock of the NI PXIe-5171 digitizer. The polarization switching and the triggering of the digitizer is implemented using the SRS DG645 delay generator, which has a jitter of < 25 ps [190]. However, the timing of the digitizer sampling is controlled by its 250 MHz internal clock, which has an acurracy of only 25 ppm [191]. Second is the chosen sampling rate of 15 MHz, initially chosen to ensure an even number of points in $\hat{X}$

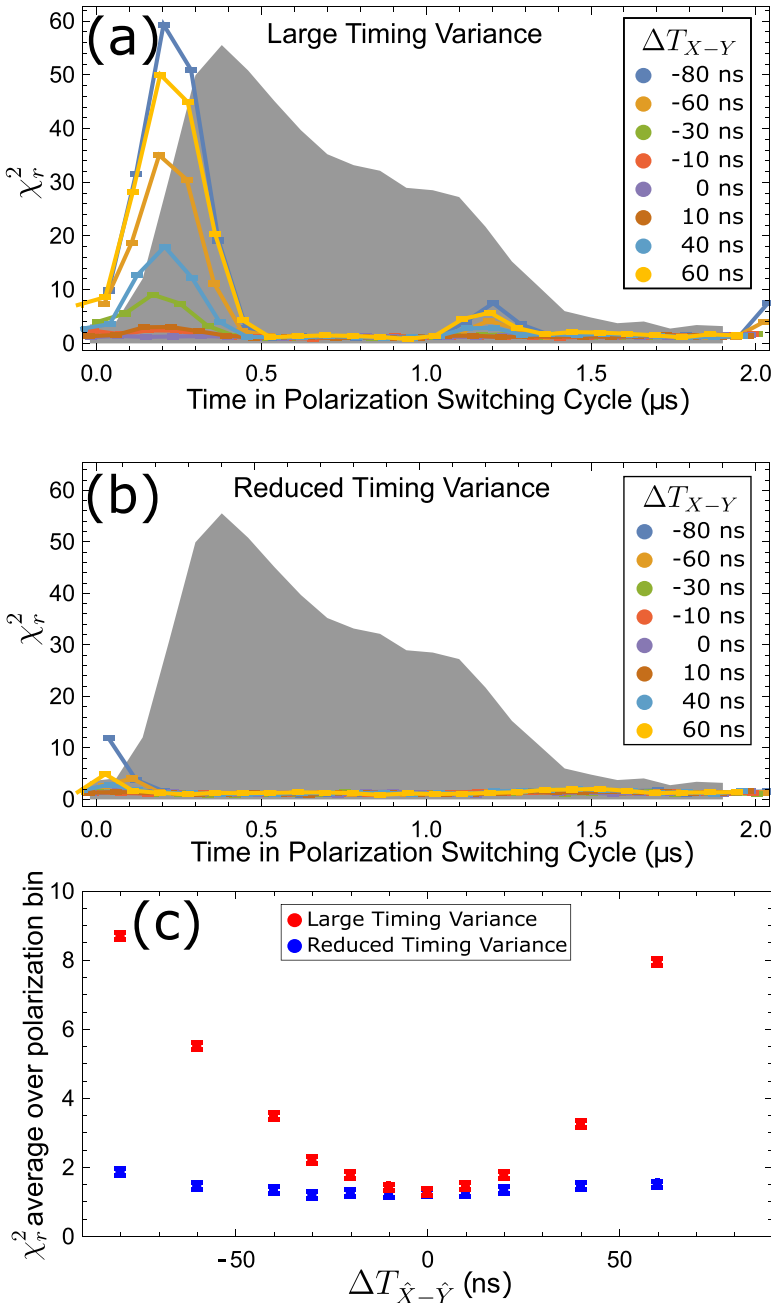

Figure 6.1.2: **Timing variance and asymmetry noise.** For these tests, we set $T = 4$ μs. a) Noise (measured by χ_r^2) within a polarization switching bin for different values of the time delay $\Delta T_{\hat{X}-\hat{Y}}$ with large timing variance in the DAQ. b) Noise after timing variance has been reduced using the methods described in §6.1.2. c) Average χ_r^2 versus $\Delta T_{\hat{X}-\hat{Y}}$ with large timing variance (red) and reduced timing variance (blue). Figure reproduced from Ref. [161].

and $\hat{Y}$ polarization bins. However, this is a non-integer divisor of the 250 MHz internal clock frequency and could not be properly processed by the digitizer firmware.

These two issues resulted in a periodic global timing offset that varied by ∼10 ns between subsequent traces, then reset to zero once it reached ∼100 ns. To suppress this jitter, we set the sampling rate to an integer divisor of 250 MHz (such as 12.5 MHz) and sync the internal clock to an external Rb frequency standard.[2] To maintain an even number of samples in the polarization switching bins, the polarization switching cycle time was changed to $T = 4.0$ μs, resulting in 50 samples per cycle.

Figure 6.1.3 illustrates this phenomenon, showing a series of 25 successive traces recorded by the digitizer when it is triggered at 50 Hz by the delay generator while being fed a synchronized rising edge from the same device. (Unlike in the ACME II dataset, the raw traces are recorded without being averaged.) When the the sampling rate is improperly set to 16 MHz and the internal digitizer clock is left to run on its own, there is noticeable jitter of ∼1-2 samples or ∼100 ns (figure a). The jitter remains when the sample rate is corrected to an integer divisor of 250 MHz without syncing the clock (figure b) or vice versa (figure b), and is only eliminated when both of these issues are corrected (figure d). This configuration was also used to produce the plots of figure 6.1.2. Similar tests were performed with the spare NI PXI-5922 DAQ system with the same results.[3]

A method to further suppress this noise is to choose the integration sub-bin in the data analysis so as not to start or end in the portion of the bin when the change in $S(t)$ is large [161]. A more ambitious possibility is to estimate the timing jitter for individual traces and adjust the integration sub-bin to correct for it [70, §4.15.4]. While applying this technique retroactively on ACME II data could have reduced the noise, it could never fully eliminate it as data was only recorded in averaged traces, not pulse-to-pulse.

After discovering the mechanism of the noise and correcting it with the techniques de-

[2] SRS FS725.

[3] For this system, the sampling rate has to be set to an integer divisor of 15 MHz.

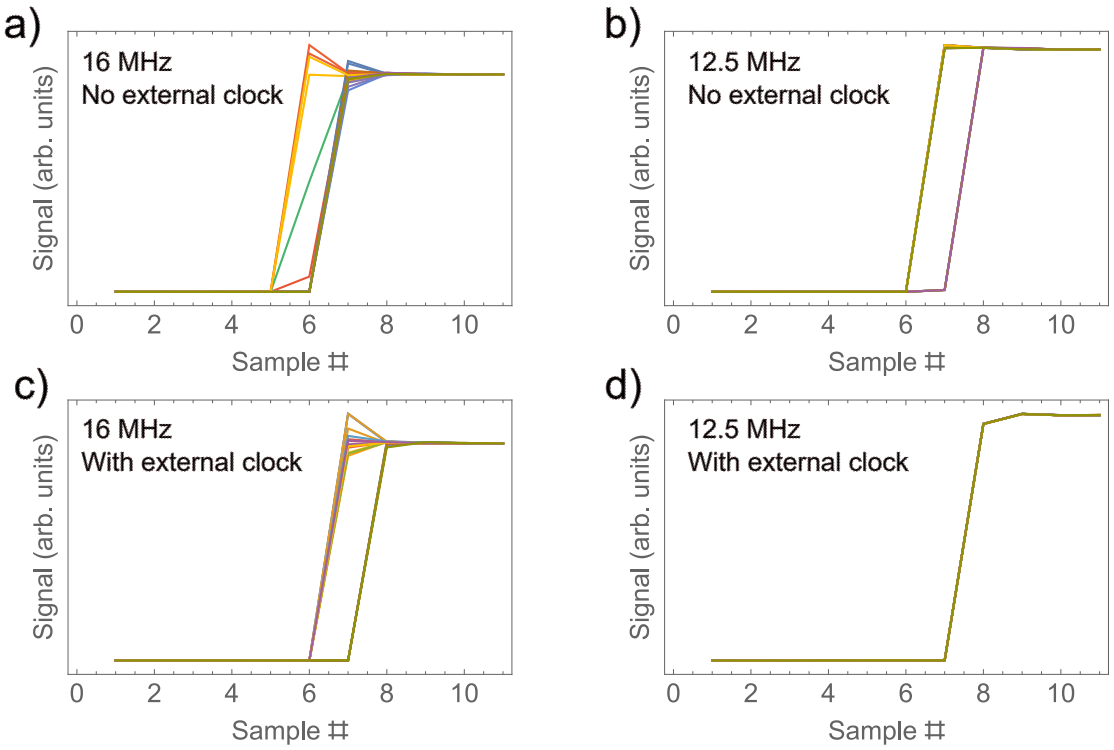

Figure 6.1.3: **Results of digitizer jitter test.** For these tests, the digitizer is triggered at 50 Hz while a synchronized square wave pulse train is sent to one of its input channels. In each plot, only the rising edge of the first square wave of 25 consecutive records are displayed. a) Traces taken at 16 MHz sampling rate, internal clock is not synced. A jitter of 1-2 samples (∼100 ns) is observed. b) Traces taken at corrected 12.5 MHz sampling rate, no syncing. c) Traces taken at 16 MHz, with external clock. d) Traces taken at 12.5 MHz sampling rate, with syncing to an external Rb frequency standard. Here, the jitter is entirely suppressed.

scribed above, a small EDM dataset was taken in the new configuration, which found χ^2_r consistent with 1, indicating that the shot noise limit had been achieved. Since then, the new DAQ system in ACME III has been set up in Northwestern incorporating the above technical improvements.

6.2 Data acquisition in ACME III

6.2.1 ACME III DAQ basic requirements

The basic requirements of the ACME III data acquisition (DAQ) system are similar to ACME II. There are eight SiPM detectors whose outputs (after all electronic summing and amplification) are in a similar dynamic range as the ACME II PMT outputs ($V_{pp} \approx 1$–2 V). The frequency of the molecular pulses stay the same at 50 Hz. The horizontal resolution required is also about the same, because the same $I - X$ decays are used. Thus, we decided to use the same DAQ digitizer as before.[4] The aging, frequently unreliable PXIe chassis used in ACME II[5] has now been replaced with a newer model.[6]

As discussed in the previous section, in order to prevent the timing variance that led to the excess noise in ACME II, the sampling rate must be adjusted to 12.5 MHz, the internal clock synced to an external frequency reference, and the polarization switching rate changed from 200 kHz to 250 kHz ($T = 4.0$ μs). Based on systematic checks we performed in ACME II, the different polarization switching rate should not cause any effect on $\omega^{\tilde{\mathcal{N}}\tilde{\mathcal{E}}}$. It will diminish the horizontal resolution of the traces by decreasing the number of samples within a polarization switching cycle from 80 to 50, or 25 per polarization. This should not be a problem as we are still able to resolve the dynamics within the readout beam (Fig. 6.1.2).

[4] NI PXe-5171.
[5] NI PXIe-1075.
[6] NI PXIe-1092.

6.2.2 Saving data from all molecular pulses

A major upgrade which we decided to pursue for ACME III is to add the capability to regularly save all records[7] of molecular pulses acquired at 50 Hz instead of saving only the average of 25 pulses. This enables more granular control over our data acquisition and analysis. Such capabilities were required to solve DAQ communication (Sec. 2.1.8) and the timing noise jitter problems (Sec. 6.1). In ACME II, it was not possible to save all records for more than short bursts. The first limitation was hardware: saving all records dramatically increases the data streaming rate, requiring it to be saved to a faster local SSD instead of the NAS device. This limits the amount of data that can be taken continuously. The second limitation was software: the LabView DAQ VI was simply unable to keep up with the data streaming rate and frequently ran out of memory after about an hour of operation, even when saving to SSD.

Let us estimate the data streaming rate (DSR) required for saving all pulses. Several factors contribute. First is the sampling rate, which we set at 12.5 MHz as previously explained. Second is the length of each record. The maximum record length is the period of the DAQ triggers (20 ms). which determined by the width of the distribution of arrival times of the molecular pulse. In ACME II, the 10 ms record length was sufficient for the 2 ms FWHM of the molecular beam, including a ∼2.5 ms region for sampling the background before the pulse. In ACME III, the molecules will travel ∼1.9× further before reaching the detection region, which may increase the pulse width and make it beneficial to increase the record length.[8] However, for now, we will assume the same 10 ms record length. Next is the number of bits per sample. In ACME II, the averaged traces were saved in 32-bit integer format. For saving the raw, unaveraged records, this is no longer required as they

[7] The term "trace" has a technical meaning in ACME as the average of signal from 25 consecutive pulses (§2.1.7). Thus for clarity we shall use "record" to refer to individual, unaveraged traces of fluorescence from each molecular pulse.

[8] This is also dependent on the specific dynamics of the variation in molecular velocities versus variation in cell exit times. Some experimental tests with longer flight times (such as the lifetime measurement and electrostatic lens tests) did not show a significant increase in pulse width.

are read from the digitizer in 16-bit integer format.[9] Finally, the triggering rate (50 Hz) and number of channels (8 detectors) remain the same as before. With these assumptions the DSR in ACME III is estimated to be

$$\begin{aligned} DSR_{\text{ACMEIII}} &\approx (12.5 \times 10^6 \text{ samples/s}) \times (10 \times 10^{-3} \text{ s/record}) \times (16 \text{ bits/sample}) \\ &\quad \times (50 \text{ records/s/channel}) \times 8 \text{ channels} \\ &= 100 \text{ MB/s}. \end{aligned} \tag{6.3}$$

This DSR is well beyond the average tabletop atomic physics experiment and needs to be maintained throughout 2-3 years of constant data-taking. In order for the ACME III DAQ system to fulfill this goal, several upgrades have been implemented:

1. Upgrading ethernet connections between the DAQ computer and the digitizer chassis.
2. Acquiring a storage solution capable of storing all the data.
3. Rewriting the LabView DAQ code for superior speed and efficiency in acquiring and saving data.
4. Enhancing real-time data analysis capabilities to assist in analyzing the large volume of data.

We shall now elaborate on each of these upgrades.

6.2.3 Connection upgrade

In ACME II, the DAQ digitizer was mounted in a chassis which was connected to the DAQ computer using a PXIe-PCIe MXI Express Gen 1 interface (Fig. 6.2.1a). All of these connections could handle a data throughput of at least 838 MB/s. However, the bottleneck in the system was the connection to the Synology DS1817+ NAS, which used a standard

[9] The digitizer has a maximum resolution of 14 bits.

the standard 1 Gigabit Ethernet (GBe) connection. Tests with standard benchmarking software[10] measured a write speed of 116 MB/s. While in theory this should be just enough to handle the DSR required to save all pulses, there were other inefficiencies and overhead in the ACME II DAQ software that prevented stable long-term operation while saving all pulses (§6.2.5). In summer 2020, tests using the same NAS but upgrading the connection to 10 Gigabit Ethernet (10GbE) found a significant speed-ups in file writing and reading (Table 6.2.1).[11] At this point, the performance was likely no longer limited by the connection but the NAS itself. They also found a significant speed-up for closing a large binary file, another time-intensive DAQ operation. With this configuration, the system was able to save batches of 25 traces at 50 Hz for at least 4 hours, superseding ACME II performance.

Operation	**1 G**	**10 G**
File read	76 MB/s	755 MB/s
File write	116 MB/s	314 MB/s
File closing time	4.6 s	1.5 s

Table 6.2.1: **DAQ system test results with 1G and 10G connections.** Results from test conducted with ACME II NAS (DS1817+) with original (1G) and upgraded (10G) connection. File write and read results are obtained using NAS benchmarking software with a 400 MB file. File closing time result is obtained by closing a 550 MB file with LabView, similar to the ACME DAQ software. Speeds vary by ∼10%.

6.2.4 Data storage

Saving data from all pulses requires a commensurate increase in data storage capacity and performance. In ACME II, the data was stored to a network attached storage (NAS) system,[12] which had 8 hard drives of 6 TB each in a RAID6 configuration, giving 32 TB of

[10] NAS Performance Tester 1.7.
[11] This required installing a separate 10G network adapter in the older NAS.
[12] Synology DS1817+.

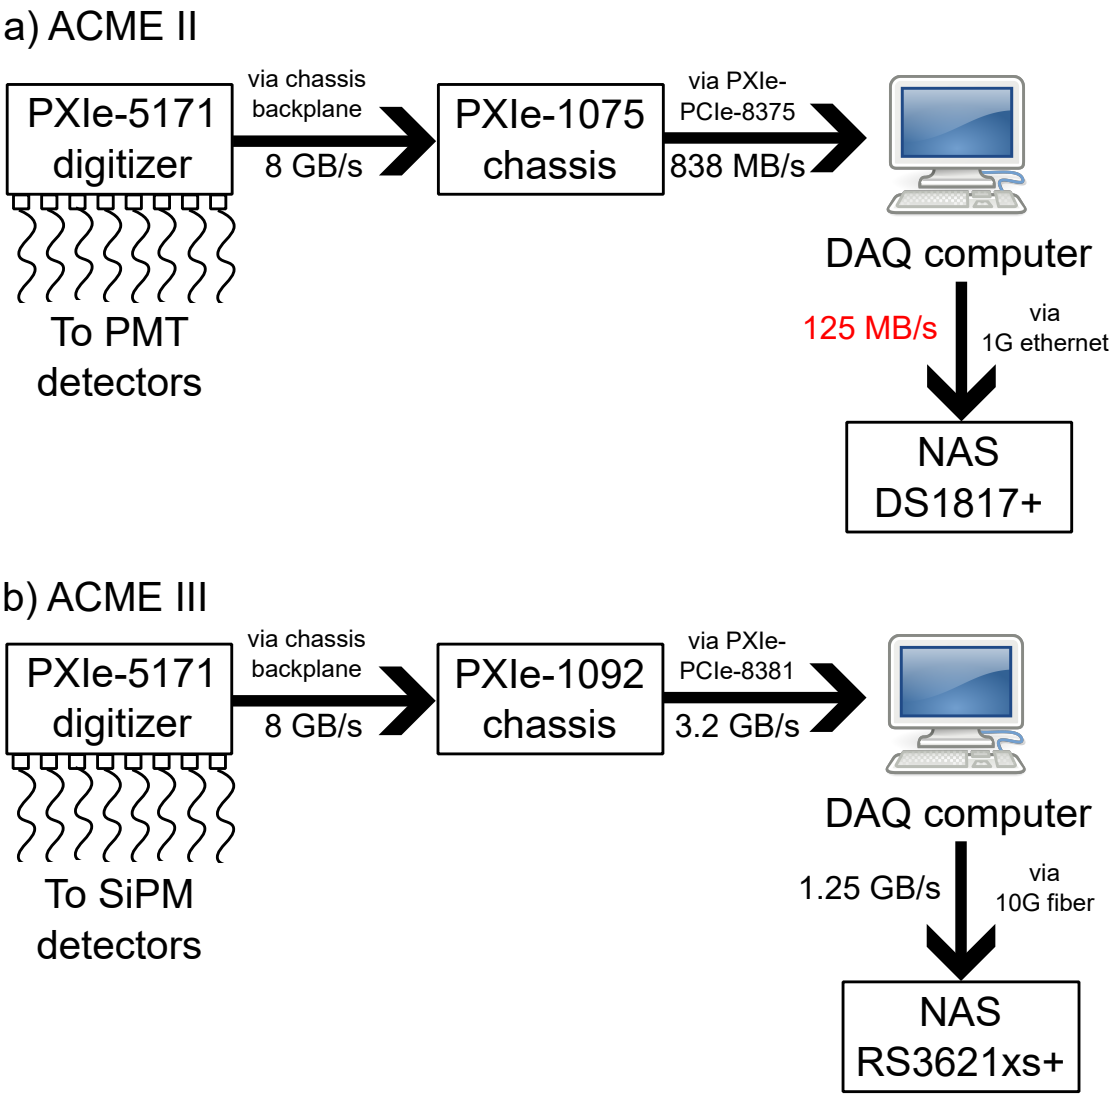

Figure 6.2.1: **DAQ system schematic.** Simplified schematic of the DAQ system in a) ACME II and b) ACME III, showing the maximum connection speeds between components. The ACME II system was limited by the 1 GBe connection with the NAS and was replaced with a 10 GBe connection in ACME III.

space.[13] By the end of ACME II (about 3 years of data-taking), this NAS was filled with about 25 TB of data. Based on the earlier calculated data streaming rate, a reasonable projection is that we will need ∼275 TB of data storage space, including space to store averaged traces alongside raw data from all pulses. After exploration of options, we decided the most cost-effective solution is to store the data in an enterprise-grade NAS[14] with 12 drives of 16 TB each in a RAID6 configuration, giving 140 TB of initial space. In the future, up to two additional expansion NAS units of the same volume can be added and integrated to the same system for a maximum of 420 TB capacity.

Unlike in ACME II where data was written to the NAS over the general ACME internal computer network, the DAQ computer now has a direct dedicated 10 Gigabit fiber connection to the NAS. Tests with benchmarking software conducted after setting up the system at NU measured read and write speeds of 1.2 GB/s and 1.0 GB/s respectively, well over the ACME III requirements.

