

SABRE: A SEARCH FOR DARK MATTER AND A
TEST OF THE DAMA/LIBRA
ANNUAL-MODULATION RESULT USING
THALLIUM-DOPED SODIUM-IODIDE SCINTILLATION
DETECTORS

EMILY KATHRYN SHIELDS

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
PHYSICS

ADVISER: PROFESSOR FRANK CALAPRICE

September 2015

© Copyright by Emily Kathryn Shields, 2015.

All rights reserved.

Abstract

Ample evidence has been gathered demonstrating that the majority of the mass in the Universe is composed of non-luminous, non-baryonic matter. Though the evidence for dark matter is unassailable, its nature and properties remain unknown. A broad effort has been undertaken by the physics community to detect dark-matter particles through direct-detection techniques. For over a decade, the DAMA/LIBRA experiment has observed a highly significant (9.3σ) modulation in the scintillation event rate in their highly-pure NaI(Tl) detectors, which they use as the basis of a claim for the discovery of dark-matter particles.

However, the dark-matter interpretation of the DAMA/LIBRA modulation remains unverified. While there have been some recent hints of dark matter in the form of a light Weakly-Interacting Massive Particle (WIMP) from the CoGeNT and CDMS-Si experiments, several other experiments, including the LUX and XENON noble-liquid experiments, the KIMS CsI(Tl) experiment, and several bubble chamber experiments, conflict with DAMA/LIBRA assuming a WIMP dark-matter model. However, these experiments use different dark-matter targets and cannot be compared with DAMA/LIBRA in a model-independent way. The uncertainty surrounding the dark-matter model, astrophysical model, and nuclear physics effects makes it necessary for a new NaI(Tl) experiment to directly test the DAMA/LIBRA result.

The Sodium-iodide with Active Background REjection (SABRE) experiment seeks to provide a much-needed model-independent test of the DAMA/LIBRA modulation by developing highly pure crystal detectors with very low radioactivity and deploying them in an active veto detector that can reject key backgrounds in a dark-matter measurement. This work focuses on the efforts put forward by the SABRE collaboration in developing low-background, low-threshold crystal detectors, designing and fabricating a liquid-scintillator veto detector, and simulating the predicted background spectrum for a dark-matter measurement. In addition, recent controversy surrounding the value of an important parameter for direct detection, the nuclear quenching factor, prompted SABRE to perform a measurement

of the quenching factor in sodium. The measurement, its results, and the implications for DAMA/LIBRA and dark matter are also described.

Acknowledgements

Writing a dissertation and earning a PhD is a monumental and humbling undertaking. The hard work and support of a large number of people made this document possible, and I would like to take the time to acknowledge their contributions and express my heartfelt gratitude to each of them.

My advisor, Professor Frank Calaprice, gave me the opportunity to help build SABRE from the ground up. He made a profound effort to ensure that I had the opportunities to delve into every aspect of experimental physics, from building detector parts in the machine shop, to analyzing data and writing simulations, to writing grant proposals and giving talks. He gave his valuable time to teach me the physics behind the experiment so that I would truly understand what I was working on at a deep level. His guidance helped form me as a scientist, and it was my profound joy to help him elevate SABRE from the conceptual phase to a full collaborative experiment. In addition, Frank took an interest in me as a person and helped me to grow in confidence and in my career. For his mentorship, his patience, and his kindness over the last five years, I would like to express my deepest gratitude.

Professor Peter Meyers served as my advisor in the first few years at Princeton, and helped me to build necessary skills I would need in my work on the SABRE project. He cultivated my interests in teaching, allowing me to take classes in the TeacherPrep program and to teach a course for the Princeton Dark Matter Summer School. He also provided thoughtful and detailed comments on this work. Peter always made an effort to help when I came to him with a question or concern, and I would like to thank him for his guidance. I would also like to thank my thesis committee members, who have taken the time to strengthen my work and listen to my defense. I would like to especially acknowledge Mariangela Lisanti, who provided helpful resources and comments on the review of the literature in this work.

I would be remiss if I did not acknowledge the crucial efforts of my colleagues in SABRE. SABRE is a small-team effort, and each member of the group has had a strong influence on the development of each component of the experiment. This thesis is, in many ways, the

result of all of our work. In addition, my colleagues in SABRE have been a joy to work with, and I could not ask for a better team of smart, interesting, and hard-working people.

I would like to especially thank Drs. Jingke Xu and Francis Froborg, who have made immense contributions to the project, but have also taken the time to mentor me as a scientist. Jingke has always given me helpful feedback on my work as well as resources and opportunities to learn new physics and new techniques. He has challenged me to become more careful, critical, and questioning as a researcher, and I believe I have benefitted greatly from our many conversations about science. I cannot express how grateful I am for his mentorship. Francis has made a tremendous contribution to the SABRE experiment and has taken on a large portion of the effort since she joined the experiment. The success of SABRE is in part due to her hard work and expertise. She has also always taken the time to provide guidance and feedback on my presentations and my writing, and I am very grateful for her advice.

I would like to also thank Burkhan Suerfu and Josh Pughe-Sanford for their contributions to the experiment and for always being a pleasure to work with. I would also like to thank Allan Nelson for his technical expertise and contributions to the experiment, as well as everything he taught me about machining and design. Bob Parcells, Bill Sands, and Paolo Lombardi provided helpful feedback and guidance as I worked on various design projects for SABRE, and furthermore taught me a number of engineering principles. Steve Lowe, Bill Dix, and Glenn Atkinson performed a large amount of sophisticated machining work for the SABRE experiment. Bill and Glenn provided a lot of feedback on our designs and fabricated a number of challenging pieces. Steve Lowe was always helpful in showing me how to machine the simpler parts myself, and taught me a lot about machining.

I would like to acknowledge our industrial partners and partners at national labs who contributed to these projects, especially Josh Tower and the people at RMD, Raj Odedra and the people at Seastar, Mike Sutherland and Patrick Saline at Allegheny Bradford Corporation, and Dr. Eric Hoppe and Cory Overman at PNNL. Without their contributions to

the powders, the crystals, the enclosures, and the liquid-scintillator veto, this work would not have been possible.

I would like to acknowledge the new members of SABRE, including our partners in Italy and Australia. I am pleased to know that after I have moved on from SABRE that the effort to test DAMA/LIBRA and to perform a dark-matter measurement with its own merits will continue stronger than when I joined. Our new collaborators promise a bright future for the SABRE experiment.

A number of people made important contributions to the sodium nuclear-recoil quenching measurement at Notre-Dame. Dr. Jingke Xu, Shawn Westerdale, and I worked closely on the experiment, but we received help from many others, especially Dr. Henning O. Back¹, Dr. Wanpeng Tan, Dr. Charles Cao, and Dr. Hugh Lippincott. Henning provided the electronics and helped during the setup of the measurement. Wanpeng was a generous host at Notre Dame, helping us with many aspects of the experiment. Charles's work on the SCENE measurement of quenching effects in liquid argon formed the inspiration for the measurement, and his help was indispensable in planning our own. Hugh provided important advice and input during the setup and the writing stages. I would also like to acknowledge the shifters who helped us run the proton accelerator, especially those who volunteered their time to run the night shift.

I have had the pleasure of working with a number of talented graduate students through the last five years, but I would especially like to acknowledge Shawn Westerdale. I count myself extremely fortunate to have worked with him on a number of projects, including the bell-jar prototype neutron veto and the University of Notre Dame quenching measurement. I especially enjoyed the insane task of wall-papering a giant sphere with him, and I consider the lining of the neutron veto one of my most rewarding achievements at Princeton. Outside of work, I enjoyed speculating about the astrophysics of Westeros and planning The Game together.

¹Full title: His Holiness, Rev. Dr. Henning O. Back, BMF

To my current and former colleagues on DarkSide: Before working on SABRE, I was a member of DarkSide and had the pleasure of working with a number of intelligent and interesting people, many of whom I am privileged to call friends. I would like to especially thank (now Dr.!) Jason Brodsky, who constantly helped me battle ROOT and Geant, and who has also been a supportive friend. I would also like to acknowledge Dr. Alex Wright, who not only helped me with a number of scientific problems and got me started on the optical Monte Carlo, but was also always infectiously cheerful, making even the most difficult research problems fun to work on. I would like to thank Professor Christiano Galbiati, who offered me a position in the group as an incoming first year, while I was still deciding what kind of physics to do in graduate school.

I would also like to acknowledge other graduate students, post-docs, researchers, and undergrads with whom I have had the pleasure to work and with whom I have enjoyed scintillating conversations², especially Christian Aurup, Charles Cao, Alvaro Chavarria, Augusto Goretti, Karla Guardado, Guangyong Koh, Aldo Ianni, Andrea Ianni, Ben Loer, Paolo Lombardi, Colin Lualdi, Jeff Martoff, Tanaz Mohayai, Pablo Mosteiro, Maria Okounkova, Hao Qian, Biagio Rossi, Brooke Russell, Jackson Ruzzo, Richard Saldanha, Chris Stanford, Will Taylor, Masayuki Wada, Susan Walker, Xin Xiang, Adam Zec, and Chengliang Zhu. I wish I could expound on the individual ways in which I am grateful to all of them, but this thesis is long enough! So, I will settle for saying that these people have provided a profoundly enjoyable community dedicated to doing groundbreaking physics, and I am so glad I had the opportunity to be a part of it.

Professors Lyman Page and Herman Verlinde have both supported me as the Physics Department chair and Director of Graduate Studies (respectively) by helping me move toward graduation. Lyman's email to me welcoming me into the program when I was still considering different graduate schools made me feel welcome in this community. They have both taken the time to meet with me regularly, both formally and informally, to make sure

²Pun most certainly intended.

I was on track to graduate. Their support in helping me found the Women in Physics group and developing the program helped me to create a support network for women physicists that will hopefully be a mainstay in this department for years to come.

The unsung heroes of SABRE are the administrative staff that help us with the day-to-day operations of SABRE. I am truly grateful for the hard work that they do to support our scientific mission. In addition, many of them have supported me personally and I consider them to be my friends. I have enjoyed their unique perspectives and getting to know them all. I would like to especially thank Helen Ju, who has supported our group and provides fresh flowers to all the graduate students from her garden in the spring, boosting our morale. I would like to thank Courtney Kohut, who was my greatest support and confidant when I struggled with starting a family and was my greatest cheerleader when I succeeded, who never tired of telling me which account an order should go on, and who has always gone the extra mile during the proposal crunch. I would like to thank Julianne Jones, whose spirit, imagination, and enthusiastic disposition has always inspired me. Geoff Gettlefinger worked hard to meet our laboratory-space needs, including helping us to find unconventional locations for our experiments. I went to Regina Savage for absolutely everything, and she always had an answer for me. Angela Lewis made beautiful posters for the physics recital every year. I would like to thank Ted Lewis, who always had something interesting to talk about when I would come down to the stock room for batteries or some unreasonably-sized hunk of metal. Jessica Heslin and Barbara Mooring provided support with the official steps needed to obtain my degrees. I sing the praises of Barbara Grunweg and Lauren Callahan, who would move heaven and earth to get an order out on time for us. I would like to thank Darryl Johnson, who not only took care of all of our shipping in a timely manner, but was nice enough to let me work in his cramped trailer one time for four hours while I waited for my engagement ring to be delivered. All of this, of course, is on top of the hard work each of them do each day to support our scientific needs.

Several people helped me put my graduate studies in context and helped me determine my future career path. During my graduate studies, I pursued a better understanding of teaching principles through the Princeton Program for Teacher Preparation. I am extremely grateful to my teachers and advisors at Princeton TeacherPrep for making me aware of the importance of communicating science to the general public, supporting me in finding my vocation, and helping me understand how people learn. I would especially like to thank my advisor, Dr. Anne Catena, and my professors, Dr. Chris Campisano, Dr. Todd Kent, and the late Dr. Gita Wilder, the last of whom I am also proud to call family. For sharing their devotion to teaching and learning with me, I would like to thank them. Kathleen Mannheimer and Amy Pszczolkowski at the career center were integral in helping me find my post-graduate path. They patiently helped me define my goals and apply my experiences to pursue the next challenge. Thanks to their guidance, I now have an amazing career ahead of me.

Since the Princeton Women-in-Physics group was founded, it has grown to become a strong support and advocacy group for women in our department as well as across campus. It gives me great joy to see this group take off so successfully. I have greatly enjoyed getting to know you all, especially those who regularly attend our meetings, including Stevie Bergman, Tejal Bhamre, Meredith Frey, Anne Gambrel, Emily Grace, Felicity Hills, Patty Ho, Lauren McGough, Christine Pappas, Mallika Randeria, Sara Simon, Katerina Visnjic, and Sophie Zhang, Professors Suzanne Staggs and Mariangela Lisanti, and the men who have attended our ally meetings. I would especially like to thank Suzanne, to whom I went for advice on multiple occasions about integrating my family plans with my graduate studies. I would also like to acknowledge the Women's Leadership Council, including Chaevia Clendinen, Genevieve Plant, Colleen Richardson, Jennifer Schieltz, Sara Simon, Angelina Sylvain, and Sophie Zhang, who I have enjoyed working with to improve the climate for graduate women in STEM across campus. I would also like to thank the administrative staff who have supported us in this important mission, including Cheri Burgess, Jacqueline Deitch-Stackhouse, Jennifer

DeLoessy, Shawn Maxam, Christine Murphy, Karen Jackson Weaver, Dale Trevino, Deborah Prentice, Mary Baum, Michelle Minter, Aly Kassam-Remtulla, and Sanjeev Kulkarni. I am extremely grateful to these people in supporting my efforts to create a more equitable environment for scientists of all identities.

Graduate school requires extreme focus, and sadly often demands the exclusion of other passions. Because of this, I would like to thank the people who helped me to keep music a part of my life. I would like to thank Laurel Lerner, who began the tradition of the Princeton Physics Recital 27 years ago and Dr. Pierre Piroué, who personally asked me to help him teach the Physics of Music course here at Princeton. I would like to thank everyone who participated in the recital while I was directing it. I would like to thank Mike Fuccillo and the other members of my husband's cohort in the Chemistry Department for allowing me to participate in writing and producing music for the Frickmas skits in 2012. Having a creative outlet went a long way to helping me maintain perspective while immersing myself in my physics studies.

In the same vein, I would like to thank the Doyle lab for making me feel welcome in their group, inviting me to their game nights and dinners, and for helping me practice interviewing for my next position.

I must, of course, mention the Bent Spoon, who has been there for me from the start. Together, we've celebrated my accomplishments, and commiserated when the going got rough. When this thesis gets published, I'll be sure to have a pint in their honor.

I could not have gone on this road without the early influence of the people who cultivated in me a love of science and mathematics. I would like to thank my undergraduate advisor, Dr. Thomas Bernatowicz, for introducing me to experimental research and for supporting me both as an undergraduate and today. I would also like to thank my high-school physics teacher, Mr. Philip Schmidt, and my high-school math teacher, Dr. Darlene Nelson, who inspired me to pursue science as a career.

I would like to thank Meredith Thomas, who was there for me during the greatest challenge I've ever faced. She helped teach me that through sheer determination I can do and endure anything, which was extremely helpful when writing this thesis.

The support of friends and family has been indispensable in my studies. I would especially like to thank Vedika Khemani; our battle with the preliminary exam was the start of a beautiful friendship that has meant a great deal to me. I would also like to thank Bonnie and Daniel Parks for the weekend gaming getaways and babysitting. It has been wonderful to have a home-away-from-home right here in New Jersey!

I was very blessed to have my sister, Torie Lebsack, in Princeton for my first few years here. Since she is no stranger to the Princeton thesis, her support and good humor during this last year also meant a great deal to me. She has definitely earned the title of Favorite Sibling, and I hope that we can someday in the future share our Saturdays again. I owe so much to my parents' continued love and support.

I would like to thank Jeanne and Todd Lebsack for cheering me on as I worked toward my doctorate. Starting the journey of parenthood myself has put their sacrifices and devotion in perspective, and I am at once humbled and extremely grateful for their choice to support and encourage my curiosity about the world and my development as a scientist. Their consistent presence in my life, their advice, and their belief in me have helped me achieve more than I thought possible. For their unwavering support, I thank them from the bottom of my heart.

To my son, Damian: You have inspired and motivated me to work harder than I ever have; you have shown me how much I am capable of achieving. You have been with me throughout this doctoral work, coming with me to Italy and the University of Notre Dame (see Figure 6.18), and it has been a joy to share my love of science with you. You are my inspiration, and I love you very much.

To my husband, Jason: It was here at Princeton that we got married and started a family, and we supported each other through the toughest of challenges. Now that we are coming out the other side, I am more grateful than ever for your steadfast support. Thank you for

taking over the chores while I was writing, thank you for listening when I struggled with the challenges of the graduate program, and thank you for helping me carve out time for fun. We could not have done this without each other. For being my greatest partner, supporter, and friend, I thank you. Words cannot express how much you mean to me, but fortunately I think you already know. I love you, and I can't wait to meet our next adventure together.

Some of this work had been published previously. Some of the material in Chapters 6, 7, and 8 came from [1], a paper authored primarily by Dr. Jingke Xu and myself, with contributions from others. Much of the material in Chapters 3, 4, and 5 was previously published in the Procedia for the 2013 TAUP Conference [2]. Though the SABRE-in-DarkSide Insertion System (SIDIS) design was a team effort to which I made a contribution, the final design was largely put together by Dr. Francis Froborg and described by her in [3]. The description of SIDIS in this work is largely based on that document.

The SABRE program has been supported by NSF Grants PHY-0957083, PHY-1103987, and PHY-1242625. The Na nuclear-recoil quenching measurement was additionally supported by NSF grant PHY-1419765.

To my husband, Jason.

Contents

Abstract	iv
Acknowledgements	vi
List of Tables	xvii
List of Figures	xviii
1 Introduction to Dark Matter	1
1.1 The case for dark matter	1
1.1.1 Missing mass and the motion of galaxies and stars	1
1.1.2 The Bullet cluster	2
1.1.3 Measurement of cosmological parameters	2
1.2 Dark matter candidates	7
1.2.1 Characteristics of dark matter	7
1.2.2 Baryonic dark matter	7
1.2.3 Neutrinos	8
1.2.4 Axions	9
1.2.5 Weakly Interacting Massive Particles (WIMPs)	10
1.3 Technical aspects of dark-matter searches	12
1.3.1 The WIMP scattering event energy	14
1.3.2 The WIMP event rate	15
1.3.3 Types of direct detection	17
1.3.4 The sensitivity of direct-detection experiments	20

1.3.5	The state of direct-detection dark-matter searches	24
2	Annual Modulation Searches and DAMA/LIBRA	28
2.1	The WIMP annual modulation	28
2.1.1	Technical aspects of annual-modulation searches	30
2.2	The DAMA and DAMA/LIBRA experiments	36
2.2.1	The DAMA/LIBRA apparatus	37
2.2.2	Radioactivity measurements of the detector components	47
2.2.3	Processing the data	51
2.2.4	The dark-matter claim	55
2.3	Controversy over the DAMA/LIBRA result	56
2.3.1	The energy spectrum	56
2.3.2	Simulating the DAMA/LIBRA background spectrum	59
2.3.3	The Phase	75
2.3.4	The modulation amplitude	75
2.3.5	The nuclear quenching factor and the channeling effect	76
2.3.6	Conflicting WIMP measurements	76
2.3.7	Possible explanations for the DAMA/LIBRA modulation and potential resolutions to the conflict	86
3	SABRE: An Independent Test of DAMA/LIBRA	94
3.1	The importance of an independent test of DAMA/LIBRA	95
3.2	Qualities of an independent test	95
3.2.1	Lowering and understanding backgrounds	96
3.2.2	Lowering the energy threshold	98
3.2.3	A new location	98
3.3	Sodium-iodide with Active Background REjection: The SABRE experiment	99
3.3.1	Low-background powders and crystals	100

3.3.2	Low-background, low-noise, high-quantum-efficiency PMTs	101
3.3.3	Low background enclosures with high light yield	102
3.3.4	Active background rejection with a liquid-scintillator veto detector . .	103
3.3.5	A deep underground site with potential for relocation	104
3.3.6	Understanding the detector response and backgrounds	104
3.3.7	Measurement of experimental parameters of relevance to direct- detection searches	105
4	The SABRE Crystal Detectors	106
4.1	Powder	106
4.1.1	The need for high-purity powder	107
4.1.2	Development of high-purity powder	108
4.1.3	Purification of powder	109
4.2	NaI(Tl) Crystals	109
4.2.1	Crystal growth	110
4.2.2	RMD facilities	111
4.2.3	Testing the Vertical Bridgman method	112
4.2.4	Growing large crystals with the Vertical Bridgman method	118
4.2.5	Future directions	120
4.3	The photomultiplier tubes	120
4.3.1	Goals	120
4.3.2	Progress	123
4.4	Enclosure	136
4.4.1	Encasement	137
4.4.2	Support structure	140
4.4.3	Reflectors and coupling	143

5	The Active Liquid Scintillator Veto and Simulations of SABRE	146
5.1	Active background rejection	146
5.1.1	The veto concept	147
5.1.2	The need for active background rejection	147
5.1.3	The DarkSide-50 neutron-veto detector	148
5.1.4	A dedicated veto detector	150
5.2	SABRE in DarkSide	150
5.2.1	The DarkSide experiment	150
5.2.2	SIDIS Requirements	151
5.2.3	The design	153
5.2.4	Progress	162
5.3	A portable veto detector	163
5.3.1	Requirements and considerations	163
5.3.2	The veto chamber design	166
5.3.3	The veto chamber construction	173
5.3.4	Electronics scheme for the liquid-scintillator veto	176
5.4	Simulations of the SABRE experiment	176
5.4.1	Backgrounds in the NaI(Tl) SABRE detectors	178
5.4.2	Light yield of the SABRE veto detector	219
5.5	The future of SABRE	224
6	Measurement of the Na Nuclear-Recoil Quenching Factor: The Experiment	226
6.1	Introduction	226
6.1.1	Prior publication of this work	226
6.1.2	The role of the ionization quenching factor in dark-matter experiments	227
6.1.3	Previous measurements of the quenching factors in NaI(Tl)	228
6.1.4	The SCENE approach	229

6.2	Preparation for the measurement	232
6.2.1	Choosing the recoil energies	232
6.2.2	Choosing the proton energy	232
6.2.3	Choosing the LiF thickness	234
6.2.4	Choosing the detector positions	236
6.3	Experimental setup	243
6.3.1	The NaI(Tl) detector	243
6.3.2	The proton beam and the LiF target	245
6.3.3	The neutron detectors	248
6.3.4	Placing the detectors	248
6.3.5	The electronics and data acquisition	260
6.3.6	Measurement Summary	265

7 Measurement of the Na Nuclear-Recoil Quenching Factor: Analysis and Results **270**

7.1	The data acquisition	270
7.1.1	Features of the waveforms	271
7.1.2	The baseline algorithm	271
7.1.3	Finding pulses	273
7.1.4	Creating event objects: the TOF module	273
7.2	Creating the energy spectrum	274
7.2.1	Finding the single photoelectron	274
7.2.2	Selecting events	276
7.2.3	Making the light-yield correction	282
7.2.4	The nuclear recoil energy spectra	285
7.3	Simulating the energy spectra	285
7.3.1	Generating the incoming neutron energy and angular spectra	287
7.3.2	The simulation	289

7.4	Evaluating the quenching factors	290
7.4.1	Standard method	291
7.4.2	Evaluating the fit for low-energy recoils	294
7.4.3	Error analysis	299
7.4.4	Other considerations	301
7.5	Iodine recoils	302
7.6	Channeling effect	302
7.7	Relation to other measurements	303
8	Measurement of the Na Nuclear-Recoil Quenching Factor: Implications for WIMP Dark Matter	305
8.1	Calculation of the WIMP dark-matter event rate	306
8.2	The Standard Halo Model	308
8.2.1	The velocity distribution	308
8.2.2	The effective cross section	310
8.3	Results for the Standard Halo Model	313
8.4	Conclusions	320
A	Principles of Scintillation	322
A.1	Radiation in matter	322
A.2	Organic liquid scintillators	323
A.2.1	Scintillation mechanism	323
A.2.2	Pulse-shape discrimination	324
A.3	Inorganic scintillators like NaI(Tl)	327
A.3.1	Scintillation mechanism	327
A.3.2	Ionization quenching	327
A.3.3	Pulse-shape discrimination	328
B	List of Abbreviations	331

C List of Symbols 336

Bibliography 342

List of Tables

1.1	SUSY dark-matter candidates	11
2.1	DAMA/LIBRA annual modulation amplitude	40
2.2	DAMA reported effects on the modulation amplitude by various environmental variables	46
2.3	Reported backgrounds in DAMA/LIBRA	47
2.4	Contribution to DAMA/LIBRA annual modulation from neutrons, as reported by DAMA	50
2.5	Simulation fit to verify DAMA radioactivity values	66
4.1	Impurity levels in the Sigma Aldrich NaI AstroGrade powder	109
4.2	Results of the small crucible tests at RMD, Inc.	117
4.3	Results of a blank test performed at RMD, Inc.	118
4.4	Radioactivity of the 3" Hamamatsu R11065-10-SEL PMT	123
4.5	Radioactivity of the high-purity alumina powder and high-purity ceramic	129
4.6	Runs for the LNGS dynode-afterglow study	133
5.1	Radioactivity of the stainless steel used for the SABRE veto vessel	174
5.2	Structure of the Xsimu simulation output	179
5.3	Expected background rates due to the crystals	190
5.4	Expected background rates due to the 3" PMTs	197
5.5	Expected background rates due to the enclosures	204

5.6	Expected background rate due to the veto PMTs	206
5.7	Expected background rates due to the veto-detector steel	209
5.8	LNGS rock and concrete composition	211
5.9	LNGS rock and concrete radioactive impurity concentrations.	212
5.10	Expected external γ background rates with no shielding	214
5.11	Total expected backgrounds for the SABRE detector	218
6.1	Calculated event rates and recoil-energy spread for various detector positions	241
6.2	Detector information for quenching measurement	243
6.3	Marker locations for the first position configuration	251
6.4	Final detector positions for Na quenching measurement	259
6.5	Initial tests of the neutron detectors	261
6.6	Final tests of the detectors	262
6.7	Run log for the SABRE Na quenching measurement	265
7.1	Time-of-flight cuts for the neutron events	278
7.2	Non-TOF Cuts used for the energy spectra	280
7.3	The ^{23}Na quenching factors from the measurement using the standard fit method	297
7.4	The ^{23}Na quenching factors from the measurement using the peak-matching method for low-energy recoils	298
7.5	Na quenching factors from Dr. Jingke Xu's analysis	299
7.6	Uncertainty in the recoil energy (ΔE_R) as a function of the angular uncertainty.	300
7.7	Quenching Factors adjusted for the non-linearity of the electron-recoil scintil- lation efficiency in NaI(Tl)	302
8.1	Dark-matter best fits to DAMA/LIBRA modulation data	320

List of Figures

1.1	The Bullet Cluster	3
1.2	Temperature fluctuations in the Cosmic Microwave Background (CMB) . . .	4
1.3	Determining baryon density from Big-Bang Nucleosynthesis	6
1.4	Axion exclusion plot	10
1.5	Feynman diagram illustrating the three types of dark-matter detection . . .	13
1.6	The state of direct-detection dark-matter searches as of 2013	25
2.1	Illustration of the annual modulation concept	29
2.2	Illustration of gravitational focusing as it applies to the WIMP annual modulation	31
2.3	^{40}K decay scheme	36
2.4	The DAMA/LIBRA annual modulation	38
2.5	The DAMA/LIBRA modulation amplitude as a function of energy, and the DAMA/LIBRA energy spectrum.	39
2.6	The DAMA/LIBRA apparatus	41
2.7	The DAMA enclosures	44
2.8	Layout of the DAMA/LIBRA experiment	45
2.9	Pulse from electronic noise vs. scintillation signal in DAMA/LIBRA	52
2.10	Electronic noise rejection in DAMA/LIBRA	53
2.11	DAMA/LIBRA's reported efficiency	55
2.12	DAMA/LIBRA energy spectrum for two crystals from 20–80 keV $_{ee}$	59

2.13 DAMA α spectra at high energies	60
2.14 Subtracting a standard WIMP spectrum from the DAMA/LIBRA energy spectrum	61
2.15 Simulated backgrounds in the DAMA/LIBRA crystals from ^{40}K , ^{238}U , ^{232}Th , and ^{87}Rb	67
2.16 Simulated backgrounds from other individual isotopes in the DAMA/LIBRA crystals	68
2.17 Simulated total backgrounds from the DAMA/LIBRA crystals	69
2.18 Simulated backgrounds from individual isotopes in the DAMA/LIBRA fused-silica light-guides	70
2.19 Simulated total backgrounds from the DAMA/LIBRA fused-silica light-guides	70
2.20 Simulated backgrounds from individual isotopes in the DAMA/LIBRA ETL PMTs	71
2.21 Simulated total backgrounds from the DAMA/LIBRA ETL PMTs	71
2.22 Simulated backgrounds from individual isotopes in the DAMA/LIBRA copper enclosures	72
2.23 Simulated backgrounds from the DAMA/LIBRA copper enclosures	73
2.24 Simulated DAMA/LIBRA single-hit spectrum using DAMA's reported radioactivity values	73
2.25 Simulated ^{40}K rejection power in DAMA/LIBRA	74
2.26 Simulated DAMA/LIBRA single-hit spectrum using best-fit radioactivity values	74
2.27 CoGeNT observed annual modulation	79
2.28 Rise time cut in the CoGeNT detector	80
2.29 Anomalous events seen by the CDMS-Si detector	81
2.30 Ionization vs. phonon timing in CDMS-Si	82
2.31 Events in the LUX detector after cuts from an 85-day exposure	84
2.32 WIMP parameter-space regions and exclusion curves in the light-WIMP region	86

2.33 Spin-response functions of various dark-matter targets	91
3.1 The active veto concept	97
4.1 Glove-box used for crystal packaging at RMD, Inc.	112
4.2 Furnaces for Vertical Bridgman crystal growth	113
4.3 Dry-room for crystal cutting at RMD, Inc.	114
4.4 RMD NaI(Tl) high-purity crystal NaI002	115
4.5 Optical properties of RMD crystal NaI002	115
4.6 Crystals grown in the small crucible purity tests at RMD, Inc.	117
4.7 Standard purity medium crystal grown with the Vertical Bridgman method .	119
4.8 Radioluminescence of the ceramic used in PMT feedthroughs	123
4.9 DarkSide-50 –HV base design	125
4.10 SABRE 3” PMT, +HV, one-cable base design	126
4.11 SABRE PMT decoupling circuit	127
4.12 Design for the PMT feedthrough plate with high-purity ceramic	128
4.13 Time difference profile in the Princeton dynode-afterglow study	131
4.14 Energy/time profile of events in Princeton dynode-afterglow study	131
4.15 Setup for the LNGS PMT dynode-afterglow measurement	132
4.16 Time profile of dynode-afterglow events at LNGS	133
4.17 Afterglow rate as a function of PMT voltage in the LNGS study	134
4.18 Comparison of afterglow rate with random coincidence in the LNGS study .	134
4.19 Afterglow energies as a function of time difference in the LNGS study	135
4.20 Afterglow light response in the LNGS study	135
4.21 SABRE detector module	137
4.22 Study of the degradation effect of copper on liquid scintillator	139
4.23 Hexagon-shaped enclosure design	141
4.24 Cap design for the SABRE enclosures	142

4.25	Crystal encasement within the SABRE detector module	143
4.26	Lumirror reflectance as a function of wavelength	145
5.1	The DarkSide-50 detector setup	149
5.2	the Sabre-in-DarkSide Insertion System (SIDIS)	155
5.3	The SIDIS moving plate	156
5.4	System for moving the worm drive in SIDIS	157
5.5	Potential design for the interface of SIDIS with the SABRE crystal detector	158
5.6	Termination of the cables in the SIDIS system	159
5.7	Bottom plate of the SIDIS chamber	160
5.8	Support panels to secure the interior SIDIS structure	161
5.9	The SIDIS chamber pieces.	161
5.10	Locations of the viewports for SIDIS	162
5.11	To-scale conceptual drawing of the SABRE portable liquid-scintillator veto vessel	167
5.12	Design for the SABRE liquid-scintillator veto vessel	168
5.13	The mu-metal shield used to protect the 8" PMTs in DarkSide	170
5.14	PMT can design for the SABRE portable veto vessel	171
5.15	Test of the PMT grip design for the portable liquid scintillator vessel	172
5.16	Installation of Lumirror reflector in DarkSide-50	174
5.17	Picture of the SABRE portable veto detector vessel	175
5.18	Electronics scheme for the SABRE detector	177
5.19	Comparison of the 3" PMT background from the upper ^{238}U chain before and after Geant4 fix	181
5.20	Geometry used in the simulations of the SABRE detector backgrounds	183
5.21	Expected background from ^{40}K in the SABRE crystals in a dark-matter measurement	187

5.22	Expected backgrounds due to ^{238}U and ^{232}Th in the crystals in a dark-matter measurement	187
5.23	Background due to individual isotopes in the crystals, other than ^{40}K , ^{238}U , and ^{232}Th in a dark-matter measurement.	188
5.24	Expected background due to the crystals in a dark-matter measurement . . .	189
5.25	Expected background for a K measurement due to the crystals	189
5.26	Expected backgrounds due to the 3" PMTs in a dark-matter measurement, from individual isotopes	193
5.27	Expected backgrounds from the currently available 3" PMTs in a dark-matter measurement	194
5.28	Expected backgrounds from the 3" PMTs with new ceramic feedthroughs in a dark-matter measurement	195
5.29	Expected backgrounds from the 3" PMTs in a K measurement	196
5.30	Expected backgrounds due to the enclosures in a dark-matter measurement, long-lived isotopes	201
5.31	Expected backgrounds due to the enclosures in a dark-matter measurement, cosmogenically-activated isotopes	202
5.32	Expected backgrounds from the enclosures in a dark-matter measurement . .	203
5.33	Expected backgrounds from the enclosures for a measurement of ^{40}K in the crystals.	203
5.34	Expected backgrounds due to the veto PMTs in a dark-matter measurement, individual isotopes	205
5.35	Expected backgrounds due to the veto PMTs in a dark-matter measurement	206
5.36	Expected backgrounds from the veto PMTs for a measurement of ^{40}K in the crystals.	207
5.37	Expected backgrounds due to the liquid-scintillator-vessel steel in a dark-matter measurement from individual isotopes	208

5.38	Expected backgrounds from the veto-detector steel in a dark-matter measurement	209
5.39	Expected backgrounds from the veto-detector steel for a measurement of ^{40}K in the crystals.	210
5.40	γ backgrounds at LNGS	212
5.41	Expected external γ background from the LNGS rock, with no shielding. . .	213
5.42	Expected external γ background from the LNGS rock in K measurement mode, with no shielding.	213
5.43	Total expected SABRE background for a dark-matter measurement	215
5.44	Total expected SABRE backgrounds with no veto, broken down by source .	216
5.45	Total expected SABRE backgrounds with veto, broken down by source . . .	217
5.46	Total expected SABRE background for a K measurement	218
5.47	Reflectance of the reflector Lumirror	221
5.48	Emission spectrum of PPO	222
5.49	Predicted light yield of the SABRE veto vessel	223
5.50	Predicted optical efficiency of the SABRE veto vessel	224
6.1	Previous measurements of the Na quenching factor	229
6.2	Experimental setup for Notre Dame quenching measurement	231
6.3	0° -incidence $^7\text{Li}(p,n)^7\text{Be}$ differential cross-section as a function of proton energy	234
6.4	Combined cross section for neutron elastic scattering with given proton energy	235
6.5	Stopping power for protons in LiF	237
6.6	Simulation to determine LiF target thickness	238
6.7	Crystal enclosure for the Na quenching measurement	246
6.8	Aligning the laser for the Na quenching measurement	250
6.9	Marking detector positions for the Na quenching measurement	250
6.10	Picture of the quenching measurement setup	252
6.11	The NaI(Tl) detector in its mount	253

6.12	The neutron detectors for the Na quenching measurement	254
6.13	Centering the 5.1-cm detectors above the floor marks	255
6.14	Marking positions for the 5.1-cm neutron detectors	256
6.15	Centering the collimator for the Na quenching measurement	260
6.16	Sample spectra from the calibration of the neutron detectors	260
6.17	Na quenching measurement electronics scheme	263
6.18	(Top) The electronics modules used in the quenching measurement. (Bottom) Testing the electronics modules. The NaI(Tl) detector and the some of the neutron detectors can be seen on the table on the left.	264
6.19	Trigger-efficiency measurement setup	269
7.1	Typical waveform for a nuclear-recoil event	272
7.2	The s.p.e. spectrum	275
7.3	Time-of-flight plot	278
7.4	Pulse-shape discrimination in the neutron detectors	280
7.5	Pulse-shape discrimination in the NaI(Tl) detector	281
7.6	First time of flight (TOF1) vs. energy for all detectors.	282
7.7	Second time of flight (TOF2) vs. energy for neutron detector 1	283
7.8	Light-yield calibration fit for Run 1 using the ^{127}I first excited state	284
7.9	Light yield of NaI(Tl) detector vs. time	284
7.10	The energy spectra for the Na quenching measurement, position configuration 1286	
7.11	The energy spectra for the Na quenching measurement, position configuration 2287	
7.12	Incoming neutron energy and angular spread	289
7.13	The quenching-measurement simulation geometry	290
7.14	Single- vs. multiple-scattering energy spectra (simulated)	292
7.15	Monte Carlo fits for the quenching factor for higher-energy recoils	293
7.16	Trigger efficiency curve for the Na quenching measurement	294
7.17	Fits for peak-matching method of fit for low-energy events	296

7.18	Pulse peak amplitude vs. energy spectrum	297
7.19	Non-linear scintillation response in NaI(Tl)	301
7.20	Results of the Na quenching measurement with previously measured values	303
8.1	Nuclear form factors for Na and I	312
8.2	Fit of the new quenching factors	315
8.3	Best fits for DAMA/LIBRA modulation with old and new quenching factors	316
8.4	DAMA/LIBRA spin-independent parameter space regions with old and new quenching values	317
8.5	DAMA/LIBRA spin-dependent parameter-space regions with old and new quenching values	318
8.6	DAMA/LIBRA reported background and maximum dark-matter signal	319
A.1	Excitation scheme for organic scintillators	325
A.2	Stopping power of electrons and protons in toluene	326
A.3	Scheme for the scintillation mechanism in NaI(Tl)	328
A.4	Scintillation efficiency in NaI(Tl) as a function of stopping power	329
A.5	Stopping power, dE/dx , for electrons and Na ions in NaI(Tl)	330

Chapter 1

Introduction to Dark Matter

1.1 The case for dark matter

Undeniable evidence has been accumulated showing that the majority of the mass in our universe is composed of non-luminous, non-baryonic matter. Recent assessments of the universe's composition conclude that this “dark” matter comprises 85% of the mass in the universe [4]. Despite dark matter's dominance, its nature remains unknown. The detection and subsequent characterization of this dark matter is one of the greatest goals in modern physics.

A variety of experimental results have made the case for the dominance of dark matter in our universe and have attempted to quantify it and constrain its properties. This evidence for dark matter stretches back for nearly a century and becomes more and more convincing with time. Some of this evidence is described below.

1.1.1 Missing mass and the motion of galaxies and stars

Non-relativistic celestial bodies follow the laws of Newtonian gravitation. Early hints of dark matter appeared during the observation of the motion of celestial bodies, which did not match expectations. The apparent non-conformity of the motions observed with Newtonian

physics suggested that more mass existed than was visible. In the 1930s, Jan Oort and Fritz Zwicky separately postulated the existence of non-luminous matter to describe the velocity distribution in the local galactic neighborhood [5] and the non-virialized motions of galaxies in the Coma Cluster [6], respectively.

In the 1970s, Vera Rubin, et al. studied galactic rotation curves [7], confirming that either stars were not exhibiting Newtonian gravitational motion, or that there was missing matter that was non-luminous. The velocity distribution of stars in the galactic plane is expected, based on the amount of visible matter, to follow the form $v(r) \propto \sqrt{M(r)/r}$, where $M(r)$ is the total mass contained in a radius r from the Galactic Center. At radii beyond the luminous disk, the velocity would be expected to fall off as $1/\sqrt{r}$; instead, Rubin observed constant v at large radii. This observation suggested a dark matter halo with mass density $\rho_{DM}(r) \propto 1/r^2$ at large radii.

1.1.2 The Bullet cluster

The Bullet cluster (1E0657-558) is a galaxy cluster whose collision with another cluster was observed through gravitational lensing [8]. The apparent difference in the behavior of this astronomical event when observed through gravitational lensing and separate luminosity measurements, as shown in Figure 1.1, suggested the existence of dark matter. Most of the mass was observed to move ballistically, while the luminous matter decelerated. This observation is compelling evidence for dark matter, but can also be explained by alternative theories of gravitation [9].

1.1.3 Measurement of cosmological parameters

More recent evidence for dark matter comes from the measurement of cosmological parameters. This method has produced the most accurate determination of the relic density of dark matter in our universe, Ω_{DM} , to date. Here, Ω_{DM} is the fractional density ρ_{DM}/ρ_{crit} , where $\rho_{crit} \equiv 3H_0^2/8\pi G$ is the critical density needed for the universe to have flat curvature.



Figure 1.1: Images of the Bullet Cluster, courtesy of [10] and [11]. The luminous matter is shown in pink, while the majority of the mass is in the blue region, according to gravitational lensing.

The value of Ω_{DM} is discussed in Section 1.3.2. H_0 is the Hubble constant that denotes the universe's rate of expansion and G is the gravitational constant.

Cosmic Microwave Background (CMB)

The hypothesized presence of cold (non-relativistic) dark matter on cosmological scales implies the existence of acoustic temperature fluctuations in the Cosmic Microwave Background (CMB), the light that first escaped when photons decoupled from matter $\sim 380,000$ years after the Big Bang [9]. The angular power spectrum of these temperature fluctuations can be calculated to high precision given certain assumptions about the laws governing cosmological evolution, and the resultant structure of these oscillations are rich enough that the shape of the spectra can be used to constrain the cosmological parameters of our universe, including the Hubble parameter¹, h , and the baryon and cold-dark-matter densities Ω_b and Ω_{DM} .

The anisotropy of the CMB has been studied with several generations of satellite probes including COBE [12] and WMAP [13, 14]. The Planck collaboration has the most accurate

¹The Hubble parameter, h , and the Hubble constant, H_0 are related to each other by the following equation: $H_0 = 100h \text{ km s}^{-1} \text{ Mpc}^{-1}$.

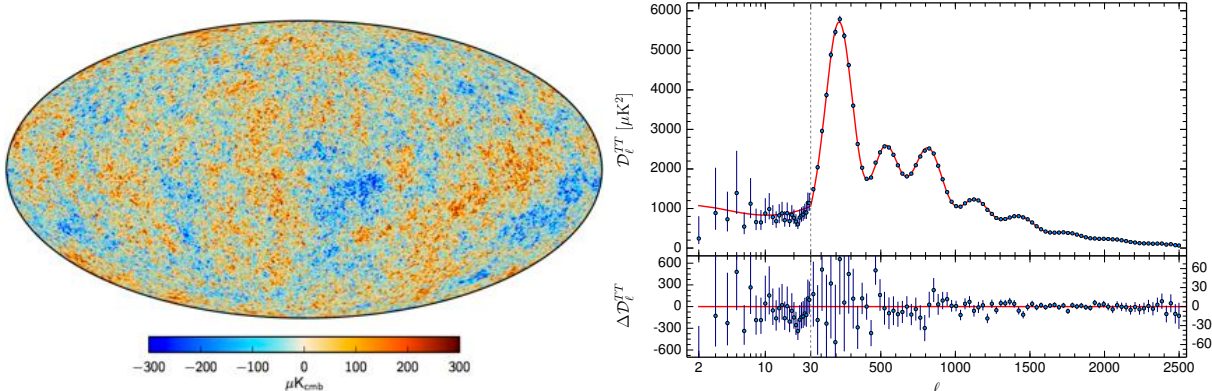


Figure 1.2: (Left) The Cosmic Microwave Background (CMB). (Right) temperature spectrum as a function of the multipole moment, l , in the CMB as predicted by a 6-parameter Λ CDM model (red) and measured by the Planck collaboration (navy blue). The top figure shows the temperature spectrum, while the bottom shows the residuals. The fit to these parameters constitutes one of the most statistically significant evidences for the existence of cold, non-baryonic dark matter. Figures from [16].

measurement of the CMB anisotropies to-date [15]. Their results, shown in Figure 1.2, fit extraordinarily well with a 6-parameter Λ CDM model that describes an expanding, flat universe following the laws of general relativity and dominated by a cosmological constant (Λ) and cold dark matter (CDM). They found that the CMB anisotropies are consistent with $\Omega_{\text{DM}} = 0.2608 \pm 0.0059$ with 68% confidence [4]. Similarly, the baryonic matter content was found to be consistent with $\Omega_b = 0.04879 \pm 0.0007$ with a Hubble parameter of $h = 0.673 \pm 0.012$. The Planck results are the most statistically significant evidentiary results for non-baryonic dark matter to-date, constituting a 42σ result for the existence of cold dark matter.

Big Bang Nucleosynthesis (BBN)

Measurements of the ratios of light elements like ${}^4\text{He}$, ${}^2\text{H}$ (deuterium), and ${}^7\text{Li}$ in low-metallicity regions of space² are used to determine the rate of Big-Bang primordial nucleosynthesis (BBN). In this process, light elements were synthesized from hydrogen when the

²Stellar nucleosynthesis alters elemental and isotopic abundances, synthesizing elements heavier than He, or in standard astrophysical nomenclature, “metals.” Areas of the universe with low metallicity have been less affected by stellar nucleosynthesis, and will therefore have elemental and isotopic abundances closer to that of the early universe.

universe was at higher temperatures, around time $t = 1$ s. The primordial ratios of the light elements to hydrogen are directly dependent on the density of baryonic matter in the early universe [9]. Observing the concentrations of these elements in regions minimally affected by stellar synthesis provides an estimate of these early abundances.

In Figure 1.3, the constraints on the baryon density for different light elements are shown. The baryon density is expressed either as $\Omega_b h^2$, or in terms of the baryon-to-photon ratio η_{10} , the baryon density normalized to the black-body photon density. Here, the subscript $_{10}$ is a multiplier such that $\eta_{10} \equiv \rho_b/\rho_\gamma \times 10^{10}$. The banded curves are the predicted abundances of the light elements as a function of the baryon density, while the yellow squares are the observed limits on the abundances. BBN is the only significant source of deuterium in the universe, which is destroyed in the stellar environment [17]. The observation of deuterium has therefore placed the most stringent limits on the baryon density, shown in the pink band. The ^4He abundance (expressed as a factor $Y \equiv \rho(^4\text{He})/\rho_b$ in the figure) agrees with the deuterium measurement. The band is also in excellent agreement with the restrictions on the baryon density from measurements of the CMB. Though the ^7Li abundance is not in agreement with the deuterium band, a discrepancy known as the ‘‘Lithium Problem,’’ the strong agreement between the CMB prediction and the deuterium observation is otherwise convincing. According to the measurements of the primordial deuterium abundance, the current baryonic matter density should be $0.021 \leq \Omega_b h^2 \leq 0.025$ (95% confidence limit) [17].

Together, these measurements of the cosmological parameters from the CMB and from BBN form extremely convincing evidence that dark matter is not only dark, but it is non-baryonic as well. These measurements use completely different techniques and depend on the conditions of the universe in very different epochs, yet are in remarkable agreement, as can be seen by the vertical bands in Figure 1.3. These measurements therefore form the most convincing evidence for a new kind of matter to-date.




Figure 1.3: Constraining the current baryon density with BBN. The baryon density is shown on the x axis in the top plot, while the baryon-to-photon ratio is shown on the x axis for the bottom three plots. The y axis is broken up for each element, and represents the elemental abundance as a fraction of the baryon density (in the case of ^4He , or Y) or the H abundance (for ^2H (D), ^3He , and ^7Li). The colored bands represent the Standard-Model predicted abundances of ^4He , D, ^3He , and ^7Li in magenta, blue, red, and green, respectively. The ^4He abundance is expressed by $Y \equiv \rho(^4\text{He})/\rho_b$, where ρ_b is the baryon density today. The yellow boxes represent the restrictions on the model by the measured elemental abundances. The heaviest constraint comes from deuterium. Finding the ^3He primordial abundance is too difficult due to a number of systematic uncertainties in the astrophysical parameters that determine the current abundance [17], so no yellow box is shown. The narrow cyan vertical band represents the baryon density according to CMB, while the wider pink band comes from BBN. The ^7Li region does not overlap with the deuterium region or the CMB region, an anomaly known as the Lithium Problem. All regions are the 95% confidence level regions. Figure from [17].

1.2 Dark matter candidates

Though ample evidence for the existence of dark matter has been accumulated, its nature remains unknown. Though there exist some alternative gravitational theories that could explain the apparent missing mass, most theories of dark matter postulate a new particle or group of particles that could explain the observations mentioned above. Several particle candidates for dark matter have been proposed; their characteristics are described below.

1.2.1 Characteristics of dark matter

From the evidence described in Section 1.1, several attributes of dark matter can be inferred. The primary requirement for dark matter is that it be dark: in other words, interacting at most very weakly with electromagnetic radiation. The dark matter must also be stable on cosmological time scales, or it would have decayed by now. Studies of structure formation show that dark matter, which has a strong influence on the formation of galaxies and galaxy clusters, must have been non-relativistic at the time of galaxy formation [18]. Requirements for the relic density of dark matter can also constrain the possible candidates.

1.2.2 Baryonic dark matter

While it is not impossible that some of the dark matter be composed of baryonic matter, BBN places a limit on the amount that baryonic matter can contribute to the total density of dark matter. It has been proposed that some of the dark matter could be composed of small concentrations of baryonic matter, called MAssive Compact Halo Objects (MACHOs). Microlensing is a gravitational lensing technique suited to find such objects. This technique, along with quasar studies, has limited the possible contribution of Compact Dark Objects (CDOs) with masses between 10^{-6} and $10^6 M_{\odot}$, where M_{\odot} is the solar mass, to $\Omega_{CDO} = 0.1$ [9].

Density fluctuations occurring in the early universe could have collapsed before the onset of BBN to create primordial black holes. These black holes would be non-relativistic and small ($\sim 10^{-8}$ cm in diameter), and therefore unlikely to collide with each other [19].

Experimental evidence suggests that primordial black holes are an unlikely candidate for several reasons. First, black holes of a mass smaller than 5×10^{14} g would not have survived to the present due to Hawking radiation³. Slightly larger black holes that would have survived would also radiate, but this radiation is not seen in the extra-galactic γ -ray background. The Energetic Gamma Ray Telescope, for example, puts a dark-matter limit of primordial black holes smaller than $m = 10^{16}$ g at 1% of the dark matter density [19]. Several other experiments have limited the density of primordial black holes to a few percent of the total dark matter density in other mass ranges, with only a small mass window remaining [19].

1.2.3 Neutrinos

Historically, neutrinos (ν) were considered a candidate for dark matter. However, neutrinos are relativistic [20], and therefore are limited as a dark matter candidate by structure formation. Furthermore, limits on the neutrino mass constrain the neutrino abundance to below $\Omega_\nu < 0.016$ [20]. Sterile neutrinos, on the other hand, remain a potential dark matter candidate.

In the Standard Model, ν are massless and the lepton numbers⁴ are conserved, but we know experimentally that they have tiny non-zero masses and that the flavors mix with each other⁵. Sterile neutrinos are a Majorana⁶ fermion that do not interact weakly like a standard

³A black hole dissipates energy through Hawking radiation. Therefore, a black hole slowly evaporates due to this power output. The time it takes a black hole of mass M to evaporate completely is given by $t_{ev} = 5120\pi G^2 M^3 / \hbar c^4$, where \hbar is the reduced Planck constant. Black holes with the mass listed above would evaporate in a time equivalent to the age of the universe.

⁴The lepton families, e , μ , and τ each have, in theory, separately conserved lepton numbers that describe the number of leptons (minus anti-leptons) in a reaction. An electron neutrino oscillating into a muon neutrino would violate this conservation law, for example.

⁵Neutrino oscillations are the observed phenomena that neutrinos oscillate through the different lepton families, existing in mixed states that can interact as different types of neutrinos with certain time-dependent probabilities [21].

⁶It is its own antiparticle.

neutrino, but do interact gravitationally. They may oscillate with the other neutrino flavors, however, and could still be detectable via this interaction. The sterile neutrino would be unstable but with a long decay lifetime. Sterile neutrinos are an attractive candidate for a new particle because they can explain dark matter, as well as baryon asymmetry⁷ and neutrino oscillations. However, this model requires fine-tuning [9]. Because of their instability, sterile neutrinos would not be completely dark, producing a characteristic decay line. Constraints on the sterile neutrino can therefore be made by looking for the decay line [9].

1.2.4 Axions

Axions were proposed as a solution to the Strong CP problem⁸. Axions are a Nambu-Goldstone boson of the broken $U(1)_{PQ}$ symmetry, a symmetry proposed by Peccei and Quinn. This global $U(1)$ symmetry, when broken by the axion, cancels the CP-violating term in the QCD Lagrangian, preserving CP conservation [22, 23].

The primary method for attempting to observe axions is through the axion-two- γ interaction, which allows for the conversion between axions and photons in an electromagnetic field. In a large-scale magnetic field, the axion-two- γ interaction operates as an oscillation similar to that of neutrino flavors [23]. Several experiments using techniques related to this axion-two- γ interaction are underway to search for axions in the mass range of less than 1 eV, some possibly reaching the 10^{-4} - to 10^{-6} eV range. Astrophysical observations can also shed light on the existence of axions. The mass regions that have not been excluded (see Figure 1.4) suggest that if axions exist they could comprise much or all of the cold dark matter. [23]

⁷The imbalance of matter and antimatter in the universe.

⁸The quantum-chromodynamics (QCD) Lagrangian contains the following term:

$$\mathcal{L}_{QCD} = \frac{\theta g^2}{32\pi^2} G_{\mu\nu}^a \tilde{G}^{a\mu\nu} + \text{other terms} \quad (1.1)$$

which violates time and parity conservation but keeps charge conjugation invariance, thus violating charge-parity (CP) conservation. This term's θ -parameter is limited to the order of 10^{-10} or less based on measured limits of the neutron electric dipole moment, while there is no reason to limit this parameter in such a way. The fact that we do not observe CP violation in nature despite the existence of this CP-violating term is known as the Strong CP problem.




Figure 1.4: Mass regions excluded for the axion by various experiments and measurements. Plot courtesy of [23].

1.2.5 Weakly Interacting Massive Particles (WIMPs)

Weakly Interacting Massive Particles (WIMPs), denoted by the symbol χ , are a class of non-Standard Model, non-relativistic particle that have a characteristic abundance determined by the thermal freeze-out method. If the WIMPs were in thermal equilibrium with Standard Model particles after inflation, they would fall out of thermal equilibrium when the temperature of the universe fell significantly below the WIMP mass (m_χ), which is typically assumed to be between 10 GeV and a few TeV. After this point, the co-moving number density of the WIMPs would remain relatively constant. The motivation for WIMPs stems from the idea that if the new particle decays with weak strength (corresponding with an annihilation cross section of $\sim 10^{-26} \text{ cm}^2$), then the current thermal relic density would match observations.

Candidate	Symbol	Spin	Composition
Neutralino	χ	1/2	$\alpha\tilde{B} + \beta\tilde{W}^3 + \gamma\tilde{H}_1 + \delta\tilde{H}_2$
Sneutrino	$\tilde{\nu}$	0	$\tilde{\nu}$
Gravitino	\tilde{g}	3/2	\tilde{g}

Table 1.1: SUSY dark-matter candidates. \tilde{H}_1 and \tilde{H}_2 are the neutral higgsinos, while \tilde{B} and \tilde{W}^3 are the bino and wino.

The fact that the required annihilation cross section happens to come out at the weak scale is known as the “WIMP miracle.”

Recent indirect detection experiments have become sensitive to the annihilation cross section posited by the WIMP miracle and have excluded this region [24], so the original motivation for WIMPs is no longer valid. However, the generic class of particles that couple weakly (in the colloquial sense) with baryonic matter remains an attractive group of dark-matter candidates. A few proposed WIMP candidates are described below:

Supersymmetric candidates

Super-symmetry (SUSY) is one of the most highly-motivated theories beyond the Standard Model (SM). It is additionally motivated with the recent discovery of the Higgs, whose mass of ~ 125 GeV is consistent, albeit with some fine-tuning, with SUSY. The quantity R-parity $R = (-1)^{3B+L+2s}$, where B is baryon number, L is lepton number, and s is the spin, is +1 for Standard Model particles and -1 for their supersymmetric partners. If R-parity is conserved, the lightest supersymmetric partner particle (LSP) would be stable and would operate by the thermal freeze-out mechanism. The lightest neutralino (a Majorana fermion) and the gravitino (the supersymmetric partner of the graviton) are potential candidates. The sneutrino, the supersymmetric partner of the neutrino, has been heavily constrained by direct detection experiments and by measurements of the invisible Z-boson decay rate at the Large Electron-Positron Collider at CERN [9]. The SUSY candidates are summarized in Table 1.1.

Kaluza-Klein dark matter and extra dimensions

Kaluza-Klein dark matter consists of excitations in extra dimensions, or Standard Model particles with excited modes in extra dimensions [9]. In general, Kaluza Klein particles would be unstable, but if extra quantities (such as the momentum in the new dimension) are conserved, there may be a new discrete symmetry, called KK parity [9]. This would allow the lightest KK particle, or LKP, to be stable.

Alternate Models

Other WIMP candidates include, but are not limited to, the Little Higgs model [25], Mirror dark matter, minimalist dark matter models that add to the Standard Model a scalar coupled only to the Higgs, and a charged dark matter bound state called the WIMPonium. The details of these models will not be discussed here, but can be found in [9].

Some models abandon the thermal freeze-out model altogether and focus instead on the similar relic densities of dark matter and baryonic matter ($\Omega_{\text{DM}}=5\Omega_b$). If dark matter consists of a particle-antiparticle species and could pass on its asymmetry to baryonic matter, then the similar relic density of dark and baryonic matter could be explained [9].

Though dark matter has been observed through its gravitational effects, the interactions of dark matter with normal matter, if any, are unknown. A narrowing of the possibilities of dark matter candidates through dark matter searches is needed to form conclusions about the models listed above.

1.3 Technical aspects of dark-matter searches

There are three main prongs in the effort to observe dark matter's interactions with baryonic matter. They are shown schematically in Figure 1.5, a Feynman diagram of the interaction between dark-matter particles and Standard-Model particles. The circle represents the in-

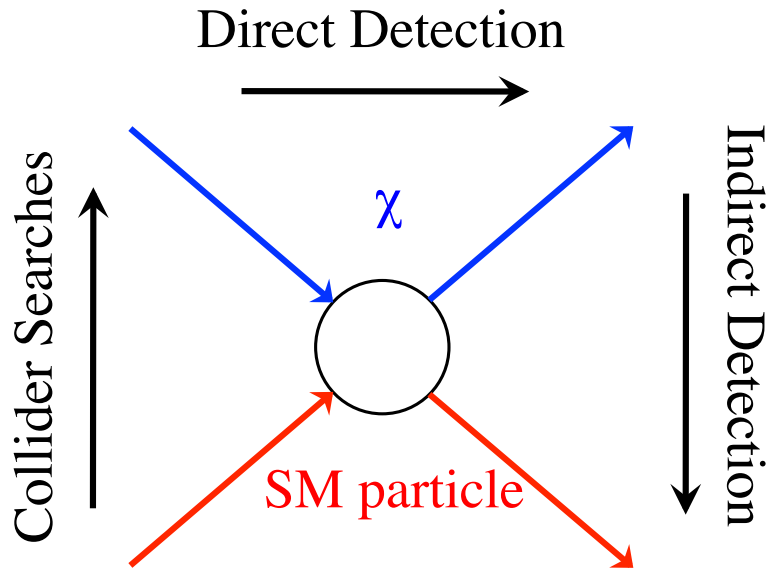


Figure 1.5: Feynman diagram illustrating the three types of dark-matter detection. The diagram represents the interaction between a dark-matter particle (blue) and a Standard-Model particle (red). The black arrows represent the flow of time with each detection method. In collider searches, where time flows from bottom to top, Standard-Model particles are annihilated in the hopes of producing WIMPs. In the direct-detection scheme, flowing from left to right, a WIMP scatters elastically with a Standard-Model particle, imparting some energy to the latter. In indirect-detection searches, WIMPs in space annihilate to produce Standard-Model particles whose signatures can be seen in cosmic-ray showers.

teraction, whose details are unspecified. The black arrows represent the flow of time in the various search modes.

In collider searches, particle colliders produce showers of high-energy particles, some of which may annihilate and produce dark-matter particle-antiparticle pairs. The dark-matter particles will escape the detector, and their characteristics can be inferred from the missing transverse momentum.

Indirect-detection searches look for the reverse interaction; dark-matter particles annihilate in space to produce Standard-Model particles. These particles can be observed to infer the mass and annihilation cross section of dark-matter particles. For example, γ -ray searches at the Galactic Center, where the dark-matter density is highest, and dwarf galaxies, which have high mass-to-light ratios, may have an enhanced ability to detect evidence

of dark-matter annihilation. Dark-matter annihilations into other Standard-Model particles would produce equal numbers of particles and antiparticles. The observation of antiparticles is therefore a potential avenue for indirect detection of dark matter. Dark matter could also thermalize and collect in the Sun; solar neutrinos can therefore also be a means for detecting dark matter indirectly.

Direct-detection searches look for scattering events between dark-matter particles and Standard-Model particles. In these interactions, the dark-matter particle imparts some energy to the Standard-Model particle; this energy can be observed with methods that have been developed for observing the energy deposits from other kinds of radiation. Generally, direct-detection experiments narrow their search to elastic scattering events with the nuclei of atoms in a detector. The technical aspects of these searches are discussed in the following sections.

1.3.1 The WIMP scattering event energy

WIMP dark matter is typically assumed to scatter elastically with atomic nuclei⁹. In this case, the recoil energy, E_R , of a nucleus with mass M after scattering with a WIMP of mass m_χ is given by:

$$E_R = \frac{\mu_\chi^2 v^2 (1 - \cos \theta_{cm})}{M}, \quad (1.2)$$

where θ_{cm} is the scattering angle in the center-of-mass frame, $\mu_\chi = m_\chi M / (M + m_\chi)$ is the reduced mass of the WIMP-nucleus system, and v is the relative velocity between the WIMP and the nucleus. If v is taken to be the characteristic velocity of the WIMPs with respect to the Sun (~ 200 km/s), E_R will be a characteristic energy of 1–100 keV for a 10–1000 GeV/ c^2 (hereafter GeV) WIMP. Dark-matter detectors must therefore be sensitive to energy deposits on this low-energy scale.

⁹In some inelastic dark-matter models, however, WIMPs in energy states with broken degeneracy can convert kinetic energy into more mass. However, these models only affect certain targets with low excited energy states and require more complicated dark-matter models, so elastic scattering is generally assumed.

1.3.2 The WIMP event rate

WIMP dark matter, if the dominant or only component of dark matter, is highly prevalent. Recent estimates place the local dark-matter density at $\rho_\chi = 0.3 \text{ GeV/cm}^3$ [26], though the uncertainty on this value are historically large [27], between 0.2 and 0.4 GeV/cm^3 . The flux of WIMPs through the earth is therefore $\sim 10^5 (100 \text{ GeV}/m_\chi) \text{ cm}^{-2}\text{s}^{-1}$ [9]. However, the rate of WIMP scattering events in a dark-matter detector will be very low because of the weak WIMP-nucleon scattering cross section. The exact event rate in a given detector is a product of several different factors: the WIMP-nucleus-scattering cross-section strength and form, the WIMP velocity distribution, and nuclear effects dependent on the target. The differential event rate per unit detector mass can be expressed as [28]:

$$\frac{dR}{dE_R}(E_R, t, m_\chi, \sigma) = \frac{n_\chi}{M} \langle v \frac{d\sigma}{dE_R} \rangle = \frac{\rho_\chi}{M m_\chi} \int_{v_{\min}}^{\infty} v f(v, t) \frac{d\sigma}{dE_R}(v, E_R) dv, \quad (1.3)$$

where ρ_χ is the local WIMP density (and n_χ the number density), $\frac{d\sigma}{dE_R}$ is the differential cross section for the WIMP-nucleus elastic scattering, and $f(v, t)$ is the normalized WIMP speed distribution, which may be time-dependent. $v_{\min} = \sqrt{M E_R / 2 \mu_\chi^2}$ is the minimum WIMP velocity to produce a scattering with recoil energy E_R . If $f(v, t)$ follows a Maxwellian distribution, the differential event rate is exponential in shape for a simplified model¹⁰.

Dark-matter rates are often considered in units of counts per day (cpd) per kg of detector per keV of the recoil energy spectrum. This unit is also known as the dru, or (cpd/kg/keV). To an order of magnitude, it can be estimated that a typical rate of dark-matter scattering events between a 100-GeV WIMP and a similarly massive target nucleus with a WIMP-

¹⁰For a detailed derivation, see [28].

nucleon cross-section of 10^{-40} cm² will have a rate of around 1 cpd/kg/keV:

$$\frac{(0.3 \text{ GeV}/c^2 \text{ cm}^{-3})(10^{-34} \text{ cm}^2)}{(100 \text{ GeV}/c^2)^3} \frac{1}{3 \times 10^7 \text{ cm/s}} \frac{6 \times 10^{26} \text{ GeV}/c^2}{\text{kg}} \frac{\text{GeV}}{10^6 \text{ keV}} \frac{10^6 \text{ s}}{\text{day}} \left(\frac{3 \times 10^{10} \text{ cm/s}}{c} \right)^2 \approx 1 \text{ day}^{-1} \text{ keV}^{-1} \text{ kg}^{-1}, \quad (1.4)$$

but this number can vary drastically with different WIMP characteristics and nuclei.

The differential event rate can be integrated over an energy region to obtain a rate within a certain window, or over all energies above the energy threshold of the experiment (E_T) to obtain the total rate, R :

$$R(t, m_\chi, \sigma) = \int_{E_T}^{\infty} dE_R \frac{\rho_\chi}{M m_\chi} \int_{v_{\min}}^{\infty} v f(v) \frac{d\sigma}{dE_R}(v, E_R) dv. \quad (1.5)$$

This relationship and that in Equation 1.3 can be further illuminated by assuming a cross-section model where the velocity-dependence is known and a velocity-independent cross-section which only depends on the recoil energy:

$$\frac{d\sigma}{dE_R}(E_R, v) = \frac{\sigma(E_R)}{4\mu_\chi^2 v^2} \Theta(E_{R,\max} - E_R). \quad (1.6)$$

Here, $E_{R,\max}$ is the maximum recoil energy possible for a collision with relative velocity v , and Θ is the Heaviside step function¹¹. $\sigma(E_R)$ is a velocity-independent cross section only dependent on E_R . This velocity-independent cross-section can further be divided into an energy-independent WIMP-nucleus cross-section, σ_0 , and an energy-dependent form factor, F , that depends on the shape and size of the target nucleus:

$$\sigma(E_R) = \sigma_0 F^2(E_R). \quad (1.8)$$

¹¹ The Heaviside step function is defined as:

$$\Theta(x) = \begin{cases} 1 & \text{for } x \geq 0 \\ 0 & \text{for } x < 0 \end{cases} \quad (1.7)$$

If an astrophysical and nuclear model can be assumed, it is evident from Equation 1.3 that the differential event rate is determined uniquely by the energy-independent cross-section and the WIMP mass. Thus, from an energy spectrum of nuclear-recoil events in a given detector, a parameter-space region can be determined in the WIMP mass/cross-section plane. A direct-detection experiment that fails to see dark-matter events will produce limit curves like those shown in Figure 1.6. The shape of the right side of the curve is determined by the fact that the mass and cross section will be proportional to each other for a given event rate; the event rate is proportional to the cross section times the flux, which is in turn inversely proportional to the WIMP mass for a fixed WIMP density. The rise of the limit curve to higher cross sections on the low-mass side is a result of experimental factors like the energy threshold.

1.3.3 Types of direct detection

When a nucleus is imparted with some recoil energy, that energy is dissipated into the medium in a number of ways. First, it can excite electrons to higher energy states, which later de-excite with a characteristic decay time, releasing optical photons. This phenomenon is called scintillation. Second, the recoiling nucleus can ionize nearby atoms, releasing free electrons into the medium. Third, the energy can be dissipated as heat. Individual dark-matter detectors take advantage of one or more of these techniques when searching for WIMP-nucleus recoils.

Scintillation

Scintillators are materials that release optical photons when energy is transferred to electrons in the material that later de-excite. In general, the number of photons released by the material is proportional to the energy deposited, though there are some non-linearities in some scintillators. A typical quantity then, is the scintillation yield, which is the number of photons released per unit energy. Due to the properties of the material, scintillators are

mostly transparent to their own scintillation light, making it possible to detect the optical photons on the periphery of the scintillator volume. A detailed discussion of scintillation detectors and the principles of scintillation is in [29], and is summarized in Appendix A.

The scintillation response of many materials depends on the type of incident particle. In particular, different particles ionize the material along the track with different ionization density, which is related to the stopping power of the ionizing radiation. Incident radiation can interact with the nuclei of atoms in the material or impart their energy to electrons. These nuclear recoils have a different stopping power than the electron recoils. The relative scintillation response between electron and nuclear recoils is characteristic of the material and is referred to as the quenching factor, Q , which can be energy dependent. Generally for dark-matter searches, nuclear recoils are the relevant interaction, but calibration of the detector light yield with mono-energetic source is easier with particles that produce electron recoils, like γ rays. Therefore, it is important to determine the detector response to nuclear recoils of different energies and compare them to the response to an electron recoil calibration, in order to find the energy-dependent quenching factor and determine the energy scale of nuclear recoils in a dark-matter search.

To characterize the scintillation response of a material independently of the type of interaction, the unit keV-electron-equivalent (keV_{ee} , or MeV_{ee} , or eV_{ee} , etc.), is used to describe the energy that produces the same scintillation response as an electron recoil with that energy¹². For electron recoils, the keV_{ee} and keV are equivalent; for nuclear recoils, they are related by the quenching factor: $1 \text{ keV} = Q(E_R) \text{ keV}_{ee}$.

Detecting scintillation light for low-energy events requires a detector sensitive to single photons. Typically, this is accomplished with a device called a photomultiplier tube (PMT). PMTs consist of a vacuum chamber with a window and an interior series of dynodes. The window is coated with a photocathode material that releases an electron when struck by

¹²Detectors are typically calibrated at a single energy, but a scintillator-detector response is typically assumed to be linear with energy. This is not necessarily the case for all materials, so the source used to calibrate the energy scale is usually specified. In this case, the electron-equivalent response is assumed to scale linearly with energy, calibrated at one point.

a photon via the photoelectric effect. This first emitted electron is called a photoelectron. A potential difference is imposed on the dynodes in the dynode chain by an externally applied high-voltage source. Typically, a high voltage is applied to a voltage-divider chain wherein each dynode is separated by a resistor, so that each successive dynode is at a higher voltage¹³. This potential difference drives the electron from dynode to dynode. At each dynode, a snowball effect occurs, producing a cascade of more electrons that travel through the chain. By the time the cascade reaches the anode, a gain in current of around 10^6 or more is achieved. This produces a detectable current, which constitutes the PMT signal.

Particular qualities of phototubes that are of concern for dark-matter experiments are the gain, the quantum efficiency, the collection efficiency, and the dark rate. The gain is the amount of current (in electrons) that results from the emission of a single photoelectron. The quantum efficiency is the chance that a photon striking the photocathode will cause the emission of a photoelectron, while the collection efficiency is the chance that the photoelectron, once emitted, will strike the first dynode and begin the cascade. Typical quantum efficiencies for current PMTs are in the 20–40% range, while the collection efficiency of a tube is typically higher, on the order of 80% or higher. Finally, the dark rate is the rate of PMT noise events when the phototube is in complete darkness. This rate can be on the order of 1 kHz but can be reduced at lower temperatures.

Ionization

A recoiling nucleus can ionize nearby atoms, generating free electrons in the material. These electrons can be gathered and detected by applying a strong electric field to the material. For example, in noble-liquid time-projection-chamber (TPC) experiments, where the dark-matter target consists of a volume of noble liquid with an array of PMTs at the top and bottom of the detector, an ionization signal is detected by applying a strong, uniform electric field drawing electrons upward to a thin gas layer at the top of the detector. The electrons

¹³A thorough discussion of the design of such a circuit can be found in [29] and [30], as well as in Section 4.3.2.

can excite or ionize atoms in the gas, producing photons through electroluminescence, thus providing a detectable signal for the PMTs.

Temperature

Cryogenic bolometric detectors take advantage of the small amount of heat that is dissipated by a recoiling nucleus. At cryogenic temperatures, this small amount of heat can be detected as a small temperature change in the material, which is equal to the nuclear recoil energy if all of the energy is dissipated as heat.

1.3.4 The sensitivity of direct-detection experiments

The achievable sensitivity of a dark-matter experiment depends on four things. The first is the exposure to a dark-matter signal, which is directly proportional to the mass of the target material and the time of operation. Typically exposures are expressed in units of mass-time, such as the kg-year.

The second is the choice of target material. Depending on the form of the WIMP coupling with matter, WIMPs may interact more frequently with certain nuclei than others. For example, if the WIMP coupling is dependent on the spin of the nucleon and is stronger for protons than neutrons or vice-versa, the sensitivity of nuclei with unpaired protons will be different than those with unpaired neutrons. Or, if the WIMP interaction adds coherently with all nucleons, the interaction cross-section will scale with the square of the atomic number, A^2 , making larger nuclei more sensitive. The third and fourth factors are the energy threshold of the experiment and the level of background events from other types of radiation, as will be described in more detail below.

Energy threshold

The third determinant of the sensitivity is the energy threshold. Because $E_R \propto \mu_\chi^2$, the characteristic recoil energy is proportional to the WIMP mass when it is much smaller than

the target mass¹⁴. The energy threshold of an experiment is therefore an important factor in its sensitivity to lighter WIMPs.

The energy threshold is determined by a large number of factors that affect the chance that a low-energy event will be detected. In scintillation experiments, a prominent factor is a quantity called the light yield (LY), which is the number of photoelectrons detected by the PMTs as a function of the energy deposited in the scintillator.

$$\begin{aligned} \text{LY} &= \frac{\text{number of photoelectrons detected}}{\text{energy deposited in scintillator}} \\ &= \text{fraction of photons detected} \times \text{scintillation yield.} \quad (1.9) \end{aligned}$$

This light yield is a function of the material’s scintillation efficiency and the efficiency of the light collection, as well as the properties of the phototubes themselves (such as the quantum efficiency). It is also dependent on the particle interaction type, which is why light yields are typically expressed in units of photoelectrons (p.e.)/keV_{ee} or p.e./MeV_{ee}, to remove this particular dependency. The light yield not only directly effects the energy threshold, but it is also important for the energy resolution of an experiment, since the observed energy is dependent on photon statistics. The higher the light yield, the less the observed energy is subject to statistical fluctuations.

The energy threshold is also dependent on the reduction of low-energy noise, since it can obscure a dark-matter signal. Some of this noise comes from the PMTs themselves; phototubes have a high rate of small-signal noise. Because the noise pulses from the phototubes are short in time compared with the rate of these events, noise can be mitigated by requiring coincident signals between multiple phototubes coupled to the same detector volume. Other noise-rejection techniques at low energies can help decrease the energy threshold, as will be described in Section 4.3.2.

¹⁴The reduced mass of a two-body system with two very different masses is approximately the mass of the lighter particle.

Background

Finally, because the WIMP-nucleus event rate is small, any detection mechanism mentioned in Section 1.3.3 will produce a lot of physical background from other particles interacting with the nuclei and electrons in the material. If this background rate is high enough, a legitimate WIMP signal can be lost in the background, and a discovery can be difficult to claim. The sensitivity of a detector to a WIMP signal is therefore dependent on the experiment's ability to block or reject backgrounds. Though there are some unique expected characteristics of a WIMP signal that can be used as a marker for dark matter in direct-detection searches (as will be discussed in later chapters), the observation of the total WIMP rate generally requires that an experiment eliminate most, if not all, backgrounds.

Physical backgrounds for dark-matter experiments come from three sources: trace radioactive isotopes in the detector materials and the surrounding environment, cosmic rays interacting in the detectors, and secondary radiation resulting from cosmic rays interacting in the vicinity of the experiment.

Most materials contain trace amounts of long-lived radioactive isotopes such as ^{40}K , ^{238}U , and ^{232}Th at the ppm, ppb, or ppt level¹⁵. These isotopes can emit background radiation at a wide range of energies from the sub-keV to MeV range. This radiation can come in the form of γ rays, neutrons, and β and α particles. If the decay of a nucleus results in a hole in its inner electron shells, a cascade can ensue, releasing X-rays and Auger electrons¹⁶ as well. In addition to these long-lived isotopes, some non-radioactive isotopes can be activated into shorter-lived radioactive states by cosmogenic radiation. Low-energy β s and α s have short interaction lengths and tend to deposit energy in the detector without escaping. γ s with high energy may not be a dark-matter background if they are absorbed in the sensitive detector material, since dark-matter signals are expected to have low energy. However, high-

¹⁵Parts-per-million, billion, and trillion by weight, respectively.

¹⁶Auger electrons are electrons emitted during the de-excitation of an atom whose inner-shell electrons have been absorbed by the nucleus in β decay or have been ejected by incident radiation. When an outer-shell electron de-excites to fill the vacancy, the energy released is imparted to another outer-shell electron, which escapes the atom. This escaping electron is called an Auger electron.

energy γ s can produce a background if they Compton scatter in the detector, leaving a small energy behind and subsequently escape. High-energy γ s have a high penetration power and can enter a dark-matter detector from outside as well. Fast neutrons have some penetrating power and can deposit energy in a dark-matter detector and subsequently escape, or they can thermalize and be captured.

Cosmic rays constantly bombard the Earth’s surface, and some, like high-energy cosmic-ray muons, can penetrate deep into the Earth’s crust. These cosmic rays can deposit energy in a detector directly or produce secondary particles that have similar effects to those listed above. Neutrinos (ν) can be a background for dark-matter experiments, but only at cross sections much lower than current sensitivities, and are generally not considered in dark-matter analyses for first-generation detectors. As detectors become more and more sensitive, like in the next round of Generation-2 (G2) and Generation-3 (G3) detectors, they will eventually reach this “neutrino floor,” at which point the ν background will have to be addressed for further progress to occur.

There are several techniques to reduce or reject backgrounds in a dark-matter experiment. Backgrounds from external radiation and cosmic rays can be reduced through passive shielding and underground operation. Because WIMPs interact so rarely with baryonic matter, they can pass through shielding and the Earth’s crust unimpeded, while some other particles can be stopped before reaching the detector.

Internal radiation can be reduced by purification, processing, and cleaning of the detector materials to reduce the prevalence of trace radioactive isotopes. The detectors can also be operated underground to prevent cosmogenic activation. Other techniques actively reject backgrounds that would otherwise mimic a dark-matter signal. One such method is to take advantage of the differing detector response for different particle types, since dark matter is expected to interact with the nuclei of atoms rather than the electrons. For example, the amount of energy that goes toward scintillation, ionization, and heat depends on the interaction type. Experiments that utilize more than one of the detection techniques in

Section 1.3.3 can distinguish between electron and nuclear recoils by observing the ratio of energy observed by the different techniques—for example, the ionization-to-scintillation ratio. In scintillation detectors, another technique is pulse-shape discrimination (PSD), which takes advantage of a scintillator’s differing timing response to electron and nuclear recoils. However, particles that induce or mimic nuclear recoils, like neutrons and sometimes α s, cannot be rejected in this way, and other methods are needed.

Both internal and external backgrounds can be reduced through a powerful method called the coincidence veto. Because dark matter is so rarely interacting, it can be assumed that the chance a single dark-matter particle will interact twice in a detector, or in two adjacent detectors, is negligible. Therefore, if two separate signals are seen in two nearby detectors at the same time, it can be assumed that the signals were caused by something else. Modular dark-matter detectors comprised of multiple detectors in close proximity can take advantage of this technique by rejecting coincident events in multiple detectors. Alternately, a detector can be placed inside of a second detector whose sole purpose is to detect incoming and outgoing radiation and provide a veto of coincident events in the dark-matter detector. The advantage of the second technique is that the background rejection power is not as dependent on the position in the dark-matter detector, since the dedicated veto detector gives 4π coverage of the sensitive volume.

