THE UNIVERSITY OF CHICAGO

DARK MATTER LIMITS FROM A 15 kg WINDOWLESS BUBBLE CHAMBER

A DISSERTATION SUBMITTED TO

THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES

IN CANDIDACY FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS

BY

MATTHEW MARK SZYDAGIS

CHICAGO, ILLINOIS

MARCH 2011

To Kel

and

to Mom and Dad

with love

ABSTRACT

The COUPP collaboration has successfully used bubble chambers, a technology previously applied only to high-energy physics experiments, as direct dark matter detectors. It has produced the world's most stringent spin-dependent WIMP limits, and increasingly competitive spin-independent limits. These limits were achieved by capitalizing on an intrinsic rejection of the gamma background that all other direct detection experiments must address through high-density shielding and empirically-determined data cuts. The history of COUPP, including its earliest prototypes and latest results, is briefly discussed in this thesis. The feasibility of a new, windowless bubble chamber concept simpler and more inexpensive in design is discussed here as well. The dark matter limits achieved with a 15 kg windowless chamber, larger than any previous COUPP chamber (2 kg, 4 kg), are presented. Evidence of the greater radiopurity of synthetic quartz compared to natural is presented using the data from this 15 kg device, the first chamber to be made from synthetic quartz. The effective reconstruction of the three-dimensional positions of bubbles in a highly distorted optical field, with ninety-degree bottom lighting similar to cloud chamber lighting, is demonstrated. Another innovation described in this thesis is the use of the sound produced by bubbles recorded by an array of piezoelectric sensors as the primary means of bubble detection. In other COUPP chambers, cameras have been used as the primary trigger. Previous work on bubble acoustic signature differentiation using piezos is built upon in order to further demonstrate the ability to discriminate between alpha- and neutron-induced events.

ACKNOWLEDGEMENTS

Above all, this thesis could not have been possible without my wife Kel, who not only ensured that it was comprehensible and not riddled with mistakes, but also put up with my moodiness as I edited from draft to draft. My parents also encouraged me to persist in the face of what felt at the time to be unsurmountable difficulties.

I am so grateful to my thesis advisor, Juan Collar, who never expected anything less than the best, and therefore encouraged me to make my best even better. His ability to juggle multiple tasks, even different experiments and collaborations, was a never-ending source of inspiration as I myself strove to juggle the many independent aspects of the experiment described in this thesis.

Fellow graduate student and friend Phil Barbeau took time out of his own experiment to help me with physics questions when I first started working with Juan, and his advice as one who had "been there, done that" was incredibly useful as I framed this thesis. Postdoc Eric Dahl has also patiently answered my many questions. Eric was a huge help in formulating the acoustic analysis and muon veto coincidence, as well as explaining the calculation of dark matter limits.

I could not have gotten everything together into a functioning dark matter experiment if it were not for a long chain of Collar lab students, Chicago students as well as summer REUs, spanning the past 5 years. I thank Ed Chen, Jacob Claussen (whose physical strength was of particular use), Adrian Culver, Maria Garzon, Luke Goetzke, Smriti Mishra, Songkiat Nutalaya, Chris Passolano, Aaron Plasek, and Hannes Schimmelpfennig. Postdoc Brian Odom helped set the stage for the windowless bubble chamber, and if it were not for Nathan Riley's master thesis, the chamber would not have had a muon veto.

The deployment of the 15 kg chamber underground in the Chicago sewers would not have been possible without the generous efforts of Tom Economou of the University of

Chicago and Mike Cavarretta of the Metropolitan Water Reclamation District. Without Pat Coleman, Bob, Jim, and all of the crane operators and ironworkers, my underground detector never would have made it underground. The daily TARP staff, including Dom, Matt, and Kevin, were friendly, supportive, and curious about the physics. Without Eric Dahl, Drew Fustin, Nicole Fields, Alan Robinson, and Steven Bogacz, the fast-paced one week decommissioning would never have been completed.

My experiment was only a part of a large collaboration, and everyone's wisdom and experience were of great help to me and my education. In particular: Mike Crisler and Ilan Levine, in addition to Juan Collar, wrote me letters of recommendation for postdoc positions as well as reviewing my personal statements. Father and son team Ed and Josh Behnke, along with Ilan Levine, bent over backwards to get me piezos and thermometers. I have also particularly enjoyed learning from Steve Brice, Peter Cooper, Jeter Hall, Erik Ramberg, and Andrew Sonnenschein. Mark Ruschman, Jerry Zimmerman, and all of the COUPP engineers were indispensable. Mark, in fact, assembled the 15 kg chamber with me in the Fermilab cleanroom.

There is a long list of University of Chicago staff who deserve my thanks. Larry Fiscelli and Jim Passolano helped load, and then months later unload, the truck full of the 15 kg chamber's accoutrements. Dave Plitt and all the machine shop staff quickly machined all the parts I needed. IT expert Valeri Galtsev helped maintain the computer and my nerves. Cristian Millang and Mary Wawro ensured that parts were ordered and shipped, and Les, Mark, Donald, O.J., Jarvis, Anthony, and all of the PSD receiving dock staff assured their prompt delivery to me. Prentice Bradford and stockroom staff also helped me get the supplies I needed. Qiti Guo helped me do surface analysis of my materials. And finally, thanks to Nobuko McNeill, I never missed a deadline and remembered to actually form a thesis committee.

My committee chair and graduate advisor, Juan Collar, was joined on the committee by my undergraduate thesis advisor, Steve Meyer, who gave me the foundation I needed to be a successful graduate student. The other two members of my committee at one point taught me in the classroom. I was truly challenged by Emil Martinec's graduate quantum mechanics course. Woowon Kang's undergraduate electromagnetism course was entertaining: I'll never forget the pigeon in the magnetic field.

To all who supported and advised me, thank you.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

PREFACE

Although I received my bachelor's degree from the University of Chicago, I was far from ready to leave. I wanted to continue in the world-class physics department I enjoyed so much, as well as to remain close to family and friends and good gyros. While I wasn't interested in diversifying geographically, I did diversify my experience by changing fields, from the CMB with Steve Meyer to dark matter with Juan Collar. I knew I wanted to stay in astrophysics, and I chose to ask Juan if I could work with him. My wife would say that I must be drawn to doing things that have no practical application, but just extend the scope of human knowledge. Over the years, I have had to deflect questions from friends and family many times about when dark matter was going to cure cancer or spark world peace or do something else remotely useful.

After many side projects and tangents, the focus of my thesis finally became clear, especially after COUPP postdoc Brian Odom, who had been steering the windowless bubble chamber project, left the KICP for a faculty position at Northwestern University. At that point it became my responsibility to complete and deploy the 15 kg windowless bubble chamber, with the help of many undergraduates, and eventually other graduate students and new postdoc Eric Dahl. I ended up coordinating almost all aspects of both hardware and software for the chamber. Because of all this engineering experience, it was easier to answer technical questions at conferences where I had the honor of representing the collaboration. Through engineering hurdles which delayed both the surface and underground runs of the 15 kg chamber, my understanding of bubble chamber physics was actually increased.

In addition to producing new dark matter limits, through the course of this experiment, Murphy's Law was inadvertently and unequivocally confirmed. One mechanical failure after another plagued the bubble chamber. However, I learned so much from being forced to constantly resuscitate the experiment, often with the help of Eric Dahl and Drew Fustin.

Equipment was replaced or repaired, or if impossible, workarounds with remaining equipment were found. There was often a silver lining. For example, piston failure led to the use of a better pressurization system. Pervasive thermometer failure led to vapor pressure measurements that yielded a surprising result, without which data analysis would have been confusing. Even the hard disk crashed, but the size of its replacement made it possible to fit more data. The heating pads gave out, but the replacement system, relying solely on a circulating pump, decreased the temperature gradient over the chamber. When the LEDs were suspected of overheating the chamber, I was forced to develop a new sound-based trigger for bubbles, which turned out to be a success despite the chamber being in a noise-filled environment and despite the distance of the piezoelectric acoustic sensors from the center of the chamber. It now offers a possible replacement for visual triggering for future chambers. Ironically, the only equipment which did not fail was that which was multiply redundant, like the LEDs (six clusters when two were needed) and the piezos (four total).

I often uncovered ways in which future COUPP chambers could be improved. Because the small number of piezos, placed only at the top and bottom of the chamber, rendered it impossible to hear bubbles effectively at all pressures in the 15 kg chamber, the 60 kg chamber design was revised to include more piezos, installed all along the length of the vessel. LED failures due to polymerization of the hydraulic fluid (mineral oil) during engineering tests for the 15 kg chamber led to more caution and care during long-term performance tests of LEDs for the 60 kg chamber. Some 3-D optical bubble position reconstruction methods used for the 15 kg anaylsis were re-applied for the 60 kg chamber. Overall, many engineering lessons were learned from the 15 kg run. That is not to say that physics lessons were not learned. Dark matter limits, though not the best in the world because of a neutron background, were still achieved despite many mechanical challenges and obstacles. A background neutron and alpha spectrum between 10 and 100 keV was measured that was

statistically a vast improvement over the 2 kg chamber data, as well as orthogonal to that of the 4 kg chamber, which only collected data at one threshold. The background spectrum may serve as a blueprint for decision-making in the future, in terms of chamber data-taking parameters.

I spent a great deal of time on what was ultimately a small part of the bigger picture: position reconstruction using cameras. I learned a lot about geometric optics in the process, and made connections across different disciplines of physics. Thanks to an old practice candidacy exam problem, I became aware of the valuable and powerful concept of Lagrangian optics. My work on 3-D reconstruction with cameras can for example be applied to PMTs or other devices. In general, much of what I have learned may be applicable to other physics problems I encounter in the future. In fact, I have already learned to apply a t_0-finding algorithm for acoustic signal analysis to finding the start times of PMT pulses for the LUX experiment.

A new physics goal appeared partway through the project when PICASSO announced the realization of alpha-neutron acoustic discrimination. First the 4 kg chamber and then the 15 kg chamber successfully confirmed this characteristic of bubble acoustics in superheated liquids. This was a game-changer for the technology, paving the way for building detectors exclusively sensitive to WIMP recoils. (In general, new ways of identifying particles can be useful for multiple disciplines.) While reading background materials as a part of researching this thesis, I discovered that recording bubbles is in fact not revolutionary. I was excited to learn from Donald Glaser's seminal papers that he had used a phonograph to trigger one of his bubble chambers. Studying the sound which bubbles make is by no means a new idea, but one stemming from the long history of the technology.

Juan, who has had to suffer through my constant summaries of all of the technical problems, has been fond of saying that you should finish what you start. Instead of shifting

to the less problem-riddled 4 kg chamber or joining in on the prep work for the 60 kg run, I persevered on the constant learning experience that was the 15 kg chamber. By the Spring of 2010, when an unstoppable oil leak at last ended data-taking for the 15 kg chamber, several hundred kg-days of engineering data had been collected, as well as several hundred more kg-days of physics data for dark matter limits and acoustic discrimination confirmation, more than had ever been collected before on COUPP.

I hope that you will find this thesis intriguing, both as a part of the COUPP dark matter search and as a new application of a decades-old technology. I pray that I will never have to see mineral oil again, having been drenched in it an umpteen number of times. And I would like to conclude by stating for the record that CF_3I is not odorless as the MSDS and Brian contend, but possesses a sickly sour smell which I will fondly remember for years to come.

CHAPTER 1
DARK MATTER, BUBBLE CHAMBERS, AND COUPP

The evidence supporting the current majority scientific consensus on the dark matter model is presented in this chapter, along with a discussion of dark matter candidates and grounds for thinking dark matter is composed of non-baryonic WIMPs (Weakly Interacting Massive Particles). Immediately following is a discussion of specific WIMP candidates, theoretical WIMP properties, and ways WIMPs interact with ordinary matter. After an overview of direct and indirect detection methods, the direct approach utilized by the Chicagoland Observatory for Underground Particle Physics (COUPP) experiment is reviewed.

COUPP utilizes the detection of single macroscopic bubbles induced by individual high-$\frac{dE}{dx}$ nuclear recoils in dense-liquid bubble chambers (Behnke et al., 2008). The bubble chamber technology was not originally conceived as a dark matter detector. It was invented by D. Glaser (Glaser, 1952, 1953, 1954), who adapted the idea from cloud chambers (Glaser, 1994), in which radiation catalyzes condensation in gas, whereas in a bubble chamber it forces a liquid to begin to boil. (He is famously alleged to have thought of the idea observing bubbles while drinking beer.) Glaser won the Nobel Prize for his discovery in 1960. This chapter will present some of the bubble chamber's applications in high energy physics following Glaser's discovery, together with its current implementation in the context of COUPP.

1.1 The Dark Matter Model

Dark matter has a decades-long wealth of indirect evidence, but conclusive direct detection still eludes physicists. Plots of the rotational velocities of stars within a galaxy as a function of the distance from the center have exhibited, and newer data continue to exhibit, behavior inconsistent with an absence of dark matter, given the current understanding of gravity.

There apparently is insufficient mass to explain the galactic rotation curves (Figure 1.1) if it is assumed that all of the matter within the galaxies is contained primarily within the constituent stars and gas (Zwicky, 1937; Rubin et al., 1980).

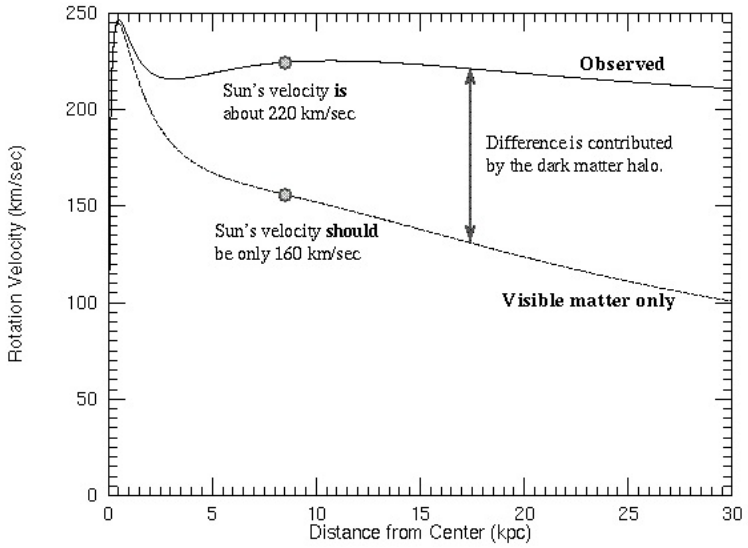

Figure 1.1: Actual and predicted (based on baryons) Galactic rotation curves for the Milky Way. Note the flatter, higher-speed behavior at high radius. According to Newtonian dynamics, the velocities should decrease if all the mass in galaxies is contained in luminous matter like stars, and other identifiable mass: past a certain point objects in the galaxy are beyond a radius containing most of the mass that can be taken into account and are getting further away, so they should be experiencing a decreasing gravitational force. Figure courtesy of Nick Strobel, Bakersfield College Physical Science Dept.

Furthermore, galactic masses can be determined by gravitational lensing measurements. In this phenomenon, the strong gravitational field produced by a large mass distorts light coming from distant objects behind it. Photons follow the shortest path along a curved space-time, creating warped images for the viewer or multiple images of the same object. Studies of gravitational lensing effects concur with the rotation curve data, pointing toward the idea that there is not enough mass in the form of visible matter (Refregier, 2003). The term "dark matter" was coined by astrophysicists to refer to this missing mass. The prevailing cosmological model is called "ΛCDM," Lambda representing the cosmological

constant thought to be behind accelerating expansion, Einstein's own usage (Riess et al., 1998), and CDM standing for "Cold Dark Matter" (Section 1.2).

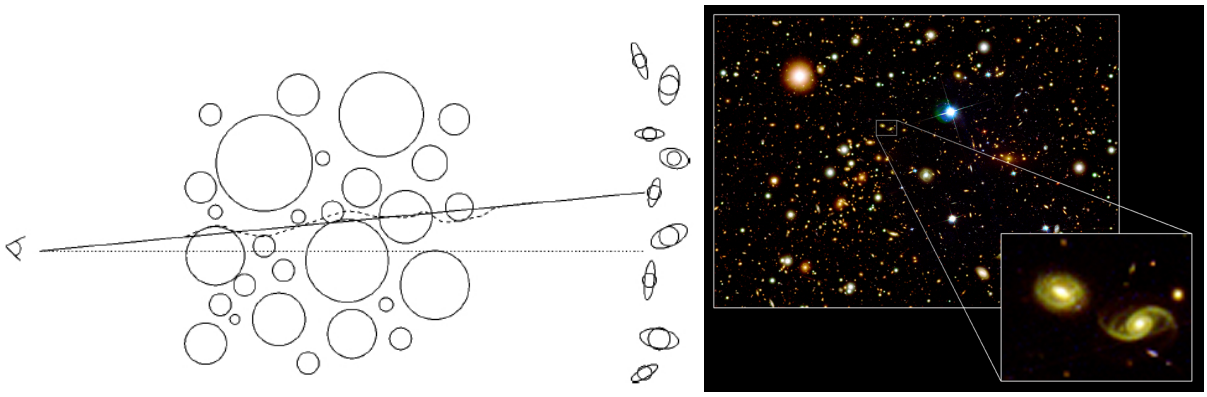

Figure 1.2: (left) Cartoon illustrating gravitational lensing. Photons experience gravitational pull of masses on the way to the observer, leading to blurry images, or magnification like in a lens. Figure from Refregier (2003).
(right) Image of the Bullet Cluster, the galaxy cluster which currently provides the best evidence of dark matter (Clowe et al., 2006). Example of galaxy duplication caused by weak gravitational lensing (inset). Mass contours were calculated using the lensing. They do not match a map of the visible matter. Public domain image courtesy of NASA.

1.2 The Case for WIMPs

Not all possible configurations of known matter radiate light. It was therefore once thought likely that certain massive bodies constituted the missing mass: black holes, protostars that had failed to ignite, etc. These baryonic dark matter candidates were termed MACHOs, or Massive Compact Halo Objects. However, empirical calculations have demonstrated that MACHOs exist in insufficient quantities to constitute the dark matter, at least within our own galaxy (Freese et al., 1999).

Astrophysical observations sensitive to the composition of the universe point to a cosmological model where $\sim$25% of the overall mass-energy content is non-baryonic matter. For instance, features of the temperature anisotropy power spectrum in the near-perfect blackbody spectrum of the Cosmic Microwave Background (CMB), the relic radiation of

the Big Bang, agree with this model when compared with predictions based on different cosmological models (Hinshaw et al., 2009). In order to explain large-scale structure, this non-baryonic matter, which is predicted by the CMB to exist in sufficient quantities to correspond with dark matter, must also be "cold" (i.e., non-relativistic). It must also be stable at least on the timescale of the current age of the universe, given the present-day existence of large-scale structures having evolved over billions of years. Lastly, it should be very weakly interacting with baryonic matter, primarily interacting through gravitational attraction caused by large "halos" of dark matter (Springel et al., 2005). The non-relativistic requirement implies a large mass.

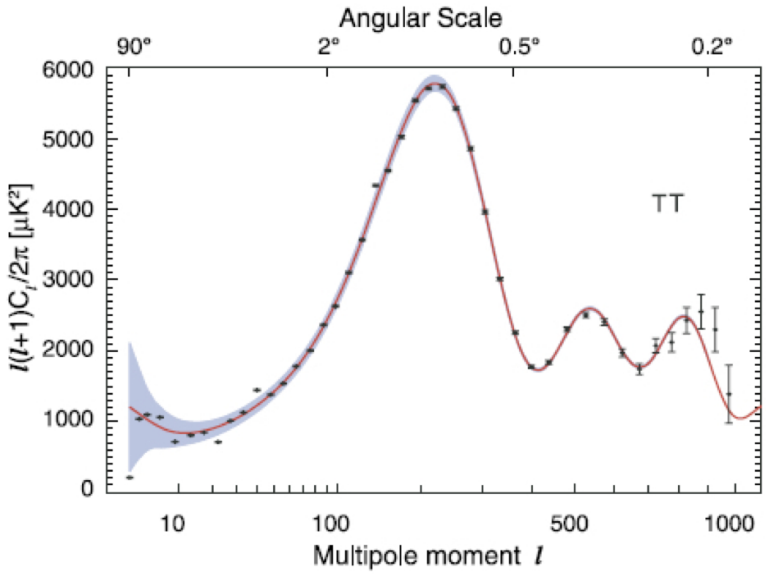

Figure 1.3: Peaks in the temperature-power spectrum of the CMB anisotropy. The colored lines indicate fits to the standard ΛCDM cosmological model, where ∼25% of the universe is cold dark matter. Figure from Hinshaw et al. (2009)

In contrast to MACHOs, the term Weakly Interacting Massive Particle (WIMP) has been applied to general particle dark matter candidates. Other theories offer candidates which are either light or possess significant non-gravitational coupling strengths such as electromagnetic for axions, or particles with both properties. Axions constitute a "natural"

explanation for why the strong force does not significantly break CP (charge-parity) symmetry as expected (Peccei & Quinn, 1977). However, such particles are beyond the scope of the WIMP detector described in this thesis.

Instead of invoking new particles, it is possible that a modification of the current understanding of gravity explains the observations in Section 1.1 of missing mass (Brownstein & Moffat, 2007). However, critics of modified gravity cite the artificial, non-relativistic modifications to Newtonian mechanics needed to explain empirical observations of galactic rotation curves (Clowe et al., 2006). In contrast, particles with the properties requisite for cold dark matter appear in natural extensions to the Standard Model of particle physics (Section 1.3).

1.3 WIMP Properties and Candidates

Experimenters quote results based on a shared set of assumptions about WIMPs in the Milky Way that comprises the Standard Isothermal Halo Model. A near-homogenous and isotropic spatial distribution of WIMPs in a spheroidal halo with $\sim$0.3 GeV/cm^3 density is capable of explaining the rotation curve. WIMPs in this model have a Maxwell-Boltzmann velocity distribution with mean of 220 km/s (Lewin & Smith, 1996). However, because the mass of an individual WIMP is unknown, such a mass density can lead to a large range of possible number densities, ranging from of order 10^4/m^3 for light WIMPs to less than 100/m^3, depending on the theory. The density of WIMPs is thought to be higher in some areas of the Galaxy and of the universe than others, with WIMPs tending to form large clumps (the aforementioned halos) under the influence of gravity. This clumping leads to baryonic matter being attracted and falling in to form galaxy clusters (Springel et al., 2005). However, the overall density of the WIMPs in the universe is thought to be static, ignoring the expansion of the universe and gravitationally bound local structures. The WIMP is

called a "thermal relic." "Freeze-out" occurred at some point during the early evolution of the universe (Kolb & Turner, 1990). As the universe expanded and cooled, the distance between WIMPs increased, while their temperature and speed decreased. According to the thermal relic model, WIMP-WIMP annihilation eventually became nearly impossible because an encounter between a WIMP and any particle with which it could interact, including another WIMP, became too unlikely. Lack of interaction with the rest of the universe, with the exception of gravitational interaction, led to the temperature freezing out relative to other particles. The WIMPs fell out of thermal equilibrium and their number density no longer decreased due to annihilation.

Known non-baryonic, weakly-interacting particles such as neutrinos have failed to be convincing explanations of dark matter. Known particles have too significant of couplings to forces other than gravity, or constitute at least mildly relativistic ("hot" or "warm") dark matter. Arguably, hot dark matter has been ruled out by the latest CMB measurements and computer simulations of large-scale structure formation with the cold dark matter model (Hinshaw et al., 2009; Springel et al., 2005). It of course does not clump as effectively as cold, challenging the dark matter explanation for large-scale structure formation. Therefore, the search for dark matter particles proceeds beyond the Standard Model of particle physics.

The WIMP could be the massive "lightest supersymmetric partner" (LSP) of a Standard Model particle. In supersymmetry (SUSY), for every fermion there exists a boson and vice versa. This so-called "superpartner" is more massive. The LSP is the lightest superpartner, but it is still heavier than most Standard Model particles and interacts only weakly with them, making it a natural WIMP (Jungman et al., 1996). Being the lightest new particle makes the LSP stable, and as discussed in Section 1.2, dark matter needs to exist long enough to draw baryonic matter into gravity wells to produce structures. In the prototypical SUSY model, the neutralino, a superposition of the superpartners of the photon, Z^0, and Higgs,

is thought to be a WIMP (Jungman et al., 1996). In addition to SUSY, the Kaluza-Klein (KK) theory of extra dimensions potentially offers candidate WIMPs. If there exist extra dimensions of space which are "rolled up" into a periodic geometry such as a cylinder, rather than possessing an infinitely linear extent like known spatial dimensions, then there would exist an infinite hierarchy of increasingly more massive copies of every particle. The lightest such particle (thus inherently stable) is called the lightest Kaluza-Klein partner (LKKP) (Appelquist et al., 2000). It is possible that dark matter is composed of more than one type of particle. Because dark matter has not yet been conclusively detected, experimentalists still assume the existence of only one WIMP to simplify data analysis.

1.4 Dark Matter WIMP Detectors

Gravity being the weakest force, the gravitational couplings of WIMPs to baryons in a laboratory setting is not a realistic approach to detecting dark matter. However, WIMPs are expected to interact, albeit very rarely, via low-energy elastic scattering off nuclei of normal matter (Lewin & Smith, 1996). The particle exchanged during this "billiard-ball collision" is model-dependent of course, but it is possible to build a detector that is simply sensitive to nuclear recoils. Thus it could search for WIMPs in a model-independent fashion. Theories which offer new particles that exhibit the correct properties to be WIMPs, like SUSY or KK, typically call for new particles between 10 GeV and several TeV in mass. These particles may not only elastically scatter, but may also annihilate each other in high WIMP density regions caused by gravity, such as the Sun or the center of the Galaxy. They would then possibly generate detectable annihilation products. Moreover, the particles could be produced in the latest, most powerful high-energy particle accelerators and colliders here on Earth and be detected in that manner.

1.4.1 Indirect and Collider WIMP Detection Methods

Indirect detection experiments involve the detection of WIMP annihilation products, such as gammas or neutrinos. Space-based telescopes, such as EGRET and FERMI-LAT, detect gammas (Cao et al., 2009). Experiments like AMANDA and ICECUBE detect neutrinos (Agrawal et al., 2010). Depending on the theory, WIMPs may annihilate into many different types of Standard Model particles with varying branching ratios. On the other hand, accelerators like the LHC should theoretically be able to produce WIMPs as proton-proton collision by-products, depending on the energies achieved and the mass of the WIMP, and on the size of the coupling of the WIMP to known particles (Agrawal et al., 2010; Cao et al., 2009).

1.4.2 Direct WIMP Detection Methods

The many existing direct detection experiments, reviewed and summarized extensively by (Gaitskell, 2004) and elsewhere, rely on translating nuclear recoils from individual particle interactions into a usable signal. The three primary techniques involve light from scintillation, charge collection from ionization, or heat (thermal phonons). Experiments will often combine multiple methods, especially to achieve particle discrimination, so that backgrounds can be understood, characterized, and rejected. Scintillation occurs when in certain substances, such as a liquid noble gas, recoiling nuclei or electrons create excited atoms or molecules which release photons when decaying back to the ground state. The number of photons is proportional to the energy of the recoiling species. These photons produce the photoelectric effect in photo-multiplier tubes (PMTs), which detect the resulting electrons. Some experiments have a dual phase liquid-gas detector and utilize primary and secondary scintillation stages. Ionization experiments rely instead upon charge collection in an electric field: electrons are knocked out by nuclear or electron recoils, leaving holes. Finally,

an experiment can focus on detecting phonons, by for instance implementing a transition-edge sensor (TES), which is a resistor with a sharp turn-on for superconductivity at low temperature. Phonons are vibrations in a crystal lattice caused by interaction at a point within the lattice. Very slight temperature variations caused by a recoil in a lattice can be detected by the dramatic change in the electrical resistance of a TES.

Depending on the mass of the target nuclei, a recoiling nucleus resulting from a WIMP collision can have kinetic energies as low as a few keV, given the lower end of the WIMP mass range mentioned above, so it is important to construct a detector not only sensitive to nuclear recoils, but also possessing a low enough energy threshold to search for WIMP candidates offered by existing theories. Still, if WIMPs exist, they may prove difficult or even impossible to find, at least given the size of present-day detectors. Supersymmetric WIMPs may for example interact at rates of lower than one recoil per ton per year, regardless of what is chosen as the target material (Jungman et al., 1996; Gaitskell, 2004). To maximize detection potential for such low-probability events, the challenge is to construct cost-effective, ton-scale detectors sensitive exclusively to WIMP-induced nuclear recoils. These should also be insensitive to other particle interactions, like elastic neutron scattering, which could be mistaken for WIMP scattering. A $\sim$1 kg Ge detector, a common radiation detector operating on the principle of charge displacement and collection in a semiconductor, if placed unshielded in a room on the surface of the Earth, fires at a rate of a few hundred Hz. This is due to cosmic rays and to terrestrial sources of radiation. Therefore, ultra-low-background detectors are necessary for the WIMP search.

Failing to detect a statistically significant signal above background, experiments set limits simultaneously on the probability for WIMP interaction, with nuclei or individual nucleons, and on the WIMP mass. Probability is quoted as cross-section, in units of area, representing the effective, not physical, area a target presents to a WIMP flux. It depends

on the size of the WIMP coupling to the scattering process. Confidence limits of 90% are traditionally presented in the field of direct WIMP detection. Because the kinematics are of course different for different nuclei, experimentalists have agreed to normalize result presentation to remove the A-dependence (Lewin & Smith, 1996). Direct detection experiments exhibit weakening sensitivity at lower masses because then the WIMP can be too light to generate a nuclear recoil above threshold. There is weakening sensitivity at higher masses as well because then the WIMP becomes too heavy to have a significant number density: the mass density in the Galaxy is theoretically fixed.

1.5 The Nature and History of Bubble Chambers

Bubble chambers are difficult to classify given the categorization of techniques outlined here. As will be discussed in Chapter 2, on bubble chamber thermodynamics, ionization is one of the processes involved. A stopping ion results in a highly localized, very sharp increase in temperature, so heat may also be the focus of understanding (though, not being solid crystals, bubble chambers are not phonon detectors, strictly speaking).

The bubble chamber technology, previously utilized in high-energy physics (HEP) for particle detection, capitalizes on the unstable, or more correctly, "metastable" equilibrium state which occurs when the pressure of a liquid is lowered (or, alternatively, its temperature is increased) adiabatically. It remains in the liquid phase despite being beneath the vapor pressure or the boiling point temperature (Glaser, 1994). The metastability of the liquid enables it to function as a particle detector. This false equilibrium collapses after enough energy is injected. The liquid seeks the now energetically-favorable gas state (see Figure 1.4).

The act of placing a liquid at a temperature and pressure where according to its phase diagram it should be gaseous, and having it nevertheless remain in the liquid phase, is known

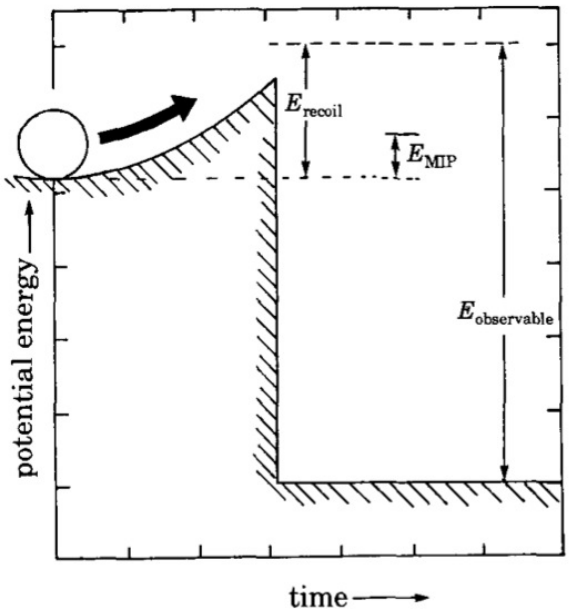

Figure 1.4: Potential energy as function of time for a superheated fluid. If an incoming particle has sufficient energy, then the nucleus (or electron) recoiling from it will cause the superheated liquid to overcome the potential barrier depicted and collapse to the lower-energy state, namely gas. Figure from Zacek (1994).

as "superheating." Liquids in such a state are called "superheated." At a fixed temperature, the difference in pressure between the vapor pressure and the operating pressure of a bubble chamber is known as the "degree of superheat." The higher the degree of superheat the less stable a liquid is. It becomes more sensitive as a radiation detector: particles of lower energy with correspondingly lower-energy recoiling nuclei are capable of generating detectable (macroscopic) bubbles (Collar, 1996). After an event, the pressure is raised to prevent further boiling, taking the chamber to a point where its fluid is "supposed" to be liquid (see Figure 1.5).

1.5.1 High Energy Physics Bubble Chamber Background

COUPP bubble chambers, including the one described in this thesis, revisit an old technology, with modifications leading to extended stability compared to traditional HEP bubble

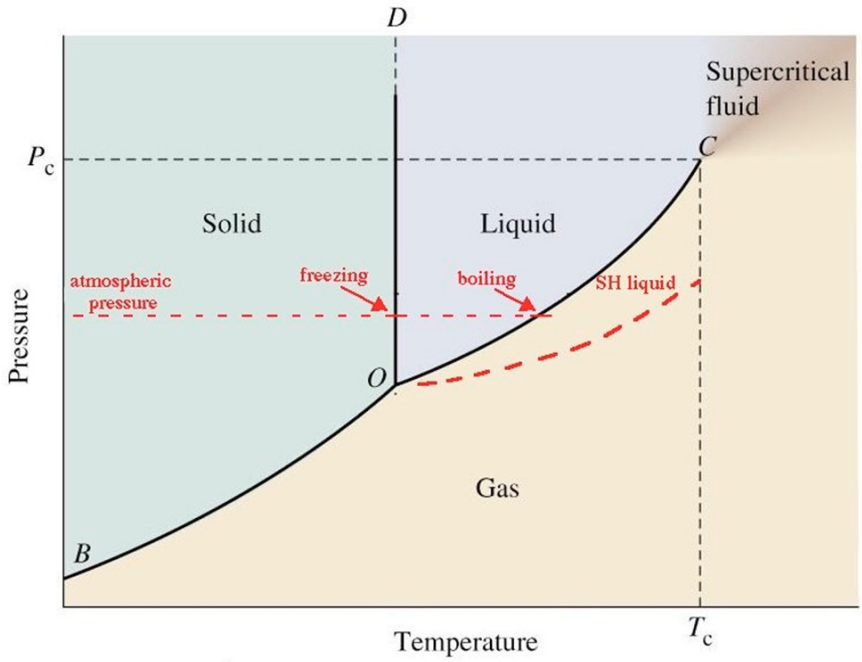

Figure 1.5: Example phase diagram illustrating the basic principle of the bubble chamber: staying liquid in the gas region of the diagram. COUPP achieves that by slowly lowering the pressure, beginning in the liquid region and entering the gas region without boiling. Boiling is avoided through use of a smooth vessel and a number of other precautions (Bolte et al., 2005). Figure courtesy of J.I.Collar.

chambers. Bubble chambers were once used as a common method of particle detection, tracking, and identification, until superseded in later years by technologies better suited to these purposes such as spark and wire chambers (Glaser, 1994). One drawback of the technology is the lack of knowledge of the energy of an interacting particle. A bubble chamber is fundamentally a threshold detector; any particle interaction above the thresholds in energy and stopping power leads to a bubble (Behnke et al. (2008) and ref. therein). Particle decay branching ratios were studied and new particles were often discovered with bubble chambers, which famously produced many intricate images of bubbly tracks in the liquid stemming from particles losing energy along their trajectories. Some were curved if charged particles were traveling through a chamber placed in a magnetic field for purposes of particle identification using the Larmor radius.

(a) 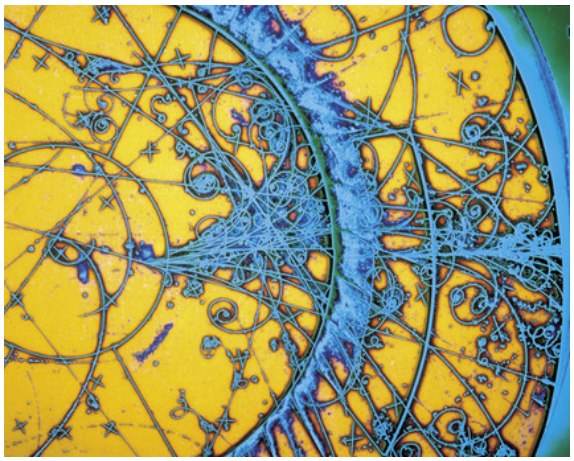 (b) 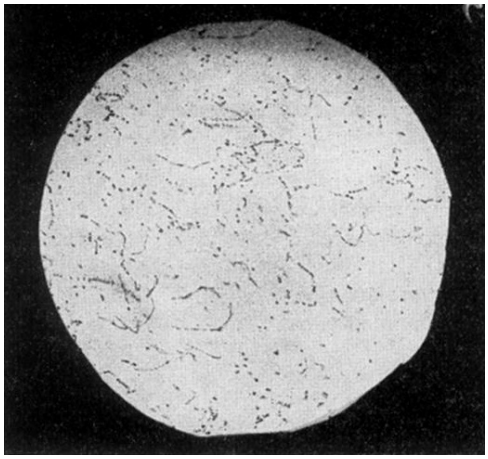

Figure 1.6: (a) Artistically-enhanced particle tracks from the the Big European Bubble Chamber (BEBC), CERN. The straight tracks are neutral particles, while the curved are from charged particles traversing a magnetic field. Tracks which disappear are from short-lived particles. New branches forming where they disappear correspond to decay products. Image from the CERN Courier (Wenniger, 2004).
(b) One of D. Glaser's early bubble chambers, filled with superheated ether. The bubble tracks are from gamma-ray-induced electron recoils (when the degree of superheat is very high). Figure from Glaser & Raim (1955).

1.5.2 Application of Bubble Chambers to WIMP Searches

To be useful as dark matter detectors, bubble chambers needed to be made more stable. In HEP, the bubble chamber operator knows when the particles of interest are due to arrive at the chamber because it is at an accelerator in a beam line. The operator can then superheat a fraction of a second (typically of order ms) before that time. Bubble chambers used to study cosmic rays need not stay superheated for long either, because the large flux of cosmic rays on Earth ensures a high reaction rate (Glaser, 1994). However, one major objective of any dark matter detector is to have a large livetime fraction. HEP chambers suffered from inhomogenous nucleation on gaskets, window ports, and other rough surfaces or edges. Their high degree of superheat, necessary for electron recoil sensitivity (see Section 2.5), exacerbated this issue and would lead to rapid boiling. Because WIMPs interact weakly, i.e. rarely, with baryonic matter, a detector must be stable and ready to

interact for as long as possible. The rarity of the interaction also changes the nature of the bubble chamber between HEP and direct WIMP detection from that of a tracking device full of multiple tracks of small bubbles from different particles traversing the detector, to a counting experiment. Regardless of degree of superheat, a successfully interacting WIMP would only scatter once, a one-bubble event already being highly improbable.

Chamber stability was addressed by COUPP with chamber material selection and cleaning methods. Every COUPP bubble chamber has been made from quartz, beginning with the earliest prototypes developed at the University of Chicago. Quartz has been chosen not only for its chemical compatibility with the superheated liquid inside (CF_3I) but also for its smoothness (Bolte et al., 2005). Surface inhomogeneities can lead to boiling, caused by pockets of gas trapped in voids. HEP bubble chambers had rough metallic wetted surfaces and seams between different parts, whereas in COUPP chambers the active liquid sees only the smooth surface of a seamless quartz jar. However, even the smoothest of possible surfaces for the container of superheated liquid cannot guarantee long-term chamber stability, so other steps must be taken to ensure it (Section 4.5).

1.5.3 Historical Trajectory of COUPP

The first dark matter limits (Figure 1.8) produced by the Chicagoland Observatory for Underground Particle Physics were achieved with a 2 kg (1 L) engineering prototype which, despite suffering from a large alpha background, produced the best spin-dependent proton limits of the time over a significant mass range (Behnke et al., 2008). This 2 kg chamber was built at the University of Chicago and tested in the subbasement lab at the Laboratory for Astrophysics and Space Research (LASR) at a depth of six meters of water equivalent (m.w.e.). This is a standard definition of depth for classifying underground sites where one meter of water is physically equivalent to $\sim$1 foot of rock in terms of cosmic ray attenuation. (This is merely an approximation; there are formulae for calculating the depth, taking

into account factors such as rock density.) The 2 kg chamber was then redeployed deeper underground at the Fermi National Accelerator Laboratory (FNAL) in the Neutrinos at the Main Injector (NuMI) beam underground hall, at a depth of 300 m.w.e. Going deeper is always better for any rare event search, because it reduces the cosmic ray flux and achieves a lower background. The results from the 2 kg chamber conflict with a spin-dependent (SD, see Section 2.6.1) explanation of the DAMA collaboration's longstanding claim of direct WIMP detection based on a signal observed in NaI[Tl] crystal scintillators (Bernabei et al., 2010). The DAMA signal changes annually in a manner consistent with what would be expected from WIMPs. As the Earth revolves around the Sun, its velocity relative to the Galactic WIMP halo changes, sinusoidally changing the WIMP interaction rate with earth-bound detectors. The SD explanation in the moderately-low mass region was the last "conventional" explanation without resorting to less popular, lower-mass WIMP theories, or to theories involving axions or axionlike electromagnetically-interacting scalar particles, ion channeling, inelastic scattering, or modifications to the standard Galactic dark matter halo model (Behnke et al., 2008). Other allowed regions in phase space for the DAMA signal, in spin-independent (SI) space, and for higher-mass SD-interacting WIMPs, had long been ruled out by other experiments, CDMS and KIMS, respectively, with 90% confidence (Akerib et al., 2004; Kim, 2008).

The sensitivity of the 2 kg device had been limited by an alpha background. The chamber was decommissioned and then redeployed in 2007 eliminating sources of radon, because an analysis of the first set of data demonstrated that it was consistent with 100% radon and radon daughters with 100% recoil-nucleation efficiency (Behnke et al., 2008). Time between single bubbles indicated their presence because of their only few-minute-long half-lives. In the new run of the 2 kg chamber, the following countermeasures were in place: metallic gaskets instead of radon-permeable o-rings (Wojcik, 1991), a custom bellows with

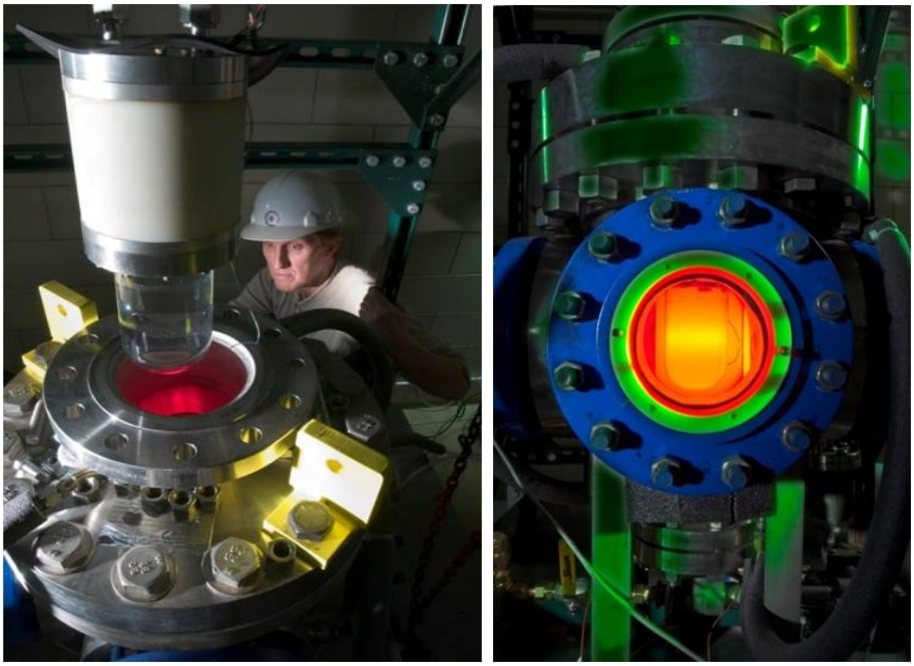

Figure 1.7: Photos from the commissioning and running of the 2 kg chamber in the NuMI gallery. The 2 kg chamber, unlike the 15 kg chamber, used windows rather than submerged cameras. The red lighting was necessary to maintain CF_3I stability (see Section 2.6.3). The quartz vessel installation for the 15 kg chamber was analogous, as shown in Figure 4.2. Images courtesy of the COUPP Collaboration.

lanthanated and electron-beam welding instead of thoriated weld lines, and, for the buffer liquid above the active liquid (necessity of buffer discussed in Section 4.1.1), SNO light water (<1 part in 10^{14} g/g uranium and thorium) (SNO, 2002). The steel outer pressure vessel from the 2 kg vessel was re-used for a 4 kg chamber (Behnke et al., 2011), while plans for a 60 kg chamber, already deployed and taking data at FNAL as of the writing of this thesis, went forward.

The focus of this thesis is a 15 kg chamber primarily constructed by the author that serves as the prototype for a new chamber design developed by J.I.Collar, and preliminarily pursued by B. Odom. The 2, 4, and 60 kg chambers have all used a steel pressure vessel with glass viewports to house a quartz chamber. In contrast, the 15 kg chamber had no viewports. The other chambers had pressure-resistant quartz windows set into viewports welded onto

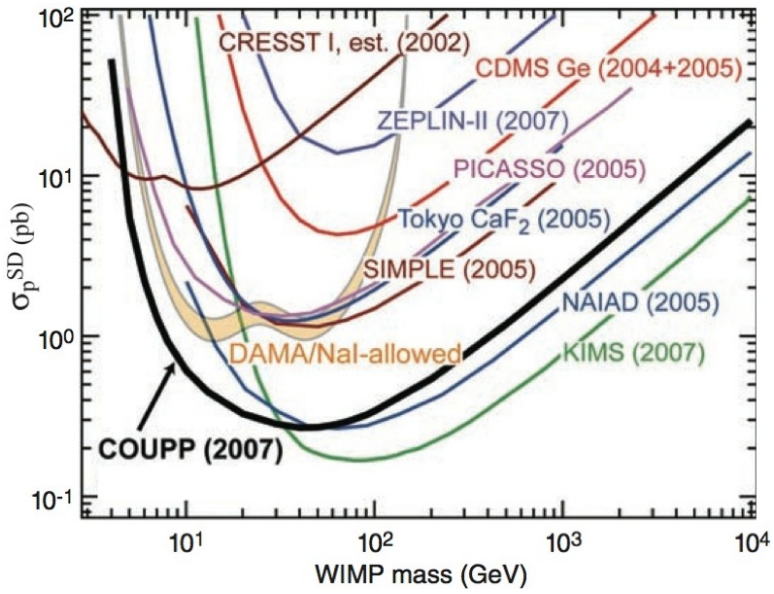

Figure 1.8: The spin-dependent WIMP-proton interaction limits from the 2 kg chamber, juxtaposed with limits from other experiments. The limits are presented as cross-section versus mass, the standard presentation for WIMP limits. The DAMA-allowed region, previously only partially ruled out by KIMS, above 50 GeV, is fully ruled out with 90% confidence by these data. Figure adapted from Behnke et al. (2008).

the steel vessel. The windows allowed bubble viewing with external cameras. The 15 kg chamber offered the advantage of not having to fashion windows capable of withstanding high pressure because its cameras were submerged under pressure in a dielectric liquid outside the active liquid. The 15 kg chamber used a standard commercial water pipe, rated for pressures exceeding those ever encountered during operation. The windowless design makes it possible to produce bubble chambers faster, more cost-effectively, and more safely, enabling other applications outside of dark matter detection (see Section 1.6). Past bubble chambers could undergo window failure due to the pressure, or worse, explosion due to the use of liquid hydrogen as the active volume (Habfast, 2010). However, the disadvantage is the need to make the cameras inaccessible by placing them inside. In the event of failure, stopping a run and temporarily decommissioning a chamber would be necessary. The camera lenses were sealed with an epoxy resistant to pressure, temperature, and oil, the external

hydraulic fluid surrounding the quartz vessel (their small indicator LEDs were sealed with a dark RTV to avoid a possible reaction with CF_3I discussed in 2.6.3). Sealed lenses were successfully pressure-tested, but as a further precaution, twice as many cameras as needed were used. The observed failure rate during actual running was two out of twelve cameras. The cameras were submerged in mineral oil, the hydraulic fluid of choice for the 15 kg chamber because it is dielectric. It does not short the naked printed circuit boards of the cameras.

1.6 Windowless Bubble Chambers in Other Fields

The compact practical design of a windowless chamber easily lends itself to a novel application: homeland security. Small-scale windowless bubble chambers could feasibly be used as portal monitors. They would be safe for this use, having no high-pressure windows. Neutron interrogation methods would be used to inspect cargo for fissile materials in a standard nuclear "car wash" configuration (Hall et al., 2007). A low energy neutron source would be directed at the cargo container, and a bubble chamber would detect the ejection of higher energy induced fission neutrons, which for a long or wide enough chamber will produce multiple bubbles more often than single (Section 2.5.1). Other backgrounds produce single bubbles alone (Section 2.5). The bubble chamber is also advantageous as a portal monitor because it is insensitive to interactions caused by the interrogating beam, thermal neutrons or gammas, while still sensitive to fast neutrons.

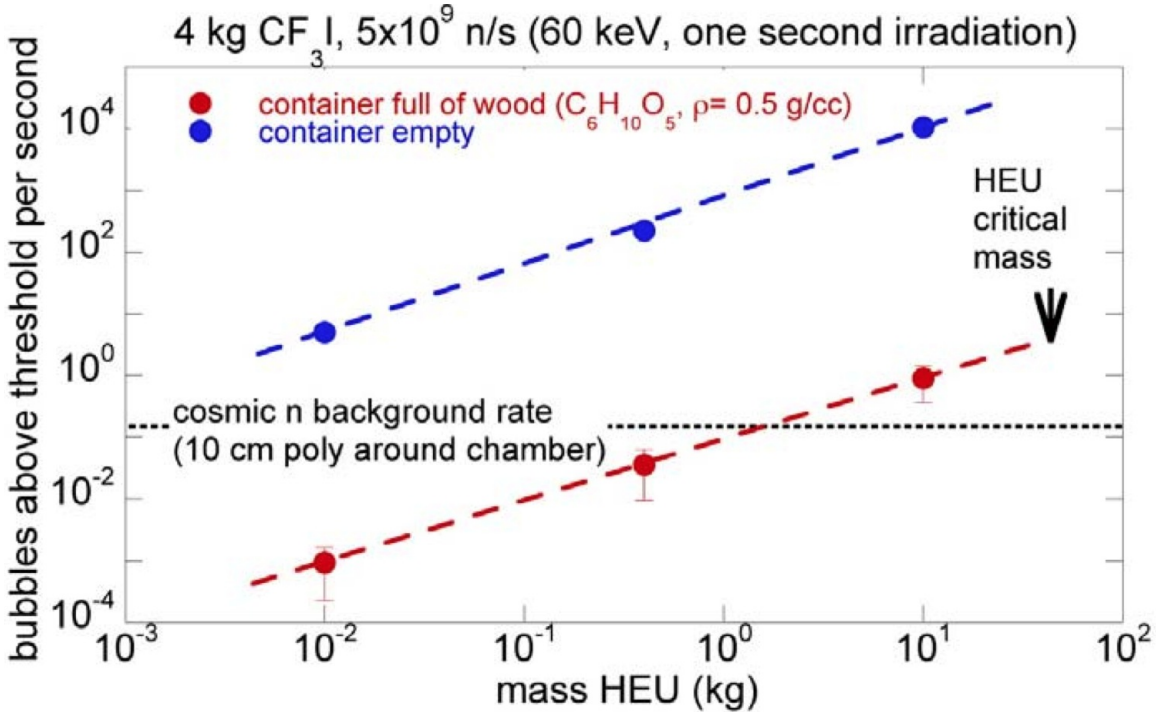

Figure 1.9: Illustration of the power of the bubble chamber to detect fission neutrons emanating from highly-enriched uranium (HEU) hidden in cargo. The response of a 4 kg bubble chamber to hidden fissile cargo using an interrogating 60 keV neutron source is simulated. A small chamber should be able to detect neutrons above background from a few kilograms of HEU, even if heavily shielded. Figure courtesy of J.I.Collar.