6.2.5 Rebuilt data acquisition software

The principle challenge of the DAQ software system is to ensure it keeps up with the hardware-timed triggering of the data acquisition, as computers can experience lag and jitter. In ACME, this is accomplished in the main DAQ LabView program by implementing a standard Producer-Consumer architecture [192]. Data is continuously acquired by the digitizer at 50 Hz and fetched at the same rate by the Producer loop. The processing and saving of the data is handled by the Consumer loop which runs in parallel. Data is transferred from the Producer to the Consumer loop using a lossless buffer with a First In/First Out (FIFO) ordering (also known in LabView as a queue). This architecture ensures that any momentary jitter or lag in processing or saving the data will not result in

[13] Note that storing data in several hard drives in a RAID (Redundant Array of Independent Disks) configuration provides not only increased storage volume and redundancy in the case of drive failure, but also higher read/write performance by making it possible to data to multiple drives simultaneously. A RAID6 configuration ensures that 2 drives can fail simultaneously without any data loss.

[14] Synology RS3621xs+.

a missed trigger in the DAQ system.

In the ACME II DAQ program, the Producer loop was able to keep up with the 50 Hz acquisition rate without issue. However, limitations in the Consumer loop program made it insufficient for the goal of saving all pulses. First, it was only able to save a running average of the data. Second, every time it saved data to a binary file, it opened the file, wrote the data, and closed it.[15] The file-closing process for larger files was a timing bottleneck in the old DAQ system.

Thus, for ACME III, while the general Producer-Consumer structure is retained, the DAQ program has been almost entirely rebuilt with significant changes. First, the structure of the Consumer Loop now saves the data to disk in every iteration. Second, instead of opening and closing the binary file every time new data is written into the file, the file is only closed when it is "full", i.e. when all intended records have been written to it. Third, the file structure is modified so that a block of data is split into two files of 32 traces (800 pulses or ∼1.5 GB) each. File-closing only occurs at the midpoint (after trace #32) and endpoints (after trace #64) of the block. The midpoint file-closing is performed simultaneously with a customary ∼4 s magnetic shield degauss, nullifying its effect on the duty cycle. Thus, only the endpoint file-closing has a small effect on the duty cycle (∼3%, assuming a worst-case 1.5 s file-closing time as in Tab. 6.2.1).

The new DAQ program is also synchronized with the Master Run VI running on the main experiment control computer. In ACME II, this was accomplished with a lossless network stream between the DAQ and main control computer running over the ACME intranet (§2.1.8). In ACME III, a similar scheme has been implemented in hardware using a pair of digital I/O modules[16] (one on each computer) connected directly with a cable. This communication bus is used to transmit the trace number, DAQ status (busy/finished saving), and file open/close commands and status. With this setup, queries between the

[15] Closing a file before opening a new one is necessary to release it from memory and prevent data corruption.

[16] NI PCIe-6259.

two computers can occur stably at frequencies up to 500 Hz. This ensures that the duty cycle is not affected by low communication latency.

Finally, the DAQ program has also been optimized in specific ways according to LabView best practices for best performance, such as decreasing the number of nested sub-VIs to improve execution speed (at the expense of more complex block diagrams) and parallelization of data-intensive tasks using separate VIs running on different execution threads.

6.2.6 Data analysis capabilities

After implementing the changes above to save all records, the experiment will generate about ∼2-3 TB of data per day. It will likely be impractical and unwieldy to download and analyze all the data with ACME analysis computers every day. Rather, the system is designed with the expectation that regular data analysis will be performed mainly on averaged and summed traces (as in ACME II). Close analysis of raw, unaveraged records will only need to occur in special cases when anomalies or systematic effects occur. To assist ACME data analysts in identifying these special cases, an online analysis VI has been developed. It consists of several loops running in parallel with the main Producer-Consumer DAQ program (Fig. 6.2.2). Each loop has a different set of functions:

- The Producer loop fetches records at 50 Hz from the digitizer and sends them to the Consumer and outlier detection loops.
- The Consumer loop saves records to the NAS, adds them to a running average of the 25 last records, and plots the running average. If 25 records have been saved, it sends their average to the Secondary Saving and Full Analysis loops.
- The Secondary Saving loop saves the the averaged trace, sums the 8 channels and saves the summed trace.
- The Outlier Detection loop computes various quantities for each of the 25 records:

center of mass of arrival time, total signal, and signal spatial ($\pm\hat{z}$) asymmetry.[17] For each of these quantities, values from the 25 records are compared and outlier records are noted in a text file to be flagged for more granular analysis.

- The Full Analysis loop computes the asymmetry $\mathcal{A}$ for averaged traces. In the future, it may be developed to perform more complex analysis, such as comparing degenerate traces within a block and computing the block EDM value.

- The Communication loop handles a communication line to interface with the Master Run VI running on the main control computer (discussed in the next section).

The high degree of parallelization of data averaging, saving, and analysis capabilities is necessary for suppressing jitter and maintaining high performance with the large data throughput. Data is transmitted between all the different loops above vith lossless queues. The DAQ and live analysis code will likely continue to be enhanced and refined based on specific future needs of ACME III data analysis. Such developmental efforts have already begun, led by Ayami Hiramoto and Maya Watts.

6.2.7 Setup, performance, and testing

In addition to the upgrades described above, the experiment has undergone a general hardware and software upgrade. New DAQ, main control, and SQL database computers have been acquired. All components of the DAQ and experimental control systems have been transferred to Northwestern. Several tests have been conducted with all the DAQ-related upgrades in place to find out the DAQ performance while taking multiple blocks of experiment data using the Master Run VI. The results are shown in Table 6.2.2, including a comparison with ACME II performance metrics, which are obtained from an analysis of Run 225 in the final ACME II dataset. The tests were conducted before all ACME III

[17] These quantities were chosen based on past experience, and will likely be updated and expanded when in accordance with the peculiarities of ACME III data.

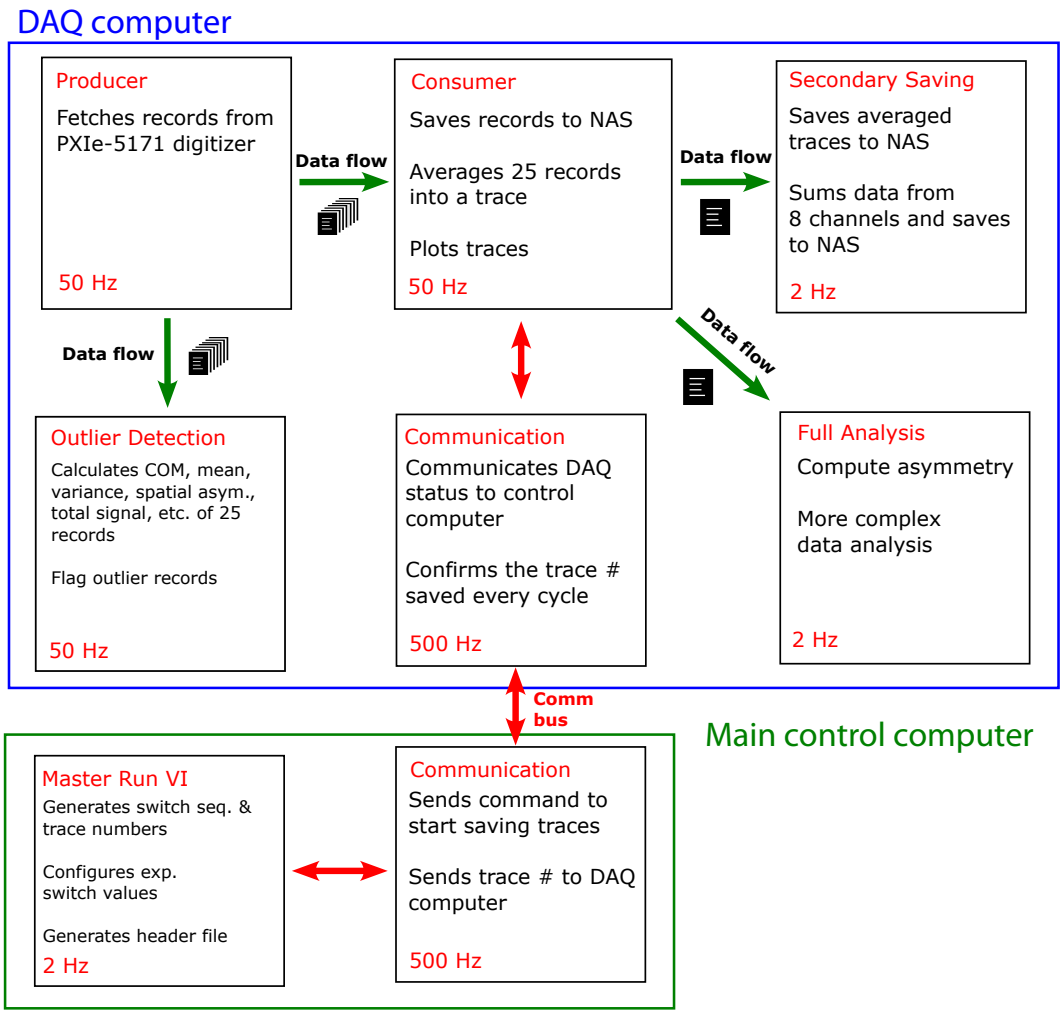

Figure 6.2.2: **DAQ software structure.** Data acquisition, saving, averaging, and analysis are spread out over multiple concurrently running loops to improve performance. Data is transmitted between loops via the use of lossless buffers (queues). Figure also shows relevant loops in the main control computer for integration with the experimental control software.

experiment devices used in switching (such as degaussing coils, electric field plates, $\tilde{\mathcal{N}}$-switch AOMs, etc.) were setup. To simulate the effects of these devices on the experiment timing, minimum settling times similar to ACME II were enforced for each switch, such as a 4.3 s wait time when switching $\tilde{\mathcal{B}}$ to simulate time for degaussing. File read/write speeds are measured separately using NAS benchmarking software and 400 MB file sizes.

Operation	**ACME II**	**ACME III**
File read	76 MB/s	1.2 GB/s
File write	116 MB/s	1.0 GB/s
File closing time (550 MB)	4.6 s	< 0.002 s
Median trace saving time	648 ms	532 ms
Median block time	68 s	57 s

Table 6.2.2: **New DAQ system test results at Northwestern.** Test results obtained with new DAQ system installed at Northwestern including all hardware and software upgrades mentioned in this chapter, and a comparison with the same metrics for ACME II. In the test, minimum wait times for each experimental switch are implemented to simulate their effects on timing.

The trace saving time is the time it takes to save to save data from 25 shots acquired at 50 Hz. The ideal trace saving time is 500 ms. A saving time longer than this indicates DAQ-related delays such as from closing large files after saving, communication between the control and DAQ computers, saving the header file, or others. From the table, we see a $\sim$20% reduction in the median trace saving time compared to ACME II, showing that these software inefficiencies have been reduced. This is despite the fact that in ACME II, only one averaged trace for each detector was saved, whereas in ACME III all 25 raw records are saved. The time to acquire one block of data has been reduced by a similar proportion. The file-closing time is now insignificant (2 ms or less). Finally, tests have also shown that the system can continuously take data for hours without any errors or instability for up to 14 ms record length, which provides flexibility in case a record length longer than 10 ms is desired. All these results mean that the DAQ system is able to fulfill the goals laid out in §6.2.2 with more efficient performance than the ACME II system.

6.3 Conclusion

In conclusion, the data acquisition system has undergone several significant improvements for the next generation. The source of excess noise in the ACME II final run dataset has been identified, characterized, and suppressed to below shot noise, resulting in a 1.7 times gain in EDM statistical sensitivity. Hardware and software upgrades to the data acquisition system enable it to save traces from all molecular pulses, enabling much more powerful and granular data analysis. The new system has been demonstrated to operate stably for extended periods of times with better performance than ACME II, despite the order of magnitude larger data throughput.

7

Magnetic fields

In this chapter, we shall describe the efforts to develop a system to control magnetic fields in the ACME III interaction region. The experiment plans to apply a magnetic field $|\mathcal{B}_z| \approx$ 100-200 μG in spin precession while having a background field $|\mathcal{B}_{nr}| < 10$ μG or better. This is required in order to suppress the excess noise arising from molecular beam velocity fluctuations and systematic effects that arise from magnetic field gradients (§2.2.5, §3.2.2). Several components are needed to accomplish this. First is sufficient magnetic shielding to suppress the earth's magnetic field by 10^5 or better. Second is a set of magnetic coils to apply a uniform field along $\hat{z}$ for precession and along other axes for systematic checks. Finally, the fields need to be adequately and reliably monitored throughout the EDM

measurement. A major challenge of designing these three components is that in ACME III, the dimensions of the interaction region have been significantly enlarged compared to before (§3.1.1).

In this chapter, we discuss all three components while devoting the most space to the magnetic coil design project in which I played a leading role.

7.1 Magnetic shielding

The ACME III interaction region is enclosed by three layers of magnetic shielding with an original design goal of 1 μG DC background field (Fig. 7.1.1). Each layer consists of 1.6 mm thick rectangular sheets of mu-metal[1] mounted on an aluminum frame. The layers are 50 mm apart from each other, with the innermost layer having dimensions 1854 mm $\times$ 1164 mm $\times$ 1164 mm. The size of the sheets were limited by the dimensions of the annealing furnace. Gaps between adjacent sheets are covered with additional overlapping sheets on top (including L-shaped sheets to cover edges). Each layer is divided into a fixed bottom face and an upper U-shaped detachable portion which can be moved into place with a crane.[2] Access holes are added to the shields for the molecular beam, lasers, magnetometers, lightpipes, and screws. Each magnetic shield layer is equipped with a set of 12 degaussing coils for each dimension, resulting in a total of 108 degaussing coils.

The shielding factor S of each layer can be estimated by using the formula [193]

$$S \approx \mu\eta\frac{d}{D}, \tag{7.1}$$

where μ is the relative permeability of the shielding material, η is a geometric factor depending on the shape of the shields, d is the shield thickness, and D is the general shield

[1] Co-Netic AA Perfection Annealed, supplied by Magnetic Shielding Corporation.

[2] This is a cleaner method to remove the shields compared to ACME II, where the shields consisted of half-cylinders that needed to be strapped into place by hand. The excessive bumping of the shields that occurred may have contributed to the observed long-term degradation in shielding (§2.1.11).

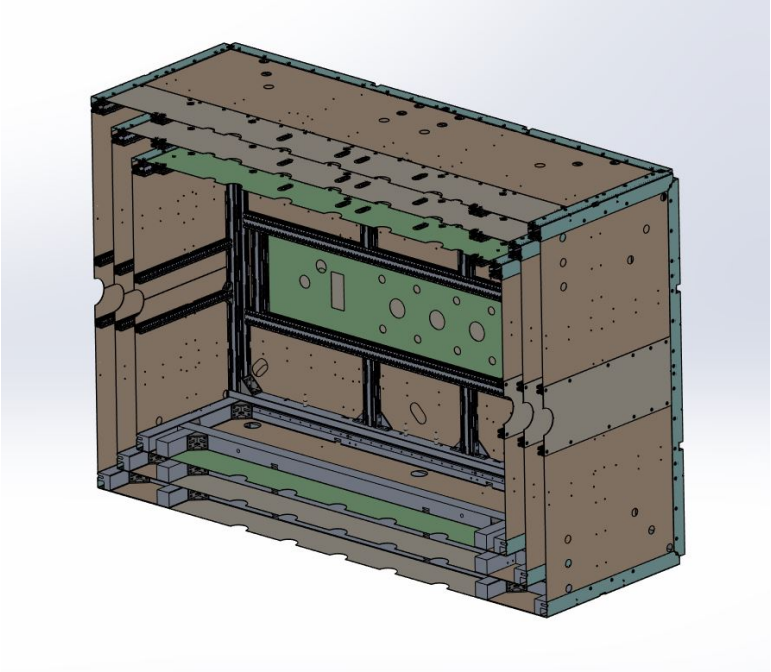

Figure 7.1.1: **ACME III magnetic shields.** Cross-section of a 3D model of the three-layered mu-metal magnetic shield. Degaussing coils are not shown.

size. For ACME III, with $d = 1.6$ mm, $D \approx 1$ m, $\mu \sim 10^5$ (based on manufacturer specifications) and $\eta \sim 1$, we obtain $S \sim 100$ per layer, and a combined shielding factor of around 10^5 to 10^6. More detailed simulations with finite element analysis (including all access holes) conducted by Siyuan Liu obtained shielding factors of 5×10^5 (transverse) and 3×10^5 (longitudinal), or a background field of several μG.[3]

The design and construction of the shields have been led by Siyuan Liu, with important contributions from Cole Meisenhelder, Daniel Lascar, and Xing Fan.[4] The shields have been manufactured, annealed, and assembled on the supporting frame. While characterization of the shielding performance and optimization of degaussing procedures are still ongoing, background fields as low as 5 μG have been achieved.

[3]Based on the lab field being about 1 G.

[4]All ACME members present at Northwestern regularly contributed to the making, installation, and testing of the shields.

7.2 Magnetometry

To monitor the magnetic fields of the experiment, three methods of magnetometry have been developed.

7.2.1 Magnetoresistive sensors

The first method of magnetometry is performed with a set of Twinleaf magnetoresistive sensors[5], which have a sensitivity of 3 mG/$\sqrt{\text{Hz}}$ [194]. Each sensor contains three magnetic sensors (one for each axis) whose resistance changes in the presence of a magnetic field. Currently, a total of 19 sensors are installed: one sensor on each face just inside each shield layer, plus one placed outside to measure the ambient field.[6] The sensors constantly measure and log the magnetic field along all three axes into the SQL database every few seconds. Initial work to calibrate, characterize, and install these sensors was performed by John Mitchell [195], while I made key contributions to the software infrastructure for logging and control.

7.2.2 Optically pumped magnetometers

The second method of magnetometry is performed with ten ultra-sensitive QuSpin optically-pumped magnetometers (OPMs).[7] These sensors will be periodically inserted into one of the many magnetometer pockets mounted on the vacuum chamber to measure the background field in the region directly adjacent to the molecules.

OPMs utilize the technique of zero-field magnetometry [196, 197]. Inside each QuSpin sensor, a set of magnetic coils produces an oscillating magnetic field that compensates the local ambient field.[8] In this environment, a laser is shone through a rubidium vapor cell

[5] Twinleaf vector GMR magnetometers.

[6] More sensors may be added in the future to better monitor and characterize changes in the laboratory ambient field.

[7] QuSpin Zero-Field Magnetometer Gen 3 (QZFM).

[8] There are three orthogonal sets of magnetic coils, allowing all three vector components of the field to be measured.

into a photodetector. The transmission has a peak at zero field, providing an error signal for feedback of the coils which is read using lock-in detection. The value of the ambient field can be determined from the magnitude of the applied compensatory field when the coils are "locked" to the transmission peak. This method of magnetometry is very sensitive ($\lesssim 0.2\ \mu\text{G}/\sqrt{\text{Hz}}$ [198]), but can only be performed in a low background field ($< 500\ \mu$G). Thus, the sensors can only be used inside the shields. In addition, they have been measured to have internal offsets of up to $\sim$40 μG, which have so far limited their absolute accuracy to several μG. More detailed characterizations of these sensors and how to use them are ongoing, led by Maya Watts and Xing Fan.

7.2.3 Q-state magnetometry

The third method of magnetometry is performed by using spin precession in the Q-state of ThO. As previously discussed (§3.1.2), the Q-state has a relatively large magnetic moment which makes it ideal for magnetometry. In addition, this method allows measurement of the magnetic field in exactly the same region as the molecules. This measurement is done by first turning off the STIRAP lasers after the electrostatic lens, such that the molecules remain in the $|Q, JM\Omega = 2, 2, -2\rangle$ after lensing. Once they enter the interaction region, the molecules can be prepared in a superposition of $|JM = 2, \pm 1\rangle$ states using a 746 nm $Q - I$ laser.[9] Spin precession then occurs as it would in the H-state. Only a Zeeman phase is acquired as the Q-state is not sensitive to d_e. In the detection region, the phase is probed by a readout laser on the same $Q - I$ transition with polarization switching, producing 512 nm photons from $I - X$ decay. Thus, to switch between H-state and Q-state precession, only the 703 nm laser needs to be switched to the 746 nm. This can be performed seamlessly by sending in the lasers from two different sides ($\pm\hat{z}$) of the experiment and using remotely controlled shutters. Furthermore, it is also possible to perform a measurement of $\partial\mathcal{B}_z/\partial x$ by

[9] State preparation and readout can also be done with an 1196 nm $Q - C$ laser, analogous to the case of H-state precession.

sending in a pulsed preparation laser counterpropagating to the molecular beam to deplete the population spatially distributed along $\hat{x}$, similar to the microwave measurement method to measure the electric fields used in ACME [113, §3.2.5].