All of these techniques can be used to drastically reduce backgrounds. Several experiments intend to operate in a “background-free” mode through the use of passive and active background reduction/rejection techniques. Others lower the background without completely eliminating it, relying instead on expected signatures of the WIMP signal, such as an expected annual modulation in the event rate, as will be described in later chapters.

1.3.5 The state of direct-detection dark-matter searches

The search for non-gravitational dark-matter interactions with normal matter is a vast effort that has been undertaken with many different target materials. The multiple-target




Figure 1.6: The state of direct-detection dark-matter searches in 2013, actual and projected. The plot shows the mass/cross-section parameter space for a spin-independent WIMP interaction in the Standard Halo Model, discussed in Section 8.2. The solid curved lines show limits for the WIMP established by existing experiments. These experiments are not consistent with WIMPs with cross sections above the line. Dotted lines are the projected limits for experiments currently in development. The colored shaded regions are those that are consistent with signals from the listed experiments. Of note are the regions for DAMA/LIBRA (light orange), CRESST (pink), CoGeNT (yellow), and CDMS-Si (blue). These experiments are discussed in more detail in Section 2.3.6. The shaded regions with no outlines correspond with some favored dark-matter interaction models. The yellow region under the red dotted line is the “neutrino floor,” the region in which neutrinos become a significant background for dark-matter searches. Figure from [31].

approach is important for dark-matter searches because 1) different materials may have different sensitivities to dark matter depending on the nature of the coupling and 2) once dark matter is found, different targets will be needed to constrain how the coupling works and what its strength is, as well as the dark-matter mass. Though dark-matter sensitivity has greatly improved over the last few decades, the existence of WIMP dark matter remains in question. The current status of the many existing dark-matter direct-detection experiments is summarized in Figure 1.6. In recent years, some hints of potential dark-matter signals have been observed, and one group, the DAMA (DARk MATter) collaboration, claims a discovery of WIMP dark matter with their DAMA/LIBRA (DARk MATter: Large sodium Iodide Bulk for RARE processes) experiment. Other direct-detection experiments seem to conflict with this result, however, so there is no consensus of a discovery.

The claim of a WIMP discovery by DAMA is an important controversy in the field of direct detection. In this document, the path toward verification or refutation of the DAMA/LIBRA claim is laid out. In Chapter 2, the principle of annual-modulation searches are discussed along with a description of the DAMA/LIBRA experiment, the controversy surrounding their result, and an attempt to understand their detector response through simulations. In Chapter 3, an approach to test the DAMA/LIBRA result is outlined, and the Sodium-iodide with Active Background REjection (SABRE) experiment is introduced as an attempt to follow this approach. In Chapters 4 and 5, the development of the components of the SABRE dark-matter crystal detectors and their proposed deployment in a liquid-scintillator active veto detector are described, respectively.

As will be described in Chapter 2, one important quantity affecting interpretations of the DAMA/LIBRA experiment is the ionization quenching factors in their scintillator. A measurement was done by SABRE to accurately measure the quenching factor in sodium (one of the target materials in DAMA/LIBRA) to shed light on the DAMA/LIBRA result and to resolve existing conflicts in the field regarding the value of this parameter. A description of the quenching measurement is outlined in Chapter 6, while the analysis of the results

is described in Chapter 7. Finally, the implications of the result for the DAMA/LIBRA controversy are discussed in Chapter 8.

Though the evidence for the existence of dark matter as a whole is incontrovertible, and despite great strides made in the WIMP sensitivity by current dark-matter direct-detection experiments, there remains no consensus as to the existence of WIMP dark matter or its properties. Hints of dark-matter scattering events are not yet convincing enough to warrant a claim of discovery. Until the DAMA/LIBRA result is tested, dark matter's nature will remain one of the deepest mysteries in modern physics.

Chapter 2

Annual Modulation Searches and DAMA/LIBRA

2.1 The WIMP annual modulation

The WIMP-nucleus interaction rate is dependent on the velocity distribution of dark-matter particles with respect to the lab frame on the Earth. Assuming that the WIMP halo has a net-zero velocity with respect to the Galactic center, the Sun has a relative velocity of ~ 220 km/s compared with the WIMP halo. The Sun therefore experiences a dark-matter “wind,” as shown in Figure 2.1. The Earth in turn orbits around the Sun at around 30 km/s, moving into the dark-matter wind in the summer and away from the wind in the winter. Because of the velocity dependence of both the flux of dark-matter particles and the scattering cross section, a dark-matter detector on Earth will experience a modulation in the event rate each year. This signature behavior can be exploited to detect dark matter, even amidst other backgrounds.

The annual-modulation signature has a number of key features that depend on the dark-matter model used. Typical analyses of dark-matter experiments make use of an astrophysical model called the Standard Halo Model (SHM), where the WIMP halo is assumed to act




Figure 2.1: Illustration of the annual modulation concept. If WIMPs are not self-interacting, the WIMP halo should have an isotropic velocity distribution with respect to the Galactic center. Due to the Sun’s velocity around the Galactic center, the WIMP velocity distribution in the Sun’s frame has a net non-zero mean velocity, called the dark-matter wind. As the Earth travels around the Sun, the Earth moves in and out of the wind at different times of the year. Because the event rate is dependent on the velocity distribution of the dark-matter halo, this produces an annual modulation effect in the WIMP-scattering event rate.

like a simple gas with a Maxwellian velocity distribution. Several other astrophysical models have been proposed, but for any distribution it is possible to expand the event rate as a function of the time of year with a Fourier series:

$$\frac{dR}{dE_R}(t) = S_0(E_R) + \sum_{i=1}^{\infty} S_{m,i}(E_R) \cos i\omega(t - t_0) + \sum_{j=1}^{\infty} S_{n,j}(E_R) \sin j\omega(t - t_0), \quad (2.1)$$

where $S_{m,i}$ and $S_{n,j}$ are the Fourier coefficients of the modulation, S_0 is the average rate, t_0 is the expected peak modulation time (\sim June 1st), and $\omega = 2\pi/\text{year}$. The $S_{n,j}$ can be ignored if the dark-matter velocity distribution is isotropic in its rest frame. In most models,

the $S_{m,i}$ can be ignored for $i > 1$ as well [32]. Thus, the modulation signature is taken, in general, to be a sinusoidal modulation of the following form:

$$\frac{dR}{dE_R}(t) = S_0(E_R) + S_m(E_R) \cos \omega(t - t_0) \quad (2.2)$$

Typically, the fractional modulation S_m/S_0 is less than 10% for many dark-matter models. Exceptions to this simplification include models where the high tail-end of the velocity distribution become important (v_{\min} is high¹) or where there is substructure in the dark-matter halo.

Gravitational focusing (see Figure 2.2) is an effect wherein dark matter is deflected by the Sun’s gravitational field on its way to the Earth. In the spring, dark-matter particles pass the Sun on the way to the Earth and are pulled closer together, increasing the flux. In the fall, this effect disappears. Gravitational focusing was previously thought to be a negligible effect [33], but new analyses show that it can shift the apparent phase of the modulation earlier in the year by around 21 days if v_{\min} is small [34]. A dark-matter halo with a slower average speed relative to the Sun is more susceptible to gravitational focusing, as would be in the case of self-interacting dark matter that forms a rotating disk like the baryonic matter in our galaxy.

The annual modulation is a signature characteristic of WIMP dark matter, and can be taken as evidence for its existence if observed in a low-background experiment. A high-mass detector with several years of exposure could be sensitive to such a modulation.

2.1.1 Technical aspects of annual-modulation searches

In any dark-matter experiment, the elimination of backgrounds is paramount to a discovery claim. Backgrounds from radioactive sources and cosmic rays can mimic dark-matter interactions and obscure their signal. Dark-matter detectors whose observations are based

¹From Equation 1.2, it is apparent that v_{\min} is high when the WIMP is light. Endothermic inelastic dark-matter models can also produce a high v_{\min} compared with elastic dark-matter models because additional kinetic energy is needed to transition the dark-matter particle into a higher energy state.




Figure 2.2: (Top) Illustration of gravitational focusing as it applies to the WIMP annual modulation, from [34]. The Sun’s gravitational field focuses the WIMP wind on the Earth in the Spring, effectively shifting the peak of the annual modulation towards March 21st. (Bottom) Effect of gravitational focusing on the phase of the annual modulation signature in NaI as a function of recoil energy and the WIMP mass, from [35]. The contours on the plot represent the time of the peak of the modulation as a function of energy for different WIMP masses. Here, the modulation is integrated over a 1-keV_{ee} window where E_{\min} is the low end of the window. The vertical purple line represents the energy threshold of DAMA/LIBRA, the current leading NaI-based dark-matter experiment. Gravitational focusing can move the peak of the modulation towards March, which could explain some aspects of the DAMA/LIBRA modulation phase (see Sections 2.2 and 2.3.3). This effect is more noticeable for higher-mass WIMPs, where v_{\min} is smaller.

on the total rate must be essentially “background-free,” with the background rate lower than the dark-matter rate, to make a credible claim of discovery. One advantage of the annual-modulation approach is that backgrounds need only be low enough to observe the modulation in a statistically significant way. Still, this modulation effect can be very small, and certain backgrounds can also modulate, so the reduction and understanding of backgrounds remains integral to annual-modulation searches. In addition, the magnitude of the fractional modulation can constrain the dark-matter model, so the careful quantification of residual backgrounds in a modulation experiment is of interest, since it can help determine the limits on the total dark-matter interaction rate.

Intrinsic backgrounds

Intrinsic backgrounds are those that result from the radioactive decay of unstable isotopes in the materials comprising a detector. There are several isotopes that are of particular concern to dark-matter searches, as they are especially difficult to remove or produce backgrounds that closely mimic a dark-matter signal. Isotopes of particular concern are ^{40}K , ^{238}U , ^{232}Th , and radon and its daughters, particularly ^{210}Pb . ^{40}K , ^{238}U , and ^{232}Th are long-lived and are present in most materials in small concentrations. They and their decay products produce a number of γ rays, low-energy X-rays and Auger electrons, α and β particles, and radiogenic neutrons.

Trace non-radioactive isotopes can be cosmogenically activated by cosmic rays, thus becoming radioactive. ^{60}Co and ^{125}I are examples of cosmogenically-activated sources of radiation. This is a particular problem for copper, iodine, and germanium, all common targets and/or detector materials used in dark-matter experiments. Cosmogenic activation can be mitigated by operating detectors deep underground after allowing the activated isotopes to decay and limiting the amount of time components spend above ground before a measurement.

Intrinsic backgrounds can be mitigated by choosing materials with already low concentrations of these isotopes or developing methods for purifying and cleaning materials. Purification of the atmosphere can help eliminate radon. Some of the remaining backgrounds can be reduced by discrimination techniques based on the differing detector response between different particle types, or by the coincidence-veto method.

Surface backgrounds

In addition to radioactivity in the bulk of the material, surfaces are a particularly difficult source of background to address. One source of surface background is radioactive isotopes in the bulk of materials in direct contact with the sensitive volume. The isotopes near the surface can emit low-energy ionizing radiation with a short range, such as α particles, that would otherwise be unable to reach the sensitive volume. Whereas similar isotopes in the bulk of the sensitive volume might be vetoed by coincident radiation that raises the detected energy of an event outside of the relevant dark-matter signal region, coincident radiation in these surface backgrounds may not make it to the sensitive volume. These surface backgrounds can then appear in the same energy window as the expected WIMP signal. This type of surface background can be limited by choosing radio-pure materials to be in contact with the sensitive volume. In detectors with position reconstruction capabilities, these backgrounds can also be addressed with fiducial volume cuts.

Surfaces pose a particularly strong challenge because some radioactive isotopes can be deposited on the surfaces of materials and cling to them. Radon is a radioactive gas present in the atmosphere that can deposit its radioactive daughters on the surfaces of objects, including ^{210}Pb , ^{210}Bi , and ^{210}Po . This type of surface background can be addressed by polishing the surfaces of materials and operating in a radon-free environment.

External backgrounds

External backgrounds come from radiogenic sources surrounding a dark-matter detector, or from cosmic rays. Deposits of ^{40}K , ^{238}U , and ^{232}Th in the rock surrounding an underground lab produce high-energy γ rays that can penetrate shielding surrounding a detector. High-energy cosmic-ray muons that penetrate deep underground can interact in a detector directly or produce showers of secondary particles, including cosmogenic neutrons, that can mimic a dark-matter signal. External backgrounds can be blocked using passive shielding or tagged using active veto techniques.

The Scintillator NaI(Tl)

Thallium-doped sodium-iodide, or NaI(Tl), is a common inorganic scintillator with a high scintillation yield of ~ 40 photons/keV $_{ee}$ [29]. In inorganic scintillators, a dopant material (in this case, thallium) creates energy states in the forbidden zone of the crystal between the valence and conduction bands. Electrons are excited into these energy states by energy depositions from incoming radiation or a recoiling nucleus, then de-excite with a characteristic time constant of the material, emitting visible photons. Because the dopant concentration is small, the activator sites are spread out, and the activation energy for the activators is smaller than the energy difference between the valence and conduction bands in the NaI material, so the scintillator can be transparent to its own emission light [29].

NaI(Tl) has a characteristic fluorescent decay time of ~ 230 ns. A slower tail with decay time $1.5 \mu\text{s}$ has also been observed [36]. The preferential decay by one mode over the other for certain particle types allows for the ability to distinguish between different particles that activate the scintillator. A phosphorescence also occurs with decay time 0.15 s contributing 9% to the total light yield [29]. NaI(Tl) can be made very pure; the crystal-growth process can segregate out impurities.

As a dark-matter target, NaI(Tl) has a number of advantages. It is a relatively cheap material that can be produced in large quantities. Because of its high light-output, it is

sensitive to small energy depositions (at the sub-keV level). It consists of one low-mass target (Na) and one high mass target (I), making it sensitive to both high-mass and low-mass WIMPs². The high mass of iodine also helps make the material more sensitive to spin-independent interactions, whose cross-sections scale with A^2 . Both ^{23}Na and ^{127}I , the most abundant isotopes of sodium and iodine, are odd-numbered nuclei with unpaired protons, making them sensitive to spin-dependent interactions³ as well. The long scintillation decay-lifetime allows for pulse-shape discrimination. Finally, their solid structure makes a modular experimental design feasible, allowing for rejection of events that occur in multiple crystals.

NaI(Tl) as a target poses some challenges as well. Some impurities are especially difficult to eliminate in NaI(Tl) experiments. Because of their position on the periodic table in the alkali metals, K and Rb are difficult to segregate from Na. ^{87}Rb (natural abundance 28%) is a β -emitter (end-point energy 283.3 keV), producing a flat background at low energies. ^{40}K is an isotope of K with natural abundance 0.011% and a long decay lifetime of 1.25×10^9 years. In 11% of decays, ^{40}K decays to an excited state of ^{40}Ar by electron capture, emitting a 1461-keV γ ray (see Figure 2.3). The hole left in the inner shells of the Ar atom (either the K or L shell) triggers a cascade of Auger electrons and X-rays with total energy of 3.2 keV. This 3.2-keV background can be rejected if the γ ray is observed⁴, but if not, this background occurs in the primary energy window of interest to WIMP searches, as will be discussed in Section 2.2.

NaI(Tl) is also susceptible to several cosmogenically activated radioactive isotopes, like ^{125}I , ^{22}Na , and ^{24}Na . These have shorter half-lives and can be reduced dramatically by leaving crystals underground for a long period. ^{129}I , however, has a long half-life of 1.57×10^7 years, and cannot be reduced in this way.

²In elastic scattering, the energy transfer and cross-section are highest when the two bodies have similar masses.

³Which only affect unpaired nucleons

⁴ In some cases, the ^{40}K can decay to the ground state of ^{40}Ar , leaving only the 3.2-keV signal, but this path has a low branching ratio, thought to be $\sim 0.2\%$ [37].




Figure 2.3: ^{40}K decay scheme, from [37].

One challenge NaI(Tl) poses as a dark-matter target is its highly hygroscopic nature. A crystal must be protected from moisture at all times to prevent a degradation of its optical properties. However, given that the detector array is modular anyway, crystals can be protected by remaining sealed in their detector modules, which can be thoroughly outgassed to remove trace amounts of water clinging to the surfaces.

Another challenge with NaI(Tl) is the non-proportionality of the scintillation efficiency as a function of energy, as well as the differing scintillation efficiency as a function of particle type. Both of these factors need to be understood thoroughly in order to determine the energy of scintillation events. Both of these effects will be discussed in Chapters 6, 7, and 8.

2.2 The DAMA and DAMA/LIBRA experiments

The DAMA experiments are annual-modulation dark-matter experiments that use an array of 25 10-kg NaI(Tl) crystals as the dark-matter target. The DAMA experiment and its successor, DAMA/LIBRA, have been observing a modulation for over a decade that the DAMA

collaboration attributes to dark matter. The modulation in the event rate for “single-hit” events⁵ is shown in Figure 2.4 and Table 2.1, and carries many of the expected signatures of a WIMP modulation, including the expected period and phase, within errors. The observed period is 0.998 ± 0.002 years, and the modulation peaks at 144 ± 7 days [38]. The dark-matter signal rate, along with its modulation, is expected to fall roughly exponentially with recoil energy. DAMA observes the modulation in their single-hit events in the energy region of 2–6 keV_{ee}, but no modulation at higher energies or in their multi-hit events, as shown in the top of Figure 2.5 and Table 2.1. The modulation rate appears to have a peak at around 3 keV_{ee}, but the falloff of the modulation below 2 keV_{ee} could be a threshold effect. The total rate of single-hit events observed by DAMA/LIBRA is also shown in the bottom of Figure 2.5. This energy spectrum of the event rate, if partially caused by WIMP-nucleus interactions, includes both a dark-matter signal and a mostly featureless background rate from other radioactive sources, as will be discussed in Sections 2.3.1 and 2.3.2. There is also a predominance of photomultiplier noise at low energies, which DAMA addresses by imposing a software energy threshold of 2 keV_{ee}.

DAMA has acquired enough data that the statistical significance of the modulation is unquestionable: 9.3σ at the beginning of 2015, according to [38]. However, their claim of a dark-matter discovery remains controversial, due in part to conflicts with other experiments with null results, as discussed in Section 2.3.6.

2.2.1 The DAMA/LIBRA apparatus

The DAMA/LIBRA apparatus is described in detail in [41]. To better qualify the physical details behind the signal that DAMA observes, the important details of the experimental apparatus are summarized here. The experiment consists of an array of 25 NaI(Tl) crystals each with a mass of 9.7 kg, for a total mass of ~ 250 kg. The crystals are long rectangular prisms with their two smallest faces coupled to photomultiplier tubes via 10-cm long synthetic

⁵“Single-hit” and “multi-hit” are the nomenclature used by DAMA for scintillation events occurring in one crystal versus in multiple crystals. This nomenclature will be adopted here as well.




Figure 2.4: The DAMA/LIBRA annual modulation of the low energy scintillation events, in the 2–4, 2–5, and 2–6 keV_{ee} energy windows, from top to bottom. Figure from [39].





Figure 2.5: (Top) The DAMA/LIBRA modulation amplitude, S_m as a function of energy, from [39]. The modulation appears to peak at around 3 keV_{ee} , which may or may not be a threshold effect. (Bottom) The energy spectrum of single-hit events in DAMA/LIBRA. From [40]. The spectrum has a featureless background characteristic of β decay and Compton scattering. In addition, a bump at $\sim 3 \text{ keV}_{ee}$ is characteristic of ^{40}K decay. Finally, at low energies the spectrum is dominated by photomultiplier noise. The energy scale for both figures is in keV_{ee} , referred to as “keV” by DAMA.

Table 2.1: DAMA/LIBRA modulation amplitude in different energy bins, reproduced from [39]. The day number is counted from January 1st. Single-hit events are defined as events that occurred in a single crystal, while multi-hit events occurred in more than one crystal.

Energy window (keV _{ee})	Modulation amplitude (cpd/kg/keV _{ee})	Period (yr)	Phase (day)
Single-hit events			
2–4	0.0183±0.0022	0.996±0.002	136±7 (May 16)
2–5	0.0144±0.0016	0.997±0.002	142±7 (May 22)
2–6	0.0114±0.0013	0.999±0.002	146±7 (May 26)
6–14	0.00007±0.00077		
Multi-hit events			
2–4	−0.0011±0.0007		
2–5	−0.0008±0.0005		
2–6	−0.0006±0.0004		

quartz (Suprasil B) light-guides. The crystals are encased in copper enclosures that are continuously flushed with high-purity nitrogen gas and surrounded by several layers of passive shielding. The environment is highly controlled for radioactivity, temperature and pressure stability, and moisture (see Table 2.2). The entire array is kept in an underground location in the Laboratori Nazionali del Gran Sasso (LNGS) at a depth of 3,100 meters-water-equivalent (m.w.e.) [42]. A schematic of the apparatus is shown in Figure 2.6.

DAMA/LIBRA became operational in 2003; its first results were reported in 2008. In 2010, DAMA upgraded their apparatus to include newer, higher-quantum-efficiency PMTs which allowed them to lower their energy threshold from 2 keV_{ee} to 1 keV_{ee}, but the data from this upgrade have not yet been released.

The crystals

As the dark-matter target, the crystals should have very few radioactive impurities that may mask or obscure the dark-matter signal, and they should have a high scintillation light yield in order to be sensitive to the low recoil energies where the modulation is expected to occur.




Figure 2.6: The DAMA/LIBRA apparatus. The crystals are each sealed in a copper box, coupled to two PMTs with 10-cm fused silica light-guides. The 25 crystal detectors are enclosed in a copper box and surrounded by several layers of radio-pure passive shielding, as described in the text. Figure from [41].

The crystals used by DAMA/LIBRA were grown by the Saint Gobain Crystals and Detectors company. The Kyropoulos method of crystal growth was used. In this method, a seed crystal is lowered into a crystal melt, the melt is slowly cooled, and the crystal forms around the seed [43]. DAMA asserts that this method produces higher-purity crystals than other methods, since the Kyropoulos method does not require that the crystal touch the crucible walls [41]. Nevertheless, the growths were performed in a high-purity platinum crucible in order to further prevent the introduction of impurities.

The residual backgrounds in the crystals were reported by DAMA within a range (e.g. ^{238}U and ^{232}Th) or with a mean value (e.g. ^{nat}K , where ^{nat}K is the inferred concentration of natural potassium, with ^{40}K having a natural abundance of 0.011%) and are summarized in Table 2.3. The most important backgrounds in the crystal are arguably ^{40}K , ^{232}Th , ^{238}U , ^{87}Rb , and ^3H , as will be discussed.

^{40}K is a particularly dangerous background for a NaI(Tl) dark-matter experiment. First, due to potassium's position in the periodic table, it is difficult to separate from sodium.

Secondly, the ^{40}K decay to ^{40}Ar produces the signature $\sim 3\text{-keV}_{ee}$ background that occurs in the same energy region as DAMA’s observed modulation. Though this background can be rejected by observing the corresponding 1461-keV γ ray that is released during the de-excitation of the ^{40}Ar nucleus, the γ ray may escape, leaving a background in the primary modulation region observed by DAMA. DAMA estimates the concentration of ^{40}K in their crystals by measuring the rate of 1461-keV $_{ee}/3\text{-keV}_{ee}$ coincidence events in their crystals and comparing it to a rate predicted by Monte Carlo. The details of this simulation have not been shared by DAMA, however, and conflict with other attempts to simulate the energy spectrum [44] [this work], as will be discussed in Section 2.3.2.

^{238}U and ^{232}Th have long decay chains that produce a variety of backgrounds, some in the low-energy region of interest to DAMA. DAMA estimates the concentrations of these isotopes by observing the peaks from characteristic α particles in the decay chain. DAMA claims that the ^{232}Th chain in their crystals is in secular equilibrium, but the ^{238}U chain is not. However, DAMA reports the relative activities for the different parts of the chain for only one crystal. The concentrations of ^{238}U and ^{232}Th assuming secular equilibrium are also given for all crystals in broad ranges.

Another important background for NaI-based scintillators is ^{87}Rb . Like K, Rb is an alkali metal and is hard to remove from NaI. ^{87}Rb is also a β emitter. Tritium (^3H) decays by β emission with an endpoint energy of 18.6 keV, thus producing a smooth background in the primary energy region of interest to DAMA. Both of these backgrounds are difficult to reject with the veto technique, as low-energy β s are not likely to escape their crystal of origin. DAMA reports limits for both of these backgrounds, but as will be discussed in Section 2.3.2, the background rate at these limits would constitute a large portion of DAMA’s signal, so it is unclear how prominent of a role they play in the DAMA experiment.

Backgrounds produced by cosmogenic activation, like ^{125}I , ^{24}Na and ^{22}Na , are diminished by allowing the crystals to sit underground for a long period before the operation of the experiment. DAMA reports limits for all of these backgrounds. For ^{22}Na , the activity at the

limits described will become relevant if DAMA lowers the energy threshold to 1 keV_{ee} , since ^{22}Na emits a $\sim 0.8\text{-keV}$ X-ray/Auger-electron with a branching ratio of 9%. Because ^{24}Na has a short half-life of 14 hours, DAMA uses this background to estimate the flux of fast neutrons in their experiment. DAMA also reports a limit for ^{85}Kr , but this isotope produces a small background compared with those previously mentioned (see Section 2.3.2).

^{129}I and ^{210}Pb produce backgrounds around $35\text{--}50 \text{ keV}_{ee}$, which is how DAMA measures the concentration of these isotopes in their crystals. However, they constitute a small background in the low-energy region at the levels reported by DAMA, as will be discussed in Section 2.3.2. However, the ^{210}Pb background can be significant for other experiments, and has been found to be a dominant background in other NaI(Tl) crystals grown for other experiments [45].

The PMTs and crystal housing

The original PMTs used for the DAMA/LIBRA experiment were produced by Electron Tubes Limited (ETL) using a bialkali photocathode. These tubes had a peak quantum efficiency of $\sim 30\%$, a dark-noise rate of 100 Hz, and a gain of around 10^6 [41]. DAMA upgraded to new, Hamamatsu R6233MOD phototubes in 2010 with a higher quantum efficiency (peak 38%) [39], which allows them to lower their energy threshold down to 1 keV_{ee} [46]. These new tubes have a dark-noise rate of $<100 \text{ Hz}$ and are operated with positive high-voltage between 900 and 1300 V.

The crystals are wrapped with Tetratex Teflon tape, which was chosen by DAMA for its radiopurity, and are packaged in oxygen-free high-conductivity (OFHC) copper enclosures. In order to protect the crystals from radioactivity from the PMTs, the PMTs are housed in a protective OFHC copper shield and separated from the crystal by a fused silica light guide, as shown in Figure 2.7. The DAMA crystal detectors have light yields in the range of $5.5\text{--}7.5 \text{ p.e./keV}_{ee}$ [41].




Figure 2.7: (Top) Shield used to protect the crystals from radioactivity in the DAMA PMTs. (Bottom) Picture of the DAMA/LIBRA PMT/light-guide setup. Figure from [41].

DAMA reports their radioactivity levels for the PMTs (see Table 2.3), but comments that the crystal exposure to these backgrounds is minimal due to the use of a light-guide to separate the crystal and PMTs, and also some shielding described below.

The Shielding

The DAMA/LIBRA crystal modules are surrounded by several layers of passive, low-radioactivity shielding. This shielding is designed to protect the crystals from external backgrounds, both cosmogenic and from the laboratory environment. The materials were kept underground before the experiment in order to allow cosmogenically activated isotopes to decay.

The crystal modules are first surrounded by 10 cm of copper bricks, which are in turn surrounded by 15 cm of low-radioactivity lead to block γ radiation. The lead array is covered in cadmium foils in order to block thermal neutrons. Finally, it is surrounded by 30 cm of

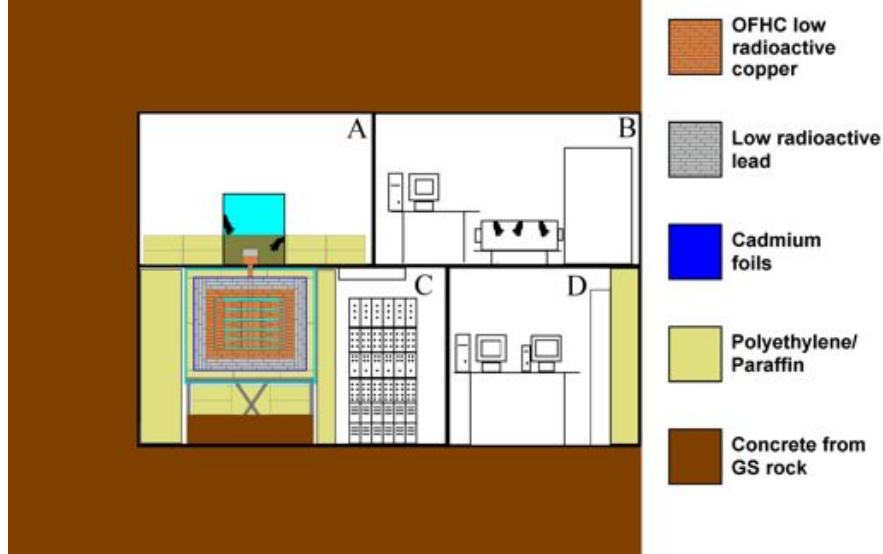


Figure 2.8: Layout of the DAMA/LIBRA experiment, from [41]. The four layers of shielding (OFHC copper, low-radioactivity lead, cadmium, and polyethylene/paraffin) surround the detector modules in room C.

paraffin and polyethylene⁶ distributed throughout the lab space. The layout of the shielding is shown in Figure 2.8.

The environment is also protected from radon contamination via two layers of sealed barriers. The first is a Supronyl covering around the room, and the second is a plexiglass box surrounding some of the shielding. The environment around the detector is filled with high-purity nitrogen gas.

DAMA has carefully controlled for the environmental conditions of their apparatus. They published the variation in various environmental variables and its expected effect on the annual modulation of scintillation events in [39], as shown in Table 2.2. They also measure environmental conditions continuously. They estimate that fluctuations in the signal rate due to environmental conditions like temperature, radon levels, changes in the energy scale or trigger efficiency, noise, background, and other reactions are each less than 1% of the observed modulation amplitude [47].

⁶These are hydrocarbons and therefore good at capturing neutrons.

Table 2.2: DAMA reported effects on the modulation amplitude by various environmental variables. The different columns represent different annual cycles of data-taking. Table from [39].

	DAMA/LIBRA-1	DAMA/LIBRA-2	DAMA/LIBRA-3	DAMA/LIBRA-4	DAMA/LIBRA-5	DAMA/LIBRA-6
Temperature	$-(0.0001 \pm 0.0061)^\circ\text{C}$	$(0.0026 \pm 0.0086)^\circ\text{C}$	$(0.001 \pm 0.015)^\circ\text{C}$	$(0.0004 \pm 0.0047)^\circ\text{C}$	$(0.0001 \pm 0.0036)^\circ\text{C}$	$(0.0007 \pm 0.0059)^\circ\text{C}$
Flux	$(0.13 \pm 0.22) \text{ l/h}$	$(0.10 \pm 0.25) \text{ l/h}$	$-(0.07 \pm 0.18) \text{ l/h}$	$-(0.05 \pm 0.24) \text{ l/h}$	$-(0.01 \pm 0.21) \text{ l/h}$	$-(0.01 \pm 0.15) \text{ l/h}$
Pressure	$(15 \pm 30)10^{-3} \text{ mbar}$	$-(13 \pm 25)10^{-3} \text{ mbar}$	$(22 \pm 27)10^{-3} \text{ mbar}$	$(1.8 \pm 7.4)10^{-3} \text{ mbar}$	$-(0.8 \pm 1.2)10^{-3} \text{ mbar}$	$(0.7 \pm 1.3)10^{-3} \text{ mbar}$
Radon	$-(0.029 \pm 0.029) \text{ Bq/m}^3$	$-(0.030 \pm 0.027) \text{ Bq/m}^3$	$(0.015 \pm 0.029) \text{ Bq/m}^3$	$-(0.052 \pm 0.039) \text{ Bq/m}^3$	$(0.021 \pm 0.037) \text{ Bq/m}^3$	$-(0.028 \pm 0.036) \text{ Bq/m}^3$
Hardware rate	$-(0.20 \pm 0.18)10^{-2} \text{ Hz}$	$(0.09 \pm 0.17)10^{-2} \text{ Hz}$	$-(0.03 \pm 0.20)10^{-2} \text{ Hz}$	$(0.15 \pm 0.15)10^{-2} \text{ Hz}$	$(0.03 \pm 0.14)10^{-2} \text{ Hz}$	$(0.08 \pm 0.11)10^{-2} \text{ Hz}$

2.2.2 Radioactivity measurements of the detector components

DAMA has published the residual backgrounds or limits thereof in many of their detector components. The backgrounds from various radioactive isotopes in the crystals, PMTs, and shielding are summarized in Table 2.3. The implications of these reported backgrounds are discussed in more detail in Section 2.3.2. In addition, DAMA reports upper limits for neutron backgrounds from various sources in Table 2.4.

Table 2.3: Backgrounds in DAMA/LIBRA. Most backgrounds are from [41], but the backgrounds from the PMTs are from [46]. The method by which each isotope’s concentration is measured is listed in the first column for the crystal backgrounds.

The crystals		
Method	Source	Activity/Concentration
γ counting	^{nat}K	13 ppb ⁷
α counting	^{238}U	0.7–10 ppt ^{8 9}
α counting/Bi-Po	^{232}Th	0.5–7.5 ppt ^{8 9}
unreported	^{87}Rb	$<300 \mu\text{Bq/kg}$
γ counting	^{125}I	unreported/negligible
γ counting	^{129}I	$^{129}\text{I}/^{nat}\text{I} = (1.7 \pm 0.1) \times 10^{-13}$
γ counting	^{210}Pb	$5\text{--}20 \mu\text{Bq/kg}$
coincidence/positron annihilation	^{22}Na	$<\text{tens of } \mu\text{Bq/kg}$
coincidence/ $\beta\gamma\gamma$	^{24}Na	$<0.26 \mu\text{Bq/kg}$
unreported	^3H	$<90 \mu\text{Bq/kg}$
unreported	^{85}Kr	$<10 \mu\text{Bq/kg}$

The PMTs and optical coupling		
Material	Source	Activity/Concentration
Old PMT (ETL)	^{nat}K	1.9 Bq/kg
Old PMT (ETL)	^{238}U	0.37 Bq/kg
Old PMT (ETL)	^{232}Th	0.12 Bq/kg
Old PMT (ETL)	^{60}Co	unreported
New PMT (Hamamatsu)	^{nat}K	0.54 Bq/kg
New PMT (Hamamatsu)	^{238}U	$<3 \text{ Bq/kg}$
New PMT (Hamamatsu)	^{228}Th	83 mBq/kg

Continued on next page...

⁷In [41], DAMA reports only that the maximum crystal concentration of ^{nat}K is 20 ppb, but in later publications, such as [48], they report an average concentration of 13 ppb.

⁸Assumes equilibrium in the chain. For the ^{232}Th chain, DAMA claims this assumption could be valid, since the longest half-life in the chain is ~ 5 years. DAMA observes evidence, however, that the ^{238}U chain is broken. However, they maintain that the assumption of secular equilibrium is a valid first approximation and quote the ^{238}U activity using this assumption. Furthermore, they only report the concentrations of the various parts of the ^{238}U chain for one crystal. Therefore, the equilibrium-assuming values are reported here.

⁹ Concentration is reported to vary by crystal

Material	Source	Activity/Concentration
New PMT (Hamamatsu)	^{60}Co	<25 mBq/kg
Light guide	^{nat}K	<1.3 ppm
Light guide	^{238}U	<1 ppb
Light guide	^{232}Th	<2 ppb
Light guide	^{60}Co	unreported

The housing

Material	Source	Activity/Concentration
Copper	^{nat}K	<0.6 ppm
Copper	^{238}U	<0.5 ppb
Copper	^{232}Th	<1 ppb
Copper	^{60}Co	unreported
Feedthroughs	^{nat}K	<1.8 ppm
Feedthroughs	^{238}U	unreported
Feedthroughs	^{232}Th	<1.6 ppb
Feedthroughs	^{60}Co	unreported
Neoprene o-ring	^{nat}K	<89 ppm
Neoprene o-ring	^{238}U	unreported
Neoprene o-ring	^{232}Th	<54 ppb

The passive shielding

Material	Source	Activity/Concentration
Copper	^{nat}K	<0.6 ppm
Copper	^{238}U	<0.5 ppb
Copper	^{232}Th	<1 ppb
Copper	^{60}Co	unreported
Lead	^{nat}K	<0.06 ppm
Lead	^{238}U	<8 ppb ¹⁰
Lead	^{232}Th	<0.042 ppb ¹⁰
Lead	^{60}Co	unreported
Polyethylene	^{nat}K	<2 ppm
Polyethylene	^{238}U	<0.3 ppb
Polyethylene	^{232}Th	<0.7 ppb
Polyethylene	^{60}Co	unreported
Plexiglass	^{nat}K	<3.3 ppm
Plexiglass	^{238}U	<0.64 ppb
Plexiglass	^{232}Th	<27.2 ppb

Continued on next page...

¹⁰DAMA/LIBRA reports three kinds of lead shielding, but it is unclear what amount of each are used, so the highest limit is reported here.

Material	Source	Activity/Concentration
Plexiglass	^{60}Co	unreported

General		
Method	Source	Activity/Concentration
	Rn	$< 3 \text{ Bq/m}^3$
^{24}Na activation	Thermal n flux	$< 1.2 \times 10^{-7} \text{ cm}^{-2}\text{s}^{-1}$

Table 2.4: DAMA estimated contributions to the modulation signal from neutrons, as reported in [38]. The flux ($\Phi_{0,k}^{(n)}$), the relative modulation amplitude (η_k), and the phase (t_k) of each component is reported in the second column. In the third column, the rate ($R_{0,k}$) and modulation (A_k) for single-hit events in the 2 to 6-keV $_{ee}$ energy region are reported. Finally, the ratio of the modulation to the purported dark-matter modulation, A_k/S_m^{exp} , is reported in the fourth column. In all of these, k stands for the individual component of the neutron signal.

Source	$\Phi_{0,k}^{(n)}$ (neutrons cm $^{-2}$ s $^{-1}$)	η_k	t_k	$R_{0,k}$ (cpd/kg/keV)	$A_k = R_{0,k}\eta_k$ (cpd/kg/keV)	A_k/S_m^{exp}	
SLOW neutrons	thermal n (10 $^{-2}$ – 10 $^{-1}$ eV)	1.08×10^{-6}	$\simeq 0$ however $\ll 0.1$	–	$< 8 \times 10^{-6}$	$\ll 8 \times 10^{-7}$	$\ll 7 \times 10^{-5}$
	epithermal n (eV-keV)	2×10^{-6}	$\simeq 0$ however $\ll 0.1$	–	$< 3 \times 10^{-3}$	$\ll 3 \times 10^{-4}$	$\ll 0.03$
FAST neutrons	fission, (α, n) \rightarrow n (1-10 MeV)	$\simeq 0.9 \times 10^{-7}$	$\simeq 0$ however $\ll 0.1$	–	$< 6 \times 10^{-4}$	$\ll 6 \times 10^{-5}$	$\ll 5 \times 10^{-3}$
	$\mu \rightarrow$ n from rock (> 10 MeV)	$\simeq 3 \times 10^{-9}$	0.0129	end of June	$\ll 7 \times 10^{-4}$	$\ll 9 \times 10^{-6}$	$\ll 8 \times 10^{-4}$
	$\mu \rightarrow$ n from Pb shield (> 10 MeV)	$\simeq 6 \times 10^{-9}$	0.0129	end of June	$\ll 1.4 \times 10^{-3}$	$\ll 2 \times 10^{-5}$	$\ll 1.6 \times 10^{-3}$
	$\nu \rightarrow$ n (few MeV)	$\simeq 3 \times 10^{-10}$	0.03342*	Jan. 4th*	$\ll 7 \times 10^{-5}$	$\ll 2 \times 10^{-6}$	$\ll 2 \times 10^{-4}$
	direct μ	$\Phi_0^{(\mu)} \simeq 20 \mu \text{ m}^{-2}\text{d}^{-1}$	0.0129	end of June	$\simeq 10^{-7}$	$\simeq 10^{-9}$	$\simeq 10^{-7}$
	direct ν	$\Phi_0^{(\nu)} \simeq 6 \times 10^{10} \nu \text{ cm}^{-2}\text{s}^{-1}$	0.03342*	Jan. 4th*	$\simeq 10^{-5}$	3×10^{-7}	3×10^{-5}

2.2.3 Processing the data

DAMA refrains from using certain active background-rejection techniques like pulse-shape discrimination in order to avoid rejecting WIMP signals. However, DAMA does employ some techniques to reject photomultiplier noise and other backgrounds. First, they require that both PMTs in a crystal module register a signal within a certain time window (50 or 100 ns [41]). DAMA rejects events with a short decay time, which helps to mitigate PMT noise. PMT noise produces a faster pulse than true scintillation, as shown in Figure 2.9. DAMA cuts on two pulse-shape parameters, equivalent to the F50 and F100 parameters¹¹ using the nomenclature in this work [41]. The parameters used by DAMA/LIBRA are called X_1 and X_2 and are defined as follows:

$$X_1 = \frac{\text{Area from 100 to 600 ns}}{\text{Area from 0 to 600 ns}}, \quad (2.3)$$

$$X_2 = \frac{\text{Area from 0 to 50 ns}}{\text{Area from 0 to 600 ns}}. \quad (2.4)$$

A histogram of electronic noise events along with true scintillation signals is shown in Figure 2.10. In the figure, the X_1 - X_2 plane is shown for a normal operation (left) and with a γ source that makes the scintillation events dominant (right) for the 2–4 (top) and 4–6 keV_{ee} (bottom) regions. Though there is good separation between the two types of events, there is some overlap. DAMA does not provide the values of the cuts, nor do they quantify the rejection power of this technique. However, they base an upper limit for the modulation fraction due to this effect of $\mathcal{O}(10^{-4})$ based on the modulation in the total hardware rate and the assumption that the distribution of noise events in X_1 and X_2 do not change, as well as the assumption that noise events constitute less than 10% of the signal after cuts [49].

These techniques can reject random PMT noise, but will not eliminate backgrounds from the dynode-afterglow effect. Dynode afterglow is a phenomenon which is believed to be caused by light emission from the PMT dynodes that can reach a second light-coupled

¹¹Here, $F[X]$ refers to the fraction of the pulse integral that occurs in the first $[X]$ ns.




Figure 2.9: Pulses for an event near the 2 keV_{ee} threshold for an electronic noise event (Top) and a scintillation event (Bottom). Noise pulses tend to have a sharper time profile, while scintillation events are more spread out.

tube, thus producing a false coincidence. These events were seen in the NAIAD NaI(Tl) experiment, and it was found that they were not rejected adequately by looking at the scintillation time-constant [50]. These types of events could play a role in DAMA/LIBRA as well.

DAMA also rejects “multi-hit events,” in which a scintillation event occurs in multiple crystals. Because some backgrounds can travel between detectors and interact twice or in coincidence with other backgrounds, but WIMP dark matter interacts only rarely, an event that is registered in multiple detectors is extremely likely to be background and can be

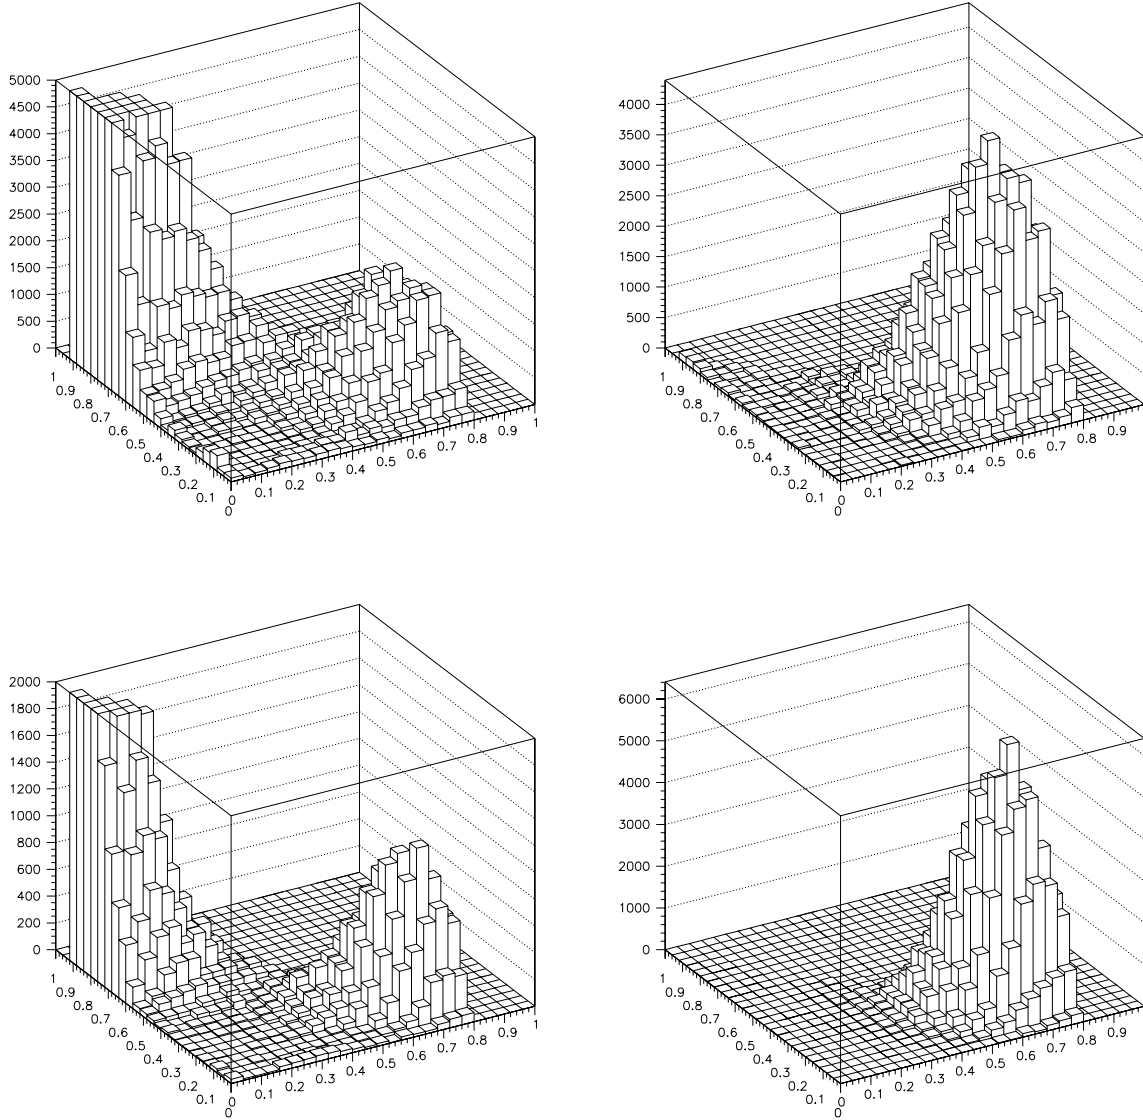


Figure 2.10: Electronic noise rejection in DAMA/LIBRA. The figures above show the X_1 (Front axis) and X_2 (Side axis) pulse shape parameters. Because noise events are faster than scintillation events, the noise events appear in the top left corner, where X_1 is small and X_2 is large, while scintillation events have a larger X_1 and a smaller X_2 . The left plots show single-hit events in DAMA's normal operation mode, while the right plots are for a run with a γ source, causing the scintillation events to be dominant. The top plots are for the 2–4 keV_{ee} region, while the bottom plots are for the 4–6 keV_{ee} region.

safely rejected. Afterpulses resulting from Cherenkov light or Bi-Po events¹² are rejected by requiring 500 μ s of dead time after each event.

DAMA reports that the hardware threshold is at the single-photoelectron level, which corresponds to ~ 0.2 keV_{ee}. However, to lower the incidence of PMT noise, they use a software energy threshold of 2 keV_{ee}. With the use of their new high-quantum-efficiency phototubes, DAMA claims that they will lower this software energy threshold to 1 keV_{ee} [39, 46].

DAMA reports the following energy resolution for their crystal detectors, which they measure periodically with ²⁴¹Am sources [41, 39]:

$$\frac{\sigma_E}{E} = \frac{\alpha}{\sqrt{E}} + \beta \quad (2.5)$$

where E is the energy in keV_{ee}, $\alpha = 0.448 \pm 0.035$ and $\beta = (9.1 \pm 5.1) \times 10^{-3}$.

The overall efficiency curve for the DAMA/LIBRA experiment is shown in Figure 2.11. The efficiency is measured periodically with an ²⁴¹Am source. This curve is used to correct their event rate for a systematic loss of low-energy events by the data acquisition system when producing the plots in Figure 2.5. DAMA attributes the loss of efficiency primarily to the coincidence time window and the pulse-timing cuts. It is notable that the loss of efficiency occurs primarily in the modulation energy region; if the cuts do not completely eliminate electronic noise, annual modulation effects in electronic noise could affect their signal.

DAMA measures the linearity of the PMT response at low energies (< 100 keV_{ee}) by measuring the light response of internal backgrounds at known energies, as shown in Figure 7.19.

¹²The near-simultaneous decay of bismuth and polonium in the ²³⁸U and ²³²Th decay chains, producing a β and an α in rapid succession.




Figure 2.11: DAMA/LIBRA’s reported detection efficiency for the old ETL tubes (black triangles) and the new Hamamatsu tubes (white circles). This detection-efficiency curve is taken into account in DAMA’s published energy and modulation spectra.

2.2.4 The dark-matter claim

On the basis of the modulation described previously, DAMA claims a positive, model-independent discovery of dark-matter particles interacting in their detector. There are enough unknown astrophysical and experimental parameters that a single, standard WIMP model cannot be used to definitively analyze their modulation. However, if one assumes the SHM and the standard vanilla WIMP with a spin-independent interaction, the DAMA modulation best fits a ~ 80 -GeV WIMP interacting primarily with iodine, with a secondary maximum at ~ 10 GeV interacting primarily with sodium. The former falls into the expected mass window for supersymmetric WIMP dark matter. The second is what would be catego-

rized as a “light WIMP.” A derivation and analysis of the DAMA parameter space regions is discussed in Chapter 8.

2.3 Controversy over the DAMA/LIBRA result

The DAMA/LIBRA modulation result has the highest statistical significance of any dark-matter experiment, yet their claim for a dark-matter discovery has not been verified conclusively by any other experiment. Significant controversy surrounds the DAMA/LIBRA result, which comes from a number of sources.

First, the published data from DAMA/LIBRA raises some questions about the energy spectrum and modulation that have not been fully addressed by DAMA. For any experiment claiming a discovery, it is crucial to have a firm understanding of the data, including backgrounds. Attempts to understand the DAMA/LIBRA detector response and backgrounds are discussed in Sections 2.3.1–2.3.5.

Second, while some experiments have observed signals commensurate (to some extent) with the DAMA/LIBRA signal, the DAMA/LIBRA result is incompatible with many experiments when using standard WIMP and astrophysical models. Though there still exist some limited models that could reconcile this conflict, other experiments are establishing more and more sensitive dark-matter limits, and it is becoming more difficult with time to reconcile the DAMA/LIBRA result with other null experiments. However, comparing DAMA with these null results is difficult because of their use of different targets. Attempts to reconcile DAMA/LIBRA with other experimental results are discussed in Section 2.3.7.

2.3.1 The energy spectrum

The DAMA/LIBRA energy spectrum, shown in Figure 2.5, has a number of important features that complicate the picture or are difficult to reconcile with a dark-matter explanation. The first notable feature is the bump at 3 keV_{ee} due to ^{40}K decay by electron capture,

which obscures the region of interest for the purported dark-matter modulation if the ^{40}K background is not well understood. Secondly, the spectrum is approximately constant from 4–10 keV $_{ee}$, indicating a smooth, near-constant background characteristic of β -decay in a region that is small compared to the end-point energy. Critics of DAMA argue that if one assumes that this near-constant background continues into the modulation region, there is no room for a WIMP signal, or that the modulation fraction would have to be so large as to be inconsistent with most dark-matter models.

Another issue has to do with the shape of the spectrum from 4–6 keV $_{ee}$, where the modulation is still observed. The total spectrum is nearly constant over this region. As discussed in Section 1.3.2, the expected WIMP rate is expected to fall roughly exponentially with energy. Though the modulation is small, if the modulation fraction is also small, the WIMP total rate will be a significant portion of the signal and will also fall rapidly over this energy range. In order to preserve the flatness of the total rate over this interval, the background spectrum would have to conveniently rise in proportion to the falling off of the WIMP signal when increasing from 4 to 6 keV $_{ee}$ in energy. If the modulation fraction is very high, however, this may not be a problem.

DAMA warns against any attempt to predict their background signal below 10 keV $_{ee}$, citing the unreliability of Monte Carlos at these energies and the unfeasibility of knowing the quantities or distributions of backgrounds in the detectors [41, 39]. Instead of trying to replicate the shape of their backgrounds with simulations, DAMA extrapolates from the slight positive slope of the rate between 4 and 10 keV $_{ee}$ down into the region that is obscured by the ^{40}K peak. They subtract a line and gaussian (for a 13 ppb ^{40}K peak) from their data, which they claim leaves room for a flat dark-matter rate of 0.25 cpd/kg/keV $_{ee}$ between 2 and 4 keV $_{ee}$ [48, 51, 38]. By contrast, [37] suggests the flat background in the modulation region is indicative of a β^- or Compton spectrum at the level of 0.85 cpd/kg/keV $_{ee}$. From this assumption, they conclude that the purported WIMP event rate would have to be lower than that reported by DAMA.

The contributions to the background spectrum below 10 keV_{ee} may be better understood if characteristic signals above 10 keV_{ee} could be used as input for a fit. Each background has its own characteristic spectrum, but many of the backgrounds reported in Section 2.2.2 have nearly indistinguishable flat backgrounds below 10 keV_{ee} . DAMA only publishes the spectrum between 20 and 80 keV_{ee} for two crystals (see Figure 2.12) and the α -event spectra at higher energies for four others (Figure 2.13). However, DAMA reports a wide variation in the radioactivity of different crystals, so these spectra are not sufficient for determining the average impurity levels for different isotopes.

Furthermore, it is notable that while the spectrum is very flat around $1 \text{ cpd/kg/keV}_{ee}$ between 4 and 10 keV_{ee} , the rate is only about $0.5 \text{ cpd/kg/keV}_{ee}$ between 20 and 30 keV_{ee} for the two crystals shown in Figure 2.12. Finally, as can be seen in Table 2.3, DAMA often reports only limits or ranges for detector contaminants, sometimes not even that, and the limits reported can sometimes produce very large background signals, as is the case with ^{87}Rb and ^3H in the crystals or the backgrounds from the fused silica light-guides.

Despite the lack of published spectra for all crystals at all energies and better estimates from contaminant concentrations in the detectors, attempts have been made to simulate the DAMA background spectrum in order to establish what the dark-matter signal might look like. In [44], a Monte Carlo simulation was done to measure the DAMA/LIBRA background spectrum and assess what rate is left over for dark matter. Initially, they used the DAMA reported contamination values reported in [41], but found that the spectra did not match the data well at high energies. They instead used a best fit to determine crystal contamination values of ^{nat}K , ^{238}U , and ^{232}Th of 20 ppb, 34 ppt, and 20 ppt respectively, but in parts of the low energy region these concentrations didn't leave room for a dark-matter signal. In another approach, they tried subtracting a best-fit WIMP from the DAMA/LIBRA energy spectrum to see what the resultant background spectrum would look like, as shown in Figure 2.14. In order to produce the expected WIMP spectrum, the backgrounds would have to have a steep drop in rate below 2 keV_{ee} .




Figure 2.12: DAMA/LIBRA energy spectrum for two crystals from 20–80 keV_{ee}, from [41], along with their fit with a Monte Carlo of the background from ¹²⁹I and ²¹⁰Pb (the solid colored lines). The sharp rise in the spectra at around 40 keV_{ee} is due to ²¹⁰Pb and ¹²⁹I. The rate below 30 keV_{ee} is significantly lower than the average values from 6–10 keV_{ee}. Either these crystals have atypically low backgrounds compared with the others, or the flat background drops significantly between 10 and 20 keV_{ee}, yet is flat below 10 keV_{ee} and from 20–30 keV_{ee}.

2.3.2 Simulating the DAMA/LIBRA background spectrum

In order to better understand the DAMA/LIBRA spectrum and the ability of Monte Carlo to simulate backgrounds for NaI(Tl) dark-matter experiments, a simulation of the DAMA/LIBRA background spectrum was performed in the same manner as described in Section 5.4. A previous attempt to simulate the DAMA/LIBRA energy spectrum by [44] used an old version of the Geant4 simulation software (4.9.2) to perform their background studies of DAMA/LIBRA. This version has known issues with energy conservation at the keV level. The simulations of the DAMA/LIBRA detector backgrounds described here were performed with the newest version of Geant4 (4.10.01). There are two main differences between the new simulations with respect to those in [44]. The first is the way Geant4 handles atomic relaxation. In 4.9.2, Geant randomly generated Auger-electron and X-ray energy values from a table. If the excited-state energy is greater than the emitted radiation,




Figure 2.13: DAMA α spectra at high energies, used to determine ^{238}U and ^{232}Th concentrations in the crystals, from [41]

the remaining energy is deposited locally in order to conserve energy. The newest version of Geant4 simulates the entire cascade, including fluorescence, particle-induced X-ray emission (PIXE), and Auger-electron emission, but only keeps track of the original hole in simulating the cascade, thus not conserving energy. The code was altered to track all holes generated with an `std::vector`, thus conserving energy.

The second difference can be found in how Geant4 handles the decay of ^{228}Ra , a β -emitter that is especially relevant to the ^{232}Th chain. In this decay, ^{228}Ra decays to an excited state of ^{228}Ac , the lowest of which is around 6 keV. The resultant de-excitation of the nucleus should contribute to the energy deposited in the crystal. In previous simulations, the ^{232}Th spectrum produced a steadily rising background with energy, whereas for the new simulations, the rate is non-existent below 6 keV_{ee} , matching better with physical expectations. The new

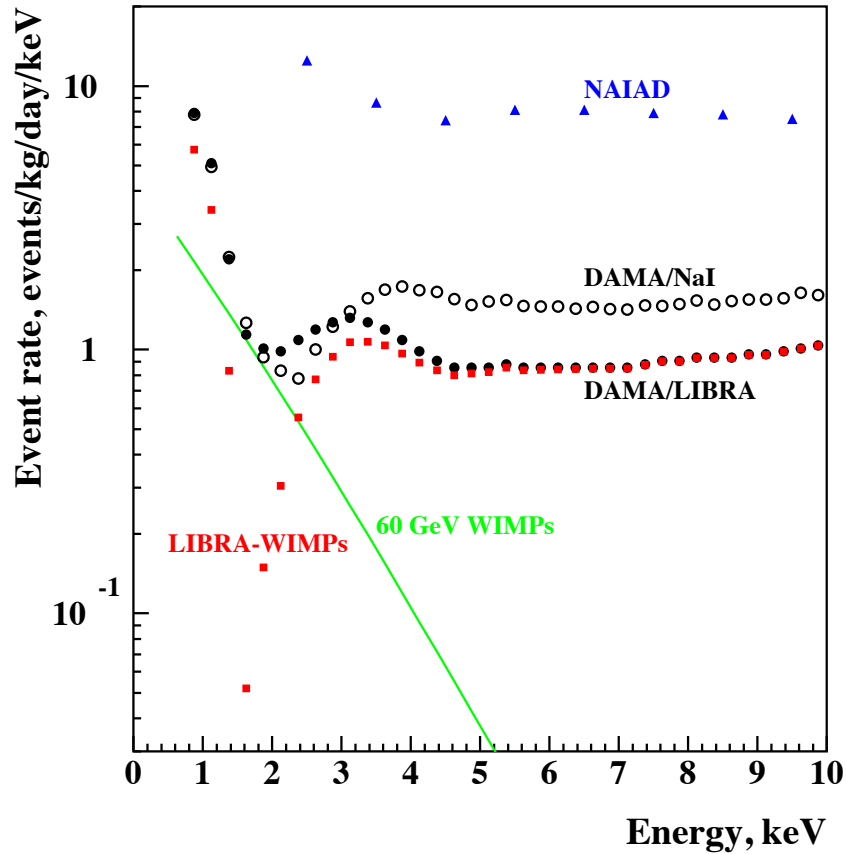


Figure 2.14: DAMA/LIBRA energy spectrum with a 60-GeV WIMP with 10^{-6} pb cross section subtracted (red squares). Also shown is the WIMP spectrum in green, the DAMA and DAMA/LIBRA energy spectra in white and black circles, respectively, and the NAIAD NaI(Tl) detector spectrum in blue. The step drop in the space for the inferred DAMA/LIBRA backgrounds around 2 keV_{ee} shows how it is difficult to account for the DAMA backgrounds in a typical WIMP picture. Plot from [44].

spectrum can help to mitigate the problem wherein the background spectrum must decrease below 6 keV_{ee} in order to make room for a WIMP signal, but only if backgrounds from ^{232}Th are high.

The DAMA/LIBRA detector was simulated using the detector geometry described in [41]. DAMA does not publish all of the details for their apparatus, however, so some details like the thickness of the copper enclosures had to be assumed. For this simulation, a 2-mm thickness was used. Furthermore, DAMA expresses that the contamination varies widely from crystal to crystal. If the crystals have very different contamination levels, it matters where they are in the array. In addition, the contaminants probably also vary within the crystal, since contaminants are generally segregated during the crystal-growth process. In the absence of any of this information from DAMA, this simulation assumes a uniform distribution of contaminants throughout the crystal volume and between crystals. This means that surface background effects were not taken into account.

Based on γ -counting measurements and measurements of the α spectra from their detectors, DAMA reports that the ^{232}Th chain is in secular equilibrium, but the ^{238}U chain is not. Given, however, that the ^{238}U chain produces a featureless background at energies below 10 keV_{ee} , and because DAMA reports a background-level range for ^{238}U assuming secular equilibrium as “a first approximation,” [41] this simulation assumes secular equilibrium for both chains. Backgrounds were simulated in the DAMA crystals, the Suprasil light-guides, the PMTs, the copper enclosures (including the shield around the PMTs), and the external γ -ray background from the LNGS rock. Individual isotopes in the crystal are shown in Figures 2.15 and 2.16. The former shows the traditionally considered isotopes of ^{40}K , ^{238}U , ^{232}Th , and ^{87}Rb , while the latter shows the other isotopes reported by DAMA. The total crystal background with and without backgrounds for which only limits were reported is shown in Figure 2.17, forming an upper and lower limit for the DAMA crystal background rate. The individual and total backgrounds for the fused-silica light-guides using the limits reported by DAMA are shown in Figures 2.18 and 2.19, respectively. The individual and

total PMT backgrounds for the DAMA ETL PMTs are shown in Figures 2.20 and 2.21, while the individual and total enclosure backgrounds at the limits reported by DAMA are shown in Figures 2.22 and 2.23. A limit for the external γ background from the LNGS concrete wall of <0.004 cpd/kg/keV $_{ee}$ from 0–10 keV $_{ee}$ was achieved in the simulation. No events in this energy range were observed. Finally, upper and lower limits of the total DAMA/LIBRA simulated background spectrum, along with the DAMA/LIBRA single-hit spectrum, are shown in Figure 2.24 with their reported values.

One key background that can be fit to the data in the 1–10-keV $_{ee}$ energy region is the 3-keV $_{ee}$ signature from ^{40}K . DAMA claims that the signal is partially rejected by the rejection of multi-hit events. This simulation confirmed that this technique does diminish the ^{40}K crystal background somewhat; assuming a uniform ^{40}K background in the crystals, a rejection power of 58% was observed, as shown in Figure 2.25. Like in [44], it was observed that the DAMA quoted ^{nat}K value of 13 ppb did not fit well with the observed peak in the spectrum (see Figure 2.24). A fit allowing the crystal radioactivity values to move over the range 1.5–5.5 keV $_{ee}$ gave a best-fit value of 23.2 ppb for ^{nat}K , similar to the value reported by [44]. This value, along with other best-fit values, is shown with the DAMA/LIBRA spectrum in Figure 2.26.

The radioactivity limits reported by DAMA for a few of their backgrounds, particularly ^{87}Rb and ^3H , would constitute a large portion of their signal if the values were at the reported limit, as shown in Figure 2.16. The ^3H background is particularly interesting, given that ^3H decays with a 100% branching ratio by β^- decay with an endpoint energy of 18.6 keV. DAMA does not report their average spectrum above 10 keV $_{ee}$, but they do show the spectra for two of their crystals in the 20–80 keV $_{ee}$ range, as shown in Figure 2.12. There, the flat background appears to be half of their average flat background from 6–10 keV $_{ee}$. If the background in these two crystals is typical, a ^3H background close to the reported limit could explain the difference between the ~ 1 cpd/kg/keV $_{ee}$ flat background below 10 keV $_{ee}$ and the ~ 0.5 cpd/kg/keV $_{ee}$ background above 20 keV $_{ee}$. A ^3H background of 90 $\mu\text{Bq/kg}$

would produce an observed difference of $0.34 \text{ cpd/kg/keV}_{ee}$. A best-fit to the flat part of the spectrum between $6\text{--}10 \text{ keV}_{ee}$ predicts a ${}^3\text{H}$ background of around $40.3 \mu\text{Bq/kg}$, which is too low to produce the observed rate difference, though given the poor quality of the fit, the uncertainty on this value is very high.

The slight positive slope of the featureless background is also of concern. DAMA’s flat part of the spectrum rises from 6 to 10 keV_{ee} . The crystal backgrounds, which are dominant at the values reported, all fall over this energy region. The PMT, light-guide, and enclosure backgrounds, by contrast, rise over this region. The conflict is especially notable if the ${}^3\text{H}$ background is high. The rise in rate from $6\text{--}10 \text{ keV}_{ee}$ could suggest a higher background from the PMTs, light-guides, and enclosures than what is reported by DAMA.

In addition to the shape of the flat region, it is difficult to fit the $6\text{--}10 \text{ keV}_{ee}$ region while also leaving room for a WIMP signal at low energies. If the ${}^{232}\text{Th}$ background, which disappears below 5 keV_{ee} , constitutes a large part of the signal, there may be room for a WIMP signal, but at the levels reported by DAMA, the ${}^{232}\text{Th}$ background only constitutes $\sim 3\%$ of the total rate.

To better quantify the backgrounds, an attempt was made to fit to the spectrum. In these fits, the adjustable parameter was the scale of the simulated spectra for each individual background. The ${}^{40}\text{K}$ background was fit in the $1.5\text{--}5.5 \text{ keV}_{ee}$ window, allowing other backgrounds to vary along with an exponentially decaying WIMP spectrum, since the WIMP signal was expected to appear in this region. A least- χ^2 method was used, allowing all of the backgrounds in Table 2.5 to vary, with the exception of those marked “fixed.” This fit was solely to find a best-fit value for the ${}^{40}\text{K}$ peak. The ${}^{40}\text{K}$ concentration matched 23.2 ppb better than the reported 13 ppb ; a similar value of 20 ppb was found in [44].

The other backgrounds were then fit in the $6\text{--}10 \text{ keV}_{ee}$ region with no WIMP value, holding the ${}^{40}\text{K}$ concentration fixed at the best-fit value from the previous fit. The quality of the fit was very poor, in part because of the rise in the DAMA signal rate with energy in the $6\text{--}10 \text{ keV}_{ee}$ interval that is not seen in the simulated component backgrounds. The

results of the fit are nevertheless summarized in Table 2.5. In order to achieve a background level at the rate seen in the 6–10 keV_{ee} window, most of the backgrounds had to be added at their upper limits, suggesting that there may be unaccounted-for backgrounds that were not mentioned in [41].

An attempt was also made to calculate a WIMP energy spectrum by subtracting the simulated backgrounds from the DAMA data. The residual WIMP spectrum was generated both for DAMA’s reported values (the upper and lower limits) and for the best-fit values in Table 2.5. This was done to evaluate the potential total WIMP rate and the shape of the WIMP spectrum that is allowed given the simulated DAMA backgrounds. The results are shown in Figures 2.24 and 2.26. The residual spectrum highlights some of the previously mentioned problems—in particular, the residual WIMP signal rises with energy above 6 keV_{ee}. This behavior does not match the observed modulation amplitude as a function of energy, nor the expectations for a WIMP signal. Additionally, the best-fit WIMP residual spectrum doesn’t allow much room for a WIMP in the 2–6 keV_{ee} region if the spectrum is fit in the 6–10 keV_{ee} region.

The challenges presented by the characteristics of the DAMA/LIBRA single-hit spectrum are difficult to overcome. A good fit for the DAMA/LIBRA singles spectrum cannot be obtained without more information about the contamination levels in their detector components. A release of the spectra for each crystal at all energies, or better estimates of the backgrounds with more sensitive limits, are necessary to calculate an expected background spectrum. In addition, there may be some issues with the way that Geant4 handles energy deposits at the keV-level, which will be further elucidated in Section 5.4. In addition to the additional information needed from the DAMA collaboration, further work will need to be done to verify the behavior of the Monte Carlo simulations at these low energies.

Table 2.5: Background levels used in the simulation of the DAMA single-hit energy spectrum. The values in parentheses in the third column were the values used as an “average” for the windows reported by DAMA. The values in the fourth column were fit in the 6–10 keV_{ee} region and allowed to vary within the ranges given by DAMA, except where noted. In the fit, most of the backgrounds went to the upper limit that was imposed. ⁴⁰K in the crystal was fit to the 3-keV_{ee} feature in a window from 1.5–5.5 keV_{ee}; the error shown is the statistical error from the fit only. Overall, the χ^2 from the fit was very poor. The fact that most of the parameters fell at the limit also makes the fit suspect.

Background volume	Isotope	DAMA reported radioactivity	Best fit value
crystal	⁴⁰ K	13 ppb ^{nat} K	23.2 ppb \pm 0.6 ^{nat} K
crystal	²³² Th	0.5–7.5 ppt (3 ppt)	7.5 ppt
crystal	²³⁸ U	0.7–10 ppt (5 ppt)	10 ppt
crystal	⁸⁷ Rb	<0.35 ppb	0.35 ppb
crystal	²¹⁰ Pb	5–20 μ Bq/kg	20 μ Bq/kg
crystal	⁸⁵ Kr	<10 μ Bq/kg	10 μ Bq/kg
crystal	¹²⁹ I	0.17×10^{-12} ¹²⁹ I/ ^{nat} I	(fixed)
crystal	²² Na	<10 μ Bq/kg	10 μ Bq/kg
crystal	²⁴ Na	<0.26 μ Bq/kg	0.036 μ Bq/kg
crystal	³ H	<90 μ Bq/kg	40.3 μ Bq/kg
pmt	⁴⁰ K	1.9 Bq/kg	(fixed)
pmt	²³² Th	0.12 Bq/kg	(fixed)
pmt	²³⁸ U	0.37 Bq/kg	(fixed)
light guide	⁴⁰ K	<1.3 ppm ^{nat} K	1.3 ppm ^{nat} K
light guide	²³² Th	<2 ppb	2.0 ppb
light guide	²³⁸ U	<1 ppb	1.0 ppb
copper	⁴⁰ K	<600 ppb ^{nat} K	600 ppb ^{nat} K
copper	²³² Th	<1 ppb	1 ppb
copper	²³⁸ U	<500 ppt	500 ppt
copper	⁶⁰ Co	unreported	1 mBq/kg (upper limit)

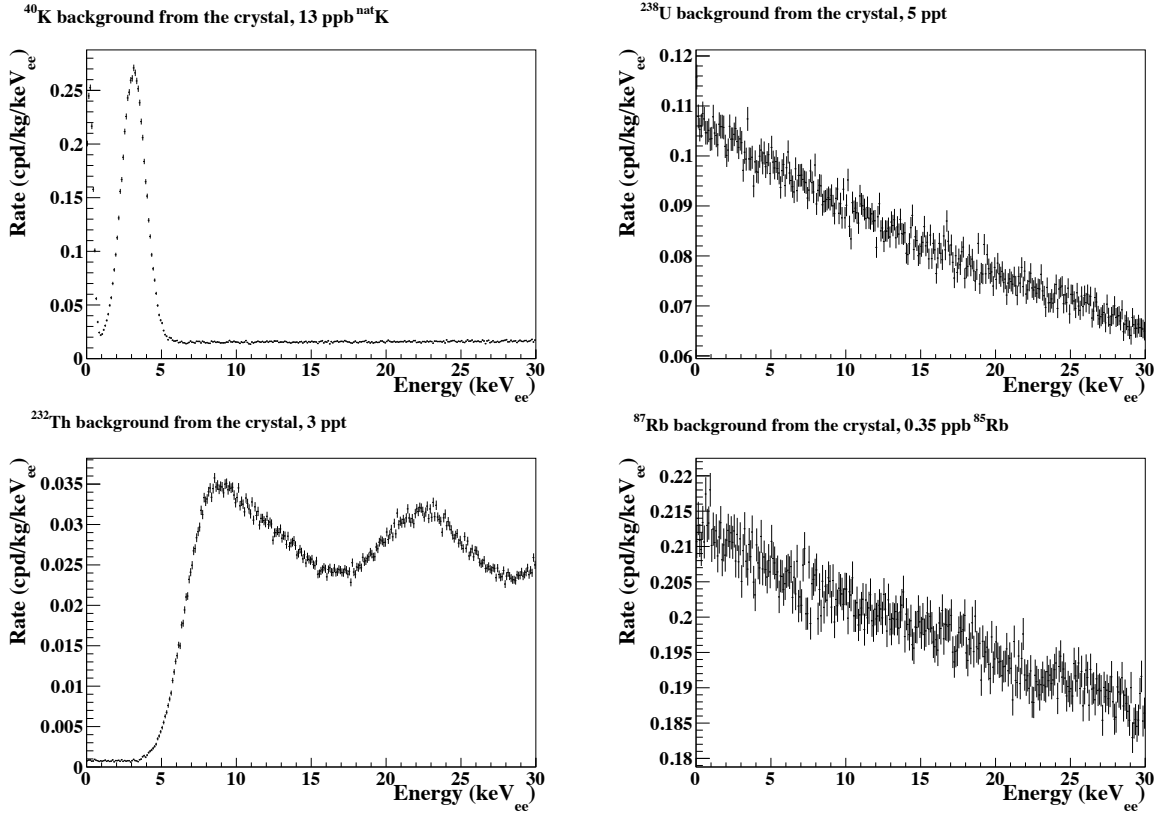


Figure 2.15: Simulated crystal backgrounds from ^{40}K (Top left, 13 ppb), ^{238}U (Top right, 5 ppt), ^{232}Th (Bottom left, 3 ppt), and ^{87}Rb (Bottom right, 0.35 ppb). The values for ^{40}K , ^{238}U , and ^{232}Th are representative of the DAMA/LIBRA reported averages, while ^{87}Rb was simulated at the reported limit. The ^{238}U and ^{232}Th chains were assumed to be in secular equilibrium.

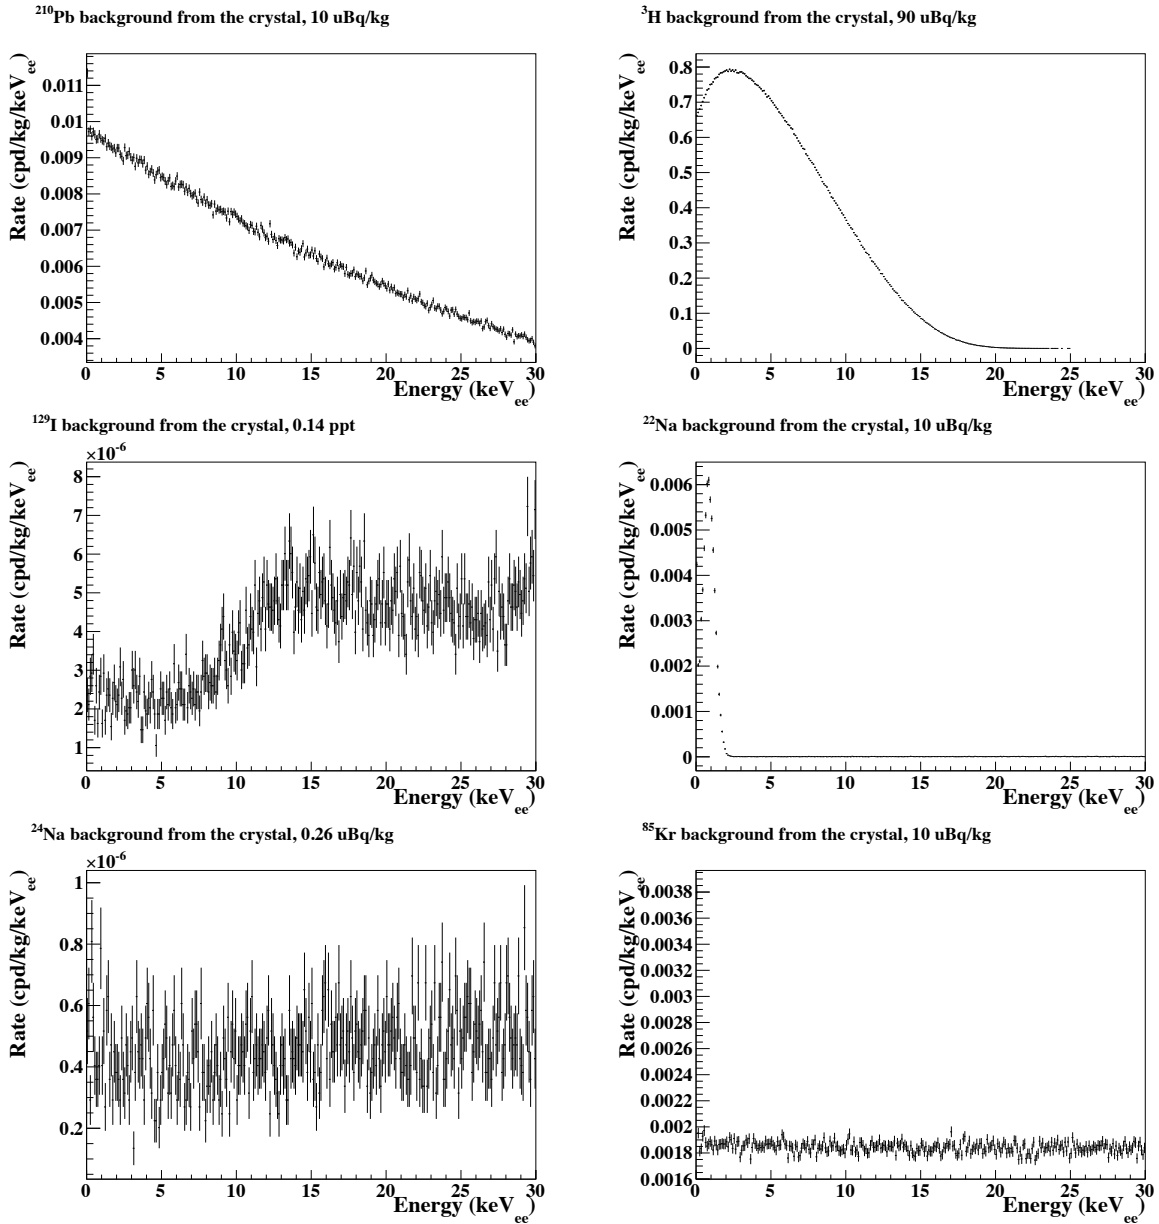


Figure 2.16: Simulated backgrounds from other individual isotopes in the DAMA crystals. (Top left) ^{210}Pb , 10 $\mu\text{Bq/kg}$; (Top right) ^3H , 90 $\mu\text{Bq/kg}$; (Center left) ^{129}I , 0.14 ppt; (Center right) ^{22}Na , 10 $\mu\text{Bq/kg}$; (Bottom left) ^{24}Na , 0.26 $\mu\text{Bq/kg}$; (Bottom right) ^{85}Kr , 10 $\mu\text{Bq/kg}$. The reported concentrations of ^{210}Pb and ^{129}I were used, while all other backgrounds were simulated at their reported limits.

