Title

Direct Excitation of the 229Th Nuclear Isomeric Transition and Investigations of its Behavior in Solid-State Hosts

Author

Elwell, Ricky

Publication Date

2024

UNIVERSITY OF CALIFORNIA

Los Angeles

Direct Excitation of the Thorium-229 Nuclear Isomeric Transition and Investigations of its Behavior in Solid-State Hosts

A dissertation submitted in partial satisfaction

of the requirements for the degree

Doctor of Philosophy in Physics

by

Richard Elwell

2024

ABSTRACT OF THE DISSERTATION

Direct Excitation of the Thorium-229 Nuclear Isomeric Transition and Investigations of its Behavior in Solid-State Hosts

by

Richard Elwell

Doctor of Philosophy in Physics

University of California, Los Angeles, 2024

Professor Eric R. Hudson, Chair

The ^{229m}Th isomeric state holds much promise as the reference for a clock of incredibly high precision and accuracy. This dissertation discusses the first direct excitation of the ^{229}Th isomer in a $LiSrAlF_6$ environment, and the laser system that was used for excitation spectroscopy. In addition, we discuss two other efforts towards determination of the energy of the isomeric transition. The first involved the implantation of ^{229m}Th into a superconducting nanowire single-photon imager (SNSPI). The second was towards the first instance of laser-based conversion electron Mössbauer spectroscopy (CEMS) in a ThO_2 sample. The results of these efforts will be discussed, as well as possible improvements that would enable detection of the isomeric transition within these environments.

The dissertation of Richard Elwell is approved.

Michail Bachtis

Paul Hamilton

Wesley C. Campbell

Eric R. Hudson, Committee Chair

University of California, Los Angeles

2024

To my grandparents,

Richard Brans Elwell and Joan Barbara Elwell

福元貞俊と藤本艶子

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I extend my deepest gratitude to my teammates on the Thorium project. The lessons they've imparted are too numerous to list in full, so I'll share a few highlights. Justin Jeet, my first mentor on the VUV laser system, taught me the value of hard work, justice, and perseverance—even when it involves putting a Viton gasket on a sideways rotatable flange. Christian Schneider remains one of the most technically capable individuals I've ever met. From him, I learned to approach work with seriousness, scrutinize data critically, and appreciate the elegance of "Lil' Bluey." James Terhune has been a fantastic addition to the lab. I'm confident he will steer the ship well in the years ahead. Through him, I remembered how to appreciate the small things in physics, the importance of reassessing your ingrained habits, and how physical strength can really expand the phase space of experiments.

Above all, I want to thank my advisor, Eric Hudson. A truly exceptional mentor, Eric has taught me not just atomic physics but also profound lessons on life and the art of being a scientist—often delivered through his unique idioms. A few of my favorites include "pretty is not a parameter," "we're thinking about the phases of the moon here," and "emotions are sometimes a shortcut your body uses to help you understand data." Eric's leadership has also been an invaluable example. I sincerely hope to emulate his approach if I ever find myself in a similar role.

Beyond my lab, I am grateful to the entire Hudson group and the wider UCLA AMO community for their warmth, kindness, and support. I always got a lot from chatting with you all in the halls over the years. From the Hudson group, I thank Justin Christensen, David Hucul, Zach Wall, Sam Vizvary, Michael Bareian, Prateek Puri, Mike Mills, Grant Mitts, Clayton Ho, Josh Rabinowitz, Elizabeth West, Hao Wu, Guozhu Zhu, Guangming Lao, Haowen Zhu, Elijah Mossman, and Steven Diaz. From the Hamilton lab, I thank Sami Khamis, Randy Putnam, Adam West, Victoria Palmaccio, Leah Phillips, and Raju Manthena. From the Campbell group, I thank Tony Ramsford, Conrad Roman, Changling Zhao, Thomas Dellaert, Hassan Farhat, Will Liu, Patrick McMillin, Drew Parks, Nicole Halwani, and Amanda Younes. From the Leibrandt group, I thank Rick Mattish, Chao Feng, Asad Contractor, Stephen Randolph, Nick Lackmann, and Ashely Kwon. From the support staff of the department, I want to thank Cameron Irvin, Jules Berger, Kyle Villa, Owen Stevenson, Peter Yu, Matt Pulos, and Brad Stewart. Special thanks to Eddie Chang, who joined the lab around the same time as me and was a close comrade for many years—I hope you're doing well.

Outside UCLA, many collaborators made our projects possible. The cryogenic SNSPI project benefited immensely from the expertise of Sae Woo Nam and Galen O'Neil's groups at NIST, Boulder. Sae Woo Nam's passing this year was a profound loss to our community, and I extend my deepest condolences to his family. The C.E.M.S. experiment was a rewarding collaboration with Peter Thirolf's group at Ludwig-Maximilians-Universität München, Jun Ye's group at CU Boulder, and Christoph Düllman's group at Johannes Gutenberg-Universität Mainz. Our recent exploration of the isomer's behavior in a solid-state environment has been greatly enriched by collaborations with Anastassia Alexandrova's group at UCLA Chemistry and Andrei Derevianko's group at the University of Nevada, Reno.

Lastly, I thank my friends and family for their unwavering support over the years. Science is a human endeavor, made possible by the people at our back. While space limits me from naming everyone, I hope you know who you are by how I've treated you. To the friends I've made from elementary school to graduate school, to my family in America and Japan, to my teachers throughout the years, and to my grandparents watching from heaven—thank you, and I love you. Whether you're in Los Angeles or Berlin, Bellingham or New York, D.C. or Takasago, please know that I always think of you as I work and always strive to make you proud.

確かに、いつも皆のために頑張ります。

VITA

2014-2018	B.S. Physics and B.A. Mathematics, University of Virginia
2018-present	University of California Los Angeles

PUBLICATIONS

Elwell, R. and Schneider, Christian and Jeet, Justin and Terhune, J. E. S. and Morgan, H. W. T. and Alexandrova, A. N. and Tran Tan, H. B. and Derevianko, Andrei and Hudson, Eric R. "Laser Excitation of the ^{229}Th Nuclear Isomeric Transition in a Solid-State Host." *Physical Review Letters 133, 013201* (2024)

Zhang, C. and von der Wense, L., and Doyle, J. F. and Higgins, J. S. and Ooi, T. and Friebel, H. U., and Ye, J. and Elwell, R. and Terhune, J. E. S. and Morgan, H. W. T., and Alexandrova, A. N. and Tran Tan, H. B. and Derevianko, A. and Hudson, E. R. "^{229}ThF$_4$ thin films for solid-state nuclear clocks." *preprint available, arXiv:2410.01753 (2024)*

CHAPTER 1

Introduction

1.1 Atomic Clocks

Much of physics concerns itself with the dynamics of systems and their behavior over time. As such, the ability to quantitatively measure the time between events, and to do so accurately, is intrinsically tied the level of precision with which physical theories can be tested. The ability to construct a clock that "ticks" the same for everyone is necessary for experiments to be replicated. A clock, at its heart, is based on the counting of periodic phenomena, and so some kind of common periodic phenomenon should serve as the fundamental unit.

In the 17th century the period of a pendulum with specified arm length served as this standard, and improvements to its design and the advent of other electromechanical devices pushed advances in physics, engineering, navigation, and communications. However, by the end of the 19th century, physicists and natural philosophers were already looking for a more "fundamental" system to base a standard of time, free of manufacturing imperfections. Even though the atomic theory was far from its modern form, already individuals like Maxwell and Thomson (Lord Kelvin) envisaged the use of atomic transitions as time standards [TT79]:

> The recent discoveries due to the Kinetic theory of gases and to Spectrum analysis (especially when it is applied to the light of the heavenly bodies) indicate to us *natural standard* pieces of matter such as atoms of hydrogen, or sodium, ready made in infinite numbers, all absolutely alike in every physical property. The time of vibration of a sodium particle corresponding to any one of its modes of vibration, is known to be absolutely independent of its position in the universe, and it will probably

> remain the same so long as the particle itself exists. - (W. Thomson & P. G. Tait 1879)

It would take until the mid-20th century for this vision to be realized, but now the second in the International System of Units (SI) is defined based on the $F = 3 \Leftrightarrow F = 4$ hyperfine transition in ^{133}Cs.

It is worth briefly discussing how these atomic clocks operate. An atomic clock consists of of three main components: an oscillator with some frequency f (in the case of the Cs clock it is the microwave field), a counter to count cycles of the oscillator, and an atomic transition that responds at some frequency f_0 (in this case the ^{133}Cs transition). The elapsed time T is the number of cycles N of the oscillator times the reciprocal of the frequency

$$T = \frac{N}{f}. \tag{1.1}$$

However, if the oscillator runs free without feedback, there is no guarantee whether the measured time is accurate. This feedback can be introduced by locking the oscillator to the atomic transition. This feedback will allow for some frequency error δf, leading to a time error

$$\frac{\delta T}{T} = \frac{\delta f}{f}. \tag{1.2}$$

The goal of the clock-maker is to make the $\sigma = \delta f/f$, the fractional frequency instability, as small as possible. This instability can be reduced by using an atomic reference with high quality factor $Q = f_0/\Delta f$, where Δf is the transition linewidth. Beyond this, one can improve the signal-to-noise-ratio S, and increase the averaging time τ. Altogether, these yield [Boy07]

$$\sigma = \frac{\chi}{2\pi Q S}\sqrt{\frac{t_c}{\tau}}, \tag{1.3}$$

where t_c is the cycling time for clock operation and χ is a constant of order one. It is customary to report σ at $\tau = 1$ s (even though many experiments take longer than a second).

As is clear from our formula for σ, one can achieve a better clock by going to higher frequency transitions to improve Q. The turn of the 21st century saw researchers move from atomic transitions in the microwave towards the optical regime. Developments that enabled this were the maturation of coherent laser technology, and the development of the optical frequency comb to bridge the gap between the optical cycling frequency of the laser and the radio-frequency regime of counting electronics [POG13]. Today, optical atomic standards based on trapped-ions and trapped atoms in optical lattices boast fractional frequency instabilities at the 10^{-17}-10^{-18} level [BCH19, BKA22].

1.1.1 Why Not Nuclei?

If large transition frequencies are a fairly sure bet to increased clock performance, it seems natural to consider using nuclear transitions as the clock reference as they are thousands to millions of times higher frequency than atomic transitions. The issue comes down to the oscillator, one needs a coherent light source (such as a laser) in order to probe the transition for clock operation. Currently, coherent laser technology is based on atomic-shell and solid-state transitions, which are limited to the eV range. Despite the promise of “nuclear clocks” based on nuclear transitions, current technology cannot access most of the nuclear domain. We say most, because there is one exception which will form the bulk of this thesis.

1.2 Nuclear Structure

Fundamentally, the nucleus of an atom is made up of quarks and gluons governed by the laws of quantum chromodynamics (QCD). However, at the present moment computing the structure of nuclei is hopelessly intractable within the framework of QCD. Experiment has shown that at low energies QCD can be described by an effective field theory (EFT) of fermionic nucleons (predominantly protons and neutrons) that interact via the strong interaction mediated by bosonic mesons (and to a weaker extent the electromagnetic interaction) [HKK20]. Nuclei are thus bound states of nucleons. Even this EFT is intractable when it comes to computing the details of bound states, energy levels, collective motions, etc., of anything but the lightest nuclei, and so further simplifications are invoked to make predictions about nuclear structure.

Broadly, these simplified nuclear models come in two classes. The first type, *collective models*, assume that nucleons interact strongly (almost like molecules in a fluid) and are used to understand collective phenomena. Famous among these is the liquid-drop model of Bohr and Mottelson [BM98]. The second type, *independent particle models*, assume that the Pauli exclusion principle restricts collisions between nucleons, and so one may begin from a framework of weak inter-nucleon interactions. Both models are necessary to obtain a full understanding of the rich behavior nuclei can exhibit, however it is the latter class of models that will be the focus of this thesis.

1.2.1 Nuclear Shell Model

We begin by considering the Hamiltonian of a nucleus made of A nucleons

$$H = \sum_i^A T_i(\boldsymbol{r}_i) + V(\boldsymbol{r}_1, \ldots, \boldsymbol{r}_A), \tag{1.4}$$

where $T_i(\boldsymbol{r_i})$ is the kinetic energy operator and V is the potential function. To good approximation, we may consider V as being made up only of two-body nucleon-nucleon interactions, so we may write

$$H = \sum_i^A T_i(\boldsymbol{r}_i) + \sum_{ij} V_{ij}(\boldsymbol{r}_i, \boldsymbol{r}_j). \tag{1.5}$$

We then imagine that to an individual nucleon, its interactions with all of the other nucleons is seen as an effective "average" potential that is a function of only its position. Given that the nucleus has no preferred direction, this effective potential is assumed to only depend on the radial distance from the origin (i.e. it is a central potential). Motivated by this physical reasoning, we add and subtract a central potential to our Hamiltonian

$$\begin{aligned} H &= \sum_i^A T_i(\boldsymbol{r}_i) + \sum_i^A U(r_i) + \sum_{ij} V_{ij}(\boldsymbol{r}_i, \boldsymbol{r}_j) - \sum_i^A U(r_i), \\ &= \left[\sum_i^A T_i(\boldsymbol{r}_i) + U(r_i) \right] + H_{res} \end{aligned} \tag{1.6}$$

where H_{res} is known as the residual interaction. The potential U can be determined from the convolution of the two-body interactions with the nucleon density

$$U(r_i) = \sum_j \int V_{ij}(\boldsymbol{r}_i, \boldsymbol{r}_j)\rho(\boldsymbol{r}_j) = \sum_j \int \psi_j^*(\boldsymbol{r}_j) V_{ij}(\boldsymbol{r}_i, \boldsymbol{r}_j) \psi_j(\boldsymbol{r}_j). \tag{1.7}$$

In much the same manner as the multi-electron problem, both the potential and wavefunctions could be determined in a self-consistent manner by Hartree-Fock (HF) or density-functional theory (DFT) methods.

Instead of trying to determine U from an interative solution to the Schrödinger equation, much progress can be made if we instead choose a model potential for U. A common choice is that of the isotropic simple harmonic oscillator (SHO)

$$U(r_i) = \frac{1}{2} m_i \omega^2 r_i^2 \tag{1.8}$$

due to its analytic simplicity. Within this potential, the eigenenergies are given by $(N + 1/2)\hbar\omega = (2n_r + l - 1/2)\hbar\omega$, where N is the oscillator shell number, n_r is the number of radial nodes of the wavefunction (counting the one at infinity) and l is the angular momentum. The nucleons, being fermions, fill in the available eignestates, much like electrons in the central potential of an atom, hence the term *nuclear shell model*. While the large distance behavior of this potential is incorrect (the strength of the potential grows without limit at large r, while the strong and electromagnetic interactions fall off with sufficient distance), it is useful for describing low energy excitations of nuclear systems.

From experiment, it is known that the nuclei display large and sudden changes in binding energy at certain "magic numbers" of nucleons $[2, 8, 20, 28, 50, 82, 126, \dots]$. Within the harmonic oscillator model this is exhibited by the filling of all degenerate levels of an oscillator shell. The degeneracy of the first few levels are given by $[2, 8, 20, 40, 70, 112, \dots]$. While the S.H.O. model reproduces the first three magic numbers correctly, it fails for higher shells. One can either attempt to find a more accurate radial potential (such as a Woods-Saxon potential), or add modifications to the S.H.O. potential. To connect with our upcoming discussion of the Nilsson model, we will opt for the latter.

It is known that the nuclear force is short range on the scale of the radius of nuclei. Binding energy per nucleon data shows saturation on the scale of Li and Be nuclei [Cas01]. Therefore, especially in heavier nuclei (such as ^{229}Th) an individual nucleon will see only the force from its neighbor nucleons. Within the nuclear volume this nucleon will thus experience no net force, and so the potential will be relatively constant for radii less than the nuclear radius. The addition of an l^2 term to the potential can effectively "flatten" the potential seen by the nucleons at larger radii. Nucleons that have large angular momentum l orbits sample larger radii, and thus "see" the higher potential energy of the S.H.O. potential. It is precisely these orbits we can target and compensate with an l^2 term.

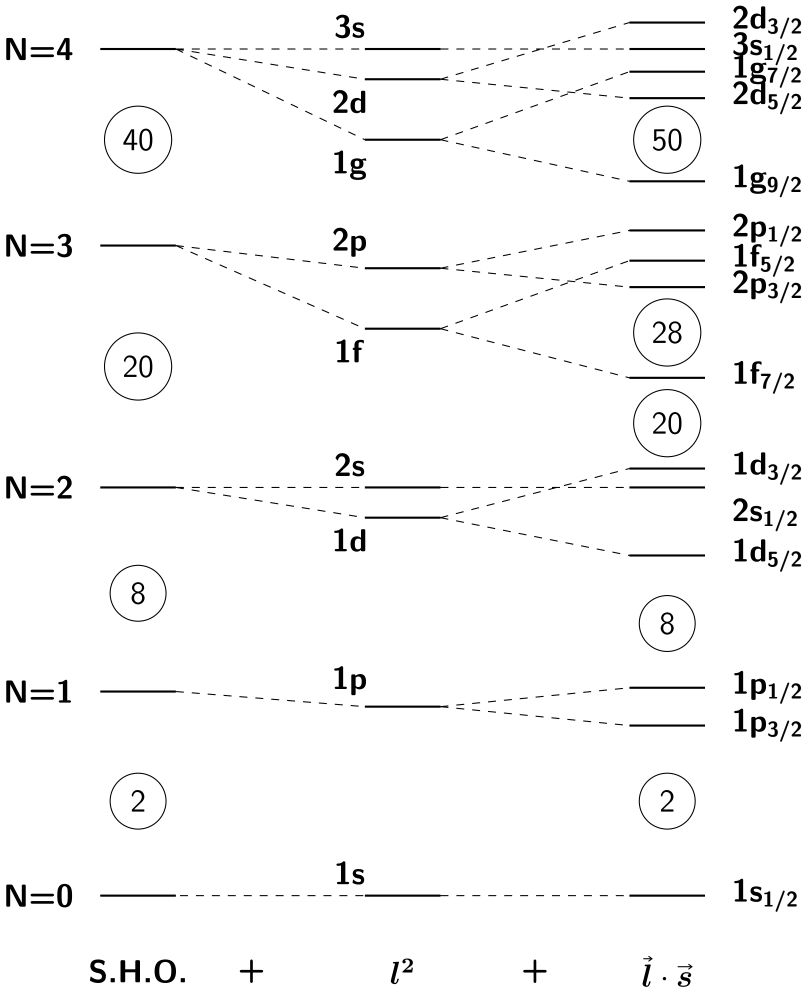

Figure 1.1: Simple harmonic oscillator eigenstates split by l^2 and spin-orbit interaction. Magic numbers shown in circles.

To fully reproduce the magic numbers, we also need to incorporate a spin-orbit interaction term of the form

$$V_{ls} = g_{ls}\frac{\partial U}{\partial r}\vec{\boldsymbol{l}}\cdot\vec{\boldsymbol{s}}, \tag{1.9}$$

where g_{ls} is an interaction constant. The nature of this term is different from the spin-orbit interaction for electrons in atomic systems. For electrons its origin is a relativistic effect that can be understood as an effective magnetic field seen by the electron in its orbit around the nucleus. While such a relativistic electromagnetic effect still plays a role for the nucleons, this term predominantly comes from the nature of the strong interaction. Deriving it from an underlying theory of the meson exchange between the nucleons is complex, and still an active area of research [BW77, CM20]. Still, we can argue for the form of Eq. 1.9 from experimental observations.

That the spin operator $\vec{\boldsymbol{s}}$ is involved is clear from considering the simplest two-body nuclear system, the deuteron (the nucleus of deuterium). The deuteron, made up of a proton and a neutron, is the only bound two-nucleon system, and possesses no stable excited states. It is a spin-1 system, meaning that that the proton and neutron have formed a spin triplet, and there is no observed spin-0 deuteron. Therefore, the strong interaction that binds the deuteron together must be spin-dependent. In addition, scattering experiments have demonstrated that when one sends a beam of nucleons that are unpolarized at a spin polarized nucleon target, the incident beam will scatter into spin-polarized channels of different outbound momenta [Kra88]. This implies an involvement of the nucleon momentum $\vec{\boldsymbol{p}}$ and its orientation with respect to $\vec{\boldsymbol{s}}$ in the nuclear force. The simplest scalar combination of $\vec{\boldsymbol{r}}$, $\vec{\boldsymbol{p}}$, and $\vec{\boldsymbol{s}}$ that respects parity and time-reversal symmetry is $(\vec{\boldsymbol{r}} \times \vec{\boldsymbol{p}}) \cdot \vec{\boldsymbol{s}} = \vec{\boldsymbol{l}} \cdot \vec{\boldsymbol{s}}$.

Thus our full "realistic" shell model Hamiltonian is given by

$$H = \left[\sum_{i}^{A} T_i(\boldsymbol{r}_i) + U(r_i) \right] + V_{l^2} l^2 + g_{ls} \frac{\partial U}{\partial r} \vec{\boldsymbol{l}} \cdot \vec{\boldsymbol{s}} + H_{res}. \tag{1.10}$$

It is this spin-orbit term that causes j, the total angular momentum, and its projection m to be good quantum numbers of the shell model (neglecting the residual interaction).

1.2.2 Nilsson Deformed Shell Model

The shell model as discussed until now is best applied to nuclei with valence nucleons occupying a single j shell. This typically only applies well to light nuclei and those heavy nuclei near closed shell configurations. As the number of valence nucleons is increased beyond a closed shell, the effect of the residual interaction leads to increasing scattering between single-particle states, and hence a great deal of configuration mixing [Cas01]. The effect of this is to cause unequal preference in the occupation of the magnetic substates m for a particular j-shell. It is this unequal filling of m

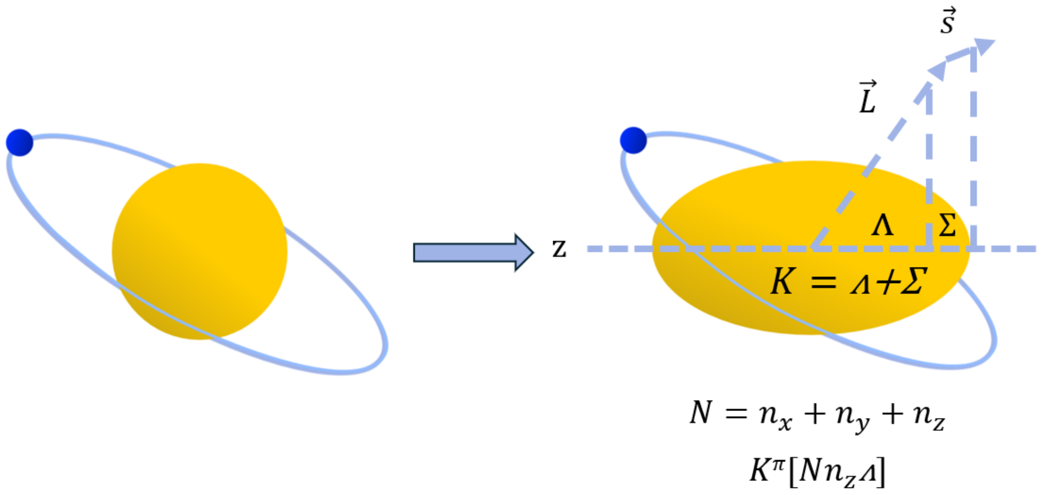

Figure 1.2: Diagram of the deformed nucleus and Nilsson quantum numbers.

states that cause the nucleus to assume a non-spherical shape, despite the fact that the underlying average potential is spherical. This is born out by experimental data which shows that nuclei exhibit quadrupole deformation in their ground states.

The Nilsson model takes these empiric facts and begins by assuming that a valence nucleon sees a potential that exhibits quadrupole deformation. Neglecting the residual interaction, Nilsson in his seminal paper [Nil55] began with a deformed oscillator Hamiltonian of the form

$$\begin{aligned} H_{Nils} = \frac{\boldsymbol{p}^2}{2m} + \frac{1}{2}m\left(\omega_\perp^2(x^2+y^2) + \omega_z^2 z^2\right) + C\vec{\boldsymbol{l}}\cdot\vec{\boldsymbol{s}} + Dl^2, \\ \omega_\perp^2 = \omega_0^2\left(1+\frac{2}{3}\delta\right), \\ \omega_z^2 = \omega_0^2\left(1-\frac{4}{3}\delta\right), \end{aligned} \tag{1.11}$$

where δ is the deformation parameter. The nucleon orbits of the Nilsson Hamiltonian are indexed by the “asympotic” Nilsson numbers

$$K^{\pi}\left[Nn_z\Lambda\right], \tag{1.12}$$

where K is the projection of the orbital and spin angular momentum on the symmetry z-axis, N is the oscillator shell number, n_z is the number of oscillator quanta along the z-axis, and Λ is the projection of the orbital angular momentum on the z-axis. In the limit of extreme deformation, the off-diagonal components of $\vec{\boldsymbol{l}}\cdot\vec{\boldsymbol{s}}$ and l^2 in

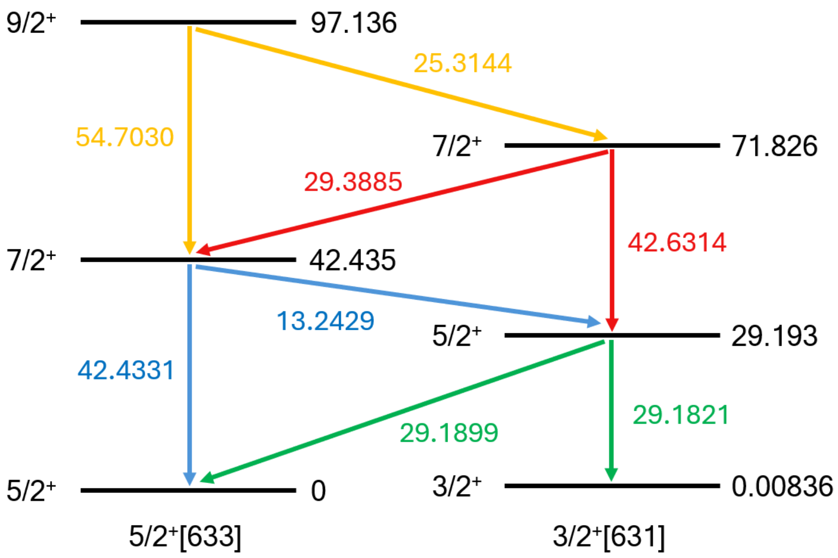

Figure 1.3: Nilsson level structure of ^{229}Th and corresponding γ-ray energies. All energies are given in keV. For comparison, energy levels are pulled from [Nat24], while the γ-ray transitions are from [SGH20].

the Nilsson basis can be neglected, and the Nilsson numbers become exact constants of the motion for the valence nucleon.

In addition to the motion of the valence nucleon, the deformed core can also undergo rotation about an axis orthogonal to the z-axis. Given the difference in time-scales between the orbit of the valence nucleon and the rotation of the core, the motions can be described by separate Hilbert spaces (in complete analogy to the Born-Oppenheimer approximation in molecular physics). To each possible Nilsson orbit (referred to as a “bandhead”) there is a “ladder” of energy states of quantized rotational motions of the core (referred to as a “Nilsson band”).

1.3 The ^{229}Th Isomeric Transition

The isomeric transition in ^{229}Th can be understood as a $\frac{5}{2}^+[633] \leftrightarrow \frac{3}{2}^+[631]$ transition in the Nilsson model. The transition is only several electron volts (8.355733554021(8) eV [ZOH24]), microscopic on the scale of nuclear transition energies which are typically keV-GeV range. Unlike other nuclear transitions, which can only be excited using incoherent gamma ray sources or synchotron facilities, the ^{229}Th transition can be driven by table-top atomic shell laser technology. This fact, in addition to the long lifetime of the transition, made ^{229}Th the intense focus of interest as a nuclear clock of high accuracy [TVL96, PT03, CRK12]. A further attractive feature of the isomer was that this clock transition would still be present upon of doping of ^{229}Th into a solid-state environment, as first discussed in [TVL96].

It is somewhat curious how such a low energy nuclear transition can exist. It can be thought of as being due to a cancellation of strong interaction and electromagnetic energy terms between the two states, $\varDelta E_{i\to f} = \varDelta E_{strong} + \varDelta E_{E\&M}$. While the individual terms are thought to change by keV-MeV scale amounts, the changes are of opposite sign leading to a near perfect cancellation. It is partially due to this cancellation that the isomer is the subject of intense theoretical interest, as it is thought that the isomeric transition will be particularly sensitive to variations of the fundamental constants governing the electromagnetic and strong forces (α, $\varLambda_{QCD}$, m_q, ...) with respect to one another [FW09, Saf19].

Another perspective by which to understand the existence of the lower energy isomer is to see it as an accidental degeneracy of two Nilsson orbits. To our knowledge, the best known value of the quadrupole moment of the $\frac{5}{2}^+[633]$ state is $Q_0 = 8.816(90)$ barns given by [BMJ88]. In the constant volume approximation, the relation between the Q_0 and the Nilsson deformation parameter δ is given by [LVH70]

$$Q_0 = \frac{4}{5} Z R_0^2 \delta \left(1 + \frac{2}{3}\delta + \frac{8}{9}\delta^2 + \frac{80}{81}\delta^3 + \cdots \right) \tag{1.13}$$

where Z is the nuclear charge and R_0=1.2$A^{1/3}$ fm is the nuclear radius for mass number A. In the case of ^{229}Th, this leads to a deformation value of $\delta \approx 0.194$. We may compare this to the Nilsson diagram in Fig. 1.4 and see this lies near a predicted crossing of two orbitals. Thus, the existence of the isomer can be seen as an accidental degeneracy due to the underlying deformation of the core nucleons.

1.3.1 Isomeric Decay Channels

The isomeric level in ^{229}Th ($\frac{3}{2}^+[631]$) will relax via different decay modes dependent upon the local electronic environment. The transition is of mixed magnetic dipole (M1) and electric quadrupole (E2) character [TSJ15, BMP18a]. If the nucleus was isolated the lifetime of the excited state would be several thousand seconds, and it

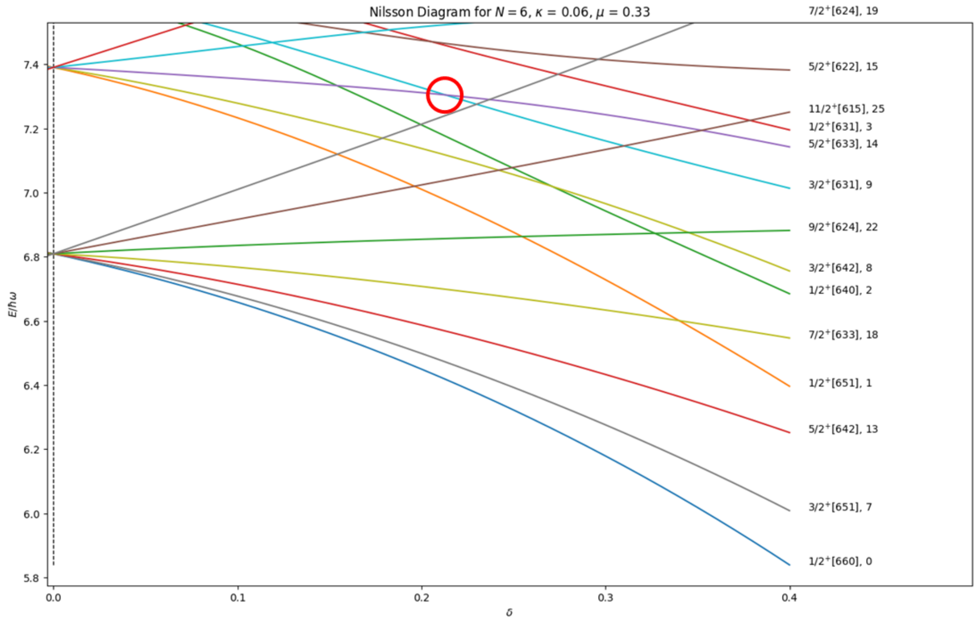

Figure 1.4: Energy of Nilsson orbitals with respect to δ for N=6 shells. Values used are $\kappa = 0.0635$ and $\mu = 0.325$, referred to as "Chasman" values in [BGQ08]. Diagram produced using code provided in [Wim20].

would undergo radiative decay by emitting a photon. For this pure radiative decay, the M1 component dominates completely. If, however, there are favorable energetics for the isomer to decay by coupling to an electron, the radiative decay channel will be near completely quenched. In these channels involving electronic coupling, the E2 component can begin to contribute significantly.

In the gas phase, the ionization energy of neutral Th is ∼6.3 eV, lower than the isomer energy. As such, in neutral ^{229}Th the isomer will decay via internal conversion (IC), a nuclear process by which an electron couples to the nuclear excited state and carries away the transition energy. It is known that the lifetime for the IC decay is on the order of microseconds, as predicted in [TSJ15] and confirmed in [SWT17]. For ionic states of Th, the ionization energy becomes >10 eV, and so the internal conversion channel is frozen out. For the 1^+-3^+ charge states, electron bridge (EB) processes may play a role in quenching the transition. During EB, an electron is promoted to a virtual state by taking the isomeric transition energy, and then relaxing to a real electronic state while simultaneously emitting a real photon. The EB rate is a sensitive function of the initial electron state and the density of final electron states, leading to calculated rates that range from the millisecond to second timescale [ST91, KT21]. Experimentally, the EB rate in Th^+ has been bounded to $\lesssim$10 ms [KT21]. When one reaches the 4^+ charge state, the first electronic excitation energy of the system is estimated be ∼13 eV [SS11], and all non-radiative quenching mechanisms are expected to be frozen out.

In the solid-state, the coupling of the isomer to the electronic environment can also lead to non-radiative quenching if the energetics are favorable. A necessary condition to observe the radiative decay of the isomer is that the band gap of the material the isomer sits in must be larger than the isomer energy. In ^{229}Thia, for example, the band gap is ∼5.8 eV [RL71], and the isomer is predicted to relax through an IC process where an electron is promoted to the conduction band (the details of this will be the subject of Ch. 4). Beyond the non-radiative quenching, the rapid attenuation of VUV light in materials makes observation of the radiated photons difficult as well. For most of the recent history of the transition, this band gap requirement has focused the efforts on the fluoride crystal family [RSD10]. Even if ^{229}Th is in a high band gap material, the possibility exists of the isomer coupling to localized defect states if it is the dopant within the environment. It may then be possible for an EB like process to occur, quenching the transition and exciting an electron to a defect state [NPB20, ESJ24]. The exact nature of the possible interactions and quenching of the isomeric transition in the solid-state environment will likely become the subject of investigation for years to come.

1.3.2 History of Measurements

In 1976, Kroger and Reich inferred the existence of the isomeric state in the gamma-ray spectra of ^{233}U α-decays [KR76]. The isomer's energy was first bounded to <100 eV, and subsequently measured to be 3.5(1.0) eV in 1994 [HR94]. This launched several unsuccessful efforts to detect a visible-UV photon from the isomeric transition. An excellent review of the history of this period can be found in the thesis of Erik Swanberg [Swa12]. The next major breakthrough for the field came in 2007, when Beck et. al. employed the National Aeronautics and Space Administration (NASA) x-ray spectrometer (XRS) to determine that the energy of the isomeric transition was 7.8(5) eV [BBB07, BWB09]. This new value made it clear that the transition lay somewhere squarely in the vacuum-ultraviolet (VUV) portion of the spectrum. While the idea of a solid-state nuclear clock had seemed attractive [PT03], the common stochiometric Th compounds (Th, ThO_2, ThF_4) were thought to have band gaps below the isomer energy.

The necessity of doping ^{229}Th into a VUV-transparent crystal as a way to search for the isomeric transition was first discussed in [HVL08, RHD10]. The strategy would be to achieve a high enough dopant concentration in a crystal to be able to observe the isomeric transition at a synchotron light source, and then to build a VUV laser system to further narrow down the transition. It was also theoretically estimated that even with broadening mechanisms present within the solid-state, a ^{229}Th-doped crystal could achieve a fractional frequency stability of 10^{-16} at 1 s averaging, which was sufficient to place strong constraints on the variation of fundamental constants [RHD10].

Several efforts began around the world to dope ^{229}Th into VUV-transparent crystals for the purpose of observing the radiative decay from the isomer. The first effort was carried out at the University of California, Los Angeles (UCLA) in cooperation with AC Materials to dope ^{229}Th into $LiSrAlF_6$, while a later effort was carried out at the Vienna University of Technology (TU Wien) to dope ^{229}Th into CaF_2. The history of both these efforts can be found in greater detail in the excellent theses of Justin Jeet and Kjeld Beeks [Jee18, Bee22]. Attempts were made by these groups to observe the radiative decay of the isomer following synchotron irradiation by looking for VUV photons, but no signal was detected [JSS15, SKS18]. The failure to observe excitation of the isomer in synchotrons prompted the production of pulsed VUV laser systems, first at UCLA and later at the Physikalisch-Technische Bundesanstalt (PTB) [Jee18, TZT23].

Concurrent to the crystalline efforts, a radioactive ion beam source at the Ludwig Maximilian University of Munich (LMU) had been modified to provide a source of ^{229}Th ions to measure the internal conversion decay of the isomer. We leave the details of these heroic experiments to the theses of Benedict Seiferle and Lars von der Wense [Sei19, Wen17]. Between 2014 and 2019, these experiments managed to directly observe the IC decay from the isomer, firmly establishing its existence (up until this point, all measurements of its existence had been through indirect gamma-ray measurements) [Wen17]. In addition, the isomer energy was measured to be 8.28(17) eV [SWB19], and its IC lifetime was determined to be 10(1) μs [SWT17], in agreement with the theoretical prediction of [TSJ15]. It was these measurements that launched a series of IC experiments that serve as the focus of Ch. 3 & 4 of this thesis.

Following these measurements further developments occurred in the study of the Nilsson bands of ^{229}Th. In 2019, a Japanese collaboration fired 29 keV X-rays from the BL19LXU beamline of the SPring-8 facility into a CaF$_2$ crystal doped with ^{229}Th. By doing so they were able to excite the cross-band transition from the $\frac{5}{2}^+[633]$ ground state to the first rotational excited state of the $\frac{3}{2}^+[631]$ band [Mas19]. In doing so, not only were they able to better measure the cross-band branching ratio, but this represented the first proven instance of the population of the isomer by a man-made light source. This x-ray pumping experiment reported a bound on the isomeric transition of [2.5,8.9] eV. Around this time, a NuClock collaboration also managed to redo the Beck et. al. measurement using a magnetic microcalorimeter to measure the x-rays from the Nilsson bands to higher precision. This yielded a value of 8.10(17) eV for the isomeric transition [SGH20].

The next flurry of activity came after a measurement was reported out of the CERN ISOLDE facility that had managed to directly observe fluorescence from the isomeric transition from ^{229}Th implanted into fluoride crystals via a ^{229}Ac beamline (more details are reserved for Ch. 5). This measurement was able to simultaneously narrow the transition uncertainty to 8.338(24) eV [KMA23], and definitively demonstrate that the isomeric transition was not quenched in a fluoride crystal environment. With this result, the ~ 5.5 THz region could be interrogated more precisly with a VUV laser system.

Not long after, 2024 saw the first observations of the direct excitation of the ^{229}Th isomer in solid-state environments via VUV pulsed laser systems. First was the result from the PTB-TU Wien collaboration in ^{229}Th:CaF$_2$ that reported a value of 2020409(7) GHz [TOZ24] with a lifetime of 630(15)s. This measurement was followed soon after by the UCLA result of 2020407.3(5)$_{\text{stat}}$(30)$_{\text{sys}}$ GHz with a lifetime of $568(13)_{\text{stat}}(20)_{\text{sys}}$ s in LiSrAlF_6 [ESJ24]. That summer then saw excitation of the isomeric transition in a CaF_2 crystal using a VUV frequency comb at JILA, which further narrowed the uncertainty in the transition to 2020407384335(2)kHz with a lifetime of 641(4) s at room temperature [ZOH24]. It is conjectured that the difference in lifetimes is primarily due to the Purcell effect (see Appendix 1.5.1 for details). As of the present moment, a similar VUV frequency comb measurement is being carried out in LiSrAlF_6, and the temperature dependence of the isomeric shift in the fluoride crystals is being investigated.

With the direct excitation of the isomer observed in ^{229}Th-doped crystals, attention was again given to the stochiometric Th compounds. Recent work had indicated, contrary to prevailing wisdom, that the band gap of ThF_4 was actually 10.2(2) eV [GEM19]. This means that the isomer can be radiatively excited and observed in bulk ThF_4. To confirm the Ye group at JILA was able to produce and characterize a thin film of $^{229}\text{ThF}_4$ on VUV-transparent substrates. The UCLA Hudson group was then able to observe radiative decay from the isomer following illumination by their VUV pulsed laser system [ZWD24]. The measured transition energy in the thin films was consistent with the doped crystal measurements, while the lifetime was significantly shorter ($\sim$150 s). Due to the large number of possible contaminants and poly-crystalline nature of the film, at the present moment it is difficult to determine whether the shortened lifetime is due to the Purcell effect or other possible quenching mechanisms.

Beyond the direct excitation experiments, two impressive measurements from Japan have added to our knowledge of the isomer. One was a continuation of the SPring-8 experiment led by Okayama University, which was able to observe the VUV photons from the decay of the isomer that was populated by x-ray pumping ^{229}Th doped into CaF_2 into the first rotationally excited state of the $\frac{3}{2}^+[631]$ band [HOB24]. From this measurement, they were able to measure the transition wavelength to $148.14(40)_{\text{stat}}(53)_{\text{sys}}$ nm using a series of transmission filters. The Okayama collaboration was also able measure the lifetime in CaF_2 to be $645(23)_{\text{stat}}(29)_{\text{sys}}$ s, consistent with the laser results. The second was a direct measurement of the isomer lifetime in ion-trap that was able to isolate $^{229m}\text{Th}^{3+}$ ions. Doing so, Yamaguchi et. al. were able to measure the lifetime of the isomer as 2000^{+900}_{-450} s [YSH24].

1.4 Light-Matter Interactions

The typical description of the interaction between a laser and the internal states of a material system are given by the semi-classical optical Bloch equations. We use the density operator $\hat{\rho}$ to model the time evolution of the occupation probabilities of the states. To simplify, we will assume a two-level nuclear system and begin with the Lindblad equation

$$\frac{\partial\hat{\rho}}{\partial t} = -\frac{i}{\hbar}\left[\hat{H}(t), \hat{\rho}(t)\right] + \mathcal{L}[\hat{\rho}] \tag{1.14}$$

Since we are dealing with a two level system, we have

$$\hat{\rho} = \rho_{gg}\left|g\right\rangle\!\left\langle g\right| + \rho_{eg}\left|e\right\rangle\!\left\langle g\right| + \rho_{ge}\left|g\right\rangle\!\left\langle e\right| + \rho_{ee}\left|e\right\rangle\!\left\langle e\right|, \tag{1.15}$$

where ρ_{gg} and ρ_{ee} denote the population in the ground and excited nuclear states, while $\rho_{ge} = \rho_{eg}^{*}$ are the coherence terms. The Hamiltonian $\hat{H}(t) = \hat{H}_N + \hat{H}_I(t)$ is made up of a static term which gives the energy of the two states (including shifts from the static external environment)

$$\hat{H}_N = \hbar\omega_g\left|g\right\rangle\!\left\langle g\right| + \hbar\omega_e\left|e\right\rangle\!\left\langle e\right|, \tag{1.16}$$

and a time-dependent term describing the interaction between the laser field and nucleus [Rin80]

$$\hat{H}_L(t) = \int d^3r\left(\hat{\rho}_N\varphi(\boldsymbol{r}, t) - \frac{1}{c}\hat{\boldsymbol{j}}_N\cdot\boldsymbol{A}(\boldsymbol{r}, t)\right), \tag{1.17}$$

where $\hat{\rho}_N$ and $\hat{\boldsymbol{j}}_N$ are the nuclear charge and current density operators, and φ and $\boldsymbol{A}$ are the scalar and vector potentials of the laser electromagnetic field. In the Coulomb gauge the laser field is described by $\varphi = 0$ [Zan12] and

$$\boldsymbol{A}(\boldsymbol{r}, t) = -\frac{\boldsymbol{E}_0}{2\omega_l}e^{i(\boldsymbol{k}_l\cdot\boldsymbol{r}-\omega_l t)} + c.c., \tag{1.18}$$

with electric field $\boldsymbol{E}_0$, angular frequency ω_l, and wave vector $\boldsymbol{k}_l$. This allows us to write Eq. 1.17 as

$$\hat{H}_L(t) = \hat{H}_L^0 e^{-i\omega_l t} + \hat{H}_L^{0\dagger}e^{i\omega_l t}, \tag{1.19}$$

where

$$\hat{H}_L^0 = \frac{1}{2\omega_l c} \int d^3r \hat{\boldsymbol{j}}_N \cdot \boldsymbol{E}_0 e^{i\boldsymbol{k}_l \cdot \boldsymbol{r}}. \tag{1.20}$$

We define the Rabi frequency Ω as

$$\Omega = \frac{2\langle e|\hat{H}_L^0|g\rangle}{\hbar}, \quad \tilde{\Omega} = \frac{2\langle e|\hat{H}_L^{0\dagger}|g\rangle}{\hbar}, \tag{1.21}$$

so that

$$\begin{aligned} \hat{H}_L(t) = &\frac{\hbar}{2}\left(\Omega e^{-i\omega_l t} + \tilde{\Omega} e^{i\omega_l t}\right)|e\rangle\langle g| \\ &+ \frac{\hbar}{2}\left(\tilde{\Omega}^* e^{-i\omega_l t} + \Omega^* e^{i\omega_l t}\right)|g\rangle\langle e|. \end{aligned} \tag{1.22}$$

For our Lindblad superoperator $\mathcal{L}[\hat{\rho}]$, we will only consider spontaneous emission of the excited state $|e\rangle$ into the ground state $|g\rangle$, which is described by the following expression

$$\mathcal{L}[\hat{\rho}] = \frac{\Gamma}{2}\left(2\sigma_-\hat{\rho}\sigma_+ - \hat{\rho}\sigma_+\sigma_- - \sigma_+\sigma_-\hat{\rho}\right), \tag{1.23}$$

where Γ is the total decay rate for the nuclear level (including non-radiative channels) and we have introduced the notation $|e\rangle\langle g| = \sigma_+ = \sigma_-^\dagger$. We will want to use the unitary operator

$$\begin{aligned} U(t) &= \exp\left(-i\omega_g t\,|g\rangle\langle g| - i(\omega_g + \omega_l)t\,|e\rangle\langle e|\right) \\ &= e^{-iPt}, \quad P = \left(\omega_g\,|g\rangle\langle g| + (\omega_g + \omega_l)\,|e\rangle\langle e|\right), \end{aligned} \tag{1.24}$$

in order to move into an interaction picture that removes the fast ω_l time dependencies (note that $[U, P] = 0$). The density matrix transforms as

$$\begin{aligned} \hat{\rho_I} = U\hat{\rho}U^\dagger &= \rho_{gg}\,|g\rangle\langle g| + \rho_{eg}e^{-i\omega_l t}\,|e\rangle\langle g| + \rho_{ge}e^{i\omega_l t}\,|g\rangle\langle e| + \rho_{ee}\,|e\rangle\langle e| \\ &= \rho_{gg}^I\,|g\rangle\langle g| + \rho_{eg}^I\,|e\rangle\langle g| + \rho_{ge}^I\,|g\rangle\langle e| + \rho_{ee}^I\,|e\rangle\langle e| \end{aligned} \tag{1.25}$$

We can then apply a time derivative to $U^\dagger\hat{\rho}_I U$ to obtain

$$\begin{aligned} \partial_t\left(U^\dagger\hat{\rho}_I U\right) &= (\partial_t U^\dagger)\hat{\rho}_I U + U^\dagger(\partial_t\hat{\rho}_I)U + U^\dagger\hat{\rho}_I(\partial_t U) \\ &= iPU^\dagger\hat{\rho}_I U + U^\dagger(\partial_t\hat{\rho}_I)U + U^\dagger\hat{\rho}_I(-iPU) \\ &= U^\dagger\left(\partial_t\hat{\rho}_I + i\left[P, \hat{\rho}_I\right]\right)U, \end{aligned} \tag{1.26}$$

where we used $[U, P] = 0$ to move around our operators. Some algebra shows that

$$[P, \hat{\rho}_I] = \rho_{eg}^I\omega_l\,|e\rangle\langle g| - \rho_{ge}^I\omega_l\,|g\rangle\langle e|. \tag{1.27}$$

We may now return to Eq. 1.14, applying U and $U^\dagger$ from the left and right to obtain

$$\begin{aligned}(\partial_t \hat{\rho}_I + i\,[P, \hat{\rho}_I]) &= -\frac{i}{\hbar} U \left[\hat{H}, \hat{\rho}\right] U^\dagger + U\mathcal{L}[\hat{\rho}]U^\dagger \\ \partial_t \hat{\rho}_I &= -\frac{i}{\hbar}\left(U\hat{H}U^\dagger U\hat{\rho}U^\dagger - U\hat{\rho}U^\dagger U\hat{H}U^\dagger\right) \\ &\quad - i\,[P, \hat{\rho}_I] + U\mathcal{L}[\hat{\rho}]U^\dagger \\ &= -\frac{i}{\hbar}\left[\hat{H}_I, \hat{\rho}_I\right] - i\,[P, \hat{\rho}_I] + U\mathcal{L}[\hat{\rho}]U^\dagger,\end{aligned} \tag{1.28}$$

where we define $\hat{H}_I = U\hat{H}U^\dagger$.

We must now compute $\hat{H}_I$, noting that U commutes with $\hat{H}_N$, so we need only consider

$$\begin{aligned}U\hat{H}_L(t)U^\dagger = &\frac{\hbar}{2}\left(\mathit{\Omega} + \tilde{\mathit{\Omega}} e^{i2\omega_l t}\right)|e\rangle\langle g| \\ &+ \frac{\hbar}{2}\left(\tilde{\mathit{\Omega}}^* e^{-i2\omega_l t} + \mathit{\Omega}^*\right)|g\rangle\langle e|\,.\end{aligned} \tag{1.29}$$

We will now apply the rotating wave approximation (RWA) and drop terms that rotate at $2\omega_l$. In matrix form, we now have

$$\hat{H}_I = \hbar \begin{pmatrix} \omega_g & \mathit{\Omega}^*/2 \\ \mathit{\Omega}/2 & \omega_e \end{pmatrix} \tag{1.30}$$

We now wish to express our Lindblad superoperator in this interaction picture, and so we find

$$\begin{aligned}U\mathcal{L}[\hat{\rho}]U^\dagger = &\frac{\mathit{\Gamma}}{2}(2U\sigma_- U^\dagger U\hat{\rho}U^\dagger U\sigma_+ U^\dagger - U\hat{\rho}U^\dagger U\sigma_+ U^\dagger U\sigma_- U^\dagger \\ &- U\sigma_+ U^\dagger U\sigma_- U^\dagger U\hat{\rho}U^\dagger) \\ \mathcal{L}[\hat{\rho}_I] = &\frac{\mathit{\Gamma}}{2}\left(2\sigma_- \hat{\rho}_I \sigma_+ - \hat{\rho}_I \sigma_+ \sigma_- - \sigma_+ \sigma_- \hat{\rho}_I\right),\end{aligned} \tag{1.31}$$

where we used the fact that $U\sigma_- U^\dagger = \sigma_- e^{-i\omega_l t}$ and $U\sigma_+ U^\dagger = \sigma_+ e^{i\omega_l t}$. We may throw all of the interaction picture operators into Eq. 1.28, and when the dust settles we find

$$\begin{aligned}\dot{\rho}^I_{gg} &= -\dot{\rho}^I_{ee} = \mathit{\Gamma}\rho^I_{ee} + \frac{i\mathit{\Omega}}{2}\left(\rho^I_{ge} - \rho^I_{eg}\right), \\ \dot{\rho}^I_{eg} &= \dot{\rho}^{I*}_{ge} = \frac{i\mathit{\Omega}}{2}\left(\rho^I_{ee} - \rho^I_{gg}\right) + \rho^I_{eg}(i\mathit{\Delta} - \frac{\mathit{\Gamma}}{2}),\end{aligned} \tag{1.32}$$

where $\mathit{\Delta} = ((\omega_e - \omega_g) - \omega_l)$ is the detuning from the transition and we have made $\mathit{\Omega}$ entirely real through the appropriate choice of an initial phase.

1.4.1 Sources of Finite Bandwidth

If the phase noise in the light field leads to significant fluctuations on the time-scale of the excitation, then Eq. 1.32 must be modified. We may model phase noise in the oscillator by a random variable $\varphi(t)$ such that the electric field can be expressed as

$$\boldsymbol{E}(\boldsymbol{r},t) = \boldsymbol{E}_0 e^{i(\boldsymbol{k}_l\cdot\boldsymbol{r}-\omega_l t-\varphi(t))} + c.c. \tag{1.33}$$

All of our previous derivation remains the same, except for the substitution $\omega_l \rightarrow \omega_l + \varphi(t)$, resulting in

$$\begin{aligned} \dot{\rho}^I_{gg} &= -\dot{\rho}^I_{ee} = \Gamma\rho^I_{ee} + \frac{i\Omega}{2}\left(\rho^I_{ge} - \rho^I_{eg}\right), \\ \dot{\rho}^I_{eg} &= \dot{\rho}^{I*}_{ge} = \frac{i\Omega}{2}\left(\rho^I_{ee} - \rho^I_{gg}\right) + \rho^I_{eg}(i(\Delta + \dot{\varphi}) - \frac{\Gamma}{2}), \end{aligned} \tag{1.34}$$

If we make the additional assumption that $\dot{\varphi}(t)$ is a white noise process

$$\langle\dot{\varphi}(t)\dot{\varphi}(t')\rangle = \Gamma_L\delta(t-t') \tag{1.35}$$

where Γ_L is the laser linewidth, then one can derive [WBS20] an effective version of Eq. 1.34 in the form

$$\begin{aligned} \dot{\rho}^I_{gg} &= -\dot{\rho}^I_{ee} = \Gamma\rho^I_{ee} + \frac{i\Omega}{2}\left(\rho^I_{ge} - \rho^I_{eg}\right), \\ \dot{\rho}^I_{eg} &= \dot{\rho}^{I*}_{ge} = \frac{i\Omega}{2}\left(\rho^I_{ee} - \rho^I_{gg}\right) + \rho^I_{eg}(i\Delta - \frac{\tilde{\Gamma}}{2}), \end{aligned} \tag{1.36}$$

where $\tilde{\Gamma} = (\Gamma + \Gamma_L)$. The full (if somewhat unwieldy) analytic solution to Eq. 1.36 is given by [NJ10], and at zero detuning ($\Delta = 0$) a simplified expression exists in the form of Torrey's solution [Tor49]. If we allow for detuning, but make the assumption that we are in the low-saturation limit ($\rho_{ee} \ll \rho_{gg} \approx 1$, $\Gamma_L \gg \Gamma$), then we may make the adiabatic approximation that the coherences relax instantaneously ($\dot{\rho}^I_{ge} = \dot{\rho}^I_{eg} = 0$), yielding

$$\rho^I_{eg} = \frac{i\Omega/2}{i\Delta - \Gamma_L/2}. \tag{1.37}$$

Plugging this into our equation for $\dot{\rho}^I_{ee}$ yields

$$\dot{\rho}^I_{ee} = -\Gamma\rho^I_{ee} + \frac{\Omega^2}{4}\frac{\Gamma_L}{\Delta^2 + \Gamma_L^2/4} = -\Gamma\rho^I_{ee} + \frac{\Omega^2}{\Gamma_L}\frac{1}{1 + 4\left(\frac{\Delta}{\Gamma_L}\right)^2}. \tag{1.38}$$

Integrating this equation with the boundary condition $\rho_{ee}(0) = 0$ yields

$$\rho_{ee}(t) = \frac{\Omega^2}{\Gamma_L}\frac{1}{1 + 4\left(\frac{\Delta}{\Gamma_L}\right)^2}\left(1 - e^{\Gamma t}\right) \tag{1.39}$$

Note the following relationship between the laser parameters and the Rabi frequency [Hil82]

$$\Omega^2 = \frac{2\pi c^2 I_l \Gamma_\gamma}{\hbar \omega_0^3}, \tag{1.40}$$

where Γ_γ is the radiative decay rate and $I_l = \frac{1}{2}\varepsilon_0 c \boldsymbol{E}_0^2$ is the laser intensity. We can insert this into the equation for $\rho_{ee}(t)$ to obtain

$$\begin{aligned} \rho_{ee}(t) &= \frac{2\pi c^2 I_l \Gamma_\gamma}{\hbar \omega_0^3 \Gamma \Gamma_L} \frac{1}{1 + 4\left(\frac{\Delta}{\Gamma_L}\right)^2} \left(1 - e^{\Gamma t}\right) \\ &= \frac{2\pi c^2 I_l}{\hbar \omega_0^3 (1+\alpha) \Gamma_L} \frac{1}{1 + 4\left(\frac{\Delta}{\Gamma_L}\right)^2} \left(1 - e^{\Gamma t}\right), \end{aligned} \tag{1.41}$$

where α is the branching ratio for non-radiative relaxation. We then can get the number of excited nuclei $N_e(t)$ by multiplying $\rho_{ee}(t)$ by $N_0 = nV_{ill}$, the total number of nuclei in the volume $V_{ill} = A_{laser} \times l$ illuminated by the laser with beam area A_{laser}. We note that near resonance $I_l \approx \hbar\omega_0 \varphi_p / A_{laser}$, where φ_p is the number of photons per second. Putting this all together we obtain

$$\begin{aligned} N_e(t) &= \frac{2\pi c^2 \hbar \omega_0 \varphi_p n A_{laser} l}{\hbar \omega_0^3 (1+\alpha) \Gamma_L A_{laser}} \frac{1}{1 + 4\left(\frac{\Delta}{\Gamma_L}\right)^2} \left(1 - e^{\Gamma t}\right) \\ &= \frac{\lambda_0^2}{2\pi} \frac{\varphi_p n l}{(1+\alpha)\Gamma_L} \frac{1}{1 + 4\left(\frac{\Delta}{\Gamma_L}\right)^2} \left(1 - e^{\Gamma t}\right). \end{aligned} \tag{1.42}$$

This is precisely the equation one would expect from the Einstein rate equation [Jee18], neglecting degeneracy factors. For the case of ^{229}Th we need to include the degeneracy factors of $g_{gnd} = 6$ and $g_{exc} = 4$ for the $\frac{5}{2}^+$ and $\frac{3}{2}^+$ states, respectively. We may re-label the laser frequency ω_l as ω and Γ by Γ_{tot} to emphasize that it is the total decay rate of radiative and non-radiative relaxation, and put in the degeneracy factors of the ground and excited state to obtain

$$N_e(t) = \frac{4}{6} \frac{\lambda_0^2}{2\pi} \frac{\varphi_p n_{Th} l}{(1+\alpha)\Gamma_L} \frac{1}{1 + 4(\frac{\omega - \omega_0}{\Gamma_{tot}})^2} \left(1 - e^{\Gamma_{tot} t}\right). \tag{1.43}$$

as the total ^{229}Th excited to the isomeric state by a laser system.

