

DEVELOPING ULTRA-LOW BACKGROUND
SODIUM-IODIDE CRYSTAL DETECTOR FOR
DARK MATTER SEARCHES

SUERFU

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: FRANK CALAPRICE

NOVEMBER 2018

© Copyright by Suerfu, 2018.

All rights reserved.

Abstract

Astronomical and cosmological observations have established that dark matter makes up about a quarter of the Universe. While the majority of underground experiments have failed to directly detect dark matter in laboratories, DAMA/LIBRA saw an annual modulation in event rate in an array of high-purity NaI (Tl) scintillation detectors, a phenomena expected from change in the relative velocity between the Earth and the dark matter halo as the Earth orbits around the Sun.

SABRE experiment aims to use ultra-high purity NaI (Tl) crystals and an active liquid scintillator veto to search for the annual modulation with sufficient sensitivity to confirm or refute the DAMA/LIBRA result. Through many years of R&D, we have successfully achieved crystal purity that has never been achieved before. We have also set up a liquid scintillator veto detector at Laboratori Nazionali del Gran Sasso (LNGS), Italy as part of a prototype detector to demonstrate the performance of low-background NaI (Tl) crystals.

In this thesis, I will present in detail the key steps and processes required to grow ultra-high purity NaI (Tl) scintillating scintillator crystals and preparation of SABRE prototype detector, including assembly of crystal detector module, commissioning of the veto detector, the trigger and the data acquisition system and a detailed analysis of projected background.

Although ultra-high purity NaI (Tl) crystals are developed as an essential component in resolving the annual modulation controversy, the techniques developed throughout this thesis have general implications for the development of low-background detectors.

Acknowledgements

I write this note of acknowledgements as the final touch of my thesis. Throughout my years at Princeton University, I have gained many skills and a lot of valuable hands-on experience. Such achievements would have been impossible without the help and support of many colleagues, friends and family members.

First and foremost I would like to express my deep and sincere gratitude to my advisor Prof. Frank Calaprice, for giving me a rare and valuable opportunity to get involved in every aspect of designing, constructing and commissioning a detector—from growing crystals to developing the trigger and writing data acquisition and analysis software. His belief—“An expert is someone who has made all the possible mistakes and learned not to make them again”—has become my motto and kept me motivated after all the failures over the course of my Ph.D. I would also like to thank Prof. Christopher Tully for introducing me to the world of particle physics and providing me a chance to work on a relic neutrino experiment as my experimental project, and Prof. Cristiano Galbiati for giving me an opportunity to work in Italy as a summer student. I am grateful to Prof. Peter Meyers who has generously agreed to be a second reader of this thesis. During a busy time of his, he went through this thesis in almost unbelievable detail and came up with many important suggestions.

I am also thankful to all the colleagues and collaborators with whom I have had the pleasure to work on this and other projects. In particular, Dr. Jingke Xu, Dr. Francis Froberg and Dr. Emily Shields introduced me to this project and helped me get started. Powder drying, filling and crystal detector assembly would have been impossible without countless hours of Dr. Masayuki Wada’s generous help. I am also grateful to Dr. Graham Giovanetti for helping setting up a glove box inside a clean-room at Princeton and proof-reading various documents and manuscripts. And I want to extend my gratitude to Dr. Chiara Vignoli, Antonio Di Ludovico, Donato

Orlandi, Dr. Paolo Montini, Dr. Giulia D’Imperio and Dr. Simone Copello for helping me during my visit to LNGS and stay in Assergi .

Life as the only graduate student in a nuclear and particle physics group is difficult and stressful. I have been saved so many times by many research assistants and visiting students. Joshua Pughe-Sanford has helped me on the early pyrolytic carbon coatings before automated DAQ system was ready. My proposal of water measurement by pressure rise in the cold trap was first proved by the tests done with the help of Grace Lawrence. Powder filling and drying for many crystal growths were done with the help of Bill Dix and Maddy Zurowski. In addition to making glassware used in many key experiments and tests, Jo Wo has came up with many interesting ideas for science demos and outreach which I enjoyed a lot. A special thanks goes to Winston Peloso for both being a great help in the lab and proof-reading this thesis.

The support from many technicians and scientists have been essential for the experiment. I have gained a lot of knowledge and skills through interactions with them, and have always enjoyed their patience and friendship. Mike Souza is a master glass blower with great sense of humor, and the No. 1 consultant for anything related to glass. He has not only produced countless glassware indispensable to the project, but also introduced me to the art and science of glass blowing. Steven Lowe has being a huge help in the student machine shop. Numerous parts and components were machined in the shop with his help, including all the custom parts used to make PyroMapper. Stanly Chidzik has not only helped with the design and manufacture of veto PMT base boards, but also generously allowed me to use his electronics shop and taught me various techniques in designing and making printed circuit boards, which almost got me addicted to circuit design and soldering. Dr. Norman Jarosik has taught me the basics of analog and digital circuits in PHY 557 and PHY 558. Without the knowledge and experience acquired in these courses, custom design of

microcontroller-based DAQ boards and development of the SABRE trigger firmware and the DAQ library would have been near impossible.

I would also like to acknowledge the kindness and help of all staff members at the Department of Physics. In particular, I am grateful to Darryl Johnson who handled all the difficult paperwork and customs documents for shipping, export and import. Research would be less enjoyable without the quick processing of purchase orders by Lauren Callahan and Ted Lewis, and a great stockroom maintained by Julio Lopez. Vinod Gupta and Sumit Saluja have been extremely helpful in setting up and maintaining DAQ computers at Princeton and planning data transfer for prototype detector operation.

Many resources and facilities provided by Princeton University have been vital for this thesis. In particular, the simulations in this thesis were conducted using the computational resources provided by PICSciE; XPS and SEM spectroscopy measurements were conducted with the help of Bert Harrop from the Department of Physics and Dr. Nan Yao, Paul Shao and John Schreiber from the IAC.

I have benefited a lot from many workshops and schools. I had a comprehensive exposure to Geant4 simulation for low-energy nuclear physics at the SWG 2015 workshop at Bratislava, and learned the fundamentals of trigger and data acquisition for particle physics at the ISOTDAQ 2017 workshop at Amsterdam. The ATI 2018 school at UCLA has been very useful in learning various dark matter theories and experiments.

Outside Princeton, I would like to thank Josh Tower and Guido Giampi at RMD for help with crystal growth and Brad McKelvey at Seastar for timely processing of ICP-MS samples. I have learned a lot about the difference between science and engineering from Jonathan Mellen and Ashley Mellen at the Mellen Company.

My life outside research is enriched by many friends and cohorts. In particular, Xin Xiang is a great housemate, officemate and travel buddy. I have been enjoying every

discussion with him regarding strange yet interesting ideas regarding both physics and life, and interesting experiments we have conducted at late hours in Jadwin Hall. Duyu Chen, Yu Shen and Zheng Ma have been great housemates that made my stay at Lakeside Apartments very comfortable. Duyu Chen is also a great lunch buddy with whom I have been exploring the Prospect House almost every single day. Yecheng Cao was among my first friends at Princeton, and has always been a reliable friend. The first few years would have been less exciting without his company.

I would like to extend my heartfelt gratitude to the Department of Physics for admitting me to the Ph.D. program, for the Joseph Henry Merit Prize and the Kusaka Memorial Prize and the Compton funds for travels and conferences. I have also enjoyed the glorious campus of Princeton University that made numerous great photographs possible.

Finally I am deeply grateful and forever indebted to my parents and grandparents for their unconditional love and everlasting faith, and for always being by my side. Without them, my Ph.D. and all other achievements would have remained daydreams.

To my parents and grandparents

Contents

Abstract	iii
Acknowledgements	iv
List of Tables	xiv
List of Figures	xvi
1 Introduction	1
1.1 A Brief Introduction to Dark Matter	1
1.1.1 Evidence of Dark Matter	1
1.1.2 Dark Matter Candidates	3
1.1.3 Laboratory Direct Search of Dark Matter	5
1.2 DAMA/LIBRA Annual Modulation	9
1.3 SABRE Experiment	10
1.4 Physical and Chemical Properties of NaI	14
1.5 Overview of This Thesis	15
2 Synthetic Fused Silica Crucible for Crystal Growth	17
2.1 Fused Quartz v.s. Synthetic Quartz	18
2.2 Crucible Design	21
2.3 Precision Cleaning of Crucible	21
2.4 Pyrolytic Carbon Coating	23
2.4.1 Coating Procedure	23

2.4.2	Measurement of Coating Thickness	29
2.4.3	Calculation of Attenuation Coefficient and Calibration by Scanning Electron Microscope	34
2.4.4	Thickness as Function of Position and Coating Duration	36
3	Precision Drying and Zone Refining of Astro-grade NaI Powder	39
3.1	Precision Drying of NaI Powder	40
3.1.1	The Apparatus for Precision Drying	40
3.1.2	The Drying Behavior of NaI Powder	42
3.2	Impurities in the Astro-grade NaI Powder	45
3.3	Zone Refining of NaI Powder	48
3.3.1	Principle	49
3.3.2	Impurity Distribution	52
3.3.3	Measuring Distribution Coefficient	53
3.3.4	Efficiency of Zone Refining	58
3.3.5	A New Apparatus for Zone Refining	59
4	Growth of Ultra-high Purity NaI (Tl) Crystal	61
4.1	Single Crystal Growth from Melt	61
4.2	Ultra-high Purity NaI (Tl) Crystal Growth	63
4.3	Major Crystal Growth Activities	64
4.3.1	NaI-014	66
4.3.2	NaI-015	69
4.3.3	NaI-025	71
4.3.4	Artoria	71
4.3.5	NaI-026	72
4.3.6	NaI-027	74
4.3.7	Gemini	75

4.3.8	NaI-031	76
4.3.9	NaI-032	78
4.3.10	NaI-033	79
4.4	Using Silicon Tetrachloride to Prevent Crystal Sticking	80
4.4.1	Test with Crystal-Grade NaI Powder	81
4.5	Conclusion on Crystal Growth	83
5	Crystal Detector Module	85
5.1	Crystal Polishing	85
5.2	Photomultiplier Tubes	89
5.2.1	Gain	89
5.2.2	Dark Rate	91
5.2.3	Afterpulsing	92
5.3	Crystal Enclosure	93
5.3.1	Design	93
5.4	Cleaning, Assembly and Test of Crystal Detector Module	95
6	Veto	99
6.1	Veto Vessel	100
6.2	Veto Photomultiplier Tubes	102
6.3	Shielding	103
6.4	Fluid Handling and Slow Control	105
6.5	Commissioning	105
6.5.1	Leak Detection at Princeton	106
6.5.2	Estimation of Veto Detector Light Yield	106
6.5.3	Commissioning at LNGS Hall B	109
6.5.4	Commissioning at LNGS Hall C	110

7	Trigger	112
7.1	Design and Implementation of Trigger Firmware	112
7.2	Measurement of Trigger Efficiency	117
8	Data Acquisition	120
8.1	Polaris	121
8.1.1	Design Philosophy	121
8.1.2	Software Architecture	121
8.1.3	Performance	125
8.2	DAQ for Powder Drying and Crucible Coating	125
8.2.1	PIC16-based DAQ Circuit	126
8.2.2	Monitoring NaI Powder Drying	126
8.2.3	Monitoring Pyrolytic Carbon Coating	127
8.2.4	Mapping Pyrolytic Carbon Coating	128
8.3	SABRE DAQ System	129
8.4	Event Reconstruction	131
8.4.1	Baseline Subtraction	132
8.4.2	Pulse Finding Algorithm	132
9	Background Simulation and Sensitivity	135
9.1	Avalanche of Auger Electron in Geant4	135
9.2	Background Calculation	136
9.2.1	Analysis of Background	138
9.2.2	Radioactive Isotopes in Low-background NaI (Tl) Crystal . . .	141
9.3	Monte-Carlo Simulation	150
10	Conclusions and Future Perspectives	152
A	Simulation of Impurity Distribution in Zone Refining	155

Glossary	157
Bibliography	177

List of Tables

1.1	Properties of NaI	15
2.1	Trace element concentration of fused quartz v.s. synthetic fused silica	18
2.2	Diffusivity constants of Li^+ , Na^+ and K^+ in crystalline quartz	19
2.3	Expected and observed increase in K and Li concentrations in NaI (Tl) crystal due to diffusion from fused quartz and synthetic fused silica crucibles	20
2.4	Comparison between PyroMapper and SEM	33
3.1	Procedure for precision drying of NaI powder	46
3.2	Impurities at different locations from a NaI ingot zone-refined by 25 passes	57
3.3	Calculated and measured distribution coefficients of LiI-NaI, NaCl-NaI, KI-NaI and RbI-NaI binary system	58
4.1	Major Crystal Growth Activities	66
4.2	ICP-MS measurements on samples from NaI-014 before and after surface cleaning	68
4.3	Concentration of impurities at different locations in NaI-014	69
4.4	Ultimate tensile strength of some high-temperature alloy	75
4.5	ICP-MS measurements of samples from NaI-031	77
4.6	ICP-MS measurements on samples from NaI-032	79

5.1	The effect of different polishing methods and reflector materials on the crystal light yield	87
5.2	Measured dark rates of R11065-20 PMTs at a gain of about 2×10^6 .	91
5.3	Probability of afterpulse for R11065-20 PMTs at 1700 V	93
6.1	Predicted veto detector light yield	108
7.1	Description of registers implemented in trigger firmware	117
9.1	Classification of background	138
9.2	Major β -emitters for a low-background NaI (Tl) crystal detector . . .	143
9.3	Major γ -emitters for a low-background NaI (Tl) detector	146
9.4	Major β -emitting surface contaminants for a low-background NaI (Tl) detector	148

List of Figures

1.1	The SABRE detector concept and the function of the liquid scintillator veto	12
1.2	Decay scheme of ^{40}K	13
2.1	A typical design of crucible for crystal growth	21
2.2	Setup for pyrolitic carbon coating	24
2.3	Furnace temperature and pressure inside crucible as functions of time in a typical pyrolitic coating procedure	26
2.4	Pyrolitic carbon coating on crucibles of various shapes and sizes	27
2.5	Thickness maps of pyrolitic carbon coating on 2'' crucibles without and with rotation.	28
2.6	CAD design of PyroMapper	31
2.7	Photo of PyroMapper	32
2.8	PyroMapper crucible rotation mechanism	33
2.9	Carbon coating thickness map for sample PyC011	34
2.10	SEM images of pyrolitic carbon coating used for calibration	35
2.11	Intensity of transmitted laser along the line connecting the two locations on PyC003 where SEM samples were taken	36
2.12	Average coating thickness as a function of position for various deposition times	37
2.13	Average thickness of carbon coating as a function of deposition time	38

3.1	Pyrolitic carbon-coated fused quartz crucible etched by molten NaOH	41
3.2	P&ID diagram and photo of the precision drying system	43
3.3	Pressure and temperature of cold trap and temperature of NaI powder over a number of water measurements	44
3.4	Water removal rate as functions of time for Astrograde NaI powder .	45
3.5	Response of vacuum gauge within first few minutes during warming up of cold trap	47
3.6	Oven temperature and immediate cold trap gauge reading as a function of time	48
3.7	XPS spectrum of Astro-grade NaI powder with and without black specks	48
3.8	An example phase diagram of a binary mixture	50
3.9	Illustration of normal gradient freeze	51
3.10	Illustration of zone refining	53
3.11	Impurity distribution after various numbers of zone passes	54
3.12	Phase diagram of NaCl-NaI and KI-NaI binary system	55
3.13	Contour plot of purification power as a function of number of zone passes and fraction of ingot kept	59
3.14	CAD design and photo of new zone refining apparatus	60
4.1	Frequently-used method for single crystal growth	62
4.2	Photo and approximate dimensions of NaI-014	67
4.3	^{137}Cs spectrum of NaI-014	67
4.4	Measured and fitted distribution of ^{39}K in NaI-014	70
4.5	Photo of NaI-015 and ^{137}Cs spectrum recorded with NaI-015	70
4.6	Photo of NaI-025	71
4.7	Photo of Artoria	72
4.8	Photo and ^{137}Cs spectrum of NaI-026	73
4.9	Growth of NaI-027 with Kanthal wire	74

4.10 Gemini- α and β	76
4.11 Photo and fitted ^{39}K distribution of NaI-031	77
4.12 Photo and fitted ^{39}K distribution of NaI-032	79
4.13 A setup to test the effectiveness of SiCl_4 in preventing NaI crystals from sticking to quartz crucibles	82
4.14 NaI crystallized with and without SiCl_4 treatment	83
5.1 Setup to test different surface treatments and reflector materials . . .	86
5.2 Photo of NaI (Tl) crystal with lapped and isopropyl alcohol-cleaned surfaces	88
5.3 Setup to characterize Hamamatsu R11065 PMTs	89
5.4 R11065-20 PMT gain as function of applied voltage	90
5.5 Average waveform corresponding to LED pulses	92
5.6 Degradation of NaI (Tl) crystal due to moisture over time	94
5.7 CAD drawing of detector module assembly	95
5.8 Surface of NaI-031 before and after sanding	97
5.9 Photo of the crystal detector module assembly	97
5.10 2D histograms of numbers of photoelectrons detected by the bottom and top PMTs in the crystal detector module	98
6.1 Photo of veto vessel at LNGS Hall B, and its Lumirror lining on the inner surface	100
6.2 CAD drawing and photo of 8" PMT, PMT holder and PMT can . . .	101
6.3 R5912 PMT pulses with ceramic and polypropylene capacitors on the divider boards	102
6.4 Hamamatsu R5912 veto PMT gain as function of applied positive high voltage	103
6.5 γ -ray flux after different shielding materials and thicknesses	104

6.6	Simplified schematics of the fluid handling system	106
6.7	^{137}Cs Compton edge measured with veto detector	107
6.8	Veto detector commissioning at LNGS Hall B	109
6.9	Majority coincidence of veto PMTs	110
6.10	Veto vessel commissioning in LNGS Hall C	111
7.1	Setup of CAEN V1495 trigger board	114
7.2	Implementation of crystal coincidence	115
7.3	Implementation of veto coincidence	115
7.4	Trigger output module	116
7.5	PMT signals and trigger signal	118
7.6	Two-channel coincidence trigger efficiency as a function of the distance between two coincidence windows	119
8.1	Polaris states, state transitions and the corresponding virtual functions	123
8.2	Two custom circuit boards for laboratory data acquisition	126
8.3	Hardware and software setup for recording NaI powder drying data .	127
8.4	Setup for monitoring pyrolytic carbon coating	128
8.5	Event visualization using local graphics library and remote Python script via internet interface module	129
8.6	The trigger and DAQ hardware setup for operating the crystal and veto detectors	130
8.7	Data flow between three key modules in <i>libsaber.so</i>	131
8.8	An example baseline of PMT signal of NaI crystal scintillation	133
8.9	Illustration of the pulse-finding algorithm	134
9.1	Difference between X-ray and Auger electron emission	137
9.2	Simulated and measured response of a ^{71}Ge proportional counter . . .	138

9.3	Probability of electrons recoiling with energy between 1-10 keV in a Compton scattering as a function of incident γ -ray energy	141
9.4	Simulated low-energy end of the spectrum of ^{129}I decay in NaI crystal	146
9.5	Energy deposit in the NaI (Tl) detector by β -decays of ^{14}C in the PTFE tape wrapping the crystal	149
9.6	Simulated background rate from various sources	150

Chapter 1

Introduction

1.1 A Brief Introduction to Dark Matter

1.1.1 Evidence of Dark Matter

Galaxy Cluster

One of the first clues of dark matter dates back to Fritz Zwicky's measurement of the mass of galaxy cluster using virial theorem [1]. According to the virial theorem, the gravitational potential energy is related to the average kinetic energy, and the galaxy cluster's total mass can be inferred by measuring the velocity dispersion of individual galaxies. Zwicky's measurement suggested that the total mass of the Coma Cluster was about one order of magnitude larger than observed luminous matter could account for. Around a similar time, the mass of another galaxy cluster, the Virgo Cluster, was estimated and similar conclusion was reached [2].

Galaxy Rotation Curve

The idea of "missing mass" in galaxies were challenged by many astronomers until the 1960s when Kent Ford and Vera Rubin measured the rotational speed of the Andromeda Galaxy as a function of radial position from the center of the galaxy with

improved measurement precision [3]. The result indicated that the rotational speed of outer part of the galaxy approached a constant value whereas Newtonian gravity predicted the speed would fall as $r^{-1/2}$ given that most of the galaxy's luminous mass is concentrated in the galactic bulge. Since then, the same phenomena had been observed with other galaxies [4]. The observed mass discrepancy can be explained either by modifying Newtonian gravity on the galactic scale [5], or by the existence of non-luminous dark matter in the galaxy.

Bullet Cluster

The existence of dark matter can also be inferred from the mass distribution during the collision of two galaxy clusters known as the “bullet cluster” [6]. The center of mass of luminous matter was calculated from photometry measurements on galaxies and X-ray imaging of inter-stellar medium, and the center of total mass was mapped using gravitational lensing. The result indicated that while luminous mass was distributed near the center due to the collision, the majority of gravitational mass simply passed through each other, a phenomenon that cannot be explained by modified Newtonian gravity.

Cosmic Microwave Background

The amount of dark matter can be quantitatively calculated by measuring the cosmic microwave background (CMB). The early Universe consisted of hot, dense plasma of electrons and baryons, and photons were essentially trapped due to short mean free path. As the Universe expanded and cooled down, electrons and protons recombined to form neutral hydrogen atoms, and photons decoupled from baryonic matter and became able to free-stream. At present, this remnant radiation has cooled down to microwave and is known as the cosmic microwave background.

The CMB temperature is mostly isotropic, with fluctuations about 10 parts per million. Such fluctuations, or anisotropies, are believed to be from structures already formed at the time of decoupling shifting the energies of the CMB photons. In the standard model of cosmology, structure formation was aided by dark matter that acted as gravitational wells for baryonic matter to fall into. Thus the pattern of the CMB temperature anisotropy is an indirect measure of ordinary and dark matter densities. According to the most recent data of Planck project, the Universe is composed of about 26% of dark matter, 5% of ordinary matter, and the rest dark energy [7].

1.1.2 Dark Matter Candidates

Astrophysical and cosmological observations have placed the following constraints on the nature of dark matter:

1. it must be electrically neutral, otherwise it would be observable on electromagnetic spectrum;
2. it must be stable enough on a cosmological scale, otherwise it would have decayed away;
3. it cannot be relativistic at the time of structure formation; and
4. if it is mostly produced in the Big Bang, its predicted density when it decoupled from the Big Bang plasma should match the density observed by the CMB measurements. This density is referred to as the relic density.¹

All known Standard Model particles are ruled out based on properties 1, 2 and 3, and various new models are proposed as dark matter candidates, including massive compact halo objects (MACHOs), sterile neutrinos, axions and weakly interacting massive particles (WIMPs).

¹This is not necessarily true if dark matter can be produced through other mechanisms.

MACHO

MACHO refers to ordinary astronomical objects that are too faint or impossible to be discovered by direct observation. These could include white dwarfs, neutron stars and primordial black holes. Recent evidences suggest that MACHO cannot account for all the dark matter observed in the galaxies [8][9].

Sterile Neutrino

Neutrinos are the only neutral and stable Standard Model particles. However, ordinary neutrinos were still relativistic at the time of structure formation, and therefore cannot account for the dark matter imprinted in the CMB patterns. If the sterile neutrino exists and has a much larger mass, it is a perfect dark matter candidate. Several years ago a few experiments reported an unidentified astrophysical X-ray line around 3.5 keV which is consistent with the existence of a sterile neutrino of mass around 7 keV [10][11]. However such interpretation is still controversial.

Axion

Axions are hypothetical pseudoscalar particles originally proposed to explain why charge-parity symmetry is not violated in quantum chromodynamics (QCD) [12]. Axions are predicted to mix with photons under external electromagnetic fields, and therefore most laboratory axion search experiments use resonant cavities to pick up the signals generated by axions passing through the cavities [13].

WIMP

WIMPs are massive, beyond-the-Standard-Model particles with weak-scale interactions. WIMPs are the most intensely searched dark matter candidate due to “the WIMP miracle”, a phenomena that a particle with electroweak-scale mass and self-annihilation cross section ($\langle\sigma v\rangle \approx 3 \times 10^{-26} \text{ cm}^3/\text{s}$) can naturally reproduce the

observed relic density via thermal freeze-out—as the Universe expands, thermally-produced dark matter particles decouple from the matter-radiation bath due to lower density and maintains their total number since then.

Experimental searches for WIMP dark matter generally make use of its hypothetical coupling to Standard Model particles: in accelerator searches, dark matter produced by collision of Standard Model particles is manifested as missing momentum in collision products; in astrophysical searches, dark matter particles annihilate with each other and produce an excess of Standard Model particles; and in laboratory searches, dark matter particles occasionally scatter off atomic nuclei, and the recoiling nuclei is detected with particle detectors.

1.1.3 Laboratory Direct Search of Dark Matter

The rate per nucleon at which WIMP particles scatter off atomic nuclei can be expressed as:

$$\frac{dR}{dE_R} = n_\chi \int_{v_{min}}^{v_{max}} v f(v) \frac{d\sigma(v, E_R)}{dE_R} d^3\vec{v}. \quad (1.1)$$

In this expression, $n_\chi = \rho_\chi/m_\chi$ is dark matter local number density. Dark matter local density is typically measured using either the Milky Way’s rotation curve or the vertical motion of nearby stars. Latest measurements indicate that the local dark matter density is likely between 0.2-0.9 GeV/cm³ [14], and 0.3 GeV/cm³ is often assumed in various calculations.

$f(v)$ is the dark matter velocity distribution in the local galaxy. Based on the galaxy rotation curve, dark matter density should be distributed as inverse square of the distance to the galactic center, and one natural choice of velocity distribution is Maxwellian assuming it is isotropic. In calculations, the most probable velocity in the distribution is set as the galactic rotation speed, which is on the order of 200 km/s [15].

The limits of the integral are the minimum dark matter velocity that can give a recoil energy of E_R and escape velocity, respectively. The escape velocity is estimated to be around 500-600 km/s based on surveys of the fastest stars in the Milky Way [16].

Velocity-dependence of interaction rate also gives rise to annual modulation: as the Earth orbits around the Sun which is moving at a constant speed relative to the dark matter halo, the relative speed between the Earth and the dark matter halo changes on a yearly basis, leading to change in the dark matter interaction rate in detectors [17]. Although processes such as cosmic muon flux [18] and temperature are known to exhibit seasonal modulation, annual modulation with the correct phase is expected to be unique to dark matter.²

The nature of the interaction is encoded in $d\sigma(v, E_R)/dE_R$, differential cross section that a dark matter particle with velocity v yields a nuclear recoil with energy E_R . In a scalar coupling, this term is proportional to the square of the number of nucleons while in an axial-vector coupling this term is related to the net spin of nucleons [19]. For collisions with high momentum transfers where the de Broglie wavelength is comparable to the size of the nucleus, the term will also depend on the nuclear form factor—the Fourier transform of nuclear density distribution [15].

Experimental Challenge

The difficulty of dark matter direct search comes from the low signal rate and overwhelming background rate. The kinetic energy of a recoiling nucleus induced by an incident dark matter particle with kinetic energy E_0 is given by

$$E_N = \frac{4m_\chi m_N}{(m_\chi + m_N)^2} E_0 \cos^2 \psi \quad (1.2)$$

where m_χ and m_N are the mass of dark matter particle and the mass of nucleus, respectively, and ψ is the angle between the direction of incidence and that of recoil.

²The performance of many types of detectors are temperature-dependent.

For a dark matter particle with mass ~ 100 GeV and a typical galactic velocity of ~ 200 km/s, the recoil energy is on the order of 10 keV, which is typically measured with ionizing radiation detectors.

Ionizing radiation (and nuclear recoil due to dark matter) can be measured with an ionization detector, a scintillator or a bolometer. In an ionization detector, an incident particle ionizes electrons in a semiconductor, liquid or gas, and these free electrons are collected by an electric field to produce electric signals. In a scintillator, an incoming particle excites electrons to excited states and photons are produced when these electrons return to the ground state. These photons are usually detected using photon detectors such as photomultiplier tubes (PMTs) or photodiodes. In a bolometer, an incident particle disrupts the crystal lattice and phonons (heat) are generated. Since less energy is required to produce a phonon than an ionization electron or a scintillation photon, bolometers can have lower thresholds and higher resolutions, but they cannot be operated at temperatures above a few Kelvin due to thermal noise.

Since these detectors are sensitive to not only nuclear recoils but also ordinary ionizing radiation, in the search of dark matter-induced recoils, these detectors can be overwhelmed by background from radioactive isotopes in the active detector medium, radioactivity in auxiliary detector components and external radiation. Assuming a weak-scale cross section (10^{-40} cm² per nucleon), a 100 GeV dark matter particle, and a moderate-size nucleon ($A \approx 30$), the expected event rate is less than 1 cpy/kg/keV.³ On the other hand, for 1 kg of naturally-occurring materials, typical radioactivity is in the mHz-Hz range. This would translate to 10^4 - 10^7 cpy/kg/keV. Therefore to be sensitive to dark matter-induced nuclear recoils, the background rate must be suppressed by many orders of magnitude.

³cpy is short for counts per year.

Methods to Reduce Background

Typically three strategies are employed to effectively suppress background: 1) shielding, 2) event-by-event rejection, and 3) improving radiopurity.

Shielding: radiation detectors on the Earth's surface are sensitive to cosmic ray. On the surface of the Earth, cosmic ray is dominated by muons with a flux of roughly $1 \text{ cm}^{-2} \text{ s}^{-1}$ [20]. Proton and neutron fluxes are orders of magnitude lower. Muons require several kilometers of water-equivalent shielding to effectively block. Therefore, dark matter detectors are often operated underground. Additional shielding made out of high-Z or proton-rich materials is often employed to block γ -rays and neutrons produced by radioactive isotopes in the rocks around the detector.

Event-by-event rejection: due to the small cross section, the probability that dark matter particles interact twice in a short time window is negligible. This makes it possible to reject background events by requiring anti-coincidence in an external active detector called a veto detector. The veto detector can tag and reject γ -emitting radioactive decays in the detector components as well as external radiations that manage to pass through passive shielding.

In addition to active veto, in certain types of detectors, the pattern of time response or relative amounts of electrons, photons and phonons produced in an interaction can be used to distinguish electron recoils from potential nuclear recoils produced in collision with dark matter.

Higher radiopurity: ultimately, the detector's radiopurity must be improved to fundamentally reduce background. This is achieved from two aspects: the detector material should be purified to reduce radioactive contaminants, and the elements that compose the detector should not have at large abundance long-lived radioactive isotopes that can contribute to background. In many cases, potassium, uranium and thorium are regarded as the most important radioactive contaminants because these isotopes have half-lives long enough to survive nucleosynthesis and they can be found

in many raw materials with varying concentrations. In terms of avoiding a detector medium with radioactive isotopes, a good example is ^{14}C : the abundance of ^{14}C in modern biological sources of carbon is about 10^{-12} . Therefore an organic scintillator made from such carbon sources will see roughly 200 ^{14}C decays per second per kg of carbon, whereas if petroleum is used as the carbon source, the ^{14}C activity can be suppressed by many orders of magnitude.

Since it is difficult to purify many common materials by oneself, in practice commercially-available high-purity products are often used. Therefore either detailed knowledge of manufacturing process or prior radiopurity assay are necessary in designing a low-background experiment.

1.2 DAMA/LIBRA Annual Modulation

Numerous detectors have been designed and constructed to search for dark matter-induced nuclear recoils. These detectors mainly include cryogenic solid state detectors, noble liquid scintillators and crystal scintillators. Among these experiments, the DAMA/LIBRA collaboration has been observing an annually modulating event rate of ~ 0.01 cpd/kg/keV around 3 keV on top of a background rate of roughly 1 cpd/kg/keV for over 14 years with a statistical significance of 9.3σ [21][22].

The DAMA/LIBRA experiment employs an array of 25 high-radiopurity NaI (Tl) scintillating crystals shielded by copper, lead, paraffin and concrete [21]. Software anti-coincidence between different NaI (Tl) crystals is imposed such that non-dark matter events can be vetoed. Each NaI (Tl) crystal, weighing about 10 kg, is coupled to two PMTs via Suprasil B light guides on both sides. PMTs coupled to the same crystal are operated in coincidence mode with a single photoelectron (p.e.) hardware threshold and a software threshold of 2 keV [21]. The NaI (Tl) crystals were purified and grown by Saint Gobain using the Kyropoulos method and platinum crucibles.⁴

⁴Details on powder purification, crystal growth and other protocols are not released to public.

Typical values of concentrations of $^{\text{nat}}\text{K}$, ^{238}U and ^{232}Th in the final crystals are 20 ppb, 0.7-10 ppt and 0.5-7.5 ppt, respectively.

The observed modulation has a period of 0.998 ± 0.002 yr and a phase of 144 ± 7 days, which are the same period and phase expected for modulation due to interaction with the dark matter halo [22]. Systematic effects such as changes in temperature, pressure, radon concentration and muon flux fail to account for the observed modulation amplitude and phase [23][24].

However, cross section and mass of the dark matter particle consistent with the observed modulation are strongly excluded by many other experiments, including those with leading sensitivities [25][26][27]. Although it is unlikely that the annual modulation observed by DAMA/LIBRA is caused by standard WIMP via standard coupling to nucleons, the inconsistency can still be explained by non-standard couplings (e.g. spin-dependent interaction with proton-only coupling [28]) or non-WIMP dark matter (e.g. mirror dark matter [29]). Moreover since the scintillation of NaI (Tl) is not yet fully understood, the modulation signal could be due to an undiscovered effect unique to inorganic crystal scintillators.

Whether the observed modulation is caused by non-trivial systematic factors, non-standard interactions of dark matter, or unknown effects of NaI (Tl) crystals, further investigation using the same detector medium is needed. Although other experiments have tried to observe the modulation, they have not being able to achieve crystal purity comparable to that of the DAMA/LIBRA experiment, resulting in backgrounds too high to search for annual modulation.

1.3 SABRE Experiment

Sodium-iodide with Active Background Rejection Experiment (SABRE) aims to design, construct and operate an array of NaI (Tl) scintillating crystals with 10-100 times

lower background compared to those used in the DAMA/LIBRA experiment to definitively test the annual modulation. The concept of the SABRE detector is illustrated in Fig. 1.1. Compared to predecessors, SABRE features

1. ultra-high purity NaI (Tl) crystals for lower intrinsic background,
2. direct coupling to high-quantum efficiency PMTs for lower threshold,
3. a liquid scintillator veto detector to suppress ^{40}K background around energy range of interest,
4. and twin detectors in the Northern and the Southern Hemispheres to eliminate seasonal effects.

At the current stage, SABRE is using the Astro-grade NaI powder from Sigma-Aldrich to grow NaI (Tl) crystals. The Astro-grade NaI powder, co-developed by Princeton University and Sigma-Aldrich, has impurities lower than the final NaI (Tl) crystals used in the DAMA/LIBRA experiment [30], and even higher purity is expected from additional purification via crystal growth.

In the development of the Astro-grade NaI powder, special care was taken to lower concentration of $^{\text{nat}}\text{K}$. Being in the same alkali group, potassium exists naturally in sodium to varying extents. ^{40}K , an radioactive isotope of potassium, decays by electron-capture with a branching ratio of about 11% to excited state of ^{40}Ar . A 1.46-MeV γ -ray is emitted as the excited nucleus returns to the ground state, and low-energy X-rays/ Auger electrons are emitted to fill the vacancy created by the electron capture (Fig. 1.2). When the high energy γ -ray escapes the NaI (Tl) crystal, the low-energy radiation appears as a peak near 3.2 keV, in the middle of the energy interval where the annual modulation is observed.

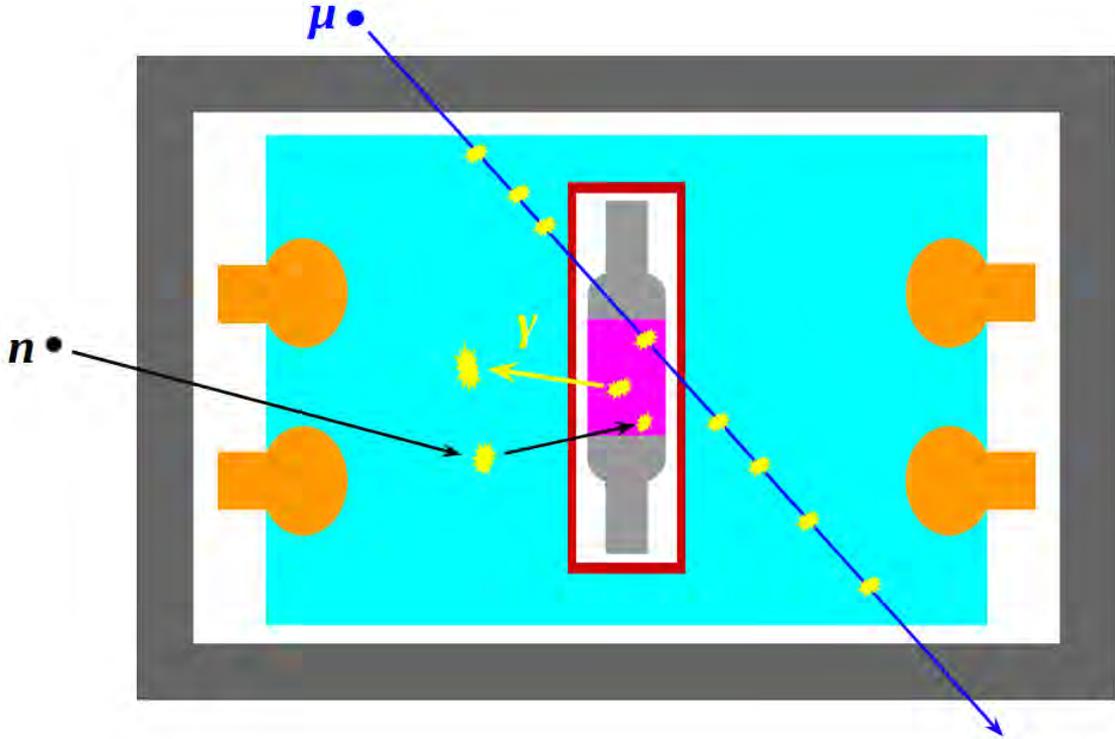


Figure 1.1: The SABRE detector concept and the function of the liquid scintillator veto. NaI (Tl) crystal (violet) is directly coupled to two PMTs (grey) and the assembly is placed inside a hermetically-sealed copper enclosure (brown). The crystal detector module is immersed in the liquid scintillator veto detector (teal). The scintillation light in the veto detector is detected by veto PMTs (orange). Another layer of passive shielding (dark grey) shields the veto detector from γ -rays and neutrons in the environment and the entire assembly is placed underground to suppress cosmic rays and prevent further cosmic activation. The liquid scintillator is effective in tagging internal crystal radioactivity with escaping γ -rays (yellow) as well as external γ -rays, neutrons (black) and muons (blue).

To enhance light-collection efficiency and to lower background, SABRE crystals will be directly coupled to Hamamatsu R11065-20 high-quantum efficiency, low-radioactivity PMTs.⁵

The liquid scintillator veto detector is designed to detect crystal radioactivity in coincidence with escaping γ -rays, such as the electron-capture decay of ^{40}K . In addition, veto detector also acts as a layer of active shielding against radioactivity in

⁵Quantum efficiency is typically around 25%.

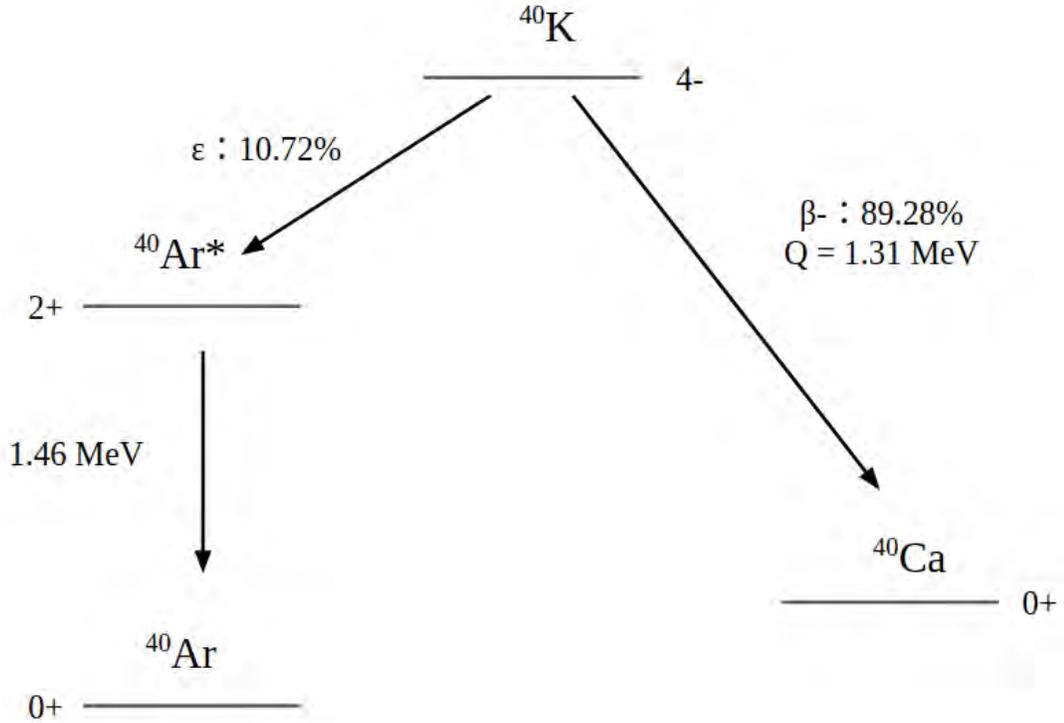


Figure 1.2: Decay scheme of ^{40}K . When ^{40}K decays by electron capture into excited state of ^{40}Ar , an atomic vacancy at either K-shell or L-shell is created. The atomic de-excitation is accompanied by a 1.46 MeV high-energy γ -ray as ^{40}Ar nucleus returns to the ground state.

the passive shielding materials and penetrating γ -rays, neutrons and cosmic muons. This principle is illustrated in Fig. 1.1.

The crystal detector module and the veto detector will be placed behind passive shielding to suppress external γ and neutron backgrounds, and the detector will be operated underground to suppress cosmic rays and cosmogenic isotopes.

SABRE will be deploying a similar detector in the Stawell Underground Physics Laboratory in Australia such that any modulation induced by seasonal effects will show a different phase in the Southern Hemisphere.

Currently the SABRE experiment is in the prototype phase aiming at demonstrating successful growth of ultra-high purity NaI (Tl) crystals and underground

operation of NaI (Tl) detector modules with liquid scintillator veto detector at Laboratori Nazionali del Gran Sasso (LNGS), Italy.

1.4 Physical and Chemical Properties of NaI

Before jumping directly into the details, a quick summary of physical and chemical properties of sodium iodide will be helpful in understanding the work of this thesis.

Sodium iodide (NaI) is an ionic compound formed by alkali element sodium and halogen element iodine. NaI powder is white in appearance and has density around 1.3-1.5 g/cm³. The melting point of NaI is around 661 °C. NaI is extremely soluble in water: 100 g of water can dissolve 184 g of NaI powder at room temperature. In addition, NaI is hygroscopic and deliquescent—it readily absorbs water in the ambient air and eventually dissolves in the absorbed water.

NaI crystal is a clear, brittle and relatively soft solid with density 3.67 g/cm³ and refractive index 1.85. Once doped with thallium (Tl), NaI crystal becomes an excellent scintillator with high light yield.⁶ NaI (Tl) is one of the most widely used scintillators in nuclear medicine, gamma spectroscopy and nuclear reactors.

Chemical properties of NaI is very similar to sodium chloride (NaCl), and NaI is very stable at room temperature. However, at elevated temperatures, NaI can react with water to form sodium hydroxide (NaOH) and hydrogen iodide (HI), or with oxygen to form sodium oxide (Na₂O) and iodine (I₂).

Table 1.1 provides a summary of these properties. Other materials are included as reference scales.

⁶Some evidences indicate that pure NaI crystal is also a good scintillator at cryogenic temperature.

Table 1.1: Properties of NaI. Some other materials are included as reference scales.

Properties	NaI	Note
Density (g/cm ³)	3.67	powder: 1.3-1.5
Melting point (°C)	661	
Hardness (Mohs)	2	graphite: 2, diamond: 10
Solubility (per 100 g solvent)	184 g (water), 43 g (ethanol), 26 g (1-propanol)	data at 25 °C, NaCl: 36 g (water)
Thermal exp. coeff. (ppm/°C)	47.4	quartz: 0.55, copper: 17
Refractive index	1.85	water: 1.3, glass: 1.4-1.6
Light yield (photons/keV)	38	Tl-activated
Peak scint. wavelength (nm)	425	Tl-activated
Primary decay time (ns)	230	Tl-activated

1.5 Overview of This Thesis

This thesis consists of two parts: the growth of ultra-high purity NaI (Tl) crystals and the development of the SABRE prototype detector.

Ultra-high Purity Crystal Growth

Large single crystals are typically grown from melt. However the ultra-high purity requirement places several extra restrictions. Chapter 2 establishes the crucible material for ultra-high purity crystal growth and the procedures for crucible cleaning and coating. Chapter 3 describes powder treatment prior to crystal growth. A procedure for precision-drying is discussed in detail, and the feasibility of additional purification by zone refining is investigated. Chapter 4 describes crystal growth activities

occurred in the past three years with emphasis on how the crystal growth methods have been refined through a series of failures.

Development of SABRE Prototype Detector

After the crystal is grown, it has to be cut and optically coupled to two PMTs and assembled into a detector module. Chapter 5 describes crystal surface cleaning, reflector material selection and the assembly and test of the crystal detector module. Chapter 6 describes the veto detector, fluid handling system and shielding. The chapter concludes with the commissioning of the veto detector in LNGS. Chapter 7 describes the design, implementation and verification of the trigger firmware required to operate the crystal detector module with the veto detector. Chapter 8 starts with the design and implementation of a general-purpose, modular data acquisition (DAQ) framework named Polaris. Since all DAQ software used throughout the studies described in this thesis is developed using the Polaris framework, the DAQ hardware and software for all tests and measurements are described in this chapter. A back-of-envelope analysis and Monte-Carlo simulations of the SABRE prototype detector background are described in Chapter 9.

Since the work described in this thesis spans multiple fields, including crystal growth, radiation detectors, electronic circuits, software development and simulations, the use of jargon is inevitable. In addition, in the context of this thesis some words and expressions have different meaning from general conventions. To make this thesis approachable to a broader audience, a glossary is compiled at the end to serve as an index and a brief dictionary of scientific terms.

Chapter 2

Synthetic Fused Silica Crucible for Crystal Growth

The crucible used to melt raw material in the growth of a single crystal is typically made out of high-temperature, chemically-inert materials such as amorphous quartz or platinum. However, the elevated temperatures at which raw materials melt often accelerate diffusion of impurities out from the crucible. As a result, to grow high-purity crystals the crucible material must be pure enough that diffusion of impurities in the crucible does not affect the crystal purity.

For this reason, amorphous quartz is a popular choice due to its good high-temperature performance and high chemical purity. Section 2.1 compares two forms of amorphous quartz (fused natural quartz and synthetic fused silica) and concludes that synthetic fused silica is the only choice in the growth of ultra-high purity NaI (Tl) crystals. A typical crucible shape is described in Section 2.2 and cleaning procedure is outlined in Section 2.3. Section 2.4 describes an improved method of applying a layer of pyrolytic carbon to the crucible as a physical barrier to prevent the crystal from sticking to the crucible.

