A Thermochemical Cryogenic Buffer Gas Beam Source of ThO for Measuring the Electric Dipole Moment of the Electron

A DISSERTATION PRESENTED
BY
ELIZABETH PETRIK WEST
TO
THE DEPARTMENT OF PHYSICS

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

HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS
MAY 2017

Thesis advisor: John M. Doyle Elizabeth Petrik West

A Thermochemical Cryogenic Buffer Gas Beam Source of ThO for Measuring the Electric Dipole Moment of the Electron

Abstract

The discovery of an electric dipole moment of the electron (eEDM) within a few orders of magnitude of the current best limit would reveal the existence of time-reversal symmetry (T) violating physics beyond the Standard Model. Certain polar molecules with unpaired electron spins possess highly advantageous qualities for eEDM searches, and recent experiments using such species [10, 103] have pushed the frontier of eEDM searches into regimes of unprecedented sensitivity.

By performing a spin-precession measurement on a beam of thorium monoxide (ThO), the ACME collaboration has shown that the eEDM is less than $10^{-28}\,e\,\mathrm{cm}$, the most stringent upper bound to date [10, 11]. This null result severely constrains many theoretically proposed T-violating mechanisms, and many of the theories that remain viable predict eEDMs within one or two orders of magnitude below this bound.

In order to probe this tantalizing regime, we have developed a new cryogenic buffer gas beam (CBGB) source that exploits a high-temperature chemical reaction [56] between thorium and thorium dioxide to produce gas-phase ThO. For a single target over a single day, this source produces ThO fluxes nearly an order of magnitude higher on average than those produced by the laser-ablation-based CBGB used in our previous measurement [106]. Other beam properties, such as forward velocity, rotational temperature, and divergence have been measured and shown to be comparable to or only marginally less favorable than those of the ablation source. By enhancing the experiment's achievable count rate, this new thermochemical beam source could improve the statistical sensitivity of a future iteration of the ACME eEDM measurement by up to a factor of 2.5. In this thesis, I discuss the background of and motivation for the new beam source, its design, and its characterization and optimization for use in a future eEDM measurement.

Contents

Listing of Figures

To the Petrik and West families, with love and gratitude

Citations to Previous Work

Portions of this work or results described herein have previously been published in the references below. Where appropriate, citations and acknowledgements are also provided in the body of the thesis for derivative or collaborative work.

1. The ACME Collaboration: J. Baron, W. C. Campbell, D. DeMille, J. M. Doyle, G. Gabrielse, Y. V. Gurevich, P. W. Hess, N. R. Hutzler, E. Kirilov, I. Kozyryev, B. R. O'Leary, C. D. Panda, M. F. Parsons, B. Spaun, A. C. Vutha, A. D. West, and E. P. West, "Methods, analysis, and the treatment of systematic errors for the electron electric dipole moment search in thorium monoxide," submitted to *NJP*, January 2017. Preprint: ArXiv:1612.09318v1.

2. 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," *Phys. Rev. A* **93**, 052110 (2016).

3. The ACME Collaboration: J. Baron, W. C. Campbell, D. DeMille, J. M. Doyle, G. Gabrielse, Y. V. Gurevich, P. W. Hess, N. R. Hutzler, E. Kirilov, I. Kozyryev, B. R. O'Leary, C. D. Panda, M. F. Parsons, E. S. Petrik, B. Spaun, A. C. Vutha, and A. D. West, "Order of magnitude smaller limit on the electric dipole moment of the electron," *Science* **343**, pp. 269–272 (2014).

4. The ACME Collaboration: W. C. Campbell, C. Chan, D. DeMille, J. M. Doyle, G. Gabrielse, Y. V. Gurevich, P. W. Hess, N. R. Hutzler, E. Kirilov, B. O'Leary, E. S. Petrik, B. Spaun, and A. C. Vutha, "Advanced cold molecule electron EDM," *EPJ Web of Conferences* **57**, 02004 (2013).

5. 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," *Phys. Chem. Chem. Phys.* **13**, pp. 18976–18985 (2011).

6. 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," *J. Phys. B* **43**, 074007 (2010).

Acknowledgments

If Newton stood on the shoulders of giants, I was hauled across the finish line by them in a groggy stupor. Any merit in this work redounds quite to the credit of my incomparable colleagues, friends, and family.

Heartfelt thanks to my advisor, John Doyle, who has always been a patient and inspiring teacher and an example of scientific leadership I hope someday to follow. John also taught me everything I know about both cryogenic engineering and fine espresso and is not to blame if either subject remains beyond my grasp.

I also gratefully acknowledge the other ACME PI's, Dave DeMille and Jerry Gabrielse, without whose guidance and vision the ACME experiment would be inconceivable.

To my labmates, all the grad students and postdocs who have worked in the Doyle group or the ACME collaboration over the years, just...thank you for everything. I'm so blessed to have gotten to work with such a clever, talented, good-hearted, and downright *fun* group of people. All credit belongs to you guys for anything of value in these pages.

I owe—and am heartily glad to pay—particular thanks to Nick Hutzler and Jacob Baron. Nick developed the first ThO beam source, modeled and understood it so thoroughly that (luckily for me) following the trail he blazed required no skill at all, and unearthed the Cold

War era radioisotope chemistry research that inspired the current work. Jacob designed and built the thermochemical ThO beam source described in these pages and performed much of its early testing. It is no exaggeration to say that this project is as much theirs as mine.[1]

All of us in the physics department would be at sea without a fantastic group of administrators and technical staff. Special thanks to Jan Ragusa for patiently untangling me from various paperwork snarls and making lots of complicated things happen with no apparent (though I'm sure lots of unobserved) effort. Thanks as well to the wonderful machinists Stan Cotreau and Steve Sansone and electrical engineer Jim MacArthur. They were the ones who actually made many parts of the apparatus I describe so proudly in these pages—and kindly explained why the bits of it I built don't work.

My dear and steadfast friends, acknowledging your many years of loving support in both bright times and dark through the medium of a dusty thesis, doomed to languish unread in the Harvard depository, feels like cheapening the sentiment. To convey more fully the depth of my feeling, how about you buy me a beer so I can thank you in person? Two exceptions: Jon Bittner and Max Lavrentovich, without you, I would have been batty (well, batti*er*) by the end of year one. You guys can buy me two beers.

Last and emphatically most, I am grateful with all my heart to my families, both old and new. To my producers, Mom and Dad, thank you for being there for me from the very beginning and encouraging me not to "be intimidated by the eggheads" out of trying to do something worthwhile. And finally to Adam, my colleague, friend, and now my husband—a sweeter course I could never have hoped to chart.

[1] So you can blame them for its mistakes.

"He said science was going to discover the basic secret of life someday," the bartender put in. He scratched his head and frowned. "Didn't I read in the paper the other day where they'd finally found out what it was?"

—Kurt Vonnegut, *Cat's Cradle*

1 Background

1.1 Introduction

At accelerators such as the Large Hadron Collider (LHC), particles are smashed together with enormous energy in order to probe the most fundamental laws of physics.[1] In this way, particle physicists seek to produce and detect massive particles that have never before been observed, thus revealing entirely new facets of nature.

On a complementary front, the precise measurement techniques of atomic, molecular, and optical (AMO) physics can reveal the subtle effects of these exotic massive particles on ordinary matter. Even in a vacuum chamber in a perfectly controlled lab environment, atoms and molecules are never truly alone. Rather, they move through a rich background of quantum fields associated with every type of elementary particle—both known and unknown. Tiny interactions with such background fields can modify the basic atomic and molecular

[1] Parts of this chapter were adapted from [35], which is licensed under Creative Commons. Many thanks to the other members of the ACME collaboration for their contributions and editing assistance.

properties we measure in the lab. One famous example of this is the electron's anomalous magnetic dipole moment (MDM). The discovery that the electron MDM is not exactly 1 Bohr magneton, but about a part in a thousand larger, represented an early and dramatic vindication of quantum field theory [197]. Today, the anomalous MDM of the electron has been measured and calculated to the 11^{th} nonzero decimal place, making it the most precisely experimentally verified theoretical result in the history of physics [14, 15, 84, 96].

Permanent electric dipole moments (EDMs), or lopsided charge distributions, are another experimentally accessible property of ordinary particles with potentially profound implications. To date, no nonzero EDM of a fundamental particle has been observed, though physicists have been searching for over 60 years. As discussed in Section 1.2, the reason for this failure is a sound one: the Standard Model of particle physics (SM), a fantastically successful theory describing all known particles and interactions, has symmetries that strongly suppress EDMs, leading to EDM predictions that are in some cases ten orders of magnitude below the reach of current experiments [115, 153]. Despite this fact, and the decades of null results, the reason for continuing the search is also an excellent one: namely, that the SM is widely held to be incomplete. Among other shortcomings, it does not appear to account for the observation of dark matter, the invisible source of mass comprising for most of the bulk of galaxies; dark energy, believed to be responsible for the accelerating expansion of the universe; or the preponderance of matter over antimatter in the observable universe. A variety of unconfirmed theories beyond the SM purport to explain such phenomena—and as a side effect, they also tend to predict EDMs that are within or near the reach of current experimental methods. These theories and their implications for EDMs, especially that of the electron, are briefly discussed in Section 1.2.2.

This thesis focuses on the efforts of the Advanced Cold Molecule EDM Experiment (ACME) [189] to measure the electron's electric dipole moment (eEDM) using thorium monoxide (ThO). Taking advantage of recent improvements in technologies and methods, including a slow, cold, and intense beam source (see Section 2.2 and [106]) and ThO's near-ideal $^3\Delta_1$

state structure (see Section 1.3.1 and e.g. [135, 183, 188, 189]), we have developed an experiment with unprecedented eEDM statistical sensitivity. In 2014, we published a new best limit of 10^{-28} $e\,\mathrm{cm}$ (90% C.L.)[2] on the magnitude of the eEDM [10], which represented an order-of-magnitude improvement upon the previous best limit [103]. The experimental setup and result are described in Chapter 2. After a few years of upgrades, described in Chapter 3, we are now poised to make a second-generation measurement of the eEDM with a shot-noise limited statistical sensitivity of $\approx 4 \times 10^{-30}$ $e\,\mathrm{cm}$ in one day of averaging time.

Whether or not Generation II of ACME observes an eEDM, there will be a strong motivation to improve the experimental sensitivity still further: If we see an EDM, we will be eager to repeat our result quickly and with improved statistics in order to confirm it. If the eEDM is still consistent with zero at the Generation II level, we will wish to extend our reach even deeper into the remaining viable parameter space of theories beyond the SM (see Section 1.2). We have therefore developed an improved buffer gas beam source of ThO, based on a high-temperature chemical reaction discussed in Chapter 4, that yields cold molecular beam intensities an order of magnitude larger that the current ablation-based source. The development and characterization of this source are the main subject of this thesis, and the apparatus and key results are described in Chapter 5.

1.2 EDM Theory

And Waldo saw that many Waldos had been this way before.

—Martin Handford, *Where's Waldo? The Fantastic Journey*

This subject has been ably covered in many places. For a detailed treatment, the reader is encouraged to consult many of the references in this chapter, especially: [21, 28, 47, 73, 115, 154]. Some useful overviews at an introductory level can be found in: [59, 80, 170] and also—honestly—Wikipedia. For just a few examples of good discussions of the relevant Standard Model and beyond-SM theory, see: [5, 138, 141]. Also, quite a few capable Ph.D.-

[2] By convention, the units used for the eEDM throughout this thesis are $e\,\mathrm{cm}$, or elementary charge $\times$ centimeters.

aspirant Waldos have covered this ground before me, including but not limited to: [91, 100, 104, 105, 143, 163, 178, 180, 188]. In this section I present my own variation on the theme of motivating EDM experiments, though readers experienced in this subject will easily recognize the much-dogeared sheet music from which I take my cues.

Permanent EDMs are interesting because they violate two fundamental discrete symmetries of nature, parity (P) and time-reversal (T), that are largely respected by the Standard Model. The observation of a nonzero EDM of a fundamental particle would reveal a breakdown of the SM.

Before proceeding, a quick aside on CPT invariance: The CPT theorem states that, given some well-motivated physical assumptions, all physical laws are invariant under the combination of charge-conjugation (C), P, and T transformations, even though each of these symmetries may be separately violated [115]. Thus, CP violation and T violation are thought to be equivalent, and the terms are often used interchangeably in the literature and in this thesis.

In the SM, only the weak force violates P and CP. Since the weak force is famously *weak*, its symmetry-violating effects generally do not show up in everyday life[3] and can only be detected by sensitive experiments. In fact, the mirror-symmetry of physical interactions was taken for granted until 1950, when Edward Purcell and Norman Ramsey pointed out that it was open to empirical challenge [158]. Six years later, the Wu experiment became the first to reveal a violation of P symmetry by detecting the preferential emission of β-particles in the direction opposite the spin of cobalt-60 nuclei [196].

CP violation proved more elusive still. In the SM, the only known source of CP violation is an order-unity complex phase in the Cabbibo-Kobayashi-Maskawa (CKM) matrix, which governs quark mixing. All instances of CP violation observed to date, such as that observed in neutral kaon decay [43], are attributable to the CKM matrix [115].[4]

[3] As far as we know now, at least! It is possible, if doubtful, that some life-sized broken symmetries, such as DNA chirality, are ultimately rooted in fundamental symmetry violations [159].

[4] It should be noted that while the CKM matrix accounts phenomenologically for CP violation in quark

As shown in Section 1.2.1, the existence of EDMs requires a mechanism for CP violation; however, because of various suppression mechanisms (see e.g. Khriplovich and Lamoreau, Chapter 5.2 [115]), the CKM matrix produces EDMs that are extremely minute. For the neutron, the SM EDM prediction is [115]:

$$d_n \sim 10^{-31}\, e\,\text{cm}. \tag{1.1}$$

Compare this to the current best experimental limit of $d_n < 2.9 \times 10^{-26}\, e\,\text{cm}$ (90% C.L.), set by a group at the Institut Laue-Langevin in 2006 [18]. For the electron, the SM EDM is suppressed still further by additional cancellations and the fact that the CKM matrix couples only indirectly to leptons. Based on the fact that all contributions to the SM eEDM vanish up to 3-loop order, it has been shown that the theoretical upper bound on the eEDM is [153][5]:

$$d_e < 10^{-38}\, e\,\text{cm}. \tag{1.2}$$

More recent estimates suggest that the SM eEDM is not just smaller but *much* smaller than this limit, as the four-loop contributions have magnitudes of $\sim 10^{-44}\, e\,\text{cm}$ [155]. For comparison, the current best experimental limit, set by the ACME collaboration in the experiment described in Chapter 2, is $d_e < 9.4 \times 10^{-29}\, e\,\text{cm}$ (90% C.L.) [10, 60, 174].

At first blush, it might seem strange to build experiments to look for an effect that is predicted by the Standard Model—one of the most successful theories in history—to be many orders of magnitude below the level of our current sensitivity. However, as will be discussed in Section 1.2.2, there is good reason to believe that the SM breaks down and new physics emerges in such a way that EDMs are greatly enhanced. In order to understand

mixing, it is not really an explanation of the effect, and it is hoped that we will eventually discover a deeper underlying physical mechanism.

[5] If eEDM experiments in atoms and molecules ever did become ten billion times more sensitive than the current state of the art, they would actually encounter a different CKM-matrix-induced effect, a CP-violating electron-nucleon interaction equivalent to an eEDM of order $10^{-38}\, e\,\text{cm}$, before they reached the precision needed to observe the SM eEDM [44, 155].

why, we must first explore why EDMs violate P and T in the first place.

1.2.1 EDMs and Fundamental Discrete Symmetries

As illustrated in Fig. 1.2.1, permanent EDMs separately violate P and T symmetries by choosing a preferred direction relative to the particle's angular momentum under mirror-reflection and time-reversal operations [80, 115, 158].

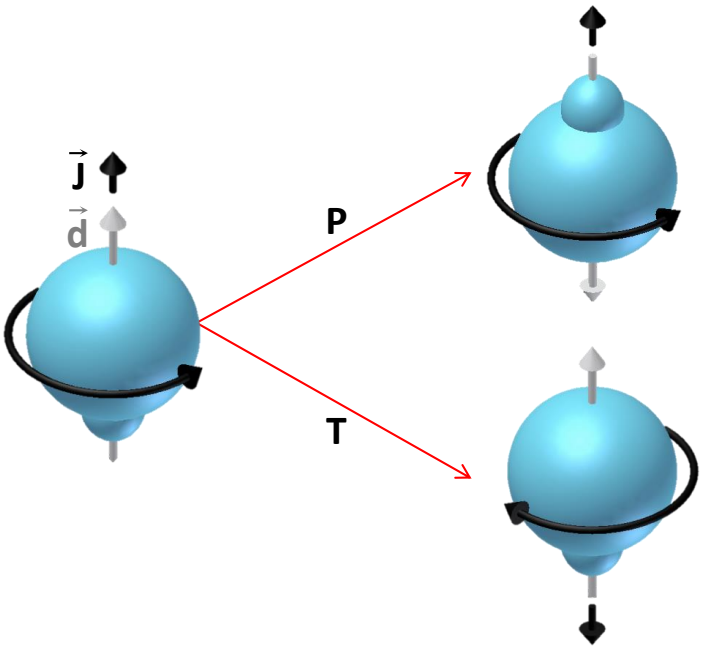

Figure 1.2.1: Behavior of an EDM under P and T. Under P, the electric dipole moment $\vec{d}_e$ (depicted as a lump of excess charge on the bottom of the particle) reverses, while the spin $\vec{J}$ remains invariant. Under T, $\vec{J}$ reverses while $\vec{d}_e$ remains invariant. The particles resulting from the P and T operations are empirically distinguishable from the original, and as discussed in the text, only one of these configurations is produced in nature. Thus, a particle with a nonzero EDM picks out a preferred mirror image and time-direction. After a rotation of 180°, the P and T transformations are identical; therefore, assuming Lorentz invariance, the P and T operations undo each other, and PT invariance is preserved.

To prove this claim more formally, we begin by invoking the Wigner-Eckart theorem (See e.g. Budker, Kimball, and DeMille, Appendix F [34]) to assert that a permanent EDM $\vec{d}$ of any system must be parallel or anti-parallel to the system's angular momentum $\vec{J}$—the spin, in the case of a fundamental particle. The classical interpretation of this fact is that any

vector components orthogonal to $\vec{J}$ would average to zero on the timescale of a rotational period.

If we try to evade the $\vec{d} \propto \vec{J}$ requirement by supposing either that the charge distribution producing $\vec{d}$ is independent of (and can change orientation with respect to) the spin $\vec{J}$, or that different particles of the same species have different (but fixed) relative orientations of $\vec{d}$ and $\vec{J}$, we run into another problem: These assumptions would introduce the extra dipole projection quantum number M_d, independent of the angular momentum projection M_J, into the system, which is contrary to experimental observations of atomic and molecular systems. For example, M_{d_e} would allow more than two electrons to occupy a single atomic orbital under the constraints of the Pauli exclusion principle. Although such behavior is in principle possible, the fact that the fields of atomic physics, nuclear physics, and chemistry have never observed such an extra degree of freedom, despite many fantastically accurate atomic and molecular structure measurements, lends strong support to the claim that a nonzero $\vec{d} \cdot \vec{J}$ must always have the same sign for a given particle type. Thus, a permanent EDM can be written as $\vec{d} = \beta \vec{J}$.

Configurations of objects, like snails, hands, and galaxies, often differ from their mirror images without necessarily arising from violations of fundamental symmetries of physics. To ensure that we are considering an explicit symmetry *violation*, as opposed to a spontaneous or *de facto* symmetry breaking, it is prudent to examine not just whether a system *looks* asymmetric, but whether the fundamental *laws* the system obeys exhibit a temporal or spatial asymmetry [32].

We therefore consider the electric dipole Hamiltonian

$$H_d \equiv -\vec{d} \cdot \vec{E} = -\beta \vec{J} \cdot \vec{E}, \tag{1.3}$$

where $\vec{E}$ is the electric field. This interaction manifestly violates P and T: $\vec{J}$ switches sign under the T transformation but not under P, and $\vec{E}$ switches sign under P but not T.

Thus, we have $PH_dP^\dagger = TH_dT^\dagger = -H_d$. Consequently, if the expectation value of H_d is nonzero for some physical system—i.e., the system possesses a permanent EDM—P and T are violated.[6]

An empirical consequence of these symmetry violations is that if we had two EDM experiments running side-by-side, and the second experiment had all fields parity-inverted or all momenta (such as currents) reversed with respect to the first, the measured dipole energy shift $\langle H_d \rangle$ would have opposite signs in the two experimental configurations. As discussed in Chapter 2, our ability to distinguish the signal due to the eEDM (and other potential P- and T-violating physics) from other effects relies on this unusual behavior under experimental switches that simulate reversals of parity and time.

Before moving on from symmetries, let's briefly consider the third fundamental discrete symmetry of nature: charge conjugation or C. Since C reverses both charge and lepton number, changing e.g. an electron into a positron, I am not aware of a simple picture for the behavior of the dipole moment under C. Nevertheless, there are powerful reasons (namely the CPT theorem, mentioned in Section 1.2) for believing that the combination of C, P, and T is a good symmetry of nature [115]. For a particle with an EDM, the P and T transformations undo each other (see Fig. 1.2.1); therefore, if nature is CPT invariant, then C is conserved. Thus, we often re-frame the T violation implied by EDMs as CP violation in order to highlight the connection to the matter-antimatter asymmetry. It is worth noting, however, that an EDM could arise from CPT violation rather than CP violation [31, 154].

[6] It's worth re-emphasizing that we are concerned only with *permanent* EDMs, i.e. those for which the dipole operator takes the form in Eq. (1.3) in the limit $|\vec{E}| \to 0$. For composite particles such as atoms and molecules, nonzero E-fields can give rise to *induced* dipole moments that do not have an intrinsically fixed orientation relative to $\vec{J}$. (See, for example, Appendix A for a discussion of induced EDMs in rigid rotor molecules.) In such cases the dipole operator takes the form $H_d^{\text{induced}} = \frac{1}{2}\alpha E^2$ for small $|\vec{E}|$, where the polarizability α is a constant of proportionality between $\vec{d}$ and $\vec{E}$. Unlike the permanent EDM case, H_d^{induced} is P and T invariant.

1.2.2 Connection to the Matter-Antimatter Asymmetry

"You're very sure of your facts," he said at last, "I couldn't trust
the thinking of a man who takes the Universe—if there is
one—for granted."

—Douglas Adams, *The Restaurant at the End of the Universe*

A notable—and fortunate for human safety—feature of the observable universe is its apparent lack of antimatter. If antimatter were intermingled with matter on the planetary scale, constant annihilation events would blow the Earth apart. If some galaxies or regions of space were composed of antimatter, we would expect to observe γ-ray emissions at 511 keV from electron-positron annihilations whenever a dust cloud made of ordinary matter drifted into one. In the absence of such observations, it appears that our universe is one in which matter is overwhelmingly predominant over antimatter.[7]

This apparent matter excess is one of the great outstanding mysteries in physics. The innate symmetry between matter and antimatter suggests that nature should have treated them practically identically: The physical processes involved in the Big Bang should have produced nearly equal numbers of particles and antiparticles, and nearly all of these should have rapidly annihilated, producing a universe filled mostly with photons and equal amounts of residual matter and antimatter.[8] Instead, one matter particle in a billion survived the explosive annihilation epoch in the first nanosecond of the cosmos's existence, and that one-billionth excess produced the matter universe we inhabit today [170].

[7] It is still possible, i.e. consistent with observational evidence, that some regions of the universe *are* composed of antimatter, but they are separated from our matter cluster by such great distances that we have so far been unable to detect annihilations from them. However, given the difficulty (though see e.g. [63]) of explaining how the matter and antimatter domains could come to be so widely separated, this possibility is disfavored [36, 92].

[8] One possible solution to the matter-antimatter asymmetry problem is that matter predominance is simply a contingent fact of nature: There happened to be an initial excess of matter in the Big Bang, and it persists to this day. This possibility has not been ruled out, but it does not seem particularly scientifically fruitful since it forecloses the question *why?* Moreover, in such a scenario, one would likely have to explain why any matter- and anti-matter-producing interactions plus thermal effects did not quickly equalize this initial excess, a result evaded in scenarios obeying Sakharov's third condition [150]. Also, if the current best models of inflationary cosmology are correct, an initial matter-antimatter asymmetry would probably not have survived without being completely washed out by the rapid expansion of the early universe [109].

In 1967, Andrei Sakharov proposed that any physical mechanism capable of producing an excess of matter from matter-antimatter-balanced conditions in the early universe must have the following three (necessary but not sufficient) features, now called the "Sakharov conditions" [166]:

1. Baryon number violation
2. C and CP violation
3. Out-of-equilibrium interactions

We will now attempt to understand these conditions and their relationship to EDM searches:

(1) In the current epoch of the universe, baryon number is conserved in the SM, meaning that any process producing baryons produces an equal number of anti-baryons, and vice versa. (The same is true for lepton number, but solutions to the matter-antimatter asymmetry problem tend to focus on baryons such as protons and neutrons because they comprise the bulk of the ordinary matter in the universe.) The first Sakharov condition states that in order to produce an excess of one over the other, some process or processes must violate this conservation law. The framework of the SM allows for such baryon-number-violating processes in the incredibly high temperatures that existed in the first hundred picoseconds after the Big Bang [36, 170].

(2) The second Sakharov condition provides that these processes should favor *matter*. Since matter and antimatter are related by C, and matter of one chirality is related to antimatter of the opposite chirality by CP, both of these symmetries must be violated in order for baryon-producing mechanisms to be favored over antibaryon-producing mechanisms [170]. The Standard Model provides insufficient sources of CP violation to allow for the observed matter-antimatter asymmetry of the universe [36]. This is the key point for EDM experiments. Section 1.2.1 showed that a measurable permanent EDM would violate CP to a degree not explained by the SM, and the second Sakharov condition calls for new

mechanisms of CP violation beyond the SM to explain the matter-antimatter asymmetry. Thus, the search for an EDM is a search for a well motivated type of new physics that could contribute to solving the mystery of our matter universe.

(3) For completeness, we touch on the third Sakharov condition: thermal non-equilibrium. This condition is required so that interactions obeying the first two conditions, which produce the matter-antimatter asymmetry, are permitted to happen at a higher rate than the inverse reactions, and so that the resulting asymmetry can be "frozen in" to the composition of the early universe instead of equilibrating away [36, 46, 75, 150, 170]. SM physics alone is believed to be insufficient to produce the necessary non-equilibrium conditions in the early universe [36, 52],[9] so Sakharov's third condition points to an additional feature we should hope to find in a successful SM extension or replacement.

1.2.3 EDMs in Beyond-SM Theories

As discussed at the beginning of Section 1.2, the SM largely respects T symmetry and therefore produces negligibly small EDMs. Only the quark flavor-mixing interaction involves a basis-independent complex phase capable of giving rise to T-violating effects. Such complex phases are associated with T violation because T is an antiunitary transformation, taking numbers to their complex conjugates as well as taking the time variable t to $-t$. (See [80, 105] for more-developed intuitive motivations of this claim.) SM extensions generically involve CP-violating complex phases [21, 154], and these could contribute to an explanation of the matter-antimatter asymmetry of the universe discussed in Section 1.2.2. Absent a physical reason for these phases to be suppressed, they are expected to be of order unity [73].

Understanding these models would require particle physics expertise far beyond this

[9] Strictly speaking, the early universe was *not* in equilibrium because it was expanding rapidly, but this expansion was slow compared to the rate of SM particle creation and annihilation at the $\sim 10^{15}$ K temperatures [170] required for baryogenesis [36], so the local conditions experienced by the particles were approximately thermal. The necessary non-equilibrium conditions *could* be produced in the SM by a first-order electroweak phase transition resulting in "bubbles" of broken electroweak symmetry (a "Higgs phase") in which the matter-antimatter asymmetry could be frozen into the early universe [36]. However, such a first-order phase transition would only be possible if the Higgs mass were below 72 GeV [52], whereas we now know it to be 125 GeV [1, 39].

writer's. However, we can use dimensional analysis to obtain a very crude estimate of the size of an EDM, d, produced by a "generic" n-loop Feynman diagram involving a new particle with rest-mass energy Λ_{New} coupling to a particle of mass m with complex phase ϕ_{CP} [10, 21, 28]:

$$\frac{d}{e} \sim \kappa \left(\frac{\alpha_{\text{eff}}}{4\pi}\right)^n \left(\frac{mc^2}{\Lambda_{\text{New}}^2}\right) \sin(\phi_{\text{CP}})\hbar c. \tag{1.4}$$

Here, every loop in the Feynman diagram suppresses the size of the EDM by a coupling constant $\alpha_{\text{eff}}/4\pi$; the need to create heavy virtual particles further suppresses the size of the effect by $m_e c^2/\Lambda_{\text{New}}^2$; $\sin(\phi_{\text{CP}})$ incorporates the CP-violating phase; κ encodes a complicated but $\mathcal{O}(1)$ function of the mass scales, which we use simply as a tuning factor here; and the remaining factors are required for dimensional consistency. An example of the type of Feynman diagram for which Eq. 1.4 is an estimate is shown in Fig. 1.2.2.

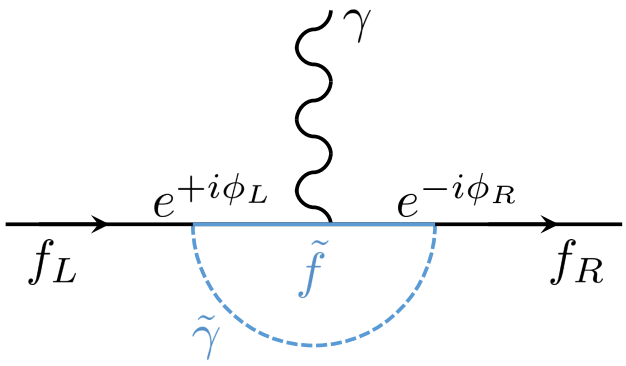

Figure 1.2.2: Example of a SUSY Feynman diagram capable of producing an EDM at the 1-loop level. A left-handed fermion f_L emits a virtual photino and sfermion, picking up a complex phase ϕ_L. The virtual sfermion interacts with an applied electric field represented by the virtual photon γ. Since SUSY interactions can violate T symmetry, the photino re-absorption process is not required to be the time-reversed version of the emission process. Thus, the phase picked up when the sfermion and photino recombine can be $\phi_R \neq \phi_L$, and the outgoing fermion state f_R can have opposite parity from the incoming state. The resulting asymmetry, which can be regarded as an asymmetry in the spatial charge distribution associated with f, permits the existence of an EDM. (Figure and explanation based on Fortson, Sandars, and Barr [80].)

There are strong reasons to suspect that the relevant new physics should arise at the electro-weak scale, i.e. with masses and coupling strengths comparable to those of the W

and Z bosons. These motivations include the "hierarchy problem," i.e. the mystery of why the Higgs boson is so light, and the "WIMP miracle," i.e. the coincidence between the electro-weak mass scale of ~ 100 GeV and our current best guess for the mass of dark matter.

One elegant class of theories that could potentially explain away the hierarchy problem, the WIMP miracle, and the matter-antimatter asymmetry is supersymmetry (a.k.a. SUSY). SUSY posits a symmetry between bosonic and fermionic matter and gives rise to an array of as-yet-undiscovered "superpartner" particles: one supersymmetric boson (called a "sfermion") for every type of fermion and one supersymmetric fermion ("bosino") for every boson. If, as has been hoped, SUSY particles have masses and couplings around the electro-weak scale, we can use Eq. 1.4 to estimate the size of the EDMs they might produce. For the electron, for example, if we take $\sin(\phi_{\rm CP}) = \kappa = n = 1$ (by naturalness), $\alpha_{\rm eff} = 4/137$ (coupling constant from electro-weak physics), and $\Lambda_{\rm New} = 100$ GeV (mass scale of W and Z bosons), we compute:

$$d_e^{\rm MSSM} \sim 10^{-24}\, e\,{\rm cm}. \tag{1.5}$$

The so-called "minimal supersymmetric standard model" (MSSM) is hence in rather severe tension with current experimental EDM limits.[10] This apparent inconsistency is known as the "SUSY CP problem," and it can be evaded by allowing higher SUSY particle masses or "unnaturally" small CP phases, or by modifying the MSSM so EDMs are less readily produced.

This last option must be left to the professionals, and a couple of beyond-SM variations developed by theorists to accommodate current EDM limits are briefly discussed below. However, even for a QFT layman, it's simple enough to plot Eq. 1.4 to show the approximate the parameter space in $\Lambda_{\rm New}$ and ϕ_{CP} ruled out by the current and projected ACME eEDM limits. These constraints are shown in Fig. 1.2.3, using the values $\kappa = 0.1$ (for conservatism) and $\alpha_{\rm eff} = 4/137$. This estimate shows that for natural CP-violating phases, i.e.

[10] Electro-weak scale SUSY is also quite gravely threatened by the fact that the LHC, which has been probing energy scales well into the TeV range over the past few years, has yet to observe any supersymmetric particles.

$\sin(\phi_{CP}) \sim 1$, the ACME Gen. I eEDM limit of $9.4 \times 10^{-29}\, e\,\text{cm}$ rules out new physics at the ~ 1 TeV (~ 0.1 TeV) energy scale with eEDM-producing couplings at the 1-loop (2-loop) level. Conversely, if the energy scale is required to be $\Lambda_{\text{New}} \sim 100$ GeV, then the CP-violating phase must be $\sim 10^{-3}$ ($\sim 10^{-1}$) at the 1-loop (2-loop) level. The projected reach of ACME Gen. II is another factor of 10 in reduced eEDM uncertainty, which corresponds to a factor of $\sqrt{10} \sim 3$ along the Λ_{New} energy axis.

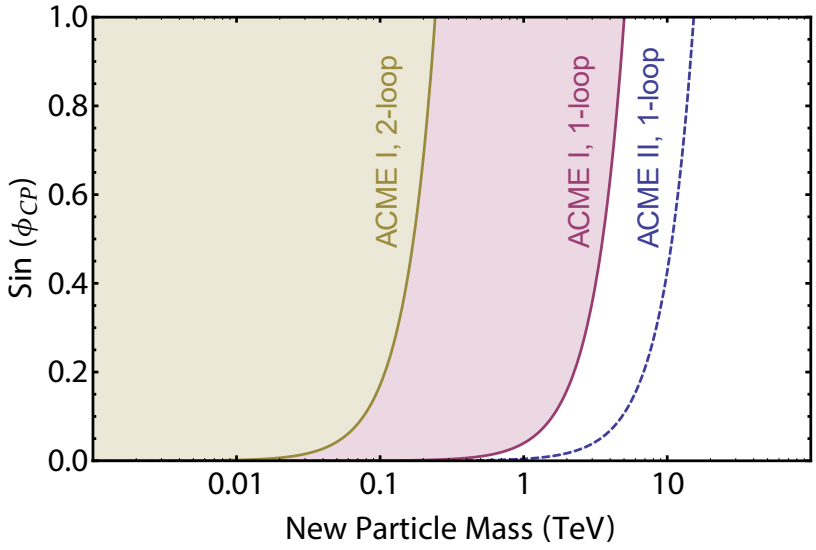

Figure 1.2.3: Estimated ACME eEDM limits on the mass (Λ_{New}) and CP-violating phase (ϕ_{CP}) of electro-weak scale new physics from a "generic" theory described by Equation 1.4 with parameter values given in the text. The region to the left of the purple curve is ruled out by the Gen. I ACME eEDM limit, assuming some new physics with electro-weak scale couplings produces an eEDM at the 1-loop level in perturbation theory. The region to the left of the gold curve is ruled out by ACME Gen. I assuming an eEDM arises at the 2-loop level. The blue dashed curve shows the region of 1-loop level new physics that ACME Gen. II is expected to probe.

If we relax the requirement of $\mathcal{O}(1)$ CP-violating phases, it is interesting to ask (see e.g. [20]) whether the bounds placed by current EDM limits on the amount of CP violation in the MSSM have ruled out electro-weak scale MSSM physics as the source of the matter-antimatter asymmetry of the universe. A recent paper by Li, Profumo, and Ramsey-Musolf [127] argues that the answer is "not quite." In a particular sector (bino–Higgsino) of the MSSM, the authors contend that a resonant enhancement of CP-violating effects can create

the necessary conditions for excess matter production while still producing EDMs small enough to be consistent with current experimental limits. Figure 1.2.4 shows the results of Li et al. for the remaining viable parameter space for matter-antimatter production from MSSM physics at the electro-weak scale. (The specific mass scales used for this benchmark are given in the paper: [127].) This calculation suggests that if an eEDM experiment 100 times as sensitive as ACME Gen. I fails to observe a nonzero eEDM, the electro-weak scale MSSM will no longer be a viable candidate for explaining the matter-antimatter asymmetry of the universe.

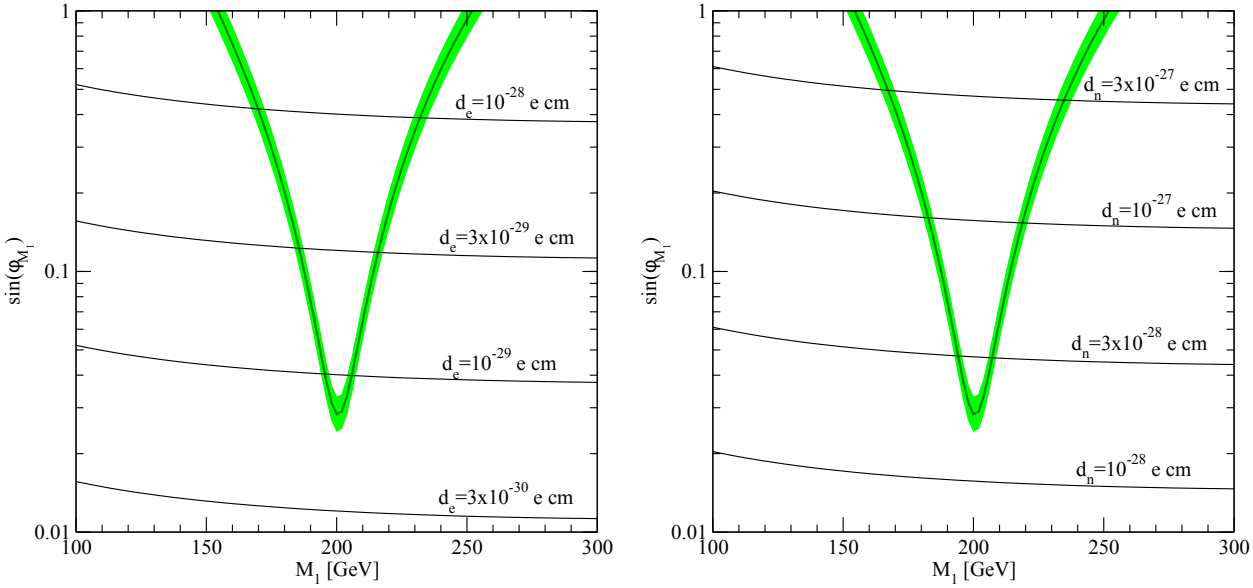

Figure 1.2.4: Viable parameter space for electro-weak baryogenesis resulting from MSSM physics and constraints from potential electron (left) and neutron (right) EDM limits. For reference, the current best eEDM limit is $9.4 \times 10^{-29}\, e\,\mathrm{cm}$ [10, 60, 174], and the current best nEDM limit is $2.9 \times 10^{-26}\, e\,\mathrm{cm}$ [18]. Figure from Li, Profumo, and Ramsey-Musolf [127], licensed under Creative Commons.

Other variations of SUSY might yet survive, however. One general extension to the MSSM, called beyond-MSSM or BMSSM, is motivated by the need to explain how the Higgs boson mass evades the upper bound set by the Z boson mass in the MSSM [138]. In order to do so, BMSSM incorporates two additional degrees of freedom in the Higgs sector, each of which has an associated CP-violating complex phase [37, 62, 138]. Nakai and Reece have recently

computed the constraints on this SUSY variant due to experimental limits on EDMs [138], and their results are shown in Fig. 1.2.5. They find that the ACME limit on the eEDM [10] and the recently improved University of Washington limit on the nuclear EDM of ^{199}Hg [87] enter at comparable levels and confirm that BMSSM CP-violating phases exhibit moderate (mild) fine-tuning for a CP-violating energy scale of 400 GeV (800 GeV).

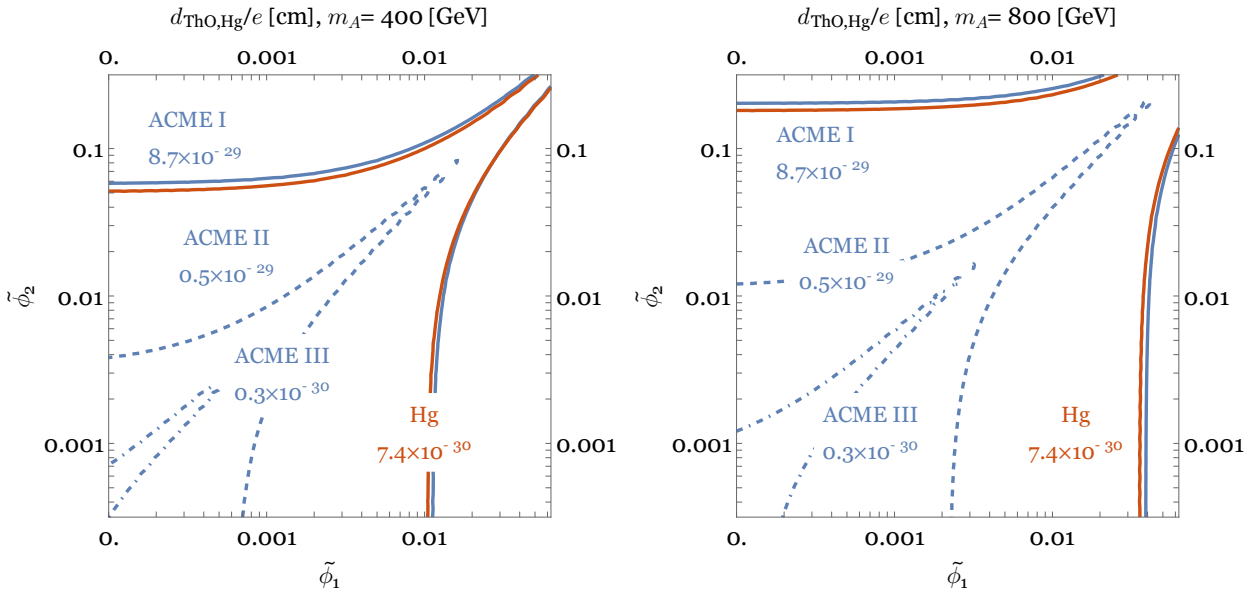

Figure 1.2.5: Electron [10] and ^{199}Hg [87] constraints on BMSSM CP-violating phases. Most of the parameter space with $\mathcal{O}(1)$ phases has been ruled out assuming the CP-violating BMSSM physics is associated with a mass scale $m_A = 400$ GeV, and future generations of ACME are expected either to reveal new physics or worsen (create) fine-tuning problems for the $m_A = 400$ GeV ($m_A = 800$ GeV) case. Figure from Nakai and Reece [138], used with the authors' permission.

The BMSSM model was chosen simply as an illustration of the power of EDM experiments in probing and narrowing the scope of the search for new physics—not because BMSSM is necessarily the best-motivated beyond-SM scenario or the one likeliest to be realized in nature. Nakai and Reece consider a variety of SUSY scenarios capable of solving the hierarchy problem in their paper [138], which is itself one among many papers exploring possible beyond-SM models in the context of existing theoretical motivations and experimental constraints (see e.g. [17, 119, 127] and the excellent recent SUSY review [78]).[11] Non-SUSY

[11] One other EDM SUSY constraint, that relating to the stop squark, is important enough to merit a brief

scenarios are also possible and are an area of active theoretical work (e.g. [73]).

A compelling—though outdated and potentially misleading—view of the constraints set by electron EDM experiments on physics beyond the SM is shown in Fig. 1.2.6, whose original version was created many years ago by David DeMille based on the work of Nir [141]. Some of the theories on this plot have since been ruled out by other considerations, some have evolved or have given birth to subclasses of successors, and most are too complex to produce simple, unconditional predictions for d_e. Nevertheless, under appropriate naturalness assumptions, it is often possible to estimate a likely range of values in which a given theory might be expected to produce an eEDM. Thus, with all the proper caveats, such plots as this can be valuable—if rough—indicators of what new physics we have explored so far and where we have left to look.

mention—even in a review as fragmentary as the present one. In many SUSY scenarios, a relatively light (electro-weak mass scale) stop squark (the supersymmetric partner to the top quark) is directly responsible for stabilizing the Higgs mass against quantum corrections, thereby solving the hierarchy problem. Nakai and Reece show that for most plausible regions of SUSY parameter space, given fairly natural CP-violating SUSY phases ($\phi_{CP} > 10^{-3}$), the ^{199}Hg EDM limit [87] sets a more stringent bound on the stop mass than the LHC, while the ACME eEDM limit [10] is competitive with the LHC [138]. (While the mercury EDM bound is found to be tighter than the electron EDM bound in this analysis, the eEDM limit is still important because of the relative simplicity of its interpretation. In particular, the ^{199}Hg nucleus is sensitive to more potential sources of CP violation, whose contributions could in principle cancel, reducing the experimental signal.) Given reasonable values for other SUSY parameters, both the ACME and the ^{199}Hg experiments are sensitive to stop masses up to $\approx$ 1–3 TeV, with the ^{199}Hg limit probing higher energy scales than ACME [138], whereas the LHC has probed stop masses up to $\approx$ 0.5 TeV [78].

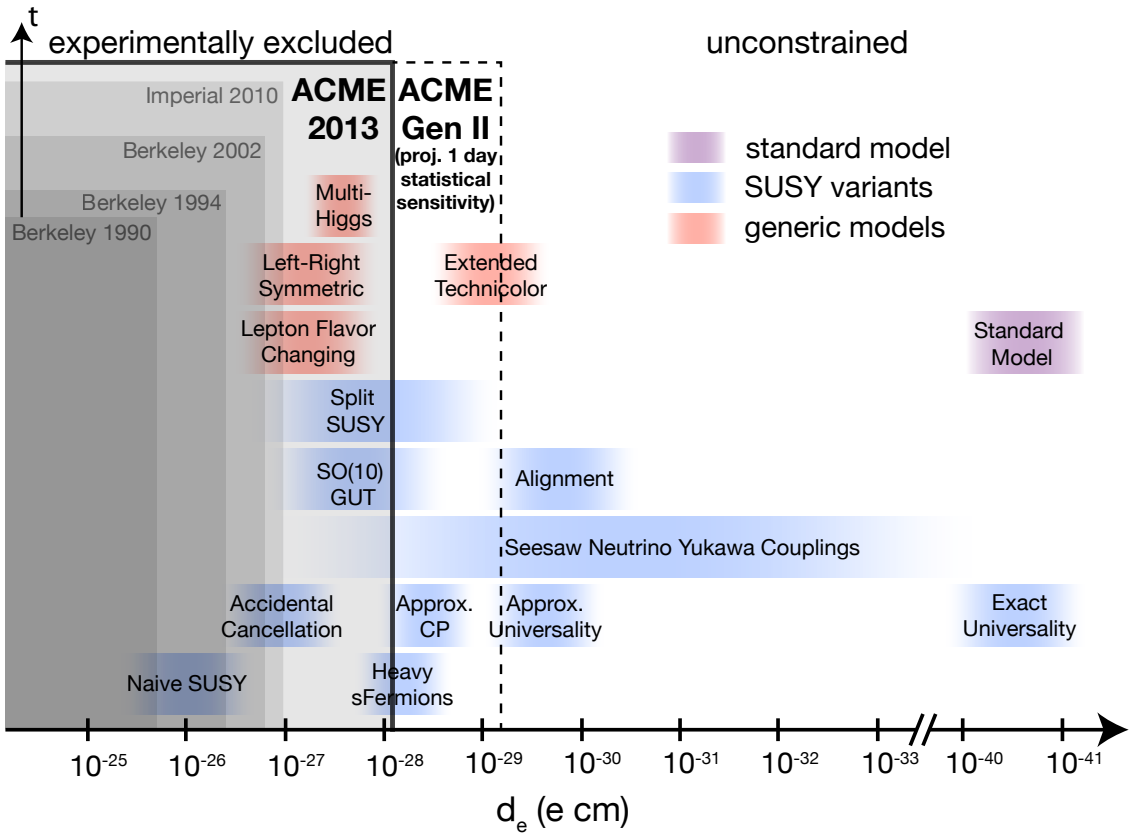

Figure 1.2.6: Estimated theoretical predictions and experimental limits on the eEDM. Limits published in: [4, 10, 48, 103, 162]. Figure originally by David DeMille based on results in [141]; restyled and updated with current experimental limits by Brendon O'Leary.

1.3 Atomic and Molecular eEDM Experiments

The subjects of EDM experiments have typically been free neutrons, diamagnetic atoms (which are primarily sensitive to CP-violating nuclear effects), and paramagnetic atoms and molecules, whose unpaired electron spins allow access to the EDM of the electron. The first two types of species are discussed extensively in the literature (see e.g. [18, 87, 115]) and will not be covered further here. The third, paramagnetic species, are the subject of the current work. In general, they are also the cleanest systems to interpret theoretically, as unlike neutrons or nuclei, electrons are point particles without internal structure.

As discussed in Section 1.2, the search for the eEDM is a sensitive probe of new physics, and this effort has long been at the forefront of fundamental physics research (see Fig. 1.3.1) [28, 115]. A high-precision measurement that discovers the eEDM or sets a stringent new limit on its size would place strong constraints on extensions to the SM. Many supersymmetric models already require fine tuning to fit the current EDM limits [5, 141]. An eEDM measurement that is 10–100 times as sensitive as the current upper bound must either observe an EDM, revealing a breakdown of the Standard Model, or set a new limit requiring such unnatural suppression of supersymmetric parameters that many supersymmetric models would have to be revised or rejected [154].

The signature of a permanent eEDM, d_e, is an energy shift U_{EDM} of an unpaired electron (or electrons) in an electric field $\vec{\mathcal{E}}$:

$$U_{\mathrm{EDM}} = -\vec{d}_e \cdot \vec{\mathcal{E}}. \tag{1.6}$$

In an eEDM measurement, the statistical uncertainty on $\vec{d}_e$ is therefore given by

$$\delta d_e = \frac{\delta U_{\mathrm{EDM}}}{\mathcal{E}}. \tag{1.7}$$

In the quantum projection noise limit, the energy-time uncertainty principle implies that for

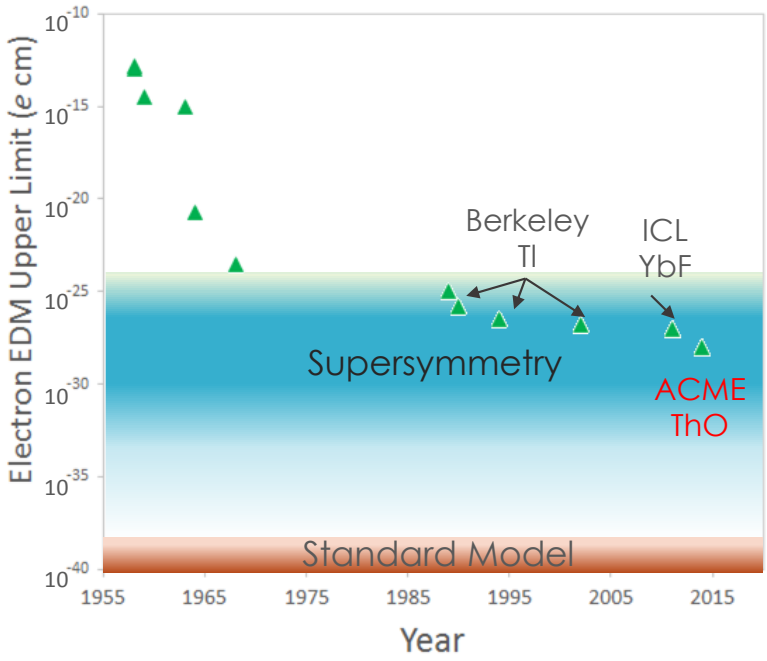

Figure 1.3.1: Best upper limits on the eEDM over time, since Ramsey and Purcell fired the starting pistol in 1950 [158]. The best limits have lately been set by experiments using heavy atomic and molecular species, with molecules taking a recent lead. The three experiments setting the most recent limits (see Table 1.3.1) are labeled on the plot. Approximate ranges of eEDM predictions from SUSY and the SM are indicated by colored bands. Based on an unpublished figure by Amar Vutha with data points compiled from: [4, 10, 48, 76, 85, 103, 137, 139, 162, 167, 169, 193].

a single measurement, the uncertainty is

$$\delta U_{\mathrm{EDM}} = \hbar/2\tau, \tag{1.8}$$

where τ is the "coherence time," or the interaction time between the electric field and the eEDM. For N independent measurements obeying Poisson statistics, the uncertainty scales as $1/\sqrt{N} = 1/\sqrt{\dot{N}T}$, where $\dot{N}$ is the count rate of independent measurements and T is the total averaging time. Substituting these results into Eq. 1.7 gives an expression for the statistical uncertainty of δd_e in terms of readily characterized experimental parameters:

$$\delta d_e = \frac{\hbar}{2\mathcal{E}\tau\sqrt{\dot{N}T}}. \tag{1.9}$$

This equation provides the basic figure of merit for any eEDM experiment: The experimental design goal is to maximize the electric field experienced by the electron, its interaction time with the field, and the measurement count rate.

In the vicinity of some atomic nuclei, electrons experience extremely strong electric fields [49, 167, 168], up to ~ 1 million times the size of fields achievable in the lab. These internal atomic and molecular fields can be partially or completely oriented by polarizing the atom or molecule, which together with relativistic effects gives the EDM of an unpaired electron a nonzero average energy shift (The shifts of *paired* electrons have opposite signs and cancel out). Per Eq. (1.6), this shift can be interpreted as an interaction between d_e and an average effective electric field $\mathcal{E}_{\text{eff}}$ produced by the atomic nucleus. The size of $\mathcal{E}_{\text{eff}}$ can be shown to scale approximately as the cube of the atomic number Z [34, 115]. Thus, the species that yield the most sensitive eEDM measurements (i.e. largest U_{EDM}) are heavy (large Z), highly polarizable, paramagnetic species whose valence electronic wavefunctions are in s-type orbitals with a large amplitude overlapping the nucleus. The $1/\sqrt{\dot{N}}$ statistical factor in Eq. 1.9 indicates that the species should be producible in large numbers, while the polarizability requirement recommends species with enough degrees of freedom to provide closely spaced levels of opposite parity that can be mixed in moderate applied electric fields. The former consideration generally favors atoms, whose simple structure makes them simple to create, while the latter consideration favors molecules, whose more complex structure affords the degrees of freedom that enhance polarizability.

These principles have guided the search for the eEDM for the last fifty years, during which time the strongest limits have consistently been set by atomic and molecular experiments (see Fig. 1.3.1). Table 1.3.1 summarizes the three most recent EDM upper bounds, obtained with atomic thallium (Tl) and the polar molecules ytterbium fluoride (YbF) and thorium monoxide (ThO).

Table 1.3.1: Comparison of the shot-noise limited statistical sensitivity of the three experiments placing the strongest limits on the eEDM. The per-day statistical sensitivity estimates are computed using the formula $\delta d_e \sqrt{T} = \hbar/(2\mathcal{C}\mathcal{E}_{\text{eff}}\tau\sqrt{\dot{N}})$, where T is the averaging time of the experiment in units of days; $\mathcal{C}$ is the "fringe contrast" discussed in Chapter 2, which accounts for technical limitations on an experiment's ability to resolve the outcome of a given measurement; $\mathcal{E}_{\text{eff}}$ is the effective electric field experienced by the electron, and $\dot{N}$ is the experimental count rate (including the in-run duty cycle, where known). The sensitivity estimate assumes that each experiment collects EDM data at its typical count rate for a full 24-hour day (This is contrary to fact: The Tl and ThO experiments, at least, had a running schedule duty cycle closer to $\approx 50\%$). In practice, technical noise can increase an experiment's statistical uncertainty over the shot-noise limit: The Tl experiment's ultimate statistical uncertainty was estimated to exceed the shot-noise limit by a factor of 1.7 [162], the YbF experiment's by a factor of 1.04 [178], and the ThO experiment's by a factor of 1.2 [11, 118].

Experiment	Commins et al. 2002	Hinds et al. 2011	ACME 2014
Species	Tl	YbF	ThO
References	[64, 128, 152, 162, 163]	[103, 123, 178]	[10, 60, 174]
$\mathcal{E}_{\text{eff}}$ (V/cm)	7.0×10^7	1.4×10^{10}	7.8×10^{10}
τ (ms)	2.4	0.642	1.1
$\dot{N}$ (s^{-1})	4.7×10^8	1.2×10^4	3×10^4
$\mathcal{C}$	0.60	0.64	0.94
Shot-noise-limited statistical sensitivity ($e\,\text{cm}\sqrt{\text{day}}$)	5.1×10^{-28}	1.7×10^{-27}	8.1×10^{-29}
Reported eEDM upper limit ($e\,\text{cm}$, 90% C.L.)	1.6×10^{-27}	1.1×10^{-27}	9.4×10^{-29}

1.3.1 Thorium Monoxide eEDM

No discussion of ACME would be complete without an encomium on its molecule of choice, ThO, which combines the aforementioned benefits of a high-Z, polar molecule with several other powerful advantages. These properties of ThO conspire to increase ACME's statistical sensitivity compared to previous eEDM experiments, mitigate the technical demands of working with molecules rather than atoms, and suppress or rule out many systematic errors [189].

ThO is a polar molecule with two valence electrons. In the $H\,^3\Delta_1$ state [134], one of these

electrons occupies a σ-orbital, and its EDM is relativistically enhanced due to the Sandars effect [49], while the other valence electron occupies a δ-orbital and allows the molecule to be easily polarized. The σ-state electron interacts with approximately 20 full atomic units of effective electric field ($\sim$ 100 GV/cm) in a molecular state that can be oriented with very modest laboratory fields ($\sim$ 1 V/cm) [190]. The interaction of this effective molecular field with a nonzero eEDM would manifest itself as a phase shift in ACME's Ramsey-type measurement protocol.

The effective internal electric field of fully polarized ThO is 77.6 GV/cm with an uncertainty of $< 10\%$ [60, 61, 79, 134, 174–177],[12] which is among the largest of any investigated species. This field is 3 times as large as the estimated field in fully polarized YbF [136], approximately 5.6 times as large as the $\mathcal{E}_{\text{eff}}$ achieved in partially polarized YbF in the Hinds experiment [171], and more than 1000 times larger than the $\mathcal{E}_{\text{eff}}$ achieved in the Tl experiment [162]. Moreover, ThO possesses a low-lying metastable state $H\,^3\Delta_1$ (see Fig. 1.3.2), which exhibits several features beneficial to an EDM experiment. Firstly, it has a measured lifetime of 1.8 ms [189], sufficient to perform our Ramsey experiment in a molecular beam with a coherence time of 1.1 ms (see Section 2.1). This is comparable to the coherence times in both the YbF (642 µs [103]) and the Tl ($\sim$2.4 ms [163]) eEDM experiments. Secondly, the spin and orbital magnetic moments of a state with $^3\Delta_1$ angular momentum cancel almost perfectly [189], and the residual g-factor is measured to be $g_{H,J=1} = -4.40(5) \times 10^{-3}$ [118, 151, 190].[13] This small magnetic moment renders the experiment highly insensitive to magnetic field imperfections.

[12]The value used for $\mathcal{E}_{\text{eff}}$ here and throughout this thesis is the unweighted mean of the two most recent values: 75.2 GV/cm, computed by Denis and Fleig [60], and 79.9 GV/cm, computed by Skripnikov [174]. These authors use different approaches for the calculation, and both report uncertainties at approximately the 5% level. The other references are to older calculations (of particular note is the 2008 paper by Meyer and Bohn [134] which confirmed at a critical stage of ACME's development that ThO was a viable choice of molecule) and technical but important discussions among the theorists about the validity of their respective approaches and error bar estimates.

[13]To avoid confusion with other, similar definitions of the molecular g-factor, we specify that in the present work's notation, the energy shift of a Zeeman sublevel of H, $J = 1$ in an applied magnetic field is given by $U_B = -g_{H,J=1}\mu_B \vec{J} \cdot \vec{B}$.

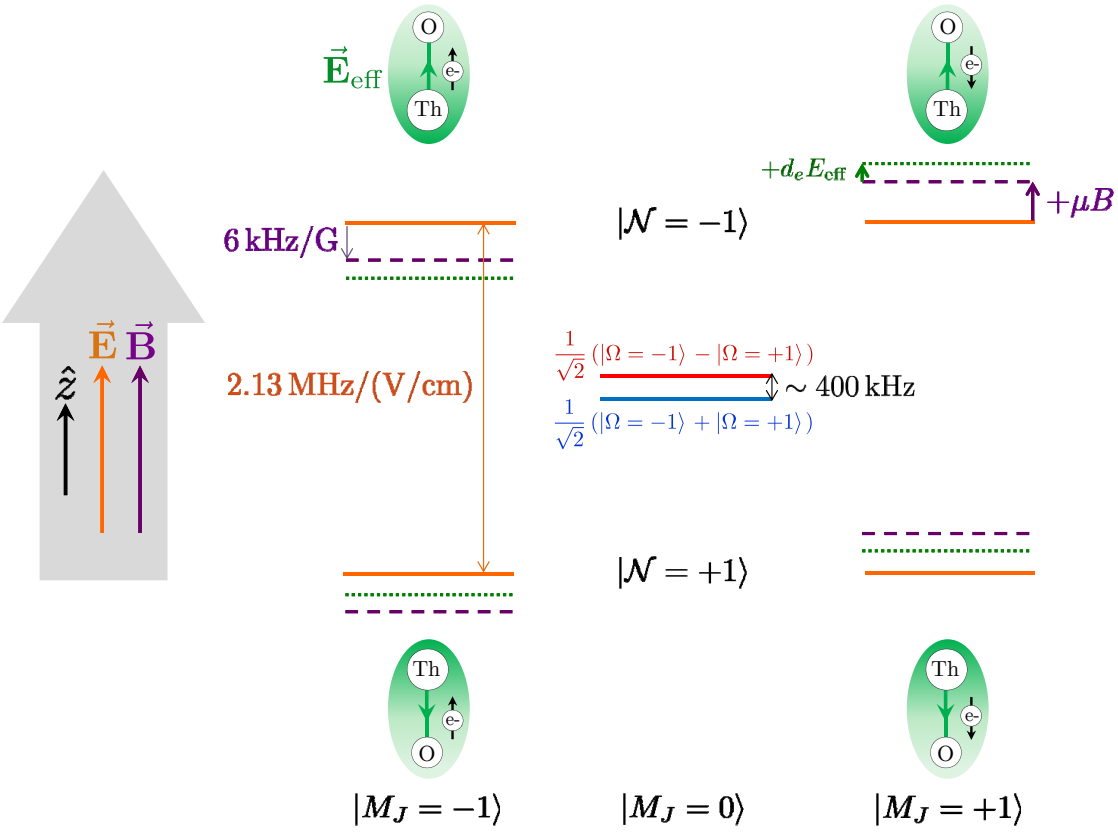

Figure 1.3.2: Sublevel structure of the $H\,^3\Delta_1$ state of ThO. In the absence of applied $\vec{E} \parallel \hat{z}$ and $\vec{B} \parallel \hat{z}$ fields, the stationary states are the Ω-doubled parity eigenstates $\frac{1}{\sqrt{2}}(|\Omega=+1\rangle \pm |\Omega=-1\rangle)$, which are split by a few hundred kHz (blue and red lines). $\vec{E}$-fields of a few V/cm fully mix these doublets in the $M_J \equiv \hat{z}\cdot\vec{J} = \pm 1$ states by resolving the aligned and anti-aligned orientations ($\mathcal{N} = \pm 1$) of the internuclear axis $\hat{n}$. The linear Stark splitting between these $\mathcal{N}$ states (solid orange lines) is measured to be 2.13 MHz/(V/cm) [190]. In an applied $\vec{B}$-field, the measured Zeeman shift (dashed purple lines) between the $M_J = \pm 1$ states of each $\mathcal{N}$ sublevel is ± 12 kHz/G [151, 190]. If $d_e \neq 0$, these M_J levels experience an additional relative shift equal to $\pm 2 d_e \mathcal{E}_{\text{eff}}$ (dotted green lines). These relative shifts are in opposite directions in the two $\mathcal{N}$ levels since $\vec{\mathcal{E}}_{\text{eff}}$ points in opposite directions. The ACME experiment is performed by measuring the energy shift between the states $|\mathcal{N}, M_J = -1\rangle$ and $|\mathcal{N}, M_J = +1\rangle$ for both $\mathcal{N}$ as a function of the electric and magnetic field and looking for a shift that depends only on the signs of $\mathcal{N}$ and $\vec{E}$. See Sections 2.1 and 2.5.

Finally, the most advantageous property of the $H\,^3\Delta_1$ state of ThO is its extremely large static electric dipole polarizability resulting from a pair of nearly degenerate, opposite-parity sublevels split by only a few hundred kHz [68, 134, 135]. This level structure gives polarizabilities on the order of 10^4 or more times larger than for a more typical diatomic molecule state, in which an applied electric field polarizes the molecule by mixing opposite-parity rotational levels typically spaced by many GHz.[14] The opposite-parity sublevels in the H, $J = 1$ state are formed by even and odd combinations of molecular orbitals with opposite signs of the quantum number $\Omega \equiv \hat{n} \cdot \vec{J}$ (the projection of the total angular momentum on the molecular bond axis $\hat{n}$, which in our convention points from the oxygen to the thorium nucleus) and are a general feature of states with $\Omega \geq 1$ in Hund's case (c) molecules [57, 99].

Such "Ω-doubled" states are immensely valuable to eEDM searches because they can be fully mixed in modest electric fields of only a few to a few tens of V/cm, completely polarizing the molecule [57, 58]. Thus, eEDM experiments on molecules with Ω-doublets can take full advantage of the molecules' effective internal fields while avoiding the technical challenges and potential systematic errors introduced by large lab fields. Furthermore, because the effective electric field in a fully polarized molecule is independent of the externally applied electric field $\vec{E}$, the eEDM signal is also independent of the magnitude of the applied field (see Eq. (2.7)), allowing such experiments to set limits on systematic effects correlated with $|E|$. Another benefit of the Ω-doublet in ThO is that the polarized H-state molecule can be spectroscopically prepared with its dipole either aligned or anti-aligned with $\vec{E}$, allowing us to switch the sign of the electric field experienced by the eEDM without physically changing the laboratory field [114]. As discussed in Section 2.7, this provides a way to rule out systematic errors correlated with the sign of the applied field, such as leakage currents, motional magnetic fields, and geometric phases [189, 191].

Besides these features, ThO also provides manifold technical advantages. Most of the

[14] The polarizability advantage such molecules have over atoms is even more pronounced: Polarizing an atom typically requires mixing opposite-parity *electronic* states spaced by tens of THz.

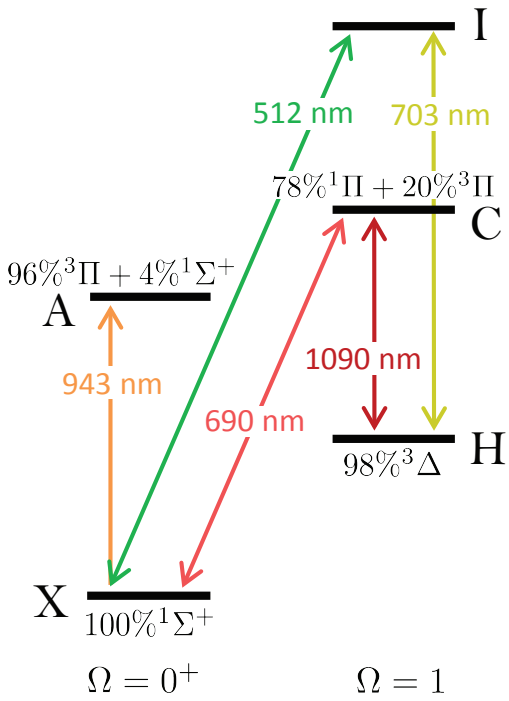

Figure 1.3.3: ThO levels and transitions used in the ACME Gen. I and II experiments, based on [69, 149, 189]. All relevant states are in the ground vibrational level. The electronic states are denoted by letters, and the angular momentum character of each state is indicated by molecular spectroscopy symbols. The wavelength of each transition is given in nanometers. The ACME measurement schemes are described in Section 2.1 and Chapter 3.

relevant optical transitions (see Fig. 1.3.3) are well studied [68, 86, 102, 131, 148, 192] and accessible to diode lasers. In addition, ThO has no nuclear spin and so evades the complexities of hyperfine structure. Finally, despite the fact that ThO is chemically reactive and its precursors are highly refractory, it can be produced in large quantities in a cryogenic buffer gas beam [106] (See Section 2.2 and Chapter 5).

When I was young, I said to God, "God, tell me the mystery of the universe." But God answered, "That knowledge is reserved for me alone." So I said, "God, tell me the mystery of the peanut." Then God said, "Well George, that's more nearly your size."

—George Washington Carver

2

ACME Experimental Methods and eEDM Limit from Generation I

The signature of an electron EDM in the ACME experiment is a tiny phase shift in a Ramsey fringe resulting from a nonzero $\vec{d}_e \cdot \vec{\mathcal{E}}_{\text{eff}}$ interaction. In order to measure this phase shift, ACME injects a bright beam of ThO into a region of highly controlled electric and magnetic fields and uses an optical state preparation and readout scheme to make a precision measurement of the electron spin precession within the molecule. In 2014, this approach produced an order-of-magnitude improvement on the upper bound of the eEDM [10]. The measurement scheme, apparatus, and results of the Gen. I ACME experiment are described here.[1]

More detailed discussions of the ACME experiment scheme, apparatus, data analysis, and

[1]Parts of this chapter were adapted from [35], which is licensed under Creative Commons. I gratefully acknowledge the contributions and editing assistance of the other members of the ACME collaboration. In particular, Nick Hutzler contributed some of the language in Section 2.2. Use of figures provided by others is acknowledged in the captions.

results are provided in references [10, 11] and the theses of Amar Vutha [188], Yulia Gurevich [91], Nick Hutzler [105], Paul Hess [100], Ben Spaun [180], and Brendon O'Leary [143].