1.5 Appendix

1.5.1 Purcell Effect

We will consider the effect of embedding the isomer in a material medium on the radiative lifetime. Consider a macroscopic medium with constant permittivity and permeability (ε, μ) that may differ from the vacuum values (ε_0, μ_0). Let $\mathit{\Gamma}_m^{EL}$ denote the decay rate in a medium of EL-radiation, and $\mathit{\Gamma}_{vac}^{EL}$ the vacuum value. The definitions are analogous for $\mathit{\Gamma}_m^{ML}$ and $\mathit{\Gamma}_{vac}^{ML}$ for ML-radiation.

Within a material the zero-point fluctuations of $\boldsymbol{E}$ and $\boldsymbol{H}$ are modified from their vacuum values due to their scaling with (ε, μ). This, in conjunction with the scaling of the photon density of states with (ε, μ), leads to the following decay rate relationships [Tka01]

$$\begin{aligned} \mathit{\Gamma}_m^{EL} &= f_L(\varepsilon)^2 (\varepsilon/\varepsilon_0)^{L-1/2} (\mu/\mu_0)^{L+1/2} \mathit{\Gamma}_{vac}^{EL}, \\ \mathit{\Gamma}_m^{ML} &= f_L(\mu)^2 (\varepsilon/\varepsilon_0)^{L+1/2} (\mu/\mu_0)^{L+3/2} \mathit{\Gamma}_{vac}^{ML}, \end{aligned} \tag{1.44}$$

where the f_L functions connect the local field at the emitter to the macroscopic field in the medium. For the case of an M1 transition in a non-magnetic medium $f_L(\mu) = 1$ and $\mathit{\Gamma}_m^{M1} = (\varepsilon/\varepsilon_0)^{3/2} \mathit{\Gamma}_{vac}^{M1} = n^3 \mathit{\Gamma}_{vac}$, where $n = \sqrt{\varepsilon/\varepsilon_0}$ is the index of refraction. It is this final expression for M1 transitions that is commonly cited as the Purcell effect in the ^{229}Th literature.

CHAPTER 2

The VUV Laser System

2.1 Discussion of ALS results

The development of the ^{229}Th:LiSrAlF$_6$ crystals and their interrogation at the Advanced Light Source (ALS) is explained in great detail in the thesis of Justin Jeet [Jee18]. The main result was that the transition was not detected and a range of possible transition energies and lifetimes were thought to be excluded. With a bandwidth of nearly 45 THz [JSS15], it was thought that the transition lifetime could be longer than anticipated [TSJ15]. In order to probe these longer lifetimes (and consequently weaker signals), a VUV source with higher spectral density was constructed. What follows is a description of this system constructed at UCLA.

2.2 Description of the VUV Pulsed Laser System

The pulsed-VUV laser system starts with a Nd:YAG pump laser at 1064 nm in the infrared (IR) with a maximum pulse energy of about 1.3 J, a pulse width of ≈ 10 ns and a repetition rate of 30 Hz (Spectra Physics Quanta Ray Pro 230); see 2.1. The pump laser is injection seeded by an internal continous-wave (cw) laser reducing the output pulse frequency spectrum. Depending on the pump scheme (see 2.1), the pump laser is either converted in a third-harmonic generation (THG) stage to 355 nm (ultraviolet, UV) with a maximum pulse energy of ≈ 350 mJ and split into two beams. Alternatively, the IR beam is split at a roughly $1:2$ ratio to independently pump a second- (SHG) and a third-harmonic generation stage to generate a maximum of > 200 mJ at 532 nm (visible, VIS) and 355 nm each, respectively.

The SHG/THG light is used to pump two pulsed-dye lasers (PDLs; Lambda Physik FL 3002; with modifications and custom control electronics). We limit the

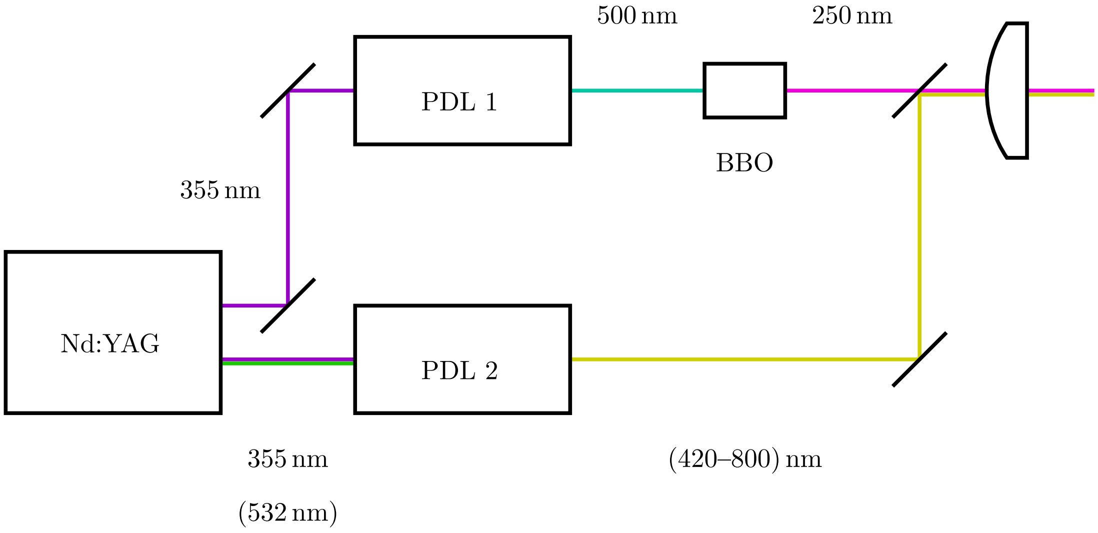

Figure 2.1: Schematic of the pump laser and pulsed-dye laser (PDL) layout.

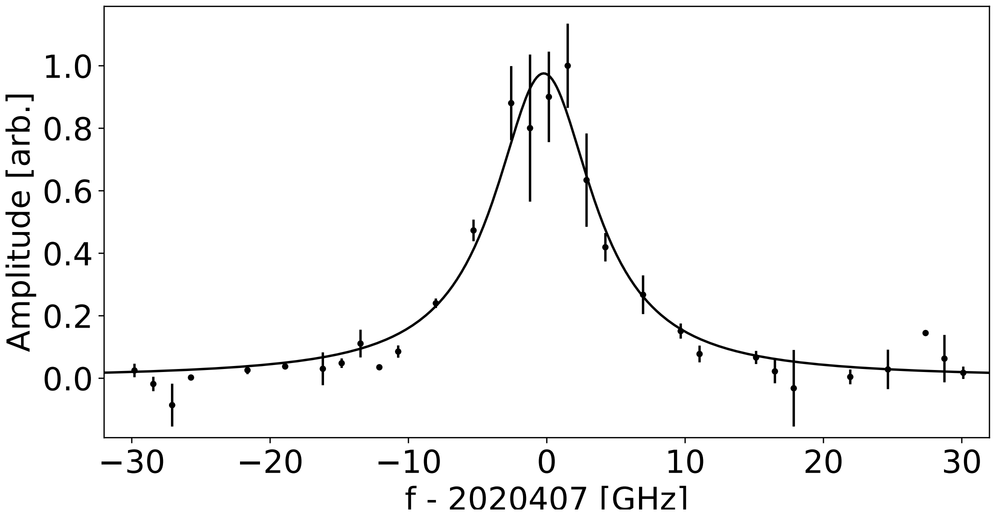

Figure 2.2: Linewidth scan of the VUV laser system over the ^{229}Th transition in a ^{229}ThF$_4$ film. Fitted Lorentzian FWHM is ≈ 9 GHz.

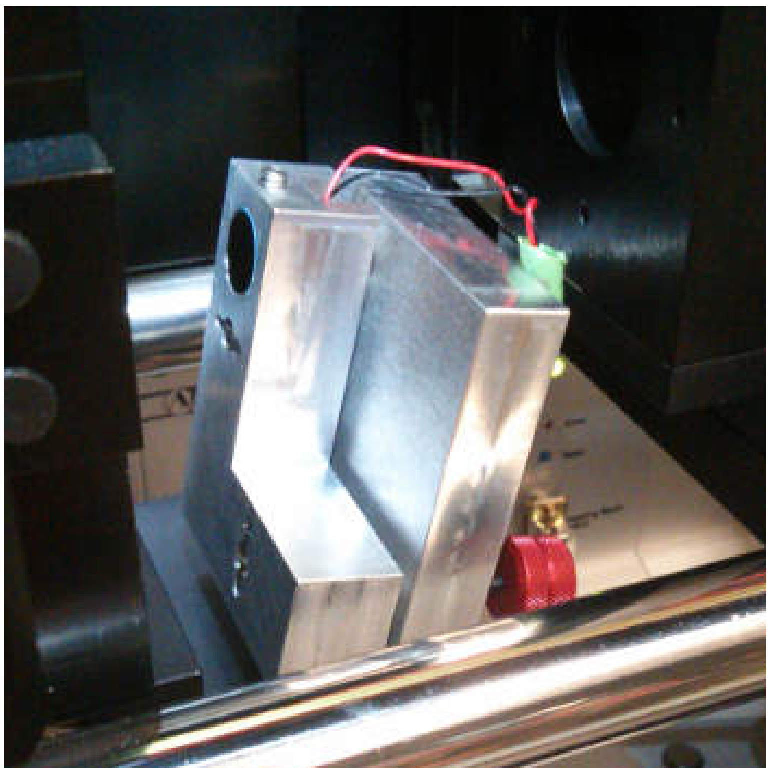

Figure 2.3: Custom oscillator-cavity endmirror mount for standard 1/2″-diameter mirrors with integrated piezoelectric adjuster inside PDL 2.

maximum pump pulse energy of PDL 1 to ≈ 135 mJ and the maximum pump pulse energy of PDL 2 to ≈ 90–125 mJ (depending on output wavelength) by detuning the input waveplates of the SHG/THG stages accordingly. First, this limit is dictated by staying below the damage thresholds of the PDL optics. Second, our scans require a constant pulse energy over extended periods of time during which the flash lamps of the Nd:YAG laser gradually degrade. The drop in IR pulse energy (and in turn SHG/THG pulse energy) can then be compensated by routinely (about weekly) adjusting these waveplates.

The PDLs were originally equipped for excimer-laser pumping and we heavily modified them for the Nd:YAG pump schemes above. We largely avoided the use of custom optics and modified the mounts inside the lasers instead to accept stock optics. The ≈ 10 % beamsplitters to split off a portion of the pump beam for oscillator and preamplifier, which require an anti-reflection (AR) coating on one side (instead of using an entirely uncoated window in the case of excimer lasers), were custom ordered. For the mirrors in the pump beams, we selected 2″-diameter optics throughout with laser-line coatings at either 355 nm or 532 nm and single angles at either 0 ° (which work sufficiently well under the required 10 °-reflection) or 45 °. The choice of coatings for a single laser line and single angle provides us with the highest damage thresholds, while being comparatively economic and widely available (typically by Edmund Optics, though some coatings have been discontinued recently). On the downside, the change of the pumping scheme of PDL 2 between 532 nm and 355 nm requires the change of the entire optics set inside the PDL and between pump laser and PDL.

Other notable modifications are the replacement of the feedback spoiler with a custom feedback-mirror assembly (consisting of a Newport U50-P mount, a standard 1/2″-diameter mirror with broadband dielectric coating, and several adapter parts); the replacement of the failing original electronics with home-developed microcontroller-based controllers with stepper-motor driver, sensor readback, and additional external trigger options for real-time control laser-grating steps; and the replacement of one of the oscillator-cavity endmirrors with a custom-built version that includes a piezo adjuster (see 2.3).

PDL 1 is operated at a fixed wavelength of $\approx$ 499.257 nm with a maximum pulse energy $\approx$ 40 mJ, which is frequency-doubled by an additional SHG stage (BBO, type I, 7.2 mm length) to $\lesssim$ 8 mJ in the UV at $\approx$ 250 nm. The SHG crystal is housed in a box on a motorized 2″-mirror mount outside of the PDL enclosure. The box itself has 3 mm tall input and output slits without windows to minimize any pulse-energy losses. Two small tubes inside the box blow dry oxygen over the input and output face of the crystal to reduce dust and humidity (BBO is slightly hygroscopic) and induce a self-healing effect of surface damage due to laser radiation. During long-term stable operation, we target $\gtrsim$ 4 mJ pulse energy in the UV, which can be maintained for about 12 h limited by the degradation of the laser dye (Coumarin 503). This light is near-resonant with a two-photon transition in xenon, see 2.4.

PDL 2 is operated with various laser dyes to span the entire wavelength range 420–800 nm with a target output pulse energy of around 10–15 mJ. Together with the UV beam, PDL 2 generates the VUV light by difference-frequency generation: $\omega_{\mathrm{VUV,-}} = 2\omega_{\mathrm{UV}} - \omega_{\mathrm{VIS}}$ (2.4).

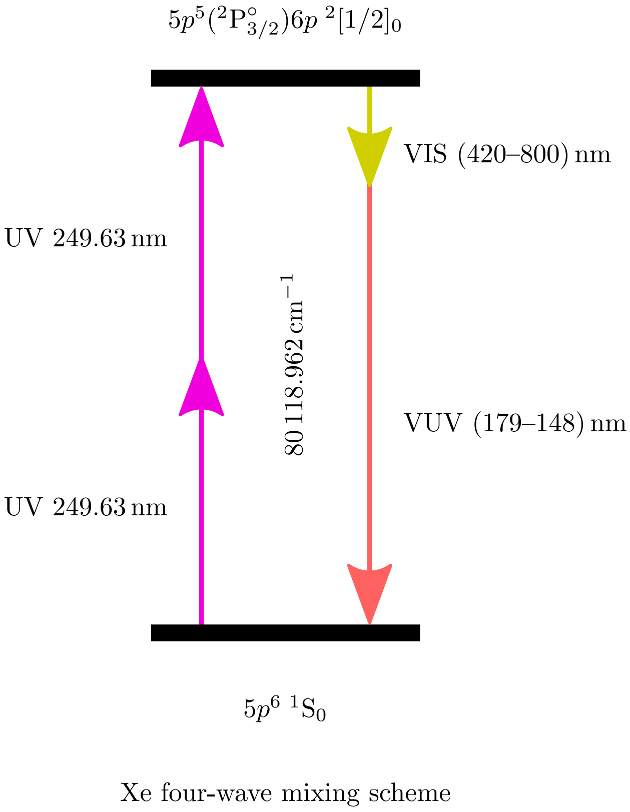

Figure 2.4: Relevant excerpt from the level scheme of xenon.

We mitigate large variations of the VUV pulse energy over the lifetime of the laser dyes by starting the Nd:YAG amplifier at 50–60 % of its maximum output power. After each data point, which takes about 1 h, we evaluate the average VUV pulse energy and increase the amplifier output as needed. As a change in the Nd:YAG pulse energy affects the subsequent THG and PDL 1 output pulse energies and possibly beam profiles, the critical phase matching of the 499 nm-SHG crystal needs to be optimized again. We deflect the 250 nm light with a motorized mirror flip mount (Thorlabs MFF101) onto a pyroelectric detector (Gentec Integra QE25LP-H-MB-QED-INT-D025) and optimize the phase-matching angle fully automatically. (It is debatable, if reducing the amplifer output power is the best strategy. Alternatively, the waveplates in front of the SHG/THG stages could again be used to control the output power, while the amplifier runs at a its maximum. For simplicity, we decided against this method.)

The output wavelength of a PDL can be controlled remotely via a stepper motor that adjusts the angle of the Littrow grating at ≈ 0.5 pm wavelength steps with a reproducibility of ≈ 1 pm, which is sufficient for wavelength scans during individual data points taking of order 1 h each. The absolute accuracy, however, is affected by the beam alignment and modifications of the PDLs causing up to several 10–100 pm shifts from the nominal wavelength. Additionally, day-to-day shifts and thermal drifts at a level of several picometers can be observed. To mitigate problems arising from these inaccuracies, we use actually measured wavelengths (first a Coherent WaveMaster and, later, a Bristol 871B wavemeter, both with ≈ 1 pm accuracy) whenever we initialize a PDL wavelength, perform larger wavelength jumps e.g. between data points, and record wavelengths during scans. Both PDLs are multiplexed onto the same wavemeter with a fiber-optical switch (Luminos S12-M4-N-ST-1-1-5H-L).

To verify both the accuracy and stability of the wavemeter, we perform spectroscopy of the two-photon resonance in xenon via a decay of the excited level in the IR; see 2.5. We use a separate spectrocopy cell for xenon out of standard Conflat (CF) parts and Swagelok VCR components for gas lines. A small fraction of the 250 nm beam from the backreflection off the input window of the xenon cell (few 10–100 µJ) is directed through a low-pressure xenon gas (few 10–100 µbar). The

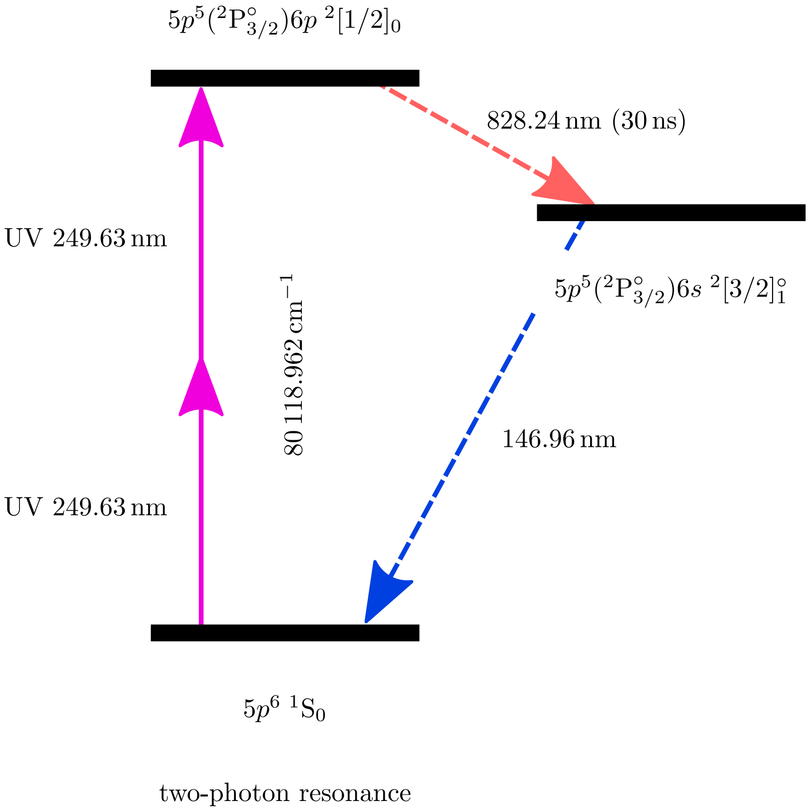

Figure 2.5: Excerpt from the level scheme of xenon including the spontaneous decay channel from the excited level. The decay occurs under emission of an infra-red photon which is detected in the spectroscopy of the two-photon resonance.

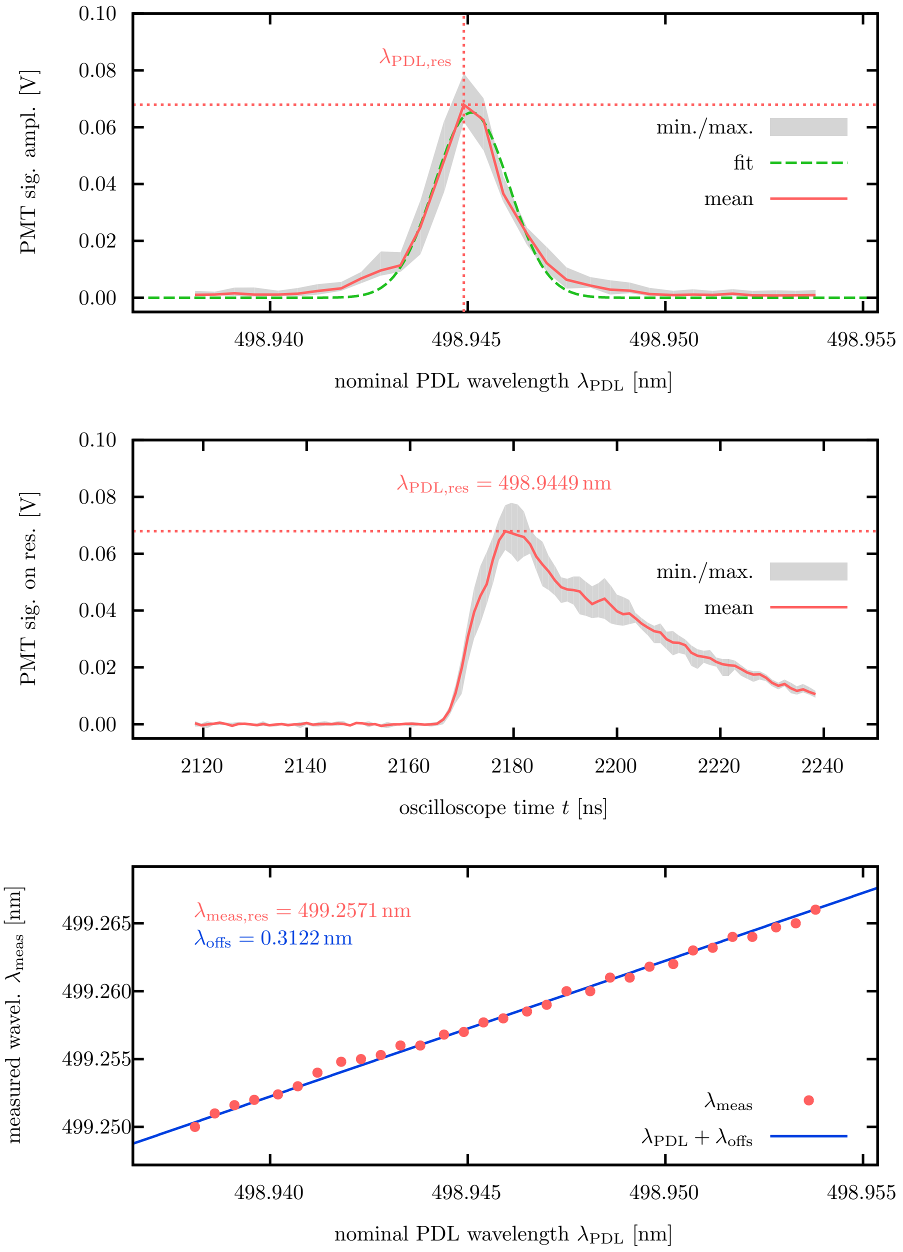

Figure 2.6: Measurement of the two-photon resonance in xenon with PDL 1. (top) Pulse height of PMT signal as a function of the nominal PDL wavelength according to its grating position λ_{PDL}. The mean is calculated from 5 PMT traces at each grating step. (center) The PMT signal on resonance as a function of time (again for a mean over 5 traces). (bottom) Relation between nominal and measured PDL wavelengths. The measured wavelength of the resonance is given together with the offset between nominal and measured wavelengths. This set of graphs is amongst those that is generated automatically for every data point to visually check the behaviour of the experiment.

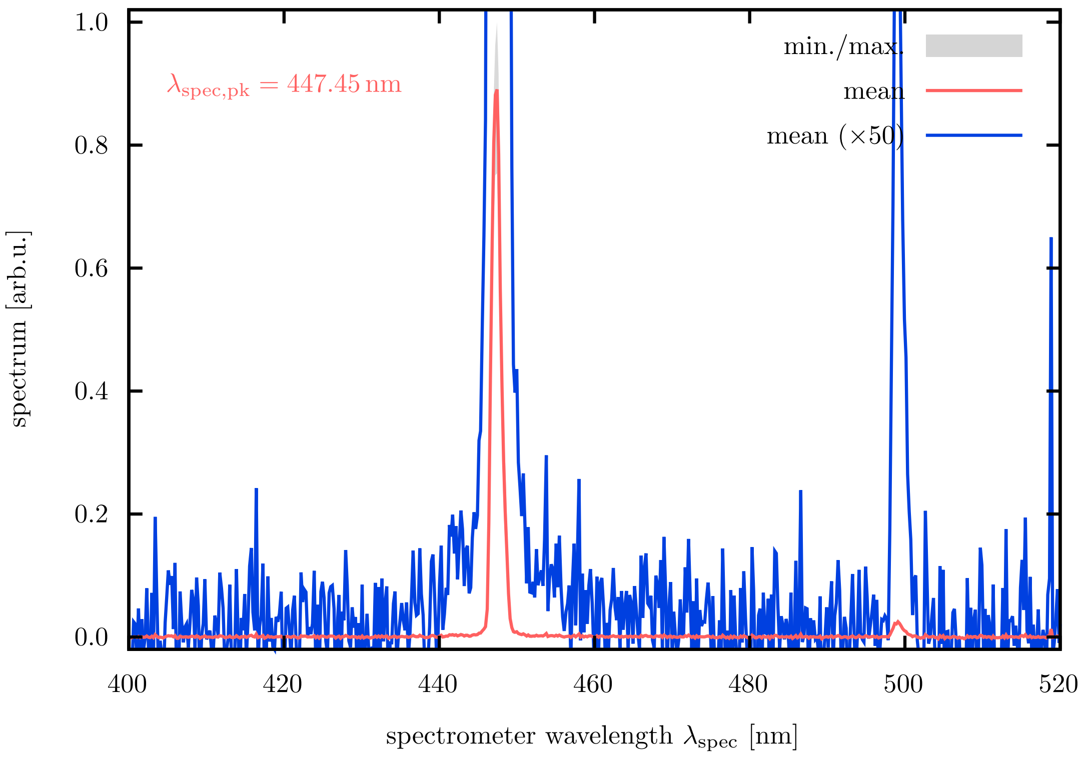

Figure 2.7: Example of a broadband spectrum averaged over 36 spectra to identify ASE in the output of PDL 2. A small portion of ASE can be identified besides the main peak at around 477 nm. This type of plot is also generated for each data point to verify proper PDL 2 performance.

spontaneous decay at ≈ 828 nm from the excited level in xenon is detected perpendicular to the beam direction with a PMT (Hamamatsu R928) in analog mode, which has an additional IR low-pass filter for background light reduction. The spectroscopy signal is recorded while stepping PDL 1 over the two-photon resonance and is repeated after every data point in our experiment. It has a typical full width at half maximum (FWHM) of a few nanometers and is fairly insensitive to power and pressure broadening effects within the aforementioned ranges; see 2.6 for a typical spectroscopy outcome. Over the course of several years of data recordings, we have not seen a single deviation of the wavelength reported by the wavemeter from the spectroscopy signal outside the datasheet accuracy of 1 pm.

An important characteristic of the output spectrum of the PDLs is their amplified-spontaneous emission (ASE), which depends critically on the exact tuning of the PDLs. The minimum ASE is typically close to but not exactly at the maximum output power. Minimization of ASE is given more weight in the fine-tuning of the PDLs than the output power to ensure a spectrally narrow VUV pulse: ASE can easily contribute to several 10 % of the pulse energy, if a PDL is not tuned carefully, and depends on the laser dye, the current grating position of the PDL with respect to the grating position at which the PDL was optimized, and other factors. While the output spectrum is impacted critically by the ASE, the pulse energy of the SHG process with the BBO crystal at 499 nm as well as the pulse energy of the VUV generation are barely affected by ASE. Hence, the presence of a large amount of ASE and, as a result, the spectral impurity in the VUV cannot be identified by solely measuring the pulse energies at 250 nm or in the VUV, respectively, as confirmed experimentally. As PDL 2 undergoes routine dye changes, reoptimizations for different wavelength ranges, and large wavelength scans, we continously monitor its output with a spectrometer (OceanOptics USB4000) (see 2.7). The fiber input of the spectrometer is fed with a backreflection from a window a few meters away from the PDL output to allow for continuous monitoring of the principal lasing signal and ASE.

The PDLs are typically quoted in their datasheets to have a linewidth of 2–6 GHz depending on the grating order of its Littrow configuration and the laser dye, respectively. This value, however, is rather the envelope of a series of much narrower lines spaced by the free-spectral range (FSR) of the PDL cavity [KLS92], which is $2\pi \times 450\,\mathrm{MHz}$ in our case. The actual spectrum of the PDLs is of utmost important for the search for the ^{229}Th nuclear isomeric transition, as the generated VUV spectrum will also contain “holes” in the spectrum. In order to analyze the full spectrum of individual PDL pulses, we constructed a Fabry–Pérot etalon setup shown in 2.8: A small portion of a PDL beam (e.g. the backreflection off an optic) is directed onto a series of diffusers followed by an adjustable iris and impinges onto the etalon under randomized angles. A ray picks up the additional phase per round trip of

$$\mathit{\Delta}\varphi = \frac{4\pi n_{\mathrm{g}} d}{\lambda}\sqrt{1-\left(\frac{n_{\mathrm{a}}}{n_{\mathrm{g}}}\sin\alpha\right)^{2}}, \tag{2.1}$$

where n_{g} is the index of refraction of the glass, n_{a} the index of refraction of air, d the thickness of the etalon, λ the wavelength of the light, and α the angle of incidence

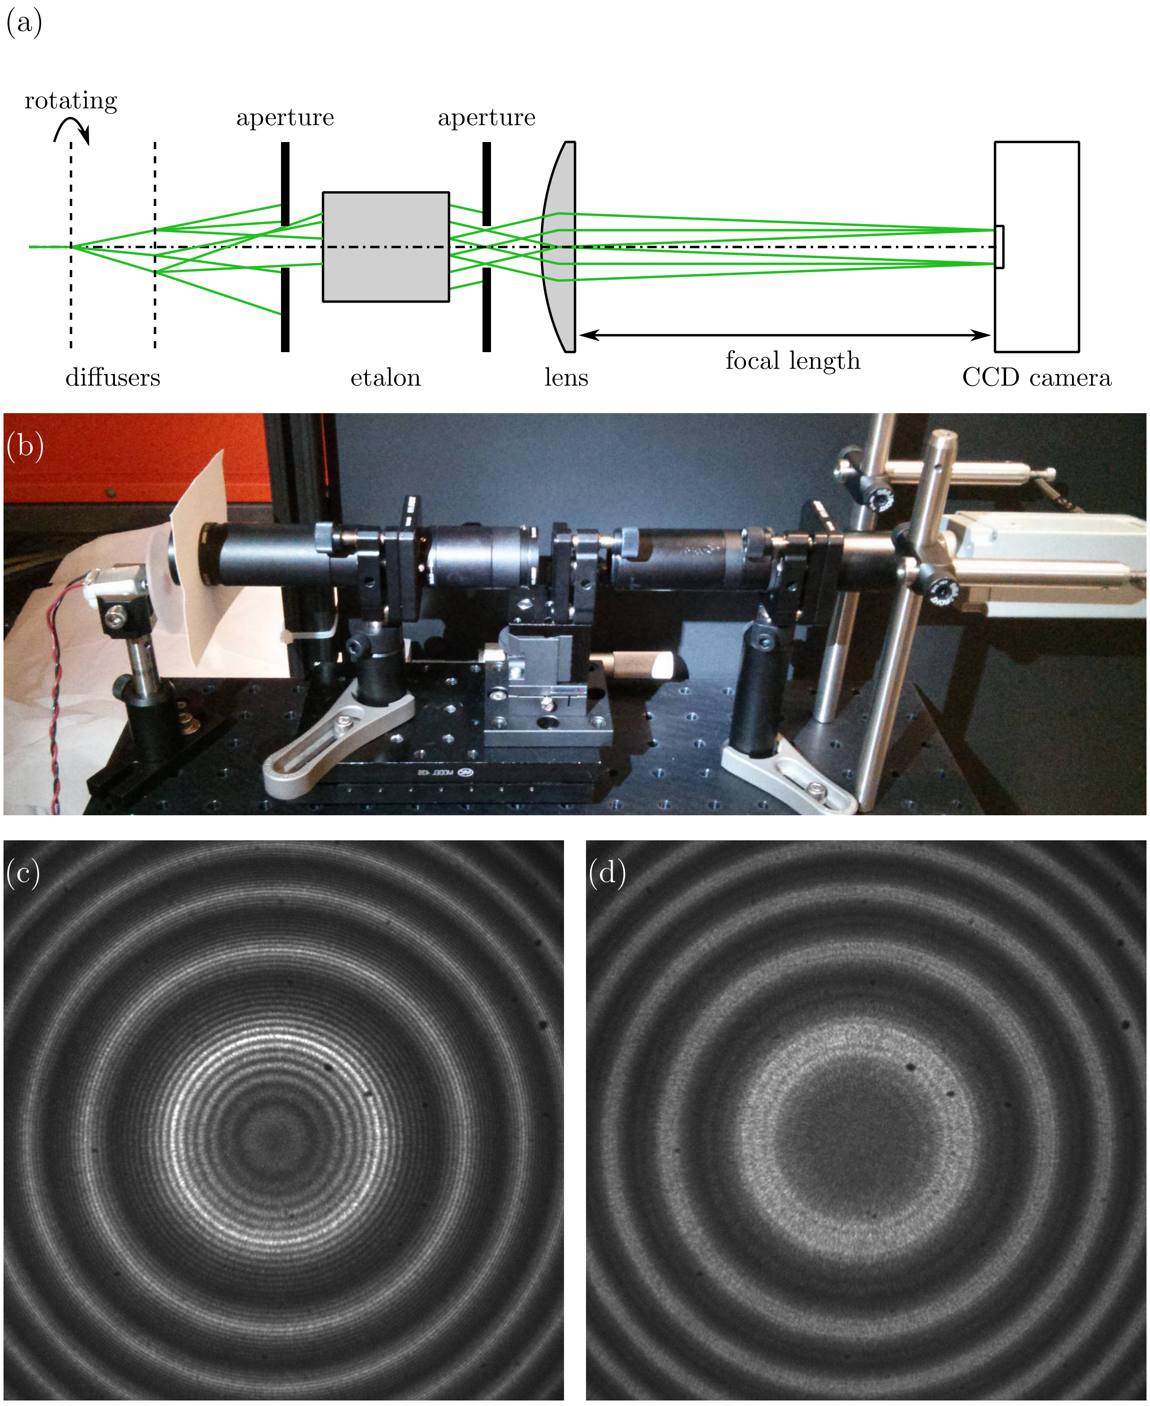

Figure 2.8: (a) Schematic of the etalon setup to analyse the PDL pulse frequency spection. (b) Photograph thereof. (c) Image of the etalon interference pattern at 462 nm with static PDL 2 cavity-end mirror and (d) with piezo-dither enabled (both averaged over 90 pulses at a fixed grating position).

onto the input face of the etalon. For the transmitted intensity, we obtain

$$I(\lambda, \alpha) = \frac{I_0(\lambda, \alpha)T^2}{1 - 2R\cos\mathit{\Delta}\varphi + R^2} \tag{2.2}$$

with the input intensity I_0, the reflectivity R of the etalon coatings assuming they are identical on the input and output face, and the transmission T. The output angle of the transmitted light is identical to the angle of incidence α. Hence, the lens with focal length f turns the angle into a radial position $r = f \sin\alpha$ in its focal plane, resulting in an image of concentric rings (see 2.8, (c) and (d)).

We achieve the highest resolution with a custom etalon (LightMachinery OP-8251-15000), which has a thickness of $d = 15\,\mathrm{mm}$, a reflectivity of $R = 98.0$–$99.6\,\%$ over a wavelength range of 400–660 nm, a surface flatness with 0.6 nm root mean square (rms), and a diameter of $1/2''$. These parameters achieve an FSR $\approx 2\pi \times 7\,\mathrm{GHz}$ and a (calculated) finesse $\mathcal{F} \approx 80$–130. For the lens, we use a standard plano-convex singlet with $f = 250\,\mathrm{mm}$ focal length and broadband-antireflection coating. We characterized the setup using a cw laser at 493 nm and obtained a minimum FWHM of the PDL cavity modes of $\approx 2\pi \times 150\,\mathrm{MHz}$, which is about a factor of 2 larger than the naïve expectation from $\mathcal{F}$. We speculate that alignment and focussing imperfections, as well as fundamental limitation due to the numerical aperture of the imaging lens, are the main contributions to this broadening. In summary, we can put an upper bound of $\approx 2\pi \times 150\,\mathrm{MHz}$ on the linewidth of the cavity modes of the PDL based on the measurements with the etalon at around 462 nm. A lower bound of $\approx 2\pi \times 40\,\mathrm{MHz}$ is estimated from a Fourier-transform-limited pulse assuming a pulse duration of 10 ns. We have other stock etalons with larger linewidths which cover wavelength ranges outside the aforementioned one. Thus, while we have no reason to believe that the PDLs behave fundamentally different in those wavelength ranges, our means to control their spectrum are more limited.

For our search, we apply a 1 Hz sawtooth voltage to the piezo adjuster of the oscillator-cavity endmirror of PDL 2 and select an amplitude such that the dither covers $\approx 1.5 \times \mathrm{FSR}$ of the PDL cavity. With the 30 Hz-repetition rate, this corresponds to a pulse-to-pulse dither of $\approx 2\pi \times 20\,\mathrm{MHz}$ and should sufficiently “smear out” the PDL and, hence, the VUV spectrum to avoid spectrally dark regions. In addition, as detailed later, a data point involves several scans over each probed wavelength at single-grating-step increments of 0.5 pm (corresponding to roughly the PDL-cavity FSR) at a rate of $1\,\mathrm{s}^{-1}$, which provides further "smearing out" of the spectrum. We record etalon images with integration times of 3 s during each data point to check the PDL spectrum. However, we have been unable to measure the spectrum of each and every individual VUV pulse directly such that a small uncertainty remains as to whether the VUV spectrum has no significant intensity variations.

The light from PDL 1 is directed with three backside-ground mirrors (Edmund Optics, 248 nm excimer laser-line mirrors, P/N 47-985) to the xenon cell. For PDL 2, we use two high-power mirrors whenever available (various types, depending on the output wavelength), or two uncoated fused-silica 90 °-prisms otherwise. Both PDL beams are overlapped on a window-substrate mirror (Spectral Products PM-1025-FS-248-45 or comparable; AR-coated or uncoated backside) with PDL 1 being reflected and PDL 2 being transmitted. (While we did not observe any significant wavelength-dependent changes in transmission, the pulse energy of PDL 2 is monitored, because the coatings are not specified for the entire wavelength range scanned by PDL 2.) After initial tests, we intentionally avoided any individual delays in the beam paths of PDL 1 and 2; tuning of the angle between them; beam shaping optics (except for those between oscillator and amplifier stage of the PDLs); or polarization control, because we could not identify any significant gains. On the contrary, additional optics typically led to losses due to reflections on imperfect AR coatings and/or absorption due to damages or accumulating contaminant layers. In addition, any increased complexity and/or wavelength dependence is undesired in long-term wide-band scans, as they require more maintenance and potentially reduce the reliability of the system. All the mirrors in the PDL beam paths require careful cleaning, as they easily accumulate contaminant layers. Still, coatings particularly in the 250 nm beam path show damages after a few 100 h of illumination. We mitigate the degradation by illuminating the mirrors off-center and rotating them slightly in their mounts on a routine basis (or replacing them entirely), while preserving their precise alignments with the help of irises and alignment-laser beams.

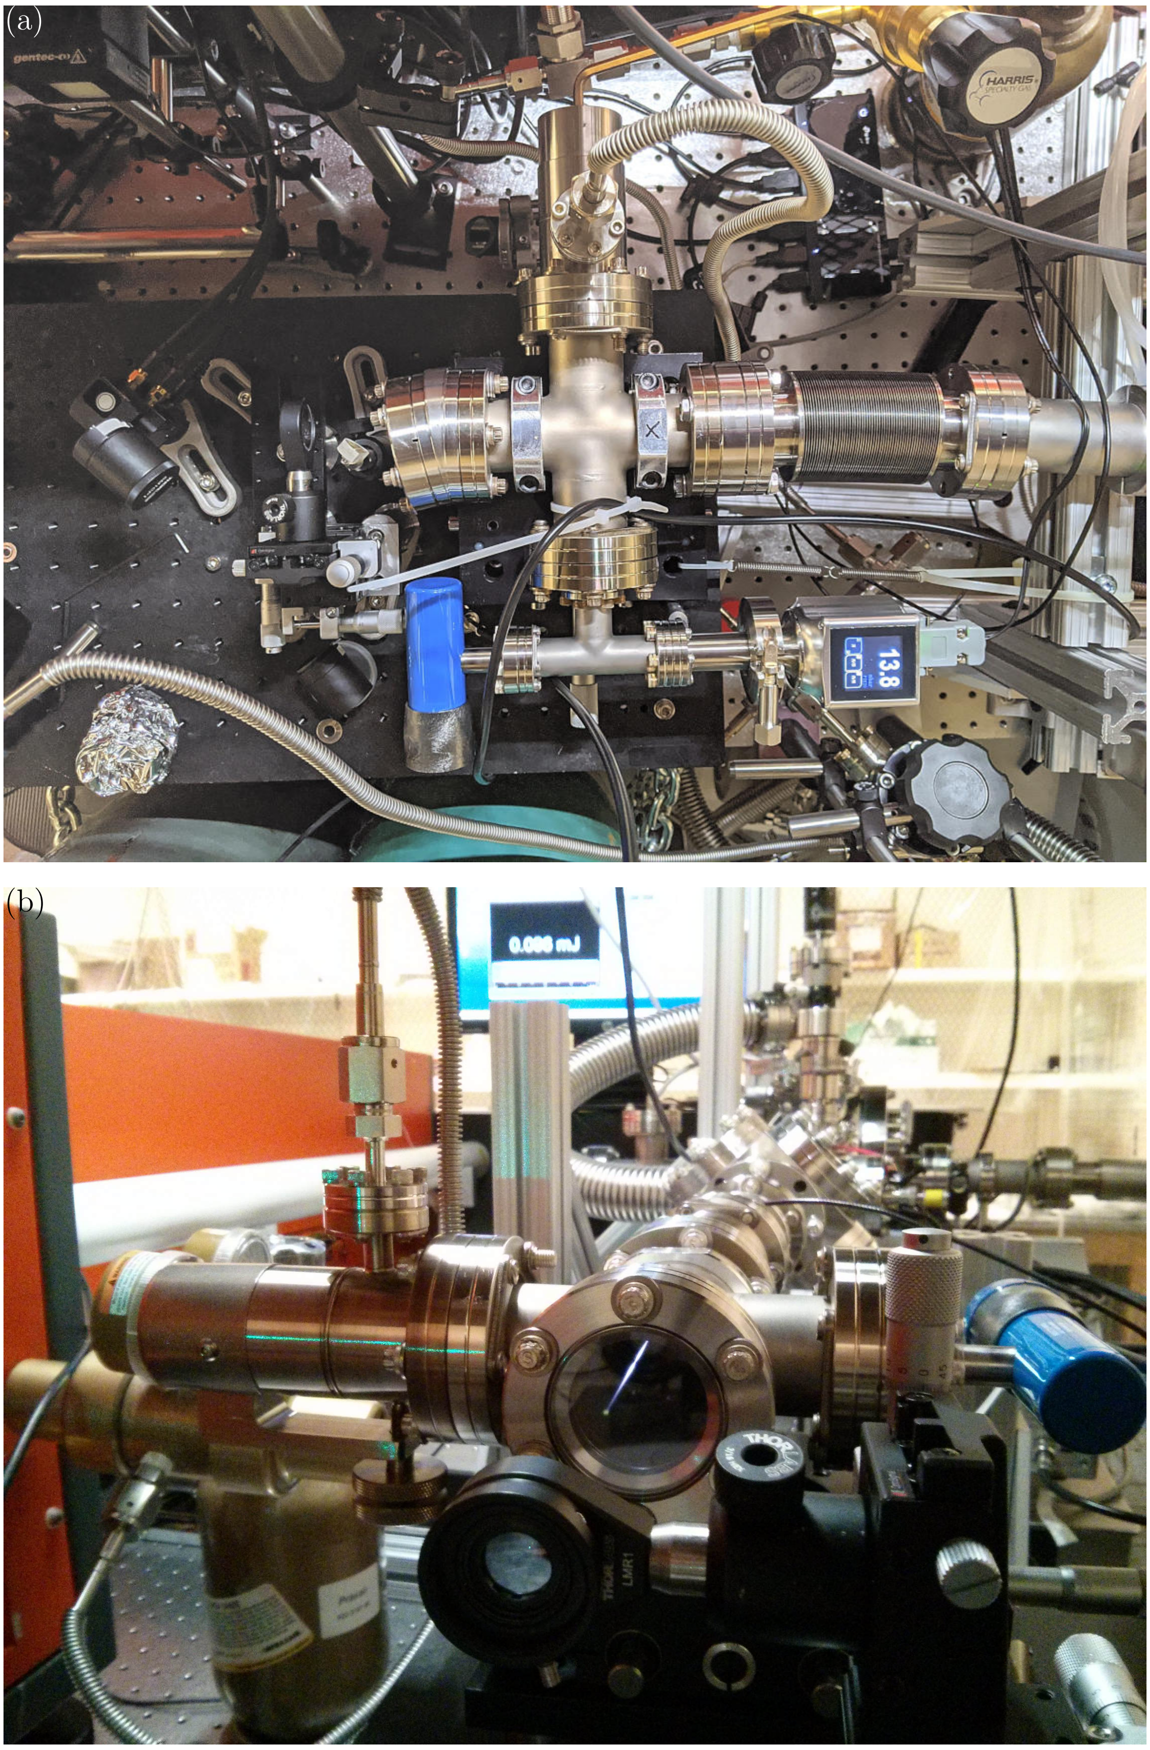

Figure 2.9: Photographs of the xenon cell with input lens; both mounted on translation stages. (a) Top view with both PDLs off. Also visible is a mirror on a motorized flip mount in front of the input lens, which is raised to reflect the PDL beams (both off) into a pulse-energy detector for power monitoring and dumping. (b) The doubled light from PDL 1 at $\approx 250\,\mathrm{nm}$ is focussed into low-pressure xenon gas and visible fluorescence can be observed (PDL 2 is off). The shape and brightness of the fluorescence give a good indication that laser and gas parameters are suitable for VUV generation. The green-ish dot on the input window is due to fluorescence in the bulk of the glass.

2.2.1 VUV Generation and Vacuum Chamber

The overlapped beams from PDL 1 and PDL 2 are focussed with a stock plano-convex lens (Edmund Optics, fused silica, $R = 68.79\,\mathrm{mm}$ with 266 nm-AR coating, P/N 38-764) into the xenon cell (see 2.9). Here, the curved side of the lens must face the incoming light to prevent a focus in the backreflected light. The xenon cell holding a low-pressure xenon gas (or mixture of xenon with other gases) is constructed using a standard CF 2 3/4″ cross with a length of 125 mm. The input window consists of deep-UV-grade or excimer-UV-grade fused silica (MPF Products A0650-1-CF and A0652-1-CF, respectively) and is mounted under an angle of $10\,^\circ$ with a custom double-sided CF adapter. The angle allows to separate backreflections off the window, use the light for the xenon spectroscopy described above, and prevent optics damages from uncontrolled beams. (Initially, we did not use the angled adapter and a focused backreflection destroyed the input lens during routine optimization.) The input lens and input window suffer from surface contamination as well as bulk and surface damage over time. As the lens has only an AR coating and the window is uncoated, these damages have a less severe effect on the transmitted power than the previously discussed damages to the high-reflectivity coatings of the mirrors. We clean the air-side surfaces routinely, rotate the input window at times, and replace both as needed on timescales of a few months to a few years (depending on operation).

The output side of the xenon cell is sealed off by the separation lens, which is plano-convex and made of VUV-grade MgF_2 (eSource Optics MF2595LCX, $R = 37.8\,\mathrm{mm}$). The lens is mounted into a custom double-sided 2 3/4″ Conflat (CF) flange with a recessed face for a Viton o-ring against which the lens is pressed using a Macor and stainless-steel ring (see 2.10(a)).

One side flange of the four-way cross is used for pumping and gas supply via a leak valve (Duniway VLVE-1000) and all-stainless-steel Swagelok VCR lines and valves. Using a scroll pump, we routinely achieve $\approx 10^{-2}$ mbar in the xenon cell prior to filling it with the low-pressure xenon gas (research grade, 99.999 % purity) or gas mixture.

The other side flange of the cross has pressure gauges. Initially, we used only a Pirani gauge (KJLC 275 series), which provided the necessary reproducibility and fast

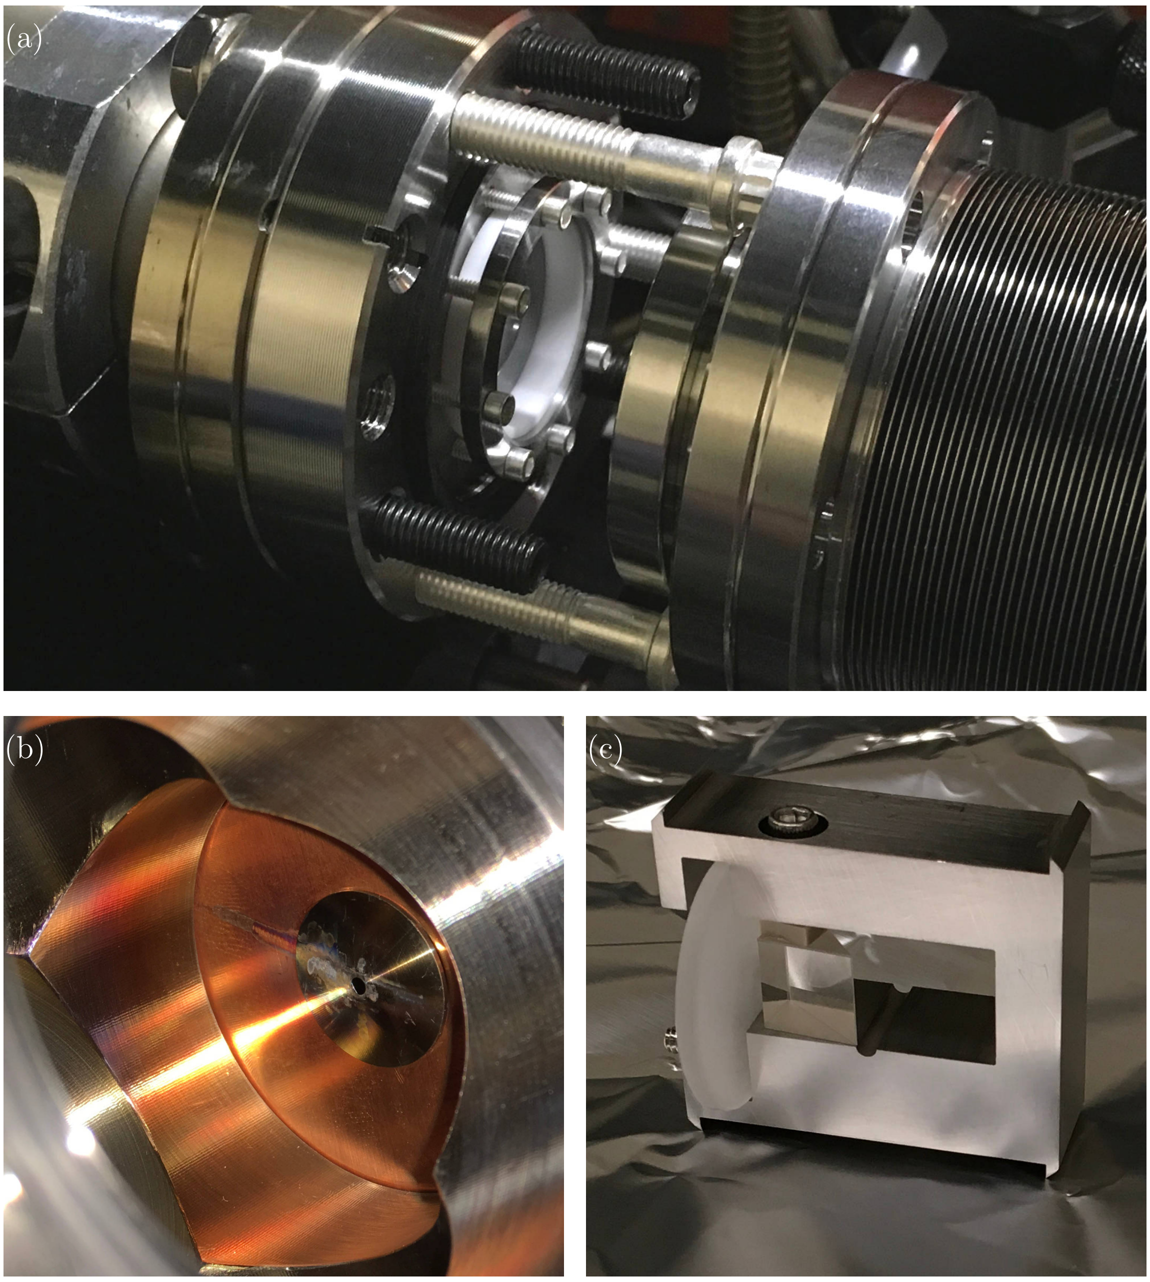

Figure 2.10: (a) Photograph of the separation-lens mount. The double-sided flange is attached to the xenon cell on the left. The lens, Macor ring and stainless-steel ring are visible in the center; the Viton o-ring is hidden in the recess. The flexible bellow, which normally connects the other side of the separaton-lens mount to the vacuum chamber (on the right), is detached and pushed back by three temporarily inserted silver-plated screws. (b) The conical stainless-steel pin hole press-fit into a CF blank gasket of the CF cube. Damage from the laser light is clearly visible. (c) Mount with prism and diffuser, which is inserted into the cube right in front of the pin hole.