2.1 Fused Quartz v.s. Synthetic Quartz

Quartz is readily available in very high purity. However, of different grades of quartz, synthetic fused silica has the highest purity. This is because synthetic fused silica is made from purified silicon-rich chemical precursors such as silicon tetrachloride (SiCl_4) while usual quartz is produced by melting naturally-occurring crystalline quartz which contains much higher levels of impurities.¹ As shown in Table 2.1, synthetic fused silica offers purity several orders of magnitude higher than fused quartz.

Table 2.1: Trace element concentration of GE 214 fused quartz [31] and Heraeus Suprasil 310 synthetic fused silica [32]. Concentrations are reported in ppm by weight.

Element	GE 214	Suprasil 310
Li	0.6	< 0.01
Na	0.7	< 0.05
K	0.6	< 0.01
Ca	0.4	< 0.05
Fe	0.2	< 0.02
Cu	< 0.05	< 0.01
Cr	< 0.05	< 0.005

The diffusion of alkali and alkaline earth elements in quartz was previously investigated by [33]. Although their study was done on crystalline quartz instead of amorphous quartz, the results can still indicate the amount of contamination introduced by using fused quartz. The temperature-dependent diffusivity constant is listed in Table 2.2.²

Assuming powder drying is carried out at 350 °C for 1 week (see Section 3.1.2) and crystal growth at 700 °C for 4 weeks, the corresponding diffusion length for

¹The word fused silica is often used interchangeably with fused quartz to refer to regular amorphous quartz glass. To avoid confusion, in this thesis quartz and fused quartz refer to regular quartz glass and synthetic fused silica is used to refer to quartz glass made from purified chemical precursors.

²The formula in Table 2.2 is verified for 300-500 °C. The result is extrapolated beyond the range.

Table 2.2: Diffusivity constants of Li^+ , Na^+ and K^+ in crystalline quartz [33]. The empirical formula is extrapolated to estimate temperature dependence above 500 °C.

Element	Diffusivity constant (cm^2/s)	Diffusion length (mm)	
		drying	growth
Li	$6.9 \times 10^{-3} \times e^{-10366/T}$	0.16	6.5
Na	$3.6 \times 10^{-3} \times e^{-12077/T}$	0.03	2.0
K	$0.18 \times e^{-15952/T}$	0.01	1.9

K^+ , given by $l = \sqrt{D\tau}$ where D is the diffusivity constant, is approximately 10 μm during drying and 1.9 mm during crystal growth. Taking the values in Table 2.1, a 4"-diameter fused quartz crucible will contribute roughly 30 ppb of K and synthetic fused silica will contribute only about 0.5 ppb averaged over the entire crystal.³ For powder drying, fused quartz and synthetic fused silica will contribute roughly 0.2 ppb and 0.003 ppb, respectively. This preliminary calculation indicates that diffusion of impurities from regular fused quartz during crystal growth will affect crystal purity while it is negligible for drying. On the other hand, diffusion of impurities from synthetic fused silica is negligible in both cases.

It is worthwhile noting that lithium has higher mobility than potassium and they both have similar concentrations in regular fused quartz. Applying the same calculation to Li^+ yields a diffusion length of 6.5 mm during crystal growth, about 3.6 times higher than that for K^+ . A naive scaling argument indicates that Li concentration should increase by approximately 110 ppb in a crystal grown in a regular fused quartz crucible. Therefore Li could be an indicator of contamination by diffusion. These results are summarized in Table 2.3.

This hypothesis is later definitively confirmed with two crystal growths, NaI-031 (Section 4.3.8) and NaI-032 (Section 4.3.9). The two crystal growths employed the same powder and cleaning procedure. The main difference was the crucible material:

³Diffusion through thin pyrolytic carbon coating on the crucible is not considered.

NaI-032 was grown in a synthetic fused silica crucible while NaI-031 was grown in a regular quartz crucible. No increase in K was observed in NaI-032, and Li was below detection limit of inductively coupled plasma mass spectroscopy (ICP-MS). On the other hand, K and Li increased by 34 ppb and 91 ppb, respectively, in NaI-031. The data are summarized in Table 2.3. The discrepancy between the expectation and the observation in Table 2.3 can be attributed to a series of assumptions and approximations made in the calculation.

Table 2.3: Expected and observed increase in K and Li concentrations in NaI (Tl) crystal due to diffusion of impurities from fused quartz and synthetic fused silica crucibles. The diameter of the crucible is 4" and mass is assumed to be 5 kg. NaI-031 was grown in a fused quartz crucible and NaI-032 in a synthetic fused silica crucible. All concentrations are in ppb.

Sample	K		Li	
	expected	observed	expected	observed
NaI-031	~ 30	14.3 ± 1.2	~ 110	91
NaI-032	0	$0.2^{+1.1}_{-0.2}$	0	0

Despite the benefit of fused quartz or synthetic fused silica, direct use is challenging due to the affinity of alkali halides to quartz. Molten alkali halides can wet, etch and diffuse into quartz.⁴ With the coefficient of thermal expansion differing by two orders of magnitude between NaI and quartz (Section 1.4), if NaI crystal sticks firmly to the crucible, either the crystal or the quartz crucible will crack upon cooling. To make the NaI crystal grow without sticking to the crucible, the crucible must be coated with a layer of pyrolytic carbon to physically separate the melt from the crucible wall.

⁴More precisely this is caused by a trace amount of alkali hydroxides in the alkali halides (Section 4.4).

2.2 Crucible Design

A typical crystal growth crucible is shown in Fig. 2.1. The body of the crucible is matched to the final diameter of the crystal, and the bottom is tapered into a cone shape with a tip (left end in the figure). During the cooling of the melt, this tip acts as a site for nucleation.

On the other side, the crucible body narrows down into a neck with an indent to hold a quartz plug for sealing the crucible. The crucible is sealed by heating the neck and the plug locally while pumping vacuum. Once the temperature exceeds the working temperature of quartz, the neck will collapse due to pressure difference and fuses with the plug inside, thus making the seal. The other end of the crucible (right end in the figure) is used to connect to vacuum system via a Pyrex compression fitting and a Viton O-ring.



Figure 2.1: A typical design of crucible for crystal growth. The crucible shown in this design has 4"-diameter, 1-ft-long body and 1"-diameter neck that expands to 2" at the end for vacuum fitting. The bottom end is tapered with a tip to provide a nucleation site during crystal growth. The neck has an indent to hold a quartz plug in place. The seal is achieved by fusing the neck with the plug using a torch.

2.3 Precision Cleaning of Crucible

Raw quartz tubings, especially high-purity ones used in the semiconductor industry, are manufactured according to strict standards such that bulk contaminations are below specified limits. However, subsequent manufacturing of crucibles from raw tubings typically takes place in glass shops where particulates and dusts in the air are

not regulated, and the crucible surface could be contaminated. To prevent contamination of the crystal, the crucible surface must be thoroughly cleaned prior to coating and crystal growth. The steps for initial cleaning prior to coating are outlined below followed by comments:

1. The crucible is first cleaned with a mixture of a few mL of Detergent 8, hydrofluoric (HF) acid and deionized water for 5-10 minutes at room temperature to remove grease and other contaminants on the surface.
2. The crucible is thoroughly rinsed with deionized water.
3. Mixture of 10% v/v nitric (HNO_3) acid and 5% v/v HF acid is used to leach out impurities at room temperature for 1 hour.⁵
4. The crucible is rinsed thoroughly with deionized water and dried under a HEPA-filtered hood.

Detergent 8 was chosen as the degreasing agent due to its relatively low potassium content compared to other detergents. After the cleaning, the crucible is sealed in a plastic bag inside a cleanroom and shipped for carbon coating.

After carbon coating, the crucible is rinsed with deionized water and leached with 6% v/v hydrochloric (HCl) acid at over 50 °C for 24 hours to remove residual metallic impurities on the surface. Subsequently the crucible is thoroughly rinsed with deionized water and dried.

The entire process is carried out in a cleanroom at SEASTAR CHEMICALS, Inc. (Seastar), and all chemicals used are ultra-high purity chemicals.

⁵ v/v represents percentage by volume. Similarly, m/m and mol/m represent mass by mass and mole-by-mass, etc.

2.4 Pyrolytic Carbon Coating

Pyrolytic carbon coating is achieved by heating hydrocarbons to above their pyrolysis temperatures to deposit atomic carbon on the surface. At the beginning, the coating was done by Sandfire Scientific, Ltd. (Sandfire), but the quality of the coating varied significantly. There are few references on pyrolytic carbon coating despite its general importance in crystal growth. Pyrolytic carbon coating of 1"-diameter quartz crucibles was previously studied by [34]. Their method works by flowing hexane vapor with a carrier gas through a probe into the crucible at 700 °C. Although we were able to reproduce their coating on 1"-diameter crucibles, coating on 2" or larger diameter crucibles failed at locations where diameter changes, such as the tip and neck of the crucible.

Building on their research, I developed an improved coating procedure using vacuum to transport vapor into the crucible instead of carrier gas. This method has been used to successfully coat quartz crucibles of various sizes and shapes. The coating procedure is outlined in Section 2.4.1. Section 2.4.2 describes PyroMapper, a novel apparatus that can automatically measure and map the thickness of pyrolytic carbon coating throughout the entire crucible. Its calibration is described in Section 2.4.3. The average thickness of coating as functions of axial position and coating time is measured with PyroMapper, and the results are discussed in Section 2.4.4.

2.4.1 Coating Procedure

Setup for Pyrolytic Carbon Coating

The setup, illustrated in Fig. 2.2, consists of a quartz crucible to be coated, a Pyrex vacuum compression fitting with two valves and a Pyrex KF25 flange, a scroll pump and a hexane-filled Erlenmeyer flask.

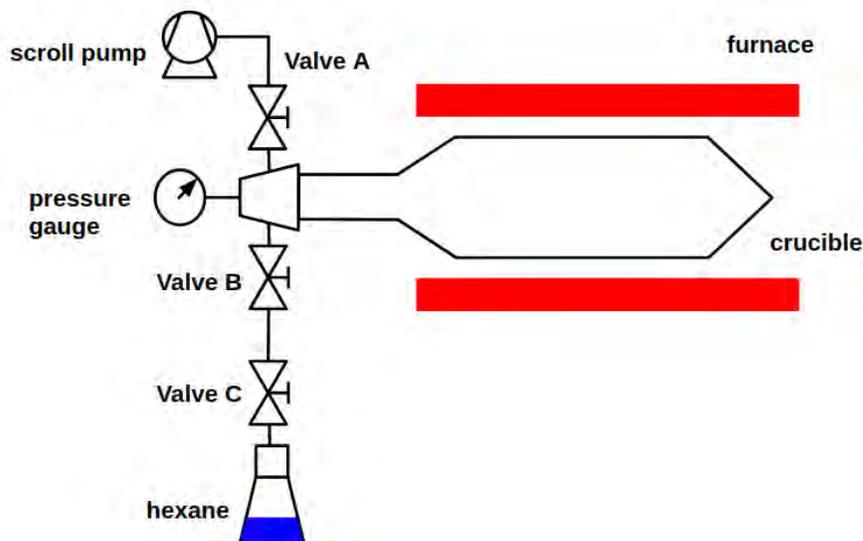


Figure 2.2: Setup for pyrolytic carbon coating. The fused quartz crucible is sealed with a Pyrex vacuum adapter via compression fitting and placed inside a horizontal furnace. The Pyrex adapter has one Pyrex KF25 flange and two valves connecting to a dry scroll pump and an Erlenmeyer flask containing hexane. The flask has another valve (Valve C) that prevents introduction of oxygen during the assembly and transportation. Pressure gauge is attached to the Pyrex KF flange on the adapter and fixed by a PTFE clamp.

Steps

The procedure for pyrolytic carbon is outlined below followed by comments.

1. **Crucible cleaning:** Prior to carbon coating, the inner surface of the crucible is rinsed with 5-10% v/v HF acid for 1-3 minutes followed by deionized water and ethanol to clean the surface.
2. **Hexane filling:** The Erlenmeyer flask is filled with hexane and attached to the Pyrex vacuum adapter. To keep away oxygen and particulates, the filling is performed inside a nitrogen-purged glovebox and during the transportation and assembly the valve on the flask (Valve C) is kept closed.
3. **Vacuum baking:** A dry scroll pump is attached to the other valve (Valve A) on the Pyrex vacuum adapter. The crucible is placed inside the furnace and

pumped to vacuum by opening Valve A and Valve B. The temperature of the furnace is raised to 600-800 °C and the crucible is vacuum-baked for at least 1 hour for degassing and removal of adsorped water.

4. **Pyrolysis:** The furnace temperature is lowered to 500 °C. Valve A is closed and Valve C is opened. Once the pressure reading stabilizes, Valve B and Valve C are closed and the hexane is detached. The temperature of the furnace is raised to 800 °C at a rate of approximately 20 °C/min. Around 700 °C, the onset of pyrolysis can be observed. To achieve azimuthally more uniform coating, the crucible is incrementally rotated by one full revolution at 45° steps in a manner that equal amount of time is spent at each orientation. The crucible is again pumped down to vacuum after 10-40 minutes after reaching 700 °C, depending on the final thickness desired. This step can be repeated multiple times to increase the thickness of coating.
5. **Annealing of carbon coating** The crucible is vacuum-baked at 1100 °C for 5 minutes to anneal the coating. Subsequently the crucible is naturally cooled down and taken out of the furnace.

The furnace temperature and crucible pressure as functions of time in a typical coating process is shown in Fig. 2.3. The data is acquired with a PIC16F1788 micro-controller and DAQ software developed using the Polaris framework (see Section 8.2.3 and 8.1).

Discussion

In the initial vacuum-baking of the crucible (Step 3), pumping for 1 hour at over 800 °C is adequate in removing adsorbed water from quartz according to [35] and [36]. A longer pumping time does not negatively impact the coating, but baking at higher

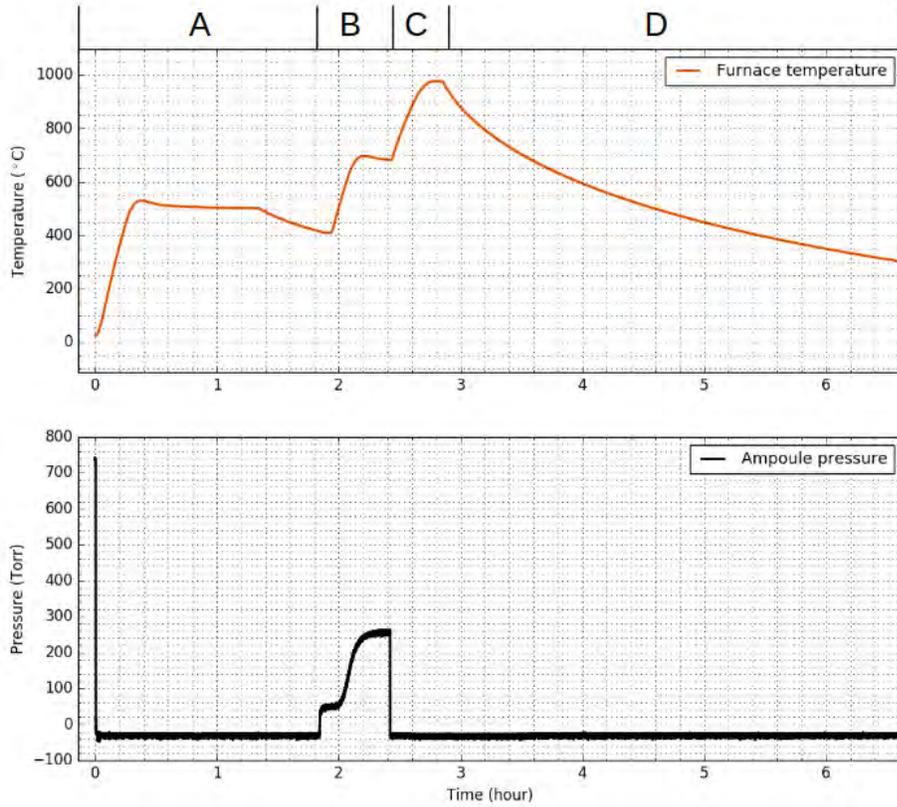


Figure 2.3: Furnace temperature and pressure inside crucible as functions of time in a typical pyrolitic coating procedure. A: vacuum-baking of the crucible. B: pyrolysis of hexane and deposition of carbon. C: annealing of the carbon coating. D: cooling of the crucible.

temperatures for extended amount of time might cause quartz to devitrify, and thus should be avoided.

The coating procedure described by [34] has an additional step of annealing of the quartz crucible prior to coating. By testing various coating conditions, we found that this step of quartz annealing had no effect on coating and therefore was omitted from the procedure, unless the quartz crucible has not already been annealed to get rid of the stress induced during manufacturing.⁶

The furnace temperature ramp-up and ramp-down rate was found to be non-critical except for the pyrolysis phase (Step 4) where it affects total duration of

⁶Extended annealing at over 1000 °C will accelerate devitrification of fused quartz and failure of carbon coating.

pyrolysis.⁷ Since there is not a sharp turn-on of pyrolysis as temperature rises, the deposition duration is counted from the moment temperature reaches 700 °C, where the hexane vapor develops a fog-like appearance.

Compared to the procedure based on carrier gas in [34], the method described in this work has less restrictions on crucible shape and size (Fig. 2.4). This is likely due to the hexane vapor more uniformly filling the crucible volume under vacuum whereas when a carrier gas is used, the flow pattern could change drastically when the size or shape of the crucible changes.



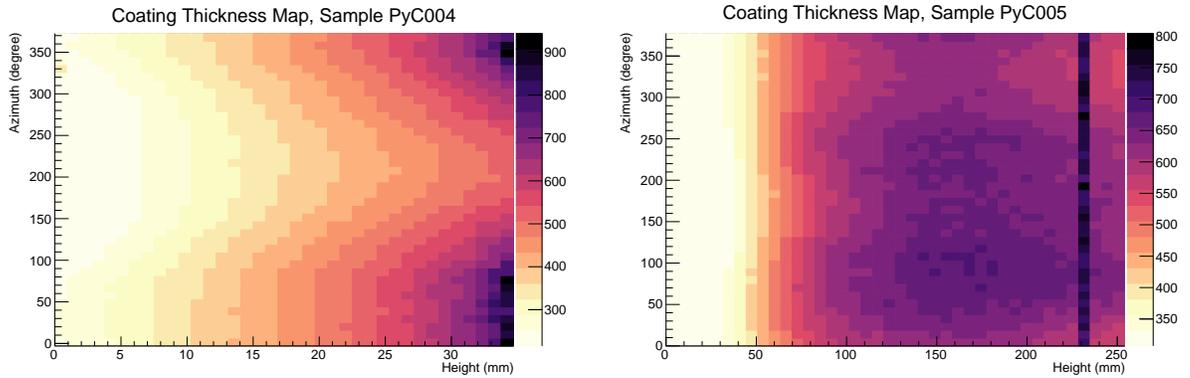
Figure 2.4: Pyrolytic carbon coating on crucibles of various shapes and sizes. From top to bottom the crucible diameters are 4", 2" and 1.5" body with 1" neck. A 12"-long ruler is included to indicate the scale. Photo courtesy of Hushu and Wuhana.

In addition, the overall risk of contamination is lower: there is no need to filter and purify carrier gas, and a simpler setup and assembly procedure lower the risk of contamination due to residues and particulates in the flow system.

The rotation of crucibles in Step 4 is found to be crucial in yielding azimuthally uniform coating. In the current study, all crucibles are rotated by one full revolution

⁷Fused quartz has very small thermal expansion coefficient and it is almost impossible to thermally shock a thin-wall quartz crucible.

at either 45° or 90° steps such that equal time is spent at each orientation. Fig. 2.5 shows the difference between carbon coatings without and with rotation. However, even with 45° rotation, the change of coating thickness in the azimuthal direction is still visible. If better uniformity is required, the crucible can be mechanically rotated continuously with a motor during pyrolysis.



(a) During deposition, the crucible is not rotated.

(b) During deposition, the crucible is rotated by one full revolution at 45° step.

Figure 2.5: Thickness maps of pyrolytic carbon coating on 2" crucibles without and with rotation. The thickness maps are obtained with PyroMapper (Section 2.4.2).

Precautions for Coating Synthetic Fused Silica

The procedure described in Section 2.4.1 was developed using regular fused quartz crucibles but the procedure works equally well on synthetic fused silica crucibles. Coating synthetic fused silica for ultra-high purity crystal growth requires additional precautions, but due to the relatively simple setup and coating procedure, the requirements can be easily met by the changes described below:

- High-purity hexane with low metallic residues upon evaporation should be selected as the carbon source. If ^{14}C is a concern, hexane obtained from petroleum should be used.
- The crucible should be precision-cleaned according to the procedure described in Section 2.3 instead of the steps outlined earlier in this section.

- All glass components including the Pyrex adapter and the Erlenmeyer flask should be cleaned with high-purity acetone, deionized water and ethanol thoroughly to remove surface particulates. At the current stage, precision-cleaning is not done since these components do not touch the inner surface of the crucible.
- The assembly of the crucible, vacuum adapter and hexane-filled Erlenmeyer flask should be carried out inside a cleanroom or in a clean glovebox and the valve to pump (Valve A) should be kept closed until the pump is running.

2.4.2 Measurement of Coating Thickness

The thickness of pyrolytic carbon coating is a crucial parameter. According to [34], the coating is expected to be a few hundred nanometers thick and is often measured using scanning electron microscope (SEM). However SEM measurements are destructive, time-consuming and expensive. To solve these challenges, I designed and constructed a new apparatus named PyroMapper to measure the thickness of carbon coating on quartz crucibles by optical attenuation.

Principle of Operation

PyroMapper works by scanning a beam of laser across the crucible surface and measuring and recording the intensity of transmitted laser. At each location as the laser light passes through the coating, the intensity is exponentially attenuated as $\exp(-d/\tau)$, where d is the thickness of the coating and τ the attenuation coefficient. Therefore the ratio of intensities between any two locations is given by

$$\frac{I_i}{I_j} = e^{-\frac{d_i - d_j}{\tau}} \quad (2.1)$$

where I_i is the intensity of the transmitted light at location i .

Rearranging Eq. 2.1, τ can be computed by measuring the intensities at two different locations with known thicknesses of coating.

$$\tau = -\frac{d_i - d_j}{\log(I_i/i_j)}. \quad (2.2)$$

After τ is known, the thickness throughout the entire scanned surface can be mapped out by Eq. 2.1, provided the absolute thickness at any of the scanned locations is known. Initial measurements of absolute thickness and calculation of attenuation coefficient can be achieved by inspecting under SEM samples taken from different locations on the crucible where intensities of transmitted laser have been pre-measured.

Once a reference thickness has been established on a crucible, this crucible can be used as a calibration source. Future laser attenuation measurements can be calibrated by repeating measurements on this calibration source at locations where coating thicknesses have been measured.

Description of PyroMapper

The CAD design of PyroMapper is shown in Fig. 2.6, and a photo of PyroMapper is shown in Fig. 2.7. Scanning is achieved with two stepper motors (M1 and M2): stepper motor M1 moves the light source and the photodetector back and forth via a guided rail along the crucible, and stepper motor M2 rotates the crucible via timing pulleys and a timing belt in the azimuthal direction. The rotation mechanism is shown in detail in Fig. 2.8. In the measurements described in this thesis, a 5-mW laser diode is used as the light source, and an off-the-shelf phototransistor is used as the photodetector. The stepper motors, the laser and the phototransistor are controlled by a custom-designed PIC16F1788-based microcontroller circuit (Section 8.2).

Data is acquired using the *pyromap* software package developed using the Polaris framework (see Section 8.1).

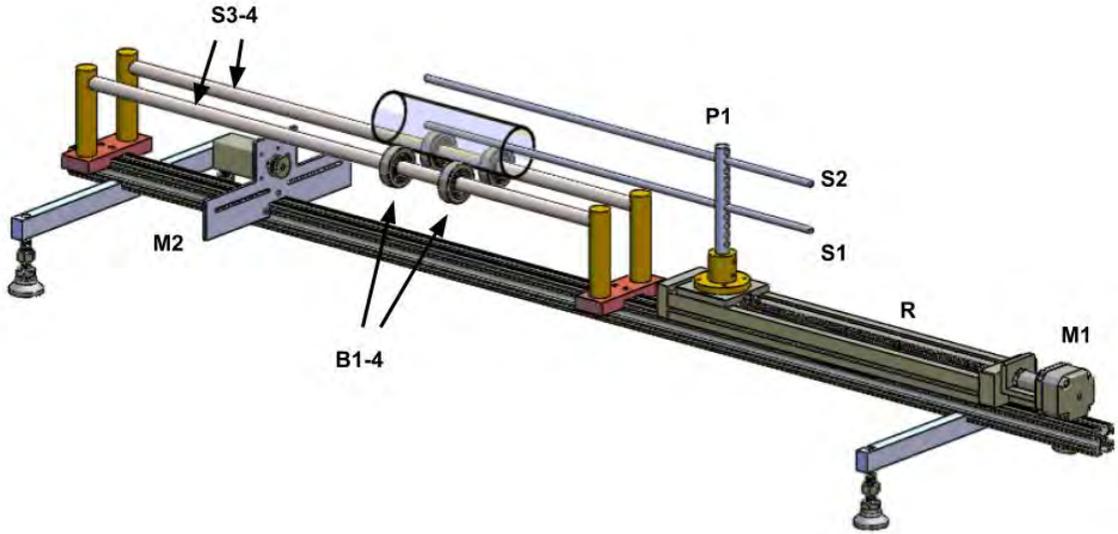


Figure 2.6: A photodetector and a light source are attached to the ends of the shaft S1 and S2, respectively, inserted through holes on the shaft pillar (P1) and fixed with set screws. The shaft pillar has an array of holes to allow positioning of each shaft at different height. The shafts S1 and S2 and the pillar P1 are moved together by a stepper motor (M1) and a guided linear rail (R). The quartz crucible is positioned between two precision shafts S3 and S4. Four ball bearings (B1-4) are used to support the crucible and to assist crucible rotation. A second stepper motor (M2) is used to rotate the crucible via timing pulleys and a timing belt. For clarity of image, the light source, photodetector and timing pulleys and timing belt on M2 are not shown.

A typical thickness map obtained with PyroMapper is shown in Fig. 2.9 and also previously in Fig. 2.5. The coating in Fig. 2.9 is done on 45×48 fused quartz crucible with 20-minute deposition time.⁸

Comparison with SEM

The comparison between PyroMapper and SEM in the context of measuring coating thickness on fused quartz crucibles is summarized in Table 2.4: SEM is expensive,

⁸Quartz tubings are usually described by ID×OD in millimeters.

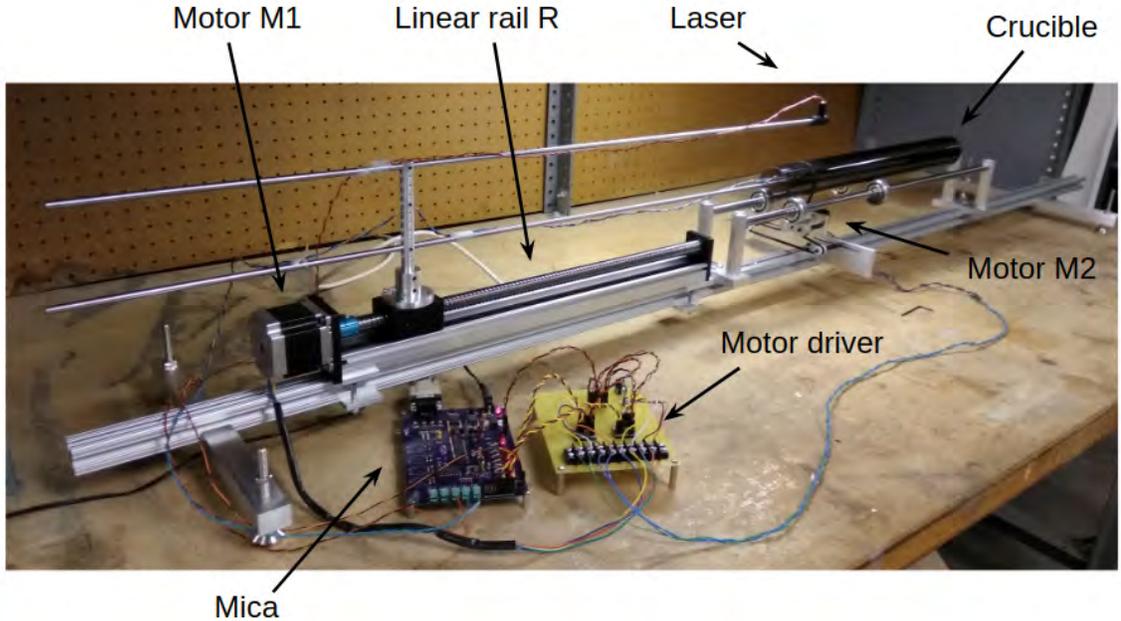


Figure 2.7: Photo of PyroMapper. Custom circuit board Mica (Section 8.2.1) is used to control the light source and photodetector as well as to provide power and control signals to stepper motors.

general-purpose equipment that requires special training and manual operation while PyroMapper is low-cost, special-purpose equipment. Once hardware and software are properly configured, the measurement is automated and sample placement is very simple. SEM samples are typically a few cm in dimension, and sample preparation is a destructive process for crucibles several feet in length. Locally, SEM has a very high spatial resolution, but part of the crucible and coating is inevitably lost when cutting the crucible to appropriate sample size, which affects global resolution. On the other hand, the local resolution for PyroMapper is mainly limited by the size of the laser beam but scanning can be performed at sub-mm steps. Since light is exponentially attenuated by the coating, there is a limit on the thickness that can be measured. With a 5-mW laser diode and an off-the-shelf phototransistor, the maximum measurable thickness is around 700 nm. The limit can be further improved by using either a more sensitive photodetector or a higher-power laser, but the gain is marginal. However, SEM can measure a much wider range of coating thickness.

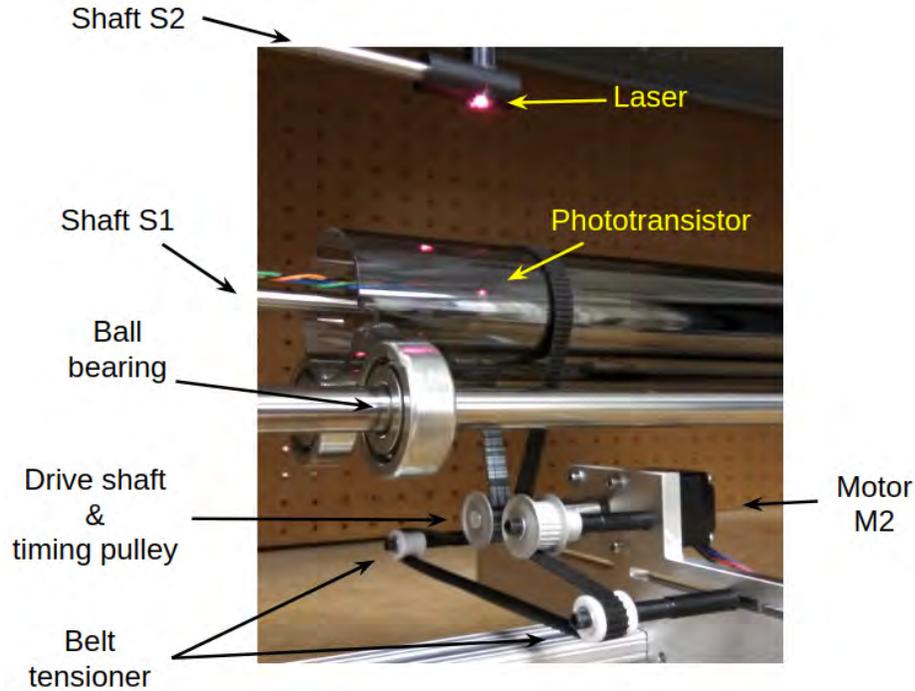


Figure 2.8: PyroMapper crucible rotation mechanism. A timing belt goes around the crucible, two timing pulleys and two PTFE tensioner pulleys (belt tensioners). The positions of the belt tensioners are adjustable such that certain tension can be maintained in the timing belt for crucibles of different diameters. One of the timing pulleys is directly driven by motor M2. As the motor shaft rotates, the timing belt also moves at a fixed rate determined by its pitch, and the friction between the crucible and the timing belt causes the crucible to rotate as well. Crucible rotation is further assisted by four ball bearings B 1-4.

Table 2.4: Comparison between PyroMapper and SEM in the context of measuring thickness of pyrolytic carbon coating on fused quartz crucibles.

	SEM	PyroMapper
Purpose	general-purpose	special-purpose
Cost	expensive	cheap (< \$200)
Operation	manual special training required	automatic no training needed
Sample preparation	destructive	non-destructive
Resolution (local)	\sim nm	100 μ m-mm
Resolution (global)	> 1 cm	100 μ m-mm
Measurement range	nm-mm	< 1 μ m

Coating Thickness Map, Sample PyC011

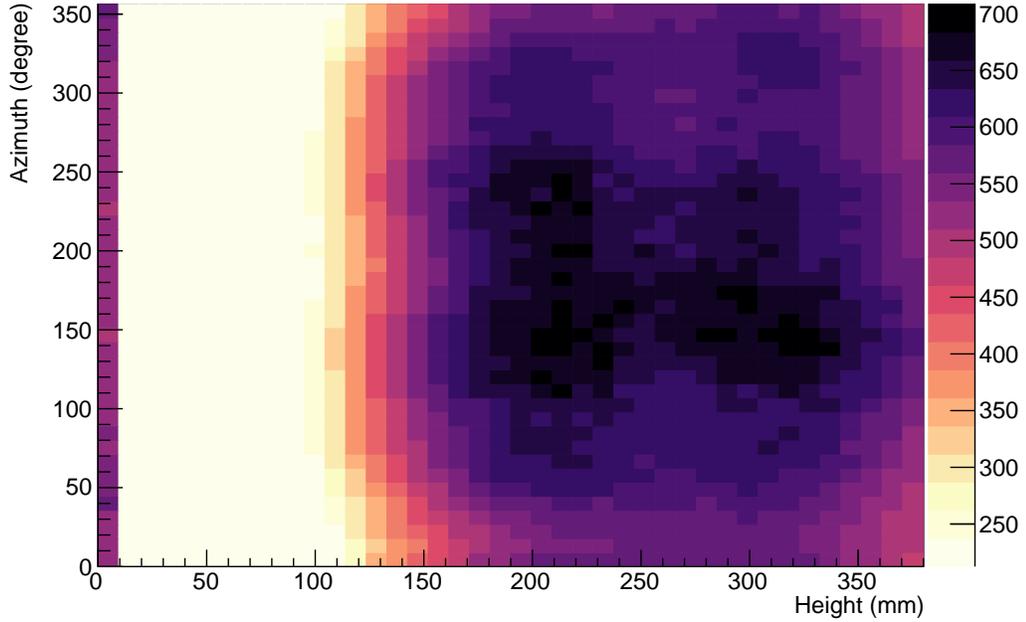


Figure 2.9: Carbon coating thickness (in nm) map for sample PyC011 (20 min deposition time). Carbon coating thickness is plotted in cylindrical coordinate. The height is measured 43 mm away from the tip, with positive direction towards the tip. The leftmost dark column corresponds to the azimuth driving belt.

2.4.3 Calculation of Attenuation Coefficient and Calibration by Scanning Electron Microscope

A 2''-diameter, 4-cm-long segment cut from a coated quartz crucible (PyC003) is used to measure the attenuation coefficient τ and to calibrate the apparatus with a reference thickness. A 5-mW, 632-nm laser diode and a phototransistor are used as the light source and the photodetector, respectively. After optical measurement on this segment, two small SEM samples (approximately 5 mm by 5 mm) are cut away from each end and inspected under SEM. The thicknesses measured using SEM are compared with laser transmission data to calculate τ .

The SEM image of these two samples are shown in Fig. 2.10a and Fig. 2.10b. The intensity of transmitted laser along the line connecting these two SEM sample

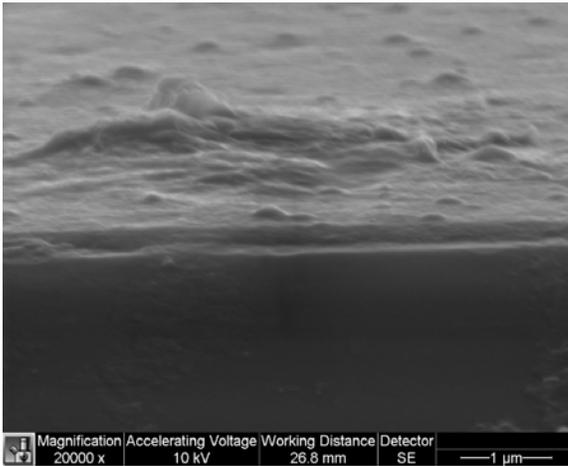
locations is shown in Fig. 2.11. From Eq. 2.2 the attenuation coefficient of 632-nm laser in pyrolytic carbon is calculated to be

$$\tau = \frac{395 - 216}{\log(0.638/0.0034)} \text{ nm} \quad (2.3)$$

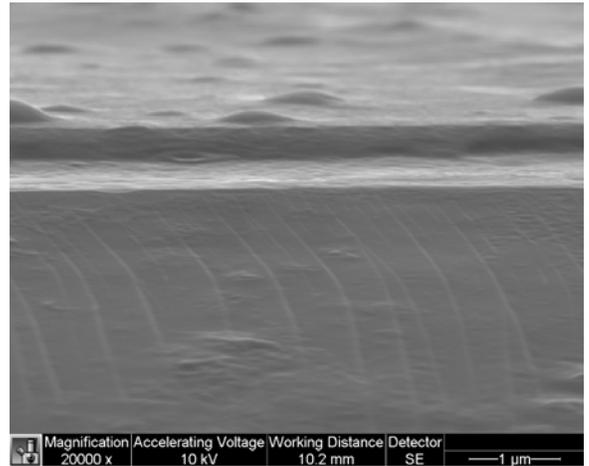
$$= 78.7 \text{ nm}, \quad (2.4)$$

in agreement with previous values reported in other optical absorption spectroscopy measurements [37].

Since the thickness of the carbon coating on this segment is known, calibration of future measurements are done by repeating optical measurements on this sample at positions where thicknesses are known. Such calibration is performed either through dedicated calibration measurement after each sample measurement, or at the same time as sample measurement by placing the calibration crucible next to the crucible to be measured.



(a) SEM image of carbon coating at the beginning of the calibration sample. The thickness averages to be 216 nm.



(b) SEM image of carbon coating at the end of the calibration sample. The measured thickness averages to be 395 nm.

Figure 2.10: SEM images of pyrolytic carbon coating at the beginning and end of calibration sample. The measurement is done on sample PyC003.

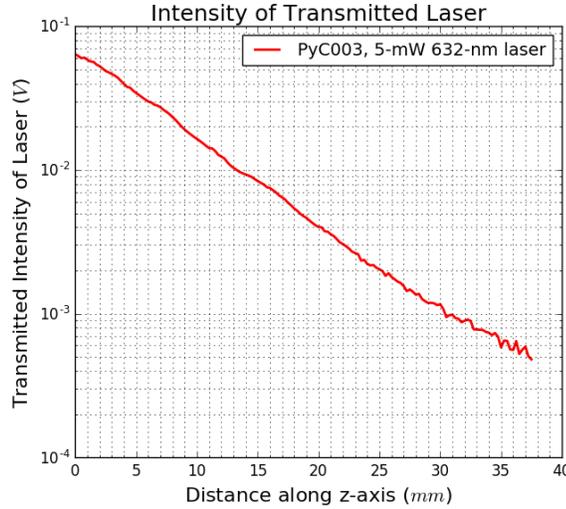


Figure 2.11: Intensity of transmitted laser along the line connecting the two locations on PyC003 where SEM samples were taken. The exponential attenuation indicates a linear increase in the thickness of carbon coating.

2.4.4 Thickness as Function of Position and Coating Duration

To investigate the thickness of carbon coating as functions of position and coating duration, four 45×48 quartz crucibles are coated with varying deposition times. Subsequently the coating thickness is mapped out and averaged azimuthally to get the average thickness of coating versus axial position (Fig. 2.12).

Spatially, the thickness as a function of height is characterized by a region of linear increase followed by a roughly flat region with a valley in the middle. As deposition time gets longer, the valley becomes smaller and the coating tends to be more uniform throughout the entire crucible. This is different from a linearly increasing thickness profile reported by [34].

Temporally, as deposition duration increases, the thickness of coating approaches a terminal value. This is likely due to the limited amount of hexane introduced for each coating: as coating progresses, hexane continues to be consumed and the density of hexane in the vapor decreases, leading to a decrease in the deposition rate.

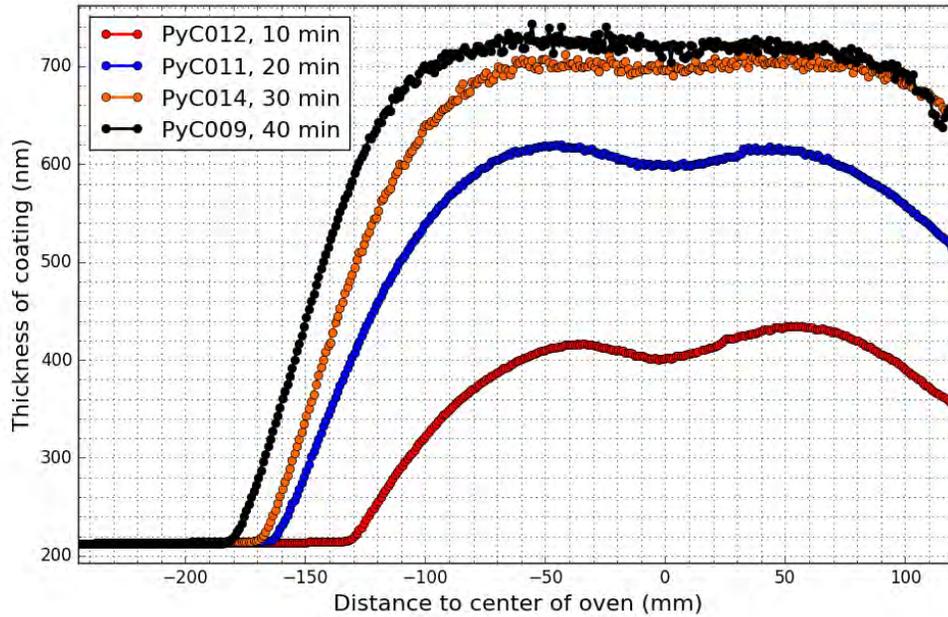


Figure 2.12: Average coating thickness as a function of position for various deposition times. For coating time over 40 minutes, the thickness exceeds the measurement range of PyroMapper and therefore is not included in the study.

To quantitatively analyze this behavior, the average thickness of coating within 5 mm from the center of the furnace is plotted as a function of coating duration. The data are fitted with the following function:

$$y = C_0(1 - e^{-x/\lambda}). \quad (2.5)$$

Eq. 2.5 follows naturally if one makes the following two assumptions:

1. the rate at which coating thickness increases is proportional to the amount of hexane in the vapor, and
2. the amount of hexane remaining in the vapor is a constant (initial amount) minus a term proportional to the thickness of existing coating (hexane consumed in pyrolysis).

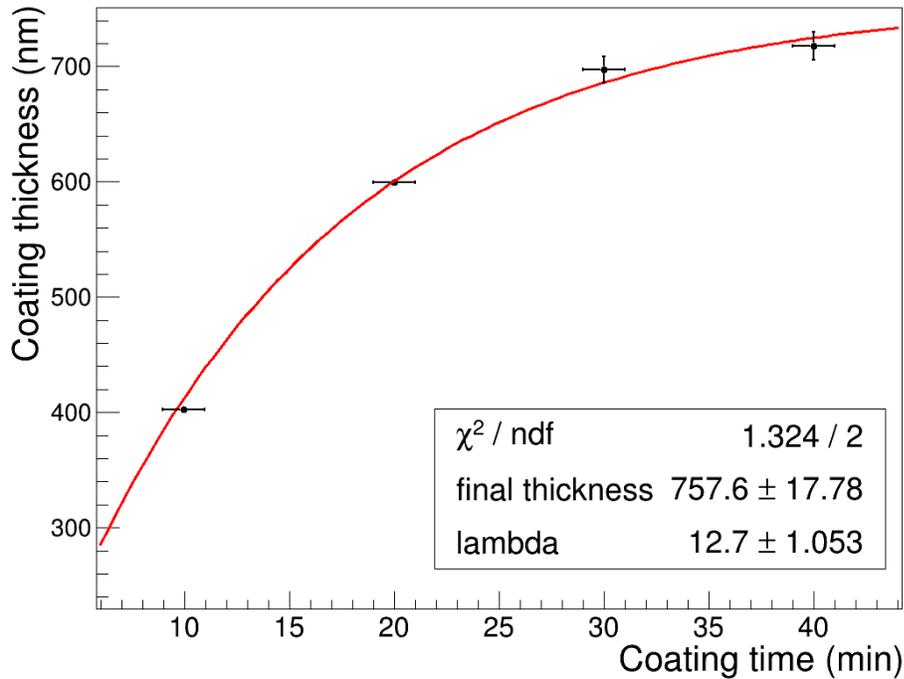


Figure 2.13: Average thickness of carbon coating as a function of deposition time. As deposition time increases, the deposition rate decreases and the thickness of coating approaches a final value around 760 nm.

The result is shown in Fig. 2.13. From the fit, it can be concluded that for 2" crucibles, the thickness of pyrolytic carbon coating approaches a terminal value of 758 ± 18 nm with an exponential time constant of about 13 ± 1 minutes.

Note that if the above assumptions are true, the terminal value will depend on the initial partial pressure of hexane and also the surface-to-volume ratio of the crucible. In the coating procedure outlined in Section 2.4.1, the initial partial pressure of hexane is automatically fixed by its natural vapor pressure, but the surface-to-volume ratio will change if the crucible diameter changes. As a result, the same relations in Fig. 2.12 and Fig. 2.13 are not expected to hold for crucibles of other sizes.

Chapter 3

Precision Drying and Zone

Refining of Astro-grade NaI

Powder

To grow ultra-high purity crystals, the starting raw powder must be of comparable purity to the final crystal. Purification of NaI powder was previously investigated by a collaboration between Sigma-Aldrich and Princeton University, and was found to be best achieved by repeated crystallizations of Na_2CO_3 hydrate, the precursor material used in making NaI. The concentrations of major impurities of interest in the resulting Astro-grade powder is comparable to or even lower than those in the crystals used in the DAMA/LIBRA experiment.

However, the powder obtained from Sigma-Aldrich still contains too much water to be used directly for high-purity crystal growth. Section 3.1 discusses potential problems associated with water in the powder and describes a technique to dry Astro-grade NaI powder in a controlled manner. Section 3.2 describes impurities present in the Astro-grade NaI powder uncovered through drying, and Section 3.3 describes zone refining—a method to further purify NaI powder.

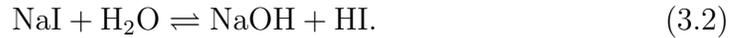
3.1 Precision Drying of NaI Powder

NaI powder is usually manufactured by mixing aqueous solutions of NaCO_3 and HI, and drying the resultant NaI solution. Although commercially available NaI powder is advertised as anhydrous, a trace level of water in the powder can still affect crystal growth and final radiopurity in a number of ways.

1. Water can consume carbon coating at elevated temperatures by



2. Water can react with NaI and produces NaOH [38]:



NaOH is confirmed to etch carbon coating and dissolve fused quartz. Fig. 3.1 shows carbon coating and fused quartz etched by molten NaOH.

3. Tritium in NaOH can become a serious background at low energies via β -decays. Tritium background is discussed in detail in Section 9.2.2.