2.1 Measurement Scheme

The Gen. I ACME apparatus and measurement scheme are illustrated in Fig. 2.1.1.[2] Traveling approximately south-to-north, in what we define to be the $+\hat{x}$ direction, a pulse of ThO molecules with a ≈ 2 ms fly-through time (i.e. a physical pulse length of ≈ 40 cm at a velocity of 180 m/s [106]) enters the magnetically shielded interaction region where the measurement takes place. (See Section 2.2 for details on the molecule beam properties and Section 2.4 for details on the interaction region.) After being collimated by a square aperture to a cross-sectional size of 1 cm $\times$ 1 cm, the beam flies in between the electric field plates and interacts with an optical pumping laser propagating along the east-west $\hat{z}$ direction that drives the transition $|X, J = 1, M_J = \pm 1\rangle \rightarrow |A, J = 0, M_J = 0\rangle$ with a wavelength of 943 nm (see Fig. 1.3.3; all levels used in the experiment are in the vibrational ground state $v = 0$). The optical pumping beam makes two passes with orthogonal polarizations through the molecule beam so that both $|M_J| = 1$ sublevels in $|X, J = 1\rangle$ are fully depleted [180]. Due to the $A \rightsquigarrow H$ branching ratio, $29 \pm 7\%$ of each $|X, J = 1, M = \pm 1\rangle$ sublevel ends up spontaneously emitted into the $H\,^3\Delta_1$ state, in which we perform the measurement (see Section 1.3.1 for further discussion of the properties of ThO's H state).

As shown in Fig. 1.3.2, the $|H, J = 1\rangle$ state consists of two opposite-parity Zeeman manifolds split by just ≈ 400 kHz [68, 134, 190]. In the presence of the $|\vec{E}| > 10$ V/cm electric field present in the interaction region, the $M_J \neq 0$ states mix with their opposite-parity counterparts and become fully polarized. In this regime, the Stark shift of the $M_J = \pm 1$ states is linear, with an electric dipole moment of 2.13(2) MHz/(V/cm) [190]. The states whose dipole is aligned along (against) the applied field are described by quantum number

[2] This description of the ACME measurement scheme is based on that in [189].

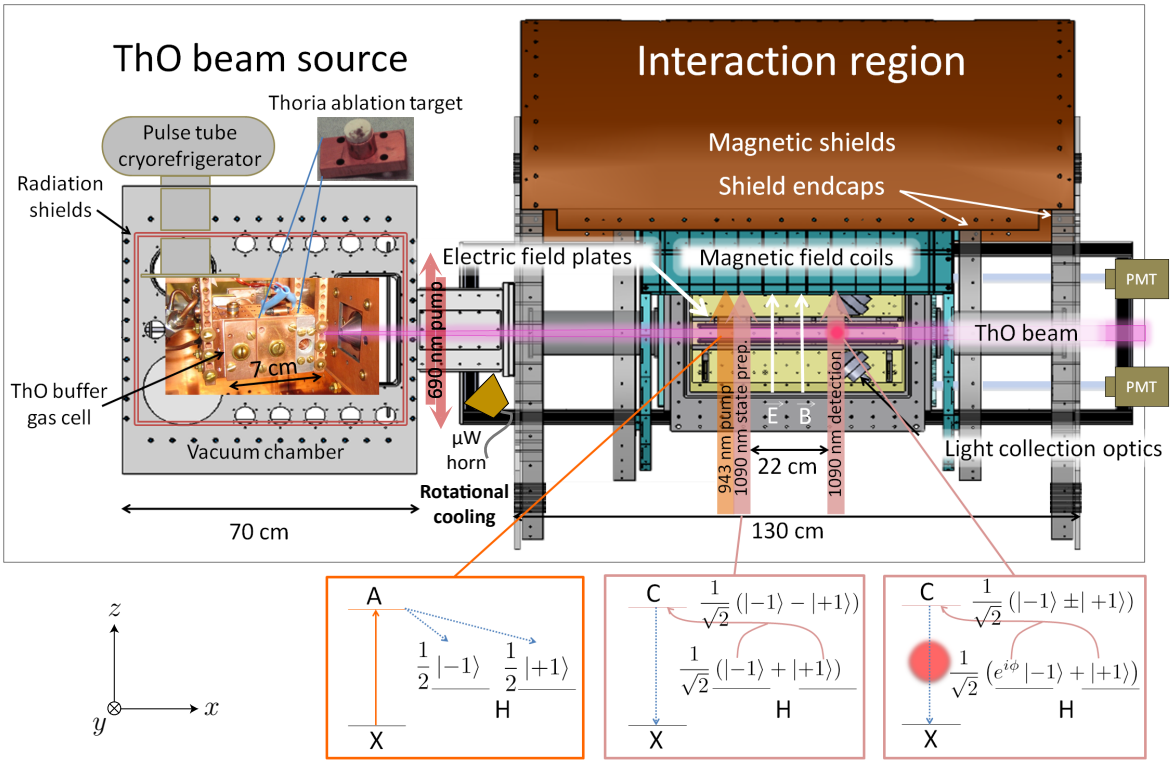

Figure 2.1.1: Schematic of the ACME apparatus and measurement scheme described in the text. On the left, a pulse of gas-phase ThO molecules is produced and cooled in a buffer gas cell (see Section 2.2) and flows out towards the right in a beam (see Section 2.3). This beam enters a magnetically shielded interaction region (see Section 2.4) where uniform, parallel E- and B-fields are applied. At the entrance of the field region, the molecules are pumped from the $|X, J = 1, M_J = \pm 1\rangle$ states to the $|A, J = 0, M_J = 0\rangle$ state, from which they spontaneously decay to the $|H, J = 1\rangle$ state, equally populating the $|M = \pm 1, \mathcal{N} = \pm 1\rangle$ sublevels (see Fig. 1.3.2). Next, a pure superposition of Zeeman sublevels $|X_\mathcal{N}\rangle$ ($|Y_\mathcal{N}\rangle$) (see Eq. (2.3)) is prepared in one of the two $\mathcal{N}$ states of H by pumping out the orthogonal superposition $|Y_\mathcal{N}\rangle$ ($|X_\mathcal{N}\rangle$) using linearly polarized light resonant with the $|H, J = 1, \mathcal{N} = \pm 1\rangle \rightarrow |C, J = 1, M_J = 0, \mathcal{P} = \pm 1\rangle$ transition. Next, the molecule state precesses in the applied E- and B-fields for approximately 1.1 ms as the beam traverses the 22 cm length of the interaction region. The relative phase accumulated between the Zeeman sublevels depends on d_e through Eq. (2.7). Near the exit of the field region, we read out the final state of the molecules: By exciting the $|H, J = 1, \mathcal{N} = \pm 1\rangle \rightarrow |C, J = 1, M_J = 0\rangle$ transition with rapidly switched orthogonal ($\hat{x}$ and $\hat{y}$) linear polarizations and detecting the $C \rightsquigarrow X$ fluorescence from each polarization, we project the population onto the $|X_\mathcal{N}\rangle$ and $|Y_\mathcal{N}\rangle$ states and compute the asymmetry $\mathcal{A}$ between these states. The phase from Eq. (2.7) depends on measured values via the expression $\cos 2\phi = \mathcal{A}$ (see Eq. (2.11)). Machine drawing adapted from Amar Vutha's thesis [188].

$\mathcal{N} = +1$ ($\mathcal{N} = -1$), defined by the relation

$$\mathcal{N} \equiv \mathrm{sgn}\left(\hat{n} \cdot \vec{E}\right) = \mathrm{sgn}\left(\hat{z} \cdot \vec{E}\right) M_J \Omega, \tag{2.1}$$

where in our convention the internuclear axis $\hat{n}$ points from the negatively charged species (O) to the positively charged species (Th).

Each of the four fully polarized states $M_J = \pm 1$, $\mathcal{N} = \pm 1$ inherits 1/6 of the total population decaying from the A state. As described below, we select a single pure state in $|H, J = 1\rangle$ in which to perform the measurement, leading to an overall $1/6 \times 29\% \approx 5\%$ state preparation efficiency from each of the starting $|X, J = 1, M_J = \pm 1\rangle$ states.

Since the H state is populated by spontaneous decay from A, it is initially in a mixed state, with all $|M_J = \pm 1, \mathcal{N}\rangle$ sublevels approximately equally populated. In order to prepare a pure state, we apply another optical pumping laser with linear polarization that addresses the transition $|H, J = 1, \mathcal{N}\rangle \rightarrow |C, J = 1, M_J = 0, \mathcal{P}\rangle$, where the molecule orientation $\mathcal{N} = \pm 1$ and the excited state parity $\mathcal{P} = \pm 1$ (C, like H, has an $\Omega = 1$ parity doublet) can be chosen independently by tuning the laser frequency to resonance with each of the four spectrally resolved transitions. This laser is at 1090 nm and is referred to as the "state preparation laser."

A given linear polarization in the xy plane couples just one equal superposition of $M_J = +1$ and $M_J = -1$ states in the selected $\mathcal{N}$ level of H to the selected $\mathcal{P}$ level of C, leaving the orthogonal superposition dark [91]. For example, if light polarized along the $\hat{y}$ direction addresses the transition to the $\mathcal{P} = +1$ state in C, the *odd* superposition of M_J sublevels in H ($|Y_\mathcal{N}\rangle$, defined below) is depleted [11], leaving behind the *even* superposition ($|X_\mathcal{N}\rangle$) as the initial pure state $|\psi_i^\mathcal{N}\rangle$ on which the measurement is performed:

$$|\psi_i^\mathcal{N}\rangle = |X_\mathcal{N}\rangle\,, \tag{2.2}$$

where we define

$$
\begin{aligned}
|X_{\mathcal{N}}\rangle &\equiv \frac{1}{\sqrt{2}}\left(|M_J=+1;\mathcal{N}\rangle + |M_J=-1;\mathcal{N}\rangle\right), \text{ and} \\
|Y_{\mathcal{N}}\rangle &\equiv \frac{1}{\sqrt{2}}\left(|M_J=+1;\mathcal{N}\rangle - |M_J=-1;\ \mathcal{N}\rangle\right),
\end{aligned} \tag{2.3}
$$

with $\mathcal{N}=+1$ ($\mathcal{N}=-1$) corresponding to the lower-energy (higher-energy) Ω-doublet component. $|Y_{\mathcal{N}}\rangle$ can be selected as the initial state either by pumping through the $\mathcal{P}=-1$ level in C using y-polarized light or by pumping through the $\mathcal{P}=+1$ level in C using x-polarized light. Pumping through $\mathcal{P}=-1$ with x-polarized light produces $|X_{\mathcal{N}}\rangle$, as in Eq. (2.2). Thus, we have the freedom both to prepare physically distinct initial states and to prepare the *same* initial state using physically distinct paths. This provides a way of searching for or ruling out potential systematic errors correlated with our choice of $|\psi_i^{\mathcal{N}}\rangle$ or its preparation method. Such experimental handles or "switches" are an important part of our data analysis and systematic error search effort and are discussed further in Sections 2.5 and 2.7.

After state preparation, the molecules fly through the interaction region, where—as the name suggests—they interact with the applied electric and magnetic fields. The relative phase of the two $|M_J\rangle$ states in $|\psi_i^{\mathcal{N}}\rangle$ is shifted by the interaction of $\vec{\mu}_{H,J=1}$ with $\vec{B}$ and $\vec{d}_e$ with $\vec{\mathcal{E}}_{\text{eff}}$. As illustrated in in Fig. 1.3.2, the energy shifts of the $M_J=\pm1$, $\mathcal{N}$ levels are given approximately by

$$
U(M_J,\mathcal{N},\vec{E},\vec{B}) = -g_{H,J=1}M_J\mu_B B\tilde{B} - d_{H,J=1}\mathcal{N}E - d_e\mathcal{E}_{\text{eff}}M_J\mathcal{N}\tilde{E}, \tag{2.4}
$$

where $g_{H,J=1}=-4.40(5)\times10^{-3}$ and $d_{H,J=1}=0.84(2)\ ea_0$ are the magnetic g-factor and electric dipole moment of the $H,J=1$ state, respectively [118, 151, 190]; μ_B is the Bohr magneton; e is the electron charge; and a_0 is the Bohr radius. The terms (from left to right) give the interaction of the magnetic dipole with the external magnetic field, the Stark shift of the $\mathcal{N}$ states in the external electric field, and the interaction of the eEDM with the effective molecular electric field. Here we assume that the H-state is fully polarized, which

occurs in external fields above ~ 1 V/cm, much smaller than the typical experimental field of 140 V/cm. In our convention, $\vec{d}_e$ is aligned along the electron spin; a negative measured value for d_e would indicate that the eEDM is anti-aligned with the spin. The magnitudes of applied field vectors are given as the variable without a vector symbol, e.g. $B = |\vec{B}|$. The tilde denotes the sign of a quantity's projection on the lab-fixed quantization axis of the experiment, e.g. $\tilde{B} = \mathrm{sgn}(\hat{z} \cdot \vec{B})$. For an explanation of all the signs in this formula, see Nick Hutzler's thesis, especially the excellent Figure 5.1. Also note that this simple formula neglects a large number of important terms, such as the E-field dependence of the g-factors [30], background fields, motional fields, etc.; nevertheless, this expression will be sufficient to explain the basic measurement procedure.

The first and third terms in Eq. (2.4) produce a relative energy shift ΔU between the M_J sublevels in $|\psi_i^{\mathcal{N}}\rangle$, which results in a relative phase accumulation $2\phi = \Delta U \tau/\hbar$ over a time τ. Thus, after the molecules have been permitted to evolve freely during their flight through the $L = 22$ cm length of the interaction region, the final wavefunction is

$$|\psi_f^{\mathcal{N}}\rangle = \frac{1}{\sqrt{2}}\left(e^{-i\phi}\,|M_J = +1; \mathcal{N}\rangle + e^{i\phi}\,|M_J = -1; \mathcal{N}\rangle\right). \tag{2.5}$$

For a molecule with velocity v along the beam axis, the accumulated phase ϕ can be expressed as

$$\phi = \int_{x=0}^{x=L} \left[U(M_J = +1, \mathcal{N}, \vec{E}, \vec{B}) - U(M_J = -1, \mathcal{N}, \vec{E}, \vec{B})\right] \frac{dx}{2\hbar v} \tag{2.6}$$

$$= -\int_{x=0}^{x=L} \left(d_e \mathcal{E}_{\mathrm{eff}} \mathcal{N} \tilde{E} + g_{H,J=1} \mu_B B \tilde{B}\right) \frac{dx}{\hbar v} \equiv \phi_{\mathcal{E}} + \phi_B. \tag{2.7}$$

Using the fact that our beam source has a narrow forward velocity distribution (with average forward velocity $v \approx 180$ m/s and spread $\Delta v_{\parallel} \ll v$, see Section 2.2), we make the approximation that all molecules experience the same phase shift as they traverse the interaction region. Furthermore, because the $\vec{E}$- and $\vec{B}$-fields are highly uniform along the

length of the interaction region (see Section 2.4), we can pull out the integrand and write:

$$\phi_{\mathcal{E}} \approx -d_e \mathcal{E}_{\text{eff}} \mathcal{N} \tilde{E} \frac{L}{\hbar v}, \text{ and} \tag{2.8}$$

$$\phi_B \approx -g_{H,J=1} \mu_B B \tilde{B} \frac{L}{\hbar v} \tag{2.9}$$

for all molecules in the beam.

We isolate the phase ϕ by projecting the final state $|\psi_f^{\mathcal{N}}\rangle$ from Eq. 2.5 onto the orthogonal "quadrature" states $|X_{\mathcal{N}}\rangle$ and $|Y_{\mathcal{N}}\rangle$, defined in Eq. 2.3. The projection measurement is performed in a manner exactly similar to the state preparation scheme described above: The state $|X_{\mathcal{N}}\rangle$ is excited by an x-polarized (y-polarized) laser coupling the $|H, J = 1, \mathcal{N}\rangle$ and $|C, J = 1, M_J = 0, \mathcal{P} = +1\rangle(|\mathcal{P} = -1\rangle)$ states, while $|Y_{\mathcal{N}}\rangle$ is excited by a y-polarized (x-polarized) laser coupling the $|H, J = 1, \mathcal{N}\rangle$ and $|C, J = 1, M_J = 0, \mathcal{P} = +1\rangle(|\mathcal{P} = -1\rangle)$ states. The C state quickly decays to the ground state with a lifetime of 490 ± 40 ns [100], emitting fluorescence at 690 nm, which we collect with an array of lenses and focus into fiber bundles and light pipes, described in Section 2.4. These, in turn, deliver the light to two photomultiplier tubes (PMTs), where it is detected. The fact that the energy of ThO's H state is appreciably higher than the ground state means that the emitted fluorescence photons are at a much shorter wavelength than those in the 1090 nm $H \rightarrow C$ "readout" laser. This allows us to reject scattered laser light efficiently using interference filters.

The probability of detecting a molecule in the quadrature state $|X_{\mathcal{N}}\rangle$ ($|Y_{\mathcal{N}}\rangle$), given by $P_X = |\langle X_{\mathcal{N}}|\psi_f^{\mathcal{N}}\rangle|^2$ ($P_Y = |\langle Y_{\mathcal{N}}|\psi_f^{\mathcal{N}}\rangle|^2$), can be expressed as $P_X = \cos^2\phi$ ($P_Y = \sin^2\phi$). The detected fluorescence signal from each quadrature state, measured in photoelectron counts, is proportional to its population. Letting the per-pulse signal from each quadrature state be S_X and S_Y, we have $S_{X(Y)} = S_0 P_{X(Y)}$, where S_0 is the total signal in a single beam pulse. Thus, when plotted against the strength of the applied magnetic field B, S_X and S_Y trace out two sinusoidal curves, or—in the parlance of separated oscillatory field spin precession measurements—Ramsey fringes [161], of opposite phase. For the highest sensitivity to d_e,

we "sit on the side of the Ramsey fringe," where small changes in $\phi_{\mathcal{E}}$ are most noticeable, i.e. where $\partial/\partial\phi_{\mathcal{E}}[S_{X(Y)}]$ is maximized. To achieve this condition, we adjust the magnetic field to yield a bias phase $\phi_B = \tilde{B}\pi/4$ and rewrite S_X and S_Y as[3]

$$S_X \approx S_0\left(-\tilde{B}\phi_{\mathcal{E}} + \frac{1}{2}\right), \text{ and } S_Y \approx S_0\left(+\tilde{B}\phi_{\mathcal{E}} + \frac{1}{2}\right) \tag{2.10}$$

to first order in small $\phi_{\mathcal{E}}$. Then the EDM phase $\phi_{\mathcal{E}}$ can be determined by constructing the asymmetry $\mathcal{A}$ between the signals from each quadrature:

$$\mathcal{A} \equiv \frac{S_Y - S_X}{S_X + S_Y} \approx 2\tilde{B}\phi_{\mathcal{E}} \tag{2.11}$$

$$\phi_{\mathcal{E}} \approx \frac{\tilde{B}}{2}\mathcal{A}. \tag{2.12}$$

Note from Eq. (2.8) that $\phi_{\mathcal{E}}$ is odd in $\tilde{E}$ and $\mathcal{N}$, even in $\tilde{B}$, and proportional to $\mathcal{E}_{\text{eff}}$, which is independent of E over a large range. In Section 2.5 we discuss how to use these correlations to isolate the EDM term from various systematic effects.

The shot-noise-limited statistical uncertainty in $\phi_{\mathcal{E}}$ is $1/(2\mathcal{C}\sqrt{N})$ (see Eq. (1.6–1.9)), where N is the total number of photoelectron counts and the quantity $\mathcal{C}$ introduced in this expressions is the Ramsey fringe contrast (or visibility), which accounts for inefficiencies in state preparation and varying precession times for different molecules [11, 188]. Therefore, the shot-noise limited uncertainty in the measured eEDM value is (from differentiating d_e with

[3]The signal slope is actually maximum with respect to $\phi_{\mathcal{E}}$ for any $|\phi_B| = \pi/4 + n\pi/2$ with integer n, and in practice we run at some of these higher B values to check for systematics; however, we restrict ourselves to the case of $n = 0$ in this discussion both for simplicity and because higher B values reduce the fringe contrast $\mathcal{C}$ due to velocity dispersion effects.

Another way to maximize the signal slope independent of B is to adjust the readout basis by rotating the polarization of the readout laser with respect to the state preparation laser. The effect of this is to add an offset phase θ to Eq. (2.6) and (2.7). This offset phase can always be tuned so that $\phi_B + \theta \approx \pi/4 + n\pi/2$. This is another valuable "switch" we used in the experiment and will be discussed briefly in Section 2.5, but we will ignore it here for the sake of keeping the discussion simple. However, we observe in passing that this approach is analogous to adjusting the phase of the readout field in Ramsey's original separated oscillatory field scheme [161].

respect to $\phi_{\mathcal{E}}$ in Eq. (2.8)) [189]

$$\delta d_e = \frac{\hbar}{2\mathcal{C}\tau\mathcal{E}_{\text{eff}}\sqrt{\dot{N}T}}, \tag{2.13}$$

where $\tau = L/v \approx 1.1$ ms is the precession time of the molecules in the fields, $\dot{N} \approx 3 \times 10^4$/s (see Section 2.6) is the time-averaged count rate at the detectors, and T is the total experimental running time, which ended up being about 2 weeks in Gen. I. The quantity $\mathcal{E}_{\text{eff}} = 77.6$ GV/cm is determined by physical properties of the H-state (see Section 1.3.1), as described above, and the large ThO fluxes achieved by the ACME beam source help to keep the uncertainty low by providing large $\dot{N}$. As described in references [188, 189], the length $L = 22$ cm of the interaction region was chosen to optimize the trade-off between having a long interaction time τ and minimizing the signal loss due to spontaneous decay of the H state, which has a lifetime of 1.8 ms.

2.2 ThO Ablation-Based Buffer Gas Beam

ACME uses a cryogenic buffer gas beam (CBGB) source to achieve high single-quantum-state intensities of the chemically reactive molecular species ThO. The heart of the cold beam apparatus, the buffer gas cell (see Fig. 5.1.3 for a schematic), is similar to those described in existing buffer-gas-cooled beam publications [22, 107, 132, 146, 147, 173]. The ACME Gen. I beam source was developed primarily by Nick Hutzler and was characterized and described in detail in [105, 106].

The cell is a small copper chamber with a cylindrical inner bore (12.7 mm in diameter and ≈ 9 cm long) whose axis points along $\hat{x}$ (see Fig. 2.1.1). Precursor targets are set into alcoves along one side of the bore, while the other side has windows for admitting the ablation laser that are offset on long tubes or "snorkels" to reduce clouding due to dust accumulation. The cell is mounted in vacuum and held at a temperature of 16 K with a Cryomech PT415 pulse tube cooler and a small ≈ 1 W resistive heater (the heater is not needed while the ablation laser is firing). Neon (Ne) buffer gas, pre-cooled by heat sinks consisting of tube

lengths thermally anchored to the warmer stage of the pulse tube and the exterior of the cell, flows into the cell through a fill line at the "upstream" end of the cylindrical volume. At the "downstream" end, an aperture 5 mm in diameter in a thin (0.5 mm) plate is open to the external vacuum, allowing the buffer gas to flow out as a beam.

The cell is surrounded by two nested chambers of metal that are also thermally connected to the pulse tube cooler (see Fig. 3.0.1). The inner chamber is a copper box held at 4 K by the PT415 cold stage that acts as a high-speed, large-capacity cryopump for neon, maintaining a high vacuum of ~ 3 μTorr in the system despite large buffer gas throughputs. "Sorbs," which are large surface areas of activated coconut charcoal glued to copper plates with a cryo-compatible, thermally conductive epoxy[4] and cooled to 4 K, also help to keep the pressure low by adsorbing residual helium and other gases in the chamber.[5] The outer chamber is an aluminum box covered in aluminized Mylar "super insulation" and kept at ≈ 60 K by the PT415 warm stage. It serves to shield the inner cryogenic regions from blackbody radiation emitted by the room-temperature vacuum chamber. The 4 K cryopump chamber, 60 K radiation shield, and vacuum chamber all have windows for optical access and apertures (described in Section 2.3.2) to transmit and collimate the buffer gas beam.

The source of ThO molecules is a ceramic target of thoria (ThO_2) fabricated in house. We make the targets using established techniques, including a recipe for ThO_2 mock nuclear fuel pellets furnished by Oak Ridge National Labs [19, 116, 189]. The details of our target-making procedure are provided in Appendix B. The general method is to prepare a mixture of 325 mesh (≤ 44 μm) ThO_2 powder and the "sintering agent" niobia (Nb_2O_5) by ball milling, pre-compacting, meshing, and purifying it in a 1000°C furnace. The prepared powder is

[4]The preferred epoxy is Stycast 2850FT mixed with 24 LV catalyst in a ratio of 100 parts epoxy to 7.5 parts catalyst and cured for at least a day.

[5]The utility of sorbs is well established for CBGBs with helium buffer gas, but it has not been proven for neon-based beams. In fact, neon that reaches the sorbs is likely to clog the pores of the charcoal and limit its effectiveness as a pump for residual helium. We nevertheless include sorbs in our system because they are unlikely to harm and may help the vacuum quality. In addition, about two-thirds of our sorb surfaces are "protected," meaning that gas must bounce multiple times off of other cryogenic surfaces before reaching the charcoal. This limits the pumping speed of the sorb but should also limit the pore-occluding detrimental effects of neon and other species, which stick to the protecting surfaces instead of the charcoal.

then mixed with a sticky binder to inhibit crumbling and cold-pressed in a 1/2" or 3/4" diameter die to a thickness of $\approx$ 1/8–1/4". The pressed pellet is then sintered at least twice at a temperature of 1150°C or above. This process produces sturdy, very hard targets with masses of 5 to 10 g and densities between 65% and 85% of the theoretical maximum ThO_2 density of 10.0 g/cm^3.

ThO molecules are introduced into the cell via laser ablation: A Litron Nano TRL 80-200 pulsed Nd:YAG laser is fired at the ThO_2 target, creating an initially hot plume of gas-phase ThO molecules, plus other detritus. The ablation pulse energy is typically set to 60–100 mJ, and the repetition rate is 50 Hz. The pulse length is $\approx$ 5 ns. Each target lasts for approximately 1 month of hard (10 h/day) running, or $\sim$ 50 million ablation shots, before becoming depleted (i.e. producing 1/3 to 1/2 of their original yields). Upon removal, the targets have typically lost 30% to 50% of their original mass and appear darkened, pitted, and spiky all over from ablation damage. Under typical running conditions, the in-cell production rate of ThO is $\sim 10^{14}$ molecules/s [105, 106]. For a month of running time at 10 h/day, this makes a total of $\sim 10^{20}$ molecules produced per target. At a mass of 248 amu per molecule, this implies that $\sim$ 40 mg or $\sim$ 1% of the ablation product from a target is converted into gas-phase ThO molecules.

On a timescale rapid compared to the emptying time of the cell into the beam region, the hot ThO molecules in the ablation plume thermalize with the 16 K buffer gas in the cell. Continuous Ne flow at $\approx$ 40 SCCM (standard cubic centimeters per minute; for reference, 1 SCCM $= 4.48 \times 10^{17}$ atoms/s) maintains a buffer gas density of $n_0 \approx 10^{15}$–10^{16} cm^{-3} ($\approx 10^{-3}$–10^{-2} Torr, where the subscript "0" indicates the steady-state value of the quantity in the cell). This is sufficient for rapid translational and rotational thermalization of the molecules and for producing hydrodynamic flow out of the cell aperture that sweeps along and extracts a significant fraction of the molecules before they can diffuse to the cell walls and stick. The measured extraction fraction is $\approx$ 10% [106]. The result is a $\approx$ 2 ms long pulsed beam of cold ThO molecules entrained in a continuous flow of buffer gas.

Just outside the cell exit, the buffer gas density is still high enough for ThO–Ne collisions to play a significant role in the beam dynamics. The average thermal velocity of the buffer gas atoms is higher than that of the molecules by a factor of $\sqrt{m_{\mathrm{mol}}/m_{\mathrm{b}}}$, where the subscripts "b" and "mol" indicate buffer gas and molecule quantities, respectively. Consequently, the ThO molecules ($m_{\mathrm{mol}} = 248$ amu) experience collisions primarily from behind, with the fast Ne atoms ($m_{\mathrm{b}} = 20$ amu) pushing the slower ThO molecules ahead of them as they exit the cell. This accelerates the molecules to an average forward velocity v_f that is larger than the thermal velocity of ThO. As the buffer gas pressure in the cell is increased, v_f approaches $v_{0,\mathrm{b}}$, the thermal velocity of the buffer gas.

The angular distribution of the molecule beam has a characteristic apex angle θ given by $\tan(\theta/2) \equiv \Delta v_\perp / 2v_f$, where $\Delta v_\perp$ is the transverse velocity spread of the beam. For the ACME beam, the apex angle is $\theta \approx 39°$, and the characteristic solid angle is $\Omega \approx 0.35$ sr. The beam velocity is measured to be ~ 180 m/s. As the gas cloud expands nearly isentropically out of the cell into the vacuum, it must also cool. The measured final longitudinal and rotational temperature of the beam is ~ 4 K, yielding a forward velocity distribution $\Delta v_\parallel$ of ~ 40 m/s FWHM (full width at half maximum) and efficiently populating low-lying rotational levels in the ground electronic state (e.g. $\sim 30\%$ in $J = 1$). The total number of molecules per pulse in one of the most populated quantum states is measured to be $\sim 10^{11}$. This slow, cold, high-intensity molecular beam provides ACME with a long interaction time τ over a short distance, low phase-decoherence due to the narrow velocity spread, and a high count rate $\dot{N}$.

In practice, we typically run the ACME ablation beam source for no more than about 10–12 hours at a time because of the necessity of regularly warming up to desorb the neon and regenerate the cryopumping surfaces. Our usual running routine is to start the beam source and tune up the experiment in the morning, run through the day, and "de-ice" the beam box in the evening. After the de-ice procedure, which takes about 1 hour and involves heating the 4 K stage to ~ 60 K, the pulse tube cooler is switched back on, and the beam

source is allowed to cool back down overnight. By the next morning, ≈ 10 hours later, the beam source is cold and ready to run again.

2.3 Beamline

While the vacuum system between the beam source and the interaction region does not (except in the most literal sense) play a central role in the experiment, there are nevertheless important effects that happen to the molecule beam in this intermediate region—including rotational cooling, collimation, and unintended attenuation due to collisions with background gas—which we discuss here. Full descriptions and diagrams of the apparatus can be found in references [11, 143, 188].

2.3.1 Rotational Cooling

After the ThO beam exits the source, a fraction of the population in the $J = 0$, 2, and 3 rotational levels is transferred into the $J = 1$ state via a combination of microwave and optical pumping. This "rotational cooling" stage, described in [11], increases the usable molecule number by a factor of about 1.5–1.8.

2.3.2 Molecule Beam Alignment and Collimation

The ThO beam cross-sectional geometry is defined by a set of apertures along the beamline, summarized in Table 2.3.1.[6] Proper alignment of these apertures is critical to ensuring

Table 2.3.1: List of apertures along the beamline that define the molecule beam cross-section. For apertures outside the beam box, the uncertainties on the distances from the source are ≈ 2 cm.

Distance (cm)	Width (mm)	Description
0	5	Circular cell exit aperture, 0.5 mm thick
2.5	6	Circular knife-edge conical collimator in 4 K shield
11.5	10	Circular aperture in 60 K shield
50	25	Circular differential pumping aperture
52	0–25	Adjustable rectangular collimator
126	10	Final beam-defining square collimator
172	22.5	Field plate exit: width defined by guard rings

[6] The rationale for the size of the final beam-defining collimator is described in Section 3.1.1.

both that the beam is transmitted into the interaction region and that it avoids depositing molecules on the electric field plates, which could result in patch potentials and reduced optical transmission. In particular, the center of the cell exit aperture must be aligned within ≈ 3 mm of the field plate centerline ≈ 1.5 m downstream to ensure that no ballistic molecule trajectories are incident on the field plate exit guard rings. We perform this alignment by adjusting the position of the beam source using a heavy-duty two-axis (y and z) translation stage designed by Emil Kirilov. To measure the beam alignment, we use a theodolite mounted on its own two-axis micrometer translation stage and positioned just outside a window at the end of the beamline so that its telescope points upstream along the center axis of the field plates (referenced to an alignment mark on the guard ring notch on the downstream end of the plates and to the final square collimator on the upstream end). The beam source is then adjusted until the cell is centered in the theodolite crosshairs, as illustrated schematically in Fig. 2.3.1(C). This approach allows the source to be aligned with a tolerance of 0.4 mm. As shown in Fig. 2.3.1(A), this tolerance is quite sufficient to avoid molecule deposition on the field plates.

For ballistic trajectories, the molecule beam cross section in the detection region is completely defined under ordinary running conditions by the cell exit aperture and the final square collimator on the field plates. The other apertures in Table 2.3.1 provide differential pumping to keep the background neon pressure in the interaction region low. However, in the interest of varying as many parameters as possible to check for unexpected signal correlations that could produce a false eEDM (see Section 2.7), we wish to be able to modify the spatial and velocity profile of the molecule beam in the interaction region. For this reason, we installed an adjustable collimator in the beamline, at 40% of the distance from the cell exit to the final collimator. This collimator consists of four straight edges, two horizontal and two vertical, each made of a stacked pair of razor blades with about 1 cm between them along the beamline (x) axis. The four edges are independently adjustable via micrometer linear vacuum feedthroughs. The two horizontal (vertical) edges are vertically (horizontally)

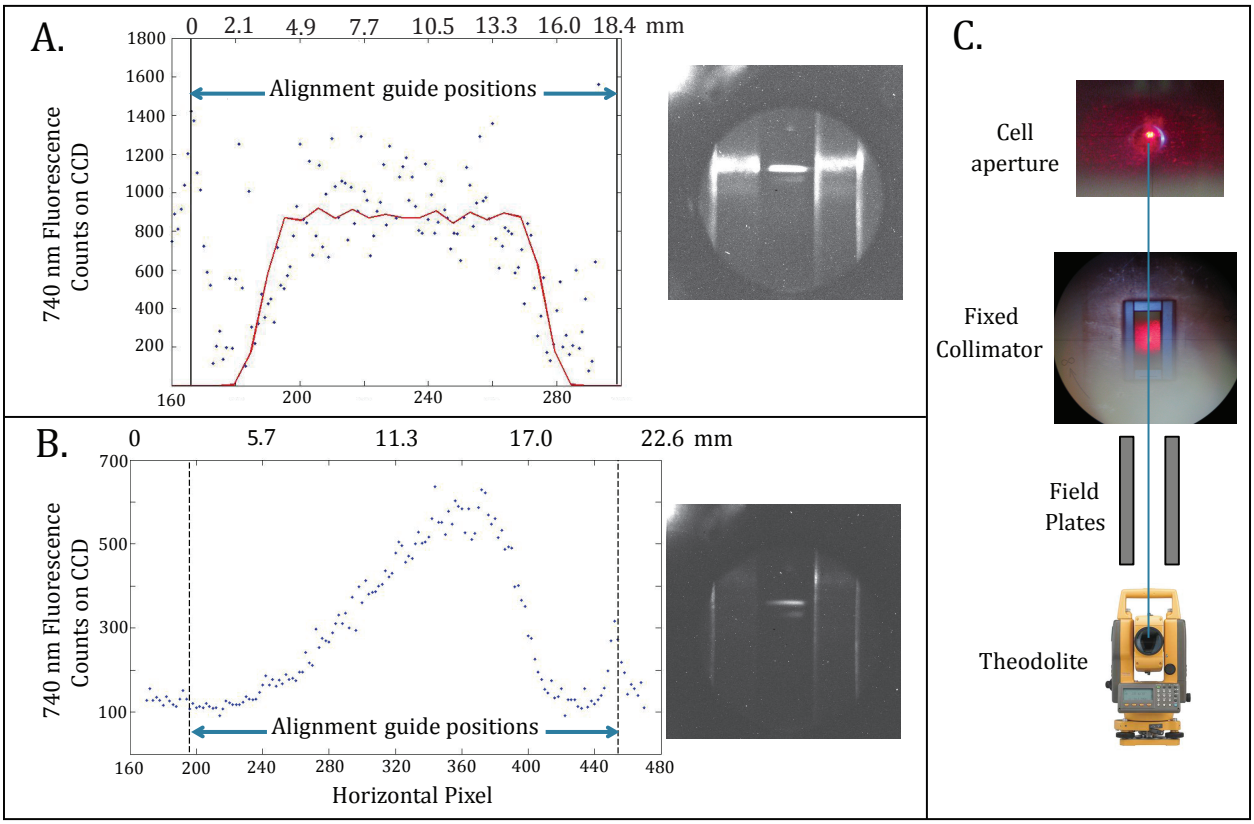

Figure 2.3.1: Beam alignment and collimation schematic and imaging data. (A–B) Laser-induced fluorescence signal from the molecule beam in the detection region with the excitation laser propagating along the z-axis, as measured on CCD camera. Plots show the signal integrated along the x-axis and plotted v. z-axis pixel. Along the top of the plot, pixel position is converted into distance in mm by means of the known, 18.4 mm spacing between alignment guides visible in the photos. (Note that the spacing between the actual field plates is 25 mm.) Beam imaging data is shown for (A) the case of a fully open adjustable collimator (The red line is a simulated beam profile generated by entering our beamline geometry into the ray-tracing software LightTools) and (B) the case where half of the beam is clipped from the left by the adjustable collimator. (C) Theodolite alignment procedure schematic with photographs of the final collimator and cell aperture as viewed through the theodolite telescope. Note that the actual fixed collimator used in the ACME Gen. I experiment is square, rather than rectangular. This data was taken with the help of Ben Spaun.

adjustable from the beam centerline to a distance of approximately ±12 mm on the y-axis (z-axis).

To verify that there were no systematic correlations between the eEDM signal and the transverse position of the ThO beam in the interaction region, we performed runs in which these collimators were moved in one at a time to cut off half of the beam. The effect, shown in Fig. 2.3.1(B), was to shift the center of mass of the beam by 2 mm ($\approx 13\%$ of the total beam width) in each direction in the interaction region. No shift was observed in the eEDM channel (or any other unexpected channel) at the $1 \times 10^{-28}\, e\,\mathrm{cm}$ level. Given that this was $2\times$ the maximum observed misalignment of the molecule beam, we were able to set a limit on any systematics due to imperfect beam positioning at the level of $5 \times 10^{-29}\, e\,\mathrm{cm}$. We also deliberately misaligned the molecule beam by moving the beam source to the west ($+z$ side) of the beamline by ≈ 2 mm. This had the effect of giving the beam an average transverse velocity of ≈ 30 cm/s ($\approx 20\%$ of the total transverse velocity width in the interaction region and a factor of about 3 larger than our alignment uncertainty under ordinary running conditions). We also saw no systematic shift during this run and were able to set a similar limit of a few $10^{-29}\, e\,\mathrm{cm}$ on possible systematics due to angling of the molecule beam.

Notably, during the beam attenuation studies described in Section 2.3.3, we observed a factor of ~ 2 drop in the signal downstream of the adjustable collimator when the blades (the top blade in particular) were close to the detected molecules but not (we believe) actually clipping them. This loss was never satisfactorily explained, and it prompted us to run with the collimator blades fully withdrawn except when performing systematic checks.

2.3.3 Molecule Beam Attenuation

A discrepancy between the observed and expected experimental count rate, discussed in Section 2.6, motivated us to investigate whether collisions with background gas in the beamline were significantly attenuating the molecule beam. As shown in Fig. 2.3.2, we performed an absorption measurement in which we observed the beam signal transmitted from a point 30 cm from the cell to a point 75 cm from the cell, as we varied the pressure on the beamline

by turning off the beamline pumps and closing the interaction region gate valve.

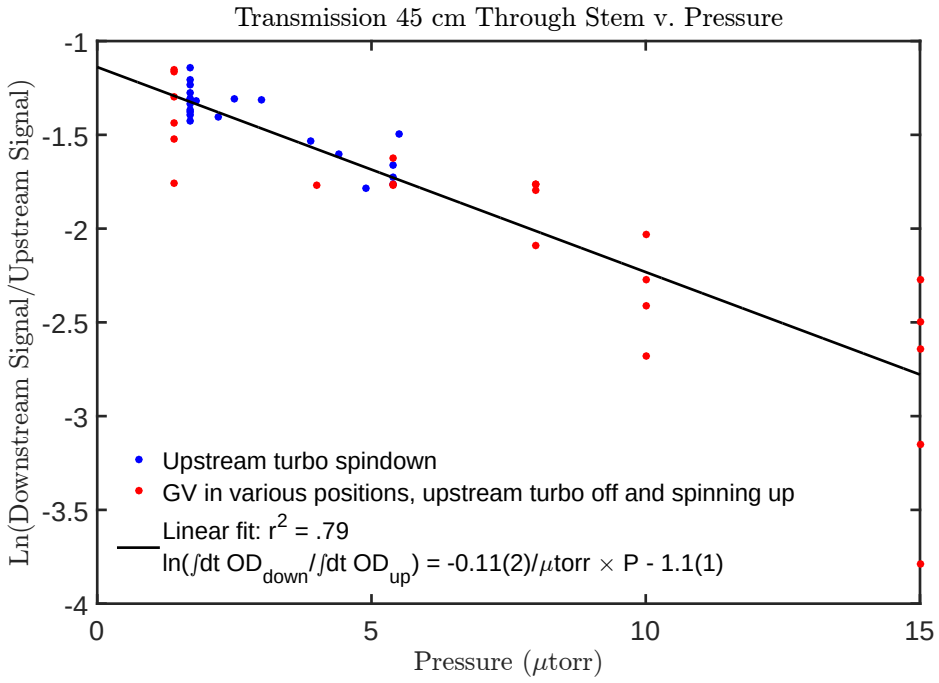

Figure 2.3.2: Beam attenuation in the "stem" region between 30 cm and 75 cm from the source. The pressure was varied by turning off the two 60 l/s stem turbo pumps and allowing them to spin down and back up while also adjusting the position of the gate valve between the stem and the interaction region. The adjustable collimators were fully withdrawn and were not clipping the beam. The molecule signals were measured using single-pass absorption on the $|X, J = 1\rangle \rightarrow |C, J = 1\rangle$ transition. The Litron YAG was firing at 50 Hz, the neon flow rate was 40 SCCM, and each data point represents 1000 pulses averaged together. Note that the "pressure" on this plot is the pressure *reading* on an ion gauge calibrated for nitrogen gas. To get the true pressure, we must divide by the gas correction factor for neon: 0.3.

The functional form of the signal ratio v. pressure curve is expected to be [160]

$$\frac{S_d}{S_u} \approx \frac{l_u}{l_u + l} e^{-lP/\kappa}, \tag{2.14}$$

where $S_{d(u)}$ is the integrated optical density of the downstream (upstream) absorption signal; $l_u = 30$ cm is the distance from the source to the upstream probe laser; $l = 45$ cm is the

distance between the probe lasers; P is the background pressure; and κ, the characteristic attenuation length $\times$ pressure, is a complicated function of the ThO–Ne cross section for small-angle scattering, species masses, temperature, and beam velocity, derived in reference [160].[7] Since the relevant scattering cross section is not well known, we treat κ as a free parameter in the fit. We also allow for an overall coefficient α, determined from the fit, to account for possible attenuation mechanisms that are not correlated with measured background pressure. The linear dependence of the signal on the reciprocal of the distance from the source is a consequence of the fact that the absorption optical density is proportional to the number density of the molecule beam times the laser path length through the molecule beam. The former scales as $1/r^2$, where r is the distance from the source, while (in the absence of intervening collimation apertures) the latter scales as r.

From the fit in Fig. 2.3.2 we find the pressure-uncorrelated attenuation factor to be $\alpha \approx 0.8$ and the characteristic attenuation constant to be $\kappa \approx 14\,\mathrm{m} \times \mu\mathrm{Torr}$, where we have used the fact that the distance between the probe beams used to make this plot is 0.45 m and the gas correction factor for neon is $P_{\mathrm{true}} = P_{\mathrm{meas}}/0.3$. The cause of the loss factor α is unknown.

By making a series of estimates of the vacuum conductance, pumping speed, and neon scatter off of apertures along the beamline, we crudely calculated the pressure as a function of distance along the beamline and deduced that the beam transmission from source to detection region in the original pumping configuration was roughly $70 \pm 15\%$. After upgrading the vacuum system by replacing one of the two 60 l/s turbo pumps on the beamline with a 500 l/s turbo and improving the vacuum conductance to the pumps via custom large-area adapters, we estimate that the transmission should have improved to about $80 \pm 10\%$.

The key takeaway is that our vacuum quality is borderline: If the background pressure were higher by a factor of order unity, we would lose a majority of molecules between the source and the interaction region due to scatter. Future generations of ACME should therefore think carefully about any proposed upgrades that might result in diminished pumping speed,

[7] For reference, the equation is $\kappa = \sqrt{\pi}\frac{v^2 m_b}{2\sigma^{\mathrm{ThO-Ne}}} \times \left[v\sqrt{\frac{m_b}{2k_BT}}\exp\left(-\frac{v^2 m}{2k_BT}\right)\int_0^{v\sqrt{m/2k_BT}} e^{-y^2}dy\right]^{-1}$.

higher gas load, or increased beamline length.

2.4 Interaction Region

The precision measurement outlined in Section 2.1 takes place in the finely controlled environment of the interaction region. Figures 2.1.1 and 3.0.1 show schematics of the relevant apparatus, and full descriptions and diagrams can be found in references [11, 143, 180, 188]. We will briefly summarize the most critical components.

Vacuum system: The interaction region is an aluminum, o-ring-sealed vacuum chamber slightly over half a meter in length. It was welded by Atlas, Inc. [188]. Everything in and on it is made of non-magnetic materials in order to avoid creating stray fields that might interact with the molecules and cause spurious precession phases.

Magnetic field coils: Surrounding the vacuum chamber is a set of coils wound on a cylinder in an approximate cosine-θ configuration. These coils apply a uniform B-field along $\hat{z}$. There are also several sets of shim coils for smoothing out the field and intentionally applying magnetic gradients. The experiment is typically run with applied magnetic fields of 0–40 mG.

Magnetic shields: The vacuum chamber and field coils are surrounded by five nested cylinders of 1/16" thick mu-metal alloy, a high-magnetic-permeability material that magnetically shields the volume it encloses. Our set of shields provides a factor of $\sim 10^5$ suppression of Earth's (~ 500 mG) magnetic field and stray DC or slowly varying fields in the lab environment. The residual field on the beamline in the interaction region is ~ 20 μG. Our shields were fabricated by Amuneal Manufacturing Corp [188].

Electric field plates: A uniform electric field along $\hat{z}$ is provided by a pair of 43 cm long, 23 cm tall, parallel Borofloat glass plates from Custom Scientific coated on their inner surfaces with a 200 nm thick layer of the transparent conductor indium tin oxide (ITO) [188]. The inner surfaces of the plates are separated by 25 mm and are aligned with respect to each other using an interferometer built by Ivan Kozyryev and a kinematic mounting structure

developed by Amar Vutha and Emil Kirilov [11, 180]. They are secured to the mounting structure by gold-coated copper guard rings that run all the way around the edges of the plates. The electrical contacts to the voltage supply leads are made underneath the guard rings. The plates are positioned in the chamber with their faces perpendicular to $\hat{z}$; one plate is to the west of the molecule beam, and the other is to the east. The final molecule beam collimator is secured to the upstream end of the field plate mounting structure. Typical applied fields are between about 40 and 140 V/cm.

Laser systems: The 943 nm optical pumping laser and the 1090 nm state preparation and readout lasers discussed in Section 2.1 are critical to the measurement scheme and must be carefully shaped and stabilized before being sent into the interaction region. All three lasers are vertically stretched to be taller than the ≈ 1 cm molecule beam on the y-axis, while along the beamline (x) axis, they are only ≈ 1 mm or so wide. The few hundred mW of 943 nm light required to saturate the optical pumping transition is produced by a diode-laser-seeded Toptica tapered amplifier. The few W of 1090 nm light required to saturate the $H \rightarrow C$ transition is provided by a pair of diode-laser-seeded fiber amplifiers from Nufern. The seed lasers are locked to a transfer Fabry-Perot cavity which is in turn referenced to an Nd:YAG laser locked to a molecular iodine clock transition via a vapor cell. The locking system was developed by Yulia Gurevich [91], and the iodine reference was built by Dan Farkas [74]. In order to measure the asymmetry (see Eq. (2.11)) with shot-noise-limited statistics [118], the readout laser polarization is rapidly switched between $\hat{x}$ and $\hat{y}$ with a frequency of 100 kHz using AOMs [11]. Additional subtleties of the laser system are covered in references [100, 143, 180].

Light collection and detection optics: The apparatus for collecting the laser-induced fluorescence that constitutes our experimental signal consists of an array of eight lens doublets, four on each side of the field plates, that are aimed at the detection region. Each of the lenses focuses fluorescence onto the end of a fiber bundle, and the four fiber bundles on either side of the field plates are married together into a single bundle, which is coupled with optical

gel into a light pipe. The two light pipes pass out of the vacuum chamber via an o-ring feedthrough. A Hamamatsu R8900U-20 PMT, fitted with appropriate bandpass filters to transmit 690 nm light while rejecting the other wavelengths used in the experiment and the background light from computer monitors, etc. in the room, is placed at the end of each of the two light pipes. Details on the design and characterization of the light collection system are found in the theses of Nick Huzler [105] and Ben Spaun [180], and details on the PMT setups are found in the thesis of Paul Hess [100].

Experiment control, data acquisition, and logging: Needless to say, much of the hardware described here must ultimately be controlled, have its status monitored, and have its outputs read and logged by computers. While this is not, strictly speaking, part of the interaction region, it certainly belongs in any discussion of the core workings of the ACME experiment. We used a LabView 2009 based software system with NI DAQs to control the ACME experiment, acquire data, and log various control states and auxiliary parameters. Much of the code was written by Paul Hess [100], and various important aspects of the software and instrumentation are described in [91, 143, 180].

2.5 Data Analysis

> [T]ruly to enjoy bodily warmth, some small part of you must be cold, for there is no quality in this world that is not what it is merely by contrast.
>
> —Herman Melville, *Moby Dick*

Figure 2.5.1 shows some example data collected using the scheme described in Section 2.1. As derived above, this measurement scheme determines the accumulated phase due to the energy shift between the two M_J levels in a given $\mathcal{N}$ state. This energy shift is given by (see Eq. (2.4)):

$$\Delta U(\mathcal{N}, \vec{E}, \vec{B}) \equiv U(M_J = +1, \mathcal{N}, \vec{E}, \vec{B}) - U(M_J = -1, \mathcal{N}, \vec{E}, \vec{B}) \tag{2.15}$$

$$= -2g_{H,J=1}\mu_B B\tilde{B} - 2d_e\mathcal{E}_{\text{eff}}\mathcal{N}\tilde{E} \tag{2.16}$$

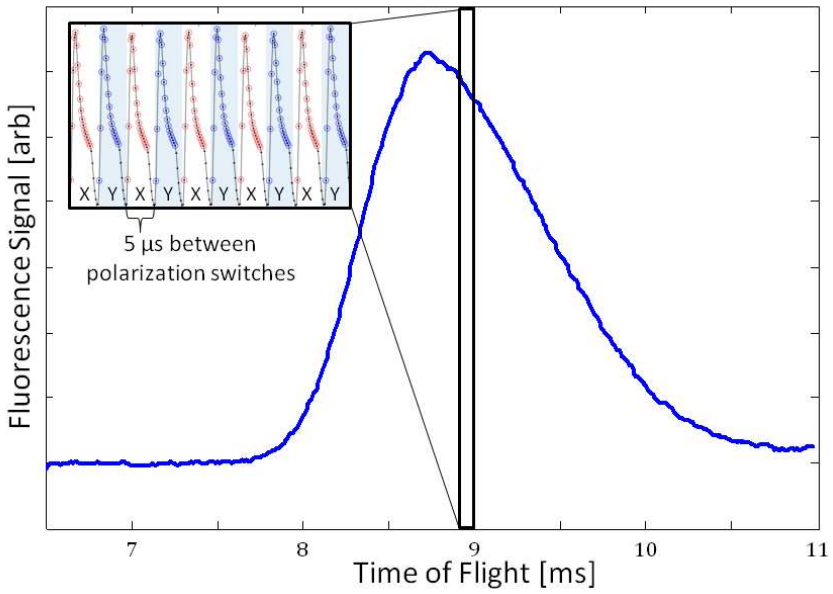

Figure 2.5.1: Average fluorescence signal profile from a molecule pulse v. time since ablation. Each pulse of molecules is ≈ 2 ms wide. Each trace on which the analysis was performed is an average of 25 molecule beam pulses and takes 0.5 s to acquire. The inset shows a zoom-in on the fluorescence signal over a 50 µs interval, revealing the 100 kHz chopping of the probe laser between $\hat{x}$ and $\hat{y}$ polarization, used to measure S_X and S_Y, respectively. An asymmetry value $\mathcal{A}$ (see Eq. 2.11) can be calculated from each adjacent pair of polarization chopping bins, and 20–30 such consecutive asymmetries are averaged together to make a measurement. These averaged asymmetries are used to construct parity sums and extract physical quantities as described in the text. Figure reproduced from reference [35], licensed under Creative Commons.

In order to isolate d_e from the Zeeman interaction term (and various other terms we have so far ignored), we perform a number of "switches." For example, we can repeat the measurement with both $\tilde{B} = \pm 1$ and take the sum of the measurements, $\Delta U(\mathcal{N}, \vec{E}, \vec{B}) + \Delta U(\mathcal{N}, \vec{E}, -\vec{B}) = -4d_e\mathcal{E}_{\text{eff}}\mathcal{N}\tilde{E}$ to eliminate terms that switch sign with $\tilde{B}$—as the eEDM term does not. We can then take the difference of the measurements to isolate the magnetic field interaction, $\Delta U(\mathcal{N}, \vec{E}, \vec{B}) - \Delta U(\mathcal{N}, \vec{E}, -\vec{B}) = 4g_{H,J=1}\mu_B B$. In other words, since the spin precession in the magnetic field is "$\tilde{B}$-odd" (reverses when $\tilde{B}$ is reversed), and the electron EDM precession is "$\tilde{B}$-even," we can distinguish these effects by taking sums or differences of precession phases measured with opposite orientations of $\tilde{B}$. Notice that we can also separate the spin and EDM precession by reversing $\mathcal{N}$ or $\tilde{E}$ since the two terms also

have opposite parity under reversal of those quantities.

In a real experiment, a number of uncontrolled effects are present, including background fields, correlated fields (e.g. magnetic fields from leakage currents which reverse synchronously with $\tilde{E}$), motional fields, geometric phases, and many more [10, 11, 115]. Despite the best experimental efforts, many of these effects cause energy shifts larger than the eEDM; however, we can isolate the eEDM from these effects using its unique "$\mathcal{N}\tilde{E}\tilde{B} = --+$" parity–i.e. odd parity under molecular dipole or electric field reversal and even parity under magnetic field reversal. Table 2.5.1 shows a sample of the types of effects that are separated from the eEDM "parity channel" by means of switches.

Table 2.5.1: Parity of energy shifts of selected effects in the ACME measurement. The difference between the g-factors of the two $\mathcal{N}$-states of H is Δg [30], and the subscript nr denotes the non-reversing component of an applied field. Products of terms denote correlations between those terms. The terms with $+--$ parity are higher-order and negligibly small. Table reproduced from reference [35], licensed under Creative Commons.

$\mathcal{N}\tilde{E}\tilde{B}$ Parity	Quantities
$+++$	Spin precession in background (non-reversing) magnetic field B_{nr}, Pump/probe relative polarization offset
$++-$	Electron spin precession in applied magnetic field
$+-+$	Leakage currents B_{leak}
$-++$	$\Delta g B_{\mathrm{nr}}$, $\Delta g B_{\mathrm{leak}} E_{\mathrm{nr}}$
$+--$	—
$-+-$	Electric-field-dependent g-factors [30]
$--+$	Electron EDM
$---$	$\Delta g E_{\mathrm{nr}}$

If we perform 8 repeated experiments, with each of the $2^3 = 8$ combinations of $\pm\mathcal{N}, \pm\tilde{E}, \pm\tilde{B}$, we can take sums and differences to compute the 8 different possible parities under $\mathcal{N}, \tilde{E}, \tilde{B}$ reversals, as shown in Table 2.5.1. Apart from higher-order terms, such as cross-terms between background electric and magnetic fields, the eEDM is the only term expected to have $\mathcal{N}\tilde{E}\tilde{B} = --+$ parity.[8] This technique of isolation by parity is how EDM experiments can

[8]To tell the whole truth, the electron-nucleon CP-violating coupling is also expected to have this behavior [10, 11]. We have ignored this term so far in the discussion because its physical interpretation is not as

perform sensitive measurements of the electron EDM with achievable levels of control of experimental parameters. We also perform a number of auxiliary switches to check for other systematic dependences of the $\mathcal{N}\tilde{E}\tilde{B} = --+$ signal, such as rotating the polarization angle of the pump and probe lasers and interchanging the positive and negative field plate voltage leads. The results of these tests are shown in Fig. 2.7.3 and discussed further in Section 2.7.

In order to avoid introducing our psychological biases into the analysis, the mean value of the $\mathcal{N}\tilde{E}\tilde{B} = --+$ eEDM channel was not revealed until the entire run was complete and we had devised a procedure for determining the systematic uncertainty. The "blind" that concealed the eEDM value was a randomly generated number drawn from a Gaussian distribution centered on zero with a width of $1 \times 10^{-27}\, e\,\mathrm{cm}$, determined by the YbF experiment eEDM limit [103]. This number, in appropriate units, was automatically added to the $\mathcal{N}\tilde{E}\tilde{B} = --+$ parity sum channel in the lowest levels of the analysis code so that researchers could only look at the eEDM value *plus* the unknown offset—never the offset or the eEDM channel alone. Significantly, this blinding method did not inhibit observation of eEDM channel *correlations* or statistical fluctuations. In order to measure and control systematic errors, it was necessary to analyze such variations in the eEDM channel; therefore, we were only blind to the critical result: the eEDM's consistency or inconsistency with zero.

The above summary is intended to give a sense for the general approach to data analysis in ACME or any EDM experiment. A thorough treatment of the ACME analysis routine and experimental switching scheme is a vast subject that is beyond the scope of this thesis. Exhaustive treatments can be found in references [11, 105, 143, 180]. Three independent analysis routines were run on the ACME Gen. I data. The three codes were developed by Nick Hutzler [105], Ben Spaun [180] and Brendon O'Leary [143]. The results of these analyses agreed well within uncertainty, and the final results reported in our limit paper [10] and discussed below were constructed from the average of the three.

intuitive as the EDM, but we will touch on it briefly in Section 2.8.

We also reserve our examination of terms that incidentally display $\mathcal{N}\tilde{E}\tilde{B}$ parity because of experimental imperfections (most of which display the "higher-order" character referenced above) for the discussion on systematic errors in Section 2.7.

2.6 Statistical Sensitivity

Averaging is the last refuge of thieves and scoundrels.

—David DeMille

The shot-noise-limited sensitivity of the ACME experiment is given by Eq. (2.13). Technical noise sources slightly reduce the achieved experimental sensitivity, but we find that our measured statistical uncertainty is within about 20% of the shot noise limit [11, 118].

Table 1.3.1 derives ACME's expected shot-noise-limited statistical EDM sensitivity from measured and calculated quantities. In this table, the average interaction time $\tau = 1.1$ ms is extracted from the Zeeman precession switch channel ($\mathcal{N}\tilde{E}\tilde{B} = ++-$): Since the magnetic field and g-factor are well known, this channel gives an accurate measurement of the precession time [11]. The contrast $\mathcal{C} = 0.94 \pm 0.02$ is determined by measuring the slope of the Ramsey fringe at $|\phi_B| = \pi/4$ by dithering the readout laser polarization [10]. The effective electric field $\mathcal{E}_{\text{eff}} = 78$ GV/cm is the unweighted mean of the two most recent theoretical calculations of this quantity [60, 174].

The count rate $\dot{N}$ can be determined directly, by converting the PMT signal to a photoelectron number, or indirectly, by starting with the measured molecule beam intensity and multiplying by the efficiency of each step in the measurement scheme. Comparing the two approaches, as in Table 2.6.1, provides an internal consistency check and allows us to determine whether there are additional signal loss factors in the experiment that we have yet to identify. In Table 1.3.1, we use the direct approach by observing that we detect ~ 1000 photoelectrons per molecule pulse [10] with a pulse rep rate of 50 Hz and a within-run duty cycle (i.e. only accounting for typical pauses and delays within a run, not the fraction of time we spend running) of $\sim 60\%$ [179] to deduce an average count rate of $\sim 3 \times 10^4$/s.

The indirect approach shown in Table 2.6.1 requires a number of additional factors. The molecule beam brightness in the absolute ThO ground state, $|X, J = 0\rangle$, was reported in [106] to be $\sim 3 \times 10^{11}$ molecules/sr/pulse. The estimated error bar on this number is a

Table 2.6.1: Observed v. expected experimental count rate.

Photoelectron count rate factors:	
Molecule beam brightness per pulse in a single $\|X, J=1, M_J\rangle$ state	$0.6\text{–}1.8 \times 10^{11}$ sr^{-1}
Solid angle subtended by detection region	$6.3 \pm 0.6 \times 10^{-5}$ sr
Pulse rate	50 Hz
Beam attenuation from background neon	0.8 ± 0.1
Rotational cooling enhancement factor	1.5–1.8
$X \to A \rightsquigarrow H$ optical pumping efficiency	0.29 ± 0.07
Number of pumped $\|X, J=1, M_J\rangle$ sublevels	2
Usable single-state population fraction in H	1/6
Spontaneous decay loss from H	0.54 ± 0.04
$C \rightsquigarrow X$ fluorescence branching ratio	0.82 ± 0.08
Fluorescence collection geometric efficiency	0.34
Losses in light collection and transmission optics	0.40
PMT quantum efficiency	0.1
Duty cycle	0.59 ± 0.06
Count rate:	
Calculated from above	$6\text{–}29 \times 10^4$ s^{-1}
Directly measured	$\sim 3 \times 10^4$ s^{-1}
Shot-noise limited statistical sensitivity:	
From calculated $\dot{N}$	$3\text{–}6 \times 10^{-29}\ e\,\mathrm{cm}\sqrt{\mathrm{day}}$
From measured $\dot{N}$	$\sim 8 \times 10^{-29}\ e\,\mathrm{cm}\sqrt{\mathrm{day}}$

factor of ~ 2. At the measured beam temperature of 4 K, the total population in the $J = 1$ state, which is used for optical pumping into H, is a factor of 2.6 larger than that in the $J = 0$ state.[9] Since the $J = 1$ state has 3 M_J sublevels, the per-quantum-state beam brightness in the $J = 1$ is $\sim 2.6 \times 10^{11}$ molecules/sr/pulse. Multiple more recent molecule flux measurements in the beam source and along the beamline, made in the spring and summer of 2012, indicated slightly lower (but still consistent with the uncertainty of the earlier measurement) $|X, J = 1, M_J\rangle$ fluxes of $0.6\text{–}1.8 \times 10^{11}$ molecules/sr/pulse.[10] Since

[9] The rotational constant in the X state is $B_e = 0.33264\,\mathrm{cm}^{-1}$ [70]. The population ratio relative to $J = 0$ is thus given by the Boltzmann factor $\exp -[hcB_eJ(J+1)/(k_BT)]$, where h is Planck's constant, c is the speed of light, J is the rotational level of interest, k_B is Boltzmann's constant, and T is the rotational temperature.

[10] One difference between the flux measurement in the ThO beam paper [106] and the 2012 measurements that might account for this small discrepancy is the ablation laser: For the beam paper measurements, we used a Continuum Minilite II YAG, fired with a rep rate of 10–15 Hz. In 2012 we switched to the 50 Hz Litron YAG used for the eEDM measurement data.

I believe the 2012 data is more representative of our actual running conditions, this is the range of values I use in Table 2.6.1.

The solid angle of the molecular beam used in the measurement is given by geometry: The final molecular beam collimator is a square aperture 10.0 ± 0.5 mm on each side and 126 cm from the beam source, so $\Omega_{\text{detect}} = (1\text{ cm})^2/(126\text{ cm})^2 = 6.3 \times 10^{-5}$ sr. The pulse rate of the YAG is set to 50 Hz. Two M_J sublevels are used for optical pumping, and a rotational population enhancement stage, briefly described in Section 2.2, increases the $J = 1$ population by 50–80%.

Of the molecules in the $|X, J = 1, M_J = \pm 1\rangle$ states that are initially on a trajectory to the detection region, the fraction that is ultimately available for detection is given by:

$$\begin{aligned}\text{Detectable molecule fraction} &= (\text{optical pumping efficiency of } X \to A \rightsquigarrow H) \\ &\times (\text{fraction of } H \text{ state population used}) \\ &\times (\exp[-\tau/(H \text{ state lifetime})]) \\ &\times (\text{Beam attenuation due to background collisions}) \qquad (2.17) \\ &= 0.29 \times 1/6 \times \exp(-1.1\text{ ms}/1.8\text{ ms}) \times 0.8 = 0.02, \qquad (2.18)\end{aligned}$$

where the first term, the optical pumping efficiency measurement, is described in [180], and the remaining terms have been discussed above.

Of the class of molecules capable of detection, about 82% emit a photon at the detection wavelength of 690 nm due to the $C \rightsquigarrow X$ branching ratio [100] (I take the uncertainty in this calculation to be about 10%). The overall 1.3% fluorescence detection efficiency is the product of the $\approx 34\%$ geometric collection efficiency of the detection lenses (which has been simulated [105] and measured [180]), the measured $\approx 40\%$ transmission through the optics that convey light out of the vacuum chamber to the PMTs ($\approx 50\%$ from the fiber bundle packing fraction and coupling losses, $\approx 85\%$ from light pipe reflections and transmission losses, and $\approx 94\%$ transmission through the PMT interference filter [105, 180]), and the

$\approx 10\%$ quantum efficiency of the PMTs [94].

The duty cycle is the fraction of the time during the run that data is being collected. ACME's Gen. I duty cycle is around 60% because of the time required to switch various parameters (e.g. laser polarization angle), degauss the magnetic shields, optimize the ablation yield, and tune up the lasers during the run.

Taken together, these factors imply an average PMT count rate between 6 and 29×10^4/s. This range is computed by adding the uncertainties from the Table 2.6.1 (dominated by the uncertainty in the molecule number and the $A \rightsquigarrow H$ branching ratio) in quadrature. Unfortunately, the measured count rate is roughly a factor of 2 below even the lower end of the range of expected count rates. The cause of this discrepancy is unknown. One possible explanation is molecule beam attenuation due to collisions within the beam or due to ballistic scatter of neon off of collimating apertures along the beamline, where we cannot measure the pressure. (This model could also help account for the unexplained attenuation α observed in Section 2.3.3 that was not correlated with pressure.) Another possibility is that the quantum efficiency of the PMTs is lower than we believe: The spec sheet [94] we use as a reference is for a slightly different PMT model than the one we actually use (R7600U as opposed to R8900U), and variations in photocathode quality from device-to-device are not unheard of, even within the same model.

As shown in Table 2.6.1, the shot-noise-limited eEDM uncertainty in the ACME Gen. I experiment computed from the measured count rate via Eq. (2.13) is about $8 \times 10^{-29}\, e\,\text{cm}$ for one full 24-hour day of averaging time. (The anticipated single-day uncertainty based on the beam source numbers and efficiency factors discussed above was about a factor of 2 smaller.) The final anticipated shot-noise-limited uncertainty can be obtained by dividing this result by the square root of the number of 24-hour days comprising the final run. As discussed in Section 2.2, we do not usually run for 24 hours at a time; however, our final data set consisted of ~ 200 hours of data taken over 2 separate $\approx$ 1-week periods in 2013 [10, 11]. This implies a shot-noise-limited uncertainty of $\sim 3 \times 10^{-29}\, e\,\text{cm}$. The actual statistical

uncertainty extracted from the scatter in the data was $4.0 \times 10^{-29}\, e\,\mathrm{cm}$ [10, 11],[11] which is consistent within uncertainty with 1.2 times the shot noise limit.

Figure 2.6.1 shows a normalized histogram of t-statistics (defined in the caption) for all the EDM data that went into our final result. The excellent fit to a normal distribution—from the center all the way out to the 4σ wings—shows that our statistics are Gaussian and unbiased to a high degree. A deep discussion of statistical subtleties in the ACME data and the various data cuts required to produce Gaussian statistics can be found in Nick Hutzler's thesis [105].

