Commissioning the DIRC Detector and Searching for Axion-like Particles at GlueX

by

Yunjie Yang

B.S., University of Michigan (2015)

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2021

© Massachusetts Institute of Technology 2021. All rights reserved.

Author .....
Department of Physics
May 14, 2021

Certified by .....
Mike Williams
Associate Professor
Thesis Supervisor

Accepted by .....
Deepto Chakrabarty
Associate Department Head

Commissioning the DIRC Detector and Searching for Axion-like Particles at GlueX

by

Yunjie Yang

Submitted to the Department of Physics
on May 14, 2021, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Physics

Abstract

This thesis centers around problems in the study of the strong nuclear force. The GlueX DIRC, a Cherenkov radiation-based detector, was proposed to upgrade the particle identification capability of the GlueX experiment, which aims to perform quantitative tests of Quantum Chromodynamics in the nonperturbative regime by searching for and studying hybrid mesons. This thesis describes the construction, commissioning, reconstruction, and calibration of the GlueX DIRC detector.

Originally proposed to solve the strong CP problem, axions and axion-like particles are hypothetical pseudoscalar particles found in many proposed extensions to the Standard Model of particle physics. This thesis presents a search for photoproduction of axion-like particles using data in photon-proton interactions collected by the GlueX experiment at Jefferson Laboratory in the $\gamma\gamma$ and $\pi^+\pi^-\pi^0$ final states of the axion-like particles.

In addition, the Monte Carlo modeling of the strong interaction at low energies leads to challenges known as the event generator tuning problem. This thesis presents a novel approach to the Monte Carlo event generator tuning problem using Bayesian optimization.

Thesis Supervisor: Mike Williams
Title: Associate Professor

Acknowledgments

First and foremost, I would like to thank my advisor, Mike Williams. Mike reached out to me when I was just starting at MIT. Immediately, I knew I wanted to work with Mike because of his passion and his constant flow of interesting ideas, which have been inspiring me ever since. Mike has taught me how to think like a physicist, and has shown me by example how to be a successful academic. Mike provided me with great flexibility and freedom regarding my projects, yet he was always available for guidance and support. Acknowledgment is also due to my other committee members Richard Milner and Xiao-Gang Wen for their feedback and patience with this thesis.

My GlueX collaborators, especially members of the DIRC team, also deserve acknowledgments. I would like to thank Matt Shepherd for making the DIRC bar box transports successful and fun. Thanks to Justin Stevens for his leadership on the DIRC project and for sharing with me his knowledge about GlueX detector and software. To Maria Patsyuk for teaching me how to make silicone cookies. To Jochen Schwiening and Roman Dzhygadlo for imparting to me a small fraction of their knowledge about the DIRC. To Cristiano Fanelli for bouncing around ideas about analyses.

I also would like to thank everyone who made LNS and MIT Physics supportive communities. Thanks to Constantin Weisser for being an inspiration from all his achievements and for being a caring friend who is always there. To John Hardin for the conversations about and beyond physics. To Rey Cruz Torres for making me want to go to Newport News other than for physics. To Tom Boettcher, Sangbaek Lee, Bobby Johnston, as well as the other past and current penthouse members, for the fun procrastination. To Field Rogers for always letting me know that she liked my talks. To Miao Hu for the delicious food we explored together. To my classmates at MIT, Taweewat Somboonpanyakul, Alfred Zong, Dahlia Klein, Huy Phan, and Ryuji Takagi, for making grad school life fun. To Cathy Modica, Sydney Miller, and Lauren Saragosa for minimizing all the non-physics hassle of which I had to take care.

I am indebted to Dan Amidei and Ryan Edgar during my undergraduate years at the University of Michigan. Dan took me as a research assistant, introduced me to the field of experimental particle physics, and encouraged me to pursue grad school. Ryan was always there to answer all my newbie questions about physics and ROOT. I would also like to thank my study partner Youngshin Kim for the company in surviving those upper-level physics classes together.

Last but not least, I would like to acknowledge my family and my friends outside of MIT. Thanks to my family, especially my mother Xiaoying Yang, for the unconditional love and support. To Boxiao Ma for being a trusted friend who I know I can turn to no matter what happens. To Xichao She for the friendship and the reflections as we navigate our lives. Lastly, I would like to thank my best friend, Qiuhan Li, for offering different perspectives, for being there with me in my good and bad times, and for making me a better individual.

Ultimately, PhD is a personal journey, but I am deeply grateful for everyone with whom I shared parts of this adventure.

Contents

Chapter 1

Introduction and Overview

The human race has always been seeking to understand the fundamental building blocks of the universe since the very beginning of its history. The quest is never-ending but at present we have come to understand that there are four fundamental forces of nature: electromagnetism, gravity, the weak, and the strong nuclear forces. With the exception of gravity, the other three forces can be elegantly described by a unified Quantum Field Theory (QFT) framework known as the Standard Model (SM) of particle physics, arguably *the* most successful scientific theory of all time. Despite its tremendous success, the Standard Model is far from either perfect or all-encompassing. The understanding of the Standard Model itself and the phenomena, such as dark matter, that cannot be explained by the Standard Model are active areas of modern physics research. The overarching theme of this thesis centers around the efforts that either aim to better understand the Standard Model, in particular the strong interaction, or to search for phenomena beyond the Standard Model (BSM).

1.1 Quantum Chromodynamics

Quantum Chromodynamics, or QCD, is believed to be the theory for the strong nuclear force, or the strong interaction. QCD describes the interactions between *quarks* and *gluons*. Quarks and gluons carry the so-called *color* charges (hence the name, Quantum *Chromo*dynamics), and there are three types of color charges, dubbed *red*, *blue*, and *green*. In principle, all the dynamics of QCD are encoded in the QCD Lagrangian

$$\mathcal{L} = \sum_q \bar{\psi}_{q,a} \left(i\gamma^\mu \partial_\mu \delta_{ab} - g_s \gamma^\mu t^C_{ab} \mathcal{A}^C_\mu - m_q \delta_{ab} \right) \psi_{q,b} - \frac{1}{4} G^A_{\mu\nu} G^{A\,\mu\nu}, \tag{1.1}$$

where repeated indices are summed over following Einstein convention. We will now unpack each of the terms in Eq. (1.1). The γ^μ are the Dirac matrices, and the $\delta_{a,b}$ is the Kronecker delta. The $\psi_{q,a}$ are quark-field spinors (the $\bar{\psi} \equiv \psi^\dagger \gamma^0$ are the adjoint spinors) for a quark of *flavor* q and mass m_q with a color-index a that runs from 1 to 3 (which is the number of color charges). In the Standard Model, there are six flavors of quarks in three *generations*: (up, down), (charm, strange), and (top, bottom). The

$\mathcal{A}_\mu^C$ are the gluon fields with the field index C running from 1 to 8 ($= 3^2 - 1$) because there are eight kinds of gluons. The gluon field strength tensor is related to the gluon fields by

$$G_{\mu\nu}^A = \partial_\mu \mathcal{A}_\nu^A - \partial_\nu \mathcal{A}_\mu^A - g_s f_{ABC} \mathcal{A}_\mu^B \mathcal{A}_\nu^C. \tag{1.2}$$

The quantity g_s is the QCD coupling constant and is related to α_S as $\alpha_S = g_s^2/4\pi$. It is the only true fundamental parameter of QCD (the quark masses have an electroweak origin through the Higgs mechanism). Finally, the t_{ab}^C are the eight 3×3 matrices that are the generators of the $SU(3)$ group, and f_{ABC} are the structure constants of the $SU(3)$ group.

There are four types of terms in Eq. (1.1) (omitting indices and numerical factors): (1) $\bar{\psi}\partial\psi$ is the kinetic term of the quark fields; (2) $m\bar{\psi}\psi$ is the mass term of the quark fields; (3) $g_s\bar{\psi}\mathcal{A}\psi$ describes the interaction between the quark fields and the gluon fields, corresponding to the quark-antiquark-gluon vertex, $\bar{q}qg$; and (4) $G_{\mu\nu}G^{\mu\nu}$ encodes the self-interactions among gluon fields, in particular, there is a 3-gluon vertex (proportional to g_s) and a 4-gluon vertex (proportional to g_s^2). This is an important distinguishing feature of QCD, which is a non-abelian gauge theory, compared to Quantum Electrodynamics (QED), which does not have photon self-interaction due to its abelian $U(1)$ group structure. This non-abelian nature of QCD is also believed to be the origin of the phenomenon known as color *confinement* – that is colored particles, *i.e.,* quarks and gluons, cannot exist in isolation below the confinement scale Λ_{QCD} of about 150 MeV, and they must form color-singlet particles, known as *hadrons*. Qualitatively, the confinement phenomenon can be understood in the following way. Unlike QED where the interaction strength becomes weaker and weaker as two electrically charged particles move apart from each other, the interaction strength between two color charged particles (such as a quark-antiquark pair) becomes stronger and stronger because of the additional self-interactions between the gluon fields connecting the two colored particles, which is absent in QED. At some point where the energy stored in the color "flux tube" between the two colored particles becomes so large that it is energetically more favorable to create additional colored particles out of vacuum to form colorless particles with the original colored particles. The description of the transition from quarks and gluons, produced in high energy collisions and described in the QCD Lagrangian, to color-neutral hadrons that are observable in experimental apparatus, is a challenging problem and its realization in Monte Carlo event generators is the subject matter of Chapter 3.

We see that a lot of the QCD dynamics is dictated by the $SU(3)$ group structure at both the quark-gluon vertex and the gluon self-interaction vertices. The interaction strength is dictated by the QCD coupling constant, whose running as a function of energy scale, shown in Fig. 1-1, is a determining factor of the QCD phenomenology in experiments. The most prominent feature of the running of of the strong coupling constant α_S is that it decreases with the energy scale. That is, at high energies, the coupling is weak, and the perturbative expansion in orders of the coupling is valid – this is the *perturbative* QCD regime (pQCD). On the other hand, the coupling becomes very large at lower energies and the perturbative expansion in orders of the coupling is no longer valid, entering the *nonperturbative* regime of QCD. Currently,

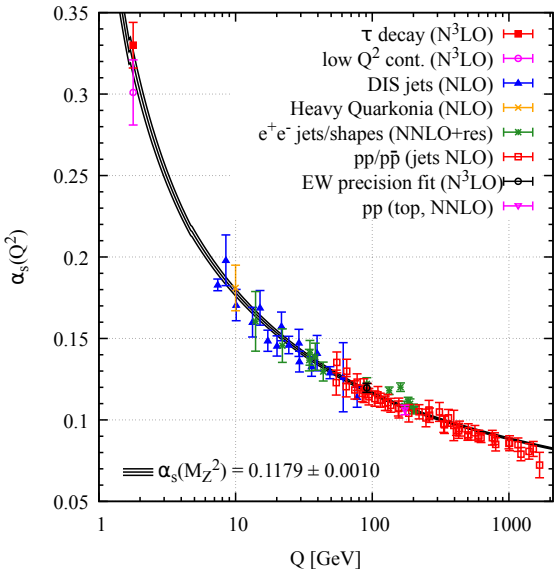

Figure 1-1: The running of the strong coupling constant α_S as a function of energy scale Q. Image source: [1].

lattice QCD (LQCD) is the only viable *ab initio* predictive method for QCD in the nonperturbative regime. There also exist other successful model frameworks that describe some aspects of low energy QCD phenomena, such as the quark model, chiral perturbation theory *etc.*.

1.2 Hadron Spectroscopy and GlueX

1.2.1 Quantitative Tests of QCD

At high energies, the strong coupling constant is small and a perturbative framework is used to make precise predictions from first principles QCD. *Factorization* techniques are used to separate the perturbative high energy phenomena from the nonperturbative low energy phenomena. Such predictions can then be compared against experimental measurements, allowing for quantitative tests of QCD at these energies. Examples of such measurements are cross sections and hadronic final-state observables, such as jets, event shapes, and jet substructures. They can be made at high energy colliders or deep inelastic scattering fixed-target experiments. Tremendous progress has been made in both the experimental and theoretical fronts over the years and together QCD has been validated to incredible precision at these high energies [1], yet another example of the success of the Standard Model.

On the other hand, the quantitative testing of QCD in the nonperturbative regime remains challenging. The large coupling strength and the phenomenon of confinement at these low energies render the perturbative toolkit invalid. However, recent advances in lattice gauge theory have enabled quantitative predictions of QCD from first principles. On the experimental side, the study of hadron spectroscopy provides

a powerful tool to provide quantitative tests of QCD in the nonperturbative regime.

1.2.2 Hadron Spectroscopy

The scientific history of spectroscopy began, arguably, with Issac Newton's optics experiment. Newton is also believed to be the first person to apply the word *spectrum* to describe the rainbow color bands of the sunlight revealed by a prism. The use of spectroscopy as a quantitative technique began with Joseph von Fraunhofer in the early 1800s. Fraunhofer replaced a prism with a diffraction grating as the source of wavelength dispersion and made quantitative observations of the solar spectrum in the optical band (the dark absorption lines he observed and characterized are still known today as the Fraunhofer lines). The use of spectroscopy to understand the underlying interaction of a bound state is best exemplified by the study of atomic spectroscopy, which laid the foundation of quantum mechanics, and ultimately QED, and has led to multiple Nobel Prizes.

In our current understanding, hadrons are color-neutral bound states of the colored particles of quarks and gluons. The quarks and gluons inside a hadron are bound by the strong interaction. Therefore, if QCD is the correct theory for the strong interaction also in this energy regime, the spectrum of hadrons, that is the pattern of ground states and excited states, should match the predictions from first principles QCD. The study of hadron spectroscopy provides a quantitative tool to test QCD in the nonperturbative regime. As mentioned, the only viable *ab initio* predictive method for QCD in the nonperturbative regime is the lattice gauge theory, or lattice QCD. The idea of lattice QCD is to discretize the four-dimensional spacetime into a lattice where quarks are put on the vertices and gluons are the linking edges of the lattice. Then the integration of the action, which contains the QCD Lagrangian, for various operators is performed over this spacetime lattice for different configurations of the quark and gluon fields. The discretization allows such integrals to be performed numerically on computers, often harnessing the power of modern supercomputers. Physical observables, such as the spectrum of hadron bound states, are then extracted. As the lattice spacing decreases and the size of the lattice box over which the computations are carried out increases, the lattice QCD result is expected to approximate the continuum QCD. The hadrons are characterized by their masses and quantum numbers, including the total angular momentum J, parity P, and charge conjugation C, written in spectroscopic notation as J^{PC}. Experimentally, the quantum numbers of those hadron states can be measured via the angular distributions of their decay products. Therefore, the spectrum of hadrons can be mapped out experimentally and compared to theoretical predictions, allowing for quantitative tests of the theory.

1.2.3 Hybrid Mesons and GlueX

Mesons can be generally defined as hadrons with integer spins and with 0 baryon quantum number. In the *constituent quark model* proposed by Gell-Mann and Zweig,

even before the formulation of QCD, a meson is a quantum mechanical system consisting of a quark-antiquark pair. Quarks and antiquarks are spin-1/2 fermions obeying $SU(2)_{\text{spin}}$ symmetry and the system can also have orbital angular momentum between the quark-antiquark pair with $SO(3)_{\text{orbital}}$ symmetry. In addition, in the light meson sector, which are mesons only made of up, down, and strange quarks, there also exists an (approximate) $SU(3)_{\text{flavor}}$ flavor symmetry (extended from the $SU(2)_{\text{isospin}}$) due to their approximately equal masses. Therefore, the light meson system can be characterized by

$$SU(3)_{\text{flavor}} \otimes SU(2)_{\text{spin}} \otimes SO(3)_{\text{orbital}}. \tag{1.3}$$

The angular momentum of the system follows the usual rules of adding angular momenta. The total intrinsic spin is the sum of the two spin-1/2 fermions: $S = 0, 1$. The orbital angular momentum can take integer values start from 0: $L = 0, 1, 2, ...$ The total angular momentum is then the vector sum of the total intrinsic spin and orbital angular momentum: $\vec{J} = \vec{S} + \vec{L}$, and can take values: $|L - S| \leq J \leq (L + S)$. In the quark model, the property of a quark-antiquark system can additionally be specified by its behavior under parity and charge conjugation operations, resulting in P and C quantum numbers specified in the following way:

$$P = (-1)^{L+1},\ C = (-1)^{L+S}, \tag{1.4}$$

giving rise to meson states with the following quantum numbers:[1]

$$J^{PC} = 0^{-+},\ 0^{++},\ 1^{--},\ 1^{+-},\ 1^{++},\ 2^{--},\ 2^{-+},\ 2^{++},\ \ldots\ . \tag{1.5}$$

For each J^P combination, there also exists a meson nonet by combining the $SU(3)_{\text{flavor}}$ quark triplet with the antiquark triplet: $3 \otimes \bar{3} = 8 \oplus 1$, resulting in nine possible combinations with one octet and a singlet in the $SU(3)_{\text{flavor}}$ space. The nonet can be organized in a very nice pattern according to the *strangeness* quantum number, the electric charge, and the z-component of the isospin as shown in Fig. 1-2 for the ground-state pseudoscalar ($J^P = 0^-$) nonet and vector nonet ($J^P = 1^-$).

Despite its simplicity, the constituent quark model has been incredibly successful in describing the phenomenology of meson systems. Fig. 1-3 shows the light meson spectrum calculated in a relativistic quark model compared with experimental observations. Recent advances in lattice QCD also allowed calculations of the light meson spectrum as shown in Fig. 1-4 that can be compared with experimental observations (although the input pion mass was made unphysically large in these specific calculations to make the computation tractable).

However, we know the quark model is incomplete because the gluons, which bind the quarks inside a hadron according to QCD, are absent in this picture. In fact, it has been a long-standing goal of the hadronic physics community to understand how the quark and gluonic degrees of freedom present in the QCD Lagrangian manifest themselves in the spectrum of hadrons. In general, mesons that do not fit the simple quark-antiquark model are called *exotics*. In particular, quark-antiquark states

[1]The concept of *naturality* of a state is sometimes used to categorize them. A state is said to have *natural parity* if $P = (-1)^J$ and *unnatural parity* if $P = (-1)^{J+1}$.

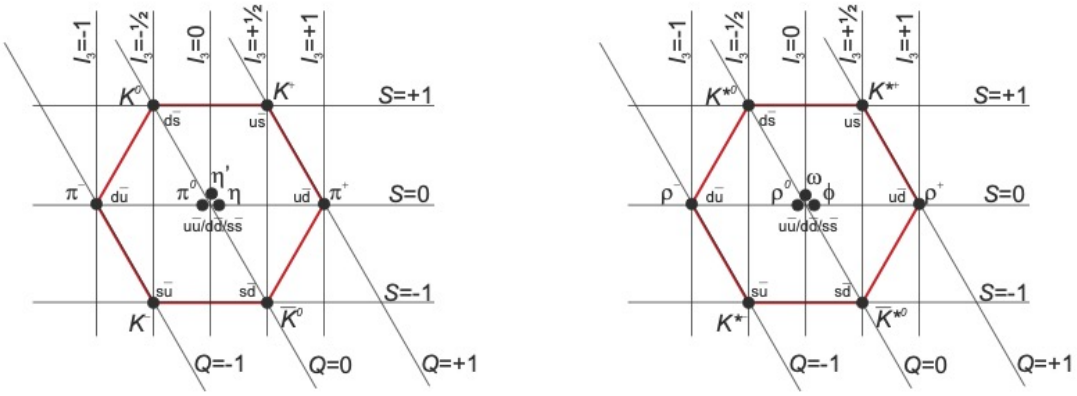

Figure 1-2: Light meson nonets of ground-state (left) pseudoscalars and (right) vectors. Image source: [2].

with an excited gluon field are called *hybrids*, which can be naively thought of as a quark-antiquark pair coupled with a valence gluon, *i.e.,* a $q\bar{q}g$ configuration. With the extra degree of freedom from the valence gluon, one could expect states with quantum numbers that cannot be realized with only a quark-antiquark configuration. Specifically, mesons with J^{PC} quantum numbers, such as

$$J^{PC} = 0^{--},\ 0^{+-},\ 1^{-+},\ 2^{+-}, ..., \tag{1.6}$$

that do not fit the description of Eq. (1.5) are called *spin-exotics*. Recent lattice QCD calculations [3] suggest that such states may very well exist because they naturally fall out of the QCD Lagrangian in the calculation, as shown in the last three column in Fig. 1-4. In addition, recent interpretation of the COMPASS results [4] provided strong evidence for the existence of the $\pi_1(1400)$ and $\pi_1(1600)$ states as the lightest exotic meson candidates which are classified by PDG [1] as spin-exotics as shown in Fig. 1-3. The primary goal of the GlueX experiment, described in Chapter 2, is to search for and study the spectrum of the hybrid mesons.

1.2.4 The Need for DIRC at GlueX

While the recent COMPASS results on the 1^{-+} $\pi_1(1400)$ and $\pi_1(1600)$ states are very encouraging, it is critically important to map out the *spectrum* of the exotics to make more definitive statements. To achieve this goal, it is required that we map out the *nonets* of the exotic J^{PC} hybrids as we did for the ordinary mesons as shown in Fig. 1-2. This necessarily involves identifying the members in the hybrid nonets with a significant strangeness content, which decay preferentially to strange-quark–containing final state particles, in particular the $K^{\pm}$ mesons.

The lowest laying exotic J^{PC} hybrids are shown in the 1^{-+} column in Fig. 1-4. Lattice calculations [3] suggest that these states have an S-wave ($L = 0$) $q\bar{q}$ pair in a color octet configuration coupled to an excited gluonic field in a color octet with effective P-wave quantum number $J_g^{P_g C_g} = 1^{+-}$. In this nonet, one can see that

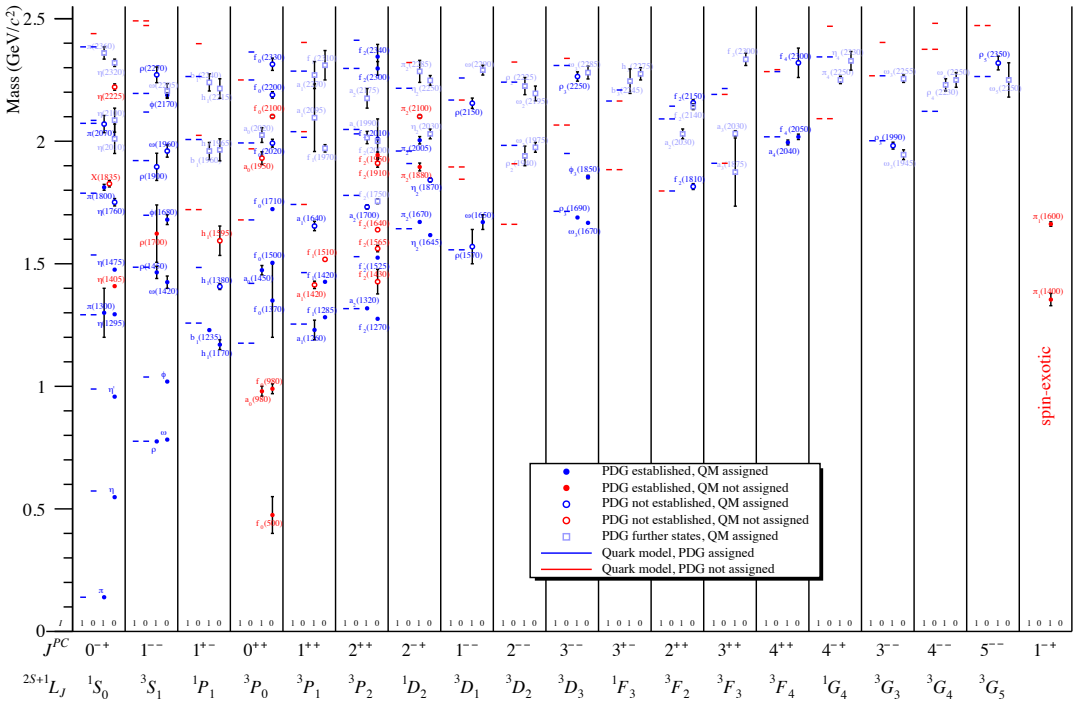

Figure 1-3: Status of the light quark meson spectrum. Image source: [2].

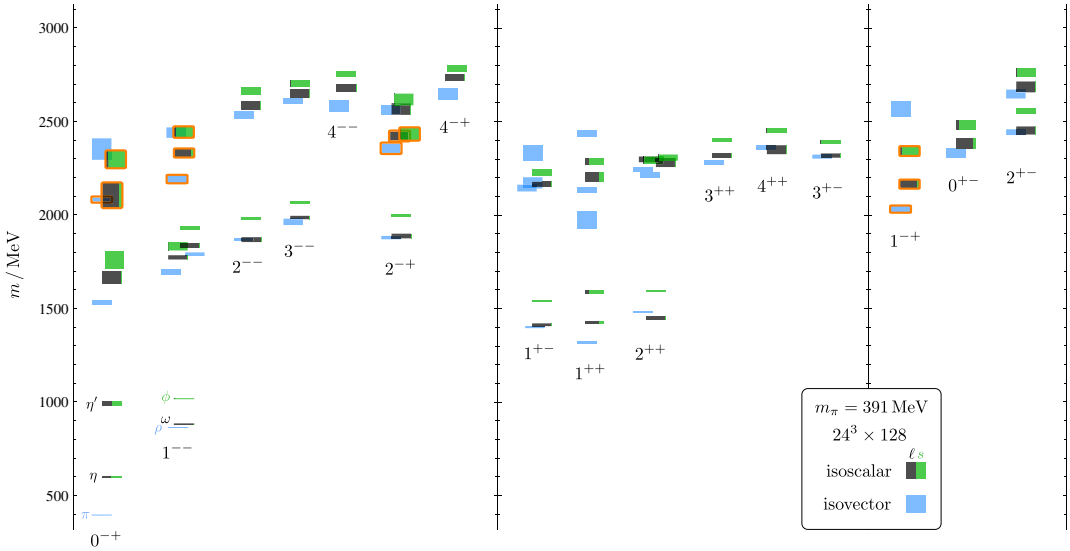

Figure 1-4: Lattice QCD prediction of the light quark meson spectrum. The computation is carried out with two flavors of quark, light (l) and strange (s). Image source: [3].

Multiplet	Approximate Mass (MeV)	Total Width (MeV)	Final States
π_1	1900	80 – 170	$\omega\pi\pi$, 3π, 5π, $\eta 3\pi$, $\eta'\pi$
η_1	2100	60 – 160	4π, $\eta 4\pi$, $\eta\eta\pi\pi$
η_1'	2300	100 – 220	$KK\pi\pi$, $KK\pi$, $KK\omega$

Table 1.1: Lowest laying exotic 1^{-+} hybrid nonet predicted from lattice QCD. Table adapted from [5].

one of the isoscalar multiplets has a significant strange flavor content. Table 1.1 shows the predicted masses, total widths, and decay final states of the lowest laying exotic 1^{-+} hybrid nonet. As one can clearly see, to map out the η_1' multiplet, it necessarily involves identification of strange quark containing final state particles. The ability to provide clear kaon identification is therefore critical to the success of the GlueX physics program. A dedicated particle identification detector based on the decommissioned BaBar DIRC components was proposed and designed to achieve this goal [5]. The commissioning effort of the GlueX DIRC detector is the subject matter of Chapter 4.

1.3 The Strong CP Problem and Axions

1.3.1 The Strong CP Problem and the Axion Solution

The Strong CP Problem

The *strong CP problem* is the puzzle of why the strong interaction preserves the charge-parity (CP) symmetry while the Standard Model as a whole does not and the strong interaction itself also does not have to (see *e.g.,* [1, 6, 7]). There is freedom to include in the QCD Lagrangian shown in Eq. (1.1) a CP-violating term

$$\mathcal{L}_{\text{QCD}} \supset \bar{\theta} \left(\frac{\alpha_S}{8\pi} \right) G^a_{\mu\nu} \tilde{G}^{a,\mu\nu}, \tag{1.7}$$

where $G^a_{\mu\nu}$ is the gluon field strength tensor, $\tilde{G}^{a,\mu\nu} \equiv \frac{1}{2}\epsilon^{\mu\nu\alpha\beta} G^a_{\alpha\beta}$ is its dual with $\epsilon^{0123} = 1$, the fully antisymmetric Levi-Civita symbol, and $\bar{\theta}$ is the effective θ parameter after diagonalizing the Yukawa quark mass matrices via chiral rotation and is given by $\bar{\theta} \equiv \theta + \arg\det(Y_u Y_d)$, where $Y_{u(d)}$ is the up (down) type Yukawa matrix and θ is the bare Lagrangian parameter. The $\bar{\theta}$ is effectively an angle parameter and in principle can assume any value from $-\pi$ to π. However, the non-observation of the neutron electric dipole moment (nEDM) so far places a stringent limit on the $\bar{\theta}$ parameter to be $\bar{\theta} \lesssim 10^{-10}$. The strong CP problem is the question of why $\bar{\theta}$ is so vanishingly small while the naive expectation is $\mathcal{O}(1)$. At its heart, the strong CP problem is also the question of why the nEDM is so vanishingly small.