CHAPTER 2
BUBBLE CHAMBER THERMODYNAMICS AND FLUID PROPERTIES

This chapter covers the theory of the physics of bubble formation and evolution in the context of a superheated liquid, and the energy and stopping power thresholds for successful nucleation. The principle of the bubble chamber is simple but strong. Even only one unseen microscopic interaction, a nuclear or electron recoil, fission, or spallation turns into the macroscopic signal of boiling: observable bubbles. This chapter also reviews the nuclear, physical, and chemical properties of the chemical CF_3I, the active superheated liquid used in the 15 kg chamber and all COUPP chambers. It enumerates the advantages of using it as a WIMP detector.

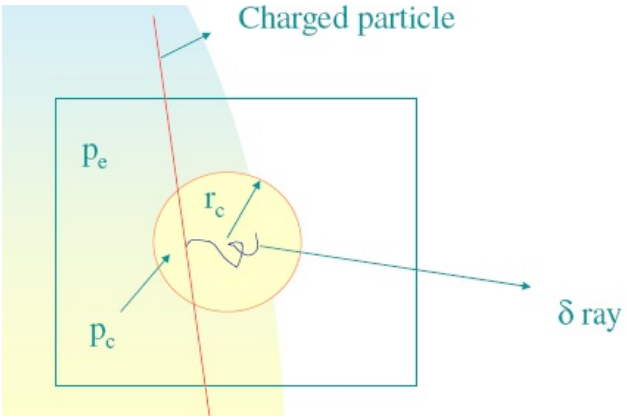

Figure 2.1: A sketch of initial bubble creation. p_e is the bubble chamber pressure, p_c (elsewhere referred to as P_v) is the vapor pressure of its liquid, and r_c represents the critical radius. A moving ion loses energy, through Coulomb scattering for example, resulting in electron ejection (δ rays). Combined with other energy deposition mechanisms, these may induce enough local heating through Coulomb interactions to create a pocket of vaporization. If the original incident particle interacts with nuclei or electrons of the liquid often enough in a distance greater than r_c, a sustainable bubble is born. A bubble larger than r_c can grow indefinitely, fueled by the pressure differential ($p_c > p_e$). Figure adapted from Lau (2008).

2.1 Underlying Principle of Bubble Chambers

A charged particle loses energy along its trajectory through a superheated liquid via ionization and other means, such as collision and radiation. The thermodynamical model developed by Seitz (1958) makes no distinction regarding the means of energy loss. Only total energy deposited is relevant. The particle is the recoiling nucleus for the case of a nuclear recoil resulting from a neutron, an alpha, or a WIMP, or the recoiling electron for the case of a colliding electron, high-energy photon, muon, or other minimally-ionizing particle (MIP). One energy deposition mechanism is ionization. Landau fluctuations are characteristic of this process. These are stochastic fluctuations in collision frequencies and energies. Energetic electrons known as delta rays, of a few keV in energy, are produced. Kinetic energy lost by the primary particle and the secondary particles it produced is converted into molecular motion within a very small volume of the liquid, causing it to undergo a phase transition and become a gas. Thus, the slowing primary particle leads to a temporary thermal excitation along its track. The temperature of the gas created is temporarily much hotter than the surrounding liquid. It becomes $\sim 10^4$–10^6 K (depending on molecular motion degrees of freedom, the pressure, and particle range), causing violent evaporation. The Seitz model is thus termed the "hot spike" model of bubble nucleation (Behnke et al., 2008). If the pressure of the hot gas is sufficient, a protobubble will overcome surface tension and grow without bound. Its growth is fueled by its own internal pressure, which is simply the vapor pressure of the liquid at the current temperature. This will of course be greater than the pressure outside the bubble because by definition a superheated liquid is below vapor pressure.

2.2 Critical Bubble Size

If the radius of the so-called protobubble thus created is smaller than a critical radius, itself a function of temperature and pressure, then it will recollapse (Blander & Katz, 1975; Seitz, 1958). Otherwise, the bubble will continue to expand indefinitely until the entire liquid boils away or becomes a frothy foam (Figure 2.2); vaporization cools the bubble surface, but not as fast as heat flows from the outside inward to sustain vaporization (Seitz, 1958). The critical radius has a concise form (Seitz, 1958; Peyrou, 1967; Apfel et al., 1985):

$$r_c = 2\gamma(T)/\Delta P, \qquad \Delta P = P_v(T) - P_e \tag{2.1}$$

where γ is the surface energy as a fluid-dependent function of temperature, T, P_v the vapor pressure at the current operating temperature, and P_e the actual current pressure of the liquid. For CF_3I, in the COUPP operating range of 30–40°C, the surface tension is 7–9 dyne/cm (NIST REFPROP). For the 15 kg chamber the vapor pressure was ∼77 psig at its mean temperature ∼30.5°C (manufacturer Ajay's vapor pressure curve). It was operated between 3–60 psig, making the resulting critical size for bubbles between ∼10–100 nm.

A bubble grows by incorporating surrounding molecules still in the liquid phase into its expanding sphere of gas. The process is slowed but never stopped by evaporative cooling. The formula derived for the post-critical size as a function of time is (Plesset & Zwick, 1952, 1954; Lau, 2008)

$$R = \frac{2(3/\pi)^{1/2}\lambda(T_c - T_e)V_m t^{1/2}}{h_v D^{1/2}}, \qquad D = \frac{\lambda}{\rho c} \tag{2.2}$$

where λ is the thermal conductivity, D the thermal diffusion constant, ρ the liquid density, V_m the gas molar volume, c the specific heat capacity, h_v the latent heat of vaporization per mole, T_c the critical temperature, T_e the ambient liquid temperature, and t is time.

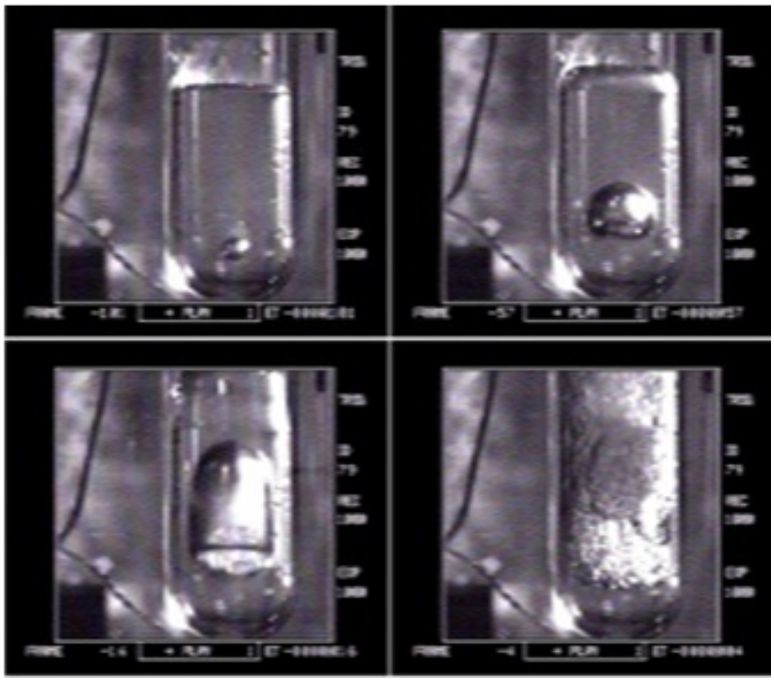

Figure 2.2: Intentional unchecked bubble expansion in an early prototype chamber (∼20 mL) at the University of Chicago for the purpose of filming the advanced stages of bubble growth, when deviation from sphericity occurred. The timespan from upper left to lower right was ∼ 0.1 s. Images courtesy of J.I.Collar.

2.3 Energy Threshold Calculation

A particle must produce a recoil with high enough energy and stopping power in order to make a bubble. Therefore, unlike other detectors, bubble chambers have two thresholds: energy and $\frac{dE}{dx}$. Both thresholds must be exceeded in order for there to be a signal, the formation of a sustainable bubble which becomes macroscopic (Roy et al., 1987; Peyrou, 1967). In other words, the energy deposition caused by a charged particle traveling through a superheated liquid must be dense, not just large. All means of energy deposition by a particle are treated equally by Seitz (1958). A particle must deposit its energy over a short enough distance, corresponding to what was referred to above as the critical radius for bubble sustainability. The thresholds can be expressed as (Collar et al., 2000; Roy et al., 1987; Peyrou, 1967; Harper, 1991, 1993a,b)

$$E > E_c = 4\pi r_c^2(\gamma - T\frac{\partial\gamma}{\partial T}) + \frac{4}{3}\pi r_c^3 \rho_v \frac{h_v}{M} + \frac{4}{3}\pi r_c^3 P_e \tag{2.3}$$

$$\frac{dE}{dx} > \frac{E_c}{ar_c} \tag{2.4}$$

E_c, critical energy (the energy threshold) is dependent upon surface tension γ, critical radius r_c, heat of vaporization h_v, mass-energy of one molecule or atom of the fluid M, temperature T, density of the gas phase ρ_v, and ambient pressure of the liquid P_e. The first term is work needed to maintain the increasing surface area (Collar et al., 2000), the second is the minimum work needed, related to the Gibbs free energy (Gibbs, 1957), and the third is the work against the external pressure provided by the ambient superheated fluid still in the liquid phase (Seitz, 1958). The threshold in energy deposited per unit distance depends on the critical energy, the critical radius, and a number scaling the critical radius known as the Harper factor. It is a single "free" parameter correction factor $\sim$O(1) (Harper, 1991, 1993a,b) that may be fluid-dependent and a function of temperature as well (Das & Sawamura, 2004). Neutron and gamma calibrations of previous COUPP chambers have shown best fits to Monte Carlo simulations with a Harper factor between 4–10 (Bolte et al., 2005; Behnke et al., 2008). This is in good agreement with expectations (Bell et al., 1974). As the temperature rises or pressure falls, and the liquid thereby becomes more superheated, both thresholds drop non-linearly because of their non-trivial dependences upon temperature and pressure through other variables.

The thresholds above are strict step functions, possibly unnatural to posit for a real-world system. In reality, there is some evidence that these are smoothed as sigmoids (Boukhira et al., 2000; Barnabé-Heider et al., 2005). However, the onsets of sensitivity are steep enough that strictly following the theory works well to predict real-life data (Bond et al., 2005; Bolte et al., 2005). In terms of neutrons at 40°C, little difference is made in using a steplike

threshold in the analysis of data, and in fact doing so is conservative, yielding slightly worse WIMP limits, at least for light WIMPs (Behnke et al., 2008). Better limits are achievable through a sigmoidal convolution with the monotonically decreasing WIMP spectrum, an exponential (see Figure 2.3). As threshold is decreased, WIMPs of increasingly smaller mass can make bubbles, but a sigmoid threshold results in sensitivity to lighter WIMPs and lower-energy recoils at a given nominal threshold than could be achieved if the sigmoid were instead a step function.

2.4 Recoil Detection Efficiency

Using a superheated liquid as the basis for a detector comes with the advantage of 100% efficiency. In the Seitz theory a recoil in the liquid surpassing the right energy and energy deposition thresholds must end in a bubble. In practice, a possible interpretation of neutron calibration data is that the efficiency is $\sim$O(50%) for CF_3I at 30°C (Bolte et al., 2005), though 100% is measured at 40°C. However, recent work (COUPP internal) suggests that a soft threshold but closer to 100% efficiency is also a good fit to the 30°C data. The efficiency has clearly been observed to approach 100% as the temperature is increased, given neutron and gamma calibrations in situ at 40°C and above (Behnke et al., 2008). These measurements led to the conclusion that bubble chambers asymptote towards ideal as temperature increases and a detector becomes more superheated and the thresholds become sharper. Onset of sensitivity to alphas is much sharper at 40°C than at 30°C (Behnke et al., 2008). All data from COUPP chambers operated at 40°C or higher have completely agreed with the Seitz formulation. The 15 kg chamber was restricted to 30°C (see section 6.4); other chambers are not, observing 100% efficiency as predicted by Seitz.

No matter the efficiency, a bubble chamber contains no exact energy information. All that is known is that the energy of the originating interaction for a given bubble surpassed

the calculable threshold. However, important information can be gleaned about distinct types of particles interacting by changing the threshold and taking more data. The spectrum of the number of bubbles produced as a function of threshold features certain simple functional forms. A WIMP signal, for instance, would be an exponential, becoming vanishingly small at high energies. The Galactic WIMPs, Maxwellian-distributed in velocity, have increasingly low probability of moving at increasingly high speed (Lewin & Smith, 1996). Neutrons are WIMP-like in spectral behavior (see Figure 2.3), becoming increasingly improbable at higher energies. Step or sigmoidal shapes are also possible, for interactions that can be abruptly "turned on" and off because of an effective maximum energy or linear energy transfer, like alpha-emitting radionuclides or alphas themselves, respectively (Behnke et al., 2008). Sensitivity to MIPs falls off as an extremely steep exponential (Section 2.5 and Figure 2.3). However, a spectrum cannot have a peak or spike, for example.

2.5 Effect of Superheat and of Backgrounds

Operating at a lower degree of superheat than used in past bubble chambers enables inherent rejection of MIP backgrounds, at levels as high as 1 in $\sim 10^{10}$–10^{11} interactions, without any data cuts. It is still possible to maintain full sensitivity to nuclear recoils for energy thresholds of interest, even for a moderately-light WIMP search (Bolte et al., 2005; Bond et al., 2005; Behnke et al., 2008). This is the best MIP rejection factor measured in the field of direct dark matter detection (Behnke et al., 2008). This makes it possible for bubble chambers to also perform measurements of cross-sections of interest in the field of nuclear astrophysics. They will be used to study a (γ, α) reaction while in an intense gamma beam (Rehm et al., forthcoming). The excellent gamma rejection benefitting both this gamma beam experiment and COUPP is made possible because recoiling electrons ejected from atomic orbitals by betas, gammas, x-rays, or muons, even when above the

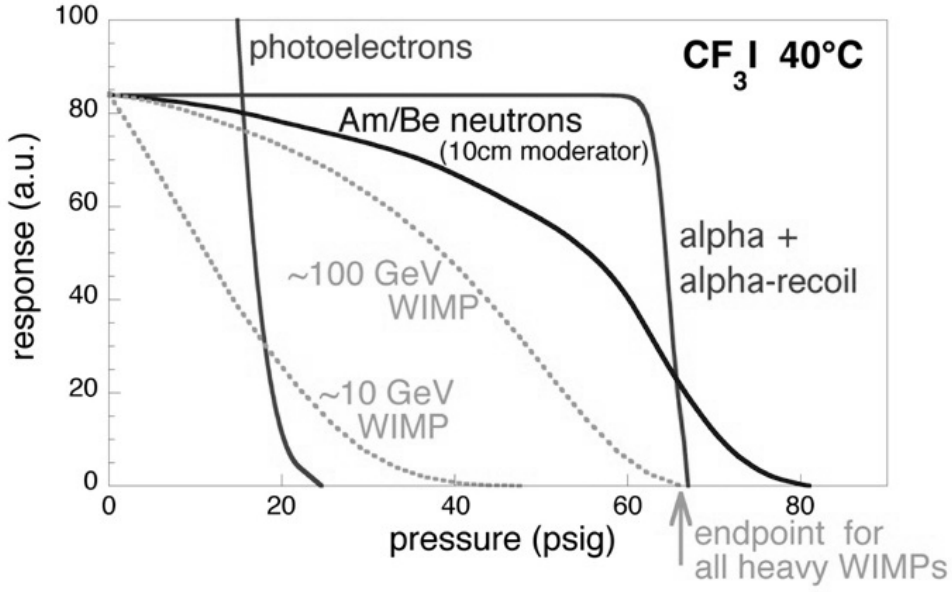

Figure 2.3: Response curves of different particles, based on fits to actual data for neutrons, alphas, and gamma rays, and calculated for WIMPs. Note that to retain the possibility of WIMP detection, any alpha background must be removed or discriminated, as the threshold cannot be raised high enough (by raising the pressure) to avoid it. Shapes appear to flatten at low threshold because at low pressure the critical energy changes very slowly, non-linearly. Figure courtesy of J.I.Collar.

energy threshold, do not exceed the required stopping power threshold (Figures 2.4 and 2.5). COUPP therefore avoids a need for high-Z materials such as lead to shield against high-energy gammas that would produce electron recoils. However, a muon veto is necessary at least at shallow depths like 300 m.w.e. because a high-energy cosmic-ray muon indirectly leads to secondary neutrons.

2.5.1 Cosmogenic and Environmental Neutrons

Gammas and betas may not be a concern, but neutron and alpha backgrounds must be addressed. An additional advantage of the COUPP technique is the short mean free path of neutrons in CF_3I, less than 15 cm even for energies in excess of 15 MeV (Chang, 2000). Moderated neutrons reaching a chamber after transport through shielding around the chamber

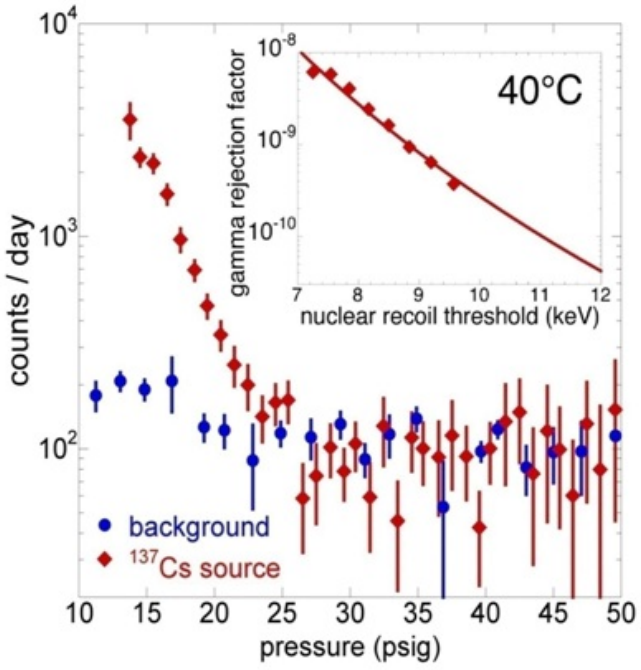

Figure 2.4: Gamma calibration of the 2 kg chamber. Results are presented with a ^{137}Cs source present or absent next to the chamber, which had a background of alphas and neutrons. Knowing the activity of the source used and the bubbling rate observed in the bubble chamber versus operating pressure, it is possible to determine the probability of a gamma interaction within the chamber, by comparing with results predicted by an MCNP (Monte Carlo Neutral Particle transport code) simulation. It is then possible to know how often gamma rays can successfully nucleate bubbles, as a function of the energy threshold. Figure courtesy of J.I.Collar.

are $\ll$1 MeV in energy and have a much shorter mean free path, of order a few cm (Figure 6.10). Large-scale wider chambers $\sim$O(100 kg) in mass, will thus be self-shielding according to simulation (Figure 2.6). Polyethylene, water shielding, or a combination thereof have been used by current and past chambers, like the 15 kg chamber, in order to thermalize environmental and cosmogenic neutrons. This places them far below even minimum thresholds. Shielding is useful because the geometry of small-scale COUPP chambers deployed thus far have allowed for the possibility of seeing single bubbles from neutrons. Moreover, even if identifiable it is best to remove as many backgrounds as possible, maximizing the livetime fraction.

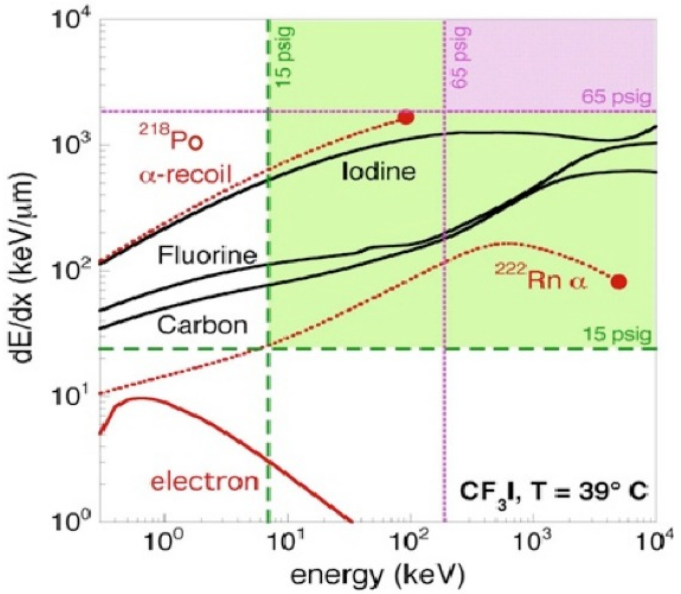

Figure 2.5: Critical $\frac{dE}{dx}$ vs. the critical energy. Nuclear recoils of all nuclei in CF_3I are plotted, plus the alpha emitted by ^{222}Rn and the recoil of its daughter nucleus. For fixed temperature, pressure regions are colored. Only species reaching into the upper-right hand quadrants thus delineated are capable of bubble nucleation. Note easy avoidance of electron recoil region while still in the nuclear recoil region. Figure from Behnke et al. (2008).

2.5.2 Alpha Particles

Alphas, the only other possible background radiation left, are addressed by studying the time correlations between bubbles caused by a series of short half-lives for radon decay products for instance (Behnke et al., 2008). This is not, of course, effective for longer-lived isotopes. Alphas form a flat background leading up to a loss in sensitivity to alpha interactions at a low degree of superheat, so this background can be subtracted if fully characterized across different pressures. Fits to data can be done using a model which includes this flat alpha background and a WIMP signal (Behnke et al., 2008). This is not ideal though, because in order to generate the best limits on WIMP coupling dark matter experiments aim to be "zero background." (Akerib et al., 2004) Mathematically, this makes the limits the strongest because then the statistical significance of number of events above a background

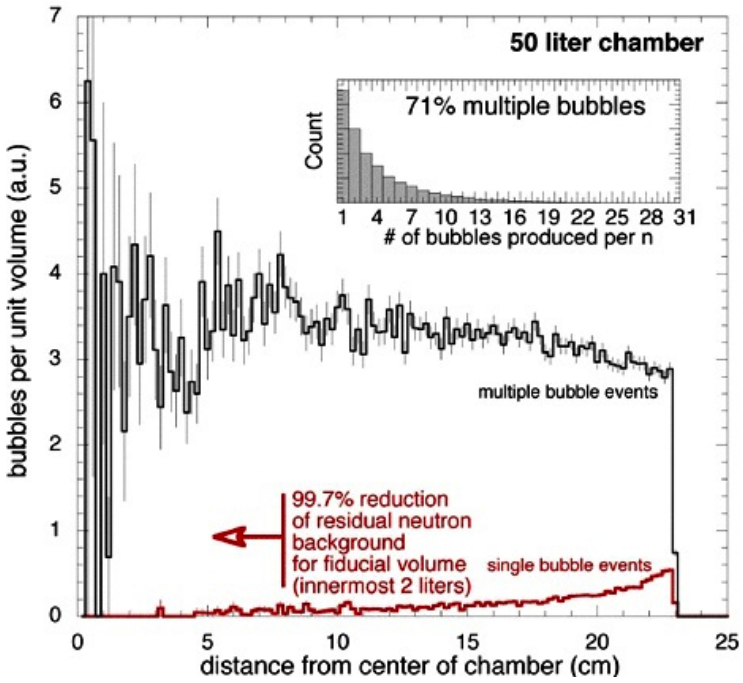

Figure 2.6: Monte Carlo simulation results for a future large-scale COUPP detector. The plot illustrates the ability to discriminate neutrons from WIMPs based on bubble multiplicity and effective fiducialization, neutrons penetrating deep into the bulk volume rapidly becoming less likely to produce single scatters. Neutrons tend to elastically scatter in collisions exceeding the thresholds more than once. Figure adapted from Bolte et al. (2006).

is not a factor (see Chapter 9). Alphas can be "turned off" by selecting a temperature and pressure at which thresholds are too high, but doing so also makes the detector insensitive to WIMPs (Figure 2.5). Even if tagged, alphas increase the dead time. The 60 kg chamber will attempt to do greater fluid purification, and the 4, 15, and 60 kg chambers all attempt to discriminate alphas, so that any remaining are completely removed from the analysis (Chapter 7). Event-by-event discrimination is of course preferable to spectral analysis. In the latter case, the type of particle responsible for individual events is unknown. This leads to more ambiguity of interpretation and worse WIMP limits mathematically.

Despite alleged identification of individual events, one may still need to account for WIMP candidates found after opening a pre-determined "box" of parameter space deter-

mined by data cuts in a blind analysis (Aprile et al., 2010). Therefore, it is important to have as many handles as possible on understanding backgrounds: those common to all direct WIMP detection experiments and those unique to a particular detector. Otherwise, a claim of discovery is all the more difficult to make.

2.6 Trifluoroiodomethane as the Target Material

CF_3I, known as Trifluoroiodomethane or Trifluoromethyl Iodide among other trade names, is a low-cost ($100–400/kg, compare to ∼$10,000/kg for the raw high-purity germanium used in detectors), fire-extinguishing industrial refrigerant related to Freon^{TM} and created as an environmentally-friendly replacement for chemicals more harmful to the ozone layer (CF_3Br and CFCs) (Yuanyuan et al., 1999; Solomon et al., 1994). COUPP purchases CF_3I for its bubble chambers from the chemical manufacturer Ajay North America, LLC.

2.6.1 WIMP Detection Properties

CF_3I is an ideal target molecule for WIMP searches in two essential aspects. By virtue of containing fluorine, whose dominant isotope's nucleus has an odd number of protons (^{19}F), and iodine, which has a heavy dominant isotope (^{127}I), it is sensitive to two different types of WIMP reactions simultaneously. The unpaired proton of the ^{19}F nucleus permits sensitivity to a possible WIMP-proton axial or spin-dependent (SD) coupling (Ellis, 1991). The WIMP scalar or spin-independent (SI) coupling, being a coherent interaction involving an entire nucleus, has a cross-section which roughly scales with the total number of nucleons squared ignoring form factor, making it generally better to use heavier nuclei (Lewin & Smith, 1996).

Because the nature of the dominant WIMP coupling remains unknown, a two-front approach to detection is prudent (Collar et al., 2000; Mayet et al., 2002). Although the undiscovered WIMP's spin and cross-sections for interaction with protons, neutrons, and nuclei are unknown, WIMP candidates found in various theories predict a variety of values

(a)

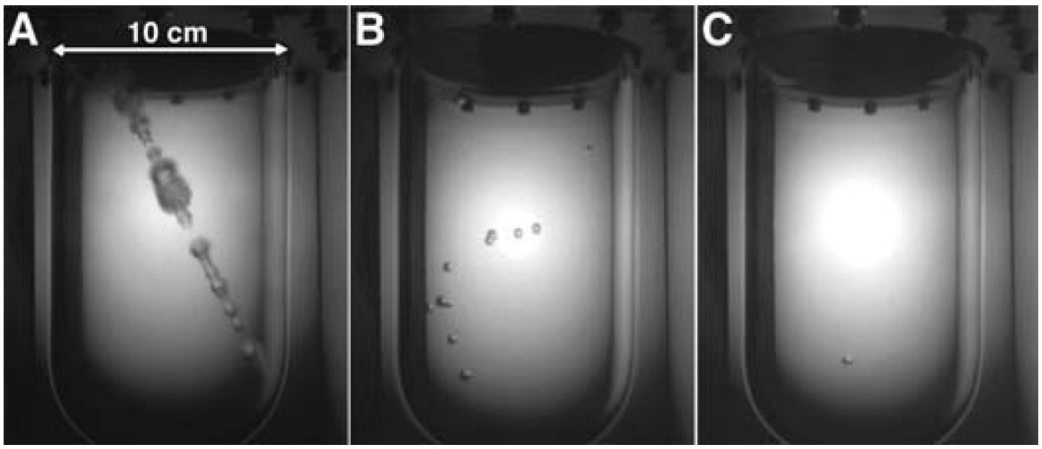

(b)

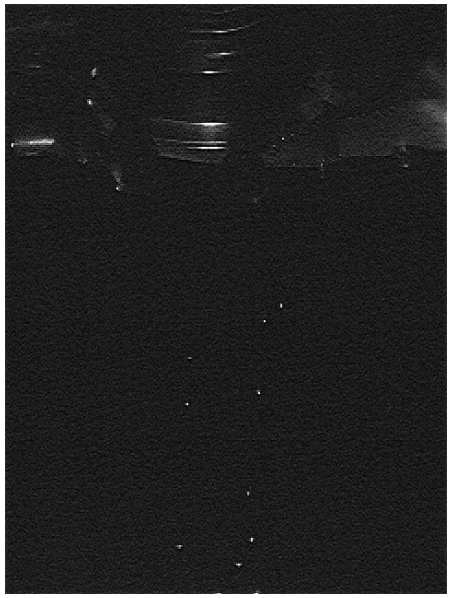

Figure 2.7: (a) Categorization of bulk COUPP events. (A) shows a track in the wake of a muon at 60°C in the 2 kg bubble chamber. This event was recorded for educational purposes only, illustrating how early bubble chambers interested in tracking recoiling electrons would function. Such an event would not occur during normal COUPP operating procedure, where the temperature never exceeds 45°C and the pressure is never below atmospheric. Normal COUPP events are either discrete multiple (B) or single (C) scatter in nature. Multiple bubbles must be caused by neutrons, but singles can come from neutrons, alphas, or WIMPs. The single-scatter neutron probability is reduced with neutron moderator, and in future chambers with self-shielding combined with a fiducial volume cut. Alphas are globally distinguished by their flat spectrum vs. minimum energy, and individually using half-life or sound (Chapter 7). Figure from Behnke et al. (2008).
(b) High-multiplicity neutron event from the 15 kg chamber (the bubbles are the scattered bright spots), easier to achieve due to its size. (See top of Figure 3.6 for examples of single bubbles in the bulk.) Notice the difference in illumination, which will be treated in Section 3.1.

for the cross-sections. The SI phase space of WIMP masses and interaction cross-sections is currently more thoroughly explored by direct detection experiments, i.e., down to lower cross-sections, because it can be reached by sheer mass. However, the SD parameter space is more compact, with predicted points of WIMP mass and cross-section not spread across as many orders of magnitudes as in the case of SI, making the SD predictions more robust (Bertone et al., 2007; Bednyakov & Klapdor-Kleingrothaus, 2001). The COUPP detector offers an ideal combination: ^{19}F is the nucleus in nature most sensitive to the SD-proton coupling, and the heavy ^{127}I surpasses or competes equally with the isotopes of other direct detection experiments for SI coupling (compared with ^{76}Ge, ^{127}Xe, etc.) (Bertone et al., 2007). Use of CF_3I also grants sensitivity to lighter WIMPs than other leading dark matter experiments by virtue of possessing such light nuclei as fluorine and carbon.

2.6.2 Physical Properties

The density of CF_3I, $\sim$2 g/cm^3 in the 30–40°C range, offers an advantage over liquid argon experiments (argon density is 1.4), though CF_3I is less dense than germanium, silicon, or xenon. However, it makes up for this with its ease of scalability. A liquid detector is easier to purify in bulk. Assuming a background-free experiment, or one in which the backgrounds do not scale up with mass, it is mass together with livetime which drives the quality of the dark matter limits produced by a detector. The exposure, measured in units of kg-days or ton-days, is the product of detector mass and livetime. Mass and time both serve to increase the number of WIMP events possible.

The density of CF_3I also provides the beneficial ability to use water as a buffer liquid "top cap" (Das et al., 1987) to fill the opaque stainless steel bellows so that bubbles do not occur where they cannot be observed via camera, nor occur due to CF_3I contact with rough surfaces. On such surfaces, CF_3I gas can get trapped in crevices. This leads to inhomogeneous nucleation, like in HEP chambers of the past (see Section 1.5.2). As outlined

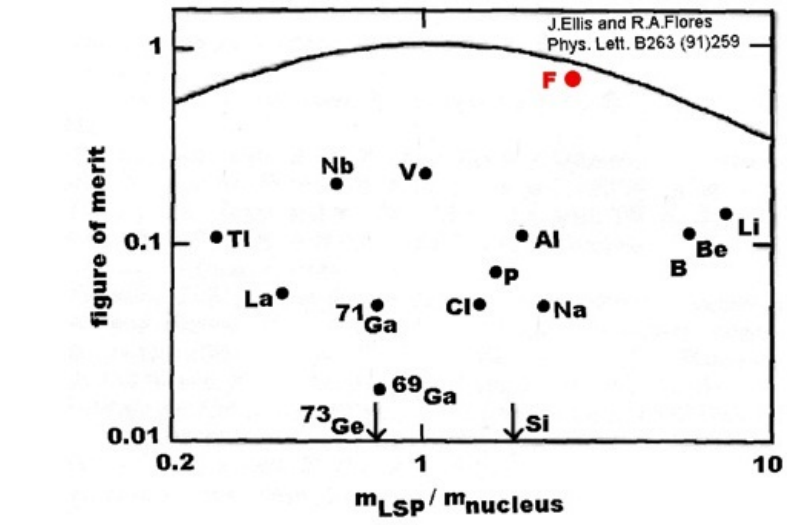

(a)

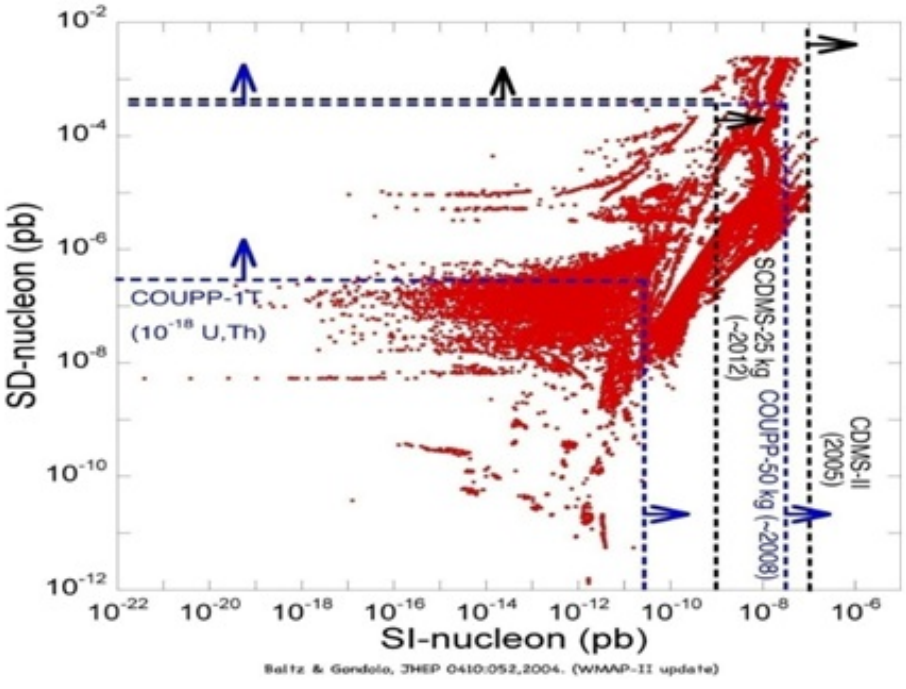

(b)

Figure 2.8: (a) Sensitivity to SD coupling of various elements to a supersymmetric WIMP, underscoring desirability of using fluorine in a WIMP detector. Figure courtesy of J.I.Collar, adapted based on information found in Ellis (1991).
(b) Graph of SD vs. SI cross-section illustrating the prudence of a two-pronged effort to directly detect WIMPs, as well as the projected ability of COUPP to compete with leading experiment CDMS, especially on the SD front. Red dots are predicted masses and cross-sections for various supersymmetric models. Top and right are excluded regions. Figure courtesy of J.I.Collar, adapted from Baltz (2004).

in Section 4.5.2, the water can be evaporated into a chamber first. As a chamber is filled with CF_3I it displaces the water upward into the bellows. In addition to being twice as dense as water, CF_3I is nearly immiscible with it (its miscibility is a few mmol/L).

The immiscibility of CF_3I and water leads to another important point. Since water and CF_3I are immiscible and since "like dissolves like" in terms of solubility (Williamson, 1994), uranium and thorium salts which dissolve readily in water should not dissolve as effectively in CF_3I. More concretely, consider a Teas diagram (Figure 4.1), which uses the polarity and the number of hydrogen bonds as "cohesion parameters" and plots relative solubilities, using proximity and sizes of dots to illustrate solubility effectiveness and time. It shows that substances very soluble in water should not be very soluble in refrigerants, which are similar to oils in terms of relevant properties. Thus, neutrino detectors with oil-based liquid scintillator like BOREXINO (Pallavicini et al., 2009) can easily achieve low backgrounds, purifying liquid that intrinsically is not very radioactive. Uranium, thorium, and their decay products (or, "daughters") can lead to an alpha background. This is one of the few significant backgrounds for a superheated liquid dark matter experiment (Section 2.5.2).

Because the superheated state is a metastable equilibrium, spontaneous homogeneous boiling caused by sufficiently strong random thermal fluctuations could initially be thought a problem. How often spontaneous boiling occurs in a superheated liquid depends on the properties of the particular liquid. However, it is possible to calculate that, due to the high critical point temperature of CF_3I (122°C), compared to the temperatures at which a COUPP bubble chamber is run (30–40°C), the probability of spontaneous boiling in CF_3I is negligible, even on timescales significantly longer than the age of the universe (Blander & Katz, 1975). Spontaneous boiling was therefore a non-issue. It cannot constitute an irreducible and unidentifiable background that could be mistaken for bubbles induced by WIMP recoils. This issue, however, refers specifically to homogenous spontaneous nucleation

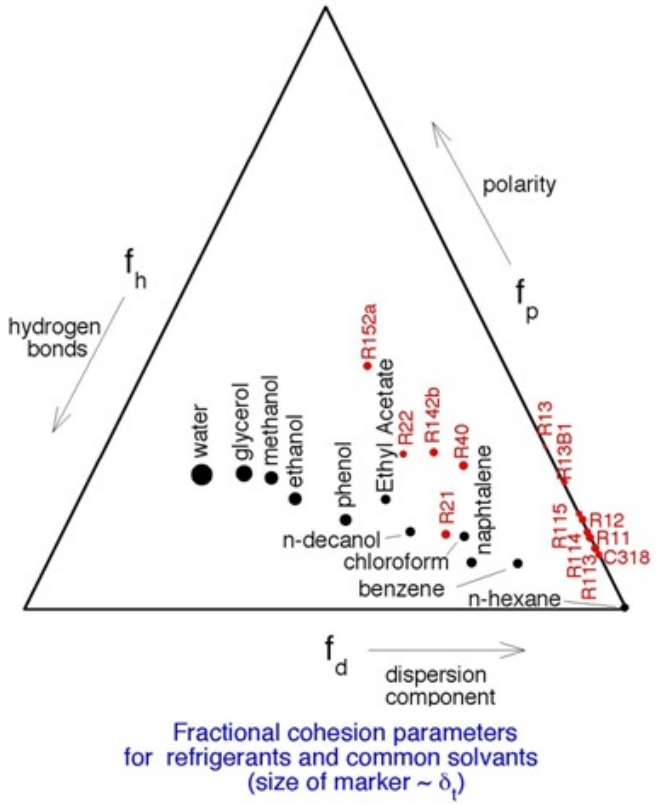

Figure 2.9: Teas diagram. Hydrogen bonding, the dispersion component (non-polarity), and polarity are three components of the Hildebrand parameter δ_t (Hildebrand, 1929). The Hildebrand parameter is related to the square root of the cohesive energy density, which is related to the enthalpy of vaporization. Solubility is similar to vaporization in the sense that atoms or molecules need to make space. The further substances are apart on the plot the less miscible they are. CF_3I is of the refrigerant family on the far right, indicated in red, contrasted with the water- and alcohol-soluble substances, in black, on the left. Note distance from water on left, implying that if uranium and thorium are easily dissolvable in water, they are thus far less readily dissolvable in CF_3I. The sizes of the dots, sometimes also depicted with different shapes, represent empirically-determined regions in the parameter space. Figure courtesy of J.I.Collar.

of bubbles in the bulk liquid, and not to inhomogeneous nucleation on the surface of the container, a separate problem addressed in Sections 1.5.2, 3.2, 3.3, and 4.1.1.

2.6.3 Disadvantageous Chemical Properties

CF_3I is an etchant gas. Being denser as a liquid only exacerbates this problem. Initial tests in the Collar lab culminating in Behnke et al. (2008) demonstrated its ability to react with brass, copper, and silver. It does not, however, react with gold, Teflon, stainless steel, or quartz. Early tests with acrylic (Collar, unpublished), a desirable material for a WIMP detector because of its low background, demonstrated the ability of CF_3I to react very strongly and dissolve it. However, no reaction with quartz has been observed in a COUPP chamber, even after years of running. These findings were consistent with Donnelly et al. (2004): long-term, high-temperature storage ascertained the stability of CF_3I in the presence of different materials. The use of CF_3I as a fire extinguisher (Section 2.6.2) motivated this intensive testing.

Another potential drawback to CF_3I is the weak C-I bond that makes it photodissociative. The cross-section for photodissociation peaks in the ultraviolet (Williamson, 1994), but sunlight (Solomon et al., 1994) and fluorescent room lights (Nyden, 1995) have sufficient fluence to compensate for their non-optimal wavelengths. The cross-section falls rapidly from UV to visible to infrared, but is always non-zero, so a large flux of any wavelength light still has potential to break the iodine bond. R&D efforts on early COUPP prototype chambers demonstrated that lighting vessels with red or infrared light ($\sim$600–1000 nm) avoids the formation of degradation products stemming from liberated iodine, on the years-long timescale of physics runs. Free iodine can discolor the CF_3I and make it pink or purple, resulting in poor illumination. The 15 kg chamber was therefore lit with LED clusters of 626 nm peak wavelength. Although infrared LEDs would be even safer for CF_3I, $\sim$900 nm LEDs were tested and rejected due to lack of power needed to make the chamber sufficiently well-illuminated to effectively trigger on the smallest bubbles at the lowest degrees of superheat (highest pressures). The power of readily commercially available and

inexpensive red multi-LED clusters is high on the other hand. This fact, combined with significantly lower sensitivity of the available Charge-Coupled Device (CCD) cameras at this wavelength, even equipped with infrared transparent lenses (Sony, 2010), on top of eye safety concerns, favored the use of red over infrared. The two-stage design of all COUPP chambers discussed in Sections 1.5.3 and 4.1 not only permits no great differential across the quartz vessel, but also allows the CF_3I to avoid exposure to room lights despite being enclosed in transparent material.

2.6.4 Target Fluid Replaceability

One of the great benefits of the bubble chamber technology is the fact that the target liquid can easily be replaced, as all liquids can in theory be superheated. All liquids for which that claim has been tested have been successfully superheated, though some more easily than others, depending on their critical point and allowed mechanisms of energy dispersion (Glaser, 1994). (For example, xenon makes a poor bubble chamber because energy can be lost through scintillation.) With the exception of the CLEAN experiment (McKinsey & Doyle, 2000), this ability to switch targets is unique in the field of direct detection, allowing rapid collection of additional corroborating empirical evidence in the case of a potential WIMP signal, leading to a more decisive claim of discovery (Bertone et al., 2007). Given the long search for WIMPs conducted by a decades-long series of direct and indirect measurements, including a continuing claim of discovery by DAMA and DAMA/LIBRA (Bernabei et al., 2010), this is an important factor in mitigating the inherently controversial nature of any discovery claim. For example, the DAMA results conflict with those of several other direct detection experiments (Akerib et al., 2004; Kim, 2008; Behnke et al., 2008), though a number of possible explanations remain.

CF_3I , C_3F_8, and C_4F_{10} all have comparable neutron scattering cross-sections but should have very different WIMP interaction cross-sections. The changing masses affect SI, and the

varying amounts of fluorine affect SD (Bertone et al., 2007). All of the refrigerants from the "family" to which CF_3I belongs have similar properties, especially vapor pressure and surface tension, implying a similar energy threshold as a function of pressure and temperature (see Section 2.3). This makes them excellent candidates to use with the same plumbing, pressure control systems, and other infrastructure as that of a CF_3I detector. Furthermore, repeating an experiment like COUPP with another liquid can aid in the differentiation of WIMP candidates (supersymmetric versus Kaluza-Klein partners, for instance), by helping to more soundly pinpoint the WIMP in the phase space of masses and cross-sections with orthogonal cuts (Bertone et al., 2007). Thus, instead of just a discovery, an actual analysis of WIMP properties becomes possible.

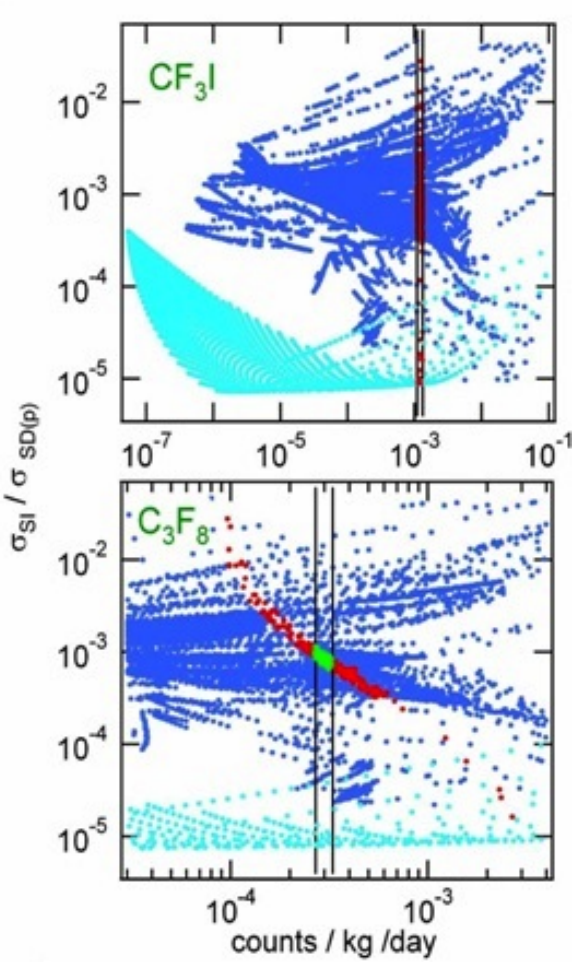

Figure 2.10: Ratio of SI to SD-p(roton) cross-section vs. interaction rate observed for potential WIMPs, in CF_3I and one its potential replacements for a bubble chamber, C_3F_8. The blue points are supersymmetric WIMPs, and the cyan Kaluza-Klein (see Section 1.3). Different masses and fluorine content significantly alter the observable WIMP signal: note the ability to pinpoint WIMP parameters in phase space if a rate is measured with both fluids in separate experiments. Red dots indicate WIMP models allowed for an example theoretical rate in each respective fluid, while green dots indicate the much more tightly constrained WIMP candidates allowed after repeating a CF_3I experiment with C_3F_8. Figure adapted from Bertone et al. (2007).

CHAPTER 3
UNIQUE FEATURES OF THE 15 KG BUBBLE CHAMBER

The 15 kg bubble chamber presented in this thesis is technically innovative not only because it is windowless, but also because the nature of its lighting scheme for bubble viewing differs from past and present bubble chambers: bottom lighting is used in lieu of backlighting. LEDs were placed at the bottom of the external steel vessel housing the quartz jar to light up bubbles, six clusters in all in order to provide redundancy due to inaccessibility. The 15 kg chamber was also the first to demonstrate the radiopurity-driven desirability of using Suprasil synthetic quartz over natural quartz for the vessel containing the active liquid, leading to the rejection of natural quartz for use in any next-generation COUPP bubble chambers.

3.1 Bottom Lighting System

Previous COUPP bubble chambers have been backlit (Behnke et al., 2008; Glaser, 1994; Glaser & Raim, 1955; Bolte et al., 2005), their bubbles appearing as dark circles on a white background. Different lighting schemes were explored for the 15 kg chamber in an attempt to find one superior to the standard backlighting. A test bed was constructed consisting of a beaker full of water into which bubbles would be introduced by depressing a syringe inserted into a copper tube ending in the water, as seen in Figure 3.1. Bubbles in the beaker were observed with a camera, while a small ring of LEDs was moved around to change the angle of the lighting. Bubbles appeared clearly as dark rings when backlit (180-degree lighting), but only with the help of a diffuser screen. Just like in the 2 kg chamber, a diffuser helped avoid light and dark spots caused by the discrete location of the LEDs and camera saturation caused by them being pointed directly into the light. Lighting from 90 degrees, by placing

LEDs below the beaker, above it, both above and below, or on the side, produced bright white filled-in circles of bubbles in camera images. Intermediate angles produced mixed results, bubbles appearing fuzzy and neither completely light nor completely dark, with bubbles completely vanishing at odd multiples of 45 degrees. Lastly, with a reflector, zero degrees (lighting coming from behind the camera) becomes equivalent to 180.

Ninety-degree lighting was selected as a promising alternative to traditional lighting angle. Side lighting was rejected because it would have called for the creation of a long line of LEDs along the side of the 26.5 in. 15 kg chamber, and because of the impossibility of having side lights simultaneously be 90 degrees away from multiple cameras in an arc around the chamber. Top lighting would prove difficult, requiring LEDs to be securely attached above the chamber somehow without risk of falling and obscuring the view. (A detached piezo partially obscured the view of the 2 kg chamber at one point.) Top lights could not be exactly at 90 degrees, but would have to be placed off axis because they could not be submerged in the active volume. Ultimately, bottom lighting was the clear choice, but it had to be paired with an optimal background. A white background was rejected because of the difficulty of seeing white against white, although for very large bubbles very close to the camera it was remotely feasible. Gray was rejected in favor of a very dark black, because the signal-to-noise ratio would increase with less reflection from the background. A woolly, non-reflective, deep black material was chosen as the ideal background.

Bottom, 90-degree lighting of bubbles with a very black background generates circular collections of bright pixels on which to trigger. Zero- or 180-degree light intrinsically provides its own (white) background, the light source itself. As light directly penetrates the transparent bubbles, they appear as dark annuli. Although in both cases bubbles are well-contrasted (white against black, or black against white) a few more triggerable pixels exist in the former scheme because the circles of differentiated pixels are filled. In hindsight,

this is a rediscovery of the cloud chamber lighting method (Das Gupta & Ghosh, 1946).

3.2 Chamber Material Selection

The original 2 kg chamber had approximately 20% dead time caused by quartz surface boiling (Behnke et al., 2008). It is possible to reject such events by using 3-D reconstruction of event positions to make a fiducial volume cut, removing the bubbles coming directly off the walls of the quartz jar. Nevertheless, every time a bubble occurs the chamber must be compressed to avoid runaway boiling. The pressure and duration of compression must be sufficient. Otherwise, pockets of supercooled gas that failed to recondense can cause loss of detector mass or lead to what is known as parasitic bubbling. This is immediate bubbling, often from the same location, which occurs upon the next attempt at an adiabatic lowering of pressure below the vapor pressure (Bolte et al., 2005). Based on the wall event rate of $\sim$0.8 events/cm^2/day observed in the 2 kg vessel (Behnke et al., 2008), the rate would have to be reduced by at least an order of magnitude in order to build chambers of mass greater than 50 kg with livetimes greater than 50% (the rate scales with surface area). The onset of the wall events corresponded with the onset of bulk alphas in threshold, leading to the conclusion that alpha emitters intrinsic to the quartz were to blame. An experiment was devised to determine the radioisotopic content of the quartz used by COUPP, which was General Electric quartz (GE-214) as distributed by Technical Glass Products (TGP).

