Lepton-Catalyzed Nuclear Fusion: Quantum Cavity Modulated Fusion

Tyler Joseph Hamm

A senior thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Bachelor of Science

David Allred and John Ellsworth, Advisors

Department of Physics and Astronomy

Brigham Young University

ABSTRACT

Lepton-Catalyzed Nuclear Fusion: Quantum Cavity Modulated Fusion

Tyler Joseph Hamm
Department of Physics and Astronomy, BYU
Bachelor of Science

Low-energy nuclear fusion rates at room temperature are extremely small due to the Coulomb repulsion between fusing nuclei. However, electron screening reduces this repulsion, making fusion rates highly sensitive to the local electromagnetic environment. In metallic systems, prior experimental studies have reported anomalously high fusion rates attributed to enhanced electron screening effects, motivating the investigation of whether microscopic surface features can account for this behavior. This thesis investigates whether microscopic cavities, such as those naturally present within metal lattices or introduced through deliberate surface modification, can enhance fusion rates by increasing electron screening near reacting deuterons. To model this effect, three theoretical approaches were developed to predict fusion rates. First, a semi-classical framework based on classical electrodynamics was explored. This framework showed very little screening in metallic environments and cavities. Second, a stochastic electrodynamics (SED) model, incorporating zero-point field contributions and cavity interactions, was analyzed. This model showed that if cavities are reduced below a certain length, determined by the characteristic frequency of a ground-state deuterium electron, the electron orbital radius would decrease, leading to increased nuclear screening. Finally, a quantum electrodynamics (QED) approach involving cavity-modified photon-mediated interactions was discussed. This model produced similar results to the SED model, but was done more qualitatively. To test these predictions experimentally, a dedicated vacuum system was constructed to form a parallel-plate cavity whose interior surfaces can be loaded with deuterium and to measure fusion products using complementary neutron and charged-particle detection methods. Experimental fusion measurements were performed using the completed system; however, the resulting data were ultimately inconclusive and did not provide definitive evidence of cavity-enhanced fusion. Nevertheless, the theoretical predictions, experimental system, and measurements presented here establish a framework for continued testing of cavity-enhanced screening as a possible contributor to anomalously high low-energy fusion rates previously observed in metallic environments.

Keywords: Electron Screening, Fusion, Cavity, Coulomb Barrier, Zero-Point Energy, Lepton-Catalyzed Fusion

ACKNOWLEDGMENTS

I am extremely grateful for the guidance, encouragement, and mentorship of my research advisors, John Ellsworth and David Allred. Their consistent support, patience, and insight throughout this project and my education were essential in shaping the direction of this work and my development as a researcher and a disciple of Christ. I especially appreciate their willingness to meet regularly, answer questions, and provide thoughtful feedback as I learned how to approach scientific problems with greater rigor and creativity.

In addition, I would like to thank the countless faculty and staff who taught me the fundamentals of physics and helped me apply these principles to meaningful research. This especially includes my first research advisor Troy Munro. Furthermore, I am grateful for the graduate students who introduced me to scientific research, Ryan Ruth and Jake Numbers. I would also like to thank the other students I worked alongside in the Laboratory Nuclear Astrophysics Research group (LNAR), including Rhett Lundell, Isaac Willden, and Blake Benefeito. Their collaboration and support made the research environment both productive and enjoyable.

I thank the College of Computational, Mathematical, and Physical Sciences. In particular, I am grateful to the Department of Physics & Astronomy for funding my work and supporting me as a research assistant throughout this project.

I am also grateful for my daughter, Lucy, who always seemed to take every moment possible to give me a break from my work and remind me to not take a single section of this work she deleted in one click of the keyboard so seriously. Finally, I wish to express my deepest gratitude to my wife, Mary, whose love, sacrifice, and support were integral in my ability to complete this thesis and my undergraduate education.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Introduction

Nuclear fusion is the fundamental process by which stars generate energy, governing their structure, evolution, and observable properties [1–4]. Since 1920, when Sir Arthur Eddington first proposed nuclear fusion as the source of stellar energy, its characteristics have been the subject of extensive theoretical and experimental investigation [5–8]. Much of the current research in nuclear physics focuses on high-energy environments such as tokamaks, particle accelerators, and inertial confinement systems. These processes seek to overcome the Coulomb repulsion between two fusing nuclei by increasing the kinetic energy of the fusion particles. However, this high-energy regime for fusion stands in contrast to astrophysical fusion, which occurs under comparatively low-energy conditions. In the Sun, for example, fusion takes place at core temperatures on the order of 10^7 K, corresponding to particle energies of approximately 1 keV [5, 9]. These energies are far below those required to classically overcome the electrostatic repulsion between positively charged nuclei (see Section 2.1). This discrepancy historically presented a contradiction between observed stellar lifetimes and predictions based on classical physics.

This apparent contradiction was resolved with the development of quantum mechanics. Within this framework, quantum tunneling allows particles to penetrate classically forbidden regions where the potential energy exceeds the particle's kinetic energy, enabling nuclear reactions at sub-barrier energies [10]. While the quantitative treatment of tunneling and barrier penetration is discussed in greater detail in Section 2.2, it is sufficient here to note that fusion rates may be enhanced either by increasing the kinetic energy of the reacting nuclei or by effectively reducing the electrostatic repulsion between them.

In modern applications, much of the study of nuclear fusion has focused on increasing the thermal energies of fusion reactants in order to overcome the Coulomb barrier. This approach is pursued primarily in laboratory-based confinement systems, including magnetic confinement devices such as tokamaks (e.g., ITER) and inertial confinement facilities such as the system found at the National Ignition Facility [11, 12]. Historically, however, some attention has also been given to increasing fusion rates through the reducing of the effective Coulomb repulsion.

The reduction of Coulomb repulsion is done primarily through the use of negatively charged particles, such as the electron or muon, collectively labeled "leptons." These leptons surround a nucleus and help shield its effective charge from nearby nuclei. This mechanism of reducing the Coulomb repulsion through the use of leptons is referred to as lepton-screening or lepton-catalysis. Although there is a slight difference between these two definitions, they will be used fairly interchangeably in this thesis. The exact differences between lepton-screening and lepton-catalysis will be discussed later in the paper. Although these effects are discussed in greater detail in Chapter 2, a brief historical overview of lepton-catalyzed fusion is provided in the following paragraphs.

Early work by scientists named Steven E. Koonin and Michael Nauenberg estimated the naturally occurring deuterium-deuterium fusion rate to be approximately $10^{-64}\,\mathrm{s}^{-1}$ under ambient conditions [13]. The unit of inverse seconds or s^{-1} refers to the number of fusion events per atom per second.

This means that a single fusion event for a given collection of deuterium nuclei is expected to occur every 10^{64} seconds. As a comparison, the life of the universe is only on the order of 10^{17} seconds.

Building upon the electrons' ability to block or screen the positive charges of a nucleus, Frederick Charles Frank theorized the concept of muon-catalyzed fusion [14, 15]. Sakharov independently proposed a similar concept shortly afterward [16]. Frank, and later Sakharov, proposed replacing an electron in the ground-state orbit of a deuterium (or any hydrogen isotope) atom with a muon (μ) [17, 18]. The muon is approximately 207 times heavier than the electron. Due to the muon's much larger mass, the resulting bound state of the muon has a characteristic radius much smaller than that of a typical electronic orbital, leading to substantially increased nuclear screening and fusion rate enhancements of several orders of magnitude [19, 20].

This process was first observed by Frank in film emulsions in 1949 and later by Luis Walter Alvarez in a hydrogen bubble chamber in 1957 [21]. By the mid-1980s, Steven E. Jones and collaborators had advanced experimental muon-catalyzed fusion to approximately 150 fusion reactions per muon, demonstrating the substantial enhancement in fusion rates that could be achieved through the altered screening environment produced by the muon [22]. Jones subsequently extended this general concept beyond muon catalysis, investigating whether modifications to the electronic and condensed-matter environment surrounding hydrogen isotopes could similarly enhance fusion at low energies.

Jones from Brigham Young University then proposed the possibility that electrons could have a similar heightened ability of screening when placed in the correct environments [23]. One of those proposed environments was metallic systems, where the free electrons in the metal could be repositioned (due to the incoming charge) and averaged to create a greater-than-average screening effect.

Subsequent research expanded to fusion processes in these metallic environments, where lattice effects and high electron densities were proposed to increase screening and enhance fusion

rates [23–25]. Although the data showed high fusion rates at low-energy levels, it quickly lost research interest. In recent decades, renewed experimental and theoretical interest in low-energy deuterium-deuterium fusion in condensed matter systems, including metallic lattices and other low work-function materials, has further motivated investigation into the physical mechanisms responsible for these reported enhancements [26–28].

1.2 Motivation

The continually advancing experimental techniques and renewed theoretical interest have prompted a reexamination of low-energy fusion phenomena, particularly mechanisms involving enhanced electron screening in dense or confined environments. Current modeling techniques and electron screening predictions have proven insufficient to fully explain observed anomalies, motivating the development of improved theoretical frameworks and a more complete physical understanding.

Notably, prior experimental work of nuclear fusion within metallic environments at Brigham Young University reported increased reproducibility of enhanced fusion signals in metallic systems when the surface of the metal was scratched using utility blades [25]. These scratches introduce surface features on the metal that can be modeled as microscopic cavities, with characteristic separations determined by the geometry and spacing of the scratches. Such features produce microscopic cavities and defects, comparable to those inherent in the metal's lattice structure, that may provide a physical origin for increased lepton-screening effects. It is therefore plausible that such cavities–either engineered or inherently present on metallic surfaces–could modify the local electromagnetic environment, leading to increased electron screening and, consequently, enhanced fusion rates.

This work investigates whether increased electron screening, arising from reduced zero-point energy density in quantum-confined cavities, can contribute to enhanced deuterium-deuterium fusion

rates in metallic systems. This is approached through theoretical modeling of the electronic structure and screening effects within microscopic cavities, alongside the design of an experimental system capable of probing these conditions. By comparing the theoretical predictions developed in this work with prior experimental observations and the anticipated results from the constructed system, this thesis aims to assess the viability of cavity-enhanced screening as a potential explanation for deviations from standard fusion rate predictions.

1.3 Organization of Thesis

This thesis first reviews the historical development of nuclear fusion theory and relevant contemporary models, providing the background necessary to understand low-energy fusion processes and electron screening effects. It then introduces the theoretical approaches developed in this work for calculating deuterium-deuterium fusion rates in quantum-confined environments. The experimental methods used to fabricate and characterize quantum cavities for testing deuterium-deuterium fusion rates are subsequently described. This is followed by a presentation and discussion of the thesis's results, including the experimental results. The thesis concludes with a summary of the findings and their implications.

Chapter 2

Background

In order to more fully understand the modeling of Chapter 3 and the experimental method discussed in Chapter 4, this section will discuss the nuclear potential model used in this work, the current mathematical understanding of lepton-screening, and the possible effects of quantum cavities on nuclear fusion.

2.1 Nuclear Interaction and the Coulomb Barrier

In nature, four fundamental interactions are known to govern physical processes: gravity, electromagnetism, the strong nuclear force, and the weak nuclear force. For the purposes of modeling nuclear fusion, however, only a subset of these interactions plays a significant role. This section will discuss the relevant simplifications and describe how each force is treated in the context of fusion.

2.1.1 Weak Nuclear and Gravitational Forces

First, the weak nuclear force can be ignored because this work focuses on stable internuclear interactions relevant to fusion. Although the weak force plays an integral part in various astrophysical

processes, such as the proton-proton chain reaction in the sun [29], neglecting the weak force allows for a much simpler calculation of a single fusion event without significantly affecting the interaction being modeled.

In addition, although gravitational forces do exert a non-zero force in nuclear interactions, especially in astrophysical systems, their magnitude is negligible in comparison to the other three forces at characteristic nuclear distances at low-energy in the laboratory. For two protons separated by a characteristic nuclear distance $r \sim 1\,\text{fm}$, the gravitational force is

$$F_g = \frac{Gm_p^2}{r^2} \approx 10^{-34}\,\text{N},$$

whereas the Coulomb force is

$$F_e = \frac{1}{4\pi\varepsilon_0}\frac{e^2}{r^2} \approx 10^2\,\text{N}.$$

Thus,

$$\frac{F_g}{F_e} \sim 10^{-36},$$

meaning gravity is roughly 36 orders of magnitude weaker than electromagnetism at the femtometer scale. Since the strong nuclear force is comparable to or stronger than the Coulomb interaction at these distances, gravitational effects are entirely negligible in laboratory-scale nuclear potentials and may be safely omitted. A general comparison of the relative strengths of these forces is shown in Figure 2.1. The weak force is omitted, consistent with the previous discussion, as its influence is confined to sub-femtometer length scales and does not significantly impact the interactions considered here.

This means that the only forces that have a non-negligible impact on nuclear fusion rates are the electromagnetic (Coulomb) force and the strong nuclear force. The Coulomb force acts as a long-range repulsive barrier between positively charged nuclei, while the strong nuclear force

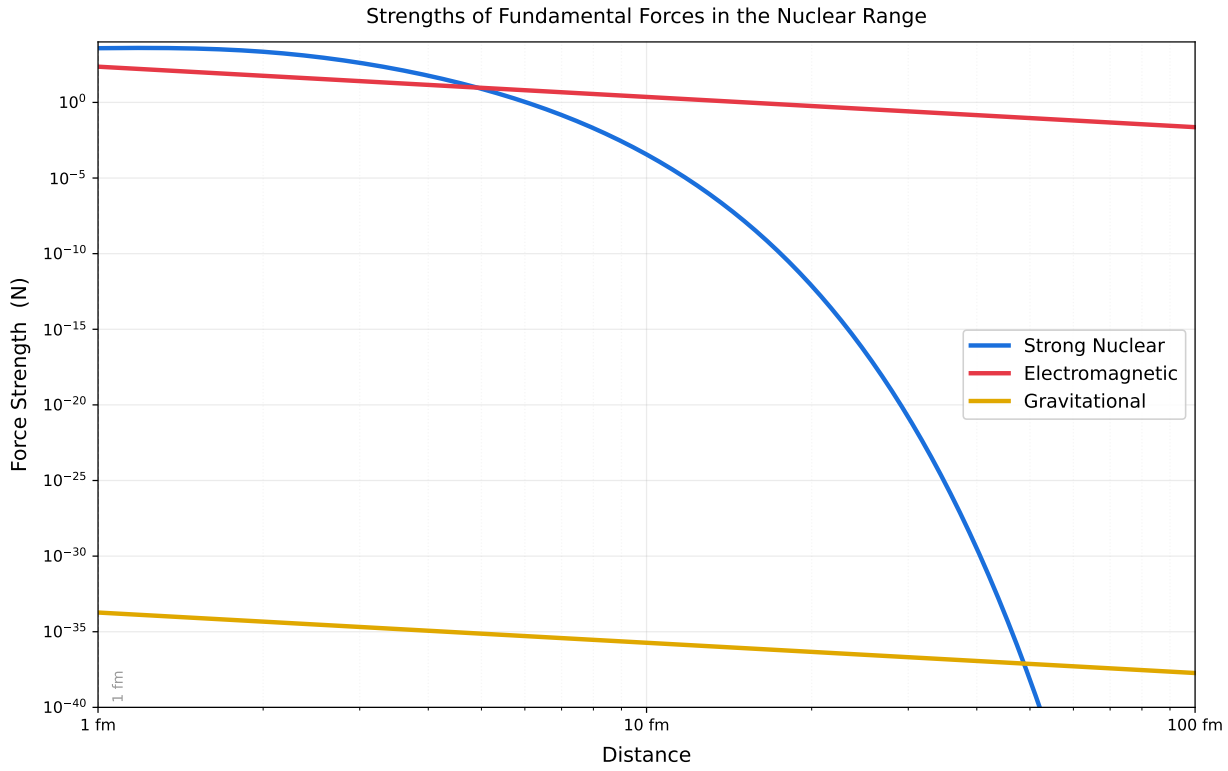

Figure 2.1 Strengths of the four fundamental forces between two protons over the nuclear fusion range (1–100 fm). The strong nuclear force (Woods-Saxon model) and Coulomb repulsion dominate across this regime; gravity and the weak force (which is already below gravity in terms of magnitude at this range) are negligible due to their extremely small coupling strength and sub-femtometer interaction range, respectively, and are excluded from further analysis.

provides a short-range attractive interaction once the nuclei are sufficiently close. The competition between these two forces gives rise to the Coulomb barrier, which must be overcome for fusion to occur.

Using these two forces, an accurate model describing what an incoming nucleus would experience while approaching an at-rest target nucleus can be designed. In this nuclear potential model, all electrons are initially neglected, with any associated screening effects introduced later. This deliberate omission allows the model to remain broadly applicable across a wide range of lepton-screening environments, as discussed in Chapter 3.

2.1.2 Electromagnetic Force

The electromagnetic force between two nuclei can be modeled as two positive point charges. In this model, the nuclei can be treated as relatively stable compared to the surrounding electronic structure. Treating the nuclei as stable is an approximation referred to as the Born-Oppenheimer approximation. This approximation is accurate because nuclear interactions occur on the order of 4-10 femtometers, whereas the nuclear radius of a hydrogen nucleus is around 1.7 femtometers. Any quantum fluctuations that may occur in the charge distribution of the nucleus can be neglected without introducing significant error, since the spatial extent of these fluctuations is small compared to the characteristic interaction distance and does not substantially alter the effective Coulomb potential between the nuclei. Modeling the nuclei as point charges within this Born-Oppenheimer approximation allows the use of the Coulomb potential

$$U_e = \frac{1}{4\pi\varepsilon_0}\frac{q_1 q_2}{r} = k\frac{q_1 q_2}{r} \ . \tag{2.1}$$

Because the nuclei of interest are hydrogen nuclei, q_1 and q_2 of equation 2.1 can both be set to a value of $+1$:

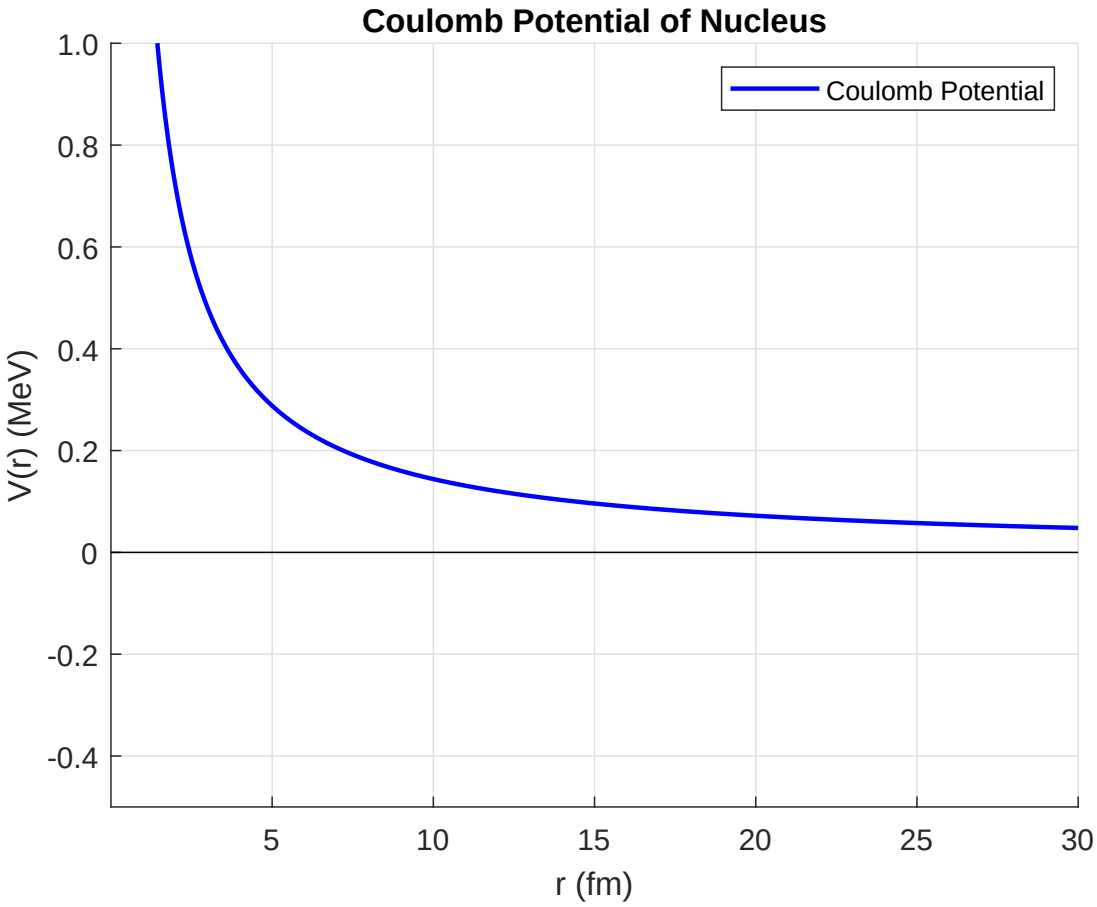

Figure 2.2 Coulomb potential between two hydrogen nuclei plotted over a part of the nuclear interaction range (1-30 fm). The $1/r$ dependence produces a rapidly increasing repulsive barrier at small separations, reaching many MeV within a few femtometers.

$$U_e = k\frac{e^2}{r}\ . \tag{2.2}$$

The Coulomb potential is very large while the nuclei are close and decreases monotonically with a $V(r) \propto 1/r$ relation as the separation between nuclei increases. This behavior is shown in Figure 2.2. The vertical axis is expressed in MeV, the natural energy scale for nuclear processes. The horizontal axis spans femtometer distances which correspond to the range of nuclear interactions.

2.1.3 Nuclear Strong Force

The nuclear strong force is a complex interaction that, in its most fundamental description, is governed by Quantum Chromodynamics (QCD). A full QCD treatment of nuclear binding, however, lies beyond the scope of this work. In order to simplify the analysis while still capturing the essential physics of nuclear attraction, this thesis adopts the use of a simplified effective potential model to approximate the strong nuclear force. Appendix A reviews several commonly used nuclear potential models. In this thesis, the Woods-Saxon potential is used due to its simplicity and success in reproducing basic features of nuclear structure.

The Woods-Saxon potential is a mean-field approximation, meaning that the complicated many-body interactions between individual nucleons are replaced by an average potential experienced by a single nucleon within the nucleus. Rather than explicitly modeling the exchange of mesons or the underlying quark-gluon dynamics, the Woods-Saxon form provides a smooth potential well that reflects two key experimental observations: the nuclear force is strongly attractive at short range, and it rapidly decreases beyond the nuclear surface. The gradual transition at the boundary also accounts for the fact that nuclei do not have perfectly sharp edges.

The Woods-Saxon potential, expressed as a function of radial distance from the nuclear center, is given by

$$U_n(r) = -\frac{V_0}{1+e^{\frac{r-R}{a}}}\ , \tag{2.3}$$

where V_0 is the depth of the nuclear potential well, R is the effective nuclear radius, and a is the surface diffuseness parameter, which characterizes how quickly the potential falls off at the nuclear boundary. These parameters are determined empirically and are chosen to match experimental nuclear binding data. The values used in this work are

$$V_0 = 35\,\mathrm{MeV}, \tag{2.4}$$

$$R = 2.12562 \times 10^{-15}\,\mathrm{m} = 2.12562\,\mathrm{fm}, \tag{2.5}$$

$$a = 0.5 \times 10^{-15}\,\mathrm{m} = 0.5\,\mathrm{fm}. \tag{2.6}$$

These values are representative of commonly used Woods-Saxon potential parameters for light nuclei and are consistent with standard nuclear physics references [30].

Using these experimentally motivated parameters, the nuclear potential can be calculated as a function of distance from the center of the nucleus. Figure 2.3 shows the Woods-Saxon potential well, illustrating the short-range attractive behavior of the strong force and the rapid decrease in magnitude beyond the nuclear radius.

2.1.4 The Combined Forces

Assuming that the nuclei are in a vacuum, the full nuclear potential can be modeled using the electromagnetic and nuclear strong potentials. Since potentials are linear, the individual contributions can be added directly to obtain the total interaction potential. This combined potential can then be given by

$$U_{tot}(r) = U_e + U_n = k\frac{e^2}{r} - \frac{V_0}{1 + e^{\frac{r-R}{a}}}\,. \tag{2.7}$$

This form of the potential highlights the competition between long-range Coulomb repulsion and short-range nuclear attraction.

Equation 2.7 is plotted in Figure 2.4. Although the full range of the potential is not completely shown due to the bounds set on the y-axis scale, the key features of the interaction are clearly illustrated within the plotted region.

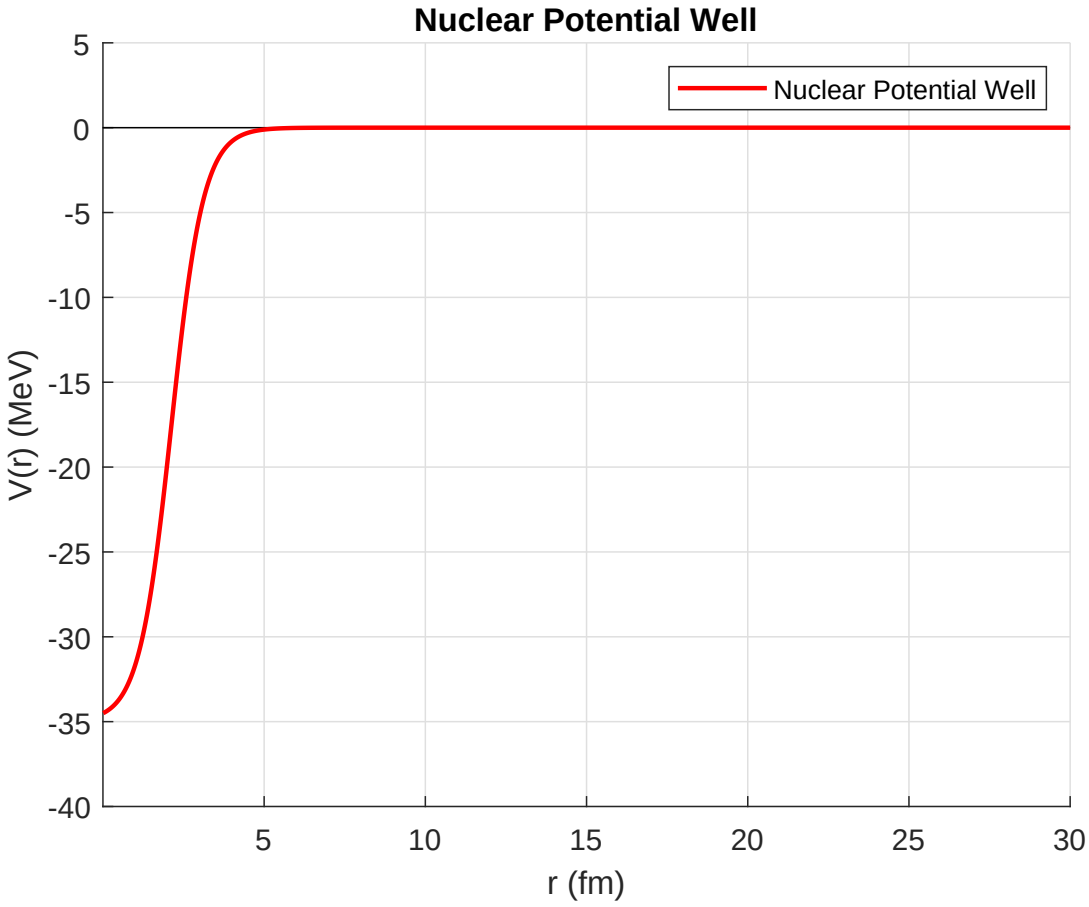

Figure 2.3 Woods-Saxon potential for a deuterium nucleus plotted over a part of the nuclear interaction range (1-30 fm). Just outside the nuclear radius, the potential rapidly increases toward zero and then remains close to zero for larger values of r.

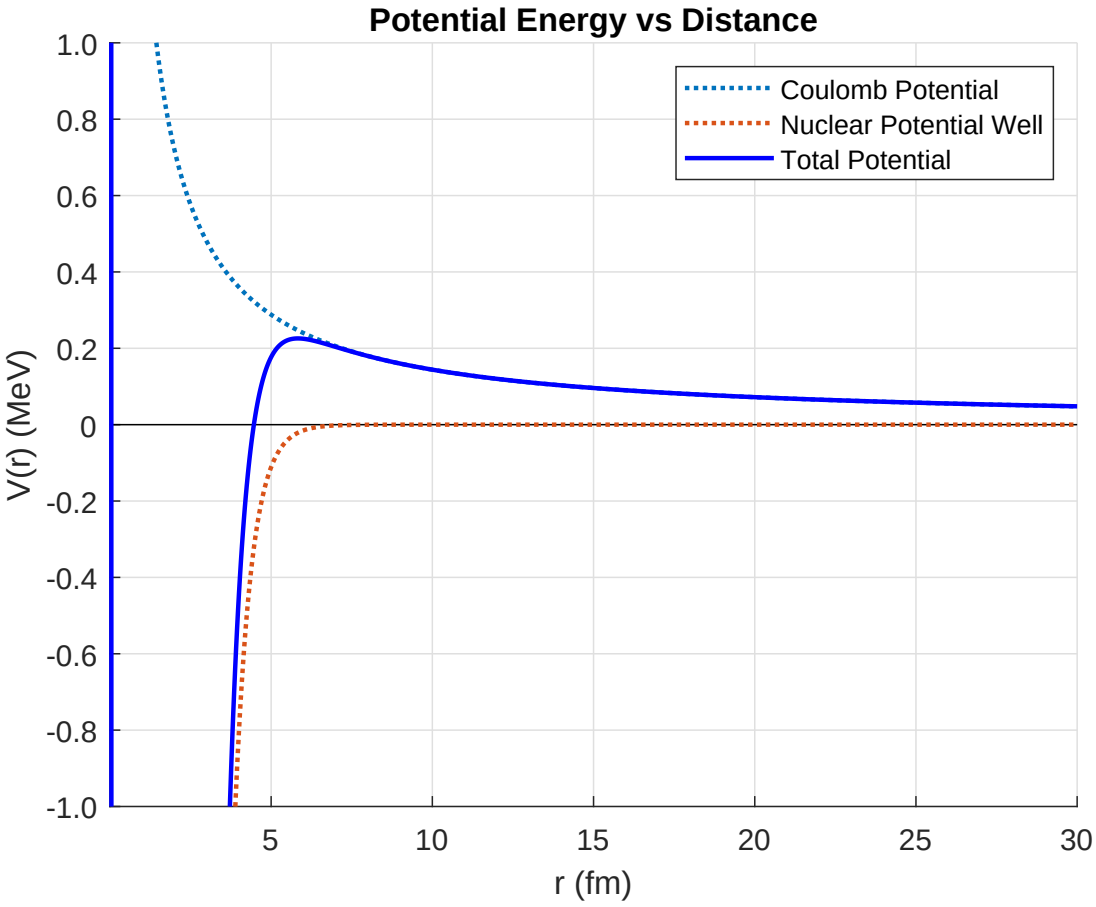

Figure 2.4 Total nuclear potential for two hydrogen nuclei, highlighting the deep nuclear well at short range and the rising Coulomb barrier that must be overcome (or tunneled through) for fusion to occur.

2.2 Quantum Tunneling

Figure 2.4 helps visualize the potential that an incoming nucleus would encounter. It also illustrates the classical behavior of a particle with a given energy approaching a target nucleus. Any incoming particle with energy less than the barrier height (approximately 0.22 MeV, occurring near $r \approx 5.5$ fm) will be reflected and unable to fuse with the target nucleus, since the region beneath the Coulomb barrier is classically forbidden.

For example, Figure 2.5 shows how an incoming nucleus (moving from the right in the perspective of the graph) with an energy of 0.1 MeV does not have sufficient energy to overcome the Coulomb barrier and is turned away from the nucleus at a separation of roughly 15 fm. In this classical picture, low-energy fusion is therefore impossible.

This is what makes nuclear fusion at low energies seemingly unattainable. The Sun, as referenced in the introduction, has typical thermal energies on the order of 1 keV, which is far below the height of the Coulomb barrier. If nuclei were required to overcome the barrier purely through classical motion, fusion could not occur at the observed rates in stellar environments. In reality, quantum mechanics provides an alternative mechanism through tunneling.

2.2.1 Conceptual Framework

Quantum tunneling arises from the wave-like nature of particles. Rather than following a single classical trajectory, an incoming nucleus is described by a quantum wavefunction that extends throughout space. When this wavefunction encounters a potential barrier, it does not vanish at the classical turning point. Instead, the wavefunction penetrates into the classically forbidden region and decays exponentially within the barrier. Because the wavefunction remains nonzero, there exists a finite probability that the particle may be found on the other side of the barrier.

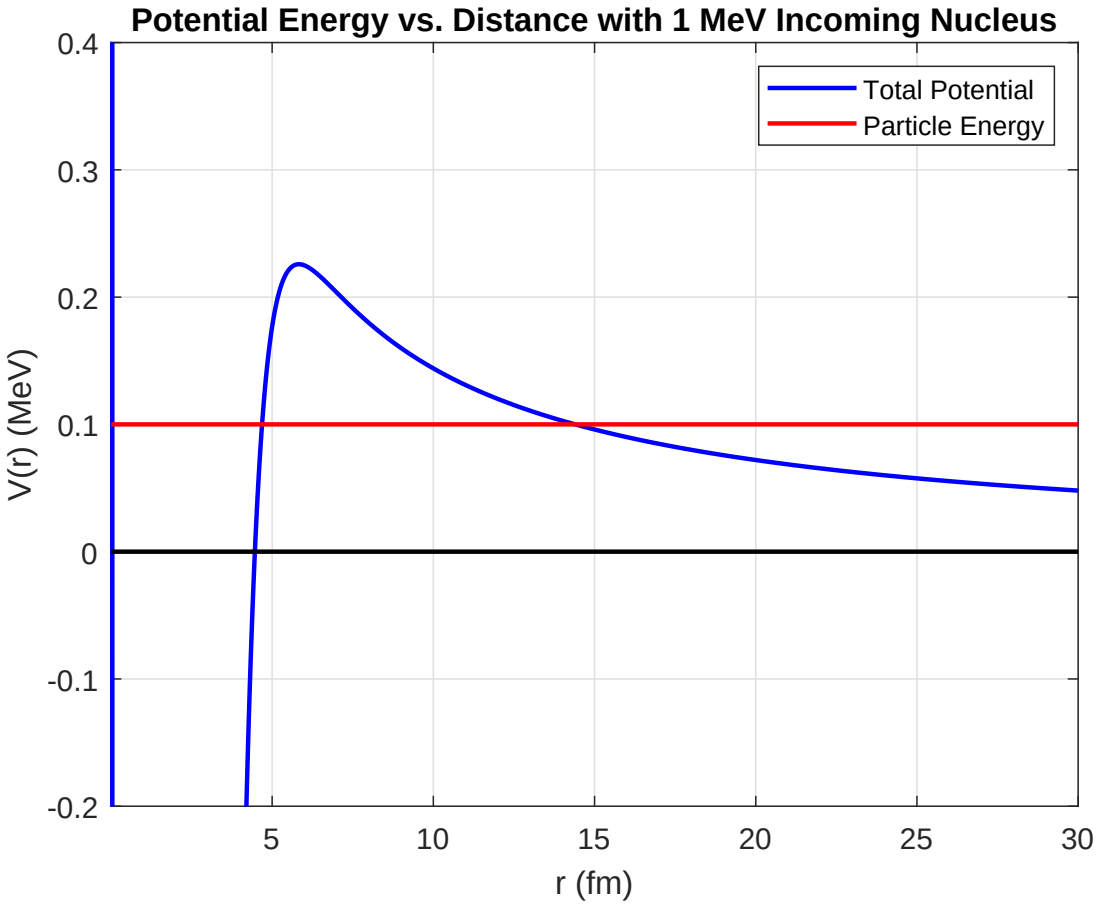

Figure 2.5 Total interaction potential between two nuclei with an incoming particle energy of 0.1 MeV. The Coulomb barrier prevents the particle from reaching the nuclear interaction region in a purely classical description, resulting in reflection at large separation distances.

In Figure 2.5, the quantum wavefunction would not vanish in the 5–15 fm range, which is classically forbidden. As a result, there is a finite probability for the incoming nucleus to be found within the 0–5 fm region, where the strong nuclear force becomes dominant. If the nuclei reach this region, fusion can occur.

Quantum tunneling is therefore the process by which a particle can pass through a classically forbidden region and access another classically allowed region without possessing sufficient classical energy to overcome the barrier. Although the tunneling probability is extremely small at low energies, even a small probability becomes significant in environments where an enormous number of particle interactions occur, such as stellar interiors.

For nuclear fusion, tunneling provides the essential mechanism that allows two positively charged nuclei to approach close enough for the attractive strong nuclear force to dominate over Coulomb repulsion. The key theoretical problem then becomes determining the probability that an incoming nucleus tunnels through the Coulomb barrier and reaches the nuclear interaction region, where fusion can take place.

2.2.2 WKB Approximation

Now that the possibility and conceptual understanding of tunneling has been discussed, it is important to note how the probability of tunneling is calculated. Stringent tunneling calculations are very tedious and can be solved exactly for only a few simple potentials. All other potentials (such as the Coulomb barrier presented in this thesis) use approximations to calculate tunneling probabilities. This work will use the Wentzel-Kramers-Brillouin (WKB) approximation. The WKB approximation is a method for finding solutions to the linear differential equations that come about due to the quantum tunneling. The detailed derivation, which appears in almost any undergraduate quantum mechanics text, is not repeated here, as it is not central to the present analysis. It is

important to note the final equations which help calculate tunneling rates. The tunneling rate is given by

$$T = e^{-2\beta} \left(1 - \frac{1}{4}e^{-2\beta}\right)^{-2} \tag{2.8}$$

where

$$\beta = \frac{1}{\hbar}\int_{x_L}^{x_R} dx' |p(x')| \tag{2.9}$$

and

$$p(x) = \sqrt{2m(V(x) - E)}\,. \tag{2.10}$$

Although the WKB approximation has more steps of simplification in order to calculate the tunneling probability, most of this simplification was done numerically using code. The code used for this work will be left for the reader in Appendix L. Now that the framework for quantum tunneling has been discussed, this thesis will continue to discuss how different physical parameters or experimental methods could modify the resulting tunneling probability from the WKB approximation.

2.3 Mechanisms for Increasing Tunneling Probability

Nuclear fusion is fundamentally limited by the Coulomb barrier, the electromagnetic repulsion between positively charged nuclei. In order for fusion to occur, nuclei must either approach with sufficient energy to penetrate this barrier, or the effective barrier itself must be reduced through environmental or quantum effects. The primary mechanisms studied in fusion research fall into these two broad categories.

2.3.1 Thermal Energy and Barrier Penetration

The most direct way to increase fusion rates is by increasing the kinetic energy of the incoming nuclei. Higher particle energies allow incoming nuclei to reach closer to the nucleus before being repulsed by the Coulomb potential. Even when the barrier is not fully overcome classically, increasing the energy shortens the classically forbidden tunneling region and therefore increases the probability of fusion. This behavior is clearly reflected in the exponential dependence of the WKB tunneling approximation, where greater energy corresponds to a smaller suppression factor and a higher tunneling probability (see equations 2.8, 2.9, and 2.10).

Although this approach forms the basis of most conventional fusion research, thermal energy alone still has many obstacles and difficulties to achieving a highly efficient fusion rate. Magnetic confinement systems such as tokamaks, as well as inertial confinement approaches, all face major challenges in reaching and sustaining the necessary temperatures. Furthermore, the energy required to heat and confine fusion fuel often exceeds the energy released. At Lawrence Livermore National Laboratory's National Ignition Facility, for example, ignition has been demonstrated, yet the total energy needed to charge and power the laser input remains far larger than the energy delivered to, or received from, the fusion target. Thus, while thermal fusion remains an important topic and goal for fusion research, it continues to face significant practical and energetic limitations.

2.3.2 Coulomb Barrier Reduction Mechanisms

An alternative approach is to reduce the effective electromagnetic repulsion between fusing nuclei. Rather than increasing the kinetic energy of the reactants, one may instead decrease the height of the Coulomb barrier through screening mechanisms. In this case, the fusion probability can increase in a similar way to the thermal approach, but without requiring the same energetic cost of heating nuclei to extreme temperatures.

The most widely studied barrier-reduction mechanism is electron screening, in which negatively charged leptons partially shield the positive nuclear charge. Screening effectively lowers the Coulomb potential experienced by approaching nuclei, thereby increasing tunneling probability at low energies. Multiple environments have been studied in which screening modifies the underlying physics of the system.

Plasma Screening

Plasma is one of the most important environments for fusion research. A plasma consists of ionized nuclei and free electrons, forming a quasi-neutral medium of charged particles. Under appropriate conditions, the free electrons collectively respond to the electric fields of interacting nuclei, producing a screening cloud that partially reduces the effective Coulomb repulsion. As two nuclei approach one another, this redistribution of electron density modifies the local electromagnetic potential, effectively lowering the barrier to fusion and increasing the tunneling probability. This screening effect is treated more rigorously in Chapter 3.

Metallic Screening

Similar to plasma environments, metals provide an environment rich with ionized nuclei and free electrons. Metals contain a sea of mobile conduction electrons that may contribute to screening effects for implanted fusion fuels. Numerous experiments have reported enhanced fusion behavior in metallic lattices [31, 32]. However, unlike plasma screening, there is not yet a complete theoretical framework that can reliably predict fusion enhancements in metal-confined systems. These challenges motivate further study of additional mechanisms that may arise in metallic environments.

Cavity Effects and Quantum Confinement

One potential contributor to enhanced screening in metals is the presence of cavities or confined surface structures. Cavities have long been known to exhibit unusual quantum behavior. Beginning with Casimir, it was shown that closely spaced conducting plates experience an attractive force due to a modification of the electromagnetic zero-point field within the confined region. If confinement alters the electromagnetic environment, it may also influence electronic structure and screening behavior near nuclei. This electromagnetic structure dependency on the zero-point field fluctuations has been theorized using various mathematical theories including stochastic electrodynamics [33](see Section 3 for more information).

These considerations have motivated and continue to motivate the study of lepton-catalysis and confinement-based screening mechanisms, including muon-catalyzed fusion, electron-based screening effects, and zero-point-energy modifications in quantum cavities. These topics will be explored in the following sections.

2.4 Screening and Catalysis Effects

One of the most well-established mechanisms for enhancing low-energy fusion rates is the reduction of the Coulomb barrier through lepton screening. In this context, a negatively charged particle partially neutralizes the repulsive Coulomb interaction between two approaching nuclei, allowing them to tunnel into the short-range region where the strong nuclear force governs the interaction. The clearest and most experimentally verified example of this effect is muon-catalyzed fusion.

2.4.1 Muon-Catalyzed Fusion

Muon-catalyzed fusion became a major topic of interest in the mid-20th century as one of the first demonstrated examples of fusion enhancement through extreme lepton screening. Muon-catalyzed

fusion is characterized by replacing an electron in a hydrogen-like atom with a muon. A muon can be simply defined as a heavy electron. Because the muon carries the same charge as the electron but has a much larger mass–almost 207 times more mass–it has a characteristic orbit significantly closer to the nucleus than the electron.

The characteristic atomic radius of a hydrogen-like bound system is given by the Bohr radius,

$$a_0 = \frac{4\pi\varepsilon_0\hbar^2}{m_e e^2}. \tag{2.11}$$

For a muonic atom, the electron mass m_e is replaced by the muon mass m_μ, giving the Bohr radius for a muon to be

$$a_\mu = a_0\frac{m_e}{m_\mu}. \tag{2.12}$$

Since the muon mass is approximately

$$m_\mu \approx 207\, m_e, \tag{2.13}$$

the orbital radius shrinks by the same factor of 207. This reduction in Bohr radius brings the muon much closer to the nucleus, producing an extremely strong screening effect. In fusion systems, this effectively lowers the Coulomb barrier and allows nuclei to approach almost within nuclear interaction distances. This would lead to a very thin barrier through which the fusing nuclei would have to tunnel, leading to greatly enhanced tunneling probabilities and corresponding fusion rates.

Muon-catalyzed fusion was first theorized by Frank based on observations in photographic emulsions. Luis Alvarez and collaborators later experimentally observed muon-catalyzed fusion in a hydrogen bubble chamber [21]. Of particular importance, Alvarez observed a two-step chain reaction in which a single muon participated in successive fusion events. The resulting image, shown in Figure 2.6, demonstrated that the muon could be released following a fusion event and subsequently participate in another reaction. Later experiments by Steven Jones at Brigham Young

University and others demonstrated that muonic screening can increase fusion probabilities by many orders of magnitude compared to normal low-energy fusion rates [20].

2.4.2 Electron-Catalyzed Fusion

Muon-catalyzed fusion provides an important conceptual framework for future work. If fusion rates can be dramatically increased by shrinking the lepton orbital radius, then enhanced screening is linked to the ability of confining negative charge near the nucleus. This led Jones and others to theorize the possibility that electrons could produce a similar catalytic effect if their effective orbital distribution could somehow be compressed inward toward the nucleus.

Instead of replacing the electron with a heavier lepton, different processes must be used to help confine the electron closer to the nucleus. The key idea remains that fusion enhancement does not require complete barrier removal, but rather a reduction in the effective Coulomb repulsion through increased electron density at small radii. If an environment could modify the electron wavefunction such that the screening potential becomes stronger at nuclear length scales, the tunneling probability, and therefore fusion rates, could increase substantially.

This motivates the study of electron screening in condensed matter systems. In metals, deuterium nuclei exist within a sea of conduction electrons, and experimental results have repeatedly shown fusion rates that exceed standard theoretical screening predictions. This high fusion rate is indicative of a physical variable that has yet to be accounted for. In particular, confined electronic environments such as metallic surface cavities may provide boundary conditions that alter the zero-point electromagnetic field and modify electron behavior in ways analogous, though far weaker, than muonic compression.

The work presented in this thesis builds directly on the motivation that while muon-catalyzed fusion demonstrates the extreme limit of lepton-induced screening, cavity-enhanced electron screen-

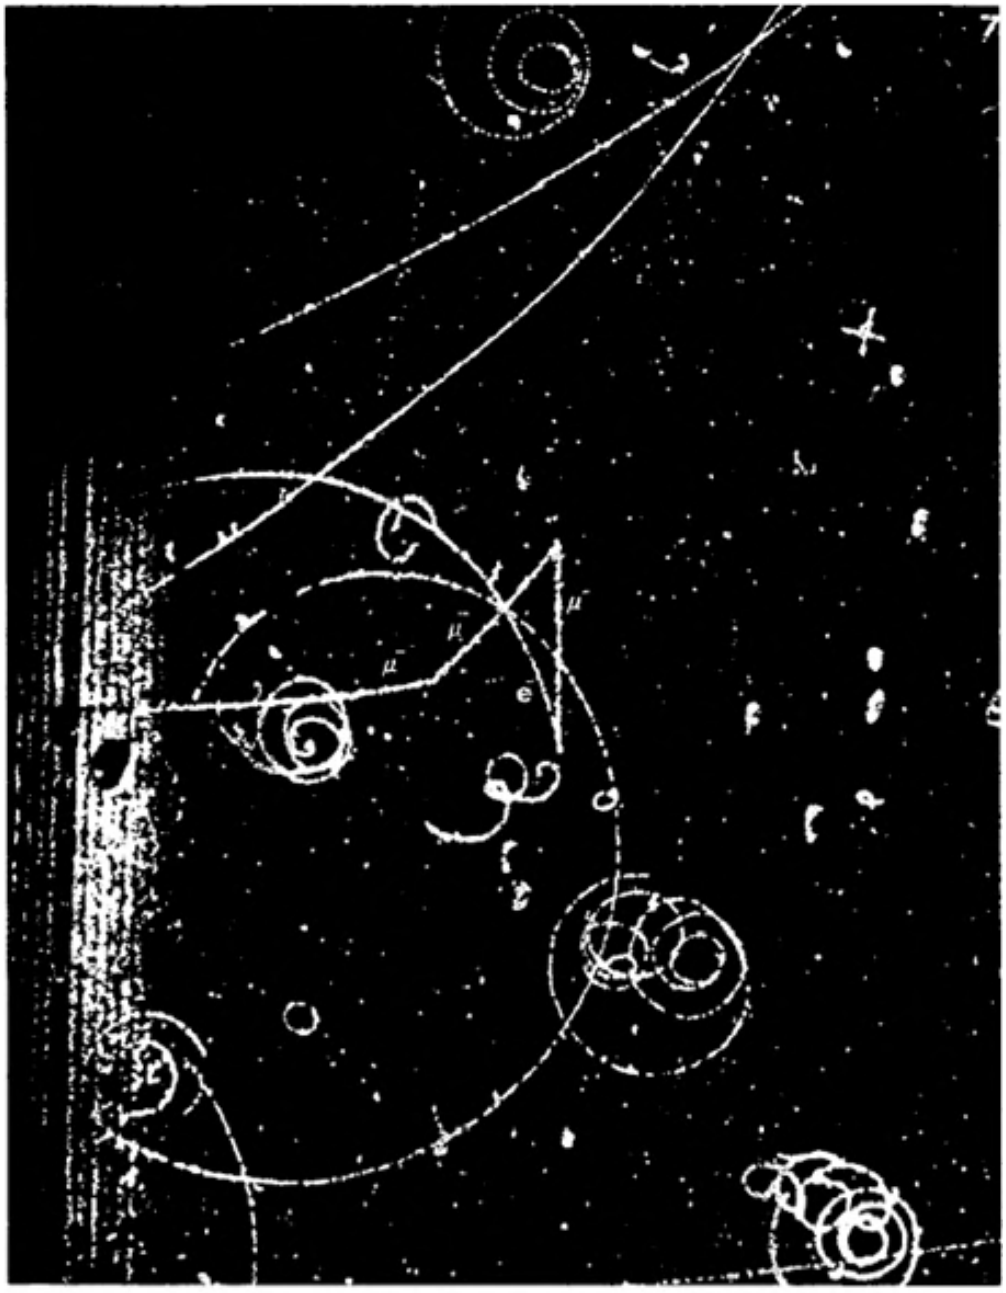

Fig. 11 . Double muon catalysis.

Figure 2.6 Bubble chamber image showing a muon-catalyzed fusion chain observed by Alvarez and collaborators. A single muon is released following the first fusion event and subsequently participates in a second fusion reaction, demonstrating the catalytic nature of the process [34].

ing may represent a physically achievable mechanism for increasing low-energy fusion probabilities through modified quantum environments.

2.5 Zero-Point Energy and Quantum Effects

Zero-Point Energy (ZPE) is the minimum energy of a quantum system, existing even at absolute zero due to quantum fluctuations. Unlike classical systems, which may remain completely at rest when cooled to zero temperature, quantum systems retain residual motion and energy. One potential argument for this origin is due to a consequence of the Heisenberg uncertainty principle. The Heisenberg uncertainty principle is

$$\Delta x \Delta p \geq \frac{\hbar}{2}. \tag{2.14}$$

If a particle is stopped ($p = \Delta p = 0$), its position would be perfectly known, violating the Heisenberg uncertainty relation in equation 2.14. Therefore, a particle that is "stopped" in the classical sense still possesses unavoidable quantum fluctuations in momentum and position, ensuring compliance with the uncertainty relation.

A variety of physical phenomena have been attributed to zero-point fluctuations. Some of the most well-known examples include the Lamb shift in atomic spectra, Hawking radiation near black hole horizons, and the Casimir force between closely spaced conducting plates. In each case, the vacuum is not truly empty, but contains fluctuating electromagnetic fields that influence the behavior of matter.

2.5.1 Zero-Point Energy in Confined Quantum Systems

In free space, the zero-point electromagnetic field is made of contributions from a continuous spectrum of allowed field modes. However, when conducting boundaries are introduced, the

boundary conditions restrict which modes may exist. This confinement alters the local vacuum energy density, producing measurable changes in the surrounding forces and interactions.

The most direct example of this effect is the Casimir force. When two parallel conducting plates are separated by a small distance, only certain electromagnetic wavelengths or modes can exist between the plates, while the vacuum outside remains unrestricted. This difference in allowed vacuum fluctuations leads to a lower zero-point energy density between the plates than outside, producing a net pressure that pushes the plates together. The Casimir effect demonstrates that the quantum vacuum is sensitive to geometry and boundary conditions, and that confined systems may experience significantly modified zero-point behavior.

This Casimir force provides direct evidence that closely spaced boundaries modify and partially suppress the vacuum electromagnetic modes comprising the zero-point field. In quantum-confined cavities, this altered vacuum structure can influence atomic and molecular behavior, motivating consideration of how reduced zero-point fluctuations may extend to and affect nuclear interactions within such environments.

2.5.2 Electron Orbital Adjustments Due to Zero-Point Energy

In modern atomic theory, electrons do not occupy classical static orbits around the nucleus, but instead reside in quantized bound states described by a wavefunction. The emergence of quantization is fundamentally tied to quantum mechanics through Planck's constant $\hbar$, which sets the natural scale of atomic structure. In hydrogen, the electron is governed primarily by the Coulomb attraction to the nucleus, yet purely classical motion would predict radiative collapse.

Electron Orbital Stability

A classically orbiting charge undergoes centripetal acceleration and continuously radiates energy according to the Larmor formula,

$$P = \frac{q^2 a^2}{6\pi\varepsilon_0 c^3} \tag{2.15}$$

implying a steady loss of orbital energy and an inward spiral toward the nucleus. In contrast, quantum mechanics prevents this collapse through wavefunction confinement and the uncertainty principle, which together establish a finite ground-state orbital size. More generally, the resulting atomic structure can be understood as a balance between electromagnetic attraction or binding and quantum kinetic pressure from quantum fluctuations.

Vacuum electromagnetic fields introduce fluctuations that can perturb bound electronic states. Within quantum electrodynamics, these fluctuations give rise to measurable effects such as energy level shifts and spontaneous emission. If the vacuum field is modified by boundary conditions–such as those imposed by a quantum cavity–the spectrum of allowed electromagnetic modes is altered. Consequently, the effective interaction between the electron and its surrounding field may also be modified.

From a complementary perspective, stochastic electrodynamics (SED) models the world using a classical framework that supplements Maxwell's equations with a Lorentz-invariant stochastic zero-point field. Within this framework, stable electron orbits arise from a dynamic equilibrium between radiative energy loss (discussed above) and continuous energy absorption from the zero-point field. This interpretation provides an intuitive mechanism by which changes in the vacuum field could influence orbital properties.

Orbital Modification

Within the stochastic electrodynamics framework, modifying the vacuum mode structure–such as through the suppression of electromagnetic modes in a confined cavity–alters the balance between radiation and absorption. A reduction in available zero-point modes would decrease the energy

absorbed by the electron, shifting the equilibrium toward smaller orbital radii. In this scenario, the electronic wavefunction contracts toward the nucleus, increasing the local electron density.

An increased electron density near the target deuteron enhances electronic screening, effectively reducing the Coulomb barrier between the interacting nuclei. This reduction in the Coulomb barrier increases the tunneling probability for low-energy fusion reactions. In the case of strong electron localization around the target nucleus, the effect is qualitatively similar to that observed in muon-catalyzed fusion, where a heavier lepton produces a significantly contracted bound state and correspondingly enhanced fusion rates.

Therefore, quantum-confined cavities provide a plausible mechanism by which modifications to the zero-point electromagnetic field can influence atomic-scale structure and, in turn, nuclear processes. By altering electron screening through changes in vacuum fluctuations, such environments may enable measurable enhancements in fusion probability. This framework establishes the theoretical foundation for the cavity-mediated fusion mechanisms investigated in this work.

Chapter 3

Screening Theory

Now that the mean-field approximation model of the nucleus has been constructed, a historical account of lepton-catalyzed fusion was recounted, and a brief overview of the potential effects of a cavity was given, this work will proceed to combine these aspects into various lepton-screening theories using three different mathematical frameworks. Each lepton-screening theory is then used to calculate a new effective interaction potential. Once this modified interaction potential has been formulated, a standard quantum mechanical approach (discussed in Section 2.2) is used to calculate the tunneling probability for a certain energy level. The tunneling probability of each theory can then be used to provide valuable information in predicting the fusion rates and cross-sections of deuterium-deuterium fusion at the specified energy levels.

Although this thesis does not dive into an exhaustive approach with complete mathematical rigors, like quantum many-body systems require, it uses three well-established frameworks that can model a single deuterium-deuterium nuclear interaction's probability of fusion. This single deuterium-deuterium fusion's probability can then be extrapolated to larger experimental scenarios with many deuterium-deuterium fusion attempts with relative certainty. The specific certainty and limits of these frameworks are discussed in Section 3.4.

Each of the three different mathematical frameworks used to develop the lepton-screening theories and modified effective interaction potentials will then be evaluated and compared to each other by the use of a "screening enhancement" factor. This screening enhancement is defined as

$$f(E) = \frac{\sigma_{screen}(E)}{\sigma_{bare}(E)}\,. \tag{3.1}$$

In other words, the screening enhancement factor is the ratio of the fusion cross-section in the presence of electron screening to the bare (unscreened) fusion cross-section. It quantifies how much the effective lowering of the Coulomb barrier, by way of the presence of electrons and their screening effects, increases the probability of tunneling.

An important fact to note is that this screening enhancement is a function of the incoming particle's energy. Lepton-screening has a different enhancement of fusion depending on the energy of the incoming nucleus. Although it would be nice to have this energy factor disappear in the ratio of cross-sections, it is too far engrossed in the WKB approximation and cross-section calculation to be easily removed. Therefore, the theory presented in this thesis will limit its mathematical predictions of fusion enhancement factors to that which is both experimentally attainable and contextually relevant. This means that the predictions will be for energies of 4 keV. For more details on this choice of energy, see Chapter 4.