An intuitive classical description of the problem as well as the axion solution is outlined here following [7]. In the classical picture, a neutron consists of one up

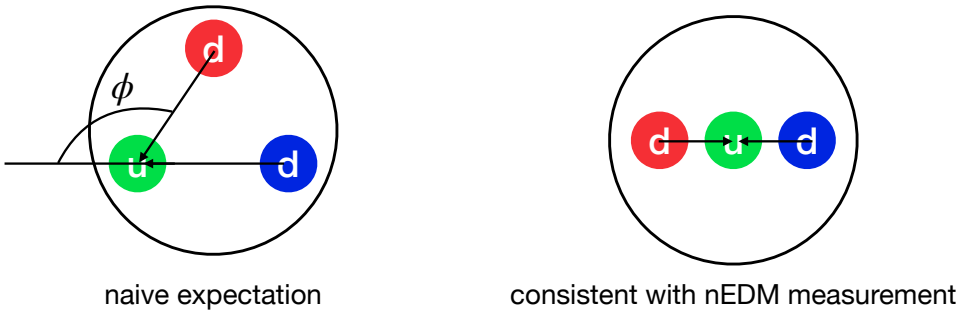

Figure 1-5: Drawings of quark distribution inside the neutron in the classical picture: (left) naive expectation and (right) configuration needed to be consistent with the non-observation of nEDM.

quark with electric charge $+2/3$ and two down quarks with electric charge $-1/3$. The electric dipole moment (EDM) depends on the spatial distribution of the three charges. Naively, one would expect the charges to be in some kind of spaced-out configuration, such as the left figure of Fig. 1-5. One could then use the formula $\vec{d} = \sum q_i \vec{r}_i$ and the fact that the size of the neutron is $\mathcal{O}(\text{fm})$ to compute the magnitude of the EDM to be

$$|\vec{d}_n| \sim 10^{-13} \sqrt{1 - \cos\phi}\, e \cdot \text{cm}, \tag{1.8}$$

where the ϕ angle is shown in the left figure of Fig. 1-5. Attempts were made to measure the nEDM using precession techniques and the non-observation of the nEDM signal could set the magnitude of the neutron EDM to be $|\vec{d}_n| \leq 10^{-26}\, e \cdot \text{cm}$ [8, 9]. This in turn means that the θ angle in Fig. 1-5 is very close to zero. In other words, we should have drawn the *udd* quarks all on the same line, shown in the right figure of Fig. 1-5. This linear configuration of quarks inside a neutron in fact bares resemblance to the molecular configuration of the familiar carbon dioxide molecules in which the two oxygen atoms are located on the exact opposite sides of the carbon atom as shown in Fig. 1-6. The atoms are arranged this way because it is the lowest energy configuration; the molecular configuration of CO_2 is realized *dynamically*. The dynamical realization scheme is the core concept behind the elegant axion solution to the strong CP problem.

Axion Solution

The axion solution elevates the θ parameter to a fully dynamical pseudoscalar field, called the *axion* field. If the axion gets its potential entirely through QCD effects, then its ground state automatically corresponds to $\theta = 0$ in a remarkable manner. The QCD axion field, a, is coupled to QCD by promoting $\bar{\theta} \to \bar{\theta} + a/f_a$:

$$\mathcal{L} \supset \frac{\alpha_S}{8\pi} \left(\bar{\theta} + \frac{a}{f_a} \right) G^a_{\mu\nu} \tilde{G}^{a,\mu\nu}, \tag{1.9}$$

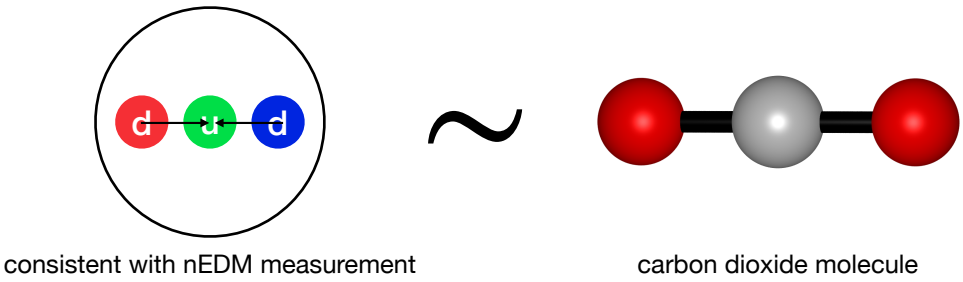

Figure 1-6: A structural illustration of the carbon dioxide molecule (right) as an analogy of the dynamical axion solution to the strong CP problem. Image courtesy: Qiuhan Li.

where f_a is the axion decay constant. The non-perturbative QCD axion potential resulting from Eq. (1.9) can be calculated using chiral perturbation theory [10, 11] as

$$\mathcal{V} \approx -m_\pi^2 f_\pi^2 \sqrt{1 - \frac{4 m_u m_d}{(m_u + m_d)^2} \sin^2\left(\frac{a}{2f_a} + \frac{\bar{\theta}}{2}\right)}, \tag{1.10}$$

where m_π and f_π are the mass and decay constant of the pion and $m_{u(d)}$ is the up (down) quark mass. Given this axion potential, the axion acquires a vacuum expectation value (VEV) of $\langle a \rangle = -f_a \bar{\theta}$, which cancels out the CP-violating term of Eq. (1.9) at low energies, hence solving the strong CP problem. The mass of the axion, m_a, as a function of the axion decay constant can then be obtained from Eq. (1.10):

$$m_a^2 \approx \frac{m_u m_d}{(m_u + m_d)^2} \frac{m_\pi^2 f_\pi^2}{f_a^2}. \tag{1.11}$$

There also exists non-axion solutions to the strong CP problem, such as the massless up quark or the running of the $\bar{\theta}$ parameter. They are beyond the scope of the discussion of this thesis and we refer the interested reader to reviews in the literature, see *e.g.*, [1, 7, 6].

1.3.2 Heavy Axions and Axion-Like Particles

Quality Problem

The axion operator in Eq. (1.9) is a dimension-5 operator; therefore, it is an effective field theory (EFT) of some high energy theory. One, and the original, UV realization of the axion is to treat the axion as a pseudo-Nambu-Goldstone boson of a new $U(1)$ Peccei-Quinn (PQ) symmetry, $e^{ia/f_a} \to e^{i\phi} e^{ia/f_a}$, which is spontaneously broken by the QCD chiral anomaly [12, 13, 14, 15]. However, if there is any other contribution to the axion potential beyond QCD, the resulting axion ground state may no longer lead to the cancellation needed to protect the QCD from large CP violation. In the far UV, quantum gravity is expected to break all global symmetries [16] such as the $U(1)_{PQ}$ symmetry; therefore, the $U(1)_{PQ}$ can at best survive as some accidental

symmetry in the low energies and needs to be protected. Following [10], we consider a $U(1)_{PQ}$-breaking Planck-suppressed operator of the form

$$\frac{\Phi^N}{M_{\rm Pl}^{N-4}} \sim \frac{f_a^N}{M_{\rm Pl}^{N-4}} e^{iNa/f_a}, \tag{1.12}$$

where $\Phi \sim f_a e^{ia/f_a}$ is a composite field from which the axion a arises as a Goldstone mode, and $M_{\rm Pl} \sim 10^{19}\,\text{GeV}$ is the Planck mass. Given the current experimental constraint $f_a > 10^9\,\text{GeV}$ [1], this requires us to forbid such $U(1)_{PQ}$-breaking operator up to a high dimension of $N = 9$ for the $U(1)_{PQ}$ symmetry to be protected. This high fragility of the axion mechanism is known as the axion *quality problem*.

Axion-Like Particles and Heavy Axions

In a more general sense, axions, or more precisely axion-like particles (ALPs), are just hypothetical pseudoscalar particles that couple to the Standard Model gauge bosons (gluons and electroweak bosons). Such couplings are suppressed at low energies by a large cut-off scale. In the minimal axion model outlined above, the axion mass, m_a, needs to follow a specific relation, dictated by the QCD axion potential in Eq. (1.10), with the axion decay constant f_a (directly related to the cut-off scale) in order to solve the strong CP problem. Beyond solving the strong CP problem, ALPs are found in many proposed extensions to the Standard Model. They could constitute all or a fraction of the dark matter content, or provide a *portal* that connects the SM particles to dark matter (see *e.g.,* [17, 18, 19, 20]). They also often appear as low energy manifestations of string theory (see *e.g.,* [21]). Therefore, the search for ALPs in all viable regions of the parameter space is an active area of investigation in the hunt for physics beyond the Standard Model.

In particular, recent model building efforts have led to considerable interest in the ALPs with GeV-scale masses (see *e.g.,* [22, 23] and references therein). One class of such efforts aims to evade the quality problem while solving the strong CP problem at the same time by introducing a mirror sector of the SM QCD with a $\mathbb{Z}_2$ symmetry [10, 24]. The axion coupling term becomes

$$\mathcal{L} \supset \frac{\alpha_S}{8\pi}\left(\bar{\theta} + \frac{a}{f_a}\right)\left(G^a_{\mu\nu}\tilde{G}^{a,\mu\nu} + G'^a_{\mu\nu}\tilde{G}'^{a,\mu\nu}\right), \tag{1.13}$$

where the primed fields are contained in the mirror sector. The confinement scale in the mirror sector, $\Lambda_{\rm QCD'}$, can be significantly larger than the SM $\Lambda_{\rm QCD}$, hence protecting the axion potential against other UV contributions, ameliorating the quality problem:

$$\mathcal{V} \supset \Lambda^4_{\rm QCD'}\left(\frac{a}{f_a} + \bar{\theta}\right). \tag{1.14}$$

At the same time, thanks to the $\mathbb{Z}_2$ symmetry, it still aligns with the potential generated by the SM QCD, hence still solving the strong CP problem. This opens up the regions of *heavy axions* with masses in the keV to TeV range and the axion decay

constant f_a in the 10–10^9 GeV range.

1.3.3 ALP-Gluon Coupling

The recent interest in the MeV-to-GeV scale ALPs has led to progress in both theoretical and experimental developments. In particular, to better understand the phenomenological impact of the ALP-gluon interactions, Ref. [22] proposed a novel data-driven method for determining the hadronic interaction strengths of such GeV-scale ALPs. We summarize the main points of this work here and refer the interested reader to [22] and its Supplemental Material for details.

The effective Lagrangian describing the ALP-gluon interactions is

$$\mathcal{L} \supset -\frac{4\pi\alpha_s c_g}{\Lambda} a G^{\mu\nu}\tilde{G}_{\mu\nu}, \tag{1.15}$$

where c_g is the dimensionless agg vertex coupling constant. The model considers $c_g \gg c_\gamma$, $c_{\rm EW}$, c_f for ALPs of masses $m_\pi \lesssim m_a \lesssim 3$GeV, but it can be generalized to any ALP couplings to the SM particles. First, a chiral transformation of the light-quark fields is performed, replacing the agg vertex by ALP-quark axial-current couplings. This leads to ALP-π^0 kinetic mixing and ALP-$\eta^{(\prime)}$ kinetic and mass mixing, which allows assigning the ALP a $U(3)$ representation at low masses. The $U(3)$-based representation is a natural choice when m_a is in the nonperturbative regime of QCD. Next, using the vector meson dominance (VMD) [25] framework, one considers an interaction vertex with two vectors and one pseudoscalar (VVP). One writes down the only valid Lorentz structure for the amplitude of the $V_1 \to V_2 P$ process and uses the available $e^+e^- \to V_1 \to V_2 P$ data to derive the unknown terms in such amplitude. Crossing symmetry is then used to derive the amplitude for $P \to V_1 V_2$ vertices, and $\Gamma_{a\to VV}(m_a)$ up to ≈ 3 GeV can be calculated. A calculation of $\eta_c \to VV$ is done using this approach and achieves $\mathcal{O}(10\%)$ accuracy, providing strong validation of the method. Finally, using this method, the authors calculated exclusive hadronic ALP decay widths and its total hadronic width. The decay branching fractions are shown in Fig. 1-7.

1.3.4 Photoproduction of Axion-Like Particles

As a further development of [22], the authors studied the photoproduction of ALPs [23] in the context of the GlueX experiment for the case where the dominant SM coupling of ALPs is to gluons. We again summarize the main points of the work here and refer the interested reader to [23] and its Supplemental Material for details.

The dominant photoproduction mechanism that provides the most sensitivity to ALP-gluon coupling is photon–vector-meson mixing and t-channel vector-meson exchange, as shown in Fig. 1-8. Authors of [23] showed that it is possible to derive a fully data-driven normalization strategy once the photoproduction mechanisms of both π^0 and η are well understood. However, as discussed in [26, 27], the production mechanism of η, while dominated by t-channel, still needs more investigation. A

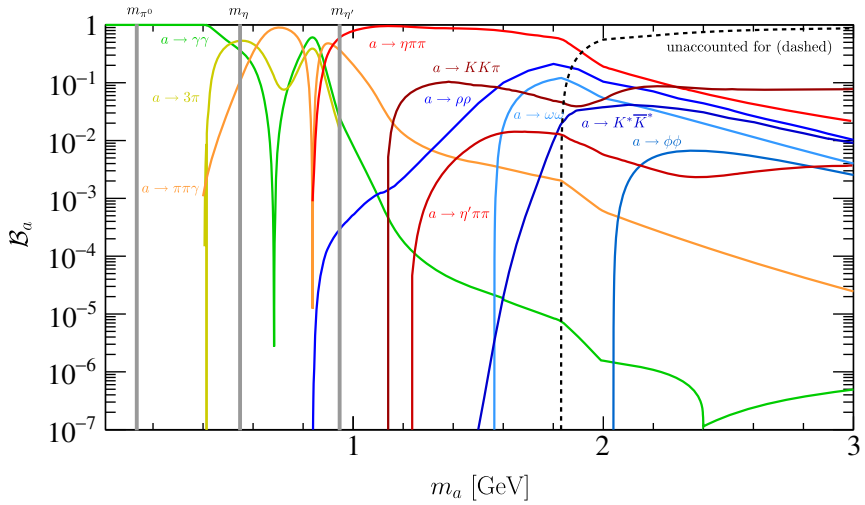

Figure 1-7: ALP decay branching fractions. Image source: [22].

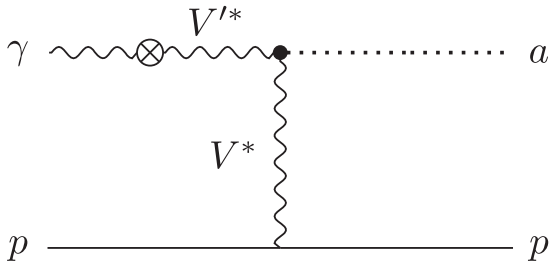

Figure 1-8: Feynman diagram of ALP photoproduction via photon–vector-meson mixing and t-channel vector-meson exchange. Image source: [23].

simplified approach was adopted in [23] to relate the differential cross section of ALP photoproduction to those of π^0 and η by the ALP-pseudoscalar mixing derived in [22]

$$\frac{d\sigma_{\gamma p\to ap}}{dt} \approx \left(\frac{f_\pi}{f_a}\right)^2 \left(|\langle a\pi^0\rangle|^2 \frac{d\sigma_{\gamma p\to \pi^0 p}}{dt} + |\langle a\eta\rangle|^2 \frac{d\sigma_{\gamma p\to \eta p}}{dt} \right), \tag{1.16}$$

where f_π and $f_a = -\Lambda/(32\pi^2 c_g)$ are the pion and ALP decay constants. This differential cross section can be further written in terms of experimental observables such as π^0, η yields and efficiencies. These relations can be used to seach for, or in the absence of signal, place limits on ALP photoproduction. We present in Chapter 5 a dedicated search for photoproduction of ALPs with masses of $\mathcal{O}$(GeV) at the GlueX experiment, described in Chapter 2.

Chapter 2

The GlueX Experiment

The *Glu*onic *Ex*citation (GlueX) experiment, located in the experimental Hall D at the US Department of Energy's Thomas Jefferson National Accelerator Facility (Jefferson Lab, JLab), aims to search and map out the spectrum of exotic mesons using a 9-GeV linearly polarized photon beam impinging on a liquid hydrogen target. The experiment includes a photon beamline facility responsible for the production and characterization of the photon beam and a detector system surrounding the target. The detector is nearly hermetic for both charged particles and photons arising from the target. Surrounded by a 2-T solenoid magnet, it consists of drift chambers in the central and forward regions for charged particle tracking, electromagnetic calorimeters in the central and forward region for photon reconstruction, and a time-of-flight detector and a DIRC detector in the forward region for particle identification. A schematic drawing of the experiment is shown in Fig. 2-1. This chapter provides a brief description of the GlueX experiment, following largely [28] and references therein.

2.1 Jefferson Lab and CEBAF

The GlueX experiment uses a linearly polarized photon beam produced from the electron beam provided by Jefferson Lab's Continuous Electron Beam Accelerator Facility (CEBAF). The facility was born out of the consensus among nuclear and particle physicists around 1980 that a new electron facility was needed to explore the transition region between the nucleon-meson and the quark-gluon descriptions of the nuclear system [29]. It began operation in the mid-1990s. The GlueX experiment was conceived in 1997 and collected first data in 2014, as part of the JLab 12 GeV upgrade [30]. A schematic drawing of the CEBAF facility, highlighting elements in the 12 GeV upgrade is shown in Fig. 2-2.

The CEBAF electron beam starts at the injector. Electrons are produced by the photoelectric effect from a diode-based laser shining on a gallium arsenide (GaAs) photocathode [29]. The laser is pulsed at 499 MHz,[1] creating a bunch every 2 ns, and

[1] Although possible, the delivery to Hall D at 499 MHz is not the norm. Instead, 249.5 MHz, with 4 ns electron bunches, is delivered to Hall D.

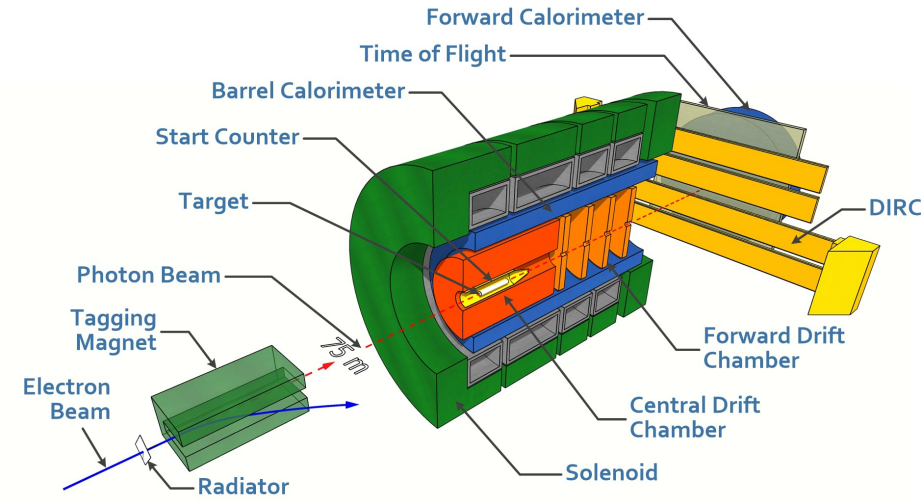

Figure 2-1: Schematic drawing of the GlueX experiment. Image source: GlueX collaboration.

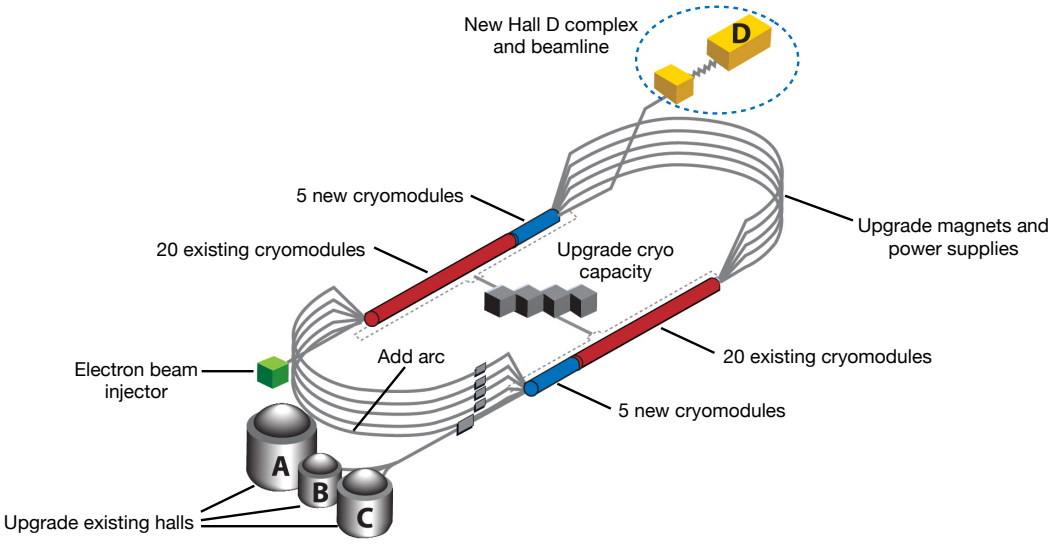

Figure 2-2: Schematic drawing of the Jefferson Lab CEBAF facility highlighting the 12 GeV upgrade. Image source: Jefferson Lab.

can be polarized (producing a polarized electron beam) or unpolarized (producing an unpolarized electron beam). The produced electrons are extracted and accelerated by superconducting RF cavities to 67 MeV before injecting into the main CEBAF recirculating linear accelerators (linacs).

The CEBAF has a *racetrack* configuration with two parallel straight acceleration sections (the North and South Linacs respectively) joined by the recirculating arcs in the bent sections. The acceleration in the linacs is provided by superconducting RF (SRF) cavities. An oscillating electromagnetic field is set up in the RF cavity. It accelerates electrons as they pass through the cavity such that the electrons always experience positive electric force, resulting in the acceleration of the electrons. The energy is transferred from that of the RF field, provided by klystrons, to the electrons. The superconducting feature minimizes energy loss in the RF field, enabling the beam quality needed for JLab physics goals.

Following the acceleration through the first straight section linacs, the electrons enter the recirculating arcs with bending power provided by the magnets before entering the other straight linac section. The recirculation continues through multiple orbits, or passes, until the desired electron energy is reached, at which point the electrons are extracted into the experimental halls A, B and C. The industrial scale adoption of the SRF technology and the use of the multipass beam recirculation were the two most important innovations of the Jefferson Lab CEBAF [29].

In the JLab 12 GeV upgrade, 5 cryomodules were added to each linac section, doubling the maximum energy from about 6 GeV to 12 GeV, or from about 1.1 GeV/pass to about 2.2 GeV/pass. In addition, a new tenth arc was added to enable a 5.5 pass acceleration for electrons into the new Hall D complex. The upgrade also included the recirculating arc magnets and power supplies to accommodate the higher beam energy and various upgrades to the existing three experimental halls.

2.2 GlueX Photon Beamline

This section provides an overview of the photon facility at GlueX. The linearly polarized photon beam that GlueX uses is produced on a thin and carefully oriented diamond crystal in the Tagger Hall, where the photon energy is determined using the photon tagging system. The polarization spectrum of the photon beam is measured by the Triplet Polarimeter and the flux of the photon beam is measured by the Pair Spectrometer and calibrated using the Total Absorption Counter. A schematic layout of the the Hall-D complex, showing the Tagger Hall, Hall D itself, and several of the key beamline devices is show in Fig. 2-3 [28].

2.2.1 Coherent Photon Source

The linearly polarized photon beam used by GlueX is produced through a process known as coherent bremsstrahlung [31, 32]. Coherent radiation is produced when the incident electron beam is aligned in special orientations with the crystal lattice of the radiator.

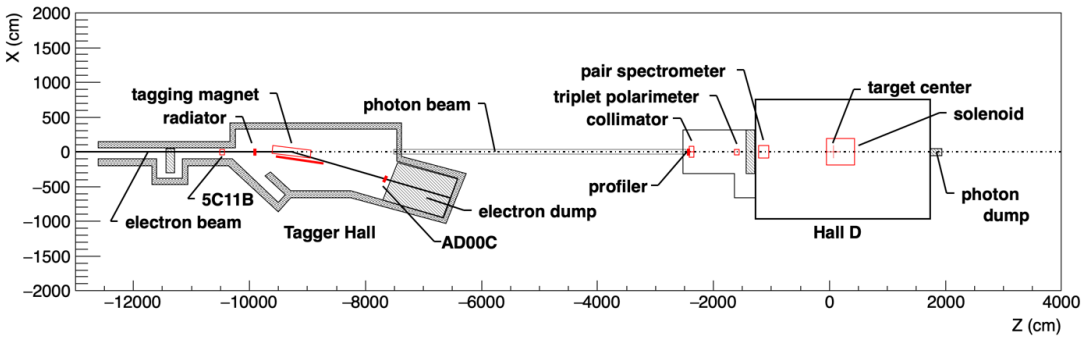

Figure 2-3: Schematic layout of the the Hall D complex, showing the Tagger Hall, Hall D itself, and several of the key beamline devices. Image source: [28].

At GlueX, $7 \times 7 \times 0.05\text{mm}^3$ diamond crystal radiators produced via chemical vapor deposition (CVD) are routinely used. Diamond is chosen both because of its crystal lattice structure with small lattice constant and its high Debye temperature. The position of the coherent edge in the photon beam intensity spectrum is a simple monotonic function of the angle between the incident electron beam direction and the normal to the (2,2,0) crystal plane. The thickness of 50 μm is found to be optimum because it is thin enough to limit the effects of multiple scattering which smears the coherent edge and it is thick enough to maintain mechanical stability from buckling due to internal stress from radiation damage.[2] The diamond radiators are mounted on a goniometer that precisely adjusts the relative orientation of the radiators with respect to the incident electron beam to the desired configuration.

In addition, the degree and direction of the linear polarization also depend on the angle between the incident electron beam direction and the normal to the $(2, 2, 0)$ crystal plane. Polarization of the photon beam is critical to the GlueX physics program for constraining the initial state in amplitude analyses. For a 12-GeV incident electron beam, a coherent edge up to 12 GeV can be achieved; however, the degree of polarization decreases as the coherent edge position increases [32]. The choice of 9 GeV for the primary peak energy at GlueX, which corresponds to a maximum of about 40% linear polarization, is found to be the optimal balance between the need for high energy photons in order to produce the predicted hybrid mesons, and the need for a large linear polarization for the identification of the produced states in the GlueX physics program [28]. The coherent radiation travels along the electron beam direction whereas the ordinary bremsstrahlung radiation is produced in a cone; therefore, a 5 mm diameter collimator placed about 75 meters downstream of the diamond radiator not only reduces the beam halo but also the incoherent component, greatly increasing the polarization of the photon beam. Fig. 2-4 [28] shows both the intensity and the polarization spectrum of the GlueX linearly polarized photon beam.

[2] A 20 μm thick crystal was also tried, but deemed not ideal due to its poorer response to radiation damage.

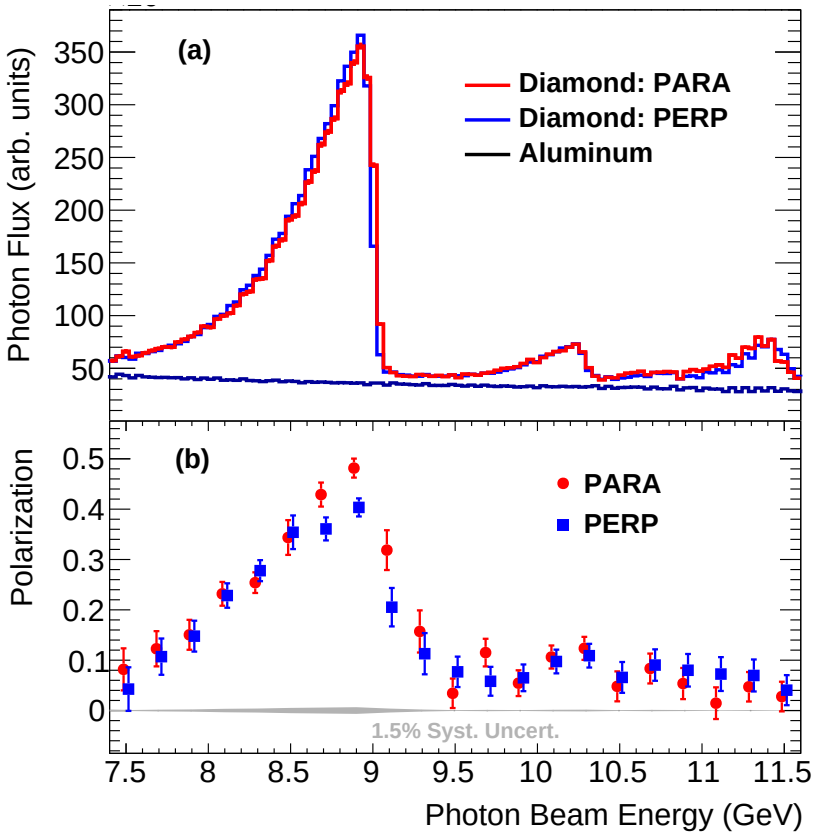

Figure 2-4: Collimated photon beam intensity spectrum (a) and polarization spectrum (b) at GlueX. The labels PARA and PERP refer to orientations of the diamond radiator that result in polarization planes that are parallel and perpendicular to the horizontal direction. The Aluminum label shows the ordinary bremsstrahlung spectrum produced by an aluminum amorphous radiator normalized by radiation lengths. Image source: [28].