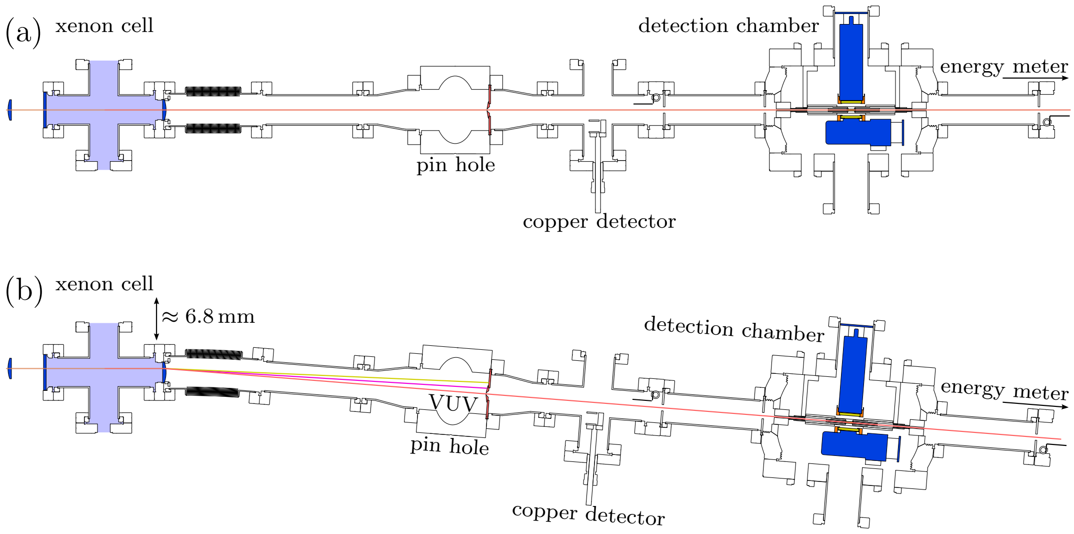

Figure 2.11: Simplified cross-section of the xenon cell and vacuum chamber of the pulsed VUV laser system (top view) in (a) the straight configuration and (b) the configuration for VUV generation.

response to pressure changes. Its pressures readings, however, are uncalibrated and highly non-linear for xenon as well as other heavy noble gases. The phase matching in xenon can be improved, in particular, toward higher VUV photon energies, by adding other gases (argon or krypton, research grade, 99.999 % purity) to the xenon. At this point, we added a piezo gauge (MKS 902B-11014), which provides an absolute gas-independent pressure measurement and significantly simplifies the optimization of the gas mixture. Xenon gas pressures and gas mixtures as a function of wavelength are published in Ref. HCS09; overall, we found the optimal parameters for our apparatus to be in agreement and produce comparable VUV photon fluxes.

An overview of the beam paths, the xenon cell, and the connected vacuum chamber is given in 2.11. In the straight configuration, the energy detector portion is removed and an alignment laser is directed against the propagation direction of the PDL beams through the (vented) vacuum chamber and xenon cell. This allows alignment of the height and tilt of the experiment chamber (with its ≈ 3 mm-diameter apertures), the pin hole (2 mm diameter), and the optical axis of the separation and input lenses. Subsequently, the vacuum chamber is rotated by ≈ 5 ° about the separation lens and the xenon cell is translated by ≈ 7 mm. Using the separation lens off-axis, its dispersion separates the VUV beam from the PDL beams. The former propagates to the experiment chamber, while the latter are dumped at the pin hole.

The pin hole is housed inside a CF 3 3/8″ cube with an oberservation window on one side and additional pumping on the other side. The pin hole itself has a slightly conical shape, is made of stainless-steel with a 2 mm-diameter aperture, and press-fit into a CF 3 3/8″ blank gasket. Damage from the laser light is clearly visible (see 2.10(b)). To prevent further damage and potential issues from ablation products, we added a prism in front of the pin hole, which reflects the high-pulse energy PDL beams by roughly 90 ° into a diffuser (see 2.10(c)). The input face of the prism is slightly tilted to avoid direct back-reflections of the PDL beams.

The light from PDL 1 and PDL 2 gives rise to other effects besides the difference-frequency generation indicated in 2.5: The sum-frequency generation process producing VUV light at $\omega_{\mathrm{VUV,+}} = 2\omega_{\mathrm{UV}} + \omega_{\mathrm{VIS}}$ is allowed and competes with the $\omega_{\mathrm{VUV,-}}$ generation. The beam at $\omega_{\mathrm{VUV,+}}$, however, does not propagate into the experiment chamber, as it is not transmitted by the separation lens and, additionally, spatially filtered out at the pin hole.

Further, the two-photon excitation with $\approx$ 250 nm alone leads to two further effects: ASE at $\approx$ 828 nm from the $6p\ ^2[1/2]_0 \rightarrow 6s\ ^2[3/2]^\circ_1$ transition (see 2.6) is generated both in forward and backward direction, which has a short pulse duration comparable to that of the PDLs [RDG89a, AS96a]. Additionally, the lifetime of the $6s\ ^2[3/2]^\circ_1$ level is increased by several orders of magnitude due to radiation trapping in the xenon gas, such that an emission at $\approx$ 147 nm can be observed with microsecond-decay times [AS96a]. The first effect can easily spoil the xenon spectroscopy signal, if beams are not aligned carefully, as we reuse a reflection of the $\approx$ 250 nm beam off the input window of the xenon cell. The second effect has little effect on the search using ^{229}Th-doped crystals, but affects another experiment, which detects electrons on timescales of $> 1\,\mu$s after a VUV pulse [WSS19]. An addition of nitrogen (N_2) to the xenon of the xenon cell (research grade, 99.999 % purity; $\approx 1 : 10$ xenon-to-nitrogen ratio) enhances the quenching rate of the $6s\ ^2[3/2]^\circ_1$ level and reduces the decay times to order 100 ns. (Not surprisingly, these effects can also be observed during regular VUV generation.)

The vacuum-facing side of the lens is prone to accumulation of hydrocarbons, where the laser beams exit the lens. These hydrocarbon deposits can be identified as a visible, brown spot and decrease the transmission of the lens. We had relatively good experiences removing them by generating VUV, while the flexible bellow is pushed back (as shown in 2.10(a)) and oxygen is blown over the (normally vacuum-facing) lens surface. We speculate that the VUV and PDL beams produce oxygen radical, ozone, and possibly a plasma, which react with the hydrocarbons and lead to their removal. Over time, we rotate the separation lens, as surface and bulk damages increase, until it is replaced after roughly 1 yr. The implementation of an argon purge described in the next section greatly reduced the hydrocarbon build-up rate.

2.2.2 Experiment Chamber

The experiment chamber is described in detail in Ref. JSS15. Modifications include the replacement of the in-vacuo solenoid-driven shutters with home-built versions that are manipulated with a pneumatically-driven linear-motion feedthrough (see Fig. 2.12). This mitigates any heat from the in-vacuo solenoids, which slightly affected the PMTs in Ref. JSS15. In addition, the side-on PMT is replaced with a second head-on PMT (Hamamatsu R6835). Further, home-built PMT mounts now produce a vacuum seal with the entrance window of the PMTs; hence, the electrical contacts of the PMTs are located on the ambient side now and there is no longer any need for in-vacuo voltage-divider boards for the supply voltages. The minimization of in-vacuo components improves vacuum conditions and simplifies maintenance. In addition, PMTs can now be continuously operated independent of the pressure inside the vacuum chamber, while operation in intermediate pressure regimes was prohibited for in-vacuo voltage-divider boards due to greatly reduced breakdown voltages.

Hydrocarbon build-ups are also observed at the entrance and exit face of the ^{229}Th-doped crystal. After $\approx$ 100 h of VUV illumination, the crystal originally had to be removed from the experiment chamber and undergo 8–24 h in an asher system, where a radio-frequency plasma is generated in a flow of mbar-level oxygen. The crystal is not exposed to the high-energy ions of the plasma directly, but only oxygen radicals, ozone, and UV radiation, which greatly lengthens the procedure but prevents surface damage.

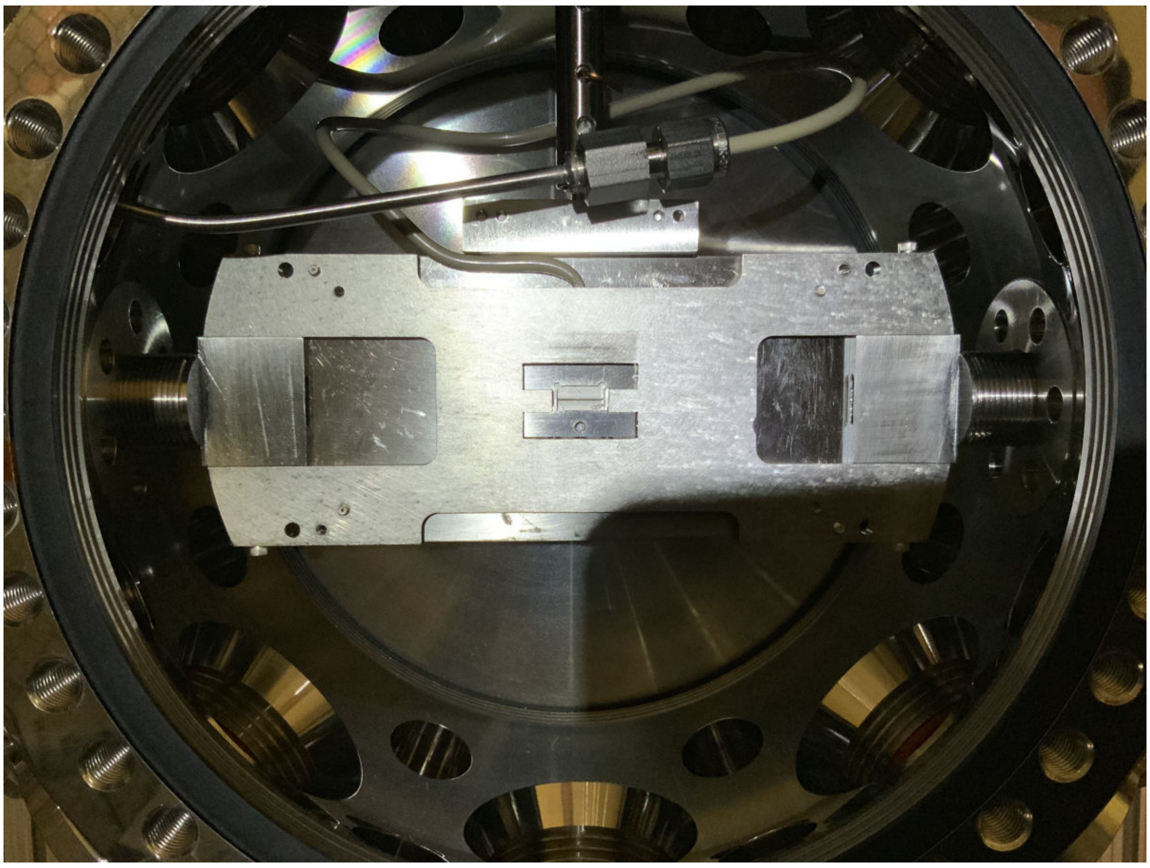

Figure 2.12: (Left) Photograph of the upgraded experiment chamber. The low-profile shutter with in-vacuo solenoids [JSS15] have been replaced with a mechanical shutter, which is attached to a pneumatic actuator. The front portion of crystal mount with front shutter flap is removed and the crystal is visible. The stainless-steel and PEEK tubings on the top can introduce an argon flow into the central portion of the crystal mount to purge the crystal faces.

To eliminate this tedious and time-consuming maintenance procedure, the vacuum chamber is upgraded with a low-pressure (sub-mbar) purge of argon (research grade, 99.999 % purity), which is directly blown onto the vacuum-side of the separation lens, crystal input and output faces, and (in a later experiment) other optical surfaces in the VUV beam path (see Fig. 2.12). The argon purge reduces the mean free path of outgassing contaminents from meters to (sub-)millimeters and likely aids in flushing them out of the vacuum chamber, before their dissociation products can reach and build up on the optical surfaces. The gas supply lines are, again, constructed with all-stainless-steel Swagelok VCR components and valves on the ambient side. The gas is fed through feedthroughs with thin stainless-steel tubes into the chamber, where they adapt to PEEK tubes via Swagelok tube fittings. The PEEK tubes have the advantage that they can be precisely thermally formed before installation; still, they are extremely durable and significant mechanical deformations during the insertion process through the vacuum ports do not lead to damage or permanent deformations. The argon is transparent for the VUV in our wavelength range and no impact from contaminants can be identified: the VUV pulse energy remains constant (within optimization limits), when the entire vacuum chamber is filled up with 1 bar argon as compared to the vacuum chamber being held at 10^{-7} mbar-vacuum levels.

2.2.3 VUV Pulse Energy Measurement

While a (quasi-)continous-wave VUV photon flux can be measured comparatively easily with photodiodes [JSS15], calibrated measurements of VUV pulse energies is more difficult and often involves knowledge of atomic or molecular cross sections Ref. YLR92, BBR17. For the laser system presented here, we developed a pyroelectric-detector assembly (see Fig. 2.13), which has a relatively flat calibration spectrum as a function of wavelength and can be assumed to have a measurement error much smaller than a factor of 2 even without any calibration for the accessible VUV wavelength range. In addition, it is sensitive to sub-1 µJ pulse energies.

Briefly, a pyroelectric detector (Gentec QS5-IL; material: $LiTaO_3$) is used, which comes in a standard TO-5 package with integrated operational amplifier. It is installed on a prototyping board (also Gentec) inside the vacuum chamber. The pyro-

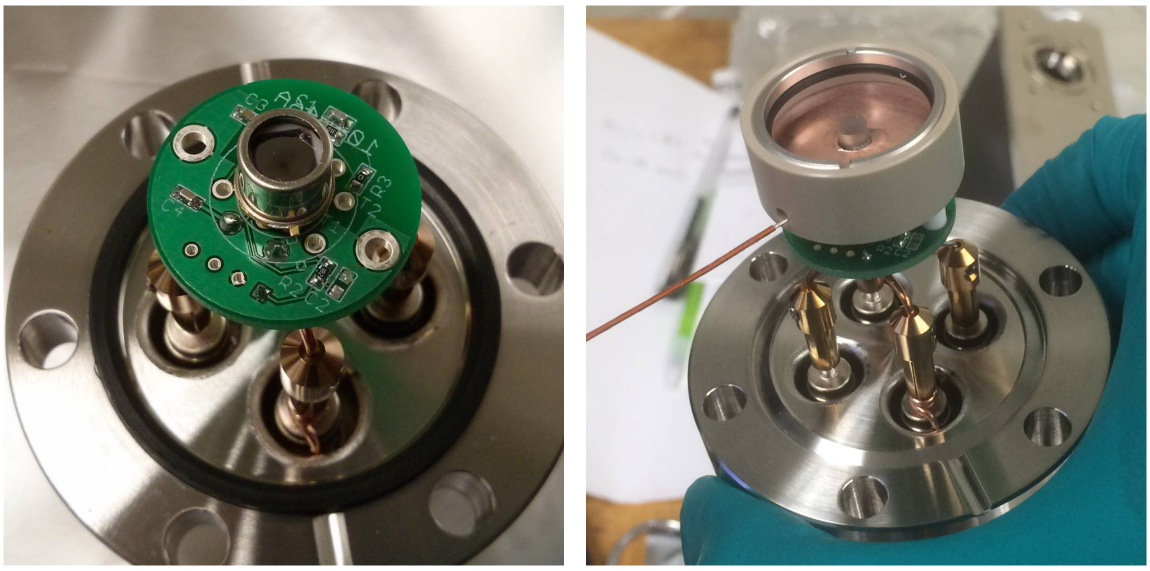

Figure 2.13: Photograph of the in-vacuo pyroelectric-detector assembly. Bare prototyping board with Gentec QS5-IL and printed-circuit board on a 4-pin vacuum feedthrough (left). The same with copper and glass aperture assembly added (right).

electric crystal is coated with thin chromium layer on its front. Per manufacturer, the chromium has a relatively flat spectral absorption of $A \approx (50\text{–}60)\,\%$ in the green to near-infrared wavelength range; it increases sharply to $A \approx 75\,\%$ at 250 nm. In the absence of accurate absorption data in the VUV and considering that the chromium surface quickly shows significant hydrocarbon deposits, we assume unity absorption, $A = 100\,\%$ in the VUV. This results in a conservative estimate of our VUV photon flux. (If a more accurate calibration was required, it would be relatively straightforward to obtain the absorption by detecting the backreflected light pulse with a second pyroelectric detector.) The operational amplifier in the QS5-IL is used as a charge amplifier and measures between the chromium-coated front and the back side of the pyroelectric crystal, which is also connected to the can of the TO-5 package. Our assembly is configured to produce a signal amplitude of 50 mV/µJ using a 10 nF feedback capacitor (see Fig. 2.14).

The use of the detector assembly in-vacuo with VUV wavelengths leads to several complications: The incident laser pulse can release photoelectrons from the chromium surface of the pyroelectric crystal which will also be detected as a competing charge compared to the pyroelectric effect by the operational amplifier. Further, stray light on the can of the TO-5 package and associated photoelectrons can amplify the mea-

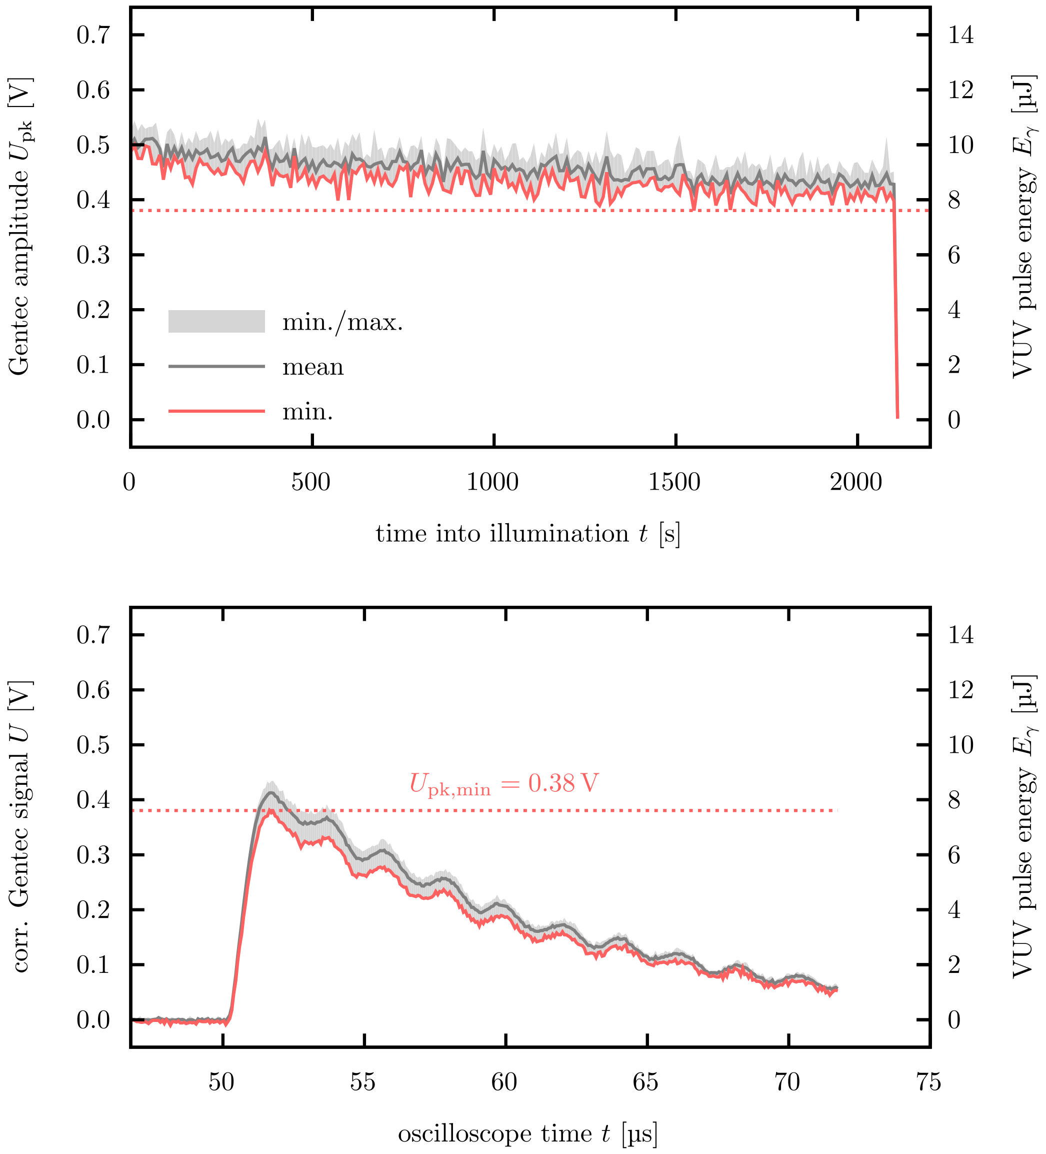

Figure 2.14: (Top) Example of the pulse energy as monitored for an entire data point of the experiment using the pyroelectric detector. (Bottom) An example of the raw pyroelectric-detector signal. Min./max./mean denote the minimum, maximum, and mean, respectively, for pulse-energy measurements belonging to the same grating position of PDL 2 during a scan.

sured charge. To mitigate these effects, our assembly has two apertures installed in front of the pyroelectric detector: First, a small aperture made of standard glass (e.g. BK7) is used to absorb any stray VUV light (or portions of the beam in case of significant misalignments). This glass aperture prevents photoelectrons from other apertures and the TO-5 can. Second, a larger copper aperture is installed between glass aperture and the pyroelectric detector, which can be biased with a voltage to surpress photoelectrons from the chromium surface of the pyroelectric crystal.

In a later experiment, we improved the pyroelectric-detector assembly to be fully UHV compatible. For this, we mounted only the pyroelectric crystal in an in-vacuo mount and placed the entire operational amplifier and other circuitry on the ambient side.

CHAPTER 3

Superconducting Nanowire Search for the Internal Conversion of the Isomer

3.1 Concept

In the past few decades, the field of superconducting nanowire single photon detectors (SNSPDs) has seen the development of systems capable of single photon detection from the far infrared [TWK23] to the extreme ultraviolet [FAN22]. SNSPDs are capable of high quantum efficiencies ($> 99\%$ [RNN20]) and extremely low dark count rates ($< 10^{-3}$ Hz [HCN19]). By fabricating the single nanowires in a meandering pattern, an active area of several thousand square microns with high filling factor can be achieved, allowing for diverse applications in quantum key distribution [VZS02], astronomical observation [YQW18], and even particle detection [ADD12].

SNSPDs operate through the breaking of Cooper pairs in the condensate. If an incident photon or particle strikes the nanowire, energy is deposited in a small region which breaks the Cooper pairs. Early so-called "normal core" or "hot-spot" models (depicted in Fig. 3.1) predicted the formation of a "normal core" of finite resistance,

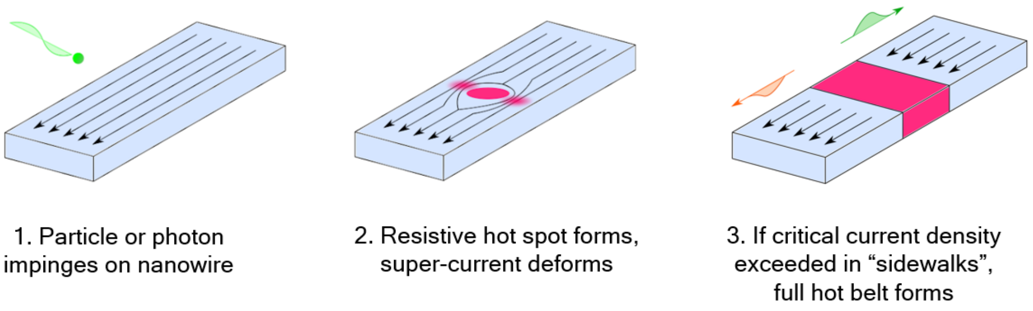

Figure 3.1: Depiction of the "hot-spot" model of detection in SNSPDs.

around which the super-current diverted to take the path of least resistance. As the supercurrent diverted into these "sidewalks", the local current density could exceed the critical current density of the nanowire, causing a breakdown of the superconducting state in the entire transverse profile of a region of the nanowire. This normal region acts as a resistor in the middle of the wire, which will have a finite voltage drop as current flows across it. This sudden voltage pulse is then picked up by detection electronics, forming the "click" signaling the detection of a photon/particle. The validity of the "normal-core" model is the subject of some debate, with other models calling for the suppression, but not total destruction of the superconducting state by the incident particle, while still others involve the generation of vortices in the supercurrent [ERI15]. What is common to all these models is a local suppression the parameter which sets the current density necessary to trigger the breakdown of the superconducting state in the region of the nanowire struck, and that the degree of suppression should scale with the energy of the incident photon or particle. The system detection efficiency (SDE) of a SNSPD setup is expressed as

$$\mathrm{SDE} = \eta_{coupling}\eta_{absorption}\eta_{internal}\eta_{trigger}, \tag{3.1}$$

where $\eta_{coupling}$ is the system dependent coupling efficiency of the source to the SNSPD; $\eta_{absorption}$ is the efficiency with which the incident photon or particle is absorbed or stopped in the detector, and will depend on the device material properties; $\eta_{internal}$ describes the efficiency with which the energy deposited will result in the formation of a normal region; and finally, $\eta_{trigger}$ describes the efficiency with which the electronics can register the electrical response of the detector to an "click" event. Typically one only has direct access to the overall SDE.

It is known that for a monochromatic source of light, the count rate (and therefore the SDE) exhibits a sigmoidal shape with bias current (as seen in Fig. 3.2). This can be understood from the detection model. If the bias current is too low, then even if the critical current density has been suppressed, the current density will not be high enough to cause a "click". Past a certain point however, the bias current will be high enough to trigger the device. This would imply the SDE (and hence the count rate) should be a sharp step function, but fluctuations in the inherent detection process, inhomogeneities in the device, and the finite energy width of the

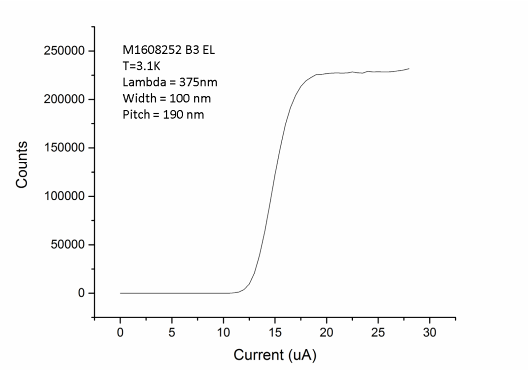

Figure 3.2: “Small” SNSPD counts vs bias current curve, adapted from the thesis of Justin Jeet [Jee18]. This calibration was performed at NIST against 375 nm photons at a temperature of 3.1 K.

source itself all contribute to smearing out the detector response. Since the SNSPD response is thought to be an integration of underlying Gaussian processes [KLM17], the detection rate $\mathit{\Gamma}$ is typically reported as

$$\mathit{\Gamma}(I_b) = \lambda \,\mathrm{erf}\left((I_b - I_c)/\sigma_I\right), \tag{3.2}$$

where we adopt the following definition of the Gaussian error function

$$\mathrm{erf}(x) = \frac{1}{\sqrt{\pi}} \int_{-\infty}^{x} e^{-t^2} dt. \tag{3.3}$$

In Eq. 3.2, λ represents the maximum count rate, σ_I width of the curve, I_b is the bias current, and I_c is the center of the detection rate curve (which we will refer to as the “center current”).

One can also consider the behavior of the detector when the bias current is kept fixed, but the energy of the incident particles/photons is varied. Assuming relatively flat detector response for absorption over the scanned energy range, the primary change will be the degree of suppression of the critical current density. In the language of the “normal-core” model, the hot-spot will be bigger. Thus, at a fixed bias current, a low energy incident particle may not cause enough of an effect to generate a “click”, but a higher energy incident particle would. What this translates to is that the center current I_c of the SNSPD will shift to lower currents as it is struck by higher energy photons and particles. It is this energy dependent behavior of the center current that unlocks the SNSPD’s potential as a spectroscopic tool.

The concept for the SNSPD experiment at UCLA was to bombard excited ^{229m}Th into a superconducting nanowire detector optimized for VUV photon detection. The excited ^{229m}Th would be sourced by a ^{233}U source, which has a 2% branching ratio into the isomeric state. As the ^{233}U undergoes α decay it emits ^{229m}Th recoil ions. These recoil ions are implanted into the nanowire with tens of keV of kinetic energy, which is gauranteed to cause a detector click. The detector then re-thermalizes with the cryostat environment, and the supercurrent is re-established. Roughly 10 μs later, the isomer should undergo internal conversion (IC) decay, depositing its energy into the nanowire [SWT17]. This IC decay will deposit the equivalent of a VUV photons energy into the nanowire, causing a click if the bias current is high enough. One can measure the rate of double click events as a function of the bias current, and hence measure the sensitivity of the detector to the energy released in the internal conversion of the isomer. If one then takes a set of calibration curves for VUV photons of different wavelengths, one is able to determine the energy released in the IC process by comparing the IC curve to the calibration curves. In other words, one is able to perform spectroscopy on the isomeric transition as well as measure its lifetime in the nanowire.

3.2 Early Experiments at UCLA

The early SNSPD experiments at UCLA are described in detail in [Jee18], but will be described briefly to discuss the main results. These experiments used a ^{233}U source purchased from Eckert and Zeigler. The source consisted of a stainless steel disk with a 19 mm active diameter and an activity of 21.22 kBq with NIST traceable calibration. A 1 mm hole was drilled through the middle of the ^{233}U source in order to directly couple in light from a deuterium lamp monochromater (McPherson 234/302). The cryostat cooled the sample region to 3.3 K to 3.8 K.

The most extensive data was taken with the WSi "thick" SNSPDs. It was found that there was a detectable fraction of events with several microsecond inter-arrival times (IAT). The rate at which these "double-clicks" occurred was at the 10^{-3}-10^{-2} Hz level, which was within the magnitude of the rate at which IC decays were expected to occur in the nanowire. Finally some VUV calibration curves were taken,

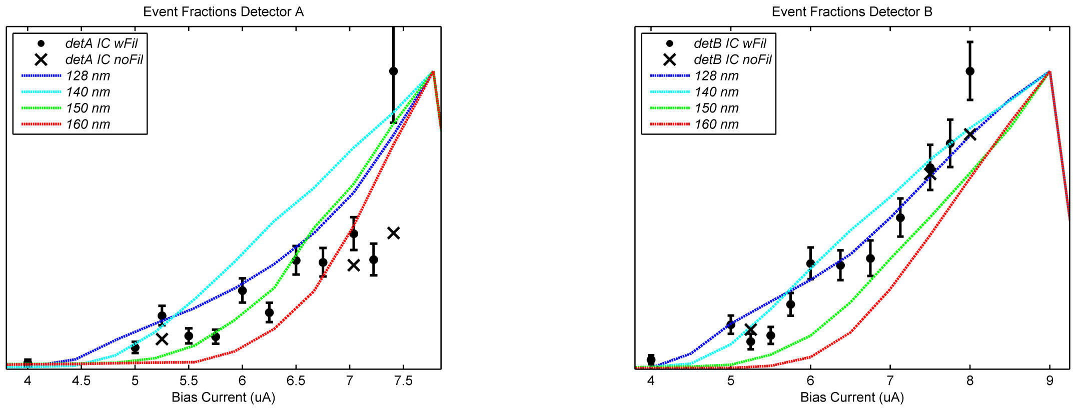

(a) “Thick” WSi detector A.

(b) “Thick” WSi detector B.

Figure 3.3: In situ VUV calibrations compared to IC event fractions. The dashed curves respresent the VUV calibration measurements with (blue, cyan, green, and red) representing the response from photons of wavelengths (128, 140, 150, and 160) nm. The black dots and x’s represent data taken with and without low pass filters, respectively. All curves have been normalized to unity. Figure adapted from [Jee18].

and early comparisons were made between the IC event rate and the VUV calibrations (see Fig. 3.3)

Concurrent to the initial analysis of the ^{233}U data, a control measurement was carried out with a 20 kBq electroplated ^{241}Am source (produced and calibrated by E&Z). The ^{241}Am α decays into ^{237}Np with energies of $\sim$ 5.6 MeV, comparable to the energies of ^{229}Th ions produced in the α decay of ^{233}U . While the half life of ^{237}Np is $\sim$ 2.144 million years, it is known to have an isomeric level with $\sim$ 0.7 μsec lifetime [SLC90]. This lifetime puts the state well within the detection window the experiment. Unlike the ^{229}Th isomer, however, the ^{237}Np isomeric level should yield a γ ray of hundreds of keV of energy. And so one should measure a flat “double clicks” curve with respect to bias current. Unfortunately, to our knowledge the population of the ^{237}Np isomer from ^{241}Am decays is unknown, and so an estimate of the double click rate from the ^{241}Am is not possible.

The result of the control measurement is that the “double click” rate curves for both the ^{233}U and ^{241}Am curves were remarkably similar (a comparison can be seen in Fig. 3.4). It was found that within a factor of 2 the “double click” curves were nearly identical between the two sources. This led us to believe that whatever we

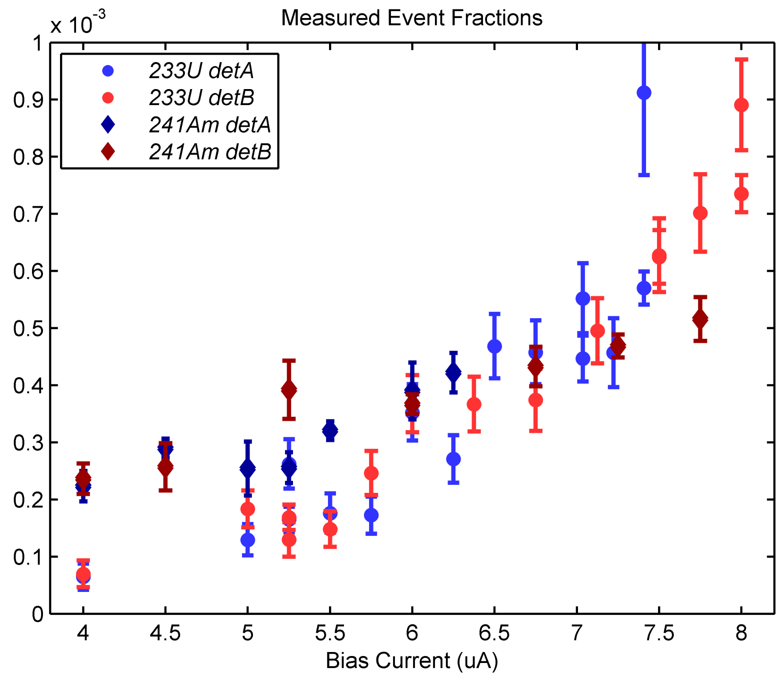

Figure 3.4: Comparison of the measured "double click" rates normalized by total activity hitting the nanowire for both ^{233}U and ^{241}Am. Figure adapted from [Jee18]

were measuring was some other process that was masking the signal from the isomer.

3.3 The SNSPI at LMU

The initial indication of a low-energy, microsecond time scale signal in the SNSPD was promising, however the results from the ^{241}Am tests indicated that a microsecond timescale background was present that would need to be eliminated. It was known from previous experiments with electroplated radio-isotopes that a considerable background of low energy electrons are emitted along with nuclear decay products [SSR18]. The recoiling nucleus generates lattice defects in the electroplated layer, and it has been speculated that the relaxation of the lattice leads to these so-called "exo-electrons". These exo-electrons can be orders of magnitude more numerous than the expected signal rate from the internal conversion decay, and they are emitted from the electroplated ^{233}U on the microsecond timescale, directly overlapping with the expected internal conversion time window.

Around this time our group, the group of Dr. Sae Woo Nam and Dr. Galen O'Neil at NIST, and the group of Dr. Peter Thirolf at Ludwig-Maximilians-Universität (LMU), began a collaboration aimed at leveraging superconducting nanowire single-photon imagers (SNSPI). Simply put, the SNSPI is just an SNSPD in a transmission line geometry, so that there are coaxial connections at both ends of the nanowire. With this geometry, if a click occurs in the detector, there is a voltage pulse that goes out both ends of the device. By measuring the timing difference between the pulses with high speed electronics, the position of the event within the detector can be determined.

The Thirolf group maintains an ion beam system capable of separating out ^{229}Th ions from a ^{233}U source that can be accelerated towards a downstream detection system. The details of this system can be found in the theses [Wen17,Sei19], so we will offer only a brief description. ^{229}Th ions (2% of which are in the isomeric state) from a ^{233}U source are cooled in a cryogenic helium buffer gas cell before going supersonic de Laval nozzle. The ions can then be held and accelerated to a few hunderd eV through an RFQ buncher. Any other contaminant ions are separated through a quadurpole mass separator before entering a downstream experiment region. In this experiment, the ions were neutralized through a graphene foil before impacting a MoSi SNSPI detector. The ^{229}Th ions that are in the isomeric state are thought to land in the downstream SNSPI, and de-excite, producing two clicks within a few microsecond time window. Due to far separation of the ^{233}U and the SNSPI, any backgrounds due to other radioactivity-induced processes should be eliminated, providing a clean IC signal.

The results of implantation of ^{229}Th ions and control measurements with ^{230}Th ions can be seen in Fig. 3.5 and Fig 3.6, respectively (all analysis presented for LMU results was performed by Dr. Justin Jeet). The striking result of this experiment was that there indeed was both a temporally and spatially correlated double-click signal present in the ^{229}Th data that was not present in the ^{230}Th data. This supported the idea that one would be able to observe the IC decay in the nanowire, and perhaps even measure its energy. Unfortunately, over the course of the LMU experiment it was found that the ion implantation rate was not stable enough to perform an accurate energy calibration. It was inferred that ice was slowly building up on the surface

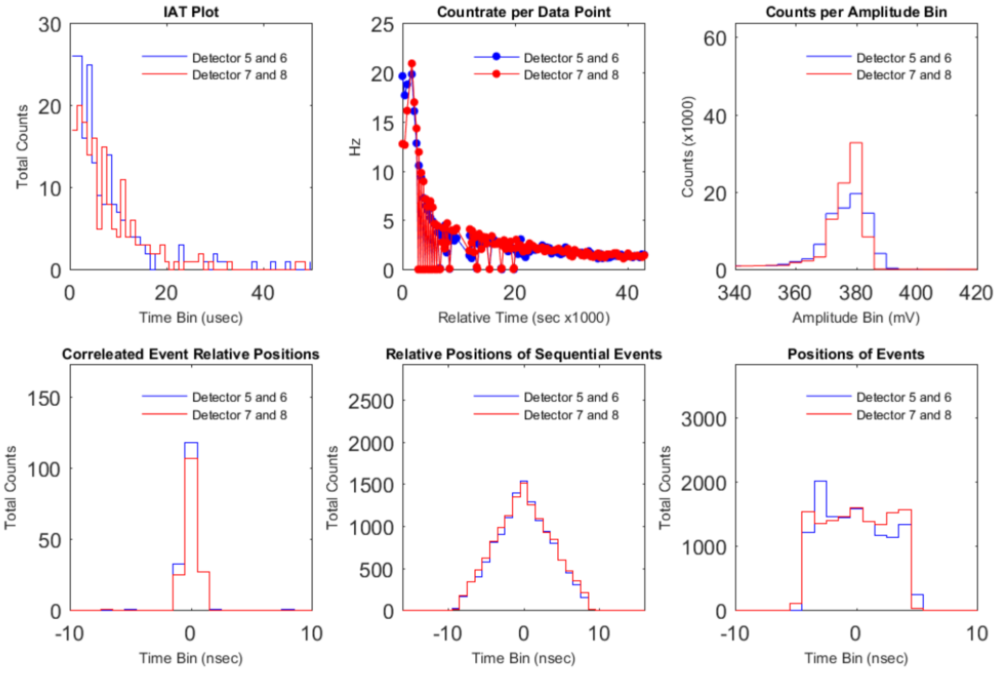

Figure 3.5: Analysis of a datapoint where $^{229}\mathrm{Th}^{3+}$ ions were implanted into the LMU SNSPI. Note the appearance of ~ 10 μs inter-arrival time clicks which occur at the same location. Analysis performed by Dr. Justin Jeet.

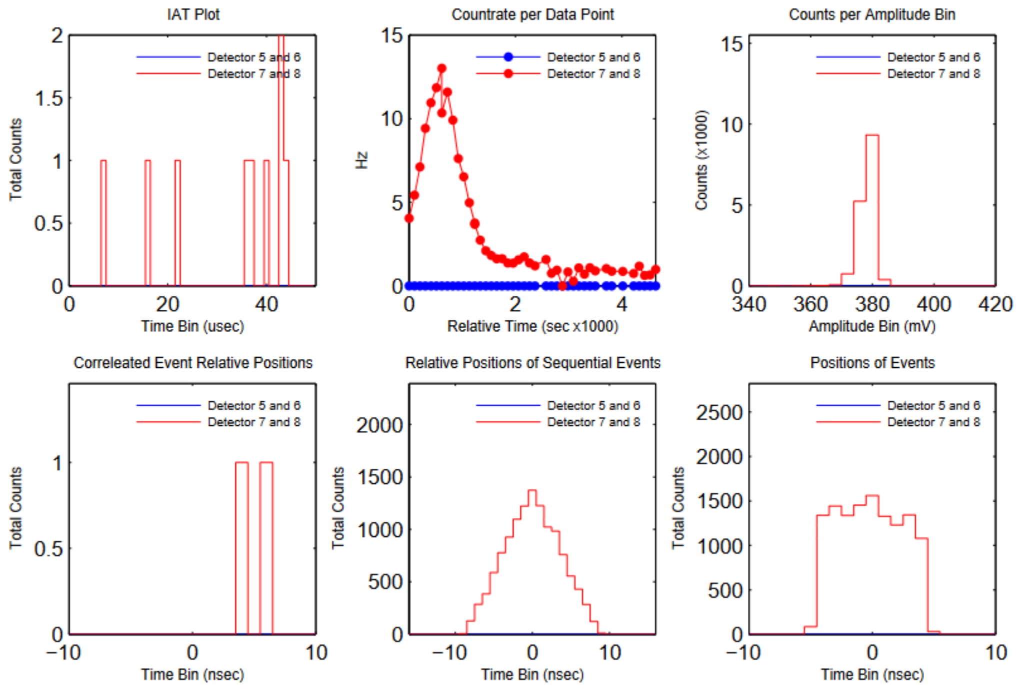

Figure 3.6: Analysis of a datapoint where $^{230}\mathrm{Th}^{3+}$ ions were implanted into the LMU SNSPI. Note the lack of ~ 10 μs inter-arrival time clicks which occur at the same location. Analysis performed by Dr. Justin Jeet.

of the SNSPI due to cryopumping, and it was difficult to obtain enough statistics for an energy measurement. In addition, it was suspected that some long term degradation was occuring in the detectors over time. Despite showing evidence of the first observation of the isomeric decay in a superconducting device, the experiment was eventually halted due to these technical limitations. To our knowledge, there is an ongoing effort at NIST to construct a copy of the LMU apparatus.

3.4 SNSPI at UCLA

The valuable insights gained from the LMU experiment were that it was possible to observe the IC decay in the nanowire, and that these could be picked out using the position resolution of the nanowire. It was decided that a simplification of the LMU setup could be achieved at UCLA, by simply taking an SNSPI and placing it near the ^{233}U source. With the SNSPI, any exo-electronic background could be eliminated by cutting any short inter-arrival time (IAT) events that did not occur at the same location. By doing so, it was thought that the main background of the original UCLA SNSPD experiments could be eliminated. In order to use the SNSPI, a new cryostat of higher cooling power would be necessary in a vacuum apparatus that could operate for months to allow for the collection of sufficient statistics.

3.4.1 Detectors

Two SNSPI detectors, fabricated by the group of Sae Woo Nam at NIST, were used in our system. The detectors were quad pixel SNSPI (giving 4 imagers total) with 100 μm x 100 μm pixels. The detectors themselves consist of several layers with the following thicknesses:

- 2 nm: Amorphous Si
- x nm: SNSPI nanowire (properties in Table 3.1)
- 150 nm: SiO2
- 100 nm: Gold

- 150 nm: SiO2 (oxidized wafer top level)
- 700 μm Si (wafer substrate)
- Unknown thickness: Elmer rubber cement
- 1.6 mm: FR-10 PCB material

Detector	Num. Pixels	thickness	wire width	pitch	fill factor	substrate
MoSi SNSPI	4	10 nm	120 nm	200 nm	0.6	Mo_3Si
WSi SNSPI	4	17 nm	120 nm	200 nm	0.6	W_3Si

Table 3.1: Detectors used in the SNSPI experiment at UCLA.

Both quad SNSPIs had 8x MMCX connector ports in order to connect to the readout transmission cables. Due to the large mismatch in the impedances between the 50 Ω coaxial cables and the superconducting nanowires, an impedance taper was used to bridge the difference (can be seen in Fig. 3.8). The amorphous nanowire materials MoSi and WSi were selected for the SNSPIs since their lack of well-defined crystal structure allows for deposition onto any substrate without degradation in properties due to lattice mismatch. In addition, longer meander features can be fabricated with minimal loss in uniformity [VKB15].

The capture efficiency $\eta_{capture}$ of the ^{229}Th ions into the nanowires was calculated using SRIM under the assumption that the density of the nanowires was given by a stochiometric weighted average of the densities of the constituent elements. The SRIM results can be seen in Fig. 3.7. The results were $\eta_{capture} \approx 0.15$ for the MoSi detector, and $\eta_{capture} \approx 0.86$ for the WSi detector.

3.4.2 Experimental Apparatus

The experiment was carried out in a vacuum chamber built around a Cryomech PT415 Cryocooler with CP1000 Helium compressor unit. The PT415 had a remote motor to minimize the coupling of vibrations and electronic noise into the experiment. The calibrated specs the of cryocooler are listed in Table 3.2. The 300 K portion of the experiment consisted of a custom aluminum dodecagon with 12x KF ports, a stainless

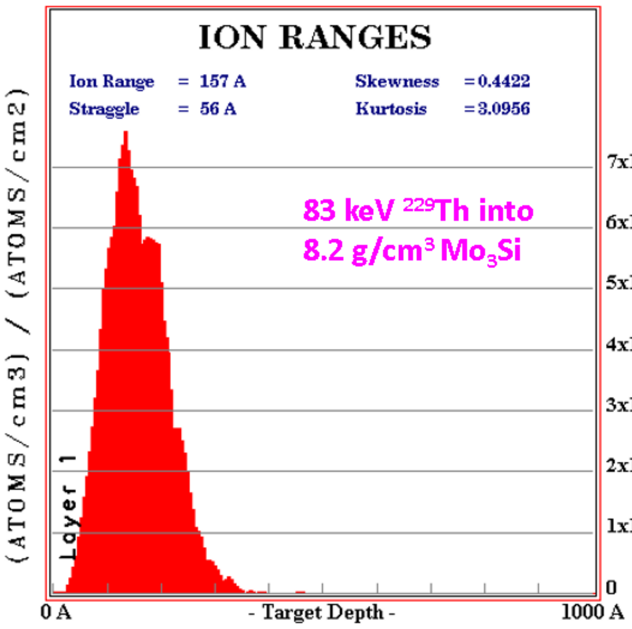

(a) Mo_3Si

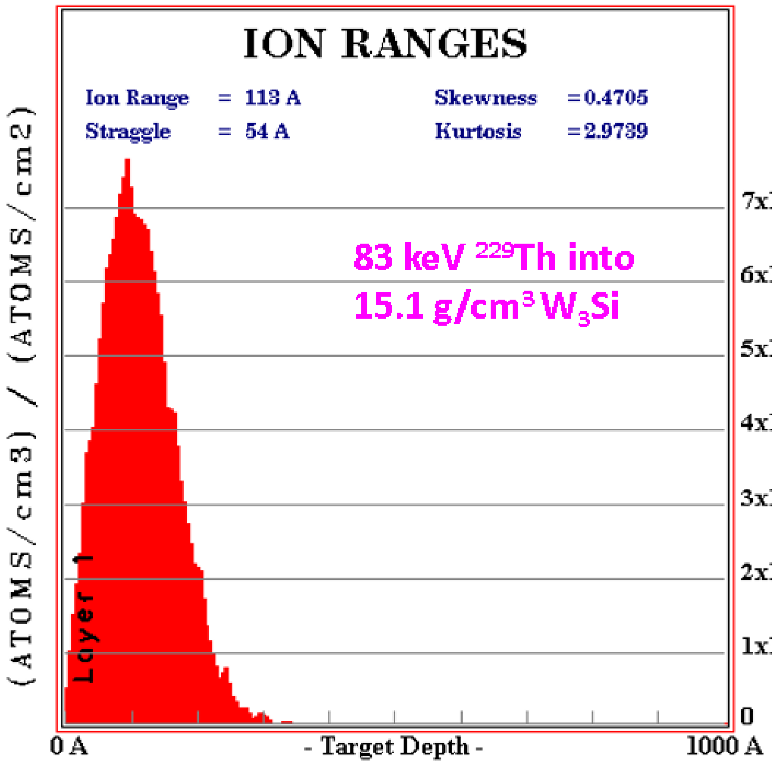

(b) W_3Si

Figure 3.7: SRIM calcuation of the stopping range of 83 keV ^{229}Th ions into nanowire substrates.

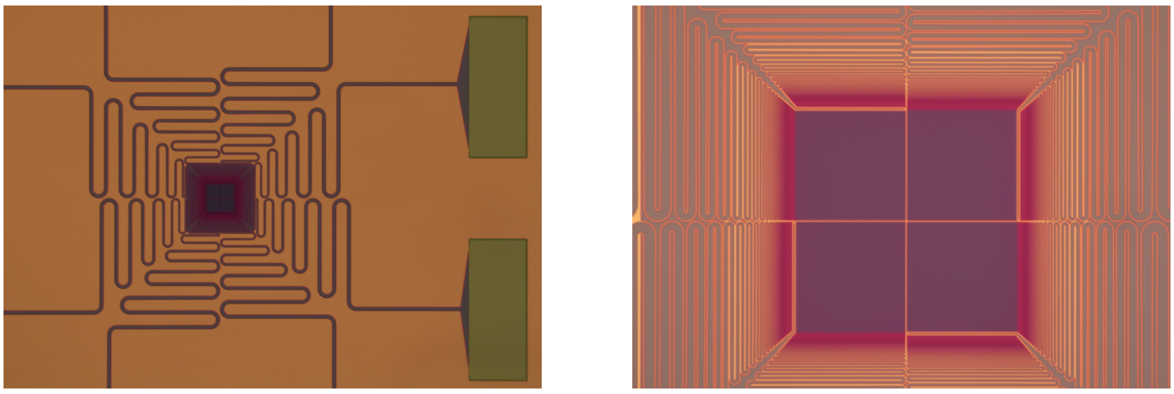

(a) The green blocks represent the coaxial connection region where the impedance taper begins.

(b) Zoom in to quad imager region. One can see the meandering wires tapering into the detector region.

Figure 3.8: Colored images of the wire layout in the quad pixel SNSPI detectors. The quad pixels themselves are meandering nanowires that fill in the square 100 μmx100 μm region. (Image Credit: Dr. Dileep V. Reddy)

steel cylinder that covered from the bottom of the dodecagon to just below the 3 K heat exchanger, and a bottom plate with an attached 2.75“ CF MgF$_2$ viewport. Between this bottom viewport was a $\sim$3” long 2.75“ CF bellow to couple the vacuum chamber to a McPherson 234/302 monochromater for in-situ VUV calibrations. The ports of the dodecagon were for the following:

- 4x Custom feedthroughs for readout of the temperature diodes
- 2x SMA feedthroughs with 4x SMA connectors
- Burst cap in case of catastrophic pressure build-up in the chamber upon warm-up
- Pirani guage for pressure measurement (MKS 275 Convectron)
- Gate valve for vacuum pumping and sealing during cooldown
- High voltage feedthrough to allow for voltage biasing of ^{233}U target

All other unused ports were sealed with brass KF blanks (the choice of brass was a matter of keeping costs minimal).

1st Stage		2nd Stage		Compressor	
Temp. [K]	Cooling Power [W]	Temp. [K]	Cooling Power [W]	High Pressure [PSIG]	Low Pressure [PSIG]
30.2	0	2.45	0	280/270	90/85
30.7	0	4.35	1.5	281/272	90/80
45	47.5	2.50	0	290/280	90/80
45	43.0	4.20	1.5	290/285	90/80

Table 3.2: Manufacturer specs for Cryomech PT415 cryocooler used in UCLA SNSPI experiment.

The main experimental stage was the 3K-can in which the ^{233}U target and SNSPI were separated by 4.5 mm and cooled down to 2.5-3.1 K. The ^{233}U target was the same described in Sec. 3.2, and was held in a copper mount. The SNSPI was screwed down to a Au-plated copper pedestal, which was in turn screwed into a Au-plated copper plate. It was this larger Au-plated copper plate that was bolted to the 3K stage of the cryostat. The ^{233}U target was also attached to this pedestal, and was separated from the SNSPI by a few washers. At the bottom of the 3K-can was a small window that was sealed with a MgF_2 window. This hole served to couple the VUV light in from the monochromator. The 3K-can was radiation shielded by a 30K-can, that was itself attached to another Au-plated copper plate which was bolted to the 30 K stage of the cryocooler. All bolting was done with brass screws and Belleville washers. The high thermal expansion coefficient of brass compared to other metals was thought to help maintain mechanical contact between the joined components during thermal cycling. The Au-plating was chosen was to ensure high thermal conductivity between the helium cold heads and the 30/3K-cans, while the low emissivity of Au minimzes the black-body radiation (BBR) load on the sample stage. Apiezon-N grease was used between the Au-plated pieces and the cold head to ensure thermal contact, and between the Au-plated pieces and the radiation shield aluminum cans to provide lubrication.

The electrical pulses from the SNSPI were be coupled out of the device using standard MMCX cabling plugged into the nanowire chip and then adapted into SMA. Each of the Au-plated copper plates featured 8x SMA electrical feedthroughs. It should be noted at all cabling and connection lengths were matched so as to minimize the introduction of any bias in the differential timing measurements. The SMA cabling in the experiment between the 30K and 3K plates consisted of custom stainless steel coaxial cables. The choice of stainless steel was made in order to minimize thermal contraction and maintain impedance matching upon cooldown, as well as limit the thermal conductivity between the different cold stages. The cabling between the 30 K stage and the 300 K chamber feedthroughs was accomplished by a chain of stainless steel cabling to again act a thermal buffer, a 90° SMA (rated to 3 GHz), and then a looped brass SMA cable that allowed for some flex during cooldown. The looped brass SMA cables were what then connected to the SMA feedthroughs on the dodecagon. On the air side, the SMA feedthroughs were connected to a DC biasing/AC amplifier box via RG-58 coaxial cables. The biasing and readout of the SNSPI was accomplished through bias tees in the amplifier box, and the pulses were amplified through two gain stages (see Fig. 3.9 for details). The biasing of the SNSPI was of order several microamps, which was sourced through a home-built low-noise bipolar power supply. Each channel of the power supply had small internal resistor and ADC for current readout. This readout was calibrated to a precision of 0.1 μA.