Each cross-section used in the screening enhancement factor is calculated by using the tunneling probability of the WKB approximation (see Section 2.2) and converting it to a cross-section through a simple conversion:

$$\sigma = \frac{\pi\hbar^2}{2m}\frac{P_{tunnel}(E)}{E} \times 10^{28} \text{ Barns}\,. \tag{3.2}$$

Using this measure of comparison, the lepton-screening theories coming from some of the mathematical frameworks can be compared very quickly and precisely. In addition, screening enhancements can quickly be determined experimentally as discussed in Chapter 4. Having both the

theoretical and experimental screening enhancement factors allows for an easy comparison between the theoretical predictions and actual experimental measurements, helping to determine the validity and predictive power of the models.

3.1 Semi-Classical Approaches

The following section develops several semi-classical models used to estimate electron screening in metallic environments. In condensed matter systems, conduction electrons partially shield the Coulomb potential between charged particles. This screening modifies the effective interaction potential and can enhance nuclear reaction cross-sections by reducing the Coulomb barrier experienced by reacting nuclei.

The models considered here treat the electron population using statistical or continuum approximations rather than a fully quantum many-body description. In particular, these approaches neglect contributions from zero-point electromagnetic fields and instead focus on the collective response of electrons within the metallic lattice. As a result, the models are well suited for describing screening effects in bulk metals, where conduction electrons behave approximately as a free or weakly interacting electron gas. However, they are not designed to capture effects associated with quantum vacuum fluctuations or cavity-modified electromagnetic modes, which may become relevant in confined metallic cavities. Consequently, the results obtained in this section should be interpreted primarily as predictions for fusion processes occurring within metallic environments rather than within engineered cavity structures.

The goal of this section is to estimate the magnitude of screening that can arise from the conduction electron population of a metal. Several models of increasing physical realism are considered. First, a classical electrostatic treatment of screening using the Debye model is introduced. While originally developed for plasmas, the Debye model provides a useful baseline estimate for the

characteristic screening length produced by a mobile charge population. The analysis then proceeds to two quantum mechanical multi-body approximations: Thomas–Fermi theory and Lindhard theory. Thomas–Fermi theory models the conduction electrons as a degenerate Fermi gas responding to a slowly varying electrostatic potential, yielding a quantum mechanical screening length. Lindhard theory further refines this picture by describing the linear response of a Fermi gas to an external perturbation. In practice, Lindhard response functions are often evaluated using electronic structure information obtained from Density Functional Theory (DFT), which allows the electron density and band structure of a specific material to be incorporated into the screening calculation.

Several simplifying assumptions are made in order to connect the theoretical models to the experimental system used in this work. The experiments described in this thesis employ titanium as the host metal, and the electronic environment is therefore approximated using parameters appropriate for titanium's conduction electron population. Electrons associated with implanted deuterium atoms are neglected in the screening calculation, since their contribution to the total electron density is small compared to that of the metallic conduction electrons. Additionally, the lattice structure of the metal is treated as providing a static background potential while the conduction electrons are modeled as a mobile charge distribution that responds to perturbations produced by the reacting nuclei.

The following subsections develop each screening model in turn, beginning with the classical Debye model and progressing toward more complete quantum mechanical descriptions.

3.1.1 Debye Model

The Debye model provides one of the earliest theoretical descriptions of electrostatic screening in systems containing mobile charge carriers. Originally developed in the early twentieth century to describe plasmas and electrolyte solutions, the model explains how free charges rearrange themselves in response to an external electric potential. When a charged particle is introduced

into such a medium, surrounding charges redistribute so that the resulting electric field is reduced beyond a characteristic distance known as the Debye length. At distances larger than this screening length, the Coulomb potential of the particle is exponentially suppressed.

Debye screening is widely used in plasma physics, astrophysical plasmas, and electrolyte theory to estimate the range over which electrostatic interactions remain significant. In these systems, the charged particles are typically weakly interacting and can often be approximated as a classical gas in thermal equilibrium. Under these conditions, the collective response of the charge carriers can be described using classical statistical mechanics, leading to a simple analytic expression for the screening length.

While the Debye model provides useful physical intuition, its applicability to metallic systems is limited. In a metal, the conduction electrons form a highly degenerate Fermi gas rather than a classical thermal plasma. The electron distribution is therefore governed by Fermi–Dirac statistics rather than the Maxwell–Boltzmann distribution assumed in the Debye derivation. Additionally, the high electron density in metals leads to screening behavior that differs quantitatively from the classical plasma case. For these reasons, the Debye model should be interpreted here primarily as a baseline estimate for screening behavior. More accurate treatments of electron screening in metals are provided by the Thomas–Fermi and Lindhard theories discussed in subsequent sections.

Mathematical Framework

The Debye length is defined by

$$\lambda_d = \sqrt{\frac{\varepsilon_0 k_B T}{n_e e^2}}\,, \tag{3.3}$$

where ε_0 is the vacuum permittivity,

$$\varepsilon_0 = 8.854 \times 10^{-12}\ \frac{\mathrm{F}}{\mathrm{m}}\,, \tag{3.4}$$

k_B is the Boltzmann constant,

$$k_B = 1.38 \times 10^{-23} \frac{\text{J}}{\text{K}} , \tag{3.5}$$

T is the temperature (assumed here to be standard temperature),

$$T = 293 \text{ K} , \tag{3.6}$$

n_e is the electron number density,

$$n_e = \frac{\rho_e}{m_e} , \tag{3.7}$$

and e denotes the elementary charge,

$$e = 1.602 \times 10^{-19} \text{ C} . \tag{3.8}$$

In metallic systems, the number density of electrons is greatly affected. Metallic bonds are often referred to as having a "sea of electrons" which move freely around the positive lattices of the metal. The fact that the electrons are free allows for many physical traits such as high electrical and thermal conductivity, material ductility, and even contributes to the luster of metallic surfaces. This sea of electrons also, as expected, contributes multiple electrons to the free electron density found in the system. As described above, the experimental method for this investigation uses titanium metal. Therefore, if standard values are taken for the density and molar mass of titanium, the titanium characteristics can be estimated as

$$\rho = 4.502 \frac{\text{g}}{\text{cm}^3} \tag{3.9}$$

and

$$M = 47.867 \ \frac{\text{g}}{\text{mol}} \ . \tag{3.10}$$

Because titanium contributes approximately two valence electrons per atom and the electron density assumes only one valence electron, the electron density estimate stated above is multiplied by a factor of two. This yields

$$n_e = 2\frac{\rho}{M}N_A \ , \tag{3.11}$$

or numerically,

$$n_e = 2\frac{4.502 \ \frac{\text{g}}{\text{cm}^3}(10^6\frac{\text{cm}^3}{\text{m}^3})}{47.867\frac{\text{g}}{\text{mol}}}(6.022 \times 10^{23} \ \frac{\text{atoms}}{\text{mol}}) \approx 1.13 \times 10^{29} \ \text{m}^{-3} \ . \tag{3.12}$$

Using this value, the Debye length becomes

$$\lambda_d \approx 3.5 \times 10^{-12} \ \text{m} \ . \tag{3.13}$$

This length represents the characteristic distance over which an electric potential is screened by the surrounding electron distribution. Mathematically, the electron screening which comes about due to the Debye length can be modeled by multiplying the Coulomb repulsion by a decaying exponential factor. This creates the modified "screened" Coulomb potential:

$$U(r) = k\frac{e^2}{r}e^{-\frac{r}{\lambda_d}} \ , \tag{3.14}$$

where

$$e^{-\frac{r}{\lambda_d}} \tag{3.15}$$

is the Debye length's decaying exponential factor added to the normal Coulomb potential.

Using this modified Coulomb potential in combination with the original nuclear well potential, a modified potential for the target nucleus can be created. In this model, the Coulomb barrier is reduced, therefore increasing the probability of tunneling through the tunneling calculations, providing a first-order estimate of screening enhancements in an electron-rich environment.

Screening Enhancement Values

Using the Debye screening length calculated above, the modified Coulomb potential given in Eq. 3.14 can be compared directly to the unscreened Coulomb potential. Figure 3.1 shows this comparison for the interaction between two deuterium nuclei. The Debye screening introduces an exponential suppression term which reduces the Coulomb potential at large distances. However, because the Debye length for the electron density of titanium is extremely small ($\lambda_d \approx 3.5 \times 10^{-12}$ m), the modification to the potential occurs primarily at distances far larger than the region where nuclear tunneling becomes significant.

At distances relevant to the nuclear barrier penetration process (typically on the order of 10^{-14}–10^{-13} m), the exponential screening factor remains very close to unity. As a result, the screened Coulomb potential is nearly identical to the unscreened Coulomb potential throughout the region that dominates the tunneling probability. This behavior can be seen in Fig. 3.1, where the screened and unscreened potentials are essentially indistinguishable on the scale relevant to nuclear interactions.

Because the tunneling probability depends exponentially on the barrier shape, only changes to the potential within the tunneling region significantly affect the fusion cross-section. Since Debye screening modifies the potential primarily at larger distances, its impact on the barrier penetration probability is minimal. Consequently, the Debye model predicts only a negligible enhancement of the deuterium–deuterium fusion cross-section in metallic systems.

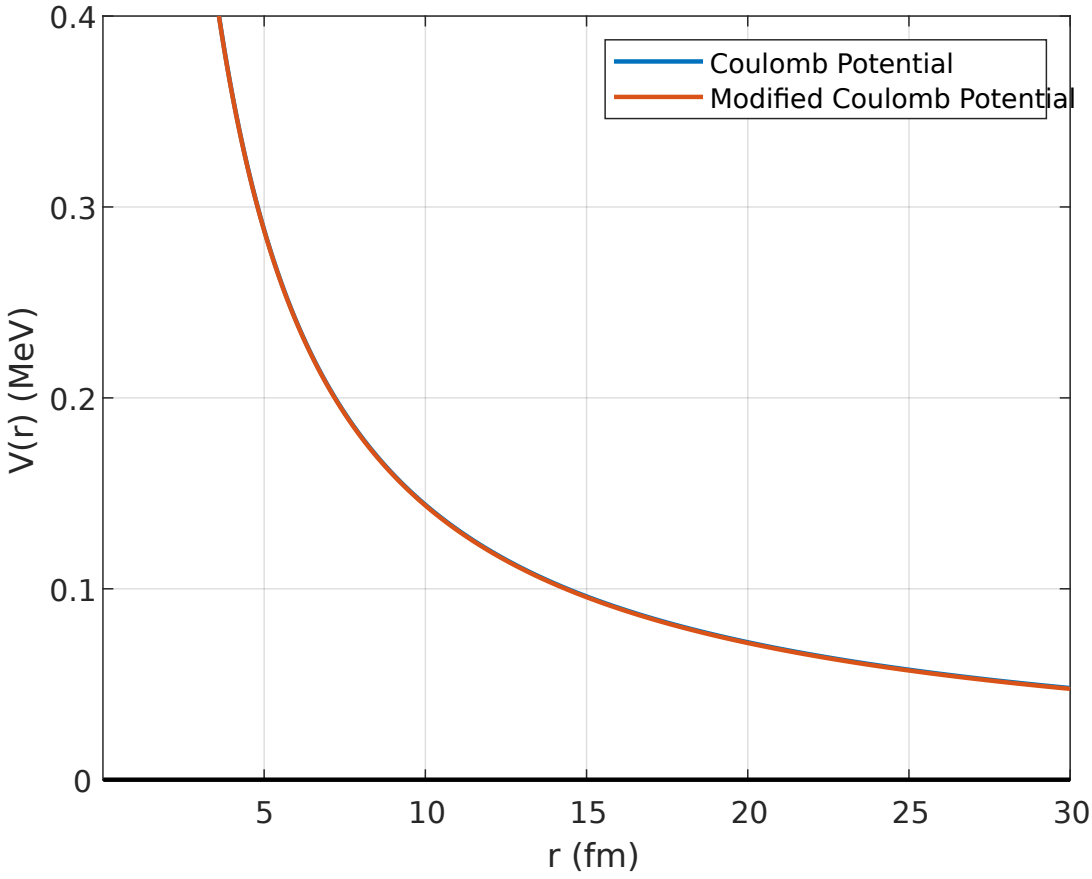

Figure 3.1 Comparison of the bare Coulomb potential and the Debye-screened Coulomb potential calculated using the electron density of titanium. Note the two potential lines overlap and are nearly indistinguishable. The Debye screening length $\lambda_d \approx 3.5 \times 10^{-12}$ m introduces an exponential suppression at large distances. However, within the short-range region that governs nuclear tunneling, the screened and unscreened potentials are nearly identical, resulting in negligible enhancement of the fusion probability.

3.1.2 Thomas-Fermi Model

While the Debye model is most appropriate for classical plasmas, metals require a different screening description because the electrons are not thermally distributed but instead form a degenerate electron gas. In this case, the screening behavior is better modeled using the Thomas-Fermi (TF) approximation. The Thomas-Fermi model treats the conduction electrons in a metal as a mean-field electron cloud which rearranges itself to screen an applied Coulomb potential. Mathematically, however, the Thomas-Fermi approximation is extremely similar to the Debye model in terms of its mathematical factors. This allows the following description of the mathematical framework to be slightly shorter than previous sections.

Mathematical Framework

Similar to the Debye model, in this approximation the screened Coulomb potential takes the form

$$U(r) = k\frac{e^2}{r}e^{-K_{tf}r}, \tag{3.16}$$

where K_{tf} is the Thomas-Fermi screening constant. This constant depends on the electron number density and is defined through the Fermi wave vector

$$K_f = \left(3\pi^2 n_e\right)^{1/3}, \tag{3.17}$$

and

$$K_{tf} = \left(\frac{e^2 m_e K_f}{\varepsilon_0 \pi^2 \hbar^2}\right)^{1/2}. \tag{3.18}$$

Using the titanium electron density calculated previously ($n_e \approx 1.13 \times 10^{29}\,\mathrm{m}^{-3}$), the Thomas-Fermi screening potential can be evaluated.

Screening Enhancement Values

The Thomas-Fermi model provides a more appropriate description of screening in metallic systems by treating the conduction electrons as a degenerate Fermi gas. Using the screening constant derived above, the modified Coulomb potential can again be compared to the unscreened Coulomb interaction between two deuterium nuclei.

Figure 3.2 illustrates the resulting screened potential obtained from the Thomas-Fermi approximation. As in the Debye case, the screening introduces an exponential suppression of the Coulomb potential at large distances. However, because the Thomas-Fermi screening length remains extremely small relative to the spatial scale over which nuclear tunneling occurs, the reduction of the Coulomb barrier within the tunneling region is limited.

When this modified potential is used in tunneling calculations, the resulting increase in barrier penetration probability is modest. For example, at an incident energy of approximately 4 keV, the Thomas-Fermi model predicts a fusion enhancement factor of

$$f(E) = \frac{\sigma_{\text{screen}}(E)}{\sigma_{\text{bare}}(E)} \approx 1.078. \tag{3.19}$$

This corresponds to only a small increase in the fusion cross-section relative to the unscreened case. While the Thomas-Fermi model provides a more realistic description of electron screening in metals than the classical Debye approach, the predicted enhancement remains far smaller than the factors reported in some experimental measurements of fusion in metal-loaded systems.

3.1.3 Lindhard Screening and Density Functional Theory

Although the Thomas-Fermi approximation provides a useful first estimate, it remains a simplified mean-field treatment of electron screening. More advanced approaches model screening through

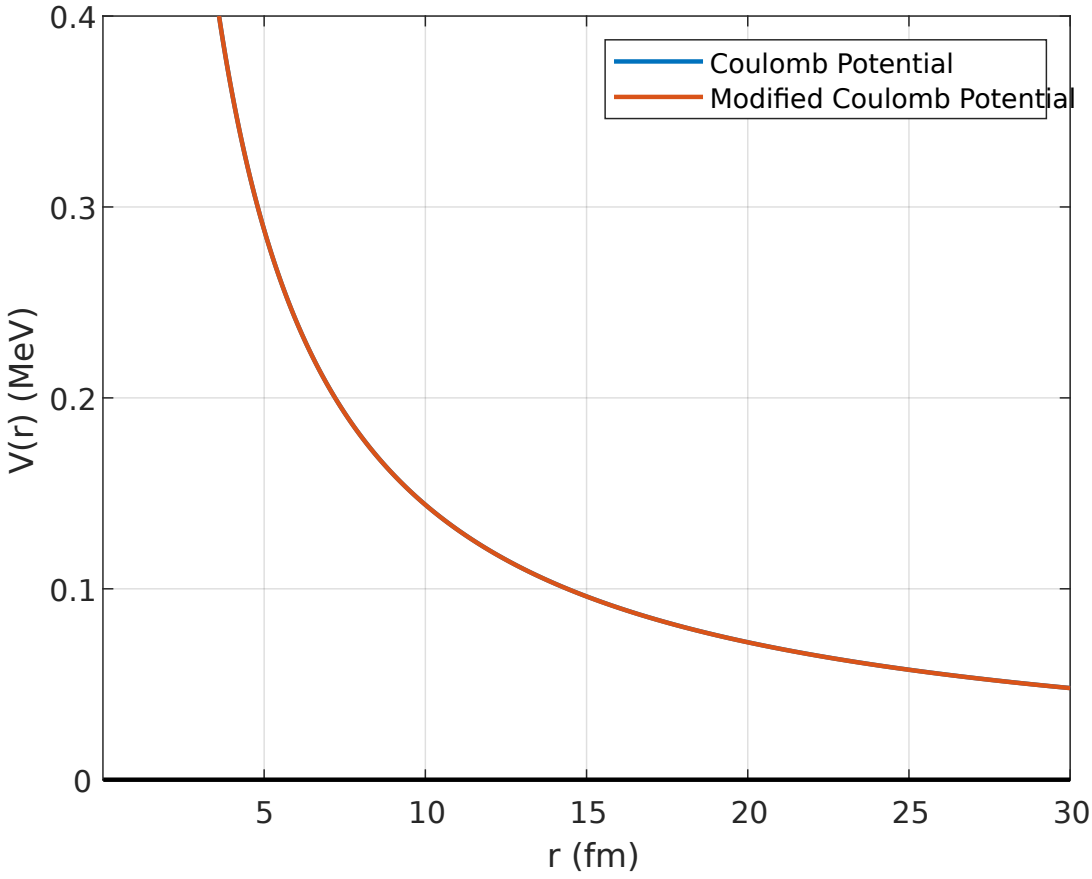

Figure 3.2 Comparison between the bare Coulomb potential and the Thomas-Fermi screened Coulomb potential calculated using the conduction electron density of titanium. Note the two potential lines overlap and are nearly indistinguishable. Although the Thomas-Fermi model incorporates the degenerate nature of the electron gas in metals, the resulting modification to the Coulomb potential remains small at the distances relevant for nuclear tunneling. This leads to only a modest increase in the predicted fusion cross-section.

quantum many-body theory, often using Density Functional Theory (DFT) or the Lindhard dielectric response of a degenerate electron gas.

In these models, the screened Coulomb potential retains the same general Coulomb potential multiplied by a decaying exponential term similar to the Thomas-Fermi model. Mathematically it is presented as

$$U(r) = k\frac{e^2}{r}e^{-K_{ftl}r}, \tag{3.20}$$

where the modified screening constant K_{ftl} is derived from the Lindhard response function. The Fermi wave vector is again defined as

$$K_f = \left(3\pi^2 n_e\right)^{1/3}, \tag{3.21}$$

and the Fermi energy is given by

$$E_f = \frac{\hbar^2 K_f^2}{2m_e}. \tag{3.22}$$

Using these quantities, the Lindhard screening constant is expressed as

$$K_{ftl} = \left(\frac{6\pi n e^2}{\varepsilon_0 E_f}\right)^{1/2}. \tag{3.23}$$

Density Functional Theory (DFT) provides a more realistic description of the electronic environment within a solid by explicitly accounting for the quantum mechanical behavior of electrons in the presence of a periodic lattice. Rather than treating the conduction electrons as a uniform free-electron gas, DFT determines the ground-state electron density of a material by solving the many-electron Schrödinger equation in an approximate form. In practice, the many-body problem is recast using the Kohn–Sham formalism, in which interacting electrons are replaced by an equivalent system of non-interacting particles moving in an effective potential that incorporates

electron–electron interactions through exchange and correlation functionals. This approach allows the spatial distribution of electron density, the electronic band structure, and the Fermi energy of a material to be calculated with considerably greater realism than simple analytic models.

For screening calculations, the electron density obtained from DFT is particularly important because the strength of electrostatic screening depends directly on the number of electrons available to respond to an external perturbation. In a real metal, the electron density is not perfectly uniform but varies depending on local bonding environments, crystal structure, and the presence of impurities or interstitial atoms such as implanted deuterium. As a result, DFT calculations typically predict a range of effective electron densities depending on the assumptions used in the electronic structure model and the specific local environment considered.

For titanium, DFT-based estimates of the conduction electron density have been reported to fall within the approximate range

$$n_{\text{small}} = 4.06 \times 10^{29}\,\text{m}^{-3} \tag{3.24}$$

and

$$n_{\text{large}} = 1.91 \times 10^{30}\,\text{m}^{-3}. \tag{3.25}$$

In the present analysis, the larger of these two values is adopted. This choice is made intentionally in order to estimate an upper bound on the possible magnitude of electron screening within the metallic system. Because the screening strength increases with electron density, using the larger value provides a conservative test of whether conventional electron screening mechanisms could plausibly account for large enhancements in fusion probability. If the predicted enhancement remains small even under this intentionally optimistic assumption, then more moderate electron densities would necessarily produce even weaker screening effects.

These increased electron densities lead to stronger screening than the Thomas-Fermi estimate derived earlier. However, as shown in the following subsection, even the enhanced screening predicted using DFT-derived densities remains insufficient to produce the large fusion-rate increases reported in some metal-confined experiments.

Screening Enhancement Factors

The Lindhard formalism provides a more complete quantum-mechanical description of screening by treating the response of the electron gas to an external perturbation using linear response theory. When combined with Density Functional Theory calculations, the model allows material-specific electron densities and band structure effects to be incorporated into the screening calculation.

For titanium, DFT studies predict conduction electron densities that can exceed the simple free-electron estimate used previously. To explore the maximum plausible screening effect, the higher density estimate is adopted in the present analysis. This increases the screening constant and therefore produces a stronger suppression of the Coulomb potential compared to the Thomas-Fermi approximation.

Figure 3.3 shows the resulting screened Coulomb potential obtained using the Lindhard response function with DFT-derived electron densities. Although the potential reduction is somewhat larger than that predicted by the Thomas-Fermi model, the modification still occurs primarily at distances larger than the nuclear tunneling region. Consequently, the reduction of the Coulomb barrier experienced by the reacting nuclei remains limited.

Using this screened potential in tunneling calculations yields a fusion enhancement factor of approximately

$$f(E) = \frac{\sigma_{\text{screen}}(E)}{\sigma_{\text{bare}}(E)} \approx 1.511 \tag{3.26}$$

for an incident energy of 4 keV. While this enhancement is noticeably larger than the Thomas-Fermi prediction, it remains far below the screening strengths required to explain the significantly larger fusion rate enhancements reported in some metal-confined experiments.

These results suggest that conventional electron screening mechanisms alone are insufficient to account for large fusion-rate increases in metallic environments. This limitation motivates the exploration of additional physical mechanisms that may modify the effective interaction potential between reacting nuclei, particularly in confined geometries such as the cavity structures investigated in this work.

3.2 Stochastic Electrodynamics Approach

Stochastic Electrodynamics (SED) is a semiclassical framework in which matter is treated using electrodynamics, while the electromagnetic vacuum is modeled as a real classical background field with a prescribed zero-point spectral energy density. Although SED has been cited by some authors as incomplete or fundamentally inconsistent when extended to fully relativistic quantum phenomena, it nevertheless reproduces a number of quantum-like predictions in restricted systems, particularly for bound harmonic systems interacting with the electromagnetic vacuum.

In the present work, SED is not used as a replacement for quantum electrodynamics, but rather as a model that provides intuition for how modifications to the vacuum mode structure, such as those induced by a Casimir cavity, may perturb the effective ground-state electron distribution. In particular, the SED balance between classical radiation damping and absorption from the zero-point field provides a useful conceptual pathway for estimating how confinement of vacuum modes can alter atomic stability conditions.

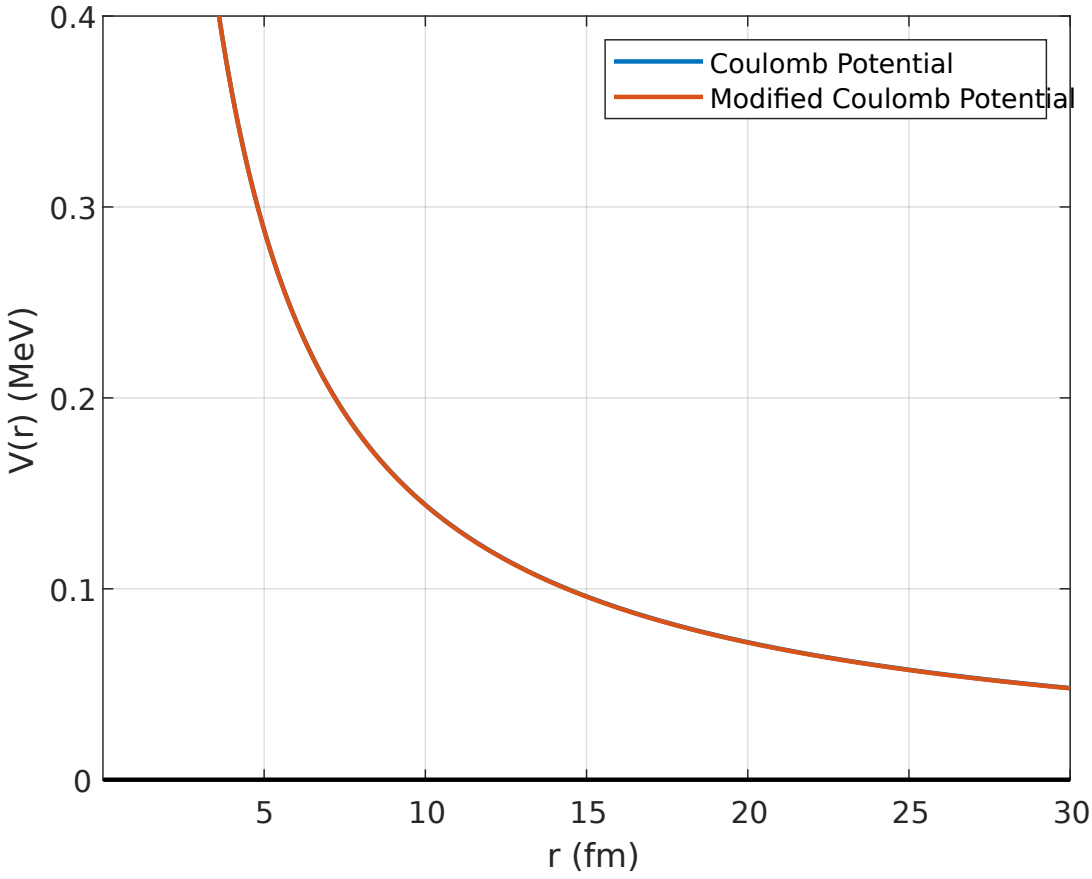

Figure 3.3 Bare Coulomb potential compared with the screened potential obtained using the Lindhard dielectric response with electron densities derived from Density Functional Theory estimates for titanium. Note the two potential lines overlap and are nearly indistinguishable. Although the higher electron density increases the strength of screening relative to the Thomas-Fermi approximation, the resulting reduction of the Coulomb barrier remains modest at the nuclear interaction scale.

3.2.1 Ground State of Hydrogen Using SED

Within the SED approach, the electron in the hydrogen ground state may be modeled as a bound charge interacting with the classical Coulomb potential while simultaneously coupled to the electromagnetic zero-point field (ZPF). The central assumption is that the electron radiates energy classically due to its acceleration, but that this radiated power is exactly balanced, on average, by energy absorbed from the surrounding ZPF. This balance condition yields a steady-state orbit which corresponds to the quantum ground state.

The average absorbed power from the vacuum field may be written in the form

$$\left\langle P^{\text{abs}} \right\rangle = \frac{e^2\hbar}{6\pi^2\varepsilon_0 mc^3}\int_0^\infty \frac{\Gamma\omega^7\, d\omega}{(\omega_0^2-\omega^2)^2+\Gamma^2\omega^6}, \tag{3.27}$$

where ω_0 is the characteristic orbital frequency and Γ is the radiative damping constant associated with the electron's coupling to the field.

For an electron, the quantity Γ is extremely small, and thus the integrand in Eq. (3.27) is sharply peaked around $\omega = \omega_0$. One may therefore use the standard resonance approximation, extending the limits of integration and replacing ω by ω_0 in all slowly varying terms. This yields the Lorentzian form

$$\left\langle P^{\text{abs}} \right\rangle = \frac{e^2\hbar\omega_0^3}{12\pi\varepsilon_0 mc^3}\int_{-\infty}^\infty \frac{1}{\pi}\frac{\gamma\, d\omega}{(\omega-\omega_0)^2+\gamma^2}, \tag{3.28}$$

where $\gamma = \Gamma\omega_0^2/2$.

The Lorentzian factor is normalized such that

$$\int_{-\infty}^\infty \frac{1}{\pi}\frac{\gamma\, d\omega}{(\omega-\omega_0)^2+\gamma^2} = 1, \tag{3.29}$$

and therefore the absorbed power reduces to

$$\left\langle P^{\text{abs}} \right\rangle_{\text{free}} = \frac{e^2\hbar\omega_0^3}{12\pi\varepsilon_0 mc^3}. \tag{3.30}$$

In the original SED treatment of the hydrogen ground state, this absorbed power is balanced against the classical Larmor radiated power, yielding an equilibrium condition that reproduces an orbital scale on the order of the Bohr radius.

3.2.2 Theoretical New Ground State Within Cavities

The previous derivation assumed that the electromagnetic vacuum supports a continuum of field modes at all frequencies. However, within a Casimir cavity the mode spectrum is modified by boundary conditions, and the ZPF is no longer characterized by a fully continuous distribution. For the case of two infinite parallel conducting plates separated by a distance L, the allowed longitudinal wave vectors are discretized according to

$$k_z = \frac{n\pi}{L}, \qquad n = 1,2,3,\ldots \tag{3.31}$$

so that the cavity mode frequencies become

$$\omega_n(k_\perp) = c\sqrt{k_\perp^2 + \left(\frac{n\pi}{L}\right)^2}. \tag{3.32}$$

Thus, the free-space integral over vacuum frequencies must be replaced by a discrete sum over cavity modes. In particular,

$$\int_0^\infty d\omega \quad \longrightarrow \quad \sum_{n=1}^{\infty} \int d^2k_\perp. \tag{3.33}$$

The absorbed power inside the cavity may therefore be written as

$$\left\langle P^{\text{abs}} \right\rangle_{\text{cav}} = \frac{e^2\hbar}{6\pi^2\varepsilon_0 mc^3} \sum_{n=1}^{\infty} \int d^2k_\perp \, \frac{\Gamma\,\omega_n^7}{(\omega_0^2-\omega_n^2)^2+\Gamma^2\omega_n^6}. \tag{3.34}$$

Under the resonance approximation, the cavity-modified absorption takes the form

$$\left\langle P^{\text{abs}} \right\rangle_{\text{cav}} = \left\langle P^{\text{abs}} \right\rangle_{\text{free}} F_{\text{cav}}(\omega_0), \tag{3.35}$$

where the cavity factor is given by

$$F_{\text{cav}}(\omega_0) = \sum_n \frac{1}{\pi} \frac{\gamma}{(\omega_n-\omega_0)^2+\gamma^2} \tag{3.36}$$

which is basically seen as a more complicated way to state

$$F_{cav}(\omega_0) = \frac{\rho_{cav}(\omega_0)}{\rho_{free}(\omega_0} \, . \tag{3.37}$$

Unlike the free-space Lorentzian normalization, the discrete cavity spectrum implies that $F_{\text{cav}}(\omega_0)$ does not generally evaluate to unity. Physically, this reflects the fact that the cavity suppresses absorption whenever no electromagnetic mode exists sufficiently close to the resonance frequency ω_0.

For a cavity separation L, the lowest allowed frequency is

$$\omega_{\text{min}} = \frac{\pi c}{L}. \tag{3.38}$$

To determine when this suppression becomes significant, the cavity cutoff frequency can be compared to a characteristic atomic frequency scale. For hydrogen, the ground-state energy of 13.6 eV corresponds to an angular frequency given by

$$\omega_0 \sim \frac{E}{\hbar} \approx 2.066 \times 10^{16} \text{ rad/s}. \tag{3.39}$$

Requiring that the cavity cutoff frequency exceeds this characteristic frequency, $\omega_{\text{min}} \geq \omega_0$, gives a condition on the cavity size:

$$L \leq \frac{\pi c}{\omega_0} \approx 45.6 \text{ nm}. \tag{3.40}$$

This result indicates that, for cavity-induced modifications to significantly affect the electron orbital structure, the cavity separation must be below approximately 50 nm. For larger cavity sizes, the electromagnetic mode spectrum remains sufficiently dense near ω_0, and the electron experiences an environment similar to free space. As the cavity dimension approaches and falls below this threshold, the suppression of available modes near the atomic frequency becomes more pronounced, leading to a measurable reduction in the effective zero-point field density and, consequently, a modification of the electron orbital radius.

3.2.3 Modification to Screening

The relevance of cavity-induced modifications becomes particularly clear when considering their effect on electron screening in comparison to muon-catalyzed fusion. In muonic hydrogen, the electron is replaced by a muon with a mass approximately 207 times larger, resulting in a corresponding reduction in the orbital radius and a significantly enhanced screening of the nuclear charge. This reduction allows the nuclei to approach much more closely, dramatically increasing the probability of fusion.

Within the cavity framework developed in this work, a similar effect may arise through modification of the electromagnetic environment rather than a change in particle mass. As shown in the previous section, the suppression of vacuum modes alters the effective zero-point field experienced by the electron, which in turn can modify the equilibrium orbital radius.

To estimate the scale at which this effect becomes comparable to the muonic case, a characteristic frequency associated with the muonic hydrogen ground state can be considered. Approximating the ground-state energy as $E \sim 13.6 \times 207 \approx 2815$ eV, the corresponding angular frequency is

$$\omega_\mu \sim \frac{E}{\hbar} \approx 4.27 \times 10^{18} \text{ rad/s}. \tag{3.41}$$

Applying the cavity cutoff condition,

$$\omega_{\min} = \frac{\pi c}{L}, \tag{3.42}$$

and requiring $\omega_{\min} \geq \omega_\mu$, the corresponding cavity size becomes

$$L \leq \frac{\pi c}{\omega_\mu} \approx 2.2 \text{ Å}. \tag{3.43}$$

This result suggests that, in order for cavity-induced modifications to produce electron orbitals comparable in scale to those in muonic hydrogen, the confinement length must approach the angstrom scale. At these dimensions, the suppression of electromagnetic modes at frequencies

relevant to tightly bound electronic states becomes significant enough to substantially reduce the effective orbital radius.

The physical consequence of this reduction is an increase in the electron density near the nucleus, leading to enhanced screening of the Coulomb repulsion between reacting nuclei. In this sense, sufficiently confined cavities may reproduce, at least partially, the screening enhancement mechanism observed in muon-catalyzed fusion, but without requiring the presence of heavy leptons.

An important implication of this result is that such length scales are not purely theoretical. Real metallic surfaces are known to contain defects, grain boundaries, and nanoscale cavities with characteristic dimensions on the order of angstroms. These naturally occurring structures may therefore act as localized regions of strong electromagnetic confinement, potentially modifying electron screening in a manner consistent with the mechanism described here. This provides a possible explanation for experimentally observed enhancements in fusion rates within metallic systems.

3.3 Quantum Electrodynamics Approach

Unlike stochastic electrodynamics, quantum electrodynamics (QED) provides a fully relativistic and internally consistent framework in which both matter and the electromagnetic field are quantized. In QED, the vacuum is not interpreted as a classical stochastic background, but rather as the ground state of the quantized photon field. Nevertheless, the vacuum exhibits measurable fluctuations, and these fluctuations produce observable radiative corrections to atomic structure.

For hydrogenic systems, the most well-known manifestation of vacuum fluctuations is the Lamb shift, which arises primarily from the electron self-energy and vacuum polarization. These corrections are fundamentally tied to the interaction of the bound electron with virtual photons and the altered structure of the electromagnetic vacuum. Importantly, because vacuum fluctuations

depend on the available field modes, boundary conditions imposed by Casimir cavities may modify the radiative corrections and thereby perturb atomic energy levels and electron distributions.

In this section, we outline how cavity-induced mode suppression enters naturally within QED through modifications of the photon propagator and the local density of vacuum states.

3.3.1 Radiative Corrections and the Hydrogen Ground State

In the absence of boundaries, the hydrogen atom is described at leading order by the Coulomb Hamiltonian,

$$H_0 = \frac{\widehat{p}^2}{2m} - \frac{e^2}{4\pi\varepsilon_0 r}, \tag{3.44}$$

with the ground-state wavefunction characterized by the Bohr radius a_0. QED corrections enter perturbatively through the interaction Hamiltonian between the electron and the quantized electromagnetic field,

$$H_{\text{int}} = -\frac{e}{m}\widehat{\mathbf{p}}\cdot\widehat{\mathbf{A}}(\mathbf{r}) + \frac{e^2}{2m}\widehat{\mathbf{A}}^2(\mathbf{r}), \tag{3.45}$$

where $\widehat{\mathbf{A}}$ is the quantized vector potential operator.

The leading radiative shift of the hydrogen ground state may be written schematically as a second-order perturbative correction,

$$\Delta E_{1s} = \sum_{\lambda}\int d^3k\,\frac{|\langle\lambda;\mathbf{k}|\,H_{\text{int}}\,|1s\rangle|^2}{E_{1s} - E_\lambda - \hbar\omega_k}, \tag{3.46}$$

where λ labels intermediate atomic states and $\omega_k = c|\mathbf{k}|$ is the photon frequency. This shift is finite only after renormalization and constitutes the self-energy contribution to the Lamb shift.

A complementary contribution arises from vacuum polarization, in which the Coulomb potential is modified by virtual electron–positron fluctuations. This yields the Uehling correction,

$$V(r) \to V(r) + V_{\text{U}}(r), \tag{3.47}$$

which perturbs the electron density primarily near the origin. Both effects reflect the fact that the hydrogen ground state is not determined solely by the Coulomb field, but also by its coupling to the quantized vacuum.

3.3.2 Cavity Modification of the Vacuum Field

The essential difference between free space and a Casimir cavity is that the photon vacuum is altered by electromagnetic boundary conditions. Between two parallel conducting plates separated by distance L, the allowed photon modes satisfy

$$k_z = \frac{n\pi}{L}, \qquad n = 1,2,3,\ldots \tag{3.48}$$

so that the photon dispersion relation becomes

$$\omega_n(k_\perp) = c\sqrt{k_\perp^2 + \left(\frac{n\pi}{L}\right)^2}. \tag{3.49}$$

In QED, radiative corrections depend explicitly on the photon propagator,

$$D_{\mu\nu}(x,x') = \langle 0|\, T\left\{A_\mu(x)A_\nu(x')\right\}|0\rangle, \tag{3.50}$$

which encodes the vacuum correlations of the electromagnetic field. Inside a cavity, this propagator is replaced by a boundary-dependent propagator,

$$D^{\text{cav}}_{\mu\nu}(x,x') \neq D^{\text{free}}_{\mu\nu}(x,x'), \tag{3.51}$$

implying that the spectrum of vacuum fluctuations sampled by the atom is modified.

Equivalently, one may express the effect in terms of the local density of photon states (LDOS). Radiative shifts are proportional to the availability of vacuum modes at the relevant atomic transition frequency. Thus, the free-space continuum integral,

$$\int d^3k, \tag{3.52}$$

is replaced by a cavity-modified sum-integral,

$$\int d^3k \quad \longrightarrow \quad \sum_{n=1}^{\infty} \int d^2k_{\perp}. \tag{3.53}$$

However, because the transverse wave vector $k_{\perp}$ remains continuous, the photon spectrum is not fully discretized. For cavity separations L larger than the atomic length scale, the LDOS near optical and ultraviolet frequencies remains close to its free-space value. This is consistent with the earlier result that significant modification requires cavity dimensions below ~ 50 nm.

3.3.3 Implications for Ground-State Screening

The relevance of cavity-modified QED fluctuations becomes apparent when considering screening and dielectric response models. In electronic structure theory, screening is often treated through response functions such as the Lindhard dielectric function, which implicitly assume a vacuum field structure consistent with free space.

In confined geometries, the modification of vacuum correlations implies that the effective electromagnetic response at short length scales may be perturbed. In QED terms, the screening potential becomes boundary dependent through vacuum polarization corrections,

$$V_{\text{eff}}(r;L) = V(r) + V_{\text{U}}(r;L), \tag{3.54}$$

where the cavity alters the virtual photon exchange responsible for polarization effects.

Although such corrections are extremely small for cavities on the order of 100 nm, the formal structure demonstrates that Casimir confinement enters atomic physics not through classical power balance, but through boundary-dependent modifications of the photon propagator and the associated radiative corrections. This provides the fully quantum mechanical analogue to the SED intuition developed in the previous section.

3.4 Limitations of the Mathematical Models

The screening models developed in this chapter provide useful first-order estimates of how conduction electrons may modify the effective Coulomb interaction between reacting nuclei. However, several important physical limitations must be acknowledged when interpreting the results. These models rely on simplifying assumptions that allow the problem to be treated analytically, but these assumptions necessarily neglect a number of effects that may become important in real condensed-matter systems.

3.4.1 Limitations of the Screening Approximations

One significant limitation is that the models treat the electron population as a static or quasi-static charge distribution characterized only by an average density. In reality, when a positively charged nucleus approaches another nucleus within a metal lattice, the surrounding electron cloud does not remain fixed. Instead, the electrons dynamically rearrange themselves in response to the changing electric field produced by the approaching charges. This dynamic response can involve collective oscillations of the electron gas, local polarization effects, and transient distortions of the electronic wavefunctions. The Debye, Thomas–Fermi, and Lindhard models approximate this behavior using simplified screening lengths or linear-response functions, but they do not fully capture the time-dependent motion of electrons during the collision process.

A related limitation arises from the mean-field nature of the screening approximations. In these treatments, the many-electron system is replaced by an averaged background electron density that produces a smooth electrostatic potential. While this approximation greatly simplifies the calculation, it neglects electron-electron correlations and local fluctuations in electron density that occur on atomic length scales. In real metallic systems, the electron distribution is strongly influenced by the crystal lattice, band structure, and the presence of defects or interstitial atoms such

as absorbed deuterium. These microscopic variations can produce local screening environments that differ from the uniform electron gas assumed in the present models.

3.4.2 Local Lattice and Structural Effects

Another important limitation is that the calculations treat the interacting nuclei as isolated particles embedded in a bulk metallic environment. In practice, deuterium atoms dissolved in metals occupy specific lattice sites and may form complex local configurations, including clusters, vacancies, or interstitial states. The electronic structure near these sites can differ substantially from that of the bulk metal, altering both the local electron density and the effective screening potential experienced by nearby nuclei. Because the models used in this chapter rely on bulk electron density estimates for titanium, they do not explicitly account for these local structural effects.

3.4.3 Approximations in the Tunneling Model

The treatment of tunneling probabilities also introduces approximations. The cross-sections in this work are calculated using the WKB approximation applied to an effective one-dimensional potential barrier. Although the WKB method provides a convenient analytical approach to barrier penetration problems, it assumes a smooth potential and neglects multidimensional effects that may occur in realistic nuclear collisions within a lattice. In addition, the nuclear interaction potential itself is represented using simplified forms that do not incorporate the full complexity of nuclear forces at very short distances.

3.4.4 Effects Beyond Bulk Electromagnetic Models

Finally, the screening models presented here do not incorporate possible modifications arising from confined electromagnetic environments such as nanoscale cavities. In particular, the Debye,

Thomas–Fermi, and Lindhard theories assume a conventional bulk electromagnetic environment and therefore neglect any changes in vacuum mode structure or electromagnetic field correlations that could occur in strongly confined geometries. If cavity-induced modifications to vacuum fluctuations alter the local electronic structure or dielectric response of the material, such effects would not be captured by the semi-classical screening models considered here.

For these reasons, the screening enhancements calculated in this chapter should be interpreted primarily as baseline predictions based on conventional electron screening mechanisms in bulk metallic systems. Any experimentally observed fusion enhancements that significantly exceed these predictions may indicate the presence of additional physical mechanisms not included in the simplified models developed here. This possibility motivates the exploration of alternative theoretical approaches discussed later in this work.

3.5 Brief Comparison of Approaches

The stochastic electrodynamics and quantum electrodynamics approaches both attribute atomic stability and radiative behavior to interactions with vacuum fluctuations, but they differ fundamentally in interpretation and mathematical foundation. In SED, the electromagnetic zero-point field is treated as a real classical stochastic background, and the hydrogen ground state emerges from a balance between classical radiation losses and absorption from the vacuum field. Casimir cavities enter naturally in this picture through a restriction of available vacuum modes, leading directly to a modified absorption condition.

In contrast, QED provides a fully quantum description in which the vacuum is the ground state of the quantized photon field, and atomic structure is determined primarily by the Coulomb interaction with small radiative corrections. Boundary conditions imposed by cavities modify these corrections not through power balance, but through changes in the photon propagator and the local

density of vacuum states. For cavity separations on the order of $L \sim 100\,\text{nm}$, the relevant vacuum modes at hydrogenic frequencies remain largely accessible, implying only minor perturbations to the ground state in both frameworks.

Ultimately, SED offers an intuitive semiclassical mechanism for visualizing cavity-induced vacuum suppression, while QED supplies the rigorous and experimentally validated foundation in which such effects appear as small boundary-dependent shifts to self-energy and vacuum polarization. Together, these perspectives motivate the expectation that Casimir confinement may perturb screening behavior, even if the resulting ground-state modifications remain subtle at experimentally realistic cavity scales.

Chapter 4

Experimental Method

The primary goal of this thesis is to investigate the hypothesis that microscopic cavities on the surface of metals contribute to the enhanced fusion rates observed in metallic environments. One proposed explanation for this enhancement is that confined cavities modify the local quantum zero-point energy field, potentially reducing the effective Coulomb barrier between reacting nuclei. The preceding chapters discussed the current challenges in accurately predicting fusion rates in metallic systems and introduced several theoretical models that could account for enhanced screening if cavity-induced effects are present.

To experimentally test this hypothesis, a dedicated system was developed to create a nanometer- to angstrom-scale cavity environment in which low-energy deuterium-deuterium fusion can occur. The system must enable the introduction of deuterium into the system, sustain a plasma to create a monolayer on the surface of the cavity walls, and finally allow a controlled formation of a microscopic-scale cavity. In addition to maintaining suitable vacuum conditions, the apparatus must also allow for reliable detection of fusion-related signals produced during operation.

Before describing the experimental apparatus in detail, it is necessary to consider the possible products of a deuterium-deuterium fusion event. Identification of these reaction products enables

the determination of fusion rates through appropriate detector measurements. The relevant reaction channels are

$$\mathrm{d} + \mathrm{d} \longrightarrow \begin{cases} t + p & \text{(i)} \\ {}^3\mathrm{He} + n & \text{(ii)} \\ \alpha + \gamma & \text{(iii)} \end{cases} . \tag{4.1}$$

In these expressions, t denotes a triton (^{3}H nucleus), p denotes a proton (^{1}H nucleus), ^{3}He denotes a helium-3 nucleus, n denotes a neutron, α denotes an alpha particle (^{4}He nucleus), and γ represents a high-energy photon.

Reaction branches (i) and (ii) occur with approximately equal probability and account for essentially all deuterium-deuterium fusion events. Branch (iii), on the other hand, is strongly suppressed and occurs with extremely low frequency. As a result, this work focuses on detection of the first two reaction channels in order to determine the fusion rate within the system. Detection of the charged fusion products and neutrons associated with these channels provides the primary experimental signature of deuterium-deuterium fusion.

The remainder of this methods section describes the experimental apparatus developed to create the required cavity environment and detect these fusion products, along with the reasoning behind the major design choices of the system. The first section (Section 4.1) will discuss the methods in making a microscopic cavity. The following section (Section 4.2) then discusses the rest of the experimental system and how system has been designed to allow for fusion in the cavity. The section following the fusion environment (Section 4.3) walks through the general procedure for the experimental trials. The final section (Section 4.4) discusses the detectors that are used to measure the fusion products during the experimental trials.

4.1 Creation of the Cavity

Building the actual cavity is an extremely important component of the experimental design. The separation distance between the cavity walls must be adjustable and controlled with high precision, as the cavity is required to operate at nanometer-scale separations. This is extremely difficult and will require a specialized system that will be described in this section.

4.1.1 Cavity Walls

As expected, the cavity has two sides. One side of the cavity is stationary and doubles as the positioning of the charged particle detector. This particle detector is made of silicon and has an extremely thin layer of aluminum over the surface. The details of the charged-particle detector will be recounted in Section 4.4.2. For this section, it is sufficient to explain that this has allowed for one side of our cavity to be a silicon detector coated by aluminum. The opposing boundary consists of an adjustable silicon wafer, referred to as the 'target'.

To establish a controlled separation between the two cavity surfaces, three small aluminum spacer columns are deposited onto the silicon wafer in a tripod configuration as seen in figure 4.1. This geometry ensures that when the target surface approaches the detector, the three spacers define a stable contact plane while maintaining a uniform gap between the two surfaces. Aluminum was selected for the spacers because the detector surface is also coated with a thin aluminum layer, ensuring material compatibility and reducing the likelihood of surface damage during contact.

The spacers are fabricated using thermal evaporation through a shadow mask designed to produce the three-column pattern on the wafer surface. Depositing the three columns simultaneously ensures that each spacer has the same thickness and therefore establishes a consistent cavity separation when contact occurs. The thickness of the deposited aluminum spacers was measured using ellipsometry techniques, confirming that the spacer height is appropriate. This thickness ultimately defines the

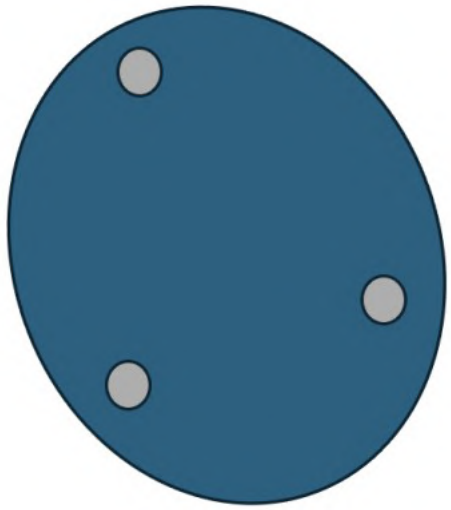

Figure 4.1 CAD model of the silicon wafer with three aluminum legs spaced in a tripod configuration for ensuring the cavity's width

minimum separation between the cavity walls when the plunger is fully advanced. For a more complete description of the deposition chamber and process, see Rhett Lundell's undergraduate thesis [35].

4.1.2 Cavity Mobility and Adjustment

During operation, the mobile wall of the cavity is formed by the silicon wafer mounted on the plunger assembly. As the plunger is advanced along the chamber axis, the wafer approaches the charged-particle detector in a face-to-face configuration. The separation between these two surfaces must ultimately be reduced to nanometer-scale distances while ensuring that the detector surface is not mechanically damaged during cavity formation. As discussed above, the controlled minimum separation between the two surfaces was accomplished through thermally-deposited aluminum columns which define the final cavity gap once contact occurs.

Additional mechanical protection is provided by a low spring-constant mount attached to the backside of the silicon wafer. The wafer is supported by this compliant spring element so that small positional errors or slight over-travel of the plunger can be absorbed elastically rather than transferring force directly to the detector. As a result, the wafer can safely approach and gently contact the detector surface through the aluminum spacers without causing mechanical damage.

A schematic representation of this configuration is shown in Fig. 4.2. The diagram illustrates the relative positions of the charged-particle detector, the cavity spacing defined by the aluminum spacers, the silicon wafer, and the spring-loaded mounting structure attached to the plunger assembly. This design allows the cavity spacing to be adjusted through controlled plunger motion while maintaining both mechanical stability and detector safety.

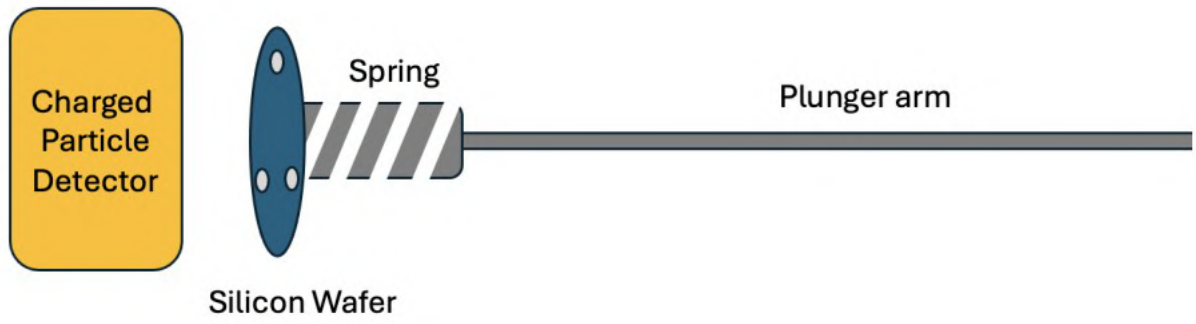

Figure 4.2 Schematic diagram of the mobile cavity wall assembly. The silicon wafer carrying three aluminum spacer columns approaches the charged-particle detector as the plunger is advanced. A low spring-constant mount on the backside of the wafer allows compliant motion, ensuring that the aluminum spacers establish the cavity gap while preventing mechanical damage to the detector surface.

4.2 Creating the Environment for Fusion

This section describes the vacuum chamber and its supporting components used to create the environment for fusion. The original vacuum system was substantially different from the one presented in this work. To see the original vacuum system, along with its issues and shortcomings, see Appendix I.

Specific details regarding the experimental procedure used to allow fusion are discussed in the following section (Section 4.3). While a detailed discussion of each component in the vacuum chamber follows, Figure 4.3 shows an overview of the completed experimental system. The following sections include various perspectives of the chamber, along with labels on each image indicating the components described in the respective section.

The central chamber is built around a six-way, 4.5-inch ConFlat flange cube. This cube serves as the primary vacuum junction and the location where the cavity is completed. The perspective shown

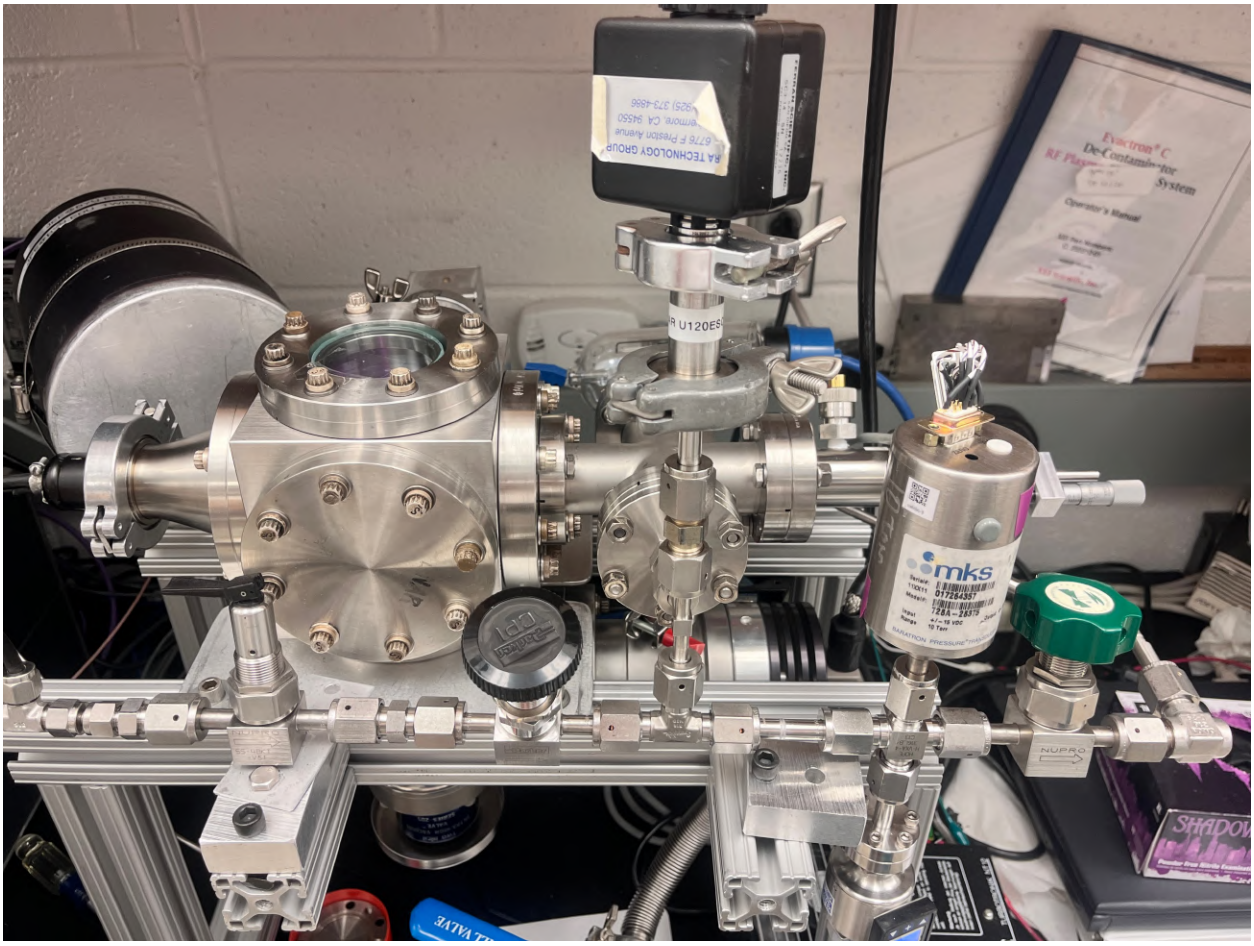

Figure 4.3 Photograph of the assembled vacuum chamber used in the experiment.

in Figure 4.3 is used to define the sides of the chamber throughout this work. The designations for the top, bottom, front, back, left, and right sides are shown in Figure 4.4.