A high magnetic field sensitivity can be achieved, as in the following estimation. We start with

$$\phi^B = \frac{\mu_{Q,J=2} B \tau}{\hbar}, \tag{7.2}$$

from which we obtain the uncertainty expression

$$\begin{aligned} \delta B &= \delta\phi^B \left(\frac{\hbar}{\mu_{Q,J=2}\tau}\right) \approx \frac{1}{\sqrt{N}} \left(\frac{\hbar}{\mu_{Q,J=2}\tau}\right) \\ &= 30 \text{ nG}/\sqrt{\text{Hz}}, \end{aligned} \tag{7.3}$$

where we have assumed that $\delta\phi \approx 1/\sqrt{N}$ [19, §3.1.1], $\dot{N} \sim 10^6$/s (based on the estimated signal in ACME III and the ∼99% loss of molecules due to the scrambling of the M-states between the lens and the entrance region), $\mu_{Q,J=2} = 0.69\ \mu_B$, and $\tau = 5$ ms. The extremely small uncertainty means the measurement will instead likely be limited by $\delta\tau$ arising from the velocity dispersion of the molecular beam, which was not taken into account in the estimation above. Tests have been performed for Q-state magnetometry using the ACME II apparatus at Harvard, led by Xing Wu. Preliminary results indicate that a sub-μG sensitivity can be achieved.[10]

7.3 Magnetic coils: $\mathcal{B}_z$ coils

7.3.1 Design goals and methods

There are two types of magnetic coils in the apparatus: first, the main $\mathcal{B}_z$ coils for normal spin precession and second, the auxiliary coils which produce fields along other axes for systematic checks. The main coil project in ACME III has the goal of designing and building

[10] More details will be published in a forthcoming paper.

a set of coils that is able to apply a uniform $\mathcal{B}_z = 200$ μG within the entire spin precession region with better than 1 μG homogeneity. (This is based on the target of the magnetic shield project.) For ACME III, the magnetic coils must be redesigned to accommodate the much longer precession region and the rectangular shape of the shields. In addition, we have the goal that the magnetic field must be confined within the coils in order to avoid magnetization of the shields. Reduced magnetization may reduce the amount of degaussing required when reversing $\tilde{\mathcal{B}}$.

Thus, we set out to build a magnetic coil with "active shielding", internally confining the magnetic fields away from the shields. Initial work on the $\mathcal{B}_z$ coil project was performed by Bingjie Hao, before I took over the project in spring 2022. The design process followed the technical papers by the Crawford group [199, 200], which outline a method of designing a coil for an arbitrary field configuration. It consists of four steps:

1. Specify the desired field configuration.
2. Determine the required boundary conditions for the magnetic scalar potential φ.
3. Numerically solve for φ using a finite element analysis software.
4. Extract the current density required to produce the field configuration.

We now explain these steps in more detail.

First, we specify the field configuration: $\vec{B} = \mathcal{B}_z\hat{z}$ inside the spin precession region, but $\vec{B} = 0$ outside (Fig. 7.3.1a). Second, we express the resulting boundary conditions. In regions where there is no current ($\vec{J} = 0$), the magnetic scalar potential φ can be used [10, §5.9B]:

$$\vec{H} = -\nabla\varphi, \tag{7.4}$$

where $\vec{H}$ is the magnetic field strength. From the field configuration, we obtain the following boundary conditions for each face (numbered ①-⑤ in the figure). (Note that n is a coordinate along $\hat{n}$, the normal vector for each face.)

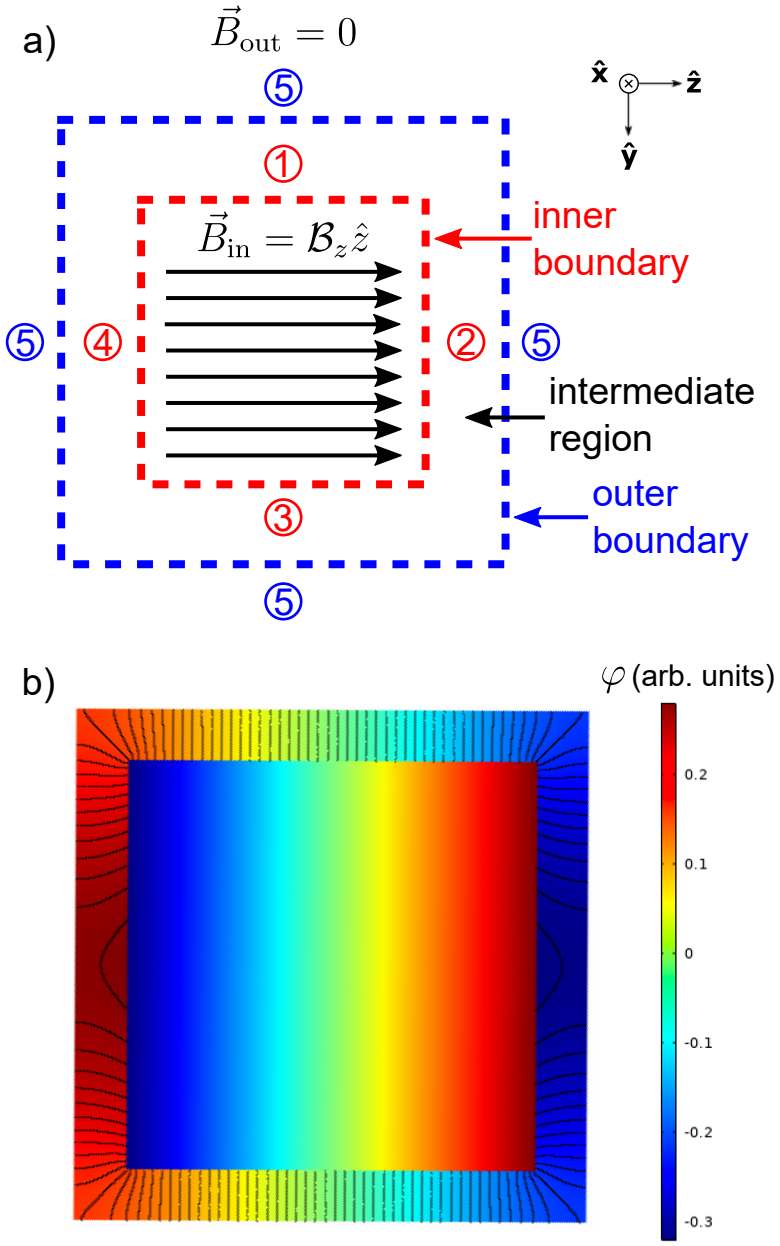

Figure 7.3.1: **Magnetic coil design method.** a) A cross section of the spin precession region with desired magnetic field configuration: a homogeneous $\mathcal{B}_z$ within the entire spin precession region which stretches along $\hat{x}$ (into the page) and zero magnetic field outside. These lead to boundary conditions (numbered ①-⑤). In order to match these boundary conditions, an intermediate region is added. b) Solution of φ from a finite element analysis program based on the boundary conditions, with equipotential lines in black, where wires are placed. Plot obtained from Bingjie Hao.

- Boundary conditions ①, ③, ⑤: $H_n = \partial\varphi/\partial n = 0$.
- ②: $H_n = \mu_0 \mathcal{B}_z$.
- ④: $H_n = -\mu_0 \mathcal{B}_z$.

In addition, we must satisfy Laplace's equation, $\nabla^2\varphi = 0$. As seen in the figure, an intermediate region between the inner and outer boundaries is added. Here we can place currents to satisfy the boundary conditions.

Next, using a finite element analysis program, φ is solved everywhere (Fig. 7.3.1b). Finally, we extract the current density in the intermediate region required to produce φ. It turns out that the currents simply lie along the equipotential lines of φ. This can be explained as follows. From Maxwell's equations, we have

$$\nabla \times \vec{H} = \vec{J}. \tag{7.5}$$

We can integrate this over an interface to obtain the boundary condition [10, p. 18]

$$\vec{n} \times \Delta\vec{H} = \hat{j}, \tag{7.6}$$

which shows that the surface current $\hat{j}$ is orthogonal to $\vec{H}$. By definition, the equipotential curves of φ are also perpendicular to $\vec{H} = -\nabla\varphi$. Since the equipotentials of ϕ and $\hat{j}$ lie on the same 2D surface (Fig. 7.3.1b), $\hat{j}$ has to flow along the equipotentials.

After some iterations and fine-tuning, this process resulted in the outer and inner coil designs shown in figure 7.3.2. For the outer coil, 198 single wire loops run along the direction of equally spaced equipotential lines (Fig. 7.3.2a). The current directions for the top and bottom halves of the coil are opposed to each other. The inner coil consists of 16 equally spaced wire loops (Fig. 7.3.2b). The direction of inner coil currents are aligned with the inner currents of the outer coil but anti-aligned with the outer currents of the

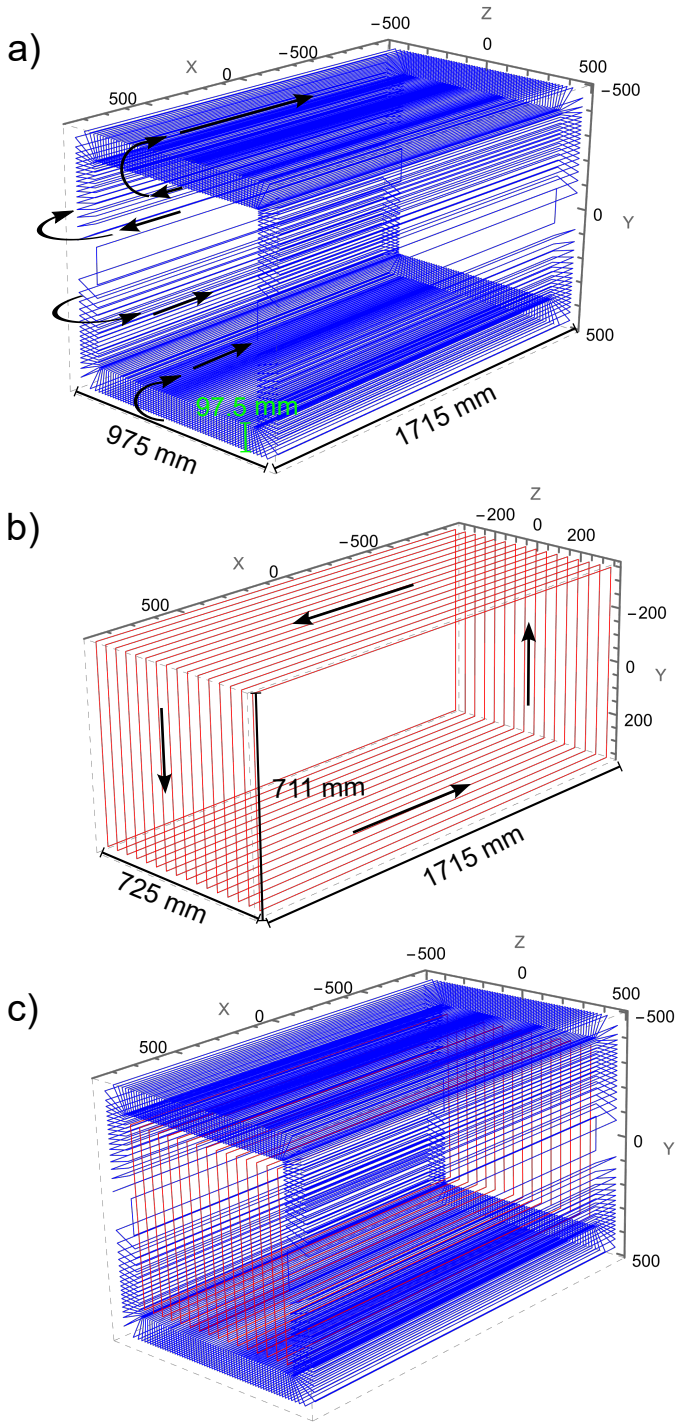

Figure 7.3.2: **Magnetic coil design** obtained after the design process described in §7.3.1. a) Outer coil, b) inner coil, c) combined coils. Black arrows denote the direction of the currents.

outer coil. This reinforces the field inside the coils but suppresses it outside, allowing the aforementioned boundary conditions to be satisfied when the two coils are superposed together (Fig. 7.3.2c).

7.3.2 Accounting for holes

The full magnetic coil design also includes holes for the molecular beam, vacuum window access, and lightpipes. These holes have dimensions of several inches and require some rerouting of the wires (Fig. 7.3.3b). Let us estimate the effect of such reroutings. A rerouting is equivalent to inserting a semicircular current loop running in the opposite direction (Fig. 7.3.3c). As the loops are relatively far away from the spin precession region in the center, we can model them as distant circular loops of current, which has magnetic field magnitude of

$$B_{\text{loop}} = \frac{\mu_0 I R^2}{r^3}, \tag{7.7}$$

where R is the size of the loop, r is the distance from it, and I is the current. Here, from the typical size of the hole we have $R \approx 10$ mm and the closest distance from the hole to the precession region is $r = 355$ mm, from which we obtain $B_{\text{loop}} \approx 0.01$ μG/mA. Even considering that there are about 40 such holes and 2-4 reroutings are required for each hole (Fig. 7.3.3), this would still be around the 1 μG target field homogeneity goal.

7.3.3 Field homogeneity

The field homogeneity in the spin precession region was then calculated from the Biot-Savart law using the Radia software package.[11] A plot of the field along $\hat{x}$ at $y = z = 0$ is shown in figure 7.3.4. It is clearly observed that the inner and outer coils create a combined field with better homogeneity. The vertical dashed lines denote the boundary of the spin

[11] Developed by the ID group of the European Synchrotron Radiation Facility.

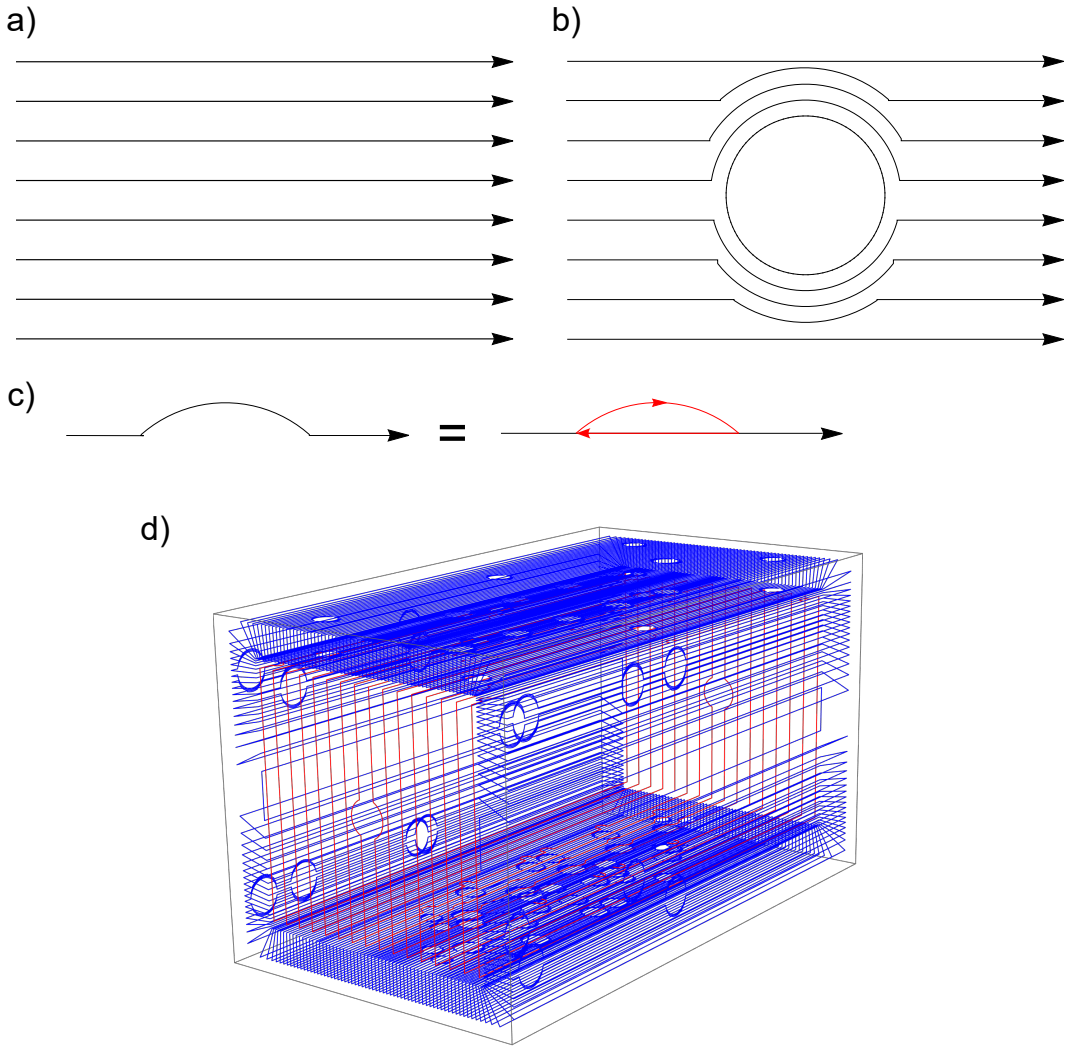

Figure 7.3.3: **Wire rerouting for holes in the magnetic coils.** a) A set of parallel magnetic coil wires. b) Wires after rerouting around an access hole. c) Rerouting is equivalent to adding a semicircle current loop whose straight portion flows in the opposite direction (red curve). d) The outer and inner coils including all wire reroutings used in the Radia simulation.

precession region.[12] Assuming $\mathcal{B}_z = 200$ μG and sampling the field within this region (a rectangular box with dimensions 530 mm $\times$ 15 mm $\times$ 15 mm), the homogeneity or spread of $\mathcal{B}_z$ is calculated to be 0.08 μG (0.04%) without accounting for hole reroutings.[13] With all hole reroutings, the homogeneity is 0.2 μG (0.09%), which is still well within the 1 μG homogeneity target. (This value is also consistent with the estimation of the effect of the holes in the previous section.) It is also possible to further improve the homogeneity in the configuration with holes by fine-tuning the ratio between the outer and inner coil currents. By setting $I_{outer}/I_{inner} = 1.01$, the homogeneity can be improved to 0.1 μG.

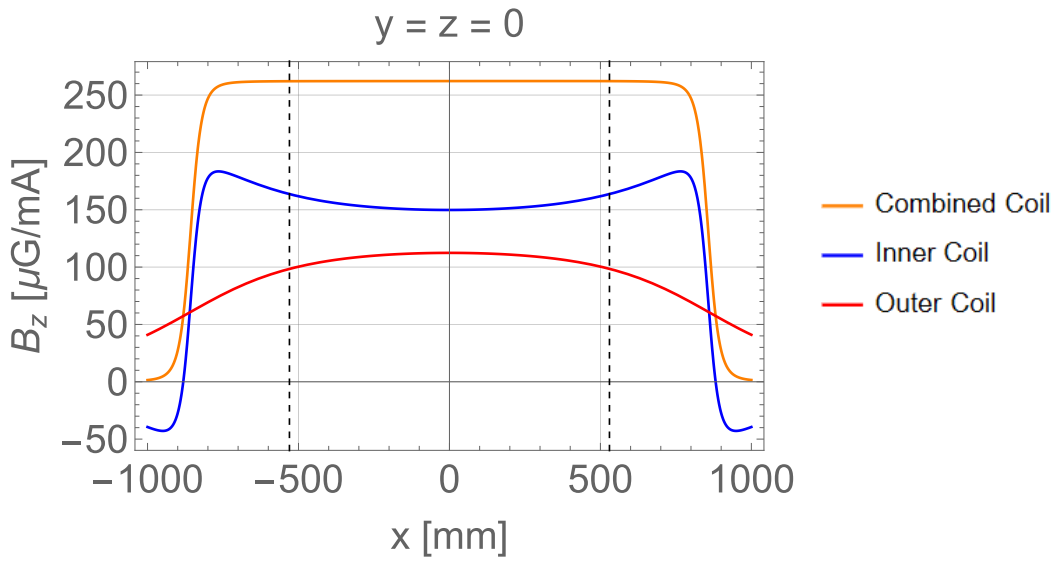

Figure 7.3.4: **Magnetic field homogeneity along $\hat{x}$.** Plot of $\mathcal{B}_z$ produced by outer coil (red), inner coil (blue), and combined coils (orange) from calculation with Radia. Vertical dashed lines denote the boundaries of spin precession region.

From the plot we also see that only a small current ($<$ 1 mA) is needed to produce the required field magnitude. Finally, some simulations were conducted to estimate the impact of small (several mm) misalignments in the wire placement, which is expected from machining tolerances and the width of the wire grooves. The effects of these misalignments

[12] The region is slightly off-center along $\hat{x}$ due to the design of the interaction region which has a shorter upstream portion (§3.1.1).

[13] There are no significant other magnetic field components.

are negligible.

7.3.4 Field confinement

The calculations also predict that the magnetic field will be well-confined within the coils. Table 7.3.1 shows the maximum and mean magnitude of the magnetic field on each face of the innermost magnetic shield when applying $\mathcal{B}_z = 200$ μG. Plots of the resulting fields are shown in figure 7.3.5. The maximum field magnitude of 17 μG, or ~8% of the applied field. This is over an order of magnitude better performance compared to the ACME II coils.[14] The maximum value occurs at the $\hat{y}$-center of the front and back faces (Fig. 7.3.5c), which is the pivot point for the reversal of current direction between the top and bottom halves of the outer coil (Fig. 7.3.2a). At other shield faces, the largest magnetic fields are the result of hole reroutings.

