

ASPECTS OF THE R&D FOR THE ENRICHED XENON
OBSERVATORY FOR DOUBLE BETA DECAY

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF APPLIED PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Kolo Wamba
December 2006

UMI Number: 3242632

INFORMATION TO USERS

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleed-through, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.



UMI Microform 3242632

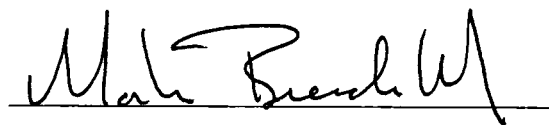
Copyright 2007 by ProQuest Information and Learning Company.

All rights reserved. This microform edition is protected against unauthorized copying under Title 17, United States Code.

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

© Copyright by Kolo Wamba 2007
All Rights Reserved

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

A handwritten signature in black ink, appearing to read "Martin Breidenbach", written over a horizontal line.

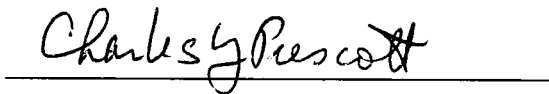
Martin Breidenbach
(SLAC) Principal Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

A handwritten signature in black ink, appearing to read "Walt Harrison", written over a horizontal line.

Walt Harrison
(Applied Physics)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

A handwritten signature in black ink, appearing to read "Charles Prescott", written over a horizontal line.

Charles Prescott
(SLAC)

Approved for the University Committee on Graduate Studies.

Abstract

Neutrinoless double beta decay is a nuclear decay process characterized by the emission of two β rays by an even-even atomic nucleus. Though required by many minimal extensions of the Standard Model, it is predicted to be extremely rare, and as such, it has yet to be detected experimentally. Its discovery might offer striking insights into some of the key parameters that describe the fundamental nature of the neutrino, not the least of which would be an absolute scale for its effective mass. The Enriched Xenon Observatory (EXO) is a proposed next-generation experiment designed to perform a high-sensitivity search for neutrinoless double beta decay in the isotope ^{136}Xe . EXO will use a 10-tonne sample of ^{136}Xe , probably condensed to the liquid state, and will identify double beta decay candidate events by the scintillation and ionization they generate in the enriched Xe sample. If a candidate event were an actual double beta decay, it would also be characterized by the formation of an ion of ^{136}Ba . Moreover, if the ^{136}Ba ion can be detected in coincidence with the ionization and scintillation, this would be a powerful means of distinguishing true double beta decay events from background. This thesis presents the results of some of the R&D studies done to establish the feasibility of operating a large liquid Xe-based ionization detector from which ions can be individually isolated and detected. Three key technologies relevant to such a system are described. The first is liquid xenon purity and its effect on electron transport in liquid xenon. The second is positive ion transport in liquid xenon. On this subject a set of theoretical and experimental studies of ion transport in liquid Xe is presented. The third technology discussed is a possible scheme for lifting an individual ion out of the liquid phase and loading it into an ion trap where it could be identified spectroscopically.

Acknowledgments

To begin with I would like to thank my research adviser, Prof Martin Breidenbach, who has been a thoughtful mentor and intellectual guide throughout this process. Despite his high profile, he was always willing to listen to my ideas, and encouraged me to independently explore whatever topics interested me. At the same time, he took my personal well-being to heart and worked tirelessly to help resolve the issues I was having with funding and juggling work and family.

I should also like to thank my EXO collaborators for offering advice, lending lab equipment, and being a sympathetic ear. In particular I would like to thank Profs Enrico Conti, Ralph DeVoe, William Fairbank Jr, Giorgio Gratta, Allen Odian, Andreas Piepke, Charles Prescott, Peter Rowson, Victor Stekhanov, and Petr Vogel. Each of them has been an inspiring scientific coach, offering important insights and assistance on issues ranging from the truly fundamental (“Can it be argued that θ_{13} must be zero?”) to the more mundane (“Why is my charge amplifier so noisy?”). Prof Gratta deserves special recognition for introducing me to EXO and for treating me like one of his own students. I also wish to thank the members of the Gratta group, Peter Fierlinger, Bjorn Flatt, Matt Green, Francisco LePort, Russel Nielson, Kevin O’Sullivan, Andrea Pocar, Sam Waldman, and Jesse Wodin, who were always eager to help and never grew tired of my requests to borrow their spot-welder. Many thanks also go to Carter Hall and Derek Mackay, for many helpful discussions, and to Knut Skarpaas for being a great resource. I also want to thank Bob Conley for his technical wizardry and for being a good mentor, confidant, and friend outside the lab.

Sincere thanks go to the members of my dissertation committee; specifically, Profs

Charles Prescott and Walt Harrison, for their scientific insights, and Profs Godfrey Mungal and Herman Winick for having agreed to serve on such short notice.

I should also like to thank the administrative staff, namely Paula Perron and Claire Nicholas in Applied Physics, and Judy Meo, Lilian DePorcel, and Natasha Haulman at SLAC, for always making me feel like I was a top priority.

Among the other SLAC staff I wish to thank are Mark Freytag and Lupe Salgado, for helping me debug electronics at 5:00pm on Friday afternoons, and Profs Ed Garwin and Robert Kirby for lending lab equipment and giving good advice.

Further thanks go to my friends and neighbors, in particular Ryan and Jackie Christensen, and especially Todd Lee and Jackie Anderson, who were there when the going was tough. I also wish to thank my Applied Physics colleagues Anu Tewary and Alan Fang for helping me through a punishing curriculum.

On a similar note I should like to acknowledge Prof Art Walker, now deceased, who was instrumental in bringing me to Stanford and showing me the ropes. He continues to be sorely missed.

The love and support of an amazing family has sustained me through a grueling 7 years, and cannot go without mention. Heartfelt thanks go to my extended clan on the Brown and Wamba sides, as well as the Depelchin-Wynters of Berkeley. I would also like to thank my brothers Philippe and James, especially Philippe, who sadly, did not live to see the completion of my studies, and I dedicate this work to his memory. I also wish to express my sincerest thanks to my incredible parents, Prof Ernest Wamba-dia-Wamba and Elaine Wamba, for their innumerable sacrifices and unwavering faith in the face of 4, 5, and then 6 years of graduate school that came and went with no end in sight. I should also like to thank my 3-year-old daughter Luezi, for helping to keep me from becoming too preoccupied with things like turbopumps and ROOT macros. Finally, I would like to thank my wife and partner in life, Alice Endamne, whose love, hope and encouragement never seemed to run out. Her support was particularly important at the times when it seemed as if the only science I was discovering was my own scientific incompetence.

To conclude, I wish to acknowledge the financial support of the Ford Foundation and the United States Department of Energy.

Contents

Abstract	v
Acknowledgments	vi
1 Introduction	1
1.1 References	3
2 Physics Motivation	4
2.1 Neutrino Basics	4
2.2 Neutrino Mass & Mixing	7
2.3 Neutrinoless Double Beta Decay	15
2.4 Nuclear Physics of Double Beta Decay	18
2.5 Two-Neutrino Double Beta Decay	19
2.6 Other possible mechanisms for $0\nu\beta\beta$	21
2.7 Detecting Double Beta Decay	21
2.8 EXO: a next-generation double beta decay experiment	25
2.9 Laser tagging	27
2.10 Conclusion	30
2.11 References	30
3 Electron Drift in Liquid Xenon	34
3.1 Basic Concepts	34
3.2 The Xenon Purity Monitor: Theory of Operation	35
3.3 The purifier	42

3.4	Experimental Procedure	44
3.5	Results	49
3.6	Discussion	56
3.7	References	60
4	Positive Ion transport in Liquid Xenon	62
4.1	Experiment	63
4.2	Results	69
4.3	Discussion	73
4.4	Conclusion	73
4.5	References	73
5	Computerized Studies of Ion Mobility in Liquid Xenon	75
5.1	Theory of Positive Ion Transport	75
5.2	Computer Calculation	78
5.3	Simulation Results	82
5.3.1	Radial Distribution Functions	82
5.3.2	Ion Mobilities	85
5.4	Discussion	91
5.5	Implications for EXO	95
5.6	Conclusion	96
5.7	References	96
6	Ion manipulation	99
6.1	The Cryogenic Probe	100
6.2	Experiment	100
6.3	Results	103
6.4	Discussion	105
6.5	Conclusion	107
6.6	References	107

7	Conclusions	109
7.1	Toward the Next Generation	109
7.2	Full-Scale EXO with Laser Tagging	110
7.3	References	112
A	Experimental Setup	113
A.1	LabView® Programs	113
A.1.1	Coolit3.vi	113
A.1.2	grabbercell.vi	117
A.1.3	newHFElowlimit.vi	117
A.1.4	FLOWMETER.vi	120
A.1.5	mobility.vi	122
A.1.6	probemover.vi	124
A.1.7	grabberVALVE.vi	124
A.2	purity.exe	125
A.2.1	Using purity.exe	127
A.2.2	Algorithm of purity.exe	129
A.2.3	Sources of systematic error on the XPM measurements	132
A.3	References	134
B	Molecular dynamics simulation	135

List of Tables

3.1	Results of initial round of purity tests using non-recirculation capable system. Error bounds are statistical only.	51
4.1	Electric fields (E), voltages (V_{drift}), and drift distances (d) used in this experiment, showing their associated systematic errors. The measured transit times (t_{tr}) and velocities (v) are the values returned by the fit. For these last two quantities, only statistical errors are given.	71
5.1	Ion mobilities in LXe, comparing a model due to Hilt, Schmidt and Khrapak (HSK) [5] to the experimental data.	78
A.1	A table of the quantities stored in the “int.txt” data file. The entries in rows 2-8 in the table list the numbers that characterize the baseline-subtracted XPM waveform in the order that they appear on the purity.exe front panel. These are the 7 fit parameters that correspond to the constants in equation 3.22 in Chapter 3.7.	131

List of Figures

2.1	Two possible neutrino mass spectra that fit the available data (not to scale), reflecting the fact that the sign of the smaller mass splitting is known while the sign of the larger is not. The spectrum on the left is commonly referred to as the “normal hierarchy” and the one on the right is “inverted”. The three red arrows to the right of the mass spectra indicate experimental limits on the neutrino masses determined from studies of the cosmic microwave background [23], kinematic measurements of the β^- decay of tritium [19], and neutrino observations of the 1987 supernova [24]. Figure adapted from Ref. [25].	15
2.2	(From Ref. [22]) ν_e effective mass ($m_{\beta\beta}$) as a function of the mass of the lightest mass eigenstate for the case of normal and inverted mass hierarchies. Note that the two scenarios become indistinguishable when the lightest mass becomes comparable in size to the mass splittings. Under these conditions, the mass hierarchy is said to be <i>degenerate</i> . The horizontal and vertical lines show different experimental limits arising from $0\nu\beta\beta$ measurements (horizontal lines) and cosmological observations (vertical lines). The white bands within the dark ones are the regions in which the third, non-Majorana phase in the PNMS matrix is nonzero. This is where CP is violated independently of the Majorana character of the neutrino.	16
2.3	Underlying diagram for neutrinoless double beta decay.	16
2.4	Summed electron energy spectrum for $2\nu\beta\beta$ and $0\nu\beta\beta$	23

2.5	The plots on the left show double beta decay electron summed energy spectra for the case of perfect energy resolution (red curve), 2 % FWHM energy resolution (black curve) and 10 % FWHM (blue curve). The abscissa in this case is given in units of the electron endpoint energy Q . The gray shaded region is the FWHM signal window for the case of 10 % energy resolution. Events within this window are counted as $0\nu\beta\beta$ signal. As shown at right, the amount of contamination by $2\nu\beta\beta$ events in the $0\nu\beta\beta$ signal window depends strongly on the detector energy resolution [40]. For these calculations, the $0\nu\beta\beta$ rate was assumed to be 100,000 times smaller than that of the $2\nu\beta\beta$, and backgrounds to $0\nu\beta\beta$ from sources other than $2\nu\beta\beta$ were not considered.	24
2.6	Liquid xenon TPC concept. Figure taken from Ref. [31].	26
2.7	EXO-200 signal scheme. The ionization electrons are drifted and collected on crossed wires, resulting in signals that give the energy and the x - and y - position of the event. Simulated examples of what these signals might look like are shown as Channel 168 and Channel 119. The z - position of the event is found using the time between the detection of scintillation signal by the avalanche photodiodes in the APD plane and the arrival of the ionization signals on the wires. This is a composite of figures found in Refs. [32, 33, 34].	28
2.8	Simplified level scheme for singly ionized barium.	29
3.1	The Xenon Purity Monitor (XPM), a device for measuring electron lifetime in liquid Xe. The three grids, going from top to bottom, are the anode grid, the cathode grid and the window grid. The anode, anode grid, window grid, and cathode are each connected via a coaxial cable to a biasing network and a preamplifier for reading out the charge signal.	37
3.2	Circuit diagram of the XPM.	39
3.3	Example of the XPM signal.	40
3.4	The SAES getter. Figure taken from Ref [12].	43

3.5	Simplified schematic of the gas handling system.	45
3.6	Equipotentials in the XPM [17], as calculated by the MAXWELL® software.	48
3.7	Electron drift speed as a function of electric field in liquid xenon. . .	49
3.8	The effect of contamination by room air on the electron lifetime. The arrows indicate the following: (1) introduction of 10 cm ³ of room air; (2) recirculation initiation; (3) re-measurement and subtraction of the baseline; (4) electron lifetime approaching the sensitivity limit.	53
3.9	Electron lifetime vs. time with Teflon® in the XPM. The green arrows indicate times when recirculation was started and red arrows indicate when it was stopped.	54
3.10	Left: Fit of Equation 3.19 to electron lifetime trend while recirculation is turned on. Right: Fit to data taken with recirculation off. Error bars are statistical only.	54
3.11	Electron Lifetime vs time for LXe sample containing an LAAPD, ribbon cables and epoxy. Green arrows indicate the start of recirculation; red, the interruption.	55
3.12	Left: Field dependence of electron attachment coefficient for the species N ₂ O (squares), O ₂ (circles), and SF ₆ (triangles) (figure taken from Ref. [8]). Right: The field dependence of the electron lifetime as measured by the SLAC XPM after pumpdown from ambient air.	56
4.1	Schematic of the probe cell. The lower half of the cell is insulated from the room-temperature laboratory by a large vacuum vessel (not shown). .	64
4.2	Left: Ion yield vs electric field for ions collected in 231 Torr of ambient Xe gas. Ions are collected even at zero field because they diffuse through the gas. Right: Ion yield vs electric field for ions collected in liquid. The source-probe distance for the gas experiment was 0.5 cm while for the liquid it was 0.025 cm.	65
4.3	Decay scheme for ²³⁰ U. Decays with branching ratios less than 5% are not shown.	67

4.4	Top left: α spectrum of the ^{230}U source in vacuum.[10] The ^{230}U peaks are at 5.82 MeV and 5.89 MeV, while ^{226}Th is identified by peaks at 6.23 MeV and 6.34 MeV. ^{222}Ra , ^{218}Rn , and ^{214}Po correspond to the peaks at 6.56 MeV, 7.13 MeV, and 7.69 MeV respectively. Top right: α spectrum of the activity collected on the probe in 0.2 bar of Xe gas. Bottom: Total α rate versus time for the activity collected on the probe. The best fit is the superposition of two exponentials with half-lives of 1773 ± 97.4 s and 37.2 ± 2.84 s. These are the ^{226}Th and ^{222}Ra half-lives, whose respective published values are 1830 s and 38 s [11].	68
4.5	Ion yield as a function of t_{drift} for the four electric fields 2.5 kV/cm, 3.9 kV/cm, 4.9 kV/cm, and 6.4 kV/cm. Error bars shown are statistical only. We take the transit time in each case to be the knee point in the fit.	70
4.6	Global fit to the ion drift velocity as a function of electric field for the entire data set.	72
5.1	Ion mobility vs temperature for four of the alkaline earth ions in LXe. The three curves are the results of calculations using the HSK model using three different estimated values for the (unknown) surface tension in the liquid Xe-solid Xe interface. σ_0 is the value estimated by HSK in Ref. [5] Figure taken from Ref. [2]	79
5.2	The liquid-liquid and ion-liquid interaction potentials used in the computer simulation. The black squares are the data points of Ref [12]. .	81
5.3	Radial distribution functions for: Ba^+ and Tl^+ in LXe at 161.4 K and 180.0 K, Xe_2^+ in LXe at 192.1 K, and pure Xe at 161.4 K. In each case, the red histogram is the initial RDF, the blue is the RDF computed after the 10 ps equilibration time and the black curve is the empirical model RDF for idealized pure LXe after Ref. [13].	83
5.4	Number of atoms in the cluster as a function of time for the ions Ba^+ and Tl^+ in LXe at 180.0 K and 161.4 K	84

5.5	Number of unexchanged atoms in the cluster peak as a function of time for: Ba^+ and Tl^+ in LXe at 180.0 K and 161.4 K, Xe_2^+ in LXe at 192.1 K, and pure Xe at 161.4 K. For the latter two cases, the “cluster” is defined as the narrow region within $\sim 0.5 \text{ \AA}$ of the first maximum in the RDF.	86
5.6	Histogram of the Ba^+ ion's displacement along x , y , or z after 3 ps for the freely diffusing ion under zero-field conditions. The diffusion constant is computed from the σ parameter of the fitted Gaussian. . .	87
5.7	Velocity autocorrelation function of ionic motion for: Ba^+ and Tl^+ , in LXe at 180.0 K and 161.4 K, Xe_2^+ in LXe at 192.1 K, and pure Xe at 161.4 K	89
5.8	Experimental and simulated ion mobilities for the five atomic species under study. The experimental data are from Ref. [2].	91
5.9	Ion mobility as a function of temperature for the simulated systems with no applied field (left panel) and with an electric field of 10^{11} V/cm . .	92
6.1	Left: Detail showing the internal structure of the cryogenic probe tip. The green arrows indicate the flow direction of the high-pressure Ar gas flow used to cool the tip. The darker arrow shows the direction of the inlet which is at high pressure. Joule-Thompson cooling occurs at the small orifice. Figure taken from Ref. [5]. Right: The Xe cell used for the cryogenic ion release probe experiment. For the tests involving LXe, the cell was insulated by a special vacuum jacket (not shown). .	101
6.2	Left: α particle spectrum for the radioactivity adhering to a bare metal probe (red histogram) and to a Xe ice-coated probe (blue histogram). Right: Count rate as a function of time corresponding to the red histogram shown at left. The smooth curve is a fit to the superposition of two decaying exponentials, whose half-lives as returned by the fit are indicated.	104

6.3	Time histograms of α counts for the ion release experiment and the null experiment. The solid black line indicates how the probe radioactivity would be expected to fall off in the absence of any ion release effect. It consists of a superposition of the exponential decay curves expected for ^{226}Th and ^{222}Ra whose half lives are, respectively 30.57 min and 38.0 s [6]. The indicated “dead” time is the time during which the probe is at the discharge station. During this time the detector is actually turned off.	105
6.4	Phase diagram of xenon [7].	106
A.1	Top half of front panel of the VI Coolit3.vi.	115
A.2	Bottom half of front panel of the VI Coolit3.vi.	116
A.3	Front panel of the VI grabbercell.vi.	118
A.4	Front panel of the VI newHFElowlimit.vi.	121
A.5	Front panel of the VI FLOWMETER.vi.	122
A.6	Front panel of the VI mobility.vi.	123
A.7	Front panel of the VI probemover.vi.	125
A.8	Front panel of the VI grabberVALVE.vi.	126
A.9	Screenshot of purity.exe.	128
A.10	(a) Example of an unaveraged XPM waveform (blue trace). The green trace in this and the other panels in which it appears is the laser Q-switch synchronous trigger. (b) Typical example of the average of 64 XPM waveforms. (c) Example of the averaged signal seen with the laser shutter closed. (d) The result of taking the difference between XPM waveforms (b) and (c). This baseline-subtracted waveform is fit to equation 3.22.	130
A.11	Ratio of cathode charge to anode charge in the XPM as a function of the ratio of the electric fields on either side of the cathode grid. . . .	133

Chapter 1

Introduction

In this thesis, some of the results and methodology of an R&D program leading to a next-generation double beta decay experiment, called the Enriched Xenon Observatory (EXO), are presented. EXO's stated purpose is a high-sensitivity search for neutrinoless double beta decay in ^{136}Xe . "Double beta decay" is the generic term for a nuclear process in which an atomic nucleus changes its nuclear charge by two units [1]. In one version of this process, two electron antineutrinos are emitted by the nucleus along with two beta electrons; we refer to this process as two-neutrino double beta decay. Much more interesting is the version of the decay in which no neutrinos are emitted, or neutrinoless double beta decay. This is because such a process, if discovered, would be clear evidence of new physics. Though EXO will ultimately be sensitive to both processes, its primary purpose is to study the latter.

EXO is currently under development and is an ambitious project involving some 50 individuals spread across 10 institutions in 4 countries. Because the final version of EXO will require as-yet undeveloped technology, a large fraction of the effort consists of R&D activities that are both planned and underway. Here, we deal with only a small subset of these studies in which the author directly participated, that were conducted at the Stanford Linear Accelerator Center (SLAC).

At present, the basic idea of EXO is to operate a detector that uses liquefied xenon isotopically enriched to 80% in its heaviest isotope, ^{136}Xe , both as the source of double beta decay and as the detection medium. double beta decay candidate events

are to be observed by using the ionization and scintillation properties of liquid Xe (LXe). In addition, EXO is to incorporate a novel set of techniques for rejecting the background to these events by detecting in coincidence the daughter ion that forms in the decay. The proposed scheme is to use single ion laser spectroscopy, or “laser tagging” as it is often called. This thesis is chiefly concerned with the R&D that is relevant to the background rejection scheme as well as some of the general techniques that are applicable to the general operation of a large scale, low background LXe ionization detector.

Essentially three double beta decay detector-related technologies are investigated in this work. The first is a method of generating and maintaining detector grade ultra-high purity liquid Xe. Closely associated with this is a means of precisely characterizing the purity of a liquid Xe sample and monitoring liquid Xe purity in real time over extended periods. These techniques are based on fairly well-established procedures in common use in the field.

The second two technologies deal with the problem of rejecting the radioactive backgrounds on an event-by-event basis, using real-time single-ion isolation and identification. Such a scheme has never been successfully implemented in any double beta decay research to date.

In Chapter 2 an overview of the physics goals of EXO is given, leading into a motivation for the basic experimental techniques whose demonstration forms the subject of this thesis. A terse summary of the relevant particle and nuclear physics at a level appropriate to the experimentalist is also included. In the following chapters a detailed discussion of the three above mentioned key EXO technologies is given.

Chapter 3 discusses techniques related to achieving long, low-loss electron drifts in liquid xenon. Essentially, this comes down to making the xenon as pure as possible and accurately measuring and characterizing the attained purity level. In this chapter a specially constructed liquid xenon purity monitor is described and the results of the initial round of liquid xenon purity measurements done at SLAC are reported. Also discussed is some of the chemical physics behind liquid xenon purity, and how it is applicable to large-scale, liquid xenon time projection chamber (TPC) operation.

Chapters 4 and 5 tackle the question of positive ion transport in liquid xenon. In

Chapter 4, the results of an experiment aimed at measuring the positive ion mobility in liquid xenon is described. This experiment was done to help establish the feasibility of isolating individual barium ions from macroscopic samples of liquid xenon for laser tagging.

Chapter 4's discussion leads into Chapter 5's, in which some models of ion transport in liquid xenon are considered, and an attempt is made at reconciling them with the available data. In particular, the mystery of why the mobilities of positive ions in liquid xenon are so small [2] is addressed. This is done by performing first-principles computer calculations of ion transport and comparing the results with experimental data and with some of the theoretical models from the literature.

In Chapter 6 the subject of ion manipulation is taken up. Ion manipulation is a requirement of a liquid xenon double beta decay detection system in which the ion must be isolated from the liquid phase in order to measure its spectroscopy. Accordingly, part of the present work involved the investigation of methods of capturing the ion in the liquid phase, transporting it to the gas phase and releasing it. While it is also possible in principle to attempt the single ion spectroscopy in the liquid phase, such a scheme has its own complications and is beyond the scope of this thesis. It is, however, in the process of being investigated by other workers [3].

The final chapter is devoted to the conclusions that can be drawn from the present work from the standpoint of liquid xenon double beta decay detector design. A proposal is given for how the techniques investigated in the present studies might or might not be incorporated into a large scale detection system.

1.1 References

- [1] V.I. Tretyak and Y.G. Zdesenko, *At. Data Nucl. Data Tables* **80**, 83 (2002).
- [2] A.J. Walters and L.W. Mitchell, *J. Phys. D: Appl. Phys.*, **36**, 11, 1323 (2003).
- [3] W.M. Fairbank, EXO Collaboration E-Log Document #327.

Chapter 2

Physics Motivation

double beta decay is the second-order nuclear process by which an atomic nucleus spontaneously changes its atomic number by two units. This phenomenon has two versions. The first, which we call two-neutrino double beta decay ($2\nu\beta\beta$), occurs with the emission of two β -ray electrons (or positrons) together with two electron-type antineutrinos (or neutrinos). This effect, although highly suppressed, has been unambiguously shown to occur in several otherwise stable elements [1]. In the second version of double beta decay, which we call neutrinoless double beta decay ($0\nu\beta\beta$), only two β -rays emerge from the disintegrating nucleus, unaccompanied by any neutrinos or antineutrinos. Such decays are expressly forbidden under the Standard Model (SM) of particle physics, and their observation would be a clear indication of some entirely new kind of physics at work. As such, $0\nu\beta\beta$ has not yet been convincingly demonstrated experimentally, although there exists one highly controversial claim [2]. Since the most popular extensions to the SM that allow for $0\nu\beta\beta$ is a modification of the theory of neutrinos, we will spend the bulk of the discussion in this chapter on neutrino physics.

2.1 Neutrino Basics

More than 50 years after their discovery, neutrinos remain among the most poorly understood of the known elementary particles. They were originally proposed as a

means to explain the apparent non-conservation of energy and angular momentum in β^- decays [4]. Up to this time, β^- decays were thought to be two body decays, which made it very difficult to understand how the associated β -ray energy spectrum was continuous, let alone how a nucleus with integer spin could remain so after emitting a spin-1/2 electron. As a way out, W. Pauli postulated, in an open letter to the participants of a 1930 Tübingen physics conference, a new, light, electrically neutral particle with spin 1/2 to carry away the missing energy and angular momentum [3]. Pauli published his proposal in 1933 and soon afterward E. Fermi picked up the idea, publishing his famous four-fermion quantum field theory of beta decay in 1934. Along the way, Fermi (and F. Perrin, who was working independently) deduced from the shape of the β^- decay spectrum that the neutrino was hundreds of thousands of times lighter than the electron, if not entirely massless altogether. Accordingly, the SM, whose complete formulation was to follow, has always assumed the neutrino mass to be identically zero.

Fermi's theory originally consisted only of a vector interaction between the fields that described the neutron (which Chadwick had discovered in 1932), the proton, the electron and the neutrino. This model adequately reproduced the decay rate of free neutrons as well as that of many different neutron-rich nuclei [5]. Using Fermi's result, H. Bethe and C. Peierls famously predicted that neutrinos with energies relevant to beta decay could propagate through 500 astronomical units of water largely unaffected [4], implying that direct neutrino detection was virtually impossible. Some 20 years later in 1953, F. Reines and C. Cowan realized that a powerful nuclear reactor could produce a high enough neutrino flux that a large detector in close proximity might observe a small but significant signal [6]. They reported a successful detection shortly thereafter [7], and Reines won the Nobel prize in 1995 for this work.

In 1959, R. Davis announced a result that implied that the neutrino accompanying a β^- decay (which was detected by Reines and Cowan's apparatus) is distinct from that produced in a β^+ (positron) decay. We now interpret the former neutrino as being the antiparticle of the latter, and we refer to these two particles, respectively, as "electron antineutrino" ($\bar{\nu}_e$) and "electron neutrino" (ν_e).