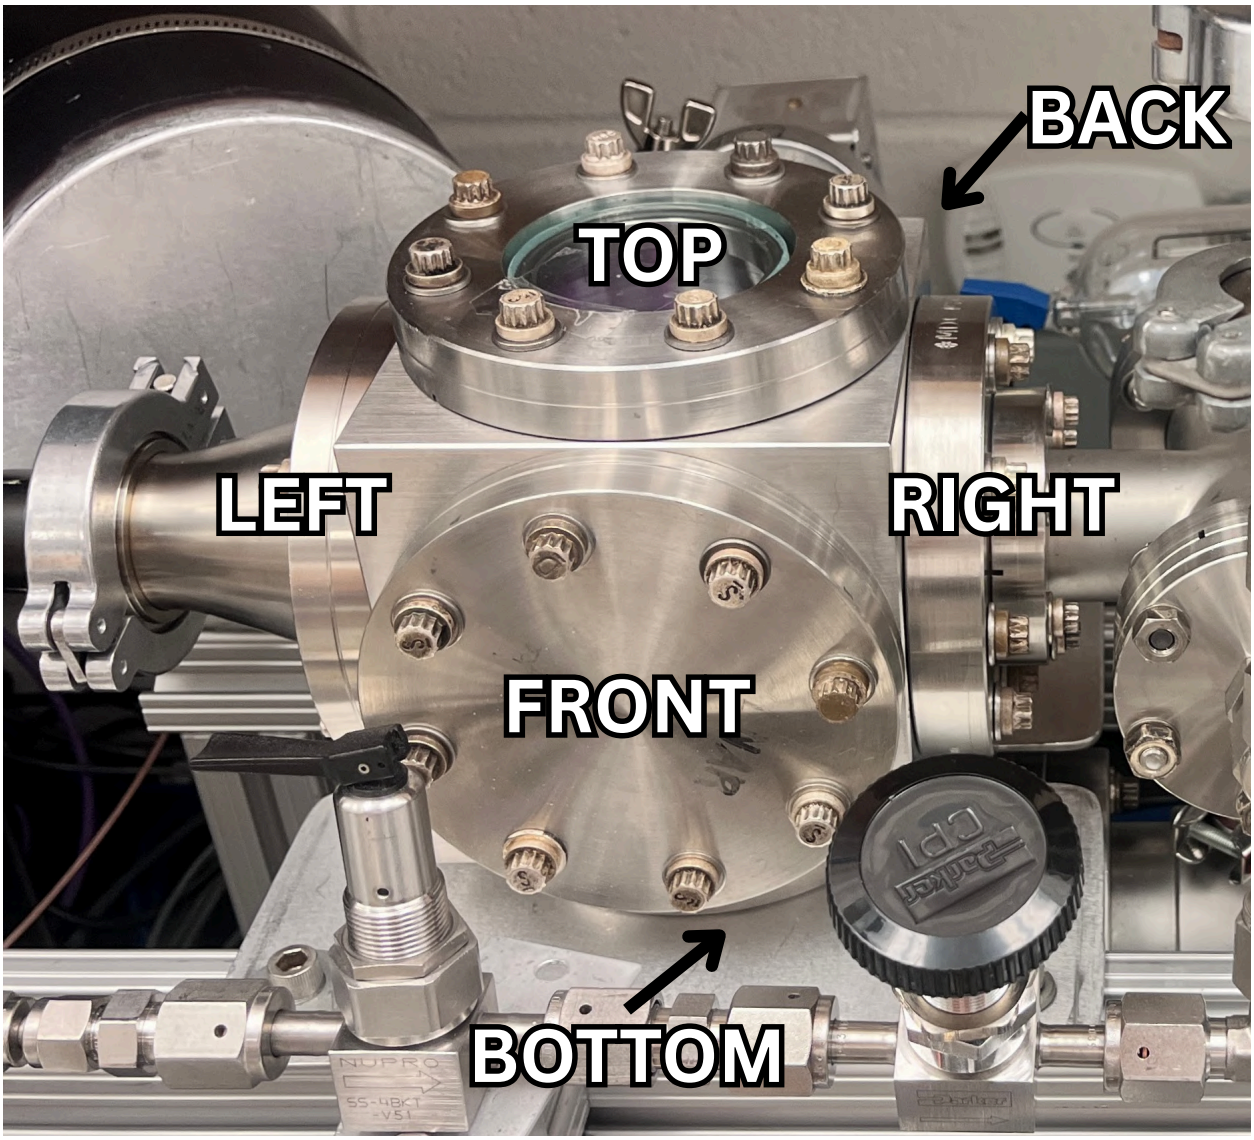

Figure 4.4 Labels for each side of the central vacuum chamber. From this perspective, the sides are designated as top, bottom, front, back, left, and right.

On the right side of the central chamber, a short ConFlat nipple is attached to connect it to a four-way, 2.75-inch ConFlat cross. This allows the system to contain more ports, a longer plunger assembly, and better diagnostics for vacuum pressure. In addition, one port is used to create a Vacuum Coupling Radiation (VCR) system. This system holds most of the gauges and a secondary outlet to the pump, allowing for more controlled vacuum moderation. This cross is discussed in greater detail in the following section.

This configuration provides a compact experimental platform in which the key functions required for the experiment—vacuum pumping, deuterium gas handling, plasma generation, cavity formation, and diagnostic access—can be integrated within a single vacuum system. The modular flange geometry also allows individual components to be installed or modified without requiring major changes to the overall chamber design.

4.2.1 Pumping

The chamber is evacuated using two pumping mechanisms: a roughing pump and a turbomolecular pump (turbopump). As is typical for a high-vacuum system, the two pumps are connected in series. The pumping system is attached to the bottom port of the central vacuum chamber.

The bottom flange of the chamber is connected directly to a bellows valve using a metal gasket. This valve allows the chamber to be isolated from both the turbopump and roughing pump.

The turbopump is attached to the opposite side of the bellows valve. It has two additional ports. One port is located on the upper side of the turbopump and connects to the VCR system discussed in Section 4.2.6. The other port connects through an open/close valve to the roughing pump. This valve provides an additional method of isolating the pumping system and is normally closed when the bellows valve is closed.

The complete pumping system is shown in Figure 4.5.

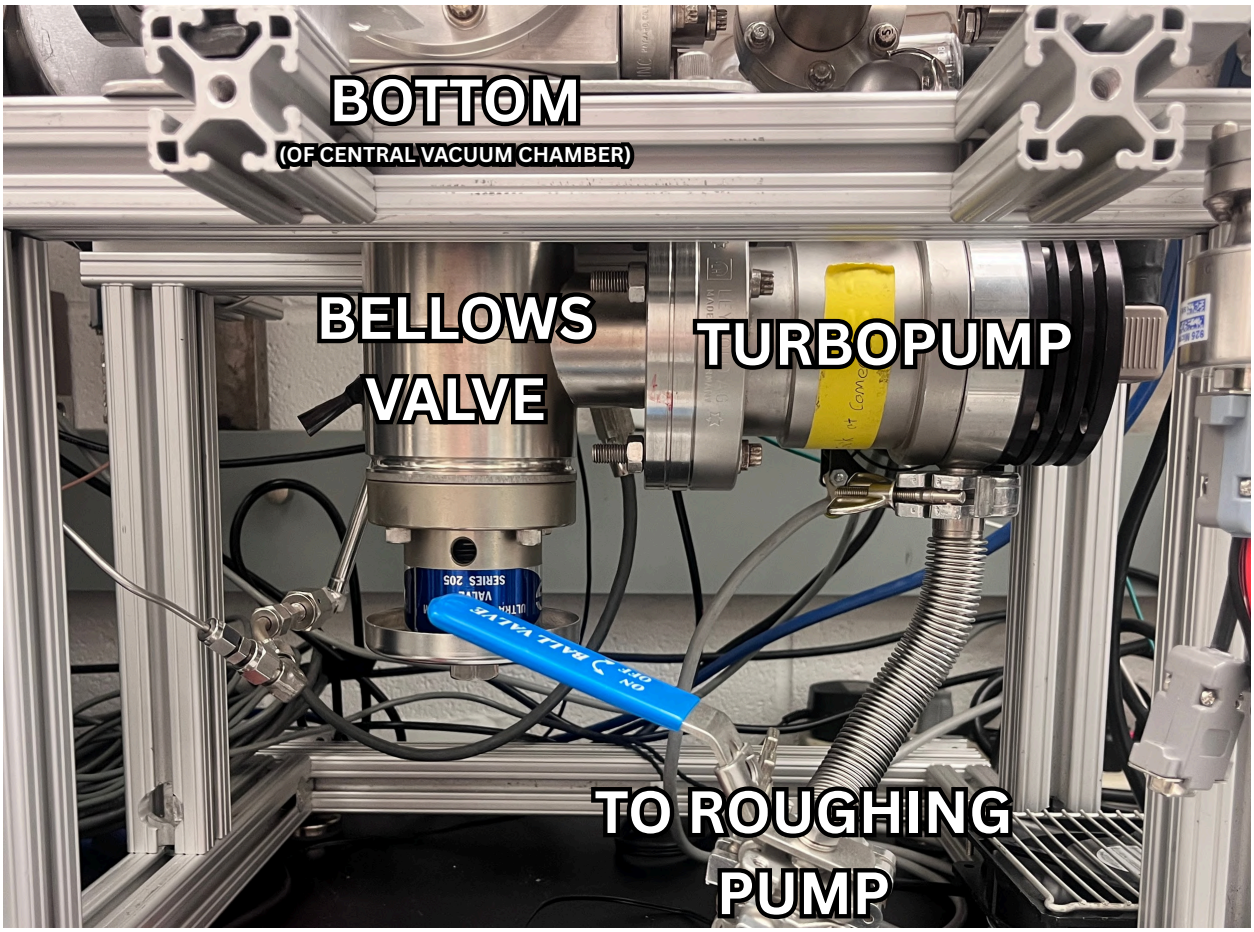

Figure 4.5 Photograph of the pumping system connected to the bottom port of the central vacuum chamber. The labeled components include the bellows valve, turbomolecular pump, roughing-pump isolation valve, and roughing-pump connection.

4.2.2 Window

A circular window is installed on the top ConFlat port of the chamber. It can be easily seen in Figure 4.3. This window enables visual alignment of the plunger during cavity formation, allowing for more precise positioning. In addition, the window facilitates placement and alignment of external diagnostic equipment, including the neutron detector.

Furthermore, the glass window is not rigidly attached to the opening of the system; instead, it rests on an O-ring. It seals due to its weight and the pressure difference between the vacuum chamber and atmosphere. This allows the glass window to act as an overpressure release. Throughout the experimental process, the chamber is filled with either argon or a hydrogen-isotope gas. Although care is taken to ensure that the chamber is not overpressurized, the window provides an additional safety measure in the event of overpressurization.

While not strictly necessary for system operation, the window provides a practical improvement for cavity control, detector positioning, and safe experimental practices.

4.2.3 Plasma Initiator

A plasma initiation assembly is attached on the backside of the central vacuum chamber. This system is seen in Figure 4.6. This assembly consists of a high-voltage power supply used to ignite and sustain a plasma within the chamber. An additional attachment modifies the incoming gas flow to ensure a steady and controlled inflow of deuterium. The plasma is initiated within a specific pressure range and can be regulated by adjusting the voltage output of the plasma source. For more information on the flow, pressure range, and voltage output used in this experiment, see Appendix C.

Once ignited, the plasma is expected to extend through the cavity region, enabling monatomic deuterium to be distributed throughout the system and onto the walls of the cavity. Throughout operation, system parameters are continuously monitored to ensure safe and stable operation. In the event of abnormal behavior, the plasma system is designed to shut off automatically.

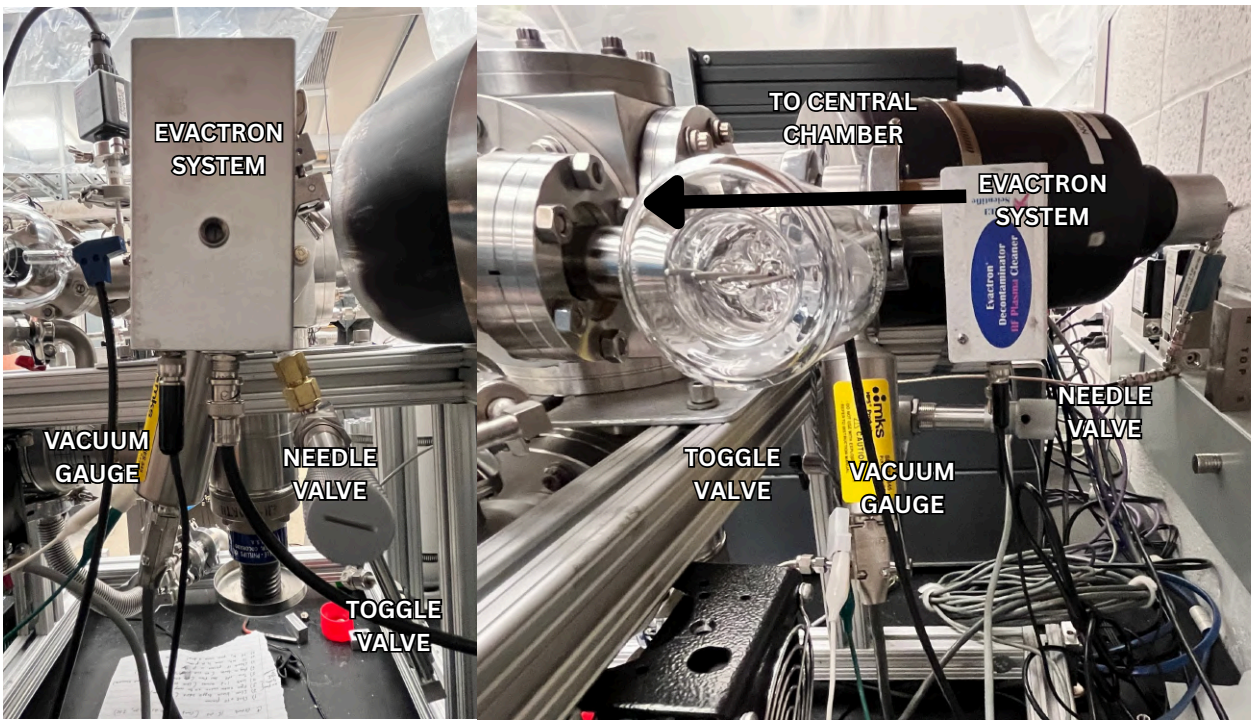

Figure 4.6 Photograph of the plasma system attached to the chamber. The Evactron system is the high voltage unit that ignites the plasma within the chamber. Attached to the walls of the vacuum system (preceding the plasma ignition point) are the various components that moderate gas inflow. In addition, a vacuum gauge is attached and its readout is on the Evactron unit seen in Figure 4.7.

After the plasma has been sustained for a sufficient period to allow deuterium loading of the cavity surfaces (see Appendix E for more information), the deuterium inlet is closed and the plasma is extinguished. At this juncture, the deuterium is either pumped out after opening the bellows valve to the turbopump or more deuterium is leaked in through the plasma initiator to reach almost atmospheric pressure (discussed in Section 4.3). At this juncture, the system can transition into the measurement phase. The cavity can then be formed using the plunger assembly while the detectors begin recording data under stable gas conditions.

In addition to the pieces attached to the system, the Evactron plasma system also has a control box to modify gas inflow, plasma power level, and overall pressure of the system. This unit is seen in Figure 4.7.

Figure 4.7 Photograph of the Evactron DECONTAMINATOR unit. This unit is used to control the plasma ignition instrumentation. This includes a toggle switch for gas flow, a plasma ignition switch, and an adjustable power output with live readout of the RF power measured.

4.2.4 Detector Positioning

Opposite the plunger, a charged-particle detector is mounted to the chamber using a custom-designed ConFlat attachment. This custom design was created by the author's advisor, John Ellsworth, and can be seen in implementation in Figure 4.8. This detector is positioned face-to-face with the target to maximize sensitivity to charged fusion products emitted from the cavity region.

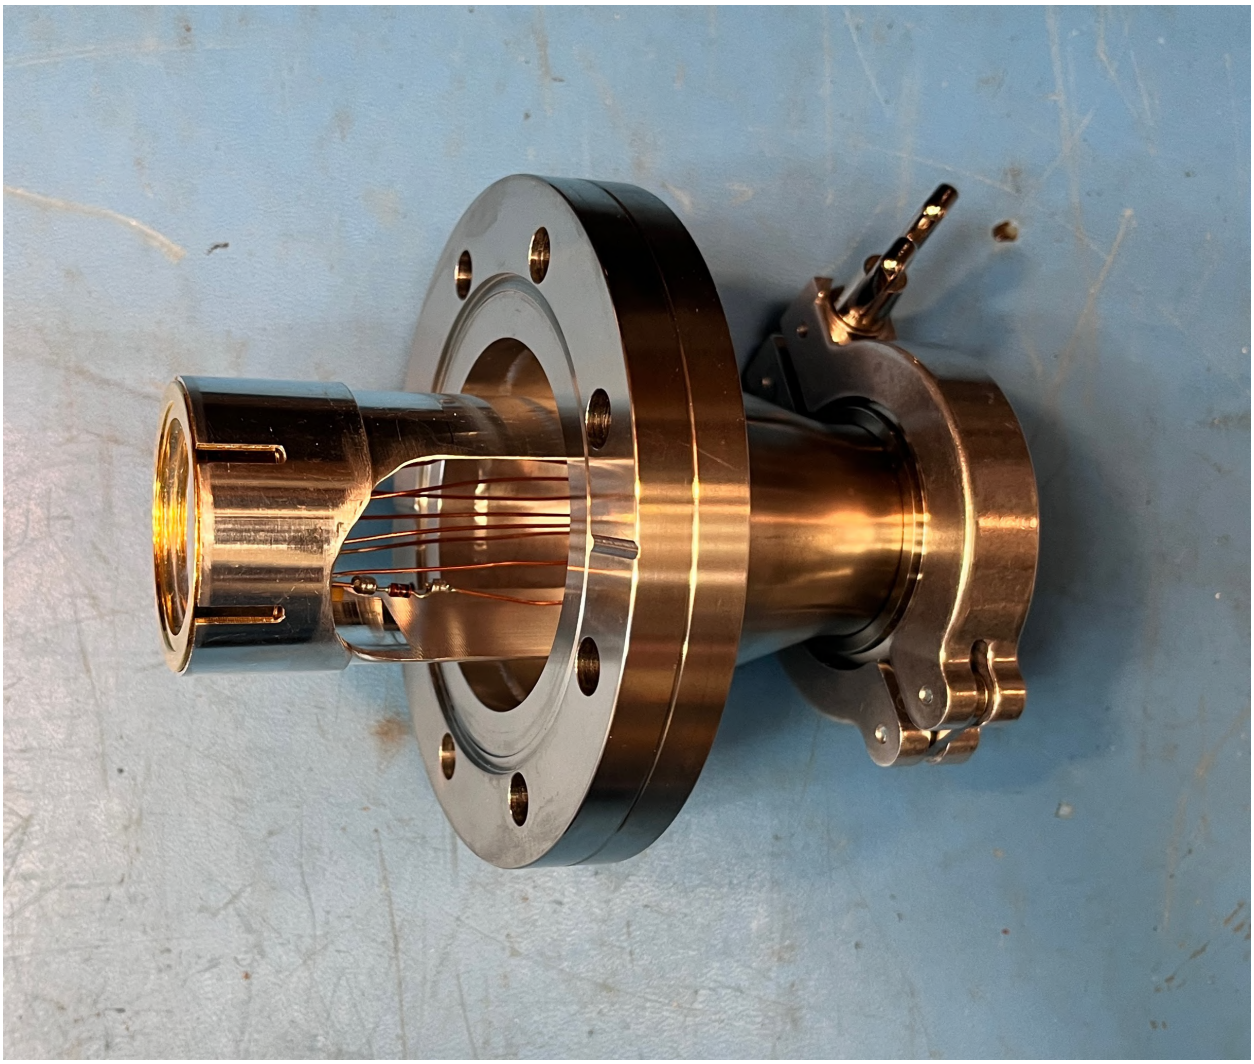

Figure 4.8 Photograph of the charged-particle detector mounted opposite the plunger using a custom ConFlat attachment (designed by J. Ellsworth). The detector faces the target directly to maximize collection of charged fusion products from the cavity.

A neutron detector is also employed in this system. Due to the high penetrating power of neutrons, precise placement of the neutron detector is less critical than for charged-particle detection. For simplicity and accessibility, the neutron detector and charged-particle detector are positioned as seen in Figure 4.9.

4.2.5 Plunger for Cavity Creation

The secondary four-way cross, in addition to providing additional ports for the system, houses the plunger assembly used to form the cavity. The internal end of the plunger holds the target prepared as described in Section 4.1. By translating the plunger along the chamber axis, the target surface can be brought toward the detector positioned on the opposite side of the chamber, allowing the separation between the two surfaces to be precisely controlled.

The mechanical design of the plunger is critical, as the detector must not be damaged during cavity formation. The assembly therefore allows smooth, controlled motion of the target while maintaining alignment between the target and detector surfaces. This configuration enables the cavity spacing to be reduced to nanometer-scale distances while minimizing the risk of mechanical contact or misalignment that could damage the detector.

In addition to enabling controlled positioning of the target, the plunger provides mechanical stability during operation so that the cavity spacing remains fixed while the plasma is sustained within the chamber. Together with the protective features described in Section 4.1, this design allows reliable formation of a confined cavity environment suitable for the fusion experiments described in this work.

4.2.6 VCR System

Coming from one port on the four-way cross is an adapter that connects the chamber to a VCR system. The VCR fittings use metal gaskets to maintain a vacuum-tight seal while providing the

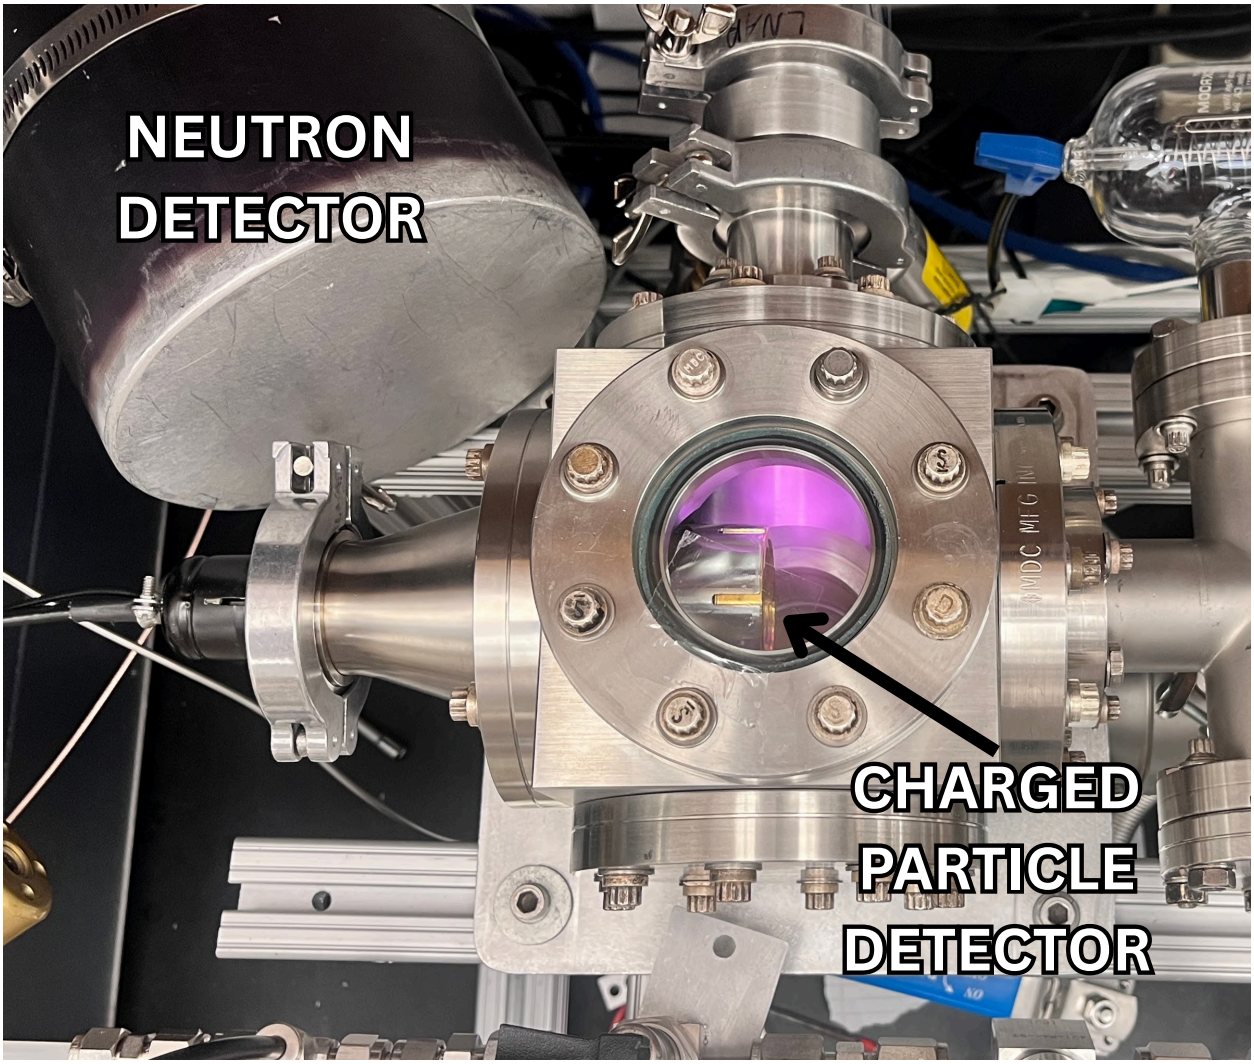

Figure 4.9 Photograph of detector positioning around the vacuum chamber. The charged-particle detector is within the chamber while the neutron detector is positioned at the back left of the central vacuum chamber.

needed space for multiple gauges, gas inlets, and pumping connections. The complete VCR system is imaged in Figure 4.10.

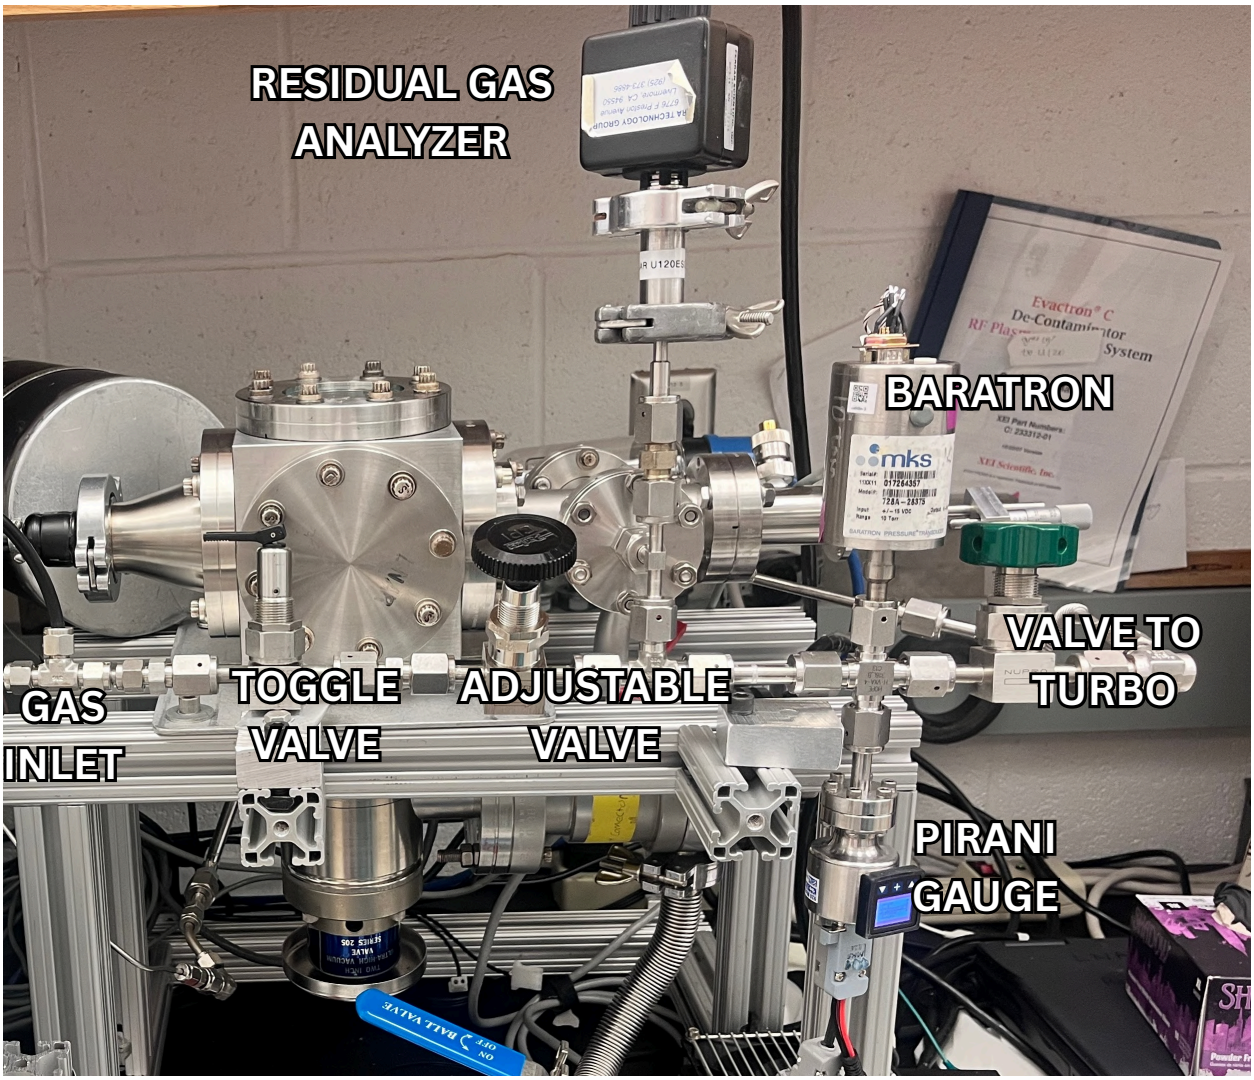

Figure 4.10 Photograph of the complete VCR system connected to the vacuum chamber. The labeled components include the MicroPirani gauge, Baratron pressure gauge, residual gas analyzer, argon gas inlet, and adjustable outlet to the turbopump.

Pirani Gauge

Attached to the system is a MicroPirani gauge that assists in determining the pressure of the vacuum system over a broad pressure range. This gauge is used during the initial pump-down to determine when the pressure is low enough to safely turn on the turbopump. It is also used while returning the chamber to atmospheric pressure and while monitoring large pressure changes during gas introduction.

The MicroPirani operates by measuring the thermal conductivity of the gas within the gauge. Because different gases transfer heat at different rates, the pressure reading depends upon the gas used to calibrate the gauge. The gauge is calibrated for nitrogen, meaning that readings taken in hydrogen, deuterium, or argon require a gas-dependent correction factor. Therefore, the MicroPirani is most useful for monitoring the general pressure range and changes in pressure rather than providing the most precise measurement during gas-specific operation.

Baratron

Similar to the Pirani gauge, the Baratron assists in describing the vacuum pressure of the system. In contrast to the Pirani, which is dependent upon the thermal properties of the gas, the Baratron measures pressure through the physical deflection of a diaphragm. As the pressure changes, the diaphragm moves and changes the capacitance measured within the gauge. This allows the Baratron reading to remain largely independent of the gas species in the chamber.

The Baratron is especially useful during plasma operation, when the pressure must be held near a specific value. For this work, the plasma is normally maintained near 0.4 Torr. The Baratron provides a more direct measurement at this pressure and allows the gas inlet and adjustable pumping valve to be modified until a stable pressure is reached. The Baratron and MicroPirani therefore complement one another by providing useful measurements during different portions of the pump-down and gas-loading process.

Residual Gas Analyzer

The residual gas analyzer (RGA) is an essential diagnostic tool for this vacuum system. Once the chamber pressure is below the specified operating limit of the RGA, it can continually measure the partial pressures associated with mass-to-charge ratios from approximately 1 to 65 atomic mass units (amu). This range includes most of the primary residual gases and contaminants expected within the vacuum chamber, including hydrogen, water, nitrogen, oxygen, argon, and carbon-containing species.

The RGA operates similarly to a mass spectrometer. Gas molecules entering the RGA are ionized by removing one or more electrons. The resulting ions then pass through a quadrupole mass filter. The electric fields within the quadrupole are varied so that only ions with a selected mass-to-charge ratio reach the detector at a given time. By sweeping across the available mass range, the RGA produces a spectrum showing ion current, or the corresponding partial pressure, as a function of mass-to-charge ratio.

A representative RGA spectrum is shown in Figure 4.11. Peaks in the spectrum can be compared with the known masses and fragmentation patterns of common residual gases. However, a single peak does not always correspond to a single gas. For example, water can produce peaks associated with H_2O, OH, and O after fragmentation within the RGA. Similarly, multiple gases may contribute to the same amu value. The complete pattern of nearby peaks must therefore be considered when identifying the chamber composition.

The displayed mass channels are separated by approximately 1 amu. This should be described as the mass resolution of the measurement rather than an uncertainty of ± 1 amu. The exact operating-pressure limit and mass resolution should be verified using the specifications for the particular RGA model before being stated numerically.

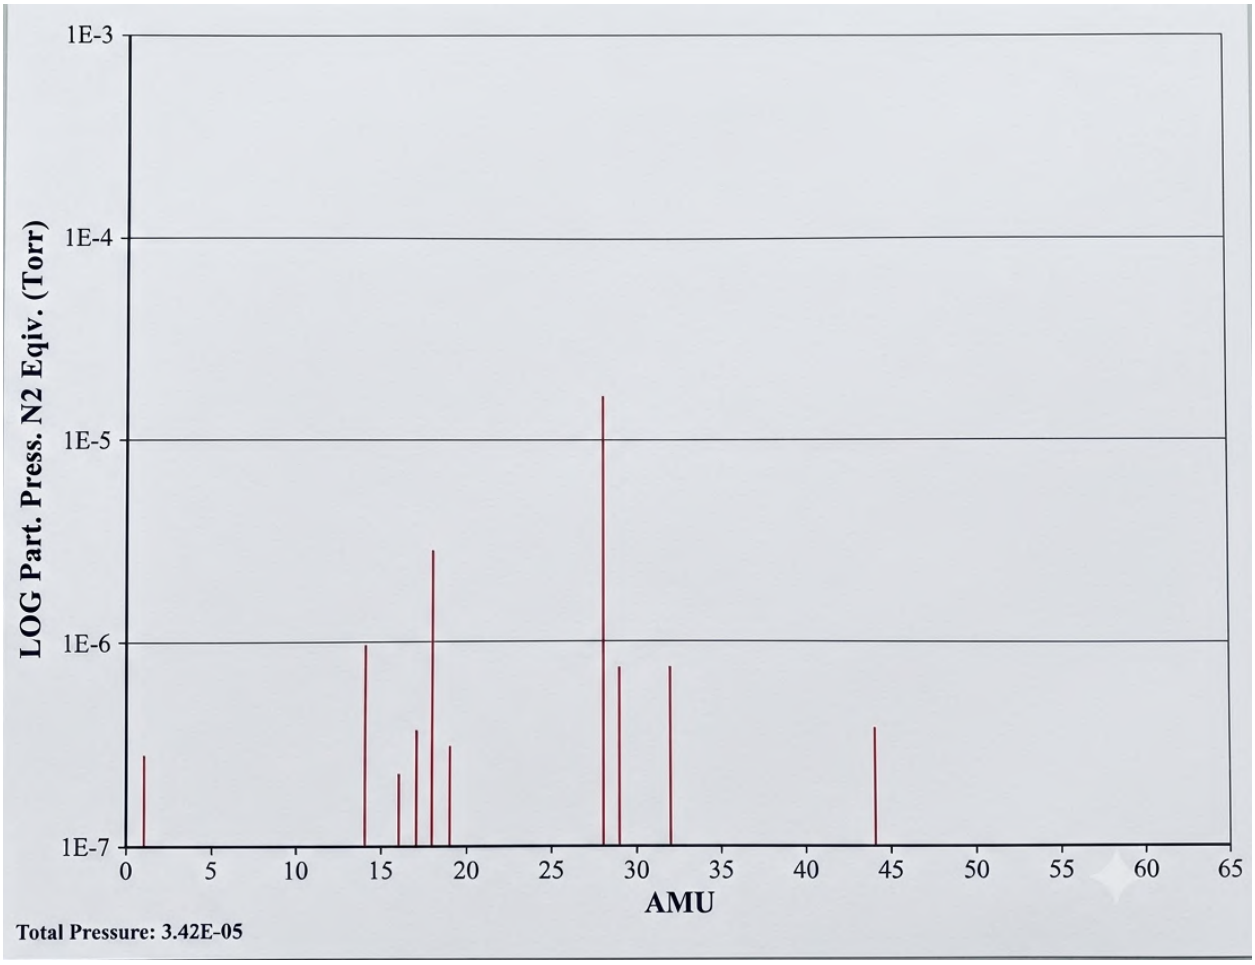

Figure 4.11 Screenshot of an RGA spectrum readout under high vacuum conditions. Because the scan progressively sweeps through AMU channels, the total pressure value displayed in the lower-left corner is skewed by an order of magnitude.

Gas Inlet

On the left side of the VCR system, tubing connects the chamber to a tank of high-purity argon. This argon is used while flushing the system to help decrease contaminants, during argon plasma cleaning, and while testing the experimental procedure without using the more expensive hydrogen and deuterium gases.

The argon inlet includes the valves needed to isolate the gas cylinder from the vacuum chamber and control the flow entering the system. These include a toggle valve and an adjustable valve. During flushing, the chamber can be partially filled with argon and then pumped down again. Repeating this process helps dilute and remove residual gases that may remain after the chamber has been opened to atmosphere.

Outlet to Turbo

On the right side of the VCR system is an adjustable valve that leads to a secondary vacuum port on the side of the turbopump. This connection allows the pumping rate to be highly moderated rather than attempting to finely control the pumping speed using the much larger bellows valve.

This control is especially important during the plasma stage, where maintaining a stable pressure is necessary for creating a consistent plasma and loading deuterium onto the target surfaces. The gas inlet and adjustable pumping valve can be changed together until the gas flowing into and out of the chamber reaches an equilibrium near the desired pressure.

During normal pump-down, the bellows valve is open and the adjustable VCR valve is normally closed. When the bellows valve is closed and additional pumping is needed, the adjustable valve can be opened in small increments or for short periods of time. This allows gas to be removed without rapidly pulling the system below the desired plasma pressure.

4.2.7 Ion Gauge

For measurements at lower pressures, a hot cathode ionization gauge tube is mounted on one of the four-way crosses that houses the plunger assembly. The ion gauge provides accurate pressure measurements in the high-vacuum regime and is used to verify the base pressure of the system prior to introducing deuterium gas. Together, these gauges allow monitoring of the chamber pressure across the full operating range required for pump-down, plasma generation, and fusion measurements.

The hot cathode ionization gauge is attached to the backside of the additional four way cross, opposite of the VCR system. Its location is pictured in Figure 4.12.

4.3 Process for Fusion

Operation of the experimental system proceeds through a sequence of stages designed to prepare the cavity environment, introduce deuterium, and subsequently measure possible fusion activity. The overall procedure consists of vacuum preparation, deuterium loading through plasma exposure, cavity formation, and long-duration data collection.

Prior to each experimental trial, the plunger assembly is positioned so that the mobile cavity wall is retracted from the detector surface, leaving a large separation between the two surfaces. In this configuration the chamber is pumped down using the roughing pump and turbomolecular pump until the desired base pressure of 1 μtorr is reached. Establishing this sufficiently low background pressure removes residual gases and contaminants that could interfere with plasma formation or introduce unwanted background signals during measurement. One of the main contaminants removed in this process is oxygen which could deteriorate the monolayer and interfere with the silicon's surface.

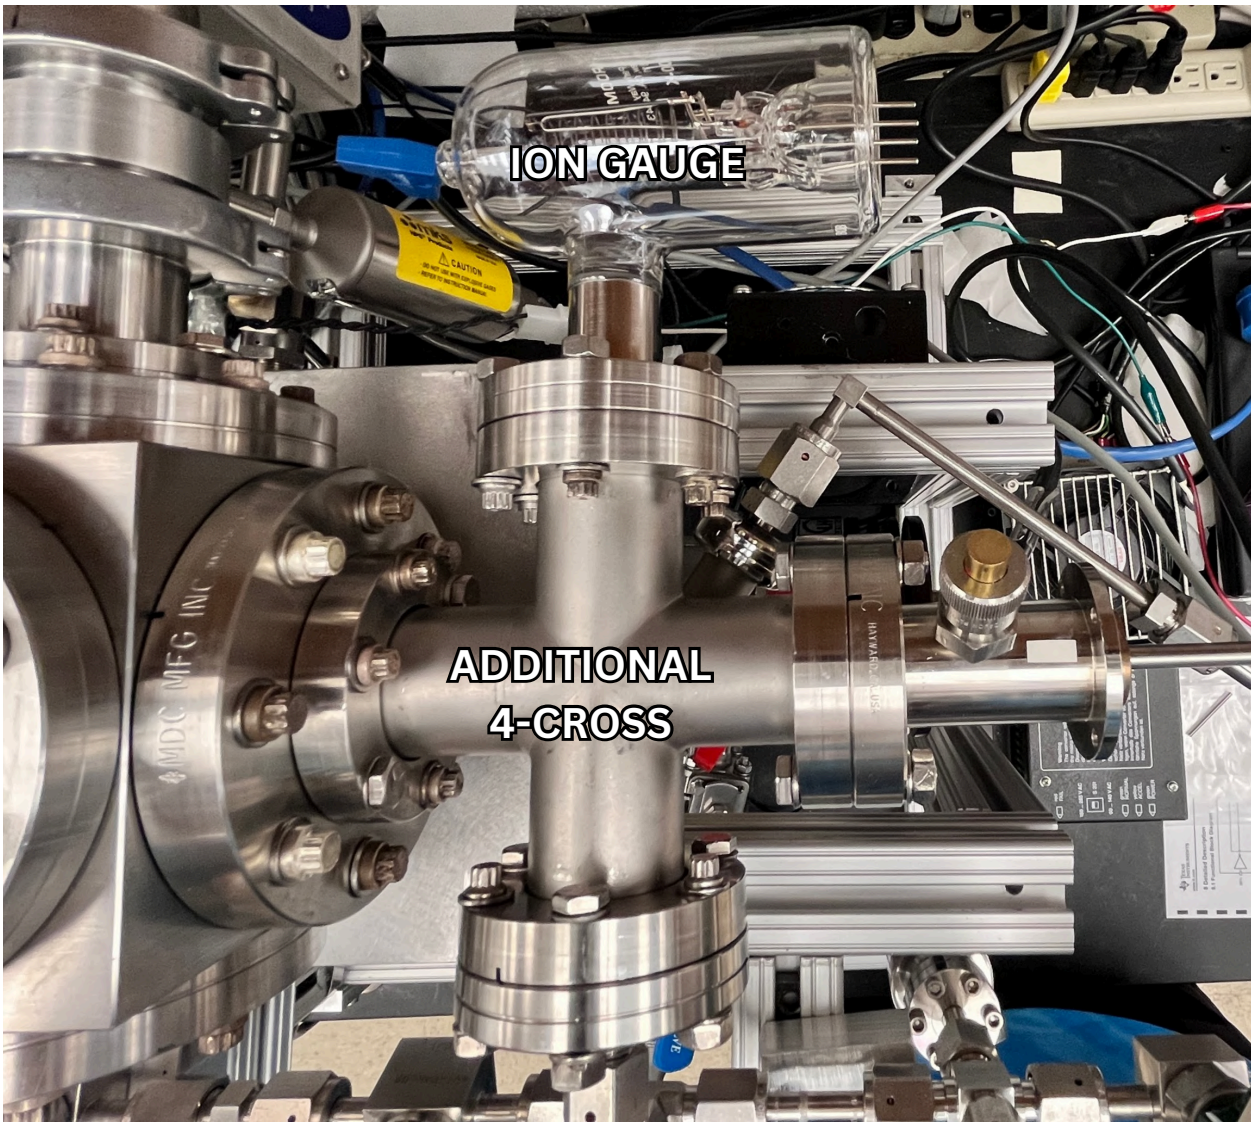

Figure 4.12 Photograph displaying the location of the hot-cathode ionization gauge mounted on the vacuum chamber.

Once the base vacuum has been established, the bellows valve to the chamber is closed tightly and deuterium gas is introduced through the gas handling system. At this stage the secondary port for pumping (Sec. 4.2.6) can be adjusted to ensure the desired pressure of 0.4 torr is maintained. This configuration allows for the majority of deuterium to remain within the chamber while maintaining controlled pressure conditions suitable for plasma formation.

After the desired deuterium pressure is reached and remains stable, the plasma initiation system is activated. The plasma is sustained within the chamber for an extended period of time, during which energetic ions interact with the exposed surfaces of the target and detector. This process deposits and implants deuterium into the surfaces that will later form the cavity walls. Because the plunger remains retracted during this stage, the two cavity surfaces are separated by a relatively large distance and the deuterium deposition occurs on the surfaces that will eventually face each other once the cavity is formed.

The plasma exposure stage serves an important role in preparing the cavity surfaces for the fusion experiment. Energetic ions and neutral species within the plasma interact with the exposed surfaces of the target and detector, allowing deuterium to become adsorbed and implanted within the near-surface region of these materials. This process effectively loads the cavity walls with deuterium, increasing the local concentration of deuterium nuclei that will later be confined within the cavity once it is formed. A sufficiently high deuterium concentration is necessary to maximize the probability of deuterium-deuterium interactions within the confined cavity region during the subsequent measurement period. For more information on the plasma pressures, timing, and system, see Appendix C.

Following sufficient plasma exposure to allow deuterium loading of the surfaces, the plasma source is shut down and the system is allowed to stabilize. Stabilization occurs either at the lowest vacuum pressure achievable by the system (< 1 μTorr) or at near-atmospheric pressure, with the chamber containing almost entirely deuterium. The reason for testing at both pressures is discussed

in Appendix B. Measurements at both pressures were performed and alternated with background data collection.

Data acquisition from the neutron and charged-particle detectors is then initiated. Shortly after data collection begins, the plunger assembly is slowly advanced toward the detector surface. As the plunger moves inward, the silicon target approaches the detector until the aluminum spacer columns described in Section 4.1 make contact with the detector surface. These spacers establish the final separation between the two surfaces, thereby forming the microscopic cavity in which the fusion measurements are conducted.

Once the cavity has been formed, the detector systems continue recording data for an extended measurement period, typically on the order of 12-36 hours. During this time the system pressure and other operational parameters are monitored to ensure stable conditions. Any fusion products generated can be detected and subsequently analyzed to determine the fusion rate, subject to the geometric coverage and intrinsic efficiency of the neutron and charged-particle detectors.

4.4 Detectors

Fusion products in this experiment are measured with two detector systems: a neutron detector and a charged-particle detector. The neutron detector allows for the counting of fusion processes which output fusion products of branch (ii) of equation 4.1. The charged-particle detector allows for the counting of fusion processes which output fusion products of branches (i), (ii), or (iii) of equation 4.1. Used together, these instruments provide a reliable determination of the deuterium-deuterium fusion rate, the primary quantity needed for comparison with previous experimental and theoretical work. A well-constrained fusion rate also enables extraction of reaction cross sections, electron-screening effects, and related parameters.

This section of the thesis describes the two detector systems used in the experiment and their roles in measuring fusion activity within the cavity. The neutron detector subsystem is presented first, including its operating principle, placement relative to the chamber, and method for determining neutron production rates. The charged-particle detector is then described, along with its positioning within the vacuum chamber and its use in measuring charged fusion products emitted from the cavity region. Together, these subsections outline how the detector systems are implemented and how their measurements are used to determine the overall fusion rate.

4.4.1 Neutron Detector

Neutrons produced through deuterium-deuterium fusion are detected using a lithium gadolinium borate (LGB) capture-gated neutron detector of the type described in [36]. Small crystals of LGB are placed in a scintillating material (see Figure 4.13). This detector operates by moderating fast neutrons to thermal energies, followed by neutron capture within the LGB scintillator material. Both the thermalization and capture emit a signal, as seen in Figure 4.14. Neutron capture occurs primarily on lithium and gadolinium nuclei, producing charged reaction products that generate scintillation light, which is subsequently converted into an electronic signal. The capture-gated detection technique allows discrimination between neutron-induced events and background signals based on characteristic timing and energy signatures.

In this experiment, the neutron detector is positioned external to the vacuum chamber, above the optical window mounted on the top ConFlat flange. Because neutrons readily penetrate the chamber walls and surrounding materials, precise placement relative to the cavity is not critical. The detector occupies a known solid angle with respect to the cavity region, which is determined geometrically based on its placement. The data from the neutron detector is digitized through a CAEN and recorded locally on a computer. The data is parsed and analyzed and graphed using a

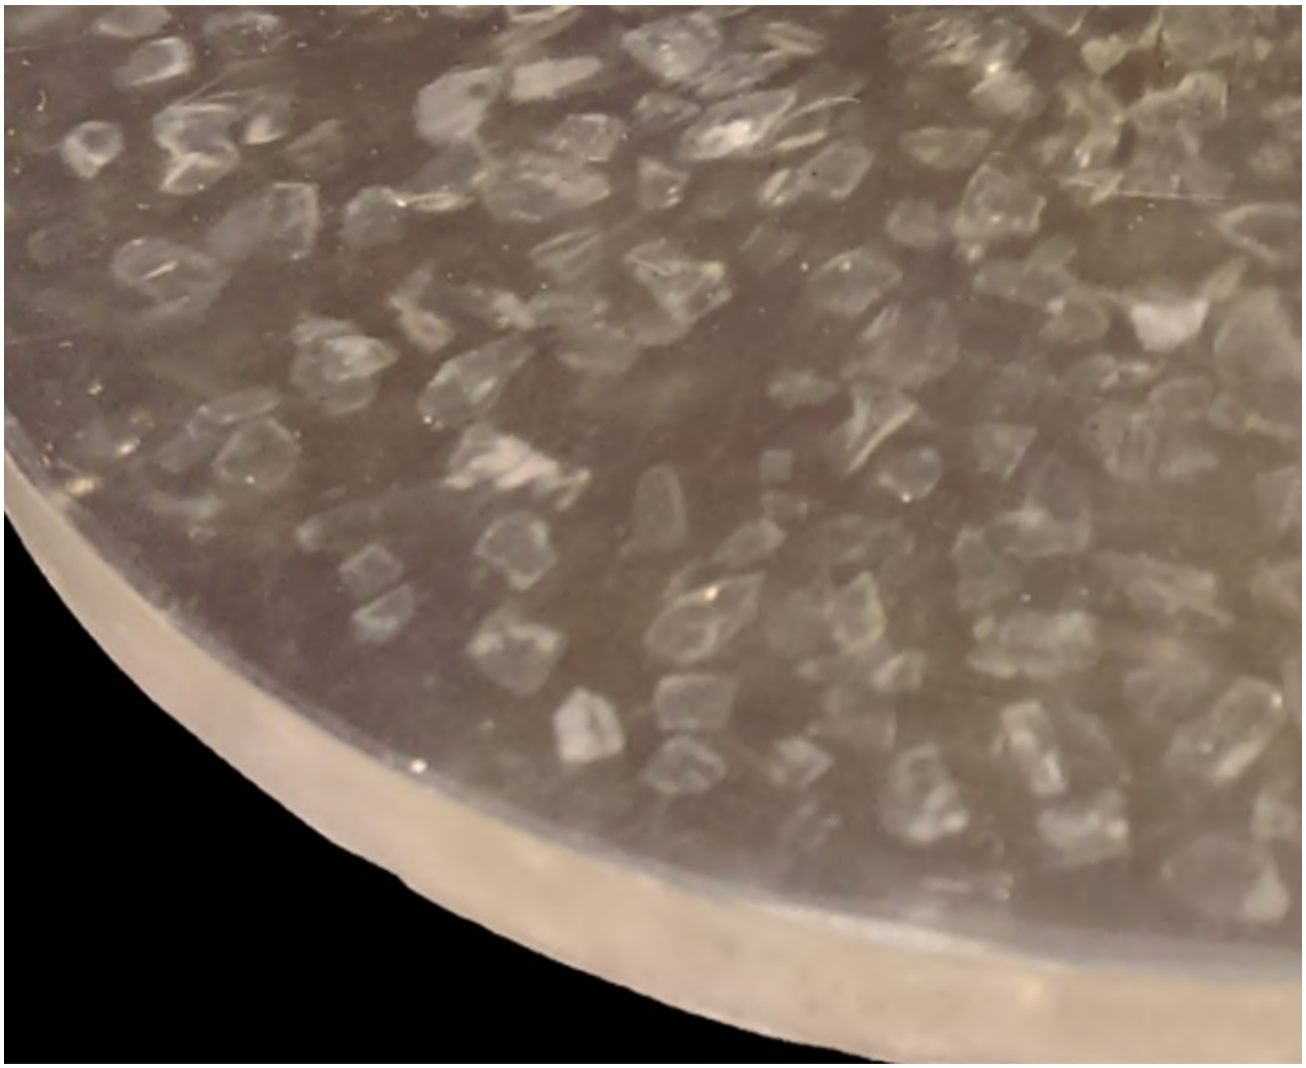

Figure 4.13 LGB crystals placed isotropically in a scintillating material.

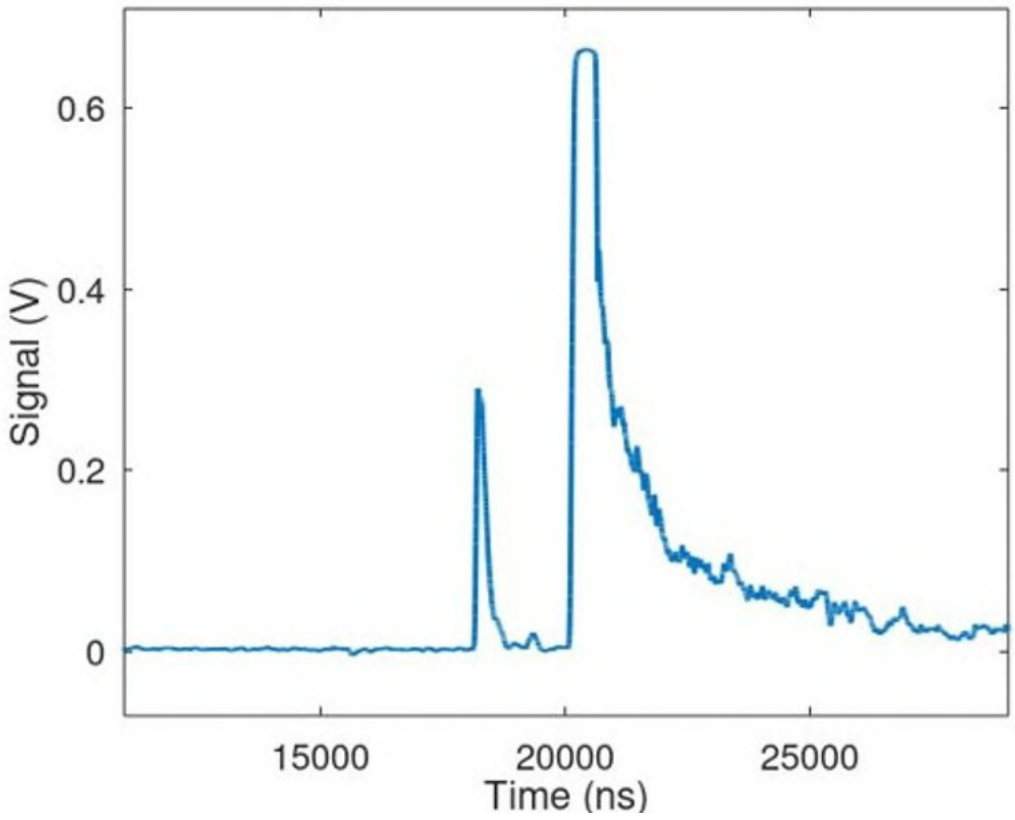

Figure 4.14 Plot of response to a high-energy neutron in the LGB neutron detector. The first pulse shows the fast neutron converting to a thermal neutron. The second shows the capture in the scintillator material.

program written in GNU Octave. The code then shows a block of neutron pulse images, similar to that pictured in Figure 4.15.