Shield face	**Max field**	**Mean field**
Top/Bottom	5	2.3
Sides	11	4.0
Front/Back	17	2.5
All	**17**	**2.7**

Table 7.3.1: **Maximum and average magnetic field produced by coils at each shield face.** Units are in μG. Results obtained from Radia simulation of coils with all holes assuming an applied field of $\mathcal{B}_z = 200$ μG.

The mechanics of the field cancellation are illustrated in figure 7.3.6. Inside the coils, the currents of the inner and outer coils flow along $+\hat{x}$, resulting in a combined homogeneous field along $+\hat{z}$. Outside the coils, the outer coil current direction is reversed, resulting in cancellation between the two coils. (See also figure 7.3.2.)

[14] This estimate was obtained by calculating the maximum field produced by the ACME II coils at the ACME II cylindrical shields.

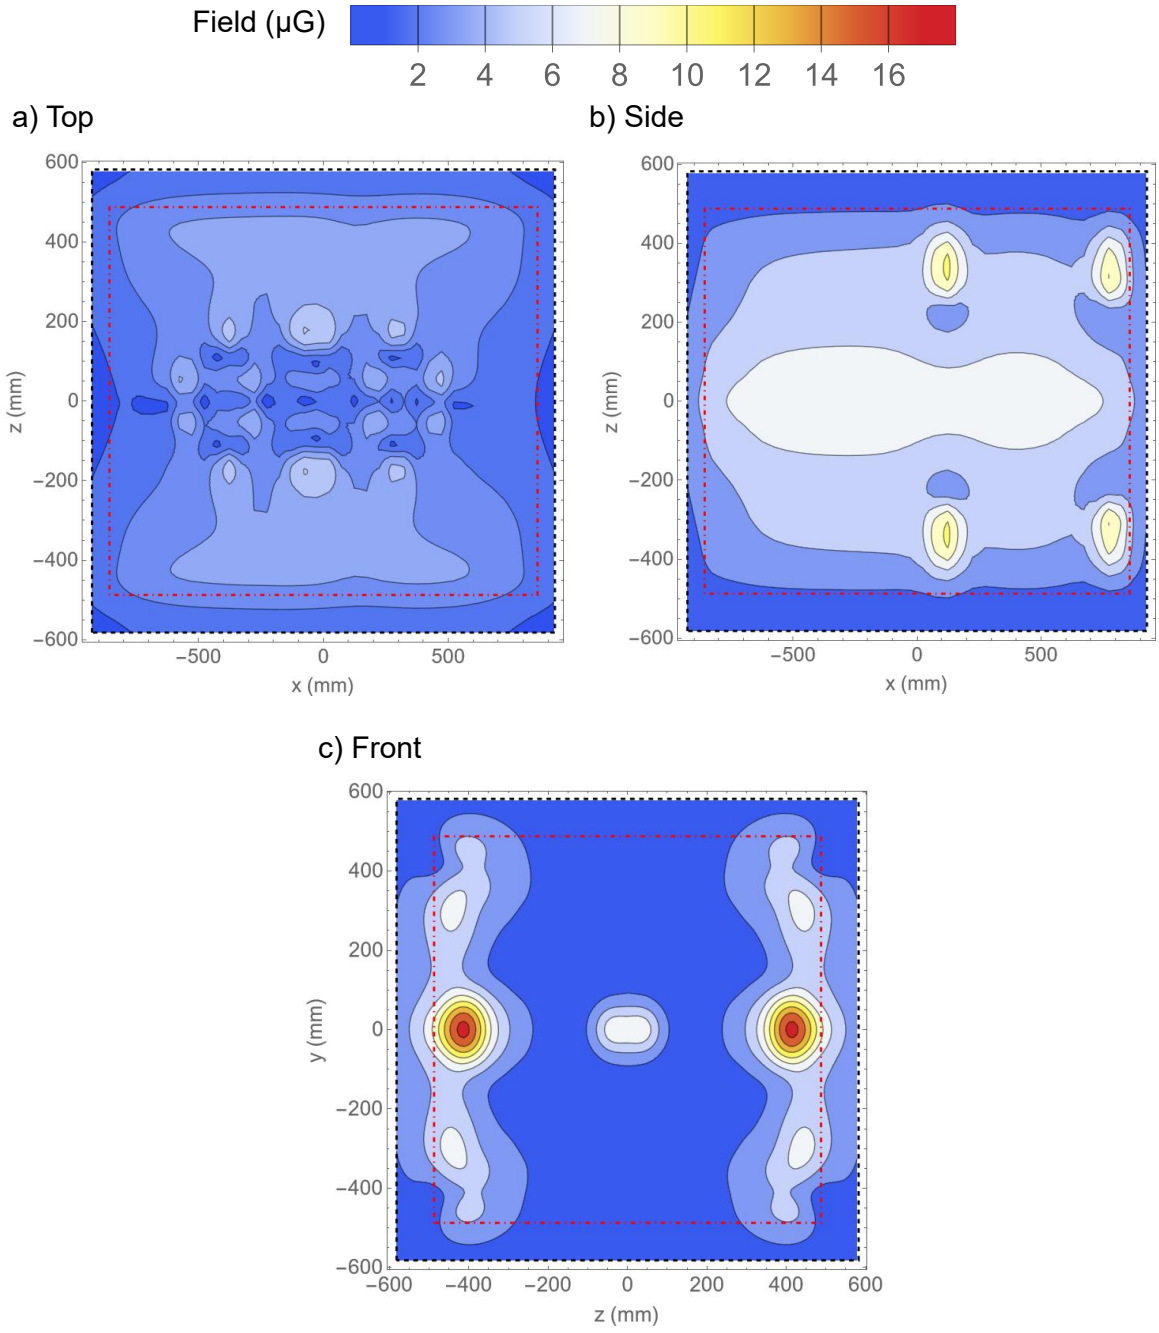

Figure 7.3.5: **Magnetic field confinement with active shielding coil.** Calculated plots of $|\mathcal{B}|$ at innermost magnetic shield when applying $\mathcal{B}_z = 200$ μG in the spin precession region: a) top face, b) side face, c) front face. Black dashed lines indicate the boundaries of the shield while red dotted-dashed lines indicate the boundaries of the outer coil.

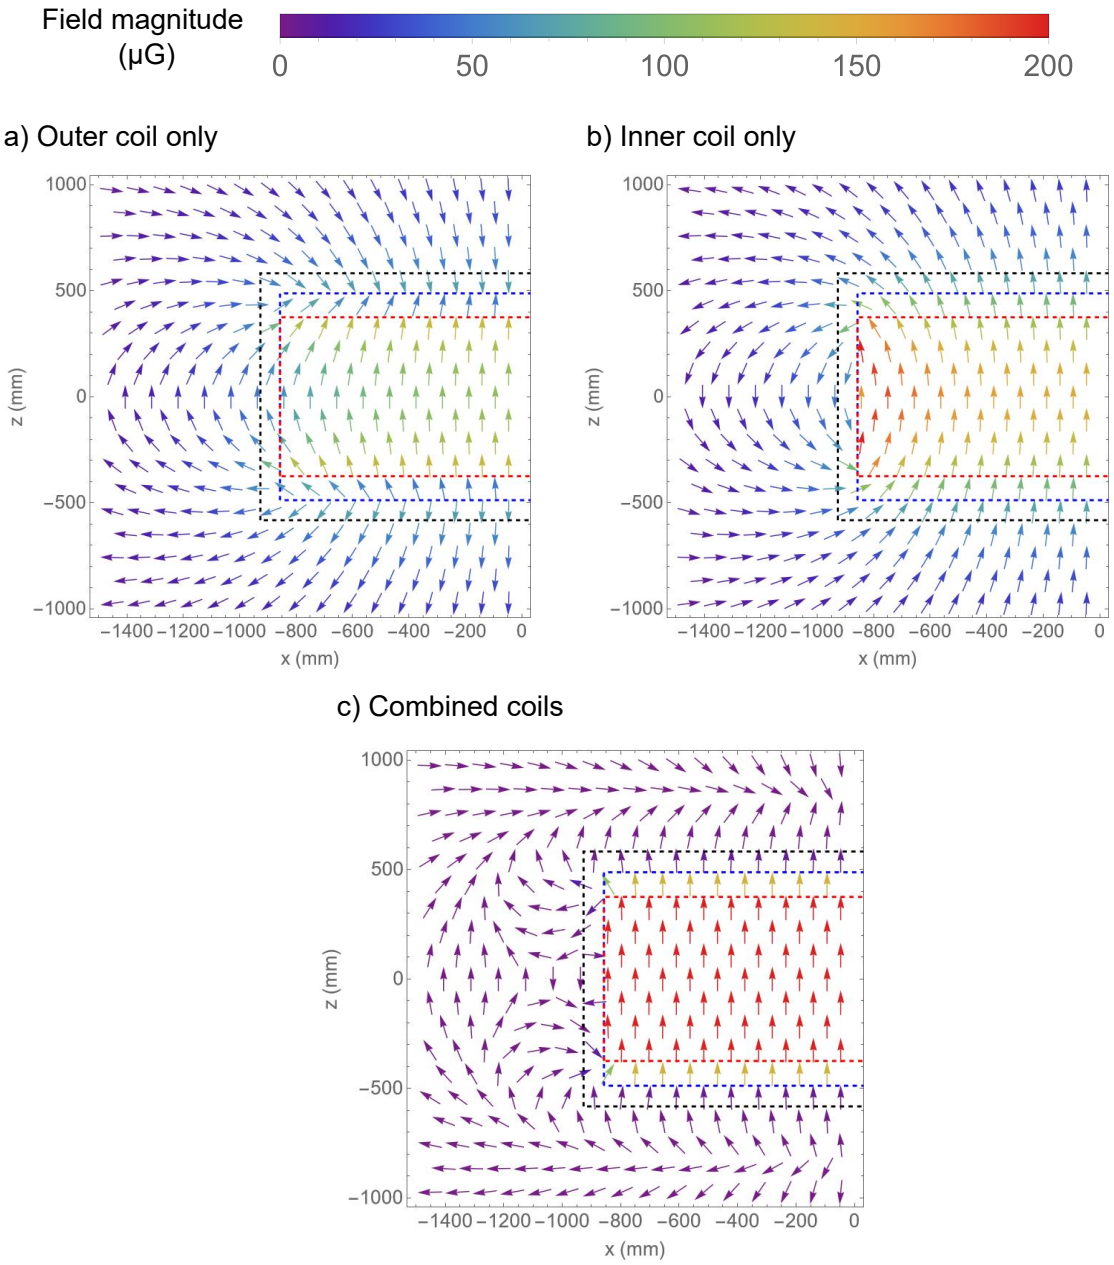

Figure 7.3.6: **Inner and outer coil active shielding.** Calculated vector plots of $\{\mathcal{B}_x, \mathcal{B}_z\}$ at $y = 0$ when applying $\mathcal{B}_z = 200$ μG with a) outer coil only, b) inner coil only, c) combined coils. Dashed lines indicate the boundaries of the innermost magnetic shield (black), outer coil (blue), and inner coil (red).

7.3.5 Construction

At the time of writing, coil construction is ongoing, led by Xing Fan and Maya Watts. Insulated copper wires (24 AWG, or 0.02" in diameter) will be implanted into HDPE sheets 1/4" thick milled with 1/16" grooves and held in place by friction and strategically placed tabs. The sheets are mounted on aluminum structures for the inner and outer coils (Fig. 7.3.7). (Note how the coil structure is divided into a bottom face and an upper-U portion to allow for straightforward vertical removal with a crane, similar to the magnetic shield structure.) Additional HDPE layers for the auxilary coils (described in the next section) are mounted on the inner coil structure.

With the specifications above, the total resistance of all the wires in the outer and inner coils is $\sim 70\ \Omega$. The current required for normal operation is ~ 0.8 mA.[15] In order to maintain 1 μG homogeneity at $\mathcal{B}_z = 200\ \mu$G, a stability of 4 μA is required. These requirements should be reachable with commercial low-noise current supplies.

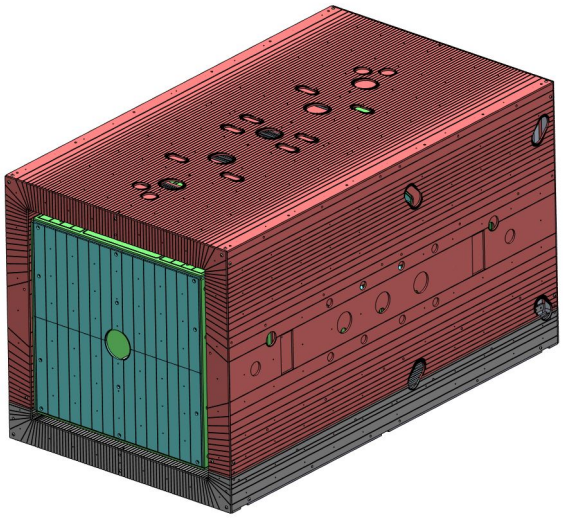

Figure 7.3.7: Computer-generated drawing of the outer and inner coil construction.

[15] Up to 25 mA will be required for operating the coils in $\partial B_z/\partial z$ mode during systematic checks as explained in the next section.

7.4 Auxiliary magnetic coils

The goal of the design of the auxiliary coils is to allow application of constant and gradient fields other than $\mathcal{B}_z$ that are much larger than their typical values while requiring a tractable amount of current (< 1 A). For constant fields, the target is to apply up to 2 mG, which is 10 times larger than the typical applied $|\mathcal{B}_z| = 200$ μG, and even larger compared to the expected ambient $\mathcal{B}_x$, $\mathcal{B}_y$ values. For field gradients, the expected ambient field gradients are 1-10 μG (§7.1). Thus, a reasonable target is to apply up to 200 μG/cm.

Unlike the main $\mathcal{B}_z$ coils, these auxiliary coils will mainly be used only for systematic checks, and thus the requirements for field homogeneity are less stringent. Field confinement is also no longer a requirement. Overall, the design of the auxiliary coils followed the basic design in ACME II (§2.1.11), with special modifications to adapt to the much larger interaction region.

7.4.1 $\partial\mathcal{B}_z/\partial z$ coils

The first gradient to apply is $\partial\mathcal{B}_z/\partial z$. This can be done using the same coils designed for $\mathcal{B}_z$, but dividing the $\pm\hat{z}$ coils and reversing the direction of the current for one group. The resulting $\mathcal{B}_z$ field is plotted in figure 7.4.1. The predicted current efficiency is 8 μG (cm mA)$^{-1}$, which means a maximum of 25 mA of current will be required to produce the desired 200 μG/cm gradients.

7.4.2 $\partial\mathcal{B}_z/\partial x$ coils

In order to apply $\partial\mathcal{B}_z/\partial x$, four additional coil sets parallel to the xz-plane are planned (Fig. 7.4.2a). Each set consists of a main coil (red, 8 turns of wire) and a shim coil (blue, 16 turns) which together produce a uniform $\mathcal{B}_z$ gradient along $\hat{x}$ (Fig. 7.4.2b).[16] All four coil

[16] Some wire rerouting around the light pipe holes is required (depicted in green in Fig. 7.4.2a). This has a negligible effect on the field produced.

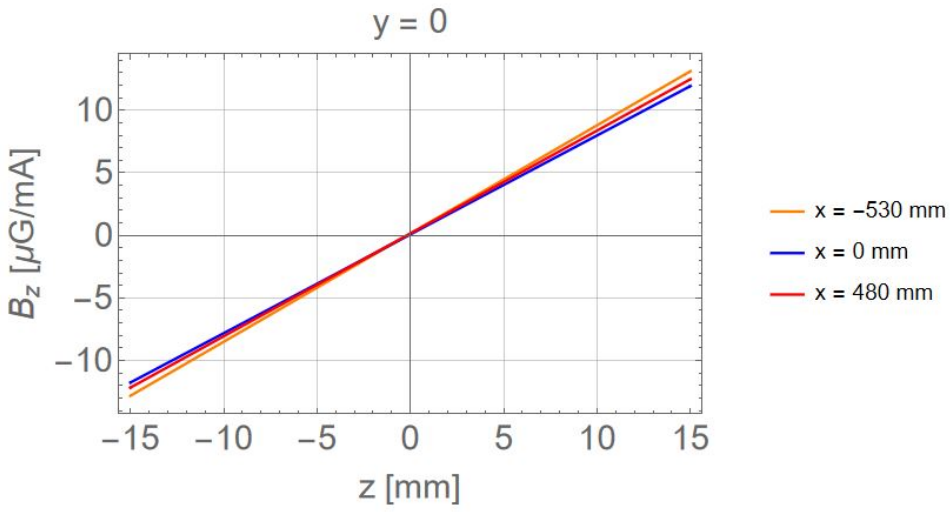

Figure 7.4.1: Calculated plot of $\mathcal{B}_z$ at $y = 0$ when operating the $\mathcal{B}_z$ coils in gradient mode. The field is plotted for three different $\hat{x}$-positions: center (blue) and the boundaries of the spin precession (orange and red).

sets can be linked up in series to the same power supply. The predicted current efficiency is 3.2 μG (cm mA)$^{-1}$, implying a maximum current requirement of 62 mA.

7.4.3 $\mathcal{B}_x$ COILS

$\mathcal{B}_x$ is produced with four pairs of square Helmholtz coils (Fig. 7.4.3a) located at $x = \pm\{125,\ 370,\ 695,\ 705\}$ mm.[17] Each coil contains a single turn of wire divided into two detachable segments (bottom face and upper-U, following the design of the main coils and magnetic shields). A larger number of turns for each coil was considered but rejected to minimize the number of connections between the segments. The currents of the four coil pairs are fine-tuned to achieve a homogeneous $\mathcal{B}_x$ field (Fig. 7.4.3b).[18] The predicted current efficiency is ≈ 43 μG/mA per coil pair, resulting in a total current requirement of 171 mA (divided among the four pairs) to produce $\mathcal{B}_x = 2$ mG.

To operate the coils in $\partial\mathcal{B}_x/\partial x$ mode, current flow is reversed in one half of the coils

[17] Here $x = 0$ is located at the center of the spin precession region.
[18] The plot in the figure assumes current ratios $\{0.89, 0.95, 0.87, 1\}$.

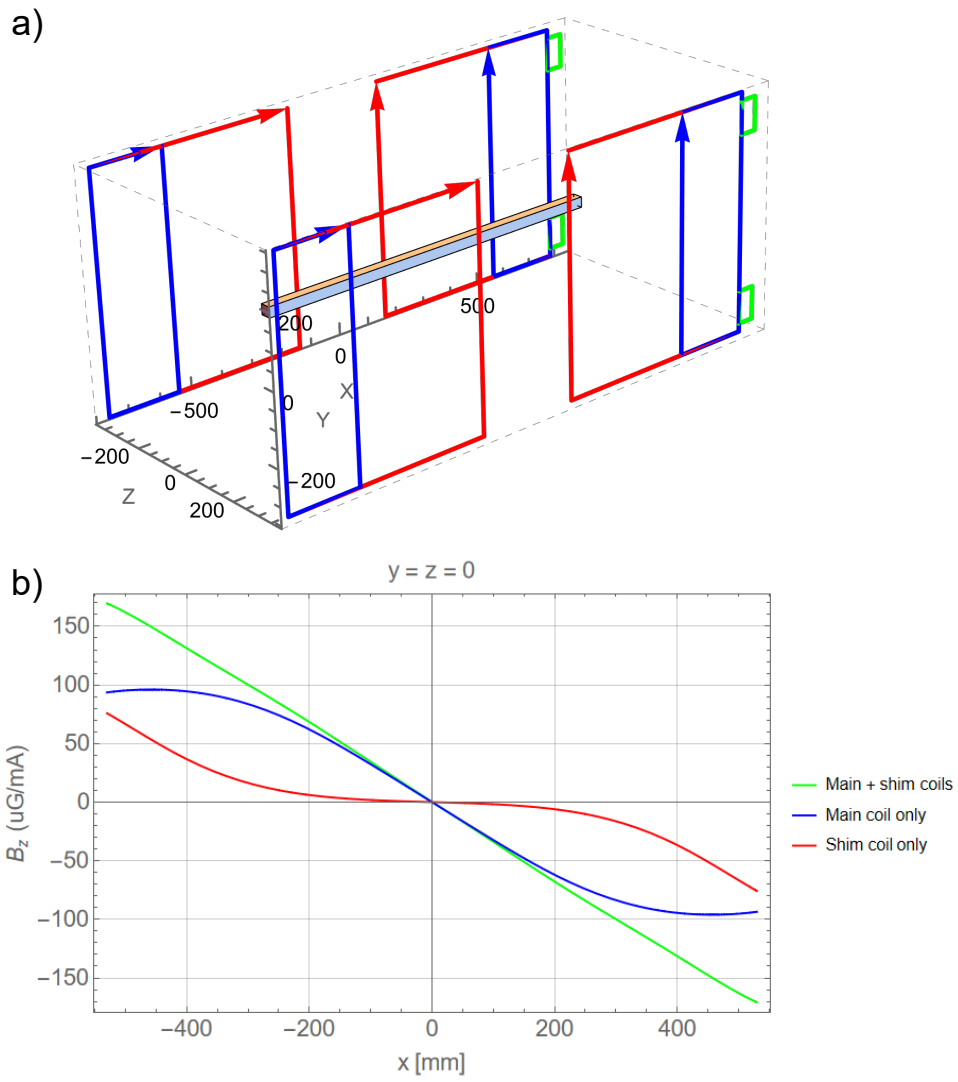

Figure 7.4.2: $\partial\mathcal{B}_z/\partial x$ **coil design and performance.** a) $\partial\mathcal{B}_z/\partial x$ coil schematic, showing four sets of main (red) and shim coils (blue). The green lines denote wire reroutings around access holes. The box-shaped spin precession region is also shown. b) Calculated $\mathcal{B}_z$ field at $\{y, z\} = 0$ along $\hat{x}$ produced by main $\partial\mathcal{B}_z/\partial x$ coil (blue), shim coil (red), and both coils (green).

and current magnitudes are changed to maintain a uniform gradient throughout the entire spin precession region (Fig. 7.4.3c). This requires larger currents for the outermost coils at $x = \pm\{695,\ 705\}$ mm.[19] A possible configuration with a single current source is shown in figure 7.4.3d, which would require a total of 593 mA to produce $\partial\mathcal{B}_x/\partial x = 200$ μG/cm. The current division can be implemented with a set of fixed resistors with low temperature coefficients for better stability.