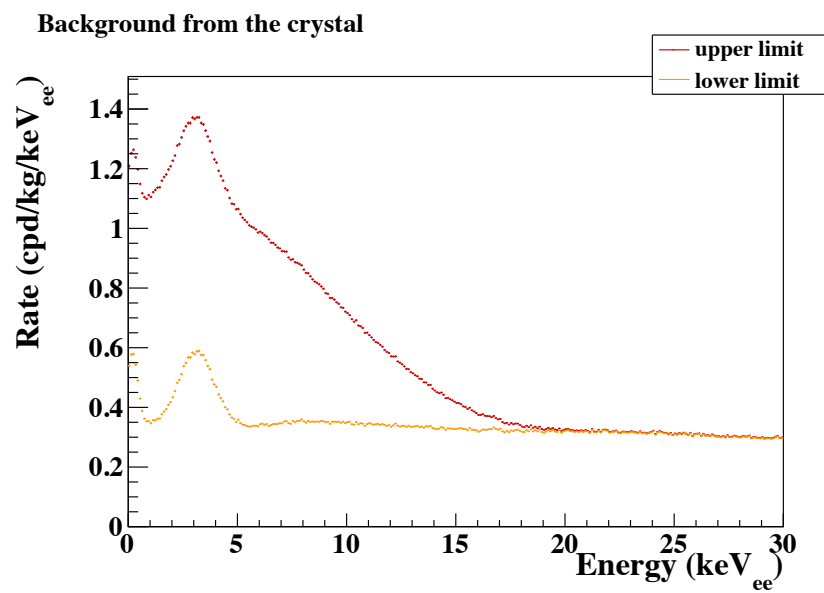


Figure 2.17: Simulated total backgrounds from the DAMA crystals, using DAMA reported values and limits. The red curve includes all backgrounds, even those for which only limits are reported. The orange curve ignores those backgrounds for which only limits are reported. The actual spectrum should be somewhere between these two curves.

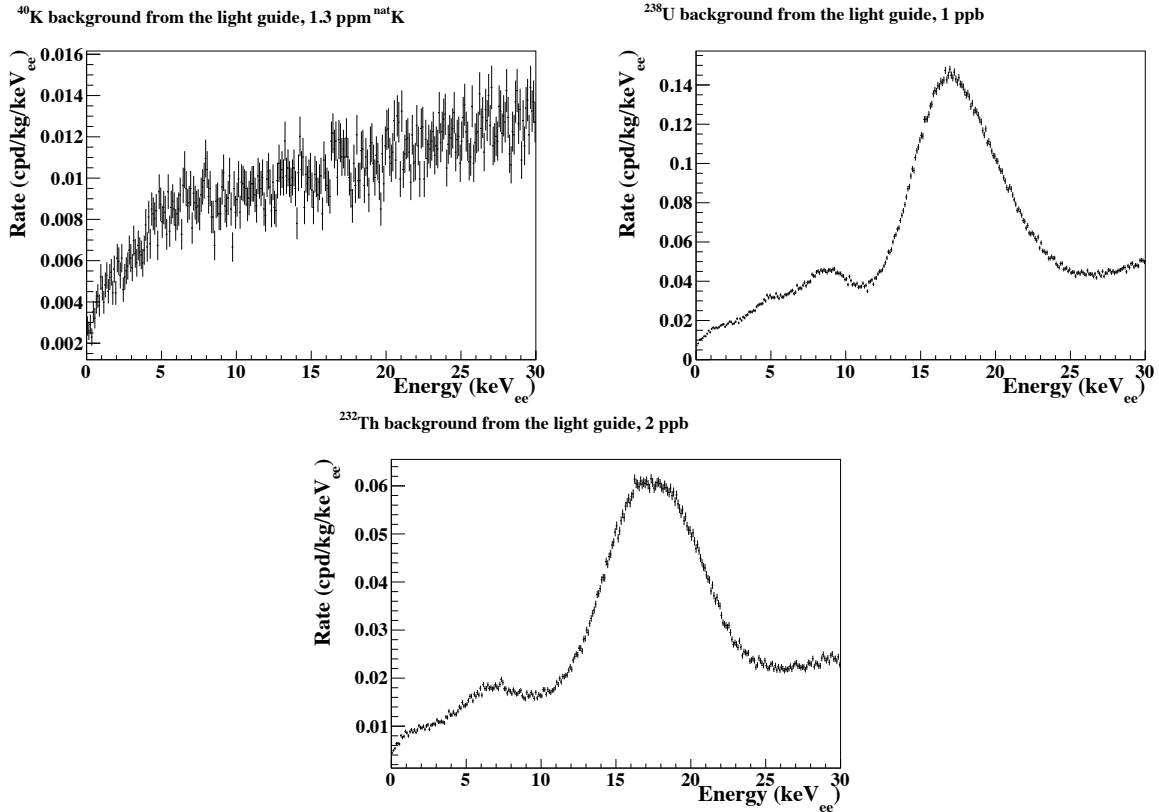


Figure 2.18: Simulated backgrounds from individual isotopes in the DAMA/LIBRA fused-silica light-guides. (Top left) ^{nat}K, 1.3 ppm; (Top right) ²³⁸U, 1.0 ppb; (Bottom) ²³²Th, 2.0 ppb.

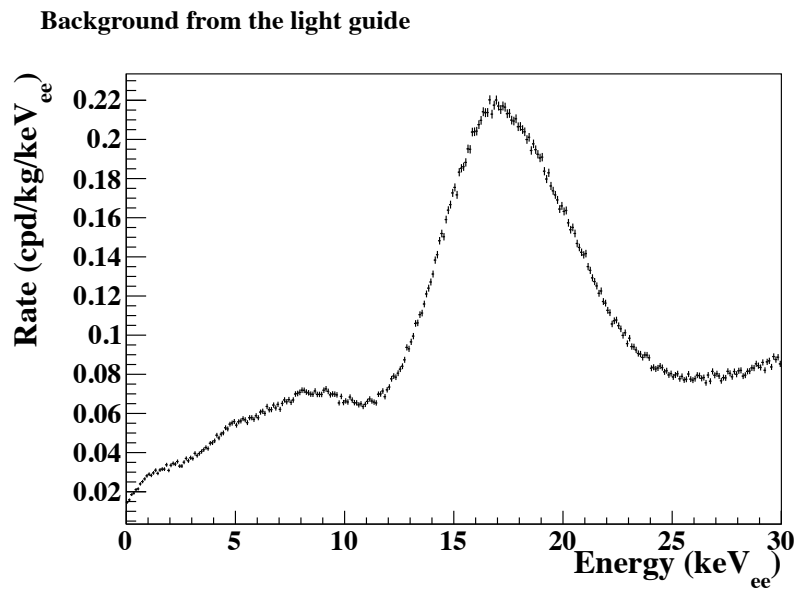


Figure 2.19: Simulated total backgrounds from the DAMA fused-silica light-guides, using the limits from [41].

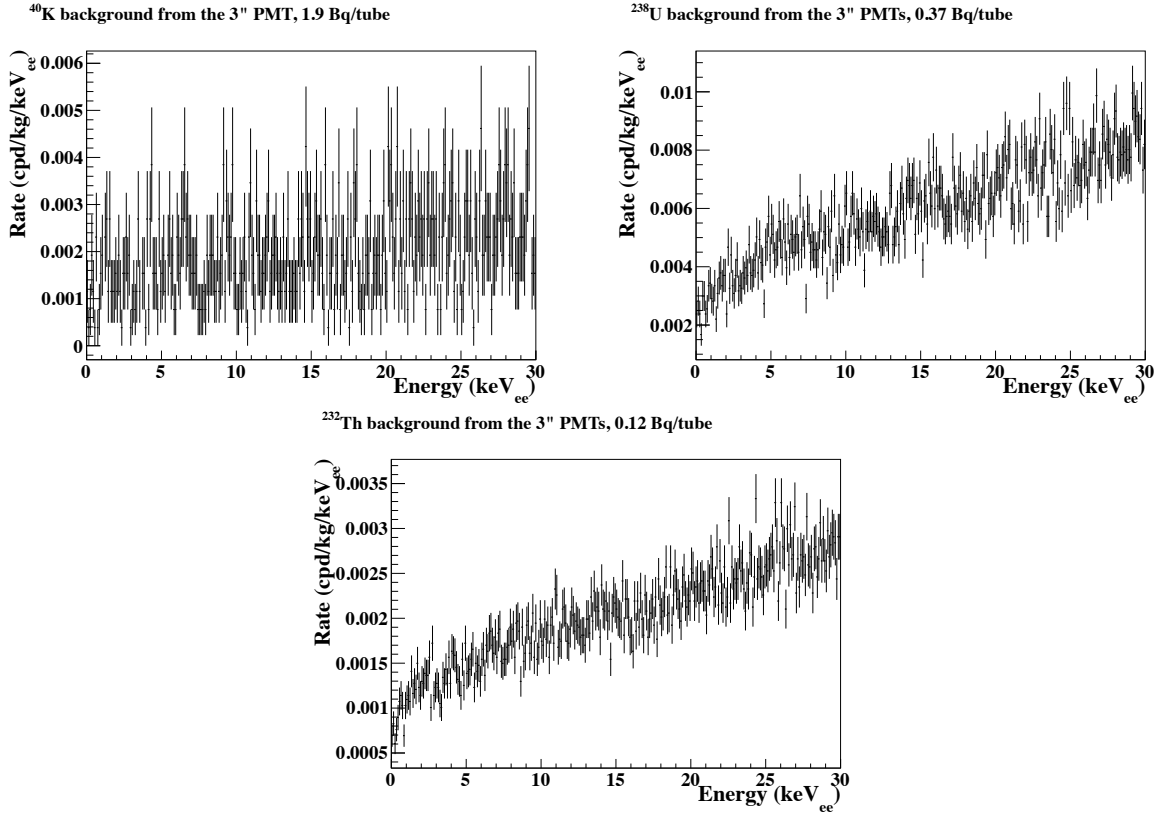


Figure 2.20: Simulated backgrounds from individual isotopes in the DAMA old ETL PMTs. (Top left) ⁴⁰K, 1.9 Bq/kg; (Top right) ²³⁸U, 0.37 Bq/kg; (Bottom) ²³²Th, 0.12 Bq/kg. The ⁶⁰Co background is not reported by DAMA.

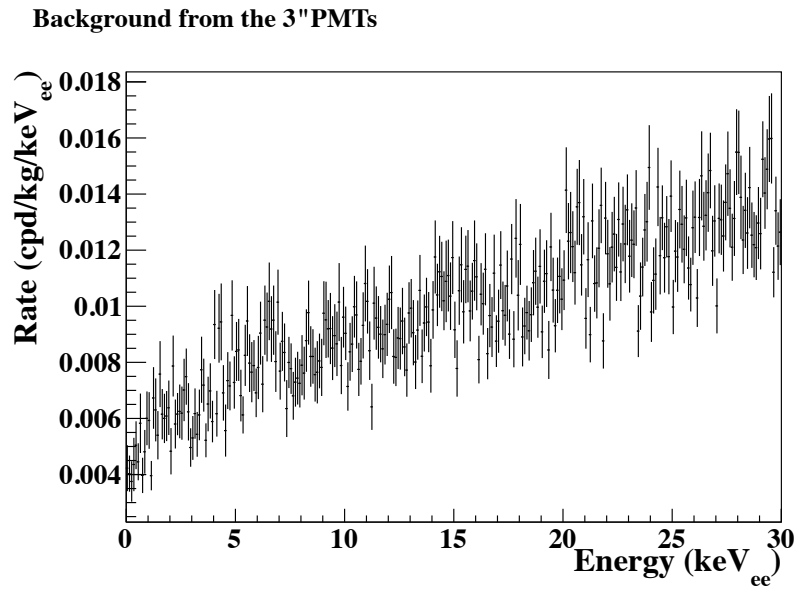


Figure 2.21: Simulated total backgrounds from the DAMA old ETL PMTs, using the activities reported in [46].

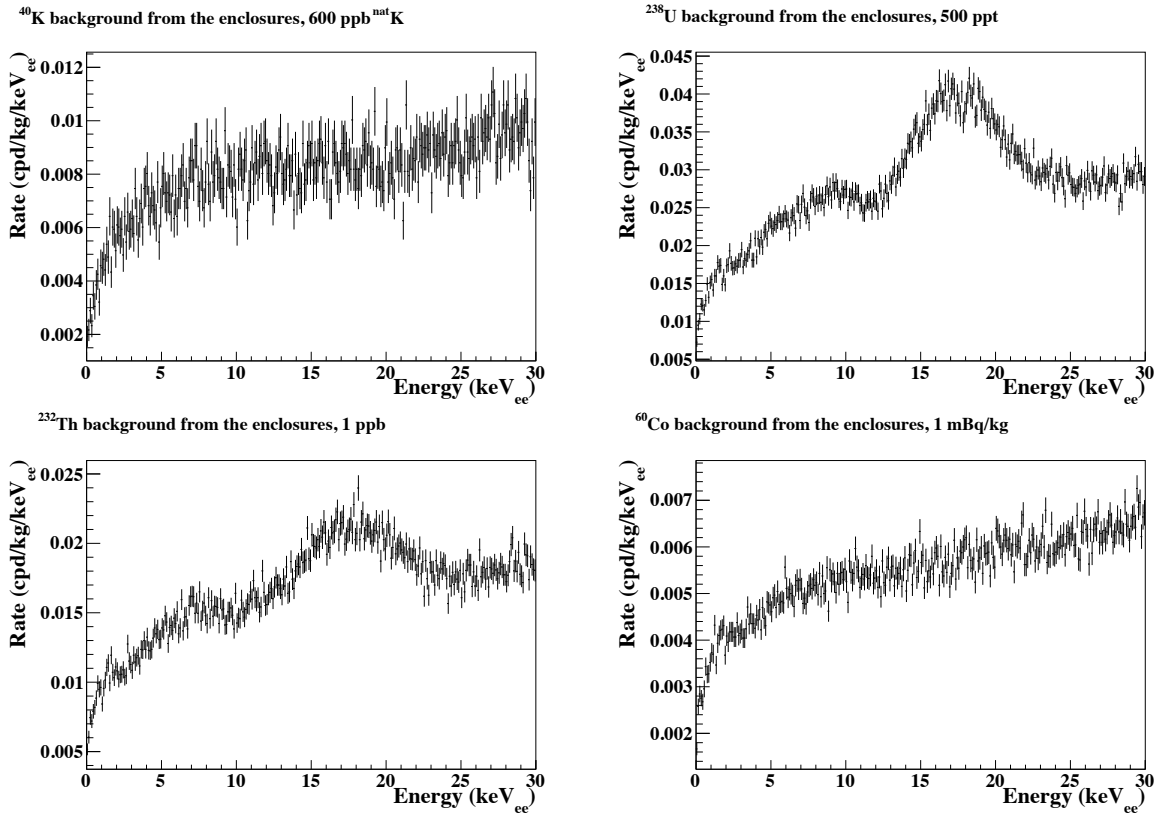


Figure 2.22: Simulated backgrounds from individual isotopes in the DAMA/LIBRA copper enclosures. (Top left) ^{40}K , 600 ppb limit; (Top right) ^{238}U , 500 ppt limit; (Bottom left) ^{232}Th , 1 ppb limit; (Bottom right) ^{60}Co , 1 mBq/kg. The ^{60}Co activity was not reported by DAMA, and an arbitrary value was assumed.

Background from the copper enclosures




Figure 2.23: Simulated backgrounds from the DAMA/LIBRA copper enclosures at the reported limits for ^{40}K , ^{238}U , and ^{232}Th . Because the ^{60}Co background was not reported by DAMA, it is not included in this spectrum.

Total DAMA background at limits

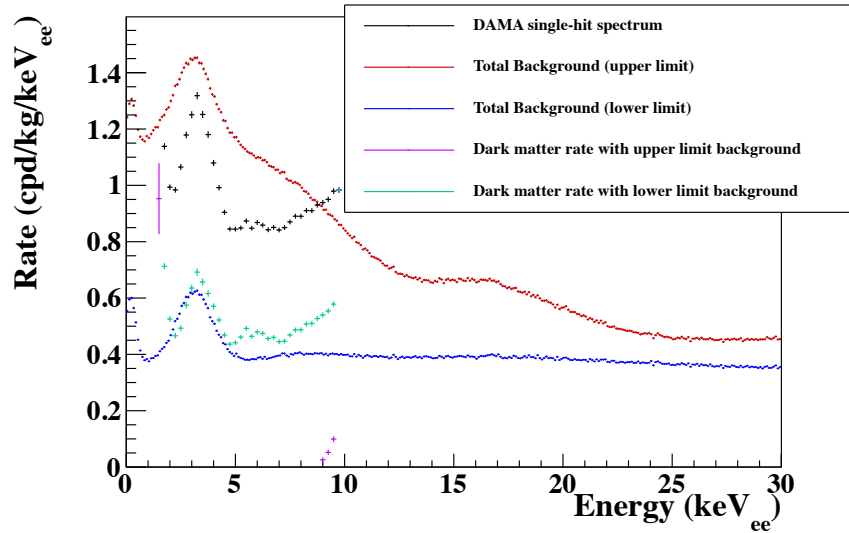


Figure 2.24: Simulated DAMA/LIBRA single-hit spectrum with DAMA/LIBRA reported radioactivity values. The red curve shows the values at the reported limits, while the blue curve ignores backgrounds reported as limits. The purple and green curves show the residual spectrum that would be attributable to WIMPs for the upper and lower-limit backgrounds, respectively. The residual spectrum for WIMPs does not follow an exponentially decreasing curve from 2–10 keV $_{ee}$ as would be expected, but rather increases with energy from 6–10 keV $_{ee}$.

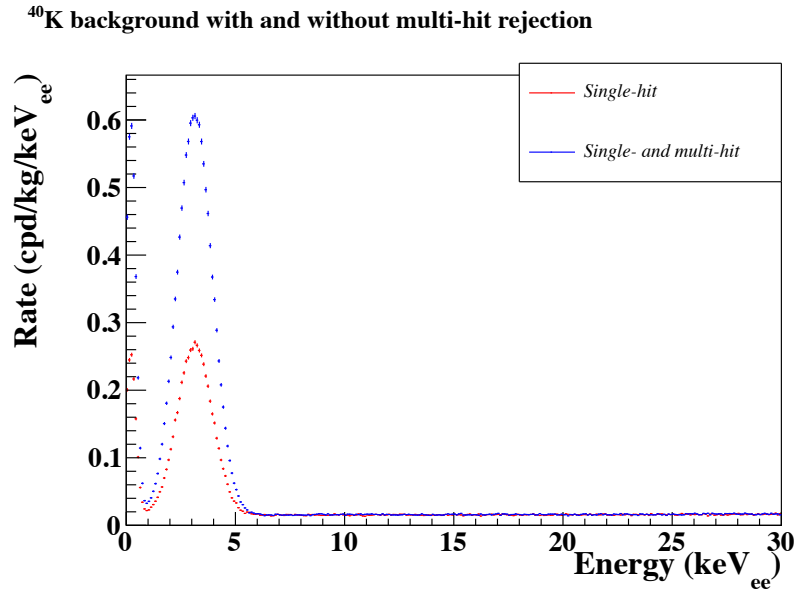


Figure 2.25: The simulation multi-hit rejection power of ^{40}K in the crystals for 13 ppb of ^{nat}K . In this simulation, the rejection power by multi-hit rejection was simulated to be 58%.

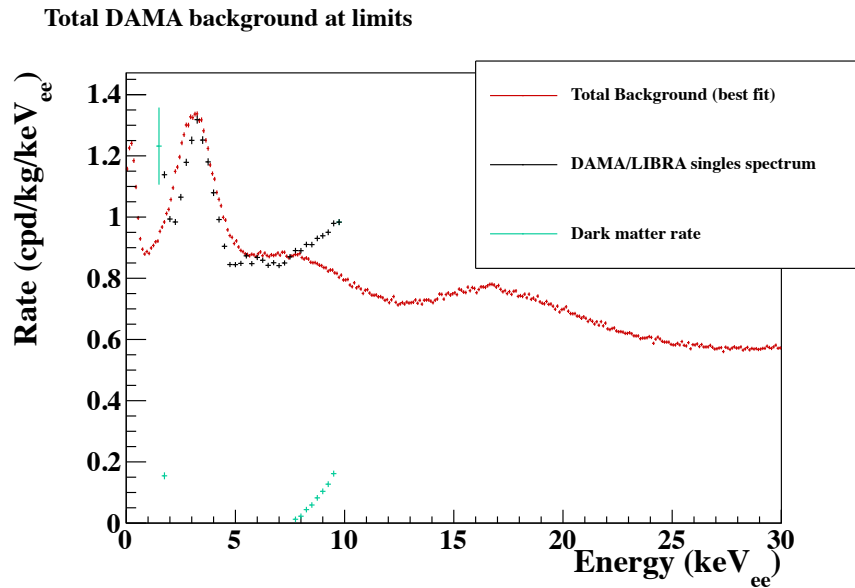


Figure 2.26: Simulated DAMA/LIBRA single-hit spectrum with best-fit radioactivity values. The ^{40}K contamination was fit over the 1.5–5.5 keV_{ee} range, while the other backgrounds were fit between 6 and 10 keV_{ee} . The ^{nat}K concentration was simulated to be ~ 23 ppb rather than the reported 13. These values leave no room for a WIMP signal above 2 keV_{ee} .

2.3.3 The Phase

The expected phase of the dark matter signal peaks on June 1. The DAMA/LIBRA peak occurs on May 16 (± 7 days) in their 2- to 4-keV_{ee} energy window, which is nearly two weeks earlier [39]. The difference in the phases has not been addressed by DAMA, who claim that the discrepancy is within errors.

The phase difference can be resolved by a number of effects, including gravitational focusing. This phenomenon would not affect a low-mass WIMP, but for a larger mass like 80 GeV, gravitational focusing can shift the peak earlier, to May 22, 21, and 19 in the energy ranges 2–6, 2–5, and 2–4 keV_{ee}, respectively, using DAMA’s quenching factor values [34]. This increase in the peak time with increasing energy is consistent with the slight differences in phase observed by DAMA in the different energy windows, as shown in Table 2.1. As can be seen in Figure 2.2, the effect of gravitational focusing on the modulation in NaI(Tl) could be tested by looking at lower threshold energies.

2.3.4 The modulation amplitude

Another unexpected behavior of the DAMA data is the apparent decreasing of the modulation amplitude with time. In the original DAMA data gathered from 1995–2001, the published modulation amplitude in the 2- to 6-keV_{ee} bin was reported to be 0.0200 ± 0.0032 cpd/kg/keV_{ee}. In the first DAMA/LIBRA data from phase 1 (2003–2007) the modulation amplitude in this energy region became 0.0107 ± 0.0019 cpd/kg/keV_{ee}. The results from phase 2 (2007–2009) can be calculated to be even lower, at around 0.0085 ± 0.0022 cpd/kg/keV_{ee} [52]. The difference in modulation amplitude between the initial results and the results from phase 2 is greater than 4σ . This trend is also the case for the 2–4- and 2–5-keV_{ee} bins [52]. This steady decrease in the modulation amplitude is a troubling signature that is not expected for a uniform (on local scales) dark-matter distribution.

2.3.5 The nuclear quenching factor and the channeling effect

The quenching factor is essential to determining the energy scale of nuclear-recoil events, and is therefore a major factor in determining the preferred dark-matter parameters and comparing experimental results with different targets. A lower quenching factor pushes the parameter space regions to higher cross sections and higher masses, which can increase the tension of DAMA/LIBRA with other experiments, as will be discussed in later chapters.

DAMA measured the quenching factor in their crystals by fitting a Monte Carlo simulation to the spectrum of single-scattering neutron events from a ^{252}Cf source. They assumed an energy-independent quenching factor, obtaining the values of 0.3 for sodium and 0.09 for iodine [53]. In this measurement, DAMA used a crystal with the same thallium concentration and growing procedure as the crystals used in DAMA/LIBRA.

Later measurements of the quenching effect in Na suggest that the quenching factor is both energy dependent and lower than the DAMA values [54, 55, 1]. However, DAMA claims that different crystals will have different quenching factors, depending on the concentration of dopants, contaminants, and differing growth procedures [39] and continues to use their energy-independent measured values. A new, more accurate measurement of the Na quenching factor performed by the SABRE collaboration is described in Chapters 6–8, and the implications for DAMA and the dark-matter picture are discussed.

It has been hypothesized that nuclear recoils aligned with the crystal lattice structure would experience no ionization quenching. This is called the channeling effect. A channeling effect would push the DAMA/LIBRA parameter-space regions downward in cross-section, potentially mitigating the conflict between DAMA/LIBRA and null experiments. This effect has not been observed, however [54, 1].

2.3.6 Conflicting WIMP measurements

In addition to the concerns about the DAMA/LIBRA data mentioned above, DAMA/LIBRA is the only dark-matter direct-detection experiment to claim a discovery of dark matter.

Some other experiments have obtained null results that exclude the DAMA region of interest, especially in the higher-mass region. However, DAMA/LIBRA’s Na target has a high sensitivity to low-mass WIMPs that other, heavier detectors have not achieved. In addition, the low-mass WIMP is attractive because of its higher sensitivity to the tail of the velocity distribution, which would allow for a larger modulation fraction, and therefore a lower total rate. As other experiments have become sensitive to the high-mass region favored by DAMA, discussion of reconciling DAMA with null experiments has focused on these light WIMPs of around 10 GeV in mass.

In 2011–2013, several experiments began to see signals that were in agreement with light WIMPs, and in certain models, with DAMA. These experiments, combined with null results from others, generated a strong interest in further study of the light-WIMP region of parameter space. Since then, some of the evidence for light WIMPs has become weaker, but there is still some interest in this region as a potential area in which the DAMA/LIBRA result can remain valid. The results from other direct-detection experiments relevant to DAMA and the light-WIMP picture are summarized below.

CoGeNT

Coherent Germanium Neutrino Technology (CoGenT) is a p-type point-contact (PPC) high-purity germanium detector that measures the ionization charge of small energy deposits in the detector (energy threshold 0.5 keV_{ee}). Because of its low energy threshold, it is sensitive to light WIMPs and has been observing a modulation in the $0.5\text{--}3 \text{ keV}_{ee}$ range with a high modulation fraction, a peak time of $t = 97 \pm 20$ days, and a period consistent with one year [56, 57]. This modulation signal, shown in Figure 2.27, is commensurate with a 13-GeV WIMP and occurs at a lower cross section compared with DAMA. Initial reports of a modulation in excess of 2σ have been later revised to a more modest 1.7σ [57]. The WIMP parameter-space region associated with this modulation is shown in Figure 2.32. In addition to the modulation signal, CoGeNT reports an excess of events in the total rate at

low energies ($0.5\text{--}2\text{ keV}_{ee}$) with an exponentially-decaying energy spectrum, as is expected of WIMP dark matter. This excess is also consistent with a WIMP with low mass [52].

However, there are some procedural complications that introduce some uncertainty in the CoGeNT results. In the processing of the data, CoGeNT uses a rise-time cut on the pulse shapes to eliminate surface events, which are known to have incomplete charge collection. At the energies where the excess is reported, the rise-time of surface and bulk events overlap and are difficult to separate, as can be seen in Figure 2.28. In addition, CDMS-Ge has a similar sensitivity to CoGeNT but does not see a modulation [52]. Finally, the PPC germanium detector, MALBEK, does not see a modulation, though their sensitivity has not yet reached the CoGeNT 2014 region [?], and Collar further argues that MALBEK's sensitivity may be further reduced, claiming that their surface-background rejection ability is overestimated [58].

CRESST

The Cryogenic Rare-Event Search with Superconducting Thermometers (CRESST) is a cryogenic dark-matter detector looking for scintillation and phonon signals from WIMP elastic scattering in their CaWO_4 crystals. CRESST initially reported an excess over background in their phase-1 detectors of CRESST-II which were consistent with light WIMPs [59] (see Figure 2.32). A later detector upgrade ruled out most of this region at low WIMP mass [60].

CDMS-Si

The Cryogenic Dark Matter Search (CDMS) uses semiconductor detectors at 40-mK temperatures to search for phonons and ionization signatures from WIMP interactions. In 2013, CDMS-Si announced a residual signal of 3 events that passed their background cuts, shown in Figures 2.29 and 2.30, which was in excess of their expected background of 0.7 events. They reported that these signals have a 5.4% chance of being due to background. Of interest




Figure 2.27: Annual modulation observed by the CoGeNT experiment. Solid lines are fits where the modulation period was constrained to one year, while dotted lines represent unconstrained fits. The modulation compatible with WIMPs is only found in bulk events at the energies where the CoGeNT excess is seen (0.5–2 keV_{ee}). The modulation is slightly out of phase with the DAMA/LIBRA modulation, whose maxima locations are indicated with arrows. Figure from [56].




Figure 2.28: (Top) Rise-time vs. energy in the CoGeNT detector. (Bottom) Rise-time profiles for surface and bulk events in 0.5–2 keV_{ee} and 2–4.5 keV_{ee} events. In the regions where the excess and modulation are reported, there is significant overlap in the rise-time profiles. Figures from [56].




Figure 2.29: (Top) Ionization yield vs. recoil energy in the CDMS-Si detectors. The top plot shows events that pass all cuts except for a phonon timing cut shown in Figure 2.30, while the bottom plot shows events passing all cuts. The black lines outline the nuclear-recoil region from calibrations. The grey band is the charge threshold of their detector. The population of electron recoils have an ionization yield near 1. The three events near the grey band are in excess of the expected background. Figure from [61].

to DAMA is that the signal region has a best fit with a WIMP spin-independent interaction where the WIMP mass is 8.6 GeV and WIMP-nucleon cross section is $1.9 \times 10^{-41} \text{ cm}^2$ [61].

There are a number of factors that weaken this result. First, it is notable that the three anomalous events are very close to the CDMS-Si charge threshold. In addition, a separate paper published by the CDMS collaboration showed no candidate WIMP events for a separate run with an expectation value of 1.1 background events. Finally, the CDMS-Ge detectors also see no candidate signals, and most of the CDMS-Si region of interest is excluded by the germanium results [62].




Figure 2.30: Events in the CDMS-Si detector showing the ionization yield (y axis) and the timing parameter (x axis) used for the cut that distinguishes the top and bottom plot in Figure 2.29. The populations are colored in the three different energy bands in Figure 2.29. The green and red histograms are the timing and ionization profiles of calibration sources for surface events (red) and nuclear recoils (green). The three anomalous events are shown in the black box. Figure from [61].

Xenon experiments

The LUX null results are currently the most sensitive published limits from liquid noble detectors. LUX is a liquid-xenon, two-phase time-projection chamber (TPC) that observes a primary scintillation signal (S1) and a secondary ionization signal (S2), the latter of which comes from electroluminescence in a gas layer at the top of the chamber that occurs after electrons are drifted to the surface with an electric field. In 2013, LUX published first results on an 85-day exposure of their 118-kg fiducial volume. After applying cuts on the ionization-to-scintillation ratio (S2/S1) and event energy (proportional to S1), LUX saw 160 events between 2 and 30 photoelectrons, shown in the red bands in Figure 2.31, which they mostly attribute to electron-recoil background. They set an upper limit on possible WIMP events in this energy region between 2.4 and 5.3 events, depending on the WIMP mass assumed. Based on this limit, LUX has set the most sensitive limits on the WIMP-nucleon spin-independent cross section to-date in the high-mass region.

Based on their null results, LUX claims to completely rule out the regions of interest for all of the experiments observing light-WIMP-like signals [63, 64] (see Figure 2.32). The XENON-10 and XENON-100 experiments have also set limits conflicting with the DAMA dark-matter claim, as well as other experiments observing potential light-WIMP signals [65, 66]. However, it is important to note that heavy targets like Xe may not be as sensitive to light WIMPs, since $v_{\min} \propto \sqrt{M}$ for a given threshold energy. Xenon-based experiments are potentially sensitive only to the high-velocity tail of the WIMP velocity distribution. Uncertainties in the escape velocity of the galaxy have a substantial effect on the sensitivity to low-mass WIMPs [67]. In addition, there is some uncertainty about the sensitivity to light WIMPs that is a result of the non-linear scintillation response of liquid xenon to nuclear recoils, \mathcal{L}_{eff} . Separate measurements of this important factor have not reached concordance at low energies [68]. Furthermore, because of the spread in the energy resolution, the detector may be sensitive to dark-matter recoils below the energy threshold, where \mathcal{L}_{eff} is unknown and must be extrapolated [68]. The uncertainty in this factor can




Figure 2.31: Events in the LUX detector during their initial 85-day exposure, after cuts. The x axis is the amount of light in photoelectrons collected in the initial scintillation signal, S1. The y axis is the ratio of secondary scintillation light from the ionization signal, S2, and the initial scintillation signal, S1. The nuclear-recoil band determined from calibrations is shown in red, while the region for electron recoils is shown in blue. A light-yield cut used by LUX is shown in the cyan vertical lines. The corresponding event energy contours are shown in gray in units of keV_{nr} , or nuclear-recoil-equivalent energy. The events in the red region are consistent with leakage of electron recoils into the nuclear-recoil band. Figure from [63].

lead to substantial uncertainty in the sensitivity of liquid-xenon detectors to WIMPs with low mass. Finally, if dark-matter is isospin-violating, the relative coupling between the dark matter and protons vs. neutrons can cause destructive interference that would make xenon much less sensitive to dark-matter interactions [69].

Bubble chamber detectors

Several experiments, such as Picasso [70], SIMPLE [71], and COUPP [72], have attempted to search for dark matter by superheating fluids and observing turbulence events that occur when energy deposited by incident radiation causes the nucleation of a bubble. These experiments are especially sensitive to light WIMPs because of their use of low-mass elements

like fluorine in the target material. All three experiments have set limits that conflict with DAMA for many dark-matter models.

KIMS

The KIMS experiment [73] is a CsI(Tl) detector array similar to the DAMA/LIBRA experiment with ~ 100 kg of target material. They recently published null results that conflict with much of the DAMA low-mass region of interest [73], as well as the high-mass region [74]. This experiment is of particular note because it shares a target with DAMA/LIBRA (iodine), and is therefore directly comparable with DAMA for the high-mass region if the quenching factors in CsI(Tl) and NaI(Tl) are the same. KIMS utilized pulse-shape discrimination methods to isolate the nuclear-recoil signals, observing a nuclear-recoil rate smaller than DAMA/LIBRA's modulation rate. Their low rate is incompatible with inelastic dark-matter models that enhance the modulation fraction, which are sometimes used to explain the apparent high modulation fraction observed by DAMA.

It is important to note that comparisons between experiments with different target materials necessarily depend on assumptions about the WIMP-nucleon interaction, form factors, and the halo model. Therefore, none of the null experiments mentioned above can provide a hard exclusion of the DAMA/LIBRA modulation result. In addition, no other experiment has performed any modulation search with the sensitivity of DAMA/LIBRA. Other experiments rely on a measurement of the total dark-matter interaction rate by applying stringent cuts on background. This difference in approach introduces dependencies on the dark-matter model that can make comparing these experiments difficult. Finally, no experiment has been conducted with NaI that has reached the sensitivity of DAMA/LIBRA. As will be discussed, there are still some models that can reconcile the DAMA/LIBRA data with other null results.




Figure 2.32: WIMP parameter-space regions and exclusion curves in the light-WIMP region. The CRESST collaboration has recently claimed that their region is not supported by current data. The most sensitive exclusion curve comes from the LUX experiment, which claims to rule out light WIMPs altogether. Figure adapted from [64].

2.3.7 Possible explanations for the DAMA/LIBRA modulation and potential resolutions to the conflict

A number of potential non-WIMP explanations of the DAMA/LIBRA signal have been suggested, as well as a number of dark-matter models that could reconcile the conflict between DAMA and other null results. None of the non-WIMP explanations are truly satisfactory, and the dark-matter models are untested given current observations. The uncertainties in understanding the DAMA/LIBRA signal, along with the lack of a satisfying explanation for the conflict, presents a strong need for an independent test of the DAMA/LIBRA experiment with a new NaI(Tl) detector.

Cosmic ray muons

Cosmic ray muons (μ) can penetrate deep into the earth's surface and interact directly with a dark-matter experiment, or they can interact in the surrounding environment or shielding materials, producing secondary particles like neutrons and γ s. The muon flux at LNGS has been measured by the Borexino experiment to be $(3.41 \pm 0.01) \times 10^{-4} \text{ m}^{-2}\text{s}^{-1}$ [75]. Other experiments have also measured the muon flux in LNGS [76].

The muon flux exhibits an annual modulation as a function of temperature, which has been measured by [75], among others. In warmer temperatures, the Earth's atmosphere is less dense, allowing more pions and kaons from cosmic-ray interactions with the upper atmosphere to decay to highly energetic muons that can penetrate to the depth of DAMA/LIBRA. The muon flux exhibits a phase with a maximum which has been measured to be in late June/early July. The recent measurement from [75] places this maximum on the 28th of June (± 6 days). Because this date is close to the expected dark-matter modulation peak (June 2nd), cosmic-ray muons and their interaction products are often proposed to be a potential source of a false dark-matter modulation signal [77, 78].

A direct signal from muon events in the crystals would not explain the DAMA signal well, since muons deposit high energies. The muon modulation should appear at higher energies and in the multi-hit events as well as the DAMA/LIBRA modulation region. However, Nygren puts forward a hypothesis in [78] that not all energy from the muons goes directly to scintillation. Rather, some of the energy could be stored for later release in the low energy ($< 10 \text{ keV}_{ee}$) region as fluorescence, which would only appear in one crystal at a time. Nygren argues that the phase and modulation are consistent within errors with the DAMA/LIBRA modulation and that the solid-state effects in NaI(Tl) are not well known, making this mechanism a possible explanation of the DAMA/LIBRA modulation.

Others have proposed that while muon-induced events in the crystals may be high in energy, nuclear recoils may result from neutrons being released from shielding and detector materials when a muon interacts in these materials. These neutron events could masquerade

as WIMP signals in the crystals. DAMA measured the flux of neutrons in the detector by measuring the activation of ^{24}Na in their crystals, and conclude that the flux of neutrons is several orders of magnitude too small to mimic their signal [76].

There are some general problems with using the muon-flux modulation as a direct or indirect explanation of the DAMA signal. The muon fractional modulation has been measured to be 1.3–1.5% [79]. Because the DAMA/LIBRA modulation is $\sim 2\%$ of their total signal at some energies, any effect due to muons would necessitate that muons make up most of the DAMA spectrum. This would mean little background below 6 keV_{ee} . It is difficult to reconcile the need for such a low background rate and the $\sim 1\text{ cpd/kg/keV}_{ee}$ rate between 6 and 10 keV_{ee} , since backgrounds are expected to have a mostly constant spectrum below 10 keV_{ee} . Also, according to simulations, the background levels quoted by DAMA produce a larger background rate below 6 keV_{ee} than would allow for such a small fractional modulation. The comparison of the modulation to the total rate worsens at lower energies between 2 and 3 keV_{ee} . [79] concludes that the total rate would have to be $1.4\text{ cpd/kg/keV}_{ee}$ in order to accommodate a muon fluctuation. Finally, DAMA estimates that the muon flux in their experiment is $2.5\ \mu/\text{day}$ [76]. In order to produce the effect Nygren describes, each muon would be responsible for hundreds of signals. The poisson fluctuations in the rate would therefore be higher than expected for scintillation events each caused by a different physical energy deposition. DAMA claims that such a fluctuation in the statistics is not observed [76].

Alternate dark-matter models

The DAMA/LIBRA parameter-space region of interest for the standard WIMP picture has been ruled out by a number of experiments. However, the simplicity of the standard WIMP picture is the typical motivation for its use in comparing experimental results. In reality, the characteristics of dark matter are unknown. Dark matter could consist of multiple particles that may even self-interact. The coupling mechanism between dark matter and

baryonic matter is also unknown. There are several existing dark-matter models that could be consistent with the results of DAMA/LIBRA and other null experiments.

Self-interacting dark matter can interact by exchanging a kinetically-mixed photon, forming bound states. These kinetically-mixed photons can transition between a dark photon¹³ and a normal photon, allowing electronic interactions in a way that can possibly reconcile DAMA with experiments that reject electron recoils [80]. There are a number of different models surrounding this idea. A recent example is [81], where it is proposed that several different species of dark matter form bound states called dark anti-atoms that bind to heavy nuclei. In DAMA's case, this would be the Tl dopant in the crystals. Other experiment targets would be light enough that this signal would look like electron recoils or would not be seen at all.

DAMA claims that dark matter could interact with leptons only, thus producing electron rather than nuclear recoils [82], or that electron recoils could arise indirectly through the WIMP interaction with the nucleus [83]. Other experiments eliminate electron-recoil signatures with pulse-shape discrimination and other cuts, and so might miss the signal in this case.

Dark matter could also contain a small mass-splitting that would allow the dark matter to transition to a higher (endothermic) or lower (exothermic) energy state when interacting with a nucleus. The models implementing this mechanism are called inelastic dark-matter models [84]. This effect can enhance the modulation fraction by making a detector only sensitive to the upper tail of the velocity distribution [37]. Inelastic models also increase the sensitivity for heavier targets in the case of endothermic scattering and lighter targets in exothermic scattering. The vanilla endothermic case is ruled out by experiments with xenon as the target, which is heavier than iodine [85]. This may be mitigated by taking advantage of the dipole moment of iodine with a magnetic coupling. This inelastic, magnetic dark matter can still be consistent with current null results if the iodine quenching factor is small

¹³A photon-like particle that does not directly interact electromagnetically with Standard-Model particles

(0.04 instead of 0.09) [86], a possibility that will be discussed in Chapter 7. The exothermic inelastic case is in conflict with bubble-chamber experiments [85].


For spin-independent couplings of dark-matter to matter, a commonly held assumption is that the coupling to the proton (f_p) and neutron (f_n) are the same. In this case, the spin-independent cross-section between the dark-matter and a nucleus with atomic number Z and atomic mass A goes as [9]:

$$\sigma_{\text{SI}} \propto [Zf_p + (A - Z)f_n]^2, \quad (2.6)$$

which goes as A^2 for $f_p \approx f_n$. However, the dark-matter could be isospin-violating. The relationship between f_p and f_n can be tuned so that the couplings destructively interfere for xenon. It has been proposed that such a model could reconcile DAMA, CoGeNT, and the Xenon experiments if $f_n/f_p = -0.7$, but the region required for this reconciliation for elastic scattering is now heavily constrained by indirect-detection experiments and collider searches [69]. Inelastic, isospin-violating dark matter was analyzed to reconcile the xenon null results with CDMS-Si, but the relationship required eliminated the DAMA/LIBRA low-mass best-fit regions [69].

The coupling could, however, be spin-dependent. In this case, the WIMP would scatter off an unpaired nucleon, rather than coherently with the entire nucleus. This case would make detectors with lighter targets comparatively more sensitive than in the spin-independent case. Some targets have an unpaired neutron while others have an unpaired proton, and the WIMP-nucleon coupling may be different for protons and neutrons. Depending on the spin-coupling ratio, Na and I, which have unpaired protons, may be sensitive while others, like Ge and Xe which have unpaired neutrons, may not. A spin-dependent WIMP coupling mostly to protons may therefore help with the tension with null results. However, [85] analyzed this possibility and found that this dark-matter model is heavily constrained by

Vector, proton coupled: $\vec{\sigma}(i)$ vs. $\vec{\ell}(i)$



Vector, neutron coupled: $\vec{\sigma}(i)$ vs. $\vec{\ell}(i)$




Figure 2.33: Spin-response functions of various dark-matter targets, with a spin-angular-momentum coupling on the left and an orbital-angular-momentum coupling on the right. The top figures represent the nuclear response for a WIMP coupling to the proton, while the bottom shows couplings for neutrons. If the underlying WIMP interaction is with the orbital-angular-momentum operator, but is assumed to be with the spin, a conflict may appear between experiments where there may be none. Figure from [87]

other experiments such as KIMS, PICASSO, and COUPP, which use fluorine as a target. This model is only reconcilable if the iodine quenching factor is higher in NaI than in CsI.

It has also been proposed that the cross-section may be dependent on orbital angular momentum or may be velocity dependent, or some other generalized angular-momentum dependence [87]. In this scenario, the response may be much greater for Na and I than for Xe, Ge, and F (see Figure 2.33). This example shows how, if misunderstood, the form of the WIMP-nucleon coupling can produce the appearance of conflict where there may be none.

Other astrophysical models

The SHM is a simple analytical model for the shape and velocity distribution of the WIMP dark matter halo. However, there is evidence from large-statistic N-body simulations that

this model may not describe our galactic halo [32]. Because the annual modulation is heavily dependent on the velocity distribution of dark matter, small changes in the halo model can make a difference in the expected dark-matter signature.

The shape of the smooth component of the velocity distribution could be different than the Maxwellian distribution. For example, self-interacting dark-matter models may produce a dark disk like the disk of the baryonic matter in the Milky Way. The dark matter halo may also have an oblate or triaxial distribution, with velocity dispersions being directionally dependent. Such a distribution can change the modulation signature on the order of 10%, depending on the model [32].

The halo may also have unvirialized components and local structure. Tidal streams of dark matter, like the one described in [88], can produce a modulation in the dark-matter rate separate from the one expected from the WIMP wind. This would loosen the phase requirement of the modulation signal. This modulation due to local structure, depending on the velocity characteristics of the stream, can be from 2–18% [88].

However, the uncertainty in the astrophysical model can only do so much to reconcile experiments with conflicting results. In [89], it is claimed that assuming only a homogeneous dark-matter distribution on the scale of the ecliptic and the stability of the halo on the scale of months, the DAMA/LIBRA modulation is excluded by the XENON, SIMPLE, and CDMS results to high significance for both the standard spin-independent and spin-dependent-coupled-to-protons cases.

The need for an independent test of DAMA/LIBRA

Because there are so many unknowns concerning the WIMP model, the galactic dark matter halo, and experimental parameters, comparisons between experiments with different targets is ultimately not possible; it will be extremely difficult to rule out the DAMA/LIBRA result in a reliable way. Because of the unresolved conflict between DAMA/LIBRA and other experiments, and because the declaration of a WIMP discovery should be supported by

multiple affirmative findings, there is a strong desire in the direct-detection community for an independent test of DAMA/LIBRA. Because of the uncertainties mentioned above, this test should use the same target material as DAMA: NaI(Tl).

Chapter 3

SABRE: An Independent Test of DAMA/LIBRA



SABRE

Original image for background from ESA/Hubble. Brightness reduced and logo drawn by Emily Shields.

3.1 The importance of an independent test of DAMA/LIBRA

In Section 2.3, it was established that the conflict surrounding the DAMA/LIBRA modulation result is an important problem in the field of direct dark-matter detection. However, uncertainties in the dark-matter model, the halo model, and the detector response itself makes it impossible to make a model-independent comparison between DAMA/LIBRA and other experiments with different targets. Furthermore, there is no other NaI(Tl) experiment that has reached the sensitivity of DAMA/LIBRA at this time.

Because of the sheer significance of the DAMA modulation signal, the high purity achieved in the construction of their apparatus, and hints of light WIMPs from other experiments, the DAMA/LIBRA result cannot be dismissed out-of-hand. For more than a decade, a large effort has been expended in an attempt to understand the DAMA result and reconcile it with conflicting experiments. A great deal of speculation has occurred on the authenticity of their signal as evidence of dark matter. An independent test of DAMA/LIBRA with the same target is strongly desired by the community so that this issue can be put to rest.

3.2 Qualities of an independent test

Because of the long-standing questions surrounding the DAMA/LIBRA result, an independent test of DAMA must be scientifically robust enough to support DAMA or rule it out. It must also accomplish this in a reasonable time frame, which is a considerable requirement given that DAMA has been gathering data for over a decade. A new NaI(Tl) experiment should therefore not simply seek to replicate the DAMA/LIBRA experiment, but to improve upon it and become a strong direct-detection experiment in its own right. There are several ways that a new experiment could accomplish this.

3.2.1 Lowering and understanding backgrounds

DAMA/LIBRA is a very low-background experiment with ample passive shielding and the purest NaI(Tl) crystals currently in existence. Improving upon their backgrounds is therefore a challenging task. However, it is a worthy goal—many of the questions surrounding the DAMA/LIBRA modulation, such as the modulation fractional amplitude and limits on the total dark-matter rate, can be better known with a lower, well-understood background. Lowering the background will make a new experiment more sensitive to the total WIMP rate and will allow an achievement of the same sensitivity to the modulation in a shorter amount of time.

To lower the background in a new experiment, there should be a strong effort to reduce the level of impurities in the detector materials. New crystal purification methods and development of other low-radioactivity materials will be paramount to a test of the DAMA/LIBRA result.

Secondly, backgrounds can be reduced by employing active veto techniques, which eliminate backgrounds that occur in multiple detectors. DAMA/LIBRA employs this technique to some degree, but the fact that the rejecting volumes are also fiducial limits this technique's effectiveness. The rejection power is comparatively reduced for crystals on the outer portion of the array, since there is less material tagging outgoing radiation. If a dedicated veto detector with 4π coverage can be employed for the sole purpose of rejecting backgrounds, and not as a WIMP target, the ability to reduce backgrounds will be improved for crystals at the outer edge of the array. The active veto is especially important for the rejection of the residual ^{40}K background. A schematic of this concept applied to a ^{40}K decay is shown in Figure 3.1.

Finally, backgrounds can be reduced by employing rejection techniques based on the characteristics of the signal. DAMA/LIBRA does this only to reduce PMT noise, but not on valid scintillation signals. NaI(Tl) has PSD capability, which can be used to lower the background. DAMA suggests that such a technique is dangerous, because if WIMPs interact




Figure 3.1: The active veto concept. Radiation deposits a small energy in the crystal detector, mimicking a dark-matter signal. The original particle or a coincident particle may escape the crystal detector, but be tagged by the veto detector. In the scheme above, the veto detector rejects a ^{40}K Auger-electron event by detecting the outgoing 1461-keV γ ray. Other background signals can also be rejected by this technique, as well.

via electron recoils, they would be rejected. However, there is no reason an analysis that can separate nuclear recoils from electron recoils could not be used to look for a modulation signal in either channel. If a modulation were found in only one type of interaction, this would provide more information about the WIMP-matter interaction. Pulse-shape discrimination can be strengthened by operating the detectors at colder temperatures, where the scintillation yield and PSD characteristics of NaI(Tl) are enhanced and the PMT noise level is lower [90]. However, the pulse-shape discrimination capability is dependent on photoelectron statistics, and is currently marginal in the DAMA/LIBRA modulation region. If a detector could be operated at cryogenic temperatures, it could take advantage of a bolometric signal in addition to scintillation, as proposed in [91].

An independent test of the DAMA/LIBRA experiment should also try to understand its backgrounds as completely as possible. This includes measuring the radioactivity of all detector components with sensitive techniques that can detect impurities at the ppb or even ppt level, preferably on short time-scales. As part of the research effort, such techniques may need to be developed. Monte Carlo simulations can be used to predict the background spectra from various isotopes, which can be calibrated at higher energies where many of

these backgrounds have distinct features. A thorough understanding of the backgrounds in a new experiment could help to avoid the current situation for DAMA, where questions about the shape of energy spectrum and a lack of published data make it difficult to conclude the veracity of a WIMP signal or to learn more about the WIMP characteristics.

3.2.2 Lowering the energy threshold

The DAMA/LIBRA modulation is very close to their software threshold and a region heavily populated by electronic noise events. DAMA attributes these noise events to fast single photoelectrons in the PMTs [49]. One way to become more sensitive to the presence of WIMPs and a modulation signal would be to lower the energy threshold of the experiment. Lowering the energy threshold is also advantageous because of the signature of the WIMP modulation. The modulation phase reverses at low energies; the location of this reversal depends on the WIMP mass [32]. For the DAMA modulation, this phase reversal happens just below DAMA's energy threshold for an 80-GeV WIMP, while the light-WIMP modulation continues to increase at lower energies (see Figures 2.2 and 8.3). Lowering the energy threshold could help distinguish between the two DAMA regions of interest.

There are several ways to lower the energy threshold. The first method is to reduce low-energy noise that can obscure legitimate scintillation signals. This can be done by noise-rejection techniques or by operating the PMTs in a low-noise configuration. The second is to increase the detector response for a given energy by amplifying signals or increasing the light yield of the detectors.

3.2.3 A new location

It is an unfortunate coincidence that the expected WIMP-modulation phase corresponds with temperature fluctuations due to the changing seasons in the Northern hemisphere. Backgrounds that vary with temperature would reverse phase if a detector was operated in

the Southern hemisphere. As previously discussed in Section 2.3.7, the muon flux also varies with seasonal temperatures, but can be mitigated by operating a detector at a deeper site.

The Stawell gold mine in Victoria, Australia, is being converted into an underground laboratory space that could host such an experiment in the Southern hemisphere. The SNOLab underground laboratory in Sudbury, Canada [92] is 6,000 m.w.e. underground, much deeper than LNGS. Though SNOLab is located in the same hemisphere as LNGS, because it is so much deeper underground, any modulation from the cosmic ray background would be drastically reduced in size. Another possibility is to place the new NaI(Tl) detector under ice in Antarctica. Placing the detectors under ice would provide a very different set of radiogenic backgrounds than the Gran Sasso rock. In addition, Antarctica's location in the Southern hemisphere can provide a test of the modulation phase. However, the cost feasibility of this last option is questionable.

3.3 Sodium-iodide with Active Background REjection: The SABRE experiment

The SABRE (Sodium-iodide with Active Background REjection) experiment is a new NaI(Tl) experiment seeking to provide a much-needed test of the DAMA/LIBRA modulation result. Through a multi-pronged approach, SABRE expects to improve upon DAMA/LIBRA in the ways described in Section 3.2, namely to achieve lower backgrounds, a lower energy threshold, and to maintain the potential for operation at different sites.

The SABRE experiment will consist of an array of NaI(Tl) crystal detectors immersed in an active liquid-scintillator veto similar to the one developed for the DarkSide-50 experiment [93]. Each crystal detector will consist of a high-purity NaI(Tl) crystal coupled to two high-purity, high-quantum-efficiency phototubes with an optical-coupling gel. The crystal will be wrapped in a reflector that will prevent scintillation light from being absorbed before it reaches the phototubes. The crystal and PMTs will be encased in a copper enclosure.

The crystal detectors will be immersed in a liquid-scintillator vessel surrounded by passive shielding and operated in a deep underground site. In its first phase, SABRE will operate a proof-of-principle experiment that will prove its capability for testing the DAMA/LIBRA result. This experiment will consist of a small array of ~ 5 -kg high-purity crystals, but later phases may expand to a larger array. The main efforts of SABRE have focused on the development of high-purity crystals from high-purity starting materials, low-radioactivity PMTs with noise-rejection techniques, low-radioactivity enclosures, and the ability to deploy crystal detectors in a liquid-scintillator veto vessel. In addition, SABRE has also made an effort to characterize basic detector performance, like the Na scintillation quenching factor. Despite SABRE's small mass compared with DAMA/LIBRA, SABRE will be fully capable of providing a reliable test of the DAMA/LIBRA modulation through these improvements.

3.3.1 Low-background powders and crystals

DAMA/LIBRA and Saint Gobain have a contract prohibiting the sale of Saint Gobain's high-purity NaI(Tl) crystals to parties other than the DAMA collaboration. It is therefore necessary for an experiment attempting to test the DAMA result to develop new high-purity crystals. Impurities in the crystals are the highest source of background for a well-shielded experiment. Therefore, a primary effort of the SABRE experiment has been to develop new high-purity crystals with the aim of achieving lower backgrounds than those of DAMA/LIBRA.

The crystal growth process offers some potential to segregate out impurities. However, the purity of the final crystal is still dependent on the purity of the starting materials. The SABRE experiment has made a concerted effort to develop high-purity materials for the SABRE crystals at every stage of the process. This involves purification of the precursors to NaI, hydroiodic acid and sodium carbonate. In addition, efforts are being made to further purify the resultant NaI powder before crystal growth. TlI powder, which is mixed with the NaI powder prior to crystal growth, can also be purified.

Techniques are being developed to limit the exposure of the crystal material to impurities during the crystal growth process. The selection of high-purity materials, the development of effective cleaning procedures, and the development of clean post-processing and packaging procedures are all necessary to developing crystals with better purity than those used by DAMA/LIBRA. The efforts to produce high-purity powders and crystals are described in detail in Sections 4.1 and 4.2.

3.3.2 Low-background, low-noise, high-quantum-efficiency PMTs

The PMTs have a large role in the background and the energy threshold of a scintillation-based dark-matter experiment. The radioactivity of the PMTs can be a problem given their close proximity to the dark-matter target. In addition, the detector response is highly dependent on the performance of the PMTs, such as their quantum efficiency and collection efficiency. SABRE is working to develop PMTs that would allow for a lower-background experiment with a lower energy threshold than DAMA/LIBRA.

SABRE is continuing the work done by the DarkSide experiment to develop low-radioactivity materials for the PMTs. The ceramic feedthroughs that insulate the wires going from the divider board to the dynode chains are a particularly high source of radioactivity. Work is being done to develop high-purity ceramic feedthroughs for existing PMT types.

The achievable energy threshold of an experiment is directly dependent on the light yield of the detectors. DAMA/LIBRA protects their crystals from the radioactivity of their phototubes by separating them with a 10-cm light-guide and packaging the PMT in a copper shield that blocks some of the PMT radiation from traveling to the crystals. This light-guide decreases the light collection of the detector by moving the photodetectors further from the source of the photons. By developing lower-background phototubes, the PMTs can be coupled directly to the crystal, allowing for better light collection.

The light collection can also be improved through the use of phototubes with a higher quantum efficiency. Since the DAMA/LIBRA experiment went online, photocathode materials like super- and ultra-bialkali photocathodes have been developed with higher quantum efficiency than the old DAMA phototubes. The Hamamatsu R11065 PMTs developed for the DarkSide experiment, for example, have a peak quantum efficiency of $\sim 35\%$ [93]. Though DAMA has installed new, higher quantum-efficiency phototubes in their detector, the current published results are based on the data collected with the old PMTs. The use of a higher-quantum-efficiency phototube will allow for a lower energy threshold than the currently published data.

Another improvement on the PMTs that can lower the energy threshold is the reduction of photomultiplier noise. DAMA/LIBRA depends on a coincidence between the two PMTs and the timing information of the pulses to reduce photomultiplier noise, which has a smaller time constant than NaI(Tl) scintillation events. These techniques cannot be used to eliminate an effect called dynode-afterglow, however. In dynode afterglow, the dynode of one PMT can emit light during the electron cascade, which can be detected by another PMT that is light-coupled to the first PMT, thus producing a false coincidence. At low energies, dynode afterglow produces a significant background. The development of external amplification for the signal can allow the PMTs to be operated at a lower voltage, which mitigates the dynode afterglow effect, as will be described in Section 4.3.2.

The SABRE efforts toward the development of low-radioactivity PMTs and low-noise modes of operation are described in Section 4.3.

3.3.3 Low background enclosures with high light yield

The enclosures are a source of background due to their close proximity to the crystals. Reflectors, optical couplings, supports for the crystal and PMTs, and the enclosures themselves should be selected from high-purity materials. DAMA encloses their crystals in OFHC copper containers, which are traditionally low in radioactivity. However, the radiopurity of the

enclosures can be further improved by developing thin enclosures by electroforming. Efforts along these lines are described in Section 4.4.

The enclosure design determines the light yield of the detectors. The optical properties of the reflector and the optical coupling in particular determine the achievable light collection. The overall size of the crystals also has an effect on the light yield, with shorter crystals allowing more light to reach the PMTs. A compromise has to be made, therefore, between target exposure, which is a direct consequence of crystal size, and energy threshold. The design considerations of the enclosures in the effort to optimize the light yield is described in Section 4.4.

3.3.4 Active background rejection with a liquid-scintillator veto detector

One of the key improvements offered by the SABRE experiment is the ability to actively reject backgrounds with a liquid-scintillator veto. A similar veto detector was developed for the DarkSide experiment to reject neutron backgrounds exiting their main detector [94]. SABRE will similarly operate inside a veto detector that will reject a number of backgrounds, including the γ ray from ^{40}K decay.

The veto power of the liquid-scintillator detector is primarily determined by its ability to stop radiation (since stopping the radiation is necessary for its detection) and its light yield. The ability to stop radiation is determined by the veto detector size and the scintillator type. A veto detector with an organic scintillator like pseudocumene (PC) or linear alkylbenzene (LAB) can detect $\sim 95\%$ of ~ 1 MeV γ rays with 50 cm of depth. With a liquid-scintillator veto detector of modest size, the background rate from ^{40}K can be reduced by a factor of 10 or more.

The veto detector should also be constructed with radio-pure materials, and be outfitted with radio-pure PMTs operated with a majority logic, so as not to introduce a heavy dead-time burden, as would be the case if only one PMT signal were required for a veto. In order

to maximize the light collection, a reflector should line the surfaces of the vessel, and a fluor should be used to shift the wavelength of the scintillation light to the primary sensitivity region of the PMTs. In the design of such a detector, many factors of safety, performance, and cleanliness are considered and discussed in detail in Section 5.3. In addition, the possibility of deploying crystal detector modules in the already-existing DarkSide neutron-veto detector is discussed in Section 5.2.

3.3.5 A deep underground site with potential for relocation

Ideally, a new NaI(Tl) detector would be portable such that it can be operated in multiple underground locations. A detector that can operate at LNGS as well as other, deeper sites or sites in the Southern hemisphere can test for any systematics inherent to the Gran Sasso site. Furthermore, a location at a deeper site will have fewer cosmogenic backgrounds. The acquisition of low-radioactivity shielding of different materials that can block neutrons as well as γ rays can also lower backgrounds in the experiment. SABRE is constructing a portable veto detector that can be relocated at several different sites, including SNOLab and the Stawell gold mine and is acquiring passive shielding to further reduce backgrounds from external radiation.

3.3.6 Understanding the detector response and backgrounds

After backgrounds are lowered through purification, active rejection, and shielding, a residual background rate may remain. To constrain possible dark-matter models, it is important to understand the residual backgrounds as much as possible. In order to accomplish this, the residual radioactivity of the detector components must be known accurately, as well as the efficiency with which the detector will register this radioactivity as a scintillation event.

The residual radioactivity in different detector components can be measured in a number of ways. One of the most accurate methods is γ counting, wherein a detector, often high-purity Ge, is placed in contact with the material, and the characteristic γ rays are counted.

This process can be very time-consuming for high-purity components with low levels of radioactivity. For materials where a small sample can be removed, like the crystals, sensitive mass-spectrometry techniques can be used to assess the residual radioactivity at very low levels.

Once the residual radioactivity levels of different detector components and the laboratory environment are known, the background spectrum can be understood with physics simulation software like the Geant4 simulation package. Recent versions of this software have been developed to simulate physics effects at sub-keV energies, making them capable of predicting the background spectrum of a NaI(Tl) experiment in the DAMA/LIBRA modulation region. Simulations of the expected SABRE detector background and performance will be discussed in Section 5.4.

3.3.7 Measurement of experimental parameters of relevance to direct-detection searches

In order to provide a definitive test of DAMA/LIBRA, a new NaI(Tl) experiment will need to accurately characterize the performance of their detectors. Important experimental quantities like the ionization quenching factors, energy resolution, trigger efficiencies, and electronic noise rates will need to be well understood in order to assess the implications of a modulation result for dark matter. Due to the importance of the low-energy region for the dark-matter search, understanding these effects near the energy threshold will be especially important. Recent controversy over the Na quenching factor, for example, prompted SABRE to perform a more accurate measurement of this quantity. This measurement will be described in Chapters 6, 7, and 8. SABRE will continue to investigate the physical characteristics of NaI(Tl) detectors to better understand the detector response to backgrounds and potential WIMP candidates.

Chapter 4

The SABRE Crystal Detectors

One of the main goals of the SABRE experiment is the development of dark-matter crystal detectors with unprecedented radiopurity and sensitivity. Progress has been made toward this goal with each component of these detector modules. The current state of the SABRE crystal detectors is described in this chapter.

Each SABRE crystal detector will consist of a NaI(Tl) crystal optically coupled to two PMTs, enclosed in a radio-pure container. The primary effort has been the development of high-purity, high-optical-quality crystals, which has been made on two fronts: the purification of the powder used to make the crystals (Section 4.1), and the crystal-growth process itself (Section 4.2). Second, the effort to produce low-noise, high-purity phototubes is described in Section 4.3. Third, the effort to design and fabricate high-purity enclosures that will encapsulate and protect the crystals is put forward in Section 4.4.

4.1 Powder

The SABRE approach to the production of high-purity crystals is to make the starting materials as pure as possible before growth. Rather than relying on the crystal-growth

process to purify the powder to the desired level, SABRE has made a strong effort with industrial partners to develop NaI and TlI powders with unprecedented purity¹.

4.1.1 The need for high-purity powder

The DAMA/LIBRA NaI powder has impurities in the following concentrations: 20 ppt ^{238}U , 20 ppt ^{232}Th , and <100 ppb ^{nat}K , where ^{nat}K is the concentration of natural potassium, with a mass fraction of 1.1×10^{-4} of ^{40}K [41]. In the TlI powder, DAMA reports that the residual contaminants are 0.8 ppb ^{238}U , 1.2 ppb ^{232}Th and <60 ppb ^{nat}K . These concentrations are several times higher than the concentration DAMA reports in their crystals; the crystal growth process purifies the material significantly [41].

In order to achieve the production of crystals with lower background than DAMA, SABRE has worked toward the goal of producing powders with impurity concentrations lower than, or at least similar to, those of DAMA's final crystals. In this way, as long as the crystal-growth process does not introduce impurities, the crystals can be made purer than those of DAMA. The initial focus was to decrease the concentration of ^{nat}K in the powders, since the background from ^{40}K produces a peak in the primary modulation region of DAMA. However, the ^{238}U and ^{232}Th chains produce a broad-spectrum background in this region as well, which is nearly on the same order of magnitude as the DAMA/LIBRA signal when ^{238}U and ^{232}Th are present at the tens of ppt level. ^{87}Rb also produces a β background in this region that is significant at the hundreds of ppt level. The concentrations of these impurities were assessed during the production of the SABRE powders as well.

¹The primary effort in the development of high-purity powders has been executed by Drs. Frank Calaprice, Jay Benziger, and Alex Wright and SABRE's industrial partners as described in the text. This section is primarily meant to be serve as an overview of their work.

4.1.2 Development of high-purity powder

NaI powder

SABRE partnered with Sigma Aldrich SAHC HiTech and Seastar Chemicals Inc./MV Laboratories, Inc. to develop and test high-purity powders. Sigma Aldrich partnered with Princeton to develop upgraded methods and purification procedures. Seastar focused on developing highly sensitive Inductively Coupled Plasma Mass Spectroscopy (ICP-MS) techniques to measure very low concentrations of impurities, particularly ^{nat}K . They also developed their own highly-pure powders. Finally, Princeton partnered with Pacific Northwest National Laboratory (PNNL) to make highly sensitive tests of the ^{238}U and ^{232}Th concentrations in powders and crystals.

Seastar achieved a powder impurity level of 13 ppb of ^{nat}K , but the powder had some organic contamination that produced black residue in crystals. The focus has since been to capitalize on their ability to perform high-sensitivity measurements of impurities and to clean materials used in crystal growth. Their proprietary ICP-MS technique allows for measurements of ^{nat}K concentrations at the ppb level.

Sigma Aldrich developed a first batch of powder in 2010, which was tested to have lower concentrations of ^{nat}K than standard NaI powders (47 ppb and 300 ppb, respectively). Since then, they have produced the highest-purity powder currently available, called NaI AstroGrade, which has a reported lowest-achieved concentration of ^{nat}K of 3.5 ppb. Measurements by Seastar, however, estimate that the impurity level may be higher (~ 18 ppb) [95]. Two major batches of AstroGrade powder were ordered for the purposes of growing test high-purity crystals. The concentrations of impurities in the AstroGrade powder at the time of purchase are summarized in Table 4.1. The ^{nat}K level, of particular interest to SABRE, was measured separately by Sigma Aldrich and by Princeton for the 2013 sample. Princeton's measurement is shown in parentheses. The Rb level was evaluated by Seastar, while the U and Th levels were measured by PNNL. The purity of the Astrograde powder has been

Table 4.1: Impurity levels in the Sigma Aldrich NaI AstroGrade powder from the two batches that have been purchased. For the ^{nat}K levels in the 2013 batch, two measurements were performed. The 3.5-ppb number comes from Sigma Aldrich, while the 18-ppb number comes from an independent measurement done by Seastar. The 9.7-ppb number for the April, 2014 batch was measured by Sigma Aldrich. The U, Th, and Rb levels were measured by PNNL. Data from [95].

Batch	^{nat}K (ppb)	^{238}U (ppt)	^{232}Th (ppt)	^{87}Rb (ppb)
DAMA	<100	20	20	unreported
August, 2013	3.5 (18)	0.6	0.5	0.2
April, 2014	9.7	–	–	–

assessed for other batches, and has been found to vary from batch to batch. This is due to Sigma Aldrich’s production procedures, which are not performed in a clean-room environment; the final purity of the powder can therefore depend on concentrations of impurities in the atmosphere, some of which may be seasonally dependent.

4.1.3 Purification of powder

NaI powder is made by combining hydroiodic acid (HI) and sodium carbonate (Na_2CO_3). In order to achieve high purity in NaI powder, it is beneficial to consider the purity of these starting materials. The development of high-purity precursors to NaI has been a focus in SABRE’s work with industrial partners.

In addition to purification of the powder during the production process, it is also possible to purify the powder materials before the growth of a crystal. Efforts in this regard have been conducted both with NaI powder and TlI powder at Radiation Monitoring Devices Inc. (RMD).

4.2 NaI(Tl) Crystals

The linchpin of the SABRE program is its crystals. The effort to produce high-purity, high-optical-quality crystals has been one of the main thrusts of the SABRE experiment. The aim is to achieve lower backgrounds and a lower energy threshold than DAMA/LIBRA. SABRE

has partnered with Radiation Monitoring Devices (RMD) Inc., a crystal-growing company specializing in purification methods and crystal growth in detector applications, to develop crystals of unprecedented purity.

4.2.1 Crystal growth

In the growth of highly-pure NaI(Tl) crystals, there are a number of considerations that complicate the process. The first is the highly hygroscopic nature of NaI. When exposed to moisture, the optical properties of the NaI degrade significantly, yellowing the crystal and introducing cracks and cloudiness. NaI(Tl) crystals must therefore be protected from all moisture. NaI(Tl) is also very brittle and cracks easily, especially when it sticks to the container in which it is grown. In the crystal growth process it is therefore important to consider the physical integrity of the crystal during and after its formation.

From a purity standpoint, the crystals must be protected from impurities before, during, and after growth. The crystal materials could be exposed to impurities during the preparation stage, from the container in which it is grown, and from surface contaminants introduced during the cutting, polishing, and packing stage. Measures must be taken to make sure all materials and the atmosphere to which the crystal is exposed are as free as possible from contamination. This may include impurities in the containers, or radon contamination in the atmosphere in which the crystal is processed, for example.

There are several methods of crystal growth that can be used for the development of high-purity crystals. DAMA/LIBRA uses the Kyropoulos method of crystal growth, which in the past has been regarded as the most pure form of growth. In the Kyropoulos method, the raw material is melted in a crucible. A seed crystal is attached to a pole, which is lowered to the surface of the melt. The melt is slowly cooled through the pole, so that the crystal forms around the seed. Because impurities are segregated at the crystal boundary, leaving many impurities in the melt, and because the crystal does not touch the walls of the crucible, the Kyropoulos method is considered to be very pure.

However, SABRE has primarily pursued the Vertical Bridgman (VB) method of growth for its crystals. In the Vertical Bridgman method, the material is packed into an ampoule: a long, sealed container typically made of quartz. Because NaI(Tl) can stick to this material, an open container shaped like a long cup called a crucible may be used as an internal container to hold the powder. The crucible is in turn sealed in the ampoule after being packed with powder and pumped down to vacuum. The material is melted from the bottom of the crucible, and the crucible is slowly lowered out of the heating element so that the crystal is formed from the bottom up. Impurities are pushed up the volume during the formation of the crystal. Crystals grown by the VB method can be sectioned after growth, removing the most impure section at the top, then regrown. This multiple-growth method can be used as an additional purification step.

The VB method has several advantages over the Kyropoulos method. First, the volume is sealed in an ampoule and pumped to vacuum, so impurities from the furnace and the atmosphere inside the furnace are not a problem. Conversely, Kyropoulos growths are open to the atmosphere inside the furnace, which has unknown cleanliness. Secondly, the VB method has fewer limits on crystal size; the crystal width is limited by the diameter of the heater, but not its length, whereas it is more difficult to develop a large Kyropoulos system at RMD. Finally, RMD has a long history of experience with the VB method, and it was advantageous to make use of their expertise in this area. Though the Kyropoulos method is superior in that the crucible walls do not have a chance to leach impurities into the crystal, this effect can be mitigated in VB growths by using clean, high-purity materials for the crucibles, as is described in 4.2.3.

4.2.2 RMD facilities

RMD has an existing infrastructure for the growth of high-purity crystals. Crystal materials are prepared and packaged in nitrogen-filled glove-boxes like the one shown in Figure 4.1 with very low humidity (~ 1 ppm moisture). Parts to which the crystal is exposed, including



Figure 4.1: Glove-box used for crystal packaging at RMD. Parts coming into contact with the crystal are baked in the furnace on the left. The final polishing and packaging occur in this glove-box, which has a humidity level of around 1 ppm. A separate glove-box like this one is used for preparing crystals before growth. Image from RMD.

crucibles, ampoules, and reflectors, can be baked out in vacuum ovens attached to the glove-boxes. The furnaces that are already in place for VB growth, shown in Figure 4.2, are sized to grow medium-sized, 5-kg crystals. Finally, post-processing of the crystal can be done in a dry room with less than 5% relative humidity. This room, shown in Figure 4.3, contains a HEPA-filtered work-bench for polishing in order to prevent surface contamination from radon daughters and other particulates.

4.2.3 Testing the Vertical Bridgman method

The twin goals of the crystal growth at RMD are to achieve lower levels of impurities than DAMA and to create structurally sound crystals with good light yield and energy resolution. Focusing on the VB effort, SABRE has worked with RMD to evaluate their crystal-growth process for ways to improve the purity of the crystals².

²The initial effort at RMD was conducted primarily by RMD. Subsequent efforts to improve RMD's method were collaborative, with all members of SABRE involved.



Figure 4.2: Furnaces for Vertical Bridgman growth, with maximum diameter of 13.3 cm. Image from RMD.

Initial effort at RMD

In the initial contract with RMD, RMD grew several crystals with high-purity powder and normal-purity beads using VB in a first attempt evaluate their growth methods and to grow high-purity crystals. Several different types of crucibles were used. Some of the material from some of the normal-purity runs was evaluated by GDMS, showing a purification of a factor of ~ 4 in potassium compared with the starting purity of the powder.

Larger runs were conducted with AstroGrade NaI powder in crucibles with different materials. Of particular note is the crystal NaI002 (shown in Figure 4.4), which was used in the SABRE measurement of the Na quenching factor described in Chapters 6, 7, and 8. It was found that the growth process introduced some amount of contamination; the concentration of ^{nat}K at the center of the tip section of the crystal (which should have the



Figure 4.3: Dry room for crystal cutting and lapping/polishing at RMD. The room has a relative humidity level of 5%. The wire saw used for cutting is shown on the right while the dry box with HEPA filtration used for polishing is shown on the left. Image from RMD.

highest purity) was around 30 ppb. The concentrations were even higher at the tail of the crystals.

The optical properties of NaI002 were found to be superior to those of standard commercial NaI(Tl). RMD measured the energy resolution of this crystal and compared it with a typical commercial NaI(Tl) crystal. The results of this comparison are shown in Figure 4.5. Princeton also separately evaluated this crystal for use in the Na quenching-factor measurement described in Chapter 6 and found its scintillation yield to be very high. In the detector used for the quenching measurement, a light yield of around 19 p.e./keV_{ee} was achieved.

From this initial effort, we assessed that RMD's crystal growth method showed great promise for growing structurally sound, appropriately-sized, high-purity crystals. However, some work would need to be done to determine the sources of contamination in the crystal-growth process and to develop ways to mitigate them or eliminate them entirely when moving to high-purity growths.




Figure 4.4: RMD NaI(Tl) high-purity crystal NaI002, grown in a pyrolytic-graphite-coated crucible. (Left) Crystal after removal from the crucible. (Right) Crystal after cutting and polishing. Images from RMD.




Figure 4.5: Comparison of the optical properties of crystal NaI002 (blue) with a standard commercial NaI(Tl) crystal (red). The large peak of the blue curve at around 1800 counts is fit with a second red curve that was used to determine the energy resolution. Figure from RMD.

Small crucible tests

After the results of initial VB growths at RMD with high-purity powder showed that some contamination was occurring during the growth and handling process, it was hypothesized that some impurities could be leaching out of the crucible or the ampoule during growth. In order to waste as little starting material as possible and in order to increase the surface-area-to-volume ratio, small crucibles and ampoules were fabricated from different high-purity materials in order to test for leaching effects. In addition, these tests would also serve to test whether crystals would stick to these different materials and crack after growth.

SABRE further purified the materials for this test through special cleaning procedures. After the crucibles and ampoules were cleaned, small crystals were grown with them; the crystal material was measured for radioactivity, along with a sample of the NaI powder for reference. This was done to test whether impurities were added to the crystal during the growth process, and whether the addition of impurities might be attributable to the crucible or ampoule material. Another measurement was made on a normal-purity NaI(Tl) crystal grown by the Kyropoulos method at RMD's sister company, Hilger Crystals. This crystal is to be used in initial tests of the liquid-scintillator veto.

The results of the small crucible tests are shown in Figure 4.6. All of the crystals had a clear appearance and good structural integrity, except for G04, which cracked badly. Samples from each crystal were measured for impurities by our industrial partners. The results of these tests are summarized in Table 4.2. Unfortunately, as can be seen in the Table, some contamination had been introduced to the powder during previous handling, and so it is unclear whether the use of these new crucible cleaning methods will achieve the desired purity levels. However, results from sample G03 suggest that impurities were not added by the crystal growth process at a detectable level. As a result of this test, SABRE and RMD have focused on growing crystals using the materials in sample G03 for large crystal growths.



Figure 4.6: Crystals grown in the small crucible purity tests at RMD. G01 and G02 are in the top left and right, respectively, G03 and G04 in the middle left and right, respectively, and G05 at the bottom. Yellow appearance of some of the crystals is due to an optical effect of the camera. All of the crystals were colorless and transparent in the bulk. In the tails, there is some discoloration due to the segregation of impurities during crystal growth.

Table 4.2: Results of the small crucible tests at RMD. The materials used for the tests are left out for proprietary reasons, and the crystals are identified by sample number. Impurity concentrations are shown in ppb.

Sample	^{nat}K	^{238}U	^{232}Th	^{85}Rb
AstroGrade Powder	34	<0.1	<0.2	<50
G01	99	0.2	0.4	<50
G02	88	0.2	1.1	<50
G03	40	<0.1	<0.2	<50
G04	97	<0.1	<0.2	<50
G05	75	<0.1	<0.2	<50
NaI Hilger	870	0.2	0.9	<50

Table 4.3: Results of a blank test preformed at RMD. A crystal growth run was performed without any powder material. A test was performed before and after the procedure to see if any impurities were introduced by the crystal growth process. Concentrations are shown in ppb in the solution used for testing. The second column shows the native concentration of each isotope in the solution used for testing. The third column shows the concentration after the solution was put in contact with the container before the blank test, while the fourth shows the concentration in the solution after it was put in contact with the container after the blank test.

Isotope	Solution (ppb)	Before (ppb)	After (ppb)
³⁹ K	0.0010	~0.02	1.5
²³⁸ U	<0.00005	<0.0004	0.00014
²³² Th	<0.0001	<0.001	0.0004
⁸⁵ Rb	< 0.0001	<0.001	0.0040

Blank test

In addition to possible contamination from the crucible/ampoule materials themselves, it is possible that contamination may occur in the preparation and growth stage from other sources, such as the handling of the materials, the baking of parts, or the packing of the raw material. RMD performed a blank test with the crucible/ampoule combination used for crystal G03, performing the entire procedure for crystal growth but with no powder. This test used the same temperature and baking time as an actual growth run. The materials could then be tested for surface impurities.

The results of this test are summarized in Table 4.3. The growth and handling procedures do introduce some impurities, so there is room for improvement, but overall, the level of impurities added is already small compared with the impurity levels reported by DAMA in their crystals.

4.2.4 Growing large crystals with the Vertical Bridgman method

Because the powder was contaminated with some potassium, it is difficult to ascertain whether the desired purity level can be achieved with the current methods and a high-purity powder. However, the G03 crucible growth is promising in that it did not introduce




Figure 4.7: Standard purity medium crystal grown with the Vertical Bridgman method. This crystal was grown in a 95-mm ampoule for a final size of 3" wide and 4" tall, 1.53 kg.

a measurable level of radioactivity. Going forward with this material, RMD is working on the growth of large VB crystals with TII doping. They grew a first test crystal with a mass of 1.53 kg, shown in Figure 4.7. There was some sticking to the ampoule, and therefore the crystal was slightly smaller than planned.