The neutron count rate measured by the LGB detector, together with the known detection efficiency and solid angle coverage, is used to calculate the neutron production rate associated with deuterium–deuterium fusion. Since neutron-producing reactions account for approximately one-half of all d-d fusion events, these measurements provide a direct estimate of the total fusion rate. In addition, the capture-gated technique offers improved background rejection, making the neutron detector a reliable diagnostic for low-rate fusion measurements in this system.

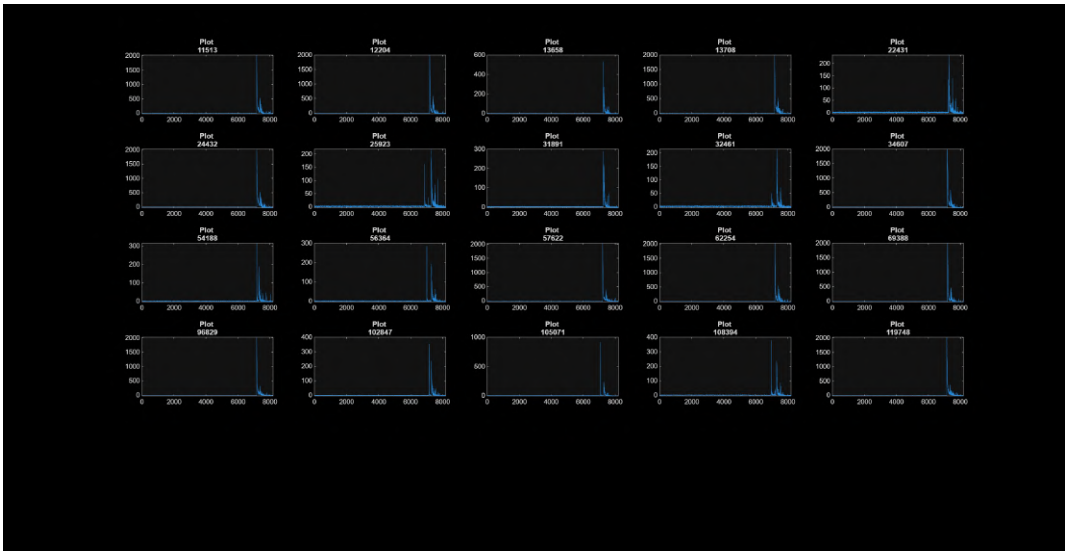

Figure 4.15 Representative neutron pulses identified by the GNU Octave neutron pulse-finder code from digitized neutron detector data. Each trace shows an individual pulse selected from the recorded detector waveform for subsequent analysis.

4.4.2 Charged-Particle Detector

Charged fusion products are detected using a DIADII-900 integrated silicon surface barrier detector (SBD) system. This detector uses a silicon surface barrier detector with an internal amplifier and single-channel analyzer, allowing for direct detection and electronic discrimination of charged particles emitted from the cavity region. An example of the charged-particle detector output is shown in Figure 4.16. Silicon surface barrier detectors are well suited for this application due to their sensitivity to charged particles in the MeV energy range and their ability to provide measurements with energy discrimination.

In the context of deuterium-deuterium fusion, this detector enables measurement of charged reaction products including protons, tritons, ^{3}He nuclei, and α-particles. These particles deposit their kinetic energy directly within the detector material. The resulting electronic signals are proportional to the particle energies, allowing for discrimination between different reaction branches through appropriate window settings. The integrated amplifier and single-channel analyzer provide

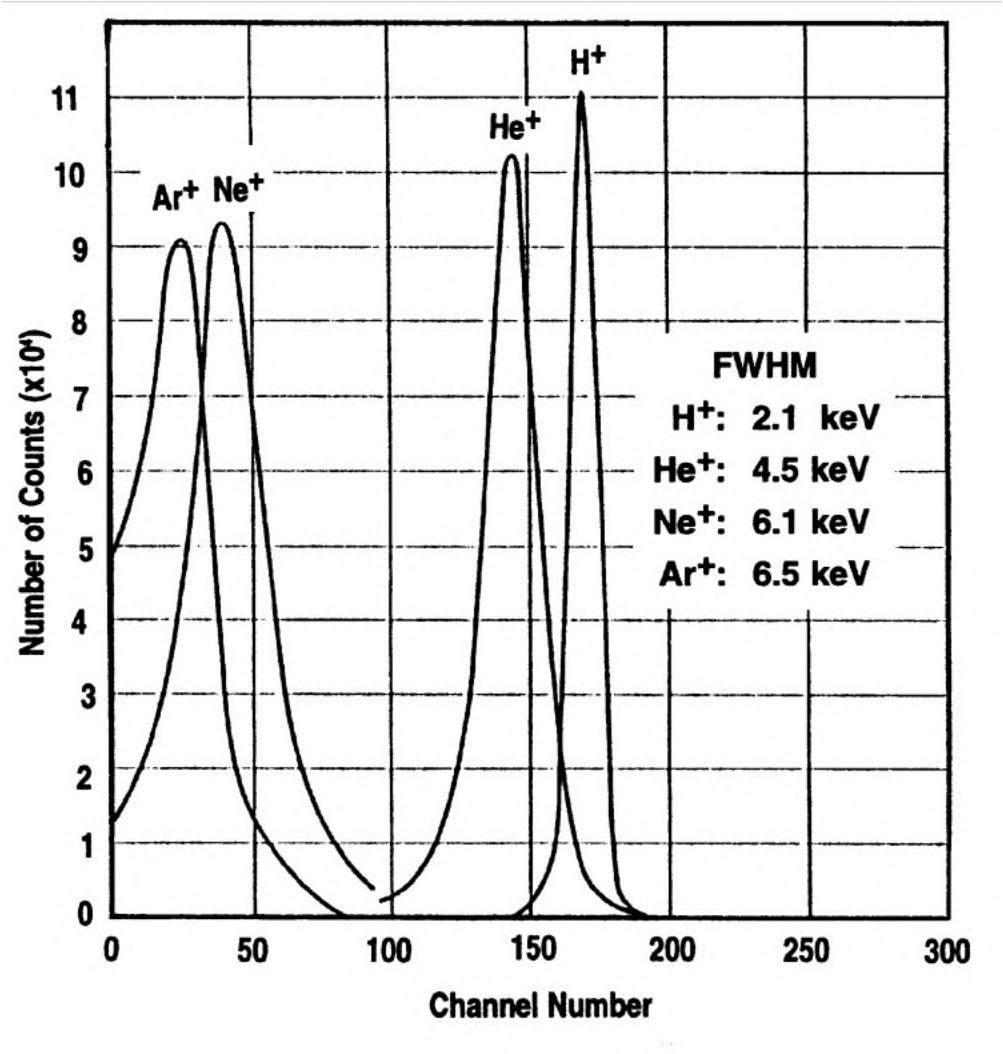

Figure 4.16 Charged-particle detector data adjusted to background, showing four particles with different charge states. Reproduced from AMETEK ORTEC's *Introduction to Charged-Particle Detectors* manual [37].

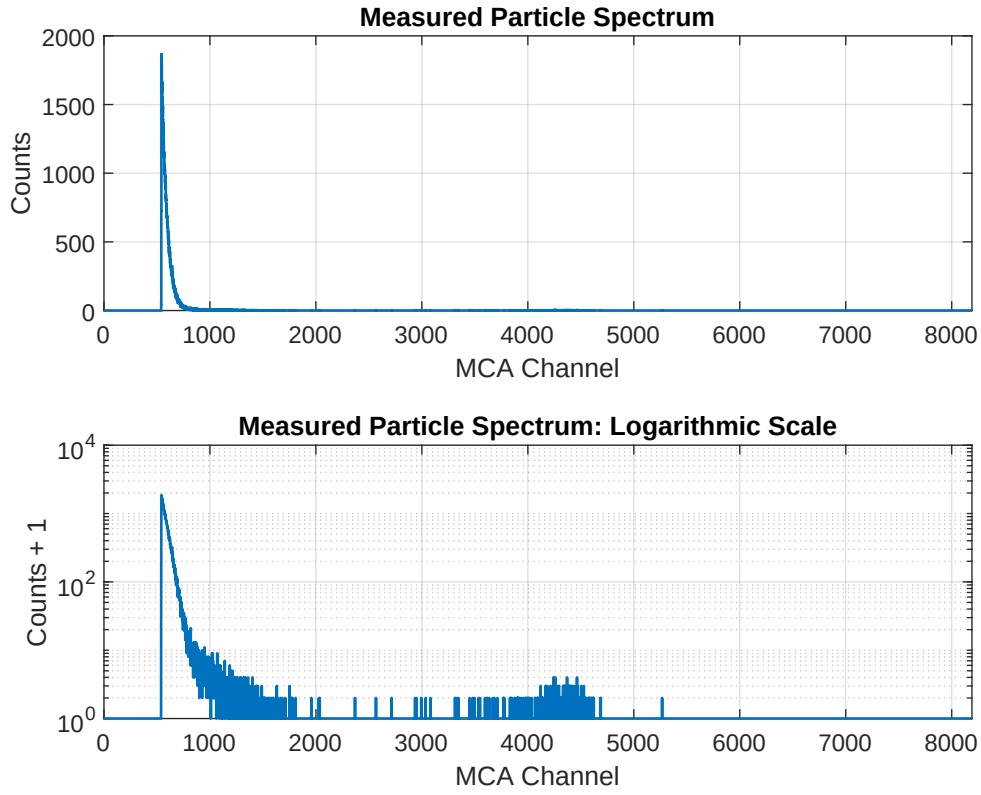

Figure 4.17 Representative charged-particle spectrum recorded using the ORTEC MAESTRO multichannel analyzer software. The upper plot shows the measured counts as a function of MCA channel number on a linear scale, while the lower plot shows the same data on a logarithmic scale. The logarithmic scale makes low-count features more visible, including a peak at higher channels corresponding to the Am source placed in the detector, which is not readily visible in the linear-scale plot. The channel axis has not yet been converted to particle energy.

stable signal processing and reduce external electronic noise, improving measurement reliability. The digitizer output is processed using the ORTEC MAESTRO MCA emulation software for data collection. A representative data set is shown in Figure 4.17.

The charged-particle detector is positioned face-to-face with the cavity region to maximize acceptance area. Once the detector efficiency and solid angle coverage are accounted for, the

measured charged-particle count rates can be used to calculate the fusion rate independently of the neutron measurements. This complementary detection method provides an important cross-check of fusion activity and enables comparison between charged-particle and neutron-producing branches of the deuterium-deuterium fusion reaction.

Chapter 5

Results

This chapter presents the results of the theoretical screening models developed in Chapter 3, as well as the experimental results obtained using the system discussed in Chapter 4. The first section will forgo the presentation of empirical fusion data and will instead focus on the predicted modifications to electron screening arising from confined geometries and the corresponding implications for low-energy deuterium-deuterium fusion.

Following the theoretical results, the experimental results will be presented. A more detailed analysis, including a comparison between the theoretical predictions and experimental measurements, will be postponed until Chapter 6.

5.1 Predicted Screening Enhancement

The screening models developed in this work provide a range of predicted modifications to the effective Coulomb interaction between reacting nuclei. Classical screening approaches, including Debye and Thomas-Fermi models, predict relatively modest reductions in the Coulomb barrier under typical metallic conditions. These reductions lead to small enhancements in tunneling probability and are insufficient to account for large deviations from standard fusion rates.

In contrast, the stochastic electrodynamics (SED) and quantum electrodynamics (QED) frameworks predict that electromagnetic boundary conditions within confined cavities can significantly alter the electronic environment. In particular, suppression of vacuum modes at frequencies near the electron orbital frequency leads to a redistribution of electron density toward the nucleus. This increased localization enhances screening and reduces the effective tunneling barrier.

5.2 Theoretical Results

This section summarizes the principal theoretical results developed in this work. Particular emphasis is placed on the characteristic cavity length scales at which modifications to electron behavior are predicted to become significant, along with the corresponding implications for electron screening and deuterium-deuterium fusion enhancement. These results also provide the theoretical basis for identifying experimentally relevant cavity dimensions.

5.2.1 Critical Cavity Length Scales

A key result of this work is the identification of characteristic cavity length scales required to produce measurable changes in electron behavior. From the SED framework, comparison between the cavity cutoff frequency and the hydrogen ground-state frequency yields a threshold condition

$$L \leq \frac{\pi c}{\omega_0} \approx 45.6 \text{ nm}, \tag{5.1}$$

below which modifications to the electron orbital structure are expected to become significant. For cavity sizes larger than this threshold, the electromagnetic mode density remains sufficiently continuous, and the electron behaves similarly to the free-space case.

Stronger effects are predicted as the cavity size is reduced further. By comparing to a muon-like orbital scale, a characteristic frequency corresponding to an energy of approximately 2815 eV yields

$$L \leq \frac{\pi c}{\omega_\mu} \approx 2.2 \text{ Å}. \tag{5.2}$$

At this scale, the predicted electron localization approaches that of muonic hydrogen, resulting in substantially enhanced screening. These results define a hierarchy of confinement regimes, ranging from negligible effects at large cavity sizes to strong modifications at angstrom-scale separations.

5.2.2 Implications for Fusion Enhancement

The predicted reduction in electron orbital radius leads directly to an increase in electron density near the nucleus, which enhances screening of the Coulomb repulsion between reacting deuterons. Because tunneling probabilities depend exponentially on the barrier width and height, even modest changes in the effective potential can produce large increases in fusion probability.

Within this framework, cavity confinement provides a mechanism for modifying fusion rates through geometric control of the electromagnetic environment. In particular, cavities with dimensions below 50 nm are expected to produce measurable deviations from conventional screening predictions, while cavities approaching angstrom-scale dimensions may produce significantly stronger enhancements.

These results also suggest that naturally occurring surface features in metals, such as defects and nanoscale cavities, may contribute to localized regions of enhanced screening. Such structures could provide a physical basis for previously reported anomalies in fusion rates within metallic environments.

5.2.3 Experimental Outlook

The theoretical results presented here define clear experimental targets, particularly the transition into the sub-50 nm regime and, ultimately, toward angstrom-scale confinement. Future experiments

will be required to determine whether the predicted scaling of screening enhancement with cavity size is realized in practice.

5.3 Experimental Results

Experimental measurements were collected under three primary conditions: background, hydrogen, and deuterium. For each condition, multiple data sets were recorded using the neutron and charged-particle detection systems described previously. These measurements included both shorter and longer collection periods, along with repeated trials intended to identify whether any measurable change in detector response occurred under the different experimental conditions.

To separate the measured signal from the background contribution, a MATLAB routine was developed to subtract a separately measured background spectrum from each foreground spectrum on a channel-by-channel basis. Because the spectra consist of discrete MCA bins, the background counts associated with each channel were first subtracted from the corresponding foreground channel. If the background counts exceeded the available foreground counts in that channel, the remaining counts were subtracted from the neighboring channels, extending up to two channels on either side of the original bin. This prevented individual bins from becoming negative when nearby counts were available to account for the background contribution. If the background contribution still exceeded the total counts available within this five-channel region, the remaining contribution was distributed across the same five-channel interval. Figure 5.1 shows a representative foreground spectrum, measured background spectrum, and the resulting background-subtracted spectrum.

The results presented in this section are primarily descriptive. Representative data from each experimental condition are shown, with comparison to the measured background used to evaluate whether any clear signal emerged during the hydrogen or deuterium experiments. More detailed

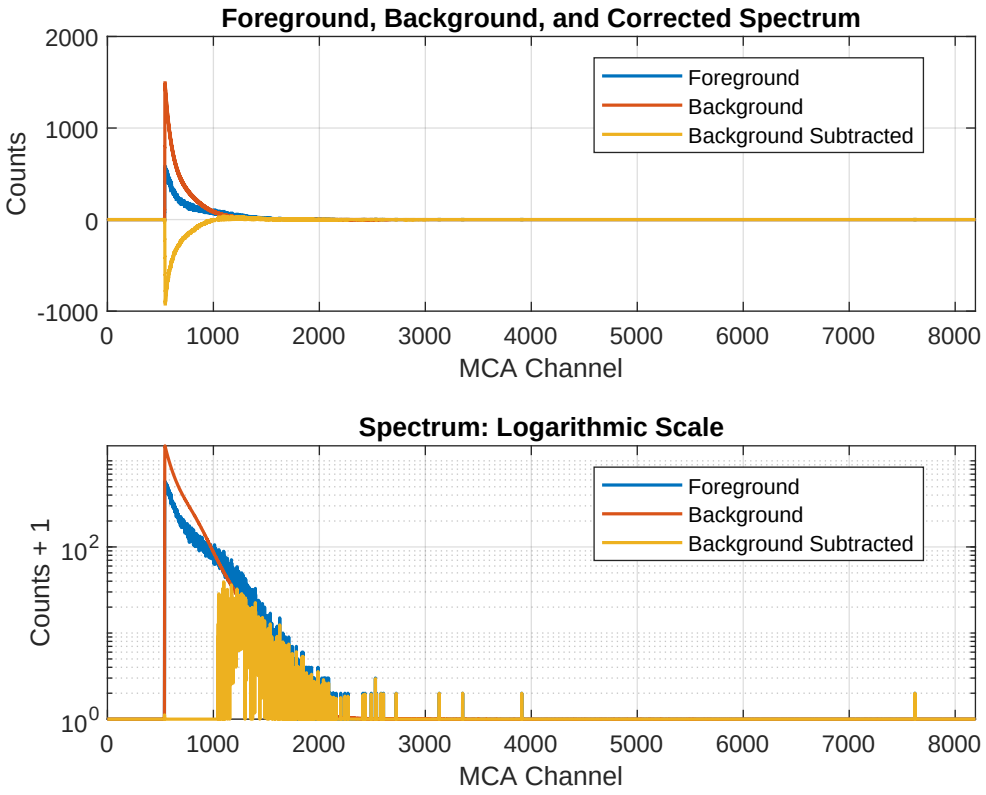

Figure 5.1 Representative particle spectrum showing the measured foreground, measured background, and resulting background-subtracted spectrum. Background subtraction was performed on a channel-by-channel basis using the procedure described in the text, with residual background counts distributed among neighboring channels when necessary.

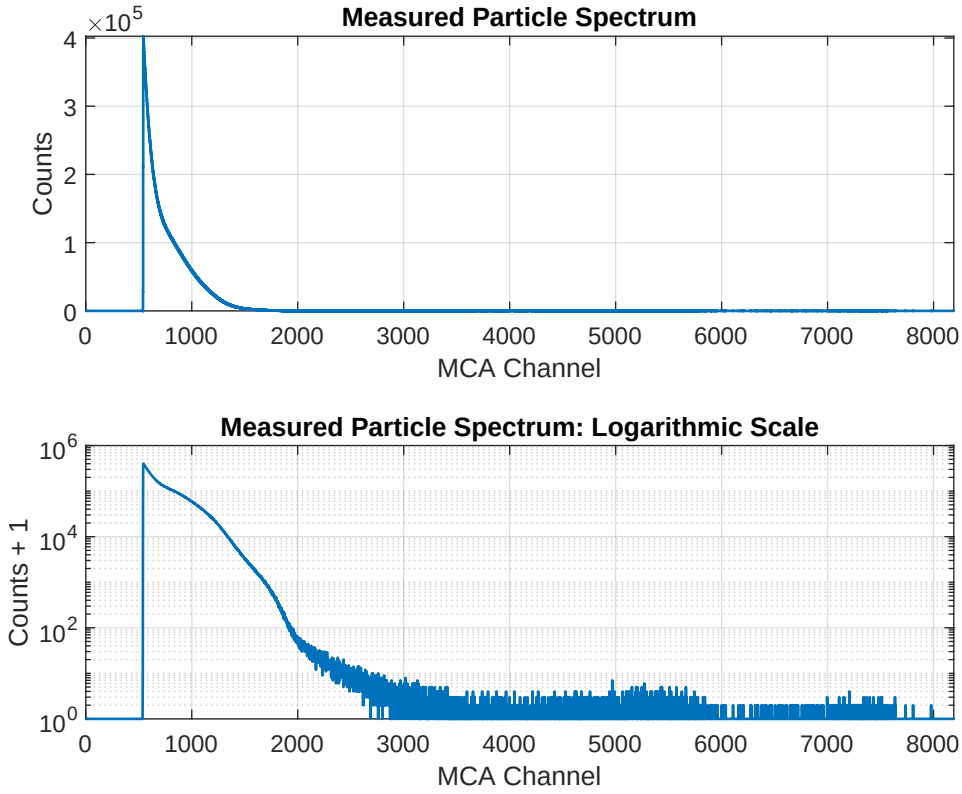

Figure 5.2 Charged-particle detector data background measurement completed before one of the deuterium experimental trials. This ran for approximately 24 hours.

interpretation of these results, including possible explanations for the observed behavior and limitations of the experimental procedure, is presented later in the discussion.

5.3.1 Background Results

Background measurements were collected prior to the hydrogen and deuterium experiments in order to characterize the normal detector response in the absence of an intentionally introduced fusion environment. An example of the background data for the charged-particle detector is shown in Figure 5.2.

The measured background consisted primarily of low-level detector noise together with occasional pulses consistent with the expected environmental and cosmic-ray background. No persistent or well-defined signal was observed that would indicate the presence of a localized source within the experimental apparatus. Variations were observed between individual background measurements, particularly for shorter collection periods, but the overall detector response remained qualitatively consistent across the measurements.

These measurements provide the baseline against which the hydrogen and deuterium experiments are compared. Because occasional pulses are present even during background collection, the observation of individual events alone is insufficient to establish the presence of a fusion-related signal. Instead, changes in the number, distribution, energy, or other characteristics of the detected events must be considered relative to the measured background.

5.3.2 Hydrogen Results

Hydrogen experiments were performed as an initial comparison to the background measurements and to evaluate the response of the experimental system under conditions similar to those later used with deuterium. Representative hydrogen data are shown in Figure 5.3.

The hydrogen measurements did not show an obvious change in detector response relative to the measured background. The observed pulses and count for the neutron detector, as well as the overall distributions of the charged particle detector, remained generally similar to those obtained during background collection, and no clearly distinguishable signal emerged from the hydrogen data.

The hydrogen run served as a functional baseline and procedure test rather than a strict quantitative control. Because the cross section for protium-protium fusion is extremely low under these conditions, fusion was not expected to occur at a detectable rate. Additionally, this phase allowed for refining the experimental protocol before using deuterium. Within the sensitivity and limitations

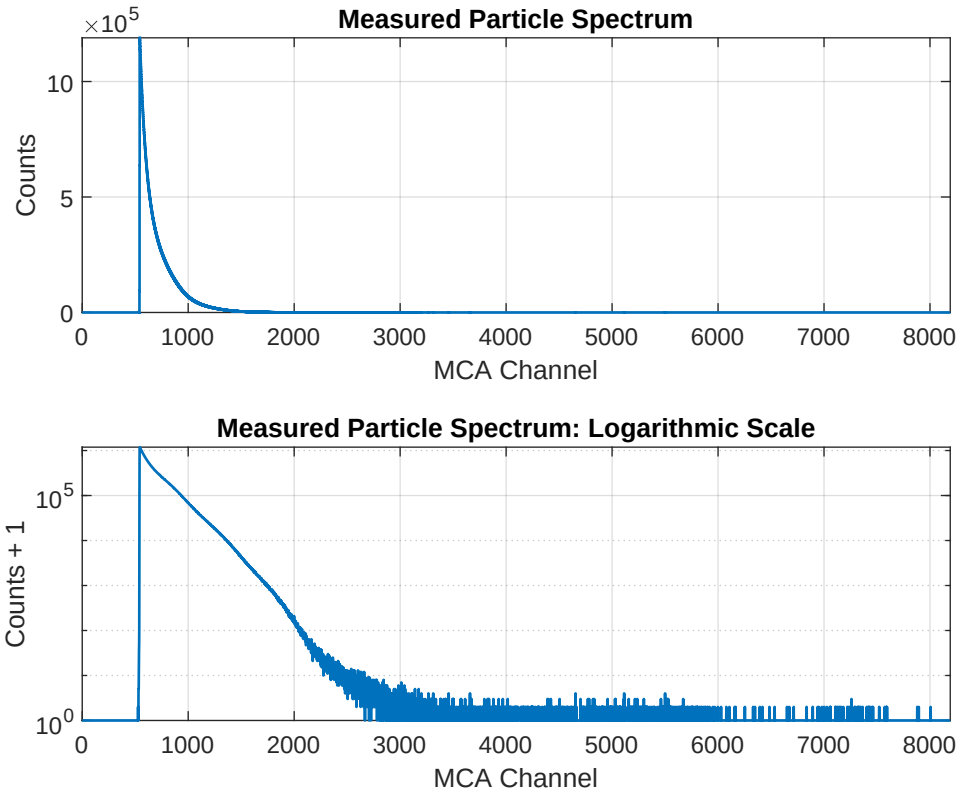

Figure 5.3 Charged-particle detector spectrum acquired during a 24-hour hydrogen trial run. No discernible peaks or significant deviations from background levels were observed.

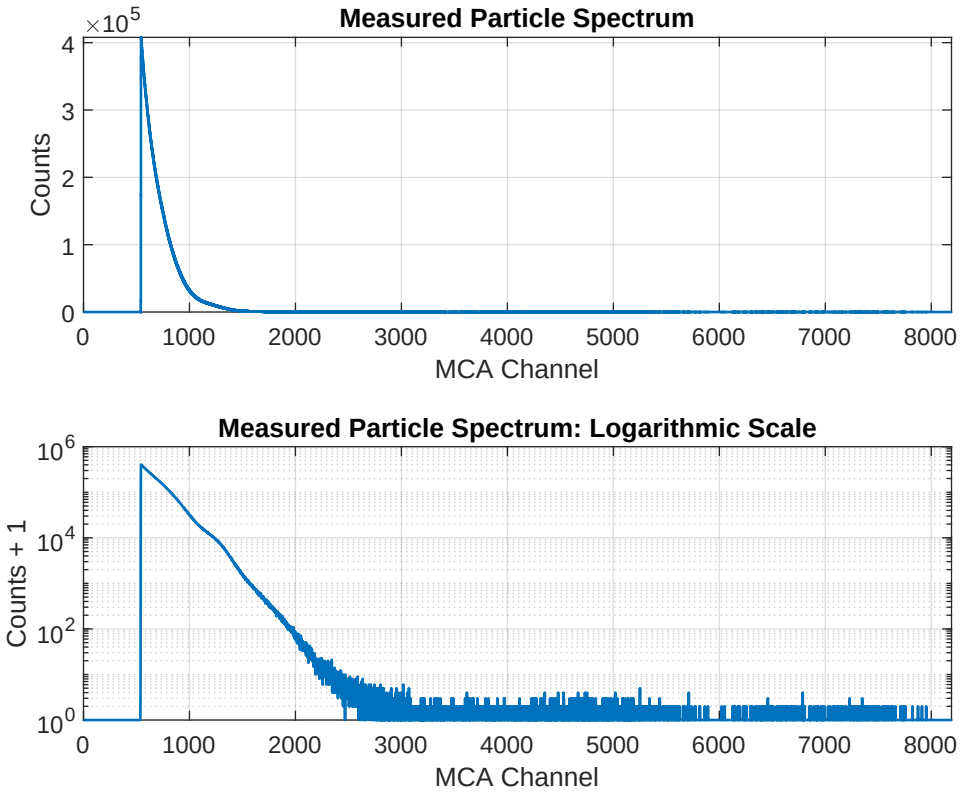

Figure 5.4 Charged-particle detector spectrum acquired during a 72-hour deuterium trial run. No discernible peaks or significant deviations from background levels were observed.

of the detector setup, no statistically significant difference was observed between the hydrogen baseline and background measurements.

5.3.3 Deuterium Results

Deuterium measurements were performed using the experimental procedure developed in this work and represent the primary experimental test of the system. Representative data are shown in Figure 5.4.

As with the hydrogen measurements, no immediately apparent or consistently reproducible change in detector response was observed relative to the measured background. Individual pulses

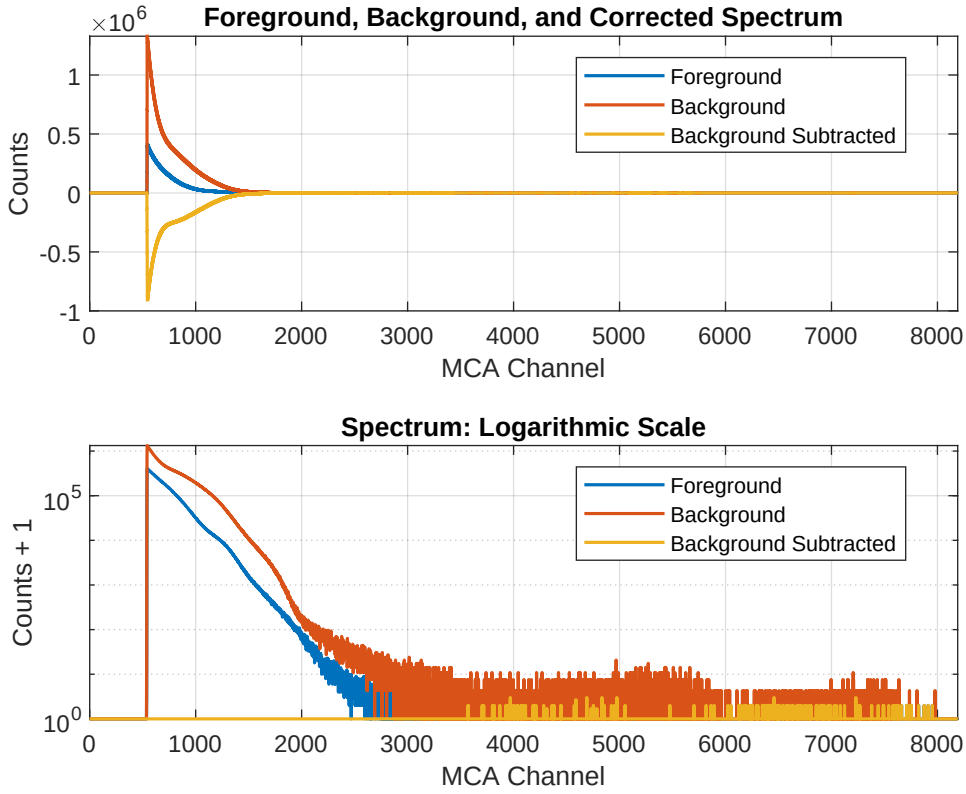

Figure 5.5 Particle spectrum obtained during a 72-hour deuterium measurement, shown with the corresponding 24-hour background measurement and the resulting background-subtracted spectrum. Background subtraction was performed on a channel-by-channel basis using the procedure described in the text.

and variations in detector activity were present throughout the measurements; however, similar behavior was also observed in the background data. This can be seen in Figure 5.5. Consequently, these events cannot be attributed directly to deuterium-deuterium fusion based solely on their presence.

Some differences between individual data sets were observed, but no clear trend emerged that consistently distinguished the deuterium measurements from the background or hydrogen measurements. In particular, the experimental data did not display an obvious new population of

events or a sufficiently distinct signal that could be confidently identified as resulting from the introduced deuterium conditions.

The deuterium measurements therefore remain inconclusive. The results do not immediately provide sufficient evidence to establish an experimentally observed increase in fusion-related detector activity under the conditions investigated. At the same time, the measurements do not independently exclude the possibility of a small effect below the sensitivity or statistical resolution of the present experiment.

5.3.4 Summary of Experimental Results

Overall, the experimental measurements did not reveal an obvious or reproducible difference between the background, hydrogen, and deuterium data. Background measurements established the expected detector response arising from environmental background, cosmic events, and electronic noise. The hydrogen measurements remained generally consistent with this background response, while the deuterium measurements similarly failed to produce a clearly distinguishable signal.

The experimental results are therefore considered inconclusive with respect to the central hypothesis of this work. No definitive evidence of enhanced deuterium-deuterium fusion was identified, and the measured data do not presently support a quantitative determination of fusion enhancement relative to background.

Further quantitative analysis is required to more rigorously compare the different experimental conditions. In particular, statistical comparison of event rates, pulse characteristics, energy distributions, and measurement durations may provide greater sensitivity to differences that are not visually apparent in the representative data. These comparisons will be used to determine whether any statistically significant deviation from the measured background is present.

The possible reasons for the inconclusive experimental outcome, along with limitations of the apparatus, experimental procedure, cavity formation, detector sensitivity, and other relevant factors, are considered in the following discussion.

Chapter 6

Discussion

This work investigates the hypothesis that quantum cavity confinement can enhance deuterium-deuterium fusion rates through increased electron screening. To investigate this possibility, a combined theoretical and experimental approach was developed. The theoretical portion of this work considered how electromagnetic confinement may modify the electronic environment surrounding reacting nuclei, while the experimental portion developed a system capable of producing the required vacuum, plasma, cavity, and detector conditions needed to search for the resulting fusion products.

The theoretical results identify cavity dimensions at which measurable modifications to electron behavior may begin to occur. In particular, the stochastic electrodynamics (SED) model predicts that suppression of zero-point electromagnetic modes within sufficiently small cavities can reduce the equilibrium electron orbital radius and increase electron density near the nucleus. The strongest predicted changes occur as the cavity dimensions become progressively smaller, with measurable effects beginning below approximately 50 nm and substantially larger modifications predicted as the cavity approaches the nanometer and angstrom scales.

Experimental measurements were also performed using the completed system. Background, hydrogen, and deuterium measurements were collected using both neutron and charged-particle detectors. However, no clear or consistently reproducible difference was observed between the

experimental and background conditions. The experimental results therefore remain inconclusive. Importantly, this does not presently provide sufficient evidence either to confirm cavity-enhanced fusion or to conclude that cavity confinement does not affect the fusion rate.

The following sections discuss the interpretation of these measurements, the experimental assumptions and limitations that prevent a definitive conclusion, the extent to which the theoretical models can be interpreted, and the improvements required for a more direct test of the central hypothesis.

6.1 Interpretation of Experimental Results

Experimental measurements were performed using the apparatus and procedures described in the preceding chapters. The primary objective of these measurements was to determine whether the detector response changed when moving from background conditions to hydrogen and ultimately to deuterium experimental conditions. Both neutron and charged-particle measurements were used in order to provide separate methods for identifying possible deuterium-deuterium fusion products.

The neutron detector data were digitized and analyzed using the pulse-finding software described previously. Individual neutron-like pulses were identified within the recorded signals and could then be compared between experimental runs. The charged-particle measurements were collected using the DIAD detector and analyzed according to the measured particle energies. Additional analysis of the DIAD spectra was performed using MATLAB in order to directly compare the distributions measured under the different experimental conditions.

Overall, neither detector produced an immediately clear or consistently reproducible increase in fusion-related activity during the deuterium experiments. Individual neutron pulses and charged-particle events were observed throughout the measurements; however, similar events were also

present within the background data. The presence of an individual event is therefore not sufficient to identify that event as a fusion product.

The charged-particle measurements similarly did not show a clearly separated population of events that appeared only after deuterium was introduced. Comparison of the resulting spectra showed that the general distributions remained similar between the background and experimental measurements. Although differences were present between individual runs, these differences were not sufficiently large or reproducible to establish a systematic relationship between the deuterium condition and the measured detector response.

The experimental measurements should therefore be considered inconclusive rather than a negative observation of cavity-enhanced fusion. A null result would only provide strong evidence against the proposed enhancement mechanism if the experimental conditions required by the theoretical model were independently known to have been produced. Several of these conditions, particularly the deuterium surface concentration and cavity separation, were not independently verified in this work.

Nevertheless, the experiments provided an important demonstration of the complete experimental procedure. The apparatus was successfully used to establish the required vacuum environment, introduce hydrogen and deuterium, generate a plasma, form the experimental cavity, and collect both neutron and charged-particle detector data. The principal remaining challenge is therefore not whether these individual components can operate, but whether the conditions within the cavity can be characterized with sufficient precision to make the resulting detector measurements conclusive.

6.2 Experimental Considerations and Limitations

Several experimental limitations prevent the measurements presented above from providing a definitive test of cavity-enhanced fusion. These limitations primarily involve uncertainty in the

amount of deuterium available to fuse, the geometry and separation of the cavity, the timing of cavity formation relative to the plasma, and the ability to distinguish a small fusion-related signal from detector background.

6.2.1 Deuterium Surface Loading

One of the largest uncertainties in the present experiment is the amount and distribution of deuterium on the two cavity surfaces. The experimental procedure uses a deuterium plasma with the intention of producing atomic and ionized deuterium species that can interact with the silicon and aluminum surfaces. The goal is to obtain approximately monolayer-scale deuterium coverage on the surfaces before the cavity is formed.

The plasma operating conditions were selected based on the stable operating range of the available system and the expected production of dissociated and ionized deuterium species. However, the amount of deuterium remaining on either surface following plasma exposure was not independently measured. The experiments therefore cannot determine whether a complete monolayer was formed, whether only partial surface coverage was achieved, or how much deuterium was lost between plasma shutdown and cavity formation.

This uncertainty directly affects the interpretation of the fusion measurements. Even if cavity confinement strongly increased the fusion probability for an individual pair of deuterons, the total measurable rate would still depend on the number and spatial distribution of deuterium nuclei available within the cavity. A sufficiently low deuterium surface concentration could therefore produce no measurable signal even if the proposed screening mechanism were physically present.

For this reason, verification of the deuterium surface concentration is one of the most important next steps in the experiment. Until the loading is independently characterized, it is difficult to distinguish between an absence of cavity enhancement and an experiment in which an insufficient number of deuterium nuclei were present within the confined region.

6.2.2 Plasma Timing and Residual Effects

The presence of plasma, which is used to deuterate the cavity walls, introduces an additional timing constraint. After plasma exposure, the cavity must be formed sufficiently quickly that substantial deuterium is not lost through desorption, diffusion, or redistribution. At the same time, the cavity should not be formed so quickly that energetic ions, electrons, or other residual plasma effects remain capable of producing detector signals that could be mistaken for fusion-related events.

These competing requirements introduce uncertainty into the experimental procedure. A long delay may reduce the amount of deuterium remaining on the surfaces, while a very short delay could make it difficult to separate cavity-related fusion from effects associated directly with the plasma.

The ideal timing therefore depends both on the lifetime of deuterium on the surfaces and the decay time of the plasma environment. Neither quantity was independently measured in the present work. Future experiments should characterize both effects so that the cavity can be formed within a known time window in which sufficient deuterium remains while plasma-related signals have become negligible.

6.2.3 Cavity Separation and Geometry

Another major uncertainty is the actual separation and geometry of the cavity surfaces. The theoretical models presented in this work predict that the magnitude of the screening modification depends strongly on cavity size. The SED framework predicts the beginning of measurable modifications below approximately 50 nm, while substantially stronger effects are expected as the separation approaches the nanometer and angstrom scales.

As a result, uncertainty in cavity separation produces a direct uncertainty in the expected magnitude of the effect. If the surfaces remain substantially farther apart than intended, the system may never enter the length-scale regime in which the theoretical model predicts a measurable modification to electron behavior.

Surface roughness and local geometry also complicate the definition of a single cavity separation. The two cavity walls are not perfectly flat at every microscopic location. Local protrusions may contact first, while neighboring regions remain separated by larger distances. The effective cavity therefore likely consists of a distribution of local separations rather than a single uniform spacing.

The experimental system developed in this work provides a method for bringing the two surfaces together and controlling their relative position, but the final microscopic spacing was not independently characterized during the fusion measurements. Future work must therefore determine the actual separation and spatial distribution of cavity dimensions. This will allow the experimental geometry to be directly compared with the length scales predicted by the theoretical models.

6.2.4 Detector Sensitivity and Background

The detector systems also place limits on the strength of the conclusions that can be drawn from the present data. Both neutron and charged-particle measurements contain background signals. Environmental radiation, cosmic-ray interactions, electronic noise, and other detector responses can produce events even when the experimental system is not operating with deuterium.

This background is particularly important when searching for low fusion rates. If the number of fusion events is comparable to or smaller than the natural variation of the background, individual events cannot be uniquely attributed to the experiment. Instead, a statistically meaningful excess in event rate, energy distribution, or another detector characteristic must be observed over repeated measurements.

The measurements collected in this work did not show such a clear and reproducible excess. Longer collection periods, additional repeated trials, improved energy calibration, and more complete statistical comparison between experimental conditions would improve the ability to detect smaller changes. The MATLAB analysis of the DIAD spectra and the pulse analysis of the neutron

detector provide a basis for such comparisons, but substantially larger data sets would provide greater statistical sensitivity.

Taken together, these experimental limitations prevent the present detector measurements from uniquely determining whether cavity confinement affected the deuterium-deuterium fusion rate. Improved experimental control is therefore required before a null detector result can be interpreted as a meaningful rejection of the cavity-enhancement hypothesis.

6.3 Interpretation of Theoretical Results

The theoretical models presented in this work provide possible mechanisms through which cavity confinement could modify electron screening. However, the models should be interpreted within the assumptions used to construct them. The theoretical results demonstrate how fusion rates could be altered if cavity boundary conditions modify the electronic structure in the predicted manner; they do not independently prove that the required modification occurs in the physical experimental system.

6.3.1 Cavity-Induced Electron Screening

The primary mechanism investigated in this work is a reduction in the characteristic electron orbital radius under confined electromagnetic conditions. If the electron becomes more strongly localized near the nucleus, the local negative charge density between two reacting deuterons increases. This increased electron density provides greater screening of the Coulomb repulsion between the positively charged nuclei.

The effect on fusion probability can be large because nuclear tunneling depends exponentially on the effective Coulomb barrier. Even a comparatively modest reduction in the barrier can therefore produce a much larger relative change in the tunneling probability.

Within the SED framework, cavity confinement modifies the available electromagnetic modes surrounding the electron. The model predicts that when the cavity becomes sufficiently small to suppress modes near the characteristic orbital frequency, the equilibrium electron distribution shifts inward. This produces the characteristic length scale of approximately 45.6 nm identified in Chapter 3, with substantially stronger effects predicted at progressively smaller cavity dimensions.

This provides a direct theoretical connection between cavity geometry and fusion enhancement. Rather than increasing the kinetic energy of the deuterons, the proposed mechanism modifies the environment through which the nuclei interact.

6.3.2 Dependence on the Zero-Point-Energy Model

An important limitation of the SED result is its dependence on the role of the zero-point electromagnetic field in determining electron structure. The model assumes that the equilibrium behavior of the electron is influenced by its interaction with the surrounding electromagnetic zero-point field. Under this assumption, changing the available electromagnetic modes through cavity confinement can modify the equilibrium electron distribution.

If the zero-point field does not contribute to electron localization in the manner described by the SED framework, then this particular mechanism for reducing the electron orbital radius would not apply. The predicted screening enhancement should therefore be interpreted as conditional on the underlying physical assumptions of the model.

This distinction is important when comparing the theoretical and experimental results. A failure to observe the predicted enhancement would not immediately determine whether the cavity itself has no effect or whether the specific SED description of that effect is incomplete. A definitive experimental result would instead provide a test of the combination of the proposed physical mechanism and the experimental realization of the required cavity conditions.

6.3.3 Other Possible Cavity Mechanisms

The SED mechanism is not the only possible means by which a cavity could influence the electronic environment. The QED approach discussed in this work provides a related but more qualitative description in which electromagnetic boundary conditions modify photon-mediated interactions and the vacuum field surrounding the electron.

Although the QED treatment developed here does not provide the same complete numerical prediction as the SED model, it similarly suggests that changes to the electromagnetic mode structure can influence electronic behavior. Therefore, even if the specific zero-point interpretation used in the SED calculation does not accurately describe the physical system, cavity-induced modifications to electron screening may still occur through other mechanisms.

The theoretical work should consequently be interpreted as identifying plausible mechanisms and experimentally relevant length scales rather than establishing a unique explanation for cavity-enhanced fusion. Additional theoretical development would be required to determine whether the different descriptions converge to the same quantitative prediction for the experimental system.

6.4 Implications for Surface-Enhanced Fusion

A broader implication of this work is the possibility that some enhancements previously associated with metallic environments may depend strongly on surface geometry rather than exclusively on bulk material properties. Metallic surfaces naturally contain scratches, defects, grain boundaries, voids, and other microscopic features. Some of these structures can form confined geometries similar in principle to the engineered cavity investigated in this work.

If cavity-induced changes to electron screening occur, these naturally occurring surface features could create localized regions in which screening differs substantially from the bulk average. Fusion

enhancement could therefore occur only within a small fraction of the material while still producing a measurable change in the overall reaction rate.

This possibility may be especially relevant to previous experiments in which modifications to a metal surface, including scratching or roughening, were associated with changes in measured fusion behavior. In such a case, the surface modification may not simply increase the available area but could also change the distribution of microscopic cavity geometries.

The present work does not establish that this mechanism explains previous observations. However, it provides a possible connection between microscopic surface structure and enhanced electron screening that can be tested experimentally. Distinguishing between surface and bulk effects will require experiments that independently control and characterize the cavity geometry while holding the remaining material properties approximately constant.

6.5 What Can Be Concluded

The central question of this thesis is whether microscopic cavity confinement can increase low-energy deuterium-deuterium fusion rates through enhanced electron screening. Based on the theoretical and experimental work completed here, this question remains unresolved.

The theoretical models indicate that an enhancement is physically possible under the assumptions considered. In particular, if electromagnetic zero-point fields contribute to electron structure in the manner described by the SED framework, cavity confinement below a characteristic length scale can increase electron localization and screening. The QED discussion provides an additional possible connection between cavity-modified electromagnetic interactions and electron behavior.

Experimentally, however, the measurements do not provide definitive evidence of an enhancement. No clear and reproducible difference was identified between the background, hydrogen, and

deuterium measurements. At the same time, several experimental quantities required for a direct comparison with theory were not independently verified.

Most importantly, the amount and distribution of deuterium on the cavity surfaces remain uncertain, and the actual microscopic cavity separation was not directly characterized during the measurements. Because both quantities strongly affect the expected fusion rate, the absence of a measurable enhancement cannot presently distinguish between two possibilities: either cavity confinement does not produce a measurable fusion enhancement, or the experimental system did not fully realize the conditions required for the predicted enhancement.

This distinction is central to the interpretation of the present work. A null experimental result becomes a strong test of the theoretical hypothesis only after the conditions required by that hypothesis have themselves been verified. Until then, the experimental measurements remain inconclusive rather than negative.

6.6 Future Work

Future work should focus primarily on removing the experimental assumptions that currently prevent a direct comparison between theory and measurement.

The highest priority is the independent characterization of deuterium surface loading. A method should be developed to determine the amount and spatial distribution of deuterium remaining on both cavity surfaces following plasma exposure. This would determine whether approximately monolayer-scale coverage is actually achieved and would provide an estimate of the number of deuterium nuclei available within the cavity.

The cavity separation should also be independently measured. The theoretical predictions depend strongly on the cavity length scale, particularly near and below approximately 50 nm. Direct characterization of the cavity geometry would determine whether the experiment actually

reaches this regime and would allow detector measurements to be compared against a known cavity separation rather than an assumed one.

Further characterization of the plasma would also improve the experiment. Measurements of the electron-energy distribution, ion density, and relative populations of molecular, atomic, and ionized deuterium could provide a more complete understanding of the surface-loading process. The time dependence of deuterium retention following plasma shutdown should also be measured so that the optimal delay between plasma operation and cavity formation can be established.

Once the deuterium loading and cavity geometry are independently known, the experimental comparison itself should be repeated with increased statistics. Background, hydrogen, and deuterium measurements should be collected under otherwise equivalent conditions using consistent collection times and detector settings. Longer data collections and repeated trials would provide greater sensitivity to a small change in the detector response.

The resulting neutron and charged-particle data should then be analyzed quantitatively rather than relying primarily on visual comparison. Count rates can be normalized to measurement time, charged-particle energy distributions can be compared directly, and statistical tests can be used to determine whether any observed differences exceed the expected variation in the background.

At that point, the experimental result would become substantially more decisive. If a reproducible increase in fusion-related detector signals occurs after the deuterium concentration and cavity spacing have been independently confirmed, the result would provide evidence supporting a cavity-induced enhancement mechanism. Alternatively, if no enhancement is observed after the required experimental conditions are verified, the null measurement would place meaningful constraints on the proposed mechanism and could indicate that cavity confinement does not provide the predicted enhancement under those conditions.

6.7 Conclusions

This work addressed the discrepancy between theoretical predictions and experimentally observed rates of low-energy deuterium-deuterium fusion in metallic environments. The central hypothesis investigated was that enhanced fusion rates may arise from electron-screening effects within microscopic cavities on metal surfaces. To explore this possibility, theoretical models were developed alongside an experimental system designed to form confined cavity environments and detect potential fusion products.

The theoretical frameworks identify a possible physical connection between cavity geometry and electron screening. In particular, the SED model predicts that suppression of zero-point electromagnetic modes below a characteristic cavity separation can reduce the equilibrium electron orbital radius and increase screening. A characteristic length scale of approximately 50 nm or smaller was identified for the onset of measurable modification, with substantially larger effects predicted as the cavity approaches the nanometer and angstrom scales. These predictions are conditional on the assumptions of the theoretical models, particularly the role of the zero-point electromagnetic field in determining electron structure.

Experimental measurements were successfully performed using neutron and charged-particle detectors. However, no clear or reproducible increase in fusion-related detector activity was observed during the deuterium experiments relative to the corresponding background and hydrogen measurements. These results do not provide evidence sufficient to confirm cavity-enhanced fusion, but they also do not provide a definitive rejection of the proposed mechanism.

The primary difficulty is that several experimental assumptions required for comparison with theory remain unverified. The amount and distribution of deuterium on the cavity surfaces were not independently measured, and the microscopic separation of the cavity surfaces was not directly characterized during the fusion measurements. As a result, it is not yet possible to determine whether

the lack of an observed enhancement reflects the underlying physics or whether the experimental system did not fully reach the conditions required for the predicted effect.

Nevertheless, substantial progress has been made toward a controlled experimental test. The vacuum system, plasma system, cavity-forming mechanism, gas-handling system, neutron detector, charged-particle detector, and associated analysis methods were successfully developed and operated together. The present work also identifies the specific assumptions that must now be experimentally verified.

Whether microscopic cavities enhance low-energy deuterium-deuterium fusion therefore remains inconclusive. The next stage of this work is not primarily the construction of an entirely new experiment, but the characterization and refinement of the existing one. By independently measuring the deuterium surface concentration and cavity spacing, improving the statistical precision of the detector measurements, and repeating the experiment under controlled conditions, future work can determine whether the absence of an observed enhancement results from limitations in realizing the intended experimental environment or from the absence of the proposed physical effect.

Appendix A

Nuclear Strong Force Potential Models

The fundamental strong interaction is governed by quantum chromodynamics (QCD). However, a direct QCD treatment of two interacting deuterons is far beyond the scope of this work. Instead, the attractive nuclear interaction is represented using several simplified, spherically symmetric effective potentials. These potentials reproduce selected qualitative and quantitative features of the nuclear interaction while remaining sufficiently simple for numerical calculations of the Coulomb barrier and tunneling probability.

The discussion in this appendix is restricted to the interaction between two bare deuterium nuclei. Electron screening, lattice effects, nuclear deformation, and other factors are neglected. Consequently, the models presented here should not be interpreted as complete descriptions of the deuterium-deuterium interaction. Their purpose is to provide approximations for comparing how the short-range nuclear potentials and electron screening affect the calculated tunneling probability.

The *Mathematica* routines used to construct these potentials, combine them with the Coulomb potential, locate the classical turning points, and calculate the resulting WKB tunneling probabilities and theoretical cross-sections are provided in Appendix L.

Many of the variables used throughout this appendix are reused across the multiple nuclear models. The common numerical parameters adopted in this work are

$$V_0 = 35 \text{ MeV}, \tag{A.1}$$

$$R = 2.12562 \text{ fm}, \tag{A.2}$$

$$a = 0.50 \text{ fm}. \tag{A.3}$$

Here, V_0 is the magnitude of the adopted nuclear-potential depth, R is the characteristic nuclear length scale, and a is the Woods-Saxon surface-diffuseness parameter. The value of R is based on the measured root-mean-square charge radius of the deuteron. Using the charge radius as the effective range of the nuclear interaction is itself a modeling approximation: a root-mean-square charge radius does not represent a sharp physical boundary. Similarly, V_0 and a are phenomenological model parameters rather than independently measured intrinsic properties of a deuteron.

A.1 Square Well

The square-well potential is the simplest model considered in this work. It approximates the attractive nuclear interaction as a region of constant potential extending from the center of the target nucleus to a fixed interaction radius. The potential is written as

$$V_{\text{SW}}(r) = \begin{cases} -V_0, & 0 \leq r \leq R, \\ 0, & r > R. \end{cases} \tag{A.4}$$

This form is analogous to the finite square-well potentials commonly used in introductory quantum mechanics and nuclear physics. The plot of the nuclear potential modeled as a square well is seen in Figure A.1. If expressed in terms of the Cartesian coordinate x rather than r and reflected across the vertical axis ($x = 0$), this potential takes the symmetric shape of a standard 1D finite square well.

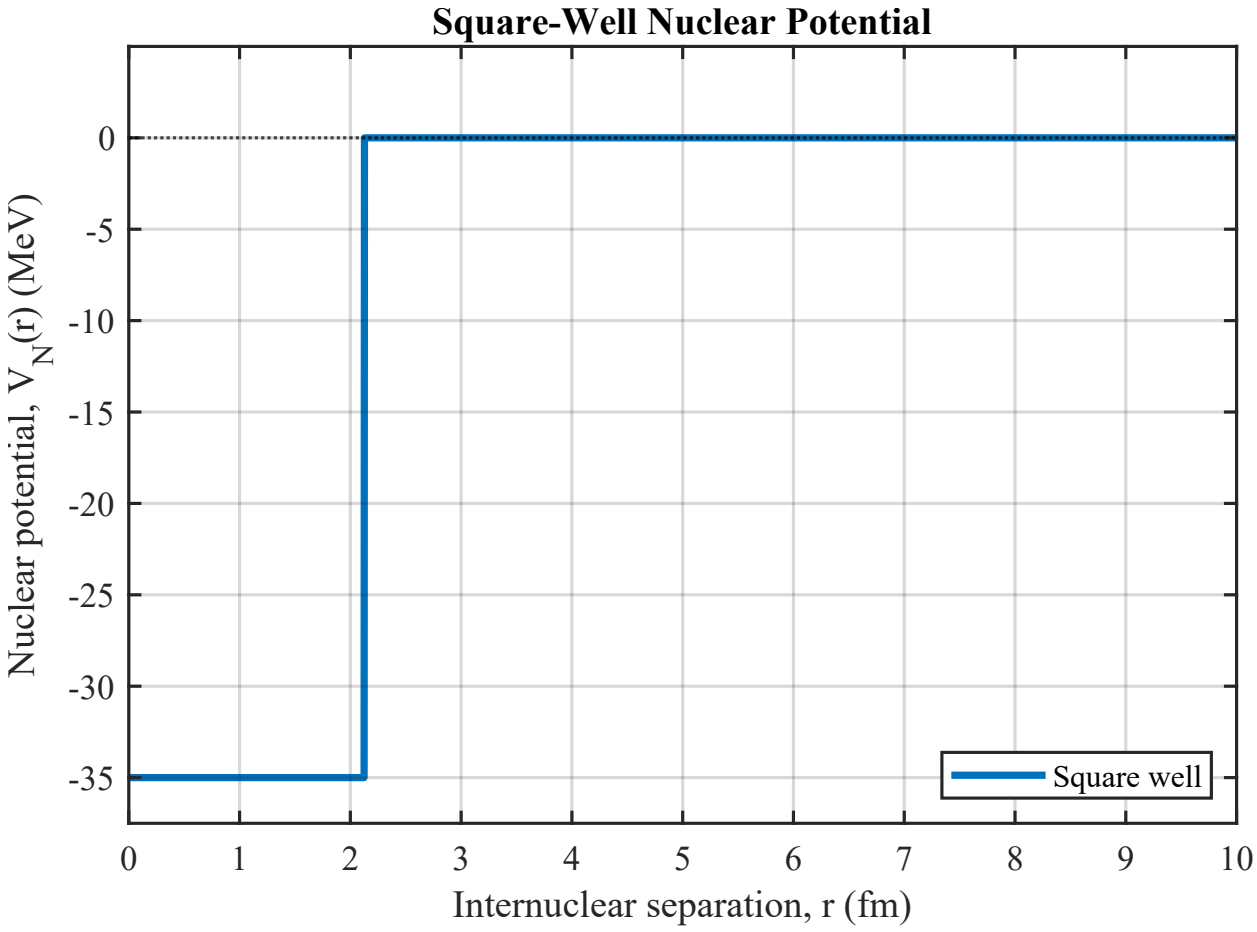

Figure A.1 Square-well approximation to the attractive nuclear potential. The nuclear term has a constant depth of $-V_0$ for $r \leq R$ and changes discontinuously to zero at $r = R$.

The principal advantage of the square-well model is its simplicity. The potential can be evaluated analytically, its parameters have an immediate geometrical interpretation, and it provides a useful baseline against which smoother potentials can be compared. It is therefore valuable for checking the numerical implementation of the tunneling calculation.

The principal limitation is the discontinuity at $r = R$. A physical nuclear density does not terminate at a perfectly sharp boundary, and the derivative of the square-well potential is undefined at the edge. Consequently, the model can produce an artificially abrupt inner boundary for the Coulomb barrier. Because the WKB tunneling probability is valid for a smoothly varying potential where the de Broglie wavelength changes slowly and it depends exponentially on the integral between the classical turning points, the calculated result may be highly inaccurate to real deuterium-deuterium tunneling potentials. The square-well model is therefore used as a limiting case rather than as the preferred potential for the calculations in this thesis.