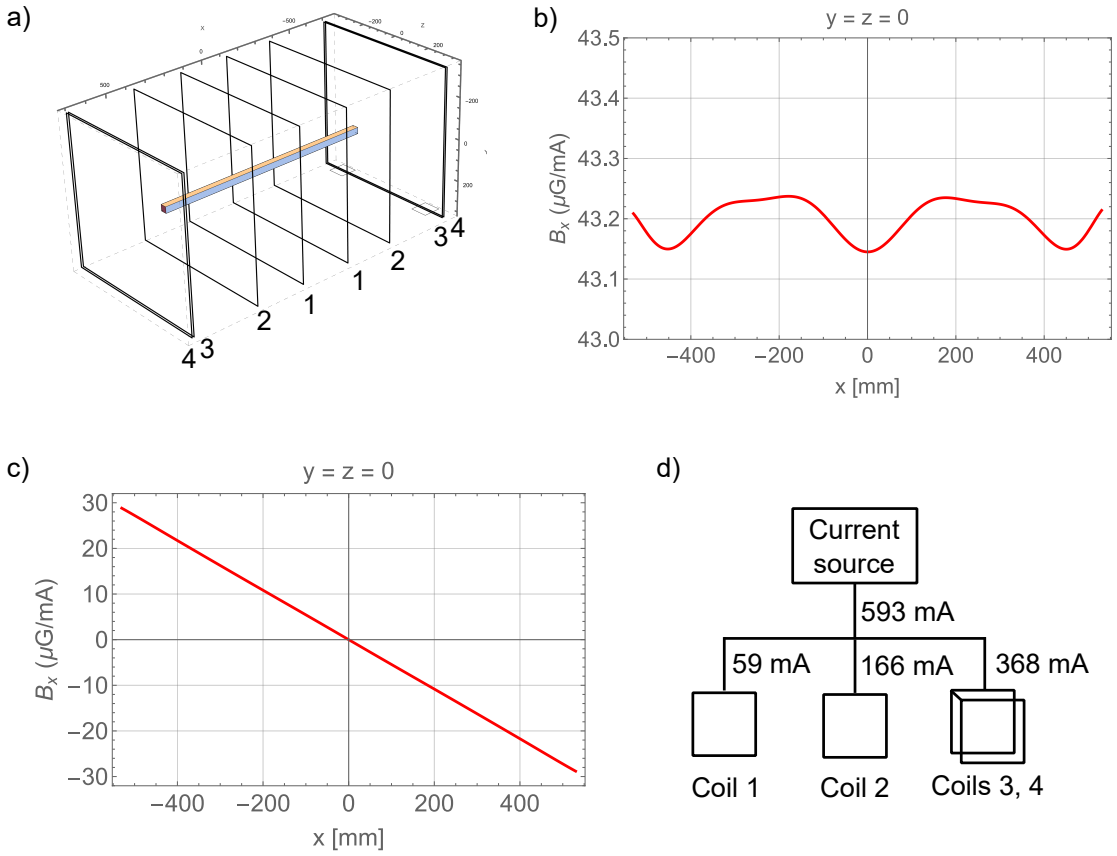

Figure 7.4.3: a) $\mathcal{B}_x$ coils, featuring four pairs of Helmholtz coils (numbered 1-4). The box-shaped spin precession region is also shown. b) Calculated plot of $\mathcal{B}_x$ field at $\{y, z\} = 0$ along $\hat{x}$. c) Calculated plot of $\mathcal{B}_x$ when operating in gradient ($\partial\mathcal{B}_x/\partial x$) mode. d) Possible configuration to produce $\partial\mathcal{B}_x/\partial x = 200$ μG/cm with a single current source. Such division of currents can be implemented with a set of resistors (not shown).

[19] The plot in the figure assumes current ratios $\{0.16, 0.45, 1, 1\}$, which is also reflected in figure d).

7.4.4 $\mathcal{B}_y$ COILS

The $\mathcal{B}_y$ coils are implemented with a set of eight lattice coil octants (green in Fig. 7.4.4a), a pair of center coils spanning the whole interaction region (red, 5 turns of wire), and two pairs of side coils (blue, 10 turns of wire). Each coil octant consists of a 20×3 lattice of square coil windings. The number of turns at each lattice site has been optimized for best field homogeneity. For construction, they can be broken down into a superposition of two simpler sets of windings (Fig. 7.4.4b), inspired by the method used in Ref. [201]. Four modes of operation are possible by switching the current polarity of the coil octants: $\mathcal{B}_x$, $\partial\mathcal{B}_y/\partial x$, $\partial\mathcal{B}_y/\partial y$, and $\partial\mathcal{B}_y/\partial z$, with different coil subsets used for each mode (Tab. 7.4.1). This result in field plots shown in figure 7.4.5. All the coils can be connected in series with a single current source and a collection of relays to switch polarities depending on the chosen operating mode. The maximum required current is 94 mA (used for $\partial\mathcal{B}_y/\partial x$ mode).

7.4.5 SUMMARY AND CURRENT STATUS

We have completed our survey of the auxiliary magnetic coil design. Due to Maxwell's equations (§2.1.11), the four sets of auxiliary coils described above are able to apply constant fields and gradients in the spin precession region along all three axes. The maximum current required for a single coil set is ≈ 600 mA (for applying $\partial\mathcal{B}_x/\partial x$). Figure 7.4.6 shows a schematic of all auxiliary coils that have been discussed. All coil sets are mounted on the same structure as the $\mathcal{B}_z$ inner coil. The designs have been carefully chosen so as to maximize homogeneity while minimizing conflicts with access holes on the apparatus which would require rerouting. At the time of writing, construction of the coils is ongoing, led by Xing Fan and Maya Watts. Some design work remains to be done for the electronics and software for controlling all the currents.

Finally, it is useful to note that the auxiliary coil system can in principle also be used to

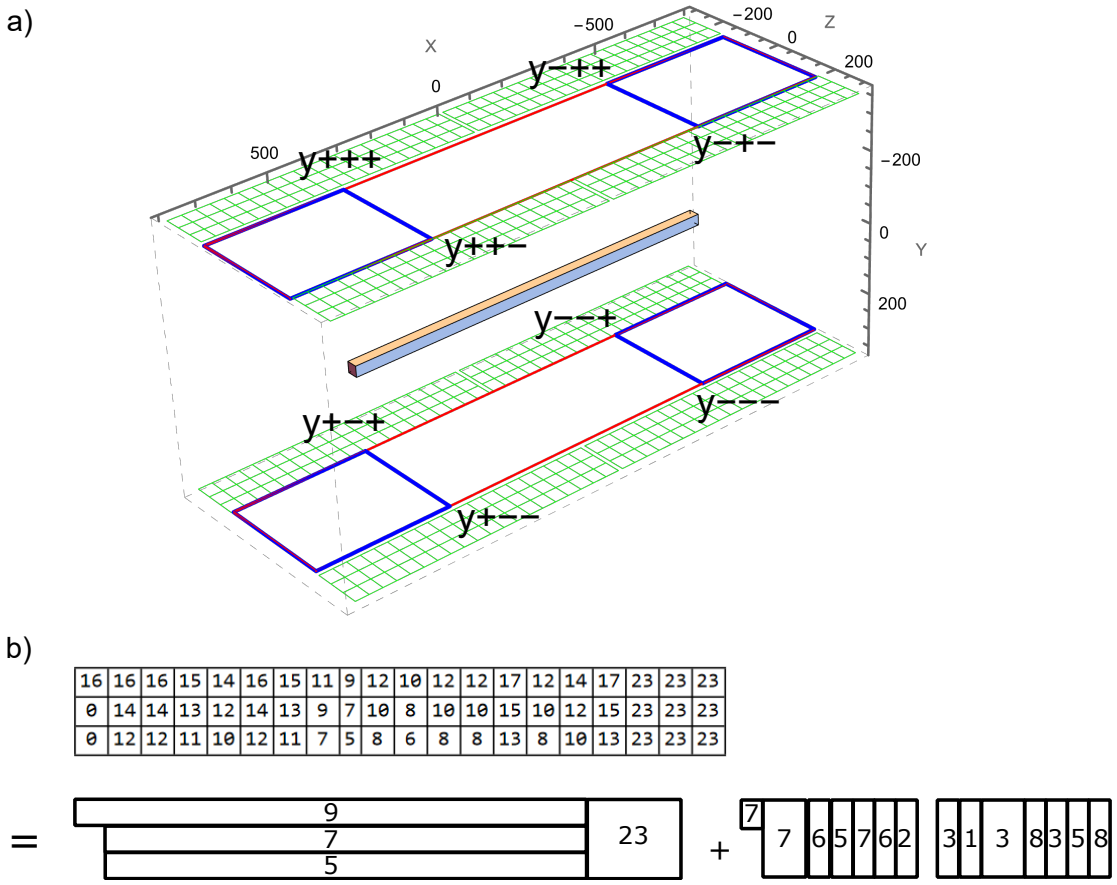

Figure 7.4.4: a) Schematic of $\mathcal{B}_y$ coils, featuring eight lattice coil octants (green, variable number of turns), a pair of center coils (red, 5 turns), and two pairs of side coils (blue, 10 turns). The box-shaped spin precession region is also shown. b) Number of windings at each octant lattice site, which is equivalent to a superposition of two layers of simpler coil structures.

Mode	**Lattice coils**	**Center coils**	**Side coils**	**Max. current (mA)**
Constant $\mathcal{B}_y$	On	On	Off	25
$\partial\mathcal{B}_y/\partial x$	On	Off	On	94
$\partial\mathcal{B}_y/\partial y$	On	On	Off	85
$\partial\mathcal{B}_y/\partial z$	On	Off	Off	67

Table 7.4.1: Coil sets used in different modes of the $\mathcal{B}_y$ coils, with the current required to produce a maximum $\mathcal{B}_y = 2$ mG or gradient field of 200 μG/cm.

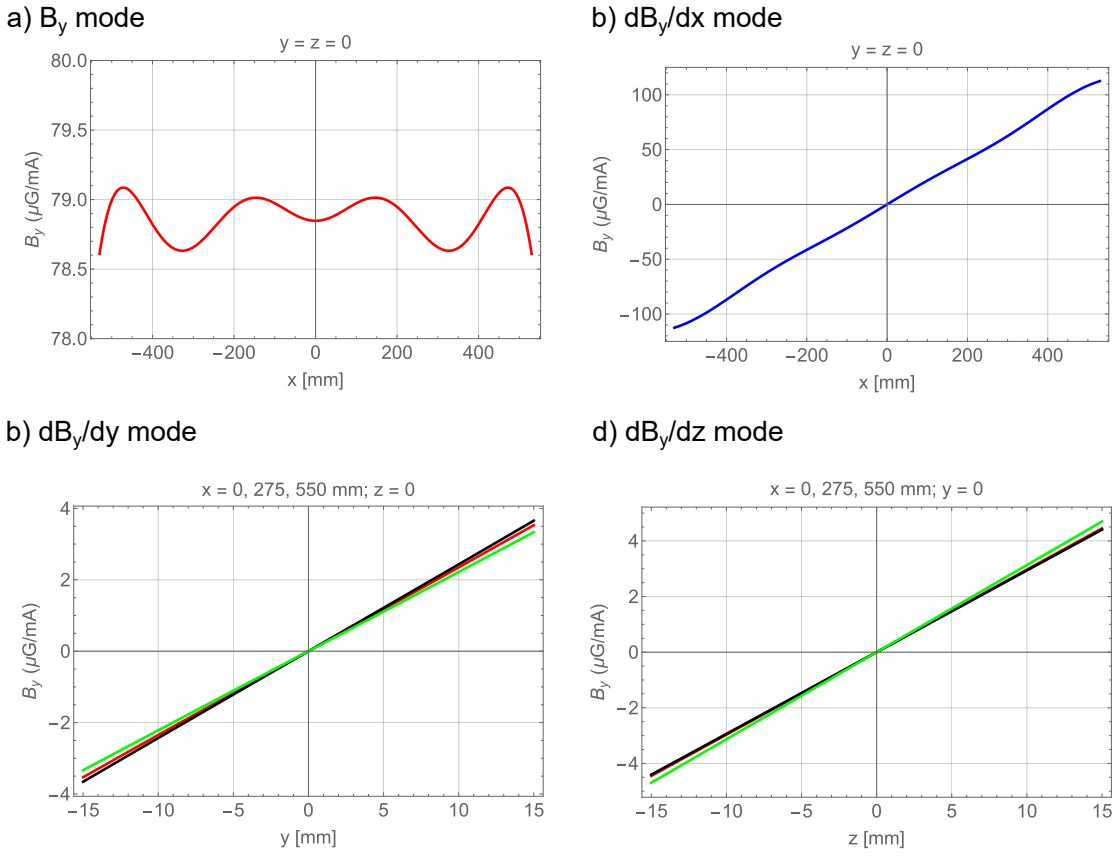

Figure 7.4.5: Calculated plots of $\mathcal{B}_y$ when operating the $\mathcal{B}_y$ coils in a) constant $\mathcal{B}_y$ mode, b) $\partial\mathcal{B}_y/\partial x$ mode, c) $\partial\mathcal{B}_y/\partial y$ mode, and d) $\partial\mathcal{B}_y/\partial z$ mode.

perform active cancellation of undesired magnetic field gradients in the interaction region that arise from sources such as holes in the magnetic shields.

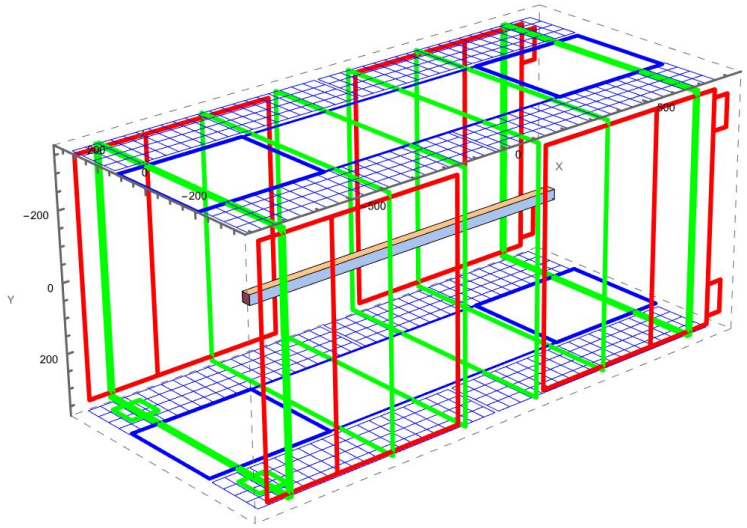

Figure 7.4.6: Schematic of all auxiliary coils: $\partial\mathcal{B}_z/\partial x$ coils (red), $\mathcal{B}_x$ coils (green), and $\mathcal{B}_y$ coils (blue). $\partial\mathcal{B}_z/\partial z$ is applied using the main magnetic coils (Fig. 7.3.2).

7.5 Conclusion

In conclusion, the next generation apparatus has been designed for improved magnetic field control and monitoring, which is crucial to suppress known magnetic field-related systematic uncertainties and excess noise. The magnetic shielding system will suppress the background fields to at most several μG. Three different methods of magnetometry have been prepared to monitor the field with greater precision, accuracy, and reliability than ever before. While assembly and characterization of the magnetic shielding system is ongoing, background fields as low as 5 μG have so far been achieved. Finally, the next generation magnetic coil system has been designed. The $\mathcal{B}_z$ coil system is designed to apply a confined, homogeneous magnetic field which will have minimal effect on the magnetic shields. The auxiliary coils are designed to allow application of constant and gradient fields along all

three axes while preserving adequate homogeneity within the much longer spin precession region. Construction and assembly of the coil system is currently underway.

8

Summary and Outlook

In this thesis, I have described the experimental progress accomplished in the ACME experiment in the last few years. In chapter 1, we discussed the background, theoretical motivations, and basic methods used to measure the electron EDM. In chapter 2, the apparatus, measurement method, systematic errors, and results of the ACME II experiment (2015-18) were described. At the time of writing, the ACME II result remains the most stringent published upper bound on the electron EDM, showing that it is possible to probe new physics at the tens of TeV level with a tabletop experiment.

Chapter 3 provided an overview of the ACME III experimental campaign (2018-present), which has a goal of performing a measurement of the eEDM with an order of magnitude

improved sensitivity compared to ACME II. It consists of efforts to increase the statistical sensitivity, suppress systematic errors, and improve the general capabilities and reliability of the experimental apparatus, while also setting up the experiment in a new site at Northwestern University. We observed that all the major statistical upgrades have been experimentally demonstrated, resulting in a projected gain in statistical sensitivity of over an order of magnitude. In addition, the apparatus has been designed with features that are projected to suppress all the known systematic effects below the statistical uncertainty goal.

At the time of writing, all the major components of the apparatus have been designed and most have already been built (the beam box, electrostatic lens, vacuum chamber, collection optics, SiPM detectors, and magnetic shields). Construction and setup of other components (electric field plates, magnetic coils, and some lasers) is ongoing, and existing components are also being assembled and integrated into their final form at Northwestern.

In the last three chapters, I described three research efforts that have been a focus of my work in ACME III. The first is the H-state radiative lifetime measurement (chapter 4), which found a lifetime value several times longer than the ACME II precession time. This opened the way towards designing the ACME III apparatus with a five times longer precession time, securing a significant boost in statistical senstivity. Second are efforts to improve the photon detection efficiency (chapter 5), most significantly designing an improved set of collection optics that has been integrated with the new SiPM detectors. Next are the efforts to upgrade the data acquisition system (chapter 6), which now has the capability to save data from all molecular pulses, boosting the experiment's ability to diagnose and analyze any future anomalies, noise sources, and systematic effects. Finally, I also described the efforts to control magnetic fields in ACME III (chapter 7), with a special focus on the design of the magnetic field coils.

With all these improvements, within the next year or two, the ACME experiment is on track to commission a new ACME apparatus with capabilities to measure the electron

EDM at an unprecedented uncertainty of $\delta d_e \approx 4 \times 10^{-31}\ e \cdot \text{cm}$.

The ACME III campaign has attempted to maximize virtually every factor that can readily reduce the statistical uncertainty, a culmination of the ACME experiment efforts which first began about 15 years ago. Beyond ACME III, some avenues remain in improving the experiment. First is to take the final dataset for longer – up to a year instead of the usual 2-3 months in ACME II. Second is the photon detection efficiency, which in ACME III has been improved to ∼20%, and could possibly be further improved with optical cycling (Appendix A) or changing the method of detection from laser-induced fluorescence to detection of ionized molecules with a microchannel plate (MCP). Third is to find ways to increase the number of ThO molecules in the beam, such as by using a thermochemical beam source [147].

A significant leap in senstivity is more realistically achieved by building a new experiment with one of the many molecules (including polyatomics) which have been identified to possess favorable properties for measuring the eEDM, especially being laser-coolable (§1.4.3). Such experiments are already under development in several groups and show that there is a bright future for the role of tabletop experiments in the quest to further our knowledge of the most basic constituents of the universe.

A

Optical Cycling

Here we describe the progress of the optical cycling project which was undertaken to improve photon detection efficiency in ACME III. Ultimately, optical cycling ended up being shelved in favor of other, more practical and cost-efficient methods to increase the efficiency (Chapter 5). However, since this technique has been mentioned several times throughout the history of ACME, it is useful to report on what was learned regarding its feasibility.

A.1 Background

As discussed previously (§3.1.3), the photon detection efficiency in ACME II was about 5%, and the excess noise factor from the PMT was $F = 1.25$. The key idea of the optical

project is to improve this by stimulating the spin-precessed molecules to produce more than one photon per molecule on average by optically pumping on a transition with a high branching ratio ξ (where $0 < \xi < 1$). Naively, assuming enough photons are cycled, one could potentially achieve a signal gain of up to 25× (or a 5× gain in EDM statistical sensitivity), in which case the experiment would be limited only by the molecule shot noise.