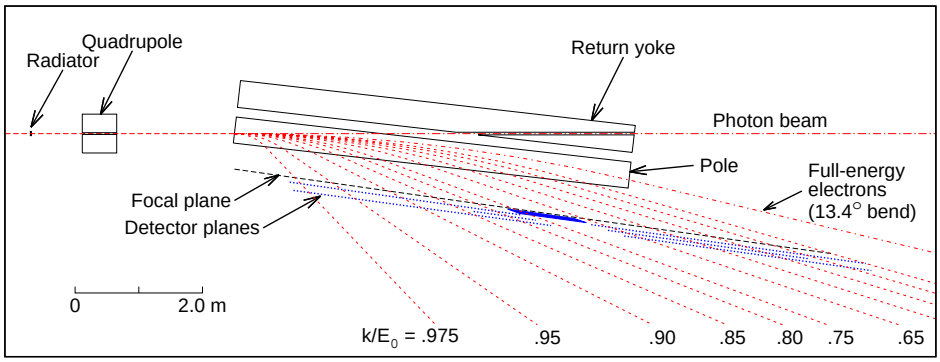

Figure 2-5: Schematic layout of the GlueX tagger system, showing the paths of the recoil electrons (red dotted lines) and the bremsstrahlung photons (red dash-dotted line). The recoil electron trajectories are identified by the energy the electron gave up to an associated bremsstrahlung photon k, as a fraction of the incident primary electron beam energy E_0. Image source: [28].

2.2.2 Photon Tagging System

The purpose of the photon tagging system, or the *tagger*, is to measure the produced bremsstrahlung photon energy by measuring the recoil electron energy. The GlueX tagger system consists of a dipole tagger magnet, a tagger microscope (TAGM), and a tagger hodoscope (TAGH). Both TAGH and TAGM are scintillating detectors which detect the recoil electrons and use the position information of the electrons to get a measurement of the energy. Fig. 2-5 shows a schematic drawing of the tagger system. More detailed information about the Tagger can be found in [28] and references therein.

Tagger Magnet

Just after the diamond radiator, the beam consists of a mixture of produced bremsstrahlung photons, recoil electrons, as well as beam electrons that did not interact. The tagger magnet is a dipole magnet with nominal operating field of 1.5 T, under

which the electrons pass through a thin window out of the beamline vacuum system and bend towards either the TAGH/TAGM tagging scintillators or the electron beam dump for those beam electrons that did not interact with the radiator to produce bremsstrahlung photons. The bremsstrahlung photon energy can be written as $E_\gamma = E_0 - E_e$, where E_0 is the nominal 12-GeV of full beam energy and E_e is the recoil electron energy. The more the beam electron loses its energy to the produced bremsstrahlung photon, the less the recoil energy, and the greater the bending; therefore, there is a one-to-one mapping between the recoil electron trajectory and the bremsstrahlung photon energy. The bremsstrahlung photons are "tagged" by the recoil electrons, and the finer the recoil electron position is determined, the better the energy resolution of the bremsstrahlung photon. There is also a permanent dipole magnet with a strength of 0.8 Tm downstream of the tagger magnet on the photon beam line to prevent the accelerator electron beam from reaching Hall D if the main tagger magnet trips.

Tagger Microscope

Tagger Microscope (TAGM) is an array of highly segmented scintillating fibers facing the recoil electrons bent by the tagger magnet. It is indicated by the dense blue region in Fig. 2-5. Each scintillating fiber is connected by a light guide and the scintillating photons are read out by a shielded silicon photomultiplier (SiPM) of matching dimensions which converts the scintillating photon energy into electronic signals which are then read out by TDCs and flash ADCs. The TAGM covers the photon energy range in the coherent peak, *i.e.,* between 8.2 and 9.2 GeV in the nominal 12-GeV running condition. It is designed to be movable should the primary beam energy change or the experiment decide to focus in a different photon energy region. The high level of segmentation of TAGM provides electron energy resolution of about 12 MeV (about 0.1% of the incident electron energy). This translates to a very small uncertainty in the photon beam energy, which, combined with the fact that the target proton is stationary, enables the kinematic fitting of a given reaction since both the initial and final states are measured. The timing resolution of TAGM is about 230 ps, which is enough to identify from which 4 ns beam bucket the reaction is initiated.

Tagger Hodoscope

The basic idea of the Tagger Hodoscope (TAGH) is similar to that of the TAGM: determining the bremsstrahlung photon energy by measuring the position of the recoil electrons using a scintillating detector. The TAGH scintillating counters are read out by photomultiplier tubes (PMTs). The TAGH scintillator counters are spaced coarser than those of the TAGM scintillating fibers, and it is designed to cover 25% to 97% of the incident electron beam energy. This broad coverage expands the GlueX physics program beyond the coherent peak. TAGH counters are indicated by the coarser blue dots in Fig. 2-5. The recoil energy resolution, depending on the location of the TAGH counter, is between 8 MeV and 30 MeV, and the timing resolution is about 200 ps.

2.2.3 Photon Beam Flux and Polarization

Triplet Polarimeter

Both the Triplet Polarimeter (TPOL) described in this section and the Pair Spectrometer (PS) described in Sec. 2.2.3 rely on the triplet photoproduction process of $\gamma e^- \to e^- e^+ e^-$. A thin 75 μm beryllium target, located as part of the TPOL detector system, is used as the atomic electron target for the triplet production. There are eight QED diagrams in total and four of them are shown in Fig. 2-6. This process produces a "slow" recoil electron (line 4), which is detected in the TPOL used to measure the photon beam polarization, and a "fast" electron and positron (lines 2 and 3) pair, which is detected in the PS further downstream used to measure the photon beam flux.

The cross section of the triplet photoproduction process for a polarized photon beam can be written as

$$\sigma_t = \sigma_0 \left[1 - P\Sigma \cos(2\phi)\right], \tag{2.1}$$

where σ_0 is the unpolarized cross section, P the photon beam polarization, Σ the beam asymmetry for the process, and ϕ the azimuthal angle of the trajectory of the recoil electron with respect to the polarization for the incident photon beam. The TPOL detector system essentially consists of the beryllium foil converter mounted on a motorized positioning assembly, a silicon strip detector (SSD) that provides the energy and azimuthal angle information of the recoil atomic electron, an electronics readout system, and a vacuum system that houses all the parts above. The SSD is segmented into 30 wedges, providing 30 azimuthal bins. The yield per azimuthal bin can be measured and fitted to a function of the form $A\left[1 - B\cos(2\phi)\right]$, where A is some normalization factor and $B = P\Sigma$ can be used to determine the polarization P once Σ is known. Σ is the beam asymmetry and can be calculated with QED from the diagrams in Fig. 2-6 but also accounting for detector geometry and detection efficiency obtained from Monte Carlo simulation and validated with measurements. More details about the TPOL can be found in [28, 33].

Pair Spectrometer

The main purpose of the Pair Spectrometer (PS) is to measure the photon beam flux by measuring the produced e^+e^- pairs from the same beryllium foil converter housed in the TPOL produced from the triplet photoproduction process described in Sec. 2.2.3. A schematic drawing of the PS detector system is shown in Fig. 2-7. A dipole magnet of 1.8 T strength bends the e^+e^- pairs in opposite directions towards the two identically instrumented arms of scintillating counters, which perform a coincidence measurement of the e^+e^- pair. There are two layers of scintillators in each arm: a higher granularity hodoscope (PS) and a set of coarse counters (PSC). The PS counters are read out by SiPMs and PSC read out by PMTs. The timing resolution of the PS is about 120 ps, which provides a coincidence measurement with the tagger as well as trigger for the TPOL. The PS also measures the energies of the e^+e^- pairs, which provides energy calibration for the TAGM and TAGH.

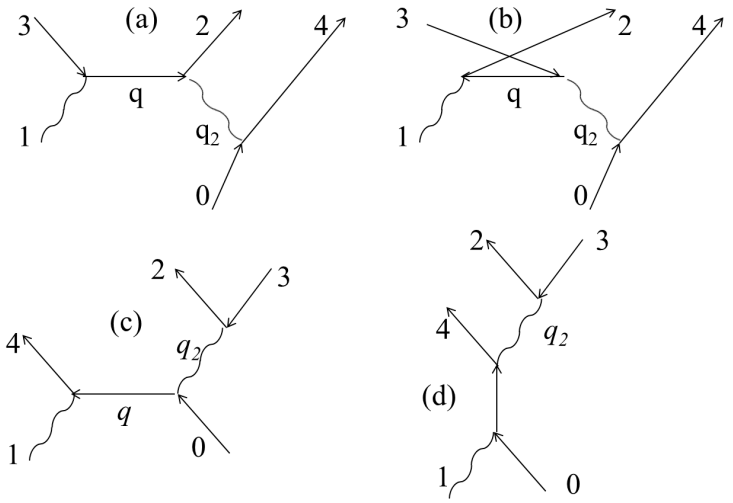

Figure 2-6: Feynman diagrams of the triplet photoproduction process. Time runs upwards. In all diagrams, the following mapping is followed: 0 – target electron, 1 – incident photon, 4–recoil atomic electron, 2,3–pair-produced electron and positron. The other 4 diagrams are obtained by exchanging lines 2 and 4. Image source: [33].

The photon beam flux is a critical input to any cross section measurement. The photon flux is determined by counting the detected e^+e^- pairs in the PS (with a dedicated PS trigger) as a function of energy and converting the raw yields into the total number of photons incident on the GlueX target using calibration parameters corresponding to the fraction of the photon beam that converts to e^+e^- pairs. The calibration parameters depend on the converter thickness, and the acceptance and efficiency of the PS detector. They are calibrated using the Total Absorption Counter described in Sec. 2.2.3. More detailed information about the PS can be found in [28, 34].

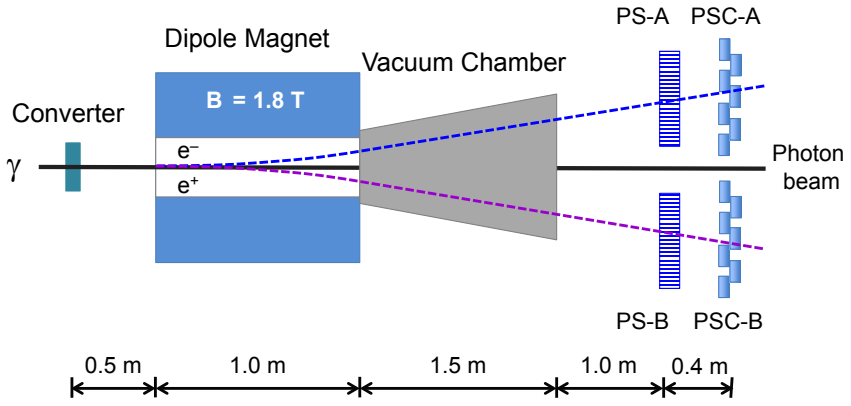

Figure 2-7: Schematic layout of the PS detector system. Image source: [34].

Total Absorption Counter

The Total Absorption Counter (TAC) is a small high-efficiency lead-glass calorimeter installed just upstream of the photon beam dump in Hall D as seen in Fig. 2-3. The TAC is used in dedicated low current runs in order to count all beam photons above a certain threshold to provide an overall normalization of the photon beam flux. During these dedicated TAC runs, the tagged photon rates in the tagger and the coincidence rate in the PS are also recorded. The relative count rates among these detectors during the dedicated TAC runs are used to convert the rates observed by the tagger and PS during normal operations to an absolute flux onto the GlueX target.

2.2.4 Other Beamline Devices

Collimators

At GlueX, a 5-mm primary collimator is used during nominal operation. Just upstream of the primary collimator, there is also an Active Collimator, which monitors the photon beam position and provides feedback to micro-steering magnets in the electron beamline, for the purpose of suppressing drifts in photon beam position. The collimation is important for GlueX as it significantly reduces the incoherent component of the photon beam, increasing the coherent fraction. The collimator cave is indicated in Fig. 2-3 about 75 meters downstream of the diamond radiator.

Beam Position Monitors and Profiler

Beam position monitors (BPM), such as the 5C11B and AD00C in Fig. 2-3, monitor the position of the primary electron beam. The beam profiler is located just upstream of the collimators. It consists of scintillating fibers read out by multi-anode PMTs. The beam profiler is only used during beam setup until the photon beam is centered on the active collimator.

Tungsten keV Filter

A 100 μm tungsten foil is installed in the beamline at the entrance of the collimator cave, with the purpose of reducing the photon flux in the 10–100 keV range. These low energy photons significantly increase the random hits in the drift chambers, especially the central drift chamber (CDC) [28]. The insertion of the tungsten foil reduced the number of random hits in the inner CDC layers by a factor of up to 8 and the anode current by 55%. The reduction of the current in the FDC was more moderate, about 25%. The greater reduction in the CDC than in the FDC was due to the fact that the CDC is farther away (10 cm) from the beamline than the FDC (3 cm), which in turn means that low energy photons have a higher impact for the CDC than for the FDC.

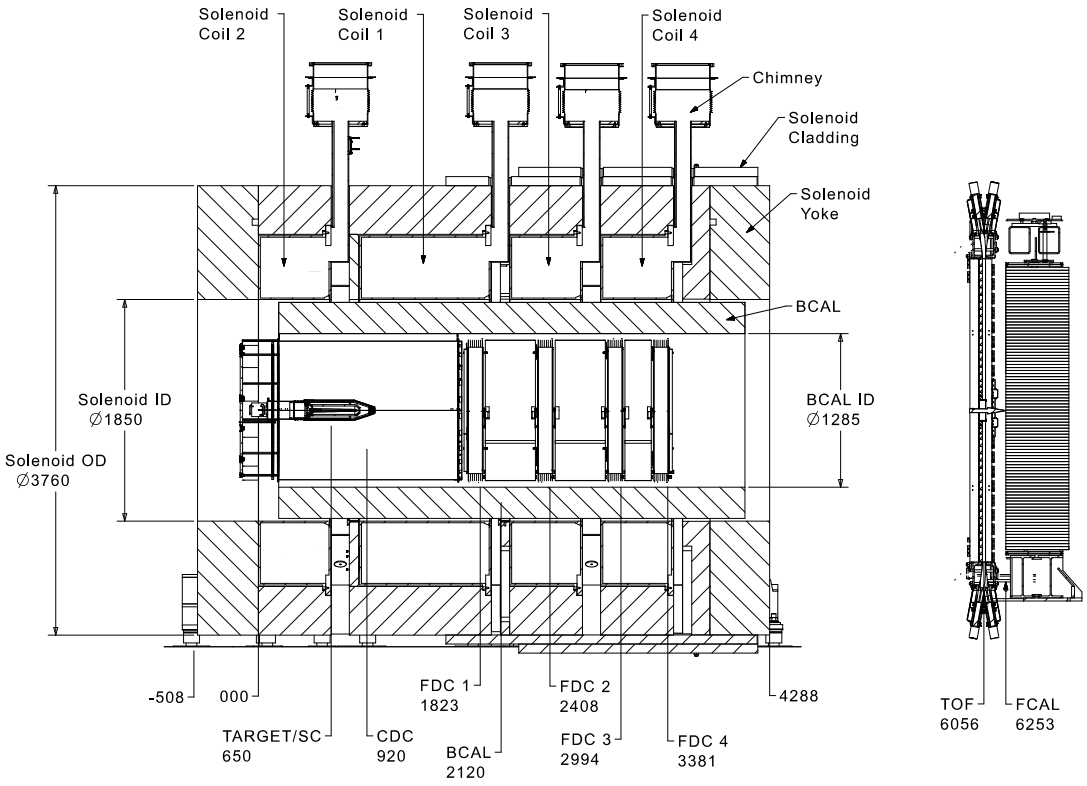

Figure 2-8: Schematic layout with dimensions in mm of the GlueX detector without the DIRC detector. Abbreviations: SC–Start Counter, CDC–Central Drift Chamber, FDC–Forward Drift Chamber, BCAL–Barrel Calorimeter, TOF–Time-of-Flight hodoscope, FCAL - Forward Calorimeter. Image source: [28].

2.3 GlueX Detector

A schematic layout of the cross sectional view of the GlueX detector is shown in Fig. 2-8. This section provides a brief description of each of the component with more details found in [28].

2.3.1 Solenoid Magnet

A superconducting solenoid surrounds the central region of the GlueX detector. The nominal operation current is 1350 A and the field strength along the axial direction is about 2 T. The superconducting coils are cooled by a liquid helium cryogenics system. The solenoid magnet was originally built for the LASS experiment and operated at SLAC in the 1980s. It was then moved to Los Alamos and used in the MEGA experiment. In the early 2000s, it was refurbished for GlueX at Indiana University and at Jefferson Lab.

Knowledge of the magnetic field map is an important input to the charged particle track reconstruction. The field map is calculated using a 2-dimensional field calculator developed at Los Alamos National Laboratory assuming axial symmetry [28] and is

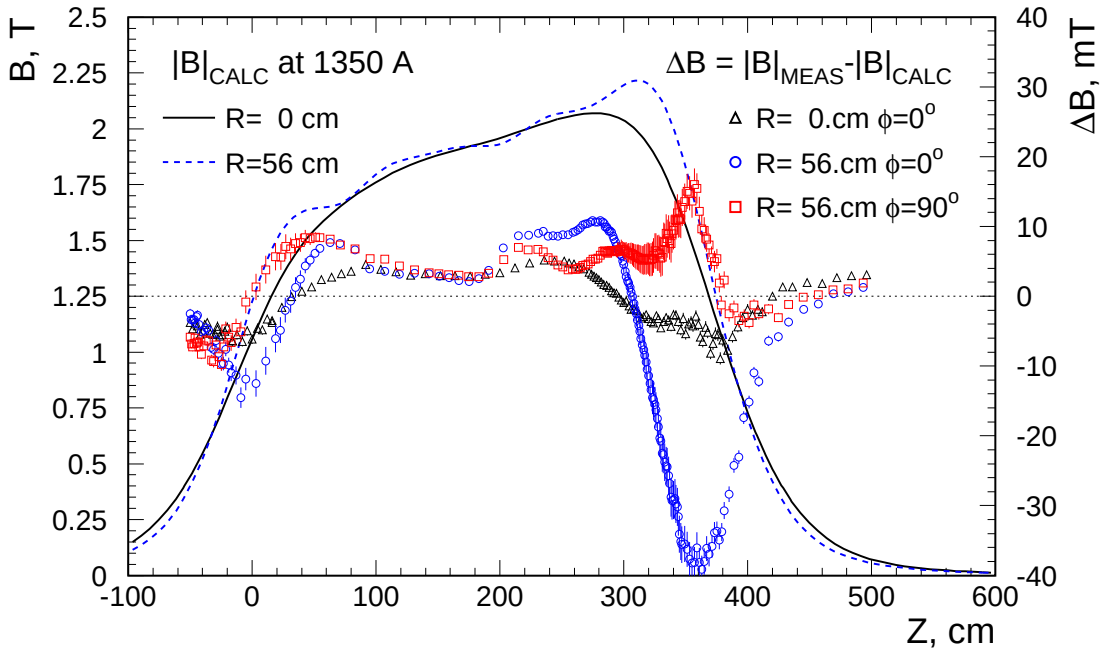

Figure 2-9: Magnetic field map of GlueX showing both the measured and calculated field strengths. Image source: [28].

validated with measurements. Fig. 2-9 shows a comparison of the magnetic field map between calculation and measurement. The largest discrepancy is about 2%, which is acceptable because the tracking resolution at GlueX is larger than 1% (see Fig. 2-12) and is dominated by multiple scattering and position measurements.

2.3.2 Target

Liquid hydrogen is used as the source for the target protons at GlueX. A staged cooling system is in place to cool the gaseous hydrogen down to its liquid phase. The cooled liquid hydrogen fills a target cell of cylindrical shape that is about 30 cm long and about 1.6 cm in diameter. The liquid hydrogen target is kept at a density of $71.2 \pm 0.3\,\mathrm{mg/cm^3}$. The target cell walls are made of aluminized 100 μm thick polyimide foils except for the entrance and exit windows which are not aluminized. The target cell tapers slightly towards the forward region forming a conical shape in order to prevent bubbles from collecting inside the cell. The target cell is enclosed within an aluminum and stainless steel vacuum chamber. Fig. 2-18 shows a schematic drawing of the target cell along with the Start Counter. The target windows as well as the vacuum chamber walls are visible in the vertex reconstruction. Dedicated empty target runs are performed occasionally or opportunistically during the nominal running periods. These empty target runs are often used for systematic studies in physics analyses. The reconstructed vertex position from an empty target run is shown in Fig. 2-13. More detailed information about the target can be found in [28] and references therein.

2.3.3 Tracking System

The GlueX tracking system consists of the Central Drift Chamber (CDC) and the Forward Drift Chamber (FDC), both enclosed in the solenoid magnet as shown in Fig. 2-8.

Central Drift Chamber

The Central Drift Chamber (CDC) consists of 3522 straw drift tubes arranged in 28 layers. Each straw tube consists of a 20-μm diameter gold-plated tungsten anode wire and 15.55-mm inner diameter Mylar tube with a 100-nm layer of vapor-deposited aluminum as the cathode. The inside volume of each straw tube is filled with 50-50 argon-carbon dioxide gas mixture at atmospheric pressure. Each straw tube is enclosed by an aluminum outer wall and supported by aluminum end plates. The active volume of the CDC has inner radius of about 10 cm and outer radius of about 56 cm measured from the beamline, covering polar angles between 6° and 168°. The 28 layers are arranged in 7 4-layer superlayers, as shown in Fig. 2-10, with 3 axial superlayers, 4 stereo layers with ±6° to provide position information along the beam direction. The anode wires are held at +2.1 kV during normal operation.

When a charged particle passes through a straw tube, it ionizes the gas molecules producing electron-ion pairs along its path. The electrons drift towards the anode wire while the ions drift towards the cathode wall. As the electrons get close to the anode wire, the electric field from the anode wire due to its small diameter becomes so strong that electrons gain enough energy to further ionize the gas molecules, creating an avalanche. The electrons from the avalanche are collected by the anode wire which is connected to preamplifier electronics which produce signal output. The information of which wires are hit as well as the hit time, related to the drift time of the electrons, is used in the reconstruction of the charged particles described later in this section. The position resolution from the CDC is a function of the distance of the hit position from the wire, and is below the design resolution of 150 μm if the tracks are at least 1.7 mm from the wire and reaches about 70 μm beyond 3.5 mm. More detailed information about the CDC can be found in [28, 35].

Forward Drift Chamber

The Forward Drift Chamber (FDC) consists of 24 disc-shaped planar drift chambers of 1 meter diameter each. They are grouped into four packages with six chambers in each package. The location of the FDC chambers can be seen in Fig. 2-8. The FDC covers charged particles of polar angles between about 1° and about 10°. Each chamber consists of four planes: the wire plane, two cathode planes on either side of the wire plane with 5 mm distance, and a ground plane; a schematic drawing of the FDC chamber is shown in Fig. 2-11. Each wire plane consists of 20-μm diameter sense wires and 80-μm diameter field wires 5 mm apart. A gas mixture of 40% argon and 60% carbon dioxide is used. The sense wires are kept at +2.2 kV whereas as the field wires at -0.5 kV. Similar to the CDC, the hit position and drift time information is used in the charged particle reconstruction. The hit position resolution from the

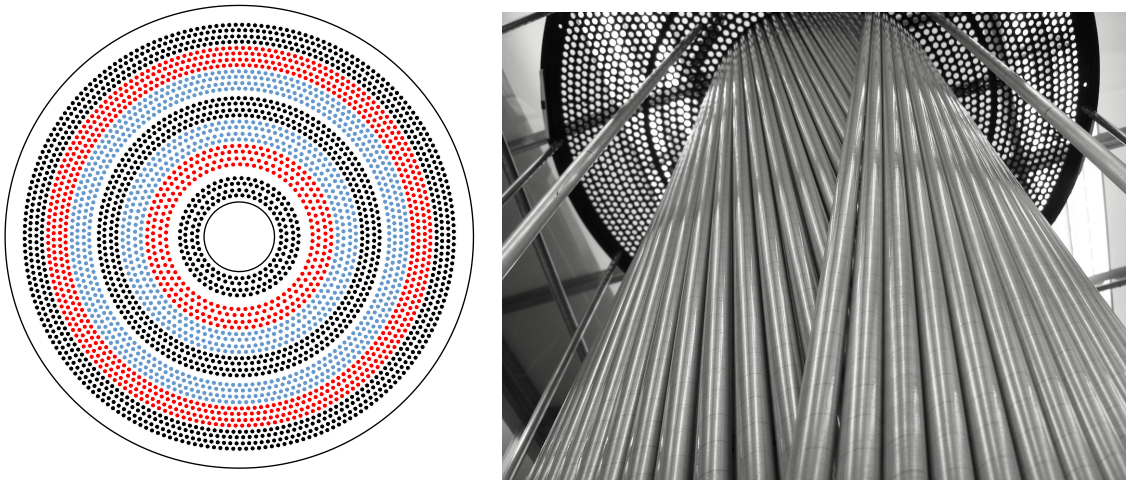

Figure 2-10: (Left) Schematic diagram of the CDC straw orientations viewed from upstream end plate; the axial straws are shown in black, the $+6°$ stereo layers are shown in red, and the $-6°$ stereo layers are shown in blue. (Right) The CDC during construction, showing the interface between two stereo superlayers. Image source: [35].

FDC is between 240 and 140 μm depending on the distance of the hit to the wire in the 0.5–4.5 mm range. More detailed information about the FDC can be found in [28, 36].

Reconstruction and Performance

The charged particle track reconstruction at GlueX is done in three stages.

- The first stage is pattern recognition. In the FDC, hits in adjacent layers in the FDC in each package are formed into *track segments* (assuming straight paths) that are then linked with track segments in other packages using a simple helical model, forming an FDC *track candidate*. In the CDC, hits in adjacent axial layers are first associated into track segments and are linked with segments in other axial layers and fitted with circles in the project perpendicular to the beamline. Intersections between these circles and the stereo wires are found and a linear fit is performed to find a z position near the beamline, forming a CDC track candidate.

- The second stage uses a Kalman filter to find the fitted track parameters $\{z, D, \phi, \tan\lambda, q/p_T\}$ at the position of the closest approach of the track to the beamline. The track candidate parameters are used as an initial guess and the fit proceeds in steps from the hits farthest from the beamline towards the beamline,[3] according to the calculated magnetic field map, accounting for energy loss and multiple scattering. Note that the drift time information from the wires is not used in the first Kalman filter pass.

[3]With the installation of the DIRC, the Kalman filter fit procedure was augmented to also include the direction from the hits innermost to the beamline outwards. See Sec. 4.6 for more details.

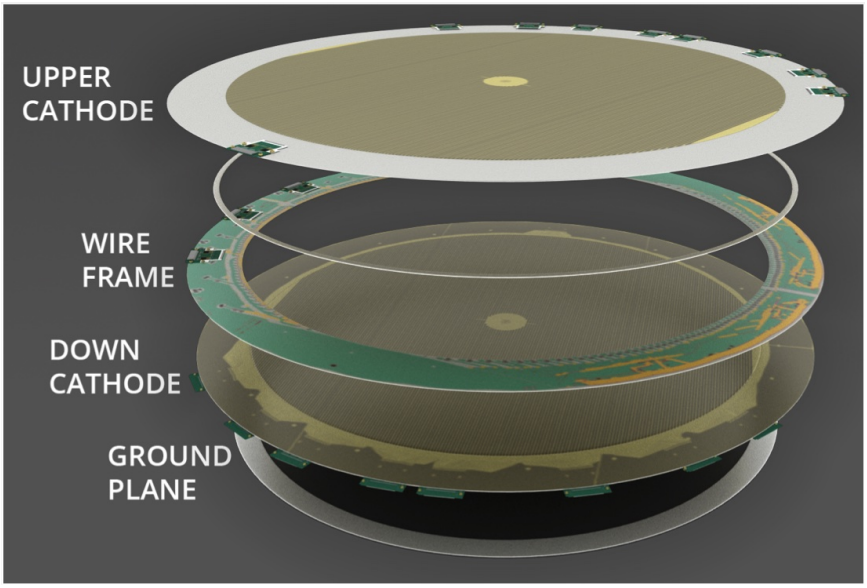

Figure 2-11: Schematic drawing of the FDC chamber showing the components. Image source: [28].