In the meantime, M. Goldhaber and coworkers had set out to measure the neutrino

helicity by studying the electron capture decay of metastable $^{152}\text{Eu}^m$. The ^{152}Eu decays to an excited state of ^{152}Sm which, in turn, decays by a prompt gamma emission. By varying the helicity of the captured electron and measuring the helicity of the ^{152}Sm gamma ray (which, by angular momentum conservation, must be the same as that of the previously emitted neutrino), the team concluded that to within experimental sensitivity the electron neutrino helicity is always negative [8].

In 1962 a second neutrino generation, consisting of the muon neutrino and antineutrino, was reported by Danby and co-workers [9], who observed that the new particles were created in pion decays of the form $\pi^+(\pi^-) \rightarrow \mu^+(\mu^-) + \nu_\mu(\bar{\nu}_\mu)$ in an accelerator experiment at Brookhaven. This was followed some years later by the discovery in 1975 by M. Perl of the τ lepton implying the existence of a third neutrino flavor associated with the τ [10].

It was not until 1991 that studies at CERN provided compelling evidence that there existed exactly three neutrino flavors that contribute to the Z decay width[11]. The τ neutrino was finally detected directly in 2001 in an accelerator experiment at Fermilab [12] called DONUT. This experiment observed the neutrinos created in the reaction $D_s^+(D_s^-) \rightarrow \tau^+(\tau^-) + \nu_\tau(\bar{\nu}_\tau)$ as well as those arising from the subsequent decays of the τ^\pm .

The past few years have seen something of a revolution in neutrino physics. Several recent experiments [13],[14],[15],[16] have produced results that demonstrate convincing evidence for cross-generational neutrino mixing, similar to the behavior of quarks in the SM. The mixing behavior of quarks is understood in the SM to be due to the fact that the quark states of definite flavor (such as up, charm or top) are actually superpositions of states of definite mass. This means that the interactions of one flavor experience interference from the interactions of the others, producing the mixing effects that we observe. In contrast to quark mixing, which is studied by measuring the time evolution or decay of the hadron containing the particular quark flavor of interest, the neutrino experiments look for neutrinos prepared in states of definite flavor to transform themselves into other flavors as they propagate from source to detector. This phenomenon is called neutrino oscillation. If its underlying mechanism is basically the same as that of quark mixing, we must conclude that all neutrinos

cannot all have vanishing masses.

Incidentally, the idea that neutrinos, by virtue of being massive, could exhibit mixing behavior analogous to the quarks is not a new one [17]. It was first proposed by B. Pontecorvo in 1957, who was interested in whether an analog to kaon oscillation existed for neutrinos. Gradually, as more and more experimental data began to accumulate, Pontecorvo was finally able in 1967 to formulate a two-component neutrino mixing scheme that closely resembled the modern theory. Within the same decade, a team of three Japanese physicists, Z. Maki, M. Nakagawa, and S. Sakata independently did the same thing [18], and the modern theory of neutrino mixing now bears the initials PNMS to honor its four architects. We will return to PNMS and their ideas presently.

To sum up, accounting for massive neutrinos requires a substantive revision of the SM. Furthermore, while it is clear that at least some of the neutrino masses are nonzero, it is still true that they must be very small indeed, as current experimental limits from kinematic measurements are in the few eV range [19]. Understanding the relative smallness of these masses will become important as we make attempts to construct theories that describe them. This will be addressed in the following section.

2.2 Neutrino Mass & Mixing

To reiterate, the recently observed phenomenon of neutrino oscillations implies that, as with quarks, the neutrino states of definite flavor are superpositions of states of definite mass. If we assume that there exist only three such mass eigenstates (there could in principle be more¹), they will be related to the flavor eigenstates by a 3×3

¹The majority of the current experimental data [13],[14],[15],[16] are compatible with the picture presented here. However, the result from the LSND experiment[20] demonstrates a set of neutrino oscillation parameters that are very different from those that are determined by the other four experiments. One way of resolving the conflict is to postulate the existence of a 4th mass eigenstate (from which it follows that there must also exist a 4th neutrino flavor that has no weak interactions [21]), but such a scenario has so far been ruled out by a simultaneous analysis of the data from LSND and other experiments to a confidence level of nearly 3σ [21]. The LSND result is now in the process of being checked by a similar experiment called miniBoone that has slightly different systematics[26]. For simplicity, we ignore the complications posed by the LSND result in the present discussion.

unitary mixing matrix U [27] (also known as the PNMS matrix after Pontecorvo, Nakagawa, Maki, and Sakata, whose ideas led to its construction). This relationship is expressed as

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{ee} & U_{e\mu} & U_{e\tau} \\ U_{\mu e} & U_{\mu\mu} & U_{\mu\tau} \\ U_{\tau e} & U_{\tau\mu} & U_{\tau\tau} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \quad (2.1)$$

where the subscripts e, μ , and τ denote the flavor eigenstates and the numerals 1, 2 and 3 label the mass eigenstates. Although the SM does not predict the magnitudes or phases of the elements in the mixing matrix U , we can deduce its main features by recalling some of the key ingredients of the SM Lagrangian.

In the SM, the fermion fields are constructed as superpositions of four-component Dirac spinors whose basis is the set of four plane-wave spinors accounting for the particle states, the antiparticle states, and their respective two possible spin orientations. Any such field can also be expressed as a superposition of two so-called chiral fields which are, for massless fermions, its two projections onto states whose spin axes, respectively, are aligned and anti-aligned with the particle momentum. (For massive fermions each chiral field consists mostly of one helicity with a small admixture of the opposite.) Making this notational change is useful because the weak interactions, which violate parity, can distinguish between the chiral fields and can therefore be written as a purely vector interaction connecting fields of the same chirality. With this in mind the SM Lagrangian density for a free fermion field is constructed as

$$\mathcal{L} = \bar{f}(i\gamma^\mu\partial_\mu - m)f \quad (2.2)$$

where f and its adjoint \bar{f} represent the properly constructed fermion field of mass m . $\{\gamma^\mu\}$ is simply the familiar Dirac gamma matrix of index μ , and all other γ that appear subsequently shall be interpreted as such unless otherwise indicated. Note that, by construction, f creates an anti-fermion and annihilates a fermion while

\bar{f} does the opposite. Because of the orthogonality of the basis spinors, the only nonzero contributions to the mass term are those that connect (anti-)fermion states to (anti-)fermion states. For the charged fermions (e.g. e , μ , and τ) the anti-fermions and fermions carry the opposite electrical charge, so the SM Lagrangian mass term implicitly conserves charge.

In the chiral basis the mass Lagrangian becomes

$$\mathcal{L}_{mass} = (\bar{f}_R + \bar{f}_L)m(f_R + f_L) \quad (2.3)$$

where the following definitions apply

$$f_R \equiv (1 + \gamma^5)f \quad (2.4)$$

$$f_L \equiv (1 - \gamma^5)f \quad (2.5)$$

$$\overline{f_{R,L}} \equiv \bar{f}(1 \pm \gamma^5). \quad (2.6)$$

It is trivial to show, using the definition 2.6, that only the terms $m\overline{f_L f_R}$ and $m\overline{f_R f_L}$ survive in 2.3. This result makes it plain that the SM fermion mass terms connect only chirally left-handed fields to chirally right-handed ones. This can be understood as an expression of the fact that a massive, interacting fermion necessarily moves slower than the speed of light, which means it can be overtaken by the particles with which it interacts. The same particle can therefore interact as a predominantly left-handed or predominantly right-handed entity depending on whether it is overtaking or being overtaken. The equation of motion that follows from the Lagrangian density 2.2 is the familiar Dirac equation and we refer to the mass term 2.2 as the Dirac mass term.

Because of the requirement of charge conservation, the SM fermion mass Lagrangian is the unique description of the charged fermion masses. Neutral fermions, i.e. neutrinos, are not subject to this constraint, so one can envision a neutrino mass term that connects fields with their charge conjugates. Such terms are physically consistent

as long as the fields involved do not carry any conserved quantum numbers whose conservation would be violated in these terms. Simply put, when constructing these mass terms for neutrinos, we cannot require that the neutrino field carry a definite lepton number, which, like electrical charge, is explicitly conserved in the SM.

A neutrino field that does not carry lepton number is a self-conjugate field, i.e. its charge conjugate and itself are identical up to a multiplicative phase factor. Such a field was first proposed by E. Majorana in 1937 [28]; we now refer to it as the Majorana neutrino. The Majorana neutrino field is expressed by making the definition

$$\nu_M \equiv (\nu + \nu^C) \quad (2.7)$$

where the neutrino field ν is constructed in an identical manner to the f fields, and ν^C is its charge conjugate. In the chiral basis, we can write the mass Lagrangian

$$\mathcal{L}_{mass} = (\overline{\nu_R} + \overline{\nu_L} + \overline{\nu_R^C} + \overline{\nu_L^C})m(\nu_R + \nu_L + \nu_R^C + \nu_L^C) \quad (2.8)$$

where we have used the convention $\nu_{(R,L)}^C \equiv (\nu_{(L,R)})^C$. It is straightforward to show, using the orthogonality of the constituent spinors and the algebraic properties of the chiral fields, that the only terms that survive from 2.8 are

$$\mathcal{L}_{mass} = m\overline{\nu_L}\nu_R + m\overline{\nu_R}\nu_L + m\overline{\nu_R^C}\nu_R + m\overline{\nu_L^C}\nu_L. \quad (2.9)$$

Note that the first 2 terms on the right hand side look like Dirac mass terms while the latter 2 are the mass terms of Majorana fermions. As it turns out, it is even more general to allow the Dirac and the Majorana terms to take on different values for the mass m . With this in mind, we can use the identity $\overline{\nu_R}\nu_L \equiv \overline{\nu_L^C}\nu_R^C$ and re-write 2.9 in matrix form as

$$\mathcal{L}_{\text{mass}} = \begin{pmatrix} \overline{\nu_R^C} & \overline{\nu_L} \end{pmatrix} \begin{pmatrix} 0 & M \\ M & m \end{pmatrix} \begin{pmatrix} \nu_L^C \\ \nu_R \end{pmatrix} + (\text{hermitian conjugate}). \quad (2.10)$$

The diagonalization of this mass Lagrangian gives the 2 mass eigenvalues

$$m_{\pm} = \frac{M \pm \sqrt{M^2 + 4m^2}}{2} \quad (2.11)$$

with eigenvectors

$$\left\{ \begin{pmatrix} 0 \\ \chi_R \end{pmatrix}, \begin{pmatrix} \chi_L \\ 0 \end{pmatrix} \right\} \quad (2.12)$$

where

$$\begin{cases} \chi_L \equiv (\nu + \nu_L^C) \\ \chi_R \equiv (\nu_R + \nu_R^C). \end{cases} \quad (2.13)$$

In order to retain the physical interpretation of the elements of the diagonalized mass matrix (which requires that particle masses be positive-definite) we use an equivalent representation of the mass matrix

$$M' = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} m_- & 0 \\ 0 & m_+ \end{pmatrix} \quad (2.14)$$

so that the -1 that multiplies m_- guarantees a positive mass eigenvalue for the $(\chi_L, 0)$ eigenfield. Now, if we take the limit $M \gg m$, then, to first order in $(m/M)^2$, the two mass eigenvalues associated with $(\chi_L, 0)$ and $(0, \chi_R)$ are, respectively m^2/M and M . We have thus an “explanation” for the observed behavior of neutrinos, which goes as follows. As was shown by Goldhaber *et al.*, only chirally left-handed neutrino

fields seem to participate in the weak interactions. Furthermore, we know from the state-of-the-art kinematic measurements [19] that the electron antineutrino rest mass is *at most* 2.3 eV. So, it is easy to see that the Majorana formulation hints that the extremely light, left-handed neutrinos from our everyday experience might actually be a manifestation of the χ_L field, while χ_R indicates some chirally right-handed object whose mass is too large to have been produced thus far in particle accelerator experiments (whose state-of-the-art center of mass energies, incidentally, are now in the neighborhood of 2 TeV). Numerologists like to assign values to m and M such that m is of order the top quark mass (~ 176 GeV), and M , the GUT scale ($\sim 10^{16}$ GeV), giving a left-handed neutrino mass of ~ 3.0 meV, entirely consistent with current experimental bounds. This scheme for suppressing the mass scale of the left handed component of the neutrino field is referred to as the “See-Saw Mechanism” [3]. Because the χ fields are self-conjugate, their phases, if present, do not cancel from their Lorentz-invariant bilinears. Put another way, a mass term such as $m^2/M\bar{\chi}_L\chi_L$ can carry a phase $e^{i\theta}$ that cannot be removed by redefining the phases of the constituent ν spinors. Such phases, if nonzero, manifest themselves experimentally as CP -violating effects. This is a simple consequence of the fact that the CP operation (parity followed by charge conjugation) transforms an expression to its hermitian conjugate.

Since the χ_R field cannot have any easily observable effects at energies small compared to its mass, we can safely ignore it in calculations involving neutrino physics at experimentally accessible energies. The remainder of the discussion, therefore, will be limited to left-handed Majorana neutrinos, with the treatment generalized to handle three lepton generations.

We now return to the PNMS matrix U . It is straightforward to show that an $n \times n$ unitary complex matrix like U is fully characterized by $n(n-1)/2$ independent amplitudes and $n(n+1)/2$ phases. This means that U can be constructed out of three mixing angles and three complex phases.² With this in mind, we give the following

²Note that the analogous matrix that describes the quark sector has only 1 phase. This is because the quarks are all Dirac particles (not self-conjugate), which means their physics is unchanged under the transformation $f \rightarrow e^{i\alpha}f$. Thus the phases of the quark fields can be redefined to absorb all but 1 of the phases in the V (CKM) matrix.

representation of U made popular by the Particle Data Group[29]:

$$U = \begin{pmatrix} c_{12}c_{13}e^{-i\alpha_1/2} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta+\alpha_2/2} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \quad (2.15)$$

where $c_{ij} \equiv \cos\theta_{ij}$ and $s_{ij} \equiv \sin\theta_{ij}$. The phase δ is analogous to the CP violating phase in the quark sector CKM matrix and can be measured by searching for CP violating effects in neutrino mixing. The other phases $\alpha_{\{i,j\}}$ are the irreducible phases from the Majorana bilinears and can only be measured by studying the Majorana character of neutrinos.

For any given neutrino flavor, we define an “effective” mass given by the

$$\langle m_\alpha \rangle = \sum_i U_{\alpha i}^2 m_i, \quad (2.16)$$

$$i \in \{1, 2, 3\} \quad (2.17)$$

$$\alpha \in \{e, \mu, \tau\}. \quad (2.18)$$

This is a kind of weighted average over the definite-mass states whose linear combination is the neutrino with flavor α .

At this time little is known about the elements of U , whose values the SM does not predict; these are among the couple-dozen or so of its so-called “free parameters”. What information we do have is from fits to the results of neutrino oscillation experiments [29], whose results (which so far have only been sensitive to the magnitudes and not the phases) indicate that the magnitudes of all elements are of order unity, save for U_{13} which is consistent with zero.³

In addition to being sensitive to the mixing angles as parametrized by matrix 2.15, the neutrino oscillation experiments measure differences between the masses of neutrino mass eigenstates. In these experiments, the mass differences are manifested as a beat

³Incidentally, this structure is in striking contrast to that of the CKM matrix, whose off-diagonal magnitudes are close to zero while the diagonal elements are close to one.

frequency arising from by the superposition of the wave functions of neutrinos having different masses. The available data seem to suggest that the three mass eigenstates are separated by two unequal gaps, which for historical reasons are called Δm_{solar} and $\Delta m_{\text{atmospheric}}$. It is known, from solar neutrino oscillation data, that of the two mass eigenstates contributing to Δm_{solar} , the state corresponding to the largest admixture of ν_e is the state of lower mass. This has to do with the fact that on their way to Earth, the solar neutrinos must propagate large distances through solar matter, making the propagation of ν_e s, which interact differently with matter than $\nu_{\mu,\tau}$, distinguishable from the others. However, it is not known whether the third state lies at higher or lower mass than the other two; only the *difference* between its mass and the next closest has been measured. The neutrino mass spectra in which this third mass eigenstate is either at the bottom or at the top are referred to respectively as the “normal” and “inverted” hierarchies (Figure 2.1). Moreover, at this time there is no data to suggest whether the lightest of the three states has a mass of zero, or if the three states are collectively offset from the origin by an amount comparable to the size of the splittings. The latest research [21] has found $\Delta m_{\text{solar}}^2 = 8.2 \pm 0.3 \times 10^{-5} \text{ eV}^2$ and $\Delta m_{\text{atmospheric}}^2 = 2.2_{-0.4}^{+0.6} \times 10^{-3} \text{ eV}^2$, implying that the heaviest of the neutrino mass eigenstates must have a mass of *at least* $\sim 30 \text{ meV}$. This means future attempts to measure the effective neutrino mass aiming for the tens-of-meV range are likely to make a positive detection *if* neutrino masses actually have an inverted hierarchy.

This idea is illustrated more precisely in Figure 2.2, where the effective mass of ν_e is plotted as a function of the mass of the lightest mass eigenstate. These plots are based on calculations given in Ref [22], which use the experimental bounds for the neutrino mixing parameters and use the published upper limit as the value of $\sin\theta_{13}$. For each case, the calculation assumes that each of the unknown phases lie somewhere in the range $[0, 2\pi]$. Clearly, the mass hierarchy has a dramatic influence over the experimental accessibility of the neutrino physics.

In the following sections we discuss how the study of neutrinoless double beta decay can provide meV-level neutrino mass sensitivity, while establishing whether the neutrino mass Lagrangian is Dirac-like, Majorana-like, or both. These issues represent some of the most pressing open questions in this field.




Figure 2.1: Two possible neutrino mass spectra that fit the available data (not to scale), reflecting the fact that the sign of the smaller mass splitting is known while the sign of the larger is not. The spectrum on the left is commonly referred to as the “normal hierarchy” and the one on the right is “inverted”. The three red arrows to the right of the mass spectra indicate experimental limits on the neutrino masses determined from studies of the cosmic microwave background [23], kinematic measurements of the β^- decay of tritium [19], and neutrino observations of the 1987 supernova [24]. Figure adapted from Ref. [25].

2.3 Neutrinoless Double Beta Decay

In neutrinoless double beta decay, an initial state containing only an atomic nucleus results in a final state with a nucleus plus two electrons. It therefore violates lepton number conservation by 2 units and is a clear signature of the involvement of massive Majorana fermions.

Taking the simplest case, which assumes no right-handed currents and “maximal mixing” of ν and ν^C (i.e. ν and ν^C are superimposed with equal weights to form χ_L), it becomes obvious that the decay amplitude $i\lambda$ for the diagram 2.3 varies as

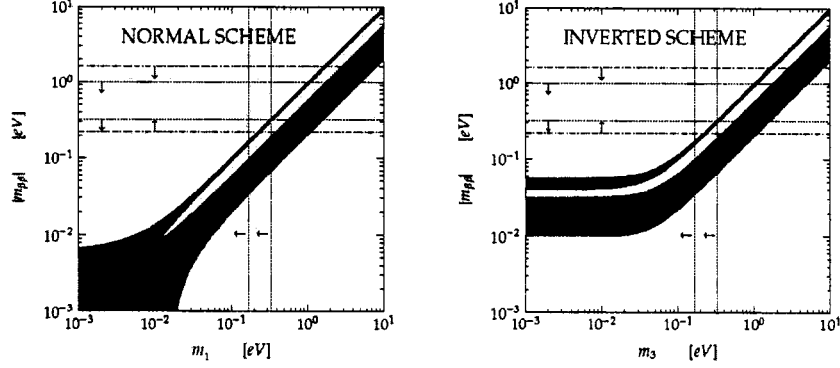


Figure 2.2: (From Ref. [22]) ν_e effective mass ($m_{\beta\beta}$) as a function of the mass of the lightest mass eigenstate for the case of normal and inverted mass hierarchies. Note that the two scenarios become indistinguishable when the lightest mass becomes comparable in size to the mass splittings. Under these conditions, the mass hierarchy is said to be *degenerate*. The horizontal and vertical lines show different experimental limits arising from $0\nu\beta\beta$ measurements (horizontal lines) and cosmological observations (vertical lines). The white bands within the dark ones are the regions in which the third, non-Majorana phase in the PNMS matrix is nonzero. This is where CP is violated independently of the Majorana character of the neutrino.




Figure 2.3: Underlying diagram for neutrinoless double beta decay.

$$i\lambda \sim \bar{e}_L \gamma^\mu \nu_L^C \bar{e}_L \gamma^\nu \nu_L \quad (2.19)$$

This is just the product of two vector interactions between the chirally left-handed electron field (e_L) and the neutrino, and it is easy to see that it must contain the factor

$$\frac{i(m_\nu + \gamma^\mu p_\mu)}{p^2 - m_\nu^2} \quad (2.20)$$

where m_ν is the neutrino mass and p is its momentum. This factor follows simply from the Feynman calculus and represents the contribution from the internal neutrino propagator. The term $\gamma^\mu p_\mu$ vanishes because of the anti-commutation properties of γ^5 with γ^μ (recall that the definition of the chiral fields contains an implicit factor of γ^5).

Momentum conservation requires that the p in the denominator be of order the momentum on the external electron lines. Since we know that m_ν must be very small, we can drop it from the denominator, and now we see that the neutrinoless double beta decay amplitude is proportional to m_ν . This means that the decay rate for this process is essentially a nearly direct measurement of the effective neutrino mass.

In practice, the graph 2.3 is only a small part of a much more complicated scheme. The quarks on the external lines are actually bound up in nucleons, which in turn are constituents of an atomic nucleus. This means that any real-life analysis of such a decay will be severely complicated by the rather messy nuclear physics involved. We take a look at some of this physics in the following section.

2.4 Nuclear Physics of Double Beta Decay

In order for double beta decay to occur for a given nucleus, the first-order process, i.e. single beta decay, must be somehow forbidden, whether by the nuclear energetics or by some selection rule. At the same time, the initial and final nuclear states must form a non-vanishing transition probability when connected by the Hamiltonian describing the decay. The rate for neutrinoless double beta decay is therefore proportional to the matrix element describing the transition probability as well as to a phase space factor to account for the kinematics of the decay. We calculate the phase space multiplier by evaluating the density of final states that can be occupied by the emerging particles which are the two emitted electrons and the daughter nucleus. If we are interested in the transition probability to a given final nuclear state, then the relevant density of states is simply that of the two beta electrons, that is, the convolution of the momentum 3-sphere with the Fermi function.

Once these various ingredients have been calculated, they are brought together to give the half-life

$$T_{1/2}^{0\nu}(0^+ \rightarrow 0^+) = [G^{0\nu}(Q, Z)|M^{0\nu}|^2\langle m_\nu \rangle^2]^{-1} \quad (2.21)$$

where Q is the maximum electron kinematic energy (the kinematic endpoint), Z is the nuclear charge, $M^{0\nu}$ is the relevant nuclear matrix element, and $\langle m_\nu \rangle$ is the neutrino mass expectation value. For simplicity, we only consider the decay between even parity, spinless states, as indicated by the expression in the parentheses on the left hand side.

In the case of 3 neutrino generations with mixing, one must include the U_{ei} matrix at each vertex in the Feynman graph 2.3. This means that the quantity $\langle m_\nu \rangle^2$, which is given by equation 2.18, is subject to cancellations that occur when the squared elements of U that are opposite in sign are added together in the sum. From the definition 2.18, it is clear that the Majorana phases, if present, are what is responsible for such a cancellation. These phases are presently unknown.

A further subtlety is the quantity $|M^{0\nu}|^2$, the nuclear matrix element. Because its

calculation requires detailed knowledge of the initial, final and intermediate nuclear states, it is highly dependent upon the choice of nuclear model. The two most commonly used models are the nuclear shell model (NSM) and the quasiparticle random phase approximation (QRPA).

The NSM assumes a structure in which the innermost nucleons form a closed, inert “shell” that acts as the source of a central potential in which the remaining nucleons move, giving rise to the nuclear physics. The nucleon wavefunctions are evaluated and used to compute the matrix elements. This is done by evaluating overlap integrals of the wavefunctions corresponding to the two nuclear states for which the matrix element is to be evaluated.

By contrast, the QRPA uses a phenomenological nucleon-nucleon interaction that is based on experimentally determined nucleon pairing energies and scattering resonances. With some approximations, the nuclear many-body problem is then solved and used to generate nuclear wavefunctions corresponding to the required quantum numbers. For a long time this procedure tended to overestimate the magnitudes of the matrix elements until it was realized that there was additional binding due to a strong proton-proton spin-isospin polarization interaction [3].

Both the NSM and the QRPA methods have steadily improved over the years in their ability to produce results that are more or less consistent with the experiments; however it is unfortunate that between the 2 models, differences by as much as a factor of 10 can occur. Such uncertainties complicate the conversion between the experimental limit on the half-life $T_{1/2}^{0\nu}$ and an upper limit on m_ν .

2.5 Two-Neutrino Double Beta Decay

Any nucleus for which the neutrinoless double beta decay is allowed must also suffer the competing decay in which two antineutrinos accompany the emission of the 2 electrons. The $2\nu\beta\beta$ process proceeds irrespective of the intrinsic neutrino properties and is fully allowed in the SM. Its decay amplitude is calculated as a second-order analogue to single beta decay and its total rate is given by

$$T_{1/2}^{2\nu}(0^+ \rightarrow 0^+) = G^{2\nu}(Q, Z)|M^{2\nu}|^2 \quad (2.22)$$

where, again, for simplicity we consider only the decay between even parity, spinless states. Here, Q is the endpoint energy, Z is the nuclear charge, and $|M^{2\nu}|$ is the relevant nuclear matrix element. The total decay rate was obtained by integrating over all electron and antineutrino momenta. To study the energy spectrum of the emitted electrons we must go back one step to the expression

$$\frac{d\Gamma}{dT} \sim F(Z)^2(T+1)^2(Q-T)^6[(Q-T)^2 + 8(Q-T) + 28]|M^{2\nu}|^2 \quad (2.23)$$

where $F(Z)$ is the Coulomb function for the final nuclear state with charge Z , and T is the kinetic energy of 1 of the electrons. Since a typical double beta decay experiment will measure only the sum of the kinetic energies of the two electrons, we must rewrite equation 2.23 in terms of $K \equiv T_1 + T_2$, where T_1 and T_2 are the respective kinetic energies of the 2 electrons. The result is

$$\frac{d\Gamma}{dK} \sim F(Z)^2 K(Q-K)^5[K^4 + 10K^3 + 40K^2 + 60K + 30]|M^{2\nu}|^2 \quad (2.24)$$

from which it is apparent that the sum-energy (K) spectrum for the electrons is continuous. By contrast, the $0\nu\beta\beta$ spectrum is simply a delta function situated at the kinematic endpoint Q . A double beta decay detector measures the superposition of the two spectra, so one must search for $0\nu\beta\beta$ by looking for a sharp spike, broadened by the detector resolution, at the endpoint of the $2\nu\beta\beta$ sum-energy spectrum. Typical (model-dependent) theoretical values for $T_{1/2}^{2\nu}$ and $T_{1/2}^{0\nu}$ are, respectively, 3×10^{19} yr and 3×10^{24} yr.

The discussion presented here is by no means exhaustive. Other forms of double beta decay also exist, such as modes in which 2 positrons are emitted, or ones which involve an electron capture. In addition, in full analogy with single beta decay, there

are inverse decays that are also possible. It is generally accepted that these decay modes are suppressed for various kinematic reasons [4]; we therefore ignore them here. The interested reader is referred to Refs. [1, 3, 4] for further information.