The detailed procedure of this experiment is covered in Section 3.3.1. The result was that uranium and thorium were present in sufficient quantities ($\sim$50 ppb and $\sim$30 ppb, respectively) to explain the observed rate. Therefore, synthetic fused silica quartz, Suprasil from Heraeus and T-6040 from Covalent, with measured concentrations of 21 ppt and of $<$100 ppt uranium, respectively, three to four orders of magnitude lower than in natural fused silica quartz (Leonard et al. (2008) and ref. therein), was procured. The 15 and 4 kg

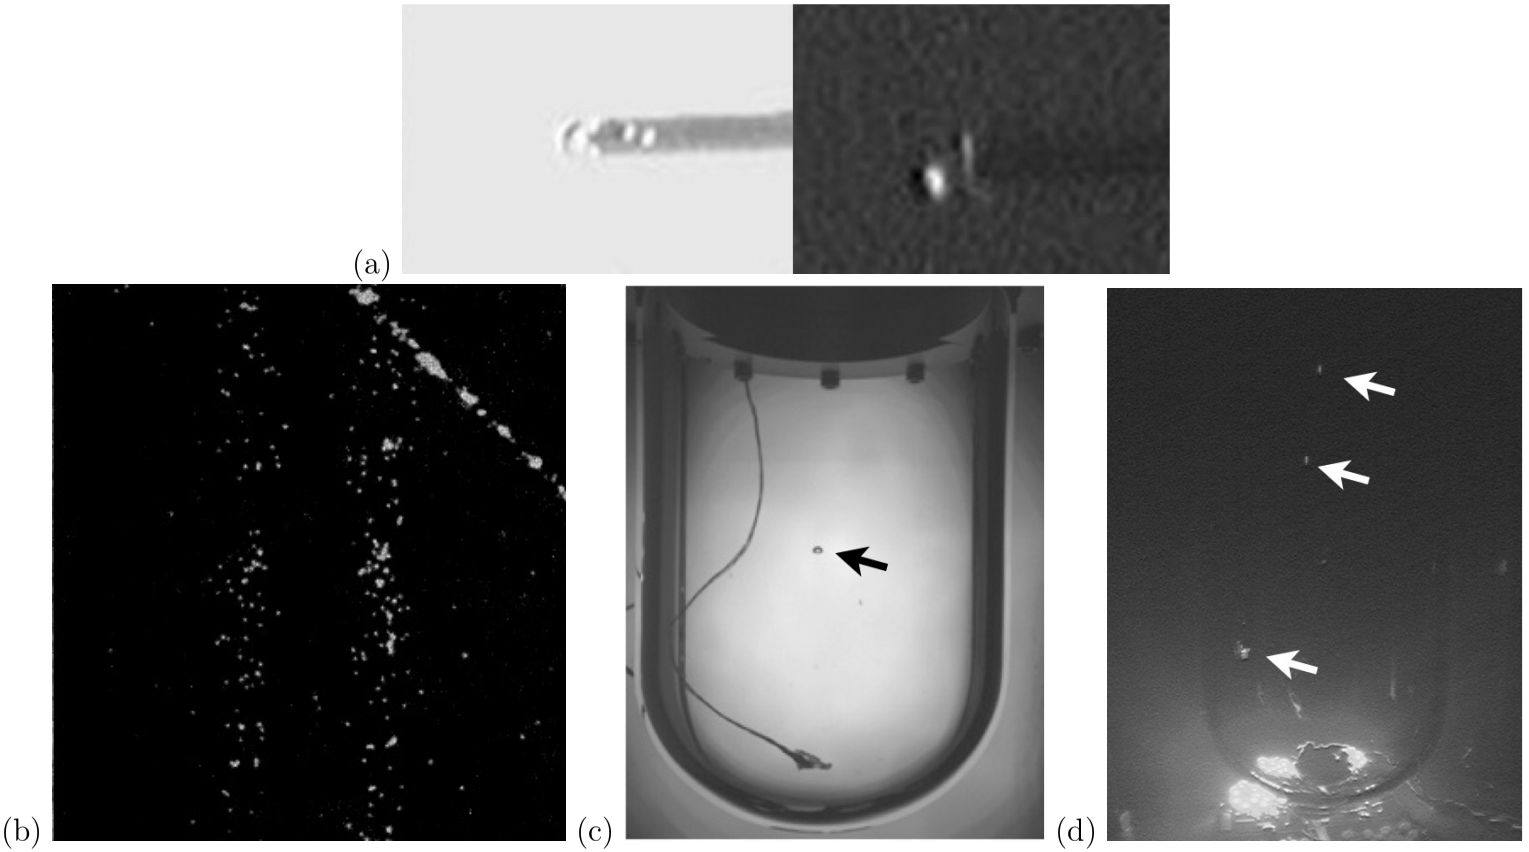

Figure 3.1: (a) Juxtaposition of the old (left) versus new (right) lighting systems for bubbles, 180-degree lighting against white (the light itself) and 90-degree lighting against a black backdrop, respectively. The old system offers only an annulus of pixels for possible detection, whereas the new system provides a more solid circle of pixels on which to trigger.
(b) Cloud chamber tracks, for comparison with (d). The "new" system for the 15 kg chamber was therefore a rediscovery of an older lighting scheme. Figure from Das Gupta & Ghosh (1946).
(c) COUPP 2 kg chamber (backlit). Image courtesy of FNAL.
(d) Three bubble event in COUPP 15 kg chamber (bottom-lit). Observe the stark contrast between bubbles and background, evident for (d) just as effectively as for (c), even at small resolution.

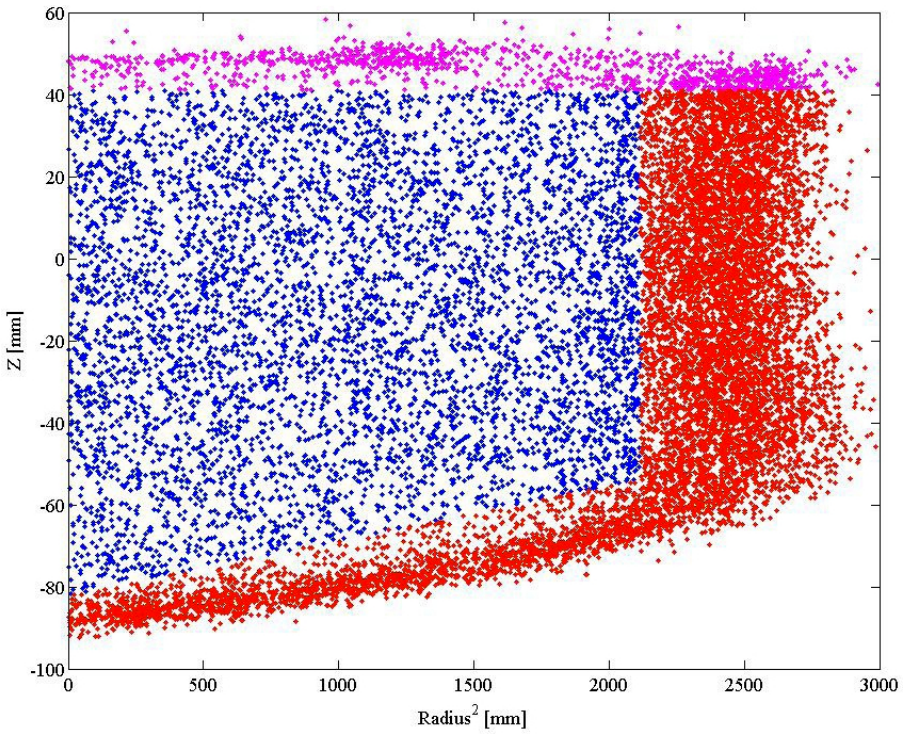

Figure 3.2: Bubble position reconstruction from the 2 kg chamber illustrating the problem of high-density wall events (red). Events in the bulk (blue) are far less dense. Events at the buffer liquid (water) interface are shown in pink (briefly elaborated upon in Section 6.5). Figure courtesy of A. Sonnenschein.

vessels were fashioned from the former synthetic quartz by TGP and the 60 kg vessel from the latter to explore multiple possibilities. Synthetic quartz is essentially produced through combination then condensation and vapor deposition of silicon and oxygen gas (Otsuka, 2007), instead of from melting together (fusing) naturally-occurring quartz crystals or silica sands. Because of this gaseous method, synthetic quartz should contain fewer metallic impurities (Quarzglas, 2010), the source of radioactive elements such as uranium and thorium. The 15 kg chamber data indicated a reduced wall rate (Section 3.3), demonstrating the effectiveness of switching to synthetic quartz.

3.3 The Quartz Bell Jar

The target liquid for a dark matter bubble chamber must be placed into an appropriate vessel. One important detail in chamber stability is the quality of the surface of the container holding the liquid to be superheated, illustrated by the well-known phenomenon of explosive boiling of water superheated in smooth mugs in a microwave. Smoothness is one relevant measure of quality. It is a question of whether there are features of order of the critical radius for bubble formation, that is, the minimum radius to which a "protobubble" must grow in order to gain the ability to grow indefinitely (Blander & Katz, 1975). This issue is avoided when water is used as a surfactant (Sections 4.1.1 and 4.5.2). Before any chamber begins taking "physics-quality" data, the walls of the quartz jar are allowed to outgas by cycling the pressure above and below the vapor pressure as though running the chamber, for a few hours or days until the rate of wall nucleation asymptotes. Any trapped CF_3I should vacate the voids to be replaced with water (Reinke, 1997, 1996). Then the released CF_3I liquefies and joins the bulk during the compression phase. Any impurities, i.e., non-CF_3I gasses unintentionally dragged in with the CF_3I during filling, should escape (Carey, 1992) to the top of the bellows in order to form a small gas pocket that should have little effect on the system, simply changing the compressibility slightly. For a further review of all these considerations, the interested reader should consult Bolte et al. (2005).

Another factor in the quality of a surface is the intrinsic radiopurity. Alphas or nuclei recoiling from alpha emission are a problem. Just like surface inhomogeneities, they cause wall nucleation and cripple livetime percentage (Zacek, 1994; Bolte et al., 2005). The high wall rate of the 2 kg chamber motivated the determination of the radiopurity of the natural fused silica quartz comprising it.

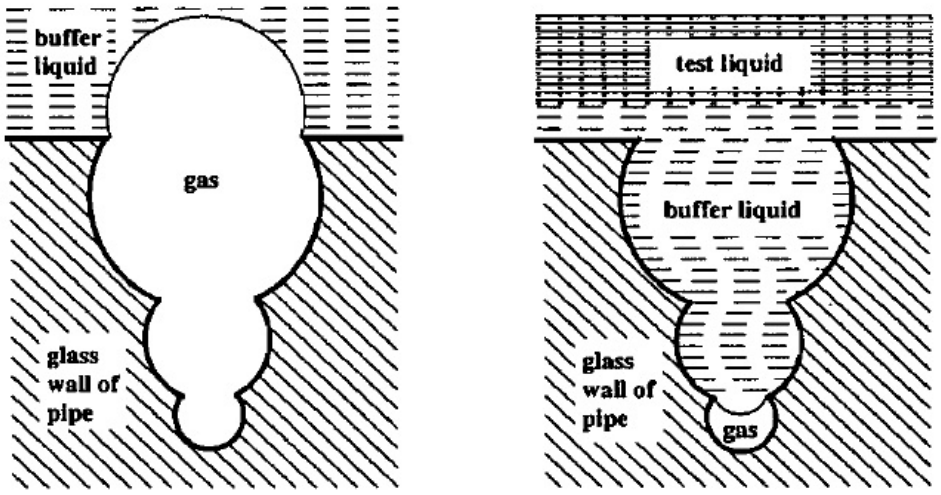

Figure 3.3: Drawing demonstrating the principle behind using a buffer fluid in a superheated detector. After non-condensable gas pockets are depleted with degassing (left), the buffer should fill in the voids, and relegate any remaining gas to an unreachable spot (right): a good buffer is immiscible and not close to boiling under operating conditions. Figure from Reinke (1997).

3.3.1 Radioisotopic Content of Natural Quartz

To determine the uranium and thorium contamination levels in TGP General Electric (GE) quartz, a commercial low-background high-purity germanium detector (HPGe) was used to count gammas from the ^{238}U and ^{232}Th decay chains, following the procedure of (Busto et al., 2002). Even though it was the alphas that COUPP was interested in, it is easier to use an HPGe detector, counting gammas, to determine radioisotopic composition. Radioisotopes from these decay chains have both alpha and gamma emissions.

An ultrasound-cleaned acrylic cylinder held shards of TGP quartz. It was placed in front of the HPGe, inside of a surrounding lead castle. Runs were an integral number of days to ensure a proper handle on any possible diurnal fluctuations in the radon concentrations in the air of the lab, unlikely as that would be due to ventilation and the low levels monitored in the lab. Data was taken both with and without the shards to make background subtraction possible. Several gamma-ray lines from different isotopes from both the uranium

and thorium chains were considered. Values for the branching ratio were all taken from an online table of isotopes at (Chu et al., 1999). The efficiency for full energy deposition of these gammas in the germanium crystal was determined with an MCNP (Monte Carlo Neutral Particle transport code) simulation. Table 3.1 lists the peaks found and corresponding mass fractions of uranium and thorium calculated assuming that the decay chains were in equilibrium within the quartz. Figure 3.4 shows a sample peak.

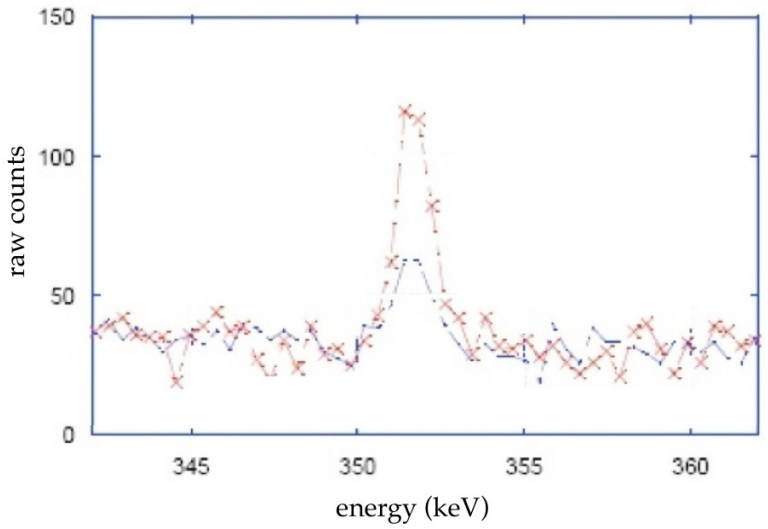

Figure 3.4: Example gamma rays line from a daughter isotope used to determine the concentration of uranium-238 within ordinary quartz: 352 keV lead peak. The data are summed from multiple runs with quartz shards alternating being in front of the HPGe (red) and removed (blue). Note clear excess in presence of quartz.

Based on the sum of the counts under six peaks, from bismuth and lead, the weighted average for the ^{238}U concentration was determined to be 52 $\pm$ 2 ppb. The result for ^{232}Th was 32 $\pm$5 ppb. (See Table 3.1.)

It is possible to estimate what these concentrations imply for the wall rate. The range of a 5 MeV alpha in quartz is 22 microns (Ziegler, 2010). After 19 microns, it still carries approximately 1 MeV, more than enough to produce nucleation in all 2 kg chamber running conditions. An alpha escaping the quartz should produce a wall event. The thinness of the

HPGe Gamma Spectroscopy Results										
Chain	Isotope	En. (keV)	B (counts)	S (counts)	S-B	Err in S-B	efficiency	bran ratio	ppb U	ppb Th
^{232}Th	^{212}Pb	238	184	417	233	25	0.01249	0.433		36 ± 4
^{238}U	^{214}Pb	295	234	604	369	29	0.01015	0.193	52 ± 4	
^{238}U	^{214}Pb	352	411	1006	595	38	0.00846	0.376	52 ± 3	
^{232}Th	^{208}Tl	583	87	146	59	15	0.00523	0.845		31 ± 8
^{238}U	^{214}Bi	609	385	862	478	35	0.00506	0.461	56 ± 4	
^{232}Th	^{228}Ac	911	40	69	29	10	0.00370	0.258		25 ± 9
^{238}U	^{214}Bi	1120	124	201	77	18	0.00308	0.151	46 ± 11	
^{238}U	^{214}Bi	1238	23	56	33	9	0.00284	0.058	55 ± 15	
^{238}U	^{214}Bi	1764	99	146	47	16	0.00212	0.154	40 ± 13	

Table 3.1: List of all lines observed above background and used for gamma spectroscopy. B is background (no quartz) counts; S is background with signal on top. These counts get divided by the cumulative livetime (7.5 days with quartz and 7.5 without) to produce rates. Gaussian error in count rates is assumed. Moreover, equilibrium is assumed along the U and Th decay chains. Efficiencies for full gamma energy deposition in the Ge detector (as determined by simulation) and branching ratios are quoted as fractions.

quartz permits a linear trajectory approximation. The solid angle fraction must be calculated: the geometric weighting factor is 1/6 (Collar et al., COUPP collaboration internal note). For a 10 cm high, 5 cm radius cylinder, the approximate dimensions of the 2 kg quartz vessel, the "crust" responsible for the wall rate is then 0.25 g of quartz. Assuming secular equilibrium along the entire chain, ^{238}U yields 9.6 alphas/(day×g×ppb) (Gratta et al., 1998). Therefore the 2 kg chamber should have had ∼0.4 wall events per cm^2 per day from the uranium chain alone. The measured amount of uranium could also have accounted for part of the bulk alpha rate, leading to a radon emanation rate of ∼50 radon atoms per day (Collar et al., COUPP collaboration internal note).

Although the wall event estimate does not account for the full 0.8/(cm^2×day) measured in the 2 kg chamber, it shows how a crude approximation yields the right order of magnitude, underscoring the high plausibility of accounting for the wall rate with uranium content. A full Monte Carlo simulation was needed to account for the originating depths of alphas of varying energy, as low as 4.2 MeV along the ^{238}U chain. A simulation written by J.I. Collar demonstrated that a 52 ppb U concentration would yield a 0.76/(cm^2-day) wall rate, in excellent agreement with the observation of 0.8/(cm^2-day) (Figure 3.5). The contribution of thorium to the wall rate was an order of magnitude lower because its decay series has half as many alpha emitters and half the specific activity as that of uranium, and because the concentration in pbb was nearly half.

Attempts to reproduce the results from the gamma-ray spectroscopy with ZnS scintillation screens and PMTs, and with Eurysis alpha spectrometers, were inconclusive because of poor statistics. In the case of the alpha spectrometers, the samples could not touch the spectrometer for fear of damaging or contaminating a thin low-background membrane, so they had to be suspended. Given the short mean free path of alphas in air (order mm) this could have prevented enough alphas from reaching. However, alpha counting performed by

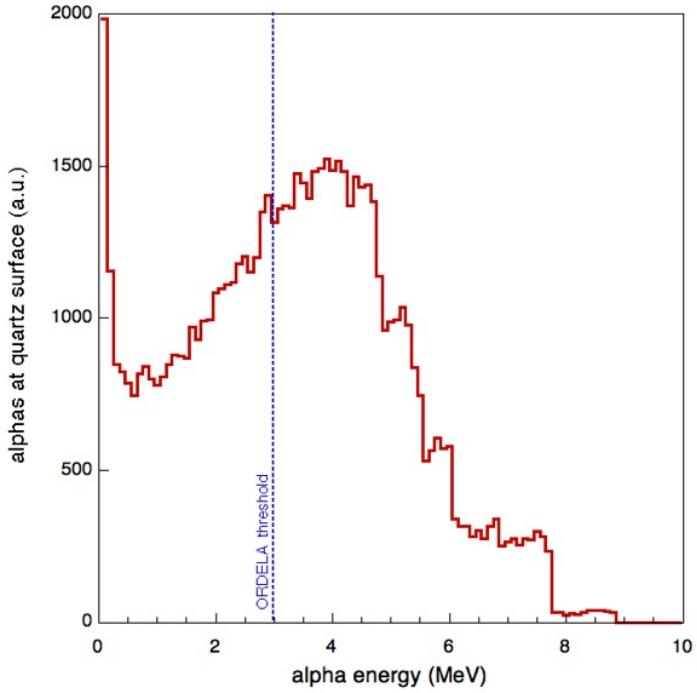

Figure 3.5: Simulation by J.I. Collar of energy spectrum of alphas exiting a natural quartz surface. The concentrations of ^{238}U and ^{232}Th are taken to be 52 and 32 ppb respectively, based on the gamma measurements, and equilibrium is assumed. Initial alphas are transported through the quartz in the simulation using the correct branching ratio and relative yields for the two chains. Based on the simulation, 0.76 alphas/(cm^2-day) were expected in the bubble chamber (12.2% of the alphas emitted successfully leave the surface). After properly correcting for detection efficiency and energy threshold (blue dashed line) the measurement performed on behalf of COUPP by ORDELA is consistent. Figure courtesy of J.I. Collar.

ORDELA agreed with the gamma spectroscopy. The final test was the running of the 15 kg chamber made with synthetically produced fused silica.

3.3.2 Advantage of Using Synthetic Quartz

When the 15 kg chamber was first turned on, it was already evident based on the overall rate (apart from an issue with the hemisphere, Section 6.4) that the wall rate was lower by at least a factor of 5–10. A visual inspection of bubbles in the cameras also showed a class of events mushroom-shaped instead of spherical like normal bubbles. In frame by frame inspection of

the bubbles' evolution, these bubbles appeared to be emanating from points along the walls. Just as for the 2 kg chamber, hand scanning proved a useful check on automated bubble finding and position reconstruction for the 15 kg chamber. A careful visual inspection of 255 kg-days of data yielded 75 candidate wall events. Given the surface area of the 15 kg chamber, this translates into $\sim$0.002 events/(cm$^2\times$day), or, $\sim$100 ppt U. The computerized three-dimensional reconstruction of bubble nucleation sites sheds more light on the issue (Section 6.5.2). The human eye is prone to bias, but because the brain excels at pattern recognition, $\sim 10^{-3}$/(cm$^2\times$day) should be taken as a good first-order estimate. The 4 kg chamber corroborated this result; it was the successor chamber made from the leftover quartz used to fashion the 15 kg vessel. A rate of $\sim 10^{-4}$–10^{-3} wall events/(cm$^2\times$day) was observed in the 4 kg chamber (Behnke et al., 2011).

Figure 3.6: (top) Examples of bona fide bulk events in the 15 kg vessel at various stages of development. Observe approximate sphericity and coherence of shape.
(bottom) Juxtaposed examples of likely wall events at various stages of development, "mushrooming" off the wall into a growing cloud, in stark contrast to bulk event growth pattern. Three cameras were needed to view the entire length of the chamber, hence the different perspectives. Different bubbles at different stages of development in different camera views are depicted here to provide a sampling of the broad range of bubble shapes in the 15 kg chamber.

CHAPTER 4
PREPARATION AND DEPLOYMENT OF THE 15 KG BUBBLE CHAMBER

Details of chamber parts and operation are included in this chapter, building upon the overview of the technology, its physics, and the aforementioned specific characteristics of the 15 kg bubble chamber. (The instrumentation and data acquisition system are covered in Chapter 5.) Although the 15 kg chamber is described here, most features are general to all COUPP chambers, if not to all bubble chambers. After an engineering run above ground in the high bay of the Accelerator building at the University of Chicago using natural quartz to debug both the software and hardware, the 15 kg bubble chamber was deployed at 330 m.w.e. underground at the Mainstream Pumping Station (MPS) in Hodgkins, Illinois, part of the Tunnel and Reservoir Plan (TARP) of the Metropolitan Water Reclamation District of Greater Chicagoland (MWRDGC).

4.1 Principle of Bubble Chamber Operation

The COUPP bubble chamber concept involves placing the vessel with the active liquid inside a larger vessel filled with a passive hydraulic fluid. The temperature is kept fixed, and a piston is used to control the pressure. The piston uniquely determines the nuclear recoil energy threshold, itself a function of temperature and pressure. It is lowered with higher temperature as the boiling point becomes further away, but raised with higher pressure as the vapor pressure is approached. A piston is also used to quickly raise the pressure of a chamber above the vapor pressure when a bubble is detected, to force the bubble to return to the liquid phase and allow the chamber to be ready for the next event. This process is known as (re)compression or (re)pressurization, followed by decompression or expansion in the pressure cycle, when the pressure is slowly lowered until the fluid returns

to a superheated state. Pistons pressurize the active liquid indirectly, compressing the hydraulic fluid: propylene glycol in the case of most COUPP chambers, and mineral oil in the case of the 15 kg chamber because of the need for an electrically non-conductive fluid in which cameras and LEDs could reside. The hydraulic fluid in turn pushes on a bellows which is located atop a quartz bell jar comprising the active volume.

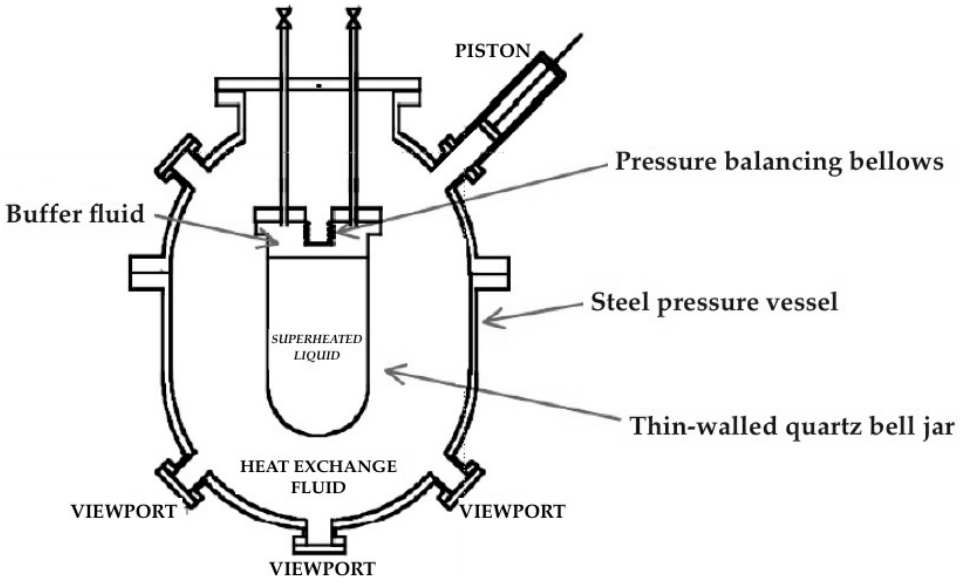

Figure 4.1: General diagram, approximately to scale for the 2 kg chamber, illustrating the basic principle behind every COUPP bubble chamber, as explained in Section 4.1. A piston presses upon the fluid in the pressure vessel (glycol or oil), which presses on the bellows, which in turn presses on the buffer fluid above the CF_3I (water), which then presses on the CF_3I itself in the quartz jar. The 15 kg chamber lacks viewports; a diagram is found in Figure 4.3. Figure courtesy of A. Sonnenschein.

4.1.1 Quartz and Support Structures

The inner vessel for the active liquid, CF_3I, described in Chapter 2, was a quartz jar with 130 mm inner diameter, 135 mm outer diameter, 26.5 in. (67.31 cm) in overall length, with a 3 cm lip and a hemispherical bottom. It was fabricated from synthetic quartz. Like all previous chambers, the 15 kg chamber was pre-filled via evaporation and condensation with

a small volume of degassed water before being filled with CF_3I. This was done in order to let the water act as a surfactant (Reinke, 1996, 1997; Bolte et al., 2005), filling in any microscopic voids in the quartz surface that could act as bubble nucleation (formation) sites by trapping pockets of gas. The water also functions as filler for the bellows volume, where bubbles cannot be seen because steel is opaque. Sources of bubble nucleation other than external particle interaction must be removed for a bubble chamber to reduce background. Minute imperfections in the wall of the container can be nucleation sites, (Sections 1.5.2 and 4.5.2). The chances of dust motes traversing the bulk volume serving as such sites are minimized via the initial cleaning of the quartz vessel with a high-powered ultrasound bath filled with distilled water (Section 4.5.1). Removing dust not only removes macroscopic inhomogeneities in the liquid, but a source of alphas as well, dust being rich in uranium and thorium at the ppm level (UNSCEAR, 2000).

The quartz vessel hung from an aluminum plate, which was in turn suspended inside the oil volume supported by an equilateral triangle of rods threaded into holes in the bottom of the top flange of the outer vessel. Nuts above and below the holes where the rods passed through the plate held the plate fixed. The same rods supported another plate, on which the LEDs were mounted, beneath the quartz jar. The entire structure of the inner volume could be lifted in and out of the outer vessel by attaching a hoist to three hooks affixed to the top flange (see Figure 4.2). A gantry was used (see Figure 4.5) for this purpose during installation and decommissioning.

4.1.2 The Pressure-Balancing Bellows

In all chambers a stainless steel bellows ensures that the thin-walled (2.5 mm thickness for the 15 kg chamber) quartz vessel never experiences a significant pressure differential which might result in explosion or implosion. Quartz is theoretically capable of surviving differential pressures even in excess of 200 psig, despite its thinness (Borodai, 1989). However, the

manufacturer, Technical Glass Products (TGP), only guaranteed survival of implosive pressure. The pressure was not to exceed one atmosphere, or 0 psig (14.7 psia). Thus a bellows was used to keep the inner CF_3I volume and outer oil volume in equilibrium at all times during chamber operation. Because the bellows must flex as the chamber pressure changes, it is a ribbed, accordion-like structure that is essentially a large spring, at equilibrium when the differential pressure across it is zero.

Where the bellows assembly was mated to the quartz vessel, five concentric rings of Teflon-encapsulated viton O-rings were used to make the seal separating the active volume from the surrounding oil. Viton rubber is a common O-ring material good for pressure seals, but it is also known to emanate radon (v.Hentig et al., 1999; Cadonati, 2001). In fact, viton was one possible source of the radon in the 2 kg chamber (Behnke et al., 2008). Therefore, the Teflon encapsulation served to seal the viton, and the redundant rings served to make a good seal against the oil, as well as against any radon in that oil or the surroundings which could diffuse into the CF_3I (Cadonati, 2001). Emanation measurements performed at SNOLAB indicated an upper limit of 1.6 radon atoms released into the bubble chamber per day from these Teflon-coated O-rings (Odom, 2010). At the top of the bellows, there was a capillary coil, only 1/8 in. in diameter and made of steel, which connected the bellows to the outer world through a feedthrough in the top flange. It was very flexible, permitting the bellows its requisite freedom of motion. The filling of the chamber was performed via this thin tube (Section 4.5).

4.1.3 The Outer Pressure Vessel

The outer vessel, filled with oil or glycol, is a steel pressure vessel capable of withstanding the pressures to which a chamber is compressed to condense bubbles. In the 15 kg chamber, this could be as high as approximately 200 psig, chosen to be well above the vapor pressure of the active liquid to minimize bubble recondensation time. The 15 kg chamber's outer

(a) 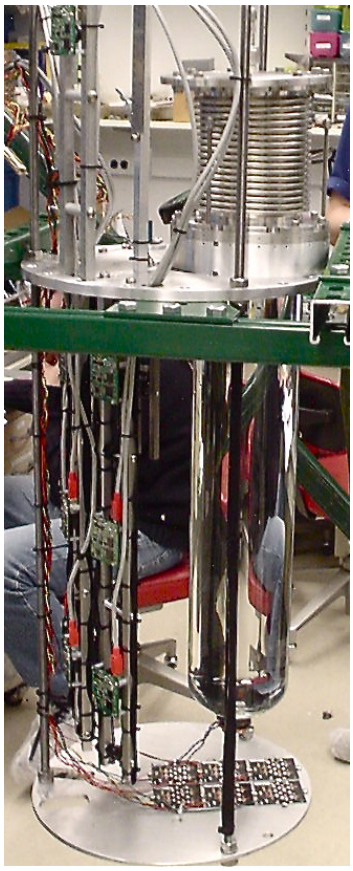 (b) 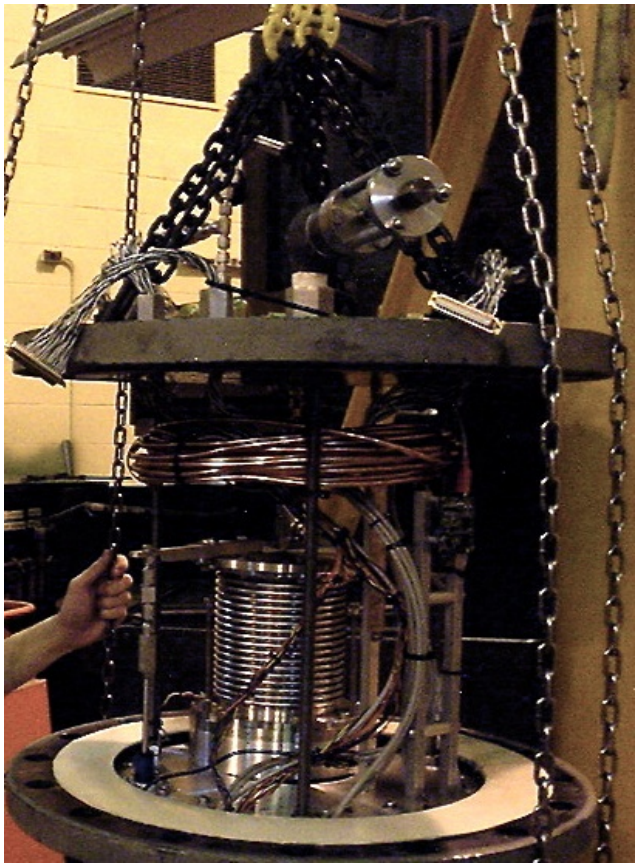

Figure 4.2: (a) Quartz vessel and supporting assembly. Note aluminum plates for supporting the quartz vessel in the middle and for mounting the LEDs at bottom (Section 4.1.1). The visible camera system for optical bubble inspection, visible as left, is discussed in Section 5.3.
(b) Demonstration of the lowering of (a) into the outer steel pressure vessel by hoist for chamber commissioning after attachment to the top flange with three threaded rods. More easily visible are the pressure-equalizing bellows and flexible capillary coil fill tube, as well as copper heat exchange coils for thermal regulation (Sections 4.1.3 and 4.5.4) and bellows camera for position monitoring (Section 4.4.2).

vessel was a VAL-FAB® low-carbon steel 47 in. tall cylinder with a 16 in. outer diameter and ∼15 in. inner diameter. JB Weld, an epoxy capable of resisting temperature, pressure, and oil, was chosen to affix the non-reflective black background necessary for the lighting system (see Section 3.1) to the inside wall. The top and bottom flanges of the outer vessel were sealed with 1/8 in. thick rubber gaskets.

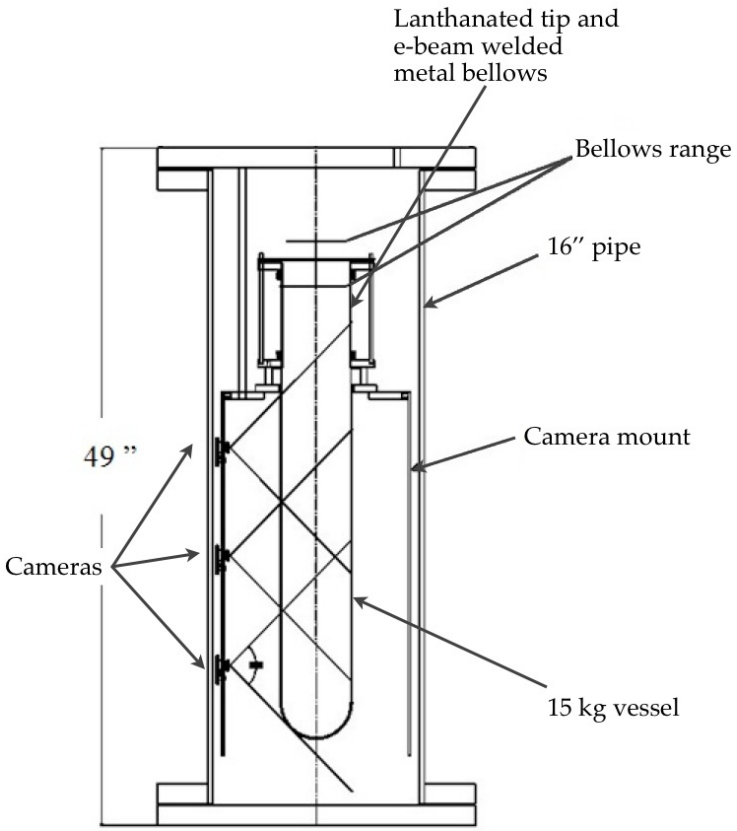

Figure 4.3: Scale diagram of the 15 kg chamber representing the quartz vessel and the bellows inside of the outer steel vessel, a pressure-stamped industrial steam pipe. Here, 16” refers to the diameter. An array of cameras was needed because of the length of the chamber, seen mounted on the left in this diagram. Camera mounts were rectangular aluminum structures, shown in Figures 4.2 and 5.1. The lines emanating from the cameras indicate their approximate field of view. The final position of the quartz jar was off center in order to provide a more generous overlap between cameras (Section 6.1). Figure adapted from B. Odom.

4.1.4 Temperature Control System

Originally, resistive heat pads, coarsely controlled by a rheostat and attached to the steel outer vessel underneath the thermal insulation, were intended to control the temperature. For fine temperature control, a relay controlled by software pulsed current through the pads for time periods determined using the Proportional-Integral-Derivative (PID) algorithm, to maintain the desired temperature setpoint. However, the heat pads, despite uniformly

covering the external area of the pressure vessel, introduced a significant thermal gradient across the inner volume. It is best to minimize the difference in temperature across the detector to ensure a uniform energy threshold, simplifying the data analysis. To ensure proper convection and uniformity of temperature in the tall pressure vessel, an intelligent temperature-regulating pump (built-in PID) circulated heat exchange fluid (ethylene glycol/water mixture) close to the setpoint through copper tubes located in the inner oil volume of the steel vessel near the bottom. The tubes were accessed through feedthroughs in the bottom flange. The setpoint of the pump would be offset from the desired setpoint of the detector by 1–2°C, an empirically determined necessity for overcoming thermal losses in the several feet of tubing leading up to the vessel.

4.1.5 Pressure Control System

The pressure is also regulated in order to have a precise knowledge of the recoil threshold for analysis. The pressure of the 15 kg chamber was varied in fixed steps to build a spectrum of bubble rates versus threshold, providing a better understanding of any backgrounds and seeking a WIMP-like (exponential) recoil spectrum above background. This was similar to how the 2 kg chamber operated, where the pressure was changed continuously (Behnke et al., 2008), but different than the 4 kg chamber (Behnke et al., 2011) which stayed at a fixed pressure. This makes the 15 kg chamber data set orthogonal to the others in terms of information gleaned. Pressure variation is also useful for energy threshold calibration with a neutron source. The spectrum from the detector and the simulation are compared (Chapter 8).

The 15 kg chamber was the first to test a modularly designed pressure control system, developed by Fermilab engineers, swappable between chambers, equipped with a Programmable Logic Controller (PLC) for communication with a computer. The system consists of a cart on wheels with the various parts mounted onto a stainless steel framework. The

cart contains two pistons, a so-called fast piston for compression operated by compressed air, and a slow piston operated by a motor for fine pressure control. A valve exhausted air in 2–3 seconds, slow enough to still be an adiabatic decompression, empirically determined for COUPP bubble chambers based on past work (Behnke et al., 2008). Decompressing too fast results in non-radiation-induced boiling because of the nature of superheated liquids (Section 1.5). A solenoid with ∼10 ms activation time would recompress for 60 s. The compression time was empirically determined to avoid "parasitic" boiling (Section 3.2), but minimized to reduce dead time. The 15 kg chamber was unusable when trying a 30 second compression time, overwhelmed by inhomogeneous nucleation. Forty-five seconds was borderline, decompression becoming impossible after a few hours. The final 60 s compression time was longer than that needed by the 2 kg chamber (30 s), possibly due to longer tubing in the compression line and the greater compressibility of mineral oil compared to ethylene glycol (Ohno et al., 2005). These factors likely contributed to the ∼300 ms response time observed for the compression. A longer response time implies that a bubble can grow larger and lead to more gas ejection, making it more difficult to recondense.

Compressed air must be provided for the fast piston. A compressor is not included with the cart. A compressed gas bottle, or air compressor, must be connected to it. For the 15 kg chamber, an air compressor set to 125 psig was connected. The 125 psig was then doubled to 250 psig with a pressure doubler (also on the cart) and reduced to 45 psig with a gas regulator. This regulated air then enters a holding tank, which in turn leads to a solenoid valve which controls access to the air above the piston. The fast piston system contains a cylinder which shrinks the surface area, effectively quadrupling the pressure. The result should be a 180 psig compressive pressure on the bubble chamber. In practice, the piston was sticky and drifted to a lower position than intended, and the compression pressure could be as high as 190–195 psig. The justification for the multiple steps in pressure was

Figure 4.4: Modular PLC-run pressure cart designed and constructed by FNAL engineers. The first prototype of this design was installed, tested, and debugged on the 15 kg bubble chamber which is the focus of this thesis.

as follows. Air is lost whenever the chamber is decompressed through an exhaust valve, and so it must be replenished to maintain the same level of compression. To decrease the mechanical strain on the air compressor by having it fire less often, the highest possible pressure is set initially so that there is a reserve of air to bring the 45 psig region back up to pressure quickly. The need for the 250 to 45 psig stage stemmed from the fact the fast-acting solenoid valve needed for fast compression could not handle pressures exceeding 150 psig.

The slow piston has a turnscrew above it operated by a motor. The PLC turns the motor, moving the piston up and down, to maintain the pressure at a setpoint. In practice, the motor did not turn quickly enough to maintain the desired pressure for the 15 kg chamber,

as the rotation speed was not a steep enough function of the difference between current pressure and setpoint. To compensate, a system was developed where the actual setpoint was significantly lower than the effective setpoint if pressure was too high, and higher if the pressure was too low, set by computer via ethernet port. With this system, the desired pressure was maintained to within better than $\sim$0.1 psi.

It was necessary to actively maintain a desired pressure to control thermal drift changing the volume of the hydraulic fluid and to combat a phenomenon called "effervescence" which has been observed while operating bubble chambers since the beginning of COUPP (Bolte et al., 2005). If decompressed to a pressure below the vapor pressure and left alone, any chamber will experience a rise in pressure over time. The lower the pressure setpoint, the steeper this rise is. J.I.Collar clearly observed continuous effervescence on the bellows of a small chamber which had a camera below it. This explained the observed pressure rise: CF_3I in the water inside the bellows nucleates on the metallic ribs of its accordion-like structure. Although the water sits above the bulk of the CF_3I, the two fluids are not perfectly immiscible. A few mmol/L of CF_3I does get dissolved in the water, saturating it. Because of this, the pressure overcorrection for the 15 kg chamber discussed above was made intentionally more aggressive if the pressure was too high compared to when it was too low.

4.2 Shielding Layers Surrounding the Experiment

The 15 kg chamber's external pressure vessel was surrounded by cubitainers, cubic cardboard containers tightly accommodating 5 gal (20 L) bladders full of water, the remaining space in the cardboard cubes having been filled with polyethylene pellets. The cubitainers served as neutron moderator, their purpose being to lower the energy of incoming neutrons below threshold to reduce the background from environmental and cosmic-ray neutrons. They

served as roof, floor, and all 4 sides of a cubic structure held together with wood and a steel connector system (Unistrut). The entire structure was an approximate cube of side length 2.5 m. The muon veto and the mineral oil in the pressure vessel also served as hydrogen-rich neutron moderator. The total thickness of hydrogenated material surrounding the detector was 44.8 cm. Surrounding the pressure vessel was a 1 in. thick layer of polyurethane foam for thermal insulation. Immediately adjacent to that was the muon veto, which helped tag cosmic-ray induced events. Its purpose was to identify, not shield against them (Chapter 5). By using the pressure to change threshold, unlike the PICASSO experiment, a superheated liquid-based experiment which ramps temperature (Piro, 2010), the recoil threshold can be changed faster but still discretely. The problem of thermal inertia of large masses is avoided: pistons are fast-acting compared to heating and cooling.

4.3 Underground Deployment

The shielding and individual components including the the CF_3I cylinder, the outer vessel, degassing vessel, gantry components, vacuum pump, and chillers were all lowered by crane from the surface to a depth of 297 ft. A second crane completed the lowering to the final level at 325 ft., which was accessible only via crane or a tight spiral staircase. Discrete items like cubitainers and muon veto panels were lowered using pallets. Wheeled devices like the muon veto electronics rack and pressure cart were first transported by elevator and then by crane to the lowest level. The computer and small DAQ devices were carried down the stairs.

The overall dimensions of the cubitainer shield, the length of the degassing vessel, the size of the gantry beams, the lengths of the steel beams for the structure for holding all the cubitainers, and sizes of all other equipment were chosen to fit together within the confined space of the room available for the experiment, which was approximately 20 ft. $\times$ 20 ft. It

Figure 4.5: The cubic structure built for housing the 15 kg chamber, and gantry for inner volume installation and extraction. The white boxes are cubitainers used for neutron moderation. These contained bladders of water, surrounded by plastic pellets to fill in empty space (Section 4.2). When complete, the cubitainers constitute a 4π shield against terrestrial and cosmogenic neutrons. As shown, the structure is in process of assembly above ground as first test of design. From left to right: Chris Passalano, Aaron Plasek, Luke Goetzke.

was also necessary to consider what could reasonably fit through the hole between the -297 and -325-ft. levels or be carried down the constrictive stairs. The gantry was chosen so that it could be raised enough to have the bottom of the inner volume structure (LED plate) clear the top of the pressure vessel, so that the inner volume could be lowered in by hoist. However, the crane also could not be too tall for a 15 ft. high room. The inner volume plumbing was transported in a steel cage specially constructed for the purpose of safely maintaining the integrity of the quartz vessel. The entire inner volume structure, including bellows and aluminum support disks for the quartz jar and LEDs alike, all interconnected with rods threaded into the top flange, was held fixed for transport with steel beams. Air-filled bags protected the quartz vessel from motion during transport, while the entire steel

cage rested on a special pallet which incorporated shock-absorbing memory foam padding.

The TARP location, about halfway between FNAL and the University of Chicago and near Argonne National Laboratory, was logistically advantageous: it was not very difficult to fit all the necessary equipment and supplies down the available shafts and conduct them into the room. The choice of location was also motivated by physics considerations: the comparable depth at which the 2 and 4 kg chambers were deployed at FNAL served as a useful check for systematic effects. At this depth, the muon rate is two orders of magnitude lower than at the surface (Riley, 2007).

4.4 Safety Precautions

The 15 kg chamber, being a moderately-high pressure system, was fitted, both in hardware and software, with features to prevent mechanical failure resulting in personal injury or equipment damage. The following is a review of all such features, intended to make a self-sufficient chamber possible. An ideal chamber collects data for the longest possible period of time without any user intervention.

4.4.1 Pressure Relief Valve System

There were numerous pressure relief valves throughout the system to prevent overpressure. There was a relief valve on the oil volume set to 220 psig. However, as the lost oil would be difficult to restore to the vessel, and could cause pistons to run out of range, there were relief valves on the pneumatic side of the pressure system as well to avoid oil loss. There was a relief valve at 135 psig on the air compressor, another at 150 psig on the input side of the doubler, and a third at 150 psig on the holding tank. As an added precaution, the software rendered it impossible to set a setpoint higher than 210 psig, and data-taking would be terminated and the slow piston would be activated to resolve the pressure problem if a pressure above 215 psig was detected. However, even without these measures explosion was

not a great concern: the compressibility of mineral oil is such that a loss of ∼100 mL would be enough to bring the oil volume (35 gal) from 200 psig down to 14.7 psig, calculated based on Ohno et al. (2005).

4.4.2 Missed Bubbles and Bellows Position

Inadvertent superheating when the data acquisition system was not ready could lead to a missed bubble. Fortunately, the bellows had enough range to accommodate the gas pocket created by a runaway bubble. The inner volume would not boil away completely in this case. Instead, enough liquid would boil until the bellows would press on the surrounding oil sufficiently hard to raise the pressure above the vapor pressure. However, in the case of a piston being able to move to accommodate the bellows, or of the temperature being mismeasured, the bellows might run out of range due to boiling, thus overexpanding until damaging or rupturing the ribs. To prevent this situation, the bellows range was limited with nuts on a tripod of threaded rods. Nuts above the bellows upper flange prevented overextension, and another set below prevented overcontraction, which could have resulted from loss of CF_3I during a leak or from thermal contraction due to heating element failure.

The nuts prevented damage to the bellows, but they could also prevent the bellows from maintaining equilibrium between the inner (CF_3I) and outer (mineral oil) volumes. This would make implosive or explosive differentials across the quartz vessel possible. To avoid this risk, the bellows position was monitored: an LED and brightness sensor that output a voltage as a function of luminosity were installed on the upper and lower flanges of the bellows. Before the vessel was sealed, this monitoring system was calibrated, determining the voltage as a function of distance, i.e., the height of the bellows. The bellows was not permitted to come within 0.2 in. of the maximum possible extension or contraction as defined by the nuts. This system was adopted rather than another involving a linear transducer, an extending rod which linearly converted distance into voltage, because this

transducer would tilt and cripple the bellows, slowing compression. In event of the LED or brightness sensor failing, moving, or becoming detached, a camera monitoring the bellows as a backup indicator of the bellows position would become the primary system. The bellows would be lit with a cluster of LEDs for this purpose. (If that cluster of eighteen LEDs were to fail, there would still be enough residual light from the LEDs lighting the chamber below to just make out the rim of the bellows.)

4.4.3 Piston Positions and Range

The positions of the pistons were also monitored, using Hall Effect magnetic sensors. The slow piston position over long time periods could betray the presence of an oil or CF_3I leak if its average decreased significantly. It could also indicate changes in temperature if the thermometers failed, or indicate the degree to which effervescence was present. When the outer vessel was being filled with oil, setting the slow piston position helped control the oil level. Furthermore, when the chamber was first filled with CF_3I and brought up to operating temperature, the slow piston position was useful for monitoring the progress of thermal expansion. The fast piston position was important because if it were anomalously low during decompression, then data-taking would be halted and the slow piston used to compress. This was because the fast piston, depending on the decompression pressure, might need to travel nearly its entire length to fully compress the chamber.

In the event of the heating elements failing, the chamber might cool down, but inadvertently become superheated, due to pressure loss caused by thermal contractions. Because the temperature is monitored, a steadily declining temperature could be recorded by the computer and data-taking could be terminated. The slow piston, which has ten times the capacity of the fast piston, then took over compression as the fast piston lost its ability to function, bottoming out as the oil contracted. When the chamber was filled with oil, care was taken to put the slow piston at a good position, such that a cooling down from an

operating temperature of 30 or 40°C to room temperature would not bottom out the slow piston. Originally, a very large piston separate from the cart was used in lieu of the cart fast piston. It had enough oil capacity to survive a significant temperature drop. However, it was rejected after its gas-oil seals failed twice, and also because it was a simple 1:1 pressure ratio device, making it impossible to compress above 150 psig with the limiting solenoid valve. Higher compression is better because it implies less time for bubble condensation. A greater pressure forces faster condensation, so that less dead time exists after events (see Section 4.1.5).

4.4.4 Power Outages and Fire

Two 1700VA, 800W APC® Uninterruptible Power Supplies (UPSs) with backup batteries would keep the entire data acquisition chain running for 5–10 min. during a brownout. In a long-term blackout the compression valve would act as fail-safe, defaulting to high pressure, which is ironically the safe condition for a bubble chamber, preventing boil-off. The observed low leak rate of the pressure cart implied that the chamber would be safe for up to a week in this condition. Loss of heat would cause contraction of the oil volume and loss of pressure control, but it could safely be balanced by the drop in vapor pressure of CF_3I with temperature. The heating of mineral oil presented an inherent danger, so the entire detector was wrapped in fire blankets. An automated fire extinguisher was fitted inside the shielding as well.