[11] This result has been adjusted from our initially reported statistical uncertainty of $3.7 \times 10^{-29}\, e\,\mathrm{cm}$ by the ratio of the old best value of $\mathcal{E}_{\mathrm{eff}}$, 84 GV/cm [177], to the new best value, 78 GV/cm [60, 174]. Since future updated $\mathcal{E}_{\mathrm{eff}}$ calculations could easily readjust these values again, a more conservative practice would be to discuss our result in terms of the measured experimental values only: i.e. $\omega^{\mathcal{N}E}$, the precession frequency that behaves like an eEDM by switching with the parameters $\mathcal{N}$ and $\tilde{E}$. For ease of interpretation, this thesis will primarily discuss our result in units of d_e and risk obsolescence (the inevitable fate of most world records and Ph.D. theses, anyway), but for reference, the theory-independent version of our result can be expressed as: $\omega^{\mathcal{N}E} = 2.6 \pm 4.8_{\mathrm{stat}} \pm 3.2_{\mathrm{syst}}\,\mathrm{mrad/s}$ or $|\omega^{\mathcal{N}E}| < 11\,\mathrm{mrad/s}$, 90% C.L. Alternatively, we can write $d_e = (-2.2 \pm 4.0_{\mathrm{stat}} \pm 2.7_{\mathrm{syst}}) \times 10^{-29}\, e\,\mathrm{cm} \times (78\,\mathrm{GV/cm})/\mathcal{E}_{\mathrm{eff}}^{\mathrm{true}}$ or $|d_e| < 9.4 \times 10^{-29}\, e\,\mathrm{cm} \times (78\,\mathrm{GV/cm})/\mathcal{E}_{\mathrm{eff}}^{\mathrm{true}}$, 90% C.L.

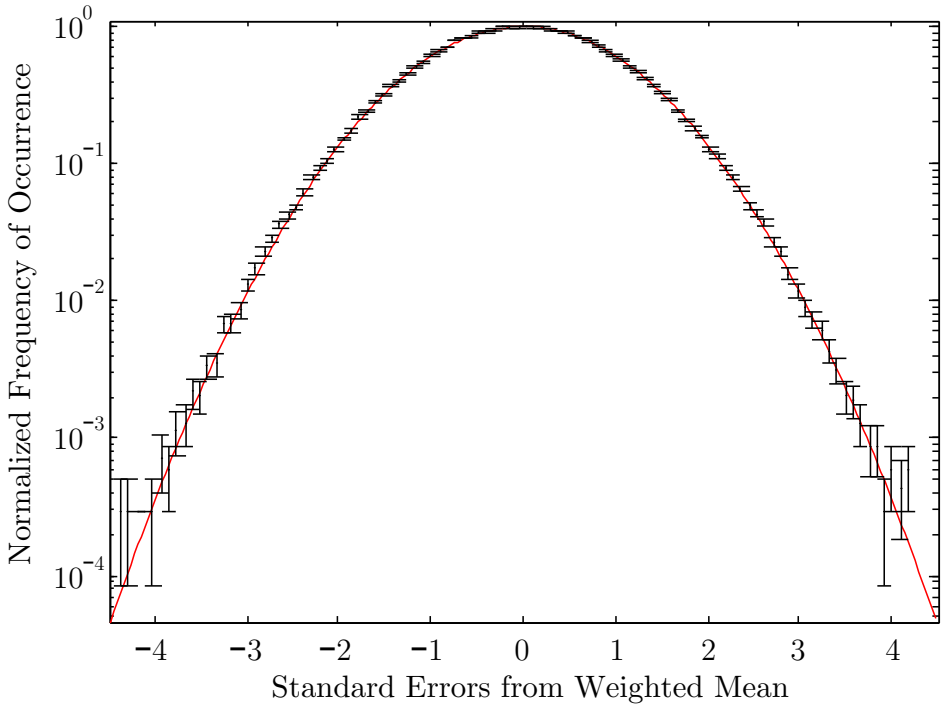

Figure 2.6.1: Log plot showing a normalized histogram of differences between the means of all $\approx 200,000$ individual eEDM measurements and the overall weighted mean in units of the sample standard error (i.e. $s/\sqrt{n}$, where s is the standard deviation of the $n \approx 20$–30 adjacent asymmetry points grouped together to make a single measurement as described in the caption of Fig. 2.5.1). In the ideal case, this distribution approaches a Gaussian with $\sigma = 1$. The red line is a Gaussian fit with $\sigma = 1.007$, showing how closely our data conforms to the ideal. Figure adapted from [105].

2.7 Systematic Errors and Uncertainty

One prefers, of course, on all occasions to be stainless and above reproach, but, failing that, the next best thing is unquestionably to have got rid of the body.

—P. G. Wodehouse, *Joy in the Morning*

As discussed above, the particular behavior of the electron EDM under reversal of applied electric field $\tilde{E}$, applied magnetic field $\tilde{B}$, and molecule electric dipole orientation $\mathcal{N}$ allows for powerful rejection of systematic effects. In order to test the limits of our ability to reject spurious eEDMs caused by experimental imperfections, we can purposely amplify these imperfections and study their effect on our measured eEDM. The key insight is that a *true* eEDM would be completely independent of experimental parameters.

Say that some physical quantity X (generally an experimental imperfection such as a stray electric or magnetic field or an uncontrolled laser polarization, frequency, or power) produces a false eEDM according to the relation $d_{e,\mathrm{false}}(X) = \alpha X$. If the uncertainty in X is δX, and the uncertainty in α is $\delta\alpha$, then the systematic uncertainty associated with X is

$$\delta d_{e,X} \approx [(\alpha\,\delta X)^2 + (\delta\alpha\, X)^2]^{1/2}. \tag{2.19}$$

The quantities X and δX can typically be determined with direct measurements (magnetometers to measure magnetic fields, spectroscopic techniques to measure electric fields, optical cavities to determine laser noise, etc.). The general technique to determine α is simply to measure d_e with intentionally varying values of X and fit the functional form of $d_{e,\mathrm{false}}(X)$. We can search for correlations of a higher order than linear by using more values of X (and thus averaging for a longer time), but for most of our purposes, the first-order term is sufficient.

Once a nonzero dependence of the eEDM channel on a parameter has been identified, we address it by (1) determining the physical origin of the effect, if possible; (2) making adjustments to the apparatus or data-taking routine to suppress or eliminate the effect (usually by minimizing α, X, or both); (3) measuring the shift αX produced by the effect under ordinary running conditions and subtracting it off in the data analysis; and (4) constructing the systematic uncertainty from the measured uncertainty and mean values of α and X as in Eq. 2.19.

Once again, the treatment provided here of the discovery, interpretation, and suppression of systematic errors in ACME is a relatively superficial one, intended to give a general sense of our approach. Thorough discussions of this subject—which is arguably the heart of any precision measurement—can be found in the theses of Ben Spaun [180], Brendon O'Leary [143], and Nick Hutzler [105], as well as ACME's recent "Methods" paper [11]. This thesis will briefly describe just one example of a systematic error that was found and corrected in

the ACME Gen. I experiment.

As shown in Fig. 2.7.1, we observed a precession frequency $\omega^{\mathcal{N}E}$ in the $\mathcal{N}\tilde{E}\tilde{B} = --+$ eEDM channel that was linearly correlated with an intentionally applied non-reversing (i.e. not switched with $\tilde{E}$) electric field $\vec{E}^{\mathrm{nr}} \parallel \hat{z}$. Following step (1) above, we varied other parameters (such as laser detunings) to see what made the effect better or worse. We eventually identified the source of the effect as an AC Stark shift induced when the molecules flew through a part of the state preparation or readout beam with a polarization ellipticity gradient. Because of angular momentum conservation, ellipticity in one of these laser beams causes it to couple more strongly to one $H, M_J = \pm 1$ sublevel than the other. If molecules interact with an elliptically polarized part of the laser that is sufficiently weak or detuned that it does not simply optically pump out an imbalanced superposition of M_J sublevels, the molecules instead experience an AC Stark shift between the two M_J levels. This generates a spurious precession-like phase. In the presence of an $\vec{E}^{\mathrm{nr}}$, this phase displays eEDM-like behavior by reversing under $\tilde{E}$ and $\mathcal{N}$ switches. The reason is that an $\vec{E}^{\mathrm{nr}}$ pointing along $\hat{z}$ *increases* the magnitude of the electric field $\vec{E}$ when it is aligned along $\hat{z}$ but *decreases* it when $\vec{E}$ points against $\hat{z}$. Thus, the linear Stark shift of the $\mathcal{N} = \pm 1$ levels is larger for one orientation of $\vec{E}$ than the other. This means that the laser detuning relative to the $H \to C$ transition changes with both $\mathcal{N}$ and $\tilde{E}$, resulting in an AC-Stark-shift-induced phase that changes with these switches, mimicking an eEDM [11].

We determined that the source of the polarization gradients producing this effect was the thermal stress-induced birefringence caused when the high-power state preparation and readout lasers passed through the glass field plates. Having identified the physical origin of the effect, we were in a good position to proceed to step (2): eliminating it. To suppress the polarization gradients, we aligned the polarization of the laser beams along the birefringence axis they were inducing in the field plates. We also introduced a chopper wheel to modulate the laser beam so that it was off between molecule beam pulses. Reducing the average laser power incident on the field plates in this fashion reduced the thermal stresses that were the

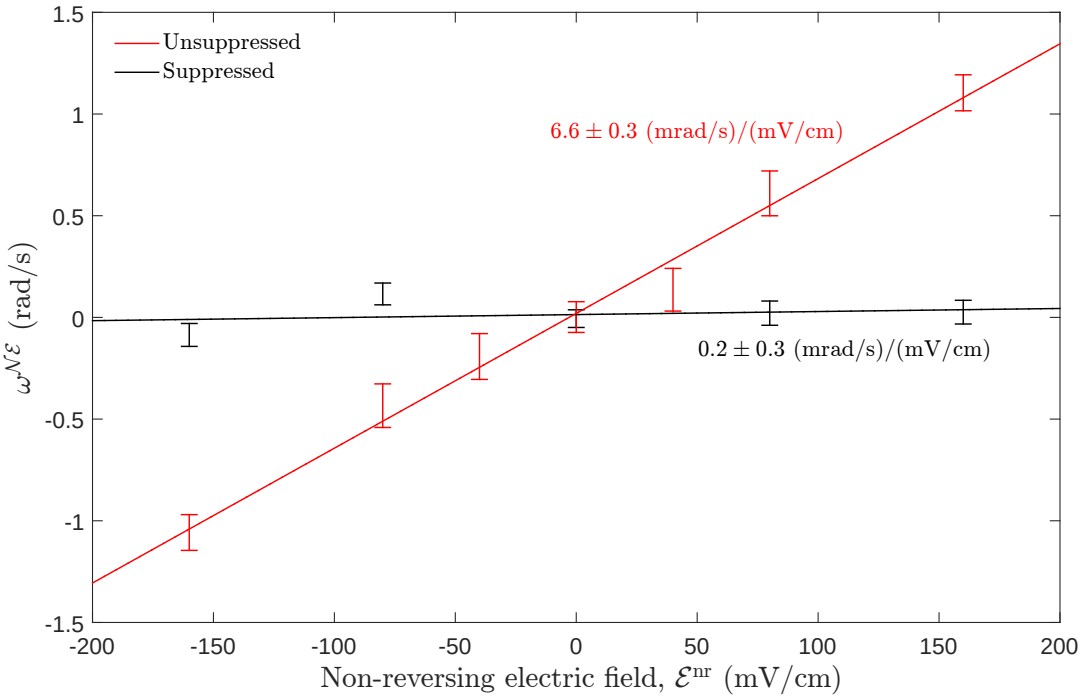

Figure 2.7.1: Dependence of eEDM channel on $\vec{E}_{\mathrm{nr}}$. Adapted from [10, 11]. ©2014, American Association for the Advancement of Science, used with permission.

source of the birefringence. The black line in Fig. 2.7.1 shows the successful result of these suppression efforts.

Having measured the slope of the effect in Fig. 2.7.1, we just needed to know the $\vec{E}^{\mathrm{nr}}$ offset during normal running conditions in order to subtract off any residual false eEDM, as prescribed in step (3). This was done via a microwave pumping measurement of the DC Stark shift, described in reference [11]. The background non-reversing electric field was measured to be about 5 mV/cm.

Finally, we completed step (4) by using the uncertainties in $\vec{E}^{\mathrm{nr}}$ and the remaining slope of $\omega^{\mathcal{N}E}$ v. $\vec{E}^{\mathrm{nr}}$ to determine the uncertainty contributed by this systematic effect. The results for this systematic and the others we investigated are shown in Fig. 2.7.2, which summarizes the systematic error budget of the Gen. I ACME experiment. The final systematic error bar was $\Delta\omega^{\mathcal{N}E}_{\mathrm{syst}} = 3.2\,\mathrm{mrad/s}$, computed by taking the quadrature sum of the individual uncertainties in the table. Using the fact that with $\mathcal{E}_{\mathrm{eff}} = 78$ GV/cm, the conversion factor from precession

frequency to eEDM is $0.849 \times 10^{-29}\, e\,\mathrm{cm}/(\mathrm{mrad/s})$, we have $\Delta d_{e,\mathrm{syst}} = 2.7 \times 10^{-29}\, e\,\mathrm{cm}$.

Parameter	Shift (mrad/s)	Uncertainty (mrad/s)	
$\mathcal{E}^{\mathrm{nr}}$ correction	-0.81	0.66	Real systematics: Caused by known experimental imperfections ameliorated in Generation II
$\Omega_{\mathrm{r}}^{\mathcal{NE}}$ correction	-0.03	1.58	
$\omega^{\mathcal{E}}$ correlated effects	-0.01	0.01	
$\omega^{\mathcal{N}}$ correlation		1.25	Expected to improve with experimental statistics and better measurements of experimental imperfections
Non-reversing $\mathcal{B}$-field ($\mathcal{B}_z^{\mathrm{nr}}$)		0.86	
Transverse $\mathcal{B}$-fields ($\mathcal{B}_x^{\mathrm{nr}}, \mathcal{B}_y^{\mathrm{nr}}$)		0.85	
$\mathcal{B}$-field gradients		1.24	
Prep./readout laser detunings		1.31	Expected to improve with experimental statistics
$\tilde{\mathcal{N}}$ correlated detuning		0.90	
$\mathcal{E}$-field ground offset		0.16	
Total Systematic	-0.85	3.24	
Statistical Uncertainty		4.80	
Total Uncertainty		5.79	

Figure 2.7.2: Gen. I systematic error budget. Adapted from [10, 11]. ©2014, American Association for the Advancement of Science, used with permission.

For three of the systematics we studied (the first three listed in Fig. 2.7.2), we applied a correcting shift. These were "true" systematic errors, in which we identified a statistically significant correlation between a false eEDM and some other variable. The dominant systematic was the E^{nr} correlation described above. The second leading effect was a false eEDM correlated with an $H \to C$ Rabi frequency $\Omega_r^{\mathcal{N}E}$ with $\mathcal{N}\tilde{E}\tilde{B} = --+$ parity. Such a Rabi frequency produces an eEDM-like precession phase via an AC Stark shift, in much the same manner as the E^{nr} effect described above [11]. We observed a nonzero $\Omega_r^{\mathcal{N}E}$ of a magnitude close to 1% of the total Rabi frequency, though we never fully understood its source, so we subtracted off the shift and included the uncertainty in our error budget.[12]

[12]For the $\Omega^{\mathcal{N}E}$ systematic, the uncertainty is larger than the shift despite $\Omega_r^{\mathcal{N}E}$ itself being statistically significant. This is largely because roughly half the data was taken with each sign of the laser propagation direction $\hat{k}\cdot\hat{z} = \pm 1$, a parameter with which $\Omega_r^{\mathcal{N}\mathrm{E}}$ was found to reverse [11] (see Section 3.2 for further discussion). Other contributing factors include the fact that the uncertainty on the correlation coefficients between $\Omega_r^{\mathcal{N}E}$ and $\omega^{\mathcal{N}E}$ is fairly large [143], and partly because in this instance, our policy of averaging the results of the three analysis codes (as discussed in Section 2.5) pushed the shift value closer to zero than some of the individual analysis results [105].

The third "true" systematic was the correlation between the eEDM channel $\omega^{\mathcal{N}E}$ and ω^{E}, the precession frequency that reverses with the $\tilde{E}$ switch (which can arise, for example, from leakage currents between the electric field plates: see Table 2.5.1). Such ω^{E} precession frequencies are largely nulled out of the eEDM channel by the $\mathcal{N}$ switch (a highly advantageous feature of eEDM experiments with Ω-doublets); however, the nulling is slightly imperfect due to a small but nonzero difference in g-factor between the two $\mathcal{N}$ states [30, 151]. Under our normal experimental running conditions, this g-factor difference is about a part per thousand, which gives a measure of the suppression of leakage currents and other $\tilde{E}$-correlated effects provided by the $\mathcal{N}$ switch. We monitored ω^{E} during our run and found that it was consistent with zero. Nevertheless, since its correlation with $\omega^{\mathcal{N}E}$ was significantly different from zero (i.e. $\alpha = \Delta g_{H,J=1}/g_{H,J=1}$), we subtracted off the shift arising from the central value of ω^{E} and included the uncertainty in our error budget. Since the shift was consistent with zero within uncertainty, we expect that this systematic error is statistics-limited: i.e., it will continue to average down as our statistical sensitivity improves and is not expected to be a limiting systematic in Gen. II.

For the other systematics in the error budget table, we did *not* observe a shift in the eEDM value, but we nevertheless included the error bars on these effects in our accounting. These effects were included for a variety of principled reasons, both physically and historically motivated, which are discussed in references [11, 105]. As indicated in Fig. 2.7.2, this latter class of uncertainties is limited by our ability to measure the imperfections with which they are correlated and by the statistics of the ACME Gen. I experiment. Thus, there is no reason to suppose that they will impose a fundamental limit on the reach of future-generation eEDM measurements. For the former class, we must take steps to reduce the systematic effects further in order to prevent them from limiting our ultimate sensitivity. Some of the measures we have taken to suppress this class of systematics are discussed in Chapter 3.

Finally, Fig. 2.7.3 shows a number of parameter variations we performed which produced

no statistically significant correlation in the eEDM channel, thus ruling out major potential classes of systematic effects. Descriptions of these parameter variations are provided in the caption. The variation in the magnitude of the electric field is notable: Since this is the first eEDM limit produced in a fully polarized system, this is the first time it has been possible to do a systematic check using vastly different E-field magnitudes without sacrificing sensitivity.

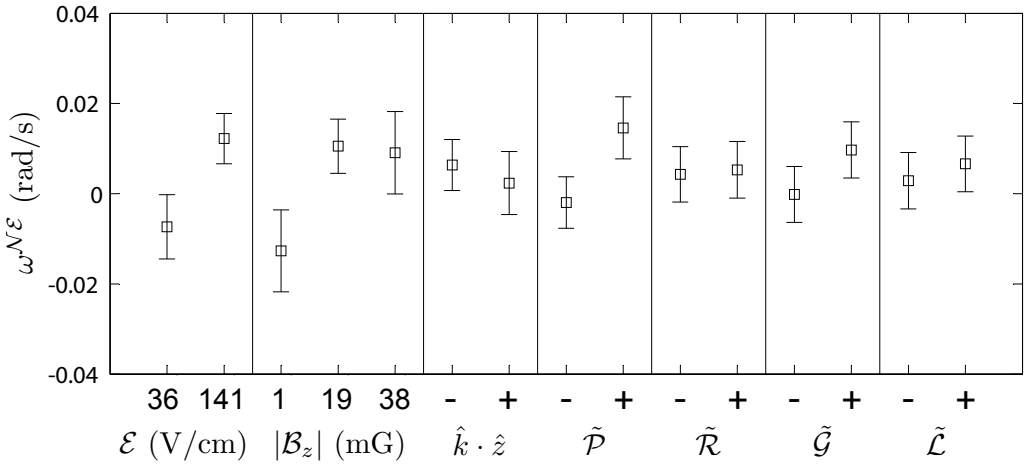

Figure 2.7.3: Average eEDM channel frequency values and error bars for different laser and field configurations. Three different values for the magnetic field magnitude and—unprecedentedly for an eEDM experiment—two widely different values for the electric field magnitude were used. The direction of laser propagation through the experiment ($\hat{k}\cdot\hat{z}$) was switched, as were the parity of the readout excited state ($\tilde{\mathcal{P}}$) and the physical connections between the voltage supply terminals and the electric field plates ($\tilde{\mathcal{L}}$ for "lead switch"). In addition, the polarization of the readout laser beam was rotated through an angle of 90°, both by itself (designated $\tilde{\mathcal{R}}$ for "readout polarization") and synchronously with the state-preparation laser beam ($\tilde{\mathcal{G}}$ for "global polarization"). Happily, no statistically significant correlation was observed between any of these switches and the measured value of the eEDM channel, allowing us to rule out many potential sources of systematic error. Figure reproduced from [10, 11]. ©2014, American Association for the Advancement of Science, reproduced with permission.

2.8 Results

No longer from head to foot than from hip to hip:
She is spherical, like a globe;
I could find out countries in her.
—Shakespeare, *The Comedy of Errors* (III.ii.105–7)

Putting the uncertainties from Sections 2.6 and 2.7 together with the central value revealed when we unblinded, the measured result was[13]

$$\omega^{\mathcal{N}E} = 2.6 \pm 4.8_{\text{stat}} \pm 3.2_{\text{syst}}\ \text{mrad/s, or} \tag{2.20}$$

$$d_e = (-2.2 \pm 4.0_{\text{stat}} \pm 2.7_{\text{syst}}) \times 10^{-29}\, e\,\text{cm}. \tag{2.21}$$

The error bars in these equations manifestly overlap with zero, indicating that we have not observed an eEDM. This result sets a significantly more stringent bound on the maximum size of the eEDM than the previous best limit [103].

Under the most general interpretation, our experiment is sensitive to any P- and T-violating interaction that produces a precession frequency shift $\omega^{\mathcal{N}E}$. The eEDM is not the only such predicted interaction for diatomic molecules [124]: a P- and T-odd nucleon-electron scalar-pseudoscalar interaction would also manifest as an $\mathcal{N}\tilde{E}$-odd spin precession in our experiment. Thus, we write

$$\omega^{\mathcal{N}E} = -d_e \mathcal{E}_{\text{eff}} + W_{\text{S}} C_{\text{S}},\ ^{14} \tag{2.22}$$

where W_{S} is a (calculated) energy scale specific to the species of study [60, 65, 66, 174, 177] and C_{S} is a dimensionless constant characterizing the strength of the T-violating nucleon-electron scalar-pseudoscalar coupling relative to the ordinary weak interaction.

We can use our measurement to set an upper limit on d_e by assuming that $C_{\text{S}} = 0$ and that $\omega^{\mathcal{N}E}$ is therefore entirely attributable to the eEDM. Adding the statistical and systematic

[13] This section is adapted from [11]. I gratefully acknowledge the ACME collaboration's contributions and editing assistance. In particular, Adam West helped to compose much of this section, and David DeMille and Nick Hutzler did yeoman's work figuring out various sign conventions in the equations.

[14] Note that the sign of the C_{S} term is opposite to that used, incorrectly, in our original paper [10]. In addition, the quantity W_{S} quoted here differs in magnitude from the related quantity $W_{\text{T,P}}$ given explicitly in [60, 174]. A detailed discussion of the sign and notational conventions for this Hamiltonian is provided in [11].

uncertainties from Eq. (2.21) in quadrature allows us to interpret our result as:

$$d_e = (-2.2 \pm 4.8) \times 10^{-29}\, e\,\mathrm{cm} \tag{2.23}$$

$$\Rightarrow |d_e| < 9.4 \times 10^{-29}\, e\,\mathrm{cm}\ (90\%\ \mathrm{C.L.}), \tag{2.24}$$

where the second line is obtained by appropriately scaling the upper bound on $\omega^{\mathcal{N}E}$ derived in Appendix D.

If, instead, we assume that $d_e = 0$, our measurement of $\omega^{\mathcal{N}E}$ in ThO can be restated as a measurement of C_{S}. Using an unweighted mean of the most recent calculations of the interaction coefficient, $W_{\mathrm{S}} = -2\pi \times 282\,\mathrm{kHz}$ [60, 174], we obtain:

$$C_{\mathrm{S}} = (-1.5 \pm 3.2) \times 10^{-9} \tag{2.25}$$

$$\Rightarrow |C_{\mathrm{S}}| < 6.2 \times 10^{-9}\ (90\%\ \mathrm{C.L.}). \tag{2.26}$$

At the time our result was published, this was an order of magnitude smaller than the existing best limit set by the ^{199}Hg EDM experiment [184], and it is still a factor of 2 smaller than the recently improved limit from the same group [87].

I knew the record would stand until it was broken.

—Yogi Berra

3

ACME Generation II

Our 2014 eEDM limit has inspired intense efforts in the theory community to account for the absence of an eEDM at the $10^{-28}\,e\,\mathrm{cm}$ level and determine which variants of proposed beyond-SM theories remain viable (some example publications from 2016 alone: [2, 3, 12, 16, 26, 41, 45, 72, 81–83, 108, 110, 120, 130, 138, 145]). Insofar as it is possible to discern a consensus view in this variegated enterprise, there seems to be a general sense that the next few orders of magnitude below the Gen. I ACME limit will prove very interesting. As discussed in Chapter 1, the search for the eEDM is already the most sensitive probe of many promising sectors of new physics, and an experiment with significantly improved sensitivity must either observe an eEDM or cut a still-deeper swath in the remaining parameter space of viable theories.

Towards this end, we have spent the past few years developing and implementing improvements to the ACME apparatus. Our goal was to improve our sensitivity by at least a factor

of 10, with a target one-day statistical uncertainty of $\delta d_e \sim 10^{-29}\, e\,\mathrm{cm}$.[1] We also took steps to ameliorate the "E^{nr}" systematic, which produced a shift at the $10^{-29}\, e\,\mathrm{cm}$ level, and to further investigate and control the still-unexplained "$\Omega^{\mathcal{N}E}$" systematic.

The upgraded apparatus is diagrammed in Fig. 3.0.1. Around the end of 2016, we began to take preliminary data to test the system and search for systematic effects. The overall measurement approach is essentially the same as that described in Section 2.1 and will not be reiterated. Some of the main improvements and early results are described in this chapter.

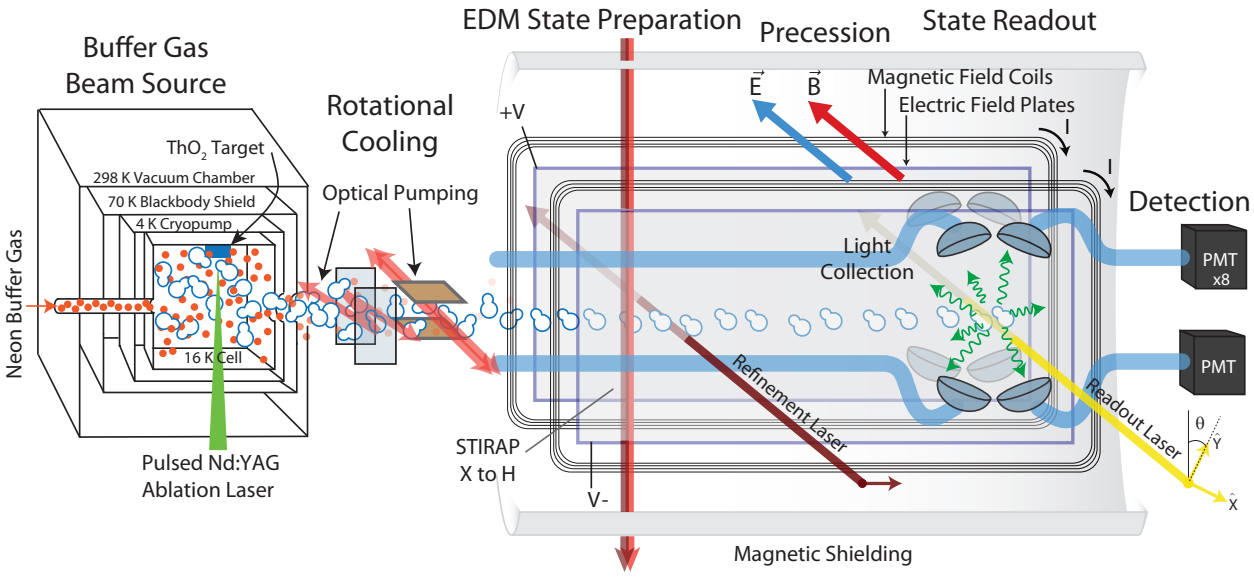

Figure 3.0.1: Schematic of the ACME Gen. II apparatus showing the modified rotational cooling, state preparation, and detection schemes. Adapted from a figure by Brendon O'Leary.

3.1 Statistical Improvements

According to Eq. 1.9, the statistical uncertainty scales as $1/\sqrt{\dot{N}}$, so increasing our sensitivity by 1 order or magnitude necessitates increasing our count rate by 2 orders of magnitude. By

[1] As with much of this thesis so far, the efforts described in this chapter were spearheaded by other members of the group. Cris Panda, Brendon O'Leary, Adam West, Zack Lasner, Vitaly Andreev, and Daniel Ang took leading roles in various projects, building on work by Paul Hess [100], Emil Kirilov, Nick Hutzler [105], Ben Spaun [180], Amar Vutha [188], Yulia Gurevich [91], and others. Jacob Baron, Christian Weber, Grey Wilburn, Cole Meisenhelder, and a number of wonderful undergraduates also contributed. Investigations by Yulia Gurevich (on G-state preparation [91]), Ana Malagon (on hybrid photodetectors), and Dave Naylor (on light collection) helped rule out promising-seeming paths that turned out to be dead ends. My own role was ancillary.

inspecting Table 2.6.1, we can see where there is room for improvement over Gen. I.

One option that we eventually ruled out was changing the beam source. Upgrading to the higher-flux, quasi-CW thermochemical source described in Chapter 5 would yield a substantial signal improvement. However, it would also require changing targets more frequently, dramatically revising our experimental timing structure, and learning the quirks of a new molecule beam, which could create new technical and statistical challenges. We have therefore decided to do without the count rate enhancement provided by the thermochemical source for the present and are still using the ablation source described in Section 2.2. The thermochemical source will be installed either in a future generation of ACME or sooner, if we determine that we need the additional signal or longer beam pulses in Gen. II.

Increasing the repetition rate of the pulsed YAG ablation laser was another rejected option: The per-pulse signals start to diminish above the current rep rate of 50 Hz, so improvements in the average count rate would be marginal up to about 100 Hz and then would trail off altogether. This slightly improved count rate would come at the cost of a modified timing structure, of which we were wary; a minor decrease in the instantaneous signal-to-noise ratio; and most decisively, the risks involved in running our finely tuned beam source and notoriously finicky YAG in an unexplored regime.

If it were possible to detect a single photon from every molecule emitted by the source, the count rate would be $\sim 10^{13}$ s^{-1}, 9 orders of magnitude larger than in Gen. I. Thus, instead of increasing the total number of molecules at our disposal, we have instead focused our efforts on using our existing beam more efficiently.

3.1.1 Beamline Geometry

The largest single loss factor arises from the small solid angle subtended by the detection region: Our beam source has a divergence of 0.35 sr FWHM [106], while the final collimator in Gen. I only subtends 6.3×10^{-5} sr, producing a loss factor of ≈ 5000. For a time, we considered refocusing the molecules from the source into the detection region by using a quadrupole electrostatic lens, and the apparatus for doing so was designed and built. Unfortunately, we

soon discovered that the voltages required to effectively focus ThO (±30 kV) also produced up to ∼ 60 keV Bremsstrahlung x-rays as an unhealthy side effect [194]. This prompted us to explore other options for increasing our capture fraction. We quickly realized that simply by moving the beam source closer to the interaction region while simultaneously increasing the spacing between the field plates and the aperture area of the final collimator, we could achieve approximately 1/3 of the gains that an electrostatic lens *plus* a similar geometry improvement would have offered, with the added bonus that it wouldn't dose everyone in the lab.

In the end, we chose to shorten the beamline by a factor of 1.07 and increase the final collimator width and height by a factor of 2.4. Since the solid angle subtended by the detection region scales as the square of these quantities, the overall signal gain from the beam geometry change was $1.07^2 \times 2.4^2 \approx 7$. The final Gen. II beamline configuration is shown in Fig. 3.1.1.

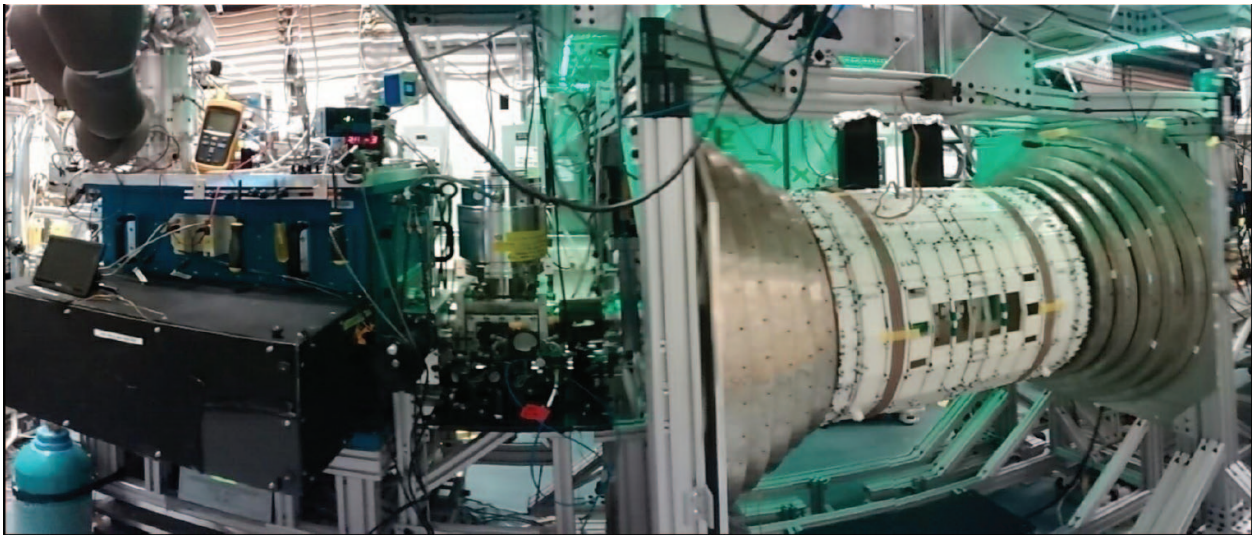

Figure 3.1.1: ACME Gen. II apparatus. Photo by Brendon O'Leary.

The beamline length was chosen to be as short as possible while still allowing for a minimum-length intermediate vacuum region (the "stem") between the source and the interaction region. The stem was required to provide (1) a bellows to mechanically isolate the in-

teraction region from the beam box and allow for precise beam alignment (see Section 2.3.2), (2) a gate vale on both the upstream and downstream ends for modularity and to isolate the cryogenic region from the interaction region during de-icing (see Section 2.2), (3) a region with ≈ 10 cm of optical access for performing rotational cooling (see Section 2.3.1), and (4) room for a large turbo pump mounted directly to the chamber to reduce beam attenuation (see Section 2.3.3; because the inlet of a 550 l/s turbo pump is typically 10 cm in diameter, this last condition is essentially automatically met if condition (3) is satisfied). The fact that we observed no unexpected signal correlations with transverse beam position and velocity in Gen. I (see Section 2.3.2) gave us confidence to omit the adjustable collimator—which would add length—from this list of requirements. After designing and installing a stem meeting these specifications, the source-to-fixed-collimator distance is now 118 cm, $\approx 7\%$ shorter than in Gen. I (see Table 2.3.1). The current Gen. II beamline aperture distances and sizes are shown in Table 3.1.1.

Table 3.1.1: List of apertures along the beamline that define the molecule beam cross-section. For apertures outside the beam box, the uncertainties on the distances from the source are ≈ 2 cm.

Distance (cm)	Width (mm)	Description
0	5	Circular cell exit aperture, 0.5 mm thick
2.5	6	Circular knife-edge conical collimator in 4 K shield
11.5	10	Circular aperture in 50 K shield
34.5	25	Square differential pumping aperture
118	24	Final beam-defining square collimator
164	≈ 40	Field plate exit: width defined by guard rings

The final collimator width was set by the geometric constraint that no molecules may hit the field plates (see Section 2.3.2). In order to meet this constraint, the collimator edges must block the line of sight from the west edge of the cell aperture to the east guard ring surface (i.e., the gold-coated copper brackets on the interior faces of the field plates that secure the plates to the mounting structure). For a field plate guard ring width f at a distance L from

a beam source with diameter a, this condition is met for

$$c \leq (f + a)\frac{l}{L} - a, \tag{3.1}$$

where c is the width of the final collimator and l is its distance from the beam source. The fixed collimator and field plate exit distances are given in the first column of Table 3.1.1. The table also shows that the beam source aperture diameter is 5 mm (optimized for buffer gas beam properties [106]). We conservatively take a to be 10 mm to allow for at least $\sim$ 1 mm of misalignment (see the discussion of alignment tolerances in Section 2.3.2) and to account for the fact that ThO–Ne collisions outside the cell exit aperture slightly enlarge the effective source size [22, 106].

In order to determine the ideal field plate spacing, we had to consider the effect of (1) fringing fields, which distort the electric field in the precession region and increase with field plate separation; (2) the transverse Doppler width, which affects our ability to optically saturate the molecule beam; and (3) the efficiency of the collection optics, which decreases with molecule cloud size. Brendon used a COMSOL simulation to determine that for field plate spacings of less than 5 cm, the fringing fields in the precession region were of order 10^{-6}, more than two orders of magnitude smaller than the distortions caused by the bowed shape of the mounted field plates [143, 180]. We also found that the beam's transverse velocity width for such a spacing produced a 1σ Doppler width of just a few MHz, well within our ability to saturate the transitions discussed below in Sections 3.1.2 and 3.1.3. Thus, the limiting consideration turned out to be the light collection efficiency, which, in simulations performed by Cris Panda, was found to fall off rapidly above a field plate spacing of 4.5 cm due to the decrease in the maximum solid angle coverage of the collection optics. We therefore chose a field plate spacing of 45 mm in Gen. II, which resulted in a slightly smaller guard ring spacing of $\approx$ 40 mm. Substituting these numbers (f = 40 mm, a = 10 mm, l = 118 cm, and L = 164 cm) into Eq. 3.1 gives a maximum collimator width of c = 26 mm. We chose a

width of 24 mm. We also made the height 24 mm because Cris found that the efficiency of the collection optics starts to fall off for molecule cloud sizes above 20 mm, so there was no reason to make the beam larger in the vertical direction.

The effect of this change in beamline geometry on the spatial and velocity profile of the molecule beam in the detection region is shown in Fig. 3.1.2. The transverse position and velocity distributions are much more flat-topped in Gen. II than in Gen. I, and the effect of transverse velocity dispersion is significantly more pronounced. The transverse spatial (velocity) width of the molecule beam in the detection region is $\approx$ 33 mm ($\approx$ 5 m/s) in Gen. II, as compared to $\approx$ 15 mm ($\approx$ 2.5 m/s) in Gen. I.

3.1.2 STIRAP State Preparation

As discussed in Section 2.6, the low $X \rightarrow A \rightsquigarrow H$ optical pumping efficiency combined with the fraction of usable H-state sublevels populated in the process renders our state preparation fairly inefficient, with only $\approx 5\%$ of addressed ground-state molecules ending up in the eEDM state.

A far less lossy approach is to use a coherent state transfer scheme such as STIRAP (STImulated Raman Adiabatic Passage) [27] to move population from one pure state to another in a controlled manner, without relying on spontaneous processes as in optical pumping. We have implemented a STIRAP scheme in the Gen. II ACME experiment that transfers population from the absolute ground state $|X, J = 0, \mathcal{P} = +1\rangle$ to the $|X_{\mathcal{N}}\rangle$ quadrature state in H (defined in Eq. 2.3) via a two-photon Λ-structure transition with $|C, J = 0, \mathcal{P} = -1\rangle$ as the intermediate excited state [144] (see Fig. 1.3.3 for level structure and transition wavelengths). The main results are summarized in Fig. 3.1.3.

The key insight in STIRAP is that in a three-level Λ-structure system where the lower states $|1\rangle$ and $|3\rangle$ are eigenstates of the unperturbed Hamiltonian, if a radiation field is turned on that couples state $|3\rangle$ to the excited level $|2\rangle$, the "dark"—or uncoupled—state $|1\rangle$ remains an eigenstate. Moreover, in the presence of coherent radiation fields coupling both $|1\rangle$ and $|3\rangle$

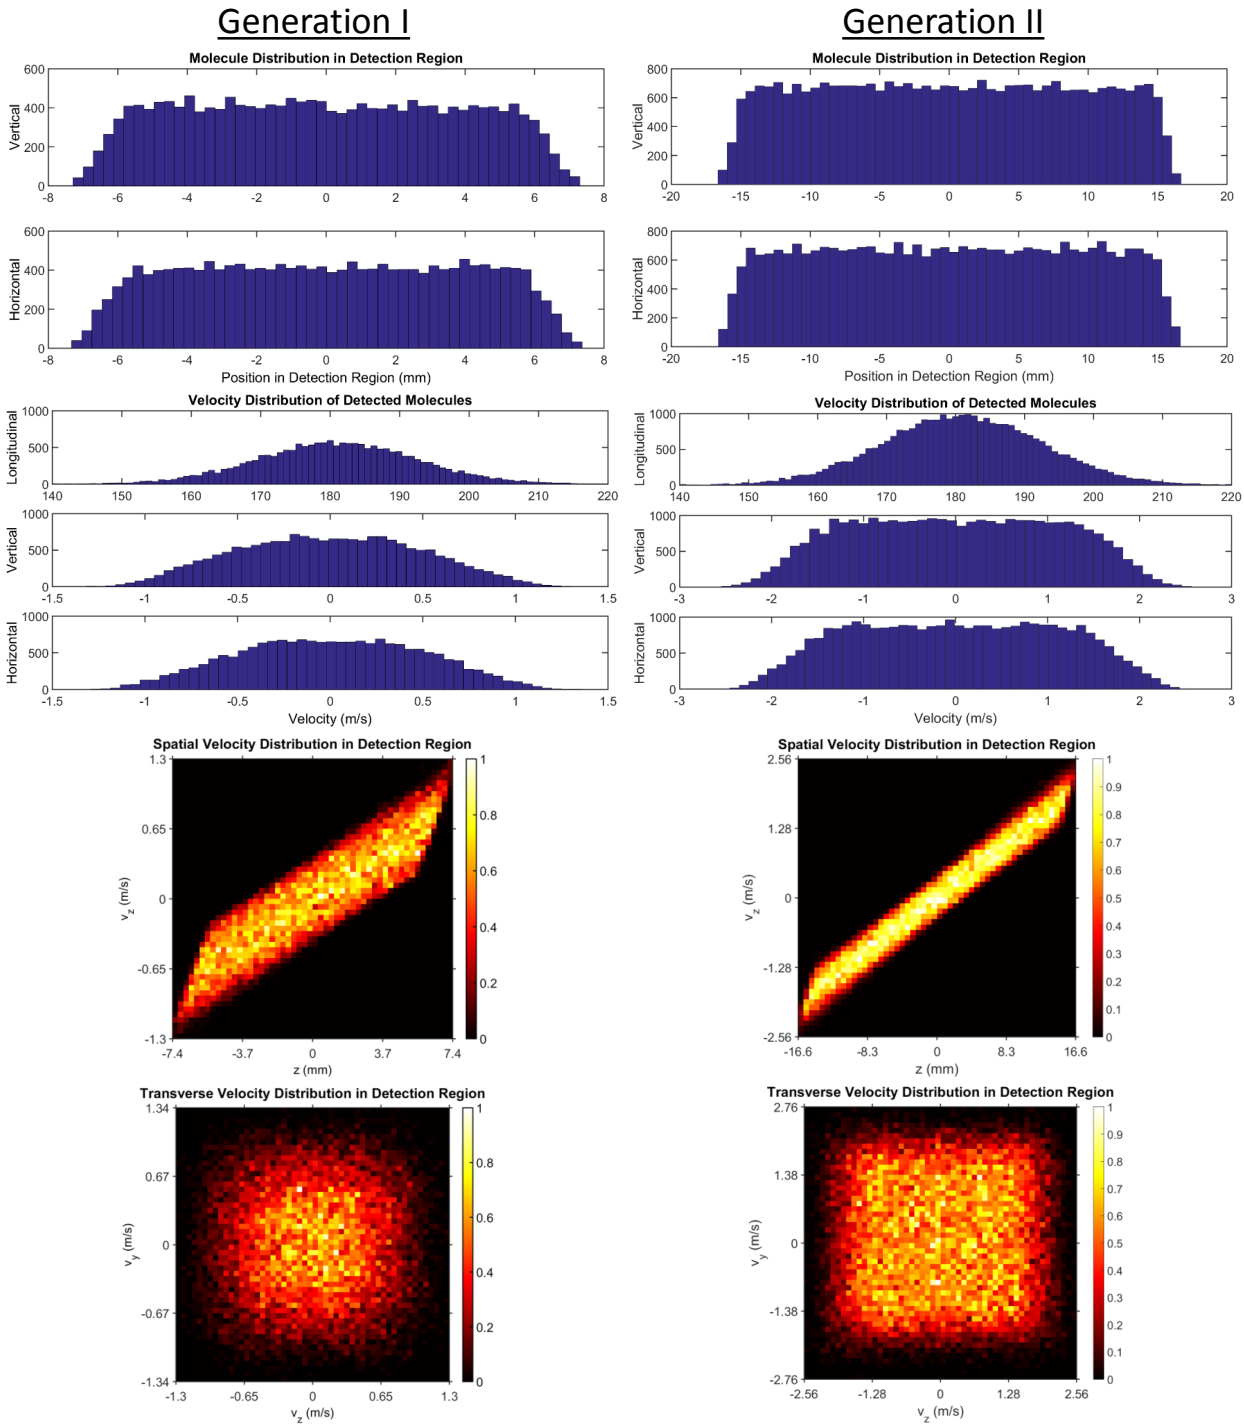

Figure 3.1.2: Simulations of position and velocity distributions in the detection region for the Gen. I (left, $\approx$ 20 thousand trajectories) and Gen. II (right, $\approx$ 30 thousand trajectories) beam geometries. The beamline aperture sizes and distances are shown in Tables 2.3.1 and 3.1.1, respectively. The detection region is $\approx$ 36 cm downstream (i.e., in the $+\hat{x}$ direction) of the final collimator. These simulations assume that the source is a circular disk of diameter 7 mm that uniformly emits an ensemble of ThO molecules with a 4 K thermal distribution in their center-of-mass frame and a center-of-mass velocity of 180 m/s.

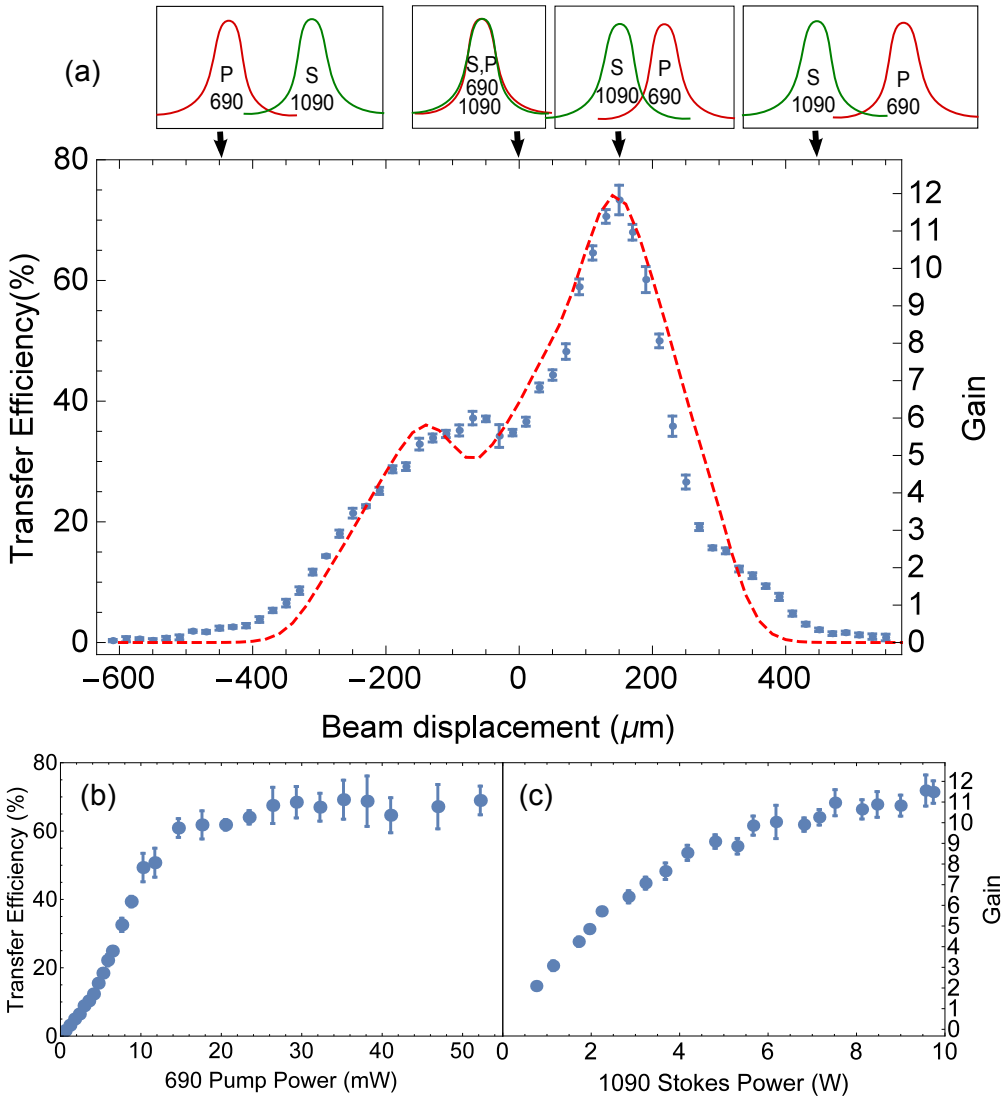

Figure 3.1.3: Efficiency of the $X \rightarrow C \rightarrow H$ STIRAP state preparation scheme as a function of (a) beam displacement between the pump (P) and Stokes (S) lasers, (b) $X \rightarrow C$ pump beam power, and (c) $C \rightarrow H$ Stokes beam power. The peak efficiency achieved is about 75%, representing a factor of ≈ 12 improvement over the efficiency of the optical pumping scheme described in Chapter 2. As expected, this peak efficiency occurs when the Stokes pulse precedes but partly overlaps with the pump pulse. In the reverse configuration (pump preceding but overlapped with Stokes), ordinary two-photon transfer dominates. The lower plots show that both pump and Stokes laser powers are near or above saturation. Figure and results from Cris Panda [144]. ©2016, American Physical Society, reproduced with permission.

to $|2\rangle$, if we choose the appropriate rotating frame where $|1\rangle$ and $|3\rangle$ are degenerate, any dark superposition of these two states is also an eigenstate. Applying this insight to the problem of coherent state transfer, the essential idea of STIRAP is to administer appropriately timed optical fields coupling $|3\rangle \leftrightarrow |2\rangle$ and $|2\rangle \leftrightarrow |1\rangle$ such that population in the initially dark state $|1\rangle$ will adiabatically follow the dark state as it smoothly transforms from $|1\rangle$ to $|3\rangle$ without ever occupying the excited state $|2\rangle$ [42]. This is done by first turning on a "Stokes" pulse to couple (typically unpopulated) states $|3\rangle$ and $|2\rangle$, and then gradually ramping it off while simultaneously ramping on a "pump" pulse that couples (populated) state $|1\rangle$ to state $|2\rangle$. The pump pulse is then ramped off, as well. If the optical fields are sufficiently strong, the transfer is sufficiently gradual, and the difference between the two radiation frequencies precisely matches the difference (in frequency units) between the non-rotating-frame energies of states $|1\rangle$ and $|3\rangle$,[2] population can be moved from state $|1\rangle$ to state $|3\rangle$ with near-perfect efficiency.

In 2015–16, building on previous work by Emil Kirilov and others, Cris Panda led an effort that demonstrated $\approx 75\%$ transfer efficiency in ThO $X \to C \to H$ STIRAP in the ACME Gen. II apparatus. The Stokes and pump pulses were applied by two tightly focused, slightly overlapped CW lasers addressing the $H \to C$ and $X \to C$ transitions, respectively (see Fig. 3.0.1). Molecules flying through the beams experience them as pulses in the time domain, as described in the previous paragraphs. A direct measurement of the gain relative to the Gen. I scheme showed a factor of 12 improvement [144].

A number of auxiliary upgrades worked on by various members of the team helped to make this achievement possible. For example, the $\Delta M_J = 0$, $|X, J = 0, \mathcal{P} = +1\rangle \to |C, J = 0, \mathcal{P} = -1\rangle$ pump transition requires z-polarized light, implying that the STIRAP laser beams must be aligned transverse to the applied electric field (see Fig. 3.0.1).[3] This

[2] This condition is required to enforce the "degeneracy in the rotating frame" condition mentioned above. It is typically expressed the literature by saying that "the two-photon detuning must be zero" for efficient STIRAP.

[3] The requirement that the STIRAP lasers must travel along $\hat{y}$ combined with the fact that the excited state parity is fixed as $\mathcal{P} = -1$ in order to have a dipole-allowed transition with the $\mathcal{P} = +1$ ground state also

meant that our lasers would have to pass vertically through the interaction region, which necessitated additional optical access through the top and bottom of the magnetic shields and the interaction region. Cutting the shields was a particularly arduous process, partly because of their size and fragility, and partly because most traditional machining approaches would cause unacceptable internal stress in the material, ruining the annealing and locking in permanent magnetic domains. Abrasive waterjet cutting is a relatively low-mechanical-stress process (or rather, it is only locally stressful, since it essentially involves rapidly grinding through the material), so Adam West led an expedition out to EBTEC Corp. in Agawam, MA, where the machinists cut the extra holes we required using a huge 5-axis waterjet.

We also needed a robust and reasonably ergonomic system to align and launch the STIRAP lasers downward through the top of the interaction region. With the help of local contractor Steve Drummey, we designed and his team built a personnel platform—dubbed the "laser lounge"—made out of 8020 T-slot framing and suspended from a pair of I-beams that we raised across the lab. We also built a separate, rigid 8020 structure supporting an optics breadboard for the STIRAP beam-launching optics. The breadboard was designed to be raised and lowered through the floor of the laser lounge platform on a set of high-precision linear guide rails. This structure and various accessories are described in the thesis of Brendon O'Leary [143].

In addition, STIRAP imposes stringent requirements for laser frequency stability and linewidth. To provide the necessary control, the lasers were referenced to a temperature-controlled, ultra-low expansion (ULE) glass cavity in vacuum via a Pound-Drever-Hall locking scheme [100, 144]. Anti-reflection (AR) coated laser diodes were also found to reduce the laser phase noise dramatically outside the two-photon linewidth of the STIRAP transfer. Switching to AR-coated diodes may have improved the coherence of the process, leading to higher ultimate signal gains [144].

means that unlike in the Gen. I optical pumping scheme, the $|Y_{\mathcal{N}}\rangle$ state (see Eq. 2.3) cannot be prepared. (See Section 2.1 for an explanation.) This should pose no fundamental problem for the measurement, but it does remove one possible experimental "switch" for ruling out systematic errors.

While our STIRAP scheme does prepare a pure state, it is also liable to produce large AC Stark shifts and other effects that could combine with other experimental imperfections to create spurious eEDM signals [143]. Therefore, after STIRAP state preparation, we re-project the state using a linearly polarized optical pumping beam (the "refinement beam") addressing the $H \to I$ transition (see Fig. 1.3.3). This beam plays an identical role to that of the $H \to C$ state preparation laser in Gen. I (see Section 2.1) by pumping out any residual population in the "wrong" initial state. Brendon O'Leary's thesis discusses the optimal refinement beam parameters and the consequences of omitting it [143]. The titanium:sapphire laser used to produce the refinement beam is the same as that used for the readout beam, and it will be discussed further in Section 3.1.3.

Finally, the rotational cooling scheme (see Section 2.3.1) was revised to pump molecules into the $J = 0$ state in X, the initial state used for STIRAP, rather than the $J = 1$ state, which was used in the Gen. I optical pumping scheme. This upgrade was developed and optimized by Zack Lasner. The overall signal enhancement it provides is comparable to or slightly higher than that observed in Gen. I (see Table 2.6.1): i.e., a factor of roughly 2.

3.1.3 Light Collection and Detection

Two significant improvements were made in the efficiency with which we collect and detect signal photons in Gen. II: (1) switching from 2 fiber bundles to 8 curved light pipes to transfer photons out of the interaction region and (2) detecting with laser-induced fluorescence on the $H \to I \rightsquigarrow X$ transition rather than on $H \to C \rightsquigarrow X$. Both of these improvements will now be described in turn.

As shown in Table 2.6.1, in Gen. I, $\approx 60\%$ of the fluorescence photons collected by the lens doublets in the detection region were lost before they were incident on the PMTs. Most of this loss derived from the $\approx 50\%$ packing fraction of the fiber bundles onto which the light was focused [11, 105, 180]. In Gen. II, in addition to re-optimizing the geometry of the collection optics for the new, larger molecule cloud, Cris Panda designed a light transfer scheme without fiber bundles in order to eliminate this loss factor. The final setup is shown

in Fig. 3.1.4.

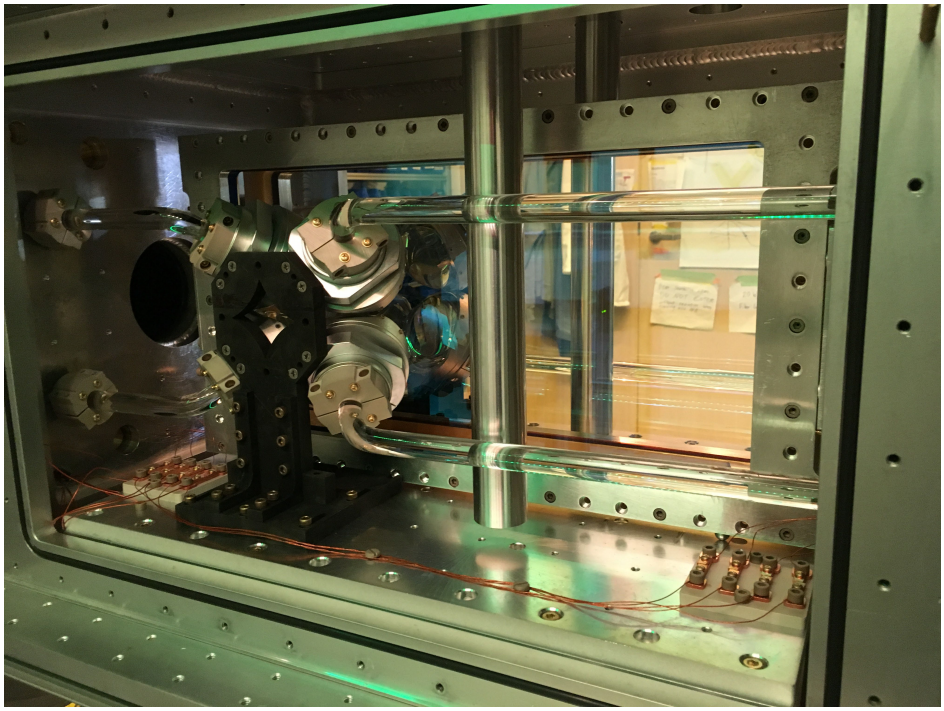

Figure 3.1.4: Light collection optics setup for Gen. II. The geometric collection efficiency of the lenses is slightly lower for the larger field plate spacing, but curved light pipes replace limited-packing-fraction fiber bundles, providing an overall factor of ≈ 1.8 in improved collection efficiency. Photograph by Cris Panda.

In the new scheme, fluorescence collected by each of the eight lens doublets is focused onto a single curved light pipe custom-made by Yale's glassblower Daryl Smith. A vacuum-compatible PEEK (polyether ether ketone) clamp secures the curved end of the light pipe to the lens mount and aligns it with the focus, while a similar clamp couples the other end to a straight light pipe, which passes via a feedthrough out of the vacuum chamber and ends at one of eight PMTs. Cris measured the light transfer efficiency of this setup to be $\approx 2.3\times$ that of the Gen. I fiber bundles or about 90%.

After re-optimization, Cris estimated the light collection efficiency of the lens array to be roughly 0.8 times that in Gen. I because of the smaller fractional solid angle subtended with the new field plate spacing. Thus, the combined gain of the solid light pipe photon transfer method and new collection optics geometry is ≈ 1.8 relative to Gen. I.

We also changed readout excitation transitions to the 703 nm $H \to I$ line investigated in references [121, 122]. The light source is an M Squared SolsTiS titanium:sapphire (Ti:Sapph) laser pumped by a Lighthouse Photonics Sprout. The Ti:Sapph beam is split into two and used for both refinement and detection. This transition provides several advantages over the Gen. I $H \to C$ line. First, it has a slightly higher branching ratio to the ground state (91% [121]) by a factor of ≈ 1.2. Second, the $I \rightsquigarrow X$ decay fluorescence is at 512 nm, a markedly more favorable wavelength for PMT quantum efficiency than the Gen. I 690 nm $C \rightsquigarrow X$ fluorescence. With the Hamamatsu R7600U-300, the quantum efficiency at this wavelength is $\approx 25\%$ [95], about a factor of 2.5 higher than the Gen. I quantum efficiency. This improvement is somewhat mitigated by the fact that the field plates transmit only about 75% as much 512 nm light as 690 nm light [143]. Combining these gain factors for the readout transition gives an overall improvement relative to Gen. I of $1.2 \times 2.5 \times 0.75 = 2.3$.

Another benefit of the I state is that the $H \to I$ transition is significantly stronger than the $H \to C$ transition, meaning that it can be saturated with several times less power. This is of great importance in suppressing the E^{nr} systematic described in Sections 2.7 and 3.2. Also, unlike the 690 nm C state fluorescence, the 512 nm I state fluorescence can easily be filtered from *all* the other laser wavelengths used in the experiment, including the 690 nm lasers employed in rotational cooling. A final advantage (or at least, potential advantage) of the I state for detection is that its higher energy gives it a decay rate $\approx 4\times$ that of the C state (whose lifetime is 490 ns [100]). This characteristic decay time sets the minimum measurement time, or equivalently, the minimum time between polarization switches. This faster decay rate thus gives us the freedom to increase the polarization switching frequency in Gen. II—which we have done, to 200 kHz—and look for unexplained effects on a shorter timescale.

3.1.4 Improvement Summary and Signal Comparison

A summary of the statistical improvements described above is given in Table 3.1.2. We anticipate an overall signal increase of ~ 350 relative to Gen. I, which corresponds to a

statistical sensitivity increase of ≈ 19 in the shot noise limit.

Table 3.1.2: Summary of statistics improvements in ACME Gen. II.

Improvement	Signal Gain
STIRAP state preparation	12
Molecule capture solid angle	7
Readout transition	2.3
Light collection	1.8
Total	≈ 350

A preliminary signal trace taken under Gen. II running conditions is shown on a log scale in Fig. 3.1.5 with a Gen. I fluorescence trace shown for comparison. Rotational cooling had not yet been implemented when this trace was taken, yet the signal gain was already over 300. Based on this early result, we tentatively expect that the ultimate signal gain may be as large as ~ 500.

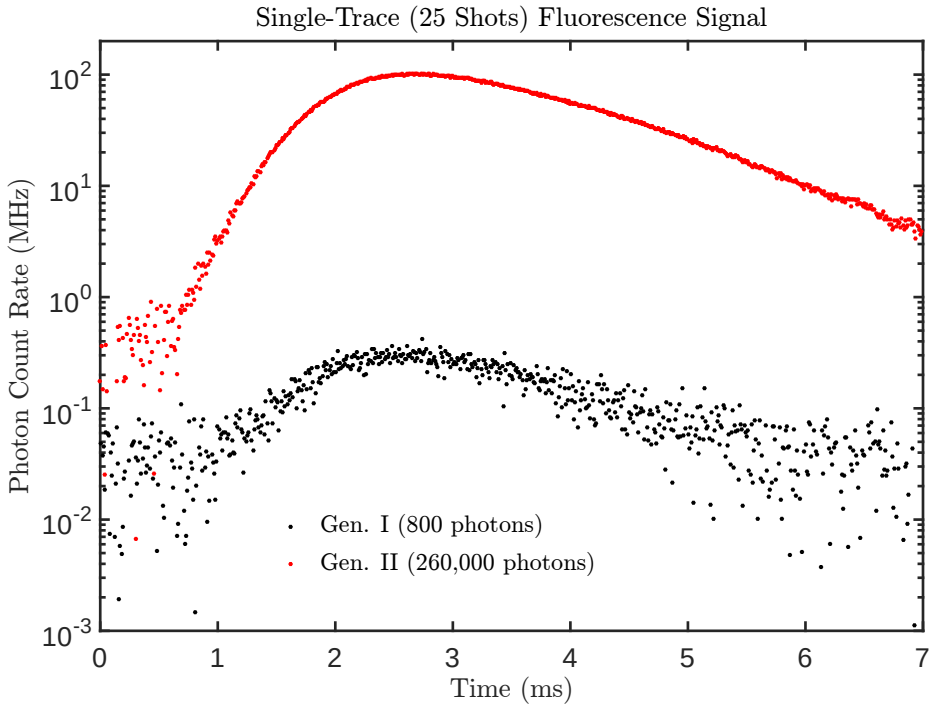

Figure 3.1.5: Single-trace count rate comparison between ACME Gen. I and Gen. II. Adapted from a figure by Cris Panda.

3.2 Systematic Error Suppression

In addition to improving our statistical sensitivity, we also took steps to suppress possible sources of systematic error in Gen. II. In the case of the dominant E^{nr} systematic in Fig. 2.7.2, we both ameliorated the known source of the effect and improved our ability to monitor the experimental imperfections that conspire to produce it. For the other effects, including the nonzero and still-mysterious $\Omega^{\mathcal{N}E}$ effect and the remaining effects that did not produce a statistically significant shift in the eEDM channel, we made improvements where possible to improve our sensitivity to the physical quantities with which they were correlated. Otherwise, we must trust that better experimental statistics will allow us to average these systematic uncertainties down—or reveal nonzero correlations that will help us identify their cause and fix them.

As discussed in Section 2.7, the E^{nr} systematic was ultimately caused by a conspiracy of two experimental imperfections: (1) a thermal stress-induced birefringence gradient in the electric field plates caused by laser absorption and (2) a stray, non-reversing electric field at the few-mV/cm level in the interaction region. Patch potentials at the few-mV level are difficult to avoid, so our efforts on imperfection (2) focused on improving our ability to measure E^{nr} regularly and accurately rather than making it go away. Thanks to efforts by Adam West and others, the microwave pumping E^{nr} measurement apparatus alluded to in Section 2.7 [11] is now set up on a fairly permanent basis, and the data-taking routine is automated so that mapping the non-reversing E-field can be a part of our regular running routine.

The physical basis for imperfection (1) above is more easily addressable. Nick Hutzler has developed an exhaustive treatment of thermal stress-induced birefringence in glass in his thesis [105]. The crucial observation is that if ϵ is the ellipticity of the state preparation or readout laser, $\Delta\epsilon$ is the change in ellipticity caused by the field plate birefringence, and w_x is the waist of the laser beam along x (the narrow axis of the elongated beam), then the figure of (de)merit for the deleterious effect is $\Delta\epsilon'(w_x)$, or the polarization ellipticity gradient

along the $\hat{x}$ at a distance of one x-axis beam waist from the center of the laser profile. This is roughly the position where the beam starts to become weak enough for AC Stark shift effects to become more important than optical pumping, as discussed in Section 2.7. Nick finds that this effect has the following dependence on the materials properties of the field plates[4]:

$$\Delta\epsilon'(w_x) \propto KE\alpha_V a/\kappa, \tag{3.2}$$

where K is the stress-optic coefficient (a constant of proportionality between mechanical stress and optical retardation), E is the elastic modulus, α_V is the coefficient of thermal expansion, a is the optical absorbance of the material at the relevant wavelength, and κ is its thermal conductivity. This model was corroborated by polarimetery measurements performed by Paul Hess [100, 105]. To minimize this quantity, Adam West and Paul Hess had a new set of field plates made by Thin Film Devices (TFD) out of Corning 7980 glass, which Nick estimated should reduce $KE\alpha_V/\kappa$ by a factor of ≈ 7 relative to the Gen. I Borofloat glass plates. Since Nick also estimated that the power absorption by the field plates was dominated by the ITO layer, we also had TFD apply an ITO coating with a thickness of 20 nm, 10$\times$ thinner than in Gen. I, in order to reduce the absorption coefficient a. These improvements, combined with the fact that the higher strength of the $H \to I$ state transition relative to the $H \to C$ transition allows us to reduce the power in the refinement and readout beams by a factor of several, should suppress the E^{nr} systematic by a factor of a few hundred.

Building on work by Paul Hess and Christian Weber, Vitaly Andreev developed a highly sensitive polarimeter for measuring small fractional polarization gradients across a laser beam [13]. Using this setup, he demonstrated that even with Gen.-I-level laser powers of 2 W at 1090 nm, the polarization gradient induced by thermal birefringence in the field plates was

[4]Note that the laser ellipticity gradient *itself* is the problem, not its (dominant Gen. I) source in the field plates. If other effects, e.g. mechanical or thermal stress-induced birefringence gradients in the vacuum windows, laser optics, etc., cause a similar-sized ellipticity gradient in the laser, these effects could also produce an E^{nr}-correlated systematic [105]. Vitaly Andreev and Paul Hess have investigated such effects and have not found any at a level of concern for our Gen. II experiment scheme [13, 100].

suppressed by at least an order of magnitude, to a level consistent with the uncertainty in the measurement. The main result is shown in Fig. 3.2.1.