A.2 Gaussian Potential

The Gaussian Potential model is extremely similar to that of the square well. Instead of a square well of defined edges and a constant bottom potential, the potential well is modeled using a parabolic shape. This smooth alternative is obtained by representing the nuclear attraction with a Gaussian function:

$$V_{\mathrm{G}}(r) = -V_0 \exp\left[-\left(\frac{r}{R_{\mathrm{G}}}\right)^2\right], \tag{A.5}$$

where R_{G} controls the radial range of the interaction. For the comparisons presented in this work, $R_{\mathrm{G}} = R$ is used. A plot of this function can be seen visually in Figure A.2.

The Gaussian potential acts much more like an actual nuclear potential compared to the square well. It eliminates the discontinuous boundaries of the square well and is straightforward to

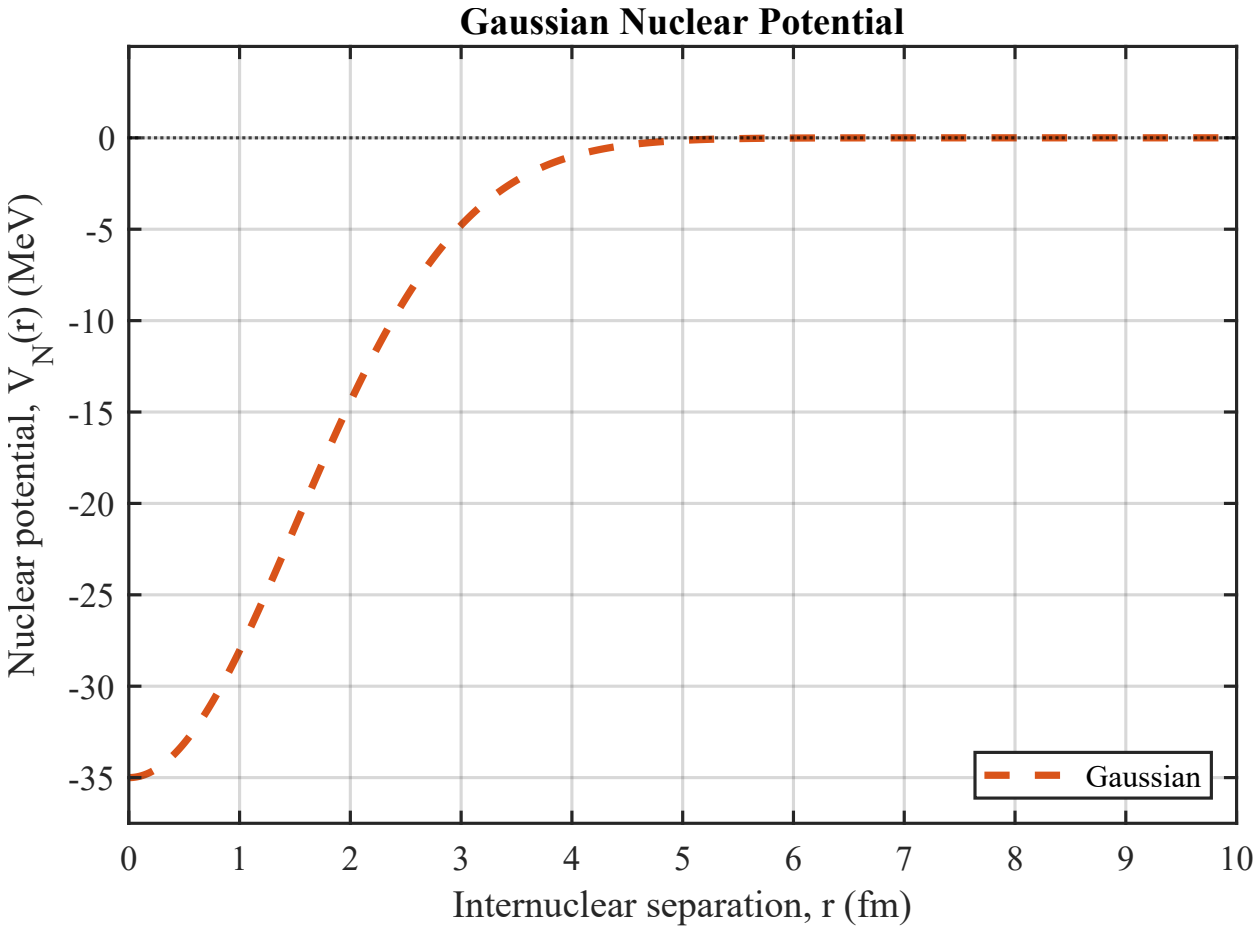

Figure A.2 Gaussian approximation to the attractive nuclear potential. The potential is finite at the origin and approaches zero smoothly at large separation.

differentiate and integrate numerically. These factors, along with its smooth behavior, generally improves the stability of numerical turning-point and WKB calculations.

However, the particular Gaussian form in Equation A.5 is not derived from the underlying deuteron-deuteron interactions. Its depth and range are selected or fitted to experimental values, but the exact characteristics of the tail and slope are arbitrary. Furthermore, similar to the square well, it excludes short-range repulsion and all spin-dependent effects. The Gaussian model is therefore best used as a more mathematically friendly and interpolated square-well model rather than as a quantitatively more accurate nuclear potential model.

A.3 Woods-Saxon Potential

The Woods-Saxon potential introduces a diffuse nuclear boundary while retaining a nearly constant potential within the nuclear potential well. It is given by

$$V_{\mathrm{WS}}(r) = -\frac{V_0}{1+\exp\left(\dfrac{r-R}{a}\right)}, \tag{A.6}$$

where V_0 is the potential-depth magnitude, R is the half-depth radius, and a determines the diffuseness of the nuclear surface. At $r = R$,

$$V_{\mathrm{WS}}(R) = -\frac{V_0}{2}. \tag{A.7}$$

This relation is easily seen in the Woods-Saxon potential, plotted in Figure A.3.

The Woods-Saxon form was originally developed as a diffuse-surface optical potential for nucleon-nucleus scattering [38]. It subsequently became widely used as a phenomenological mean-field potential in nuclear structure and reaction calculations.

The Woods-Saxon potential model combines several useful features of the preceding models. Similar to the square well, it has an approximately constant interior potential depth. Like the

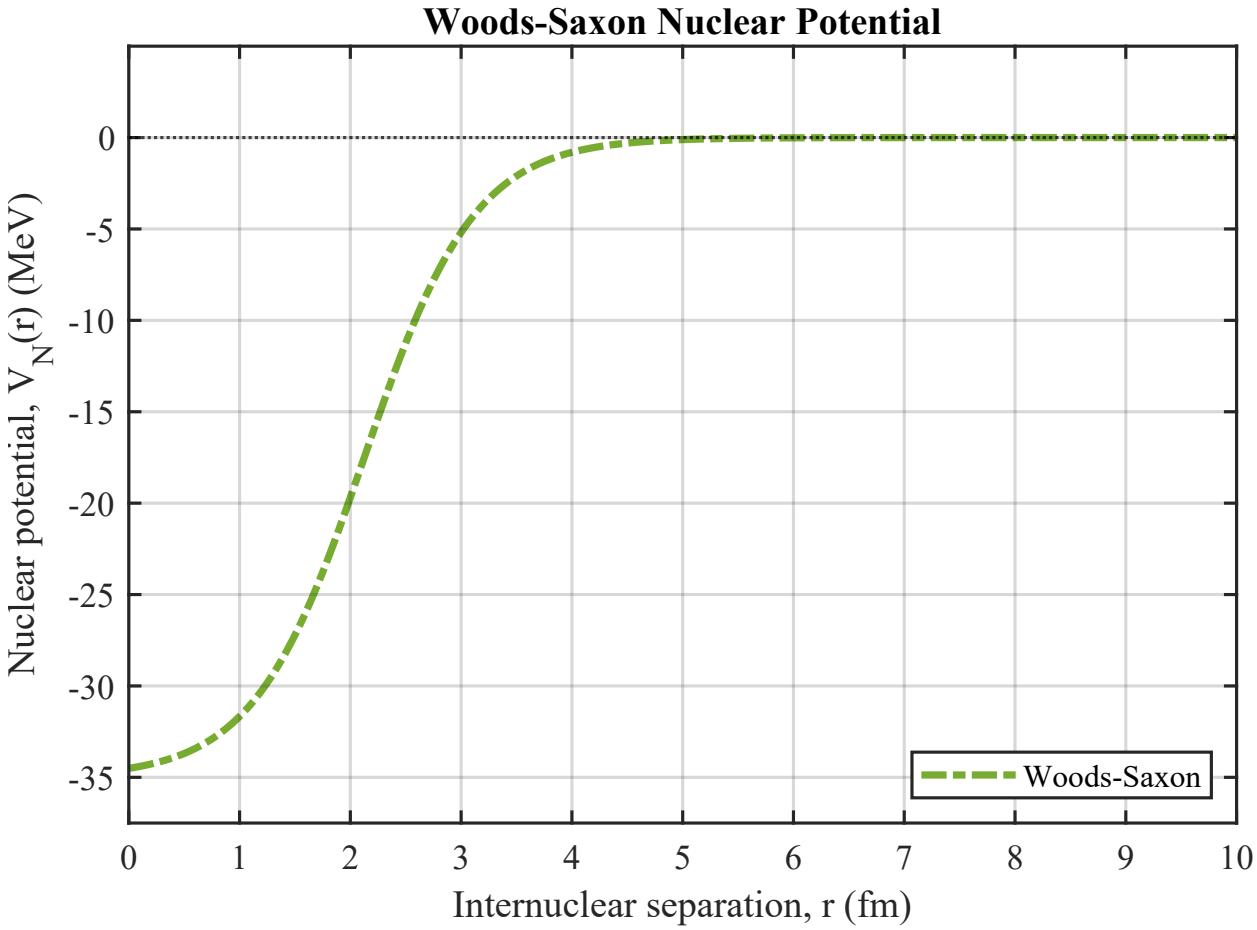

Figure A.3 Woods-Saxon approximation to the attractive nuclear potential. The diffuseness parameter *a* produces a smooth transition between the interior potential and the region outside the nucleus.

Gaussian potential, it changes smoothly and has no discontinuous surface. The varied parameters (V_0, R, a) also provide separate control over the depth, radial extent, and surface transition of the potential. These properties make it convenient for numerical evaluation of the combined Coulomb and nuclear potentials.

However, the Woods-Saxon potential is not a fundamental deuteron-deuteron interaction. Its original mean-field and optical-model interpretation applies most directly to the interaction of a nucleon with a larger nucleus. A deuteron is a weakly bound two-nucleon system without a well-defined, nearly uniform nuclear interior like that of a heavier, many-nucleon system. Therefore, an exhaustive deuteron-deuteron calculation would normally require an alternative treatment using a few-body model.

A.4 Yukawa Potential

The Yukawa potential provides an important description of a finite-range interaction generated by the exchange of a massive mediator. A simple attractive central Yukawa potential may be written as

$$V_{\mathrm{Y}}(r) = -A\frac{\exp(-r/\lambda)}{r}, \tag{A.8}$$

where A determines the interaction strength and λ is its characteristic range. In the original meson-exchange interpretation,

$$\lambda = \frac{\hbar}{m_{\mathrm{med}}c}, \tag{A.9}$$

where m_{med} is the mass of the exchanged particle. The Yukawa potential is plotted in Figure A.4.

The Yukawa model has greater historical physical motivation than the square well or the Gaussian interpolation because its exponential range follows naturally from massive-particle exchange. It therefore reproduces the qualitative expectation that the nuclear interaction should decrease rapidly outside the nuclear region.

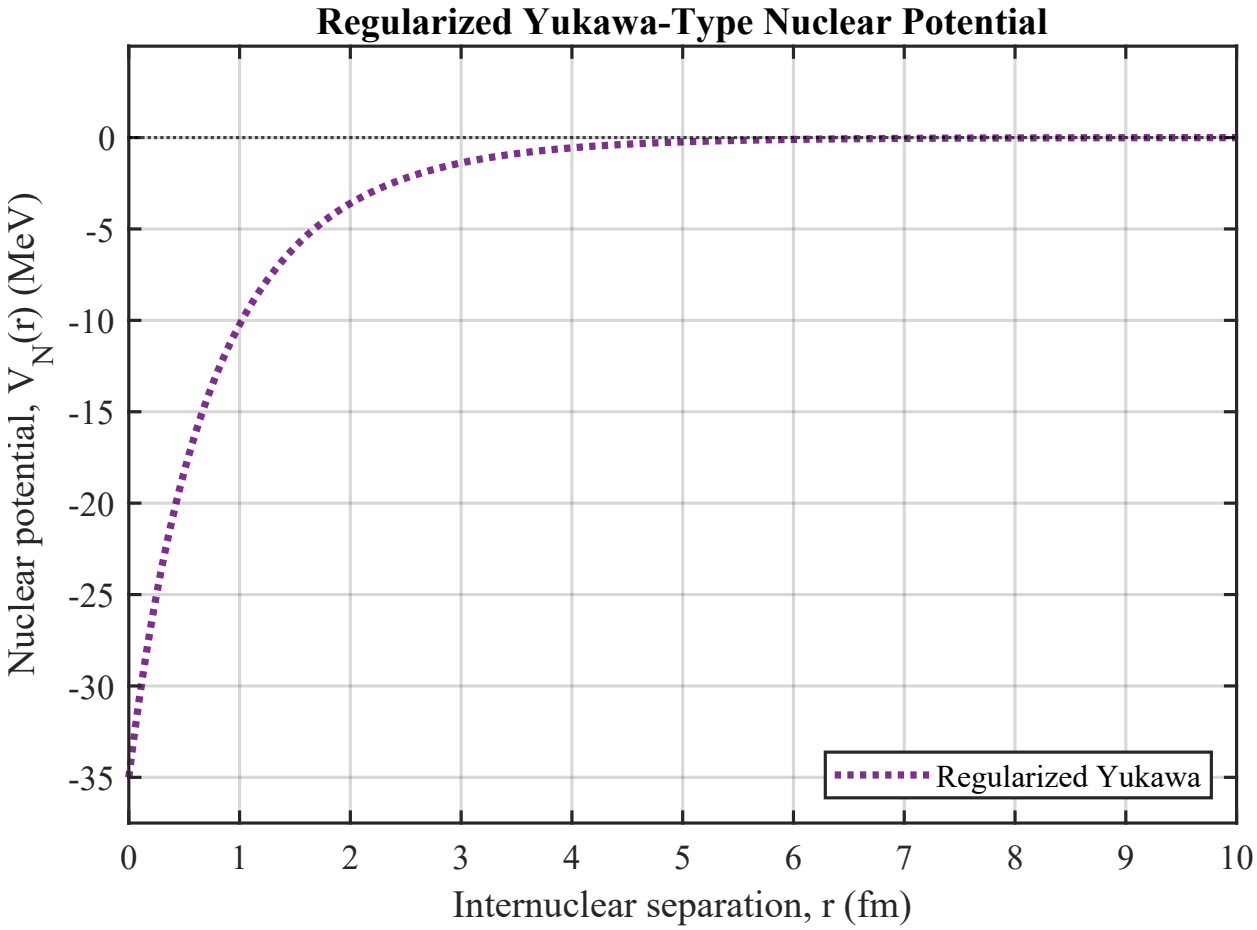

Figure A.4 Illustrative Yukawa-type nuclear potential. A short-distance regularization must be specified before the simple central form can be used over the complete radial domain.

However, the simple central form in Equation A.8 diverges as $r \to 0$ and must be regularized or supplemented with a short-range repulsive component. A single Yukawa term also omits many other nuclear interactions such as spin or isospin. Including more Yukawa term would not only make it more difficult to calculate the tunneling probabilities, but it would also require more guesswork for the various parameters and variables. Finally, the parameters A and λ in the first Yukawa term cannot be determined solely from the deuteron charge radius and the assumed Woods-Saxon depth.

A.5 Comparisons and Model Selection

The four potentials considered here represent different levels of mathematical and physical complexity. A side by side comparison of each model is plotted in Figure A.5 where many of the weaknesses and strengths of each can be seen. The square well provides the simplest baseline but introduces an unphysical discontinuity. The Gaussian potential removes this discontinuity but remains largely an empirical interpolation. The Yukawa potential incorporates the finite-range behavior expected from massive-particle exchange, but a single central Yukawa term is insufficient for a realistic short-range interaction. The Woods-Saxon potential provides a smooth finite-depth well with independently adjustable radius and surface diffuseness, making it the most convenient model for the numerical methods used in this work.

Accordingly, the Woods-Saxon potential is used in the calculations presented in the main body of this thesis. The resulting tunneling probabilities should be interpreted as predictions within this effective-potential approximation. A more complete treatment would test the sensitivity of the results to V_0, R, and a, compare multiple effective potentials, and eventually incorporate experimentally constrained deuteron-deuteron interaction models.

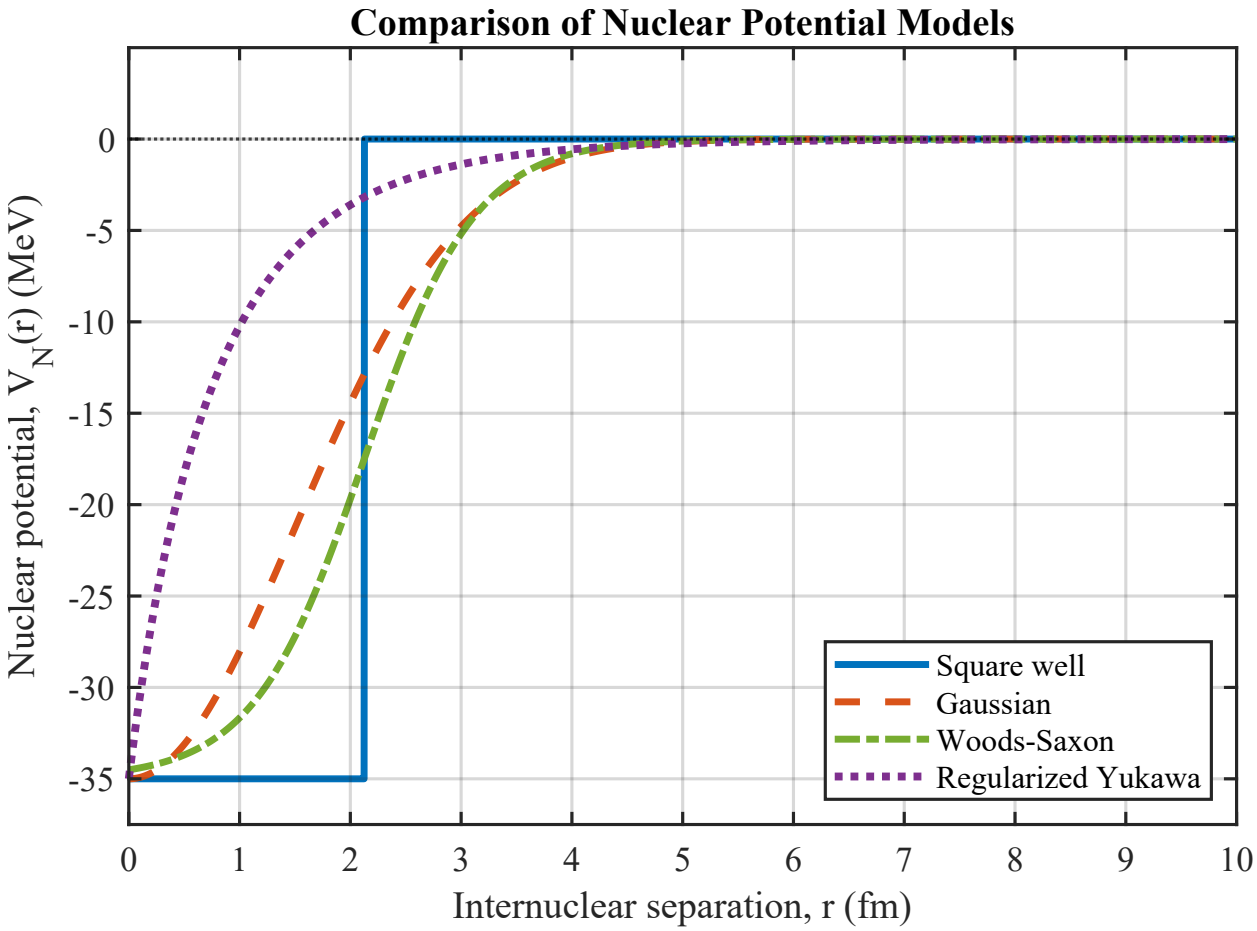

Figure A.5 Comparison of the square-well, Gaussian, Woods-Saxon, and regularized Yukawa-type models for the attractive nuclear potential. Each model uses a characteristic potential depth of 35 MeV, but differs in how the potential approaches zero with increasing internuclear separation. The solid black line represents $V_{\mathrm{N}}(r) = 0$. No Coulomb potential is included.

Appendix B

Vacuum Level During Testing

The pressure at which fusion testing should occur required the comparison of several competing physical effects. The two conditions considered in this work were the lowest vacuum pressure achievable by the system and a near-atmospheric pressure consisting almost entirely of deuterium. Both conditions provided possible advantages and disadvantages for preserving the loaded deuterium surface and limiting unwanted interactions during data collection.

The parameters that mattered most were the partial pressures of contaminating gases and the density and collision environment of the deuterium gas. Contaminants could alter the deuterium-loaded surface, increase the chamber background, or introduce additional particles into the detector measurements. Oxygen and water vapor were of particular concern because they can react with the silicon surface and contribute to the growth or modification of the native oxide layer.

The partial pressure of each contaminant is more important than its fraction of the total gas. A gas may have a very small impurity concentration while still producing a meaningful impurity partial pressure when the total pressure is large. The following sections compare these quantities for the high-vacuum and near-atmospheric testing conditions. The differences in the particle-collision environment are also discussed.

B.1 High-Vacuum Testing

For high-vacuum testing, the system was pumped to its lowest achievable pressure after the plasma-loading process was completed. The residual gas analyzer (RGA), discussed in Chapter 4, was then used to estimate the composition of the remaining gas.

A recreated RGA spectrum from a representative measurement is shown in Figure B.1. The primary contaminant signals were associated with water and oxygen, including peaks corresponding to H_2O, OH, O, and O_2. The combined pressure associated with these oxygen- and water-related peaks was estimated to be below approximately

$$2\ \mu\text{Torr} = 2 \times 10^{-6}\ \text{Torr}.$$

Because molecules may fragment during ionization within the RGA, the O and OH peaks do not necessarily represent independent gas species. They may instead result partly from the fragmentation of water or oxygen molecules. The summed value should therefore be understood as an approximate measure of the oxygen- and water-related background rather than an exact sum of independent partial pressures.

Although a partial pressure below 2 μTorr is small, it is not zero. At high vacuum, oxygen- and water-related molecules may make up a significant fraction of the residual gas even when their absolute pressures are low. These molecules may still collide with the target during the extended data-collection period and gradually alter the surface.

B.2 Near-Atmospheric Testing

For near-atmospheric testing, deuterium or hydrogen from a compressed-gas cylinder was introduced into the chamber after the plasma-loading process. The chamber was filled to a total pressure near

$$P_{\text{tot}} = 760\ \text{Torr}.$$

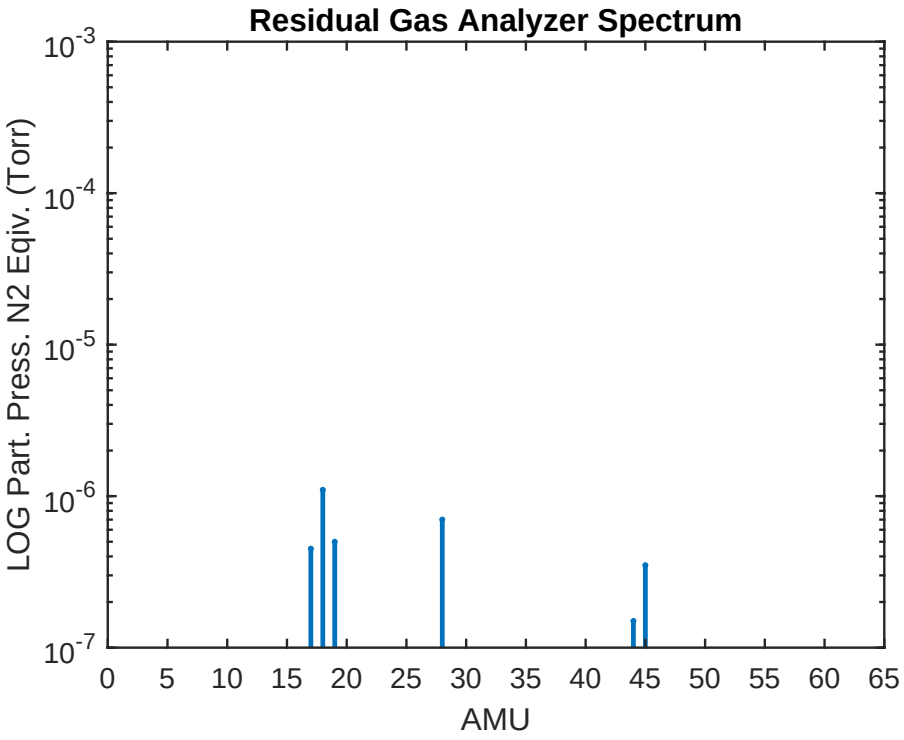

Figure B.1 Representative residual gas analyzer spectrum measured under the high-vacuum testing condition. The labeled peaks identify the main residual gases and fragments present within the chamber.

Any oxygen or water already remaining in the chamber at high vacuum was still present when the additional gas was introduced. Assuming a fixed chamber volume and no removal of the residual gas during filling, adding deuterium increases the total pressure but does not directly reduce the absolute partial pressure of the molecules already present. Additional contaminants were also introduced through the specified impurities in the gas supply.

The supplied gas was rated as ultra-high-purity gas. The maximum listed impurity concentrations are summarized in Table B.1. All impurity concentrations are given in parts per million by volume.

Table B.1 Maximum impurity specifications for the ultra-high-purity gas used to fill the chamber. Concentrations are given in parts per million by volume. The combined concentration of CO and CO_2 was specified not to exceed 1 ppm.

Purity (%)	O_2	H_2O	THC	CO	CO_2	Ar	N_2
99.999	1	2	0.5	1	1	–	5

Here, THC represents the total hydrocarbon concentration. The partial pressure of an impurity is calculated using

$$p_i = x_i P_{\text{tot}}, \tag{B.1}$$

where p_i is the impurity partial pressure, x_i is its volume fraction, and P_{tot} is the total chamber pressure. For a concentration c_i given in parts per million,

$$x_i = c_i \times 10^{-6}. \tag{B.2}$$

Using the maximum oxygen concentration of 1 ppm gives

$$p_{O_2} = \left(1 \times 10^{-6}\right)(760 \text{ Torr}) \tag{B.3}$$

$$= 7.6 \times 10^{-4} \text{ Torr}. \tag{B.4}$$

Similarly, the maximum water-vapor concentration of 2 ppm gives

$$p_{H_2O} = \left(2 \times 10^{-6}\right) (760 \text{ Torr}) \tag{B.5}$$

$$= 1.52 \times 10^{-3} \text{ Torr}. \tag{B.6}$$

The combined oxygen and water-vapor partial pressure introduced by the gas supply is therefore

$$p_{O_2} + p_{H_2O} = 2.28 \times 10^{-3} \text{ Torr} \tag{B.7}$$

$$= 2.28 \text{ mTorr}. \tag{B.8}$$

This value is much larger than the oxygen- and water-related pressure measured under high vacuum. The residual contamination originally present at high vacuum, below approximately 2×10^{-6} Torr, is small compared with the approximately 2.28×10^{-3} Torr introduced by the gas supply. Consequently, the impurity specifications of the supplied gas dominate the estimated oxygen and water partial pressures under near-atmospheric conditions.

Despite the larger absolute impurity partial pressures, the chamber remains overwhelmingly filled with deuterium. The near-atmospheric condition therefore provides a deuterium-rich environment with only ppm-level impurity fractions, whereas the high-vacuum condition contains much less gas but may have a larger relative fraction of residual water and oxygen.

B.3 Particle Energy and Collision Environment

The plasma source was turned off before data collection began. Therefore, particles were not continuously accelerated by the RF plasma during either testing condition. After the system stabilized, the gas particles were expected to have approximately thermal energies.

For an ideal gas at temperature T, the average translational kinetic energy is

$$\langle E_{\mathrm{k}} \rangle = \frac{3}{2} k_{\mathrm{B}} T. \tag{B.9}$$

This average energy depends primarily on temperature rather than pressure. At room temperature, it is on the order of a few hundredths of an electronvolt. Changing the pressure therefore does not directly give each molecule substantially more or less thermal energy.

Pressure does, however, strongly affect the number density and collision environment. The number density is given by

$$n = \frac{P}{k_{\mathrm{B}} T}. \tag{B.10}$$

At near-atmospheric pressure, the chamber contains a much larger number of deuterium molecules, and collisions between gas molecules occur frequently. The mean free path is correspondingly short, causing particles released from the target or chamber walls to undergo many collisions before traveling across the system.

Under high vacuum, the number density is much smaller and the mean free path is much longer. Molecules released from a surface may travel a substantial distance without colliding with another gas molecule. Thus, the main difference between the two pressure conditions is not the average thermal energy of each particle, but the frequency of particle collisions and the amount of gas available to interact with the target surfaces.

B.4 Pressure Comparison

The high-vacuum condition minimized the absolute partial pressures of oxygen, water vapor, and other residual gases. Based on the RGA measurements, the combined oxygen- and water-related pressure was below approximately 2 μTorr. This condition reduced the total number of gas

molecules available to interact with the target but did not completely eliminate surface contamination or chamber outgassing.

The near-atmospheric condition provided a much larger density of deuterium gas. However, the ppm-level impurities in the gas supply corresponded to an estimated combined oxygen and water-vapor partial pressure of approximately 2.28 mTorr. This was more than three orders of magnitude larger than the estimated oxygen- and water-related pressure under high vacuum.

Neither condition was therefore assumed to be ideal in every respect. High vacuum minimized the absolute contaminant pressure, while near-atmospheric testing placed the targets in an environment dominated by a large density of deuterium. Testing under both conditions allowed the possible effects of total pressure, deuterium density, residual contamination, and gas-surface collision frequency to be compared.

For this reason, measurements were performed under both high-vacuum and near-atmospheric conditions and were alternated with background-data collection. This approach reduced the need to assume beforehand which pressure condition provided the more favorable environment for preserving the loaded target surfaces or observing a possible fusion signal.

Appendix C

Plasma System

This appendix briefly describes the research and reasoning used to determine appropriate pressure and power conditions for a protium or deuterium plasma intended to produce approximately monolayer-scale hydrogen loading on the surfaces forming the cavity. In the experimental configuration used in this work, one cavity surface consists of silicon on the target side, while the opposing surface consists of an aluminum coating deposited on the detector.

Because the available plasma system was originally designed for plasma cleaning rather than controlled hydrogen or deuterium surface loading, the operating conditions required for this application were not directly specified by the manufacturer. The plasma parameters were therefore selected based on the established operating range of the system, the conditions required to dissociate and ionize molecular hydrogen or deuterium, and the expected interaction of the resulting species with the cavity surfaces.

C.1 Plasma Cleaning Background

The plasma cleaning system used in this experiment has a recommended pressure and power for operation with nitrogen. These operating conditions are specified in the Evactron system manual

and provide a useful baseline for establishing a stable low-pressure plasma within the experimental chamber. The preferred operating pressure is approximately 0.4 Torr, while the recommended forward power is approximately 12 W. The power referred to throughout this appendix corresponds to the measured forward power rather than the reflected power.

These values were originally selected for nitrogen plasma cleaning and therefore cannot be assumed to represent optimal conditions for hydrogen or deuterium loading. However, they demonstrate that the system can sustain a stable plasma within this pressure and power regime. The nitrogen operating conditions were consequently used as a starting point when determining appropriate parameters for protium and deuterium plasma operation.

C.2 Ionizing Particles

Based on the established plasma operating conditions, two primary processes were considered when selecting the protium and deuterium plasma parameters. First, sufficient electron energy must be present within the plasma to ionize the gas. Second, because the supplied protium and deuterium gases are predominantly molecular, sufficient energy must also be available to dissociate the molecules and produce atomic hydrogen or deuterium species.

These two processes are important because the desired surface interaction is not expected to depend solely on intact molecular hydrogen or deuterium. Dissociation of the molecular gas produces atomic species that can more readily interact with and adsorb onto the exposed silicon and aluminum surfaces. Ionization additionally produces charged species that may be accelerated through local plasma potentials and directed toward the surfaces.

The characteristic energies associated with molecular dissociation and ionization were therefore considered when determining whether the available plasma power was sufficient to produce the desired species. Although the externally applied RF power does not correspond directly to the

kinetic energy of an individual electron or ion, it determines the overall plasma conditions and electron-energy distribution from which these processes occur.

For the experiments presented in this work, plasma operation was tested using 0.4 Torr and 12 W. However, stable plasma behavior was observed under a much broader range of pressures and powers. These values were selected to maintain a stable discharge while providing sufficient energy for the production of dissociated and ionized protium or deuterium species.

C.3 Monolayer Creation

Using the plasma conditions described above, it was assumed that the resulting environment could produce the atomic and ionic species necessary for hydrogen or deuterium adsorption onto the cavity surfaces. The objective was not to produce a thick hydride or deuteride layer, but instead to provide sufficient surface coverage such that approximately monolayer-scale concentrations could be present when the two cavity surfaces were brought together.

During plasma exposure, molecular protium or deuterium within the chamber is partially dissociated and ionized. These species interact with the silicon and aluminum surfaces through adsorption, implantation, and other surface-interaction processes. After the plasma is extinguished, a fraction of the hydrogen or deuterium is expected to remain associated with the surfaces. The cavity is then formed by bringing the two surfaces into close proximity, ideally trapping deuterium on opposing cavity walls within the confined region.

The intended process can therefore be summarized as follows. Molecular deuterium is introduced into the chamber and a plasma is established under the selected pressure and power conditions. Electron collisions within the plasma produce atomic and ionized deuterium species. These particles interact with the silicon target and aluminum-coated detector surfaces, producing surface loading.

The plasma is subsequently extinguished, after which the two surfaces are brought together to form the cavity.

An important timing consideration is the period between plasma shutdown and cavity formation. Cavity formation must occur sufficiently soon after plasma exposure that substantial desorption, diffusion, or redistribution of the deuterium does not occur. At the same time, sufficient time must be allowed for energetic plasma particles and residual electrical effects to decay so that any subsequently measured detector response cannot be attributed directly to the active plasma. The experimental procedure therefore attempts to balance retention of surface-bound deuterium against removal of plasma-related effects.

C.4 Limitations and Future Work

Although the operating conditions used in this work were selected using the available physical and experimental information, significant uncertainty remains regarding the formation of a controlled deuterium monolayer using a low-density plasma. Most established plasma-processing literature focuses on surface cleaning, etching, implantation, or bulk hydrogen loading rather than the formation of a reproducible single monolayer under the specific pressure, power, material, and geometric conditions used in this experiment.

Consequently, the presence of exactly one monolayer of deuterium was not independently verified in this work. The plasma conditions were instead selected with the intention of producing sufficient dissociation and surface interaction to create substantial deuterium coverage on both cavity surfaces. The actual surface concentration may vary with plasma duration, pressure, forward power, surface cleanliness, material composition, surface roughness, temperature, and the time between plasma exposure and cavity formation.

Additional uncertainty arises from differences between the two cavity materials. Silicon and aluminum may exhibit different adsorption, diffusion, implantation, and desorption behavior for protium and deuterium. As a result, identical plasma exposure does not necessarily produce identical deuterium coverage on the two opposing cavity surfaces.

Future work should directly characterize the amount and distribution of hydrogen or deuterium remaining on each surface following plasma exposure. Surface-sensitive techniques could be used to determine whether monolayer-scale coverage is achieved and how that coverage depends on plasma pressure, power, and exposure time. Such measurements would allow the plasma parameters to be optimized experimentally rather than inferred primarily from plasma stability and known dissociation and ionization processes.

A more complete study should also characterize the plasma itself, including the electron-energy distribution, ion density, and relative populations of molecular, atomic, and ionized deuterium species. These measurements would provide a stronger connection between the externally controlled plasma parameters and the actual particle populations responsible for surface loading. Such characterization would substantially reduce one of the major uncertainties in determining the number and spatial distribution of deuterium nuclei available within the cavity.

Appendix D

Target Preparation and Cleaning

The targets used to form the experimental cavity in this thesis were silicon wafers. Because the experiment depends on bringing two surfaces into close proximity, the condition of the silicon surfaces was important. Particles, oils, organic residues, and other surface contaminants could interfere with the cavity spacing, increase outgassing within the vacuum chamber, or introduce impurities during deuterium gas loading.

Titanium targets were also used extensively during the broader development of the experiment. In particular, titanium was used while developing the hydrogen- and deuterium-loading procedures and while testing different surface-cleaning methods. Although the final cavity described in this thesis used silicon targets, the titanium work is included here because it contributed to the development and evaluation of the target-preparation procedure.

The same general cleaning procedure was used for both silicon and titanium targets. This procedure combined distilled-water and isopropyl-alcohol rinsing, plasma cleaning, piranha solution cleansing, and a diluted sulfuric-acid treatment. The purpose of combining these methods was to remove several different types of contamination while providing a consistent preparation process for each target.

The procedure was not intended to produce an atomically clean surface. Instead, it reduced visible and organic contamination and limited differences between targets caused by handling or storage. The construction, operation, and physical basis of the plasma-cleaning system are discussed separately in Appendix E. This appendix focuses on the order in which the silicon and titanium targets were cleaned and handled as well as the cleaning procedure's effectiveness.

D.1 Importance of Target Cleaning

Surface contamination can affect multiple parts of the experiment. Oils and organic residues can block portions of the target surface, alter the local interaction between the target and its surrounding environment, and introduce carbon-containing material into the vacuum chamber. During electrochemical loading, contamination can also interfere with the current flow and result in uneven hydrogen or deuterium loading.

For the silicon targets, particles and residue were especially important because the targets formed the two sides of the experimental cavity. Material trapped between the surfaces could prevent the targets from reaching the intended separation or produce a cavity with an uneven spacing. Contaminants could also outgas after the targets were placed under vacuum and make it more difficult for the system to reach its desired base pressure.

For the titanium targets, contamination could interfere with electrochemical loading and introduce impurities onto or into the surface. Therefore, both materials were handled carefully, contact with the active surfaces was limited, and the cleaned targets were stored in isolated containers until use.

D.2 Target Materials

Two target materials were cleaned using this procedure: silicon and titanium. The silicon targets were the primary targets used in the cavity experiment described in Chapter 4. These targets formed the parallel surfaces between which the confined cavity was created.

The titanium targets were used during earlier loading and surface-preparation tests. Titanium readily absorbs hydrogen and deuterium under appropriate conditions, making it useful for developing the loading process. The titanium targets also provided a convenient material for evaluating stronger chemical-cleaning methods before those methods were considered for the silicon targets.

Although the two materials have different surface chemistry, the same cleaning sequence was used for both in this work. No claim is made that the procedure was independently optimized for either material.

D.3 Cleaning Procedure

The complete target-cleaning sequence was

1. hold the target in piranha solution to dissolve any large organic compounds;
2. rinse the target with distilled water then isopropyl alcohol;
3. plasma clean the target;
4. repeat the distilled-water and isopropyl-alcohol rinse;
5. chemically clean the target using diluted sulfuric acid;
6. repeat the distilled-water and isopropyl-alcohol rinse;
7. plasma clean the target a second time; and

8. store the target in a clean, isolated glass container until use.

This procedure was applied to both the silicon and titanium targets. Each stage was intended to remove contamination that may not have been completely removed by the preceding stage.

D.3.1 Piranha Solution Cleaning

The first chemical-cleaning step used a piranha solution to remove larger organic contaminants from the target surface. Piranha solution is a strong oxidizing mixture commonly used for removing organic residues from silicon and other laboratory surfaces. In this work, the solution was prepared and handled according to the laboratory chemical-safety procedure.

The target was placed into the piranha solution and allowed to remain submerged while the solution reacted with organic material present on the surface. This process was done for approximately 5-10 minutes or until no further gas was observed being released. This treatment was intended primarily to remove oils, residues, and other carbon-containing contamination that could remain from fabrication, storage, or previous handling.

Following the piranha treatment, the target was removed and thoroughly rinsed with distilled water before proceeding to the isopropyl-alcohol rinse described below. The piranha treatment was used as an initial cleaning step rather than as a method for producing an atomically clean surface. Its purpose was to reduce substantial organic contamination before the target underwent the subsequent solvent, plasma, and diluted sulfuric-acid cleaning stages.

D.3.2 Initial Water and Solvent Rinse

Each target was first rinsed with distilled water and then with isopropyl alcohol (IPA). The distilled water removed loose particles and water-soluble residue, while the IPA helped remove oils and other organic material from the surface.

During the IPA rinse, the target was held at an angle so that the remaining liquid collected along one edge. That edge was then brought into contact with a Kimwipe to draw away the liquid without wiping directly across the active surface. The IPA was allowed to evaporate completely before the target was placed into the plasma-cleaning chamber.

The targets were handled at their edges whenever possible. Direct contact with the central silicon or titanium surfaces was avoided in order to reduce the introduction of oils, particles, or other contamination.

D.3.3 First Plasma Cleaning

Following the initial rinse, the target was placed near the center of the plasma-cleaning chamber. The chamber was closed and pumped down. Gas was then introduced until the pressure stabilized at approximately

0.4 Torr.

The plasma was operated at approximately

12 W

for one to two hours, depending on the size of the material and the expected level of contamination.

This stage was used primarily to remove organic material remaining after the water and IPA rinse. Once the plasma treatment was complete, the plasma and gas flow were turned off, the chamber was returned to atmospheric pressure, and the target was removed. The target was then rinsed again with distilled water and IPA.

The detailed design and operation of the plasma-cleaning system are discussed in Appendix E.

D.3.4 Diluted Sulfuric-Acid Cleaning

After the second water and IPA rinse, the target was chemically cleaned using a solution consisting of 80% water and 20% sulfuric acid by volume. The solution was prepared according to the laboratory procedure for diluting strong acids.

The target was placed into the diluted sulfuric-acid solution for 10 minutes. During this period, the target was gently agitated every 30–60 seconds. The agitation helped move fresh solution across the surface and prevented removed material from remaining in one location.

After 10 minutes, the target was removed from the solution and rinsed again with distilled water followed by IPA. The IPA was allowed to evaporate completely before the target proceeded to the final plasma-cleaning stage.

The sulfuric-acid solution could be stored and reused approximately four to six times. Used solution was placed in a labeled, chemically compatible container and disposed of according to the laboratory chemical-waste procedure.

D.3.5 Final Plasma Cleaning and Storage

After chemical cleaning and rinsing, the target underwent a second plasma treatment using the same general procedure as the first treatment. This final plasma clean removed organic contamination that may have been introduced during the chemical-cleaning, rinsing, or handling stages.

Following the final plasma treatment, the target was placed into a clean, isolated glass container, preferably one with a lid. The container limited exposure to dust, oils, and accidental contact while the target was waiting to be installed into the loading apparatus or experimental cavity.

The targets were not assumed to remain completely unchanged during storage. Both silicon and titanium can develop or modify surface oxide layers when exposed to air. The purpose of the container was therefore to reduce unnecessary contamination rather than to preserve an atomically clean surface.

Appendix E

Plasma Cleaning System

A separate plasma-cleaning system was constructed in order to provide a repeatable method for removing surface contamination from the titanium and silicon targets used throughout this work. The system was designed around an available Evactron plasma-cleaning unit and a small vacuum chamber assembled specifically for target preparation.

The purpose of this system was not to produce an atomically clean surface or maintain the targets continuously under vacuum. Instead, the goal was to remove oils, organic material, dust, and other contaminants that remained after the chemical and solvent-cleaning procedures. Because the targets were subsequently handled and transferred through atmosphere, some re-adsorption and surface oxidation were unavoidable. The plasma-cleaning system was therefore used as one part of a larger cleaning procedure rather than as a method for preserving an ideal surface indefinitely.

The complete target-preparation sequence is discussed in Appendix D. This appendix instead focuses specifically on the construction, operation, design choices, and verification of the plasma-cleaning system.

E.1 Constructing the Chamber

The chamber used for plasma cleaning centers around a four-way ConFlat cross. The cross provided a compact vacuum chamber while allowing separate ports for gas control, pumping, observation, and returning the system to atmospheric pressure.

On the left side of the chamber, an adjustable valve opens the system to atmosphere. This valve allows the chamber to be slowly returned to atmospheric pressure after plasma cleaning so that the targets can be removed without rapidly venting the system.

The front side of the cross contains a viewport. This window provides direct visual access to the interior of the chamber and serves two purposes. First, it allows the placement and condition of the targets to be observed before the chamber is closed. Second, it allows the plasma itself to be viewed during operation. The color, intensity, and spatial distribution of the plasma provide a simple visual indication that a discharge has formed and that it extends through the region containing the target.

The right side of the cross contains both the gas-input assembly and the connection to the roughing pump. The gas-input components were supplied as part of the Evactron system. Gas flow into the chamber is controlled using a needle valve, allowing the operating pressure to be adjusted gradually. The Evactron unit also controls the valve connecting the plasma source to the chamber and only permits operation when the chamber is within an appropriate pressure range.

The roughing pump provides the vacuum required for plasma operation. Because the plasma cleaner operates at pressures much higher than the base pressure of the primary experimental vacuum chamber, a high-vacuum turbomolecular pumping system was not required. The simpler roughing-pump arrangement was sufficient to reach and maintain the approximately 0.4 Torr operating pressure used for cleaning.

The completed plasma-cleaning system is shown in Figure E.1.

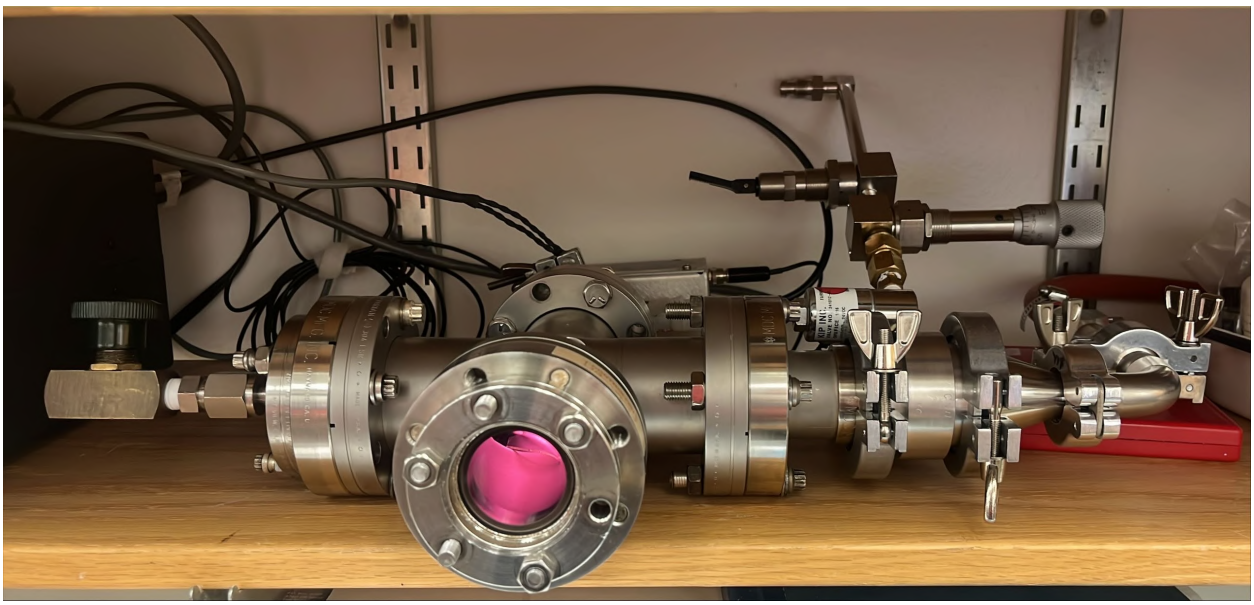

Figure E.1 Photograph of the completed plasma-cleaning chamber and associated vacuum components. The system is centered around a four-way ConFlat cross and includes a viewport for observing the plasma and targets, gas-flow controls connected to the Evactron plasma source, a roughing-pump connection, and a valve used to return the chamber to atmospheric pressure following cleaning.

E.2 Operating Conditions and Design Choices

The plasma-cleaning system was operated using the pressure and power range recommended for the Evactron unit. The preferred operating pressure was approximately

$$0.4 \text{ Torr}, \tag{E.1}$$

with an applied forward power of approximately

$$12 \text{ W}. \tag{E.2}$$

These conditions were originally specified for nitrogen plasma cleaning and provided a reliable starting point for the target-cleaning system.

The selected pressure represents a compromise between maintaining sufficient gas density for frequent electron-neutral collisions and maintaining a low enough pressure for electrons to gain energy between collisions. At substantially higher pressures, the mean distance between collisions decreases and the electrons may not acquire sufficient energy between collisions to efficiently sustain the discharge. At substantially lower pressures, there may not be enough neutral gas present to maintain a stable plasma using the available power supply and chamber geometry.

The Evactron system was useful because it combined plasma generation with control of the gas-input conditions. Rather than manually attempting to ignite and maintain a plasma over a wide range of pressures, the chamber could first be pumped down and then slowly filled through the needle valve until the appropriate operating regime was reached. The Evactron unit could then sustain the discharge at the selected forward power.

The relatively low power was also desirable for the targets used in this work. The goal was to remove contamination from the surface rather than significantly heat, sputter, or otherwise modify the titanium or silicon itself. A stable low-power plasma therefore provided a method for

exposing the surface to reactive plasma species for extended periods while limiting more aggressive modification of the target.

E.3 Plasma-Cleaning Procedure

For plasma cleaning, a target was placed near the center of the chamber where it would be exposed to the visible plasma. The chamber was then closed and pumped down using the roughing pump.

Once a sufficiently low pressure was reached, cleaning gas was introduced through the Evactron gas-input assembly. The needle valve was adjusted until the pressure stabilized near approximately 0.4 Torr. The plasma source was then operated at approximately 12 W forward power.

Targets were typically exposed to the plasma for approximately one to two hours, depending on the target size and expected level of contamination. For several of the cleaning-verification tests described below, the target was exposed for approximately 1.5 hours.

After plasma cleaning was complete, the plasma source and gas flow were shut off. The chamber was then slowly returned to atmospheric pressure using the vent valve, after which the target could be removed.

E.3.1 Atmospheric Exposure

The plasma-cleaning chamber was not designed as a load-lock or continuously evacuated target-transfer system. As a result, both the targets and the chamber were exposed to atmosphere before and after cleaning.

This atmospheric exposure was not considered a major limitation for the intended purpose of the plasma cleaner. The targets had already undergone water, solvent, and chemical cleaning before or after the plasma-cleaning stages, depending on their position within the complete cleaning sequence.

These procedures removed the majority of loose material and bulk contamination before the plasma was required to remove more strongly bound organic material.

The plasma-cleaning step was therefore primarily intended to remove contamination that remained attached to the surface after normal rinsing and chemical cleaning. It was not intended to prevent all subsequent adsorption of atmospheric molecules.

Additionally, oxidation of the titanium, silicon, or aluminum surfaces could not be completely avoided in the overall experimental procedure because the targets were later handled and transferred through atmosphere. Maintaining the target continuously under vacuum following plasma cleaning would therefore have provided limited benefit unless the remainder of the preparation and experimental system were similarly isolated from atmosphere.

For these reasons, the cleaning procedure focused on producing a consistently clean surface rather than an atomically clean or oxide-free surface. The combination of rinsing, chemical treatment, and plasma cleaning was considered sufficient for removing the types of contamination most likely to interfere with cavity formation or increase outgassing in the experimental vacuum system.

E.4 Cleaning Verification

The effectiveness of the plasma-cleaning system was evaluated by deliberately introducing visible contamination onto test targets and comparing the surfaces before and after plasma exposure. Fingerprints, oils, dust, and other normal handling contamination were placed on the target surface so that the ability of the plasma to remove these materials could be observed directly.

Microscope images were collected from several locations on the target before plasma cleaning. These locations are labeled A, B, C, and D in Figure E.2. The target was then plasma cleaned for approximately 1.5 hours using the standard cleaning conditions.

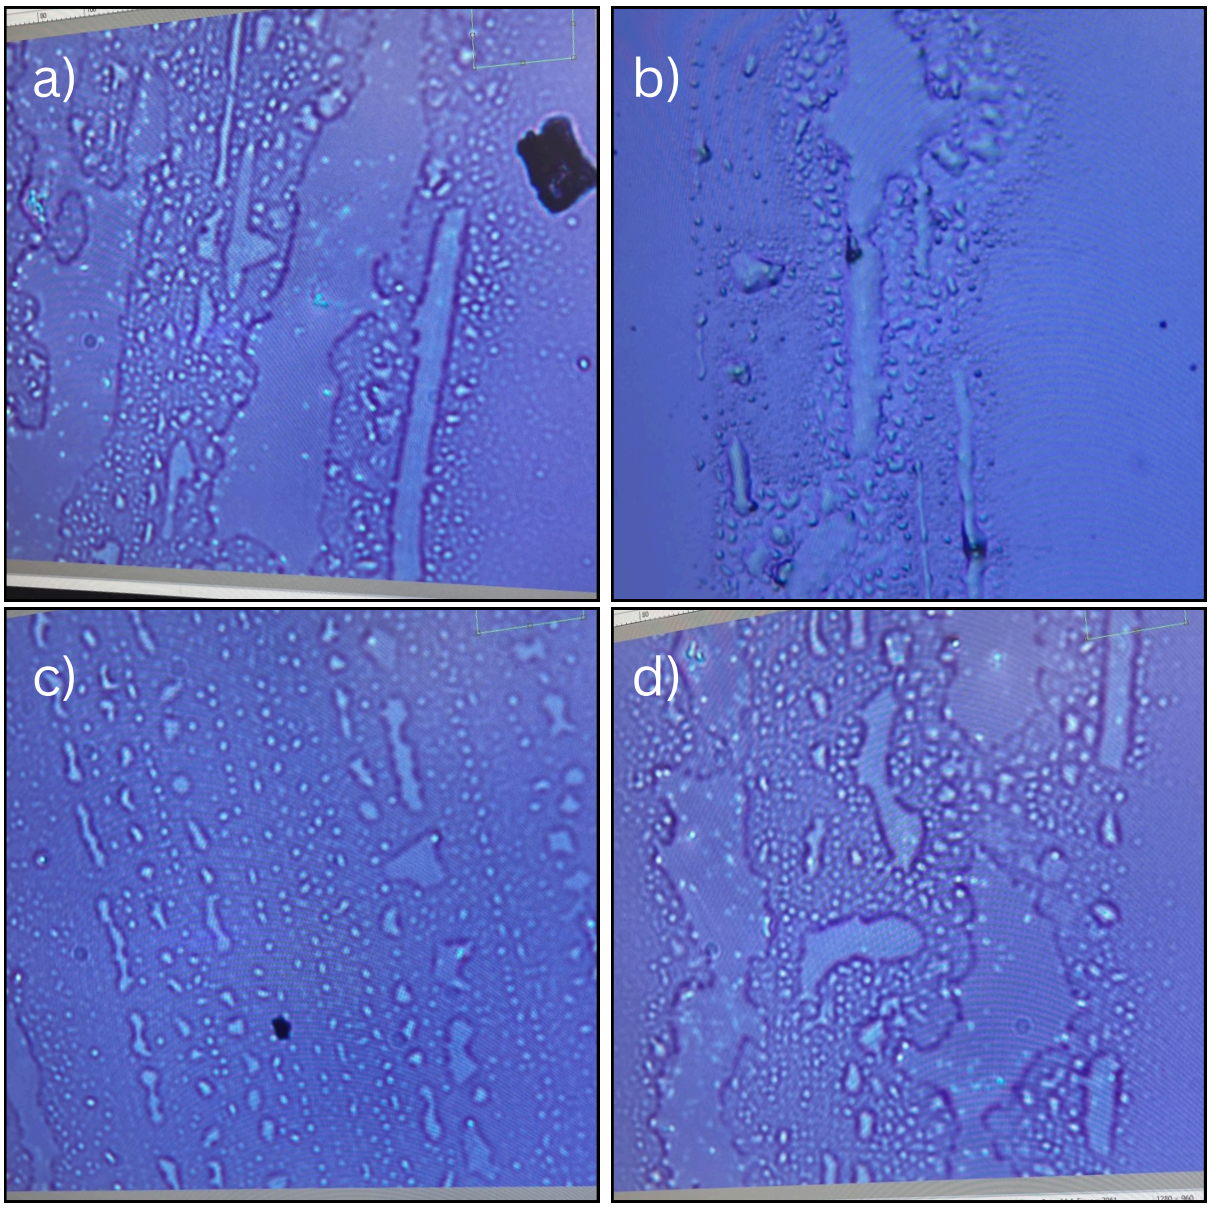

Figure E.2 Microscope images of selected regions of a deliberately contaminated target before plasma cleaning. Regions A, B, and C correspond directly to the same target locations shown after cleaning in Figure E.3. Visible oils, particles, and other surface contamination were intentionally present to evaluate the effectiveness of the plasma-cleaning procedure.

Following plasma cleaning, the same regions were imaged again. The corresponding post-cleaning microscope images are shown in Figure E.3. Regions A, B, and C correspond to the same approximate locations before and after cleaning, allowing the change in visible contamination to be compared directly. The exact corresponding location for region D in the original image could not be confidently identified and is therefore not used for direct comparison.