4.5 Cleaning and Filling Procedures

Because WIMP detection requires a low-background experiment (Chapter 1), a bubble chamber must be cleaned to the highest feasible degree and assembled in a clean environment, just like any other detector for a rare event search. A dark matter detector needs to be as low background as possible, ideally zero, so that the strictest possible limits can be

placed on WIMP cross-section, and so that any alleged signal can be unmistakably above background. After cleaning, a bubble chamber is filled with the buffer water, then CF_3I. The steel outer pressure vessel for the 15 kg chamber was filled with hydraulic recompression fluid after the water filling, but before the CF_3I fill, because without it pressure could not be effectively regulated.

4.5.1 Quartz Vessel Cleaning

The quartz vessel needs to be cleaned prior to chamber deployment and operation to remove radioactive contaminants from the surface, specifically alpha-emitting isotopes, as well as to remove any dust motes that could lead to inhomogeneous nucleation, either on the wall or in the bulk. The quartz vessel for the 15 kg chamber, like all previous vessels for past COUPP chambers, was cleaned in an ultrasound bath. Harsh cleaning, with a cloth for example, is not acceptable because a cloth can introduce additional particles and the risk of scratching. An ultrasound cleaner uses sound waves to bombard the material to be cleaned, forcing trapped particulates to fall off the surface (Bardina, 1988), into a bath whose water is recirculated through a filter. However, the best cleaning parameters (time, frequency, power, temperature) had not been not systematically determined in the past. If quartz is cleaned to excess in an ultrasound bath, the ultrasound may cause surface damage. Sinusoidal surface pitting is possible at high power, long cleaning times, and at certain frequencies. If large enough holes are created, then surface nucleation may result, causing dead time in the detector. At the same time, the vessel must be made as clean as possible. Therefore, an experiment was undertaken with samples of natural TGP quartz with the expectation that synthetic quartz would behave similarly, to determine the longest times and highest powers possible without damage to the quartz surface. An atomic force microscope (AFM) at the Physics Department of University of Chicago was used to determine the surface quality given distinct ultrasound treatments (Appendix A). It was ultimately decided to clean for

5 minutes at 45 kHz, 10% power, 60°C. Bellows, O-rings, and all CF_3I fill plumbing were cleaned in the ultrasound as well, at the FNAL A0 lab. Parts were then given a high-velocity distilled water rinse at FNAL and allowed to dry.

Electrostatic surface implantation of radon daughters (charged after alpha emission), especially long-lived polonium, is a concern because they can contribute to the wall rate. Therefore, between manufacture and delivery, and during transit between cleaning and the assembly stage, the quartz vessel was capped off with an acrylic flange and placed in a metalized bag. Acrylic is well known for its low radon diffusivity. The prevalence of high indoor ^{222}Rn concentrations (ICRP, 1994) are a problem where any direct detection experiment or other low-background experiment is being assembled. Radon levels are significantly lower in a cleanroom environment with good ventilation. As a further precaution, after being finished by TGP, the quartz jars destined for COUPP chambers are given an HF bath to uniformly etch the surface and remove implanted radioisotopes. Radiacwash^{TM} is used as the detergent at the sonication step. It is a cleaning agent widely used for lifting and removing actinides from surfaces.

4.5.2 Cleanroom Assembly and Water Filling

After being dried, all cleaned internal parts were bagged and taken to the class 10 cleanroom at Lab 3 at FNAL for immediate assembly. Leak checking of the quartz-bellow seal was performed. Vacuum was established in the quartz-bellows assembly and was maintained to avoid plating of radon daughters onto the quartz surface.

Light water from the Sudbury Neutrino Observatory (SNO) was then distilled into the quartz vessel. SNO water is some of the cleanest possible (SNO, 2002). Precautions were taken to avoid corrupting the water purity with radon or dust. The water was kept sealed in its original, electropolished container until immediately before use. In preparation for distillation, the water was heated to 80°C while the quartz vessel was chilled to 1°C. The

water was not completely boiled to avoid micron-sized droplets of water dragging dust along into the chamber. Enough water was transferred (∼4 L) to cover the volume of the bellows, and parts of the active volume obscured by the aluminum holder plate and other opaque parts. The opportunity was taken to calibrate the 15 kg chamber pressure sensor first using vacuum and then the vapor pressure of water.

The water had to be degassed to avoid a gas-filled chamber which would make pressure control difficult. It was degassed by evacuating the space above it in the quartz vessel and bellows with a vacuum pump, until the vapor pressure of water was achieved and held for one day. This was significantly longer than the forty minutes allotted for previous chambers, because of the presence of a thin tube (capillary coil, Section 4.5) through which the water had to be degassed. The water was heated from the bottom (an effective degassing technique, Section 4.5.3) but not boiled, to avoid excessive loss of water. After degassing, the water volume was isolated for the several weeks during which the chamber could not yet be deployed underground. No significant rise in pressure was measured, implying a proper degassing.

4.5.3 Oil Degassing and Filling

The oil serving as the hydraulic fluid for the 15 kg chamber (Penreco® Drakeol 7 light-viscosity mineral oil) had to first be degassed before being transferred into the steel outer pressure vessel. The purpose of the degassing was two-fold. First, the liquid must be as incompressible as possible, so that the pistons acting upon it will have enough range to decompress the chamber to any desired low pressure as well as up to a high level of compression, allowing a margin of error in each direction. Because a gas is extremely compressible in contrast to liquid, any dissolved gas will enhance its compressibility. Second, dissolved gas may go into and out of solution as pressure is respectively raised and lowered, causing bubbles to appear in the cameras' fields of view. Although such bubbles could be

position-reconstructed as external to the quartz vessel, and rejected as events, the visual trigger can trip, degrading livetime. The degassing of mineral oil was especially important because the solubility of air within it is high, roughly ten times the solubility of air in water (Logvinyuk et al., 1970; Kubie, 1927).

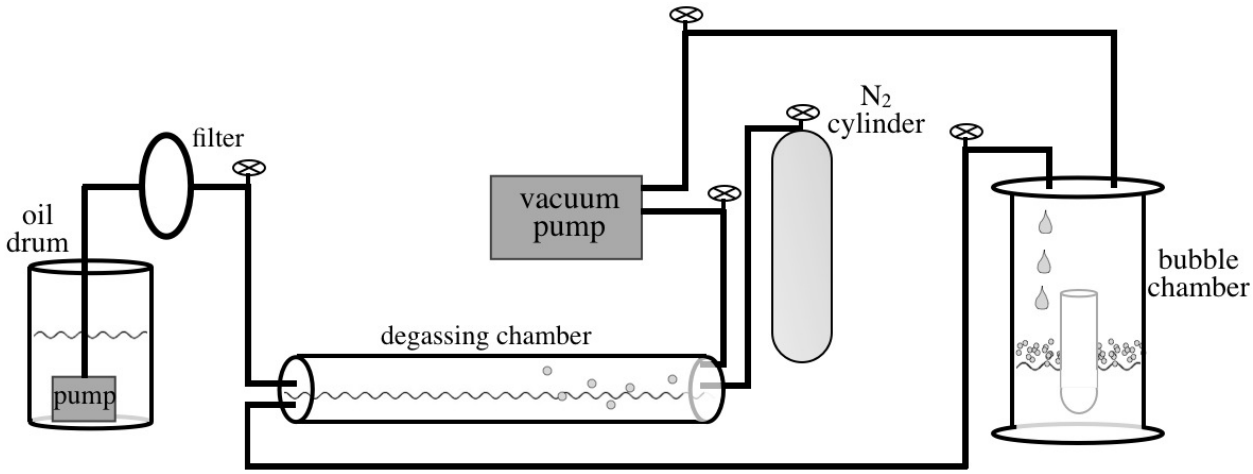

Figure 4.6: Schematic of the degassing process. Oil was pumped, filtered, and degassed before being routed into the bubble chamber by means of a pressure gradient provided by a compressed gas cylinder, as described in Section 4.5.3. The oil fill was performed from the top (far right) to reduce the risk of regassing by the nitrogen used to push the oil into the pressure vessel.

Before degassing, the mineral oil was filtered as it was pumped out of a large source drum. Any particles of dirt or dust, especially white, translucent, or shiny particles, of even a few microns in size, could create a bright point of light of at least a pixel in area. Even floating outside the quartz vessel they could trigger a camera. The mineral oil was pumped from its drum through cheesecloth to remove the coarsest particulates, and then through a 3-micron oil filter into the degassing vessel, an 8 ft. long steel cylinder of diameter 1.5 ft. Degassing is less effective when the dissolved gas must bubble up through a tall column of liquid. Therefore, the cylinder was placed on its side and filled to only a shallow depth. Furthermore, degassing is the most effective when liquid is heated from the bottom. By forcing convection, this brings bottom liquid with dissolved gas upwards while decreasing

the solubility of gasses in the liquid by increasing their temperature. An oil-based motor vacuum pump evacuated the volume above the oil. The degassing process was considered finished after the ultimate pressure of this pump (0.03 mmHg) was reached and held for at least one day. Degassing took ∼3 days. For safety reasons, the heat pads, which covered the bottom third of the cylinder, were brought up to ∼115°C (still far from the minimum boiling point of 150°C for one of the oil components), and only powered when someone was present. The degassing was demonstrated to be effective unintentionally: during the above-ground engineering run of the chamber, at one point the outer vessel was opened up and exposed to air so that needed repairs could be effected to the inner volume plumbing. The oil was not degassed again. The following day, the chamber was unusable, with the pistons traveling one order of magnitude further than they did originally. The piston travel distance for a given change in pressure after degassing is within a factor of 2–3 of what would be expected, given the compressibility of mineral oil (Ohno et al., 2005). This was sufficient for the pistons to have enough usable range, and given the uncertain and variable composition of mineral oil, a mixtures of different oils, it may have even been consistent with an essentially complete degassing.

After degassing, the mineral oil was forced out of the degassing chamber and into the pressure vessel with ultra-high-purity nitrogen gas. Nitrogen was used instead of air because oxygen is a known aggravating factor in oil aging, discoloration, and polymerization or solidification (Faust, 2007). So as to not reintroduce gas into the oil, the filling was done from the top, allowing the oil to fall straight down and form a head of nitrogen-rich surface foam (compare to beer) which would be removed by the vacuum pump, which continuously pumped on the pressure vessel during oil filling. Originally, the oil was to be filled from the bottom, but when this was attempted, it was clear from the cameras that a significant amount of gas was again dissolving in the oil, evident from excessive bubbling throughout

the entire volume. The pressure vessel was evacuated prior to the transfer, so that the nitrogen could be set to 0 psig and still create an effective push, utilizing a one-atmosphere differential. A translucent Teflon tube at the top of the pressure vessel was the highest point in the plumbing, and was thus used to determine when the oil fill was complete. However, two fills were necessary due to the relative sizes of the pressure and degassing vessels. The first (partial) fill was known to be complete once a vacuum gauge at the top of the pressure vessel began to rise sharply from near-vacuum toward 1 atm.

4.5.4 Active Volume Filling

After the pressure vessel had been filled with oil, the quartz jar, pre-filled with evaporated water, was filled with CF_3I. A custom-ordered cylinder electropolished internally for cleanliness and containing 22.5 kg of CF_3I was placed into a bucket of a 50/50 ethylene glycol and water mixture and heated to 30°C. The oil in the pressure vessel was cooled down to ∼5–10°C. A chiller was attached to the copper spiral coils near the top and bottom flanges of the vessel. The copper tubing was a separate and sealed system that was submerged within, but not communicating with, the oil (Section 4.1.4). A thermal siphon was effectively set up, boiling off CF_3I inside its source cylinder and condensing it into the quartz vessel. The CF_3I gas passed through two consecutive filters, 10 micron and 0.03 micron, to remove any dust-sized contaminants. The entire fill line was made from stainless steel, as Donnelly et al. (2004) and collective Collar lab experience have shown that CF_3I can corrode brass and copper. Teflon, as well-known as stainless steel for its lack of reactivity, could have be used as well, but since Teflon is translucent and CF_3I can react with light (see Section 2.6.3), stainless steel was chosen. The fill line was also heated and angled to prevent premature liquefication clogging the constrictive filters. The pressure varied during the fill from 40 to 75 psig as the source cylinder heated up, and fell rapidly as the fill finished. Another sign of a complete fill was when the scale on which the CF_3I cylinder had been

placed reached an asymptote value of 14.67 ± 0.01 kg. The mass must be known in order to normalize reaction rates within the detector, and can later be used as a check on the visual reconstruction of the volume with cameras by knowing the density of the liquid (Chapter 6).

Using the heat pads attached to the pressure vessel in combination with the chiller, the temperature of the entire system was quickly raised to the initial operating point temperature of 30°C after completion of the fill. Although liquid CF_3I is already superheated at room temperature (its boiling point is -22.5°C at 1 atm), the higher temperature provides a lower threshold, increasing sensitivity to lower-mass WIMPs. However, an upper limit on the operating temperature was necessary in order to avoid beta and gamma backgrounds (Section 2.5). The bellows expanded in height as the CF_3I changed volume with temperature, its progress monitored by camera. The rising bellows caused the volume available to the mineral oil (also thermally expanding) to shrink, causing a rise in pressure. Once the pressure rose above 150 psig, the slow piston was employed to commence pressure regulation, and the chamber could now begin data acquisition.

CHAPTER 5
INSTRUMENTATION AND DATA ACQUISITION

This chapter describes all of the active sensors with which the 15 kg chamber was arrayed and the means by which data was acquired with them.

5.1 Pressure Transducer Readout

The pressure of the CF_3I was recorded by means of an ultra-high-purity Setra® pressure transducer. This device contains a thin membrane which flexes with pressure. The membrane is a polarizable material which generates a potential difference when deformed. In this manner pressure is linearly determined by voltage. The pneumatic side of the pressure cart also had its own transducer. This way the internal and external pressures of the chamber could be monitored independently. Because of the spring constant of the bellows, they differed by $\sim$5 psig during chamber operation. The ability to monitor both pressures also helped to prevent an unnoticed bellows malfunction resulting in loss of pressure balance. An extra transducer on the oil volume, located at the bottom of the pressure vessel, was used to quantify pressure gradients.

5.2 Thermometry Readout

Diodes with a temperature-sensitive voltage drop were used to monitor the temperature. One was placed outside the pressure vessel epoxied to its steel surface under the thermal insulation, and three were placed within the oil: one near the top of the top flange, one near the top of the quartz vessel, and one at the very bottom of the quartz vessel and oil. This setup had the purpose of monitoring thermal gradients. Three of the four thermometers failed early on in the experiment. Fortunately the bottom thermometer, placed in the oil close to the hemisphere, remained functional. Together with vapor pressure measurements

(Chapter 8), this was sufficient for determining the temperature and gradient, so that the energy threshold could still be determined.

5.3 Image Acquisition with Cameras

An array of CCD video cameras was placed into the oil volume, affixed with screws to aluminum bars suspended from the middle plate supporting the quartz vessel (Section 4.1.1). Twelve cameras were arrayed in four columns of three in a semicircular arc on one side of the chamber. Three cameras were necessary to see the entire length of the jar without blind spots. This was because of the long, thin geometry of the chamber and the proximity of the cameras to the jar. Two cameras were needed for each row to do stereoscopic reconstruction of bubble positions in three dimensions. An additional two rows were added to the system for backup in the event of failure, through oil inundation of the sealed lenses, for example. This indeed occurred to two cameras on the timescale of months. Other COUPP chambers, which have external cameras and viewports (Section 1.5.3), require only two cameras to view the entire volume and to perform optical reconstruction.

Only black-and-white images were taken even though the cameras were color-capable. This simplifies the analysis of raw images, both for triggering purposes and later for bubble number and position determination. Eight-bit mono (i.e., where 0=black and 255=completely white) was used, as bubbles were not bright enough to justify use of 16-bit or higher. Ten images of each event were saved in an attempt to get a full picture of bubble evolution. The effort was made to record images both before and after bubble birth, although images could no longer be taken beforehand once the visual trigger was no longer the primary method used (Section 5.6). Images were taken by the cameras at a frame rate of 30 fps for the 15 kg chamber. The other chambers use cameras with one to two orders of magnitude greater frame rates. However, in order to not risk damage to high-speed cameras

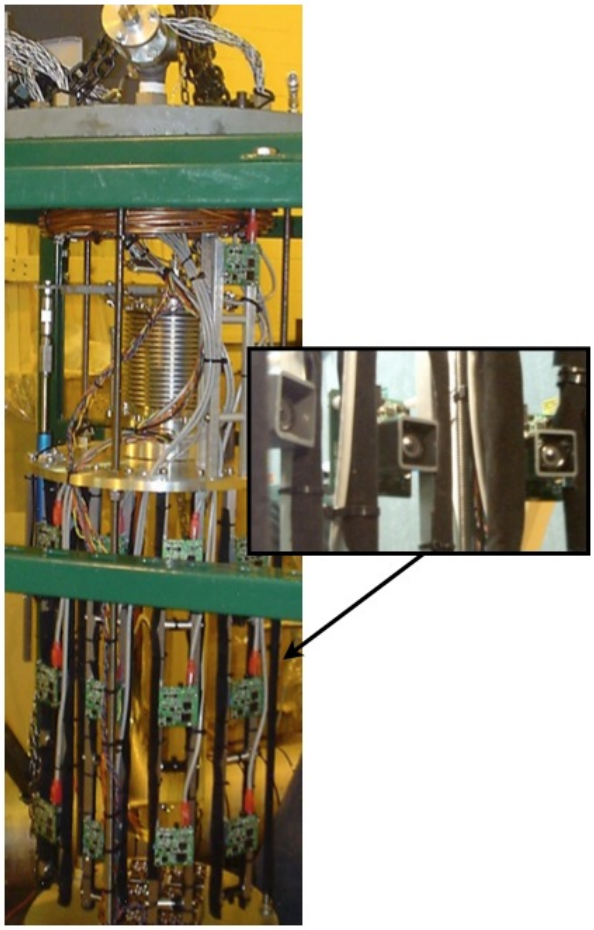

Figure 5.1: Camera array surrounding the 15 kg chamber, in the context of its support structure. The linear transducer for bellows position determination (Section 4.4.2) is also visible (blue rod with silver extension at center left).
(detail) Front of cameras. The gray PVC cubes provide glare reduction.

submerged in oil, inexpensive naked CCD board cameras were used for the 15 kg chamber, simple 640×480 resolution IEEE 1394 Firewire-A 400 Mbps cameras from Unibrain, Inc. The framerate was sufficient to catch bubbles of order mm in diameter (Lau, 2008) as they grow as the square root of time (Chapter 2). One challenge was to strip the firewire cables and propagate them through high-pressure feedthroughs into the chamber without signal degradation (Appendix B).

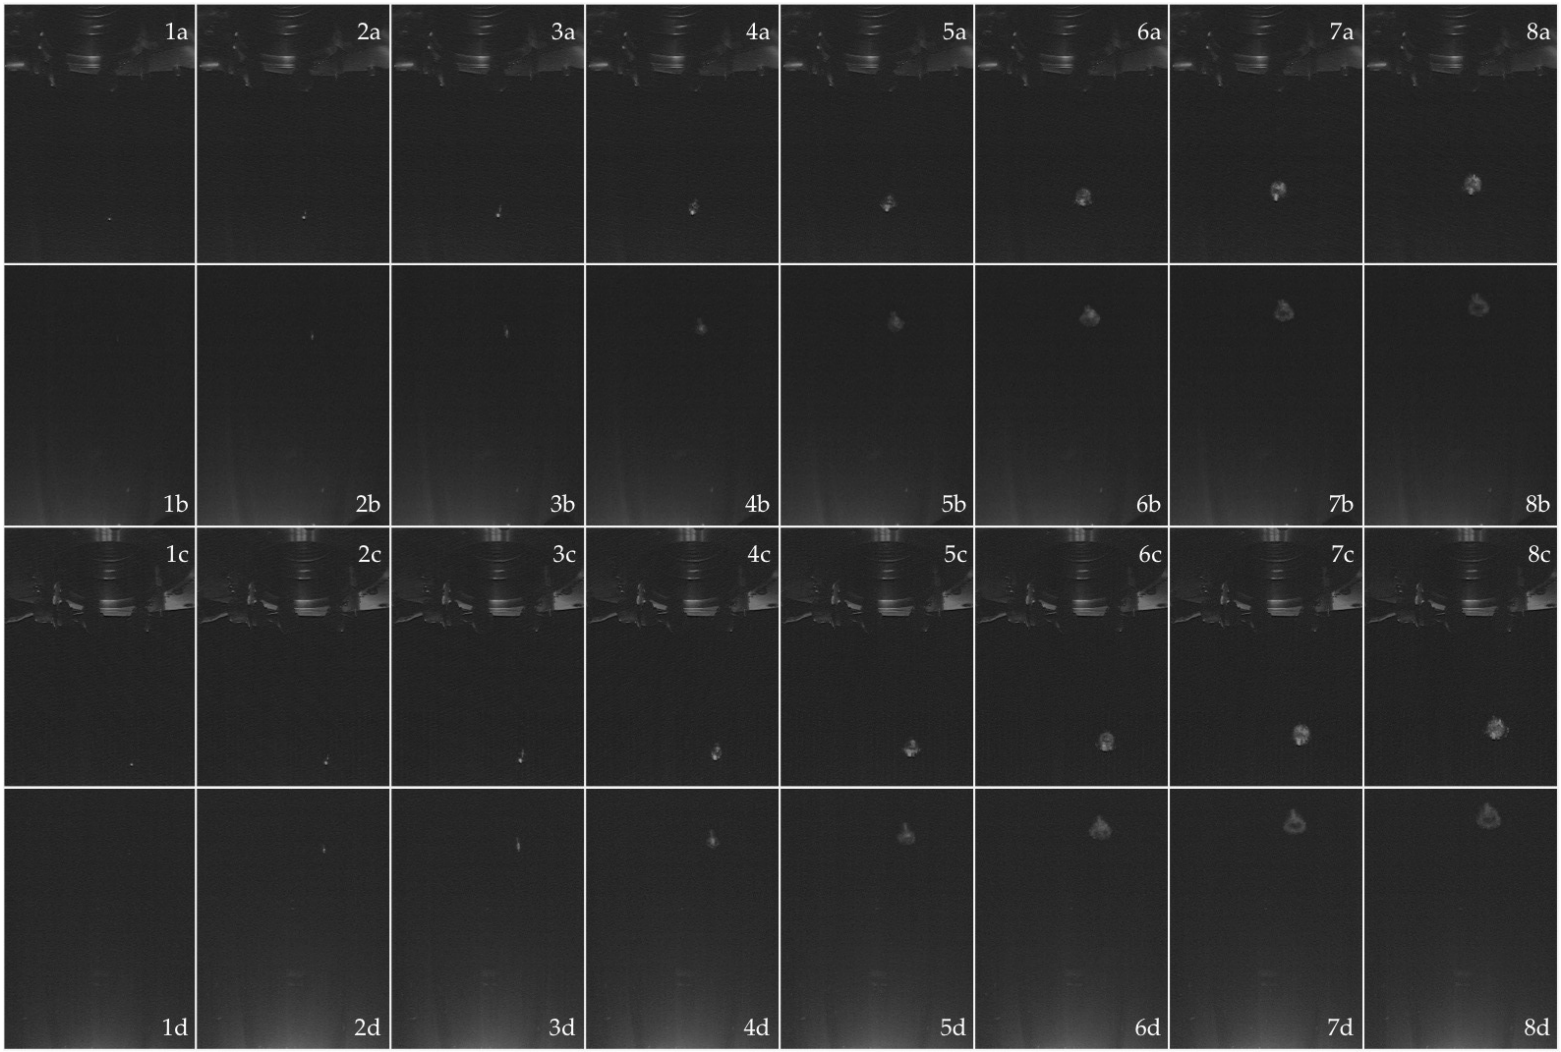

Figure 5.2: Samples of bubble evolution images as taken by the 15 kg firewire cameras at 3 psig, 30°C. The same bubble is seen by the overlapping visual area of two camera rows, top (rows a and c) and middle (rows b and d). Individual frames are separated by ∼33 ms. Bubble deformation and fading at late stages was caused by the onset of recompression.

5.4 Piezoelectric Acoustic Sensors

Four piezoelectric sensors were epoxied to the quartz jar with JB Weld in order to record the sound emitted by bubbles. A piezo contains a polarizable material which, when biased with voltage below ground at one end and above at the other, generates sinusoidal voltages in response to compressions and rarefactions caused by an incoming sound wave. Piezos turn sound into voltage, similar to speakers, which instead turn voltage into sound. Because of the slow framerate of the cameras compared to the muon veto rate (Section 5.7), the piezos, digitized at 0.25–2 MHz, were used to do muon veto coincidence timing (Figure 5.6). That made the possibility of false coincidence negligible. When a bubble forms, the expanding pocket of gas spherically pushes the surrounding liquid outward, producing an audio signal that extends into the ultrasound range (Apfel et al., 1983). The bubbles are audible to humans as well: in fact, Glaser used a phonograph to record the audible "plinks" of bubbles in some of his later chambers (Glaser & Raim, 1955). The SIMPLE experiment uses piezos to "listen in" on bubbles (Collar et al., 2000), and so does PICASSO (Barnabé-Heider et al., 2005; Piro, 2010), both similar to COUPP. These experiments operate superheated droplet detectors (SDDs) instead of bubble chambers. The COUPP collaboration, which shares members with PICASSO and has former members of SIMPLE, began utilizing piezos with the 2 kg chamber.

Three of the piezos of the 15 kg chamber were placed in a semicircular arc at the top of the quartz vessel, opposite the cameras so as to not block their view. The fourth was placed at the center of the bottom on the hemisphere. This was done to maximize the chances of the piezos being able to pick up bubbles originating in any part of the chamber, while ensuring that they were not substantially obscuring the black backdrop. The piezos were directly attached to the 15 kg quartz vessel because in the 2 kg chamber run, the piezos, which were not attached to the quartz, recorded faint or non-existent acoustic signals at

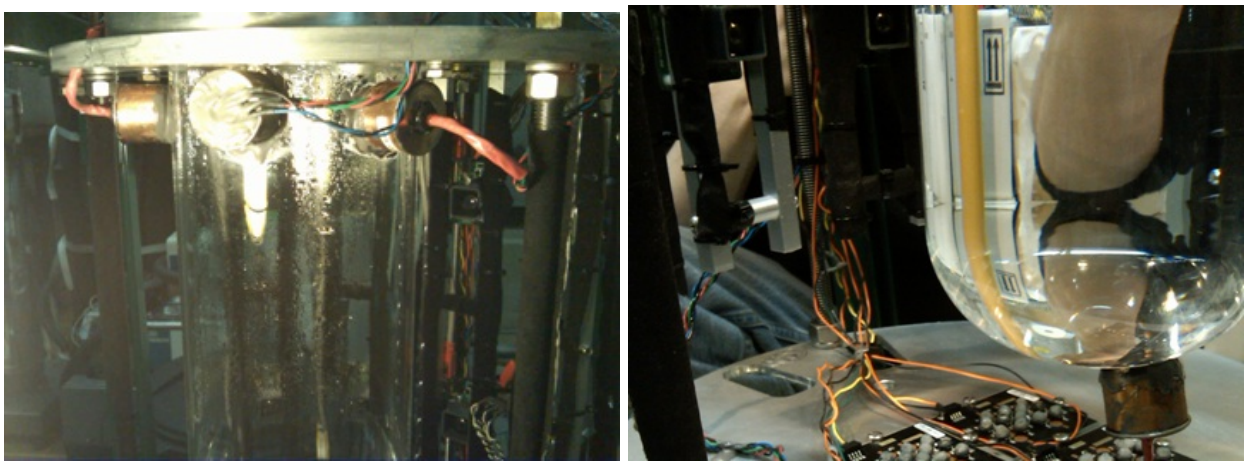

Figure 5.3: Piezo transducer positions at the top (left) and bottom (right) of the 15 kg chamber. They were used to trigger on the sounds of the bubbles in lieu of a optical trigger, converting sound waves into voltage (Section 5.6.3). They were also used to distinguish between neutrons and alphas (Chapter 7).

high pressure (low superheat) due to sound attenuation in the recompression fluid (ethylene glycol).

5.5 Computer Digitizers for Data Acquisition

The pressure transducers were read in by a 200 kHz maximum sampling rate National Instruments PCI-6035E card, which also interfaced with the pressure cart to do fast compression. The pressure cart PLC (Programmable Logic Controller) transmitted piston positions via ethernet cable and operated the slow-control piston. The thermometers were read by the 48 kHz NI USB-6009 card. This card controlled the heatpads and all LEDs, operated and monitored the neutron source, recorded the bellows height, and performed other "slow" DAQ (data acquisition) tasks. The cameras were plugged into firewire PCI cards. A 2 MHz PCI-6133 digitizer captured piezo traces and the muon veto signal. The general philosophy was to place "like with like," so that interdependent systems were handled by the same hardware. Software written by the author in National Instruments LabVIEW 8.5 coordinated all aspects of the DAQ.

5.6 Triggering on Bubble Detection

The COUPP chamber design is inherently scalable because larger masses are easily monitored. Three triggering mechanisms exist: pressure rise due to the bubble growth, acoustic (ultrasound) detection, and pixel subtraction and comparison methods using camera images.

5.6.1 Pressure Trigger

In early test vessels for COUPP, pressure triggering was as reliable as the visual trigger. A pressure trigger involves a DAQ card sending pressure transducer voltages to a computer, which waits for sudden, prolonged (to avoid noise triggers) upward shifts above a threshold determined by the average noise in the transducer. This method for triggering is relied upon less and less in larger chambers, having been used for the 15 kg chamber only in the event of other triggers failing. The larger the volume of a chamber, the smaller the difference in pressure caused by a bubble and the slower it increases. The reliability of this trigger is coupled very tightly to the degree of superheat, as bubbles are less violent at low superheat.

5.6.2 Visual Trigger

This was the primary trigger for both runs of the 2 kg chamber, for the 4 kg chamber, and now for the 60 kg chamber. Likewise, the 15 kg chamber used visual triggering at first. CCD video cameras constantly monitor the chamber for this trigger. Grayscale image frames adjacent in time are subtracted from each other: pixel differences above noise are sought. If there are enough pixels above threshold or a single pixel above a higher threshold, then the computer monitoring the experiment signals there is a probable bubble and compresses the chamber. Images are written to file for inspection and further analysis. Bubble position and multiplicity are later determined.

In the 2 kg chamber, the sources of light were LEDs mounted behind the quartz jar,

shining directly into the cameras. For the 4 and 60 kg chambers, the LEDs are behind the cameras and shine upon a retroreflective screen. Either way, the effect is the backlighting discussed in Section 3.1. A computer determines whether the trigger condition has occurred by finding dark pixel hits, i.e., pixels below a threshold. The 15 kg chamber, on the other hand, has LEDs underneath the chamber shining upward. With this method, although the signal-to-noise ratio is comparable, there are possibly more differing pixels on which to trigger (Section 3.1). For the 15 kg chamber, the speed with which the computer analyzed the images had to be optimized. In previous chambers, only two cameras had been necessary to monitor the entire volume and reconstruct location, but the 15 kg quartz vessel required at least six. The speed of reading six cameras was optimized by simplifying the algorithm for determining the trigger condition. A fast histogramming algorithm (LabView IMAQ Histogram) was used to determine whether the single brightest pixel was above a threshold, instead of converting images to arrays and considering multiple pixels. This method takes only 1–2 ms per pair of camera images to perform subtraction and histogram of the resulting image with LabVIEW. Therefore, given the framerate, there was ample time to check for a trigger on all cameras during the 33 ms before the next frame. Parallel processing on a quad-core PC was used to simultaneous juggle all DAQ tasks.

LEDs were flashed in time with the cameras to avoid excessive heating, decreased lifetime, and ill effects on the mineral oil such as light-catalyzed polymerization (Appendix C). The LEDs used were high-luminosity ∼15 W 625-nm peak wavelength Philips Lumileds® LUXEON® LED clusters of 12–18 individual LEDs each divided in half between parallel circuits (cluster models used: LXHL-MDCB, LXHL-MDJA). Two clusters lit the chamber, not centered on the chamber. They were placed closer to the back away from the cameras to minimize glare and maximize the brightness of the distant bubbles. The voltage was set low enough so the CCDs were not light-saturated, but high enough to see bubbles, and

optimized to have the highest signal-to-noise ratio. This was a non-trivial task because the signal and noise would both increase with more light, each at different rates.

The pixel threshold had to be optimized to avoid time-wasting triggers on digital noise yet enable triggering on small bubbles far away from the cameras in their earliest stages of growth at the lowest degrees of superheat. The algorithm had to be kept simple so it would not be time-consuming. Therefore, contrast, shutter speed, etc. were optimized on the cameras while small bubbles were watched. These bubbles were introduced in a controlled fashion with a copper tube and a compressed air bottle with regulator. This was done prior to deployment with a mock setup consisting of a dummy quartz jar filled with water and placed in a plastic drum full of oil. The index of refraction of CF_3I was realized to be identical to that of water (Appendix D). It was decided to err on the side of sensitivity, so that bubbles were not missed, at the expense of false triggers adding $\sim$1% dead time. During underground operation, a new significant source of false triggers was pinpointed, found also by the DM-TPC experiment, termed "worms" (Dujmic et al., 2010). These were perfect squares (4×4 or 6×6) of pixels brightly lit for one frame only. Because worms were uncorrelated between cameras, a minimum two-camera coincidence requirement addressed this issue without hindering the computer's ability to trigger on bona fide bubbles.

5.6.3 Audio Trigger

When the windowless bubble chamber idea was first conceived, the plan was to rely solely upon an array of piezos for a sound trigger, instead of submerging cameras. This simple windowless concept has uses outside of dark matter research, such as in homeland security (Section 1.6). After Monte Carlo simulations of sound propagation through a bubble chamber were performed by this author, however, it became clear that multiple bubble events would hamper time and position determination. Even a large number of piezos covering the outside surface of a chamber would fail to provide a unique result for the number of

bubbles, their times of birth, and their individual positions, creating a large system of non-linear equations. Bubble position resolution might be poor due to other conditions. The positions of the piezos could not realistically be known to high enough precision, and quartz vessel ringing could cause a different piezo to be the first to register the bubble than naïvely expected assuming linear sound propagation.

The 15 kg chamber at first used the visual trigger like all other chambers, but evidence mounted that the LEDs, despite being flashed and not set to maximum intensity, were causing the CF_3I to be heated above the desired temperature setpoint. Vapor pressure measurements and neutron rates from a calibration source were both higher than anticipated. A trigger based on pressure or bellows position was considered and rejected for being much too slow.

A slow trigger leads to numerous problems. First, images are taken of bubbles that are so large that it is impossible to tell whether the event was a single or a multiple, because several smaller bubbles may have had the chance to merge into a larger one. Furthermore, slow triggering disrupts the ability to pinpoint bubble origin, which is important for rejecting surface nucleation and for demonstrating a homogenous distribution in the case of a potential WIMP signal (WIMPs, of course, have no preferred interaction site). A slow trigger can also allow CF_3I gas to wet the surface of the quartz vessel and cause more wall events.

However, the underground chamber in which the 15 kg chamber was deployed was noisy, with construction ongoing outside, making the one remaining option, a sound trigger, difficult to implement. Because of the size of the vessel, a unrealistically large soundproof box would have had to be constructed, and because the visual trigger was to be the primary trigger, this was not considered beforehand. This issue was addressed by implementing two thresholds, one so high that a bubble was certain, and another lower, as shown in Figure 5.6. The lower threshold was used to determine whether to turn the LEDs on or not. If

the second trigger threshold was not crossed within 100 ms, then the LEDs were shut off. Otherwise, the chamber was compressed and pictures were taken of the event. Triggering was not permitted for the 25 seconds following decompression due to piston and chamber settling noises. In conjunction with bandpass (Butterworth) filtering in software to remove high-frequency noise, careful threshold tuning made the false trigger rate negligible. The voltage across the LEDs was also greatly reduced, but the apparent brightness did not suffer dramatically. Images of bubbles were still of an acceptable quality, but ease of bubble finding was moderately affected (Section 6.3). This was because when turning the LEDs on suddenly after long periods of darkness, the cameras would be temporarily more sensitive to light during an initialization period of self-adjustment.

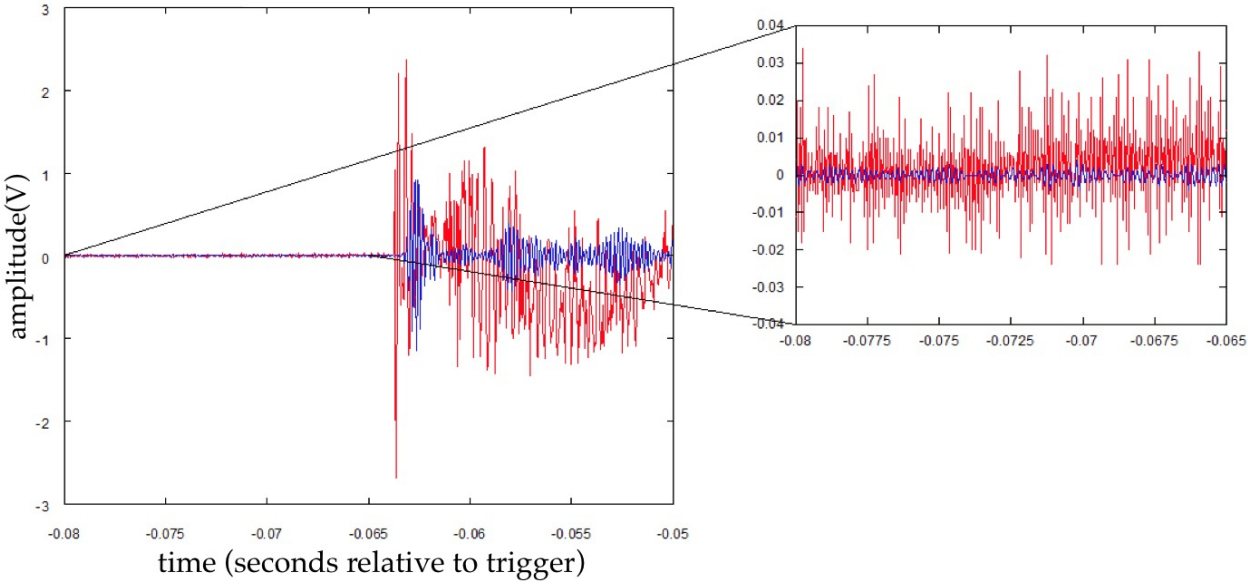

Figure 5.4: Unfiltered (red) versus bandpass (4–6 kHz) filtered (blue) piezo signal, 25 kHz. Signal and noise are both reduced, but noise more than signal (inset), allowing for easier t_0 determination and more effective triggering. Slow pressure rise, manifested in a piezo as a negative slope, is absent after filtering. (See also Figure 5.5.) However, this feature of the traces is not necessary for triggering, except when at low degree of superheat, i.e., high pressure.

The time of the trigger (t=0 in Figures 5.5 and 5.6) does not correspond with the

crossing of the second, stricter threshold because of the time it would take for the computer to retrieve the acoustic data from the PCI-6133 card and process it. The 6133 card comes with on-board buffering and low-level thresholding capability, but the need to do a higher-level analysis of the signal to remove room noise made it such that the computer had to attempt a continuous retrieval of the acoustic data. To make this as smooth and fast as possible, and to permit retention of enough data far enough back in time to observe the beginning of a bubble at all pressures, the digitization rate was lowered to 25 kHz from the maximum of 2 MHz for the NI 6133 card. The filter illustrated in Figure 5.4 was an 8th-order bandpass filter from 4 to 6 kHz. The time between points (40 μs) was small enough to be far below coincidence timing algorithm errors (Section 5.7.4). It was still possible to use these lower-frequency signals to do particle discrimination (Chapter 7).

The 15 kg chamber was the first large COUPP bubble chamber to rely primarily on the sound-based trigger, although the precedent for this type of trigger was already established by D. Glaser in the first bubble chambers, where he used a phonograph to trigger the lights and cameras (Glaser & Raim, 1955). More recently, SDDs (Superheated Droplet Detectors) like those used by the PICASSO collaboration have relied on piezos for triggering (Boukhira et al., 2000; Barnabé-Heider et al., 2005). SDDs use droplets of superheated liquid dispersed in a passive gel which essentially act like miniature bubble chambers, of order 10–100 microns in size (Collar, 1996). Bubbles are difficult to see deep in the opaque medium used. PICASSO evenly distributes piezos around the outside of their detectors, to ensure they pick up bubbles located anywhere in the detector, and then places them in sound-insulated boxes.

The amount of acoustic energy generated by a bubble depends on the degree of superheat to the power of 3/2 (Martynyuk & Smirnova, 1991). Nucleation is a more violent event the less stable (farther from equilibrium) a superheated liquid is. At high pressures near the

vapor pressure, the piezos mounted on the 15 kg chamber did not produce any discernible acoustic emission above the white noise of the background. However, they would still show a change in slope in the flat baseline at all degrees of superheat, due to the rise in pressure from a bubble. This made the piezo trigger equivalent to a pressure trigger at high pressure, and better because the "official" pressure transducer on the chamber was noisier and slower to respond to changes in pressure. Following the rationale given in Section 5.4, the piezos were not mounted along the side of the quartz vessel, which could have alleviated this problem. Before assembly, the piezos were not under consideration as the primary trigger.

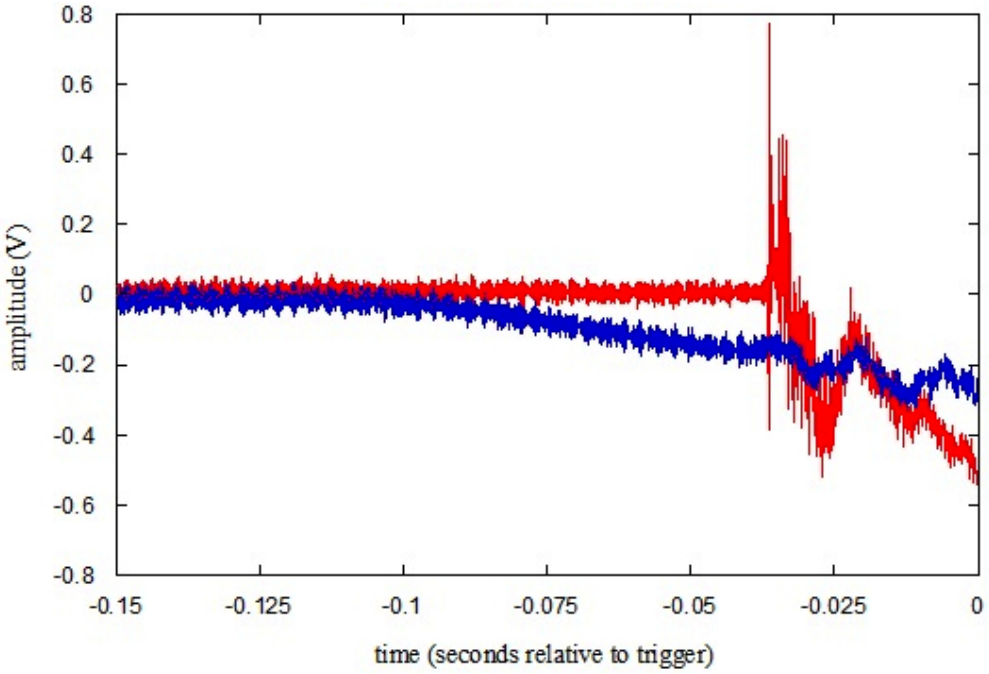

Figure 5.5: A "healthy," loud piezo trace (3 psig, red) and a poor-quality piezo trace at high pressure (27 psig, blue): no discernible acoustic emission from the bubble. Such traces would not be filtered, so that the change in pressure as observable from the drift in the baseline could be used as a trigger.

As long as no cameras were malfunctioning, the visual trigger was inherently 100% efficient. Bubbles become very large and bright and must eventually be detected even given a poorly chosen pixel threshold. Moreover, the buffer of previous images contains the past

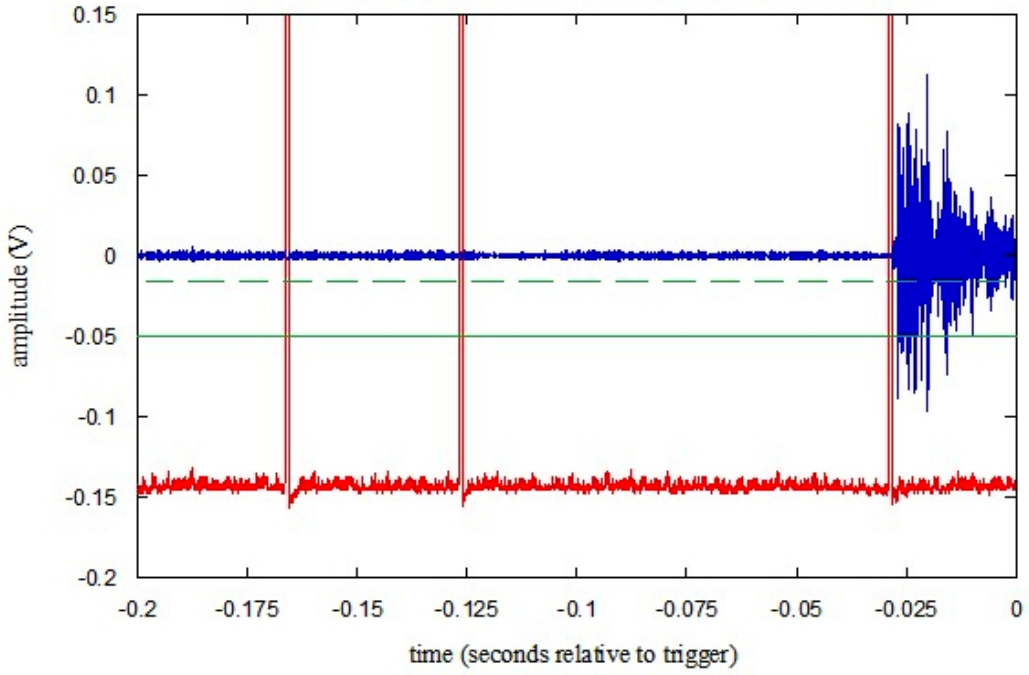

Figure 5.6: Sample muon veto coincident trace with bubble detection thresholds displayed. The piezo signal is shown in blue, and the veto logic signal in red. Dashed and solid green lines respectively represent the first (LED turn on at -0.015 V) and second (confirmation of bubble at -0.050 V) trigger thresholds (see Section 5.6.3).

growth (Section 5.6.2). However, the efficiency of the new sound trigger was in question. To address this concern, during the development of this trigger the visual trigger was left active as a cross-check. Bubbles nearest to the water interface and forming at the lowest degree of superheat utilized in the 15 kg chamber run, i.e., its highest pressure (60 psig), were especially scrutinized because they were likely to be missed: bubbles stop growing once they cross into the water. The low probability of events at such high pressure (low superheat) was enhanced by using the neutron calibration source (Chapter 8). It was observed that even for bubbles at the water interface at 60 psig, the piezo trigger would still be effective. Even such bubbles could cause $\sim$1 psi increases in the pressure of the chamber and a very slow further increase towards the vapor pressure if pressure control and compression were disabled. Signals like those in Figure 5.5 were observed for such a minute incline in pressure,

which could take up to 1 second to cross the thresholds. As a backup, the pressure trigger would still be able to register these difficult-to-detect bubbles because pressure stability was controlled to better than 0.1 psi (Section 4.1.5). Two seconds of acoustic traces were saved for each piezo to ensure that the bubble start time was included. This was the maximum that could be achieved with the 6133 card at 25 kHz without causing a memory overflow error due to the amount of data being read from five channels (four piezos and the veto).

5.7 Cosmic Muon Veto

At shallow depths ∼O(100) m.w.e. cosmic rays are still a concern. Relativistic muons in particular have a long enough lifetime to traverse the atmosphere and reach a detector, and can constitute a significant background (Bogdanova et al., 2006). Although the chamber runs at a low enough degree of superheat where muon interactions cannot lead to bubbles, neutrons may be byproducts of muon-induced particle showers (Riley, 2007; Kozlov et al., 2010). In small chambers, fiducialization and bubble multiplicity are insufficient to tag every neutron, and at the ∼330 m.w.e. depth of the 15 kg chamber, a ∼1 muon/(m^2×s) flux was expected (Bugaev, 1998; Bazzotti et al., 2009; Bogdanova et al., 2006). Therefore, a muon veto was built by Nathan Riley (Riley, 2007) to surround the chamber.

As delineated in Kozlov et al. (2010), muons can lose energy by four different means: ionization, electron-positron pair production, bremsstrahlung radiation, and deep inelastic scattering. Hadronic and electromagnetic showers may develop in all cases. Although neutrons can be indirectly generated in electromagnetic showers, the hadronic showers are more likely to lead to neutrons interacting in the detector, stemming from high-energy neutron spallation in dense substances or heavy nuclei (Kozlov et al., 2010), such as the steel of the outer pressure vessel or the iodine in the CF_3I, respectively.

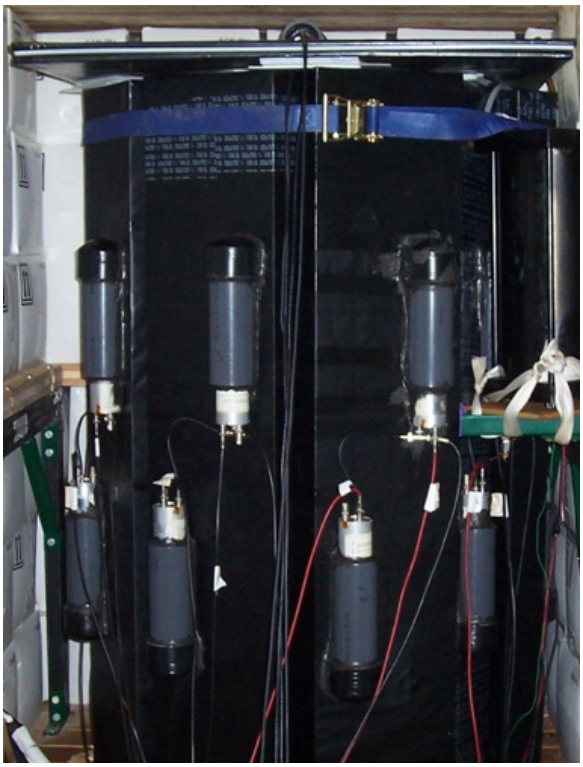

Figure 5.7: The complete muon veto setup, 4π except for the bottom, for the 15 kg chamber, deployed at TARP around the steel vessel. Eight (60 in. $\times$ 12 in. $\times$ 4 in.) polyvinyltoluene scintillator panels, with one pair of PMTs (silver bases visible) affixed to each, surrounded the vessel, in addition to two top panels with generous side overhang and one PMT each. All of the panels were carefully wrapped with several interleaving layers of opaque black electrical tape to minimize light leaks. All of the PMTs were protected by stiff gray PVC tubes and faced light guides at the ends of the tubes, affixed to the plastic panels with optical cement.

5.7.1 Physical Setup of the Muon Veto

A muon ionizes the scintillator plastic (polyvinyltoluene). The residual kinetic energy of these ionized electrons results in excited states upon recombination. When an excited state decays back to ground, a photon is released: the plastic is said to scintillate. PMTs are used to collect this light. One or more photons eject electrons from a photocathode via the photoelectric effect. Stages of electrodes at ever higher voltages lead to an electron avalanche which results in a detectable current.

The muon veto system for the 15 kg bubble chamber consisted of eight scintillator panels forming an octagon around the outer vessel and two panels stacked on top of each other just above the vessel. Each of the side plastic panels comprising the octagon included two PMTs, and each of the top panels had one. All panels were covered with liberal quantities of an opaque black tape to shield the scintillator from stray light in the room. Light leaks are a background, and if bright enough, can overwhelm and damage the PMTs. The PMTs were held in place on the panels with opaque gray PVC tubes and a black silicone RTV (Room Temperature Vulcanization rubber acting like an epoxy). The PMTs faced light guides attached to the plastic panels with optical cement, and the panels were covered with reflector underneath the black tape.