2.6 Other possible mechanisms for $0\nu\beta\beta$

For completeness, we emphasize here that there are other scenarios that lead to $0\nu\beta\beta$ that involve other mechanisms than the one described in the previous sections. Examples include processes in which the decay is mediated by the exchange of supersymmetric particles, or by W bosons coupling to a charged Higgs, and so on [35, 36]. Additionally, it is possible that the emerging electrons may be accompanied by the emission of the Goldstone boson, called the majoron, that is associated with the breaking of a putative space-time symmetry associated with the conservation of lepton number [3, 4]. Such an effect would give a continuous electron sum-energy spectrum with a shape different from that associated with $2\nu\beta\beta$.

It has been shown [35], however, that *any* double beta decay process that does not emit neutrinos along with the two electrons must, at some level, contain a Majorana-like neutrino propagator. It can therefore be concluded that if $0\nu\beta\beta$ is observed it would amount to unambiguous evidence that neutrinos possess a Majorana-like mass term. Even so, additional information on how, if at all, the alternative mechanisms contribute must be obtained before Equation 2.21 can be used to transform a measured $0\nu\beta\beta$ half-life into a ν_e effective mass.

2.7 Detecting Double Beta Decay

Because the half-lives are so long, detecting double beta decay requires a substantial amount of the decaying isotope together with a long integration time. In addition, as with any search for an extremely rare event, background considerations are an important driver of the experimental design. Thus, while many nuclei possess the

appropriate structure for double beta decay, only a handful are considered practical from an experimental standpoint. Two of the more important considerations that go into designing a double beta decay experiment are (1) large Q value, and (2) high natural abundance. It is also highly desirable but not essential that the isotope be usable in such a way that the decaying sample also serves as the detection medium. Backgrounds must be controlled by operating the detectors deep underground to reduce their susceptibility to cosmogenic radioactivity, and the materials in the experimental setup must be ultra-clean. Furthermore, the utmost care must be taken not to introduce radioactive contaminants into any part of the system that is in close proximity to the detector. Using an isotope with a high Q value (≥ 2) MeV is also necessary because, as we saw in the previous section, the phase-space integral in the decay rate scales rapidly with Q . In addition a higher Q allows better separation of the $0\nu\beta\beta$ from radioactive backgrounds. This is because the backgrounds tend to consist of the gamma ray lines associated with the members of the uranium decay chain most of whose strongest lines are at energies less than about 2.0 MeV.

Roughly speaking, one expects the maximum detectable half-life to scale as $T \propto Nt$ where N is the number of double beta decay candidate nuclei in the sample and t is running time of the experiment. When radioactive contaminants are present within the sample at some fraction r , we expect to observe roughly $(\ln 2)rNt\Delta E$ background events where r is the contamination ratio, and ΔE is the detector energy resolution. These background events are subject to a fluctuation $\sqrt{rNt\Delta E}$ which corrupts the signal, compromising our ability to measure T . When this is taken into account, the half-life sensitivity scales as $T \propto Nt/\sqrt{rNt\Delta E} = \sqrt{Nt/r\Delta E}$. In other words, the detector sensitivity scales directly with energy resolving power and inversely with radioactive contamination [39]. Figure 2.5 shows the result of an exact calculation that illustrates this. In this case, only the background due to $2\nu\beta\beta$ was considered, but this background scales the same as any constant-fraction contaminant and the same principles apply. Of course, the entire above discussion assumes that the background gamma spectrum does not contain any sharp lines that coincide exactly with the $0\nu\beta\beta$ endpoint; if it does, then the energy resolution of the detector offers no advantage and other background rejection techniques are needed.




Figure 2.4: Summed electron energy spectrum for $2\nu\beta\beta$ and $0\nu\beta\beta$.

To date, the best success has been realized with ^{76}Ge , which is now the state-of-the-art, having set a world-record $T_{1/2}^{0\nu}$ limit of 1.2×10^{25} yr (90% CL) [30]. To achieve this result, the Ge was isotopically enriched and fashioned into large ultra-radiopure calorimeters, with energy resolution approaching 10 keV in the neighborhood of 2 MeV. The combination of high energy resolution together with the use of pulse-shape discrimination [41] makes these detectors especially powerful at distinguishing the signal from the backgrounds.

As it turns out, xenon's intrinsic properties are not quite as favorable. The energy resolution obtained in xenon gas-filled detectors is of order 5 times worse than germanium's. This is because, for as-yet undetermined physical reasons, the ionization process in xenon suffers from a larger Fano factor [37]. The Fano factor of a given detection medium is a measure of the extent to which the fluctuations in the ionization process are described by Poisson statistics [38]. In the case of germanium, the fluctuation is much smaller than what it would be for a purely Poisson process. This is because essentially all of the energy from an ionizing event is converted into free electrical charge and there is a tight constraint on the amount of charge an event with a given energy can liberate. In xenon the fluctuations are larger, presumably because the ionization process is competing with other, unseen, random energy loss




Figure 2.5: The plots on the left show double beta decay electron summed energy spectra for the case of perfect energy resolution (red curve), 2 % FWHM energy resolution (black curve) and 10 % FWHM (blue curve). The abscissa in this case is given in units of the electron endpoint energy Q . The gray shaded region is the FWHM signal window for the case of 10 % energy resolution. Events within this window are counted as $0\nu\beta\beta$ signal. As shown at right, the amount of contamination by $2\nu\beta\beta$ events in the $0\nu\beta\beta$ signal window depends strongly on the detector energy resolution [40]. For these calculations, the $0\nu\beta\beta$ rate was assumed to be 100,000 times smaller than that of the $2\nu\beta\beta$, and backgrounds to $0\nu\beta\beta$ from sources other than $2\nu\beta\beta$ were not considered.

mechanisms. These effects are even more pronounced in liquid Xe, where the resolution relative to germanium is worse by more than a factor of 20.

The current generation of double beta decay experiments have Nt that are measured in kg-years of exposure. Realizing significant gains in sensitivity will require that the next generation experiments contain many tons of material that can be observed over several years. In addition, for most candidate isotopes, the natural abundance is of order 10%, which means that the next generation of experiments will probably require some form of isotopic enrichment, techniques for which exist only for a subset of the available candidate isotopes.

Though it is relatively inexpensive to isotopically enrich xenon, a large scale xenon experiment would almost necessarily be done in the liquid phase for compactness' sake. This means that the poor energy resolution in LXe must be somehow overcome in order for a ton-scale xenon experiment to be competitive.

Fortunately, xenon appears to afford the possibility of rejecting the backgrounds on an event-by-event basis by allowing coincident detection of the decay daughter. This idea was first proposed by M.K. Moe in 1991 [40]. If this is done, virtually all radioactive backgrounds can be rejected, leaving only the $2\nu\beta\beta$ spectrum whose rejection is achieved with only modest energy resolution.

2.8 EXO: a next-generation double beta decay experiment

EXO will ultimately consist of a several-ton sample of xenon, enriched to 80% in the double beta decaying isotope ^{136}Xe , whose natural abundance is 8.9% and whose Q -value is also quite favorable at 2.48 MeV. In addition, EXO aims ultimately to exploit the possibility of laser tagging the double beta decay daughter as a way of rejecting the backgrounds. EXO is now in its R& D phase with the bulk of the effort currently devoted toward developing a prototype liquid Xe time projection chamber (TPC), without laser tagging. The EXO prototype is henceforth referred to as EXO-200. It

consists of a 40 cm long right cylindrical chamber holding about 200kg of liquid Xe. The inner volume of the cylinder is split symmetrically into two sub-volumes, each bounded on one end by a set of crossed anode grid wires, together with a planar array of ultraviolet-sensitive avalanche photodiodes (APDs), and a shared cathode at the other end (Figure 2.6). double beta decay candidate events are identified by the ionization and scintillation they generate in the liquid xenon. The 175 nm scintillation light is picked up by the APDs, and a high electric field is used to drift the ionization electrons over to the anode grid. These electrons generate signals in the grid wires which are individually read out to give the x and y (i.e., lateral) coordinates of the event. The z coordinate is determined by the time delay between the arrival of the




Figure 2.6: Liquid xenon TPC concept. Figure taken from Ref. [31].

ionization charge at the anode wires and the scintillation. Fig. 2.7 shows the complete xenon TPC principle schematically. The three-dimensional reconstruction offered by the TPC design is an important tool for fighting backgrounds, particularly in a system that does not incorporate laser tagging. This is because the spatial information of each event can be used to distinguish multiple gamma ray scatters as well as to perform accurate fiducial volume cuts.

The energy of double beta decay candidate events is given by the relative size of the ionization signal recorded on the anode grid wires. Measuring these signals with ultra-low noise charge-sensitive amplifiers gives an estimated energy resolution at the kinematic endpoint of roughly 3% [37]. In principle, one can do better by using the scintillation light generated in the event to perform an energy correction. This is due to an event-by-event anti-correlation of scintillation with ionization arising from certain details of the atomic-level physics of the ionization process in liquid xenon. This effect was first demonstrated by the EXO collaboration in Ref. [37]. By using this technique, EXO-200 may achieve 2% energy resolution or better. The ultimate sensitivity of the EXO prototype will ultimately be limited by the radioactive backgrounds. It is therefore planned to include laser tagging in all future EXO experiments.

2.9 Laser tagging

To reiterate, EXO will search for the reactions

$$^{136}\text{Xe} \rightarrow ^{136}\text{Ba}^{++} + 2e^- + 2\bar{\nu}_e \quad (2\nu\beta\beta) \quad (2.25)$$

$$^{136}\text{Xe} \rightarrow ^{136}\text{Ba}^{++} + 2e^- \quad (0\nu\beta\beta) \quad (2.26)$$

for which only limits on the half lives have been reported in the literature to date [42, 43]. The most stringent lower bounds on $T_{1/2}^{0\nu}$ and $T_{1/2}^{2\nu}$ are given in Ref. [43] and are, respectively, $T_{1/2}^{2\nu} > 1.1 \times 10^{22}$ yr and $T_{1/2}^{0\nu} > 7.0 \times 10^{23}$ yr (both reported at 90% CL).

For each decay mode, the two emerging electrons are to be detected by the ionization and scintillation light they produce in the LXe sample. Candidate events can in principle be distinguished from fake events due to radioactive contamination by coincident detection of the Ba^{++} daughter.

Performing such a single-ion detection is possible because of the unique atomic properties of the Ba ion and because this ion will probably be quite accessible in the LXe




Figure 2.7: EXO-200 signal scheme. The ionization electrons are drifted and collected on crossed wires, resulting in signals that give the energy and the x - and y - position of the event. Simulated examples of what these signals might look like are shown as Channel 168 and Channel 119. The z - position of the event is found using the time between the detection of scintillation signal by the avalanche photodiodes in the APD plane and the arrival of the ionization signals on the wires. This is a composite of figures found in Refs. [32, 33, 34].

environment. It is generally assumed [44] that when a Ba^{++} ion forms in liquid Xe, it neutralizes to form Ba^+ which is metastable against further neutralization because of the energetics. Ba^+ , in turn, has a unique and very well studied single-ion optical spectroscopy [45]. The EXO approach therefore is to detect the LXe ionization and scintillation followed by the isolation and spectroscopic identification of the resulting Ba^+ ion.

The essentially three-level structure of Ba^+ (Figure 2.8) is due to the unpaired s valence electron which makes Ba^+ quite analogous to the hydrogen atom. The strongly allowed dipole transitions between the s , p and d states (the latter of which is metastable with a lifetime of order 60 s) are accessible using fairly standard, commercially available 493 nm and 650 nm solid-state laser technology, and a saturated ion can radiate up to $\sim 10^7$ photons per second. Such a signal is easily detectable using conventional optoelectronic devices.




Figure 2.8: Simplified level scheme for singly ionized barium.

While possible in principle to perform the single ion spectroscopy in an LXe system *in situ*, the detailed spectroscopic physics of Ba^+ surrounded by a liquid is essentially unknown at this time. There has, however, been much work reported in the literature on Ba^+ detection in various background gases [46, 47]. EXO's scheme, therefore, will be to attempt to isolate the Ba^+ ion from the liquid so that the spectroscopy can be done in a low- pressure gas environment. It was therefore imperative that some

research on ion transport in LXe along with single ion manipulation in LXe/Xe be included in EXO's R&D program.

2.10 Conclusion

By now, it should be clear that the study of double beta decay in xenon has created a need for understanding some of the physics of electron drift in liquid xenon as well as the dynamics of ion transport and individual ion manipulation. The R&D that has been performed for EXO in these areas is addressed in the following chapters.

2.11 References

- [1] V.I. Tretyak and Y.G. Zdesenko, *At. Data Nucl. Data Tables* **80**, 83 (2002).
- [2] This claim was first presented in H.V. Klapdor-Kleingrothaus *et al.*, *Mod. Phys. Lett. A* **16**, 2409 (2001). See, however, Y.G Zdesenko *et al.*, *Phys. Lett. B* **546**, 206 (2002) for a lucid discussion of how this claim is most likely erroneous.
- [3] See, for example, F. Boehm and P. Vogel, *Physics of Massive Neutrinos*, Cambridge University Press (1987) for a good review of this subject.
- [4] This section draws heavily from the excellent textbook by C.W. Kim and A. Pevsner, *Neutrinos in Physics and Astrophysics*, Harwood Academic Publishers (1993).
- [5] P.E. Hodgson, E. Gadioli and E. Gadioli-Erba, *Introductory Nuclear Physics*, Oxford University Press (2000).
- [6] F. Cowan and C. Reines, *Phys. Rev.* **92**, 462 (1953).
- [7] F. Cowan *et al.*, *Phys. Rev.* **117**, 159 (1959).
- [8] *Phys. Rev.* **109**, 1015 (1958).

- [9] G. Danby, *et al.*, Phys. Rev. Lett. **9**, 36 (1962).
- [10] M.L. Perl, *et al.*, Phys. Rev. Lett. **35**, 1489 (1975).
- [11] Burkhardt and Steinberger, Ann. Rev. Nucl. Part. Sci. **41**, 55 (1991).
- [12] K. Kodama, *et al.*, Physics Letters B **504**, 218 (2001).
- [13] S.N. Ahmed *et al.*, Phys. Rev. Lett. **92**, 181301 (2004).
- [14] Y. Ashie, *et al.*, Phys. Rev. Lett. **93**, 101801 (2004).
- [15] T. Araki *et al.*, Phys. Rev. Lett. **94**, 081801 (2005).
- [16] E. Aliu, *et al.*, Phys. Rev. Lett. **94**, 081802 (2005).
- [17] S.M. Bilenky, Physica Scripta **121**, 17 (2005).
- [18] Z. Maki, *et al.*, Prog. Theor. Phys. **28**, 870 (1962).
- [19] Ch. Kraus, *et al.*, arXiv:hep-ex/0412056 (2005).
- [20] Aguilar, A. *et al.* (LSND), Phys. Rev. D **64**, 112007 (2001).
- [21] M.C. Gonzalez-Garcia, arXiv:hep-ph/0410030 (2004).
- [22] C. Giunti, Nucl. Phys. B (Proc. Suppl.), **145** 231236 (2005).
- [23] D.N. Spergel, *et al.*, Ap. J. (Suppl. Ser.), **148**, 175 (2003).
- [24] K.S. Hirata *et al.*, Phys. Rev. D **38**, 448 (1988).
- [25] R.N. Mohapatra *et al.*, arXiv:hep-ph/0412099 (2004).
- [26] E. Church, *et al.*, FERMILAB-P-0898 (1998).
- [27] The discussion in this section draws heavily from the review article by B. Kayser, in G. Altrarelli and K. Winter (eds.), *Neutrino Mass*, Springer Tracts in Modern Physics (2005).

- [28] E. Majorana, *Nuovo Cimento* **14**, 171 (1937) (In Italian).
- [29] S. Eidelman *et al.*, *Phys. Lett. B* **592**, 1 (2004).
- [30] Y.G Zdesenko *et al.*, *Phys. Lett. B* **546**, 206 (2002).
- [31] J. Wodin, EXO Collaboration E-Log Document #560.
- [32] C. Hall, EXO Collaboration E-Log Document #76.
- [33] R. Neilson and C. Hall, EXO Collaboration E-log Document #109.
- [34] C. Hall, EXO Collaboration E-log Document #221.
- [35] J. Schechter and J. W. F. Valle, *Phys. Rev. D* **25**, 289 (1982).
- [36] H. V. Klapdor-Kliengrothaus, *First Int. Symp. on Lepton and Baryon Number Violation (Proceedings)*, 251, Trento (1998).
- [37] E. Conti *et al.*, *Phys. Rev. B* **68**, 054201 (2003).
- [38] W.R. Leo, *Techniques of Nuclear and Particle Physics Experiments: A How-To Approach* (2nd Ed.), Springer-Verlag (1994).
- [39] M. Danilov *et al.*, *Phys. Lett. B* **480**, 12 (2000).
- [40] M. Moe, *Phys. Rev. C* **44**, R931 (1991).
- [41] L. Baudis *et al.*, *Phys. Rev. Lett.* **83**, 41 (1999).
- [42] R. Luescher *et al.*, *Nucl. Phys. B (Proc. Suppl.)* **66**, 195 (1998).
- [43] R. Bernabei *et al.*, *International Workshop on Technique and Application of Xenon Detectors*, Kashiwa, Japan (2001).
- [44] A.J. Walters and L.W. Mitchell, *J. Phys. D: Appl. Phys.*, **36**, 11, 1323 (2003).
- [45] W. Neuhauser, *et al.*, *Phys. Rev. Lett.* **41**, 233 (1978).
- [46] S. Waldman, PhD thesis, Stanford University (2005), and references therein.

- [47] W. Neuhauser, *et al.*, Phys. Rev. Lett. **41**, 233 (1978).

Chapter 3

Electron Drift in Liquid Xenon

In the previous chapter, we motivated the need for an understanding of the charge transport processes in liquid xenon (LXe) as part of a liquid xenon double beta decay detector engineering program. In this chapter we describe some studies of electron transport in LXe that were conducted at SLAC. The specific aim of this work was threefold. First, it was necessary to establish the role that electronegative impurities have in limiting the transport of free electrons. Second, it was needed to qualify purification schemes for the removal of the electronegative impurities from a LXe detection medium. Finally, it was used to qualify detector candidate materials based upon the extent to which they introduced electronegative contaminants into LXe.

3.1 Basic Concepts

It is very well documented that the electron mobility in LXe (as well as in many other liquefied noble gases) is similar to that in semiconductors, approaching $3 \times 10^2 \text{ cm}^2 \text{ V}^{-1}\text{s}^{-1}$ at large applied electric fields [1]. The detailed physics of the electron transport process is at least somewhat understood and has been calculated in the literature using various free-carrier conduction models based on an assumed band structure for LXe [1, 2]. The mobilities for ions of either sign are typically 10^5 to 10^6 times smaller [3], and the physics of why ion conduction is so much slower is an area of current research [4, 5, 6, 7]. This issue is taken up in greater detail Chapters 4 and

5.

From the standpoint of detector physics, the enormous mismatch between the transport parameters of ions and electrons is of concern because it impacts how efficiently the charge from ionizing events in the liquid can be collected. This is because electronegative impurities, which are always present to some degree in a given LXe sample, are known to attach to free electrons to form negative ions [8]. Once this happens the negative charge that is now associated with these ions propagates through the liquid about one million times more slowly and is therefore effectively lost.

This phenomenon can be used to characterize the relative purity of a sample of LXe, or indeed any liquefied noble gas, by injecting free electrons into the sample and drifting them through it on timescales short compared to the drift time of the negative ions. The purity is then determined by measuring the relative amount of charge that makes it across on these timescales. A dedicated system for performing exactly this type of experiment, called the Xenon Purity Monitor (XPM) has been developed and installed at SLAC. Its design is based on a similar apparatus used by the ICARUS collaboration [9] and it has been instrumental in drawing important conclusions on both practical considerations and the physical chemistry of liquid xenon purity.

3.2 The Xenon Purity Monitor: Theory of Operation

The XPM consists of a 6.5 cm ID cylindrical copper vessel with a stainless steel top flange that has a Xe gas inlet and the electrical feedthroughs. The outlet consists of a braised-on piece of 3/8" stainless tubing. The entire unit is immersed in ~40L of HFE-7000 [10] heat transfer fluid, and the XPM copper wall is maintained at ~163 K. The HFE-7000 is cooled by a PolyCold® 550HC [11] series industrial cryogenic refrigerator. The PolyCold® unit is capable of producing up to 1 kW of cooling power, although in practice only a small fraction of this actually cools and condenses the Xe. This is primarily because the HFE-7000 is inefficient as a heat transfer fluid at low temperatures. As Xe gas is introduced into the upper part of the XPM it

condenses on the copper walls, filling the XPM and part of the exit tube with LXe. The interior of the XPM, whose electrostatics will be discussed presently, contains the elements that are necessary for drifting small amounts of electric charge along its length. There is also a space at the bottom for placing samples that are to be tested. The XPM operates on the principle that a charge Q injected into a sample of LXe will persist for a characteristic time τ before becoming attached to an impurity molecule such as O_2 . Expressed mathematically, the time evolution of Q is

$$\frac{dQ}{dt} = -knQ \quad (3.1)$$

where $\frac{dQ}{dt}$ is the temporal rate of change of charge Q , n is the number concentration of electronegative impurities and k is the (effective) attachment rate coefficient. Here we have assumed that n is a constant in time, however, as we shall see shortly, it will in general vary slowly over long periods. If the charge is injected in the form of a finite pulse whose centroid moves at a constant drift velocity v_d , we can make the change of variables $t \rightarrow x/v_d$, where x is the drift distance. This gives

$$v_d \frac{dQ}{dx} = -knQ. \quad (3.2)$$

The solutions to these equivalent equations are given below, with definitions for the parameters λ and τ which are, respectively the characteristic attenuation distance and attachment time for injected charge.

$$Q(t) = Q_0 e^{-t/\tau} \quad (3.3)$$

$$Q(x) = Q_0 e^{-x/\lambda} \quad (3.4)$$

$$\tau \equiv \frac{1}{nk} \quad (3.5)$$

$$\lambda \equiv \frac{v_d}{nk}. \quad (3.6)$$

Thus, λ and τ are related to the concentration of impurities, assuming the effective

mean attachment coefficient for these impurities is known. In a typical XPM experiment the *electron lifetime* τ is measured and an “oxygen equivalent” impurity concentration is inferred, using the attachment coefficient for oxygen, as determined by Bakale *et al.* [8], to invert relation 3.6. In reality the k of our system is some effective average weighted by the relative concentration of

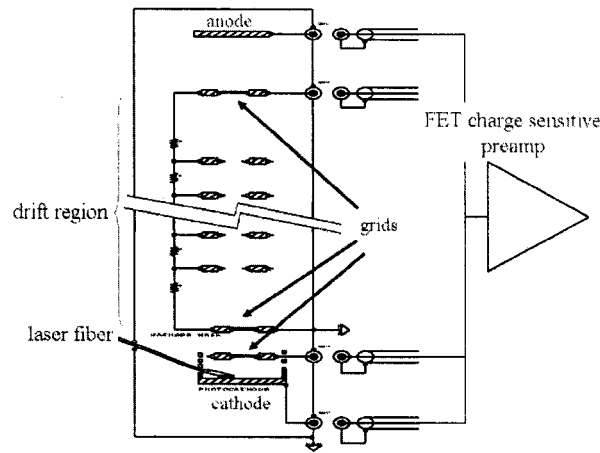


Figure 3.1: The Xenon Purity Monitor (XPM), a device for measuring electron lifetime in liquid Xe. The three grids, going from top to bottom, are the anode grid, the cathode grid and the window grid. The anode, anode grid, window grid, and cathode are each connected via a coaxial cable to a biasing network and a preamplifier for reading out the charge signal.

each electronegative species.

Figure 3.1 shows a diagram of the SLAC XPM. Electrons are created at the photocathode by means of a 266nm laser pulse which is guided by a multi-mode quartz optical fiber. Located 0.6cm above the cathode are two grids, separated by 0.5cm. These are, respectively, the window and cathode grids. The photocathode voltage is provided by a DC high voltage power supply and a resistive divider is used to set the window grid to a potential halfway between that of the photocathode and ground. The cathode grid is grounded and connected via a 300 M Ω resistor chain to another

grid, 10.9cm away, that is 2.5mm below the anode. The cathode grid and the anode grid define the electron drift region for the experiment. Figure 3.2 shows a circuit schematic of the XPM.

The anode, window, and cathode grids are all AC coupled, via 4.7 nF blocking capacitors, to a low-noise charge amplifier whose rise time is about $3\mu\text{s}$ and whose decay time is roughly $400\mu\text{s}$. As the electrons drift between the window grid and the cathode, the induction signals they create cancel, so that no signal is produced by the preamp until the electrons enter the region between the window grid and the cathode grid. This is desirable because it causes a small time delay between the laser flash lamp pulse (which produces a great deal of electromagnetic interference) and the beginnings of a visible charge signal.

The signal from the preamp rises quickly to some value, where it remains while the electrons drift past the cathode grid and into the drift region. They then travel the $\sim 11\text{ cm}$ to the anode and create an induction signal at the anode once they drift past the anode screening grid. This signal is recorded as a pulse with a falling edge. Fig. 3.3 shows an example of the oscilloscope trace corresponding to the preamp output. To work out the electron lifetime τ , the pulse-height at the rising edge of the XPM signal is interpreted as the amount of charge Q_c *entering* the drift region. This is then compared with the charge Q_a *exiting* the drift region, represented by the XPM signal falling edge. Taking time $t = 0$ to be the time at which Q_c passes the cathode grid, we use the solution 3.6 to equation 3.1 to write

$$Q_c = Q_0 \quad (3.7)$$

$$Q_a = Q_0 e^{-t_a/\tau} \quad (3.8)$$

where t_a is the time that Q_a arrives at the anode. We then solve for τ to give

$$\tau = \frac{t_a}{-\log|\frac{Q_a}{Q_c}|} \quad (3.9)$$

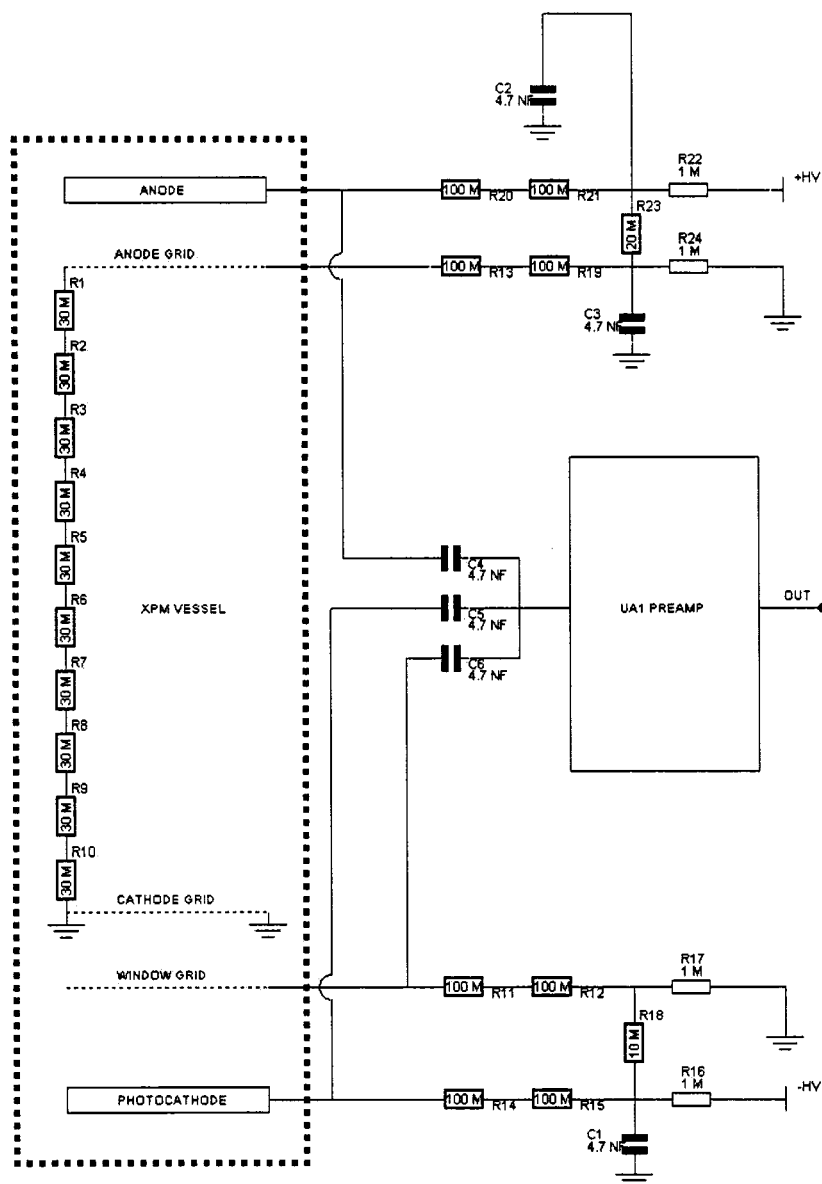


Figure 3.2: Circuit diagram of the XPM.