Furthermore, due to high electronegativity, it is difficult to measure the amount of water or hydroxyl group (OH^-) by ICP-MS. Therefore thorough removal of water in a controlled way is an essential step in successful growth of high radiopurity NaI crystals.

3.1.1 The Apparatus for Precision Drying

To remove residual water remaining in the powder, we developed a technique for precision drying of NaI powder. The procedure works by vacuum-baking of NaI powder at several stages of increasing temperatures, first starting from room temperature.

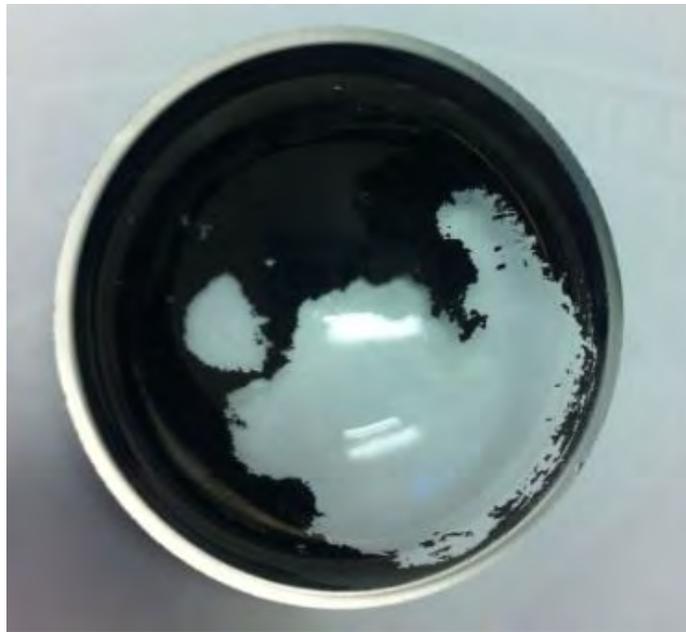


Figure 3.1: Pyrolytic carbon-coated fused quartz crucible etched by molten NaOH.

At each stage, temperature is not raised until the water removal rate—defined as the amount of water removed per unit time—decreases to a terminal value. This guarantees water is removed as efficiently as possible at as low a temperature as possible to avoid production of NaOH by Eq. 3.2.

The piping and instrumentation (P&ID) diagram for the setup is shown in Fig. 3.2a, and a photo of the drying system is shown in Fig. 3.2b. To guarantee high throughput, vacuum lines, valves and connectors used have inner diameters no less than 1". The borosilicate glass cold trap is connected to the rest of the vacuum system by Pyrex-to-Kovar seals and KF25 flanges.

Water removal rate is measured by trapping and collecting water vapor with the borosilicate glass cold trap in liquid nitrogen for certain amount of time (typically 1 hour), since direct measurement is difficult due to low amount of water present after free water is removed. After water has been collected, the cold trap is shut off from the rest of the vacuum system and warmed up to vaporize condensed water. The

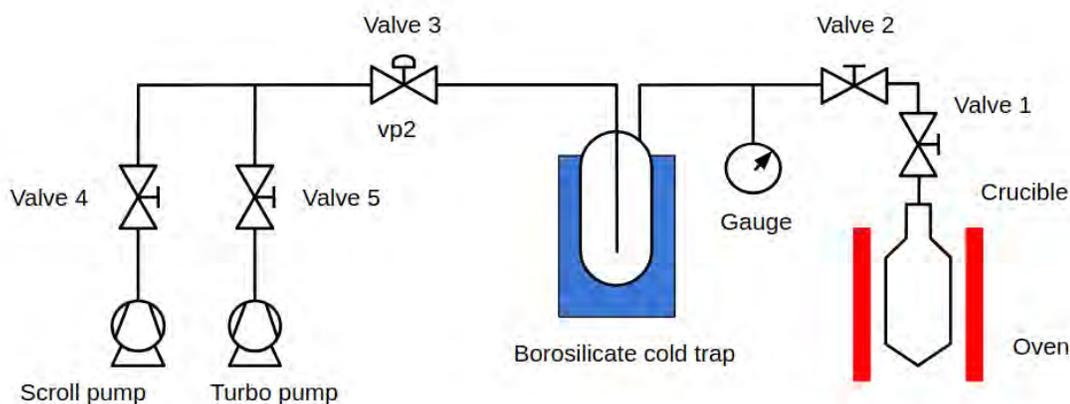
final pressure relative to baseline is used as an indicator of amount of water removed during the interval.

Fig. 3.3 shows data corresponding to a series of water measurements conducted on June 20, 2018. The top plot (red) shows the temperature of the oven; the middle plot shows temperature of the cold trap (blue) and the bottom plot (green) shows the pressure at the cold trap. The oven temperature is raised from room temperature to 120 °C after 1.5 hours since start of DAQ. Each rise in the cold trap temperature and the corresponding peaks in the cold trap pressure indicate a water measurement. Due to limited thermal conductivity under high vacuum, the increase of water removal rate happens after the increase of oven temperature with a delay. Subsequently, the water removal rate decreases until reaching a final constant value. Pressure gauge and sensors are controlled with custom PIC16F1788-based circuit, and data is collected with DAQ software developed using the Polaris framework (see Section 8.2.2).

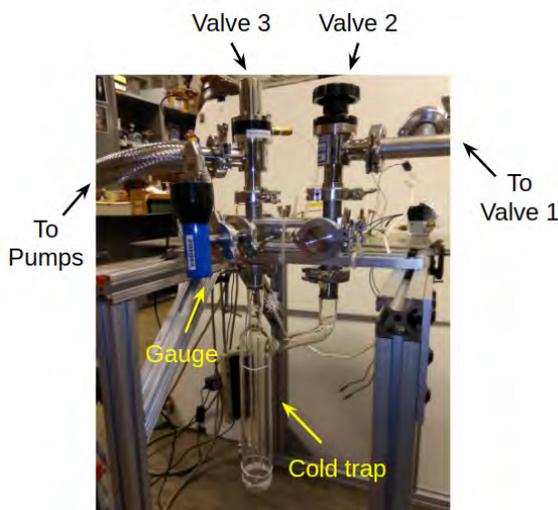
3.1.2 The Drying Behavior of NaI Powder

Water removal rate as functions of time and temperature in Astro-grade NaI powder, shown in Fig. 3.4, is characterized by two major peaks, one at room temperature and one around 120 °C followed by a series of smaller peaks at higher temperatures. The room temperature peak corresponds to free water being removed in the raw powder, which also accounts for the majority of water removed. For a 2-kg batch it takes about one week's time before water removal rate reaches its terminal value.¹ The peak at 120 °C corresponds to decomposition of NaI hydrate. The decomposition temperature of NaI hydrate is around 80 °C [38] but a slightly higher oven temperature is chosen to make sure the temperature throughout the entire crucible is above the decomposition temperature. Above 120 °C, there are no more significant processes that release significant amount of water, and temperature of oven is slowly increased to the final

¹Crystal-grade NaI powder has a similar behavior, but due to larger amount of water present, a significantly longer time is needed.



(a) P&ID diagram of the precision drying system. The pressure gauge is a convection-enhanced Pirani gauge with sensitivity down to 0.01 mbar. The scroll pump (Edwards nXDS6i) has peak pumping speed of 1.7 L/s and can reach an ultimate pressure of about 0.02 mbar. The turbomolecular pump (Pfeiffer HiPace 80) has a pumping speed of 67 L/s and an ultimate vacuum below 10^{-8} mbar.



(b) Photo of the precision drying system. In the vacuum-baking phase, the cold trap is immersed in liquid nitrogen to trap and collect water. In the water measurement phase, Valve 2 and 3 are closed and the cold trap is warmed up and pressure change is recorded.

Figure 3.2: P&ID diagram and photo of the precision drying system. NaI powder is loaded in a quartz drying crucible inside a glovebox. During transportation, Valve 1 is closed to keep away ambient moisture and particulates. Valve 4 and 5 are used to select between a scroll pump and a turbomolecular pump, depending on the level of moisture and type of gas being pumped. In the earlier stages where water might condense in the pump, the scroll pump is used with gas ballast open. As the powder gets drier, the turbomolecular pump is used instead.

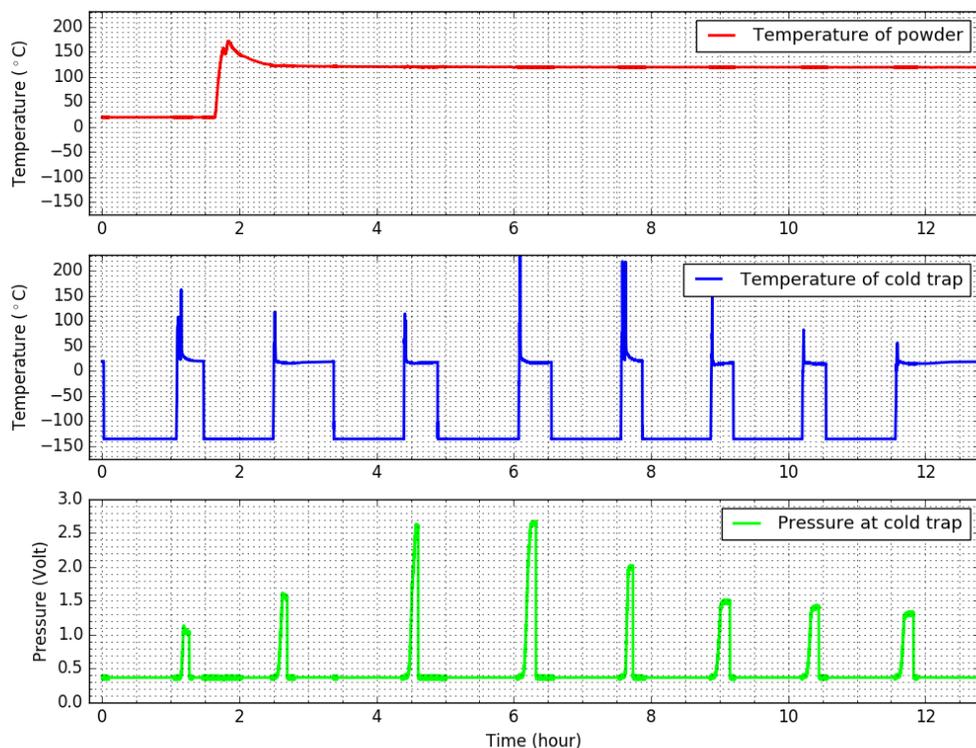


Figure 3.3: Pressure and temperature of cold trap and temperature of NaI powder over a number of water measurements. As indicated by the top plot, data is taken during and after temperature of the oven is raised to 120 °C from room temperature. In the middle plot, each well in the cold trap temperature corresponds to water collection with liquid nitrogen and the subsequent spike in temperature rise is due to temporary overheating of the trap with heat gun. Water collection time for each measurement is 1 hour. The bottom plot shows pressure of the cold trap as a function of time. After the temperature of the oven is raised, the pressure due to trapped water increases, indicating a higher water removal rate. As the remaining amount of water decreases, the amount of collected water also decreases until final equilibrium is achieved.

temperature of 380 °C. Each steps of precision drying and the characteristic behavior of water are summarized in Table 3.1.

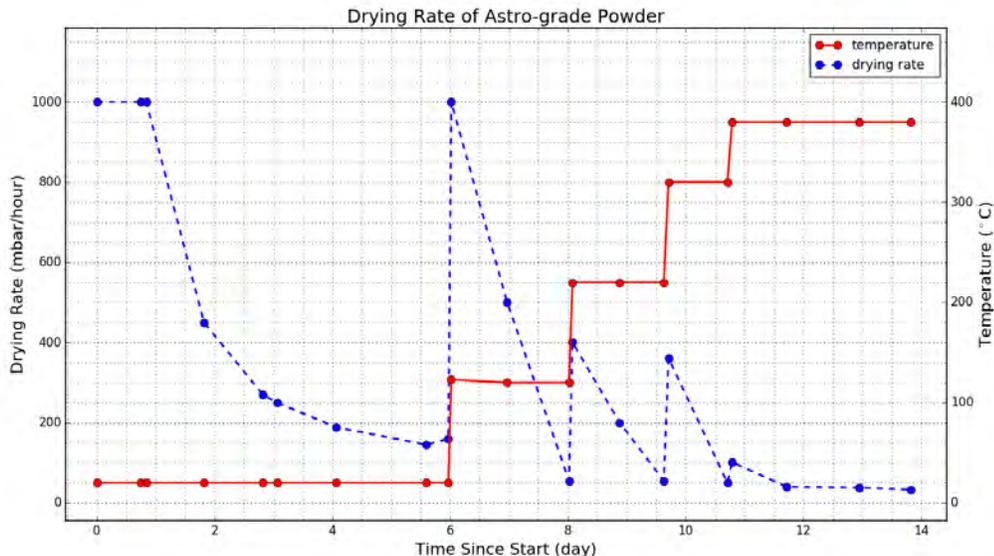


Figure 3.4: Water removal rate as functions of time for Astro-grade NaI powder. In the first few measurements, liquid water is present in the trap after warming up, and the maximum value of the vacuum gauge is taken instead.

3.2 Impurities in the Astro-grade NaI Powder

It was observed that for the first few water measurements above 300 °C, a different gas species is released from the powder. The presence of this gas component is indicated by a prompt response of the vacuum gauge as the cold trap is warmed up (Fig. 3.5).

This can be explained by the gas having a lower melting point than water: upon heating the outer surface of the cold trap, water vapor will condense around the cold inner column and it takes some time before convection can warm up the inner column, causing a delay in the response of vacuum gauge. However, if the inner column is not cold enough to condense the unknown gas, it will vaporize immediately, causing a prompt response in the pressure gauge.

The amount of immediate gauge response is plotted as a function of time and temperature in Fig. 3.6. The temperature threshold and finite amount suggest that this unknown gas is associated with thermal reactions of impurities in the NaI powder.

Table 3.1: Procedure for precision drying of NaI powder

Temperature	Duration	Characteristics
20 °C	1 week	Free water is pumped out. Cold trap or gas ballast is required to prevent water from condensing in the pump. A large amount of condensed water is observable in the cold trap after the cold trap is warmed up.
120 °C	1-3 days	NaI hydrate begins to decompose, contributing to a high water removal rate. As the amount of hydrate decreases, water removal rate also goes down quickly.
240 °C	1-2 days	Water removal rate slightly increases at the beginning, but quickly goes down.
320 °C	1-2 days	Drying rate increases again at the beginning. Decomposition of other impurities (e.g. NaHCO ₃ and NaOH hydrate) contributes to an increase in the water removal rate.
380 °C	1-2 days	No significant change in the water removal rate is observed.

One hypothesis is thermal decomposition of NaHCO₃. An alternative possibility is water released from NaOH hydrate [39] further reacting with NaI and producing HI by Eq. 3.2.

In addition, occasionally Astro-grade NaI powder develops black specks during vacuum-baking at above 200 °C. X-ray photoelectron spectroscopy (XPS) measurement determined that the contaminant is carbon. Fig. 3.7 compares aluminum K α XPS spectrum of white NaI powder to that of powder with black specks. It can be clearly seen that carbon 1s valence electron peak at 285 eV and carbon KVV auger electron peak at 1223 eV are present in the spectrum of powder with black specks but are absent from the spectrum of white powder.

Judging from the fact that 1) it only appears at temperatures above 200 °C and 2) it is not removed by prior thorough pumping, the carbon is likely from the decomposition of low-vapor pressure hydrocarbons present in the raw powder. The

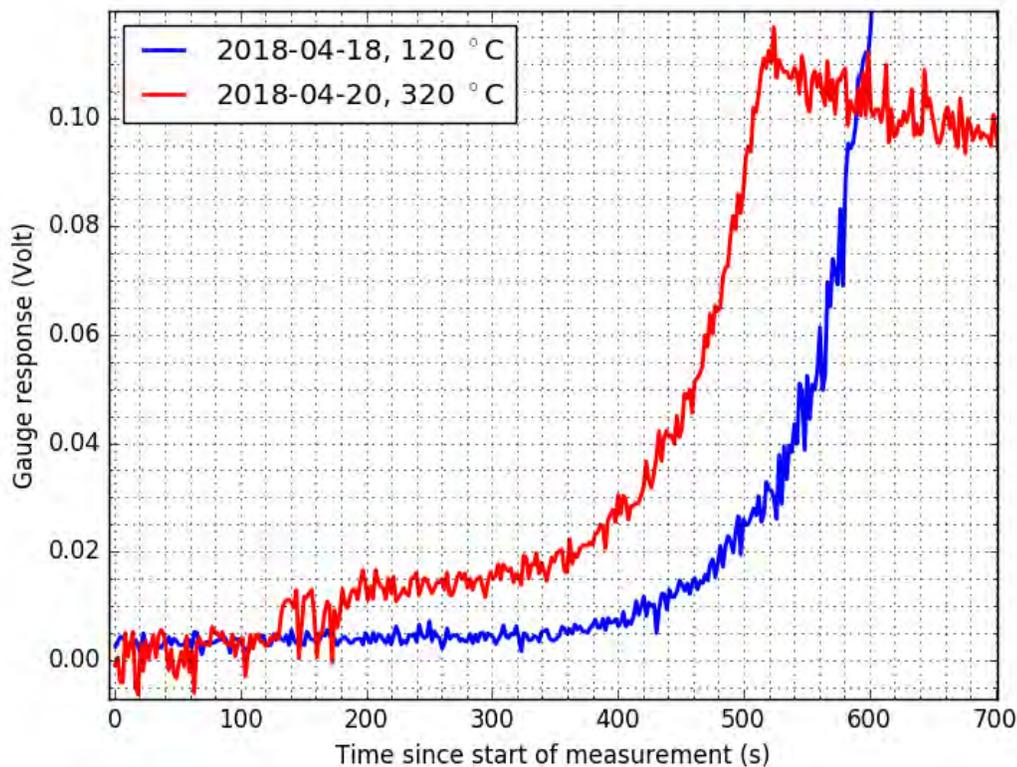


Figure 3.5: Response of vacuum gauge within the first few minutes as cold trap is heated. Red curve shows the response for data taken at 320 °C and blue curve shows data collected at 120 °C for comparison. Between 100 to 200 seconds, the red curve has a change in gauge voltage absent in the blue curve, indicating the presence of vapor other than water.

possibility of back-streaming of pump oil can be ruled out since both pumps used in the precision drying are oil-free dry pumps. However the Astro-grade NaI powder may have been vacuum-dried with oil-based pumps at the factory.

The presence of these impurities along with the benefits from improved powder purity motivates further purification of Astro-grade NaI powder prior to crystal growth. One such method—zone refining—is discussed in the next section.

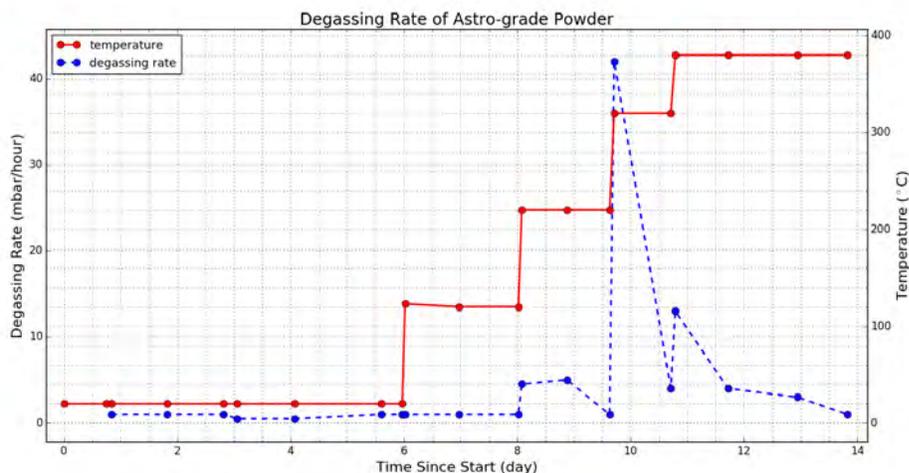


Figure 3.6: Oven temperature (red) and immediate cold trap gauge reading (blue) as a functions of time. The immediate cold trap gauge reading indicates the amount of gas species with melting point lower than water. Below 300 °C this is not significant, but once temperature exceeds 300 °C, the amount of released gas immediately rises.

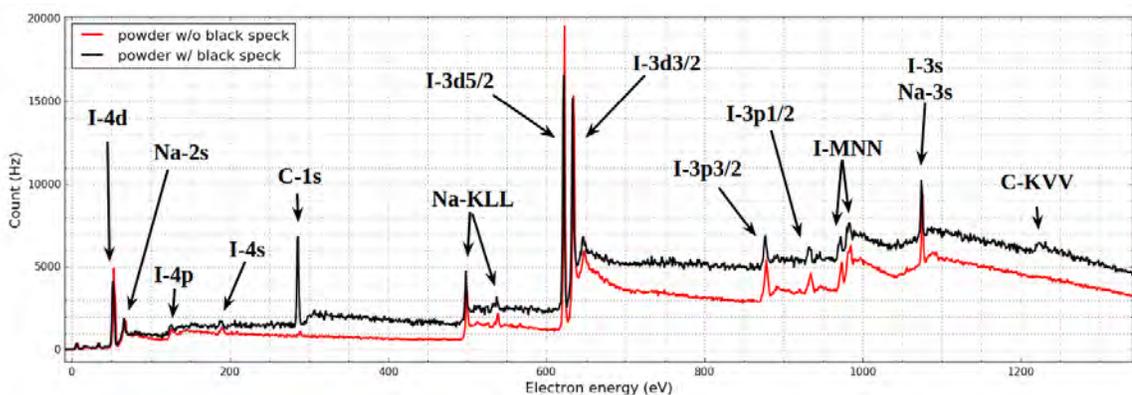


Figure 3.7: XPS spectrum of Astro-grade NaI powder powder with and without black specks. It can be clearly seen that 1s valence electron peak and KVV auger electron peak is present in the spectrum of powder with black specks but are absent from that of white powder.

3.3 Zone Refining of NaI Powder

Zone refining is a crystal purification process in which impurities are moved to the ends of the ingot by molten zones. The principle and process of zone refining is explained in detail in Section 3.3.1, and Section 3.3.2 describes the impurity distribution after crystal growth/zone refining. Section 3.3.3 discusses different ways to calculate the

distribution coefficient—a key parameter describing the effectiveness of zone refining. Optimization of zone refining in the context of purifying NaI powder is described in Section 3.3.4. Finally, Section 3.3.5 presents the design and construction of a new zone refining system.

3.3.1 Principle

When a crystal grows from melt or solution, the crystal lattice will purify itself by attracting certain types of impurities and repelling other types. This phenomena is best explained from thermodynamics. Consider the phase diagram of a binary mixture shown in Fig. 3.8. Vertical axis is the temperature of the mixture and the horizontal axis the concentration of the impurity. The initial state when everything is molten is represented by a point above the red liquidus line. As the temperature drops, the system stays as a homogeneous liquid mixture until past the liquidus line where homogeneous liquid mixture is energetically unfavorable and the system spontaneously separates into a liquid phase with concentration C_{L1} and a solid phase with concentration C_{S1} . According to this phase diagram, the impurity concentration in the solid phase is lower than that in the liquid phase, thereby the crystal is purified. The equilibrium distribution coefficient k_{eq} is defined as the ratio of impurity concentration in the solid phase to that in the liquid phase:²

$$k_{\text{eq}} = C_S/C_L. \quad (3.3)$$

Pictorially k_{eq} is also the ratio of the slope of the liquidus line to that of the solidus line. k_{eq} is a measure of the strength of phase separation: the lower k_{eq} is, the stronger impurities are repelled from the crystal matrix.

²Such definition assumes the liquidus and solidus curves to be linear, but such assumption is always true near 0, or when impurity concentration is small.

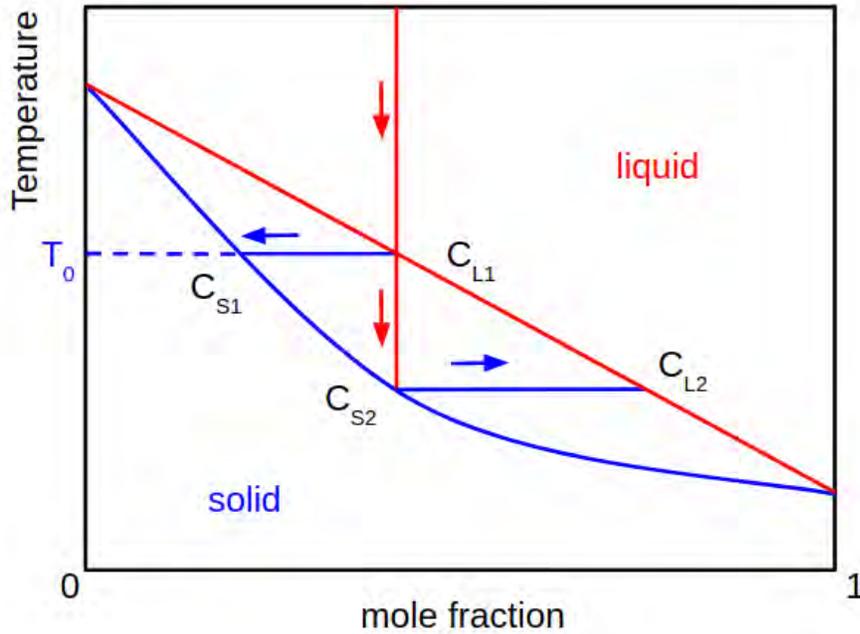


Figure 3.8: An example phase diagram of a binary mixture. Above the liquidus curve (red), everything melts as a homogeneous liquid. Below the solidus curve (blue), everything freezes to a homogeneous solid.

In the phase diagram shown in Fig. 3.8 the impurity has effectively lowered the melting point. If the slope of the liquidus and solidus lines are positive instead, the impurity has effectively raised the melting point. Such case can be analyzed in a similar manner and the distribution coefficient takes a value greater than 1, indicating that impurities tend to stay in the crystal lattice. Due to the similarity, following discussions will mainly focus on the case of $k_{\text{eq}} < 1$.

It must be pointed out that k_{eq} applies only to systems at or near equilibrium. If the process happens too fast for the different phases to reach equilibrium (e.g. rapid crystal growth), k_{eq} must be replaced with effective distribution coefficient k_{eff} . For simplicity of discussion, the notation k_{eq} will be used throughout the rest of the thesis, but k_{eq} is understood to be k_{eff} for non-equilibrium systems.

Normal Gradient Freeze

When a single crystal is grown from melt, the raw powder is first melted entirely, and the melt is allowed to crystallize slowly from one end. When the first bit of crystal forms, the impurities in the crystal is suppressed by a factor of k_{eq} compared to the remaining molten liquid, and the extra impurities are pushed into the liquid phase, slightly raising the average impurity concentration in the liquid. As crystal grows, the impurity concentration in the liquid keeps rising due to additional impurities being pushed into the melt from the solid crystal and newly formed crystal has a higher impurity concentration (Fig. 3.9). Therefore the impurity concentration in the final crystal will increase from the tip which crystallizes first to the tail which crystallizes last. This process is called normal gradient freeze.

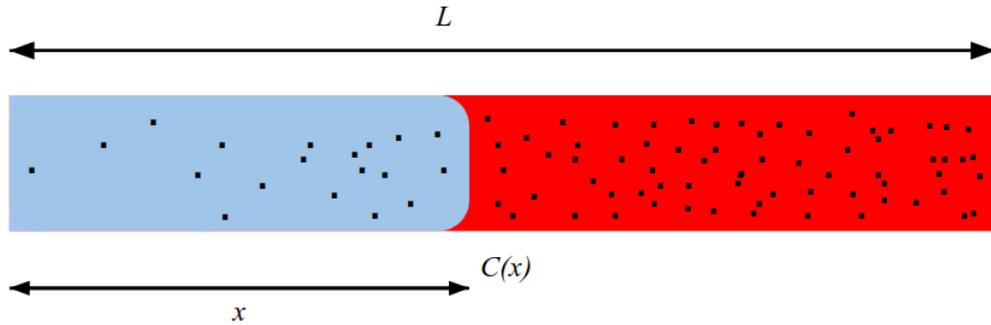


Figure 3.9: Illustration of normal gradient freeze. As the solid-liquid boundary advances, impurity concentration in the solid is k_{eq} times that of the remaining liquid. Due to additional impurities from the solid being pushed into the liquid phase, as the boundary advances the impurity concentration in the solid phase also increases.

It is very easy to find the distribution of impurities in this scenario: assume an ingot of length L . The impurity concentration at distance x from tip is given by

$$C(x) = k_{\text{eq}} \int_x^L C(y) dy / (L - x). \quad (3.4)$$

It states that the impurity concentration for crystal newly formed at position x is k_{eq} times the impurity concentration of the liquid, which is expressed as an average of

the impurity concentration over the remainder of the ingot. The differential equation obtained by taking derivative with respect to x on both sides of Eq. 3.4 has a simple solution:

$$C(x) = k_{\text{eq}}C_0(1 - x/L)^{k_{\text{eq}}}, \quad (3.5)$$

where the boundary condition of $C(0) = k_{\text{eq}}C_0$ is imposed.

Zone Refining

In zone refining, instead of melting everything, part of the ingot, also referred to as a zone, is melted. This molten zone is moved across the entire ingot repeatedly in the same direction. As the zone is moved, newly formed crystal has a smaller impurity concentration compared to the molten liquid. Thus impurities are pushed towards the ends along with the molten zone. This concept is shown in Fig. 3.10. Unlike normal freezing in which everything melts and homogenizes at the beginning, zone refining occurs locally: as the zone makes the next pass, the molten liquid has less impurities due to purification by previous pass, thus the solid gets further purified. Such process is repeated multiple times before the dirty end of the ingot is cut away to make a crystal with higher purity.

Since its invention by Pfann [40], zone refining has remained a key technique in producing high-purity silicon and germanium for semiconductor applications.

3.3.2 Impurity Distribution

Unlike normal directional freeze (Section 3.3.1), the impurity distribution after arbitrary number of zone refining passes is complicated, and no analytical solution has been found. However this problem can be solved numerically with ease [41]. An

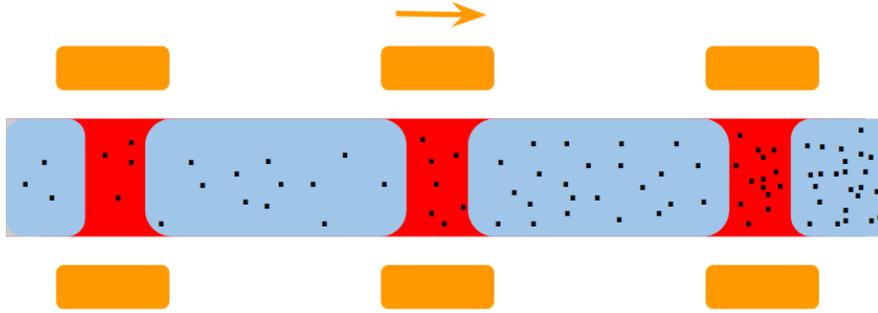


Figure 3.10: Illustration of zone refining with 3 zones. After the first zone (rightmost) pushes impurities towards the right end by same principle as normal crystal growth, the second zone sees less average impurity concentration. Therefore as the second zone passes, further purification is achieved. Such zone passes can be repeated until equilibrium distribution of impurities is reached.

improved algorithm based on [41] that takes into account crucible shape is described in Appendix A.

Impurity distribution after 1, 2, 5, 10 and 25 zone passes are shown in Fig. 3.11. From this plot, it is not difficult to see two limitations of zone refining:

1. similar to any other purification techniques, higher purity comes at the cost of more sacrifice of raw material, and
2. there is an ultimate equilibrium distribution, achieved when each part of the ingot has the right amount of impurity such that ratio of impurity concentration in solid to that in liquid is k_{eq} throughout the entire ingot.

3.3.3 Measuring Distribution Coefficient

In the purification of NaI powder for ultra-high purity crystal growth, ^{36}Cl , ^{40}K , ^{87}Rb and ^{210}Pb are of concern because they are major sources of background (Section 9.2.2). To quantify how well zone refining removes these impurities, k_{eq} for each impurity must be known. The coefficient between alkali halides has been previously measured by [42], but the accuracy is not known and a new estimate is necessary.

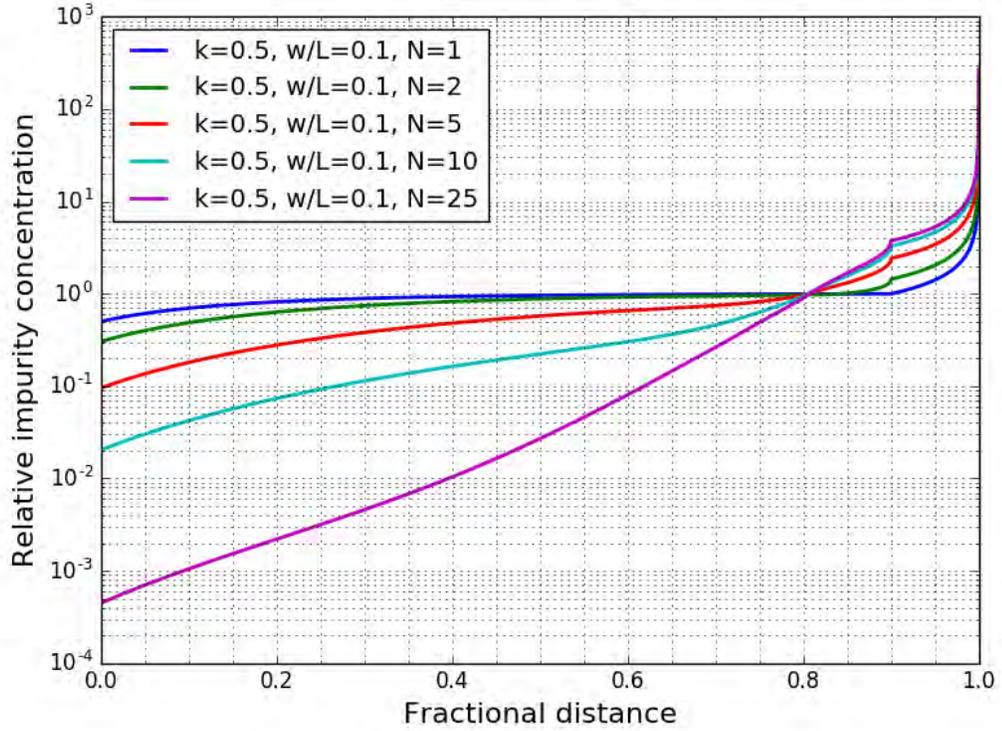


Figure 3.11: Impurity distribution after various numbers of zone passes for distribution coefficient of 0.5 and zone width of 1/10 the ingot length. The distance is measured from the beginning as fraction of total ingot length. The relative impurity concentration is normalized to an initial uniform distribution of 1. The singularity around 0.9 is due to molten zone reaching the far end and impurity having no where else to go.

Based on previous discussions, distribution coefficient can be estimated in 3 ways:

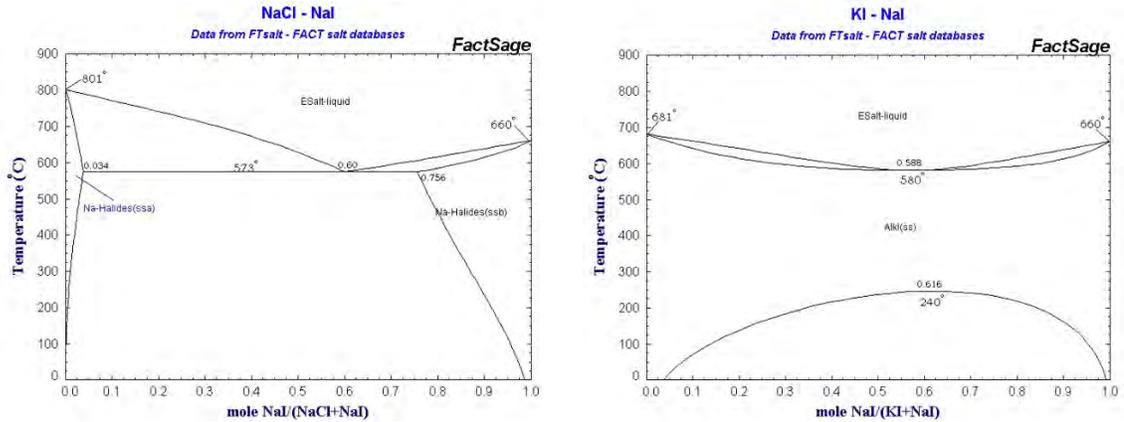
1. ratio of impurities in the solid phase to that in the liquid phase in a phase diagram of binary mixture,
2. measuring impurity distribution after crystal growth, and
3. measuring impurity distribution after dedicated zone refining tests.

Calculated k_{eq} are summarized in Table 3.3 at the end.

Calculating k_{eq} from Phase Diagram

k_{eq} is calculated from phase diagrams of binary system by dividing the equilibrium impurity molar concentration on solidus curve by that of liquidus curve, both evaluated at a temperature close to the melting point of NaI. This process is repeated at different temperatures to estimate errors.

Phase diagrams of NaCl-NaI, KI-NaI, RbI-NaI and LiI-NaI are taken from [43]. Phase diagrams of NaCl-NaI and KI-NaI are reproduced in Fig. 3.12 as illustration. The phase diagram of PbI_2 -NaI system can be found at [44], but the data is not conclusive enough to allow calculation of k_{eq} .



(a) Phase diagram of NaCl-NaI system.

(b) Phase diagram of KI-NaI system.

Figure 3.12: Phase diagram of NaCl-NaI and KI-NaI binary system [43]. The equilibrium coefficient is calculated by taking the ratio of molar fraction in solid phase over that in liquid phase at a temperature near 661 °C.

Calculation of k_{eq} using phase diagrams is quick and cost-effective, but it is subject to the availability of phase diagram of particular binary system of interest. The molar concentrations displayed in the phase diagrams are usually on the order of 0.1%, and the accuracy of phase diagrams when extrapolated to the sub-ppm region is not known.

Measuring k_{eq} by Crystal Growth

As crystal grows from melt by gradually freezing from one end (gradient freeze), the impurity will be distributed according to Eq. 3.5.³ Since impurity concentrations at different locations of the final crystal are determined by only two parameters—amplitude (or original impurity concentration in the powder C_0) and k_{eq} , impurity measurements at two different positions are needed.

To estimate k_{eq} , ICP-MS measurements of crystal impurities in samples from NaI-014, NaI-031 and NaI-032 are used (Section 4.3). Two samples from each crystal, one near the tip of the crystal and the other near the tail, are analyzed using ICP-MS.⁴ Measured impurity concentrations as functions of position is fitted with simulated impurity distribution, keeping overall impurity concentration and k_{eq} as free parameters. The simulation algorithm is described in Appendix A, with zone width set as the length of ingot.

k_{eq} of KI in NaI is obtained from NaI-014 (Section 4.3.1) since it provides the cleanest data without contamination from the fused quartz crucible. LiI is obtained from NaI-031 (Section 4.3.8) because LiI in NaI-014 and NaI-032 are below detection limit due to the ultra-high purity of synthetic fused silica.⁵

Since crystal growth is carried out very slowly, the solid-liquid mixture is more likely to be in equilibrium. k_{eq} calculated from crystal growth might be unrealistic for zone refining which is usually done more than 10 times faster than crystal growth.

Measuring k_{eq} by Zone Refining

To measure k_{eq} in the most realistic setting, dedicated zone refining tests were conducted at Radiation Monitoring Devices, Inc. (RMD). NaI powder was loaded in a

³The same distribution can also be calculated by numerical simulation using the algorithm outlined in Appendix A.

⁴Tip is the end that crystallizes first during crystal growth.

⁵ Li diffusing out from crucible is not modeled in the simulation and there could be corresponding error in the estimate.

2-cm-diameter, 30-cm-long crucible. Powder was doped with KI, RbI and PbI, respectively, in 3 different runs to enhance signal-to-background ratio. The crucible was zone-refined for 25 passes using a ring oven with ID slightly larger than the OD of the crucible at a rate similar to that used by [42]. Samples from 5 different locations were analyzed using glow discharge mass spectroscopy (GDMS). Sample locations and measurement results are summarized in Table 3.2.

Table 3.2: Impurities at different locations from a NaI ingot zone-refined by 25 passes. It can be seen that K and Rb are mainly pushed in the direction of motion of oven while Pb goes in the other direction. Data provided by RMD.

Element	Concentration (ppm)					k
	5 mm	80 mm	150 mm	225 mm	295 mm	
K	1.5	11.1	122	224	5459	0.57
Rb	20	0.3	0.21	4.8	36000	0.37
Pb	0.14	0.083	0.072	0.042	0.029	1.11 ± 0.01

Summary of Measurement

In general there is good agreement between k_{eq} obtained by different means. k_{eq} obtained from zone refining are expected to be higher than that obtained from crystal growth whereas values calculated from phase diagrams are consistent with those obtained from crystal growth.

Although k_{eq} obtained from zone refining tests are reliable only at ppm level due to doping, data from crystal growth indicates that phase separation of K still occurs at the ppb-level. Therefore zone refining of NaI powder is also expected to be effective at reducing K impurities at ppb level.

Note that for PbI_2 , k_{eq} is greater than 1, indicating that PbI_2 will move in the opposite direction during zone refining. This is true only for PbI_2 -NaI mixture. Unlike highly reactive alkali metals such as Li and K, Pb might not react with HI when

Table 3.3: Calculated and measured distribution coefficients of LiI-NaI, NaCl-NaI, KI-NaI and RbI-NaI binary system. Data from [42] is included for comparison. Data not available is marked with n\a. The value for PbI₂ is larger than 1, indicating that PbI₂ will concentrate in the solid phase.

Impurity	Distribution coefficient			
	Phase diagram	Crystal growth	Zone refining	Literature ([42])
LiI	0.38 ± 0.06	0.42 ± 0.07	n\a	0.45
NaCl	0.52 ± 0.07	n\a	n\a	n\a
KI	0.45 ± 0.10	0.35	0.57	0.39
RbI	0.18 ± 0.02	n\a	0.37	0.16
PbI ₂	n\a	1.2 ± 0.3	1.11 ± 0.01	n\a

Na₂CO₃ is mixed with HI to make NaI. If Pb exists in NaI powder in other forms, phase separation will occur differently.

3.3.4 Efficiency of Zone Refining

Impurity distribution after zone refining depends only on two parameters: k_{eq} and ratio of zone width to ingot length w/L . Ideally the effect of w/L should be considered as well, however in practice, it is mostly determined by the zone refining oven (width and range of motion of each ring oven). Typical values range from 0.05-0.2. For simplicity of analysis, w/L is taken as 0.1.

To study the effectiveness of purification of NaI by zone refining, the projected final purity achieved is plotted as a function of the number of zone refining passes (N) and the fraction of the ingot kept (f). Two contour plots are shown in Fig. 3.13 for $k = 0.3$ and $k = 0.5$, inferred from Table 3.3.

The plot indicates that with 20 zone passes, more than 70% of the raw material can be purified by a factor of 100. However, subsequently the contour line becomes horizontal, indicating that additional zone refining will not significantly improve the final purity.

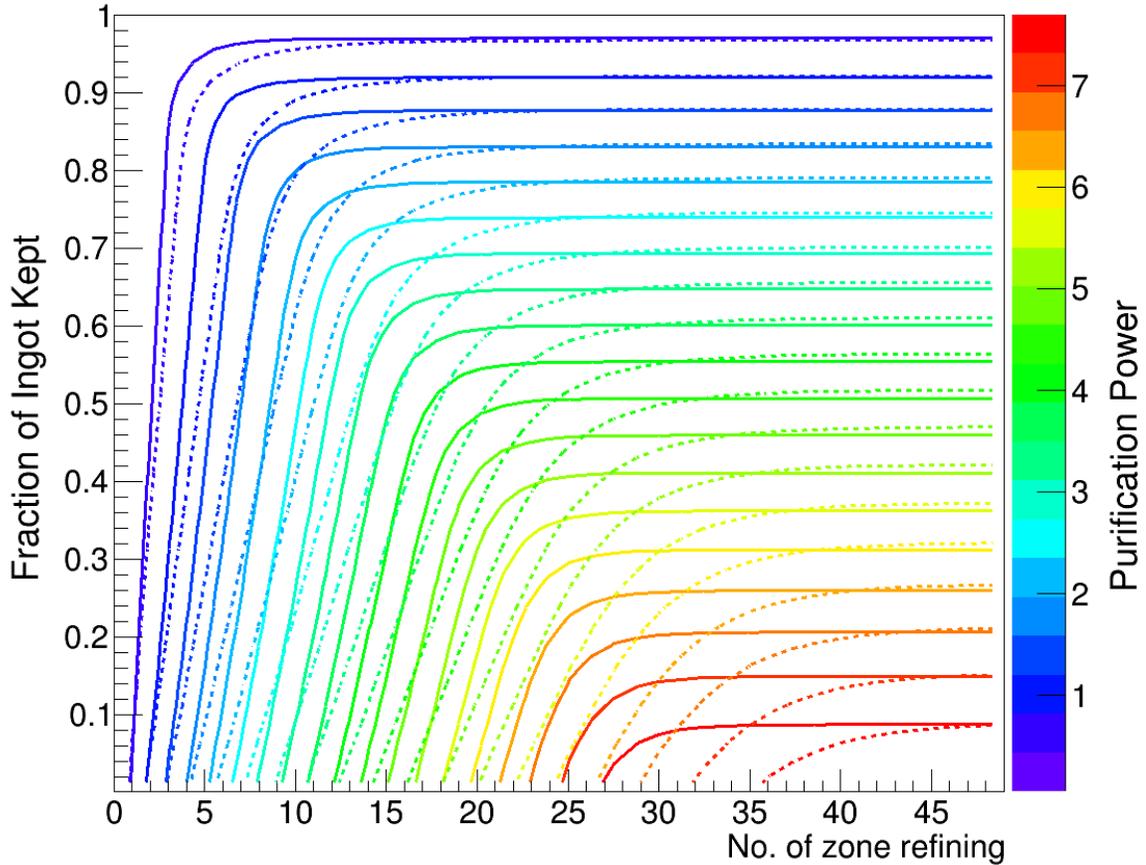
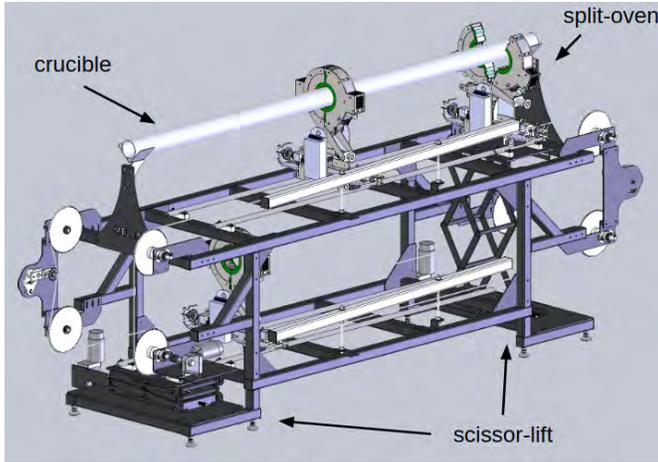


Figure 3.13: Contour plot of purification power as a function of number of zone passes and fraction of ingot kept. Purification power is defined as negative log of ratio of impurities after and before zone refining. The solid line corresponds to $k = 0.3$, and the dashed line corresponds to $k = 0.5$.

3.3.5 A New Apparatus for Zone Refining

In zone refining, multiple ring ovens are often moved across the ingot simultaneously (Fig. 3.10) such that in the time one oven makes one pass, the ingot has been zone-refined multiple times. However since ring ovens all move along the same line, after the first ring oven finishes the travel, it needs to wait for all other heaters to finish before they can move back to the starting location for next run. To overcome this limitation, together with Mellen Company, Inc. (Mellen), we co-invented a novel zone refining system that features continuous operation of the ring ovens.