5.7.2 Muon Veto Characteristics

The voltage threshold for recorded PMT pulses had to be set high enough to avoid gamma background and overwhelmingly high rates, but low enough to permit muons to be detected with high efficiency. The continuously ionizing muon deposits more energy per distance (about 2 MeV/cm in plastic scintillator) and leads to more scintillation than a gamma ray does (Bichsel et al., 2006). A significant contribution from gammas to a muon veto rate could lead to a high false coincidence rate: veto pulses due to gammas might be mistaken for muon-induced. A set of old PMTs from the kTeV experiment at FNAL were recovered for use in the muon veto of the 15 kg bubble chamber. Biases need to be individually set for PMTs because gain can vary from unit to unit even for the same model. The bias voltage affects the veto trigger rate by changing gain. The muon efficiency may be low if the gain is insufficient. The gains were set in such a manner that all side-panel PMTs could share the same threshold in voltage and approximately the same trigger rate (roughly 0.5–1 Hz underground). Two PMTs were attached to each side panel in order to ensure an even coverage in light collection (Riley, 2007).

The pulses from all the side panels were combined using a Boolean "or" operation on a NIM (Nuclear Instrumentation Module) logic module. The PMTs on the top panels were run in coincidence in order to properly identify muons. Being only 1 cm thick, it was impossible for the panels to individually distinguish between gamma- and muon-based energy deposition (Bichsel et al., 2006). A 0.25 in. thick slab of polyethylene between them suppressed spurious coincidence from electrons ejected by gammas (Riley, 2007). The top panels extended beyond the muon veto side panels to increase efficiency.

5.7.3 Muon Veto Calibration

The rate from the muon veto was stable for the duration of the experiment. It was consistently 6–8 Hz total, 1–2 Hz from the top panels. This was expected for a 330 m.w.e. depth (Bugaev, 1998; Bazzotti et al., 2009; Bogdanova et al., 2006). The efficiency was determined to be 95% when the muon veto was tested prior to deployment at the LASR subbasement lab (6 m.w.e.). The efficiency was determined by the coincidence of muon veto pulses with pulses from a NaI[Tl] scintillator crystal, placed at the center of the veto. The veto was setup in its correct octagon shape, including the top panels (Riley, 2007). When the veto was deployed at TARP, efficiency was checked again, primarily because the LASR mockup was not an emulation of the final geometry: to allow some plumbing and cables to come through, a gap had to be created between two of the side panels. A longer NaI[Tl] calibration was needed than at LASR to get good statistics because of the much lower muon flux underground at TARP. The NaI[Tl] was placed next to the pressure vessel outer wall at TARP. The efficiency was determined to be $\sim$90%, a conservative value because of the possibility of alpha contamination in the NaI[Tl] crystal. Given the loss of efficiency produced by having to create a gap between panels, a $\sim$5% loss in veto efficiency compared with the value obtained at LASR was not unexpected (Riley, 2007, 2009). Weak ^{88}Y and ^{137}Cs gamma sources were used for the energy calibration of the NaI[Tl].

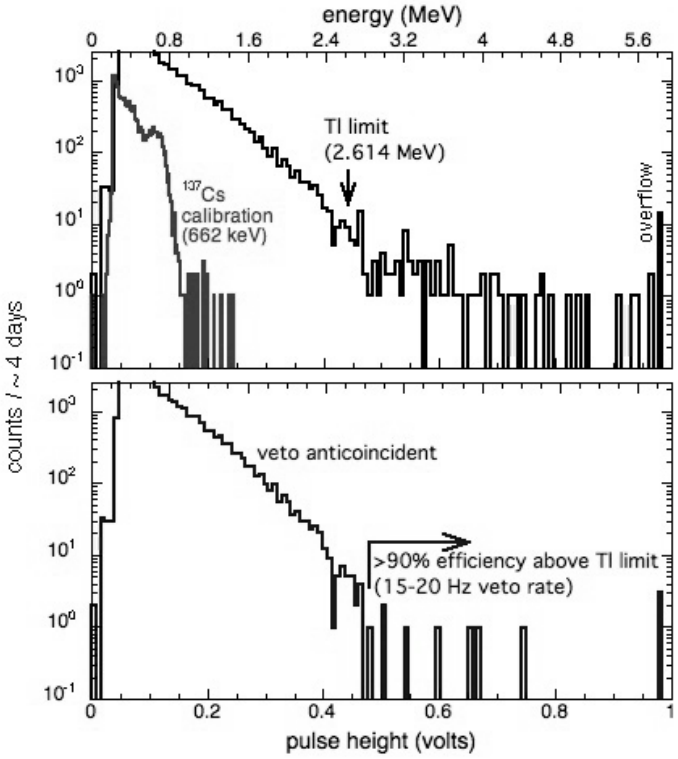

Figure 5.8: The energy and efficiency calibrations of the muon veto. In actual running after this calibration the veto stabilized to 7 Hz, but this was not a sudden increase in inefficiency: the difference in rate was traced to one aging PMT which would sporadically increase in rate. (The ^{208}Tl limit is the maximum gamma energy of the^{232}Th and ^{238}U radioactive decay chains. Because of the prevalence of these chains, it essentially provides a minimum energy pedestal, above which other particles like muons can more easily be observed and studied.) $\sim$90% efficiency for cosmic muons was measured, determined by checking for coincidence between the veto and a NaI[Tl] scintillator suspended just outside the vessel, within the octagon of the veto. Figure courtesy of J.I.Collar.

5.7.4 Application of the Muon Veto to the Experiment

The piezoelectric transducers were used for muon veto coincidence timing. Given the veto triggering rate of $\sim$7 Hz, using the 30 fps cameras instead for this purpose was less than ideal. At low degrees of superheat, the bubbles produced a loud "plink" easily discernible above the background noise of $\pm$100 mV peak-to-peak, likely caused by ambient room sounds. Above about 18 psig, the signal-to-noise ratio became too poor to observe clear acoustic

emissions from bubbles. However, a change in slope in the baseline was still evident, caused by the pressure rise a bubble induces, but this did increase the error in finding the start time of a bubble (see Figure 5.5). LabVIEW was used to analyze the traces, to find the start time of statistically significant deviations from the baseline noise of a piezo. Because the piezos were at different locations throughout the chamber (three at the top in a semicircle, one at the bottom of the quartz hemisphere), their minimum distances to the quartz vessel varied and the signal propagation time to each piezo varied. The earliest determined start time was used for veto coincidence determination. Time determination was too unreliable to do muon veto coincidence timing outside the range of 3–15 psig.

It was important to find the correct time marking the beginning of a bubble, with $\sim$O(1) ms precision, because of its use in the veto coincidence cut (as in Figure 5.6). The muon veto rate $\sim$O(10) Hz already implies a 1% spurious coincidence rate even at such high precision (Figure 5.9). The earliest moment of bubble existence must be found, while concurrently not mistaking noise for signal. The first point above a fixed voltage threshold (set empirically by hand-scanning pulses) would be found in a given trace. The threshold would be chosen to be high, to be sure it was already after the bubble, as it was used to coarsely determine the start time. The subsequent 1 ms window was checked for another point above threshold. If none were found, then the next point above the original threshold would be sought. However, if another point above threshold were found, then the time window between 1 and 4 ms before that point would be scanned, seeking the first voltage 5-sigma above the root-mean-square of the absolute value of a flat, pre-bubble section of the trace containing only white noise. This algorithm was redundantly verified using built-in "transition measurement" and "basic level trigger" detection routines in LabVIEW, and vetted on over 200 events by eye.

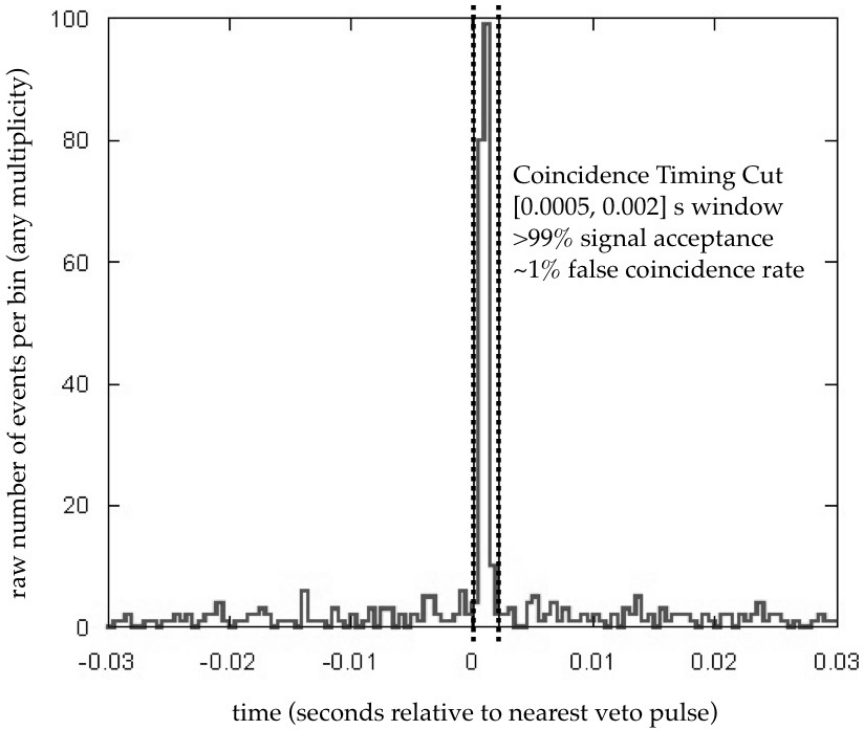

Figure 5.9: Determination of the muon veto coincidence timing window cut, using a histogram of the time between the start of a bubble sound in an acoustic trace and the time of the closest muon veto pulse (rising edge). The peak observed is slightly to the right of t=0 because of the time that it takes the prompt neutron produced by a muon to reach the bubble chamber (up to 20 μs), added to the subsequent time for sound propagation from the neutron-induced bubble (O(100 μs) in CF_3I). The estimated false coincidence rate quoted is based on the total veto rate. The signal acceptance estimate is based on an assumption of Gaussian error. "Any multiplicity" refers to events with any number of bubbles. (Neutrons are capable of scattering multiple times within a chamber because of a high cross-section in CF_3I.) Only combined data for the 3–15 psig range are depicted because above 15 psig the piezo time resolution becomes poor (see Section 5.7.4).

CHAPTER 6
3-D BUBBLE POSITION RECONSTRUCTION

It is important to properly reconstruct the locations of all bubbles in three dimensions to reject wall events, prove a homogeneous distribution of suspected WIMP-induced bubbles, and formulate a distance correction for the amplitudes of the acoustic traces if they are used for particle identification and discrimination. Spatial reconstruction is also useful for comparing the actual position distribution of bubbles occurring during a neutron source calibration with the predicted distribution from a Monte Carlo simulation. In previous COUPP chambers, no spatial calibration was done beforehand, but reconstruction was not hard to perform because the wall events provided a "crust" of bubbles clearly defining the cylinder and hemisphere of the quartz jar. With synthetic quartz creating far fewer wall events, this was deemed to no longer be a realistic approach. Thus, preparations were made to calibrate the 15 kg chamber with reference coordinates prior to deployment. This chapter reviews the trial and error which led up to the final method successfully used for the spatial reconstruction.

6.1 Position Calibration in Three Dimensions

Luckily, the index of refraction of CF_3I at 30°C is essentially identical to that of water at room temperature (1.33), allowing easy position calibration in a test vessel filled with water. Although the chamber could be run at higher temperatures, correction for the changing index of refraction of CF_3I is possible: its temperature dependence is known (see Appendix D). As reference points, stainless steel spheres of ∼2 mm radius were used. They were suspended in long chains of ∼1 mm inter-sphere spacing, spanning the entire length of the quartz vessel. Only one chain was used, but its position was changed and images taken again with all twelve cameras. It had 53 possible positions, being hung from holes in an

acrylic disk set on top of the quartz vessel, bolted into place on the same flange from which the quartz jar was suspended by its top lip. There was a center hole, and 52 holes arranged in concentric circles of ever-increasing radius, from 0.5 in. up to 2.5 in. (recall the quartz vessel has an inner radius of 65 mm, or 2.559 in.), in order to cover most of the interior of the quartz jar effectively. The quartz vessel and cameras, hanging from the same aluminum disk, were carefully leveled to within 0.01 in. with a digital level to ensure the spheres were hanging as perfectly vertically as possible within the vessel.

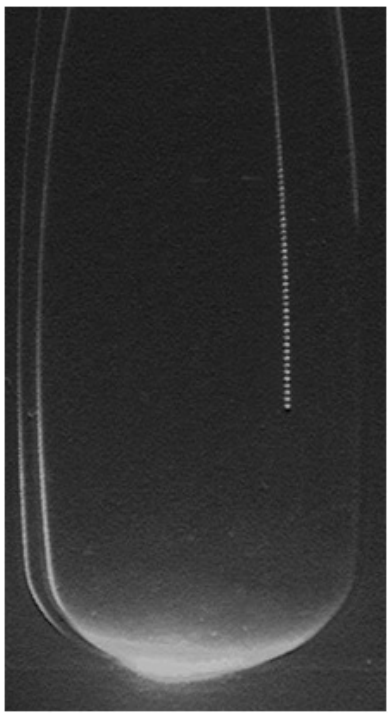

Figure 6.1: Example image of steel spheres used for position calibration in the 15 kg chamber hanging underwater within the quartz vessel. Note the bend caused by lens distortion; in reality, this chain of spheres hung straight down. The spheres were ∼4 mm in diameter.

A copy of the final jar made of natural instead of synthetic quartz, filled with tap water, was used to submerge the spheres, so as to not introduce dirt or scratches to the final quartz vessel. The size of the spheres had been chosen so that in all cameras and at all distances, the individual spheres could be resolved in the 640×480 images, but just barely, because

the smaller the spheres, the finer and thus better the calibration. Small white Teflon flags tied around links in the chain served as additional points of reference, useful for pinpointing the same location as viewed by different cameras.

The cameras, mounted on pairs of parallel aluminum bars, one fixed and the other mobile, the mobile one moving through an arc-shaped slot, were set to their final positions before beginning calibration. This approximately centered the chamber in the field of view of all cameras, lending greater simplicity to the optical reconstruction. The quartz vessel itself was mobile, in an oval-shaped hole through its support plate. The quartz vessel had to move back as far as possible to ensure no blind spots in the cameras. This made it off-center inside the pressure vessel, which was taken into account in the neutron calibration simulation (Chapter 8).

6.2 Mathematical Challenge of Reconstruction

The problem of three-dimensional reconstruction is one of taking a pair of two-dimensional coordinates from two cameras, and extracting the three-dimensional information. Thus, there are three unknowns but four input variables, so the problem is solvable, if not in fact over-constrained. One camera is of course insufficient because the images are 2-D. One pixel corresponds not to a point in space, but to a ray (or a 1-D curve, depending on lens distortion) originating at the center of the camera, or in actuality a tube since pixel resolution is not infinite. With two cameras, the intersection of two rays defining a unique point in three dimensions can be found.

6.2.1 Review of Attempted Methods

Initial sphere position reconstruction attempts using ray-tracing were unsuccessful. Assuming no quartz or using the thin-wall approximation, hoping for a slight, almost parallel shift in incoming rays caused by the wall, was invalid, the quartz evidently being significant. Tak-

ing the inner and outer walls of the quartz vessel fully into account still resulted in failure, most likely due to the non-linear distortion of the fish-eye lenses. Such lenses were chosen in order to see as much of the chamber as possible, so that as few cameras as possible could be used per column. The increase in refractive index of oil over air (the oil used had n=1.4655 at 20°C) already caused a severe reduction in field of view. An empirically-based model suggested by Trucco & Verri (1998) using a 4×4 matrix to take the four numbers from the two cameras and convert them into three numbers had inaccuracies as could be as high as 2–3 cm for events near the wall, important events to reconstruct properly. Therefore, an alternative method was attempted: Lagrangian optics.

Fermat's Principle states that light always minimizes travel time. Lagrangian optics is the application of the same formulae from Lagrangian mechanics to optics, applicable because a minimization of light travel time is mathematically equivalent to the minimization of action in mechanics. This new method worked well (Figure 6.2). The non-trivial application of Snell's Law of Refraction to the cylindrical surfaces of the quartz vessel inner and outer edges was handled following the treatment of Saltiel & Sokolov (1982) and Gerald Onuoha (see Appendix D). Unfortunately, it proved too time-consuming to nail down all of the free parameters for each camera, such as subtle angular offsets and CCD misalignment vis-à-vis the lens.

For the next strategy, the horizontal sphere position in the 2-D image as a function of the vertical 2-D position (the vertical axis was along the length of the cylindrical chamber and the 640 pixel-long axis of the 640×480 images) was fit well with a quadratic or quartic function. The coefficients of the fits were functions of the 2-D horizontal angle between a line of spheres and the central axis of a camera. Toward the center of a camera, real-life straight lines appear nearly straight, but distortion increases toward the edges of the field of view (Figures 6.1 and 6.2). Cubic fits were found to be best for the 3-D real-world vertical

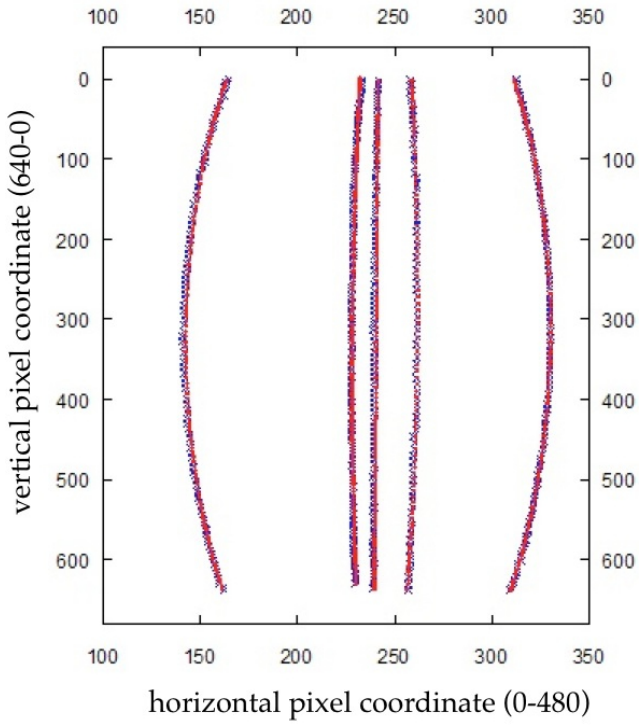

Figure 6.2: Sample result from LabVIEW program written for applying the Fermat/Lagrangian optical reconstruction method to one of the cameras comparing prediction (red) with data (blue). All discrepancies are only a few pixels at most. Note curvature of what should be straight lines, as in Figure 6.1. Five example chain positions are displayed out of the 53 possible: the center of the chamber and four points near the walls, approximately equidistant from each other and from the center.

coordinate versus the 2-D vertical pixel coordinate. The coefficients were a function of the 2-D distance from the center of the line of spheres to the center of a camera. This strategy of utilizing polynomial fits with variable coefficients was found to be much less time-consuming than an application of Fermat's Principle, but also less accurate, leading to errors as high as $\sim$1 cm. So it too was abandoned in favor of finding a method both efficient and accurate, though parts of this method were still valuable (Section 6.2.2).

6.2.2 Optimum Reconstruction Solution

The final method attempted was to use images of calibration data with the fewest number of preconceptions, instead of a first-principles approach to match calibration data with prediction. Given two pairs of pixel coordinates, one from each camera in a row, for each pair the calibration data was used to create a set of possible 2-D points in the chambers that would appear with those coordinates in that camera. These corresponded to the 2-D coordinates where the chains were placed and photographed during calibration. A quadratic spline connected that set of points (one set per camera) and the intersection point of the two splines was found. (Again, the fish-eye distortion made linear ray-tracing a naïve approach.) This would only provide the horizontal real-world coordinates (x and y). The vertical (z) coordinate was calculated in the same manner as the previous method. This divorcing of the flat coordinates from the height allowed this method to be very simple and non-computer-intensive, so that the locations of the tens of thousands of bubbles from the actual running of the chamber could be efficiently pinpointed, but accuracy comparable to the best achieved with the Lagrangian approach was retained.

6.3 Computerized Bubble Finding

An automated bubble finding and counting algorithm was written for the 15 kg chamber. It was similar to what had been used previously on the 2 kg and 4 kg chambers, and adapted from the sphere-finding algorithm, but uniquely suited to address the challenges of the new lighting scheme (Section 3.1). Furthermore, due to a need to decrease illumination to prevent overheating, the bubbles were not as bright as originally planned, and the implementation of the piezo trigger as a replacement for the original optical trigger (Section 5.6.3) caused the bubbles to always be captured at least one camera frame late. Therefore, the bubbles were not always well-formed circles in the images, and the bottom lighting could lead to bright

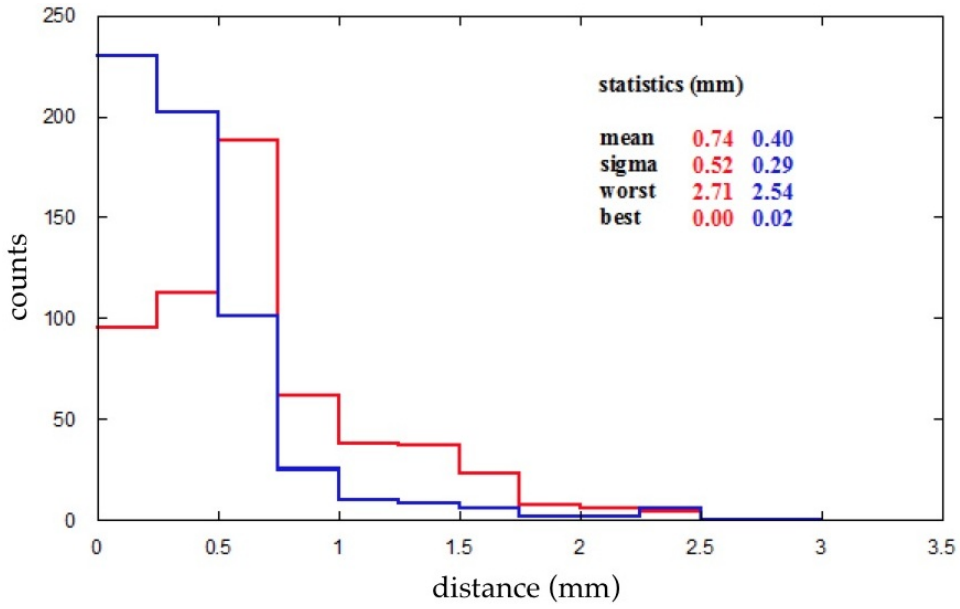

Figure 6.3: Histogram of the radial error in the position reconstruction of the centers of the calibration spheres, using the Lagrange method (red), and using the spline method (blue). The spline method is not only more accurate by several different measures, but it is computationally more than an order of magnitude faster, and is free of a need to set any free parameters.

reflections within the quartz wall that could be mistaken for additional bubbles. The length of the chamber and bottom illumination style created the further problem of brightness gradient, bubbles most easily seen at the bottom. The bubble finder software therefore had to be carefully tuned to work for each individual camera. A built-in LabVIEW image-acquisition (IMAQ) function for counting objects in images was implemented, with two free parameters, threshold in pixel brightness and number of pixels comprising the object.

6.3.1 Features of Bubble-Finding Algorithm

Given the challenges enumerated here, the basic plan had to be augmented. A reasonable set of ranges for the free parameters, determined by manual inspection, was systematically examined by means of nested program loops. The mode, i.e., most common number of bubbles found, was chosen as the bubble number in a particular image. The use of mode

bubble number as opposed to mean reduced the risk of a non-integer bubble number being returned by the program and decreased error because this was most probable number of bubbles. Similarly, the use of the mode bubble coordinates prevented outlying errors from being included, as they would have been in an average. These parameters and conditions would be tweaked by the program itself in the event of self-detected errors. Examples included differing numbers in two cameras in the same row capturing the same event, and bubbles found at disparate vertical locations. Cameras within the same row were level to within 5–20 pixels, so bubbles should have been observed at the same height. The program was therefore self-corrective to a degree.

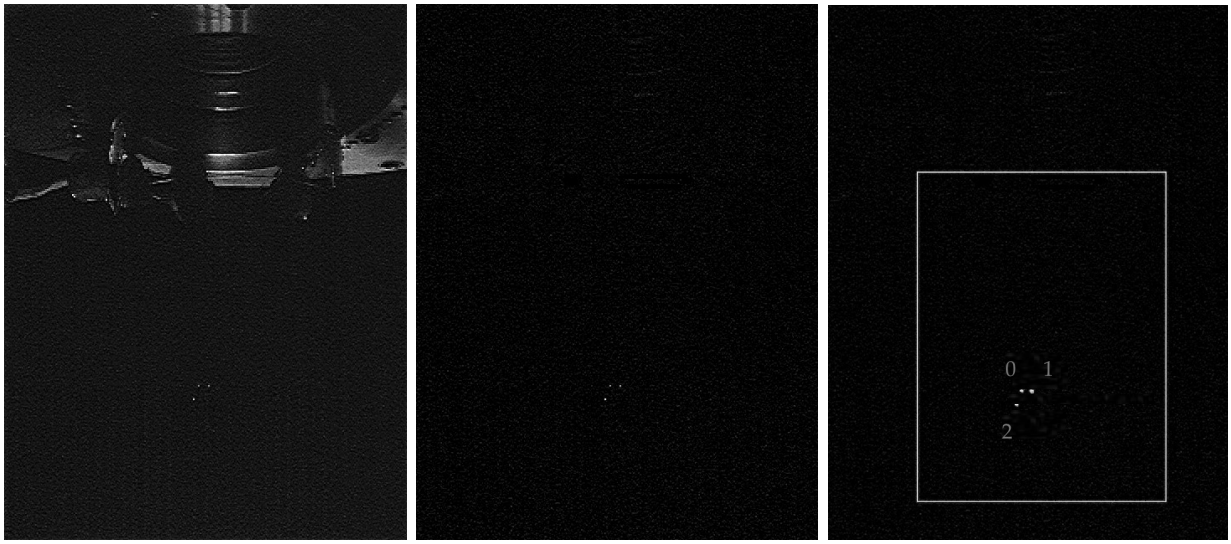

Figure 6.4: (left) Image of a closely-spaced three-bubble event in top camera, under unfavorable conditions: deep in the chamber far from the cameras and at high pressure. High pressure corresponds to a low degree of superheat: the superheated liquid is closer to equilibrium and so bubbles grow more slowly.
(center) Same image after subtraction of no-bubble template. Subtraction makes the bubbles the only significant clusters of bright pixels clearly remaining above the digital noise.
(right) Bubble locations as successfully found in the noise-subtracted image with bubble finder program written in LabVIEW. Bubbles are numbered by their height in the chamber for purposes of 3-D reconstruction. Bubbles assigned the same numerical tag for different camera images should be the same bubble.

To find bubbles with pixel subtraction it was necessary to create template images with

no bubbles (one per camera). Such templates were created by utilizing images from zero-bubble events caused by false triggers. This template creation was necessary because the LEDs would not be active until a bubble was detected by the piezos: images saved before the existence of a bubble were completely black, so that it was not possible to subtract adjacent images, with and without a bubble, as originally had been planned. The template images were slightly brightened to minimize the chance of digital noise being counted as bubbles. To avoid the opposite problem, failing to find bubbles, all images in the ten-image buffer were inspected, so that bubbles could be found late but never missed. The buoyant bubbles rise straight up through the CF_3I until they hit the water at the top in late frames. Early frames are of use in avoiding counting multiple bubbles as singles due to mergers: the short neutron mean free path can cause multiples to be only centimeters apart. Furthermore, early times do not show signs of bubble fragmentation and shape distortion, which occur sooner at a high degree of superheat, when nucleation is more violent.

6.3.2 Vetting of Bubble Finder

Over 300 events were selected to vet the software, spanning all classes: bulk single, multiple, cylindrical wall, hemispherical wall, water interface, false triggers from interface undulation, and unknown-cause false triggers. Ultimately, however, after analyzing thousands of bubbles it became clear that some hand-scanning would still be necessary. Problematic events included high-multiplicity events, multiple bubble events where one bubble was much closer to the camera than the other and caused a single bubble classification, and false positives caused by LED flickering. Images from the bottom cameras were cropped to remove LEDs, but the LEDs could still cast faint reflections high up in the quartz vessel or on the aluminum plate holding it up. These problems could prevent the accurate determination of bubble location.

6.3.3 Bubble Finder Result

Pixel coordinates calculated by the bubble finder were then fed into the 3-D reconstruction algorithm outlined above, camera by camera, bubble by bubble. Coordinates of bubbles visible to adjacent camera rows were averaged together, so a seamless map could be made across the top, middle, and bottom rows.

The first attempt to analyze real data resulted in anisotropy in an x-y (collapsed-z) view, and in reconstruction of bubbles outside the chamber. However, this was easily corrected by shifting both horizontal and vertical coordinates inward by 3%. This took into account mineral oil refractive index change between 30°C and room temperature and the difference in refractive index of quartz because of temperature and use of synthetic instead of natural quartz. The center of the vessel also had to be corrected: care was taken not to shift camera positions relative to calibration, and the final quartz vessel should have been placed in the same location as its natural counterpart, but the cameras or quartz jar still managed to shift.

6.4 Hemisphere Activity Issue

As will be seen in the fully reconstructed chamber maps (Figure 6.7), there were far more events in the hemisphere than elsewhere. During the operation of the 15 kg chamber the author observed the hemisphere was much more active in terms of nucleation than any other region, just based on a visual inspection of events live on a computer screen during running. The history of the hypotheses of causes, and attempts to solve the issue, unfolded as follows.

6.4.1 Realization of Hemisphere Problem

In order to achieve the best WIMP limits, going lowest in WIMP mass, it is advantageous to achieve the lowest possible degree of superheat, right above the beginning of significant

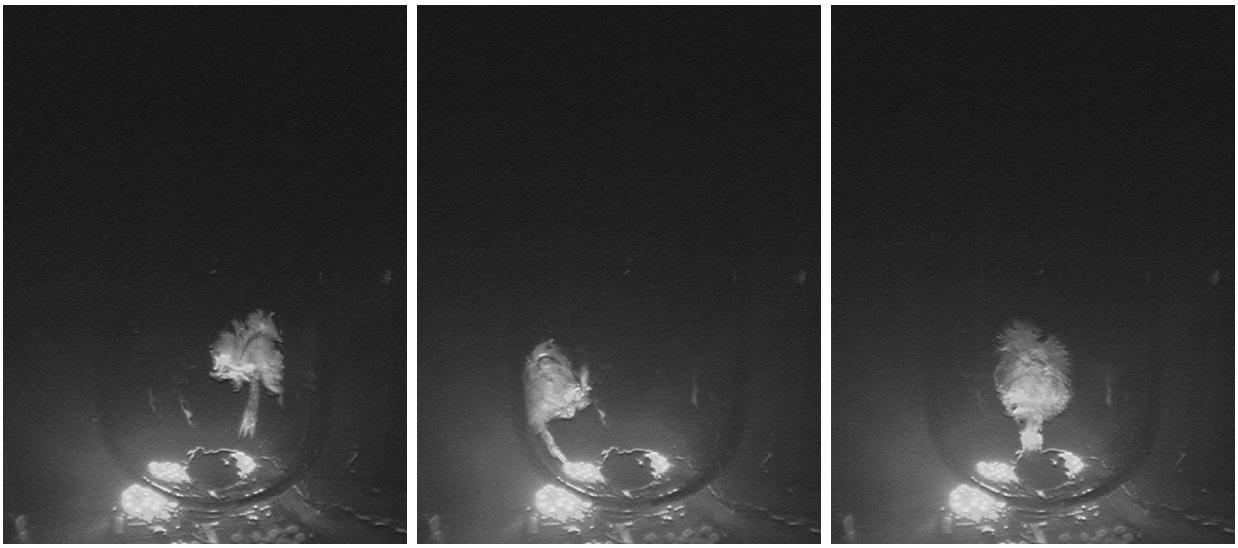

Figure 6.5: Examples of the appearance of hemisphere events. The exact nature of this class of events remains unknown. 3-D positioning confirmed the visual intuition that these most likely originated exclusively from the surface of the quartz hemisphere (see Figure 6.6). Bubbles formed in regions of LED glare are still reconstructed. They must simply rise and shift prior to successful detection by the bubble finder program.

MIP sensitivity. The bubble chamber acts more ideally the higher the degree of superheat, but more specifically the higher the temperature as opposed to the lower the pressure. Ideal behavior is defined by Seitz model predictions, which include 100% efficiency and sharp thresholds. Greater superheat also makes bubbles grow larger and faster, making them easier to trigger on visually as well as audibly, because the bubbles are more violent events. Therefore, the 15 kg chamber was originally meant to be operated at 40°C.

During attempts to run the 15 kg chamber at 40°C, incessant boiling apparently coming off the wall of the hemisphere, mushroom-cloud shaped (Figure 6.5), prevented decompression to a desired pressure. It was impossible to go below ∼30 psig without a hemisphere event occurring during decompression. The first hypothesis was that this was simply chamber outgassing, either a release of CF_3I, or of contaminant non-condensable gasses, trapped in cracks (Section 3.3). If this were true, then the hemisphere event rate would go down rapidly over time as the gas was depleted. However, this was not the case during several

days of observation. It was possible that the surface was somehow damaged, covered with cracks of order the critical radius. For instance, in the 2 kg chamber a dropped tool left a scratch that lead to obvious inhomogeneous nucleation (Behnke et al., 2008). However, there is no record of such an occurance for the 15 kg chamber. The synthetic quartz jar was sonicated, rinsed, and assembled without incident at FNAL, and a natural quartz dummy jar was used for the position calibration outlined in Section 6.1 as well as the surface engineering run at the University. Visual inspection of the vessel during the assembly did not yield any evidence of the hemisphere appearing different from the majority cylindrical portion of the vessel. Furthermore, the problem was not observed during surface running with the dummy jar.

6.4.2 Search for Cause of Hemisphere Activity

The chamber was inactive for weeks after water filling prior to CF_3I filling due to scheduling delays. It was proposed that a great deal of water vapor had become trapped in the wall. However, this cannot explain the issue. First, the level of the 4 L of water rose significantly above the hemisphere, yet only the hemisphere was overactive, as seen visually and later confirmed by reconstruction of the bubble position (Figures 6.6 and 6.7). Second, the 4 kg vessel was filled with water and inactive for over a year, but did not experience a hemisphere problem. Third, as discussed above, trapped gas is finite and is eventually depleted. Even if a great deal is trapped, so much that it cannot be depleted over the course of even months and years of running, then at the very least the rate of activity should have declined over time. Such a decline was not observed in the 15 kg chamber. Fourth, water was the buffer fluid above the CF_3I and was intended to coat the walls regardless, as this should prevent CF_3I gas being trapped behind liquid (see Figure 3.3).

A seam between hemisphere and cylinder, like any seam, could be a source of inhomogeneous nucleation. Though the chamber was supposed to be seamless, imagining that

the hemisphere was somehow separate, accidentally made of natural instead of synthetic quartz, the wall rate for such a scenario was calculated. However, it could not account for the high hemisphere event rate. The hemisphere events, qualitatively appearing like other wall events, along the cylinder (contrast Figure 3.3 with Figure 6.5), appeared to be more violent than the other wall events, thus something different in nature, especially at low pressure (high superheat). Finally, the 2 kg wall event rate was nearly entirely explicable with alphas (Section 3.3.2), yet 15 kg hemisphere events were recorded even at 60 psig, which, according to Behnke et al. (2008), corresponds with too high of a threshold for alpha nucleation at 30°C (Chapter 9), and the effect appeared to worsen lower in the hemisphere (Figure 6.6). That was impossible to explain appealing to uniform alpha emission from bulk quartz.

Instead of alphas, perhaps dust or other large particulates falling down from the top and settling at the bottom of the vessel were responsible, although this should have been impossible given the filtration system used for filling (Section 4.5.4). Even if possible, the violent gaseous ejection of a hemisphere event (Figure 6.5), especially large at low pressure (high superheat), should have been sufficient to eject the offending objects into the bulk, yet no evidence of inhomogeneous bulk nucleation was observed (Figure 6.6). Surface damage may be the correct explanation instead of discrete particulates: the high-pressure water rinse discussed in Section 4.5.1, pointed at the hemisphere with jar inverted, could have been responsible for such damage. AFM (Atomic Force Microscope) inspection of rinsed quartz samples proved inconclusive, though.

6.4.3 Solutions Enacted for the Hemisphere

The hemisphere problem appeared to scale with degree of superheat, because running at 30°C ameliorated it to the point that the chamber was usable. Strictly 30°C data sets are used to build the WIMP limits (Chapter 9), removing the hemisphere with a fiducial

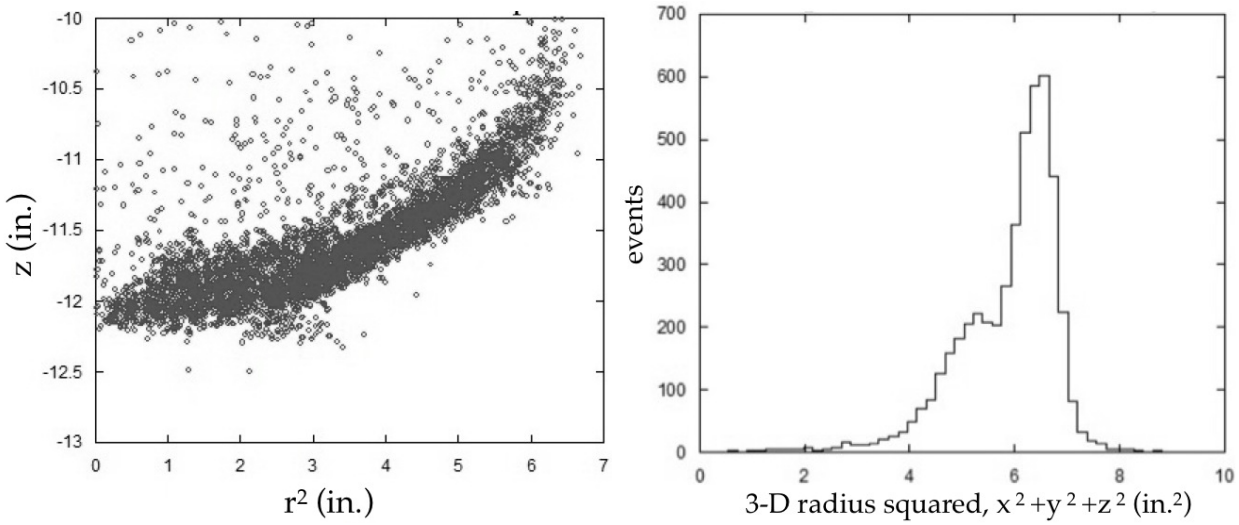

Figure 6.6: (left) Side-view profile of the reconstructed hemisphere. This plot of bubbles vs. squared radius provides evidence for the claim that the events are a "skin" on the surface and not in the bulk. The thickness of this skin is due to imperfect reconstruction: the mushroom shape and rapid growth cause these bubbles to be pinpointed by the bubble finder after having already risen slightly off the surface of the quartz vessel.
(right) Histogram of hemisphere events in radius squared, providing more evidence that the problematic hemisphere events are not in the bulk fluid. Notice clear excess at high radius, corresponding well with the general wall cut, for the cylinder, in Figure 6.9. The drop-off at the left is due to the hemispherical shape resulting in a vanishingly small radius lower in the chamber, corresponding to a density reduction towards the center.

volume cut (Figure 6.8). Furthermore, to explore the possibility that the LEDs, so close to the hemisphere, were overheating it, the LEDs were turned down and a piezo trigger was implemented to reduce their duty cycle so much so that they were almost never on (Section 5.6.3). This had the effect of nearly doubling the livetime of the 15 kg chamber, which was poor due to the hemisphere problem. It increased from 40 to 75%, allowing for a greater exposure for the duration of the TARP run. At the lowest pressure data point alone (3 psig), the mean superheat time increased from ∼1 to ∼30 s. This would apparently indicate that the LEDs were overheating the chamber and creating a thermal gradient. However, Figure 6.6 shows that hemisphere bubbles appeared to be coming from the surface of the quartz vessel, not the bulk. An increased temperature would have lowered

the recoil energy threshold, leading to a higher bulk event rate. At a sufficiently high degree of superheat, the superheated fluid becomes sensitive to betas and gammas (Section 2.5). Nevertheless, hemisphere activity increased with degree of superheat: the rate increased almost exponentially with lower pressure.

In the end the hemisphere phenomenon, a concern because of the possibility of recurrence in future chambers and a worsening with larger surface area, could not be satisfactorily explained. However, the 4 kg chamber did not suffer from the problem, and the 60 kg chamber does not appear to either as of time of writing. For the case of the 15 kg chamber, the hemisphere was placed under control given the measures outlined above, and then the remaining events were handled with a fiducial cut.

6.5 Spatial Reconstruction Results

Ultimately the 15 kg chamber was well-reconstructed, showing cylinder and hemisphere with correct radii. A thick skin of events was observed at the water interface (Figure 6.7) but this, like the hemisphere, was removed with fiducialization (Figure 6.8). Unlike the hemisphere problem, this was expected, as it is a feature of all COUPP chambers. Alpha-emitting isotopes could easily become trapped by surface tension at the CF_3I-water boundary: commonly formed amphiphilic (attracted to both water and oil-like substances) uranium and thorium salts are known to accumulate at such an interface (Collar et al., 2000; Pan et al., 1999). Another fact about the reconstructed interface is that it displays a slight tilt in the reconstruction, possibly due to an uneven floor in the room, a tilt in the cameras, or undulations at the water interface caused by the presence of a bubble. Other chambers have encountered similar tilts (Behnke et al., 2011; Szydagis, 2008).

Assuming 2.046 g/cm^3 to be the density for CF_3I (NIST REFPROP software) at 30.59°C, the average temperature of the 15 kg chamber during the TARP run, the reconstructed

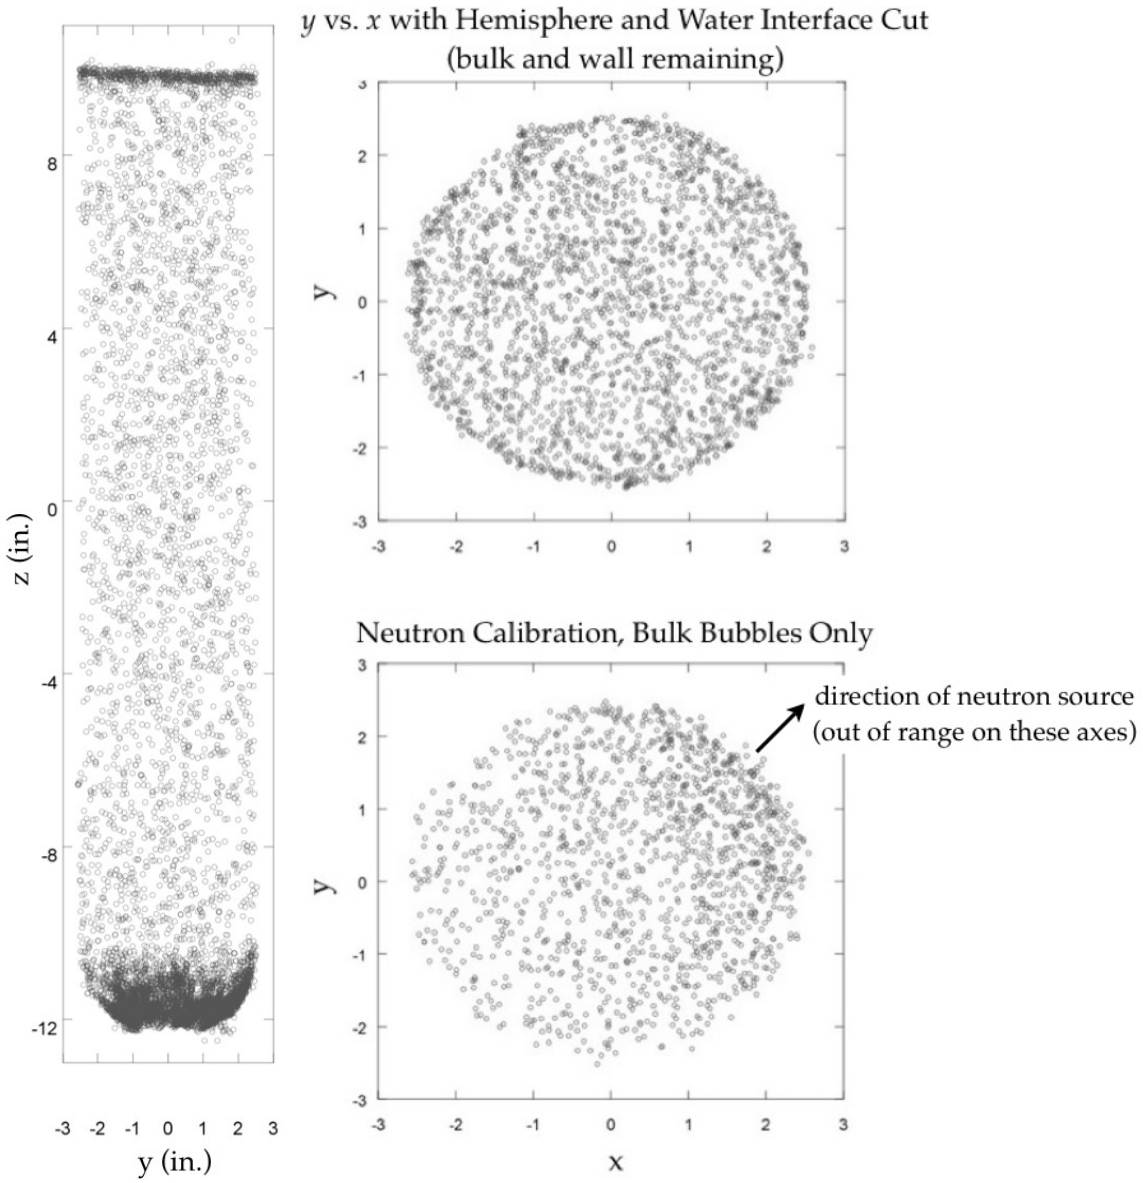

Figure 6.7: Reconstruction map of all single bubbles from a 400 kg-day exposure: z vs. y (left), with water boundary evident at $z\simeq$+10 in. and hemisphere beginning at $z\simeq$-10.5 in., and y vs. x (upper right; collapsed radial view without regard to z, not a slice), after fiducial cuts to remove hemisphere and top layer. A slight wall event crust is barely visible. Statistical checks yield no anisotropy indicative of poor reconstruction or inhomogeneous nucleation. y vs. x at lower right is a display of events occurring during calibration with a neutron source (Chapter 8). Bubbles occurred more often near the source, located outside of the chamber and pressure vessel in the upper right quadrant: neutrons have a mean free path of a few cm in CF_3I. The center of the z-axis is defined approximately as the center of the field of view of one of the middle cameras, determined by the calibration in Section 6.1.

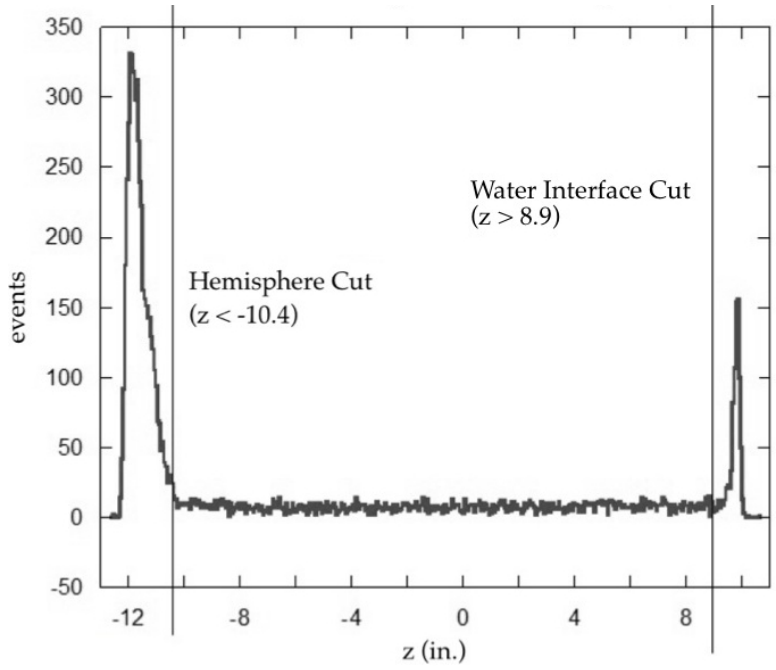

Figure 6.8: 1-D, vertical-axis fiducial volume cuts applied to histogram of single bubble events. The cut for the interface appears more conservative than the hemispherical cut, but that is because the hemispherical cut is supplemented by the general wall event cut to remove unwanted events (Figure 6.9, compare to Figure 6.6, right).

chamber (Figure 6.7) has a gross mass of 14.76 kg, in excellent agreement with the estimated 14.67 kg transferred during filling (Section 4.5.4), despite the errors in reconstruction, worst for the z-axis (Section 6.5.1). This density based on mean temperature was used to calculate the fiducial volume based on the cuts defined by the optical reconstruction (Figure 6.6). Because only the bottom thermometer was functional for the duration of the run (Section 5.2), the vapor pressure helped determine the accuracy of the temperature (Figure 8.6). The fiducial volume was thus determined to be 10.75 $\pm$ 0.20 kg. (The source of the estimated uncertainty is the position resolution discussed next in Section 6.5.1.) The CF_3I density calculated using the NIST REFPROP program was again used when an additional fiducial cut became necessary to ensure the quality of the audio signal used for particle discrimination (Chapter 7).

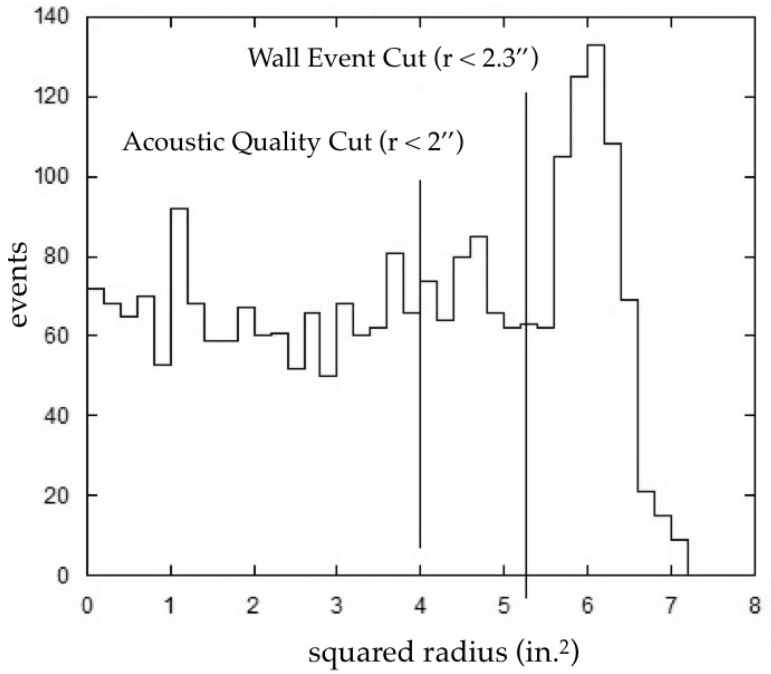

Figure 6.9: 2-D, radial fiducial volume cuts applied to histogram of single bubble events. Hemisphere and interface have already been removed. The flat portion (within statistical fluctuation) of this histogram, plotted versus radius squared, is consistent with equal probability of events per unit volume. Neutrons may cause a slight deviation, being less probable towards the center (see Figure 2.6), but alphas in the bulk should not be unequally distributed. The so-called "acoustic quality cut" is defined next, in Chapter 7, adopted because of poorer signal quality close to the quartz wall.