Figure 3.3: Example of the XPM signal.

implying

$$\lambda = \frac{D}{-\log \left| \frac{Q_a}{Q_c} \right|} \quad (3.10)$$

where D is the length of the drift region, about 10.9 cm for the SLAC XPM.

The XPM's main function is to characterize the contamination rate of different candidate materials for the construction of the EXO detector. It also serves the purpose of providing an understanding of the purification rates associated with the purification system. When the XPM is in operation, many successive measurements of τ are made over the course of several hours to several days. On these timescales the impurity concentration n is given by

$$\frac{dn}{dt} = \frac{dN_i}{dt} - \frac{n}{\tau_r} - \frac{n}{\tau_f} \quad (3.11)$$

where $\frac{dN_i}{dt}$ is the rate at which impurity species enter the LXe and τ_r is the time constant that characterizes the rate of their removal by purification. Note that there may also be additional, accidental mechanisms for contaminant removal, such as freezeout, whose effect is represented generically by the term $-n/\tau_f$. We have no control over these processes, but their effect can be characterized by studying how the LXe purity behaves when the artificial recirculation and purification are turned off. If we suppose the impurities originate from some reservoir that supplies a constant fraction of its material per unit time, then the first term on the right hand side of 3.11 must be given by

$$\frac{dN_i}{dt} = \frac{N_i}{\tau_i} \quad (3.12)$$

where τ_i is the time constant of contamination from the reservoir. The solution to 3.12 is

$$N_i(t) = N_0 e^{-t/\tau_i} \quad (3.13)$$

where N_0 is the initial concentration of contaminant molecules available in the reservoir. Plugging 3.13 into 3.11 gives the solution

$$n(t) = e^{-t(\frac{1}{\tau_r} + \frac{1}{\tau_f})} \left[\frac{N_0}{1 - \tau_i(\frac{1}{\tau_r} + \frac{1}{\tau_f})} + n_0 \right] + \frac{N_0}{\tau_i(\frac{1}{\tau_r} + \frac{1}{\tau_f}) - 1} e^{-\frac{t}{\tau_i}} \quad (3.14)$$

where n_0 is the starting impurity concentration in LXe [15]. It is clear from 3.14 that when $\tau_i \gg \tau_{r,f}$, the impurities in the LXe build up, approaching N_0 , and in the opposite limit they equilibrate at close to zero concentration. The essential feature of expression (3.14) is that it is the superposition of two decaying exponentials, one with a time constant characterizing the introduction of impurities; the other, their removal. We thus make the definitions

$$R \equiv \frac{N_0}{1 - \tau_i(\frac{1}{\tau_r} \tau_f)} + n_0 \quad (3.15)$$

$$I \equiv \frac{N_0}{\tau_i(\frac{1}{\tau_r} + \frac{1}{\tau_f}) - 1} \quad (3.16)$$

$$\tau_d \equiv \frac{\tau_r \tau_f}{\tau_r + \tau_f} \quad (3.17)$$

and re-write Equation 3.14 as

$$n(t) = R e^{-\frac{t}{\tau_d}} + I e^{-\frac{t}{\tau_i}} \quad (3.18)$$

which will be simpler to use in analyzing purity data. Finally, we use 3.6 to re-write 3.18 in terms of the measured electron lifetime,

$$\tau(t) = \frac{k N_0}{R e^{-\frac{t}{\tau_d}} + I e^{-\frac{t}{\tau_i}}} \quad (3.19)$$

which we will use to directly interpret the results of some of the XPM tests.

3.3 The purifier

In the SLAC system, the Xe is purified entirely in the gas phase by means of a Saes® MonoTorr Phase I 3000 getter (fig. 3.4). It consists of an electropolished canister containing a finely divided zirconium alloy that becomes activated when heated to ~ 400 °C. As the Xe gas is passed through the getter material, its electronegative impurities react with the hot zirconium and remain behind. The unit in current use is the model PF3-C3-R-1 and is rated to extract, down to the ≤ 1 ppb level, a long list of substances, including O₂, CO, CO₂, H₂ and even N₂. This level of purification, as specified, is achieved in a single pass, provided the gas flow rate is ≤ 5 standard liters per minute, and the pressure drop is less than 1 bar. Optimum

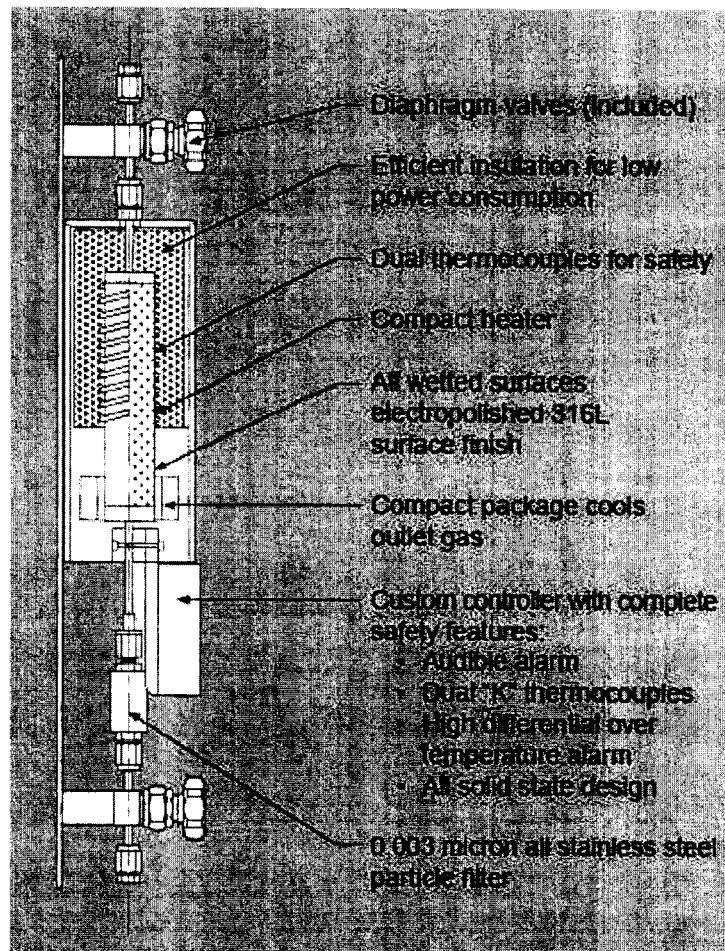


Figure 3.4: The SAES getter. Figure taken from Ref [12].

performance is realized only in one flow direction.

Though the purifier's single-pass performance is satisfactory, multiple passes are often necessary, particularly when there are sources of impurities downstream. To address this issue, a continuous recirculation scheme was implemented. The recirculation system consists of a custom-built mechanical bellows pump [16] which forces gaseous Xe out through the XPM exit tube, through the purifier, and back to the XPM through its input tube. This process is facilitated by a 30W heater placed against the XPM exit tube to help the Xe boil. The Xe flow rate can be varied by adjusting the rotational frequency of the electric motor that actuates the bellows, but there are some operational constraints that prevent Xe from being recirculated much faster than ~ 2 SLPM. This system can be operated essentially indefinitely while the XPM is read out continuously, monitoring gradual changes in purity over time.¹

3.4 Experimental Procedure

To perform a purity test, the sample(s) to be studied are placed in the bottom of the XPM which is then sealed and pumped to about 1.0×10^{-7} mbar. The XPM and its accompanying vacuum system is bakeable to 200 °C. For some of the tests the entire system was baked out before performing the measurements. Figure 3.5 shows a simplified schematic of the gas handling system. Pumping is performed by two Varian turbomolecular pumps (only one is shown in Figure 3.5), and the vacuum is monitored by two Varian T-NUDE-F ion gauges and an SRS 300 AMU residual gas analyzer. The XPM is then immersed in a large, sealed dewar containing about 40 L of HFE-7000 and cooled with the PolyCold® refrigeration unit. The temperature of the XPM copper walls are monitored by three type T thermocouples. In addition, the Xe gas pressure is read out using a set of capacitance manometers. The XPM temperature is maintained at a setpoint of 163 ± 1 K by means of an on/off control algorithm (see Appendix) implemented in LabView™. With this algorithm, the PC

¹An earlier version of the setup used a liquid nitrogen reservoir for cooling and had neither Xe recirculation capabilities, nor automated fitting and data capture. Some of the data that are presented in the following section were taken on this apparatus.

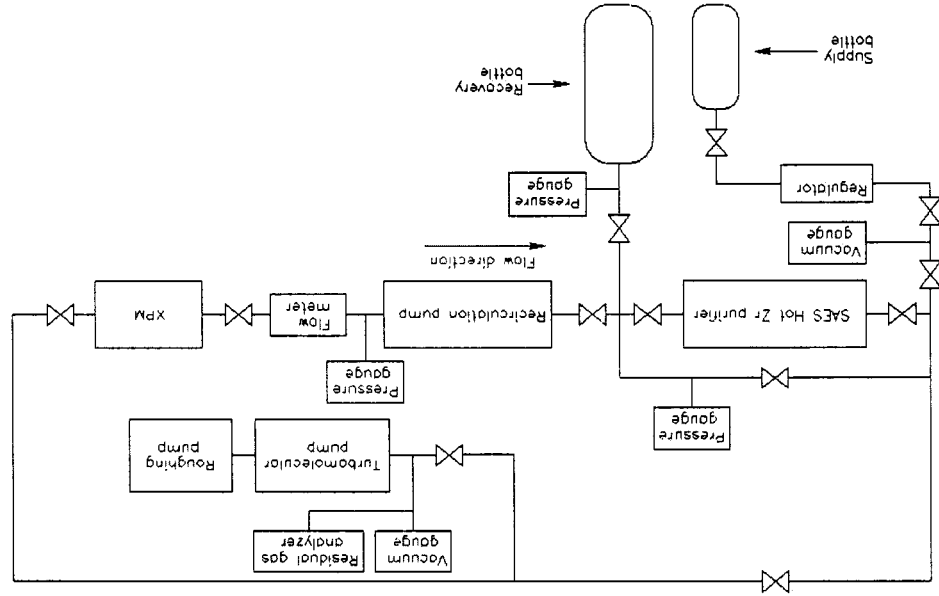


Figure 3.5: Simplified schematic of the gas handling system.

can modulate the flow of the PolyCold® refrigerant by changing the state of a solenoid valve.

When the XPM wall temperature reaches the setpoint, Xe gas is introduced from either the supply bottle or the recovery bottle, passing through the SAES® getter. LXe then begins to condense in the XPM. When the LXe level, which is monitored by a capacitive level sensor inside the XPM, reaches the top flange, purity measurements can begin. Although some initial tests were done by varying the drift field while measuring the purity, the bulk of the measurements were done at a drift field of about 200 V/cm. This field is achieved by biasing the anode at 3 kV. The photocathode in turn must always be biased to ensure the transparency of the window and cathode grids, i.e.

$$E_w < E_c < E_d < E_a \quad (3.20)$$

where E_w is the electric field between the cathode and the window grid, E_c is that between the window grid and the cathode grid, E_d is the drift field, and E_a is the field at the anode. More precisely, it has been shown in Ref [18] that the transparency condition for electrons crossing a given grid to go from, say, region 1 to region 2, is given theoretically by

$$E_2 \geq \frac{2\pi r_g}{d} E_1 \quad (3.21)$$

where r_g is the radius of the wire in the grid, and d is the pitch. In practice, for a real system with non-ideal geometry the inequality 3.21 must be over-satisfied in order to achieve full transparency. For typical XPM runs, the photocathode is biased at a voltage of -50 V, while the anode is biased at 3000 V. The grids in the XPM consist of 50 μm wire with a 1.5 mm pitch, and their transparency is ensured by the resistor network external to the XPM (see Appendix), provided the photocathode and anode voltages maintain the ratio $\geq 60:1$. A computer model of the XPM using the MAXWELL® software package whose qualitative results are shown in fig. 3.6

was also constructed and used to help estimate the electric fields.

Electrical charge is injected into the XPM using a Continuum® MiniLite pulsed UV laser, which delivers about 3 mJ of 266 nm radiation per pulse. This deep ultraviolet (UV) laser light is actually generated by quadrupling the 1064 nm Nd:Yag primary emission. Photoelectrons are generated when the UV laser pulse strikes the gold photocathode inside the XPM.

Custom data-acquisition software is used to download the signal waveforms and to analyze them in real time by performing a fit to the following function

$$Q(t) = \frac{Q_c}{2} \operatorname{erfc}\left(\frac{t_c - t}{\sigma_c}\right) e^{-\frac{t-t_c}{\tau_h}} - \frac{Q_a}{2} \operatorname{erfc}\left(\frac{t_a - t}{\sigma_a}\right) e^{-\frac{t-t_a}{\tau_h}} + Q_{\text{off}} \quad (3.22)$$

where $Q_{(c,a)}$ are defined as before, $t_{(c,a)}$ are the respective arrival times of the photoelectrons at the cathode grid and anode grids, τ_h is the charge amplifier hold time, $\sigma_{(c,a)}$ are the charge transit times in the second and fourth gaps in the XPM, and Q_{off} is the overall offset on the baseline. Note that for this waveform electrical charge and voltage amplitude have the same units. In this system the charge amplifier gain is assumed to be unity. If SI units for electrical charge are desired, simply divide the voltages by the amplifier gain, which is about 20 mV per pC and varies somewhat from device to device.

In typical XPM runs, the electron drift time ($t_a - t_c$) is of order 70 μs , which means a single measurement is instantaneous on the hours-to-days timescales relevant to changes in LXe purity. In practice, the XPM readout electronics are intensely susceptible to low-frequency noise, microphonic pickup being a particular problem. To get around this, data taking must be synchronized to the firing of the laser Q-switch (whose repetition rate has a maximum setting of 10 Hz) and many (about 100) pulses must be averaged for optimum signal-to-noise. Since the laser pulse rate under normal operation is 1 Hz, this means that there must elapse at least ~ 100 s between data points. Unfortunately, there is an additional complication arising from the electromagnetic interference generated by the laser flash lamp circuit. The preamp picks this up as a wide ($\sim 200\mu\text{s}$), 5mV pulse that severely distorts the waveform baseline.

This is dealt with by periodically recording an oscilloscope trace consisting of just the baseline, storing it, and subtracting it from the signal trace. The baseline trace is obtained by closing the shutter at the laser output port. More details of the baseline-subtraction procedure and the fitting algorithm are discussed in the Appendix. In




Figure 3.6: Equipotentials in the XPM [17], as calculated by the MAXWELL® software.

a typical run, τ is measured every 5 min for several days. After the first several data points are taken, the recirculation pump is started and periodically stopped and re-started in order to study how the LXe purity evolves.

To date, the following materials have been tested using this procedure: Lexan® (polycarbonate), PTFE Teflon®, MasterBond® EP29LPSP epoxy, Kapton® (polyimide) ribbon cables, Advanced Photonix® large area avalanche photodiodes, and *p*-terphenyl wavelength shifter. All of these materials were at one point or another considered for use as part of the EXO prototype in direct contact with the LXe. Other minor tests were also performed to determine the effect of exposing the XPM to sources of gaseous contaminants. These included an air leak, outgassing hydrogen, and both baked and unbaked vacuum plumbing.

3.5 Results

As an initial test of the system, the electron drift time as a function of electric field was measured, the drift time being taken as the difference between the two waveform fit parameters (equation 3.22) t_c and t_a . The drift time determines the electron drift speed, which is plotted as a function of electric field in fig. 3.7. These measurements




Figure 3.7: Electron drift speed as a function of electric field in liquid xenon.

do not seem to agree precisely with those of other workers [1, 13], but it has been documented that the presence of trace impurities can either enhance or reduce the mobility of free electrons [14], depending on the identity of the contaminant. With this in mind it seems plausible that differences between the various results might be entirely due to varying Xe purity conditions among the three experiments. Of course, this assertion is difficult to verify, given that neither of Refs. [1] or [13] report the relative purity of their samples. Nonetheless, the general trend as well as the overall order-of-magnitude seen in our result matches fairly well. This provides some qualitative assurance that the XPM is working as it should, and that our understanding of its construction is correct.

The first few materials and contamination tests were performed using the earlier version of the system. Since these tests were rather basic they are easily summarized by simply reporting the final electron lifetime that was measured for each. This is shown in Table 3.1. For each material, the reported “preparation” indicates whether, and if so how, the system was baked for the test. The “test period” is the length of time between the first and the final electron lifetime (τ) measurements. From these results it is apparent that *p*-terphenyl is the only material that is categorically incompatible with electron drift in LXe. In fact, its effects were found to be so severe that the cathode signal was only a few millivolts at maximum electric field and laser power. The electron lifetime did not recover immediately after the sample was removed. Instead 10 manual Xe fills through the purifier were performed, resulting in a modest increase in the size of the cathode signal. This implies that in addition to the XPM, the *p*-terphenyl had also contaminated the Xe supply during Xe recovery. The anode signal was not recovered until after the XPM copper cell had been replaced and the plumbing downstream of the purifier had been baked at 200 °C, flushed with ethanol, and baked again. Following this, the electron lifetime as measured immediately after filling through the purifier was > 1 ms. This suggests that the purifier was able to remove the remaining contamination in a single pass.

Unbaked Lexan®, however, shows only borderline performance, which means that applications involving this material might require a bakeout. None of the other materials appear incompatible with > 1 ms electron lifetimes; this corresponds to an impurity concentration of 0.5 ppb oxygen equivalent.

The following results are from the recirculation-equipped system. In Figure 3.8 we show the behavior of the LXe purity after leaking in 10 cm³ of room air. To obtain these data, the XPM was first filled and the electron lifetime was measured repeatedly over a day, giving a final result of 2.5 ± 0.2 ms. Then an upstream valve was opened to a 10 cm³ reservoir of room air, and we see the electron lifetime plummet over the course of several minutes. Several hours later Xe gas recirculation is commenced and the electron lifetime is seen to improve slowly at first and then more rapidly over the next few days. For this dataset, the electron lifetime was recorded once per 5 minutes. On this timescale, the individual measurements are basically independent,

Material	Form	Preparation	Initial τ (ms)	Test period	Final τ (ms)
Lexan®	190 cm ² area disks	baked@150 °C	0.52±0.03	6.5 d	0.79±0.04
hydrogen	outgassing from purifier	baked@150 °C	0.64±0.05	24 h	0.64±0.05
p-terphenyl	~ 3 μ g on Lexan® disk	baked@150 °C	no anode signal – test aborted		
air	outgassing from plumbing at 10 ⁻⁷ mbar	unbaked	0.8±0.1	10 h	2.0±0.5
ethanol	ethanol rinse; pump to 10 ⁻⁷ mbar	unbaked	1.8±0.3	10 h	2.0±0.4
Teflon®	190 cm ² area disks	unbaked	1.8±0.4	10 h	1.9±0.5
Lexan®	190 cm ² area disks	unbaked	0.03±0.001	11 d	0.3±0.01
EP29LSP epoxy	65 cm ² area disk	baked@50 °C	0.3±0.01	10 h	1.5±0.5

Table 3.1: Results of initial round of purity tests using non-recirculation capable system. Error bounds are statistical only.

which means that the scatter in the datapoints represents the statistical error. As the purity increases, it becomes clear that the error on the measurement does so also, exceeding 10% at electron lifetimes greater than ~ 2.0 ms.

The lifetime also suffers from a large systematic error stemming from the sensitivity of the waveform fit to the subtraction of the baseline. Essentially, what happens is that the specific baseline distortions drift gradually, resulting in a shift in the fitted signal pulseheights. The effect of this is also apparent in Figure 3.8, manifested as a discontinuity showing up precisely where the baseline is re-measured. This problem was remedied in later measurements by performing a dynamic baseline subtraction (see Appendix for details).

Figure 3.9 shows a test with 190 cm² surface area PFE Teflon® disks in the XPM. At the start of the experiment, the residual air outgassing from the Teflon was sufficient to give an electron lifetime close to zero. Recirculation improves the situation over the next few days. When recirculation is stopped, however, we see the purity begin to fall again.

A subset of the recirculation data were fit to Equation 3.19 in Figure 3.10, and we see that the model appears to describe the main qualitative features fairly well under conditions in which the purification process dominates. In Figure 3.10 the fit to the purification cycle gives $\tau_d = 4.7 \pm 0.2$ h and τ_i immeasurably large, with a reasonable χ^2 . However, the fit to the contamination cycle is very poor; it actually appears as though the data exhibit a superposition of at least two decay constants. This implies multiple sources of impurities that diffuse into the liquid at different rates. This effect is currently under study. With smaller errors it might even be possible to perform these fits as a way of characterizing the outgassing rates associated with different test samples.

The final test we describe is that of a sample consisting of a large area avalanche photodiode (LAAPD), some polyimide ribbon cable and some additional epoxy. These results are shown in Figure 3.11. Again, here we see that with recirculation, satisfactory purity is readily and repeatably achieved.




Figure 3.8: The effect of contamination by room air on the electron lifetime. The arrows indicate the following: (1) introduction of 10 cm³ of room air; (2) recirculation initiation; (3) re-measurement and subtraction of the baseline; (4) electron lifetime approaching the sensitivity limit.




Figure 3.9: Electron lifetime vs. time with Teflon® in the XPM. The green arrows indicate times when recirculation was started and red arrows indicate when it was stopped.




Figure 3.10: Left: Fit of Equation 3.19 to electron lifetime trend while recirculation is turned on. Right: Fit to data taken with recirculation off. Error bars are statistical only.




Figure 3.11: Electron Lifetime vs time for LXe sample containing an LAAPD, ribbon cables and epoxy. Green arrows indicate the start of recirculation; red, the interruption.

3.6 Discussion

The XPM data seem to indicate that vast majority of the materials tested are viable LXe detector materials. However, this conclusion comes with a few caveats. Bakale




Figure 3.12: Left: Field dependence of electron attachment coefficient for the species N_2O (squares), O_2 (circles), and SF_6 (triangles) (figure taken from Ref. [8]). Right: The field dependence of the electron lifetime as measured by the SLAC XPM after pumpdown from ambient air.

and co-workers [8] have shown that the attachment coefficients of some LXe contaminants exhibit a marked dependence on the magnitude of electric field (fig. 3.12). The construction of the SLAC XPM is such that in our investigations we were limited to operating electric fields of 250 V/cm or less. This means that it is possible that the present studies are mis-characterizing the performance of the 200 kg EXO prototype, which will actually operate at a field of some ~ 3.0 kV/cm. Fig. 3.12 shows the field dependence for some electronegative molecules, from which it is apparent that at least one kind of impurity (N_2O) has an attachment coefficient that increases by as much as 2 orders of magnitude when the electric field strength is similarly increased. Early on in the experimental program, a set of field scans were performed in the SLAC XPM, an example of which is shown on the right hand side in fig. 3.12. If N_2O were a significant component of the impurities in the XPM, one would expect a factor of

~ 3 decrease in the measured τ over this range of electric fields. Since the data do not exhibit such a trend, it appears plausible that the SLAC system is N_2O free. However, it is difficult to draw any firm conclusions without additional data at higher electric fields. Furthermore, it cannot be ruled out that there might exist other, non- N_2O contaminants whose effect seems benign at XPM-compatible electric fields, only to become severe at fields approaching a few kV/cm. These issues notwithstanding, it is probably a safe bet that the XPM performance is a reasonable qualitative predictor of the performance of the full-scale EXO.

The 200 kg EXO prototype is expected to consist of a copper pressure vessel filled with LXe and immersed in 4 tonnes of HFE 7000. If we take the vessel to be a cylinder whose height equals its diameter and whose axis is oriented along the z direction, then, using the density of LXe (~ 3.3 g/cm³) we find a maximum electron drift distance of 44 cm, assuming that the drift field is parallel to the cylinder axis. If we further require that the cathode be placed in the midplane to form 2 sub-detectors, then the maximum electron drift distance is 22 cm. When electronegative impurities are present, the pulse height associated with the charge collected at the anode will depend on the z -position of the ionizing event. For an actual double beta decay experiment, therefore, the z of each event must be measured and an energy correction made, provided the attenuation length λ is accurately known. λ can be determined empirically by using the detector as an XPM. This is realized by using a calibration source that gives, say, a 2.6 MeV gamma line to generate ionizing events at several different z locations. If it is demanded that the correction factor introduce an error on the reconstructed energy no larger than the detector's intrinsic energy resolution, we can translate this requirement into a specification for the detector grade LXe purity. Expressed mathematically, this specification follows from

$$A = cA_m \tag{3.23}$$

where A_m is the measured pulseheight for a given event, A is the corrected pulseheight and c is the correction factor. Of course, c is given by

$$c = e^{-\frac{z}{\lambda}} \quad (3.24)$$

with z representing the reconstructed z coordinate of the event and λ the attenuation length as determined by running the detector with the calibration source. Since λ and therefore c are empirically determined they will be subject to experimental error. If we assume that the error on c is uncorrelated with the intrinsic detector energy resolution we can express our LXe purity requirement as

$$\frac{\Delta c}{c} = \frac{z_{\max}}{\lambda^2} \Delta \lambda \leq \frac{\Delta E}{E} \quad (3.25)$$

where $\Delta \lambda$ is the error on the attenuation length, Δc is the error on the correction factor $\Delta E/E$ is the intrinsic detector energy resolution, and z_{\max} is the maximum drift distance.

Here we have assumed that the empirical determination of z is considerably more accurate than the measurement of c . This assumption is justified because z is determined by measuring the time delay between the LXe scintillation and the signal from the ionization electrons. The EXO electronics amplify the signals with a shaping time of 3 μ s and sample at a rate of 1 MHz. This means that without doing a fit to the ionization waveform the time delay is determined to an accuracy of at worst ~ 1 -2 μ s. Performing a waveform fit will improve on this by a factor of almost 100. At 3 kV/cm the electron drift speed is about 2.0×10^3 m/s [1], so for the longest drift (22 cm) the z reconstruction accuracy using a fit is better than 1 part in 500.

We now invert equation 3.25 to find the LXe purity requirement

$$\lambda \geq \frac{z_{\max}}{R} \frac{\Delta \lambda}{\lambda} \quad (3.26)$$

where we have made the substitution $R \equiv (\Delta E/E)$.

Provided the readout electronics do not contribute overwhelmingly to R , we can

assume that R is strictly proportional to \sqrt{E} , where E is the energy of the event. To see whether this is a fair assumption, one must first consider the number of ionization electrons produced in each event. The W_i value, i.e. the average energy needed to create an electron-ion pair, for LXe has been measured and is known to have a value of 15.6 eV [19] at high electric fields. This means that a 2.5 MeV event will produce about 1.6×10^5 electrons on average. If we assume a Fano factor of 20 for the LXe [20], we estimate that the fluctuation on this number should be about 1800 electrons RMS. Comparing this with the design value of the EXO electronics noise of ~ 800 electrons RMS, we see that the amplifier noise contribution can be neglected to a fair approximation.