Once the procedure for normal purity large growths are established, RMD will grow a first medium-sized high-purity crystal for SABRE to test in a liquid scintillator veto.

4.2.5 Future directions

SABRE will continue to work with RMD to produce crystals with higher purity than the DAMA/LIBRA crystals. Future efforts may focus on the handling procedures in the powder-packing stage and post-processing of the crystals, polishing techniques that can reduce surface contamination, and upgrading equipment to reduce radon contamination and to grow larger crystals.

In addition to efforts to improve the growth process, SABRE will work toward achieving better assessments of the crystal backgrounds. Seastar's ICP-MS techniques are not sensitive to the ^3H concentration, which could be an important background, for example. It may be necessary to find other methods that are sensitive to this impurity. SABRE will also work toward using the liquid-scintillator veto as a γ counter to measure residual ^{40}K , ^{238}U , and ^{232}Th backgrounds in the crystal, as will be discussed in Chapter 5.

4.3 The photomultiplier tubes

4.3.1 Goals

The photomultiplier tubes are an essential component of the SABRE experiment. They are the means by which the scintillation signal is detected, and therefore must be as free as possible from noise and false signals. The light collection in the crystal detectors is of paramount importance to achieving a low energy threshold. One way to accomplish this is to place the PMTs in direct proximity to the crystals, a strategy that requires that the PMTs be of very low background. Because of this, SABRE has worked to develop new phototubes with high efficiency, low noise, and low background.

When reading out a PMT signal, there are several considerations in the design of the voltage-divider circuit. In powering the dynodes, the dynodes further down the chain must be at a higher voltage than earlier dynodes, but the absolute voltage is not constrained.

Thus, PMTs can be operated in a positive high-voltage (+HV) scheme, in which the photocathode is grounded and the signal is produced at a high positive voltage, or a negative high-voltage scheme (−HV), in which the DC voltage at the anode is at ground potential, but the photocathode is at a negative high voltage.

There are advantages and disadvantages to each scheme. In the −HV scheme, the photocathode must be insulated carefully from other detector parts that may conduct a current, and also from each-other, in order to maintain a constant high voltage on the photocathode. This can make coupling PMTs directly to a NaI(Tl) crystal difficult. However, in the +HV scheme, the signal comes out with a high DC offset. In order to read out the fast AC signal, it must be decoupled from the DC offset with a decoupling capacitor circuit. Because these capacitors must have a high voltage rating, they are generally fairly bulky and can be a source of radioactive background if located on the PMT base³. Another alternative is to transfer the signal far away from the PMT in the same cable that powers the tube, and decouple the signal at a far-away location. This technique (called the one-cable design) is especially vulnerable to signal degradation and reflection effects; however, it bypasses the radioactivity problem from the decoupling capacitor and also makes the insertion of the crystal detector modules into a secondary veto detector simpler from a mechanical standpoint.

PMTs must pass an electrical current out from the vacuum interior into the outside world. In order to electrically isolate the dynodes from each-other, the feedthroughs must be composed of an insulating material, generally ceramics or borosilicate glass. However, these materials have a high radioactivity when compared to other components of the phototube. Ceramic is also problematic, as it has been shown to scintillate in the near-infrared region, to which the photocathode has a small amount of sensitivity [96] (see Figure 4.8). For these reasons, there is a strong impetus to both develop high-purity ceramics and to reduce the amount of ceramic material in the feedthrough plate. SABRE has worked to develop pin-

³The divider chain is often mounted on a “base” circuit that is attached directly to the PMT. Pins connected to the dynodes are fed through the back of the PMT and can be connected to the base on the outside of the PMT.

style feedthroughs that would reduce the amount of ceramic material in the feedthrough plate on the back of the PMTs.

Another potential concern for the PMT signal is noise and the dynode afterglow effect. Dynode afterglow is an effect in which a signal in one PMT is followed by a signal in another light-coupled PMT with a time difference of the PMT's characteristic transit time. These events can be at the single photoelectron level in both tubes, or can be larger, as described in [97]. It is believed to be caused by light emission of the dynodes in the cascade chain under electron bombardment [97]. This light can escape the tube of origin and trigger a cascade in another tube if they are light-coupled. Operating light-coupled PMTs in coincidence can drastically reduce known photomultiplier noise sources such as single-photoelectron (s.p.e.) noise and afterpulses, since these sources of noise should not be correlated between different PMTs. This coincidence requirement does not help with dynode afterglow.

In SABRE, this effect can in principle produce a coincidence event that is not caused by scintillation in the crystal, or could raise the apparent energy of an actual scintillation event. Though some techniques based on the asymmetry of the pulse size have been used to reduce these kinds of background noises, the effectiveness of these methods are reduced at low energies, where photon statistics create large variation in the relative scintillation pulse size [97]. In addition, a cut on the time difference between the phototube signals is impractical given the long scintillation time of NaI(Tl). Additional quality cuts based “steppiness,” a pulse-shape parameter describing the discreteness of a signal have been proposed by [97]; the effectiveness of these cuts has been demonstrated in signals larger than 4 keV_{ee} . However, these quality cuts do not have enough discriminatory power to adequately lower the afterglow rate for a dark-matter search, and their effectiveness at lower energies is unknown. In order to achieve an energy threshold below 1 keV_{ee} , the dynode afterglow problem has to be addressed (see Section 4.3.2).




Figure 4.8: Radioluminescence of the ceramic used in PMT feedthroughs. Data from Nerine Cherepy, figure from [96].

Table 4.4: Radioactivity level of the 3” Hamamatsu R11065-10-SEL PMT, as measured by γ counting at LNGS. Data from [98].

Isotope	Activity (mBq/tube)
^{40}K	17
^{232}Th	2.2
^{238}U	4.2
^{60}Co	5.4

4.3.2 Progress

The Hamamatsu 3” R11065 PMT was developed in collaboration with Hamamatsu Phototonics for the DarkSide experiment for low-background, cryogenic operation. This PMT has a very low radioactivity, at the level of ~ 10 mBq/tube, as summarized in Table 4.4. The low radioactivity is in part due to the use of a cobalt-free metal for the casing and the replacement of the feedthrough plate, which was previously made of borosilicate glass, with a high-purity (99.9%) ceramic. These tubes were used as the basis for further study into improvements for the PMTs.

Developing a low-noise base design

The Darkside-50 experiment uses a $-HV$ base design. For initial tests of the SABRE crystal detectors, SABRE is developing a $+HV$ scheme with one cable. The $+HV$ scheme was chosen to avoid electrically coupling $-HV$ PMTs to each other and to avoid having to insulate the phototube casing, which is electrically coupled to the photocathode on the R11065 PMTs. The one-cable design was chosen for ease of inserting the detectors into the liquid-scintillator veto detector with as few cables as possible and for the reduction of background contamination from the cables or decoupling capacitors. The Darkside base design is shown in Figure 4.9 while the new SABRE design is shown in Figures 4.10 and 4.11⁴.

The differences between the circuits were kept to a minimum in order to be able to modify existing DarkSide bases for the $+HV$ scheme. This was due to the high expense of Cirlex, the chosen material for the SABRE low-radioactivity bases⁵. With a minimally modified scheme, the old bases can be used to create the new divider circuit.

One difference between the two circuits is the capacitor, C7, which allows the output impedance of the base to be $50\ \Omega$ for high-frequency signals, while allowing the anode to operate at high voltage. This was necessary in order to eliminate reflections that occur due to impedance mismatching, since the cables are $50\ \Omega$. The high-pass filter consisting of components R22, R24, C5, and C6 was added to reduce noise from the later dynodes. Finally, a decoupling circuit is added in a decoupling box far from the PMT in order to separate the large DC offset from the fast signal. The decoupling circuit is shown in Figure 4.11. Resistor R26 was added in order to prevent signal from being lost due to any stabilizing circuit components that might exist in the HV source. R27 prevents ground loops. The capacitor C8 is the decoupling capacitor. R28 is set at a high value of $10\ \text{k}\Omega$ in order to

⁴The SABRE base design was developed and tested primarily by Dr. Jingke Xu and myself.

⁵Cirlex is easy to solder components to, is mechanically robust, and has a similar thermal expansion coefficient to copper, which allows for the operation of the NaI(Tl) detectors at lower temperatures if desired. The radioactivity levels are also acceptable.

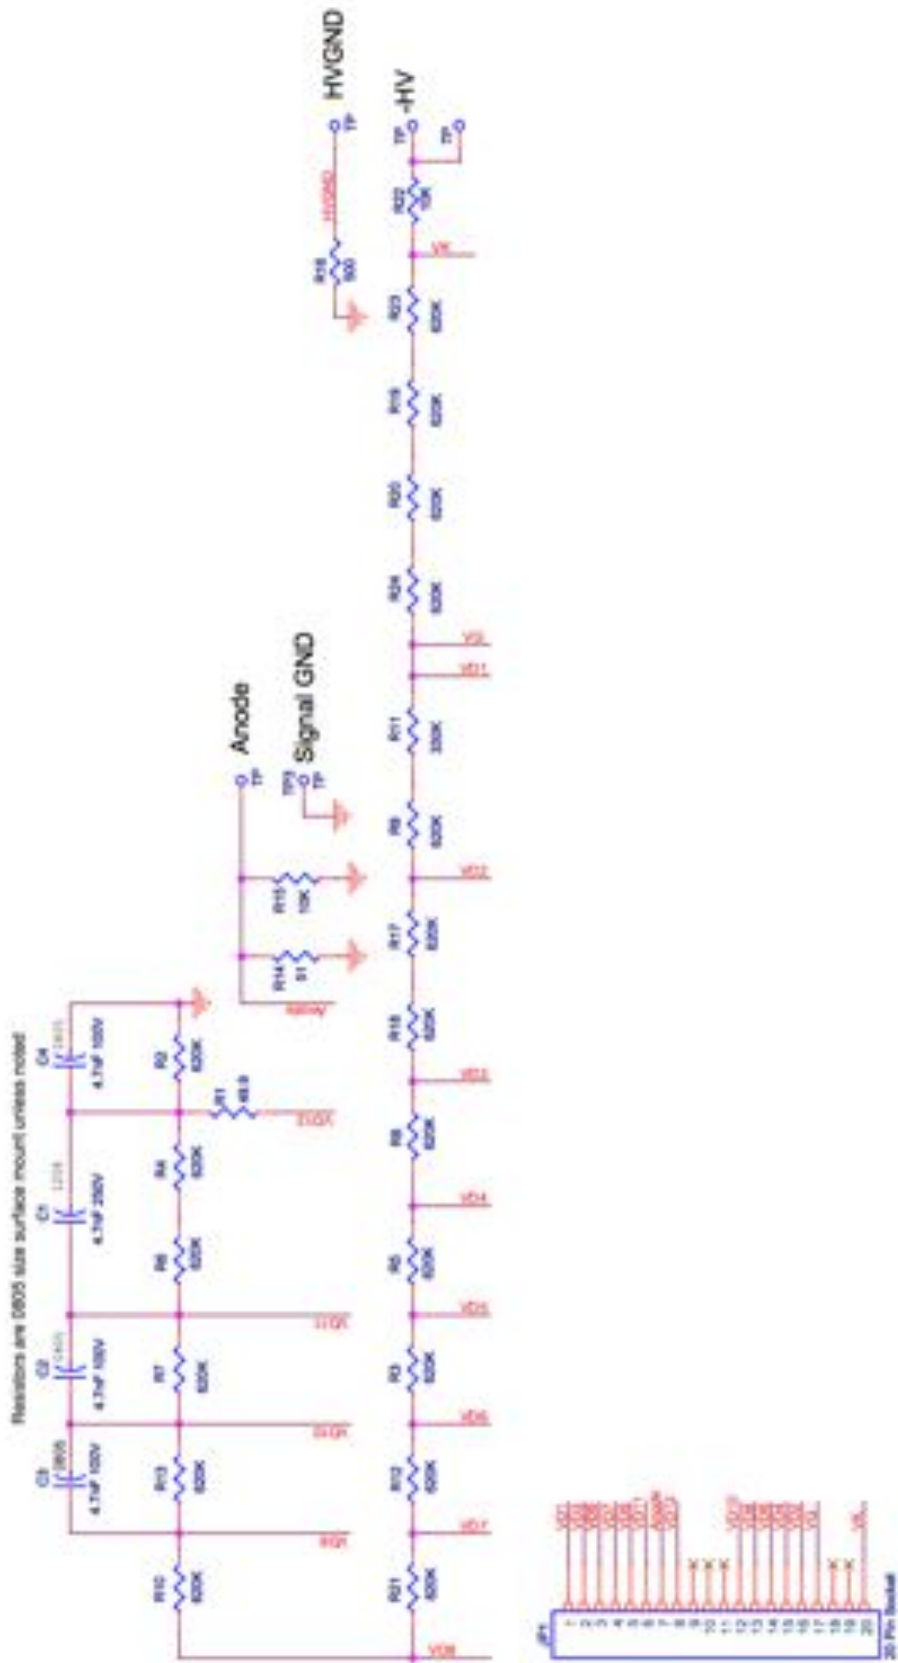


Figure 4.9: DarkSide-50 –HV base design, which was used as a template for the SABRE 3" PMT base design shown in Figure 4.10.

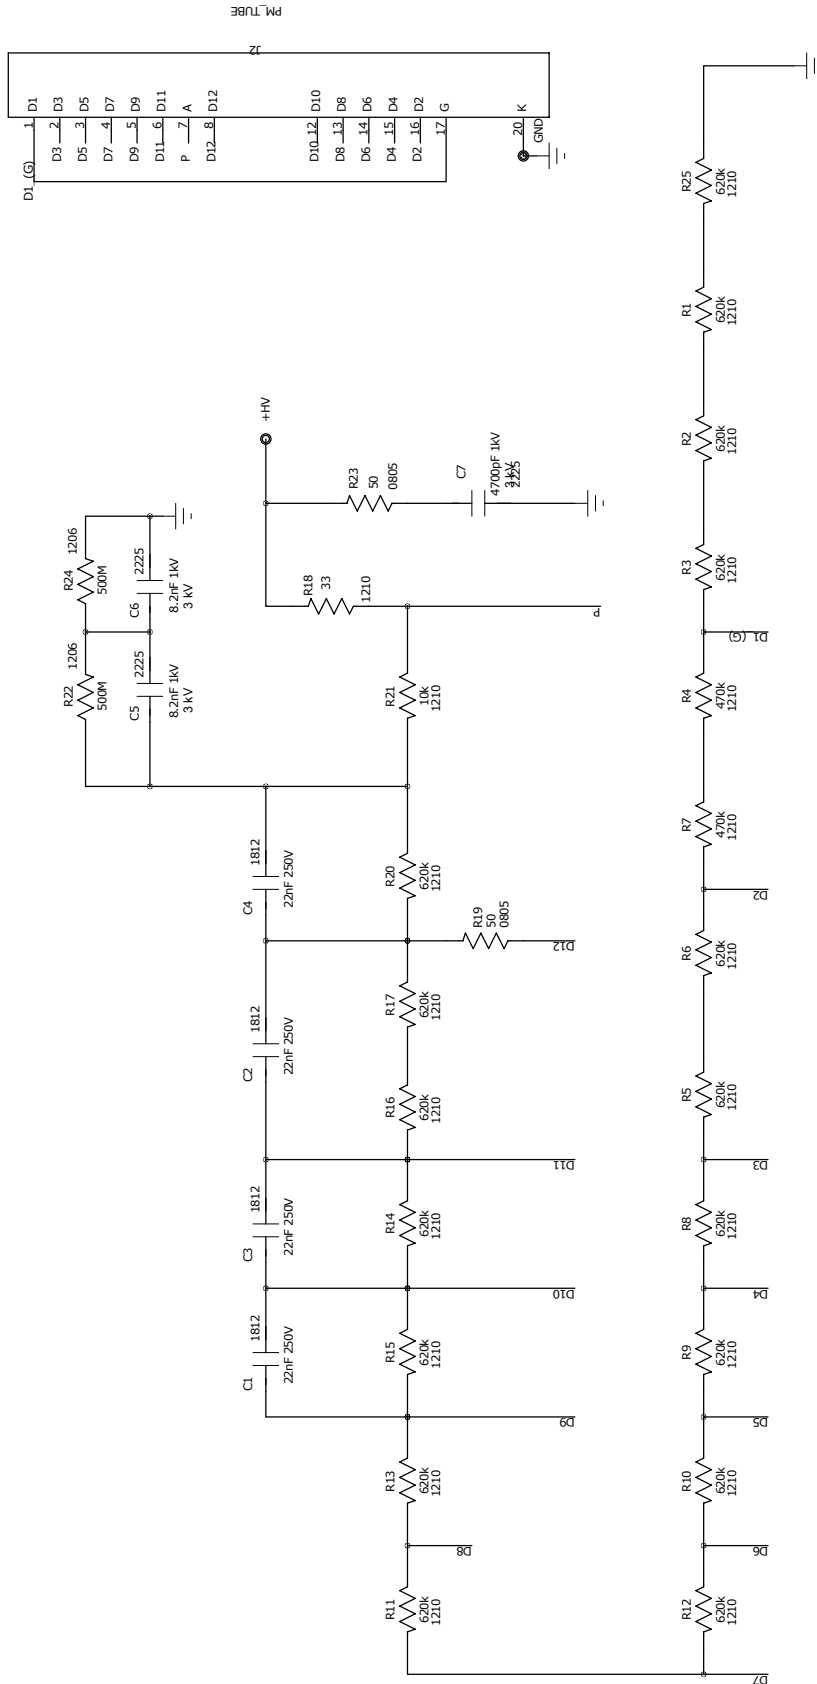


Figure 4.10: SABRE 3" PMT, +HV, one-cable base design, modified from the DarkSide scheme shown in Figure 4.9. Differences between the two circuits include the capacitor, C7, the high-pass filter consisting of components R22, R24, C5, and C6, and a decoupling circuit in a decoupling box far from the PMT, shown in Figure 4.11.

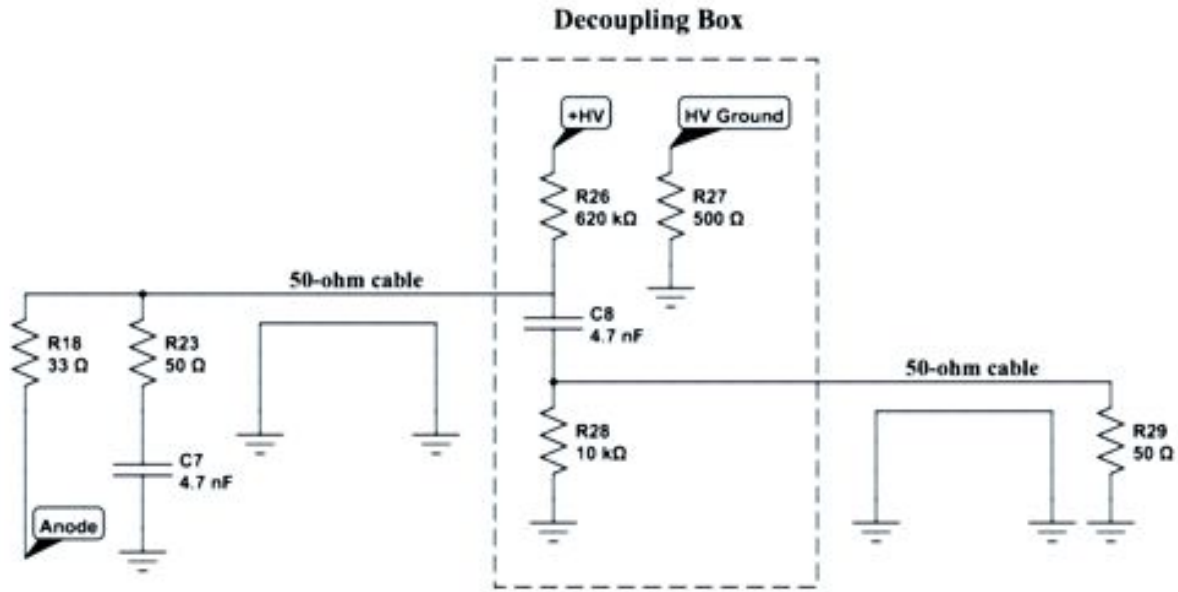


Figure 4.11: Preliminary decoupling circuit used for the SABRE one-cable design. Resistor R26 was added in order to prevent interference from HV source. R27 prevents ground loops, while C8 is the decoupling capacitor.

prevent the loss of signal and also prevents a DC offset in the circuit after the decoupling capacitor.

This base design, combined with the decoupling circuit, was tested with RG58, 50-Ω cables. Because R23 had to be given a low resistance to reduce ringing, the signal is split in half between the path through this resistor and the path through the decoupling circuit, thus effectively reducing the gain of the PMT. Because of this, the single-photoelectron peak was not easily resolved at low voltages. This problem can be ameliorated by appropriate amplification. A fast linear $\times 10$ amplifier was developed at LNGS, and can be used for this purpose.

For a dark-matter measurement, SABRE may return to the $-HV$ scheme in order to make use of the cold, on-base preamplifiers developed for the DarkSide experiment. These preamplifiers allow for the single-photoelectron peak to be resolved with an operating voltage as low as 900 V, which can be beneficial for reaching lower energy thresholds by reducing




Figure 4.12: Design for the PMT feedthrough plate with high-purity ceramic (courtesy of Dr. Francis Froberg). The ceramic feedthroughs, shown in white, allow for a pin to be fed through the plate to the dynodes. The feedthroughs can be brazed to the Kovar plate on the back of the tube.

dynode-afterglow noise, as described in Section 4.3.2. In addition, the $-HV$ scheme may also reduce signal loss from noise in the one-cable scheme.

Ceramic feedthroughs

A high-purity ceramic was developed by an industrial partner that uses high-purity alumina powder (99.999%) as its basis. The measured radioactivities of this alumina powder and the ceramic are shown in Table 4.5. This material was successfully brazed to a Kovar plate, demonstrating the feasibility of using this material for the pin feedthroughs. SABRE is attempting to further reduce the radioactivity by developing small, individual feedthroughs for each pin, rather than a single large plate⁶. The current design for such a feedthrough is shown in Figure 4.12. The design features individual ceramic feedthroughs to reduce the overall volume of ceramic material used in the PMTs. With this design, the total mass of ceramic per PMT is reduced from 16 g [99] to 5 g, assuming 16 feedthroughs⁷.

A study was done by Dr. Jingke Xu to demonstrate that a pin-style feedthrough would not introduce reflections in the output signal. In this study, the PMT signal output and

⁶This effort has been primarily led by Dr. Francis Froberg, though I helped in the initial design phase.

⁷The feedthroughs have been measured on a scale; each feedthrough is 0.33 ± 0.01 g [100].

Table 4.5: Measured radioactivity of the high-purity alumina powder and ceramic by γ -counting.

Isotope	Alumina powder		Ceramic	
	Activity (mBq/kg)	Concentration (ppb)	Activity (mBq/kg)	Concentration (ppb)
^{232}Th chain				
^{228}Ra	1.6 ± 0.4	0.4 ± 0.09	<7.4	<1.8
^{228}Th	2.0 ± 0.3	0.48 ± 0.07	8 ± 2	1.9 ± 0.5
^{238}U chain				
^{226}Ra	3.77 ± 0.35	0.305 ± 0.028	3.2 ± 1.6	0.26 ± 0.13
^{234}Th	< 44	< 3.6	< 36	< 2.9
$^{234}\text{Pa}^m$	36 ± 13	2.9 ± 1.0	< 72	< 5.8
^{235}U	0.8 ± 0.4	1.4 ± 0.7	< 4.3	< 7.5
Other				
^{40}K	18 ± 3	580 ± 90	260 ± 40	9000 ± 1000
^{137}Cs	0.35 ± 0.12		3 ± 1	
^{60}Co	< 0.032		< 0.86	
^7Be	0.32 ± 0.13			

ground were fed through two separated wires, between which were placed several different insulating materials, including Teflon and ceramic. The other components of the circuit had impedances that were carefully matched. No reflection of the signal was observed.

Reducing dynode afterglow

The dynode afterglow effect and its hypothesized cause is explained in Sections 2.2.3 and 4.3.1.

To study the rate of dynode afterglow, the dark coincidence rate at different operating voltages was measured⁸. A first test was done at Princeton. In this test, two Hamamatsu R11065 PMTs from the DarkSide experiment were coupled with optical coupling grease. Lumirror reflector (described in Section 4.4.3) was wrapped around the heads of the PMTs to block light from the outside and also reflect any light escaping the coupling. The PMTs

⁸Two measurements of the dynode afterglow effect are described in the text. The first measurement, done in Princeton, was conducted by Dr. Jingke Xu. The second, done at LNGS, was performed by myself along with Dr. Alessandro Razeto, George Korga, and Davide Sablone. An initial analysis was done by me, followed by a more detailed analysis by Dr. Xu. A final analysis with the updated data-acquisition software was conducted by me and is presented here.

were operated in a dark box with the same voltage, and the rate was observed at different voltages. The time-profile⁹ of the coincidence events is shown in Figure 4.13. A tall peak in random coincidence events with no delay was present at all voltages. These “prompt” events have a smooth, exponentially-decreasing energy profile. In addition, a “delayed” signal with a difference of ~ 50 ns was observed. In these signals, one channel would have an s.p.e. pulse, followed by a smaller pulse 50 ns later in the other tube. The rate of delayed coincidences was highly dependent on the operating voltage. The energy and time profiles of these events are shown at 1,600 V in Figure 4.14.

In a dynode-afterglow event, the initial cascade in one PMT could cause the release of a photon from one of the later dynodes. This photon could travel to the second PMT, triggering a cascade that would take the transit time of the tube to be registered. Because the transit time in the R11065 tubes is reported to be ~ 46 ns, this delayed signal was interpreted to be dynode afterglow. This observation prompted further study at LNGS.

In the LNGS study, the PMTs were coupled with a ring between them so that the PMT faces were not touching directly, as shown in Figure 4.15. The PMTs were separated to avoid voltage difference effects, since the PMTs were operated at $-HV$. No light-guide was used. The inside of the ring was coated with different reflectors for different runs. Teflon tape seemed to produce strange signals in the PMTs, so was not used. A copper foil was used in Run 1 and Lumirror thereafter. A cold pre-amplifier was used directly on the PMTs inside the dark-box to amplify the signal and compensate for the low gain used in some of the runs. The PMTs were then hooked to a custom fast linear $\times 10$ amplifier developed at LNGS. The output of this amplifier was sent to the digitizer. Coincident events were recorded and analyzed with the same data acquisition software described in Section 7.1.

The results of the LNGS afterglow study are summarized in Table 4.6. For each run, the random coincidence rate was determined by counting all events for $0.1 < |t| < 0.14 \mu\text{s}$,

⁹For all of the timing information, the “half-max time” was used to denote the start of a pulse. The half-max time is defined as the time at which the pulse reaches half its maximum amplitude during the sharp rise.




Figure 4.13: Time difference between the half-max time, the time at which the signal reaches half its maximum value, of coincident pulses in the Princeton dynode-afterglow study. The highest rate of events was seen in the central “prompt” events wherein events happened simultaneously in both PMTs. A secondary feature is a “delayed” signal which differs by ~ 50 ns. Plot courtesy of Dr. Jingke Xu.




Figure 4.14: Pulse integral of the PMT plugged into the digitizer channel 1 (left) and channel 3 (right) vs. time difference (channel 1–channel 3) in the Princeton dynode-afterglow study. The prompt signal has a smooth, exponentially-decaying energy profile, while the delayed events (shown in the red and blue circles) correspond with a s.p.e. in one tube (~ 150 ADC counts) followed by a smaller pulse in the other tube 50 ns later. The red circles correspond with events where channel 1 preceded channel 3, while the blue circle corresponds with events that started in channel 3. These events are interpreted as dynode afterglow. Plot courtesy of Dr. Jingke Xu.

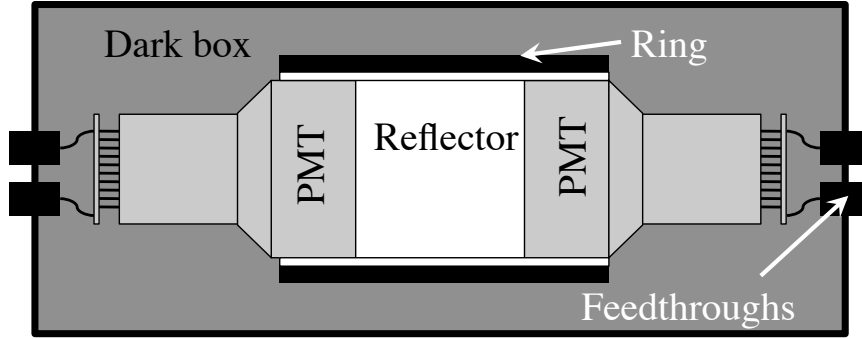


Figure 4.15: Setup for the LNGS PMT dynode-afterglow measurement. The PMTs were separated with an insulating ring, shown in black, that was coated with reflector on the inside. The setup was sealed in a dark box.

where t is the time difference between the two PMT signals, and dividing by the total run time. This time-difference region was determined to be sufficiently outside of the afterglow region. The rate of afterglow events was determined by subtracting the flat background from the peaks at ± 50 ns and integrating over the peak, and again dividing by the run time. The total afterglow rate as a function of voltage is shown in Figure 4.17.

The energy distribution of the afterglow events are shown in Figures 4.19 and 4.20. In Figure 4.20, the energies are shown for afterglow events initiated by PMT 1 and PMT 2 separately. The afterglow events are typically a single photoelectron in both PMTs, with some events having a larger signal in the first PMT to fire. This observation is in contrast to the Princeton study, where the secondary PMT produced a signal smaller than a single photoelectron. The reason for this difference is unknown.

In order to provide an estimate of the significance of the afterglow rate, the afterglow rate was compared to the rate of random-coincidence backgrounds in a similar time-difference window. In this comparison, the height of the afterglow peak was estimated by integrating over a small (4–8 ns) window around its highest point. This was compared to the flat random-coincidence background rate over a similar window. This comparison is shown in Figure 4.18. The afterglow peak was comparable to the random coincidence background rate when the tubes were operated at a voltage lower than 1100 V. Therefore, with the use of the cold, on-board pre-amplifier and operation at low voltages, the afterglow rate can be suppressed.




Figure 4.16: Typical time profile of coincidence events in the LNGS dynode-afterglow study, from Run 2 (Left) and Run 9 (Right). The right figure is operated at a higher voltage ($\sim 1450\text{V}$) while the left figure is at a lower voltage ($\sim 1000\text{V}$). Both runs exhibit a prompt peak where both PMTs see radiation simultaneously. The afterglow peaks at $\pm 50\text{ns}$ are visible only in the high-voltage run.

Table 4.6: Run information for the LNGS dynode-afterglow study.

Run	PMT1		PMT2		Coincidence rate	
	Voltage (V)	Afterglow rate (Hz)	Voltage (V)	Afterglow rate (Hz)	Prompt (Hz)	Random (Hz/ μs)
2	1055.0	0.030	954.0	0.020	2.19	2.08
3	1159.9	0.069	1052.4	0.034	2.26	3.78
4	1251.8	0.162	1149.1	0.073	2.33	6.19
5	1251.8	0.151	1149.1	0.073	2.33	5.38
6	1350.2	0.267	1251.5	0.139	2.36	5.89
7	1451.3	0.522	1350.9	0.261	2.39	7.30
8	1451.3	0.493	1350.9	0.245	2.38	6.75
9	1452.2	0.491	1450.4	0.479	2.39	7.10
10	1452.2	0.474	1450.4	0.455	2.40	6.70

It is notable that the DAMA/LIBRA PMTs are operated between 900 and 1300 V [46], so they may experience some dynode afterglow in some of their phototubes.

It is important to note that this study was performed in a dark environment with no scintillating material. The afterglow rate and signal size may increase with the number and size of real scintillation events. The afterglow effect can therefore distort the energy of real events if real scintillation signals produce secondary afterglow signals. Further study of this phenomenon and the development of additional means of suppressing it may be necessary.




Figure 4.17: Total afterglow rate in Hz as a function of the voltage of the first PMT to produce a signal in the LNGS afterglow study.




Figure 4.18: Comparison of the afterglow rate with random-coincidence background in the LNGS study. The integral around the peak of the afterglow was compared to the average random-coincidence rate for a similar time window. The dynode-afterglow rate becomes comparable to the random-coincidence background (y -axis value of 1) below $\sim 1,100$ V.

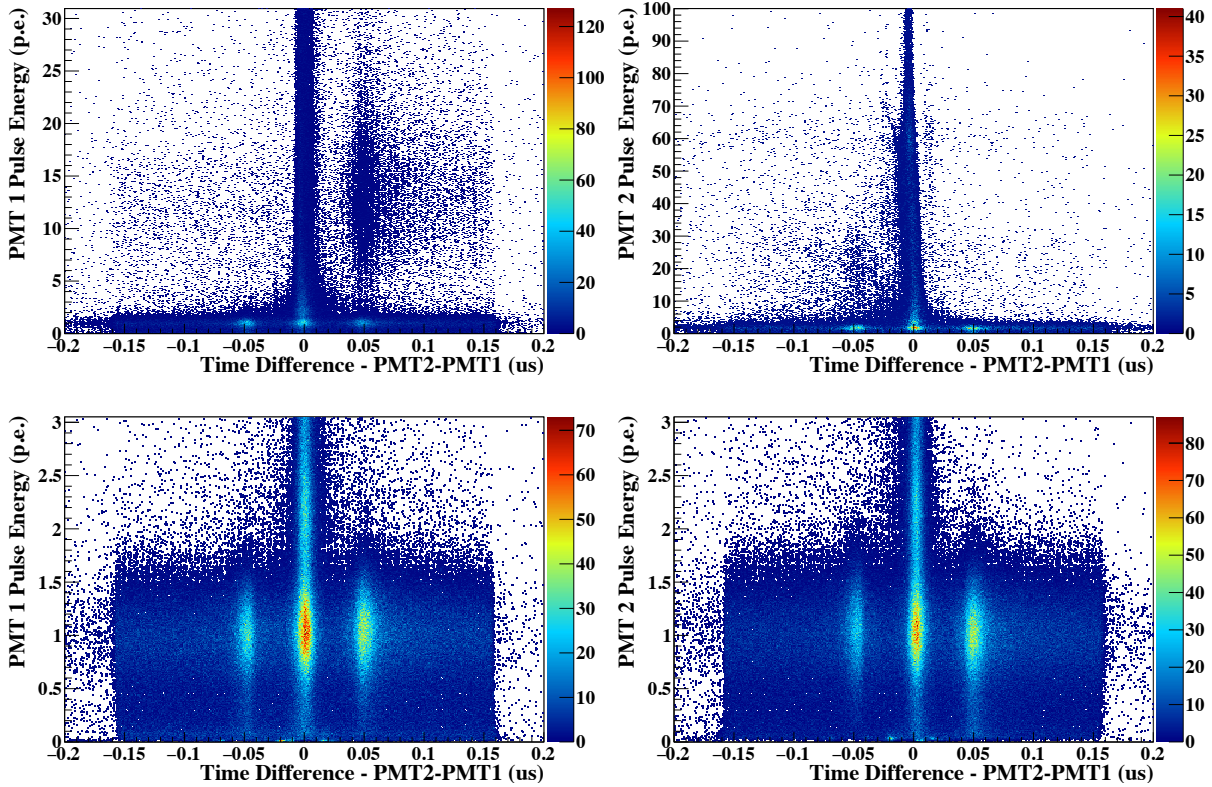


Figure 4.19: Energy of events in photoelectrons for the PMTs in the LNGS dynode afterglow study for Run 8. The bottom figures are zoomed into the 0–3 p.e. region. Most events have a single photoelectron in each PMT. Some events have a higher energy in the first phototube to fire, producing the asymmetry in the histogram, but the secondary signal is most often one photoelectron in size.




Figure 4.20: Afterglow light response in the two PMTs at LNGS. The energy seen in each PMT is shown for afterglow events where PMT 1 sees the event first (Left) or where PMT 2 sees the event first (Right). Regardless of the energy of the initial PMT, the secondary signal tends to be around one photoelectron. The exception lies in a small population of events for which PMT 2 fires first with one photoelectron, and PMT 1 fires later with a larger signal, as can be seen in the right figure.

4.4 Enclosure

Because of the hygroscopic nature of NaI, the need to prevent light from traveling from one crystal to another, and the need to protect the crystals from contamination, mechanical stress, and the liquid-scintillator, the NaI(Tl) crystals in the SABRE experiment must be enclosed in a light-tight, sealed container. The SABRE enclosure design must account for a number of complicating factors.

First, this container must protect the crystals from moisture. The enclosure must be able to be cleaned and dried, so as not to introduce moisture to the crystals. The enclosure must be sealed and leak-tight. The components holding the crystal and PMTs should not put much mechanical stress on the crystal or the PMT window.

A second role of the enclosure is to optimize the light collection in the phototubes, so that the energy threshold can be reduced. The crystal should be optically coupled to the PMTs in a way that maximizes light transfer, and the crystal should be surrounded by a reflector that allows more photons to reach the PMTs.

Because the enclosure will be immersed in a liquid-scintillator veto detector, all exposed components must be compatible with the scintillator. The enclosures must be heavy enough not to float in the scintillator, but light enough such that they do not block a significant amount of outgoing radiation from the crystals. The enclosure should contain feedthroughs for gas flow (to flush the crystals with nitrogen gas) and electrical signals. Ideally, the veto detector would be large enough that the enclosures could be spaced far apart, to maximize the veto power. For a limited veto size, however, the enclosures should be shaped so that they can be packed efficiently to minimize the volume of scintillator dead-space, since the light-collection in small spaces between detectors is drastically reduced compared with large spaces.

Finally, like all other detector components, the enclosure should be radio-pure, in order to not introduce background scintillation events in the NaI(Tl). All materials used for the enclosure must be screened and selected for radiopurity.




Figure 4.21: SABRE detector module constructed for initial tests. The crystal (a mock crystal in this picture) and the PMTs are held by a clamp structure that screws into a cap, which is sealed to the enclosure via a Viton o-ring seal. No reflector is shown.

The design for the SABRE enclosures consists of a cylindrical or hexagonal copper tube with one cap on each end¹⁰. One cap is brazed to the tube, while the other is removable and is sealed with an o-ring or other mechanical seal. The non-welded cap has a support structure attached that holds the crystal and the PMTs. The crystal is surrounded by a reflector and coupled to the PMTs by a thin optical coupling of a yet-to-be-determined material. Such a detector was built for initial tests of the crystal radiopurity and is shown in Figure 4.21.

4.4.1 Encasement

The enclosure will be made with copper, which can be made very radio-pure. Electroforming is a process that can be used to create enclosures with very pure chemical composition. In this

¹⁰The initial design for the SABRE enclosure was led by me with input from the other members of SABRE. Later incarnations were primarily developed by Dr. Francis Froberg with input from myself and Dr. Jingke Xu. Some of the parts for the various enclosures or prototypes were fabricated by me.

process, the material is deposited on a mandrel formed in the desired shape, which can later be removed. Pacific Northwest National Laboratory (PNNL) has produced electroformed copper enclosures for the Majorana experiment with very low levels of radioactivity, achieving a limit of $<1 \mu\text{Bq/kg}$ in ^{238}U and ^{232}Th [101, 102]. Electroforming is a slow process, forming at the rate of $\sim 0.04 \text{ mm/day}$, so it should only be used to produce the thin walls of the chamber. Because the walls of the chamber are closest to the crystal, they must be lowest in radioactivity and should also be thin. Electroforming is therefore ideally suited to the creation of the enclosure walls.

The caps, which must be more structurally sound, will be made from C10100 oxygen-free high-conductivity copper (hereafter OFHC). This material has a low radioactivity ($\sim 10 \mu\text{Bq/kg}$ ^{238}U and ^{232}Th) but not as low as electroformed copper. In order to pack the enclosures closely, the cap may be sealed to the enclosure via a shaft seal rather than a flange seal.

The crystal modules would ideally pack with no space in-between them, in order to reduce dead-space in the veto detector. It was proposed that the detectors could be formed in a hexagon shape. An initial test was conducted, in which a small hexagonal mandrel was constructed with a slight taper for ease in removing the copper shell. A layer of electroformed copper was deposited on the surface of the mandrel, and was found to be uniform, even at the corners, demonstrating the feasibility of creating polygonal enclosures by this method.

In order to ensure the safety of deploying the enclosures in the liquid-scintillator veto detector, tests were conducted by Dr. Francis Froberg to ensure that copper, unlike stainless-steel, has no degradation effect on the liquid-scintillator cocktail used in DarkSide-50. In this study, small scintillator samples were given exposure to copper pieces, representing a long period of exposure (>100 years) for the surface-area-to-volume ratio in the DarkSide-50 neutron veto detector. The transmission coefficient was measured with a Lambda-650 spectrophotometer. No degradation of the scintillator was observed, as can be seen in Figure 4.22.




Figure 4.22: Measurements of the attenuation length at different wavelengths as a function of time for liquid scintillator exposed to copper (Cu) and copper cleaned with sulfamic acid (Cu*). In the top figure, the scintillator used is pure pseudocumene (PC) while in the bottom figure, it is a mix of pseudocumene and trimethyl-borate (TMB), the scintillator mix used in the DarkSide-50 experiment. In both cases, a long exposure showed no degradation effect. Figures courtesy of Dr. Francis Froberg.

A conceptual design for the encasement is shown in Figures 4.23 and 4.24. These designs are not finalized, but demonstrate the concept for how the detectors might be sealed. The final design considerations will depend on the shape and accessibility of the veto detector and how the enclosures are deployed. Two possibilities for the deployment of the SABRE detectors are discussed in Sections 5.2 and 5.3. The hexagonal design with a shaft-sealed flange was proposed for optimal packing of the crystal detectors in the SABRE portable liquid-scintillator vessel for an initial deployment of seven crystals. The six peripheral crystal modules in the hex-packing arrangement would face alternating directions, such that the cylindrical sections where the caps seal do not overlap. The central crystal module would either have a smaller cylindrical section or an interior flange seal, with a smaller crystal inside. This design would allow for the crystal modules to directly touch each-other, with no dead-space in between. If optimal packing is considered a lesser priority, the enclosure design may exhibit different features; flange seals, for example, are easier to seal and machine, and may be considered more advantageous. In any case, the final design will feature the electroformed walls, which can be done on both cylindrical and polygonal surfaces, and an OFHC cap of some kind on each end.

The feedthroughs for the current prototype enclosure are bulkhead feedthroughs, but these feedthroughs may leak slowly over time. Future designs may include custom feedthroughs that are brazed to the cap, which may need to be made of a different material, such as Kovar, since copper may be difficult to braze to ceramic. In addition, the cap may include multiple feedthroughs for gas circulation, allowing the volume to be flushed with nitrogen gas, either at room temperature or cold. It may also include feedthroughs for monitoring environmental variables, such as temperature.

4.4.2 Support structure

The support structure for the crystal and PMTs used in our prototype is shown in Figure 4.21. Half-rings were machined from highly-crystalline Teflon to form a clamp around the




Figure 4.23: Hexagon-shaped enclosure design that allows for maximal packing of the detectors. The enclosure walls are made from electroformed copper, while the rounded section is machined from OFHC copper. The pieces are e-beam welded together, and a cap, shown in Figure 4.24, seals to the inner diameter via an o-ring seal.




Figure 4.24: Basic cap design for the SABRE enclosures. The cap will fit into the rounded section of the enclosure in Figure 4.23 with an o-ring shaft seal. Feedthroughs for gas and electrical cables (shown in beige) are arrayed on the top of the cap. The support structure for the crystal and PMTs screws into the bottom of the cap.

walls of the crystal. The pressure is distributed by the Teflon-tape reflector. Two more Teflon pieces cradle the PMTs. These clamps are held in place azimuthally and radially by three OFHC copper rods, and vertically by a series of OFHC copper sleeves, springs, and nuts. The rods screw into the OFHC copper cap discussed in Section 4.4.1. The crystal is wrapped with Lumirror reflector (not shown in the figure).

Several design changes may be made in the future. If a rigid reflector like Spectralon is used, as discussed in Section 4.4.3, the clamp system may not be necessary, or may clamp to the reflector. If a polygonal crystal is used, the mechanism for holding the crystal may change. Finally, for better protection of the crystal from moisture, it may be the case that a second, smaller enclosure holding only the crystal inside of the primary enclosure is warranted. This enclosure would have two fused silica windows coupling the crystal to the PMTs. Such an inner enclosure may better protect the crystal and make the packaging




Figure 4.25: Crystal encasement proposed by PNNL. An inner enclosure within the crystal detector module may better protect the crystal, but also may introduce more radioactivity and may decrease the efficiency of the veto by introducing dead space. Figure from [103].

process easier, but may also introduce more dead space and radioactivity, thus limiting the efficiency of the veto and increasing backgrounds. The additional optical surfaces that would need to be coupled together could also decrease the light yield. An example of such an inner chamber is shown in the PNNL-proposed design in Figure 4.25

4.4.3 Reflectors and coupling

In order to lower the energy threshold, the attainment of a high light-yield is imperative. The choice of a good reflector makes a large difference in the achievable light yield, since a typical photon can reflect many times before interacting with the phototube. Traditionally,

RMD packages their crystals by wrapping them in Teflon tape. Teflon must be very thick in order for the amount of transmitted light to become negligible, however. Lumirror reflector, a polyester-based reflector with a protective coating [104], is a promising alternative, with a high reflectivity at many wavelengths (see Figure 4.26). However, Lumirror has high concentrations of U and Th, and therefore may not be suitable as a reflector in close proximity to a high-purity crystal. Gore, another reflector with extremely high reflectance, may also scintillate. Spectralon is a rigid material that is also highly reflective and very radio-clean, and may reduce the mechanical stress on the crystal. However, it is also very expensive, and may be difficult to couple directly to the crystal surface.

The choice of reflector may depend on other factors, such as the type of optical coupling; optical gels tend to make Teflon tape and other void-based reflectors transparent. This does not affect Lumirror, which has a protective coating.

In order to increase the light yield of the detector, SABRE has chosen not to use a light-guide to separate the PMTs from the crystal. An optical coupling between these surfaces is necessary to reduce reflections at the surface boundaries. Finding an ideal coupling is difficult because of the high refractive index of NaI(Tl), 1.85, when compared with quartz, ~ 1.5 . Typical optical couplings have lower indices of refraction, around 1.4 or less, or are colored and lack transparency. The optical coupling between the crystal and the PMTs will be done in one of two ways. SABRE has used a transparent optical-coupling gel from Cargille Labs [105] with a refractive index of 1.52 to package the NaI002 crystal for the quenching measurement discussed in Chapter 6. Another option is optical gels that cure into a flexible pad, like Sylgard [106] or Angstrom optical gel AL-3252, which has a refractive index of 1.52 [107].

A preliminary enclosure has been constructed to test SABRE's ability to measure residual backgrounds, which will be tested with already-grown normal-purity crystals. SABRE is in the process of using the information described above to design an enclosure for the high-purity crystals being developed at RMD. Once built, these crystal detectors can be




Figure 4.26: The Lumirror reflector reflectance as a function of wavelength, as measured on a Lambda650 spectrophotometer. Lumirror is highly reflective over a broad range of wavelengths. It also uniquely maintains its high reflectance when soaked in liquid. The black spectrum shows the reflectance of Lumirror dry, while the blue curve shows the reflectance when the Lumirror has been soaked in liquid scintillator for 10 months. Figure from measurements done by me with formatting by Shawn Westerdale.

deployed in an active liquid-scintillator veto detector.

SABRE has made substantial progress toward the development of highly-pure crystal detectors, developing highly pure powders with industrial partners, establishing clean crystal-growing methods, establishing the operation of the PMTs and developing cleaner PMT components, and designing appropriate enclosures for the crystals. Work will continue in this area to improve upon the crystal detectors, while the current crystals and enclosure designs will be tested in an active liquid scintillator veto. The current plans for the deployment of the SABRE detectors are discussed in the next Chapter.

Chapter 5

The Active Liquid Scintillator Veto and Simulations of SABRE

5.1 Active background rejection

The three prongs of the SABRE approach to testing the DAMA/LIBRA result are the development of highly pure detector components, the lowering of the energy threshold, and the active rejection of residual backgrounds through the use of a liquid-scintillator veto detector. The DAMA/LIBRA experiment rejects some of its backgrounds through the selection of what they call “single-hit” events, defined as events in which energy depositions only occur in one crystal. However, the rejection power of this technique is highly dependent on the location of the primary scintillation event in the crystal array. Because of this, a 3-keV_{ee} feature from ⁴⁰K decay remains in the DAMA/LIBRA energy spectrum, among other backgrounds. The SABRE approach is to emulate the active rejection techniques of the DarkSide experiment [94], rejecting residual backgrounds by deploying the main detectors inside a liquid-scintillator detector. With this approach, lower backgrounds than those of DAMA/LIBRA may be achievable.

5.1.1 The veto concept

A common technique for reducing the background event rate in dark-matter detectors is to surround them with a large mass of material that can stop incoming radiation. This passive shielding technique is employed by the DAMA/LIBRA experiment. However, passive rejection techniques cannot reduce backgrounds that are caused by radioactive decay in the detectors themselves. By surrounding the detector materials with a liquid-scintillator vessel, outfitted with photodetectors to detect the scintillation light, particles that deposit energy in the main detector and then escape can be tagged on the way out. In addition, an active veto detector can tag incoming radiation on the way to the detector, making it unnecessary to stop the radiation completely. In this way, less shielding is needed to achieve the same rejection power. This technique has been successfully utilized by the DarkSide experiment [93].

In addition to rejecting backgrounds in a dark-matter measurement, the liquid-scintillator detector can be used as a γ counter to measure the concentration of impurities in the crystal and other detector components. This mode allows for quick (on the order of weeks) feedback on crystal-growth methods and the purity of the detector modules.

5.1.2 The need for active background rejection

The DAMA/LIBRA modulation amplitude peaks at around 3 keV_{ee} . In addition, the total energy spectrum also has a characteristic increase in the event rate around 3 keV_{ee} , which is commonly attributed to the decay of ^{40}K . ^{40}K decays by electron capture to ^{40}Ar in 10.7% of cases. In this decay, a γ ray with energy 1461 keV is emitted at the same time as the atomic transition causes the absorption of an inner-shell electron. This transition triggers a cascade that includes K- and L-shell X-rays or Auger electrons, some with an energy of $\sim 3\text{ keV}$. One of the main purposes of deploying the crystals in an active veto detector is the rejection of low-energy, internal ^{40}K backgrounds by a detection of the high-energy γ ray as it escapes the crystal volume.

In addition to ^{40}K in the crystal, external backgrounds from cosmogenic and radiogenic sources may enter the crystals and deposit a small amount of energy as they pass through. Of particular importance is the background of cosmic-ray muons that interact in the surrounding laboratory environment and shielding and produce neutrons. These neutrons can interact in the crystal via nuclear recoils, and are therefore immune to pulse-shape discrimination techniques. Though this background has been discussed in the literature [77, 79, 78, 76], and the modulation is unlikely to be entirely explained by muons, the possibility remains that this background may contribute an interfering modulation signal that could obscure a dark-matter modulation or change its apparent phase [77]. It is therefore extremely important to tag neutrons as they enter the crystal volume, a task which the liquid-scintillator veto could perform, especially if the scintillator is loaded with an isotope with a high neutron-capture cross section, like ^6Li , ^{10}B , and ^{157}Gd [29]. Hydrogen is also an effective target for fast neutrons, so scintillating hydrocarbons can be sufficient for tagging neutrons [29]. In addition to neutrons, the veto detector could also reject γ rays entering the crystal volume from the rock material, which contains ^{40}K , ^{238}U , and ^{232}Th .

5.1.3 The DarkSide-50 neutron-veto detector

A large, operational veto detector is already available in the form of the DarkSide-50 neutron-veto detector, shown in Figure 5.1. The DarkSide-50 veto is a ϕ 4-m stainless-steel sphere filled with a liquid-scintillator cocktail, lined with Lumirror reflector, and instrumented with 110 8-inch Hamamatsu R5912 PMTs. The total photocathode coverage is $\sim 1\%$. The scintillator cocktail contains the primary scintillator pseudocumene (PC), is loaded with trimethyl-borate (TMB), which has a high neutron-capture cross-section, and is doped with (2,5)-diphenyloxazole (PPO), a fluor that shifts the wavelength of scintillation photons to the primary detection region of the PMTs.

The DarkSide-50 neutron veto, shown in Figure 5.1, has a number of advantages for use as a veto detector for SABRE. The first is its location; it is situated inside a 10-m-tall, 11-




Figure 5.1: The DarkSide-50 detector setup. The DarkSide-50 liquid-argon detector is situated at the center of a $\phi 4\text{-m}$ spherical liquid-scintillator chamber called the neutron-veto detector. The neutron-veto detector is filled with a boron-loaded liquid scintillator for the efficient capture of neutrons and is outfitted with 110 8" PMTs to detect the scintillation light. The neutron veto is in turn encapsulated in the 10-m-tall muon veto detector, an extra layer of active water shielding. Four “organ pipes” (labeled as “pipes” in the drawing) extend from the neutron-veto detector to a clean room above the water tank, making the veto-detector volume accessible from the outside [94]. Figure from the DarkSide collaboration.

m-diameter cylindrical water tank, which serves as a muon-veto detector and also provides a large amount of passive shielding¹. Secondly, the veto chamber is outfitted with four “organ pipes” that extend from the neutron detector up to a clean room situated on top of the water tank. These pipes were designed to provide access to the liquid-scintillator volume for the insertion of calibration sources and, importantly, the SABRE crystal modules. With the DarkSide-50 infrastructure already in place, the DarkSide-50 veto detector provides an attractive option for housing the SABRE crystals, at least for initial internal-background

¹The muon-veto detector is equipped with 100 8" PMTs to collect Cherenkov light and is coated with a Tyvek-polyethylene reflective film to enhance the detection efficiency.

tests. The current plan for interfacing with the DarkSide-50 detector is described in Section 5.2.

5.1.4 A dedicated veto detector

Though the DarkSide-50 veto detector is a sound location for the crystal detectors, there are two main reasons to consider a dedicated, portable veto detector for the SABRE experiment. The first is location; with DarkSide-50 established at LNGS, the same location as the DAMA experiments, using DarkSide would make SABRE subject to many of the same environmental systematic effects as DAMA. A second benefit to a dedicated veto detector is lack of any compatibility issues with a second experiment. If SABRE uses the DarkSide-50 veto, SABRE is subject to the operations of DarkSide, which could potentially limit the possible exposure of a SABRE dark-matter run or cause delays in the schedule.

An independent veto detector of this kind was constructed for a proof-of-principle dark-matter experiment. Its design and plans for its implementation are described in Sections 5.3. In addition, simulations of this detector's performance and backgrounds are discussed in 5.4.

5.2 SABRE in DarkSide

5.2.1 The DarkSide experiment

DarkSide-50 is a liquid-argon dark-matter detector situated at LNGS in Assergi, Italy. The DarkSide-50 detector setup is shown in Figure 5.1. It has a well-established infrastructure with a significant active-background-rejection effort, notably in the form of its muon- and neutron-veto detectors.

The already established infrastructure of the DarkSide-50 veto detector, its location inside the muon-veto detector, and its accessibility through the four organ pipes shown in Figure 5.1 make the DarkSide-50 veto ideally suited to house the SABRE detectors. The Darkside-

50 veto detector can house the crystal detectors for initial tests of internal backgrounds, and possibly for a full-scale dark-matter detector. The primary effort toward deploying the SABRE crystals in the DarkSide-50 veto has been on the design and construction of an insertion system that would lower the SABRE detectors into the neutron-veto detector while allowing for the continued operation of DarkSide-50.

This insertion system will couple to a gate valve already in place in the clean room (CRH) on top of the water tank, which allows access to one of the organ pipes. Such a system must meet a number of requirements in order to maintain the integrity of the DarkSide veto detector and minimally impact its operation. The Sabre-in-DarkSide Insertion System (SIDIS) has been developed for this purpose and is currently being constructed.

5.2.2 SIDIS Requirements

The goal of SIDIS is to lower the crystal detectors through the organ pipes into the DarkSide-50 veto. It was designed to lower one 10-kg NaI(Tl) detector in a sealed module with the potential to be upgraded for two modules daisy-chained one above the other. SIDIS will lower the detector(s) 5.5 m into the main body of the liquid-scintillator vessel. A detailed description of all the requirements that were considered in the design is in [108]; the main points from that document are outlined here².

SIDIS must interface with the crystal detectors. It must be able to guide four cables (two signal and two high-voltage) from the crystal module out to the electronics in the clean room. These cables must be protected from mechanical stress as well as the liquid scintillator. In a dark-matter measurement, the system must be able to transfer a line of nitrogen gas to the detector to flush the volume around the crystal. A line for monitoring the crystal temperature is also desired.

The smooth and unimpeded motion of the detectors must be ensured. SIDIS must take into account the possibility that the detector would get stuck when re-entering the organ

²The SIDIS requirements document was primarily authored by Dr. Francis Froberg with input from myself and Dr. Jingke Xu.

pipe on the way out of the veto detector. There must be a smooth, tapered surface so the detector is guided into the organ pipe without the potential to catch. The position of the detector should be lockable, regardless of the detector position. It is also important to ensure that the gate valve cannot be closed when the crystal is below it, since the cables and whatever weight-bearing system holding the detectors will be in the way.

The liquid-scintillator cocktail used by DarkSide is hazardous to many materials and also to humans. In addition, it is sensitive to exposure to certain materials and atmospheric conditions like humidity, and could degrade optically if exposed to the laboratory environment. One of the primary requirements of SIDIS, therefore, is that it protect the scintillator from the laboratory environment and vice-versa. The system must be designed so that personnel and the laboratory environment are not exposed to liquid scintillator, and that any materials exposed to the scintillator should be tested for compatibility and should be precision cleaned. Because SIDIS will be dismantled after operation, a mechanism for draining liquid that enters SIDIS is needed, but due to exposure to the materials in SIDIS, the liquid cannot be drained back into the veto detector. Gas flow inside the system must maintain a dry, oxygen-free atmosphere, which should be monitored.

There are also mechanical requirements concerning the stability of the system and its compatibility with CRH. One complicating feature of the scintillator is that DarkSide fills the neutron-veto detector to the level of an overflow tank some distance above the floor of the clean room. Therefore, in order to keep the liquid level below the gate valve, the system must be operated at positive pressure.

SIDIS must also have accommodations for the possibility that the gate valve is not level. It must also be compatible with the existing infrastructure, fitting through doors, etc. Because of the underground operation at Gran Sasso, which is prone to earthquakes, the system must be able to withstand strong seismic activity.

The system must also ensure a minimal effect on the operation of DarkSide. The system must be light-tight so that the liquid-scintillator veto can remain in operation during the

measurement. The radioactivity of materials entering the veto should be low, not only to keep the background in the crystal due to insertion-system parts low, but also to not produce a background in the DarkSide-50 experiment.

The system should be monitored to ensure safe operation. The pressure of the system, the liquid level (should it rise above the gate valve), the atmospheric content inside SIDIS, the detector position, and any potential leaks should be monitored.

The system is not designed to be used for a vacuum, nor is it a pressure vessel beyond the requirement that it hold 150 mbar of pressure. It should therefore not be required that the insertion system be used in this way.

5.2.3 The design

SABRE settled on a linear-motion scheme wherein the crystal detector is lowered into the veto detector with a hose guided around a series of pulleys, shown in Figure 5.2. A previous design involving a wheel was abandoned due to the inability to control the winding of the hose. Another design with a removable glove-box was abandoned due to the pressure requirement and the inability to find a material that would be transparent, scintillator-compatible, and strong enough to withstand the over-pressure required to push the liquid level below the gate-valve³.

In the current design, the pulleys are brought closer together by moving a plate attached to two of the pulleys toward the third, thereby lowering the detector. This design avoids any twisting or length contraction effects that occur when winding a hose with internal cables. This motion is achieved by driving the plate with a worm drive. The electrical cables are contained in a convoluted PTFE hose wrapped in braided stainless steel, which holds the weight of the detector. The moving parts are contained in a sealed chamber that separates

³The initial efforts to design the insertion system, including previous designs, were led by me. Bob Parsells provided the initial conceptual design for the current system and helped choose components, and was heavily involved in the revision process. Once SABRE switched to the pulley design, I was initially heavily involved in developing the details of the design along with Dr. Francis Froborg, Joshua Pughe-Sanford, and Dr. Jingke Xu. After an initial design phase, the design effort was primarily headed by Dr. Froborg.

the scintillator vapor from the laboratory environment and provides a light-tight seal for the veto detector. Exterior interfaces with the system allow for the driving and monitoring of the system from outside of the chamber. The system is shown in Figure 5.2. A detailed description of the system is available [3], but the basic features of the system are described here following the main points in that document.

The motion

The convoluted hose containing the cables is wrapped around three pulleys and attached to the bottom of the SIDIS chamber on one end and the detector on the other. The pulleys are 7" cast-nylon sheaves⁴ with a round groove capable of accommodating a 1.25" hose. One pulley is attached to the bottom of the chamber (hereafter called the "bottom plate"). The other two pulleys are attached to an aluminum plate allowed to move up and down along three shafts extending from the bottom to the top of SIDIS. This aluminum plate is shown in Figure 5.3 along with the pulleys and shafts. The total range of motion of the moving plate is 2.0 m, allowing for the crystal detector to be lowered the full 5.5 m into the veto detector.

The moving plate is allowed to travel smoothly up and down along the three shafts thanks to linear bearings attached to each. These shafts prevent the rotation of the moving plate. They are held in place at the bottom of the chamber and at another aluminum plate near the top of the chamber, called the "top plate." A bumper on the moving plate prevents the pulleys from hitting the top of the chamber.

The motion of the plate is driven by a worm drive that allows the moving plate to move up and down along the central screw. The worm drive is a $\phi 1.5$ " screw with the capacity to bear 15 tons of centralized load. This value was calculated to be necessary given that the load is off-center compared to the axis of the worm drive, and because the worm drive is supported only at the bottom. The motion of the worm drive is driven from the outside of

⁴The pulleys were tested for compatibility with the scintillator and were found to be compatible.




Figure 5.2: The SABRE-in-DarkSide Insertion System (SIDIS). A moving plate provides the linear motion of two pulleys that lower a detector into the DarkSide liquid-scintillator veto detector. The motion of the plate is driven by a worm drive, and the entire system is contained in a sealed chamber filled with gaseous nitrogen at slight overpressure (~ 50 mbar). Drawing courtesy of Dr. Francis Froberg, from [3].




Figure 5.3: The moving plate that lowers the crystal detectors into the DarkSide-50 veto detector. The pulleys are attached to a moving plate that travels along the three support rods and is driven up and down by the worm drive. Figure courtesy of Dr. Francis Froberg, from [3].

the chamber, transmitted through a ferro-fluid rotary feedthrough attached to a port on the side of the chamber, as shown in Figure 5.4.

The signal transfer

The PTFE hose is connected to the detector on one end and the bottom plate of the chamber on the other. It holds the weight of the detector, and has been tested with a load of 23 kg. The hose has a convoluted PTFE interior, which is chemically resistant and highly flexible, and has a braided stainless-steel exterior for tension strength adequate to hold the weight of the crystal detectors. The hose protects the four electrical cables needed to transmit the signal from the two PMTs to the electronics in CRH in a $-HV$ configuration. The hose also contains a stainless-steel, nylon-coated support cable that serves as a back-up mechanism for holding the weight of the detector in the case of hose failure. It has additional space for a capillary tube to transmit nitrogen gas for the flushing of the enclosure.

The hose connects to the detector through an adapter protecting the connections of the electrical cables with the crystal enclosure. A concept of this adapter setup is shown in




Figure 5.4: System for moving the worm drive. The motion is transferred through a ferrofluid rotary feedthrough. The position of the detector is tracked via a rotary counter attached to the feedthrough. Figure courtesy of Dr. Francis Froberg, from [3].

Figure 5.5. The adapter itself or another piece will be constructed with a rounded or sloped edge to guide the detector into the organ pipe on its way out of the veto detector.

On the other side, the electrical cables and the backup support cable are terminated in a tee junction attached to the bottom plate of the SIDIS chamber, as shown in Figure 5.6. This feature allows the volume inside the hose to be isolated from the external laboratory environment.

The chamber/infrastructure

The SIDIS chamber consists of a bottom plate that attaches to the gate valve and supports most of the interior components, three aluminum cylindrical wall pieces stacked one on top of the other, and a plate that is welded to one of the top cylinder, forming the top of the chamber. The bottom plate, shown in Figure 5.7, connects to the gate valve with a bellows, to allow for the possibility that the gate valve may not be level. The bottom plate is supported by three adjustable legs. The worm drive and the support rods are bolted to




Figure 5.5: Potential design for the interface of SIDIS with the SABRE detector. A separate adapter (shown in grey) bolts to the enclosure cap (yellow) and connects to the PTFE hose (red). Inside, the steel support cable is terminated on the enclosure cap, and the electrical cables are connected to gas-tight feedthroughs (shown in beige). The adapter can be rounded on the edges to ease the reentry of the enclosure into the organ pipes, or a separate piece for this purpose can be attached to the hose above the crystal detector.

the bottom plate. They are also secured at the top plate to prevent tipping, though the top plate does not bear any of the weight of the worm drive or the support rods.

The cylindrical walls of the SIDIS chamber are outfitted with flanges and are connected to each other with Viton o-ring seals⁵. The lowest piece contains an extra flange face designed to attach the support panels shown in Figure 5.8. These panels bolt on the other end to the interior top plate to prevent it from rotating. The panels are outfitted with oblong bolt holts allowing for their alignment. The cylindrical walls of the chamber have been built and are shown in Figure 5.9.

⁵Viton is compatible with the scintillator, unlike other o-ring materials like Buna.




Figure 5.6: Termination of the electrical cables and backup support cable exiting the SIDIS system. A tee junction is attached to the bottom of the SIDIS chamber. The backup support cable (blue) is terminated without turning from the vertical by an adjustable turnbuckle. The electrical cables are turned to the side and terminated at electrical feedthroughs for easy access from the outside. Figure courtesy of Dr. Francis Froberg, from [3].

The chamber is outfitted with several exterior features providing an interface with the interior components, including viewports (discussed in the next section), gas feedthroughs, and the rotary feedthrough for the worm drive. There is one gas feedthrough at the bottom of the chamber beneath the bottom plate, which allows for nitrogen gas to be flowed into the chamber at the location closest to the scintillator, and another at the top for exhaust.

Monitoring

SIDIS contains several features for monitoring and safety. Several viewports (see Figure 5.10) are installed at the bottom and top of SIDIS to allow visual inspection of the SIDIS interior. In addition, the position of the detector can be quantifiably monitored by a rotary counter attached to the crank. A mirror can be installed on the floor of the chamber at a 45° angle to view up into the chamber, to manually inspect the detector location.




Figure 5.7: Bottom plate of the SIDIS chamber. The worm-drive, bottom pulley, support rods, and feedthrough tee from Figure 5.6 are bolted to this plate. An o-ring groove allows for the chamber walls to be attached separately. A bellows at the bottom attaches to the gate valve. The gas feedthrough above the gate valve allows nitrogen gas to be fed into the chamber close to the scintillator level, to protect the scintillator from humidity and oxygen. Three feet allow the angle of the chamber to be adjusted. A moat carved into the plate guides scintillator that may drip off the hose away from the organ pipe and out of the chamber. A Viton flap (shown on the left of the bottom figure) wipes scintillator off the detector as it rises above the gate valve. Figures courtesy of Dr. Francis Froberg, from [3].




Figure 5.8: Support panels to secure the interior SIDIS structure. The panels are designed to prevent the rotation of the top aluminum plate, securing the worm drive and the three support rods. The panels are bolted to the bottom chamber piece and also to the top plate. Figure courtesy of Dr. Francis Froberg, from [3].



Figure 5.9: The SIDIS chamber pieces.




Figure 5.10: Locations of the viewports for SIDIS. Once the detector is lowered, the viewports will be covered with light-tight cover plates. Figure courtesy of Dr. Francis Froberg, from [3].

5.2.4 Progress

The SIDIS system is currently being built. A test system has been set up in Princeton allowing a mock detector to be lowered several meters under a grate outside of the physics building. After all components are built and thoroughly tested in this setup, SIDIS will be disassembled, precision cleaned, and installed at LNGS. An initial test of the DarkSide veto detector's ability to measure the ^{40}K rate in a crystal will be conducted by lowering a standard purity crystal detector into DarkSide with SIDIS. The crystal and enclosure for this test have been obtained. Later, a measurement will be conducted with a high-purity crystal.

5.3 A portable veto detector

The SABRE veto vessel is a portable detector designed to actively reject a number of backgrounds associated with a NaI(Tl) dark-matter measurement⁶. It was designed to be deployed in multiple locations and contain a 50–100-kg dark-matter experiment. It was designed to reject 95% of ~ 1 -MeV γ backgrounds. In addition, it can also be used as a coincidence counter to estimate the radioactive impurity levels in the detectors, to provide quick feedback on the developing crystal-growth methods.

The SABRE veto vessel consists of a cylindrical, steel tank 1.52 m long and 1.37 m in diameter sitting on its rounded side. The crystal detectors will hang in the center from a port at the top. 10 PMTs, five on each end, will be installed in ports on the flat ends of the vessel to collect the scintillation light. The vessel will be filled with a liquid scintillator, either linear-alkylbenzene (LAB) or pseudocumene (PC), with the wavelength shifter PPO. It will be lined with Lumirror reflector to increase the light collection. The vessel will in turn be surrounded by passive lead and polypropylene shielding to keep the veto rate low.

5.3.1 Requirements and considerations

The SABRE veto vessel had a number of requirements that were considered in its design. In addition to the performance requirements determined by the physics of radioactive backgrounds, the vessel had a number of operation and safety requirements. The requirements that were considered in the design are as follows:

Physics requirements

1. The veto vessel should have a volume big enough to reject 95% of scattering events from γ s with an incoming energy in the ~ 1 -MeV range, particularly the γ ray from the decay of ^{40}K by electron capture.

⁶The design effort for the portable veto detector and overseeing its construction was led by myself.

2. The veto should maximize the light collected by the PMTs from radioactivity events in the liquid scintillator. The light yield should ideally be high enough to efficiently detect energy depositions of about 30 keV_{ee} , in order to potentially catch the fast-neutron signal from thermalization. The vessel must preserve the optical properties of the scintillator by protecting it from exposure to materials/gasses that may degrade its scintillation performance. Finally, the detection efficiency should be minimally affected by the event position to achieve a good energy resolution and avoid dead space.
3. The veto detector material should have a radioactivity that produces a sufficiently small background in the NaI(Tl) detectors (preferably less than 5% of internal background from the NaI(Tl) crystal, the 3" PMTs, and the copper enclosure).

Movement-and-installation requirements

1. The vessel must be transportable (but only while empty) and should be compatible with the infrastructure of existing laboratory spaces. In particular, it was desired that the vessel fit in the SNOLAB elevator so that the vessel could be deployed at SNOLab without being assembled underground.
2. The top-port cover plate should be liftable by one to two people and strong enough to hold the weight of a full-sized crystal array (up to 100 kg in a 19-crystal array).
3. The crystal detectors should be held in place by a yet-to-be-determined support structure.
4. The crystal detectors should be able to be removed from the vessel and replaced without disturbing or contaminating the vessel. We anticipate draining the vessel to change or access the crystals, but it may be done without emptying the vessel.

Feedthrough requirements

1. Electrical feedthroughs are needed to transfer the crystal-detector signals out of the experimental apparatus.
2. The vessel should be accessible for the insertion of calibration sources.
3. The pressure, temperature, and liquid level of the vessel and the temperatures of the crystal detectors should be monitored.

PMT requirements

1. The veto PMTs should be held in place against their tendency to float.
2. The veto PMT bases should be able to be inspected without emptying the vessel; therefore, they must be accessible.
3. The veto PMTs should be shielded from magnetic fields.

Safety requirements

1. The vessel must handle fluid pressure of 2 p.s.i.g. on the flat side walls.
2. The vessel should be designed to allow for variation in liquid level due to thermal expansion. This will require an overflow tank due to the large volume of scintillator and the desire for a large margin of safety in temperature variation ($\pm 20^{\circ}\text{C}$).
3. The vessel should be equipped with a means for handling leaks; there should be a means to contain leaks before they are otherwise addressed, and the vessel should be capable of being quickly drained into a separate container. This draining system should be both active and passive.
4. The vessel should be protected with passive systems such as burst discs, etc.

5. The laboratory environment should contain monitoring and equipment to deal with potential hazards, such as fire or earthquakes.

Non-requirements

1. The vessel was not required to be certified as a pressure vessel.
2. The vessel was not designed to hold vacuum.

5.3.2 The veto chamber design

The vessel was designed to meet all of the requirements listed in 5.3.1. The design for the vessel and how the design meets the requirements is described as follows.

Detector dimensions and orientation

The vessel is 1.52 m long and 1.37 m in diameter. The dimensions were chosen with the purpose of making the vessel fit in the entrance elevator to the SNOLAB facility (movement-and-installation requirement 1). To address the physics requirements 1 and 2, the efficacy of a detector of this size as a veto was determined by Monte Carlo simulation, as described in 5.4. To summarize, a calculation based on the attenuation of γ rays in scintillator shows that a vessel of this size would allow $\sim 5\%$ of 1461-keV γ rays to traverse its volume undetected, either from the crystals to the outside or vice versa, a distance of ~ 50 cm. This veto efficiency calculation assumes an array of 7 crystal detectors in a hex pattern, but does not include the detection efficiency of the γ ray by other crystal detectors. Rather, this number is an estimate of what fraction of γ rays entering the liquid scintillator volume leave undetected.

The vessel has a cylindrical wall thickness of 1/4" and a side-wall thickness of 3/8", for an overall mass of ~ 600 kg. The overall radioactivity of a vessel of this size is discussed in 5.4, to address physics requirement 3.

The vessel will lie on its side for more complete draining and for easier access to the PMTs, which will be installed on the vessel's flat faces. The supports for the vessel contain




Figure 5.11: To-scale conceptual drawing of the SABRE portable liquid-scintillator veto vessel. The crystal detectors, shown in brown, hang at the center of the cylindrical veto vessel. A port at the top of the veto vessel allows for the insertion of the detectors and also contains feedthroughs for necessary electronic cables and other sensors. The PMTs are inserted in ports on the sides of the vessel. A catch-basin is shown in red, which can hold 20% of the scintillator volume in the case of a leak. The vessel is surrounded by lead or steel shielding, which will provide a passive shield for γ rays and other radiation.

notches for a fork lift. Two lifting lugs on top also allow the vessel to be lifted by crane. These features make the scintillator vessel portable, as per movement-and-installation requirement 1. A schematic of the vessel is shown in Figure 5.11 and a final drawing in Figure 5.12.

The top flange

The vessel has one large flange at the top of the vessel, hereafter referred to as the “top flange.” The top flange allows access to the interior for the crystal detectors and other smaller components, as well as the installation of the reflector. The top flange has a 23.5” bore, so that a single crystal detector could be inserted horizontally into the detector, or a larger




Figure 5.12: Final drawing of the SABRE portable veto vessel, from Allegheny Bradford Corporation.

array can be installed vertically and then rotated, addressing movement-and-installation requirement 4.

The top flange has two bolt circles. The innermost bolt circle allows a cover-plate to be attached to the vessel. The outermost bolt circle can be bolted to a removable glove box or a flexible isolation system like a nylon bag for atmospheric control. This isolation feature will allow for operations to be performed in the detector without draining or contaminating it—for example, the insertion of calibration sources or the removal and reinsertion of the crystal detectors.

The top flange cover-plate is 3/8”-thick stainless steel and has been modeled to be strong enough to hold the crystal array in a Solidworks simulation, meeting movement-and-installation requirement 2. It contains several small ports that can be used as feed-throughs for electrical cables, sensors, calibration sources, etc., meeting feedthrough requirements 1, 2, and 3. A central 6” Conflat flange is big enough to house many of the needed electrical feedthroughs. Eight additional 2 3/4” Conflat ports surround the central port for remaining feedthroughs as well as a pressure sensor, temperature sensors, a potential connection to an expansion tank (for safety requirement 2), and any other necessary monitoring or control devices. The flanges are well-spaced so that tees or additional flanges could be installed if necessary. The cover-plate also has eye-loops on the bottom to secure a fixture that would hold the crystal detectors, and on the top for lifting, meeting movement-and-installation requirement 3.

The PMT ports

The side walls of the vessel are outfitted with 5 ports on each side for the attachment of PMT cans housing 8” Hamamatsu R5912 PMTs. This model of PMT was used in the DarkSide-50 neutron veto. In DarkSide, the phototubes were protected from magnetic interference with a mu-metal shield⁷, as shown in Figure 5.13. Each of the ports on the liquid-scintillator vessel

⁷Mu metal is a nickel-iron alloy with high magnetic permeability used in shielding electronics from static or low-frequency magnetic fields.




Figure 5.13: The mu-metal shield used to protect the 8" Hamamatsu R5912 PMTs in DarkSide-50. The veto vessel was designed such that a similar shield could be used on the veto PMTs in SABRE.

will have an inner bore of 10", which should be wide enough for the insertion of the PMT plus a mu-metal shield like those used in DarkSide (PMT requirement 3).

The ports were spaced far apart and were staggered on each side of the detector to minimize positional effects on the veto detection efficiency (physics requirement 2).

The PMTs themselves will be affixed to a can that will be bolted to the ports, so that the PMT photocathode is just inside the main volume of the veto chamber. Placing the PMTs all the way in the chamber would reduce the light-collection capability of the veto detector along the periphery, thus decreasing the effective fiducial veto volume. The cans feature a shaft seal on the neck of the PMT so that the base of the PMT could be accessed while the vessel is full (PMT requirement 2). The design is shown in Figure 5.14. A test of this seal was performed with a stainless-steel tube, a Teflon or stainless-steel insert, a PMT, and two o-rings, as shown in Figure 5.15. The tube was filled with water and watched for several weeks, and was found not to leak. Furthermore, the seal provided by the o-rings is such that the PMT does not easily slip out of its place in the can. Further tests of the integrity of




Figure 5.14: PMT can design for the SABRE portable veto vessel. The PMT is gripped by a Teflon or stainless steel insert, which cradles the PMT and provides a liquid-tight seal so that the base of the PMT can be accessed from the outside of the vessel. The port is wide enough to allow for the insertion of a mu-metal shield.

this seal under liquid pressure or under vacuum will be needed before trusting the integrity of the seal with liquid scintillator.

Fluid handling

The thermal expansion of the liquid will be accounted for by an expansion tank stored above the main vessel. The expansion tank can either be attached to the vessel at the top flange or at the Conflat port on the top of one of the side walls. The liquid level in this expansion tank can be monitored to detect leaks in the volume and to trigger draining of the vessel (safety requirements 3 and 4). The expansion tank can also be backfilled with nitrogen during the draining process to ease the draining of the liquid.



Figure 5.15: Test of the PMT grip design for the portable liquid scintillator vessel. A Teflon insert with o-ring grooves on both the PMT side and the can side creates a liquid-tight seal. This setup was tested with water for several weeks with no leaking observed.

The vessel will sit inside a catch basin that can hold 20% of the liquid's volume. A small Conflat port at the bottom of the vessel allows for liquid to be drained or pumped out of the vessel into other storage. Another Conflat port is placed at the top of the vessel on the opposite side to backfill the vessel with gas as it is drained. The vessel can be lifted or shimmed to create an incline for the complete draining of the vessel.

Coating and reflectors

Stainless steel has a known degradation effect on PC [109]. Therefore, it was decided that the interior of the vessel would be coated with ETFE, in order to prevent this degradation from occurring (physics requirement 2). The coating involves a sand-blasting process, in which a rough bead material is fired at the walls of the chamber to roughen the surface for better adhesion of the coating material. This process could embed some of the blasting material in the steel. In order to ensure that the blasting process would not introduce radioactive material into the veto chamber, we provided pellets of high-purity alumina (UHPA-AC2-OOS from Sasol Chemicals Co.) to serve as the blasting material (physics requirement 3).

A test of this coating was done on a bell-jar prototype veto chamber that had been used to predict the performance of the DarkSide-50 veto detector [110, 111]. This test revealed that the coating process would interfere with the ability to form a seal if done on a sealing surface. Therefore, sealing surfaces were masked off during the blasting and coating process. The thickness of the coating is 0.060”.

The vessel walls have welded studs for the hanging of the reflector, Lumirror, similar to those used in the DarkSide-50 experiment. Three staggered circles of eight studs on the round wall, two layers of circles on either flat face, four studs inside the bore of the top flange, and four on the bottom face of the top flange have been installed to fix the Lumirror to the walls. Two more studs were placed just outside the top flange bore for additional support along the ceiling of the vessel.