(a) CAD design of new zone refining system.



(b) New zone refiner during test.

Figure 3.14: CAD design and photo of new zone refining apparatus.

Continuous operation is realized by two-level design and a split oven; zone refining takes place at the upper level. Once an oven finishes a zone refining pass, it is split open and brought down to the lower level by a scissor-lift and moved back to below the starting position. The heater is then brought up again by another scissor lift and begins the next zone refining pass.

Currently, the apparatus is in the final tuning phase, after which 3 kg of Astro-grade NaI powder will be zone-refined in a synthetic fused silica crucible as a final test.

Chapter 4

Growth of Ultra-high Purity

NaI (Tl) Crystal

To grow large single crystals, raw material is first melted in a crucible and slowly cooled down in a controlled manner such that the entire crystal consists of a single crystal lattice. Section 4.1 gives an overview of three major growth methods used industrially to produce large crystals, and compares them in the context of ultra-high purity crystal growth. Steps and special considerations in growing NaI (Tl) crystal are described in Section 4.2. Section 4.3 summarizes major crystal growth activities that happened over the course of my Ph.D., and this chapter concludes with a brief discussion of using SiCl_4 to improve crystal growth (Section 4.4).

4.1 Single Crystal Growth from Melt

Growth of large single crystals from melt is often achieved by either the Czochralski process, Kyropoulos process, or Bridgman-Stockbarger process. The difference between these processes lies in how the melt is cooled. In the Czochralski process, the raw material for the crystal is melted in a crucible and a rod with a seed crystal attached at the end is inserted into the melt and slowly pulled out. As the rod is

being pulled out, the end of the rod cools down and the melt crystallizes around the rod. In the Bridgman-Stockbarger process, a crucible filled with raw material is melted and slowly cooled by passing the entire melt through a thermal gradient. The Kyropoulos process is similar to the Czochralski process, except the crystal is slowly rotated and pulled out at a much slower rate such that the entire melt can crystallize into a single crystal.



(a) The Bridgman-Stockbarger method. (b) The Czochralski/Kyropoulos method.

Figure 4.1: Illustration of frequently-used methods for single crystal growth. Blue color indicates temperatures below the melting point of the raw material and red temperatures above the melting point. In the Bridgman-Stockbarger method, the part that first enters the cold zone crystallizes and provides nucleation sites for the molten materials in the hot zone as the crucible advances into the cold zone. The crucible is usually designed with a tip in the bottom to select a single crystal lattice out of randomly-formed poly-crystals. In the Czochralski/Kyropoulos method, a rotating seed crystal is used instead as nucleation sites for the molten materials.

Typically the Kyropoulos method has the highest yield, and the final single crystal does not contact the crucible wall, leading to less stringent requirements on the crucible material. Commercially, the Kyropoulos method is frequently used to produce large alkali halide crystals, including NaI (Tl). However in the Kyropoulos method, it is difficult to completely seal the melt from outside, and the melt can potentially become contaminated by particulates in the ambient air and furnace refractory materials. The process also involves many parameters and the growth conditions must be constantly monitored and controlled. On the other hand, in the Bridgman-

Stockbarger method, the growth process is less complex and can be automated. Since the crucible containing the melt can be completely sealed to prevent additional contaminants from being introduced during crystal growth, the purity of the crystal, which depends only on the quality and purity of the crucible and the initial raw material, can be strictly controlled. This feature makes it very suitable for growing ultra-high purity crystals. In this thesis I will be investigating growth of ultra-high purity NaI (Tl) crystals using the Bridgman-Stockbarger method.

4.2 Ultra-high Purity NaI (Tl) Crystal Growth

The procedure for growth of ultra-high purity NaI (Tl) crystal is summarized below followed by a few comments.

1. The crucible is precision-cleaned prior to coating to make sure the surface is free of contaminants. Cleaning procedure is described in detail in Section 2.3.
2. The crucible is coated with a layer of pyrolytic carbon with thickness ranging from a few hundred nm to 1 μm . Coating is either done at Sandfire or at Princeton using the method described in Section 2.4.1.
3. The crucible is precision-cleaned again to remove potential contaminants introduced during and after coating. Then the crucible is vacuum-baked to remove water adsorbed to the surface during cleaning.
4. The raw NaI powder is precision-dried to remove water. Precision drying is described in detail in Chapter 3.1.
5. TlI powder is mixed with the NaI powder at 0.2% mol/mol concentration and loaded into the crucible inside a nitrogen-purged glovebox. Subsequently the crucible is pumped down with a turbomolecular pump and sealed with a quartz plug (see Fig. 2.1).

6. The crucible is heated up in the hot zone of a vertical Bridgman furnace to around 750 °C for a few days to melt and mix the NaI powder and the TII powder.¹ The crucible is then slowly lowered into the cold zone with temperature around 500 °C at a rate of 7-10 mm/day.
7. After growth is finished, the crucible is cooled down to room temperature over one week and cut open. Subsequently the crystal is cut to its final size. Both the crucible and the crystal are cut using a diamond wire saw.

According to the literature, best performance is achieved when the concentration of TII is about 0.1% mol/mol. However, similar to other impurities in NaI, added TII will separate in the course of crystal growth. The distribution coefficient of TII in NaI was investigated by [45] and the authors found a value near 0.35, which is similar to KI during crystal growth. For crystal sizes of concern, K concentration in the final crystal is roughly half of initial concentration, justifying the empirical practice of adding twice as much TII as desired in the final crystal to compensate for separation during growth.

Due to the hygroscopic nature of NaI powder and crystals, all powder handling is carried out in nitrogen-purged glovebox inside a cleanroom, except the crucible and the crystal cutting, which takes place in a dry room at RMD. In the future, this step will be carried out in a softwall dry room inside a radon-free cleanroom in Jadwin Hall, Princeton University.

4.3 Major Crystal Growth Activities

In the past 3 years, a number of crystals have been grown using different grades of NaI powder, different crucible materials and under different growth conditions. Many

¹TII can be thoroughly mixed for small crystals, but the uniformity of TI in larger-size crystals obtained in this way needs further quantitative studies.

conclusions in earlier chapters are based on crystal growth activities outlined in this section. These conclusions are summarized below and a list of major crystal growth activities is summarized in Table 4.1.

1. As mentioned in Section 2.1, impurities diffusing out of fused quartz will contaminate the crystal. Therefore synthetic fused silica crucibles should be used for the growth of ultra-high purity crystals. This statement is based on a comparison of NaI-014 (Section 4.3.1) and NaI-032 (Section 4.3.9) with NaI-031 (Section 4.3.8).
2. Water must be thoroughly removed at as low a temperature as possible, and the powder should be free of NaOH. In addition, Crystal-grade NaI powder from Sigma-Aldrich contains too much NaOH and/or other impurities to be used with fused quartz or synthetic fused silica crucibles, even with carbon coatings. This conclusion is based on NaI-014 (Section 4.3.1), NaI-025 (Section 4.3.3) and Artoria (Section 4.3.4).
3. Pedestal support of crucible during growth would alter the temperature distribution inside the molten NaI, leading to defects in the final crystal.² This conclusion is based on NaI-026 (Section 4.3.5) and NaI-031 (Section 4.3.8).
4. Suspension of the crucible from the top using quartz rod or wire during crystal growth runs the risk of mechanical failure, especially for crucibles with large mass (Section 4.3.1 and 4.3.6). An ideal alternative material is Inconel [46] (Section 4.3.8).

²This can be potentially solved by careful profiling and tuning of temperature distribution prior to growth.

Table 4.1: A list of major crystal growth activities. AG and CG in the table refer to Astro-grade and Crystal-grade NaI powder, different grades of NaI powder sold by Sigma-Aldrich. Typical impurities in Astro-grade powder are at the ppb-level while those for Crystal-grade powder are at the ppm-level.

Name	Date	Powder	Crucible Type	Diameter	Coating
NaI-014	2015-10	2-kg AG	synthetic silica	3.5 inch	Sandfire
NaI-015	2015-12	2-kg AG	fused quartz	3.5 inch	Sandfire
NaI-025	2017-03	6-kg CG	fused quartz	4 inch	Sandfire
Artoria	2017-08	1-kg CG&AG	fused quartz	2 inch	Princeton
NaI-026	2017-08	6-kg AG	fused quartz	4 inch	Sandfire
NaI-027	2017-10	8-kg AG	synthetic silica	4 inch	Sandfire
Gemini	2018-04	2 × 200-g AG	fused quartz	1.5 inch	Princeton
NaI-031	2018-06	6-kg AG	fused quartz	4 inch	Sandfire
NaI-032	2018-08	6-kg AG	synthetic silica	4 inch	Princeton
NaI-033	2018-11	6-kg AG	synthetic silica	4 inch	Sandfire

4.3.1 NaI-014

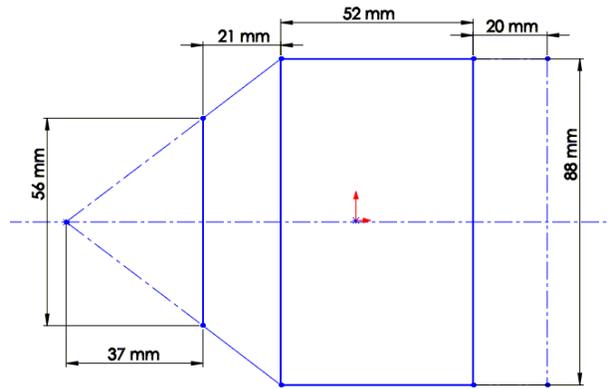
NaI-014 was grown with 2 kg of Astro-grade NaI powder and a 4"-diameter synthetic silica crucible. The pyrolytic carbon coating was done at Sandfire, and the crucible was precision-cleaned after the carbon coating.³ Raw powder was precision-dried for the first time using an earlier approach based on pump&purge over 4 days at Princeton.⁴ The crystal was grown using a vertical Bridgman oven at RMD. However, during cooling a quartz rod used for suspension broke. Held by only one thermocouple wire, the crucible lost balance and hit the oven wall, causing a crack in the final crystal. The approximate final dimensions are shown in Fig. 4.2.

³Regular cleaning using Detergent 8, deionized water and isopropyl alcohol was performed prior to coating.

⁴In this method, instead of continuous vacuum baking, inert gas (helium or nitrogen) is periodically introduced to facilitate removal of water. Preliminary tests indicate that pump&purge is about 25% more effective at removing water, but requires a series of filters and a getter to purify the gas.



(a) Photo of NaI-014. The crack is not visible from this angle due to total internal refraction of NaI crystal.



(b) Sketch of measured dimensions of NaI-014. Dashed parts at the tip and tail are cut away during crystal fabrication to improve crystal purity.

Figure 4.2: Photo and approximate dimensions of NaI-014. The crack is not visible due to internal refraction.

Fig. 4.3 shows the ^{137}Cs spectrum measured by NaI-014. Due to the crack, the crystal performed poorly as a scintillator. To take the data, γ -rays from ^{137}Cs were collimated to minimize the effect of the crack.

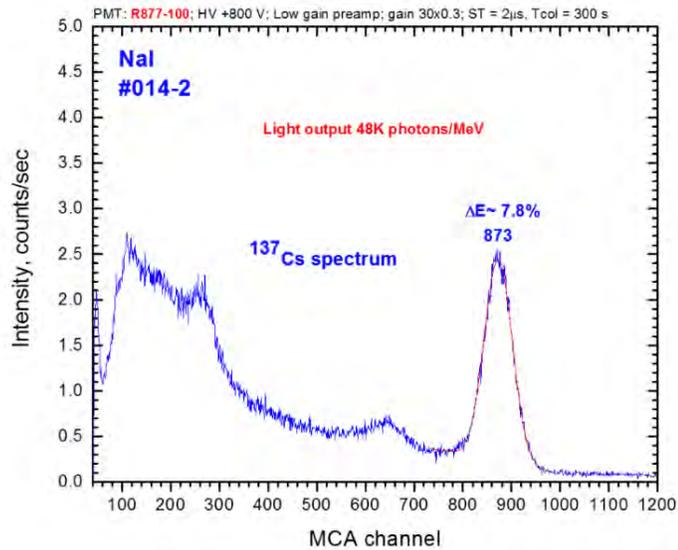


Figure 4.3: ^{137}Cs spectrum of NaI-014. The crystal was cut and polished into a 88-mm-diameter 73-mm-long cylinder. The resolution at 661.7 keV is roughly 7.8%. Data provided by RMD.

To measure the crystal’s purity two samples approximately 37 mm and 110 mm away from the tip were sent to Seastar for ICP-MS measurement. The initial measurement showed a high level of K, but the presence of yttrium indicated the possibility of cross-contamination when the sample was cut with a diamond-wire saw that has been used to cut other crystals. To get bulk impurities, the surface of the crystal sample was incrementally dissolved with ethanol 5 times in synthetic quartz glassware before measurements. New result summarized in Table 4.2 indicated a lower level of K contamination.

Table 4.2: ICP-MS measurements by Seastar on samples from NaI-014 before and after surface cleaning. [^{89}Y] is back to normal after cleaning, and [^{39}K] dropped by 2.

Element	Powder	Tip (27 mm, in ppb)		Tail (110 mm, in ppb)	
		pre-cleaning	post-cleaning	pre-cleaning	post-cleaning
^7Li	5	4	< 0.5	7	< 0.5
^{39}K	10	11	5	26	12
^{85}Rb	22	24	< 0.5	22	< 0.5
^{89}Y	0.2	260	0.8	81	0.2
^{208}Pb	0.2	1.2	0.7	3.5	0.4
^{232}Th	< 0.2	< 0.2	< 0.2	< 0.2	< 0.2
^{238}U	< 0.03	< 0.03	< 0.03	< 0.03	< 0.03

To obtain more accurate result, the 52-mm-long cylindrical body was sliced into 5 equally-spaced 5-mm-thick disks and four of these disks were carefully cleaned with high-purity ethanol and sent to Seastar and PNNL for independent impurity measurements by ICP-MS. The results are summarized in Table 4.3.

To estimate the average concentration of K, the position-dependent ^{39}K concentration is fitted with simulation (Section 3.3.2) and distribution coefficient k is extracted. Fig. 4.4 shows the combined measurements and best fit. The distribution coefficient k is estimated to be 0.35. The average K concentration in the entire ingot is estimated to be 17.6 ± 1.2 ppb. Assuming the fitted average concentration and distribution co-

efficient, K concentration in the cylindrical part of the final crystal is 10.3 ± 0.7 ppb, and 9.5 ± 0.7 ppb if the half-cone in Fig. 4.2b is included. The slight increase in the overall average impurity compared to the powder could be due to lack of precision cleaning.

Table 4.3: ICP-MS measurements of impurity concentration in ppb at different locations in NaI-014. Distance is measured from the tip of the crystal. Initial measurements at 27 mm and 110 mm are also listed for comparison. Since measurements were done in two batches, there could be systematic error between the two measurements.

Isotope	Concentration (ppb)					
	27 mm	58 mm	68 mm	88 mm	98 mm	110 mm
^7Li	< 0.5			< 9		< 0.5
$^{39}\text{K}(\text{Seastar})$	5	9	7	13	9	12
$^{39}\text{K}(\text{PNNL})$	n\a	10.0	9.1	9.7	9.8	n\a
^{85}Rb	< 0.5			< 0.4		< 0.5
^{89}Y	0.8			< 0.8		< 0.2
^{208}Pb	0.7	1.3	1.4	1.5	1.1	< 0.4

These results indicated that NaI-014 had the highest purity ever achieved for NaI crystals, and proved the effectiveness of the crucible preparation, drying processes and growth method.

4.3.2 NaI-015

To test the effectiveness of precision-drying, RMD dried 2 kg of Astro-grade powder using a vacuum-baking approach with temperature profile similar to that employed for NaI-014. The crystal was grown in a fused quartz crucible at RMD. Although no sticking or damage to crystal was observed, the surface was not as smooth as that of NaI-014. The crystal was fabricated into a $3'' \times 3''$ cylinder, and exhibited a good resolution. Due to the limitation of the DAQ system at RMD, the light yield is not known.

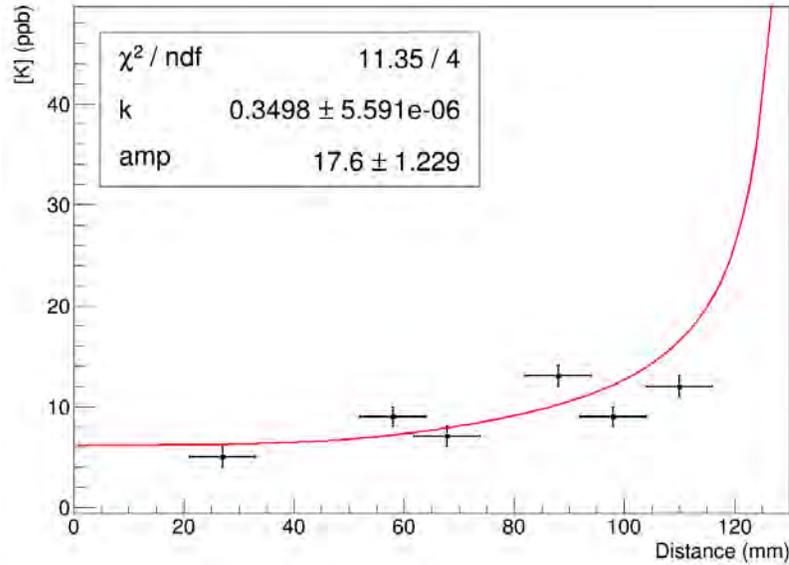
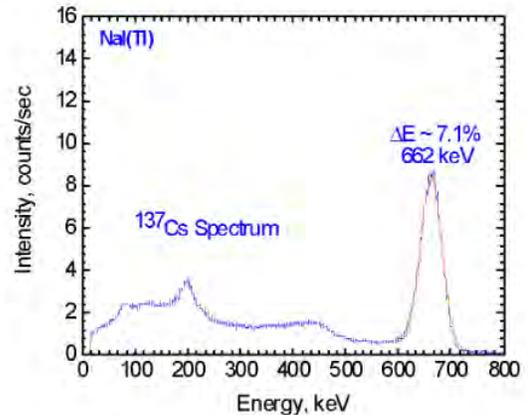


Figure 4.4: Measured and fitted distribution of ^{39}K in NaI-014. The distribution coefficient is estimated to be 0.35 and average ^{39}K in the entire ingot 17.6 ± 1.2 ppb. After cutting away the tip and tail samples, average ^{39}K concentration in the crystal is 9.5 ± 0.7 ppb. If the conical section is entirely cut away, the average ^{39}K concentration is 10.3 ± 0.7 ppb.



(a) Photo of NaI-015.



(b) NaI-015 ^{137}Cs spectrum. The crystal was fabricated into a $3'' \times 3''$ cylinder. The resolution at 661.7 keV is roughly 7.1%.

Figure 4.5: Photo of NaI-015 and ^{137}Cs spectrum recorded with NaI-015. To measure the spectrum the crystal is fabricated into a $3'' \times 3''$ standard size. Photo and data provided by RMD.

ICP-MS measurements on samples from the tip and tail of the ingot indicated a ^{39}K level of 21 ppb and 75 ppb respectively. Such a high level of contamination is due

to the use of fused quartz instead of synthetic silica and possibly lack of precision-cleaning.

4.3.3 NaI-025

To scale up the crystal mass and dimension without risking more expensive Astro-grade NaI powder, in February 2017 an effort was made to test crystal growth with 6 kg of precision-dried Crystal-grade NaI powder. A 4"-diameter crucible made out of regular quartz was used since contamination was not a concern in this test. However during crystal growth, the carbon coating peeled off and the crucible was found to be cracked after cooling (Fig. 4.6).



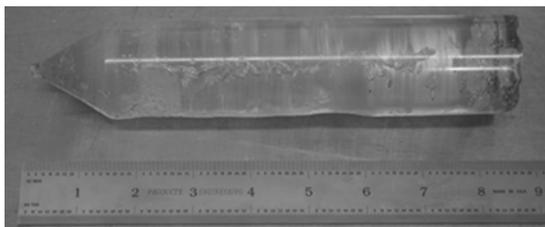
Figure 4.6: NaI-025 failed due to carbon coating peeling off and crucible cracking during cooling. Pieces of pyrolytic carbon is embedded in the crystal. The coating on the top part of the crucible (the right end in picture) is completely consumed due to reaction with oxygen in the air at the elevated temperature after the crucible cracked open. Photo courtesy of RMD.

The failure of NaI-025 lead to development of the new pyrolytic carbon coating technique described in Chapter 2 and investigations into cause of failure in Section 4.3.4.

4.3.4 Artoria

Artoria was grown with 1 kg of Astro-grade NaI powder in a 2" fused quartz crucible to test the pyrolytic carbon coating procedure developed at Princeton University.

Instead of physically moving the crystal, the crystal is placed at the natural gradient created by a regular two-zone clam-shell oven and the oven temperature is slowly and uniformly lowered. The crystal came out without sticking to the crucible, and the carbon coating was intact. However, the crystal had a chain of bubble-like cavities. The cause is most likely a nonlinear temperature gradient inside the oven.⁵ A photo of the crystal is shown in Fig. 4.7a.



(a) Photo of Artoria. Grown with 1-kg precision dried Astro-grade powder and quartz crucible coated at Princeton. Photo courtesy of Hushu and Wuhana.



(b) Crystal grown with 1-kg precision-dried Crystal-grade powder as a control group.

Figure 4.7: Photo of Artoria and control group using Crystal-grade powder. Everything else is exactly the same in these two crystal growths.

An identical crystal growth using precision-dried regular Crystal-grade NaI powder from Sigma-Aldrich resulted in the carbon coating peeling off during growth and the crucible cracking during cooling, leading to the conclusion that impurities present were the main cause of failure with regular Crystal-grade powder. Based on small-scale tests and other literatures ([42], [47], [39]), the impurity of consequence is likely NaOH introduced during improper drying at Sigma-Aldrich.

4.3.5 NaI-026

After Artoria (Section 4.3.4) pointed out the impurities in the powder as a major cause of failure in crystal growth using Crystal-grade powder, all crystal growth were carried

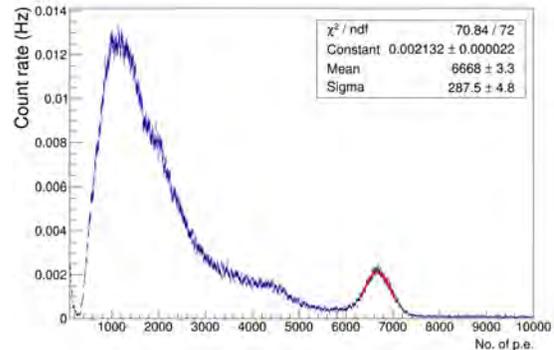
⁵Moving at a constant rate through a stationary linear temperature gradient is equivalent to uniformly decreasing the oven temperature without any motion only if the gradient is linear.

out using Astro-grade NaI powder. NaI-026 was grown with 6 kg of Astro-grade NaI powder in a precision-cleaned, fused quartz crucible coated by Sandfire. Instead of suspension from top, the crystal was supported from bottom with a quartz pedestal. Although the resultant crystal did not stick to the crucible, the crystal had many small bubble-like cavities in the bulk (Fig. 4.8a). A later crystal growth (Section 4.3.8) indicated that these bubbles were likely due to presence of the bottom quartz pedestal altering the temperature distribution of the melt.

The conical section of the tip was free of bubbles and it was fabricated into a $1.5'' \times 1.5'' \times 1''$ cube. The light yield of the cube was measured with Hamamatsu R11065 PMT at Princeton and exhibited a relatively high light yield of 10.1 p.e./keV. Assuming typical quantum efficiency of 25% for R11065 PMTs, the crystal has a light yield of about 40 p.e./keV, consistent with the high light yield reported in the literature [48]. The cube was subsequently used to study the effect of different reflector materials (see Section 5.1). The bulk of the crystal could not be used as scintillator due to poor light propagation caused by the bubbles.



(a) Photo of NaI-026 cross section. NaI-026 had many small bubble-like cavities in the bulk of the crystal. Photo courtesy of RMD.



(b) ^{137}Cs spectrum measured with a $1.5'' \times 1.5'' \times 1''$ cubical sample near the tip of NaI-026. Light yield at 661.7 keV is measured to be around 10.1 p.e./keV.

Figure 4.8: Photo of NaI-026 and ^{137}Cs spectrum measured with a $1.5'' \times 1.5'' \times 1''$ cubical sample near the tip of NaI-026.

Samples were taken for impurity assay, and ICP-MS results from Seastar indicated ^{39}K concentrations of 14 ppb and 21 ppb in the tip and tail regions, respectively. The increase in ^{39}K is consistent with other crystals grown using quartz crucibles.

4.3.6 NaI-027

After realizing the possibility of the pedestal causing the bubbles in NaI-026, another effort was made in October 2017 to grow an 8 kg crystal using top suspension and synthetic fused silica. For this growth, the powder was precision-dried at Princeton and coating was done at Sandfire. During the crystal growth, RMD used folded multi-stranded Kanthal wire to support the crucible, but the Kanthal wire lost mechanical strength rapidly as the temperature rose, and eventually broke after 4 days of crystal growth. The crucible fell onto the floor and the tip of the crucible was broken.



(a) Kanthal wire and indented quartz flange for supporting crucible.



(b) After 4 days of growth, the Kanthal wire broke right at the quartz flange, causing the crucible to fall onto the floor.

Figure 4.9: Multi-stranded Kanthal wire was used in the growth of NaI-027. However after 4 days of growth, the wire broke and crystal growth failed. Photo courtesy of RMD.

The failure of NaI-027 led to the use of Inconel 625 [46], a nickel-based alloy with good corrosion-resistance and mechanical properties at high temperature. Table 4.4 compares Inconel 625 with other high-temperature alloys. Nichrothal 80 [49] and

310 stainless steel—a major high-temperature stainless steel in the austenitic steel family—are also listed to give a baseline.⁶ It can be seen that Inconel has an ultimate tensile strength almost one order of magnitude greater than kanthal [50], and twice higher than 310 stainless steel [51]. Based on this finding, crystal growths after NaI-027 were all carried out by suspension from top using Inconel 625-based structures.

Table 4.4: Ultimate tensile strength of some high-temperature alloy.

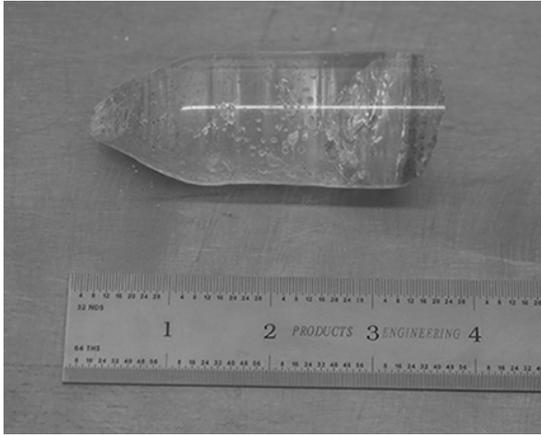
Material	Ultimate tensile strength (MPa)		Strength remaining
	20 °C	900 °C	
Kanthal A-1	680-760	34	5%
Nichrothal 80	660-700	100	14%
Inconel 625	827-965	223	23-27%
310 stainless steel	600-710	130	18-22%

4.3.7 Gemini

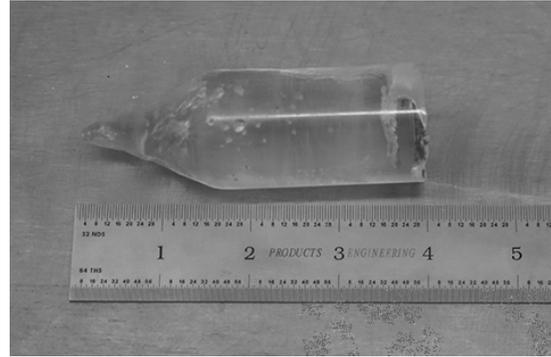
Gemini- α and β are two small 100-g crystals grown with 1.5" fused quartz crucibles. The purpose was to test whether a new batch of Astro-grade NaI powder with lower concentration of K but higher concentration of water was suitable for crystal growth. The crucibles were coated at Princeton University, and the crystals were grown at a rate of 7.5 mm/day by passing them through temperature gradients created with two different two-zone clam-shell ovens. Both crystals came out without sticking to the crucibles. Unlike Artoria (section 4.3.4), Gemini β did not have any defects in the bulk. Gemini α had similar “bubbles” to Artoria, which is possibly due to nonuniform temperature distribution caused by partially damaged thermal insulation in the oven wall.

⁶Nichrothal 80 is a nickle-chromium alloy made by Kanthal[®] with same chemical composition as nichrome 80/20 alloy.

However, it must be pointed out that the presence of NaOH can still be a problem for larger crystals—a smaller surface-to-volume ratio leads to more NaOH per unit surface area when the same NaI powder is used.



(a) Gemini- α .



(b) Gemini- β .

Figure 4.10: Gemini- α and β were grown in two different clamb-shell ovens of the same model. The defects in α are suspected to be due to damaged thermal insulation on the oven used. Photo courtesy of Hushu and Wuhan.

4.3.8 NaI-031

NaI-031 was grown with 6-kg Astro-grade NaI powder in a precision-cleaned carbon-coated quartz crucible. The powder was dried at RMD using a similar procedure to the precision drying described in Section 3.1. The crucible was suspended from the top using Inconel [46] structures instead of a pedestal in contrast to NaI-026. The crystal did not stick to the crucible and didn't have any defects inside.

Samples from tip and tail were sent to Seastar for ethanol cleaning and ICP-MS measurement. The results are summarized in Table 4.5. By fitting the measured distribution with model, the average concentration of ^{39}K throughout the entire ingot is estimated to be 26 ppb and that of ^7Li to be 91 ppb. Comparing to NaI-014 for which synthetic fused silica was used (Table 4.2), the concentrations of ^{39}K and ^7Li have increased by ~ 20 ppb and ~ 90 ppb, respectively. This is also consistent with

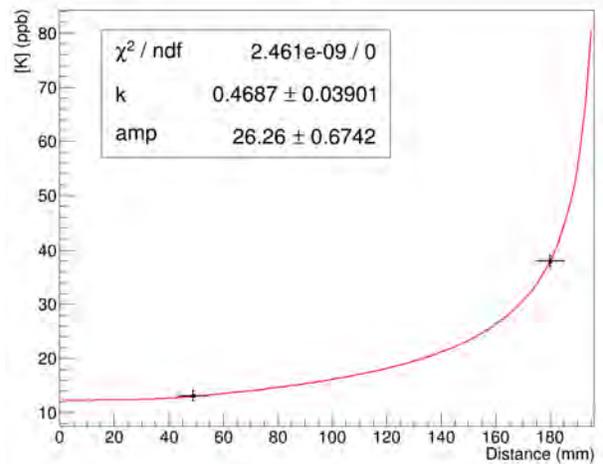
the back-of-envelope calculations of alkali metal diffusion from quartz crucible (Section 2.1), especially the big increase of ${}^7\text{Li}$, which is well-consistent with its greater mobility.

Table 4.5: ICP-MS measurements of samples from NaI-031. Tip and tail samples were taken approximately 49 mm and 180 mm away from tip end respectively.

Isotope	Powder	Concentration (ppb)	
		Tip Sample (49 mm)	Tail Sample (180 mm)
${}^7\text{Li}$	< 10	40	130
${}^{39}\text{K}$	8	13	38
${}^{85}\text{Rb}$	< 0.2		< 0.2
${}^{89}\text{Y}$	1.5	1	< 0.5
${}^{208}\text{Pb}$	2.1	3.0	2.4
${}^{232}\text{Th}$	< 0.08		< 0.08
${}^{238}\text{U}$	< 0.02		< 0.02



(a) Photo of NaI-031. Photo courtesy of RMD.



(b) ${}^{39}\text{K}$ distribution in NaI-031 as function of position. The average concentration in the entire ingot is about 26 ppb.

Figure 4.11: Photo and fitted ${}^{39}\text{K}$ distribution of NaI-031 as function of position.

After cutting off the tip and tail, the remaining 3.6 kg crystal, 98 mm in diameter and 131 mm in length, was coupled with two PMTs and packaged into a crystal

detector module (see Section 5.4). Based on the fit to impurity distribution in tip and tail, the average concentration of ^{39}K in the final crystal is estimated to be 16.5 ± 1.0 ppb.

4.3.9 NaI-032

NaI-032 was grown with 6 kg of Astro-grade NaI powder in a precision-cleaned synthetic silica crucible coated at Princeton. The powder was precision-dried at Princeton. However the crystal growth failed due to the carbon coating peeling off and the crystal sticking to and eventually breaking the crucible (Fig. 4.12a).

The cause of failure is still unknown. Coating procedure has been previously tested with Artoria (Section 4.3.4), and Gemini (Section 4.3.7). Powder from the same batch as NaI-032 was also tested with Gemini. However this does not rule out the possibility of NaOH in the powder being the cause.

Samples from tip and tail were analyzed with ICP-MS and the concentrations of Li and K are summarized in Table 4.6. ^{85}Rb , ^{238}U , and ^{232}Th were all below detection limit. K concentration in the tip sample was very close to the detection limit of ICP-MS (about 3 ppb). By fitting with the distribution model, the K concentration in the entire ingot is estimated to be 8.9 ± 0.1 ppb (Fig. 4.12b). Comparing to the concentration in the powder (8.7 ppb), it can be seen that little to no contamination was introduced in the crystal. The crystal section between the two sample locations had an average K concentration of 4.7 ± 1.4 ppb. If it could be used as a scintillator, it would have an unprecedented purity in terms of K.

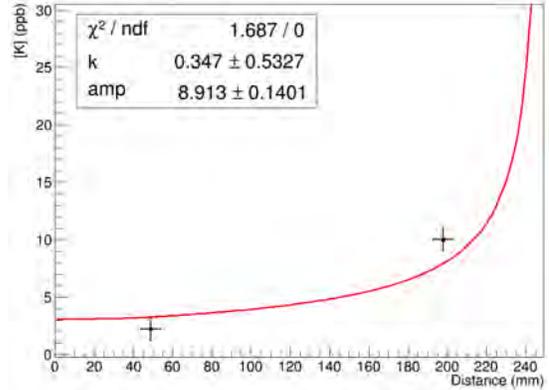
Although the crystal growth was not successful, purity analysis not only proved ion diffusion from the crucible being a problem in regular fused quartz, but also showed that the procedures and protocols adopted (powder handling, drying, crucible cleaning and coating) are compatible with ultra-high purity requirements.

Table 4.6: ICP-MS measurements on samples from NaI-032. Tip and tail samples were taken approximately 49 mm and 198 mm away from the tip end respectively.

Isotope	Powder	Concentration (ppb)	
		Tip Sample (49 mm)	Tail Sample (198 mm)
^7Li	< 10	< 10	< 10
^{39}K	8.7	2.2	10



(a) Photo of NaI-032. The carbon coating detached from the crucible and the crystal stuck firmly to the crucible, causing the crucible to eventually break. The yellow color on the right end is due to TII being purified out as crystal grows.



(b) ^{39}K distribution in NaI-032 as function of position. The average concentration in the entire crystal is about 8.9 ± 0.1 ppb, consistent with the K concentration in the raw powder prior to growth.

Figure 4.12: Photo and fitted ^{39}K distribution of NaI-032 as function of position. Photo courtesy of RMD.

4.3.10 NaI-033

Comparing with NaI-031, the failure of NaI-032 can be attributed to different coating procedures between Sandfire and Princeton, different drying procedures between RMD and Princeton, and different amounts of NaOH in the powder used. To eliminate the former two as potential causes and to get a high-purity crystal for underground counting, a new crystal growth (NaI-033) was started in September 14, 2018 with identical conditions to NaI-031, except synthetic fused silica is used instead of regular fused quartz. The crystal growth is expected to finish on October 31, 2018.

4.4 Using Silicon Tetrachloride to Prevent Crystal Sticking

NaOH is known to adversely affect growth of alkali halide single crystals by causing sticking to quartz crucible. A previous test (Section 3.1) indicates that NaOH can permeate through the carbon coating and etch the quartz underneath the coating, causing the coating peeling off and the crystal sticking to the crucible. Precision drying of NaI powder is effective in removing water and preventing further formation of NaOH, but cannot remove NaOH that has already formed during production in the factory. To fundamentally solve the problem of sticking and improve yield of crystal growth, NaOH must be completely removed prior to crystal growth.

Based on past literatures, hydrogen halide (HX) and silicon tetrachloride (SiCl_4) are two agents that can effectively remove NaOH. Powder treated with HX/ SiCl_4 is reported to not stick to even uncoated fused quartz [42][47]. Hydrogen halide works by neutralizing NaOH:



Presence of excess acid (HX) will push the equilibrium to the right and favors the formation of water and halide instead of NaOH. If HI is used as the hydrogen halide, the product is NaI and no foreign substance is introduced. However this method only prevents the formation of NaOH and does not remove OH^- anions from the raw material.

SiCl_4 works by binding with oxygen atoms and forming very stable silicon dioxide (SiO_2). SiCl_4 vapor can further react strongly with water to form HCl and SiO_2 :



HCl will further react with NaOH according to Eq. 4.1 and produce water that in return feeds Eq. 4.2 until all NaOH is consumed.⁷ Therefore SiCl₄ will completely consume the water and possibly NaOH, and SiO₂ and NaCl proportional to initial amount of NaOH are produced as by-product.

SiO₂ will not dissolve in molten NaI and will float on top of the melt, but NaCl will dissolve in molten NaI (Fig. 3.12a). Although NaCl will separate from NaI during zone refining or crystal growth, its effect on the light yield of final NaI (Tl) crystal must be understood.

Another potential approach is to use SiI₄ instead such that NaI is formed as a by-product. However, with a melting point of 120.5 °C, SiI₄ is solid at room temperature, and it is more difficult to introduce into NaI powder than SiCl₄. It is also inferred that SiI₄ reacts less strongly than SiCl₄ [47].

4.4.1 Test with Crystal-Grade NaI Powder

To test the feasibility of using SiCl₄ to prevent NaI crystal from sticking to the crucible, the apparatus shown in Fig. 4.13 was constructed. Approximately 100 g of precision-dried crystal-grade NaI powder is loaded into 1" uncoated fused quartz crucible and about 25 mL of SiCl₄ is loaded into the Erlenmeyer flask under a fume hood. Prior to attaching the SiCl₄ to the vacuum adapter, air is removed from the Erlenmeyer flask by immersing it in liquid nitrogen and pumping with a scroll pump for 5 minutes. Valve 2, Valve 3 and Valve 4 are opened to pump the crucible to vacuum. After the pressure drops below 20 μTorr, Valve 3 is closed and Valve 1 and Valve 2 are opened to let in the SiCl₄ vapor. The powder is then heated slowly in an oven to 450 °C over 20 hours and then melted at 700 °C for 30 minutes and naturally cooled to room temperature.

⁷Although SiCl₄ has a strong affinity for oxygeneous anions, it is not known whether SiCl₄ vapor can react directly with NaOH.

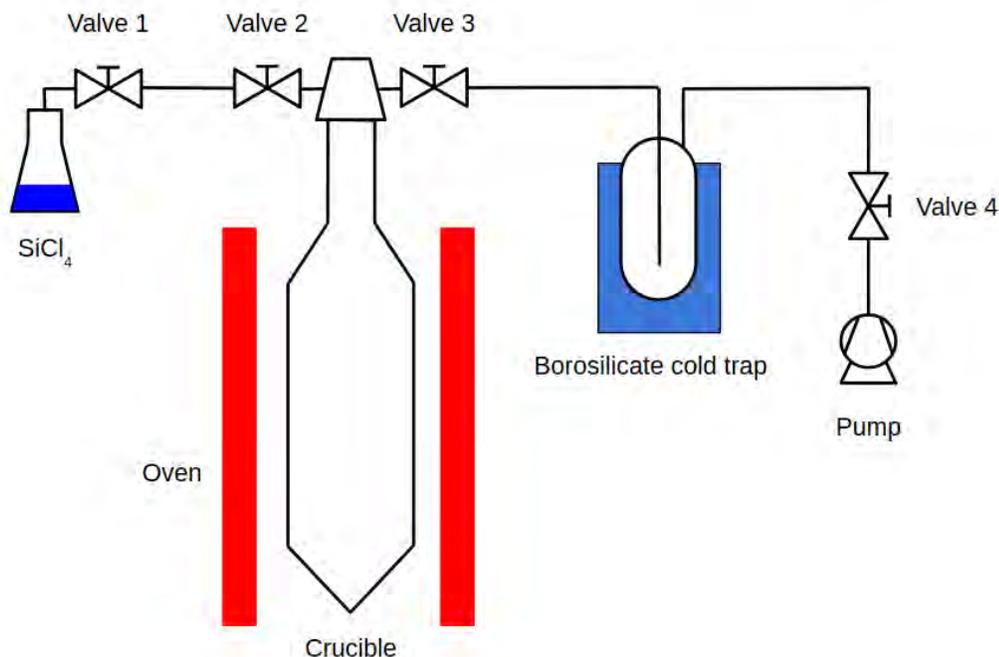


Figure 4.13: A setup to test the effectiveness of SiCl_4 in preventing NaI crystals from sticking to quartz crucibles. NaI powder is loaded in the fused quartz crucible and the crucible is connected to a vacuum adapter. The vacuum adapter has two valves that connects to an Erlenmeyer flask containing SiCl_4 and a pump through a borosilicate cold trap. After the crucible is pumped to vacuum, Valve 3 is closed and Valve 1 is opened to let SiCl_4 vapor into the crucible. The NaI powder is subsequently heated to above its melting point in the SiCl_4 vapor.

As the powder is heated and melted, violet vapor appears in the part of crucible outside the oven. This is likely due to Cl^- in HCl replacing I^- in NaI and the decomposition of HI into elemental iodine (I_2):



This could be an indication that NaOH is present in precision-dried powder. The fact that the amount of I_2 does not increase further shows that SiCl_4 does not directly consume NaI .

After cooling down to room temperature, the crucible is taken out for inspection. The crystal is free to slide within the crucible, which can be seen from the gap between bottom of crucible and crystal in Fig. 4.14a. A control group without SiCl_4 treatment resulted in the crystal firmly sticking to the crucible, thereby confirming the effectiveness of using SiCl_4 to prevent NaI crystals from sticking to uncoated quartz crucibles.



(a) SiCl_4 -treated Crystal-grade NaI powder crystallized in an uncoated quartz crucible without sticking. The crystal is a polycrystal due to rapid cooling. Compared to the control group, the quartz surface is also cleaner. The gap at the left end is created by sliding the crystal inside the crucible.



(b) Crystal-grade NaI powder without SiCl_4 treatment adheres strongly to fused quartz upon crystallization. The quartz is also etched by NaI/NaOH vapor.

Figure 4.14: NaI crystallized with and without SiCl_4 treatment. When the powder is treated with SiCl_4 , the crystal is free to slide without sticking after cooling. If the same powder is not treated with SiCl_4 , the crystal strongly adheres to the fused quartz crucible, indicating the importance of SiCl_4 in preventing sticking.

4.5 Conclusion on Crystal Growth

The successful growth of large NaI (Tl) single crystal with ultra-high purity and good scintillation performance is very difficult, and has never been achieved. However, through trials and errors, we have determined a number of key steps and procedures, and have achieved unprecedented purity compared to other experiments [23][52][53]. Yet intrinsic NaOH present in the Astro-grade NaI powder before precision-drying still remains a big threat to crystal growth using synthetic fused silica.

Currently, we have identified two potential solutions to this problem. Section 4.4 has shown that SiCl_4 treatment is a promising approach. However the amount of chloride introduced into the NaI crystal and its influence on scintillation (light yield, resolution, etc.) is still under investigation. HI treatment is another approach that does not introduce foreign atoms to the final NaI (Tl) crystal. However, HI is only available as an aqueous solution and obtaining anhydrous HI gas requires a more complex setup. The effectiveness of HI treatment will be studied in the near future.

Chapter 5

Crystal Detector Module

After the NaI (Tl) crystal has grown from melt, it must be cut to the appropriate size. The surface must be cleaned to remove contamination introduced during crystal cutting without seriously affecting the light collection efficiency. To operate the crystal as a scintillation detector, it must be optically coupled to a photon detector and packaged in a light-tight and leak-tight dry container to prevent degradation of the NaI (Tl) crystal due to moisture.

Section 5.1 describes surface cleaning with isopropyl alcohol and its potential impact on light collection efficiency. Section 5.2 summarizes tests of the PMTs to be used with the NaI (Tl) crystal. Section 5.3 describes how a crystal and two PMTs are packed inside a copper enclosure as a detector module. Finally, Section 5.4 outlines the cleaning and assembly of the detector module and scintillation tests performed with NaI-031.

5.1 Crystal Polishing

Cutting will not only roughen the crystal surface, but also introduce surface contamination. The traditional polishing method using sandpaper of various grit sizes may yield a better surface, but will make it more difficult to achieve high purity. On the

other hand, NaI has moderate solubility in isopropyl alcohol and isopropyl alcohol is readily available in very high purity. Therefore, two cubical NaI (Tl) crystals, NaI-013 and NaI-026, were used to study the effect of isopropyl alcohol polishing and different reflector materials.

The setup used in this study is shown in Fig. 5.1. NaI-013 is $1.5'' \times 1.5'' \times 1''$ and NaI-026 is $1'' \times 1'' \times 1$ cm in dimension. The four side surfaces and the bottom surface of the crystal sample are covered with reflectors and the top surface is optically coupled to a Hamamatsu R11065 PMT using optical grease.

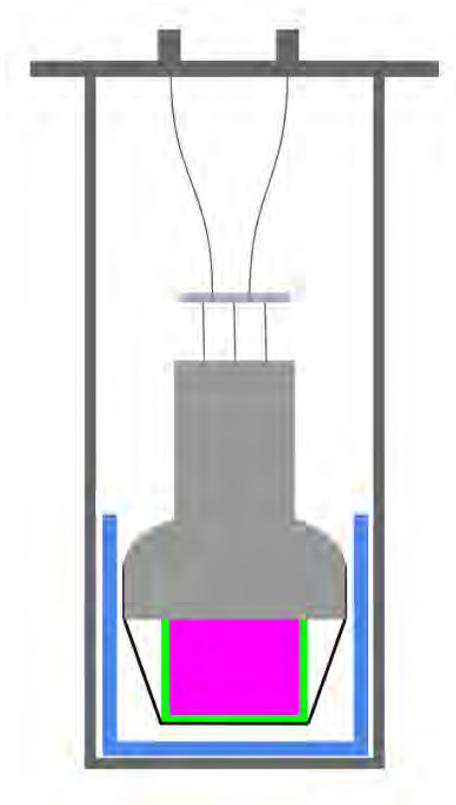


Figure 5.1: Setup to test different surface treatments and reflector materials. Cubical NaI (Tl) crystal (purple) is covered with reflectors (green) on four sides and bottom. The top surface of the crystal is optically coupled to a R11065 PMT (grey) using optical grease. The crystal is fixed against the PMT by electrical tape (black). The entire assembly is placed inside a PTFE cup (blue) in a stainless steel enclosure. The PMT is operated with negative high voltage and the cathode is insulated from the stainless steel enclosure by the PTFE cup.

Different surface treatments, reflector materials and results are summarized in Table 5.1. For lapped surface, the crystal is sanded with 2000-grit sandpaper followed by 8000-grit sandpaper to create a translucent frosted-glass-like finish (Fig. 5.2a). The isopropyl alcohol surface finish is created by wiping with isopropyl alcohol-soaked Kimwipes followed by dry Kimwipes. isopropyl alcohol will dissolve a thin layer of the crystal surface, yielding an optically-transparent surface (Fig. 5.2b). If the isopropyl alcohol on the surface is not wiped away but left to dry naturally, evaporation of isopropyl alcohol leaves dissolved NaI as powder residues on the surface, thereby creating a surface similar to but more transparent than a lapped surface (Fig. 5.2c).