With this in mind, we can use the definition 3.10 to find

$$\frac{\Delta\lambda}{\lambda} = R \frac{\lambda}{(z_2 - z_1)} \sqrt{1 + \frac{Q_1}{Q_2}} \quad (3.27)$$

where Q_1 and Q_2 are the charges collected from two different z locations in the detector and $(z_2 - z_1)$ is their separation in z . To consider the worst-case conditions, i.e. the largest variation of the collected charge Q , we take Q_1 to be the charge gathered from a position close to the cathode ($z = 0$) and Q_2 to be at the anode ($z = z_{\max}$). In addition, λ will be the most difficult to measure accurately when the purity is high, i.e. when Q_1 and Q_2 are within a detector's resolution of one another. This leads us to

$$\frac{\Delta\lambda}{\lambda} = R \log \left(\frac{1}{1 - R} \right) \sqrt{2 - R} \quad (3.28)$$

where now we see that the fractional accuracy of λ is determined only by the accuracy R with which charge can be measured in the detector. This is true as long as the measurement of z is sufficiently accurate that its effect can be neglected.

Plugging 3.28 into 3.26 we obtain

$$\lambda \geq z_{\max} \log \left(\frac{1}{1-R} \right) \sqrt{2-R} \quad (3.29)$$

as the final expression for the requirement of the LXe purity. When we put in numbers for the 200kg LXe prototype, we obtain the inequality $\lambda \geq 123$ cm, or correspondingly, $\tau \geq 495$ μ s.

As shown in the previous section, this requirement has already soundly been met. In practice, however, there may be additional systematic contributions to $\Delta\lambda/\lambda$, so equation 3.29 should be taken with some caution. Nonetheless, it is encouraging that a simple-minded calculation such as this has produced a plausible result.

It should be pointed out that the purity constraints become more severe for future versions of EXO with larger fiducial masses; note that the lower bound on λ scales with the z dimension of the detector. Even so, the proposed maximum size of EXO is 10 tonnes, which means that the z dimension will probably scale up by a factor of ~ 4 from the 200 kg prototype. Similarly scaling up the constraint 3.25 yields the values 492 cm and 1980 μ s for λ and τ respectively. Again, the data seem to show that these levels of purity are routinely achievable using the techniques presented in this chapter.

3.7 References

- [1] Miller *et al.*, Phys. Rev. **166**, 871 (1968).
- [2] A. Vertes, J. Phys. Chem. **88**, 3722 (1984), and references therein.
- [3] A.J. Walters and L.W. Mitchell, J. Phys. D: Appl. Phys., **36**, 11, 1323 (2003).
- [4] L. Palleschi, S. Sacchetta and F.P. Ricci, Mol. Phys, **42**, 4, 961 (1981).
- [5] H.T. Davis, S.A. Rice and L. Meyer, J. Chem. Phys., **37**, 5, 947 (1962).
- [6] G. Ciccotti and G. Jacucci, Phys. Rev. Lett. **35**, 12, 789 (1975).

- [7] K. Wamba *et al.*, Nucl. Instr. and Meth. **A555**, 205(2005).
- [8] G. Bakale *et al.*, J. Phys. Chem. **80**, 2556 (1976).
- [9] G. Carugno *et al.*, Nucl. Instr. and Meth. **A292**, 580 (1990).
- [10] 3M™ Novec™ engineered fluid HFE-7000 from 3M, Inc., St. Paul, MN 55144
- [11] From ICG PolyCold Systems International, San Rafael, CA 94903
- [12] SAES Pure Gas, Inc., MonoTorr® Phase I Purifier Product Manual, SAES Pure Gas, San Luis Obispo, CA.
- [13] H. Wang, “ZEPLIN-II detector for direct dark matter search,” talk presented at SAGENAP 2000.
- [14] H. D. Pruett and H. P. Broida, Phys. Rev. **164**, 1138 (1967).
- [15] The discussion in this section is essentially a generalization of the discussion given in C. Hall, EXO Collaboration E-log Document #249
- [16] K. Skarpaas, *private communication*.
- [17] R. Neilson, *private communication*.
- [18] O. Bunemann *et al.*, Can. J. Res. **A27**, 191 (1949).
- [19] T. Takahashi *et al.*, Phys. Rev. A **12**, 1771 (1975).
- [20] E. Conti *et al.*, Phys. Rev. B **68**, 054201 (2003).

Chapter 4

Positive Ion transport in Liquid Xenon

In this chapter a study of the positive ion mobility in LXe is presented. The main aim of this study was to demonstrate the feasibility of repeatably transporting ions across short distances in LXe and moving them from the liquid phase to the gas phase with an electrostatic probe. This procedure is considered as part of a scheme for isolating and detecting the $^{136}\text{Ba}^{++}$ daughter ion that is produced in neutrinoless double beta decay.

Prior work has shown [1] that ions whose first ionization potential (IP) is less than, and whose second ionization potential is greater than, the local band gap energy in liquid Xe (9.5eV [2]) are stable in the +1 charge state in a pure (< 1 ppb oxygen equivalent) liquid Xe environment. Since barium's first and second IP's are 5.2eV and 10.0eV respectively, the implication is that the Ba^{++} ion that is created in a double beta decay should form Ba^+ when it interacts with the surrounding Xe atoms, and remain so indefinitely. This is important because it is technologically more practical to perform single ion spectroscopy on Ba^+ than Ba^{++} [3]. If it is indeed true that the Ba^+ ion is stable in liquid Xe, then it seems plausible that Ba^+ extraction and identification could be used to tag individual double beta decays in liquid Xe.

Part of EXO's current R&D is to develop a probe to deliver the Ba^+ to an ion trap, and a prototype probe has been used to measure the Th^+ mobility in LXe. The ion

mobility is important because it determines the ion collection time. While ion mobilities in nonpolar liquids have been studied for some time now ([1],[4]), a measurement of Ba^+ in liquid Xe had not been reported until after this work had been completed [5].

The present study uses $^{226}\text{Th}^+$ because it has similar chemical properties to Ba^+ and is experimentally much more attractive. This is because its radioactivity makes it straightforward to produce in an ionized state and to detect with good signal-to-noise. This chapter reports a measurement of the mobility of ^{226}Th ions in liquid Xe at 163.0 K and 0.9 bar pressure.

4.1 Experiment

Our system, schematically shown in Fig. 4.1, consists of a vacuum-insulated cell which we fill with about 10cm^3 of LXe. The cell is cooled by a cold finger that forms a thermal link with a liquid nitrogen reservoir. The temperature is maintained at 163.0 ± 0.1 K by a resistive heater. The cell, built almost entirely out of UHV-compatible materials, is baked at 400 K for several days, after which it is pumped to a 7.0×10^{-8} mbar vacuum by means of two turbomolecular pumps. When filling the cell, we pass research-grade Xe (which, as supplied, is quoted as 99.999% pure) through a hot Zr getter [6], which is rated to remove reactive impurities to the <1 ppb level. These precautions are taken to minimize any effect of impurities that might influence ion transport characteristics.

^{226}Th ions are produced by a ^{230}U source mounted inside the cell. The ^{230}U was obtained as a decay product of ^{230}Pa , which is manufactured using a cyclotron by means of the reaction $^{232}_{90}\text{Th} + p \rightarrow ^{230}_{91}\text{Pa} + 3n$ on natural thorium [7]. After irradiation, the ^{230}Pa was left to β^- decay for a few of its 17.4-day half lives, allowing some build-up of ^{230}U , which was then chemically separated and electroplated onto a platinum disk. Fig. 4.3 shows the decay scheme for the final ^{230}U source.

^{226}Th ions are formed as a result of the 20.8-day α decay of ^{230}U . The average recoil energy of $\sim 100\text{keV}$ extracts ^{226}Th atoms from the surface in a highly ionized state.




Figure 4.1: Schematic of the probe cell. The lower half of the cell is insulated from the room-temperature laboratory by a large vacuum vessel (not shown).

As the ^{226}Th ion enters the liquid it exchanges charge with the surrounding Xe atoms and reaches equilibrium in some final charge state. Although the energetics of the charge exchange in principle favor a final state containing Xe^+ and Th^+ [1], a large number of the emerging Th ions can also recombine with free ionization electrons that arise from the original α decay [1]. This effect presents a major problem as it greatly suppresses the ion yield.

In fig. 4.2 data that highlights the difference in behavior between ions in liquid and ions in gas is shown. In 231 Torr of Xe gas (the left hand plot), the ion yield seems to have no discernible dependence on the magnitude of electric field, but in liquid the ion yield seems to increase with increasing field strength. This can be understood as being due to the effect that the electric field has in reducing the recombination by drawing free electrons out of the liquid before they have a chance to attach to the Th^+ ions. Although the overall scale of the ion yield appears similar in the two data sets, it should be pointed out that the gas data were taken using a shorter ion collection time and count time. In addition, the probe was situated much farther from the detector than it was for the liquid test. When these issues are corrected for,




Figure 4.2: Left: Ion yield vs electric field for ions collected in 231 Torr of ambient Xe gas. Ions are collected even at zero field because they diffuse through the gas. Right: Ion yield vs electric field for ions collected in liquid. The source-probe distance for the gas experiment was 0.5 cm while for the liquid it was 0.025 cm.

we find that the ratio of ion yield in gas to liquid is roughly 200:1 for the range of electric fields used in the mobility measurement.

We collect ions on a stainless steel probe which we lower into the liquid. Ions are attracted to the probe when it is grounded while applying a positive bias to the source electrode. The tip of the probe is fashioned into a 5 mm \times 16 mm, 2 mm thick rectangular stainless steel plate, ensuring a spatially uniform electric field in the vicinity of the source.

The ions that land on the probe are detected by retracting the probe vertically to a counting station, located above the liquid surface. It consists of a 50 mm² active area ion-implanted Si detector[8] that is read out using standard pre- and post-amplification electronics. The ion yield is determined by counting the α particles emitted in the 30.6 min ^{226}Th decay, along with those produced by the subsequent α decays of ^{222}Ra , ^{218}Rn , and ^{214}Po . Fig. 4.3 schematically shows the complete decay chain.

It should be pointed out that in addition to the ^{226}Th the ^{230}U source also produces ions arising from the other α decays in the chain. These can also drift to the probe and be detected at the counting station. However, because of their considerably shorter lifetimes, they present a negligible background to the ^{226}Th signal provided some short time elapses between the probe extraction and the counting. In addition, the low angular acceptance of our α particle detection system ($\sim 5\%$) ensures that the probability of detecting more than one α particle from the 4- α decay sequence of a given ^{226}Th ion is negligibly small, so the ions are not double-counted.

We measure the ion mobility in liquid Xe by determining the ion transit time from source to probe at a given probe distance and electric field. This is achieved by biasing the source with an alternating potential produced by a computer driving a high voltage amplifier.[9] The system is set to produce a square wave switching between V_{drift} and V_{sup} with respect to the probe, where V_{drift} can be varied from 1.0kV to 4.0kV, and V_{sup} is fixed at -4.0kV. V_{drift} is set for a time $t_{drift} \approx 1\text{s}$, very long compared to the $\sim 20\mu\text{s}$ switching time of the high voltage system. During each full cycle, the voltage is reversed for a time t_{sup} . When this occurs, the ions that do not reach the probe during t_{drift} are drawn back toward the source. This occurs when t_{drift} is too short for the ions to complete the journey to the probe. We maximize our total ion integration time by having $t_{sup} \leq t_{drift}$ and $|V_{drift}| < |V_{sup}|$. [1] During each




Figure 4.3: Decay scheme for ^{230}U . Decays with branching ratios less than 5% are not shown.




Figure 4.4: Top left: α spectrum of the ^{230}U source in vacuum.[10] The ^{230}U peaks are at 5.82 MeV and 5.89 MeV, while ^{226}Th is identified by peaks at 6.23 MeV and 6.34 MeV. ^{222}Ra , ^{218}Rn , and ^{214}Po correspond to the peaks at 6.56 MeV, 7.13 MeV, and 7.69 MeV respectively. Top right: α spectrum of the activity collected on the probe in 0.2 bar of Xe gas. Bottom: Total α rate versus time for the activity collected on the probe. The best fit is the superposition of two exponentials with half-lives of $1773 \pm 97.4 \text{ s}$ and $37.2 \pm 2.84 \text{ s}$. These are the ^{226}Th and ^{222}Ra half-lives, whose respective published values are 1830 s and 38 s [11].

data taking run, V_{drift} is set to a particular value and the field is switched for 1800 s. The constant duration of collection runs ensures that the amount of ion loss due to radioactive decay during collection is uniform. We then measure the ion yield at the counting station over a period of 2 hours to allow sufficient time for the collected ^{226}Th ions to decay before the next run. Data are collected at four different electric fields, using four source-probe voltages and two source-probe separation distances.

We determine the background by performing 2-hour counts between data runs in which no ions are collected. In these runs the source-probe voltage remains reverse-biased. The observed rate during the background runs is averaged over all runs, giving the mean background rate and an associated statistical error. The resulting value of 4.50 ± 0.87 counts per 2-hour run is subtracted from the number of counts observed at each t_{drift} . When ions are collected, the typical rate is ~ 30 counts per run. Sources of background include cosmic rays and neutral Th atoms that evaporate from the surface of the liquid and diffuse through the gas.

For the purpose of characterizing the system a test in Xe gas was also done. The cell was filled with 0.2 bar of room-temperature Xe gas and the probe was alternately brought close to the ^{230}U source and raised to the counting station, where ions were respectively collected and counted. This was repeated many times over several hours to build up counting statistics. The electric field between source and probe for ion collection was roughly ~ 200 V/cm for this test. Shown in Fig. 4.4 are the α spectrum and time histogram for the collected Th and its decay daughters. For comparison, the ^{230}U spectrum, as measured in vacuum at LLNL,[10] is also shown. An important conclusion from the gas test was that ion recombination with free ionization electrons is a factor of ~ 200 times more efficient in LXe than in gaseous Xe.

4.2 Results

Shown in Fig. 4.5 are plots of the ion yield vs t_{drift} . For each data point, the ion yield is determined by dividing the number of detected α particle counts by the number of cycles, giving the ion yield per cycle. The background contribution to each of the

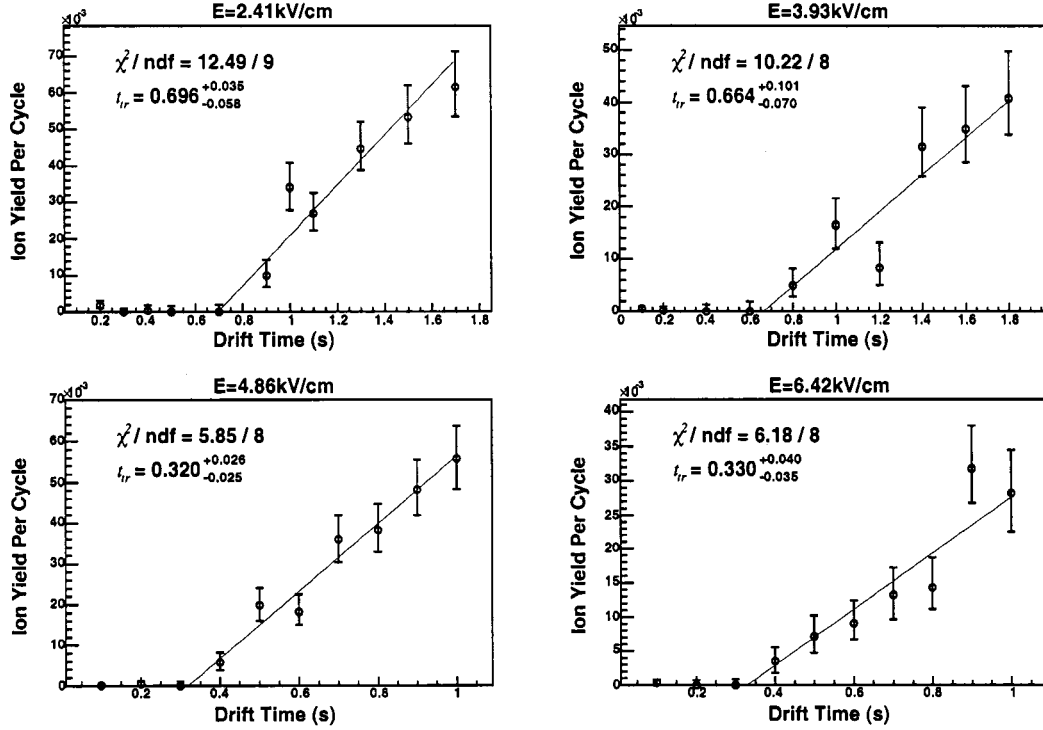


Figure 4.5: Ion yield as a function of t_{drift} for the four electric fields 2.5 kV/cm, 3.9 kV/cm, 4.9 kV/cm, and 6.4 kV/cm. Error bars shown are statistical only. We take the transit time in each case to be the knee point in the fit.

data points in these units is ≤ 0.003 . We find that the ion yield is consistent with zero (after background subtraction) when the t_{drift} is short, while at longer times, the ion yield increases linearly with t_{drift} . The abscissa of the join point between these two behaviors is the ion transit time. We fit the ion yield data to the following two-parameter piecewise linear function,

$$\begin{aligned}
 Y &= 0 & t_{drift} < t_{tr} \\
 Y &= m(t_{drift} - t_{tr}) & t_{drift} \geq t_{tr}
 \end{aligned}$$

d ($\times 10$ cm)	V_{drift} ($\times 10$ kV)	E (kV/cm)	t_{tr} ($\times 10$ s)	v ($\times 10$ cm/s)
3.87 ± 0.08	9.32 ± 0.14	2.41 ± 0.06	$6.96^{+0.35}_{-0.57}$	$0.56^{+0.27}_{-0.46}$
5.87 ± 0.08	23.10 ± 0.24	3.93 ± 0.07	$6.64^{+1.01}_{-0.70}$	$8.84^{+1.13}_{-0.93}$
3.87 ± 0.08	18.88 ± 0.10	4.86 ± 0.11	$3.20^{+0.26}_{-0.25}$	$12.1^{+0.98}_{-0.94}$
5.87 ± 0.08	37.71 ± 0.32	6.42 ± 0.11	$3.30^{+0.40}_{-0.35}$	$17.8^{+2.15}_{-1.89}$

Table 4.1: Electric fields (E), voltages (V_{drift}), and drift distances (d) used in this experiment, showing their associated systematic errors. The measured transit times (t_{tr}) and velocities (v) are the values returned by the fit. For these last two quantities, only statistical errors are given.

where Y is the ion yield per cycle, m is the slope (ion yield per cycle per unit t_{drift}), and t_{tr} is the ion transit time. The value of t_{tr} yielded by the fit is used to calculate the ion drift velocity. To perform the fits, and to calculate the statistical error on t_{tr} in each case, we use an algorithm due to D.J. Hudson [12] that is designed specifically to handle the difficulties that arise when performing a piecewise fit. One peculiarity of such a fit is that the error bounds are almost always asymmetric [12]. The results of the fits are summarized in Table 4.1.

Fig. 4.6 shows a plot of the computed ion drift velocity as a function of electric field. The data are consistent with the linear relation $v = \mu E$, where v is the ion drift speed, μ is the mobility, and E is the electric field. From a fit to the data in Fig. 4.6 we find an ion mobility $\mu = 0.240 \pm 0.011$ cm²/(kV-s). A fit not constrained to go through the origin results in $\mu = 0.268 \pm 0.031$ cm²/(kV-s), a y-intercept of -0.095 ± 0.088 cm/s, and $\chi^2 = 1.14$. This is fully consistent with the original result.

The mobility and its statistical error are determined largely by the data taken at the lowest electric field. This point is more precise for two reasons. First, the ion transit time at this electric field is the longest of the four datasets, meaning that the fractional error on the velocity is the smallest. Secondly, in this dataset the ion transit time happens to be within ~ 0.01 seconds of the nearest t_{drift} at which we collected data, and this results in a precise measurement. While, in principle, better accuracy

could be obtained by repeating the measurements and always taking a point near this optimal time, the 20.8-day half life of the ^{230}U source makes it impractical to repeat the entire experiment. It is also true that the data from which the first point is determined was recorded earlier in the lifetime of the ^{230}U source; this means that the corresponding ion yield statistics were relatively high. This fact also contributes to the smaller fractional error, but its effect is only of secondary importance.

Systematic uncertainties on the mobility derive from small mechanical vibrations of the probe when the electric field is switched and from the accuracy of the voltage gain in the amplifier. These contributions are estimated to give a combined error of $\pm 0.011 \text{ cm}^2/(\text{kV}\cdot\text{s})$. Table 4.1 shows the absolute errors on the voltages, drift distances and electric fields. The systematic error on the ion mobility was computed by shifting the electric field and drift distance values by their corresponding errors and re-doing the fit.




Figure 4.6: Global fit to the ion drift velocity as a function of electric field for the entire data set.

4.3 Discussion

Here we see a direct manifestation of the large mismatch between the mobilities of electrons and those of ions in LXe. Although it should come as no surprise that the ions drift more slowly than electrons a difference of a factor of about one million is a little unexpected at first glance. The physics behind this is commonly attributed to the formation of a “snowball” structure of liquid atoms that adhere to the ion and travel with it [13]. We will delve more deeply into this physics in the following chapter.

4.4 Conclusion

We have measured the ^{226}Th ion mobility in LXe at 163.0K and standard pressure and have found it to have a value of 0.240 ± 0.011 (stat) ± 0.011 (syst) $\text{cm}^2/(\text{kV}\cdot\text{s})$. In the following chapter this result is compared with similar measurements with other ions by other workers, along with a set molecular dynamics simulations.

4.5 References

- [1] A.J. Walters and L.W. Mitchell, J. Phys. D: Appl. Phys., **36**, 11, 1323 (2003).
- [2] U. Asaf *et al*, Phys. Rev. B, **10**, 10, 4464 (1974).
- [3] M. Danilov *et al.*, Phys. Lett. B **480**, 12 (2000).
- [4] H.T. Davis, S.A. Rice and L. Meyer, J. Chem. Phys., **37**, 5, 947 (1962).
- [5] S. Jeng, W.F. Fairbank Jr and M. Miyajima, submitted to Phys. Rev. B.
- [6] From SAES Getters USA, Inc. Colorado Springs, CO 80906.
- [7] H. Tewes, Phys. Rev. **98**, 25 (1955).
- [8] ORTEC®ULTRA ion-implanted Si detector, from ORTEC, Oak Ridge, TN 37830.

- [9] From Trek, Inc. Electrostatic Instrumentation & Devices, Medina, NY 14103.
- [10] K. Roberts, private communication.
- [11] R. Firestone *et al.*, *Table of Isotopes*, John Wiley & Sons, Inc. 1996.
- [12] D.J. Hudson, J. Am. Stat. Assoc. **61**, 1097 (1966).
- [13] W.F. Schmidt, K.F. Volykhin and A. Khrapak, J. Electrostat. **47**, 83, (1999).

Chapter 5

Computerized Studies of Ion Mobility in Liquid Xenon

In this chapter we delve into the physics of positive ion transport. The bulk of the discussion will center around a recent model that we test by comparison with a computerized molecular dynamics simulation. We also check the simulation results against experimental data from other workers.

5.1 Theory of Positive Ion Transport

To understand ion transport theoretically, a naive first step might be to treat the ion as a tiny, rigid sphere moving through a continuous fluid with a well-defined viscosity. In this case we would use the well-known Stokes' law [1], which states that a spherical object pulled through a fluid with a constant force moves at a constant velocity whose magnitude is directly proportional to that of the applied force. The law is expressed mathematically as

$$F_d = 6\pi\eta Rv_d \tag{5.1}$$

where F_d is the applied force, η is the liquid viscosity, R is the radius of the sphere, and v_d is the drift velocity.

Substituting qE for F_d and rearranging the terms we conclude that an ion with charge q and radius R has the mobility

$$\mu = \frac{q}{6\pi\eta R}. \quad (5.2)$$

Unfortunately, this equation overestimates the mobility by more than an order of magnitude for ions like Ba^+ whose radii lie in the range 1-2 Å. This problem led to the development of an alternative model first proposed by K.R. Atkins who was studying the transport of He ions in liquid He [4]. In his treatment, Atkins suggested that the ion's strong electric field exerts a pressure on the liquid of sufficient magnitude for a microscopic ice layer to form around the ion. The ion thus becomes part of a large "snowball," giving it an effective radius that is much bigger, and the mobility is suppressed and brought into closer agreement with Equation 5.2.

Hilt, Schmidt and Khrapak[5] (henceforth referred to as HSK) adapted the Atkins model to LXe and predicted a very similar effect. Like Atkins, they used thermodynamic arguments in which a balance is struck between the volume work in compressing the liquid to the solid phase, the electrostatic work and the solid-liquid surface tension. This results in a temperature-dependent snowball radius that diverges as the liquid triple point is approached. The divergence is controlled only by the small surface energy term, which, in effect limits the cluster radius to about 25 Å. Experimentally, this effect would be manifested by a sharp dip in the ion mobility at the liquid triple point.

Shown in Table 5.1 are some selected experimental mobilities for atomic ions in LXe. From these results it is clear that the HSK model does quite well in predicting the overall scale of the observed mobilities at these temperatures. However, the significant difference between the $^{208}\text{Tl}^+$ and $^{226}\text{Th}^+$ numbers suggests that these species differ enough microscopically that the HSK model, when applied to them, begins to break down.

Further problems with the HSK model are exposed in Figure 5.1, where the measured temperature dependences of several of the alkaline earth ions is plotted [2]. Superimposed on the plot are three HSK calculations that each use a different estimated value

for the (unknown) solid-liquid surface tension. Clearly, the data do not exhibit the same trend. Moreover, the predicted dip at the triple point of 161.4K is conspicuously absent.

Overall, it can be concluded that HSK provides a reasonable rough approximation of the ion mobility only under certain conditions. In addition, it is worth understanding the situations in which the ion identity becomes important because this would offer clues as to how HSK could be modified to account for this.

A theory put forward by Davis, Rice and Meyer (henceforth referred to as DRM) might offer some clues as to how this might be accomplished.[3] These authors have derived an expression for the positive ion mobility in which the internal structure of the ion and its interaction with the liquid are taken into account. In this model the drag force experienced by the ion is determined by the microscopic interaction potentials rather than the bulk properties of the liquid. Essentially, DRM propose replacing the denominator in Equation 5.2 with a so-called “friction constant”. This quantity is constructed out of two terms, the first reflecting the hard collisions between liquid atoms and the ion and the other deriving from the long-range attraction between them. It is conceivable, at least in principle, that as the ion-liquid interaction increases in strength, bound states between the ion and several liquid atoms can occur, particularly as the temperature of the liquid is lowered. Nevertheless, the treatment reported in Ref. [3] is specific to conditions in which this effect is not present, i.e., the ion-liquid interaction is too weak and the liquid temperature is too high. Indeed, the DRM model has only been compared with experimental data from systems involving molecular ions in liquids at temperatures well above the liquid triple point, namely He_2^+ in LHe (Ref. [7]), Ar_2^+ in LAr (Ref. [3]), and Kr_2^+ in LKr (Ref. [6]). While these systems are satisfactorily described by the model, it is clear that some modification will be necessary in order to treat the case of atomic ion mobilities close to the triple point. This is because the atomic ions under consideration have interaction strengths that are tens of times stronger[12, 19] than those of the molecular ions previously considered.