The Lumirror will be affixed to the walls by poking a hole in the material and pinning it to the wall with a Teflon nut and washer. The Lumirror will have several inches of overlap, as Lumirror loses its reflectance at the edges slowly when immersed in scintillator. This method was used in DarkSide-50⁸, as shown in Figure 5.16.

5.3.3 The veto chamber construction

The veto chamber was constructed by Allegheny Bradford Company (ABC) in Lewis Run, Pennsylvania. They specialize in the fabrication of clean storage vessels, and have fabricated other vessels for liquid scintillator processing for the Borexino experiment in the past. The blasting and coating of the vessel was performed by Fisher Company in Salt Lake City, Utah.

The vessel was constructed according to the atmospheric vessel standards in the American Society of Mechanical Engineers (ASME) code, Section 8 [112], and the welding standards in Section 9 [113].

⁸The installation of the reflector in DarkSide-50 was led by me and performed by Shawn Westerdale and myself.



Figure 5.16: Installation of the Lumirror reflector in the DarkSide-50 neutron veto. The Lumirror was held in place by a series of studs welded to the interior wall of the sphere and some Teflon nuts and washers. A similar procedure will be used for the SABRE veto detector.

Table 5.1: Radioactivity of the stainless steel used for the SABRE veto vessel, as measured by GDMS at Evans Analytical Group.

Lot Number	Thickness (in.)	U		Th		K		Pb
		ppb	mBq/kg	ppb	mBq/kg	ppb	mBq/kg	ppb
S536	3/8	0.3	3.7	<0.1	<0.4	4	0.12	10
T915	1/4	0.04	0.5	0.02	0.08	<1	0.03	9.3

Steel radioactivity

The radioactivity of the vessel steel must be low in order to not introduce a high level of backgrounds in the experiment. The steel used for the vessel was ordered from Stainless Plate Products USA Inc. (SPPUSA). Before the steel was ordered, samples from the available lots used were measured for radioactivity by GDMS. The measured radioactivity for the steel used for the SABRE vessel is very low and is summarized in Table 5.1. Samples of the steel have been sent for a γ counting measurement as well, which should give a more sensitive estimation of the steel radioactivity. The implications for this level of radioactivity for the SABRE background is discussed in 5.4.



Figure 5.17: The SABRE portable veto-detector vessel, constructed at Allegheny Bradford Corporation. The vessel has been placed on a custom-made track so that it can be rotated freely.

Testing of the vessel

Once fabricated, the vessel was tested for its ability to seal with a helium leak test and a water leak test. Blank flanges were fabricated to cover the PMT ports in place of the PMT cans for this test. The helium leak test at ABC was unsuccessful due to a non-functional leak tester. Before it was realized that the leak-tester was not working, the vessel was inadvertently over-pressurized and the flat walls bowed outward slightly. To compensate for this, an additional circle of studs was placed on each wall, to better secure the Lumirror against the now-curved surface. The vessel failed initial water leak tests on three of the ports, but additional material was welded into the weld joints and subsequent water tests showed that the vessel was leak-tight using 4 psi of pressure for 120 minutes. The constructed vessel is shown in Figure 5.17.

5.3.4 Electronics scheme for the liquid-scintillator veto

The liquid-scintillator detector will be initially deployed in Princeton, where it will be filled with water and tested. The detector will be lined with reflector, and an insertion system for the crystal-detector modules will be constructed.

Electronics for the data acquisition of this detector have been purchased, and a scheme has been conceptually developed for both a dark-matter measurement and a ^{40}K measurement. The scheme for the data acquisition for a ^{40}K measurement is shown in Figure 5.18. The waveforms from the PMTs both in the crystal detectors and the veto detectors will be digitized by CAEN V1720-E digitizers after being appropriately amplified. The digitizers send a low-voltage-differential-signaling (LVDS) discriminator signal to a V1495 logic unit, which provides a fast trigger for the digitizers. There, the PMTs for each crystal detector are operated in coincidence. In a K-measurement mode, an OR logic is performed on the crystal detectors and ANDed with a majority signal from the veto-detector PMTs to produce the trigger. In a dark-matter measurement, the OR signal from the crystal modules (composed of the ANDed output of the two PMTs for each crystal) would provide the trigger. Once the digitizers are triggered, they will send data for a fixed acquisition window to the computer. The veto-PMT signals could be acquired separately. Alternately, to reduce the data-storage requirements, a majority discriminator output from the veto PMTs could be recorded instead. It is also possible to operate without a trigger and use zero-suppression techniques to make the data manageable, performing all logical analysis offline. This technique would eliminate any dependence on a trigger efficiency, but would be data-intensive.

5.4 Simulations of the SABRE experiment

The success of the SABRE portable veto detector will be dependent on its ability to shield the crystal detectors from and actively reject backgrounds. This will depend on the the ability of the veto to tag or stop incoming and outgoing radiation, which is primarily determined

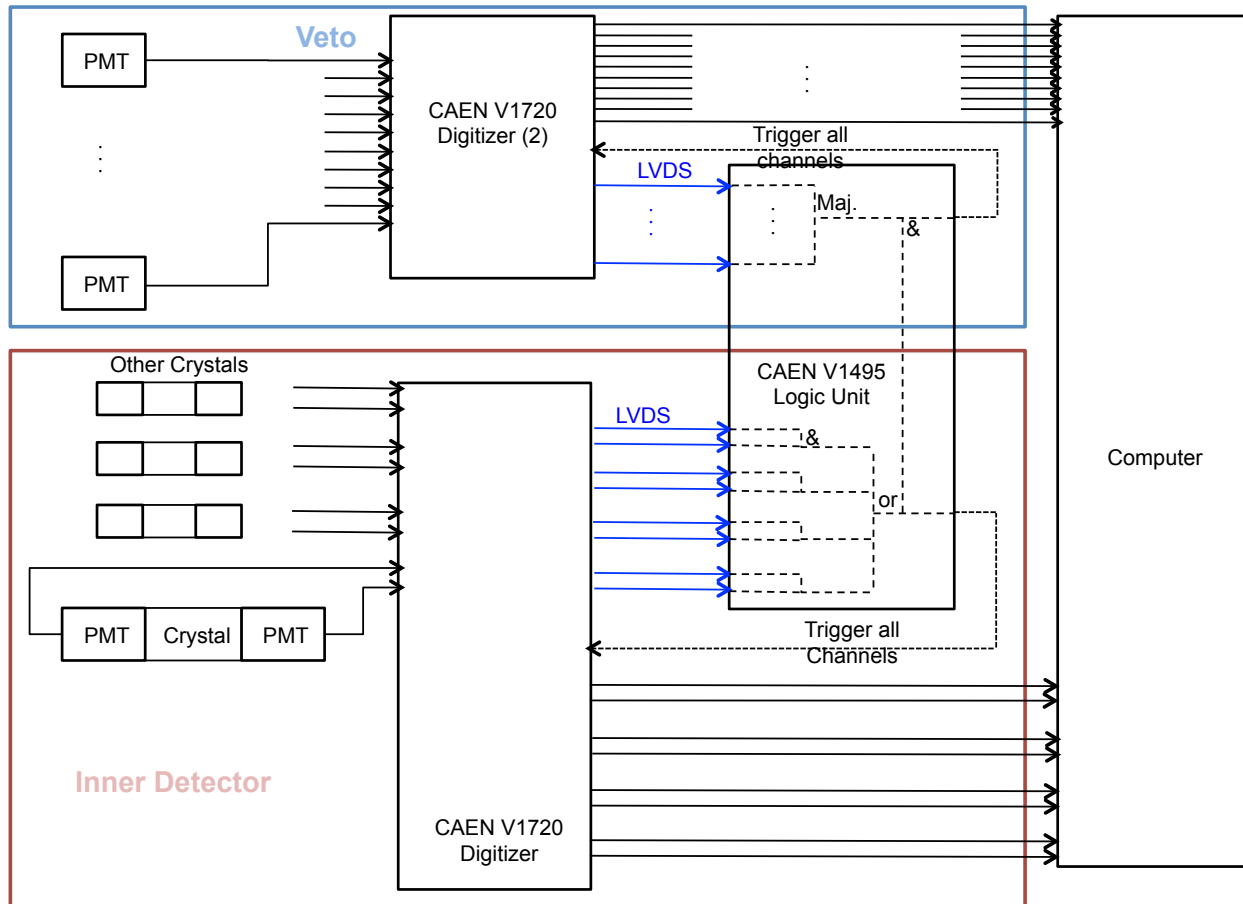


Figure 5.18: Electronics scheme for the SABRE detector in a high-data K-measurement mode of operation. The PMTs from the crystal detectors and the veto detector are fed into CAEN V1720-E 250 Msample/s digitizers after appropriate amplification. The digitizers send a discriminator signal to a V1495 logic unit that would perform the trigger logic for the digitizers.

by its mass. It also depends on the efficiency with which it detects energy depositions within its volume. Simulations were performed to predict both parameters. Geant4 Monte Carlo simulations predicted the backgrounds with and without the veto in a dark-matter-measurement mode, as well as in a coincidence-counting mode meant to measure the ^{40}K concentration in the crystals. These simulations are discussed in Section 5.4.1. A custom optical Monte Carlo predicted the light yield of the veto detector, as described in Section 5.4.2.

5.4.1 Backgrounds in the NaI(Tl) SABRE detectors

A Monte Carlo simulation was performed to predict the backgrounds in the SABRE experiment given the radioactivity in various components⁹. The simulation was performed in the most recent version of Geant, 4.10.01. The framework, called Xsimu, was developed by Dr. Jingke Xu. In this framework, a primary particle is generated (a primary event) and tracked by the simulation software. Sub-events that occur in defined sensitive volumes are recorded as hit objects in a ROOT TTree. For each primary event, the primary-event-level information is also stored in the form of three hit objects. The first gives the event ID, the second the position of the primary event, and the third the direction of the primary particle. The information recorded is shown in Table 5.2.

There are a few notable features of the simulation framework that are designed to prevent potential pitfalls. The timing information is especially important for radioactive decay chains in which the decay time may be much longer than the time for which the resolution of a double variable becomes relevant to distinguishing events that are nano-seconds apart (10^{15} nanoseconds, or 11.6 days). If the global time of each sub-event was stored, sub-events resulting from a radioactive decay with a long half-life may appear as if they happen at the same time when in reality they do not. In order to compensate for this, particles that are created with a high global time (>1 s by default) are killed, the particles are stored as a hit with identifiers showing they are killed particles, and they are recreated with event_time equal to 0.

The physics list used for these simulations is a modified version of the Shielding Physics list recommended for low-energy, underground applications by Geant4. Atomic relaxation fluorescence, Auger electron emission, and particle-induced X-ray emission were enabled, but it was found that the energy spectrum due to these effects were skewed due to a lack of conservation of energy. It was found that this effect was due to Geant4 only tracking a single hole in the atomic shells when in many Auger electron cascades, the Auger process

⁹These simulations were performed by myself and the results were validated by Dr. Jingke Xu.

Table 5.2: Structure of the Xsimu simulation output. The output is a TTree with the following branches. Each primary event has three events that serve as a header for the entire event, followed by sub-events associated with individual energy depositions. The third column describes the contents of each branch in the header events, while the fourth column describes the content of each branch for the individual sub-events. Some of the values in the third column are placeholder strings or constants to mark the event as a header, such as -1 or “<-EventID.” An int32_t is a signed-integer data type with size 32 bits, used to correctly align the data in the TTree object.

Branch name	Data type	Primary event-level hits contain	Sub-events contain
particle_id	int32_t	event ID	ID of the particle producing the interaction, where 0 is a primary particle
particle	char[16]	“<-EventID”	Name of the particle (e.g. “gamma”)
parent_id	int32_t	-1	ID of the particle producing this particle
prod_proc	char[16]	“<-Flag”	Production process for this particle (e.g. “primary,” “compt”)
event_time	double	-1	Local time of this sub-event (described in text)
pos_x_cm	float	x position or direction	x position of this sub-event in cm
pos_y_cm	float	y position or direction	y position of this sub-event in cm
pos_z_cm	float	z position or direction	z position of this sub-event in cm
interaction	char[16]	“New”	Type of interaction (e.g. “RadioactiveDeca,” “Transportation”)
e_inc	float	0	Incident energy of the particle (keV)
e_dep	float	0	Energy deposited in the volume (keV)
volume	char[16]	“Event”	Name of the volume where the subevent occurred

produces multiple holes. A fix was implemented by Burkhant Suerfu and Dr. Jingke Xu, allowing Geant4 to track all of the holes created by the Auger process. This fix was found to recover energy conservation in the simulation.

In addition to the Auger electron fix, another bug in the Geant4 code was discovered in the way that Geant4 reads the tables storing the radioactive decay information. The bug was first discovered by observing the decay of ^{234m}Pa in the ^{238}U decay chain. The meta-stable excited state of ^{234}Pa has an energy of 73.92 keV and decays directly to ^{234}U with a 99.84% branching ratio (BR). In the other 0.16% of cases, it de-excites to the ground state before decaying to ^{234}U . In the simulations, ^{234m}Pa was de-exciting to the ground state in 100% of cases. In addition, the de-excitation to the ground state was occurring in a non-physical way. The de-excitation most commonly occurs by releasing a conversion electron with an energy of 52.8 keV (0.11% BR) or 68.6 keV (0.027% BR). This should trigger a cascade of multiple X-rays and Auger electrons with low energy, but instead Geant4 generates a single γ ray with the remaining energy, in this case producing a prominent peak in the energy spectra at ~ 21 and ~ 5 keV_{ee}, respectively. These peaks can be seen most prominently in the background spectrum for ^{238}U in the PMT window, shown in Figure 5.19.

In older versions of Geant, the radioactive decay module, G4RadioactiveDecay, would find the correct entry in the table for a given parent particle by comparing the parent particle's energy with the energies in the table. Once a match was found to within a 2 keV tolerance, the branch information was retrieved from the table. When Geant4 added multi-threading to their simulation software, they switched to a model wherein a quantity called an “isomer level” is assigned to each excited state of an isotope. The radioactive decay information is retrieved by comparing the isomer level with the entry number in the table.

However, in Geant 4.10.01, the isomer levels are not set correctly when the table of isotopes is created, and all isotopes have an isomer of 0 (ground state) or 9 (any excited state). When G4RadioactiveDecay does not find a 9th entry in the data table, it defaults to a de-excitation to the ground state (called an isomeric transition by Geant) and decays




Figure 5.19: Comparison of the 3" PMT background with the liquid scintillator veto from the upper ^{238}U chain before (Red) and after (Black) applying a fix to the G4RadioactiveDecay module. In the red curve, a non-physical 21-keV_{ee} peak and a smaller one at 5 keV_{ee} can be seen due to the incorrect decay of ^{234m}Pa . In the black curve, these peaks have been eliminated, causing a slight increase in the smooth background below 10 keV_{ee}. The removal of these large peaks expose smaller peaks from other decays in the chain, such as ^{230}Th (the 19-keV_{ee} peak). These peaks are due to a problem with “isomeric transitions” as described in the text.

from there. This first problem explains the incorrect application of the branching ratios in the decay of ^{234m}Pa —it furthermore affects all decays from excited states, meaning the problem was further-reaching than originally thought. This problem was fixed by returning to comparing the energies rather than the isomer level, which, according to Dr. Laurent Desorgher, one of the primary authors of the G4RadioactiveDecay module, is the solution which will be implemented in an upcoming patch to Geant 4.10.01 [114]. After implementing this fix using code provided by Dr. Desorgher, the decays from excited isotopes now follow the correct branching ratios. The 21-keV_{ee} feature is now absent, and the 5-keV_{ee} feature reduced, as can be seen on the right in Figure 5.19.

However, the second problem, wherein the isomeric transitions to the ground state release a single γ ray with a fixed energy, rather than executing a cascade where appropriate, remains an issue that distorts the energy spectrum. In Figure 5.19, a new, smaller peak

at ~ 19 keV can be seen that results from the isomeric transition of ^{228}Th or ^{224}Ra . The isomeric transitions use the photon evaporation module, `G4PhotonEvaporation`, to de-excite to the ground state, which does not trigger a cascade through `G4AtomicRelaxation`, even if conversion electrons are ejected from the atom. This has the effect of concentrating the expected background at specific peaks, which will slightly lower the featureless background in the region of interest and distort the energy spectra. Future simulations should attempt to address this issue for a more accurate estimation of the background. The discovery of these anomalies suggest that the Monte Carlo physics in these low-energy regions may be somewhat unreliable and should be investigated further for any yet undiscovered issues when determining the level of residual contaminations in SABRE by Monte Carlo. However, for a general prediction of the expected background in SABRE, simulations with the current decay physics can be informative.

The geometry of the SABRE background simulation is shown in Figure 5.20. The seven 6.7-kg crystals¹⁰, shown in blue, are coupled to the two 3" PMTs, which are composed of a Kovar shell, quartz windows, and a ceramic feedthrough plate. The crystals and PMTs are enclosed in copper cans, shown in brown. The veto PMTs (green) are situated such that their bulbs are halfway into the body of the vessel. They are surrounded by 10" ID cans. Surrounding the vessel is 8" of lead shielding which could be turned on or off.

For most components of the detector, the primary radioactive isotope was generated uniformly throughout the volume of interest and was allowed to decay. Events that had an energy deposition in any of the crystals were recorded. A separate macro was written to reconstruct the detector response, taking into account the quenching factors from Chapter 7, the energy and spatial resolution of the detectors (the DAMA values were used), the light yield of the veto detector (discussed in Section 5.4.2), and the event timing information. This macro could generate spectra for individual backgrounds in single-hit and multi-hit events with and without the veto from the liquid-scintillator detector. The spectra were analyzed to

¹⁰Seven crystals, which can be arranged in a hex pattern were assumed. The 6.7-kg size is the current expected size for the crystals.




Figure 5.20: Geometry used in the simulations of the SABRE detector background. The crystals, shown in blue, are coupled to the two 3" PMTs, which are composed of a Kovar shell, quartz windows, and a ceramic feedthrough plate. The crystals and PMTs are enclosed in copper cans, shown in brown. The veto PMTs (green) are situated halfway in the body of the vessel. Surrounding the vessel is 8" of lead shielding.

determine the background rate in both a dark-matter mode (DMM) and a ^{40}K measurement mode (KMM).

In the dark-matter mode, crystal events occurring simultaneously with any veto signal (to within $1\ \mu\text{s}$) above a certain energy ($30\ \text{keV}_{ee}$) were rejected. In the ^{40}K measurement mode, events with an energy between 1200 and $1700\ \text{keV}_{ee}$ in the veto were kept as ^{40}K events. The latter range is chosen to produce a reasonable acceptance for true ^{40}K events while rejecting other backgrounds, simulating a potential cut that could be used with the actual data. The choice of the $1200\text{--}1700\ \text{keV}_{ee}$ window in the veto will reduce backgrounds that mimic ^{40}K by a factor of about 5. By contrast, the ^{40}K signal rate is cut by a factor of 3 total. Because non-potassium backgrounds in the crystal and veto should have a flat energy

spectrum, looking at events where 1–10 keV_{ee} are deposited in the crystal reduces the non-⁴⁰K background by a factor of roughly 100. Therefore, a relative reduction in background of about 130 can be expected by applying these cuts [115].

For many of the backgrounds simulated, the concentrations of impurities are known only with an upper limit. For these backgrounds, the simulation was performed assuming a concentration at the measured limit to form an “upper-limit” background spectrum. A second spectrum was generated ignoring these backgrounds to represent a “lower limit.” Spectra were also generated with and without the liquid-scintillator veto using a 1- μ s coincidence window. The upper-limit spectra with no veto is shown in blue in this document, the upper limit with the veto is shown in red. When the spectrum contains both backgrounds for which only limits are available and quantified backgrounds, the lower limit with no veto is shown in green, and the lower limit with the veto is shown in orange.

Throughout the following sections, the rates of various backgrounds are shown for the 2–4 keV_{ee} and 6–10 keV_{ee} regions. This is to estimate the backgrounds in the primary modulation region for DAMA/LIBRA and the location of the ⁴⁰K feature in the former case, and to estimate the backgrounds in the region of the essentially flat background of the DAMA/LIBRA energy spectrum where no modulation is observed in the latter. Determining the rate in these two regions can provide a basis of comparison between SABRE’s background-rejection potential compared with the backgrounds achieved by DAMA.

Backgrounds due to the crystal radioactivity

The crystal backgrounds were simulated using the measured values from the crystals grown by RMD, the Astrograde NaI powder, or the DAMA values, whichever is lowest. The isotopes measured were ⁴⁰K, ²³⁸U, ²³²Th, ⁸⁷Rb, ³H, ²²Na, ²⁴Na, ¹²⁹I, ²¹⁰Pb, and ⁸⁵Kr. The DAMA values and a description of how each value was measured is in Section 2.2.2, particularly in Table 2.3. The values for the backgrounds due to radiation in the SABRE crystals are shown in Table 5.3, while the spectra are shown in Figures 5.21–5.23. The crystal backgrounds

were found to be the dominant background in the predicted energy spectra at the current radioactivity levels, similar to the simulations of the DAMA/LIBRA experiment. The ^{40}K , ^{238}U , and ^{232}Th backgrounds are still significant at the current levels achieved in the crystals, but could be made lower through further purification along the lines of the efforts described in Section 4.2.

The ^{40}K background is shown in Figure 5.21. The 3-keV $_{ee}$ feature is due to the K-shell X-rays and Auger electrons. The L-shell feature can be seen below 1 keV $_{ee}$. Both features are drastically reduced by the veto. If the same level of ^{40}K is achieved as that of DAMA/LIBRA, the 3-keV $_{ee}$ feature is reduced to less than 0.05 cpd/kg/keV $_{ee}$, compared with the 0.5 cpd/kg/keV $_{ee}$ seen by DAMA and the 0.3 cpd/kg/keV $_{ee}$ expected in DAMA/LIBRA based on the simulations in Section 2.3.2. The rejection power of the veto can be seen in Figure 5.21, where the veto provides a factor of ~ 10 reduction in the ^{40}K background in addition to the multi-crystal rejection power. This large expected reduction in the ^{40}K background speaks to the power of the liquid-scintillator veto technique.

The ^{22}Na background, shown in Figure 5.23, has a K-shell Auger electron/X-ray at ~ 0.9 keV $_{ee}$, which appears as a peak in the figure. This cascade is coincident with high-energy γ rays from positron annihilation, making this feature effectively rejected by the liquid-scintillator veto, similar to ^{40}K . Without the peak, the flat background is low compared to other crystal backgrounds.

The ^{238}U and ^{232}Th backgrounds are shown in Figure 5.22. The background assumes secular equilibrium of both chains. The ^{238}U background is smoothly descending from 0–30 keV $_{ee}$, while the ^{232}Th background has more structure. The ^{232}Th background below 6 keV $_{ee}$ is very low, due to the treatment of the decay of ^{228}Ra , which is the dominant source of low-energy β particles in the ^{232}Th chain, as described in Section 2.3.2. The feature at ~ 9 keV $_{ee}$ is due to a number of Auger electrons and X-rays in the 8–10-keV $_{ee}$ range. The feature at 22 keV $_{ee}$ does not have an obvious physical origin and seems to result from the way that Geant4 generates conversion electrons during isomeric transitions. Neither of

these backgrounds are reduced very well by the liquid scintillator veto, so a reduction in the crystal ^{238}U and ^{232}Th concentrations is the best way to reduce these backgrounds. At the levels currently achieved in SABRE crystals, the ^{238}U and ^{232}Th backgrounds are no longer dominant.

With the reduction of the ^{40}K background through the liquid-scintillator veto and ^{238}U and ^{232}Th through crystal purification, the low-energy β -emitters ^{87}Rb and ^3H (see Figure 5.23) come to the forefront. The tritium concentration is not known in the SABRE crystals, since it cannot be measured with ICP-MS. However, because of the distinctive shape of the tritium background, SABRE may be able to measure this background by observing the shape of the energy spectrum. DAMA only reports a limit for the tritium background in their crystals, and do not comment on their method of generating this limit [41]. The ^{87}Rb background is shown at the currently-achieved levels in the SABRE crystals. SABRE is working to reduce the ^{87}Rb background through purification techniques.

Less prominent backgrounds include ^{85}Kr , the primordial ^{129}I , the neutron-activated ^{24}Na , and the radon daughter ^{210}Pb . At DAMA's reported limits, the ^{85}Kr , ^{129}I , and ^{24}Na backgrounds are small compared with the other crystal backgrounds. The ^{24}Na background could be further reduced if the liquid scintillator vessel is doped with a chemical with a high neutron-capture cross-section, or if passive shielding with a high cross section for neutron capture is used. The ^{210}Pb background here assumes a bulk concentration, but the event rate and spectrum may be different if there is a significant ^{210}Pb surface background. An effort should be made to eliminate surface backgrounds through polishing techniques.

Overall, the approximately constant rate from the crystals in SABRE is predicted to be about $0.1 \text{ cpd/kg/keV}_{ee}$ lower than the simulated DAMA rate, due to the lower contributions from ^{238}U and ^{232}Th and the significant reduction in the ^{40}K bump.




Figure 5.21: Expected background due to ^{40}K in the crystals, with and without veto, for 13 ppb ^{nat}K .




Figure 5.22: Expected background due to ^{238}U (Left) and ^{232}Th (Right) in the crystals, with and without veto. The ^{232}Th background has more structure, as described in the text.

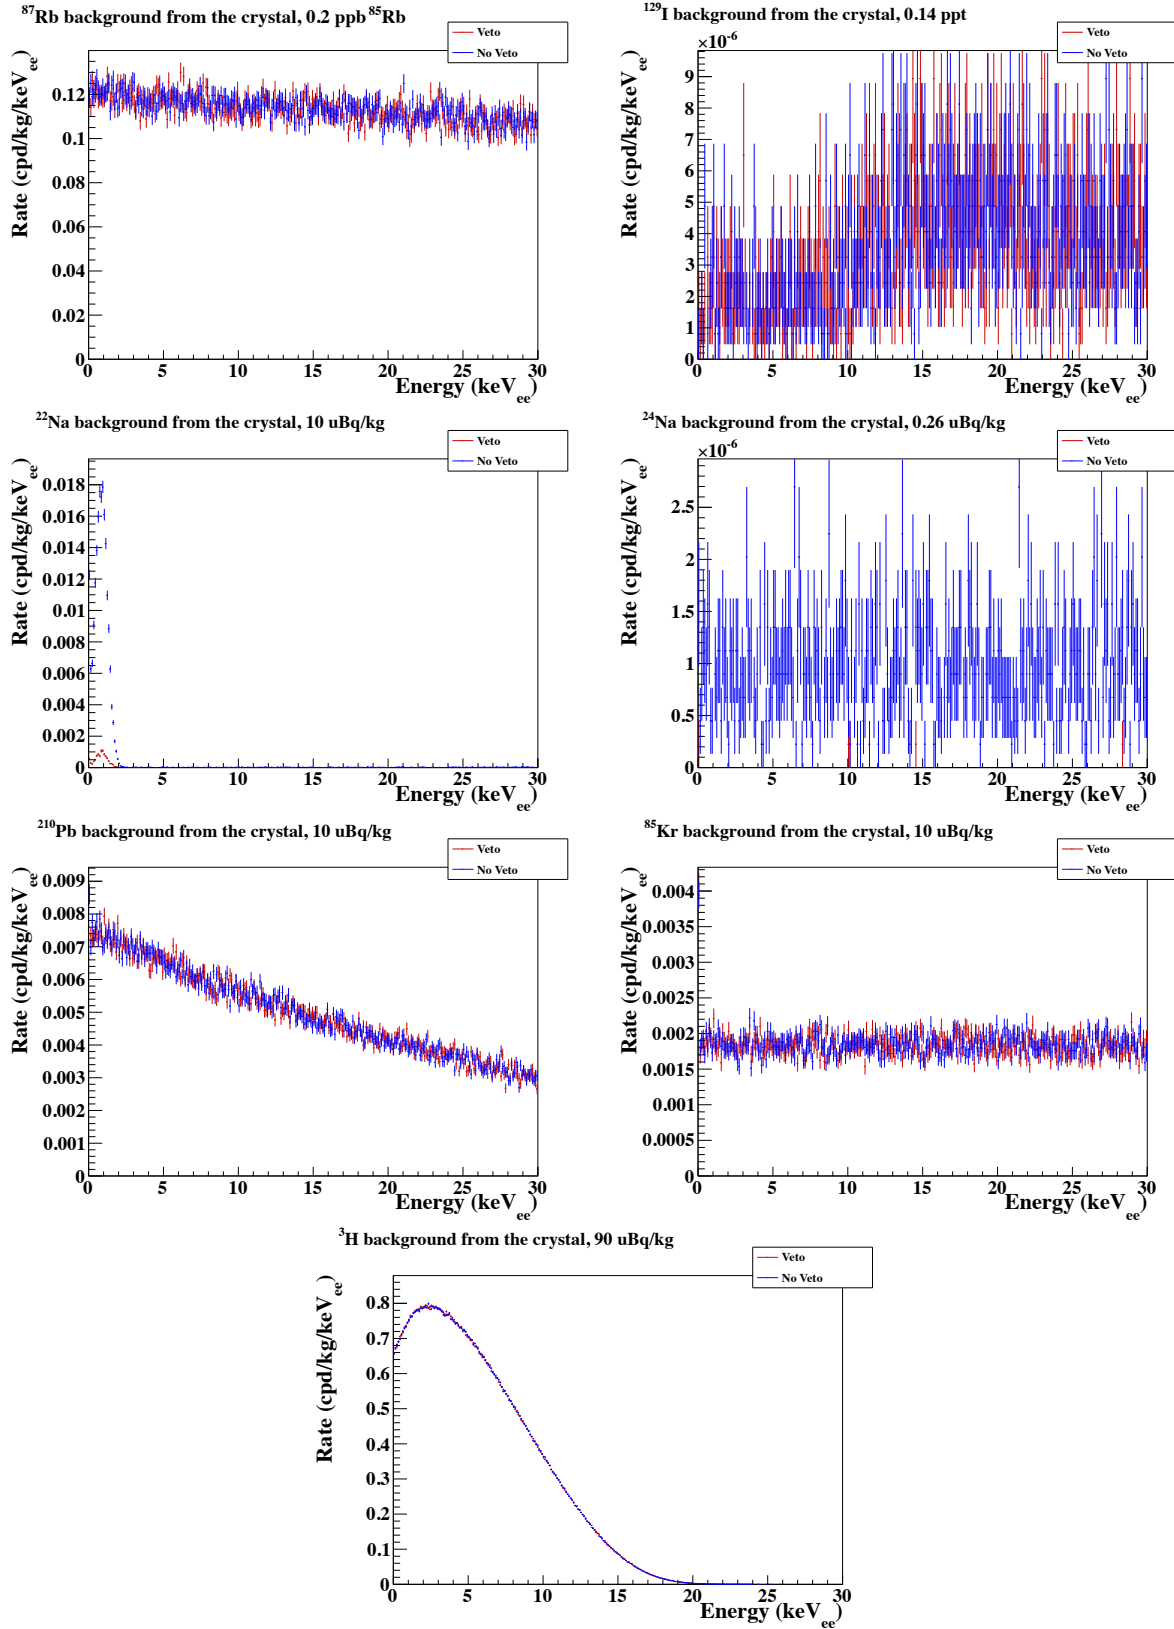


Figure 5.23: Background due to individual isotopes in the crystals other than ^{40}K , ^{238}U , and ^{232}Th , with the concentrations shown in Table 5.3. From top to bottom, the spectra are ^{87}Rb , ^{129}I , ^{22}Na , ^{24}Na , ^{210}Pb , ^{85}Kr , and ^3H . ^{87}Rb and ^3H are especially significant backgrounds at the powder value and DAMA limit, respectively.




Figure 5.24: Total expected background caused by the crystals for a dark-matter measurement, with and without the liquid-scintillator veto. The upper-limit spectra (red and blue with and without the veto, respectively) use measured upper limits for ^{22}Na , ^{24}Na , ^3H , and ^{85}Kr , while the lower-limit spectra (orange and green with and without the veto, respectively) leave out these backgrounds.




Figure 5.25: Total expected background caused by the crystals for a measurement of ^{40}K in the crystals. The upper-limit spectrum (red) uses measured upper limits for ^{22}Na , ^{24}Na , ^3H , and ^{85}Kr , while the lower-limit spectrum (orange) leaves out these backgrounds.

Table 5.3: Rates of events due to the presence of impurities in the crystal in the K-measurement mode (KMM) and dark-matter-measurement mode (DMM), with (v) and without (No veto) the liquid-scintillator veto. Rates are shown in cpd/kg/keV_{ee}.

Impurity	Concentration (Source)	Rate		
		KMM(v)	DMM(v)	No veto
2–4 keV _{ee}				
^{nat} K	13 ppb (DAMA)	1.41e-01	5.83e-02	3.43e-01
²³² Th	0.6 ppt (PNNL)	1.65e-04	1.27e-04	1.67e-04
²³⁸ U	0.5 ppt (PNNL)	1.28e-02	1.27e-02	1.29e-02
⁸⁷ Rb	0.2 ppb (powder)	1.25e-01	1.24e-01	1.24e-01
²¹⁰ Pb	10 μBq/kg (DAMA)	7.31e-03	7.24e-03	7.26e-03
⁸⁵ Kr	10 μBq/kg (DAMA limit)	1.96e-03	1.93e-03	1.92e-03
¹²⁹ I	0.14 ppt (DAMA)	1.95e-06	2.03e-06	2.11e-06
²² Na	10 μBq/kg (DAMA limit)	1.60e-05	1.30e-06	3.67e-05
²⁴ Na	0.26 μBq/kg (DAMA limit)	8.09e-07	<1.12e-08	8.31e-07
³ H	90 μBq/kg (DAMA limit)	8.17e-01	8.18e-01	8.17e-01
6–10 keV _{ee}				
^{nat} K	13 ppb (DAMA)	1.57e-02	1.54e-02	1.58e-02
²³² Th	0.6 ppt (PNNL)	5.01e-03	4.98e-03	5.02e-03
²³⁸ U	0.5 ppt (PNNL)	1.14e-02	1.13e-02	1.15e-02
⁸⁷ Rb	0.2 ppb (powder)	1.18e-01	1.19e-01	1.19e-01
²¹⁰ Pb	10 μBq/kg (DAMA)	6.08e-03	6.08e-03	6.06e-03
⁸⁵ Kr	10 μBq/kg (DAMA limit)	1.89e-03	1.88e-03	1.88e-03
¹²⁹ I	0.14 ppt (DAMA)	2.38e-06	2.36e-06	2.38e-06
²² Na	10 μBq/kg (DAMA limit)	7.99e-06	1.51e-06	1.32e-05
²⁴ Na	0.26 μBq/kg (DAMA limit)	8.54e-07	<5.62e-09	1.03e-06
³ H	90 μBq/kg (DAMA limit)	5.18e-01	5.19e-01	5.19e-01

PMT Backgrounds

The PMT background in SABRE is important to understand, given that the SABRE 3" PMTs will be directly coupled to the crystals, as opposed to DAMA/LIBRA, where the PMTs are separated by a 10-cm light-guide. The SABRE 3"-PMT background is therefore expected to be higher than that of DAMA/LIBRA. However, efforts to lower the intrinsic radioactivity of the PMTs, including the use of low-radioactivity ceramic feedthroughs, can help reduce the background rate added by this decision.

The PMT was split into three parts: the body of the tube, the window (which included backgrounds from the photocathode as well as the fused silica), and the feedthrough plate. The radioactivity values used were those that had been measured by [99] for the Xenon R11410-21 PMTs, as shown in Table 5.4. These PMTs are identical to the R11065-20 PMTs that will be used by SABRE, except for their photocathode material. The PMT window, the PMT body, and the feedthroughs are the dominant sources of backgrounds for the PMTs; the electrodes and other components of the PMT were not simulated.

In addition, the radioactivity of the new ceramic feedthroughs were also separately investigated. The radioactivity values from Table ?? were used, along with the total mass of the ceramic pieces from the current design shown in Figure 4.12.

The most notable feature in the simulation of the PMT backgrounds is the 19- and 16-keV_{ee} features from the upper- and lower-²³⁸U chains in the PMT window, respectively. These backgrounds occur in the simulation because of the mono-energetic γ rays released in isomeric transitions in Geant, and because the PMT window is directly coupled to the crystal (see Figure 5.26). The γ ray can escape into the crystal, whereas less energetic X-rays or Auger electrons that should carry the extra energy would not. Not only is this peak unphysical, but the ²³⁸U content in the PMT window is only reported as a limit by [99], so it is anticipated that a peak of this kind would simply not exist in an actual measurement.

The overall rate from the PMTs is expected to be on the order of the DAMA modulation after the veto is applied (see Figure 5.27), and could be smaller if the backgrounds for which

only limits are known are significantly below their limits. The new ceramic feedthrough stems can also be helpful in lowering backgrounds, as is shown in Figure 5.28.

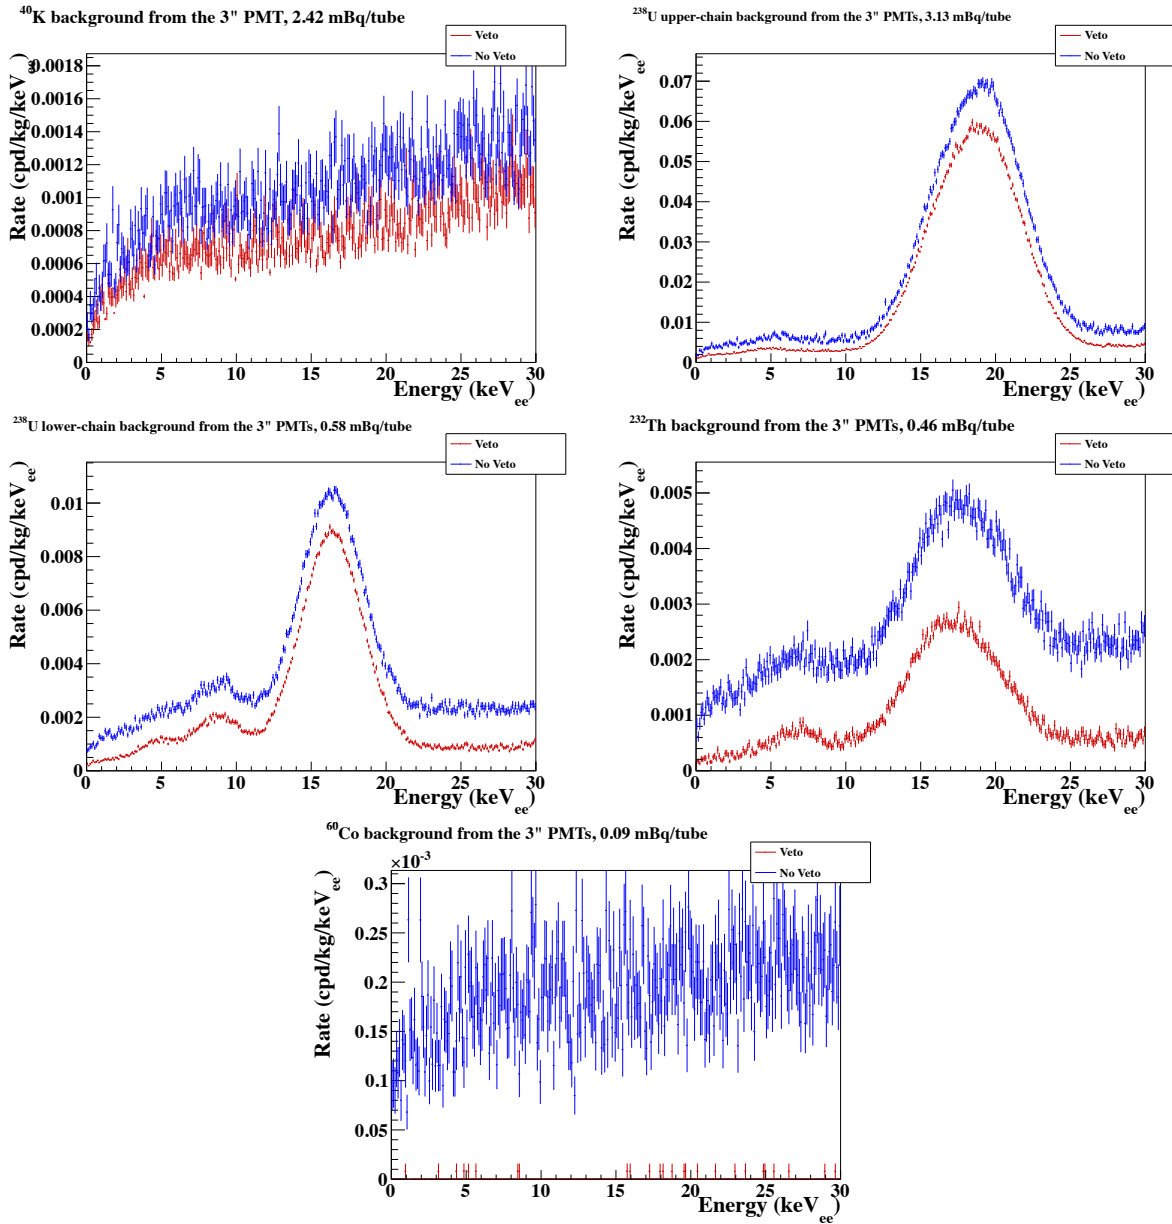


Figure 5.26: Expected backgrounds due to the 3" PMTs for individual isotopes, with and without the liquid-scintillator veto, with the old ceramic feedthrough plate. The backgrounds are from ⁴⁰K (Top left), the upper chain of ²³⁸U (Top right), the lower chain of ²³⁸U (Middle left), ²³²Th (Middle right), and ⁶⁰Co (Bottom). The liquid-scintillator veto highly suppressed the ⁶⁰Co background, so that only a few simulated events were missed by the veto detector.

Background from the 3" PMTs with and without veto




Figure 5.27: Expected backgrounds from the currently available R11065-20 3" PMTs for a dark-matter measurement, with and without the liquid-scintillator veto.




Figure 5.28: Expected backgrounds from the 3" PMTs for a dark-matter measurement with the new ceramic feedthrough stems, with and without the liquid-scintillator veto.

Background from the 3" PMTs in K measurement mode




Figure 5.29: Expected background from the current 3" PMTs for a measurement of ⁴⁰K in the crystals.

Table 5.4: Rates of background due to the 3" PMTs from 2–4 keV_{ee} and 6–10 keV_{ee} in a dark-matter measurement (DMM) and a ⁴⁰K measurement (KMM), with (v) and without (No veto) the liquid-scintillator veto. Radioactivity measurements are based on those made on R11410 PMTs by [99], which are identical to the R11065-20s that will be used by SABRE, except for the photocathode material. Radioactivity values for the new feedthrough plates came from Table ???. Rates are shown in cpd/kg/keV_{ee}.

2–4 keV _{ee}				
The PMT body				
Impurity	Activity (mBq/tube)	Rate		
		KMM(v)	DMM(v)	No veto
^{nat} K	<0.99	2.45e-04	1.40e-04	2.74e-04
²³⁸ U upper chain	<0.095	5.74e-05	2.19e-05	7.15e-05
²³⁸ U lower chain	<0.26	5.92e-04	1.52e-04	7.28e-04
²³² Th	<0.34	1.07e-03	1.75e-04	1.15e-03
⁶⁰ Co	0.07	8.61e-05	3.93e-07	1.24e-04
The PMT window				
Impurity	Activity (mBq/tube)	Rate		
		KMM(v)	DMM(v)	No veto
^{nat} K	<0.081	3.87e-04	3.75e-04	3.97e-04
²³⁸ U upper chain	<0.33	2.99e-03	2.58e-03	3.08e-03
²³⁸ U lower chain	0.036	7.35e-04	5.39e-04	7.75e-04
²³² Th	<0.011	2.16e-04	1.38e-04	2.21e-04
⁶⁰ Co	<0.0045	1.07e-05	6.94e-08	1.45e-05
The PMT feedthrough plate - old plates				
Impurity	Activity (mBq/tube)	Rate		
		KMM(v)	DMM(v)	No veto
^{nat} K	1.35	6.54e-05	2.20e-05	9.90e-05
²³⁸ U upper chain	<2.7	1.91e-03	1.26e-04	1.96e-03
²³⁸ U lower chain	0.28	2.19e-04	2.76e-05	2.34e-04
²³² Th	0.11	1.60e-04	1.54e-05	1.58e-04
⁶⁰ Co	<0.016	1.22e-05	4.98e-08	1.71e-05
The PMT feedthrough plate - new plates				
Impurity	Activity (mBq/tube)	Rate		
		KMM(v)	DMM(v)	No veto
^{nat} K	1.36	6.61e-05	2.23e-05	1.00e-04
²³⁸ U upper chain	0.023	1.59e-05	1.05e-06	1.64e-05
²³⁸ U lower chain	0.017	1.32e-05	1.66e-06	1.40e-05
²³² Th	0.042	6.09e-05	5.88e-06	6.01e-05
⁶⁰ Co	< 4.5 × 10 ⁻⁶	3.45e-09	1.41e-11	4.83e-09

Continued on next page...

Total for old phototubes ¹¹				
Impurity	Activity (mBq/tube)	Rate		
		KMM(v)	DMM(v)	No veto
<i>nat</i> K	<2.42	6.97e-04	5.37e-04	7.70e-04
²³⁸ U upper chain	<3.13	4.95e-03	2.72e-03	5.11e-03
²³⁸ U lower chain	<0.58	1.55e-03	7.19e-04	1.74e-03
²³² Th	<0.46	1.44e-03	3.28e-04	1.53e-03
⁶⁰ Co	<0.09	1.09e-04	5.12e-07	1.55e-04
Total for phototubes with new feedthrough plates ¹¹				
Impurity	Activity (mBq/tube)	Rate		
		KMM(v)	DMM(v)	No veto
<i>nat</i> K	<2.43	6.98e-04	5.38e-04	7.71e-04
²³⁸ U upper chain	<0.45	3.07e-03	2.60e-03	3.17e-03
²³⁸ U lower chain	<0.32	1.34e-03	6.93e-04	1.52e-03
²³² Th	<0.39	1.35e-03	3.19e-04	1.44e-03
⁶⁰ Co	<0.08	9.68e-05	4.63e-07	1.38e-04

6–10 keV _{ee}				
The PMT body				
Impurity	Activity (mBq/tube)	Rate		
		KMM(v)	DMM(v)	No veto
<i>nat</i> K	<0.99	2.44e-04	1.70e-04	3.15e-04
²³⁸ U upper chain	<0.095	8.33e-05	3.83e-05	9.54e-05
²³⁸ U lower chain	<0.26	9.49e-04	3.93e-04	1.12e-03
²³² Th	<0.34	1.38e-03	3.22e-04	1.46e-03
⁶⁰ Co	0.07	1.14e-04	3.93e-07	1.54e-04
The PMT window				
Impurity	Activity (mBq/tube)	Rate		
		KMM(v)	DMM(v)	No veto
<i>nat</i> K	<0.081	4.95e-04	4.83e-04	5.05e-04
²³⁸ U upper chain	<0.33	3.36e-03	2.84e-03	3.44e-03
²³⁸ U lower chain	0.036	1.50e-03	1.27e-03	1.55e-03
²³² Th	<0.011	3.63e-04	2.73e-04	3.70e-04
⁶⁰ Co	<0.0045	1.14e-05	7.22e-08	1.57e-05

Continued on next page...

¹¹Simulations were not conducted on the electrodes and other smaller parts, so the total activity quoted is larger than the sum of the previously mentioned activities.

The PMT feedthrough plate - old plates				
Impurity	Activity (mBq/tube)	Rate		
		KMM(v)	DMM(v)	No veto
<i>nat</i> K	1.35	1.00e-04	4.97e-05	1.43e-04
²³⁸ U upper chain	<2.7	2.16e-03	1.47e-04	2.43e-03
²³⁸ U lower chain	0.28	2.85e-04	6.47e-05	2.99e-04
²³² Th	0.11	1.98e-04	3.64e-05	2.18e-04
⁶⁰ Co	<0.016	1.47e-05	5.19e-08	2.04e-05
The PMT feedthrough plate - new plates				
Impurity	Activity (mBq/tube)	Rate		
		KMM(v)	DMM(v)	No veto
<i>nat</i> K	1.36	1.01e-04	5.02e-05	1.44e-04
²³⁸ U upper chain	0.023	1.81e-05	1.23e-06	2.03e-05
²³⁸ U lower chain	0.017	1.71e-05	3.88e-06	1.79e-05
²³² Th	0.042	7.54e-05	1.39e-05	8.33e-05
⁶⁰ Co	$< 4.5 \times 10^{-6}$	4.14e-09	1.46e-11	5.75e-09
Total for old phototubes ¹¹				
Impurity	Activity (mBq/tube)	Rate		
		KMM(v)	DMM(v)	No veto
<i>nat</i> K	<2.42	8.39e-04	7.03e-04	9.63e-04
²³⁸ U upper chain	<3.13	5.61e-03	3.03e-03	5.96e-03
²³⁸ U lower chain	<0.58	2.74e-03	1.73e-03	2.97e-03
²³² Th	<0.46	1.94e-03	6.31e-04	2.05e-03
⁶⁰ Co	<0.09	1.40e-04	5.17e-07	1.90e-04
Total for phototubes with new feedthrough plates ¹¹				
Impurity	Activity (mBq/tube)	Rate		
		KMM(v)	DMM(v)	No veto
<i>nat</i> K	<2.43	8.40e-04	7.03e-04	9.65e-04
²³⁸ U upper chain	<0.45	3.46e-03	2.88e-03	3.55e-03
²³⁸ U lower chain	<0.32	2.47e-03	1.67e-03	2.69e-03
²³² Th	<0.39	1.82e-03	6.08e-04	1.91e-03
⁶⁰ Co	<0.08	1.25e-04	4.65e-07	1.70e-04

Enclosure Backgrounds

Backgrounds for the enclosure were simulated using the limits achieved for different kinds of copper. For ^{40}K , the concentration used was that of the OFHC copper used in the first SABRE enclosure built for the SABRE-in-DarkSide tests as measured by GDMS. The ^{238}U and ^{232}Th values were taken from electroformed copper as described in [102].

In addition to the limits for long-lived isotopes, cosmogenically-activated backgrounds were simulated for ^{56}Co , ^{57}Co , ^{58}Co , ^{60}Co , ^{46}Sc , ^{54}Mn , and ^{59}Fe , as shown in Figure 5.31. The saturated concentrations at sea level from [116] were used, averaging the values from the Ziegler [117] and Gordon et al. [118] cosmic neutron spectra. The enclosure backgrounds are low at these limits compared to the crystal and 3"-PMT backgrounds, about an order of magnitude lower than the 3"-PMTs for a dark-matter measurement with the liquid-scintillator veto and a factor of 2 lower for a ^{40}K measurement. As can be seen in Figure 5.32, the veto is very effective at reducing the copper backgrounds.

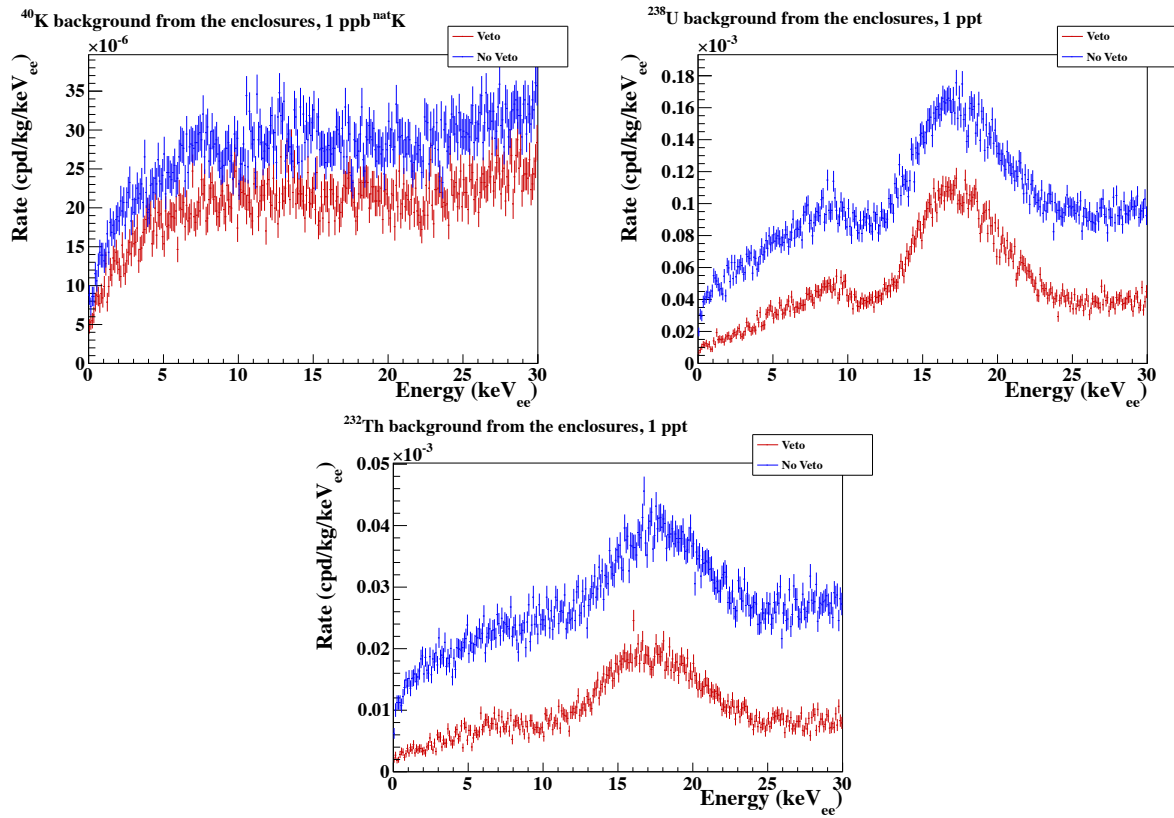


Figure 5.30: Expected backgrounds due to the enclosures from ^{40}K (Top left), ^{238}U (Top right), and ^{232}Th (Bottom), with and without the veto.

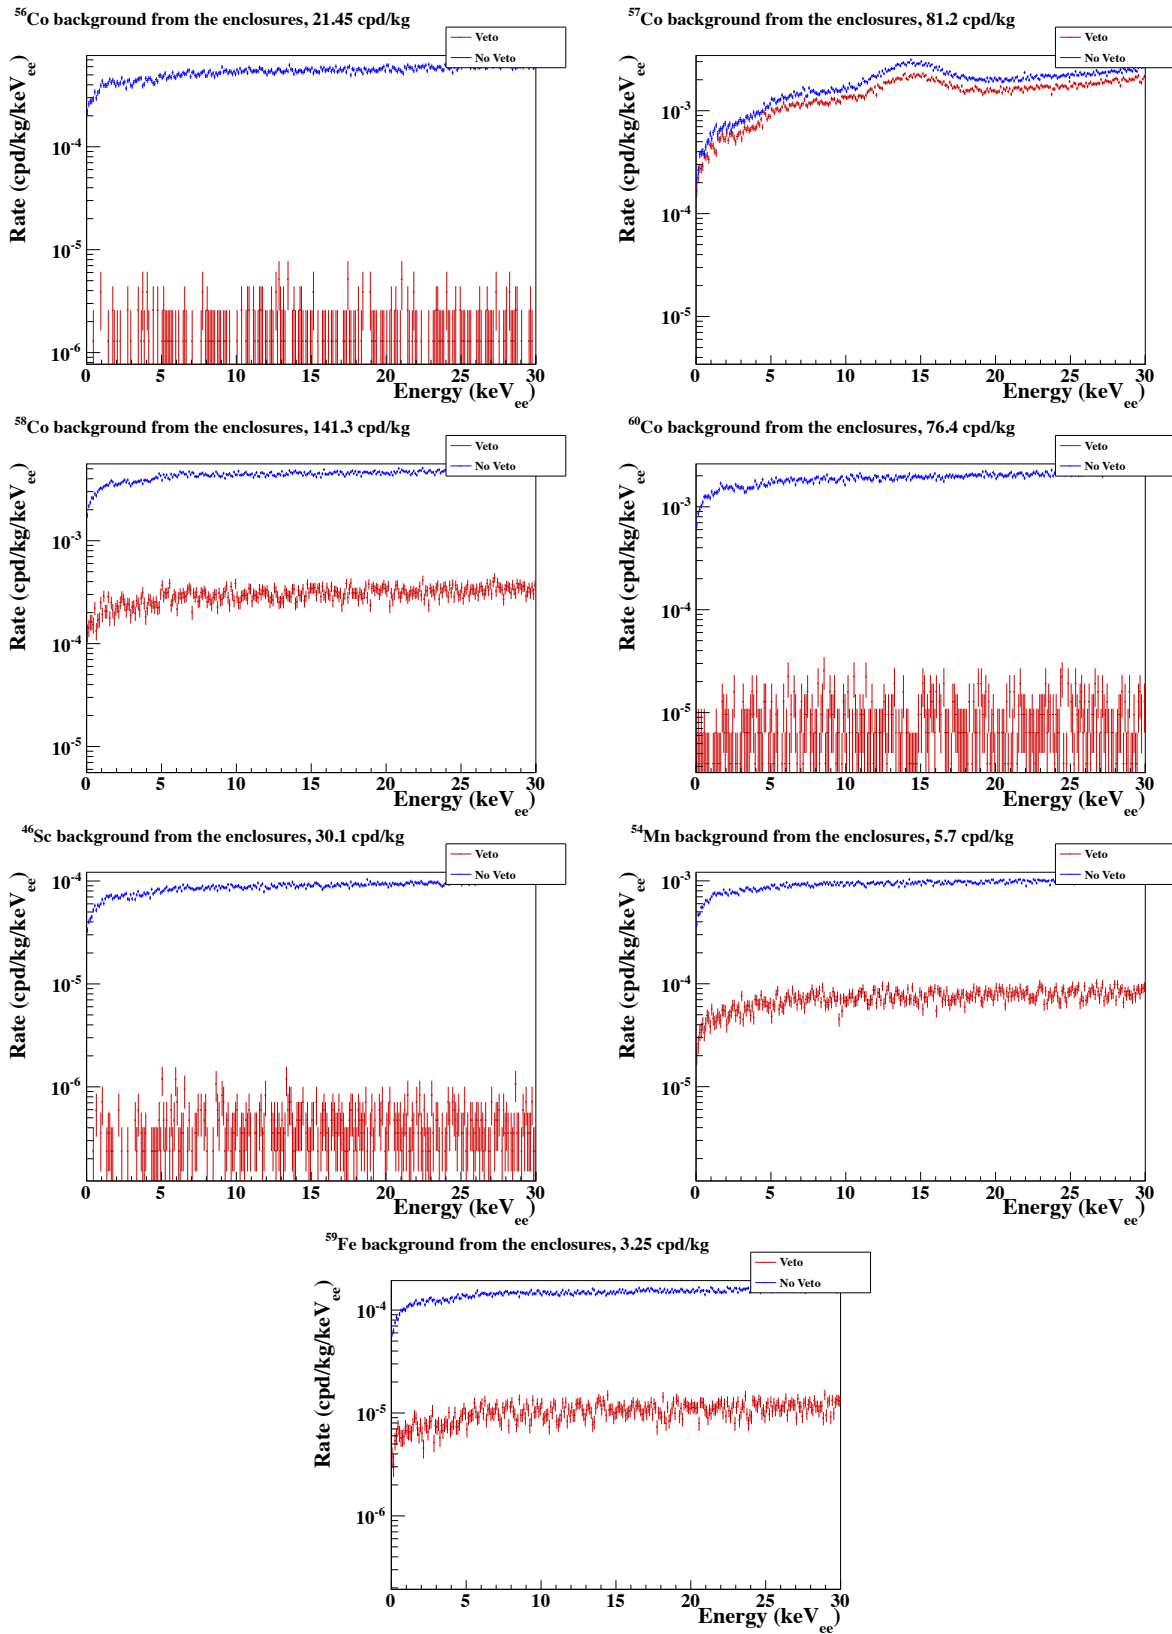


Figure 5.31: Expected backgrounds due to the enclosures from cosmogenically-activated isotopes. Reading from top-left to bottom-right, the isotopes simulated are ^{56}Co , ^{57}Co , ^{58}Co , ^{60}Co , ^{46}Sc , ^{54}Mn , and ^{59}Fe , with and without the veto. Radioactivity values were the saturated values from [116].

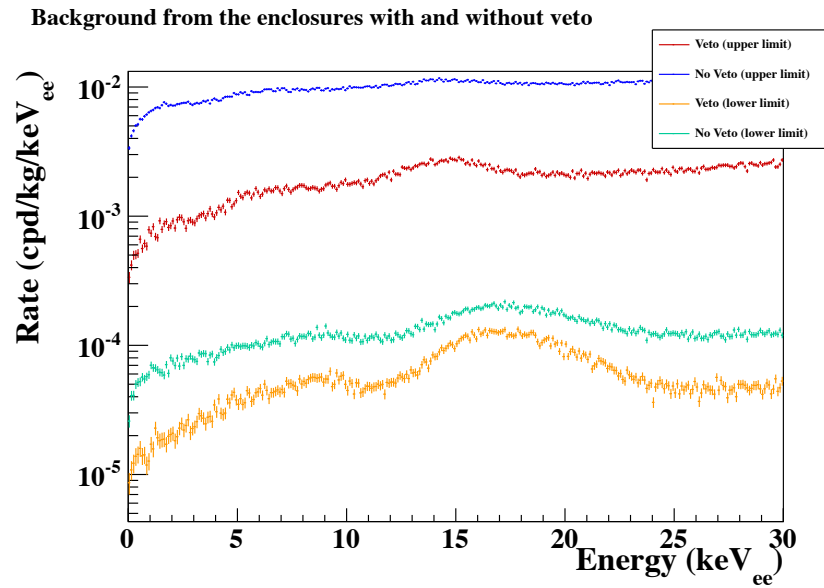


Figure 5.32: Expected backgrounds from the enclosures in a dark-matter measurement, with and without the liquid-scintillator veto.




Figure 5.33: Expected backgrounds from the enclosures for a measurement of ^{40}K in the crystals.

Table 5.5: Rates of background due to the enclosure for a dark-matter measurement (DMM) and a K measurement (KMM) with (v) and without (No veto) the liquid-scintillator veto. The limits for OFHC and electroformed copper are used. Rates are shown in cpd/kg/keV_{ee}.

Impurity	Concentration	Rate		
		KMM(v)	DMM(v)	No veto
2–4 keV _{ee}				
^{nat} K	1 ppb (OFHC limit)	1.92e-05	1.59e-05	2.20e-05
²³² Th	1 ppt (electroform)	1.80e-05	4.96e-06	1.90e-05
²³⁸ U	1 ppt (electroform)	5.47e-05	2.15e-05	6.41e-05
⁵⁶ Co	21.45 cpd/kg Cu, from [116]	3.96e-04	9.69e-07	4.48e-04
⁵⁷ Co	81.2 cpd/kg Cu, from [116]	8.34e-04	6.45e-04	8.40e-04
⁵⁸ Co	141.3 cpd/kg Cu, from [116]	3.59e-03	2.49e-04	3.89e-03
⁶⁰ Co	76.4 cpd/kg Cu, from [116]	1.27e-05	2.66e-08	1.06e-05
⁴⁶ Sc	30.1 cpd/kg Cu, from [116]	5.25e-05	2.20e-07	7.58e-05
⁵⁴ Mn	5.7 cpd/kg Cu, from [116]	8.35e-04	5.90e-05	8.37e-04
⁵⁹ Fe	3.25 cpd/kg Cu, from [116]	1.09e-04	7.79e-06	1.30e-04
6–10 keV _{ee}				
^{nat} K	1 ppb (OFHC limit)	2.50e-05	2.12e-05	2.82e-05
²³² Th	1 ppt (electroform)	2.27e-05	7.70e-06	2.41e-05
²³⁸ U	1 ppt (electroform)	8.54e-05	4.41e-05	9.42e-05
⁵⁶ Co	21.45 cpd/kg Cu, from [116]	4.80e-04	8.40e-07	5.33e-04
⁵⁷ Co	81.2 cpd/kg Cu, from [116]	1.56e-03	1.23e-03	1.55e-03
⁵⁸ Co	141.3 cpd/kg Cu, from [116]	4.14e-03	2.99e-04	4.50e-03
⁶⁰ Co	76.4 cpd/kg Cu, from [116]	1.53e-05	9.31e-08	1.27e-05
⁴⁶ Sc	30.1 cpd/kg Cu, from [116]	6.14e-05	3.77e-07	8.84e-05
⁵⁴ Mn	5.7 cpd/kg Cu, from [116]	9.48e-04	7.39e-05	9.48e-04
⁵⁹ Fe	3.25 cpd/kg Cu, from [116]	1.27e-04	1.07e-05	1.50e-04

Veto PMT Backgrounds

The Hamamatsu R5912 veto PMT activities have been previously measured for the DarkSide experiment [119]. These activities were used in the simulation. The veto PMT backgrounds are very efficiently rejected by the liquid-scintillator veto. Despite the low statistics achieved, it is apparent that this background is small when compared with the crystal, 3" PMT, and external backgrounds.

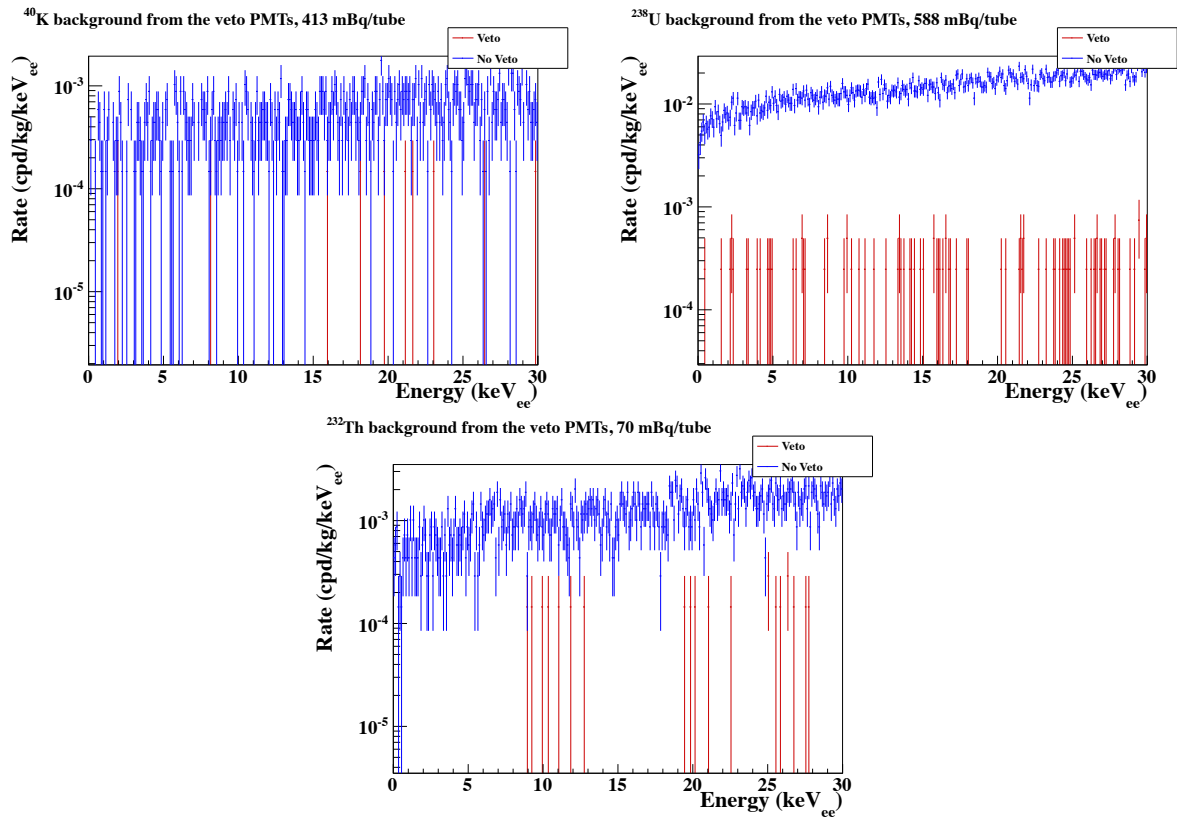


Figure 5.34: Expected backgrounds due to the veto PMTs due to ⁴⁰K (Top left), ²³⁸U (Top right), and ²³²Th (Bottom), with and without the veto.




Figure 5.35: Expected backgrounds from the veto PMTs for a dark-matter measurement, with and without the liquid-scintillator veto.

Table 5.6: Rates of background due to the veto PMTs in a dark-matter-measurement (DMM) and a K-measurement mode (KMM), with (v) and without (No veto) the liquid-scintillator veto. Radioactivity measurements by DarkSide. Rates are shown in cpd/kg/keV_{ee}.

Impurity	Activity (mBq/tube)	Rate		
		KMM(v)	DMM(v)	No veto
2–4 keV _{ee}				
<i>nat</i> K	413	1.47e-04	<1.74e-05	3.61e-04
²³² Th	70	5.79e-04	<7.24e-06	6.59e-04
²³⁸ U	588	6.50e-03	8.67e-05	9.16e-03
6–10 keV _{ee}				
<i>nat</i> K	413	1.21e-04	3.68e-06	4.42e-04
²³² Th	70	9.34e-04	1.09e-05	1.10e-03
²³⁸ U	588	8.40e-03	7.43e-05	1.22e-02

Background from the veto PMTs in K measurement mode




Figure 5.36: Expected backgrounds from the veto PMTs for a measurement of ^{40}K in the crystals.

Steel Backgrounds

Radioactive sources were placed uniformly throughout the steel. The activities used were the weighted average of the activities measured by GDMS in Table 5.1. GDMS is not a very accurate measurement tool for radioactivity, so there is some uncertainty to these values. A sample of the veto-vessel steel should be measured by γ -counting for a more accurate measurement of the radioactivity levels. Using the GDMS numbers, the background from the steel, like the veto PMTs, is heavily rejected by the liquid-scintillator veto, and is similar in rate to the veto-PMT background.

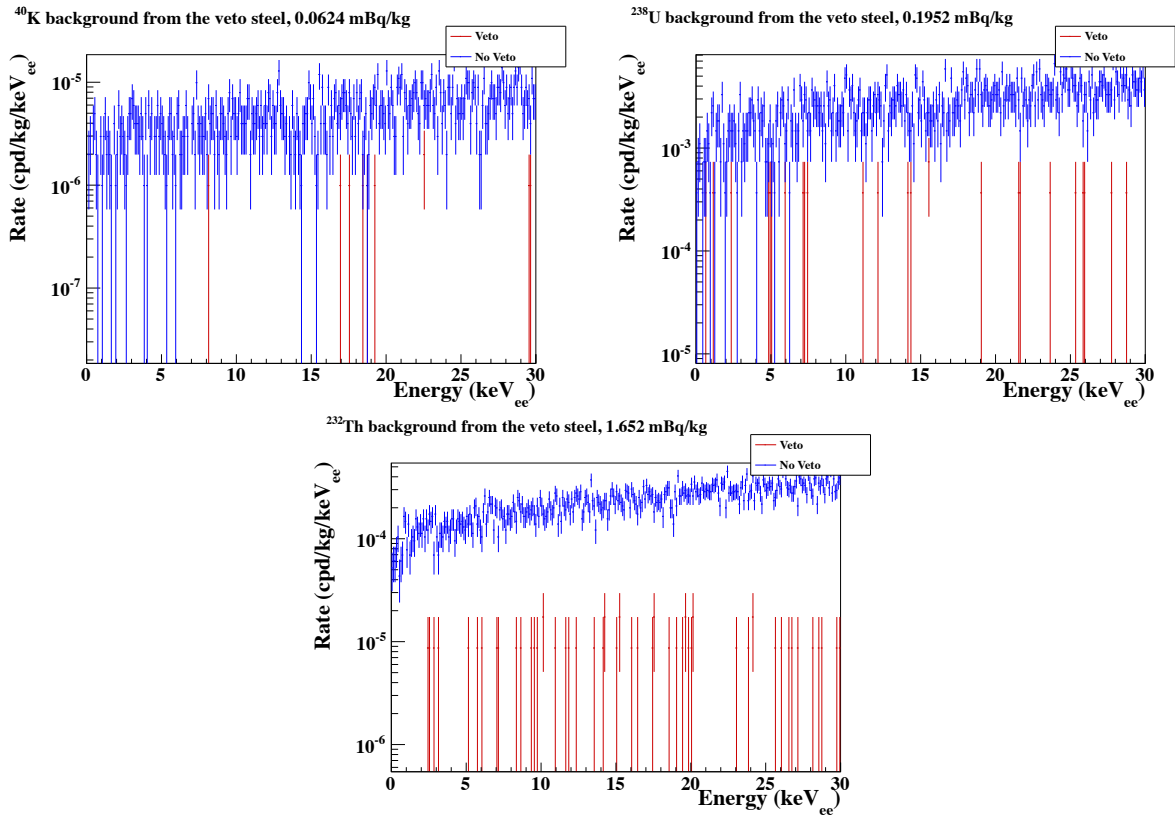


Figure 5.37: Expected background due to the liquid-scintillator-vessel steel from ^{40}K (Top left), ^{238}U (Top right), and ^{232}Th (Bottom), with and without the veto.




Figure 5.38: Expected backgrounds from the veto-detector steel in a dark-matter measurement, with and without the liquid-scintillator veto.

Table 5.7: Rates of background due to the veto-detector steel in a dark-matter measurement (DMM) and K-measurement mode (KMM) with (v) and without (No veto) the liquid-scintillator veto. The radioactivity of the steel used for the vessel was measured by GDMS (see Table 5.1). Rates are shown in cpd/kg/keV_{ee}.

Impurity	Activity (mBq/kg)	Rate		
		KMM(v)	DMM(v)	No veto
2-4 keV _{ee}				
^{nat} K	0.0624	1.68e-06	<1.71e-07	3.81e-06
²³² Th	0.1952	1.25e-04	1.74e-06	1.35e-04
²³⁸ U	1.652	1.23e-03	1.84e-05	1.62e-03
6-10 keV _{ee}				
^{nat} K	0.0624	1.68e-06	2.48e-08	4.06e-06
²³² Th	0.1952	1.77e-04	1.74e-06	1.89e-04
²³⁸ U	1.652	1.93e-03	2.75e-05	2.50e-03

Background from the veto steel in K measurement mode




Figure 5.39: Expected backgrounds from the veto-detector steel for a measurement of ^{40}K in the crystals.

Table 5.8: LNGS rock and concrete composition, according to [120].

Element	Rock	Concrete
H	–	0.89
C	11.88	7.99
O	47.91	048.4
Na	–	0.60
Mg	5.58	0.85
Al	1.03	0.90
Si	1.27	3.86
P	–	0.04
S	–	0.16
K	1.03	0.54
Ca	30.29	34.06
Ti	–	0.04
Fe	–	0.43

External Backgrounds

For external γ radiation from the laboratory rock and concrete, a simulation was performed with 3 m of LNGS rock (see Table 5.8) in which ^{40}K , ^{232}Th , and ^{238}U were allowed to decay. The outgoing γ radiation leaving the rock was recorded, and a spectrum was created with the energy and angular distributions of the γ rays. This spectrum was scaled to the rock radioactivity at LNGS in Hall B shown in Figure 5.40, as measured by [42], as shown in Table 5.9. The total γ flux from the LNGS rock in the simulation was $0.27 \text{ } \gamma/\text{s}/\text{cm}^2$ between 0 and 5 MeV_{ee} , which is similar to the flux measured in Hall B in [42] of $0.33 \text{ } \gamma/\text{s}/\text{cm}^2$. The angular spectrum was found to fit well with an analytical distribution where the probability that a γ was emitted with an angle θ from the normal is proportional to $\sin(\theta/2)$.

The energy spectrum from the rock simulation as well as the analytical expression for the angular distribution were used as input for the primary particles generated in the external-background simulation. The external γ background was simulated without the use of any passive shielding outside of the veto, and could be later scaled with a given shielding power. This simplification distorts the energy spectrum, but was nonetheless used in order to gather adequate statistics.




Figure 5.40: γ backgrounds at LNGS, as measured with a NaI(Tl) detector in [42]. The radioactivity values for Hall B were used in the simulations for the SABRE detector.

Table 5.9: LNGS rock and concrete radioactive impurity concentrations.

Location	^{238}U (Bq/kg)	^{232}Th (Bq/kg)	^{40}K (Bq/kg)
Hall A [42]	11.7 ± 3.9	14.8 ± 2.8	62 ± 14
Hall B [42]	19.6 ± 4.9	13.2 ± 2.7	52 ± 10
Concrete (Floor) [42]	26 ± 5	8 ± 2	170 ± 27
Concrete (Wall) [42]	15 ± 2	3.8 ± 0.8	42 ± 6
Rock Hall A [120]	84.0	8.8	–
Rock Hall B [120]	5.2	0.3	–
Rock Hall C [120]	8.1	0.3	–
Concrete [120]	13.0	2.7	–

The external γ background is very high, and will need to be reduced through shielding. In the 2–4 keV $_{ee}$ range, the background without shielding is about 50 cpd/kg/keV $_{ee}$, while the crystal background is around 1 cpd/kg/keV $_{ee}$ in this region. Therefore, a reduction through shielding of ~ 500 is needed to achieve a background from external γ rays on par with the lower-limit background from the crystals themselves for a measurement of ^{40}K in the crystals. The dark-matter-measurement shielding requirements are less stringent because the liquid-




Figure 5.41: Expected external γ background from the LNGS rock, with no shielding.




Figure 5.42: Expected external γ background from the LNGS rock in K measurement mode, with no shielding.

scintillator veto rejects a large portion of the external γ background. A conservative shielding factor of 10^4 brings the external γ background down to the level of the 3" PMTs and the DAMA/LIBRA modulation amplitude.

However, another issue relevant to dark-matter searches is the dead-time imposed on the experiment whenever the veto detector is triggered. A simulation was performed to calculate

Table 5.10: Rates of background from 2–4 and 6–10 keV_{ee} due to external γ background with no shielding in a dark-matter measurement (DMM) and K-measurement mode (KMM) with (v) and without (No veto) the liquid-scintillator veto. Radioactivity measurements are based on [42] for Hall B. Rates are shown in cpd/kg/keV_{ee}.

Impurity	Rate		
	KMM(v)	DMM(v)	No veto
2–4 keV _{ee}			
No shielding	5.80e+01	1.83e+00	6.93e+01
6–10 keV _{ee}			
No shielding	7.42e+01	2.50e+00	8.89e+01

the rate of the liquid scintillator veto with 8" (~ 20 cm) of lead shielding. The rate due to the veto PMTs, the veto-detector steel, and the external γ radiation from the LNGS rock were considered. According to the simulation, a shielding factor of 10^4 would not only make the external γ background much lower than the crystal background, but the rate in the liquid-scintillator veto would be a fairly reasonable 25 Hz.

Total expected SABRE background

The total expected backgrounds for the SABRE detector are shown in Figure 5.43, and broken down by source in Figures 5.44, and 5.45 for a dark-matter measurement. The total simulated backgrounds are shown for a ^{40}K measurement in Figure 5.46. The total backgrounds are also summarized in Table 5.11.

The liquid-scintillator veto is a strong mechanism for lowering the backgrounds of the experiment. In particular, it reduces the ^{40}K 3-keV_{ee} signature feature by nearly an order of magnitude. It also provides powerful rejection for external γ radiation and enclosure backgrounds, particularly cosmogenically-activated radioactive isotopes.

There are some uncertainties to the background spectra measured in this simulation. First, the geometries of the enclosures are not known in detail; a more detailed geometry will be needed as the enclosure design is finalized. Some materials were not simulated that are very close to the crystal, like the reflector and the optical coupling gel. Measurements

of some of the material are not as precise as they could be. For example, the material to be used for the copper is not yet known, and the steel used for the vessel has not been measured by γ counting, but has only been measured by GDMS, a less reliable counting method. The achievable purity levels in the crystals are not yet known; these may be the dominant background in the experiment. Finally, issues with the simulation software itself, such as the way Geant4 handles isomeric transitions, may cause distortions in the simulated energy spectra.

While the achievable impurity levels in the SABRE crystals are not yet known, if SABRE can reproduce the purity of the DAMA crystals where they haven't already been surpassed, the total background will be much lower than the total rate in the DAMA/LIBRA experiment ($\sim 1\text{--}1.5$ cpd/kg/keV $_{ee}$). If dominant crystal backgrounds can be further reduced, a background on the order of the DAMA/LIBRA modulation is achievable.