Comparison of the images shows a substantial reduction in visible oils and other contamination after plasma cleaning. The improvement was especially apparent in the regions containing intentional fingerprints and other organic material. Although some particles and surface features remained visible, the overall amount of removable contamination was dramatically reduced.

These measurements were primarily qualitative and were not intended to determine the remaining surface contamination at an atomic or molecular level. However, they demonstrated that the plasma cleaner was effective at removing the types of visible oils and organic contamination expected from normal target handling.

When combined with the distilled-water rinses, IPA rinses, piranha cleaning, and sulfuric-acid cleaning described in Appendix D, the plasma-cleaning results provided confidence that the targets were substantially cleaner than untreated material before being installed into the experimental system.

E.5 Use During Target Preparation

The plasma-cleaning system was incorporated directly into the complete target-preparation procedure described in Appendix D. In general, the targets first underwent bulk chemical and solvent cleaning to remove loose particles, oils, and other contaminants. The target was then plasma cleaned to remove organic material that remained attached to the surface.

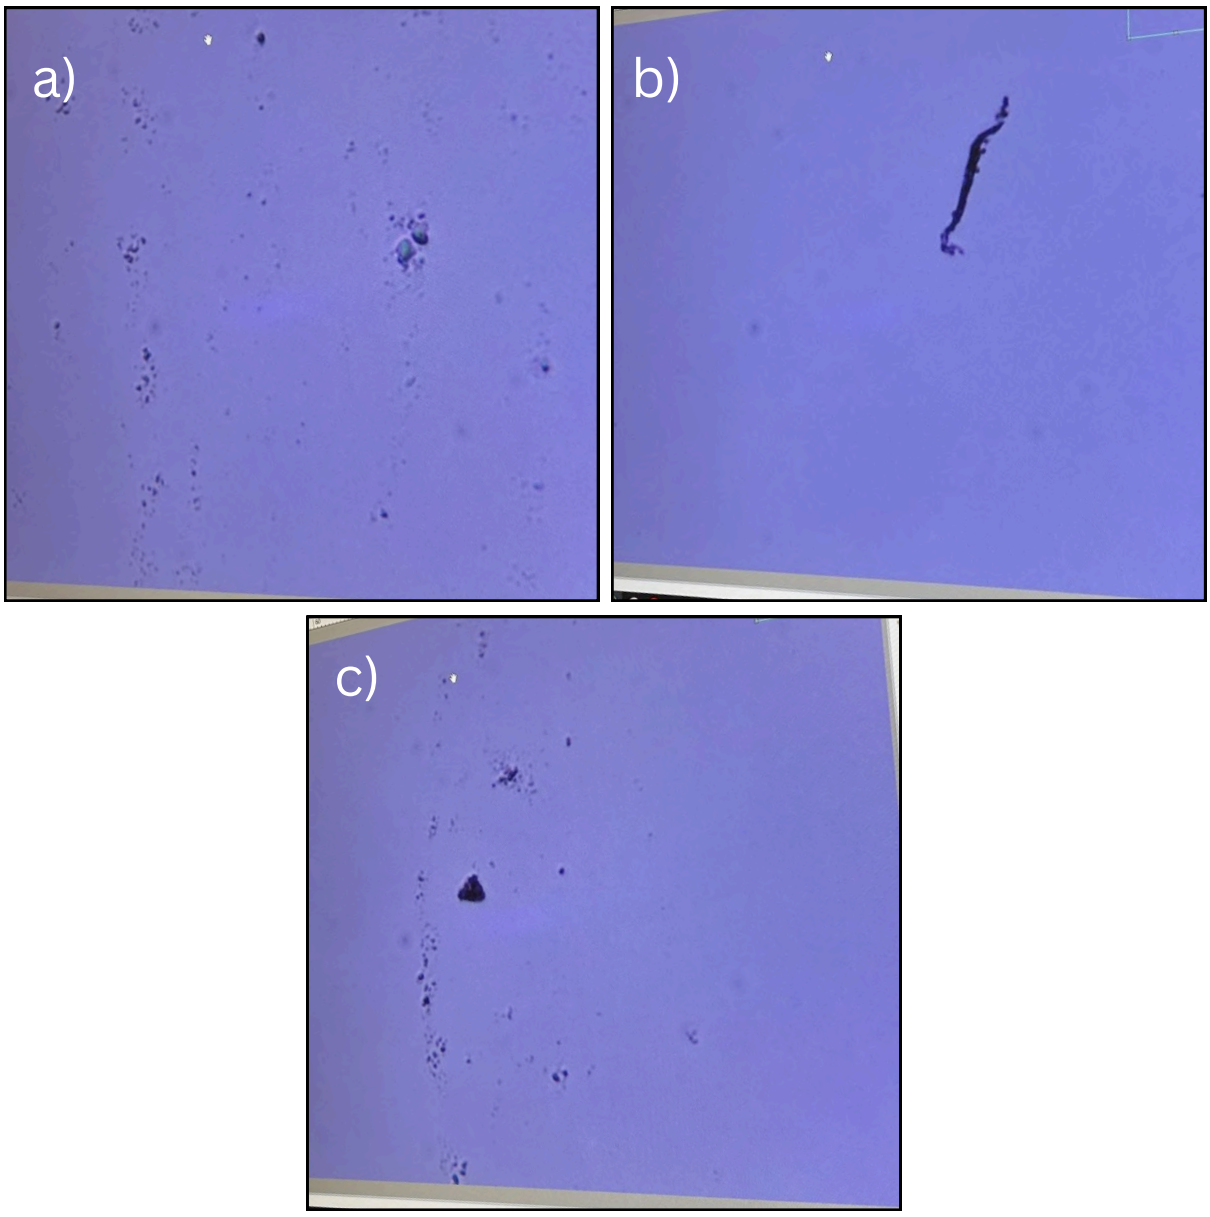

Figure E.3 Microscope images of the test target following approximately 1.5 hours of plasma cleaning. Regions A, B, and C correspond to the similarly labeled locations in Figure E.2. A substantial reduction in visible oils, particles, and other surface contamination can be observed following plasma exposure.

Additional chemical cleaning and rinsing were performed after the initial plasma treatment, followed by a final plasma-cleaning step before storage or use. The repeated plasma treatments reduced the possibility that contamination introduced during one of the intermediate chemical or handling stages would remain on the target.

Following the final plasma treatment, targets were stored in clean, isolated containers until they were required for gas loading or installation in the experimental cavity. Although the targets were necessarily exposed to atmosphere during this process, unnecessary handling and exposure to dust or oils were minimized.

The plasma-cleaning system therefore served as one component of a broader target-preparation procedure rather than as an independent cleaning technique. The combination of chemical cleaning, solvent rinsing, plasma exposure, and controlled handling provided the most practical method available in this work for producing consistently clean titanium and silicon target surfaces.

Appendix F

Hydrogen Gas Loading

In order for the experimental targets to be useful in studying hydrogen- or deuterium-related effects, it was necessary to develop a method for introducing hydrogen or deuterium into the target material. Several different methods were investigated throughout the development of this experiment. These methods included electrolysis using sulfuric acid, electrolysis using lithium hydroxide, and thermal gas infusion.

The majority of the initial development was performed using protium rather than deuterium. This allowed the loading procedure to be tested using less expensive materials before transitioning to deuterated chemicals. Once a loading procedure had been developed, the same general process was modified to minimize protium exposure and instead load the targets using deuterium.

An additional challenge throughout this process was determining whether the loading procedure had actually introduced a measurable quantity of hydrogen or deuterium into the target. Several verification methods were investigated, including measurements of target mass, mechanical changes, spectroscopy, and gas released from the target. These verification methods are discussed separately in Appendix G. This appendix instead focuses on the development and implementation of the loading procedures themselves.

F.1 Sulfuric Acid Electrolysis

The first method attempted to hydrogenate the titanium targets was electrolysis using a diluted sulfuric-acid solution. Several different acid concentrations, applied voltages, and exposure times were tested in an attempt to determine conditions that would introduce hydrogen into the titanium without significantly altering the target itself.

F.1.1 Sulfuric-Acid Concentration

The sulfuric acid available for these experiments was approximately 96-98% concentrated. For the initial concentration calculations, the sulfuric acid was approximated as 100% concentrated to simplify preparation of the solutions.

The first solutions were prepared in 100 mL baths. The amount of sulfuric acid required for a 0.2 M concentration was calculated as

$$(0.2 \text{ mol/L})(0.1 \text{ L}) = 0.02 \text{ mol}. \tag{F.1}$$

Using the molar mass and density of sulfuric acid,

$$(0.02 \text{ mol}) \left(\frac{98.079 \text{ g}}{1 \text{ mol}} \right) \left(\frac{1 \text{ mL}}{1.85 \text{ g}} \right) \approx 1.06 \text{ mL}. \tag{F.2}$$

Therefore, approximately 1.06 mL of concentrated sulfuric acid was required for every 0.2 M increase in concentration for a 100 mL solution. This resulted in approximately 4.24 mL of concentrated sulfuric acid for a 0.8 M solution and approximately 5.3 mL for a 1.0 M solution.

F.1.2 0.8 M Initial Test

The electrolysis setup was first tested using two small squares of titanium in a 0.8 M sulfuric-acid solution. One titanium square was connected by an alligator clip to the positive terminal of a DC

power supply and acted as the anode. The second titanium square was connected to the negative terminal and acted as the cathode.

Both pieces were submerged in approximately 100 mL of diluted sulfuric acid. The solution was contained in a beaker placed on a hot plate with a magnetic stirrer. The stirrer was operated at a very low rate simply to maintain motion within the solution.

The applied voltage was initially kept low and gradually increased until bubbles could be observed forming on the surfaces of both electrodes. This occurred at approximately 0.6 V. The electrolysis was then allowed to continue for approximately one hour.

During the experiment, bubble formation would occasionally decrease or stop. When this occurred, the applied voltage was increased slightly until visible bubbling resumed. Following the experiment, portions of the titanium surfaces appeared duller, but no other noticeable physical changes were observed.

F.1.3 1.0 M Test

A second experiment was performed using a 1.0 M sulfuric-acid solution. In order to more carefully evaluate whether hydrogenation was occurring, two new pieces of titanium were weighed before and after electrolysis.

The same general experimental setup was used, although the applied voltage was increased to approximately 1.2 V and was increased slightly later in the experiment. The measured masses before and after the experiment are shown in Table F.2.

The changes in mass suggested that processes other than simple hydrogen loading may have been occurring during the experiment. In particular, the opposing changes in the masses of the two pieces raised the possibility of material transfer or other electrochemical changes between the electrodes. The goal of the procedure was to introduce hydrogen into the titanium while minimizing

Table F.1 Titanium target masses before and after electrolysis in a 1.0 M sulfuric-acid solution.

Piece	Before	After
Larger $(-)$	0.101 g	0.100 g
Smaller $(+)$	0.070 g	0.071 g

alteration of the target itself. Therefore, a lower sulfuric-acid concentration and lower initial voltage were tested again.

F.1.4 Second 0.8 M Test

A second experiment was performed using a 0.8 M sulfuric-acid solution. The procedure was similar to the previous experiments but was divided into two stages. First, the titanium pieces underwent electrolysis at approximately 0.6 V for 30 minutes. The targets were then removed and weighed.

The targets were subsequently returned to the sulfuric-acid bath and underwent electrolysis for an additional 1.5 hours at a substantially higher voltage of approximately 15 V. The measured masses throughout this experiment are shown in Table F.2.

Table F.2 Titanium target masses during the second 0.8 M sulfuric-acid electrolysis test.

Piece	Before	30 min at 0.6 V	1.5 hr at 15 V
Larger $(-)$	0.100 g	0.101 g	0.101 g
Smaller $(+)$	0.070 g	0.070 g	0.070 g

The lack of any substantial additional mass change was useful because it suggested that significant material transfer between the electrodes was not occurring under these conditions. However, it also showed that measuring the target mass was not sufficiently sensitive to verify hydrogen loading.

Any hydrogen incorporated into the titanium was too small relative to the mass of the target and the precision of the available balance to be distinguished using this method.

Therefore, although the sulfuric-acid experiments demonstrated that electrolysis could be maintained and that gas formation occurred at the titanium surfaces, they did not provide a reliable method for confirming that the titanium had been successfully hydrogenated. In addition, the possibility of unwanted changes to the titanium during exposure to the sulfuric-acid solution motivated the investigation of a different electrolyte. The loading procedure was therefore transitioned to lithium hydroxide, as discussed in the following section.

F.2 Lithium Hydroxide Electrolysis

Following the sulfuric-acid experiments, lithium hydroxide was investigated as an alternative electrolyte for hydrogen loading. Lithium hydroxide provided a milder environment than the sulfuric-acid bath while still allowing hydrogen to be produced at the titanium surface during electrolysis. The goal was to obtain similar hydrogen loading while reducing deterioration or unwanted chemical alteration of the titanium target.

F.2.1 Lithium-Hydroxide Concentration

The initial lithium-hydroxide experiments used a 0.2 M solution. For a 100 mL bath, the amount of lithium hydroxide required for a 0.1 M concentration was calculated from

$$(0.1\ \text{mol/L})(0.1\ \text{L}) = 0.01\ \text{mol}, \tag{F.3}$$

and

$$(0.01\ \text{mol})\left(\frac{23.95\ \text{g}}{1\ \text{mol}}\right) = 0.2395\ \text{g}. \tag{F.4}$$

Therefore, approximately 0.2395 g of LiOH was required for every 0.1 M increase in concentration for a 100 mL solution. A 0.2 M solution was consequently prepared using approximately 0.479 g of LiOH.

F.2.2 Initial 0.2 M Test

The same general electrolysis arrangement used during the sulfuric-acid experiments was retained for the initial lithium-hydroxide test. Two titanium samples were placed into the 0.2 M LiOH solution and connected to the positive and negative terminals of the DC power supply.

The electrolysis was performed for approximately two hours while the applied voltage ranged from approximately 6 to 13 V. In order to determine whether hydrogenation could be observed through a change in target mass, both titanium samples were weighed before and after electrolysis. The results are shown in Table F.3.

Table F.3 Titanium target masses before and after the initial 0.2 M lithium-hydroxide electrolysis test.

Piece	Before	After
$(-)$	0.073 g	0.073 g
$(+)$	0.061 g	0.061 g

The results were again inconclusive. No measurable mass change was observed at the resolution of the available balance. Therefore, although electrolysis was occurring, the mass measurements could not determine whether a significant amount of hydrogen had been incorporated into the titanium.

F.2.3 Larger Titanium Target

Because the expected mass of hydrogen within the small titanium samples was very small, a larger titanium target was investigated in an attempt to produce a more easily measurable change. The electrolysis bath was increased from approximately 100 mL to approximately 1800 mL, and the small titanium squares were replaced with a substantially larger sheet of titanium.

The intention was to hydrogenate the larger sheet and subsequently cut smaller targets from the loaded material. By increasing the total quantity of titanium exposed to the electrolysis process, a larger absolute amount of hydrogen could potentially be incorporated into the material. This would increase the possibility of detecting hydrogenation through a measurable change in total target mass.

For a 0.2 M LiOH concentration in a 1.8 L bath, the required lithium-hydroxide mass was calculated as

$$(0.2\ \text{mol/L})(1.8\ \text{L})\left(\frac{23.95\ \text{g}}{1\ \text{mol}}\right) = 8.622\ \text{g}. \tag{F.5}$$

This larger-bath configuration allowed a substantially greater area of titanium to be exposed to the loading process. However, due to different students doing the various parts of the experiment, the before-and-after mass measurements were not able to be correctly taken and recorded for this larger target, so the success of the loading could not be determined from weight measurements alone.

F.2.4 Palladium Anode Tests

Later lithium-hydroxide experiments replaced the titanium anode with palladium foil while retaining titanium as the material being hydrogenated. Several experiments were performed using this configuration in an attempt to improve the repeatability of the electrolysis process.

Multiple tests were performed using the palladium anode, including additional tests in which the voltage was increased. However, complete numerical results for these experiments were not recorded. Consequently, these tests contributed to the development of the loading procedure but cannot be used here to provide a quantitative comparison with the earlier titanium-anode measurements.

F.2.5 Later Lithium-Hydroxide Test

A later loading attempt used a higher-concentration LiOH solution. This solution was at approximately 0.4 M. The titanium target had an initial mass of approximately 0.073 g.

During this experiment, the electrolysis was operated at approximately

$$0.02 \text{ A} \tag{F.6}$$

and

$$3.0 \text{ V}. \tag{F.7}$$

Overall, the lithium-hydroxide experiments provided a less aggressive alternative to the sulfuric-acid loading procedure and allowed stable electrolysis to be performed over a range of voltages, currents, bath sizes, and target configurations. However, the available weight measurements remained insufficient to independently verify the amount of hydrogen incorporated into the titanium.

The lithium-hydroxide procedure nevertheless became important in the development of the later deuterium-loading method. Because the same general electrolysis process could be performed using lithium deuteroxide and deuterium oxide, the procedure could be transferred from protium loading to deuterium loading while minimizing the introduction of ordinary hydrogen. This transition is discussed in Section F.4.

F.3 Thermal Gas Infusion

In addition to electrochemical loading, thermal loading was investigated as a method for introducing hydrogen or deuterium more deeply into the target material. The motivation for this approach was to improve the probability of gas being incorporated into the bulk of the metal rather than relying only on electrochemical loading near the surface.

A furnace was therefore investigated for use in the loading procedure. Before using the furnace for target loading, its temperature regulation was tested using an independent type-K thermocouple. The thermocouple was inserted into the center region of the furnace assembly with the thermocouple bead touching the wall of the metal cylinder. The furnace's indicated temperature was then compared with the independent thermocouple measurement.

Initial operation showed that simply setting the furnace to its automatic mode did not immediately produce the intended heating behavior. The furnace control and temperature-limit settings were therefore adjusted and tested first near 100 °C.

During this test, the furnace indicated 60 °C while the thermocouple initially measured approximately 39.3 °C. After approximately 13.5 minutes, the furnace display reached 100 °C while the thermocouple measured approximately 79.9 °C. After approximately 19 minutes, the thermocouple had increased to approximately 94 °C.

The furnace was then tested at a higher setting near 200 °C. At this temperature, significantly better agreement was obtained between the furnace display and the independent thermocouple. After approximately three minutes, both measurements were near 120 °C. After approximately eight minutes, both were near 180 °C, and after approximately 11 minutes the furnace displayed 200 °C while the thermocouple measured approximately 202 °C.

These measurements confirmed that the furnace could reach and regulate temperatures in the range required for later thermal-loading experiments. They also demonstrated the importance of

allowing the system sufficient time to equilibrate and independently verifying the target-region temperature rather than relying only on the furnace display.

The available development notes primarily document the furnace calibration rather than a complete finalized thermal gas-infusion procedure. Therefore, the exact gas pressure, loading temperature, and exposure duration used for thermal loading are not established here. The furnace experiments instead demonstrated that controlled heating could be incorporated into future hydrogen- and deuterium-loading procedures.

F.4 Deuteration

Once the hydrogen-loading procedure had been developed using ordinary hydrogen-containing chemicals, the process was modified for deuterium. The general electrolysis procedure remained similar, but materials containing protium were replaced wherever practical with their deuterated equivalents.

Throughout preparation and loading, the target was kept in contact with deuterated materials as much as possible. Distilled water was replaced by deuterium oxide, D_2O, and sulfuric acid was replaced by deuterated sulfuric acid, D_2SO_4. During the electrolysis loading process, LiOH was replaced by lithium deuteroxide, LiOD.

Several normal cleaning steps were also altered. IPA cleaning was removed from the deuteration procedure because ordinary isopropyl alcohol contains protium. Piranha cleaning was also avoided because the available hydrogen peroxide was not deuterated. These changes were made to reduce the amount of ordinary hydrogen brought into contact with the target during preparation.

The first deuteration tests used larger targets measuring approximately 1.5 in $\times$ 1.5 in rather than the approximately 1 in $\times$ 1 in targets used previously. The larger target reduced the fraction of

the surface covered by Kapton tape and helped maintain a good electrical connection between the target and its supporting frame.

The targets were first rinsed using distilled D_2O and then dried using a dry-gas blower rather than IPA. Both sides of the targets were then plasma cleaned for approximately 1.5 hours. Following plasma cleaning, the targets were placed into a D_2SO_4 bath for approximately 10 minutes and gently agitated approximately once per minute. They were then removed, rinsed with D_2O, dried, and plasma cleaned again for approximately 1.5 hours on each side.

After cleaning, the targets proceeded to the LiOD electrolysis loading step. A 0.2 M LiOD solution was selected. Because deuterated chemicals were comparatively limited, the smallest practical liquid volume was used. For the approximately 1.5 in $\times$ 1.5 in target geometry, a 50 mL beaker was selected.

The 0.2 M LiOD solution was prepared using approximately 0.2495 g of LiOD in 50 mL of D_2O.

During the first deuteration attempt, a relatively large current of approximately 1 A was used for an extended period. Over time, however, evaporation of the solution became substantial and made the loading process increasingly difficult to maintain. Although the target was exposed to deuterium throughout the process, the degree of deuteration produced by this first attempt could not be determined conclusively.

A second loading attempt was therefore performed with several modifications intended to reduce evaporation and provide a more stable electrolysis environment. A container with a lid was used, and Parafilm was placed near the top of the container to reduce vapor loss. The current was also reduced to approximately 0.2 A.

This second process was performed over two days. The target was deuterated for approximately four hours during the first day and approximately 24 hours during the second day. The measured target mass increased from approximately

$$0.140 \text{ g} \tag{F.8}$$

before loading to approximately

$$0.144 \text{ g} \tag{F.9}$$

after loading. However, the post-loading mass measurement was taken before the target had been fully rinsed. Therefore, the measured increase could not be attributed uniquely to deuterium incorporated into the titanium.

Overall, the deuteration procedure represented an improvement over the earlier hydrogen-loading experiments because the chemical environment was more carefully controlled and protium-containing materials were removed where practical. Nevertheless, verification of the exact amount of deuterium incorporated into the target remained difficult. For this reason, the loading procedure was evaluated using the additional verification techniques described in Appendix G.

Appendix G

Verifications of Gas Loading

In order to confirm that hydrogen or deuterium gas had been successfully integrated into the metal, several separate verification methods were considered and tested. Some of these methods were unsuccessful in determining whether hydrogen or deuterium had been loaded into the material, while others provided useful evidence. This appendix describes each method that was considered, the reasoning behind its use, and the results that were obtained.

The purpose of these tests was not necessarily to determine the exact concentration of hydrogen or deuterium within the metal. Instead, the primary goal was to verify that loading had occurred and that the gas was retained within the target after the loading procedure. Several methods were explored because no single inexpensive and readily available technique initially provided a clear answer.

Of the methods attempted, the most useful verification was obtained through out-gassing measurements using a helium leak tester. This method provided a direct way to detect hydrogen or deuterium leaving the material after loading and therefore gave the clearest experimental evidence that gas had been incorporated into the target.

G.1 Weight

The first and most direct method considered was measuring the change in mass of the target before and after hydrogen or deuterium loading. During electrolysis, hydrogen or deuterium is expected to enter the metal. If a sufficiently large quantity of gas is absorbed, this should produce a corresponding increase in the total mass of the target.

The targets were therefore weighed before and after the loading procedure in an attempt to identify this change. In principle, the measured mass difference could be compared with the expected mass of hydrogen or deuterium introduced into the material.

In practice, however, the expected mass change was very small compared with the resolution and repeatability of the available balance. Additional effects, including residual water, surface contamination, oxide formation, and small differences in target handling, could also produce changes comparable to or larger than the expected increase from hydrogen or deuterium loading.

As a result, the weight measurements were not sufficiently sensitive to provide a reliable confirmation of loading. Although this method was simple and nondestructive, it was ultimately not useful for determining whether a meaningful amount of hydrogen or deuterium had been incorporated into the targets.

G.2 Stress Test

A second verification method considered was a mechanical stress test. Hydrogen and deuterium can alter the mechanical properties of certain metals and may cause hydrogen embrittlement after sufficient loading. A successfully loaded target could therefore be expected to become more brittle or otherwise behave differently under mechanical stress.

Loaded targets were compared with unloaded or differently treated samples in order to determine whether a noticeable change in mechanical behavior had occurred. Increased brittleness or easier fracture could provide qualitative evidence that hydrogen or deuterium had entered the material.

However, this method contained several important ambiguities. The targets were exposed to other processes that could also alter their mechanical properties. In particular, heating in furnaces, extended electrolysis baths, chemical cleaning, and other preparation steps could introduce stress or otherwise weaken the material.

Consequently, even when a change in brittleness was observed, it could not be uniquely attributed to hydrogen or deuterium loading. The stress test therefore provided, at most, qualitative supporting evidence and was not considered a reliable independent verification method.

G.3 Spark Spectroscopy

Spark spectroscopy was also considered as a possible method for identifying hydrogen or deuterium within the loaded targets. In this method, a sufficiently energetic electrical discharge can remove and excite material from the surface. The emitted light can then be measured spectroscopically, allowing characteristic emission features to be identified.

In principle, hydrogen or deuterium present in the target could produce identifiable spectral features. This would provide a more direct method of confirming that the gas was present within or near the target surface.

Preliminary attempts indicated that this method could potentially be useful. However, the available experimental equipment was not sufficient to perform the measurement with the required accuracy. In particular, the available voltage was not high enough to consistently produce the desired discharge conditions, and the spectrometer did not provide sufficient resolution to confidently distinguish the relevant spectral features.

For these reasons, spark spectroscopy was not developed further in this work. With a higher-voltage discharge system and a higher-resolution spectrometer, this method may provide a useful independent verification technique in future experiments.

G.4 XDF

An additional method considered was XDF. This technique could potentially provide detailed information regarding changes in the material caused by hydrogen or deuterium loading and could therefore serve as a strong verification method.

The primary limitation of this approach was accessibility. The necessary equipment was not readily available for the experiments described in this work and would have significantly increased the cost and complexity of the loading verification procedure. Maintaining relatively inexpensive and accessible verification techniques was an important consideration during the development of the experiment.

As a result, XDF was not used as a primary verification method in this work. It remains a possible future technique for independently characterizing loaded targets if appropriate equipment becomes available.

G.5 Out-gassing

The most successful verification method used in this work was measurement of gas released from the target through out-gassing. If hydrogen or deuterium has been incorporated into a metal during the loading procedure, some fraction of that gas should subsequently leave the material when it is placed under vacuum. Detecting this released gas therefore provides direct evidence that hydrogen or deuterium had previously been present within the target.

A helium leak tester was used for these measurements. Although the instrument is normally used to locate helium leaks in vacuum systems, its mass-spectrometer-based detection system can also be used to monitor gases at selected mass-to-charge ratios. This made it possible to configure the instrument to search for deuterium-related signals rather than helium.

A probe arrangement was constructed so that gas released from the loaded target could be directed toward the leak tester. The target was mounted across two electrical leads via screw terminals to ensure a robust connection. These feedthrough leads were welded through a vacuum flange, which sealed a vacuum tube equipped with an auxiliary port for the leak detector. After assembly, the system was pumped down, a current on the order of 10 amps was run through the target, and the detector response was recorded as a function of time.

Loaded and unloaded targets were compared using the same general procedure. The loaded samples produced a larger deuterium signal during pump-down than the corresponding background or unloaded samples. The measured signal then decreased with time as gas was removed from the system and released from the target at a progressively lower rate.

After the loaded target was introduced to the measurement system, the leak-tester response increased to a peak and then gradually decreased with time. This behavior was consistent with hydrogen or deuterium being released from the target as the system was pumped down. As the available gas within and near the surface of the metal was depleted, the measured out-gassing rate correspondingly decreased.

The decay portion of the measured signal was fit as a function of time. By integrating the fitted out-gassing curve, an estimate could be made of the total amount of hydrogen or deuterium released from the target during the measurement. Although the calculated quantity of retained gas was not especially large, it was sufficient to demonstrate that measurable hydrogen or deuterium had been incorporated into the metal during the loading procedure. The estimated quantity also indicated

that a meaningful population of deuterium nuclei could be present within the target, providing a sufficient number of potential interactions for the fusion measurements performed in this work.

Because the original numerical data from this verification measurement were not retained, the exact fitted parameters and integrated gas quantity are not reproduced here. Nevertheless, the observed rise, peak, and subsequent decay in the mass-specific leak-tester signal provided the clearest experimental verification that the loading procedure successfully introduced hydrogen or deuterium into the target.

Unlike the weight and mechanical tests, the out-gassing measurement directly detected the gas species released from the target and provided a method for estimating the total amount of retained gas. It also used equipment that was already available as part of the vacuum system. For these reasons, the helium leak tester became the preferred method for verifying that hydrogen or deuterium had been successfully loaded into the targets.

G.6 Summary of Verification Methods

Several methods were investigated in an attempt to verify hydrogen and deuterium loading. Weight measurements were simple to perform but were not sufficiently sensitive to detect the expected mass change. Mechanical stress testing produced some potentially useful qualitative observations, but changes in material strength could not be uniquely attributed to hydrogen or deuterium loading. Spark spectroscopy showed potential but was limited by the available high-voltage equipment and spectrometer resolution. XDF could provide more detailed material characterization, but the required equipment was not readily available within the scope of this work.

Out-gassing measurements using the helium leak tester ultimately provided the clearest verification. The detection of hydrogen- or deuterium-related gas released from previously loaded targets provided direct evidence that gas had been incorporated into the material. Although this method did

not independently determine the exact hydrogen or deuterium concentration within the target, it provided a practical and repeatable verification that the loading procedure successfully introduced gas into the material.

Appendix H

Vacuum System Cleaning Methods

In order to ensure that the vacuum system remained as free of contamination as reasonably possible, several different cleaning and handling procedures were used throughout its construction and operation. These procedures were intended both to remove contamination already present on the vacuum components and to prevent additional contamination from being introduced during assembly.

The cleanliness of the system was particularly important because oils, machining residue, water, dust, and other contaminants can significantly increase outgassing and make it difficult to reach the desired vacuum pressure. The procedures described in this appendix were developed throughout the construction of the experimental system and ultimately allowed the chamber to routinely reach pressures below the microtorr range.

The cleaning process can generally be divided into three stages. First, individual vacuum components were cleaned before assembly. Second, the vacuum system was assembled while minimizing exposure to contamination. Finally, after assembly, extended pumping, gauge degassing, argon flushing, and plasma cleaning were used to remove remaining gases and surface contamination from the interior of the chamber.

H.1 Cleaning Parts

Before assembly of the vacuum system began, each component was thoroughly cleaned. Many of the components had either been newly machined or stored for extended periods before use. As a result, machining oils, grease, dust, fingerprints, and other surface residues could potentially remain on the parts.

Removing these contaminants before assembly was significantly easier than attempting to remove them after the pieces had been installed into the vacuum system. Each component therefore underwent an initial chemical and solvent cleaning procedure before being placed onto the chamber.

H.1.1 Initial Cleaning

The first stage of cleaning used a low-concentration sulfuric-acid solution. Each component was sprayed with the solution and allowed to sit for several minutes. A soft-bristled cleaning tool was then used to gently scrub accessible surfaces and remove oils or other residues that remained attached to the material.

Following the sulfuric-acid treatment, the component was thoroughly rinsed with distilled water to remove remaining acid and loosened contamination. Immediately following the distilled-water rinse, the component was rinsed again using isopropyl alcohol (IPA). The IPA assisted in removing remaining organic material and displaced much of the water remaining on the surface.

The component was then dried using a dry compressed-air hose. Once completely dry, the cleaned part was placed inside a clean plastic bag and stored until assembly. In general, cleaned components were installed within approximately one or two days in order to reduce the amount of time available for additional contamination to accumulate.

This cleaning procedure was not intended to create an atomically clean surface. Instead, it removed the bulk oils, machining residue, dust, and other contaminants most likely to contribute significantly to outgassing once the system was placed under vacuum.

H.1.2 Final IPA Cleaning

Immediately before assembly, each component was removed from its storage bag and given a final IPA rinse. This additional cleaning step was used in case the storage bag, surrounding air, or handling during the previous cleaning procedure had introduced any new residue onto the component.

If visible contamination remained after the IPA rinse, a Kimwipe was used carefully to remove the material. Wiping of vacuum surfaces was minimized whenever possible in order to avoid leaving fibers or introducing additional particles. After any visible material was removed, the component was rinsed again with IPA and dried using the dry-air hose.

Once this final cleaning step was complete, the component was kept covered until immediately before installation into the vacuum system.

H.2 Vacuum Assembly

The vacuum system was assembled with particular care given to maintaining the cleanliness established during the component-cleaning procedure. Open vacuum ports were kept covered whenever they were not actively being worked on. Depending on the size and location of the opening, ports were temporarily protected using clean covers, aluminum foil, or other barriers until the corresponding component was ready to be installed.

Nearly all of the ConFlat connections in the system used copper gaskets. New copper gaskets were preferred whenever a connection was opened. Before installation, the gaskets were handled

carefully and were kept free from fingerprints, oils, and other contamination. When additional cleaning was necessary, they underwent a procedure similar to that used for the other vacuum components.

Gloves were worn during assembly, and direct contact with internal vacuum surfaces was avoided whenever possible. Once all components required for a particular connection were prepared, the port was opened and assembled quickly. This reduced the amount of time that cleaned internal surfaces were directly exposed to the surrounding laboratory environment.

The same approach was followed as the chamber was progressively assembled. Components were cleaned and prepared before opening the system, and exposed ports were closed as soon as practical. Although complete elimination of prolonged atmospheric exposure was not possible, minimizing both the duration and frequency of exposure substantially reduced the amount of contamination introduced during construction.

H.3 Outgassing the Chamber

Even after careful cleaning and assembly, gases remain adsorbed on the internal surfaces of a vacuum system. Water vapor is particularly persistent after atmospheric exposure, although other residual gases and contaminants can also remain on chamber walls, gauges, seals, and other internal components.

After assembly, several procedures were therefore used to reduce this remaining gas load. These included extended pump-down periods, degassing of the hot-cathode ionization gauge, and repeated flushing of the chamber with argon.

H.3.1 Pump-Down Time

One of the simplest but most important methods was allowing the system to remain continuously under vacuum for extended periods. The chamber was pumped using the turbomolecular pump for periods extending beyond one week when possible.

During the initial pump-down, pressure decreased rapidly as the majority of the free gas within the chamber was removed. At lower pressures, the pump-down became increasingly dominated by gases slowly desorbing from the internal surfaces of the system. Maintaining the chamber under vacuum for extended periods allowed these molecules to gradually leave the surfaces and be removed by the pumping system.

Following sufficiently long pump-down periods, the chamber could reach the microtorr region and continue decreasing as the internal surfaces became progressively cleaner.

H.3.2 Ion-Gauge Degassing

Once the chamber reached the appropriate low-pressure range, the degassing function of the hot-cathode ionization gauge was periodically used.

The ionization gauge itself can become an important source of outgassing because gases can adsorb onto its filament, grid, and surrounding components. The degassing procedure heats portions of the gauge and increases the rate at which these adsorbed species are released from its surfaces. Once released, the gases can be removed by the vacuum pumps.

During degassing, the chamber pressure could temporarily increase as gases were driven from the gauge components. Continued pumping following the degassing cycle then removed these molecules from the chamber. Repeated use of this procedure helped reduce the gauge's contribution to the residual gas load and allowed the system to reach progressively lower pressures.

H.3.3 Argon Flushes

After the completed system had been closed, exposure to ordinary atmosphere was avoided whenever possible. Argon was regularly used to flush the chamber before and after experimental tests and was also used during plasma-cleaning procedures.

For a typical argon flush, the chamber was first isolated from the turbomolecular pump. Argon was then introduced until the chamber pressure reached approximately three-quarters of atmospheric pressure. This allowed the residual low-pressure gas within the chamber to be substantially replaced by argon.

The chamber was then slowly pumped down using the secondary pumping path. Once the pressure had decreased sufficiently, the turbomolecular pump could again be connected to the main chamber and used to continue the pump-down into the high-vacuum region.

Repeated argon flushing reduced the fraction of atmospheric gases remaining in the chamber and provided a controlled gas environment before subsequent plasma cleaning or experimental operation. This procedure was particularly useful following periods when the system had accumulated residual water or other gases.

H.4 Plasma Cleaning

One of the most useful in-situ cleaning procedures after vacuum assembly was plasma cleaning. Argon was introduced into the chamber and a plasma was generated using conditions similar to those discussed for the separate plasma-cleaning system in Appendix E.

The purpose of this procedure was to expose the interior chamber surfaces to the plasma and assist in removing adsorbed contamination. Energetic plasma species interact with material present on the chamber walls, allowing some surface contaminants to be released into the gas phase where they can subsequently be removed during pump-down.

A visible change in the plasma was commonly observed during extended cleaning periods. An example is shown in Figure H.1. The three images were taken at different times during approximately one hour of plasma operation.

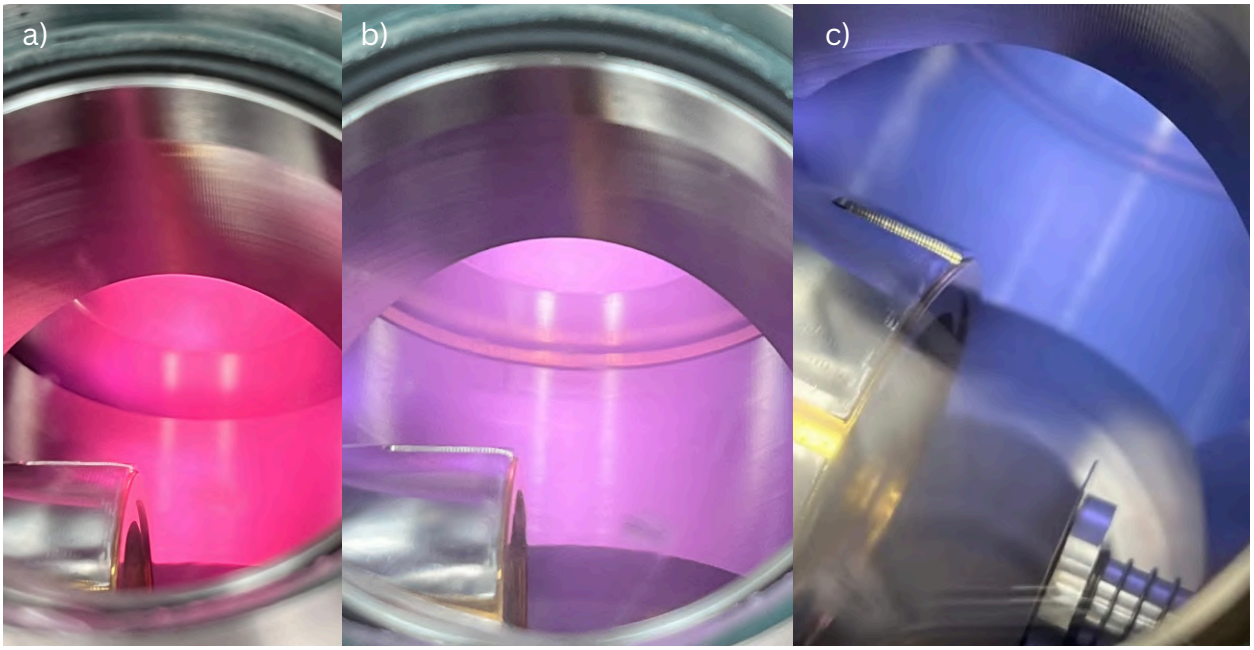

Figure H.1 Photographs of the chamber plasma during an approximately one-hour plasma-cleaning cycle. Image (a) was taken shortly after the plasma was ignited, image (b) after approximately 20 minutes, and image (c) after approximately one hour. The plasma changed from a lighter pink color shortly after ignition to a darker purple color after extended operation. This change was qualitatively consistent with a reduction in residual atmospheric and water-related species and an increasing dominance of the introduced argon within the plasma.

At the beginning of the cleaning process, the plasma frequently appeared substantially lighter and more pink than the characteristic darker purple observed when operating primarily with argon. As plasma cleaning continued, the color gradually shifted toward the darker purple appearance.

The plasma color by itself does not provide a quantitative measurement of the gas composition within the chamber. However, the repeatable change observed during cleaning provided a useful

qualitative indication that the gas environment was changing as residual contamination was removed and the introduced argon became increasingly dominant.

Following plasma cleaning, the chamber was again pumped down for an extended period. Combining plasma exposure with continued pumping provided a practical method for cleaning surfaces that could no longer be physically accessed after assembly.

H.5 Results

The combined cleaning procedures produced a substantial improvement in the performance of the vacuum system. Earlier chamber configurations were significantly more difficult to pump down and could remain in substantially higher pressure ranges. They also exhibited faster pressure increases after isolation, indicating greater leakage, outgassing, or both.

After implementing the component cleaning, careful assembly, extended pump-down, gauge degassing, argon flushing, and plasma-cleaning procedures described above, the completed chamber routinely reached the microtorr region much more rapidly. Under favorable conditions, the system could pump from its initial operating state to below approximately 0.5 μTorr within approximately ten minutes once the high-vacuum pumping system was operating.

The pressure-rise behavior of the closed chamber also improved substantially. When isolated from the pumping system, the measured pressure increase was approximately 1 $\mu\mathrm{Torr}/\mathrm{s}$ or less under the tested conditions. This slower pressure increase allowed the chamber to remain within the required vacuum range during experimental procedures in which portions of the pumping system needed to be isolated.

Overall, these cleaning and handling methods transformed the vacuum system from one that was difficult to maintain at low pressure into a system capable of repeatedly reaching the sub-microtorr range. This improvement was particularly important for experiments using hydrogen and

deuterium, where limiting atmospheric contamination reduced unwanted gas species and allowed the comparatively expensive experimental gases to be introduced into a better-controlled environment.

Appendix I

Previous Experimental Design

I.1 Creating the Environment for Fusion

This section describes the vacuum chamber and its supporting components used to create the environment for fusion. Specific details regarding the experimental procedure to allow fusion is discussed in the following section (Section 4.3). While a detailed discussion of each component in the vacuum chamber follows, a CAD rendering of the theorized system is shown in Fig. I.1. The actual system created is shown in figure I.2. Slight modifications from the CAD design were required though all basic components remain the same.

The central chamber is built around a standard ConFlat four-way cross that serves as the primary vacuum junction. An additional smaller ConFlat flange is welded to the center of one face of the cross to provide a dedicated mounting port for components positioned along the experimental axis. On the opposite side, four VCR fittings are welded in a square pattern around the center of the flange to provide ports for gas handling, pressure diagnostics, and other auxiliary connections. During construction, a second four-way ConFlat cross was attached to one of the primary flanges in order to increase the number of available ports for pumping, plasma generation, and diagnostics.

This configuration provides a compact experimental platform in which the key functions required for the experiment–vacuum pumping, deuterium gas handling, plasma generation, cavity formation, and diagnostic access–can be integrated within a single vacuum system. The modular flange geometry also allows individual components to be installed or modified without requiring major changes to the overall chamber design.

Throughout this section, the chamber is described with the primary four-way cross oriented vertically, with the turbomolecular pump attached to the bottom flange.

I.1.1 Pumping

With the 4-cross chamber positioned upright, a Varian turbomolecular pump (turbopump) is attached to the bottom ConFlat flange. This pump serves two primary purposes: assisting the roughing pump during initial pump-down and enabling continuous circulation of deuterium gas through the system during fusion operation. To accomplish this, the turbopump exhaust is routed into a T-junction controlled by two bellows valves. One valve connects the system to the roughing pump during pump-down, while the other redirects the exhaust back into the chamber through a flange attached to the secondary four-way cross. This allows the system to have controlled gas circulation while fusion is underway.

I.1.2 Plunger for Cavity Creation

The secondary four-way cross, in addition to providing additional ports for the system, houses the plunger assembly used to form the cavity. The internal end of the plunger holds the target prepared as described in Section 4.1. By translating the plunger along the chamber axis, the target surface can be brought toward the detector positioned on the opposite side of the chamber, allowing the separation between the two surfaces to be precisely controlled.

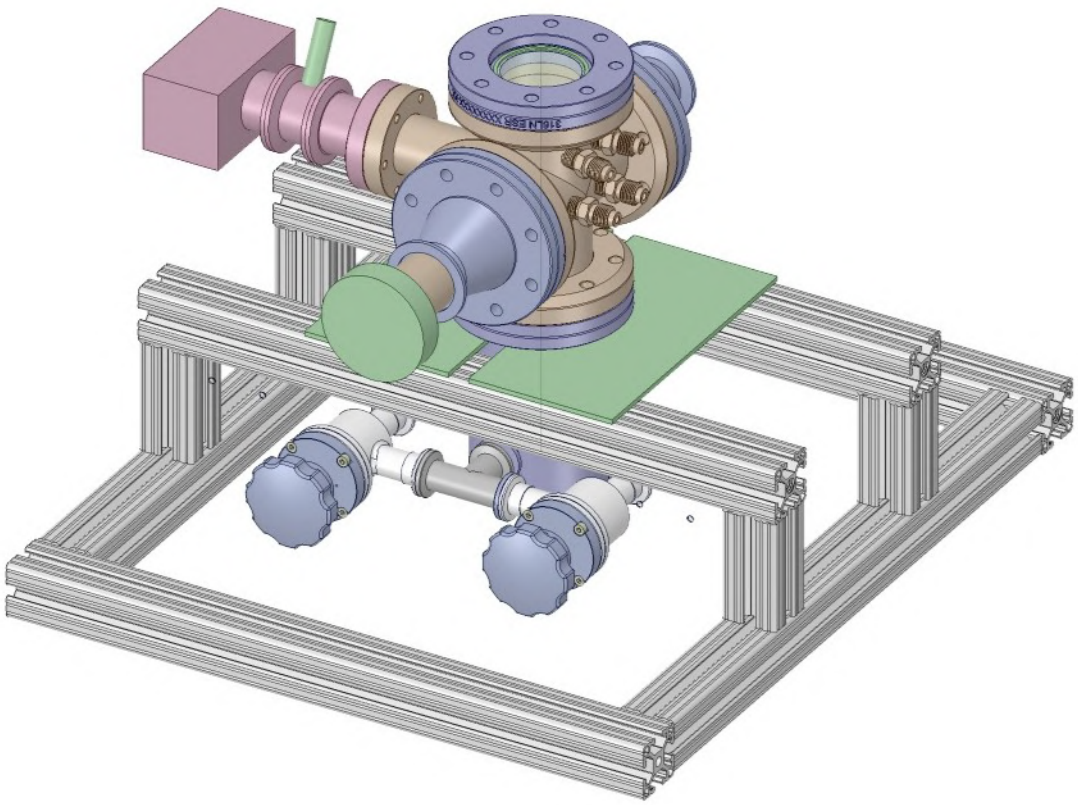

Figure I.1 CAD model of the vacuum chamber used in the experimental system. The design illustrates the chamber geometry and component layout used to support cavity formation, plasma generation, and fusion product detection.

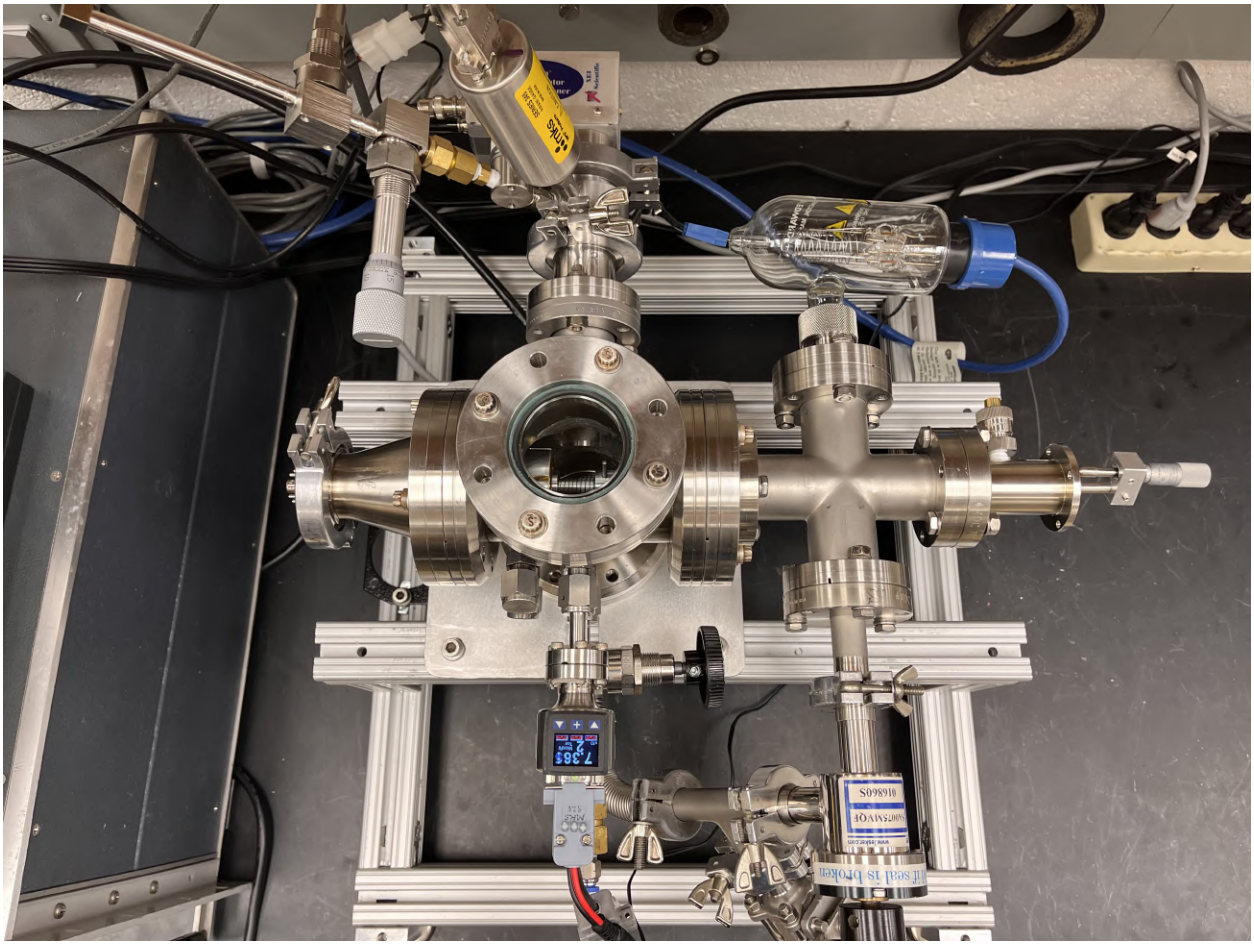

Figure I.2 Photograph of the assembled vacuum chamber used in the experiment, showing the physical implementation of the system based on the CAD design in Fig. I.1.

The mechanical design of the plunger is critical, as the detector must not be damaged during cavity formation. The assembly therefore allows smooth, controlled motion of the target while maintaining alignment between the target and detector surfaces. This configuration enables the cavity spacing to be reduced to nanometer-scale distances while minimizing the risk of mechanical contact or misalignment that could damage the detector.

In addition to enabling controlled positioning of the target, the plunger provides mechanical stability during operation so that the cavity spacing remains fixed while the plasma is sustained within the chamber. Together with the protective features described in Section 4.1, this design allows reliable formation of a confined cavity environment suitable for the fusion experiments described in this work.

I.1.3 Window

A circular window is installed on the top ConFlat port of the chamber. This window enables visual alignment of the plunger during cavity formation, allowing for more precise positioning. In addition, the window facilitates placement and alignment of external diagnostic equipment, including the neutron detector.

Furthermore, the glass window is not rigidly attached to the opening of the system; instead, it rests on an O-ring. It seals due to its weight and the pressure difference between the vacuum chamber and atmosphere. This allows the glass window to act as an overpressure release. Throughout the experimental process, the chamber is filled with either argon or a hydrogen-isotope gas. Although care is taken to ensure that the chamber is not overpressurized, the window provides an additional safety measure in the event of overpressurization.

While not strictly necessary for system operation, the window provides a practical improvement for both cavity control, detector positioning, and safe experimental practices.

I.1.4 Plasma Initiator

A plasma initiation assembly is attached to the smaller ConFlat flange welded onto the chamber. This assembly consists of a high-voltage power supply used to ignite and sustain a plasma within the chamber. An additional attachment modifies the incoming gas flow to ensure a steady and controlled inflow of deuterium. The plasma is initiated within a specific pressure range and can be regulated by adjusting the voltage output of the plasma source.

Once ignited, the plasma is expected to extend through the cavity region, enabling monatomic deuterium to be distributed throughout the system and onto the cavity walls. Throughout operation, system parameters are continuously monitored to ensure safe and stable operation. In the event of abnormal behavior, the plasma system is designed to shut off automatically.

After the plasma has been sustained for a sufficient period to allow deuterium loading of the cavity surfaces, the deuterium inlet is closed and the plasma is extinguished. At this stage the deuterium gas remains confined within the chamber while the plasma source is turned off, allowing the system to transition into the measurement phase. The cavity can then be formed using the plunger assembly while the detectors begin recording data under stable gas conditions.

I.1.5 Detector Positioning

Opposite the plunger, a charged-particle detector is mounted to the chamber using a custom-designed ConFlat attachment. This detector is positioned face-to-face with the target to maximize sensitivity to charged fusion products emitted from the cavity region.

A neutron detector is also employed in this system. Due to the high penetrating power of neutrons, precise placement of the neutron detector is less critical than for charged-particle detection. For simplicity and accessibility, the neutron detector is positioned above the glass window on the top ConFlat port.

I.1.6 Other Instrumentation

Several additional diagnostic instruments are incorporated into the vacuum system to monitor operating conditions during the experiment. A Pirani vacuum gauge is mounted on one of the VCR ports to provide pressure measurements during the initial pump-down and during operation in the higher pressure regime used for plasma ignition. This gauge provides real-time feedback for adjusting the gas flow and ensuring that the chamber pressure remains within the range required for stable plasma formation.

For measurements at lower pressures, an ionization gauge tube is mounted on one of the four-way crosses that houses the plunger assembly. The ion gauge provides accurate pressure measurements in the high-vacuum regime and is used to verify the base pressure of the system prior to introducing deuterium gas. Together, these gauges allow monitoring of the chamber pressure across the full operating range required for pump-down, plasma generation, and fusion measurements.

Appendix J

Previous Inertial Electrostatic Confinement Fusion Work

Before beginning the cavity-based fusion experiment discussed in this thesis, I participated in the construction and operation of a Farnsworth fusor. This work began as a collaborative between the Laboratory of Nuclear Astrophysics Research and a final project with Viktor Horvath and Joshua Brady in PHSCS 245 at Brigham Young University [39].

The fusor project was separate from the primary experiment presented in this thesis, and its measurements are not included in the results of the present work. It is included in this appendix because it provided previous experience with deuterium-deuterium fusion, vacuum systems, plasma generation, high-voltage operation, deuterium gas handling, leak checking, and nuclear radiation detection. Many of these skills were later applied during the design and construction of the cavity-based fusion system.

J.1 Project Motivation

The original objective of the project was to assemble and characterize an inertial electrostatic confinement device that could eventually be used to study proton-boron fusion. Proton-boron fusion was of interest because its primary products are energetic alpha particles rather than the high-energy neutrons associated with many other fusion reactions.

Because the proton-boron reaction requires more demanding experimental conditions and specialized materials, deuterium was used during the initial operation and testing of the system. Deuterium-deuterium fusion provided a more accessible reaction for evaluating the vacuum chamber, cathode grid, high-voltage system, gas handling, and radiation detectors.

This project differed from the primary work in this thesis in the mechanism used to increase the fusion probability. The fusor accelerated deuterium ions electrostatically to increase their kinetic energy. The present thesis instead investigates whether electron screening in confined environments can reduce the effective Coulomb barrier at much lower incident energies. Despite this difference, both experiments required the creation and detection of deuterium-deuterium fusion under controlled laboratory conditions.

J.2 Farnsworth Fusor

A Farnsworth fusor is an inertial electrostatic confinement device. It consists of a vacuum chamber containing a negatively biased cathode grid. Deuterium gas is introduced into the chamber at low pressure and ionized to form a plasma. The negatively charged grid accelerates the positive deuterium ions toward the center of the chamber.

Many of the accelerated ions pass through the openings in the grid and continue across the central region. The electric field then decelerates and redirects the ions, causing them to pass

repeatedly through the center of the device. Collisions between sufficiently energetic deuterium ions can result in deuterium-deuterium fusion.

The principal components of the system were

- a stainless-steel vacuum chamber;
- a roughing pump and turbomolecular pump;
- thermocouple, capacitance-manometer, and ionization vacuum gauges;
- a high-voltage feedthrough and central cathode grid;
- a regulated deuterium gas supply;
- a neutron detector;
- an EG&G ORTEC DIAD charged-particle detector; and
- a Geiger-Müller detector.

J.3 Experimental Procedure

The chamber was first assembled and checked for leaks. Leak checking was important because stable operation required a continuous flow of deuterium while maintaining the chamber within the millitorr pressure range. The system was pumped to a base pressure of approximately

$$2.8 \times 10^{-5} \text{ Torr}$$

before deuterium was introduced.

Deuterium gas was then leaked into the chamber. The turbomolecular-pump flow was throttled so that a nearly constant chamber pressure could be maintained while deuterium continued to enter the system. Test pressures ranged from approximately 1 to 50 mTorr.