A limiting factor is that ThO does not have many favorable transitions for optical cycling. The number of additional photons cycled N_{cyc} can be calculated from ξ as $N_{cyc} = 1/(1-\xi)-1$. The commonly used $C-X$ transition has $\xi = 0.74$ [141] which would give $N_{cyc} = 2.8$, a modest gain given the additional lasers and setup required. However, in 2014, Kokkin et al. characterized several new electronic transitions of ThO including the $I - X$ transition, which was measured to have $\xi = 0.91$ and thus $N_{cyc} \approx 10$, a potential order of magnitude gain in signal [152]. Thus, efforts were pursued to explore the possibility of using this transition for optical cycling in ACME III.

A.2 Phase measurement schemes with optical cycling

Two possible schemes for performing a phase measurement incorporating optical cycling detection are shown in Fig. A.2.1. Both schemes require changing from polarization switching detection to *parity switching detection*. Instead of probing the state using a rapidly switched pair of laser beams tuned to the same transition with orthogonal polarizations (as in ACME I and II), the schemes use laser beams tuned to two different parity sublevels of the excited state I. Parity switching detection has been previously experimentally demonstrated [133] and is physically equivalent to performing the $\tilde{\mathcal{P}}$ switch in ACME II.

Both schemes consist of two steps: a "state-shelving" step and an optical cycling step. In scheme a), a pair of state shelving lasers tuned to $|H, J = 1\rangle \rightarrow |I, J = 1, \mathcal{P} = \pm 1\rangle$ are shone onto the molecular beam at the end of the precession region. As the X state has

$|\Omega| = 0$, the parity of a rotational sub-level J is expressed by $\mathcal{P} = (-1)^J$. The decays

$$|I, J = 1, \mathcal{P} = +1\rangle \rightarrow |X, J = 1^-\rangle\,,$$
$$|I, J = 1, \mathcal{P} = -1\rangle \rightarrow |X, J = 0^+, 2^+\rangle \tag{A.1}$$

occur, "encoding" the phase information in the relative sizes of the populations of different rotational levels.[1] The molecules then proceed to the optical cycling region, which can be located either immediately adjacent to the shelving region or in a separate chamber downstream. Here, the phase is read out by sending in two rapidly switched sets of 512 nm optical cycling lasers: ($|X, J = 1^-\rangle \rightarrow |I, J = 1, \mathcal{P} = -1\rangle$ and $|X, J = 0^+, 2^+\rangle \rightarrow |I, J = 1, \mathcal{P} = -1\rangle$) and comparing the relative 512 nm fluorescence signal obtained using the two sets. Polarization switching of the lasers will also be required to avoid magnetic dark states (§A.4.1).

Scheme b) uses a pair of state-shelving lasers tuned to $|H, J = 1\rangle \rightarrow |I, J = 1^+, 2^-\rangle$, resulting in decays to $|X, J = 1^-, 2^+\rangle$. The phase is read out using a pair of rapidly switched optical cycling lasers tuned to $|X, J = 1^-\rangle \rightarrow |I, J = 1^+\rangle$ and $|X, J = 2^+\rangle \rightarrow |I, J = 2^0-\rangle$. While this scheme requires only two cycling lasers, it has a weakness in that the $|I, J = 1, 2\rangle$ states are separated by tens of GHz, requiring two 703 nm shelving lasers. In contrast, the parity doublet in of the excited state $|I, J = 1\rangle$ in scheme a) is only separated by 92.5 MHz, such that only a single 703 nm laser shifted with an AOM is required. Thus, both schemes require a total of four lasers for state shelving and cycling, assuming that cycling on the $I - X$ transition alone produces a sufficient number of photons. Otherwise, additional lasers will be needed to repump the population decaying to other states such as H, Q, and $X,\ \nu = 1$.

[1] Note that any pre-existing molecule population in rotational sub-levels of X would reduce the measurement contrast and thus may need to be transferred into another state with a long lifetime (e.g. $|X, \nu = 1\rangle$ states) before state shelving.

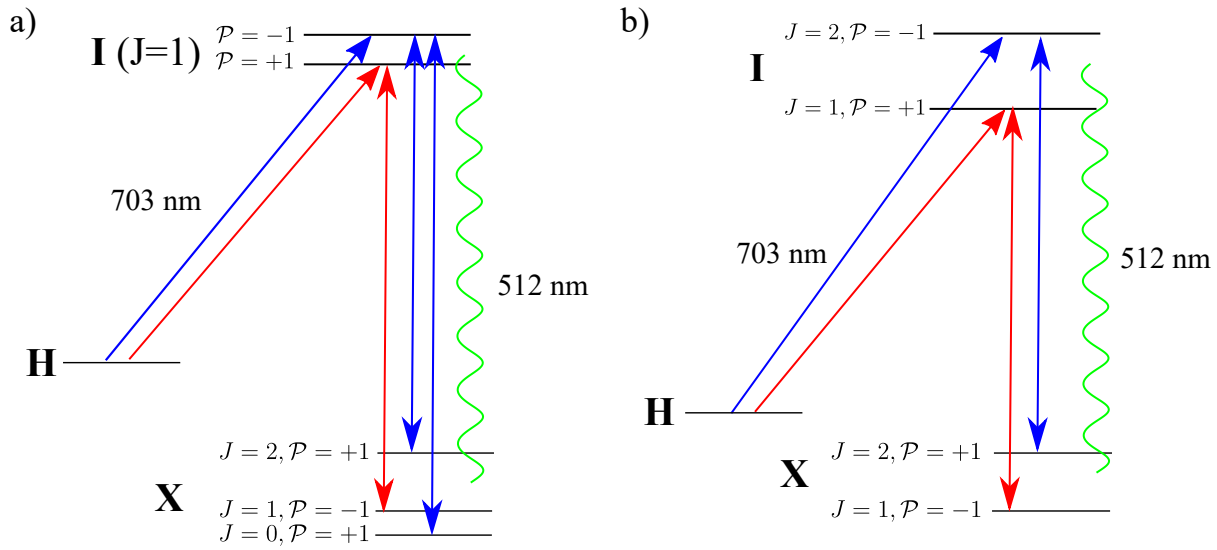

Figure A.2.1: **Two possible schemes to implement optical cycling detection.** Both schemes consist of a "state-shelving" step where the phase information is stored in different rotational levels of X and a cycling step where the information is extracted using parity-switching detection on the $X - I$ cycling transition.

A.3 Stray light in optical cycling detection

As the wavelength of the detected photons (512 nm) is the same as the wavelength of the cycling lasers, there will be a large background from scattered photons in the chamber which cannot be filtered out with a bandpass filter. Experimental tests estimate that the typical stray photon rate for a laser sent into a vacuum chamber orthogonal to the molecular beam is $\sim 10^{10}$ Hz/mW per detector, or $\sim$1 per 10^6 photons in the laser. Assuming 8 detectors, 100 mW laser light for cycling, an ACME III fluorescence signal of 10^8 Hz, and the requirement that the scatter needs to be at most two orders of magnitude smaller than the signal, one needs to suppress the stray photon rate to $\sim$1 in 10^9 per detector. This is not infeasible, as previous molecular beam setups have accomplished stray light suppression as low as 1 in 10^{11} [202, §5.9.2]. A typical suppression method is to blacken the vacuum chamber walls and use a series of apertures or baffles to collimate the laser beam [203, §5.6]. This is more readily accomplished by building a separate chamber after

the interaction region dedicated for optical cycling detection. Without the presence of field plates, such a chamber could also allow even better collection optics to collect more molecular fluorescence [204].[2]

A.4 Experimental tests

In 2018-19, tests of optical cycling were performed to explore its feasibility. In these tests, ThO molecules in the ground X state are produced with a standard cryogenic buffer gas beam source and optically pumped with a 512 nm linearly polarized laser tuned to the $|X, J' = 1^-\rangle \rightarrow |I, J = 1, \mathcal{P} = +1\rangle$ transition.[3] The polarization of the laser is alternated between $\hat{x}$-polarization, $\hat{z}$-polarization, and fast polarization switching between $\hat{x}$ and $\hat{z}$. A PMT records the laser-induced fluorescence at 512 nm. This results in signals F_X, F_Z, and F_{XZ}. The number of photons cycled can be measured by comparing these three signals, as explained below.

A.4.1 Expected fluorescence when optical pumping

Due to parity selection rules, the only single photon decays are back to $|X, J' = 1\rangle$. For the case of optical pumping with a single polarization $\hat{z}$, molecules can be excited from $M' = 1 \rightarrow M = 1$ or $M' = -1 \rightarrow M = -1$. ($M' = 1 \rightarrow M = 0$ is forbidden for $\Delta J = 0$.) Let us consider the population excited to $M = 1$. They have have a probability ξ of decaying back to $|X, J' = 1\rangle$, where ξ is the electro-vibrational branching ratio of $I - X$. Decays to $M' = 0$ and $M' = 1$ sub-levels occur with equal probability (which can be verified by calculating the Clebsch-Gordan coefficients). Those which decay to $M' = 1$ remain bright to the laser and will emit a second photon. The probability of this occurring

[2] One possible complication is that the parity switching time bins from the state-shelving step (which occurs in the interaction region) will be "smeared out" by the time the molecules arrive in the optical cycling chamber due to the velocity dispersion in the beam.

[3] Toptica DL Pro HP, locked to a high finesse ULE cavity.

is

$$p_2 = \left(\frac{\xi}{2}\right)\xi = \frac{\xi^2}{2}. \tag{A.2}$$

The cycle then repeats, resulting in continuous emission of photons until all molecules are in the dark state $M' = 0$. The total fluorescence produced can be calculated by summing the geometric series:

$$\begin{aligned} F_{+1,Z} &= P_1\left[(\xi) + \left(\frac{\xi}{2}\right)\xi + \left(\frac{\xi}{2}\right)^2\xi + ...\right] \\ &= P_1\xi\frac{1}{1-\frac{\xi}{2}} \\ &= \frac{2P_1\xi}{2-\xi}, \end{aligned} \tag{A.3}$$

where P_1 is the initial population in $|X, J' = 1, M' = 1\rangle$. The same result obtains for the case of $F_{-1,Z}$ i.e. decay from $M = -1$. Thus the total fluorescence from optical pumping with $\hat{z}$-polarized light is

$$F_Z = 2F_{+1,Z} = \frac{4P_1\xi}{2-\xi} = F_X, \tag{A.4}$$

where by symmetry we have also argued that the fluorescence obtained from using x-polarized light would be given by the same expression.

For the case where the polarization is switched rapidly between $\hat{x}$ and $\hat{z}$, there are no magnetic dark states in $|X, J' = 1\rangle$, so molecules in all 3 M' sub-levels will continuously cycle photons until they decay to other electic and/or vibrational states. The total fluorescence produced is

$$\begin{aligned} F_{XZ} &= 3P_1[\xi + \xi^2 + ...] \\ &= \frac{3P_1\xi}{1-\xi} \end{aligned} \tag{A.5}$$

Assuming that F_X and F_Z are detected with roughly equal efficiency, then one can calculate the ratio

$$R = \frac{F_{XZ}}{F_Z + F_X} = \frac{\frac{3P_1\xi}{1-\xi}}{\frac{8P_1\xi}{2-\xi}} = \frac{3}{8}\left(\frac{\xi-2}{\xi-1}\right). \tag{A.6}$$

From R, one can determine the number of photons cycled and the branching ratio ξ. To ensure equal detection efficiency between F_X and F_Z, a waveplate to rotate the laser polarization such that both polarizations are 45 degrees from the normal vector of the detector. About 7-10 mW of 512 nm laser light was typically used for each polarization, resulting in $\sim 80-90\%$ saturation.

A.4.2 Results

An example trace from the optical cycling tests is shown in Fig. A.4.1. This trace was taken with polarization switching at 400 kHz and performing 4 passes of the ~ 6 mW laser through the molecular beam. From this result, the ratio R is calculated to be about 1.5, implying an $I-X$ branching ratio of $\xi \approx 0.67$, markedly lower than the value of 0.91 measured by the Steimle group in Ref. [152]. This implies only about 2 additional photons generated from optical cycling.

Multiple methods were attempted to increase the number of photons cycled: varying the number of laser beam passes, polarization switching rate, laser power, collection optics, applying an electric field, and implementing the polarization switching using a Pockels cell (which has a better rise time compared to an AOM). The measured value of R remained about the same, changing by no more than $\sim 10\%$. By the end of the optical cycling project, the source of the discrepancy between these tests and the Steimle result was not yet well-understood.

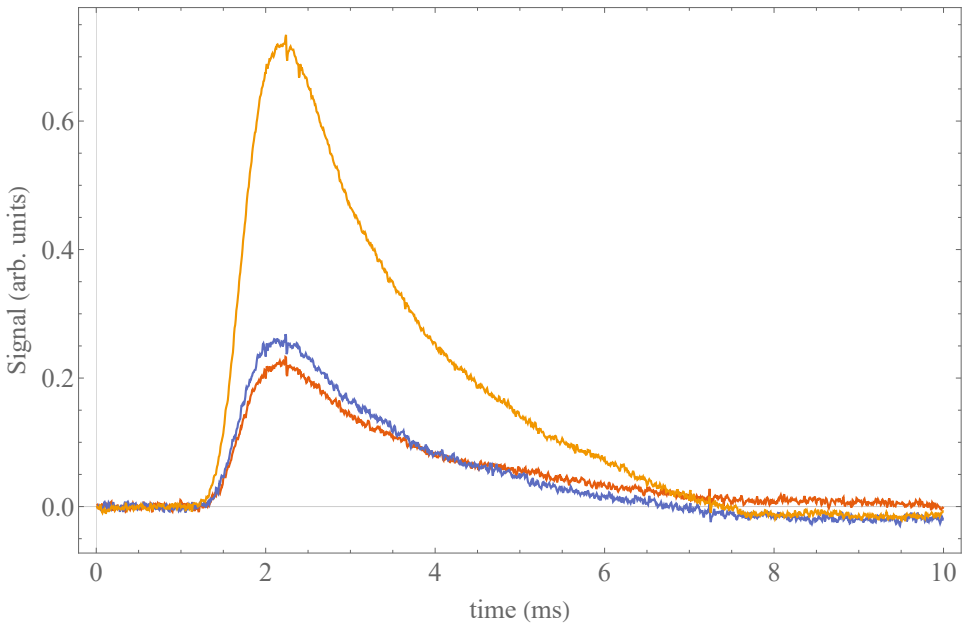

Figure A.4.1: **Optical cycling test results.** Fluorescence traces when optically pumping on the $X - I$ transition with $\hat{x}$-polarization (red), $\hat{z}$-polarization (blue), and rapidly switching between the two polarizations at 400 kHz (orange). The $I - X$ branching ratio ξ and the number of photons cycled can then be extracted.

A.5 Projected improvement in statistical sensitivity from optical cycling

Even if the experiment signal can be significantly increased with optical cycling, this would not translate to a proportional increase in statistical sensitivity of the phase measurement. The stochastic nature of optical cycling means that the number of photons a molecule emits before it decays to a dark state follows a statistical distribution which adds excess noise above the shot-noise limit [70, 205]. This noise is dependent on the branching ratio ξ, photon scattering rate r, photon scattering time T, and photon detection efficiency ϵ.[4] As calculated by Lasner and DeMille [205], the statistical uncertainty in a phase measurement with optical cycling is

$$\sigma = \frac{1}{2\sqrt{N}}\sqrt{F}, \tag{A.7}$$

where N is the number of molecules participating in the spin precession[5] and the excess noise factor F is defined as

$$F = 1 + \frac{1}{1 - e^{-rTb}}\left(\frac{bf}{\epsilon} + \frac{1 - 2b + 2be^{-rTb}(1 - rT(1-b)) - e^{-2rTb}}{1 - e^{-rTb}}\right), \tag{A.8}$$

where $b = 1 - \xi$ is the branching ratio to a dark state, f is the photodetector excess noise factor, and r, T, ϵ have been defined above. F incorporates all statistical gains from the increased number of photons in optical cycling and can thus be used for direct comparisons of σ with and without cycling.

For simplicity, let us assume that $rT \to \infty$ (i.e. the width of the cycling laser beam can be extended to ensure that cycling is performed "to completion"). Then the expression for

[4] If the rT is small then some molecules may not have decayed to a dark state by the end of the cycling process. If ϵ is small, then by the end of the cycling process some molecules may not have emitted any detected photons, while others may have emitted multiple detected photons. More detailed explanations can be found in Ref. [70, §6.4].

[5] This is slightly different from earlier expression of EDM uncertainty such as Eq. 1.23, where N denoted the number of *detected* molecules. Here, the photon detection efficiency is encoded in the excess noise factor F.

F simplifies to

$$F = 2 + b\left(\frac{f}{\epsilon} - 2\right). \tag{A.9}$$

In the case where there is no optical cycling performed (i.e. as in previous generations of ACME), then $b = 1$ and $rT \to \infty$, such that $F = f/\epsilon$.

Based on the improvements described in the preceding chapters, the approximate gain in EDM statistical sensitivity from optical cycling can be estimated as follows. We have

- $f = 1.07$ for the SiPM detectors (§5.2),
- $\epsilon \approx 0.15$, based on collection optics efficiency of ≈ 0.3 and SiPM detection efficiency of ≈ 0.5 (Chapter 5),
- $b = 0.09$ (assuming the best case scenario that the $\xi = 0.91$ measured in Ref. [152] is correct), or ~ 10 photons cycled.

This results in $F \approx 2.5$ with optical cycling and $F \approx 7$ without. Thus, the reduction in EDM statistical uncertainty implementing optical cycling using the ACME III apparatus would be $\sqrt{7/2.5} = 1.7$, which is a limited gain considering the large amount of engineering effort required (multiple additional lasers for state shelving and cycling, dealing with high levels of stray light, and changing detection scheme to parity switching detection). It is possible that by performing cycling in a dedicated chamber, one could increase the detection efficiency (from improved collection optics) by another factor of ~ 2. In this case, the combined statistical gain from cycling would only increase to 1.8. Even in the extreme case where one added multiple repumping lasers to reduce $b \to 0$, we would still have $F = 2$, or a statistical gain of 1.9. For comparison, the theoretical maximum statisical gain one could accomplish from conventional detection improvements is $\sqrt{1.07/0.15} = 2.7$. Thus, the stochastic nature of cycling limits its attractiveness as a method of improving photon detection in ACME.

A.6 Conclusion and future outlook

Given the limited numbers of cycled photons achieved experimental tests (§A.4) and the limited statistical gain that would be possible given the excess noise present in cycling (§A.5), the decision was made to discontinue the optical cycling project for ACME III and focus on conceptually simpler methods of improving the detection efficiency which are reported in Chapter 5. In future generations of the experiment, however, it is possible that such statisical gains would be worth pursuing. In that case, the priority would be to continue exploring the causes of the low number of cycled photons in the earlier experimental tests. This could be done by dramatically increasing the laser power, the number of laser beam passes, and progressively adding repumping lasers for the H and Q transitions to trace where the population ends up.

B

X-A and H-A transition dipole moment measurement

B.1 Background

The $X - A$ and $H - A$ transitions are of interest because they offer a STIRAP scheme requiring less power than the $X - C - H$ scheme used in ACME II. Such a scheme can then be used to perform horizontal STIRAP through the field plates, which provides more options to suppress polarization gradient-related systematic errors (§3.2.1). Knowledge of the $X - A$ and $H - A$ transition dipole moments would allow a more concrete estimate of

the laser powers required for STIRAP. Previously, the $X - A$ transition dipole moment had been measured in a previous dissertation [151, Eq. 4.22][1], obtaining $d_{A,0-X,0} = (0.071 \pm 0.038)\ \mathrm{ea_0}$. This measurement used a classical rate equation model with several simplifying assumptions (such as that the A-state lifetime is roughly equal to the C-state lifetime). In contrasts, the measurements in the present chapter were performed to obtain the transition dipole moments of the $X - A$ and $H - A$ transitions with improved precision and utilized a model where each input variable is experimentally determined.

B.2 Experimental method

The measurement technique used is the same used in recent measurements of other transition dipole moments such as $X - C$ [141, §B]. Optical pumping of the $X - A$ transition was performed using a locked 943 nm laser at different powers. The resulting saturation curve was then fit to a numerical model of the optical pumping dynamics in the molecular beam using optical Bloch equations [206, §2.3] An input parameter into the model is the $A - H$ branching ratio, which was previously measured to be [151, §4.3][2]

$$\xi_{A,0-H,0} = 0.29 \pm 0.07. \tag{B.1}$$

(Note that the subscript denotes the electronic states and the vibrational quantum numbers of the energy levels in the transition.) Assuming a diagonal Franck-Condon factor of 0.8665 [144], this implies that $\xi_{A,0-X,0} = 0.62$. The branching ratio ξ of a transition can be related to the transition dipole moment d by the equation [141, 152]

$$\xi = 3.317 \times 10^{-13} \left[\frac{\mathrm{m}^3}{\mathrm{D}^2\mathrm{s}}\right] \frac{d^2}{\lambda^3}\tau, \tag{B.2}$$

[1] Note that in this equation in the dissertation, $X - A$ is mistakenly referred to as $H - A$.