Figure 5.43: Total expected SABRE background for a dark-matter measurement with a 10^4 shielding factor. The total background without the veto is shown in blue (upper limit) and green (lower limit), while the background with the veto is shown in red (upper limit) and orange (lower limit). This plot assumes the radioactivity values for the currently available 3" PMTs. With this amount of shielding, the crystal background is the dominant background in the experiment.

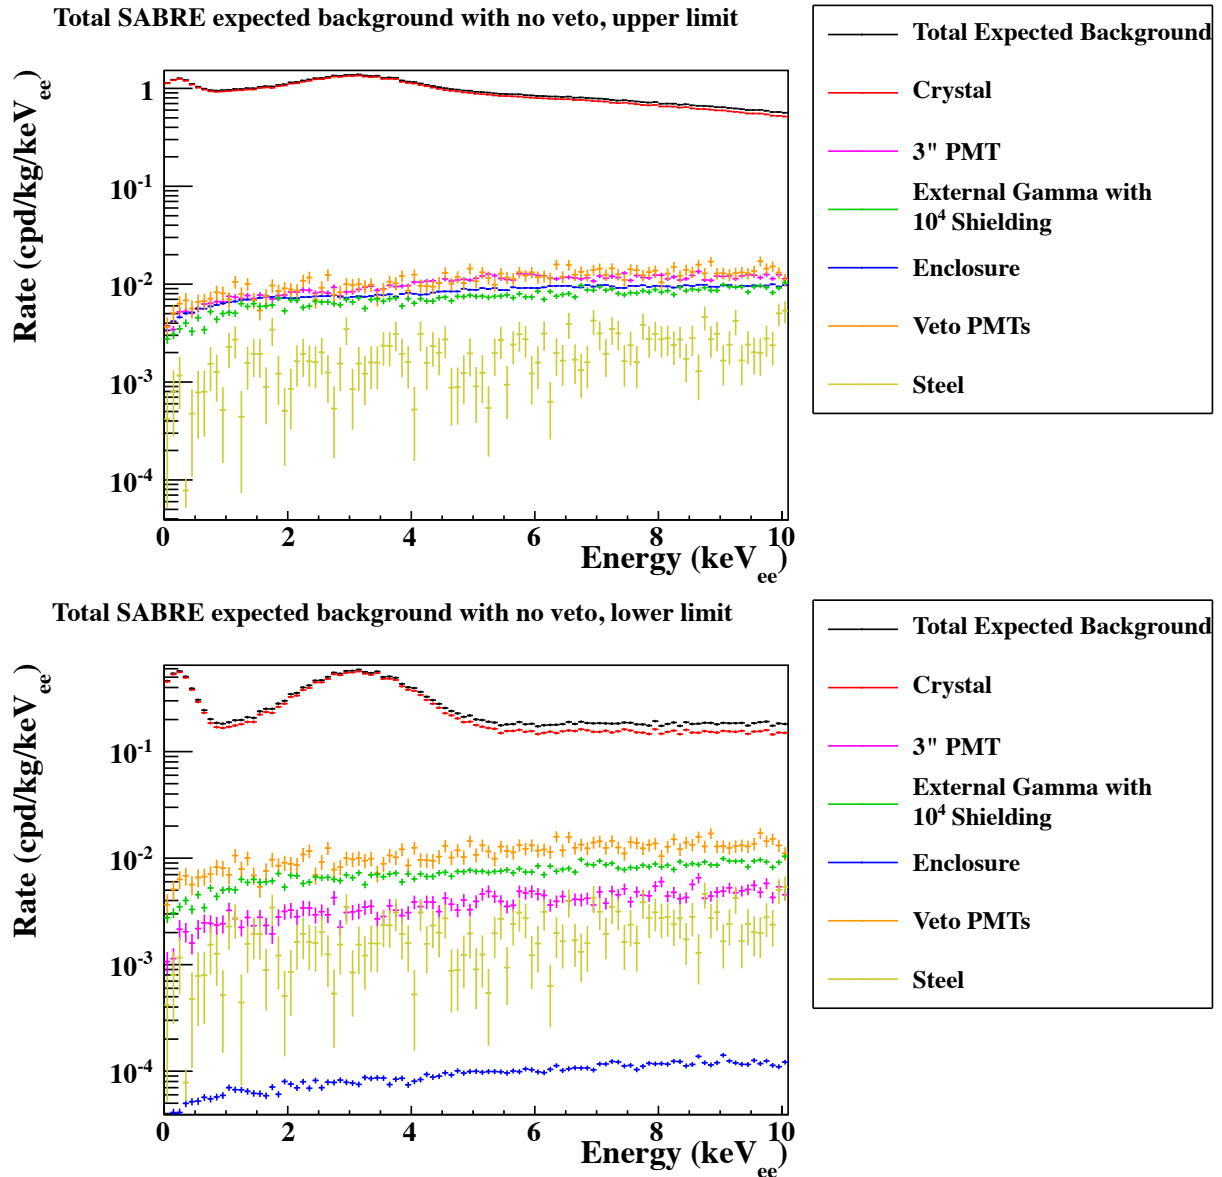


Figure 5.44: Total expected SABRE backgrounds for a dark-matter measurement without the veto from 0–10 keV_{ee}, broken down by source. The upper limit is shown on top, while the lower limit is on bottom. The highest contributor to the background is the crystal, followed by the veto PMTs, the 3" PMTs (with the old ceramic feedthrough plates), and the external γ background. The copper enclosures have a wide variation in how much they contribute, depending on the concentration of cosmogenically-activated radioactive isotopes. A shielding power for the external background of 10^4 is assumed.

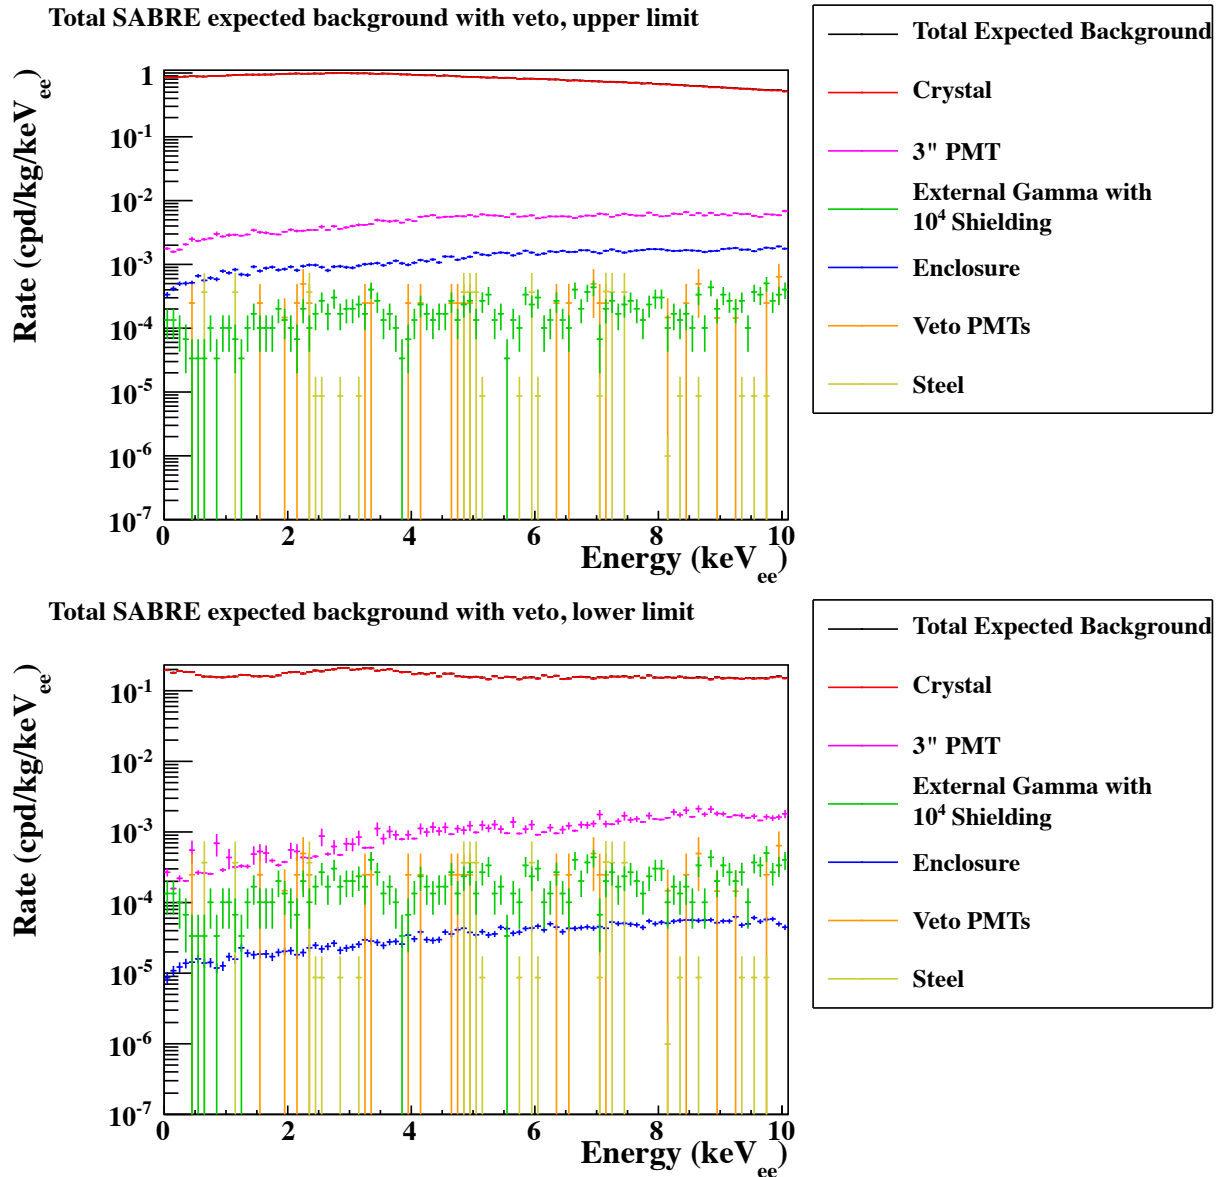


Figure 5.45: Total expected SABRE backgrounds for a dark-matter measurement with the veto from 0–10 keV_{ee}, broken down by source. The upper limit is shown on top, while the lower limit is on bottom. The highest contributor to the background is the crystal. The copper enclosures may be the next-highest contributor, depending on the rate of cosmogenic activation. The next largest contribution comes from the 3" PMTs (with the old feedthrough plates). With the veto active, it was difficult to get enough statistics to infer the rate from the veto PMTs, but it is unlikely to be higher than the 3" PMTs. A shielding power for the external background of 10⁴ is assumed.

Total SABRE expected background in K measurement mode, with 10^4 shielding power




Figure 5.46: Total expected SABRE background for a K measurement with 10^4 shielding factor. The crystal and external γ backgrounds are dominant. The upper limit is shown in red, while the lower limit is shown in orange. The ^{40}K signal is shown in green for reference. A factor of 10^4 reduction in shielding is sufficient to make the total background comparable to the ^{40}K signal.

Table 5.11: Total expected backgrounds for the SABRE detector in the 2–4 and 6–10 keV_{ee} windows for a dark-matter measurement (DMM) and a K measurement (KMM) with (v) and without (No veto) the liquid-scintillator veto.

Detector volume	KMM(v)	DMM(v)	No veto
2–4 keV_{ee}			
Crystal	1.11e+00	1.02e+00	1.31e+00
3" PMTs	1.75e-02	8.62e-03	1.86e-02
Enclosure	5.92e-03	1.00e-03	6.34e-03
Veto PMTs	7.23e-03	8.67e-05	1.02e-02
Veto steel	1.36e-03	2.01e-05	1.75e-03
External (no shielding)	5.80e+01	1.83e+00	6.93e+01
6–10 keV_{ee}			
Crystal	6.77e-01	6.77e-01	6.78e-01
3" PMTs	2.25e-02	1.22e-02	2.43e-02
Enclosure	7.46e-03	1.69e-03	7.93e-03
Veto PMTs	9.45e-03	8.88e-05	1.37e-02
Veto steel	2.11e-03	2.93e-05	2.69e-03
External (no shielding)	7.42e+01	2.50e+00	8.89e+01

5.4.2 Light yield of the SABRE veto detector

The light-collection efficiency of the SABRE veto detector determines the detection efficiency of low-energy depositions. While certain higher-energy backgrounds like the 1461-keV γ ray from the decay of ^{40}K are easily detected, lower-energy signals like the neutron-thermalization signal may be missed by a detector with a low light yield. The light yield of the SABRE veto detector was predicted with a custom-written optical Monte Carlo that had previously been used to successfully predict the light yield of the DarkSide-50 neutron veto detector¹².

This optical Monte Carlo tracks photons after they are emitted in point-like events as they are reflected or absorbed by reflectors, absorbed and reemitted by the scintillator cocktail used in DarkSide-50, and detected or not by PMTs. The wavelength of the light is tracked, and many interactions are wavelength-dependent. The deposited-energy spectrum and the location of events can be used as input parameters. Relevant information on the time and location of detected photons, as well as some statistics on the light-collection efficiency of the detector are outputted into a ROOT file, which can be examined later.

Photon creation

In the simulation, a number of photons were created at a single point in the detector with an isotropic velocity vector, representing a scintillation event. The number of photons created was determined by the energy deposited, the light yield of the scintillator, and a randomized factor to simulate Poisson statistics. A scintillation yield of 11.5 photons/keV_{ee} was used, assuming a scintillator composed of PC with 1.5 g/L PPO, a value measured by [121] with an error of 10%. The start-time of the photons was distributed with an exponential decay with the lifetime of the scintillator. Photons were assigned a wavelength with a distribution similar to the emission spectrum of PPO.

In the simulation of the SABRE veto vessel, events with random energy were produced with the distance from the crystal following an exponentially decaying distribution with a

¹²This simulation was written by myself based on a code written by Dr. Alex Wright.

decay distance of 26 cm, which is the attenuation length for 1.5-MeV γ rays in the liquid scintillator.

Photon propagation

For purposes of tracking the time of propagation, the index of refraction of the scintillator was assumed to be 1.5044, the refractive index of PC [122]. Scattering, absorption, and reemission were all accounted for in this simulation. The scattering length of the scintillator was assumed to be 2 meters¹³. Photons were reemitted isotropically, as this process has a lower scattering length than anisotropic scattering [123].

Absorption could occur in two ways. First, the photon could be absorbed by the PPO and reemitted with a new wavelength. The probability of reemission was 0.82, the quantum yield of PPO. The absorption length curve for this process was taken from [124]. Secondly, a global attenuation length accounted for attenuation due to impurities in the scintillator. From experiments with a prototype neutron veto this value was tuned to 16 m, which is consistent with measurements done by Aldo Ianni for the DarkSide collaboration [125].

The photon could also hit a barrier, like the reflector. If this occurred, the photon was absorbed or reflected with a probability dependent on the wavelength. Reflection off the reflector was taken to be Lambertian, while reflection off the PMT or the crystal detectors was specular. The reflectance spectra were taken from measurements using a Lambda650 spectrophotometer. Notably, the measurement of the reflector was Lumirror soaked in PC and TMB for 10 months (see Figure 5.47).

Whichever mechanism (reflection, scattering, absorption by PPO, attenuation from impurities) was calculated to act first on the photon was executed.

¹³This value was adopted as a guess by Dr. Alex Wright in an earlier version of the optical Monte Carlo, and was used for the simulations for the prototype neutron veto, the DarkSide-50 veto and the SABRE veto. Subsequently, it was discovered that a better value for PC is closer to 7–11 m depending on the wavelength, as measured by [123]. Because the photons reflect several times on average before being detected, this difference is likely inconsequential to the predicted light yield.




Figure 5.47: Reflectance of the reflector Lumirror soaked in scintillator (PC+TMB), used in the simulation. The insert is zoomed into the high-reflectivity region.

Interaction with the PMT

When the photon hit a photocathode, it was first checked whether it passed the quantum efficiency of the PMT. The collection efficiency (85%) was then taken into account. If the photon passed the quantum efficiency check but not the collection efficiency, it was lost. If it failed the quantum efficiency check, it was reflected with a probability of 20% or otherwise absorbed. If the photon hit the dead region of the PMT, it reflected with a probability of 80%, which is similar to most shiny metals.

The detection efficiency was measured by looking at the time information of detections by individual PMTs and requiring a simple majority of 3 within a time-window of $1 \mu\text{s}$.

Wavelength-dependent quantities

The wavelength of an emitted photon followed the distribution in Figure 5.48, which is the emission spectrum of PPO in PC, digitized from [124]. The distribution was digitized at

10 nm intervals, and photons were emitted uniformly in the range of each wavelength bin. If a photon was absorbed and reemitted, it was reemitted at a longer wavelength according to this distribution.

The reflector spectra were taken with a Lambda-650 spectrophotometer and were digitized at 1-nm intervals (see Figure 5.47). The quantum efficiency of the Hamamatsu R5912 PMTs was sampled from the Hamamatsu literature at 10-nm intervals [126].




Figure 5.48: Emission spectrum of PPO used in the optical Monte Carlo simulation of the SABRE veto vessel. The output shown is the randomly generated emissions from the simulation, but the input parameters were the digitized spectrum of PPO in PC from [124].

Geometry of the simulation

The geometry of the simulation consisted of a cylinder representing the veto. At the center, seven cylindrical containers 4" in diameter in a hex packing represented the crystal-detector modules. The inner detector is assumed to reflect with the same reflectance as electropolished stainless steel (58%). None of the PMT ports were represented in this simulation. Neither was the mechanism for holding the crystals or any cables, since the dimensions of these features are not yet known. The PMTs were represented by a 13.1-cm-radius hemisphere

on the wall of the veto. The photocathode subtended an angle according to the diagrams for the 8" tubes by Hamamatsu.

Results for the SABRE portable veto

The predicted light yield for the SABRE portable veto is $0.21 \text{ p.e./keV}_{ee}$, as shown in Figure 5.49. This light yield is adequate to reject the absorbed γ rays from ^{40}K decay with high efficiency. It may also be able to detect the signal from the capture of thermal neutrons, which is around $30\text{--}40 \text{ keV}_{ee}$ [93]. The detection efficiency of the veto detector was simulated by counting the number of PMTs that detect at least one photoelectron. For a majority requirement of 3 PMTs, the detection efficiency as a function of energy is shown in Figure 5.50. The predicted light yield has a high optical detection efficiency above 40 keV_{ee} .




Figure 5.49: Predicted light yield of the SABRE veto vessel of $0.21 \text{ p.e./keV}_{ee}$. This light yield was simulated with a custom optical Monte Carlo that has been verified with the DarkSide-50 neutron veto [93, 110].




Figure 5.50: Predicted optical efficiency of the SABRE portable veto vessel as a function of deposited energy in keV_{ee} . Here, the requirement for detection is the detection of one or more photoelectrons in at least 3 PMTs.

5.5 The future of SABRE

SABRE is continuing its effort to develop low-radioactivity crystals that will surpass the purity of the DAMA crystals. Future efforts will include the additional purification of Astro-grade powder and an effort to further reduce key backgrounds in the crystal-growth process. The ^{40}K , tritium, and ^{87}Rb backgrounds will be of paramount importance in this effort.

While the primary focus of the SABRE effort has been the creation of highly pure crystal detectors, the deployment of those detectors in an underground setting with active shielding is also an important step to testing the DAMA/LIBRA result. There are a number of possibilities for housing the SABRE detectors. The deployment of SABRE in DarkSide has already been discussed in Section 5.2, but there are still a number of potential scenarios for the deployment of the portable liquid-scintillator vessel.

Initial operation of the liquid-scintillator vessel will be performed in Princeton wherein the vessel will be filled with water. This initial deployment is primarily to establish the operation of the veto electronics and data acquisition, and to establish procedures for a

liquid-scintillator run. After initial tests have been conducted in Princeton with water, a first deployment of the detector will occur with scintillator at LNGS. The veto vessel will need to be surrounded by passive shielding to be effective. The amount of rejection needed is between 10^3 and 10^4 , depending on whether the operation is a dark-matter measurement or a coincidence-counting measurement. The OPERA experiment is currently being decommissioned, freeing up a large mass of lead sheets that can be reformed for the purposes of the fabrication of lead shielding. This lead can be used to achieve the required level of passive shielding. A twin experiment can also take place at the Stawell gold mine in Australia, providing a check for seasonal effects.

With the dual efforts of crystal purification and the rejection of residual backgrounds through active rejection techniques, the SABRE experiment is working toward the achievement of new sensitivities to dark matter and a test of the controversial DAMA/LIBRA modulation result.

Chapter 6

Measurement of the Na

Nuclear-Recoil Quenching Factor:

The Experiment

6.1 Introduction

6.1.1 Prior publication of this work

At the time of this writing, a paper written primarily by Dr. Jingke Xu and myself on the Notre Dame Na quenching measurement is published on arXiv [1] and has been accepted for publication by Physical Review C. The content of Chapters 6, 7, and 8 are derived from that publication, and are discussed in more detail here. Some of the language in this work is directly taken from that publication. The analysis described in Chapter 7 is my own, independent analysis performed to separately validate the analysis done by Dr. Xu in the paper. This analysis was found to agree with Dr. Xu's analysis. Similarly, the analysis on the implications of this measurement for the dark-matter interpretation of DAMA/LIBRA, described in Chapter 8, is an independent analysis done by me separately from that in the paper.

6.1.2 The role of the ionization quenching factor in dark-matter experiments

While it is expected that WIMP dark matter can be observed interacting with the nuclei of atoms, dark-matter detectors' energy scales are most often calibrated with mono-energetic γ sources, which produce electron recoils. Scintillating detectors respond differently to nuclear and electron recoils, as described in Appendix A. The fraction of nuclear-recoil energy that is released as scintillation compared with that of electron recoils is known as the quenching factor.

A precise knowledge of the quenching factors for different nuclei is essential to understanding the detector response and ultimately producing a correctly-scaled energy spectrum for dark-matter-induced nuclear recoils. The scale of the energy spectrum is essential to determining the WIMP characteristics, particularly its mass and cross-section with nucleons. An accurate knowledge of the quenching factor is also essential to compare results between detectors with different targets. The Na quenching factor is especially important when evaluating the possibility that the DAMA/LIBRA modulation signal is attributable to light-WIMP dark matter.

In this and subsequent chapters, I will adopt the convention of using the unit keV_{ee} to describe the electron-equivalent energy of a recoil event, and keV_{nr} to describe the true energy of nuclear-recoil events. While keV_{ee} is a calibrated energy scale for nuclear-recoil events and is dependent on the energy of the calibration used, in this work, keV_{nr} reflects the actual energy of the nuclear recoil. I will therefore not use keV_{nr} to describe the nuclear-recoil-equivalent energy of an electron recoil event, as is sometimes done in the literature. The units of keV_{nr} and keV_{ee} are related by the quenching factor, $Q(E_R)$:

$$1 \text{ keV}_{nr} = Q(E_R) \text{ keV}_{ee}. \quad (6.1)$$

6.1.3 Previous measurements of the quenching factors in NaI(Tl)

The DAMA/LIBRA experiment assumes energy-independent quenching factors for both Na and I [53]. DAMA measured their reported values of 0.3 for ^{23}Na and 0.09 for ^{127}I by observing the neutron-induced nuclear-recoil spectrum with a ^{252}Cf source and fitting it to a simulated spectrum. They measure the light yield of their detector with a ^{137}Cs source during this measurement to determine the energy scale. In their experiment, they continue to use ^{137}Cs for some calibrations, but they measure their light yield with ^{241}Am [41].

Since this measurement, several other groups have measured the Na quenching factor with mono-energetic neutron sources to test the energy-independence of the quenching factor [127, 128, 129, 130, 131], primarily by the $d(d, ^3\text{He})n$ reaction. These experiments measure the Na quenching factor at higher recoil energies ($\gtrsim 20 \text{ keV}_{nr}$) and carry large uncertainties.

These uncertainties stem from, in part, the high energy of the incoming neutrons (2.85–11 MeV), which require small scattering angles to produce low-energy recoils. The relative angular uncertainty caused by the finite size of the detectors, $\Delta\theta/\theta$, is larger for smaller angles, and the spread in the recoil energies is also higher. The angular uncertainty in some of these experiments is also compounded by a larger crystal size [130, 55]. This larger crystal size also increases the rate for multiple-scattering events, which distort the energy spectrum at a given scattering angle. Finally, the light-collection of some of these experiments was limited by the capabilities of their photomultiplier tubes, which had lower quantum efficiencies than are available today. These measurements were fairly consistent with the DAMA values at higher recoil energies. The results of these measurements are shown in Figure 6.1.

Later measurements by Collar [54] and Chagani [55] found quenching values that deviated significantly from the DAMA values in their modulation region, and also differ dramatically from each other, as can be seen in Figure 6.1. These recent measurements prompted SABRE to conduct an independent measurement of the Na quenching values in the DAMA/LIBRA-modulation energy region.




Figure 6.1: Previous measurements of the Na quenching factor, from Collar et al. [54]. Most previous experiments found Na quenching factors consistent with the DAMA/LIBRA values at energies above $\sim 30 \text{ keV}_{nr}$, but Collar [54] (marked “this measurement” in the figure) and Chagani [55] observed lower values for the Na quenching factor at lower energies. These values also conflict with each other but have large error bars.

6.1.4 The SCENE approach

The measurement of the quenching factor in Na was inspired by similar measurements of quenching effects in liquid argon conducted by the SCENE collaboration [132, 133] at the University of Notre Dame Institute for Structure and Nuclear Astrophysics. Using the pulsed beam facility, SABRE conducted measurements of the quenching factor that improved on the previous measurements in the following ways:

1. While previous experiments used a time-coincidence between NaI(Tl) nuclear-recoil events and subsequent interaction of the neutron in a neutron detector, SABRE used a triple time-coincidence between a pulsed neutron beam, the NaI(Tl) detector, and

a neutron detector to select neutron elastic-scattering events and to reduce random-coincidence backgrounds, providing a more stringent cut than previous double coincidences.

2. The SABRE measurement used low-energy neutrons (~ 690 keV) so that low-energy nuclear recoils (< 50 keV_{nr}) could be obtained at large neutron-scattering angles so that relative angular uncertainties could be suppressed compared with previous measurements.
3. A small NaI(Tl) crystal (a 25-mm cube) was used to suppress multiple-scattering backgrounds.
4. A high-quantum efficiency PMT was used to enhance light collection (~ 18 photoelectrons/keV_{ee} achieved).
5. Pulse-shape discrimination (PSD) was used to select neutron events among background and noise.

In our measurement, a pulsed beam of protons from the accelerator interacted with a LiF target via the ${}^7\text{Li}(p,n){}^7\text{Be}$ reaction (Q-value: -1.644 MeV), emitting neutrons with a nominal energy of 690 keV at 0° incidence. A detector consisting of an enclosed NaI(Tl) crystal and a coupled PMT was placed on the beam line. The neutrons traveling in this direction could produce nuclear-recoil scintillation events in the crystal and be subsequently detected by a liquid-scintillator-based neutron detector at a fixed recoil angle. The kinematics of this interaction determined the nuclear-recoil energy.

The recoil energy, determined by kinematics, combined with the scintillation light collected by the NaI(Tl) detector, provided a measure of the light yield of the detector for nuclear-recoil events. Calibrating this measurement with an electron recoil of known energy allowed for a determination of the quenching factor.

SABRE used a single NaI(Tl) detector in two position configurations and a stationary array of six neutron detectors, thereby measuring 12 nuclear recoil energies. In the first




Figure 6.2: To-scale, bird's-eye view of the experimental setup in the first position configuration with a side view of the NaI(Tl) detector. A LiF target on the proton beam path emits neutrons when struck by the beam. The neutrons travel to the NaI(Tl) detector, where they scatter off a sodium or iodine nucleus. If they scatter with the correct energy transfer, they may travel to one of the neutron detectors positioned at the relevant angle. Two different types of neutron detector were used with different sizes, as described in the text. A polyethylene collimator blocks neutrons from hitting the neutron detectors directly from the LiF target. Also shown is a side view of the NaI(Tl) detector, with the crystal shown in white and the PMT in green.

position configuration, the NaI(Tl) detector was placed 50 cm from the LiF target, and 91 cm in the second. A scheme of the experimental setup in the first position configuration is shown in Figure 6.2.

6.2 Preparation for the measurement

Before the measurement was performed, several parameters of the experimental setup needed to be chosen to optimize the rate and energy spread in the quenching data. Several calculations were performed to choose the optimal proton energy, the recoil energies to be measured, the LiF thickness, and the positions of the detectors¹.

6.2.1 Choosing the recoil energies

Previous measurements of the Na quenching factor are well established for nuclear recoils around 30 keV_{nr} . The Collar measurement measured nuclear recoils in Na down to 8 keV_{nr} [54]. The aim of the SABRE measurement was to reach recoil energies at least as low as that of the Collar measurement, and in addition, to completely cover the DAMA/LIBRA region of interest. Given that DAMA has recently lowered their energy threshold to 1 keV_{ee} [46], and the assumption that the Na quenching factor would be 0.3, SABRE aimed to reach recoil energies that would be equivalent to 1 keV_{ee} , or $\sim 3 \text{ keV}_{nr}$. Six energies were initially chosen between 3 and 30 keV_{nr} . Sufficient data for these energies were gathered before the beam time was over, so an additional position was selected for the NaI(Tl) detector in order to measure 6 additional recoil energies. The NaI(Tl) second position was chosen to maximize the diversity in the recoil energies measured below 15 keV_{nr} .

6.2.2 Choosing the proton energy

The proton energy was chosen to optimize the event rate for neutron-induced nuclear recoils while reducing as much as possible the spread in the recoil energy for a given neutron detector. There are several ways in which the chosen proton energy affects this spread. The relative spread in the recoil angle and the spread in the recoil energy is smaller for larger angles. This is due to the finite detector size and the characteristics of the recoil energy's

¹These calculations were performed by me while Dr. Jingke Xu developed the data acquisition software for the measurement.

angular dependence. Because six different energies were to be measured at a time, the energy of the proton beam was chosen to obtain a large angular spread between different detectors for the recoil energies of interest. A calculation of the event rate and the quenching factor uncertainty was performed to choose the optimal proton energy as well as other parameters.

Because the neutron energy is dependent on the scattering angle off the LiF target, the position of the NaI(Tl) detector with respect to the beam line had to be specified before determining the proton energy. There is some evidence that the γ background from the LiF target is lower at a 15° scattering angle [134]. However, the neutron flux is slightly forward-peaked, and the γ background can be strongly suppressed with time-of-flight cuts. Therefore, the differential cross-sections as a function of proton energy for the ${}^7\text{Li}(p,n){}^7\text{Be}$ reaction for both angles were digitized from [135], shown in Figure 6.3. The total cross-sections for elastic and inelastic neutron scattering off Na were also digitized from [136].

For different neutron energies, the recoil angle was calculated for several energies of interest, and a weight was given to the neutron elastic-scattering cross-section based on the differential scattering cross-sections at that angle, as determined by [136]. A polynomial fit to these values was applied so that the total elastic cross-section could be weighted by the differential cross-section for the desired recoil energies. For each incident neutron energy, the weighted cross-section for the neutron elastic scattering was multiplied by the ${}^7\text{Li}(p,n){}^7\text{Be}$ cross-section at the proton energy that would produce that neutron. This produces, in effect, a “combined cross section” for the interaction in which an incident proton results in a neutron-induced elastic-scattering event of a desired energy in the NaI(Tl) detector. Because elastic scattering is desired, the inelastic scattering cross section for the chosen incoming neutron energy should be low—this was taken into account when choosing the proton energy. A higher proton energy will produce neutrons with energy high enough to excite ${}^{23}\text{Na}$ to an excited state (the first excited state is 440 keV). It was later decided that the detector would be placed on the beam-line to simplify calculations. For a NaI(Tl) detector




Figure 6.3: (Top) 0° -incidence ${}^7\text{Li}(p,n){}^7\text{Be}$ cross-section as a function of proton energy, from Burke et al. [135]. (Middle) The ratio of the Burke data to other measurements. (Bottom) The errors on the Burke data.

on the beam line, the combined elastic-scattering cross sections for several recoil energies of interest as a function of the incident proton energy are shown in Figure 6.4.

From this calculation, it can be seen that an energy of 2.44 MeV has multiple advantages. The peak in the combined cross-section is the highest and is fairly broad at this energy. The combined cross section falls off slowly as the proton energy decreases, which is important because the proton may lose some energy in the LiF target before it interacts with the ${}^7\text{Li}$. Finally, higher neutron energies have a higher cross-section for inelastic scattering, which would produce some unwanted background.

6.2.3 Choosing the LiF thickness

As protons travel through the LiF target, they continuously lose energy predominantly due to interactions with electrons in the material, as described in Equation A.1 in Appendix A. This energy loss causes a spread in the proton energy at the time when it undergoes the ${}^7\text{Li}(p,n){}^7\text{Be}$ reaction, thereby broadening the incoming-neutron energy spectrum. To demonstrate this effect, the stopping power for protons in LiF, defined as the energy loss




Figure 6.4: Combined cross section of a neutron-induced Na recoil of the recoil energy given in the legend as a function of the incoming proton energy for an on-beam NaI(Tl) detector. The ${}^7\text{Li}(p,n){}^7\text{Be}$ cross section and neutron-elastic-scattering cross sections are combined to calculate the probability that a proton with a given energy will produce a neutron-induced nuclear-recoil in the NaI(Tl) detector with the desired energy. From this calculation, a proton energy of 2.44 MeV was chosen for the measurement.

of the particle as a function of distance traveled, is given as a function of energy in Figure 6.5. A 2.5-MeV proton in LiF has a stopping power of $\sim 100 \text{ MeV cm}^2/\text{g}$. Since the stopping power is approximately constant over this energy range, the energy loss over a $0.5\text{-mg}/\text{cm}^2$ target is $100 \text{ MeV cm}^2/\text{g} \times 0.0005 \text{ g}/\text{cm}^2 = 0.05 \text{ MeV}$, or 50 keV.

The rate of the ${}^7\text{Li}(p,n){}^7\text{Be}$ reaction directly scales with the target thickness, since the interaction length for this reaction and the proton-penetration depth are long compared to the thickness of the target. Therefore, the LiF target needed to be as thick as possible given an acceptable amount of spread. In this measurement, the budgeted spread was $\sim 5\%$, since the aim for the total spread was 20%, and some spread would be introduced from the finite size of the detectors. To calculate this spread in the neutron energy, a Geant4 simulation was done. Unfortunately, Geant4 does not model the ${}^7\text{Li}(p,n){}^7\text{Be}$ reaction. To accommodate for this, the ${}^7\text{Li}(p,n){}^7\text{Be}$ reaction was assumed to have an equal probability of occurring as a function of proton depth, and the energy of the protons were tracked as a function of

depth in the LiF. The resultant 0° neutron energy was calculated according to a relativistic kinematic calculation. A histogram was filled with the outgoing neutron energy as a function of depth in the LiF, weighted by the ${}^7\text{Li}(p,n){}^7\text{Be}$ cross-section as a function of proton energy from [135], since the interaction probability to emit a neutron is dependent on the proton energy (Figure 6.6a). This histogram was integrated to produce a neutron-energy spread as a function of the LiF thickness. In order to predict the spread in neutron energy for LiF targets of various thicknesses, a target much thicker than the one eventually used was simulated. For the 2.44-MeV proton chosen (as described in Section 6.2.2), the energy of the outgoing neutron was found to be 690 keV with a spread of 4% at the full-width-half-maximum for a LiF target 0.5 mg/cm^2 thick, as shown in Figure 6.6.

6.2.4 Choosing the detector positions

The choice of detector positions, given the proton energy and LiF thickness, was a compromise between event rate and the spread in the nuclear-recoil energy. The finite size of the detectors introduces some spread in the nuclear-recoil energy which can be mitigated by placing the detectors farther apart. However, doing so decreases the event rate. A calculation of the expected event rate and the nuclear-recoil spread for given detector positions was performed.

The event rate calculation

The event rate of the nuclear recoils is the product of neutron flux incident on the NaI(Tl) detector and the cross section for a recoil at an angle that will be detected by a given neutron detector. The flux of neutrons at the distance of the NaI(Tl) detector from the LiF, Φ_n , is the product of the flux of protons into the LiF target and the chance that a neutron will be emitted in the solid angle of the detector:

$$\Phi_n = R_p \left(\frac{d\sigma}{d\Omega_p} \right)_0 \frac{DA_g}{m_{\text{LiF}} x_1^2}, \quad (6.2)$$




Figure 6.5: Stopping power for protons in LiF as a function of proton energy, from the NIST pstar database [137]. From the stopping power listed for a ~ 2.5 -MeV proton, ~ 100 MeV cm^2/g , the expected energy loss of the proton over a $0.5\text{-mg}/\text{cm}^2$ target is small: about 50 keV. The stopping power of the material is therefore approximately constant over this range, making the proton mean energy loss linear with distance traveled through the target.

where R_p is the rate of incoming protons in protons/s, $\left(\frac{d\sigma}{d\Omega_p}\right)_0$ is the differential cross-section of the ${}^7\text{Li}(p,n){}^7\text{Be}$ reaction at 0° scattering angle, D is the LiF target depth in g/cm^2 , A_g is Avogadro's number, m_{LiF} is the LiF mass in a.m.u., and x_1 is the distance from the LiF target to the NaI(Tl) detector in cm. The event rate for recoils into a neutron detector with recoil energy E_R , R_{E_R} is:

$$R_{E_R} = \Phi_n \frac{d\sigma}{d\Omega_{n,\text{EL}}}(\theta) \frac{N_{\text{NaI}} A_{\text{ND}}}{x_2^2} \epsilon_{\text{ND}}, \quad (6.3)$$




Figure 6.6: Results of the simulation to determine the LiF target thickness for a 2.44-MeV proton. (Top) Histogram of the neutron energy as a function of depth in the target, weighted by the ${}^7\text{Li}(p,n){}^7\text{Be}$ cross-sections from [135]. The z -axis is counts. (Bottom left) Distribution of neutron energies as a function of LiF thickness. This plot was achieved by integrating over the top plot. (Bottom right) Projection of bottom-left plot at a LiF thickness of 0.5 mg/cm^2 , showing the expected neutron-energy spread at that thickness. This thickness is what was used in the measurement.

where $\frac{d\sigma}{d\Omega_{n,EL}}(\theta)$ is the differential cross-section for neutron elastic scattering for the given scattering angle θ , N_{NaI} is the number of NaI molecules in the crystal, A_{ND} is the cross-sectional area of the neutron detector, x_2 is the distance from the NaI(Tl) detector to the neutron detector, and ϵ_{ND} is the detection efficiency of the neutron detector, which was assumed to be a constant 30%² for this calculation.

The energy spread calculation

The energy spread in the nuclear recoils is a product of two factors: the angular spread and the spread of the incoming neutron energy. The uncertainty in the quenching factor is also dependent on photon statistics, which in turn is a function of the light yield of the detector. The uncertainty in the quenching factor was calculated as a function of the detector positions, given the previously discussed neutron-energy spread and the measured light yield of the NaI(Tl) detector.

To calculate the angular spread, the minimum and maximum scattering angles were calculated given the positions for the NaI(Tl) and neutron detectors. A Geant4 simulation³ showed that the mean angular spread was approximately 1/3 of the maximum angular difference given the finite detector sizes. A linear fit was done with the Monte Carlo data to obtain the angular resolution as a function of the maximum angular difference for a given geometrical configuration.

To calculate the spread in the recoil energy of the Na nucleus, the dependency of the recoil energy on experimental parameters was derived from relativistic kinematics (equations derived from other relations in [138]):

$$E_R = \sqrt{p_{cm}^2 + m_{\text{Na}}^2} \cosh(\zeta) - \tanh \zeta \left[p_n \cos(\theta_n) - \sqrt{p_{cm}^2 + m_n^2} \sinh(\zeta) \right], \quad (6.4)$$

²Based on the typical peak quantum efficiency of PMTs.

³performed by Dr. Jingke Xu.

where ζ is the center-of-mass rapidity:

$$\zeta = \ln \left[\frac{p_{cm} + \sqrt{p_{cm}^2 + m_{\text{Na}}^2}}{m_{\text{Na}}} \right]. \quad (6.5)$$

p_{cm} is the momentum of the neutron and the Na nucleus in the center-of-mass frame, m_{Na} is the mass of the Na nucleus, θ_n is the neutron scattering angle in the lab frame, m_n is the mass of the neutron, and p_n is the momentum of the neutron after the collision in the lab frame, which has the form:

$$p_n = \frac{\sqrt{m_n^2 + p_{cm}^2} \cos(\theta_n) \sinh(\zeta) \pm \cosh(\zeta) \sqrt{p_{cm}^2 - m_n^2 \sin(\theta_n)^2 \sinh(\zeta)^2}}{1 + \sin(\theta_n)^2 \sinh(\zeta)^2}, \quad (6.6)$$

and

$$p_{cm} = \sqrt{\frac{(2m_n m_{\text{Na}} + 2m_{\text{Na}} T_k)^2 - 4m_n^2 m_{\text{Na}}^2}{4((m_n + m_{\text{Na}})^2 + 2m_{\text{Na}} T_k)}}, \quad (6.7)$$

where T_k is the kinetic energy of the incoming neutron. The spread in the nuclear recoil energy E_R depends on the spread in T_k as well as the spread in the scattering angle of the neutron, θ_n . The analytical form of the partial derivatives of E_R with respect to these variables were calculated based on the equations above and combined with the spread in T_k and θ_n to obtain an expected spread in E_R before photon statistics. Finally, a quenching factor of $Q = 0.3$ or 0.2 in Na was assumed to predict the spread due to a poisson distribution of photons with mean number of photons detected, $\langle N_\gamma \rangle = LQE_R$, where L is the light yield and Q is the assumed quenching factor. This spread was combined with the others to determine an overall spread in the recoil energy, as well as an uncertainty in the quenching factor. In choosing the detector positions, the goal was to achieve a relative quenching-factor uncertainty of less than 30% at the lowest energies, and a rate greater than 0.01 Hz for each detector⁴. A table of the parameters calculated, including the energy spread, and which ones were chosen⁵ is shown in Table 6.1.

⁴0.01 Hz would give a few thousand events over two days of beam time.

⁵These values were used for the first six positions measured. The second set of six recoil energies were chosen during the measurement to achieve more energies without moving the neutron detectors.

Table 6.1: Calculated event rates and recoil-energy spreads for various detector configurations, assuming a beam current of 100 nA with a bunch-selection ratio of 1:6, a neutron detector efficiency of 0.3 and a Na quenching factor of 0.3 or 0.2 (for the purposes of applying Poisson statistics). These calculated quantities were used to choose the detector positions for the Na quenching-factor measurement. Here, x_1 is the distance from the LiF target to the NaI(Tl) detector, and x_2 is the distance from the NaI(Tl) detector to the neutron detector. $\Delta E_R/E_R$ and $\Delta Q/Q$ are the calculated relative spread in the recoil energy and the quenching factor, respectively. A designator of “Y” indicates that this position was used in the measurement for position configuration 1.

E_R (keV _{nr})	Ang. (deg)	ND size (in)	x_1 (cm)	x_2 (cm)	Event Rate (Hz)	$\Delta E_R/E_R$	$\Delta Q/Q$ Q=0.3	$\Delta Q/Q$ Q=0.2	Used?
3	18	2"	50	70	0.028	0.16	0.37	0.44	Y
3	18	2"	100	70	0.007	0.12	0.35	0.42	
3	18	5"	50	150	0.038	0.15	0.36	0.43	
3	18	5"	100	150	0.009	0.11	0.35	0.42	
3	18	5"	50	200	0.021	0.11	0.35	0.42	
3	18	5"	100	200	0.005	0.07	0.34	0.41	
6	25	2"	50	70	0.026	0.11	0.26	0.31	
6	25	2"	100	70	0.007	0.09	0.25	0.30	
6	25	5"	50	150	0.035	0.11	0.26	0.31	
6	25	5"	100	150	0.009	0.08	0.25	0.30	
6	25	5"	50	200	0.02	0.08	0.25	0.30	Y
6	25	5"	100	200	0.005	0.06	0.24	0.29	
10	33	2"	50	70	0.024	0.09	0.20	0.24	Y
10	33	5"	50	150	0.033	0.09	0.20	0.24	
10	33	5"	50	200	0.018	0.07	0.19	0.23	
15	41	2"	50	70	0.019	0.07	0.16	0.20	
15	41	5"	50	150	0.026	0.07	0.16	0.19	Y
20	48	2"	50	70	0.018	0.06	0.14	0.17	Y
20	48	5"	50	150	0.024	0.06	0.14	0.17	
20	48	5"	50	200	0.013	0.05	0.14	0.17	
30	60	2"	50	70	0.012	0.05	0.12	0.14	
30	60	5"	50	150	0.017	0.05	0.12	0.14	Y
30	60	5"	50	200	0.009	0.04	0.11	0.14	

The distance between the LiF target and the NaI(Tl) target was chosen to be 50 cm. For a larger distance, the neutron flux would be unsuitably low. For a smaller distance, the angular spread would be too high. Given this distance, the event rate and spread in E_R were plotted as a function of the distance between the neutron detectors and the NaI(Tl) detector. Two sizes of neutron detector were available for the experiment, a 5.1-cm (2") detector and a 12.7-cm (5") detector. It was determined that the trade-off between event rate and energy spread was optimized at 150 cm for the 5" detectors and 70 cm for the 2" detectors. For the 6-keV $_{nr}$ recoil, we chose to place the detector at a further distance (200 cm) to decrease the spread at the cost of some event rate.

6.3 Experimental setup

Twelve recoil energies at twelve different scattering angles were measured, as shown in Table 6.2. The details about the detectors, their positioning, and the electronics are given in this section.

Table 6.2: Detector information and positions for position configurations 1 (Top) and 2 (Bottom). The “flight distance” for the neutron detectors (ND) refers to the distance from the NaI(Tl) detector to the neutron detector, while for the NaI(Tl) detector it is the distance from the LiF target to the NaI(Tl) detector. The neutron detectors are cylindrical Eljen detectors with EJ-301 as the scintillator and EJ-510 as the reflector. The detector size given for the neutron detectors is both the diameter and length of the cylindrical scintillator volume, while the size given for the NaI(Tl) detector is the length of each side of the cube.

Detector	Detector size (cm)	Scattering angle (deg)	Recoil energy (keV _{nr})	Flight distance (cm)
NaI(Tl)	2.5	0		50
		0		91
ND1	12.7	59.1	29.0	150
		74.2	43.0	135
ND2	12.7	41.3	15.0	150
		54.4	24.9	122
ND3	12.7	24.9	5.7	200
		31.1	8.8	164
ND4	5.1	47.9	19.4	70
		84.0	51.8	52
ND5	5.1	32.2	9.1	70
		64.6	33.3	41
ND6	5.1	18.2	2.9	70
		41.1	14.3	33

6.3.1 The NaI(Tl) detector

The NaI(Tl) detector consisted of a 25-mm cubical NaI(Tl) crystal optically coupled to a 76-mm super-bialkali Hamamatsu R6233-100 PMT. The crystal was grown at Radiation Monitoring Devices Inc. with high-purity Astro-grade NaI powder from Sigma Aldrich, shown in Figure 4.4. The crystal has a Tl concentration of $\sim 0.1\%$, similar to the DAMA/LIBRA

crystals [41]. The small crystal size was chosen to minimize the probability of neutron multiple scattering.

The enclosure used in the quenching measurement⁶ is shown in Figure 6.7. This enclosure was designed to seal the crystal in a light-tight environment that would also protect the crystal from moisture. Another design consideration was the desired minimization of material in the path of the neutron, to avoid multiple-scattering of the neutron in detector materials. Finally, the components of the detector were chosen to maximize the light yield, so that low-energy recoils would be detected with a high efficiency.

The enclosure consisted of a stainless-steel can with nominal 1/16" wall thickness with a flange welded on one end and a sealed cap on the other. The middle section of the enclosure was machined down to 0.5 mm to minimize the material between the LiF target and the crystal. The can was sealed with a cover-plate with a 1/8" viton o-ring. The cover-plate contained two half-inch bulkhead feedthroughs for high-voltage and two for signal cables, as well as a gas feedthrough which could be used to pump and purge the enclosure. The cover-plate also contained three threaded holes where stainless-steel rods could be affixed to hold the inner components.

The inner components of the detector were held in place in the center of the enclosure, as shown in Figure 6.7. Three stainless-steel sleeves were placed on the rods to set the vertical location of the PMT holder. The PMT holder consisted of a PVC piece that cradled the curved bulb of the phototube, a teflon cover plate for the PMT with a crystal-sized hole in the center, and three Teflon screws that held the pieces together around the PMT. The crystal was held in place with an aluminum holder designed to have a minimal mass while holding the crystal firmly in place. It consisted of a square piece with a square hole in which the crystal sat, and three legs extending to the rods. The aluminum holder was held in place in turn by another set of sleeves and a set of springs that pressed everything together.

⁶The enclosure was designed by Allan Nelson, Dr. Jingke Xu, and myself and constructed by Allan Nelson for an initial measurement of the Na quenching factor conducted in 2013. Modifications for this measurement, including the construction of the new aluminum holder, were performed by myself. The packaging of the crystal was done by Dr. Xu and myself, and the leak-testing and initial light-yield tests were done by Dr. Xu.

Before packaging, the crystal was found to have suffered degradation due to moisture exposure since a previous, preliminary quenching measurement conducted in the previous year. The crystal had yellowed and cracked at the edges. The crystal was sent to RMD to be polished, and the yellow material was removed. This process restored the crystal to a state of high clarity. The crystal was polished smooth on one face, while the other faces were left rough. The crystal was coupled to the PMT with a transparent optical gel from Cargille Labs with a refractive index of 1.52 (number 081160). No light-guide was used in this experiment in order to maximize the light collection. A piece of Lumirror reflector was made into a 5-faced box and wrapped around the crystal, and was further sealed in place with 4–5 layers of Teflon tape inside the reflector, and 2–3 layers outside. The Lumirror served as a back-up reflector for the Teflon tape. The PMT had a high peak quantum efficiency of $\sim 35\%$ and a fast rise-time, making it a suitable choice for this experiment.

During the assembly, all parts were cleaned with ethanol and were assembled in a nitrogen-filled glove box with a humidity level between 0.6% and 1% relative humidity. After assembly, the enclosure was pumped overnight and purged with nitrogen for several cycles. The enclosure was also leak-tested with a helium leak-tester and was found to be acceptably leak-tight. In Princeton, the light yield was initially tested to be ~ 19 p.e./keV_{ee}, measured with a ^{137}Cs source.

6.3.2 The proton beam and the LiF target

The proton beam comes from an 11-MV FN Tandem accelerator. Negative ions are produced by cesium-sputtering, which are accelerated to high energy and passed through a thin carbon foil terminal. The foil strips electrons from the ions in the beam, leaving the beam positively charged. This positively charged beam then accelerates away from the terminal and exits the accelerator at even higher energy [139].

The protons produced by this accelerator, incident on a LiF target, cause the emission of neutrons through the $^7\text{Li}(p,n)^7\text{Be}$ reaction (Q-value: -1.644 MeV). Ignoring energy loss

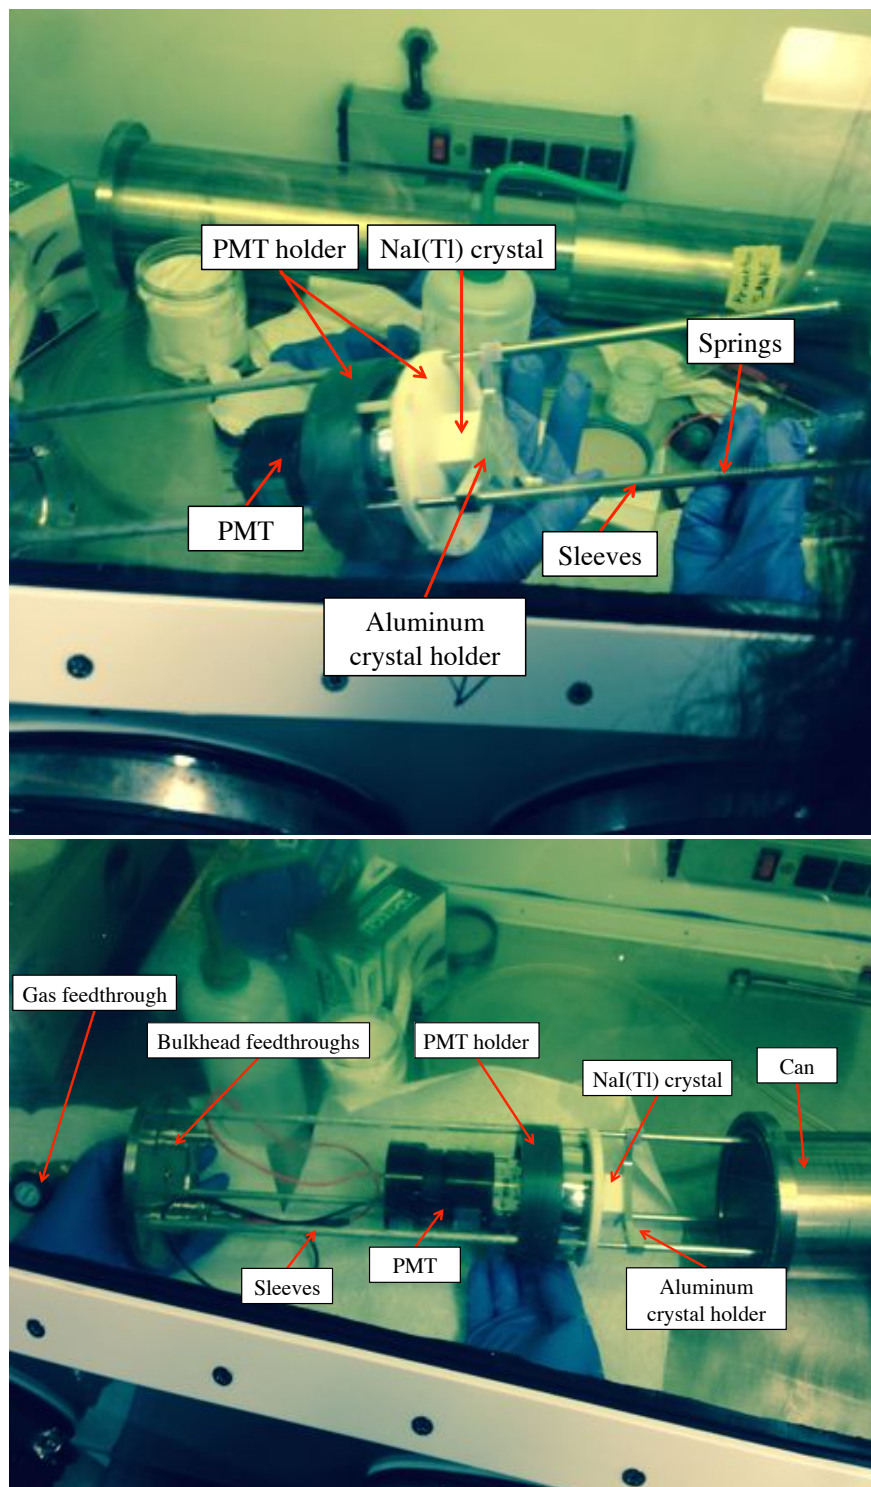


Figure 6.7: Enclosure used for the SABRE Na quenching measurement. (Top) The internal components to the enclosure. (Bottom) The internal components as they are inserted into the stainless-steel can. Both pictures were taken through the window of the glove box used for assembly.

in the target, the neutrons emitted at a nominal scattering angle of 0° have an energy of 690 keV.

The beam was separated into time-bunched pulses by a three-part pulsing system with a timing resolution of 2 ns and an intrinsic period of 101.5 ns. The pulse-selecting system consists of a buncher, a sweeper, and a pulse selector. The buncher alternately accelerates and decelerates the beam as it passes through using an alternating high voltage, thereby bunching the beam into pulses. The sweeper uses another alternating voltage to sweep away protons not in the bunches. The pulse selector, which is upstream of the buncher, deflects the beam from reaching the buncher on a separate timer. The pulse-selector can be used to allow only one out of every n pulses to pass through, effectively increasing the period [139]. The beam cross section was 3 mm in diameter. The beam was also highly stable, with a variation in proton energy of around 1 keV_{nr}.

The energy of the beam, 2.44 MeV, was chosen to increase the event rate while reducing the spread in the nuclear recoil energy, as was described in Section 6.2.2.

A higher pulse frequency can lead to a higher neutron event rate, but it may also cause pileups in the NaI(Tl) scintillation time window, especially since the scintillation time of NaI(Tl) is 200–300 ns. Based on the calculation of the neutron flux in Section 6.2.4, a pulse separation of 609 ns was thought to be enough to reduce the pileup rate. Based on this assumption, the bunching ratio for the pulser was set at first to 1 in 6, for an effective pulse period of 609 ns, but was later changed to 1 in 8, for a pulse period of 812 ns, after a higher event rate than expected was observed (see Table 6.7). Each pulse carried $\sim 1 \times 10^4$ protons.

The LiF target was deposited on a 0.4-mm tantalum backing, which stops the proton beam and minimizes the γ background. Incoming protons lose energy as they travel through the LiF target, leading to a broadening of the outgoing neutron energy spectrum. The LiF target thickness was chosen to be 0.52 mg/cm² LiF in order to compromise between the event rate and the spread in the neutron energy, as was described in Section 6.2.3. The mean neutron energy for a target thickness of 0.52 mg/cm² was calculated to be 690 keV with a

spread of 4% at the full-width-half-maximum. At that thickness, and with a bunching ratio of 1 in 8, the neutron flux was calculated to be around 300 neutrons/s at the NaI(Tl) detector for the first position configuration and about 100 neutrons/s in the second.

A ϕ 22-cm, 22-cm-long, cylindrical polyethylene collimator with a ϕ 2.5-cm hole was used to prevent neutrons from traveling directly from the LiF target to the neutron detectors.

6.3.3 The neutron detectors

Two types of neutron detectors were used for the measurement: 5.1-cm Eljen 510-20x20-9/301 and 12.7-cm Eljen 510-50x50-1/301 liquid-scintillator detectors. Both types have the reflector EJ-510 and the liquid scintillator EJ-301, a xylene-based scintillator with organic fluors. This scintillator has pulse-shape discrimination capability, which allows for the selection of events induced by desired particle types. A typical light-yield response of these detectors was measured to be ~ 1 p.e./keV_{ee}.

The angles chosen for the neutron detectors, between 18 and 84 degrees, allowed for data to be collected for Na nuclear-recoil energies between 3 and 52 keV_{nr}. The detector distances were chosen to maximize the event rate while maintaining a recoil energy uncertainty due to finite detector size of less than 5%. Their positions are summarized in Table 6.2, and their performance is summarized in Section 6.3.5.

6.3.4 Placing the detectors

The NaI(Tl) detector was placed on the beam-line axis to maximize the event rate. The NaI(Tl) detector was placed 50 cm from the LiF target in the first position configuration and 91 cm in the second configuration, as described in Table 6.2. The 50 cm position was chosen in order to produce a high event rate while keeping the total angular spread below 5%. The 91 cm position was chosen to measure a variety of recoil angles with the existing neutron detector positions. The placement of the detectors was done with the aim of producing a

positional uncertainty of around 1 cm for each detector⁷. The procedures for placing the detectors and measuring their positional uncertainty are described below.

A self-leveling laser was used to establish the location of the beam line. The laser was situated at the height of the beam. This laser produces both a horizontal and vertical line. The horizontal location of the beam line at the back of the room (where the beam enters the room) was previously marked by the SCENE collaboration. This mark, along with the LiF target and a similar mark on a beam in front of the LiF target, were used to align the laser horizontally with the proton beam line. Tape marked with the vertical location of the beam line was also placed on a wall on the side of the room and a pipe on the other side. These marks were used to align the laser vertically with the beam. Some of the marks used to align the laser are shown in Figure 6.8. The laser projected a vertical plane on the floor. A long piece of tape was placed on the floor along this line, and a pen was used to mark the line. A plumb-bob was used to mark the location of the front of the holder for the LiF target along this line. A caliper was used to measure the depth of the LiF target in the holder, and this distance was marked on the tape to reflect the horizontal location of the LiF target. This location was used as the origin for further measurements.

The desired x and y coordinates for the central locations of the detectors were calculated. The y positions along the beam line were marked for each detector. A large T-square and a meter stick were used to draw perpendicular lines along the floor at each of these locations. The desired positions for the detectors were marked (see Figure 6.9). The positions of these markers were measured with a 2-m-long meter stick and a tape measure, both the distance from the target as well as the x and y positions to confirm the accuracy of the marker placement. The marker locations are shown in Table 6.3.

The NaI(Tl) detector was held by a mount provided by Fermilab, shown in Figure 6.10 and Figure 6.11, which consisted of a tripod with a custom enclosure for the detector. The enclosure consisted of a ring of wooden panels held together with a metal ring in the middle,

⁷This effort was led by myself with help from Shawn Westerdale and Dr. Henning O. Back.



Figure 6.8: Aligning the laser. The laser was aligned with two pieces of tape at beam height on either side of the beam (Top left and top right) and a vertical line in front of (Bottom left) and behind the beam, as well as the target itself (Bottom right). The marks had been previously made by the SCENE collaboration.




Figure 6.9: The markers denoting the desired detector positions. The central axis was marked using the laser level such that it was parallel to the beam line. The detector locations along this axis were marked. A T-square was used to draw a second axis for each detector emanating from the beam line.

Table 6.3: Measured marker positions for the first position configuration. The calculated distances use the subtracted x and y values, while the measured distances were measured directly with a long ruler.

Detector ID	Nominal Energy (keV _{nr})	x (cm)	y (cm)	Distance from NaI (calculated)	Distance from NaI (measured)
NaI(Tl)	N/A	0.0	50.00	50.00	50.00
1	30	130.05	125.15	150.20	150.3
2	15	98.33	163.25	149.98	150.00
3	6	84.35	231.25	199.92	200.20
4	20	-52.10	96.85	70.07	70.05
5	10	-38.05	108.75	70.00	70.05
6	3	-21.50	116.63	70.01	70.00

a floor on the bottom, and a clamp at the top. The floor was connected to a metal tube which fit in the tripod. The tripod held the tube in place with a set screw. The detector was marked so that it could be oriented such that the flat crystal face was perpendicular to the beam line and that one rod was on the beam line behind the crystal while the other two straddled the beam line.

The tripod had a circular ring in the center, which was used to align the detector with the floor. The diameter of the ring was measured to be 10.8 cm. A mark was made along the beam line at half this distance away from the central mark denoting the desired position for the front of the ring. The mount was aligned with this mark with a plumb-bob. First, the center of the mount was aligned with the beam by eye. The plumb bob was held along the edge of the circle and was moved slowly along the edge to assure that the point closest to the target in the y direction was on the beam axis. Holding the plumb-bob at this point, the mount was aligned in the y direction with the mark. The stand was taped down to the floor with duct tape.

To measure the final location of the NaI(Tl) detector, a circle made of stock paper was cut with the diameter of the ring, and a hole was punched in the center. Three marks were made on the floor with the plumb-bob around the perimeter of the ring. The paper circle was aligned with these points, and the center was marked through the hole. This was used to




Figure 6.10: The experimental setup for the quenching measurement. The NaI(Tl) detector in the red box sits in a custom height-adjustable mount in a tripod. Behind it is the LiF target, the source of neutrons. The polyethylene collimator that blocks the beam of neutrons from passing directly from the LiF target to the neutron detectors is the white circle enclosed in the red box behind the NaI(Tl) detector. On either side of the NaI(Tl) detector are the neutron detectors, in the bright green boxes. The 5.1-cm detectors are shown on the right, while the 12.7-cm detectors are shown on the left. In this picture, the 12.7-cm detector positions are being marked using a plumb-bob as described in the text.

mark the central position of the stand. This same procedure was repeated when the detector was later moved to 91 cm from the LiF target.

However, it was later noticed that the tube that fit into the tripod had a diameter that was ~ 4 mm smaller than the inner diameter of the tripod. The set screw that holds the detector in place, therefore, offsets the detector by a small amount with respect to the tripod that was aligned with the floor. This offset was estimated to be ~ 2 mm. However, the detector had already been moved to its second location by the time this was discovered, so the orientation of the stand during the measurement is unknown. It was therefore conservatively estimated that the total uncertainty in the NaI(Tl) detector position was 1 cm.




Figure 6.11: The NaI(Tl) detector in its mount. The locations of the crystal and the PMT are marked in red and green, respectively. In the picture is Dr. Henning O. Beck, who provided equipment for the detector setup and helped with their installation.




Figure 6.12: (Left) The 5.1-cm Eljen 510-20x20-9/301 neutron detectors on their mounts. (Right) The 12.7-cm Eljen 510-50x50-1/301 detectors on their mounts. The location of the scintillator cell and the PMT are marked with the red and green arrows.

The neutron detectors were mounted on detector mounts provided by Fermilab. The stands consisted of a tripod with a PVC tube connected to the side. A cap with a tapped hole wherein a holder could be affixed was also available for the 5.1-cm detectors. The 12.7-cm detectors could sit directly in the tube. Because the floor was slightly uneven, the detectors were aligned with the beam in the vertical direction using the laser level.

The 5.1-cm detectors were taped to a holder that screwed into the cap for the PVC tube on the detector mounts, as shown in Figure 6.12. The cap was taped down with duct tape to secure it and prevent rotation. The stand was placed in the approximately correct location by eye. The detector was turned to face the laser that had previously been aligned with the proton beam. The z position of the detector was set to the laser height by aligning the laser with the machined lines on the front of the detector, which formed concentric circles about the center of the detector. The stand was then turned so the neutron detector would face the NaI(Tl) detector. A plumb-bob was held in the approximate center of the 5.1-cm detector while the position of the stand was adjusted to be directly over the mark on the floor, as shown in Figure 6.13. The legs of the detector mounts were taped down to the floor to secure their locations.



Figure 6.13: Centering the 5.1-cm detectors above the floor marks. In the picture is Shawn Westerdale and myself.

Using a ruler, the length of the cylinder containing the neutron detector was measured to be 4.45 cm. Including the flange, the length was measured to be 5.21 cm. The thickness of the face was assumed to be 0.16 cm. A line was marked 2.67 cm from the face of the detector along the curved wall, to represent the center of the scintillator volume. A string holding a plumb-bob was tied loosely around the marked line, so that the height of the plumb-bob could be adjusted. The looseness of the loop assured that the plumb-bob would hang from the center in the radial direction. The height was adjusted to just above the floor, and the actual location of the center of the 5.1-cm detector was marked using the plumb-bob (see Figure 6.14).



Figure 6.14: Marking the location of the 5.1-cm neutron detectors using a plumb-bob. The central location along the cylinder's z axis was measured and marked. A string with a plumb-bob on one end was tied around the detector loosely at this location, and the plumb-bob was allowed to hang from the center. The height was adjusted so the plumb bob was just above the floor, and the position of the detector was marked.

The placement of the 12.7-cm detectors was more complicated due to their center being directly above the PVC tubes. The PVC tubes were oriented above one of the legs of the tripod for stability. Because of this, the center of the detector could not be directly aligned with the floor. Because the outer diameter of the neutron detector flange was 15.2 cm, a mark on the floor was placed 7.6 cm from the centering mark in the direction of the NaI(Tl) detector, to mark the desired position of the side of the neutron detector. The detectors were placed on the mounts and temporarily secured with electrical tape. The heights were approximately adjusted by the laser (at the time, there was no mark at the vertical center of the detector).

The 12.7-cm detector mounts were placed so that the edge of the detector lined up with the 7.6-cm mark on the floor using a plumb-bob. The detector was rotated so that the center of the detector was along the neutron path by eye. One person held the plumb-bob

against the detector while another stood at the position where the NaI(Tl) detector would be located. This person directed the plumb-bob holder to center the plumb-bob in the azimuthal direction around the detector. Once the plumb-bob was in position, the detector stand was moved until the plumb-bob was aligned with the secondary mark on the floor. The legs of the stand were taped down to the floor for stability. The detectors were removed from their mounts for the night before the run to keep them out of harm's way. The next day, the center of the detector was marked in the vertical direction by measuring 6.35 cm (half the height of the scintillator cell) down from the top. The mount height was adjusted until this line was aligned with the laser, and the detectors were taped securely to the mount to keep them in place.

To mark the actual x - y locations of the detectors, a circular piece of stock paper the same diameter as the PVC tube holding the detector was made. A small hole was punched in the center of the paper. The plumb-bob was held against the tube at three locations around the diameter, and the corresponding positions on the floor were marked. The circular paper was aligned with these three points (see Figure 6.10), and the center of the circle was marked through the previously punched hole.

However, this location corresponded with the center of the bottom of the PVC pipe. The detector mounts were later found to be slightly non-vertical. After the measurement, but before moving the neutron detectors, the effect of this non-vertical alignment of the tube was measured. The laser level was removed from its alignment with the proton beam. For each detector, the laser level was placed on the floor some distance from the mount in two positions. The first position was to the side of the detector, with the vertical beam aligned parallel with the axis in the x direction that corresponded with the given detector, but passing through the center mark on the floor. For the second position, the T-square and meter stick were placed on the floor to measure a point the same distance away from the proton beam line as the detector, but with a different y position. In this case, the laser beam was aimed at the center mark.

For both the NaI(Tl) detector and the neutron detectors, the “center” of the detector in the x and y directions were marked by a person straddling the vertical beam and holding a ruler up to the detector. The distance between this mark and the laser beam was noted for all three 12.7-cm detectors. It was estimated that this process introduced a 1-cm uncertainty in the 12.7-cm detector positions. Because the 5.1-cm detectors were aligned with the floor directly and not through the mounts, no adjustment due to the tilt of the mounts was needed. For all of detectors, the difference between the initial desired positions and the actual positions were measured in the x and y directions. These differences were used as an adjustment to note the final positions of the detectors. The final distances between detectors were also measured along the floor after the detectors were removed. These values, along with the final positions of the detectors in x - y coordinates, are shown in Table 6.4.

The polyethylene collimator was centered in front of the LiF target in order to block neutrons from traveling directly to the neutron detectors. A movable cart was placed just in front of the LiF target. The collimator was too high on this cart. The height difference was measured between the central hole and the vertical location of the LiF target, and the collimator was cut such that the hole would be level with the LiF target. The collimator was raised to approximately the correct height by placing it on thin stainless steel sheets. These sheets were also used as shims to level the collimator. The laser level was used to center the collimator, as shown in Figure 6.15. The vertical offset of the collimator was -0.3 cm with respect to the LiF target. Also, the collimator was placed about 0.6 cm away from the target in the beam direction, since the target was receded into its holder.

Table 6.4: The final detector positions with their uncertainties. The x and y values are the positions from the LiF target, with y being along the beam axis. For each neutron detector, the recoil angle and energy and the distance from the NaI(Tl) detector are shown for position configuration 1 in the top row and position configuration 2 in the bottom row. The distance from the NaI(Tl) detector to the neutron detectors was a measured value. The uncertainty for the 12.7-cm detectors is higher than the 5.1-cm detectors because of the non-level mounts, as described in the text. The uncertainty in the NaI(Tl) detector positions was also higher than that of the 5.1-cm detectors because the set-screw in the mount offset the detector slightly from the mount.

The NaI(Tl) Detector			
Position configuration	x (cm)	y (cm)	Uncertainty (cm)
NaI(Tl) Pos 1	-0.2	50.0	1
NaI(Tl) Pos 2	-0.5	91.0	1

The Neutron Detectors						
Det. ID	x (cm)	y (cm)	Uncertainty (cm)	Recoil angle (deg)	Recoil energy (keV _{nr})	Dist. from NaI(Tl) (cm)
ND1	131.1	128.7	1	59.1	29.0	150.3
				74.2	43.0	134.6
ND2	98.8	162.7	1	41.3	15.0	148.8
				54.4	25.0	121.1
ND3	84.6	232.4	1	24.9	5.7	200.0
				31.1	8.8	163.3
ND4	-52.1	96.9	0.5	47.9	19.4	69.9
				84.0	51.5	51.8
ND5	-37.3	108.9	0.5	32.2	9.1	69.6
				64.6	33.2	40.7
ND6	-22.0	116.0	0.5	18.2	2.9	69.7
				41.1	14.3	32.7




Figure 6.15: Using the laser level to center the polyethylene collimator with respect to the LiF target. The collimator had to be cut on one side to be at the correct height.




Figure 6.16: Sample spectra from the neutron-detector calibrations for ND6. On the left is the calibration of the single-photoelectron response with mean 37.2 ADC counts and $\chi^2/\text{NDF} = 334/130$, performed by looking at the tails of large pulses. Here, the single-photoelectron response is ~ 45 ADC counts. On the right is the energy spectrum when exposed to a ^{241}Am source, scaled to units of photoelectrons. Here, the 60-keV γ peak can be seen at around 56 photoelectrons.

6.3.5 The electronics and data acquisition

Testing the detectors

Five 5.1-cm neutron detectors and three 12.7-cm detectors were available for the measurement. The detectors were tested for their gain⁸. Three of each type of detector were selected for the measurement. The results of tests of the neutron detectors are shown in Table 6.5.

⁸The preliminary tests of the neutron detectors were performed by Dr. Jingke Xu and Shawn Westerdale. The light-yield values reported here are from subsequent analysis done by me.

Table 6.5: Initial tests of the neutron detectors for the measurement. The detectors were evaluated for the separation of the single photoelectron (s.p.e.) peak from the pedestal (small fluctuations in the baseline causing false signals smaller than a single photoelectron) in ADC counts, the quantum efficiency, and the light yield in photoelectrons (p.e./keV_{ee}) to select the optimal detectors for the measurement. The light yield was measured with an ²⁴¹Am source.

Det. ID	Serial number	Voltage (V)	s.p.e. (counts)	Separation	Light yield (p.e./keV _{ee})
12.7-cm detectors					
ND1	5079-01-01a	1000	29	Good	1.58 ± 0.05
ND2	5079-01-02	1000	16		1.57 ± 0.20
ND3	5079-01-03	1000	28	Good	1.64 ± 0.05
5.1-cm detectors					
ND4	3900-01-02	1300	50	Fair	1.07 ± 0.03
ND5	3205-02-01	1360	31	Good	0.95 ± 0.04
ND6	3900-01-01	1460	41	Good	0.94 ± 0.11
N/A	3900-01-04	1500	NONE	NONE	N/A
N/A	3205-02-02	1400	70	Bad	N/A

The electronics

The electronics scheme was designed to digitize the signals from neutron-induced scintillation events in the NaI(Tl) and the neutron detectors, as well as the signal from the accelerator (hereafter referred to as the proton pulser)⁹. This would allow for an accurate determination of the event energy and timing, as well as the particle type through pulse-shape discrimination. The scintillation time for NaI(Tl) is $\tau \approx 230$ ns, while the scintillation time for the neutron detectors is much shorter. In order to capture the entirety of the scintillation signals, a region of 8 μ s was digitized with 2 μ s before the trigger. A series of logical operations were performed to select windows of time in which events happened in all three locations: the proton pulser, the NaI(Tl) detector, and one of the neutron detectors.

The electronics scheme used for the measurement is shown in Fig. 6.17. The modules themselves are shown in Figure 6.18, and the final configuration for the detector signals and their thresholds are shown in Table 6.6. The PMT signal from the NaI(Tl) detector was

⁹The initial design of the electronics was done by myself with input from Dr. Jingke Xu. This design was modified right before the run by Dr. Xu.

Table 6.6: Results from the final tests of the detectors before the measurement.

Det. ID	Voltage/max voltage (V)	Amplifier	Discriminator/level (mV)	Rate (kHz)	s.p.e. (mV)
NaI(Tl)	1350/1500	Gran Sasso FE	Phillips 711/11	0.5–1	~7
ND1	1200/1500	Phillips 779	Lecroy 621 AL/67	1	~25
ND2	1200/1500	Phillips 779	Phillips 711/40	1	~15
ND3	1150/1500	Phillips 779	Phillips 711/50.8	1	~20
ND4	1450/2000	Phillips 779	Phillips 711/57.6	1	~25–30
ND5	1550/2000	Phillips 779	Phillips 711/31.4	1	~20
ND6	1600/2000	Phillips 779	Phillips 711/12	0.8	~20
Proton	–	–	LeCroy 621 AL/423	–	–

amplified with a $\times 10$ front-end amplifier module developed at LNGS while the neutron-detector signals were sent to a Phillips 779 $\times 10$ amplifier. Amplified signals from both detectors were sent to a LeCroy 428F linear fan-out. One copy of each signal was digitized with a CAEN V1720E digitizer (12 bit, 250 MS/s, 2-V range), while the other copy was used to produce a trigger for the digitizer. The signals to be used for the trigger first went to low-threshold discriminators (Phillips 711 and LeCroy 621 AL), whose discrimination levels were set at 1.5 p.e. in order to reach a low energy threshold while reducing the random coincidence rate. The Phillips 711 had a minimum threshold of 10 mV, so most channels were sent there. Neutron Detector 1 and the proton pulser were sent to the LeCroy 621 AL, which had a minimum threshold of 30 mV. The discriminator outputs for the neutron detectors were combined by a NIM logical fan-in whose OR output was subsequently combined with the NaI(Tl) discriminator output in a logical AND (NIM 375L). The coincidence window for this AND logic was set to 400 ns to conservatively include the neutron-coincidence events with the longest time of flight. Subsequent triggers within the acquisition window were discarded. The signal from the pulsed proton beam was not used in the trigger, but was recorded for off-line analysis. Due to the degradation of the proton-pulser signal in long transmission lines, the signal was fed through a discriminator before being digitized.

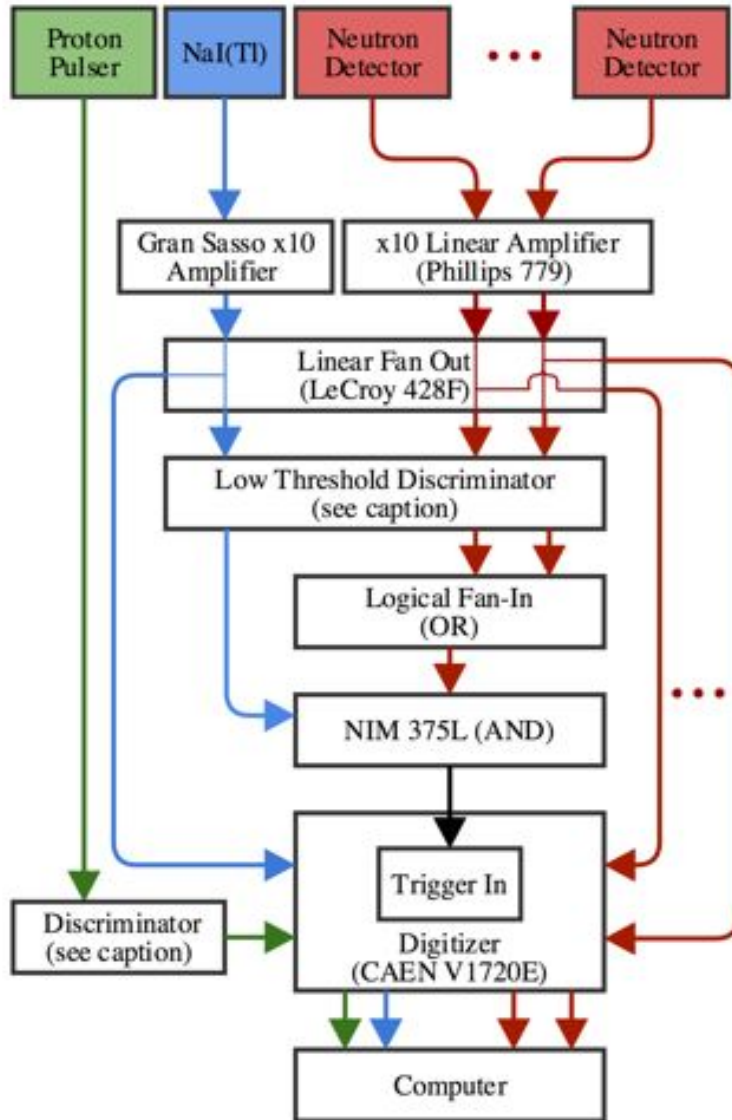


Figure 6.17: Electronics scheme for the measurement. The NaI(Tl) signal was fed through a front-end amplifier module developed at LNGS. The NaI(Tl) and neutron detectors with lower gain were passed through a Phillips 711 discriminator with a threshold of 10 mV, while other neutron detectors and the proton pulse selector were passed through the LeCroy 621 AL discriminator with a higher threshold equivalent to 1–2 photoelectrons in the neutron detectors.