- The third and final stage matches the fitted tracks to the other subdetectors including the Start Counter, the Time-of-Flight scintillators, the Barrel Calorimeter, or the Forward Calorimeter to determine a start time. Then each track is refitted with the drift time information separately for each mass hypothesis from the set $\{e^{\pm}, \pi^{\pm}, K^{\pm}, p^{\pm}\}$.

The momentum and angular resolution as a function of momentum and polar angle are shown in Fig. 2-12. The charged particle momentum resolution can be parameterized in the following form

$$\frac{\delta p}{p} = \frac{1}{B \sin\theta} \left(\frac{a}{\beta} \oplus bp \right), \tag{2.2}$$

where B is the magnetic field, θ the polar angle of the particle, β particle velocity, a, b parameters, and $\oplus$ indicating addition of the two terms in quadrature. The a term is due to multiple scattering and the b term is due to uncertainty from the position measurement. Due to the solenoid field configuration, the momentum resolution is poor for small angles because only the projection of the particle momentum perpendicular to the beamline is subject to the bending. For pions, we see that the resolution is the best around 1 GeV/c and becomes worse for lower momenta due to the multiple scattering term and for higher momenta due to the position term. For protons, due to its large mass and hence low velocity, the momentum resolution is dominated by the multiple scattering term; therefore, it improves as the momentum increases between 0.5 GeV/c and 1.5 GeV/c. The region around 10° is the transition region between the FDC and CDC which results in degradation in the momentum resolution.

The reconstructed z vertex from an empty target run is shown in Fig. 2-13, where

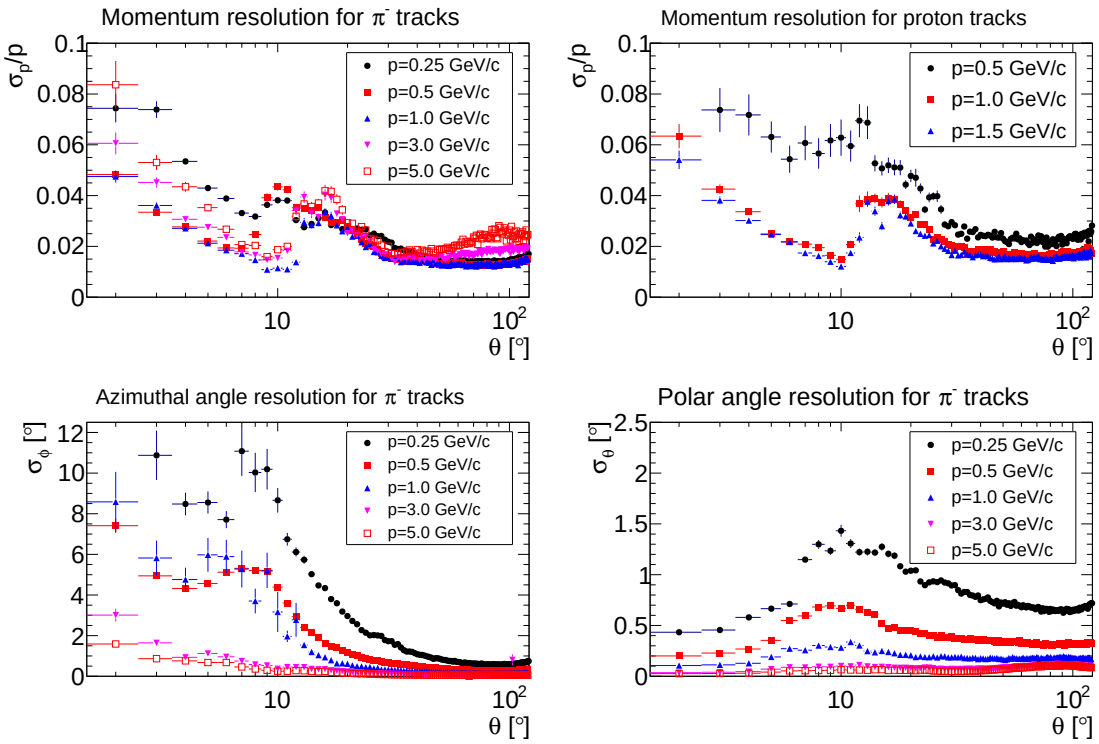

Figure 2-12: (Top) Momentum resolution for π^- tracks (left) and proton tracks (right). (Bottom) Angular resolution for π^- tracks in azimuthal angle, ϕ (left) and in polar angle, θ (right). The resolutions are plotted as a function of track polar angle θ and for a few momentum categories. Image source: [28].

the target windows as well as the vacuum chamber wall are clearly visible and are used to estimate the vertex resolution of the tracking system. The z position resolution is determined to be about 3 mm.

2.3.4 Calorimeter System

The calorimetry system at GlueX consists of the Barrel Calorimeter (BCAL) in the central region inside the solenoid magnet bore and the Forward Calorimeter (FCAL) in the forward region, as shown in Fig. 2-8.

Barrel Calorimeter

The Barrel Calorimeter (BCAL) is an electromagnetic sampling calorimeter in the shape of an open cylinder. In the radial direction away from the beamline, it is located between the CDC and the magnet. The BCAL covers photons with energy from about 50 MeV to a few GeV, and from about 11° to about 126° in polar angle, and all 2π of the azimuthal angle. The BCAL consists of 48 modules in total each covering 7.5° in the azimuthal direction. Each module is essentially a matrix of 0.5 mm thick lead sheets and 1 mm diameter scintillating fibers. Each module has approximately 185

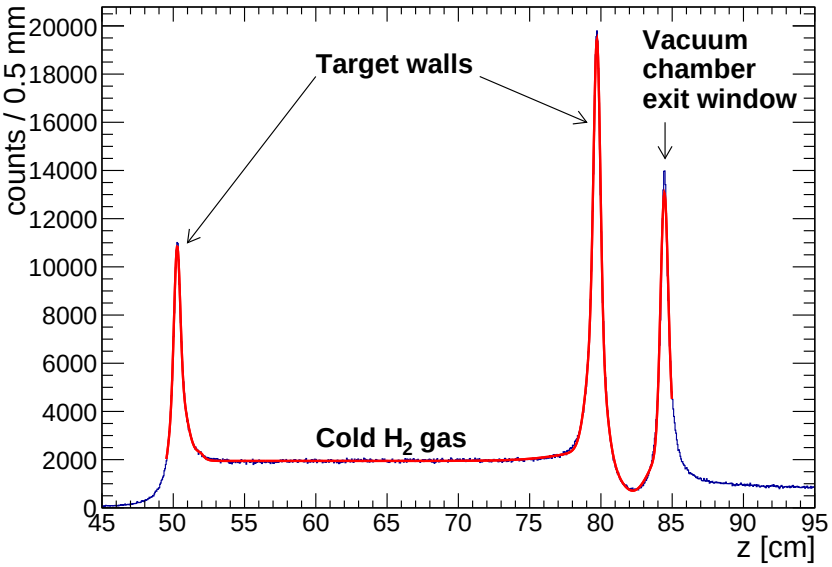

Figure 2-13: Reconstructed vertex positions within 1 cm of radial distance around the beamline from an empty target run. Image source: [28].

lead sheet layers and 15,000 scintillating fibers. The scintillating fibers are grouped and read out at the upstream end by SiPMs via light guides with matching dimensions to the SiPM surface. SiPMs are used because they are insensitive to the magnetic field. A schematic rendition of the BCAL is shown in Fig. 2-14.

As a high energy photon enters the BCAL, it interacts with the electric field of the lead nuclei through primarily pair production at GlueX relevant energies. The produced e^+e^- pair interacts again with the lead nuclei to produce bremsstrahlung photons, which then further pair produce e^+e^- pairs, creating an electromagnetic *shower*. As the created electrons and positrons pass through the scintillating fibers interspaced between the lead sheets, they excite the molecules which emit optical photons when the molecules de-excite. The collected scintillation light is proportional to the energy loss of electrons and positrons in the shower, which can then be related to the high energy photon that initiated the shower. It is called a sampling calorimeter because a lot of energy lost by the initial photon happens in the lead and the scintillating fibers only *sample* a fraction of the shower energy. Depending on the incident angle of the photon, the BCAL has a thickness of about 15.3 up to 67 radiation lengths. Ideally, all the light from one shower is read out by the SiPMs which can then be used to measure the photon energy by applying calibration parameters. The layered readout system also provides radial information of the shower shape which helps in energy reconstruction and particle identification. If the photon is not fully stopped in the BCAL, energy leakage happens which degrades the energy resolution. More detailed information about the BCAL can be found in [28, 37].

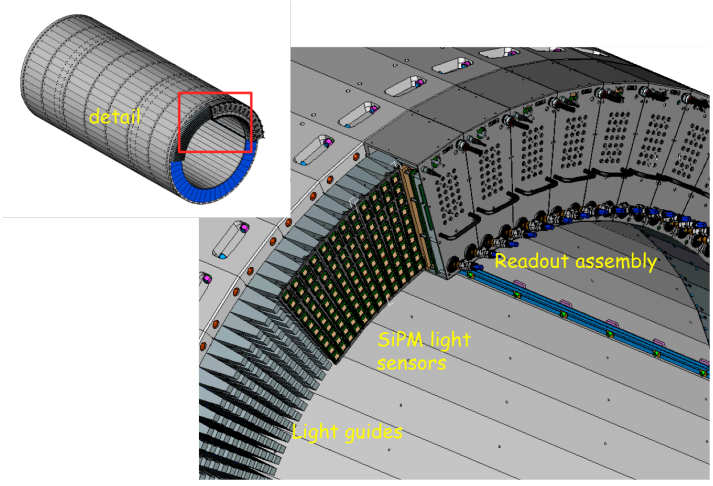

Figure 2-14: Schematic drawing of the BCAL, showing the wedge modules, light guides, SiPMs, and the read out electronics. Image source: [28].

Forward Calorimeter

The Forward Calorimeter (FCAL) is a homogeneous electromagnetic calorimeter in the shape of a wall in the forward region. It consists of 2,800 lead glass blocks stacked in a circular manner around the beamline with a diameter of about 2.4 meters. Each lead glass block has a dimension of $4 \times 4 \times 45\,\mathrm{cm}^3$. Each block is read out by a PMT via an acrylic light guide which is connected to the block via an optical interface ("cookie"). The detector is enclosed in a dark room. The blocks and most of the PMTs were taken from the decommissioned E853 experiment at Brookhaven National Laboratory and the RadPhi experiment at Jefferson Lab. The FCAL covers photons with energy from about 100 MeV to a few GeV and about 1° to 11° in polar angle. A photo of the FCAL during construction is shown in Fig. 2-15 and a schematic drawing of an FCAL block-readout module is shown in Fig. 2-16.

A lead glass calorimeter, such as the FCAL, detects the Cherenkov light produced by the fast electrons and positrons in the electromagnetic shower. The number of produced Cherenkov photons is proportional to the path length of the charged particle, in this case electrons and positrons. The path length of the electrons and positrons in the shower is in turn proportional to the energy loss, which is proportional to the high energy photon that initiated the electromagnetic shower. A high energy photon typically initiates a shower that expands multiple FCAL modules. Photons and hadrons also have different signatures in the spatial energy deposition. A reconstruction algorithm takes such effects into account and sums energies deposited in different blocks in a systematic way. More detailed information about the FCAL can be found in [28, 38].

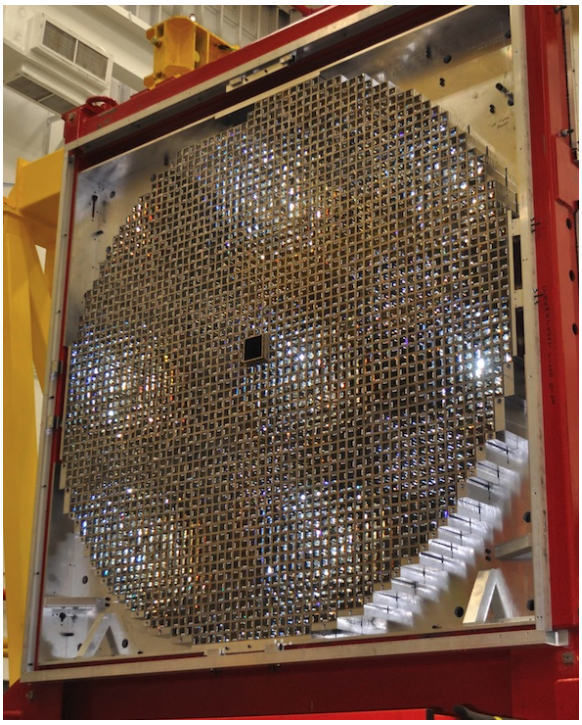

Figure 2-15: Picture of the FCAL wall during construction. The illumination was due to LED light testing. Image credit: GlueX collaboration.

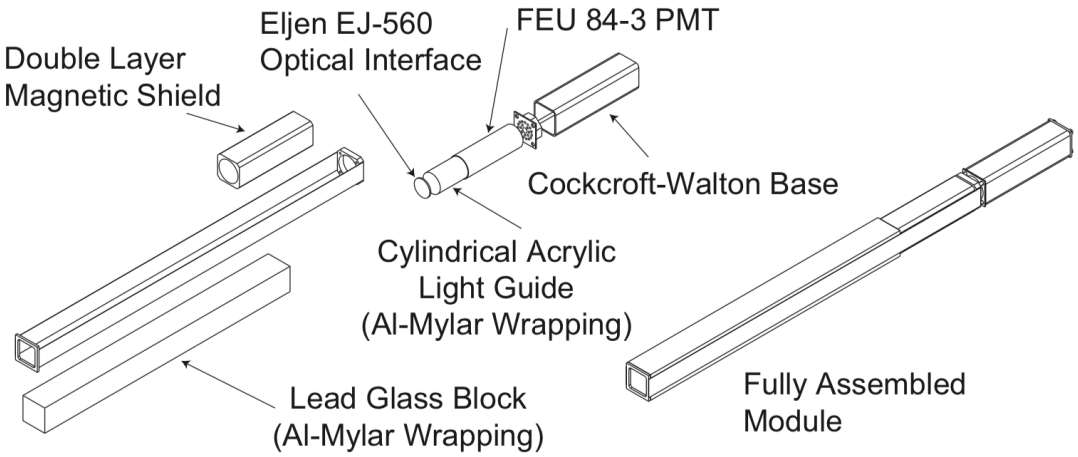

Figure 2-16: Schematic view of a single FCAL module. Image source: [28].

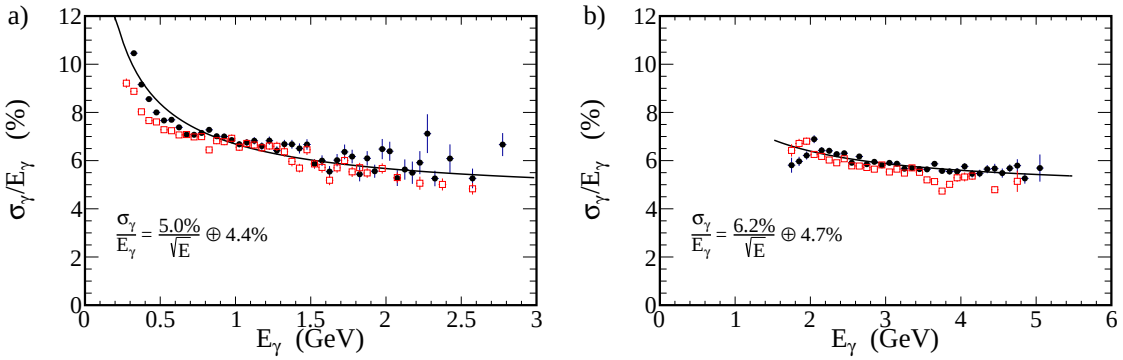

Figure 2-17: The energy resolution, σ_γ/E_γ, as a function of photon energy for single photons in the (a) BCAL and (b) FCAL. Solid black circles are data and open red squares are simulation. The fits to data including the stochastic and constant terms are indicated as the black lines.

Performance

The photon energy resolution can be parameterized in the following form

$$\frac{\delta E}{E} = \frac{a}{\sqrt{E[\text{GeV}]}} \oplus b, \tag{2.3}$$

where E is the photon energy, and a, b are parameters. The a term is known as the stochastic term due to the light generation, sampling, the statistical nature in photon detection and so on. The b term is called the constant term which takes into effects such as shower leakage and measurement systematics.

The photon resolution can be evaluated in data using the so-called *symmetric decays*, where either $\pi^0 \to \gamma\gamma$ or $\eta \to \gamma\gamma$ meson decays are used. By selecting a clean sample of π^0 or η mesons and requiring the two daughter photons have similar energies, the single-photon energy resolution can be related to the π^0 or η invariant mass resolution and the opening angle resolution, which may or may not be relevant depending on the experiment. The single-photon resolution obtained in this way can then be compared with Monte Carlo simulation, as shown in Fig. 2-17. More details on this analysis can be found in [28, 37].

2.3.5 Particle Identification System

The particle identification (PID) at GlueX Phase I is done primarily through energy loss, dE/dx, in the drift chambers, the Start Counter (SC) and the Time-of-Flight, as well as timing, with primarily the Time-of-Flight (TOF) hodoscope and the calorimeters. The DIRC detector was added for GlueX Phase II, which will be discussed extensively in Chapter 4.

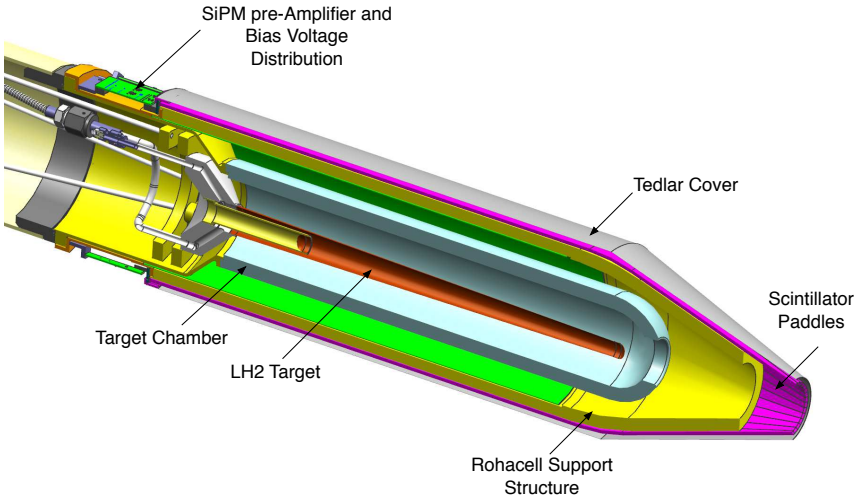

Figure 2-18: Schematic drawing of the Start Counter and target cell. Image source: [28].

Start Counter

The Start Counter (SC) is essentially an array of scintillating paddles surrounding the target cell arranged in an open cylinder shape. The SC consists of 30 scintillator paddle modules around the beamline. Each paddle is about 3 mm thick and about 600 mm in length. The paddles also taper down in the forward region to increase the acceptance. Each paddle is read out from the upstream end using SiPMs, which then connect to the read out electronics including the flash ADCs and TDCs. A schematic drawing of the SC is shown in Fig. 2-18.

The main purpose of the SC is to have high enough timing resolution to cleanly identify which of the 4 ns beam buckets initiated the reaction. Once the beam bucket is identified, the start time of any given reaction is determined from propagating the accelerator RF time from the diamond radiator to the target cell.[4] The timing resolution of the SC is a function of the hit position along the paddle with an average of about 230 ps, good enough to distinguish between different 4 ns tagged photon bunches. More detailed information about the SC can be found in [28, 40].

Time-of-Flight

The Time-of-Flight (TOF) detector is essentially a wall of scintillators located about 5.5 meters downstream from the target. The detector has two planes of scintillator paddles stacked in the horizontal and vertical directions. Most paddles, except a few

[4]Each electron bunch in the accelerator has a time spread of about 1 ps [39], which is much better than any other component in the GlueX system. The beam photon four-momentum vector, including the arrival time and the beam photon energy, is treated as known in the kinematic fitting.

Figure 2-19: Picture of the TOF wall. The transparent block near the beamline is made of lucite in order to reduce the count rate from the electromagnetic beam background in that region. Image credit: GlueX collaboration.

near the beamline, are 252 cm long, 6 cm wide, and 2.54 cm thick. Each paddle is read out by a PMT via a UV transmitting plastic light guide. A picture of the TOF wall is shown in Fig. 2-19.

The main purpose of the TOF is to measure the flight time of charged particles. Combined with the momentum measurement from the tracking, the flight time provides a measurement of the particle mass as a means to achieve particle identification. The average time resolution of the TOF reaches about 100 ps. This is good enough to provide π/K separation up to about 2 GeV/c, while the DIRC aims to extend this up to about 3.5 to 4 GeV/c. Fig. 2-20 shows the measured velocity β as a function of momentum p for positive tracks. We can clearly see the bands from protons, kaons, pions, and electrons; it is also clear that the pion and kaon bands merge after about 2 GeV/c. More information about the TOF can be found in [28].

2.4 Triggering, Data Acquisition, and Operation

This section provides a brief overview of how the raw analog signals from the various GlueX detector systems, such as those from the PMTs, SiPMs, and anode wires, are processed and eventually become the physically-meaningful quantities, such as the collection of four-momenta for all reconstructed particles in an event, available to the analysts for physics analyses. More detailed information can be found in [28]

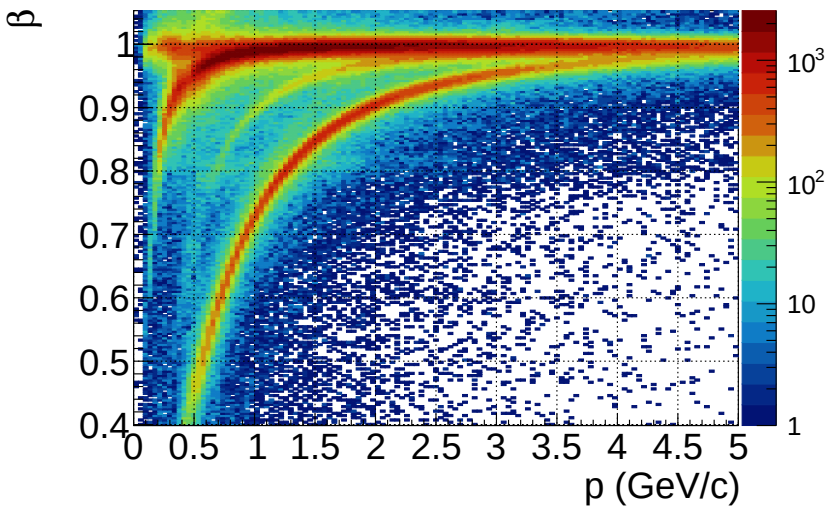

Figure 2-20: Particle velocity β of positive tracks as a function of track momentum p measured by TOF. Image source: [28].

and references therein.

2.4.1 Triggering and Data Acquisition

Flash ADC

One of the enabling technologies for the GlueX experiment and for many modern nuclear and particle physics experiments is the flash ADC. An analog-to-digital converter (ADC) is a device that digitizes an analog signal, such as an electrical current in the case of a PMT. A conventional old-school ADC often first performs an analog integration of the current using capacitors and then converts the charge into digital bits. The charge integration is interrupted if the ADC is being read out, resulting in dead time. The live time drops significantly past a few kHz and the ADC becomes useless [41]. A flash ADC (FADC) solves this problem in two ways. First, it samples the electrical current directly at a high rate and buffers the digitized current waveform in circular memory so that the integration can be done later. Second, it decouples the digitization of the data from the readout. As long as the readout signal, *e.g.,* from a trigger decision, is made within the length of the circular memory buffer, only the selection regions in the buffer are read out and this readout process is independent of the digitization. Fig. 2-21 is an illustration of the working principles of a 250 MHz FADC used in many subdetector systems at GlueX.

Trigger Types

The photons are incident on the target every 4 ns producing physics events of interest and low energy electromagnetic interactions in the detector. The detector is also under constant bombardment of cosmic rays. We cannot record all the events, but we also

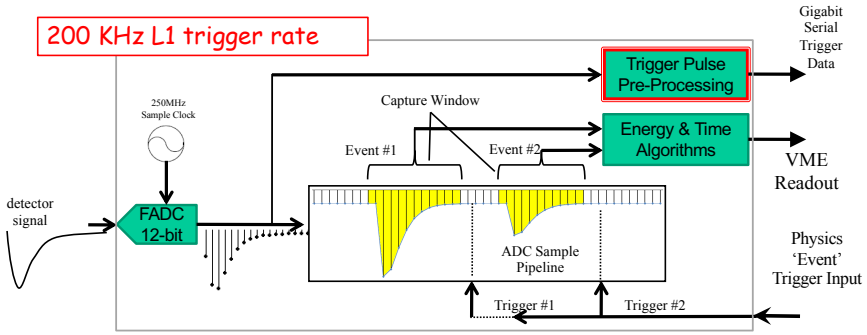

Figure 2-21: Graphical illustration of the working principles of a flash ADC. Image credit: [42].

do not need to because most of those events are not interesting. We need to *trigger* on signatures that are deemed potentially interesting and record those. At GlueX, a so-called Level-1 hardware trigger is used. The nominal physics triggers at GlueX are from combinations of the calorimeter energies if one of the following conditions is satisfied:

1. $2 \times E_{\mathrm{FCAL}} + E_{\mathrm{BCAL}} > 1\,\mathrm{GeV}$ and $E_{\mathrm{FCAL}} > 0$; or
2. $E_{\mathrm{BCAL}} > 1.2\,\mathrm{GeV}$.

The first condition uses the fact that most events produce energy in the forward region due to the fixed target set-up of GlueX and the second condition aims to capture events with large transverse energy deposition in the BCAL, such as J/ψ meson decays. There are several other specialized triggers using information from the PS, TAGH/TAGM, or monitoring LEDs. There is also a random trigger that records the activity in the detector when it is not associated with any physically meaning events as a means to assess the background activity in the detector system.

Trigger Architecture

A schematic layout of the GlueX trigger architecture is shown in Fig. 2-22. Detector signals come in to the front-end crates where they are digitized by the FADCs. For each channel for the BCAL and the FCAL,[5] the FADC amplitude is summed every 4 ns (for the 250 MHz FADCs). The crate processor module (CTP) sums the amplitudes from all the relevant FADCs and transmits to their sub-system processor (SSP) module in the Global Trigger Crate where the energy is computed from the digitized amplitudes. The global trigger processor (GTP) module collects the relevant SSP module results and makes a trigger decision as described above. The trigger decision is transmitted to trigger supervisor (TS) module located in the Trigger Distribution Crate. The TS then distributes the trigger signal to all trigger distribution (TD) modules, which then sends the trigger signal back to trigger interface (TI) module

[5]Only BCAL and FCAL signals are used because we describe the main physics triggers here. The other trigger types follow a similar process.

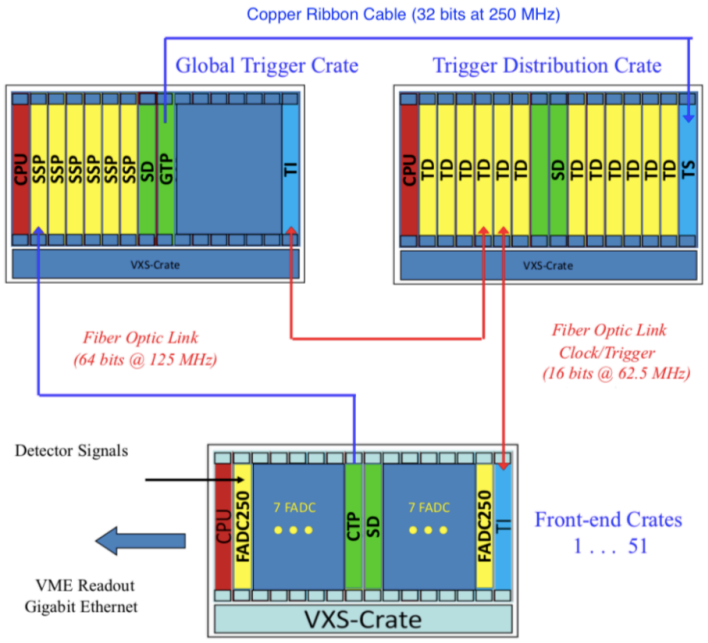

Figure 2-22: Schematic layout of the trigger architecture at GlueX. Image source: [28].

in the front-end crates. The trigger signal received by the TI initiates the readout process, where relevant chunks of digitized amplitudes are read out from the stored buffer of the FADCs to the data acquisition system.

Data Acquisition

The data acquisition (DAQ) layout and data flow are shown in Fig. 2-23. The data from the detector front-end crates are read out by the readout controllers (ROC), which are single-board computers. The ROCs transfer data to the data concentrators (DC) which are programs that build partial events from a subset of the crates.[6] The event builder (EB) then builds the whole events from data from all the detector crates. The built events are then written to local storage disks by an event recorder (ER). The ER also communicates with a separate network which transmits data to an online computing farm for monitoring, thus decoupling the DAQ and monitoring networks. More detailed information about the trigger and data acquisition system can be found in [28] and references therein.