Temperature sensing was done via 4x Lake Shore Cryotronics DT-670 CU uncalibrated silicon temperature diodes, and the readout of theses diodes was accomplished via a Lake Shore Cryotronics Model 336 temperature controller. 3x diodes were attached to different portions of the 3K stage for local monitoring, while one was mounted to the base plate of the 30K stage. A more detailed diagram of the placement can be seen in Fig. 3.10.

3.4.3 Cleaning Procedure

The following cleaning procedure was adopted for the smaller components (< 4" diameter):

- An initial wipe down of the components was done with isopropanol to remove gross contamination.

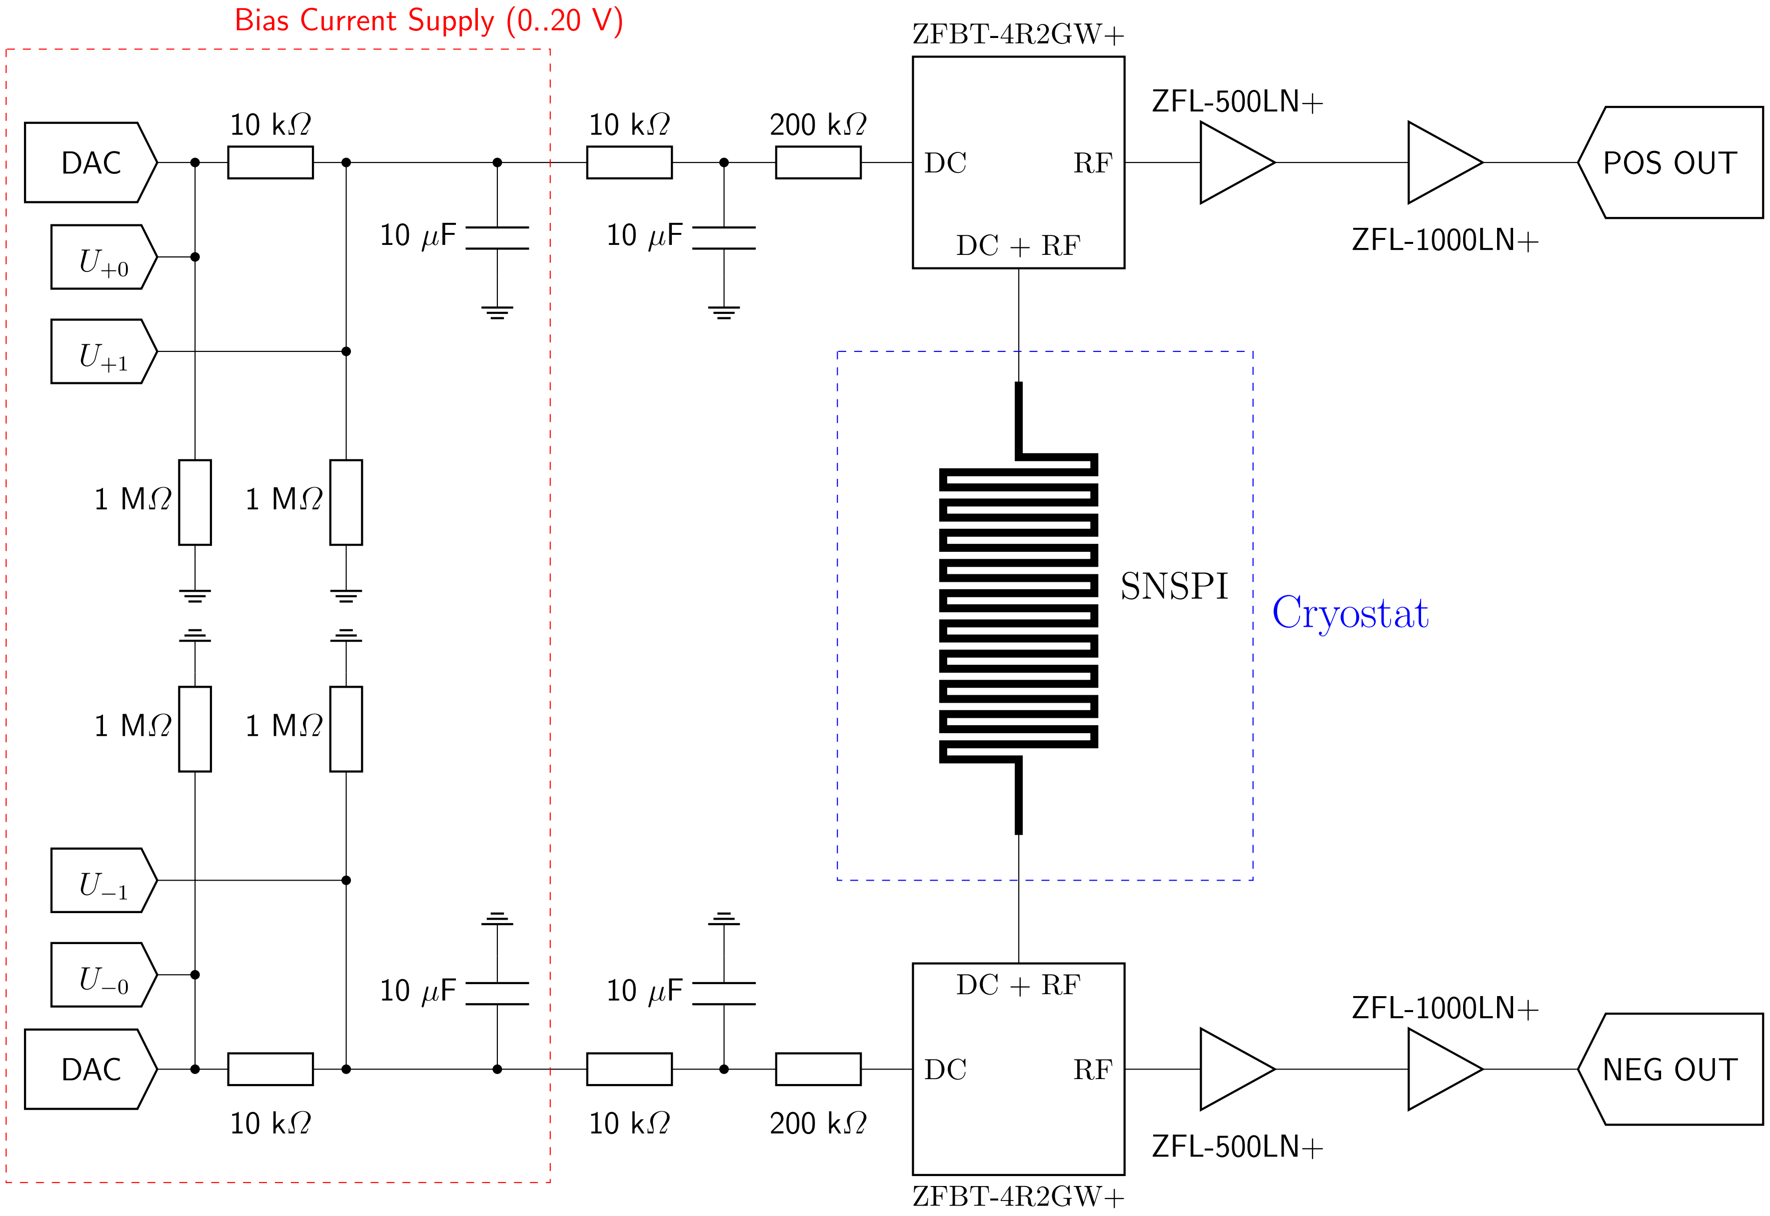

Figure 3.9: Diagram of the biasing and readout circuit for a single SNSPI pixel. The DAC's supplied the biasing voltages which were of equal magnitude and opposite polarity. ADCs (labelled U_{+0},U_{+1},U_{-0},U_{-1}) measured the voltage across a precision 10 kΩ resistor to calculate the bias current (imput impedances were 10 MΩ). POS OUT and NEG OUT were two channels of the CAEN digitizer.

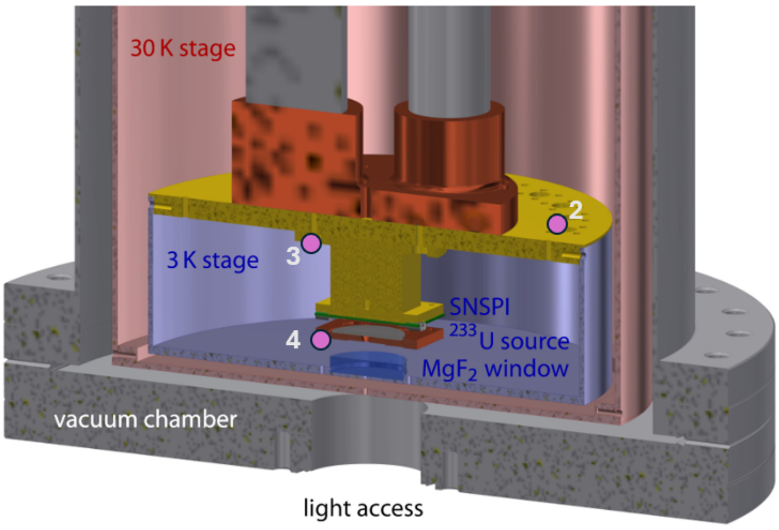

Figure 3.10: Placement of temperature diodes 2-4 in the 3K stage area. Diode 1 was attached to the gold plate that formed the 30K can (not pictured).

- Parts were sonicated in a room temperature of water and detergent. Citranox was used for copper and Au-plated copper, while Alconox was used for all other metals.
- Parts were then sonicated in isopropanol for ∼30 min.
- Parts were then removed from isopropanol, wiped with lint-free tissue, then dried in air.

For larger components such as the main 30/3 K Au-plated copper plates and the dodecagon, the components were sponged down with the appropriate water+detergent, sprayed with water, and then wiped with isopropanol wipes.

3.4.4 Pump Out and Cooldown Procedure

Prior to cooldown, the chamber would be pumped out for ∼1 day using a turbopump to pump out as much water as possible. The pressure would reach around the low 10^{-6} mBar as read out by the turbopump ion gauge. The cryostat compressor would then be turned on, which would cause a slight pressure increase in the chamber due to the initial heating of the helium. After several minutes the temperature inside the chamber would drop such that cryopumping would start to occur. This cyropumping was observed as a drop in the turbopump ion gauge pressure below the point just before the compressor was engaged. At this point, the gate valve to the cryostat is closed. After several hours the chamber reaches its low temperature point of ~ 3 K. A Pirani guage attached to the dodecagon monitors for any potentially large gas leaks while the compressor is active.

3.5 MoSi SNSPI

The MoSi SNSPI data-runs were a proving ground for many of the techniques that would mature during the WSi data-runs, including the VUV calibrations and the ^{148}Gd calibrations. As can be seen in Fig. 3.11, the calibration curves did not achieve a saturation plateau before the device switched into a normal state. This lack of saturation makes the fitting of the error function model unreliable, which in turn precludes an effective energy calibration of the device.

Despite the sense that the energy calibrations would not work out, initial efforts were made to look at the rate of events occuring in the MoSi nanowire with an inter-arrival time (IAT) of [2,20] μs at the same location while it was exposed to the ^{233}U source. It was found that the rate of these correlated events seemed to increase with bias current, though the shape of this curve was far to the red of any of the calibrations taken. A comparison was made with an electroplated 20 kBq source of ^{148}Gd produced by Eckert & Zeigler. As a pure α emitter with a half-life of $\sim$87 yr, the only products of the decay of ^{148}Gd are a 3.2 MeV α particle and a $\sim$89 keV ^{144}Sm recoil ion [Nic14]. The nearly identical kinematics to ^{233}U α decays make ^{148}Gd an ideal control source to see the effects of on the nanowire purely from bombardment by α particles, recoil ions, and exo-electrons. A potential complication comes in the form of an isomeric state in ^{144}Sm with a half-life of ~ 880 ns [KR71], however the chance of it being populated is exceedingly rare, and cutting on events with inter-arrival times longer than 2 μs should eliminate the majority of these decays.

The ^{148}Gd source was loaded, and curve was taken to see the rate of events occurring at the same location in the nanowire within a [2,20] μs IAT window. To our surprise, the curve seemed to be almost the same as the one taken for ^{233}U (see Fig. 3.12). No feature distinctive of the ^{229}Th isomer could be seen in the data. It was theorized that the low capture efficiency of the MoSi for recoil ions may also be allowing the IC electrons to escape, or their energy to leak out from the superconductor. Between that and the poor calibration curves, we decided to move on to the WSi detector.

3.6 WSi SNSPI

Given that the WSi nanowire was thicker and boasted higher stopping power than MoSi, it was thought that a clearer IC signal would be observable in this device. However, given the background of seemingly "correlated" events within the IC time window, it seemed clear that we would have to be expect a signal on the order of 10^{-4} Hz on top of a 10^{-3} Hz background. Therefore, we would need to take data for several months to collect sufficient statistics at a set of bias currents.

3.6.1 Calibration Curves and Temperature Normalization

In stark contrast to the MoSi SNSPI, the WSi detector exhibited clean saturation in its calibration curves that were well fit by Eq. 3.2 (see Fig. 3.13). Over the course of months of data taking it was found that the temperature would drift by hundreds of mK, which would cause the calibration curves to shift.

In the literature of superconducting nanowires, a recurrent quantity is the normalized bias current $i = I_b/I_s(T)$, where $I_s(T)$ is the temperatuere dependent switching current of the device [ES13,ERI15]. One could consider using the temperature dependent switching current as a way to correct for temperature drifts in our experiment. In order to test this idea an experiment was carried out in which calibration curves were taken at different temperatures. Variac power heater tape was wrapped around the outer shell of the chamber to heat the system and ultimately vary the temperature at the SNSPI. An I-V curve of the current flowing through the SNSPI circuit versus the applied voltage was measured. These curves exhibit a sudden drop off when the

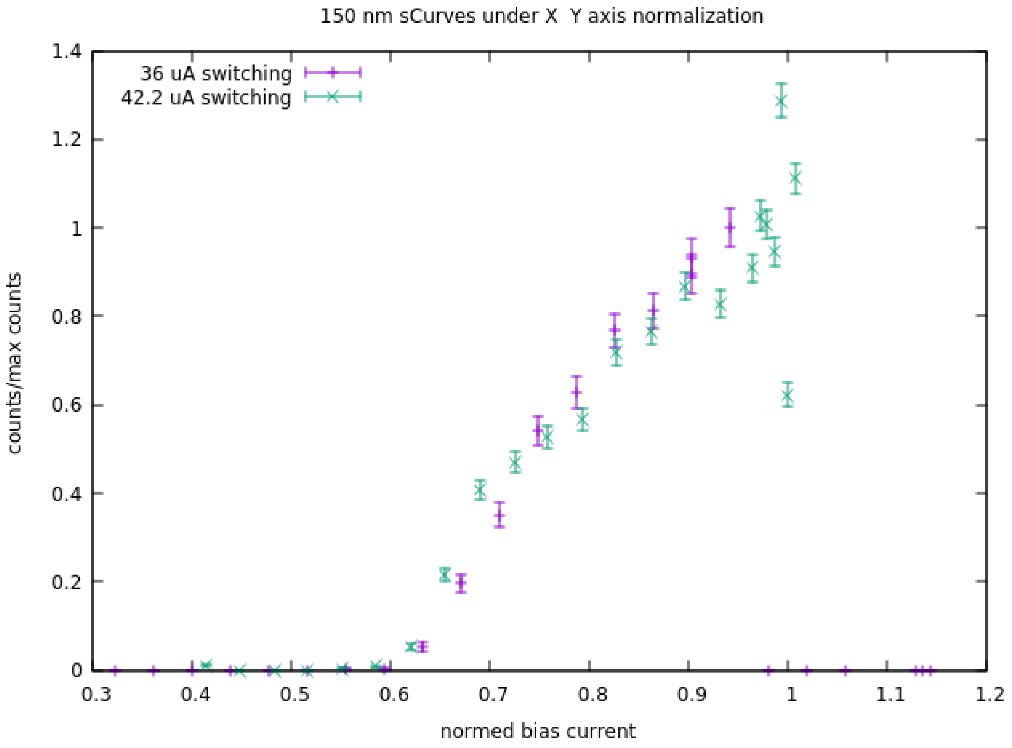

(a) 150 nm calibration curve for MoSi SNSPI.

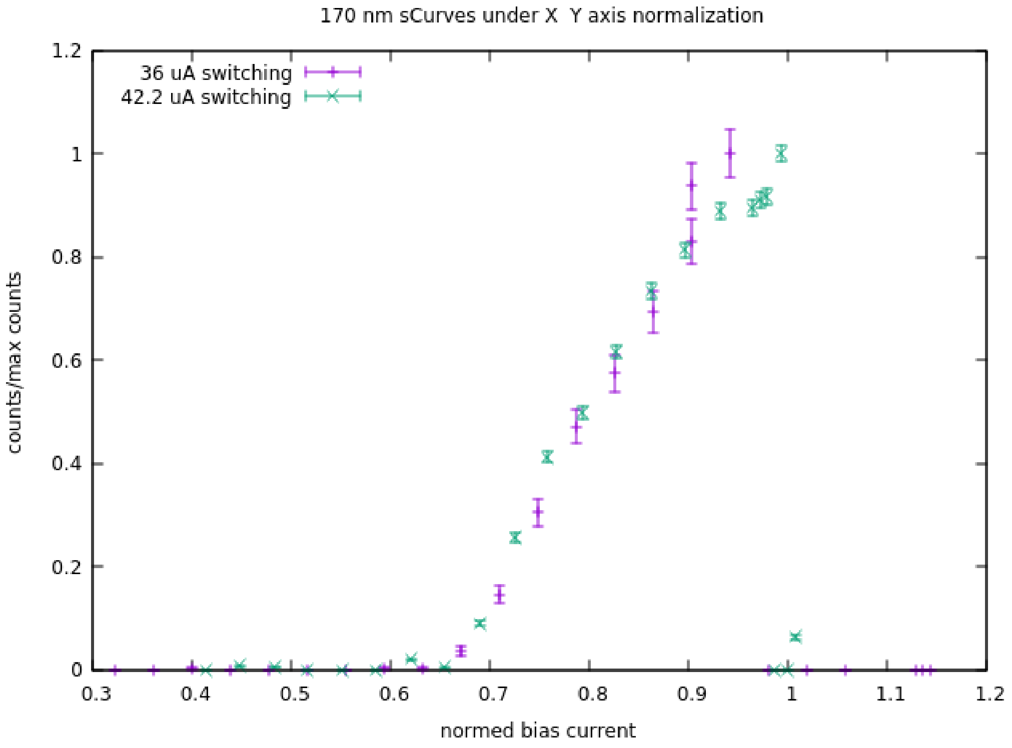

(b) 170 nm calibration curve for MoSi SNSPI.

SNSPI switches from it superconducing to its normal state (an example can be seen in Fig. 3.14), which is recorded as the switching current I_s. This switching current is then associated with an average temperature $\langle T \rangle$ during the I-V ramp. After many

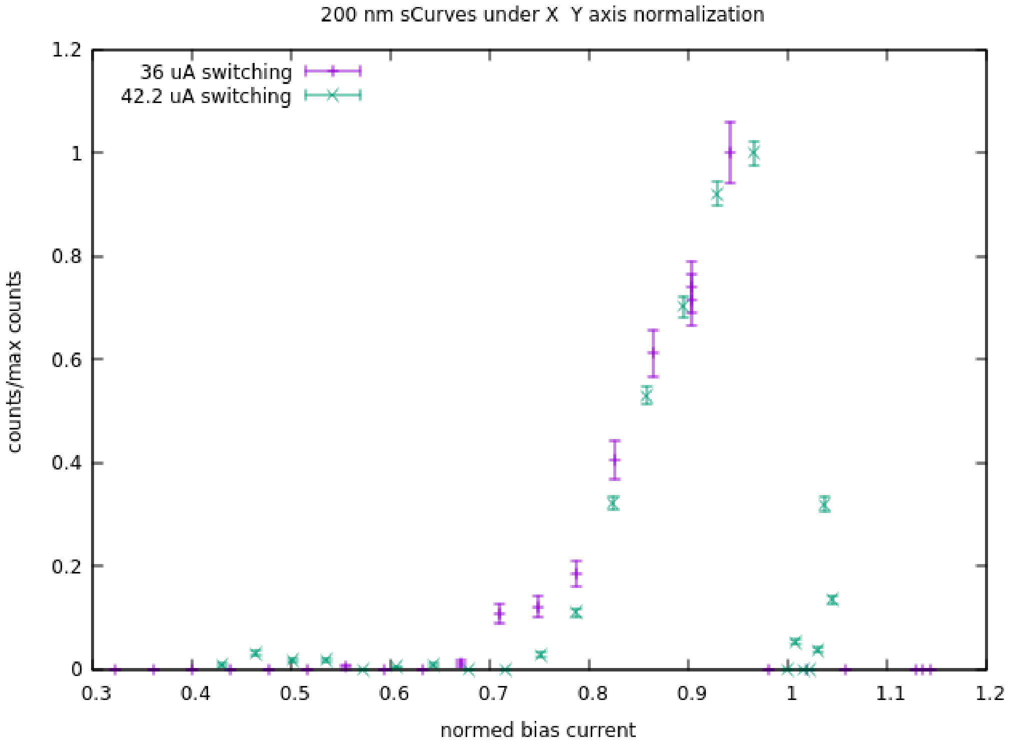

(c) 200 nm calibration curve for MoSi SNSPI.

Figure 3.11: Calibration curves taken for SNSPI pixel using a D2 lamp monochromator. The x-axes is the bias current normalized by the switching current (listed in legend). The y-axis was normalized to the maximum counts achieved for each wavelength.

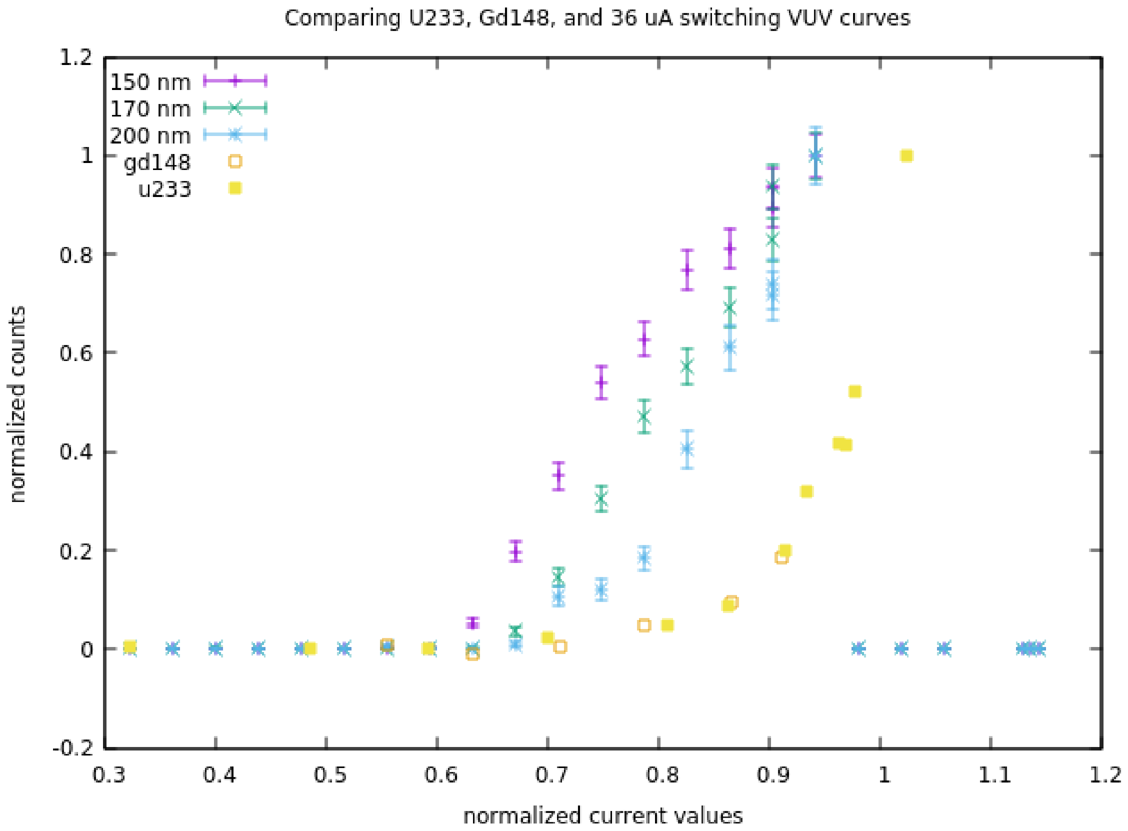

Figure 3.12: ^{233}U and ^{148}Gd curves in comparison to the calibration curves for the MoSi SNSPI. The the curves were normalized on the x-axis by dividing the bias current by the switching current.

repetitions and by varying the heater tape power, a large set of I_s and $\langle T \rangle$ pairs is obtained. One may then perform a linear regression between I_s and $\langle T \rangle$ (it should be noted that this regression was performed with ordinary least squares) in order to obtain a calibration $\hat{I}_s = a\langle T \rangle + b$. The utility of this is because there are many more temperature measurements for a data point than there are switching current measurements, making $\langle T \rangle$ a much more well known quantity for for a data point in a VUV calibration or experiment data point. The result of a linear regression can be seen in Fig. 3.15.

The next step was to perform VUV calibrations at different temperatures. As can be seen in Fig. 3.16, the calibration curves for a single energy of VUV photon were shifted with respect to each other. However for each calibration 30 s datapoint $\langle T \rangle$was recorded, and the bias current could be normalized as $i = I_b/\hat{I}_s$, where $\hat{I}_s$ was given by the regression. All of the calibration curves for a single VUV photon energy

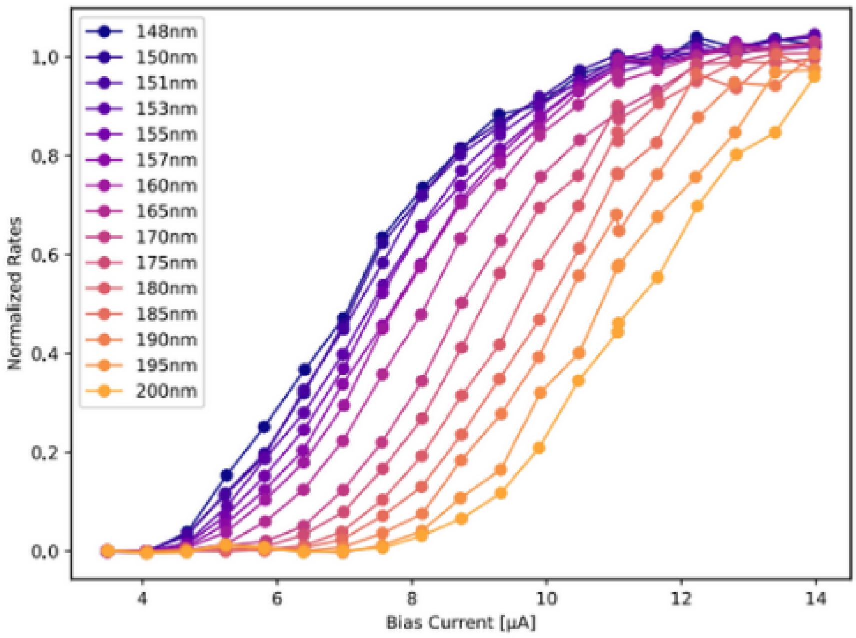

Figure 3.13: VUV calibration curves taken with the WSi SNSPI. All count rates were normalized to their maxima.

could then be plotted on a normalized x-axis with this procedure. The result can be seen in Fig. 3.17, where it is clear that the procedures aligns the 160nm curves quite well. This capability to correct the calibration curves for temperatures shifts was tested at 150nm, 160nm, and 170nm, and it was demonstrated that to within experimental uncertainty the shift of calibration curves due to temperature could be accounted for by this "switching current normalization" across a wide range of photon energies (this can be seen in Fig. 3.18, where the normalized error function centers of different wavelengths is plotted against temperature).

A natural question would then be under what conditions this sort of calibration would hold. VUV calibration sets had been taken were spaced out by several months. When the switching current normalization was performed on the VUV calibration curves, it was found that even if they were spaced out by roughly a month, they could be made to agree on the normalized axis so long as the cryostat was not temperature cycled (i.e. the system was warmed up and cooled down). However, if the system was temperature cycled, the calibration curves would shift to new positions on the normalized axis (this can be seen for 160nm data in Fig. 3.19). Evidently, other properties that governed the VUV response of the SNSPI changed upon warm-up

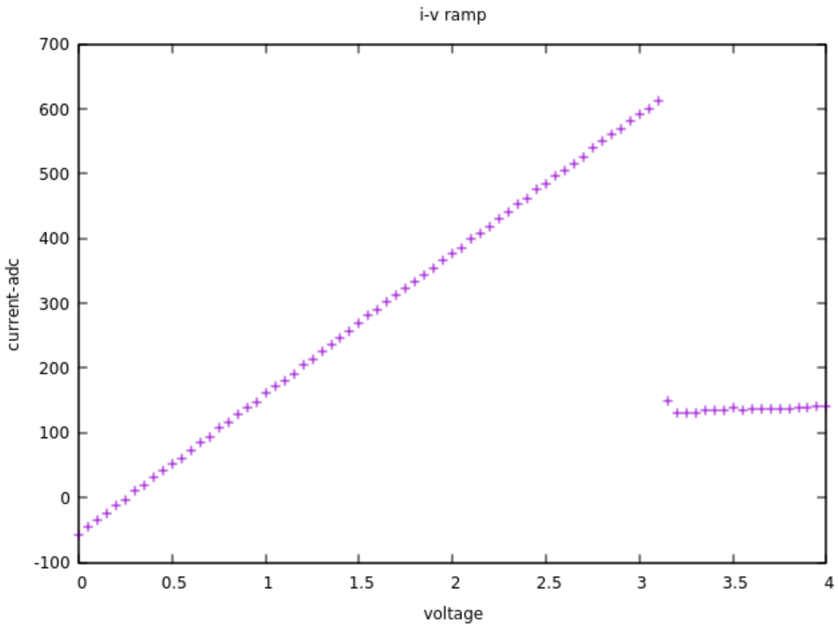

Figure 3.14: Example I-V curve of the current flowing through the SNSPI circuit versus the applied voltage from the bias power supply.

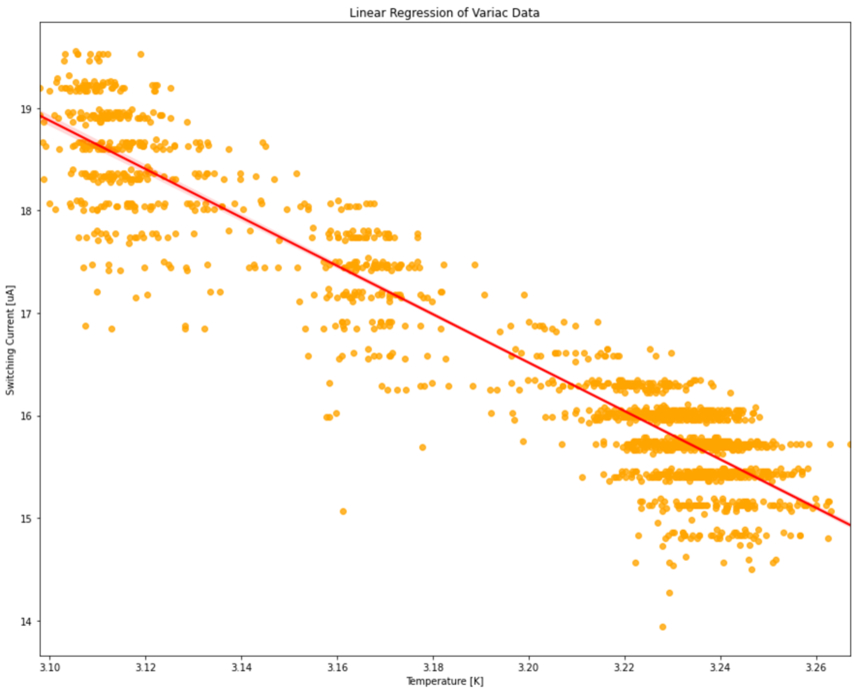

Figure 3.15: Results of the testing with the Variac powered heater tape. The red line represents an ordinary least squares linear regression.

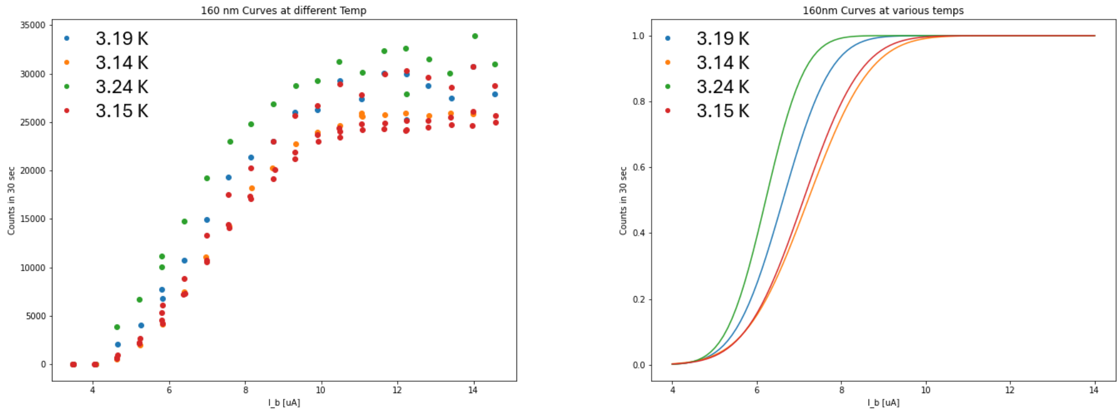

Figure 3.16: 160nm VUV calibration curves at different temperatures. The left plot shows the raw data, while the right plot shows the fitted error functions. Note the centers of the error functions are not aligned.

and cool-down.

From these tests, it was concluded switching current normalization could correct temperature drifts, and data runs would be divided by instances of temperature cycling. In fact, it was these instances of temperature cycling that would come to define the data runs. For each 1000s datapoint looking for the IC decay, an initial I-V curve would record a single measurement of I_s, and $\langle T \rangle$ would be recorded. After several months, the regression between I_s and $\langle T \rangle$ could be performed, and all datapoints within a datarun could then be plotted on a normalized bias current axis.

3.6.2 Extracting the Correlated Events

At the simplest level one is looking for events that occur within a [2,20] μs IAT window (henceforth called the IC time window) and binning them according to their relative position. In reality acquiring the events requires filtering out noise waveforms and correcting the time tags that the data acquisition system assigns. We will leave the details to Sec. 3.9.1,3.9.4 in the appendix, and assume we can acquire the waveforms and assign them correct time tags. We can then associate time tags sufficiently close in time coming from opposite sides of the nanowire as coming from the same "click", and treat the time difference between the pulses, δt, as the position. The observable we then care about is the position difference Δ between subsequent

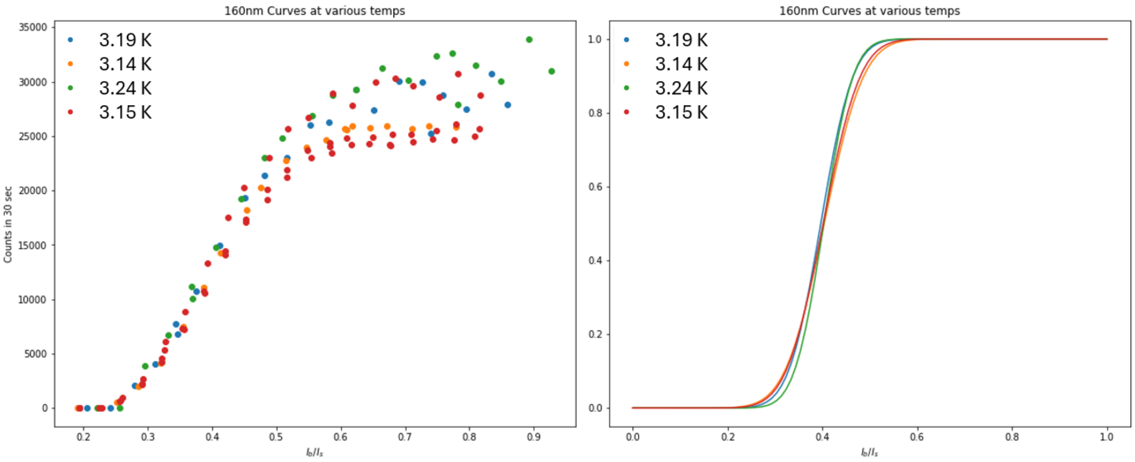

Figure 3.17: 160nm VUV calibration curves at different temperatures with the $\hat{I}_s$ normalization. The left plot shows the raw data, while the right plot shows the fitted error functions. Note the centers of the error functions are in agreement.

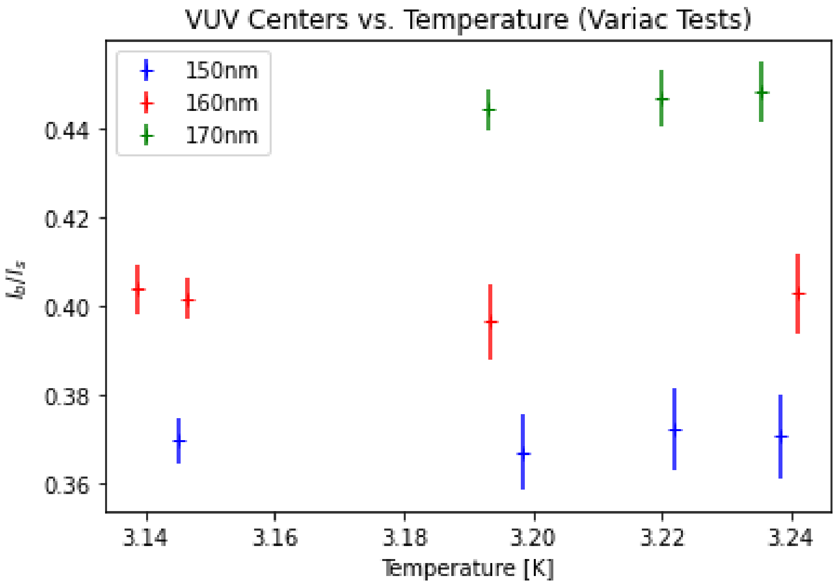

Figure 3.18: The center of the calibration curves on the normalized $I_b/\hat{I}_s$ scale for different photon energies as a function of temperature.

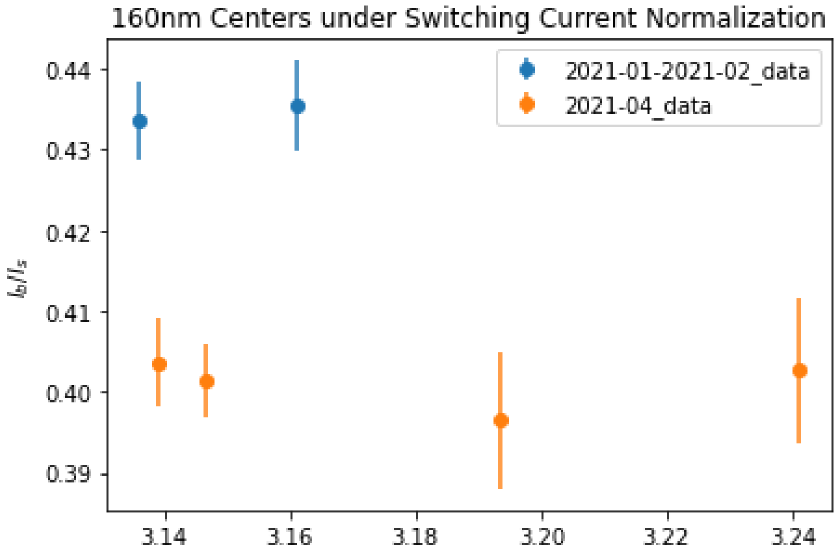

Figure 3.19: 160 nm calibration curves taken at different times. The blue points represent data taken in 2021-01 and 2021-02, while the orange represent data taken in 2021-04. Between the blue and orange points an instance of temperature cycling had occured.

events given by

$$\mathit{\Delta} = \delta t_1 - \delta t_2. \tag{3.4}$$

If the events within the time window are uncorrelated (i.e. from random impacts of particles on the nanowire) the distribution of $\mathit{\Delta}$ should be broad and triangular (see Sec. 3.9.2 in the appendix for details). We will refer to this uncorrelated distribution as $U(\mathit{\Delta})$. On the other hand, for the isomeric transition $\mathit{\Delta} = 0$ as the ^{229}Th will de-excite at the same position in the nanowire that it landed. In reality, due to smearing of the time stamps due to the internal clock jitter of the CAEN acquisition system, the true single-site distribution $\alpha(\mathit{\Delta})$ is given by

$$C(\mathit{\Delta}) = \begin{cases} 1/6 & \mathit{\Delta} = -1, \\ 2/3 & \mathit{\Delta} = 0, \\ 1/6 & \mathit{\Delta} = 1. \end{cases} \tag{3.5}$$

The derivation of Eq. 3.5 is left to the Appendix Sec. 3.9.3.

For a particular value of $i = I_b/I_s$ and observation time T, we would observe some number of event N_{obs} within the IC time window. There would be some number of events α (β) associated with the correlated (uncorrelated) distribution $C(\mathit{\Delta})$ ($U(\mathit{\Delta})$).

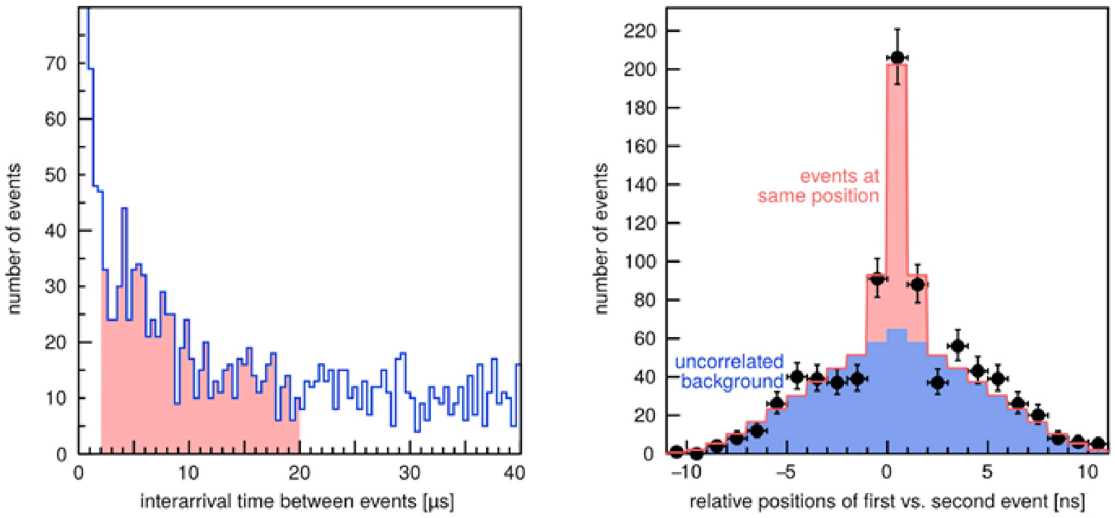

Figure 3.20: (Left) [2,20] μs IAT cut used to look for correlated IC events. Note the buildup of events at very short timescales (likely due to a large electron background). (Right) The extended MLE fits of the correlated (red) and uncorrelated (blue) distributions integrated over the time cut.

In other words, we have two distributions with unique shapes, which together yield a Poisson fluctuating total of events. For such a case, we apply the extended maximum likelihood estimation (extended MLE) method [Bar90], and seek to maximize the likelihood.

$$\mathcal{L}\left(\alpha, \beta; \mathcal{N}\right) = \frac{e^{-\mathcal{N}}\mathcal{N}^{N_{obs}}}{N_{obs}!} \times \prod_{\mathit{\Delta}} \left(\frac{\alpha C(\mathit{\Delta})}{\mathcal{N}} + \frac{\beta U(\mathit{\Delta})}{\mathcal{N}}\right), \tag{3.6}$$

where $\mathcal{N}$ is the expected total number of events and the product runs over all position differences $\mathit{\Delta}$ possible on the nanowire. Intuitively Eq. 3.6 has two major components, the right product component which describes the probability of events from the correlated and uncorrelated distributions being distributed into the $\mathit{\Delta}$ bin, and the left part which describes the fact that the normalization $\mathcal{N}$ is itself a Poisson fluctuating variable. As is customary in likelihood analysis, instead of maximizing 3.6, we instead minimize the negative log-likelihood,

$$-\ln(\mathcal{L}) = \mathcal{N} - N_{obs}\ln(\mathcal{N}) - \sum_{\mathit{\Delta}} \ln\left(\frac{\alpha C(\mathit{\Delta})}{\mathcal{N}} + \frac{\beta U(\mathit{\Delta})}{\mathcal{N}}\right) \tag{3.7}$$

where we have dropped the $\ln(N_{obs}!)$ term as an overall constant. By minimizing this equation we can obtain a point estimate of the number of correlated and uncorrelated evenets $(\hat{\alpha}, \hat{\beta})$. However, we really only care about the number of correlated events, and moreover we care about the rate λ that they occur. The relation between α, the observation time T, and λ is given by $\alpha = \lambda T$, and so our point estimate $\hat{\alpha}$ can be turned into a point esimtate $\hat{\lambda}$. We assume that λ is some function of the normalized bias current, and that for the next stage of model fitting we will need not only point estimates $\hat{\lambda}$ of the correlated event rate, but also a likelihood function associated to that normalized bias current value.

The method of profiling out the nuisance parameters $(\beta, \mathcal{N})$ from Eq. 3.7 generates the function we need. The method of profiling involves the repeated optimization of Eq. 3.7 by fixing $\alpha = \alpha'$ and varying only β and $\mathcal{N}$. Whatever the minimum value $-\ln(\mathcal{L})|_{\alpha=\alpha'}$ achieved for that fixed value of α' is recorded, and a new function is generated from the $(\alpha', -\ln(\mathcal{L})|_{\alpha'})$ pairs. It is this new function that is known as a profile likelihood. Note that the minimum of this function is identical to that of the original log-likelihood. This profile likelihood is a standard technique for estimates of a confidence interval on a parameter of interest in the presence of nuisance parameters. Finally, we may divide α' by the observation time T to obtain a profiled log-likelihood function for our rate parameter λ. An example of a profiled log-likelihood function can be seen in Fig. 3.21. The set of all of these profiled log-likelihoods for each measured normalized current i will be denoted as $\ln(\mathcal{L}_i(\lambda))$. In our analysis, both the extended MLE and profiling of histograms were done using the RooFit v6.14 software within the CERN ROOT analysis package.

3.6.3 Fitting the Correlated Event Rate

With our profile likelihood procedure we could now obtain a set of log-likelihood functions for the correlated event rate (CER) λ corresponding to each measured normalized bias current i. Ideally, the CER would only be composed only of events corresponding to the internal conversion of the isomer, however testing with ^{148}Gd again showed that despite all of the temporal and spatial filtering, there was still a background of correlated events that increased with bias current (see Fig. 3.22). Our goal now was to model this background so that it could be accounted for in our ^{233}U data.

A model for the correlated event rate consists of a function $f(i, \boldsymbol{\xi})$, where i is the normalized bias current and $\boldsymbol{\xi}$ are a set of parameters that yields a prediction

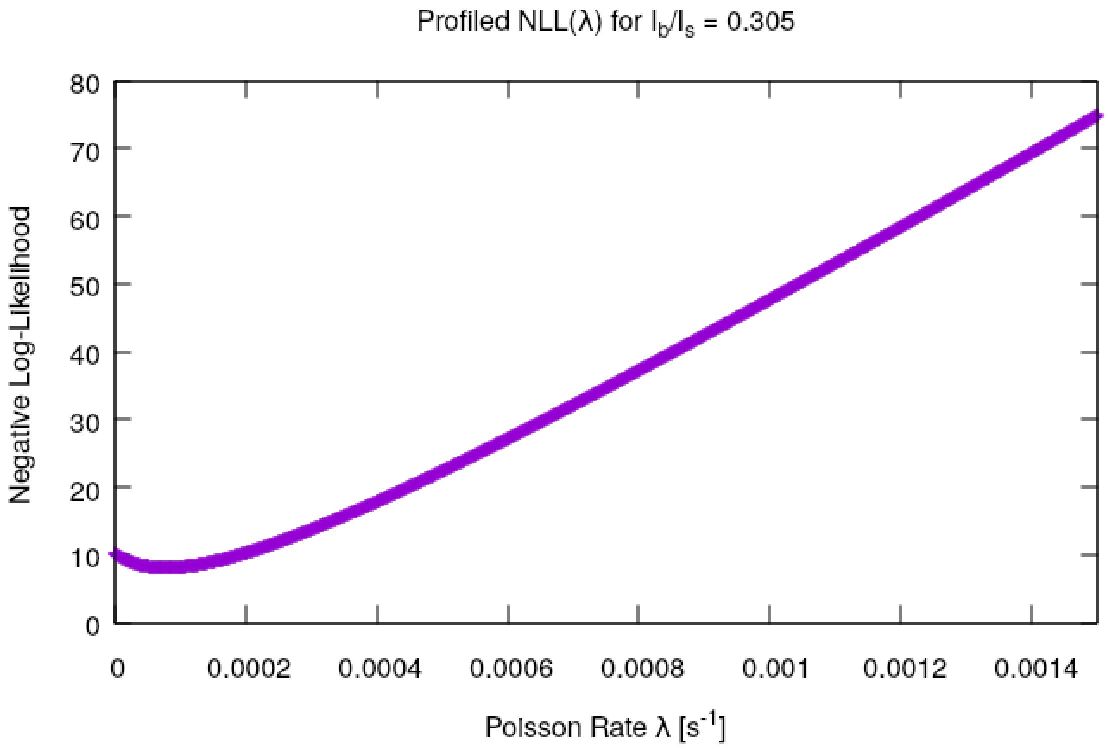

Figure 3.21: An example of a profiled likelihood function for the correlated event rate λ for a normalized bias current value of $i = 0.305$.

for λ at that value of i. Our task of model fitting then becomes minimizing a larger likelihood object given by

$$-2\ln\mathcal{L}_{tot}(\boldsymbol{\xi}|f) = \sum_i -2\ln(\mathcal{L}_i(f(i,\boldsymbol{\xi}))), \tag{3.8}$$

where the factor of 2 is inserted as a customary scaling in the statistics literature.

Before continuing, it is useful to consider what we are doing. Within the language of statistics, we are performing a form of MLE in order to yield point estimates of $\boldsymbol{\xi}$, and we will find this will yield both confidence intervals for the parameters as well as a goodness-of-fit test. To give a physical analogy, we can imagine the function $f(i,\boldsymbol{\xi})$ as describing a curve and each of the profile log-likelihoods $-2\ln(\mathcal{L}_i)$ as describing a potential energy curve at a set of points i. Then what we are doing is equivalent to finding the shape of a curve $f(x,\boldsymbol{\xi})$ that is subject to forces $\partial_\lambda\mathcal{L}_i$ at a discrete set of points.

A few more concepts we will need are the *deviance* and the *saturated model*. If one has N observations at values $\{i_1, i_2, ..., i_N\}$ that yields N results $\{y_1, y_2, ..., y_N\}$, the saturated model is the model f_{sat} that has N parameters such that $f_{sat}(i_k) = y_k$ for $k \in \{1..N\}$. In other words, it is the "fill-in-the-dots" model where one makes it exactly equal to the data at each observation point. It is interesting to note a general principle that with N parameters, the model can fit perfectly, but will yield no physical explanation/insight. We can throw the saturated model into Eq. 3.8 to get a value of $-2 \ln \mathcal{L}_{tot}(f_{sat})$, and obtain the lowest possible value of any model. The total deviance of any other model f_0 is given by

$$D(f_0) \equiv 2(\ln \mathcal{L}_{tot}(f_{sat}) - \ln \mathcal{L}_{tot}(f_0)), \tag{3.9}$$

where it is implied that f_0 is MLE optimized. It is useful to re-define all of our $\ln(\mathcal{L}_i)$ so that they are zero at their minima, making $D(f_{sat}) = 0$. This is allowed since shifting each log-likelihood by a constant does not affect the minimization problem. For a more physical picture, shifting the potential energy at point by an overall constant does not change the physical forces on the curve, and hence cannot influence the final shape it takes. With this renormalization in place, the deviance now becomes a metric of goodness-of-fit [BC84]. In fact, in the special case where the likelihood functions describe N gaussian distributed data points, the deviance of a model that well describes the data will have

$$\langle D(f_0) \rangle \approx N - p \tag{3.10}$$

where p is the number of degrees of freedom (if this seems familiar then one should work out that this is equivalent to the method of reduced χ^2). For more general types of data, one has to peform Monte Carlo simulations to turn the deviance into a goodness-of-fit metric. This involves generating fake datasets based on the model f_0 that varies according the probability distribution the data generating process is thought to follow (Poissonian, negative binomial, etc.), perform the MLE fit of f_0, and then look at the the resulting distribution of $D(f_0)$. If the experimental value of $D(f_0)$ is far in the tail of the distribution, then the fit is considered poor (or you have poorly captured the behavior of your data).

3.6.4 Fitting the ^{148}Gd Data

The increase of the CER in the ^{148}Gd data was unexpected, but physically we reasoned that it may be due to some sort of low energy process in the nanowire. This was evidenced by the fact that it seemed to increase without saturation as the bias current was increased. Since ^{148}Gd exposes the nanowire to the same environment as ^{233}U , minus the IC decays, we assumed that if we could model this background properly we could "subtract" it from the ^{233}U data in order to observe a signal from the isomer. As such, we turned to the familiar expression of Eq. 3.2, and decided to model the background as

$$f_1(i) = \lambda_b \operatorname{erf}\left((i - i_c^b)/\sigma_b\right), \tag{3.11}$$

where $(\lambda_b, i_c^b, \sigma_b)$ were free parameters to be optimized. This came to be called the single erf model. All point estimates of parameters were obtained in this and later sections through minimizing Eq. 3.9 using an ensemble version of the Nelder-Mead method (details can be found in Sec. 3.9.5 of the appendix). To see if this model captured the data, we computed $D(f_1)$ and obtained values on the order of $\approx$120. From Monte Carlo analysis this was found to be the average deviance one would expect if the data truly did follow f_1 over the normalized bias currents that were measured. From this we felt assured that f_1 could capture the CER background.