Figure 6.18: (Top) The electronics modules used in the quenching measurement. (Bottom) Testing the electronics modules. The NaI(Tl) detector and the some of the neutron detectors can be seen on the table on the left.

A basic online analysis was performed to show the scintillation waveforms, the coincidence event rate, and the time-of-flight spectra for all neutron detectors. The waveform data from all channels were saved to disk to be used for offline analysis, as discussed in Chapter 7.

6.3.6 Measurement Summary

Data were collected in the first position configuration for 26 hours and in the second position configuration for 20 hours, giving approximately 1,000–4,000 coincidence events per energy. A list of the runs with the beam parameters is shown in Table 6.7. Each day, a calibration of the light yield was performed with ^{241}Am and ^{133}Ba sources placed ~ 15 cm from the NaI(Tl) detector. These are Runs 12–14, 27–28, and 57–58. In addition to this light-yield calibration, a beam-on calibration of the light yield was performed by observing the inelastic scattering of neutrons with ^{127}I , as described below. A final determination of the detector light yield was done in this manner in Runs 59–63. A measurement of the trigger efficiency in the NaI(Tl) detector was also performed, as described in Section 6.3.6.

Table 6.7: Run log for the SABRE Na quenching measurement. “Current” is the accelerator beam current. The proton energy was determined by a script written by Notre Dame to calculate the beam energy.

Run	Date/ time	Duration	Current (nA)	E_p (MeV)	Bunch ratio	Rate (Hz)	Notes
1	9/18 12:27	2:02:34	24.5	2.4395	1:6	70	
2	9/18 14:39	1:58:34	25	2.4400	1:6	85	
3	9/18 16:37	2:18:11	26	2.4400	1:6	90	
4	9/18 18:59	2:30:02	27	2.4400	1:6	85	
5	9/18 21:32	1:15:32	23.5	2.4400	1:6	70	
6	9/18 22:52	2:08:56	21	2.4400	1:6	65	
7	9/19 1:15	2:00:25	20	2.4400	1:6	60	
8	9/19 3:30	2:06:43	20	2.4400	1:6	55	
9	9/19 5:30	1:49:16	20	2.4400	1:6	55	
10	9/19 7:20	1:28:35	20	2.4400	1:6	55	
11	9/19 8:51	2:12:59	16	2.4400	1:8	45	
12	9/19 11:16	0:09:25	–	–	–	–	^{133}Ba

Continued on next page...

Run	Date/ time	Duration	Current (nA)	E_p (MeV)	Bunch ratio	Rate (Hz)	Notes
							calibration
13	9/19 11:33	0:08:02	–	–	–	–	²⁴¹ Am calibration
14	9/19 11:46	0:06:43	–	–	–	–	²⁴¹ Am calibration FE amp only
16	9/19 12:37	2:15:21	17	2.4400	1:8	70	
17	9/19 14:54	0:49:56	17	2.4400	1:8	70	
18	9/19 15:55	0:17:22	17	2.4410	1:8	65	
19	9/19 16:36	5:16:34	17	2.4410	1:8	25	Pos. Cfg. 2
20	9/19 21:56	2:17:22	20	2.4406	1:8	28	
21	9/20 00:10	2:26:16	19.5	2.4406	1:8	27	
22	9/20 02:37	1:44:35	20	2.4406	1:8	29	
23	9/20 04:21	2:07:55	20	2.4406	1:8	28	
24	9/20 06:30	3:26:41	20	2.4406	1:8	28	
25	9/20 09:57	2:48:19	20.5	2.4405	1:8	30	
26	9/20 12:46	2:10:59	21	2.4404	1:8	28	
27	9/20 15:08	0:11:47	–	–	–	–	²⁴¹ Am calibration
28	9/20 15:25	0:09:42	–	–	–	–	¹³³ Ba calibration
29	9/20 15:52	0:15:19	–	–	–	–	Trigger Eff.
30	9/20 16:09	0:27:41	–	–	–	–	Trigger Eff.
31	9/20 16:37	0:25:46	–	–	–	–	Trigger Eff.
32	9/20 17:02	0:25:25	–	–	–	–	Trigger Eff.
33	9/20 17:28	0:24:41	–	–	–	–	Trigger Eff.
34	9/20 17:53	0:20:42	–	–	–	–	Trigger Eff.
57	9/22 02:27	0:04:18	–	–	–	–	²⁴¹ Am calibration
58	9/22 02:34	0:03:31	–	–	–	–	¹³³ Ba calibration
59	9/22 02:55	0:06:36	15	2.440(?)	1:10	–	Beam-On calibration
60	9/22 03:02	0:14:13	15	2.440(?)	1:10	–	Beam-On calibration
61	9/22 03:20	0:16:32	15	2.440(?)	1:10	–	Beam-On calibration
62	9/22 03:37	0:25:13	15	2.440(?)	1:10	–	Beam-On calibration

Continued on next page...

Run	Date/ time	Duration	Current (nA)	E_p (MeV)	Bunch ratio	Rate (Hz)	Notes
63	9/22 04:02	0:29:17	15	2.440(?)	1:10	–	Beam-On calibration

Calibrating the light yield

Calibrations of the light yield of the NaI(Tl) detector were performed in two different ways. The first was in-beam, observing the 57.6 keV γ ray that comes from the first excited state of ^{127}I , which can be induced through inelastic scattering of the neutrons with an iodine nucleus; the energy of the recoiling nucleus is very small in comparison, as will be described in Section 7.2.3. The second was a series of separate source calibration runs performed by placing either a ^{133}Ba or ^{241}Am source near the detector and measuring the location of the full-energy peaks in the γ spectra.

The light yield of this detector was observed to decrease over the run due to some degradation of the crystal from moisture exposure, and possibly also a degradation of the optical coupling. The cause of this exposure is unknown, but one possibility is damage to the gas feedthrough during shipping, or a leak caused by the bulkhead feedthroughs on the enclosure. Another possibility is a shock during shipping that may have displaced the crystal slightly. In any case, the degradation of the crystal necessitated an accurate determination of the light yield of the detector as a function of time. The in-time measurement of the 57.6 keV peak was used to calibrate our detector performance in our analysis, as discussed in Section 7.2.3.

Measurement of the trigger efficiency

In the data acquisition setup, the acquisition of an event is dependent on the discriminator output for the NaI(Tl) signal; the signal can have a small amplitude, which may cause the discriminator not to fire. It is necessary to understand the efficiency of the discriminator's

response to small pulses¹⁰, hereafter referred to as the “trigger efficiency.” After the data were collected for the measurement of the quenching factor, a separate run was conducted to measure the trigger efficiency of the NaI(Tl) detector at low energies. Ideally, to perform such a measurement, the NaI(Tl) detector signal would be observed at random times along with the output of the discriminator, but this method would produce a large amount of empty signal. In order to increase the chance that an event would be observed in the NaI(Tl) detector, sources were placed near the detector and a second detector was used as a trigger in the hopes that coincident signals would also occur in the NaI(Tl) detector.

The setup is shown in Fig. 6.19. A Bicron 76-mm NaI(Tl) detector was set up at a reasonable distance away from the 25-mm NaI(Tl) detector. A ^{22}Na source was placed directly between the two detectors, which produced back-to-back 511-keV γ rays. ^{133}Ba and ^{241}Am sources were also put near the 25-mm detector opposite the 76-mm detector to increase the event rate. The 76-mm detector was used as a trigger for the measurement with a high threshold. The 25-mm detector signal from the amplifier and, separately, from the discriminator were fed into the digitizer to measure the efficiency of the trigger for low-energy recoils. The correction to the quenching-factor measurement as a result of this trigger-efficiency measurement is discussed in Section 7.4.2.

¹⁰Dr. Hugh Lippincott recommended this measurement be conducted.

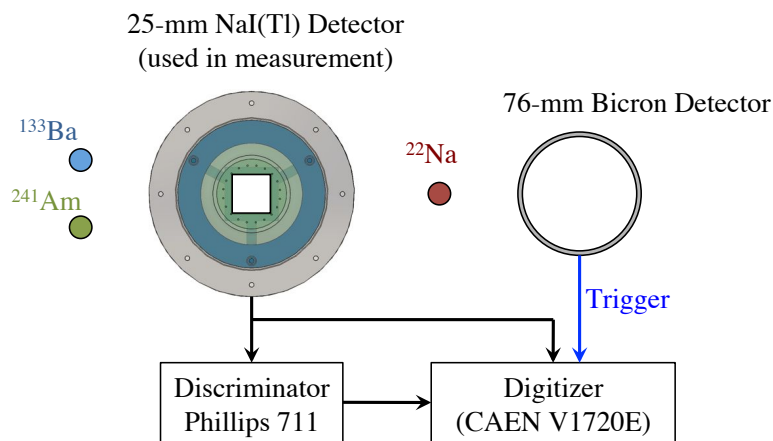


Figure 6.19: Setup for the trigger-efficiency measurement of the 25-mm NaI(Tl) detector. The analog signal from the NaI(Tl) detector was compared to the output of the discriminator to determine the response as a function of pulse amplitude. A ^{22}Na source was placed between the NaI(Tl) detector and a separate 76-mm Bicron NaI(Tl) detector so that the back-to-back 511-keV $_{nr}$ γ rays could be detected by both detectors. The signal was fed both directly into the digitizer and also through the discriminator, while the Bicron detector was used as a trigger to reduce the acquisition of data with no pulses. Additional sources were placed near the 25-mm detector to increase the event rate.

Chapter 7

Measurement of the Na

Nuclear-Recoil Quenching Factor:

Analysis and Results

7.1 The data acquisition

The scintillation lifetime of NaI(Tl) is 200–300 ns. The voltage waveforms from the PMTs in the NaI(Tl) and neutron detectors were digitized in 4-ns samples in an 8- μ s window for each trigger. Triggers that occurred within an acquisition window after a previous trigger were ignored. The waveform from the proton pulse selector was fed through a discriminator, whose signal was digitized as well. The waveforms from each channel were stored in binary format for later analysis.

The data acquisition software used for the measurement, called daqman, was written by Dr. Ben Loer for the Darkside-10 detector and modified by Dr. Jingke Xu. This software analyzes the pulses in the different channels and produces relevant metadata about each event. For this analysis, the pulse integral corresponding with the event energy was of particular interest. The peak amplitude of the pulses was also determined for an evaluation of

the trigger efficiency of the discriminators. The timing of the pulses and the pulse shapes were also analyzed for the selection of neutron-induced nuclear-recoil events among background¹.

7.1.1 Features of the waveforms

A typical waveform of a neutron-induced nuclear-recoil event is shown in Figure 7.1. Each event contains eight waveforms: one from the NaI(Tl) detector, six from the neutron detectors, and one from the proton pulse selector. Each channel has a baseline with some amount of variation and drift. Due to the discrete nature of scintillation events, which can consist of a small number of single photoelectrons detected over a long scintillation time, the NaI(Tl) and neutron channels can contain a number of pulses that comprise a single scintillation event. The signal is also complicated by some electronic ringing that was observed, as can be seen in Figure 7.1. The origins of this ringing are unknown, but their effects on the calculated pulse integral and the baseline were mitigated in the analysis.

7.1.2 The baseline algorithm

An accurate determination of the waveform baseline is necessary to determine the other quantities of interest. The baseline of each acquisition window was established using an algorithm that identified possible pulses and areas of ringing and excluded them. After excluding these regions, a moving average was used to establish the baseline, with a linear interpolation in pulse and ringing regions, as can be seen in Figure 7.1.

The stability of the baseline was important for calculating the energies of events, since any integral over the baseline should be 0 if the baseline is interpreted correctly. If a “non-zero” cumulative integral over the baseline is observed, the estimate of the event energy will

¹An initial analysis to determine the quenching factors, including a Monte Carlo fit, was done by me while Dr. Jingke Xu worked on pulse-shape analysis and, with Shawn Westerdale, on a neutron energy/angular spectrum for the simulation input to refine the analysis. While I was on maternity leave, Dr. Xu performed a complete analysis independently (with some reference to my Monte Carlo code) for publication in [1]. Upon my return, I performed a complete, independent analysis of the quenching factors based on my prior work and following the approach of Dr. Xu, as reported here.

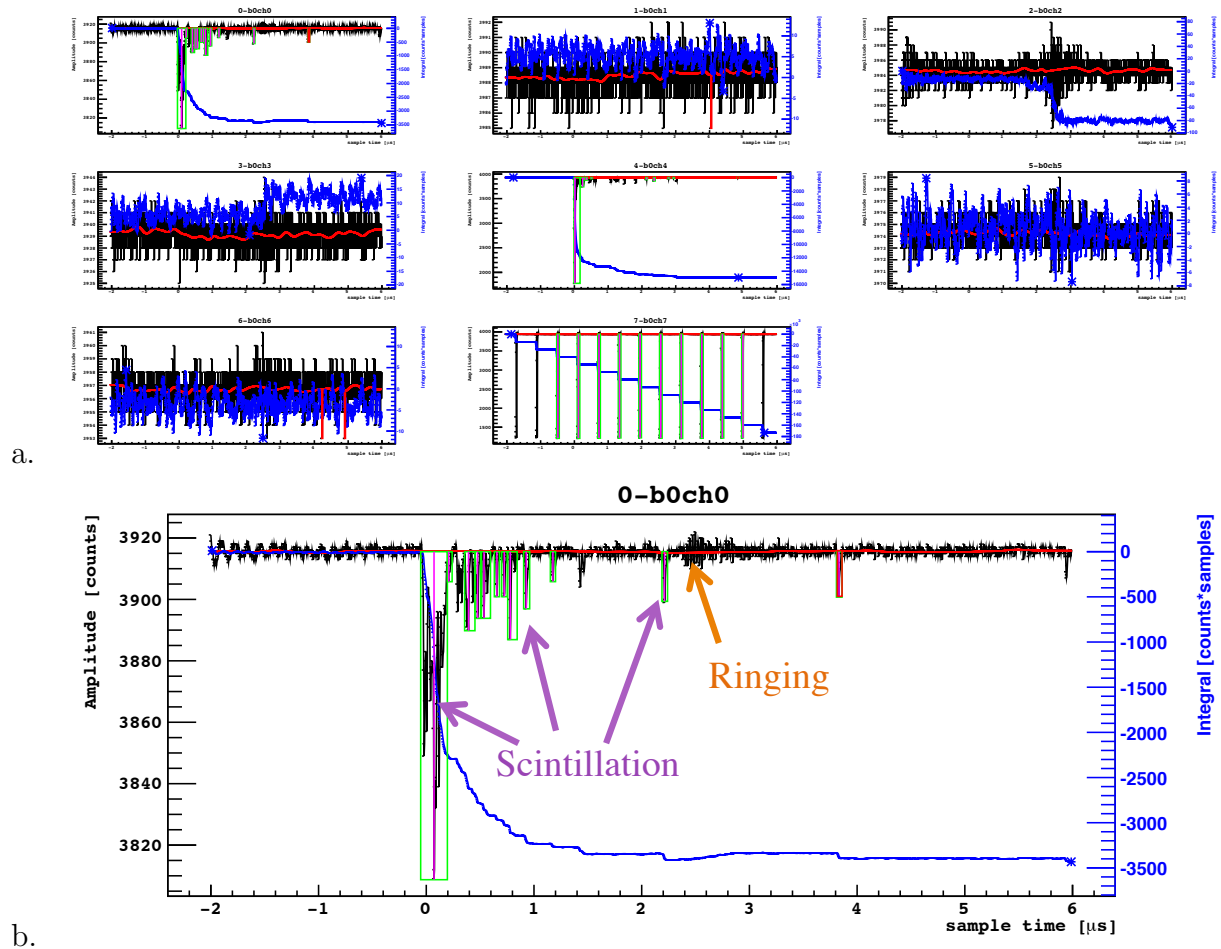


Figure 7.1: (a) Typical waveforms for a 20 keV_{nr} nuclear-recoil event. Signals are present in the NaI(Tl) detector (Top left) and Neutron Detector 4 (Center). The NaI(Tl) waveform is shown in (b). This event contains both scintillation pulses (marked with purple arrows) and electronic ringing (marked with an orange arrow), which can be seen around the $2.5\text{-}\mu\text{s}$ mark. The extrapolated baseline is shown in red, while the waveform baseline-subtracted integral is shown in blue. Pulses are outlined with green boxes, and the peak is highlighted with a magenta-colored line.

be skewed. The baseline algorithm was tested such that the integral over the baseline did not vary by more than 10–20 count-samples. The trigger threshold for the NaI(Tl) detector was around 20 ADC counts, so a variation in the baseline of less than 20 count-samples over the acquisition window would be less than the smallest possible scintillation pulse. Typical pulse integrals in the energy range of interest are $\sim 2,000$ – $10,000$ count-samples, so the drift in the baseline can contribute less than 1% to the energy uncertainty.

7.1.3 Finding pulses

The pulse-finding algorithm incorporates a discriminator search. Pulses begin and end when the baseline-subtracted waveform crosses a threshold which is a certain multiple of the baseline variance. Pulses that overlap or are too close together are separated by finding the point between the peaks that is closest to the baseline. Pulses with a small width are combined. The software also requires that pulses be a certain width before recording them. The thresholds for these operations are set in a configuration file.

The Pulse object contains relevant data about the pulse, including the maximum baseline-subtracted amplitude, the integral, and the timing information of the pulse. The most important timing variable is the `half_max_time` variable ($t_{1/2}$), which is defined as the time at which the signal first drops below half the peak amplitude². This value was used as the official “start time” of the pulse for purposes of time-of-flight calculations.

7.1.4 Creating event objects: the TOF module

The TOF (short for Time-of-Flight) module was custom-written for this measurement by Dr. Jingke Xu to take a list of pulses for each event and create an event object with useful metadata about each event. The liquid-scintillator neutron-detector channels, the NaI(Tl)-detector channel, and the proton pulse selector were each processed differently.

²The pulses have a negative amplitude.

Because the trigger occurs when the second coincident channel (which should be the neutron detectors in a true neutron-recoil event) fires, the neutron-detector event should begin near time 0. Pulses in the neutron channel were searched for between -0.1 and $0.1 \mu\text{s}$. If a pulse was found in this region, all subsequent pulses for the next $4 \mu\text{s}$ were combined with it to produce a scintillation event. A new TOF object with type Pulse was created, its integral being the sum of all the pulse integrals in the $4\text{-}\mu\text{s}$ window of that channel. $t_{1/2}$ of the TOF object is the half-max-time of the first pulse found within the search window. NaI(Tl) events were processed similarly, except that the search window began at $-1 \mu\text{s}$.

The proton pulse selector, which is periodic, has an unknown but constant offset in time from the NaI(Tl) and neutron channels. Therefore, a pulse needed to be chosen as the one that was associated with the events in a consistent way. Choosing the nearest pulse to $t_{1/2} = 0$ or the nearest pulse before time 0 arbitrarily splits the events in half. Therefore, an offset variable was created to choose the nearest pulse to a certain point in time, so that neutron- and γ -induced events could be clustered together in a time-of-flight analysis, as discussed in Section 7.2.2.

7.2 Creating the energy spectrum

For each of the measured recoil angles, an electron-equivalent energy spectrum of neutron-induced nuclear-recoil events was created with the measured data. This was done to find the electron-equivalent energy of the nuclear recoils with known recoil energy. This section describes the process by which these spectra were generated.

7.2.1 Finding the single photoelectron

In order to quantify the light detected in a nuclear-recoil event, the response of the NaI(Tl) detector to single photoelectrons (s.p.e.) should be well understood. The s.p.e. spectrum is created by searching for single photoelectrons in the tails of larger pulses using a low-

Single photoelectron spectrum for Run 1




Figure 7.2: The single photoelectron (s.p.e.) spectrum for Run 1. A double-gaussian was fit to the spectrum to determine the s.p.e. mean. Here, p_0 and p_2 are the amplitude and width, respectively, of the pedestal gaussian (the mean, p_1 , was fixed at 0). p_3 , p_4 , and p_5 , are the s.p.e. amplitude, mean, and width, respectively.

threshold discriminator search. The sharp scintillation signals that can be seen in Figure 7.1 at later times could be single photoelectrons of the kind used to generate the s.p.e. spectra. Such a search will find single photoelectrons along with larger multiples and large variations in the baseline. A spectrum of this kind should therefore contain three features. The most noticeable will be a spread around the mean of the s.p.e., and a “pedestal” centered around 0, which is a result of the normal variation in the baseline. There should also be spreads for pulses with n photoelectrons, where $n > 1$, but these features should be small if the integrating window is far enough away from the main pulse.

The daqman software was used to produce a spectrum of single photoelectron pulses, as shown in Figure 7.2. A fit was performed with a gaussian centered at 0 for the pedestal and a second gaussian for the s.p.e. response. This fit was performed for each run separately. The s.p.e. was ~ 39.20 ADC counts with a variance of 0.08.

7.2.2 Selecting events

Because of the low discriminator threshold, the above-ground location, and the characteristics of the beam, the measurement was subject to a large background, especially at low energies. Though the neutron-induced recoils were clearly visible even without cuts, the lower-energy recoils were obscured by low-energy backgrounds. Several cuts were used in the creation of the recoil-energy spectra to clean the signal.

In the generation of the energy spectrum for a given neutron detector, the first requirement was that the TOF module found an event in all three detectors: the NaI(Tl), the neutron detector of interest, and the proton pulse selector. If the event passed this basic cut, it was analyzed for additional characteristics. It was also required that for each detector (both the NaI(Tl) and neutron detectors) a clean baseline of $0.2\ \mu\text{s}$ preceded the triggering pulse, to assure that the triggering pulse was not part of a larger, preceding event.

Due to the well-known energy of the neutron before and after its interaction in the NaI(Tl) detector, the timing information of the events could be used to reject random-coincidence and γ backgrounds. Because the proton-pulse signal had an undetermined time offset from the NaI(Tl) and neutron detectors, the time difference between the NaI(Tl) detector and the proton pulse for neutron events had to be resolved by comparing $t_{1/2}$ of each event to $t_{1/2}$ of a prominent peak due to the essentially instantaneous γ background, shown in the red box in Figure 7.3. The time difference between the NaI(Tl) detector and the proton pulse was approximately 43 ns in the first position configuration and 86 ns in the second. The time difference between the neutron detector and the proton pulse was dependent on the neutron detector, because they were placed at different distances, but typically fell in the 100–200 ns range.

To establish the cuts, for each detector a spectrum of the timing information was made with the time difference between the NaI(Tl) detector and the proton pulse on the x axis (TOF1) and the time difference between the neutron detector and the proton pulse on the y axis (TOF2). A typical time-of-flight plot is shown in Figure 7.3. The brightest region in

the bottom left (shown in a red square) is the background of coincident γ s coming from the proton pulser, as well as from single γ rays that come from the proton pulser and interact in both detectors. The region in the blue box contains the neutron events. The horizontal line of events passing through the γ peak corresponds with γ rays that travel directly to the neutron detector from the LiF target in coincidence with random background in the NaI(Tl) detector. Similarly, the vertical line through the γ peak is caused by γ rays traveling directly to the NaI(Tl) in coincidence with random events in the neutron detector. The vertical line through the neutron region is caused by neutrons from the LiF hitting the NaI(Tl) and an unrelated random event occurring in the neutron detector. The diagonal line passing through the γ peak is caused by events that happen simultaneously in the neutron and NaI(Tl) detectors not caused by the proton beam, but rather by independent events like cosmic-ray showers.

The neutron region on these plots was confirmed with a calculation of the expected time of flight; neutrons with an energy of ~ 690 keV should traverse the distance from the LiF target to the NaI(Tl) detector in ~ 40 ns in the first configuration and ~ 80 ns in the second configuration. Subtracting the time of flight of the neutron peak and the γ peak for TOF1 yields transit times that match the calculated ones. Similarly, TOF2 matches the expected values based on the expected neutron energy. Conservative cuts were used to maximize the acceptance of neutron events. The time-of-flight cuts used for the different detectors are shown in Table 7.1.

The shape of the scintillation pulses in NaI(Tl) and in organic liquid scintillators are different for different types of events. Pulse-shape discrimination can therefore be used to separate neutron elastic-scattering events from other backgrounds. PSD was applied with conservative cuts in both the neutron detectors and the NaI(Tl) detector. For both detectors, the F50 parameter, defined as the fraction of the pulse integral that occurred within the first 50 ns of the pulse, was used as the PSD parameter.




Figure 7.3: Typical time-of-flight plot. The x axis, TOF1, is the time difference between the NaI(Tl) detector and the proton pulse, while the y axis, TOF2, is the time difference between the neutron detector and the proton pulse. Both values have an arbitrary offset but can be re-centered by observing the γ peak, the bright region in the red square. The neutron events are in the blue square. A detailed description of the features of this plot are in the text.

Table 7.1: Time-of-flight cuts for the different measurements of the Na quenching factor. The time of flight for the NaI(Tl) detector is TOF1, while the time of flight for the neutron detectors is TOF2. The timing information for the γ peak is also shown, which was essentially the time at which the proton pulse was fired.

Detector ID	TOF position 1 (μs)		TOF position 2 (μs)	
	Min	Max	Min	Max
NaI(Tl)	-0.15	-0.09	-0.11	-0.05
1	0	0.04	0.02	0.07
2	0	0.05	0.01	0.05
3	0.04	0.09	0.04	0.09
4	-1.0	-0.06	-0.08	-0.03
5	-1.0	-0.06	-0.09	-0.04
6	-1.0	-0.05	-0.09	-0.06
Gamma peak time	TOF1= -0.176, TOF2= -0.16			

The F50 parameter is shown against the energy deposited in the neutron detector for one of the detectors in Figure 7.4. TOF cuts were applied to this figure so that the neutron elastic-scattering events, shown in the red box, could be seen more easily. Events with higher F50 are faster events; because nuclear recoils from neutrons have a slower scintillation profile than electron recoils from γ s (see Appendix A), the γ band appears closer to F50=1 than the neutron events. The low-energy signals with F50 around 1 (the single red pixel) are due to Cherenkov radiation, which are not true scintillation events, but are near-instantaneous flashes of light from relativistic charged particles traveling through the medium.

NaI(Tl) has a slow scintillation time, ~ 230 ns. Therefore, the F50 parameter can also be used to reject fast Cherenkov radiation, which would have a F50 parameter close to 1^3 . Because of this, events with an F50 parameter higher than 0.45 in the NaI(Tl) detector were rejected, as can be seen in Figure 7.5. The nuclear- and electron-recoil bands were too close to separate without losing many of the neutron elastic-scattering events. The cuts used on the F50 parameter for both the NaI(Tl) and neutron detectors are summarized in Table 7.2.

After the time-of-flight and PSD cuts, cuts were made on the energy of the events in the neutron detectors to isolate the neutron events as shown in Figure 7.4. The energy cuts are summarized in Table 7.2.

³The pulse-finding algorithm gives some buffer to the beginning of the pulse, so some pulses with a smaller width than 50 ns will have an F50 parameter slightly smaller than 1, as can be seen in Figure 7.5. This means that fast pulses may produce a broad spread in F50 close to F50=1.




Figure 7.4: Pulse-shape discrimination in Neutron Detector 3, position configuration 1 (nominal recoil energy 15 keV_{nr}) after the application of TOF cuts. The F50 parameter is the fraction of the pulse that occurs within the first 50 ns. The cuts on this parameter and the energy in the neutron detector are shown in the red box. The electron recoil band sits slightly higher in F50. Cherenkov events are characterized by the single red pixel at F50=1.

Table 7.2: Cuts used on the energy spectra after the application of the time-of-flight cut. The NaI(Tl) detector had only a PSD cut, while the neutron detectors had both a PSD cut and an energy cut.

Det.	Detector size (cm)	Energy min (p.e.)	F50 cut	
			Min	Max
NaI(Tl)	2.5	–	0	0.45
1	12.7	40	0.55	0.8
2	12.7	40	0.55	0.8
3	12.7	50	0.55	0.82
4	5.1	10	0.5	0.84
5	5.1	10	0.5	0.85
6	5.1	4	0.6	0.90




Figure 7.5: Pulse-shape discrimination in the NaI(Tl) detector applying no cuts. The F50 parameter is the fraction of the pulse that occurs within the first 50 ns. Events above the red line, which include Cherenkov radiation (the narrow peak at $F50=1$) and any other non-scintillation light, were rejected. The neutron-induced nuclear recoils are indicated with the pink arrow. The electron recoil band is indicated with the orange arrow, which points, in particular, to the 57.6-keV_{ee} γ from the ^{127}I first excited state.




Figure 7.6: Time of flight from the proton pulser to the NaI(Tl) detector vs. energy for all detectors. The blue box corresponds with neutron elastic-scattering events, while the red box corresponds with inelastic-scattering events that excite ^{127}I to the first excited state of 57.6 keV. γ rays coming from the LiF target when being struck by the proton beam are indicated with the orange arrow.

7.2.3 Making the light-yield correction

As discussed in Section 6.3.6, two different measures of the light yield were performed. The first was an in-run calibration by observing the inelastic scattering of neutrons on ^{127}I , which produces a 57.6-keV γ ray, and the second a series of separate source runs performed daily with ^{133}Ba and ^{241}Am sources.

Incoming neutrons can scatter inelastically off an ^{127}I nucleus, causing it to excite to the 57.6-keV state and also to recoil. The 57.6 keV is released as a γ ray when the nucleus decays to the ground state. Because the ^{127}I nucleus is very massive compared with ^{23}Na , the recoil energy of the nucleus is very small for the same neutron scattering angle compared with a Na recoil. For example, for a 690-keV neutron, a 33-degree scatter will induce a nuclear recoil of 10 keV_{nr} for ^{23}Na , but 1.7 keV_{nr} for ^{127}I . Therefore, the skewing of the peak due to the kinetic energy of the iodine recoil itself is small. This effect was estimated to be at or below 1% by a Geant4 simulation. The average iodine recoil energy was 4.25 keV_{nr} .




Figure 7.7: Time of flight from the proton pulser to the neutron detector (TOF2) vs. energy observed in the NaI(Tl) detector for coincidences with neutron detector 1 (a nominal 30 keV_{nr} recoil). The blue circle is around the neutron-induced elastic-scattering events. The blue band corresponds to the TOF2 cut isolating neutron events. The red band represents the 57.6-keV_{ee} γ events from inelastic scattering of a neutron off ^{127}I . Though the intersection of the blue and red bands in the figure represent 57.6-keV γ events wherein the neutron goes on to hit the given detector, all events with TOF1 matching that of a neutron scattering event in the NaI(Tl) were accepted for the light-yield calibration. Figure made by Dr. Jingke Xu for [1] and modified for this work.

Assuming a quenching factor of 0.09 for iodine, as reported by DAMA, the recoil of the nucleus should contribute 0.4 keV_{ee} to the calibration. The detection of the 57.6-keV γ ray, therefore, provides a convenient mono-energetic signal for the purposes of calibrating the light-response to electron recoils.

The in-run observation of the ^{127}I excited state was done on a run-by-run basis. For each run, the energy was plotted for all NaI(Tl) events that had the correct time of flight from the LiF target, which corresponds to the red square/band in Figures 7.6 and 7.7, respectively. A gaussian plus flat background was fit to the spectrum, as shown in Figure 7.8 for Run 1. The mean of this gaussian was taken to correspond with the light yield in p.e. for the time at the middle of the run. A linear fit of this data was taken to obtain the light yield as a




Figure 7.8: Light-yield calibration fit for Run 1 using the ^{127}I first excited state. Here, a gaussian+constant is fit to the energy spectrum. p_0 is the gaussian amplitude, p_1 is the gaussian mean, p_2 is the gaussian spread, and p_3 is the constant.




Figure 7.9: Light yield of the NaI(Tl) detector as a function of time in hours, based on the on-beam light-yield calibration from observing neutron inelastic-scattering events with ^{127}I .

function of time, as shown in Figure 7.9. Higher-order polynomial fits were tried, but the higher-order coefficients were too small to be significant, so the linear fit was used.

The source calibration was done by hanging a ^{241}Am or ^{133}Ba source about 15 cm away from the NaI(Tl) detector at the level of the crystal. Simulated spectra for these runs were generated with Geant4. A fit was created by smearing the simulated spectra with a gaussian and scaling it to match the data. Because the NaI(Tl) scintillation response is non-linear

with energy, the simulated spectra were not fit over the whole energy range, but rather at the lowest-energy peak that could be resolved.

A few-percent systematic difference in the light yield between the two measurements (between 3% and 5% for ^{133}Ba) was observed; the source calibrations showed a lower light yield than the real-time calibration with ^{127}I . One potential reason for this difference stems from the position distribution of the scintillation events; the γ s from the first excited state of ^{127}I are evenly distributed throughout the crystal, whereas the γ s from the external sources will interact within a few mm of the crystal edge. This effect can cause a systematic decrease in the light yield, as the light yield may be position dependent, especially since some cracks in the crystal remained after the polishing before the run.

The recoil-energy spectra were corrected for the changing light yield event-by-event in the offline analysis. As the energy spectrum was filled, the integral of the pulse was divided by the s.p.e. mean for the run, as well as the light yield at the time of the event according to the calibration fit, to produce an energy spectrum in keV_{ee} .


7.2.4 The nuclear recoil energy spectra

After the data was processed, the cuts were applied, and the pulse integral was scaled with the s.p.e. and light-yield correction, the energy spectra were created for all 12 nuclear recoil angles. The spectra are shown in Figures 7.10 and 7.11 for position configurations 1 and 2, respectively.


7.3 Simulating the energy spectra

In principle, the recoil energy of the events in each detector can be calculated from kinematics. However, the finite detector sizes and the spread in the incoming neutron energy complicate matters and introduce some uncertainty in the mean energy. Therefore, it was necessary to


Neutron Detector 1, 30 keV




Neutron Detector 2, 15 keV




Neutron Detector 3, 6 keV



Neutron Detector 4, 20 keV



Neutron Detector 5, 10 keV



Neutron Detector 6, 3 keV




Figure 7.10: Experimental energy spectra for all 6 recoil angles in position configuration 1. The detector ID and nominal energy in keV_{nr} is listed at the top of each plot.

simulate the neutrons' paths through the detectors to create an expected recoil spectrum that could be compared to the observed one.

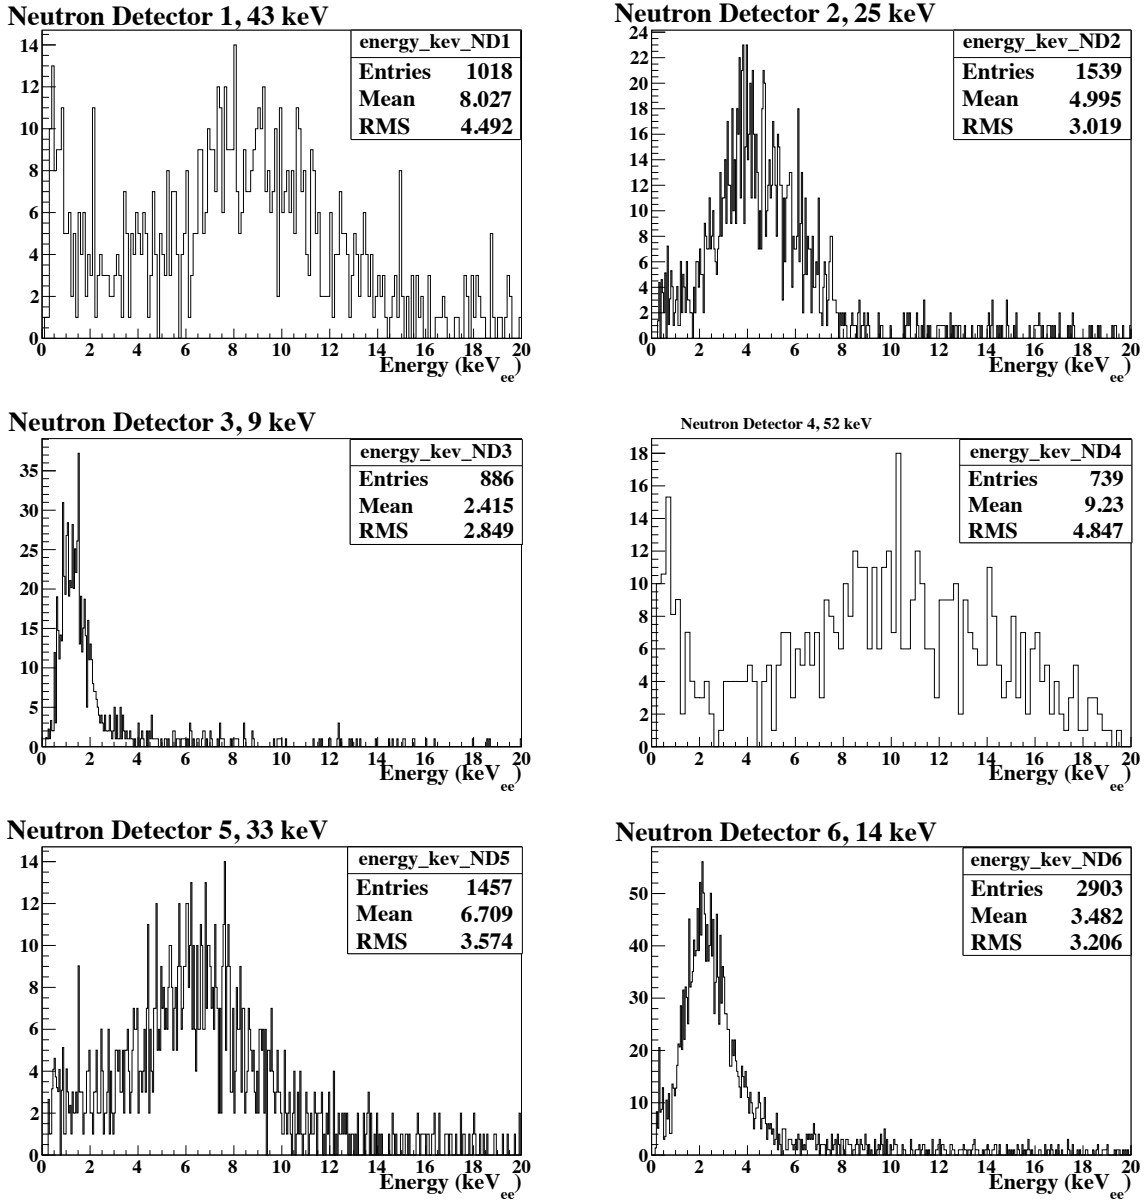


Figure 7.11: Energy spectra for all 6 recoil angles in position configuration 2. The detector ID and nominal energy in keV_{nr} is listed at the top of each plot.

7.3.1 Generating the incoming neutron energy and angular spectra

The Geant4 package was used to simulate the incoming neutron energy spectrum. While the simulation described in Section 6.2.3 provided an approximate energy spectrum of the

neutrons leaving the LiF target at a 0° angle from the proton beam, this simulation provided a more sophisticated, angular-dependent estimation of the neutron spectrum.

In order to characterize the behavior of protons moving through the LiF, 2.44-MeV protons were sent into the LiF target with a 0.52-mg/cm^2 thickness. The protons were tagged each time they traveled a small but defined distance through the LiF. This provides a more accurate estimation of the proton's chance to interact than the approach used in Section 6.2.3. The angle between the proton's velocity and the beam line as well as the proton energy at this position were recorded. For a large sampling of scattering angles⁴, a large number of neutrons were generated in a cone about the proton's velocity vector. The neutron was assumed to escape the LiF without interacting, so the neutron's outgoing energy was calculated based on the proton energy and the scattering angle using the relativistic kinematic calculation in Section 6.2.4. The neutron energies were entered into a histogram, weighted with the ${}^7\text{Li}(p,n){}^7\text{Be}$ differential cross sections digitized from [135]. The resulting energy and angular spectrum was saved in a histogram and used as input for the main simulation. This spectrum is shown in Figure 7.12.

A simulation was conducted by generating neutrons with this spectrum in front of the polyethylene collimator. This was to reduce the simulation time in the main simulation by determining a maximum-starting-angle cut without biasing the simulation. Neutrons were observed at the distance of the NaI(Tl) detector to determine how many neutrons would make it past the collimator to potentially interact in the detector as a function of the starting angle. It was found that less than 2% of events in the crystal were caused by neutrons with a starting angle greater than 6° . Of these, many do not pass a time-of-flight cut because they interact in the PMT coupled to the NaI(Tl) crystal, losing much of their energy before scattering in the crystal.

⁴The angle between the initial proton velocity vector and the final neutron velocity vector.




Figure 7.12: Simulated energy and angular spectrum of the neutrons leaving the LiF target. The z -axis scale is counts, proportional to the probability of a neutron emission at that energy and angle.

7.3.2 The simulation

The simulation used the Geant 4.9.6 package. A framework and user interface were written by Dr. Jingke Xu. The framework allows for the generation of neutrons with the energy and angular spectra simulated previously. The neutrons are tracked as they travel through the detector setup. Events in which energy is deposited in both the crystal and one of the neutron detectors are recorded with information about each interaction including the particle type, the incoming energy, and the deposited energy.

The detector geometry, shown in Figure 7.13, was constructed from scratch using the G4NistManager for materials. The detectors and the collimator were simulated. The NaI(Tl) and neutron detectors were simulated as completely as possible using drawings for the neutron detectors and our own measurements for all the parts of the NaI(Tl) detector. The neutrons were generated with a maximum angle cutoff of 6° , as determined by the collimator simulation.




Figure 7.13: The geometry for the simulation of the Na quenching measurement in the first position configuration. The polyethylene collimator is shown in white on the right. The NaI(Tl) detector is shown in its enclosure (red). The 5.1-cm detectors are shown near the bottom of the figure, while the 12.7-cm detectors are at the top. The scintillator volume in the neutron detectors is in blue, while a gas bubble is in green.

7.4 Evaluating the quenching factors

Fitting simulated spectra accurately requires that the statistics of the simulated spectra be at least 10 times higher than that of the data [140], so ~ 50 – $100,000$ events were simulated for each detector. The simulated spectra were created with the intent to fit them to the data with the quenching factor as a fit parameter. For the majority of the data energy spectra, the peaks were well separated from noise with the TOF and PSD cuts, but some of the lower-energy events were affected by a loss in trigger efficiency. This loss had the effect of skewing the energy spectrum upwards. Therefore, two different techniques were needed to evaluate the quenching factors for the higher-energy recoils and those affected by the trigger-efficiency loss. The two techniques are described in the following sections.

7.4.1 Standard method

For the high-energy events, the general approach was to create a simulated energy spectrum and use it as a fit function to the energy spectrum of the data.

The main challenge was to establish which simulated events to include in the fit. The simulation included energy depositions from neutrons, ^{23}Na and ^{127}I nuclei, and γ and β backgrounds. ^{127}I recoils are subject to a stronger quenching effect than ^{23}Na , so the electron-equivalent energy for these events is low. They were not observed in the data, and so were not included in the simulated spectrum. Electron-recoil events from β s and γ s have no quenching effect, and so could be combined with the ^{23}Na events with appropriate scaling. However, this combination would have to be done on an event-by-event basis, reevaluated with every potential quenching factor, and would require too much computational time to implement. It was instead established that electron-recoil events constituted a very small, low-energy background that would have a negligible effect on the simulated spectrum.

Finally, the simulation also produced multiple-scattering events. It was established that the multiple-scattering background was smooth, featureless, and small compared with the rate for single-scattering events, and so these were ignored in the fit. It was also established that fitting while including these events did not significantly change the resultant location of the simulated energy peak. The single-scattering and multiple-scattering spectra are shown in Figure 7.14.

The energy spectra, once created, were fit to the data using three fit parameters. These parameters are referred as $p0$, $p1$, and $p2$, respectively in Figure 7.15. The first, $p0$, was the quenching factor. The second, $p1$, was an arbitrary amplitude adjustment to match the height of the two spectra. The third, $p2$, was the spread of a gaussian convolved with the simulated spectrum representing photon statistics adjusted by the Fano-factor of NaI. The Fano-factor is a material-dependent parameter describing the difference in the energy resolution from predictions due to pure Poisson statistics, which is due to the discrete nature of the electron levels. Because the Fano factor is not well-known in NaI(Tl), it was used as




Figure 7.14: (Left) Simulated single and single+multiple-scattering energy spectra for one scattering angle, shown in red and black, respectively. The ^{23}Na nuclear-recoil events are contained in the peak around 15keV_{nr} . The iodine recoils are shown in the sharp peak around $3\text{--}4\text{keV}_{nr}$. (Right) Spectrum of simulated events that are excluded by using single-scattering events only, equivalent to the subtraction of the two curves in the left figure. These events include the peak from iodine scattering and the multiple-scattering background. The multiple-scattering background was found to be featureless and smooth over the energy region of interest. It was also small compared with the single-scattering peak, and was therefore not used in the fit.

a fit parameter. The fits for the scattering angles that were not affected by trigger-efficiency loss are shown in the figure.

The error of the fit was evaluated by varying the fit range and observing the variation in the quenching parameter. In addition to this systematic error, an error of 1.5% was added for the γ calibration of the light yield, and 3–12% for the positions of the detectors. The errors are described in detail in Section 7.4.3. The results of these fits are shown, along with the low-energy recoils, in Tables 7.3 and 7.4, respectively.

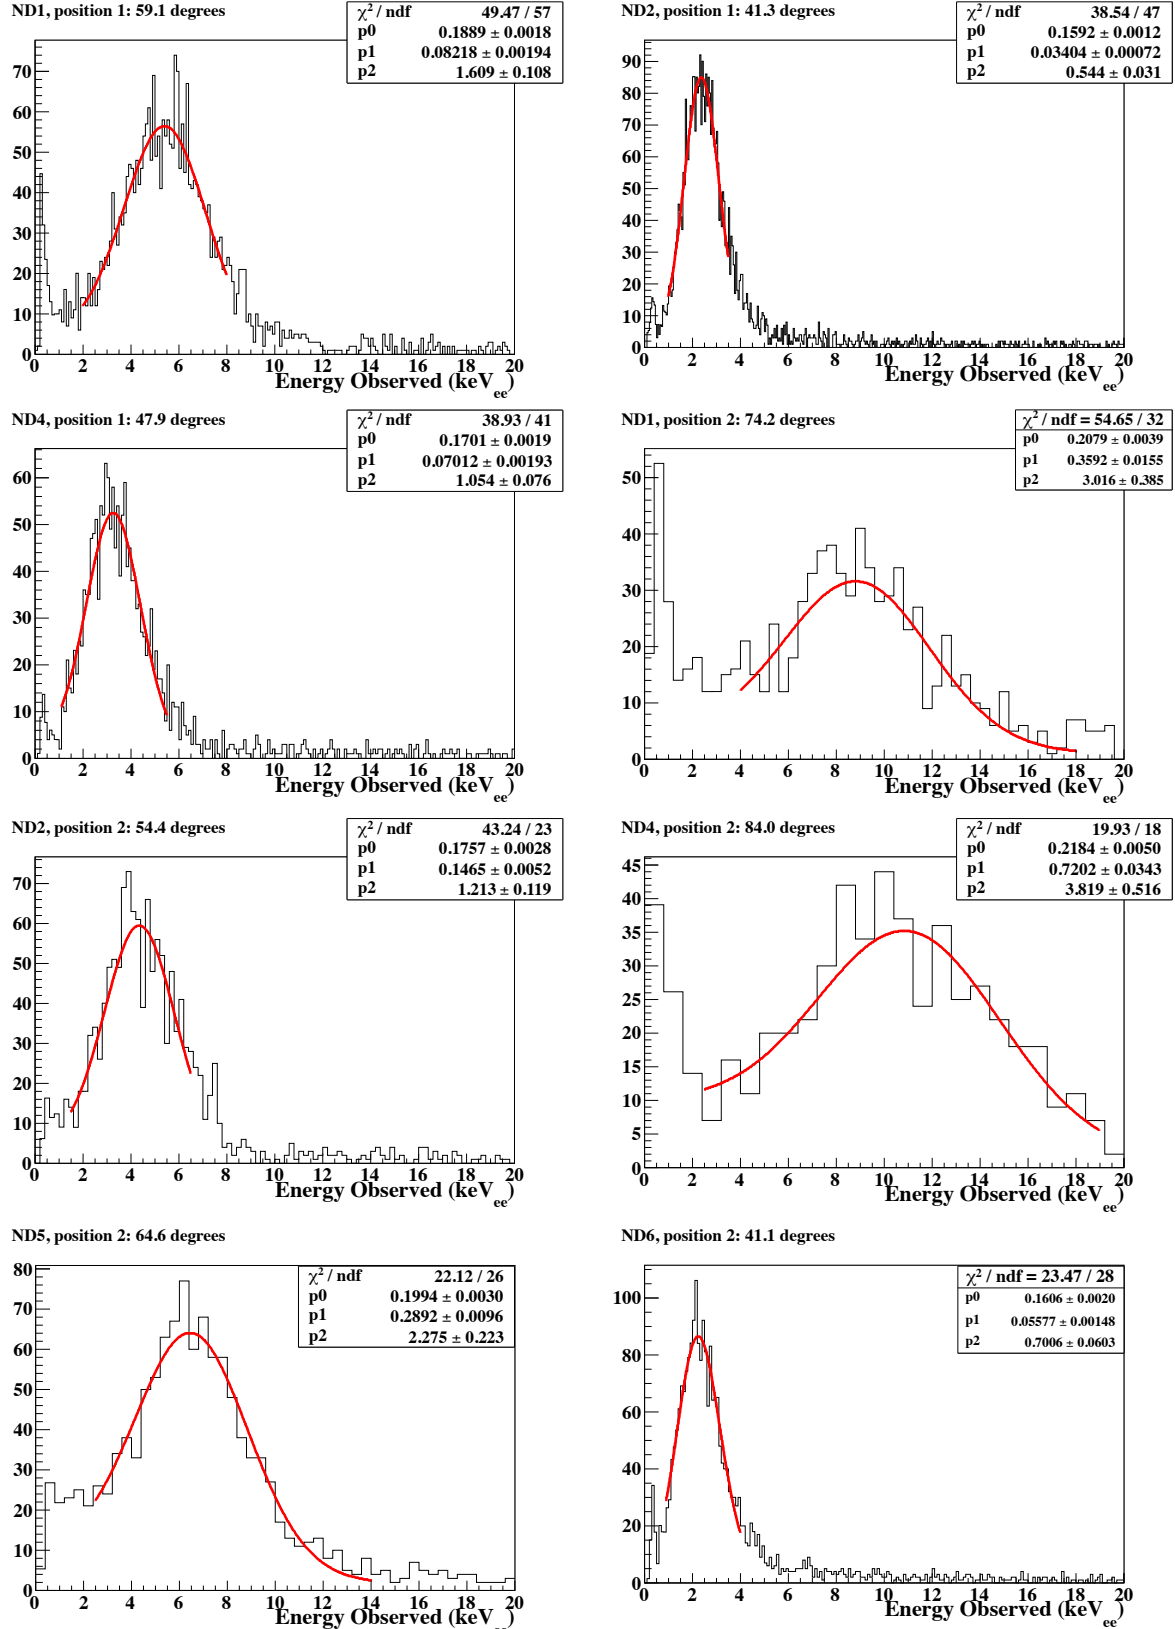


Figure 7.15: Fits of the Monte-Carlo-simulated spectra of single-scattering neutron-recoil events to the data gathered at the Notre Dame pulsed-neutron facility. This method of fitting the Monte Carlo with the data was used for nuclear recoils above 10 keV_{nr} in energy. The fit parameters are as follows: p_0 is the quenching factor, p_1 is an amplitude adjustment, and p_2 is the effective Fano-factor, which represents the spread of the simulated spectrum.




Figure 7.16: Trigger-efficiency curve for the Na-quenching measurement, evaluated from the runs described in Section 6.3.6.

7.4.2 Evaluating the fit for low-energy recoils

The low-energy recoil spectra were obscured by a loss in trigger efficiency. Because the recording of scintillation events was dependent on the digitization of the signal through the discriminators, if a pulse fell below the threshold of the discriminator, a nuclear-recoil event with low energy could be missed. The response of the discriminators should be dependent on the pulse amplitude, but the threshold is not a hard cutoff. The trigger efficiency as a function of amplitude was evaluated in a separate measurement by observing random scintillation events at the same time as the discriminator output, to see what fraction of events triggered the discriminator (see Section 6.3.6). For each event, the peak amplitude of each pulse was evaluated, and it was determined whether a signal from the discriminator followed. Pulses that occurred within an already firing discriminator signal were ignored.

The results of this trigger-efficiency measurement are shown in Figure 7.16. Events with an amplitude of less than 20 ADC counts were almost uniformly missed. This effect biases the energy spectra strongly for recoil energies less than 6 keV_{nr} and slightly for recoil energies less than 10 keV_{nr} , and necessitated a different analytical approach for evaluating the quenching factors for these energies.


In the peak-matching method, an amplitude spectrum of the data was created and corrected bin-by-bin with the trigger-efficiency curve when the trigger efficiency was greater than 10%, as shown in Figure 7.17. From this corrected spectrum, a peak could be resolved for all but the lowest energy measured (2.9 keV_{nr}). The peak was determined by fitting a gaussian to the spectrum, varying the fit range and binning and taking the mean value. The variation in this fit was added to the error in the observed energy.

For amplitudes between 27 and 81 ADC counts, a histogram of the energies (see Figure 7.17) was made by projecting the amplitude-energy histogram of the neutron events for each position configuration. The spectrum for position configuration 1 and a projection for 41 ADC counts are shown at the top of Figure 7.18. A gaussian was fit to each projection to determine the peak energy as a function of the peak amplitude. A linear fit to these projections was used to determine the peak energy for each neutron detector, as shown in the bottom of Figure 7.18. This method was confirmed to produce the correct peak energy for the higher energies, thus validating this method. The peak energy was then compared to the simulated energy to produce the quenching factor. Using this method on higher-energy recoils produced quenching factors within 3% of the values produced by the Monte Carlo-fit method.


For the lowest recoil energy, the apparent peak in the spectrum was within the energy range for which the trigger efficiency was significantly less than unity. The upper limit of the peak amplitude was taken to be 22 ADC counts, below which the trigger efficiency is less than 0.5. This amplitude corresponds with an energy of $0.65 \pm 0.01 \text{ keV}_{ee}$ in position configuration 1, according to Dr. Jingke Xu's analysis. The upper bound was used as a limit for the observed energy, which was compared with the simulated peak energy to produce a limit on the quenching factor.

Combining these results with the ones from the standard analysis, the quenching values are shown in Tables 7.3 and 7.4. The analysis was performed separately by Dr. Jingke Xu


Amplitude Spectrum for ND6, position 1: 18.2 degrees



Amplitude Spectrum for ND3, position 1: 24.9 degrees



Amplitude Spectrum for ND3, position 2: 31.1 degrees



Amplitude Spectrum for ND5, position 1: 32.2 degrees




Figure 7.17: Fits for the peak-matching method of fit for low-energy events. The green curve is the collected data. The blue curve is the trigger-efficiency curve with arbitrary scale. The black curve is the data adjusted for the trigger efficiency. The red curve is a gaussian fit to the adjusted data.

and myself; my values are shown here, while Dr. Xu's values are in [1] and in Table 7.5. The error analysis is discussed in Section 7.4.3.

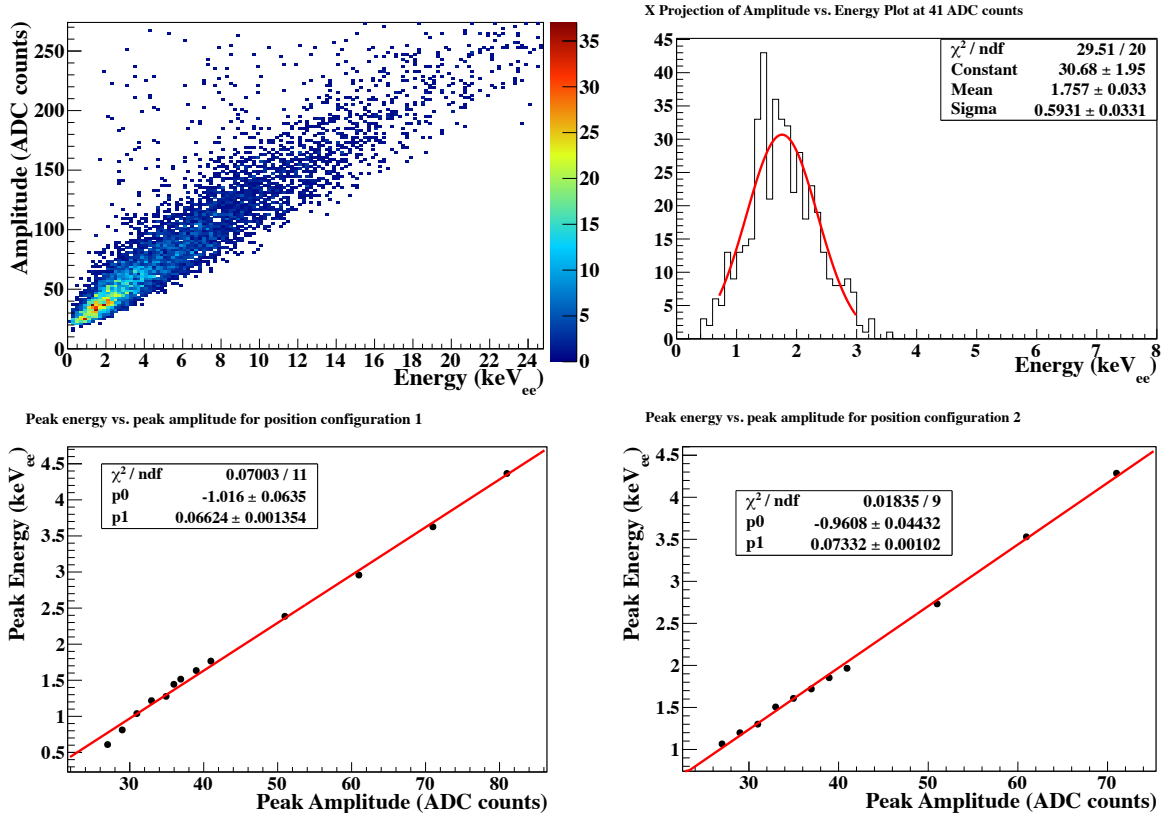


Figure 7.18: (Top left) Pulse peak amplitude vs. energy for the data. (Top right) X-projection of the spectrum for the determination of a relationship between peak energy and peak amplitude for 41 ADC counts, position configuration 1. (Bottom) Linear fits to determine the relationship between peak amplitude and peak energy for positions 1 (Left) and 2 (Right).

Table 7.3: ^{23}Na quenching factors from this analysis using the energy-spectrum fit method (described in Section 7.4.1). These values are consistent with the analysis done by Dr. Jingke Xu described in [1] and reproduced in Table 7.5. The simulated energy is the peak energy of the simulated spectrum, while the observed energy is the peak energy of the data. The quenching factor shown is a result of the Monte Carlo fit to the data in percent.

High-energy recoils, standard method			
Recoil angle	Simulated energy (keV_{nr})	Observed energy (keV_{ee})	Quenching factor (%)
41.1	14.31 ± 0.53	2.21 ± 0.05	15.9 ± 0.9
41.3	15.03 ± 1.79	2.36 ± 0.03	15.8 ± 2.0
47.9	19.36 ± 0.92	3.13 ± 0.04	16.6 ± 1.1
54.4	25.04 ± 0.80	4.27 ± 0.06	17.2 ± 1.2
59.1	28.99 ± 0.71	5.37 ± 0.08	18.9 ± 1.0
64.6	33.21 ± 1.94	6.43 ± 0.15	19.8 ± 1.4
74.2	43.01 ± 0.92	8.62 ± 0.12	20.4 ± 1.0
84.0	51.51 ± 1.70	10.42 ± 0.15	20.4 ± 1.8

Table 7.4: ^{23}Na quenching factors from this analysis using the peak-matching method (described in Section 7.4.2). The simulated energy is the peak energy in a simulated distribution. The observed energy is the energy derived from the peak amplitude after a trigger-efficiency correction. These values are consistent with the analysis done by Dr. Jingke Xu described in [1] and reproduced in Table 7.5. The limit for the 2.9-keV $_{nr}$ recoil was evaluated by Dr. Jingke Xu.

Low-energy recoils, peak-matching method			
Recoil angle	Simulated energy (keV $_{nr}$)	Observed energy (keV $_{ee}$)	Quenching factor (%)
18.2	2.93 ± 0.41	$< 0.65 \pm 0.01$	< 22.2
24.9	5.70 ± 0.26	0.76 ± 0.17	13.4 ± 3.0
31.1	8.82 ± 0.38	1.29 ± 0.10	14.7 ± 1.4
32.2	9.08 ± 0.69	1.42 ± 0.12	15.6 ± 1.9

Table 7.5: The Na quenching factors from Dr. Jingke Xu’s analysis, and modified from [1]. Quenching factors were evaluated by spectral fits between observation and simulation above 10 keV_{nr} and by comparing the peak energy positions at lower energies after correcting for the trigger-efficiency loss.

Dr. Jingke Xu’s quenching values, from [1].			
Scattering angle (deg)	Simulated Na recoil energy (keV _{nr})	Observed recoil energy (keV _{ee})	Quenching factor (%)
18.2	2.9 ± 0.7	<0.65	< 22
24.9	5.7 ± 0.7	0.76 ± 0.4	13.3 ± 1.8
31.1	8.8 ± 1.2	1.13 ± 0.5	12.9 ± 1.4
32.2	9.1 ± 1.2	1.46 ± 0.5	16.2 ± 1.2
41.1	14.3 ± 2.4	2.21 ± 0.9	15.9 ± 1.9
41.3	15.0 ± 1.4	2.36 ± 0.8	16.0 ± 1.0
47.9	19.4 ± 1.6	3.21 ± 1.0	16.8 ± 0.9
54.4	24.9 ± 2.4	4.10 ± 1.5	17.1 ± 1.0
59.1	29.0 ± 1.9	5.36 ± 1.9	18.8 ± 0.8
64.6	33.3 ± 2.8	6.19 ± 2.1	19.1 ± 1.1
74.2	43.0 ± 2.2	8.53 ± 2.7	20.4 ± 0.8
84.0	51.8 ± 2.6	10.59 ± 4.5	20.7 ± 1.0

7.4.3 Error analysis

The error on the quenching factors comes from statistical and systematic sources, which were added in quadrature. The statistical error of the fit from the spectrum-fit method was determined by varying the fit range and taking the maximum variation of the resulting quenching factors. For the peak-matching method, the error in the fit to the energy spectrum was used. For higher energies, the two methods produced a slightly different value for the quenching factor (<3%). This difference was also included as a statistical uncertainty and added in quadrature for all energies.

The systematic error on the simulated energy is determined by the uncertainty in the recoil angle from the detector positions and the uncertainty in the mean of the incoming neutron energy. A conservative estimate for the latter was placed at 10 keV, which is small compared to the simulated 4% spread in the nominal energy of 690 keV. The error for the positions was determined by assuming a displacement of the detectors by the amounts shown

Table 7.6: Uncertainty in the recoil energy (ΔE_R) as a function of the angular uncertainty.

Recoil energy (keV _{nr})	Recoil angle (deg)	ΔE_R (keV _{nr})
2.9	18.2 ± 1.2	0.4
5.7	24.9 ± 0.6	0.3
8.8	31.1 ± 0.7	0.4
9.1	32.2 ± 1.2	0.7
14.3	41.1 ± 0.7	0.5
15.0	41.3 ± 2.6	1.8
19.4	47.9 ± 1.1	0.9
24.9	54.4 ± 0.9	0.8
29.0	59.1 ± 0.7	0.7
33.3	64.6 ± 2.1	1.9
43.0	74.2 ± 0.7	0.9
51.8	84.0 ± 1.6	1.7

in Table 6.4 in directions that would maximize and minimize the recoil angle displacement, $\Delta\theta$. The larger of the two were taken as the uncertainty in the angle. The resultant uncertainty in the recoil energy was calculated analytically using the relation in Equation 6.4. The values for the angular uncertainties and the resultant uncertainties in the recoil energy are shown in Table 7.6.

The uncertainty in the observed energy comes primarily from the uncertainty in the light yield, and the uncertainty in the s.p.e. mean. The standard deviation for the s.p.e. mean was taken as the uncertainty, which was 0.14%. The residuals in Figure 7.9 were taken as the uncertainty in the light yield, which added another 1.4% uncertainty. The difference between the peak-matching method and the Monte Carlo method for high-energy recoils was at most 3%, so this difference was included in the error as well.

For the peak-matching method, there was additional uncertainty added to account for the variance in the peak amplitude as a function of the fit range and binning and the residual error in the linear fit of the peak amplitude to the peak energy.




Figure 7.19: Previous observations of non-linearity in the electron-recoil scintillation output in NaI(Tl). (Left) an observed non-linearity by Moses [141], with a peak around 10 keV_{ee} . (Right) γ calibrations by DAMA/LIBRA, observing no non-linearity in the scintillation output.

7.4.4 Other considerations

A non-linearity in the scintillation output for electron recoils in NaI(Tl) has been observed by Moses [141]. However, DAMA/LIBRA claims to see no non-linearity in the output for their crystals by observing several γ full-energy peaks in their background [41]. The results of both of these observations are summarized in Figure 7.19. Though these observations disagree, it is worth noting that the DAMA/LIBRA calibration does not have any points in the region where the non-linearity is most pronounced, around 10 keV_{ee} . The quenching factors, should such a non-linearity exist, should be adjusted given that the calibration in our measurement was done at 57.6 keV_{ee} . The difference could be on the order of 20%. The adjusted quenching values are shown in Table 7.7. These values were determined by deriving an adjustment factor for each energy based on the curve in Figure 7.19 and the relative light output at 57.6 keV_{ee} .

Table 7.7: Quenching Factors adjusted for the non-linearity of the electron-recoil scintillation efficiency in NaI(Tl). The adjustment factors were derived by digitizing the curve for NaI(Tl) on the left of Figure 7.19 and comparing with the relative light output at 57.6 keV_{ee} .

Recoil energy	Adjustment factor	Non-linearity adjusted quenching factor
2.9	1.07	<20.7
5.7	1.17	11.5
8.8	1.22	12.0
9.1	1.22	12.8
14.3	1.18	13.5
15.0	1.18	13.4
19.4	1.14	14.6
25.0	1.1	15.6
29.0	1.08	17.5
33.2	1.06	18.7
43.0	1.03	19.8
51.5	1.01	20.2

7.5 Iodine recoils

In choosing the scattering angles to be measured in this experiment, the observation of ^{127}I recoils was not prioritized. The height and width of the expected iodine signal could be predicted from simulations combined with the Na spectrum. A 3σ upper limit of 0.065 on the ^{127}I quenching factor at 10 keV_{nr} was set based on the absence of an observed iodine peak above 0.65 keV_{ee} ⁵. This upper limit is lower than DAMA's reported 0.09 [53].

7.6 Channeling effect

It has been theorized that a channeling effect in NaI(Tl), as described in Section 2.3.5, could account for some of the disagreement between the DAMA/LIBRA observation and exclusions from other experiments. The orientation of the lattice in the crystal used in this measurement is unknown, so a sensitive measurement of a potential channeling effect could not be made. However, because 12 different recoil angles were measured, there was some

⁵This limit was set in Dr. Jingke Xu's analysis and was not checked by me.

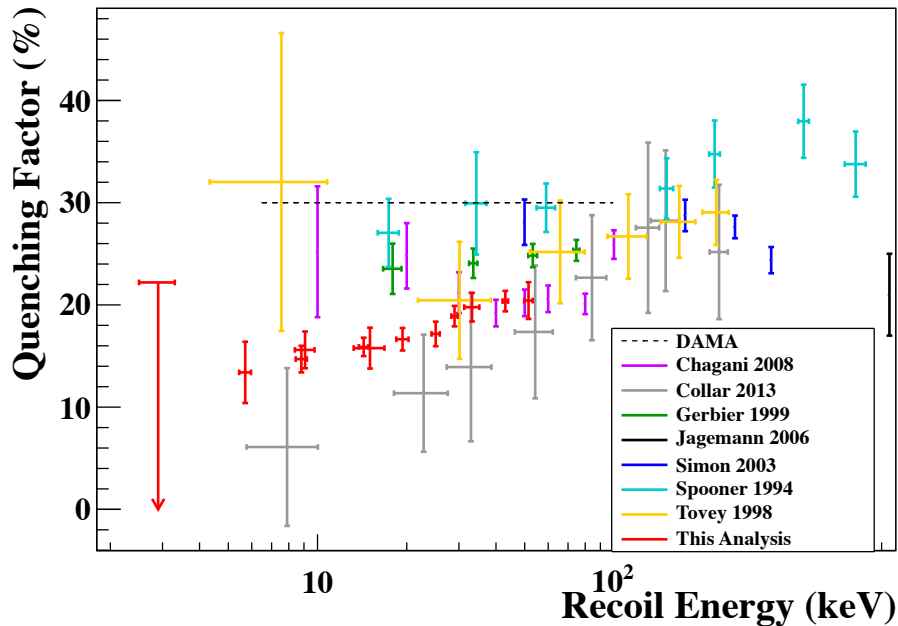


Figure 7.20: Results of the Na quenching measurement, along with previously measured values from other experiments. The new quenching factors show a strong energy dependence, and fall between previous measurements from [54] and [55] at low energy. They are also statistically consistent with these previous measurements.

possibility that such an effect may be observed. The quenching values measured followed a continuous trend with no major outliers, so no channeling effect was observed.

7.7 Relation to other measurements

The new quenching factors for Na, along with those of the previous measurements, are shown in Figure 7.20. Comparisons between measurements are complicated by the observed non-linearity in NaI(Tl) electron recoil response, which some other experiments do not account for in their reported values. Different experiments use different calibration sources, making a direct comparison difficult. The quantities shown for the SABRE measurement in Figure 7.20 are the non-adjusted values.

The quenching factors show a strong energy dependence, and fall between previous measurements at low energies. Because of the large uncertainties in previous measurements, the

SABRE measurement is consistent with these values, but has a much lower uncertainty. The SABRE measurement constitutes the most precise measurement of the Na quenching factor to date.

Chapter 8

Measurement of the Na

Nuclear-Recoil Quenching Factor:

Implications for WIMP Dark Matter

Within the Standard Halo Model (SHM), the scale of the energy spectrum for WIMP dark-matter interactions is uniquely determined by the WIMP mass (m_χ) and the WIMP-nucleon cross section. The quenching factor, which determines the observed energy scale, therefore factors heavily in the determination of the WIMP parameters that are compatible with the observed DAMA/LIBRA modulation spectrum. The comparison of DAMA/LIBRA with other null experiments is dependent on this value.

Generally, when comparisons are made between direct-detection experiments, the Standard Halo Model (SHM) is assumed, but there is some evidence that this model may not accurately describe the astrophysical state of our galaxy [32]. Furthermore, the nature of the dark-matter/nucleus interaction is also not well-established. However, due to the large number of alternative models, the implications for the DAMA/LIBRA parameter-space preferred regions are only evaluated in the SHM here.

DAMA/LIBRA uses the 59.5-keV_{ee} peak of the ²⁴¹Am spectrum to calibrate their detector light yield. This energy is very close to the energy of our 57.6-keV_{ee} calibration. Therefore, there is no adjustment needed to correct the quenching factors reported in Table 7.3 and 7.4 for the non-linearity in the NaI(Tl) electron-recoil response in order to apply the new quenching results to the DAMA/LIBRA spectrum, since all of their results are reported in keV_{ee} as defined by their calibration at 59.5 keV_{ee}.

8.1 Calculation of the WIMP dark-matter event rate

The WIMP dark-matter elastic-scattering event rate is derived here following the derivations of [32, 28, 142].

In the standard WIMP model, a WIMP of mass m_χ scatters off of a nucleus with mass M , causing the nucleus to recoil with energy $E_R = (\mu_\chi^2 v^2 / M)(1 - \cos \theta_{cm})$, where μ_χ is the WIMP-nucleus reduced mass, equal to $m_\chi M / (m_\chi + M)$. v is the relative velocity, and θ_{cm} is the scattering angle in the center-of-mass frame. The differential rate of WIMP-nucleus interactions per unit detector mass, typically given in units of cpd/kg/keV_{nr}, is therefore:

$$\frac{dR}{dE_R}(E_R, t, m_\chi, \sigma) = \frac{n_\chi}{M} \langle v \frac{d\sigma}{dE_R} \rangle = \frac{2\rho_\chi}{m_\chi} \int_{v_{min}}^{\infty} d^3v v f(\mathbf{v}, t) \frac{d\sigma}{dq^2}(q^2, v), \quad (8.1)$$

where $n_\chi = \rho_\chi / m_\chi$ is the WIMP number density, $f(\mathbf{v}, t)$ is the probability distribution function for the WIMP velocity, \mathbf{v} as a function of time, t . $q = \sqrt{2ME_R}$ is the momentum transfer of the interaction, and $\frac{d\sigma}{dq^2}(q^2, v)$ is the velocity-dependent differential cross section. In this analysis, ρ_χ is the mean dark-matter density in the local neighborhood, 0.3 GeV/cm³ [143].

In typical cross-section models, $\frac{d\sigma}{dq^2}(q^2, v)$ is expressed by:

$$\frac{d\sigma}{dq^2}(q^2, v) = \frac{\sigma(q)}{4\mu_\chi^2 v^2} \Theta(q_{max} - q), \quad (8.2)$$

where $\sigma(q)$ is an effective cross-section, discussed later, $\Theta(q_{\max} - q)$ is the Heaviside step-function, and $q_{\max} = 2\mu_\chi v$ is the maximum momentum transfer for a collision of velocity v . The differential event rate can therefore be expressed as:

$$\frac{dR}{dE_R}(E_R, t, m_\chi, \sigma) = \frac{\sigma(q)\rho_\chi}{2m_\chi\mu_\chi^2}\eta(v_{\min}, t), \quad (8.3)$$

where v_{\min} is the minimum velocity that can produce a recoil of energy E_R , which is equal to $\sqrt{\frac{ME_R}{2\mu_\chi^2}}$ for elastic scattering. $\eta(v_{\min}, t)$ is an expression denoting the mean inverse speed of the WIMPs in the lab frame at time t , and is given by:

$$\eta(v_{\min}, t) = \int_{v>v_{\min}} d^3v \frac{f(\mathbf{v}, t)}{v}. \quad (8.4)$$

The total rate in a direct-detection experiment between two energies E_1 and E_2 is dependent on detector characteristics in addition to the nuclear and astrophysical characteristics of the WIMP-nucleus interaction. Typically, detectors have a detection efficiency, $\epsilon(E')$, for a given electron-equivalent energy, $E' = Q(E_R)E_R$, where $Q(E_R)$ is the quenching factor and E_R is the recoil energy. The detection efficiency is the chance that a detector will detect an event with quenched energy E' .

Detectors also have an energy resolution, $\sigma_E(E')$ which denotes the spread in the measured energies of events. $\Phi(E', E_1, E_2)$ is a response function denoting the chance that a recoil of energy E_R and with quenching factor $Q(E_R)$ is observed between electron-equivalent energies E_1 and E_2 . If the observed energy E' is normally distributed about E' with a standard deviation $\sigma_E(E')$, the response function can be expressed as:

$$\Phi(E', E_1, E_2) = \frac{1}{2} \left[\operatorname{erf} \left(\frac{E_2 - E'}{2\sigma_E(E')} \right) - \operatorname{erf} \left(\frac{E_1 - E'}{2\sigma_E(E')} \right) \right]. \quad (8.5)$$

In the case of DAMA/LIBRA, the energy resolution reported by the collaboration is $\sigma_E(E') = 0.448\sqrt{E'} + 0.0091E'$ where E' is the quenched energy in keV_{ee} [41]. Though

DAMA reports an energy-dependent efficiency, they claim that their modulation spectrum is already adjusted for this efficiency, and therefore in this analysis of the DAMA rate, $\epsilon(E')$ is taken to be equal to 1.

The rate per unit mass of detector between energies E_1 and E_2 , then, can be expressed in the following way [142]:

$$R(t, m_\chi, \sigma) = \int_0^\infty dE_R \epsilon(E') \Phi(E', E_1, E_2) \frac{\rho_\chi}{2m_\chi \mu_\chi^2} \sigma(q) \eta(\mathbf{v}_{\min}, t). \quad (8.6)$$

The forms for the velocity distribution and the effective cross-section are discussed in the following sections.

8.2 The Standard Halo Model

8.2.1 The velocity distribution

The SHM assumes an isothermal, isotropic, spherical dark-matter halo with a Maxwellian distribution of velocities with velocity dispersion σ_v in the galactic rest-frame. The distribution of velocities in the galactic rest frame, $\tilde{f}(\mathbf{v})$ is assumed to be:

$$\tilde{f}(\mathbf{v}) = \begin{cases} \frac{1}{N_{\text{esc}}} \left(\frac{3}{2\pi\sigma_v^2} \right)^{3/2} e^{-3|\mathbf{v}|^2/2\sigma_v^2}, & \text{for } |\mathbf{v}| < v_{\text{esc}}, \\ 0, & \text{otherwise.} \end{cases} \quad (8.7)$$

where the normalization factor is:

$$N_{\text{esc}} = \text{erf}(z) - 2 \frac{z}{\sqrt{\pi}} e^{-z^2}. \quad (8.8)$$

Here, $z \equiv v_{\text{esc}}/\bar{v}_0$, where $\bar{v}_0 = \sqrt{\frac{2}{3}}\sigma_v = 235$ km/s is the most probable speed of the WIMP at the radius of the solar system [32].

The velocity distribution in the lab frame is dependent on the lab velocity with respect to the galactic rest frame, \mathbf{v}_{obs} . This velocity is dependent on the Sun's velocity relative to the galactic rest-frame, \mathbf{v}_{\odot} , and the Earth's velocity relative to the Sun, \mathbf{V}_{\oplus} . In galactic-frame coordinates, $\hat{\mathbf{x}}$ is the direction toward the Galactic Center, $\hat{\mathbf{y}}$ is the direction of the disk rotation, and $\hat{\mathbf{z}}$ is the North Galactic Pole. In these coordinates, $\mathbf{v}_{\odot} = (11, 247, 7)$ km/s [142]. The magnitude of \mathbf{V}_{\oplus} is taken to be 29.8 km/s [142], while its directional vector is defined with $\hat{\mathbf{e}}_1$ and $\hat{\mathbf{e}}_2$, the directions of the Earth's velocity at the Spring Equinox (March 21) and the Summer Solstice (June 21), respectively, defined according to [142] as:

$$\hat{\mathbf{e}}_1 = (0.9931, 0.117, -0.01032) \quad (8.9)$$

$$\hat{\mathbf{e}}_2 = (-0.0670, 0.4927, -0.8676). \quad (8.10)$$

With this, the velocity of the lab with respect to the dark-matter halo is:

$$\mathbf{v}_{\text{obs}} = \mathbf{v}_{\odot} + |\mathbf{V}_{\oplus}| [\hat{\mathbf{e}}_1 \cos \omega(t - t_1) + \hat{\mathbf{e}}_2 \sin \omega(t - t_1)], \quad (8.11)$$

where $\omega = 2\pi/\text{year}$ and t_1 is March 21.