Table 5.1: The effect of different polishing methods and reflector materials on the crystal light yield. A lapped surface is created by sanding with 2000-grit sandpaper followed by 8000-grit sandpaper. An isopropyl alcohol surface finish is created by wiping with isopropyl alcohol-soaked Kimwipes and left to dry naturally. In PTFE runs, 10 layers of plumber’s PTFE tape is used. In 3M foil and LumirrorTM (Lumirror) runs, reflectors are cut to the size of crystal surfaces and fixed in position by additional 5-6 layers of PTFE tapes. The optical grease in the surface finish refers to optical grease between the crystal surface and the reflector, and IPA refers to isopropyl alcohol.

ID	Surface Finish	Reflector	Light yield (p.e./keV)
NaI-013-1	lap	PTFE	12.0 ± 1.0
NaI-013-2	IPA	PTFE	11.14 ± 0.5
NaI-013-3	IPA+optical grease	PTFE	11.44 ± 0.01
NaI-026-1	lap	PTFE	10.08 ± 0.01
NaI-026-2	lap	3M foil	7.73 ± 0.01
NaI-026-3	lap	Lumirror	11.94 ± 0.01
NaI-026-4	lap+optical grease	Lumirror	10.27 ± 0.01

The highest light yield is achieved with a lapped surface and a high-reflectivity specular reflector such as Lumirror. Although 3M foil also has a relatively high reflectivity, the cutoff wavelength is very close to the scintillation wavelength of NaI (Tl). When optical grease is applied on the lapped surface, the grease fills microscopic frac-

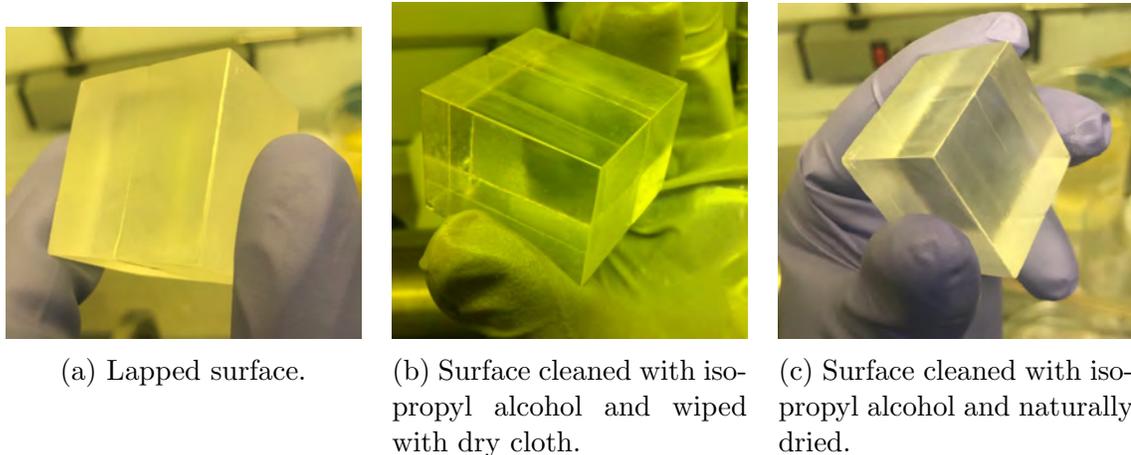


Figure 5.2: Photo of NaI (Tl) crystal with lapped and isopropyl alcohol-cleaned surfaces. Lapped surface has a translucent appearance similar to frosted glass. When isopropyl alcohol is quickly wiped away after cleaning, a transparent surface is achieved. Leaving the isopropyl alcohol to dry by itself yields an intermediate surface.

tures and the surface becomes transparent. The fact that light yield has decreased in NaI-026-4 indicates that a diffuse reflector is necessary. The comparison of NaI-013-2 and NaI-013-3 supports this claim: since PTFE is already a good diffuse reflector, the application of optical grease does not make much difference.

isopropyl alcohol can be used to clean the crystal by dissolving and removing a surface layer of material. If high-purity isopropyl alcohol is used, very limited contamination can be introduced. Although light yield is adversely affected according to NaI-013-2, it is still sufficient for a keV threshold provided the final crystal's larger dimension does not further degrade the light yield significantly. Therefore wiping with isopropyl alcohol is a surface cleaning technique that can preserve a relatively high light yield without introducing significant surface contamination.

It must be pointed out that the results reported in Table 5.1 could be subject to a large systematic error. To do a preliminary analysis of reproducibility of these surface finishes, NaI-031-1 and NaI-031-2 were repeated. For the isopropyl alcohol cleaning of the surface, two runs were consistent to within one photoelectron, but the lapped surface finish varied by about 2 photoelectrons.

5.2 Photomultiplier Tubes

To characterize the Hamamatsu R11065-20 PMTs to be used in the detector module, the test PMT is placed facing a PTFE cup in a dark box. A blue light emitting diode (LED) with 420 nm peak wavelength is used as light source. Two polarizing films with adjustable relative angle are placed between the PMT cathode and the LED to attenuate the LED light as needed. The setup is shown in Fig. 5.3.

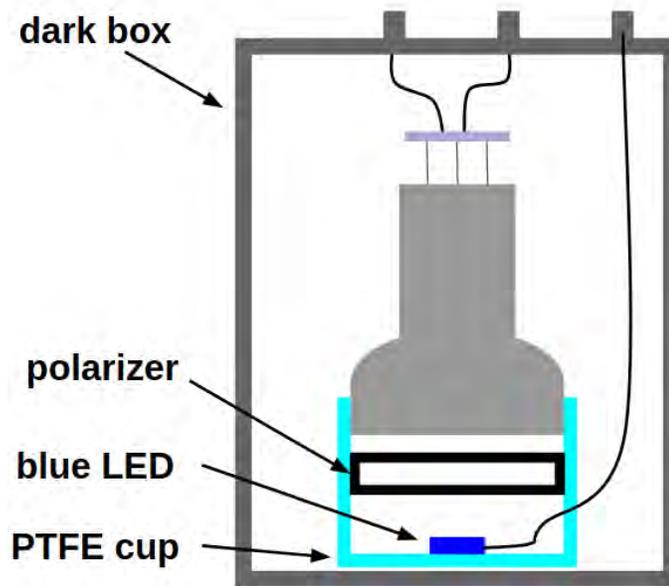


Figure 5.3: Setup to characterize Hamamatsu R11065-20 PMTs. A blue LED is placed inside a PTFE cup, and two linear polarizers are placed between the PMT and the LED to adjust the intensity of the light. The LED is pulsed with transistor-transistor logic (TTL) signals from a function generator.

5.2.1 Gain

To measure PMT gain as a function of applied high voltage, the PMT high voltage is increased from 1100 V to 1700 V at 50 V steps. Before each measurement, the PMT is allowed to stabilize for 30 minutes. The LED is pulsed at 5 kHz frequency by 5-V, 100-ns pulses from a function generator. The PMT signal is directly digitized by a CAEN V1720 analog-to-digital converter (ADC) board and the data acquisition is

triggered by the PMT waveform crossing the threshold which is set at about halfway between baseline and the minimum of the pulse. PMT gain at different voltages is computed as the total number of electrons in a single photoelectron pulse at 1700 V scaled by the relative gain determined by the LED pulse. The result is shown in Fig. 5.4. This information will be used in the future to select pairs of PMTs and determine their operation voltages.

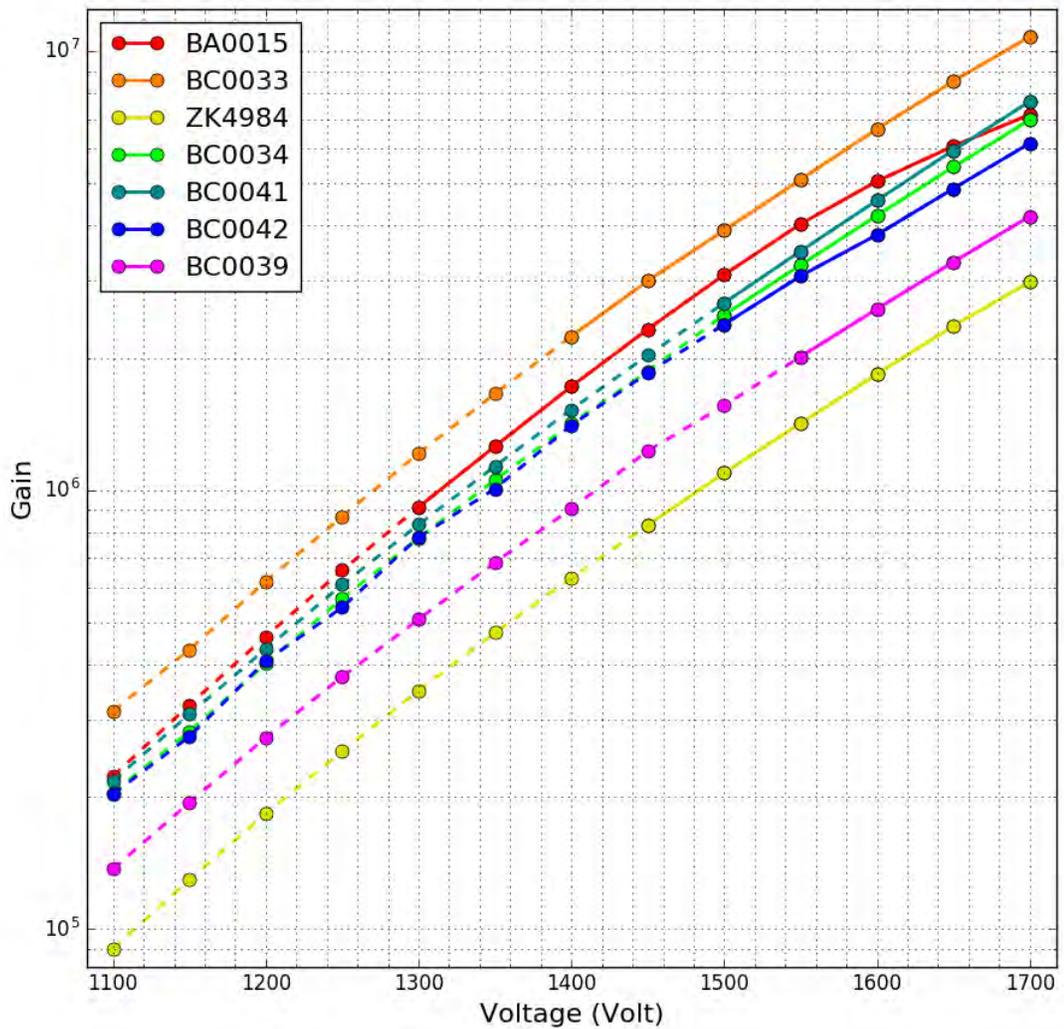


Figure 5.4: R11065-20 PMT gain as function of applied voltage. Dashed lines indicate single photoelectron is not distinguishable from noise without amplification.

5.2.2 Dark Rate

To measure the PMT dark rate, the cathode of the test PMT is covered with a black plastic cap and the PMT is placed inside a dark box for at least 30 minutes. After each high voltage adjustment, the PMT is allowed to stabilize for another 30 minutes. For each run, 2000 periods of 5-ms-long continuous waveforms are recorded and the number of single photoelectrons is counted offline by software using the algorithm described in Section 8.4. Since a PMT of this model has a typical dark rate of a few hundred Hz at a gain of 10^6 , a 10 second live window will record a few thousand samples and the computed dark rate can have an error of a few percents. However this method is not biased by the choice of threshold or the drift of baseline.

The measured dark rates of R11065-20 PMTs at a gain of about 2×10^6 is summarized in Table 5.2. The pulse height criteria in identifying single photoelectron is deviation of at least 5 ADC counts from baseline. The measurement voltage guarantees that unamplified single photoelectron pulses are still distinguishable from baseline noise.

Table 5.2: Measured dark rates of R11065-20 PMTs at a gain of about 2×10^6 . Since no amplifiers are used, identification of single photoelectrons rely completely on pulse-finding algorithm. The listed voltages at which measurements are conducted guarantee that single photoelectrons are distinguishable from baseline noise.

Serial No.	Voltage (V)	Dark rate (Hz)
BA0015	1400	366 ± 6
ZK4984	1600	1068 ± 12
BC0033	1450	543 ± 7
BC0034	1550	540 ± 7
BC0039	1600	656 ± 8
BC0041	1550	850 ± 9
BC0042	1600	911 ± 10

5.2.3 Afterpulsing

The location of PMT afterpulsing is obtained by averaging approximately one million PMT pulses obtained at an operation voltage of 1700 V. A constant baseline is used throughout the whole process. PMT pulses are aligned by the first main-pulse sample 2σ away from the baseline, where σ refers to the standard deviation of the baseline. Averaged PMT waveforms are shown in Fig. 5.5.

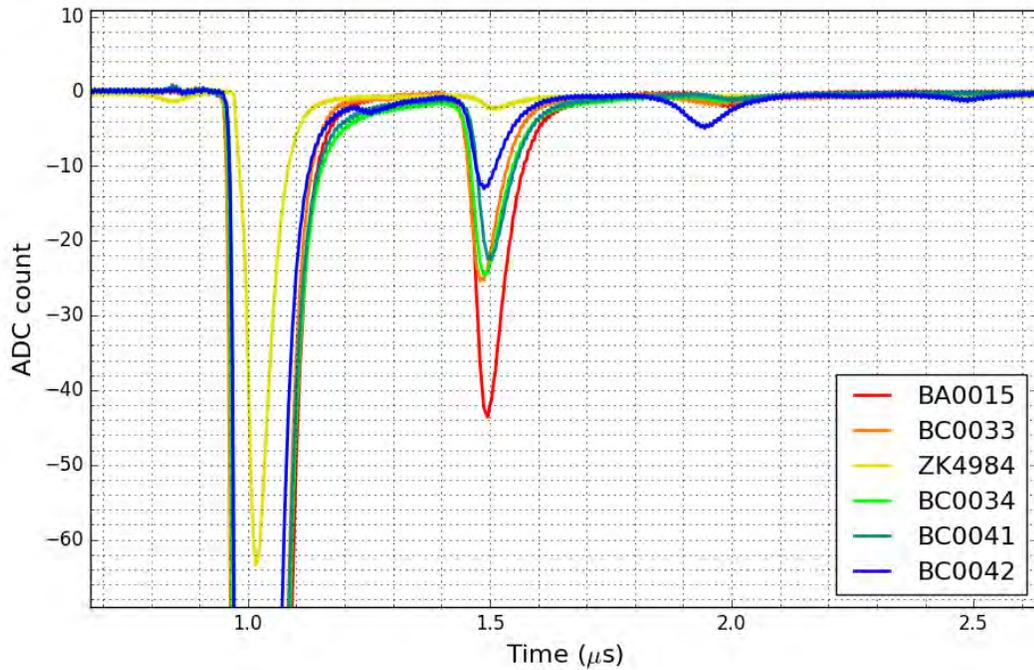


Figure 5.5: Average waveform corresponding to LED pulses. Data is taken at 1700 V.

PMT afterpulsing is characterized as the ratio of afterpulse area to that of main pulse. This ratio roughly reflects the probability of having an afterpulse in each of the regions of concern. The result is summarized in Table 5.3

Table 5.3: Probability of afterpulse for R11065-20 PMTs at 1700 V.

Serial No.	Afterpulse region (μs)		
	1.4-1.7	1.9-2.2	2.4-2.7
BA0015	0.027	-	-
ZK4984	0.032	0.0048	-
BC0033	0.045	0.0034	-
BC0034	0.087	0.069	0.0056
BC0039	0.038	0.012	0.0064
BC0041	0.041	0.0043	-
BC0042	0.044	0.0045	-

5.3 Crystal Enclosure

5.3.1 Design

The design of a crystal detector module must address the following factors:

1. Material with low radioactivity should be used to reduce background radiation to the crystal.
2. The enclosure should use minimal material so as not to absorb excessive γ -rays and affect veto efficiency.
3. The enclosure should be leak-tight and have low humidity inside. Components used in the assembly should be thoroughly baked out. Otherwise the NaI (Tl) crystal will degrade due to absorption of outgassed moisture (Fig. 5.6).
4. The PMTs coupled to the crystal employs negative high voltage for a cleaner baseline, therefore the PMT bodies must be well-insulated from other components.

5. During shipping and commissioning, the PMTs and the crystal must maintain good mechanical contact for optical coupling. In addition, the crystal-PMT assembly should be spring-loaded to absorb shocks during shipment.

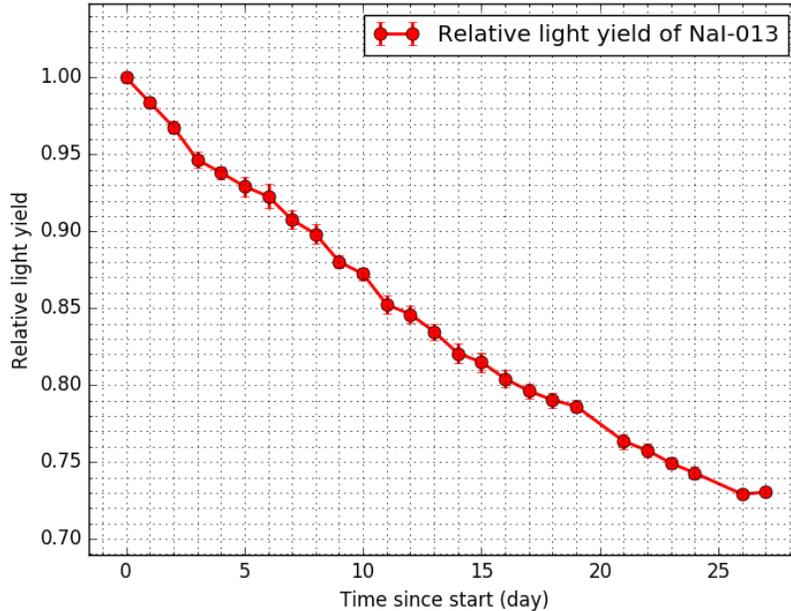


Figure 5.6: A $1.5'' \times 1.5'' \times 1''$ NaI (Tl) crystal is coupled to a PMT in a leak-tight stainless steel structure. The crystal is wrapped in 10 layers of PTFE tapes and the assembly is insulated from the container by PTFE structures. The PTFE components were not previously vacuum-baked to remove moisture. The light yield of the NaI (Tl) crystal is degrading over time due to moisture.

Based on these requirements, a crystal enclosure system was designed and manufactured. A CAD drawing of the detector module assembly is shown in Fig. 5.7.

To keep the background low, the copper tube is made out of special low-radioactivity copper while all other copper parts are constructed out of C10100 copper (also known as oxygen-free high-conductivity copper). The crystal holders, PMT holders and PMT sleeves are made out of Delrin for mechanical strength and electric insulation. Springs can have relatively high radioactivity per unit mass, but the mass is small and it is farther away from the crystal. Its contribution to the background is expected to be small.

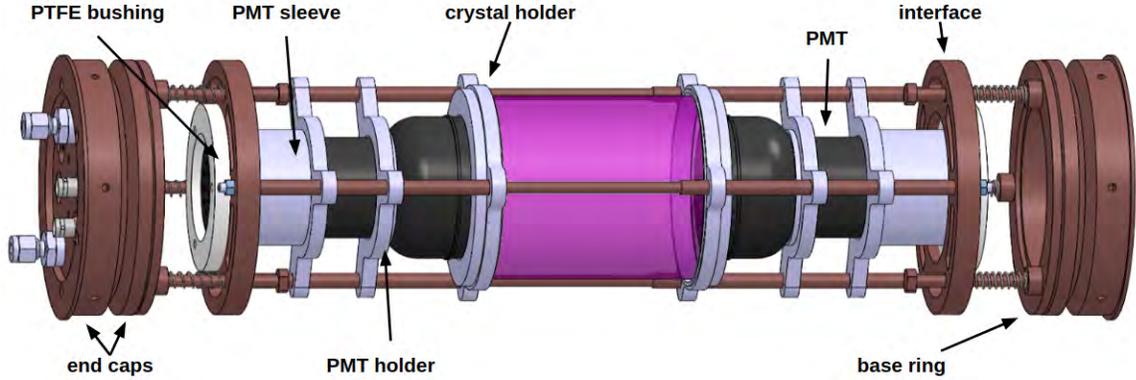


Figure 5.7: CAD drawing of detector module assembly. The crystal is held in place by two crystal holders and three copper rods; each copper rod comprises three segments joined by threads and the crystal holder is pinched at the joint. Two PMTs are fixed against the crystal by a PMT holder, a PMT sleeve and three equally-spaced springs. The PMT sleeves are pushed forward by the PTFE bushings. The copper interface plate has threaded holes to support the three copper rods and the PTFE bushing. The interface is spring-loaded against the copper base ring sitting next to the end caps. The entire assembly sits inside a copper tube (not shown). At each end of the tube, two end caps sandwich an O-ring by their chamfered edges. When the two end caps are tightened with screws, the O-ring expands laterally and makes the seal against the copper tube.

Electrical and fluid feedthroughs are located on the end cap on one side. Two Swagelok connectors are used to purge the detector module with dry nitrogen. High voltages for PMTs are passed from the SHV bulkhead feedthroughs to the PMT bases via Kapton cables, and PMT signals are transmitted using RG-178 cables to the SMA (SubMiniature version A) bulkhead feedthroughs.

5.4 Cleaning, Assembly and Test of Crystal Detector Module

The copper enclosure and all internal components are thoroughly cleaned before assembly. Copper parts are cleaned with Detergent 8 to remove grease on the surface. After thorough rinsing with deionized water, the parts are cleaned with 3% v/v HCl acid, 1% m/m citric acid and deionized water. The HCl removes surface contaminants

by dissolving away a thin layer of copper surface, and the citric acid passivates the copper surface to prevent further oxidization. After cleaning, the copper parts are baked at 80 °C for 5 days under constant flow of dry nitrogen to remove moisture.¹

The PTFE tape, PTFE bushing, Delrin PMT sleeves, Delrin PMT holders and crystal holders are cleaned with Detergent 8 and rinsed thoroughly. These parts are then leached with 10% v/v HCl acid and vacuum-baked at 80 °C for 5 days.

To prevent degradation of the NaI (Tl) crystal and contamination by particulates, the assembly is carried out in a glovebox inside a class-100 cleanroom at Jadwin Hall, Princeton University.

It was noticed that there were some residual carbon flakes embedded on one side of the crystal (see Fig. 5.8a) and the other side had developed a yellow-green color, possibly due to TlI (Fig. 5.8b). These carbon flakes and discoloration will affect light yield by absorbing incident scintillation photons. To remove the embedded carbon, the crystal surface was sanded with nylon wool and SiC-embedded nylon mesh. Nylon and SiC are chosen because they are less likely to introduce radioactive contamination compared to alumina-based ones. Although carbon flakes were effectively removed (see Fig. 5.8c) , the other half of the cylindrical surface was too rough and discoloration was too deep beneath the surface and cannot be removed by sanding with nylon.

After sanding, the surface is cleaned with isopropyl alcohol a few times to remove the debris. The crystal is wrapped in 10 layers of PTFE tape and optically coupled to two R11065 PMTs via optical grease. The assembly of the crystal, two PMTs and internal supporting structures are shown in Fig. 5.9.

After the crystal detector module is assembled, it is placed vertically behind walls of lead bricks which is 4" thick on the bottom and four sides, and at least 2" on the top. Initial characterization of scintillation was carried out with uncollimated ¹³⁷Cs,

¹Vacuum-baking was not permitted since the thin wall of the external copper shell was not designed to withstand vacuum.

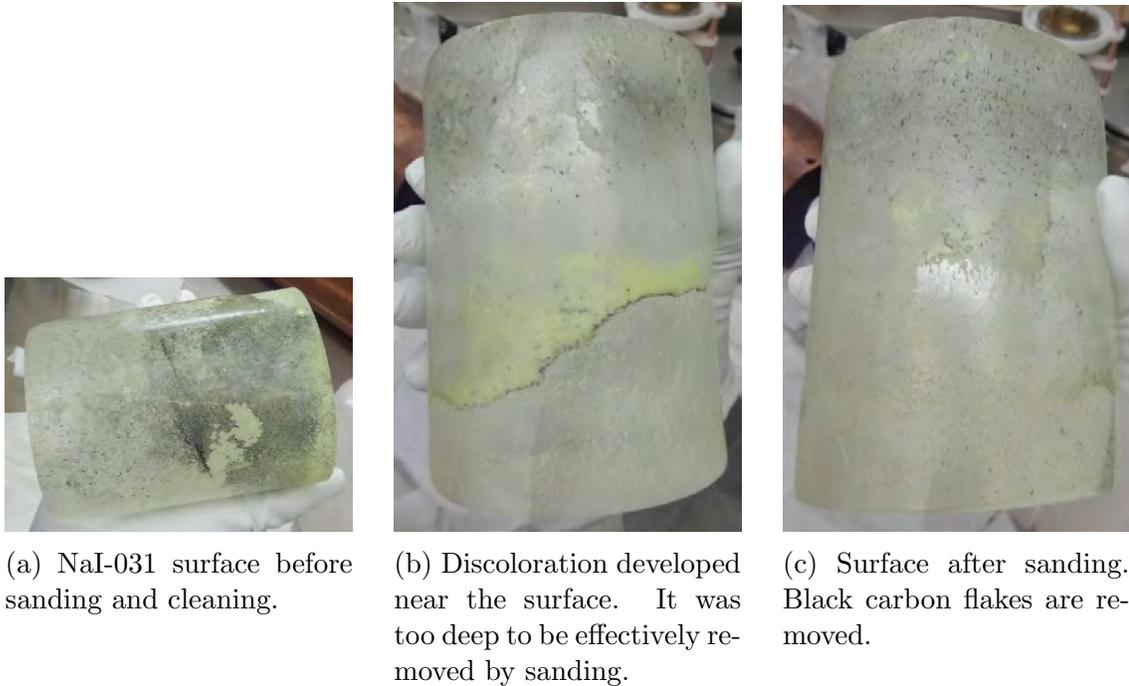


Figure 5.8: Surface of NaI-031 before and after sanding.



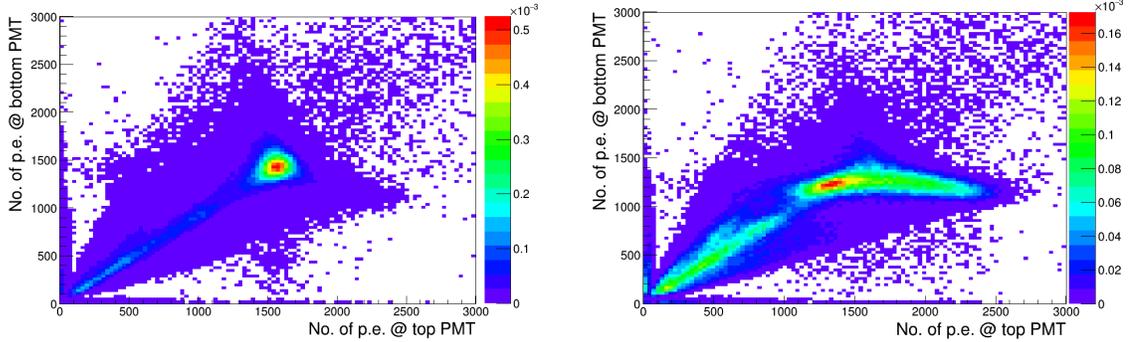
Figure 5.9: Photo of the NaI (Tl) crystal, two PMTs and internal supporting structures. The crystal is wrapped with 10 layers of PTFE tape as reflector. The structure is placed inside a copper tube and sealed on both ends to make a detector module.

but light yield was found to be about 5 p.e./keV, rather low compared to that of typical NaI (Tl) crystals. The 661.7 keV photopeak was also very broad.

To see if this is related to the nonuniformity of the crystal, the ^{137}Cs source was collimated through a $1/4''$ hole on a $2''$ -thick, $4'' \times 4''$ lead brick.

The crystal's response at different axial positions is studied by moving the collimated γ -source vertically at $1''$ steps. It was discovered that when the source was placed within $4''$ away from the bottom of the crystal, both PMTs recorded similar

spectra (Fig. 5.10a). However when the source was moved closer to the top PMT by $1''$, the top PMT began to see almost twice as much light (Fig. 5.10b).



(a) Collimated ^{137}Cs source is placed approximately $4''$ above bottom surface of crystal.

(b) Collimated ^{137}Cs source is placed approximately $5''$ above the bottom surface of the crystal, close to the top surface of the crystal.

Figure 5.10: 2D histograms of numbers of photoelectrons detected by the bottom (BA0015) and the top (ZK4984) PMTs. The color scale indicates count rate in Hz. The position of collimated calibration source differs by only $1''$.

It is worth pointing out that when the γ source is placed near the bottom surface, the spectra recorded by both PMTs are similar. Azimuthal scans were also done at 4 equally spaced angles at both $4''$ and $5''$ above bottom, and the responses were identical along the circumference of the crystal.

If the asymmetry is due to surface defects, a relatively large azimuthal dependence is expected. The absence of azimuthal dependence indicates that the asymmetry is likely due to intrinsic crystal nonuniformity, which could be caused by non-ideal crystal surface, or insufficient mixing of NaI with TlI, which is also hinted by the green-yellow color on the surface.

Chapter 6

Veto

The veto detector consists of a stainless steel cylindrical vessel filled with 2 tons of pseudocumene liquid scintillator and ten R5912 photomultiplier tubes. The veto detector functions as an active shielding that stops and tags escaping γ -rays from radioactivity in the crystal detector and penetrating γ -rays and neutrons from the experimental hall.

External γ -rays and neutrons are further suppressed by a combination of lead, water and polyethylene shielding. To suppress radon, the vessel is placed inside a hermetically-sealed radon box, and both the radon box and the crystal detector module are constantly purged with dry nitrogen.

To prevent accidental leakage of pseudocumene, the veto vessel is seated in a catch basin inside the radon box. In addition, fluid handling and slow control infrastructure is implemented to circulate, monitor and, in case of leak, drain pseudocumene.

Section 6.1 describes the veto vessel and Section 6.2 describes the PMTs used to detect the scintillation light in the veto detector. The shielding and fluid handling system are presented in Section 6.3 and Section 6.4, respectively. Finally, Section 6.5 discusses estimation light yield of the veto detector at Princeton and installation of

the veto PMTs, trigger and electronics in LNGS. Discussion of the trigger algorithm and data acquisition is reserved for Chapter 7 and Chapter 8.

6.1 Veto Vessel

The veto vessel, shown in Fig. 6.1a, is 54" in diameter and 60" long. It is designed to hold 2 tons of pseudocumene purified by Borexino pseudocumene purification plant. 3 g/L of 2,5-diphenyloxazole (PPO) is added to the pseudocumene as wavelength shifter. The vessel wall is constructed out of 1/4"-thick low-radioactivity stainless steel. The inner surface of the vessel is coated with a layer of ethylene tetrafluoroethylene to prevent degradation of pseudocumene due to the stainless steel. The inner surface also has a number of threaded studs to hold highly-reflective Lumirror lining (Fig. 6.1b) to enhance light collection efficiency.



(a) The liquid scintillator veto vessel in LNGS Hall B.

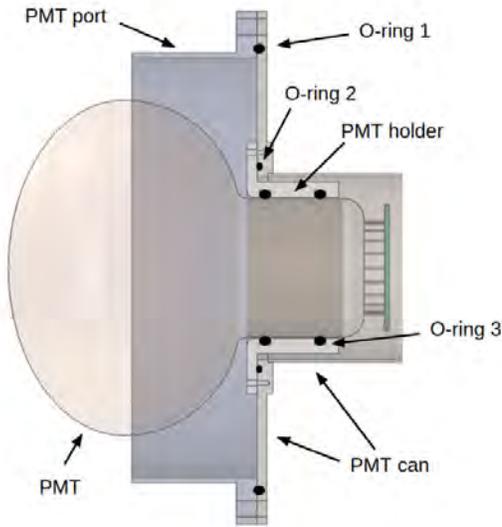


(b) Lumirror lining to improve scintillation light collection efficiency.

Figure 6.1: Photo of veto vessel at LNGS Hall B, and its Lumirror lining on the inner surface.

To collect scintillation light, five 8-inch Hamamatsu R5912 PMTs are installed at each end of the veto vessel. The borosilicate glass PMTs are mounted inside stainless steel PMT cans via polyethylene PMT holders, and the PMT cans seal to the ASA

flange flanges on the ends of the veto vessel (Fig. 6.2). Viton O-rings are used for all O-ring seals for chemical compatibility with pseudocumene. Two half nipples with KF16 flanges are welded onto the sides of each PMT can to circulate dry nitrogen. The end cap of the PMT can has electrical feedthroughs to apply high voltage bias to and collect signal from the veto PMT. An SHV bulkhead connector with solder cup termination and PTFE-insulated single wire are used for high voltage bias while an SMA (SubMiniature version A) connector and RG-174 cable are used inside the PMT can for signal transmission.



(a) CAD drawing of the veto PMT, PMT can and PMT holder. The PMT holder is used to prevent borosilicate glass from contacting stainless steel directly.



(b) Photo of 8" R5912 PMT mounted onto the PMT holder and PMT can.

Figure 6.2: CAD drawing and photo of 8" PMT, PMT holder and PMT can. The PMT is placed inside a polyethylene PMT holder which is in turn fixed onto the stainless steel PMT can by screws. The PMT can is mounted onto the veto vessel via Viton O-ring face seal. The PMT is sealed against the PMT holder via two Viton O-ring piston seals in series for redundancy.

The top cover plate on the veto vessel is welded to a blind-end high-purity copper tube that extends to the center of the vessel. This copper tubing is used to place the crystal detector module in the center of the veto detector. The flange that seals the copper tube has gas feedthroughs such that the copper tube can be purged to reduce

radon and water around the crystal detector module. Additional details regarding the design and construction of the veto vessel are given in [30].

6.2 Veto Photomultiplier Tubes

The veto PMTs have a relatively high gain and a high quantum efficiency (around 25%), and are designed to be operated with high voltage bias without any further amplification. In testing the divider boards, it was found that ceramic capacitors would cause the baseline to oscillate after each pulse. Such “ringing behavior” has been previously reported by [54]. After the ceramic capacitors were replaced with polypropylene capacitors, the noise disappeared. This is shown in Fig. 6.3. Similar ringing phenomena was also observed for a NaI (Tl) crystal using Hamamatsu R11065 3" PMT with positive high voltage and ceramic decoupling capacitors and was solved by switching to polypropylene capacitors.

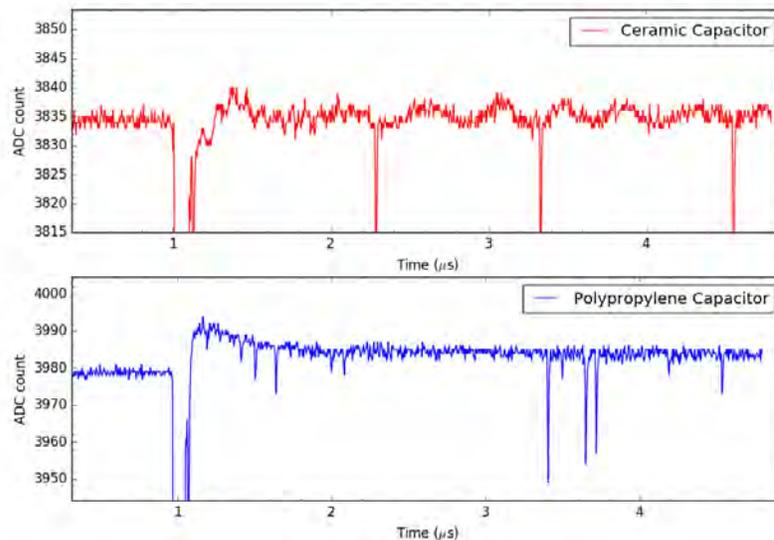


Figure 6.3: R5912 PMT pulses recorded with CAEN V1720 digitizer board. The top (red) waveform is recorded with a divider board using ceramic capacitors; the bottom (blue) waveform is recorded with a divider board using polypropylene capacitors. The ringing present in the top waveform after $2.5 \mu\text{s}$ is absent from the bottom waveform.

To measure the gain of each veto PMT, varying high voltage is added to the divider board of each PMT at 50 V steps and single photoelectron pulses from dark counts are recorded. The number of electrons in each single photoelectron pulse is computed by dividing the integral of the pulse by the resistance of the readout resistor. The gain of each PMT is then computed as the mean number of electrons in single photoelectron pulses. Veto PMT gains as functions of applied voltages are shown in Fig. 6.4, and will be used in the future to determine the operational voltages of veto PMTs. The pulse-finding algorithm is discussed later in Section 8.4.

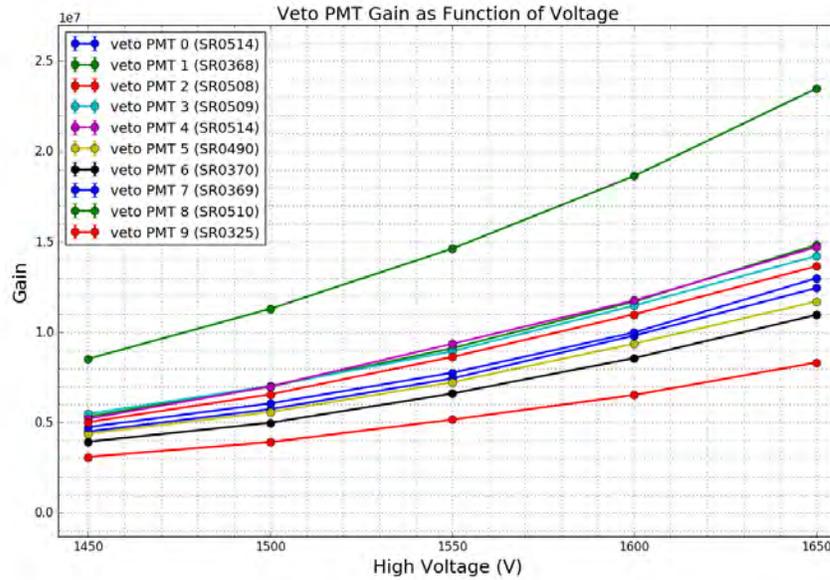


Figure 6.4: Hamamatsu R5912 veto PMT gain as function of applied positive high voltage. For this measurement, single photoelectrons from dark counts are used.

6.3 Shielding

To determine the amount of necessary shielding, Geant4 simulation is performed to evaluate the shielding power of water and lead of various thicknesses.

The flux of γ -rays from the surface of the rocks in the experimental hall is obtained by simulating the radioactive decays of ^{40}K , ^{238}U and ^{232}Th and propagation of decay

products in the rock. The concentrations for these isotopes are taken from [55]. Shown in Fig. 6.5, the flux of γ -rays that escape the rock surface is used to predict the flux after 20 cm of lead and 90 cm and 2 m of water.

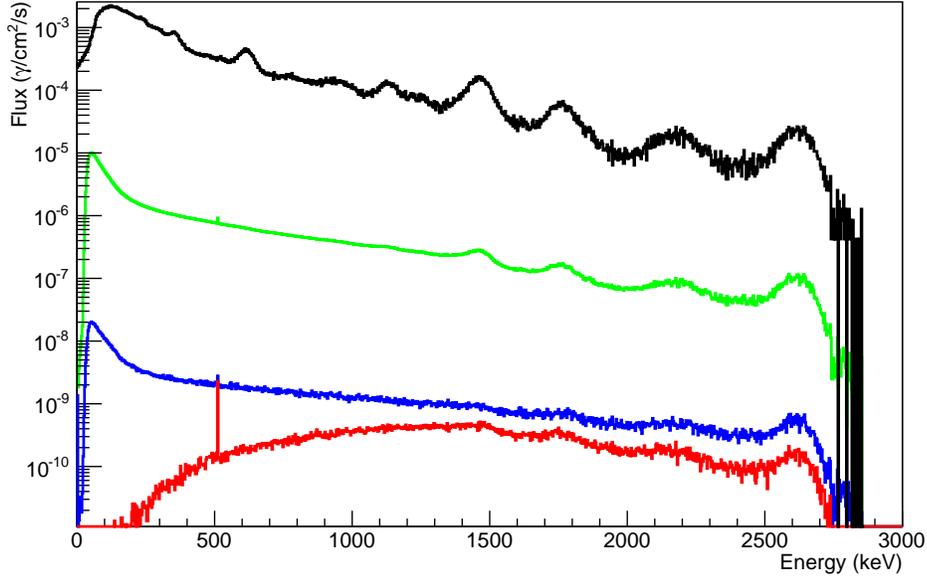


Figure 6.5: γ -ray flux after different shielding materials and thicknesses: no shielding (black), 90 cm water (green), 2 m water (blue) and 20 cm lead (red).

To determine the amount of shielding required, the γ -ray flux in Fig. 6.5 is used to predict the background rate in the dark matter search due to external γ -rays hitting the crystal without being detected by the veto detector. Although the lowest background is achieved with 20 cm of lead, 90 cm of water yields a background as low as 1×10^{-3} cpd/kg/keV, two orders of magnitude lower than the projected background due to the crystal radioactivity (~ 0.1 cpd/kg/keV).

Due to the availability of materials and engineering constraints, the actual shielding consists of 10 cm of polyethylene and 80 cm of water on the top, 40 cm of polyethylene and 91 cm of water on the four sides, and 15 cm of lead on the bottom. This new configuration still guarantees approximately 90 cm of water-equivalent shielding

power and the external background is similar to that obtained with 90 cm of water on every side. Different shielding components are shown and labeled in Fig. 6.10.

6.4 Fluid Handling and Slow Control

During normal operations, the fluid handling infrastructure monitors and regulates about 2 tons of pseudocumene and two grades of dry nitrogen. Fig. 6.6 shows a simplified view of the fluid handling system. Regular high-purity nitrogen is used to purge each veto PMT cans and the radon box while the volume near the NaI (Tl) crystal is purged with special low-argon, low-krypton nitrogen to suppress radon and to prevent the NaI (Tl) crystal from degrading due to outgassed moisture. Used gas is filtered with charcoal before it is vented. Pseudocumene leakage is detected with flammable gas sensors at various locations and a liquid sensor at the catch basin. Pseudocumene is filled and drained with two hydraulic pumps and a third hydraulic pump is used to pump pseudocumene in the catch basin into a waste tank in case of emergency.

6.5 Commissioning

After the veto vessel was manufactured at Allegheny Bradford Company, the vessel was leak-checked and the veto detector's light yield was measured at Princeton before shipment to LNGS in November, 2016. Subsequently in March, 2017, all 10 veto PMTs and the trigger and data acquisition systems were installed and an integral test was performed at LNGS Hall B.

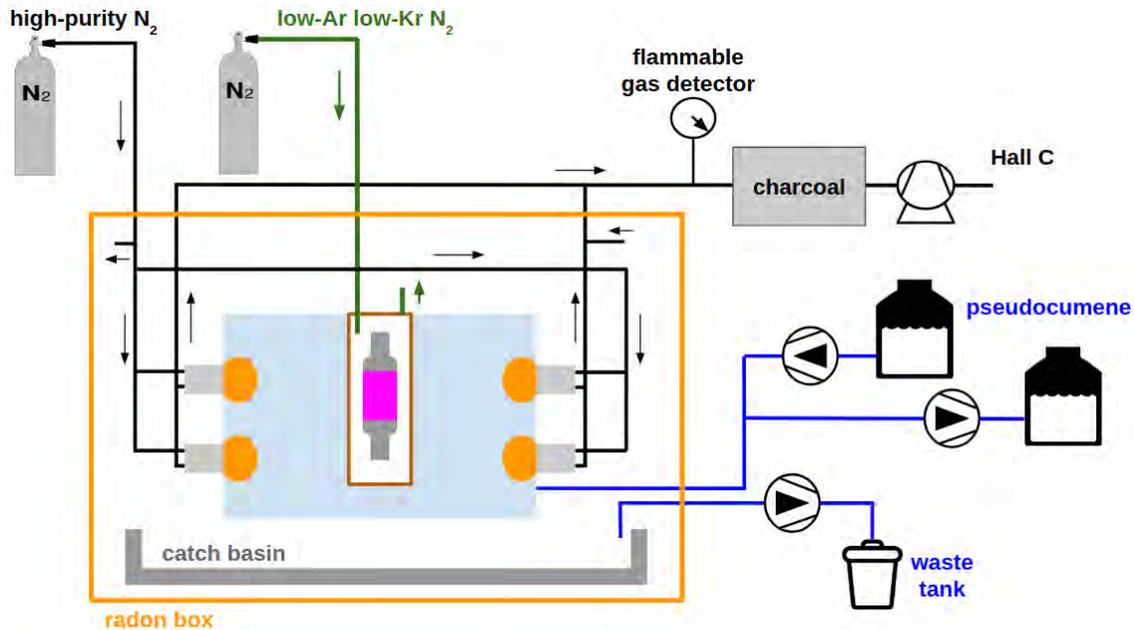


Figure 6.6: Simplified schematics of the fluid handling system. For simplicity, various valves, pressure indicators, liquid level sensors, etc. are not shown.

6.5.1 Leak Detection at Princeton

Since the vessel is not designed for vacuum, initial leak detection was done with helium pressure decay. However, the results were inconclusive due to a lack of precision, the vessel's relatively large volume and fluctuations in ambient pressure. To overcome these challenges, the PMT ports were sealed from the vessel body by O-ring piston seals and helium vacuum leak detector was used to check the leak-tightness of each port locally. No leak bigger than 2.7×10^{-7} mbar L/s was observed, compatible with leak-tightness requirement of pseudocumene.

6.5.2 Estimation of Veto Detector Light Yield

Instead of installing all 10 veto PMTs and filling the veto vessel with 2 tons of flammable and carcinogenic liquid, a quartz scintillation vial and one veto PMT are used to estimate the light yield by scaling: a 3" quartz cell filled with a mixture of pseudocumene and PPO was placed about 20 cm away from one veto PMT while all

other PMT ports were covered with Lumirror foil. A ^{137}Cs source was placed next to the quartz cell and the Compton edge of 661.7-keV γ -ray is fitted with Geant4-simulated spectrum (Fig. 6.7). The light yield under this configuration is estimated to be 0.66 photoelectrons per keV (p.e./keV).