[2] In Ref. [151], this branching ratio value was measured separately from the earlier $X - A$ transition dipole measurement, preserving the independence of the current measurement.

where λ is the wavelength of the transition and τ is the excited state lifetime.

Other model input parameters include the laser beam profile, molecular beam height, longitudinal velocity distribution, and Doppler width. All of these were independently measured before being entered into the model, leaving only the transition dipole moment d as the free parameter to be fitted to the power scan data.

B.3 Results

The resulting power scan is shown in Fig. B.3.1. From here, we obtain $d_{A,0-X,0} = 0.35$ D.

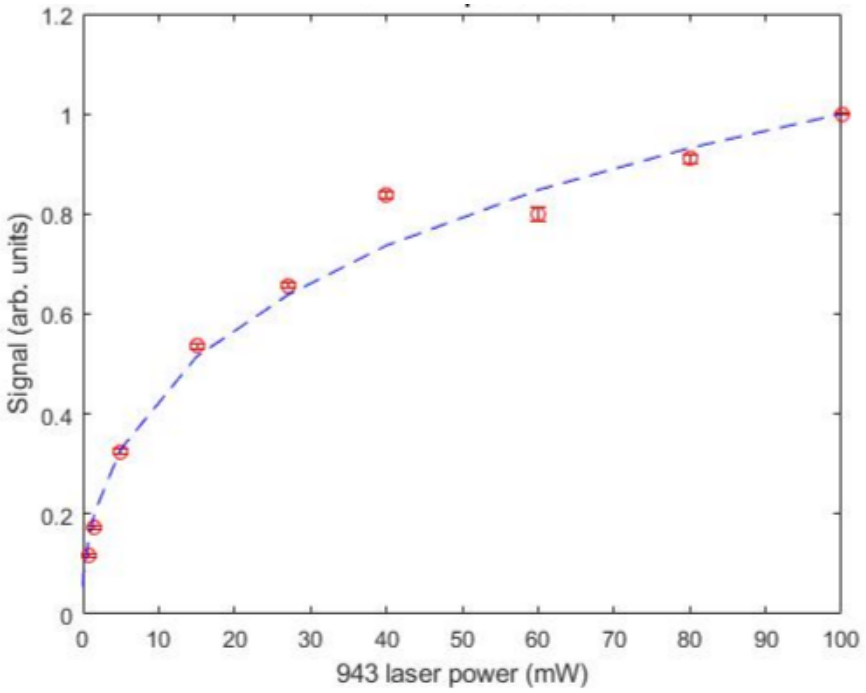

Figure B.3.1: **Results of the $X-A$ transition dipole moment measurement.** X-A laser power scan data (red circles) and fitted model (blue dashed line).

The systematic uncertainty of the measurement is estimated by fitting the model to a second, earlier set of data which was taken using half the laser beam height, resulting in less saturation. From this, an uncertainty of 0.1 D is obtained. The inputs to the model

were also varied over the range of their uncertainties. These varied the results of the measurement by much less than the systematic uncertainty. Hence, we obtain the result

$$d_{A,0-X,0} = (0.4 \pm 0.1)\ \mathrm{D} = (0.15 \pm 0.04)\ \mathrm{ea_0}. \tag{B.3}$$

from which one can use Eq. B.2 to obtain the lifetime of the A-state, $\tau = (10 \pm 5)\ \mu$s. This relatively longer lifetime may partially explain the discrepancy between this measurement and the earlier less precise measurement which assumed the lifetime to be the same as the C-state lifetime of 490 ns [151].

Using the previously known branching ratio $\xi_{A,0-H,0}$, from this result we can also deduce that

$$d_{A,0-H,0} = (0.8 \pm 0.2)\ \mathrm{D} = (0.31 \pm 0.09)\ \mathrm{ea_0}. \tag{B.4}$$

B.4 Implications for STIRAP

To estimate the laser powers required to perform $X - A - H$ STIRAP,[3] one can assume that the same Rabi frequency needs to be achieved in the new system for each transition:

$$\begin{aligned} \Omega_{X-C} &= \Omega_{X-A} \\ \implies d_{X-C}\sqrt{P_{X-C}} &= d_{X-A}\sqrt{P_{X-A}} \\ \implies P_{X-A} &= \frac{d_{X-C}^2}{d_{X-A}^2} P_{X-C}. \end{aligned} \tag{B.5}$$

Using $d_{X-C} = 1.27(5)$ D [141], $d_{X-A} = 0.4(1)$ D (this work), and $P_{X-C} = 15$ mW [140], we obtain $P_{X-A} = 151$ mW. A similar equation can be written to estimate for the $H - A$ and $H - C$ transitions. We use $d_{H-C} = 0.05(2)$ D [151, Eq. 4.19], $d_{H-A} = 0.8(2)$ D, and $P_{H-C} = 10$ W, obtaining $P_{H-A} = 40$ mW. If the most conservative values of the transition

[3] This section is based on a write-up by Cole Meisenhelder.

dipole moments within their quoted uncertainty ranges are used,[4] one obtains $P_{X-A} = 290$ mW and $P_{H-A} = 140$ mW.

[4]In other words, the lowest values for d_{H-A} and the highest for d_{H-C}, and similarly for d_{X-A} and d_{X-C}.

Bibliography

[1] 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).

[2] CMS Collaboration, "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC," Physics Letters B **716**, 30–61 (2012).

[3] D. Hanneke, S. Fogwell, and G. Gabrielse, "New measurement of the electron magnetic moment and the fine structure constant," Physical Review Letters **120801**, 1–4 (2008).

[4] X. Fan, T. G. Myers, B. A. D. Sukra, and G. Gabrielse, "Measurement of the electron magnetic moment," Phys. Rev. Lett. **130**, 071801 (2023).

[5] L. Canetti, M. Drewes, and M. Shaposhnikov, "Matter and antimatter in the universe," New Journal of Physics **14**, 095012 (2012).

[6] G. Bertone and D. Hooper, "History of dark matter," Reviews of Modern Physics **90**, 045002 (2018).

[7] M. Li, X.-D. Li, S. Wang, and Y. Wang, "Dark energy: A brief review," Frontiers of Physics **8**, 828–846 (2013).

[8] T. S. Virdee, "Beyond the standard model of particle physics," Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **374**, 20150259 (2016).

[9] Y. Gouttenoire, *Beyond the Standard Model Cocktail: A Modern and Comprehensive Review of the Major Open Puzzles in Theoretical Particle Physics and Cosmology with a Focus on Heavy Dark Matter*, Springer Theses (Springer International Publishing, Cham, Switzerland, 2022).

[10] J. Jackson, *Classical Electrodynamics* (Wiley, 1998).

[11] D. Budker, D. Kimball, and D. DeMille, *Atomic Physics: An Exploration through Problems and Solutions* (OUP Oxford, 2008).

[12] 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).

[13] M. Sozzi, *Discrete Symmetries and CP Violation* (Oxford University Press, New York, USA, 2008).

[14] T. D. Lee and C. N. Yang, "Question of parity conservation in weak interactions," Phys. Rev. **104**, 254–258 (1956).

[15] 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).

[16] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, "Evidence for the 2π decay of the K_2^0 meson," Physical Review Letters **13**, 138–140 (1964).

[17] Belle Collaboration, "Observation of large CP violation in the neutral B meson system," Phys. Rev. Lett. **87**, 091802 (2001).

[18] LHCb Collaboration, "Observation of CP violation in charm decays," Phys. Rev. Lett. **122**, 211803 (2019).

[19] I. B. Khriplovich and S. K. Lamoreaux, *CP Violation without Strangeness* (Springer, New York, 1997).

[20] R. Streater and A. Wightman, *PCT, Spin and Statistics, and All That*, Mathematical Physics Monograph Series (W.A. Benjamin, 1964).

[21] G. Gabrielse, A. Khabbaz, D. S. Hall, C. Heimann, H. Kalinowsky, and W. Jhe, "Precision mass spectroscopy of the antiproton and proton using simultaneously trapped particles," Phys. Rev. Lett. **82**, 3198–3201 (1999).

[22] S. Ulmer, C. Smorra, A. Mooser, K. Franke, H. Nagahama, G. Schneider, T. Higuchi, S. Van Gorp, K. Blaum, Y. Matsuda, W. Quint, J. Walz, and Y. Yamazaki, "High-precision comparison of the antiproton-to-proton charge-to-mass ratio," Nature **524**, 196–199 (2015).

[23] ALPHA Collaboration, "Characterization of the 1s–2s transition in antihydrogen," Nature **557**, 71–75 (2018).

[24] V. A. Kostelecký and N. Russell, "Data tables for Lorentz and CPT violation," Reviews of Modern Physics **83**, 11–31 (2011).

[25] AMS Collaboration, "Search for antihelium in cosmic rays," Physics Letters B **461**, 387–396 (1999).

[26] G. Steigman, "When clusters collide: constraints on antimatter on the largest scales," Journal of Cosmology and Astroparticle Physics **2008**, 001 (2008).

[27] 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 (2020).

[28] A. Dolgov, "Non-GUT baryogenesis," Physics Reports **222**, 309–386 (1992).

[29] M. Dine and A. Kusenko, "Origin of the matter-antimatter asymmetry," Rev. Mod. Phys. **76**, 1–30 (2003).

[30] A. D. Sakharov, "Violation of CP invariance, C asymmetry and baryon asymmetry of the universe," Pisma Zh. Eksp. Teor. Fiz. **5**, 1967 (1967).

[31] R. L. Workman and Others (Particle Data Group), "Review of Particle Physics," PTEP **2022**, 083C01 (2022).

[32] M. B. Gavela, P. Hernández, J. Orloff, and O. Pène, "Standard model CP-violation and baryon asymmetry," Modern Physics Letters A **09**, 795–809 (1994).

[33] J. M. Cline, "Baryogenesis," in *Particle Physics and Cosmology: The Fabric of Spacetime: Lecture Notes of the Les Houches Summer School 2006*, edited by F. Bernardeau, C. Grojean, and J. Dalibard (Elsevier Science, Amsterdam, Netherlands, 2007) 1st ed., pp. 53–116.

[34] D. E. Morrissey and M. J. Ramsey-Musolf, "Electroweak baryogenesis," New Journal of Physics **14**, 125003 (2012).

[35] J. L. Barrow, L. Broussard, J. M. Cline, P. S. B. Dev, M. Drewes, G. Elor, S. Gardner, J. Ghiglieri, J. Harz, Y. Kamyshkov, J. Klaric, L. W. Koerner, B. Laurent, R. McGehee, M. Postma, B. Shakya, R. Shrock, J. van de Vis, and G. White, "Theories and experiments for testable baryogenesis mechanisms: A Snowmass white paper," (2022), arXiv:2203.07059 [hep-ph] .

[36] D. Bödeker and W. Buchmüller, "Baryogenesis from the weak scale to the grand unification scale," Rev. Mod. Phys. **93**, 035004 (2021).

[37] Y. Li, S. Profumo, and M. Ramsey-Musolf, "Bino-driven electroweak baryogenesis with highly suppressed electric dipole moments," Physics Letters B **673**, 95–100 (2009).

[38] D. Egana-Ugrinovic, "The minimal fermionic model of electroweak baryogenesis," Journal of High Energy Physics **2017**, 64 (2017).

[39] Y. Ema, T. Gao, and M. Pospelov, "Standard model prediction for paramagnetic electric dipole moments," Phys. Rev. Lett. **129**, 231801 (2022).

[40] Y. Yamaguchi and N. Yamanaka, "Large long-distance contributions to the electric dipole moments of charged leptons in the standard model," Phys. Rev. Lett. **125**, 241802 (2020).

[41] Muon $g-2$ Collaboration (Muon $g-2$ Collaboration), "Measurement of the positive muon anomalous magnetic moment to 0.46 ppm," Phys. Rev. Lett. **126**, 141801 (2021).

[42] G. Dorsch, S. Huber, T. Konstandin, and J. No, "A second higgs doublet in the early universe: baryogenesis and gravitational waves," Journal of Cosmology and Astroparticle Physics **2017**, 052 (2017).

[43] J. L. Feng, "Naturalness and the Status of Supersymmetry," Annual Review of Nuclear and Particle Science **63**, 351–382 (2013).

[44] J. L. Feng, "Dark matter candidates from particle physics and methods of detection," Annual Review of Astronomy and Astrophysics **48**, 495–545 (2010), https://doi.org/10.1146/annurev-astro-082708-101659 .

[45] ATLAS collaboration, "Search for squarks and gluinos in final states with jets and missing transverse momentum using 139 fb^{-1} of $\sqrt{s} = 13$ TeV pp collision data with the ATLAS detector," Journal of High Energy Physics **2021**, 143 (2021).

[46] W. Bernreuther and M. Suzuki, "The electric dipole moment of the electron," Reviews of Modern Physics **63**, 313–340 (1991).

[47] N. Fortson, P. Sandars, and S. Barr, "The search for a permanent electric dipole moment," Physics Today **56**, 33–39 (2003).

[48] 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, Vol 63, No 1 **71**, 21–74 (2013).

[49] ACME Collaboration, V. Andreev, D. G. Ang, D. DeMille, J. M. Doyle, G. Gabrielse, J. Haefner, N. R. Hutzler, Z. Lasner, C. Meisenhelder, B. R. O'Leary, C. D. Panda, A. D. West, E. P. West, and X. Wu, "Improved limit on the electric dipole moment of the electron," Nature **562**, 355–360 (2018).

[50] 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).

[51] 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).

[52] S. A. R. Ellis and G. L. Kane, "Theoretical prediction and impact of fundamental electric dipole moments," Journal of High Energy Physics **2016**, 77 (2016).

[53] Y. Nakai and M. Reece, "Electric dipole moments in natural supersymmetry," Journal of High Energy Physics **2017**, 31 (2017).

[54] C. Cesarotti, Q. Lu, Y. Nakai, A. Parikh, and M. Reece, "Interpreting the electron EDM constraint," Journal of High Energy Physics **2019**, 59 (2019).

[55] G. Panico, A. Pomarol, and M. Riembau, "EFT approach to the electron electric dipole moment at the two-loop level," Journal of High Energy Physics **2019**, 90 (2019).

[56] T. Appelquist and G.-H. Wu, "Electroweak chiral lagrangian and CP-violating effects in technicolor theories," Phys. Rev. D **51**, 240–250 (1995).

[57] T. Appelquist, M. Piai, and R. Shrock, "Lepton dipole moments in extended technicolor models," Physics Letters B **593**, 175–180 (2004).

[58] T. S. Roussy, D. A. Palken, W. B. Cairncross, B. M. Brubaker, D. N. Gresh, M. Grau, K. C. Cossel, K. B. Ng, Y. Shagam, Y. Zhou, V. V. Flambaum, K. W. Lehnert, J. Ye, and E. A. Cornell, "Experimental constraint on axionlike particles over seven orders of magnitude in mass," Phys. Rev. Lett. **126**, 171301 (2021).

[59] O. Grasdijk, O. Timgren, J. Kastelic, T. Wright, S. Lamoreaux, D. DeMille, K. Wenz, M. Aitken, T. Zelevinsky, T. Winick, and D. Kawall, "CeNTREX: a new search for time-reversal symmetry violation in the ^{205}Tl nucleus," Quantum Science and Technology **6**, 044007 (2021).

[60] B. Graner, Y. Chen, E. G. Lindahl, and B. R. Heckel, "Reduced limit on the permanent electric dipole moment of ^{199}Hg," Physical Review Letters **116**, 161601 (2016).

[61] Allmendinger, Fabian, Engin, Ilhan, Grasdijk, Olivier, Heil, Werner, Jungmann, Klaus, Karpuk, Sergei, Krause, Hans-Joachim, Niederländer, Benjamin, Offenhäusser, Andreas, Repetto, Maricel, Schmidt, Ulrich, Willmann, Lorenz, and Zimmer, Stefan, "A new limit of the ^{129}Xenon electric dipole moment," EPJ Web Conf. **219**, 02003 (2019).

[62] C. Abel, S. Afach, N. J. Ayres, C. A. Baker, G. Ban, G. Bison, K. Bodek, V. Bondar, M. Burghoff, E. Chanel, Z. Chowdhuri, P.-J. Chiu, B. Clement, C. B. Crawford, M. Daum, S. Emmenegger, L. Ferraris-Bouchez, M. Fertl, P. Flaux, B. Franke, A. Fratangelo, P. Geltenbort, K. Green, W. C. Griffith, M. van der Grinten, Z. D. Grujić, P. G. Harris, L. Hayen, W. Heil, R. Henneck, V. Hélaine, N. Hild, Z. Hodge, M. Horras, P. Iaydjiev, S. N. Ivanov, M. Kasprzak, Y. Kermaidic, K. Kirch, A. Knecht, P. Knowles, H.-C. Koch, P. A. Koss, S. Komposch, A. Kozela, A. Kraft, J. Krempel, M. Kuźniak, B. Lauss, T. Lefort, Y. Lemière, A. Leredde, P. Mohanmurthy, A. Mtchedlishvili, M. Musgrave, O. Naviliat-Cuncic, D. Pais, F. M. Piegsa, E. Pierre, G. Pignol, C. Plonka-Spehr, P. N. Prashanth, G. Quéméner, M. Rawlik, D. Rebreyend, I. Rienäcker, D. Ries, S. Roccia, G. Rogel, D. Rozpedzik, A. Schnabel, P. Schmidt-Wellenburg, N. Severijns, D. Shiers, R. Tavakoli Dinani, J. A. Thorne, R. Virot, J. Voigt, A. Weis, E. Wursten, G. Wyszynski, J. Zejma, J. Zenner, and G. Zsigmond, "Measurement of the permanent electric dipole moment of the neutron," Phys. Rev. Lett. **124**, 081803 (2020).

[63] PanEDM collaboration, "The PanEDM neutron electric dipole moment experiment at the ILL," EPJ Web Conf. **219**, 02006 (2019).

[64] nEDM@SNS collaboration, "The neutron electric dipole moment experiment at the Spallation Neutron Source," EPJ Web Conf. **219**, 02005 (2019).

[65] T. Feder, "CERN considers a 100 TeV circular hadron collider," Physics Today (2019), 10.1063/PT.6.2.20190205a.

[66] N. F. Ramsey, "A new molecular beam resonance method," Phys. Rev. **76**, 996–996 (1949).

[67] N. F. Ramsey, "A molecular beam resonance method with separated oscillating fields," Physical Review **78**, 695–699 (1950).

[68] E. D. Commins and D. DeMille, "The electric dipole moment of the electron," in *Lepton Dipole Moments*, edited by B. L. Roberts and W. J. Marciano (World Scientific, Singapore, 2010) Chap. 14, pp. 519–581.

[69] 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 **43** (2010), 10.1088/0953-4075/44/7/079803.

[70] Z. Lasner, *Order-of-Magnitude-Tighter Bound on the Electron Electric Dipole Moment*, Ph.D. dissertation, Yale University (2019).

[71] J. H. Smith, E. M. Purcell, and N. F. Ramsey, "Experimental limit to the electric dipole moment of the neutron," Phys. Rev. **108**, 120–122 (1957).

[72] G. Feinberg, "Effects of an electric dipole moment of the electron on the hydrogen energy levels," Phys. Rev. **112**, 1637–1642 (1958).

[73] E. E. Salpeter, "Some atomic effects of an electronic electric dipole moment," Physical Review **112**, 1642–1648 (1958).

[74] 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).

[75] D. T. Wilkinson and H. R. Crane, "Precision measurement of the g factor of the free electron," Phys. Rev. **130**, 852–863 (1963).

[76] G. R. Burleson and H. Kendall, "An experimental search for dipole structure of the electron," Nuclear Physics **19**, 68–78 (1960).

[77] J. Goldemberg and Y. Torizuka, "Upper limit of the electric dipole moment of the electron," Phys. Rev. **129**, 2580–2581 (1963).

[78] P. G. H. Sandars, "The electric dipole moment of an atom," Physics Letters **14**, 194–196 (1965).

[79] L. I. Schiff, "Measurability of nuclear electric dipole moments," Physical Review **132**, 2194–2200 (1963).

[80] 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).

[81] E. R. Meyer, J. L. Bohn, and M. P. Deskevich, "Candidate molecular ions for an electron electric dipole moment experiment," Phys. Rev. A **73**, 062108 (2006).

[82] E. A. Hinds, "Testing time reversal symmetry using molecules," Physica Scripta **1997**, 34 (1997).

[83] N. R. Hutzler, *A New Limit on the Electron Electric Dipole Moment: Beam Production, Data Interpretation, and Systematics*, Ph.D. dissertation, Harvard University (2014).

[84] D. P. DeMille, "Searches for new, massive particles with AMO experiments," in *Current Trends in Atomic Physics*, edited by A. Browaeys, T. Lahaye, T. Porto, C. S. Adams, M. Weidemuller, and L. F. Cugliandolo (Oxford University Press, Oxford, United Kingdom, 2019) illustrated edition ed., pp. 29–81.

[85] P. Sandars, "Enhancement factor for the electric dipole moment of the valence electron in an alkali atom," Physics Letters **22**, 290–291 (1966).