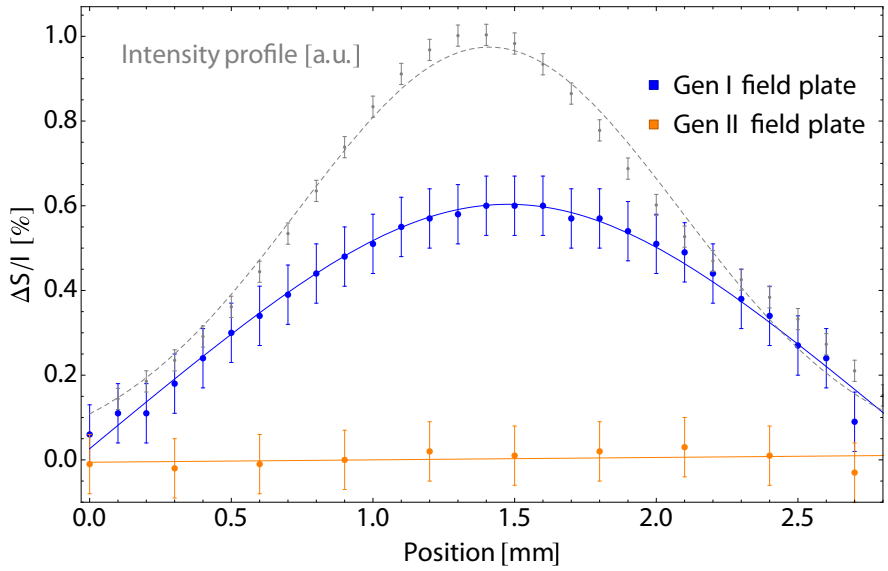

Figure 3.2.1: Stress-induced polarization ellipticity comparison between ACME Gen. I and Gen. II field plates. The intensity profile of a 2 W 1090 nm laser, elongated so that its waist is 30 mm along $\hat{y}$ and 1.4 mm along $\hat{x}$, is plotted in gray with a Gaussian fit, and the polarization profile of the laser beam after passing through the Gen. I (Gen. II) electric field plates is plotted in blue (orange). S and I are the Stokes parameters characterizing polarization circularity and beam intensity, respectively [29], so that S/I is a measure of the fractional ellipticity of the laser beam. Constant ellipticity offsets are unimportant and have been subtracted off. The blue and orange curves are fits to the photo-thermoelastic law discussed in references [100, 105]. For the Gen. II field plates, both the magnitude of $\Delta S/I$ and its gradients are suppressed by more than a factor of 10. Figure and results from Vitaly Andreev [13].

As discussed in Section 2.7, the mechanism producing an $\mathcal{N}\tilde{E}\tilde{B} = --+$ correlated Rabi frequency systematic, $\Omega^{\mathcal{N}E}$, is still unexplained. One valuable but rather confounding hint was that the effect showed a dependence on the direction of laser propagation through the interaction region, i.e. on whether $\hat{k}\cdot\hat{z} = +1$ or -1 [11, 143]. In Gen. I, reversing the laser propagation direction was a time-consuming process requiring realignment of the entire optics breadboard for launching lasers into the interaction region. Therefore, we only performed the $\hat{k}\cdot\hat{z}$ switch once during Gen. I. In Gen. II we would like to perform this switch more

frequently to gain a better understanding of its correlations with our experiment phases and to average them down more effectively, so we are building a duplicate optical setup that will allow us to reverse $\hat{k} \cdot \hat{z}$ simply by moving a few optical fibers.

Several of the larger systematic uncertainties in Fig. 2.7.2 were related to various B-field imperfections. These were included in our error budget because experiments conceptually similar to ours saw unexplained systematic shifts in the eEDM channel that were correlated with these fields (although we did not see any such effects) [11]. In the Gen. I experiment, the magnetic field was characterized by inserting a fluxgate magnetometer down the beamline once the measurement was complete [11, 143]. In Gen. II, deep pockets intrude into the interaction region so that we can insert an array of fluxgates to monitor the magnetic fields more frequently without breaking vacuum, potentially even while the experiment is running [143]. A more accurate measurement of the B-field imperfections will help to reduce the uncertainty on the correlated systematics.

3.3 Other Improvements

> The Golux groaned. "If there's one thing in the world I hate," he said, "it is amendments."
>
> —James Thurber, *The 13 Clocks*

Besides the upgrades described above, a few additional modifications were made to the ACME apparatus. These are briefly cataloged here.

- *Room magnetic field compensation coils*: In Gen. I, the outer two magnetic shields were fully saturated by the Earth's ~ 500 mG magnetic field. If the background field were nulled out, we could potentially remove these shield layers and have more flexibility with the apparatus design. For example, we could consider shortening the beamline further. As an intermediate step along this path, we have kept all the shields in Gen. II but also built an array of six magnetic field compensation coils around the walls of the lab made of 10–20 turns of welding cable. These coils are provided with a constant, passive DC current that is tuned to cancel the field around the interaction region to the

level of < 50 mG. One major benefit of this is that it provides a low ambient magnetic field for performing spectroscopy and magnetometry in and around the interaction region *without* the magnetic shields installed.[5]

- *FPGA data acquisition*: The additional six PMTs and faster fluorescence timescale of the I state discussed in Section 3.1.3 make our data acquisition requirements significantly more stringent. The extra PMTs mean that we need to record at least eight DAQ channels, while, as the fastest relevant physics timescale in our detection scheme, the fluorescence lifetime sets the minimum data acquisition rate. We have therefore upgraded our DAQ to a National Instruments (NI) PXIe-5171R field programmable gate array (FPGA) scope. Daniel Ang and Adam West led the work on testing the device, configuring it, and integrating it into our computer system.

- *Upgraded data-taking, control, logging, and notification software*: As Labview 2009 was becoming obsolescent, and NI and others were ceasing to make hardware retro-compatible with it, we migrated all of our data taking and control software to LabView 2014. In addition, Brendon O'Leary revamped many of our software systems, including the SQL server logging database and the Master Run VI, which controls experiment switches [143]. Daniel Ang also contributed many improvements, including a notification system which can alert users by email of undesirable events or status conditions.

[5] A number of fantastic undergraduate students took the lead on various aspects of this project. These students included Joe Greenberg, Joy Wang, Aaron Markowitz, and Zack Soule. Adam West, Brendon O'Leary, and I provided significant help.

I started by looking at some papers from industrial food production that simulated heat flow through various pieces of meat. It took me a while to realize there was a much easier way to learn what combinations of time and temperature will effectively heat the various layers of a steak: Check a cookbook.

—XKCD's Randall Munroe, *What If?*

4

ThO Thermochemical Source Background

We are beholden to the Cold War era U.S. obsession with nuclear technology for laying the groundwork for our thermochemical ThO source.[1] In the mid-twentieth century, thorium-232, the dominant isotope of thorium (Th, half-life 1.4×10^{10} years), was found to be fertile: Neutron capture transforms thorium-232 into fissile uranium-233, making thorium a promising breeder reactor fuel candidate [53, 54]. This discovery seized the attention of the nuclear physics community and prompted a multitude of detailed investigations into the physical [125] and chemical [7, 24, 25, 55, 56, 142] properties of thorium and its compounds. One particularly intensive research program focused on developing mechanically robust thorium dioxide (ThO_2) fuel pellets [19, 53, 93, 97, 111, 112, 125]. We benefited greatly from this body of work when fabricating the ablation targets described in Appendix B for our

[1] I wish to thank Nick Hutzler for performing the initial literature search that brought this reaction to our attention. I also wish to thank Rémi Louf, who visited our group from École Normale Supérieure in 2010: During his brief time with us, he contributed greatly to the ideas, apparatus, and results described in this chapter. His work is credited in the text where possible.

original beam source (see Section 2.2).

Our new approach to making ThO exploits a chemical reaction discovered in the course of this exhaustive mid-century investigation into thorium. At sufficiently high temperatures, thorium metal in the solid or liquid phase reacts with thorium dioxide in the solid phase to produce ThO in the gas phase. We drive this reaction in a cryogenic buffer gas beam source (CBGB) like that described in Section 2.2 by locally heating a spot on the surface of a target made of a mixture of pressed Th and ThO_2 powder to temperatures above 2000 K using a focused 50 W laser beam. The new source, which is intended to be implemented in a future iteration of the ACME experiment, is described in detail in Chapter 5. In the current chapter we discuss the chemical kinetics and thermodynamics of the reaction in Section 4.1, estimate the dominant heat transfer mechanisms and the laser power required to produce ThO in a CBGB in Section 4.2, and describe two preliminary experiments performed to test the feasibility of this production method in Section 4.3.

4.1 ThO Thermochemical Reaction

In 1959 Darnell et al. [56] set out to fill a gap in the mushrooming body of literature on metallic thorium-232 by measuring its vapor pressure. Using a radio-frequency induction coil, they heated a cylindrical slug of thorium metal of known surface area in vacuo to temperatures up to 2000 K. At various intervals, they switched off the induction coil and weighed the cylinder to measure the mass loss. A nuisance chemical reaction was soon found to be creating a large systematic effect: When ThO_2 impurities were present above the level of 30 ppm, the initial rate of mass loss from the slug was dominated by a transient lasting for a few dozen hours before the vaporization rate settled to the constant $\mathrm{Th(s)} \rightarrow \mathrm{Th(g)}$ value. The vapor pressure of ThO_2 was already known to be many orders of magnitude too small to account for this effect. Additional experiments involving deliberate oxidation of the slug revealed that the transient was caused by the previously unobserved chemical reaction

$$\mathrm{Th(s)} + \mathrm{ThO_2(s)} \rightarrow 2\mathrm{ThO(g)}. \tag{4.1}$$

With an initial ThO_2 impurity concentration of 0.8%, Darnell et al. observed a mass loss rate at 1883 K of

$$0.8\%\ \mathrm{ThO_2},\ 1883\,\mathrm{K} \implies F(\mathrm{ThO}) = 5\,\mathrm{mg/cm^2/h} = 3\times10^{11}\,\mathrm{mol.}/(100\,\mu\mathrm{m})^2/\mathrm{s}, \quad (4.2)$$

where $F(\mathrm{ThO})$ represents the flux of ThO from the surface. In the final expression, I have used the ThO mass of 248 amu and a benchmark surface area of 100 μm to convert to units convenient for comparing with a focused-laser-based source ("mol." here stands for "molecules," not "moles"). In another of Darnell's experiments, performed at the same temperature but with an initial impurity concentration of 0.14% ThO_2, the mass loss rate was

$$0.14\%\ \mathrm{ThO_2},\ 1883\,\mathrm{K} \implies F(\mathrm{ThO}) = 0.8\,\mathrm{mg/cm^2/h} = 5\times10^{10}\,\mathrm{mol.}/(100\,\mu\mathrm{m})^2/\mathrm{s}. \quad (4.3)$$

In a system like this, in which a high-temperature solid-state reaction produces a highly condensable gas-phase product that will stick to the first cool surface it encounters, the reaction does not proceed to equilibrium, because the product is lost to the environment. Under these conditions, the reaction rate is determined by the chemical *kinetics* rather than the *thermodynamics* [198]. In this case, the reaction rate takes the following form:

$$-[\dot{\mathrm{ThO_2}}] = k[\mathrm{Th}]^m[\mathrm{ThO_2}]^n, \quad (4.4)$$

where brackets indicate concentrations of the chemical they enclose (or partial pressures in the case of gaseous species), the dot indicates a time derivative, and m and n are "orders of reaction" which depend on the detailed reaction mechanism. For elementary (single-step) reactions, the orders of reaction are equal to the stoichiometric coefficients (i.e., the number of molecules of each reactant species involved in the reaction), though they often differ from this for more complex reactions [198]. The rate constant k is given by the Arrhenius equation,

$$k = Ae^{-\frac{E_a}{k_B T}}, \tag{4.5}$$

where A is the "frequency factor" accounting for the rate at which the reactants collide with a relative orientation favorable to the reaction, E_a is the "activation energy" barrier that must be overcome for the reaction to proceed, k_B is the Boltzmann constant, and T is the temperature.

In the Darnell experiment, we can assume [Th] = 1 to a very good approximation because it comprises the bulk of the solid material [198]. Equation (4.4) then reduces to the statement that the time derivative of the ThO_2 concentration is proportional to the n^{th} power of the [ThO_2] concentration. If the reaction in Eq. (4.1) is a single-step reaction involving an exchange of an oxygen atom between ThO_2 and Th, then the reaction order is the same as the stoichiometric coefficient, or $n = 1$ [198]. This model is corroborated by the Darnell [56] results shown in Eq. (4.2–4.3). We observe that the ratios between the ThO_2 concentrations and the production rates of ThO are the same (6 mg/cm^2/h per % ThO_2) for both equations and conclude that for small [ThO_2], $\dot{N}(\text{ThO}) \propto [\text{ThO}_2]$, where $N(X)$ is the number of molecules of species X in the system. The stoichiometry of Eq. (4.1) also implies that $\dot{N}(\text{ThO}) = -2\dot{N}(\text{ThO}_2)$. Furthermore, since $N(\text{ThO}_2) \ll N(\text{Th})$, we have $[\dot{\text{ThO}_2}] \equiv \dot{N}(\text{ThO}_2)/[N(\text{ThO}_2) + N(\text{Th})] \approx \dot{N}(\text{ThO}_2)/N(\text{Th})$. Putting all these results together gives $[\dot{\text{ThO}_2}] \propto [\text{ThO}_2]$; therefore, the rate scaling in Eq. (4.4) is linear and $n = 1$.[2] For small [ThO_2], we can then write the production rate of ThO as:

$$F(\text{ThO}) = A' e^{-\frac{E_a}{k_B T}} [\text{ThO}_2], \tag{4.6}$$

Where A' is the product of A from Eq. (4.5) and the constant of proportionality relating F(ThO) and [ThO_2].

[2]This conclusion is further bolstered by the exponential decay fit to the ThO vaporization rate F(ThO) in Fig. 1–2 of Darnell [56]: We have shown in the text that F(ThO) is proportional to [$\dot{\text{ThO}}$] when [ThO] is small, and exponential time-dependence arises from a linear ($n = 1$) relationship between a quantity (in this case the concentration of the ThO_2 impurity) and its time derivative.

This result is valuable because it allows us to linearly extrapolate the rate of ThO vaporization with the ThO_2 concentration, as long as the latter remains small compared to 1. For example, by scaling results in Eq. (4.2–4.3), we estimate the ThO production at 1883 K with $[ThO_2] = 10\%$ to be:

$$10\%\ \mathrm{ThO_2},\ 1883\,\mathrm{K} \implies F(\mathrm{ThO}) = 62\,\mathrm{mg/cm^2/h} = 4\times10^{12}\,\mathrm{mol./(100\ \mu m)^2/s}. \tag{4.7}$$

For comparison, the in-cell ThO production with the ablation CBGB source running at 50 Hz is estimated to be $\sim 10^{14}$ molecules ThO(g)/s [106]. Thus, with a ~ 100 μm laser spot heating a target made out of 10% ThO_2 + 90% Th to a temperature of 1883 K, we would anticipate the in-cell yields from thermally driving the reaction in Eq. (4.1) to be roughly 20 times smaller than the ablation yield.

At first blush, this result does not make our proposed approach sound terribly promising, but a closer look at the physics gives reasons for hope: According to Eq. (4.6), the rate of ThO production is proportional to e^{-E_a/k_BT}, a quantity that increases rapidly with T when $T \ll E_a/k_B$. At temperatures $T \gg E_a/k_B$, the reaction rate approaches a maximum of A'. We therefore need to know the activation energy E_a in order to evaluate whether this reaction can produce yields competitive with the ablation source at temperatures above 1883 K. In Section 4.3 we describe an experiment that provides a rough activation energy estimate of $E_a/k_B \approx 6 \pm 1 \times 10^4$ K (see Eq. (4.18)), which implies that the ThO signal should increase rapidly with T for temperatures around 1883 K. By substituting this activation energy and the Darnell experiment [56] result from Eq. (4.2) (i.e., $F(\mathrm{ThO}) = 3\times10^{11}\,\mathrm{mol./(100\ \mu m)^2/s}$, $[ThO_2] = 0.8\%$, and $T = 1883$ K) into Eq. (4.6), we can solve for A' to obtain the estimate $A' \approx 3\times10^{27}\,\mathrm{mol./(100\ \mu m)^2/s}$ under the Darnell experiment running conditions.[3] Substituting these results for A' and E_a/k_B into Eq. (4.6), we can then estimate the temperature

[3] These estimates are only valid for a limited temperature range around the values used in the experiments. For solid state reactions, the reaction rate also depends on a number of highly variable physical parameters, such as the degree of contact between the reactants in the mixture. The estimated value of A' for the Darnell experiment should therefore not be taken as a general physical result.

that a 100 μm laser spot on a 10% ThO_2 + 90% Th target is required to achieve in order to produce a continuous ThO flux 10 times larger than the time-averaged ablation source flux:

$$10\%\ ThO_2,\ F(ThO) = 10^{15}\ \text{mol.}/(100\ \mu\text{m})^2/\text{s} \implies T = 2300\ \text{K}. \tag{4.8}$$

This is roughly the temperature of an incandescent light bulb.

Before we move on to investigating the feasibility of achieving such high temperatures in a cryogenic buffer gas beam source, it is worthwhile briefly considering the equilibrium position of the reaction in Eq. (4.1). I.e., under conditions where the reaction can proceed in both directions, $Th(s/l) + ThO_2(s) \leftrightharpoons 2ThO(g)$, what is the equilibrium partial pressure of ThO as a function of temperature? The answer to this question will indicate the fluxes achievable using an oven source with a small hole to leak out ThO gas: In this scenario, the particle loss rate is small compared to the total reaction rate, so the two directions of the reaction will reach equilibrium. In this case, the equilibrium position is determined by the standard Gibbs free energy difference $\Delta G^\circ(T)$ between reactants and products, according to the equation [172, 198]:

$$K(T) = e^{-\Delta G^\circ(T)/RT}, \tag{4.9}$$

where $R = 8.314$ J/K/mole is the ideal gas constant, and K is the equilibrium constant, which for the reaction in Eq. (4.1) is given by

$$K(T) = \left(\frac{P_{ThO}(T)}{P^\circ}\right)^2. \tag{4.10}$$

In this expression, P° is the standard pressure, generally taken to be 1 atmosphere, $P_{ThO}(T)$ is the partial pressure of ThO, the exponent comes from the stoichiometric coefficient of ThO in Eq. (4.1), and we have ignored the activities of the bulk solid or liquid Th and ThO_2, taking them to be unity [6]. $\Delta G^\circ(T)$ can be computed from the standard values of entropy, enthalpy of formation, and heat capacity as a function of temperature for the species involved. These values are measured empirically and tabulated in reference [142]. A

calculation of $P_{\mathrm{ThO}}(T)$ was carried out by visiting grad student Rémi Louf, and the result is plotted in Fig. 4.1.1. Three direct measurements of $P_{\mathrm{ThO}}(T)$, performed by Hildenbrand et al. [101], Ackermann et al. [6], and Darnell et al. [55] are shown for comparison. Their measurements were carried out by placing a mixture of Th and ThO_2 in a heated cell and measuring the output from a small aperture. Over the temperature range of interest, the vapor pressure of Th (ThO_2) above the pure substances is approximately 10^2–10^3 (10^3–10^4) times smaller than the partial pressure of ThO [6, 8, 9, 55, 56]. Thus, ThO is by far the dominant gas-phase product of the reaction.

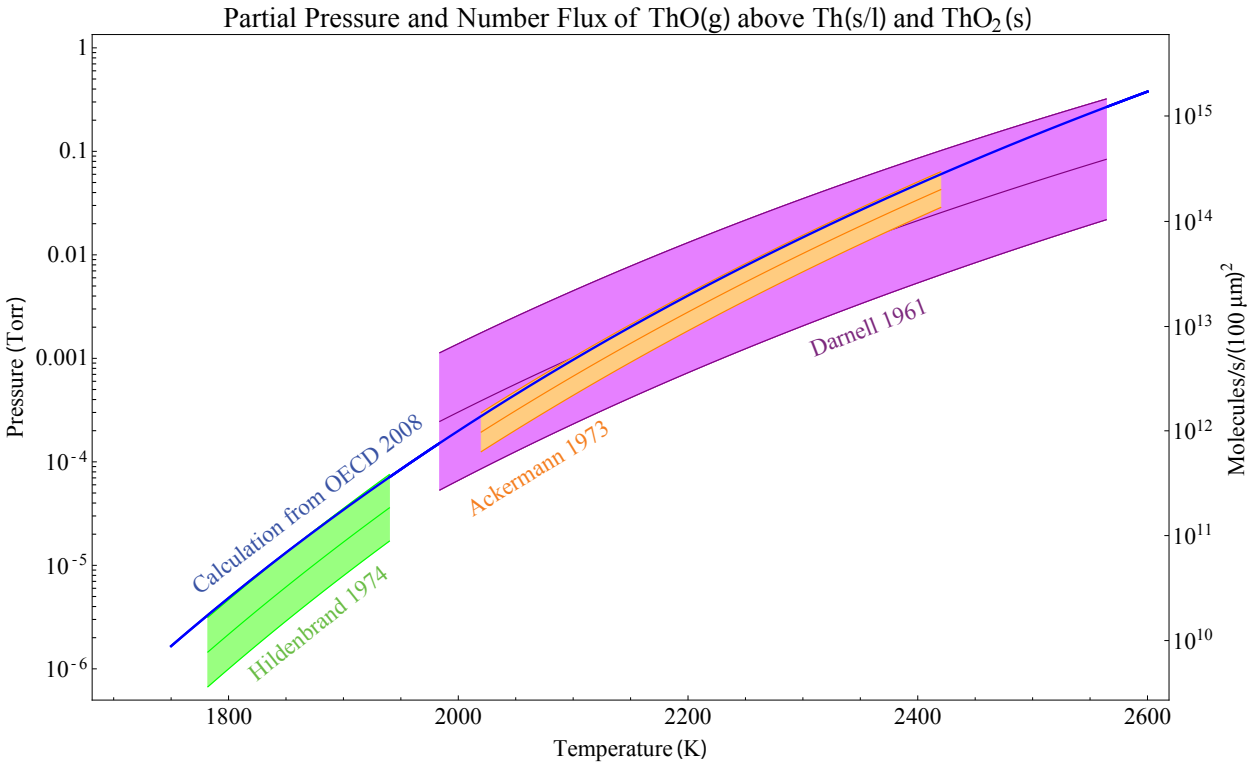

Figure 4.1.1: Literature values for the equilibrium partial pressure of ThO(g) above Th(s/l) and ThO_2(s), shown over the temperature ranges in which the measurements were performed. The results plotted here are: Darnell 1961 [55]: $\log_{10}[P_{\mathrm{ThO}}(T)/\mathrm{atm.}] = (8.16 \pm 0.47) - (35,500 \pm 1,100)/T$, Ackermann 1973 [6]: $\log_{10}[P_{\mathrm{ThO}}(T)/\mathrm{atm.}] = (7.58 \pm 0.09) - (28,630 \pm 200)/T$, and Hildenbrand 1974 [101]: $\log_{10}[P_{\mathrm{ThO}}(T)/\mathrm{atm.}] = (8.386 \pm 0.164) - (30,480 \pm 306)/T$. There is no simple analytic expression for the blue curve $P_{\mathrm{ThO}} = P^\circ \exp[-\Delta G^\circ(T)/2RT]$ (see Eq. (4.9–4.10)), calculated by Rémi Louf from data in reference [142], but for comparison, $\Delta G^\circ(2000\,\mathrm{K})/2R = 30,300\,\mathrm{K}$. On the right-hand y-axis, these partial pressures are converted into number fluxes $\dot{N}_{\mathrm{ThO}}$ from an $A = (100\,\mu\mathrm{m})^2$ orifice using Eq. (4.16).

Figure 4.1.1 suggests that an oven source of ThO at 2000 K with an $A = 1$ cm aperture would produce $\sim 10^{16}$ ThO molecules/s. Rémi Louf investigated the possibility of coupling such a source to a buffer gas cell in order to produce a cold beam of ThO, but the challenges inherent in (a) constructing an oven that would withstand such temperatures and run stably without failing and (b) dealing with the blackbody radiation power ($\propto AT^4$) such a source would emit into the cryogenic apparatus ultimately dissuaded us from this approach.

4.2 Heat Load Estimates

We propose to produce ThO(g) by heating a target composed of a mixture of Th and ThO_2 using a continuous wave (CW) laser with a focal diameter of $\sim$ 100 μm.[4] We have seen in Section 4.1 that in order to drive the thermochemical reaction (Eq. (4.1)) at a rate that compares favorably to the ablation yields, the target surface must achieve temperatures above $\approx$ 2300 K. In order for the thermochemical source molecule beam to have similar advantageous properties to the ablation source beam (see Section 2.2), this laser heating must take place in a neon buffer gas cell held at 16 K with a buffer gas density of up to $\sim 10^{17}/\text{cm}^3$ [105]. We must now inquire into the feasibility of this approach given reasonable technological constraints: Pulse tube cryo-refrigerators currently on the market have a cooling power up to $\approx$ 15 W at 16 K [51], which sets an upper limit on the laser power we may use to achieve these target temperatures.

The change in the temperature of the target surface is $\Delta T = \Delta Q/C$, where C is the heat capacity of the material and Q is the net heat flow into the target surface, i.e. $Q = Q_{\text{in}} - Q_{\text{out}}$. Because of the aforementioned cryo-refrigerator cooling power limitations, the average heat load must be $Q_{\text{in}} < 15$ W. For temperatures between 16 K and 2300 K, we must have $Q_{\text{out}} < Q_{\text{in}}$ to achieve the desired temperature increase. The target surface can lose heat via

[4] This benchmark focal size is chosen as approximately the minimum size for which the laser spot heats more than one of the $\approx$ 40 μm diameter grains of Th and ThO_2 that compose the target (see Appendix B). It is also a reasonably achievable minimum beam waist given that the focusing optics must be $\approx$ 30 cm from the target because of the vacuum chamber geometry, and that the laser must pass through a $\approx$ 2 cm diameter cell window. The expected dependence of the signal on the laser spot size will be addressed in the current section, and the measured dependence will be discussed in Section 5.5.

three mechanisms: conduction through the target body, conduction and convection in the buffer gas, and blackbody radiation from the target surface. We now estimate the loss Q_{out} due to each of these mechanisms.

4.2.1 Target Conduction

To limit the thermal conduction through the targets, we fabricate them predominantly from insulating ThO_2 and endeavor to keep the Th fraction below the percolation threshold, as discussed in Appendix B. Percolation thresholds have been calculated for a multitude of particle site configurations [195]. For randomly packed ("jammed") spheres of equal size, the percolation threshold is 31% [156]. For the case where the non-conducting grains are 1.5 times larger than the conducting grains, as in our targets (see Appendix B), simulations show that the percolation threshold is slightly lower: $\approx 27\%$ [98]. We choose a Th fraction of 25% and a ThO_2 fraction of 75%. The fact that the grains are irregularly shaped rather than spherical is ignored for the purpose of these estimates.

Before considering the conductive heat loss in the steady-state condition of a target region heated to 2300 K, let us first determine whether it is plausible for the region to *reach* these temperatures in the first place. The energy required to heat a volume of target with a characteristic length scale 100 μm can be computed by integrating the weighted mean of the Th and ThO_2 heat capacities $C(T)$ from 16 to 2300 K. The heat capacity formulas are taken from reference [142] and approximated by a linear relationship extrapolated down to 16 K. We estimate $C_{\text{Th}} \approx C_{\text{ThO}_2} \approx [0.01 \times T/(\text{K}) + 3]$ J/K/mole, so that the energy required to heat 1 mole of target material from 16 to 2300 K is ~ 30 kJ. A 100 μm diameter hemispheric region of target with the theoretical density of ThO_2 (10 g/cm^3) contains $\approx 1 \times 10^{-8}$ moles of ThO_2, so the energy required to heat just this region is ~ 0.3 mJ. For a 10 W laser, the time required to deposit this much energy is ~ 0.03 ms.

In order to determine whether the energy will be lost from the target region faster than it is

deposited, thereby quenching the heating process, we must estimate the thermal diffusivity,

$$\alpha = \frac{k}{\rho C}, \tag{4.11}$$

where k is the thermal conductivity of the target and $\rho \approx 4 \times 10^4$ mole/m^3 is the density. Both k and C depend on temperature, so for conservatism, we choose to evaluate the heat loss at a temperature where the diffusivity is high, i.e. k is large while C is small. The thermal conductivity of Th has been measured from 5–1000 K: k_{Th} increases from 2 to 60 W/m K between 0 and 20 K and then flattens off around 50 W/mK above 100 K [157]. The thermal conductivity data on ThO_2 does not extend into the cryogenic realm: $k_{\mathrm{ThO_2}}$ falls from 15 W/m K at 200 K to ~ 2 W/m K at 2000 K [186]. Since k_{Th} and $k_{\mathrm{ThO_2}}$ both have maxima in their measured values around 200 K, we evaluate the diffusivity at this temperature.[5] A weighted mean of the thermal conductivities gives $k \approx 0.75\, k_{\mathrm{ThO_2}}(200K) + 0.25\, k_{\mathrm{Th}}(200K) \approx 24$ W/m K, while the heat capacity is $C \approx 5$ J/K/mole. We therefore estimate the diffusivity in the regime of interest to be $\alpha \lesssim 1 \times 10^{-3}$ m^2/s. The characteristic time for heat to flow out of the volume of interest is then $SA/\alpha \sim 0.1$ ms, where $SA = 2\pi(50\ \mu\mathrm{m})^2/3$ is the surface area of the hemisphere being heated by the laser. Thus, the characteristic timescale of heat loss from the volume is longer than the time required to deposit enough energy to heat the volume to the desired temperature. This suggests that the laser power will win, and the target spot will be heated.

A crude, back-of-the-envelope estimate of the heat loss due to conduction through the target in the steady state can be obtained by applying Newton's law of cooling in one dimension: We assume that a spot on the target 100 μm in diameter is at a temperature $T = 2300$ K, that the direction of heat flow is primarily straight into the target, and that the length scale over which the temperature falls off is set by the beam waist. We can then

[5] According to the third law of thermodynamics, $C \to 0$ as $T \to 0$. At very low temperatures, the diffusivity is therefore necessarily large, but the energy required to heat a body at this temperature is commensurately small. We will therefore assume that the laser will rapidly heat the target surface out of the regime where the diffusivity blows up.

estimate $Q_{\text{conduction}} \sim -kT\pi w_0^2/w_0$. Taking the average value of the thermal conductivity to be $k(800\,\text{K}) \approx 17\ \text{W/m K}$ for a 25% Th + 75% ThO_2 target, we estimate $Q_{\text{conduction}} \sim 12$ W. Since this is smaller than the 15 W pulse tube cooling power, we conclude that heat loss due to conduction through the target will not prevent favorable rates of thermochemical production. Moreover, this is likely a conservative estimate for conductive heat loss since, as will be shown below, the actual temperature gradient is less steep than we have assumed here.

To get a more detailed understanding of heat diffusion through the target, we can model it as a semi-infinite solid impinged upon by a circular disk of uniform, constant laser power, as illustrated in Fig. 4.2.1. For this model, the heat equation has an analytic solution given by [38]:

$$T(r,z,t) = \frac{Q_{\text{laser}}}{2\pi w_0 k}\int_0^\infty J_0(\lambda r)J_1(\lambda w_0) \times \left[e^{-\lambda z}\text{erfc}\left(\frac{z}{2\sqrt{\alpha t}} - \lambda\sqrt{\alpha t}\right) - e^{\lambda z}\text{erfc}\left(\frac{z}{2\sqrt{\alpha t}} + \lambda\sqrt{\alpha t}\right)\right]\frac{d\lambda}{\lambda}, \quad (4.12)$$

where r and z are the radial and axial coordinates defined in Fig. 4.2.1, Q_{laser} is the incident laser power, w_0 is the radius of the laser beam at its focus, J_0 and J_1 are Bessel functions, and λ is a dummy variable that is integrated out.

In Fig. 4.2.1, we use this formula with $Q_{\text{laser}} = 9$ W and w_0 varied from 1 to 125 μm to estimate the surface temperature of a target with a thermal conductivity $k = 10\ \text{W/m K}$ and a diffusivity of $\alpha = 2.5 \times 10^{-5}\ \text{m}^2/\text{s}$. These are approximately the properties of an 800 K target composed of 90% ThO_2 and 10% Th. This model neglects the temperature dependence of k and α (The thermal conductivity decreases by a factor of ~ 2 and the diffusivity decreases by a factor of ~ 4 between 800 and 2000 K) as well as potential phase changes in the target. Nevertheless, the top plot shows that with 9 W of power in a 250 μm diameter beam, we expect to achieve temperatures of 2000 K if conduction through the body

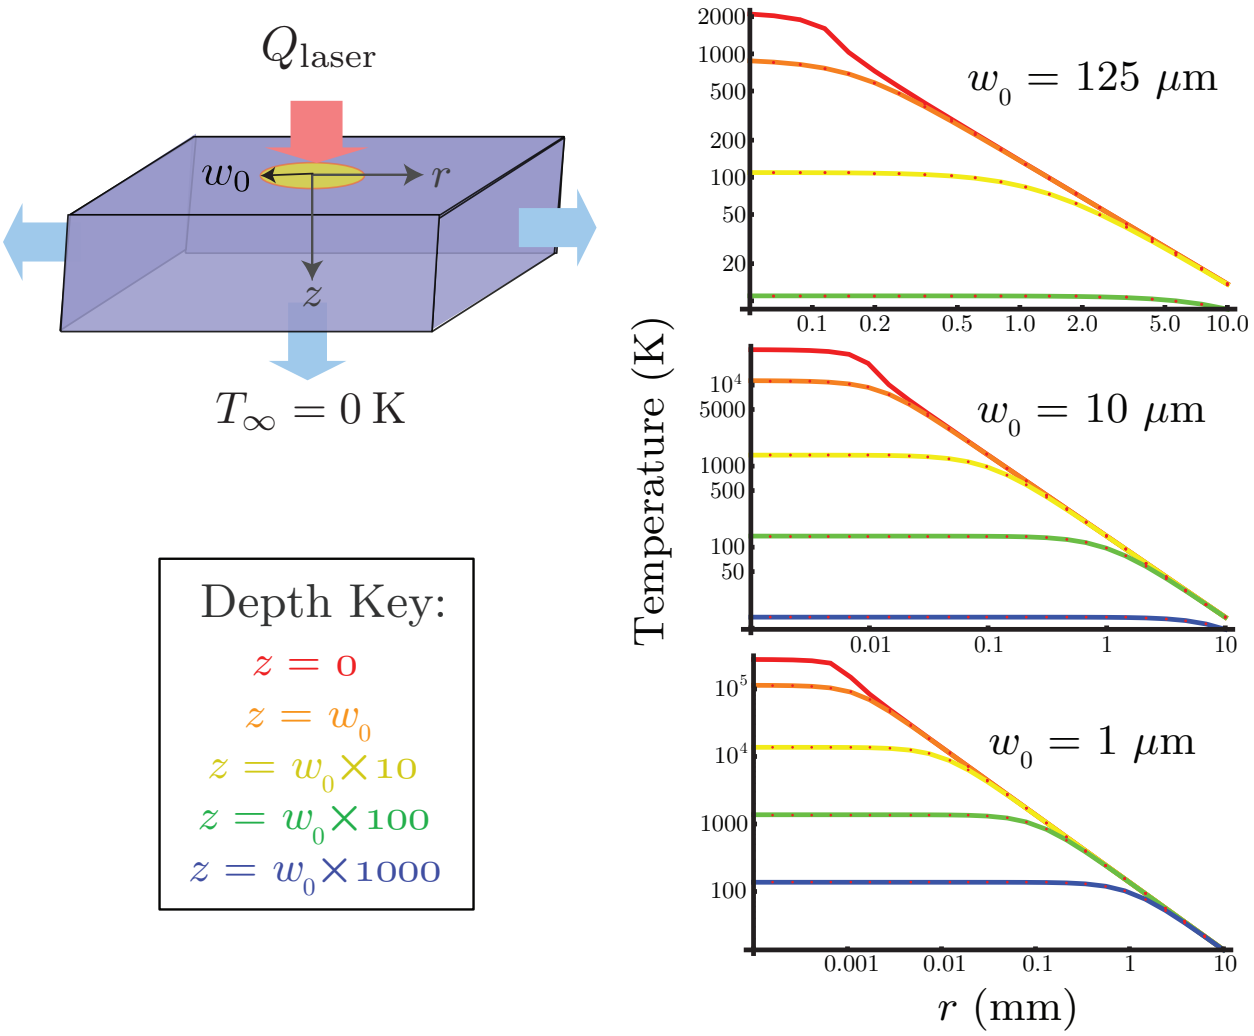

Figure 4.2.1: Top left: Illustration of the target, modeled as a semi-infinite solid, with a uniform, circular laser beam incident upon its surface. The temperature at $z = \infty$ is taken to be 0. Right: Steady-state temperature profiles in the model target at various beam waists, calculated using the highly simplified model illustrated at left and described in the text. In Eq. (4.12), we have let $Q_{\text{laser}} = 9$ W, $k = 10$ W/m K, and $\alpha = 2.5 \times 10^{-5}$ m^2/s. Each plot shows the resulting temperature profile as a function of radius and depth for a different value of the beam waist. Heat loss mechanisms other than thermal conduction through the target body are ignored for these plots.

of the target is the dominant heat loss mechanism.

The physical conditions used in this calculation are slightly different from those considered up to this point: Specifically, the temperature is lower, the spot size is larger, and the target is more insulating. Using Eq. (4.12), however, it is straightforward to scale these results to the conditions of interest, namely, a 2300 K, 100 μm diameter laser spot on a 75% ThO_2 + 25% Th target. Since the spot temperature scales linearly with the incident laser power and inversely with beam waist and thermal conductivity, and the thermal conductivity of a 75% ThO_2 + 25% Th target is ≈ 70% higher than that of a 90% ThO_2 + 10% Th target, the required power to heat the target surface to 2300 K is about $9\,\mathrm{W} \times 1.7 \times (2300\,\mathrm{K})/(2000\,\mathrm{K}) \times (100\ \mu\mathrm{m})/(250\ \mu\mathrm{m}) \approx 7$ W. We should therefore be able to achieve the intended spot temperatures without exceeding the 15 W limit set by the pulse tube cooling power, thus affirming the result of the back-of-the-envelope Newton's law calculation above.

Another observation from Fig. 4.2.1 is that for $r > w_0$ or $z > w_0$, the temperature falls rapidly below the level at which ThO production occurs at an appreciable rate. Thus, we expect target surface depletion (i.e., exhaustion of the chemical reactants by driving the reaction to completion) to occur on the length scale of a beam waist, an expectation borne out by the results in Section 5.5. Moreover, the size of the ThO production region scales as w_0^2, while the temperature scales as $1/w_0$. Since the yield scales much faster than quadratically with temperature in the vicinity of 2000 K, as discussed in Section 4.1, this suggests that the trade-off between the area and the temperature of the laser spot will favor smaller beam waists. The limit on this trade-off is expected to occur when the beam waist approaches the ≈ 40 μm diameter of the Th and ThO_2 grains: If only one of the reactants is subject to heating, the reaction will not be driven efficiently. The data in Chapter 5 do *not* show a very clear preference for smaller beam waists, suggesting that high local temperatures may produce other effects that are detrimental to the yield. We have not determined what these effects are, but some possible culprits include melting and "laser cutting" of the target surface or thermal ejection of material from the target.

Finally, by studying the time dependence of Eq. (4.12), we find that after 1 ms, a point at the 125 μm beam waist has reached 75% of its final temperature, and after 10 ms, it has reached 90% of its final temperature. These estimates are quite rough and probably only trustworthy within about an order of magnitude. If we wish to model the temporal behavior of the laser heating more carefully in the future, we will have to take into account the dependence of the physical properties on temperature, the finite size and actual thermal boundary conditions of the target, and phase transitions in the material. These estimates do suggest, however, that the heating timescale under the conditions described here will probably not exceed ~ 100 ms. This is confirmed by the data in Chapter 5.

Using two different methods, we have now seen that a 100 μm spot on the target surface that is held at 2300 K will lose heat through the body of the target at a rate of ~ 10 W. As long as the other two heat transfer mechanisms, blackbody radiation and conduction through the buffer gas, do not exceed a few watts, we may conclude that the proposed thermochemical source is feasible.

4.2.2 Blackbody Radiation

We next consider blackbody radiation, which is given by the Stefan-Boltzmann law as

$$Q_{\text{blackbody}} = \epsilon\sigma\pi(SA)^2T^4, \tag{4.13}$$

where ϵ is the emissivity, $\sigma = 5.67 \times 10^{-8}\ \text{W/m}^2/\text{K}^4$ is the Stefan-Boltzmann constant, and SA is the heated surface area of the target. The rapid scaling of the blackbody emission rate with temperature ensures that the heat loss will be dominated by the hottest part of the target surface, which (as shown in Fig. 4.2.1) is the region on which the laser is incident; therefore, $SA = \pi w_0^2$. For conservatism, we will assume $\epsilon = 1$ for all wavelengths.[6] In this

[6]While most materials *are* nearly perfect blackbodies ($\epsilon \approx 1$) at high temperatures, the spectral emissivity of ThO_2 is unusually low in the infrared, while remaining high at visible wavelengths. For example, at 2500 K, the emissivity at 960 nm is 0.6 [165]. For this reason, thorium dioxide is often used to make lantern mantles that cast a bright light when heated by a flame without wasting a lot of energy in the invisible infrared wavelengths.

case, the power emitted by a 100 μm diameter laser spot at 2300 K is $Q_{\text{blackbody}} \approx 10$ mW, much less than the heat loss due to conduction through the target.

In order for a 100 μm diameter target spot to emit 1 W of blackbody radiation, its temperature would have to be 7000 K. At a temperature of $\sim 1.5 \times 10^4$ K, the blackbody emission (22 W) would become comparable to the conduction loss rate through the target. From Eq. (4.12–4.13), we estimate that achieving temperatures this high would require ≈ 60 W of average laser power. Because of the strong expected dependence of signal on temperature (see Section 4.1), it may be advantageous to achieve a higher instantaneous spot temperature at the cost of a lower duty cycle. That is, if the laser spot reaches a steady state on a timescale rapid compared to the thermalization time of the buffer gas cell, we may choose to modulate the laser using an instantaneous laser power of ~ 60 W and a duty cycle of $\sim 25\%$ to heat the target spot intermittently to $\sim 10^4$ K while keeping the time-average laser power below the 15 W pulse tube cooling power threshold. In practice, this is roughly the approach we adopt to optimize the signals in Chapter 5. Achieving temperatures above $\sim 1.5 \times 10^4$ K becomes rapidly more difficult as blackbody radiation starts to dominate the heat loss with its T^4 scaling.

4.2.3 Buffer Gas Conduction

The final cooling mechanism we consider is heat loss to the buffer gas. We estimate this by assuming that the buffer gas atoms incident on the laser spot carry away an average amount of energy consistent with the gas having thermalized with the hot spot. By the equipartition theorem, the average energy of each quadratic degree of freedom is $k_B T/2$ [172]. Since neon is a monatomic gas, and its first electronic excited state (≈ 17 eV) is well above the thermal energy associated with the target temperature ($k_B \times 2300\,\text{K} \approx 0.2$ eV), its internal degrees of freedom are "frozen out." We can therefore treat the buffer gas as a structureless particle with 3 kinetic energy degrees of freedom (in x, y, and z) and obtain that the average energy carried away by each neon particle incident on the target is $3k_B T/2$. The total power

delivered to the buffer gas is therefore:

$$Q_{\text{buffer}} \sim Z_c \pi w_0^2 \frac{3}{2} k_B T_{\text{buffer}} \tag{4.14}$$

$$\sim \frac{3}{2} P_{\text{buffer}} w_0^2 \sqrt{\frac{\pi k_B T_{\text{buffer}}}{2 m_{\text{buffer}}}}, \tag{4.15}$$

where Z_c is the collisional flux (rate of collisions per unit area), P_{buffer} is the pressure of the neon buffer gas, and $m_{\text{buffer}} = 20$ amu is its atomic mass. In Eq. (4.15), we have used the result from kinetic gas theory that the collisional flux is

$$Z_c = \frac{P}{\sqrt{2\pi m k_B T}}. \tag{4.16}$$

We take the temperature of the incident buffer gas to be that of the cold part of the cell, $T_{\text{buffer}} = 16$ K; the beam waist to be $w_0 = 50$ μm; and the pressure in the buffer gas cell to be at most about 1 Torr or $P_{\text{buffer}} \sim 130$ Pa. Plugging these values in gives $Q_{\text{buffer}} \sim 50$ μW. Heat transfer through the buffer gas is therefore clearly negligible.

To summarize, for a 75% ThO_2 + 25% Th target with a 100 μm diameter laser spot heated to 2300 K in a 16 K cell with 1 Torr of neon buffer gas—a scenario in which we expect the ThO production rate to exceed that of the ablation source by a factor of ~ 10 (see Section 4.1)—the heat loss mechanisms are:

$$Q_{\text{conduction}} \sim 10\,\text{W} \gg Q_{\text{blackbody}} \sim 10\,\text{mW} \gg Q_{\text{buffer}} \sim 50\,\mu\text{W}. \tag{4.17}$$

4.3 Feasibility Demonstrations

In order to test the rate of the thermochemical reaction (Eq. (4.1)) for ourselves, we built the apparatus described in Fig. 4.3.1, named the "Diving Bell" for its amusing shape. A stoichiometric mixture of Th and ThO_2 was placed in a resistive tungsten evaporation boat and heated using currents of up to ≈ 125 A at powers of up to ≈ 500 W. The output from

the boat was detected by an RGA with an electron multiplier attachment at a distance of ≈ 25 cm from the boat.

The results from two runs using two different tungsten boats are shown in Fig. 4.3.2. The pressure shown on the plot is the ThO + fractionalized Th pressure reported by the RGA, multiplied by a geometric correction factor (i.e. the ratio of the hemispheric area at the position of the RGA to the exposed surface area of powder in the boat) to give an estimate of the ThO pressure directly above the boat. These pressures and temperatures have a significant uncertainty that is difficult to estimate: The RGA sensitivity to ThO is unknown [181], the exposed powder surface area is difficult to measure accurately, and the measured boat temperature may be higher than the actual temperature of the powder.

While we have no way to account for the uncertainty in the powder area or temperature, we can normalize out the first source of error, the unknown RGA sensitivity, by using the data from the ThO collection plate shown in Fig. 4.3.1. After each run, the collector plate was surveyed using a GM pancake probe (Ludlum Model 44-9) or an alpha-scintillation probe (Ludlum Model 43-1), and the count rates were converted into molecule numbers. During the first run, $\sim$ 1.2 mg/cm^2 were deposited in 1 hour, while during the second run, $\sim$ 2.5 mg/cm^2 were deposited in 2.5 hours. The typical deposition rate is therefore $\sim$ 1 mg/cm^2/h. Given that the collector plates are held at a distance of $\approx$ 5 cm from the boat surfaces, which contain a ~ 0.25 cm^2 area of exposed powder, this implies a total emission rate of 600 mg/cm^2/h $= 4 \times 10^{13}$ mol./s/(100μm)^2 at 2000 K.[7] I estimate that this result is accurate within about an order of magnitude. This normalization factor was used to produce the right-hand y-axis in Fig. 4.3.2.

The Arrhenius Law (see Eq. (4.5)) fits shown on the plot provide a rough measurement of the activation energy of this reaction. Given the slight disagreement between the two

[7] Compare to Eq. (4.7), which gives an order-of-magnitude smaller emission rate for a 10% ThO_2 substrate at 1883 K.

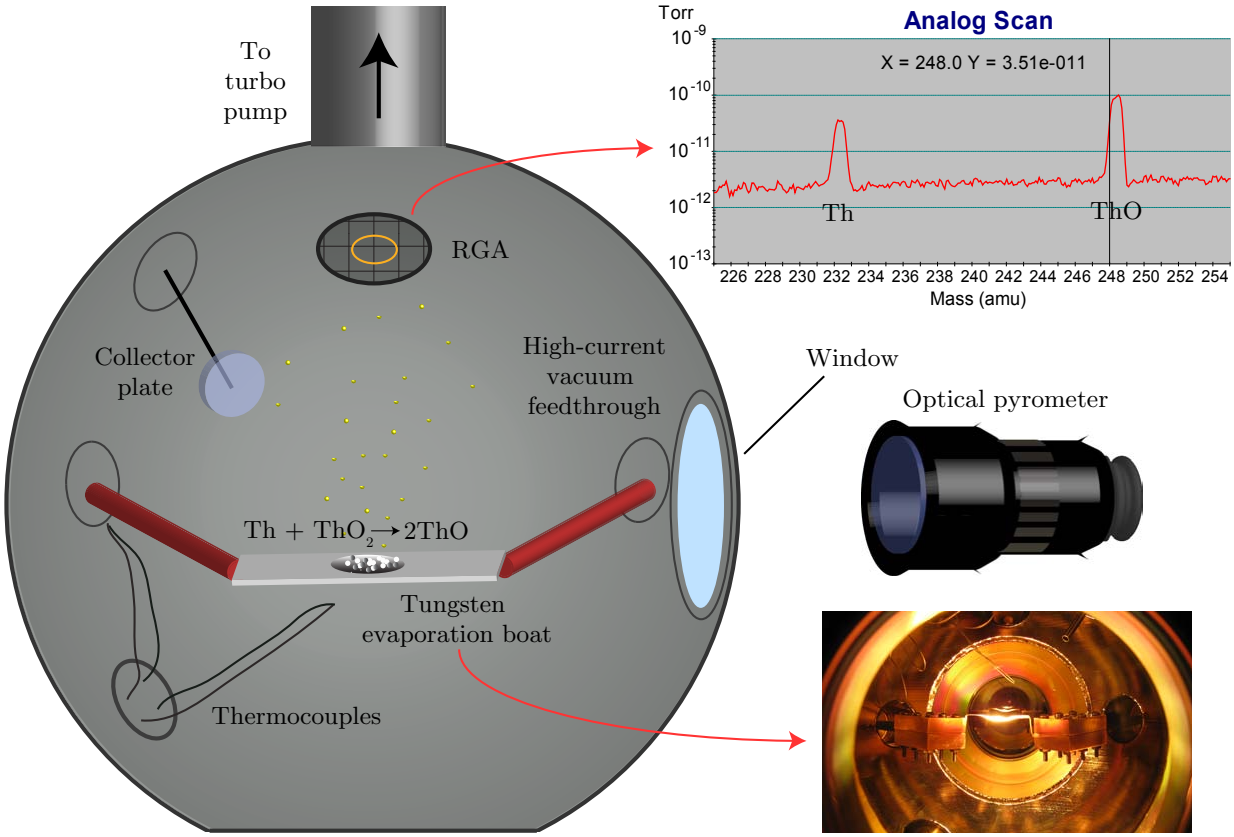

Figure 4.3.1: Measurement of the $\mathrm{Th(s/l) + ThO_2(s) \rightarrow 2ThO(g)}$ reaction rate. A resistive tungsten evaporation boat containing a $\sim$ 1 g mixture composed of equal parts Th and $\mathrm{ThO_2}$ was ohmically heated in vacuum. The temperature of the boat was measured by eye using a disappearing filament optical pyrometer (Pyrometer Instrument Co., Model 87C). The emission from the boat was detected on a Stanford Research Systems Model RGA300 residual gas analyzer. A glass plate was mounted at a distance of 5 cm from the boat to collect the total vaporized thorium-containing material emitted into a well defined solid angle for later measurement using a radiation survey meter. Thermocouples were placed near the boat and the vacuum chamber wall to guard against overheating. An example RGA trace taken at a boat temperature of 2000 K is shown at the top right. The ^{232}Th peak is attributable to fractionalized ThO; the thorium vapor pressure is negligible at these temperatures. A photograph of the boat at 2070 K is shown on the bottom right. The vacuum pressure measured on an ion gauge was $\sim 10^{-9}$ Torr when the boat was off and $\sim 10^{-6}$ Torr when it was on at 2000 K. This apparatus was referred to as the "Diving Bell" because of the appearance of the round, knobbly ConFlat vacuum chamber.

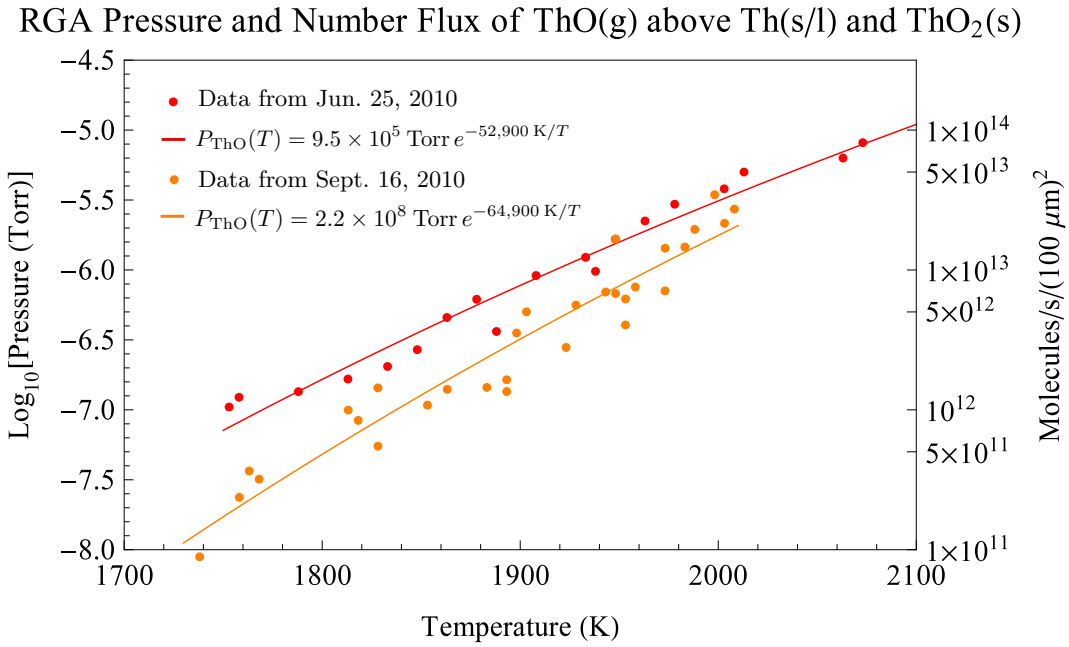

Figure 4.3.2: Data from the RGA measurements of the $\mathrm{Th} + \mathrm{ThO_2} \rightarrow 2\mathrm{ThO}$ reaction in the Diving Bell. The orange and red datapoints were taken using different tungsten boats, and a fit of each result to the Arrhenius Law (Eq. (4.5)) is shown on the plot. The left-hand y-axis shows the estimated ThO pressure above the surface of the boat, calculated as described in the text. These pressures and temperatures have a large uncertainty that is difficult to estimate. Nevertheless, we can obtain the rate of molecule emission per unit area from the surface of the boat within about an order of magnitude by using the thorium accumulated on the collector plate (see Fig. 4.3.1) to normalize the total production from the boat. The radioactivity on the collector plate divided by the total time at which the boat was run at its highest temperature indicates a production rate of $\sim 4 \times 10^{13}$ mol/s/(100 μm)^2 at 2000 K. This normalization factor is used to convert nominal pressures to estimated molecule fluxes on the right-hand y-axis.

measurements, we estimate

$$E_a \sim 6 \pm 1 \times 10^4 \text{ K} \times k_B \approx 5 \pm 1 \text{ eV}. \tag{4.18}$$

This result was used in Section 4.1 to obtain an estimate for the ThO production rate as a function of temperature.

Our second feasibility measurement, illustrated in Fig. 4.3.3, was a prototype version of the thermochemical cryogenic buffer gas source: In order to ensure that good ThO production

rates were achievable in a cryogenic apparatus, we used a liquid helium bath cryostat from IR Labs to test our proposed production method. (By analogy with the Diving Bell, this test apparatus was called the Scuba Tank.) A closed copper cell was fixed to the 4 K cold plate in the cryostat, and four targets of different composition were mounted inside. The cell was filled with a few hundred millitorr of helium buffer gas, and ThO was produced either by laser ablation using a Continuum Minilite II pulsed YAG, or by laser heating using a Millenia Xs continuous-wave (CW) YAG by Spectra-Physics. The in-cell ThO production was measured via absorption on the $X \rightarrow C$ 690 nm transition. We observed per-pulse signals a few times larger with laser heating than with laser ablation, mostly due to the longer duration of the pulses; the peak heights were similar or slightly smaller than those produced by ablation.

Once we had tested the feasibility of thermochemical ThO production in a cryostat, we halted the Scuba Tank tests to move on to a more realistic testbed: a pulse tube cooled neon CBGB, which is the subject of the next chapter.

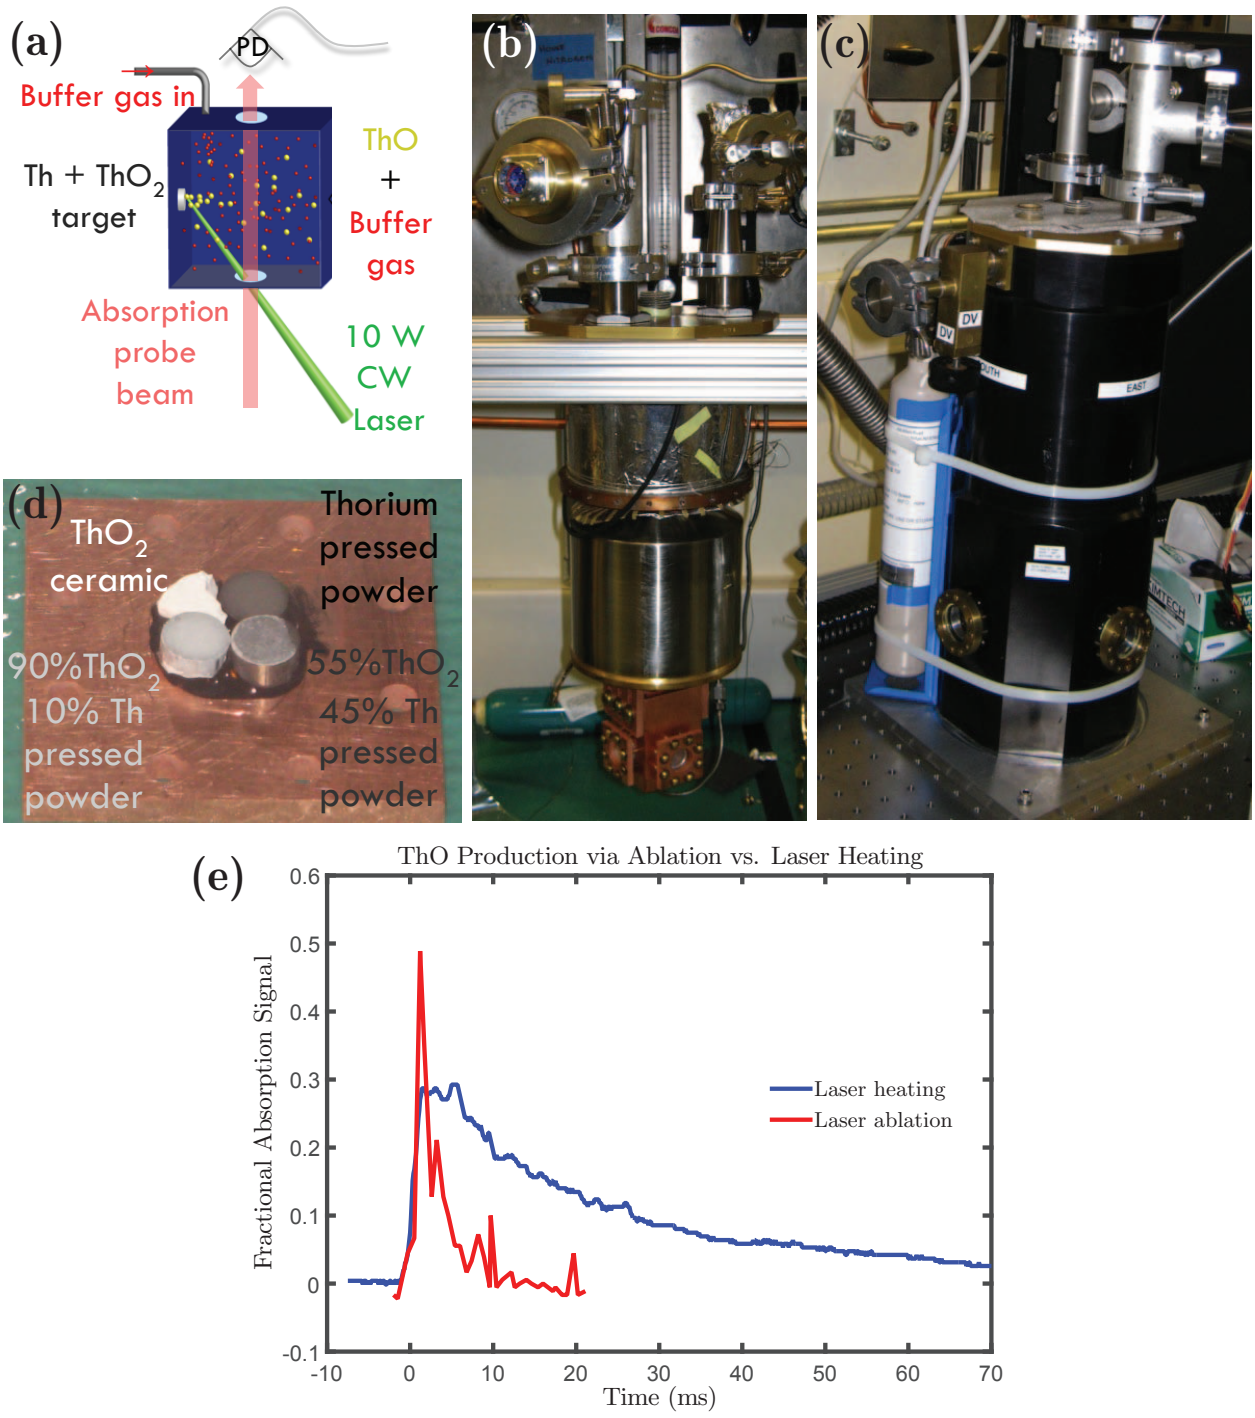

Figure 4.3.3: Test of ThO production in "Scuba Tank" IR Labs cryostat: (a) Schematic of the cell and measurement approach. (b) Photograph of the dewar innards, showing (from top to bottom) thermometer and gas feedthroughs, the LN2 tank, the LHe tank, and the copper cell. (c) Cryostat exterior after assembly shown with a helium buffer gas bottle and a valved vacuum pumping port. (d) Four targets of different composition mounted inside the cell. (e) Data comparing the in-cell ablation and laser heating absorption signals.

It was one of those cases where you approve the broad, general principle of an idea but can't help being in a bit of a twitter at the prospect of putting it into practical effect. I explained this to Jeeves, and he said much the same thing had bothered Hamlet.

—P. G. Wodehouse, *Jeeves in the Morning*

5

Thermochemical Buffer Gas Beam Source

The thermochemical cryogenic buffer gas beam source (CBGB) of ThO that grew out of the ideas in the previous chapter has now been designed, built, tested, and optimized over a wide range of parameters. We describe the results in this chapter.[1]

5.1 Apparatus

The thermochemical source is closely modeled on the ablation-based ThO CBGB apparatus briefly described in Section 2.2 and presented in detail in references [105, 106]. The designs for the old beam source were adapted for the thermochemical source by Jacob Baron. The welded aluminum vacuum chamber is custom-made by Precision Cryo and is approximately

[1] Jacob Baron was a prime mover in this project, designing the thermochemical source, building it, setting up its supporting instrumentation, writing LabView control software and MATLAB analysis software, and performing many of the initial tests. I am deeply grateful to him for making this thesis possible. Much of the work he led is credited explicitly in the text, but his influence pervades this entire chapter. Other collaborators who contributed substantially to this project include Nick Hutzler, Grey Wilburn, and Shuya Yamanaka. I also thank rest of the ACME collaboration for their contributions and technical support.

69 cm long by 64 cm wide by 64 cm tall. Two Cryomech pulse tube coolers are installed in the top plate: A Model PT415 with a cooling power of 1.5 W at 4 K [50] keeps the cryopumping "shields" at 4 K (see Section 2.2), while a Model PT810 with a cooling power of 15 W at 16 K [51] cools the cell. The first (warmer) stage of both pulse tubes is connected to the outer blackbody radiation shield, which achieves a base temperature of $\approx$ 60 K.[2] Flexible ropes of copper 101 welded to brackets of the same alloy serve as thermal links between the pulse tube stages and the apparatus while keeping them mechanically isolated. The total cooldown time is 20 hours.

Thermal tests, the results of which are shown in Fig. 5.1.1, indicate that a maximum heat load of 12 W can be applied to the cell while keeping it below 18 K. These tests also show that the 4 K cryopump is well insulated from the cell so that its temperature does not increase noticeably below a heating power of 14 W. Therefore, the cryopumping speed will not be affected by the laser heating of the targets.

In order to transmit as large a fraction of the beam as possible out of the beam source, we configured the 4 K and 60 K shields to minimize the distance between the cell exit aperture and the exit port of the beam box vacuum chamber. The relevant distances are shown in Table 5.1.1.

The conical collimator in the 4 K shield deflects excess buffer gas emitted from the cell out of the beam so that it can be cryopumped by the 4 K stage. It is kept above 24 K to prevent neon from sticking to it. This collimator must be placed far enough from the source that collisions within the beam are frozen out (or equivalently, far enough that all molecule trajectories are ballistic); otherwise, it may disrupt the hydrodynamic processes near the cell

[2] We initially had only the warm stage of the PT415 coupled to the 60 K radiation shield, while the PT810 warm stage was disconnected and thermally isolated in order to give it a larger cooling power at the cold stage [51]. Unfortunately, the warm stage of the PT415 exhibited the strange misbehavior (not observed in other PT415's in the lab) of having a very low cooling power in the vicinity of 100 K. This made our cooldown times unacceptably long: $\approx$ 40 h to reach steady state. We therefore added heat links between the PT810 warm stage and the 60 K radiation shield, which greatly reduced the cooldown time. The cost was a cold stage cooling power that was diminished by $\approx$ 1 W at a cell temperature of 16 K. (Looking at the data another way, the cell temperature increased by $\approx$ 1 K at a heat load of 12 W.)

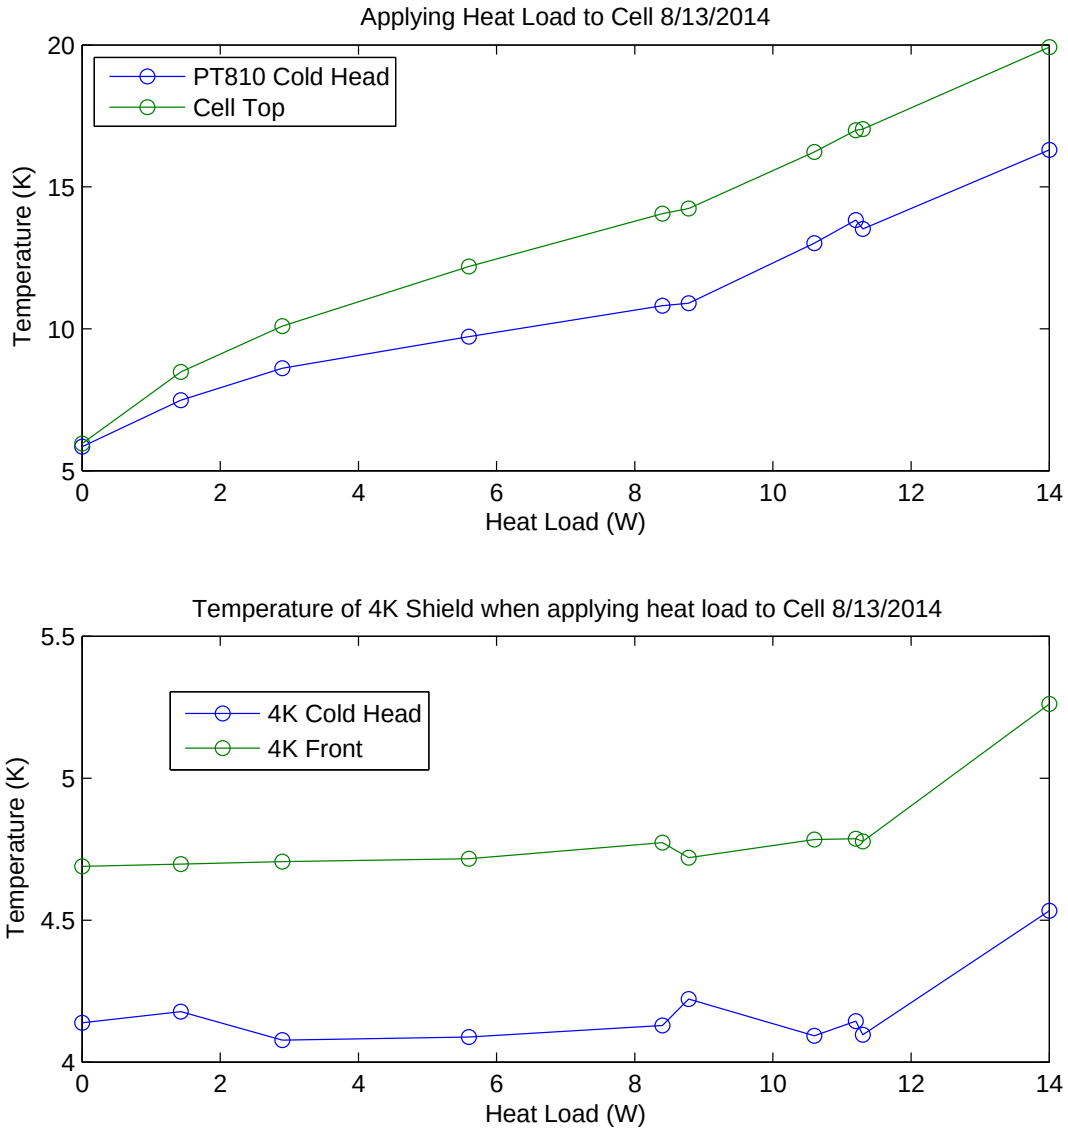

Figure 5.1.1: Cell and 4 K shield temperature as a function of applied heater power. The top plot suggests that the fiber laser heater power must be less than about 12 W in order to keep the cell temperature below 18 K. The bottom plot shows that the temperature of the 4 K shield is fairly independent of the laser power applied to the cell, indicating good thermal isolation between the two. Data and figure from Jacob Baron.

exit and change the final beam properties [106, 107]. The minimum distance between the cell and the tip of the 4 K collimator is therefore determined by the gas dynamics of the source, so it is kept the same as in the ablation source. The other distances are determined by cryogenic mechanical engineering constraints: the minimum distance between surfaces at different temperatures is kept above ≈ 1 cm to avoid accidental thermal touches. Altogether, the distance from the cell exit aperture to the source vacuum chamber beam exit port is shorter by ≈ 5.5 cm than in the ablation source.