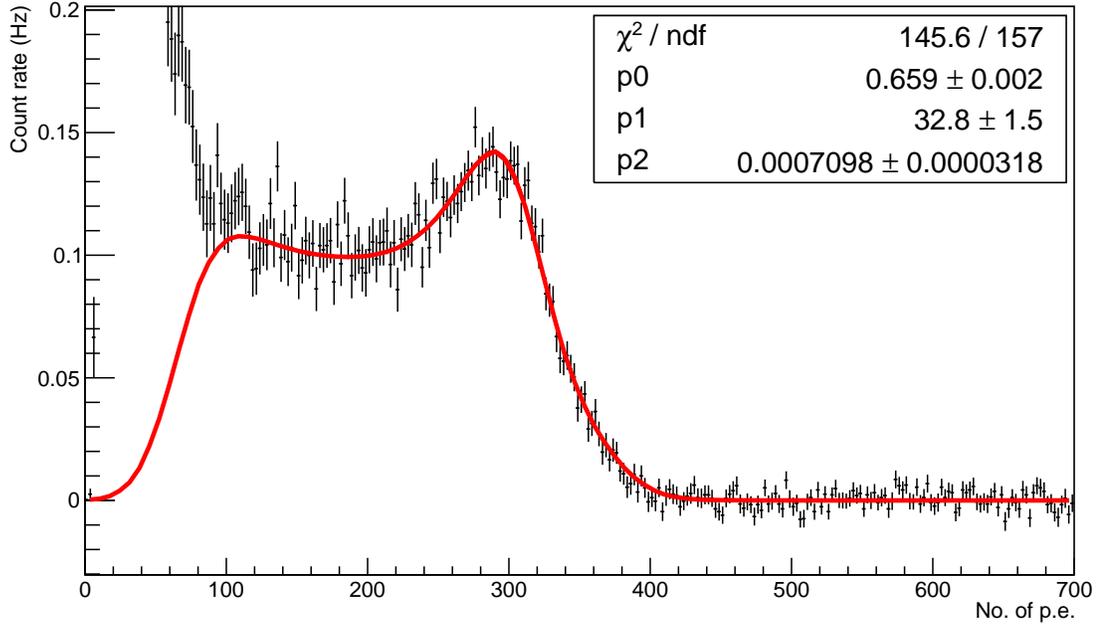


Figure 6.7: ^{137}Cs Compton edge measured with a single veto PMT and a pseudocumene-filled 3" quartz cell inside the Lumirror-coated veto vessel. The measured spectrum (black) is fitted with Geant4-simulated spectrum (red). The light yield in this configuration is calculated to be 0.66 p.e./keV and the resolution around the Compton edge is 33 keV. The deviation at low energy is due to gamma-rays scattering off the steel vessel, which is not considered in the simulation.

The result in Fig. 6.7 is scaled accordingly to estimate the light yield with 10 PMTs and 2 tons of pseudocumene. PPO's average emission wavelength is estimated to be about 385 nm and reflectivity of Lumirror at this wavelength is 0.956 [56]. The attenuation length in pseudocumene is taken as 4.5 m [57]. Let Ω be the solid angle of one PMT viewed from the source, r the reflectivity of Lumirror at 385 nm, and l attenuation length in pseudocumene, the fraction of scintillation photons arriving at

veto PMTs is

$$\eta = \eta_{\text{direct}} + \eta_{\text{reflected}} \quad (6.1)$$

$$= N_{\text{pmt}} \Omega / 4\pi \times e^{-L/l} + r^{\bar{n}} \times e^{-\bar{n}L/l} \quad (6.2)$$

where \bar{n} is average number of reflections before a scintillation photon hits a PMT. It is estimated as the ratio of Lumirror-lined surface area (65669 cm² minus photocathode total surface area) to the total surface area of installed veto PMTs (324 cm² each). For the case with one veto PMT and all other ports covered with flanges, this number is adjusted accordingly. L —the average distance of travel between reflections—is assumed to be approximately 100 cm, which is on the same order of magnitude as the dimensions of the vessel. The first term in Eq. 6.2 accounts for fraction of photons detected directly after attenuation, and the second term fraction of photons surviving after reflections at Lumirror foil and attenuation in pseudocumene.

Table 6.1: The measured light yield of veto detector with air and one veto PMT is scaled to predict light yield with pseudocumene and 10 veto PMTs.

	One PMT, air	Ten PMTs, pseudocumene
\bar{n}	202	19
η_{direct}	0.065	0.037
$\eta_{\text{reflected}}$	0.0001	0.0062
η	0.065	0.043
Light yield (p.e./keV)	0.66	0.44

The estimated light yield with 10 veto PMTs and liquid scintillator is 0.44 p.e./keV. In reality, photons will likely travel longer than the diameter of the vessel, therefore the actual light yield is expected to be lower than 0.44 p.e./keV. This conclusion is in rough agreement with 0.21 p.e./keV predicted by optical Monte-Carlo [30].

6.5.3 Commissioning at LNGS Hall B

The veto vessel was shipped to LNGS in November, 2016, and in March, 2017, ten veto PMTs, high voltage power supply, trigger board and data acquisition electronics were delivered and installed at LNGS Hall B and an integral test of the veto detector was performed. All 10 veto PMTs were mounted onto the vessel with new divider boards using polypropylene capacitors. In this phase pseudocumene was not used due to regulations in Hall B.

The setup is shown in Fig. 6.8. Fig. 6.9 shows an event recorded with a $1'' \times 1'' \times 6''$ bismuth germanate (BGO) crystal scintillator suspended in the center of the vessel on the first day of commissioning at Hall B. Data acquisition was triggered by the veto majority coincidence with a coincidence level of 5, which was a preliminary indication that the veto PMTs, trigger firmware and DAQ system are functioning as expected. A dedicated test of trigger firmware is discussed in Chapter 7 and DAQ electronics and software are described in detail in Chapter 8.



Figure 6.8: Veto detector commissioning at LNGS Hall B. The veto vessel is on the right and the author and the electronics are on the left.

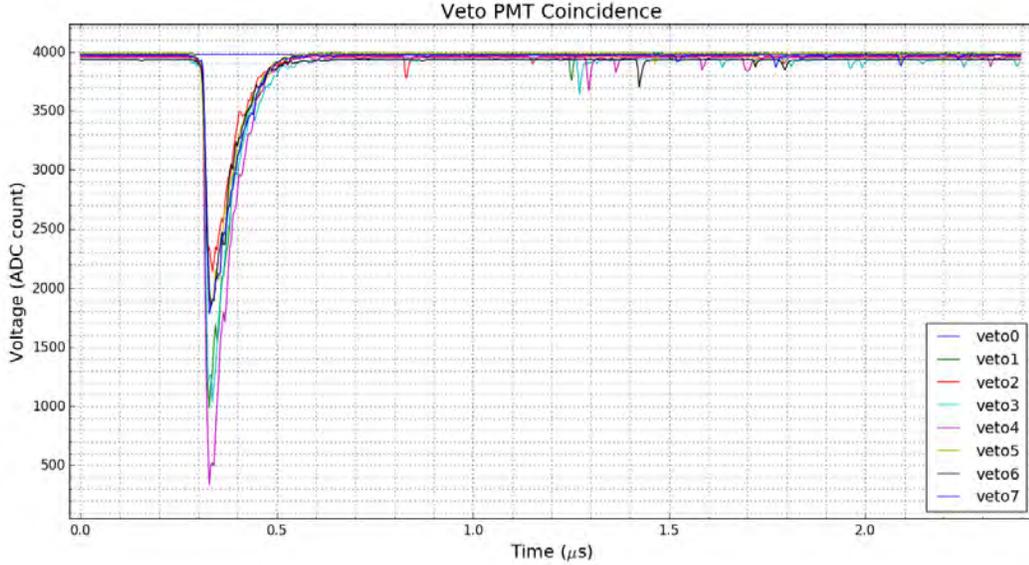


Figure 6.9: Majority coincidence of veto PMTs. Data taken on the first day of commissioning at LNGS Hall B. The majority level was set at 5. The event corresponds to a cosmic muon hitting a BGO crystal scintillator suspended in the center of the veto vessel. Later it was realized that the contact pin for channel 0 on the digitizer had been broken due to an earlier incident, and as a result there is no response in channel 0.

6.5.4 Commissioning at LNGS Hall C

After the integral test in Hall B, the veto vessel was relocated to the final site at Hall C next to the Borexino scintillator storage tank for convenient access to pseudocumene. In June, 2018, PMTs were dismantled from the vessel and the vessel was precision-cleaned according to the Borexino cleaning procedure. The resistance of the deionized water used to rinse the vessel was used as a control parameter, and after cleaning, the water had a resistivity of 15.18 M Ω -cm.

The PMT cans, PMT holders and veto PMTs were separately cleaned inside a cleanroom using acetone, deionized water and isopropyl alcohol. The previously-used Lumirror foils were dismantled prior to the vessel cleaning and used as a pattern to cut new Lumirror lining. The newly-cut Lumirror foils were cleaned in a cleanroom. Along with the veto PMTs, the Lumirror foils were mounted onto the veto vessel in

a temporary HEPA-filtered cleanroom. Fig. 6.10 shows the shielding around the veto vessel and the temporary cleanroom built next to the vessel.

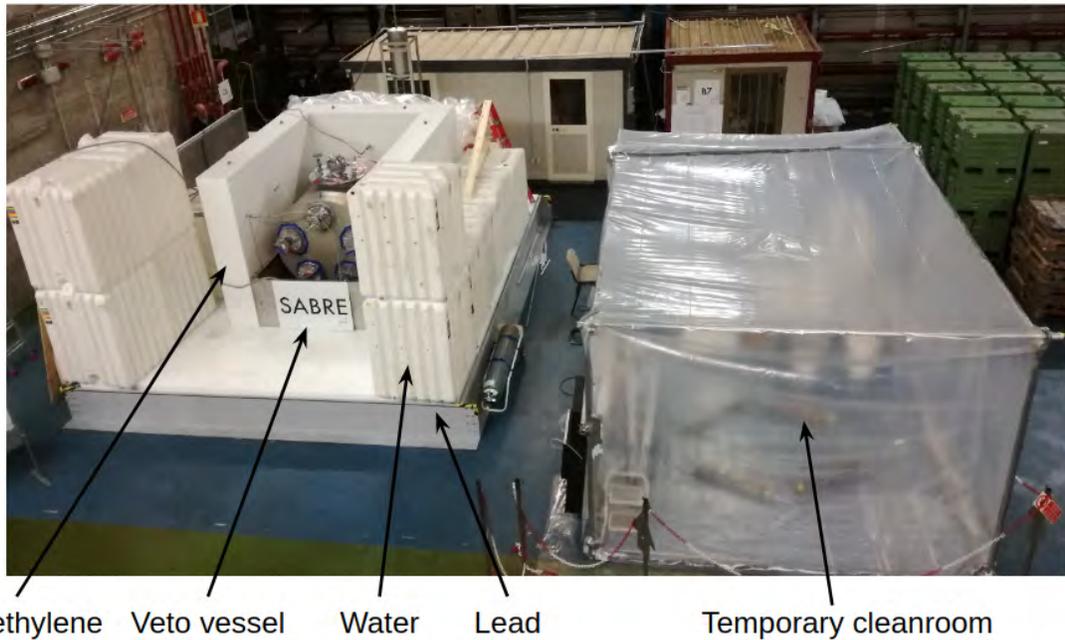


Figure 6.10: Veto vessel commissioning in LNGS Hall C. The veto vessel is shielded by 10 cm of polyethylene and 80 cm of water on the top, 40 cm of polyethylene and 91 cm of water on four sides, and 15 cm of lead on the bottom. A temporary cleanroom with HEPA-filtered air was built near the final site for cleaning the vessel and installing the veto PMTs and the Lumirror lining. Photo courtesy of Chiara Vignoli.

Chapter 7

Trigger

In nuclear and particle physics experiments, trigger is a system that monitors the status of every signal channel for potential event of interest and decides when to start recording data. Trigger is usually implemented as various logical combinations of digital signals that reflect the status of the ADC channels. Considering the performance, scalability, and flexibility requirements of modern experiments, a natural choice of hardware for implementing the trigger is a field-programmable gate array (FPGA)—a chip on which the user can electrically program different digital circuits.¹

The trigger firmware I developed for operating the NaI (Tl) detector together with the veto detector underground is described in Section 7.1. The measurement of trigger efficiency and verification of the trigger firmware is discussed in Section 7.2.

7.1 Design and Implementation of Trigger Firmware

To support various modes of detector operation, the trigger firmware is designed to offer four different trigger channels: *CRYSTAL*, *COINCIDENCE*, *VETO* and *PC*.

- The *CRYSTAL* channel is triggered by the coincidence between two PMTs coupled to the same crystal, regardless of the status of the veto detector. This

¹The circuit is referred to as firmware.

trigger channel is useful for both dark matter search and ^{40}K counting: in both cases there is a signal in the crystal leading to the coincidence of the two crystal PMTs, and the only difference is the presence or absence of a signal in the veto detector, which can be identified offline by software.

- The *COINCIDENCE* channel looks for coincident signals between the crystal detector module and the veto detector. This channel is useful in detector tests and calibrations using sources with coincident γ -rays such as ^{22}Na .²
- The *VETO* channel triggers upon the coincidence signals in two PMTs coupled to the same crystal and the absence of majority coincidence in the liquid scintillator veto. This channel was designed for preliminary dark matter searches in high-event-rate environments.³
- *PC* channel looks for a signal in the pseudocumene liquid scintillator, regardless of the status of the crystal detector module. This channel was designed such that the veto detector can be operated as a stand-alone detector to measure the purity of pseudocumene.

The trigger firmware is written in VHDL and is programmed on an Altera Cyclone FPGA in a CAEN V1495 Programmable Logic VME board. The overall setup is shown in Fig. 7.1. The low-voltage differential signaling (LVDS) ports of the ADC board is connected to the digital input pins of the trigger board. Once the waveform in a channel crosses the threshold, a signal is sent from the LVDS output to the corresponding input pin of the FPGA chip. Upon satisfying the trigger conditions, a TTL signal is sent from the trigger output of the FPGA to the external trigger input port of the ADC board to start data recording.

²Although this channel can also be used to count ^{40}K , this channel would miss potential dark matter events and the *CRYSTAL* channel should be used instead in physics runs.

³When event rate is low, *CRYSTAL* channel could be used instead and events with majority coincidence in the veto detector can be filtered offline.

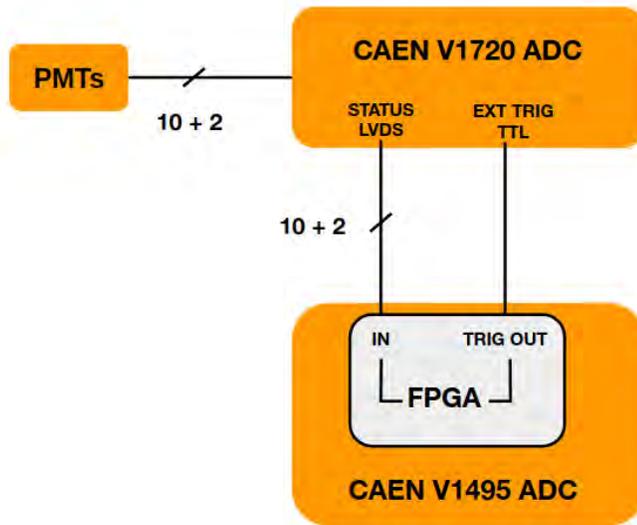


Figure 7.1: Setup of CAEN V149 trigger board. The cross-threshold information is sent from the LVDS pins of the ADC boards to the FPGA in the trigger board. When trigger condition is met, an output signal is sent from the FPGA to the external trigger inputs of the ADC boards.

Inside the FPGA chip, the trigger algorithm is divided into three functional blocks: *Window Generator*, *Coincidence* and *Output*. Identical *Window Generator* blocks detect incoming TTL signals from the LVDS outputs of the ADC channels, synchronize the pulses with the FPGA’s clock through two stages of D flip-flops, and output pulses of programmable length to the *Coincidence* block.

The *Coincidence* block runs different algorithms for the crystal and the liquid scintillator veto: for the crystal, output signals from two independent *Window Generator* blocks are passed to an AND-logic gate and the output is taken as the crystal coincidence signal. This is schematically illustrated in Fig. 7.2. For the liquid scintillator veto, the output of ten *Window Generator* blocks are passed to a programmable majority comparator. The majority comparator will count the number of active coincidence windows and compare this number to a programmable majority level. If the number of active coincidence windows is no less than the specified majority level,

the majority comparator outputs a liquid scintillator majority trigger signal. The schematic for the veto majority trigger is shown in Fig. 7.3.

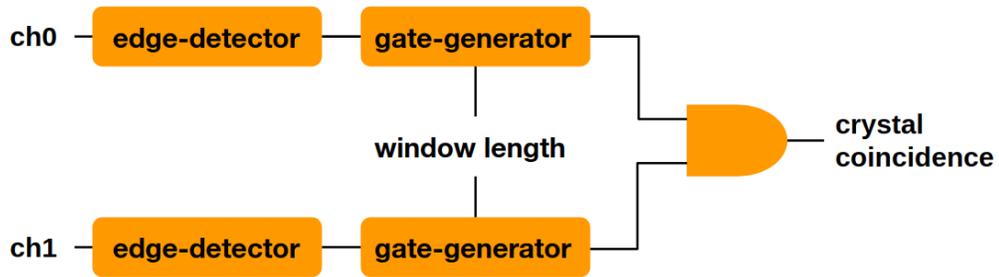


Figure 7.2: Implementation of crystal coincidence. Input LVDS signals from crystal PMT channels are used to generate a coincidence window of programmable length. When there is an overlap between the two coincidence windows, a crystal coincidence signal is generated.

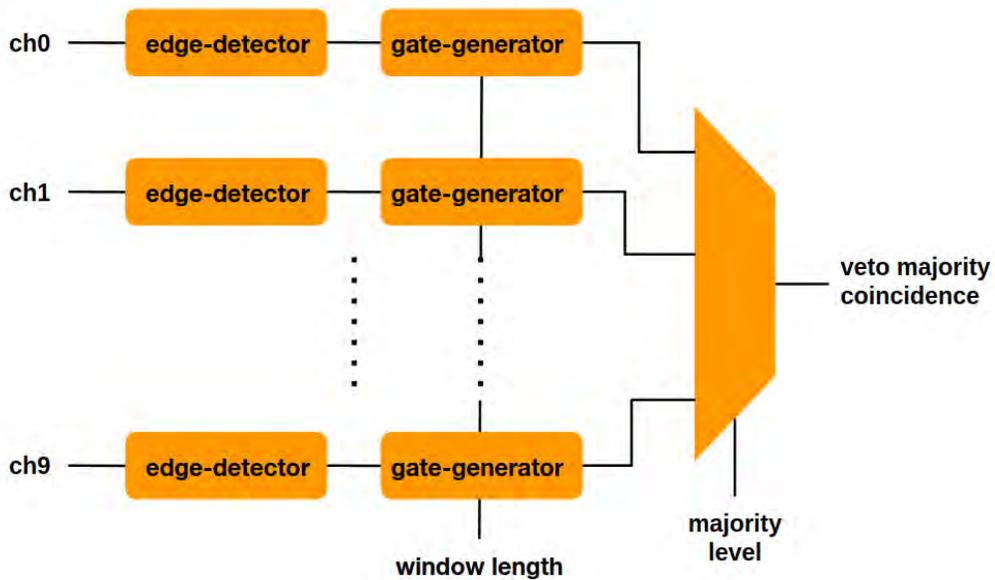


Figure 7.3: Implementation of veto coincidence. The LVDS signals from the veto channels are processed by *Window Generator* blocks identical to those used for crystal channels. Instead of directly using an AND-gate, the outputs are passed onto a majority comparator which counts the number of active windows for the input signals, and if this number is greater than a pre-programmed majority level, a veto majority coincidence signal is generated.

As illustrated in Fig. 7.4, the *Output* block consists of a multiplexer to select different trigger channels and an output module to generate a trigger signal of programmable length. In the *CRYSTAL* or *PC* trigger channels, the crystal coincidence or veto majority coincidence signal is directly passed to the trigger output module. In the *COINCIDENCE* channel, crystal coincidence and veto majority coincidence signals are passed into an AND-gate. In the *VETO* channel, the crystal coincidence signal is delayed by a duration of programmable length, and if no veto majority coincidence occurs in this time interval, the signal is passed to the trigger output module.

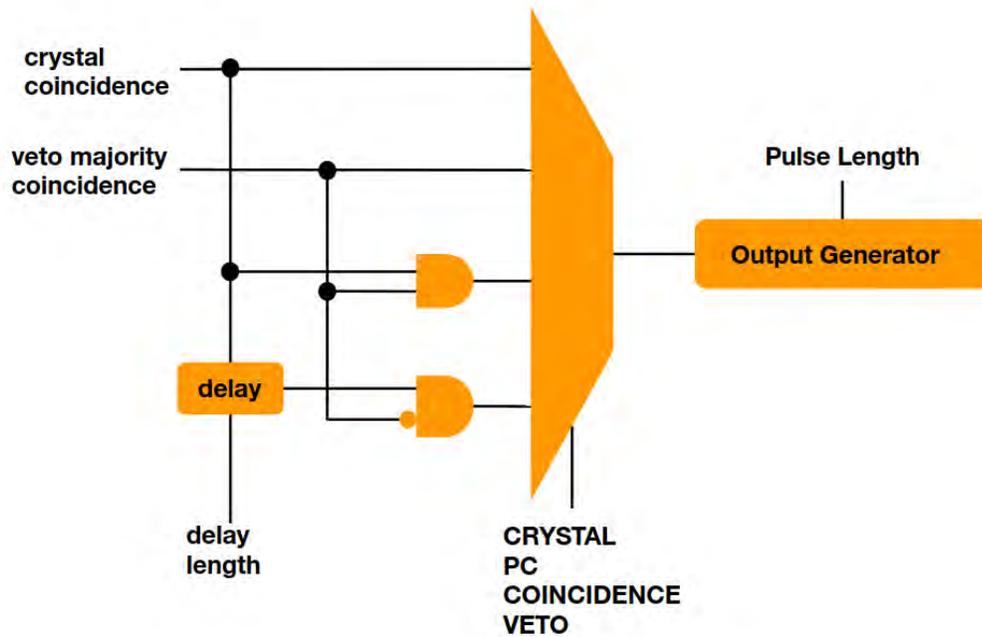


Figure 7.4: Trigger output module. Four different trigger channels are selected by a multiplexer and the output of the multiplexer is sent to the trigger output generator with programmable pulse length. In the *VETO* channel the crystal coincidence signal is delayed by a duration of programmable length such that no false positive is generated when a crystal coincidence signal arrives before a veto majority signal.

Table 7.1 summarizes available registers, their addresses and their functions. These registers are accessed through CAEN's protocol provided with the VME board.

Table 7.1: Description of registers implemented in trigger firmware. All time length refers to counts of FPGA clock cycles (25 ns per cycle).

VME Address	Bits	Name	Description
0x1030	0	Enable	1: enable trigger 0: disable trigger
	1	$\overline{\text{Reset}}$	1: no effect 0: reset trigger circuit
	3:2	Trigger channel	00: CRYSTAL 01: PC 10: VETO 11: COINCIDENCE
	4	Crystal retrigger	1: enable crystal retrigger 0: disable crystal retrigger
	5	Veto retrigger	1: enable veto retrigger 0: disable veto retrigger
	15:6	Veto mask	
	19:16	Veto majority level	
	22:20	Delay in crystal coincidence under veto mode	
	30:23	Trigger output length	
	31	Reserved	
0x1034	15:0	Crystal coincidence window length	
	31:16	PC coincidence window length	
0x1038	15:0	Dead time imposed after trigger	
	31:16	Veto window length	
0x1042	31:0	Software trigger	

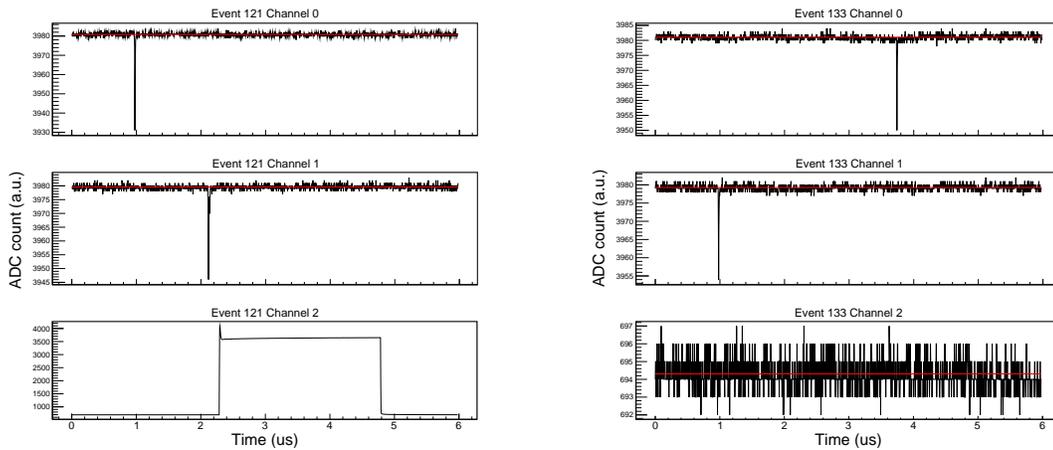
7.2 Measurement of Trigger Efficiency

Due to the complexity of the trigger system, it is challenging to measure trigger efficiency for every combination of trigger channels and parameters. On the other hand, the trigger firmware is mostly digital logic that can be simulated realistically and reliably using FPGA simulation environments provided by the vendors. In the trigger design, identical *Window Generator* circuit blocks are used for both crystal coincidence and veto majority coincidence. The remaining unknown factor is the two-channel coincidence trigger efficiency for different relative ADC cross-threshold times. The trigger efficiency of the *CRYSTAL* channel is directly influenced by this

two-channel coincidence efficiency, and the multi-channel majority coincidence can be easily estimated from the two-channel coincidence efficiency.

Dark counts from two Hamamatsu R5912 PMTs are used to randomly sample different relative cross-threshold times. Signals from the two PMTs are sent to a CAEN V1720 ADC board, and their corresponding LVDS outputs are connected to the coincidence inputs of a CAEN V1495 trigger board. The trigger output is sent to the third ADC channel to record the output trigger signal. The data acquisition is started by the ADC board's local cross-threshold trigger for either of the two PMTs.

The trigger efficiency is computed as the number of events with trigger output signals divided by the total number of events for each relative cross-threshold time. Fig. 7.5a shows a successful trigger and Fig. 7.5b shows an unsuccessful trigger due to two pulses being too far apart.



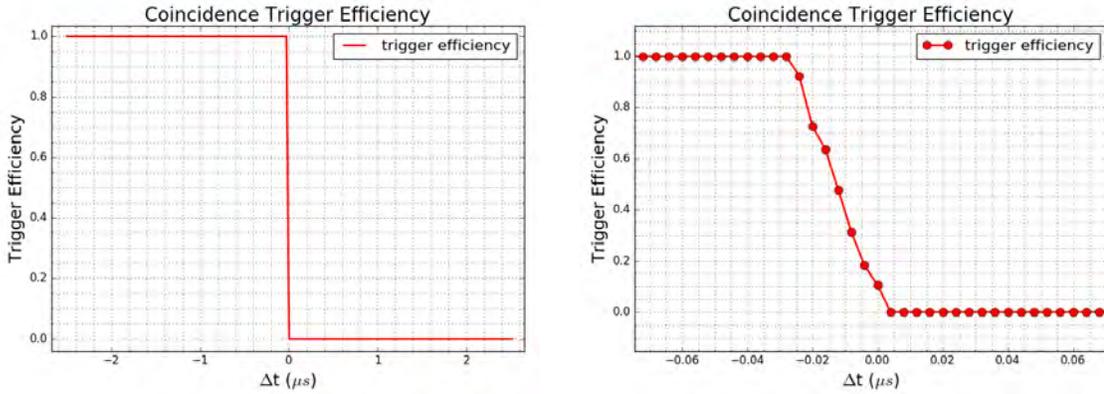
(a) The trigger is generated slightly after the second waveform crosses the threshold.

(b) The cross-threshold times of the two waveforms are separated by more than the coincidence window.

Figure 7.5: The top and middle channels are PMT waveforms and the bottom channel is digitized trigger signal. Data acquisition is started by either waveform crossing the threshold in the first two PMT channels. The coincidence window is set to be $2.5 \mu\text{s}$.

The measured coincidence trigger efficiency is shown in Fig. 7.6. The trigger efficiency makes a sharp transition from 100% to 0% around $0 \mu\text{s}$ overlap with an uncertainty of 25 ns , the trigger board's clock period. This uncertainty arises from

pulse synchronizer in the firmware: if an external input signal arrives during a clock transition, the output signal is updated in either the current clock cycle or the next.



(a) Unzoomed two-channel coincidence trigger efficiency.

(b) The region around $0 \mu s$ is enlarged for a detailed view.

Figure 7.6: Two-channel coincidence trigger efficiency as a function of the distance between two coincidence windows. Negative number means there is an overlap between two coincidence windows.

Chapter 8

Data Acquisition

Typical DAQ software comes in two categories: ready-to-run programs and application-specific programs. Ready-to-run programs are easy to use, but are largely limited to simple predefined tasks and are difficult to interface with other programs in the same application. On the other hand, application-specific programs can achieve better performances but the development process is often expensive and time-consuming.

To solve the challenge, I developed an open-source, general-purpose modular framework for data acquisition called Polaris. Polaris is written in C++ and can meet a wide range of DAQ requirements from routine laboratory measurements to mid-scale nuclear and particle physics experiments.¹

Section 8.1 describes the design philosophy and software architecture of Polaris. To support various data acquisition and control tasks, two custom electronic circuit boards were developed. These boards were used to monitor precision-drying and pyrolytic carbon coating and to operate PyroMapper. The hardware and corresponding Polaris software libraries are described in Section 8.2. Section 8.3 discusses the Polaris library used to operate the SABRE detector. Finally, the baseline subtraction and reconstruction algorithms used in this thesis is presented in Section 8.4.

¹Polaris can be downloaded at <https://github.com/suerfu/polaris>.

8.1 Polaris

The challenge of developing a general-purpose DAQ framework lies in the balance between flexibility and the amount of customization required. Almost all tasks differ to some extent in the hardware used, the types of data to be acquired, and how data is logged and displayed. On the other hand, different DAQ tasks have common behaviors: all DAQ applications need to go through an initialization phase, a hardware and software configuration phase and a run phase. Polaris exploits the common features of DAQ and leaves application-specific features to user implementation.

8.1.1 Design Philosophy

In the Polaris framework, common DAQ features are managed by Polaris while application-specific features are implemented as runtime dynamic libraries called modules. There are three advantages to a modular design:

1. by specifying and loading different sets of modules, the same program can be used on different measurements without having to recompile the main program;
2. features are grouped into different modules, which simplifies debugging and maintenance; and
3. new features can be easily added—the user only has to develop a corresponding new module and load it at runtime.

8.1.2 Software Architecture

The software architecture adopts a finite-state machine-based modular design. State transition is managed by a controller object called *plrsController* and every user module must derive from the base class *plrsStateMachine*. *plrsStateMachine* defines virtual functions to be called at each state and at each state transition, guaranteeing

that every module adheres to the state transition rules and the right functions are invoked at the right moment.

Polaris DAQ States

There are five main states: *INIT*, *CONFIG*, *RUN*, *PAUSE* and *END*. The main features of each state is summarized below:

- ***INIT***: After each module is loaded, data and command first-in-first-out (FIFO) buffers are allocated by *plrsController* to user modules. User modules acquire parameters from a configuration file and initialize their own variables.
- ***CONFIG***: User modules perform hardware and software configurations. *plrsController* compiles a look-up table of registered modules. User modules configure the data flow path by looking up the indexes of other modules' data buffers.
- ***RUN***: Main event loop is executed.
- ***PAUSE***: The *PAUSE* state is only accessible from the *RUN* state. *Pause()* and *Resume()* functions are invoked when entering and exiting the *PAUSE* state.
- ***END***: Modules deconfigure their hardware and release their resources. Threads are terminated and dynamically-loaded libraries are closed. The main program is ready to exit at this stage.
- ***ERROR***: Transition to the *ERROR* state can be made from any state. Functions on the path from the current state to the *END* state will be called successively to make sure hardware is properly deconfigured before exiting.

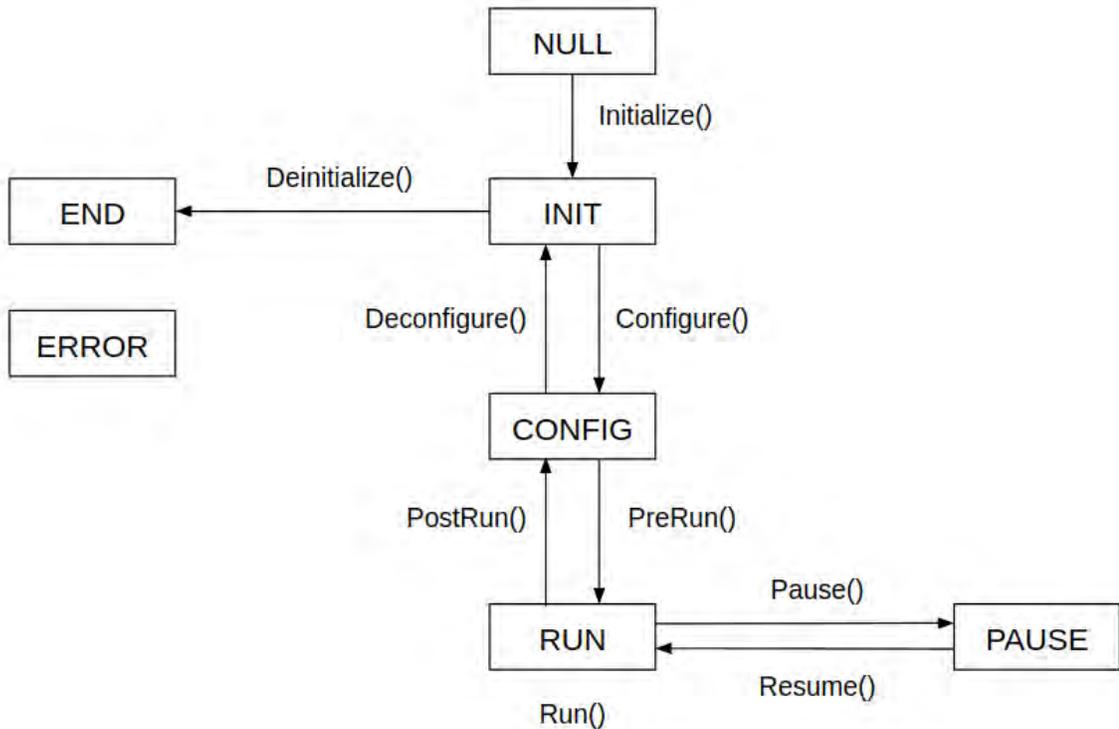


Figure 8.1: Polaris states, state transitions and the corresponding virtual functions. The state transitions are managed by the *plrsController* and virtual functions are defined by the base class *plrsStateMachine* from which every user module should be derived. The function next to the arrow are called at each state transition.

State Transition

Internally, two variables (*state* and *status*) are used to manage the state transition of every module. *plrsController* changes the value of *state* as a signal to make a state transition. *plrsStateMachine* polls the value of *state* and compares it to its local variable *status*. When *state* differs from *status*, *plrsStateMachine* makes a state transition by invoking appropriate functions and updates *status* upon success. *plrsController* polls regularly to check if every module is in the same status. The next state transition will not occur until all modules are in the same state and status. This mechanism is implemented on the base class *plrsStateMachine* and derived classes (user modules) do not have to worry about it.

Once state transition signal is issued, *plrsController* also starts a watchdog timer that can be extended by user modules as needed. If a module fails to update the timer for any reason, *plrsController* will set the state to *ERROR* and ask every module to terminate.

Data Flow

When each module is loaded, *plrsController* assigns two FIFO buffers, one for data and the other for commands. The FIFO buffers of all modules are kept in an indexed array by *plrsController*. Upon configuration, *plrsController* compiles a look-up table with which user modules can look up the indexes of other registered modules. Once the index is known, every module can push data and commands to any other module as needed. However, each module can only pull from its own buffers. With such architecture, it is possible to realize versatile and dynamic flow patterns. This is useful, for example, when data flow has to bypass certain modules.

The data buffer is implemented as a custom, thread-safe circular buffer with a static address instead of a standard template queue which dynamically re-allocates new memories. For demanding applications, the address instead of the data itself should be stored in these buffers to avoid wasting time copying large amounts of data.

Interface

plrsModuleInput manages command-line input in a non-blocking, thread-safe way. This is the only module loaded by default. This behavior can be disabled by command-line options upon starting.

In addition to command-line interface, inter-process interface and inter-host interface are implemented by *plrsModuleInterface*. *plrsModuleInterface* uses a Unix socket to establish the connection between Polaris and other programs running on the same

computer (host), and a TCP/IP socket for connection with programs running on remote hosts.

With an inter-process interface, it is very easy to visualize event data and develop graphical interfaces. Inter-host communication allows for remote monitoring of the experiment data, and the development of a multi-host distributed DAQs system. As an example, in Section 8.2, *plrsModuleInterface* is used to send data to a Python script running on a remote host for monitoring the drying/coating process.

8.1.3 Performance

The acquisition bandwidth of Polaris depends on the host running the application, the number of modules loaded, the depth of the FIFO buffer and the complexity of operations on the data. In the simplest case in which integer dummy data is read and stored on disk, an acquisition rate of a few GB/s is easily achieved on the author's Dell Latitude E6530 laptop computer. For comparison, a CAEN A3818 PCIe card has a data transfer bandwidth of about 80 Mb/s. Therefore, in most small- to mid-scale experiments, the limiting factor is likely the data transfer rate from DAQ hardware to PCs.

8.2 DAQ for Powder Drying and Crucible Coating

The DAQ software for powder drying and crucible coating is developed as Polaris modules. Together with Polaris modules for operating the SABRE crystal modules and veto detectors, these modules are compiled into one single shared library *lib-saber.so*, available at <https://github.com/suerfu/saber>.²

²Although this DAQ software system can be used on the final SABRE experiment, currently it is used mainly for tests at Princeton. Migration from previous old DAQ software (also developed by Suerfu) to a Polaris-based DAQ system is in progress.

8.2.1 PIC16-based DAQ Circuit

Custom-designed PIC16-based DAQ boards are used to control various devices and digitize and record signals from different sensors. The DAQ boards, shown in Fig. 8.2, use PIC16F1788 microcontrollers for main processing. Peripherals include 8 ADC ports, 8 general-purpose I/O ports and 6 control ports that can be configured as SPI ports or PWM output ports. Communication with a computer is done via UART.³



(a) Photo of Mica v0. It uses RS-232 cables to communicate with a PC's serial port.



(b) Photo of Lyra v0. Communication to the PC is done using a USB cable and an on-board USB-to-serial converter chip.

Figure 8.2: Two custom circuit boards for laboratory data acquisition. Mica is powered from the wall with AC adapters and has on-board voltage regulators, relays and power MOSFETs to power up and control other devices directly. Lyra is designed to be powered from a PC's USB port and its function is limited to data acquisition and digital I/O with other independently powered devices.

Although these DAQ boards are designed as general-purpose interfaces between computer software and physical signals, dedicated special firmware can be uploaded for complex tasks (e.g. controlling a stepper motor).

8.2.2 Monitoring NaI Powder Drying

For monitoring NaI powder drying, Lyra v0 (Fig. 8.2b) is used to read the temperatures of the oven and the cold trap, and the pressure of the cold trap. The setup is shown in Fig. 8.3. Temperature is measured using thermocouples and the volt-

³Upgrade to Universal Serial Bus (USB) is in progress.

age signals are amplified before digitization. The signal from the pressure gauge is attenuated by a factor of 3 to fit in the dynamic range of the microcontroller.

Two essential Polaris modules in recording drying data are a *SerialDAQ* module and an *AsciiRecorder* module. During run, the *SerialDAQ* module periodically polls to read digitized input signals and sends available data to the *AsciiRecorder* module, which writes data onto the output file in ASCII format. Optionally the data can be visualized by a graphics module or sent to another host by a network interface module. Example data is shown in Fig. 3.3.

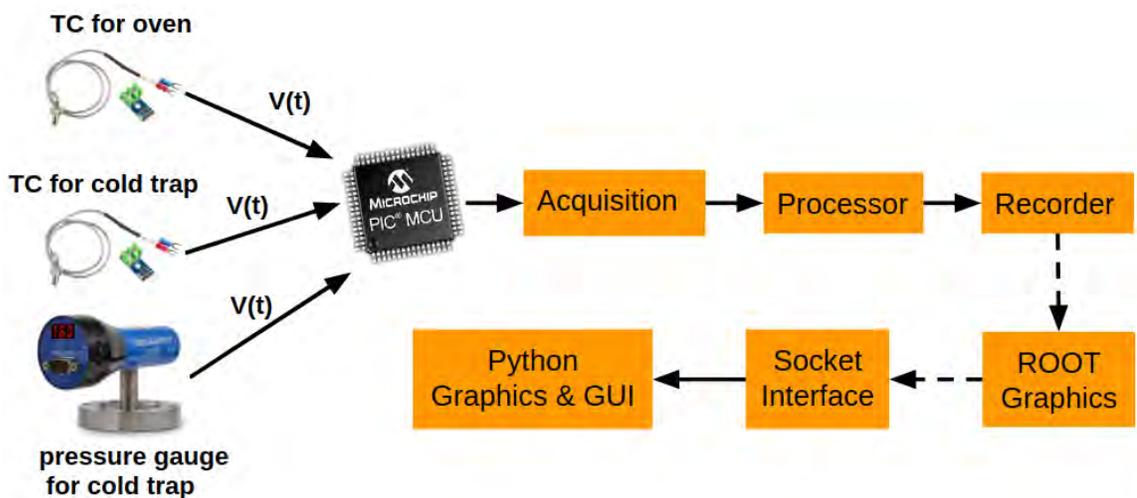


Figure 8.3: Hardware and software setup for recording NaI powder drying data. Temperature signals from thermocouples are amplified by IC chips while voltage from pressure gauge is attenuated by a factor of 3. *SerialDAQ* and *AsciiRecorder* are two essential polaris modules for recording data.

8.2.3 Monitoring Pyrolytic Carbon Coating

Pressure and temperature are the two most important parameters in a pyrolytic carbon coating process. The setup of the monitoring system is shown in Fig. 8.4. Data is acquired by a *SerialDAQ* module and then sent to a processor module. The processor module averages and converts the data to physical units from ADC counts based on the calibration characteristics of each sensor. After written to disk, data is sent to a

ROOT graphics module for local display or over the internet to a Python visualization script running on remote hosts. Fig. 8.5a shows a screenshot of local ROOT graphics during the execution of the program and Fig. 8.5b shows a screenshot of Python visualization script running on a remote host.

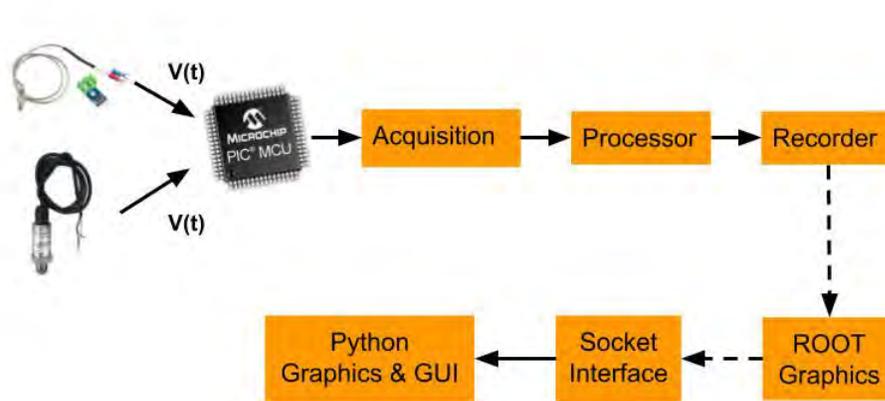
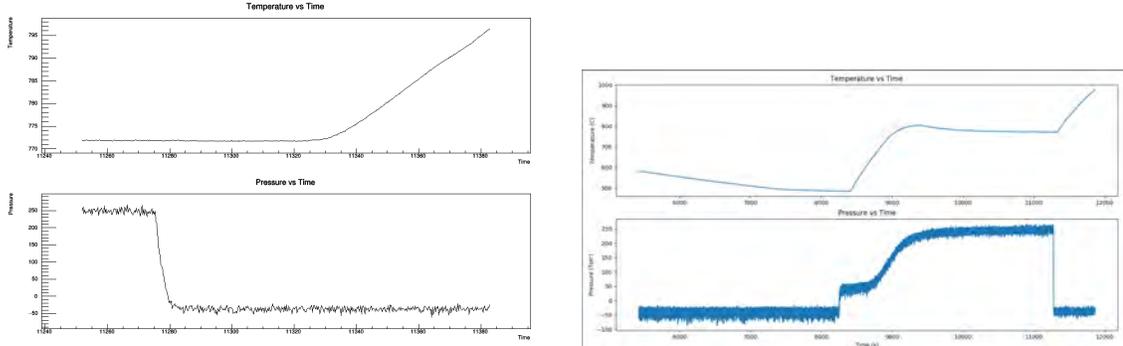


Figure 8.4: Setup for monitoring pyrolytic carbon coating. The pressure is obtained from a Dwyer 628 pressure transducer and the temperature is obtained by amplifying the signal from a K-type thermocouple. This application uses a processor module to convert data from ADC counts to physical units before storing, and an Internet Interface Module to send data to a remote Python visualization script.

8.2.4 Mapping Pyrolytic Carbon Coating

The thickness of pyrolytic carbon coating can be mapped by measuring laser attenuation and the apparatus is described in Section 2.4.2. Custom circuit board Mica is used to measure the intensity of the transmitted laser and to power up and control the laser and the stepper motors. Since stepper motors with specific step angles and shaft diameters are used in this application, a dedicated firmware is implemented to drive



(a) Screenshot of local ROOT graphics during the execution of the program. The graphics displays short-term temperature and pressure data.

(b) Screenshot of Python visualization script on remote host during the execution of the program. Graphics shows long-term data since connection is established.

Figure 8.5: Event visualization using local graphics library and remote Python script via internet interface module.

and keep track of rotational positions of each stepper motor at the microcontroller level.

For data taking, a dedicated *PyroMapDAQ* module is used to interface with the special firmware. The module turns the laser diode on and off at the beginning and end of each run, and issues commands to move the laser-sensor assembly to the correct position using the location feedback from Mica. After readout, an *AsciiRecorder* module is used to record the data and a *RootGraphics* module is used to visualize the data.

8.3 SABRE DAQ System

The trigger and DAQ system for the SABRE detector consists of a trigger board with custom FPGA firmware and two ADC boards to digitize and record signals from 10 veto PMTs and 2 PMTs coupled to the crystal. All boards are mounted on a VME crate and communication with the PC is managed by the VME bridge board through an optical fiber. The setup is illustrated in Fig. 8.6.

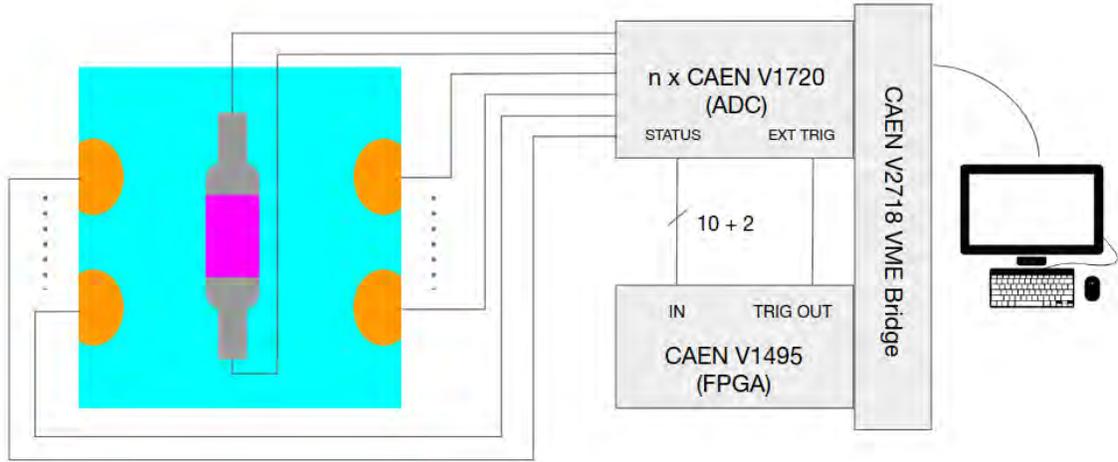


Figure 8.6: The trigger and DAQ hardware setup for operating the crystal and veto detectors is illustrated here.