6.5.1 Resolution of the Position Reconstruction

To determine the effective resolution of position reconstruction, the correct position of the jar edges was compared with the reconstructed position. For this purpose a small number of points definitively on the wall of the chamber were selected based on observations of bubbles clearly emanating from the wall (see Figure 3.6). The systematic error in determining the radius of the cylinder was quite low, 0.003 in. However, the standard deviation of the radius in cylindrical coordinates of these points was 0.077 in., implying the reconstruction had a negligible systematic error in determining bubble position, but a significant statistical error. Another handle on the statistical error was an examination of the bubbles in the field of

view of more than one camera row. Based on such measurements, the error was determined to be 0.011 in. Thus, the reconstruction had a resolution of ∼0.1 in., radially. Vertically, the resolution was poorer, ∼0.3 in., evident by the apparent thickness of the water boundary. However, this was acceptable because the radial coordinate was more important. It was necessary to classify wall events and remove them with a fiducial volume cut, without the benefit of the dense skin of wall events in the 2 kg chamber data (see Figure 1.8). The thick crust of interface and hemisphere events made it possible to easily remove these classes of events even given the poor vertical (labeled z-axis) resolution.

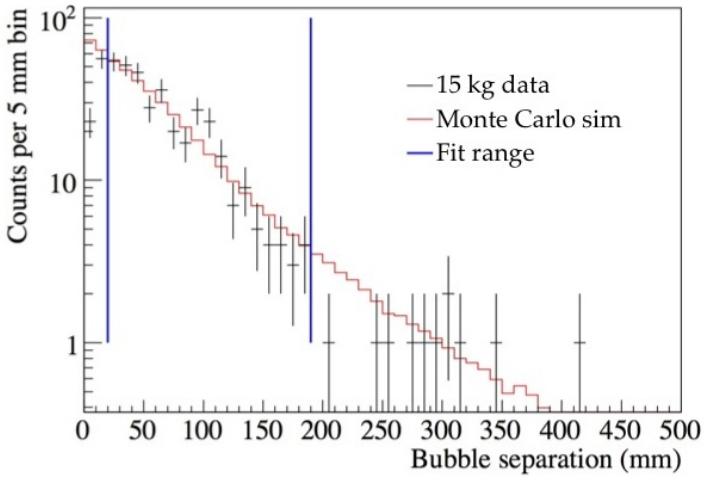

Figure 6.10: Histogram of the distance between double bubbles, both with and without the neutron source, with best exponential fit determined by Monte Carlo simulation shown for the portion excluding the highest distances with poor statistics and the aberrant first bin. According to this data the neutron scattering length was ∼5 cm. The first low bin is explained by mergers (Section 6.5.1). Multiple bubble events are due to a single neutron scattering multiple times. Figure courtesy of Hugh Lippincott.

One caveat to the statement of a radial resolution of ∼2.5 mm was the fact that two bubbles events where the pair of bubbles was less than ∼3.5 mm apart were reconstructed as single, most likely from a physical merger, or from too great a sharing of pixels. This is evident from the unphysical dip at short distance in the histogram of double bubble separation that would otherwise be exponential (Figure 6.10).

6.5.2 Quartz Surface (Wall) Nucleation Rate

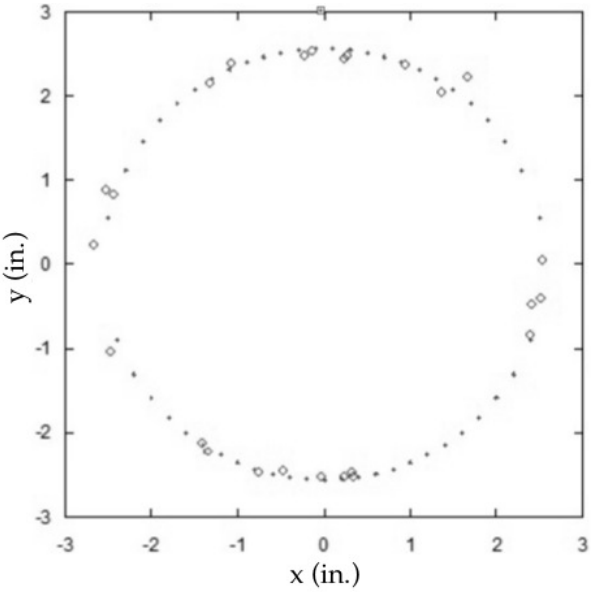

Figure 6.11: Wall events used to determine spatial resolution (large annuli). The dotted circle represents the actual position of the inner edge of the quartz vessel. The edge is the most reliable position reference. The same approach was used for the 2 kg chamber, but with less difficulty because of the greater number of wall events in that chamber. Only bubbles definitely evolving from a clear fixed point on the quartz surface (see Figure 6.7), or observed to touch the surface and leave an artifact (a malformed half-bubble stuck to the wall), were hand-selected from all data. Several of the events depicted were visible for multiple camera rows, chosen in order to determine the difference in bubble position determined by different cameras. In this manner the mean error was found to be ∼0.1 in. (Section 6.5.1), the high end of the expected error range (Figure 6.3).

Histogramming number of events versus radius squared, which should be a flat distribution if bubbles are equally probable in equally-sized slices of volume, showed a clear excess near the walls (Figure 6.9). However, an estimate of this excess yields a wall rate too high given the uranium concentration in synthetic quartz, but within a factor of two of estimates performed by visual inspection (Section 3.3.2). A surface rate of ∼0.004/(cm^2×day) was observed, whereas 10^{-4}/(cm^2×day) was expected based upon the 21 ppt U measured in Suprasil (Leonard et al., 2008). However, it is possible that the ∼0.1 in. radial resolution,

poorer than in previous chambers (e.g., $\sim$O(1mm) for the 2 kg chamber), caused by the distorted optical view, allowed bulk events to bleed into the wall event candidate region. In addition, as discussed in previous chapters (Chapters 1, 2, and 4), wall events may have origins other than intrinsic alpha-emitting nuclei, such as microscopic surface deformations or implanted alpha-emitters originating from atmospheric radon. Therefore, it is not surprising that the measured rate exceeds the predicted rate because the predicated rate was only based upon one knowable cause. Ultimately, given the apparent 2–3 order of magnitude reduction in wall rate over the 2 kg chamber, the 15 kg chamber demonstrated that, barring the re-appearance of the unexplained hemisphere issue, COUPP will be able to construct chambers of ever larger surface area without dead time from wall events being a serious issue. This was a positive step forward for the bubble chamber technology as a viable form of direct dark matter WIMP detection.

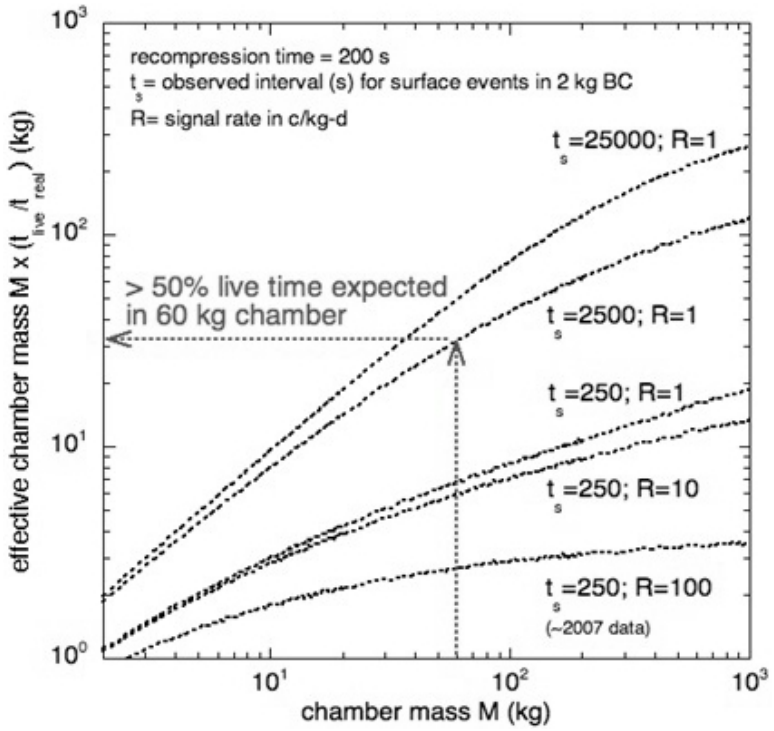

Figure 6.12: Plot of "effective" chamber mass (based on the ratio of livetime to total time) vs. actual gross mass, made after the 2 kg chamber run to extrapolate the wall event problem to chambers of increasing surface area. Emphasized here was the need to find a solution to the dead time problem generated by wall nucleation that reduced the rate by at least an order of magnitude in order to achieve 50% livetime for a 60 kg chamber. The 15 kg chamber achieved and exceeded this goal, with the exception of the hemisphere (Section 6.4), but the 60 kg chamber does not exhibit that unusual behavior thus far. Figure courtesy of J.I.Collar.

CHAPTER 7
ACOUSTIC DISCRIMINATION

Being extraordinarily MIP-insensitive, the only significant background a dark matter bubble chamber must either eliminate or identify is alphas. This assumes cosmogenic neutron events are vetoed, and environmental neutrons are rejected by bubble multiplicity, or by fiducialization for the case of large chambers. The PICASSO collaboration has recently demonstrated that the ultrasound signature of a bubble may be useful in differentiating alpha events from neutron or WIMP recoils (Aubin et al., 2008; Archambault et al., 2009). The 15 and 4 kg chambers have had marked success in reproducing these results, although presently only at low degrees of superheat. The 15 kg chamber was built before the seminal PICASSO publication, so the piezos were not optimally placed. Alpha-neutron discrimination was simply not under consideration. Nevertheless, excellent discrimination was demonstrated at 3–15 psig. Beyond those pressures, bubbles could no longer be "picked up," especially when distant relative to the piezo mounting locations.

7.1 Theoretical Background on Acoustics

Acoustic discrimination relies on heightened acoustic energy for multiple bubbles, and the premise that alphas have the potential for making more than one protobubble. The theory behind this is that the light (compared with C, F, and I) and charged alpha travels distances of order tens of microns in the liquid. It leaves a dense ionization track in its wake probabilistically leading to formation of multiple microscopic protobubbles (Aubin et al., 2008), which then merge into the macroscopic one captured on camera. Via Coulomb encounters along its microns-long path, the alpha particle deposits energy within several critical radii ($r_c \lesssim 100$ nm in CF_3I, Section 2.2). The multiple individual protobubbles generate more high-frequency sound corresponding with the early stages of bubble development than a

single protobubble would. That makes alpha-induced nucleation louder. Neutrons just lead to recoiling nuclei of carbon, fluorine, or iodine, which are heavier and so have shorter ranges (e.g., ∼50 nm for a 10 keV iodine recoil), making only single protobubbles per each scatter overwhelmingly likely. Alphas lead to at least two protobubbles, one for the alpha, and one for the recoiling nucleus which emitted it. This is true as long as the nuclear recoil threshold is low enough. Any recoiling nucleus has about the same energy after an alpha decay: recoil energies of an alpha-emitting contaminant after ^{220}Rn and ^{222}Rn in the ^{232}Th and ^{238}U decay chains range from 101 keV to 169 keV. For example, the daughter ^{218}Po recoils with ∼100 keV after a ^{222}Rn emits an alpha. When the recoil of the nucleus after alpha emission is below the energy threshold, the few-MeV alphas themselves can still produce at least one protobubble. However, for a sufficiently high pressure (low degree of superheat), alpha stopping power does not exceed the $\frac{dE}{dx}$ threshold (Figure 2.5).

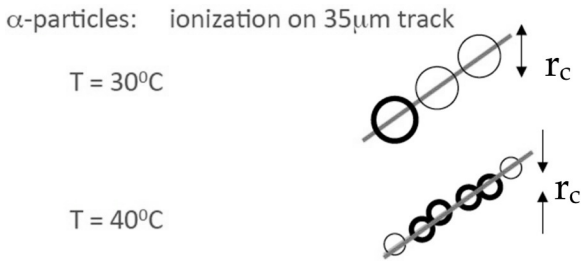

Figure 7.1: Illustration of changing alpha-particle ionization profiles with increasing superheat (higher temperature). As the critical radius (r_c) decreases, alphas are capable of generating larger numbers of individual protobubbles. They generally produce at least two (Section 7.1). Alpha particle track lengths calculated using SRIM (Ziegler, 1984, 2010). Figure adapted from J.I.Collar.

Martynyuk & Smirnova (1991) offer more insight into the physical underpinnings of acoustic discrimination. The pressure of a volume of liquid in which a particle has interacted is intimately related to the energy imparted to this volume. The energy deposition profile of the impinging particle should therefore relate to the acoustic qualities of the bubble formed in a bubble chamber, as the gas will force the surrounding liquid out of its path in different

ways during the earliest stages of protobubble development, which correspond with high frequencies because of the small sizes of the protobubbles at those times.

Given the differences between bubble chambers and SDDs (Superheated Droplet Detectors) (Section 5.6.3) the COUPP collaboration anticipated observing better discrimination than PICASSO. A comparison of Figures 7.4 and 7.5 confirms this expectation. In an SDD, alphas originating in the inactive gel substrate surrounding the superheated liquid, or traversing both the gel and the liquid, will undergo energy loss in the gel. A partial trajectory through a droplet leaves an incomplete energy deposition, increasing the probability of an alpha generating only one protobubble and behaving as a neutron. On the other hand, a bubble chamber has 100% active volume. Only alphas originating from radioactive impurities in the quartz can have a partially liquid trajectory, with the parent nucleus recoiling deeper into the quartz and the alpha exiting into the CF_3I (or vice versa). However, this is not a difficulty since wall events are already removed with fiducialization (Section 6.5.2). The transparency of a bubble chamber constitutes an additional advantage. It enables position reconstruction, which simplifies spatial correction of acoustic signals for amplitude as a function of distance from piezos. The bubble chamber also offers a homogeneous medium, making acoustic dispersion from impedance mismatch inapplicable, whereas SDDs have numerous interfaces between gel and liquid droplets. Lastly, COUPP uses amplified piezos with more bandwidth (Levine, 2010).

7.2 Practical Application of Acoustics

Analysis for the 4 kg chamber was begun before that of the 15 kg chamber. At first, the amplitude of acoustic pulses, corrected for position, was histogrammed parallel to the PICASSO treatment. However, like the PICASSO results (Figure 7.2), this showed only weak separation of neutron and alpha populations. Therefore, a new method was devised

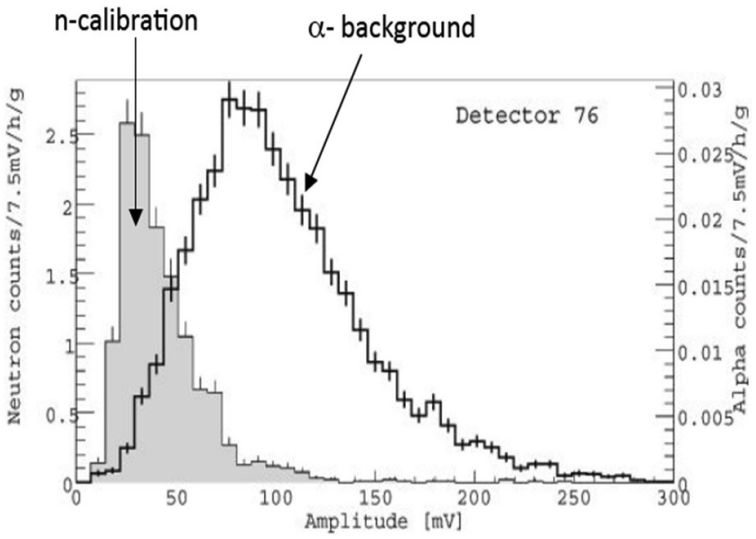

Figure 7.2: Alpha-neutron discrimination as first shown by the PICASSO collaboration in a superheated droplet detector (SDD). This is a histogram of the amplitudes of piezoelectric acoustic traces. The alpha background is thought to appear at an average higher amplitude by generating more protobubbles per event. Figure adapted from Aubin et al. (2008).

where the squared product of the amplitude spectrum and frequency was examined over various frequency ranges. In the 15 kg chamber analysis, the acoustic energy was measured by integrating this product over the first 10 ms after the beginning of a bubble in piezo traces. A longer timescale implies better statistics but may also lead to loss of discrimination power due to inclusion of secondary acoustic effects such as resonances of the piezos or quartz jar. However, it is unclear whether quartz ringing is a significant contribution to discrimination ability. For the 4 kg chamber, there was no difference between utilizing 5 or 10 ms windows. Too long of a timescale also affects discrimination by allowing time for protobubbles to merge and for sound to decay (observe the exponentially decaying sinusoid in Figures 5.4 and 5.6). Seminal work on the 4 kg chamber which was concurrent with the preliminary data-taking stages for the 15 kg chamber established 10 ms as an optimal choice. This was used for the 15 kg chamber, which shared the same radius and wall thickness.

The energy in various frequency bins was considered, and the frequency band with the

best discrimination was selected. However, all bands where a bubble had a good signal to noise ratio provided some evidence for discrimination. The frequency bands 3–25 kHz and 100–250 kHz were examined in particular, comparable to those studied in the 4 kg chamber analysis. For reasons having to do with data acquisition provided in Chapter 6, 25 kHz was the practical limit for the 15 kg chamber, for most data (Section 7.3). Although the discrimination should be better at higher frequencies by incorporating contributions from the sounds of smaller protobubbles, the result of the 3–25 kHz band analysis is itself interesting. It provides more evidence that acoustic discrimination extends into the audible frequencies. A carefully attuned ear close enough to a bubble chamber could thus distinguish between alpha and neutrons. However, the raw amplitude scales with the square root of the number of protobubbles. More bubbles implies more energy, and energy scales with voltage squared.

7.2.1 Position Correction to Acoustic Parameter

"Acoustic parameter" is defined as the energy in units of squared volts per square root of frequency in the bin selected for discrimination. It was normalized so that the single-macroscopic-bubble neutron population, as determined by neutron calibration and veto coincidences, would form a peak at one. To have a narrow, symmetric peak, a correction is necessary for the change in absolute acoustic energy with position caused by unique piezo distances.

The acoustic parameter was a function of radius in the chamber, height, and polar angle. A separation of the position correction into three components was assumed because it was known to work well for the 4 kg chamber. (The radial correction dominated, followed by height.) To check whether such separation was valid, after applying the corrections the acoustic parameter was again plotted as a function of radius, height, and angle to check if these distributions were now flat. At first it was expected that good statistics would permit

a full-fledged 3-D correction. However, in the 15 kg chamber the lack of a discernible initial burst of sound from a bubble in acoustic traces above ∼18 psig limited the accurate determination of t_0. At high pressure (low superheat) the pressure change from a bubble could be observed, but this greatly increased the error in determination of time. That in turn limited the ability to select an appropriate time window for a power spectrum analysis. Therefore, discrimination was achieved for only a subset of data (3–15 psig).

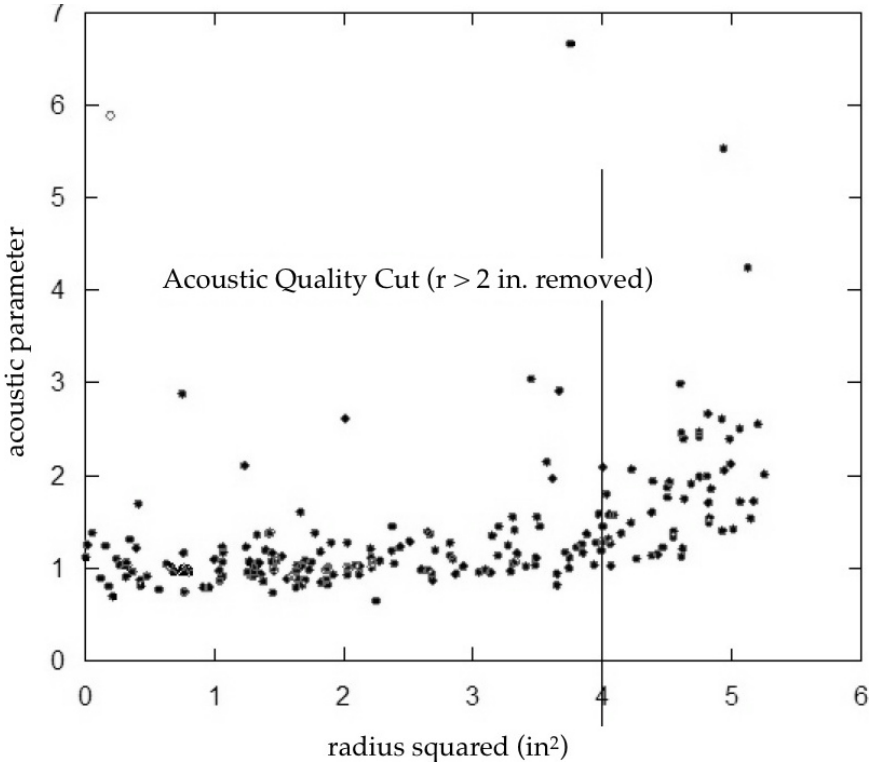

Figure 7.3: The acoustic parameter of likely neutron-induced single-bubble events (as determined by veto coincidence or neutron calibration) versus the square of the cylindrical radius within the quartz vessel. The acoustic parameter has been normalized to the median. The acoustic quality cut shown removes the volume where the width of the band of neutron events has dramatically widened and increased in mean. The origin of this effect is not understood. Proximity to the quartz wall is one possible explanation.

The signals from the four piezos in the 15 kg chamber were averaged together to reduce the position dependence, so only small corrections needed to be made. If the possibility of efficient acoustic discrimination had been known beforehand, more piezos would have been

included, placed along the entire length of the vessel. The camera interference issue would still require a solution (Section 5.4). Fortunately, piezo placement (Figure 5.3) minimized both height and angular effects. The distance between a bubble and the closest point on the quartz vessel was important, more so than the distance between bubble and piezo. This was evident in the radial dependence, likely because sound could have traveled first through the quartz to the piezos instead of directly through the liquid. The sound speed in quartz is ∼5,000 m/s, while it is only 350 m/s in CF_3I (NIST REFPROP). In the 15 kg chamber analysis, events had to be cut several cm from the walls because of a signal degradation that is not understood (Figure 7.3).

7.2.2 Pressure and Temperature Correction to Acoustic Parameter

One can also normalize the acoustic parameter with respect to the pressure. A polynomial spline with a sufficiently high degree was utilized so that the individual pressure corrections were essentially independent for each bin, though an exponential was also a close fit. The median acoustic parameter of neutron-only singles (not mean due to a high-parameter tail discussed in 7.4.1) and median pressure were used in each pressure bin. The temperature varied slightly over the course of the experiment (30.59 ± 0.46°C), but no significant dependence on it was observed. Further discussion is found in Section 7.4.2.

7.3 Data Sets of the 15 kg Chamber Run

The 15 kg chamber accrued several months of engineering data while DAQ parameters were being finalized and glitches in hardware and software or physical problems like the hemisphere and veto malfunctions were isolated. Two major data sets were used for analysis in this thesis: the pressure scan data set (3–60 psig in 3 psig steps) and the dedicated, lowest-pressure data set (3 psig). The "physics-quality" data sets used are distinguished

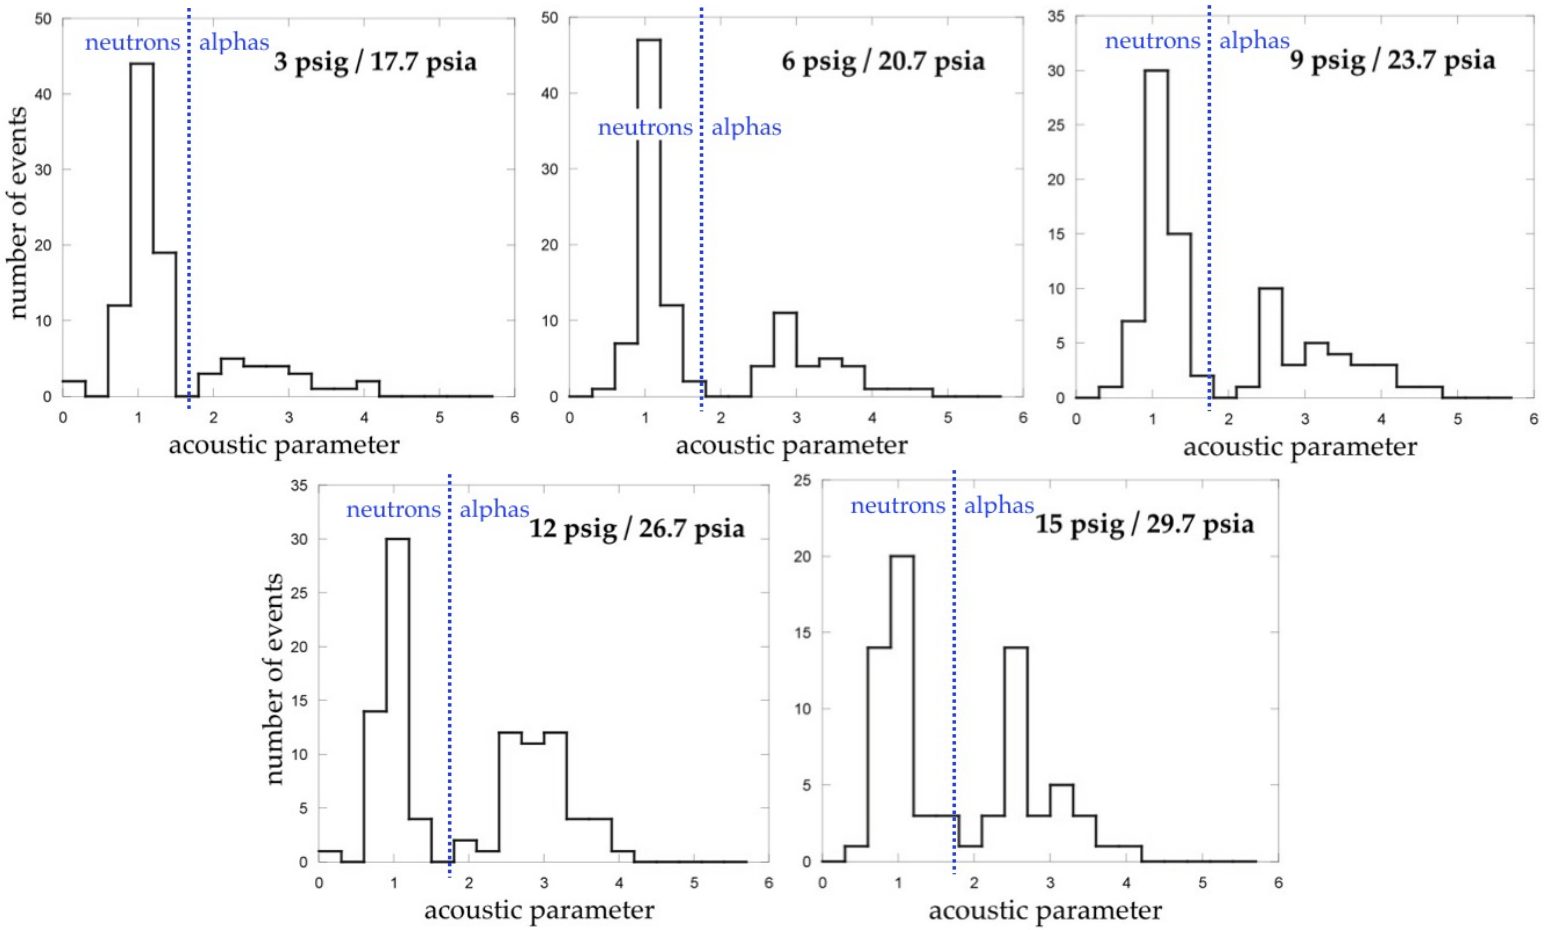

Figure 7.4: Evidence of two populations of events at different pressures shown to likely be neutrons (AP=1) and alphas in Figures 7.5 and 7.6 using veto coincidence to extract known neutron events. Position corrections have already been applied here, in addition to wall, interface, and hemisphere fiducial cuts. The wall event cut has been extended as per Figure 7.3. Neutron calibration data are not included in these histograms, though both vetoed and unvetoed events are. The normalized acoustic parameter should be an indication of the number of protobubbles (Section 7.2 and Figure 7.7). Data presented are exclusively from events comprised of single macroscopic bubbles.

from engineering runs based on the following criteria:

1. Only data using the piezo trigger is studied. The visual trigger involved the LEDs being brighter and on longer, overwhelming the data with hemisphere events. Only data accumulated after the finalization of piezo trigger parameters like thresholds, LED on-time, and hold-off for decompression noise (Section 5.6.3), are included.

2. A glitch in the communication between veto and DAQ causing lost pulses was removed by lengthening the pulses and increasing the digitization rate.

3. Pressure scanning was performed with an attempt to normalize superheat times: initially the same number of events was recorded at each pressure, but this created bias with greater exposure at high pressure (low superheat), the less interesting region where an experiment could only detect a high-mass WIMP with high cross-section likely to already be ruled out. This was a consequence of great low-pressure incidence of hemisphere and interface events. In the new system, more events were recorded at lower pressure, but this served to equalize exposure across all pressures.

These criteria apply to the results presented in this chapter as well as Chapter 9. Neutron source calibration data was collected in dedicated runs, and by building in a 5% chance in all data for the neutron source to activate (Chapter 8).

For the pressure scan, 25 kHz was the digitization rate so that acoustic traces spanning 2 s could be recorded. Thus, even at high pressure (low superheat) where the bubbles grew slowly, t_0 would be included in the traces (see Section 5.6.3). This was unnecessary for the 3 psig run where the bubble would grow rapidly, so that it was sufficient to save acoustic data for only 200 ms, allowing for an increase in digitization to 250 kHz.

7.4 Acoustic Particle Identification

There are several ways to tell whether the two populations in acoustic parameter space truly correspond to neutrons and alphas, at low and high parameter, respectively. First, when the neutron source is activated, the base bulk rate in the chamber increases by one to two orders of magnitude. Neutron calibration data is scrutinized to determine whether all the events cluster around low acoustic parameter. Second, neutron production by muons (Section 5.7) implies that the events tagged by the muon veto should also be neutrons and so be in the low parameter population. Third, the acoustic energy should linearly scale not only with the number of protobubbles, but also with the number of observable bubbles. Therefore, multiple bubble events are checked for the proper scaling (Figure 7.7).

7.4.1 Alpha-Neutron Discrimination

The muon veto and neutron source events were used as the control set. This population of events, which should almost all be caused by neutrons, was used to normalize the acoustic parameter so that one bubble or protobubble would register as AP=1.0. Alphas were not expected to necessarily always produce the exact same number of protobubbles (Section 7.1), so a wide peak was possible for them. They may generate varying numbers of protobubbles along stochastic ionization trails. Alternatively, alphas coming from distinct decays may create unique numbers of protobubbles. The narrower neutron peak in acoustic energy was useful for making the distance and temperature corrections (Section 7.2). Then all of the data could be compared, including unvetoed and neutron source off.

The 15 kg chamber data clearly shows two populations of events when both unvetoed events and events from times when the neutron source (SAMBE) was inactive are included (Figure 7.5). There is a narrow peak at 1.0 corresponding to the neutron population: it matches the solitary peak at 1.0 evident for veto and neutron source data combined. The

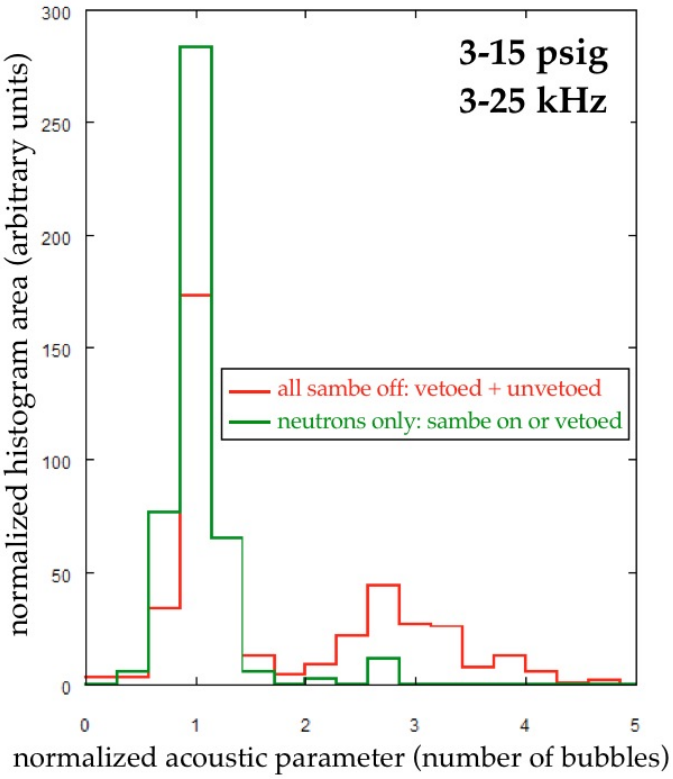

Figure 7.5: Compilation of 3–15 psig data. The acoustic parameter is based on the 3–25 kHz band. A neutrons-only data set based on neutron source (SAMBE) and muon veto coincident bubbles is contrasted with vetoed and unvetoed data (with the source off), which should contain both neutrons and alphas. The areas under the histograms for these two sets of data are normalized to serve as a better gauge of the discrimination power. The appearance of the second peak when unvetoed events are included is linked to alphas by process of elimination. Potential interactions are restricted by the MIP insensitivity (Section 2.5). Pressure bins can be merged because of the weak dependence of alpha protobubble number on pressure (Figure 7.4).

histogram is surprising in that only single bubble events, as determined visually, were used to produce it: yet, when veto anti-coincident data and data for dark matter limits (neutron source inactive) are included, a peak appears at roughly three times the acoustic energy contained within single bubbles. Based on the theory presented in Section 7.1, this is interpreted as one alpha creating three protobubbles on average, which then merge into the single visible bubble. This alleged alpha peak occurs at the same acoustic parameter in the

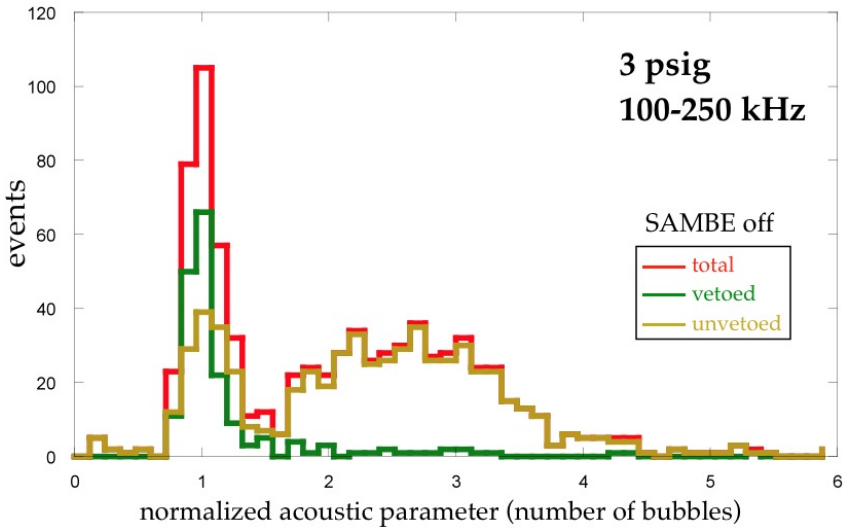

Figure 7.6: Analysis of dedicated 3 psig data. Data from times when the neutron source (SAMBE) was switched off are divided into vetoed and unvetoed events. Vetoed events are likely cosmic muon-induced neutrons. Note the obvious lack of the wide peak at the right in the distribution of vetoed-only data. The remaining unvetoed neutrons will be discussed during Chapter 9.

3–25 kHz and 100–250 kHz bands (Figures 7.4, 7.5, and 7.6).

The soundest explanation for the neutron tail to high acoustic parameter is that an (n,α) reaction is occurring. The ^{19}F cross-section for this reaction is of sufficient magnitude to explain the data (Shibata et al., 2002). The false veto coincident rate is not high enough to explain the integrated rate under the tail. Significant leakage of alpha particles into the neutron population is unlikely, extrapolating from the acoustic behavior of observable multiple bubble nucleation: two bubble events appear to be at least twice as loud as one bubble events, and are often even louder, but never as soft as one, and similarly for three or more bubbles (Figure 7.7). Before the above explanation for the tail was formulated, nuclear spallation by high-energy muons or gammas was considered. After the parent nucleus breaks up, the individual fragments may each produce a protobubble.

The 4 kg chamber had an additional way to produce neutron-induced bubbles, the NuMI neutrino beam at FNAL. Among other particles, this beam can indirectly produce neutrons

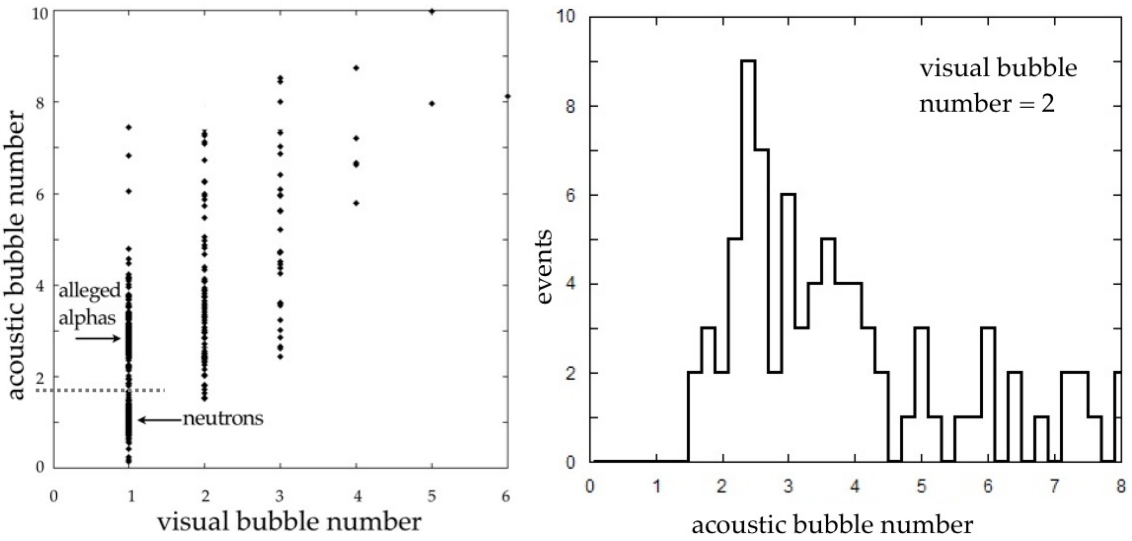

Figure 7.7: (left) Bubble number as determined by acoustic parameter versus bubble number as determined by camera. The acoustic energy clearly increases as the number of bubbles does. The scaling is not perfectly linear, and the visual bubble number is only a rough lower endpoint for the acoustic bubble count.
(right) Histogram of visual two-bubble events alone. The median acoustic bubble number is slightly more than twice that of a single neutron.

in the surrounding rock in the MINOS hall cavern which then interact with the detector. The timing between bubbles was also checked and it was possible to identify certain short-lived isotopes of ^{222}Rn in the chamber early on in the run which were injected during the CF_3I fill. Because the 15 kg chamber was unusable at the beginning of its TARP run due to the worse hemisphere issue at 40°C, there is no early data which could be used to look for radon injection at time of fill. Therefore, these additional checks on the alpha population could not be performed. Timing correlation between bubbles was nonetheless checked for the 15 kg chamber but nothing conclusive was found. The run may have been dominated by long-lived isotopes. Under consideration by the COUPP collaboration is the possibility of doping future bubble chambers with alpha emitters in order to better study the discrimination phenomenon. The current plan is to find short-lived isotopes which will not hinder a dark matter run, or to build a dedicated chamber.

7.4.2 Superheat Dependence of the Discrimination

Nothing conclusive was measured regarding the number of alpha protobubbles versus degree of superheat based on acoustic emission. Because of poor statistics there is no strong evidence that it increased with lower pressure as expected (Figure 7.4). Looking at the temperature dependence of the pressure scan data, there is again no statistically significant evidence that the alpha protobubble count increases with temperature (Figure 7.8).

7.5 Acoustic Signatures of Non-Particle-Induced Events

It is possible to acoustically reject non-particle-induced nucleation as PICASSO does (Archambault et al., 2009). PICASSO rejects events stemming from mine blasts at SNOLAB. For the 15 kg chamber, the hemisphere, water interface, and wall events have unique acoustic signatures compared to bulk events. Some are even louder than alpha-induced events; others are softer than even neutron-induced events.

7.5.1 Hemisphere Events

The hemisphere sounds appear to comprise one large population, both bulk and wall, louder than alphas (Figure 7.9). Two factors may account for this: the inherent explosiveness of hemisphere events and the focusing of sound by the hemisphere onto the one piezo attached to it. This may be useful for future chambers should the hemisphere problem re-emerge. In the 60 kg chamber, the camera angle does not permit complete view of the hemisphere, but hemisphere events that cannot be visually tagged as such may perhaps lend themselves to an acoustic identification.

7.5.2 Water-CF_3I Interface Events

Interface events were found to be quieter than even single-bubble neutron bulk events, creating a rough peak at $\sim$0.6 "protobubbles." This is surprising for three reasons. First,

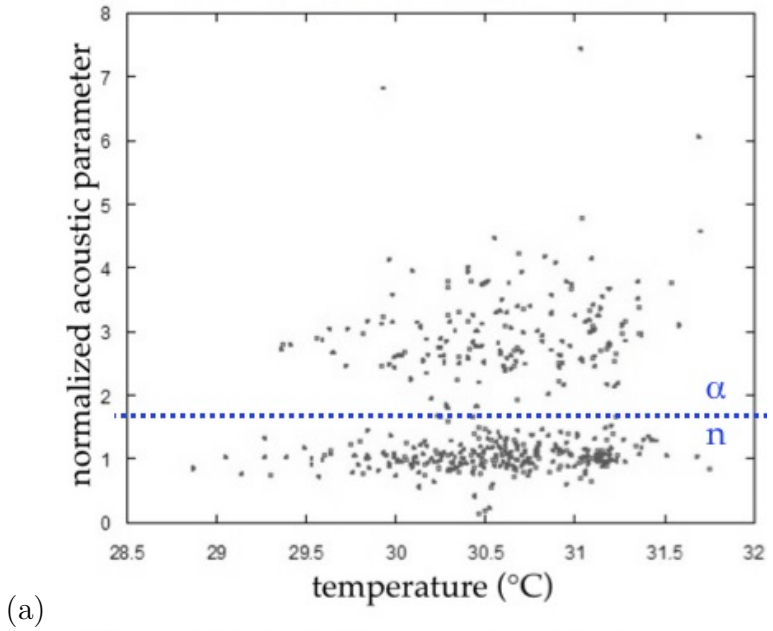

(a)

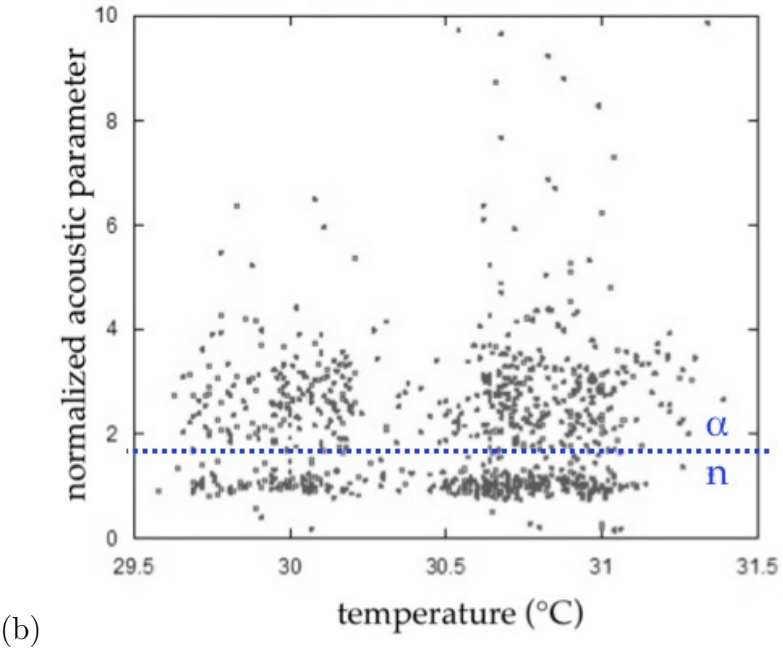

(b)

Figure 7.8: (a) The temperature dependence of the separation of alphas and neutrons for the pressure scan data. The highest density of points occurs at 30.6°C, which was the average temperature of the 15 kg chamber during the pressure scan run. The temperature was actively maintained as close as possible to that value.
(b) The temperature dependence of the separation of alphas and neutrons at 3 psig. The greater number of points appearing at AP > 4 at the higher temperatures is not significant and can easily be explained by the fact that the high-temperature population is greater.

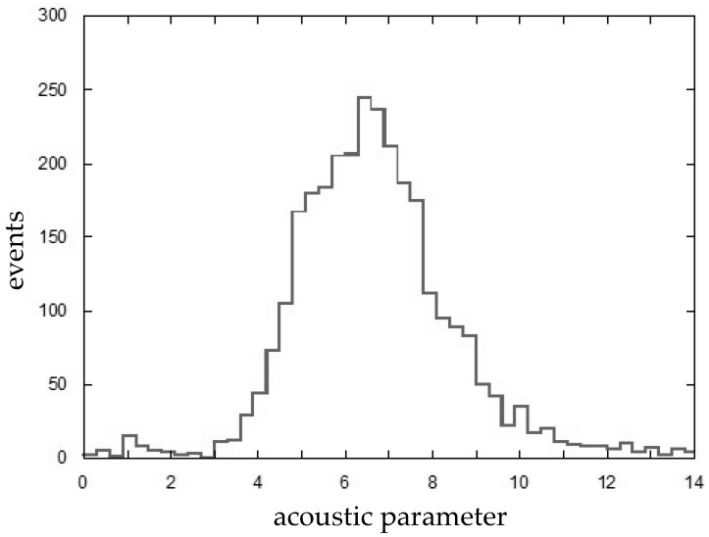

Figure 7.9: Histogram of acoustic parameter for the events (both wall and bulk) occurring in the hemisphere. The acoustic parameter shares the same scale as in all previous figures. Therefore, these events are more than twice as "loud" as typical alphas.

three piezos were immediately outside the quartz vessel at the interface. Second, interface events often appear to be explosive in nature, similar to hemisphere events. Third, one explanation for interface activity is alphas, which should be loud. On the other hand, the low energy is plausible because interface events cannot grow as normal (bulk) bubbles do; once in the water they no longer have more CF_3I vapor to integrate. Perhaps the value of 0.6 comes from roughly half of the acoustic energy going into the water if a bubble precisely on the water interface is half in the water and half in the CF_3I at birth. This may help future chambers, as position resolution may get poorer in larger vessels, making it difficult to isolate the interface location by optical position reconstruction alone. It serves as an additional check on whether hand-scanning can be trusted. If a camera is positioned at an angle with respect to the interface, bulk events interfere. In larger chambers, at least in the preliminary stages of setting up the parameters for the visual trigger, it may be possible to miss bulk events near the interface and count them as interface events once they escape into the water. This is already a concern for the 60 kg chamber and this may be the reason

for the high-AP tail observed in Figure 7.10.

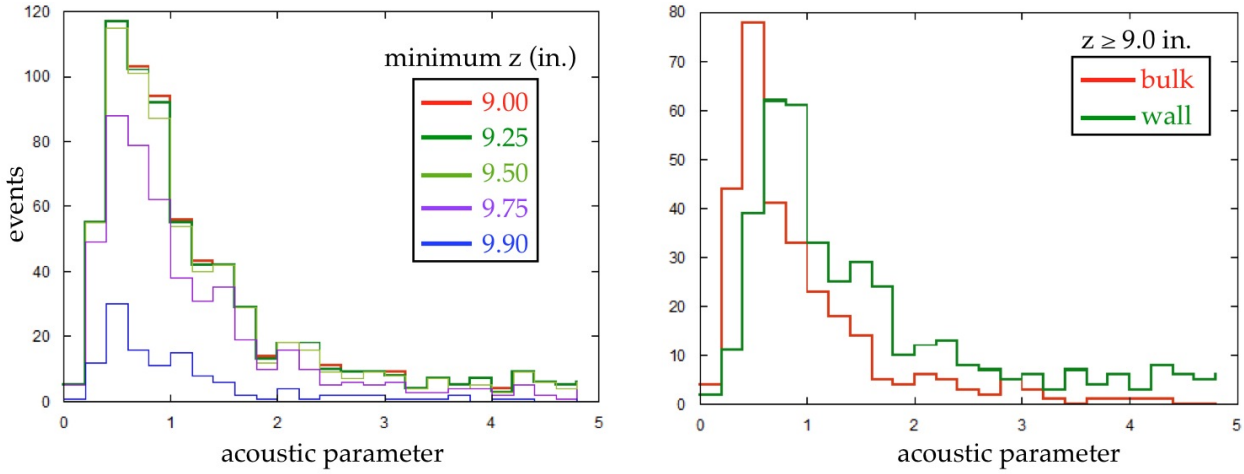

Figure 7.10: (left) Acoustic histogram of interface events as a function of different definitions for the minimum level of the interface. Observe the peak at AP < 1 for all cases.
(right) Division of the interface histogram into bulk (r < 2.3 in.) and wall (r > 2.3 in.). Wall events are louder than bulk, but not as loud as non-interface wall events (Figure 7.11) and still at AP < 1. A significant portion of the interface events was near the wall, consistent with new data from the 60 kg chamber. The phenomenon is not yet understood, but is likely due to iodine crystals or other solids in the water.

7.5.3 Wall Events

Bubbles near or on the cylindrical wall, like hemisphere bubbles, appear to be louder than alphas (Figure 7.11). This may be due not to proximity to the walls, but in fact to actual alpha events, which exist even in synthetic quartz (Section 3.3.2). This consideration goes beyond the fiducial cut made due to degradation of signal near the wall in Figure 7.3 manifested as higher acoustic parameter in a wider range. This is why any remaining bulk neutron events in Figure 7.11 form an asymmetric peak with significant contribution above AP=1. The signal degradation implies that this analysis is merely an indicator of general trend, and is not concrete. However, given the uncertainties in optical position resolution, which may worsen with future, larger bubble chambers, this is a tantalizing prospect for a future acoustic discrimination of the wall. It would permit easier fiducialization, even if the

degradation is not entirely comprehended and only partially solved (Figure 7.11).

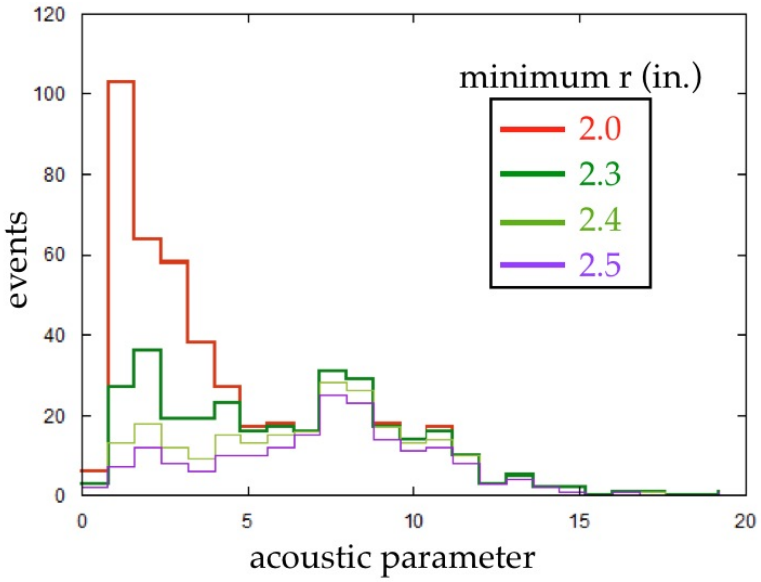

Figure 7.11: Attempt at performing acoustic discrimination of quartz wall events. Different definitions for the wall are plotted. As the definition becomes stricter, the feature at low acoustic parameter, no doubt stemming from bulk neutrons, becomes increasingly less prominent. Meanwhile, the apparent peak at AP=8 is much less affected. This may mean that wall events constitute a high-AP population similar to the hemisphere events (Figure 7.9), or mean that they are exclusively alphas, which would not be surprising given alpha contamination in quartz (Section 3.3). The sum of all the events comprising the high-AP peak appears to be too large when compared with the results shown in Sections 3.3.2 and 6.5.2, implying a contribution from bulk alphas.