Table 5.1.1: List of apertures along the beamline that define the molecule beam cross-section in the thermochemical source. In this apparatus, we prioritized minimizing the distance from the cell to the the beam box exit port. This distance is 12 cm in the thermochemical source, compared to 17.5 cm in the ablation source.

Distance (cm)	Width (mm)	Description
0	5	Circular cell exit aperture, 0.5 mm thick
2.5	6	Circular knife-edge conical collimator in 4 K shield
6.7	10	Circular aperture in 60 K shield
12	—	Vacuum port on beam exit plate (green-painted exterior)

In addition, in light of the beam attenuation measurements discussed in Section 2.3.3, we have added three removable "chevron baffle vents" (designed and made by Jacob Baron) to the front of the 60 K shield [71]. These blackened copper pieces are shaped like overlapping V's to form bent channels that transmit gas while absorbing blackbody photons. The three chevron vents provide an estimated total conductance of 120 L/s through the 60 K shield. Since any transmitted gas particles will, with very high probability, stick to the 4 K shield, this conductance determines the pumping speed of the 4 K shield on the room temperature part of the apparatus. At present, we have only one of the three vents installed, providing a pumping speed of 45 L/s.

A photo of the new beam source during its final assembly is shown in Fig. 5.1.2. To summarize, the main differences between the old Gen. I/II cryogenic chamber and the new thermochemical chamber are:

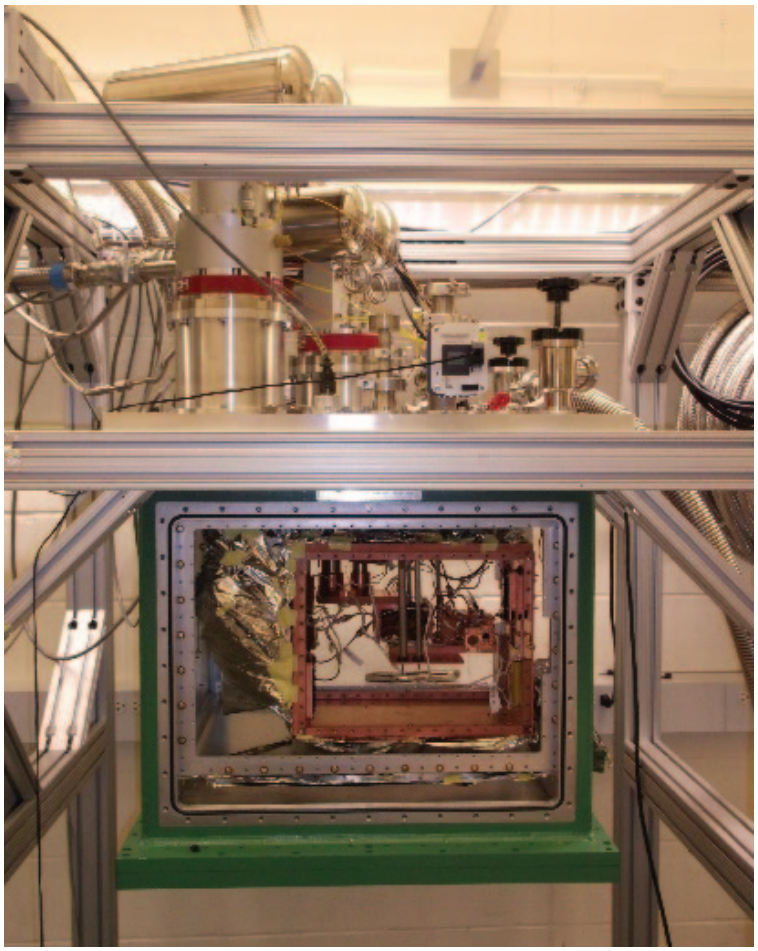

Figure 5.1.2: Photograph of thermochemical source.

- *Two pulse tube coolers*: A PT415 cools the 4 K shield, and a PT810 cools the cell, which is thermally insulated from 4 K stage. Both pulse tubes cool the 60 K shield.

- *Reduced beam length*: The distance from the cell exit to the vacuum exit is 5 cm shorter in the thermochemical than in the ablation source.

- *Chevron baffle vents*: The front of the 60 K shield has spots for three chevron vents with a total pumping speed of $\approx$ 120 L/s. Only one of the three is currently installed.

- *Optical access after 4 K collimator*: A few millimeters of optical access have been provided just downstream of the 4 K shield exit ($\approx$ 4.7 cm from the cell) in order to permit absorption spectroscopy on the collimated beam.

- *Extra superinsulation and thermally matched screws*: During the period when we were investigating and trying to improve the slow cooldown time, we added multiple layers of aluminized Mylar superinsulation to the shields, covered all light leaks in the 60 K shield with aluminum tape, and replaced a number of brass screws in critical locations (including on the pulse tube cold head and the flexible heat link brackets) with stainless steel screws accompanied by titanium washer stacks to match the differential thermal contraction of the screws and the copper.

- *4 K shield heaters*: Two of the time-limiting operations in the ablation source were de-icing the beam box and warming it up to room temperature in preparation for changing targets. To help speed this process up, Jacob added five 15 Ω, 100 W power resistors to the 4 K shield. Unfortunately, when we tried to use them to perform a rapid de-ice, the heaters seemed to outgas a lot of material (possibly water) that condensed onto the cell windows, forcing us to warm up nearly to room temperature in order to remove it. Any future attempts to use these heaters must be very cautious to avoid contaminating the surrounding cold surfaces.

- *Paint*: The old beam source chamber is painted a dull and staid "Nitro Blue," while the new chamber is a vibrant and youthful "Vista Green."

In addition to the cryostat changes described above, we have also made several modifications to the cell and its contents. A comparison between the two cell designs is shown in Fig. 5.1.3. The main changes are:

- *Mixed Th + ThO_2 pressed powder targets*: In contrast to the ceramic ThO_2 ablation targets, the thermochemical targets are a 75% ThO_2 + 25% Th pressed powder mixture, fabricated as described in Appendix B. A photograph of a used thermochemical target is shown in Fig. 5.1.4.

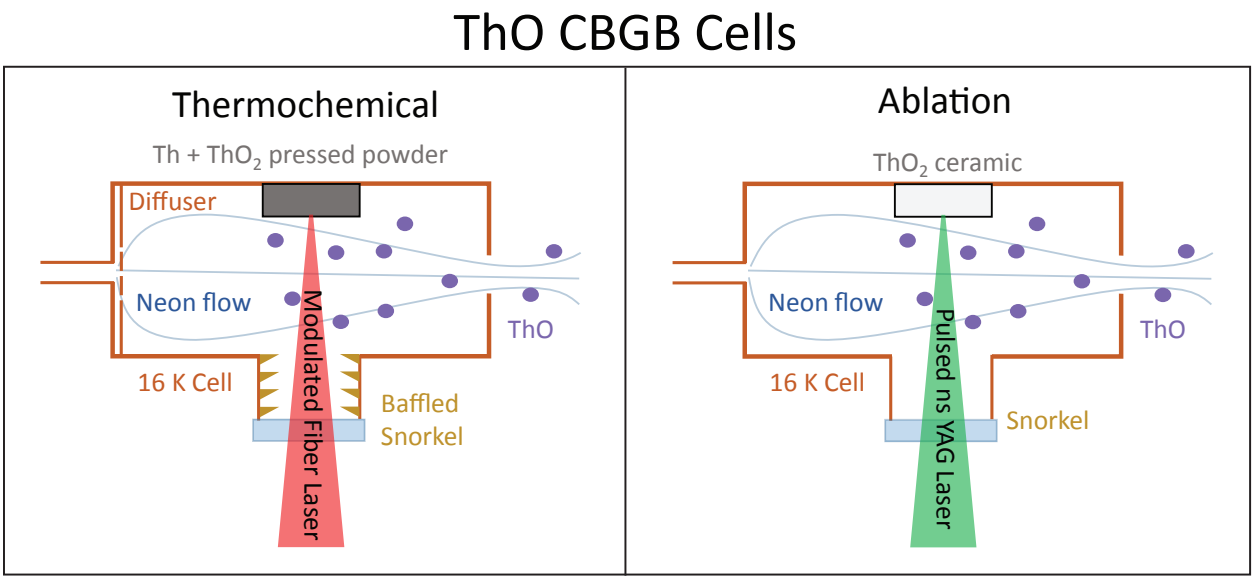

Figure 5.1.3: Comparison between ablation and thermochemical buffer gas cells.

- *Modified snorkels*: The windows on the thermochemical cell become opaque faster than the windows on the ablation cell due to buildup of material ejected from the target. Surveys of the snorkels indicate a fairly uniform dust distribution along the length of the snorkel, up to a distance of 7 in. from the cell body. This suggests that the dust may not simply be diffusing or propagating ballistically to the cell windows, but rather ricocheting off the snorkel walls until it is stopped by the window. We therefore recently added knife-edge circular baffles and 1 mm thick copper wool to the snorkel walls in an effort to trap the dust. In addition, we use windows made of sapphire (an extremely inert material), left uncoated on the inside face, to avoid having chemicals from the beam source react with the window and degrade the optical quality. We have no direct evidence yet that either of these changes has helped with the window clouding problem, but we plan to survey the used baffled snorkels to find out.

- *Multi-hole diffuser*: Just inside the fill line opening into the cell, we have now added a thin plate with an array of small, $\sim$ 0.5 mm holes on a 10 mm diameter. This is intended to interrupt any strong flow lines that might give rise to eddies in the cell or prevent efficient entrainment of ThO in the beam, especially from upstream targets

near the fill line. We have no evidence that the diffuser has had any effect on the signals.

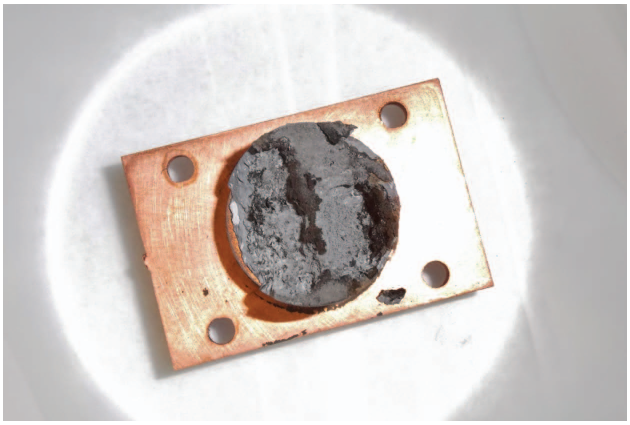

Figure 5.1.4: Mostly depleted 0.75 in. diameter thermochemical source target composed of 75% ThO_2 + 25% Th. The target surface appears to undergo complex physical and chemical changes as it becomes depleted. The snorkel for this target had a horizontal copper divider down the center in front of where the dark ridge appears on this target.

The heating laser hardware is also a crucial component of the thermochemical source. The key items are a 50 W maximum power diode-pumped 1070 nm ytterbium fiber laser (Model YLR-50-AC from IPG Photonics, purchased and set up by Nick Hutzler) and a two-axis scanning galvo with a 0.2 radian/ms scanning speed (Model XLR8-10 from Nutfield Technology, purchased and set up by Jacob Baron). The fiber laser, which can be modulated with a rise time of ~ 40 μs via its emission gate, is aligned through the galvo and focused onto the target. The two mirrors of the galvo allow us to rapidly move the laser spot around the target to find regions of good yield in a programmatic way. Jacob has written galvo control LabView software that enables us to run in any of the following modes:

1. Rastering the laser around a defined target area with a pre-programmed or custom waveform
2. Hopping the laser randomly around a defined target area and dwelling on each target spot for a predetermined amount of time
3. Hopping the laser randomly around a defined target area as in 2 and rastering at each point as in 1, or

4. Running in mode 2 or 3, but hopping to a new spot only when the molecule signal falls below a set threshold.

The following sections describe our measurements of the beam properties and our exploration of the parameter space of running conditions with the goal of optimizing the ThO beam flux.

5.2 Buffer Gas Flow Rate

A hydrodynamic CBGB like that used in the ACME experiment relies on the species of interest experiencing enough collisions with the cold buffer gas to thermalize (see Appendix C), as well as on the buffer gas flow rate being high enough to entrain the molecules in the flow out of the cell before they can diffuse to the walls. The relevant parameter for hydrodynamic entrainment is [106, 107]:

$$\gamma_0 = \frac{t_{\text{diff}}}{t_{\text{pump}}} \sim \frac{f_{0,\text{b}}\sigma}{\bar{v}_{0,\text{b}} L_{\text{cell}}}, \tag{5.1}$$

where t_{diff} is the diffusion time to the cell walls, t_{pump} is the characteristic emptying or "pump-out" time of the cell, $f_{0,\text{b}}$ is the neon flow rate in SCCM, σ is the Ne–ThO collision cross section, $\bar{v}_0$ is the mean thermal velocity of the in-cell neon, and L_{cell} is the distance from the target to the cell exit. As in Section 2.2, the subscript "0" indicates the steady-state value of an in-cell quantity, and the subscript "b" indicates a buffer gas quantity. In order to achieve significant extraction from the cell into the beam, we require the pump-out time to be faster than the diffusion time or $\gamma_0 \gtrsim 1$. For a cell temperature of 16 K, a cell length scale of $L_{\text{cell}} \sim 1$ cm, and a collision cross section of $\sigma \approx 3 \times 10^{-15}$ cm^3 [106], we find that the neon flow rate must be $f_{0,\text{b}} \gtrsim 10$ SCCM.

In order to determine the optimum neon flow rate in the thermochemical source, we performed an absorption measurement using the setup illustrated in Fig. 5.2.1. A test beamline constructed of standard ISO-100 and ISO-80 vacuum parts was connected to the beam source, and square apertures for collimation and differential pumping were installed at the positions indicated in the figure. The molecule beam signal was simultaneously measured

at a distance of 0 and 77.5 cm from the cell exit via absorption spectroscopy on the $X \rightarrow C$ $Q(1)$ line (i.e., $|X, J = 1\rangle \rightarrow |C, J = 1\rangle$). The integrated absorption optical depth (OD) was converted into a molecule flux using the beam geometry and an estimate of the absorption cross section [34, 105]. The results from this measurement are shown in Fig. 5.2.2.

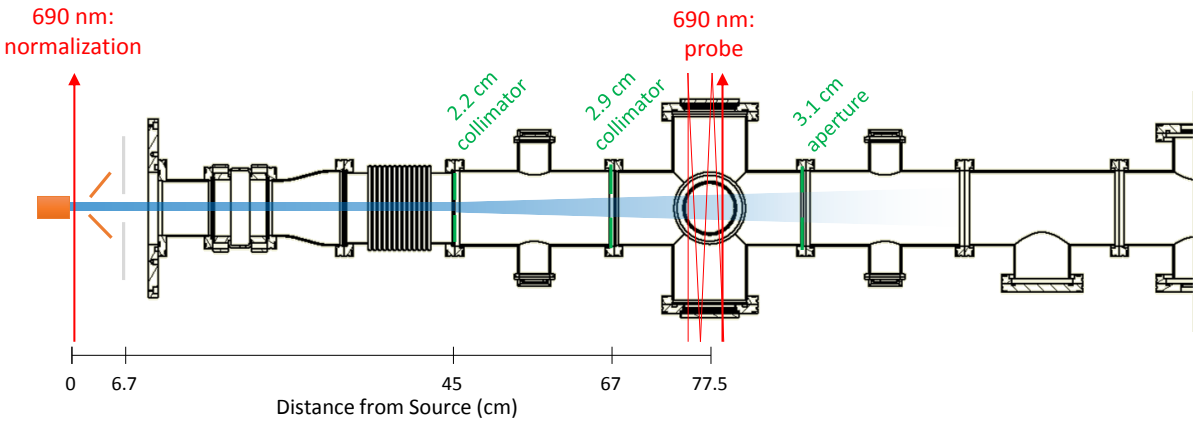

Figure 5.2.1: Schematic of the test beamline multi-pass absorption setup used for the data in Fig. 5.2.2. The beamline consists of standard ISO-80 and ISO-100 parts with optical access for spectroscopy. From left to right, the relevant components are: the buffer gas cell (orange), the conical 4 K collimator (orange), the 60 K shield aperture (gray), the beam box exit plate, a nipple and gate valve, a conical ISO-80 to ISO-100 adapter, a bellows for mechanical isolation of the beamline, three consecutive ISO crosses with optical access for absorption and fluorescence spectroscopy, and additional downstream components. The large ISO cross also holds a 550 l/s turbo pump to maintain the beamline vacuum. The square collimator and differential pumping aperture positions and sizes are indicated in green, and the collimated ThO beam is shown in blue. A single-pass absorption laser measures the beam signal outside the cell, and a multi-pass absorption laser measures the beam signal at a distance of 77.5 cm downstream.

At a later time, we performed another measurement of the beam signal and transmission efficiency v. flow rate under different running conditions. The main difference in this second experiment was the high cell temperature, which was around 29 K instead of our usual range of 15–17 K. Other differences include a longer fiber laser pulse length, a higher fiber laser current, and a different downstream absorption transition: $X \rightarrow C$ $Q(2)$ (i.e., $|X, J = 2\rangle \rightarrow |C, J = 2\rangle$) instead of $Q(1)$. For this data, the absorption lasers were at 0 and 6 cm from the cell exit, with the downstream laser positioned just past the exit of the 4 K conical

collimator. The results of this measurement are shown in Fig. 5.2.3.

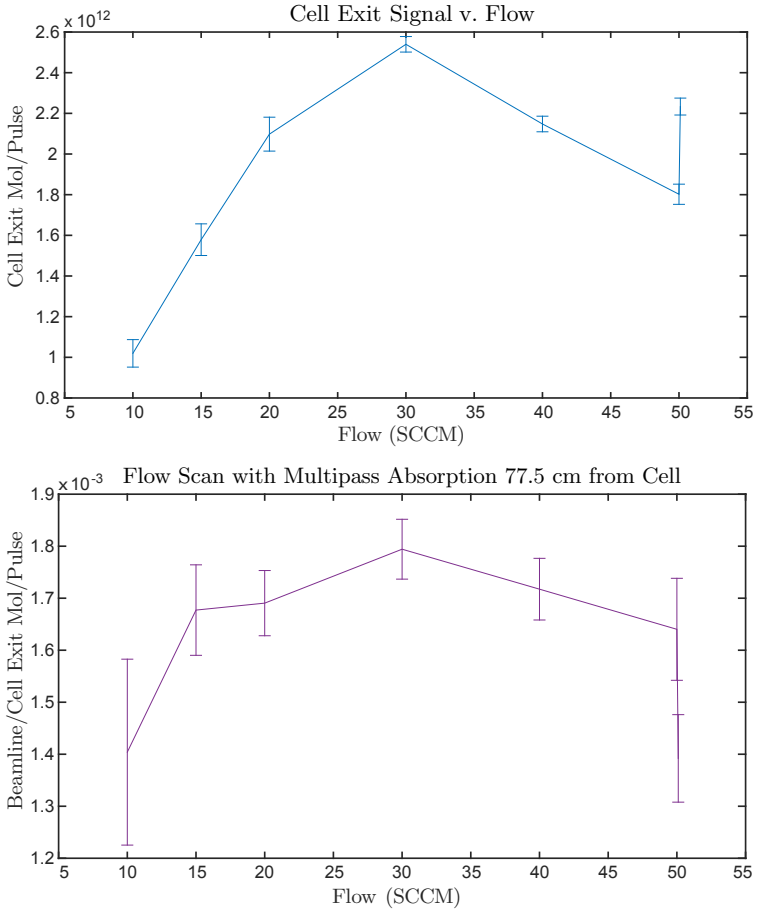

Figure 5.2.2: Beam flux and transmission v. neon flow rate. Fiber laser firing at upstream target with 20 ms pulses at a repetition rate of 10.7 Hz. Fiber laser current at 70%. Cell top temperature typically 15.5 K. **Top:** Total beam flux in $|X, J = 1\rangle$ determined by measuring the absorption signal on the $X \rightarrow C$ $Q(1)$ line just outside the cell. **Bottom:** Transmission to the beamline at a distance of 77.5 cm from the cell.

Results: During the initial tests shown in Fig. 5.2.2, a broad peak in extraction and transmission efficiency appeared around 30 SCCM. The cell exit beam signal was fairly independent of flow between 20 SCCM and 40 SCCM, and it became distinctly worse below 15 SCCM. At 50 SCCM, Fig. 5.2.2 shows some indication that both the extraction from the cell and transmission to the beamline may begin to decrease. These results are fairly consistent with the results and beam dynamics model in reference [106]. In this work the cell extraction increased up to neon flow rates of $\approx$ 30 SCCM and then began to diminish for flow

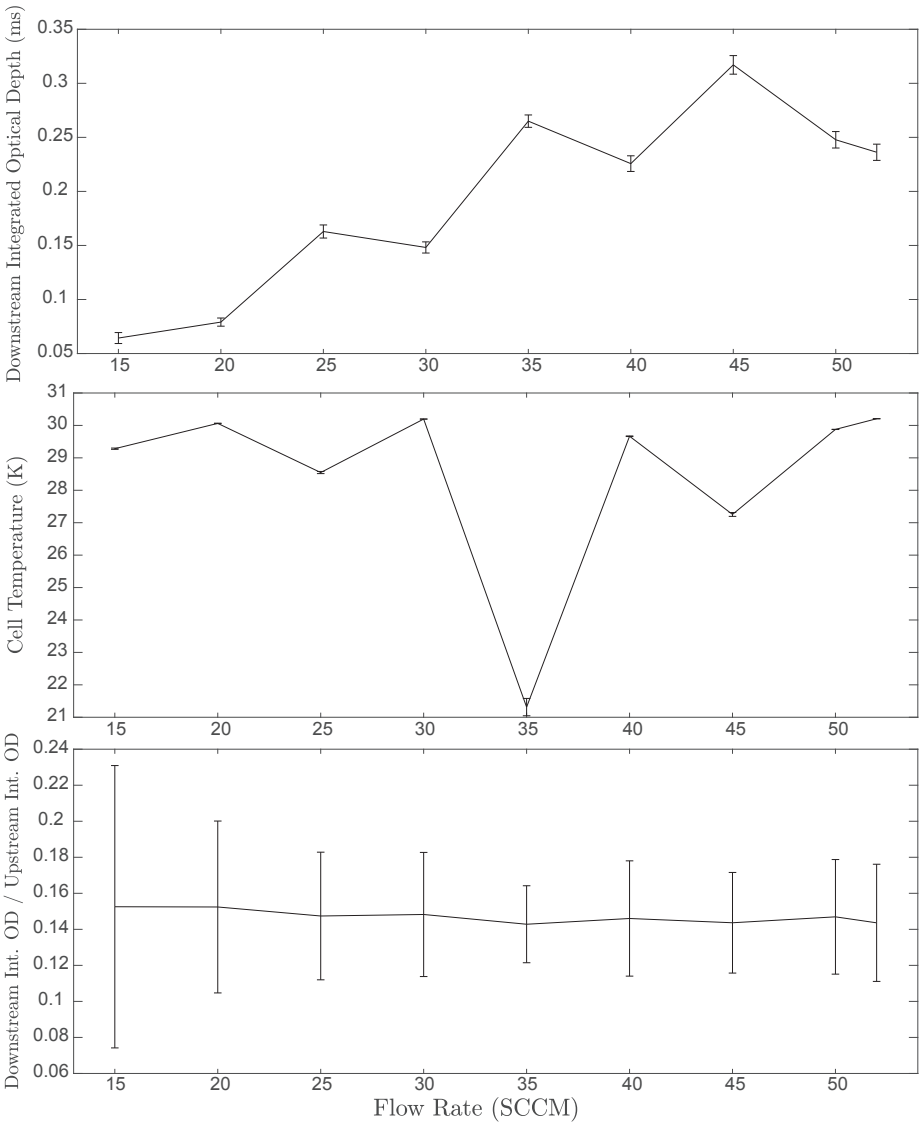

Figure 5.2.3: Beam flux and transmission v. neon flow rate at high cell temperatures. Firing at downstream target (2.5 cm from cell exit) with 75 ms pulses modulated at 800 Hz with a duty cycle of 80% (1.0 ms on, 0.25 ms off) at a rep rate of 9.3 Hz with the fiber laser current at 100%. Using an f = 400 mm lens positioned at about 420 mm from the target and spot hopping randomly around most of the target with a frequency of 5 Hz. **Top:** Downstream integrated OD on the $X \rightarrow C$ $Q(2)$ line just after the 4 K shield exit aperture. **Middle:** Average cell temperature for each flow point in the top plot. This data suggests part of the cause of the signal fluctuations in the top plot: The signal is generally anti-correlated with the cell temperature, and the high signal spikes are seen for data points where the cell temperature fluctuates downward. **Bottom:** Transmission signal from just outside the cell exit to just outside the 4 K shield as a function of flow rate. The cell exit signal is measured on the $X \rightarrow C$ $Q(1)$ absorption line. This transmission signal inside the beam box is quite flat within uncertainty, in contrast to the beamline transmission data in Fig. 5.2.2, where the correlation with neon flow is on the borderline of statistical significance. Data analysis by Adam West.

rates above ≈ 50 SCCM. The fall-off at low flow rates is attributable to the cell pump-out time exceeding the diffusion time to the walls at low densities [106]. The fall-off at high flow rates may be attributable to the high buffer gas density creating high-order diffusion modes (i.e. ThO concentrations localized near the target spot) that decay away rapidly, causing most of the molecules to return to the target surface and nearby walls rather than entering the flow [173].

Conversely, the more recent test with a high-temperature cell shown in Fig. 5.2.3 did not show a clear signal peak as a function of buffer gas flow. Instead, there was a generally increasing trend in signal v. flow with significant fluctuations, where the fluctuations seemed to be correlated with the fluctuations in cell temperature. As discussed in Section 5.3, however, running at temperatures above 20 K is generally undesirable, so we ignore the latter data set and consider only the former for the purpose of selecting a flow rate.

Conclusion: Run at 30 SCCM.

5.3 Cell Temperature

The temperature of the cell affects the final rotational and translational temperatures of the beam, its divergence, and its forward velocity. Lower cell temperatures produce lower (i.e., more favorable) values for all three of these quantities [106, 107]. We therefore typically run the cell at the lowest possible temperature at which the buffer gas has significant vapor pressure. For neon, this is approximately 16 K.

In the thermochemical source, we expect the ThO production to scale rapidly with the laser power applied to the target, as discussed in Chapter 4. Since the cell temperature also increases with this applied heat load, we wished to investigate the dependence of the single-quantum-state beam signal on the cell temperature in order to optimize the trade-off between keeping the cell temperature low and maximizing the beam yield.

Figure 5.3.1 shows a measurement of the per-pulse $|X, J = 1\rangle$ beam signal (in units of time-integrated optical depth on the $X \rightarrow C$ $Q(1)$ transition) as a function of cell tem-

perature. All other settings are held constant. The cell temperature was controlled using a resistive heater bolted to the cell exterior.[3] The $|X, J = 1\rangle$ beam signal at the cell exit shows an approximately linear dependence on temperature between 16 and 23 K, falling by a factor of 2 between 16 and 21 K. This decrease is likely caused by a combination of effects, including an increased rotational temperature spreading out the ThO population among more J levels, an increased translational temperature broadening the Doppler profile (typically $\approx$ 140 MHz FWHM at a cell temperature of 16 K, which is much larger than the absorption laser linewidth of $\approx$ 1 MHz) and thereby decreasing the optical depth on resonance, and a higher forward velocity decreasing the fly-through time of the molecules. Since all of these effects are detrimental to the eEDM sensitivity (see Eq. (6.2)), we take the integrated OD as a rough proxy for the beam figure of merit. We can therefore infer from Fig. 5.3.1 that lower cell temperatures are strictly preferable, down to 16 K.

Below a cell temperature of about 15.5 K, however, the beam signals are no longer stable: At these temperatures, neon gradually freezes inside the fill line, forming a plug of ice. When this happens, we have to halt the run in order to thaw it out. This effect sets a lower limit on the ideal cell temperature.

The data in Fig. 5.3.1 indicate that the signal degrades with increasing cell temperature when all other parameters are held constant. However, as shown in Chapter 4, increasing the heat load from the fiber laser is expected to give higher production rates of ThO. In order to get a better sense for the trade-off between these effects, we monitored the cell temperature during a run in which many beam parameters were varied in an effort to optimize the signal. This data is shown in Fig. 5.3.2. For these running conditions, the molecule flux is roughly independent of temperature up to $\approx$ 19.5 K, at which point the signals start to drop. Although this measurement plainly has a number of confounding variables, it does suggest that at sufficiently high temperatures, the signal decrease with temperature observed

[3]The temperatures of the cell and other cryogenic components in the source are measured using DT-670-CU (uncalibrated) Lake Shore Cryotronics silicon diode thermometers.

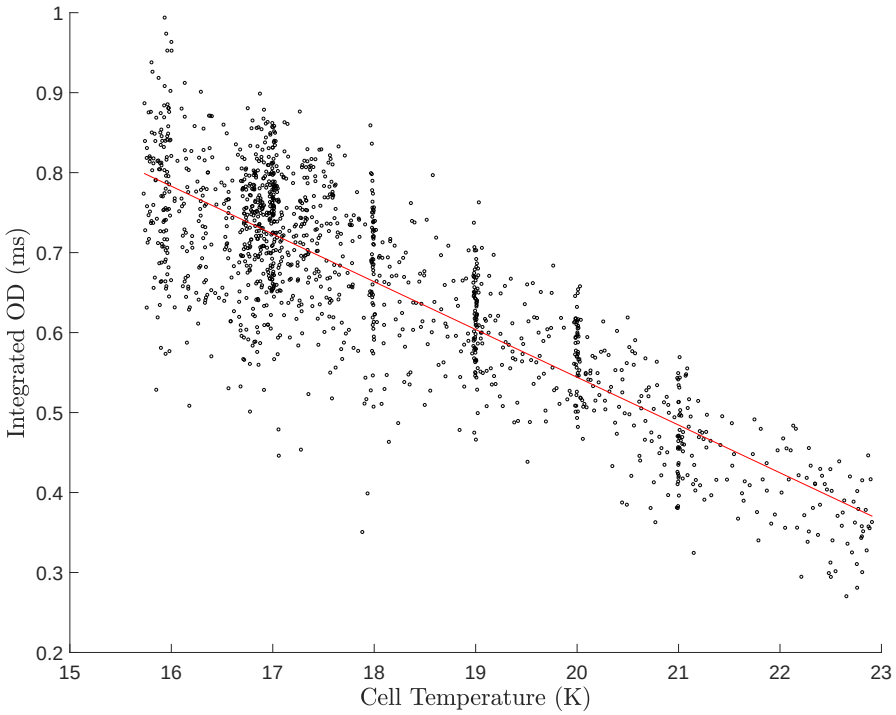

Figure 5.3.1: Beam signal v. cell temperature. Fiber laser firing at the 3$^{\text{rd}}$ upstream target (7 cm from the cell exit). Galvo programmed to spot hop randomly around most of the target surface with a frequency of 5 Hz. The lens has a 500 mm focal length and is positioned at $\approx$ 460 mm from the target. The fiber laser has its current set at 100% and is pulsing at a rep rate of 13.7 Hz with a 25 ms pulse width modulated at 400 Hz with a duty cycle of 70% (1.75 ms on, 0.75 ms off). The absorption signal is measured on the $X \rightarrow C$ $Q(1)$ line just after the cell exit. Each point represents an average of 16 pulses. Over this range, the data shows a clear linear decrease in signal v. temperature with a rate of about 8% per K. The red line is a linear fit which serves as a guide to the eye. Data analysis by Adam West.

in Fig. 5.3.1 is no longer compensated by the larger ThO yields effected by a higher time-averaged fiber laser power.

Results: All else being equal, the single-quantum-state signal falls off monotonically with cell temperature above 16 K. Above 19.5 K, this fall-off is no longer overcome by most other advantages of increasing the heat load from the fiber laser.

Conclusion: Choose a time-averaged fiber laser power less than 13 W, such that the average cell temperature is at most 18 K when the entire heat load comes from the fiber laser (see Fig. 5.1.1). Fix the cell temperature setpoint at 16 K, and servo using a resistive heater. This will keep the cell warm enough to prevent neon ice from clogging the fill line during

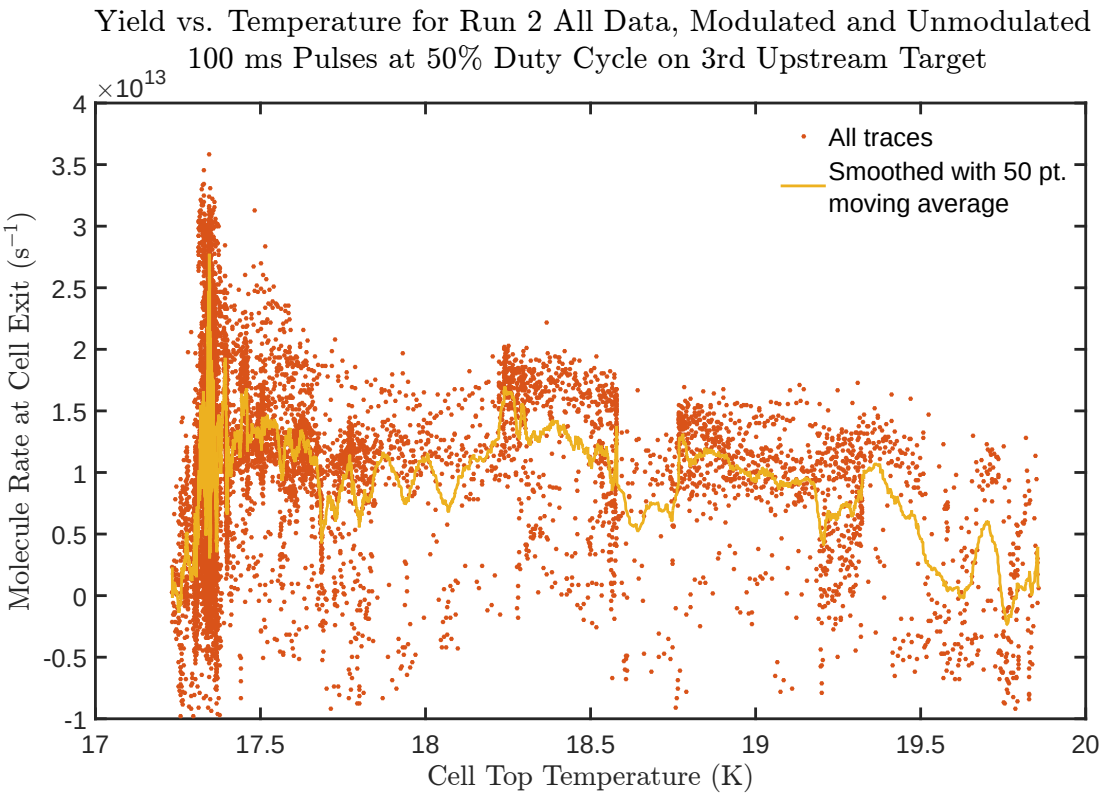

Figure 5.3.2: This data shows the correlation between the beam signal (population in $|X, J = 1\rangle$ at the cell exit) and cell temperature during a run in which the cell temperature setpoint was kept at 17 K while the fiber laser power, modulation status, and modulation frequency were varied. This data and the conditions under which it was taken are presented more fully in Fig. 5.4.11. This plot illustrates that temperature correlations like the one shown in Fig. 5.3.1 can be compensated or more than compensated by other advantages of increasing the heat load on the cell. Nevertheless, even under these rather catholic running conditions, a universal fall-off in the signal appears around 19.5 K, suggesting a limit to these advantages.

times when the fiber laser deposits too little power.

5.4 Laser Pulses

The timing and intensity of the fiber laser present a large parameter space for exploration. A critical part of the ACME experiment routine is the periodic measurement of the background (i.e. the signal with no molecules present) that is made in order to obtain good statistics on our signals. In order to leave time for such background measurements, we rule out running the thermochemical source in true CW (continuous-wave) mode, and instead create our beam with a series of fiber laser pulses of a given pulse width and repetition rate. Within a pulse, we can vary the instantaneous fiber laser power to produce "pulse modulation." These various timescales are illustrated in Fig. 5.4.1.

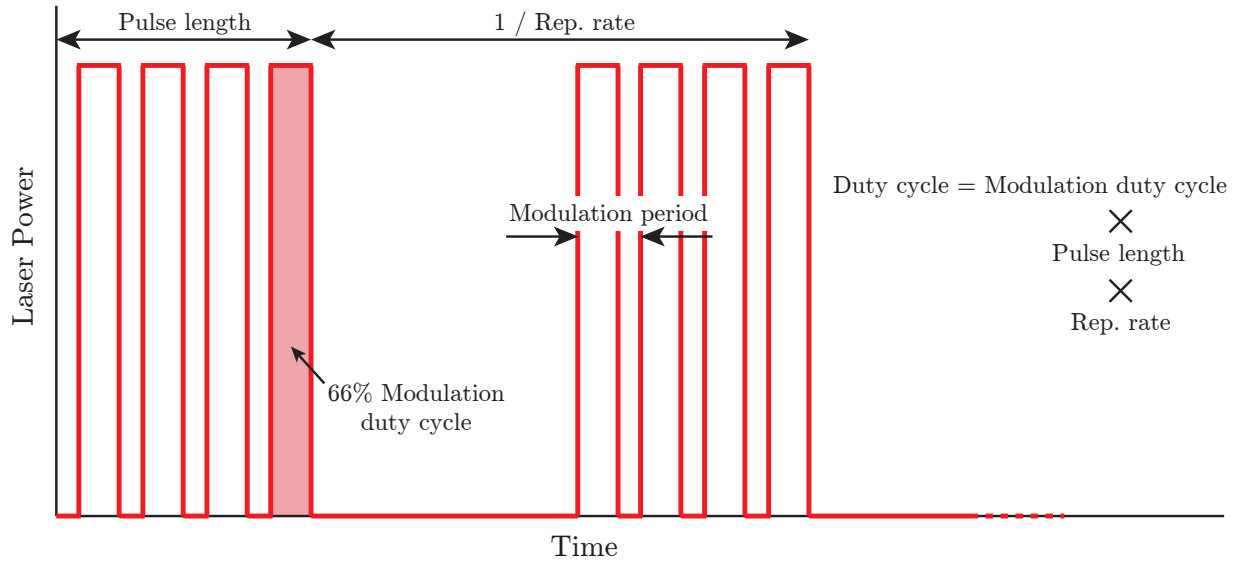

Figure 5.4.1: Fiber laser pulse timing diagram showing the pulse length, pulse repetition rate, modulation period, and modulation duty cycle. Figure by Adam West.

In this section, we investigate the ThO beam flux dependence on the instantaneous fiber laser power, pulse width and repetition rate, and modulation of the laser within a pulse. Many of these effects are interdependent, and determining the overall optimum in this multi-dimensional parameter space presents a challenge. By studying the effect of each parameter independently, we can make some deductions about the most likely ideal running conditions. We first summarize the conclusions and the final selected running conditions for the fiber

laser and then present the results of our studies in detail in the subsections below.

Results: The dependence of the signal on the instantaneous fiber laser power is presented in Section 5.4.1. As shown in Fig. 5.4.8, the peak molecule flux is only weakly dependent on the laser power between 10 and 50 W. This result is somewhat surprising, given the expected strong dependence of ThO production rate on spot temperature derived in Section 4.1. This suggests either that the spot temperature is reaching a threshold where the heat loss is dominated by blackbody radiation (which is unlikely, given the analysis in Section 4.2) or that other effects cause the yield rate to saturate at sufficiently high temperatures. Such possible effects include depletion of the reactants at the laser spot, thermal ejection of target material, collisions near the target surface, or detrimental effects in the buffer gas dynamics caused by the high target surface temperatures.

Although the dependence of the peak ThO flux on the heating laser power is unexpectedly weak, the flux rise time is typically (though not always) shorter at higher laser power, as shown in Fig. 5.4.8. This is likely due to the heating timescale of the target, discussed in Section 4.2.1: A higher laser power heats the target surface more rapidly to a temperature at which ThO production rates are favorable. At high laser power, therefore, the time-averaged signal is comparatively independent of pulse length when the beam duty cycle (i.e., pulse length times repetition rate) is held constant. Conversely, at lower power, longer pulses are favored because of the additional time required to reach peak signal output. Furthermore, at low instantaneous laser power, we find that more target spots are dead on arrival, especially when we have previously been running at higher power. This effect is illustrated in Fig. 5.4.3 and 5.4.4. There is also a target aging effect wherein as the target surface gets used, higher laser powers are needed to achieve the same signal maxima previously achieved with lower powers: see Fig. 5.4.6. All of these observations argue for running with the maximum available (≈ 50 W) instantaneous laser power.

This picture is somewhat complicated by the data shown in Fig. 5.4.7, which summarizes the data on molecule flux rate v. instantaneous laser power at constant duty cycle. The

ultimate constraint on production is the time-averaged heat load on the cell, which is limited to $\approx$ 13 W, as discussed in Section 5.3. Under some running conditions, Fig. 5.4.7(b) suggests that slightly ($\sim$ 30%) higher time-averaged beam fluxes may be achievable by running a lower-power fiber laser beam with a higher duty cycle, whereas under other running conditions, Fig. 5.4.7(a) shows that higher instantaneous laser powers are strictly better. Since these results are contradictory, we judge that they do not overrule the aforementioned considerations in favor of higher instantaneous laser power.

Another variable over which we sought to optimize is the fiber laser pulse width, discussed in Section 5.4.2. When running with a near-maximum laser power, Fig. 5.4.9 shows a broad peak in time-averaged single-state molecule flux at a pulse width of 80 ms. For pulse widths below about 40 ms, the flux has not reliably reached its peak steady-state value, while for pulse widths above about 120 ms the instantaneous flux starts to decrease, either because of depletion of the target region, overheating of the cell or buffer gas, or other unknown effects.

Finally, we investigated modulating the fiber laser with a rectangle wave within each pulse. The data on these studies is presented in Section 5.4.4. At a fixed heat load (i.e. time-averaged fiber laser power), we observe some evidence that a modulation duty cycle of $\approx$ 60% or lower provides better beam signals than continuous pulses. For modulation duty cycles below $\approx$ 40%, the pulses start to overlap when the fiber laser heat load is held constant, which runs counter to our background acquisition requirement discussed above. The highest modulation frequency we tried (and the one giving the smoothest beam signals) was 800 Hz. Modulation frequencies much higher than a few kHz would start to encroach on the $\approx$ 40 μs rise time of the fiber laser.

Conclusion: Run with 100% fiber laser current and an overall pulse width of $\approx$ 80 ms. Within each pulse, modulate the laser with an 800 Hz rectangular wave with a duty cycle of 50% (i.e. a square wave with 0.625 ms on time, 0.625 ms off time per cycle). In order to keep the total heat load below 13 W, the pulse repetition rate under these running conditions must be $\approx$ 7 Hz.

5.4.1 Fiber Laser Current

Figure 5.4.2 shows the dependence of the ThO beam signal on the fiber laser power under the particular set of running conditions stated in the caption. This data was taken using absorption on the $X \rightarrow C$ $Q(1)$ line using 20 ms laser pulses. Both the 1σ spread in the data and the standard error are shown to give a sense for the large shot-to-shot fluctuations in the signal. The positive x-intercept in this plot suggests that the beam flux exhibits "threshold" behavior as a function of laser power, i.e., the molecule beam only turns on when the laser power is sufficiently high. This is in contrast to the result in Fig. 5.4.8, discussed below. Unless indicated, the laser spot size on the target for the data shown in this section is ≈ 100 μm in diameter.

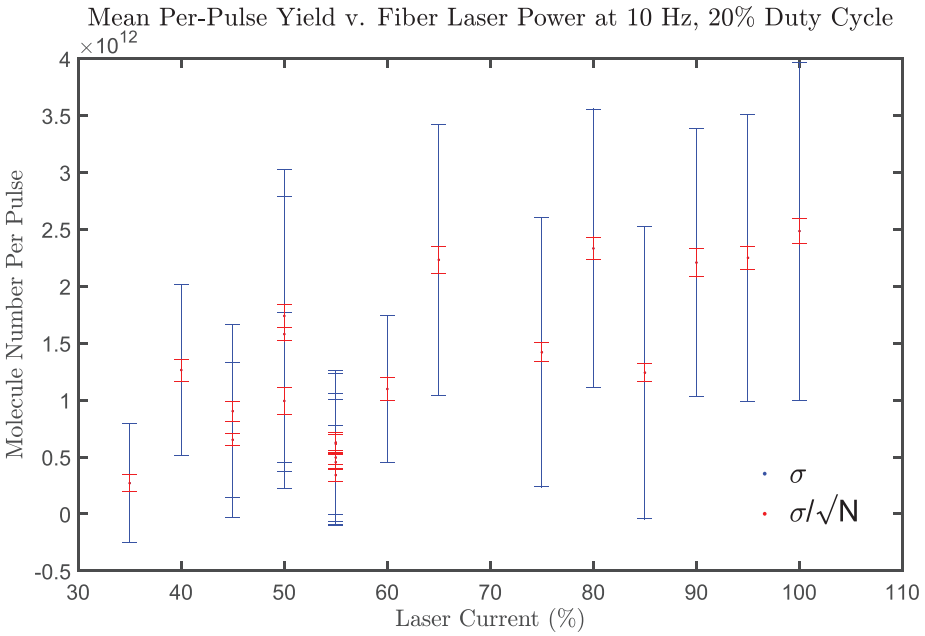

Figure 5.4.2: Cell exit signal (in $J = 1$ state molecules per pulse) v. fiber laser current. Firing at downstream target (2.5 cm from the cell exit) with 20 ms pulses at a repetition rate of 10.2 Hz. Cell temperature setpoint and typical value = 15.5 K.

Additional data taken on this same day indicate that the signal v. laser power correlation is much less clear for 40 ms pulses at 5.1 Hz. This corroborates the model that the peak beam flux saturates v. laser power, but the rise time is faster for higher laser power.

Figure 5.4.3 shows some additional time-dependent behavior of the signal throughout the run presented in Fig. 5.4.2. This plot reveals a few different kinds of hysteresis. The x-axes are trace numbers, each of with represents 0.39 s of data at 10.2 Hz with 4× averaging (150 traces ≈ 1 minute). Focusing first on the bottom plot, at lower laser power ($\leq 60\%$), the signal at first increases on a new target spot, then reaches a maximum, then gradually decreases. At higher laser power, ($\geq 80\%$), the signal starts out higher than at lower power, then gradually decreases. When switching from a high power ($\geq$ 80%) to a lower one ($< 60\%$), the signal generally vanishes and takes a few minutes to recover, even if signals were previously good at the lower power. The top plot shows that an average fiber laser power of ≈ 10 W suffices to keep the cell at 15.5 K with no cell heater input. It also reveals that it takes a few minutes for the cell to reach a thermal steady state.

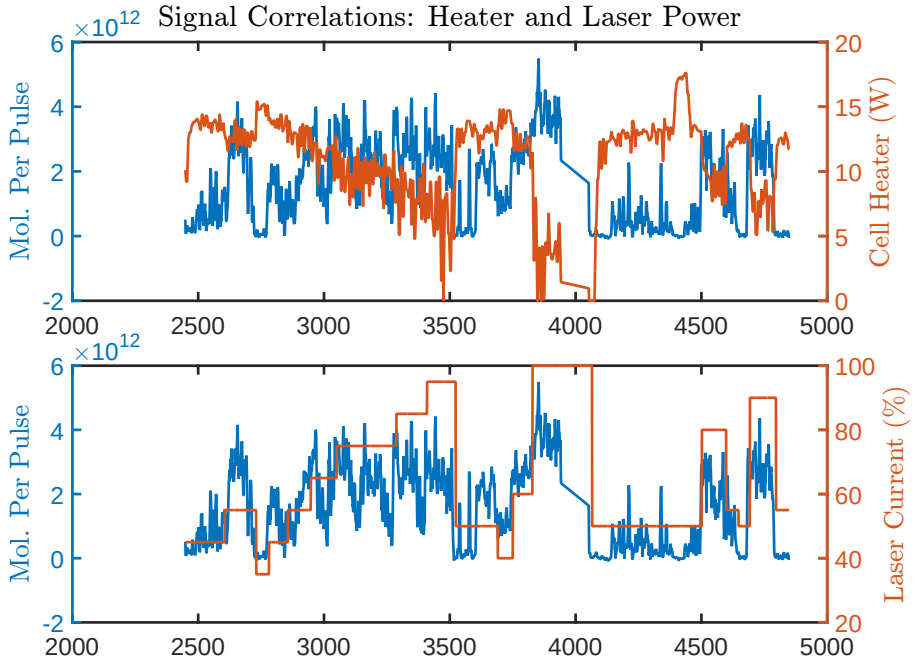

Figure 5.4.3: Additional time-dependent trends in the data shown in Fig. 5.4.2. **Top:** Correlation between molecule production (in $J = 1$ state molecules per pulse) and nominal cell heater power. **Bottom:** Correlation between molecule production and laser current.

Figure 5.4.4 shows a histogram of the cell exit signals from Fig. 5.4.2 at three different laser powers. This plot emphasizes that not only are the average signal sizes smaller at lower

laser power, but the fraction of spots that are dead on arrival is also much higher.

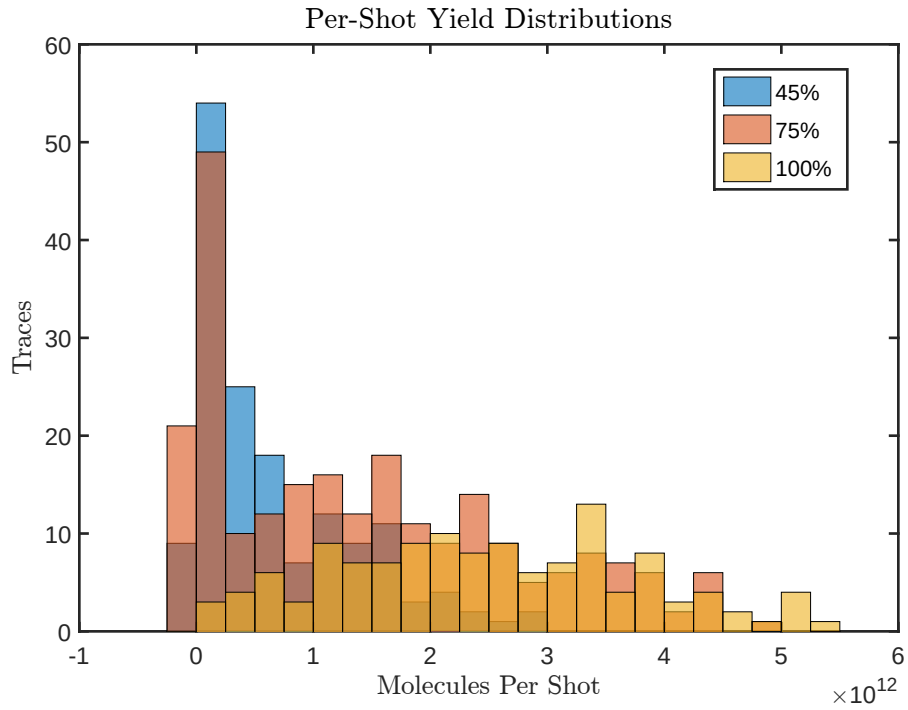

Figure 5.4.4: Histograms of the cell exit signals (in $J = 1$ state molecules per pulse) plotted in Fig. 5.4.2 for three different values of the laser current: 45% of the maximum in blue, 75% in orange, and 100% in yellow.

Figure 5.4.5 shows the absorption signal dependence on fiber laser power measured simultaneously outside the cell and on the beamline (see Fig. 5.2.1) at a distance of 77.5 cm from the cell. As in Fig. 5.4.2, the cell exit data reveal a roughly linearly increasing trend with a positive x-intercept, again suggesting that the trade-off between fiber laser power and duty cycle at constant heat load should favor high laser power under these running conditions. This stands in contrast to the data shown in Fig. 5.4.6.

The beamline signal in Fig. 5.4.5 shows a very similar trend to the cell exit signal. The overall transmission is about 0.2%, and the ratio between the upstream and downstream signals is independent of laser power within uncertainty, suggesting that the choice of fiber laser power (at roughly constant cell temperature) does not significantly affect the beam transmission or divergence.

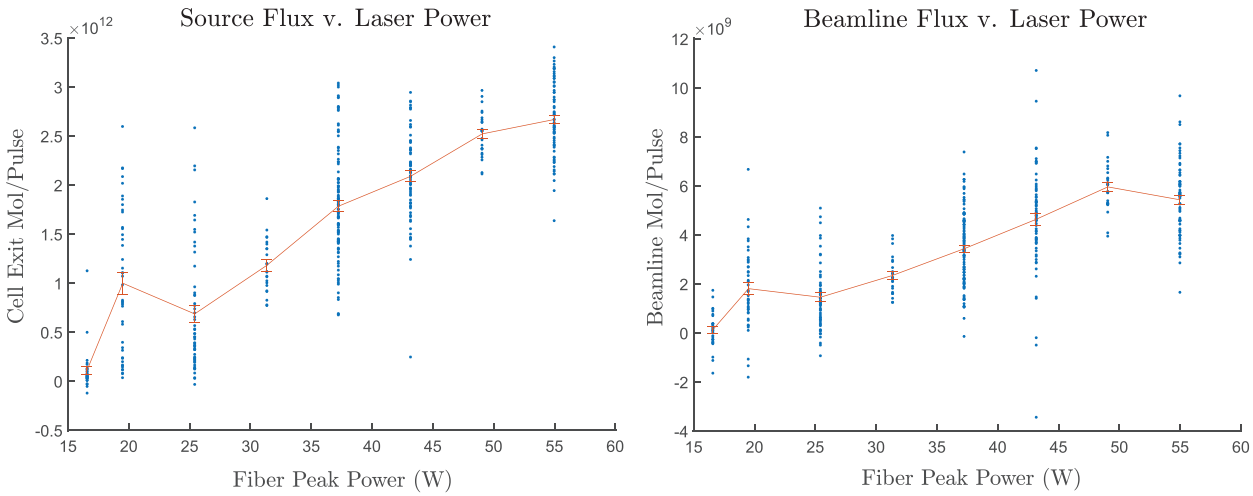

Figure 5.4.5: Signals (in $J = 1$ state molecules per pulse) at two different positions along the beamline as a function of fiber laser power. Firing at bottom half of upstream target with 20 ms pulses at a rep rate of 10.7 Hz. Cell top temperature was typically 15.5–16 K with a setpoint of 15.5 K. Measuring absorption on the $X \rightarrow C$ $Q(1)$ line. **Left:** Cell exit signal. **Right:** Simultaneous beamline signal measured using 5-pass absorption signal at 77.5 cm from the beam source.

Figure 5.4.6 shows the beam signal as a function of fiber laser power under a different set of running conditions from the previous data. The data was taken over six different days, and the decrease in signal during that time is shown in the right-hand plot. The right-hand plot shows that as the target becomes depleted, higher fiber laser powers are required to achieve the same signal levels observed at lower laser powers on a fresh target. The left-hand plot reveals a roughly linearly increasing trend in signal v. laser power with a positive y-intercept, suggesting that the trade-off between fiber laser power and duty cycle at constant heat load would favor high duty cycle under these running conditions, in contrast to the data shown in Fig. 5.4.5 and 5.4.2. The reason for the difference between this result and the previous ones remains a mystery.

Figure 5.4.7 shows a comparison between the three signal v. fiber laser power data sets presented in this section. For the orange and blue data sets, the yield is strictly better at higher power. For the purple data set, Fig. 5.4.7(b) shows that if the signal per pulse is

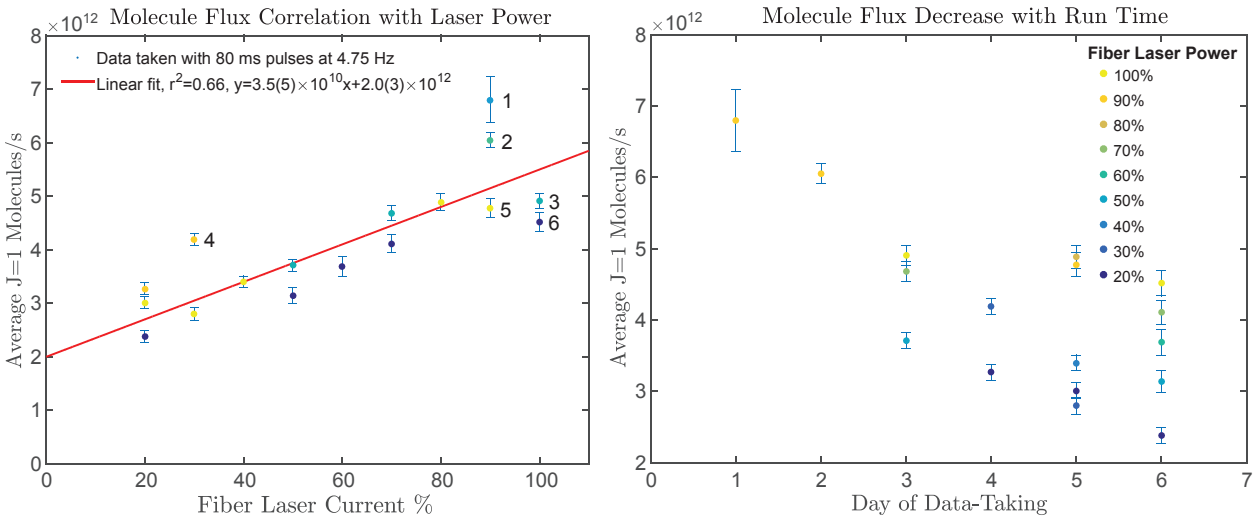

Figure 5.4.6: Beam signal ($J = 1$ state molecules/s) measured at the 4 K shield exit plotted as a function of the fiber laser current and run day. Firing at upstream target 5 cm from the cell exit with 80 ms pulses at a repetition rate of 4.75 Hz. Measuring the absorption signal on the $X \rightarrow C$ $Q(1)$ line just outside the 4 K shield conical collimator. The two plots show the same data with different x-axes. **Left:** Signal v. fiber laser current and linear fit. The different days on which this data was taken are color coded and labeled. **Right:** Signal v. run day. The different fiber laser powers at which this data was taken are color coded and labeled.

independent of repetition rate, a lower laser power at a higher duty cycle is favored.

Figure 5.4.8 shows the pulse shape variation with pulse power using average data traces from the days labeled #3 and #5 in Fig. 5.4.6. With all else equal, pulses taken at higher laser power typically have a faster rise time and reach a higher peak signal—though not always, as can be seen by comparing the 20% and 30% traces and the 80% and 90% traces.

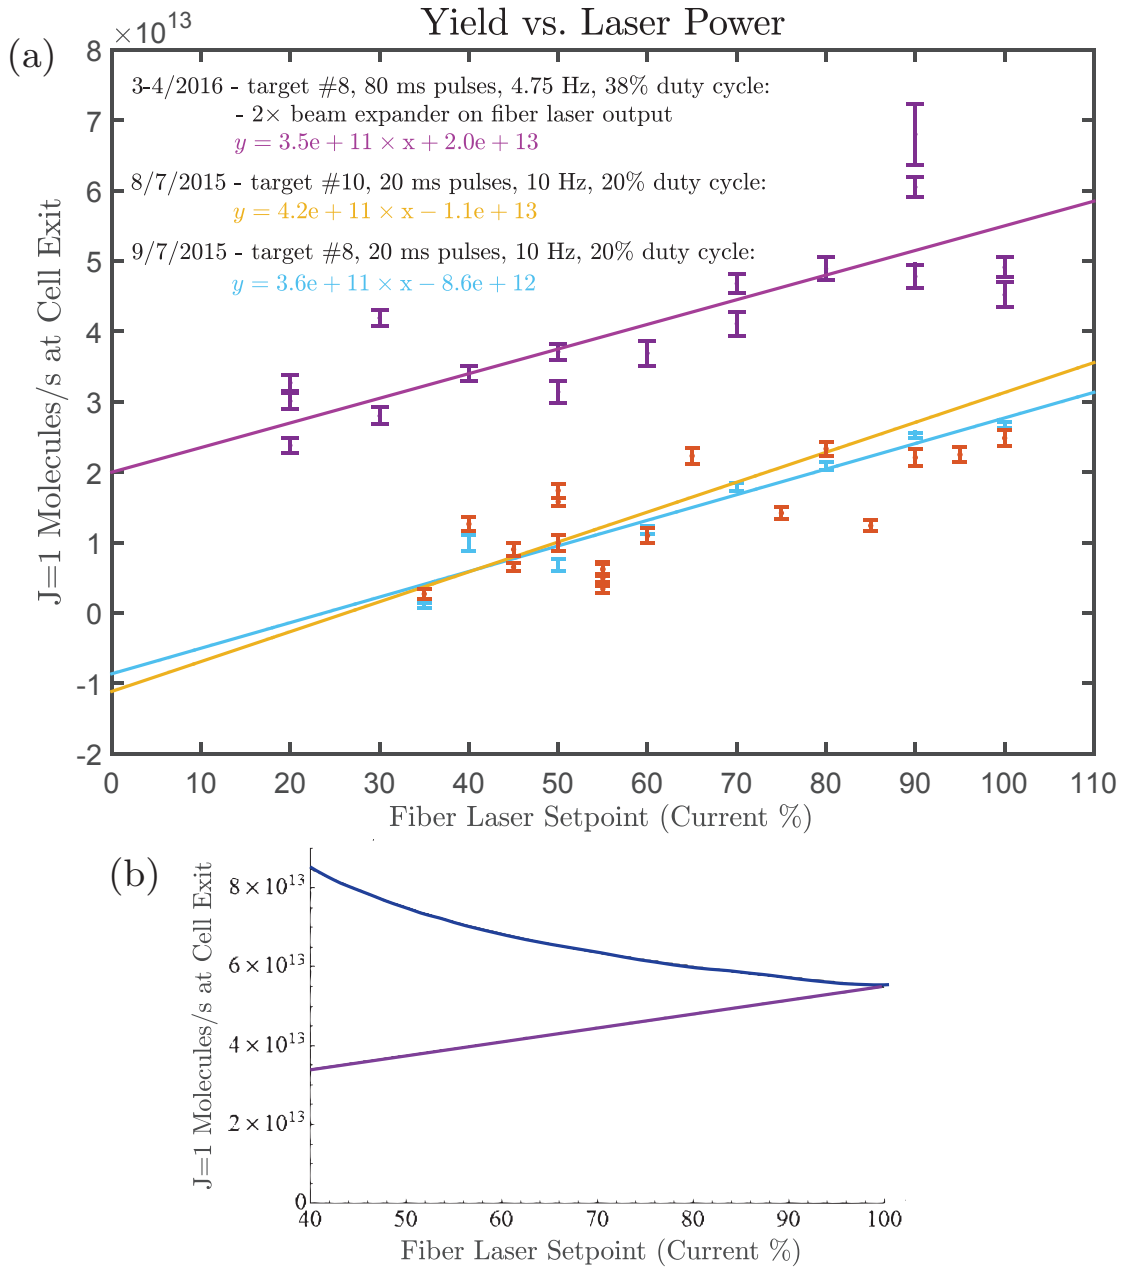

Figure 5.4.7: Comparison between signal v. fiber laser power results acquired under different running conditions. **(a)** Data from Fig. 5.4.2 (blue), 5.4.5 (orange), and 5.4.6 (purple) plotted with linear fits. Because the purple data was taken outside the 4 K shield, it had to be scaled by the transmission loss factor (≈ 10) for comparison with the other two data sets; however, the large uncertainty on the transmission factor may distort the comparison. **(b) Purple:** Fit line from data in Fig. 5.4.7 (purple line in (a)) showing the molecule beam flux v. fiber laser current for 80 ms pulses at 4.75 Hz. **Blue:** Calculated signal v. fiber laser current with a constant time-averaged fiber laser heat load (i.e., fiber laser power × duty cycle) of 19 W (the actual heat load for the purple line at 100% current) assuming the same per-pulse signals as the purple line. It is interesting to compare this curve with the data in Fig. 5.4.12, which was taken under running conditions approximating these assumptions.

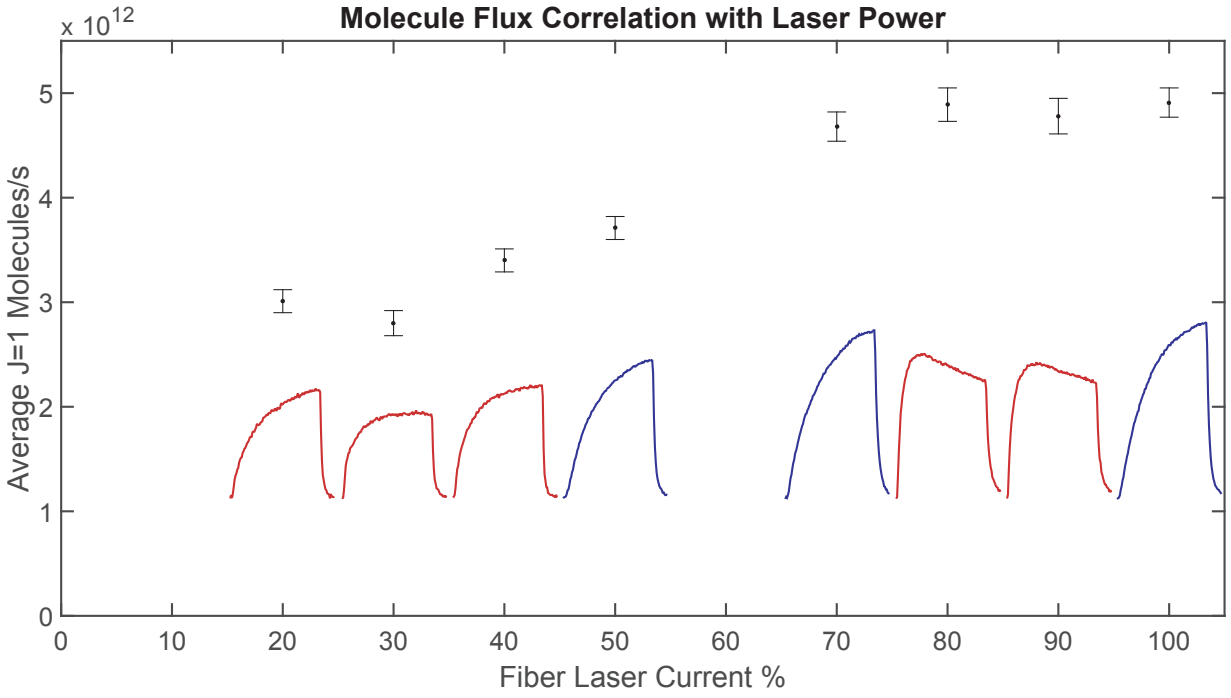

Figure 5.4.8: Average 4 K exit molecule flux and pulse shapes from the days labeled #3 (blue) and #5 (red) in Fig. 5.4.6. Apart from the different run days, these data were taken under identical running conditions.

5.4.2 Pulse Width

To determine the optimal pulse width, we ran the ThO beam with a constant instantaneous laser power and duty cycle and varied only the pulse duration while measuring the absorption signal outside the 4 K shield exit. The resulting measured beam fluxes (in $J = 1$ molecules/s) and pulse shapes are shown in Fig. 5.4.9. We find that the optimum pulse width is about 80 ms; however, the peak is broad, and the signals at 40 and 120 ms are within typical signal size fluctuations of the maximum. By examining the pulse shapes, we can explain the position of this optimum: The instantaneous signal flux reliably reaches a peak by around 40 ms and begins to drop off by around 120 ms (though there are large variations in these timescales). This means that on average the time spent at peak instantaneous signals is maximized for pulse widths of ≈ 80 ms.

Figure 5.4.10 illustrates the variability in the rise times and pulse shapes, even for data

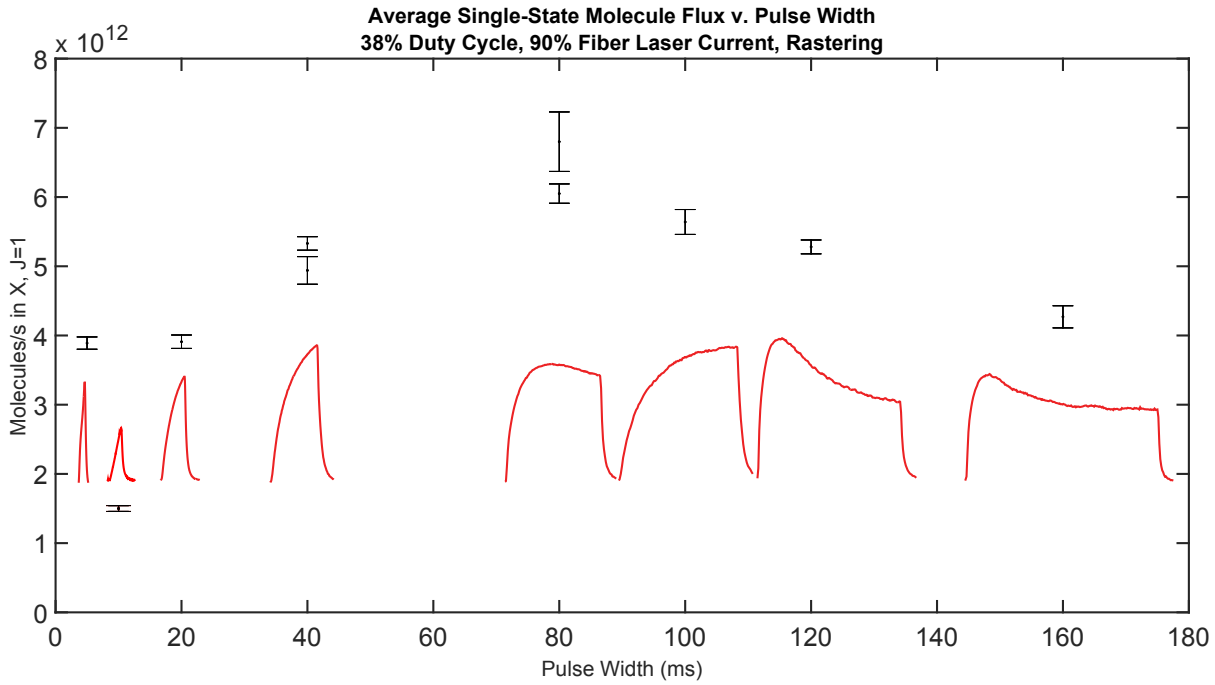

Figure 5.4.9: Average 4 K exit molecule flux and pulse shapes plotted as a function of pulse width. Firing at upstream target (5 cm from the cell exit) with a constant duty cycle of 38% and the fiber laser current set to 90% while measuring the absorption signal on the $X \rightarrow C$ $Q(1)$ line just outside the 4 K shield conical collimator.

taken under essentially the same running conditions. The cause of the dramatically variable signal rise times remains mysterious.