2.4.2 Reconstruction and Monitoring

Event reconstruction is the process in which the raw detector data, recorded through triggering and DAQ systems described above, are processed with dedicated detector-

[6]Event-building refers to the process where raw detector data from one physics event are collected, associated, and formatted.

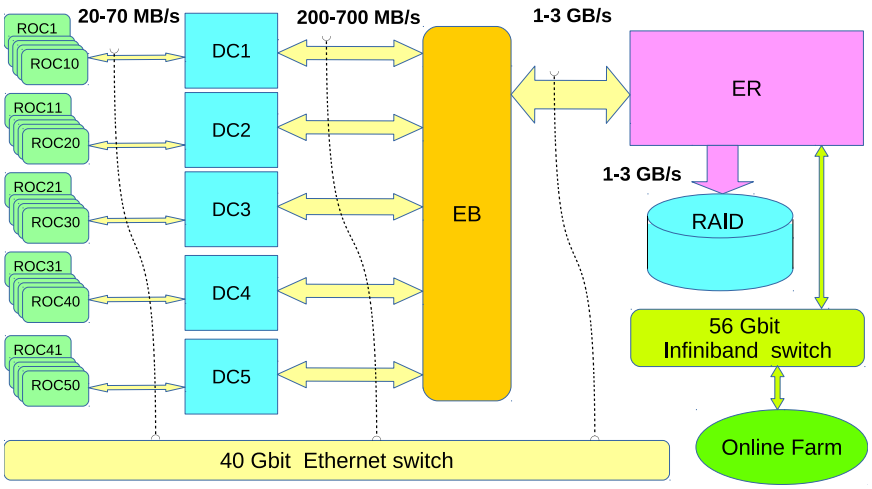

Figure 2-23: Schematic layout of the data acquisition system showing the data flow at GlueX. Image source: [28].

dependent algorithms into physically meaningful quantities. We briefly described how particles are reconstructed in Sec. 2.3. An overview of the data production and monitoring flow is shown in Fig. 2-24. In addition to the main data production flow, there is also a monitoring flow, where a subset of the raw data is reconstructed and monitored both in almost real-time during data taking and with $\mathcal{O}$(hours) delay by subdetector experts to ensure the detector is running and the data are being recorded properly. This subset of raw data is also used to perform calibration tasks of the subdetectors, often resulting in calibration parameters stored in a calibration database. Those calibration parameters are then used in the reconstruction of the full datasets, which are then further processed to produce the so-called Analysis Trees which are manageable in size, *e.g.,* $\mathcal{O}$(TB), on which analysts can perform their desired physics analyses.

2.4.3 Monte Carlo Simulation

In addition to the actual data, Monte Carlo simulation is an important tool to help understand the collected data and provide information which can be difficult to measure in data. A schematic diagram of the Monte Carlo simulation workflow is shown in Fig. 2-25. The Monte Carlo simulation is done in the following steps: event generation, detector simulation, detector response simulation, and reconstruction. A brief description of each of the steps is provided here with more detailed information found in [28] and references therein.

Event Generation

Event generation is often the first step in the simulation workflow done by programs known as event generators. They simulate specific physics processes of interest and produce a list of particles coming out of the collision. At GlueX, there are event gen-

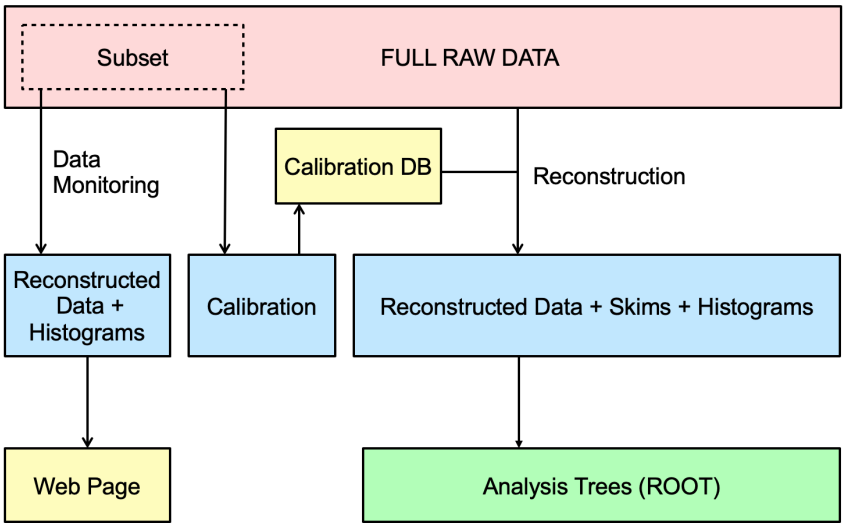

Figure 2-24: Overview of GlueX data production and monitoring. Image source: [28].

erators that simulate specific reactions of interest. For instance, the `gen_omega_3pi` generator simulates the reaction $\gamma p \to p\omega, \omega \to \pi^+\pi^-\pi^0$ including the full angular distribution information of final state particles. There is also a Pythia 6 based total photoproduction generator `bggen`, which simulates all photoproduction processes according to their corresponding cross sections. It is used to study backgrounds and develop analysis tools for signal extraction. This step is in principle general and not experiment-dependent.

Detector Simulation

Once the physics events are generated, they are passed to a full detector simulation. At GlueX, this is done via a Geant4-based tool `hdgeant4`. It simulates the interactions of the produced particles with the detector material and detection elements, *e.g.,* anode wires, SiPMs, and PMTs. The geometry specification, including the physical dimensions and locations of the detector elements and the material properties of relevant detector volumes, is done through an XML-based tool called Hall D Detector Specification, or `hdds`. The output of this step is often energy deposition and time information in the detector elements of interest.

Detector Response Simulation

This step converts the time and energy deposition information from the detector simulation step into electronic signals that match the characteristics of the readout from the actual experiment. This process is sometimes known as *smearing* or *digitization*, and is done via the tool `mcsmear` at GlueX. The output of this step is digitized signals that would have been produced by the detector, with the addition of the so-called *truth information* which are the raw time and energy deposition informa-

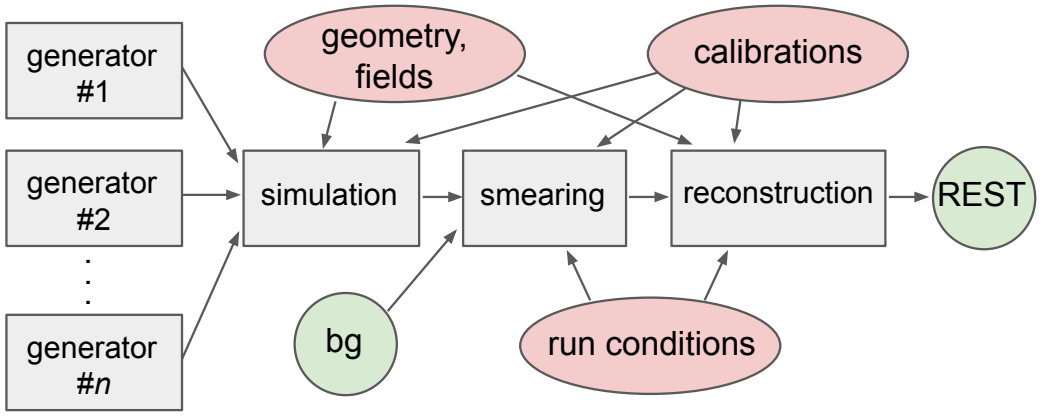

Figure 2-25: Schematic diagram of the Monte Carlo simulated data flow at GlueX. Image source: [28].

tion from the previous step. At this step, relevant experiment-related information is included to ensure that the detector response matches the experimental conditions. Additionally, experimental data events from random triggers are also included at this step to account for the background noise that might be present in data but not in the simulation.

Reconstruction

The reconstruction of the simulated data is done in exactly the same way as that for data. At this step, relevant experimental condition-related information is again included. There is often weak dependence of the reconstruction on the reconstruction software used. In simulation, the exact same software stack that is used to reconstruct the data is used to reconstruct the simulated dataset.

Chapter 3

Event Generator Tuning Using Bayesian Optimization

The phenomenology of the strong interaction at low energies is challenging to describe analytically. Models have been developed and implemented in Monte Carlo event generators to enable the simulation from partons produced in high energy hard scatterings to hadrons observed in experimental apparatus. Since these models are not derived from first principles, they often contain parameters that need to be matched to experimental data. This is known as the *event generator tuning problem*, which we discuss in more detail in Sec. 3.1. It is a challenging computation problem because large amounts of Monte Carlo events need to be generated and compared with data. We present a novel approach to this problem using Bayesian optimization, whose general principles are discussed in Sec. 3.2. We present a proof-of-concept closure study of this approach in Sec. 3.3 and discuss aspects regarding the application of this method to real-world problems in Sec. 3.4.

3.1 Event Generator Tuning Problem

3.1.1 The Need for Tuning

Monte Carlo event generators are widely used tools to simulate particle collisions and interactions. They incorporate our understanding of the underlying physics and use Monte Carlo methods to simulate the manifestations of the underlying physics. They often contain a large number of parameters that must be determined, or *tuned*, by comparing the output of the generator with experimental data. This is due to our imperfect knowledge of the underlying physics, often those involving the strong interaction. I will use the simulation of high energy e^+e^- collisions as an example for the purpose of this discussion. The input is almost as simple as it can be because it only involves elementary particles, *i.e.,* electrons and positrons. The hard scattering interactions they participate in are well described in the Standard Model and their matrix elements can be calculated with high precision. In the case of final states involving a quark-antiquark pair, *i.e.,* $e^+e^- \to \gamma^*/Z \to q\bar{q}$, (anti-)quarks which are

colored come out of the hard scattering. They then undergo further splitting into quarks and gluons, in the process known as *fragmentation*, and eventually the colored quarks and gluons combine to form hadrons which are colorless, in the process known as *hadronization*, that can then be detected by the experimental instruments. Both the fragmentation and hadronization processes require the modeling of the strong interaction in the non-perturbative regime where analytical first-principle descriptions do not exist, and instead various phenomenological models have been developed to achieve this goal. One such model is the Lund string model [43], which is the fragmentation model used in Pythia [44, 45]. As a concrete example, the scaling function $f(z, m_T)$ that describes the probability for a quark to combine with an antiquark of the same flavor to give a meson with transverse mass m_T and energy-momentum fraction z to yield symmetric fragmentation has the following form [43]

$$f(z, m_T) = N\frac{1}{z}(1-z)^a \exp\left(b\frac{m_T^2}{z}\right), \tag{3.1}$$

where N is a normalization constant and a, b are two (highly correlated) free parameters that need to be tuned to data. They correspond to the `StringZ:aLund` and `StringZ:bLund` parameters in Pythia [45]. There are many more such parameters, often related to fragmentation and hadronization, that need to be compared with experimental data.

3.1.2 Existing Solutions

Generating enough events with a fixed set of parameter values to enable making a comparison between the simulated data and the experimental data is extremely CPU intensive. For example, it takes $\mathcal{O}$(hour) on a modern CPU core to generate 1M events for a single set of parameter values using the Pythia 8 event generator [44, 45]. A full tune of Pythia to e^+e^- data involves optimizing ≈ 20 parameters, which clearly cannot be performed using a brute-force grid-based approach. Even a tune of only a small subset of parameters, *e.g.*, the 6 parameters that control fragmentation, would take $\mathcal{O}(100)$ CPU years using a coarse 10-bins-per-parameter scheme.

All available tunes provided with the Pythia 8 package were obtained either manually or parametrically.

Manual Approach

The manual approach, as the name suggests, requires an expert to manually choose how to vary the parameters based on extensive knowledge and insight, guided by comparing generated and experimental distributions. We believe that only a few physicists are capable of performing a large-scale manual tune of Pythia, and even for such an expert it takes considerable effort. The manual approach does not scale to larger parameter sets, and is not well suited to less-intuitive models or to producing many experiment-specific tunes (or a large number of dedicated tunes in general). That said, the manual approach is less prone to finding an *unphysical* local minimum

in the parameter space that happens to provide a decent description of the data distributions being compared to during the tuning process; the expert can intuitively identify such situations. The default tune, known as the *Monash tune* of PYTHIA 8 was obtained this way as described in detail in [46].

Parametrical Approach

In the parametrical approach, the generator response to changes in the parameters is itself parameterized based on a large set of reference generator data sets, which then facilitates optimizing the parameters via minimization of an objective function, *e.g.,* a χ^2. For instance, in [47], the histogram bin content was parameterized with respect to the parameters under consideration with some functional form. The parametrization approach is easily parallelized, but requires that the generator response—including multi-parameter correlations—is well approximated by the chosen parametric function within the parameter hypercube to be explored. Furthermore, the optimal working point must be included in the parameter hypercube, though this can be achieved by first doing a coarse scan of the parameter space.

In this study [48], we propose treating Monte Carlo event-generator tuning as a black-box optimization problem to be addressed using the framework of Bayesian optimization. We choose to use the SPEARMINT package [49, 50] (a python implementation of Bayesian optimization) in our study and document the code in [51].

3.2 Bayesian Optimization

Bayesian optimization (BO) is a sequential model-based approach for global optimization. It is particularly advantageous for problems where the function is a black-box and expensive to evaluate. Many excellent reviews on BO exist in the literature, see *e.g.,* [52, 53, 54], and we summarize the main features of BO here.

Bayesian optimization is a class of optimization methods aiming to solve the problem

$$x^* = \arg\max f(x), \text{ for } x \in A, \tag{3.2}$$

where the input x is in $\mathbb{R}^d$, and $f(x)$ is the function, and x^* is the solution that maximize (or minimize) the function $f(x)$. The following characteristics of the function $f(x)$ can make the problem particularly suitable for Bayesian Optimization:

- $f(x)$ is a black-box function: no parametric form of $f(x)$ exists, nor is any special structure such as concavity or linearity known about the function;
- $f(x)$ is derivative-free: we only observe the function value of $f(x)$, but no first- or second-derivatives;
- $f(x)$ is expensive to evaluate: this often means that each evaluation of $f(x)$ takes a substantial amount of time, *e.g.,* minutes to hours (given the computing resources available), limiting the the number of evaluation calls of $f(x)$ to a few

hundred.[1]

Other characteristics in a problem suited for BO include: (1) $f(x)$ is a continuous function; (2) the dimension d of the input space A is not too high, typically $d \leq 20$ in most successful applications of BO [52]; (3) the evaluation of $f(x)$ may have noises. Conversely, if some of the above characteristics are absent, there likely exists other, potentially better, optimization strategies. Algorithm 1 shows the pseudocode of a Bayesian optimization loop while Fig. 3-1 shows a graphical illustration of the process. As illustrated in Algorithm 1 and Fig. 3-1, a Bayesian optimization procedure consists of two main components: a *statistical model* of the objective function f, and an *acquisition function* for deciding where to sample next. In essence, BO belongs to the class of optimization methods known as *surrogate methods*: it maintains a surrogate that models the objective function, from which an acquisition function is used to choose where to evaluate. It distinguishes itself from other surrogate methods in that the surrogate is built and interpreted in Bayesian ways.

Algorithm 1: Pseudocode for Bayesian Optimization loop.

Query function f for n_0 points as the initial dataset
$\mathcal{D}_n = \{\{x_1, y_1\}, \{x_2, y_2\}, ..., \{x_{n_0}, y_{n_0}\}\}$ (setting $n = n_0$);
while *stopping criteria not met* **do**
    Update the statistical model of f using the available dataset $\mathcal{D}_n$;
    Select the next point x_{n+1} according to the acquisition function;
    Query the function f to obtain $y_{n+1} = f(x_{n+1})$;
    Augment the observed dataset $\mathcal{D}_{n+1} = \{\mathcal{D}_n, \{x_{n+1}, y_{n+1}\}\}$;

3.2.1 Gaussian Process

The statistical model assigns a prior probability distribution over the unknown objective function. This is often done using a Gaussian Process (GP), which assumes that the set of n observations forms a multinormal distribution (see *e.g.,* [55] for a detailed discussion):

$$f(x_{1:n}) \sim \text{Normal}(\mu_0(x_{1:n}), \Sigma_0(x_{1:n}, x_{1:n})), \tag{3.3}$$

where $x_{1:n}$ indicates the input vector $[x_1, ..., x_n]$, the function value vector $f(x_{1:n}) = [f(x_1), ..., f(x_n)]$, the mean vector $\mu_0(x_{1:n}) = [\mu_0(x_1), ..., \mu_0(x_n)]$, and the covariance matrix $\Sigma_0(x_{1:n}, x_{1:n}) = [\Sigma_0(x_1, x_1), ..., \Sigma(x_1, x_n); ...; \Sigma_0(x_n, x_1), ..., \Sigma_0(x_n, x_n)]$. The μ_0 function is called the mean function, and Σ_0 is called the covariance function, or kernel, for each point pair (x_i, x_j). One can compute the posterior probability distri-

[1] This expensiveness can also come from the literal monetary cost associated with each evaluation, *e.g.,* purchase of materials in an engineering design optimization problem.

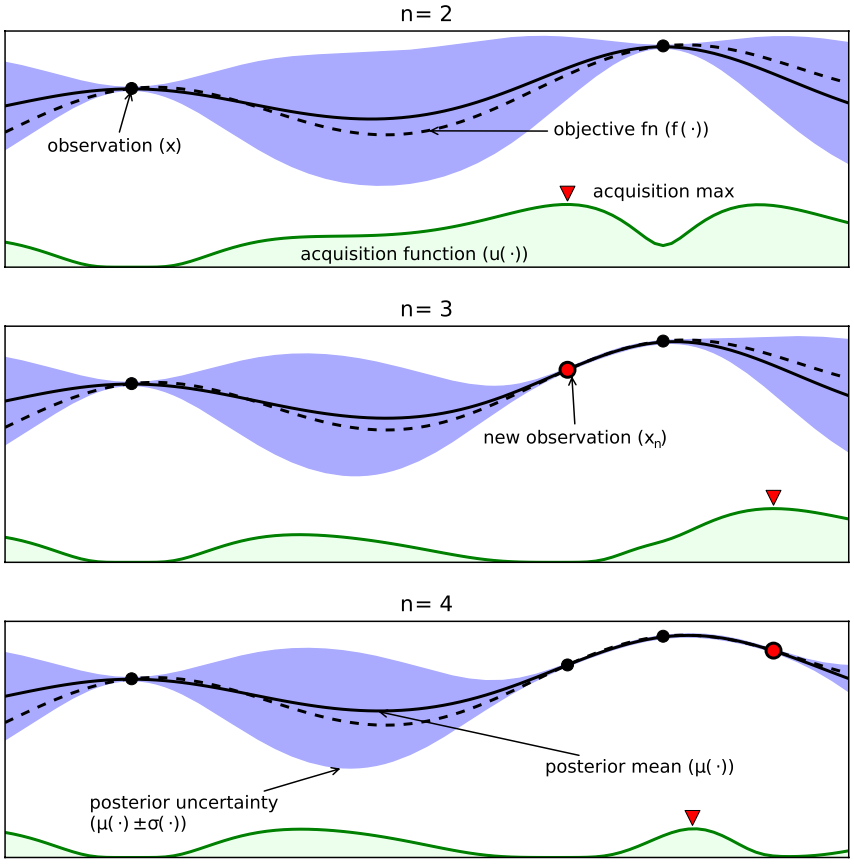

Figure 3-1: Illustration of the Bayesian optimization process over three iterations, showing the (dots) observations, (dashed black) the mean and (violet) confidence intervals of the statistical model, and (green) the acquisition function. The objective function is also shown for illustrative purposes although it is unknown in practice. Image source: [54].

bution as a conditional distribution of f at a new point x, in closed form, as [54, 55]

$$\begin{aligned} f(x)|f(x_{1:n}) &\sim \text{Normal}(\mu_n(x), \sigma_n^2(x)), \text{ where} \\ \mu_n(x) &= \Sigma_0(x, x_{1:n})\Sigma_0(x_{1:n}, x_{1:n})^{-1}(f(x_{1:n}) - \mu_0(x_{1:n})) + \mu_0(x), \\ \sigma_n^2(x) &= \Sigma_0(x, x) - \Sigma_0(x, x_{1:n})\Sigma_0(x_{1:n}, x_{1:n})^{-1}\Sigma_0(x_{1:n}, x). \end{aligned} \tag{3.4}$$

A common choice of the mean function is a constant function $\mu_0(x) = \mu$, because it acts somewhat like an overall normalization. What really matters is $\mu_n(x)$, which is determined by the existing data and the covariance matrix, as shown in Eq. (3.4). The covariance function, or kernel, typically has the property that points closer in the input space should be more strongly correlated. A common choice is the Gaussian kernel

$$\Sigma_0(x, x') = \alpha_0 \exp(-||x - x'||^2), \tag{3.5}$$

where $||x-x'||^2 = \sum_{i=1}^{d} \alpha_i (x_i - x_i')^2$ and α_i are hyperparameters which control the "correlation distance", and can be integrated internally in the posterior calculation [54]. As shown in Fig. 3-1, the model has minimal uncertainty at the observed points and the confidence level bands grow wider as one moves further away from the observed points.

3.2.2 Acquisition Function

The acquisition function specifies the process of how to obtain the next point to sample based on the posterior distribution. There are many choices available [54] and we briefly discuss the expected improvement (EI) method, which is a common choice and also the default implementation in Spearmint. The expected improvement, at point x given n previous observations, is defined as

$$\text{EI}_n(x) := E_n[[f(x) - f_n^*]^+], \tag{3.6}$$

where $E_n[\cdot] = E[\cdot|x_{1:n}, f(x_{1:n})]$ indicates the expectation over the posterior distribution given observed points, $f_n^* = \max f(x_{1:n})$ is the observed maximum, and $[f(x) - f_n^*]^+ = \max(0, f(x) - f_n^*)$. The next point to sample x_{n+1} can then be defined as the point that maximizes the expected improvement

$$x_{n+1} = \arg\max \text{EI}_n(x). \tag{3.7}$$

The EI can be solved analytically in a closed form and the EI as well as its first- and second-derivatives are inexpensive to evaluate; therefore, the maximization can be implemented using gradient-based techniques [54]. The EI depends on both the distance of the input points and the variance at the point of interest, as illustrated in the examples in Fig. 3-1. This tradeoff between the high expected performance (in the case of $n = 4$ in Fig. 3-1) and the high uncertainty (in the cases of $n = 2$ and $n = 3$ in Fig. 3-1) is known as the "exploration vs. exploitation tradeoff."

We choose to mostly utilize the default implementation of the Bayesian optimization in Spearmint in this study and discuss possible directions for future improve-

ment in [48].

3.3 Tuning Pythia

We perform a closure test to demonstrate how to apply Bayesian optimization to Monte Carlo event generator tuning and to validate its performance. The full description of the study is found in [48] and we summarize the study and the important findings in this thesis and refer the interested reader to [48].

3.3.1 Closure Test Overview

The basic set-up of the closure test is outlined below:

- a 10M-event e^+e^- data sample is generated using Pythia 8 with its default parameter values, collectively referred to as the *Monash* tune [46];
- various observable distributions are built from the Monash simulated data sample and treated as experimental data;
- a set of 20 parameters in Pythia is chosen for tuning;
- each parameter is allowed to vary freely within a large pre-defined range and the true Monash values are treated as unknown;
- and, finally, the Bayesian optimization framework is applied using Spearmint to obtain the 20 optimal (tuned) parameter values.

Treating the Monash data sample as experimental data permits validating the performance by comparing the Monash parameter values to the optimal ones found by Spearmint. This treatment ensures that each distribution can be perfectly modeled by Pythia. Closure is achieved if the Bayesian optimization procedure can successfully recover the Monash parameter values. In reality, Monte Carlo event generators often cannot model experimental data perfectly. Considerations of applying this method to real-world tunes are briefly outlined in Sec. 3.4 and discussed in more detail in [48].

Bayesian optimization is a sequential optimization strategy. First, observable distributions (see Sec. 3.3.3) are constructed from the 10M-event pseudo-experimental data using Monash parameters. Then the optimization loop proceeds as follows:

- for each query, the optimizer provides a set of parameter values which are passed to Pythia 8 and used to generate 1M events;
- once the Pythia sample is generated for each parameter set, the observable distributions are constructed and used to calculate the objective function (see Sec. 3.3.2), which is provided to the optimizer and used to update its internal model from which the next set of parameters to query is determined and passed to Pythia;

- the query steps are repeated until some chosen stopping criterion (see Sec. 3.3.4) is met.

Upon exiting the optimization loop, the optimal parameter set is evaluated. In our study, the optimal parameter set is taken to be the one that the optimizer's internal model predicts is the best, and not the set for which a PYTHIA data sample was generated and found to have the best objective function value (in practice this makes little difference). Partial scans of the parameter profile using the optimizer's internal model are used to assign error bars to parameters whose details, along with other technical aspects of the tuning procedure are given in [48].

In the study, we considered two approaches to tuning the 20 parameters: a block-diagonal strategy and a global strategy. We use results from block 3 (related to hadron-type) as an example for discussion in this thesis, summarize the main findings, and refer the interested reader to [48] for more details regarding the other blocks and the global tune.

3.3.2 Objective Function

We define our objective function as a pseudo-χ^2 in a similar way to the one used in producing the Monash tune [46]:

$$\chi^2 \equiv \sum_{i=1}^{n_{\text{bins}}} \frac{(\text{Monash}_i - \text{MC}_i)^2}{\sigma^2_{\text{Monash},i} + \sigma^2_{\text{MC},i}}, \tag{3.8}$$

where σ denote the statistical uncertainties on the Monash and Monte Carlo values in the i^{th} bin. Any σ_i value that corresponds to less than a 1% relative uncertainty is set to be 1%. The choice of setting a minimum value avoids having a few bins with large occupancies dominating the tuning. In practice, the systematic uncertainties on the experimental distributions implicitly accomplish this. The sample size of each Monte Carlo data set is chosen such that the 1% value is used in most bins. We ignore correlations between bins in our definition of χ^2, since this information is often not available for experimental data. It may be desirable to alter the χ^2 to include weight factors for each bin. Incorporating correlations or bin weights into the tuning procedure is straightforward, as only the definition of the χ^2 needs to be modified.

3.3.3 Parameters and Observables

We choose to tune a large set of 20 continuous parameters, which roughly corresponds to the full set of PYTHIA 8 parameters constrained by the observable distributions from e^+e^- data that were used in the Monash tune, and that enter into the χ^2 defined in Eqn. (3.8). The full list of parameters included in the tune can be found in [48]. We place a uniform prior over each parameter within the specified range, *i.e.,* parameters are allowed to vary freely within these ranges.

Since the goal of the study was to demonstrate the power of the Bayesian optimization process, we chose to use minimal expert knowledge. In addition, the physical meaning of the parameters and the observable distributions is also not important for

Table 3.1: Tuning results for block 3.

Parameter	Monash Value	Tune Value	Range Considered
probStoUD	0.217	$0.219^{+0.001}_{-0.002}$	$[0, 1]$
probQQtoQ	0.081	0.082 ± 0.01	$[0, 1]$
probSQtoQQ	0.915	$0.892^{+0.014}_{-0.018}$	$[0, 1]$
probQQ1toQQ0	0.0275	0.0276 ± 0.0009	$[0, 1]$
etaSup	0.6	0.59 ± 0.02	$[0, 1]$
etaPrimeSup	0.12	0.12 ± 0.01	$[0, 1]$
decupletSup	1	$1^{+0}_{-0.04}$	$[0, 1]$
mesonUDvector	0.5	$0.51^{+0.01}_{-0.02}$	$[0, 3]$
mesonSvector	0.55	0.55 ± 0.01	$[0, 3]$
mesonCvector	0.88	$0.89^{+0.04}_{-0.05}$	$[0, 3]$
mesonBvector	2.2	2.1 ± 0.1	$[0, 3]$

the study. In this thesis, I provide a concrete example for completeness and will omit further detailed description of the parameters and the observable distributions, and refer the interested reader to [46] and [45]. One of the parameters in block 3 is `probStoUD`. It belongs to the `StringFlav` class, which handles the choice of a new flavor in the fragmentation process and the production of a new hadron from a set of input flavors. As the class name suggests, it follows from the string fragmentation model [45]. Specifically for `probStoUD`, it handles "the suppression of s quark production relative to ordinary u or d" production as stated in the PYTHIA user manual [45]. One would expect that it affects the relative abundance of the strangeness-containing hadrons to those that do not contain strange quarks. The abundances of various types of identified hadrons are therefore included in the observable distributions for block 3, as shown in Fig. 3-2.

3.3.4 Results and CPU Usage

The results of tuning the block 3 parameters on the hadron-type distributions are presented in Figs. 3-2 and 3-3, and in Table 3.1. The block 3 tune involves 11 parameters (largest among the three blocks considered), and we find that all observable distributions and PYTHIA parameter values are consistent with Monash. We observe similarly good agreement for the other two blocks. We also performed a global tune of all 20 parameters. While the global approach provides no improvement for this particular example as expected, the fact that performing a 20 parameter tune is possible is both novel and exciting. Detailed results from the study and further discussions are found in [48].