CHAPTER 8
EFFICIENCY AND THRESHOLD CALIBRATIONS

A neutron source was used to calibrate the detector in order to learn the efficiency for nuclear recoils above threshold to nucleate bubbles. In doing so, it could help pinpoint the one "free" parameter in the Seitz model, the Harper factor, which simply scales the critical radius (Section 2.3). A calibration with a known source also helps verify whether the temperature and pressure are correct. The response of the detector to radiation is a function of the recoil thresholds, which depend strongly on the degree of superheat (Section 2.3). A neutron source calibration is also important because of the similarity of the falling spectrum for neutrons to a WIMP spectrum (Section 2.5.1). Alphas, on the other hand, produce a sigmoid. They constitute a flat background up to a cut-off pressure at which the chamber is no longer sensitive to them (Section 2.5.2).

8.1 Switchable Americium-Beryllium Neutron Source

A neutron source was developed for the 15 kg chamber, duplicating the design of an existing source used for the 2 and 4 kg chambers: a Switchable Americium-Beryllium Source, or, sAmBe/SAMBE. It consisted of three copper disks inside an aluminum canister. One disk held two squares of beryllium foil, and the other two held fifty small americium sources recovered from smoke detectors arranged in rough squares matching the location and area of the beryllium. A solenoid would rotate the disk with the beryllium, so that when the source was "on," the alpha sources would be in line with the beryllium. An alpha-n reaction takes place, converting beryllium into carbon and releasing a neutron:

$$ {}^{9}_{4}Be + \alpha \rightarrow {}^{13}_{6}C \rightarrow {}^{12}_{6}C + {}^{1}_{0}n \tag{8.1} $$

The ability to turn SAMBE on and off was important because it needed to be close to the chamber to produce reactions, yet it would be inconvenient to remove it from behind the muon veto and water shield only to replace it later. Leaving a neutron source permanently next to the chamber was never an option because the high neutron background would inhibit the search for the WIMP.

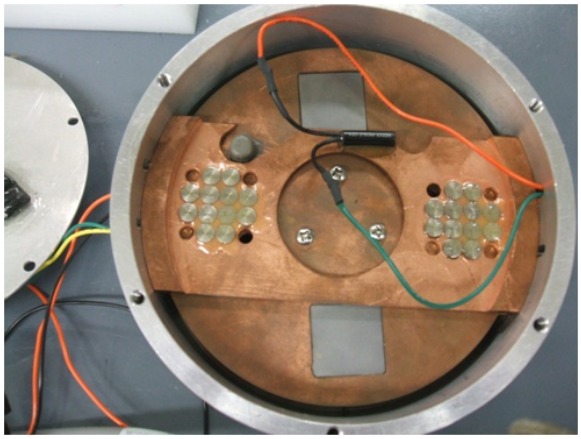

Figure 8.1: The Switchable Americium-Beryllium source (SAMBE), a compact neutron source with negligible neutron yield in the "Off" position. When the copper plate with the embedded Be foils is rotated 90 degrees the ^{241}Am alpha sources line up with the foils, generating ∼4 n/s neutrons through an (α,n) reaction. Small anisotropies in neutron emission with respect to alpha direction are minimized by isotropically bathing Be foils with alphas. The long black cylinder with attached wires is a Hall effect magnetic sensor. The gray spot in the copper notch at left is a magnet. When SAMBE is engaged, alignment occurs, resulting in a voltage across the sensor.

8.1.1 Improvement over Past Sources

The 15 kg chamber's version of SAMBE was an improved design in several ways. A fan and a solenoid saver were added to prevent the solenoid from overheating. The solenoid could draw too much current over a prolonged period of time while holding the disk in place. The solenoid saver is a device which pulses the voltage in a square-wave pattern for an adjustable duty cycle. Set at the minimum of 20%, it was sufficient to allow the solenoid to keep the disk fixed while reducing the chance of overheat. Another new feature

was the addition of a magnet and a magnetic sensor, so that it was possible to tell whether SAMBE was engaged. The magnet was too far from the sensor to register a voltage unless the disk had been fully rotated into place. It was thus possible to detect whether SAMBE was malfunctioning even if the correct current was still being drawn. This can happen if the solenoid becomes jammed in its original position.

SAMBE data was not only collected in contiguous days-long runs, but interleaved with WIMP search data. For every decompression, there was a 5% chance of activating SAMBE built into the DAQ software. This random 5% component stood in contrast to past chambers. It served to be better calibration for the entire data set by being more representative of the entire run, during which systematic changes in behavior could have occurred that would not be reflected in a one-time-only dedicated calibration.

8.1.2 MCNP Simulations of Neutron Source

A ^{3}He counter was used to characterize the SAMBE neutron yield. This detector capitalizes on the reaction

$$ {}^3_2He + {}^1_0 n \rightarrow {}^3_1H + {}^1_1 p \tag{8.2} $$

^{3}He is used because of its high absorption cross-section for thermal neutrons. The resulting protons become detected as they ionize the stopping gas of a proportional counter. The Los Alamos National Laboratory-maintained program MCNP (Monte Carlo Neutral Particle transport code) was utilized to create a simulation of the geometry of the setup in order to determine the efficiency, so that counts in the ^{3}He, a cylinder sitting adjacent to SAMBE, could be translated into an overall neutron rate in 4π solid angle. SAMBE and the ^{3}He counter were placed into a cubic structure composed of high-density polyethylene bricks in order to reduce the ambient neutron background and increase the signal-to-noise ratio for

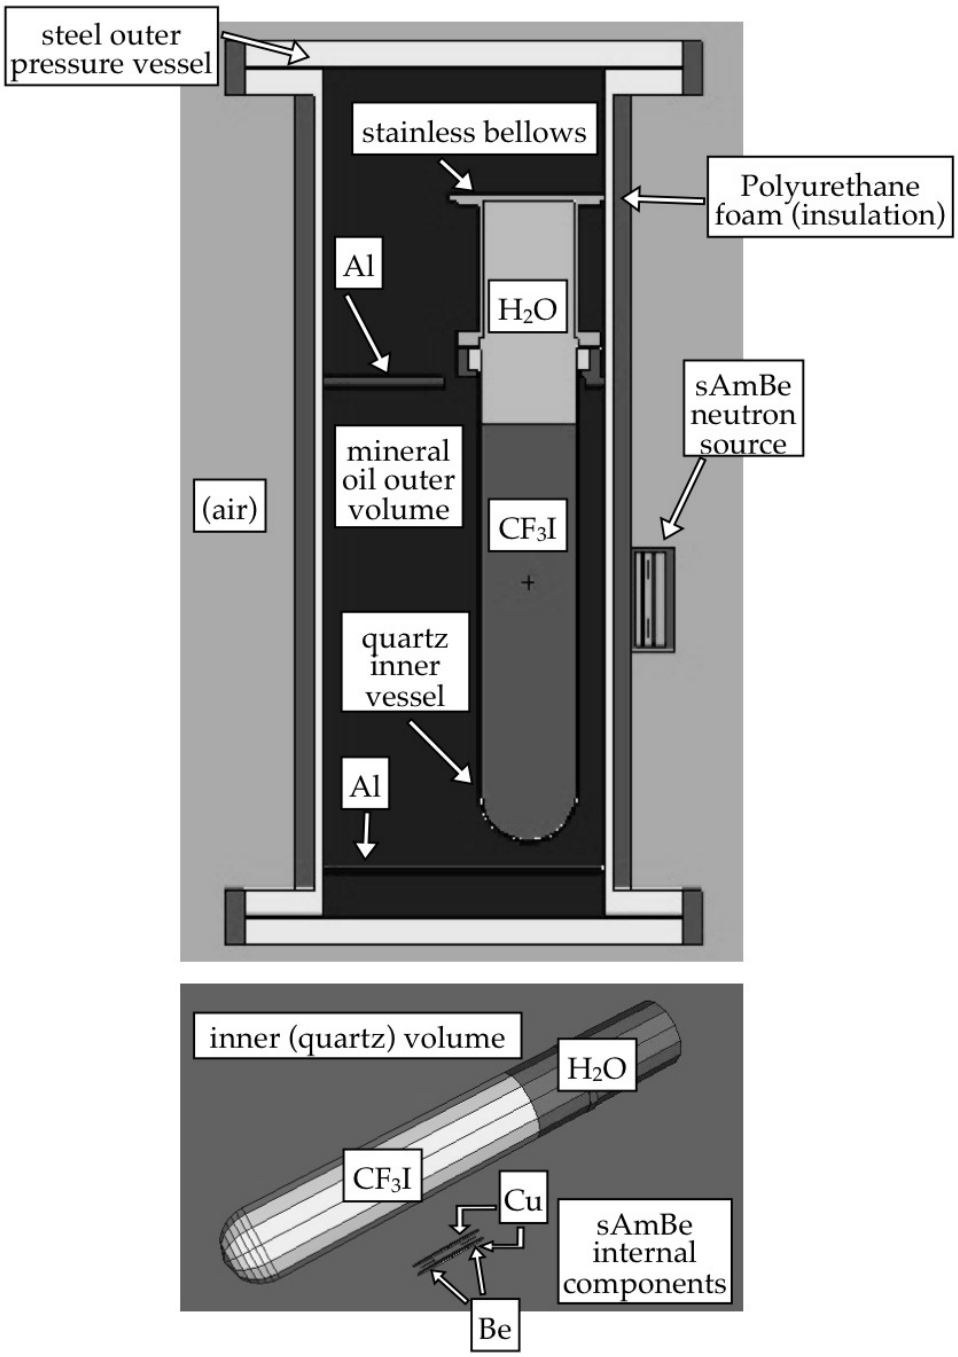

Figure 8.2: 2-D cross-sectional (top) and 3-D (bottom) representations of the 15 kg chamber geometry as simulated in MCNP (Monte Carlo Neutral Particle transport code). As depicted at top, the SAMBE was approximately centered on the CF_3I volume, positioned immediately outside the pressure vessel, pressed against the heat insulation. The asymmetric quartz vessel placement was necessary to ensure full visual coverage by the cameras (Section 6.1). At the bottom only certain volumes are depicted to maintain figure clarity (CF_3I, water, Be, Cu).

the neutron rate. SAMBE neutrons were moderated by surrounding the ^{3}He counter with polyethylene. The rate was determined to be 4.4 $\pm$ 0.2 n/s, agreeing better than 10% with a previous study of the original SAMBE unit for the 2 and the 4 kg chambers. The original study included a measurement of the hardness of the neutrons emitted by the source as well (Collar et al., COUPP collaboration internal note).

Another MCNP simulation was written for the 15 kg chamber geometry, for modeling interactions of neutrons with the chamber after being moderated by mineral oil and passing through the steel and other materials present. More specifically, the MCNP-Polimi variant of the software was used, which is an analog simulation, tracing individual neutrons and recording their interactions (Pozzi et al., 2003). The standard version of MCNP, a digital Monte Carlo, utilizes the Russian Roulette method for particle fates. (For example, if a neutron has a 50% probability of scattering off a particular nucleus, then MCNP creates a particle of weight 0.5 which scatters and another of 0.5 that does not. Those weights are used for tallying purposes.) MCNP-Polimi, by virtue of tracking individual simulated particles, has the ability to output the energies of their interactions, locations, timing, and reaction type. This is necessary for finding the number of neutrons capable of making bubbles. The simulation output, combined with a threshold calculation for CF_3I based on vapor pressure, the masses of the individual nuclei, and other physical properties (Section 2.3) was used to generate predictions of the number of events of various multiplicities produced by SAMBE per kg-day. These predictions were then compared with real data.

8.2 Results and Discussion

The best fit to the data was found by varying the Harper factor, the single free parameter of the Seitz model (Section 2.3), and the efficiency. The fit is poorer if the first two pressure bins (3 and 6 psig) are included. They are clearly anomalous (Figure 8.3). Excluding them

is conservative, however, because doing so leads to a lower best-fit efficiency, hence worse WIMP limits. Capitalizing on the interspersion of SAMBE data amongst all data (Section 8.1.1), subsets of data from different times during the run were checked, but the high-rate low-threshold anomaly was not tied to any particular time period. A possible explanation for these high outliers is a systematic uncertainty in the SAMBE livetime driven by its mechanical switching speed. Nucleation in the presence of the source was so prompt at low pressure (high superheat) that the time between SAMBE being engaged and nucleation was often less than one second. Barring this anomaly, an excellent fit to the data is achieved assuming constant (<100%) efficiency versus pressure. All previous COUPP chambers have successfully used this approach for 30°C data points. The efficiency approaches 100% with higher temperatures. At ≥40°C, all past data was successfully fit assuming 100% efficiency, as expected in the Seitz hot spike model.

As an additional complication to temperature-dependent efficiency, changes in threshold may actually be sigmoid in nature (Section 2.4), as the PICASSO collaboration assumes when analyzing their data (Archambault et al., 2009; Aubin et al., 2008). COUPP is planning more in-depth measurements in the near future with dedicated bubble chambers reserved for neutron calibrations in order to study this issue. For this thesis (and for all COUPP publications) a conservative approach is taken in adopting a constant efficiency for all pressures with step-function thresholds. This leads to less stringent dark matter limits than assuming a pressure-dependent efficiency and soft thresholds would.

The predicated and actual SAMBE data are shown in Figure 8.3. This result is consistent with previous COUPP results at 30°C. Statistics are better, but there is still error. One contribution to the vertical error bars is Gaussian, while another is the best estimate of the error in the determination of the livetime, ∼0.3 s based on observations made by the operator, plus audio traces where the sound of the mechanical movement associated with the

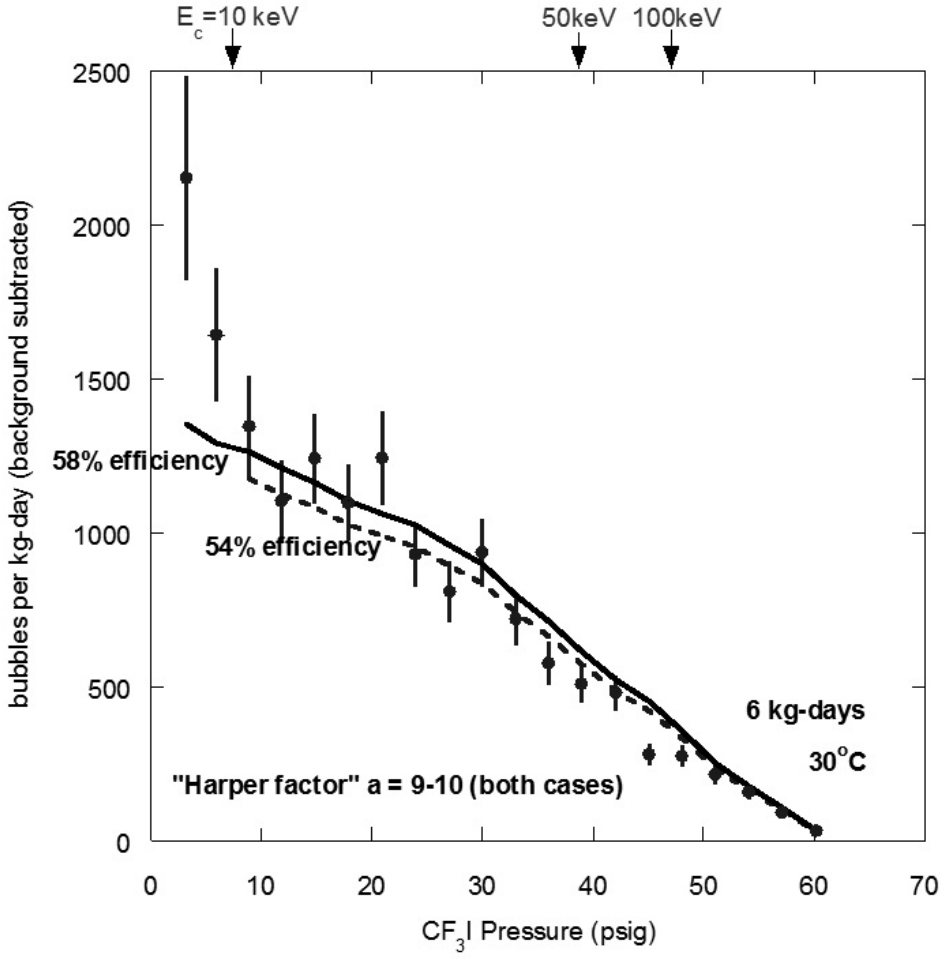

Figure 8.3: Comparison of the neutron rates obtained with the Switchable Am-Be source with the predicted rates based on an MCNP-Polimi simulation. The two free parameters for the fits were the efficiency and Harper factor (Section 2.3). The best-fit efficiency as determined with a Neumann chi-squared test, which uses the theoretical error (estimated error of the simulation in this case) is lower than 100% at 30°C, a low degree of superheat. A 92% confidence limit (C.L.) is calculated when including all data (solid, 58% efficiency), or 99% when excluding the first two bins (dashed, 54% efficiency). (The Seitz theory predicts a constant 100% efficiency. Behnke et al. (2011) demonstrates agreement with this at 40°C.) It is conservative to choose a constant, lower efficiency, which yields worse WIMP limits (Chapter 9).

switching on of the SAMBE was evident. Horizontal error bars (not visible) are based on the standard deviation of all pressures contributing to each data point on the plot. For the fits, the simulation was assumed to be the dominant source of error, not the experimental errors assigned to the data, primarily due to systematic uncertainty in relative position causing a neutron rate uncertainty. Blindly determining the exact SAMBE placement relative to the quartz jar from outside the pressure vessel for the purpose of accurate reproduction in simulation geometry proved difficult. The simulation was estimated to have a ∼30% error, following a similar assessment of uncertainties for the 2 kg SAMBE run, and this was used in Neumann chi-squared tests of the fits, yielding confidence levels (C.L.'s) of >90% with or without the anomalous first two points. If a stricter 20% error is assumed, then these points cause the C.L. to drop to 22%, while removing them keeps it above 90%. The background rate of events not induced by the neutron source, as estimated from all of the data taken when the source was off, was individually subtracted at each point.

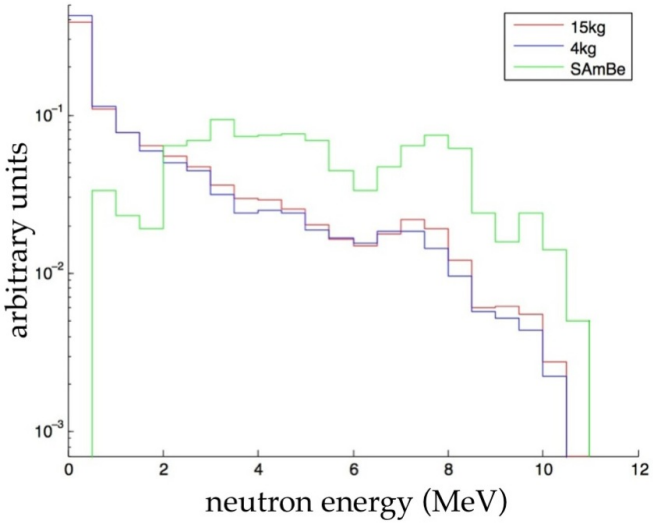

Figure 8.4: The energy spectrum of neutrons as they exit the source canister compared to their energies upon arrival and initial interaction at the CF_3I volume of the 15 kg bubble chamber, and the 4 kg chamber for contrast. SAMBE initial spectrum (green) has been adapted from Collar et al. (COUPP collaboration internal note). Figure courtesy of Drew Fustin.

In Figure 8.3, predicted nucleation rates for events of any bubble multiplicity are compared with actual rates. The comparison is not made using solely single bubble events, in order to sidestep the issue of bubble mergers yielding inaccurate bubble numbers. The best fit Harper factor is 9 or 10 (the fits are degenerate for these two values), which plays a significant role in determining the cut-off pressure where the nucleation rate becomes negligible, while the best-fit efficiency is close to 50%. The total SAMBE exposure was 6 kg-days, slightly more than in previous COUPP bubble chambers.

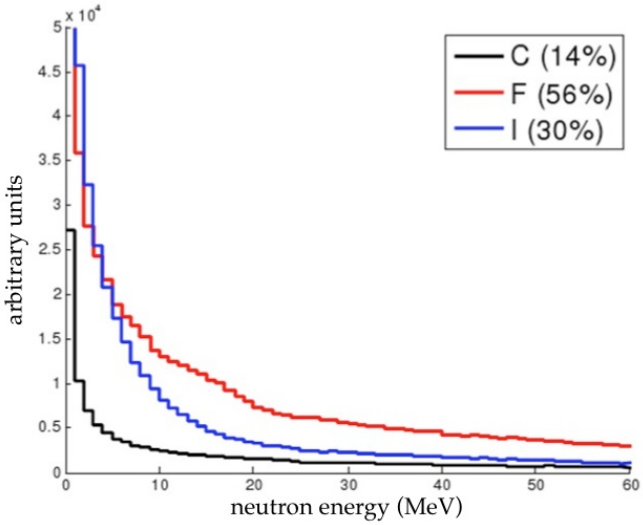

Figure 8.5: Histograms of the the recoil energies of carbon, fluorine, and iodine nuclei in the 15 kg vessel interacting with SAMBE neutrons. Percentages indicate how many neutrons interact with the different nuclei for a given fixed reaction rate. The WIMP-like decreasing spectrum of these neutron-induced recoils, heavily biased to low energy, implies that neutrons serve as a useful calibration particle for a WIMP detector (Section 2.5.1). Figure courtesy of Drew Fustin.

8.3 Vapor Pressure Measurement

After the failure of all thermometers except for the bottom one, it became necessary to measure the vapor pressure in order to properly determine the mean temperature of the detector. The measurements were performed by permitting the chamber to boil, purposely

not compressing during a bubble. The SAMBE source was activated to more effectively boil the CF_3I.

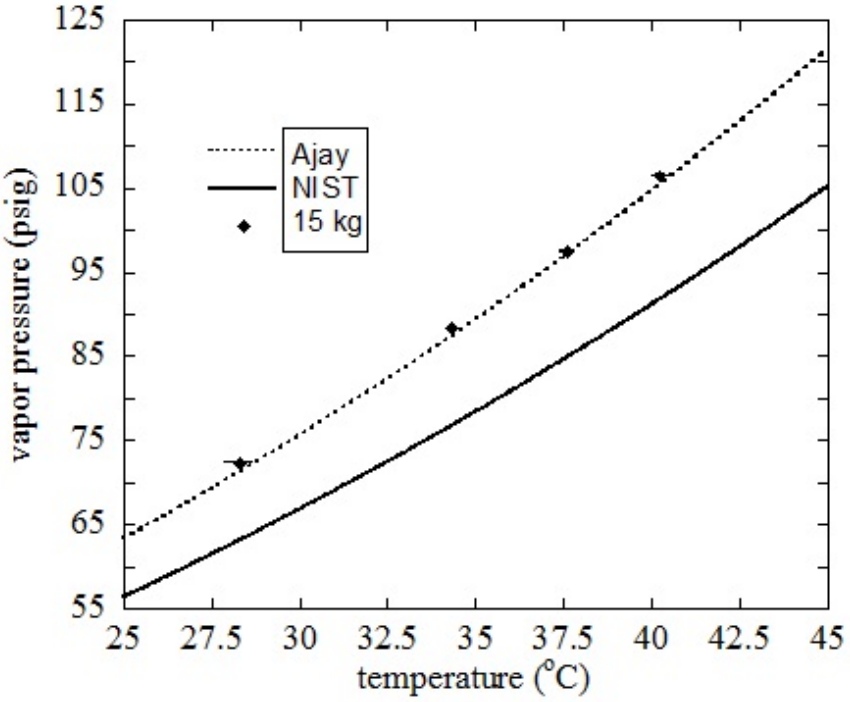

Figure 8.6: Comparison of 15 kg vapor pressure data to the predicated Ajay and NIST curves. Uncertainty in pressure does not appear due to its size (<0.5 psi). The temperature was measured using the single surviving thermometer. Error in temperature was driven primarily by diurnal thermal drift during day-long spans when the chamber pressure was given time to asymptote to its final value. Other chambers have failed to reproduce these results: this may have been a consequence of their use of a two-stage distillation process to remove impurities.

The CF_3I in the 15 kg chamber was found not to follow the NIST vapor pressure curve, one created by compiling all relevant experiments and weighting them based on experimental uncertainty, but instead followed the vapor pressure curve provided by the manufacturer, Ajay (Figure 8.6). Further evidence is the pressure onset of alpha sensitivity, especially when compared to the 2 kg chamber results (Chapter 9). This may have been a result of volatile impurities in the Ajay product acting to raise the vapor pressure. Future chambers will pursue aggressive purification techniques. Although for the 15 kg chamber the mea-

sured vapor pressure meant a lower minimum energy threshold (8.7 instead of 12.1 keV as expected based on the NIST curve), thus marginally improving the WIMP limits (Chapter 9), impurities are in general nonetheless undesirable. They may lead to inhomogeneous nucleation by creating pockets of non-condensable gas in surface cavities, react chemically with internal chamber parts, or produce other unpredictable effects.

In the SAMBE results quoted in Section 8.2, the Ajay vapor pressure curve is used to calculate the nuclear recoil thresholds. The NIST vapor pressure curve, indicating a lesser degree of superheat for the same pressure compared to the Ajay curve, cannot reproduce the $\sim$60 psig cut-off for neutron reactions observed in the 15 kg chamber data. Less superheat forces the cut-off to occur earlier in pressure in a manner entirely inconsistent with the data in the high-pressure region, where the uncertainties in rate are by far the smallest. Using the NIST curve leads to fits with C.L. $\ll 1\%$.

8.4 Gamma Calibration

Due to the powerful gamma rejection measured in previous COUPP chambers (Figure 2.4), a gamma calibration was not performed. At the lowest threshold achieved by the 15 kg chamber (8.7 keV), the gamma rejection factor was still $>10^9$ (Figure 2.4, inset). For this reason, environmental background radiation is not sufficient to produce a measurable rate resulting from gammas. A hot source must be placed directly next to a COUPP vessel in order for data to exhibit gamma-induced nucleations. However, future work on the 4 kg chamber at SNOLAB will involve a high-intensity gamma-ray re-calibration, as a function of both temperature and pressure, as a further demonstration and confirmation.

CHAPTER 9
DARK MATTER WIMP LIMITS

The calculation of the 90% confidence limit (C.L.) excluding a region of WIMP masses and cross-sections is presented in this final chapter. The method used to extract a limit from the raw data of event rate versus pressure, the indicator of nuclear recoil energy threshold, is detailed. The flat background at high pressure (high energy threshold) likely dominated by alphas is subtracted from lower pressure points. All equations in this chapter closely follow the derivation in Lewin & Smith (1996) unless otherwise noted.

9.1 Final Pressure Scan Results

The pressure scan data set covered pressures from 3–60 psig in 3 psig bins, spanning recoil energy thresholds of 8.7 keV to ∼1 MeV. Over 400 kg-days of physics-quality, dark matter search data was ultimately collected, after several months were spent resolving all DAQ issues (Section 7.3). The data set was reduced to 287 kg-days after fiducial volume cuts necessary to remove the water interface, the hemisphere, and the wall, i.e., the quartz vessel surface (Section 6.5). The resulting spectrum is compared and contrasted in Figure 9.1 with 2 kg COUPP data taken at the same temperature and over a comparable range of pressures.

9.1.1 Accomplishments and Lessons Learned

Compared to its immediate predecessor, the 15 kg chamber collected significantly more data (compared with 15.3 kg-day exposure for the 2 kg chamber at 30°C) with superior statistics (10–20% vs. ∼50% uncertainty). This enables the 15 kg chamber data to serve as a blueprint for all future chambers. The 15 kg chamber achieved this through sheer mass and runtime. A livetime for the ten-times greater fiducial mass (10.75 vs. 1 kg) of >70% was made possible with wall rate reduction accomplished using synthetic quartz (Sections

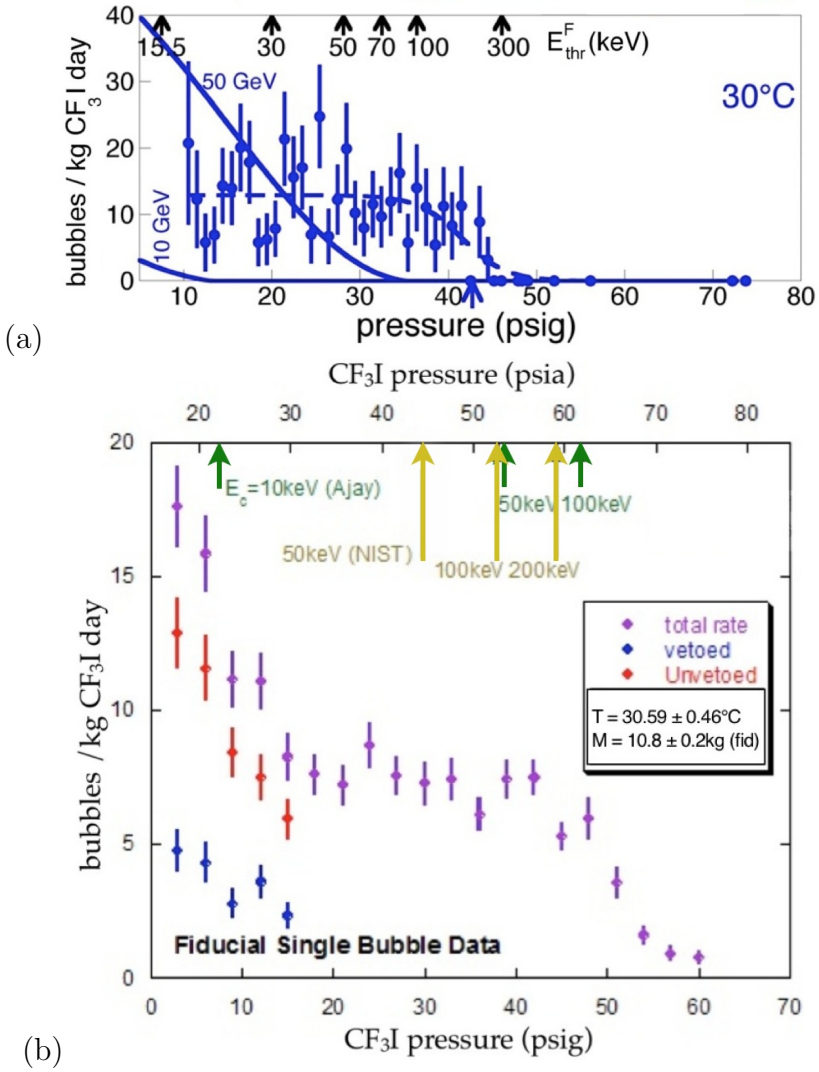

Figure 9.1: (a) The 30°C raw data from the 2 kg chamber (15.3 kg-days of exposure). The alternative x-axis at top is the energy threshold for fluorine recoils, in keV. (Below 1 MeV, tied with carbon, it has the lowest threshold of all three constituents of CF_3I.) The arrow at one point along the bottom x-axis is the calculated onset of sensitivity to alpha recoils: it is the "cut-off" for the recoils that accompany alpha emissions. Example spectra from WIMPs of masses 10 and 50 GeV and 3 pb cross-section indicated in the figure are ruled out by those data. Plot adapted from Behnke et al. (2008).
(b) The full pressure scan results from the 15 kg run (287 effective kg-days). Note the higher statistics (smaller error bars) and smaller background than in (a). Vetoed and unvetoed events are individually plotted where possible (where the bubble acoustic signal is sufficiently large to use). The energies labelled are the critical energies (energy thresholds) calculated using NIST and Ajay vapor pressure curves. Note that agreement with the known alpha cut-off of ∼100 keV is achieved only using the Ajay curve.

3.3.2 and 6.5.2), despite the increase in surface area necessary to accommodate this larger mass. The published 2 kg chamber result was not achieved with a muon veto or piezos (Behnke et al., 2008), both of which were used to great effect in the 15 kg chamber for an enhanced understanding of the results. When taking into account the apparent different vapor pressure of the 15 kg chamber (Ajay vs. NIST curves), the alpha cut-off observed in both chambers agree, occurring at $\sim$100 keV as anticipated (Section 7.1).

The inability to extract useful acoustic data above 15 psig motivated the placement of a greater number of piezos upon the 60 kg vessel, with better positioning, and the vapor pressure disagreement with official NIST data ensured that greater attention would be paid to fluid handling in future COUPP deployments.

9.1.2 Spectrum of Acoustically Tagged Events

The unvetoed events of Figure 9.1 are sorted acoustically into neutron or WIMP-like events, and likely alpha candidates (Figure 9.3). The dividing line was chosen to be AP=1.75, containing 95% of the neutron population based on the SAMBE and veto-coincident data, and $<$5% of alphas, given the alpha peak centroid of AP$\simeq$3, assuming a roughly symmetric peak (Figure 7.5). This is a conservative estimate, given the lack of a low-AP tail in the events with multiple observable bubbles (Figure 7.7), which motivates the conclusion that multiple bubbles must be louder than singles. Alpha events should generate at least two protobubbles, one corresponding to the recoiling nucleus and one to the alpha itself (Section 7.1). Few or no unvetoed neutrons were expected due to the presence of the water shield (Section 4.2), but an unvetoed neutron-like background is observed to remain. It is indistinguishable event-by-event from WIMPs. Given the $\sim$90% efficiency of the muon veto, with its rate meeting expectations for TARP's depth of 300 m.w.e.(Section 5.7), veto inefficiency cannot adequately explain this background. The most likely explanation is degradation in the cubitainers (water containers used as neutron moderator around the

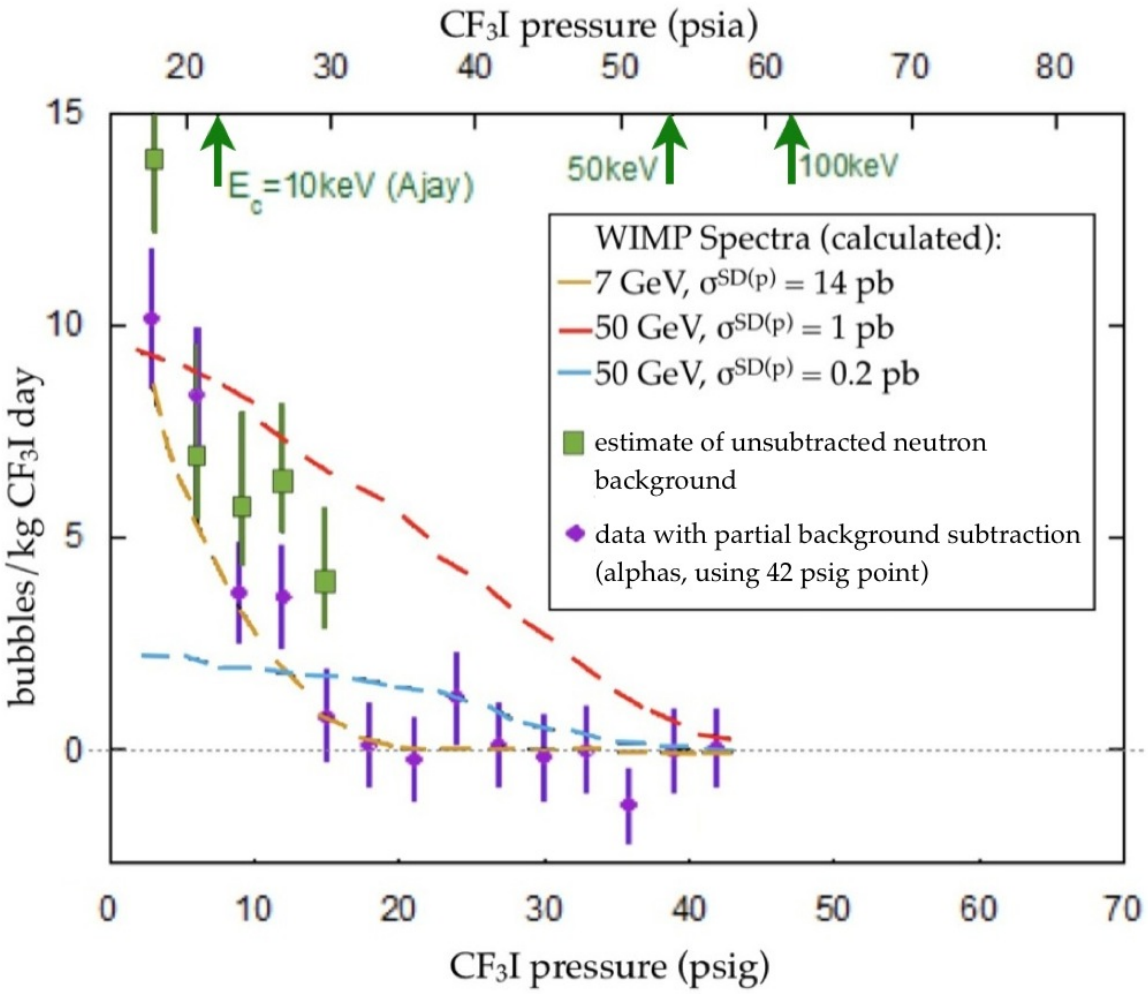

Figure 9.2: Final 15 kg chamber data, with background subtraction. This background is assumed to be dominated by alpha emitters, generating a constant background as a function of pressure. Three sample spectra for WIMPs of different masses and cross-sections are plotted: one in blue close to not being ruled out, one in red strongly unsupported by the data, and one in orange whose spectral shape is comparable to the data and cannot be ruled out. The plausibility of a neutron contribution at low threshold is illustrated by an extrapolation based on the muon-induced neutron rate as determined with the muon veto. Neutrons have conservatively not been included in the model generating the WIMP limit curves in Figures 9.4, 9.5, and 9.6 (refer to Section 9.1.2 for details).

experiment). The originally cubic containers bowed over time, creating large gaps that could serve as direct lines of sight for environmental or cosmogenic neutrons. It was not possible to write a reliable simulation to estimate the neutron leakage through these gaps, due to

the unknown amount of solid angle subtended, which was considerable, easily perceptible by visual inspection.

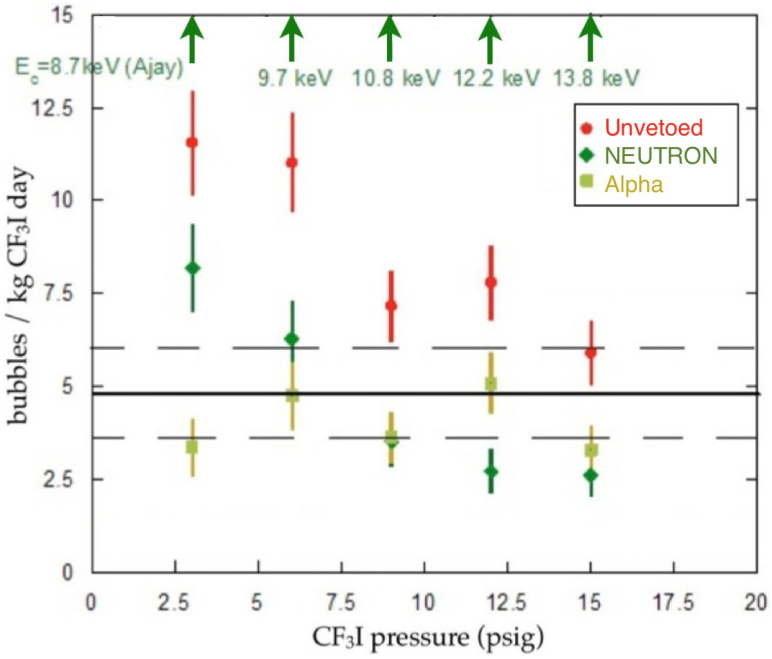

Figure 9.3: Total unvetoed rate (red) with two constituents identified as alphas (gold) or neutrons (green) acoustically. Treating the $\lesssim$50% inefficiency in vetoing multiple bubble events (necessarily neutrons because of their short mean free path) as implying the same value for the inefficiency in vetoing single neutrons, one can calculate an estimated "corrected" total unvetoed rate (thick horizontal line in figure, with dashed lines delineating the 1-sigma error region). Observe the resulting consistency between the total unvetoed rate and the alpha rate. This makes it more plausible that the WIMP-like decreasing signal (green) is a neutron background, but given the uncertainty in extrapolating knowledge from multiples to singles, and the inconsistency with the 90% measured muon veto efficiency discussed in Section 5.8, this observed consistency is conservatively not used to calculate a stronger WIMP limit.

The multiple bubble events provided a handle on this neutron background. Multiples can only be generated by neutrons because of their short mean free path (Section 2.5.1), and the percentage of vetoed multiple bubble events did not correspond with the measured veto efficiency of 90%, instead varying between ~30–50%. Assuming that this was due to neutron leakage through the shield, the measured multiple bubble inefficiency can be used as

a new overall neutron inefficiency. A correction to the unvetoed neutron rate relying on the percentage of unvetoed multiples results in a total unvetoed rate with excellent overlap with the alpha rate. This motivates the treatment of the falling spectrum of neutron-like events in Figure 9.3 as neutron background instead of a WIMP signal. However, the "corrected" total unvetoed rate is not subtracted from the acoustically tagged alphas and used as the maximum possible WIMP signal, even though this would result in significantly improved WIMP limits. This was conservatively not done because of the uncertainity associated with extrapolating the percentage of unvetoed multiple bubble neutron events to single neutron events.

Total neutron rates are estimated by dividing the vetoed single-bubble rates by the measured inefficiency of the veto for multiple-bubble events (Figure 9.2). Bubble acoustic emission is too weak above 15 psig to use for determining coincidence between muon veto pulses and acoustic traces, so the estimate does not extend beyond that pressure in the figure. The similarity in shape and size of the estimated neutron background to the data in this figure further justifies not treating the low threshold points as evidence of WIMP detection. Due to the uncertainty in extrapolating information about singles bubbles from multiples with poor statistics, this estimated neutron background is not subtracted from the 15 kg chamber data to improve the resulting WIMP limit. The limit improves with subtraction of additional backgrounds because the single bubble rates measured serve as upper bounds: WIMPs which would yield rates lower than those observed cannot be excluded by the data. The data are treated point by point as independent experiments, however (Section 9.2). For instance, the spectrum of a 50 GeV WIMP with 0.2 pb cross-section in blue in Figure 9.2 is ruled out as a consequence of the low rate measured at 21 psig, despite lying beneath the data at lower pressures. This is nevertheless conservative: this example WIMP is nearly not ruled out, because its calculated rate for 21 psig lies near the data point for that pressure. If

instead best fits are applied to the data varying WIMP mass and cross-section, the spectrum of such a WIMP is clearly a poor fit to it, thus leading to a better limit.

The flat portion of the spectrum at high pressure (low superheat) in Figure 9.1b is best explained by alpha recoils, which are approximately monochromatic (Section 7.1). However, it disagrees with the flat alpha rate measured using acoustics at low pressure (high superheat) by a factor of two. What is most likely the case is that alphas are not the only contribution to the flat spectral shape in Figure 9.1b. A neutron component slowly decreasing in rate with rising threshold is likely present. Alpha tagging inefficiency is unlikely given the clear, sharp decrease observed in the number of probable-alpha events in the acoustic parameter distributions below AP=2 (e.g. in Figure 7.5).

If all three unvetoed neutron-like events observed in the 4 kg chamber are conservatively interpreted as alphas, then the alpha tagging efficiency was only 70%. However, high-energy, so-called "punch-through" neutrons which can exist at shallow depths (Behnke et al., 2011) can adequately explain those events as neutrons.

9.2 Method of Dark Matter Limit Determination

A conservative two-point background subtraction is used to calculate the limits on the coupling of nucleons to WIMPs. It capitalizes on the nature of the bubble chamber, which is a threshold detector (Section 1.5.1). This implies that the spectra in Figures 9.1 and 9.3 are integral, not differential. All interactions which can successfully nucleate a bubble at a high threshold (high pressure) must necessarily also contribute to the total event rate at lower threshold (lower pressure). Therefore, it is possible to subtract a high pressure point from a low pressure one, calling the high pressure point the background and the low pressure point background plus potential WIMP signal. The flatness of the high-pressure spectrum implies that it is alpha-dominated (Section 2.5.2) and this flatness lends itself effectively

to background subtraction. However, given the inconsistency with the acoustic alpha rate implying neutron events still accounting for at least some of the high-pressure data, a concrete understanding of the flat region is compromised. It is therefore not averaged. This would improve the WIMP limit, but would not be conservative. Instead, only one flat-rate bin is selected (42 psig) to represent a background so that its corresponding rate can be subtracted from earlier bins. The bin is chosen to be as far away as possible in threshold, but before loss of alpha sensitivity, so that there is a sizable background which can be subtracted.

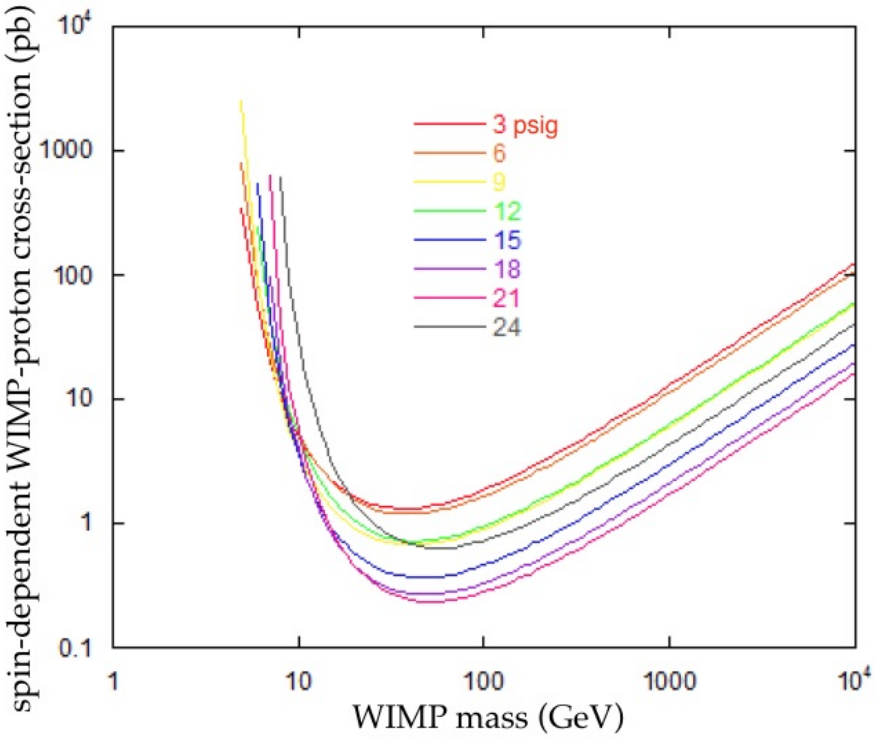

Figure 9.4: The 90% C.L. dark matter limits individually determined through background subtraction of the 42 psig (64.3 keV) rate, treating the data at various pressures as independent experiments and taking the best result (Figures 9.5 and 9.6). The Ajay vapor pressure curve is used to calculate the thresholds because of the strong evidence in support of it (Section 8.2 and Figure 9.1). Limit curves are generated here taking into account the measured 54% nuclear recoil efficiency found in Chapter 8.

9.2.1 Reasons for Rejecting Other Limit Calculation Methods

Choosing two bins close together in threshold is not conservative because the thresholds may be so similar that the subtraction may be invalid, especially at low pressure where the threshold changes slowly, or if the threshold is not perfectly step-like. Furthermore, uncertainty in the threshold arises from the temperature uncertainty. It is therefore safest to choose points well-separated in energy threshold for a background subtraction, and for this reason the 42 psig (64 keV) rate was not subtracted from any pressure bin higher than 24 psig (21 keV).

The slowly-changing threshold at low pressure, the only region where the acoustic discrimination worked because of the loss of signal at high pressure, that is, low superheat (Section 5.6.3), was unsuitable for supplying high and low points for the subtraction for the above reason. Therefore it was not feasible to capitalize on the acoustic discrimination to improve the WIMP limits. Points are not only too close together, but the falling background spectrum with increasing threshold at low pressure is non-conducive to background subtraction. Lastly, it was not conservative to perform fits to the data with a WIMP spectra of varying mass and cross-section included with the background: this approach would greatly improve the limits compared to treating each pressure point individually, but a neutron background model was lacking given the unknown size of the gaps in the neutron shield.

9.2.2 Mathematical Framework for WIMP Interactions

The general cross-section for a WIMP-nucleus interaction may be expressed as

$$\sigma = \sigma_0 \left(\frac{\mu_T}{\mu_N} \right)^2 F^2 I \tag{9.1}$$

where σ_0 is the zero-momentum cross-section per nucleon, which is what is actually graphed in WIMP limit plots to normalize all experiments with target nuclei of varying masses. The

variables μ_T and μ_N are the reduced masses for the entire target nucleus and one nucleon within it, respectively. F is the form factor. It causes an effective reduction in cross-section with increasing momentum transfer, due to finite nuclear radius ($|F| < 1$). I is the interaction factor, differing for spin-dependent (SD) and -independent (SI) interactions.

The form factor also depends on interaction type. For SI, the solid-sphere approximation is used because SI is a coherent interaction with the nucleus as a whole (Section 2.6.1):

$$F(qr_n) = 3\ \frac{\sin(qr_n) - qr_n\cos(qr_n)}{(qr_n)^3} \times e^{-(qs)^2/2} \tag{9.2}$$

where q is the momentum of the recoiling nucleus ($q = \sqrt{2M_T E_R}$, where M_T is its mass and E_R is its recoil energy), r_n is the radius of the nucleus, and s is its so-called "skin depth," which is related to the charge density (Lewin & Smith, 1996). The SD form factor usually used is piece-wise defined, with partially-filled in Bessel function zeroes as a result of considering all relevant, interacting nucleons (consult Section 2.6.1)

$$\begin{aligned} F(qr_n) &= j_0(qr_n) = \sin(qr_n)/qr_n && (qr_n < 2.55 \text{ or } qr_n > 4.5) \\ F(qr_n) &= \sqrt{0.047} && (2.55 \leq qr_n \leq 4.5) \end{aligned} \tag{9.3}$$

The functional form is based on the thin-shell approximation in this case, because the SD interaction involves individual nucleons with unpaired (uncancelled) spins in nuclei with an odd number of protons or neutrons. The skin thickness s is taken to be $\sim$1 fm, and the nuclear size r_n is approximated as

$$r_n = \sqrt{c^2 + \frac{7}{3}\pi^2 a^2 - 5s^2} \tag{9.4}$$

where $a = 0.52$ and $c = 1.23A^{1/3} - 0.60$, where A is the atomic number of the target nucleus.

$I = A^2$ for SI interactions. This is known as the coherence bonus, but is less of an advantage for nuclei recoiling from large-mass or fast WIMPs because of form factor. The value of the interaction factor I for SD couplings is instead equal to $\frac{4}{3} < S_{p,n} >^2 \frac{J+1}{J}$, where $< S_{p,n} >$ is the expectation value of the spin of the proton or neutron when bound within the nucleus in question, and J is the total nuclear spin. The values of S that are used in this thesis are taken from Tovey et al. (2000), from which the $\frac{4}{3}$ factor also comes (it is $\frac{8/\pi}{6/\pi}$).