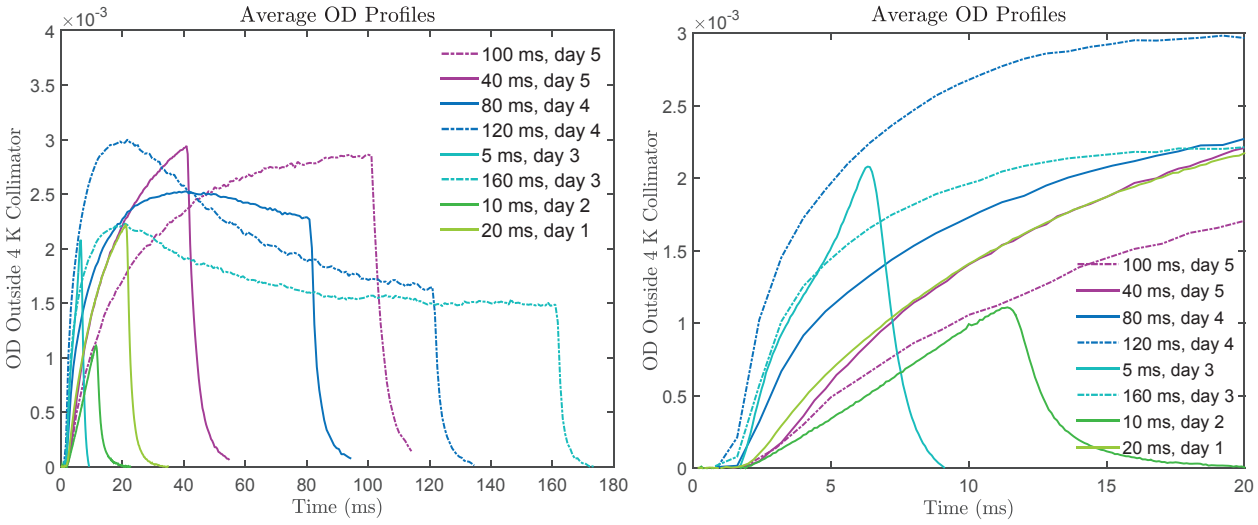

Figure 5.4.10: **Left:** Pulse shape variation with laser pulse width using average data traces from Fig. 5.4.9. Different days of data-taking are color-coded and labeled in the legend. **Right:** Zoom-in on early part of pulses from the left plot. Since the data are taken under identical conditions except for the fiber laser pulse duration, it is unknown why the rise times are so variable.

5.4.3 Repetition Rate

The data in this chapter was taken at many different repetition rates. Since the pulses are largely independent, the selected repetition rate is determined by the requirement that the cell temperature remain below 18 K, as discussed in Section 5.3. See the summary at the head of Section 5.4 for further discussion.

5.4.4 Modulation Frequency and Duty Cycle

In Section 5.4.1, we saw that high instantaneous laser powers tended to produce larger and more reliable signals; however, as discussed in Section 5.3, high continuous laser powers also produce large heat loads and limit the duty cycle at which we can run the beam. In an effort to optimize this trade-off, we investigated the possibility of modulating the fiber laser pulse with a rectangular wave so that it would be on at full power for only a fraction of the time that the molecule beam was on.

If the local thermalization timescale of the target surface is rapid compared to the modulation frequency, we would expect this method to produce larger instantaneous temperatures

and therefore larger signals than an un-modulated laser pulse with the same average (and therefore lower instantaneous) power. The results in this section do not conclusively support this model, but they do give some indication that modulated pulses are preferable to un-modulated ones.

In Section 5.4.4, we discuss studies of the modulation frequency with a fixed 50% modulation duty cycle (i.e. square wave modulation), and in Section 5.4.4 we discuss studies where the modulation duty cycle was varied while the overall heat load was held constant by compensating with the pulse repetition rate.

Varying modulation frequency at 50% modulation duty cycle

Figure 5.4.11 shows the beam signal (in $J = 1$ molecules/s outside the cell exit) as a function of laser power for three different modulation settings, 200 Hz with 50% modulation duty cycle, 400 Hz with 50% modulation duty cycle, and no modulation. Although the signals are noisy (1σ signal fluctuations are given by the large error bars, while the standard error in the mean is given by the small error bars), this plot hints at some advantages of using modulation: The maximum signals are comparable with and without modulation at 200 or 400 Hz, and they occur when the time-averaged heat load from the fiber laser is approximately 12 W. The fact that the modulated signals have no distinct maximum even at 100% fiber laser power suggests that gains might still be won by varying the modulation duty cycle or the pulse repetition rate.

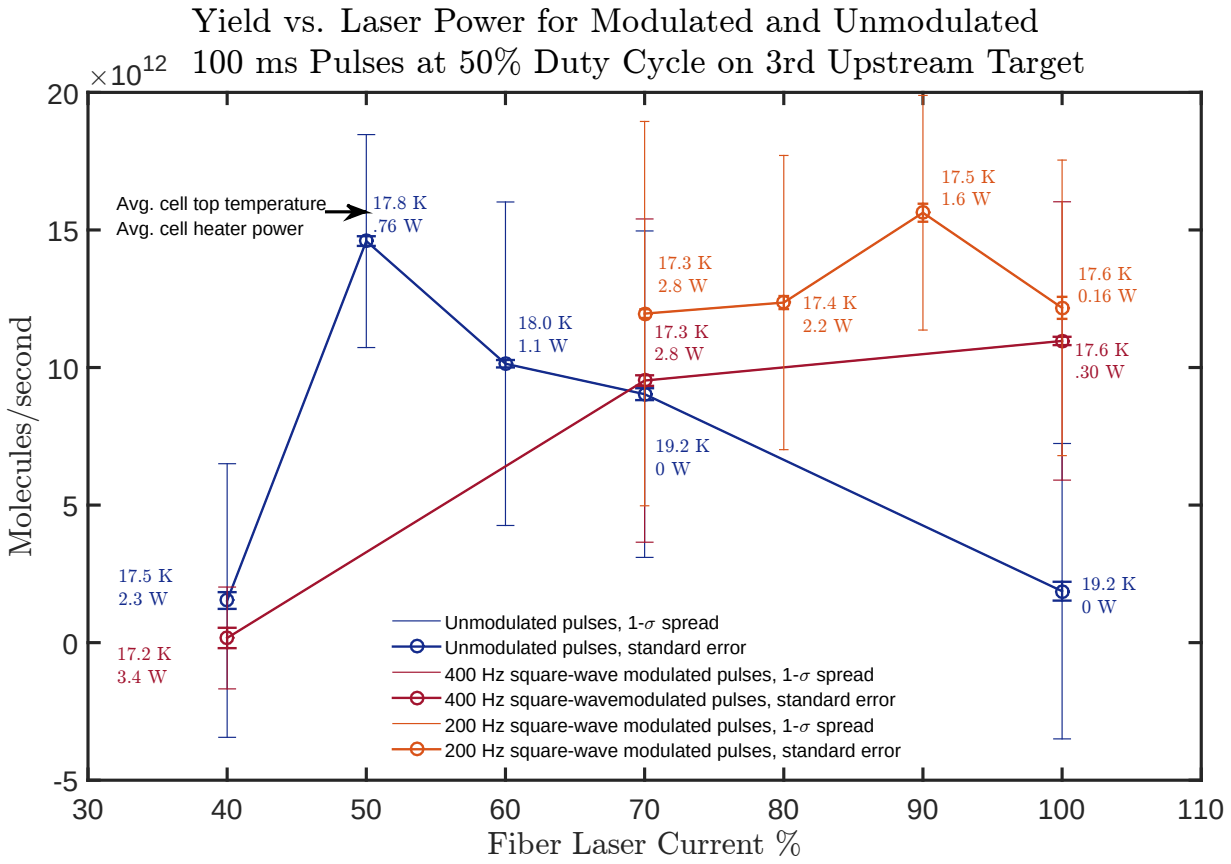

Figure 5.4.11: Cell exit signal in $J = 1$ molecules/s v. fiber laser power with different modulation settings. Measured via absorption on the $X \rightarrow C\ Q(1)$ line. Using 3rd upstream target (7 cm from the cell exit) with 100 ms pulses at a repetition rate of 5 Hz. The modulated pulses have a modulation duty cycle of 50%. The cell temperature setpoint was 17 K, while the actual average cell top temperatures and corresponding cell heater powers used to servo the temperature are labeled on the plot.

Varying modulation duty cycle at 24% overall duty cycle

In this section, we explore varying the modulation duty cycle while keeping the overall heat load constant. For this data, the fiber laser current was kept at 100%, the fiber laser modulation was a 400 Hz frequency rectangular wave, and the pulse repetition rate was chosen to keep the overall duty cycle (i.e., fractional on-time of the laser) at 24%.

Figure 5.4.12 shows the cell and 4 K shield exit signals (in units of integrated OD) as a function of modulation duty cycle for both 25 ms and 50 ms overall pulse widths. The discrepancy between the downstream signal transmission fractions for the two pulse widths is surprising, and the transmission fractions are also both a factor of almost 3 too low (Other data show downstream $Q(2)$/upstream $Q(1)$ ratios of $\approx 14\%$), suggesting that either the downstream laser was poorly tuned for both runs or that it was misaligned and only catching a narrow sliver of the beam. Ignoring the downstream signal, the upstream results suggest that the signals decrease with modulation duty cycle and are optimal for duty cycles below 60%.

The explanation for this trend is indicated in Fig. 5.4.13, which shows the average pulse shapes for the data in Fig. 5.4.12. The 25 ms pulse width data was taken the day before the 50 ms pulse data and shows a faster rise time. Such rise-time drifts correlated with running day are also observed in Fig. 5.4.8 and 5.4.10. The signal rise time and peak signal are also roughly correlated with modulation duty cycle, just as these quantities are observed to be correlated with laser power in Fig. 5.4.8; however, the scaling between signal per pulse and modulation duty cycle is slower than linear, causing lower duty cycles with higher repetition rates to be generally favored. When the modulation duty cycle is below 40%, however, the pulses begin to overlap, which is undesirable. In addition, the rise times at lower duty cycle are not as reliably fast as those at higher duty cycle. A modulation duty cycle of about 50–60% is therefore roughly optimal.

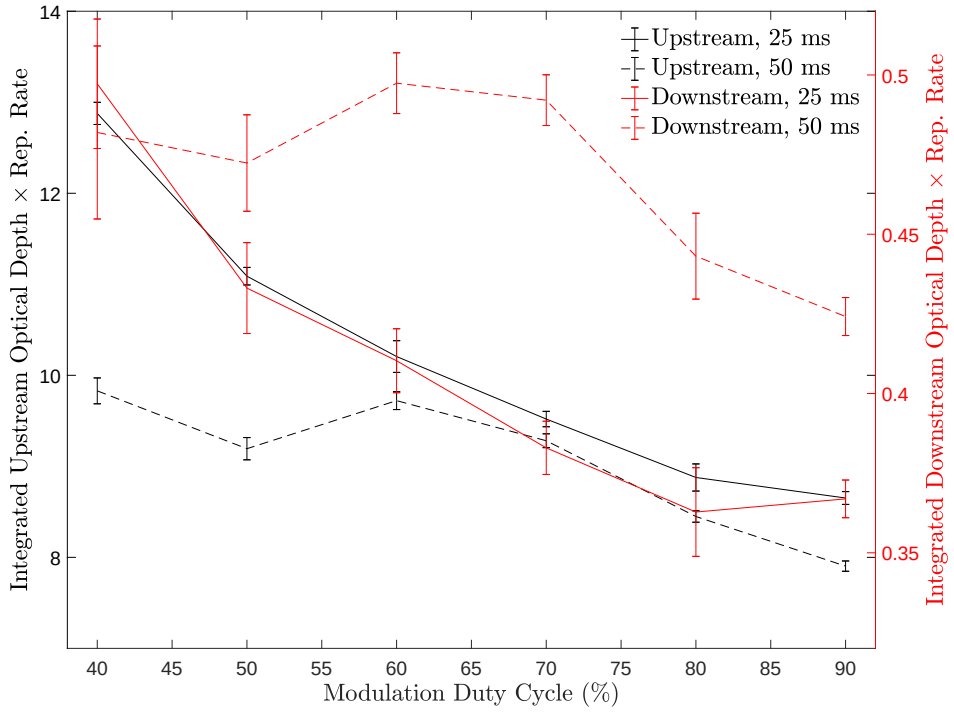

Figure 5.4.12: Cell and 4 K shield exit signals v. modulation duty cycle at different pulse widths: Measuring absorption on the $X \rightarrow C$ $Q(1)$ line just after the cell exit and the $X \rightarrow C$ $Q(2)$ line just after the 4 K shield exit aperture. Firing at the 3rd upstream target (7 cm from cell exit) with a peak fiber laser current of 100%. The fiber laser pulses were modulated at 400 Hz with a variable modulation duty cycle and a repetition rate chosen to keep the overall duty cycle at 24%. The focusing lens had a 500 mm focal length and was positioned at about 490 mm from the target, and the laser spot was randomly hopped around most of the target surface with a frequency of 5 Hz. The cell temperature was typically near the 17 K setpoint. Figure and data analysis by Adam West.

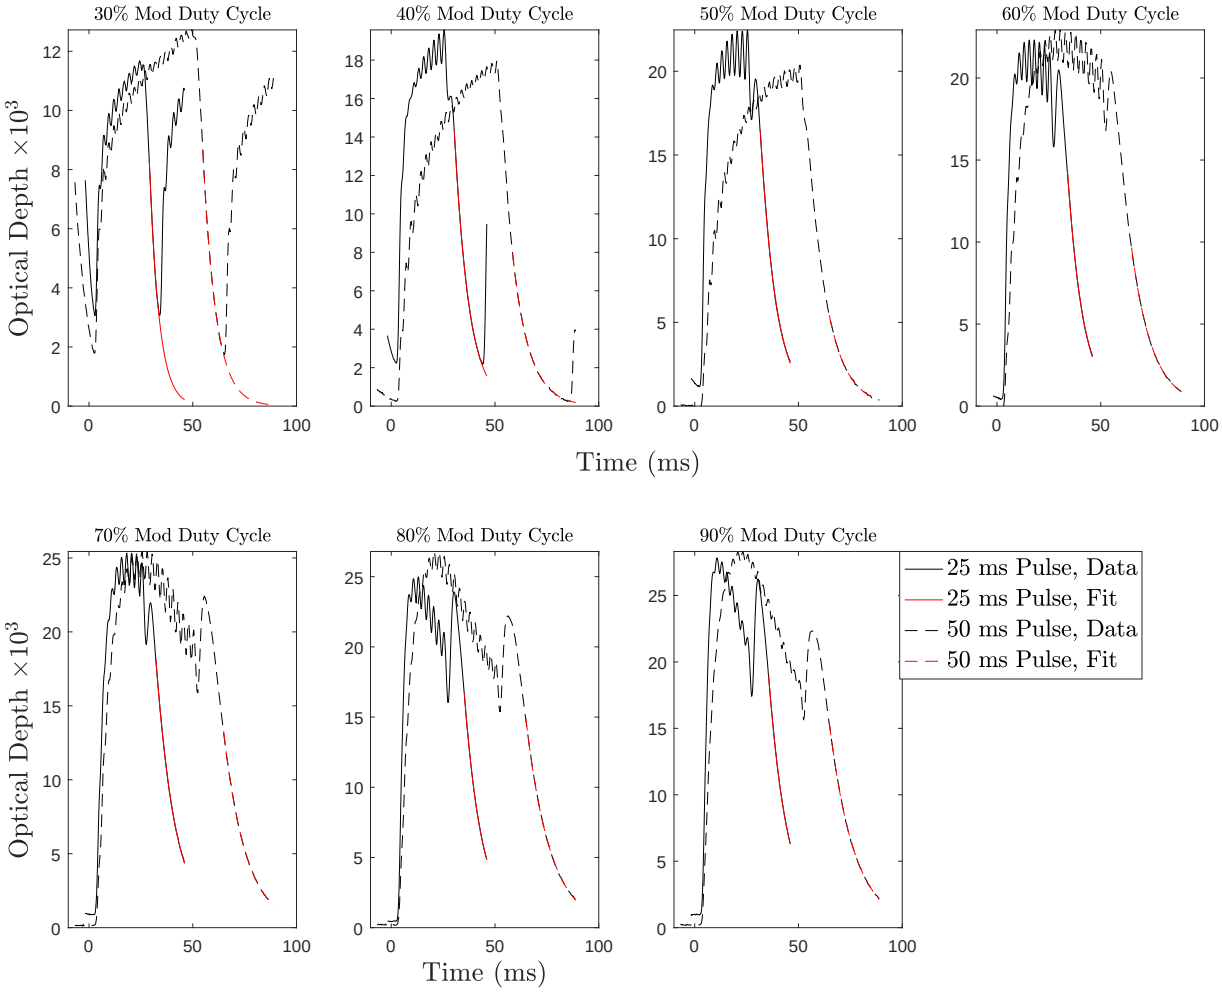

Figure 5.4.13: Comparison between the 25 and 50 ms pulse shapes at varying modulation duty cycle described in Fig. 5.4.12. The exponential fits (red) are used to determine the background level for each pulse. Data analysis and figure by Adam West.

5.5 Laser Spot

Another set of beam source parameters we sought to optimize are those relating to the laser spot on the target. These parameters include the fiber laser focal size, the amount of time spent firing at a single target spot, whether the laser spot position is rastered during a pulse, and the program for exploring the surface area of the target. We found that the behavior of the target surface was one of the most complicated things to understand about this source, especially because the beam properties change so much as the target surface ages. The data from the "stress-test" run in Section 5.7 shows the dramatic difference between signals from a fresh target surface and those from a used surface. New targets produce beam signals that are consistent, large, and relatively independent of running parameters, while older targets exhibit more shot-to-shot fluctuations, more dead and short-lived spots, and more stringent requirements for e.g. fiber laser power and focal size.

In order to better understand the target surface usage, Jacob Baron performed the analysis shown in Fig 5.5.1. For this analysis, the target surface was divided into an imaginary grid of 100 μm squares, and each time the galvo aimed the laser into one of the squares, it was counted as a "visit." The fiber laser focus was 140 μm FWHM during this run. The galvo was programmed to hop randomly to a new location on the target whenever the signal fell below a fixed threshold (as in item 4 in Section 5.1), but it could return to the same spot later. Figure 5.5.1 shows that while the target spots do recover nearly to their previous signal levels on each sequential visit, the target spots become slightly shorter-lived and less productive each time. After 10 visits, the total yield produced from a spot has fallen by nearly 50%. If we take this as the criterion for "spot death," we conclude that each $(100\,\mu\text{m})^2$ region on the $3\,\text{cm}^2$ target surface can be visited 10 times with an average of 25 shots before it is depleted. If the source is run with 20 ms pulses at 25 Hz, then we should be able to run for ~ 80 hours on a single target. This estimate is only slightly higher than the 70 hours of total run time achieved during the "stress-test" run described in Section 5.7.

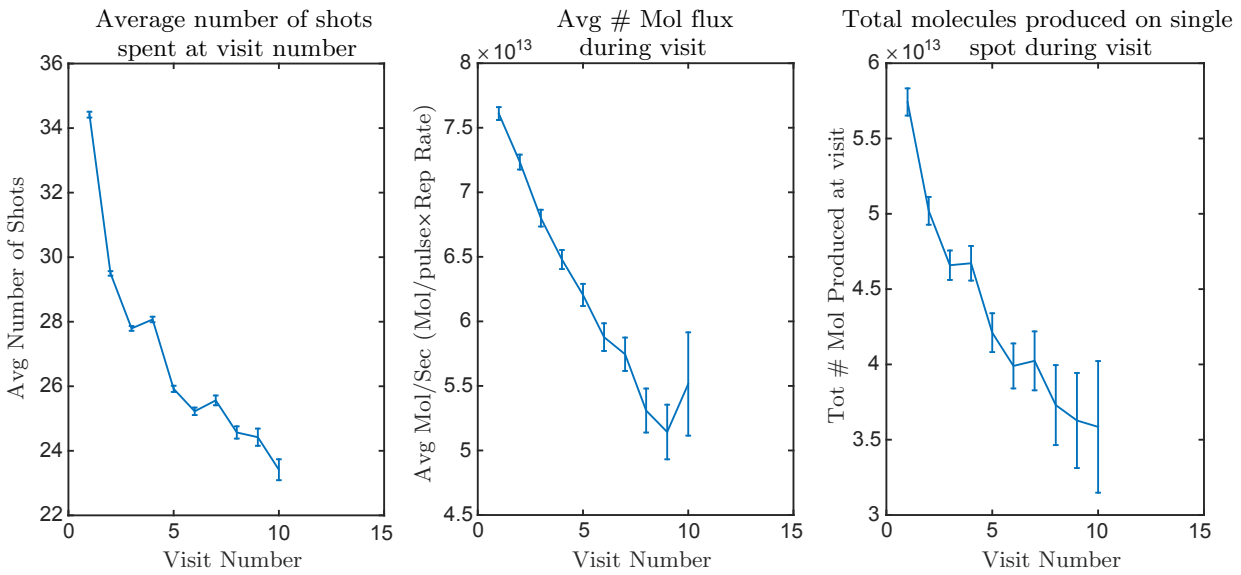

Figure 5.5.1: Target spot depletion upon repeated visits. All molecule signals refer to the total population of $|X, J = 1\rangle$. **Left:** This plot shows that the target spots do recover between visits but deplete more rapidly with each sequential visit. The number of shots providing good signal decreases by 25% after 5 visits. **Middle:** This plot shows that the average molecule flux from a spot also decreases slightly with each subsequent visit and falls by 20% after 5 visits. **Right:** This plot shows the decrease in the total molecule production from a spot in each visit, which has fallen by 30% after 5 visits. Figure, data, and analysis from Jacob Baron.

Figure 5.5.2 shows the length scale over which the target surface is depleted by the laser pulse. As expected, this length scale is comparable to the beam diameter at the focus, ≈ 100 μm.

Results: Regarding spot positioning, we have found in data not shown here that rastering (i.e. continuously scanning the laser) tends to add noise to the signal on short timescales while smoothing out longer-timescale drifts. In some cases, adding a small "micro-raster" with a ~ 1 mm amplitude and a ~ 1 Hz frequency while spot hopping may help with the instantaneous signals and target region longevity on a depleted target (see Section 5.7), but on a fresh target face, spot hopping without rastering tends to give more consistent signals.

In principle, spot hopping with a signal threshold should be strictly better than spot hopping without one; however, on a fresh target face, simply hopping to a new spot at

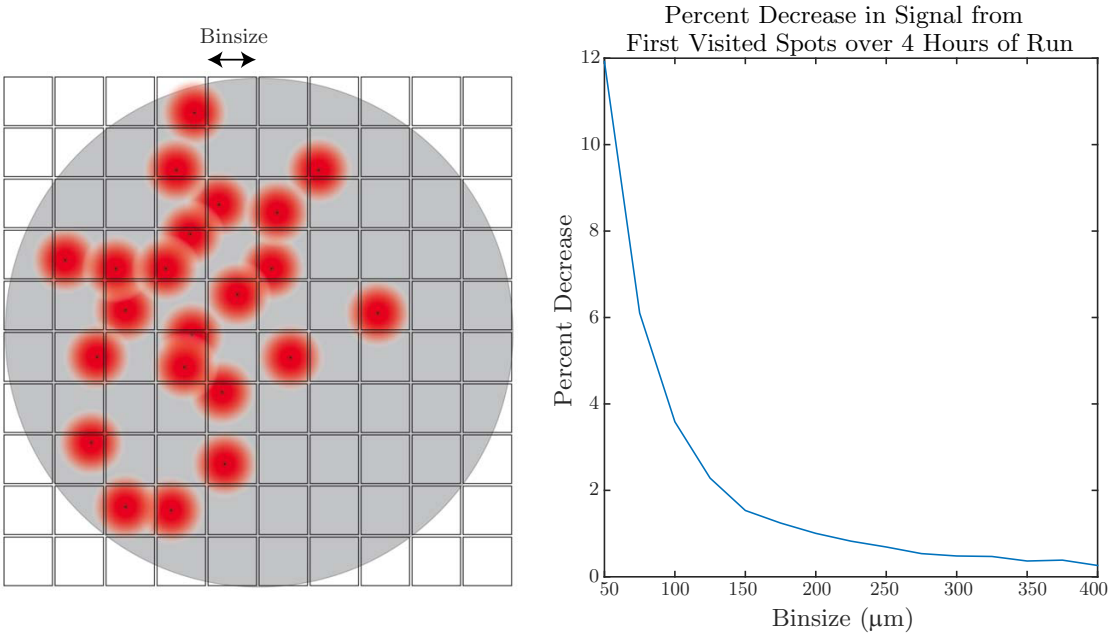

Figure 5.5.2: Target spot depletion correlation length scale. **Left:** The target surface is divided into a grid with regions of width "binsize," and each bin is considered "first visited" the first time the center of the fiber laser spot falls into a bin. **Right:** The length scale of target surface depletion can be defined as the bin size below which spots are significantly less productive on their "first visit" as the run progresses. Such behavior indicates that the "first visit" is not truly the "first," as previous laser spots must have overlapped significantly with these bins in order to deplete them. This plot shows that the relevant length scale for target depletion is ≈ 100 μm, which is comparable to the diameter of the laser spot. Figure, data, and analysis from Jacob Baron.

every pulse gives the maximum signal levels, as we found during the long run discussed in Section 5.7. As the target surface starts to age and deplete, turning on the signal threshold helps to find and stay on the remaining good spots.

Regarding the optimum spot size, the data is at present inconclusive. We performed several runs in which we translated lenses of different focal length (300 mm, 400 mm, and 500 mm) through their focal position and for several cm on either side. In some instances, the optimum signals occurred with the target at the focus, while at others, a spot size of up to approximately 1 mm appeared to give the best signals. When the beam size became much larger than 1 mm, the signals fell off rapidly.

Conclusion: On a fresh target, hop to a new target spot with every pulse in order to maximize the beam flux. Turn on the signal threshold spot-hopping algorithm as the target starts to age and the signals decrease. Turn on "micro-rastering" as desired to smooth out slow drifts.

Before running on a new target, optimize the lens position. Re-optimize periodically as the target ages. Expect that the target signals will fall by a factor of about 2 after 1 week of run time.

5.6 Beam Properties

In this section, we describe our measurements of the forward velocity and velocity spread, rotational temperature, and divergence of the ThO beam. These properties are summarized in Table 6.1.1. We find that the thermochemical beam is slightly hotter, faster, and more divergent than the ablation source beam, but we do not expect any of these changes to significantly degrade the beam performance or eEDM sensitivity.

5.6.1 Beam Velocity

Figure 5.6.1 illustrates the method used to measure the forward velocity of the molecule beam. A 690 nm laser tuned to address the $X \rightarrow C$ $Q(1)$ line is split into a strong and a weak beam. The weak beam is used to perform absorption spectroscopy downstream of the

4 K shield collimator, which serves to normalize beam fluctuations and keep the laser on resonance with the zero-forward-velocity molecule class. The strong beam is sent through a fiber EOM (Jenoptik PM705) and then through the molecule beamline 1 m from the source at an angle of 45° with respect to the beam. The EOM applies tunable sidebands to match the first-order Doppler shift of the molecule beam. A PMT collects fluorescence from the velocity probe beam. When the reference beam and the probe beam are both on resonance, their relative detuning (set by the RF frequency applied to the fiber EOM) can be used to determine the forward velocity of the molecule beam.

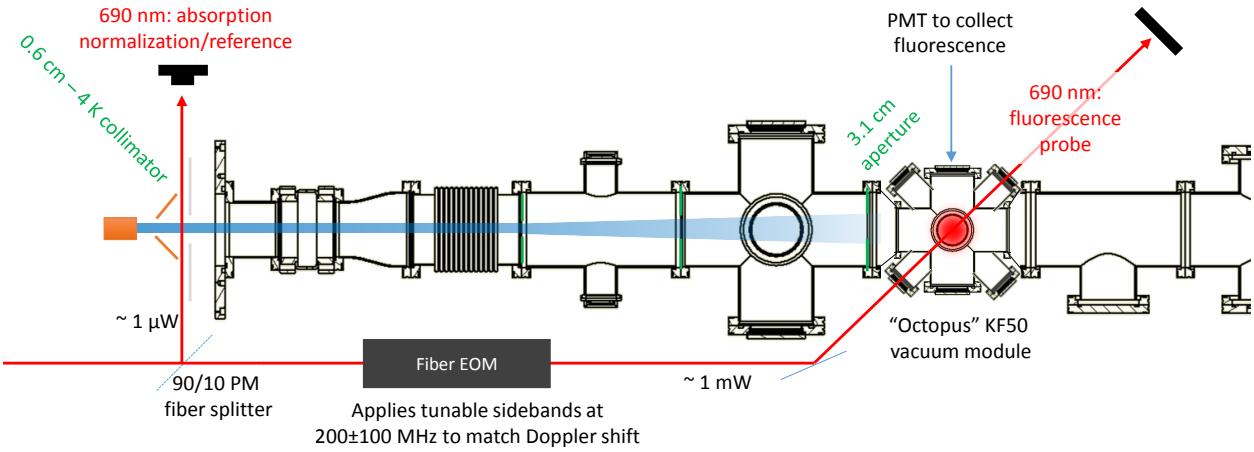

Figure 5.6.1: Schematic of the test beamline 45° fluorescence spectroscopy setup used for the data in Section 5.6.1. The vacuum components and apertures are described in Fig. 5.2.1. The measurement scheme is described in the text.

Figure 5.6.2 shows a measurement of the forward velocity and velocity spread as a function of the fiber laser power, while Fig. 5.6.3 shows the forward velocity and velocity spread as a function of the fiber laser pulse width.

Results: The mean beam velocity is typically 195 m/s, and the velocity spread is 45 m/s FWHM (standard deviation 19 m/s). The drifts and shot-to-shot fluctuations in the mean velocity (velocity spread) are of order 10% (20%). We found a small, nonzero correlation between the forward velocity and the fiber laser power with a slope of 2 m/s per 10% increase

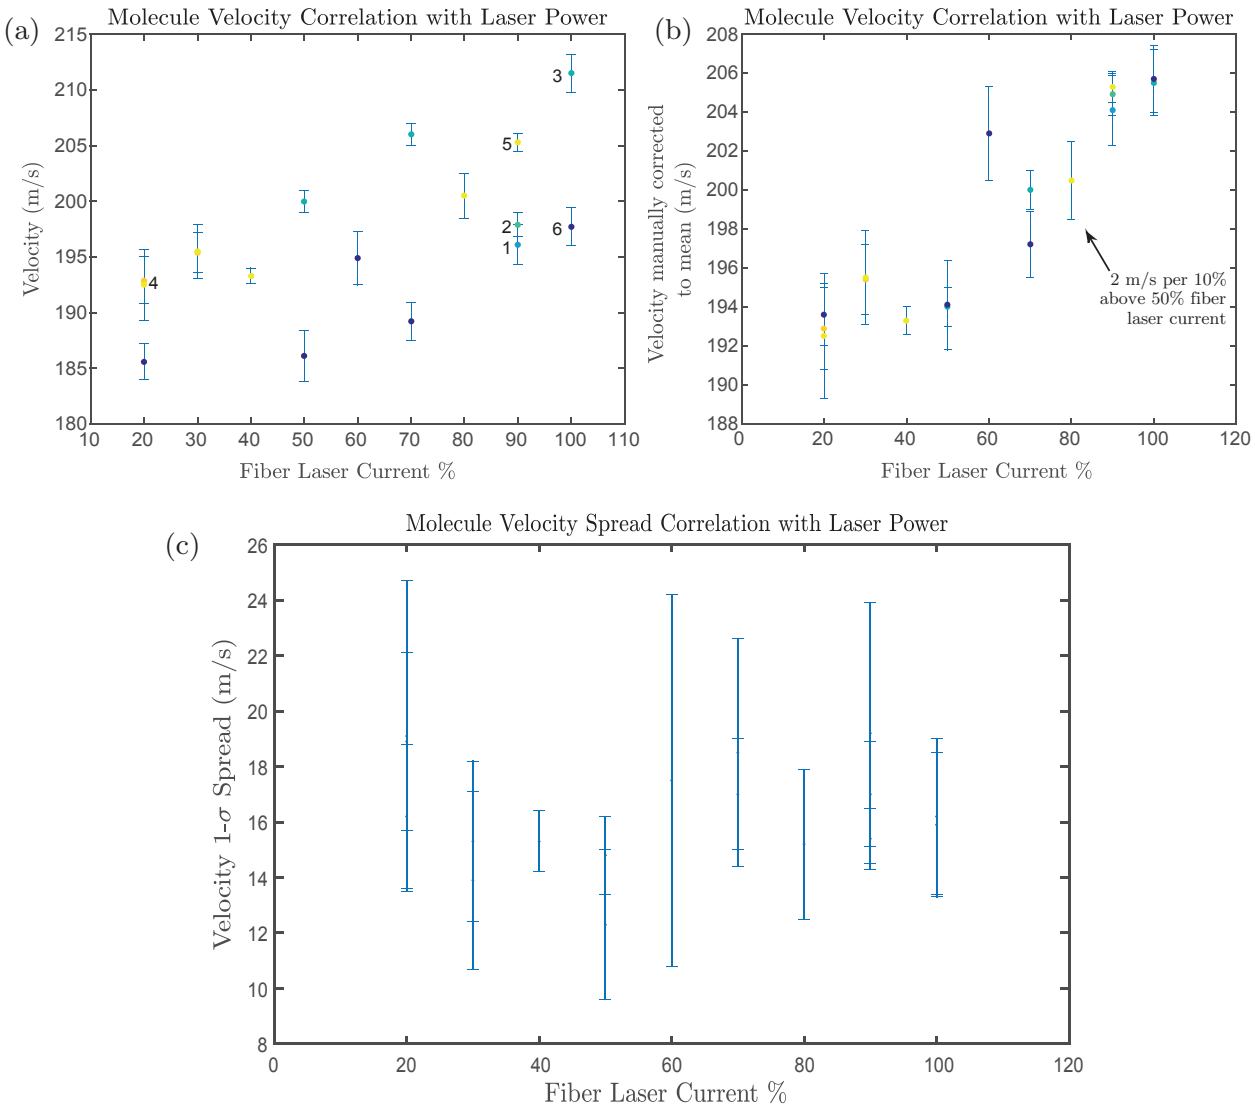

Figure 5.6.2: Beam velocity and velocity spread as a function of fiber laser power. **(a)** Mean molecule velocity with statistical error bars. The different days on which the data were taken are color coded and labeled. The generally increasing trend in the data is obscured by the day-to-day velocity drifts. **(b)** Same data as in (a) but with the velocity measurements for each day manually adjusted so that their mean approximately overlaps with the overall mean. This adjustment more clearly reveals the small increasing slope of velocity with fiber laser current. **(c)** Forward velocity 1σ Gaussian width and statistical uncertainty as a function of laser current. No correlation is observed.

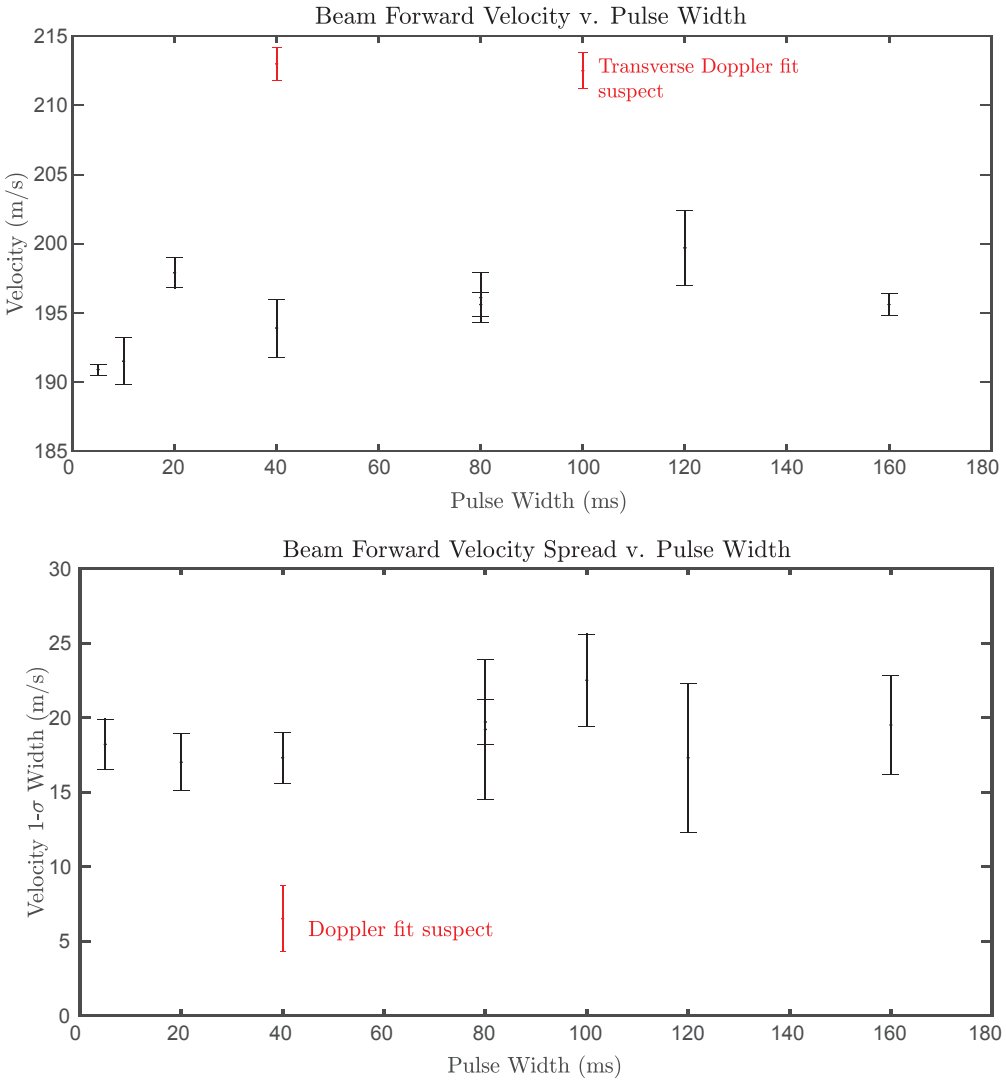

Figure 5.6.3: Beam velocity (top) and 1σ Gaussian velocity width (bottom) v. fiber laser pulse width. Apart from a few outliers (red), which are correlated with and likely attributable to noise in either the normalization or the probe beam Gaussian fit, these quantities are independent of pulse width within uncertainty. Analysis and figure from Adam West.

in fiber laser current.

Conclusion: The velocity is independent of all studied parameters within reasonable ranges, so we should not need to worry about optimizing the velocity for EDM sensitivity.

Velocity behavior during pulses

Beam velocity variations within a pulse produce varying precession times in the eEDM experiment. These variations are a source of technical noise. We have therefore studied the beam velocity as a function of time within a pulse to look for correlations.

Figure 5.6.4 shows the average beam velocity profile v. time for a 20 ms pulse width, while Fig. 5.6.5 shows the same results for an 80 ms pulse width. Whereas the 20 ms pulse width data does not exhibit a statistically significant time-dependence of the beam velocity, the peak velocity class in the 80 ms pulse width data has an increasing slope of ≈ 250 cm/s per ms. These measurements suggest that for shorter pulse lengths, the velocity is independent of time within a pulse, while for longer pulse lengths, the velocity gradually increases as a function of time. This is consistent with our expectation that the forward velocity should vary approximately as the square root of the buffer gas temperature [106, 107], which gradually increases as the cell is heated by the laser. Notably, the increasing trend in Fig. 5.6.5 is not always repeatable: Some of our measurements for pulse widths of 100 ms and longer do not show a correlation between velocity and time within the pulse.

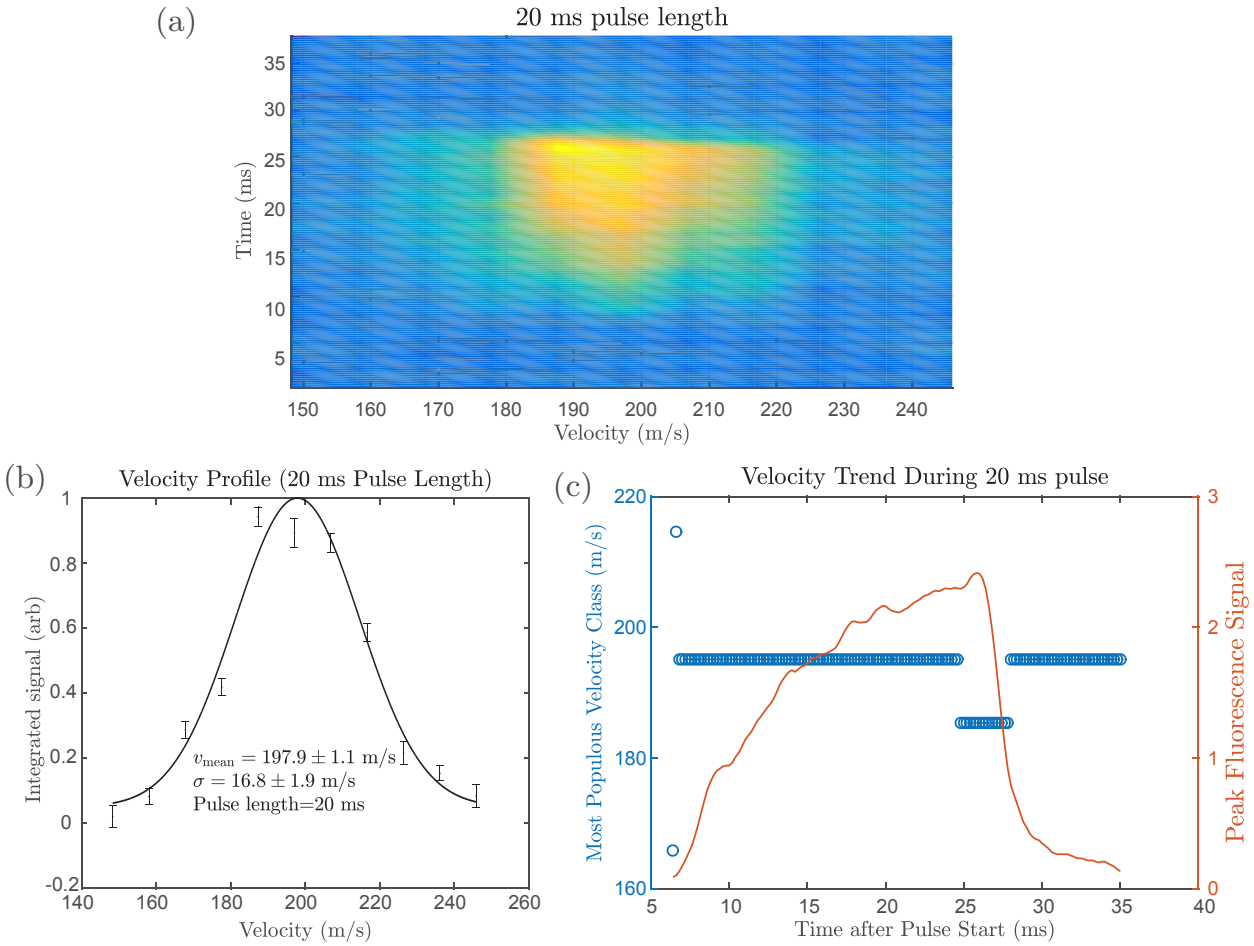

Figure 5.6.4: Average beam velocity profile for a 20 ms laser pulse as a function of time after the start of the laser pulse. Velocity is measured at a distance of 1 m from the beam source. **(a)** Velocity population surface plot v. time after the start of the laser pulse. **(b)** Velocity profile averaged over time at each velocity class and fit to a Gaussian distribution. **(c)** Peak fluorescence signal in the probe region v. time after the start of the laser pulse (orange) and most populous velocity class in each time interval (blue).

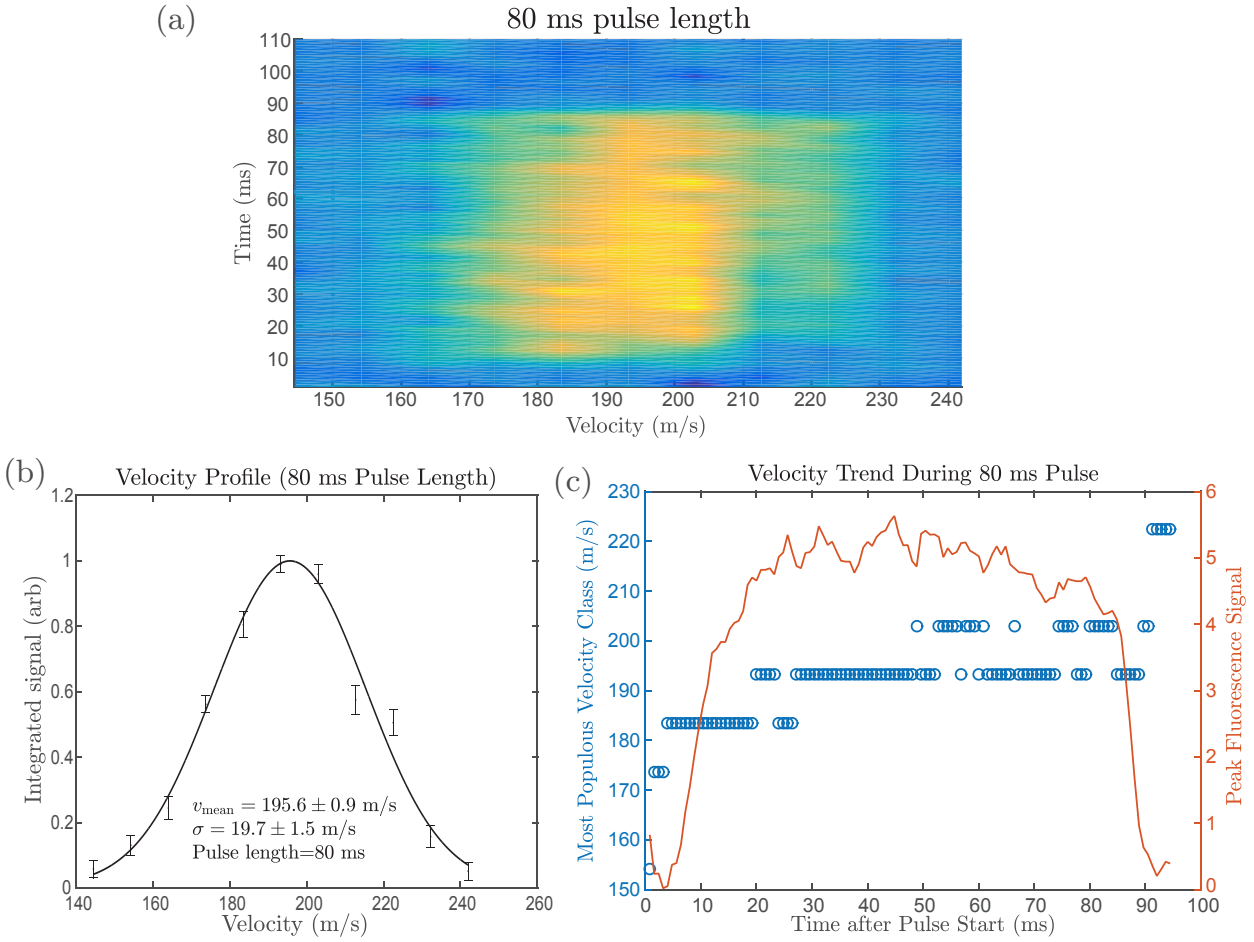

Figure 5.6.5: Average beam velocity profile for an 80 ms laser pulse as a function of time after the start of the laser pulse. Velocity is measured at a distance of 1 m from the beam source. Running parameters: 90% fiber laser current, 4.75 Hz repetition rate, upstream target (5 cm from cell exit), with 2× beam expander installed. **(a)** Velocity population surface plot v. time after the start of the laser pulse. **(b)** Velocity profile averaged over time at each velocity class and fit to a Gaussian distribution. **(c)** Peak fluorescence signal in the probe region v. time after the start of the laser pulse (orange) and most populous velocity class in each time interval (blue). For this pulse length, the peak velocity class shows an increasing slope of ≈ 250 cm/s per ms after the start of the laser pulse.

5.6.2 Rotational Temperature

We have measured the rotational distribution by tuning a downstream absorption laser to different Q-branch (i.e. $\Delta J = 0$) transitions in the $X \rightarrow C$ rotational manifold while normalizing to an absorption signal on the $X \rightarrow C$ $Q(1)$ transition just outside the cell exit. The population in the addressed rotational level is proportional to the absorption OD with the same constant of proportionality for all Q-branch lines [105].

Figure 5.6.6 shows the rotational distribution for 25 ms and 50 ms pulse widths measured just after the collimator in the 4 K shield ($\approx$ 60 mm from the cell exit). By fitting the results to a Boltzmann distribution, we find that the rotational temperatures agree within uncertainty for these two running conditions.

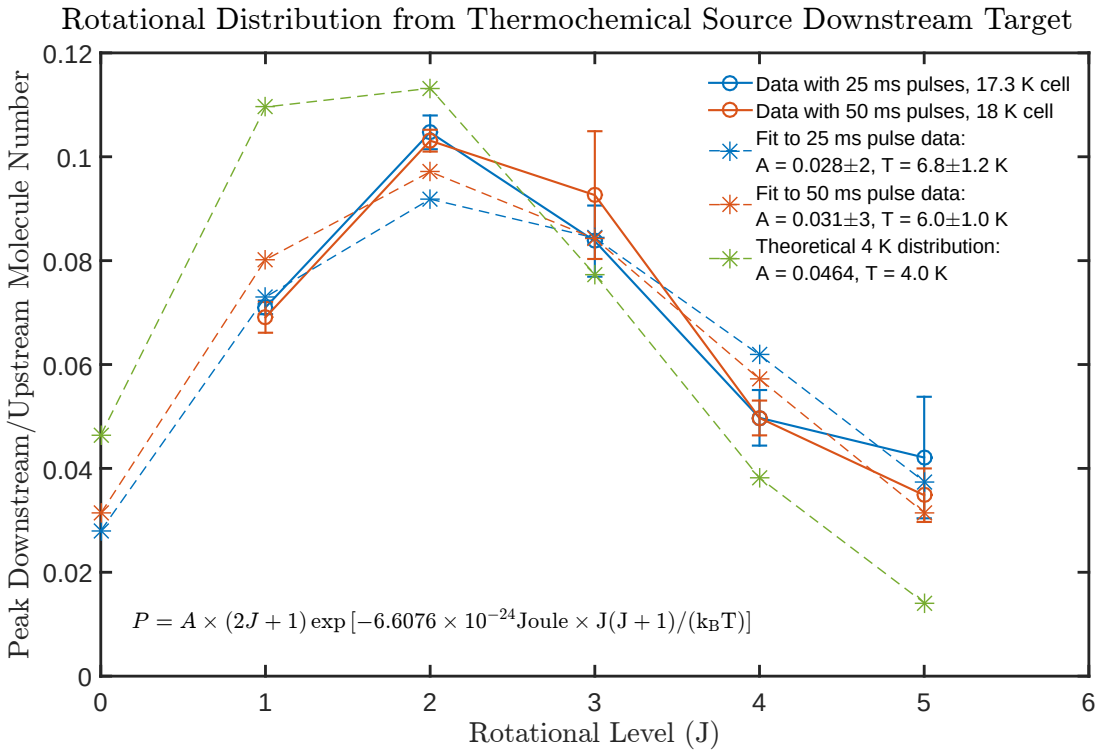

Figure 5.6.6: Rotational distributions measured outside the 4 K shield for 25 and 50 ms pulses, normalized to the absorption signal outside the cell. Firing the laser at the downstream target, 2.5 cm from the cell exit. A 4 K thermal distribution (approximately the temperature of the ablation-based beam) with the same total population is plotted for comparison.

Figure 5.6.7 shows a different analysis of the data in Fig. 5.6.6. The average pulse shapes are plotted together with the rotational temperature, derived from a Boltzmann distribution fit, as a function of time throughout each pulse. The temperature is approximately constant for the shorter, 25 ms pulse width and increases gradually throughout the pulse for the longer, 50 ms pulse width.

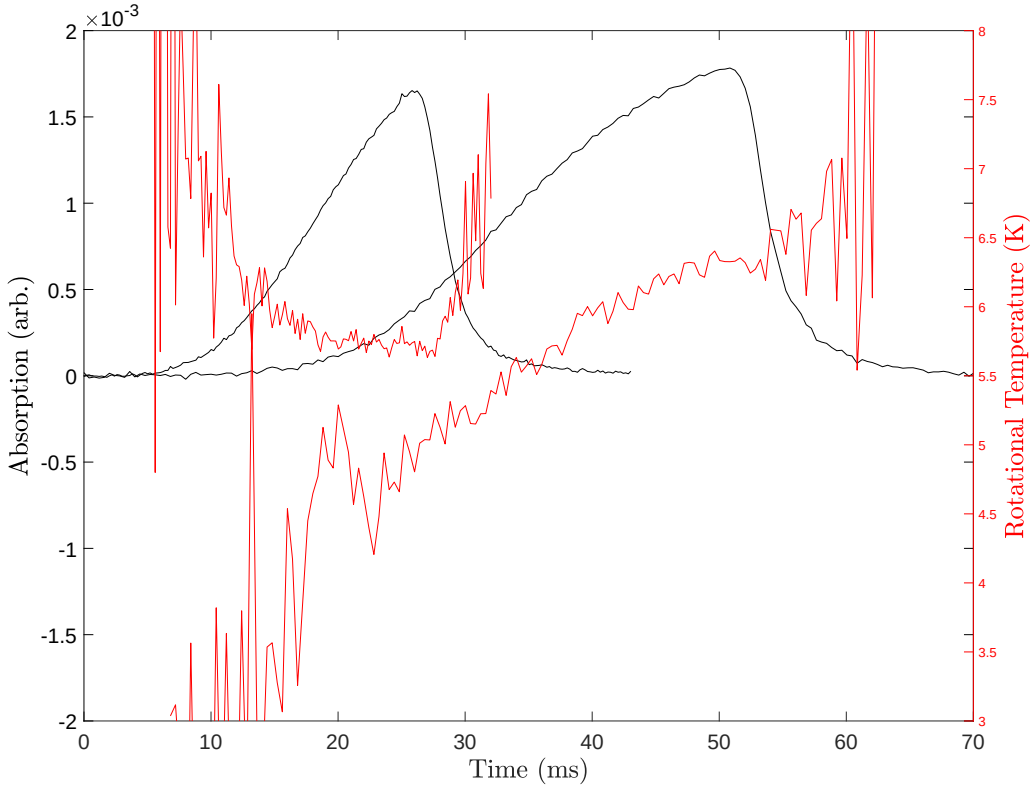

Figure 5.6.7: Average absorption pulse shapes (black) and rotational temperature fits (red) plotted as a function of time for 25 ms (left) and 50 ms (right) pulse widths. Figure and analysis from Adam West.

We also measured the rotational distribution at a distance of 11 mm from the cell during the long, stress-test run described in Section 5.7. The average result and its fit temperature are shown in Fig. 5.6.8, together with the results from Fig. 5.6.6. The rotational temperature is the same within uncertainty at 11 mm and 60 mm from the cell exit, indicating that rotational cooling collisions are frozen out by the shorter distance.

Results: Rotational thermalization occurs within 11 mm from the cell exit, and the average

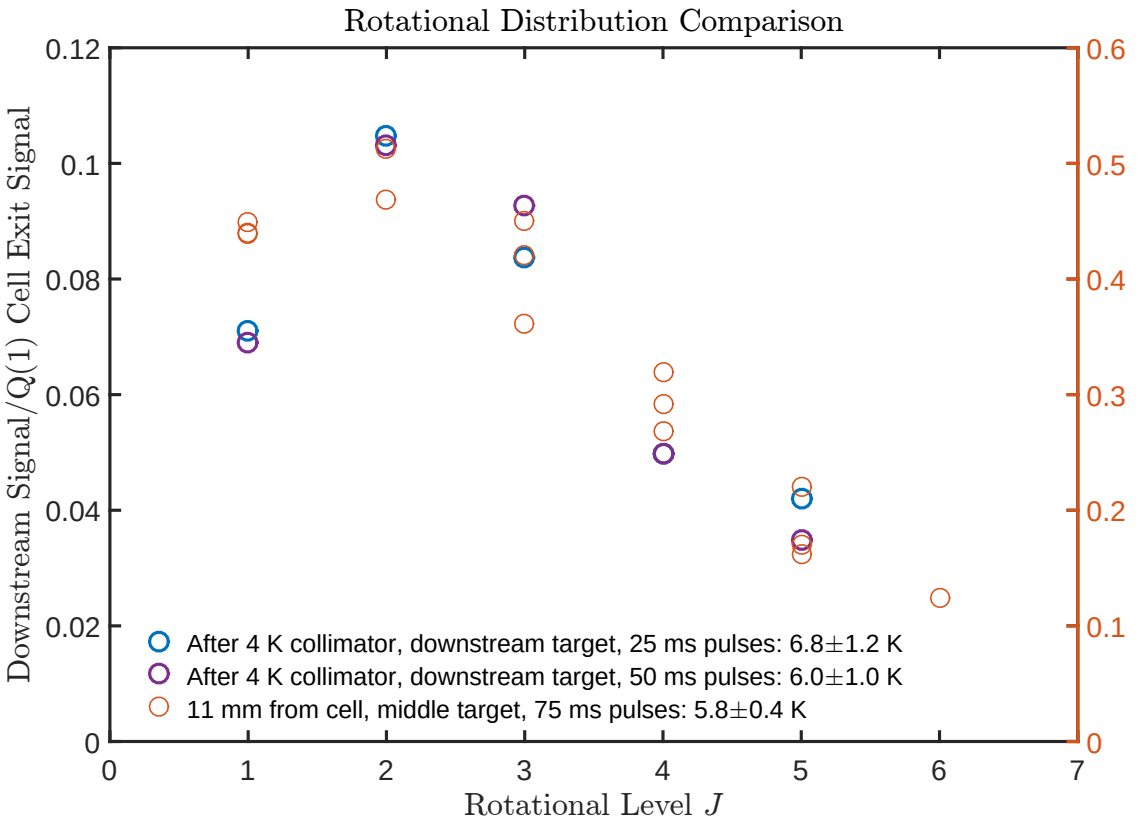

Figure 5.6.8: Rotational distribution comparison at 11 mm (orange, right axis) and 60 mm (blue and purple, left axis) from the beam source. The blue and purple data points and fits given in the legend are from Fig. 5.6.6, while the orange data points were taken during the long run described in Section 5.7.

final rotational temperature of the beam is 6.5 ± 1.5 K. During longer pulses, there is some evidence that the instantaneous beam temperature increases at a rate of about 0.1 K/ms after the start of the pulse.

Conclusion: The average rotational temperature is consistent within uncertainty under the three disparate sets of running conditions used in Fig. 5.6.8. This suggests that the rotational distribution will be independent of most reasonable choices of running parameters.

5.6.3 Divergence

The results presented in Section 5.6.2 indicate that collisions within the beam are largely frozen out by a distance of 11 mm from the cell exit. Molecule trajectories after this point can be assumed to be ballistic.

In order to measure the divergence of the beam, we performed a Doppler scan on the $X \rightarrow C$ $Q(2)$ line at 11 mm from the cell during the long run described in Section 5.7. For comparison, we also took a Doppler scan on the $X \rightarrow C$ $Q(1)$ line just outside the cell exit. While one laser was tuned, the other was held on resonance to normalize out fluctuations in the molecule beam signal. The results of these measurements are plotted with Gaussian fits in Fig. 5.6.9.

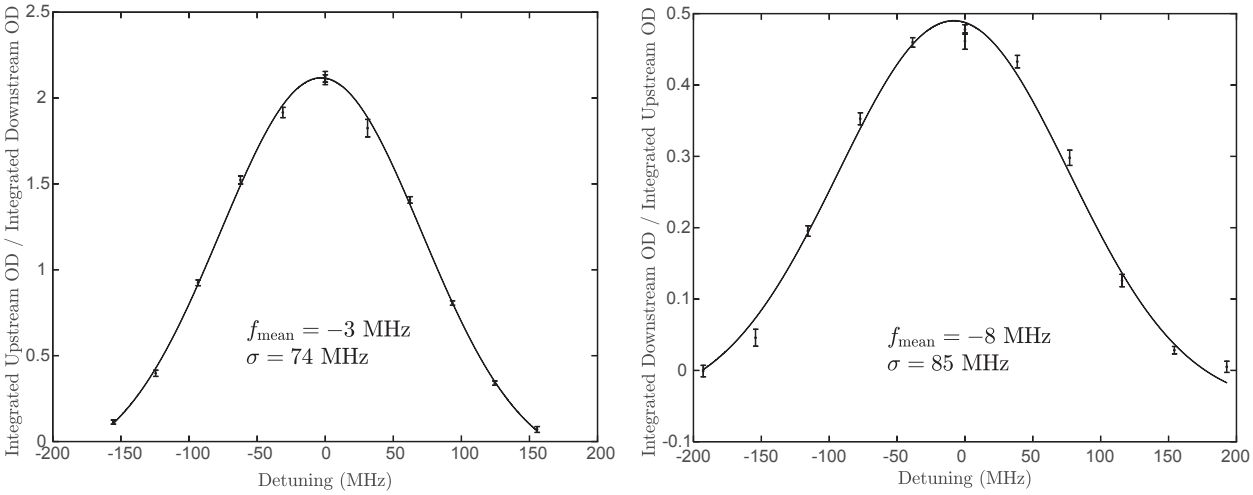

Figure 5.6.9: Transverse Doppler profiles taken with a laser wavelength of 690 nm at 0 mm (left) and 11 mm (right) from the cell exit.

Results: The signal at 11 mm has a 1σ Gaussian width of 85 MHz, which is equivalent to $2\sqrt{2\ln(2)}\sigma = 200$ MHz FWHM (full width at half maximum). With a laser wavelength of 690 nm, this corresponds to a velocity width of $690\,\text{nm} \times 200\,\text{MHz} = 138\,\text{m/s}$ FWHM. Given that the forward velocity of the beam is typically about 195 m/s, the divergence angle is $\theta_{\text{FWHM}} = 2 \times \arctan[138\text{ m/s}/(2 \times 195\text{ m/s})] = 39°$ FWHM. The FWHM divergence solid angle is therefore $\Omega_{\text{div.}} = 2\pi[1 - \cos(39°/2)] = 0.36$ sr. This is consistent within uncertainty with the result from the ablation source [106].

The cell exit Doppler width has been observed to vary between about 110 MHz and 180 MHz FWHM from day to day and under various running conditions. This suggests that the divergence may also vary by a comparable fraction. Note from Fig. 5.6.9 that the Doppler width at 11 mm is broader than that just outside the cell. Evidently, the molecules undergo collisions in this region that increase the transverse velocity width.

5.7 Long-Term Beam Flux

Our final experiment was a long "stress test" of a single, previously unused thermochemical source target. The running conditions were designed to simulate those of the eEDM experiment as closely as possible. We ran continuously for several hours each day over the course of about 2 weeks for a total run time of 70 hours. After approximately every 10 hours of cumulative run time, we performed a de-ice of the beam source (see Section 2.2). During the run, we monitored the resonant absorption signal on the $X \to C$ $Q(1)$ line (i.e. $J = 1 \to J' = 1$) in order to measure the total beam flux. The results, shown in Fig. 5.7.1 and 5.7.2, reveal a dramatic decrease in the beam flux as the "fresh" (un-fired-at) spots on the target surface become rare.

For this run we used the 19 mm diameter target 50 mm from the cell exit. The running conditions were chosen based on the optimization efforts described in previous sections. The selected settings were as follow ("*" indicates that the setting was changed at some point in

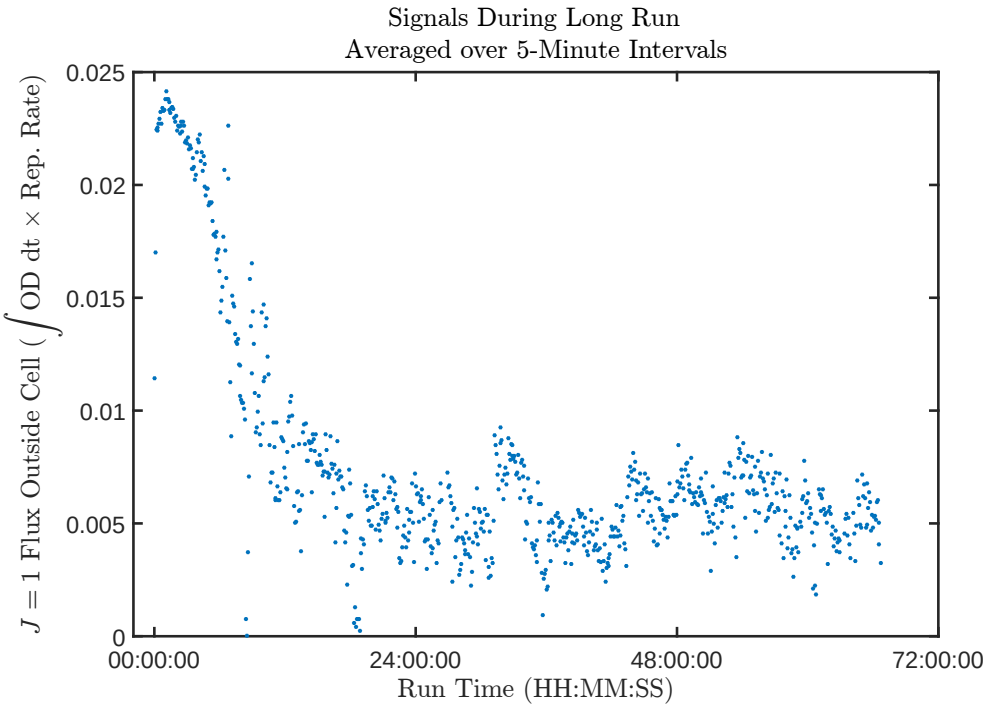

Figure 5.7.1: Measured flux trend (in integrated optical depth on the $X \to C$ $Q(1)$ line $\times$ repetition rate) during the long run. This plot shows the flux in terms of directly measured quantities, which are not subject to revision due to updated OD to number flux conversion factors. For comparison, in June 2012, Nick Hutzler measured the ablation source integrated OD outside the cell to be $6 \times 10^{-5}\,\mathrm{s/pulse} \times 50\,\mathrm{Hz} = 3 \times 10^{-3}$. The molecule flux gain is given by the ratio of these measurements times the product of the forward velocity ratio and the transverse Doppler width ratio, or, for the first 10 hours shown on this plot: $(2 \times 10^{-2}/3 \times 10^{-3}) \times (195\,\mathrm{m/s} \,/\, 180\,\mathrm{m/s}) \times (140\,\mathrm{MHz}/110\,\mathrm{MHz}) = 9$.

the run, as described below):

1. 100% Fiber laser current
2. 800 Hz rectangle wave fiber laser emission modulation
3. *50% fiber laser modulation duty cycle
4. 75 ms fiber laser pulse envelope (60 modulation cycles)
5. *7.8 Hz fiber laser pulse repetition rate
6. 16.5 K cell temperature setpoint; heat load adjusted to keep the actual cell temperature between 16.5 and 18 K
7. *30 SCCM neon flow
8. *400 mm focal length fiber laser lens at $\approx$ 405 mm from the target surface

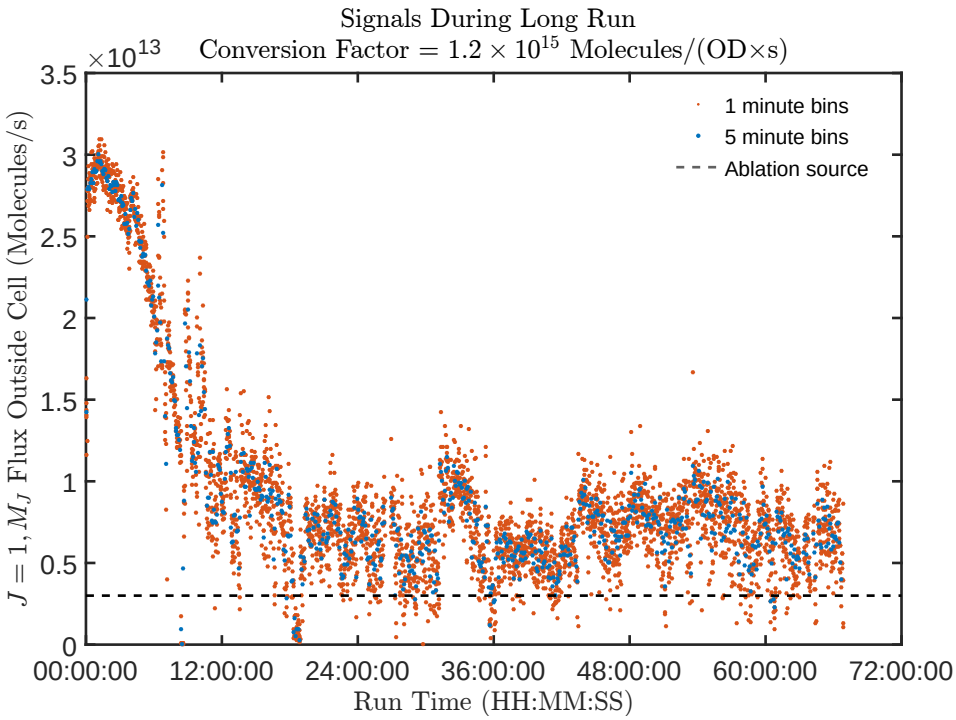

Figure 5.7.2: Molecule flux per Zeeman sublevel in $|X, J = 1\rangle$ during the long run, calculated from the result in Fig. 5.7.2. The estimated average ablation flux is represented by the dashed line.