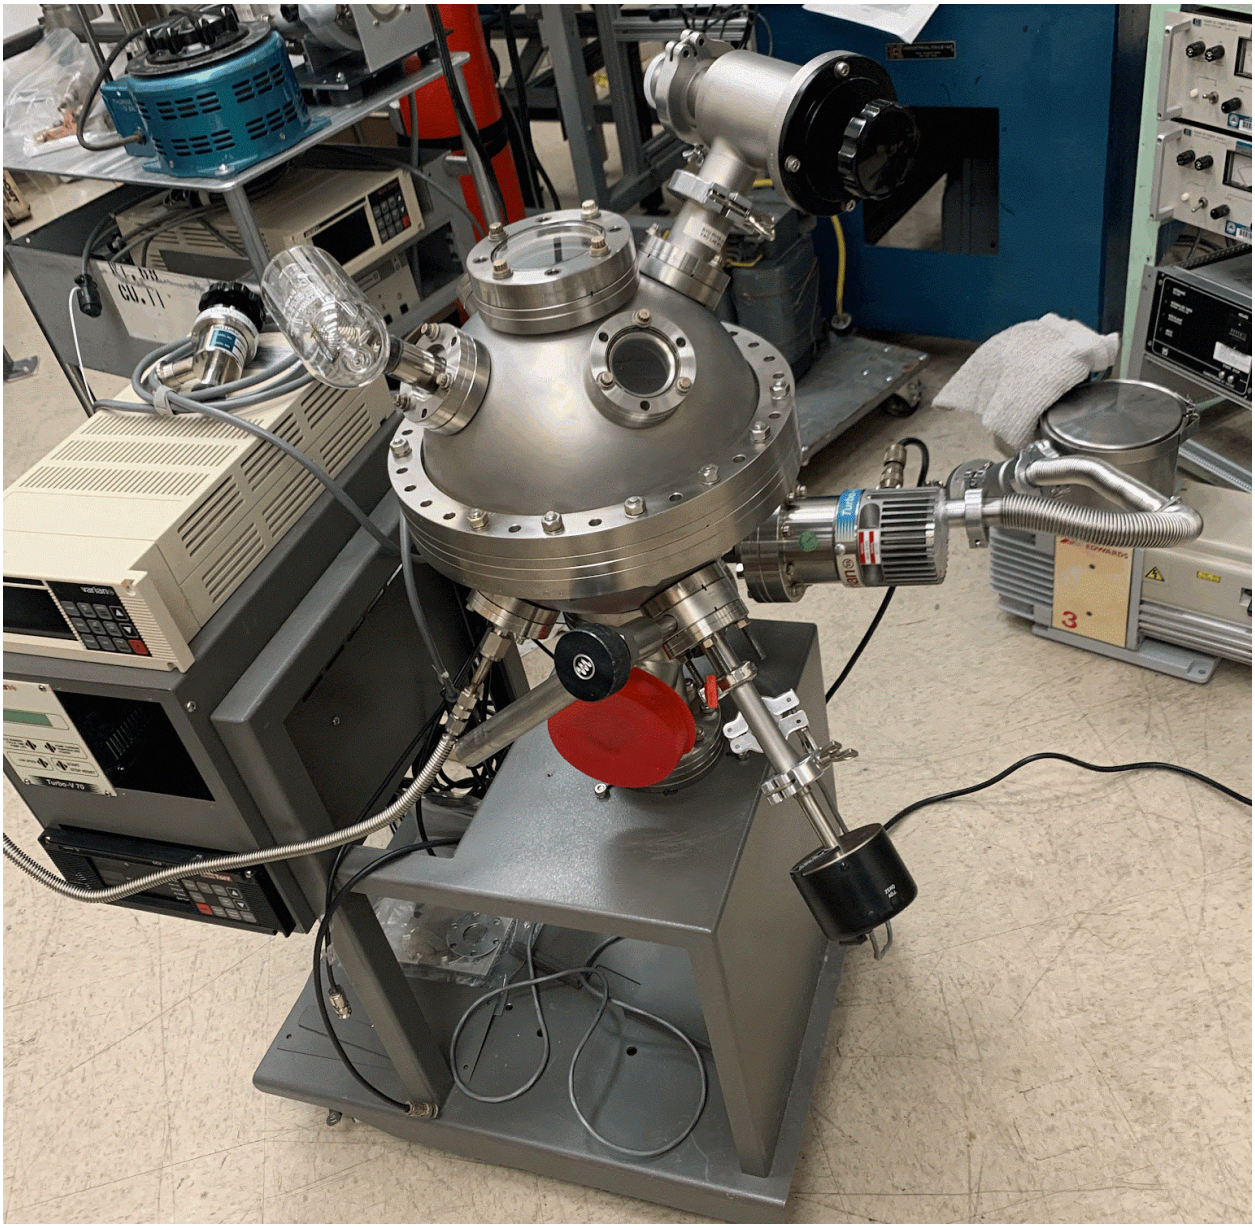

Figure J.1 Farnsworth fusor used during the previous deuterium-deuterium fusion project. The system included a stainless-steel vacuum chamber, central cathode grid, high-voltage feedthrough, vacuum pumps, pressure gauges, and radiation detectors.

After a stable pressure was reached, a negative voltage was applied to the central cathode grid. The increasing voltage ionized the gas and accelerated the deuterium ions toward the central region. The plasma was observed visually while the neutron, charged-particle, and Geiger-Müller detectors were monitored.

The general operating sequence was therefore

1. pump and outgas the vacuum chamber;
2. introduce deuterium to the desired operating pressure;
3. throttle the pumping system to maintain a continuous deuterium flow;
4. apply negative high voltage to the central grid;
5. observe the plasma and monitor the radiation detectors; and
6. compare the measured detector counts with separately measured background values.

J.4 Initial Fusor Configuration

The first configuration used a tungsten-wire cathode grid with a diameter greater than approximately 3 inches. The system was operated over a pressure range of approximately 5-30 mTorr for a measurement period of approximately 1000 s. See Figure J.2 for an image of the attempted fusion with the large cathode grid.

This configuration was unable to reach the desired high voltage and operated at only approximately 3 kV. The chamber pressure also fluctuated as the applied voltage increased. The large tungsten grid was identified as a likely contributor to both problems. Its size and geometry increased the electrical and plasma-loading demands on the high-voltage system.

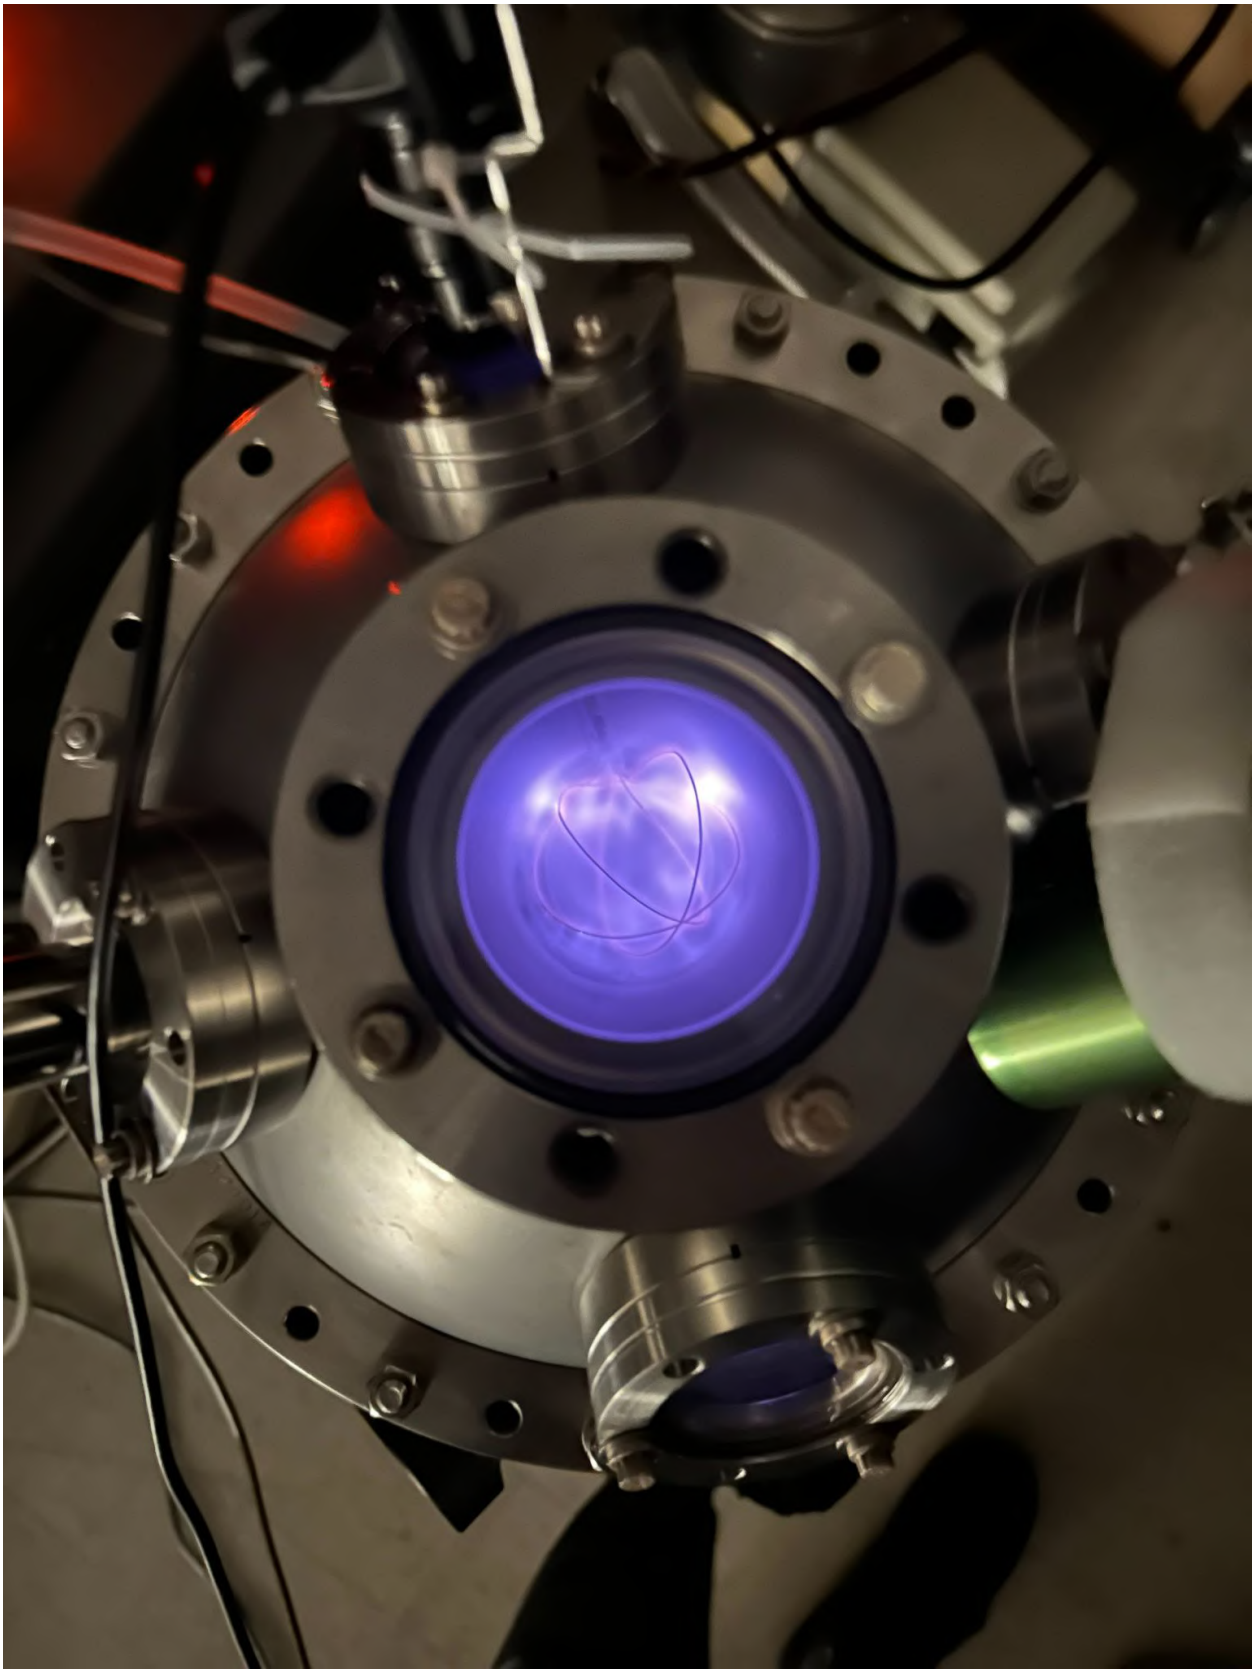

Figure J.2 Photograph of plasma produced using a large grid in the Farnsworth fusor. This light around the cage indicated a high density plasma but no clear evidence of fusion.

Although neutron and charged-particle detector counts were recorded during this attempt, the unstable pressure and limited applied voltage prevented a strong conclusion from being drawn from the measurements. The first attempt was therefore primarily useful for identifying changes needed in the grid geometry and operating procedure.

J.5 Revised Fusor Configuration

For the second configuration, the tungsten grid was replaced by a smaller stainless-steel grid with a diameter below approximately 2 inches (see Figure J.3. The smaller grid allowed the system to reach a substantially larger applied voltage while maintaining a more stable plasma.

The revised system was operated at a pressure of approximately

$$6.2 \text{ mTorr}$$

and an applied voltage of approximately

$$34 \text{ kV}.$$

The measurement continued for approximately 500 s.

A bright, localized plasma structure was visible near the center of the cathode grid, as shown in Figure J.4. The concentration of the plasma near the grid center was consistent with the expected behavior of an inertial electrostatic confinement device. In addition, the star pattern seen through the openings in the grid also give evidence for deuterium fusion.

J.6 Detector Measurements

The higher-voltage run produced detector counts above the separately measured background values. The recorded values are summarized in Table J.1.

Figure J.3 Photograph of the smaller stainless-steel cathode grid.

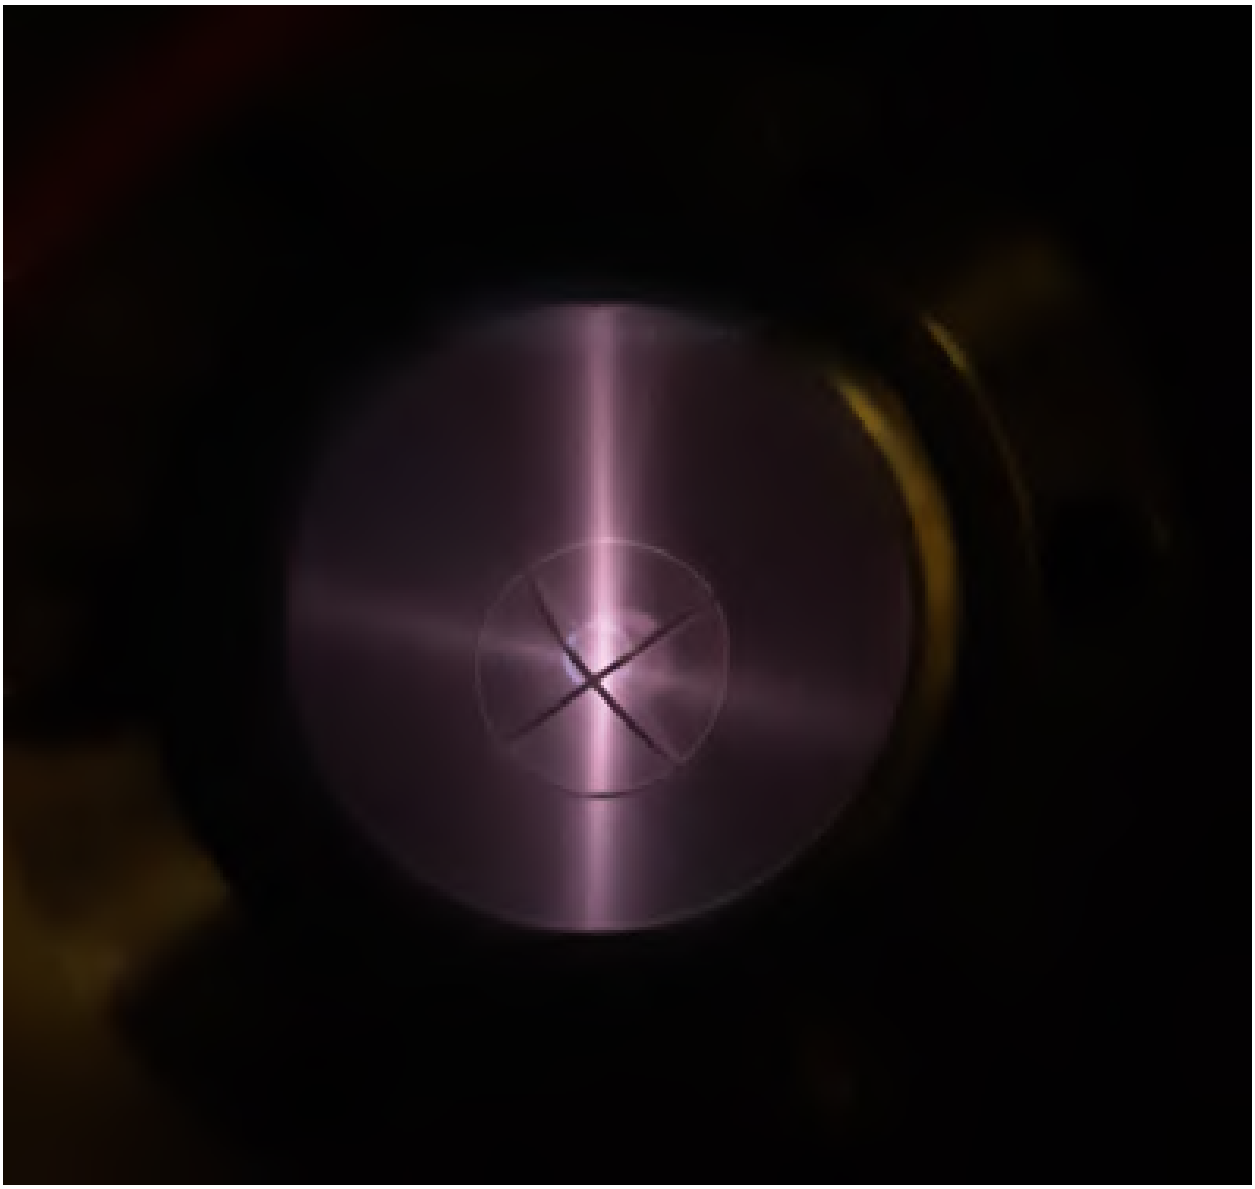

Figure J.4 Photograph of the plasma produced using the smaller stainless-steel cathode grid in the Farnsworth fusor. The dense central core and the star pattern visible through the grid openings provide evidence of fusion.

Table J.1 Detector measurements from the revised fusor configuration. The neutron column gives the total counts recorded during the approximately 500 s measurement period.

Measurement	Neutron-detector counts	Geiger-Müller rate (counts/s)
Fusor operating	101	40.0
Background	18	0.3
Difference above background	83	39.7

The neutron detector recorded 101 counts during operation, compared with a background of approximately 18 counts over the corresponding measurement period. The Geiger-Müller detector recorded approximately 40 counts/s, compared with a background rate of approximately 0.3 counts/s.

To estimate the total neutron-production rate, the measured detector counts were corrected for both the detector efficiency and the limited geometric coverage of the detector. The detector's fractional solid-angle coverage was approximated by comparing its active area, A_{det}, with the surface area of a sphere centered on the plasma:

$$f_{\Omega} = \frac{A_{\mathrm{det}}}{4\pi d^2}, \tag{J.1}$$

where d is the distance from the approximate fusion region to the detector. Assuming that the neutrons were emitted isotropically, the total neutron-production rate was then estimated using

$$\dot{N}_n = \frac{N_{\mathrm{det}}}{\eta f_{\Omega}}, \tag{J.2}$$

where N_{det} is the background-subtracted detector count rate and η is the neutron-detector efficiency. Using an estimated detector efficiency of approximately 5% and the calculated geometric coverage produced an estimated neutron rate of approximately

$$\dot{N}_n \approx 995.976 \text{ neutrons/s}.$$

This calculation provided an approximate correction for the large number of neutrons that traveled in directions outside the detector's field of view. However, the result depended strongly on the assumed location and size of the fusion region, isotropic neutron emission, the detector efficiency, and the accuracy of the background subtraction. It should therefore be interpreted as an order-of-magnitude estimate rather than a precisely measured absolute fusion rate.

These measurements showed a clear increase in detector response during the higher-voltage plasma operation. However, the measurements were preliminary. The detector efficiencies, solid-angle coverage, energy response, and possible electromagnetic interference were not characterized with sufficient precision to determine an accurate absolute fusion rate. For this reason, the results are presented as evidence of detector response above background rather than as a precise measurement of the deuterium-deuterium fusion yield.

J.7 Experimental Limitations

Several limitations prevented the project from producing a complete quantitative characterization of the fusor.

First, only a limited number of background and operating measurements were performed. Additional repeated runs would have been required to determine the statistical significance and reproducibility of the measured count excess.

Second, the neutron-detector efficiency and its geometric acceptance were only approximately known. Any calculated neutron-production rate would therefore contain substantial uncertainty.

Third, the Geiger-Müller detector was sensitive to ionizing radiation but did not independently identify the type or energy of the detected radiation. Its increased count rate could not by itself distinguish neutrons, charged-particle products, x rays, or electrical interference.

Finally, the first cathode-grid configuration produced unstable pressure and was unable to reach the intended voltage. Although the second configuration substantially improved the system, the project duration did not allow a full study of the dependence on voltage, pressure, grid size, or deuterium flow rate.

The results should therefore be understood as an initial demonstration and system-development project rather than a precision measurement of the fusor reaction rate.

J.8 Relevance to the Present Work

The fusor project provided several forms of experience that directly contributed to the experiment described in this thesis. The construction and operation of the fusor required the assembly and leak checking of a vacuum system, operation of roughing and turbomolecular pumps, interpretation of multiple vacuum gauges, controlled introduction of deuterium gas, generation of a plasma, and safe operation of high-voltage equipment.

The project also introduced several of the difficulties associated with detecting low-rate nuclear reactions. These included measuring detector background, accounting for detector efficiency and solid-angle coverage, distinguishing real particle signals from electrical noise, and avoiding conclusions based only on a single detector.

The present cavity-based experiment uses a different physical method for increasing the fusion probability, but many of the experimental requirements are the same. Both systems require a clean vacuum environment, controlled deuterium introduction, reliable plasma operation, complementary radiation detectors, and careful comparison with background measurements.

The previous fusor work therefore served as an experimental foundation for the more specialized system developed in this thesis. It also demonstrated how changes to apparently simple components,

such as the cathode-grid size and material, can strongly affect the pressure stability, achievable voltage, plasma structure, and detector response of a fusion experiment.

Appendix K

Hydrogen and Deuterium Detection Paper

The following paper is included as part of this appendix because it was critical in ensuring safe gas handling throughout the experimental method and procedure. The work describes how commercially available carbon monoxide alarms can be used to detect hydrogen gas before and at its explosive limits. This allowed the laboratory to maintain a large supply of hydrogen-detection safety devices to monitor hydrogen gas tanks and guard against leaks and other hazards.

Performance evaluation of commercial carbon monoxide alarms for hydrogen and deuterium detection

Tyler J. Hamm[1, a)] and John E. Ellsworth[1, b)]

Department of Physics and Astronomy, Brigham Young University, Provo, UT 84604

(Dated: 2 July 2026)

Hydrogen detection is critical for safety in laboratory environments due to its flammability. We report that commercially available carbon monoxide (CO) detectors, specifically those incorporating Figaro sensor elements, can meet most regulatory and performance standards required for hydrogen and deuterium gas detection. These standards are defined by Underwriters Laboratories (UL), the National Fire Protection Association (NFPA), the Occupational Safety and Health Administration (OSHA), and the International Organization for Standardization (ISO). To evaluate detector performance, a sealed chamber was constructed to introduce known concentrations of hydrogen gas isotopes (protium and deuterium) to commercial CO detectors that have passed relevant safety compliance standards. A series of experiments were conducted to measure sensitivity thresholds and alarm response times. We find that CO detectors using Figaro TGS5042 sensors consistently produced an alarm at the hydrogen and deuterium concentrations stated in the standards. Their average response time, however, exceeded a 30-second standard for hydrogen detection. In conclusion, these results suggest that while Figaro-based CO detectors are not optimized for rapid hydrogen detection, they still serve as a cost-effective and largely compliant alternative for hydrogen (protium and deuterium isotopes) detection in controlled settings.

[a)]hammtyler66@gmail.com

[b)]jell@byu.edu

I. INTRODUCTION

Throughout the last century, hydrogen has become an increasingly integral part of scientific studies. Hydrogen isotopes are often referred to by their specific names: protium (H_2), deuterium (D_2), and tritium (T_2), though general references to "hydrogen" typically encompass all isotopic forms unless otherwise specified. These various hydrogen isotopes have many productive uses in chemistry, industry, and energy. Along with the productive uses of hydrogen are its many dangers. Its explosive limit ranges from 4% to 75% and only requires 0.02 millijoules of ignition energy[1]. For comparison, natural gas—commonly used for domestic cooking, heating, and transportation fuel—has a flammable range of 5% to 15% and an ignition energy of 0.29 millijoules. Thus, hydrogen has a flammable range that is over 7 times greater than natural gas while requiring an ignition energy only 7% that of natural gas. Furthermore, its lower explosive limit (LEL) of 4% is very small and could be reached within seconds to minutes if it leaked into a medium-sized room depending on the flow rate and ventilation. To minimize safety concerns, many organizations have set forth different standards, codes, and regulations of hydrogen storage, use, and monitoring. These regulations, along with the increased casual usage of hydrogen, including H_2 and D_2 isotopes, in homes, labs, and workplaces, have led to the need for cheap, plentiful hydrogen detection for homeowners, hobbyists, and other basic scientific needs.

A. Current standards

In the current literature, there are a variety of organizations and standards for chemical gas detectors and their sensitivity limits. In the United States, some of the most common are Underwriters Laboratories (UL), The Occupational Health and Safety Administration (OSHA), The National Institute for Occupational Safety and Health (NIOSH), and The National Fire Protection Association (NFPA)[2]. Internationally, UL and the International Organization for Standardization (ISO) are used frequently to benchmark gas detection. Currently, OSHA and NIOSH do not refer to any specific sensitivity level for hydrogen gas. Instead, they aim to control the volume of H_2 contained by a workplace. Furthermore, most organizations and literature turn to UL and the NFPA for specifics on the standards of hydrogen gas sensors[3]. Accordingly, this paper focuses on the standards established by the

NFPA, UL, and ISO.

NFPA regulations for the alarm threshold of gaseous hydrogen are put forth in NFPA 715 Section 5.8.5.3.3. They are stated to align with the standards set forth in UL 2075[4]. All other details of the sensor (location, alarm sound, etc.) are cited as similar to any other gaseous detection devices, such as carbon monoxide alarms. All other NFPA regulations (such as NFPA 52 about hydrogen gas environments) only mention the need to be in conformance with UL 2075. Thus, UL 2075 is the common standard for hydrogen detection[5–7]. Section 15 of this code outlines the sensitivity standards for sensors, detectors, and alarms. The specific limit of the sensing is put forth in section 15.1. The most critical criteria (unique to UL standards) for hydrogen detection is placed in part C of section 15.1 and states in part: "A detector must produce an alarm signal at or below 25% of the LEL. This performance requirement applies to residential and commercial and industrial applications for detectors and/or their control assemblies that are intended for use in ordinary (non-hazardous) locations"[7]. Since hydrogen has an LEL of 4%, a hydrogen detector must produce an alarm at 1% hydrogen concentration. ISO 26142:2010 is the only standard that specifies a detection-time requirement, stating that a hydrogen detector must produce an alarm within 30 seconds of initial exposure to hydrogen at 25% of the LEL[8]. Thus, current standards set forth by NFPA and UL require hydrogen detectors to alarm at or below 25% of the LEL, corresponding to 1% hydrogen, while ISO 26142:2010 additionally specifies that this alarm occur within 30 seconds of initial exposure.

B. Carbon monoxide alarms

Carbon monoxide (CO) alarms are a possible inexpensive substitution for commercially dedicated hydrogen detectors. In contrast to top-of-the-line hydrogen detectors, which cost anywhere from hundreds to thousands of dollars, most high-end commercially available carbon monoxide alarms cost less than $50. Furthermore, as a commercial gas alarm, it has already gone through a testing process to ensure it meets the basic standards for commercial alarms. Clearly, the alarm is in compliance with the bulk of NFPA standards since, as a carbon monoxide alarm, it necessarily successfully passed tests proving it has a sufficiently loud alarm, durable enclosure, and so forth. Therefore, the only standards a carbon monoxide alarm has not innately met for being a qualified hydrogen detection device is the detection

of hydrogen at 25% of the LEL and producing an alarm within 30 seconds. Notably, certain electrochemical carbon monoxide sensors are documented to exhibit cross-sensitivity to hydrogen[9].

Many electrochemical CO sensors operate amperometrically, where CO is oxidized at a catalytic working electrode, commonly based on noble-metal catalysts, and the resulting current is used to calculate the gas concentration. This description includes the type of sensing element used in the Figaro TGS5042 sensor, which is the primary sensor considered in this work. Because hydrogen can also be oxidized at similar catalytic electrodes, exposure to H_2 or D_2 can produce a current response that appears as a false CO signal. In most CO-sensing applications, this cross-sensitivity is treated as an error source and can be reduced using filters, compensation electrodes, or sensor designs that subtract the hydrogen contribution. For hydrogen-isotope detection, however, this same cross-sensitivity may be useful if the alarm response is sufficiently repeatable and sensitive. The novelty of this work is therefore not the existence of cross-sensitivity itself, but the quantitative evaluation of intact commercial CO alarms for H_2 and D_2 alarm thresholds, response times, and practical agreement with the relevant hydrogen detection standards.

Common electrochemical carbon monoxide alarms include models produced by First Alert, Kidde, X-Sense, and Fire-X. Among these manufacturers, First Alert and X-Sense both incorporate the Figaro TGS5042 sensor, whereas Kidde and Fire-X rely on proprietary sensing elements for which detailed specifications are not publicly available. A Kidde alarm using a proprietary sensing element was also tested in preliminary trials and showed no observable alarm response to hydrogen under the tested conditions. This preliminary result further motivated the focus on alarms containing sensors with known hydrogen cross-sensitivity. First Alert alarms were selected as the primary test platform because they are low-cost, widely available consumer alarms that use the Figaro TGS5042, a common electrochemical CO sensor with specification documents listing sensitivity to both hydrogen and carbon monoxide[10]. Its presence inside a CO alarm is easily confirmed once the front cover is removed. Figure 1 shows the First Alert Model CO1210 with and without its cover, and the Figaro TGS5042 sensor is visible in Fig. 1(b).

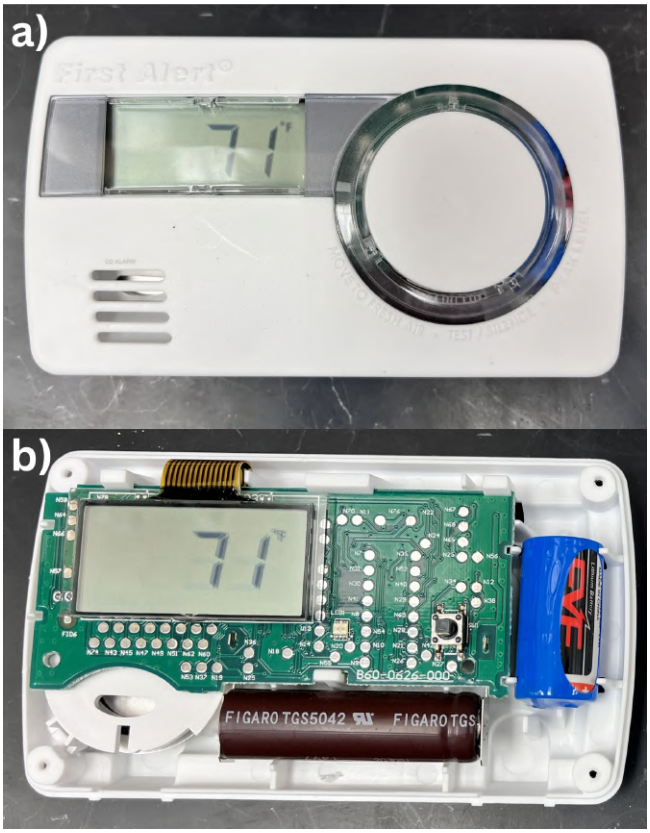

FIG. 1. Images of the First Alert CO detector model CO1210: (a) closed detector and (b) opened detector, with the Figaro TGS5042 electrochemical sensor visible.

II. TESTING METHOD

This testing procedure has two objectives. First, it determines whether the selected CO alarms detect H_2 and D_2 at the regulatory threshold of 25% of the LEL. Second, it evaluates the concentration range over which the alarms produce consistent responses by exposing the detectors to concentrations above and below this threshold. These additional tests identify both the lower concentration range where the response becomes unreliable and the higher concentration range where the alarm time no longer decreases.

A. Detector selection

This test uses three different models of First Alert carbon monoxide detectors: CO1210, CO400, and Explosive Gas. The use of multiple First Alert models allowed the response to be compared across different alarm configurations while keeping the same primary sensor

type. Since the Figaro TGS5042 sensor is the component expected to provide the hydrogen response, the results are expected to be relevant to other commercial alarms using the same sensor, including some X-Sense detectors. However, differences in internal circuitry, filtering, and alarm logic may affect the exact response time between products.

Furthermore, two CO400 models were employed to see if repeated testing altered the reactivity of a sensor. One CO400 detector was reserved for testing only at 25% of the LEL, while the other CO400 detector was exposed across multiple H_2 concentrations to check whether repeated exposure produced any measurable change in the alarm response. No measurable change was observed, as discussed in Section II F.

B. Chamber setup

To test the CO detectors' sensitivity to hydrogen, the detector must be confined in a gas-tight system, exposed to a known concentration of hydrogen gas, and monitored to measure the time required to alarm. This required a system that is gas-tight, large enough to fit the CO detector, and immune to repeated usage. The design seen in Fig. 2(a) meets these criteria. The chamber was assembled with the detector inside. Support structures were employed to create the final product, shown in Fig. 2(b). The window on the top allows observation of the detector's face inside the chamber. This was necessary when detectors with a displayed PPM were tested. In addition to the window, this vacuum chamber contains a port on one end which allows for careful control of gas concentrations. This port remained sealed during normal testing and was opened only during controlled gas-removal or gas-injection steps. When it was necessary to remove or add specific quantities of gas, a syringe was attached and the port was opened.

C. Gas concentration preparation

To accurately set the hydrogen gas concentration, specific withdrawals of normal air and controlled inputs of H_2 or D_2 gas were required. For each trial, the amount of air removed was calculated from the desired final hydrogen partial pressure and the known chamber volume. The removed air volume was larger than the hydrogen injection volume because a small amount of air was intentionally reintroduced after the hydrogen injection as a backfill. This

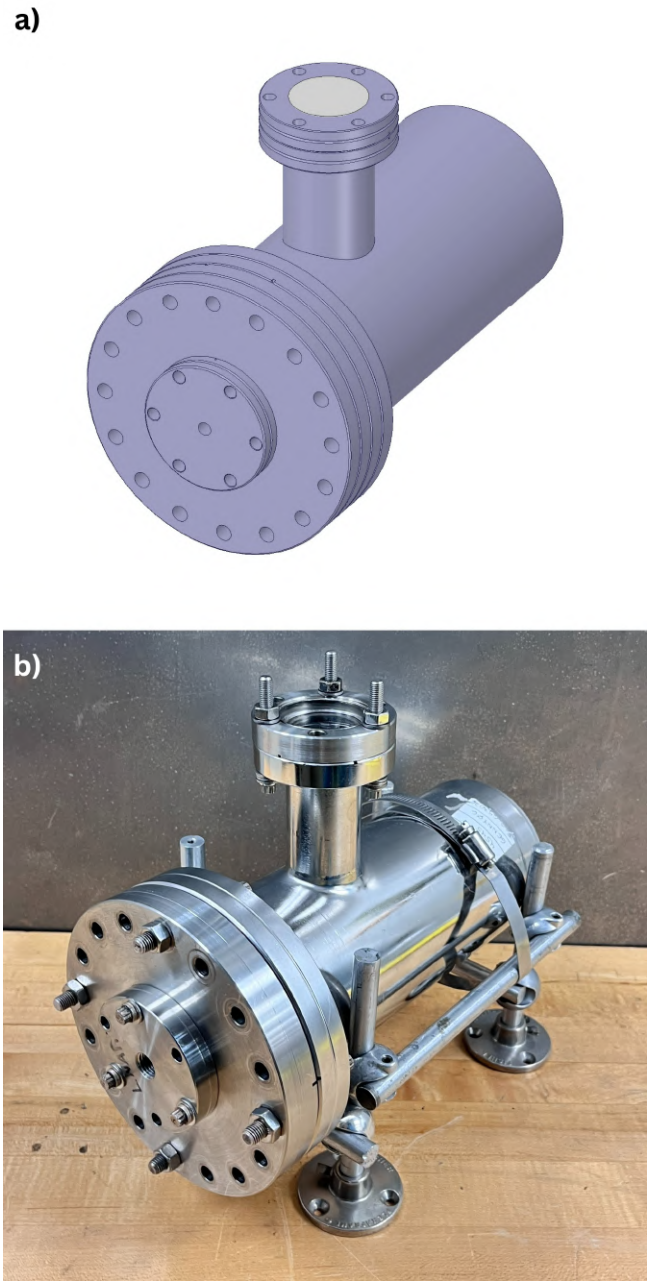

FIG. 2. Hydrogen testing chamber used for detector exposure experiments: (a) CAD design of the chamber and (b) assembled chamber with support structures attached.

backfill step restored the chamber to atmospheric pressure and promoted convective mixing of the gas inside the chamber. Treating the gas inside the chamber as ideal, a computational script was used to calculate the required air-removal, hydrogen-injection, and air-backfill volumes for each target concentration. The injected H_2 or D_2 volume was the quantity used to set the desired final concentration, and the total volume of air removed was calculated

as this gas-injection volume plus an additional backfill volume, typically about 50% of the gas-injection volume.

D. Gas mixing and partial-pressure assumptions

Gas mixing was an important experimental consideration because H_2 or D_2 introduced into the chamber may initially rise toward the highest points of the chamber rather than being uniformly distributed throughout the system. In the limiting case of molecular diffusion without convective mixing, the slowest-decaying one-dimensional diffusion mode scales as

$$exp(-\pi^2 Dt/L^2)\,, \tag{1}$$

where D is the diffusion coefficient, t is time, and L is the relevant chamber length scale. In the chamber geometry, H_2 or D_2 could initially rise toward either the window region or the upper surface of the cylindrical chamber volume. To minimize this effect, the chamber could be oriented so that the window region was on its side, reducing the relevant vertical separation between the highest likely gas location and the sensor to approximately 5 cm. Using $D \approx 7 \times 10^{-5}\,\mathrm{m^2/s}$ for H_2 in air and $L \approx 5\,\mathrm{cm}$ gives an estimated 90% mixing time of about 8 s. If the larger 11 cm separation is used as a conservative upper-bound length scale, the corresponding estimate is about 40 s. These estimates are conservative because the hydrogen injection and subsequent air backfill introduced convective mixing, and because the full hydrogen volume was not introduced at the maximum possible height of the chamber. Since the measured alarm times were generally on the order of several minutes, gas diffusion and mixing were not expected to be the dominant contribution to the observed alarm times. The reported times should therefore be interpreted as the practical response time of the complete commercial alarm under this exposure procedure, including any internal signal averaging or alarm-threshold logic used by the detector.

The dependence of the response on gas partial pressures was also considered. If the detector response is produced by catalytic oxidation at the electrochemical working electrode, the reaction rate may depend on both the hydrogen-isotope partial pressure and the oxygen partial pressure.

For the low-concentration tests, however, oxygen was present in large excess. At the standard test concentration of 1% H_2, the remaining gas was approximately 99% air, corre-

sponding to about 20.8% O_2 by volume. Complete oxidation of 1% H_2 would require only about 0.5% O_2 by volume. Oxygen was therefore present at roughly 40 times the stoichiometric amount required at the 1% H_2 test concentration. For this reason, detector response near the regulatory threshold was treated primarily as a function of the H_2 or D_2 partial pressure, while changes in oxygen partial pressure were assumed to be a small correction.

At higher H_2 concentrations, oxygen displacement becomes more significant, so those tests are interpreted as practical alarm-response tests under hydrogen-rich conditions rather than as isolated measurements of hydrogen partial pressure alone. A separate study in which oxygen and hydrogen partial pressures are independently varied would be required to isolate this dependence directly.

E. Exposure procedure and timing

At each gas concentration, each detector underwent three exposure trials before testing began at a different concentration. The repeated trials were used to evaluate the consistency of the alarm response at each concentration. Throughout the testing process, the chamber, the gas source (H_2 or D_2), and the port remained unaltered. For each test, the appropriate gas concentration was introduced into the chamber, and the time required to alarm was recorded.

The same procedure was repeated for each test. First, the detector was placed inside the chamber and the system was sealed. The required volumes of air to be removed, hydrogen gas to be injected, and air to be reintroduced were calculated in advance. Two syringes were then prepared: one containing the necessary amount of hydrogen gas, and a second, referred to as the backfill syringe, containing the required volume of air for reintroduction. Next, the calculated volume of air was removed using a third syringe. The hydrogen gas was then injected using the first syringe, followed by injection from the backfill syringe to promote thorough mixing and restore the system to standard pressure and concentration.

To limit exchange with room air during this process, the syringe attachment valve was kept closed while syringes were attached or exchanged and was opened only during the withdrawal or injection step. This minimized bulk exchange with room air, although a small amount of ambient air may have remained in the syringe-fitting dead volume. Because this trapped volume was small compared with the chamber volume and the injected gas volumes,

it was not expected to significantly affect the final gas concentration. The chamber pressure also remained below ambient pressure until the final backfill step, further limiting outward gas loss during the removal and injection sequence.

Timing was performed manually using a stopwatch. The stopwatch was started as soon as the hydrogen syringe was fully depressed. The backfill syringe was fully depressed approximately 2 s later. Alarm times were rounded to the nearest 5 s. Thus, the timing uncertainty for each individual trial was taken to be approximately ± 5 s, which conservatively includes the ± 2.5 s rounding uncertainty and the short delay between hydrogen injection and completion of the backfill injection. When the detector produced an alarm, the elapsed time was noted.

F. Chamber reset and repeated-exposure check

Between trials, the chamber was reset by fully opening the system to room air and flushing it with a compressed-air hose three times. This procedure was used to remove residual H_2 or D_2 before the detector was exposed to a new gas concentration.

A repeated-exposure check was also included to determine whether prolonged exposure altered the detector response. One detector was exposed to 5% H_2 for approximately 12 hours. After the chamber was reset using the same flushing procedure described above, this detector was tested again and its alarm time was compared with that of a reserved detector of the same model that had not undergone the extended exposure. This check verified that repeated testing did not produce an observable change in detector response within the concentration levels of the repeated trials.

III. RESULTS AND DISCUSSION

For each detector and gas condition, the repeated alarm times were averaged and the standard deviation was calculated from the individual repeated trials. Unless otherwise stated, reported uncertainties represent one standard deviation of these repeated measurements. The manual timing uncertainty for each individual trial was estimated to be approximately ± 5 s. This value includes the ± 2.5 s uncertainty from rounding alarm times to the nearest 5 s, as well as the short delay between hydrogen injection and completion of the backfill step.

TABLE I. Alarm response times for each detector exposed to 1% H_2 by volume, corresponding to 25% of the hydrogen LEL.

Detector (Model Number)	$t_1\,(min)$	$t_2\,(min)$	$t_3\,(min)$	$t_{avg}(min)$
CO1210	2.75	3.5	2.75	3
CO400(1)	3	2.25	3.25	2.83
CO400(2)[a]	2.25	2.0	3.0	2.42
Explosive Gas	2.75	2.75	2.5	2.67
Overall Average				**2.73**

[a] This detector's first tests were at 25% of the LEL

Since ±5 s corresponds to only ±0.083 min, it was smaller than the standard deviation of the repeated trials, which was typically several tenths of a minute or larger. For this reason, the standard deviation of the repeated trials was used to represent the uncertainty in the reported alarm times. Statistical comparisons were used only as supporting tests because the number of repeated trials was small.

A. Alarm response at 25% of the LEL

The primary performance test was whether the CO alarms produced an alarm at 1% H_2 or D_2 by volume, corresponding to 25% of the hydrogen LEL. All tested Figaro-based alarms produced an alarm at this concentration for both H_2 and D_2. The individual trial response times are shown visually in Fig. 3, where each point represents a single exposure trial and the square markers show the mean response time for each detector. The numerical results are summarized in Table I for H_2 and Table II for D_2. Across all 1% H_2 trials, the average alarm time was 2.73 ± 0.43 min, where the uncertainty represents one standard deviation of the repeated trials. Across all 1% D_2 trials, the average alarm time was 2.84 ± 0.65 min. The larger standard deviation for D_2 indicates greater trial-to-trial variation, but the overall response remained on the same several-minute time scale as H_2. These results demonstrate that the alarms responded at the required concentration threshold.

However, none of the 1% H_2 or 1% D_2 trials produced an alarm within 30 s. Even when the observed trial-to-trial variation is considered, the alarm times remained well above

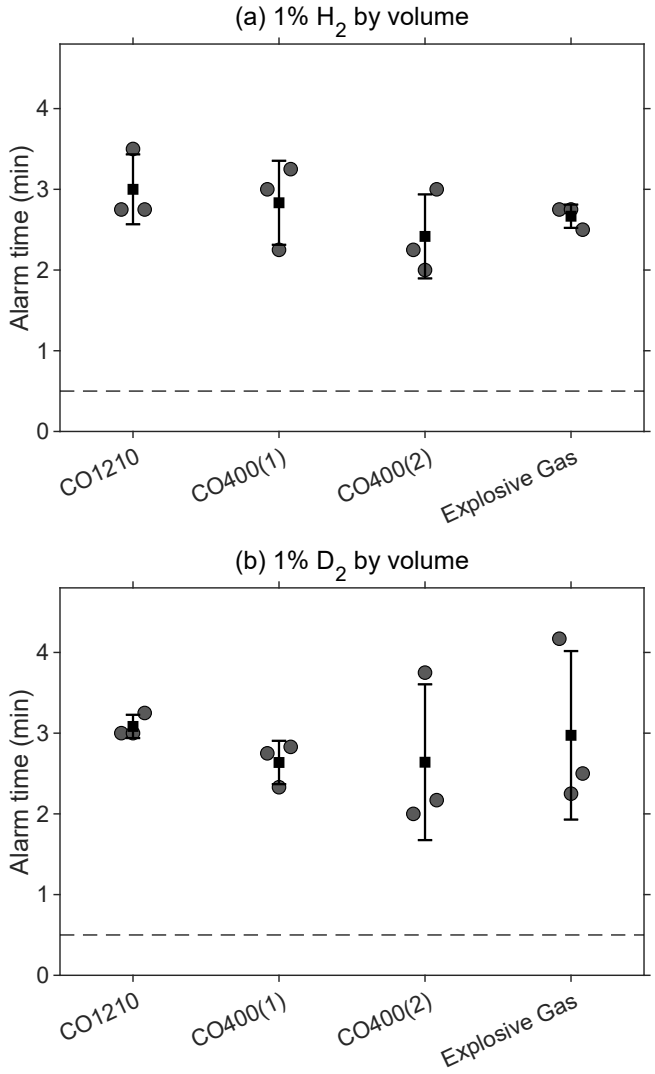

FIG. 3. Alarm response times for commercial CO alarms exposed to 1% gas by volume, corresponding to 25% of the hydrogen LEL: (a) H_2 and (b) D_2. Individual circles show repeated trials, and square markers show the mean response time for each detector. Error bars represent one standard deviation of the repeated trials. The dashed horizontal line indicates the 30 s response time specified by ISO 26142:2010.

the 30 s mark. Thus, while the detectors satisfied the concentration-response portion of the hydrogen alarm requirement, they did not satisfy the ISO 26142:2010 response-time requirement. This distinction is important because the alarms responded reliably, but not rapidly.

TABLE II. Alarm response times for each detector exposed to 1% D_2 by volume, corresponding to 25% of the hydrogen LEL.

Detector (Model Number)	$t_1\,(min)$	$t_2\,(min)$	$t_3\,(min)$	$t_{avg}(min)$
CO1210	3.0	3.0	3.25	3.08
CO400(1)	2.75	2.33	2.83	2.64
CO400(2)[a]	2.0	3.75	2.17	2.67
Explosive Gas	4.17	2.25	2.5	2.97
Overall Average				**2.84**

[a] This detector's first tests were at 25% of the LEL

B. Comparison of H_2 and D_2 response

The detector-level mean response times were also compared to determine whether the alarms responded differently to H_2 and D_2. The average response times for the two gases differed by approximately 0.11 min. A paired comparison, or paired t-test, of the detector-level means showed no statistically significant difference between H_2 and D_2 response times ($p \approx 0.41$). Therefore, within the scatter of these measurements, the alarms responded similarly to H_2 and D_2 at the 1% test concentration.

Because of the limited number of detector models tested, the 95% confidence interval for the paired mean difference remained relatively wide, spanning approximately -0.25 to 0.46 min. Therefore, this result should be interpreted as evidence that no large isotope-dependent response difference was observed in Figaro-based CO alarms, rather than as proof that the two response distributions are identical.

C. Detector-to-detector variation

The use of multiple First Alert alarm models also allowed the response to be compared across different commercial alarm configurations containing the same primary Figaro sensor type. Detector-to-detector variation was evaluated using a one-way ANOVA comparing the repeated alarm times across the four detector models. This comparison did not show a significant model-dependent difference in response time for either H_2 ($p \approx 0.44$) or D_2

($p \approx 0.83$). Because only three trials were performed for each detector, this test was used only as a supporting comparison. However, the result suggests that the central conclusion is not specific to a single alarm model: Figaro-based CO alarms responded at the 1% threshold, but their response times were much longer than 30 s.

D. Hydrogen response range

The H_2 response-range tests were used to determine whether the alarms behaved like calibrated concentration meters or threshold alarms, and to identify the practical lower response limit and upper response plateau. Fig. 4 shows individual alarm times at each H_2 concentration, along with the mean response time and standard deviation of the repeated trials. If the alarms were acting as calibrated hydrogen concentration meters, the response time would be expected to decrease systematically as the H_2 concentration increased. Instead, the alarm times remained clustered on the same several-minute scale for concentrations from approximately 0.5% to 20% H_2. A linear regression of the mean alarm time against $\log_{10}$ of the H_2 concentration over this range gave a slope of approximately -0.075 min per decade, with $p \approx 0.66$, indicating no significant concentration-dependent decrease in alarm time. This indicates that, once the response exceeds the internal alarm threshold, the measured alarm time is likely dominated by the commercial alarm's internal signal processing, averaging, or alarm-threshold logic rather than by hydrogen concentration alone. This interpretation applies to the concentration range tested here. At substantially higher H_2 concentrations, oxygen displacement may become more important, and separating the effects of hydrogen partial pressure from oxygen partial pressure would require a different experimental system and procedure.

At lower H_2 concentrations, Fig. 4 shows that the alarm response became slower and more scattered. This behavior is consistent with the hydrogen-generated electrochemical signal approaching the effective alarm threshold of the commercial CO alarm. At 0.1% H_2, only one of the three trials produced a measured alarm response. Therefore, these lower-concentration measurements should be interpreted as evidence of a practical lower response range rather than as a calibrated lower detection limit.

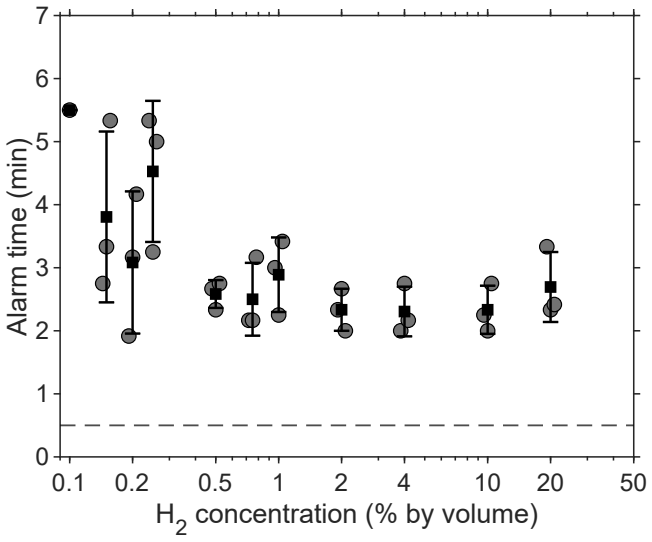

FIG. 4. Alarm response time as a function of H_2 concentration by volume. Individual circles show repeated trials, and square markers show the mean response time at each concentration. Error bars represent one standard deviation of the repeated trials. The dashed horizontal line indicates the 30 s response time specified by ISO 26142:2010.

E. Practical interpretation

Taken together, these results show that Figaro-based CO alarms can provide a low-cost indication of H_2 or D_2 reaching the 1% threshold concentration, but they are not suitable when rapid hydrogen detection is required. Their several-minute response time means they may be useful for supplemental monitoring of slowly developing hydrogen-isotope leaks, especially in low-cost or noncritical laboratory applications. However, because none of the tested alarms met the 30 s ISO response-time requirement, they should not replace dedicated hydrogen detectors in applications where rapid warning, calibrated concentration measurement, or formal ISO 26142:2010 compliance is required.

IV. CONCLUSION

Commercial carbon monoxide alarms containing Figaro TGS5042 electrochemical sensors were evaluated as low-cost hydrogen-isotope alarms for protium (H_2) and deuterium (D_2). Hydrogen has a lower explosive limit (LEL) of 4% by volume. Under UL 2075, *Standard for Gas and Vapor Detectors and Sensors*, hydrogen detectors are required to alarm at or below

25% of the LEL, which corresponds to 1% hydrogen by volume[7]. ISO 26142:2010, *Hydrogen Detection Apparatus—Stationary Applications*, further specifies that a hydrogen detection apparatus should alarm within 30 s when exposed to hydrogen at 25% of the LEL[8]. These two requirements provide the central comparison points for this work: whether the alarms respond at 1% H_2 or D_2, and whether that response occurs within 30 s.

Each CO alarm equipped with a Figaro TGS5042 sensor alarmed at 1% H_2 and 1% D_2 by volume, demonstrating sensitivity at the regulatory threshold concentration (See Fig. 3). At this concentration, the average alarm response time was 2.73 $\pm$ 0.43 min for H_2 and 2.84 $\pm$ 0.65 min for D_2. Although the alarms met the concentration-response requirement, none of the tested detectors alarmed within the 30 s response time specified by ISO 26142:2010.

The hydrogen response-range tests, which examined alarm behavior below and above the 1% threshold concentration, also showed that these alarms should be interpreted as threshold alarms rather than calibrated hydrogen concentration meters. The alarm time did not systematically decrease at concentrations above 1% H_2, suggesting that the response was limited by the internal signal processing, averaging, or alarm-threshold logic of the commercial CO alarm rather than by hydrogen concentration alone. At lower concentrations, the alarms continued to respond until the concentration became low enough that response times increased or the alarm response became less reliable. This lower-limit behavior is visible near 0.1% H_2 in Fig. 4, where the response became more scattered and some trials did not produce a measured alarm response.

Therefore, Figaro-based commercial CO alarms provide a low-cost and widely accessible means of monitoring gradually increasing hydrogen-isotope concentrations. In this context, "gradually increasing" refers to leaks for which the hydrogen concentration increases slower than about 1% H_2 by volume per minute. For faster leaks, or for applications where ISO 26142:2010 response-time compliance is required, dedicated hydrogen detection remains necessary.

ACKNOWLEDGMENTS

This work was supported in part by the College of Computational, Mathematical and Physical Sciences, specifically by the Department of Physics and Astronomy at Brigham Young University. The first author was supported as a research assistant by the College of

Computational, Mathematical, and Physical Sciences. The authors gratefully acknowledge David Allred for his guidance as a research colleague, including reviewing research progress and encouraging the development of this manuscript. The authors also thank the technical staff at X-Sense, First Alert, and Kidde Global Solutions who provided valuable information about the electrochemical sensors through personal communications.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

[1]U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, "Hydrogen safety fact sheet," `https://www1.eere.energy.gov/hydrogenandfuelcells/pdfs/h2_safety_fsheet.pdf` (2004), accessed 17 August 2025.

[2]K. O'Malley, H. Lopez, J. Cairns, R. Wichert, C. Rivkin, R. Burgess, and W. Buttner, "Overview of north american hydrogen sensor standards," Technical Report NREL/TP-5400-62062 (National Renewable Energy Laboratory, Golden, CO, 2015) accessed 18 January 2026.

[3]N. D. Marsh and T. G. Cleary, "Towards a test method for hydrogen sensor performance," in *NHA Annual Hydrogen Conference 2008* (Sacramento, CA, 2008) accessed 18 January 2026.

[4]National Fire Protection Association, "NFPA 715, standard for the installation of fuel gases detection and warning equipment," Standard NFPA 715 (National Fire Protection Association, Quincy, MA, 2023) accessed 25 July 2025.

[5]A. Joshi, F. Sattari, M. A. Khan, and L. M. Lefsrud, "Gap analysis of codes and standards for hydrogen refueling stations," Technical Report (University of Alberta, 2024) accessed 18 January 2026.