3.6.5 Fitting the ^{233}U Data

There were two separate data runs that were analyzed. Run 1 spanned from 2020-12-22 to 2021-04-02. Run 2 spanned from 2021-05-21 to 2021-06-28. In between these two runs there is a gap of "bad" data due to a variety of reasons. Run 1 was terminated due to a warm-up cycle as we attempted to add voltage biasing between the nanowire and the ^{233}U source. This biasing ended up being unhelpful, as we ended up attracting so many exoelectrons towards the SNSPI that it overwhelmed any correlated signal. Then in April the nanowire underwent some sort of "phase transition" where both the total measured radioactivity dropped and the calibration curves shifted within the span of a day. It is unclear if this was sign of long term radioactive damage after being near continuously exposed to radioactivity for $\sim$5

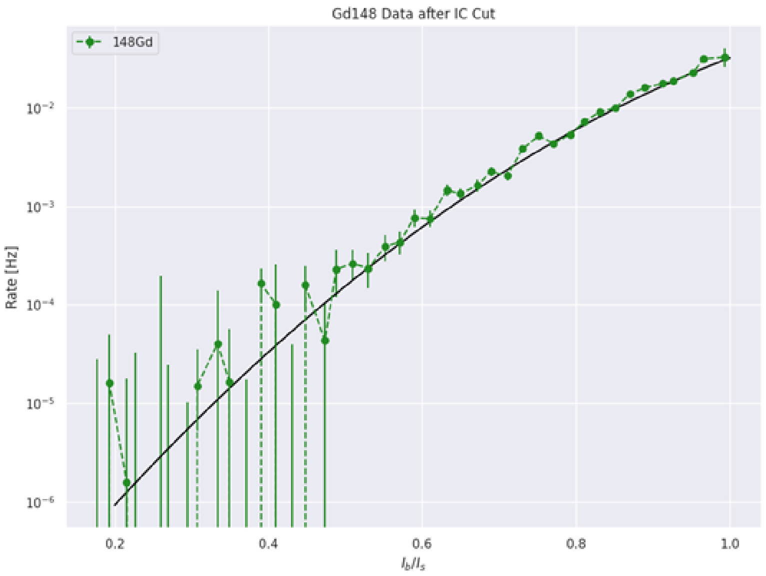

Figure 3.22: ^{148}Gd data after filtering for temporally and spatially correlated events. Black curve shows the optimized single error function. The y-errorbars represent $\varDelta \ln \mathcal{L} = \frac{1}{2}$ from the minimum value of λ for each datapoint.

months. In any case we wanted to now try to fit our ^{233}U data for the IC signal and demonstrate that there was some significant difference between the ^{233}U and ^{148}Gd data.

For the ^{233}U data we adopted the following double erf model

$$f_2(i) = \lambda_{IC}\,\mathrm{erf}\left((i - i_c^{IC})/\sigma_{IC}\right) + \lambda_b\,\mathrm{erf}\left((i - i_c^{b})/\sigma_b\right) \tag{3.12}$$

where the first term is supposed to represent the sigmoidal curve (s-curve) corresponding to the correlated events due to ^{229}Th IC decays, while the second term is the CER background. If the value of $D(f_2)$ is much smaller than $D(f_1)$, then one can say that $D(f_2)$ is a better fit, and hence there is some sort of signal in the ^{233}U data. A more typical way to state this is that we are computing the log-likelihood ratio (abrr. $2\ln\varLambda$) between the models

$$-2\ln\varLambda = -2(\ln\mathcal{L}_{tot}(\boldsymbol{\xi}_1|f_1) - \ln\mathcal{L}_{tot}(\boldsymbol{\xi}_2|f_2)) \tag{3.13}$$

and claiming discovery when $-2\ln\varLambda$ is large. How much larger $2\ln\varLambda$ needs to be is a now the question.

Our background model f_1 is actually nested in f_2, in the sense that one can turn $f_2(i, \boldsymbol{\xi}_2)$ into any possible realization of $f_1(i, \boldsymbol{\xi}_1)$ by appropriate choice of the parameters (the possible values of $\boldsymbol{\xi}_1$ are in a subspace within the parameter space of $\boldsymbol{\xi}_2$). The f_1 is said to be *nested* inside f_2. One then can invoke Wilks's theorem [Wil38], which states:

Theorem (Wilks). *Let the model f_1 be nested within f_2 and let $\ln \Lambda$ be the log-likelihood ratio between them. If certain regularity conditions are met, then as the number of observation $N \to \infty$, -$2\ln \Lambda$ will follow the distribution χ^2_ν where ν is the difference in the number of parameters between f_2 and f_1.*

The power of the theorem comes from the fact that significance can now be determined by looking at the value of the $-2\ln \Lambda$ and comparing it to standard tables of χ^2_ν and reporting a p-value. The first issue with this approach is that it is an asymptotic result, and only holds in the limit that the number of observations approaches infinity (though this is a common issue of most statistical tests). The more troubling issue comes from the fact that the regularity conditions are fairly restrictive, and many experimental situations can fail to meet them. The full formal discussion of the regularity criteria can get quite technical, but an important set for most experimentalists [AAM20] is as follows:

1. The parameters of f_1 and f_2 should be in the interior of the the parameter space. This fails for example if some parameter $\mu = 0$ for f_1 and is only allowed to be non-negative ($\mu \geq 0$) in f_2.
2. Each value of the parameters yields a unique model. This fails for example in the case of a gaussian with both varying amplitude A and center μ. In this case there would be many degenerate models for different μ but $A = 0$.
3. There are no other backgrounds or variations in the data that are not captured by the models. This point, while somewhat obvious, should always be remembered when analyzing one's data.

For our double erf model, both points 1 and 2 are not met. Our erf and Poisson rate is only defined if the rate amplitudes $(\lambda_{IC}, \lambda_b)$ are positive, there is a subspace of degenerate models if one of the error function amplitudes is zero but the centers can vary, and the model has a symmetry under the interchange of the sets of (λ, i_c, σ).

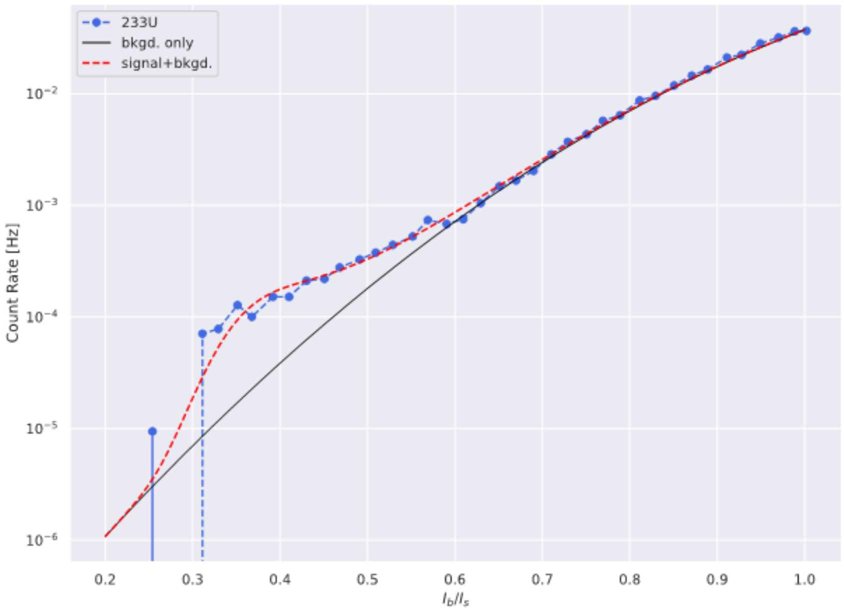

Figure 3.23: ^{233}U data after filtering for temporally and spatially correlated events. Red curve shows optimized double error function. Black curve shows the background single error function component. The points represent the point-wise maximum likelihood estimates of λ.

Thus, while we will still use Λ as our test statistic to look for a significant signal, we cannot make very strong statements without performing a Monte Carlo analysis. Occasionally, in these sorts of non-regular conditions, $-2\ln\Lambda$ may follow a linear combination of χ^2_k distributions where k runs over multiple degrees of freedom [Sha85, Sha88]. In general, the coefficients of such linear combination can only be obtained by simulation of the experiment.

For both Run 1 and Run 2, we optimized both the single and double erf models on the data, and obtained what looked like a significant departure from the single erf model in the data (an example can be see in Fig. 3.23). The difference in log-likelihood between the two models -2Λ was computed for each run. To obtain a significance level, the following Monte Carlo simulation was carried out:

1. Let $f_1(i, \hat{\boldsymbol{\xi}})$ be the fitted single erf model for the data set (Run 1 or Run 2) where $\hat{\boldsymbol{\xi}}$ were the fitted parameters. Each measured bias current i was measured for some total time T_i, so draw a Poisson random variable number of events N_i for each bias current given by $f_1(i, \hat{\boldsymbol{\xi}}) \times T_i$. **We are assuming the null**

hypothesis that only the single erf model describes our data.

2. For each bias current, compute the Poisson negative log-likelihood function for the rate parameter λ given one observed N_i events. This will be given by $-\ln \mathcal{L}(\lambda|N_i) = \lambda T_i - N_i \ln(\lambda T_i) + \ln(N_i!)$. It should be noted that this is an approximation, the profile likelihood function the correlated events follow deviates from a true Poisson distribution.
3. Optimize f_1 and f_2 on this dataset, and compute the likelihood ratio $-2\ln\Lambda$ between them.
4. Repeat many times to obtain a histogram of $-2\ln\Lambda$ over many simulated experiments (see Fig. 3.24).

We may then use this empiric distribution of $-2\ln\Lambda$ under the null hypothesis and compute the p-value of seeing the value observed in the experiment [Nav15]. For both Run 1 and Run 2, the p-value obtained was $\lesssim 4\times10^{-4}$ (or more than 4σ). Thus we had strong evidence to reject the single erf model and a statistically significant signal had been found. A similar procedure was conducted on the ^{148}Gd data, and no statistically significant signal was observed.

3.6.6 Analysis of the Signal

Once we had determined that there was indeed a signal in the ^{233}U data, we wanted to study its lifetime. Using double erf fit to the ^{233}U data, we could integrate all the datasets that were in the region where the IC signal was present, but the CER background was still small (see Fig. 3.25). Upon doing so we could construct an integrated IAT histogram for both the ^{233}U and ^{148}Gd datasets and compare their features.

As one can see in Fig. 3.26, the ^{233}U data had a distinct microsecond timescale decay feature that was absent from the ^{148}Gd data. In many ways this was seen as the smoking gun of the presence of a signal in the ^{233}U data. The rate of the signal clicks was found to be (1-2)$\times10^{-4}$ Hz and the lifetime was found to be $\tau = 5.3(6)$ μs.

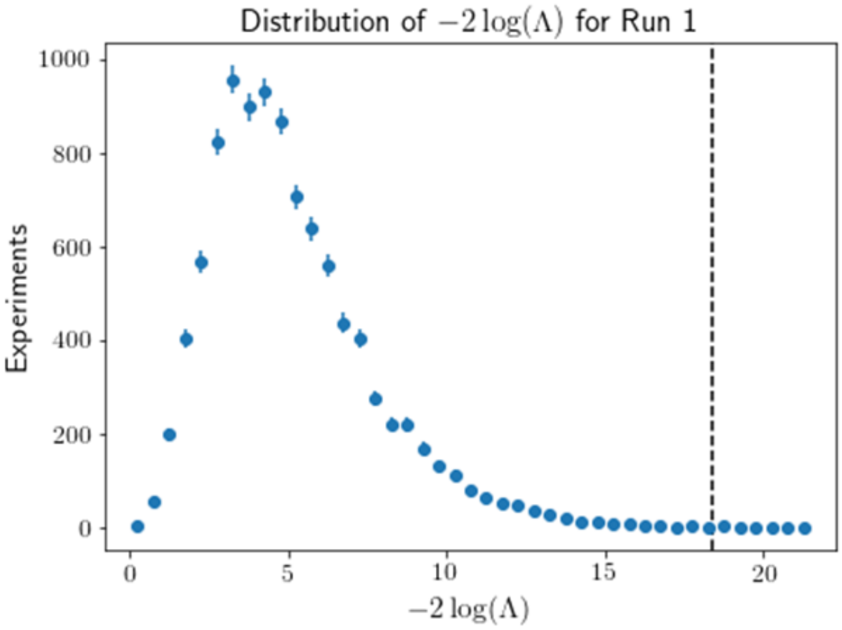

Figure 3.24: Distribution of likelihood ratios obtained by Monte Carlo simulation. The black line represent the actual obtained value of $-2\ln \Lambda = 18.8$ in Run 1. This corresponds to a p-value of $\approx 4 \times 10^{-4}$ for the null hypothesis.

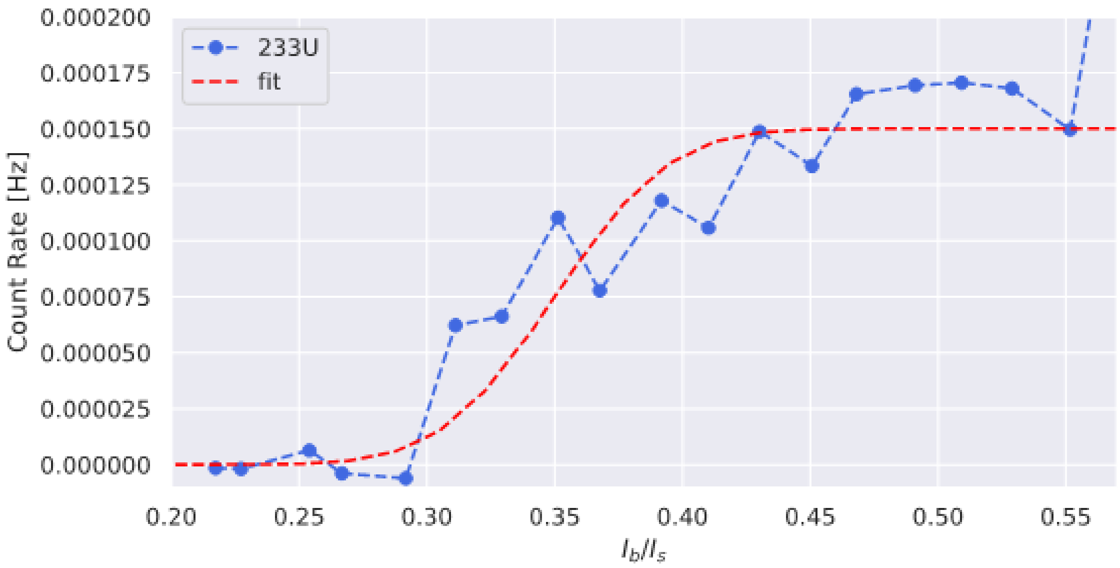

Figure 3.25: The IC signal observed in the ^{233}U data with the background CER curve subtracted off. The plots in Fig. 3.26 we made by integrating over the normalized bias currents in this plot.

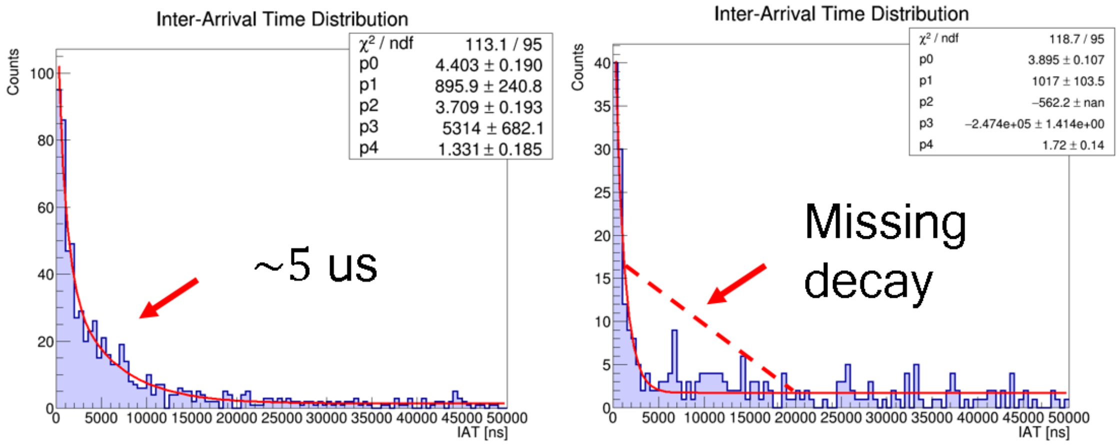

Figure 3.26: Comparison of the IAT histograms for the two datasets within the IC signal region. Fits were performed using binned log-likelihood in ROOT package with double exponential model. (Left) ^{233}U data showing ~5 μs decay. (Right) ^{148}Gd data showing lack of similar feature.

3.7 ^{213}Po Studies

While it was odd that the signal rate was a factor of 10x lower than expected, and the lifetime seemed shorter than the values for the isomeric decay observed previously, this did not raise any alarm until it was realized that the 5.3(6) μs lifetime was very close to the known lifetime of 5.37 μs for one of the ^{233}U decay daughters, ^{213}Po [Bas07]. The process shown in Fig. 3.27 could have been masquerading as the isomeric decay, where a high energy β decay from ^{213}Bi causes a click, and 5 μs later an α decay of the leftover ^{213}Po ion generates a large number of phonons that also cause a breakdown of the superconducting state. Since most of the kinetic energy is carried away by the β particle, the recoil ^{213}Po would stay in the same spot, so it would look like temporally and spatially correlated decays. To make matters worse, none of these daughters would have to actually be in the nanowire, since it is known that superconducting devices can detect radioactive decays in the substrate purely through the generated phonons [YCL90].

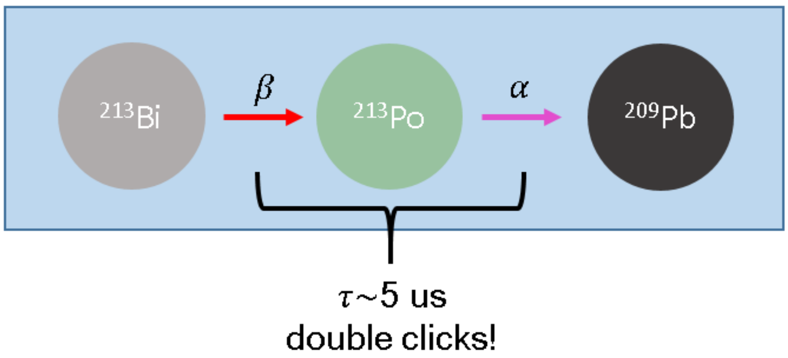

Figure 3.27: The ^{213}Po process that would mimic the IC decay from ^{229}Th. Most of the kinetic energy in the β decay is carried away by the electron, so the β and α decays would occur in the same location.

3.7.1 Simulation of ^{213}Po Decays

Our task was to determine the amount of these ^{213}Po decays we could expect in the nanowire during the experiment via simulation. Due to the large discrepancy in time scales between all of the decay daughters as well as the possibly for recoil ions to leave the nanowire, the problem was divided into a particle transport simulations and rate equations.

3.7.1.1 Modeling of the ^{233}U Target

The first task was to determine the probability with which ^{229}Th and its daughters escaped the ≈20 nm thick layer of UO_2 that made up the ^{233}U target. SRIM was used to determine whether the ions were stopped by the oxide layer. For the first SRIM simulation, ^{229}Th could be started the UO_2 layer which was assumed to be a uniform cylindrical distribution. Each ion was started in a random direction within the full 4π solid angle. Some ^{229}Th would leave the ^{233}U target, forming the escape fraction, while the remaining would terminate in the target at some final position, forming the embedded fraction. The spatial distribution of the embedded fraction would then be fed into a text file that could be fed back into SRIM to form the starting positions for the next daughter isotope (in this case ^{225}Ra). The distribution of the daughters in terms of the depth into the target can be seen in Fig. 3.28. This process would repeat until we terminated at the daughter ^{213}Bi which immediately precedes ^{213}Po.

For each isotope in the decay chain, the kinetic energy of the ions was calculated using the dominant decay mode.

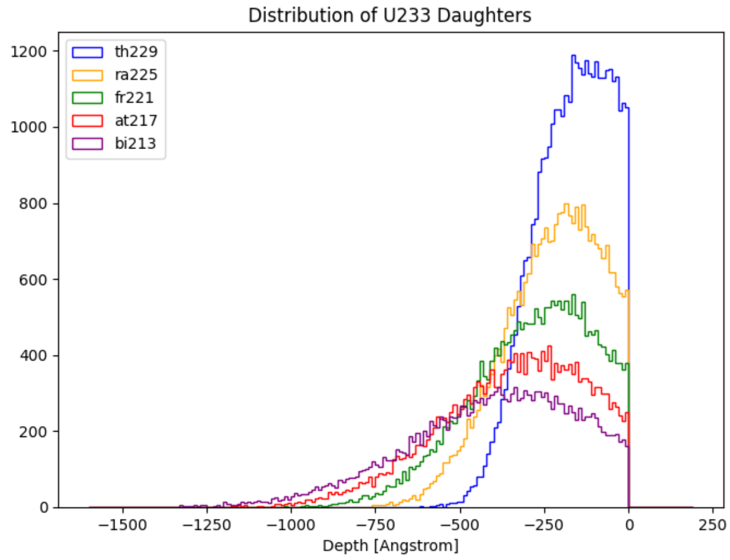

Figure 3.28: Depth distribution of the daughters of ^{233}U in the UO$_2$ layer.

We then calculated the radioactivity of the source using the following Bateman equations:

$$\begin{aligned} \dot{N}_{Th} &= \mathit{\Gamma} - \lambda_{Th} N_{Th} \\ \dot{N}_{Ra} &= (1 - p^{esc}_{Th})\lambda_{Th} N_{Th} - \lambda_{Ra} N_{Ra} \\ &\vdots \\ \dot{N}_{Po} &= (1 - p^{esc}_{Bi})\lambda_{Bi} N_{Bi} - \lambda_{Po} N_{Po}, \end{aligned} \tag{3.14}$$

where $\mathit{\Gamma}$=20 kBq is the ^{233}U source activity and N_k, p^{esc}_k, and λ_k are the population, escape probability, and decay rate of the k-th isotope, respectively.

A ray tracing simulation was used to determine the probability that an ion would actually make it to the SNSPI nanowire region, given that it had escaped from the source. Given the target radius, distance, and fill fraction of the nanowire, this was found to be $\eta = 1.2 \times 10^{-5}$. Thus, the rate Q_k at which the k-th isotope arrived from the target to the nanowire was given by

$$Q_k = \eta p^{esc}_k \lambda_k N_k, \tag{3.15}$$

except for the ^{229}Th which was given by $Q_{Th} = b\eta p^{esc}_{Th}\mathit{\Gamma}$.

3.7.1.2 Modeling Isotope Populations in the SNSPI

If a daughter ^{233}U actually makes it into the SNSPI it will continue to undergo recoils as it decays, potentially allowing it to scatter out of the surface or deeper into the substrate (Fig. 3.29 gives as schematic of the potential processes). We denote the probability that a particle of the k-th isotope is detectable by the SNSPI by the time it reaches ^{213}Po as the survival probability p_k^{sur}. If the decaying particle embeds itself deep enough into the substrate, it may no longer be detectable by the SNSPI. This maximum detectable depth d is not known a priori and therefore is a free parameter in calculating p_k^{sur} (we will sometimes refer to d as the "cut-plane depth").

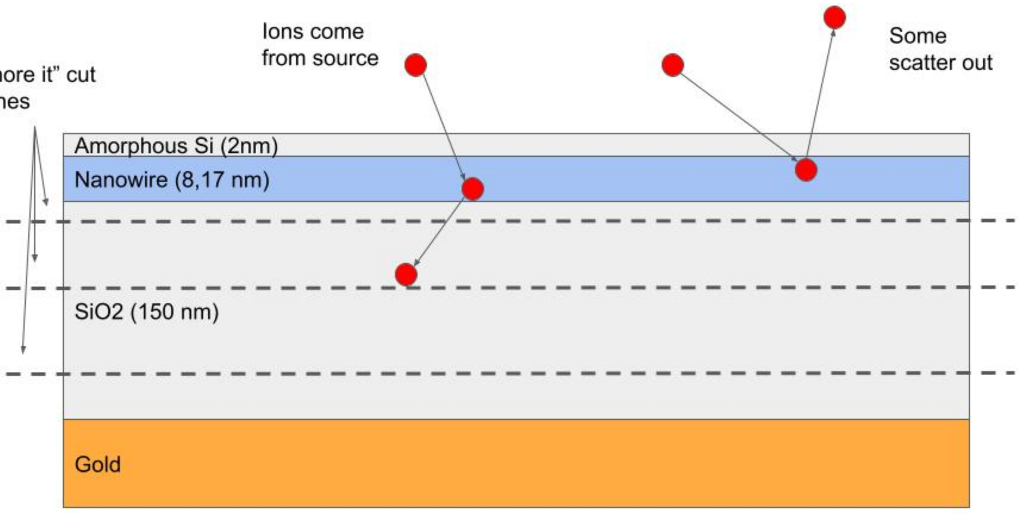

Figure 3.29: The layers in the SRIM simulation. Ions would impact from the source and then be evolved through stages of radioactive decay. Some would backscatter out of the system and be removed from the simulation. The dotted lines show possible choices of the maximum depth that a final ^{213}Po decay could be detected.

To calculate p_k^{sur} we performed a separate chain of SRIM simulations for each daughter emitted from the source. The simulation procedure was as follows:

1. Start with 3000 ions of the k-th isotope from the source, have them impact the SNSPI. From the solid angle study the distribution of velocity vectors that actually make it to the SNSPI is known, so draw from that distribution.
2. Obtain the final positions of the ions. Remove any ion positions that backscattered out of the surface of the nanowire.
3. Give the ions a randomized angle drawn from 4π.
4. Have that ion "decay" into its next isotope and give it a kinetic energy associated with the dominant decay mode. Evolve the ion position in SRIM.
5. Repeat steps 2-4 until one reaches ^{213}Po. The fraction of ions that remain in the detection volume form p_k^{sur}. Since we may define various maximum detection depths d, the survival fraction becomes a function $p_k^{sur}(d)$ (see Fig. 3.31).

One we had obtained the individual survival probabilities, we could solve for a series of Bateman equations. For the k-th isotope, the Bateman equations would

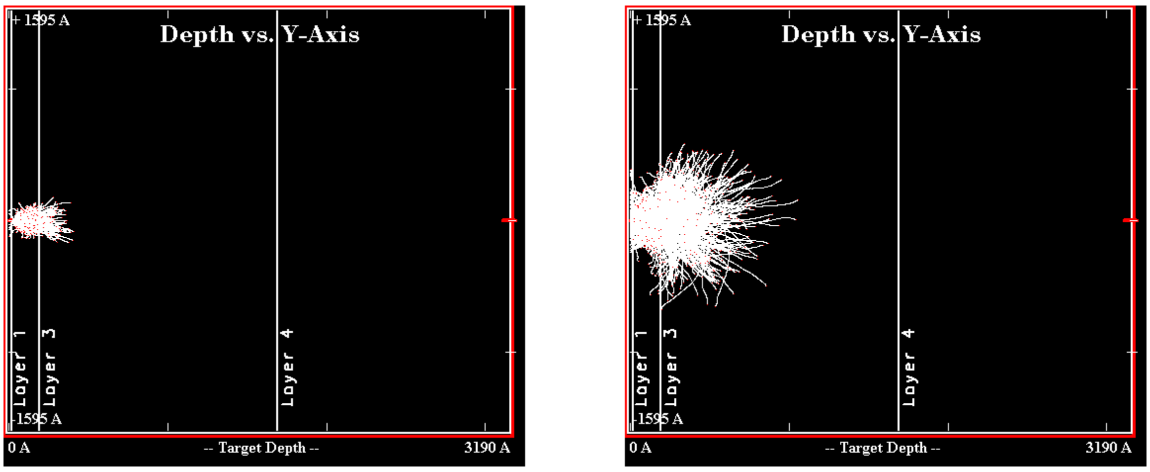

Figure 3.30: An example of a portion of the particle transport simulation. (Left) Initial impact of an isotope form the nanowire into the system creates a spread of positions which are fed into a RANGE_3D.dat file. (Right) RANGE_3D.dat file is converted to TRIM.dat file to create initial positions of next portion of simulation. Ions are then given random velocity direction and transported.

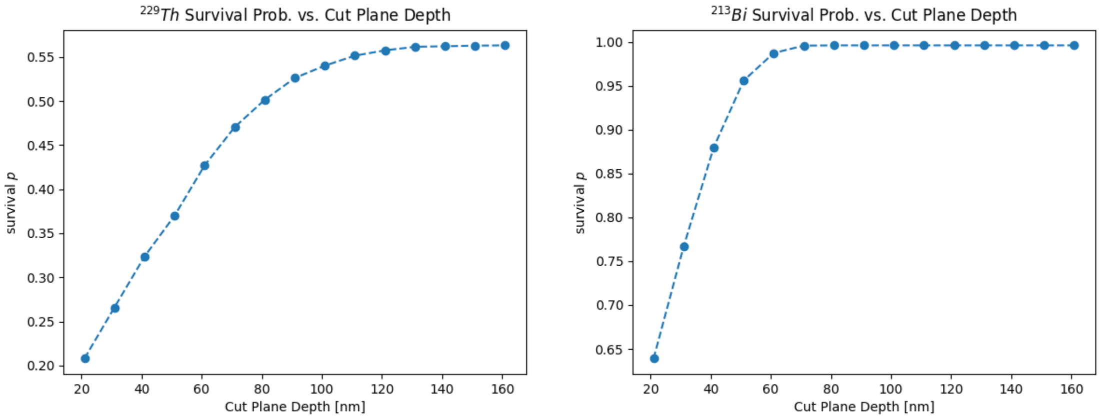

Figure 3.31: Comparison of the $p_k^{sur}(d)$ between ^{229}Th and ^{213}Bi as a function of cut plane depth.

start from the k-th isotope, ignoring those before it. This would of the form

$$\begin{aligned} \dot{N}_k^k &= -\lambda_k N_k^k + Q_k \\ \dot{N}_{k+1}^k &= \lambda_k N_k^k - \lambda_{k+1} N_{k+1}^k \\ &\vdots \\ \dot{N}_{Po}^k &= \lambda_{Bi} N_{Bi}^k - \lambda_{Po} N_{Po}^k . \end{aligned} \tag{3.16}$$

For the k-th isotope, one would expect that its partial contribution to the ^{213}Po signal rate would be

$$A_k(d) = p_k^{sur}(d) \lambda_{Po} N_{Po}^k \tag{3.17}$$

and that the total ^{213}Po signal rate as a function of the detection depth would be given by

$$A_{Po}^{total}(d) = \sum_k A_k(d), \tag{3.18}$$

where the sum runs over all k isotopes. A plot of A_{Po}^{total} vs. d can be seen in Fig. 3.32. The two free parameters were that the age of the ^{233}U target and the maximum detection depth. We define the age of the source by saying that at $t = 0$, the source is pure ^{233}U with no daughters present. The rates appeared to converge by allowing for a detection depth of 60 nm, and the source data sheet for the ^{233}U target indicated 42 Bq of possible contaminants, which bounds the age of the source from 12-20 years. This puts the expected "double-clicks" from ^{213}Po to be $\approx 1 \times 10^{-4}$ Hz, which is in good agreement with the observed maximum rate in the ^{233}U datasets.

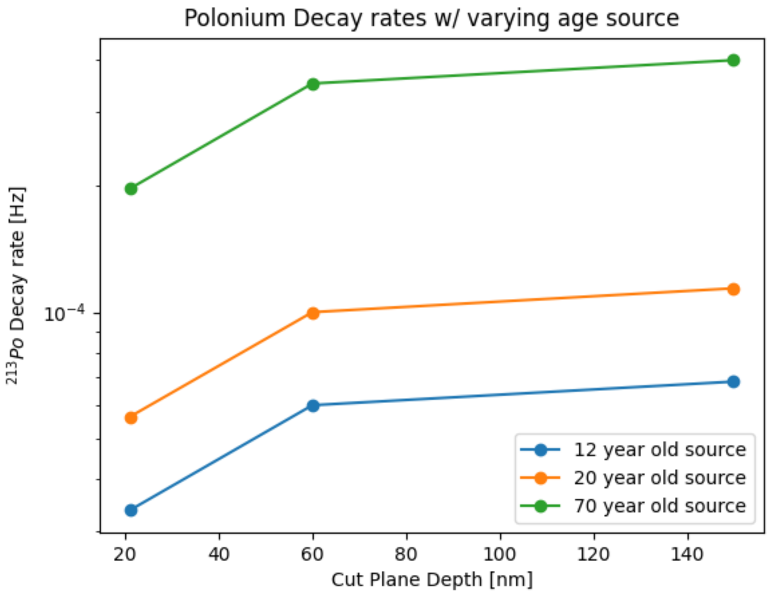

Figure 3.32: The estimated rate of ^{213}Po decays that would occur within the SNSPI for different values of the maximum detection depth and age of the ^{233}U source.

3.7.2 Loading ^{229}Th Target

To verify the SNSPI could detect ^{213}Po, we loaded a 5.1 kBq ^{229}Th target into the cryostat. This target emits only ^{225}Ra onwards in the ^{233}U decay chain, so it has no isomeric signal. Therefore, if a signal was observed within the IC time window, it would have to be due to the ^{213}Po embedded in the nanowire.

Within a week double clicks were observed in the nanowire with the IAT plot showing a ∼5 μs decay lifetime. Moreover, the rate of these double clicks was ∼ 2×10^{-3} Hz, which was predicted by a similar version of the analysis described in Sec. 3.7.1.2. A scaled version of the CER curves obtained from the ^{233}U and ^{229}Th targets can be seen in Fig. 3.33. As one can see the ^{233}U and ^{229}Th data sets have very similar excess CER at low bias currents. This test led us to conclude that the majority of the significant signal which had been observed had been due to ^{213}Po.

3.8 Conclusion

After it was concluded that ^{213}Po was making up some of the signal at low bias currents, we continued to look for any possible signature of the IC signal buried in

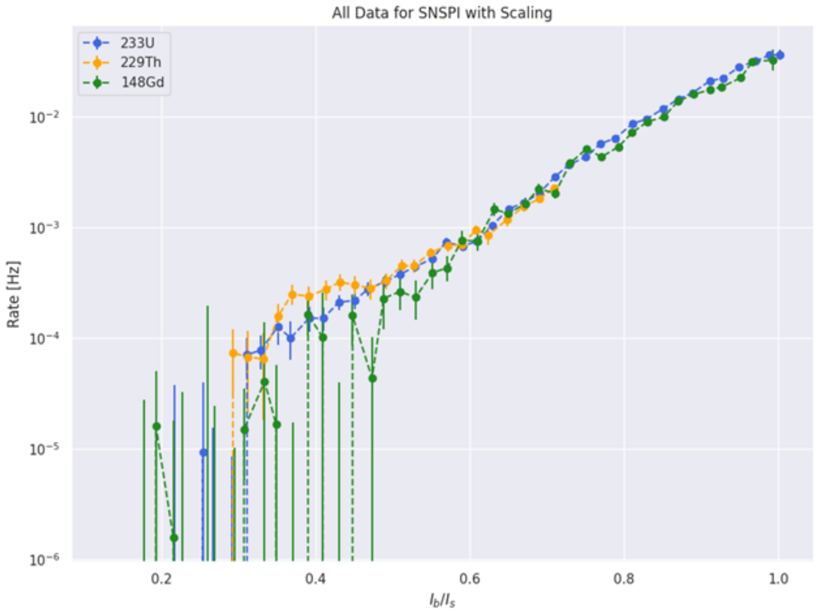

Figure 3.33: CER data for the ^{233}U (blue), ^{148}Gd (green), and ^{229}Th (yellow) targets in the WSi SNSPI experiment. All data has been scaled to line up the high bias current tails. Note the similar excess CER at lower currents present in the ^{233}U and ^{229}Th data, but absent in the ^{148}Gd curve. This excess was attributed to ^{213}Po contamination.

the data. This amounted to trying a triple err model of the form

$$f_3(i) = \lambda_{IC}\,\mathrm{erf}\left((i - i_c^{IC})/\sigma_{IC}\right) + \lambda_{Po}\,\mathrm{erf}\left((i - i_c^{Po})/\sigma_{Po}\right) + \lambda_b\,\mathrm{erf}\left((i - i_c^{b})/\sigma_b\right), \quad (3.19)$$

but the log-likelihood ratio between this and the double erf model yielded p-values ∼0.5. We concluded that the experiment had not detected any significant signal of the IC decay.

Several ideas were put forth as to why the signal was not observed in our experiment, but seemingly was observed at LMU:

- The stopping range calculations were wrong and the ^{229}Th penetrating through the SNSPI, and not releasing its energy into the nanowire
- The ^{229}Th was stopping in the nanowire, but at a location deep enough that the heat generated by IC was coupling out of the nanowire through the substrates.
- The IC rate may be modified by the local electronic environment when it

implants deeper into the nanowire. This may cause it to quench more rapidly or more slowly than expected.

It was proposed that a next stage of the experiment would combine a stopping buffer gas of static He in order to slow down the ^{229}Th ions so that they would be in a soft implantation regime. In addition, plans were under way to use a double layer SNSPI that would be able to distinguish between an IC event that would only be detected in its top layer, versus a ^{213}Po event that would generate enough energy to be detected in both layers. While improvements were underway, the uncertainty in the isomeric energy was reduced from 8.28(17) eV to 8.338(24) eV [KMA23]. This was thought to be beyond the energy resolution capabilities of the SNSPI and monochromator calibrations of our experiment, and thus further work on this project was halted to focus efforts on the laser based efforts.

3.9 Appendix

3.9.1 Waveform Classification

There were many noise events that would trigger the digitizer and these needed to be filtered from our analysis. To do this efficiently we opted to use a support vector machine (SVM) classifier (implemented in Scikit-learn v1.5.1). Several hundred waveforms of different amplitudes were hand separated into four categories: *singles* (single normal waveform), *doubles* (two nanowire clicks in one 240 ns acquisition), *noise* (clear noise in the detection system), and *weirds* (where the classification was unclear). After sorting the waveforms, the classifier would be trained. An example of its ability to separate waveforms can be seen in Fig. 3.34. It was decided to reject "weird" waveforms as noise events. The classifier was benchmarked to be able to correctly classify a waveform with $> 99\%$ success.

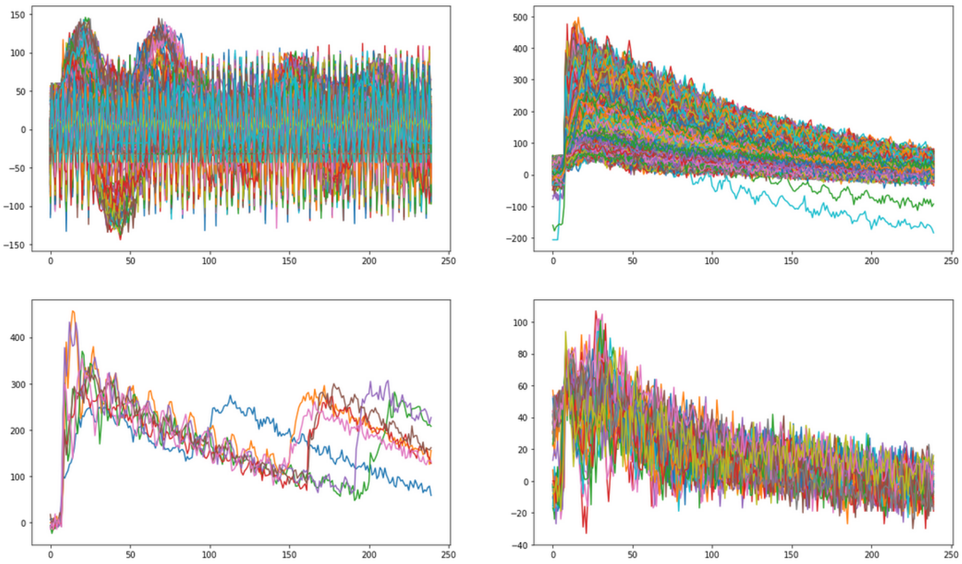

Figure 3.34: SVM classifier separating the waveforms. (Top left) Noise waveforms. (Top Right) Single clicks. (Bottom Left) Two events in a single acquisition window. (Bottom Right) "Weird" waveforms that are difficult to classify.

3.9.2 Derivation of the Uncorrelated Background Distribution

At a particular value of the normalized bias current, the histogram of impact positions for events with an inter-arrival time ≥ 100 us was treated as a distribution of uncorrelated events. From this histogram a probability distribution for which SNSPI pixel an event would occur could be obtained, $P(X)$. The following procedure was used to calculate the position difference distribution for the uncorrelated events:

Consider two random variates δt_1 and δt_2 drawn from $P(\delta t_1)$ and $P(\delta t_2)$, respectively. The probability of a particular value of their difference, $U\,(\mathit{\Delta} = \delta t_1 - \delta t_2)$, is given by the following sum over the joint distribution

$$U(\mathit{\Delta}) = \sum_{\delta t_1 - \delta t_2 = \mathit{\Delta}} P(\delta t_1) P(\delta t_2) \tag{3.20}$$

3.9.3 Derivation of the Single-Site Distribution Function

When Thorium hits some portion of the nanowire, two pulses are detected from either side of the circuit. The true time difference between the pulses is given by some $t_1 \in [-L, L]$. When the isomeric Thorium then decays, it creates two more pulses that should also have a true timing difference $\delta t_1 = \delta t_2 = \delta t$. Our CAEN has a nanosecond timing jitter, which we can model as an additional uniform random variable $\varphi_1, \varphi_2 \in [0, 1)$.

Thus, the timing difference we measure is given by

$$\mathit{\Delta} = \lfloor \delta t + \varphi_1 \rfloor - \lfloor \delta t + \varphi_2 \rfloor \tag{3.21}$$

where we have used a floor function since our device can only measure the arrival times of the pulses with respect to an internal digital clock. We note that for any $x, y \in \mathbb{R}$ and $k \in \mathbb{Z}$

$$\lfloor x + k \rfloor - \lfloor y + k \rfloor = \lfloor x \rfloor - \lfloor y \rfloor \tag{3.22}$$

which means that without loss of generality, we can consider $\delta t \in [0, 1)$.

The possible values $\mathit{\Delta}$ can take on are $-1, 0, 1$. If we look at a particular value of δt we get the regions depicted in Fig. 3.35. We can then note that the two $\mathit{\Delta} = 0$ regions trace out two oblique pyramids in the t-φ_1-φ_2 cube. The volume of each of these pyramids is $1/3$, and so $C(\mathit{\Delta} = 0) = 2/3$. By symmetry, $C(\mathit{\Delta} = 1) = C(\mathit{\Delta} = -1) = 1/6$.

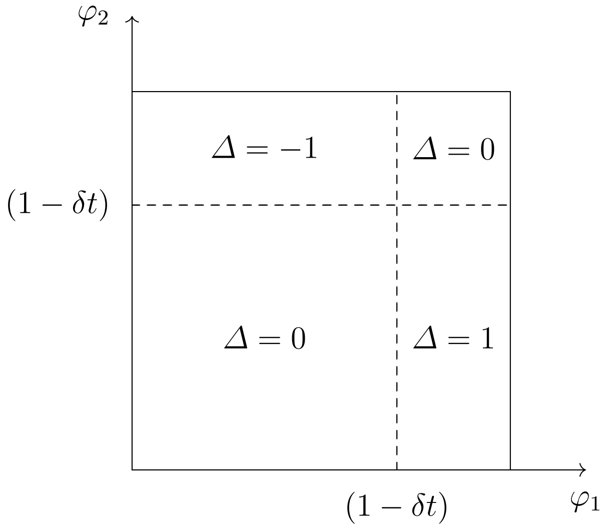

Figure 3.35: The possible values of $\varDelta$ for a particular value of δt under different jitter values (φ_1, φ_2).

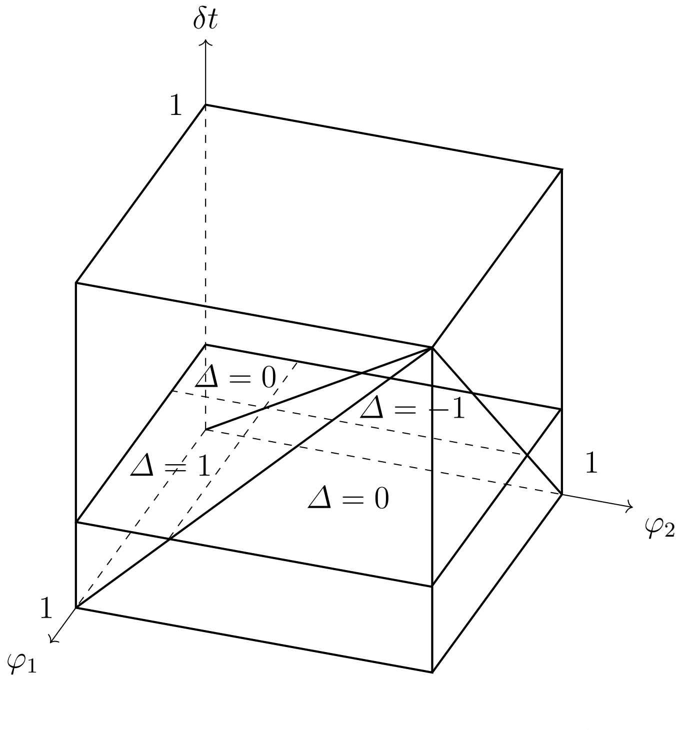

Figure 3.36: A δt slice in the t-φ_1-φ_2 space. Note that as δt is varied, both $\varDelta = 0$ regions trace out an oblique pyramid (one is shown by the diagonal lines).

Equivalently, one can obtain the same answer by carrying out the following integrals:

$$\begin{aligned}
C(\mathit{\Delta}=0) &= \int_0^1 d(\delta t) \int_0^{1-(\delta t)} d\varphi_1 \int_0^{1-(\delta t)} \varphi_2 + \int_0^1 d(\delta t) \int_{1-(\delta t)}^1 d\varphi_1 \int_{1-(\delta t)}^1 \varphi_2 \\
C(\mathit{\Delta}=1) &= \int_0^1 d(\delta t) \int_{1-(\delta t)}^1 d\varphi_1 \int_0^{1-(\delta t)} \varphi_2 \\
C(\mathit{\Delta}=-1) &= \int_0^1 d(\delta t) \int_0^{1-(\delta t)} d\varphi_1 \int_{1-(\delta t)}^1 \varphi_2,
\end{aligned} \tag{3.23}$$

obtaining the same results.

3.9.4 Time Tag Correction

3.9.4.1 Noisy Waveform Modeling

The CAEN DT5751 time tags are determined by when a waveform crosses a set voltage threshold (subtracting the average baseline). In our experiment this was set to 40 mV; this threshold was the lowest trigger that the detection system could handle without being flooded by noise events. An analysis was carried out by Christian that found that the SNSPI waveforms had similar sharp rising edges to about 40% of their total amplitudes, but then had a variety of different patterns. Examples of the different waveform types can be seen in Fig. 3.37.

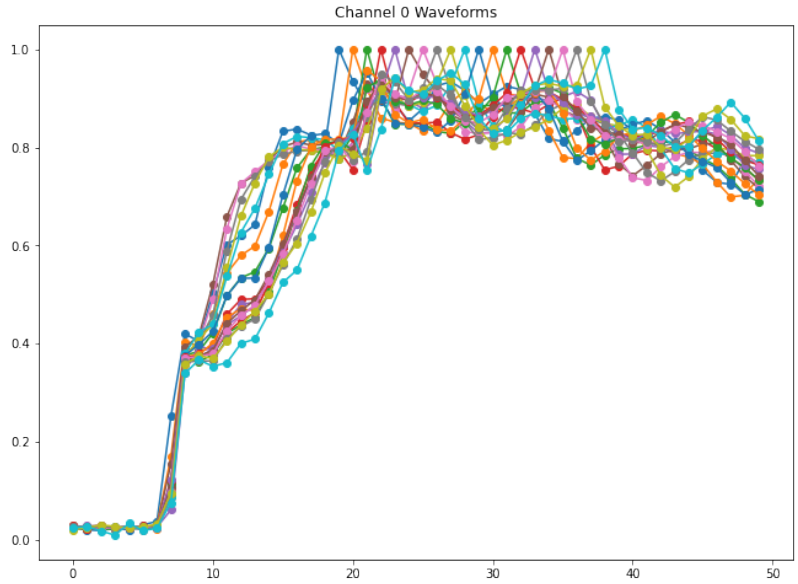

Figure 3.37: Examples of different waveform shapes without additional electronic noise (normalized to unit amplitudes).

These waveforms would then also be distorted by electronic noise; a separate analysis found that the noise of the device could be approximated as 9 mV RMS white noise. This distortion can be seen in Fig. 3.9.4.1. The black horizontal line in the aforementioned figure serves to demonstrate that if one naively chooses a triggering threshold without regard to the waveform amplitudes, the uncertainty in the timing of a waveform easily be on the order of 10 ns.

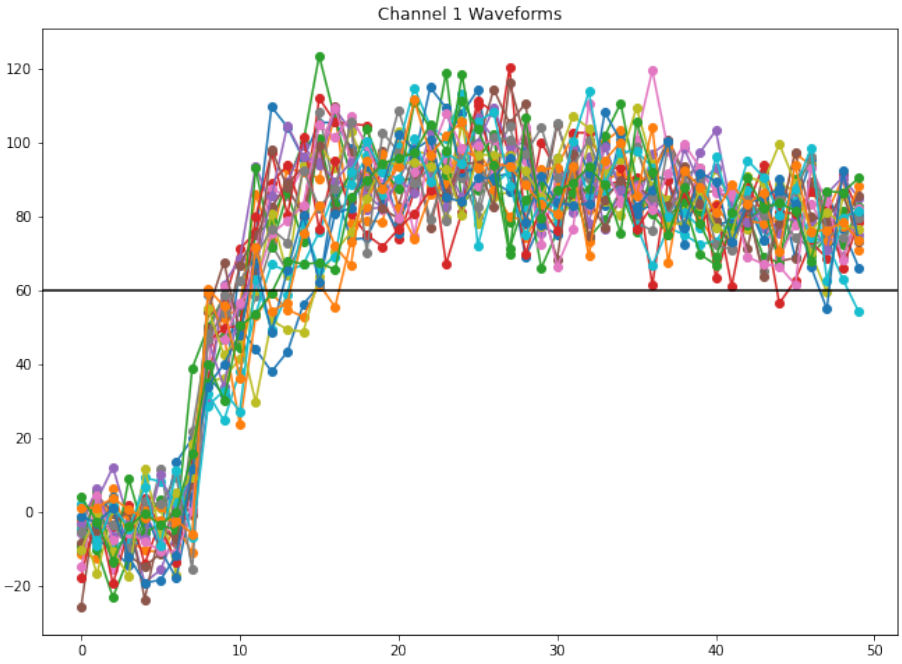

Figure 3.38: Examples of different waveform shapes with additional electronic noise (average amplitude is 105 mV). The black line serves as an example of a poor choice of trigger threshold.

3.9.4.2 Threshold vs. Amplitude

This systematic uncertainty distorts the measured time difference between positive and negative waveforms from the ends of the SNSPI (as in the Single-Site Click Model discussion we refer to these differences as t_i). For the IC events, these distortions propagate to the position difference of subsequent events $\mathit{\Delta} = \delta t_1 - \delta t_2$, which causes the distribution of single site events to differ from the ideal case derived in our discussion of the Single-Site Click model (see Fig. 3.39b)

In order to minimize the effect of this systematic on the rest of the data analysis, the following time tag correction procedure was used:

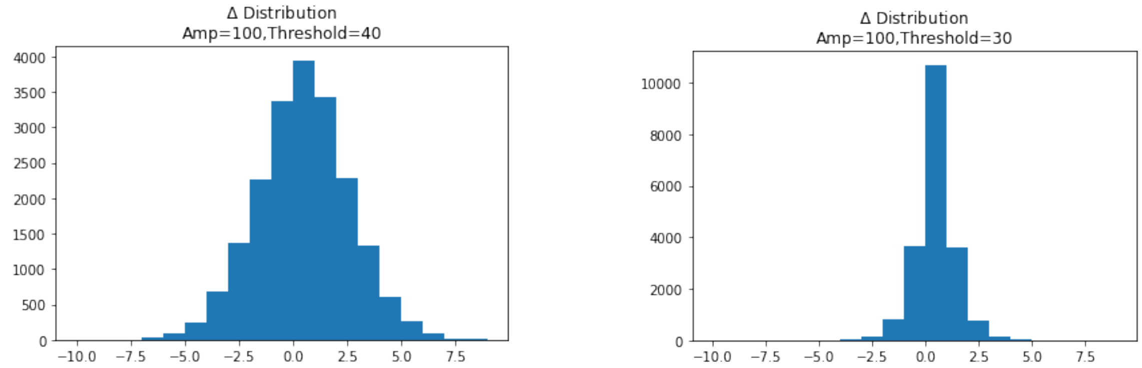

(a) A single site distribution that has been smeared by a poor choice of trigger threshold.

(b) A single site distribution that closer approaches the ideal distribution limited only by the CAEN's 1 ns jitter.

Figure 3.39: $\mathit{\Delta}$ distributions under different threshold choices.

1. For a given amplitude, multiply a pair of idealized waveforms by that amplitude and add electronic noise (one is the "positive" pulse and the other is the "negative pulse").

2. Choose a trigger threshold (e.g. 60 mV) and define the time difference of the pulses as when both pulses cross this threshold (call this t_1). To this timing difference add an additional 1 ns timing jitter φ_1 to simulate the CAEN's jitter.

3. Repeat this procedure to obtain another time tag t_2.

4. Repeat this process many times with a given amplitude and threshold, and one should obtain a distribution of position difference $\mathit{\Delta}$. The ideal case is given by the distribution

$$C(\mathit{\Delta}) = \begin{cases} 1/6 & \text{if } \mathit{\Delta} = -1 \\ 2/3 & \text{if } \mathit{\Delta} = 0 \\ 1/6 & \text{if } \mathit{\Delta} = 1 \end{cases} \tag{3.24}$$

   One must compute the difference/distance between the ideal case and the obtained distribution of position differences $\mathit{\Delta}$. In this work, the distance was defined using the Earth Mover's Distance [RTG98].

5. Repeat this process at a fixed amplitude for several values of the threshold to obtain the optimal threshold (i.e. the threshold that minimizes the distance between the Monte Carlo distribution and the ideal distribution).

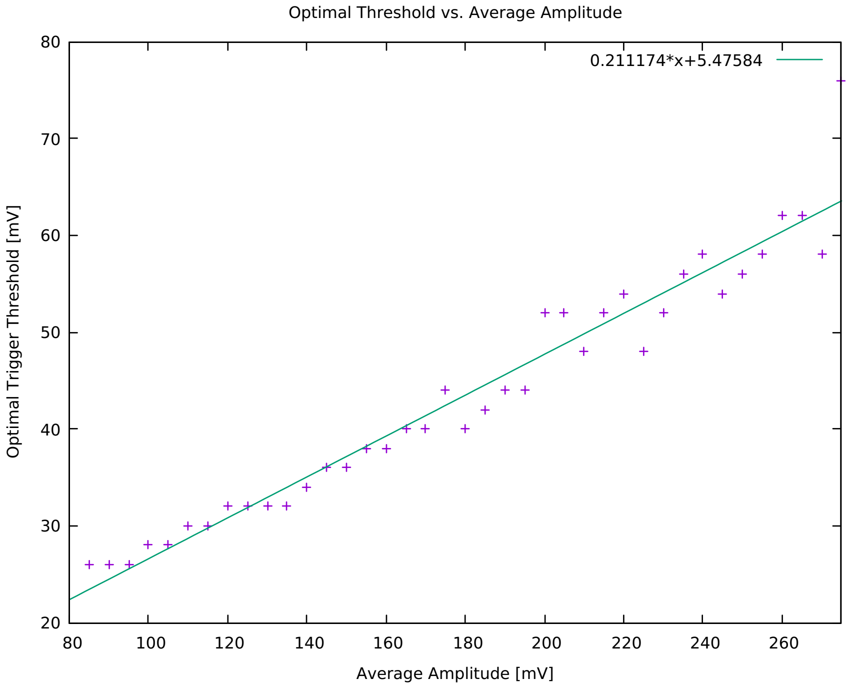

Figure 3.40: The optimal trigger threshold that most closely approximates the ideal single-site distribution, as a function of average waveform amplitude (as determined by Monte Carlos with Christian's waveform templates and 9 mV RMS white noise). The line that best fits the above dataset is given by $0.211174x + 5.47584$.