9. *Galvo randomly spot hopping (as described in Section 5.1 in item 2) with a frequency of 5 Hz about a square region inscribed in the circular target face.

After the signals began to drop at ≈ 10 h of cumulative run time, various changes were made in an effort to re-optimize the flux. The timing and details of these changes are described below:

1. At 13.25 hours into the run, we began using a threshold with the galvo spot hopping (as described in Section 5.1 in item 2 and 4). This allowed us to dwell on the remaining good target spots and avoid wasting time on the depleted spots. The spot hopping threshold was varied to optimize the signals, but its typical level was $\approx 6 \times 10^{12}$ molecules/s in a single M_J sublevel of $|X, J = 1\rangle$.

2. At 18.75 hours into the run, the 400 mm focal length lens was replaced by a 300 mm focal length lens at ≈ 325 mm from the target. A 2× beam expander was also installed at the output of the fiber laser, and its focal length was adjusted to maximize the signals.

3. At 19 hours into the run, the neon flow rate was increased to 40 SCCM.

4. After 19 hours, the fiber laser pulse repetition rate was adjusted as needed between 5.5 and 7.8 Hz to keep the cell temperature below 18 K.

5. After 19 hours, the pointing of the fiber laser into the cell was intermittently adjusted to reach undepleted target regions. The focal length of the beam expander was also adjusted as needed.

6. At 20.5 hours into the run, the fiber laser modulation duty cycle was increased to 90%.

7. At 31 hours into the run, a small-amplitude, 1 Hz "micro-raster" was applied to the galvo to move the laser spot around a small target region between spot hops (as described in Section 5.1 in item 3 and 4). This allowed for more efficient exploration and usage of the remaining "good" regions on the mostly depleted target.

Of the adjustments described above, only numbers 1, 5, and 7 seemed to produce any improvement, as shown in Fig. 5.7.1 and 5.7.2: The slope of the signal fall-off decreases around 13 hours, when the galvo thresholding is turned on, and the signals become more stable on short timescales around 31 hours, when the micro-raster is turned on. The small signal jumps that occur later in the run, e.g. around 32 and 45 hours, are the result of adjusting the beam pointing to find still-undepleted target spots. The run was stopped when we could no longer reliably find spots that produced yields above the $\approx 6 \times 10^{12}$ molecules/s (in $|X, J = 1, M_J\rangle$) threshold.

In the shot noise limit, the eEDM statistical sensitivity is determined (in part) by the average count rate, which is proportional to the average molecule flux. The time-averaged molecule flux in a single M_J sublevel of the $J = 1$ state in X is plotted as a function of time in Fig. 5.7.3. In Fig. 5.7.4, this quantity is divided by the single-quantum-state ablation source flux 3×10^{12}/s, taken to be at the high end of the range of values in Table 2.6.1 (and assuming a beam divergence of 0.35 sr [106]), to give the thermochemical source flux gain as a function of time.

In Fig. 5.7.5, the data from Fig. 5.7.2 are used to produce a histogram of fluxes in 1-minute time bins. The distribution is distinctly bimodal, with the early high fluxes from spot hopping over a fresh target face forming the smaller peak, and the later, diminished

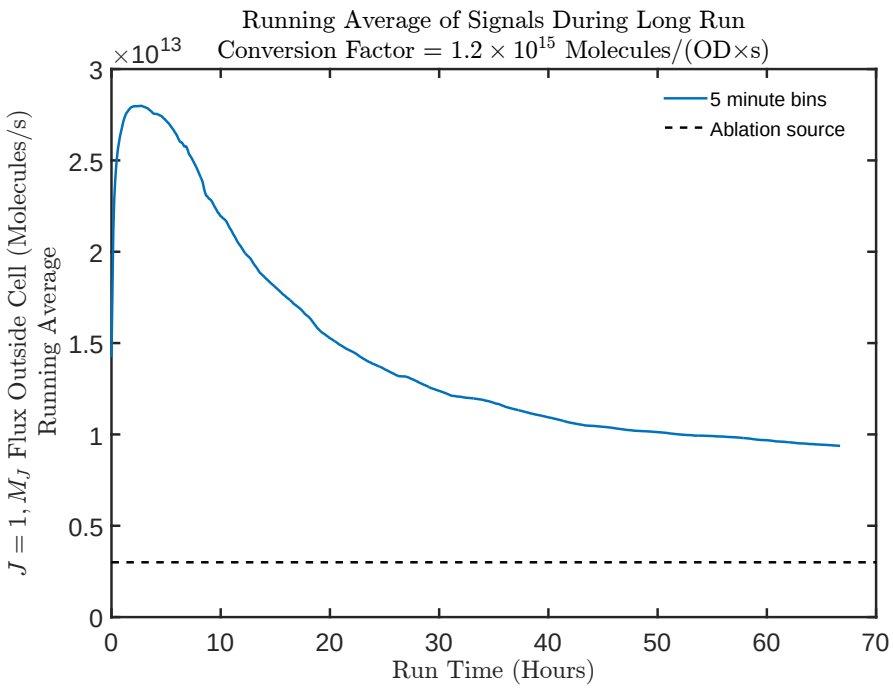

Figure 5.7.3: Running average of the molecule flux per Zeeman sublevel in $|X, J = 1\rangle$ during the long run. The estimated average ablation flux is represented by the dashed line.

fluxes produced by flux-thresholded spot hopping over a partially depleted target forming the larger peak. In Fig. 5.7.6 we plot the cumulative sum of the gain distribution from high gain to low gain to obtain the total amount of run time at or above each gain level. Thus, we see that the flux gain is above 8 for a total of 5 hours, above 4 for 10 hours, and above 2 for 50 hours.

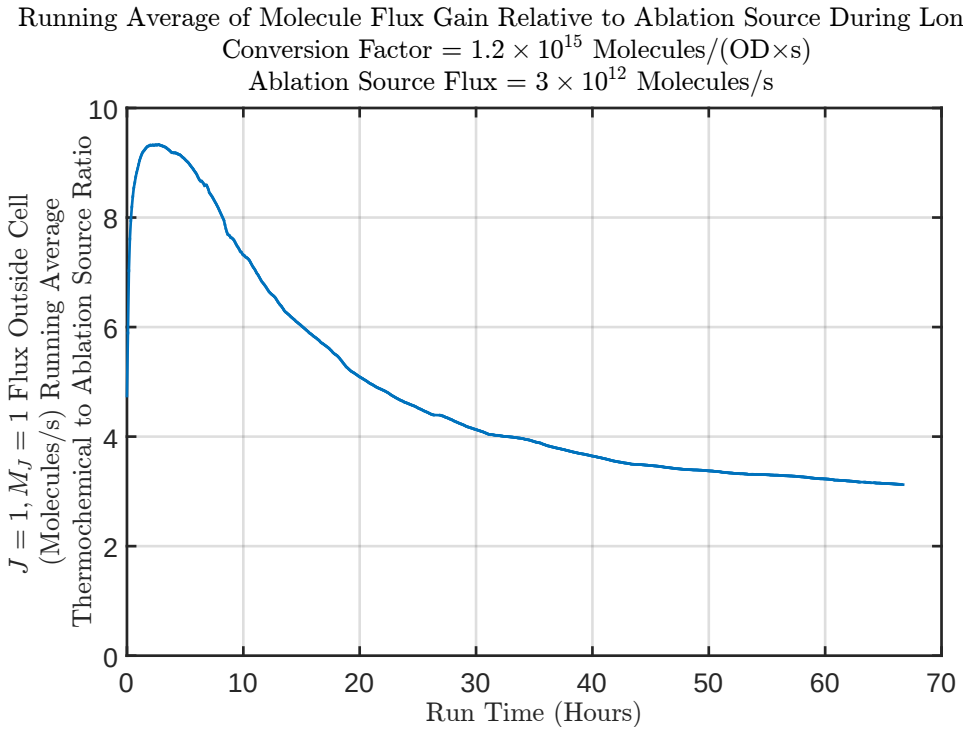

Figure 5.7.4: Running average of the molecule flux gain during the long run, calculated by dividing the result in Fig. 5.7.3 by the average ablation flux per Zeeman sublevel in $|X, J = 1\rangle$.

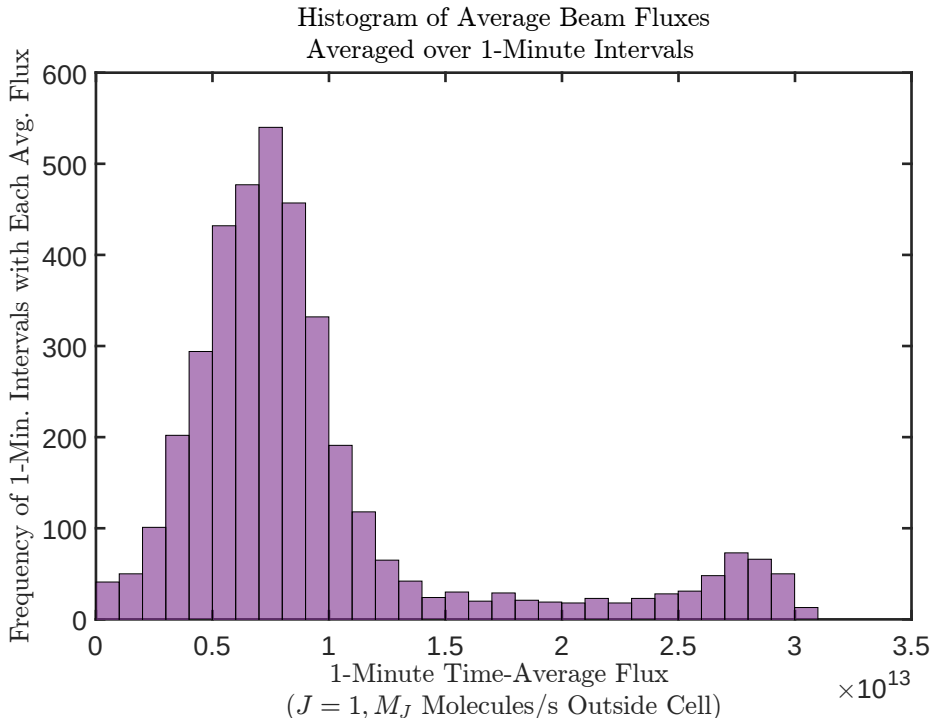

Figure 5.7.5: Histogram of molecule fluxes per Zeeman sublevel in $|X, J = 1\rangle$ during the long run, calculated from the 1-minute binned data in Fig. 5.7.2.

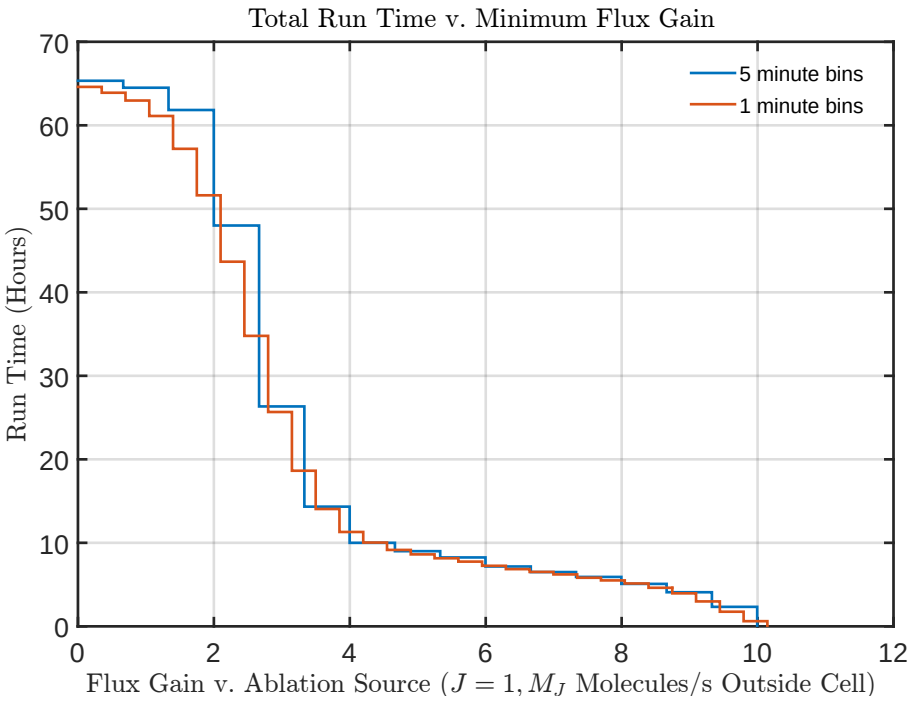

Figure 5.7.6: Cumulative distribution of molecule flux gains during the long run, calculated by summing the gain distributions in 5 minute and 1 minute bins from $+\infty$.

The labours of men of genius, however erroneously directed, scarcely ever fail in ultimately turning to the solid advantage of mankind.

—Mary Shelley, *Frankenstein*

While none of the work we do is very important, it is important that we do a great deal of it

—Joseph Heller, *Catch 22*

6 Conclusion

The aim of this work was to develop a new, high-flux thermochemical ThO CBGB source for the purpose of improving the statistical sensitivity of the ACME electron EDM experiment. In Chapter 5, we describe such a source and demonstrate flux gains of 7× (3×) in 10 hours (50 hours) of run time on a single target. Other properties of this beam source have also been measured and are found to differ from the ablation source properties within an acceptable range. The expected statistical sensitivity gains from this new source are summarized in Section 6.1.

For all its advantages, the proposed implementation of the thermochemical source in ACME's measurement scheme presents significant practical challenges. In particular, the short target lifetime compared to that of the ablation source necessitates warming up the source and replacing the targets on approximately a weekly basis during heavy running. While this difficulty is not unexpected—the relative mechanical weakness and higher fluxes of the thermochemical targets must inevitably reduce their longevity—it does present a real

hurdle in the context of the ACME running routine. Six months of hard running could require up to 25 target changes and 100 targets. Several possible approaches to avoiding these difficulties, e.g. by outsourcing or streamlining the target making process, developing a system for replacing targets in situ without warming up or opening the source, and extending the target longevity, are discussed in Section 6.2. Thermochemical source maintenance will need to be made less time-consuming and labor-intensive before the source is fully usable in the ACME experiment.

6.1 eEDM Sensitivity Gain from Thermochemical Source

From the shot-noise-limited eEDM uncertainty expression in Eq. (2.13), we can write the beam source figure of merit for eEDM sensitivity in the following convenient form:

$$\frac{1}{\delta d_e} \propto \tau \times \sqrt{\dot{N}} \times \mathcal{C} \tag{6.1}$$

$$\propto \frac{1}{v} \times \sqrt{\frac{\dot{N}_{\text{source}}}{\Omega_{\text{div.}}}} \times \exp\left[-\frac{\pi}{2}\left(\frac{\Delta v}{v}\right)^2\right]. \tag{6.2}$$

The first factor in Eq. 6.2 expresses that the interaction time is proportional to the forward velocity v of the molecule beam, while the second factor says that the experimental count rate is proportional to the beam intensity, i.e., the beam flux $\dot{N}_{\text{source}}$ in the relevant quantum states divided by the beam divergence $\Omega_{\text{div.}}$. The third term gives the reduction in fringe contrast caused by the fractional forward velocity spread $\Delta v/v$. This relationship is derived in reference [188].

In Table 6.1.1, we use Eq. (6.2) together with the measured thermochemical source properties described in Chapter 5 to compute the shot-noise-limited sensitivity gain from replacing the ablation source with the thermochemical source. The time-averaged cell exit flux after 10 and 50 hours of running on a single target is read from Fig. 5.7.3. The ablation source cell exit flux is assumed to be at the high end of the range quoted in reference [35] and Table 2.6.1. The forward velocity spread that enters the contrast term is the *instantaneous*

forward velocity spread in the interaction region, which was measured by Brendon to be $\Delta v \approx 31$ m/s for the ablation source. Because of velocity dispersion, this is smaller than the total forward velocity width of the beam [106]. The other ablation source numbers are taken from references [105, 106].

The rotational temperature in Table 6.1.1 is needed to determine the available source flux $\dot{N}_{\text{source}}$: Because we use rotational cooling to concentrate population into the desired rotational sublevel in X, a higher-temperature beam will have a higher fraction of its usable population in excited rotational states and a smaller fraction in the $J = 1$ level, whose population we measure. To correct for this population shift, Zack Lasner has calculated that with the rotational cooling scheme used in the Gen. II ACME experiment, we can get $\sim 10\%$ more population from non-($J = 1$) rotational levels in a 6 K beam than in a 4 K beam. By adding an additional rotational cooling laser (addressing $J = 4$), I estimate that we can add another 10% to this number. Altogether, a 6 K beam has about 50% more population than a 4 K beam in levels that are not addressed by our current rotational cooling scheme. For the purposes of the sensitivity gain estimate in Table 6.1.1, I assume a signal gain of 20% relative to the ablation source, implying a statistical sensitivity gain of $\sqrt{1.2} = 1.1$.

Table 6.1.1: Comparison between thermochemical and ablation source properties. The sensitivity gain is computed via Eq. (6.2) under the assumption that our statistics are shot-noise limited.

	Ablation	Thermochemical (50 h/10 h)	Sensitivity Gain (50 h/10 h)
$\lvert X, J = 1, M_J\rangle$ cell exit flux (mol./s)	3×10^{12}	$1.0 \times 10^{13}/2.2 \times 10^{13}$	1.8/2.7
Forward velocity (m/s)	180	195	0.92
Fractional forward velocity width (FWHM)	0.17	0.23	0.96
Divergence (sr, FWHM)	0.35	0.36	0.97
Rotational temperature (K)	4	6	1.1
Total (50 h/10 h)			1.7/2.5

6.2 Future Directions

Here are some of the thermochemical-source-related upgrades and investigations I would suggest to future beam box jockeys. They are arranged roughly in order of practicality, starting with the "necessary," passing through the "recommended" and finishing up with the "wild-eyed":

1. *Compare thermochemical source yields to ablation source yields* measured under ordinary Gen. II running conditions. Monitor the absorption signal on the $X \rightarrow C$ $Q(1)$ line outside the ablation cell during the Gen. II run to obtain a more accurate measurement of the thermochemical source gain. Also compare with signal levels on a fresh ablation target to measure the gain relative to the case where we change ablation targets as frequently as we will have to change thermochemical targets. Finally, run for a full day on an ablation target with a YAG repetition rate of 100 Hz to compare the signal gains under those conditions. Keep in mind that the thermochemical beam source can also be run as an ablation source with a higher heat load limit.

2. *Test the large (1" bore) cell* with four targets. If it works well, this may increase the time between thermochemical source target changes by a factor of 2.

3. *Add an ion sweeper* to the Beam Box II exit plate. Copy the design of the Gen. I source ion sweeper.

4. *Improve the thermal conductivity between the PT810 and cell.* Figure 5.1.1 suggests that it may be possible to gain ≈ 3 W more cooling power at 16 K with improved heat links, permitting a $\approx 25\%$ higher duty cycle.

5. *Reinstall additional two chevron baffles* on the 70 K shield front plate to help with pumping speed in the beamline and room temperature beam box region.

6. *Run with a higher neon flow rate* to clarify the correlation between signal and neon flow. We have so far run with flow rates up to about 50 SCCM, and under some

running conditions, the signal seemed to increase monotonically with flow up to that value. In addition, changing the flow rate dramatically could help elucidate the cause of the "after-pulsing" effect observed in some of the pulse shapes and could improve the extraction efficiency for a third upstream target, increasing the usable target-holding cell length.

7. *Purchase a higher-power IPG fiber laser.* The data in Section 5.4 shows that the instantaneous signal still has an increasing trend at 50 W, the maximum power of our current laser. The analysis in Section 4.2 suggests that the marginal gains should start to diminish around 60 W. It would be interesting to verify this, and I would suggest purchasing a 100–150 W laser to do so. This might also come in handy if we eventually thermally decouple the target from the cell: In this case, the maximum heat load on the target will be much higher.

8. *Survey the test beamline* to determine the dust distribution. This will help us assess how close the field plates can be placed to the beam source and whether dust escaping the beam box will pose a problem in the future.

9. *Build a copy of the heavy-duty translation stage* on the ablation beam source and put the thermochemical beam box on it so that it can be swapped into the experiment more easily.

10. *Purchase a cold isostatic press* for target making. This could help improve the target turnaround time and simplify what is currently the most annoying target making step: extracting the pellet from the die without breaking it or contaminating the press.

11. *Investigate the possibility of outsourcing the target fabrication.* Depending on the frequency with which we decide to change thermochemical targets, making them could become a serious burden. See if anyone (e.g. IBI Labs) has the facilities to make them for us in bulk.

12. *Design a load-lock mechanism for changing the targets* in situ without warming up or opening the source. This would make the target changing procedure much less painful and perhaps enable us to change to fresh, more-productive targets on a daily or weekly instead of a monthly basis.

13. *Test possible nozzle configurations on the cell exit.* Some of the pulse shape data shows long "after-pulsing" tails on the signal, which suggests that in-cell gas dynamics may be complex. We may be able to suppress possible detrimental effects related to eddies inside the cell that entrain ThO by smoothing the sharp edges and corners in the cell interior and adding a flow-shaping nozzle to the cell exit aperture.

14. *Vary the target composition.* Apart from one or two targets with a non-standard thorium-thoria ratio, which had a limited but perhaps non-negligible effect on the yield, we have not carefully investigated the correlation between target composition and yield. It would be interesting to learn more about this dependence.

15. *Investigate better target-making procedures.* Reference [23] shows that by melting and re-solidifying a mixture of thorium and thoria, it is possible to produce theoretically dense, pore-free solids with interleaved ~ 10 μm microstructures of the two materials. Since the reaction rate in a solid depends on the contact surface area between the reactants, this would appear to be a potentially promising route to improving our ThO yields, as well as possibly making our targets stronger and less dusty. Sintering thorium and thoria together did not produce these results because the reaction to ThO was driven too rapidly at high temperatures, thereby depleting the target before we could use it. However, if the materials were completely melted—perhaps in an induction furnace—and then rapidly quenched, the reaction rate would likely be irrelevant because the liquid thorium and thoria would continuously flow to maintain contact.

16. *Investigate in-situ-replaceable targets.* If the target could be mounted on a movable contraption outside the cell without detriment to the thermals or in-cell gas dynamics,

one could imagine designing a panel or a rolling wheel of targets (or even a single, large target) that could be swapped out as the signals deplete.

17. *Investigate thermally decoupling the target from the cell.* Many of the thermal constraints discussed in Section 4.2 would be substantially loosened if the target were thermally coupled to, say, the 60 K shield instead of the cell. In this case, we might be able to increase our yields by using much higher heating powers. By reducing the need for a thermally insulating target, this might also allow us to make our targets out of a better material e.g. oxidized thorium, which might be more mechanically robust and have greater longevity than the pressed powder targets. Such a target might also be made "refreshable"....

18. *Investigate "refreshable" target surfaces.* In reference [56], Darnell et al. take a pure sample of thorium, deliberately oxidize it, and then observe ThO(g) production at high temperatures. It is possible that we could use such a scheme to reduce our target changing frequency: If an oxidized sample of thorium could be re-oxidized in situ—or in an oven region to which we could periodically withdraw it without opening the chamber—then we might be able to produce a refreshable target surface with significantly less effort than that involved in a target change.

6.3 Afterword

> No human face is exactly the same in its lines on each side, no leaf perfect in its lobes, no branch in its symmetry. All admit irregularity as they imply change; and to banish imperfection is to destroy expression, to check exertion, to paralyze vitality. All things are literally better, lovelier, and more beloved for the imperfections which have been divinely appointed, that the law of human life may be Effort, and the law of human judgment, Mercy.
>
> —John Ruskin, *The Stones of Venice*

The decades-long search for a flaw in the ineffably symmetrical electron has required by turns ingenious innovations, jury-rigged hacks, cautious reticence, bold strokes of action,

and—almost always—sheer blind luck. Like the experiments in whose steps it follows [4, 48, 57, 67, 103, 113, 137, 162, 171] and those that travel with it [88, 126, 140, 185], the ACME experiment has had its share of all these moments on the road from its inception, through its record-setting first-generation result, to its new incarnation as a fledgling second-generation experiment. Looking back upon the impressive progress of modern precision measurements, it is humbling to realize how much we can still learn from the generation that began the quest for the eEDM nearly seventy years ago. If ACME's new thermochemical thorium monoxide source comes to fruition, it will mark another delightful twist in the history of eEDM searches: A systematic error in a mid-20^{th} century nuclear materials experiment caused by the nuisance production of ThO will have inspired an improved molecule source for a 21^{st} century precision measurement. I'd like to think that Norman Ramsey would be pleased.

A

Dipole Interaction of the Rigid Rotor

In this appendix, we derive the permanent and induced dipole moments and the perturbative Stark shift of a rigid rotor molecule with a fixed (molecule-frame) electric dipole moment $\vec{d}_{\text{mol}}$ and moment of inertia I.[1]

A.1 Rigid Rotor

First, let's review the solution to the rigid rotor problem (For more details see, e.g. Brown and Carrington chapter 6.8.1 [33] and Townes and Schawlow chapter 1.1 [187]). Consider the rigid rotor illustrated in Fig. A.1.1, which consists of a pair of masses M_1 and M_2 separated by a fixed distance R. In the center-of-mass frame, this system reduces to a single mass $\mu = \frac{M_1 M_2}{M_1+M_2}$ constrained to the surface of a sphere of radius R. The energy is just the kinetic energy of a system with angular degrees of freedom (θ, ϕ), which is given by

$$H_{\text{rot}} = \vec{J}^2/2I, \tag{A.1}$$

[1] This writeup was originally prepared for a course on molecules taught by Prof. David DeMille in the spring of 2014. Thanks to him and to my classmates for useful insights and feedback.

where $\vec{J}$ is the angular momentum operator, and $I = \mu R^2$ is the moment of inertia. The solutions to the time-independent Schrödinger equation $H_{\text{rot}}\psi(\theta,\phi) = E\psi(\theta,\phi)$ are the spherical harmonics

$$\psi(\theta,\phi) = Y_J^m(\theta,\phi), \tag{A.2}$$

with eigenenergies

$$E_{J,m} = \frac{\hbar^2}{2I} J(J+1). \tag{A.3}$$

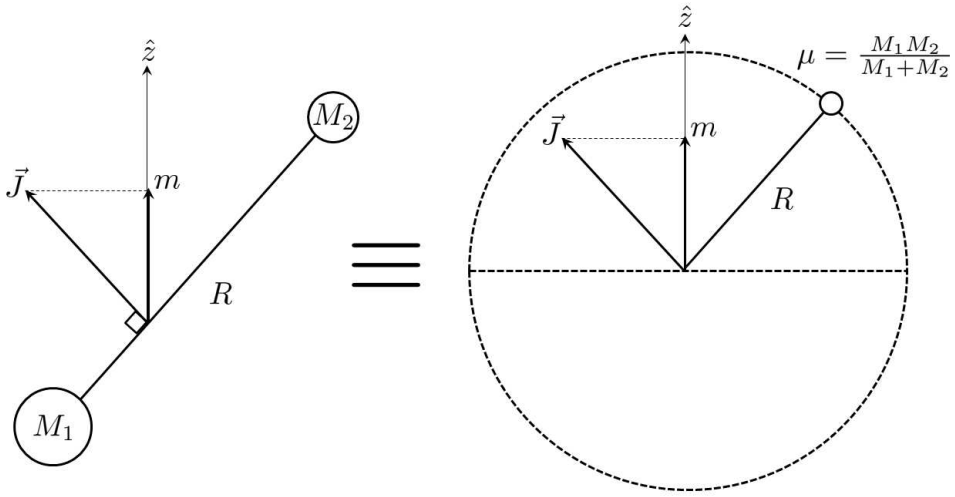

Figure A.1.1: Rigid rotor. In its center-of-mass frame, the dumbbell model on the left is mathematically equivalent to a single particle of reduced mas μ constrained to the surface of a sphere of radius R, as depicted on the right.

In the frame of reference that rotates with the molecule, the rigid rotor dipole moment $\vec{d}_{\text{mol}}$ is equal to $\Delta q \vec{R}$, where Δq is the absolute value of the charge excess per atom and depends on the internal structure of the atoms and on the properties of the molecular bond.

A.2 Permanent EDM

The dipole operator changes the parity of the state it acts upon, so it cannot couple two parity eigenstates with the same parity. Specifically, it cannot connect $|J,m\rangle$ to itself. Therefore,

the expectation value of the electric dipole moment is zero in the lab frame, even though the molecule may have a nonzero dipole moment $\vec{d}_{\text{mol}}$ in the frame of reference that rotates with the molecule. Note that this also means that the linear Stark shift $E_{\text{St}}^{(1)} = \langle \vec{d} \rangle \cdot \vec{\mathcal{E}}$ vanishes.

To see explicitly that a permanent zero-field electric dipole moment (EDM) is forbidden by parity, we use the transformation properties of $|J, m\rangle$ and $\vec{d}$ under the parity operator P:

$$P\vec{d}P^\dagger = qP\vec{r}P^\dagger = -q\vec{r} = -\vec{d}, \text{ and} \tag{A.4}$$

$$P\,|J, m\rangle = (-1)^J\,|J, m\rangle\,. \tag{A.5}$$

In Eq. (A.4), we have used the fact that the dipole operator is equal to the charge q times the displacement operator $\vec{r}$, and the displacement transforms into its opposite under parity. Equation (A.5) describes the parity properties of the spherical harmonics (see, e.g. Merzbacher chapter 11.4 [133]).

Now, with malice aforethought, we calculate the negative expectation value of the dipole operator:

$$-\langle J, m|\,\vec{d}\,|J, m\rangle = \langle J, m|\,P^\dagger\vec{d}P\,|J, m\rangle \tag{A.6}$$

$$= \langle J, m|\,(-1)^J\vec{d}(-1)^J\,|J, m\rangle = [(-1)^2]^J\,\langle J, m|\,\vec{d}\,|J, m\rangle \tag{A.7}$$

$$= +\langle J, m|\,\vec{d}\,|J, m\rangle \tag{A.8}$$

$$= 0. \tag{A.9}$$

In Eq. (A.6), we have used Eq. (A.4) and the hermicity of the parity operator, and in Eq. (A.7) we have substituted Eq. (A.5) for the parity operator acting on the spherical harmonics.

This proof can also be performed in position space using integrals over the spherical harmonics. See Budker, Kimball, and DeMille chapter 7.6 [34].

Note that this proof works if you substitute any eigenstate of parity for $|J, m\rangle$. In the absence of applied fields that fix a preferred direction, the Hamiltonian of a system that respects parity commutes with P, so the energy eigenstates can be written as eigenstates of parity. Thus, as long as the Hamiltonian has no degenerate eigenstates of opposite parity, *there can be no permanent EDMs without parity violation.* This is a slightly more fleshed-out version of the proof given in Section 1.2.1.

A.3 Dipole Matrix Elements

We can express the dipole matrix elements in terms of integrals over products of spherical harmonics:

$$\langle J', m'| \vec{d} \cdot \hat{z} |J, m\rangle = \langle J', m'| d_{\text{mol}} \cos\theta |J, m\rangle \tag{A.10}$$

$$= d_{\text{mol}} \int d\Omega \, [Y_{J'}^{m'}(\theta, \phi)]^* \cos\theta \, Y_J^m(\theta, \phi) \tag{A.11}$$

$$= d_{\text{mol}} \sqrt{\frac{4\pi}{3}} \int d\Omega \, [Y_{J'}^{m'}(\theta, \phi)]^* \, Y_1^0(\theta, \phi) \, Y_J^m(\theta, \phi) \tag{A.12}$$

where d_{mol} is the molecule-frame EDM, θ is the angle between the z-axis and the dipole moment, and ϕ is the azimuthal angle about the z-axis. In Eq. (A.12), we have used the formula $Y_1^0(\theta, \phi) = \sqrt{3/4\pi} \cos\theta$ to write the angular dependence of the dipole operator as a spherical harmonic.

To solve this integral, we introduce a useful identity (from Merzbacher chapter 17.6 [133]) that turns integrals over products of three spherical harmonics into Clebsch-Gordan coeffi-

cients[2]:

$$\int d\Omega \, [Y_{J_3}^{m_3}(\theta,\phi)]^* \, Y_{J_1}^{m_1}(\theta,\phi) \, Y_{J_2}^{m_2}(\theta,\phi) \\ = \sqrt{\frac{(2J_1+1)(2J_2+1)}{4\pi(2J_3+3)}} \times \langle J_2 0; J_1 0 | J_3 0 \rangle \, \langle J_2 m_2; J_1 m_1 | J_3 m_3 \rangle \, . \tag{A.13}$$

Substituting Eq. (A.12) into Eq. (A.13), we obtain:

$$\langle J', m' | \, \vec{d} \cdot \hat{z} \, | J, m \rangle = d_{\text{mol}} \left(\frac{2J+1}{2J'+1} \right)^{\frac{1}{2}} \langle J0; 10 | J'0 \rangle \, \langle Jm; 10 | J'm' \rangle \, . \tag{A.14}$$

Now we can use the angular momentum conservation properties of the Clebsch-Gordan coefficients to constrain the possible values of J' and m'. Note that the final Clebsch-Gordan coefficient in Eq. (A.13) vanishes unless the z-components of the angular momenta satisfy $m_1 + m_2 = m_3$ and the total angular momenta satisfy the triangle condition $|J_3 - J_2| \leq J_1$. Since $m_1 = 0$ and $J_1 = 1$, we obtain the usual dipole selection rules $m' = m$, and $J' = J \pm 1$. (For the total angular momentum selection rule, the triangle condition tells us that J and J' differ by *at most* 1, but from Section A.2, we know that the matrix element vanishes when they differ by 0; therefore, J and J' must differ by *exactly* 1.)

Thus we can write the nonvanishing matrix elements as:

$$\langle J' = J \pm 1, m' = m | \vec{d} \cdot \hat{z} \, | J, m \rangle \\ = d_{\text{mol}} \left(\frac{2J+1}{2(J \pm 1)+1} \right)^{\frac{1}{2}} \langle J0; 10 | (J \pm 1)0 \rangle \, \langle Jm; 10 | (J \pm 1)m \rangle \, . \tag{A.15}$$

[2]David DeMille pointed out when he read this that instead of looking up the slightly obscure identity in Eq. (A.13), it is possible to derive this result using the Wigner-Eckart theorem. I will not perform this calculation here, but essentially, one would proceed by writing down the Wigner-Eckart theorem for the matrix element between $|J', m'\rangle$ and $|J, m\rangle$ (noting that $\cos\theta$ is proportional to T_1^0), and then eliminate the reduced matrix element by solving for it in terms of the dipole matrix element between $|J', 0\rangle$ and $|J, 0\rangle$, which is an integral over spherical harmonics that can be performed in Mathematica. In the end, all unknown constants of proportionality cancel, and one is left with some Clebsch-Gordan coefficients to calculate, as above.

Next, we can use Mathematica (or the recursion relations, if you're bolder than I) to calculate the Clebsch-Gordan coefficients for the two cases $J' = J+1$ and $J' = J-1$. After just a line or so of algebra, we obtain the solution:

$$\langle J', m'| \vec{d} \cdot \hat{z} |J, m\rangle = d_{\text{mol}} \times \begin{cases} \left[\frac{(J-m+1)(J+m+1)}{(2J+3)(2J+1)}\right]^{\frac{1}{2}} & \text{if } J' = J+1 \text{ and } m' = m. \\ \left[\frac{(J-m)(J+m)}{(2J-1)(2J+1)}\right]^{\frac{1}{2}} & \text{if } J' = J-1 \text{ and } m' = m. \\ 0 & \text{otherwise.} \end{cases} \tag{A.16}$$

A.4 Quadratic Stark Shifts

The Hamiltonian for the Stark shift is

$$H_{\text{St}} = -\vec{d} \cdot \vec{\mathcal{E}} = -\mathcal{E}\vec{d} \cdot \hat{z} = -\mathcal{E} d_{\text{mol}} \cos\theta. \tag{A.17}$$

In the perturbative limit we assume that the dipole interaction is much smaller than the rigid rotor energy level splitting, i.e. $\mathcal{E} d_{\text{mol}} \ll \hbar^2/2I$.

The first-order Stark shift is $E^{(1)}_{J,m} = \langle H_{\text{St}} \rangle = -\langle \vec{d} \rangle \cdot \vec{\mathcal{E}}$, which we know vanishes from Section A.2.

We next examine the quadratic Stark shift. From non-degenerate second-order perturbation theory,[3] we have (see, e.g. Griffiths chapter 6.1.3 [90]):

$$E^{(2)}_{J,m} = \sum_{J'=J\pm1} \frac{|\langle J', m| H_{\text{St}} |J, m\rangle|^2}{E_{J,m} - E_{J',m}}, \tag{A.18}$$

where I have excluded all terms in the sum ($J' \neq J \pm 1$, $m' \neq m$) that were found to vanish in Section A.3. We now explicitly substitute the rigid rotor energies from Eq. (A.3) and

[3] Even though rigid rotor eigenstates with the same J are degenerate, we can use non-degenerate perturbation theory because as shown in Section A.3, the dipole matrix elements only connect states of different J.

write out $\langle J', m| H_{\mathrm{St}} |J, m\rangle$ in terms of dipole matrix elements:

$$E^{(2)}_{J,m} = \frac{\mathcal{E}^2 d^2_{\mathrm{mol}}}{\hbar^2/2I} \sum_{J'=J\pm1} \frac{|\langle J', m| \cos\theta |J, m\rangle|^2}{J(J+1) - J'(J'-1)}. \tag{A.19}$$

Now we can substitute in the formula for the matrix elements calculated in Eq. (A.16), perform the sum over $J' = J \pm 1$, and after a few lines of algebra, we obtain:

$$E^{(2)}_{J,m} = \frac{\mathcal{E}^2 d^2_{\mathrm{mol}}}{\hbar^2/2I} \times \begin{cases} -\frac{1}{6} & \text{if } J = 0. \\ \frac{1}{2}\left[\frac{(J-m)(J+m)}{J(2J-1)(2J+1)} - \frac{(J-m+1)(J+m+1)}{(J+1)(2J+1)(2J+3)}\right] & \text{if } J \geq 1. \end{cases} \tag{A.20}$$

(This result agrees with Eq. 10.8 and 10.9 in [187] for the Stark shift of a linear molecule.)

The quadratically Stark-shifted energy levels are plotted in Fig. A.4.1. Let's attempt a physical interpretation of a few aspects of this figure and the formulae on which it is based.

1. The Stark shift is symmetric in m (Notice that changing every instance of m to $-m$ leaves the formula unaltered). This makes physical sense because parity symmetry dictates that it shouldn't matter whether the wavefunction is circulating clockwise or counter-clockwise with respect to the direction of the electric field.

2. The magnitude of the Stark shift decreases for increasing J because there are more powers of J in the denominator than in the numerator of Eq. (A.20). The physical cause of this diminished perturbation is that as J increases, the rigid rotor level spacing also increases (in proportion to J), and mixing with adjacent J levels is therefore suppressed by the energy denominator in Eq. (A.18).

3. For the extremal values of m, given by $|m| = J$, the first (positive) term in Eq. (A.20) vanishes, so that the quadratic Stark shift is always negative. This happens because the dipole interaction causes adjacent energy levels to mix and repel, and the state $|J, \pm J\rangle$ can only mix with the energy level above it, $|J+1, \pm J\rangle$; the $J-1$ level has no

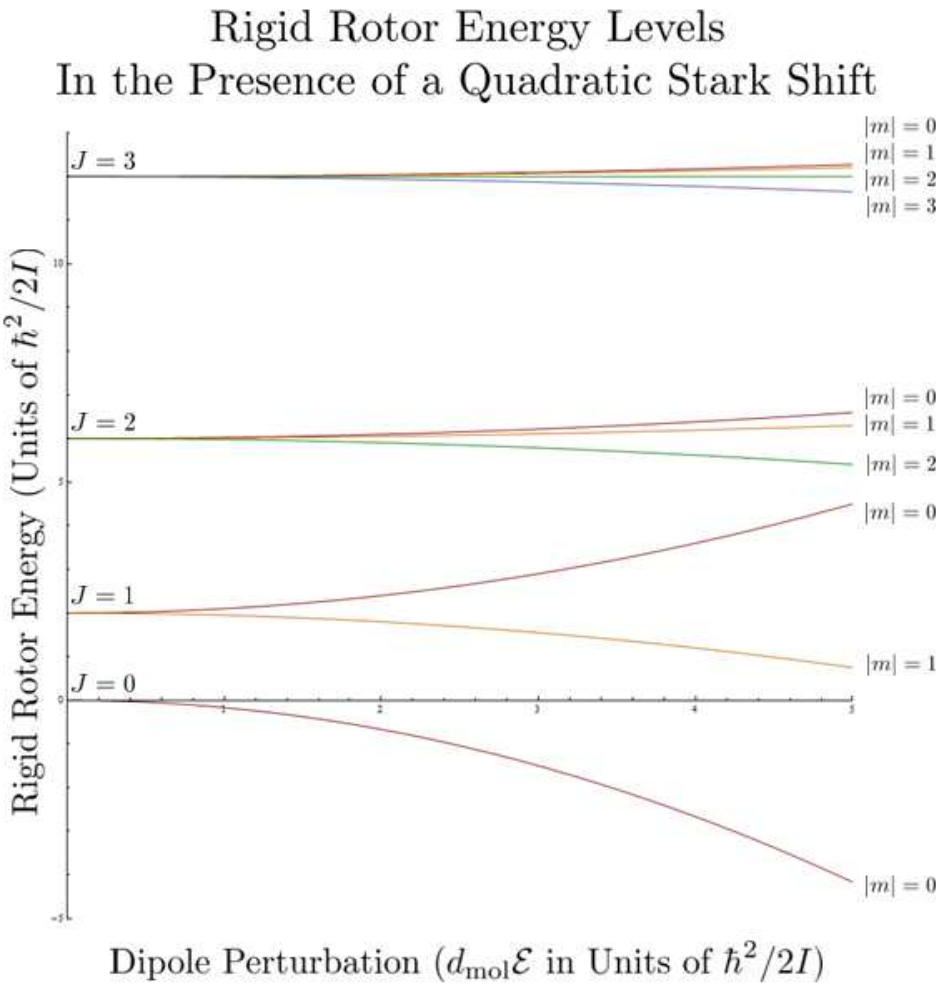

Figure A.4.1: Energy levels of a rigid rotor molecule with dipole moment d_{mol} in an applied electric field $\mathcal{E}$. The quadratic Stark shifts are given as a function of electric field in units of rotational constant $\hbar^2/2I$ divided by d_{mol}. Technically, these perturbative results are valid only at $\mathcal{E}$ values much less than 1 on the x-axis, but the Stark shifts are shown here out to large values of $\mathcal{E}$ so that the shifts at higher J are visible.

$m = J$ sublevel with which the $|J, \pm J\rangle$ state can interact under dipole selection rules.

4. For a given rotational level, the quadratic Stark shift decreases monotonically with $|m|$; some examination of Eq. (A.20) reveals that the second (negative) term gradually wins out over the first (positive) term as $|m|$ approaches J. I find this behavior a bit tricky to explain from a physical perspective, but I'll make two attempts: one taking the classical point of view and one using quantum mechanics.

   In the classical picture, this pattern can be understood by imagining a dipole with a given angular momentum $\vec{J}$ whose axis of rotation points at an angle θ relative to the external field $\vec{\mathcal{E}}$. As θ approaches 0 or π (equivalent to m approaching $\pm J$), the dipole becomes oriented more and more orthogonal to the field. Therefore, the torque $|\vec{d} \times \vec{\mathcal{E}}| = d \cdot \mathcal{E} \sin(\theta + \pi/2)$, which tends to align the dipole and lower its energy, becomes stronger. Thus molecules with large values of $|m|$ tend to experience a negative Stark shift. Conversely, as θ approaches $\pi/2$, the torque that tends to align the dipole vanishes on average, and the dipole spins so that it is alternately aligned and anti-aligned with the applied field. As the spinning dipole approaches alignment with the applied field, the torque from the field accelerates its rotation so that the dipole's maximum angular velocity occurs when it is aligned with $\vec{\mathcal{E}}$. For similar reasons, the minimum angular velocity occurs when the dipole is anti-aligned with the applied field. Thus a spinning molecule with small m spends more time pointing against $\vec{\mathcal{E}}$ than with it and experiences a positive Stark shift. (Thanks to Townes and Schawlow chapter 10.1 [187] for help refining this argument.)

   Returning to the quantum mechanical picture, I believe we can understand this effect somewhat intuitively if we stare long enough at the amplitudes of the spherical harmonics. Note that for a given value of m, Y^m_{J+1} has one more node along the z-axis than Y^m_J. Thus, mixing with $|J+1, m\rangle$ adds a piece of wavefunction that interferes constructively with $|J, m\rangle$ along the positive z-axis but destructively along the negative

z-axis, thereby aligning the dipole and lowering its energy. The larger the value of $|m|$ for a given J, the fewer nodes already exist along the z-axis, and so the more effective the interference pattern becomes at aligning the dipole.

A.5 Induced Dipole Moment

To calculate the energy shift due to an induced dipole, imagine beginning with the molecule in zero electric field where the dipole moment is zero and ramping the field up to its final value of $\mathcal{E}$. Then the energy shift is given by

$$\Delta E = -\int_0^{\mathcal{E}} \langle d_z(\mathcal{E}')\rangle_{(J,m)} d\mathcal{E}', \tag{A.21}$$

where the induced dipole $\langle d_z(\mathcal{E})\rangle_{(J,m)}$ is the expectation value of the dipole operator in the perturbed eigenstates. The energy shift ΔE due to the induced dipole moment is equivalent to the quadratic Stark shift calculated in Section A.4. With the benefit of hindsight, we can express this Stark shift as $E^{(2)}_{J,m} = -\alpha(J,m)/2 \times \mathcal{E}^2$, where $\alpha(J,m)$ is a constant of proportionality given by Eq. (A.20). By substituting this into Eq. (A.21) differentiating both sides with respect to $\mathcal{E}$, we obtain

$$\frac{\partial}{\partial \mathcal{E}}\left[-\frac{1}{2}\alpha(J,m)\mathcal{E}^2\right] = \frac{\partial}{\partial \mathcal{E}}\left[-\int_0^{\mathcal{E}} \langle d_z(\mathcal{E}')\rangle_{(J,m)} d\mathcal{E}'\right] \tag{A.22}$$

$$\alpha(J,m)\mathcal{E} = \langle d_z(\mathcal{E})\rangle_{(J,m)}. \tag{A.23}$$

Thus an induced dipole moment is proportional to the applied electric field with a constant of proportionality α, known as the "polarizability."

We can use Eq. (A.23) to express the quadratic Stark shift in terms of the dipole moment as $E^{(2)}_{J,m} = -\langle d_z(\mathcal{E})\rangle_{(J,m)}/2 \times \mathcal{E}$. Solving for the induced dipole moment, we have

$$d_{z(J,m)}(\mathcal{E}) = -\frac{2E^{(2)}_{J,m}}{\mathcal{E}}. \tag{A.24}$$

Finally, we can use Eq. (A.20) to calculate the right-hand side of the equation above for all the states in the $J = 0$ and $J = 1$ manifolds and obtain the solutions:

$$d_{z(0,0)}(\mathcal{E}) = \frac{d_{\text{mol}}\mathcal{E}}{\hbar^2/2I} \times \frac{d_{\text{mol}}}{3} \tag{A.25}$$

$$d_{z(1,0)}(\mathcal{E}) = -\frac{d_{\text{mol}}\mathcal{E}}{\hbar^2/2I} \times \frac{d_{\text{mol}}}{5} \tag{A.26}$$

$$d_{z(1,\pm1)}(\mathcal{E}) = \frac{d_{\text{mol}}\mathcal{E}}{\hbar^2/2I} \times \frac{d_{\text{mol}}}{10}. \tag{A.27}$$

Note that the induced dipole is negative (i.e., anti-aligned with the applied field) for the $|J = 1, m = 0\rangle$ state, where the Stark shift is positive.

B

ThO Source Targets

B.1 ThO_2 Ablation Targets

The goal of our Gen. I target-making program was to create hard, dense, mechanically robust ThO_2 targets that could withstand the violent ablation process with a long lifetime while producing consistently high ThO yields. The nanosecond ablation process, which involves the local, instantaneous deposition of gigawatts of power per square millimeter, is extremely complex (see e.g. reference [40]), involving light-matter interactions, solid, liquid, and gas-phase physics, plasma dynamics, and chemistry. Nevertheless, both the existing literature and the collective experience of the atomic physics community provide some insight into optimization strategies.

Anecdotal evidence suggests that pure crystalline solids tend to produce sub-optimal yields, possibly because their good thermal conductivity enlarges the volume over which the energy of the ablation laser is deposited, resulting in more energy dissipation via vibrations within and damage to the crystal lattice, and fewer chemical bond ruptures producing

the desired gas-phase species. Pressed powder targets are also not ideal because they are mechanically weak and tend to disintegrate rapidly under laser fire. Ceramics, which have a lower thermal conductivity than single-crystal solids due to their granularity and porosity, but which are mechanically stronger than pressed powder because the individual grains are fused together, seem to represent a happy medium.

Our ThO_2 ceramic ablation target recipe is based on a paper by Balakrishna et al. [19] and a formula for mock ThO_2 "nuclear fuel pellets" made for testing purposes that was emailed to us by a researcher at Oak Ridge National Labs [116]. For details on individual target-making runs, visit the Google Sheet "ThO2 Target Summary" at `https://goo.gl/XsPjUf`.

As a tightly regulated nuclear material, ThO_2 is sold in a limited number of forms, so we began with a precursor material that we could easily obtain: 325 mesh (44 μm) powder from Materion/Cerac. Our step-by-step target-making procedure is provided below and illustrated in Fig. B.1.2. Here I briefly describe the rationale behind some of the more potentially alchemical-seeming steps.

Transforming a "green" material made of individual grains into a ceramic requires firing or "sintering" the material at high temperature so that the grains soften and flow together, resulting in a reduction in porosity, an increase in strength, and shrinking and densification of the material. As a general rule, the ideal temperature for sintering a ceramic is a little below the melting point of the material. Because ThO_2 is extremely refractory, with a melting point of 3390°C, well above that of most heating elements, sintering it in an ordinary furnace poses particular challenges. Balakrishna et al. [19] found that the sintering temperature of ThO_2 could be dramatically reduced to $\approx$ 1200°C by the addition of a small amount of niobium pentoxide (Nb_2O_5), which acts as a "sintering agent." The mechanism, known as "activated sintering," is illustrated in Fig. B.1.1.

In order for sintering to proceed at all, the grains must initially be in good mechanical contact with each other. We create this condition by compacting the powder into a pellet before sintering it. The compacting pressures, given in the procedure below, are approxi-

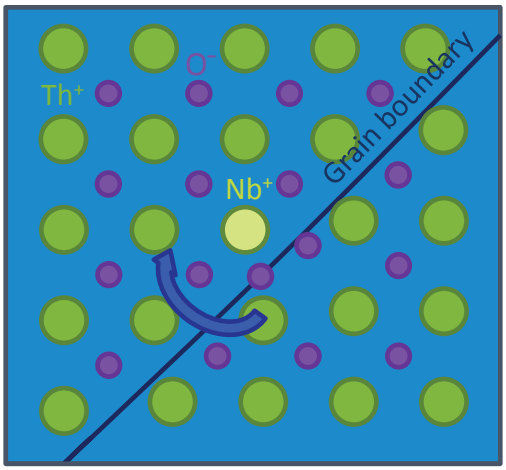

Figure B.1.1: Cartoon illustration of one mechanism of activated sintering: Cations are represented by large green dots and anions by small purple dots. Impurities create defects in the lattice structure which reduce the energy barrier to ion diffusion. This added mobility allows grain boundaries to flow more easily below the melting point and facilitates ceramic densification at lower temperatures than in the absence of such impurities [19].

mately those recommended in the references [19, 116]. We found that varying the pressure or the pressing time over a fairly wide range did not significantly affect the final product. We did find that unless the pellet thickness-to-diameter ratio was $\lesssim 1$, the pellets came out weak and crumbly, probably because of force gradients created by powder jamming within the pressing die. Crumbling was also more common with dies 1/4" in diameter than with those 1/2" and above, perhaps because the higher surface-area-to-volume ratio in the small dies resulted in more friction.

Even with more ideal target aspect ratios and die sizes, the "green" (un-sintered) pellets were initially fragile and would often break during die extraction or handling. It is perhaps intuitive to see why: ThO_2 is a hard crystalline substance with very little "give." This makes it hard for the grains to achieve sufficient contact surface area to form the van der Waals bonds that allow them to adhere together. (As a point of reference in our everyday experience, it is easy to make a small "patty" by pressing squishy, moisture-laden flour together between your fingers, but try the same thing with very fine, dry sand, and it will fall apart.) To overcome this problem, we added a small amount of "binder," a sticky organic polymer, to the powder. We tried and rejected stearic acid before settling on polyethylene

glycol with molecular weight 8000.[1] Targets with roughly 5–6% by weight of this binder were much more mechanically robust and easier to handle in the "green" state.

The quality of the final ceramic depends on a large number of factors, including the purity and homogeneity of the precursor and the initial grain configuration and size distribution. In order to improve homogeneity, we mix the powders together for several hours in a ball mill with a few 1/4" diameter stainless steel ball bearings. To help the powder find a "good" grain configuration, we perform a pre-compaction step, in which we press a pellet that we then break up and rub through a mesh. This step can be thought of as a way of exploring the grain configuration space to help the grains find local energy minima in which they are clumped together. In order to improve the powder purity and promote grain growth, we also include a "calcining" step, in which the pre-compacted powder is heated to a temperature below its sintering point to burn out any volatile substances and allow the grains to begin to fuse. It should be noted that in our small sample size, we did not observe a difference between targets in which the precursor was ball milled, pre-compacted, and calcined and those in which these steps were omitted. Nevertheless, because they are recommended by many ceramics-making references, we continued to include these steps.

Procedure for ThO_2 ablation target preparation:

1. Mix 2.52 parts niobia (Nb_2O_5) with 1000 parts thoria (ThO_2) by weight (0.25 mole % Nb_2O_5)
2. Mix in 5 weight % unexpired PEG 8000 binder
3. Add distilled water to make a slurry, then ball mill for at least 4 h
4. Dry powder in $\approx 70°C$ furnace
5. Pre-compact at 128 MPa $\approx$ 18.5 ksi in 3/4" diameter die (3.28 mV readout on Futek button-top load cell)
6. Break up the pre-compacted powder and rub through 1 mm nylon mesh
7. Calcine in a crucible at 1000°C for at least 1 h

[1] The Oak Ridge recipe used PVP K-15, which we did not try [116].

8. Mix in 6 weight % unexpired PEG 8000 (polyethylene glycol) binder and ball mill for at least 4 h

9. Press slowly to 180 MPa ≈ 26.1 ksi (4.612 mV load cell readout for 3/4" diameter die)

10. Sinter for 3 hours at 1150°C, ramping to the target temperature at a rate of 300°C/h

11. Repeat sintering as necessary until densification has reached desired level

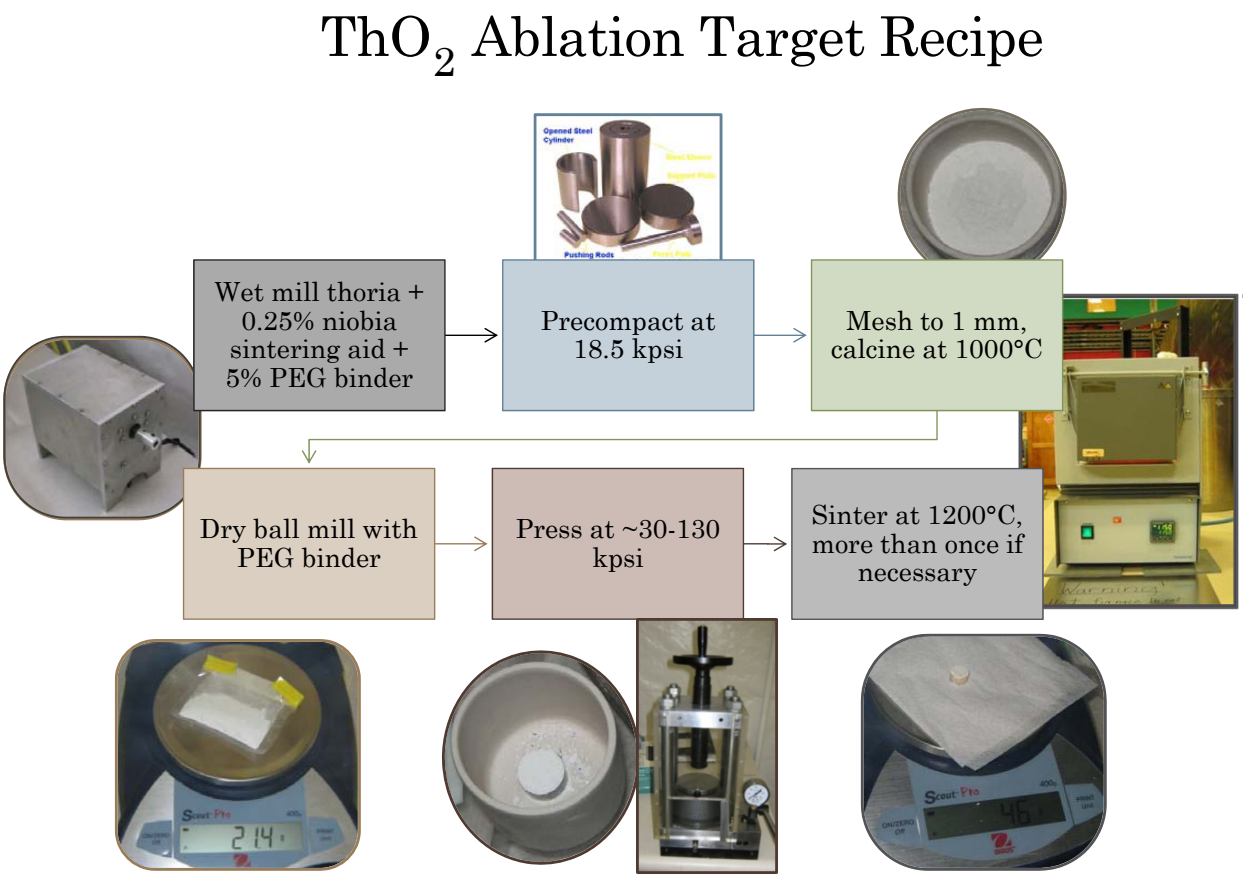

Figure B.1.2: Ablation target recipe with equipment and photographs of target appearance at each stage. The furnace was a benchtop 1200°C box furnace from MTI, the 40 metric ton hydraulic press was also a benchtop MTI model, and the gauge used to verify the compacting pressure was a high-capacity load button load cell from FUTEK with a stable voltage supply made by Jim MacArthur. The ball mill was a homebuilt setup consisting of a canister attached to the horizontal drive shaft of a small 60 RPM motor. The crucibles were high-purity alumina or zirconia rated to over 1500°C.

In 2015 we received a ThO_2 fuel pin from Idaho National Labs that should obviate the need to fabricate our own ablation targets. Details on this fantastic acquisition are provided

in my lablog entries "Opening ThO2 INL Fuel Pin" (July 24, 2015) and "INL Shipment of Thorium-232 Materials" (January 20, 2015) and Jacob Baron's lablog entry "Thorium Fuel Pin XRD" (July 27, 2015).

B.2 Mixed ThO_2 + Th Thermochemical Source Targets

The recipe for making a mixed ThO_2 + Th target is essentially a subset of the ThO_2 ceramic target recipe. The steps are enumerated below. For details on individual target-making runs, visit the Google Sheet "Mixed Pressed Powder Target Summary" at `https://goo.gl/wbUeZ0`.

The ThO_2 powder is the same 325 mesh powder from Cerac as that used for the ablation targets, while the Th powder is from MP Biomedicals and has unspecified grain size and purity. When the Th powder was initially sourced by MP Bio, its purity was 99.1%, but by the time they sold it to us, they could no longer guarantee this purity level due to the tendency of thorium to oxidize slowly, even though the bottles it was stored in were sealed in an argon atmosphere. We found that about 3/4 of the thorium powder passed through a size 325 mesh, indicating a typical powder diameter of < 44 μm, and examination under a microscope confirmed that the average grain size was around 30 μm, as shown in my lablog entry "Thorium metal powder" (January 13, 2010). More recently, we have acquired some 325 mesh (44 μm) thorium powder from International Bio-Analytical Industries, Inc. (IBI Labs), which has not yet been used in targets.[2]

The ratio 75% ThO_2 to 25% Th was chosen as a compromise between chemistry and physics: We wanted a target with thermal conductivity close to that of insulating ThO_2 in order to minimize the power required to produce the high local temperatures needed for the reaction (see Section 4.2); however, we also wanted to use the target material efficiently to produce ThO, which argued for a stoichiometric 50-50 ratio of thorium to thoria. A target composition with 75% ThO_2 was roughly as close to the stoichiometric ratio as we could go

[2]For future reference, we have also purchased a 1 kg block of thorium metal from IBI Labs, which they are currently storing for us in their Florida facility. Our contact is Alex Besenyo.

while keeping the thorium concentration below the "percolation threshold" [98, 156]. The percolation threshold is the concentration at which a species in a solid mixture is guaranteed to form chains that go all the way through the composite. When the concentration of metallic species in a metal-insulator composite is raised above the percolation limit, the composite's *electrical* conductivity increases dramatically [98], and since thermal and electrical conductivities often track each other (albeit very imperfectly: see e.g. [129]), we elected to keep our targets below this threshold. A target that we made with 50% ThO_2 and 50% Th to test the hypothesis that a stoichiometric ratio would produce worse yields than a more insulating ratio did indeed give instantaneous beam fluxes a factor of 2 or so worse than the typical 75% ThO_2 + 25% Th targets.

In an effort spearheaded by Jacob Baron, we attempted to sinter the mixed targets to see if this would improve their strength or yield. Unfortunately, the sintered target we tried produced poor yields. We believe that this is because the sintering process involved temperatures high enough to drive the $Th + ThO_2 \rightarrow 2ThO$ reaction, so that the targets were basically already depleted by the time we tried to use them. Interestingly, the sintered targets also appeared to show an x-ray diffraction peak indicating the presence of crystalline ThO. In the end, we settled on using "green" pellets of pressed powder in the thermochemical source.

Procedure for Th + ThO_2 thermochemical source target preparation:

1. Mix 1 part Th by weight with 3 parts ThO_2
2. Mix in 3 weight % unexpired PEG 8000 binder
3. Add distilled water to make a slurry, then ball mill for at least 3 h
4. Dry powder in $\approx 40°C$ furnace
5. Break up the powder and rub through 1 mm nylon mesh
6. Press slowly to 640 MPa $\approx$ 93 ksi (16.3 mV load cell readout for 3/4" diameter die)

Proper radiation protocol (maintained by the current radiation safety officer) must always be followed when performing either of these target making procedures, and the target-making

and radiation safe inventory logs must be kept up to date. The final products are shown in Fig. B.2.1

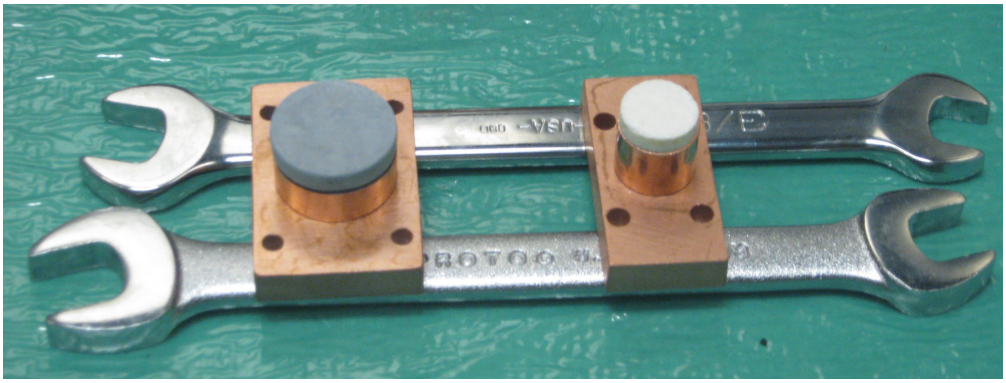

Figure B.2.1: Mounted thermochemical (left) and ablation (right) targets.

C

Buffer Gas Cooling Time and Length Scales

To obtain an estimate of the number of collisions required to cool the molecules, we follow the argument in [117]. We assume that the collisions are elastic and isotropic and consider a single collision between the species of interest with mass M and a buffer gas atom with mass m, as illustrated in Figure C.0.1. Then we can write energy and momentum conservation in the center-of-mass frame as follows:

$$\vec{p} = 0 = M\vec{V}_i + m\vec{v}_i = M\vec{V}_f + m\vec{v}_f \tag{C.1}$$

$$E = \frac{M}{2}\vec{V}_i^2 + \frac{m}{2}\vec{v}_i^2 = \frac{M}{2}\vec{V}_f^2 + \frac{m}{2}\vec{v}_f^2, \tag{C.2}$$

where $\vec{V}_i$ and $\vec{V}_f$ are, respectively, the initial and final velocity of the species of interest, and $\vec{v}_i$ and $\vec{v}_f$ are, respectively, the initial and final velocity of the buffer gas atom. (Variables without primes will be used to designate quantities in the center-of-mass frame. Primes will denote lab-frame variables.) Then from momentum conservation in the CM frame

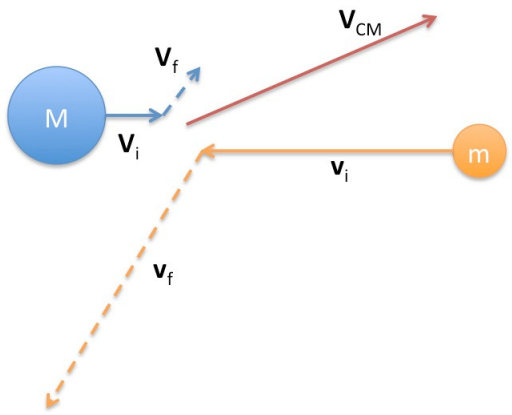

Figure C.0.1: A collision between a buffer gas atom and species of interest in the center-of-mass frame.

(Eq. (C.1)), we have:

$$\vec{v}_i = -\frac{M}{m}\vec{V}_i \quad \text{and} \quad \vec{v}_f = -\frac{M}{m}\vec{V}_f. \tag{C.3}$$

Substituting these relations into the conservation of energy relation in Eq. (C.2) gives

$$\frac{M}{2}\left(1+\frac{M}{m}\right)\vec{V}_i^2 = \frac{M}{2}\left(1+\frac{M}{m}\right)\vec{V}_f^2 \tag{C.4}$$

$$\vec{V}_i^2 = \vec{V}_f^2. \tag{C.5}$$

So for an elastic collision in the CM frame, no energy is exchanged between the particles. To determine the average kinetic energy lost by the hot species in the lab frame, we take an ensemble average over all the particles and do some algebra. The goal will be to write the difference between the mean initial and final kinetic energies of the species of interest in the lab frame in terms of the mean initial kinetic energies of the species and buffer gas atom. We will then be able to relate these quantities to the temperatures. For the following, we

will need the center-of-mass velocity in the lab frame:

$$\vec{V}_{\rm CM} = \frac{1}{M+m}\left(M\vec{V}_i' + m\vec{v}_i'\right) = \frac{1}{M+m}\left(M\vec{V}_f' + m\vec{v}_f'\right). \tag{C.6}$$

The primed variables denote lab frame quantities. $\vec{V}_{\rm CM}$ is unprimed for notational simplicity but is also, of course, a lab-frame quantity. We will also need the ensemble average of $\vec{V}_{\rm CM}^2$:

$$\langle \vec{V}_{\rm CM}^2 \rangle = \frac{1}{(M+m)^2}\left(M^2\langle(\vec{V}_i')^2\rangle + m^2\langle(\vec{v}_i')^2\rangle\right). \tag{C.7}$$

For this ensemble average, we have taken the first equality from Eq. (C.6) and squared it, then eliminated all cross terms proportional to $\langle \vec{V}_i' \cdot \vec{v}_i' \rangle$. These cross terms express correlations between the incoming particle velocities in the lab frame, which vanish on the assumption that collisions are isotropic.