The CPU cost of performing these tunes depends on how many queries are made by the optimizer SPEARMINT; therefore, determining when to terminate the optimization process governs how much total CPU is required. Fig. 3-4 presents the evolution of the SPEARMINT model χ^2 value versus query number for block 3. The SPEARMINT model χ^2 converges to a value close to the mean χ^2 value under the null hypothesis

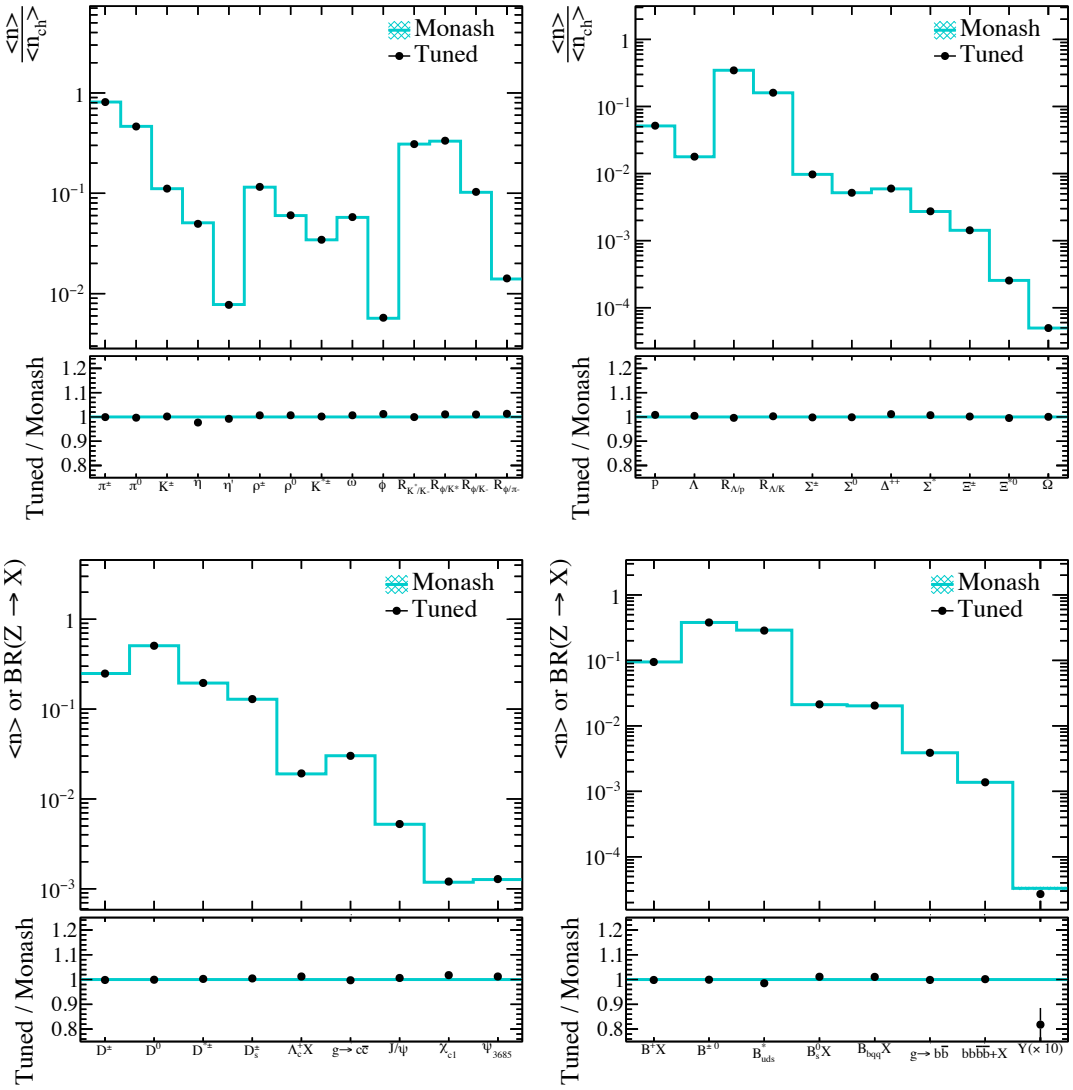

Figure 3-2: Hadron-type distributions obtained from the Monash pseudo-data sample compared to those obtained from the optimal tune of the parameters in block 3. Both samples used here have 10M events.

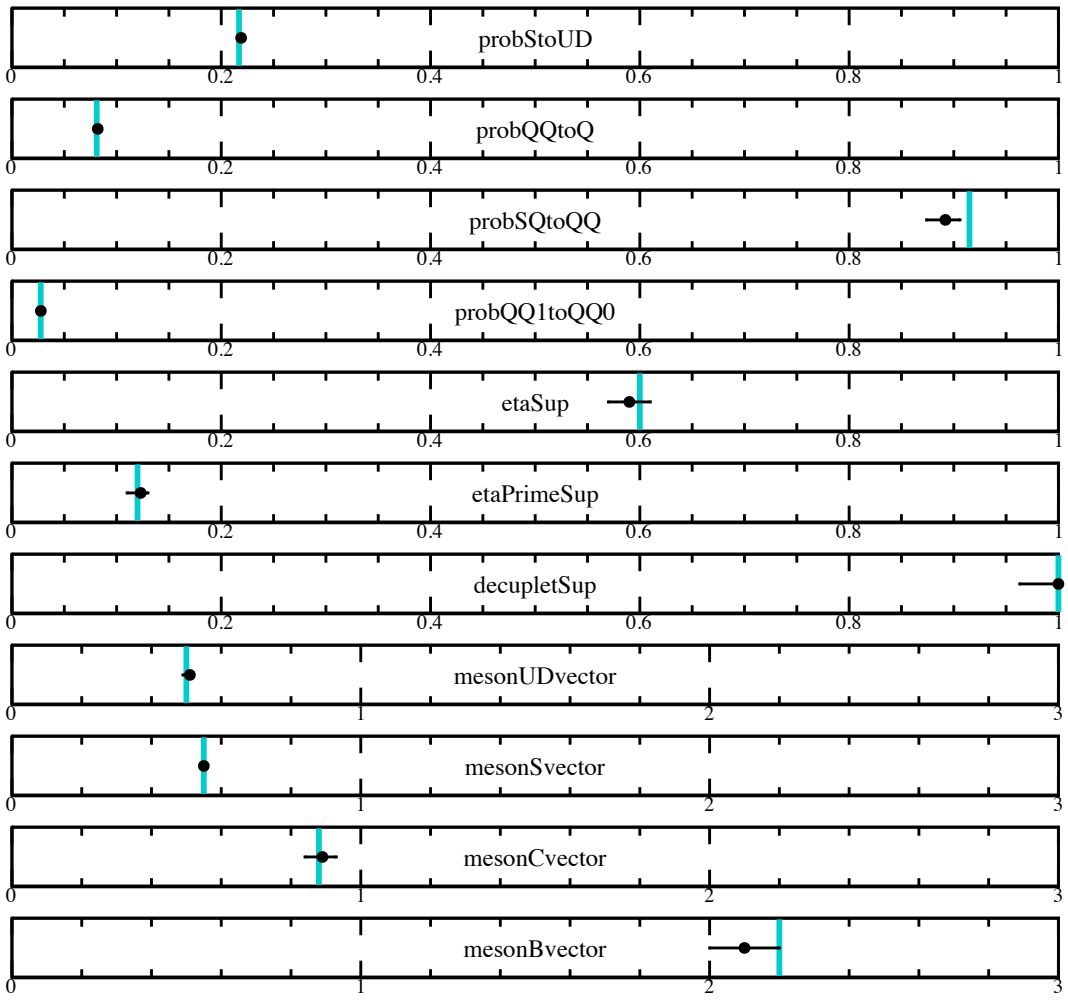

Figure 3-3: (black points) Block 3 parameters from our optimal tune compared to their (vertical cyan lines) Monash values. The horizontal-axis ranges are the regions considered by the optimizer during tuning.

(see [48] for details on how the mean value is obtained); *i.e.,* the SPEARMINT model χ^2 converges to the mean χ^2 value expected using the true parameter values. Therefore, by first computing the null mean χ^2 using a Monte Carlo data sample constructed to have per-bin errors that match the experimental data distributions—where the true parameters are known—it is possible to obtain an estimate of what the SPEARMINT model χ^2 value should converge to. Figure 3-4 shows that the SPEARMINT model χ^2 value is unstable until about $25 \cdot n(\text{par})$ queries are made, and that each block[2] has fully stabilized by $50 \cdot n(\text{par})$ queries ($n(\text{par})$ denotes the number of parameters being tuned). This also holds true for the parameter values. Based on the results of the tunes performed in this study, we devised several stopping criteria [48].

The wall time required to perform these tunes on a quad-core i7 2.8 GHz 2015 Macbook Pro laptop is about 6, 14, and 25 hours for blocks 1, 2, and 3, respectively.[3] In total, 45 hours of wall time is required to perform the full 20 parameter block-diagonal e^+e^- tune; therefore, a full e^+e^- tune of 20 PYTHIA parameters can be performed on a laptop in less than 2 days using SPEARMINT. The event-generation processes dominate the total CPU required to tune each block. Since event generation is trivial to do in parallel, the tunes of each block could be performed much faster using more computing power. This is one of the novel and exciting aspects of this study.

Bayesian optimization implementations like SPEARMINT are not designed for the case where the parameters are too large, although we do observe good performance for the 20 parameter tune. Further improvements of both this study and the Bayesian optimization implementations are discussed in [48].

3.4 Towards a Real-World Tune

There are a number of aspects present in a real-world tune that are absent from the closure test in our study. Some of them are straightforward to include within the framework proposed above, including the inclusion of expert knowledge, weighting of observable distributions, and the inclusion of multiple beam energies and different types of particle collisions (*e.g.,* proton-proton collisions); see [48] for details. The solutions mostly involve the modification of the objective function or the sampling strategy, which are both straightforward to incorporate in the proposed framework. Such solutions likely require more computing resources but these problems factorize from the optimization procedure itself. In my view, the main difference between the closure test and a real-world tune is the fact that the existing Monte Carlo model almost certainly does not capture all the features present in the real data. Even once the optimal parameter values are found, one still expects that discrepancies between the Monte Carlo and data will remain. First, we argue that discovering such

[2]Only results from Block 3 are shown in this thesis. Results from the other two blocks can be found in [48].

[3]The whole study was performed on my personal laptop. It has 8 virtual cores. We run SPEARMINT on one core, and PYTHIA event generation is performed in parallel on the remaining 7 cores.

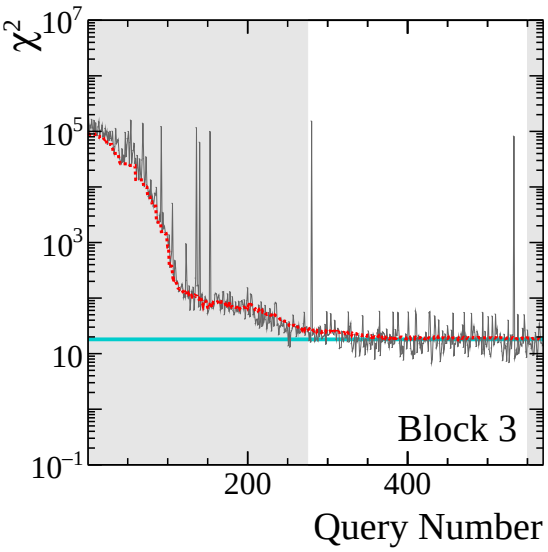

Figure 3-4: The χ^2 value for each SPEARMINT query obtained using (black) the PYTHIA sample produced for the current query and (red) the SPEARMINT χ^2 model. The white regions show $25 \cdot n(\text{par})$–$50 \cdot n(\text{par})$ queries for block 3.

situations as quickly as possible should be viewed as one of the goals of the parameter-optimization process because the output of the optimization process, when performing as expected, should be the set of parameters that describes the data as well as the Monte Carlo model allows. In [48], we performed a proof-of-principle tune with real data. The tune converges in about the same number of queries as in the closure test and demonstrates successful application of the optimization process to experimental data. However, we argue that some care is needed for interpretation and that expert knowledge is critical to inform the next steps in action. Regardless, by combining the Bayesian optimization approach with expert-level knowledge, it should be possible to produce better tunes in the future by making it much faster and easier to both optimize the generator parameters and to study discrepancies between Monte Carlo and experimental data.

Chapter 4

The DIRC Detector at GlueX

The GlueX DIRC detector was proposed to upgrade the particle identification capability of the GlueX experiment in order to fully exploit its physics potential, as discussed in Sec. 1.2. The acronym DIRC stands for Detection of Internally Reflected Cherenkov light. We present the general principles of Cherenkov radiation and the DIRC technology in Sec. 4.1. The next sections present various aspects of the realization of the GlueX DIRC project to which I made contributions: the transport of the BaBar bar boxes in Sec. 4.2, the construction of the optical box and commissioning of the full DIRC detector in Sec. 4.3, the simulation and reconstruction of the detector in Sec. 4.4 and Sec. 4.5, respectively, and the calibration efforts in Sec. 4.6.

4.1 Introduction to DIRC

4.1.1 Cherenkov Radiation and Particle Identification

Cherenkov Radiation

Cherenkov radiation is the electromagnetic radiation emitted when a charged particle travels through some dielectric medium at a speed greater than the phase velocity of light in that medium. It is named after Soviet physicist Pavel Alekseyevich Cherenkov, who later shared the Nobel Prize in Physics in 1958 with Soviet physicists Ilya Mikhailovich Frank and Igor Yevgenyevich Tamm "for the discovery and the interpretation of the Cherenkov effect." [56, 57, 58]. Intuitively, the emission mechanism of Cherenkov radiation can be understood in the following way [59]. When a charged particle passes through the dielectric medium, it polarizes the medium along its path. The time variation of the dipole field leads to the emission of electromagnetic radiation. When the particle speed is small, $v < c/n$, the polarization is symmetric and the overall dipole field vanishes. When the particle speed is large enough, $v > c/n$, the polarization becomes asymmetric and constructive interference from the dipoles occurs, resulting in the characteristic cone-shaped radiation wavefront, as illustrated in the cartoon in Fig. 4-1.

There are two important relations regarding Cherenkov radiation. The first is

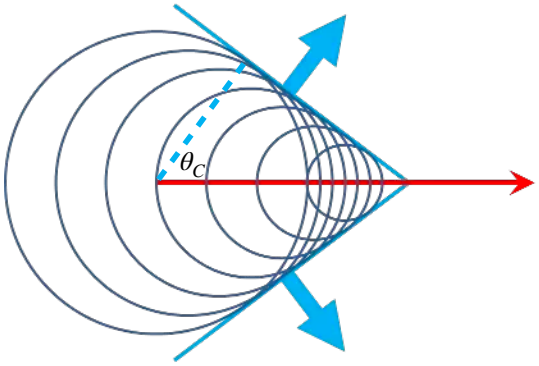

Figure 4-1: Schematic figure showing Cherenkov radiation. The red arrow shows the charged particle path; blue arrows indicate the Cherenkov radiation wavefront; θ_C is the characteristic Cherenkov angle between the charged particle path and the emitted Cherenkov photons.

known as the *Cherenkov angle relation*,

$$\cos\theta_C = \frac{1}{n\beta}, \tag{4.1}$$

where θ_C is the characteristic Cherenkov angle between the charged particle path and the emitted Cherenkov photons, as illustrated in Fig. 4-1, n is the index of refraction of the medium, and β is the speed of the particle relative to the speed of light, $\beta \equiv v/c$, where v is the particle speed. This is the basis for using Cherenkov radiation in particle identification.

The other relation is known as the *Frank-Tamm formula*, which describes the frequency spectrum of the Cherenkov radiation per unit path length of the charged particle in the original form of the formula [57]. It can be rewritten in the more convenient form which relates the number of Cherenkov photons to the path length of the charged particle, dx, and the wavelength, λ, of the those photons as

$$\frac{d^2N}{dxd\lambda} = \frac{q^2}{4\pi}\frac{c\mu(\lambda)}{h}\left(1 - \frac{1}{\beta^2 n(\lambda)}\right)\frac{1}{\lambda^2}, \tag{4.2}$$

where q is the electric charge of the particle, h is the Planck constant, c is the speed of light in vacuum, and $\mu(\lambda)$ and $n(\lambda)$ are the wavelength dependent permeability and index of refraction. The $1/\lambda^2$ dependence gives rise to the characteristic blue-purple glow of the Cherenkov radiation as seen in nuclear reactors. The proportionality of the number of Cherenkov photons to the path length of the particle is the basis for using Cherenkov radiation in calorimetry as in the GlueX Forward Calorimeter and in dosimetry in some medical applications [60].

Particle Identification

In the context of experimental particle and nuclear physics, Particle IDentification (PID), in generic terms, often refers to the process in which the experimenter makes an inference of the identity of the particle based on the experimental signatures in the detector instruments left behind by the particle. This generally means the determination of the charge and the mass of a particle and it is often achieved by combining information from several detectors in a modern experiment.

Depending on the specifics of a particular experiment, such as the energy scale of the particles of interest and the primary goal of the experiment, different types of PID detectors are used that exploit different particle interactions with matter. In the context of the GlueX experiment, the particles of interest are photons, electrons, pions, kaons, protons, and neutrons. For charged particle PID, the basic idea is the following. A charged particle follows a helical trajectory in a magnetic field. The tracking system can make a measurement of the three momentum and the electric charge of the particle based on the trajectory and the knowledge of the magnetic field. If, in addition, one could make a measurement of the particle velocity, one could use the relation between momentum, mass and velocity, $|\mathbf{p}| = \gamma m v$, to make a measurement of the mass of the particle. Therefore, the task of PID reduces to measuring particle velocity. Common principles that PID detectors leverage to achieve this goal include the energy loss, time-of-flight, transition radiation and Cherenkov radiation [59].

Cherenkov Radiation for Particle Identification

The principle of using Cherenkov radiation for PID is illustrated in Fig. 4-2. It leverages the Cherenkov angle relation of Eq. (4.1). Cherenkov photons are emitted when the charged particle passes through some dielectric medium at a speed greater than the speed of light in the medium. The medium is sometimes referred to as the radiator because it is responsible for the radiation of Cherenkov photons. If the medium is transparent to the Cherenkov photons, those Cherenkov photons can escape the medium and be detected downstream by some photodetectors such as photomultiplier tubes (PMTs) and silicon photomultiplier (SiPMs). One can make a measurement of the Cherenkov angle θ_C from the measured photon hit locations and the direction of the charged particle path, which is measured by the tracking system. From θ_C, one could make a measurement of the particle velocity given the particle momentum, achieving the goal of PID. A detector designed based on this principle is known as the ring-imaging Cherenkov detector (RICH).

4.1.2 DIRC Technology

The acronym DIRC stands for the **D**etection of **I**nternally **R**eflected **C**herenkov light. The technology was pioneered by the BaBar experiment at the SLAC National Accelerator Laboratory [61]. The DIRC is a special type of RICH detector which leverages the total internal reflection of the Cherenkov photons inside the radiator medium, as

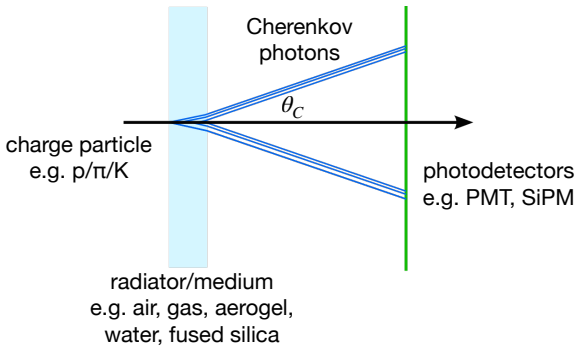

Figure 4-2: Schematic figure that illustrates the principle of using Cherenkov radiation for the purpose of particle identification.

illustrated in Fig. 4-3. Specifically, if the radiator medium has a rather high index of refraction compared to the environment, typically air, the produced Cherenkov photons can satisfy the total internal reflection condition, as illustrated in the bottom figure of Fig. 4-3. As a result, instead of escaping the radiator medium, they internally reflect inside the medium until the total internal reflection condition is not met. At that point, typically at the end of the radiator, the Cherenkov photons exit the radiator and enter some optical expansion volume, which directs the Cherenkov photons towards some photo-detection system, *e.g.,* a photodetector array plane, where they are measured. In this process, the medium serves not only as the radiator which is responsible for the production of Cherenkov photons but also as the light guide which is responsible for the transport of the produced Cherenkov photons.

A DIRC system is more complex in terms of the optical paths of the Cherenkov photons compared to a RICH system and a naive geometric measurement of the Cherenkov angle from the photodetector hits directly as illustrated in Fig. 4-2 is not possible. However, since the Cherenkov photons undergo either total internal reflection in the radiator bar or ordinary reflection off of mirrors in the optical box, the Cherenkov angle information is in principle preserved all the way from the production to the photodetection plane despite a much more complex pattern. If the detector geometry is known, one could still make a measurement of the Cherenkov angle, which depends on the charged particle speed, thus achieving the goals of particle identification. Reconstruction of the DIRC will be discussed in more detail in Sec. 4.5.

For the GlueX DIRC (to be discussed in detail in the following sections), the radiator is made of fused silica with an index of refraction of about 1.473. At the radiator-air interface, the threshold angle for total internal reflection can be calculated to be about 43°. As shown in Fig. 4-4, the Cherenkov angles θ_C of pions and kaons in the momentum range of interest of above 2 GeV/c are mostly greater than the threshold total internal reflection angle, meaning that most of the Cherenkov photons of interest undergo total internal reflection. The Cherenkov angle resolution is an important characteristic when discussing a RICH/DIRC system. It determines the momentum range over which the system can perform effective particle identification.

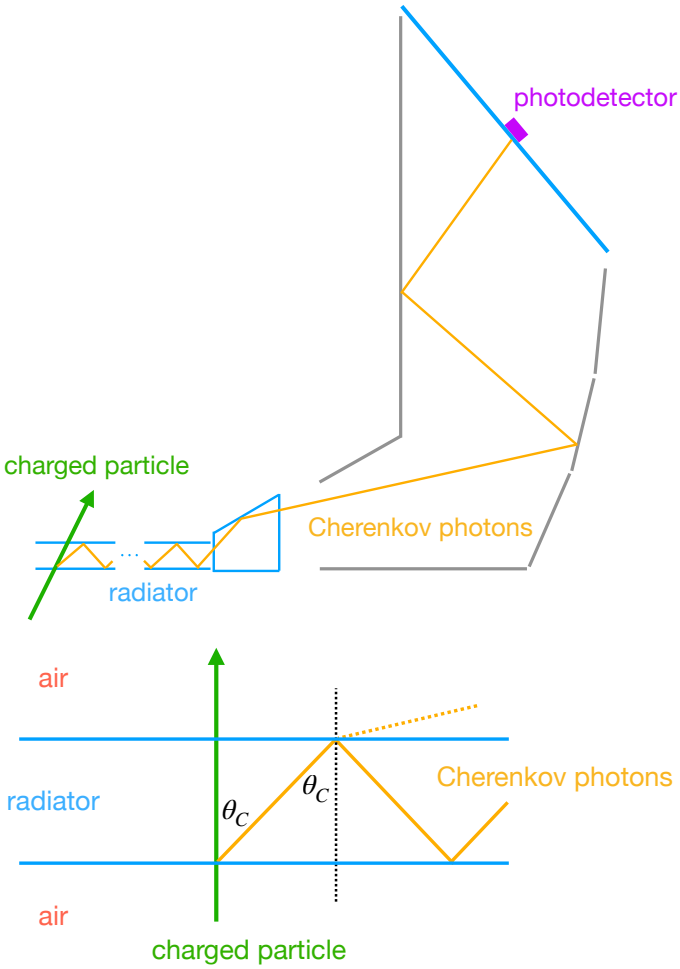

Figure 4-3: Schematic figure of a DIRC system (top) and a detailed illustration of the total internal reflection.

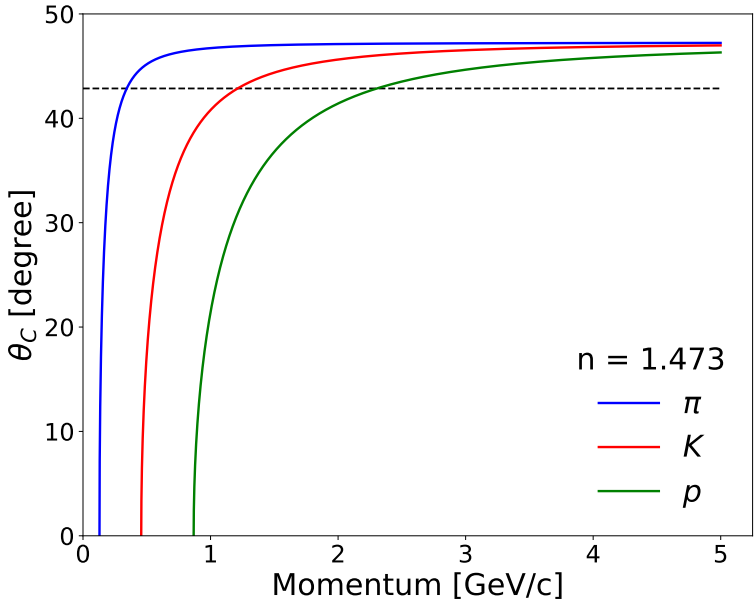

Figure 4-4: Cherenkov angle θ_C as a function of particle momentum of pions, kaons, and protons in fused silica ($n = 1.473$). The dashed line represents the threshold total internal reflection angle for fused silica-air interface of about 43°.

Fig. 4-5 shows the difference of Cherenkov angles in mrad for pions and kaons as a function of charged particle momentum in fused silica assuming $n = 1.473$ from 2 GeV/c to 5.5 GeV/c. Supposing that the Cherenkov angle resolution of a detector system achieves 2–2.5 mrad (which is the design goal of the GlueX DIRC [62]), it would achieve π/K separation at about 5σ level at 3 GeV/c and at about 3σ level at 4 GeV/c.

4.1.3 The GlueX DIRC Detector

A DIRC detector system consists of two major components: the radiator and the optical system. The GlueX DIRC uses the decommissioned radiators from the BaBar DIRC detector, whereas the optical system for the GlueX DIRC is designed specifically for the GlueX DIRC detector [62]. The fully installed GlueX DIRC detector is shown in Fig. 4-6 with annotations for some of the visible components. We provide an overview of the system here with many components described in detail in the following sections.

GlueX DIRC radiator: the bar boxes

The GlueX DIRC detector re-uses the decommissioned radiators from the BaBar DIRC. The design, construction, and performance are described in detail in [61] and summarized here. The radiators are made of fused silica with an index of refraction of

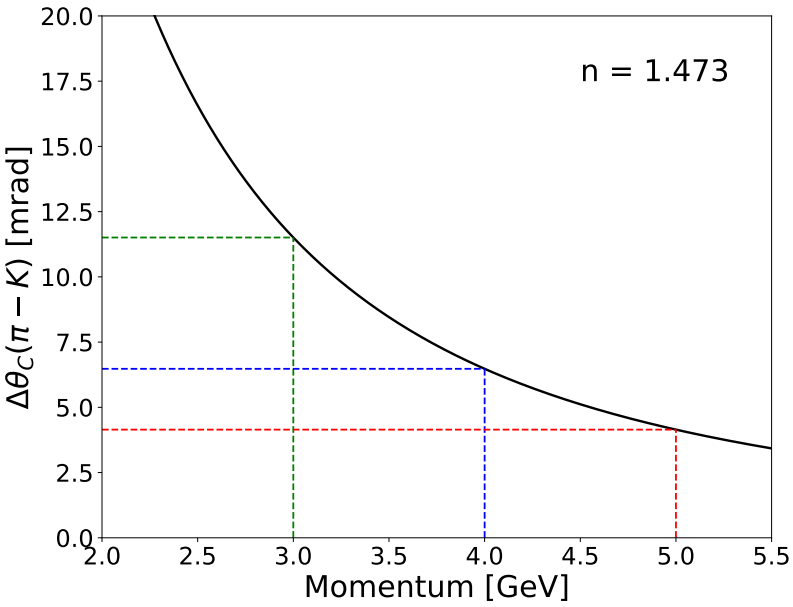

Figure 4-5: The difference of Cherenkov angles for pions and kaons as a function of charged particle momentum in fused silica (with $n = 1.473$) from 2 GeV/c to 5.5 GeV/c.

about 1.473. Each radiator bar is about 4.9 meters long (in the x-axis direction in the GlueX coordinate system), about 3.5 centimeters wide (in the y-axis direction), and about 1.725 centimeters thick (in the z-axis direction). At one end of a radiator bar, there is a reflective mirror with an air gap between the bar and the mirror to reflect the Cherenkov photons off in the other direction. At the other end of the radiator bar, it is glued to an extension wedge also made of fused silica. Twelve such radiators are grouped together, enclosed by a thin aluminum casing to form a so-called "bar box." The wedges are then glued to a thin fused silica window, which can then be interfaced with the optical expansion system. A schematic drawing of the bar box structure is shown in Fig. 4-7. At GlueX, four of the decommissioned BaBar bar boxes are used, with a total of 48 individual radiator bars.