The results of these Monte Carlo can be found seen in Fig. 3.40

3.9.4.3 Amplitude vs. Bias

At this point, we have determined that there is an optimal threshold for a particular waveform amplitudes, such that one obtains the optimal single-site position difference distribution. In the real data however, each waveform is subject to noise and fluctuations, and as such the measured amplitude differs from the true "noise-free" amplitude. We can use the average amplitude as a function of normalized bias current as an estimate of the "noise-free" amplitude and use this quantity as the basis for our threshold correction scheme.

Plots of the amplitude versus normalized bias current are seen in Fig. 3.41. These were fit with a polynomial of the form $A = p_1 i + p_0$, where A is the average amplitude and i is the normalized bias current.

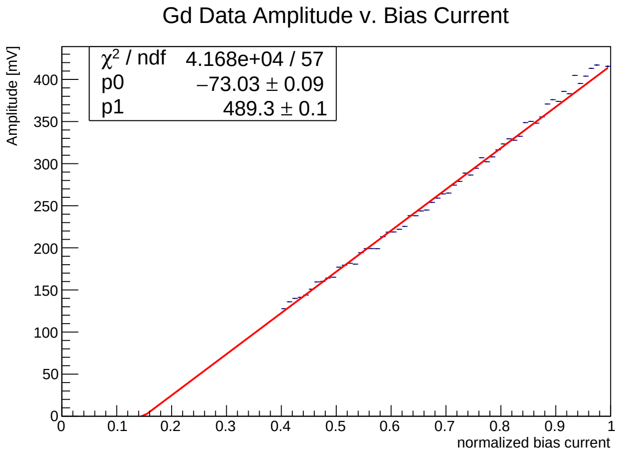

Gd Data Amplitude v. Bias Current

(a) Gd148 data (control).

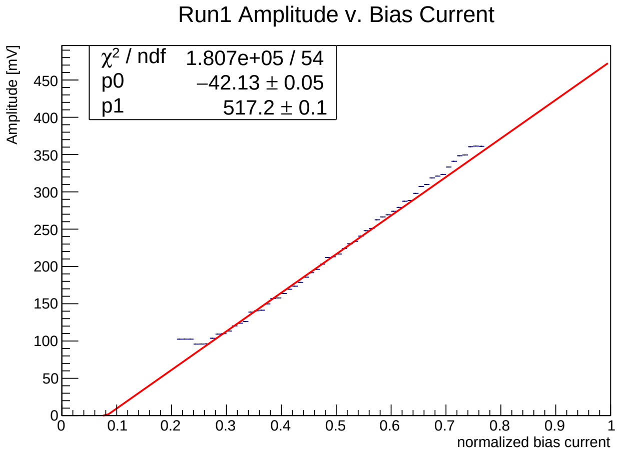

Run1 Amplitude v. Bias Current

(b) Run 1 U233 data (no source mount voltage biasing).

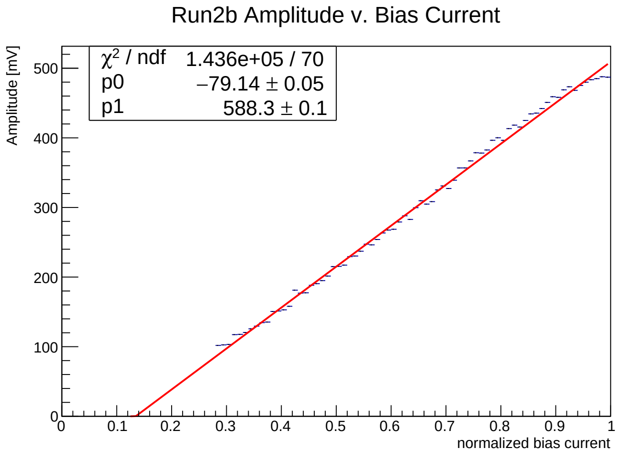

Run2b Amplitude v. Bias Current

(c) Run2b data (May/June).

Figure 3.41: Amplitude vs. Normalized Bias current for relevant datasets. Lines of best fit obtained by linear regression.

Dataset	p_0	p_1
Gd148	-73.03	489.3
Run 1	-42.13	517.2
Run 2	-79.14	588.3

Table 3.3: Linear regression coefficients of average amplitude vs. normalized bias current.

3.9.4.4 Threshold vs. Bias

With Table 3.3 we have a way to assign an amplitude to a bias current. With our analysis of the trigger thresholds, we have a method to assign an optimal trigger threshold to each amplitude. Thus, we can use both these linear models to assign an optimal threshold to each normalized current in the form of a linear polynomial $t = q_1 i + q_0$ where t is the optimal threshold and i is once again the normalized bias current. The coefficients of the linear models for each dataset are in Table 3.4.

Dataset	q_0	q_1
Gd148	-9.946	103.3
Run 1	-3.421	109.2
Run 2	-11.236	124.2

Table 3.4: Linear regression coefficients of optimal threshold vs. normalized bias current.

From this calibration between threshold and normalized bias current, we could go back through the data and correct assigned time tags of the lower bias current data.

3.9.5 MLE Point Estimates via Simplex Ensemble

Given the high dimensionality and non-linear nature of the models that needed to be fit, often a single instance of a minimization routine would fail to find the global minimum of the negative log-likelihood. At the same time, a grid search of a 3-6 dimensional space with enough sampling is costly as well. The solution was to combine the approaches, a grid would define an ensemble of initial conditions from which Nelder-Mead simplices [NM65] would begin their optimization routine. Some member of the simplex ensemble would converge to the global minimum, while at the same time the rest of the ensemble would tend to sample the likelihood function around that minimum. This large sample of values of the likelihood function around its global minimum could be used later for profiling methods and computation of confidence intervals.

CHAPTER 4

Towards Laser-based Conversion Electron Mössbauer Spectroscopy

4.1 Concept

When earlier efforts had failed to definitively observe the radiative decay of the isomer in a crystalline environment, new avenues were sought to observe the excitation of ^{229m}Thusing a pulsed VUV laser system. In 2017, the idea of laser-based conversion-electron Mössbauer spectroscopy (CEMS) in ^{229}Th had been proposed in [WSS17] as a way to resolve the transition, and the details of its implementation in our laser system were fleshed out in [WSS19]. Unlike the VUV-transparent crystal approach, which seeks to keep ^{229}Th in a high band gap material (>9eV) to freeze out any potential quenching mechanisms, the CEMS approach instead relies on the quenching of ^{229m}Thby the internal coversion (IC) channel in a low (<8eV) band gap environment. This low band gap environment is provided by ThO_2 electroplated onto a stainless steel target.

Within the low band gap environment, an excited ^{229}Th nucleus will de-excite via the internal conversion channel, and promote an electron across the band gap of ∼5.8 eV [RL71]. If this promoted IC electron was from a shallow enough state in the valence band, the electron will have enough energy to overcome the work function barrier of ThO_2. These IC electrons may then emerge from the ThO_2 surface and be guided to a detection system to serve as the spectroscopic signal for excitation of the transition. However, there is the potential that the initial burst of photoelectrons from the VUV pulse on the ThO_2 target will overwhelm the detectors, and so a system of fast voltage pulsing needs to be used to suppress the initial photoelectrons for the first ∼100 ns after the laser pulse, and then the polarity must be flipped to

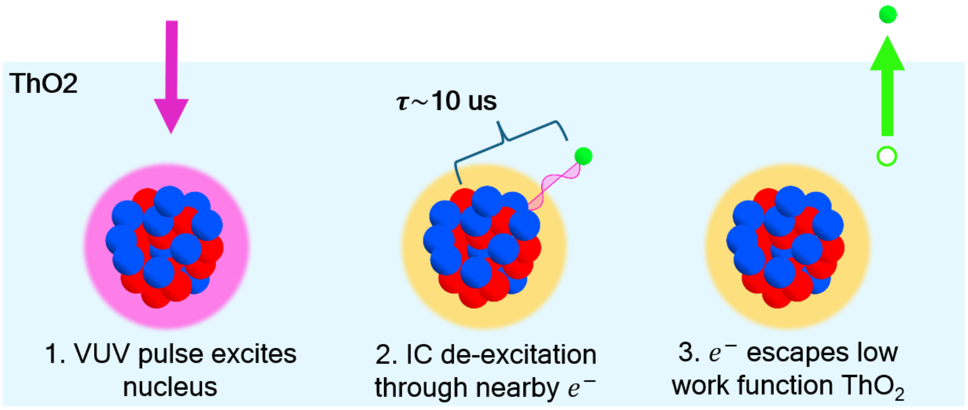

Figure 4.1: Schematic concept of the CEMS experiment.

help in the extraction of IC electrons.

Assuming that the lifetime of the isomer has been reduced to 10 μs within the solid-state environment [SWT17], the search can be performed much more rapidly, where the spectroscopic signal can be searched for between each pulse of the VUV pulsed-laser system.

The number of excited ^{229}Th atoms is given by

$$N_{exc}(t_{exc},\omega) = \frac{g_{exc}}{g_{gnd}} \frac{\lambda_0^2}{2\pi} \frac{\varphi_p n_{Th} l}{(1+\alpha_{IC})\Gamma_L} \frac{1}{1+4(\frac{\omega-\omega_0}{\Gamma_L})^2} \left(1-e^{\Gamma_{IC} t_{exc}}\right), \tag{4.1}$$

while the number of detected IC electrons is given by

$$N_{det} = \eta_{esc}\eta_{coll}N_{exc}, \tag{4.2}$$

where η_{esc} is the efficiency with which IC electrons can emerge from the ThO_2 surface, and where η_{coll} is the efficiency with which those emerging IC electrons are collected into our detection system. While η_{coll} can be directly measured by determining the collection efficiency of photoelectrons, there is no direct way to determine η_{esc}. For this experiment, we used the efficiency with which photoelectrons were extracted from the ThO_2 surface as an estimate of this value, but it is unlikely there is complete parity between the processes that underlie the excitation of the electrons across the band gap by photons versus internal conversion.

4.2 Experimental Apparatus: Construction, Design, and Data Acquisition

4.2.1 Vacuum Chamber Components

The experimental chamber (hereafter referred to as the IC chamber) was based around a Kimball octagon attached to the end of the VUV crystal beamline. The two systems were joined by removing the pyroelectric detector, and placing a MgF_2 window in between the two sections. An Ar purge was introduced right at the MgF_2 window on the crystal chamber side. A sketch of the chamber including the major components is shown in Fig. 4.2. The major components are listed as follows:

- Two multi-channel plate detectors (MCPs) for electron detection. The MCPs were Photonis 30305 matched plates.
- SAES NexTorr D200-5 combined NEG/Ion pump for pumping of the chamber.
- Fused silica viewport to monitor to allow for optical access to the chamber.
- VAT gate valve for initial pump during bakeout
- Plasma cross made up of the following:
  - Four-way 2.75" CF cross
  - KJLC EFT0122053Z feedthrough for the plasma RF
  - Window to view the plasma ignition
  - Duniway VLVE-1000 leak valve to leak in the oxygen.
- An InstruTech BA601 nude hot ion gauge for UHV pressure measurement independent of the gauge integrated into the SAES pump.
- The ThO_2 target and target mount, which was attached to a custom KJLC 8-pin high voltage electrical feedthrough. More details about the target mount will be discussed in Section 4.2.3.

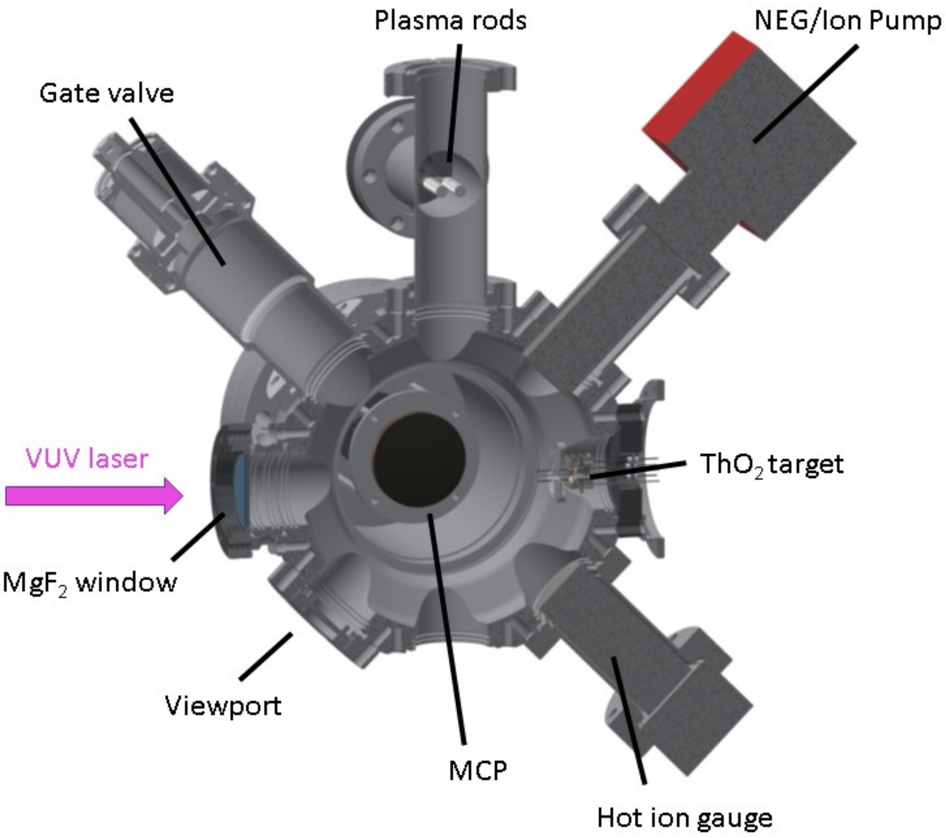

Figure 4.2: Design drawing of the IC chamber with an early version of the ThO_2 target. Note that certain components (such as the hot ion gauge) were moved in the final version of the chamber.

If one has to decide between trying to deflect the VUV beam vertically in our experiment, versus moving the entire chamber with some kind of mechanical contraption, it turns out that the latter is far simpler. We went with a stepper motor based approach to move the entire chamber up and down on our optical table. The chamber sat on a large aluminum 80-20 bar, which in turn was attached to a L-arm swing structure and supported by a stepper motor assembly (see Fig. 4.3). The stepper motor assembly consists of a stepper motor linear actuator (Ametek E57H4A-3.25-900) inside a aluminum housing and an aluminum plate with dowel pins. The aluminum housing allows the stepper motor to rotate with respect to the optical table normal. Freedom of movement in the aluminum housing is necessary as the stepper motor needs to be able to rotate as it move the bar at the bottom of the chamber up and

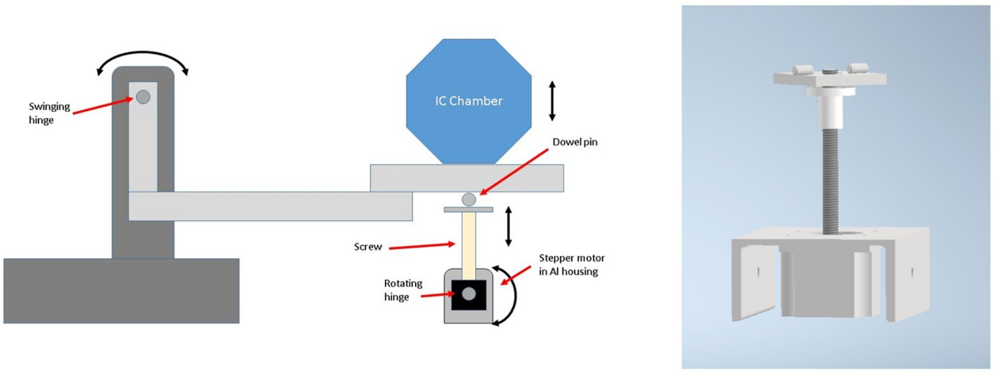

Figure 4.3: Sketches of the components that allowed for movement of the IC chamber up and down above the optical table. On the left is a diagram of the hinge system and stepper motor that allowed for movement of the chamber. The range of movement is indicated by the black arrows. On the right is a CAD mockup of the stepper motor and dowels that were below the chamber that moved it up and down.

down. At the same time there needs to be some freedom of movement at the point of contact between the linear actuator and the bottom bar holding the chamber. This is achieved by attaching an aluminum plate to the end of the actuator, and establishing the contact with smooth dowel pins (see right image of Fig. 4.3).

4.2.2 Cleaning and Bakeout Procedure

All metallic and ceramic components were separated by material and cleaned using the following procedure:

- An initial wipe down of the components was done with isopropanol to remove gross contamination.
- Parts were sonicated in a hot bath (∼30°C) of water and detergent. The sonicator was set for an initial 5 min. of degassing, followed by ∼30 min in the normal sonication mode. Degassing was critical for removing air bubble from various small holes. Citranox was used for copper and beryllium-copper, while Alconox was used for all other metals.

- Parts were rinsed in Type 1 de-ionized water heated to ~30°C. Clean spray bottles of de-ionized water were used to spray into blind taps until soap was seen to come out.
- Parts were then sonicated in Type 1 de-ionized water. The sonicator was set for an initial 5 min. of degassing, followed by ~30 min in the normal sonication mode.
- Parts were removed from water bath and sprayed dry with filtered nitrogen using clean PTFE spray gun.

Assembly of the chamber was conducted in a positive pressure laminar flow hood with integrated HEPA filter. For the bakeout portion, a Pfeiffer HiCube turbopump stand was attached via a bellow to the IC chamber by the gate valve. On the turbopump side an ion gauge and Standford Research Systems residual gas analyzer (RGA) were attached to monitor the bakeout.

Prior to the bakeout, the chamber was first cleaned using an oxygen plasma flow. The oxygen was sourced from a cylinder of 5N oxygen that was introduced through the leak valve port of the chamber. Plasma ignition was achieved by a combination of adjusting oxygen flow rate via the speed of the turbopump and the conductance of the Duniway leak valve, impedance matching the RF frequency into the RF rods, and adjusting the input RF power. Once plasma ignition was achieved in the RF cross, the plasma was left on for 2 hours to remove large mass ($\gtrsim$55 amu) contaminants in the chamber. The removal of these contaminants was verified using the RGA (see Fig. 4.4).

Bakeout of the chamber was conducted in a home-built oven that was able to achieve ~170° C. Before the initial heating began, the SAES NEG pump was set to conditioning mode. The temperature was then slowly raised (~ 5° C per 15 min) until the system reached ~165° C. The system was pumped by the turbopump for ~1 week at this temperature, meanwhile the NEG was kept in conditioning mode. A power outage then occurred during the bakeout, during which the pumps shut off while the oven was on, causing a coating of the chamber with high mass contaminants. It was then necessary to cool the chamber down to room temperature and perform another

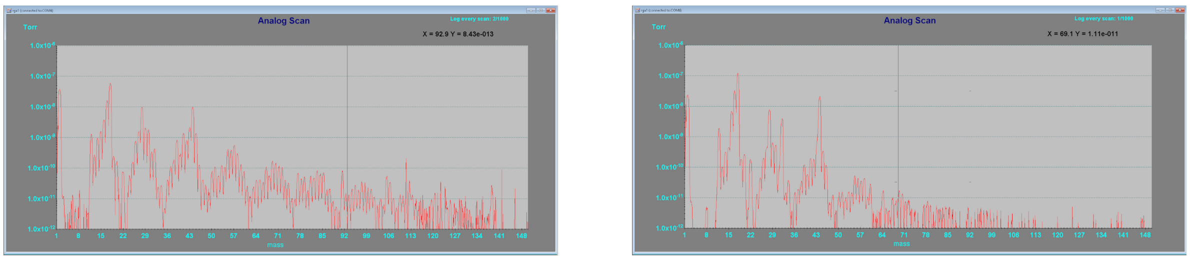

(a) Before plasma cleaing at T=20° C, ∼1-e7 mBar.

(b) After ∼2 hours of oxygen plasma cleaning at T=20° C, ∼2-e7 mBar.

Figure 4.4: RGA scans taken before bakeout. Note the disappearance of high mass ($\gtrsim$55 amu) peaks after plasma cleaning. Pressure readings taken from ion gauge.

Figure 4.5: (Left) IC chamber sitting in home built oven. The chamber was hooked up to oxygen bottle and RF tank circuit for plasma cleaning. (Right) Visible oxygen plasma ignition in the RF cross.

~2 hour plasma cleaning in order to remove these contaminants. After this second plasma cleaning, the system was raised again to ~165° and the bakeout continued for another ~1 week until the pressure stopped dropping on the turbopump ion gauge. We then lowered the temperature slowly, and in the last few hours before reaching room temperature we set the NEG pump to its timed activation mode, in order to have it release any contaminants it had absorbed during the bake. Once room temperature was reached, the SAES NEG was set to normal operation with the ion pump activated, and the chamber was sealed off.

The final pressure was inferred using both the ion pump gauge of the SAES system, as well as the InstruTech hot ion gauge. The SAES ion pump reported a pressure of <1e-10 mBar, however the hot cathode gauge reported a pressure of ~2.5e-10 mBar. It is likely that the hot cathode gauge is a more reliable measure of the overall system pressure, given that the SAES ion pump reading is based on the pressure behind the getter material, which will naturally be lower than other parts of the chamber. As will be discussed later, this pressure was sufficiently low as to have minimal impact on the collection efficiency of electrons from the ThO_2 target in our system. In addition, this pressure seemed to have been sufficient to mitigate any hydrocarbon buildup on our target, as no visual degradation of the target was observed over the duration of the scan.

4.2.3 ThO_2 Target

The ThO_2 targets were produced by Eckert & Ziegler (E&Z). The production of the targets was done by electroplating of ThO_2 onto an ~25 mm disk of stainless steel. The electroplated region was ~5 mm in diameter. Four targets were produced with activities 5.1, 6.4, 16.3, and 29 kBq. Selection of the target was informed by two criteria:

- The attenuation length of VUV light in the region of 150 nm in materials is of order tens of of nanometers [CHH64]. Thus, increasing thickness at a certain point will only increase the background due to radioactivity without increasing the signal.

- The mean free path of electrons a few eV above the Fermi level (as would be the case for our promoted IC electrons) is of of order tens of nanometers [SD79], again meaning that too thick a target is detrimental to one's signal-to-noise ratio.

Assuming the target was completely made of ThO_2, the 6.4 kBq target corresponded to ∼25 nm thickness, making it the highest signal candidate within the ideal thickness range of a few ten nanometers. This thickness estimate was corroborated by an alpha spectrum taken by E&Z.

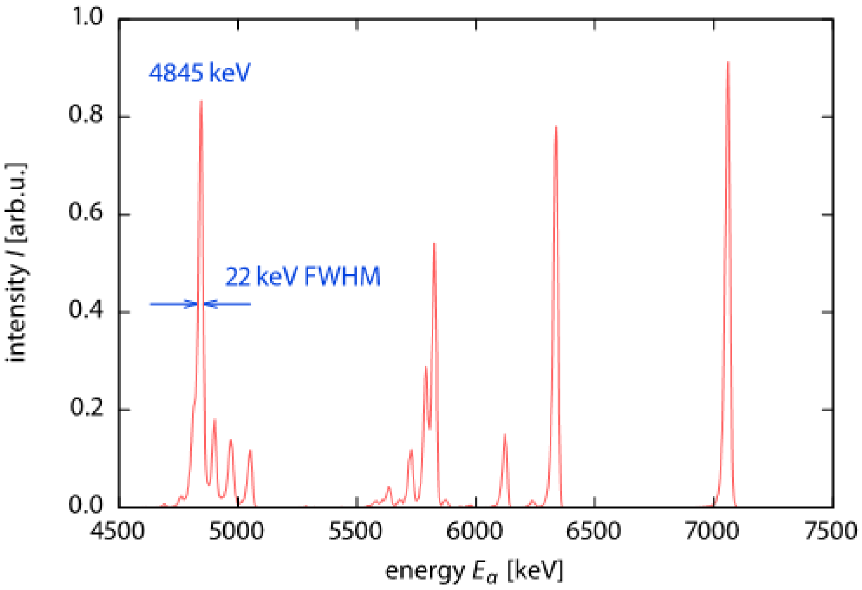

Figure 4.6: Alpha spectrum taken by E&Z of the ThO_2 target. The 22 keV FWHM of the peaks is consistent with 10-20 nm thickness of the target.

4.2.4 Target Mount

The ThO_2 mount was designed to allow for individual electrical biasing of the ThO_2 target and front aperture, while also allowing for readout of the voltage difference across a bare pyroelectric crystal (Gentec-EO catalogue #201351) sandwiched between the front aperture and a contact. The electrical isolation of the metallic components was established using Macor parts. As can be seen in Fig. 4.7, a spring contact was used to contact the back of the bare pyroelectric crystal, while the other side of the crystal was in contact with the front aperture. The mount was designed to have the ThO_2 target at an angle to leave open the possibility of measuring the reflected VUV light during the target illumination sequence.

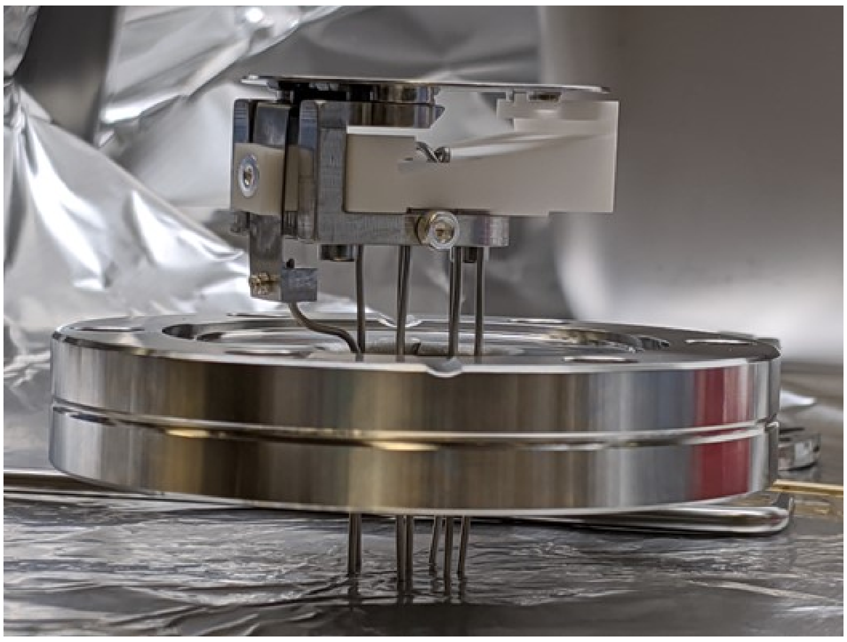

Figure 4.7: Photo of the ThO_2 mount without the radioactive target loaded. Macor pieces are used to keep the metal components electrically isolated. On the left a thin spring contact snakes its way around to contact the rear of a bare pyroelectric crystal. Also visible are the bent central feedthrough wires, meant to allow for voltage biasing of the ThO_2 target, as well as holding it in place mechanically.

A conical aperture was placed in front of the ThO_2 target (see cutway in Fig. 4.12b). It was found in SIMION simulations that this piece generated additional curvature in the equipotential surfaces near the target. This curvature, in turn, helped to slightly focus the electrons emerging from the target, and ensure that they would not crash into the front aperture piece (somewhat similar in spirit to an Einzel lens). With this conical piece added, our simulated efficiency for directing emerging electrons from the target past the front aperture was near unity.

4.2.5 Electronic Pulsing of the MCP and Target

From the outset, we had assumed that there would be a significant amount of photoelectrons and scattered VUV light in the first few hundred nanoseconds of the experiment. To deal with this, two things were needed. The first was to hold the ThO_2 target at a positive bias with respect respect to the front aperture so as to

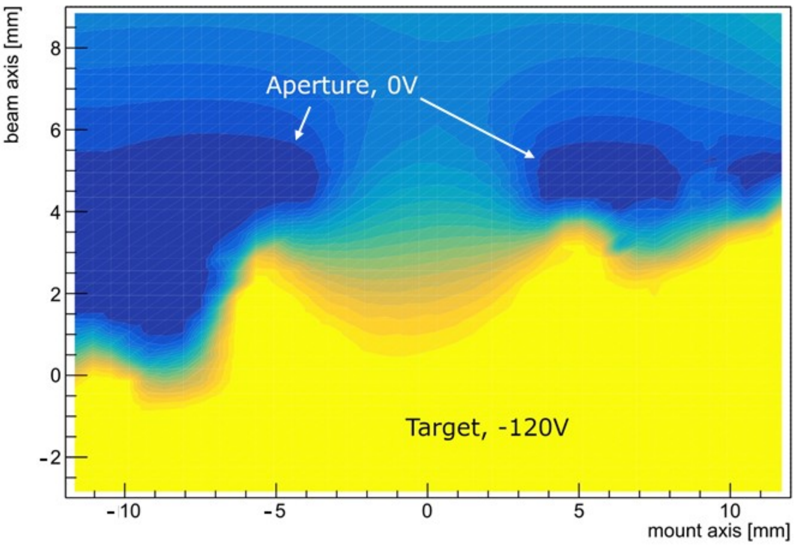

Figure 4.8: Potential contours for the ThO_2 mount region, as extracted from SIMION. These contours are from the solution to the Laplace equation with the target+conical aperture biased at -120V, while the front aperture and chamber walls are held at ground. The target+conical aperture, and the front aperture have been labelled for clarity.

suppress photoelectrons, and then pulse it negative with respect to the front aperture to extract IC electrons. The second would be to hold the front plate of the detection MCP at a high positive voltage such that there would be no gain between it and the second MCP plate while the chamber was flooded with background, and then to pulse it down to a lower positive voltage to turn on the MCP gain.

To accomplish these tasks two high-voltage (HV) pulser circuits were designed and constructed. These pulsers were built and tuned to be able to smoothly pulse voltages at a rate of up to 25 V/ns. The pulse timing was set by the Q-switch TTL of the YAG pump laser, and a delay circuit stage was in both pulsers to allow us to continuously vary the delay in the pulsing for up to 1.5 μs.

4.2.6 MCP Readout

As the anode of the MCP was biased to +2500 V, the electron signal was readout through a bias tee. Because of the high-voltage pulsing in our experiment, there were quite significant oscillations on the AC-coupled anode readout (see Fig. 4.11). These

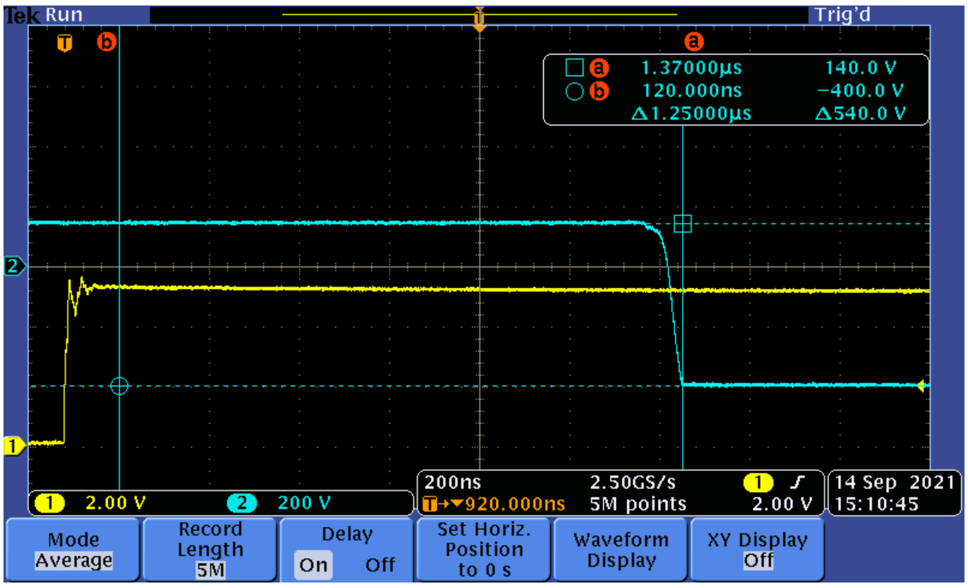

Figure 4.9: Scope trace of the pulsing of the ThO_2 target from +135 V to -405 V and its offset from the YAG Q-switch TTL. Voltage on the target is the blue trace, while laser Q-switch TTL is shown in yellow.

traces were then sent to our waveforms digitizer (CAEN DT5751) for processing.

These MCP traces would contain the small spikes from electron impacts over this large oscillating waveform. Because the pulsing of the MCP was well-defined in time, we could use an average waveform of the MCP signal with no electrons, and subtract that off the MCP traces acquired during the actual experimental runtime. This background subtraction worked remarkably well to flatten out the oscillations within the IC detection time window of [2,38] μs (see Fig. 4.16 for examples of traces with the oscillating waveform subtraction in action).

4.3 Calibrations

4.3.1 Stepper Motor Position and Beam Pointing

Originally we had expected to use fluorescence from the small fused silica window and the pyro-electric detector to determine if we were aligned along the axis with the ThO_2 target. However, it turns out that even the cheapest fused silica windows from Edmund Optics are good enough these day to be free of VUV fluorescing color centers (for reference, ten years ago one would even expect to see contaminant fluorescence

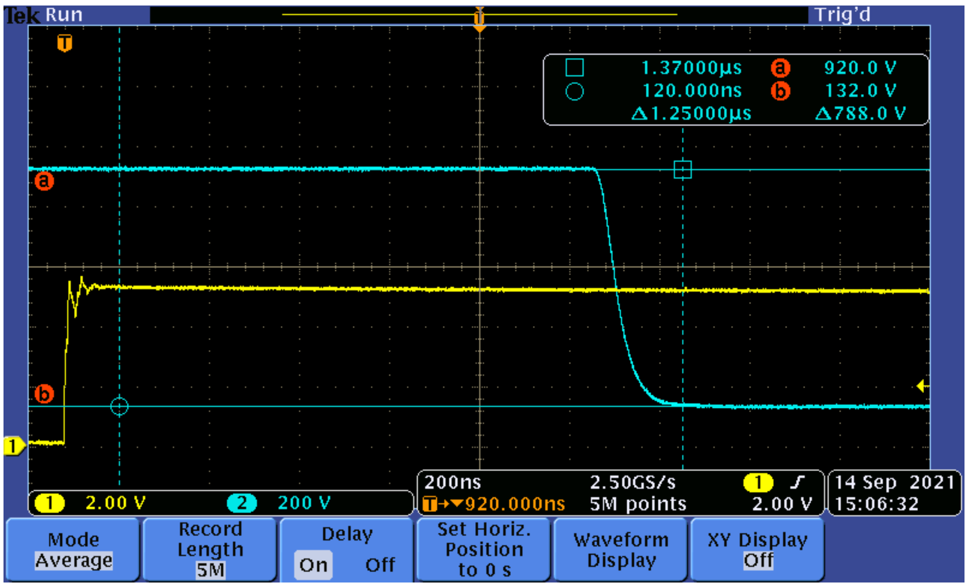

Figure 4.10: Scope trace of the pulsing of the MCP front plate from +750 V to +140 V and its offset from the YAG Q-switch TTL. Voltage on the MCP is the blue trace, while laser Q-switch TTL is shown in yellow.

from “high-quality” MgF_2). While this is a remarkable feat of the market, it meant we had to take a somewhat circuitous route to center the chamber.

The first step was to move the stepper motor some well defined number of steps (say 10,000), and measure how far that moved the chamber. From this we were able to determine the calibration between stepper motor encoder position and actual distance, which turned out to be approximately 1.1 mm per 1000 encoder steps.

The next step was to let in the VUV laser and monitor two signals on the readout board, the pyro voltage signal and the combined photoelectron signal from the ThO_2 target and front aperture. As can be seen in Fig. 4.12, clear features could be seen in both signals that corresponded to the various components of the ThO_2 mount. We could thus push our chamber around laterally while comparing the length of these voltage features to CAD drawings of the mounting structure. Centering was achieved when the measured feature lengths were in agreement with the center lengths of the drawings.

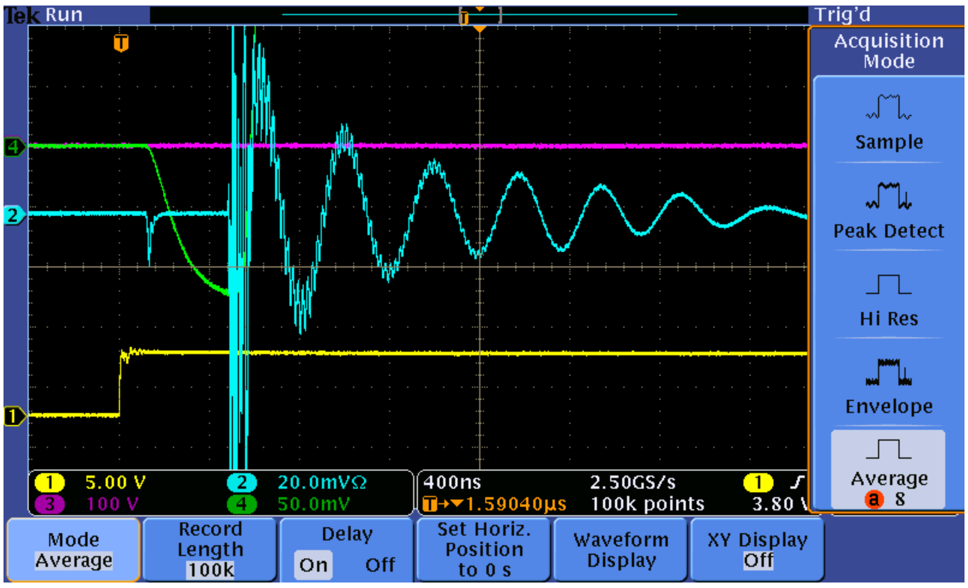

Figure 4.11: AC-coupled anode readout from the MCP (blue trace). Note the large oscillations last for several microseconds. The signal spike in the leftmost part of the blue trace is from the laser pulse prior to the MCP gain being pulsed on.

4.3.2 Photoelectron Extraction Efficiency

As noted previously, the closest one can get to an experimental measurement of η_{esc} is the photoelectron exctraction efficiency. The extraction efficiency was determined by measuring the number of electrons generated by a VUV pulse and dividing that by the number of photons in the pulse. The number of photoelectrons was determined by measuring the voltage pulse on AGND, which serves as an effective charge reservoir for the ThO_2 target (schematic of the pyro/target board circuit can be see in Fig. 4.13). This pulse could then be divided by the 50 Ω resistance of the oscilloscope while the laser was hitting the ThO_2 target. The number of photons per pulse could be determined by the energy measurement on the target mount pyro detector. It was found that biasing the front aperture at +135 V and the target at -405 V yielded an extraction efficiency of 1.0(5)e-3. It should be noted that no saturation was seen in the extraction efficiency with respect to the differential biasing of the front aperture and target, the limit in our experiment was purely what our circuit board could handle.

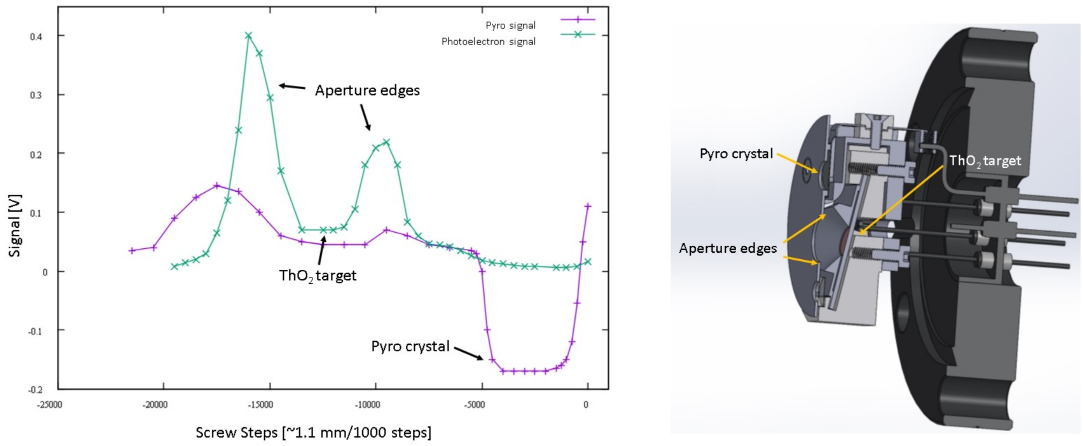

(a) Features in the photoelectron signal (green) and pyro signal (purple) with assigned features.

(b) Corresponding features in CAD mockup.

Figure 4.12: Photoelectron and pyro signal mapping of the ThO_2 mount. The characteristic features in the signals could be unambiguously assigned to parts of the mounting structure.

4.3.3 Photoelectron Collection Efficiency

The IC electron collection efficiency was determined by using photoelectrons generated at the ThO_2 target. We determined this value by measuring the amount of charge that hit the detection MCP and dividing by the number of photoelectrons that were generated by the VUV pulse (the latter value is determined in the same manner as for the photoelectron extraction efficiency). The amount of charge hitting the MCP was determined by measuring by measuring the current from the front and back MCP plates into an oscilloscope with no gain applied. The measured collection efficiency for the combined voltage biasing and magnetic field configuration used in the experiment was $\sim$0.5.

4.4 Backgrounds: Studies and Suppression

In the initial proposal for the experiment, the only expected backgrounds were the initial photoelectron pulse which would have a timescale on the order of a few hundred

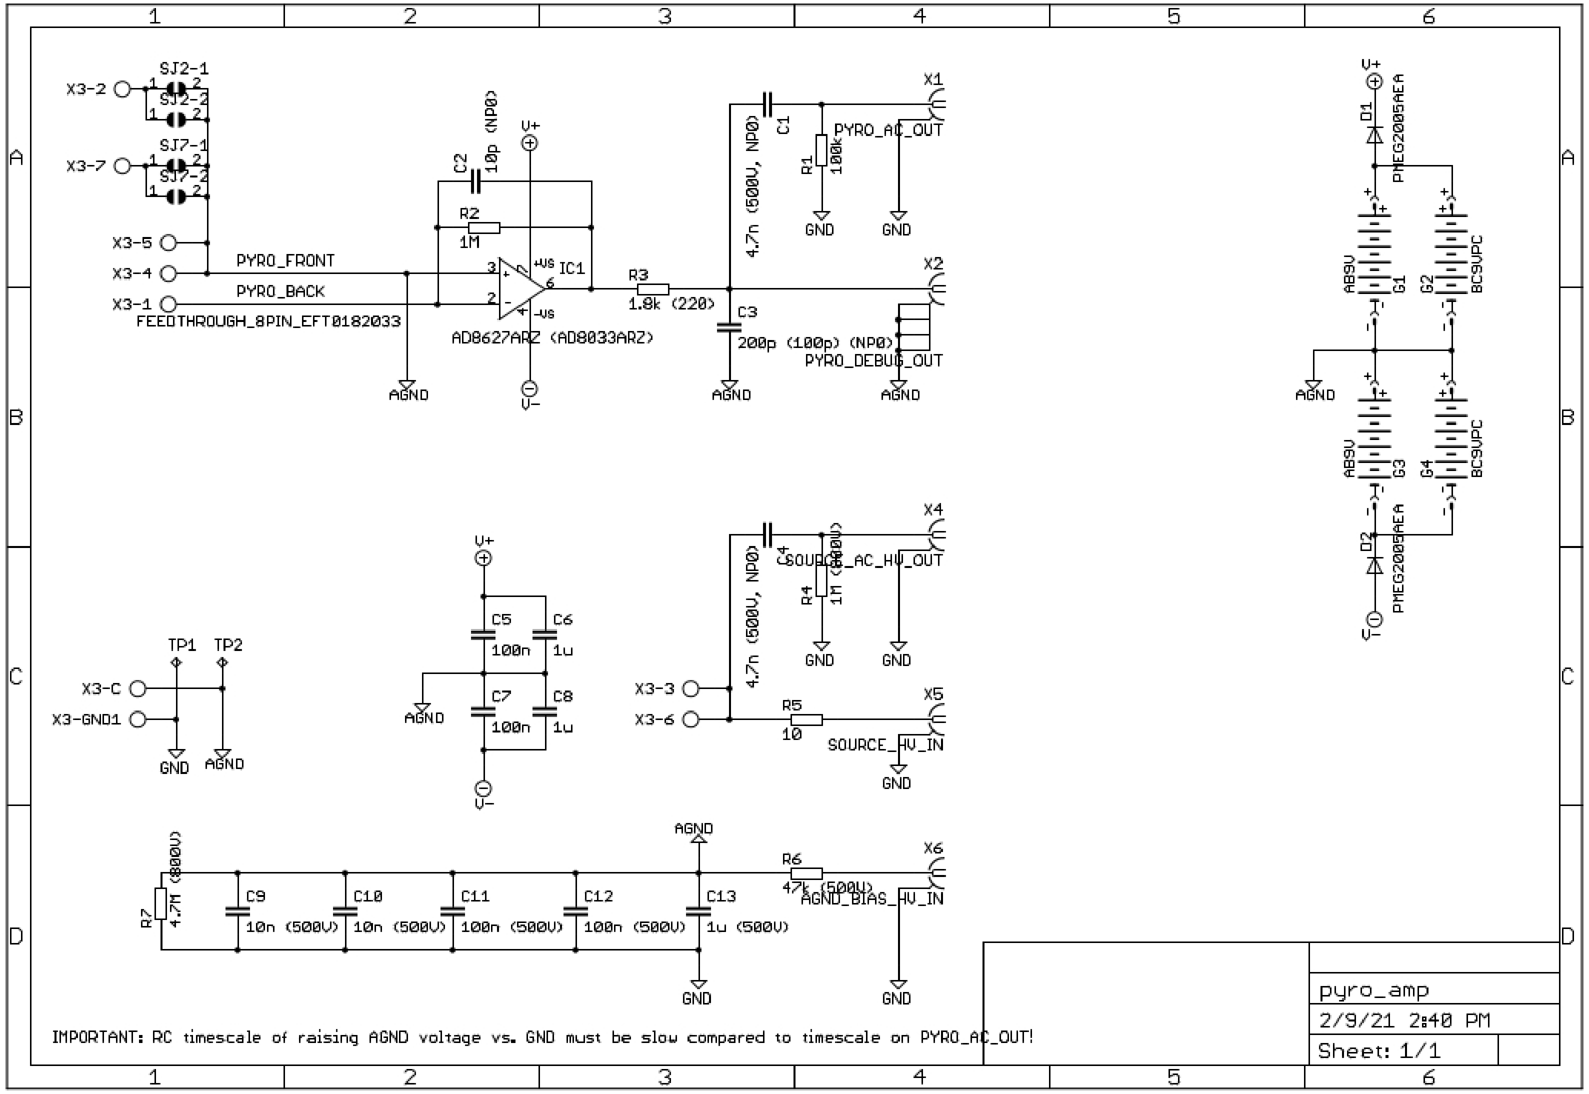

Figure 4.13: Schematic of pyro readout board used in CEMS experiment (Designed by Dr. Christian Schneider).

nanoseconds, and the small constant stream of electrons from the target due to the radioactivity. However, it turned out there was a large amount ($\gtrsim$ 1e7 s^{-1}) of electrons squarely within the IC detection window of tens of microseconds. It was imperative to both understand the origin of and suppress this background.

4.4.1 Radiation Trapping and ASE in Xe

A large portion of the background was found to be due to a fairly well-documented effect in Xe four-wave mixing. When the 249.6 nm laser is on resonance with the $5p^6\ ^1S_0 \rightarrow 5p^5\left(^2P^\circ_{3/2}\right)6p\ ^2[1/2]_0$ transition, whatever Xe does not participate in the four-wave mixing process will be left in the excited two-photon state. This $5p^5\left(^2P^\circ_{3/2}\right)6p\ ^2[1/2]_0$ (abbreviated as $6p\ ^2[1/2]_0$) state will then decay to the $5p^5\left(^2P^\circ_{3/2}\right)6s\ ^2[3/2]^\circ_1$ (abbreviated as $6s\ ^2[3/2]^\circ_1$) by 828 nnm emission in ≈30 ns, and the $6s\ ^2[3/2]^\circ_1$ state will decay to the $^1S_0$ ground state by emission of a 147 nm photon in ≈3.7 ns.

If the Xe pressure is in the several millibar range, as it is for efficient four-wave mixing, there are enough nearby Xe atoms that they may re-absorb these 828 nm and 147 nm photons. Effectively, this leads to radiation trapping which extends the effective fluorescence lifetime of the Xe spontaneous emission [AS96b]. At the same time the 828/147 nm spontaneous emission will experience gain as it stimulates Xe in the excited $6p\ ^2[1/2]_0$ and $6s\ ^2[3/2]^\circ_1$ to emit, yielding amplfied spontaneous emission (ASE). This gain will be highly directional, since the excited Xe will essentially lie in a column defined by the propagation of the 249.6 nm pulse of the pump laser. This interplay between radiation trapping and ASE will then yield bi-directional emission from the Xe cell along the pumping axis, which will have a timescale much longer than the spontaneous emission lifetime of the excited states [RDG89b]. In other words, you get a forward and back propagating "laser" that lasts for several microseconds (see for example Fig. 4.14)!

For our experiment the 147 nm ASE beam presents a problem, as it is essentially aligned along the beam axis high enough energy to excite photoelectrons on the ThO_2 in the IC detection time window. Already previous work had found that mixing another buffer gas of molecules into the Xe cell could suppress this the bi-directional

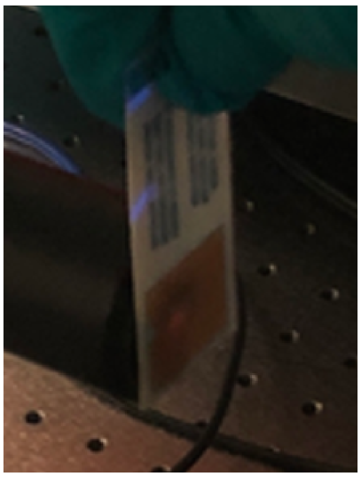

Figure 4.14: Back-propagating 828nm ASE seen on IR card (small orange dot). First evidence that biconical emission was occuring in our system.

emission through a variety of quenching mechanisms [AS96b]. Despite its relatively low quenching strength compared to other molecular gases, the relatively inert N_2 gas could be introduced into our system in large quantities without complication. It was found that a 19:1 mixturem of N_2:Xe was sufficient to suppress the background on our detection MCP from the $\gtrsim$1e7 s^{-1} level to $\sim$1e6 s^{-1} (a comparison of the background with and without quenching gas can be see in Fig. 4.16. The introduction of the buffer gas reduced the VUV power by roughly 10%, but this was seen as an acceptable trade.

4.4.2 Chamber Wall Emission

It turned out that the VUV light scattered from the ThO_2 target onto the stainless steel of the chamber yielded yet another background of photoelectrons. It is an unfortunate coincidence that this "chamber wall emission" also happens to have a time scale on the order of a few tens of microseconds. In the original proposal for the experiment, it was expected that the front plate of the detection MCP could be held at a large positive voltage compared to the rest of the chamber in order to suck up all the IC electrons that emerged from the target region. However, this would also suck up all of these chamber wall electrons.

From simulations it was determined that combination of voltage biasing of our MCPs and a static magnetic field could solve this issue. Essentially, one of the MCPs

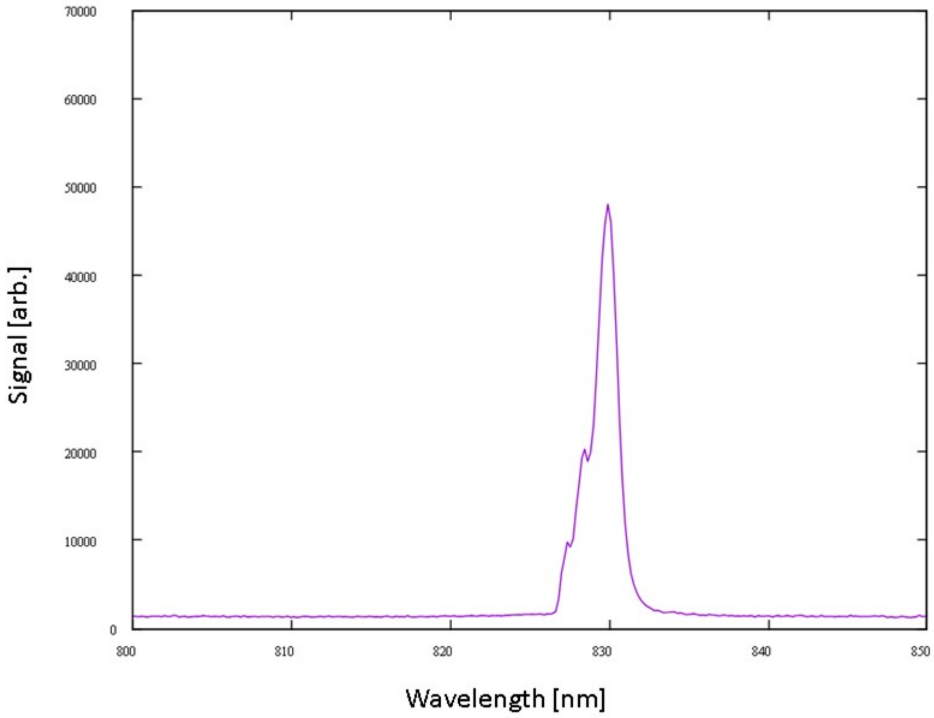

Figure 4.15: Ocean optics spectrum of the back-propagating beam to confirm that it is from the Xe $6p\ {}^2[1/2]_0 \rightarrow 6s\ {}^2[3/2]^\circ_1$ transition.

would be "sacrificed" and biased to a very high positive voltage to attempt to suck up all of the chamber wall emission, meanwhile the IC electrons that would emerge from the target region would be in the right location, as well as have the right velocity vectors, to be guided through a saddle point in the potential landscape that would lead them to the detection MCP, as show in Fig. 4.17. The magnetic field would need to be a few gauss pointed up with respect to the plane of the table. This would be provided by an internally water cooled copper coil with 6 wraps and 20" diameter on the winding. This coil was placed on top of the chamber, roughly 8" above the plane the electrons travel in.

After an iterative process of simulations and measurements of the photoelectron collection efficiency, we settled on the following configuration for our "detection mode":

- The ThO_2 target+conical aperture would be held at -405 V and the front aperture would be held at +100 V.

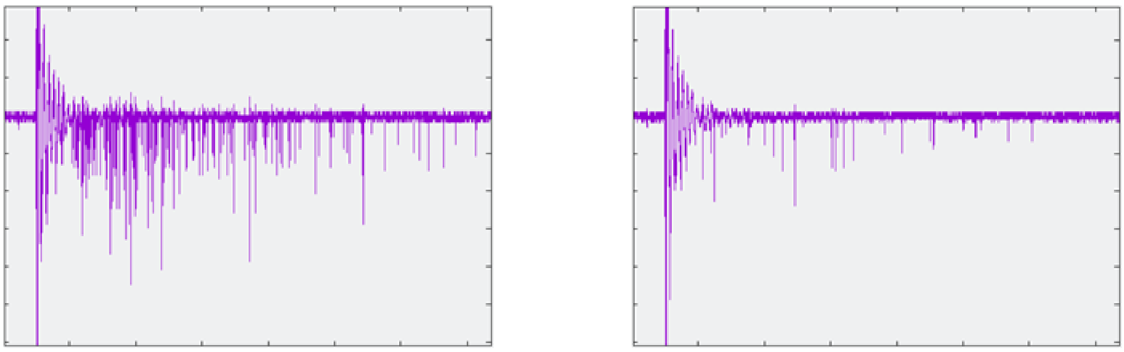

(a) MCP trace with no N_2 quenching of the Xe gas during four-wave mixing.

(b) MCP trace with 19:1 mixture of N_2:Xe.

Figure 4.16: Comparison of electron background on MCP with and without N_2 quenching of the Xe cell. Each horizontal tick represents 1 μs.