If one additionally defines $x \equiv v_{\text{min}}/\bar{v}_0$, and $y \equiv v_{\text{obs}}/\bar{v}_0$, the mean-inverse speed in the lab frame becomes:

$$\eta(v_{\text{min}}, t) = \begin{cases} \frac{1}{\bar{v}_0 y} & \text{for } z < y, x < |y - z| \\ \frac{1}{2N_{\text{esc}}\bar{v}_0 y} \left[\text{erf}(x + y) - \text{erf}(x - y) - \frac{4}{\sqrt{\pi}} y e^{-z^2} \right] & \text{for } x > y, x < |y - z| \\ \frac{1}{2N_{\text{esc}}\bar{v}_0 y} \left[\text{erf}(z) - \text{erf}(x - y) - \frac{2}{\sqrt{\pi}} (y + z - x) e^{-z^2} \right] & \text{for } |y - z| < x < y + z \\ 0 & \text{for } y + z < x \end{cases} \quad (8.12)$$

8.2.2 The effective cross section

The effective cross section, $\sigma(q)$, consists of a spin-independent (SI) component and a spin-dependent (SD) component:

$$\sigma(q) = \sigma_{\text{SI}} + \sigma_{\text{SD}}. \quad (8.13)$$

The spin-independent cross-section is given by the zero-momentum-transfer cross-section, $\sigma_{0,\text{SI}}$, and an energy-dependent form factor, $F(q)$, that accounts for the finite size of the nucleus:

$$\sigma_{\text{SI}} = \sigma_{0,\text{SI}} F^2(q) \quad (8.14)$$

The spin-independent zero-momentum-transfer cross-section is given by:

$$\sigma_{0,\text{SI}} = \frac{4\mu_\chi^2}{\pi} [Zf_p + (A - Z)f_n]^2, \quad (8.15)$$

where μ_χ , again, is the WIMP-nucleus reduced mass, Z is the number of protons, A is the atomic mass, and f_p and f_n are the WIMP couplings to the proton and neutron, respectively. Commonly, it is assumed that $f_p \approx f_n$, so that the spin-independent cross-section coherently scales with A^2 , such that the cross section at zero-momentum transfer is

$$\sigma_{0,\text{SI}} = \frac{\mu_\chi^2}{\mu_p^2} \sigma_{0,p} A^2, \quad (8.16)$$

where μ_p is the WIMP-proton reduced mass and $\sigma_{0,p}$ is the WIMP-proton cross-section (which is the same as the WIMP-neutron cross-section) at zero momentum transfer. Therefore, the total spin-independent cross-section is given by:

$$\sigma_{\text{SI}}(q) = \sigma_{0,p} \frac{\mu_\chi^2}{\mu_p^2} A^2 F^2(q). \quad (8.17)$$

The most realistic shape for the nucleus is considered to be the Fermi distribution [28]:

$$\rho(r) = \rho_0 \left[1 + e^{\frac{r-c}{a}} \right]^{-1}, \quad (8.18)$$

where ρ_0 , a , and c are constants in r that depend on the nucleus. However, there is no analytical expression for the form factor for a distribution of this kind. One traditional form factor which is a close approximation to the Fermi distribution is the Helm distribution, which assumes a uniform spherical nucleus [28]. This analytical fit to the Fermi distribution breaks down for large A and large E_R , but does provide a good fit for the energies of interest to the DAMA modulation, as shown in Figure 8.1. The Helm distribution is given by:

$$F(q)_H = 3e^{-q^2 s^2/2} \frac{\sin(qr_n) - qr_n \cos(qr_n)}{(qr_n)^3}, \quad (8.19)$$

where s is the nuclear skin thickness, and is generally taken to be 0.9 fm, and r_n is the effective radius of the nucleus. In this analysis, the fit to the nuclear radius outlined in [28] and the following form for r_n are used:

$$r_n = \sqrt{c^2 + \frac{7}{3}\pi^2 a^2 - 5s^2}, \quad (8.20)$$

where

$$c = (1.23A^{1/3} - 0.6) \text{ fm}, \quad (8.21)$$

and $a = 0.52$ fm.

The spin-dependent form factor for the WIMP takes the following form:

$$\sigma_{\text{SD}}(q) = \frac{32\mu_\chi^2 G_F^2}{2J+1} [a_p^2 S_{pp}(q) + a_p a_n S_{pn}(q) + a_n^2 S_{nn}(q)], \quad (8.22)$$




Figure 8.1: Nuclear form factors for Na (Left) and I (Right), from [28]. The Fermi distribution is shown as a solid line, while the Helm distribution is shown with a dashed line with $r_n = 1.14A^{1/3}$ and $s = 0.9$ fm.

where J is the total spin of the nucleus, G_F is the Fermi constant, a_p and a_n are the WIMP-nucleon couplings to the proton and neutron, respectively, in units of $2\sqrt{2}G_F$. The $S(q)$ are the nuclear-structure functions specific to the dark-matter target.

a_p and a_n can be related to the isoscalar (a_0) and isovector (a_1) couplings, as follows:

$$a_0 = a_p + a_n \quad (8.23)$$

$$a_1 = a_p - a_n. \quad (8.24)$$

It is useful to define a unit-less parameter, $y = qb/2$, where $b = 1.6864$ fm. In this form, the nuclear-structure functions used in this analysis for ^{23}Na are:

$$S_{00}^{Na}(y) = 0.0380 - 0.1743y + 0.3783y^2 - 0.3430y^3 \quad (8.25)$$

$$S_{01}^{Na}(y) = 0.0647 - 0.3503y + 0.9100y^2 - 0.9858y^3 \quad (8.26)$$

$$S_{11}^{Na}(y) = 0.0275 - 0.1696y + 0.5077y^2 - 0.6180y^3, \quad (8.27)$$

as calculated in [144]. From the same source, the structure functions for ^{127}I are:

$$S_{00}^I(y) = e^{-2y}(0.0983 - 0.4891y + 1.1402y^2 - 1.4717y^3 + 1.1717y^4 - 0.5646y^5 + 0.1583y^6 - 0.0239y^7 + 0.0015y^8) \quad (8.28)$$

$$S_{01}^I(y) = e^{-2y}(0.1199 - 0.6184y + 1.5089y^2 - 2.0737y^3 + 1.7731y^4 - 0.9036y^5 + 0.2600y^6 - 0.0387y^7 + 0.0024y^8) \quad (8.29)$$

$$S_{11}^I(y) = e^{-2y}(0.0366 - 0.1950y + 0.5049y^2 - 0.7475y^3 + 0.7043y^4 - 0.3930y^5 + 0.1219y^6 - 0.0192y^7 + 0.0012y^8). \quad (8.30)$$

In traditional analyses of the WIMP parameter-space regions in the SHM, the spin-independent cross-section is evaluated separately from the spin-dependent ones. This approach will be used in the following sections, such that no mixing of WIMP-nucleon interaction types is considered.

8.3 Results for the Standard Halo Model

The implications of the new energy-dependent quenching factor measurements on the prospects for a WIMP-dark-matter explanation of the DAMA/LIBRA signal were evaluated for the SHM¹ through the process described in Sections 8.1 and 8.2. The new quenching factors were fit with a square-root plus a constant below $E_R = 43\text{keV}_{nr}$, as shown in Figure 8.2. The quenching factor was set to $Q = 0.20$ for higher energies. The DAMA/LIBRA iodine quenching value of 0.09 was used. The combined DAMA-DAMA/LIBRA modulation data were fit in 0.5keV_{ee} bins below 10keV_{ee} to the prediction for a SHM WIMP with the WIMP mass and the WIMP-nucleon cross-section as free parameters. Though typically only spin-independent cross-sections are considered when comparing experiments, because

¹These evaluations were performed with a fitting code written by myself.

of the lightness of the Na target, spin-dependent interactions were considered as well, to which heavier targets may not be as sensitive.

Three separate evaluations were made with no mixing of the cross-section types: a spin-independent WIMP-nucleon cross-section, $\sigma_{0,SI}$, a spin-dependent cross-section with the proton, $\sigma_{0,SD,p}$, and a spin-dependent cross-section with the neutron, $\sigma_{0,SD,n}$. The χ^2 were evaluated for each fit, creating confidence regions at 1σ , 3σ , and 5σ for the SHM WIMP and the DAMA data.

The best fits to the DAMA/LIBRA spectrum below 10 keV_{ee} with both DAMA's reported quenching values and the SABRE values are shown in Figure 8.3. The best fit to the data in both cases is a 65-GeV WIMP, with a secondary maximum around 10 GeV. The ~ 10 -GeV WIMP fits much better with the DAMA data using the DAMA quenching value than with the SABRE values.

The DAMA/LIBRA parameter space regions evaluated for their reported quenching values are shown in the top plot in Figures 8.4 and the left plots in 8.5, while the regions with the new quenching factors are shown on bottom and the right, respectively. These plots were generated by performing a χ^2 fit of the predicted modulation in the SHM with the 36-bin DAMA data below 10 keV_{ee} [40]. The regions shown correspond with the 1σ (red), 3σ (blue), and 5σ (green) contours. The parameter-space regions using DAMA's quenching values match well with the published regions in the literature, e.g. those in [142], validating the calculation.

Also shown in the Figures 8.4 and 8.5 is the mass-dependent cross-section limit curve based on the total rate observed by DAMA, shown in magenta. This curve was defined by the region at which the theoretical rate is higher than the rate observed by DAMA in one or more bins, assuming that the entire signal is attributable to WIMPs and not to background. A second, more stringent limit curve, shown in black, required that no bin have a rate higher than $0.3 \text{ cpd/kg/keV}_{ee}$ from $2\text{--}6 \text{ keV}_{ee}$. This requirement was based on DAMA's estimation of the fraction of their signal not attributable to background, shown in Figure 8.6. DAMA

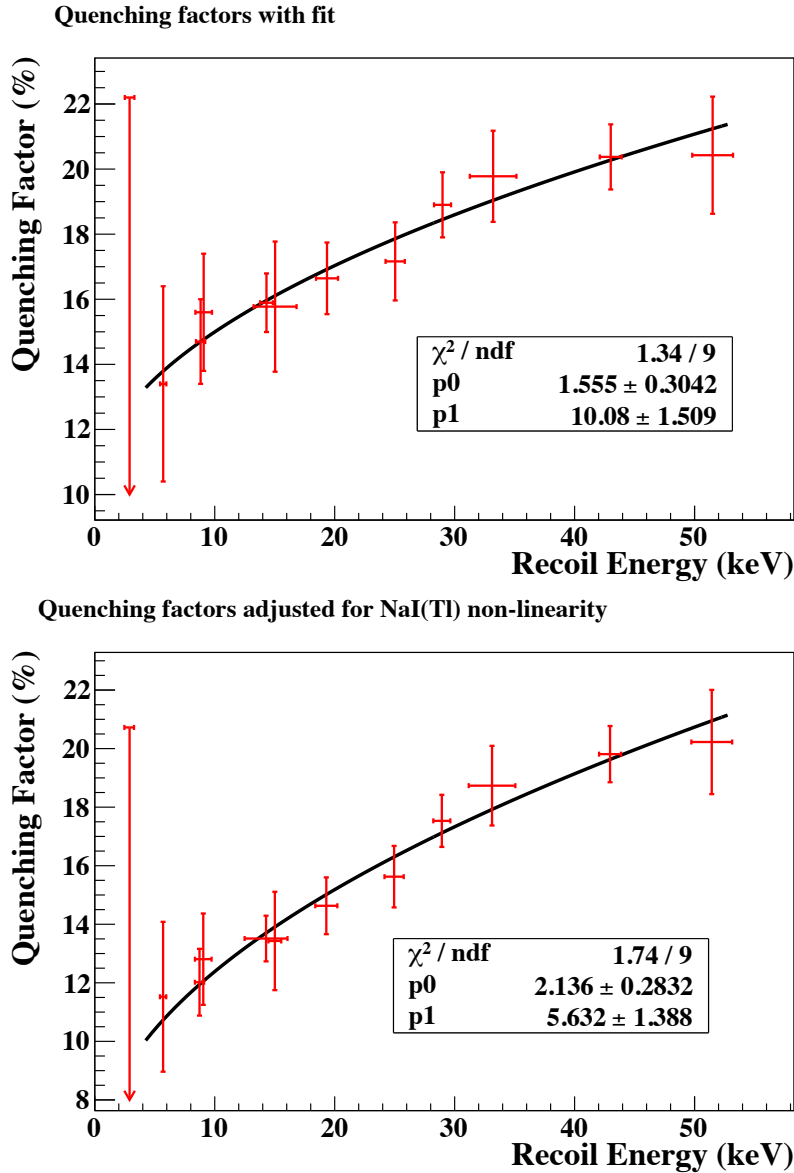


Figure 8.2: Fit for the new quenching factors. The fit follows the form $Q = p_0\sqrt{E_R} + p_1$ below $E_R = 43$. For higher energies, a value of $Q=0.20$ was used. The top figure shows the fit to the quenching values assuming a linear scintillation response, while the bottom figure adjusts for the non-linearity of NaI(Tl) described in Section 7.4.4.




Figure 8.3: Local best-fits for the DAMA/LIBRA modulation spectrum at ~ 10 GeV and ~ 65 GeV for a spin-independent WIMP, using the DAMA/LIBRA-reported quenching factors of 0.3 for Na and 0.09 for I in the green and magenta curves, and the new quenching values for Na in the blue and red curves, with and without the adjustment for the NaI(Tl) non-linearity. The best fit for the high-mass WIMP did not change substantially (since the DAMA iodine quenching factor was used), but the fit for the low-mass WIMP has significantly worsened. The χ^2/NDF for the light-WIMP fit went from 18.7/14 to 34.9/14 from the green curve to the red curve. The fit is even worse if the non-linearity of the NaI(Tl) scintillation response is taken into account, with a χ^2/NDF of 40.0/14.

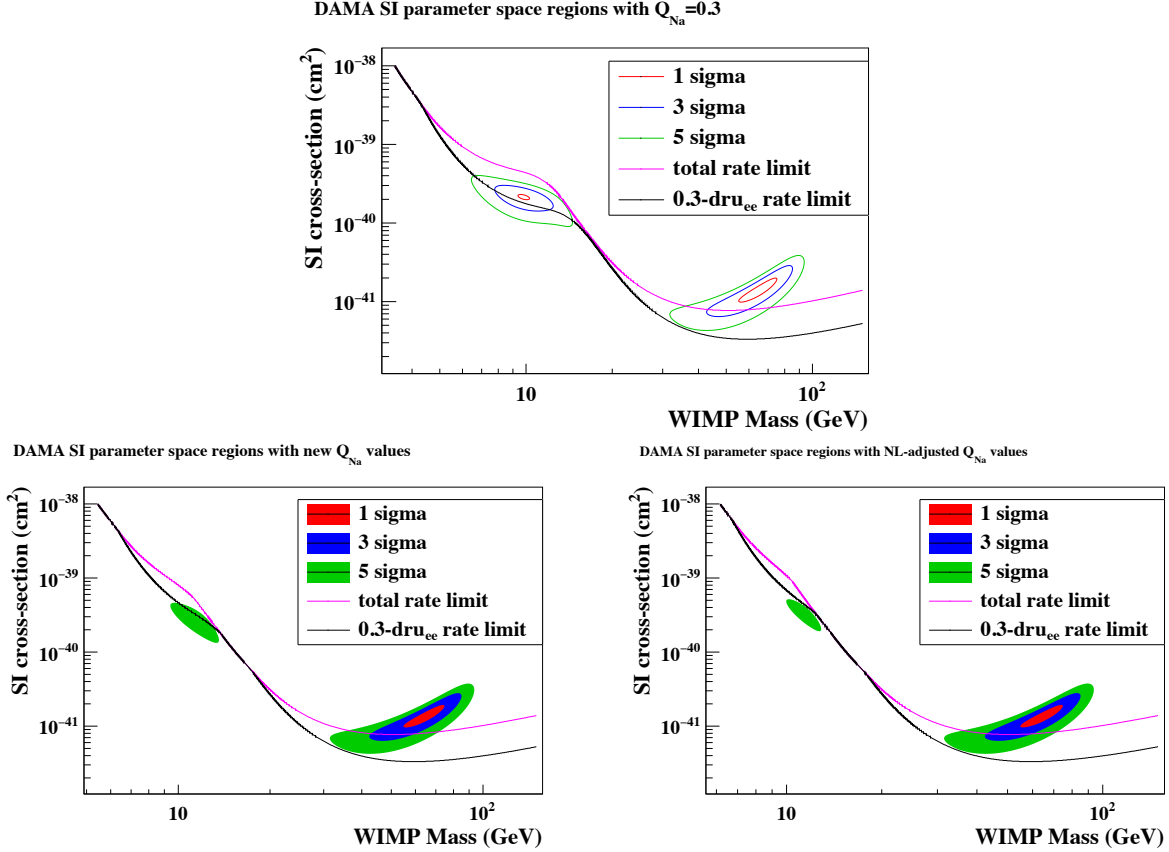


Figure 8.4: DAMA/LIBRA spin-independent parameter-space regions of interest, as evaluated by this analysis. The contours shown are the 1, 3, and 5σ regions in red, blue, and green, respectively. Also shown is the 3σ limit curve based on the total rate of the DAMA signal, in magenta. Above this line, the dark matter would produce too high of a rate in at least one energy bin. A more stringent limit barring the total dark-matter rate from exceeding $0.3 \text{ cpd/kg/keV}_{ee}$ between 2 and 6 keV_{ee} is shown in black. The overlap between the total-rate limits and the higher-mass region prompted the interest in a light-WIMP picture. The values for the old quenching factors are shown in contour lines in the top figure, while the regions for the new values are shown with filled regions in the bottom-left figure. The parameter-space regions for the non-linearity-adjusted quenching values are shown in the bottom-right figure. The light-WIMP no longer fits the data well—the 1- and 3σ regions for the light WIMP ($\sim 10 \text{ GeV}$) have disappeared.

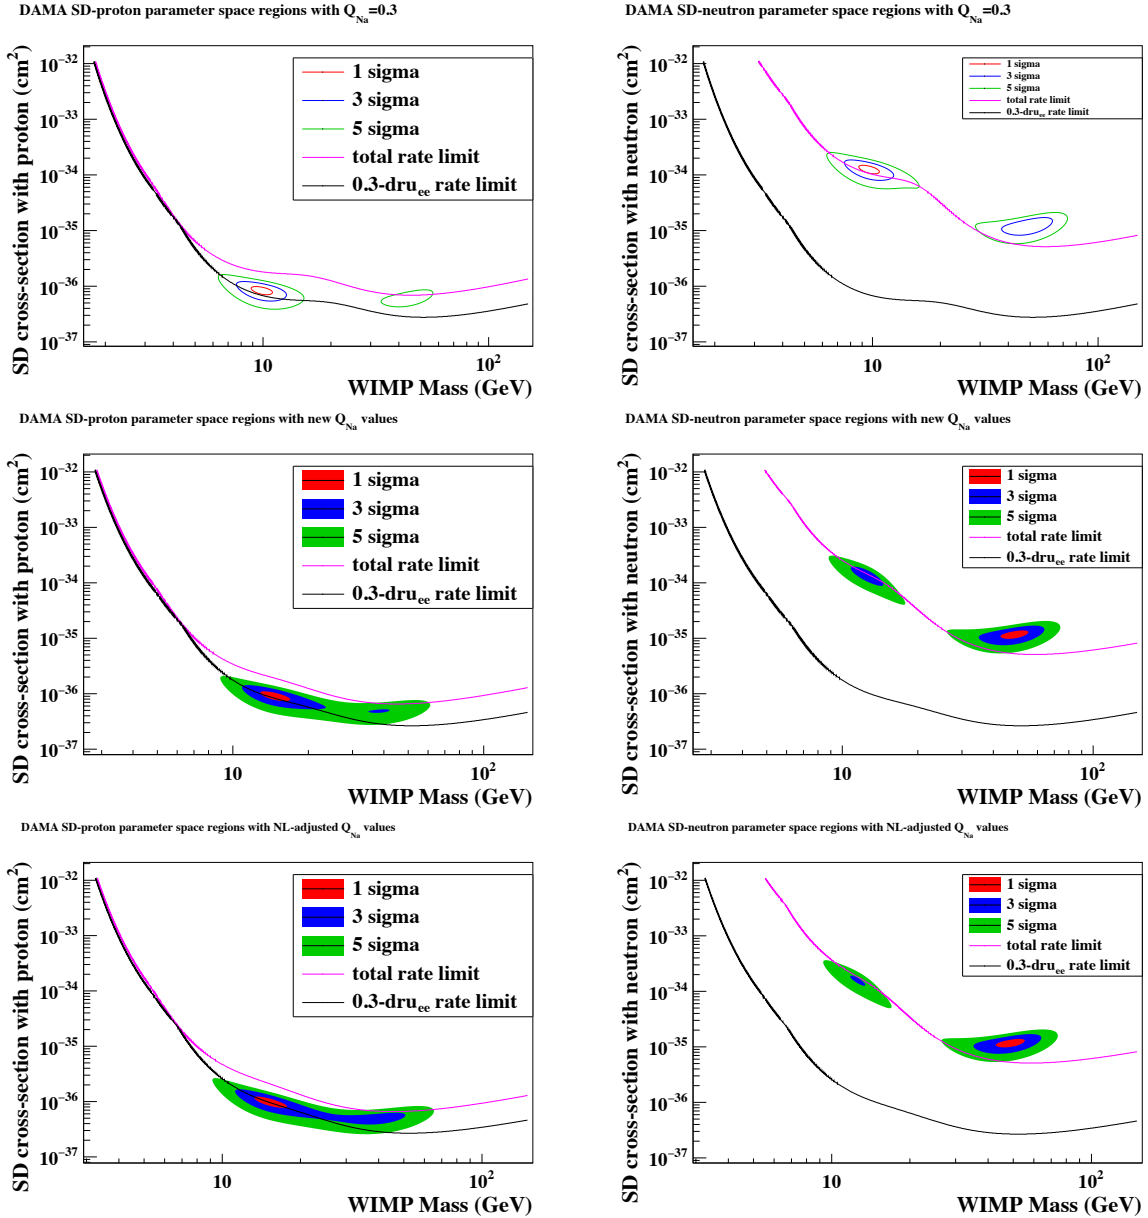


Figure 8.5: DAMA/LIBRA spin-dependent parameter-space regions of interest, as evaluated by this analysis. The contours shown are the 1, 3, and 5σ regions in red, blue, and green respectively. Also shown is the 3σ limit curve based on the total rate of the DAMA signal, in magenta, and the limit curve based on a requirement that the dark-matter total rate not exceed 0.3cpd/kg/keV $_{ee}$ between 2 and 6 keV $_{ee}$, in black. The top figures show the regions with the old quenching values, the new quenching values are used on the middle, and the new quenching values adjusted for the NaI(Tl) non-linearity are shown on the bottom. The left figures show the couplings to the proton, while the right figures show the coupling to the neutron.




Figure 8.6: DAMA/LIBRA reported background and maximum signal. The background is extrapolated down to their software threshold by fitting a line to the data at higher energies and a gaussian to the ^{40}K peak assuming a K concentration of 13 ppb. This fit is shown in red. The green curve is the quoted maximum possible dark-matter signal, which DAMA claims is $0.25 \text{ cpd/kg/keV}_{ee}$ between 2 and 4 keV_{ee} . Note that in DAMA’s convention, keV refers to the electron-equivalent energy, which is referred to as keV_{ee} in this text. Figure from [51].

claims that the rate unattributable to background is $0.25 \text{ cpd/kg/keV}_{ee}$ on average between 2 and 4 keV_{ee} [51], but they do not give this value as a function of energy. From the figure, it was estimated that the maximum value in this window was $0.3 \text{ cpd/kg/keV}_{ee}$ where the signal crosses the software energy threshold of 2 keV_{ee} .

In the SI case, the conflict between the high-mass WIMP region, in which most of the interactions are with ^{127}I , and the limit curve based on the total rate, prompted initial interest in the light-WIMP region around and below 10 GeV. However, the new quenching values heavily disfavor the light-WIMP region; the 1 and 3σ contours have disappeared entirely at low mass, meaning the low-mass region no longer fits the data to within 3σ . This increases the tension of the DAMA/LIBRA result with other null experiments in the SI case,

Table 8.1: Dark-matter best-fits to DAMA/LIBRA modulation data, with new and old quenching values.

Quenching factors	m_χ (GeV)	$\sigma_{0,\text{SI}}(\text{cm}^2)$	χ^2	NDF
DAMA	65.5	1.41×10^{-41}	16.8	14
	9.8	2.13×10^{-40}	18.7	14
SABRE	65.0	1.37×10^{-41}	16.9	14
	11.5	2.67×10^{-40}	34.9	14
SABRE (non-linearity adjusted)	65.1	1.37×10^{-41}	16.9	14
	11.4	3.25×10^{-40}	40.0	14

since the high-mass region is already in strong tension with other null experiments, and much of recent attention had been paid to the low-mass region, where other experiments had set less stringent limits. The limit based on the total rate from the black curve also heavily disfavors the high-mass WIMP, making a simple spin-independent WIMP very unlikely as an explanation for the DAMA modulation.

For the spin-dependent case, shown in Figure 8.5, the favored region shifts to slightly higher mass from the old quenching values to the new ones. For the coupling with the proton, the low-mass WIMP still has a good fit with the data which is not completely eliminated by their total rate.

8.4 Conclusions

The calculation of the preferred parameter-space regions in the SHM highlight the importance of understanding the Na and I quenching factors before drawing any conclusions about dark matter from the DAMA/LIBRA modulation signal. The quenching factors measured by SABRE imply a profound difference in the DAMA modulation’s implications for the dark-matter mass and cross section compared to what is generally assumed.

The picture for reconciling DAMA with conflicting experiments remains complicated, even in light of the new quenching values reported in this study. Though comparisons between experiments of different targets generally assume the standard WIMP interaction

model and the SHM, and though the prospects for a reconciliation in this model are not promising, the assumptions that are integral to the model are not necessarily true. As mentioned in Chapter 2, there is some evidence from high-resolution N-body dark-matter simulations that the SHM may not accurately reflect the local WIMP velocity distribution, especially at the high-velocity tail, which is important for light WIMPs [32]. In addition, some have theorized that alternate WIMP models may help reconcile the conflict between DAMA/LIBRA and other experiments.

Ultimately, an independent dark-matter measurement with NaI(Tl) is needed to resolve the question of whether the DAMA/LIBRA modulation signal can be attributed to dark matter. The SABRE collaboration will continue the effort to perform such a measurement and to further elucidate other open questions about the relevant parameters that affect it (like the quenching factors in Na and I). With continued effort, SABRE may be able to resolve the long-standing controversy surrounding the DAMA/LIBRA result.

Appendix A

Principles of Scintillation

This appendix is a review of scintillation mechanisms and effects that occur in both organic liquid scintillators and in inorganic scintillators like NaI(Tl). The information here follows the arguments in Knoll [29] with, where listed, some from Birk [36].

A.1 Radiation in matter

Charged particles in a medium lose their energy by interacting with the electrons in the material. The differential energy lost over a differential path length in the material, dE/dx , is known as the stopping power, which depends on the particle type. The classical expression describing this energy loss is known as the Bethe formula, and is as follows:

$$-\frac{dE}{dx} = \frac{4\pi e^4 z^2}{m_0 v^2} NB, \quad (\text{A.1})$$

where e is the electron charge, z is the integer charge of the particle, m_0 is the electron rest mass, v is the particle speed, N is the number density of the absorber atoms in the bulk material, and B is a parameter defined as:

$$B \equiv Z \left[\ln \frac{2m_0 v^2}{I} - \ln \left(1 - \frac{v^2}{c^2} \right) - \frac{v^2}{c^2} \right]. \quad (\text{A.2})$$

Z is the atomic number of the absorbing material, and I is an experimentally defined parameter intrinsic to the absorbing material. For non-relativistic particles, only the first term in Equation A.2 is significant.

Neutral particles like γ s and neutrons do not lose their energy continuously and smoothly like charged radiation. Rather, they undergo “catastrophic” events like scattering or absorption that cause them to lose part or all of their energy in short bursts. These events produce or release secondary charged radiation that deposits energy in the material in the manner described above.

Scintillation occurs when electrons are excited to higher energy states, or alternately when ionization results in recombination to excited energy states, and the electrons subsequently fluoresce by de-excitation to the ground state. Scintillation is triggered either directly by charged radiation traveling through the scintillating material and exciting electrons, or indirectly through the catastrophic event mechanism mentioned above. For example, a neutral particle like a neutron or a WIMP could scatter off the nucleus of an atom in the material, which recoils and causes ionization in the surrounding material.

A.2 Organic liquid scintillators

A.2.1 Scintillation mechanism

In organic liquid scintillators, energy imparted to individual molecules serves as the primary scintillation mechanism. In one large subset of liquid scintillators, the π -electrons have energy levels like those shown in Figure A.1. Incident radiation excites individual molecules to a variety of different energy states, most of which de-excite non-radiatively to the S_{10} excited singlet (spin 0) state¹. From there, the molecule can decay to the ground state via prompt fluorescence, releasing an optical photon with an energy of around 3–4 eV. Sometimes, the S_{10}

¹The first integer refers to the spin-excitation state, while the second refers to the vibrational state of the molecule. The S stands for singlet, while T stands for the triplet (spin 1) state.

state can transition through intersystem crossing to a much longer-lived triplet (spin 1) state, T_1 . From here, the molecule can decay directly with a longer wavelength and a long decay time through phosphorescence, or it can be re-excited to S_{10} and de-excite through delayed fluorescence. This latter process needs to be initiated by an energy deposition from another source. Often, this takes the form of bi-molecular interactions, in which two molecules in the triplet excited state interact to excite one to the S_{10} state and de-excite the other to the ground state.

A.2.2 Pulse-shape discrimination

Prompt fluorescence in organic scintillators is very fast, on the order of several ns. Delayed fluorescence and phosphorescence contribute to the total scintillation light by producing a long tail in scintillation pulses. Not only does the proportion of light that goes to delayed fluorescence and phosphorescence depend on the number of molecules in the triplet state, but the chance that delayed fluorescence will occur through a bimolecular interaction also depends on the triplet-state population. The concentration of triplet states along the ionization track in turn depends on the stopping power for the incident particle, dE/dx . For organic scintillators, the slow component fraction should be highest for particles with large stopping power, like α particles, and lowest for particles with lower stopping power, like electrons.

Neutrons interact primarily with atomic nuclei rather than electrons. In many organic scintillators, the primary scattering interaction is with hydrogen. As can be seen in Figure A.2, at the 10-keV level, protons have a higher stopping power than electrons, so neutrons should produce a slower scintillation pulse than electrons and γ s.




Figure A.1: Excitation scheme for organic scintillators with π -electronic structure. Molecules are excited into a number of different spin/vibrational states, which quickly de-excite to the S_{10} state non-radiatively. Prompt fluorescence occurs when the molecule de-excites to the lowest-energy singlet (spin 0) state. Longer-lived phosphorescence occurs when molecules transition to the triplet (spin 1) state, T_{10} and de-excite from there. Delayed fluorescence can also occur if the molecule is re-excited to S_{10} before it has a chance to phosphoresce. Figure from [29].





Figure A.2: Stopping power of electrons (Top) and protons (Bottom) in toluene, an organic scintillator similar to pseudocumene, the scintillator used by the DarkSide experiment in their neutron veto. Graphs from [137].

A.3 Inorganic scintillators like NaI(Tl)

A.3.1 Scintillation mechanism

Unlike organic liquid scintillators, the scintillation mechanism in inorganic crystals depends on the excited states of the crystal lattice itself, rather than for individual molecules. In a perfect crystal, electrons can be found in one of two energy regions: the low-energy valence band, where they are bound to their lattice sites, and the conduction band, where they are free to move throughout the crystal. Between them is a “forbidden” band of disallowed energy states. If a crystal is intentionally doped with impurities, several activator sites will become available in the forbidden band. A charged particle moving through the crystal would produce electron-hole pairs, exciting electrons to the conduction band. When this happens, the hole will drift to an activator site and ionize it if it has a lower ionization energy than the bulk material. Meanwhile, the electron is free to move to an ionized activator and fall into the site, producing a sometimes-excited neutral activator. In other cases, the electron and hole move together in a bound state called an exciton. The exciton can travel to an activator, exciting it. If this activator can transition to the ground state, an optical photon can be emitted. Because the scintillation wavelength is longer than what is required to excite an electron to the conduction band, doped crystal scintillators are transparent to their own scintillation light. Because the creation of the excited activator states takes some time, the scintillation times of inorganic scintillators are typically longer than organic scintillators—on the order of several hundred ns.

A.3.2 Ionization quenching

The theory of ionization quenching in NaI(Tl) is outlined here as described in [36]. The Murray-Meyer model of scintillation efficiency relates the scintillation efficiency of a crystal to the stopping power of different kinds of radiation. According to this model, the theoretical




Figure A.3: Scheme for the scintillation mechanism in NaI(Tl). Dopants create activator sites between the conduction and valence bands. Electron-hole pairs created by radiation traveling through the crystal travel to the activator sites in an excited state. If the transition to the activator ground state is allowed, the activator de-excites, releasing a scintillation photon. Figure from [29].

scintillation efficiency in NaI(Tl), dL/dE , as a function of dE/dx is given by:

$$\frac{dL}{dE} \propto \frac{\alpha K \frac{dE}{dx}}{1 + \alpha K \frac{dE}{dx}} f'_{xy}, \quad (\text{A.3})$$

where αK is an adjustable parameter based on the properties of the crystal and is best fit to $2 \text{ mg keV}^{-1} \text{ cm}^{-2}$ based on experimental results. f'_{xy} is the quantum efficiency of the energy transfer from the lattice to the activator, which is numerically independent of dE/dx when $dE/dx < 20 \text{ keV cm}^2/\text{mg}$, but is dependent on dE/dx for higher stopping powers. f'_{xy} is also dependent on the thallium dopant concentration. The predictions of the model for different thallium concentrations are shown in Figure A.4. The relationship of the scintillation efficiencies for various particles, according to this model, is complicated and energy dependent.

A.3.3 Pulse-shape discrimination

In NaI(Tl) and other inorganic crystal scintillators, a slow component of scintillation has been observed, which may result from the reactivation of “trap” states², [36]. The proportion

²Activator excited states that are forbidden from decaying to the ground state.





Figure A.4: Calculated scintillation efficiency in NaI(Tl) as a function of stopping power, dE/dx , for different concentrations of Tl by weight (C), using the Murray-Meyer model. Figure from [145].

of scintillation that goes to this slow component ($\tau=1.5 \mu\text{s}$ in NaI(Tl) [36]) is dependent on the stopping power for the incident radiation. Therefore, pulse-shape discrimination can be used to some extent in NaI(Tl) to distinguish between electron and nuclear recoils. In Figure A.5, the stopping power for electrons and Na ions in NaI are shown. For the latter, the plot was generated by retrieving the proton stopping power from [137] and scaling the x -axis for the difference in mass using Equation A.1. Because the stopping power only depends on the particle velocity and charge, and because, like a proton, a Na ion likely has a charge of $z = +1$, this method produces an approximate measure of the stopping power for Na ions, may be ionized by a nuclear recoil. A higher level of ionization would produce an even higher stopping power. The difference in stopping power between the electrons and the Na ions contributes to the differences in the pulse shapes.



Stopping power of Na ions in NaI, calculated




Figure A.5: (Top) Stopping power, dE/dx , for electrons in NaI(Tl) as a function of electron energy, from [137]. (Bottom) Stopping power for Na ions in NaI, derived from the graph for protons from [137] and a scaling following the Bethe formula (Equation A.1) as described in the text.

Appendix B

List of Abbreviations

Λ CDM	The Λ -cold-dark-matter model, a model parametrizing the amount of dark energy (in the form of a cosmological constant), dark matter, and baryonic matter in the universe.
ABC	Allegheny Bradford Corporation.
AC	Analog current.
ADC	Analog-to-digital converter.
a.m.u.	Atomic mass units.
ASME	American society of mechanical engineers.
BBN	Big-bang nucleosynthesis.
Bq	Unit describing a decay rate of 1 s^{-1} .
BR	Branching ratio.
CAEN	CAEN Technologies, an electronics company.
CDMS	Cryogenic dark-matter search, a series of cryogenic bolometric dark-matter detectors.
CDO	Compact dark objects.
CMB	Cosmic-microwave background.
COBE	Cosmic background explorer, a satellite mapping the CMB.

CoGeNT	Coherent germanium neutrino technology, a germanium-based detector.
COUPP	Chicagoland observatory for underground particle physics, a bubble-chamber dark-matter detector.
CP	Charge-parity.
cpd	counts per day.
CRESST	Cryogenic rare-event search with superconducting thermometers, a dark-matter experiment.
CRH	Clean-room H, the clean room on top of the DarkSide-50 water tank.
CVD	Chemical vapor deposition.
DAMA	Dark matter, a NaI(Tl) dark-matter direct-detection collaboration.
DAMA/LIBRA	Dark Matter: large sodium-iodide bulk for rare processes, a direct-detection dark-matter experiment.
DC	Direct current.
DMM	Dark-matter measurement mode.
dru	Unit equal to counts per day per kg of detector per keV of energy in the recoil-energy spectrum.
dru _{ee}	Unit equal to counts per day per kg of detector per keV _{ee} of energy in the recoil-energy spectrum.
EAG	Evans Analytical Group.
ETFE	Ethylene tetrafluoroethylene, a fluorine-based plastic.
ETL	Electron Tubes Limited, a phototube manufacturer.
eV	Electron-volt, the energy required to raise an electron by one volt of potential. Can also be used to describe a mass equivalent to $1 \text{ eV}/c^2$.

eV _{ee}	Unit of electron-equivalent energy, used to define the scintillation response of detectors in a way that is independent of the type of incident radiation.
eV _{nr}	Unit of nuclear-recoil-equivalent energy.
GDMS	Glow-discharge mass spectroscopy.
HEPA	High-efficiency particulate arrestance, a type of air filter.
HV	High-voltage.
ICP-MS	Inductively coupled plasma mass spectroscopy.
ID	Inner-diameter.
KIMS	Korea invisible mass search, a CsI(Tl) dark-matter experiment.
KK	Kaluza-Klein.
KMM	Potassium-measurement mode.
LAB	Linear alkyl-benzene, an organic liquid scintillator.
LEP	Large electron-positron collider at CERN, the Conseil Européen pour la Recherche Nucléaire.
⁷ Li(p,n) ⁷ Be	Li-p,n reaction ${}^7\text{Li}+p \rightarrow {}^8\text{Be}+n$.
LKP	Lightest Kaluza-Klein particle.
LNGS	Laboratori Nazionali del Gran Sasso.
LSP	Lightest supersymmetric partner particle.
LUX	Large underground xenon experiment, a liquid-xenon dark-matter experiment.
LVDS	Low-voltage differential signaling.
LY	Light-yield.
MACHO	Massive compact halo object.
m.w.e.	Meters-water-equivalent.
NAIAD	NaI advanced detector, a discontinued NaI(Tl) dark-matter direct-detection experiment.

NaI(Tl)	Thallium-doped sodium-iodide, a common scintillator.
ND	Neutron detector.
NDF	Number of degrees of freedom.
NIM	Nuclear instrumentation module, an electronics module standard.
OD	Outer diameter.
OFHC	Oxygen-free, high-conductivity. A type of low-radioactivity copper.
OPERA	Oscillation project with emulsion-tracking apparatus, a neutrino detector at LNGS.
PC	Pseudocumene, an organic liquid scintillator.
p.e.	Photoelectron.
PICASSO	Projet d'Identification de candidats supersymétriques sombres, a bubble-chamber dark-matter experiment.
PIXE	Particle-induced X-ray emission.
PMT	Photomultiplier tube.
PNNL	Pacific Northwest National Laboratory.
ppb	Parts per billion by mass.
PPC	P-type point contact, a type of ionization-charge-collecting detector.
ppm	Parts per million by mass.
PPO	(2,5) diphenyl-oxazole, a wavelength shifter.
ppt	Parts per trillion by mass.
PSD	Pulse-shape discrimination.
p.s.i.g.	Pounds per square-inch (gauge pressure).
PTFE	Polytetrafluoroethylene, commonly known as Teflon.
PVC	Polyvinyl chloride, a type of plastic.

QCD	Quantum chromodynamics.
QE	Quantum efficiency.
RMD	Radiation Monitoring Devices, Inc.
S1	Primary scintillation signal in a TPC detector.
S2	Secondary scintillation signal in a TPC detector that provides a measure of the ionization.
SABRE	Sodium-iodide with Active Background REjection.
SCENE	Scintillation efficiency of noble elements, experiment measuring the scintillation properties of argon and xenon.
SD	Spin-dependent.
SHM	Standard-halo model.
SI	Spin-independent.
SIDIS	SABRE-in-DarkSide insertion system.
SIMPLE	Superheated instrument for massive particle experiments, a bubble-chamber dark-matter detector.
SM	Standard model of particle physics.
SNOLAB	Sudbury neutrino observatory.
s.p.e.	single-photoelectron.
SPPUSA	Stainless Plate Products USA Inc.
SUSY	supersymmetry.
TMB	Trimethyl-borate.
TOF	Time of flight.
TPC	Time-projection chamber, a type of direct-detection detector that measures a scintillation and ionization signal.
VB	Vertical Bridgman method of crystal growth.
WIMP	Weakly-interacting massive particle.
WMAP	Wilkinson microwave anisotropy probe.

Appendix C

List of Symbols

α	The alpha particle.
β	The beta particle.
γ	The gamma particle.
Δ	Uncertainty or spread, E.g. $\Delta\theta/\theta$ is the fractional spread in the scattering angle.
ϵ	Detector efficiency.
$\hat{\epsilon}_1, \hat{\epsilon}_2$	Directional vectors describing the Earth's velocity around the Sun.
ζ	The center-of-mass rapidity in a relativistic kinematic calculation.
$\eta(v_{\min}, t)$	Mean inverse velocity at time t .
η_{10}	The baryon density normalized to the black-body photon density.
$\Theta(x)$	The Heaviside step function equal to 1 if $x \geq 0$ and 0 otherwise.
θ	Scattering angle.
Λ	The cosmological constant.
μ	The muon.
μ_χ	WIMP-nucleus reduced mass.

μ_p	WIMP-proton reduced mass.
ν	The neutrino.
$\tilde{\nu}$	The sneutrino, the supersymmetric partner to the neutrino.
ρ_{crit}	The critical density needed for a flat universe.
ρ_{DM}	The dark-matter density.
ρ_χ	The local WIMP density.
σ	Standard deviation.
$\sigma(E_R), \sigma(q)$	Velocity-independent WIMP-nucleus cross section.
σ_0	Velocity-independent WIMP-nucleon cross section at zero momentum transfer.
$\sigma_{0,SD,p}$	Velocity-independent, spin-dependent WIMP-proton cross section at zero momentum transfer.
$\sigma_{0,SD,n}$	Velocity-independent, spin-dependent WIMP-neutron cross section at zero momentum transfer.
$\sigma_{0,SI}$	Spin- and velocity-independent WIMP-nucleon cross section at zero momentum transfer.
σ_E	Energy resolution.
$\sigma_{SD,p}$	Velocity-independent, spin-dependent WIMP-proton cross section.
$\sigma_{SD,n}$	Velocity-independent, spin-dependent WIMP-neutron cross section.
σ_{SI}	Spin- and velocity-independent WIMP-nucleon cross section.
σ_v	WIMP velocity dispersion.
$\frac{d\sigma}{dE_R}(v, E_R)$	Energy- and velocity-dependent WIMP-nucleus interaction cross section.
$\frac{d\sigma}{d\Omega_{n,EL}}(\theta)$	Differential neutron elastic-scattering cross section at scattering angle θ .

$\left(\frac{d\sigma}{d\Omega p}\right)_0$	Differential cross section of the ${}^7\text{Li}(p,n){}^7\text{Be}$ reaction at 0° incident angle.
τ	Decay time or scintillation time.
$\Phi(E', E_1, E_2)$	Response function of a detector for quenched energy, $E' = Q(E_R)E_R$, between energies E_1 and E_2 .
Φ_n	Neutron flux.
χ	Symbol for the generic Weakly Interacting Massive Particle (WIMP). Or the neutralino.
χ^2	Weighted sum of the squared errors of a distribution, equal to $\chi^2 = \sum_n (x_n - \bar{x})/\sigma_x$.
Ω	Ohm (unit).
$\Omega_b, \Omega_{\text{DM}}$	Fraction of the critical density for baryons and dark-matter, respectively. Ω with other subscripts represents the fraction of the critical density for the species in the subscript.
A	The axion. Or mass number.
A_g	Avogadro's number, 6.02×10^{23} .
A_{ND}	Cross-sectional area of the neutron detectors in the SABRE Na quenching measurement.
a_0, a_1	Isoscalar/Isovector couplings.
c	The speed of light in a vacuum, equal to $\sim 3.0 \times 10^8$ m/s.
D	Deuterium. Or LiF depth in the SABRE Na quenching measurement.
E_R	Nuclear recoil energy.
$E_{R,\text{max}}(v)$	Maximum possible nuclear recoil energy for a scattering event with a WIMP-velocity v .
$F^2(E_R)$	Nuclear form-factor.
$F(q)_H$	Helm distribution.

F50	Pulse-shape discrimination parameter describing the fraction of the pulse integral occurring in the first 50 ns.
f_A	Axion decay constant.
$f_{p,n}$	WIMP coupling to the proton/neutron.
$f(\mathbf{v}), f(v)$	Normalized WIMP velocity/speed distribution.
$\tilde{f}(\mathbf{v})$	Normalized WIMP velocity distribution in the Galactic rest frame.
G	The gravitational constant, equal to $6.7 \times 10^{-11} \text{ N m}^2 \text{ kg}^{-2}$.
G_F	Fermi constant, equal to $(\sqrt{2}g)/(8m_W) = 1.166 \times 10^{-5} \text{ GeV}^{-2}$, where g is the weak-interaction coupling constant and m_W is the mass of the W boson that mediates the interaction.
\tilde{g}	The gravitino, the supersymmetric partner to the graviton.
H_0	The Hubble constant, describing the rate of expansion of the universe.
h	The Hubble parameter, where $H_0 \equiv 100h \text{ km s}^{-1} \text{ Mpc}^{-1}$.
\hbar	Reduced Planck constant, equal to $h/2\pi = 1.05 \times 10^{-34} \text{ J s}$, where h is the Planck constant.
^{nat}I	Natural iodine.
J	Total nuclear spin.
^{nat}K	Natural potassium.
L	Light yield.
\mathcal{L}_{eff}	Scintillation response of xenon to nuclear recoils.
M_\odot	The Solar mass.
m_χ	The WIMP mass.
m_{LiF}	LiF mass in a.m.u.
m_x	Mass of species x .
N_{esc}	Normalization factor in $\tilde{f}(\mathbf{v})$.

N_γ	Number of observed photons in a scintillation event.
N_{NaI}	Number of NaI molecules in the crystal in the SABRE Na quenching measurement.
n_χ	WIMP number density.
P_{E_R}	Probability of a neutron-scattering event with energy E_R in the NaI(Tl) detector in the SABRE Na quenching measurement.
p_{cm}	Center-of-mass momentum.
p_x	Momentum of particle x .
$Q(E_R)$	Ionization quenching factor.
Q	The Q-value of a nuclear decay, which is equal to the difference in mass between the reactants and products.
q	Momentum transfer, $q = \sqrt{2ME_R}$, where E_R is the recoil energy and M is the target nucleus mass.
R	R-parity, equal to $(-1)^{3B+L+2s}$, where B is the baryon number, L is the lepton number, and s is the spin.
$R(t, m_\chi, \sigma)$	Total rate in a dark-matter detector.
R_p	Rate of incoming protons in the Notre Dame quenching measurement.
$\frac{dR}{dE_R}(E_R, t, m_\chi, \sigma)$	Differential nuclear-recoil event rate.
r_n	Effective nuclear radius.
$S(q)$	Nuclear structure functions.
$S_0(E_R)$	Time-averaged WIMP-interaction rate in a detector as a function of recoil energy.
$S_m(E_R)$	WIMP-interaction modulation amplitude.
s	Nuclear skin-thickness.
T_k	Kinetic energy.
t_0	Peak WIMP modulation time (\sim June 1st).

$t_{1/2}$	Half-max time, the time at which a scintillation pulse signal reaches half its peak amplitude. This quantity is used as a measure of the start-time of the pulse.
V_{\oplus}	Earth's velocity relative to the Sun.
\mathbf{v}, v	WIMP velocity/speed.
\bar{v}_0	Most probable WIMP speed, around 235 km/s.
v_{\odot}	Sun's velocity relative to the Galactic rest frame.
v_{esc}	Galactic escape velocity.
$v_{\text{min}}(E_R, m_{\chi}, M)$	Minimum WIMP velocity required to produce a nuclear recoil with energy E_R , WIMP-mass m_{χ} , and nuclear mass M .
\mathbf{v}_{obs}	Velocity of the observatory with respect to the Galactic rest frame.
x	Unit-less factor equal to v_{min}/\bar{v}_0 .
x_1	Distance from LiF target to NaI(Tl) detector in the SABRE Na quenching measurement.
x_2	Distance from NaI(Tl) detector to a neutron detector in the SABRE Na quenching measurement.
Y	Measure of the ${}^4\text{He}$ density, equal to $\rho({}^4\text{He})/\rho_b$, where ρ_b is the baryon density.
y	Unit-less factor equal to v_{obs}/\bar{v}_0 .
Z	Atomic number.
z	Unit-less factor equal to v_{esc}/\bar{v}_0 .

Bibliography

- [1] J. Xu, E. S. F. Calaprice, S. Westerdale, *et al.*, “Scintillation efficiency measurement of Na recoils in NaI(Tl) below the DAMA/LIBRA energy threshold,” *Phys.Rev.C (accepted)* (2015) , [arXiv:1503.07212](https://arxiv.org/abs/1503.07212) [physics.ins-det].
- [2] E. Shields, J. Xu, and F. Calaprice, “SABRE: A New NaI(Tl) Dark Matter Direct Detection Experiment,” *Phys.Procedia* **61** (2015) 169–178.
- [3] F. Froberg, “SABRE in DarkSide project procedure: Description of insertion system,” tech. rep., Princeton University, 2015.
- [4] **Planck** Collaboration, P. Ade *et al.*, “Planck 2013 results. XVI. Cosmological parameters,” *Astron.Astrophys.* **571** (2014) A16, [arXiv:1303.5076](https://arxiv.org/abs/1303.5076) [astro-ph.CO].
- [5] J. Oort, “The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems,” *Bulletin of the Astronomical Institutes of the Netherlands* **6** (1932) 249.
<http://adsabs.harvard.edu/abs/1932BAN.....6..249O>.
- [6] F. Zwicky, “On the Masses of Nebulae and of Clusters of Nebulae,” *Astrophys.J.* **86** (1937) 217–246.
- [7] V. Rubin, N. Thonnard, and J. Ford, W.K., “Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 /R = 4kpc/ to UGC 2885 /R = 122 kpc/,” *Astrophys.J.* **238** (1980) 471.
- [8] D. Clowe, M. Bradac, A. H. Gonzalez, M. Markevitch, S. W. Randall, *et al.*, “A direct empirical proof of the existence of dark matter,” *Astrophys.J.* **648** (2006) L109–L113, [arXiv:astro-ph/0608407](https://arxiv.org/abs/astro-ph/0608407) [astro-ph].
- [9] G. Bertone, ed., *Particle Dark Matter: Observation, Models and Searches*. Cambridge University Press, 2010.
- [10] M. Markevitch, “Chandra observation of the most interesting cluster in the universe,” *ESA Spec.Publ.* **604** (2006) 723, [arXiv:astro-ph/0511345](https://arxiv.org/abs/astro-ph/0511345) [astro-ph].
- [11] NASA/CXC/CfA, “Astronomy picture of the day.”
<http://apod.nasa.gov/apod/ap060824.html>, Aug., 2004.

- [12] G. F. Smoot, C. Bennett, A. Kogut, E. Wright, J. Aymon, *et al.*, “Structure in the COBE differential microwave radiometer first year maps,” *Astrophys.J.* **396** (1992) L1–L5.
- [13] **WMAP** Collaboration, C. Bennett *et al.*, “First year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Preliminary maps and basic results,” *Astrophys.J.Suppl.* **148** (2003) 1–27, [arXiv:astro-ph/0302207](#) [astro-ph].
- [14] **WMAP** Collaboration, D. Spergel *et al.*, “First year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Determination of cosmological parameters,” *Astrophys.J.Suppl.* **148** (2003) 175–194, [arXiv:astro-ph/0302209](#) [astro-ph].
- [15] **Planck** Collaboration, P. Ade *et al.*, “Planck 2015 results. XIII. Cosmological parameters,” [arXiv:1502.01589](#) [astro-ph.CO].
- [16] **Planck** Collaboration, R. Adam *et al.*, “Planck 2015 results. I. Overview of products and scientific results,” [arXiv:1502.01582](#) [astro-ph.CO].
- [17] **Particle Data Group** Collaboration, J. Beringer *et al.*, “Review of Particle Physics: 22. Big-bang nucleosynthesis,” *Phys. Rev.* **D86** (2012) 010001.
- [18] M. Drees and G. Gerbier, “Mini-Review of Dark Matter: 2012,” [arXiv:1204.2373](#) [hep-ph].
- [19] F. Capela, M. Pshirkov, and P. Tinyakov, “Constraints on primordial black holes as dark matter candidates from capture by neutron stars,” *Phys.Rev.* **D87** (2013) no. 12, 123524, [arXiv:1301.4984](#) [astro-ph.CO].
- [20] K. Garrett and G. Duda, “Dark matter: A primer,” *arXiv* (January, 2011) .
- [21] V. Barger, D. Marfatia, and K. L. Whisnant, *The Physics of Neutrinos*. Princeton University Press, 2012.
- [22] R. Peccei, “The strong cp problem and axions,” *Lect. Notes Phys.* **741** (2008) 3–17, [arXiv:hep-ph/0607268](#) [hep-ph].
- [23] **Particle Data Group** Collaboration, J. Beringer *et al.*, “Review of Particle Physics: Axions and other similar particles,” *Phys. Rev.* **D86** (2012) 010001.
- [24] G. Steigman, B. Dasgupta, and J. F. Beacom, “Precise Relic WIMP Abundance and its Impact on Searches for Dark Matter Annihilation,” *Phys.Rev.* **D86** (2012) 023506, [arXiv:1204.3622](#) [hep-ph].
- [25] J. Reuter and M. Tonini, “Can the 125 GeV Higgs be the Little Higgs?,” *JHEP* **1302** (2013) 077, [arXiv:1212.5930](#) [hep-ph].
- [26] J. Bovy and S. Tremaine, “On the local dark matter density,” *Astrophys.J.* **756** (2012) 89, [arXiv:1205.4033](#) [astro-ph.GA].

- [27] M. Weber and W. de Boer, “Determination of the Local Dark Matter Density in our Galaxy,” *Astron.Astrophys.* **509** (2010) A25, [arXiv:0910.4272 \[astro-ph.CO\]](#).
- [28] J. Lewin and P. Smith, “Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil,” *Astropart.Phys.* **6** (1996) 87–112.
- [29] G. F. Knoll, *Radiation Detection and Measurement*. John Wiley and Sons Inc., 4th ed., 2010.
- [30] *Photomultiplier Tubes: Basics and Applications*. Hamamatsu Photonics K.K., 3rd ed., 2007.
https://www.hamamatsu.com/resources/pdf/etd/PMT_handbook_v3aE.pdf.
- [31] L. Baudis, “WIMP Dark Matter Direct-Detection Searches in Noble Gases,” *Phys.Dark Univ.* **4** (2014) 50–59, [arXiv:1408.4371 \[astro-ph.IM\]](#).
- [32] K. Freese, M. Lisanti, and C. Savage, “Colloquium: Annual modulation of dark matter,” *Rev.Mod.Phys.* **85** (2013) 1561–1581, [arXiv:1209.3339 \[astro-ph.CO\]](#).
- [33] K. Griest, “Effect of the Sun’s Gravity on the Distribution and Detection of Dark Matter Near the Earth,” *Phys.Rev.* **D37** (1988) 2703.
- [34] S. K. Lee, M. Lisanti, A. H. G. Peter, and B. R. Safdi, “Effect of Gravitational Focusing on Annual Modulation in Dark-Matter Direct-Detection Experiments,” *Phys.Rev.Lett.* **112** (2014) no. 1, 011301, [arXiv:1308.1953 \[astro-ph.CO\]](#).
- [35] M. Lisanti, “The effect of gravitational focusing on annual modulation,” 2013.
<https://indico.cern.ch/event/221841/session/8/contribution/35/material/slides/0.pdf>.
- [36] J. B. Birks, *The Theory and Practice of Scintillation Counting*. The Macmillan Company, New York, 1964.
- [37] J. Pradler, B. Singh, and I. Yavin, “On an unverified nuclear decay and its role in the DAMA experiment,” *Phys.Lett.* **B720** (2013) 399–404, [arXiv:1210.5501 \[hep-ph\]](#).
- [38] R. Bernabei, “Dark Matter Particles in the Galactic Halo,” [arXiv:1412.6524 \[astro-ph.CO\]](#).
- [39] R. Bernabei, P. Belli, S. d’Angelo, A. Di Marco, F. Montecchia, *et al.*, “Dark Matter investigation by DAMA at Gran Sasso,” *Int.J.Mod.Phys.* **A28** (2013) 1330022, [arXiv:1306.1411 \[astro-ph.GA\]](#).
- [40] **DAMA** Collaboration, R. Bernabei *et al.*, “First results from DAMA/LIBRA and the combined results with DAMA/NaI,” *Eur.Phys.J.* **C56** (2008) 333–355, [arXiv:0804.2741 \[astro-ph\]](#).
- [41] **DAMA** Collaboration, R. Bernabei *et al.*, “The DAMA/LIBRA apparatus,” *Nucl.Instrum.Meth.* **A592** (2008) 297–315, [arXiv:0804.2738 \[astro-ph\]](#).

- [42] M. Haffke, L. Baudis, T. Bruch, A. Ferella, T. Marrodan Undagoitia, *et al.*, “Background Measurements in the Gran Sasso Underground Laboratory,” *Nucl.Instrum.Meth.* **A643** (2011) 36–41, [arXiv:1101.5298 \[astro-ph.IM\]](#).
- [43] “Kyropoulos.” <http://rmdinc.com/kyropoulos/>.
- [44] V. Kudryavtsev, M. Robinson, and N. Spooner, “The expected background spectrum in NaI dark matter detectors and the DAMA result,” *Astropart.Phys.* **33** (2010) 91–96, [arXiv:0912.2983 \[hep-ex\]](#).
- [45] H. Lee, G. Adhikari, P. Adhikari, S. Choi, I. Hahn, *et al.*, “Pulse-shape discrimination between electron and nuclear recoils in a NaI(Tl) crystal,” [arXiv:1503.05253 \[astro-ph.IM\]](#).
- [46] R. Bernabei, P. Belli, A. Bussolotti, F. Cappella, V. Caracciolo, *et al.*, “Performances of the new high quantum efficiency PMTs in DAMA/LIBRA,” *JINST* **7** (2012) P03009.
- [47] R. Bernabei, “Physics and astrophysics with low background scintillators,” *Eur.Phys.J.* **ST163** (2008) 207–226.
- [48] R. Bernabei, P. Belli, F. Cappella, V. Caracciolo, R. Cerulli, *et al.*, “Comment on ‘On an unverified nuclear decay and its role in the DAMA experiment’ ,” [arXiv:1210.6199 \[hep-ph\]](#).
- [49] R. Bernabei, P. Belli, R. Cerulli, F. Montecchia, M. Amato, *et al.*, “On the investigation of possible systematics in WIMP annual modulation search,” *Eur.Phys.J.* **C18** (2000) 283–292.
- [50] M. Robinson, *The NAIAD Dark Matter Experiment: Backgrounds and Analysis*. PhD thesis, University of Sheffield, 2003. <http://www.hep.shef.ac.uk/research/dm/theses/thesis-robinson.pdf>.
- [51] F. Nozzoli, “Technical aspects and dark matter searches,” in *Topics in Astroparticle and Underground Physics*. 2009. <http://taup2009.lngs.infn.it/slides/jul3/nozzoli.pdf>.
- [52] **Particle Data Group** Collaboration, J. Beringer *et al.*, “Review of Particle Physics: 24. Dark Matter,” *Phys. Rev.* **D86** (2012) 010001.
- [53] R. Bernabei, P. Belli, V. Landoni, F. Montecchia, W. Di Nicolantonio, *et al.*, “New limits on WIMP search with large-mass low-radioactivity NaI(Tl) set-up at Gran Sasso,” *Phys.Lett.* **B389** (1996) 757–766.
- [54] J. Collar, “Quenching and channeling of nuclear recoils in NaI(Tl): Implications for dark-matter searches,” *Phys.Rev.* **C88** (2013) no. 3, 035806, [arXiv:1302.0796 \[physics.ins-det\]](#).

- [55] H. Chagani, P. Majewski, E. Daw, V. Kudryavtsev, and N. Spooner, “Measurement of the quenching factor of Na recoils in NaI(Tl),” *JINST* **3** (2008) P06003, arXiv:0806.1916 [physics.ins-det].
- [56] **CoGeNT** Collaboration, C. Aalseth *et al.*, “Search for An Annual Modulation in Three Years of CoGeNT Dark Matter Detector Data,” arXiv:1401.3295 [astro-ph.CO].
- [57] C. Aalseth, P. Barbeau, J. Colaresi, J. D. Leon, J. Fast, *et al.*, “Maximum Likelihood Signal Extraction Method Applied to 3.4 years of CoGeNT Data,” arXiv:1401.6234 [astro-ph.CO].
- [58] J. Collar, “Comments on ”A Dark Matter Search with MALBEK”,” arXiv:1407.2581 [physics.ins-det].
- [59] G. Angloher, M. Bauer, I. Bavykina, A. Bento, C. Bucci, *et al.*, “Results from 730 kg days of the CRESST-II Dark Matter Search,” *Eur.Phys.J.* **C72** (2012) 1971, arXiv:1109.0702 [astro-ph.CO].
- [60] **CRESST-II** Collaboration, G. Angloher *et al.*, “Results on low mass WIMPs using an upgraded CRESST-II detector,” *Eur.Phys.J.* **C74** (2014) no. 12, 3184, arXiv:1407.3146 [astro-ph.CO].
- [61] **CDMS** Collaboration, R. Agnese *et al.*, “Silicon Detector Dark Matter Results from the Final Exposure of CDMS II,” *Phys.Rev.Lett.* **111** (2013) no. 25, 251301, arXiv:1304.4279 [hep-ex].
- [62] **SuperCDMS** Collaboration, R. Agnese *et al.*, “Maximum Likelihood Analysis of Low Energy CDMS II Germanium Data,” *Phys.Rev.D* **D91** (2015) 052021, arXiv:1410.1003 [astro-ph.CO].
- [63] **LUX** Collaboration, D. Akerib *et al.*, “First results from the LUX dark matter experiment at the Sanford Underground Research Facility,” *Phys.Rev.Lett.* **112** (2014) 091303, arXiv:1310.8214 [astro-ph.CO].
- [64] C. Ghag, “The LUX direct dark matter search experiment,” arXiv:1412.4660 [hep-ex].
- [65] **XENON10** Collaboration, J. Angle *et al.*, “A search for light dark matter in XENON10 data,” *Phys.Rev.Lett.* **107** (2011) 051301, arXiv:1104.3088 [astro-ph.CO].
- [66] **XENON** Collaboration, S. Orrigo, “Direct Dark Matter Search with XENON100,” arXiv:1501.03492 [astro-ph.CO].
- [67] E. Del Nobile, G. B. Gelmini, P. Gondolo, and J.-H. Huh, “Update on Light WIMP Limits: LUX, lite and Light,” *JCAP* **1403** (2014) 014, arXiv:1311.4247 [hep-ph].

- [68] P. Sorensen, “A coherent understanding of low-energy nuclear recoils in liquid xenon,” *JCAP* **1009** (2010) 033, [arXiv:1007.3549](#) [[astro-ph.IM](#)].
- [69] N. Chen, Q. Wang, W. Zhao, S.-T. Lin, Q. Yue, and J. Li, “Exothermic isospin-violating dark matter after supercdms and {CDEX},” *Physics Letters B* **743** (2015) no. 0, 205 – 212.
<http://www.sciencedirect.com/science/article/pii/S0370269315001343>.
- [70] **PICASSO** Collaboration, S. Archambault *et al.*, “Constraints on Low-Mass WIMP Interactions on ^{19}F from PICASSO,” *Phys.Lett.* **B711** (2012) 153–161, [arXiv:1202.1240](#) [[hep-ex](#)].
- [71] **SIMPLE** Collaboration, M. Felizardo *et al.*, “The SIMPLE Phase II Dark Matter Search,” *Phys.Rev.* **D89** (2014) no. 7, 072013, [arXiv:1404.4309](#) [[hep-ph](#)].
- [72] **COUPP** Collaboration, E. Behnke *et al.*, “First Dark Matter Search Results from a 4-kg CF_3I Bubble Chamber Operated in a Deep Underground Site,” *Phys.Rev.* **D86** (2012) no. 5, 052001, [arXiv:1204.3094](#) [[astro-ph.CO](#)].
- [73] H. Lee, H. Bhang, J. Choi, S. Choi, I. Hahn, *et al.*, “Search for Low-Mass Dark Matter with CsI(Tl) Crystal Detectors,” *Phys.Rev.* **D90** (2014) no. 5, 052006, [arXiv:1404.3443](#) [[astro-ph.CO](#)].
- [74] S. Kim, H. Bhang, J. Choi, W. Kang, B. Kim, *et al.*, “New Limits on Interactions between Weakly Interacting Massive Particles and Nucleons Obtained with CsI(Tl) Crystal Detectors,” *Phys.Rev.Lett.* **108** (2012) 181301, [arXiv:1204.2646](#) [[astro-ph.CO](#)].
- [75] **Borexino** Collaboration, G. Bellini *et al.*, “Cosmic-muon flux and annual modulation in Borexino at 3800 m water-equivalent depth,” *JCAP* **1205** (2012) 015, [arXiv:1202.6403](#) [[hep-ex](#)].
- [76] R. Bernabei, P. Belli, F. Cappella, V. Caracciolo, R. Cerulli, *et al.*, “No role for muons in the DAMA annual modulation results,” *Eur.Phys.J.* **C72** (2012) 2064, [arXiv:1202.4179](#) [[astro-ph.GA](#)].
- [77] J. H. Davis, “Fitting the annual modulation in DAMA with neutrons from muons and neutrinos,” *Phys.Rev.Lett.* **113** (2014) 081302, [arXiv:1407.1052](#) [[hep-ph](#)].
- [78] D. Nygren, “A testable conventional hypothesis for the DAMA-LIBRA annual modulation,” [arXiv:1102.0815](#) [[astro-ph.IM](#)].
- [79] J. Klinger and V. Kudryavtsev, “Muon-induced neutrons do not explain the DAMA data,” *Phys.Rev.Lett.* **114** (2015) no. 15, 151301, [arXiv:1503.07225](#) [[hep-ph](#)].
- [80] J. M. Cline, Z. Liu, and W. Xue, “Millicharged Atomic Dark Matter,” *Phys.Rev.* **D85** (2012) 101302, [arXiv:1201.4858](#) [[hep-ph](#)].

- [81] Q. Wallemacq and J.-R. Cudell, “Dark antiatoms can explain DAMA,” *JCAP* **1502** (2015) no. 02, 011, [arXiv:1411.3178](#) [hep-ph].
- [82] R. Bernabei, P. Belli, F. Montecchia, F. Nozzoli, F. Cappella, *et al.*, “Investigating electron interacting dark matter,” *Phys.Rev.* **D77** (2008) 023506, [arXiv:0712.0562](#) [astro-ph].
- [83] R. Bernabei, P. Belli, F. Cappella, R. Cerulli, C. J. Dai, *et al.*, “Technical aspects and dark matter searches,” *J.Phys.Conf.Ser.* **203** (2010) 012040, [arXiv:0912.4200](#) [physics.ins-det].
- [84] D. Tucker-Smith and N. Weiner, “Inelastic dark matter,” *Phys.Rev.* **D64** (2001) 043502, [arXiv:hep-ph/0101138](#) [hep-ph].
- [85] E. Del Nobile, G. B. Gelmini, A. Georgescu, and J.-H. Huh, “Reevaluation of spin-dependent WIMP-proton interactions as an explanation of the DAMA data,” [arXiv:1502.07682](#) [hep-ph].
- [86] G. Barello, S. Chang, and C. A. Newby, “A Model Independent Approach to Inelastic Dark Matter Scattering,” *Phys.Rev.* **D90** (2014) no. 9, 094027, [arXiv:1409.0536](#) [hep-ph].
- [87] N. Anand, A. L. Fitzpatrick, and W. Haxton, “Model-independent Analyses of Dark-Matter Particle Interactions,” *Phys.Procedia* **61** (2015) 97–106, [arXiv:1405.6690](#) [nucl-th].
- [88] K. Freese, P. Gondolo, and H. J. Newberg, “Detectability of weakly interacting massive particles in the Sagittarius dwarf tidal stream,” *Phys.Rev.* **D71** (2005) 043516, [arXiv:astro-ph/0309279](#) [astro-ph].
- [89] J. Herrero-Garcia, T. Schwetz, and J. Zupan, “Astrophysics independent bounds on the annual modulation of dark matter signals,” *Phys.Rev.Lett.* **109** (2012) 141301, [arXiv:1205.0134](#) [hep-ph].
- [90] K. Ianakiev, B. Alexandrov, P. Littlewood, and M. Browne, “Temperature behavior of NaI (Tl) scintillation detectors,” *Nucl.Instrum.Meth.* **A607** (2009) 432–438, [arXiv:physics/0605248](#) [physics.ins-det].
- [91] P. Nadeau, M. Clark, P. Di Stefano, J.-C. Lanfranchi, S. Roth, *et al.*, “Sensitivity of alkali halide scintillating calorimeters with particle identification to investigate the DAMA dark matter detection claim,” *Astropart.Phys.* **67** (2015) 62–69, [arXiv:1410.1573](#) [physics.ins-det].
- [92] SNOlab. <https://www.snolab.ca/>.
- [93] **DarkSide** Collaboration, P. Agnes *et al.*, “First Results from the DarkSide-50 Dark Matter Experiment at Laboratori Nazionali del Gran Sasso,” *Phys.Lett.* **B743** (2015) 456–466, [arXiv:1410.0653](#) [astro-ph.CO].

- [94] A. Wright, P. Mosteiro, B. Loer, and F. Calaprice, “A Highly Efficient Neutron Veto for Dark Matter Experiments,” *Nucl.Instrum.Meth.* **A644** (2011) 18–26, arXiv:1010.3609 [nucl-ex].
- [95] **The SABRE Collaboration** Collaboration, J. Xu, “SABRE: A new NaI dark matter experiment,” 2014. http://www.pa.ucla.edu/sites/default/files/webform/jingke_xu.pdf.
- [96] N. Cherepy, “Scintillating ceramics.” email to Frank Calaprice, 02/25/2014, cc’ed to Emily Shields.
- [97] M. Robinson, P. Lightfoot, M. Carson, V. Kudryavtsev, and N. Spooner, “Reduction of coincident photomultiplier noise relevant to astroparticle physics experiment,” *Nucl.Instrum.Meth.* **A545** (2005) 225–233, arXiv:physics/0503044 [physics].
- [98] **DarkSide Material Qualification Group** Collaboration, M. Laubenstein, “R11065 pmt background.” DarkSide wiki.
- [99] **XENON** Collaboration, E. Aprile *et al.*, “Lowering the radioactivity of the photomultiplier tubes for the XENON1T dark matter experiment,” arXiv:1503.07698 [astro-ph.IM].
- [100] F. Froborg. Email to emily shields, 07/02/2015.
- [101] E. Hoppe, E. Mintzer, C. Aalseth, D. Edwards, I. Farmer, O.T., J. Fast, D. Gerlach, M. Liezers, and H. Miley, “Microscopic evaluation of contaminants in ultra-high purity copper,” *JRNC* **282** (2009) no. 1, 315–320. <http://dx.doi.org/10.1007/s10967-009-0241-1>.
- [102] **Majorana** Collaboration, N. Abgrall *et al.*, “The Majorana Demonstrator Neutrinoless Double-Beta Decay Experiment,” *Adv.High Energy Phys.* **2014** (2014) 365432, arXiv:1308.1633 [physics.ins-det].
- [103] J. Orrell. email to Frank Calaprice, 05/23/2014, cc’ed to Emily Shields.
- [104] Toray Plastics, Inc. www.toraytpa.com.
- [105] Cargille Labs. www.cargille.com.
- [106] Dow Corning Chemicals, Sylgard. <http://www.dowcorning.com/applications/search/products/Details.aspx?prod=01064291>.
- [107] Fiber Optic Center. <http://focenter.com/shop/uncategorized/angstromlink-al-3252-encapsulation-gel-1-52-2-pt/>.
- [108] F. Froborg, “SABRE in DarkSide requirements document,” tech. rep., Princeton University, 2015.
- [109] J. Benziger, “Environmental effects on the optical properties of pseudocumene,” tech. rep., Princeton Borexino Group, 1998.

- [110] E. Shields, “Adapting active veto technology to SABRE, a sodium-iodide detector for WIMP dark matter,” Pre-thesis, Princeton University, November, 2012.
- [111] E. Shields, S. Westerdale, and F. Calaprice, “A prototype neutron veto for dark matter detectors.” In progress.
- [112] American Society of Mechanical Engineers, “ASME Section 8: Rules for construction of pressure vessels.” <https://www.asme.org/shop/standards/new-releases/boiler-pressure-vessel-code-2013>.
- [113] American Society of Mechanical Engineers, “ASME Section 9: Welding and Brazing Qualifications.” <https://law.resource.org/pub/us/code/ibr/asme.bpvc.ix.2010.pdf>.
- [114] L. Desorgher, “G4radioactivedecay loaddecaytable question.” email to Emily Shields, 06/04/2015.
- [115] E. Shields and J. Xu, “Geant4 background simulations for sabre, phase 1,” tech. rep., Princeton University, July, 2013.
- [116] S. Cebrin, H. Gmez, G. Luzn, J. Morales, A. Toms, and J. Villar, “Cosmogenic activation in germanium and copper for rare event searches,” *Astroparticle Physics* **33** (2010) no. 56, 316 – 329. <http://www.sciencedirect.com/science/article/pii/S0927650510000642>.
- [117] J. Ziegler, “Terrestrial cosmic ray intensities,” *IBM J. Res. Develop.* **42** (1998) 117.
- [118] M. Gordon, P. Goldhagen, K. P. Rodbell, T. H. Zabel, H. H. K. Tang, J. M. Clem, and P. Bailey, “Measurement of the flux and energy spectrum of cosmic-ray induced neutrons on the ground,” *IEEE Trans. Nucl. Sci.* **51** (2004) 3427.
- [119] **DarkSide Material Qualification Group** Collaboration, P. Lombardi, “Veto pmt bacground.” email to Peter Meyers, 04/17/2014, DarkSide wiki.
- [120] N2 – Network of the EU Deep Underground Laboratories, “Status report for 2004 of LNGS infrastructure, logistics and scientific activities,” tech. rep., LNGS/ILIAS, 2004. http://www-lsm.in2p3.fr/ilias/n2/Docs/InfrasN2_LNGS_04.pdf.
- [121] G. Alimonti, C. Arpesella, M. Balata, G. Bellini, J. Benziger, S. Bonetti, B. Caccianiga, L. Cadonati, F. Calaprice, G. Cecchet, M. Chen, N. Darnton, A. de Bari, M. Deutsch, F. Elisei, F. von Feilitzsch, C. Galbiati, F. Gatti, M. Giammarchi, D. Giugni, T. Goldbrunner, A. Golubchikov, A. Goretti, T. Hagner, F. Hartmann, R. von Hentig, G. Heusser, A. Ianni, M. Johnson, M. Laubenstein, P. Lombardi, S. Magni, S. Malvezzi, J. Maneira, I. Manno, G. Manuzio, F. Masetti, U. Mazzucato, E. Meroni, M. Neff, L. Oberauer, A. Perotti, R. Raghavan, G. Ranucci, E. Resconi, C. Salvo, R. Scardaoni, S. Schnert, O. Smirnov, R. Tartaglia, G. Testera, R. Vogelaar, S. Vitale, and O. Zaimidoroga, “Light propagation in a large volume liquid scintillator,” *Nuclear Instruments and Methods in Physics Research*

- Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **440** (2000) no. 2, 360 – 371.
<http://www.sciencedirect.com/science/article/pii/S0168900299009614>.
- [122] Sigma Aldrich, “Products: 1,2,4-Trimethylbenzene.” <http://www.sigmaaldrich.com/catalog/product/aldrich/t73601?lang=en®ion=US>.
- [123] M. Wurm, F. von Feilitzsch, M. Goeger-Neff, M. Hofmann, T. Lachenmaier, *et al.*, “Optical Scattering Lengths in Large Liquid-Scintillator Neutrino Detectors,” *Rev.Sci.Instrum.* **81** (2010) 053301, arXiv:1004.0811 [physics.ins-det].
- [124] M. Johnson, *Scintillator Purification and Study of Light Propagation in a Large Liquid Scintillation Detector*. PhD thesis, Princeton University, 1998.
http://borex.princeton.edu/public-docs/theses/johnson_phd_by_chapter/.
- [125] A. Ianni, “Attenuation measurements in PC and TMB.” email to Frank Calaprice, 04/19/2011, forwarded to Emily Shields on 10/7/2011.
- [126] Hamamatsu Photonics, Japan, “R5912 data sheet.” <http://pdf1.alldatasheet.com/datasheet-pdf/view/62744/HAMAMATSU/R5912.html>.
- [127] N. Spooner, G. Davies, J. Davies, G. Pyle, T. Bucknell, *et al.*, “The Scintillation efficiency of sodium and iodine recoils in a NaI(Tl) detector for dark matter searches,” *Phys.Lett.* **B321** (1994) 156–160.
- [128] D. Tovey, V. Kudryavtsev, M. Lehner, J. McMillan, C. Peak, *et al.*, “Measurement of scintillation efficiencies and pulse-shapes for nuclear recoils in NaI(Tl) and CaF-2(Eu) at low-energies for dark matter experiments,” *Phys.Lett.* **B433** (1998) 150–155.
- [129] G. Gerbier, J. Mallet, L. Mosca, C. Tao, B. Chambon, *et al.*, “Pulse shape discrimination with NaI(Tl) and results from a WIMP search at the Laboratoire Souterrain de Modane,” *Astropart.Phys.* **11** (1999) 287–302.
- [130] T. Jagemann, F. Feilitzsch, and J. Jochum, “Measurement of the scintillation light quenching at room temperature of sodium recoils in NaI(Tl) and hydrogen recoils in {NE} 213 by the scattering of neutrons,” *Nucl.Instrum.Meth.* **A564** (2006) no. 1, 549–558.
<http://www.sciencedirect.com/science/article/pii/S0168900206005377>.
- [131] E. Simon *et al.*, “SICANE: a detector array for the measurement of nuclear recoil quenching factors using a monoenergetic neutron beam,” *Nucl.Instrum.Meth.* **A507** (2003) no. 3, 643 – 656.
<http://www.sciencedirect.com/science/article/pii/S0168900203014384>.
- [132] **SCENE** Collaboration, T. Alexander *et al.*, “Observation of the dependence on drift field of scintillation from nuclear recoils in liquid argon,” *Phys.Rev.* **D88** (2013) no. 9, 092006, arXiv:1306.5675 [physics.ins-det].

- [133] **SCENE** Collaboration, H. Cao *et al.*, “Measurement of Scintillation and Ionization Yield and Scintillation Pulse Shape from Nuclear Recoils in Liquid Argon,” arXiv:1406.4825 [physics.ins-det].
- [134] W. Aslam, V. Prestwich, and F. E. McNeill, “Thin target ${}^7\text{Li}(p, p'){}^7\text{Li}$ inelastic gamma-ray yield measurements,” *JRNC* **254** (February, 2002) 533–544.
- [135] C. A. Burke, M. T. Lunnion, and H. W. Lefevre, “ ${}^7\text{Li}(p, n_0){}^7\text{Be}$ angular distributions to $E_p=3.8$ MeV,” *Phys.Rev.C* **10** (Oct., 1974) 1299–1308.
- [136] D. Brown, “Evaluated Nuclear Data File (ENDF),” 2011. <http://www.nndc.bnl.gov/exfor/endl00.jsp>.
- [137] M. Berger, J. Coursey, M. Zucker, and J. Chang, “NIST:Stopping-Power and Range Tables for Electrons, Protons, and Helium Ions,” 8, 2005. <http://www.nist.gov/pml/data/star/>.
- [138] S. K. L. Sjue. <http://skisickness.com/2010/04/25/>.
- [139] W. Tan. email to Jingke Xu, 01/28/2015, forwarded to Emily Shields on 02/15/2015.
- [140] R. J. Barlow and C. Beeston, “Fitting using finite Monte Carlo samples,” *Comput.Phys.Commun.* **77** (1993) 219–228.
- [141] W. Moses, “Current trends in scintillator detectors and materials,” *Nucl.Instrum.Meth.* **A487** (2001) 123–128.
- [142] C. Savage, G. Gelmini, P. Gondolo, and K. Freese, “Compatibility of DAMA/LIBRA dark matter detection with other searches,” *JCAP* **0904** (2009) 010, arXiv:0808.3607 [astro-ph].
- [143] R. Catena and P. Ullio, “A novel determination of the local dark matter density,” *JCAP* **1008** (2010) 004, arXiv:0907.0018 [astro-ph.CO].
- [144] V. Bednyakov and F. Simkovic, “Nuclear spin structure in dark matter search: The Finite momentum transfer limit,” *Phys.Part.Nucl.* **37** (2006) S106–S128, arXiv:hep-ph/0608097 [hep-ph].
- [145] R. B. Murray and A. Meyer, “Scintillation response of activated inorganic crystals to various charged particles,” *Phys. Rev.* **122** (May, 1961) 815–826. <http://link.aps.org/doi/10.1103/PhysRev.122.815>.