Next, we take the ensemble (bracketed) average of the CM frame quantities in order to express them as measurable lab frame quantities:

$$\langle \vec{V}_f^2 \rangle = \langle \vec{V}_i^2 \rangle = \left\langle \left(\vec{V}_i' - \vec{V}_{\rm CM}\right)^2 \right\rangle \tag{C.8}$$

$$= \langle(\vec{V}_i')^2\rangle + \langle \vec{V}_{\rm CM}^2 \rangle - 2\langle \vec{V}_i' \cdot \vec{V}_{\rm CM} \rangle \tag{C.9}$$

$$= \langle(\vec{V}_i')^2\rangle + \langle \vec{V}_{\rm CM}^2 \rangle - \frac{2M}{M+m}\langle(\vec{V}_i')^2\rangle - \frac{2m}{M+m}\langle \vec{V}_i' \cdot \vec{v}_i' \rangle \tag{C.10}$$

$$= \left(1 - \frac{2M}{M+m}\right)\langle(\vec{V}_i')^2\rangle + \langle \vec{V}_{\rm CM}^2 \rangle \tag{C.11}$$

$$= \left(1 - \frac{2M}{M+m} + \frac{M^2}{(M+m)^2}\right)\langle(\vec{V}_i')^2\rangle + \frac{m^2}{(M+m)^2}\langle(\vec{v}_i')^2\rangle. \tag{C.12}$$

In the next-to-last step, we have again used isotropy of collisions to eliminate the final term in Eq. (C.10).

Next, using the previous results, we write the ensemble average of the lab frame measurables in terms of each other. We will need to use the fact that the particle velocities in the

center-of-mass frame are independent of $\vec{V}_{\mathrm{CM}}$ to eliminate correlations of the form $\langle \vec{V}_f \cdot \vec{V}_{\mathrm{CM}} \rangle$.

$$\langle (\vec{V}_f')^2 \rangle = \left\langle \left(\vec{V}_f + \vec{V}_{\mathrm{CM}} \right)^2 \right\rangle \tag{C.13}$$

$$= \langle \vec{V}_f^2 \rangle + \langle \vec{V}_{\mathrm{CM}}^2 \rangle \tag{C.14}$$

$$= \left(1 - \frac{2M}{M+m} + \frac{2M^2}{(M+m)^2} \right) \langle (\vec{V}_i')^2 \rangle + \frac{2m^2}{(M+m)^2} \langle (\vec{v}_i')^2 \rangle \tag{C.15}$$

$$= \frac{M^2+m^2}{(M+m)^2} \langle (\vec{V}_i')^2 \rangle + \frac{2m^2}{(M+m)^2} \langle (\vec{v}_i')^2 \rangle. \tag{C.16}$$

To obtain Eq. (C.15), we substituted in Eq. (C.7) for $\langle \vec{V}_{\mathrm{CM}} \rangle$ and Eq. (C.12) for $\langle \vec{V}_f^2 \rangle$.

Now we can derive the average change in temperature of the species of interest ensemble due to each particle undergoing a single collision. According to the equipartition theorem, for a system in thermal equilibrium at temperature T, each quadratic degree of freedom of the system carries an average energy of $k_{\mathrm{B}}T/2$. Therefore, the average translational kinetic energy $M\langle \vec{V}^2 \rangle/2$, which is a quadratic degree of freedom, is related to the temperature by

$$M\langle \vec{V}^2 \rangle = k_{\mathrm{B}} T. \tag{C.17}$$

Anticipating this equipartitional proportionality between $M\langle (\vec{V}')^2 \rangle$ and temperature, we use Eq. (C.16) to write:

$$\langle (\vec{V}_f')^2 - (\vec{V}_i')^2 \rangle = -\frac{2Mm}{(M+m)^2} \langle (\vec{V}_i')^2 \rangle + \frac{2m^2}{(M+m)^2} \langle (\vec{v}_i')^2 \rangle \tag{C.18}$$

$$\langle M(\vec{V}_f')^2 - M(\vec{V}_i')^2 \rangle = -\frac{2Mm}{(M+m)^2} \left(M\langle (\vec{V}_i')^2 \rangle - m\langle (\vec{v}_i')^2 \rangle \right). \tag{C.19}$$

Finally, using Eq. (C.17), we can write the change in temperature of the particle per collision with the buffer gas in the following delightfully simple form:

$$T_f - T_i = -\frac{2Mm}{(M+m)^2} \left(T_i - T_b \right), \tag{C.20}$$

where T_b is the temperature of the buffer gas.

Taking the continuum limit, we can write Eq. (C.20) as a differential equation,

$$\frac{dT}{dn} = -\frac{2Mm}{(M+m)^2}\left(T - T_b\right), \tag{C.21}$$

where n is the number of collisions undergone by the species of interest. Note that the buffer gas temperature is assumed not to change since its density is so much higher than that of the introduced species.

By changing variables to $u = T - T_b$, with initial condition $u = T_i - T_b$, the differential equation is easily solved:

$$T = (T_i - T_b)\exp\left[-\frac{2Mm}{(M+m)^2}n\right] + T_b \tag{C.22}$$

Thus, the characteristic number of collisions for thermalization is $\frac{(M+m)^2}{2Mm}$, which, given that we ordinarily have $M \gg m$, is approximately $\frac{M}{2m}$. For typical masses $40 < M < 400$ for atomic and molecular species of interest and $4 < m < 20$ for buffer gases, 100 collisions is usually enough to bring the species into thermal equilibrium with the buffer gas.

For a typical cryogenic buffer gas cell mean free path of ~ 0.1 mm [105], this gives a "thermalization length" of ~ 1 cm. For a typical few-kelvin thermal velocity of ~ 100 m/s, the thermalization time is $\sim 100 \times 0.1\,\text{mm}/(100\,\text{m/s}) = 0.1$ ms.

D

Feldman-Cousins Confidence Intervals

A classical (i.e. frequentist) confidence interval [164] is a natural choice for reporting the result of an eEDM measurement.[1] For repeated and possibly different experiments measuring the eEDM, the frequency with which the confidence intervals include or exclude the value $d_e = 0$ suggests whether the results are consistent or inconsistent, respectively, with the Standard Model. Furthermore, the confidence level (C.L.) represents an objective measure of the *a priori* probability that the confidence interval assigned to any one of these measurements, selected at random, includes the unknown true value of the eEDM $d_{e,\text{true}}$. Since no statistically significant eEDM has yet been observed, the recent custom has been for electron eEDM experiments to report an upper limit at the 90% C.L. [103, 162]. The proper interpretation of such limits is that if the experiment were performed a large number of times, and the confidence interval were *computed in the same way* for each experimental trial, $d_{e,\text{true}}$ would fall within the interval 90% of the time.

[1] This section is adapted from [11]. I thank the ACME collaboration for their contributions and editing assistance. In particular, Adam West helped with the figure design and Brendon O'Leary checked the calculations and helped elucidate their interpretation.

Feldman and Cousins pointed out that in order for this interpretation to be valid, the confidence interval construction must be independent of the result of the measurement [77]. If the procedure for constructing 90% confidence intervals is chosen contingent upon the measurement outcome, the resulting intervals may "undercover," i.e. fail to include the true value more than 10% of the time. Feldman and Cousins termed this inconsistent approach "flip-flopping."

For example, suppose we were to report a two-sided central confidence interval if our measured eEDM, $d_{e,\mathrm{meas}}$, fell two or more standard deviations from zero and an upper bound based on a folded Gaussian distribution otherwise. The resulting confidence band would be that shown in Fig. D.0.1. The confidence interval for a particular measurement outcome is given by the vertical range of the shaded region at the position of the measured value $x = d_{e,\mathrm{meas}}$. If the true value of the eEDM were in the range $1.6\sigma \leq d_{e,\mathrm{true}} \leq 3.7\sigma$, our claimed "90% confidence intervals" would include the true value in fewer than 90% of our experimental trials.

In order to avoid flip-flopping, we chose a confidence interval construction, the Feldman-Cousins method described in reference [77], that consistently unifies these two limits. We applied this method to a model with Gaussian statistics, in which the measured magnitude of the eEDM channel, $x = |\omega^{\mathcal{N}E}_{\mathrm{meas}}|$, is sampled from a folded Gaussian distribution

$$P(x|\mu) = \frac{1}{\sigma\sqrt{2\pi}}\left(\exp\left[-\frac{(x-\mu)^2}{2\sigma^2}\right] + \exp\left[-\frac{(x+\mu)^2}{2\sigma^2}\right]\right), \tag{D.1}$$

where the location parameter is the unknown true magnitude of the eEDM channel, $\mu = |\omega^{\mathcal{N}E}_{\mathrm{true}}|$, and the scale parameter σ is equal to the quadrature sum of the statistical and systematic uncertainties given in Eq. 2.20 and at the bottom of Fig. 2.7.2.

The central idea of the Feldman-Cousins approach is to use an ordering principle which, for each possible value of the parameter of interest μ, ranks each possible measurement outcome x by the "strength" of the evidence it provides that μ is the true value. The values of x

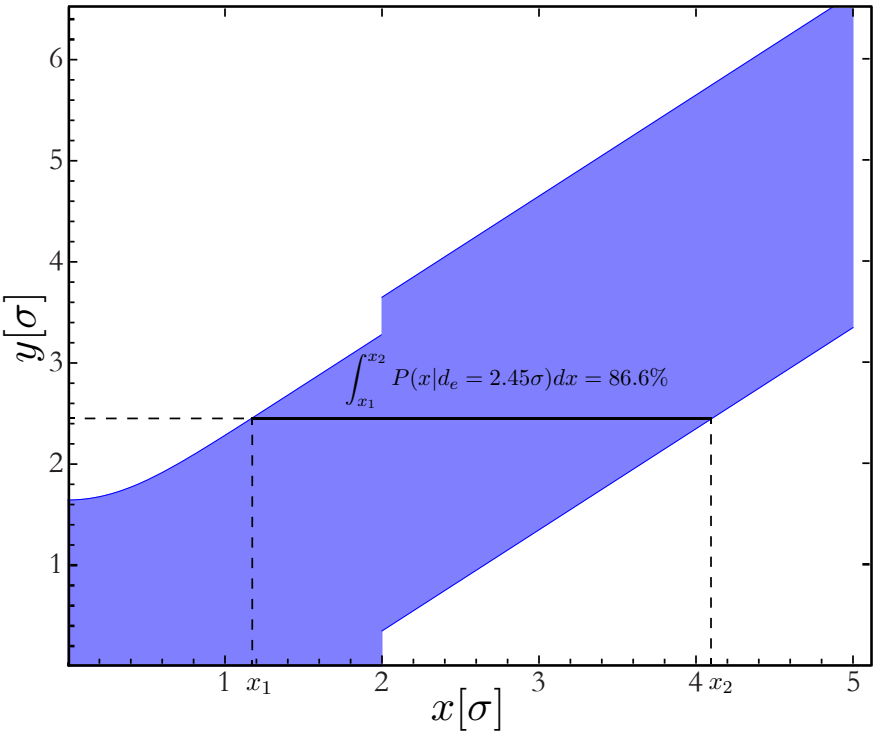

Figure D.0.1: Flip-flopping "90% confidence" band described in the text. These intervals are not true 90% confidence intervals because if the true value of the eEDM falls between about 1.6σ and 3.7σ, fewer than 90% of experimental trials will contain the true value. An example of this undercoverage is shown in the plot for the case $d_{e,\text{true}} = 2.45\sigma$.

that provide the strongest evidence for each value of μ are included in the confidence band for that value. In the Feldman-Cousins method, the metric for the strength of evidence is the likelihood of μ given that x is measured (i.e. $\mathcal{L}(\mu|x) = P(x|\mu)$), divided by the largest probability x can possibly achieve for any value of μ. The denominator in this prescription takes into account the fact that an experimental result that is somewhat improbable under a particular hypothesis can still provide good evidence for that hypothesis if the result is similarly improbable under even the most favorable hypothesis. This approach has its theoretical roots in likelihood ratio testing [182].

Our specific procedure for computing confidence intervals was a numerical calculation performed using the following recipe (cf. Fig. D.0.2):

1. Construct the confidence bands on a Cartesian plane, of which the horizontal axis represents the possible values of x and the vertical axis the possible values of μ. Divide

the plane into a fine grid with x-intervals of width Δ_x and μ-intervals of height Δ_μ. We will consider only the discrete possible values $x_i = i\Delta_x$ and $\mu_j = j\Delta_\mu$, where the index $i(j)$ runs from 0 to $n_x(n_\mu)$.

2. For all values of i, maximize $P(x_i|\mu_j)$ with respect to μ_j. Label the maximum points $\mu^{\max,i}$.

3. For some value of j, say $j = 0$, compute the likelihood ratio $R(x_i) = P(x_i|\mu_j)/P(x_i|\mu^{\max,i})$ for every value of i.

4. Construct the "horizontal acceptance band" at μ_j by including values of x_i in descending order of $R(x_i)$. Stop adding values when the cumulative probability reaches the desired C.L. of 90%, i.e., $\sum_{x_i} P(x_i|\mu_j)\Delta_x = 0.9$.

5. Repeat steps (3)–(4) for all values of j.

6. To determine the reported confidence interval, draw a vertical line on the plot at $x = |\omega^{\mathcal{NE}}_{\text{meas}}|$. The 90% confidence interval is the region where the line intersects the constructed confidence band.

The left-hand plot in Fig. D.0.2 was generated using the prescription above at several different C.L.'s. Note that the 90% confidence intervals switch from upper bounds to two-sided confidence intervals when the value of $|\omega^{\mathcal{NE}}_{\text{meas}}|$ becomes larger than 1.64σ. This is the level of statistical significance required to exclude the value $d_e = 0$ from a 90% C.L. central Gaussian confidence band.

From Eq. (2.20), we find $|\omega^{\mathcal{NE}}_{\text{meas}}| = 0.46\sigma$ with $\sigma = 5.79\,\text{mrad/s}$. In our confidence interval construction, this corresponds to an upper bound of $|\omega^{\mathcal{NE}}| < 1.9\sigma = 11\,\text{mrad/s}$ (90% C.L.). A comparison between three different 90% confidence interval constructions for small values of μ is shown in the right-hand plot of Fig. D.0.2. The black dashed lines represent the central confidence band for the signed values (rather than the magnitude) of μ and x, where

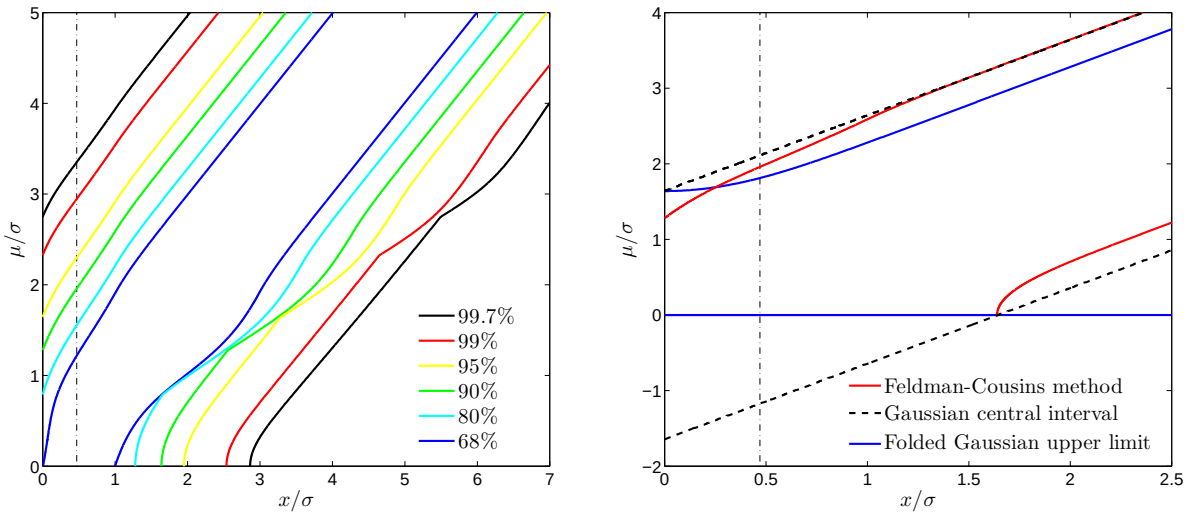

Figure D.0.2: Left: Feldman-Cousins confidence bands for a folded Gaussian distribution, constructed as described in the text, for a variety of confidence levels. Each pair of lines indicates the upper and lower bounds of the confidence band associated with each C.L. To the left of the x-intercepts, the lower bounds are zero. Confidence bands are plotted as a function of the possible measured central values x scaled by the standard deviation σ, and our result is plotted as a vertical dot-dashed line. The μ-value of the point at which our result line intersects with each of the colored lines gives the upper limit of our measurement at different C.L.'s. Right: Comparison between 90% confidence intervals computed using three different methods, described in the text. Confidence bands are plotted as a function of the possible measured central values of a quantity x scaled by the standard deviation σ. Our result, $|\omega_{\text{meas}}^{\mathcal{N}E}|/\sigma = 0.46$, is plotted as a vertical dot-dashed line. The μ-values of the points at which our result line intersects the upper and lower line for each method give the upper and lower bounds of three possible 90% confidence intervals for our measurement. To avoid invalidating the confidence interval by flip-flopping, our result should be interpreted using the Feldman-Cousins method, which we chose before unblinding.

μ is the mean of a Gaussian probability distribution in x. The blue lines give an upper bound constructed by computing the the value of μ such that the cumulative distribution function for the folded Gaussian in Eq. D.1 is equal to 0.9 for each value of x. It should be noted that this upper bound is more conservative than a true classical 90% confidence band, as it overcovers for small values of μ (e.g., if the true value were $\mu_{\text{true}} < 1.64\sigma$, the confidence intervals of 100% of experimental results would include μ_{true}). We nevertheless include this construction for comparison because we believe that previous experiments have reported EDM upper bounds using this method [89, 103, 162]. These intervals have a valid

interpretation as Bayesian "credible intervals" conditioned on a uniform prior for μ [77]. Finally, the red lines represent the Feldman-Cousins approach described here, which unifies upper limits and two-sided intervals. For our measurement outcome, indicated by the vertical dot-dashed line, the Feldman-Cousins intervals yield a 7% larger eEDM limit than the folded Gaussian upper bound would have produced.

References

[1] Georges Aad, T Abajyan, B Abbott, J Abdallah, S Abdel Khalek, AA Abdelalim, O Abdinov, R Aben, B Abi, M Abolins, et al. 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):1–29, 2012.

[2] Asmaa Abada and Takashi Toma. Electric dipole moments of charged leptons with sterile fermions. *Journal of High Energy Physics*, 2016(2):174, 2016. ISSN 1029-8479. doi: 10.1007/JHEP02(2016)174. URL http://dx.doi.org/10.1007/JHEP02(2016)174.

[3] Asmaa Abada and Takashi Toma. Electron electric dipole moment in inverse seesaw models. *Journal of High Energy Physics*, 2016(8):79, 2016. ISSN 1029-8479. doi: 10.1007/JHEP08(2016)079. URL http://dx.doi.org/10.1007/JHEP08(2016)079.

[4] K. Abdullah, C. Carlberg, E. D. Commins, H. Gould, and S. B. Ross. New experimental limit on the electron electric dipole moment. *Phys. Rev. Lett.*, 65:2347, 1990. doi: 10.1103/PhysRevLett.65.2347.

[5] S. Abel, S. Khalil, and O. Lebedev. EDM constraints in supersymmetric theories. *Nuclear Physics B*, 606(1–2):151 – 182, 2001. ISSN 0550-3213. doi: http://dx.doi.org/10.1016/S0550-3213(01)00233-4. URL http://www.sciencedirect.com/science/article/pii/S0550321301002334.

[6] R J Ackermann and E G Rauh. High-temperature properties of the thorium-oxygen system: A revision of the thermodynamic properties of ThO(g) and ThO_2(g). *High Temp. Sci.*, 5:463, 1973.

[7] R J Ackermann and E G Rauh. The preparation and characterization of the metastable monoxides of thorium and uranium. *Journal of Inorganic Nuclear Chemistry*, 35:3787–3794, 1973.

[8] R. J. Ackermann, E. G. Rauh, R J Thorn, and M. C. Cannon. A thermodynamic study of the thorium-oxygen system at high temperatures. *Journal of Physical Chemistry*, 67(4):762–769, April 1963. ISSN 0022-3654. doi: 10.1021/j100798a010. URL http://pubs.acs.org/cgi-bin/doilookup/?10.1021/j100798a010.

[9] R J Ackermann, E G Rauh, Chemistry Division, and Argonne National. The thermodynamics and vaporization of thorium, hafnium, and zirconium. *Journal of Chemical Thermodynamics*, 4:521–532, 1972.

[10] ACME Collaboration, J. Baron, W. C. Campbell, D. DeMille, J. M. Doyle, G. Gabrielse, Y. V. Gurevich, P. W. Hess, N. R. Hutzler, E. Kirilov, I. Kozyryev, B. R. O'Leary, C. D. Panda, M. F. Parsons, E. S. Petrik, B. Spaun, A. C. Vutha, and A. D. West. Order of magnitude smaller limit on the electric dipole moment of the electron. *Science*, 343:269–272, 2014. doi: 10.1126/science.1248213.

[11] 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. *arXiv*, October 2016. URL https://arxiv.org/abs/1612.09318.

[12] Tommi Alanne, Kimmo Kainulainen, Kimmo Tuominen, and Ville Vaskonen. Baryogenesis in the two doublet and inert singlet extension of the standard model. *Journal of Cosmology and Astroparticle Physics*, 2016(08):057, 2016. URL http://stacks.iop.org/1475-7516/2016/i=08/a=057.

[13] Vitaly Andreev. Polarimetry on the advanced cold molecule electron electric dipole moment experiment. Master's thesis, Technische Universität München, December 2016.

[14] Tatsumi Aoyama, Masashi Hayakawa, Toichiro Kinoshita, and Makiko Nio. Complete tenth-order qed contribution to the muon $g-2$. *Phys. Rev. Lett.*, 109:111808, September 2012. doi: 10.1103/PhysRevLett.109.111808. URL http://link.aps.org/doi/10.1103/PhysRevLett.109.111808.

[15] Tatsumi Aoyama, Masashi Hayakawa, Toichiro Kinoshita, and Makiko Nio. Quantum electrodynamics calculation of lepton anomalous magnetic moments: Numerical approach to the perturbation theory of QED. *Progress of Theoretical and Experimental Physics*, 01A(1):107, 2012. doi: 10.1093/ptep/pts030.

[16] Jonas B. Araujo, Rodolfo Casana, and Manoel M. Ferreira Jr. General cpt-even dimension-five nonminimal couplings between fermions and photons yielding {EDM} and {MDM}. *Physics Letters B*, 760:302 – 308, 2016. ISSN 0370-2693. doi: http://dx.doi.org/10.1016/j.physletb.2016.06.070. URL http://www.sciencedirect.com/science/article/pii/S0370269316303306.

[17] N. Arkani-Hamed, S. Dimopoulos, G.F. Giudice, and A. Romanino. Aspects of Split Supersymmetry. *Nuclear Physics B*, 709(1-2):3–46, March 2005. ISSN 05503213. doi: 10.1016/j.nuclphysb.2004.12.026. URL http://linkinghub.elsevier.com/retrieve/pii/S0550321304010259.

[18] C. Baker, D. D. Doyle, P. Geltenbort, K. Green, M. G. D. van der Grinten, P. G. Harris, P. Iaydjiev, S. N. Ivanov, D. J. R. May, J. M. Pendlebury, J. D. Richardson, D. Shiers, and K. F. Smith. Improved Experimental Limit on the Electric Dipole Moment of the Neutron. *Physical Review Letters*, 97(13):131801, September 2006.

ISSN 0031-9007. doi: 10.1103/PhysRevLett.97.131801. URL http://link.aps.org/doi/10.1103/PhysRevLett.97.131801.

[19] P. Balakrishna, B. P. Varma, T. S. Krishnan, T. R. R. Mohan, and P. Ramakrishnan. Low-temperature sintering of thoria. *Journal of Materials Science Letters*, 7(6):657–660, June 1988. ISSN 0261-8028. doi: 10.1007/BF01730326. URL http://www.springerlink.com/index/10.1007/BF01730326.

[20] C. Balázs, M. Carena, A. Menon, D. E. Morrissey, C. E. M. Wagner, and C. Bala. Supersymmetric origin of matter. *Physical Review D*, 71(7):1–16, April 2005. ISSN 1550-7998. doi: 10.1103/PhysRevD.71.075002. URL http://link.aps.org/doi/10.1103/PhysRevD.71.075002.

[21] S. M. Barr. A review of CP violation in atoms. *International Journal of Modern Physics A*, 08(02):209–236, January 1993. ISSN 0217-751X. doi: 10.1142/S0217751X93000096. URL http://www.worldscientific.com/doi/abs/10.1142/S0217751X93000096.

[22] J. F. Barry, E. S. Shuman, and D. Demille. A Bright, Slow Cryogenic Molecular Beam Source for Free Radicals. *Phys. Chem. Chem. Phys.*, 13(42):18936, June 2011. ISSN 1463-9084. doi: 10.1039/c1cp20335e. URL http://www.ncbi.nlm.nih.gov/pubmed/21706119.

[23] R Benz. Thorium-thorium dioxide phase equilibria. *Journal of Nuclear Materials*, 29:43–49, 1968.

[24] R Benz. Thorium-thorium dioxide phase equilibria. *Journal of Nuclear Materials*, 29 (1):43–49, January 1969. ISSN 00223115. doi: 10.1016/0022-3115(69)90125-1. URL http://linkinghub.elsevier.com/retrieve/pii/0022311569901251.

[25] R. Benz, C. G. Hoffman, and G. N. Rupert. Some phase equilibria in the thorium-nitrogen system. *Journal of the American Chemical Society*, 89(2):191–197, 1967.

[26] Joshua Berger, Matthew W. Cahill-Rowley, Diptimoy Ghosh, JoAnne L. Hewett, Ahmed Ismail, and Thomas G. Rizzo. CP-violating phenomenological MSSM. *Phys. Rev. D*, 93:035017, Feb 2016. doi: 10.1103/PhysRevD.93.035017. URL http://link.aps.org/doi/10.1103/PhysRevD.93.035017.

[27] K. Bergmann, H. Theuer, and B. W. Shore. Coherent population transfer among quantum states of atoms and molecules. *Rev. Mod. Phys.*, 70(3):1003, 1998.

[28] W. Bernreuther and M. Suzuki. The electric dipole moment of the electron. *Reviews of Modern Physics*, 63(2):313, 1991.

[29] H. G. Berry, G. Gabrielse, and A. E. Livingston. Measurement of the Stokes parameters of light. *Applied Optics*, 16(12):3200–5, December 1977. ISSN 0003-6935. doi: 10.1364/AO.16.003200. URL http://www.ncbi.nlm.nih.gov/pubmed/20174328.

[30] S. Bickman, P. Hamilton, Y. Jiang, and D. DeMille. Preparation and detection of states with simultaneous spin alignment and selectable molecular orientation in PbO. *Physical Review A*, 80:023418, August 2009. ISSN 1050-2947. doi: 10.1103/PhysRevA.80.023418. URL http://link.aps.org/doi/10.1103/PhysRevA.80.023418.

[31] Pavel A. Bolokhov, Maxim Pospelov, and Michael Romalis. Electric dipole moments as probes of CPT invariance. *Phys. Rev. D*, 78:057702, Sep 2008. doi: 10.1103/PhysRevD.78.057702. URL http://link.aps.org/doi/10.1103/PhysRevD.78.057702.

[32] Mary-Anne Bouchiat and Claude Bouchiat. Parity violation in atoms. *Reports on Progress in Physics*, 60:1351–1396, 1997.

[33] J. M. Brown and A. Carrington. *Rotational Spectroscopy of Diatomic Molecules*. Cambridge University Press, New York, 2003.

[34] Dmitry Budker, Derek F. Kimball, and David P. DeMille. *Atomic Physics: An Exploration Through Problems and Solutions*. Oxford University Press, New York, 2004.

[35] Wesley C. Campbell, Cheong Chan, David Demille, John M. Doyle, Gerald Gabrielse, Yulia V. Gurevich, Paul W. Hess, Nicholas R. Hutzler, Emil Kirilov, Brendon R. O'Leary, Elizabeth S. Petrik, Ben Spaun, and Amar C. Vutha. Advanced cold molecule electron EDM. *EPJ Web of Conferences*, 57:02004, August 2013. ISSN 2100-014X. doi: 10.1051/epjconf/20135702004. URL http://www.epj-conferences.org/10.1051/epjconf/20135702004.

[36] Laurent Canetti, Marco Drewes, and Mikhail Shaposhnikov. Matter and antimatter in the universe. *New Journal of Physics*, 14:095012, September 2012. URL http://iopscience.iop.org/article/10.1088/1367-2630/14/9/095012/meta.

[37] Marcela Carena, Eduardo Pontón, and José Zurita. Beyond the MSSM higgs bosons at the Tevatron and the LHC. *Phys. Rev. D*, 82:055025, September 2010. doi: 10.1103/PhysRevD.82.055025. URL http://link.aps.org/doi/10.1103/PhysRevD.82.055025.

[38] H. S. Carslaw and J. C. Jaeger. *Conduction of Heat in Solids*. The Clarendon Press, Oxford, 1947.

[39] S. Chatrchyan, V. Khachatryan, A.M. Sirunyan, A. Tumasyan, W. Adam, E. Aguilo, T. Bergauer, et al. Observation of a new boson at a mass of 125 gev with the {CMS} experiment at the {LHC}. *Physics Letters B*, 716(1):30 – 61, 2012. ISSN 0370-2693. doi: http://dx.doi.org/10.1016/j.physletb.2012.08.021. URL http://www.sciencedirect.com/science/article/pii/S0370269312008581.

[40] B. N. Chichkov, C. Momma, S. Nolte, F. von Alvensleben, and A. Tunnermann. Femtosecond, picosecond and nanosecond laser ablation of solids. *Appl. Phys. A*, 63:109–115, 1996.

[41] Y. T. Chien, V. Cirigliano, W. Dekens, J. de Vries, and E. Mereghetti. Direct and indirect constraints on CP-violating Higgs-quark and Higgs-gluon interactions. *Journal of High Energy Physics*, 2016(2):11, 2016. ISSN 1029-8479. doi: 10.1007/JHEP02(2016)011. URL http://dx.doi.org/10.1007/JHEP02(2016)011.

[42] Lily Childress and Mikhail Lukin. Lily notes - 285b mikhail lukin, October 2013. URL https://www.openrev.org/paper/lily-notes-285b-mikhail-lukin. Lecture notes for AMO I, published in Open Rev.

[43] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay. Evidence for the 2π decay of the K_2^0 meson. *Phys. Rev. Lett.*, 13:138–140, Jul 1964. doi: 10.1103/PhysRevLett.13.138.

[44] D. V. Chubukov and L. N. Labzowsky. $\mathcal{P},\mathcal{T}$-odd electron-nucleus interaction in atomic systems as an exchange by Higgs bosons. *Phys. Rev. A*, 93:062503, Jun 2016. doi: 10.1103/PhysRevA.93.062503. URL http://link.aps.org/doi/10.1103/PhysRevA.93.062503.

[45] V. Cirigliano, W. Dekens, J. de Vries, and E. Mereghetti. Is there room for CP violation in the top-Higgs sector? *Phys. Rev. D*, 94:016002, Jul 2016. doi: 10.1103/PhysRevD.94.016002. URL http://link.aps.org/doi/10.1103/PhysRevD.94.016002.

[46] James M. Cline. Baryogenesis. *ArXiv*, pages 1–64, November 2006. URL https://arxiv.org/abs/hep-ph/0609145v3.

[47] E. D. Commins and D. DeMille. The Electric Dipole Moment of the Electron. In B. L. Roberts and W. J. Marciano, editors, *Lepton Dipole Moments*, chapter 14, pages 519–581. World Scientific, 2010.

[48] Eugene Commins, Stephen Ross, David DeMille, and B. Regan. Improved experimental limit on the electric dipole moment of the electron. *Physical Review A*, 50(4): 2960–2977, October 1994. ISSN 1050-2947. doi: 10.1103/PhysRevA.50.2960. URL http://link.aps.org/doi/10.1103/PhysRevA.50.2960.

[49] Eugene D. Commins, J. D. Jackson, and David P. DeMille. The electric dipole moment of the electron: An intuitive explanation for the evasion of Schiff's theorem. *American Journal of Physics*, 75(6):532, 2007. ISSN 00029505. doi: 10.1119/1.2710486. URL http://link.aip.org/link/AJPIAS/v75/i6/p532/s1&Agg=doi.

[50] Cryomech. PT415 Cryorefrigerator Capacity Curve . Technical report, Cryomech, Inc., July 2007. URL http://www.cryomech.com/capacitycurve/PT415_cc.pdf.

[51] Cryomech. PT810 Cryorefrigerator Capacity Curve. Technical report, Cryomech, Inc., July 2007. URL http://www.cryomech.com/capacitycurve/PT810_cc.pdf.

[52] F. Csikor, Z. Fodor, and J. Heitger. End point of the hot electroweak phase transition. *Phys. Rev. Lett.*, 82:21–24, Jan 1999. doi: 10.1103/PhysRevLett.82.21. URL http://link.aps.org/doi/10.1103/PhysRevLett.82.21.

[53] C. E. Curtis. Thorium oxide in ceramic applications. In *Progress in Nuclear Energy, Series V, Metallurgy and Fuels*, volume 2, pages 223–236. Pergamon Press, London, 1959.

[54] Nuclear Fuel Cycle and Materials Section. Thorium fuel cycle—Potential benefits and challenges. Technical report, International Atomic Energy Agency, May 2005. URL http://www-pub.iaea.org/mtcd/publications/pdf/te_1450_web.pdf.

[55] A J Darnell and W. A McCollum. High Temperature Reactions of Thorium and Thoria and the Vapor Pressure of Thoria. Technical report, Atomics International, 1961.

[56] A. J. Darnell, W. A. McCollum, and T. A. Milne. Vapor pressure of thorium. *Journal of Physical Chemistry*, 64(3):341–346, March 1960. ISSN 0022-3654. doi: 10.1021/j100832a014. URL http://pubs.acs.org/cgi-bin/doilookup/?10.1021/j100832a014.

[57] D. DeMille, F. Bay, S. Bickman, D. Kawall, D. Krause, S. Maxwell, and L. Hunter. Investigation of PbO as a system for measuring the electric dipole moment of the electron. *Physical Review A*, 61(5):52507, April 2000. ISSN 1050-2947. doi: 10.1103/PhysRevA.61.052507. URL http://link.aps.org/doi/10.1103/PhysRevA.61.052507.

[58] D. DeMille, F. Bay, S Bickman, D. Kawall, L. Hunter, D. Krause, S. Maxwell, and K. Ulmer. Search for the electric dipole moment of the electron using metastable PbO. In *AIP Conference Proceedings*, volume 596, pages 72–83. AIP, 2001. doi: 10.1063/1.1426795. URL http://link.aip.org/link/?APC/596/72/1&Agg=doi.

[59] David DeMille. Diatomic molecules, a window onto fundamental physics. *Physics Today*, 68(12):34, December 2015.

[60] Malika Denis and Timo Fleig. In search of discrete symmetry violations beyond the standard model: Thorium monoxide reloaded. *The Journal of Chemical Physics*, 145 (21):214307, 2016. doi: 10.1063/1.4968597. URL http://dx.doi.org/10.1063/1.4968597.

[61] Malika Denis and Timo Fleig. Comment on "Theoretical study of thorium monoxide for the electron electric dipole moment search: electronic properties of $H(3)\Delta(1)$ in ThO". *ArXiv*, pages 1–5, May 2016. URL https://arxiv.org/abs/1605.03091.

[62] Michael Dine, Nathan Seiberg, and Scott Thomas. Higgs physics as a window beyond the MSSM. *Phys. Rev. D*, 76:095004, November 2007. doi: 10.1103/PhysRevD.76.095004. URL http://link.aps.org/doi/10.1103/PhysRevD.76.095004.

[63] A. D. Dolgov, S. I. Godunov, A. S. Rudenko, and I. I. Tkachev. Separated matter and antimatter domains with vanishing domain walls. *Journal of Cosmology and Astroparticle Physics*, 2015(10):027, 2015. URL http://stacks.iop.org/1475-7516/2015/i=10/a=027.

[64] V. A. Dzuba and V. V. Flambaum. Calculation of the (T,P)-odd electric dipole moment of thallium and cesium. *Physical Review A*, 80(6):062509, December 2009. ISSN 1050-2947. doi: 10.1103/PhysRevA.80.062509. URL http://link.aps.org/doi/10.1103/PhysRevA.80.062509.

[65] V. A. Dzuba, V. V. Flambaum, and C. Harabati. Relations between matrix elements of different weak interactions and interpretation of the parity-nonconserving and electron electric-dipole-moment measurements in atoms and molecules. *Phys. Rev. A*, 84:052108, Nov 2011. doi: 10.1103/PhysRevA.84.052108. URL http://link.aps.org/doi/10.1103/PhysRevA.84.052108.

[66] V. A. Dzuba, V. V. Flambaum, and C. Harabati. Erratum: Relations between matrix elements of different weak interactions and interpretation of the parity-nonconserving and electron electric-dipole-moment measurements in atoms and molecules [phys. rev. a **84**, 052108 (2011)]. *Phys. Rev. A*, 85:029901, Feb 2012. doi: 10.1103/PhysRevA.85.029901. URL http://link.aps.org/doi/10.1103/PhysRevA.85.029901.

[67] S. Eckel, P. Hamilton, E. Kirilov, H. W. Smith, and D. DeMille. Search for the electron electric dipole moment using Ω-doublet levels in PbO. *Physical Review A*, 87 (5):052130, May 2013. ISSN 1050-2947. doi: 10.1103/PhysRevA.87.052130. URL http://link.aps.org/doi/10.1103/PhysRevA.87.052130.

[68] G. Edvinsson and A. Lagerqvist. Rotational analysis of yellow and near infrared bands in ThO. *Physica Scripta*, 30(5):309–320, 1984.

[69] G Edvinsson and A Lagerqvist. A low-lying Ω=2 state in the ThO molecule. *J. Mol. Spectrosc.*, 113:93, 1985.

[70] G. Edvinsson, L.-E. Selin, and N. Aslund. On the band spectrum of ThO. *Ark. Phys.*, 30(22):283–319, 1965.

[71] D. Egorov. *Buffer-Gas Cooling of Diatomic Molecules*. PhD thesis, Harvard University, 2004. PhD Thesis, Harvard University.

[72] Sebastian A. R. Ellis and Gordon L. Kane. Theoretical prediction and impact of fundamental electric dipole moments. *Journal of High Energy Physics*, 2016(1):77, 2016. ISSN 1029-8479. doi: 10.1007/JHEP01(2016)077. URL http://dx.doi.org/10.1007/JHEP01(2016)077.

[73] Jonathan Engel, Michael 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*, 71:21–74, July 2013. ISSN 01466410. doi: 10.1016/j.ppnp.2013.03.003. URL http://linkinghub.elsevier.com/retrieve/pii/S0146641013000227.

[74] Daniel Farkas. *An Optical Reference and Frequency Comb for Improved Spectroscopy of Helium*. PhD thesis, Harvard University, 2006. (advisor: G. Gabrielse).

[75] Glennys R. Farrar and M. E. Shaposhnikov. Baryon asymmetry of the universe in the minimal standard model. *Phys. Rev. Lett.*, 70:2833–2836, May 1993. doi: 10.1103/PhysRevLett.70.2833. URL http://link.aps.org/doi/10.1103/PhysRevLett.70.2833.

[76] G. Feinberg. Effects of an electric dipole moment of the electron on the hydrogen energy levels. *Phys. Rev.*, 112:1637–1642, 1958. doi: 10.1103/PhysRev.112.1637.

[77] Gary J. Feldman and Robert D. Cousins. Unified approach to the classical statistical analysis of small signals. *Physical Review D*, 57(7):3873–3889, April 1998. ISSN 0556-2821. doi: 10.1103/PhysRevD.57.3873. URL http://link.aps.org/doi/10.1103/PhysRevD.57.3873.

[78] Jonathan L. Feng. Naturalness and the status of supersymmetry. *Annual Review of Nuclear and Particle Science*, 63:351–382, October 2013. doi: 10.1146/annurev-nucl-102010-130447. URL http://www.annualreviews.org/doi/abs/10.1146/annurev-nucl-102010-130447.

[79] T. Fleig and M. K. Nayak. Electron electric dipole moment and hyperfine interaction constants for ThO. *J. Mol. Spectrosc.*, 300:16–21, 2014.

[80] Norval Fortson, Patrick Sandars, and Stephen Barr. The search for a permanent electric dipole moment. *Physics Today*, 56:33–39, 2003.

[81] Takeshi Fukuyama and Koichiro Asahi. Filling the gaps between model predictions and their prerequisites in electric dipole moments. *International Journal of Modern Physics A*, 31(14n15):1650082, 2016. doi: 10.1142/S0217751X16500822. URL http://www.worldscientific.com/doi/abs/10.1142/S0217751X16500822.

[82] Takeshi Fukuyama, Koji Ichikawa, and Yukihiro Mimura. Revisiting fermion mass and mixing fits in the minimal SUSY $SO(10)$ GUT model. *Phys. Rev. D*, 94:075018, October 2016. doi: 10.1103/PhysRevD.94.075018. URL http://link.aps.org/doi/10.1103/PhysRevD.94.075018.

[83] Kaori Fuyuto, Junji Hisano, and Eibun Senaha. Toward verification of electroweak baryogenesis by electric dipole moments. *Physics Letters B*, 755:491 – 497, 2016. ISSN 0370-2693. doi: http://dx.doi.org/10.1016/j.physletb.2016.02.053. URL http://www.sciencedirect.com/science/article/pii/S0370269316001490.

[84] G. Gabrielse. The standard model's greatest triumph. *Physics Today*, 66:64–65, dec 2013. doi: 10.1063/PT.3.2223. URL http://dx.doi.org/10.1063/PT.3.2223.

[85] J. Goldemberg and Y. Torizuka. Upper limit on the electric dipole moment of the electron. *Phys. Rev.*, 129:2580–2581, 1963. doi: 10.1103/PhysRev.129.2580.

[86] Vasiliy Goncharov, Jiande Han, Leonid A. Kaledin, and Michael C. Heaven. Ionization energy measurements and electronic spectra for ThO. *J. Chem. Phys.*, 122(20):204311, May 2005. ISSN 0021-9606. doi: 10.1063/1.1898222. URL http://link.aip.org/link/JCPSA6/v122/i20/p204311/s1/html.

[87] B. Graner, Y. Chen, E. G. Lindahl, and B. R. Heckel. Reduced limit on the permanent electric dipole moment of ^{199}Hg. *Phys. Rev. Lett.*, 116:161601, April 2016. doi: 10.1103/PhysRevLett.116.161601. URL http://link.aps.org/doi/10.1103/PhysRevLett.116.161601.

[88] Daniel N Gresh, Kevin C Cossel, Yan Zhou, Jun Ye, and Eric A Cornell. Broadband velocity modulation spectroscopy of ThF+ for use in a measurement of the electron electric dipole moment. *Journal of Molecular Spectroscopy*, 319:1–9, 2016.

[89] W. Griffith, M. Swallows, T. Loftus, M. Romalis, B. Heckel, and E. Fortson. Improved limit on the permanent electric dipole moment of ^{199}Hg. *Physical Review Letters*, 102(10):101601, March 2009. ISSN 0031-9007. doi: 10.1103/PhysRevLett.102.101601. URL http://link.aps.org/doi/10.1103/PhysRevLett.102.101601.

[90] David J. Griffiths. *Introduction to Quantum Mechanics*. Pearson Education, New Delhi, 2005.

[91] Yulia V. Gurevich. *Preliminary Measurements for an Electron EDM Experiment in ThO*. PhD thesis, Harvard Univ., 2012. (advisor: G. Gabrielse).

[92] Alan Guth and Paul Steinhardt. The inflationary universe. In Paul Davies, editor, *The New Physics*, chapter 3, pages 34–60. Cambridge University Press, August 1992.

[93] George P. Halbfinger and Morris Kolodney. Activated sintering of ThO_2 and ThO_2-Y_2O_3 with NiO. *Journal of the American Ceramic Society*, 55(10):519–524, 1972.

[94] Hamamatsu. Hamamatsu Photomultiplier Tube R7600U Series. Technical report, Hamamatsu Photonics K. K., 2010. URL http://www.doylegroup.harvard.edu/files/internal/Hamamatsu_R7600U.pdf.

[95] Hamamatsu. Hamamatsu R7600-300 Spectral Response Characteristics. Technical report, Hamamatsu Photonics K. K., 2015. URL http://www.doylegroup.harvard.edu/wiki/images/8/85/R7600-300_cathode_spectral_response-2.pdf.

[96] D. Hanneke, S. Fogwell, and G. Gabrielse. New measurement of the electron magnetic moment and the fine structure constant. *Physical Review Letters*, 100(12): 120801, March 2008. ISSN 0031-9007. doi: 10.1103/PhysRevLett.100.120801. URL http://link.aps.org/doi/10.1103/PhysRevLett.100.120801.

[97] Y. Harada, Y. Baskin, and J. H. Handwerk. Calcination and sintering study of thoria. *Journal of the American Ceramic Society*, 45(6):253–257, 1962.

[98] Da He and N. N. Ekere. Effect of particle size ratio on the conducting percolation threshold of granular conductive–insulating composites. *Journal of Physics D: Applied Physics*, 37(13):1848, 2004. URL http://stacks.iop.org/0022-3727/37/i=13/a=019.

[99] G. Herzberg. *The Spectra and Structures of Simple Free Radicals*. Cornell University Press, Ithaca, 1971.

[100] P. W. Hess. *Improving the Limit on the Electron EDM: Data Acquisition and Systematics Studies*. PhD thesis, Harvard Univ., 2014. (advisor: G. Gabrielse).

[101] D. L. Hildenbrand. Mass spectrometric studies of gaseous ThO and ThO2. *The Journal of Chemical Physics*, 61(3):1232, 1974. ISSN 00219606. doi: 10.1063/1.1682000. URL http://link.aip.org/link/?JCP/61/1232/1&Agg=doi.

[102] K. P. Huber and G. Herzberg. *Molecular Spectra and Molecular Structure IV: Constants of Diatomic Molecules*. Van Nostrand Reinhold Company, 1979.

[103] 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(7348):493–6, May 2011. ISSN 1476-4687. doi: 10.1038/nature10104. URL http://www.ncbi.nlm.nih.gov/pubmed/21614077.

[104] Jonathan James Hudson. *Measuring the electric dipole moment of the electron with YbF molecules*. PhD thesis, University of Sussex, September 2001.

[105] N. Hutzler. *A New Limit on the Electron Electric Dipole Moment: Beam Production, Data Interpretation, and Systematics*. PhD thesis, Harvard University, 2014. (advisor: J. Doyle).

[106] Nicholas R. Hutzler, Maxwell F. Parsons, Yulia V. Gurevich, Paul W. Hess, Elizabeth Petrik, Ben Spaun, Amar C. Vutha, David DeMille, Gerald Gabrielse, and John M. Doyle. A cryogenic beam of refractory, chemically reactive molecules with expansion cooling. *Phys. Chem. Chem. Phys.*, 13:18976–18985, 2011. doi: 10.1039/C1CP20901A. URL http://dx.doi.org/10.1039/C1CP20901A.

[107] Nicholas R. Hutzler, Hsin-I Lu, and John M. Doyle. The buffer gas beam: an intense, cold, and slow source for atoms and molecules. *Chemical Reviews*, 112 (9):4803–27, September 2012. ISSN 1520-6890. doi: 10.1021/cr200362u. URL http://pubs.acs.org/doi/abs/10.1021/cr200362u.

[108] Satoru Inoue, Grigory Ovanesyan, and Michael J. Ramsey-Musolf. Two-step electroweak baryogenesis. *Phys. Rev. D*, 93:015013, Jan 2016. doi: 10.1103/PhysRevD.93.015013. URL http://link.aps.org/doi/10.1103/PhysRevD.93.015013.

[109] Koji Ishiwata, Kwang Sik Jeong, and Fuminobu Takahashi. Moduli-induced baryogenesis. *Journal of High Energy Physics*, 2014(2):62, 2014. ISSN 1029-8479. doi: 10.1007/JHEP02(2014)062. URL http://dx.doi.org/10.1007/JHEP02(2014)062.

[110] Minyuan Jiang, Ligong Bian, Weicong Huang, and Jing Shu. Impact of a complex singlet: Electroweak baryogenesis and dark matter. *Phys. Rev. D*, 93:065032, Mar 2016. doi: 10.1103/PhysRevD.93.065032. URL http://link.aps.org/doi/10.1103/PhysRevD.93.065032.

[111] J. R. Johnson and C. E. Curtis. Note on sintering of ThO_2. *Journal of the American Ceramic Society*, 37(12):611, 1954.

[112] P. J. Jorgensen and W. G. Schmidt. Final stage sintering of ThO2. *Journal of the American Ceramic Society*, 53(1):24–27, 1970.

[113] D M Kara, I J Smallman, J J Hudson, B E Sauer, M R Tarbutt, and E A Hinds. Measurement of the electron's electric dipole moment using YbF molecules: methods and data analysis. *New Journal of Physics*, 14(10):103051, October 2012. ISSN 1367-2630. doi: 10.1088/1367-2630/14/10/103051. URL http://stacks.iop.org/1367-2630/14/i=10/a=103051?key=crossref.935bac343ba2f6bf853baf0bc5df44bc.

[114] D. Kawall, F. Bay, S. Bickman, Y. Jiang, and D. DeMille. Progress towards measuring the electric dipole moment of the electron in metastable PbO. In *AIP Conference Proceedings*, volume 698, pages 192–195. AIP, 2004. doi: 10.1063/1.1664224. URL http://link.aip.org/link/?APC/698/192/1&Agg=doi.

[115] I. B. Khriplovich and S. K. Lamoreaux. *CP Violation without Strangeness*. Springer, Berlin, 1997.

[116] J. Kiggans. private communication, June 2007.

[117] J. Kim. *Buffer-Gas Loading and Magnetic Trapping of Atomic Europium*. PhD thesis, Harvard University, 1997.

[118] E. Kirilov, W. C. Campbell, J. M. Doyle, G. Gabrielse, Y. V. Gurevich, P. W. Hess, N. R. Hutzler, B. R. O'Leary, E. Petrik, B. Spaun, A. C. Vutha, and D. DeMille. Shot-noise-limited spin measurements in a pulsed molecular beam. *Physical Review A*, 88(1):013844, July 2013. ISSN 1050-2947. URL http://arxiv.org/abs/1305.2179http://link.aps.org/doi/10.1103/PhysRevA.88.013844.

[119] Ryuichiro Kitano, Hirosi Ooguri, and Yutaka Ookouchi. Supersymmetry breaking and gauge mediation. *Annual Review of Nuclear and Particle Science*, 60:491–511, July 2010. doi: 10.1146/annurev.nucl.012809.104540. URL http://www.annualreviews.org/doi/10.1146/annurev.nucl.012809.104540.

[120] Ryuichiro Kitano, Ryuji Motono, and Minoru Nagai. Mssm without free parameters. *Phys. Rev. D*, 94:115016, Dec 2016. doi: 10.1103/PhysRevD.94.115016. URL http://link.aps.org/doi/10.1103/PhysRevD.94.115016.

[121] D. L. Kokkin, T. C. Steimle, and D. DeMille. Branching ratios and radiative lifetimes of the U, L, and I states of thorium oxide. *Phys. Rev. A*, 90:062503, 2014.

[122] D. L. Kokkin, T. C. Steimle, and D. DeMille. Characterization of the $I(|\Omega| = 1)X^1\Sigma^+(0,0)$ band of thorium oxide. *Phys. Rev. A*, 91:042508, 2015.

[123] 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(18):L607, 1997. URL http://stacks.iop.org/0953-4075/30/i=18/a=003.

[124] M. G. Kozlov and L. N. Labzowsky. Parity violation effects in diatomics. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 28(10):1933–1961, May 1995. ISSN 0953-4075. doi: 10.1088/0953-4075/28/10/008. URL http://stacks.iop.org/0953-4075/28/i=10/a=008?key=crossref.2547c28c7d7ac4a7fa9f914888c84fc6.

[125] S. M. Lang and F. P. Knudsen. Some physical properties of high-density thorium dioxide. *Journal of the American Ceramic Society*, 39(12):415–525, 1956.

[126] Aaron E Leanhardt, John L Bohn, Huanqian Loh, Patrick Maletinsky, Edmund R Meyer, Laura C Sinclair, Russell P Stutz, and Eric A Cornell. On Measuring the Electron Electric Dipole Moment in Trapped Molecular Ions. *ArXiv*, page 39, August 2010. URL http://arxiv.org/abs/1008.2997.

[127] Yingchuan Li, Stefano Profumo, and Michael Ramsey-Musolf. Bino-driven electroweak baryogenesis with highly suppressed electric dipole moments. *Physics Letters B*, 673(1):95 – 100, 2009. ISSN 0370-2693. doi: http://dx.doi.org/10.1016/j.physletb.2009.02.004. URL http://www.sciencedirect.com/science/article/pii/S037026930900149X.

[128] Z. W. Liu and Hugh 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, 4 1992. doi: 10.1103/PhysRevA.45.R4210. URL http://link.aps.org/doi/10.1103/PhysRevA.45.R4210.

[129] Ye.P. Mamunya, V.V. Davydenko, P. Pissis, and E.V. Lebedev. Electrical and thermal conductivity of polymers filled with metal powders. *European Polymer Journal*, 38(9):1887 – 1897, 2002. ISSN 0014-3057. doi: http://dx.doi.org/10.1016/S0014-3057(02)00064-2. URL http://www.sciencedirect.com/science/article/pii/S0014305702000642.

[130] W. J. Marciano, A. Masiero, P. Paradisi, and M. Passera. Contributions of axion-like particles to lepton dipole moments. *Phys. Rev. D*, 94:115033, December 2016. doi: 10.1103/PhysRevD.94.115033. URL http://link.aps.org/doi/10.1103/PhysRevD.94.115033.

[131] Christel M. Marian, U. L. F. Wahlgren, O. D. D. Gropen, and P. Pyykko. Bonding and electronic structure in diatomic ThO: quasirelativistic effective core potential calculations. *J. Mol. Struct.*, 169:339, 1987.

[132] S. E. Maxwell, N. Brahms, R. DeCarvalho, D. R. Glenn, J. S. Helton, S. V. Nguyen, D. Patterson, J. Petricka, D. DeMille, and J. M. Doyle. High-Flux Beam Source for Cold, Slow Atoms or Molecules. *Physical Review Letters*, 95(17):173201, October 2005. ISSN 0031-9007. doi: 10.1103/PhysRevLett.95.173201. URL http://link.aps.org/doi/10.1103/PhysRevLett.95.173201.

[133] Eugen Merzbacher. *Quantum Mechanics*. John Wiley & Sons, Inc., 3rd edition, 1998.

[134] E. R. Meyer and J. L. Bohn. Prospects for an electron electric-dipole moment search in metastable ThO and ThF^+. *Phys. Rev. A*, 78(1):10502, 2008. doi: 10.1103/PhysRevA.78.010502.

[135] Edmund R. Meyer, John L. Bohn, and Michael P. Deskevich. Candidate molecular ions for an electron electric dipole moment experiment. *Phys. Rev. A*, 73(6):62108, June 2006. ISSN 1050-2947. doi: 10.1103/PhysRevA.73.062108. URL http://link.aps.org/doi/10.1103/PhysRevA.73.062108.

[136] N. S. Mosyagin, M. G. Kozlov, and A. V. Titov. Electric dipole moment of the electron in the YbF molecule. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 31(19):L763–L767, October 1998. ISSN 0953-4075. doi: 10.1088/0953-4075/31/19/002. URL http://stacks.iop.org/0953-4075/31/i=19/a=002?key=crossref.fc40d2593364d311c6a64905cd54e3ef.

[137] S. A. Murthy, D. Krause, Z. L. Li, and L. R. Hunter. New limits on the electron electric dipole moment from cesium. *Phys. Rev. Lett.*, 63:965–968, 8 1989. doi: 10.1103/PhysRevLett.63.965. URL http://link.aps.org/doi/10.1103/PhysRevLett.63.965.

[138] Yuichiro Nakai and Matthew Reece. Electric dipole moments in natural supersymmetry. *ArXiv*, December 2016. URL https://arxiv.org/abs/1612.08090.

[139] D. F. Nelson, A. A. Schupp, R. W. Pidd, and H. R. Crane. Search for an electric dipole moment of the electron. *Phys. Rev. Lett.*, 2:492–495, 1959. doi: 10.1103/PhysRevLett.2.492.

[140] K.-K. Ni, H. Loh, M. Grau, K. C. Cossel, J. Ye, and E. A. Cornell. State-specific detection of trapped HfF^+ by photodissociation. *Journal of Molecular Spectroscopy*, 300:12–15, June 2014. doi: 10.1016/j.jms.2014.02.001.

[141] Y. Nir. CP violation in and beyond the standard model. Lectures given in the XXVII SLAC Summer Institute on Particle Physics, July 1999. URL https://arxiv.org/abs/hep-ph/9911321.

[142] OECD. *Chemical Thermodynamics of Thorium*. OECD Publishing, 2008. doi: http://dx.doi.org/10.1787/9789264056688-en.

[143] Brendon 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*. PhD thesis, Yale University, December 2016.

[144] 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. *Phys. Rev. A*, 93:052110, May 2016. doi: 10.1103/PhysRevA.93.052110. URL http://link.aps.org/doi/10.1103/PhysRevA.93.052110.

[145] Paride Paradisi. On the interrelationship among leptonic g − 2, edms and lepton flavor violation. *EPJ Web of Conferences*, 118:01026, 2016. doi: 10.1051/epjconf/201611801026. URL https://doi.org/10.1051/epjconf/201611801026.

[146] David Patterson and John M. Doyle. Bright, guided molecular beam with hydrodynamic enhancement. *The Journal of Chemical Physics*, 126(15):154307, April 2007. ISSN 0021-9606. doi: 10.1063/1.2717178. URL http://www.ncbi.nlm.nih.gov/pubmed/17461626.

[147] David Patterson, Julia Rasmussen, and John M. Doyle. Intense atomic and molecular beams via neon buffer-gas cooling. *New Journal of Physics*, 11(5): 055018, May 2009. ISSN 1367-2630. doi: 10.1088/1367-2630/11/5/055018. URL http://stacks.iop.org/1367-2630/11/i=5/a=055018?key=crossref.ada50c0bc5df340d553560bba982f04c.

[148] J. Paulovic, T. Nakajima, and K. Hirao. Third-order Douglas–Kroll ab initio model potential for actinide elements. *J. Chem. Phys.*, 117:3597, 2002. doi: doi:10.1063/1.1483850.

[149] J. Paulovic, T. Nakajima, K. Hirao, R. Lindh, and P. A. Malmqvist. Relativistic and correlated calculations on the ground and excited states of ThO. *J. Chem. Phys.*, 119(2):798–805, 2003.

[150] Dennis V. Perepelitsa. Sakharov conditions for baryogenesis. Online, November 2008. URL phys.columbia.edu/~dvp/dvp-sakharov.pdf.

[151] 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. *Phys. Rev. A*, 89:062505, 6 2014. doi: 10.1103/PhysRevA.89.062505. URL http://link.aps.org/doi/10.1103/PhysRevA.89.062505.

[152] S. G. Porsev, M. S. Safronova, and M. G. Kozlov. Electric dipole moment enhancement factor of thallium. *Phys. Rev. Lett.*, 108:173001, Apr 2012. doi: 10.1103/PhysRevLett.108.173001. URL http://link.aps.org/doi/10.1103/PhysRevLett.108.173001.

[153] M. E. Pospelov and I. B. Khriplovich. Electric dipole moment of the w boson and the electron in the kobayashi-maskawa model. *Yad. Fiz.*, 53:1030–1033, 1991.

[154] Maxim Pospelov and Adam Ritz. Electric dipole moments as probes of new physics. *Ann. Phys.*, 318(1):119–169, July 2005. ISSN 00034916. doi: 10.1016/j.aop.2005.04.002. URL http://linkinghub.elsevier.com/retrieve/pii/S0003491605000539.

[155] Maxim Pospelov and Adam Ritz. Ckm benchmarks for electron electric dipole moment experiments. *Phys. Rev. D*, 89:056006, Mar 2014. doi: 10.1103/PhysRevD.89.056006. URL http://link.aps.org/doi/10.1103/PhysRevD.89.056006.

[156] M. J. Powell. Site percolation in randomly packed spheres. *Physical Review B*, 20 (10):4194–4198, November 1979. URL http://journals.aps.org/prb/pdf/10.1103/PhysRevB.20.4194.

[157] R. W. Powell, C. Y. Ho, and P. E. Liley. *Thermal Conductivity of Selected Materials*. U.S. Govt. Print. Off., Washington, D. C., 1966.

[158] E. M. Purcell and N. F. Ramsey. On the possibility of electric dipole moments for elementary particles and nuclei. *Phys. Rev.*, 78(6):807, June 1950. doi: 10.1103/PhysRev.78.807.

[159] M. Quack. Molecular parity violation and chirality: The asymmetry of life and the symmetry violation in physics. In Kiyoshi Nishikawa, Jean Maruani, Erkii J. Brändas, Gerardo Delgado-Barrio, and Piotr Piecuch, editors, *Quantum Systems in Chemistry and Physics, Progress in Theoretical Chemistry and Physics Volume 26*, chapter 3, pages 47–76. Springer Science, 2012. doi: 0.1007/978-94-007-5297-9. URL http://www.ir.ethz.ch/preprints/MQ354%20Quack%20Kanazawa%20Workshop.pdf.

[160] Norman Ramsey. *Molecular Beams*. Oxford University Press, 1985. ISBN 0198520212. URL http://books.google.com/books?id=T_7Hg08X7CMC&pgis=1.

[161] Norman F. Ramsey. A Molecular Beam Resonance Method with Separated Oscillating Fields. *Physical Review*, 78(6):695, 1950. doi: 10.1103/PhysRev.78.695.

[162] B. Regan, Eugene Commins, Christian Schmidt, and David DeMille. New limit on the electron electric dipole moment. *Physical Review Letters*, 88(7):18–21, February 2002. ISSN 0031-9007. doi: 10.1103/PhysRevLett.88.071805. URL http://link.aps.org/doi/10.1103/PhysRevLett.88.071805.

[163] Brian Christopher Regan. *A Search for Violation of Time-Reversal Symmetry in Atomic Thallium*. PhD thesis, University of California at Berkeley, 2001.

[164] K. F. Riley, M. P. Hobson, and S. J. Bence. *Mathematical methods for physics and engineering: a comprehensive guide*. Cambridge Univ. Press, Cambridge, 2006.

[165] C. Ronchi and J. P. Hiernaut. Experimental measurement of pre-melting and melting of thorium dioxide. *Journal of Alloys and Compounds*, 240(1):179 – 185, 1996. ISSN 0925-8388. doi: http://dx.doi.org/10.1016/0925-8388(96)02329-8. URL http://www.sciencedirect.com/science/article/pii/0925838896023298.

[166] A. D. Sakharov. Violation of CP Invariance, C Asymmetry, and Baryon Asymmetry of the Universe. *JETP Lett.*, 5:27–30, 1967.

[167] E. E. Salpeter. Some atomic effects of an electronic electric dipole moment. *Phys. Rev.*, 112:1642–1648, 1958. doi: 10.1103/PhysRev.112.1642.

[168] P. G. H. Sandars. The Electric Dipole Moment of an Atom. *Physics Letters*, 14(3): 194, February 1965. ISSN 00319163. doi: 10.1016/0031-9163(65)90583-4. URL http://dx.doi.org/doi:10.1016/0031-9163(65)90583-4.

[169] P. G. H. Sandars and E. Lipworth. Electric Dipole Moment of the Cesium Atom. A New Upper Limit to the Electric Dipole Moment of the Free Electron. *Physical Review Letters*, 13(24):718, 1964. doi: 10.1103/PhysRevLett.13.718.

[170] Eric Sather. The mystery of the matter asymmetry. *Beam Line*, pages 31–37, 1996.

[171] B. E. Sauer, J. J. Hudson, D. M. Kara, I. J. Smallman, M. R. Tarbutt, and E. A. Hinds. Prospects for the measurement of the electron electric dipole moment using YbF. *Physics Procedia*, 17:175–180, January 2011. ISSN 18753892. doi: 10.1016/j.phpro.2011.06.034. URL http://linkinghub.elsevier.com/retrieve/pii/S1875389211003658.

[172] Daniel V. Schroeder. *An Introduction to Thermal Physics*. Addison Wesley Longman, San Francisco, 2000. ISBN 0-201-38027-7.

[173] S M Skoff, R J Hendricks, C D J Sinclair, J J Hudson, D M Segal, B E Sauer, E A Hinds, and M R Tarbutt. Diffusion, thermalization, and optical pumping of YbF molecules in a cold buffer-gas cell. *Physical Review A*, 83(2):23418, February 2011. ISSN 1050-2947. doi: 10.1103/PhysRevA.83.023418. URL http://link.aps.org/doi/10.1103/PhysRevA.83.023418.

[174] L. V. Skripnikov. Combined 4-component and relativistic pseudopotential study of ThO for the electron electric dipole moment search. *J. Chem. Phys.*, 145:214301, 2016. doi: 10.1063/1.4968229. URL http://dx.doi.org/10.1063/1.4968229.

[175] 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. *J. Chem. Phys.*, 142:024301, 2015.

[176] L. V. Skripnikov and A. V. Titov. Reply to the comment on "theoretical study of thorium monoxide for the electron electric dipole moment search: electronic properties of $h(3)\delta(1)$ in ThO". *arXiv*, pages 1–4, October 2016. URL https://arxiv.org/abs/1610.02928.

[177] 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(22):221103, December 2013. ISSN 1089-7690. doi: 10.1063/1.4843955.

[178] Ian Joseph Smallman. *A New Measurement of the Electron Electric Dipole Moment Using Ytterbium Fluoride*. PhD thesis, Imperial College London, 2013.

[179] Benjamin N. Spaun, September 2012. Private communication.

[180] Benjamin Norman Spaun. *A Ten-Fold Improvement to the Limit of the Electron Electric Dipole Moment*. PhD thesis, Harvard University, 2014. (advisor: G. Gabrielse).

[181] SRS. Operating Manual and Programming Reference: Models RGA100, RGA200, and RGA300 Residual Gas Analyzer. Technical report, Stanford Research Systems, 2009. URL http://www.thinksrs.com/downloads/PDFs/Manuals/RGAm.pdf.

[182] Alan Stuart, J. Keith Ord, and Steven Arnold. *Classical Inference and the Linear Model*, volume 2A of *Kendall's Advanced Theory of Statistics*. Arnold, London, 6 edition, 1999.

[183] Russell Stutz and Eric Cornell. Search for the electron EDM using trapped molecular ions. *Bulletin of the American Physical Society*, 89:76, 2004.

[184] M. D. Swallows, T. H. Loftus, W. C. Griffith, B. R. Heckel, E. N. Fortson, and M. V. Romalis. Techniques used to search for a permanent electric dipole moment of the ^{199}Hg atom and the implications for *CP* violation. *Phys. Rev. A*, 87:012102, Jan 2013. doi: 10.1103/PhysRevA.87.012102. URL http://link.aps.org/doi/10.1103/PhysRevA.87.012102.

[185] M. R. Tarbutt, B. E. Sauer, J. J. Hudson, and E. a. Hinds. Design for a fountain of YbF molecules to measure the electron's electric dipole moment. *New Journal of Physics*, 15:0–17, 2013. ISSN 13672630. doi: 10.1088/1367-2630/15/5/053034.

[186] Y. S. Touloukian. *Thermophysical properties of matter*. IFI/Plenum, New York, 1970.

[187] C. H. Townes and A. L. Schawlow. *Microwave Spectroscopy*. Dover Publications, New York, 1975.

[188] A. C. Vutha. *A search for the electric dipole moment of the electron using thorium monoxide*. PhD thesis, Yale University, December 2011.

[189] 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*, 43(7):74007, April 2010. ISSN 0953-4075. doi: 10.1088/0953-4075/43/7/074007. URL http://stacks.iop.org/0953-4075/43/i=7/a=074007?key=crossref.114fcc8fa117ce235d58bc66e7812c5f.

[190] A. C. Vutha, B. Spaun, Y. V. Gurevich, N. R. Hutzler, E. Kirilov, J. M. Doyle, G. Gabrielse, and D. DeMille. Magnetic and electric dipole moments of the $H^3\Delta_1$ state in ThO. *Physical Review A*, 84(3):034502, September 2011. ISSN 1050-2947. doi: 10.1103/PhysRevA.84.034502. URL http://link.aps.org/doi/10.1103/PhysRevA.84.034502.

[191] Amar Vutha and David DeMille. Geometric phases without geometry. *ArXiv*, July 2009. URL http://arxiv.org/abs/0907.5116.

[192] Y. Watanabe and O. Matsuoka. All-electron Dirac–Fock–Roothaan calculations for the ThO molecule. *J. Chem. Phys.*, 107(9):3738–3739, 1997. doi: 10.1063/1.474487.

[193] 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, 1968. doi: 10.1103/PhysRevLett.21.1645.

[194] Adam D. West, Zack Lasner, David DeMille, Elizabeth P. West, Cristian D. Panda, John M. Doyle, Gerald Gabrielse, Adam Kryskow, and Corinne Mitchell. An underappreciated radiation hazard from high voltage electrodes in vacuum. *Health Physics*, 112:33–41, January 2017.

[195] Wikipedia. Percolation threshold. Online, 2017. URL https://en.wikipedia.org/wiki/Percolation_threshold. Visited on Mar. 22, 2017.

[196] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Phys. Rev.*, 105:1413–1415, Feb 1957. doi: 10.1103/PhysRev.105.1413. URL http://link.aps.org/doi/10.1103/PhysRev.105.1413.

[197] A. Zee. *Quantum Field Theory in a Nutshell*. Princeton University Press, 2003. ISBN 0691010196.

[198] Steven S. Zumdahl and Susan A Zumdahl. *Chemistry*. Houghton Mifflin Company, Boston and New York, 7$^{\text{th}}$ edition, 2007. ISBN 978-0-618-52844-8.

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.