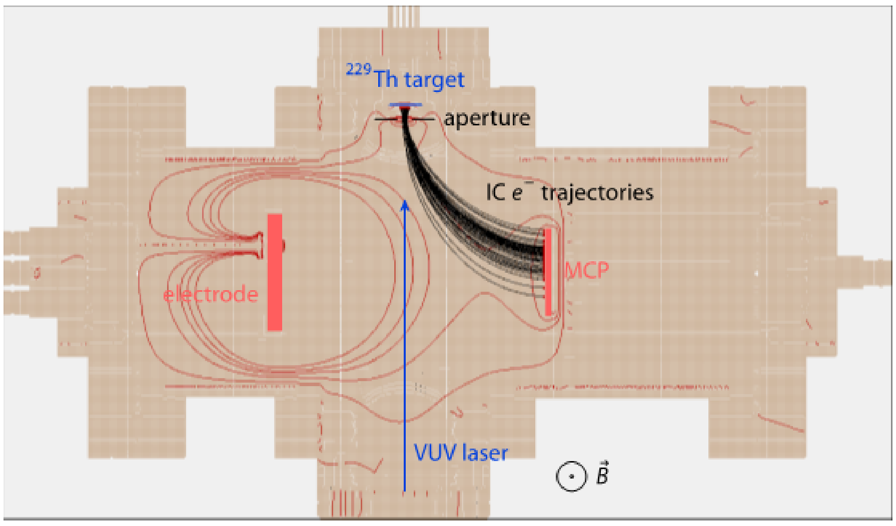

Figure 4.17: SIMION trajectories for electrons accelerated out of the target region. Note how they are guided through a saddle region in the potential landscape due to the Lorentz force from the static magnetic field (pointed out of the plane towards the viewer). The “suction” MCP is labelled as “electrode”.

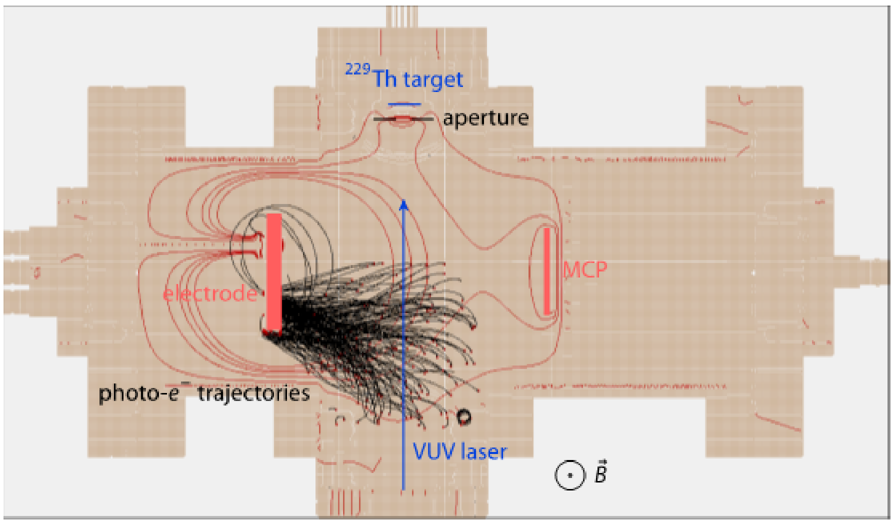

Figure 4.18: SIMION trajectories for electrons generated by scattered VUV light. The combination of the large biasing of the "suction" MCP and the magnetic field either causes them to be sucked up into the "suction" MCP or crash into the chamber walls.

- The "suction" MCP would be held at +2500 V.
- The front plate of the detection MCP would be held at +140 V.
- The magnetic coil would be operated at 96 A, yielding a magnetic field at the center of the chamber of ≈5.5 gauss.

With this configuration the background was suppressed to the $\lesssim$1e6 s^{-1} level.

4.4.3 ThO_2 Fluorescence

The final background seemed to be from the target region, and was tied to laser illumination of target. This background generated electrons in the IC time window, and had a linear relation with the laser power. This background was also found to have a $1/t$ time dependence with respect to when the laser pulse arrived, meaning that the signal-to-noise ratio was now a function of the integration window used in the experiment. The exact behavior of this background was better understood over the course of the experiment, as its dependence on laser power could be scrutinized. More details about this background will be given in the analysis section.

4.5 Experimental Sequence and Data Acquisition

Each datapoint involved a single direction sweep of the laser grating over 500 grating positions. This constituted a scan over ≈ 500 pm, or ≈250-270 GHz within the wavelength range scanned by the dye laser. Each data point was overlapped by 40 GHz so as to avoid the chance of missing the transition in any particular datapoint due to power instabilities. During the duration of the scan both the UV and VIS PDL wavelengths would be measured by a wavemeter (Coherent Wavemaster). In addition, the spectrum of the VIS PDL would be monitored using a spectrometer (Ocean Optics USB4000).

For the data acquisition the following sequence was done for each data point:

1. The UV PDL beam would be scanned over the Xe two-photon resonance and emission at 828 nm would be monitored in the spectroscopy cell. This was to check the calibration of the wavemeter.

2. The chamber was moved to have the VUV laser hit the pyro detector. An initial sweep of the VIS PDL stepper motor would be done from the maximum grating position to the bottom with one laser pulse per grating position (≈ 16.6 s). This was to check for any wavelength dependence in the VUV generation. Each pulsed signal from the pyro detector was acquired by a oscilloscope (Digilent Analog Discovery 2) that was externally triggered by the YAG Q-switch TTL.

3. The chamber would be moved to overlap the VUV beam with the ThO_2 target. The acquisition of the MCP traces would then take place. Each grating position would be illuminated for 120 pulses before the grating was stepped. The entire time of this sequence would be 2000 s.

4. From the VUV energy readings the YAG amplifier would be incremented if the VUV power fell below 1 μJ/pulse.

Between 12 sets of datapoints the dye changes, laser maintenance, and realignment would be carried out. This allowed for 24 datapoints per day at stable operation.

4.5.1 Daily Reports

Each day the datapoints were ran through an analysis suite to make sure that the scan occured, there was VUV power, the MCP traces were acquired, etc. An example of such a daily report can be seen in Fig. 4.19. If there were any abnormalities in the report, the datapoint would be flagged for re-scanning.

4.6 Data Analysis and Results

Over the course of scanning the 3σ region defined in [SWB19], a new result from the ISOLDE collaboration narrowed the uncertainty in the transition from 149.7(3.1) nm to 148.71(42) nm [KMA23]. Since we had scanned over this region, and having not seen any clear signal by eye in the daily reports, we decided to perform a more careful analysis of the data to see if there were any points of interest.

4.6.1 Background Analysis

We first wanted to see the effect of the integration time window on the background (what was dubbed earlier as the "$\mathrm{ThO_2}$ fluorescence"). To do this, we took selected five possible time windows: [2,38],[4,38],[6,38],[8,38], and [10,38] μs. For a particular time window, all of the electron counts in a datapoint would be integrated and then binned by grating position. For each datapoint this would yield 5 traces, as shown in Fig. 4.20.

The average counts (μ) and standard deviation (σ) of the bin-wise counts of each integration window would then be calculated for each datapoint. We could then generate scatter plots of the μ versus σ, as well as μ versus laser power to look for any trends. The following facts about the background were determined:

- The background exhibited fluctuations that were super-Poissonian ($\sigma > \sqrt{N}$). The excess fluctuations could be captured by scale factor s, such that $\sigma = s\sqrt{N}$, but this scale factor was a function of the lower integration bound.
- The background μ was found to be linearly correlated with the laser power, such that $\mu = \beta E$, where E is the average energy per VUV pulse. It was found

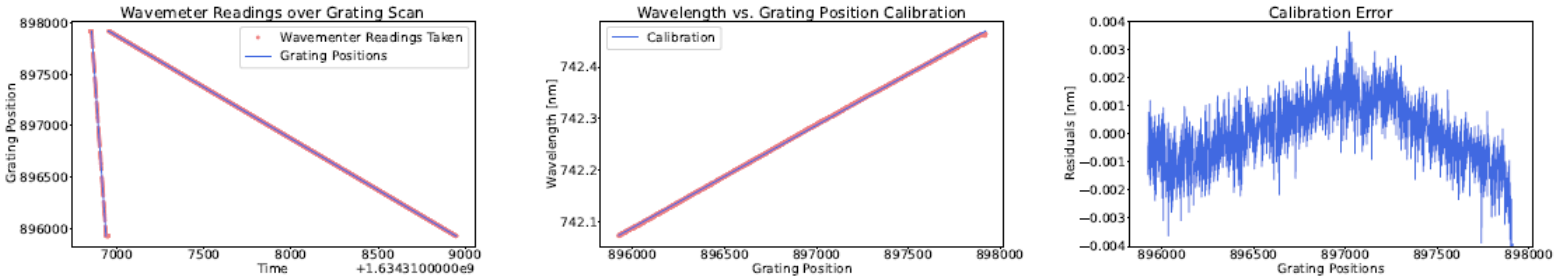

(a) Wavemeter Readings over VIS PDL Scan (b) Wavelength vs. Grating Position Calibration (c) Error on Calibration

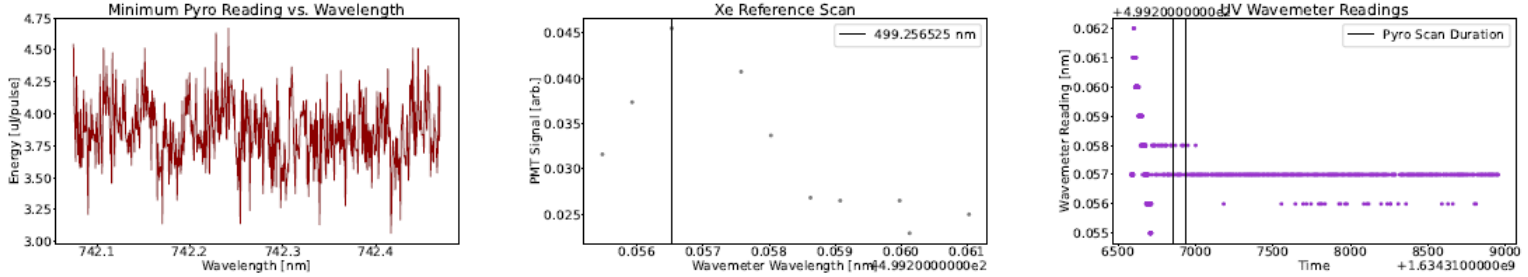

(d) Minimum Energy over Wavelength Scan (e) Xe Reference Scan (f) UV Wavelength Readings

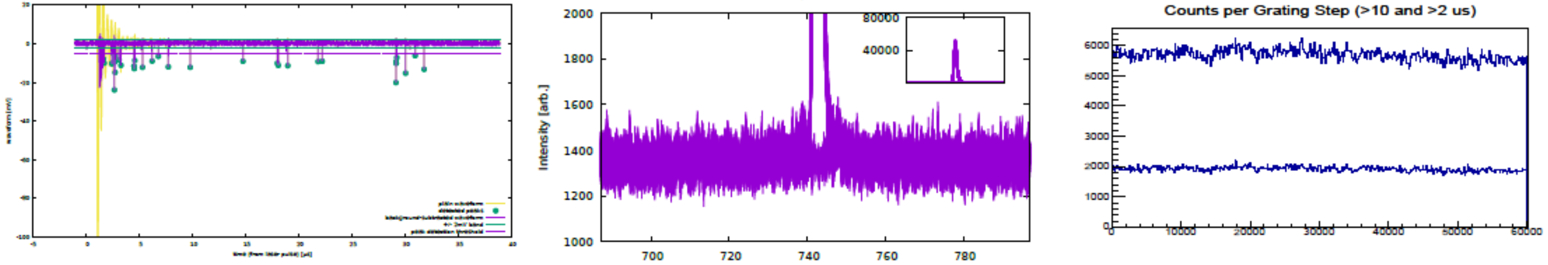

(g) MCP Trace (h) Ocean Optics Readout (i) Raw Counts

Data Point: .//2021-10-15_ic_data_run/000001_224480

Figure 4.19: Example of a daily scan report for the IC metal experiment.

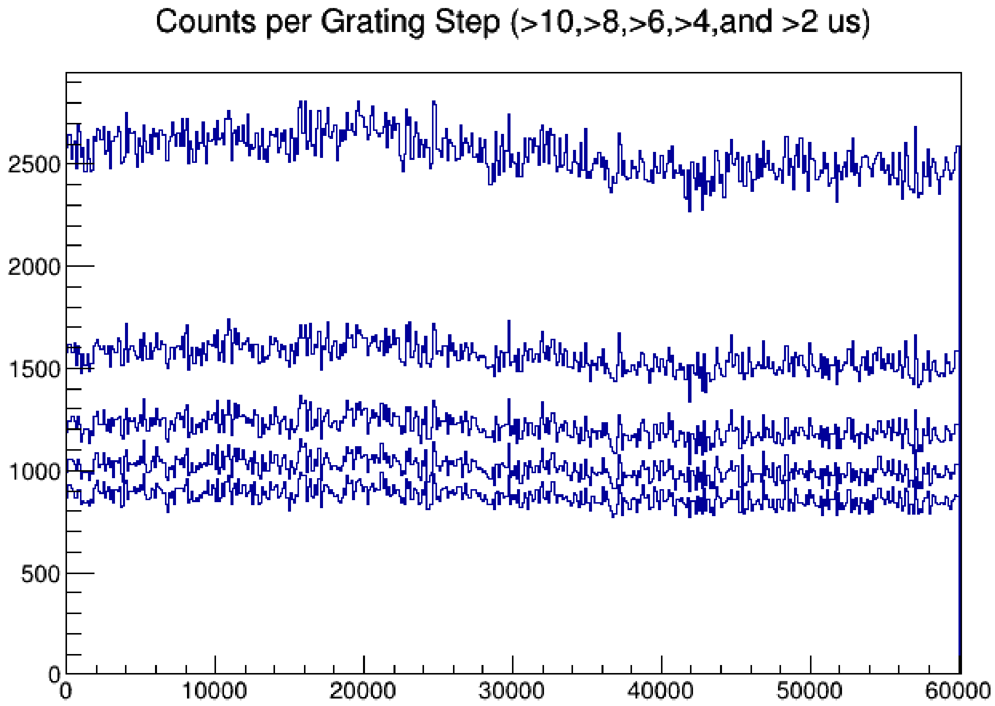

Figure 4.20: Example of the integrated counts in each grating step for [2,38],[4,38],[6,38],[8,38], and [10,38] μs integration windows. The x-axis is in units of laser shots, and so needs to be divided by 120 to get the grating step. With 500 grating steps, this represents roughly 350 pm.

that β was also a function of the lower integration bound.

The results of these background studies could then be summarized in Table 4.1.

Integration Window [μs]	s [#]	β [kHz/μJ]
$[2, 38]$	2.03	458
$[4, 38]$	1.79	286
$[6, 38]$	1.59	230
$[8, 38]$	1.52	202
$[10, 38]$	1.47	185

Table 4.1: Variation in the background s and β parameters as a function of the integration window used.

4.6.2 Optimum Integration window

With a model of the background, one can now begin to analyze what time window would yield the highest statistical power to detect the signal. In other words, for a given test statistic with a set false alarm rate, we seek the time window has the highest chance of flagging the IC signal. Our tool of choice to determine this is a Monte Carlo analysis.

The Monte Carlo analysis proceeded as follows:

1. Choose a particular integration window from $[t_1, t_2] = [\text{x},38]$ μs. One must also assume a certain value for unknown experimental parameters (η_{esc} and $\mathit{\Gamma}_L$).

2. Generate 500 bins covering 250 GHz (in accordance with the experimental protocol).

3. In each of these frequency bins, generate an amount of background counts B from a Gaussian distribution. The mean $\langle B \rangle$ should equal to the time window width $\beta(t_2 - t_1)$, and the width should be equal to $s\sqrt{\langle B \rangle}$, with β and s pulled from the appropriate row of Table 4.1 (or by linear interpolation from the table).

4. Generate signal counts N_{sim} in each of the frequency bins according to

   $$N_{sim} = N_{exc}\left(\exp(-\mathit{\Gamma}_{IC}t_1) - \exp(-\mathit{\Gamma}_{IC}t_2)\right), \tag{4.3}$$

   where N_{exc} is calculated based on Eq. 4.1.

5. Perform two fits on the data, with both fits seeking to minimize

   $$\chi^2 = \sum_i \frac{(y_i - f_i)^2}{y_i}, \tag{4.4}$$

   where f_i is the model prediction for a grating position i, and y_i is the counts in a bin. The first fit is a simple linear background given by $ai + b$ where a and b were allowed to float. The second is a Lorentzian plus background model of the form

   $$f(i) = \frac{A}{2\pi} \frac{\mathit{\Gamma}}{(i - x_0)^2 + (\frac{\mathit{\Gamma}}{2})^2} + a'i + b', \tag{4.5}$$

where $(A, x_0, \mathit{\Gamma}, a', b')$ are allowed to float. $\mathit{\Gamma}$ was bounded to not exceed a width of 40 GHz, to prevent degenerate fits. The second fit is essentially looking for an excess of counts above background that had the shape of a resonance.

6. Form the signal statistic by computing the difference in χ^2 between the background only and signal+background models (dubbed $\mathit{\Delta}\chi^2$).
7. Record this value of $\mathit{\Delta}\chi^2$, and then repeat from Step 2.

The Monte Carlo procedure was peformed multiples times to be build up a distribution of $\mathit{\Delta}\chi^2$ for each time window, as well as for different values of η and $\mathit{\Gamma}_L$. A particular value of $\mathit{\Delta}\chi^2_{th}$ was chosen to yield a false alarm rate of 5% (a detailed discussion of the false alarm rate (FAR) and the determination of $\mathit{\Delta}\chi^2_{th}$ is reserved for the appendix). The statistical power of a time window was defined as the fraction of the $\mathit{\Delta}\chi^2$ distribution that was beyond a this value of $\mathit{\Delta}\chi^2_{th}$, i.e. its chance to detect the IC signal with a given threshold. As can be seen in Fig. 4.21, the ideal lower integration bound is around 5-6 μs, balances cutting off the large initial background while catching enough of the 10 μs IC signal.

4.6.3 Combing for a Signal

With the ideal integration windows determined, then applied the $\mathit{\Delta}\chi^2$ statistic to all our data within the ISOLDE region. In either integration window, 11 points stood out for further examination. For the [5,38] μs window they stood out with $\mathit{\Delta}\chi^2 > 60$. Several points had "resonance-like" feature, i.e. small bumps on the order of 5-40 GHz, but none had time-scale that was consistent with the internal conversion channel. In fact, all the "features" observed decayed with timescales that were consistent with the time variation of the background. It was concluded that all features were consistent with background fluctuations due to intermittent fluctuations in the laser power over the duration of the scan.

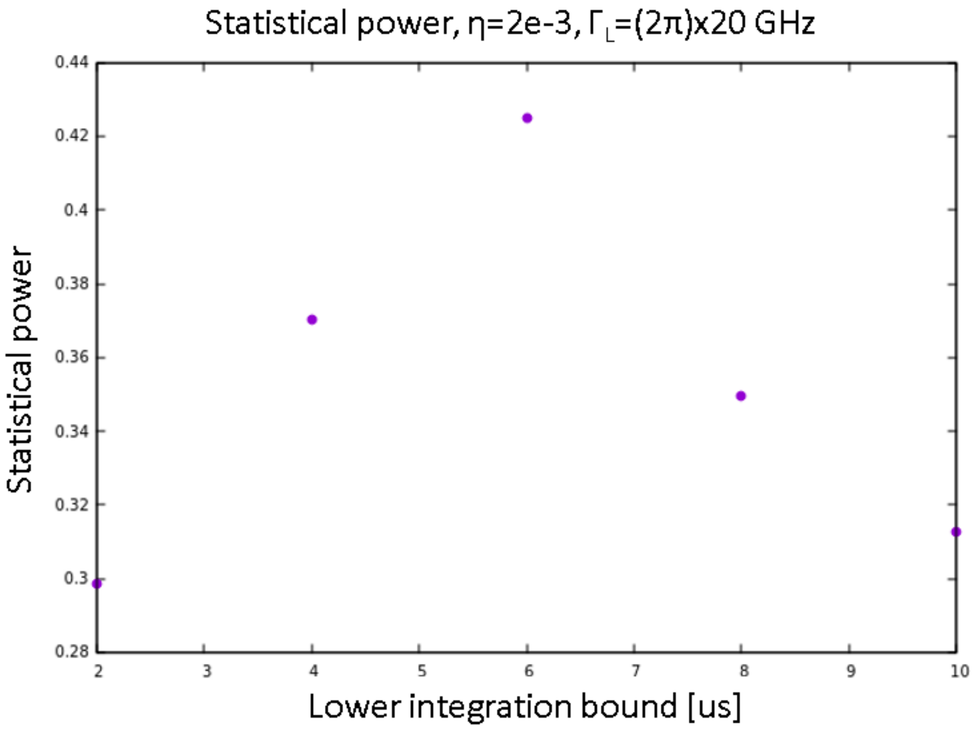

Figure 4.21: Statistical power for integration time windows of the form [x,38] μs and false alarm rate of 0.05. The assumed IC electron extraction efficiency was $\eta_{esc} = 2 \times 10^{-3}$ and the laser bandwidth was assumed to be $\Gamma_L = 2\pi \times 20$ GHz.

4.7 Conclusions and Speculation

Before the transition was known to < 1 GHz certainty, the lack of any observed signal in the ISOLDE region led us to conclude that we would be unable to observe the transition with the current scan parameters. It is likely that the IC electron extraction efficiency from the $\mathrm{ThO_2}$ is lower than anticipated. During the course of the experiment, the group of Jun Ye had discovered that the E&Z targets were not as flat as was originally thought, and seemed to be host to a series of contaminants (private communication with Lars von der Wense and Chuankun Zhang). These contaminants may impede the IC electrons escape from the surface. Another concern about the contaminants is that they may have contributed to the photoelectron signal, making the the photoelectron extraction efficiency even more disparate from the IC electron extraction efficiency.

Another concern about is whether the lifetime of 10 μs is actually a valid estimate of the internal conversion rate for ^{229m}Th in the $\mathrm{ThO_2}$ environment. In all exper-

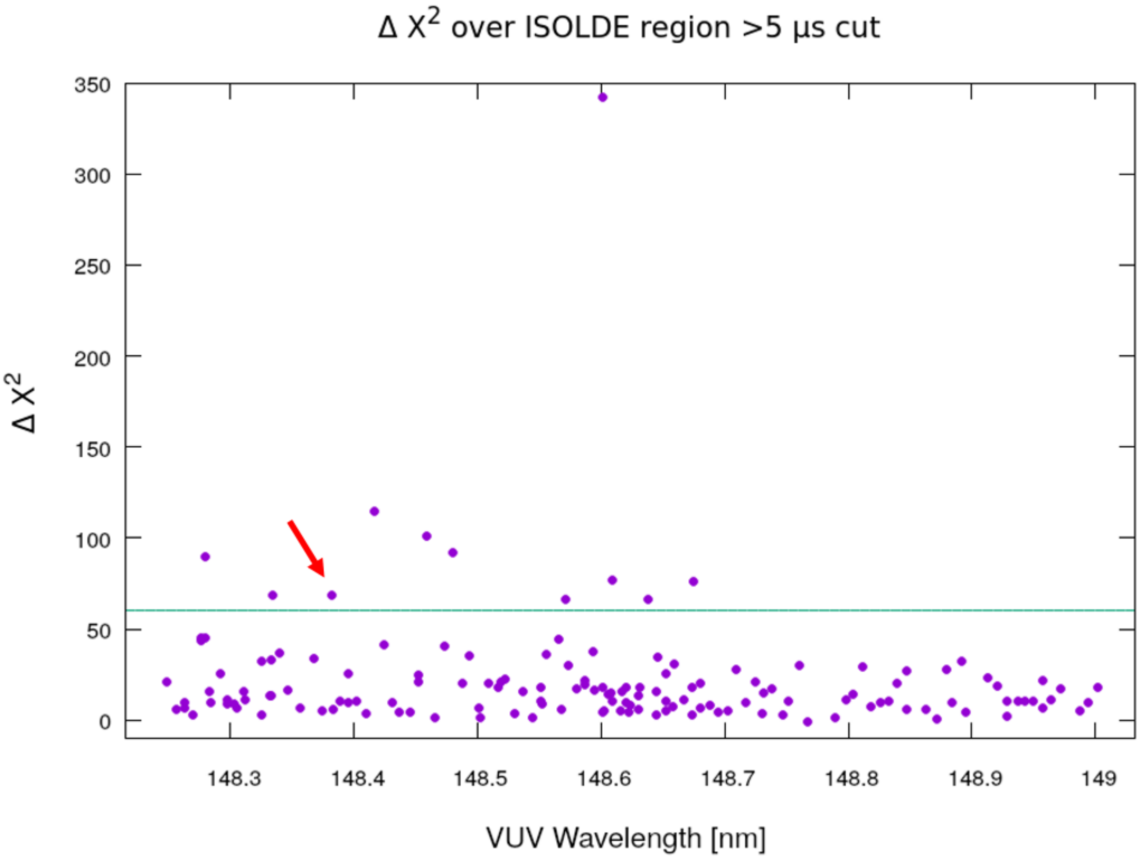

Figure 4.22: $\Delta\chi^2$ for the datapoints collected within the ISOLDE region. The red arrow marks the datapoint that overlaps with where the isomeric transition would be later found.

iments were the internal conversion lifetime has been measured, ^{229m}Th has been soft-implanted onto a surface. The local electronic environment of the soft implated ^{229m}Th could potentially be quite different from the electronic environment in the ThO_2 matrix. Since the internal conversion coefficient depend quite sensitively on the electronic environment [SAZ90, BMP18b], it may be possible that the internal conversion rate lies outside the 10μs time window.

Now that the transition has been resolved to a precision of < 1 GHz, it is likely worth it now to retry this experiment, increasing the integration time to increase the sensitivity to lower IC electron extraction efficiency and looking at different detection time windows. Improved theoretical calculations of the internal conversion coefficients in ThO_2 may also help in the search, and would serve as a good benchmark for a calculation of hyperfine interaction matrix elements in the solid state.

4.8 Appendix

4.8.1 False Alarm Rate and $\Delta\chi^2_{th}$

To pick out a signal in a datapoint, we would compute $\Delta\chi^2$ between the signal+background and background only fits, and then flag the point as containing a signal if $\Delta\chi^2$ exceeds some threshold value. The false alarm rate (FAR) is the probability that a datapoint containing no signal will be flagged at a given threshold value of $\Delta\chi^2$. The threshold can be set as low as one wants, but there is always a trade-off. As you increase your detection power you will inevitably start having more and more false alarms. One typically has to set some value of the false alarm rate they are willing to accept, and determine their detection power with that constraint. For our experiment, a FAR of 0.05 is acceptable as it consitutes about one false alarm per day.

To determine the threshold value $\Delta\chi^2_{th}$ that yields a FAR of 0.05, one must perform a shortened version of the Monte Carlo procedure explained in the detection power study:

1. Choose a particular integration window from $[t_1, t_2] = [\text{x},38]$ μs.
2. Generate 500 bins covering 250 GHz (in accordance with the experimental protocol).
3. In each of these frequency bins, generate an amount of background counts B from a Gaussian distribution. The mean $\langle B\rangle$ should equal to the time window width $\beta(t_2 - t_1)$, and the width should be equal to $s\sqrt{\langle B\rangle}$, with β and s pulled from the appropriate row of Table 4.1 (or by linear interpolation from the table).
4. Perform two fits on the data, with both fits seeking to minimize

$$\chi^2 = \sum_i \frac{(y_i - f_i)^2}{y_i}, \tag{4.6}$$

where f_i is the model prediction for a grating position i, and y_i is the counts in a bin. The first fit is a simple linear background given by $ai + b$ where a and

b were allowed to float. The second is a Lorentzian plus background model of the form

$$f(i) = \frac{A}{2\pi} \frac{\mathit{\Gamma}}{(i - x_0)^2 + (\frac{\mathit{\Gamma}}{2})^2} + a'i + b', \tag{4.7}$$

where $(A, x_0, \mathit{\Gamma}, a', b')$ are allowed to float. $\mathit{\Gamma}$ was bounded to not exceed a width of 40 GHz, to prevent degenerate fits.

5. Form the signal statistic by computing the difference in χ^2 between the background only and signal+background models (dubbed $\mathit{\Delta}\chi^2$).

6. Record this value of $\mathit{\Delta}\chi^2$, and then repeat from Step 2.

One performs this Monte Carlo multiple times on simulated datasets containing no signal. One then obtains a histogram of $\mathit{\Delta}\chi^2$ values for each detection window. For a particular time window, the value of $\mathit{\Delta}\chi^2_{th}$ is equal to the value of $\mathit{\Delta}\chi^2$ such that 5% of Monte Carlo trials lie above that value.

CHAPTER 5

Short-Timescale Fluorescence in ^{229}Th:LiSrAlF$_6$

5.1 Prelude: ISOLDE Results

Having not observed the expected signal from the isomeric transition in the LiSrAlF$_6$ crystals with our VUV laser system, we began to consider whether the transition was being non-radiatively quenched. Around early 2022, the first reports from the group of Piet Van Duppen were publicized that the radiative decay of the transition had been observed in MgF$_2$ and CaF$_2$. The group had used the ISOLDE radioative beam line at CERN to implant ^{229}Ac into crystal samples, which were then rotated into a precision VUV spectrometer system. This ^{229}Ac undergoes β-decay into ^{229}Th. Of these β-decays, $\sim$14% populate the isomeric state. The isomeric state will then radiatively decay via VUV photons in the high band-gap crystal systems into which they have been implanted. By observing the emitted VUV photons through a spectrometer, they may be separated from the background and the energy of the transition was measured to 148.71(42) nm (8.338(24) eV) [KMA23]. It was noted in this experiment that given an estimate of the experimental efficiencies, it appeared that only a fraction of the implanted isomeric ^{229}Thwas participating in the radiative decay channel.

5.2 Short-Timescale Fluorescence

Given these results, we began investigations on the data collected with our VUV laser system for any signatures of the transition in the 1σ region ISOLDE had determined the transition to lie (hereafter referred to as the ISOLDE region). Combing through the data it was found that as one crossed over the ISOLDE region with the excitation wavelength, there was a dramatic increase in a fluorescence emitted from the crystal that had a lifetime of a few seconds (see Fig. 5.1). The number of fluorescent photons in the first 50s normalized to laser power N_γ/P_L was plotted for all the data taken in the 2016-2019 scan, and forms the blue dataset seen in Fig. 5.2

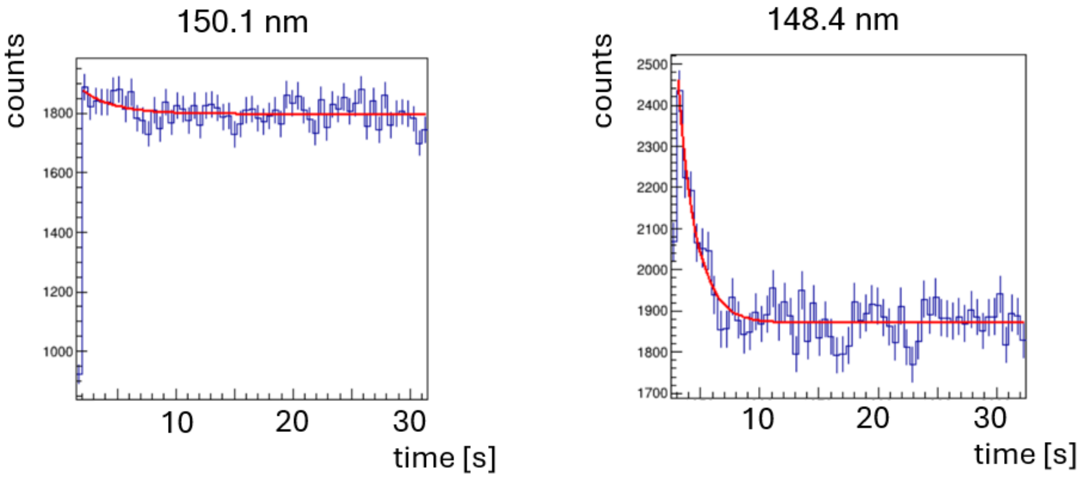

Figure 5.1: Comparison of the seconds timescale fluorescence observed in the ^{229}Th :LiSrAlF$_6$ crystal 2.2. (Left) 150.1 nm excitation red of the ISOLDE region. (Right) 148.4 nm excitation blue of the ISOLDE region.

A clear excess of detected photons can be seen in the blue dataset of Fig. 5.2, which shows N_γ/P_L in the first 50 s after the laser is extinguished as a function of excitation wavelength. This fluorescence appears to grow considerably as the excitation wavelength approaches the region identified by the ISOLDE experiment [KMA23]. Analysis suggests that this fluorescence possesses multiple timescales and is reasonably fit by a two-lifetime model of the form

$$f(t) = c_1 e^{-t/\tau_1} + c_2 e^{-t/\tau_2} + b, \tag{5.1}$$

as shown in Fig. 5.3(Top). The profiled χ^2 contours for these two lifetimes are shown in Fig. 5.3(Bottom). These contours were obtained by profiling the out the parameters (c_1, c_2, b) [Paw13].

Following this observation, several additional experiments were performed to ascertain the origin of the short-timescale fluorescence. For these experiments, the crystals were illuminated by the VUV laser for 60 s, the shutters opened in 100 ms, and the resulting fluorescence observed for roughly 60 s. First, a second ^{229}Th:LiSrAlF$_6$ crystal (batch 3.1) was investigated to confirm the existence of the spectral feature.

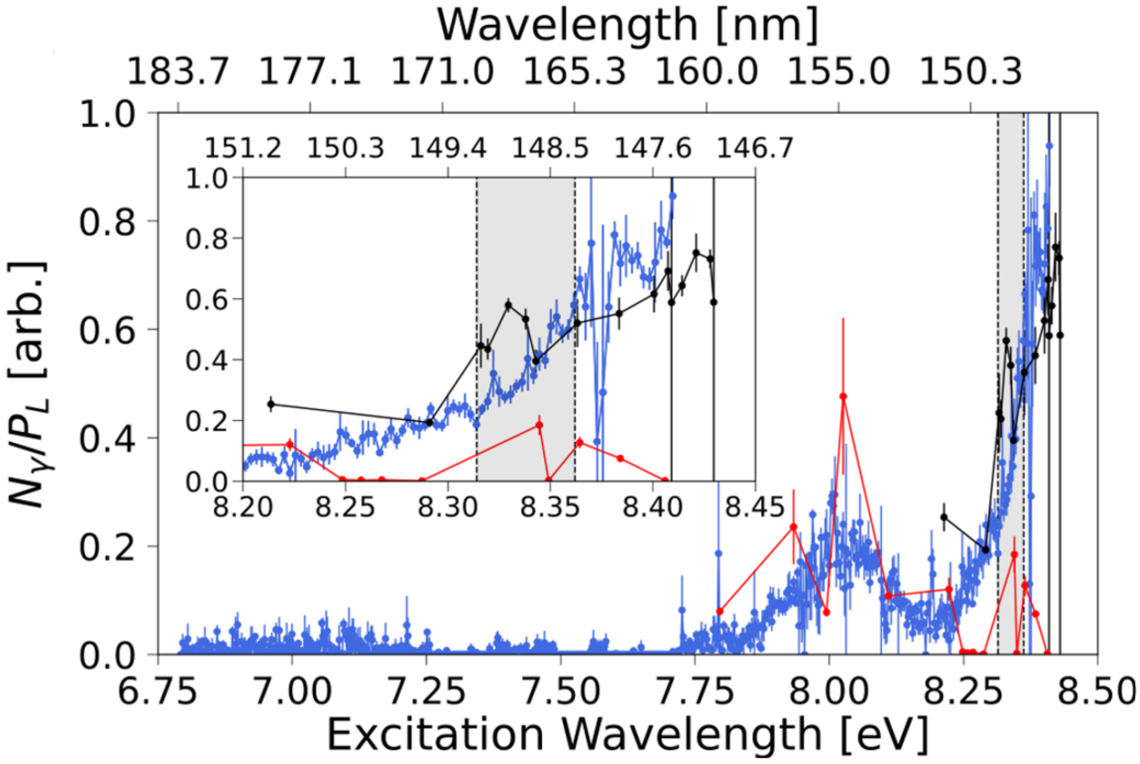

Figure 5.2: Short timescale fluorescence normalized by laser power vs. excitation energy in ^{229}Th (2.2 in blue, 3.1 in black) and ^{232}Th (red) $LiSrAlF_6$ crystals. The inset shows the region near the ISOLDE 68% confidence region (shaded). Each of the blue data points required roughly 1 hr to collect. The drop in signal near 8.37 eV is due to a resonance in the Xe gas used for generating the VUV.

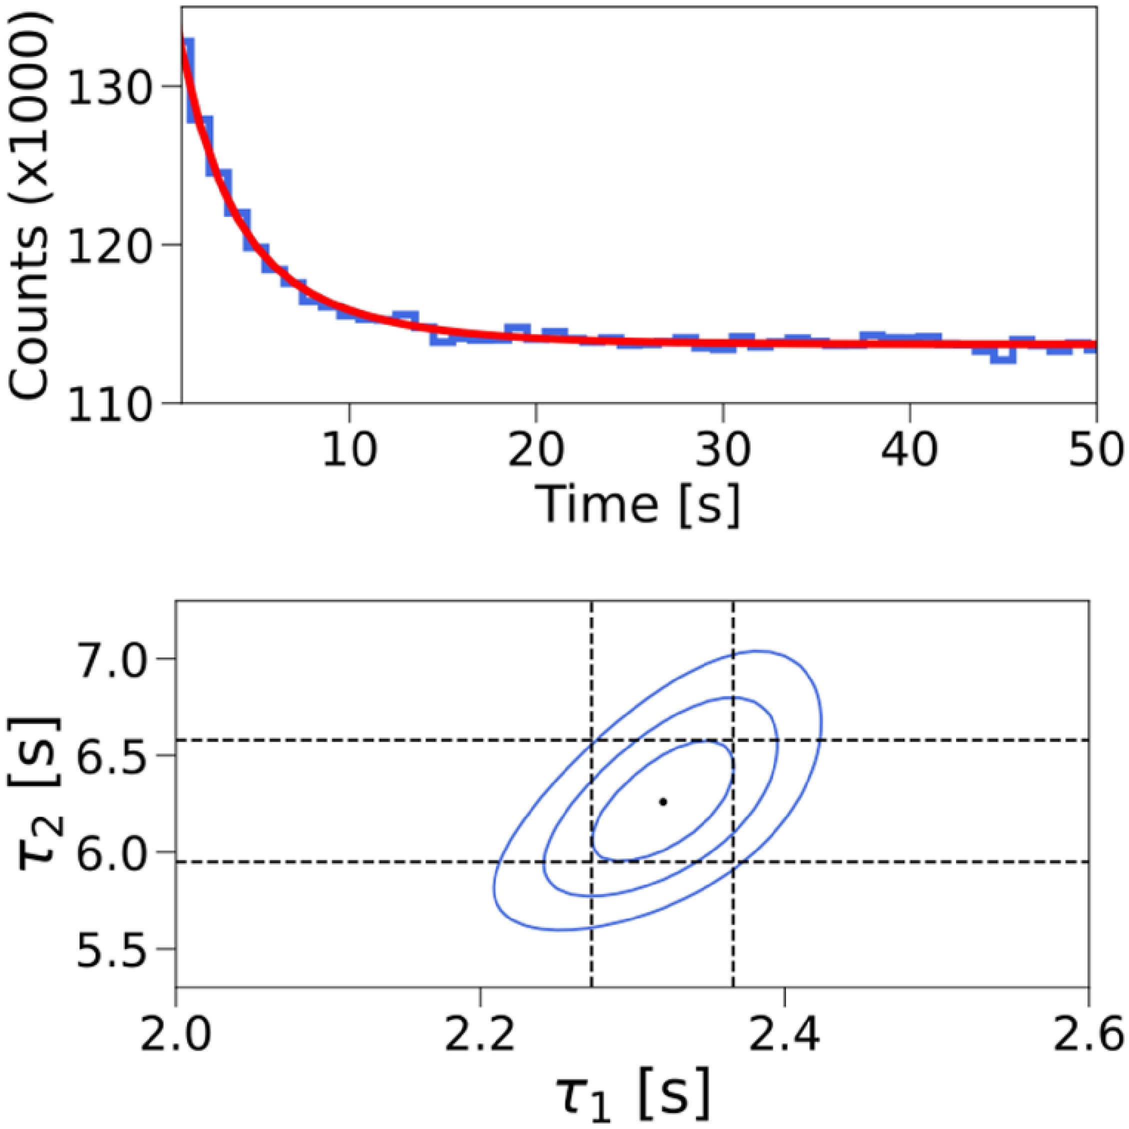

Figure 5.3: (Top) Time binned histogram of the short time fluorescence from ^{229}Th crystal 2.2 fitted with a two timescale exponential decay model. (Bottom) The likelihood contours of the two timescales found from the data in (Top). Each contour represents the [68%, 95%, 99.7%] confidence regions around χ^2_{min}.

The resulting fluorescence is shown as the black line in Fig. 5.2 and follows closely the observed fluorescence in crystal 2.2; the fluorescence exhibits a two-timescale behavior that is indistinguishable from that of crystal 2.2. Following these measurements, the experiment was repeated with a ^{232}Th:LiSrAlF$_6$ crystal (red line). As can be seen in Fig. 5.2, the ^{232}Th:LiSrAlF$_6$ and ^{229}Th:LiSrAlF$_6$ crystals yield a similar spectroscopic feature around 8 eV, but only the ^{229}Th:LiSrAlF$_6$ crystals yield a significant signal near the nuclear transition. Assuming this feature arises from nuclear excitation, we estimate from the known ^{229}Th density in these crystals [JSS15] and the detection efficiency of the system that only a portion of order 1 ppm of the doped ^{229}Th atoms participate in this short-timescale process. Other explanations, such a radiation-induced color centers cannot be ruled out.

This data suggests the observed short-timescale fluorescence could have a nuclear origin and it is therefore interesting to determine the wavelength of emission. However, the number of emitted photons was sufficiently small ($\sim 10^4$ photons into 4π) that it could not be analyzed with available monochromators. Therefore, in order to determine the spectrum of the fluorescence emitted from the highest activity ^{229}Th:LiSrAlF$_6$ crystal (crystal 2.2), the transmission of the fluorescence was measured through a set of band-pass filters. An imaging system was used to collimate the collected light before passing through the filters in order to eliminate angle effects [LHK11]. The filtered light was then focused onto a photomultiplier tube for detection. The VUV laser wavelength was chosen to be ≈ 147.4 nm to maximize the amount of fluorescence. The resulting fluorescence through various filter configurations is plotted in Fig. 5.4, where the horizontal axis is the central transmission wavelength of each configuration and the horizontal 'error bar' denotes the FWHM of the filter transmission (see Appendix 5.4.1). As can be seen, the fluorescence is clearly in the VUV region as it primarily transmits through only two filters nominally centered at 154 nm and 180 nm. To better determine the wavelength of the emitted light, a Gaussian deconvolution procedure [MFS20] was used to fit a model spectrum to the data (see SI). This analysis is most consistent with a narrow emission at ≈ 182 nm with a long-tailed background centered red of 280 nm. The long-tail in the UV is expected and has previously been understood as due to fluorescence induced by radioactive decay of the ^{229}Th in the crystal (see Fig. 2(b) of Ref. [JSS15]). As

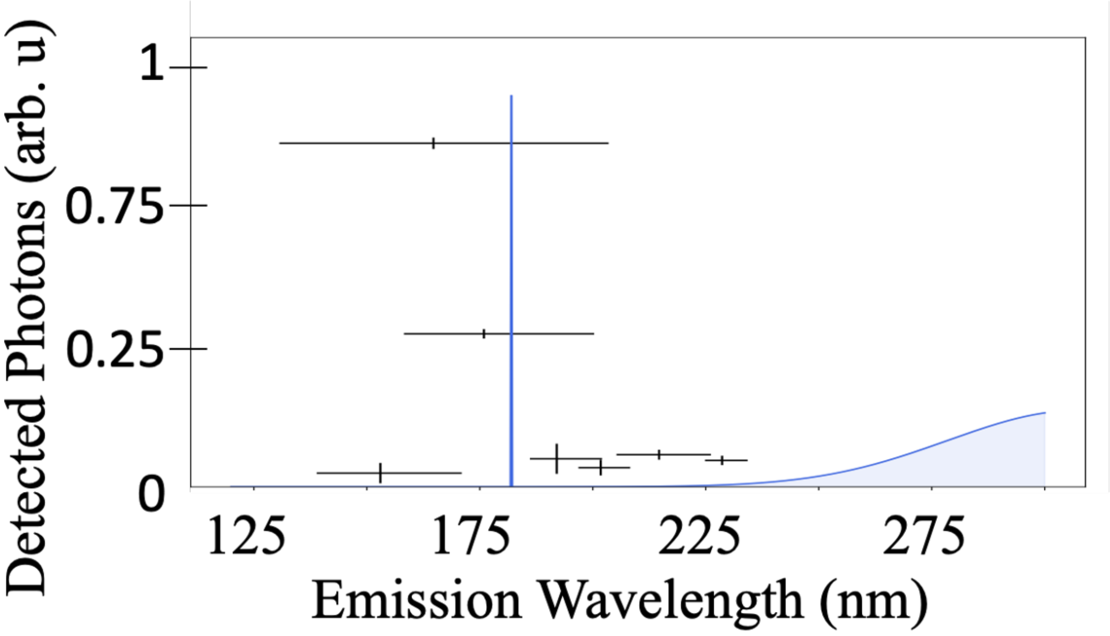

Figure 5.4: Fitted spectrum for the fluorescence from crystal 2.2 under excitation at 147.4 nm using a Gaussian deconvultion (see SI). The black points represent the recorded photon numbers; the vertical line represents the statistical uncertainty while the horizontal line represents the full width at half maximum (FWHM) of the transfer function of that particular filter configuration. Each point is centered at the peak of the transfer function. The shaded blue region represents the spectrum inferred from the Gaussian deconvolution. The feature near 182 nm has a width of $\lesssim 1$ nm.

will be discussed in the next section the narrow VUV emission could be attributed to the isomeric state quenched by excitation of a defect electron-hole pair in the crystal.

5.3 Theory of Nuclear Quenching by Defect States

To better understand these results, we performed density functional theory (DFT) calculations using the modified Becke-Johnson [TB09, BJ06] functional to estimate the electronic properties of the ^{229}Th-doped crystals (see SI). Fig. 5.5 shows the projected density of states (PDOS) for Th in Th:LiSrAlF$_6$, in a $3 \times 3 \times 2$ supercell containing one Th atom and 2 interstitial F atoms for charge balancing; this is the lowest-energy configuration in our calculations (see SI). We find a manifold of the

Th $5f$ states at 6.8 eV (182 nm) inside the insulator bandgap (d_i states). This is consistent with the observed fluorescence wavelength in Fig. 5.4, indicating the possibility of nuclear decay via excitation of a valence-band electron to a d_i state. Subsequent photon emission via the recombination of the d_i state with the hole, which has moved to the top of the valence band on ps timescales typical for non-radiative relaxation of hot holes [WEK92], would produce the observed, red-shifted 182 nm photons.

Such coupling of the electron and nuclear sub-systems is mediated by the hyperfine interaction (HFI). In the absence of the HFI, the eigenstates of the system are the product states $|\mathrm{el}\rangle|\mathrm{nuc}\rangle$; the HFI leads to admixtures of the excited nuclear state $|e\rangle$ with the product states involving the ground $|g\rangle$ nuclear state. In the first part of the process, the laser excites the occupied valence band state $\approx |\Omega\rangle|g\rangle$ to a particle-hole eigenstate $|\Psi_L\rangle = a^\dagger_{d_i} a_h |\Omega\rangle|g\rangle + |\delta\Phi_{\mathrm{el}}\rangle|e\rangle$ resonant with the laser frequency (here we use the creation and annihilation operators $a^\dagger$ and a; $|\Omega\rangle$ is the valence band quasi-vacuum state). The HFI-mediated admixture $|\delta\Phi_{\mathrm{el}}\rangle|e\rangle$ can be substantial due to a contribution of the $a^\dagger_{d_i} a_{h_r}|\Omega\rangle$ states resonant with the nuclear transition. It also contains the $|\Omega\rangle|e\rangle$ state. Such laser excitation exhibits a step-like character of Fig. 5.2. The components of this coherently excited state decay due to a combination of vibronic couplings to the crystal on a ps timescale and radiative particle-hole recombination on a ns timescale. While the majority of the decays result in a photon emission on the ns timescale, some population may evolve into the state $|\Omega\rangle|e\rangle$. The resulting $|\Omega\rangle|e\rangle$ state is embedded into the continuum of particle-hole states $a^\dagger_{d_i} a_h |\Omega\rangle|g\rangle$ and can thereby decay into this continuum with the rate $\Gamma = 2\pi\rho|\langle e|\langle\Omega|V_{\mathrm{HFI}}|a^\dagger_{d_i} a_{h_r}|\Omega\rangle|g\rangle|^2$. Here the hole h_r is such that the energy difference $\varepsilon_{d_i} - \varepsilon_{h_r}$ matches the nuclear excitation energy ω_{nuc} and ρ is the valence band density of states. The rate has to be summed over all allowed final states, including phonon degrees of freedom which bring in Frank-Condon factors in the rate evaluation. This internal conversion process quenches the nuclear transition. The vibronic couplings further "float" the hot resonant hole state h_r to the top of the valence band. Finally, a radiative particle-hole recombination emits a fluorescence photon that is red-shifted compared to ω_{nuc}. We find that this model semi-quantitatively explains the observed short-timescale fluorescence spectrum of

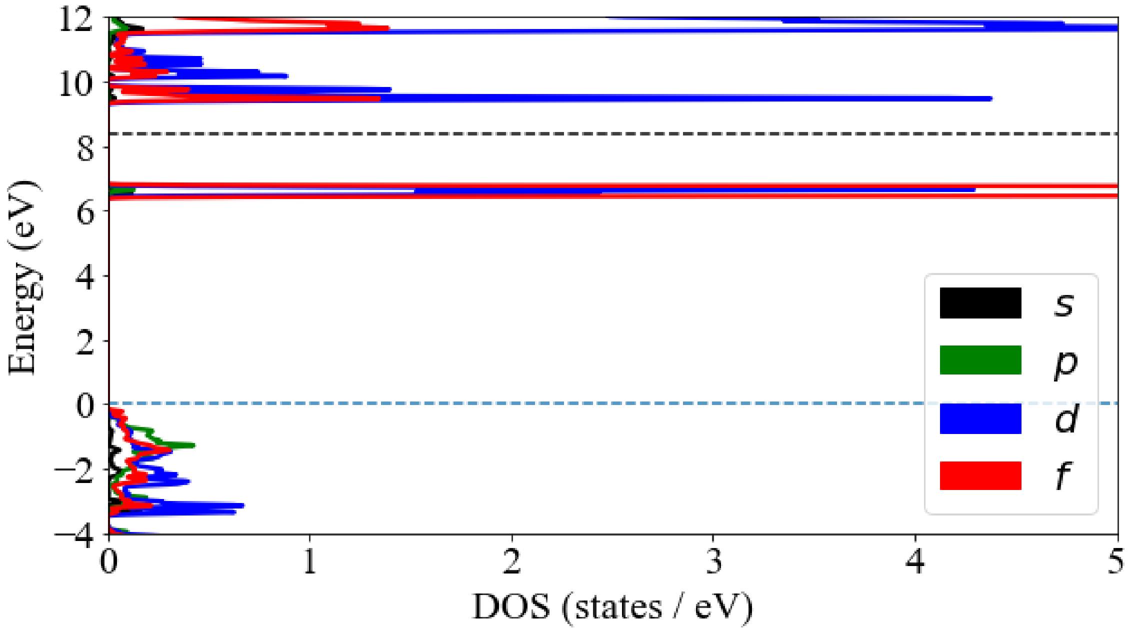

Figure 5.5: Thorium PDOS for ^{229}Th:LiSrAlF$_6$. The dashed black line represents the energy of the ^{229}Th nuclear excited state relative to the top of the valence band.

Fig. 5.2, 5.3, and Fig. 5.4. Full numerical simulations are still underway that seek to properly incorporate effects such as changes in the configuration geometry between the ground and excited states, as well as Franck-Condon factors.

It should be noted that this mechanism relies on the energy of the $5f$ defect states ε_{d_i} being lower than $\omega_{\rm nuc}$. Thus, the long-timescale signal observed from the nuclear isomeric transition is potentially due to ^{229}Th atoms that are substituted into the crystal in arrangements that have $\varepsilon_{d_i} > \omega_{\rm nuc}$ or vanishing ρ near the isomeric energy.

5.4 Appendix

5.4.1 Obtaining the Fluorescence Spectrum via Gauassian Deconvolution

The spectrum of the fluorescence was obtained by measuring its transmission through various filter configurations. The number of fluorescent photons (normalized to the amount of excitation photons), Y_i that make it through a particular filter configuration is given by

$$Y_i = \int \rho(\lambda)\ T_i(\lambda)\ \eta_i(\lambda)\ d\lambda\ , \tag{5.2}$$

where $\rho(\lambda)$ is the spectrum of the fluorescence, T_i are the filter transmission functions (provided by the manufacturer and measured when possible), and η captures the remaining quantum efficiencies of the detector and geometric acceptance of the lens system.

In order to determine ρ, a functional form for the spectrum was chosen, $\rho^*(\lambda, \vec{\xi})$, which is determined by a set of parameters $\vec{\xi}$. For a particular value of $\vec{\xi}$ we may then compute the following

$$Y_i^*(\vec{\xi}) = \int \rho^*(\lambda, \vec{\xi})\ T_i(\lambda)\ \eta_i(\lambda)\ d\lambda\ , \tag{5.3}$$

and from that compute

$$\chi^2 = \sum_i \frac{\left(Y_i - Y_i^*(\vec{\xi})\right)^2}{\sigma_i^2}\ , \tag{5.4}$$

where the sum runs over all filter configurations, and σ_i are the uncertainties in the normalized fluorescence counts. Minimizing χ^2 with respect to the parameters $\vec{\xi}$ will converge to the true spectrum if ρ is of a functional form described by $\rho^*(\lambda, \vec{\xi})$.

Given the observed filter data, it was clear that a major contribution to the signal was a sharply peaked component in the region of 150 nm to 200 nm, with a ‘redder’ background component. This was reflected in the improvement of χ^2 from ≈ 600 with one Gaussian function to ≈ 8 with the addition of another Gaussian. Adding more Gaussians did little to improve the value of χ^2. The choice of Gaussians to describe the spectrum is only reflective of their simplicity as peaked functions. It is likely the true spectrum is not exactly described by our parameterization, and the value quoted by our fit for the “sharp” part of the spectrum should only be seen as a rough estimate.

CHAPTER 6

Direct Laser Excitation of the ^{229}Th Isomeric Transition in $LiSrAlF_6$

6.1 Confirmation From ISOLDE

Our group managed to get in contact with the collaboration that had carried out the ISOLDE measurement discurssed in the previous chapter. We were able to arange for a piece of $LiSrAlF_6$to be placed in the ^{229}Ac beamline to see whether radiative decay from the isomer could be observed in $LiSrAlF_6$. Preliminary results from this run were shared with us almost immediately after the start of beamtime, which indicated that photons from the isomeric decay were observable and the lifetime was similar to that seen in MgF_2 and CaF_2. Data from this run has recently been released in pre-print [PCB24], and the relevant $LiSrAlF_6$ data can be seen in Fig. 6.1.

6.2 Fine Scan for the Isomeric Transition

With this confirmation we concluded that it was indeed possible to observe the isomeric transition in a $LiSrAlF_6$ crystal, but the fraction of emitters that could be expected was perhaps far lower than expected. The first step was to re-analyze the data in the ISOLDE region, where we fit an exponential with a 900s to look for any long-timescale fluorescence above background. The result of this analysis can be seen in Fig. 6.2, several points were flagged as being 2σ above the mean. While not definitive by any means, these were the most likely candidates to start the scan on.