[6]T. G. Cleary, W. L. Grosshandler, and A. A. Chernovsky, "Smoke detector response to nuisance aerosols," in *Proceedings of the 11th International Conference on Automatic Fire Detection, AUBE '99*, edited by H. Luck (Duisburg, Germany, 1999) pp. 32–41, accessed 18 January 2026.

[7]Underwriters Laboratories, "UL 2075, standard for gas and vapor detectors and sensors," Standard UL 2075 (UL Standards & Engagement, Northbrook, IL, 2013) revised 23 January 2023, accessed 25 July 2025.

[8]International Organization for Standardization, "ISO 26142:2010, hydrogen detection apparatus—stationary applications," Standard ISO 26142:2010 (International Organization for Standardization, Geneva, Switzerland, 2010) accessed 1 September 2025.

[9]M. Majder-Łopatka, T. Wesierski, A. Dmochowska, Z. Salamonowicz, and A. Polańczyk, Sustainability **12**, 14 (2020).

[10]Figaro Engineering Inc., "TGS5042 product information," `https://www.figarosensor.com/product/docs/tgs5042_productinfomation(fusa)_rev09.pdf` (2024), rev. 09, accessed 5 July 2025.

Appendix L

Nuclear Modeling Code

This appendix contains the Wolfram Language code used to construct the nuclear potential models and calculate the tunneling probabilities and cross sections discussed in this thesis. The calculations were developed in *Mathematica* 14.2. Only input code required to document the models is included here; notebook outputs, repeated numerical results, front-end metadata, and diagnostic messages have been omitted.

The code has been separated according to the physical model being calculated. Each listing begins by clearing the global workspace and then defines the constants, model parameters, potential functions, and numerical calculations needed for that section. Variable names containing Greek characters in the original notebook were changed to equivalent ASCII names to improve source-code portability and LaTeX compatibility. These changes do not intentionally alter the mathematical expressions.

The full working notebook should be retained with the research files. The listings below are intended to provide a readable record of the calculations used in the thesis rather than reproduce every exploratory calculation that appeared during development.

L.1 Nuclear Potential Models

The following code defines the square-well, Gaussian, and Woods-Saxon potentials used to represent the nuclear interaction. The uncalibrated Yukawa implementation was not included because it was not used for the final calculations.

```
(* ::Package:: *)
(* No-Screening Nuclear Potential Models *)
(* Extracted and reformatted from LNAR_Models.nb (Mathematica 14.2). *)
(* Only input cells are included; notebook outputs and messages are omitted. *)
(* Repeated Clear cells were consolidated into the single line below. *)

ClearAll["Global`*"];

(* NOTE: The uncalibrated Yukawa code was moved to Code/Exploratory and is not
    included here. *)


(* SUBSECTION: Values - Run First *)

(*Values to apply *)
(* Constants *)
hbar = 1.0545718*(10)^(-34); (* reduced Planks constant [J*s]  *)
k = 8.99*(10)^(9);               (* (1)/(4*Subscript[Piepsilon, 0]) Boltmann constnt
      [(N*(m)^(2))/((C)^(2))] *)
e = 1.602*10^-19;    (* Charge of an electron    [C]   *)

(* Experimental D2 Values *)
V0eV = 35*(10)^(6);             (* Nuclear Potential of Deuterium [eV] *)
V0J = (V0eV)/(6.242*(10)^(18));    (* Nuclear Potential of Deuterium [J] *)
m = 3.3435*(10)^(-27);        (* Experimental Mass of deuterium nucleus [kg] *)
```

```
24 R0exp = 2.12562*(10)^(-15); (* Experimental charge radius of deuteron nucleus [m] *)
25
26 (* Analytical D2 Values *)
27 a = 0.5*(10)^(-15);    (* other nuclear dist. values *)
28 r0 = 1.25*(10)^(-15);(* Nuclear dist. constant [m] *)
29 R = r0*(2)^(1/3);        (* Theoretical nuclear charge radius of detuerium [m]*)
30
31 (* ========================================================================= *)
32 (* SECTION: Potential Modeling *)
33 (* ========================================================================= *)
34
35 (* SUBSECTION: Square-Well *)
36
37 (* Variable *)
38 EngeV = .1*(10)^(6);                      (* Energy of incoming particle [eV] *)
39 Eng =EngeV/(6.242*(10)^(18));   (* Energy of incoming particle [J] *)
40
41 (*Define Potential Functions*)
42 CoulombPoteV[r_]:=((k*(e)^(2)))/(r)*(1)/((1.602*(10)^(-19))) (* Coulomb Potential in
       eV *)
43 (* Define Nuclear Potential Well as a finite square well *)
44 SquareWelleV[r_]:=Piecewise[{{-V0eV,r<R0exp},{0,r>R0exp}}] (* Nuclear Potential in eV
       *)
45
46 (*Total potential energy*)
47 TotalSQRPotentialeV[r_]:=CoulombPoteV[r]+SquareWelleV[r]
48
49 (*Plot with line for energy *)
50 Plot[{CoulombPoteV[r],SquareWelleV[r],TotalSQRPotentialeV[r],EngeV},{r,0,20*(10)^(-15)
       },PlotLegends->{"Coulomb Potential","Nuclear Potential Well","Total Potential", "
```

```
      KE"},AxesLabel->{"r (m)","V(r) (eV)"},PlotRange->{-1*(10)^(6),1*(10)^(6)},
      PlotStyle->{Dashed,Dotted,Thick}]

(* SUBSECTION: Gaussian Nuclear Potential *)

(* Variable *)
EngeV = .1*(10)^(6);                        (* Energy of incoming particle [eV] *)
Eng =EngeV/(6.242*(10)^(18));    (* Energy of incoming particle [J] *)

(*Define Potential Functions*)
CoulombPoteV[r_]:=((k*(e)^(2)))/(r)*(1)/((1.602*(10)^(-19))) (* Coulomb Potential in
      eV *)

(*Define Strong Potential Function*)
GaussPoteV[r_]:=-V0eV*Exp[-(((r)/(R0exp)))^(2)] (* Nuclear Potential in eV *)

(*Total potential energy*)
TotalGSSPotentialeV[r_]:=CoulombPoteV[r]+GaussPoteV[r]

(*Plot*)
Plot[{CoulombPoteV[r],GaussPoteV[r],TotalGSSPotentialeV[r], EngeV},{r,0,30*(10)^(-15)
      },PlotLegends->{"Coulomb","Nuclear (Gaussian)","Total Potential", "KE (req)"},
      AxesLabel->{"r (m)","V(r) (eV)"},PlotRange->{-1*(10)^(6),1*(10)^(6)},PlotStyle->{
      Dashed,Dotted,Thick}]

(* SUBSECTION: Woods Saxon *)

(* Variable *)
EngeV = .004*(10)^(6);                          (* Energy of incoming particle [eV] *)
Eng =EngeV/(6.242*(10)^(18));    (* Energy of incoming particle [J] *)
```

```
75
76 (*Define Potential Functions*)
77 CoulombPoteV[r_]:=((k*(e)^(2)))/(r)*(1)/((1.602*(10)^(-19))) (* Coulomb Potential in
       eV *)
78
79 (* Strong force modeled by the Woods Saxon: a mean field potential for nucleons inside
        atomic nucleus *)
80 WoodsPoteV[r_]:= -(V0eV)/(1+Exp[(r-R0exp)/(a)]); (* Nuclear Potential in eV *)
81
82 (*Total potential energy*)
83 TotalWDSPotentialeV[r_]:=CoulombPoteV[r]+WoodsPoteV[r]
84
85 (*Plot*)
86 Plot[{CoulombPoteV[r],WoodsPoteV[r],TotalWDSPotentialeV[r], EngeV},{r,0,30*(10)^(-15)
       },PlotLegends->{"Coulomb","Nuclear (P Well)","Total Potential", "KE (req)"},
       AxesLabel->{"r (m)","V(r) (eV)"},PlotRange->{-1*(10)^(6),1*(10)^(6)},PlotStyle->{
       Dashed,Dotted,Thick}]
```

Listing L.1 Implementation and plotting of the no-screening nuclear potential models.

L.2 WKB Tunneling and Cross-Section Calculations

The next listing defines the combined nuclear and Coulomb potentials, determines the classical turning points, evaluates the WKB integral, and calculates the corresponding tunneling probabilities and cross sections.

```
1 (* ::Package:: *)
2 (* WKB Tunneling and Cross-Section Calculations *)
3 (* Extracted and reformatted from LNAR_Models.nb (Mathematica 14.2). *)
4 (* Only input cells are included; notebook outputs and messages are omitted. *)
```

```
(* Repeated Clear cells were consolidated into the single line below. *)

ClearAll["Global`*"];

(* NOTE: This file preserves the numerical method and parameter choices from the
    notebook. *)


(* SUBSECTION: Values - Run First *)

(*Values to apply *)
(* Constants *)
hbar = 1.0545718*(10)^(-34); (* reduced Planks constant [J*s]  *)
k = 8.99*(10)^(9);               (* (1)/(4*Subscript[Piepsilon, 0]) Boltmann constnt
      [(N*(m)^(2))/((C)^(2))] *)
e = 1.602*10^-19;    (* Charge of an electron    [C]   *)

(* Experimental D2 Values *)
V0eV = 35*(10)^(6);             (* Nuclear Potential of Deuterium [eV] *)
V0J = (V0eV)/(6.242*(10)^(18));    (* Nuclear Potential of Deuterium [J] *)
m = 3.3435*(10)^(-27);        (* Experimental Mass of deuterium nucleus [kg] *)
R0exp = 2.12562*(10)^(-15); (* Experimental charge radius of deuteron nucleus [m] *)

(* Analytical D2 Values *)
a = 0.5*(10)^(-15);    (* other nuclear dist. values *)
r0 = 1.25*(10)^(-15);(* Nuclear dist. constant [m] *)
R = r0*(2)^(1/3);        (* Theoretical nuclear charge radius of detuerium [m]*)

(* ========================================================================== *)
(* SECTION: WKB, sigma(E), and Rate *)
```

```
(* ========================================================================== *)

(* SUBSECTION: Square-Well *)

(* STEP: Functions and Variables *)

(* Variable *)
EngeV = .1*(10)^(6);                        (* Energy of incoming particle [eV] *)
Eng =EngeV/(6.242*(10)^(18));   (* Energy of incoming particle [J] *)

(*Define Potential Functions*)
CoulombPoteV[r_]:=((k*(e)^(2)))/(r)*(1)/((1.602*(10)^(-19))) (* Coulomb Potential in
    eV *)
CoulombPotJ[r_]:=((k*(e)^(2)))/(r) (* Coulomb Potential in J *)

(* Define Nuclear Potential Well as a finite square well *)
SquareWelleV[r_]:=Piecewise[{{-V0eV,r<R0exp},{0,r>R0exp}}] (* Nuclear Potential in eV
    *)
SquareWellJ[r_]:=Piecewise[{{-V0J,r<R0exp},{0,r>R0exp}}]   (* Nuclear Potential in J
    *)

(*Total potential energy *)
TotalSQRPotentialJ[r_]:=CoulombPotJ[r]+SquareWellJ[r] (* Total Potentials Joules *)
TotalSQRPotentialeV[r_]:=CoulombPoteV[r]+SquareWelleV[r] (* Total Potentials eV *)

(* Tunneling Fuctions *)
tunnelProbability[beta_]:=Exp[-2*beta]*((1-(1)/(4)*Exp[-2*beta]))^(-2);
tunnelSimp[beta_]:=Exp[-2*beta];

(* STEP: Tunneling Prep *)
```

```
60
61 (* Turning Points *)
62 xRsquare=r/.FindRoot[TotalSQRPotentialeV[r]== EngeV,{r,1.4*(10)^(-14)}]//Quiet
63 xLsquare=2.4*(10)^(-15)
64
65 (*Plot*)
66 SQRTurnPointPlot = ListPlot[{{xRsquare, EngeV},{xLsquare,EngeV}}, Filling->Axis,
       PlotMarkers->{Automatic, Medium}];
67 SQRPotentialPlot = Plot[{CoulombPoteV[r],SquareWelleV[r],TotalSQRPotentialeV[r], EngeV
       },{r,0,30*(10)^(-15)},PlotLegends->{"Coulomb Potential","Nuclear Potential Well","
       Total Potential", "KE"},AxesLabel->{"r (m)","V(r) (eV)"},PlotRange->{-1*(10)^(6)
       ,1*(10)^(6)},PlotStyle->{Dashed,Dotted,Thick}];
68 Show[SQRPotentialPlot, SQRTurnPointPlot]
69
70 (* STEP: Tunneling Calculation *)
71
72 (* Calculate momentum *)
73 psquare = Sqrt[2*m*(Eng-TotalSQRPotentialJ[r])];
74 (* Calculate beta integral *)
75 betasquare = (1)/(hbar)*NIntegrate[Abs[psquare], {r,xLsquare,xRsquare}]
76
77 (* Plug into tunneling approximation Formula *)
78 tunnelProbability[betasquare]
79 (* Beta is not >> 1 and therefore we must keep the above whole probability *)
80 tunnelSimp[betasquare];
81
82 (* SUBSECTION: Gaussian Nuclear Potential *)
83
84 (* STEP: Functions and Variables *)
85
```

```
(* Variable *)
EngeV = .1*(10)^(6);                        (* Energy of incoming particle [eV] *)
Eng =EngeV/(6.242*(10)^(18));   (* Energy of incoming particle [J] *)

(*Define Potential Functions*)
CoulombPoteV[r_]:=((k*(e)^(2)))/(r)*(1)/((1.602*(10)^(-19))) (* Coulomb Potential in
    eV *)
CoulombPotJ[r_]:=((k*(e)^(2)))/(r) (* Coulomb Potential in J *)

(* Nuclear Potentials *)
GaussPoteV[r_]:=-V0eV*Exp[-(((r)/(R0exp)))^(2)] (* Nuclear Potential in eV *)
GaussPoteJ[r_]:=-V0J*Exp[-(((r)/(R0exp)))^(2)]

(*Total potential energy*)
TotalGSSPotentialeV[r_]:=CoulombPoteV[r]+GaussPoteV[r]
TotalGSSPotentialJ[r_]:=CoulombPotJ[r]+GaussPoteJ[r]

(* Tunneling Fuctions *)
tunnelProbability[beta_]:=Exp[-2*beta]*((1-(1)/(4)*Exp[-2*beta]))^(-2);
tunnelSimp[beta_]:=Exp[-2*beta];

(* STEP: Tunneling Prep *)

(* Turning Points *)
xRgauss=r/.FindRoot[TotalGSSPotentialeV[r]== EngeV,{r,1.4*(10)^(-14)}]//Quiet
xLgauss=r/.FindRoot[TotalGSSPotentialeV[r]==EngeV,{r,4.7*(10)^(-15)}]//Quiet

(*Plot*)
GSSTurnPointPlot = ListPlot[{{xRgauss, EngeV},{xLgauss,EngeV}}, Filling->Axis,
    PlotMarkers->{Automatic, Medium}];
```

```
GSSPotentialPlot = Plot[{CoulombPoteV[r],GaussPoteV[r],TotalGSSPotentialeV[r], EngeV
    },{r,0,30*(10)^(-15)},PlotLegends->{"Coulomb Potential","Nuclear Potential Well","
    Total Potential", "KE"},AxesLabel->{"r (m)","V(r) (eV)"},PlotRange->{-1*(10)^(6)
    ,1*(10)^(6)},PlotStyle->{Dashed,Dotted,Thick}];
Show[GSSPotentialPlot, GSSTurnPointPlot]

(* STEP: Tunneling Calculation *)

(* Calculate momentum *)
pgauss = Sqrt[2*m*(Eng-TotalGSSPotentialJ[r])];
(* Calculate beta integral *)
betagauss = (1)/(hbar)*NIntegrate[Abs[pgauss], {r,xLgauss,xRgauss}]

(* Plug into tunneling approximation Formula *)
tunnelProbability[betagauss]
(* Beta is not >> 1 and therefore we must keep the above whole probability *)
tunnelSimp[betagauss];

(* SUBSECTION: Woods Saxon *)

(* STEP: Functions and Variables *)

(* Variable *)
EngeV = .004*(10)^(6);                          (* Energy of incoming particle [eV] *)
Eng =EngeV/(6.242*(10)^(18));   (* Energy of incoming particle [J] *)

(*Define Potential Functions*)
CoulombPoteV[r_]:=((k*(e)^(2)))/(r)*(1)/((1.602*(10)^(-19))) (* Coulomb Potential in
    eV *)
CoulombPotJ[r_]:=((k*(e)^(2)))/(r) (* Coulomb Potential in J *)
```

```
140
141 (* Strong force modeled by the Woods Saxon: a mean field potential for nucleons inside
        atomic nucleus *)
142 WoodsPoteV[r_]:= -(V0eV)/(1+Exp[(r-R0exp)/(a)]); (* Nuclear Potential in eV *)
143 WoodsPotJ[r_]:= -(V0J)/(1+Exp[(r-R0exp)/(a)]); (* Nuclear Potential in J *)
144
145 (*Total potential energy*)
146 TotalWDSPotentialeV[r_]:=CoulombPoteV[r]+WoodsPoteV[r]
147 TotalWDSPotentialJ[r_]:=CoulombPotJ[r]+WoodsPotJ[r]
148
149 (* Tunneling Fuctions *)
150 tunnelProbability[beta_]:=Exp[-2*beta]*((1-(1)/(4)*Exp[-2*beta]))^(-2);
151 tunnelSimp[beta_]:=Exp[-2*beta];
152
153 (* STEP: Tunneling Prep *)
154
155 (* Turning Points *)
156 xRwoods=r/.FindRoot[TotalWDSPotentialeV[r]== EngeV,{r,1.0*(10)^(-14)}]//Quiet
157 xLwoods=r/.FindRoot[TotalWDSPotentialeV[r]==EngeV,{r,4.7*(10)^(-15)}]//Quiet
158
159 (*Plot*)
160 WDSTurnPointPlot = ListPlot[{{xRwoods, EngeV},{xLwoods,EngeV}}, Filling->Axis,
       PlotMarkers->{Automatic, Medium}];
161 WDSPotentialPlot = Plot[{CoulombPoteV[r],WoodsPoteV[r],TotalWDSPotentialeV[r], EngeV
       },{r,.36*(10)^(-12),10*(10)^(-12)},PlotLegends->{"Coulomb","Nuclear (P Well)","
       Total Potential", "KE (req)"},AxesLabel->{"r (m)","V(r) (eV)"},PlotRange
       ->{0,10000},PlotStyle->{Dashed,Dotted,Thick}];
162 Show[WDSPotentialPlot, WDSTurnPointPlot]
163
164 (* STEP: Approximation *)
```

```
(* Calculate momentum *)
pwoods = Sqrt[2*m*(Eng-TotalWDSPotentialJ[r])];
(* Calculate beta integral *)
betawoods = (1)/(hbar)*NIntegrate[Abs[pwoods], {r,xLwoods,xRwoods}];

(* Plug into tunneling approximation Formula *)
Twoods = tunnelProbability[betawoods]
(* Beta is not >> 1 and therefore we must keep the above whole probability *)
tunnelSimp[betawoods];

(* STEP: sigma(E) *)

(* Cross Section calculation, assume Subscript[P, nuc] is one for maximum value *)
sigmaWoods = (Pi*(hbar)^(2))/(2*m)*(1)*(Twoods)/(Eng) *(10)^(28)(* NOTE THE CONVERSION
    TO BARNS *)
```

Listing L.2 WKB tunneling and cross-section calculations for the no-screening potential models.

L.3 Atomic Electron Screening Model

The following code constructs the ground-state electron wavefunction, calculates its charge density, and estimates the resulting screening potential. This model was used to examine the magnitude of screening produced by an isolated atomic electron.

```
(* ::Package:: *)
(* Atomic Electron Screening Model *)
(* Extracted and reformatted from LNAR_Models.nb (Mathematica 14.2). *)
(* Only input cells are included; notebook outputs and messages are omitted. *)
```

```
(* Repeated Clear cells were consolidated into the single line below. *)

ClearAll["Global`*"];

(* NOTE: The notebook describes this model as exploratory and reports very little
    screening. *)


(* ========================================================================== *)
(* SECTION: Electron Wavefunction (unmodified to incoming nucleus) *)
(* ========================================================================== *)

(* SUBSECTION: Electron Plots *)

(* STEP: Creating the Basics *)

Asquare=Asquare/.Solve[Assuming[a0>0,Integrate[Asquare*(Exp[(-r)/(a0)])^(2)*(r)^(2), {
    r,0,Infinity}]==1], Asquare][[1]]
Anorm = Sqrt[Asquare]

(*
Wavefunction:
psi = Subscript[R, nl]*(r) Subscript[(Y)^(m), l]*(theta,phi)
*)

(* For the ground state, n = 1, l = m = 0 *)
elecR0= (Anorm)*Exp[(-r)/(a0)];
elecY0 = (1)/(2*Sqrt[Pi]);

(* Finalized wave function *)
```

```
33 psielec[r_] := elecR0*elecY0/.{a0->5.29177*(10)^(-11)};
34
35 (* STEP: Plot Wavefunction *)
36
37 Plot[psielec[r], {r,0,0.25*(10)^(-9)}, PlotRange->All]
38
39 (* Normalization is to ensure the sum of all probabilities is 1. Note the addition of
       the Jacobian (r)^(2)*Sin[theta] *)
40 Integrate[Integrate[Integrate[(Abs[psielec[r]])^(2)*Sin[theta]*(r)^(2), {r,0,Infinity
       }],{theta,0,Pi}],{phi,0,2*Pi}]
41
42 (* STEP: Radial and Charge *)
43
44 (* Radial Probability Density *)
45 P[r_] := Assuming[r \[Element] Reals, (Abs[psielec[r]])^(2)*4*Pi*(r)^(2)]
46 Plot[P[r],{r,0,0.25*(10)^(-9)}, PlotLabel->"Radial probability", AxesLabel->{"r (m)","
       P (1/m)"}]
47
48 (* Charge Density *)
49 e=1.602*10^-19; (*C*)
50 rho[r_] := (Abs[psielec[r]])^(2)*(-e)
51 Plot[rho[r],{r,0,0.15*(10)^(-9)}, PlotLabel->"Charge Density",AxesLabel->{"r (m)","rho
        (C/m^3)"}, PlotRange->All]
52
53 (* SUBSECTION: Finding the Coulomb Forces *)
54
55 (* STEP: Finding the potential *)
56
```

```
(* this is to modify the various fuctions so that I can have primes already and take out the real r which may be messing with our system since we already will use real values *)

rhodens = -3.441198773476452`*^11 (E)^(-3.779453755548711`*^10 r);(* use better units for it *)
rhodensInt = rhodens/.r->rp;
rhodensCart = rhodens/.r->Sqrt[(x)^(2)+(y)^(2)+(z)^(2)];
rhodensIntCart = rhodensCart/.{x->xp,y->yp,z->zp};

(* Just for constants *)
epsilon0 = 8.8541878*(10)^(-12);

(* This is the integral for just the theta, remember that we already took out the phi which just adds a factor of 2 Pi. Here is how we took it out:

The only portion that had its factor is in the bottom:
Sqrt[(r)^(2)+(rp)^(2)-(2*r*rp*(Sin[theta]*Sin[thetap]*Cos[phi-phip]+Cos[theta]*Cos[thetap])

To take this out, we took the incoming particle (non-prime) to be in a straight line at theta = 0, this would make the entire first term above 0 and the second term's Cos[theta]=1 therefore we simplify to:

Sqrt[(r)^(2)+(rp)^(2)-2*rrpCos[thetap]]
The integral over phi would then just add a factor of 2*Pi


Assumptions have been added to ensure a quick and simply calculation *)

```

```
ThetaIntegral=(2*Pi)/(4*Pi*epsilon0)*Integrate[(rhodensInt)/(Sqrt[(r)^(2)+(rp)^(2)-(2*r*rp*Cos[thetap])])*(rp)^(2)*Sin[thetap], {thetap,0,Pi}, Assumptions -> {r>=0 ,r \[Element] Reals, rp>=0 , rp \[Element] Reals, r!=rp}]//FullSimplify (* add exlusion for the equals to look up more details but we also have the other thing*)

(* Now this is the radial integral. This is taken only to 1 since beyond that point there is minimal additions to the system. I do not print the result because it is a messy piece wise, but we can plot it. *)

Vfin[r_]:=Assuming[r>=0  && r \[Element] Reals, Integrate[ThetaIntegral, {rp,0,1}]]

Plot[Vfin[r],{r,0,1.5*(10)^(-9)}, PlotRange->All]

(* Turn the potential into a coulomb by multiplying by the incoming charge (e) *)

eCoulombJ[r_]:=+e*Vfin[r]

Plot[eCoulombJ[r],{r,0,1.5*(10)^(-9)}, PlotRange->All]

(* STEP: Plotting effective screening *)

(*Constants*)
k=8.99*10^9; (*N***m^2/C^2*)
e=1.602*10^-19; (*C*)
V0eV=35*10^6; (*depth in eV*)
V0J=V0eV/(6.242*(10)^(18)); (*Convert to joules*)

(*Gaussian Potential Parameters*)
r0=1.25*10^-15; (*fm scaling*)
R0=r0*(2)^(1/3); (*Effective interaction range for deuteron*)
```

```
105 R0exp = 2.12562*(10)^(-15); (* experimental charge radius of deuteron nucleus *)
106
107 (*Define Potential Functions*)
108 classCoulombeV[r_]:=((k*(e)^(2)))/(r)/(1.602*(10)^(-19))
109 classStrongWelleV[r_]:=Piecewise[{{-V0eV,r<R0exp},{0,r>R0exp}}]
110
111 eCoulombeV[r_] := eCoulombJ[r]/(1.602*(10)^(-19))
112
113 (*Total potential energy*)
114 totalPotentialeV[r_]:=classCoulombeV[r]+classStrongWelleV[r]+eCoulombeV[r]
115
116 (*Plot*)
117 Plot[{classCoulombeV[r],classStrongWelleV[r],eCoulombeV[r],totalPotentialeV[r],
        0.1*(10)^(6)},{r,0,20*(10)^(-15)},PlotLegends->{"Coulomb","Nuclear (P Well)","
        Electron Screening","Total Potential", "KE (req)"},AxesLabel->{"r (m)","V(r) (eV)"
        },PlotRange->{-100,1*(10)^(6)},PlotStyle->{Dashed,Dashed,{Thick,Dashed},Thick,
        Dotted}]
118
119 (* SUBSECTION: Finding the Coulomb Forces V2.0 *)
120
121 (* STEP: Finding the potential *)
122
123 Efield =(1)/(epsilon0*(rp)^(2))*Integrate[(-e)/(4*Pi)*(4)/((a0)^(3))*Exp[(-2*r)/(a0)
        ]*(r)^(2),{r,0,rp}, Assumptions -> {r \[Element] Reals, rp \[Element] Reals}]
124
125 V = Assuming[r>0 && r<1000,Integrate[Efield, {rp,r,1000}]]
126
127 coloumb2[r_]:=e*V/.{a0->5.29177*(10)^(-11)}
128
129 (* Plotting just the potential *)
```

```
Plot[eCoulombeV[r],{r,0,10}]

(* STEP: Plotting effective screening *)

(*Constants*)
k=8.99*10^9; (*N***m^2/C^2*)
e=1.602*10^-19; (*C*)
V0eV=35*10^6; (*depth in eV*)
V0J=V0eV/(6.242*(10)^(18)); (*Convert to joules*)

(*Gaussian Potential Parameters*)
r0=1.25*10^-15; (*fm scaling*)
R0=r0*(2)^(1/3); (*Effective interaction range for deuteron*)
R0exp = 2.12562*(10)^(-15); (* experimental charge radius of deuteron nucleus *)

(*Define Potential Functions*)
classCoulombeV[r_]:=((k*(e)^(2)))/(r)/(1.602*(10)^(-19))
classStrongWelleV[r_]:=Piecewise[{{-V0eV,r<R0exp},{0,r>R0exp}}]

eCoulombeV[r_] := coloumb2[r]/(1.602*(10)^(-19))

(*Total potential energy*)
totalPotentialeV[r_]:=classCoulombeV[r]+classStrongWelleV[r]+eCoulombeV[r]

(*Plot*)
Plot[{classCoulombeV[r],classStrongWelleV[r],eCoulombeV[r],totalPotentialeV[r],
    0.1*(10)^(6)},{r,0,20*(10)^(-15)},PlotLegends->{"Coulomb","Nuclear (P Well)","
    Electron Screening","Total Potential", "KE (req)"},AxesLabel->{"r (m)","V(r) (eV)"
    },PlotRange->{-1*(10)^(6),1*(10)^(6)},PlotStyle->{Dashed,Dashed,{Thick,Dashed},
    Thick,Dotted}]
```

Listing L.3 Ground-state electron-density and atomic-screening calculations.

L.4 Debye Screening Model

The Debye model estimates the screened Coulomb interaction using the electron density calculated for titanium and the corresponding Debye length.

```
(* ::Package:: *)
(* Debye Screening Model for Titanium *)
(* Extracted and reformatted from LNAR_Models.nb (Mathematica 14.2). *)
(* Only input cells are included; notebook outputs and messages are omitted. *)
(* Repeated Clear cells were consolidated into the single line below. *)

ClearAll["Global`*"];

(* NOTE: The intermediate Clear cell was removed so the calculated Debye length
    remains defined. *)


(* SUBSECTION: Values - Run First *)

(*Values to apply *)
(* Constants *)
hbar = 1.0545718*(10)^(-34); (* reduced Planks constant [J*s]  *)
k = 8.99*(10)^(9);               (* (1)/(4*Subscript[Piepsilon, 0])  Coulomb constant
     [(N*(m)^(2))/((C)^(2))] *)
e = 1.602*10^-19;    (* Charge of an electron    [C]   *)
kb = 1.38*(10)^(-23); (*J/K*)
```

```
epsilon0 = 8.854*(10)^(-12); (* F/m *)


(* Experimental D2 Values *)
V0eV = 35*(10)^(6);              (* Nuclear Potential of Deuterium [eV] *)
V0J = (V0eV)/(6.242*(10)^(18));    (* Nuclear Potential of Deuterium [J] *)
m = 3.3435*(10)^(-27);        (* Experimental Mass of deuterium nucleus [kg] *)
R0exp = 2.12562*(10)^(-15); (* Experimental charge radius of deuteron nucleus [m] *)

(* Analytical D2 Values *)
a = 0.5*(10)^(-15);    (* other nuclear dist. values *)
r0 = 1.25*(10)^(-15);(* Nuclear dist. constant [m] *)
R = r0*(2)^(1/3);        (* Theoretical nuclear charge radius of detuerium [m]*)

(* ========================================================================= *)
(* SECTION: Titanium Debye Model *)
(* ========================================================================= *)

(* SUBSECTION: Woods Saxon *)

(* STEP: Titanium Calculations *)

rhoTi = 4.502 (* g/(cm)^(3) *);
MTi = 47.867 (* g/mol *);

NaTi = (rhoTi)/(MTi)*(6.022*(10)^(23))*((100))^(3); (* Atoms/m^3 *)

(* To find the electron density, we will take NaTi and multiply by 2 since the 2
     valence electrons of titanium are the only that are assuemd to contribute (others
     remain around Ti atom) *)
```

```
48 ne = 2*NaTi
49
50 (* STEP: Calculating the Debye length lambdaD *)
51
52 (* We will assume standard temp *)
53 T = 293 (*K*);
54
55 (* Debye length*)
56 lambdaD = (((epsilon0*kb*T)/(ne*(e)^(2))))^(1/2) (* m *)
57
58 (* STEP: WKB Prep *)
59
60 (* Variable *)
61 EngeV = .004*(10)^(6);                        (* Energy of incoming particle [eV] *)
62 Eng =EngeV/(6.242*(10)^(18));   (* Energy of incoming particle [J] *)
63
64 (*Define Potential Functions*)
65 CoulombPoteV[r_]:=((k*(e)^(2)))/(r)*(1)/((1.602*(10)^(-19))) (* Coulomb Potential in
       eV *)
66 CoulombPotJ[r_]:=((k*(e)^(2)))/(r) (* Coulomb Potential in J *)
67
68 (* Debye Function *)
69 Debye[r_]:= Exp[(-r)/(lambdaD)];
70
71 (* Screened Potentials *)
72 ScrnCoulombPoteV[r_] := Debye[r]*CoulombPoteV[r]
73 ScrnCoulombPotJ[r_] := Debye[r]*CoulombPotJ[r]
74
75 (* Strong force modeled by the Woods Saxon: a mean field potential for nucleons inside
        atomic nucleus *)
```

```
76 WoodsPoteV[r_]:= -(V0eV)/(1+Exp[(r-R0exp)/(a)]); (* Nuclear Potential in eV *)
77 WoodsPotJ[r_]:= -(V0J)/(1+Exp[(r-R0exp)/(a)]); (* Nuclear Potential in J *)
78
79 (*Total potential energy*)
80 TotalWDSPotentialeV[r_]:=CoulombPoteV[r]+WoodsPoteV[r]
81 TotalWDSPotentialJ[r_]:=CoulombPotJ[r]+WoodsPotJ[r]
82 ScrnWDSPotentialJ[r_]:=ScrnCoulombPotJ[r]+WoodsPotJ[r] (* Total Potentials Joules *)
83 ScrnWDSPotentialeV[r_]:=ScrnCoulombPoteV[r]+WoodsPoteV[r] (* Total Potentials eV *)
84
85 (* Tunneling Fuctions *)
86 tunnelProbability[beta_]:=Exp[-2*beta]*((1-(1)/(4)*Exp[-2*beta]))^(-2);
87 tunnelSimp[beta_]:=Exp[-2*beta];
88
89 (* STEP: Tunneling Prep *)
90
91 (* Turning Points *)
92 ScrnxRwoods=r/.FindRoot[ScrnWDSPotentialeV[r]== EngeV,{r,8*(10)^(-12)}]//Quiet
93 ScrnxLwoods=r/.FindRoot[ScrnWDSPotentialeV[r]==EngeV,{r,4.7*(10)^(-15)}]//Quiet
94
95 (*Plot*)
96 ScrnWDSTurnPointPlot = ListPlot[{{ScrnxRwoods, EngeV},{ScrnxLwoods,EngeV}}, Filling->
       Axis, PlotMarkers->{Automatic, Medium}];
97 ScrnWDSPotentialPlot = Plot[{CoulombPoteV[r],WoodsPoteV[r],ScrnWDSPotentialeV[r],
       EngeV},{r,0,5*(10)^(-13)},PlotLegends->{"Coulomb Potential","Nuclear Potential
       Well","Total Potential", "KE"},AxesLabel->{"r (m)","V(r) (eV)"},PlotRange
       ->{-1*(10)^(5),.5*(10)^(6)},PlotStyle->{Dashed,Dotted,Thick}];
98 Show[ScrnWDSPotentialPlot, ScrnWDSTurnPointPlot]
99
100 (* STEP: Tunneling Calc *)
101
```

```
(* Calculate momentum *)
Scrnpwoods= Sqrt[2*m*(Eng-ScrnWDSPotentialJ[r])];
(* Calculate beta integral *)
Scrnbetawoods = (1)/(hbar)*NIntegrate[Abs[Scrnpwoods], {r,ScrnxLwoods,ScrnxRwoods}];

(* Plug into tunneling approximation Formula *)
Twoods = tunnelProbability[Scrnbetawoods]
(* Beta is not >> 1 and therefore we must keep the above whole probability *)
tunnelSimp[Scrnbetawoods];

(* STEP: Plot to visualize *)

Plot[{CoulombPoteV[r],ScrnCoulombPoteV[r]},{r,0,1*(10)^(-12)}, AxesLabel->{"r(m)", "E(
    eV)"}, PlotLegends->{"Normal Coulomb", "Screened Potential"}]

(* STEP: sigma(E) *)

(* Cross Section calculation, assume Subscript[P, nuc] is one for maximum value *)
sigmaWoods = (Pi*(hbar)^(2))/(2*m)*(1)*(Twoods)/(Eng) *(10)^(28)(* NOTE THE CONVERSION
     TO BARNS *)

(* STEP: Comparison to No Screen *)

(sigmaWoods)/(4.830398989596194`*^-7)
```

Listing L.4 Debye-screened potential, WKB tunneling, and cross-section calculations.

L.5 Thomas-Fermi Screening Model

The Thomas-Fermi calculation uses the titanium electron density to estimate the electronic screening length and applies the resulting screened potential to the WKB calculation.

```
(* ::Package:: *)
(* Thomas-Fermi Screening Model for Titanium *)
(* Extracted and reformatted from LNAR_Models.nb (Mathematica 14.2). *)
(* Only input cells are included; notebook outputs and messages are omitted. *)
(* Repeated Clear cells were consolidated into the single line below. *)

ClearAll["Global`*"];

(* NOTE: The intermediate Clear cell was removed so the calculated screening
    parameters remain defined. *)


(* SUBSECTION: Values - Run First *)

(*Values to apply *)
(* Constants *)
hbar = 1.0545718*(10)^(-34); (* reduced Planks constant [J*s]  *)
k = 8.99*(10)^(9);               (* (1)/(4*Subscript[Piepsilon, 0])  Coulomb constant
     [(N*(m)^(2))/((C)^(2))] *)
e = 1.602*10^-19;    (* Charge of an electron    [C]   *)
kb = 1.38*(10)^(-23); (*J/K*)
epsilon0 = 8.854*(10)^(-12);(* F/m *)
me = 9.109*(10)^(-31); (* Kg *)


(* Experimental D2 Values *)
```

```
25 V0eV = 35*(10)^(6);              (* Nuclear Potential of Deuterium [eV] *)
26 V0J = (V0eV)/(6.242*(10)^(18));    (* Nuclear Potential of Deuterium [J] *)
27 m = 3.3435*(10)^(-27);         (* Experimental Mass of deuterium nucleus [kg] *)
28 R0exp = 2.12562*(10)^(-15); (* Experimental charge radius of deuteron nucleus [m] *)
29
30 (* Analytical D2 Values *)
31 a = 0.5*(10)^(-15);    (* other nuclear dist. values *)
32 r0 = 1.25*(10)^(-15);(* Nuclear dist. constant [m] *)
33 R = r0*(2)^(1/3);         (* Theoretical nuclear charge radius of detuerium [m]*)
34
35 (* ========================================================================= *)
36 (* SECTION: Titanium *)
37 (* ========================================================================= *)
38
39 (* SUBSECTION: Woods Saxon Attempt *)
40
41 (* STEP: Titanium Calculations *)
42
43 rhoTi = 4.502 (* g/(cm)^(3) *);
44 MTi = 47.867 (* g/mol *);
45
46 NaTi = (rhoTi)/(MTi)*(6.022*(10)^(23))*((100))^(3); (* Atoms/m^3 *)
47
48 (* To find the electron density, we will take NaTi and multiply by 2 since the 2
       valence electrons of titanium are the only that are assuemd to contribute (others
       remain around Ti atom) *)
49 ne = 2*NaTi
50
51 (* STEP: Calculating the Thomas-Fermi correction *)
52
```

```
(* We will assume standard temp *)
T = 293 (*K*);

(* kf *)
kf = ((3*(Pi)^(2)*ne))^(1/3);

(* Debye length*)
Kft = ((((e)^(2)*me*kf )/((Pi)^(2)*epsilon0*(hbar)^(2))))^(1/2)(* m *)

(* STEP: WKB Prep *)

(* Variable *)
EngeV = .004*(10)^(6);                          (* Energy of incoming particle [eV] *)
Eng =EngeV/(6.242*(10)^(18));   (* Energy of incoming particle [J] *)

(*Define Potential Functions*)
CoulombPoteV[r_]:=((k*(e)^(2)))/(r)*(1)/((1.602*(10)^(-19))) (* Coulomb Potential in
    eV *)
CoulombPotJ[r_]:=((k*(e)^(2)))/(r) (* Coulomb Potential in J *)

(* Debye Function *)
ThomFerm[r_]:= Exp[-Kft*r];

(* Screened Potentials *)
ScrnCoulombPoteV[r_] := ThomFerm[r]*CoulombPoteV[r]
ScrnCoulombPotJ[r_] := ThomFerm[r]*CoulombPotJ[r]

(* Strong force modeled by the Woods Saxon: a mean field potential for nucleons inside
     atomic nucleus *)
WoodsPoteV[r_]:= -(V0eV)/(1+Exp[(r-R0exp)/(a)]); (* Nuclear Potential in eV *)
```

```
81  WoodsPotJ[r_]:= -(V0J)/(1+Exp[(r-R0exp)/(a)]); (* Nuclear Potential in J *)
82
83  (*Total potential energy*)
84  TotalWDSPotentialeV[r_]:=CoulombPoteV[r]+WoodsPoteV[r]
85  TotalWDSPotentialJ[r_]:=CoulombPotJ[r]+WoodsPotJ[r]
86  ScrnWDSPotentialJ[r_]:=ScrnCoulombPotJ[r]+WoodsPotJ[r] (* Total Potentials Joules *)
87  ScrnWDSPotentialeV[r_]:=ScrnCoulombPoteV[r]+WoodsPoteV[r] (* Total Potentials eV *)
88
89  (* Tunneling Fuctions *)
90  tunnelProbability[beta_]:=Exp[-2*beta]*((1-(1)/(4)*Exp[-2*beta]))^(-2);
91  tunnelSimp[beta_]:=Exp[-2*beta];
92
93  (* STEP: Tunneling Prep *)
94
95  (* Turning Points *)
96  ScrnxRwoods=r/.FindRoot[ScrnWDSPotentialeV[r]== EngeV,{r,8*(10)^(-15)}]//Quiet
97  ScrnxLwoods=r/.FindRoot[ScrnWDSPotentialeV[r]==EngeV,{r,4.7*(10)^(-15)}]//Quiet
98
99  (*Plot*)
100 ScrnWDSTurnPointPlot = ListPlot[{{ScrnxRwoods, EngeV},{ScrnxLwoods,EngeV}}, Filling->
        Axis, PlotMarkers->{Automatic, Medium}];
101 ScrnWDSPotentialPlot = Plot[{CoulombPoteV[r],WoodsPoteV[r],ScrnWDSPotentialeV[r],
        EngeV},{r,0,30*(10)^(-15)},PlotLegends->{"Coulomb Potential","Nuclear Potential
        Well","Total Potential", "KE"},AxesLabel->{"r (m)","V(r) (eV)"},PlotRange
        ->{-1*(10)^(6),1*(10)^(6)},PlotStyle->{Dashed,Dotted,Thick}];
102 Show[ScrnWDSPotentialPlot, ScrnWDSTurnPointPlot]
103
104 (* STEP: Tunneling Calc *)
105
106 (* Calculate momentum *)
```

```
Scrnpwoods= Sqrt[2*m*(Eng-ScrnWDSPotentialJ[r])];
(* Calculate beta integral *)
Scrnbetawoods = (1)/(hbar)*NIntegrate[Abs[Scrnpwoods], {r,ScrnxLwoods,ScrnxRwoods}];

(* Plug into tunneling approximation Formula *)
Twoods = tunnelProbability[Scrnbetawoods]
(* Beta is not >> 1 and therefore we must keep the above whole probability *)
tunnelSimp[Scrnbetawoods];

(* STEP: Plot to visualize *)

Plot[{CoulombPoteV[r],ScrnCoulombPoteV[r]},{r,0,1*(10)^(-12)}, AxesLabel->{"r(m)", "E(
    eV)"}, PlotLegends->{"Normal Coulomb", "Screened Potential"}]

(* STEP: sigma(E) *)

(* Cross Section calculation, assume Subscript[P, nuc] is one for maximum value *)
sigmaWoods = (Pi*(hbar)^(2))/(2*m)*(1)*(Twoods)/(Eng) *(10)^(28)(* NOTE THE CONVERSION
     TO BARNS *)

(* STEP: Comparison to No Screen *)

(sigmaWoods)/(4.830398989596194`*^-7)
```

Listing L.5 Thomas-Fermi screening and WKB tunneling calculations.

L.6 Density-Functional Screening Approximation

The final listing applies the electron-density values adopted for the density-functional approximation to calculate the screening correction and the resulting WKB tunneling probability. The source and physical meaning of the input electron densities should be identified in the main thesis text.

```
(* ::Package:: *)
(* Density-Functional Screening Approximation *)
(* Extracted and reformatted from LNAR_Models.nb (Mathematica 14.2). *)
(* Only input cells are included; notebook outputs and messages are omitted. *)
(* Repeated Clear cells were consolidated into the single line below. *)

ClearAll["Global`*"];

(* NOTE: The electron-density inputs must be cited and justified in the thesis text.
    *)


(* SUBSECTION: Values - Run First *)

(*Values to apply *)
(* Constants *)
hbar = 1.0545718*(10)^(-34); (* reduced Planks constant [J*s]  *)
k = 8.99*(10)^(9);               (* (1)/(4*Subscript[Piepsilon, 0])  Coulomb constant
     [(N*(m)^(2))/((C)^(2))] *)
e = 1.602*10^-19;    (* Charge of an electron    [C]   *)
kb = 1.38*(10)^(-23); (*J/K*)
epsilon0 = 8.854*(10)^(-12);(* F/m *)
me = 9.109*(10)^(-31); (* Kg *)


```

```
(* Experimental D2 Values *)
V0eV = 35*(10)^(6);              (* Nuclear Potential of Deuterium [eV] *)
V0J = (V0eV)/(6.242*(10)^(18));    (* Nuclear Potential of Deuterium [J] *)
m = 3.3435*(10)^(-27);         (* Experimental Mass of deuterium nucleus [kg] *)
R0exp = 2.12562*(10)^(-15); (* Experimental charge radius of deuteron nucleus [m] *)

(* Analytical D2 Values *)
a = 0.5*(10)^(-15);    (* other nuclear dist. values *)
r0 = 1.25*(10)^(-15);(* Nuclear dist. constant [m] *)
R = r0*(2)^(1/3);         (* Theoretical nuclear charge radius of detuerium [m]*)

(* ========================================================================== *)
(* SECTION: Titanium *)
(* ========================================================================== *)

(* SUBSECTION: Woods Saxon Attempt *)

(* STEP: DFT Calculations of Electron Density *)

nsmall = 1.91*(10)^(30);

nlarge = 4.06*(10)^(29);

(* STEP: Titanium value calcs *)

(* Kf vals *)
kfsm = ((3*(Pi)^(2)*nsmall))^(1/3);
kflg = ((3*(Pi)^(2)*nlarge))^(1/3);

(* Engs *)
```

```
54 Efsm = ((hbar)^(2)*(kfsm)^(2))/(2*me);
55 Eflg = ((hbar)^(2)*(kflg)^(2))/(2*me);
56
57 (* Debye length*)
58 Kftsm = (((6*Pi*(e)^(2)*nsmall )/(epsilon0*Efsm)))^(1/2)(* m *)
59 Kftlg = (((6*Pi*(e)^(2)*nlarge)/(epsilon0*Eflg)))^(1/2)(* m *)
60
61 (* STEP: WKB Prep *)
62
63 (* Variable *)
64 EngeV = .004*(10)^(6);                        (* Energy of incoming particle [eV] *)
65 Eng =EngeV/(6.242*(10)^(18));   (* Energy of incoming particle [J] *)
66
67 (*Define Potential Functions*)
68 CoulombPoteV[r_]:=((k*(e)^(2)))/(r)*(1)/((1.602*(10)^(-19))) (* Coulomb Potential in
       eV *)
69 CoulombPotJ[r_]:=((k*(e)^(2)))/(r) (* Coulomb Potential in J *)
70
71 (* Debye Function *)
72 DFT[r_]:= Exp[-Kftsm*r];
73
74 (* Screened Potentials *)
75 ScrnCoulombPoteV[r_] := DFT[r]*CoulombPoteV[r]
76 ScrnCoulombPotJ[r_] := DFT[r]*CoulombPotJ[r]
77
78 (* Strong force modeled by the Woods Saxon: a mean field potential for nucleons inside
        atomic nucleus *)
79 WoodsPoteV[r_]:= -(V0eV)/(1+Exp[(r-R0exp)/(a)]); (* Nuclear Potential in eV *)
80 WoodsPotJ[r_]:= -(V0J)/(1+Exp[(r-R0exp)/(a)]); (* Nuclear Potential in J *)
81
```

```
(*Total potential energy*)
TotalWDSPotentialeV[r_]:=CoulombPoteV[r]+WoodsPoteV[r]
TotalWDSPotentialJ[r_]:=CoulombPotJ[r]+WoodsPotJ[r]
ScrnWDSPotentialJ[r_]:=ScrnCoulombPotJ[r]+WoodsPotJ[r] (* Total Potentials Joules *)
ScrnWDSPotentialeV[r_]:=ScrnCoulombPoteV[r]+WoodsPoteV[r] (* Total Potentials eV *)

(* Tunneling Fuctions *)
tunnelProbability[beta_]:=Exp[-2*beta]*((1-(1)/(4)*Exp[-2*beta]))^(-2);
tunnelSimp[beta_]:=Exp[-2*beta];

(* STEP: Tunneling Prep *)

(* Turning Points *)
ScrnxRwoods=r/.FindRoot[ScrnWDSPotentialeV[r]== EngeV,{r,8*(10)^(-15)}]//Quiet
ScrnxLwoods=r/.FindRoot[ScrnWDSPotentialeV[r]==EngeV,{r,4.7*(10)^(-15)}]//Quiet

(*Plot*)
ScrnWDSTurnPointPlot = ListPlot[{{ScrnxRwoods, EngeV},{ScrnxLwoods,EngeV}}, Filling->
    Axis, PlotMarkers->{Automatic, Medium}];
ScrnWDSPotentialPlot = Plot[{CoulombPoteV[r],WoodsPoteV[r],ScrnWDSPotentialeV[r],
    EngeV},{r,0,5*(10)^(-13)},PlotLegends->{"Coulomb Potential","Nuclear Potential
    Well","Total Potential", "KE"},AxesLabel->{"r (m)","V(r) (eV)"},PlotRange
    ->{-.01*(10)^(6),.4*(10)^(6)},PlotStyle->{Dashed,Dotted,Thick}];
Show[ScrnWDSPotentialPlot, ScrnWDSTurnPointPlot]

(* STEP: Tunneling Calc *)

(* Calculate momentum *)
Scrnpwoods= Sqrt[2*m*(Eng-ScrnWDSPotentialJ[r])];
(* Calculate beta integral *)
```

```
Scrnbetawoods = (1)/(hbar)*NIntegrate[Abs[Scrnpwoods], {r,ScrnxLwoods,ScrnxRwoods}];

(* Plug into tunneling approximation Formula *)
Twoods = tunnelProbability[Scrnbetawoods]
(* Beta is not >> 1 and therefore we must keep the above whole probability *)
tunnelSimp[Scrnbetawoods];

(* STEP: Plot to visualize *)

Plot[{CoulombPoteV[r],ScrnCoulombPoteV[r]},{r,0,1*(10)^(-12)}, AxesLabel->{"r(m)", "E(
    eV)"}, PlotLegends->{"Normal Coulomb", "Screened Potential"}]

(* STEP: sigma(E) *)

(* Cross Section calculation, assume Subscript[P, nuc] is one for maximum value *)
sigmaWoods = (Pi*(hbar)^(2))/(2*m)*(1)*(Twoods)/(Eng) *(10)^(28)(* NOTE THE CONVERSION
     TO BARNS *)

(* STEP: Comparison to No Screen *)

(sigmaWoods)/(4.830398989596194`*^-7)
```

Listing L.6 Density-functional electron-density approximation and WKB tunneling calculations.

Bibliography

[1] C. E. Rolfs and W. S. Rodney, *Cauldrons in the cosmos: Nuclear astrophysics* (University of Chicago press, 1988).

[2] D. D. Clayton, *Principles of Stellar Evolution and Nucleosynthesis* (University of Chicago Press, 1983).

[3] C. J. Hansen, S. D. Kawaler, and V. Trimble, "An Overview of Stellar Evolution," Stellar Interiors: Physical Principles, Structure, and Evolution pp. 43–144 (2004).

[4] R. Kippenhahn, A. Weigert, and A. Weiss, *Stellar structure and evolution* (Springer, 1990), Vol. 192.

[5] A. S. Eddington, *The internal constitution of the stars* (Cambridge University Press, 1926).

[6] H. A. Bethe and C. L. Critchfield, "The formation of deuterons by proton combination," Physical Review **54,** 248 (1938).

[7] H. A. Bethe, "Energy production in stars," Physical Review **55,** 434 (1939).

[8] G. Gamow, *The quantum theory of the atomic nucleus* (US Atomic Energy Commission, Division of Technical Information Extension, 1963).

[9] B. W. Carroll and D. A. Ostlie, *An introduction to modern astrophysics* (Cambridge University Press, 2017).

[10] A. B. Balantekin and N. Takigawa, "Quantum tunneling in nuclear fusion," Reviews of Modern Physics **70,** 77 (1998).

[11] G. H. Miller, E. I. Moses, and C. R. Wuest, "The National Ignition Facility: enabling fusion ignition for the 21st century," Nuclear fusion **44,** S228–S238 (2004).

[12] B. Coppi and J. Rem, "The Tokamak approach in fusion research," Scientific American **227,** 65–75 (1972).

[13] S. E. Koonin and M. Nauenberg, "Calculated fusion rates in isotopic hydrogen molecules," Nature **339,** 690–691 (1989).

[14] F. Frank, "Hypothetical Alternative Energy Sources for the 'Second Meson' Events," Nature **160,** 525–527 (1947).

[15] Y. B. Zeldovich, Doklady Akademii Nauk SSSR **95,** 493 (1954).

[16] A. D. Sakharov, unpublished (unpublished).

[17] S. E. Jones, A. N. Anderson, A. J. Caffrey, J. B. Walter, K. D. Watts, J. N. Bradbury, P. A. M. Gram, M. Leon, H. R. Maltrud, and M. A. Paciotti, "Experimental Investigation of Muon-Catalyzed $d-t$ Fusion," Physical Review Letters **51,** 1757–1760 (1983).

[18] B. L. Altshuler, "Andrei Sakharov's research work and modern physics," Physics-Uspekhi **64,** 427 (2021).

[19] S. E. Jones *et al.*, "Observation of unexpected density effects in muon-catalyzed d-t fusion," Physical Review Letters **56,** 588–591 (1986).

[20] S. E. Jones, A. N. Anderson, A. J. Caffrey, J. B. Walter, K. D. Watts, J. N. Bradbury, P. A. M. Gram, M. Leon, H. R. Maltrud, and M. A. Paciotti, "Experimental Investigation of Muon-Catalyzed $d-t$ Fusion," Physical Review Letters **51,** 1757–1760 (1983).

[21] L. W. Alvarez *et al.*, "Catalysis of Nuclear Reactions by μ Mesons," Physical Review **105,** 1127–1128 (1957).

[22] S. E. Jones *et al.*, "Observation of unexpected density effects in muon-catalyzed d-t fusion," Phys. Rev. Lett. **56,** 588–591 (1986).

[23] S. E. Jones, E. P. Palmer, J. B. Czirr, D. L. Decker, G. L. Jensen, J. M. Thorne, S. F. Taylor, and J. Rafelski, "Observation of cold nuclear fusion in condensed matter," Nature **338,** 737–740 (1989).

[24] C. Rolfs, "Electron screening in metallic environments: a plasma of the poor man," AIP Conference Proceedings **847,** 245–248 (2006).

[25] F. W. Keeney, S. E. Jones, A. C. Johnson, P. L. Hagelstein, G. Hubler, D. B. Buehler, F. E. Cecil, M. R. Scott, and J. E. Ellsworth, "Charged-Particle Emissions From Deuterided Metals," Condensed Matter Nuclear Science (2005).

[26] J. Cruz, H. Luís, M. Fonseca, and A. P. Jesus, "Electron screening effects in nuclear reactions: still an unsolved problem," Journal of Physics: Conference Series **337,** 012062 (2012).

[27] M. Lipoglavsek and A. Cvetinović, "Electron Screening in Nuclear Astrophysics," EPJ Web of Conferences **227,** 01012 (2020).

[28] M. Lipoglavšek, S. Markelj, M. Mihovilovič, T. Petrovič, S. Štajner, M. Vencelj, and J. Vesić, "Observation of electron emission in the nuclear reaction between protons and deuterons," Physics Letters B **773,** 553–556 (2017).

[29] E. Grohs, A. R. Howe, and F. C. Adams, "Universes without the weak force: Astrophysical processes with stable neutrons," Physical Review D **97,** 043003 (2018).

[30] K. S. Krane and D. Halliday, *Introductory nuclear physics* (Wiley, New York, 1987).

[31] V. Pines *et al.*, "Nuclear fusion reactions in deuterated metals," Physical Review C **101,** 044609 (2020).

[32] K.-Y. Chen *et al.*, "Electrochemical loading enhances deuterium fusion rates in a metal target," Nature **644,** 640–645 (2025).

[33] H. E. Puthoff, "Ground state of hydrogen as a zero-point-fluctuation-determined state," Phys. Rev. D **35,** 3266–3269 (1987).

[34] L. W. Alvarez, "Recent Developments in Particle Physics,", Nobel Lecture, 1968, nobel Lecture, December 11, 1968, The Nobel Foundation.

[35] R. Lundell, Senior thesis, Brigham Young University, Provo, Utah, 2026.

[36] K. Terry, Senior thesis, Brigham Young University, Provo, Utah, 2016.

[37] AMETEK ORTEC, *Introduction to Charged-Particle Detectors*, AMETEK ORTEC, 2025.

[38] R. D. Woods and D. S. Saxon, "Diffuse Surface Optical Model for Nucleon-Nuclei Scattering," Physical Review **95,** 577–578 (1954).

[39] V. Horvath, J. Brady, and T. Hamm, PHSCS 245 Final Project, Brigham Young University (unpublished).

Index