GlueX DIRC optical system: the optical box

The design of the optical system of the GlueX DIRC detector is largely based on the fDIRC design for the SuperB experiment [63]. Details of the design can be found in [62, 64] and are summarized here. In essence, it is a system of reflective mirrors immersed in distilled water to match the index of refraction of the radiators to minimize photon loss at the interface. The Cherenkov photons are read out through a fused silica window by an array of Multi-Anode PMTs (MaPMTs). The water-filled "mirror box" is enclosed by a stainless steel outer structure, which also contains a dark box that encloses the PMTs and the readout electronics system, together forming the

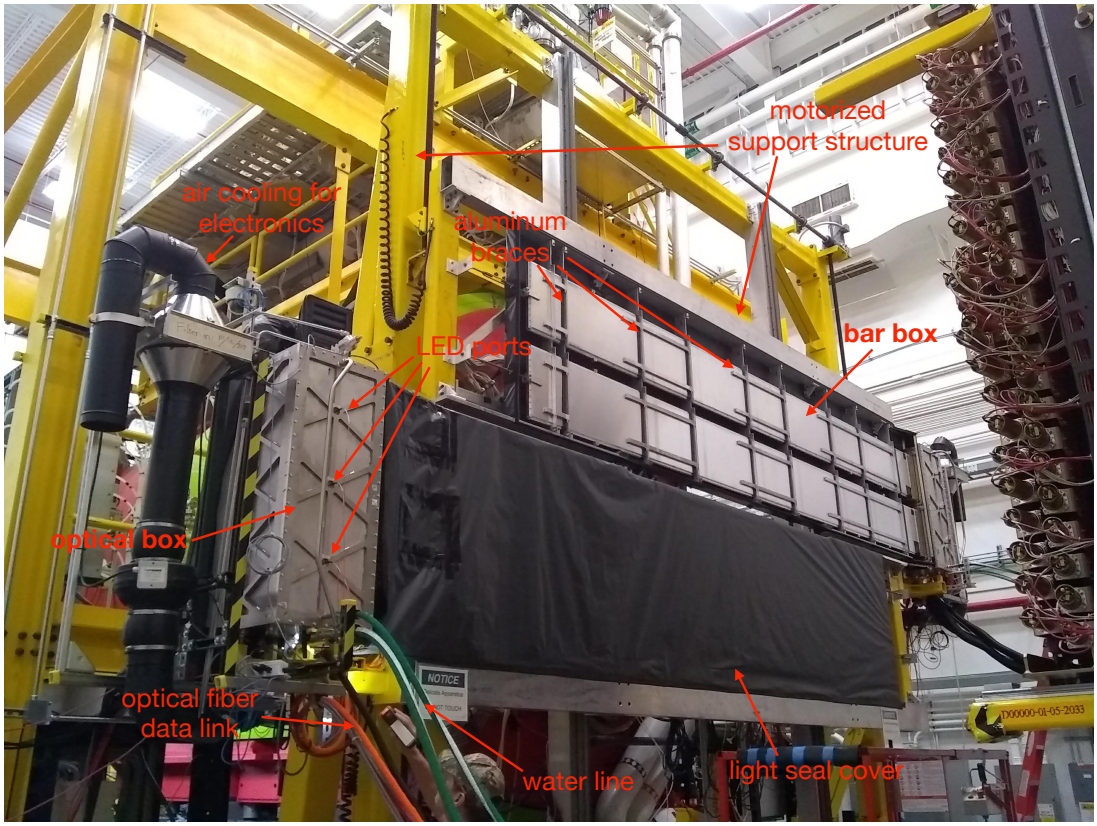

Figure 4-6: The fully installed GlueX DIRC detector in Hall D of Jefferson Lab.

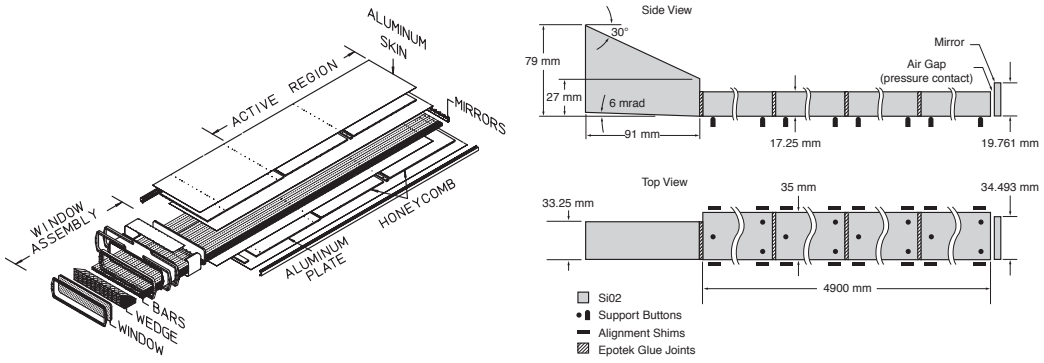

Figure 4-7: Schematic of the BaBar bar box. Images are taken from [61].

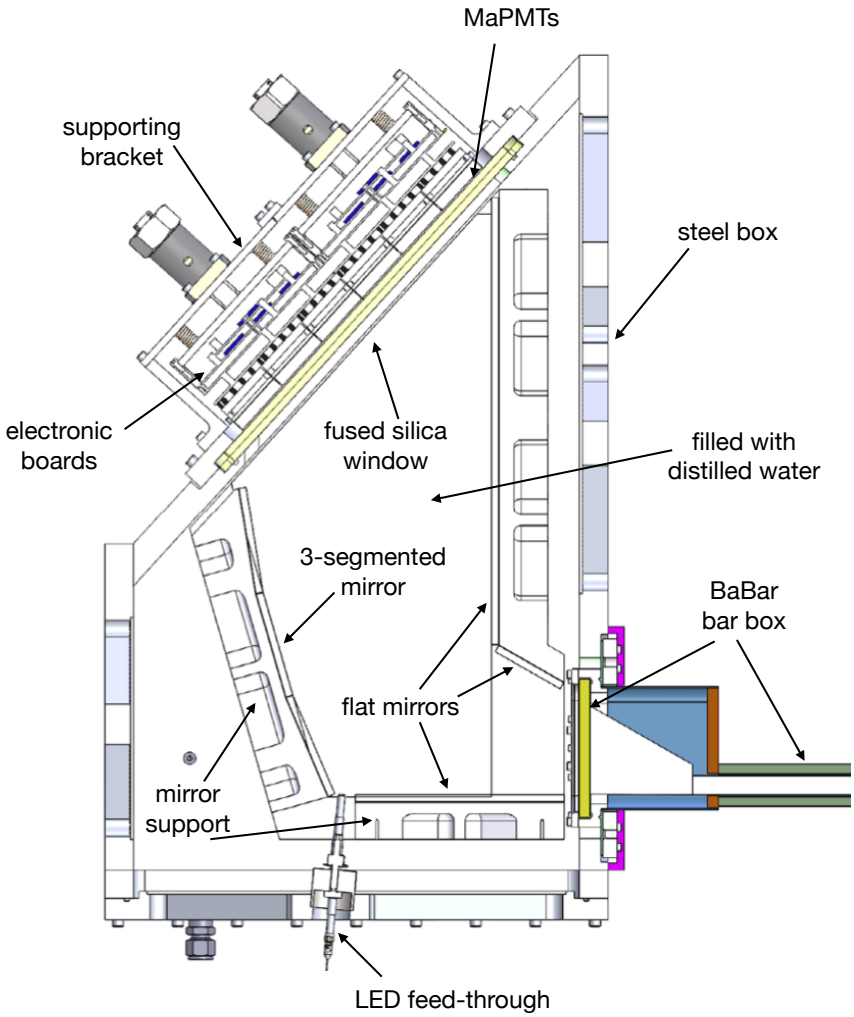

Figure 4-8: Engineering drawing of the GlueX DIRC optical box design.

so-called "optical box" (OB). A schematic drawing of the GlueX DIRC optical box design is shown in Fig. 4-8. The GlueX DIRC detector contains two such optical boxes, with each serving 2 bar boxes, or 24 individual radiator bars.

4.2 Transport of Bar Boxes

The GlueX DIRC uses four out of the twelve radiator bar boxes from the decommissioned BaBar DIRC, which had been stored at SLAC since the decommissioning of the BaBar experiment. They needed to be safely transported from SLAC to Jefferson Lab. We made two such cross-country trips that successfully transported one bar box in November 2017 and the other three in June 2018. This section describes the planning, execution and some lessons learned from those trips.

4.2.1 Considerations and Transport Plan

Mechanical Considerations

The radiator bars are long yet thin with about 4.9 meters along the long axis and about 1.7 centimeters in thickness. Each bar consists of 4 sub-bars joined together by Epotek optical glue. The bars were optically isolated and mechanically supported by so-called support buttons made of nylon [61]. In the transverse direction to the long axis, considerable mechanical stress can occur if the bars experience substantial torque due to their length and weight (the fused silica inside a box weighs about 200 lbs). Along the long axis, there are springs opposite the fused silica window behind the far end mirrors that provide a constant force to keep the glue joints in compression and to compensate for the hydrostatic load on the window from the BaBar water tank. Too much excess force in this direction risks window rupture. In addition, since decommissioning, the bar boxes had been stably stored at SLAC and had not been moved. The mechanical impact from aging of the fused silica itself, the optical glue joints, as well as the nylon support buttons was not well known. A finite element analysis was carried out in the planning stage of the transport and mechanical limits were established [62]. The ALARA (As Low As Reasonably Achievable) principle was adopted regarding minimizing the mechanical stress on the bar boxes.

In consultation and collaboration with Art Crating Los Angeles (ACLA), a delicate art service company, collaborators at Indiana University designed and built container crates for the bar boxes. A picture of the crate with various shock-absorbing features is shown in Fig. 4-9. It is essentially a "crate-in-crate" structure that can be disassembled for the ease of loading and unloading if necessary. A mock bar box was constructed and the shock-absorbing effects were tested. During the transports, shocks on the outer crate, inner crate, and bar box itself were monitored. The air spring pressure was also monitored.

Optical Considerations

The optical integrity of the radiator bar surface is critical for the DIRC detector to work because imperfections or damage of the fused silica surface could lead to diffuse reflection or photon loss. During storage, the bar boxes were placed in a climate controlled trailer and the bar boxes were under constant dry nitrogen purge to keep an overpressure with respect to the ambient atmospheric pressure to keep out moisture and dust deposition on the bar surface. There is a gas inlet on each bar box and the bar boxes are not air-tight, so the dry nitrogen can naturally leak out. The leak property of each bar box is also different. During the transport, a nitrogen cylinder was installed and provided constant flow of dry nitrogen gas into the bar boxes. The differential pressure between the bar boxes and the ambient atmospheric pressure as well as the flow rate from the cylinder into the bar boxes were also monitored during the trip. Fig. 4-10 shows the dry nitrogen purge system along with some other parts of the monitoring system.

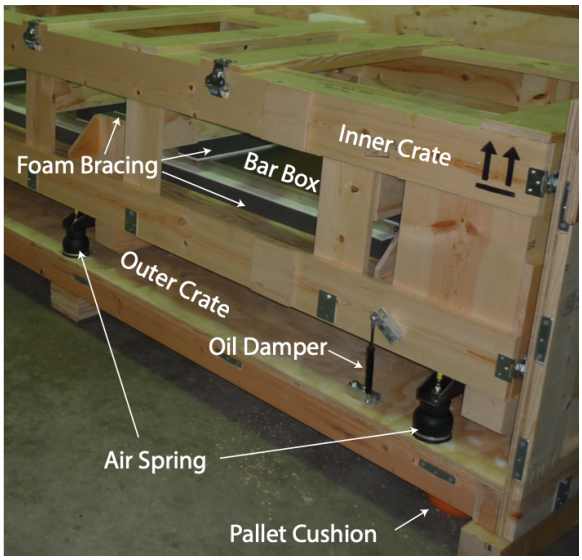

Figure 4-9: Picture of the bar box shipping crate showing key shock-absorbing components. Image taken from [62].

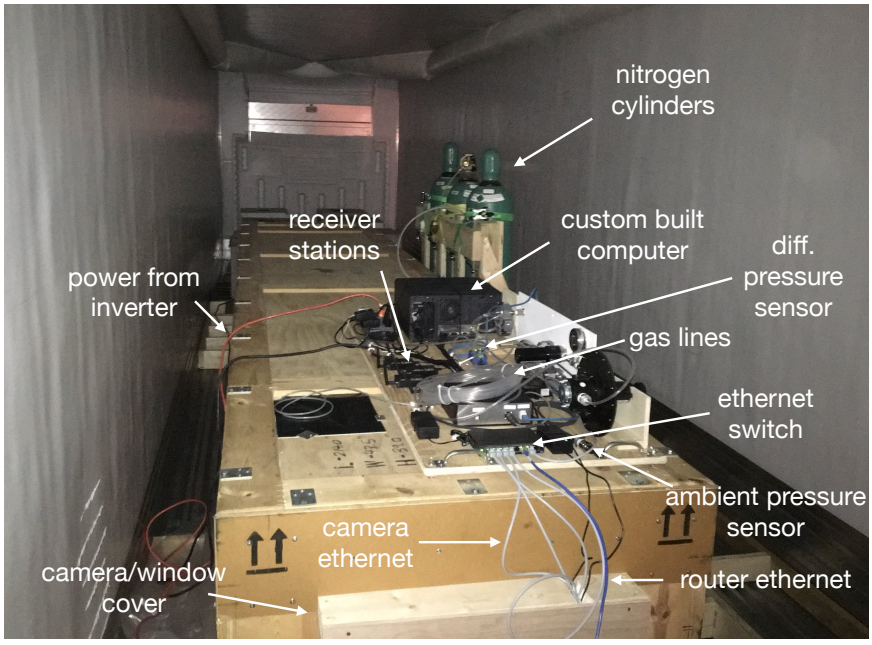

Figure 4-10: Picture of the truck trailer interior showing the crates, monitoring system, and nitrogen cylinder.

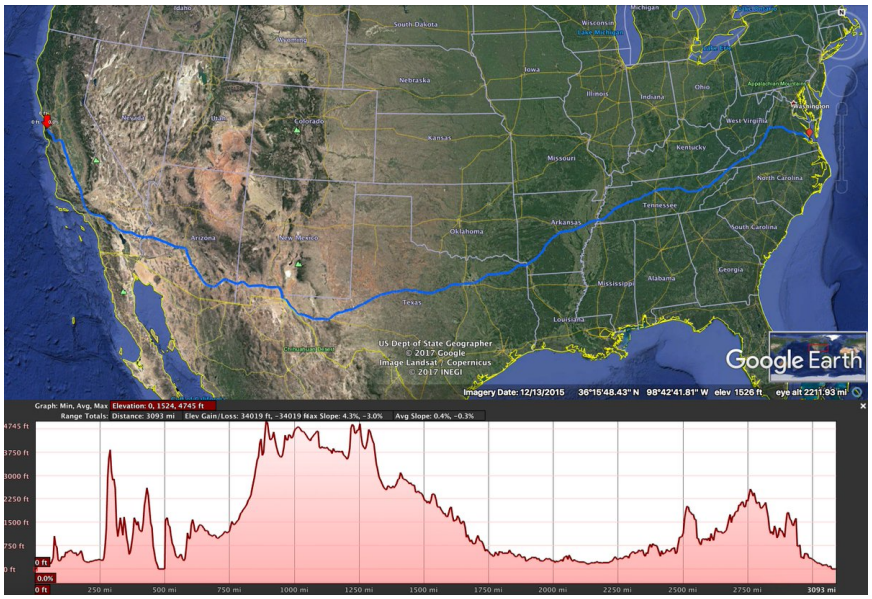

Figure 4-11: Route with elevation profile for the DIRC bar box transport trips from SLAC, Menlo Park, CA to Jefferson Lab, Newport News, VA.

Transport Plan

With the mechanical and optical constraints in mind, it was decided an overland cross-country transport was the best option. The crates containing the bar boxes together with the monitoring system and gas purge system would be placed inside of a climate controlled air-bed truck trailer. The transport crew would follow the truck in a *chase car* along the trips to perform near real-time monitoring. Transport plans via air or water were considered but ultimately rejected because the greatest risk would be during loading and unloading, which would require loading into a truck for either plan. Fig. 4-11 shows a map of the transport route along with the elevation profile. The route was chosen to minimize altitude variation in order to maintain dry nitrogen purge overpressure. The trips were chosen to happen in November 2017 and June 2018 to avoid seasonal extremes.

4.2.2 Monitoring System

Overview

The ultimate goal of the monitoring system was to support the mechanical crating system and the nitrogen purge system to ensure that the bar boxes arrived at Jefferson Lab intact, mechanically and optically. We wanted to ensure the bars were mechanically stable and were not broken during the transport. To this end, a camera system was set up to optically monitor the kaleidoscope pattern from the reflections of the bars. Accelerometers were placed on various locations of the crates to monitor the shocks and the air spring pressure was also monitored to ensure the shock-absorbing features of the crate function as designed. In addition, to best preserve the surface

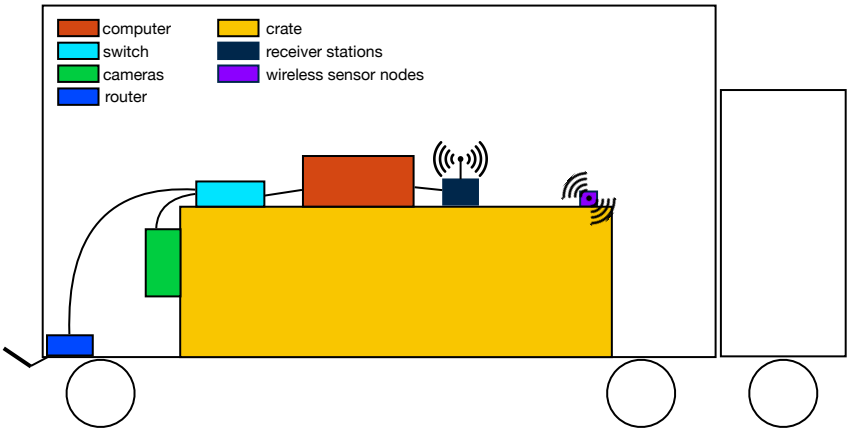

Figure 4-12: Schematic drawing of the monitoring system layout during the DIRC bar box transport.

optical property of the radiator bars, constant dry nitrogen gas purge was maintained and monitored. The cameras and various sensors communicated to a custom built desktop computer placed inside the truck trailer. The computer set up a local wireless network via a router.

Data Acquisition and Network

A custom desktop computer was built with sufficient storage space and processing capability. The monitoring system was powered by an inverter inside the truck. The camera and sensor systems communicated with the computer via either wire or wireless connections (via USB receiver stations). The computer collected and processed data from various sensors. A wireless router was connected to the computer to set up a local wireless network. The router antenna was extended outside via the trailer's drain holes for optimal signal range for the chase car. Fig. 4-12 shows a schematic drawing of the layout of the monitoring system, largely based off a similar drawing in [64] but with the addition of the wireless system including the accelerometry, nitrogen purge system as well as temperature and pressure.

Camera System

A good proxy of the mechanical integrity was determined to be the kaleidoscope pattern from the internal reflections of the bars. If an unfortunate catastrophic event such as a glue joint or bar breakage were to occur, the pattern would change in a manner significant enough to be visible to the human eye.

For each bar box, four cameras were required to provide full coverage. The cameras were powered and data communicated via an Ethernet switch which was connected with the computer. An LED strip was mounted just above the camera which would illuminate the bars and the cameras were focused at the window looking into the kaleidoscope pattern. Still pictures were taken at a rate of 1 frame per second. Every 15 minutes, the pictures were assembled into a time-lapsed video and archived on the

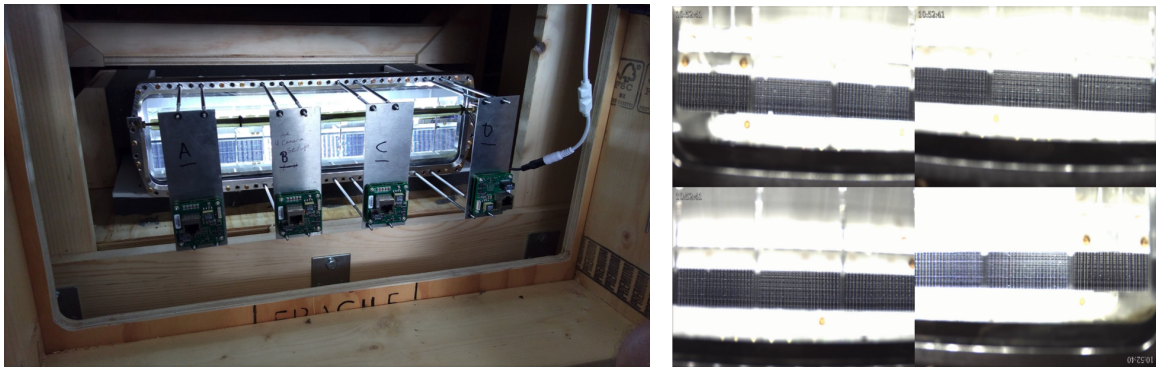

(a) (b)

Figure 4-13: Camera system during transport: (a) camera system set-up and (b) an example of the kaleidoscope pattern of the radiator bars as viewed by the cameras.

computer. The time-lapsed videos could then be downloaded via the local network by the monitoring crew in the chase car for viewing. Under normal conditions, those videos were "boring" because they looked like still pictures with barely noticeable vibrations. Occasionally, when large shocks were registered by the accelerometers, the video-viewing personnel would be notified to pay special attention to check if there was visible change in the kaleidoscope pattern around the time when the shock happened. Fig. 4-13 shows an example of such pattern. The archived videos from the trips and more details about the camera system can be found in [64].

Shock Absorption and Accelerometry

The accelerometers used in the transport were the G-LINK-200 versions purchased from the LORD MicroStrain® Sensing Systems [65], which provided acceleration measurements of up to $8g$ along the three axes. They were battery-powered and communicated wirelessly to the so-called "gateway" receivers which were connected via USB to the computer for data acquisition and processing. A schematic drawing of the system is shown in Fig. 4-14. There were essentially three layers in the "crate-in-crate" design: the outer crate, the inner crate, and the bar box itself. Accelerometers were attached via screws to the outer and inner wooden crates and via double-sided tape to the bar boxes.

The sampling rate of the accelerometers can be configured by the user from 1 sample per hour up to 4096 Hz. Depending on the placement of the accelerometers and the gateway base station bandwidth, 128 and 256 Hz sampling rates were used. The vendor provided several ways for data collection via a PC application, a web application, or an API library [65]. We chose to use the API library because it provided the maximum flexibility. Accelerometry data were streamed constantly to the computer. Every 10–15 minutes, a separate process was activated to retrieve and plot the data for the past 10–15 minutes. Both the raw data in text format and the monitoring plots were archived on the computer. The monitoring personnel in the chase car could then download the monitoring plots and inspect any abnormality. Fig. 4-15 shows

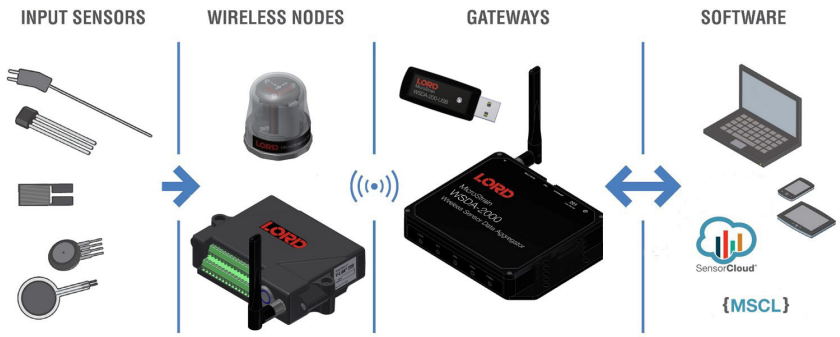

Figure 4-14: Schematic drawing of the LORD MicroStrain® Sensing Systems.

example monitoring plots of the accelerometry data that the monitoring personnel were looking at during the June 2018 transport. Such plots are typically uneventful like the one shown here.

Although not expected to fail, the air spring pressures of the crates were also monitored at a rate of 1 Hz. The pressure measurement itself was provided by a pressure sensor from Omega™ [66]. The analog signal from the pressure sensor was passed to the V-LINK-200 LORD MicroStrain® Sensing Systems ADC [67] which would convert the analog signal to digital signal and then wirelessly communicate with one of the gateway stations. The digital data were then archived and processed on the computer and monitoring plots were made in similar 10–15-minute intervals similar to the accelerometer data stream. During the second trip, we experienced data loss, likely due to RF interference. It was not clear if it was an inherent limitation of the system or a result of the challenging environment in which we operated it (the refrigerated truck trailer is a giant Faraday cage as far as the RF system is concerned). Nonetheless, a wired system would likely be a more robust choice if there are any similar shipments in the future.

Nitrogen Purge System

It was important to keep an overpressure of the bar box with respect to the ambient atmospheric pressure to keep out moisture and dust. This was achieved by a nitrogen purge system. The differential pressure between inside the bar box and the ambient atmospheric pressure as well as the flow rate from the nitrogen cylinder into the bar boxes were both monitored. Both analog signals of the differential pressure and the flow rate were passed to the V-LINK-200 ADC and sampled at a rate of 1 Hz. The rest of the data stream and monitoring routine were similar to that of the air spring pressure.

In addition to the usual monitoring routine described above, the nitrogen purge data stream (differential pressure and flow rate) was duplicated and written to a separate text file at the sampling rate of 1 Hz. This allowed the monitoring crew in the chase car to achieve almost real time monitoring of those parameters via `tail -f` so long as the network connection permitted. This functionality proved to be useful especially during steep descents where the ambient atmospheric pressure increased

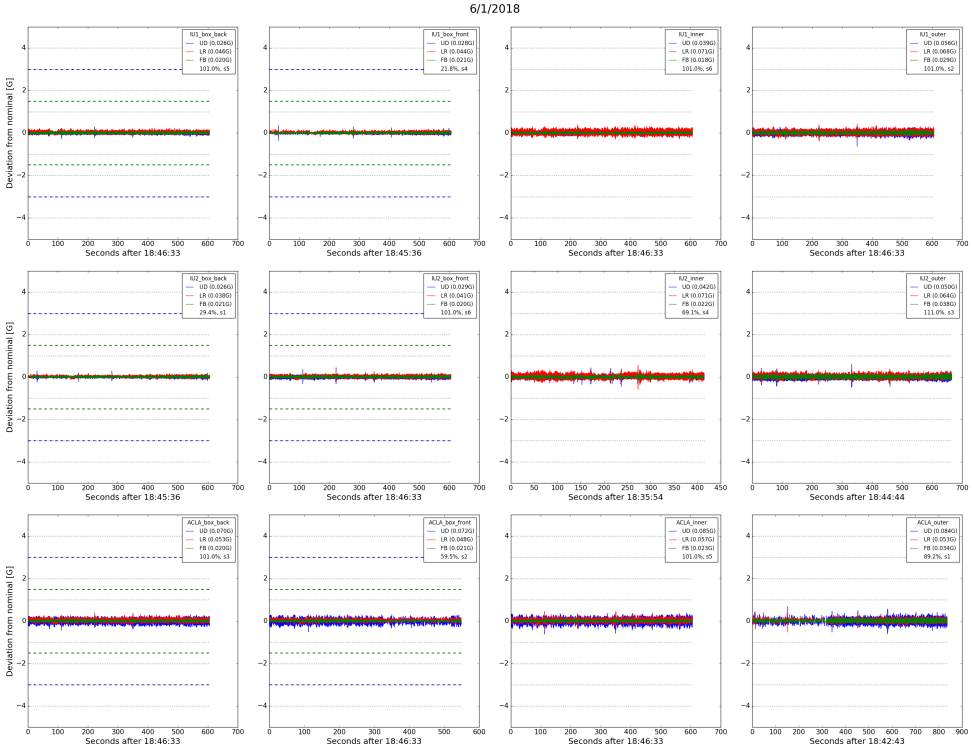

Figure 4-15: Example monitoring plots of the accelerometer data during the transport that the monitoring crew was inspecting.

rapidly and the nitrogen flow from the cylinder struggled to keep up in order to maintain constant overpressure. During descents, if the monitoring crew observed rapid differential pressure drop, the chase car could communicate to the truck driver via two-way radio to request slow-down whenever possible. During another instance, the monitoring crew observed the flow rate had dropped to zero and immediately communicated to truck driver to stop at the next convenient location. The crew went into the truck trailer and had discovered that one gas-line connection fell loose.