We began a new search, optimized for detection of the isomeric transition, assuming only a few percent of the doped ^{229}Th atoms contributed. This search was performed over the spectral regions within the ISOLDE region that showed an excess

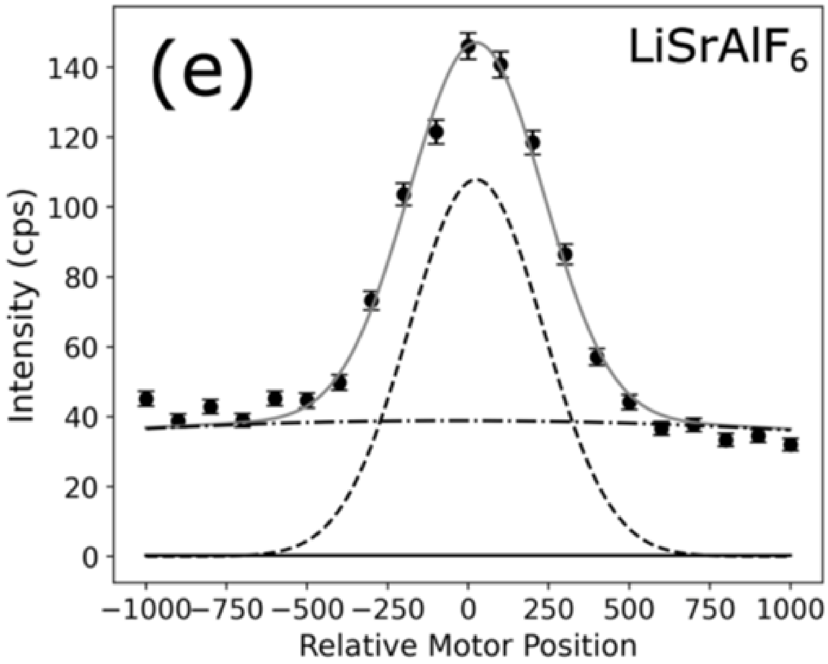

Figure 6.1: Spectrometer data from [PCB24] demonstrating radiative decay from the isomer in LiSrAlF$_6$.

of photon counts relative to their neighbors (see Fig. 6.2).

Initially, each datapoint consisted of a single sweep over 24 grating steps of the VIS PDL. At each grating position the laser would fire 2500 shots, the PMT shutters would open, and we would collect photons for 1000 s. This constituted 2000 s of illumination time for each datapoint over a frequency sweep of $\sim$ 8-10 GHz. Each datapoint would at a grating position overlapped with the half-way point of the previous datapoint, to minimize the chance of missing the transition due to improper overlap. Each datapoint was monitored for issues in laser power, grating calibration, the Xe reference scan, etc. An example of the diangostics that were carried out for each datapoint can be seen in Fig. 6.3. In addition, each datapoint was fit with an emitter fraction p_0 with a fixed lifetime of τ = 966 s (see Appendix 6.3.1). If any datapoint was noticed to have an emitter fraction of 0.02 or higher it was flagged for re-scanning.

Near a region of excess photons counts in our 2019 data (see dashed blue line in Fig. 6.2), the emitter fraction showed an increase to $p_0 \sim 0.05$, an observed lifetime on the order of $\sim$ 600 s. We then changed the scan parameters to illuminate the crystal

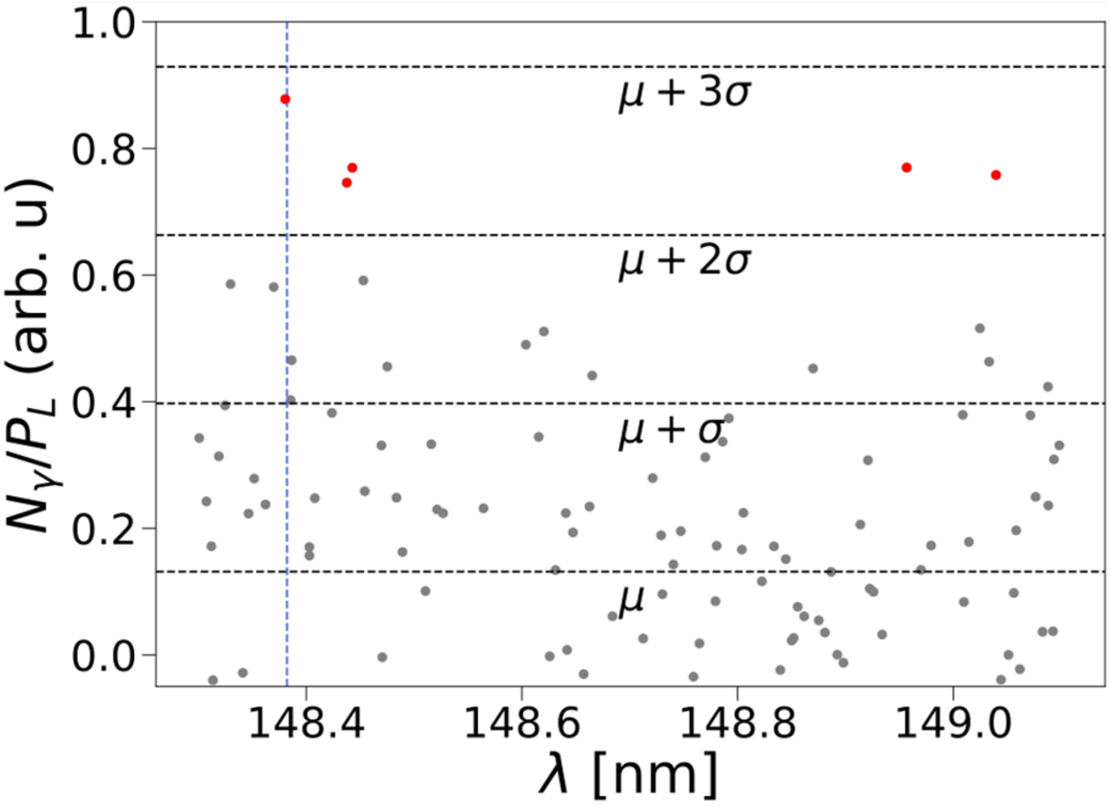

Figure 6.2: Long timescale fluorescence from ^{229}Th:LiSrAlF$_6$ crystal 2.2 normalized by laser power after ≈ 2000 s illumination shown over the 68% confidence interval of the ISOLDE experiment [KMA23]. Dotted black lines denote the mean (μ) and mean plus multiples of the standard deviation (σ). Though the expected signal is roughly two-orders of magnitude larger than what is observed, points above $\mu + 2\sigma$ were denoted as potential detection events and slated for further study. The vertical blue dashed line denotes the location of the narrow feature observed in Fig. 6.4(a).

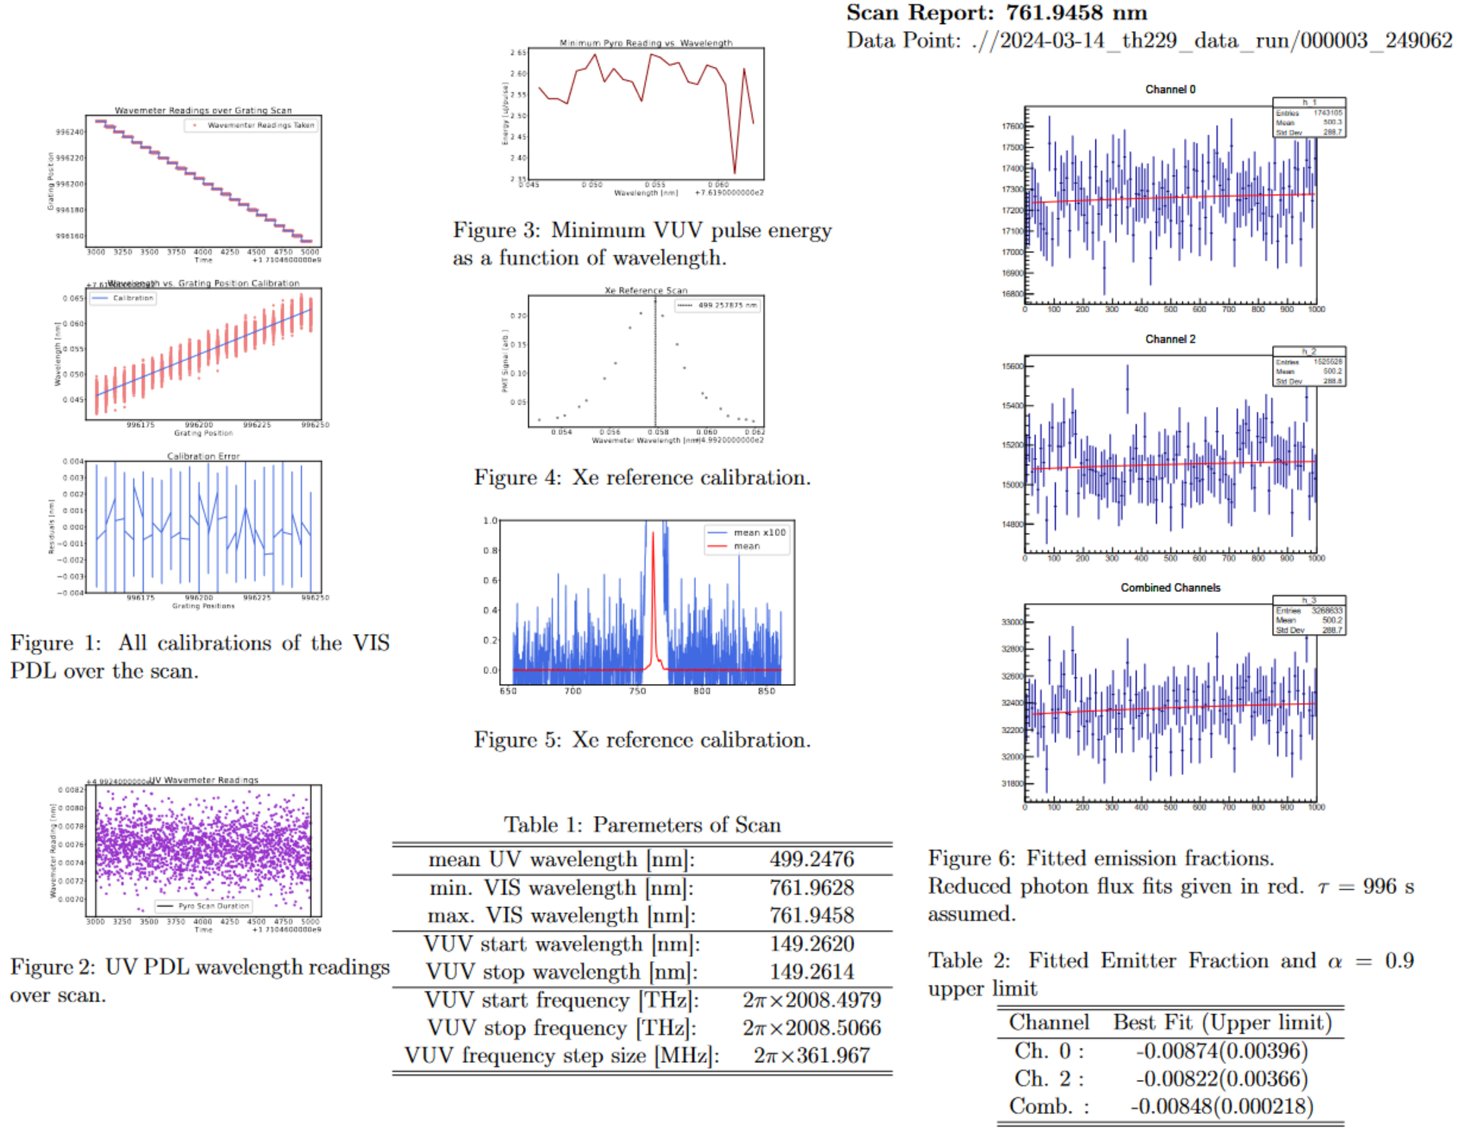

Scan Report: 761.9458 nm
Data Point: .//2024-03-14_th229_data_run/000003_249062

Figure 1: All calibrations of the VIS PDL over the scan.

Figure 2: UV PDL wavelength readings over scan.

Figure 3: Minimum VUV pulse energy as a function of wavelength.

Figure 4: Xe reference calibration.

Figure 5: Xe reference calibration.

Table 1: Paremeters of Scan

mean UV wavelength [nm]:	499.2476
min. VIS wavelength [nm]:	761.9628
max. VIS wavelength [nm]:	761.9458
VUV start wavelength [nm]:	149.2620
VUV stop wavelength [nm]:	149.2614
VUV start frequency [THz]:	$2\pi\times2008.4979$
VUV stop frequency [THz]:	$2\pi\times2008.5066$
VUV frequency step size [MHz]:	$2\pi\times361.967$

Figure 6: Fitted emission fractions. Reduced photon flux fits given in red. $\tau = 996$ s assumed.

Table 2: Fitted Emitter Fraction and $\alpha = 0.9$ upper limit

Channel	Best Fit (Upper limit)
Ch. 0 :	-0.00874(0.00396)
Ch. 2 :	-0.00822(0.00366)
Comb. :	-0.00848(0.000218)

Figure 6.3: Example of a scan report produced for every datapoint during the fine scan. Fitted emitter fractions can be seen in the bottom right table.

for 1200 s at a single grating position, followed by 1800 s of long-lifetime fluorescence collection. The resulting power-normalized, long-timescale fluorescence (determined by the scaled emitter fraction) are shown in Fig 6.4(a). The normalized fluorescence in this region reveals a narrow spectral feature at $148.38219(4)_{\text{stat}}(20)_{\text{sys}}$ nm ($2020407.3(5)_{\text{stat}}(30)_{\text{sys}}$ GHz) in both crystals 2.2 and 3.1; the primary source of systematic uncertainty is wavemeter calibration (see Appendix 6.3.2). The full width at half maximum (FWHM) of this feature is roughly 15 GHz and in agreement with a measurement of the laser linewidth. The central wavelength of this feature is identical, within error, in both crystals 2.2 and 3.1 and to that observed in Ref. [TOZ24]. It is, thus, attributed to the excitation of the ^{229}Th isomeric state. From this data, we estimate that approximately 1% (50%) of the doped ^{229}Th nuclei in crystal 2.2 (3.1) are able to be excited on the nuclear isomeric transition. We can translate these emitter percentages into an effective "participant density" n_{par}, which is listed in Table 6.1.

Crystal	n_{Th}	n_{par} [atoms/cm^{-3}]	Background [s^{-1}]
2.2	1×10^{16}	1×10^{14}	1600
3.1	4×10^{15}	8×10^{14}	190

Table 6.1: Relevant crystal parameters for computing signal-to-noise ratios. The ^{229}Th density n_{Th} was estimated based on gamma ray spectrometer measurements.

The measured isomeric decay lifetime is $\tau = 568(13)_{\text{stat}}(20)_{\text{sys}}$ s as shown in Fig.6.4(b), where the systematic error is estimated based on long-term drifts in the PMTs background count rates. The LiSrAlF$_6$ index of refraction is not well-characterized in the VUV but can be estimated as 1.485 [Jee18], leading to an estimate of the isomeric lifetime in vacuum of $\tau_{is} \approx 1860(43)_{\text{stat}}(66)_{\text{sys}}$ s [?] via the Purcell effect. This estimate is larger than that observed in ^{229}Th:CaF$_2$ [TOZ24] (2510(72) s), and is likely attributed to a poor estimate of the index of refraction of these crystals in the VUV. Beyond the Purcell effect, the nuclear radiative lifetime may still be shortened by effects of the crystal electric field, which mix Th electronic states of opposite parity. The hyperfine interaction couples the nuclear degree of freedom to these mixed-parity states, allowing for a competing $E1$ decay channel for the isomeric state. Such effects may complicate efforts to determine the radiative lifetime of the bare ^{229}Th isomeric state from that observed in crystalline hosts. Nonetheless, both values are within the range predicted in Ref. [TSJ15].

6.3 Appendix

6.3.1 Calculating the Expected Isomeric Transition Signal

We slightly modify Eq. 1.43, so that the number of excited ^{229}Th nuclei is given by

$$N_0 = p_0 \frac{2}{3} \frac{\lambda^2}{2\pi} \frac{\varphi_p n_{Th} l}{\mathit{\Gamma}_L} \frac{1}{\left(1 + 4\left(\frac{\omega - \omega_0}{\mathit{\Gamma}_L}\right)^2\right)} (1 - \exp(-T_e/\tau)) \tag{6.1}$$

where p_0 is the fraction of ^{229}Th nuclei that relax radiatively, λ (ω) is the excitation wavelength (frequency), φ_p is the photon flux per second, n_{Th} is the density of Th nuclei, l is the length of the crystal, $\mathit{\Gamma}_L$ is the laser bandwidth, ω_0 is true transition frequency, T_e is the illumination time, and τ is the lifetime of the transition in the fluoride crystal environment. The lifetime of the isomer in the crystal environment will necessarily be shorter than the vacuum lifetime due to the increased density of electromagnetic states, which leads to a shortening of the lifetime by n^3 [**?**], where n is the crystal index of refraction.

The photons emitted by the excited nuclei are collected on two PMTs positioned with a few mm of the crystal. Calibration of the laser energy is done with a custom detector [Jee18]. The solid angle of collection of the detection system is calculated via ray tracing software and the quantum efficiencies of the PMTs is done with a NIST calibrated photodiode [JSS15, Jee18]. The total detection efficiency is plotted as C2a in Fig. 2(a) of Ref. [JSS15]. It is ≈1% at the isomeric transition wavelength.

6.3.2 Isomeric Transition Frequency Systematic Uncertainty

The chief source of systematic uncertainty in the determination of the isomeric transition frequency is the calibration of the VUV laser frequency. The VUV frequency is given by the difference mixing relation $\omega = 2\omega_u - \omega_v$ and therefore depends on the measurement of ω_u and ω_v, the UV and visible laser frequency, re-

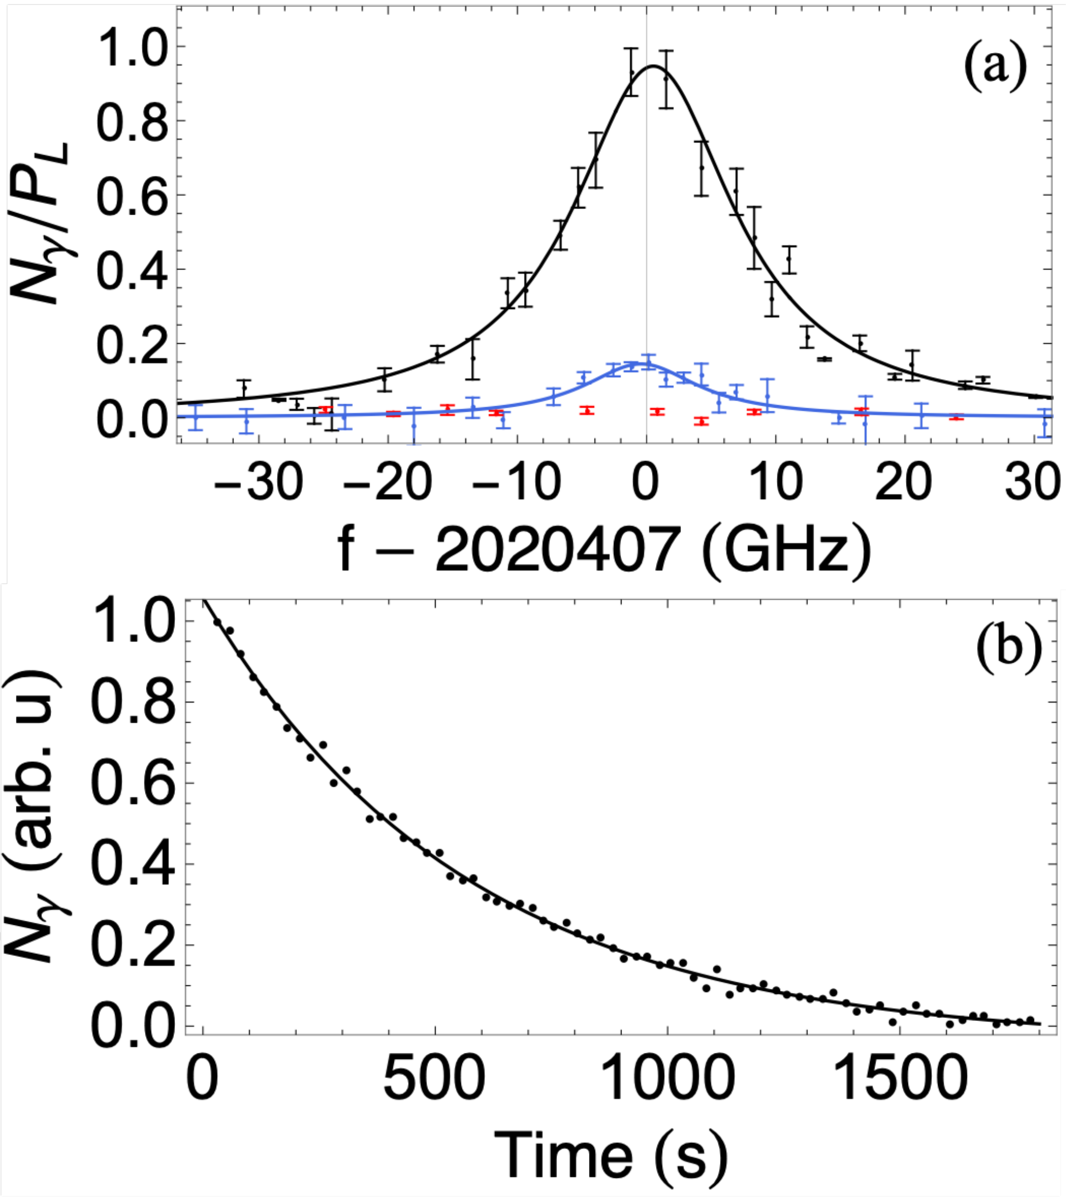

Figure 6.4: (a) Spectrum of the ^{229}Th isomeric transition. The total power-normalized photon count in 1800 s following illumination is plotted versus excitation wavelength for crystal 2.2 (blue), 3.1 (black) and a ^{232}Th:LiSrAlF$_6$ crystal (red). The solid lines represent fits of a Lorentzian lineshape, which result in statistically identical center frequencies. The linewidth of the fitted Lorentzians are also statistically identical and consistent with the linewidth of the VUV laser system, indicating a narrow line. (b) Observed photon count rate versus time from crystal 3.1 alongside an decaying exponential fit which determines the isomeric state lifetime within the crystal to be $\tau = 568(13)_{\text{stat}}(20)_{\text{sys}}$ s. A statistically identical lifetime is observed in crystal 2.2.

spectively. These frequencies are measured using a Bristol 871B wavemeter and augmented with a calibration via a Xe reference line. Specifically, the wavelength of the doubled pulsed dye laser (UV PDL) is calibrated by scanning the two-photon $5p^6\ {}^1S_0 \rightarrow 5p^5\left({}^2P^\circ_{3/2}\right)6p\ {}^2\left[1/2\right]_0$ transition in Xenon and comparing to the value on the NIST Database. The peak was always within 0.5 pm of the expected value (499.2576 nm for the fundamental). Thus, the UV PDL wavelength was accurate to within $\approx$ 1.2 GHz. The visible dye laser (VIS PDL) wavelength was measured by the same wavemeter, whose stated calibration uncertainty is 1 pm, or $\approx$ 500 MHz at 785.7 nm. Assuming the errors are anti-correlated provides an estimate of the systematic uncertainty as $\sigma_{VUV} = 2\sigma_u + \sigma_v = 3$ GHz.

CHAPTER 7

Conclusion

Nearly a half century has passed since Kroger and Reich first discovered the ^{229}Th isomeric transition, and its direct excitation and resolution by a laser represents the dawn of a new era of applications and experiments. Most obvious is the potential to create a portable precision clock for geodesy and tests of the variation of fundamental constants. One can also envision using the clock transition as a probe of the chemical environment, as in the tradition of Mössbauer spectroscopy. It is this last point that most lies within the spirit of this thesis, and one which I would like to offer some thoughts.

In fluoride crystal systems, the spectroscopic signature would be the VUV fluorescence, and its associated splittings and lineshapes. Already work is emerging that associates these observables with the electric field gradients and isomer shifts in a crystal [ZOH24, HOD24]. One could also imagine re-visiting the low band gap systems (e.g. ThO_2) and attempting to observe the lifetime and emission of IC electrons emitted by the isomer to learn about the local electron density of states. With sufficient integration times and laser power, one could even attempt to excite phonon sidebands of the isomer transition in solids to learn about the local phonon behavior, as in nuclear resonance vibrational spectroscopy (NRVS). All of these experiments have the potential to push forward theories that lie at the intersection of crystal engineering, nuclear theory, relativistic atomic calculations, and quantum chemistry.

Beyond learning about the solid-state environment, these experiments will be necessary to learn more about the systematics that may limit the ultimate precision and accuracy of a solid-state ^{229}Th clock. There is also the possiblity of observing collective effects in a densly doped environment, such as nuclear forward scattering [NLP18] and collective nuclear spontaneous emission [Tka11]. Suffice it to say, much more work needs to be done, but we are closer than ever to the promise of what may be the ultimate time-keeping device, the nuclear clock.

[AAM20] Sara Algeri, Jelle Aalbers, Knut Dundas Morå, and Jan Conrad. “Searching for new phenomena with profile likelihood ratio tests.” *Nature Reviews Physics*, **2**(5):245–252, May 2020.

[ADD12] Hatim Azzouz, Sander N. Dorenbos, Daniel De Vries, Esteban Bermúdez Ureña, and Valery Zwiller. “Efficient single particle detection with a superconducting nanowire.” *AIP Advances*, **2**(3):032124, 07 2012.

[AS96a] V. A. Alekseev and D. W. Setser. “A pulsed source for Xe(6s$[3/2]_1$) and Xe(6s’$[1/2]_1$) resonance state atoms using two-photon driven amplified spontaneous emission from the Xe(6p) and Xe(6p’) states.” *J. Chem. Phys.*, **105**(11):4613–4625, 1996.

[AS96b] V. A. Alekseev and D. W. Setser. “A pulsed source for Xe(6s[3/2]1) and Xe(6s’[1/2]1) resonance state atoms using two-photon driven amplified spontaneous emission from the Xe(6p) and Xe(6p’) states.” *The Journal of Chemical Physics*, **105**(11):4613–4625, 09 1996.

[Bar90] Roger Barlow. “Extended maximum likelihood.” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **297**(3):496–506, 1990.

[Bas07] M.S. Basunia. “Nuclear Data Sheets for A = 213.” *Nuclear Data Sheets*, **108**(3):633–680, 2007.

[BBB07] B. R. Beck, J. A. Becker, P. Beiersdorfer, G. V. Brown, K. J. Moody, J. B. Wilhelmy, F. S. Porter, C. A. Kilbourne, and R. L. Kelley. “Energy Splitting of the Ground-State Doublet in the Nucleus ^{229}Th.” *Phys. Rev. Lett.*, **98**:142501, Apr 2007.

[BBR17] Arne Baumann, Sophia Bazzi, Dimitrios Rompotis, Oliver Schepp, Armin Azima, Marek Wieland, Daria Popova-Gorelova, Oriol Vendrell, Robin Santra, and Markus Drescher. “Weak-field few-femtosecond VUV photodissociation dynamics of water isotopologues.” *Phys. Rev. A*, **96**:013428, Jul 2017.

[BC84] Steve Baker and Robert D. Cousins. “Clarification of the use of CHI-square and likelihood functions in fits to histograms.” *Nuclear Instruments and Methods in Physics Research*, **221**(2):437–442, 1984.

[BCH19] S. M. Brewer, J.-S. Chen, A. M. Hankin, E. R. Clements, C. W. Chou, D. J. Wineland, D. B. Hume, and D. R. Leibrandt. “^{27}Al$^+$ Quantum-Logic Clock with a Systematic Uncertainty below 10^{-18}.” *Phys. Rev. Lett.*, **123**:033201, Jul 2019.

[Bee22] K. Beeks. *The nuclear excitation of Thorium-229 in the CaF2 environment.* PhD thesis, Technischen Universität Wien, 2022.

[BGQ08] D.G. Burke, P.E. Garrett, Tao Qu, and R.A. Naumann. “Nuclear structure of 229,231Th studied with the 230,232Th(d,t) reactions.” *Nuclear Physics A*, **809**(3):129–170, 2008.

[BJ06] Axel D. Becke and Erin R. Johnson. "A simple effective potential for exchange." *J. Chem. Phys.*, **124**(22):221101, 06 2006.

[BKA22] Tobias Bothwell, Colin J. Kennedy, Alexander Aeppli, Dhruv Kedar, John M. Robinson, Eric Oelker, Alexander Staron, and Jun Ye. "Resolving the gravitational redshift across a millimetre-scale atomic sample." *Nature*, **602**(7897):420–424, Feb 2022.

[BM98] Aage Bohr and Ben R. Mottelson. *Nuclear Structure.* New World Scientific, 1998.

[BMJ88] C E Bemis, F K McGowan, J L C Ford Jr, W T Milner, R L Robinson, P H Stelson, G A Leander, and C W Reich. "Coulomb excitation of states in 229Th." *Physica Scripta*, **38**(5):657, nov 1988.

[BMP18a] Pavlo V. Bilous, Nikolay Minkov, and Adriana Pálffy. "Electric quadrupole channel of the 7.8 eV ^{229}Th transition." *Phys. Rev. C*, **97**:044320, Apr 2018.

[BMP18b] Pavlo V. Bilous, Nikolay Minkov, and Adriana Pálffy. "Electric quadrupole channel of the 7.8 eV ^{229}Th transition." *Phys. Rev. C*, **97**:044320, Apr 2018.

[Boy07] Martin M. Boyd. *High Precision Spectroscopy of Strontium in an Optical Lattice: Towards a New Standard for Frequency and Time.* PhD thesis, University of Colorado, Boulder, 2007.

[BW77] R. Brockmann and W. Weise. "Spin-orbit coupling in a relativistic Hartree model for finite nuclei." *Phys. Rev. C*, **16**:1282–1284, Sep 1977.

[BWB09] B R Beck, C Wu, P Beiersdorfer, G V Brown, J A Becker, K J Moody, J B Wilhelmy, F S Porter, C A Kilbourne, and R L Kelley. "Improved Value for the Energy Splitting of the Ground-State Doublet in the Nucleus 229mTh." 7 2009.

[Cas01] Richard F. Casten. *Nuclear Structure from a Simple Perspective.* Oxford University Press, 2nd edition, 2001.

[CHH64] Canfield, L.R., Hass, G., and Hunter, W.R. "The optical properties of evaporated gold in the vacuum ultraviolet from 300 Å to 2 000 Å." *J. Phys. France*, **25**(1-2):124–129, 1964.

[CM20] Guy Chanfray and Jérôme Margueron. "Contribution of the ρ meson and quark substructure to the nuclear spin-orbit potential." *Phys. Rev. C*, **102**:024331, Aug 2020.

[CRK12] C. J. Campbell, A. G. Radnaev, A. Kuzmich, V. A. Dzuba, V. V. Flambaum, and A. Derevianko. "Single-Ion Nuclear Clock for Metrology at the 19th Decimal Place." *Phys. Rev. Lett.*, **108**:120802, Mar 2012.

[ERI15] A Engel, J J Renema, K Il'in, and A Semenov. "Detection mechanism of superconducting nanowire single-photon detectors." *Superconductor Science and Technology*, **28**(11):114003, sep 2015.

[ES13] Andreas Engel and Andreas Schilling. “Numerical analysis of detection-mechanism models of superconducting nanowire single-photon detector.” *Journal of Applied Physics*, **114**(21):214501, 12 2013.

[ESJ24] R. Elwell, Christian Schneider, Justin Jeet, J. E. S. Terhune, H. W. T. Morgan, A. N. Alexandrova, H. B. Tran Tan, Andrei Derevianko, and Eric R. Hudson. “Laser Excitation of the ^{229}Th Nuclear Isomeric Transition in a Solid-State Host.” *Phys. Rev. Lett.*, **133**:013201, Jul 2024.

[FAN22] Silvio Fuchs, Johann J. Abel, Jan Nathanael, Julius Reinhard, Felix Wiesner, Martin Wünsche, Slawomir Skruszewicz, Christian Rödel, Detlef Born, Heidemarie Schmidt, and Gerhard G. Paulus. “Photon counting of extreme ultraviolet high harmonics using a superconducting nanowire single-photon detector.” *Applied Physics B*, **128**(2):26, Jan 2022.

[FW09] V. V. Flambaum and R. B. Wiringa. “Enhanced effect of quark mass variation in ^{229}Th and limits from Oklo data.” *Phys. Rev. C*, **79**:034302, Mar 2009.

[GEM19] T. Gouder, R. Eloirdi, R. L. Martin, M. Osipenko, M. Giovannini, and R. Caciuffo. “Measurements of the band gap of ThF_4 by electron spectroscopy techniques.” *Phys. Rev. Res.*, **1**:033005, Oct 2019.

[HCN19] Yonit Hochberg, Ilya Charaev, Sae-Woo Nam, Varun Verma, Marco Colangelo, and Karl K. Berggren. “Detecting Sub-GeV Dark Matter with Superconducting Nanowires.” *Phys. Rev. Lett.*, **123**:151802, Oct 2019.

[HCS09] S. J. Hanna, P. Campuzano-Jost, E. A. Simpson, D. B. Robb, I. Burak, M. W. Blades, J. W. Hepburn, and A. K. Bertram. “A new broadly tunable (7.4–10.2eV) laser based VUV light source and its first application to aerosol mass spectrometry.” *Int. J. Mass Spectrom.*, **279**(2):134–146, 2009.

[Hil82] Robert C. Hilborn. “Einstein coefficients, cross sections, f values, dipole moments, and all that.” *American Journal of Physics*, **50**(11):982–986, 11 1982.

[HKK20] H.-W. Hammer, Sebastian König, and U. van Kolck. “Nuclear effective field theory: Status and perspectives.” *Rev. Mod. Phys.*, **92**:025004, Jun 2020.

[HOB24] Takahiro Hiraki, Koichi Okai, Michael Bartokos, Kjeld Beeks, Hiroyuki Fujimoto, Yuta Fukunaga, Hiromitsu Haba, Yoshitaka Kasamatsu, Shinji Kitao, Adrian Leitner, Takahiko Masuda, Ming Guan, Nobumoto Nagasawa, Ryoichiro Ogake, Martin Pimon, Martin Pressler, Noboru Sasao, Fabian Schaden, Thorsten Schumm, Makoto Seto, Yudai Shigekawa, Kotaro Shimizu, Tomas Sikorsky, Kenji Tamasaku, Sayuri Takatori, Tsukasa Watanabe, Atsushi Yamaguchi, Yoshitaka Yoda, Akihiro Yoshimi, and Koji Yoshimura. “Controlling 229Th isomeric state population in a VUV transparent crystal.” *Nature Communications*, **15**(1):5536, Jul 2024.

[HOD24] Jacob S. Higgins, Tian Ooi, Jack F. Doyle, Chuankun Zhang, Jun Ye, Kjeld Beeks, Tomas Sikorsky, and Thorsten Schumm. “Temperature sensitivity of a Thorium-229 solid-state nuclear clock.”, 2024.

[HR94] R. G. Helmer and C. W. Reich. “An excited state of ^{229}Th at 3.5 eV.” *Phys. Rev. C*, **49**:1845–1858, Apr 1994.

[HVL08] Eric R. Hudson, A. C. Vutha, S. K. Lamoreaux, and D. DeMille. “Investigation of the optical transition in the 229Th nucleus: Solid-state optical frequency standard and fundamental constant variation.” *Int. Conf. Atm. Phys., poster M028*, 2008.

[Jee18] J. Jeet. *Search for the low lying transition in the ^{229}Th Nucleus.* PhD thesis, University of California, Los Angeles, 2018.

[JSS15] Justin Jeet, Christian Schneider, Scott T. Sullivan, Wade G. Rellergert, Saed Mirzadeh, A. Cassanho, H. P. Jenssen, Eugene V. Tkalya, and Eric R. Hudson. “Results of a Direct Search Using Synchrotron Radiation for the Low-Energy ^{229}Th Nuclear Isomeric Transition.” *Phys. Rev. Lett.*, **114**:253001, Jun 2015.

[KLM17] A. G. Kozorezov, C. Lambert, F. Marsili, M. J. Stevens, V. B. Verma, J. P. Allmaras, M. D. Shaw, R. P. Mirin, and Sae Woo Nam. “Fano fluctuations in superconducting-nanowire single-photon detectors.” *Phys. Rev. B*, **96**:054507, Aug 2017.

[KLS92] T. T. Kajava, H. M. Lauranto, and R. R. E. Salomaa. “Mode structure fluctuations in a pulsed dye laser.” *Appl. Opt.*, **31**(33):6987–6992, November 1992.

[KMA23] Sandro Kraemer, Janni Moens, Michail Athanasakis-Kaklamanakis, Silvia Bara, Kjeld Beeks, Premaditya Chhetri, Katerina Chrysalidis, Arno Claessens, Thomas E. Cocolios, João G. M. Correia, et al. “Observation of the radiative decay of the 229Th nuclear clock isomer.” *Nature*, **617**(7962):706–710, May 2023.

[KR71] J. Kownacki and K.G. Rensfelt. “At 880 ns isomeric state in 144Sm.” *Physics Letters B*, **35**(2):153–155, 1971.

[KR76] L.A. Kroger and C.W. Reich. “Features of the low-energy level scheme of 229Th as observed in the α-decay of 233U.” *Nuclear Physics A*, **259**(1):29–60, 1976.

[Kra88] Kenneth S. Krane. *Introduction to Nuclear Physics.* John Wiley & Sons, 1988.

[KT21] F.F. Karpeshin and M.B. Trzhaskovskaya. “A proposed solution for the lifetime puzzle of the 229mTh+ isomer.” *Nuclear Physics A*, **1010**:122173, 2021.

[LHK11] M.G. Lödahl, V. M. J. Henriques, and D. Kiselman. “A tilted interference filter in a converging beam.” *A&A*, **533**:A82, Sep. 2011.

[LVH70] K.E.G. LÖBNER, M. VETTER, and V. HÖNIG. “NUCLEAR INTRINSIC QUADRUPOLE MOMENTS AND DEFORMATION PARAMETERS.” *Atomic Data and Nuclear Data Tables*, **7**(5):495–564, 1970.

[Mas19] Takahiko Masuda et al. “X-ray pumping of the^{229}Th nuclear clock isomer.” *Nature*, **573**(7773):238–242, 2019.

[MFS20] George H. Major, Neal Fairley, Peter M. A. Sherwood, Matthew R. Linford, Jeff Terry, Vincent Fernandez, and Kateryna Artyushkova. “Practical guide for curve fitting in x-ray photoelectron spectroscopy.” *J. Vac. Sci. Technol. A*, **38**(6):061203, 10 2020.

[Nat24] National Nuclear Data Center. “NuDat 3.” `https://www.nndc.bnl.gov/nudat/`, 2024. Accessed: 2024-10-05.

[Nav15] William Navidi. *Statistics for Engineers and Scientists 4th Edition*. McGraw-Hill, 2015.

[Nic14] N. Nica. “Nuclear Data Sheets for A = 148.” *Nuclear Data Sheets*, **117**:1–229, 2014.

[Nil55] S G Nilsson. “BINDING STATES OF INDIVIDUAL NUCLEONS IN STRONGLY DEFORMED NUCLEI.” *Kgl. Danske Videnskab. Selskab., Mat.-fys Medd.*, **Vol. 29, No. 16**, 1 1955.

[NJ10] Heung-Ryoul Noh and Wonho Jhe. “Analytic solutions of the optical Bloch equations.” *Optics Communications*, **283**(11):2353–2355, 2010.

[NLP18] Brenden S. Nickerson, Wen-Te Liao, and Adriana Pálffy. “Collective effects in ^{229}Th-doped crystals.” *Phys. Rev. A*, **98**:062520, Dec 2018.

[NM65] John A. Nelder and Roger Mead. “A Simplex Method for Function Minimization.” *Comput. J.*, **7**:308–313, 1965.

[NPB20] Brenden S. Nickerson, Martin Pimon, Pavlo V. Bilous, Johannes Gugler, Kjeld Beeks, Tomas Sikorsky, Peter Mohn, Thorsten Schumm, and Adriana Pálffy. “Nuclear Excitation of the ^{229}Th Isomer via Defect States in Doped Crystals.” *Phys. Rev. Lett.*, **125**:032501, Jul 2020.

[Paw13] Yudi Pawitan. *In All Likelihood*. Oxford University Press, 2013.

[PCB24] S. V. Pineda, P. Chhetri, S. Bara, Y. Elskens, S. Casci, A. N. Alexandrova, M. Au, M. Athanasakis-Kaklamanakis, M. Bartokos, K. Beeks, C. Bernerd, A. Claessens, K. Chrysalidis, T. E. Cocolios, J. G. Correia, H. De Witte, R. Elwell, R. Ferrer, R. Heinke, E. R. Hudson, F. Ivandikov, Yu. Kudryavtsev, U. Köster, S. Kraemer, M. Laatiaoui, R. Lica, C. Merckling, I. Morawetz, H. W. T. Morgan, D. Moritz, L. M. C. Pereira, S. Raeder, S. Rothe, F. Schaden, K. Scharl, T. Schumm, S. Stegemann, J. Terhune, P. G. Thirolf, S. M. Tunhuma, P. Van Den Bergh, P. Van Duppen, A. Vantomme, U. Wahl, and Z. Yue. “Radiative Decay of the ^{229m}Th Nuclear Clock Isomer in Different Host Materials.”, 2024.

[POG13] N. Poli, C. W. Oates, P. Gill, and G. M. Tino. “Optical atomic clocks.” *Rivista Del Nuovo Cimento*, **36**(12), 2013.

[PT03] E. Peik and Chr. Tamm. “Nuclear laser spectroscopy of the 3.5 eV transition in Th-229.” *Europhysics Letters*, **61**(2):181, jan 2003.

[RDG89a] Michael B. Rankin, Jon P. Davis, Christopher Giranda, and Lloyd C. Bobb. “Two-photon-excited bidirectional emission in xenon.” *Opt. Commun.*, **70**(4):345–349, March 1989.

[RDG89b] Michael B. Rankin, Jon P. Davis, Christopher Giranda, and Lloyd C. Bobb. “Two-photon-excited bidirectional emission in xenon.” *Optics Communications*, **70**(4):345–349, 1989.

[RHD10] Wade G Rellergert, Eric R Hudson, D DeMille, R R Greco, M P Hehlen, and J R Torgerson. “Constraining the Evolution of the Fundamental Constants with a Solid-State Optical Frequency Reference Based on the sup 229Th Nucleus.” *Physical Review Letters*, **104**(20), 5 2010.

[Rin80] Peter Ring, Peter Schuck. *The Nuclear Many-Body Problem.* Springer Berlin, Heidelberg, 1st edition, 1980.

[RL71] Elward T. Rodine and Peter L. Land. “Electronic Defect Structure of Single-Crystal ThO_2 by Thermoluminescence.” *Phys. Rev. B*, **4**:2701–2724, Oct 1971.

[RNN20] Dileep V. Reddy, Robert R. Nerem, Sae Woo Nam, Richard P. Mirin, and Varun B. Verma. “Superconducting nanowire single-photon detectors with 98% system detection efficiency at 1550nm.” *Optica*, **7**(12):1649–1653, Dec 2020.

[RSD10] Wade G. Rellergert, Scott T. Sullivan, David DeMille, Richard R. Greco, Markus P. Hehlen, Ricky A. Jackson, Justin R. Torgerson, and Eric R. Hudson. “Progress towards fabrication of 229Th-doped high energy band-gap crystals for use as a solid-state optical frequency reference.” *IOP Conference Series: Materials Science and Engineering*, **15**:012005, 2010.

[RTG98] Yossi Rubner, Carlo Tomasi, and Leonidas J. Guibas. “A metric for distributions with applications to image databases.” *Sixth International Conference on Computer Vision (IEEE Cat. No.98CH36271)*, pp. 59–66, 1998.

[Saf19] Marianna S. Safronova. “The Search for Variation of Fundamental Constants with Clocks.” *Annalen der Physik*, **531**(5):1800364, 2019.

[SAZ90] C. Schober, V. N. Antonov, and P. Ziesche. “Internal Conversion Coefficient for Valence Electrons of Metals. Theory and Application to 110Ag.” *physica status solidi (b)*, **159**(2):771–779, 1990.

[SD79] M. P. Seah and W. A. Dench. “Quantitative electron spectroscopy of surfaces: A standard data base for electron inelastic mean free paths in solids.” *Surface and Interface Analysis*, **1**(1):2–11, 1979.

[Sei19] B. Seiferle. *Characterization of the* ^{229}*Thnuclear clock transition.* PhD thesis, Ludwig–Maximilians–Universität München, 2019.

[SGH20] Tomas Sikorsky, Jeschua Geist, Daniel Hengstler, Sebastian Kempf, Loredana Gastaldo, Christian Enss, Christoph Mokry, Jörg Runke, Christoph E. Düllmann, Peter Wobrauschek, Kjeld Beeks, Veronika Rosecker, Johannes H. Sterba, Georgy Kazakov, Thorsten Schumm, and Andreas Fleischmann. "Measurement of the ^{229}Th Isomer Energy with a Magnetic Microcalorimeter." *Phys. Rev. Lett.*, **125**:142503, Sep 2020.

[Sha85] Alexander Shapiro. "Asymptotic Distribution of Test Statistics in the Analysis of Moment Structures Under Inequality Constraints." *Biometrika*, **72**(1):133–144, 1985.

[Sha88] A. Shapiro. "Towards a Unified Theory of Inequality Constrained Testing in Multivariate Analysis." *International Statistical Review / Revue Internationale de Statistique*, **56**(1):49–62, 1988.

[SKS18] Simon Stellmer, Georgy Kazakov, Matthias Schreitl, Hendrik Kaser, Michael Kolbe, and Thorsten Schumm. "Attempt to optically excite the nuclear isomer in ^{229}Th." *Phys. Rev. A*, **97**:062506, Jun 2018.

[SLC90] M. Steinmayer, K. E. G. Löbner, L. Corradi, U. Lenz, P. R. Pascholati, U. Quade, K. Rudolph, W. K. Schomburg, and S. J. Skorka. "Isomeric state at 945.3 keV in237Np." *Zeitschrift für Physik A Atomic Nuclei*, **337**(4):371–375, Dec 1990.

[SS11] U. I. Safronova and M. S. Safronova. "Correlation and relativistic effects in actinide ions." *Phys. Rev. A*, **84**:052515, Nov 2011.

[SSR18] Simon Stellmer, Yudai Shigekawa, Veronika Rosecker, Georgy A. Kazakov, Yoshitaka Kasamatsu, Yuki Yasuda, Atsushi Shinohara, and Thorsten Schumm. "Toward an energy measurement of the internal conversion electron in the deexcitation of the ^{229}Th isomer." *Phys. Rev. C*, **98**:014317, Jul 2018.

[ST91] V. F. Strizhov and E. V. Tkalya. "Decay channel of low-lying isomer state of the 229Th nucleus Possibilities of experimental investigation." *Soviet Physics - JETP (English Translation)*, **72**(3):387–390, 1991.

[Swa12] Erik Swanberg. *Searching for the Decay of* ^{229}*Th.* PhD thesis, University of California, Berkely, 2012.

[SWB19] Benedict Seiferle, Lars von der Wense, Pavlo V. Bilous, Ines Amersdorffer, Christoph Lemell, Florian Libisch, Simon Stellmer, Thorsten Schumm, Christoph E. Düllmann, Adriana Pálffy, and Peter G. Thirolf. "Energy of the 229Th nuclear clock transition." *Nature*, **573**(7773):243–246, Sep 2019.

[SWT17] Benedict Seiferle, Lars von der Wense, and Peter G. Thirolf. "Lifetime Measurement of the ^{229}Th Nuclear Isomer." *Phys. Rev. Lett.*, **118**:042501, Jan 2017.

[TB09] Fabien Tran and Peter Blaha. "Accurate Band Gaps of Semiconductors and Insulators with a Semilocal Exchange-Correlation Potential." *Phys. Rev. Lett.*, **102**:226401, Jun 2009.

[Tka01] E. V. Tkalya. "Spontaneous multipole radiation in a condensed medium." *Journal of Experimental and Theoretical Physics*, **92**(1):61–68, Jan 2001.

[Tka11] E. V. Tkalya. "Proposal for a Nuclear Gamma-Ray Laser of Optical Range." *Phys. Rev. Lett.*, **106**:162501, Apr 2011.

[Tor49] H. C. Torrey. "Transient Nutations in Nuclear Magnetic Resonance." *Phys. Rev.*, **76**:1059–1068, Oct 1949.

[TOZ24] J. Tiedau, M. V. Okhapkin, K. Zhang, J. Thielking, G. Zitzer, E. Peik, F. Schaden, T. Pronebner, I. Morawetz, L. Toscani De Col, F. Schneider, A. Leitner, M. Pressler, G. A. Kazakov, K. Beeks, T. Sikorsky, and T. Schumm. "Laser Excitation of the Th-229 Nucleus." *Phys. Rev. Lett.*, **132**:182501, Apr 2024.

[TSJ15] E. V. Tkalya, Christian Schneider, Justin Jeet, and Eric R. Hudson. "Radiative lifetime and energy of the low-energy isomeric level in ^{229}Th." *Phys. Rev. C*, **92**:054324, Nov 2015.

[TT79] W. Thomson and P. G. Tait. *Elements of Natural Philosophy Part One - Science.* P. F. Collier and Son, 1879.

[TVL96] E V Tkalya, V O Varlamov, V V Lomonosov, and S A Nikulin. "Processes of the nuclear isomer 229mTh(3/2+, 3.5 ± 1.0 eV) resonant excitation by optical photons." *Physica Scripta*, **53**(3):296, mar 1996.

[TWK23] Gregor G. Taylor, Alexander B. Walter, Boris Korzh, Bruce Bumble, Sahil R. Patel, Jason P. Allmaras, Andrew D. Beyer, Roger O'Brient, Matthew D. Shaw, and Emma E. Wollman. "Low-noise single-photon counting superconducting nanowire detectors at infrared wavelengths up to 29µm." *Optica*, **10**(12):1672–1678, Dec 2023.

[TZT23] J Thielking, K Zhang, J Tiedau, J Zander, G Zitzer, M V Okhapkin, and E Peik. "Vacuum-ultraviolet laser source for spectroscopy of trapped thorium ions." *New Journal of Physics*, **25**(8):083026, aug 2023.

[VKB15] V. B. Verma, B. Korzh, F. Bussières, R. D. Horansky, S. D. Dyer, A. E. Lita, I. Vayshenker, F. Marsili, M. D. Shaw, H. Zbinden, R. P. Mirin, and S. W. Nam. "High-efficiency superconducting nanowire single-photon detectors fabricated from MoSi thin-films." *Opt. Express*, **23**(26):33792–33801, Dec 2015.

[VZS02] A. Verevkin, J. Zhang, Roman Sobolewski, A. Lipatov, O. Okunev, G. Chulkova, A. Korneev, K. Smirnov, G. N. Gol'tsman, and A. Semenov. "Detection efficiency of large-active-area NbN single-photon superconducting detectors in the ultraviolet to near-infrared range." *Applied Physics Letters*, **80**(25):4687–4689, 06 2002.

[WBS20] Lars von der Wense, Pavlo V. Bilous, Benedict Seiferle, Simon Stellmer, Johannes Weitenberg, Peter G. Thirolf, Adriana Pálffy, and Georgy Kazakov. “The theory of direct laser excitation of nuclear transitions.” *The European Physical Journal A*, **56**(7):176, Jul 2020.

[WEK92] M. Woerner, T. Elsaesser, and W. Kaiser. “Relaxation processes of hot holes in p -type germanium studied by picosecond infrared spectroscopy.” *Physical Review B*, **45**(15):8378–8387, apr 1992.

[Wen17] L. von der Wense. *On the Direct Detection of* ^{229m}*Th.* PhD thesis, Ludwig–Maximilians–Universität München, 2017.

[Wil38] S. S. Wilks. “The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses.” *The Annals of Mathematical Statistics*, **9**(1):60–62, 1938.

[Wim20] Kathrin Wimmer. “Nilsson.” https://github.com/wimmer-k/Nilsson, 2020.

[WSS17] Lars von der Wense, Benedict Seiferle, Simon Stellmer, Johannes Weitenberg, Georgy Kazakov, Adriana Pálffy, and Peter G. Thirolf. “A Laser Excitation Scheme for ^{229m}Th.” *Phys. Rev. Lett.*, **119**:132503, Sep 2017.

[WSS19] Lars C. von der Wense, Benedict Seiferle, Christian Schneider, Justin Jeet, Ines Amersdorffer, Nicolas Arlt, Florian Zacherl, Raphael Haas, Dennis Renisch, Patrick Mosel, Philip Mosel, Milutin Kovacev, Uwe Morgner, Christoph E. Düllmann, Eric R. Hudson, and Peter G. Thirolf. “The concept of laser-based conversion electron Mössbauer spectroscopy for a precise energy determination of 229mTh.” *Hyperfine Interactions*, **240**(1):23, Mar 2019.

[YCL90] B.A. Young, B. Cabrera, A.T. Lee, C.J. Martoff, B. Neuhauser, and J.P. McVittie. “Phonon-mediated detection of X-rays in silicon crystals using superconducting transition edge phonon sensors.” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **288**(1):119–124, 1990. Proceedings of the Fifth European Symposium on Semiconductors Detectors.

[YLR92] Xue-Feng Yang, Jean-Louis Lemaire, François Rostas, and Joëlle Rostas. “VUV laser absorption study at 110.6 nm of the rotationally structured $1\mathit{\Pi}3$g, 3p$\mathit{\Pi}_u$ $3\mathit{\Sigma}_u$ Rydberg state of CO2.” *Chemical Physics*, **164**(1):115–122, 1992.

[YQW18] Lixing You, Jia Quan, Yong Wang, Yuexue Ma, Xiaoyan Yang, Yanjie Liu, Hao Li, Jianguo Li, Juan Wang, Jingtao Liang, Zhen Wang, and Xiaoming Xie. “Superconducting nanowire single photon detection system for space applications.” *Opt. Express*, **26**(3):2965–2971, Feb 2018.

[YSH24] Atsushi Yamaguchi, Yudai Shigekawa, Hiromitsu Haba, Hidetoshi Kikunaga, Kenji Shirasaki, Michiharu Wada, and Hidetoshi Katori. “Laser spectroscopy of triply charged 229Th isomer for a nuclear clock.” *Nature*, **629**(8010):62–66, May 2024.

[Zan12] Andrew Zangwill. *Modern Electrodynamics.* Cambridge University Press, 2012.

[ZOH24] Chuankun Zhang, Tian Ooi, Jacob S. Higgins, Jack F. Doyle, Lars von der Wense, Kjeld Beeks, Adrian Leitner, Georgy A. Kazakov, Peng Li, Peter G. Thirolf, Thorsten Schumm, and Jun Ye. "Frequency ratio of the 229mTh nuclear isomeric transition and the 87Sr atomic clock." *Nature*, **633**(8028):63–70, Sep 2024.

[ZWD24] Chuankun Zhang, Lars von der Wense, Jack F. Doyle, Jacob S. Higgins, Tian Ooi, Hans U. Friebel, Jun Ye, R. Elwell, J. E. S. Terhune, H. W. T. Morgan, A. N. Alexandrova, H. B. Tran Tan, Andrei Derevianko, and Eric R. Hudson. "^{229}ThF_4 thin films for solid-state nuclear clocks.", 2024.