To investigate these details experimentally, one could perform an accurate measurement of the temperature dependence of the ion mobility in several different ion-liquid

Ion	T (K)	P (bar)	Mobility ($\text{cm}^2 \text{kV}^{-1} \text{s}^{-1}$)	
			HSK	Expt.
$^{208}\text{Tl}^+$	163.0	0.9	0.31	0.133 ± 0.004^1
$^{226}\text{Th}^+$	163.0	0.9	0.31	$0.240 \pm 0.011 \pm 0.011^2$
Ba^+	168.0	1.2	0.33	$0.211^{+0.020}_{-0.012}^3$

Table 5.1: Ion mobilities in LXe, comparing a model due to Hilt, Schmidt and Khrapak (HSK) [5] to the experimental data.

systems. This would require multiple, well-understood and well-controlled sources of ions of known identity. A complementary approach would be to use the ^{226}Th ion source in several different liquefied noble gases at various temperatures.

The next best thing to such a series of experiments is the computer simulation, the mechanics of which will be discussed in the following section.

5.2 Computer Calculation

A set of computer simulations of the ions Ca^+ , Sr^+ , Mg^+ , Ba^+ , and Tl^+ in LXe at various temperatures was performed. For comparison's sake the charged Xe dimer Xe_2^+ was also studied at conditions far from the LXe triple point. The simulations were done using a modified version of the Verlet algorithm [9]. L. Verlet's algorithm was originally developed for modeling atomic liquids such as liquid argon and xenon. It described these liquids as collections of atoms interacting in pairs according to the Lennard-Jones (L-J) potential. The L-J potential is given by

$$V_{L-J} = 4\epsilon \left[\left(\frac{\sigma}{r} \right)^{12} - \left(\frac{\sigma}{r} \right)^6 \right] \quad (5.3)$$

where ϵ is the potential minimum (the “well-depth”), r is the center-to-center displacement of the two interacting atoms, and σ is the center-to-center displacement at which the potential equals zero (also known as the atomic “hard-sphere diameter”).




Figure 5.1: Ion mobility vs temperature for four of the alkaline earth ions in LXe. The three curves are the results of calculations using the HSK model using three different estimated values for the (unknown) surface tension in the liquid Xe-solid Xe interface. σ_0 is the value estimated by HSK in Ref. [5] Figure taken from Ref. [2]

The parameters σ and ϵ are specific to the identities of the atoms in the interacting pair and must be measured experimentally. The potential 5.3, though simple, accurately describes many atomic systems, and molecular dynamics simulations like Verlet's have been shown to agree quite well with experimental data [9].

In the present calculations, the L-J potential was used to describe the Xe-Xe interaction. On the other hand, the Xe-ion interaction was described in one of four ways. For the alkaline earth ions, the ion-liquid interaction was modeled by a set of potentials kindly provided by W. Fairbank [10]. These potentials had been determined either directly from spectroscopic data (for Mg^+ , Sr^+ and Ca^+) or by extrapolation (for Ba^+). For Tl^+ , the available experimental data [11, 12] only covered the long-range, attractive part of the potential, and the repulsive part was guessed at by fitting the published data to a standard form. The fourth kind of potential was that for Xe_2^+ which was described by DRM in Ref. [3]. This potential is essentially the L-J, plus an extra term to account for the interaction of ion with the induced atomic dipole moment that it provides to the Xe atom.

To save on computer time, the Xe-Xe interaction potential was truncated at a separation equal to 2.6 times the Xe-Xe hard core radius σ (4.07 Å). The ion-Xe interaction was truncated at 3.65σ . All of the relevant potentials are plotted, in atomic units, in Figure 5.2. We stress that all of the calculations make an implicit approximation by neglecting any dipole-dipole contribution to the interactions between Xe atoms that are polarized by the electric field of the ion. We shall return to this important point later.

For each case, the computer model consists of one ion surrounded by 499 Xe atoms inside an imaginary cube whose side is chosen to give the desired density. For the systems in the present study the cube side was of order 35 Å. The entire system is subject to periodic boundary conditions, which is to say that a particle that passes through the side of the imaginary bounding box reappears at the opposite wall. The initial conditions are chosen such that the 500 particles are arranged as they would be in pure LXe. This is done by assigning them positions picked at random from the phenomenological radial distribution function of E. Matteoli and G. Mansoori [13] corresponding to the desired temperature and pressure. Finally, each atom

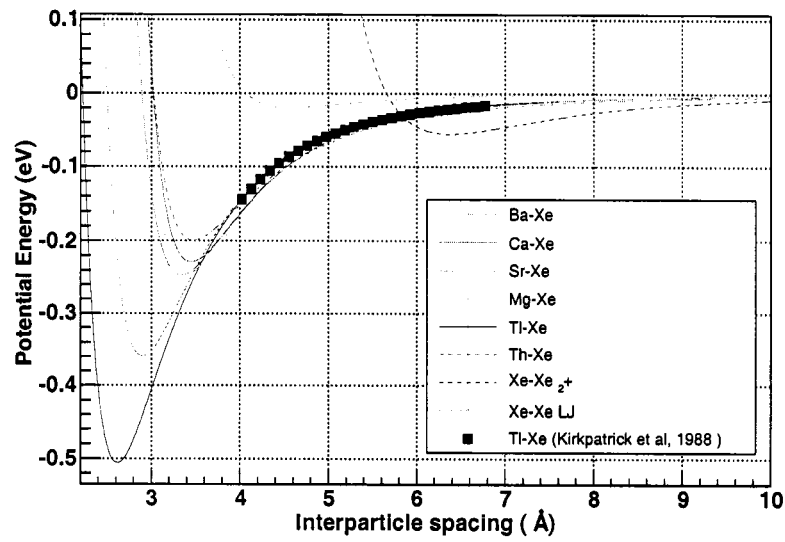


Figure 5.2: The liquid-liquid and ion-liquid interaction potentials used in the computer simulation. The black squares are the data points of Ref [12].

is assigned a random thermal velocity taken from a Maxwell-Boltzmann distribution of the appropriate temperature.

With these initial conditions in place the system is time-evolved in discrete 2.0×10^{-15} s steps. The force on each particle is computed by taking the gradient of the interaction potential between each particle and its neighbors. Using Newton's second law, the particle motion is then found to second order in the time step. This computational technique is commonly referred to as the "velocity Verlet" algorithm [14].

For the alkaline earth ions and Tl^+ , simulations were run at the temperatures 161.4 K, 163.0 K, 168.0 K, 172.0 K, 176.0 K, and 180.0 K at densities corresponding to the saturation curve. For Xe_2^+ , the temperature/pressure combination 192.1 K/55.9 atm was used to match the conditions studied by Davis, Rice and Meyer in Ref. [3].

The ion mobility in each case was determined in two ways. First, the ion mobility was inferred from the ion's diffusion under zero field conditions. This was done by setting up the system and taking 5000 time steps to reach equilibrium, followed by another 145000 steps to study the ionic diffusion. In the second method, the equilibration period was followed by a 20000-step run in which the system was subjected to an electric field of 1.0×10^{11} V/cm. For the simulation at high field, an approximation was made in which the effect that the applied field has in polarizing the xenon atoms was completely neglected.

5.3 Simulation Results

5.3.1 Radial Distribution Functions

The radial distribution function (RDF or $g(R)$) for a system of particles is essentially the average number density as a function of distance from the center of any one particle in the ensemble. It is evaluated by finding the number of atoms per shell of radius R_i and thickness dR_i , dividing by the shell volume $4\pi R_i^2 dR_i$, where R_i is the distance from the i th atom, and averaging over all the atoms in the ensemble.

In our case, because we are only interested in the distribution of Xe atoms around the ion (of which there is only one), the ensemble average is replaced by an average over

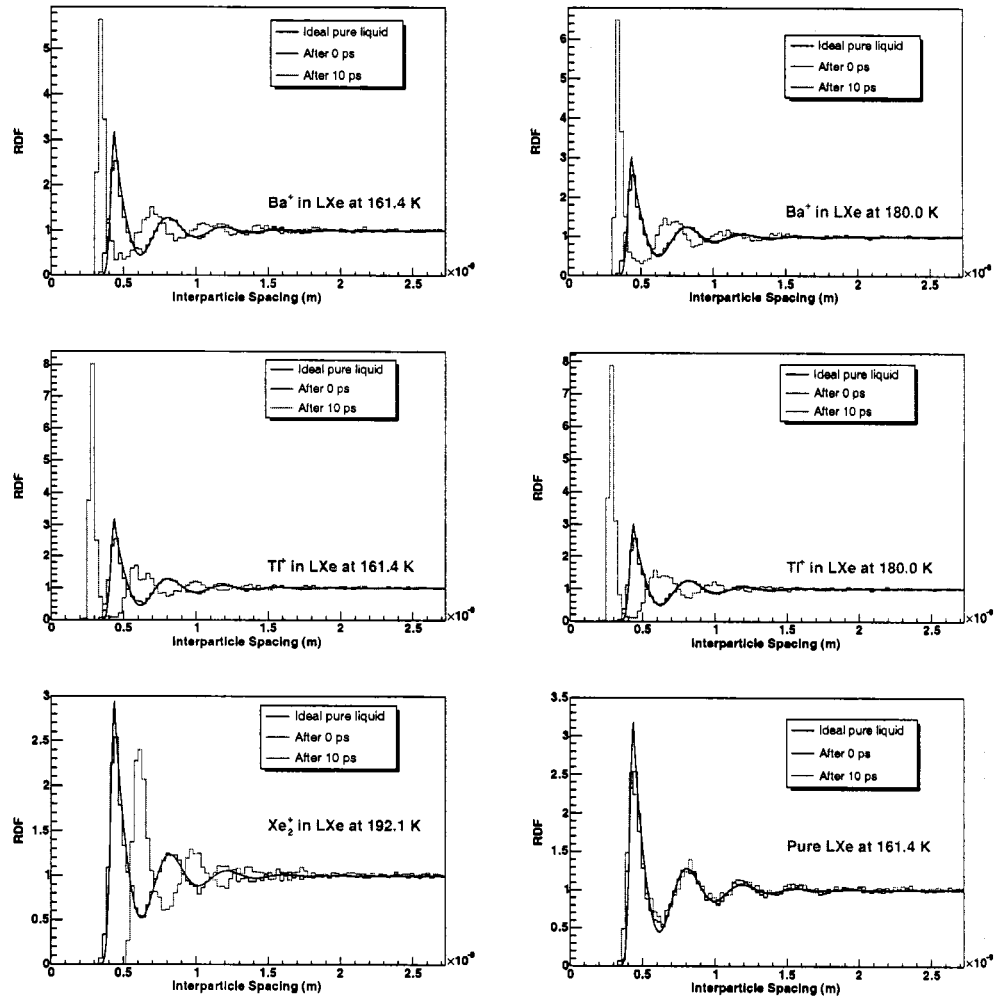


Figure 5.3: Radial distribution functions for: Ba^+ and Tl^+ in LXe at 161.4 K and 180.0 K, Xe_2^+ in LXe at 192.1 K, and pure Xe at 161.4 K. In each case, the red histogram is the initial RDF, the blue is the RDF computed after the 10 ps equilibration time and the black curve is the empirical model RDF for idealized pure LXe after Ref. [13].

many snapshots of the system at times separated by a long enough time interval (in this case, 1 ps) that they can be considered uncorrelated. This follows the Liouville theorem [1].

A representative sample of the computed RDFs are shown in Figure 5.3. Because the results for the monatomic ions were qualitatively very similar, we only show plots for Ba^+ and Tl^+ . The RDFs were computed by histogramming the number of atoms

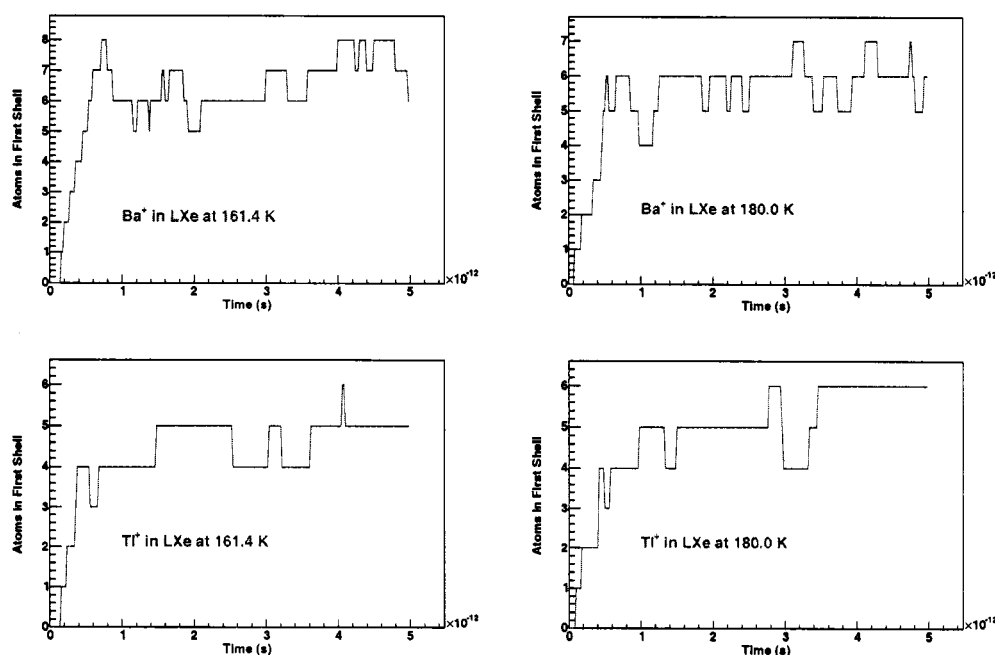


Figure 5.4: Number of atoms in the cluster as a function of time for the ions Ba^+ and Tl^+ in LXe at 180.0 K and 161.4 K

whose centers are in each spherical shell of thickness 0.3 \AA as we move away from the ion. This is done once per picosecond and averaged starting at 10 ps. For comparison, we show the starting RDF in each case along with the a calculation of the pure liquid RDF using the empirical formula of Ref. [13].

As shown, the equilibrium RDFs for the atomic ions are qualitatively quite different

from those of pure LXe and Xe_2^+ . All of the simulations begin with the 500 particles arranged such that their radial distribution function is that of the pure liquid. However, within 2 ps, the RDFs of the monatomic ions begin to exhibit a sharp spike at the equilibrium separation in their interaction potential with the liquid. This corresponds to a clustering of Xe atoms around the ion as they sit in the deep well characterizing the ion-Xe interaction. On average, the cluster contains between 6-10 atoms, and the question that comes to mind is whether this feature is, in fact, the putative snowball of HSK.

To answer this question, we can look at the time evolution of the cluster. If it were truly a solid layer of Xe ice, one would expect it to form very rapidly, and, once formed, it should be permanently frozen onto the ion, exhibiting essentially no time dependence. In Figures 5.4 and 5.5 we show the time evolution of the cluster. Clearly, for the atomic ions, the cluster forms very quickly and is quite stable. However, it is also extremely dynamic. Its atoms are not at all permanently attached and are swapped with fresh atoms from the bulk roughly twice per picosecond. For comparison we also show the time evolution of the swarm of Xe atoms immediately surrounding both neutral Xe and Xe_2^+ . In this case we see very similar behavior, only the exchange seems to occur about a factor of 4 more quickly. This seems to indicate that the behavior of the Xe cluster around the atomic ion is more analogous to that of a very dense liquid as opposed to a true solid.

5.3.2 Ion Mobilities

The ion mobility is calculated using two methods. The first is based on inferring the mobility from the diffusion coefficient D for the ion. To calculate D we use that fact that for long times τ , the distance traveled by the ion along x , y or z is Gaussian distributed with mean 0 and $\sigma = \sqrt{2D\tau}$ [15]. By “long times” we mean time intervals over which the ionic motion is no longer correlated with itself. Once we have established how long a time interval τ is sufficient, we may evaluate D by fitting a Gaussian to a histogram of the ion’s one-dimensional displacement taken between regular time intervals τ . To ensure that τ was sufficiently long the velocity

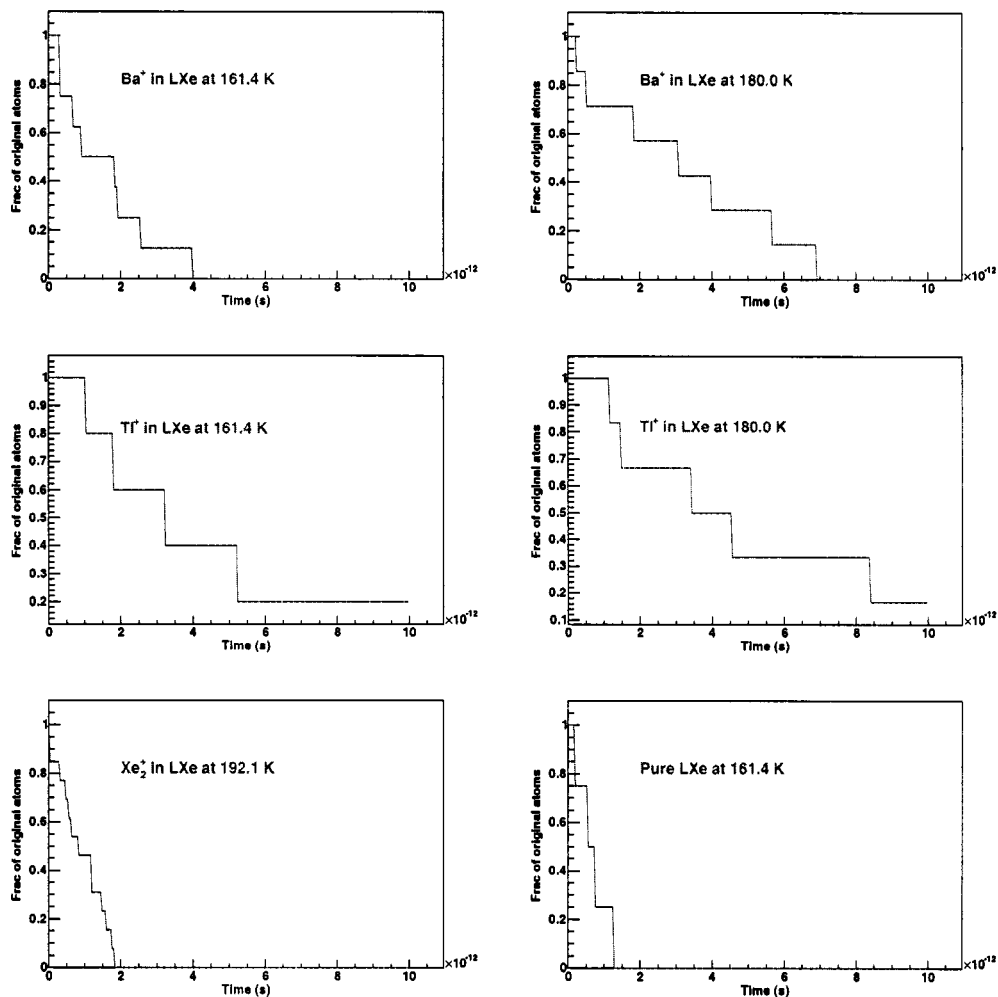


Figure 5.5: Number of unexchanged atoms in the cluster peak as a function of time for: Ba^+ and Tl^+ in LXe at 180.0 K and 161.4 K, Xe_2^+ in LXe at 192.1 K, and pure Xe at 161.4 K. For the latter two cases, the “cluster” is defined as the narrow region within ~ 0.5 Å of the first maximum in the RDF.




Figure 5.6: Histogram of the Ba^+ ion's displacement along x , y , or z after 3 ps for the freely diffusing ion under zero-field conditions. The diffusion constant is computed from the σ parameter of the fitted Gaussian.

autocorrelation function for the ion was studied. This function is given by

$$U(t) = \frac{1}{|\mathbf{v}(0)|} \langle \mathbf{v}(0)_i \cdot \mathbf{v}(t)_i \rangle \quad (5.4)$$

where $\mathbf{v}(t)$ is the (3-dimensional) velocity vector evaluated at time t and again, the angle brackets imply the ensemble average.

The velocity autocorrelation function is plotted in Figure 5.7 for the systems under study. Note that the baseline in each case is quite noisy. This is at least partly due to the low statistics (recall that the average is only taken over time and not also over the 500-particle ensemble). Nonetheless, it is apparent that the time correlations in the system in each case die off after a few picoseconds. Based on this analysis, the interval 3.0 ps was universally chosen for τ for the determination of D . In Figure 5.6 we show a typical example of a Gaussian fit to the histogrammed one-dimensional displacement after 3 ps. As shown, the histogram shape is very well described by a Gaussian.

In well-behaved systems the ion mobility is directly proportional to the diffusion constant and is given by

$$\mu = \left(\frac{q}{k\kappa T} \right) D \quad (5.5)$$

where μ is the ion mobility, q is the ionic charge, κ is the dielectric constant of the liquid, and k and T have their usual meanings. This equation is commonly known as the Nernst-Einstein relation [16].

After having calculated the mobilities using this first method, they were calculated again by running simulations in which the ion was subjected to a constant electrical force oriented along the z axis. For these calculations, the system was first allowed to equilibrate for 10 ps, or 5000 time steps. Then, a further 20000 time steps were taken with the electrical force turned on. The mobility was found by averaging the z -component of the ion's velocity over the final 12500 time steps and dividing by the product of the applied electric field strength (1.0×10^{11} V/cm) and the LXe dielectric constant (1.94). Despite the fact that the applied electric field is unphysically large

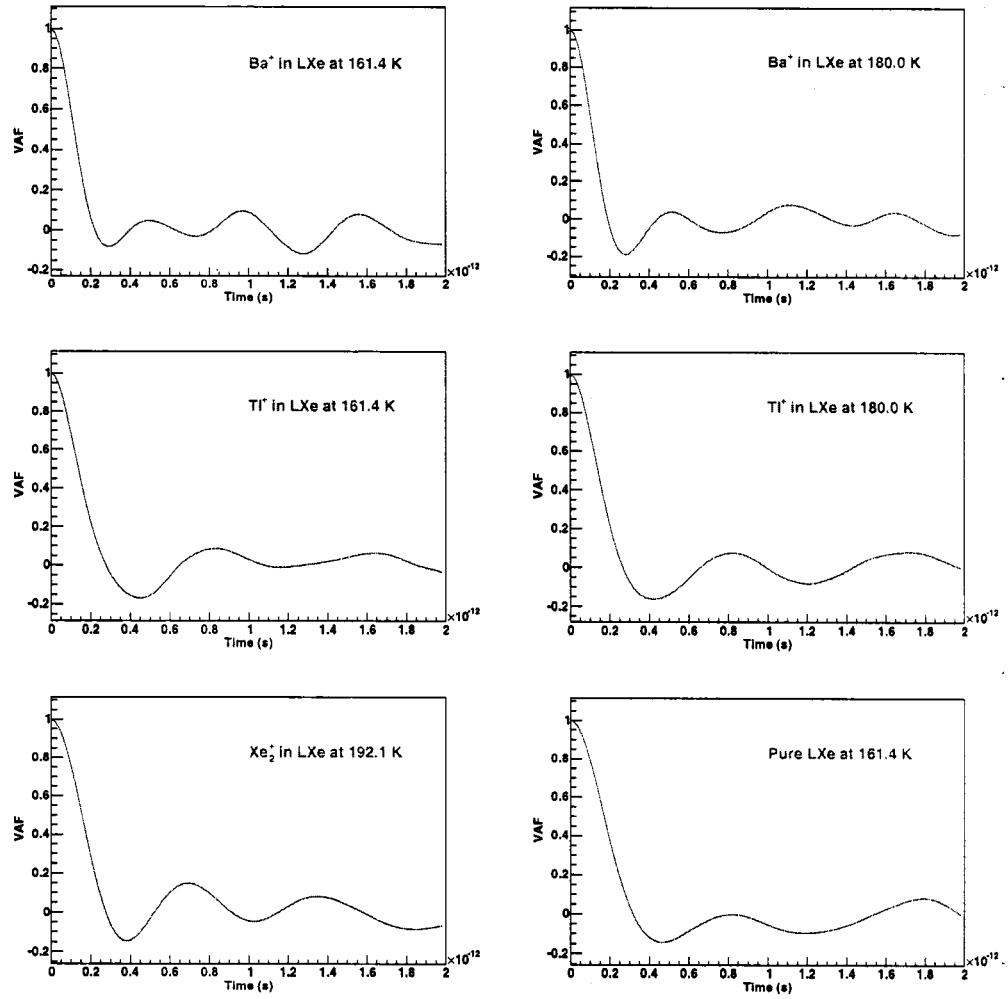


Figure 5.7: Velocity autocorrelation function of ionic motion for: Ba^+ and Tl^+ , in LXe at 180.0 K and 161.4 K, Xe_2^+ in LXe at 192.1 K, and pure Xe at 161.4 K

(it needed to be so because of the small size of the 500-particle system), there was no indication that structure and dynamics of the system were affected.

Incidentally, a calculation using the perturbation-and-subtraction (P-S) technique of Ciccotti and co-workers [17] was also attempted, but it failed to produce reliable results. In the P-S algorithm, the trajectory of the ion is computed twice, once with an infinitesimally weak electric field turned on and once with it turned off. The ion's response to the applied field is then computed by subtracting the two trajectories. In principle, this method is very sensitive and accurate because it effectively subtracts the "signal" from the thermal "noise" and allows one to study the system using realistic electric fields. However, the trouble is that this only works when the ion's random motion remains correlated with itself for times that are long compared with the time it takes the system to reach a steady state. As we have seen, the correlations die off quite rapidly, which causes the P-S method to break down.

In Figure 5.9 we plot the calculated ion mobility as a function of temperature for each case. Unfortunately, the simulation appears to over-estimate the mobilities by about a factor of 2. For the sake of computational expediency, the interactions between induced atomic Xe dipoles were left out of the simulation. If this actually constitutes a poor approximation, then it might explain the magnitude mismatch between experiment and simulation. A further problem is that the low statistics have translated to large errors on the individual results. This issue is particularly evident in the high field data which suffers not only from the low statistics, but also the fact that despite the enormous electric field, the ion drift velocity is small enough that it is severely corrupted by the random thermal motion. The result is worse scatter in the data points.

Despite these problems, it is apparent that the simulated temperature dependence qualitatively matches the data more closely than it does the HSK model. In other words, we may provisionally state that according to the simulation results, nothing special happens at the triple point. This assertion is supported by the fact that the RDFs and their temporal behavior do not change much with the temperature as we saw in the previous section.

If we take this to indicate that the ion mobility actually is independent of temperature, we may average the results over temperature for each ion in order to make a comparison between ions. When this is done, we again see a hint that the simulation is correctly supporting the experiments while contradicting HSK. Recall that in the HSK model the mobility is strictly independent of ion identity. A plot showing the experimental and computational results for the different ions is shown in Figure 5.8. As shown, both the simulation and the experimental data seem to suggest a downward trend in mobility as we go from Ca^+ to Tl^+ . We discuss why the HSK model




Figure 5.8: Experimental and simulated ion mobilities for the five atomic species under study. The experimental data are from Ref. [2].

may be failing to correctly describe the systems under study in the following section.

5.4 Discussion

To see how the HSK model might break down, we first consider it in some more detail. HSK calculate the size of the snowball by minimizing the thermodynamic potential




Figure 5.9: Ion mobility as a function of temperature for the simulated systems with no applied field (left panel) and with an electric field of 10^{11} V/cm.

ΔQ . This quantity is given by [8]

$$\Delta Q = (P_m - P_l) \left(\frac{n_s - n_l}{n_l} \right) \frac{4\pi}{3} R^3 + 4\pi\sigma_{ls}R^2 - \left(\frac{\kappa_l - \kappa_s}{8\pi\epsilon_0^2\kappa_s\kappa_l} \right) \left(\frac{q^2}{R} \right) \quad (5.6)$$

where R is the radius; P_m is the melting pressure; P_l is the pressure in the liquid; $n_{l,s}$ are the respective number densities in the liquid and solid and $\kappa_{l,s}$ their dielectric constants; q the ionic charge and σ_{ls} is the surface tension in the liquid-solid interface. The three terms on the right hand side of this equation respectively describe the work done in compressing the liquid into a solid, the surface energy required to form the liquid-solid interface, and the electrostatic work done on the liquid by the ion's field. When this function is differentiated with respect to R and solved for $d(\Delta Q)/dR = 0$ we find that R increases rapidly with decreasing temperature, resulting in $R \sim 25$ Angstroms at the triple point, where $(P_m - P_l) = 0$.