Given existing derivations of WIMP kinematics in Lewin & Smith (1996); Tovey et al. (2000), and elsewhere within numerous other papers and theses, the following is a concise summary, but one designed to provide sufficient information so that the reader can reproduce the WIMP limit presented at the end of this chapter. The basic principle is that the event rate per particle (nucleus) is

$$\frac{dR}{dn} = \frac{N_A}{A}\sigma v \tag{9.5}$$

where N_A is Avogadro's number. The event rate per particle is proportional to the product of cross-section and velocity. Assuming that WIMPs have a (Maxwellian) most probable velocity v_0, mass M_D, and energy density ρ_D, the definition for the rate of recoil in the case of the velocity of the Earth in the Galactic frame being zero and the escape velocity of the WIMPs being infinite (with corrections to follow) is:

$$R_0 = \frac{2}{\pi^{1/2}}\frac{N_0}{A}\frac{\rho_D}{M_D}\sigma_0 v_0 \tag{9.6}$$

Normalizing to the case where $\rho_D = 0.4$ GeV/cm^3 and $v_0 = 230$ km/s converts equation (9.6) into

$$R_0 = \frac{503}{M_D M_T}\left(\frac{\sigma}{1\text{ pb}}\right)\left(\frac{\rho_D}{0.4\text{ GeV/cm}^3}\right)\left(\frac{v_0}{230\text{ km/s}}\right) \tag{9.7}$$

in units of events per kg-day if the mass-energies are in GeV. For this thesis, as for all other COUPP results, $\rho_D = 0.3$ GeV/cm^3 and $v_0 = 230$ km/s, which within a few percent have become standard assumptions in the field (Behnke et al., 2011). The differential rate is then

$$\frac{dR(v_E, \infty)}{dE_R} = \frac{R_0}{E_0 r} \times \frac{\pi^{1/2}}{4} \times \frac{v_0}{v_E} \times \left[\mathrm{erf}\left(\frac{v_{min} + v_E}{v_0} \right) - \mathrm{erf}\left(\frac{v_{min} - v_E}{v_0} \right) \right] \tag{9.8}$$

still assuming an infinite escape velocity. $E_0 = \frac{1}{2} M_D v_0^2$ is the most probable WIMP kinetic energy, $v_{min} = \sqrt{\frac{E_R}{E_0 r}} v_0$ ($r = \frac{4 M_D M_T}{(M_D + M_T)^2}$) is the WIMP velocity corresponding to the minimum WIMP energy capable of generating a nuclear recoil of energy E_R, and $v_E = 244 + 15 \sin(2\pi y)$ km/s, where y is the time (in years) elapsed since March 2$^{\mathrm{nd}}$. The constant v_E is the Earth velocity through the Galaxy as it moves with the solar system. It is approximately equal to the WIMP velocity, varying as the Earth revolves around the Sun. For the 15 kg chamber analysis, v_E is approximated as a constant (244 km/s) because data was taken over the course of only one month. As in Behnke et al. (2011), the experiment was not long enough to warrant taking this modulation into account. The differential rate for a finite escape velocity is

$$\begin{aligned} &\frac{dR}{dE_R} = \frac{dR(v_E, v_{esc})}{dE_R} = \frac{1}{k} \left[\frac{dR(v_E, \infty)}{dE_R} - \frac{R_0}{E_0 r} e^{-v_{esc}^2 / v_0^2} \right], \text{ where} \\ &k = \mathrm{erf}\left(\frac{v_{esc}}{v_0} \right) - \frac{2}{\pi^{1/2}} \frac{v_{esc}}{v_0} e^{-v_{esc}^2 / v_0^2} \end{aligned} \tag{9.9}$$

where v_{esc} is the velocity at which WIMPs can break free of the Galactic gravitation field. Following Behnke et al. (2011), 650 km/s is used as the escape velocity. It is typical to assume $v_{esc} > 600$ km/s on theoretical grounds (Lewin & Smith (1996) and ref. therein). All incoming WIMP angles are treated as equally likely (consult Lewin & Smith (1996)

for how this is accomplished). This is a valid approximation because the experiment was not a multi-year one nor directionally sensitive, unlike specifically directional efforts, such as those presented in Ahlen et al. (2010), sensitive to the "WIMP wind" coming from the constellation Cygnus as a result of the motion of the Sun through the Galaxy, taking the Earth along. This is only a few-percent effect for the differential rate regardless (Lewin & Smith, 1996). (The gently-superheated bubble chamber cannot track particles the way HEP chambers did, as explained in Section 1.5).

To find the integrated rate, the differential rate $\frac{dR}{dE_R}$ is integrated from the recoil energy threshold (equation (2.3)) to infinity. Carbon, fluorine, and iodine are treated separately. The total rate is then calculated by multiplying the individual rates by the correct mass fractions and summing. The WIMP rate at a given cross-section is reduced with a multiplication by the recoil efficiency (Section 8.2). The background-subtracted rate (see Section 9.2) plus 1.645σ (corresponding to a 90% confidence level) is treated as the total potential WIMP signal at each pressure. The cross-section corresponding to this rate is found by looping over cross-sections in a C program and calculating their associated rates, starting high. This loop was contained within another loop, for WIMP mass. In this fashion, for each given WIMP mass, the cross-section was found which would make that WIMP interact at the remaining rate after subtraction plus the correct multiple of a standard deviation, so that WIMPs at that mass but with a higher cross-section are ruled out with $\geq$90% confidence. The WIMP limit curve is thus constructed. Because rejection is being performed instead of discovery being claimed, 5σ (99.9999%) is not the standard, but 90% (Behnke et al. (2011) and ref. therein).

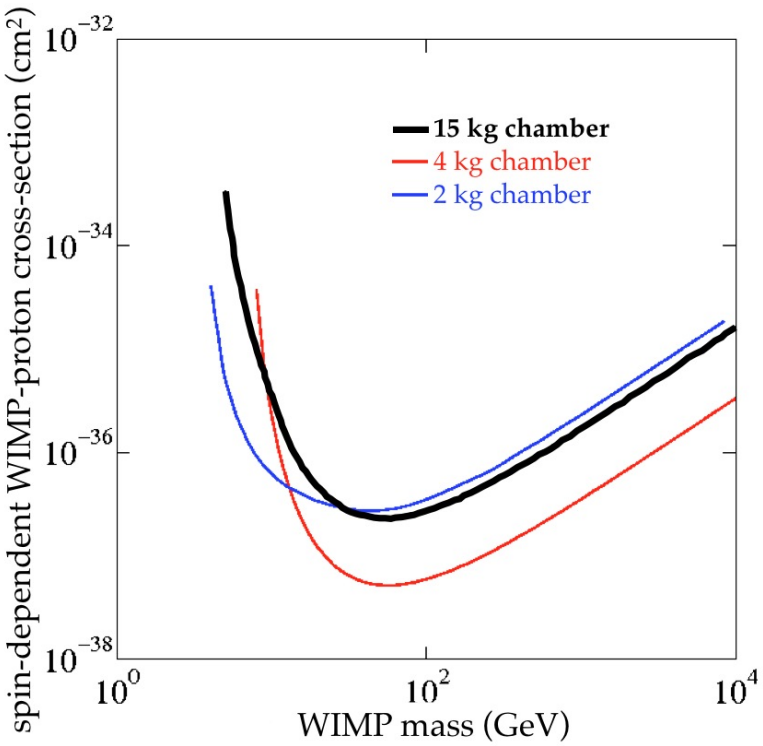

Figure 9.5: The limits of SD (spin-dependent) WIMP coupling to protons are depicted for all COUPP chambers to date. The result from the 15 kg chamber (the subject of this thesis) improves upon the limit achieved by its immediate predecessor, the 2 kg chamber. Figure generated using online limit-plotting software at http://dmtools.brown.edu (Gaitskell et al., 2010). The 2 kg and 4 kg chamber limits are taken from Behnke et al. (2008, 2011), respectively.

9.3 WIMP Limit and Discussion

The individual 3, 6, 9, ..., 24 psig points are treated as signal and the 42 psig point subtracted as background. The resulting limits are plotted in Figure 9.4. The lowest point in cross-section is selected at each mass, in order to create the best possible limit. This is still conservative because each pressure point can be treated as an individual dark matter experiment. The vast majority of experiments have a fixed energy threshold (Gaitskell (2004) and ref. therein). Taking into account multiple bins at once, on the other hand, would involve a spectral fit. This would greatly improve the limit, but be an unsound approach

because of the background shaped like a WIMP spectrum (Section 9.2.1). (SI limits are not depicted. They are poor compared to leading limits from CDMS and XENON given that problem.) This rising low-pressure (high superheat) background made it unhelpful to use the dedicated 3 psig data set (see Section 7.3).

The limit obtained for SD coupling to protons is compared with past COUPP results in Figure 9.5. (The limit on the equivalent SD coupling to neutrons is not depicted because CF_3I does not contain an odd-n isotope with significant coupling strength to generate a competitive limit.) In Figure 9.6 this limit is compared with those obtained by other leading experiments in the field. Although the limit presented here is not the world's best at the time of completion of this thesis due to the neutron background from poor shielding and the inability to raise the temperature to 40°C in order to achieve better efficiency (Section 6.4.3), it is comparable to the latest limits. Whereas most experiments are concerned solely with the SI interaction, this fluorine-rich bubble chamber successfully sets good SD WIMP limits. The 15 kg chamber that is the subject of this thesis also provided additional confirmation of the acoustic alpha-neutron discrimination technique. Lastly, the feasibility of the windowless bubble chamber design was demonstrated, as were the strength of the bubble chamber technology in general and its potential for improvement in the future.

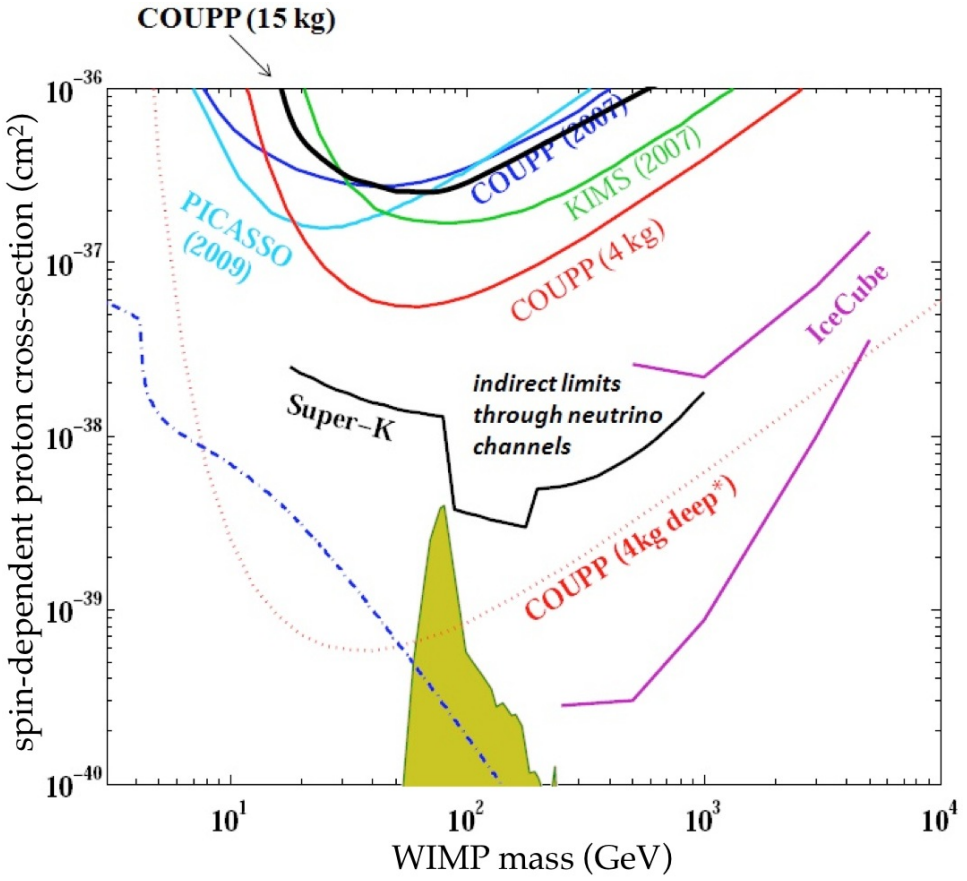

Figure 9.6: The SD-proton WIMP limit from the 15 kg chamber in the context of all other leading experiments which produce the same type of limits. Fluorine-rich experiments do best because an odd number of protons leaves a nucleus with a strongly uncanceled spin. The phase space above a curve is the area of cross-section ruled out with 90% confidence. "COUPP 4 kg deep" shows what should be achieved by that chamber gathering 400 kg-days of data at 40°C at the 6000 m.w.e. depth of SNOLAB, where it has already begun data-taking at time of writing. The Super-Kamiokande and IceCube limits (IceCube limits stemming from two different theoretical assumptions) assume dominant WIMP annihilation into neutrinos and are therefore not as general as limits set by direct detection experiments despite appearing stronger. The yellow-shaded region constitutes a supersymmetric WIMP parameter space and the blue dashed line is the new Maverick WIMP model developed at the University of Chicago. Figure adapted from Behnke et al. (2011), with 15 kg limit added. See references therein for supersymmetric and Maverick models.

Appendix A
Determination of Ultrasound Cleaning Parameters

Three variables were considered for the ultrasound cleaning of the quartz: frequency, power, and time. Temperature was fixed at 60°C, recommended for quartz in water (O'Donoghue, 1984). Combinations of 25 and 45 kHz, 1 and 10 minutes, and 10 and 100% power were explored. One quartz sample served as the control, and one additional sample was treated at 25 kHz (more cleaning and more damage than at 45 kHz according to O'Donoghue (1984) and Cole-Parmer (2010)) at 100% for 60 minutes, i.e., a "full-blast" treatment. Samples were all cleaned in the Cole-Palmer A-08871-30 ultrasound bath power with 3.7 kW maximum power.

Samples came from fragments of GE 214 quartz. Prior to testing, they were first manually agitated in a beaker of tap water to remove coarse dust motes. They were then cleaned with isopropyl and acetone, but allowed to dry to avoid scratching. All samples were immersed at the center of the bath and a motor was used to rotate them to ensure equal sampling of ultrasound nodes. Past Collar lab students had observed clear macroscopic patterns of periodic damage when this precaution was not taken. Afterwards, samples were rinsed with still water to remove detergent. Each was then wrapped in clean room cloth and bagged awaiting analysis.

The Atomic Force Microscope (AFM) at the Physics Department of the University of Chicago was used to analyze each sample. In an AFM a probe is lowered so close to a surface that the van der Waals interaction between the atoms in the tip and atoms in the sample becomes measurable. In the mode used for this analysis, the probe was set to harmonically oscillate. Knowledge of the alternatingly repulsive and attractive force provides knowledge of the distance. An electron microscope provides a faster scanning technology, but is compatible only with electrically conducting surfaces. Quartz can be coated with a

conductor, but this process could potentially mask any surface damage present.

Scan areas of 200 $\times$ 200 nm, 500 nm, 1 μm, 4 μm, 20 μm, 40 μm, and 100 $\times$ 100 μm were explored, in addition to one-dimensional segments. Using the optical microscope built into the AFM, areas appearing relatively smooth were sought to avoid the bias of examining rough features only local in extent.

A rough trend was evident: damage in the form of holes that exhibited a periodicity of 20–50 nm increased in severity with increasing power or time. To err on the side of caution, the treatment prescribed for the 15 kg chamber was as follows. The vessel was soaked in 60°C Radiacwash solution for 30 minutes, with the bath recirculating through a filter. This allowed the Radiacwash chemistry to work (Biodex, 2009). This was followed by a minimally aggressive ultrasound treatment of 10 min. at 45Hz and 10% of the power to facilitate detachment of remaining particulates.

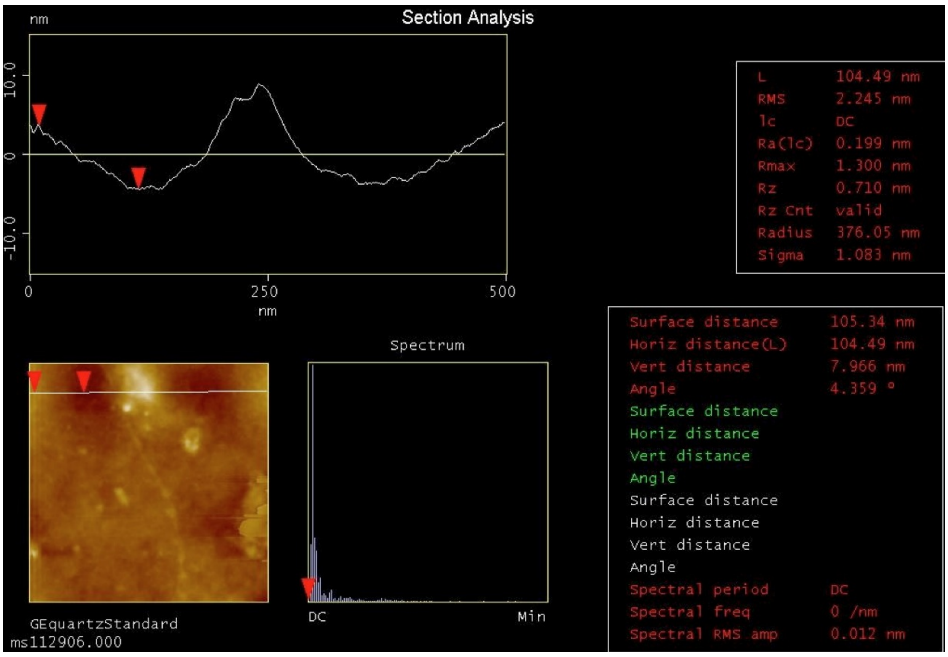

(a)

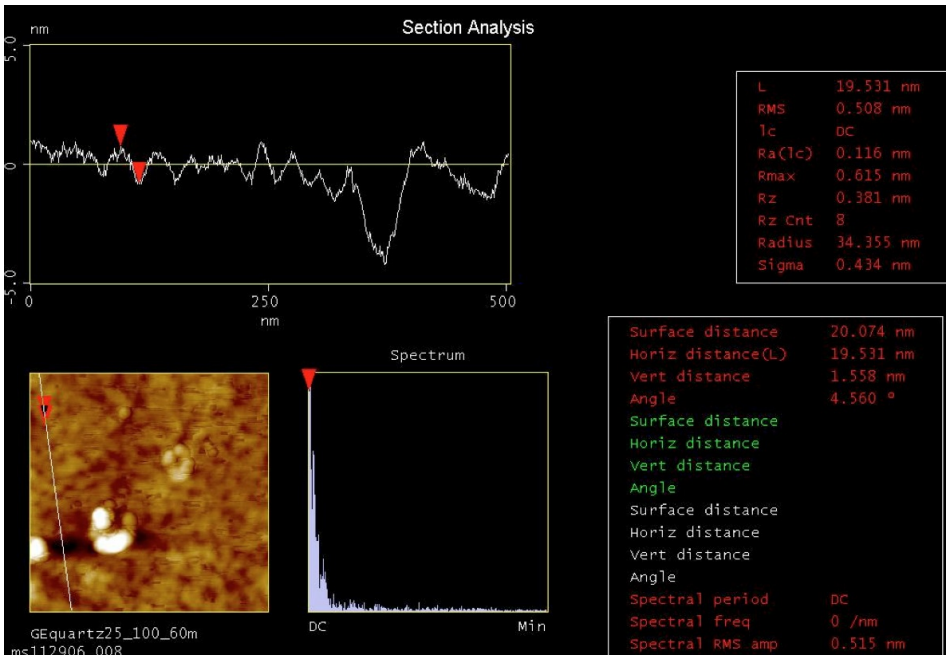

(b)

Figure A.1: (a) 1-D trajectory across the face of the no-ultrasound reference sample (trace at upper left). Insets at lower left corners are 500×500 nm segments with 30 nm height scale. Test conditions (frequency, power, time) are summarized beneath the insets. The 1-D slices used can be seen there as white lines.
(b) Further evidence of small-scale scoring of the "full blast" sample, given the longest and harshest treatment. Intense pitting is evident here, absent in (a). Typical peak-to-peak of the "dents" is ∼O(1 nm). The period of the trace at upper left is in rough agreement with a visual impression of the "landscape" at lower left.

Appendix B
Firewire Camera Feedthroughs for the Pressure Vessel

Camera firewire cables were stripped and tested because wires had to be sent through the top flange of the outer vessel of the 15 kg chamber via feedthroughs. There do not exist high-pressure feedthroughs which keep firewire cables intact.

The plastic insulation and braided metallic grounding shield were removed by razor. A Unibrain Fire-i 2056 camera like the ones to be used in the chamber was connected via an untouched cable to a repeater from which two different firewire cables led to a computer, one untouched as well, the other stripped. This permitted easy cable-swapping between tests without disturbing the position of the camera, which was screwed into place in a plastic cylinder (Figure B.1). During actual testing, a brick atop the cylinder (only partially obstructing the view) made absolutely certain that no motion could corrupt the results.

Figure B.1: Camera for test, in black nylon cup.

Per each distinct setup, a total of 500 black and white 480×640 images were collected. There were three test conditions: normal cable, as the control; stripped cable with untwisted

wire pairs; and, lastly, stripped cable with twisted pairs intact. Pixel map subtractions were performed, with one row selected for analysis. This row had the most variation in brightness, containing parts of the room, ceiling lights, and the black cylinder.

The average intensity was calculated for the chosen row in all image differences and histogrammed. The average and standard deviation of this set of averages was in turn found. Table 1 summarizes the resulting data, in chronological order.

Firewire Cable Testing		
	Mean (0–255)	σ (0–255)
Control1	2.2	0.1
Mangled1	2.2	0.1
Mangled2	2.2	0.2
Mangled3	2.6	0.8
Control2	2.2	0.1
Retwist1a	2.1	0.1
Retwist1b	2.1	0.1
Retwist1c	2.2	0.1
Retwist2	2.2	0.1
Retwist3	2.2	0.1

Table B.1: A control data point was taken (Control1), followed by three data points with the mangled cable, the originally twisted pairs having been pulled far apart (Mangled1-3), another control point (Control2), and three data points after reestablishing the twists (Retwist1a-c). An attempt to then untwist to take more data failed, as the camera was no longer recognized. Surprisingly, retwisting again was successful (Retwist2), but untwisting once more led to the same error (timeout). Finally, retwisting one last time, the camera worked yet again (Retwist3).

Prior to even beginning image analysis, there was a noticeable "jitter" in the image when the camera was attached to the computer via the cable with untwisted, completely separated wires. By moving the cable one could control the effect. Three cases (Mangled 1, 2, and 3) demonstrated varying degrees of jitter, evidenced by the appearance of a secondary

peak (Figure B.2). The image shifted between frames, changing what the selected row was displaying. This effect was perceived live by eye.

Ultimately, using insulating firewire cable did not have a measurable effect upon noise. However, untwisting the twisted pairs resulted in a noticeable pixel jumping phenomenon or temporary operational failure. Retwisting restored normal operation. This conclusion led to the purchase of twisted-pair feedthroughs instead of the less expensive feedthroughs designed for individual wires.

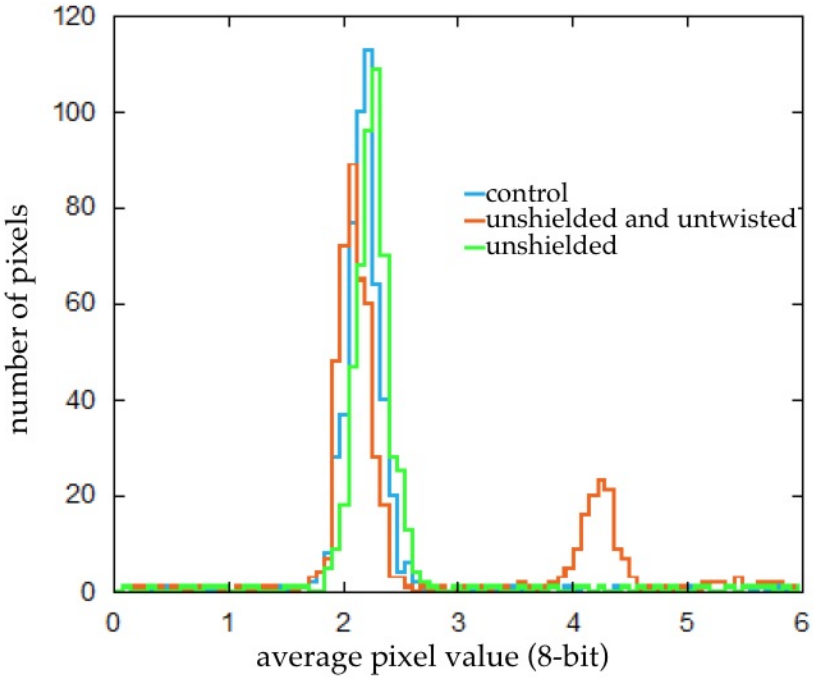

Figure B.2: The comparable histograms for control and unshielded twisted pairs, which appear to be single-peak, normal distributions. Contrast with the "mangled" example (orange) where an additional feature appears.

Appendix C
Mineral Oil Degradation Tests

During the first tests of small prototype chambers with mineral oil as the hydraulic fluid, it became evident that mineral oil degrades under several conditions (including photodegradation; compare with CF_3I in 2.6.3), all of which must be addressed in order to allow the cameras to clearly see bubbles.

Black butyl rubber gaskets caused the clear mineral oil to turn first yellow, then increasingly dark shades of orange. After research and testing of several materials, new gaskets were fashioned from a light-colored, oil-resistant mixture of Buna-N and vinyl rubbers (Section 4.1.3). Virgin oil could still yellow after the gasket change, though not as severely.

A free-radical inhibitor (Vitamin E) included with the commercial oil would absorb oxygen from the air and turn yellow in the presence of no other liquids or solids, given sufficient heat or visible or UV light, and this would occur faster the more intense the heat or light (Faust, 2007). Mineral oil which sat exposed to fluorescent light at room temperature remained clear after one year, but oil heated to 60°C would perceptibly yellow within hours. The maximum planned temperature for the 15 kg chamber, 40°C, did not cause significant yellowing; however the LEDs slated for placement at the bottom of the chamber were a cause for concern.

At maximum intensity, an LED cluster of the type intended for use in the 15 kg chamber exposed to air without a heat sink would reach 100°C within minutes. Submerged in mineral oil, which would help conduct heat away from the LEDs, and placed on a metal surface to act as heat sink (Al plate in 4.1.1), individual LED lenses could still produce temperatures in excess of 50°C. Not only did this lead to oil discoloration, but light output suffered because LED lenses became covered with a soft white gel. At first, it was thought that a chemical reaction was taking place between the plastic lenses and the mineral oil. However, a close

inspection confirmed that the solid lenses were still intact underneath the gel build-up. This problem had to be addressed to prevent loss of all light in the chamber over time.

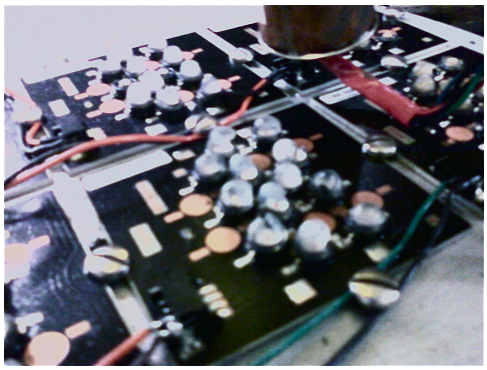

Figure C.1: Example of the buildup of the opaque white substance apparently due to oil solidification at the hotspots corresponding to individual LEDs. The more transparent lenses have less buildup.

LEDs encapsulated under 3 cm of epoxy demonstrated a similar build-up on the flat top surface of the cube of epoxy. Gel formed at points corresponding to the locations of individual LEDs underneath. This showed that even when there were no local hot spots in contact with oil, the light from the LEDs caused this strange formation in the brightest areas. LED encapsulation was clearly not a sufficient solution. It slowed but did not prevent the gel formation, and over the months-long course of a dark matter experiment it could still build up significantly. It also reduced light output from refractive index mismatch with the epoxy and made the light less parallel to the chamber, reducing the effectiveness of the lighting scheme (Section 3.1).

Tests of LEDs in both inhibited and uninhibited oil proved that the inhibitor was preventing what was apparently a light and heat-catalyzed polymerization of the mineral oil. In the case of a single LED lit in a 10 mL container of uninhibited oil, a 1 cm tower of white substance appeared overnight, visibly lowering the oil level. The inhibited oil with submerged active LED did not show any signs of polymerization for a week, and after one month still did not reach the level measured in the uninhibited sample (which essentially

remained the same). Different inhibitor chemicals (recommended by Sims et al. (1973), for a different oil) and varying concentrations in the oil were tested and accelerated aging tests were performed by heating the oil and inundating it with UV light to speed polymerization. (In the presence of no LEDs, uninhibited oil would turn milky white during such a test and contain floating fragments of a flaky white solid.) No advantage was observed in using an inhibitor other than the Vitamin E already present in commercially-available oil. Increasing the concentration would slow the polymerization, but not significantly, and was not an ideal solution because it was the inhibitor which turns the oil yellow, not the oil itself (Faust, 2007). More inhibitor meant darker, more opaque oil.

The new solution devised and ultimately implemented for the 15 kg chamber was to reduce the intensity of the LEDs and flash them in time with the cameras to minimize the light exposure of the oil. A low duty cycle would naturally also decrease the excess heat. The camera model used was not advanced enough to permit good synchronization among all six cameras used at one time, so the LED duty cycle had to be much higher ($\sim$50%) than the original 3% thought possible. However, use of a piezo trigger (Section 5.6.3) made it possible to have the LEDs on $\ll$1% of the time at low intensity. The piezo trigger did not make the efforts expended for oil integrity preservation unimportant, however, because the cameras were still needed for position reconstruction, and for triggering during the decompression phase.

The oxygen-induced yellowing, aggravated by oil's affinity for air (Logvinyuk et al., 1970; Kubie, 1927), was addressed with aggressive degassing, already necessary for the experiment for other reasons. The fact that the 15 kg chamber outer volume was a closed hydraulic system also meant that there would be no environmental oxygen issue, and nitrogen was used to force the oil into the pressure vessel during filling (Section 4.5.3). At the end of the experiment, the 15 kg chamber oil was only minimally yellowed after one year and

polymerization on the lenses was not well advanced. During operation, no significant drop in lighting quality was observed.

Appendix D
CF_3I Refractive Index Measurement

In order to perform an accurate three-dimensional reconstruction of bubble position in a chamber, it is beneficial to know the index of refraction of the active liquid, CF_3I in this thesis, so that one is not forced to rely solely upon strictly empirical methods (Chapter 6). In 2005, summer REU (Research Experience for Undergraduates) student Gerald Onuoha, under the direction of Prof. Collar, measured the refractive index of CF_3I, not found in the literature to date, by pointing the beam of a laser through a glass flask of CF_3I (Onuoha, 2005). Figure D.1 is a photo of his setup. The angle of incidence of the laser on the flask was changed, and the position of the beam at the other side was measured. Simple geometric ray tracing can show that the slope of the dependence of the latter on the former was the index of refraction. Red LEDs illuminate COUPP bubble chambers, so a red laser was used in order to measure the index at the relevant color (refractive indices are known to be functions of wavelength). Moreover, values for the index for other compounds for red light were readily available as references.

The method used here was validated blindly by determining the known indices of several substances (Figures D.2 and D.3). The results were in excellent agreement with manufacturer data when compared. The measured temperature dependence of the index of refraction for CF_3I is shown in Figure D.4. The vapor pressure was used as the primary means of determining the temperature, but it was cross-checked against a thermometer. It is fortuitous that the index of CF_3I is so close to that of water. This is convenient for open-air photography tests, where water can be used in place of CF_3I (Section 6.1).

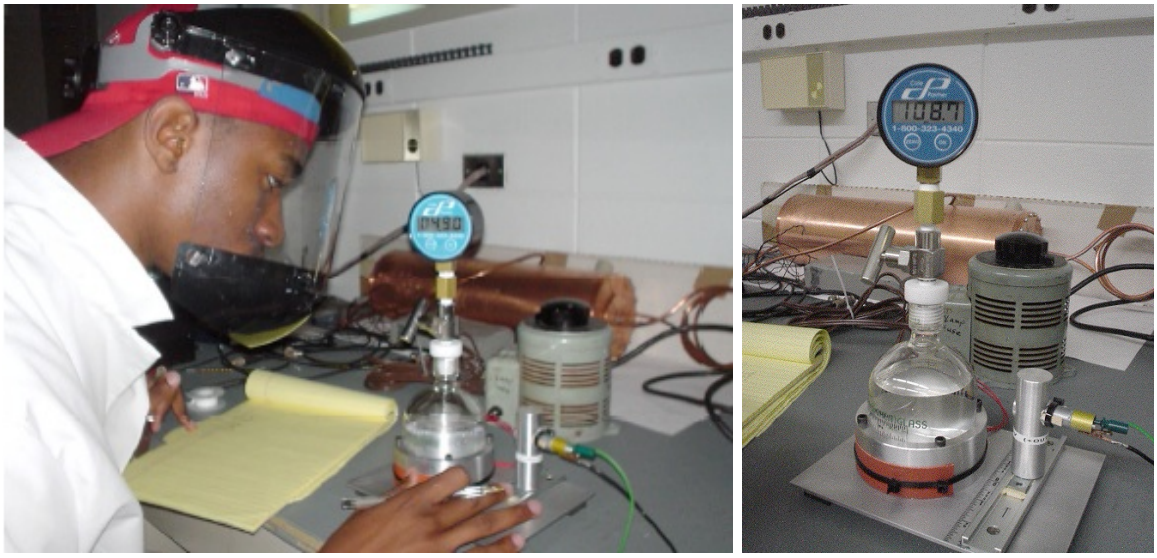

Figure D.1: The setup for the index measurement experiment with quartz vial (of known dimensions and index) containing CF_3I. Notice the laser (right of vial), ruler, and pressure gauge (for determining the vapor pressure). Figure from Onuoha (2005).

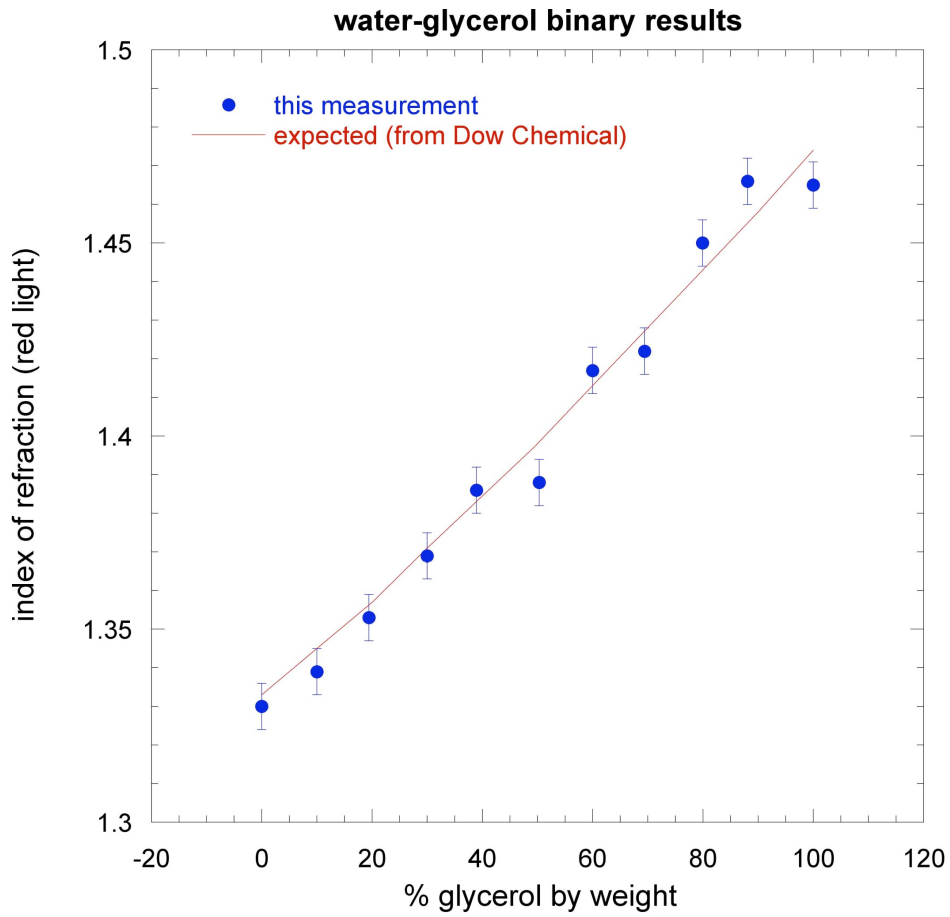

Figure D.2: One of the many validating calibrations of the refractive index-determining experimental setup demonstrating an accurate reproduction of the refractive index of a known substance: water with varying amounts of glycerol added. Figure from Onuoha (2005).

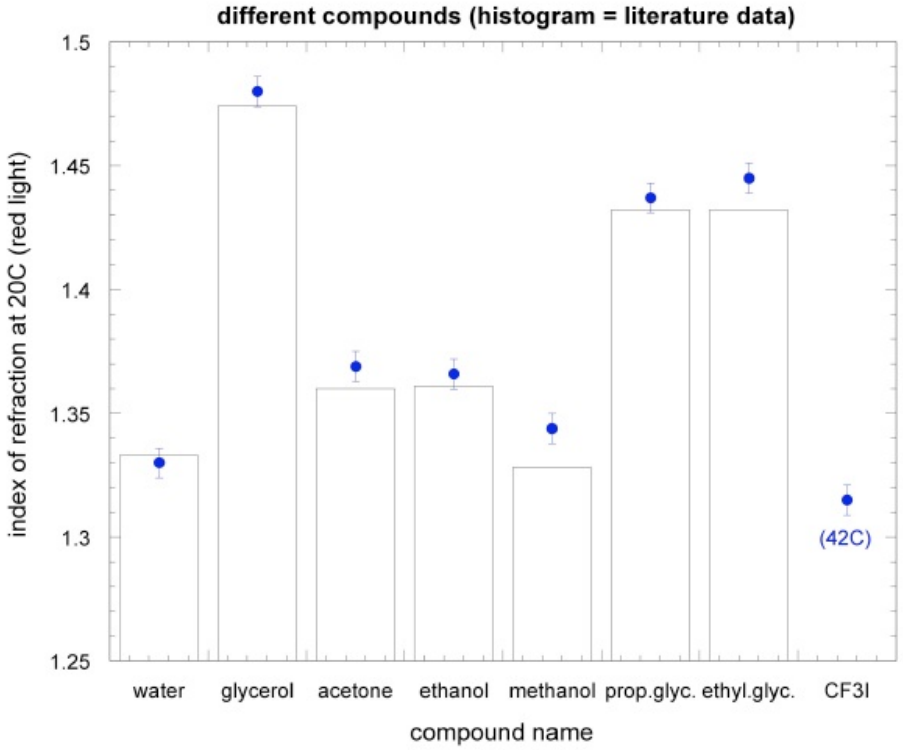

Figure D.3: Blind comparison of measurements of the index of refraction of several compounds to their known values. CF_3I data point is included at the far right. The other compounds constituted a preliminary test. Figure from Onuoha (2005).

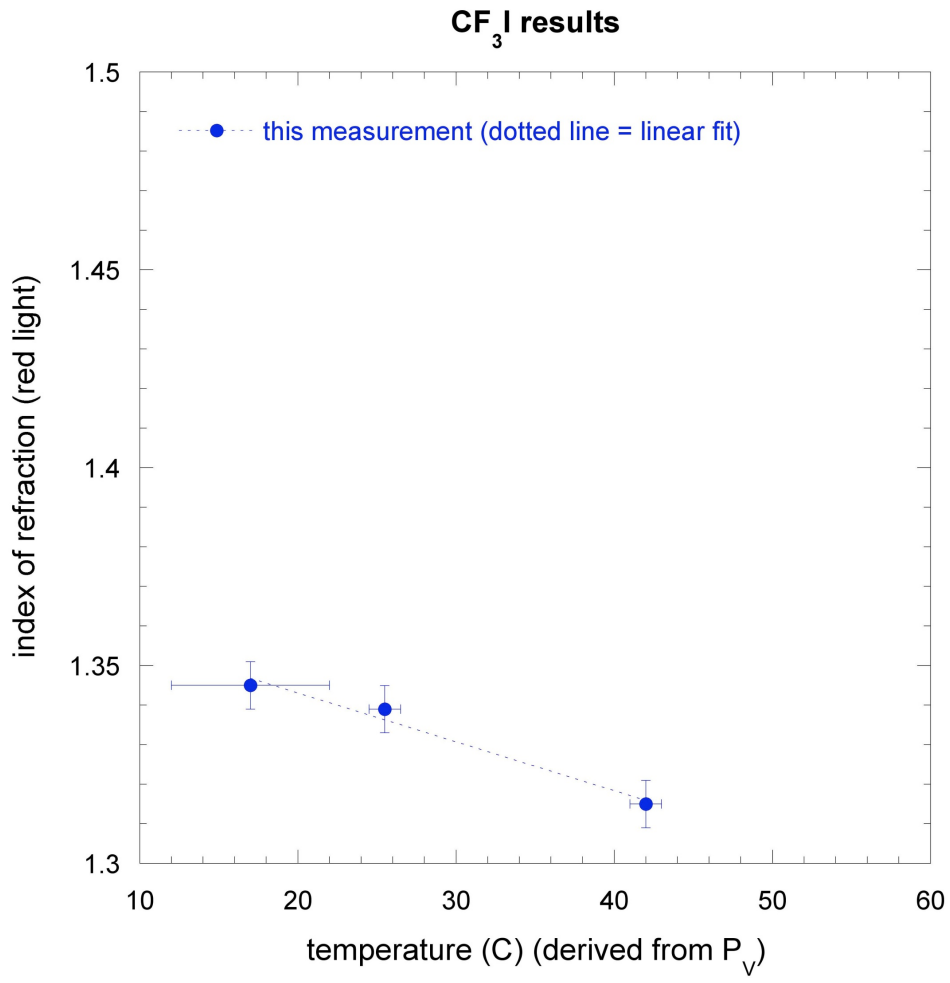

Figure D.4: Final result: the dependence of the index of refraction of CF_3I on the temperature. Temperature was derived from the vapor pressure and was in agreement with a thermocouple. The first data point comes from thermalization of the system with the lab. Others comes from active heating. Figure from Onuoha (2005).

REFERENCES

Agrawal, P., et al. 2010, arXiv:1003.5905v1 [hep-ph]

Ahlen, S., et al. 2010, IJMPA, 25, 1

Akerib, D. S., et al. 2004, Phys. Rev. Lett., 93, 211301

Apfel, R. E., et al. 1983, Rev. Sci. Instrum., 54, 1397

—. 1985, Phys. Rev. A., 31, 3194

Appelquist, T., et al. 2000, Phys. Rev. D., 64, 035002

Aprile, E., et al. 2010, arXiv:1001.2834v1 [astro-ph.IM]

Archambault, S., et al. 2009, Phys. Lett. B, 682, 185

Aubin, F., et al. 2008, New J. Phys., 10, 1

Baltz, E. 2004, J. High En. Phys., 10, 1

Bardina, J. 1988, Particles on Surfaces, ed. K. L. Mittal (New York: Plenum Press)

Barnabé-Heider, M., et al. 2005, Phys. Lett. B, 624, 186

Bazzotti, M., et al. 2009, arXiv:0910.4259v1 [hep-ph]

Bednyakov, V. A., & Klapdor-Kleingrothaus, H. V. 2001, Phys. Rev. D., 63, 095505

Behnke, E., et al. 2008, Science, 319, 933

—. 2011, Phys. Rev. Lett., 106, 021303

Bell, C. R., et al. 1974, Nucl. Sci. Engng., 53, 458

Bernabei, R., et al. 2010, arXiv:1007.0595v1 [astro-ph.CO]

Bertone, G., et al. 2007, Phys. Rev. Lett., 99, 151301

Bichsel, H., et al. 2006 (Phys. Lett. B), 258–270

Biodex. 2009

Blander, M., & Katz, J. L. 1975, AIChE J., 21, 833

Bogdanova, L. N., et al. 2006, Phys. Atom. Nucl., 69, 1293

Bolte, W. J., et al. 2005, Nucl. Instrum. Meth. A., 577, 569

—. 2006, J. Phys.: Conf. Ser., 39, 126

Bond, L., et al. 2005, Nucl. Phys. B (Proc. Suppl.), 138, 68

Borodai, F. Y. 1989, Glass and Ceramics, 46, 463

Boukhira, N., et al. 2000, Astroparticle Phys., 14, 227

Brownstein, J. R., & Moffat, J. W. 2007, Mon. Not. R. Astron. Soc., 382, 29

Bugaev, E. V. 1998, Phys. Rev. D., 58, 054001

Busto, J., et al. 2002, Nucl. Instrum. Meth. A., 492, 35

Cadonati, L. 2001, PhD thesis, Princeton University

Cao, Q. H., et al. 2009, arXiv:0912.4511v1 [hep-ph]

Carey, V. P. 1992, Liquid-Vapor Phase-Change Penomena (Washington: Hemisphere)

Chang, J. H. 2000, Online plotter for MCNP and ENDF cross section data http://atom.kaeri.re.kr/cgi-bin/endfplot.pl

Chu, S. Y. F., et al. 1999, WWW Table of Radioactive Isotopes http://nucleardata.nuclear.lu.se/nucleardata/toi/radSearch.asp

Clowe, D., et al. 2006, Astrophys. J., 648, L109

Cole-Parmer. 2010

Collar, J. I. 1996, Phys. Rev. D., 54, R1247

Collar, J. I., et al. 2000, New J. Phys., 2, 14.1

—. COUPP collaboration internal note

Das, M., & Sawamura, T. 2004, Nucl. Instrum. Meth. A., 531, 577

Das, P. K., et al. 1987, Int. J. Heat Mass Transfer, 30, 631

Das Gupta, N. N., & Ghosh, S. K. 1946, Rev. Mod. Phys., 18, 225

Donnelly, M., et al. 2004, CF_3I Stability under Storage, Tech. Rep. Technical Note 1452, NIST

Dujmic, D., et al. 2010, J. Phys.: Conf. Ser., 203

Ellis, J. 1991, Phys. Lett. B, 263, 259

Faust, H. 2007, Private Communication with M. Szydagis

Freese, K., et al. 1999, arXiv:astro-ph/9904401v1

Gaitskell, et al. 2010, The Dark Matter Community Website http://dmtools.brown.edu/

Gaitskell, R. 2004, Annu. Rev. Nucl. Part. Sci., 54, 315

Gibbs, J. W. 1957, The Collected Works of J. W. Gibbs, Vol. 1, p. 254 (New Haven: Yale University)

Glaser, D. A. 1952, Phys. Rev., 87, 665

—. 1953, Phys. Rev., 91, 762

—. 1954, Nuovo Cimento, 11, Suppl 2, 361

—. 1994, Nucl. Phys. B (Proc. Suppl.), 36, 3

Glaser, D. A., & Raim, D. 1955, Phys. Rev., 97, 474

Gratta, G., et al. 1998, Estimates of KamLAND Backgrounds from Radon Emanated by Photomultiplier Tube Glass

Habfast, C. 2010, Phys. in Perspective, 12, 219

Hall, J. M., et al. 2007, Nucl. Instrum. Meth. B, 261, 337

Harper, M. 1991, PhD thesis, University of Maryland

—. 1993a, Nucl. Sci. Engng., 114, 118

—. 1993b, Nucl. Instrum. Meth. A., 336, 220

Hildebrand, J. H. 1929, J. Am. Chem. Soc., 51, 66

Hinshaw, G., et al. 2009, Astrophys. J. Suppl. Ser., 180, 225

ICRP. 1994, International Commission on Radiological Protection (ICRP): Protection against Radon-222 at Home and at Work, 23rd edn., Vol. ICRP Publ. 65, Annals of the ICRP 23 (2) (Oxford: Pergamon Press)

Jungman, G., et al. 1996, Phys. Rep., 267, 195

Kim, S. K. 2008, J. Phys.: Conf. Ser., 120, 1

Kolb, E. W., & Turner, M. 1990, The Early Universe (Westview Press)

Kozlov, V. Y., et al. 2010, Astroparticle Phys., 34, 97

Kubie, L. S. 1927, J. Biol. Chem., 72, 545

Lau, K. 2008, OCPA Workshop, University of Houston

Leonard, D. S., et al. 2008, Nucl. Instrum. Meth. A., 591, 490

Levine, I. 2010, COUPP collaboration private communication

Lewin, J. D., & Smith, P. F. 1996, Astroparticle Phys., 6, 87

Logvinyuk, V. P., et al. 1970, Chem. Tech. of Fuels and Oils, 6, 353

Martynyuk, Y., & Smirnova, N. 1991, Sov. Phys. Acoust., 37, 376

Mayet, F., et al. 2002, arXiv:astro-ph/0201097v1

McKinsey, D. N., & Doyle, J. M. 2000, J. Low Temp. Phys., 118, 361

Nyden, M. 1995, Fire Suppression System Performance of Alternative Agents in Aircraft Engine and Dry Bay Laboratory Simulations, Vol. 1, ed. R. Gann, Vol. 1 (NIST), 77

Odom, B. 2010, COUPP collaboration private communication

O'Donoghue, M. 1984, Microcontamination, 2, 63

Ohno, N., et al. 2005, Tribology Trans., 48, 165

Onuoha, G. 2005, Measurement of index of refraction in liquid CF_3I

Otsuka, H. 2007, Patent 7232778, http://www.freepatentsonline.com/7232778.html

Pallavicini, M., et al. 2009, arXiv:0910.3367v1 [astro-ph.SR]

Pan, L. K., et al. 1999, Nucl. Instrum. Meth. A., 420, 345

Peccei, R. D., & Quinn, H. R. 1977, Phys. Rev. Lett., 38, 1440

Peyrou, C. 1967, Bubble and Spark Chambers, ed. R. P. Shutt (New York: Academic)

Piro, M.-C. 2010, arXiv:1005.5455v1 [astro-ph.IM]

Plesset, M. S., & Zwick, S. A. 1952, J. Appl. Phys., 23, 95

—. 1954, J. Appl. Phys., 25, 493

Pozzi, S. A., et al. 2003, Nucl. Instrum. Meth. A., 513, 550

Quarzglas, H. 2010, Chemical Purity

Refregier, A. 2003, Annu. Rev. Astron. Astrophys., 41, 645

Reinke, P. 1996, PhD thesis, Swiss Federal Institute of Technology

—. 1997, Exp. Heat Transf., 10, 133

Riess, A., et al. 1998, Astron. J., 116, 1009

Riley, N. 2007, Master's thesis, University of Chicago

—. 2009, COUPP collaboration private communication

Roy, S. C., et al. 1987, Nucl. Instrum. Meth., A255, 199

Rubin, V., et al. 1980, Astrophys. J., 238, 471

Saltiel, C., & Sokolov, M. 1982, Appl. Optics, 21, 4033

Seitz, F. 1958, Phys. Fluids, 1, 2

Shibata, K., et al. 2002, J. Nucl. Sci. Technol., 39, 1125

Sims, R., et al. 1973, Lipids, 8, 337

SNO. 2002, The SNO Detector http://www.sno.phy.queensu.ca/sno/sno2.html

Solomon, S., et al. 1994, J. Geophys. Res., 99, 20,929

Sony. 2010, http://www.sony.net/Products/SC-HP/datasheet/90203/confirm.html

Springel, V., et al. 2005, Nature, 435, 629

Szydagis, M. 2008, Rencontres de Moriond

Tovey, D. R., et al. 2000, MPI-PhT, 16

Trucco, E., & Verri, A. 1998 (New Jersey: Prentice Hall)

UNSCEAR. 2000, Sources and Effects of Ionizing Radiation, Annex B: Exposures From Natural Radiation Sources, http://www.unscear.org/docs/reports/annexb.pdf (United Nations)

v.Hentig, R., et al. 1999, Nucl. Phys. B (Proc. Suppl.), 78, 115

Wenniger, H. 2004, CERN Courier

Williamson, K. J. 1994, Macroscale and Microscale Organic Experiments, 2nd edn. (Lexington, Mass.: D. C. Heath)

Wojcik, M. 1991, Nucl. Instrum. Meth. B, 61, 8

Yuanyuan, D., et al. 1999, J. Therm. Sci., 8, 73

Zacek, V. 1994, Nuovo Cimento, 107A, 291

Ziegler, J. 1984, Ion Implantation Science and Technology, Vol. 1 (Academic Press, Inc.)

—. 2010, Particle Interactions with Matter

Zwicky, F. 1937, Astrophys. J., 86, 217