[86] W. R. Johnson, D. S. Guo, M. Idrees, and J. Sapirstein, "Weak-interaction effects in heavy atomic systems," Phys. Rev. A **32**, 2093–2099 (1985).

[87] Z. W. Liu and H. P. Kelly, "Analysis of atomic electric dipole moment in thallium by all-order calculations in many-body perturbation theory," Phys. Rev. A **45**, R4210–R4213 (1992).

[88] E. Lipworth, A. Adler, J. Carrico, T. Stein, and P. Sandars, "Electric-dipole moment of the electron," in *Bulletin of the American Physical Society*, Vol. 11 (American Physical Society, 1966) p. 403.

[89] C. O. Thornburg, Jr. and J. G. King, "Search for an Electric Dipole Moment in the Cesium Atom," in *Bulletin of the American Physical Society*, Vol. 11 (American Physical Society, 1966) p. 329.

[90] T. S. Stein, J. P. Carrico, E. Lipworth, and M. C. Weisskopf, "Electric dipole moment of the cesium atom. a new upper limit to the electric dipole moment of the free electron," Phys. Rev. Lett. **19**, 741–743 (1967).

[91] 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," Phys. Rev. Lett. **21**, 1645–1648 (1968).

[92] J. P. Carrico, E. Lipworth, P. G. H. Sandars, T. S. Stein, and M. C. Weisskopf, "Electric dipole moments of alkali atoms. a limit to the electric dipole moment of the free electron," Phys. Rev. **174**, 125–138 (1968).

[93] 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).

[94] E. D. Commins, S. B. Ross, D. DeMille, and B. C. Regan, "Improved experimental limit on the electric dipole moment of the electron," Phys. Rev. A **50**, 2960–2977 (1994).

[95] 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).

[96] 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 (New York, N.Y.) **343**, 269–72 (2014).

[97] 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," (2022), arXiv:2212.11841 [physics.atom-ph] .

[98] P. 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).

[99] J. Angel, P. Sandars, and M. Tinker, "Observation of a $v \times E$ effect in an electric dipole moment experiment using a reversible atomic beam machine," Physics Letters A **25**, 160–161 (1967).

[100] 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, Molecular and Optical **3**, 1620–1635 (1970).

[101] 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).

[102] M. G. Kozlov, "Enhancement of the electric dipole moment of the electron in the YbF molecule," Journal of Physics B: Atomic, Molecular and Optical Physics **30**, L607 (1997).

[103] 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).

[104] M. Denis, M. S. Nørby, H. J. A. Jensen, A. S. P. Gomes, M. K. Nayak, S. Knecht, and T. Fleig, "Theoretical study on ThF^+, a prospective system in search of time-reversal violation," New Journal of Physics **17**, 043005 (2015).

[105] A. N. Petrov, N. S. Mosyagin, T. A. Isaev, and A. V. Titov, "Theoretical study of HfF^+ in search of the electron electric dipole moment," Phys. Rev. A **76**, 030501 (2007).

[106] 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^+," Phys. Rev. A **105**, 022823 (2022).

[107] J. Lee, E. Meyer, R. Paudel, J. Bohn, and A. Leanhardt, "An electron electric dipole moment search in the $X^3\Delta_1$ ground state of tungsten carbide molecules," Journal of Modern Optics **56**, 2005–2012 (2009).

[108] M. G. Kozlov and D. DeMille, "Enhancement of the electric dipole moment of the electron in PbO," Phys. Rev. Lett. **89**, 133001 (2002).

[109] 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: Atomic, Molecular, and Optical Physics **87**, 1–19 (2013).

[110] M. Abe, V. S. Prasannaa, and B. P. Das, "Application of the finite-field coupled-cluster method to calculate molecular properties relevant to electron electric-dipole-moment searches," Phys. Rev. A **97**, 032515 (2018).

[111] P. Aggarwal, H. L. Bethlem, A. Borschevsky, M. Denis, K. Esajas, P. A. B. Haase, Y. Hao, S. Hoekstra, K. Jungmann, T. B. Meijknecht, M. C. Mooij, R. G. E. Timmermans, W. Ubachs, L. Willmann, A. Zapara, and The NL-eEDM collaboration, "Measuring the electric dipole moment of the electron in BaF," The European Physical Journal D **72** (2018), 10.1140/epjd/e2018-90192-9.

[112] 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).

[113] 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 (2016).

[114] Y. Zhou, Y. Shagam, W. B. Cairncross, K. B. Ng, T. S. Roussy, T. Grogan, K. Boyce, A. Vigil, M. Pettine, T. Zelevinsky, J. Ye, and E. A. Cornell, "Second-scale coherence measured at the quantum projection noise limit with hundreds of molecular ions," Phys. Rev. Lett. **124**, 053201 (2020).

[115] L. Caldwell, T. S. Roussy, T. Wright, W. B. Cairncross, Y. Shagam, K. B. Ng, N. Schlossberger, S. Y. Park, A. Wang, J. Ye, and E. A. Cornell, "Systematic and statistical uncertainty evaluation of the HfF^+ electron electric dipole moment experiment," (2022), arXiv:2212.11837 [physics.atom-ph] .

[116] C. J. Ho, J. A. Devlin, I. M. Rabey, P. Yzombard, J. Lim, S. C. Wright, N. J. Fitch, E. A. Hinds, M. R. Tarbutt, and B. E. Sauer, "New techniques for a measurement of the electron's electric dipole moment," New Journal of Physics **22**, 053031 (2020).

[117] X. Alauze, J. Lim, M. Trigatzis, S. Swarbrick, F. Collings, N. Fitch, B. Sauer, and M. Tarbutt, "An ultracold molecular beam for testing fundamental physics," Quantum Science and Technology **6** (2021), 2058-9565/ac107e.

[118] A. L. Collopy, S. Ding, Y. Wu, I. A. Finneran, L. Anderegg, B. L. Augenbraun, J. M. Doyle, and J. Ye, "3d magneto-optical trap of yttrium monoxide," Phys. Rev. Lett. **121**, 213201 (2018).

[119] D. J. McCarron, M. H. Steinecker, Y. Zhu, and D. DeMille, "Magnetic trapping of an ultracold gas of polar molecules," Phys. Rev. Lett. **121**, 013202 (2018).

[120] H. J. Williams, L. Caldwell, N. J. Fitch, S. Truppe, J. Rodewald, E. A. Hinds, B. E. Sauer, and M. R. Tarbutt, "Magnetic trapping and coherent control of laser-cooled molecules," Phys. Rev. Lett. **120**, 163201 (2018).

[121] S. L. Campbell, R. B. Hutson, G. E. Marti, A. Goban, N. D. Oppong, R. L. McNally, L. Sonderhouse, J. M. Robinson, W. Zhang, B. J. Bloom, and J. Ye, "A Fermi-degenerate three-dimensional optical lattice clock," Science **358**, 90–94 (2017).

[122] N. J. Fitch, J. Lim, E. A. Hinds, B. E. Sauer, and M. R. Tarbutt, "Methods for measuring the electron's electric dipole moment using ultracold ybf molecules," Quantum Science and Technology **6**, 014006 (2020).

[123] I. Kozyryev and N. R. Hutzler, "Precision measurement of time-reversal symmetry violation with laser-cooled polyatomic molecules," Phys. Rev. Lett. **119**, 133002 (2017).

[124] A. C. Vutha, M. Horbatsch, and E. A. Hessels, "Oriented polar molecules in a solid inert-gas matrix: A proposed method for measuring the electric dipole moment of the electron," Atoms **6** (2018), 10.3390/atoms6010003.

[125] 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," Phys. Rev. A **98**, 032513 (2018).

[126] S. Shin, Y. Kim, E.-s. Moon, D. H. Lee, H. Kang, and H. Kang, "Generation of strong electric fields in an ice film capacitor," The Journal of Chemical Physics **139**, 074201 (2013).

[127] W. Demtröder, *Molecular Physics: Theoretical Principles and Experimental Methods*, Physics Textbook (Wiley, 2008).

[128] 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).

[129] E. R. Meyer and J. L. Bohn, "Prospects for an electron electric-dipole moment search in metastable ThO and ThF^+," Physical Review A: Atomic, Molecular, and Optical Physics **78**, 010502 (2008).

[130] L. V. Skripnikov, A. N. Petrov, and A. V. Titov, "Communication: Theoretical study of tho for the electron electric dipole moment search," The Journal of Chemical Physics **139**, 221103 (2013).

[131] T. Fleig and M. K. Nayak, "Electron electric dipole moment and hyperfine interaction constants for ThO," Journal of Molecular Spectroscopy **300**, 16–21 (2014).

[132] L. V. Skripnikov and A. V. Titov, "Theoretical study of thorium monoxide for the electron electric dipole moment search: Electronic properties of $H^3\Delta_1$ in ThO." Journal of Chemical Physics **142**, 024301 (2015).

[133] 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: Atomic, Molecular, and Optical Physics **88**, 013844 (2013).

[134] J. M. Brown and A. Carrington, *Rotational Spectroscopy of Diatomic Molecules*, Cambridge Molecular Science (Cambridge University Press, 2003).

[135] A. C. Vutha, *A Search for the Electric Dipole Moment of the Electron Using Thorium Monoxide*, Ph.D. dissertation, Yale University (2011).

[136] N. Solmeyer, *Progress Towards An Electron Electric Dipole Moment Measurement With Laser-Cooled Atoms*, Ph.D. dissertation, Pennsylvania State University (2012).

[137] P. Hamilton, *Progress toward Searching for Electron Electric Dipole Moment using PbO*, Ph.D. dissertation, Yale University (2008).

[138] N. R. Hutzler, M. F. Parsons, Y. V. Gurevich, P. W. Hess, E. Petrik, B. Spaun, A. C. Vutha, D. DeMille, G. Gabrielse, and J. M. Doyle, "A cryogenic beam of refractory, chemically reactive molecules with expansion cooling," Physical Chemistry Chemical Physics **13**, 18976 (2011).

[139] I. M. Rabey, *Improved shot noise limit of the YbF EDM experiment*, Ph.D. dissertation, Imperial College London (2016).

[140] 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 ThO $H^3\Delta_1$ states for an improved electron electric-dipole-moment measurement," Physical Review A: Atomic, Molecular, and Optical Physics **93**, 1–9 (2016).

[141] 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).

[142] 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).

[143] G. Edvinsson and L.-E. Selin, "The band spectrum of thorium oxide," Physics Letters **9**, 238–239 (1964).

[144] T. Wentink and R. J. Spindler, "The isoelectronic series ScF through ThO—I notes on the band spectra of TiO, HfO and ThO," Journal of Quantitative Spectroscopy and Radiative Transfer **12**, 1569–1590 (1972).

[145] 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 ThO $H^3\Delta_1$ for the electron electric-dipole-moment search," Physical Review A: Atomic, Molecular, and Optical Physics **89**, 062505 (2014).

[146] 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. dissertation, Yale University (2017).

[147] E. Petrik West, *A Thermochemical Cryogenic Buffer Gas Beam Source of ThO for Measuring the Electric Dipole Moment of the Electron*, Ph.D. dissertation, Harvard University (2017).

[148] C. D. Panda, *Order of Magnitude Improved Limit on the Electric Dipole Moment of the Electron*, Ph.D. dissertation, Harvard University (2019).

[149] K. Bergmann, H. Theuer, and B. W. Shore, "Coherent population transfer among quantum states of atoms and molecules," Reviews of Modern Physics **70**, 1003–1023 (1998).

[150] B. W. Shore, "Picturing stimulated Raman adiabatic passage: a STIRAP tutorial," Adv. Opt. Photon. **9**, 563–719 (2017).

[151] B. N. Spaun, *A Ten-Fold Improvement to the Limit of the Electron Electric Dipole Moment*, Ph.D. dissertation, Harvard University (2014).

[152] 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: Atomic, Molecular, and Optical Physics **90**, 062503 (2014).

[153] P. W. Hess, *Improving the Limit on the Electron EDM : Data Acquisition and Systematics Studies in the ACME Experiment*, Ph.D. dissertation, Harvard University (2014).

[154] *Lossless Communication with Network Streams: Components, Architecture, and Performance*, National Instruments Corporation (2022).

[155] P. J. Huber, *Robust statistics*, 2nd ed., Wiley Series in Probability and Statistics (Wiley, Hoboken, New Jersey, 2009).

[156] B. Efron and R. Tibshirani, "Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy," Statistical Science **1**, 54–75 (1986).

[157] M. Brownnutt, M. Kumph, P. Rabl, and R. Blatt, "Ion-trap measurements of electric-field noise near surfaces," Rev. Mod. Phys. **87**, 1419–1482 (2015).

[158] V. Andreev, *Polarimetry on the Advanced Cold Molecule Electron Electric Dipole Moment Experiment*, Master's thesis, Technische Universität München (2016).

[159] V. Wirthl, C. D. Panda, P. W. Hess, and G. Gabrielse, "Simple self-calibrating polarimeter for measuring the Stokes parameters of light," OSA Continuum **4**, 2949–2969 (2021).

[160] *Corning HPFS 7979, 7980, 8655 Fused Silica Optical Materials Product Information*, Corning Inc. (2015).

[161] 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).

[162] G. J. Feldman and R. D. Cousins, "Unified approach to the classical statistical analysis of small signals," Physical Review D: Particles and Fields **57**, 3873–3889 (1998).

[163] T. Chupp and M. Ramsey-Musolf, "Electric dipole moments: A global analysis," Phys. Rev. C **91**, 035502 (2015).

[164] V. V. Flambaum, I. B. Samsonov, and H. B. Tran Tan, "Limits on CP-violating hadronic interactions and proton EDM from paramagnetic molecules," Journal of High Energy Physics **2020**, 77 (2020).

[165] B. K. Sahoo, "Improved limits on the hadronic and semihadronic cp-violating parameters and role of a dark force carrier in the electric dipole moment of ^{199}Hg," Phys. Rev. D **95**, 013002 (2017).

[166] D. G. Ang, C. Meisenhelder, C. D. Panda, X. Wu, D. DeMille, J. M. Doyle, and G. Gabrielse, "Measurement of the $H^3\Delta_1$ radiative lifetime in ThO," Physical Review A: Atomic, Molecular, and Optical Physics **106**, 022808 (2022).

[167] 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," Opt. Express **31**, 1943–1957 (2023).

[168] 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).

[169] J. J. Everdij, A. Huijser, and N. F. Verster, "Improved space focusing of polar diatomic molecules in a system of quadrupole and hexapole fields," Review of Scientific Instruments **44**, 721–725 (1973).

[170] V. A. Cho and R. B. Bernstein, "Tight focusing of beams of polar polyatomic molecules via the electrostatic hexapole lens," The Journal of Physical Chemistry **95**, 8129–8136 (1991).

[171] 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), arXiv:1607.01389 .

[172] 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**, 16914–16926 (2021).

[173] R. Zhou, T. Liu, L. Zhang, J. Hu, L. Zhu, A. Zeng, and H. Huang, "Strict mathematical model of mechanical stress birefringence for optical plates," Applied Optics **60**, 9117–9121 (2021).

[174] B. E. Kruschwitz, J. H. Kelly, I. M. J. Shoup, L. J. Waxer, E. C. Cost, E. T. Green, Z. M. Hoyt, J. Taniguchi, and T. W. Walker, "High-contrast plasma-electrode Pockels cell," Applied Optics **46**, 1326–1332 (2007).

[175] 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 and Electro Optics and the Pacific Rim Conference on Lasers and Electro-Optics* (2009) pp. 1–2.

[176] J. A. Salem, N. A. Smith, and A. Ersahin, *Mechanical Properties of a High Lead Glass Used in the Mars Organic Molecule Analyzer*, Tech. Rep. (NASA, 2015).

[177] H. Pauly, *Atom, Molecule, and Cluster Beams II: Cluster Beams, Fast and Slow Beams, Accessory Equipment and Applications*, Atom, Molecule, and Cluster Beams (Springer, 2000).

[178] Y. S. Au, *Inelastic Collisions of Atomic Thorium and Molecular Thorium Monoxide with Cold Helium-3*, Ph.D. dissertation, Harvard University (2014).

[179] 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), https://doi.org/10.1063/1.1149573 .

[180] 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).

[181] L. Lyons, *Statistics for Nuclear and Particle Physicists* (Cambridge University Press, 1989).

[182] CVI, "LAG: Aspheric Glass Condenser Lenses," https://web.archive.org/web/20200807173449/https://www.cvilaseroptics.com/products/spherical-lenses/aspheric-glass-condenser-lenses/lag (2020), accessed: 2023-1-24.

[183] L. Jones, "Reflective and Catadioptric Objectives," in *Geometrical and Physical Optics, Polarized Light, Components and Instruments*, Handbook of Optics, Vol. 1, edited by M. Bass, C. DeCusatis, J. Enoch, V. Lakshminarayanan, G. Li, C. MacDonald, V. Mahajan, and E. Van Stryland (McGraw Hill LLC, 2009) 3rd ed., pp. 29.1–29.41.

[184] *Photomultiplier tubes: R7600U Series*, Hamamatsu Photonics (2019), accessed: 2023-1-26.

[185] S. Gundacker and A. Heering, "The silicon photomultiplier: fundamentals and applications of a modern solid-state photon detector," Physics in Medicine & Biology **65**, 17TR01 (2020).

[186] *Multi-Pixel Photon Counter arrays: S13361-6050 series*, Hamamatsu Photonics (2022), accessed: 2023-1-26.

[187] W. C. J. Hunter, R. S. Miyaoka, L. R. MacDonald, and T. K. Lewellen, "Measured temperature dependence of scintillation camera signals read out by geiger-Müller mode avalanche photodiodes," in *2009 IEEE Nuclear Science Symposium Conference Record (NSS/MIC)* (2009) pp. 2662–2665.

[188] M. R. 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).

[189] A. Gola, C. Piemonte, and A. Tarolli, "Analog circuit for timing measurements with large area SiPMs coupled to LYSO crystals," IEEE Transactions on Nuclear Science **60**, 1296–1302 (2013).

[190] *DG645 Digital Delay Generator User Manual*, Stanford Research Systems (2008).

[191] NI, "PXIe-5171 Specifications," https://www.ni.com/docs/en-US/bundle/pxie-5171-specs/page/specs.html (2021), accessed: 2023-1-30.

[192] *Producer/Consumer Architecture in LabVIEW*, National Instruments Corporation (2022).

[193] T. J. Sumner, J. M. Pendlebury, and K. F. Smith, "Convectional magnetic shielding," Journal of Physics D: Applied Physics **20**, 1095 (1987).

[194] Twinleaf, "VMR Magnetoresistive Vector Magnetometer," https://twinleaf.com/vector/VMR/ (2023), accessed: 2023-2-8.

[195] J. Mitchell, *Magnetic Field Measurement System for the ACME III Electron EDM Search*, Master's thesis, Northwestern University (2022).

[196] J. Dupont-Roc, S. Haroche, and C. Cohen-Tannoudji, "Detection of very weak magnetic fields (10^{-9} gauss) by ^{87}Rb zero-field level crossing resonances," Physics Letters A **28**, 638–639 (1969).

[197] QuSpin, "Introduction to Zero-Field Magnetometer," http://quspin.com/products-qzfm-gen2-arxiv/zero-field-magnetometer-description/ (2023), accessed: 2023-2-8.

[198] QuSpin, "QZFM Gen-3," https://quspin.com/products-qzfm/ (2023), accessed: 2023-2-8.

[199] C. Crawford and Y. Shin, *A method for designing coils with arbitrary fields*, Tech. Rep. (University of Kentucky, 2009).

[200] E. Martin and C. Crawford, *A Double Cosine Theta Coil Prototype*, Tech. Rep. (University of Kentucky, 2010).

[201] M. Rawlik, A. Eggenberger, J. Krempel, C. Crawford, K. Kirch, F. M. Piegsa, and G. Quéméner, "A simple method of coil design," American Journal of Physics **86**, 602–608 (2018).

[202] H. Pauly, *Atom, Molecule, and Cluster Beams I: Basic Theory, Production and Detection of Thermal Energy Beams*, Atom, Molecule, and Cluster Beams (Springer, 2000).

[203] P. Hobbs, *Building Electro-Optical Systems: Making It all Work*, Wiley Series in Pure and Applied Optics (Wiley, 2011).

[204] K. Shimizu and F. Shimizu, "Laser induced fluorescence spectra of the a $a^3\Pi_u - X^1\Sigma_g^+$ band of Na_2 by molecular beam," The Journal of Chemical Physics **78**, 1126–1131 (1983).

[205] Z. Lasner and D. DeMille, "Statistical sensitivity of phase measurements via laser-induced fluorescence with optical cycling detection," Phys. Rev. A **98**, 053823 (2018).

[206] H. Metcalf, P. Van der Straten, J. Birman, J. Lynn, and H. Stanley, *Laser Cooling and Trapping*, Graduate texts in contemporary physics (Springer, 1999).

Colophon

This thesis was typeset using LaTeX, originally developed by Leslie Lamport and based on Donald Knuth's TeX. The formatting for this thesis is adapted from a template released under the permissive MIT (X11) license that be found online at github.com/suchow/ or from the author at suchow@post.harvard.edu.