The problems with this analysis stem from the fact that the thermodynamic limit is implicitly assumed, and in the immediate vicinity of the ion, this is not the case. The electrostatic term in Equation 5.6 is actually correct at distances relatively far from the ion. At very close range, the ion-liquid interaction becomes repulsive, and because it is electrostatic in nature its effect is presumably slightly different depending upon the density in its immediate environment.

In any case, these arguments are sort of moot because it is difficult to make sense of concepts like pressure and density when dealing with distance scales comparable to atomic sizes and regions containing only a dozen-or-so particles. Instead an atomistic analysis is probably more appropriate.

S. Patil [18] has solved the problem of determining the maximum number N of rare gas atoms that can form a bound cluster around an ion with a given interaction potential. For the singly charged alkaline earth ions, we estimate, based on Patil's analysis, that N is of order 15 atoms, much less than the 1013 triple-point atoms estimated by HSK. In the computer simulation, the number of atoms in the cluster was found on average be less than N because interactions with the liquid cause atoms to be knocked out of the cluster after having been there for about 1 ps. In a sense, from the ion's point of view, it moves through an extremely dense, viscous fluid, experiencing an enhanced collision rate with liquid atoms that are slow to disengage. The DRM expression for the mobility μ is expressed as

$$\mu = \frac{q}{\zeta + \frac{8}{3}N_0\sqrt{kT}\sigma_i^2g(\sigma_i)\sqrt{\frac{2\pi m_i m}{m_i+m}}} \quad (5.7)$$

where ζ is the contribution to the friction constant from the long-range attraction between the ion and the liquid atoms, N_0 is the (bulk) number density of the liquid, σ_i is hard-core diameter in the ion-liquid interaction, $g(\sigma_i)$ is the radial distribution function of liquid atoms around the ion evaluated at σ_i , and m_i and m are the respective masses of the ion and liquid atom. For ζ , DRM give the expression

$$\zeta = \frac{1}{3}N_0 \int_{\sigma_i}^{\infty} \nabla^2 V(R)g(R) dR \quad (5.8)$$

where $V(R)$ is the ion-liquid interaction potential and $g(R)$ is, once again, the radial distribution function around the ion.

It is clear by inspection that the DRM formulae are along the right track because they properly take into account the microscopic graininess of the liquid as well as the large variations in density close to the ion. However, because of their unusual shape, the RDFs in present study give $g(\sigma_i) = 0$ which leads to an incorrect determination

of μ . Clearly, the DRM formulae must be re-derived to handle the uniquely shaped RDFs that characterize the monatomic ion-LXe systems. Moreover, a scheme for predicting $g(R)$, given a particular ion and liquid temperature, will also be required. Both problems are quite difficult, so instead of solving them explicitly, we will make some qualitative arguments that will lay some of the groundwork.

Let us consider first the RDF $g(R)$. Because the integrand in Equation 5.8 falls off rapidly with increasing distance from the ion it is clear that ζ is determined mainly by the value of RDF at distances close to the ion. Therefore, ζ can be found by evaluating the number of atoms in the cluster around the ion and approximating $g(R)$ as a narrow peak containing that many atoms. To see how this might be accomplished we write

$$\frac{dn(t)}{dt} = \nu(\Omega_1 - \Omega_2)n \quad (5.9)$$

where $n(t)$ is the number of atoms in the cluster at any given time, ν is the mean collision rate between Xe atoms and the ion and $\Omega_{1,2}$ are rate coefficients that describe, respectively the likelihood of an atom being added to the cluster and the likelihood of it being removed. In general the Ω coefficients will depend on the parameters of the ion-liquid interaction as well as the temperature, and we neglect the fact that the likelihood of adding additional atoms actually decreases as the available real estate around the ion fills up.

It is easy to imagine that Ω_1 might scale with the σ parameter of the interaction potential because as σ increases, the larger a “target” the ion is. Similarly, the deeper the well, the more attractive the interaction and the more likely an atom will (temporarily) stick to the ion. One might expect the converse to be true of Ω_2 . In effect, $\Omega_{1,2}$ are cross sections, and there actually exists a theoretical industry devoted to their evaluation from a given, arbitrary potential [19]. Most often this is done numerically.

The solution to Equation 5.9 is

$$n(t) = b_0 \frac{\Omega_1}{\Omega_2} (1 - e^{-\nu\Omega_2 t}) \quad (5.10)$$

where b_0 is a constant. The steady-state value of $n(t)$ is therefore our number of atoms in the cluster.

We now briefly consider the collisional term in the denominator in Equation 5.7. In the original derivation of this expression the ion-liquid interactions were approximated as instantaneous, hard collisions. This approximation probably needs to be rethought to allow for the temporary trapping of liquid atoms in the ion's vicinity. However, exactly how to do this in a tractable way is unclear and warrants further study.

In closing, despite the problems, the approach taken by DRM appears to be the most attractive. This is because in the DRM model, there are two key ingredients that are consistent with the experimental data, namely a very weak temperature dependence and a non-negligible dependence on ion identity. For this reason it is certainly worth studying how the model can be adapted to the systems that are the subject of the present study.

5.5 Implications for EXO

The EXO laser-tagging scheme hinges on the as-yet unverified assumption that the Ba^{++} ion that forms as a result of a double beta decay will neutralize to form Ba^+ as it comes to equilibrium. However, suppose there are mechanisms whereby the Ba^+ re-ionizes, such as through charge exchange with impurities in the liquid? Or, what if the original assumption about the band structure in the liquid is simply wrong? In simulating the ions with and without an electric field and obtaining ion mobilities that agree between the two cases, we have shown in effect that the Nernst-Einstein relation holds for these systems. If it applies equally real ions as it does to simulated ones, then then it may provide a means by which the Ba^+ charge state may be determined experimentally. This would be done by measuring the rate of diffusion of the ions in LXe and comparing this with the measured ion mobility. (Ideally both measurements would be done on the same apparatus.) The measurements could be repeated at several different temperatures to generate a plot of μ/D versus κkT . It is easy to see, using Equation 5.5, that the charge of the ion could then be calculated by simply taking the inverse slope of this plot.

Such an experiment, though challenging, is probably doable provided one has a good source of easily detectable Ba ions. Such an ion source may be available in ^{137}Cs . About 70% of the time, ^{137}Cs β -decays to an excited state of ^{137}Ba , which then decays by emitting an easily detected 661-keV γ ray [20]. A means of turning a sample of ^{137}Cs into a source of ^{137}Ba ions is now under investigation by other workers [21].

5.6 Conclusion

We have performed a comprehensive set of computer simulations of various singly-charged positive ions in LXe. The results indicate the formation of a dynamic, though stable, cluster of Xe atoms around the ion. Although the simulated ion mobilities have come out too high by a factor of 2, within errors they exhibit the correct temperature dependence. In addition, there is a hint that the simulated ions show the correct trend when comparing the mobilities of different ionic species. The simulation results qualitatively contradict the analytical calculations of Hilt, Schmidt and Khrapak. This seems to hint that an atomistic approach along the lines of the theory of Davis, Rice and Meyer is more accurate.

These statements can be made stronger by performing a new simulation with higher statistics. In addition, it would be interesting to see if the factor of 2 discrepancy can be explained by the absence of the dipole term in the simulated interaction between Xe atoms. It is presently being studied whether a future version of the calculation can feasibly incorporate this additional sophistication.

A possible application of this analysis is the determination of the charge state of Ba^+ in LXe. Work is now underway to test the feasibility of this idea.

5.7 References

- [1] E. Reichl, *Statistical Physics: A Modern Course* (5th Ed.).

- [2] S. Jeng, W.F. Fairbank Jr and M. Miyajima, submitted to Phys. Rev. B.
- [3] H.T. Davis, S.A. Rice and L. Meyer, J. Chem. Phys., **37**, 5, 947 (1962).
- [4] K.R. Atkins, Phys. Rev. **116**, 1339 (1959).
- [5] W.F. Schmidt, K.F. Volykhin and A. Khrapak, J. Electrostat. **47**, 83, (1999).
- [6] L. Palleschi, S. Sacchetta and F.P. Ricci, Mol. Phys, **42**, 4, 961 1981.
- [7] H.T. Davis, L. Meyer and S.A. Rice, J. Chem. Phys, **37**, 7, 1521 1962.
- [8] O. Hilt, W.F. Schmidt and A. Khrapak, IEEE 11th Conference on Conduction and Breakdown in Dielectric Liquids (ICDL '93), pp 111 (1993).
- [9] Phys. Rev. **159**, 98 (1967).
- [10] W.F. Fairbank Jr., *private communication*.
- [11] M.S. Byers *et al.*, J. Chem Phys, **78**, 5, 2796 (1983).
- [12] C. Kirkpatrick and L. Viehland, Chem. Phys. **120** 235 (1988).
- [13] E. Matteoli and G. A. Mansoori, J. Chem. Phys. **103**, 4672 (1995).
- [14] M.P. Allen and D.J. Tildesley, *Computer Simulation of Liquids* (7th Ed.), Oxford University Press, 1994.
- [15] A. Einstein, *Investigations on the Theory of Brownian Movement*, Dover, 1956.
- [16] F. Reif, *Fundamentals of Statistical and Thermal Physics*, McGraw-Hill, 1965.
- [17] G. Ciccotti, *et al.*, J. Stat. Phys. **21**, 1 (1979).
- [18] S. Patil, J. Chem. phys. **94**, 3586 (1991).
- [19] E. A. Mason and Earl W. McDaniel, *Transport Properties of Ions in Gases*, Wiley 1988.
- [20] R. DeVoe, EXO Internal Note #248.

[21] C. Hall and R. DeVoe, *private communication*.

Chapter 6

Ion manipulation

In Chapter 4 it was shown that ions can be easily drifted across short distances in LXe, extracted with a probe, and detected. In this chapter we address the manipulation of ions, i.e. procedures by which individual ions can be captured in the liquid phase and released into the gas phase without neutralization. Releasing the ion is necessary for laser tagging schemes in which the spectroscopy is done in an ion trap. It may be possible to perform spectroscopy on the adsorbed ion, but this technology has never been investigated. By contrast, the physics of individual trapped Ba^+ is well documented in the literature, both under high vacuum and in inert background gases at low pressure [1, 2].

Releasing an ion that has been captured against a material surface means overcoming the binding energy between the ion and the surface, which is usually nontrivial. Part of the strong binding arises from the electromagnetic interaction of the ion with the surface. The relative strength of this interaction varies as the square of the ionic charge divided by the ionic radius. Many ions, of which Ba^+ , Th^+ and Ra^+ are examples, experience additional binding due to the covalent action of unpaired electrons in their outermost orbitals [3]. Taking only the electrostatic binding into account, it is easily shown that removing an ion from a surface with an electric field requires, for an ionic radius of, say, a few angstroms, at least several megavolts per centimeter. Such a high field is very difficult to obtain by conventional means, although it might

be achievable using carefully chosen, exotic geometries [4]. Notwithstanding, this difficulty prompted the investigation of alternative strategies, one of which is presented here.

The ion release technique that will be discussed in this chapter is called the “cryogenic probe”. This is a device in which the ion is made to adhere to a layer of Xe ice and subsequently released by causing the ice layer to rapidly melt or sublime.

6.1 The Cryogenic Probe

In this section we describe the experiment that was done as a very basic proof-of-principle for the cryogenic probe concept. The first goal of this study was to demonstrate how a layer of Xe ice might be grown and destroyed at will on the tip of an electrostatic probe. The second was to see if ions could be collected onto a Xe ice-coated tip and released from the tip by removing the ice. This study does not, nor was it intended to, address how this technique might deliver ions to a trap; this will probably be the domain of future research.

6.2 Experiment

This experiment was performed on the same apparatus as that described in the previous chapter. In this case, however, the probe used (cf. Figure 6.1) is a modified version of the original, consisting of a 16” long, 0.095” diameter stainless steel tube that is welded shut at one end to form the probe tip. Interior to this tubing is an ultra-thin (0.025”) stainless steel tube with one opening placed close to the inside of probe tip. The system is set up so that 100-200 bar of Ar gas can be forced through the ultra-thin inner tubing. This is achieved by means of a 350 bar supply bottle outfitted with a special regulator designed for operation at elevated gas pressures. When the high-pressure Ar reaches the tiny opening of the inner tube, it suffers a rapid and severe adiabatic expansion as it escapes into the inside volume of the outer tube. This causes the Ar gas to undergo the Joule-Thompson effect, which is the process whereby gas atoms exchange thermal energy with their surroundings when

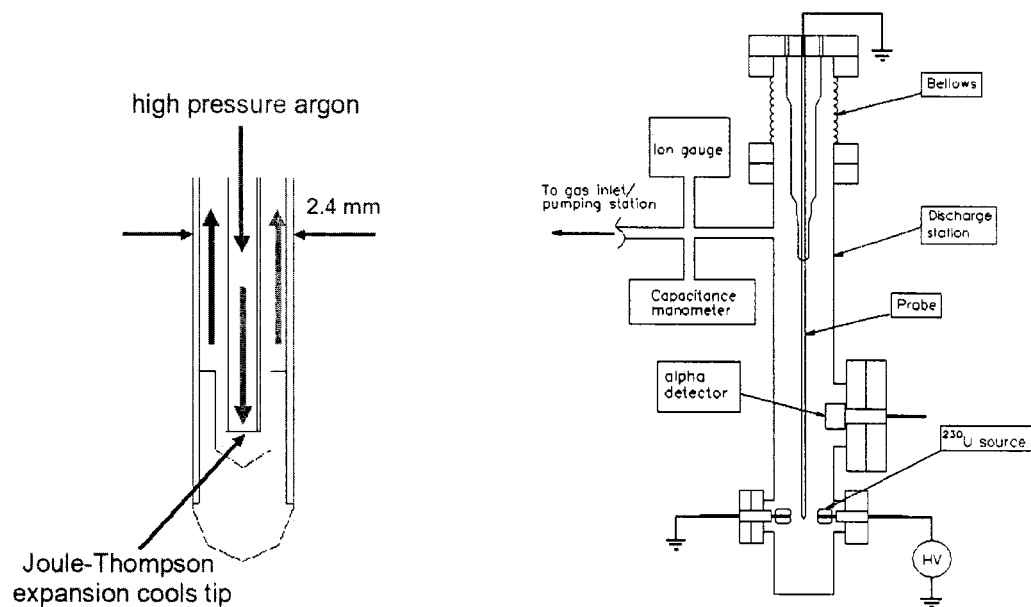


Figure 6.1: Left: Detail showing the internal structure of the cryogenic probe tip. The green arrows indicate the flow direction of the high-pressure Ar gas flow used to cool the tip. The darker arrow shows the direction of the inlet which is at high pressure. Joule-Thompson cooling occurs at the small orifice. Figure taken from Ref. [5]. Right: The Xe cell used for the cryogenic ion release probe experiment. For the tests involving LXe, the cell was insulated by a special vacuum jacket (not shown).

their average inter-atomic spacing is adiabatically increased. In this case, the end result is that massive cooling ensues. The probe tip temperature is thus made to drop ~ 200 degrees in a matter of seconds.

Surrounding the outer 0.125" stainless tube is a thin (0.002") electrically nonconducting Kapton® jacket. The purpose of this jacket is to insulate the probe from a second stainless steel tube placed concentrically around it. This second tube serves as an electrostatic sheath that can be biased independently of the probe itself. Both probe and sheath are biasable to ± 5 kV.

The probe is mounted inside the vacuum-insulated test cell that was described in the previous chapter. As before, the ion source in this experiment is a ^{230}U -plated electrode, and the movable tip can be brought close to the ion source to collect ions and retracted to the counting station for ion detection. In addition, for this experiment, the probe can be retracted to a third position which we refer to as the discharge station. This position is located about 4" above the counting station and is where ion release is attempted.

The experimental procedure is conducted with the cell filled with 0.3 bar of room-temperature, purified Xe gas. Operations are commenced with the probe at the counting station, with the sheath and the tip both grounded. The radioactivity of the probe tip is then measured over the course of a few minutes to obtain an estimate of the backgrounds. During this time, the ^{230}U source electrode is held at -500 V. This prevents large numbers of ions from leaving the source electrode. Next, the probe is lowered to the ^{230}U source location, refrigeration is commenced, and Xe ice forms within several seconds. At the same time, the Xe gas pressure in the chamber drops by about 10 mbar. The ice formation process is monitored through a set of viewports. Once a thin (roughly 2 mm thick) ice layer is visible, the source and the probe sheath voltages are switched to +500 V, thus initiating ion collection onto the ice. During this process, the high-pressure Ar is slowly cycled on and off in order to maintain the ice layer at a roughly constant thickness. This is desirable because if the ions become buried in too thick an ice coating, the α particles they emit are stopped, making the ions undetectable.

After ion collection, the probe is retracted to the counting station, where its radioactivity is measured over the course of 15 s. Following this, the probe is brought to the discharge station where the high pressure Ar flow is stopped and exchanged for about 20 bar of He gas. The inter-atomic interaction parameters of He are such that when it undergoes an adiabatic expansion at these pressures, it rapidly heats up its surroundings. In this way, the temperature of the probe tip is quickly raised back to room temperature, which causes the Xe ice coating to sublime. The probe is then sent back to the counting station and any radioactivity that remains on it is counted for the next 200 s.

6.3 Results

In Figure 6.2 we show the recorded α spectra for ions that are collected onto the bare metal probe and the ice covered probe. These data are a sum over several individual operations, each of which consisted of lowering the probe, gathering ions for 30 min, and counting the probe activity for a 200 s counting period immediately afterward. Each spectrum has been normalized to 663 counts in 200s to match the average rate given in Figure 6.3. By comparison, the background is about 3.0 counts in 200s. In all of the tests we found the activity to be on average about 200 times the detector background which was measured at 3.0 ± 1.2 counts per 200 s. In contrast to the bare probe, the ice-covered probe does not exhibit any clearly discernible peaks in its α spectrum. This is presumably because the ions that are counted are uniformly distributed in the outermost microscopic ice layers, wherein the α particles suffer significant energy losses.

The count rate as a function of time was fit to the sum of two decaying exponentials. The smaller of the two half-lives returned by the fit agrees quite well with the known half-life of ^{222}Ra (38.0 s [6]). The longer half-life has a large error primarily because of the comparatively short counting period (200 s). However, a fit omitting the second, slower decay results in $\chi^2 = 79.07$ (for 10 degrees of freedom), compared to $\chi^2 = 8.91$ for the dual exponential fit. This seems to suggest that the activity collected on the probe is consistent with contributions from ^{222}Ra and ^{226}Th , as expected. This




Figure 6.2: Left: α particle spectrum for the radioactivity adhering to a bare metal probe (red histogram) and to a Xe ice-coated probe (blue histogram). Right: Count rate as a function of time corresponding to the red histogram shown at left. The smooth curve is a fit to the superposition of two decaying exponentials, whose half-lives as returned by the fit are indicated.

seems to indicate that ^{226}Th and ^{222}Ra , are the dominant species in whatever ions are recoiled off of the ^{230}U source and attracted to the probe.

Figure 6.3 shows the results of the ion-release experiment. Plotted in blue is the time histogram of the radioactivity on the (Xe ice-coated) probe after ion collection and the final activity after the ion thawing procedure in which the ice is forced to sublime. The triangular points show the control experiment, i.e., the trial in which the probe is retracted to the discharge station and brought back without destroying the ice. It should be stressed that these data are actually the cumulative results of several trials of each procedure; in other words the demonstrated effect is quite reproducible. Focusing on the portion of the data taken immediately after the probe's return from the discharge station, we see that a total of 290 ± 40 counts was recorded over 200 s for the control experiment, after background subtraction. For the ion-release experiment, we give a background-subtracted limit of < 3.95 counts in 200 s (90 % C.L.). This clearly demonstrates that under these conditions, the Xe ice layer actually prevents the ions from reaching and adhering to the metallic probe surface.

We also note that in our setup, the Xe ice sublimation process takes place in such




Figure 6.3: Time histograms of α counts for the ion release experiment and the null experiment. The solid black line indicates how the probe radioactivity would be expected to fall off in the absence of any ion release effect. It consists of a superposition of the exponential decay curves expected for ^{226}Th and ^{222}Ra whose half lives are, respectively 30.57 min and 38.0 s [6]. The indicated “dead” time is the time during which the probe is at the discharge station. During this time the detector is actually turned off.

a way that the layers of Xe ice closest to the probe surface are the first to warm up and come off. This turns the remaining ice into a hollow shell that simply slips off the probe and falls to the bottom of the cell, apparently taking the ions with it. This process is easily witnessed visually through the cell viewports.

6.4 Discussion

This simple experiment is, of course, only a first step in a long list of tests that must be performed before the cryogenic probe can be qualified for ion extraction and manipulation in a real Xe double beta decay experiment. A chief concern is the problem whereby the ice and the ions fall off the probe because such behavior is not consistent with precise loading of an ion trap. One possible remedy might be to




Figure 6.4: Phase diagram of xenon [7].

warm the probe tip rapidly from the outside, rather than using the current internal Joule-Thompson method.

In addition, this study was unable to address the transfer of Xe ice with ions between liquid and gaseous Xe environments. Unfortunately, it turns out to be impossible, with the current setup, to repeat our procedure in liquid Xe, which, of course, is much more relevant to our ultimate application. This is because the cooling effect of the high-pressure Ar gas flow is far too great, easily overwhelming our ability to maintain a thin ice layer when the probe is immersed in the liquid. Figure 6.4 demonstrates in part why this is so. As shown, at 0.3 bar and 300 K, forming Xe ice requires a temperature excursion of at least 150 K, while at LXe temperatures corresponding to 1.0 bar ambient pressure, only about 4 K is required. Accordingly, the small heat load from the 10 cm^3 of LXe in the cell is easily overcome, and the probe tip quickly freezes almost the entire LXe sample, making it impossible to retract the probe, let alone count the collected ions. These problems can hopefully be remedied by redesigning the probe to provide more modest and localized cooling.

Another problem, illustrated by Figure 6.4, is that at Xe gas pressures approaching 1.0 bar, the Xe ice no longer undergoes a direct solid-gas transition when heated. This leads, in the present setup, to dripping LXe off of the probe which is also undesirable.

This can in principle be remedied by implementing an external heater with sufficient power to cause the Xe to rapidly boil off of the tip.

A still further issue is that the present data cannot rule out the possibility that the ions neutralize once they come into contact with the Xe ice. That they remain ions must be demonstrated in future research; unfortunately, the current setup is not equipped to do this.

6.5 Conclusion

A small but very important first step has been demonstrated in the development of an ion release technique based on the cryogenic probe. It has been shown that a ~ 2 mm thick temporary layer of Xe ice in 0.3 bar room temperature Xe gas can serve as a barrier to positive ions. It therefore seems plausible that a small amount of Xe ice is a possible medium for the transfer of captured Ba^+ ions out of a LXe detector and into an ion trap. Further research is needed to refine the technique and to determine how the ions are affected by their interaction with the ice. In addition, some consideration must be given to the design parameters of an ion trap that can accommodate loading by a cryogenic probe.

Most of the refinements to the cryogenic probe that have been discussed in the previous section are currently in the process of being engineered and tested. In addition, future experiments involving charge state identification and trap loading are now being designed. Once a viable ion capture, transport and release strategy is developed its complete efficiency will have to be characterized. Planning for such work is also underway.

6.6 References

- [1] W. Neuhauser *et al.*, Phys. Rev. Lett. **41**, 233 (1978).
- [2] S. Waldman, PhD thesis, Stanford University (2005), and references therein.
- [3] M.D. Scheer and J. Fine, J. Chem. Phys. **39** 1752 (1963).

- [4] B. Flatt, EXO Collaboration E-log Document #262.
- [5] P. Rowson, EXO Collaboration E-log Document #7.
- [6] R. Firestone *et al.*, *Table of Isotopes*, John Wiley & Sons, Inc. 1996.
- [7] R. Bernabei *et al.*, *Nucl. Inst. Meth. A* **482**, 728 (2002).

Chapter 7

Conclusions

The results presented in this thesis allow one to draw some important conclusions relevant to R&D for large ultra-low background LXe detectors. By way of conclusion, we summarize the main findings and offer a proposal for how to incorporate them into future detector designs.

7.1 Toward the Next Generation

In some sense the main findings presented here confirm what was already known to workers in the liquid rare gas detector community. Much prior work (see, for example, Ref. [1]) has already shown that continuous recirculation is a necessary means of maintaining a high electron mobility in the liquid. The positive ion mobility in LXe had also been studied previously, and though ^{226}Th is a new ion, our result is hardly surprising. Finally, although this is probably the first time a temporary layer of Xe ice has been used as a protective dielectric insulator, doing so did not require the discovery of any fundamentally new physics.

However, what is completely unknown is whether the properties of LXe that were studied allow the creation of a large-scale detector technology that is simultaneously optimized for TPC performance and for single-ion identification by laser-tagging. Hopefully, the studies presented here have provided a better idea of the constraints, if not the precise path to follow to realize such a system. In the following section we

describe a proposed full-scale EXO that incorporates LXe TPC performance together with daughter ion retrieval for tagging.

7.2 Full-Scale EXO with Laser Tagging

So far it has proven impossible to test laser tagging in the 1.0 atm saturated vapor above the LXe because of the operational constraints of ion traps and laser optics. Furthermore, it has been shown [2] that a rough Xe vacuum renders the ion trap extremely unstable. It therefore appears likely that the actual tagging would have to occur in a separate chamber that can be rapidly pumped to about 10^{-6} mbar. We therefore propose that the TPC reside in a volume that is connected to the trap chamber via a rapidly cycling gate valve.

To retrieve an ion from the TPC and place it in the trap, a device similar to the cryogenic probe of Chapter 6 would be used. The probe would be mounted on a set of rails via a mechanism that gives it x -, y - and z - degrees of freedom.

The TPC would have to be constructed so that an electric field of ~ 3 -4 kV/cm was maintained throughout its sensitive volume. 10 tonnes of LXe occupies a volume of just over 3 m³, so the fiducial volume would be segmented into about 12 sub-volumes, each bounded by an anode and a cathode, some of which would be shared between adjacent volumes. Uniformity of the electric field would be achieved by means of ~ 1 cm-spaced field shaping rings.

When a $0\nu\beta\beta$ candidate event occurs, the probe would dip into the TPC, accessing the LXe volume by going in between the field shaping rings to get within ~ 1 cm of the location of the daughter ion. The TPC drift field, which remains energized during this process could be used to help drift the ion to the probe tip. If this is done, however, the probe sheath must be held at a potential above that of the nearest field-shaping rings, which may prove to be nontrivial. Alternatively, the TPC drift field could be cycled every time the probe is dipped.

Once the ion had been frozen onto the tip, the probe would be retracted and moved laterally until it was just above the ion trap chamber. There it would detach from the

$x - y - z$ rail system and attach to a rigid rod that retracts into the trap chamber, and the gate valve would close behind it. This transferal of the probe between rail systems could presumably be done through the use of electromagnets.

Presumably, the lines for power and cooling of the cryogenic probe would be long and flexible and access the probe through the trap chamber. This implies that the cooling system ought to be based on a scheme that does not require high pressure in the flexible line. One option might be ultra-cold, low-pressure helium gas.