The clock signal from one ADC board is sent to the PLL of the second ADC board such that the two boards have synchronized clocks and the same time reference for events. During data acquisition, the ADC boards continuously write digitized samples to internal circular buffers. Once the input signal crosses the ADC threshold, a LVDS signal for the corresponding channel is sent to the input on the trigger board. Upon satisfying the trigger condition, a TTL trigger signal is sent from the trigger board to the external trigger inputs of the two ADC boards. The ADC boards will further write a certain number of samples to the buffer before freezing it for readout. The size of each circular buffer and number of samples to acquire after the trigger are configurable. Such a scheme allows acquiring a complete waveform including pre-trigger samples and post-trigger samples.

The Polaris modules managing the trigger and ADC boards are compiled in *libsaber.so*. *SaberDAQ*, *SaberDiskRecorder* and *SaberGraphics* are three key Polaris modules (see Fig. 8.7).

All trigger and DAQ settings are specified in a configuration file which is loaded at the beginning of the execution of Polaris. Based on the information in this configuration file, *SaberDAQ* initializes connection to the VME crate using an appropriate

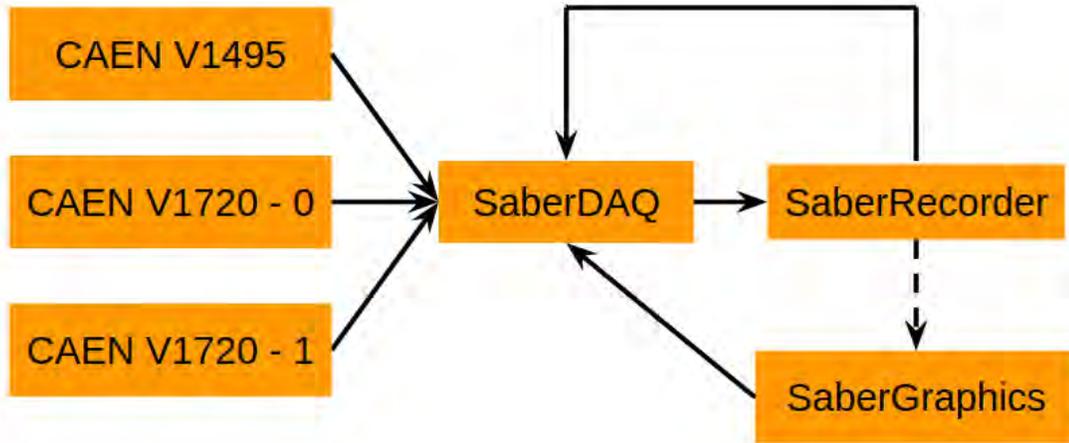


Figure 8.7: Data flow between three key modules in *libsaber.so*. *SaberDAQ* is responsible for low-level communications with DAQ hardware. The address of acquired data is sent to *SaberRecorder*, which depending on the trigger rate, passes the address to either *SaberGraphics* for event display, or back to *SaberDAQ* to record new events.

communication method, and creates and initializes *CAENV1495* and *CAENV1720*, two C++ objects that directly manage the corresponding boards.

At the beginning and end of each run, *SaberDAQ* compiles parameters and other useful information of the run into a header object to be included in the output file. During run, *SaberDAQ* module continuously polls ADC boards for new data until specified run time or event number is reached.

After event data is read, it is assembled into a *SaberRawData* object and sent to a *SaberDiskRecorder* module for output to a file. To save disk space, event data is stored in binary format. Subsequently data can be optionally sent to a *SaberGraphics* module for real-time event display using ROOT graphics library.

8.4 Event Reconstruction

After an event is recorded, it must be processed or reconstructed to extract information about the total energy deposited in the detector medium and possibly the type

of incoming radiation. Since PMT pulses have negative polarity, to enhance dynamic range, the input PMT signal is offset by a digital-to-analog converter on the digitizer board such that the baseline is around 4000 ADC counts, slightly below 4095, the maximum output of a 12-bit ADC. To integrate the pulse to get the total charge, the baseline must be subtracted from the waveform. After baseline subtraction, the pulse-finding algorithm identifies the beginning and end of pulses in the interval of interest and computes pulse heights and integrals.

8.4.1 Baseline Subtraction

The baseline is often not constant due to drifts in electronics or long acquisition windows. However, for NaI (Tl) crystals, a constant baseline seems to be sufficient: the acquisition window is relatively short, and negative high-voltage divider boards are used for crystal PMTs to reduce baseline fluctuations by direct coupling to the ground.

To find the baseline, the entire waveform is scanned with a window of programmable-width to identify regions with variance below a pre-defined threshold. This criteria will select regions with clean baselines and exclude pulse regions where the waveform varies rapidly. All samples in qualified regions are then averaged to give the best-estimate of the baseline. An example baseline is shown in Fig. 8.8.

8.4.2 Pulse Finding Algorithm

After baseline subtraction, the pulse-finding algorithm searches the waveform for the main scintillation pulse and trailing single photoelectron pulses. The algorithm exploits the fact that the negative peak of the pulse is also the minimum of the waveform. The algorithm is described below and illustrated in Fig. 8.9.

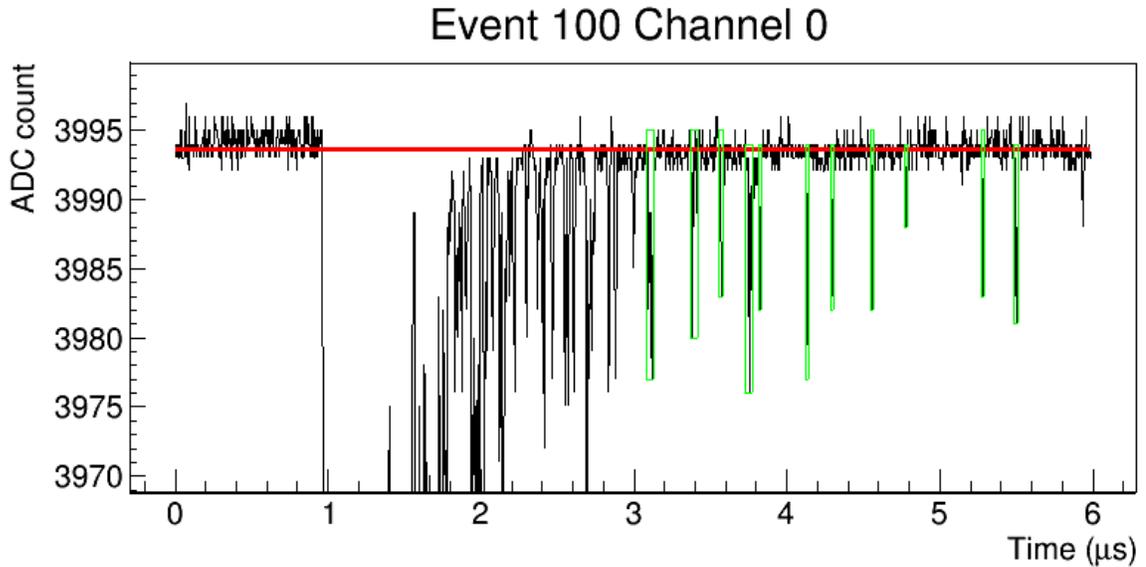
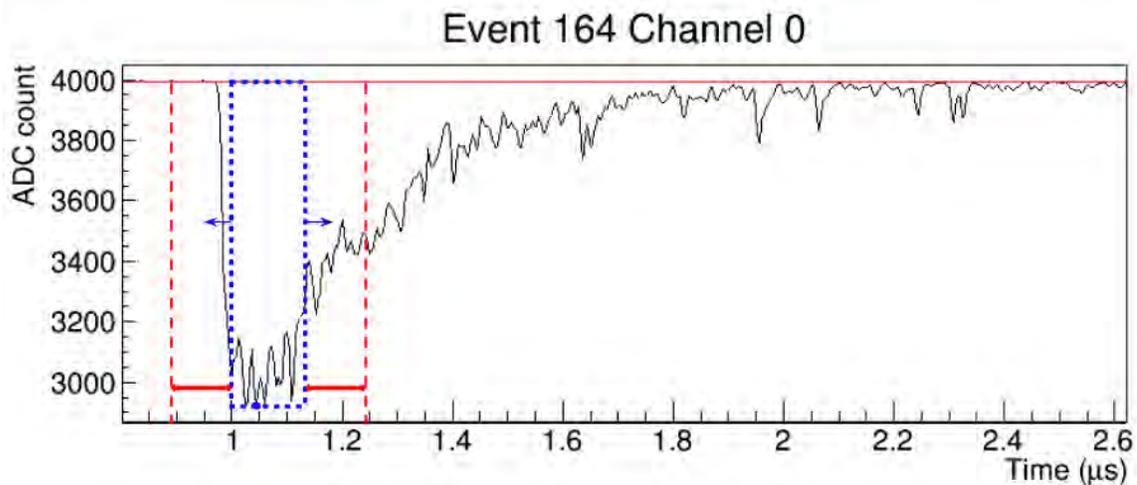
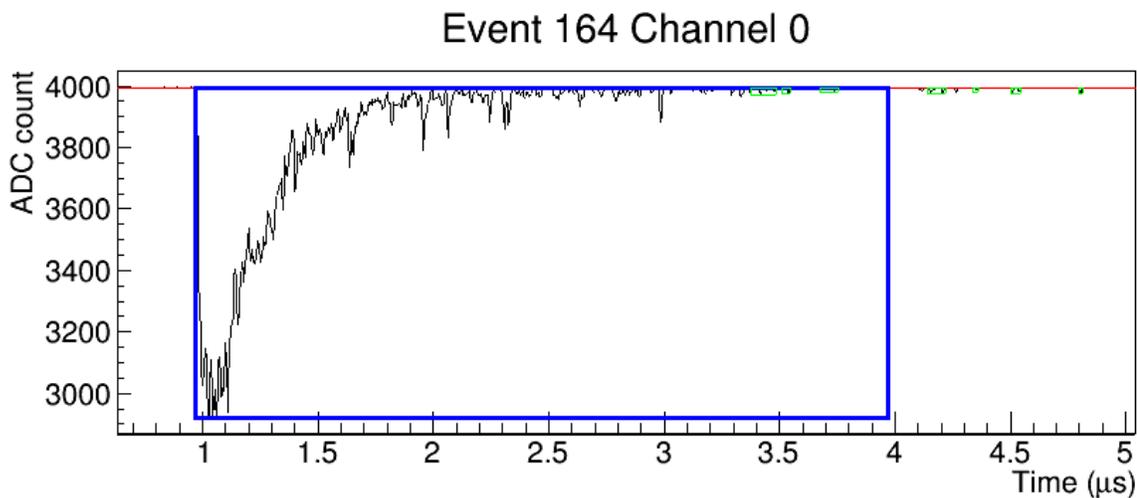


Figure 8.8: An example baseline (in red) of PMT signal of NaI crystal scintillation. The baseline is found by selecting regions with small variance and averaging samples in these regions.

1. The search region is defined by the two variables *search_begin* and *search_end*, and the minimum within this range is identified.
2. A sliding window of length *pre_window* is moved backwards in time to locate the beginning of the pulse. When the variance in this sliding window is smaller than *pre_win_var*, the beginning of the pulse is identified.
3. In the fixed-window mode, pulse region is identified as samples within a window of length *window* from the beginning of the pulse. In automatic-window mode, a similar sliding window is used to locate the end of the pulse at least *window* samples away from the beginning of the pulse.
4. Identified pulse regions are masked from future searches and the algorithm is repeated to identify the next pulse until the pulse height of a newly-found pulse is below the preset *threshold*. The algorithm then terminates.



(a) The pulse-finding algorithm works by defining a peak searching window (dashed blue box). Variance of waveform within a sliding search window (dashed red) is calculated to determine the beginning of a pulse (and the end of the pulse in automatic mode).



(b) The search window keeps sliding until the variance is below the pre-defined threshold. The identified pulse is masked and excluded from future searches. These steps are repeated until all pulses have been identified.

Figure 8.9: Illustration of the pulse-finding algorithm.

Chapter 9

Background Simulation and Sensitivity

Geant4-based Monte-Carlo simulations are used to estimate backgrounds from various sources. To accurately model low-energy atomic de-excitation, a bug in the Geant4 source code was fixed to correctly handle Auger electron emission (Section 9.1). Potential backgrounds are analyzed in Section 9.2.1 and Monte-Carlo simulations are used to verify the analysis (Section 9.3).

9.1 Avalanche of Auger Electron in Geant4

When an ionizing radiation particle frees atomic electrons in a medium, a vacancy is created and the atom is excited. Transition to the ground state is often accompanied by emission of electromagnetic radiation or Auger electrons—electrons from higher energy levels. In Geant4 versions earlier than 10.2, the two processes are managed by the same piece of code: given an inner shell vacancy, an experimentally-determined probability matrix is used to calculate which of the outer-shell electrons to fall into the vacancy and the energies of the outgoing X-rays and/or Auger electrons, and the same code is recursively applied to propagate the vacancy outwards. However,

this algorithm neglects a key fact: when an Auger electron is emitted, in addition to the initial vacancy, an additional vacancy corresponding to the Auger electron is created. Since this additional vacancy is absent if X-ray is emitted, Auger electron emission must be handled separately. Otherwise binding energies corresponding to new vacancies due to Auger electrons are lost and the total de-excitation energy will appear to be smaller. The difference between X-ray and Auger electron emission is illustrated in Fig. 9.1.

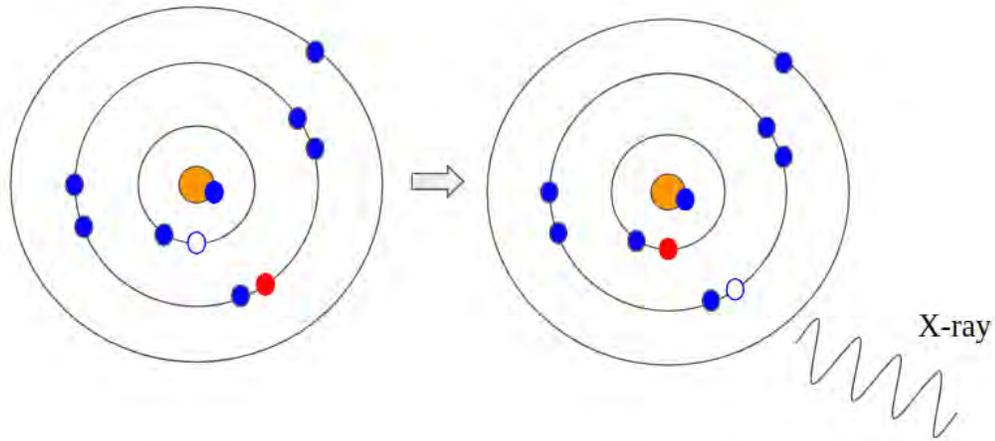
To solve this inconsistency, I proposed a bug fix that keeps track of additional Auger electron vacancies in a FIFO buffer: each time additional vacancies are created due to Auger electrons, the vacancies are pushed into the FIFO buffer. The de-excitation code is repeatedly executed until no more vacancies remain in the FIFO buffer.

This bug fix is indispensable in the simulation of crystal background as well as low-energy/high-resolution radiation detectors. Fig. 9.2 shows the simulated and measured response of a ^{71}Ge miniature proportional counter before and after the bug fix [58]. Before the bug fix, the simulated response (Fig. 9.2a) hardly agrees with the measurement while after the bug, fix the simulation agrees very well with the measurement.

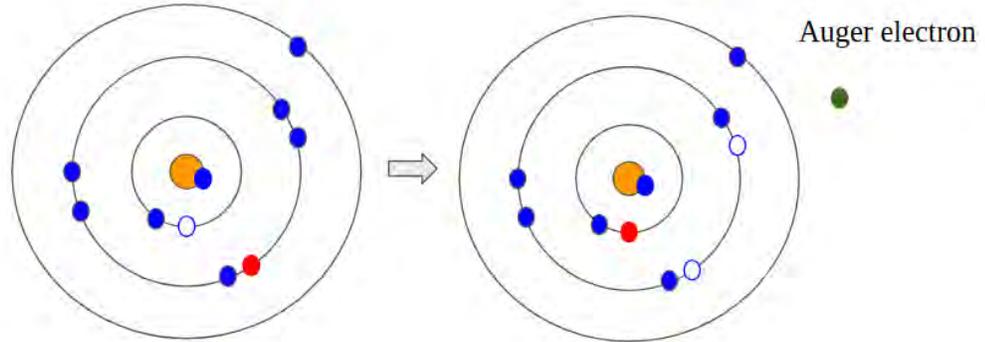
The avalanche of Auger electron has been officially implemented since Geant4 version 10.2 patch 02.

9.2 Background Calculation

Background refers to false events that can mimic features of signals from true events. Potential dark matter events are characterized by the presence of a scintillation signal in the NaI (Tl) crystal detector and the absence of signals in the liquid veto detector. Since we are mainly interested in events of a few keV_{ee} recoil energy, events that



(a) Illustration of emission of an X-ray. The red circle indicates an electron that changes orbit.



(b) Illustration of emission of an Auger electron. The red circle indicates an electron that changes orbit. The emitted Auger electron is marked with dark green.

Figure 9.1: Difference between X-ray and Auger electron emission. When an atom returns to the ground state by X-ray emission, the number of vacancies stay the same. However for Auger electron emission, one additional vacancy is created per Auger electron emitted.

deposit a few keV energy only in the crystal detector can become background to dark matter searches. Therefore, events that either have low rate, tend to deposit excessive energy in the crystal, or interact with the veto detector are not likely to be background.

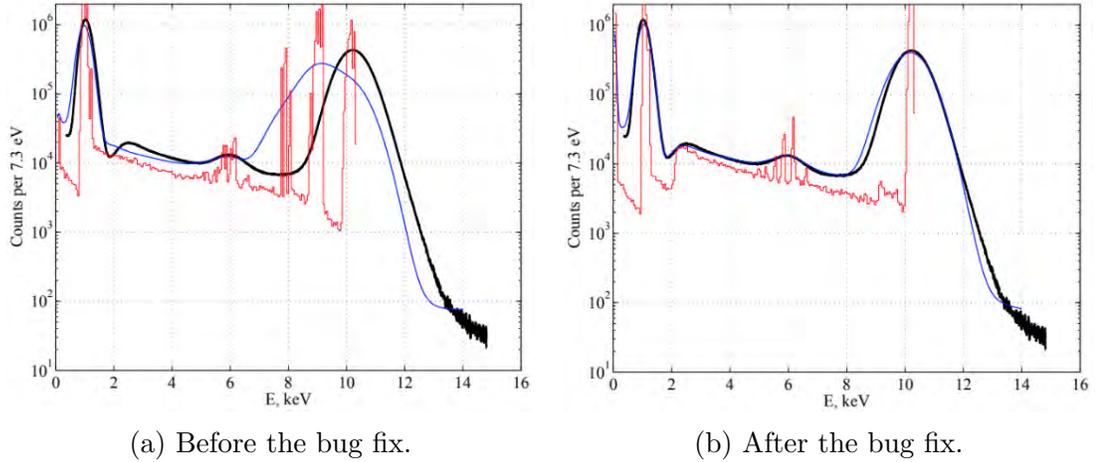


Figure 9.2: Simulated and measured response of a ^{71}Ge proportional counter. The simulated model response is shown as the solid red line while the blue line shows model response with detector resolution. The black line shows measured response. Proper handling of Auger electron avalanche yields a better consistency between the simulation and the measurement.

9.2.1 Analysis of Background

Table 9.1 classifies background in terms of source, type of interacting particle and how it can be suppressed. Each entry is discussed in detail below.

Table 9.1: Classification of background in terms of source, type of interacting particle and suppression strategy. In the table E means the interaction tends to deposit too much energy in the crystal detector to be background; V means the interaction can be vetoed by the veto detector; S means the interaction can be effectively suppressed by shielding provided by other components or small solid angles. The letter is repeated if the background can be strongly suppressed. Inherent event rate is not considered in this table.

Src\Type	Low-energy		High-energy		n	α	μ
	β	γ	β	γ			
crystal bulk			EE	V, E	V, E	EE	—
crystal surface			E	V, E	V, E	E, S	—
enclosure	SS	SS	S	V, E	V, E	SS	—
veto detector	SS	SS	SS, VV	S, V, E	S, V, E	SS, VV	—
shielding	SS	SS	SS, V	S, V, E	S, V, E	SS	—
external	SS	SS	SS, VV	SS, VV, E	SS, VV	SS	VV, EE

Cosmic ray: Cosmic rays are heavily suppressed by the overburden. Those deep-penetrating cosmic muons reaching the underground site have high mean energy loss on the order of MeV/cm in the detector medium, and can be effectively rejected by the veto detector or by excessive energy deposit in the crystal. Less frequent potential background due to cosmic muon activation, spallation and high-energy neutrons produced by muons in the rocks and the fluorescence of NaI induced by muons are not precisely calculated yet.¹

α particle: α particles from radioactive decays have high ionization density and a short range. Unless produced within a narrow distance from the crystal surface, α will not become background due to either excessive energy deposit or too short a range. This statement excludes background from secondary processes such as neutrons produced from α - n reactions.

Neutron: Neutrons can scatter off atomic nuclei and produce nuclear recoils similar to that expected in collisions with dark matter particles. The mean-free path of a slow neutron is relatively short, and interaction with atomic nuclei depends strongly on the nuclear structure of the target nuclei. Therefore, slow neutrons can be effectively shielded by proton-rich materials (e.g. polyethylene and water) or isotopes with large neutron-capture cross section (e.g. ^{157}Gd , ^6Li and ^{10}B).

However, high-energy fast neutrons have longer mean-free distances and interaction with nuclei is dominated by scattering instead of neutron-capture. To effectively suppress fast neutron background, a neutron veto system of enough thickness is needed.

High-energy β and γ : High-energy γ -rays produced outside the passive shielding can be effectively shielded by passive layers of lead, water and polyethylene, and an active layer of liquid scintillator. Along with those produced in the crystal and

¹Production of high-energy neutrons by muons is usually accompanied by a shower of electromagnetic radiation. If the interaction happens in the rock, the shower will be shielded and only deep-penetrating neutrons will escape.

enclosure components, γ -rays that penetrate passive and active shielding are further suppressed or constrained by the differential scattering cross section or kinematics of interaction. For instance, γ -rays of a few MeV predominantly interact by Compton scattering. There are three competing processes in estimating the background due to Compton scattering of high energy γ -rays.

1. As the photon energy gets higher, the range of electron recoil energy gets larger, and the maximum scattering angle that can yield a low-energy electron recoil gets smaller.
2. The overall probability for small-angle scattering angle gets smaller due to decrease in the solid angle.
3. According to Klein-Nishina formula [59], small-angle forward scattering has a larger cross-section compared to a large-angle scattering.

Although this problem can be solved analytically, it is easier to study it numerically. Fig. 9.3 shows the relative probabilities of electron recoiling with energy in the 1-10 keV range in Compton scattering with γ -rays of various energies. It can be seen that as the energy of the γ -ray increases, the probability of electron recoiling with a few keV kinetic energy decreases.

High-energy β particles do not travel very far and will either deposit excessive energy in the crystal or not reach the crystal. The induced bremsstrahlung can be suppressed in a manner similar to γ -rays.

Low-energy β and γ : Low-energy β -particles and γ -rays have very short ranges and will be strongly shielded from the crystal by the reflector materials or the PMT's quartz window. For instance, the range corresponding to the mean energy of ^{14}C β -decay is 0.12 mm while one layer of plumber's PTFE tape has a thickness of 0.09 mm, indicating that only first few layers can contribute to background.

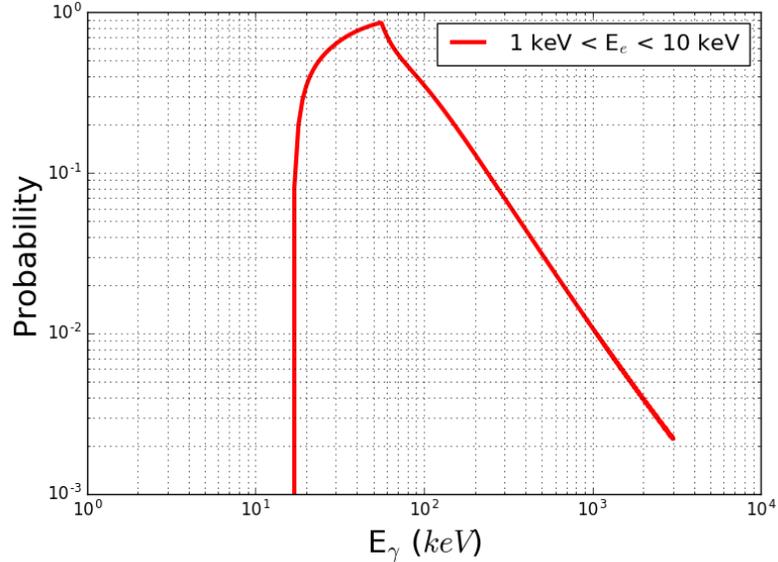


Figure 9.3: Probability of electrons recoiling with energy between 1-10 keV in a Compton scattering as a function of incident γ -ray energy. As the energy gets higher, the probability of low-energy electron recoil is suppressed. The low-energy cutoff is due to the γ -ray not having sufficiently high energy to produce a 1-keV electron recoil.

However, when the low-energy β or γ is produced in the crystal or on the crystal surface, the veto detector will not see any signal unless there is an accompanying high-energy γ -ray, and the crystal has a relatively high chance of seeing a few keV energy deposit, thereby mimicking a dark matter interaction.

9.2.2 Radioactive Isotopes in Low-background NaI (Tl) Crystal

From the discussion in Section 9.2.1, low-energy radioactive isotopes in or near the crystal poses the biggest threat to low-background dark matter searches. For a certain amount of impurity embedded in the crystal lattice, potential contribution to background can be predicted in terms of its radioactive isotope's specific activity, natural abundance and shape of spectrum. The element's tendency to stay in the crystal lattice is also critical, yet in most cases such data is difficult to obtain.

Isotopes with very high specific activities will decay away quickly and cannot persist in the crystal. On the other hand, isotopes with low specific activities will not decay enough times within the time window to be a big background, unless this is compensated by a high natural abundance (e.g. ^{87}Rb) or its presence in the crystal in large quantities (e.g. ^{129}I).

The shape of the decay spectrum determines the probability that the released energy lies within the interval of interest. For β -decays in the bulk, the calculation is straight-forward since in most cases crystal will see all the energy from the β -decay and the probability is simply the ratio of integral of spectrum over the interval of interest to that over the entire range up to the Q-value. Back-of-envelope calculations for γ -rays and surface contaminants are complicated due to the need to consider propagation of γ -rays or β -particles through the detector medium and the ability to veto accompanying high-energy γ -rays with veto detector.

β -decay Isotopes in the Crystal

Of all isotopes and impurities present in the crystal, ^3H , ^{22}Na , ^{40}K , ^{87}Rb and ^{210}Pb are particularly dangerous. These isotopes can emit low energy radiation as a result of β -decay or electron capture. Their half-lives are long enough such that their activities do not decrease significantly during the time scale of the experiment, and their decay rate combined with their natural abundance can yield non-negligible activities.

^3H : Tritium is a pure β -emitter with a small Q-value of 18.59 keV and a half-life of 12.32 years. The low-energy β particles from ^3H cannot be vetoed by the liquid scintillator, and a large fraction of the β -decay spectrum falls onto the region of interest. A simple calculation using Fermi's theory predicts about 39% of tritium decays would have energy between 2-6 keV.

^3H can get into NaI (Tl) crystal not only as spallation product of cosmic rays, but also as NaOH produced from reaction with water. However, the concentration of

Table 9.2: Summary of important pure β -emitters for low-background NaI (Tl) crystal detector. The natural abundance (N.A.) is expressed as negative log to compare the order of magnitudes between different isotopes. The fraction column summarizes fraction of β -decays with energy in the 2-6 keV interval per keV. It is a measure of the overlap between β -decay spectrum and the energy region of interest. Potential background contribution is expressed as background rate for 1 ppb of particular impurity, except ^{129}I which constitutes the crystal lattice. Half-lives and Q-values are taken from [60].

	$T_{1/2}$ yr	Spec. activity Bq/g	N.A. $-\log_{10}$	Q-value keV	Fraction keV^{-1}	Rate (per ppb) cpd/kg/keV
^3H	12.3	3.59×10^{14}	19-16	18.59	0.1	10^{-6} - 10^{-3}
^{14}C	5730	1.65×10^{11}	18-12	156.5	0.013	10^{-10} - 10^{-4}
^{36}Cl	3.01×10^5	1.22×10^9	14-12	709.6	0.0059	10^{-9} - 10^{-6}
^{87}Rb	4.81×10^{10}	3.2×10^3	0.55	282.2	0.008	0.62
^{129}I	1.57×10^7	6.5×10^6	14-12	149.3	0.00047	0.0024-0.24

^3H varies significantly depending on the source of water. Mainly produced by cosmic rays, nuclear reactors and atmospheric tests of nuclear weapons since 1952 [61], ^3H activity in precipitated water can be anywhere between 5 – 100 TU [62][63][64][65].² Production of tritium in underground water is dominated by $^6\text{Li}(n,\alpha)^3\text{H}$ reaction, and its concentration can be as low as 0.1 TU [61]. Without detailed knowledge of how NaI is produced from solution, it is hard to predict the amount of tritium background.

^{14}C : Produced mainly in interactions with cosmic rays, ^{14}C exists in the Earth's atmosphere and in all living organism with a natural abundance of about 10^{-12} . Once the organism stops intake of carbon, concentration of ^{14}C begins to decrease with a half-life of 5730 years. Therefore, $^{14}\text{C}/^{12}\text{C}$ ratio in petroleum products can be many orders of magnitude lower.

Once in the crystal, the low-energy end of the ^{14}C β -spectrum can become background that cannot be vetoed. ^{14}C can be originally present in the powder (Section 3.2) or get introduced into the crystal through contact with pyrolytic carbon

²TU, or tritium unit is defined as the ratio of 1 tritium atom to 10^{18} hydrogen atoms. Sometimes it is also used as a measure of tritium activity in water, in which case 1 TU= 0.118 Bq/L.

coating during crystal growth. The solubility of carbon in molten NaI is not precisely known, but given the relatively large bond dissociation energies of carbon allotropes and based on the fact that carbon tends to be pushed to the tail end of the crystal during crystal growth, the solubility could be very small.

According to Table 9.2, assuming the worst scenario of atmospheric carbon, 1 mg of carbon is needed per 1 kg of NaI powder to dominate background. On the other hand, the mass of 500 nm-thick carbon coating in contact with 1 kg of molten NaI in a 4"-diameter crucible is 12 mg. Assuming all the carbon coating end up in the crystal, the background from ^{14}C is easily suppressed by using hexane produced from petroleum.

^{36}Cl : Compared to ^{14}C , ^{36}Cl has a lower specific activity due to a longer half-life of 3.01×10^5 years. Its branching ratio of β -decay is about 98% and the probability of overlap between its β -spectrum and the interval of interest is about 0.16%/keV. However, with about 2% branching ratio, ^{36}Cl can also decay by electron capture, where the atomic relaxation will produce low-energy radiations in the interval of interest. This channel contributes about twice as much probability compared to the β -decay channel.

Although not present in the NaI (Tl) crystals in large quantities, ^{36}Cl can be introduced to the crystal by SiCl_4 used to prevent crystal sticking (Section 4.4). Unlike carbon, NaCl can dissolve in NaI, and zone refining is needed to remove it from the crystal (Section 3.3.3).

In the atmosphere, ^{36}Cl is mainly produced by spallation of ^{36}Ar by protons. In the lithosphere, neutron-capture by ^{35}Cl stays an important process in both shallow and deep sites while for shallow sites, spallation of ^{40}Ca and ^{39}K is another important process. Near the surface, $^{36}\text{Cl}/\text{Cl}$ depends on the amount of ^{39}K and ^{40}Ca in the ore and it can be as high as 10^{-11} . For deep sites, cosmic ray-induced processes are negligible and the production is dominated by $^{35}\text{Cl}(n,\gamma)^{36}\text{Cl}$ reaction. $^{36}\text{Cl}/\text{Cl}$ ratio

is essentially independent of the mineral composition and is on the order of 10^{-14} . Should chlorine be present in the final crystal lattice, it would be important to choose SiCl_4 with underground origin instead of surface or marine origin.

^{87}Rb : ^{87}Rb is a primordial isotope with half-life longer than the age of the Universe. Thus it has the lowest specific activity of all β -decay isotopes discussed here, yet this is counterbalanced by its low Q-value and a high natural abundance. ^{85}Rb has an upper limit of 0.2 ppb measured by ICP-MS. Projected background at this level is 0.17 cpd/kg/keV. Measurement with higher sensitivity is challenging and direct counting is needed.

^{129}I : ^{129}I is primarily produced as a nuclear fission product. The natural abundance of ^{129}I has a large variation depending on when and where the sample is taken: earlier measurements reported values in 10^{-14} - 10^{-12} range [66][67][68], but values in the 10^{-10} - 10^{-9} have been observed [68], especially after the Fukushima nuclear incident in 2011 [69]. If iodine is taken from ore instead of from marine sources, the abundance tends to be lower. Previous NaI (Tl) detector has measured $^{129}\text{I}/^{127}\text{I}$ to be about 10^{-13} [21] and similar values have been obtained for the Astro-grade NaI powder using accelerator mass spectroscopy (AMS).

Compared to other β -isotopes, iodine exists in the crystal as a constituent of the crystal lattice. However, ^{129}I decays to the 39.6-keV excited state of ^{129}Xe . Since this energy will be summed onto the β -spectrum, the crystal will not see as much background from low-energy β -particles as other β -emitters. Decay of ^{129}I directly to the ground state of ^{129}Xe is a unique second-forbidden decay that has never been observed. The only background from ^{129}I comes from events where the 39.6-keV X-ray manages to escape the crystal via Rayleigh scattering ($\sigma_{\text{Rayl}} \approx 0.02\sigma_{\text{tot}}$) or Compton scattering ($\sigma_{\text{Comp}} < 0.01\sigma_{\text{tot}}$). From the Geant4-simulated low-energy part of ^{129}I spectrum in a 5-kg NaI (Tl) crystal (Fig. 9.4), the spectrum near 2-6 keV is suppressed by about 2 orders of magnitude.

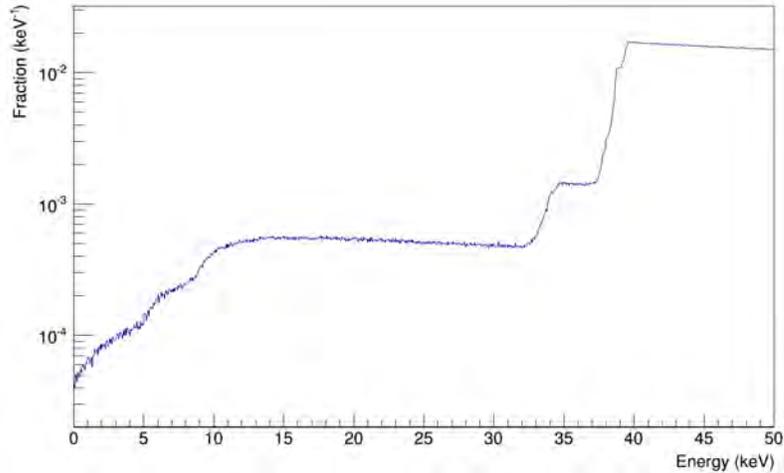


Figure 9.4: Simulated low-energy end of the spectrum of ^{129}I decay in NaI crystal. The spectrum begins to rise after around 35 keV due to energy released during the atomic relaxation. This summing effect moves the low-energy part of the β -spectrum away from the interval of interest.

γ -decay Isotopes in the Crystal

Often γ -rays from nuclear decay come from relaxation of daughter nuclei. If the nuclear decay is followed by a rearrangement of orbital electrons, low-energy X-rays and/or Auger electrons of a few keV will be produced. When the γ -rays escape the, the crystal sees only the low-energy radiations. Liquid scintillator veto outside the crystal is a powerful way to suppress such background. In the NaI (Tl) crystal, two main sources of γ -background are ^{22}Na and ^{40}K .

Table 9.3: Summary of important γ -emitting isotopes for a low-background NaI (Tl) crystal detector. The natural abundance for ^{22}Na is obtained by germanium γ -spectroscopy. BR_{EC} lists the branching ratio of electron capture.

	$T_{1/2}$ yr	Spec. activity Bq/g	N.A. $-\log_{10}$	E_{γ} MeV	$E_{\text{X-ray}}$ keV	BR_{EC} %	Rate cpd/kg
^{22}Na	2.6	2.3×10^{14}	20	1.27	0.85	9.6	2.93
^{40}K	1.25×10^9	2.65×10^5	4	1.46	2.9 – 3.2	10.55	0.023/ppb

^{22}Na : ^{22}Na has a half-life of 2.6 years and undergoes β^+ -decay or electron capture. In the latter case, along with a 1.27 MeV γ -ray, low-energy X-rays/Auger electrons

are produced with energy around 0.8 keV which is typically outside the interval of interest, but will become the dominating background with lower threshold.

The high-energy γ -ray and annihilation radiation from e^+ make it detectable by high-purity germanium detector. In 2010, γ -spectrum of Astro-grade NaI powder was measured underground at LNGS and ^{22}Na activity per 1 kg of NaI powder was determined to be 0.8 ± 0.1 mBq/kg. This corresponds to a natural abundance of roughly one part in 10^{20} , but its high specific activity counterbalances its small natural abundance. The electron-capture decay rate of ^{22}Na is about 3 cpd/kg, which can seriously affect low-energy sensitivity.

^{22}Na is mainly produced by cosmic activation of ^{23}Na . Since Na is one of the constituents of the NaI crystal lattice, the only way to suppress the background from ^{22}Na is to use underground sodium as the starting material for growing NaI (Tl) crystals and minimize its time above ground. Due to a relatively short half-life, the production rate and decay rate of ^{22}Na near the surface must be equal. If underground sodium has negligible ^{22}Na , about 10^4 ^{22}Na will be produced during 3 months of crystal growth and fabrication. This corresponds to an isotopic abundance of $\sim 10^{-21}$ and henceforth background 10 times lower.

^{40}K : ^{40}K is a primordial isotope with a very long half-life. The dominating decay channel of ^{40}K is β -decay with a Q-value of 1.31 MeV. According to the analysis in Section 9.2.2, such high Q-value will suppress its contribution to low-energy background by $\sim 10^3$. However, about 10% of ^{40}K will decay by electron capture to ^{40}Ar , with emission of a high-energy γ -ray and low-energy X-rays/Auger electrons. When the accompanying γ -ray escapes the crystal without interaction, the low energy radiation becomes background unless the γ -ray is tagged and vetoed.

The presence of K in the Earth's mantle and its similar chemical properties to sodium make it relatively difficult to separate K from NaI. One approach is to exploit

the difference between solubilities of Na_2CO_3 and K_2CO_3 . Another possibility is zone refining (Section 3.3).

β -decay Isotopes on the Surface

Contaminants on the crystal surface and in the reflector material are slightly different from bulk contaminants because the interaction happens only within a sub-millimeter distance from crystal surface. In this section, ^{14}C and ^{210}Pb are analyzed.

^{14}C is present in the PTFE tape wrapping the crystal, and ^{210}Pb —a radon daughter—can get deposited on the surface if the crystal is handled in an environment with radon.

Table 9.4: Summary of important β -emitting surface contaminants for a low-background NaI (Tl) crystal detector. In the case of ^{14}C , the fraction measures probability of β -particles penetrating the PTFE tape and depositing an energy in the 2-6 keV range. The rate is calculated for the entire PTFE reflector for a 4''-diameter 5-kg crystal.

	$T_{1/2}$ yr	Spec. activity Bq/g	N.A. $-\log_{10}$	Q-value keV	Fraction keV^{-1}	Rate cpd/kg/keV
^{14}C	5730	1.65×10^{11}	18 – 12	156.5	0.000026	$10^{-5} - 10$
^{210}Pb	22.3	2.83×10^{12}	–	63.55	–	–

^{14}C : low-energy electrons from β -decay of ^{14}C in the PTFE tape wrapping the crystal can contribute to the background. The range of these electrons are very short and only first 1-2 layers of PTFE tape can contribute to the background.

The simulation assumes 10 layers of PTFE tape with density approximately 1 g/cm^3 wrapping a 4''-diameter 5-kg NaI (Tl) crystal. The thickness of each PTFE layer is taken to be 0.18 mm. The result shows roughly 26 electrons per 1 million ^{14}C decays in the entire 1.8-mm PTFE tape end up in the 2-6 keV window. If the carbon in the PTFE comes from contemporary carbon sources, the background due to ^{14}C is on the order of 10 cpd/kg/keV. On the other hand, if the carbon comes from

petroleum, the background will be negligible. Therefore it is important to measure $^{14}\text{C}/^{12}\text{C}$ ratio in advance for the PTFE tape to be used.

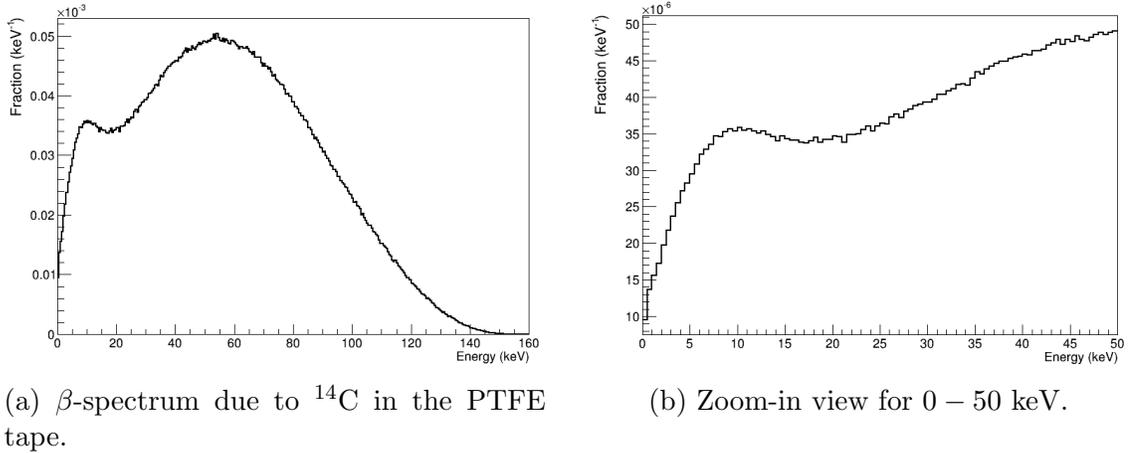


Figure 9.5: Energy deposit in the NaI (Tl) detector by β -decays of ^{14}C in the PTFE tape wrapping the crystal. Compared to normal β -spectrum, low-energy end is heavily suppressed. Roughly 26 electrons per 1 million ^{14}C decays in the 1.8-mm PTFE tape end up in the 2-6 keV window.

^{210}Pb : ^{210}Pb has the longest half-life among radioactive isotopes of lead. ^{210}Pb mainly decays by β decay, becoming a background when present either in the bulk or on the surface of the crystal. ^{210}Pb can be introduced into the crystal as a trace element in the water used to make NaI solution. Since it is mainly produced from radon-222 as part of ^{238}U decay chain, handling the crystal in a radon-rich environment can also contaminate the crystal surface with ^{210}Pb .

Quantitative determination of ^{210}Pb is difficult via mass spectroscopy, and the only option is direct underground counting. Other NaI (Tl)-based dark matter experiments have been reporting ^{210}Pb contamination [52][53]. To avoid surface contamination as much as possible, most of powder and crystal handling has been carried out in a glovebox inside a radon-free cleanroom.

9.3 Monte-Carlo Simulation

Geant4-based Monte-Carlo simulation was used to estimate the background in dark matter search. The NaI (Tl) crystal, crystal PMTs, enclosure, liquid scintillator, veto vessel, veto PMTs were modeled in the simulation. External shielding assumes 90 cm of water shielding (see Section 6.3).

Concentration of ^{40}K is assumed to be 9 ppb (Table 4.2) and that of ^{87}Rb 0.2 ppb (upper limit). Concentrations of ^{238}U (0.5 ppt) and ^{232}Th (0.5 ppt) are obtained from ICP-MS measurements by PNNL and ^{22}Na is obtained from underground counting with germanium at LNGS (Section 9.2.2). Abundance of ^{14}C is assumed to be 10^{-15} , approximately the detection limit of ^{14}C in AMS. ^{129}I concentration is measured by AMS. ^{36}Cl is not considered since the impact of SiCl_4 is unknown.

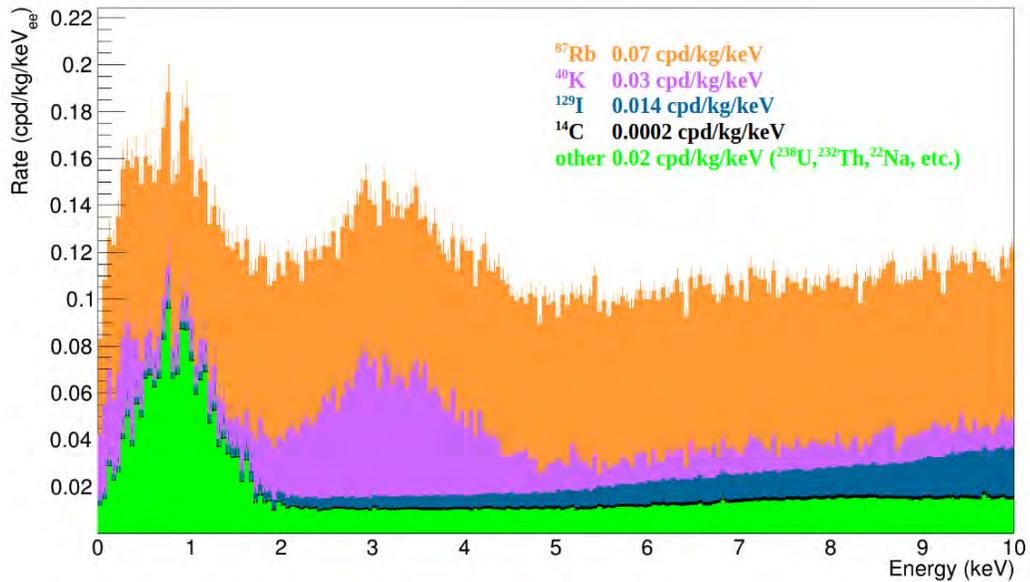


Figure 9.6: Simulated background rate from various sources. At low energies, ^{87}Rb , ^{40}K and ^{129}I dominate the background rate. The peak near 1 keV is due to electron-capture decay of ^{22}Na .

The estimated total background in 2-6 keV interval is 0.14 cpd/kg/keV. The simulated background spectrum, shown in Fig. 9.6, is dominated by ^{87}Rb and ^{40}K in the crystal to the first order. As discussed in Section 9.2.2, ^{22}Na becomes more

important for low energies. Background from other components are sub-dominant. It is worth pointing out that the ^{87}Rb background is only an upper limit based on ICP-MS measurements. Actual ^{87}Rb concentration could be significantly lower.

Later an updated Monte-Carlo was done to evaluate backgrounds from new copper enclosure design, cables and electronics in the crystal detector module, and cosmogenic isotopes. Although it predicts a slightly higher background of 0.22 cpd/kg/keV due to cosmogenics, the result of updated simulation is consistent and basic conclusions remain unchanged [70].

Chapter 10

Conclusions and Future Perspectives

The annual modulation seen by the DAMA/LIBRA experiment still remains a major controversy in the quest of direct detection of dark matter. To resolve it, a NaI (Tl)-based dark matter detector with even lower background is needed, and ultra-high purity NaI (Tl) scintillating crystal is an essential component of such a detector.