If we assume that the ion drifts and is collected on the probe within a second, then the bulk of the time spent to tag a single event will be during the transfer of the probe between the TPC and the trap and during pumping of the trap volume. If motion plus pumping happen within 120 s and performing the spectroscopy takes another 60 s, this implies a “dead” time for the detector of order 3 minutes. This means that the cumulative rate of all the radioactive backgrounds producing activity in the $0\nu\beta\beta$ window cannot exceed 1 count in 3 min, which may or may not be feasible depending upon the types of materials available for constructing the rails, valves, bearings, etc. However, it should be pointed out that this criterion is far, far less stringent than the background requirement for an untagged EXO whose rate is quoted in counts per *year* [3].

Of course, this whole scheme seems woefully complicated. Indeed, it does not seem like the type of system that can be operated repeatably and reliably hundreds of times a day for years on end, mainly because of all the moving parts and extra pumping. Its only attractiveness is that it is utterly consistent with the R&D findings presented here.

An alternative, simpler approach might be to attempt a detection of the ion in the gas phase above the liquid using as yet undeveloped techniques, albeit ones that are fairly well motivated. One idea involves dispensing with the probe entirely; instead the ion is drifted to and collected on a stationary cathode made of up a fine wire grid. Once it is certain that the ion is stuck to the grid, the liquid level is lowered (by simply applying pressure externally to force some liquid out of the exit pipe), exposing the grid to the gas phase. The ion is then desorbed from the grid (as a neutral Ba atom) using a laser-induced desorption technique [4] and subsequently

detected by resonantly ionizing it and using the Xe gas multiplication processes to observe its charge. This scheme has the advantage that all of the complicated moving parts are outside the TPC, there are fewer of them and they would require fewer degrees of freedom.

Assuming that lowering the liquid level and detecting the ion take a combined 10 s, this leaves 170 s to drift the ion (so that the total tagging time matches the 3 min quoted for the $x - y - z$ system). Assuming an ion mobility of $\sim 2.0 \times 10^{-3} \text{ cm}^2 \text{ V}^{-1} \text{ s}^{-1}$ this implies that the drift distance can be no longer than $\sim 20 \text{ cm}$. This in turn means that a 3 m^3 fiducial volume could have a cylindrical shape with a radius of 2.2 m.

Although this option seems a great deal simpler, much R&D is needed to establish the mechanics of the desorption and resonant detection, as these techniques have not been investigated in a high pressure Xe-gas environment. These issues are presently under investigation by other workers.

7.3 References

- [1] G. Carugno *et al.*, Nucl. Instr. and Meth. **A292**, 580 (1990).
- [2] S. Waldman, PhD thesis, Stanford University (2005), and references therein.
- [3] C. Hall, EXO E-log document # 471.
- [4] J.A. Misewich, *et al.*, Phys. Rev. Lett. **68**, 3737 (1992).

Appendix A

Experimental Setup

This section gives additional details about the experimental setup, focusing mainly on the software that was developed for system monitoring/control and data acquisition. Copies of all source code and some executables have been made available on the Web at <http://www.slac.stanford.edu/~kwamba/geant/code/>.

A.1 LabView® Programs

Several LabView® programs were developed and maintained to support this research. We describe the relevant ones here. All of these files can be found on the Windows® machine `sldnt44.win.slac.stanford.edu` (which is at the SLAC EXO lab) in the directory `C:\RECOVERED\oldCdrive\EXO`, or at the website mentioned above.

A.1.1 Coolit3.vi

This virtual instrument (VI) is used to control and monitor the temperatures of the XPM when it is operated with liquid nitrogen (LN_2) cooling. This method is seldom used now that a high-power cryogenic refrigerator is available.

To use this method, a copper cold finger is bolted onto the bottom of the XPM. The XPM is then wrapped in yellow thermally insulating foam leaving only the cold

finger exposed and the whole thing is inserted into a large stainless steel dewar. The mouth of the dewar is then closed off with a specially shaped plug made of solid blue insulating foam. Strategically placed holes in the foam plug allow penetrations for some cables and a LN₂ fill tube.

LN₂ flow from the supply dewar is controlled by an AC-powered solenoid valve which opens to allow LN₂ to flow and closes to cut it off. The LN₂ level inside the XPM dewar is measured by a special custom-made level sensor. This device consists of a 9.5" by 1" pair of 0.125" thick aluminum plates spaced 0.035" apart to form a parallel-plate capacitor. When dipped vertically into the LN₂, the capacitance of the two plates increases linearly with the liquid level because of the large difference in dielectric constant between liquid and gaseous nitrogen.

The capacitance is measured using a simple circuit based on the 555 chip that outputs a voltage proportional to the capacitance across its inputs. This voltage is read out and digitized by a Measurement Computing® PCI-DDA-04/12 ADC board. The VI Coolit3.vi then monitors the digitized output from the capacitance-measuring circuit, interpreting this as the LN₂ fill level.

Users of Coolit3.vi may specify upper and lower LN₂ fill-level setpoints on the front panel. When the fill level drops below the smaller of these two numbers, Coolit3.vi sends the signal for the solenoid valve on the fill line to be opened. The valve is closed when the fill level exceeds the upper setpoint.

In addition to monitoring the LN₂ level, Coolit3.vi also looks at the temperature as measured by a set of T-type thermocouples placed on the outside of the XPM body. These are plugged into a Measurement Computing® PCI-DAS-TC thermocouple board.

Coolit3.vi allows the user to select one of these temperatures as the "control channel" and to specify the desired setpoint. The XPM body temperature is then regulated using a PID control sub-VI provided by Measurement Computing®. This sub-VI takes the control temperature and the setpoint as inputs and outputs the appropriate response, which is bounded by user-specified upper and lower limits. This response is used, via the PCI-DDA-04/12 board, to control a set of heaters which counteract the cooling provided by the LN₂. These and other details are summarized in Figures




Figure A.1: Top half of front panel of the VI Coolit3.vi.

A.1 and A.2. With this system temperature stability of order 0.1 degree are routinely achieved. The main disadvantage of this system, however, is that it requires periodic replenishment of the LN₂ supply whenever it runs dry, or about once every 12-16 hours.

A.1.2 grabbercell.vi

This VI is actually a version of Coolit3.vi (which is the subject of the preceding section) that has been modified to cool the probe cell described in Chapter 4. Everything about the underlying algorithm is the same; the only real difference between the two VIs has to do with the fact that they use different digital and analog channels (for reading out thermocouples, controlling heaters, etc). Figure A.3 shows the main features of this VI.

A.1.3 newHFElowlimit.vi

As with Coolit3.vi, the main purpose of newHFElowlimit.vi is also to control and monitor the body temperature of the XPM. The main difference between the two is that while Coolit3.vi relies on a liquid nitrogen based cooling system, newHFElowlimit.vi makes use of the PolyCold® 550HC series industrial refrigerator. To use the latter VI, the (uninsulated) XPM is placed in a system in which it is attached to the stainless steel lid of a sealable dewar that can hold up to 40 L of the HFE-7000 heat transfer fluid (HFE). This dewar is connected, via a set of (valved) 1/2" stainless steel pipes, to a second, identical dewar which acts as a storage/supply reservoir. Because HFE is highly volatile it is important that no part of its handling system be open (or leaking) to the air; this would cause the entire ~40 L supply to evaporate in a matter of hours.

When the XPM has been properly sealed inside its dewar, the appropriate valves are opened and the HFE is pumped in from the supply dewar using about 20 PSI of argon from a gas bottle to supply the required pressure. Care must be taken not to




Figure A.3: Front panel of the VI grabbercell.vi.

overfill the XPM dewar; ideally, the HFE liquid surface should be just below the top of the upper flange on the XPM. The HFE level can be inspected visually through a viewport on the top flange of the XPM dewar. In addition, there is a set of electronic liquid level sensors inside the XPM dewar that are read out by the newHFElowlimit.vi program.

Cooling of the XPM begins once the newHFElowlimit.vi is started. The algorithm works as follows. First, the temperature reading from one of the thermocouples in contact with the XPM body is compared to a user-specified setpoint. If too high, then the computer waits approximately¹ 10 seconds, and then gives the command to the PolyCold® refrigerator for the solenoid valve on the refrigerant lines to be opened. This causes the refrigerant to circulate through a set of copper pipes inside the XPM dewar in thermal contact with the HFE, cooling the HFE and with it the XPM. When the temperature dips below the setpoint, the computer waits another 10 seconds (again, see Footnote 1) and then gives the command for the refrigerant solenoid valve to be shut. This allows the XPM to slowly warm up again to just above the setpoint, where the process repeats. In this way, a temperature stability of about 2 degrees is maintained.

An additional feature of the VI is that it is programmed not to allow the temperature recorded at the lowermost point in the XPM dewar to fall below a user-specified limit. If this happens all cooling is disabled until this temperature exceeds the limit again. This is done to keep the HFE from freezing in the bottom of the dewar, where it is considerably colder.

Besides the temperature control, newHFElowlimit.vi also provides system monitoring and alarms. It monitors the Xe pressure and sounds an audible alarm if it exceeds 45 psi. The alarm also sounds if the saturated Xe vapor pressure above the liquid in the XPM dips below 12 psi.

Because the cryogenics used in this system are sealed, we are able to leave the XPM

¹In this VI, all times are measured not in seconds, but in “loop iterations,” one loop iteration being the length of time it takes LabView to execute one iteration of the VI’s main while loop. Although the LabView function Wait(ms) is used in an attempt to make this time always equal exactly 1 second, in practice, this time is usually slightly longer and is highly variable, depending on the processor load at any given time. At some point, this unfortunate property of this VI probably ought to be fixed.

running, essentially unattended, indefinitely. To facilitate such running, the system is outfitted with an external control box that makes sure the computer is functioning properly. It does this by accepting a “keep-alive” signal generated by newHFElowlimit.vi via a National Instruments® model 6602 counter-timer board. This signal consist of a 1 kHz square wave lasting 1 s. If the external control box fails to detect this signal for two full minutes, it sounds an audible alarm and sends a command to another unit, a Zetron® SentiDial callbox, whose job is to telephone or page whoever is on duty. The external control box also monitors the PolyCold® refrigerator and will sound the alarm and activate the Zetron® callbox if the refrigerator should shut down. To disable these alarms (for example, when deliberately shutting down the refrigerator) the control box can be switched to “bypass” mode. Finally, it should also be pointed out that the control box is also responsible for converting the LabView valve control command (which comes in form of a high/low TTL logic signal) to a 0-26V supply voltage suitable for actuating the refrigerant solenoid valve. Please see Figure A.4 for additional information.

A.1.4 FLOWMETER.vi

FLOWMETER.vi is responsible for reading out and displaying the Xe gas flow rate as measured by an MKS® type 179A series all-metal mass flow meter. Figure A.5 shows the essential features of the front panel. This VI is meant to be used during Xe recirculation while operating the XPM. It is a passive monitor of the flow rate, displaying both its instantaneous value and a measurement averaged over a specified length of time² The readout is done by measuring an analog voltage whose value is proportional to the measured flow rate with a National Instruments® model 6024E I/O board.

²As with newHFElowlimit.vi this time is in loop iterations which in this case equals approximately 0.1 second and is subject to change depending on processor load conditions.




Figure A.4: Front panel of the VI newHFElowlimit.vi.




Figure A.5: Front panel of the VI FLOWMETER.vi.

A.1.5 mobility.vi

This VI is used to take the data for determining the mobility of ^{226}Th in liquid Xe. The fields on the front panel (Figure A.6) allow the user to enter the desired voltages to be used on the probe, its sheath and the source electrode, along with the parameters required for the high-voltage (HV) switching waveform. In addition, the user may specify the length of time that the radioactivity collected onto the probe (and the background) will be counted and supply a file path in which to store the data. The user also can specify the “boil time” which is the length of time that the computer waits for LXe boiling to stop before commencing HV switching. This is important because if Xe gas bubbles are present while the high voltage is turned on they can break down.

When the VI is executed, the computer begins by counting the background for the designated amount of time. It then lowers the probe into the liquid.

The probe is moved by a set of pneumatic actuators that are controlled by solenoid valves that either interrupt or allow the flow of compressed air. The solenoid valves are controlled by the computer, which changes the valve state by sending a TTL logic signal to a box that converts it to the 24 V supply voltage needed to power the solenoids. The valves are such that they close when powered and are open by default.




Figure A.6: Front panel of the VI mobility.vi.

Once the probe is in the liquid, the computer waits for the specified “boil time,” and then sends the switching waveform to the Trek® high voltage amplifier whose output is connected to the source electrode. Next, the computer raises the probe to the counting station and counts the activity on the probe for the specified time. It then plots the results, a time histogram and a pulseheight histogram of the radioactivity on the probe. To measure the ion mobility, this process is repeated several times while varying the HV switching parameters and the voltage between the probe and the source electrode.

The format of the data that gets written to disk is a two column ASCII file, where the first column contains the time stamp in seconds for each recorded event and the second column contains the pulseheight in volts. The timestamps are assigned such that events with timestamps less than or equal to the value of the background count time are only the ones that were recorded during the initial background measurement.

A.1.6 probemover.vi

This VI allows the user to manually move the probe in the probe cell between the three stations: the source station, which next to the radioactive ion source and is near the bottom of the cell; the counting station, which is where the α -particle counters are located, about halfway to the top of the cell; and the discharge station, which is near the top of the cell. Figure A.7 summarizes these features.

A.1.7 grabberVALVE.vi

This VI (Figure A.8) is responsible for manually controlling the flow of high-pressure argon gas for the purpose of rapidly cooling the tip of the cryogenic probe to form Xe ice. The gas flow is controlled by a solenoid valve specially designed to tolerate the high gas pressure. To close (open) this valve the computer sends a high (low) logic signal to a box that converts it to the 24V (0V) supply voltage required by the valve. To use helium gas (for warming the probe tip), simply switch the feed hose over to a

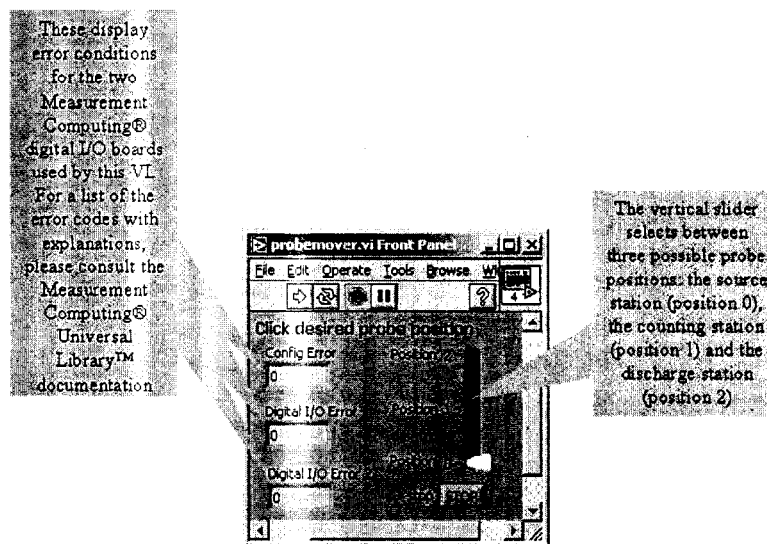


Figure A.7: Front panel of the VI probemover.vi.

helium supply bottle.

A.2 purity.exe

This section is devoted to the software that was developed for the purpose of reading out the LXe purity from the XPM. This software was created using the National Instruments® LabWindows™ development system. This system was chosen because its ANSI C-based compiler conveniently allowed the use of Numerical Recipes algorithms [1] for the complex calculations.

The standalone executable can be found on the machine `slcnt44.slac.stanford.edu` in the directory `C:\Documents and Settings\pearl\Desktop`, and the source code is contained within the LabWindows™ project `scopeMODv8.prj` in the same directory. These files, of course, are also available on the Web.

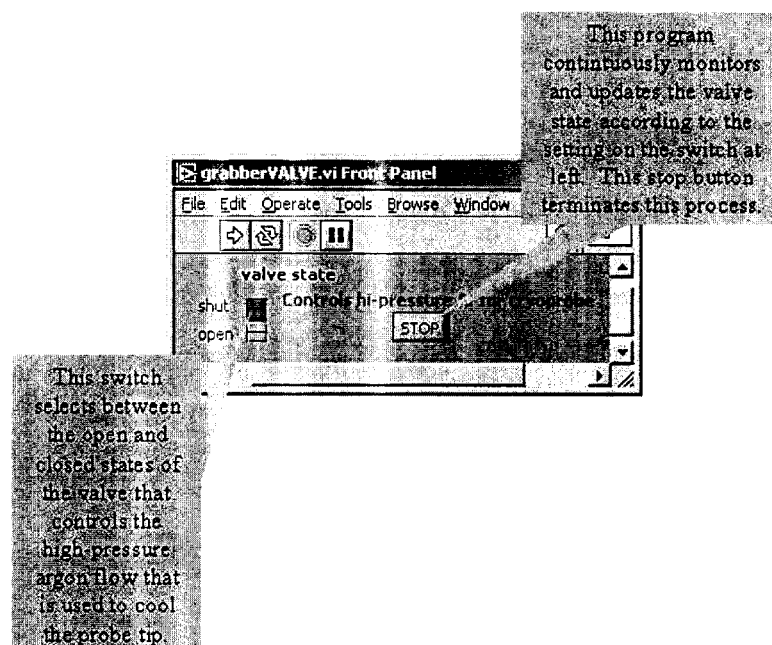


Figure A.8: Front panel of the VI grabberVALVE.vi.

A.2.1 Using purity.exe


Figure A.9 shows an annotated screenshot of purity.exe. To use the program, first enter the desired time between data points (which can be no smaller than 32 seconds) and input a file path for storing the data. By default the data will consist of a single ASCII file whose name will be the entered file path with the letters “int.txt” appended to the end. The file contains a time history of the various measured quantities separated by commas into columns with one measurement per row. Table A.1 gives the detailed format of the “int.txt” file. If the “Save every event” box is checked, purity.exe will also write out an ASCII-encoded version of each recorded XPM waveform. The program names these files by simply appending to the entered path an integer indicating the order in which the waveform was recorded and the extension “.txt”.

If pulse train mode is desired, choose the appropriate pulse parameters and connect the laser external trigger input to the appropriate output on the National Instruments® model 6602 counter-timer board. In this mode, the computer will use the 6602 board to fire the laser Q-switch remotely a specified number of times at a given repetition rate, and the resulting signals will be averaged together.

When ready, click the “Begin” button, and purity.exe will display the time that data taking began and enter the “Running” state. Under these conditions one measurement of the XPM LXe purity will be recorded each time the system clock equals a multiple of the specified time interval.

If desired, the program will also record the laser intensity levels, as determined by two photodiodes placed inside the laser optics bench. To use this option, simply check the “Laser Intensity Data” box. The photodiodes are positioned so that they respectively measure stray reflected 1064 nm fundamental light and its 266 nm fourth-harmonic. The corresponding signals are displayed on the oscilloscope and can also be averaged, downloaded, and shown on the same axes as the XPM signal. In addition, the pulse-heights of the two signals can be computed, scaled appropriately, and displayed as a time history on the same axes as the lifetime-versus-time plot (see Figure A.9).

At the end of the run, data taking is stopped by clicking the “End” button. Running can also be abruptly aborted by clicking “Abort.” This will cause any data that has

Figure A.9: Screenshot of `purity.exe`.

been recorded so far to be discarded. The program may be exited at any time by clicking “Quit” or simply closing the panel window.

A.2.2 Algorithm of purity.exe

The program works by communicating over Ethernet[2] with a Tektronix® model TDS3034 digital phosphor oscilloscope that continuously monitors the XPM waveforms. To record a measurement the program first sends a TTL signal from the computer (via a National Instruments® PCI-MIO-16E-1 multifunction DAQ board) to a device that mechanically closes the shutter at the output of the laser. This prevents photoelectrons from being created at the XPM cathode and allows the XPM waveform baseline to be measured in the absence of a signal. The computer then sends a command via Ethernet to the oscilloscope for it to take 64 triggers and to average the waveforms from each. When this is done, the computer downloads the averaged baseline waveform and stores it in its memory. Following this, the computer opens the laser shutter, commands the oscilloscope again to average 64 waveforms and downloads the result (which now should look like a superposition of the XPM photoelectrically induced signal and the baseline— see Figure A.10). To make a purity measurement, the computer subtracts the stored baseline waveform from the acquired signal+baseline waveform, fits the result to the standard form (see Chapter 3.7), and uses the best-fit parameters to compute the electron lifetime. The waveform, the fit result, and the computed electron lifetime are displayed on various plots (Figure A.9), along with a table of the waveform best-fit parameters (see Table A.1).

Though in principle all 7 parameters in Equation 3.22 could be floated, in practice it has proved to be more efficient to fix the ones that are assumed not to depend significantly on the LXe purity. These are the cathode departure time t_c ; the anode arrival time t_a ; the charge amplifier hold time τ_h ; and the cathode and anode rise times σ_c and σ_a . Except for τ_h , all of these quantities depend only on the electron drift speed (provided there are no significant impurity-related space charge effects). For large signals (or for catastrophically high impurity levels), the accumulation of space charge results in a non-Gaussian spatial profile for the drift electrons, causing




Figure A.10: (a) Example of an unaveraged XPM waveform (blue trace). The green trace in this and the other panels in which it appears is the laser Q-switch synchronous trigger. (b) Typical example of the average of 64 XPM waveforms. (c) Example of the averaged signal seen with the laser shutter closed. (d) The result of taking the difference between XPM waveforms (b) and (c). This baseline-subtracted waveform is fit to equation 3.22.

Quantity	Description	units	Symbol	Column in "int.txt" file
Lifetime	computed best-fit electron lifetime	μs	$\tau = \frac{t_a - t_c}{\log \frac{Q_c}{Q_a} }$	N/A
Cathode	pulse height on rising edge of XPM waveform	mV	Q_c	5
Anode	pulse height on falling edge of XPM waveform	mV	Q_a	6
Tc	time after Q-switch trigger that cathode charge reaches drift region	μs	t_c	1
Tcrise	rise time of rising edge of XPM waveform	μs	σ_c	3
Ta	time after Q-switch trigger that cathode charge reaches anode	μs	t_a	2
Tarise	rise time of falling edge of XPM waveform	μs	σ_a	4
Offst	offset from 0 of the baseline subtracted XPM waveform	mV	Q_{off}	7
Current time	Time elapsed since Jan 1 1900	s	N/A	8
UV	UV (4th harmonic) laser intensity	V	N/A	9
IR	IR (fundamental) laser intensity	V	N/A	10

Table A.1: A table of the quantities stored in the "int.txt" data file. The entries in rows 2-8 in the table list the numbers that characterize the baseline-subtracted XPM waveform in the order that they appear on the purity.exe front panel. These are the 7 fit parameters that correspond to the constants in equation 3.22 in Chapter 3.7.

the rising and falling edges of the waveform to deviate from an error-function shape. Fixing σ_a and σ_c effectively excludes the edges from the fit, sidestepping this complication. The quantity τ_h is measured offline by fitting the step response of the charge amplifier to a decaying exponential. Whenever τ_h and the other fixed parameters are changed, the purity.exe source code must be recompiled with the new values. This would occur for example when replacing a damaged charge amplifier or taking data at a different drift field.

A.2.3 Sources of systematic error on the XPM measurements

After the dynamic baseline subtraction procedure was implemented the remaining systematic errors on the XPM measurements stemmed from two sources. These are uncertainties in the transparencies of the grids, and gradual drifts in the output power of the UV laser.

The grid transparency is controlled by setting the anode and cathode voltages with the Bertan® power supplies and using the resistor network to partition the voltage between the grids. The circuit diagram for the XPM is shown in Chapter 3.7.

As mentioned in Chapter 3.7, electrons will drift across a grid without loss only when the difference in electric field on either side of the grid is large enough. In theory the required field ratio across the grid can be found simply by applying the Bunemann [3] formula. However, in the XPM, which suffers from geometric imperfections that make it non-ideal, the Bunemann formula must be significantly over-satisfied in order to achieve complete transparency. Some of the very early XPM data actually underestimated the electron lifetimes because the electric field ratio was insufficient for full transparency, and this was not taken into account.

In Figure A.11 we show the results of a systematic study of the ratio of anode and cathode charge as a function of the ratio between the electric fields on either side of the cathode grid. The plot shows the ratio of charge leaving the cathode to that collected on the anode as a function of the field ratio. These data were taken when relatively high LXe purity levels were present and there was no test sample in the

XPM.

As shown, the Bunemann formula predicts that the cathode grid becomes fully transparent (i.e. the ratio of anode charge to cathode charge equals 1.0) when the ratio of electric fields is about 2. However, in our non-ideal system, a field ratio of 2 is not enough; a ratio exceeding 5 is more appropriate. For most of the XPM tests, a ratio of 5.2 was used.

As already mentioned the UV laser power also has a systematic effect. There is some qualitative evidence that the reconstructed electron lifetime is correlated with the UV laser power. The exact reason for this is poorly understood and is a current subject of study.




Figure A.11: Ratio of cathode charge to anode charge in the XPM as a function of the ratio of the electric fields on either side of the cathode grid.

A.3 References

- [1] W.H. Press *et al.*, *Numerical Recipes in C*, Cambridge University Press, 1992.
- [2] Tektronix, Inc., *TDS3000 Series Digital Phosphor Oscilloscopes 071-0381-01 Programmer Manual*.
- [3] O. Bunemann *et al.*, Can. J. Res. **A27**, 191 (1949).

Appendix B

Molecular dynamics simulation

The numerical calculations of ion mobilities in LXe were undertaken using the software package `gen_run`. The source code can be found in the files `gen_run.cc` and `gen_run_main.cc` which are now located at the www.slac.stanford.edu/~kwamba web site. The code is in ANSI C++ and compiled using the Linux GNU C and C++ compiler, also known as `g++`. It makes heavy use of the ROOT library functions and classes and seems to run best on the Enterprise Linux systems.

The `gen_run` executable is responsible for running a 500-particle simulation for 500×10^{-15} s time steps to give exactly 1.0 ps of simulated dynamics. It can be invoked either as a fresh (initialization) run from a standard set of initial conditions or as a continuation of an existing run (in order to simulate dynamics for >1 ps). When executing `gen_run` a set of 6 command line arguments must be given. These are as follows. The first is a string specifying whether or not the external electric field is applied and whether this is an initialization run or a continuation run. The syntax of this string is, e.g., “`noiniteoff`” to indicate “no initialization (i.e. this is a continuation run); external field is turned off.” This string must always be in lower-case letters.

The second command line argument specifies the filename to store the data. The format of the file is that of the ROOT TFile class, so it is customary to specify a filename with the suffix “`.root`”. For a continuation run the specified file is assumed to contain data from a previous run to use as initial conditions from which to start the present one. Molecular dynamics data from a run and all of its continuations are

all stored in the same data file. There are 50 cycles of each key in the `.root` file for every picosecond simulated.

The third, fourth, fifth and sixth command line arguments specify the liquid being simulated (e.g. Xe for LXe), the ion being simulated (e.g. Mg^+), the liquid temperature in K, and the liquid density (in number of atoms per m^3). The non-numeric arguments are all case-sensitive.

So, for example, to simulate 100 ps of Mg^+ in LXe at the LXe triple point, starting from standard initial conditions with no electric field, one might first issue the command: `"gen_run initeoff mg_lxe.triple.pt.root Xe Mg+ 161.4 1.36049024e+28"`. This would then be followed by 99 repetitions of the command: `"gen_run noiniteoff mg_lxe.triple.pt.root Xe Mg+ 161.4 1.36049024e+28"`. These commands can be executed sequentially by the SLAC batch system by means of a properly constructed shell script.

Runs with a nonzero electric field always use the endpoint of the most recent zero-field run as their starting point. This means that after each 1 ps nonzero-field run, a zero field continuation run must be performed before running continuing the run at nonzero field.

To analyze the simulation results, the `.root` data file is opened using the ROOT TFile constructor. The keys in the file consist of two `TNtupleD` objects containing the three components of the 500 particles positions and velocities, respectively. In addition, `TGraph` objects displaying the ion's displacement, its 3 velocity components, and the system's total kinetic and potential energies are also stored.