In this thesis, I have demonstrated that ultra-high purity in NaI (Tl) crystals can be achieved by a combination of high-purity NaI powder, precision-drying of the powder and precision-cleaning and pyrolytic carbon coating of high-purity synthetic fused silica crucibles. The crystal purity can be improved further by zone refining.

One of the crystals grown, NaI-031, is optically-coupled with two PMTs and packaged into a crystal detector module and tested at Princeton. The result indicates that the light yield of NaI-031 is not as good as typical small samples taken from the tips of crystals. In addition, a test with collimated γ source indicates that the crystal is subject to large nonuniformity.

The veto detector needed for low-background scintillation counting has been commissioned in LNGS Hall B, and subsequently cleaned and moved to its final location

in Hall C. The fluid handling system has been installed and currently the vessel is ready for liquid scintillator filling. The trigger and data acquisition system needed for operating the prototype detector has been designed, implemented and verified.

Finally, potential background due to various crystal impurities are analyzed and Monte-Carlo simulations are used to assess the projected background. Based on currently available information, the projected background is on the order of 0.1 cpd/kg/keV, lower than any other NaI (Tl)-based dark matter experiments.

However, ultra-high purity NaI (Tl) detector development is still an on-going effort, and a number of improvements are necessary for more reliable crystal growth and better performance of the crystal detector module.

Zone Refining of NaI Powder

So far, this thesis has solved “how not to introduce more contaminants”, but in this case the best that can be achieved is a reduction in impurities by a factor of 2.

Zone refining is a powerful technique that addresses the question of “how to remove existing contaminants”. Zone refining holds the key to even higher-purity crystals with K concentration below the detection limit of ICP-MS.

We have co-invented a new zone-refining system with Mellen. However, before zone-refined powder can be used for crystal growth, a series of detailed studies are needed to fine-tune the zone-refining system and study how to cut the zone-refined ingot and fill in crystalline NaI into the crystal growth crucible.

Silicon Tetrachloride Treatment

A new breakthrough in crystal growth is achieved by treating NaI powder with SiCl_4 such that the crystal doesn't stick to uncoated quartz. However, the use of SiCl_4 comes with other risks, notably the potential effect of chlorine on the scintillation of the final crystal and the background introduced by radioactive isotopes of chlorine.

To evaluate the risk of SiCl_4 , a small-scale crystal growth test is currently underway to directly measure the light yield and other scintillation properties of a SiCl_4 -treated NaI (Tl) crystal grown using the Astro-grade powder.

Fine-tuning of Crystal Growth

In crystal growth, the scale-up of crystal size is accompanied by new challenges. Currently, growth of large-diameter single crystals are mainly based on past experience and empirical practice. However, to ensure quality and consistency of the final crystals, crystal growth should be carefully studied and fine-tuned in terms of temperature distribution within the molten NaI, growth rate, and the uniformity of Tl.

Improvement on Crystal Detector Module

Although current crystal enclosure has been successfully assembled, there are a number of shortcomings.

1. The enclosure is too heavy.
2. Deformation of plastic parts after acid treatment and/or vacuum baking should be considered and tolerance should be increased accordingly.
3. A better sanding and polishing technique is needed in case the crystal has surface defects.

Preparation of Full-scale Experiment

In a full-scale experiment searching for annual modulation, DAQ hardware and software and trigger firmware must be scaled up to accommodate more ADC channels. However, this can be done easily under the Polaris framework by implementing added ADC channels as add-on modules. The trigger is also modular in design and currently it uses less than 1% of FPGA's resources, and thus is easily extendable.

Appendix A

Simulation of Impurity

Distribution in Zone Refining

1. Partition the ingot into N cells of equal length, and let C_i and R_i represent the impurity concentration and the radius at location i (i ranges from 0 to $N - 1$). If the crucible does not have rotational symmetry, in subsequent steps, R_i^2 should be replaced with the corresponding cross-sectional area. Let another integer w represent the width of the zone in terms of the number of cells.¹ The total volume of the molten zone and total impurities are denoted as V and Q , respectively.
2. Iterate the index i from w to $N - 1$. As the zone moves through each index, add impurities from the newly-molten cell $i + 1$ and increase the volume accordingly:

$$Q_i = Q_i + R_{i+1}^2 C_{i+1}, \quad (\text{A.1})$$

$$V = V + R_{i+1}^2. \quad (\text{A.2})$$

¹The impurity distribution depends only on the ratio of zone width to the total length of the ingot. Using integers will simplify the computer program. For higher resolution, simply increase N and w while keeping w/N fixed.

3. Freeze the $(i - w)$ th zone. Assign the impurity concentration in the solid phase according to the distribution coefficient: $C_{i-w} = kQ/V$. Decrease the total impurity Q and the volume of molten zone V accordingly:

$$Q = Q - kC_{i-w}R_{i-w}^2, \quad (\text{A.3})$$

$$V = V - R_{i-w}^2. \quad (\text{A.4})$$

4. Once i reaches $N - 1$, the zone has reached the boundary. At this point, no more impurities should be added to the melt from newly-molten cells. Therefore iterate i from $N - 1$ to $N + w - 1$ and apply only Step 3 at each cell.

Glossary

Accelerator mass spectroscopy a type of mass spectroscopy in which isotopes are accelerated and separated by a particle accelerator. It is frequently used to detect rare isotopes among neighbouring abundant isotopes (e.g. ^{14}C and ^{129}I).

145

ADC analog-to-digital converter. 89, 91, 112–114, 117, 118, 126–132, 154, *Glossary*: analog-to-digital converter

Afterpulse a signal in a PMT after the main pulse. Afterpulse is mainly caused by residual gas ionized by accelerated photoelectrons and drifting back to the photocathode under the electric field. In PMTs with good timing, the drift time is roughly the same for different drift paths. Therefore, afterpulse also tends to appear after certain fixed times relative to the main pulse. 92, 93

AMS accelerator mass spectroscopy. 145, 150, *Glossary*: accelerator mass spectroscopy

Analog-to-digital converter an electrical circuit component that converts analog signals to a digital number that reflects the amplitude of the analog signal. 89

Annealing a heat treatment that changes physical and chemical properties. In this thesis, annealing mostly refers to quartz glass, and is achieved by heating the glass to high temperature for a short amount of time and letting it cool down to remove internal stresses. 25, 26

Annual modulation As the Sun moves in the dark matter halo at a constant speed and the Earth orbits around the Sun, there is a change in the relative speed between the Earth and the dark matter halo with a period of one year. This change in relative speed induces a yearly change of dark matter interaction rate in detectors. 6, 9–11, 152, 154

Anti-coincidence a method frequently used to suppress unwanted background in nuclear and high energy physics by requiring the presence of signals in some channels and the absence of signals in other channels. 8, 9

ASA flange an asymmetric flange designed to seal with elastomers. One of the pair of the flanges has a groove to hold an elastomer (usually an O-ring) and the other flange has a flat surface. The elastomer is squeezed by tightening the flanges with bolts using holes located outside the groove. 100

Astro-grade a brand name of NaI powder by Sigma-Aldrich featuring very low amounts of impurities. 11, 39, 42, 45–48, 60, 66, 69, 71–73, 75, 76, 78, 83, 145, 147, 154

Baryon a type of subatomic particle consisting of odd numbers of quarks. Protons and neutrons fall into this category. 2, 3

BGO bismuth germanate. 109, 110, *Glossary*: bismuth germanate

Bismuth germanate a crystal scintillator. 109

Bolometer a device for measuring the energy of incoming ionizing radiation or electromagnetic radiation by sensing temperature change. 7

Borexino a neutrino experiment located at LNGS Hall C. 100, 110

Borosilicate a type of glass with silica and boron trioxide as main constituents. Borosilicate glass has a relatively low thermal expansion coefficient compared

to normal soda-lime glass (borosilicate glass: 3 ppm/°C, soda-lime glass: 10 ppm/°C, quartz: 0.55 ppm/°C). It softens at around 800 °C and thus are much easier to work with compared to quartz glass. Borosilicate is often used to construct laboratory glasswares. Often the word Pyrex is used as a synonym to borosilicate glass. 41, 82, 100, 101

Bridgman a crystal growth method. See Section 4.1. 61–64, 66

Buffer memory used temporarily to store input or output data while the previous data is still being processed. 122, 124, 125, 130, 136

CAEN a company that manufactures nuclear electronics. The VME crate, trigger and digitizer boards used in this thesis were bought from this company. 89, 102, 113, 114, 116, 118, 125, 131

Cathode a negatively charged electrode of a device. For some PMTs (e.g. R11065), the cathode of the PMT is connected to the photocathode of the tube, and cathode sometimes refers to the photocathode of the PTMs. 86, 89, 91, 108

Cleanroom a controlled environment with low level of particulates. 22, 29, 64, 96, 110, 111, 149

Clock a digital signal that oscillates between a high-state and a low-state at a fixed frequency. Clock signals are important, especially for synchronous circuits, because many digital operations (e.g. reading a memory) are driven by clock transitions. 114, 117–119, 130

CMB cosmic microwave background. 2–4, *Glossary*: cosmic microwave background

Coating a process by which the surface property of a material is changed by an added layer of another material. In the context of this thesis, coating refers to

pyrolytic carbon coating, unless explicitly stated otherwise. 15, 19, 22–38, 40, 63, 65, 66, 71, 72, 74, 78–80, 120, 125, 127, 128, 144, 152

Coincidence a method frequently used to pick up events of interest in nuclear and high energy physics by requiring the presence of signals in two or more channels. Coincidence can effectively lower random backgrounds. 9, 12, 112–119

Cold trap a device (often made out of stainless steel or borosilicate glass) that condenses the gases with low vapor pressure at the temperature of the cold trap. In the laboratory, ice water, dry ice or liquid nitrogen are often used to cool the cold trap. 41–48, 82, 126

Compton scattering the process by which a high-energy photon collides with an electron, transferring photon’s energy and momentum to the electron. 107, 140, 141, 145

Cosmic Microwave Background remnant electromagnetic radiation created when electrons and protons recombined and photons decoupled from matter as a result. 2

Cosmic muon muons found in cosmic rays. These muons are deeply-penetrating. At sea level, cosmic muon flux is approximately 1 muon/cm²/s. See muon. 6, 10, 13, 110, 139

Cross section a measure of the likelihood for a process to happen, expressed in units of area. Intuitively, it is simply the area transverse to motion within which particles must meet in order to interact.

In classical inelastic hard spheres, cross section is simply the cross sectional area of the spheres. In quantum mechanics, cross section is correspondingly weighed by probability. 4, 6–8, 10, 139, 140

Crystal-grade a brand name of NaI powder by Sigma-Aldrich with impurity concentrations on the ppm-level. Earlier crystal growths have all failed, most likely due to an excessive amount of NaOH impurity in the crystal. 42, 65, 66, 71, 72

Czochralski a crystal growth method. See Section 4.1. 61, 62

DAMA/LIBRA a dark matter experiment. See Section 1.2. 9–11, 39, 152

DAQ (data acquisition) the process of digitizing and storing signals. 16, 25, 42, 69, 109, 120, 121, 125, 126, 129–131, 154

Dark matter a hypothetical form of non-luminous matter that can explain motion of galaxies and structure formation. 1–10, 104, 113, 136, 137, 139, 141, 149, 150, 152, 153

Deionized water water from which ions have been removed. Deionized water is a poor electrical conductor, and the resistivity is about 18.2 M Ω ·cm. 22, 24, 29, 66, 95, 110

Delrin a brand name for polyoxymethylene, an engineering thermoplastic that is often used in machining precision parts. 94, 96

Detergent 8 a detergent developed by Alconox. 22, 66, 95, 96

Devitrification growth of crystalline structures on the surface of glass. Devitrification are often caused by surface contamination at high temperatures, and will lower the thermal and mechanical strengths of the glass. 26

D flip-flop a circuit made from digital logic gates that have two stable states. Therefore, D flip flops can be used as memory to store information. Update of the state is triggered by a clock transition. 114

Erlenmeyer a laboratory glassware with a flat bottom, a conical body and a cylindrical neck. 23, 24, 29, 81, 82

Ethanol also called alcohol. Ethanol is a hydrocarbon with two carbon atoms that is frequently used as a cleaning agent and a solvent in laboratories. 15, 24, 29, 68, 76

Field-programmable gate Array an integrated circuit with programmable logical units and inter-connections. Various custom logic can be electrically programmed on site. 112

FIFO a type of data structure or buffer for queuing. In a FIFO, the oldest (first) entry is processed the first. 122, 124, 125, 136

Finite-state machine also called finite-state automaton. It is a computational model in which system is represented by a number of discrete and distinct states, with strictly defined state transition conditions. In this architecture, it is easy to guarantee that 1) system is always in one of the allowed states, and 2) state transition is strictly triggered by certain well-defined conditions. 121

Firmware permanent or semi-permanent software programmed into a read-only memory of hardware. 16, 109, 112, 113, 117, 119, 126, 128, 129, 154

FPGA field-programmable gate array. 112–114, 117, 129, 154, *Glossary*: field-programmable gate array

Fused quartz an amorphous form of quartz. Usually made by fusing naturally-existing quartz minerals or silica sand. The difference with synthetic fused silica is crucial in this thesis. 17–20, 24, 26–28, 31, 33, 40, 41, 56, 65, 66, 69, 71, 73, 75, 78–83

Fused silica Historically fused silica and fused quartz both refer to quartz glass made from silica sand or quartz mineral. Today in glass industry fused quartz refers to quartz made from natural sources of SiO_2 and fused silica is reserved for quartz glass made from purified synthetic source of silicon, such as silicon tetrachloride. Due to significantly higher-purity, synthetic fused silica has even higher transmissivity in the UV spectrum compared to quartz, and thereby used extensively to make optical fibers. The word synthetic will always be added when referring to synthetic fused silica to emphasize the difference. 17–20, 28, 65, 66, 71, 74, 76, 79, 83

GDMS glow discharge mass spectroscopy. 57, *Glossary*: glow discharge mass spectroscopy

Geant4 a nuclear and particle physics simulation toolkit. It can be used to simulate the propagations of radiations in detectors and other materials. 103, 107, 135, 136, 145, 150

Glovebox a hermetically-sealed box with gloves to manipulate oxygen- or moisture-sensitive materials inside. The glovebox is often purged with dry nitrogen or argon. 24, 29, 43, 63, 64, 96, 149

Glow discharge mass spectroscopy a type of mass spectroscopy that ionizes sample atoms by passing electric current (glow discharge) and analyzes atomic composition by a subsequent mass spectrometer. 57

HEPA (high-efficiency particulate absorber) a type of filter that forces the gas through a fine mesh to trap particulates. 22, 111

Hexane an alkane with six carbon atoms. In this thesis, it is used as a carbon source for pyrolytic carbon coating. 23–29, 36–38, 144

Hydrochloric acid a strong acid with maximum concentration around 30%. 22, 80–82, 95, 96

Hydrofluoric acid a weak and toxic acid often used for cleaning and etching in the semiconductor and glass industry. 22, 24

I/O input/output. 126

ICP-MS inductively coupled plasma mass spectroscopy. 20, 40, 56, 70, 74, 76, 79, 145, 150, 151, 153, *Glossary*: inductively coupled plasma mass spectroscopy

Inductively coupled plasma mass spectroscopy a type of mass spectroscopy in which sample atoms are ionized by inductive heating and separated by a subsequent mass spectrometer. It offers high sensitivity to atoms that are easy to ionize. 20

Ionization detector a type of ionizing radiation detector that detects incoming radiation by collecting ionized electrons in gas, liquid or semiconductor using electric fields. 7

Ionizing radiation particles and electromagnetic waves with sufficient energy to free electrons from atoms and molecules. In the context of this thesis, ionizing radiation includes part of ultraviolet radiation, X-rays, γ -rays, α -particles, β -particles, neutrons and other heavy ions. 7

Isopropyl alcohol also known as iso-propanol. At room temperature, it is a colorless and flammable liquid with properties very similar to ethanol. It is a structural isomer of 1-propanol (same chemical formula but different structures). In this thesis, isopropyl alcohol is used for crystal polishing as well we cleaning of detector components. 66, 85–88, 96, 110

Kapton a brand name for a polyimide film developed by DuPont. It has a high dielectric strength and a wide range of service temperatures. Since a thinner insulation is sufficient for high voltage, Kapton-insulated high-voltage cables are more flexible. 95

KF a vacuum flange. Also known as Klein Flange or QF (Quick Flange). The flange is sealed by squeezing an O-ring with two symmetric flange surfaces tightened by a circular clamp. In case there are minor scratches or unevenness on the flange surface, vacuum grease can be applied onto the O-ring to improve vacuum seal. 23, 24, 41, 101

Kovar a trademark of a nickel-cobalt alloy designed to match the thermal expansion coefficient of borosilicate glass. It can be sealed directly to borosilicate glass with relatively high mechanical strength. 41

Kyropoulos a crystal growth method. See Section 4.1. 9, 61, 62

LED light emitting diode. 89, 90

Library (software) a collection of computer codes and resources that can be used by different applications. For a static library, the resources are loaded into the program (the library is linked) at the time of compiling the program. For a shared library, the contents of the library is loaded at runtime, although the names of the functions to be loaded in the library must be known at compilation time. For a dynamic library, the resources to be used are specified at runtime. Polaris uses a dynamic library to load different modules at runtime. 120–122, 125, 129, 131

Light yield a property of scintillators that measures the amount of light (in number of photons) per unit incoming energy. In this thesis, the unit of light yield is pho-

toelectrons per keV (p.e./keV) while in literature, other units (e.g. p.e./MeV) are frequently used. 14, 69, 73, 81, 84, 87, 88, 94, 96, 97, 99, 105–108, 152, 154

LNGS an underground laboratory located in L’Aquila, Italy. 14, 16, 100, 105, 109–111, 147, 150, 152

Low-voltage differential signaling an electric communication standard that utilizes the difference of a signal and its inverse to boost signal and reject common-mode noise during signal transmission. 113

Lumirror a trade name of a polyester film developed by Toray Plastic, Inc. Owing to its high reflectivity, Lumirror is often used as reflectors in scintillator applications and other physics experiments. 87, 100, 107, 108, 110, 111

LVDS low-voltage differential signaling. 113–115, 118, 130, *Glossary*: low-voltage differential signaling

Majority coincidence coincidence between multiple channels. The minimum number of channels that need to be in coincidence is specified by the majority level. 109, 110, 113, 115–118

Majority level see majority coincidence. 110, 114, 115, 117

Mellen a manufacturer of high-temperature furnaces and refractory materials based in Concord, New Hampshire. 59, 153

Microcontroller a small computer on a single integrated circuit equipped with analog and digital peripherals to communicate with and issue control signals to other devices. Microcontrollers can be programmed to execute certain tasks on a firmware level. 25, 30, 126, 127, 129

Multiplexer a circuit component or a device that selects one of several input signals and forward it to its output. A multiplexer has an address line that specifies which input signal to select. 116

Muon an elementary particle similar to, but heavier than an electron. Muons have an electric charge of -1 and a mass of around 106 MeV. It decays into an electron and two neutrinos with a half-life of about $2.2 \mu\text{s}$. Since muons are heavier than electrons, it loses much less energy due to bremsstrahlung in the interaction with nuclei. 8, 12, 139, 160

NaI (Tl) thallium-activated sodium-iodide. An inorganic crystal scintillator with high light yield that is widely used in γ -spectroscopy. For its properties, see Section 1.4. 9–15, 17, 20, 61–63, 81, 83–88, 93, 94, 96, 97, 102, 105, 112, 132, 136, 142–150, 152–154

NaOH a strong base formed by sodium and hydroxide. 14, 40, 41, 46, 65, 72, 76, 78–83, 142

Neutrino an elementary particle that interacts only via weak interaction and gravity. Neutrinos come in three families: electron neutrinos, muon neutrinos and tau neutrinos. 3, 4

Nitric acid a strong and oxidizing acid. 22

Nucleosynthesis the process of making new atomic nuclei. It can include Big Bang nucleosynthesis (mostly hydrogen, helium and trace lithium), stellar nucleosynthesis (up to iron), supernova nucleosynthesis, cosmic ray spallation (nuclear reaction with cosmic rays) and radiogenesis (production of nuclei from radioactive decay). 8

Nylon a tough and lightweight elastic synthetic polymer. Compared with other plastics, it has relatively high mechanical strength. 96

O-ring an elastomer in the shape of a torus that is frequently used for seals by compressing it in a groove. For proper seals, the grooves have to be designed and manufactured in the right dimensions. 21, 95, 101, 106

Outgassing release of absorbed or adsorped gas. 93, 105

PCIe (Peripheral Component Interconnect express) a circuit board that can be inserted to designated slots on the computer motherboard to expand the computer's functionality. 125

Phonon a quanta of elementary vibrational motion in a lattice. It is associated with heat and sound waves in solids. 7, 8

Photodiode a semiconductor device that can convert light into electric current. 7

Photoelectron electrons produced from light (e.g. in a PMT) or X-ray (e.g. XPS and photoelectric absorption). In some figures and tables, photoelectron will be abbreviated as p.e. 9, 88, 90, 91, 98, 103, 107, 132

Photomultiplier tube a vacuum tube with a light-sensitive photocathode and several stages of that can convert light to electrical signal. The light is first converted to a photoelectron by photoelectric effect of the photocathode, and the photoelectron is amplified by secondary emission of electrons by the dynodes.
7

Phototransistor a semiconductor device that uses a photodiode to convert light into an electric current which is amplified by several stages of transistors before output. 30, 32, 34

PIC16 8-Bit microcontrollers developed by Microchip Technology. 25, 42, 126

PLL (phase-locked loop) a circuit component that generates output signals with fixed phase with respect to a input signal. It is often used to synchronize the clocks between different circuits. 130

PMT photomultiplier tube. 7, 9, 11, 12, 16, 73, 77, 85, 86, 89–103, 105–113, 115, 118, 129, 132, 133, 140, 150, 152, *Glossary*: photomultiplier tube

PPO an organic scintillator, often used as a wavelength shifter to convert UV light to longer wavelength to match spectral response of photodetectors. 100, 106, 107

Pseudocumene an organic liquid scintillator. In this thesis it is sometimes abbreviated as PC. 99–101, 105–110, 113

PTFE short for polytetrafluoroethylene. PTFE is a fluorocarbon polymer and Teflon is the most well-known trademark of PTFE. PTFE is known for chemical inertness, low coefficient of friction, high dielectric strength and good mechanical performance at cryogenic temperatures, and often finds application as lubricant, fluid and vacuum sealant, electrical insulator, component of laboratory equipment in corrosive environment. PTFE is also an excellent diffuse reflector. PTFE will cold flow over time when subject to mechanical stress, and has a maximum service temperature of 200 – 300 °C. 24, 33, 86–89, 94–97, 101, 140, 148, 149

PWM (pulse width modulation) a technique in which the width of the output pulse is changed or modulated, thereby changing the average of the signal over a period. When such signal is fed into circuits with slower response (e.g. a motor), the modulated signal is essentially averaged, thereby an analog signal is effectively created from digital signals. 126

Pyrex a brand name for a line of clear, low-thermal-expansion borosilicate glass introduced by Corning, Inc. Pyrex glass is used mainly for laboratory glassware and some kitchenware. Later the trademark (PYREX) was transferred from Corning. Today commercially PYREX mainly refers to kitchenware made from tempered soda-lime glass.

In the context of this thesis, Pyrex is used as a synonym for original borosilicate glass, as many glassblowers still do. 21, 23, 24, 29, 41

Pyrolitic carbon atomic carbon obtained by pyrolysis of carbon-rich organic molecules. 17, 19, 20, 23, 24, 27–29, 33, 35, 38, 41, 63, 66, 71, 120, 127, 128, 143, 152

Pyrolysis thermal decomposition of molecules. Pyro means fire and lysis means to separate. The process takes place at high-temperature in an inert atmosphere. 23, 25, 28

PyroMapper a device used to measure and map the thickness of pyrolitic carbon coating on fused quartz crucibles. 23, 28, 29, 31–33, 37, 120

Quartz In the usual context, quartz refers to a crystal or mineral consisting of silicon and oxygen. It is widely used as decoration and gemstone, and in electronics industry as oscillators to provide accurate timing.

In glass industry, quartz (or fused quartz) refers to amorphous quartz glass, a glass consisting of pure silicon dioxide (SiO_2). Contrary to crystalline quartz, quartz glass has extremely low thermal expansion coefficient and thermal conductivity, low electrical conductivity and high transmissivity over visible spectrum. Quartz can withstand very high temperature ($> 1200\text{ }^\circ\text{C}$) and are also chemically inert, with the exception of alkaline base and fluorine. Owing to

these properties, quartz glass is often used as crucibles for semiconductors and windows for various equipment.

In this thesis, the word quartz will be used to refer to amorphous quartz glass, and the word crystalline will be mentioned explicitly when referring to crystalline quartz. In addition, the difference between quartz and synthetic fused silica is crucial in the research described in this thesis.

Glossary: fused silica

Radiation detector In the context of this thesis, it is interpreted as a device used to measure dosage, flux, or energy spectrum of incoming ionization radiations. Non-ionizing radiation of electromagnetic spectrum is not considered. 8

Radiopurity the absence of radioactive impurities and contaminants. In the context of this thesis, low-background is strongly related to radiopurity, and high purity is used almost as a synonym of high radio-purity. 8, 9, 40

Rayleigh scattering the scattering of photons by particles. Only the direction is changed and no energy is lost. 145

RG-xxx coaxial cable standards that specifies the mechanical dimensions and electrical properties of different types of coaxial cables. 95, 101

RMD a research company based in Watertown, Massachusetts. 56, 57, 64, 66, 67, 69–71, 73, 74, 76, 77, 79

SABRE a dark matter annual modulation experiment that features a liquid veto and ultra-high purity NaI (Tl) crystals. See Section 1.3. 10–13, 15, 16, 120, 125, 129

Sandfire a manufacturer of scientific glassware based in British Columbia, Canada. 23, 63, 66, 73, 74, 79

Scanning electron microscope a microscope that uses a focused beam of electrons to probe sample surface topography and elemental compositions. 29

Scintillator a material that emits light when excited by ionizing radiation. To use a scintillator as a radiation detector, it needs to be coupled to photon detector such as photomultiplier tubes.

Although theoretically every material can scintillate to some extent, ordinary materials are opaque to its own scintillation light due to self-absorption. In many organic scintillators, the wavelength of scintillation light is shifted by vibrational states in the energy level structure of single molecules whereas in many inorganic crystal scintillators this is achieved by energy states created by activators in the middle of band gap. 7, 9–12, 14, 67, 73, 78, 83–85, 87, 96, 99, 100, 106–110, 113–115, 132, 133, 136, 139, 142, 146, 150, 152–154

Scroll pump a vacuum pump consisting of two spiral vanes (scrolls), one of which is fixed while the other moves eccentrically to compress fluid towards the center of the spiral. Scroll pumps typically have ultimate pressure below 0.1 mbar, have high pumping speed and does not require pump oil as lubricant. 23, 24, 43, 81

Seastar a manufacturer of high purity reagents based in British Columbia, Canada. 22, 68, 74, 76

SEM scanning electron microscope. 29–36, *Glossary*: scanning electron microscope

SHV Connector SHV stands for Safe High-Voltage: an RF connector used to terminate coaxial cables. Compared to most other connectors, SHV connectors have thicker and extended inner insulators to prevent accidental contact with the inner conductor at high-voltage. In this thesis, SHV connectors are mainly used to apply high-voltage bias to photomultiplier tubes. 95, 101

SiC (silicon carbide) also known as carborundum. It is an extremely hard (Mohs 9) material that is often used to manufacture abrasive and cutting tools. 96

Sigma-Aldrich a chemical, life science and biotechnology company. Most chemicals used in this research were bought from Sigma-Aldrich. 11, 39, 65, 66, 72

SMA (SubMiniature version A) a small electrical radio-frequency connector for coaxial cables. It usually has a 50Ω impedance, and the mechanical connection is maintained by threading a hex nut onto a threaded barrel. 95, 101

SPI (serial peripheral interface) a synchronous serial communication interface. It is often used for communication between a master (generates the clock signal and initiates data transfer) and multiple slaves. The interface has four lines: CLK (serial clock from master), DO (master data output), DI (master data input), and SS (slave select, used to select among multiple slave devices). 126

Standard Model a theory describing the interactions between known elementary particles. The elementary particles include six quarks (u, d, s, c, t, b), six leptons (e, μ , τ , ν_e , ν_μ , ν_τ), four gauge bosons (g, γ , Z, W) and one scalar boson (h). 3–5

Stepper motor a brush-less DC motor that divides a full rotation into a number of equal steps. The motor works by energizing toothed electromagnets to attract and align a toothed magnetic gear. 30–32, 126, 128, 129

Swagelok a company that manufactures and sells fluid systems components. Swagelok fitting (a fitting in which two ferrules bites into and seals the piping) was invented by Swagelok. 95

Synchronization the process of aligning the frequencies and phases of the clocks or the input signals from other circuits. A phase-locked loop is frequently used

for clock alignment while multiple D flip flops can be used to synchronize input signals. 114, 119, 130

TCP/IP a network communication protocol that defines the format of data structures to be exchanged. It uses handshaking and error detection algorithms to achieve a loss-less transmission with minimal error. In Unix operation systems, TCP/IP interface is managed by TCP/IP sockets. 125

Thermocouple a temperature sensor that uses the temperature-dependent voltage created by joining two wires of different materials. The small voltage produced as a function of temperature depends on the type of metals used, and usually requires amplification before it can be read. 66, 126, 127

Thread a small sequence of programs. Multi-threading refers to the concurrent execution of multiple threads from a single program or multiple programs such that multiple processes appear to be running at the same time. 122, 124

Transistor-transistor logic a logic family implemented with bipolar-junction transistors (BJTs). The logic levels, etc. are closely related to the characteristics of BJTs. 89

TTL transistor-transistor logic. 89, 113, 114, 130, *Glossary:* transistor-transistor logic

Turbomolecular pump a vacuum pump with a series of turbines. The blades on the turbine are designed such that gas molecules are efficiently and preferentially bounce away from the vacuum chamber when gas pressure is sufficiently low (molecular flow).

Although turbomolecular pump has high ultimate vacuum and high pumping speed, it is designed to operate neither at pressure above approximately 1 mbar,

nor with a pressure difference of one atmosphere. Therefore turbomolecular pump requires a forepump to pump down to rough vacuum at the beginning of operation and a backing pump to reduce the pressure at the exhaust. This can be achieved by a single pump installed at the exhaust of turbomolecular pump. 43, 63

UART (universal asynchronous receiver-transmitter) an asynchronous serial communication protocol. Physically it only consists of three lines: input, output, and ground. The two devices agree to sample the input at a particular frequency (16 times the output), and no clock signal is used (therefore asynchronous). UART is often used in the early communication between hardware and computers through the computer's serial ports. 126

VHDL a programming language used to describe low-level physical electronic circuits. VHDL stands for VHSIC Hardware Description Language, where VHSIC stands for Very High Speed Integrated Circuit. 113

Virial theorem a general equation that relates time-averaged total kinetic energy to time-averaged total potential energy in a stable system bound by conservative forces. 1

Virtual function an inheritable and overridable function in object-oriented program design. The base object defines a set of virtual functions, and the objects derived from the base objects provide the detailed implementation of these functions. When the program is executed, these virtual functions will be invoked using the protocols and formats defined by the base object, but will execute the functions implemented in the derived objects. 121, 123

VME also called VMEbus (Versa Module Europa bus). VMEbus is a computer standard that defines a set of mechanical dimensions, connectors and electrical

signalling protocols. Boards complying with this standard can be mounted on a crate and communicate with other boards on the same crate via shared electrical lines on the back-plane. 113, 116, 117, 129, 130

Watchdog timer a timer (either hardware or software) that generates a system reset or error signal if the main program fails to extend it in time. It is often used to reset devices and programs in case they hangs due to software or hardware fault. 124

Weakly Interacting Massive Particle a type of non-baryonic particle dark matter candidate with weak-scale cross section. 3

WIMP weakly interacting massive particle. 3–5, 10, *Glossary*: weakly interacting massive particle

X-ray photoelectron spectroscopy a surface analysis technique based on X-ray-induced photoelectron emission. Elemental compositions and chemical states of surface elements can be measured by comparing with characteristic binding energies of different elements. 46

XPS X-ray photoelectron spectroscopy. 46, *Glossary*: X-ray photoelectron spectroscopy

□

Bibliography

- [1] Fritz Zwicky. Die rotverschiebung von extragalaktischen nebeln. *Helvetica Physica Acta*, 6:110–127, 1933.
- [2] Sinclair Smith. The mass of the virgo cluster. *The Astrophysical Journal*, 83:23, 1936.
- [3] Vera C Rubin and W Kent Ford Jr. Rotation of the andromeda nebula from a spectroscopic survey of emission regions. *The Astrophysical Journal*, 159:379, 1970.
- [4] N Krumm and EE Salpeter. Rotation curves, mass distributions and total masses of some spiral galaxies. *Astronomy and Astrophysics*, 56:465–468, 1977.
- [5] Mordehai Milgrom. M. milgrom, *astrophys. j.* 270, 365 (1983). *Astrophys. J.*, 270:365, 1983.
- [6] Douglas Clowe, Maruša Bradač, Anthony H Gonzalez, Maxim Markevitch, Scott W Randall, Christine Jones, and Dennis Zaritsky. A direct empirical proof of the existence of dark matter. *The Astrophysical Journal Letters*, 648(2):L109, 2006.
- [7] Peter AR Ade, N Aghanim, M Arnaud, M Ashdown, J Aumont, C Baccigalupi, AJ Banday, RB Barreiro, JG Bartlett, N Bartolo, et al. Planck 2015 results-xiii. cosmological parameters. *Astronomy & Astrophysics*, 594:A13, 2016.
- [8] Charles Alcock, RA Allsman, David R Alves, TS Axelrod, Andrew C Becker, DP Bennett, Kem H Cook, N Dalal, Andrew John Drake, KC Freeman, et al. The macho project: microlensing results from 5.7 years of large magellanic cloud observations. *The Astrophysical Journal*, 542(1):281, 2000.
- [9] Patrick Tisserand, L Le Guillou, C Afonso, JN Albert, J Andersen, R Ansari, É Aubourg, P Bareyre, JP Beaulieu, X Charlot, et al. Limits on the macho content of the galactic halo from the eros-2 survey of the magellanic clouds. *Astronomy & Astrophysics*, 469(2):387–404, 2007.
- [10] Esra Bulbul, Maxim Markevitch, Adam Foster, Randall K Smith, Michael Loewenstein, and Scott W Randall. Detection of an unidentified emission line in the stacked x-ray spectrum of galaxy clusters. *The Astrophysical Journal*, 789(1):13, 2014.

- [11] Alexey Boyarsky, Oleg Ruchayskiy, Dmytro Iakubovskiy, and Jeroen Franse. Unidentified line in x-ray spectra of the andromeda galaxy and perseus galaxy cluster. *Physical review letters*, 113(25):251301, 2014.
- [12] RD Peccei et al. Rd peccei and hr quinn, phys. rev. lett. 38, 1440 (1977). *Phys. Rev. Lett.*, 38:1440, 1977.
- [13] Stephen J Asztalos, G Carosi, C Hagmann, D Kinion, K Van Bibber, M Hotz, LJ Rosenberg, G Rybka, J Hoskins, J Hwang, et al. Squid-based microwave cavity search for dark-matter axions. *Physical review letters*, 104(4):041301, 2010.
- [14] JI Read. The local dark matter density. *Journal of Physics G: Nuclear and Particle Physics*, 41(6):063101, 2014.
- [15] Peter F Smith and John David Lewin. Dark matter detection. *Physics Reports*, 187(5):203–280, 1990.
- [16] Martin C Smith, Gregory R Ruchti, Amina Helmi, Rosemary FG Wyse, Jon P Fulbright, Kenneth C Freeman, Julio F Navarro, George M Seabroke, Matthias Steinmetz, Mary Williams, et al. The rave survey: constraining the local galactic escape speed. *Monthly Notices of the Royal Astronomical Society*, 379(2):755–772, 2007.
- [17] Andrzej K Drukier, Katherine Freese, and David N Spergel. Detecting cold dark-matter candidates. *Physical Review D*, 33(12):3495, 1986.
- [18] G Bellini, J Benziger, D Bick, G Bonfini, D Bravo, M Buizza Avanzini, B Caccianiga, L Cadonati, F Calaprice, C Carraro, et al. Cosmic-muon flux and annual modulation in borexino at 3800 m water-equivalent depth. *Journal of Cosmology and Astroparticle Physics*, 2012(05):015, 2012.
- [19] Gerard Jungman, Marc Kamionkowski, and Kim Griest. Supersymmetric dark matter. *Physics Reports*, 267(5-6):195–373, 1996.
- [20] Keith A Olive, Particle Data Group, et al. Review of particle physics. *Chinese Physics C*, 38(9):090001, 2014.
- [21] R Bernabei, P Belli, A Bussolotti, F Cappella, R Cerulli, CJ Dai, A dAngelo, HL He, A Incicchitti, HH Kuang, et al. The dama/libra apparatus. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 592(3):297–315, 2008.
- [22] Rita Bernabei, P Belli, F Cappella, V Caracciolo, S Castellano, R Cerulli, CJ Dai, A dAngelo, S dAngelo, A Di Marco, et al. Final model independent result of dama/libra–phase1. *The European Physical Journal C*, 73(12):2648, 2013.

- [23] R Bernabei, P Belli, F Cappella, R Cerulli, CJ Dai, A d'Angelo, HL He, A Incicchitti, HH Kuang, JM Ma, et al. First results from dama/libra and the combined results with dama/nai. *The European Physical Journal C*, 56(3):333–355, 2008.
- [24] R Bernabei, P Belli, F Cappella, V Caracciolo, R Cerulli, CJ Dai, A d'Angelo, A Di Marco, HL He, A Incicchitti, et al. No role for muons in the dama annual modulation results. *The European Physical Journal C*, 72(7):2064, 2012.
- [25] E Aprile, J Aalbers, F Agostini, M Alfonsi, FD Amaro, M Anthony, F Arneodo, P Barrow, L Baudis, Boris Bauermeister, et al. First dark matter search results from the xenon1t experiment. *Physical review letters*, 119(18):181301, 2017.
- [26] DS Akerib, S Alsum, HM Araújo, X Bai, AJ Bailey, J Balajthy, P Beltrame, EP Bernard, A Bernstein, TP Biesiadzinski, et al. Results from a search for dark matter in the complete lux exposure. *Physical review letters*, 118(2):021303, 2017.
- [27] Xiangyi Cui, Abdusalam Abdukerim, Wei Chen, Xun Chen, Yunhua Chen, Binbin Dong, Deqing Fang, Changbo Fu, Karl Giboni, Franco Giuliani, et al. Dark matter results from 54-ton-day exposure of pandax-ii experiment. *Physical review letters*, 119(18):181302, 2017.
- [28] Christopher Savage, Graciela Gelmini, Paolo Gondolo, and Katherine Freese. Compatibility of dama/libra dark matter detection with other searches. *Journal of Cosmology and Astroparticle Physics*, 2009(04):010, 2009.
- [29] R Foot. Mirror dark matter and the new dama/libra results: a simple explanation for a beautiful experiment. *Physical Review D*, 78(4):043529, 2008.
- [30] Emily Kathryn Shields. *SABRE: A search for dark matter and a test of the DAMA/LIBRA annual-modulation result using thallium-doped sodium-iodide scintillation detectors*. PhD thesis, Princeton University, 2015.
- [31] Momentive. Chemical composition of fused quartz. <https://www.momentive.com/en-US/categories/quartz/chemical-composition/#>.
- [32] Heraeus. Suprasil 310—premium synthetic quartz glass tubes for deep uv applications. https://www.heraeus.com/en/hqs/fused_silica_quartz_knowledge_base/properties/properties.aspx.
- [33] John Verhoogen. Ionic diffusion and electrical conductivity in quartz. *American Mineralogist: Journal of Earth and Planetary Materials*, 37(7-8):637–655, 1952.
- [34] Mark J Harrison, Adam P Graebner, Walter J McNeil, and Douglas S McGregor. Carbon coating of fused silica ampoules. *Journal of crystal growth*, 290(2):597–601, 2006.
- [35] K Erents and G Carter. The evolution of water vapour from glass after atmospheric exposure. *Vacuum*, 15(12):573–575, 1965.

- [36] T Akermark, G Hultquist, and Q Lu. Release of water and hydrogen during outgassing of some materials. *Journal of materials engineering and performance*, 5(4):516–520, 1996.
- [37] Galyna I Dovbeshko, Volodymyr R Romanyuk, Denys V Pidgirnyi, Vsevolod V Cherepanov, Eugene O Andreev, Vadim M Levin, Polina P Kuzhir, Tommi Kaplas, and Yuri P Svirko. Optical properties of pyrolytic carbon films versus graphite and graphene. *Nanoscale research letters*, 10(1):234, 2015.
- [38] J Simon, S Gál, and Erdey L. Thermal reactions of sodium halides. *Acta Chimica Academiae Scientiarum Hungaricae*, 53:105, 1967.
- [39] Dmitriy S Sofronov, Konstantin A Kudin, A Yu Voloshko, Alexander M Kudin, and Oleg V Shishkin. Origin of the thermal desorption peaks of gases in nai above 180 c. *Inorganic Materials*, 45(11):1314, 2009.
- [40] William G Pfann. Principles of zone-melting. *JOM*, 4(7):747–753, 1952.
- [41] William G Pfann. *Zone melting*. John Wiley and sons, 1966.
- [42] Ulrich Gross. Zone purification of alkali iodides. *Materials Research Bulletin*, 5(2):117–127, 1970.
- [43] CW Bale, E Bélisle, P Chartrand, SA Decterov, G Eriksson, K Hack, I-H Jung, Y-B Kang, J Melançon, AD Pelton, et al. Factsage thermochemical software and databases recent developments. *Calphad*, 33(2):295–311, 2009.
- [44] RJM Konings, A Kok-Scheele, and EHP Cordfunke. On the phase diagrams of the systems pb pbi₂, pbi₂ nai and pbi₂ zni₂. *Thermochimica acta*, 261:221–225, 1995.
- [45] V.A. Isaenko, A.N. Kirgintsev, and N.N. Smirnov. Ravnovesnye koehffitsienty raspredeleniya tli pri napravlennoj kristallizatsii nai. *Izvestiya Sibirskogo Otdeleniya Akademii Nauk SSSR, Seriya Khimicheskikh Nauk*, 4:146, 1977.
- [46] Special Metals. Inconel alloy 625. <http://www.specialmetals.com/assets/smc/documents/alloys/inconel/inconel-alloy-625.pdf>.
- [47] J Eckstein, U Gross, and E Rubinová. Growth of alkali halide crystals with a lowered concentration of oxygenous anions using six4 as a removing agent. *Kristall und Technik*, 3(4):583–587, 1968.
- [48] Glenn F Knoll. *Radiation detection and measurement*. John Wiley & Sons, 2010.
- [49] Kanthal. Nichrothal 80 resistance heating wire and resistance wire. <https://www.kanthal.com/en/products/material-datasheets/wire/resistance-heating-wire-and-resistance-wire/nikrothal-80/>.

- [50] Kanthal. Kanthal a-1 resistance heating wire and resistance wire. <https://www.kanthal.com/en/products/material-datasheets/wire/resistance-heating-wire-and-resistance-wire/nikrothal-80/>.
- [51] American Iron and Steel Institute. *High Temperature Characteristics of Stainless Steels. A Designers Handbook Series No. 9004*. American Iron and Steel Institute, 2011.
- [52] J Amaré, S Cebrián, C Cuesta, E García, C Ginestra, M Martínez, MA Oliván, Y Ortigoza, A Ortiz de Solórzano, C Pobes, et al. Preliminary results of anais-25. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 742:187–190, 2014.
- [53] P Adhikari, G Adhikari, E Barbosa de Souza, N Carlin, S Choi, WQ Choi, M Djamal, AC Ezeribe, C Ha, IS Hahn, et al. Background model for the nai (tl) crystals in cosine-100. *The European Physical Journal C*, 78(6):490, 2018.
- [54] Jiang Wen-Qi, Gu Shu-Di, John Joseph, Liu Da-Wei, Kam-Biu Luk, Herbert Steiner, Wang Zheng, and Wu Qun. Suppressing ringing caused by large photomultiplier tube signals. *Chinese Physics C*, 36(3):235, 2012.
- [55] PG Catalano. Pg catalano et al., mem. soc. geol. it. 35, 647 (1986). *Mem. Soc. Geol. It.*, 35:647, 1986.
- [56] P Agnes, L Agostino, IFM Albuquerque, T Alexander, AK Alton, K Arisaka, HO Back, B Baldin, K Biery, G Bonfini, et al. The veto system of the darkside-50 experiment. *Journal of Instrumentation*, 11(03):P03016, 2016.
- [57] Michael Clinton Johnson. *Scintillator purification and study of light propagation in a large liquid scintillation detector*. PhD thesis, Princeton University, 1998.
- [58] DN Abdurashitov, Yu M Malyshkin, VL Matushko, and B Suerfu. Response of a proportional counter to 37ar and 71ge: Measured spectra versus geant4 simulation. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 373:5–9, 2016.
- [59] Oskar Klein and Yoshio Nishina. Über die streuung von strahlung durch freie elektronen nach der neuen relativistischen quantendynamik von dirac. *Zeitschrift für Physik*, 52(11-12):853–868, 1929.
- [60] John Cameron, Jun Chen, Balraj Singh, and Ninel Nica. Nuclear data sheets for a= 37. *Nuclear Data Sheets*, 113(2):365–514, 2012.
- [61] Bernhard E Lehmann, Stanley N Davis, and June T Fabryka-Martin. Atmospheric and subsurface sources of stable and radioactive nuclides used for ground-water dating. *Water Resources Research*, 29(7):2027–2040, 1993.

- [62] M Palomo, A Penalver, C Aguilar, and F Borrull. Tritium activity levels in environmental water samples from different origins. *Applied radiation and isotopes*, 65(9):1048–1056, 2007.
- [63] IAEA/WMO. The gnip database - global network of isotopes in precipitation. <http://www.iaea.org/water>.
- [64] M Hadžišehović, N Miljević, D Golobočamn, M Župančić, V Šipka, and S Kudia. Tritium abundance in belgrade natural waters. *Isotopenpraxis Isotopes in Environmental and Health Studies*, 26(1):5–10, 1990.
- [65] Illinois Environmental Protection Agency. Use of tritium in assessing aquifer vulnerability. <http://www.iaea.org/water>.
- [66] R-R. Edwards. Iodine-129: its occurrence in nature and its utility as a tracer. *Science*, 137(3533):851–853, 1962.
- [67] J Fabryka-Martin, H Bentley, D Elmore, and PL Airey. Natural iodine-129 as an environmental tracer. *Geochimica et Cosmochimica Acta*, 49(2):337–347, 1985.
- [68] U Fehn, GR Holdren, D Elmore, T Brunelle, R Teng, and PW Kubik. Determination of natural and anthropogenic ^{129}I in marine sediments. *Geophysical Research Letters*, 13(2):137–139, 1986.
- [69] C Stan-Sion, M Enachescu, and AR Petre. Ams analyses of ^{129}I from the fukushima daiichi nuclear accident in the pacific ocean waters of the coast la jolla–san diego, usa. *Environmental Science: Processes & Impacts*, 17(5):932–938, 2015.
- [70] M Antonello, E Barberio, T Baroncelli, J Benziger, LJ Bignell, I Bolognino, F Calaprice, S Copello, D D’Angelo, G D’Imperio, B. Suerfu, et al. Monte carlo simulation of the sabre pop background. *arXiv preprint arXiv:1806.09344*, 2018.