Fig. 4-16 shows an example of the monitoring plots including those related to the nitrogen gas purge system that the monitoring crew was inspecting during the June 2018 transport. This particular profile was captured during the descent into Las Cruces, New Mexico on U.S. Interstate-10 from the west. As the descent happened, the atmospheric pressure rose (measured in inches of water column), which resulted in a decrease in the differential pressure between inside the bar box and the ambient, which in turn triggered the increase in the flow rate to maintain a positive differential pressure.

Environmental

Ambient atmospheric pressure and temperature inside the truck trailer were also monitored. The ambient atmospheric pressure was monitored by an Omega™ pressure sensor [66] and the temperature monitoring was provided by a thermocouple by Omega™ [68] connected with a LORD MicroStrain® Sensing Systems wireless node [69]. The thermocouple was placed near the cameras to make sure any heat produced by the camera system or the LED light did not heat up locally near the bar box windows. The analog signals from the pressure sensor and the thermocouple were sampled at 1 Hz and passed to the V-LINK-200 ADC. The rest of the data stream was similar to that of the air spring pressure and the nitrogen purge system.

4.2.3 Summary

Immediately after the bar boxes arrived at Hall D of Jefferson Lab, optical inspections of the kaleidoscope patterns were conducted and no noticeable differences were observed. Subsequently, the analysis of the DIRC commissioning data showed comparable performance in photon yield and single photon resolution (see below) as expected from BaBar data after accounting for differences in the optical system design, indicating the success of the transports. We do not expect that damages incurred during the transport, if any, would gradually appear over time although long-term monitoring of the DIRC performance is part of the normal operation protocol. An internal summary including the lessons learned of the transports was documented in [70].

4.3 Installation and Commissioning

With the bar boxes successfully transported to Hall-D of Jefferson Lab, the next phase involved the installation and integration of the full DIRC system into GlueX

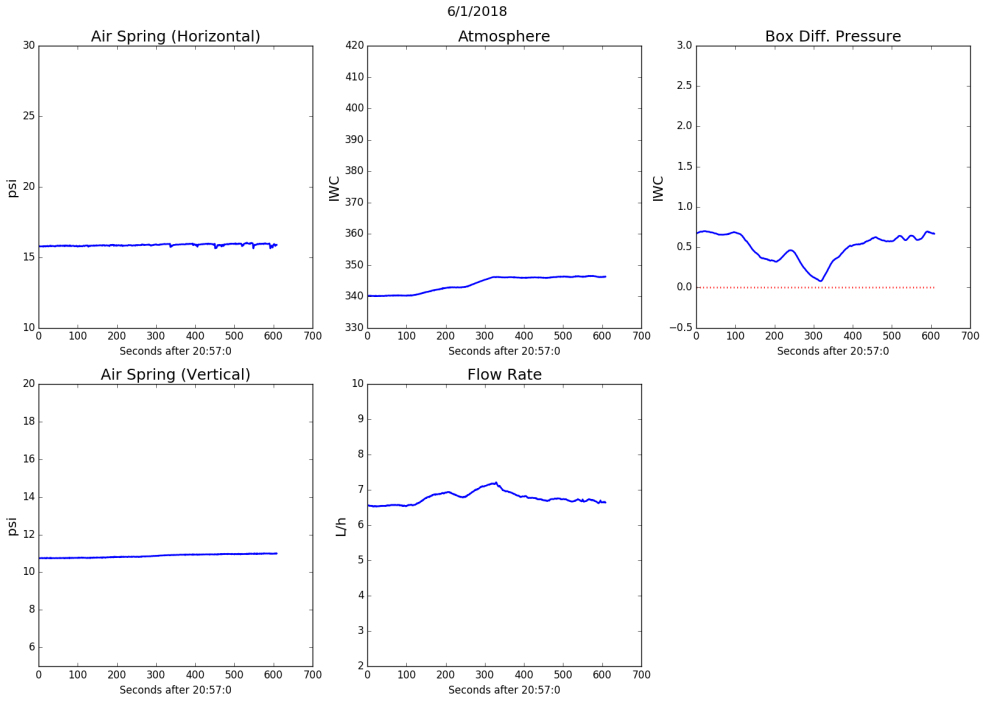

Figure 4-16: Example monitoring plot of the V-LINK-200 ADC node data: air spring pressures, nitrogen flow rate, differential pressure between bar box and ambient, and ambient pressure. See text for the explanation of the variations.

and its commissioning. The mechanical support structure of the DIRC bar boxes was designed, manufactured, and delivered by the Indiana University group with details found in [62]. The design and manufacturing of the optical boxes were led by the MIT group with details found in [62, 64]. The assembly of the optical boxes including the reflective mirrors, the water and gas system, and the final assembly of the DIRC detector along with many other tasks were all handled by the Hall-D engineering staff. The downstream integration of the DIRC readout into the rest of the GlueX trigger and DAQ scheme was handled by the Hall-D electronics group. During the DIRC installation, the physics group was primarily focused on the photodetection electronics.

4.3.1 Photodetection Assembly

The detection of Cherenkov photons is achieved by an array of Hamamatsu H12700 multi-anode photomultiplier tubes (MaPMTs) [71] instrumented behind the optical box fused silica window. Each H12700 MaPMT is an 8×8 array of 64 pixels. The total effective area is 48.5 mm $\times$ 48.5 mm with each pixel of 6 mm $\times$ 6 mm in size. The PMT plane of each optical box consists of an array of 18×6 PMT locations. A total of 18 such locations are not currently instrumented with PMTs due to a combination of their low occupancy and budget/replacement considerations. Such locations were chosen such that the impact on reconstruction was minimal due to their low occupancy. They were instead instrumented with 3D-printed plastic *dummies* which have the exact same dimensions as the actual H12700 MaPMTs and were treated identically during installation. Each MaPMT has $8 \times 8 = 64$ channels. Therefore, each optical box has a total of $[(18 \times 6) - 18] \times 64 = 5760$ channels or a total of 11520 readout channels for the GlueX DIRC system. The MaPMTs are read out by a series of electronics boards which digitize the analog PMT signals and pass to the downstream trigger and DAQ stream. A schematic drawing of the readout scheme is shown in Fig. 4-17a [62]. Each *PMT module* consists of three MaPMTs, an adapter board, an ASIC board and an FPGA board as shown schematically in Fig. 4-17b [72]. The adapter board supplies high voltage to the PMTs and routes the PMT outputs to the ASIC board. Each ASIC board hosts three multi-anode readout chips (MAROC), each serving one MaPMT, which amplify and digitize the PMT current output into hit time and time-over-threshold information. The processed pixel-by-pixel information is then passed to the FPGA board which communicates with the trigger and DAQ system downstream in the data flow. Two sets of the PMT modules (each housing three PMTs) are assembled in a stack under the accompanying mechanical support bracket as shown in red in Fig. 4-17c [73]. Fig. 4-17d shows a picture of the fully assembled PMT module brackets being attached onto the fused silica window of the optical box in a clean tent set up in Hall-D during installation [74].

4.3.2 Optical Cookies

In the GlueX DIRC, there are on average about 30 to 40 Cherenkov photons per charged particle that are detected by the PMTs, taking into account detection effi-

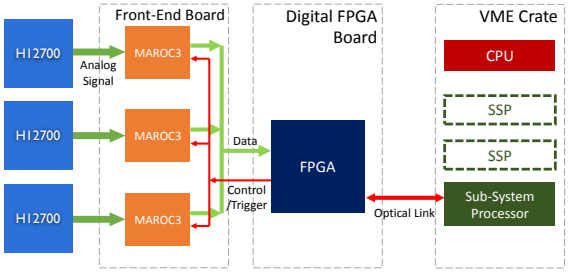

(a) Readout flow chart.

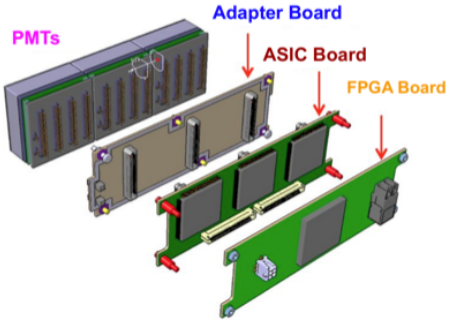

(b) Schematic drawing of the PMT module electronics boards stack.

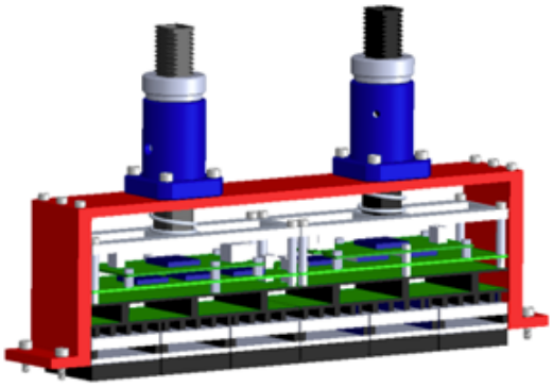

(c) Schematic drawing of the full PMT module including the mechanical support structures (in red).

(d) PMT modules during installation.

Figure 4-17: DIRC readout scheme and PMT module.

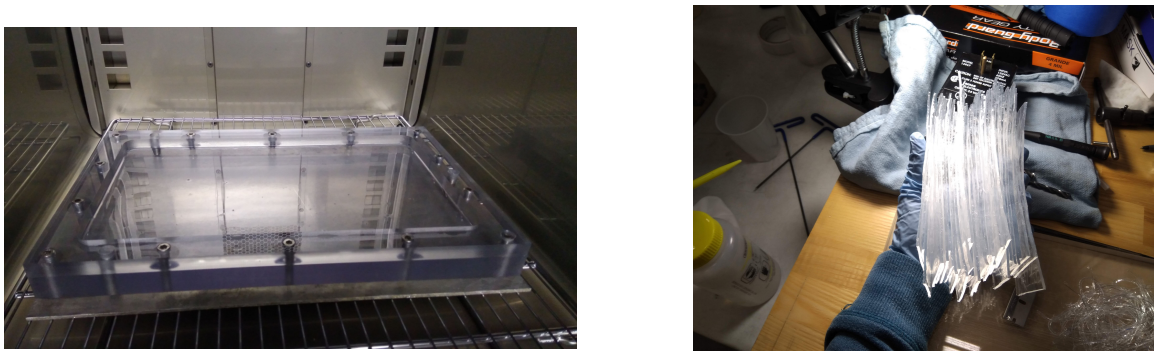

(a) (b)

Figure 4-18: Pictures of DIRC cookies during production and installation: (a) RTV-hardener mixture in a custom made mold before curing in oven; (b) cookies ready to be installed.

ciency and all the attenuation in the radiator bars, in the water in the optical box, and in various interfaces. Therefore, any increase in the number of detected photons, or photon yield, is favorable to the reconstruction. Along the path of a Cherenkov photon from the production in the fused silica radiator to the PMT photocathode, it is desirable to match the indices of refraction when it crosses different interfaces to minimize photon loss. One potential place for large refraction index mismatch is from the fused silica window ($< n > \approx 1.473$) to the PMT window, which is made of borosilicate glass [71] ($< n > \approx 1.5$ in the optical range) because it would have to go through air, which has n of about 1. Simulation suggests that a 0.1 cm air gap relative to the best case scenario, where the photons go directly from the fused silica window into the PMT window, results in a photon loss of about 15–20% [73].

A solution, inspired by Belle II TOP detector experience [75], was proposed to add a layer of the so-called "optical cookies" which provide a better refraction index matching to replace the air between the OB window and the PMT window. For the GlueX DIRC, the cookie is a thin pad made of Room-Temperature Vulcanizing (RTV) silicone. The RTV silicone used for DIRC cookies is in liquid form at room temperature and the production of the cookies involves mixing RTV with some hardening agent silicone, degassing, and curing at a higher temperature (*e.g.,* 100°C) in a custom made mold before extraction and cutting to proper sizes. Each cookie serves three PMTs. Tests of different RTV-hardener ratios and curing temperatures were done to optimize the optical performance in terms of transmission and the production procedure. To expedite the application of cookies to the PMT modules and eliminate air bubbles from forming between the OB window and the cookies or between the cookies and the PMT window, optical greasing oil was applied during the installation. A detailed descripiton of the installation procedure, as well as more details about the cookie technology, production procedure and tests can be found in [73, 64]. With the addition of the cookies (and the optical greasing oil layer), the photon loss improves to about 2–5% [73].

4.3.3 Commissioning

There had been two commissioning periods for the GlueX DIRC. The first commissioning period was a ten-day run in February 2019 under nominal GlueX Phase I (low intensity) running conditions. Half of the DIRC was commissioned: two bar boxes (the two below the beamline, #10 and #11) and one optical box (the South OB). The second commissioning was a two-week run in December 2019. During the second commissioning period, the full GlueX DIRC system was commissioned and it was also tested under GlueX Phase II (high intensity) running conditions.

4.4 DIRC Simulation

Accurate Monte Carlo (MC) simulation of the DIRC system is critical to the success of the detector because the reconstruction algorithms often depend on aspects of the simulation, especially the geometry. More details are discussed in Sec. 4.5. The Monte Carlo simulation of the GlueX DIRC detector is realized within the standard Geant4-based GlueX HDGeant4 framework discussed in Sec. 2.4.3. In addition to the specification of the DIRC geometry and materials, two primary generators have also been added: one for the generation of a look-up table (LUT) for the geometrical reconstruction method (to be discussed in more detail in Sec. 4.5) and the other for the LED calibration system.

4.4.1 As-Built Geometry

There are two major components to a DIRC system – the radiators and the optical system. In our case, they are called the bar boxes and the optical boxes. During the design stage, all the components have their nominal dimensions and are aligned in the nominal way, that is perfectly perpendicular or vertical with no rotations or offsets. In practice, this *nominal geometry* is never achieved and the actual so-called *as-built* geometry may not be knowable. There are trade-offs in the implementation of the as-built geometry one needs to consider. On the one hand, one would obviously like to implement the as-built geometry as close to reality as one could. On the other hand, one would like to make sure that (1) the as-built geometry is known to a good enough precision so that it actually represents the reality and (2) it does not slow down the simulation by too much. Both concerns are relevant for the DIRC. For (1), studies have shown that we need to know the unknown misalignment of the mirrors to a few mrad in order to ensure reconstruction performance [62, 64]. This level of precision would be very challenging to survey and implement in simulation. This type of misalignment would need to be calibrated out with real data either by estimating the misalignment and implementing it in the simulation or by applying some correction factors determined from data. The optimal calibration strategy depends on the reconstruction algorithm. For (2), we do not want to over-complicate the geometry specification because optical photon simulation is already slow in Geant4.

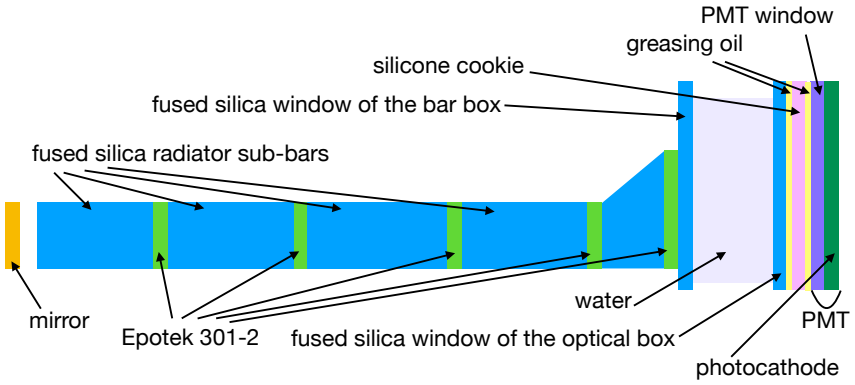

Figure 4-19: Schematic of the materials used in the GlueX DIRC simulation. Components are not drawn to scale. Image source: [76].

Material Specification

The definition of the Geant4 geometry includes not only the physical dimensions of the objects but also the material specification. At GlueX, they are specified within the Hall-D Detector Specification HDDS framework. A summary schematic of all the volumes of different material relevant for a Cherenkov photon is shown in Fig. 4-19.

In addition to physical properties such as elemental composition and density, optical properties are also very important for the modeling of the DIRC, which include the energy/wavelength-dependent index of refraction and reflectivity in the case of the mirrors. These are important in the modeling of the detected photons, including the photon yield and photon timing. Fig. 4-20a shows the dependence of the absorption length of the different materials used in the simulation on photon wavelength. Fig. 4-20b shows the photon wavelength spectra from different contributions in the simulation. Black is the spectrum in use which includes everything. The red histogram shows the underlying Cherenkov spectrum where the red line is a $1/\lambda^2$ fit. Blue adds the PMT quantum efficiency (QE). Cyan adds the materials whose optical properties are more well known, whereas magenta adds in other materials including the Epotek optical glue in the bar box [61, 63], borosilicate PMT window, the cookies and the optical greasing oil for the cookies. There is a strong cutoff near around 300 nm due to the Epotek glue and borosilicate PMT window.

Bar Box As-Built Geometry

The bar boxes are enclosed and protected by a thin aluminum case and internal honeycomb support structure as shown in Fig. 4-7. In addition, the bar boxes are mechanically supported by thin aluminum braces on the support structure in Hall D as shown in Fig. 4-6. They add to the material budget on top of the fused silica radiator bars themselves. Therefore, such support structures needed to be modeled, as shown in Fig. 4-21a, for the accurate modeling of the detectors downstream of the DIRC, mainly the Forward Calorimeter.

The GlueX DIRC uses the decommissioned BaBar DIRC bar boxes as the ra-

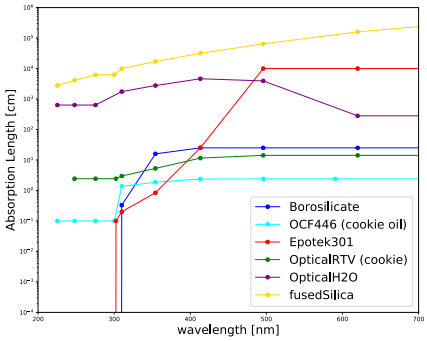

(a) Absorption length as a function of photon wavelength in the DIRC simulation.

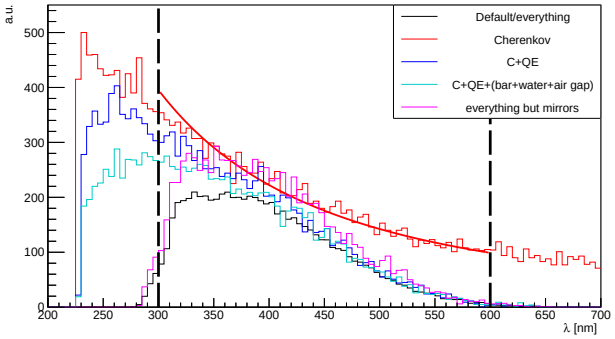

(b) Photon wavelength spectra showing the impact from different contributions in the simulation. More details about the different curves are described in the text.

Figure 4-20: Material absorption length and the photon wavelength spectra.

diators. The dimensions of the individual radiator sub-bars (each about 1.25 meters long) inside a bar box were documented by the BaBar DIRC group and this information was passed on to GlueX [77]. During construction, the bars were also aligned following a defined convention. The individual bar dimensions and their alignment were implemented in the standard GlueX Geant4 simulation. Although we do not expect them to have a substantial impact on our reconstruction, the information was available, precise, and easy to implement. The implementation also did not slow down the simulation in any appreciable manner. Fig. 4-21b shows a zoomed-in view of the interface between two individual sub-bars showing different dimensions and their alignment.

Optical Box As-Built Geometry

After the DIRC detector was installed, both the bar boxes and the optical boxes were surveyed by the Jefferson Lab survey group. Measurements of the translation and rotations were done for both. In principle, we could implement the survey results to the best of our ability in the DIRC simulation and such attempts were made [78]. Fig. 4-22 shows a zoomed-in view near the 3-segmented mirror. Rotational offsets can be seen as a result of the implementation of survey results. Afterwards, a new look-up table was generated with the surveyed results. However, the new LUT resulted in worse performance than the nominal configuration in preliminary studies. It was not unexpected for two reasons. First, it was a nontrivial task to implement such rotations and offsets in the first place as can be seen in [78], so it would not be surprising that errors were made. More importantly, the required precision for the knowledge of the mirror position and offsets is of order a few mrad, which is at least the same order of precision the survey could achieve if not even smaller. The implementation of such survey results was, therefore, not too different from random Gaussian noise. Therefore, the alignment calibration would need to be done with

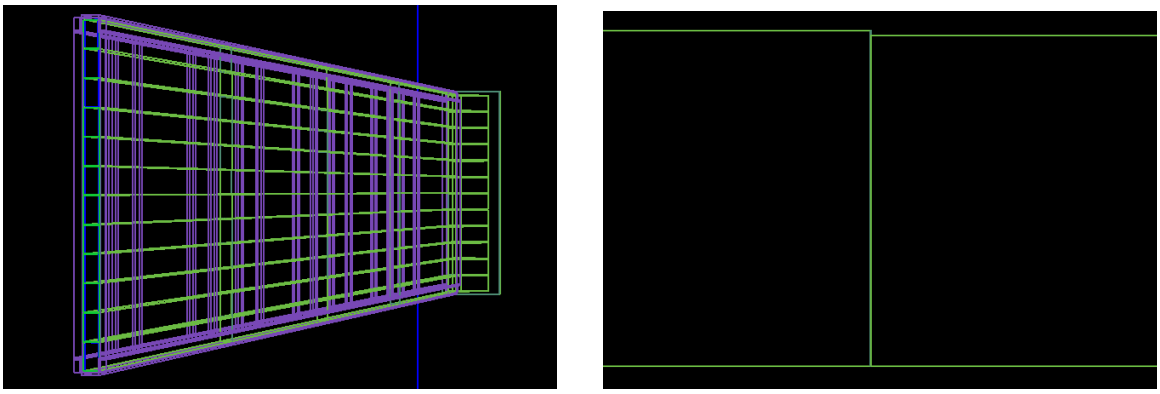

(a) (b)

Figure 4-21: Implementation of the bar box as-built geometry in the DIRC simulation: (a) aluminum casing of the bar boxes and the aluminum mechanical support braces in simulation; (b) individual bar dimensions implemented in simulation.

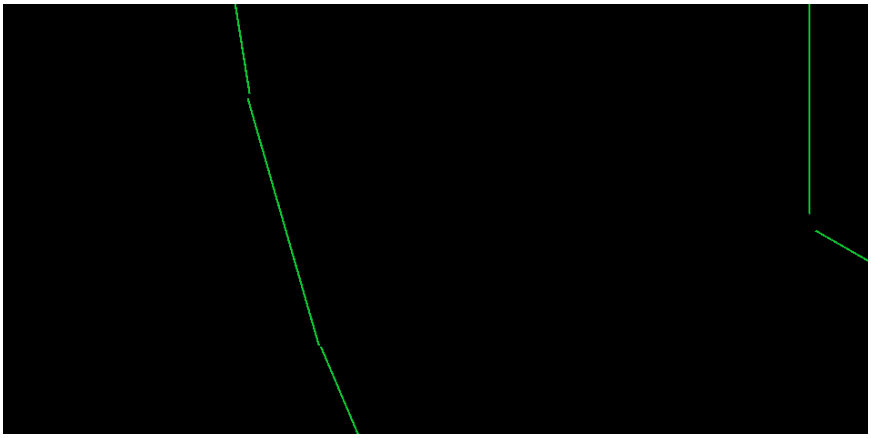

Figure 4-22: Example attempted implementation of the mirror survey results. Rotational offsets from the nominal can be seen in this picture.

actual data, potentially in slightly different ways depending on the reconstruction algorithm. The issue of calibration and alignment will be further discussed in Sec. 4.5.

4.4.2 Additional Primary Generators

LED Generator

In each optical box, there are three LED feed-through ports that can direct LED light at the PMT plane, as shown in Fig. 4-8, to provide monitoring and calibration. The LED system primarily consists of an LED source with about 405 nm wavelength, a reference Silicon Photomultiplier (SiPM), and the LED ports into the OBs. It contributes to another trigger running in parallel with the main physics DIRC trigger under normal operating conditions at 150 Hz (and adjustable) trigger rate. The three LED ports in an OB are separated 10 ns apart in time. Using the timing from the

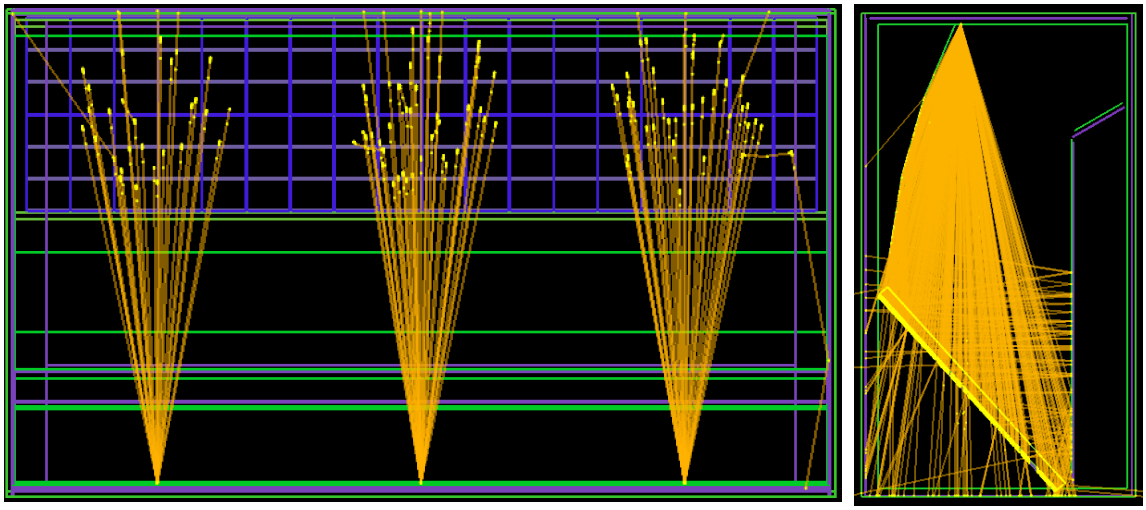

Figure 4-23: Visualization of the LED system in simulation.

reference SiPM and the detected per-pixel time, the LED system provides a per-pixel timing calibration of the PMTs with respect to each other, while the global timing calibration of the DIRC system is obtained from an overall shift of all the pixels with respect to some physics trigger provided in conjunction with other parts of the GlueX system. In order to achieve uniform coverage of the full PMT plane and minimize overlap (that is to avoid one pixel being hit by two LED ports because different ports have different timing offsets), diffusers of different shapes and diffusing angles were explored and studied in simulation. In addition, the LED system can provide insight into the long term performance of the detector, such as water quality and detector efficiency. The LED system is described in more detail in [72, 76].

Simulation of the LED system was implemented to provide validation of the timing calibration procedure as well as different diffuser shapes and inclination angles. Fig. 4-23 shows example visualizations of the DIRC LED system.

LUT Generator

For the geometrical reconstruction (described in Sec. 4.5), a look-up table (LUT) of all possible photon paths in the optical box from the end of each radiator bar is used in the reconstruction procedure. This look-up table is pre-generated with a large amount of photons emerging from the end surface of a radiator bar into the wedge and then the optical box before they are lost or hit a certain PMT pixel. A generator for producing the LUTs was written. Modifications of essentially careful bookkeeping were also implemented to ensure that after rotations and offsets of the bar box or the optical box, photons were still always emerging from the end surface of the radiator bars. Fig. 4-24 shows example visualizations of the LUT generator, verifying that the generator behaved as expected when rotations and offsets were present.

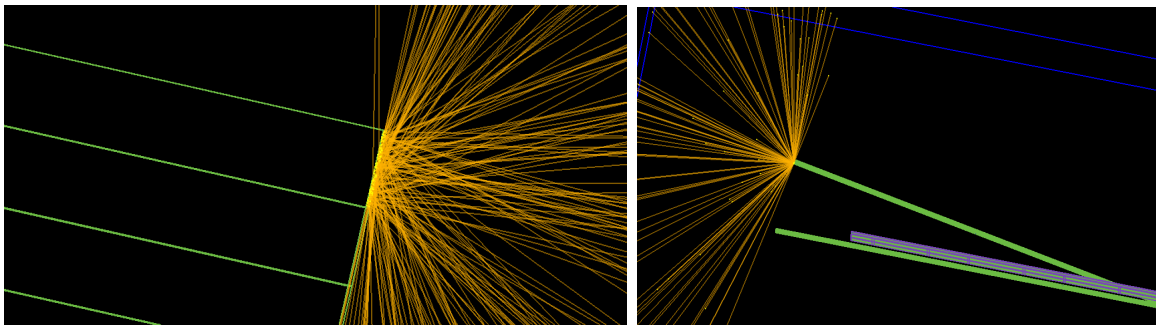

Figure 4-24: Visualizations of the LUT generator verifying that the generator behaved as expected even when rotations and offsets are present.

4.5 DIRC Reconstruction

4.5.1 DIRC Data and Overview of Reconstruction Strategies

For a track/charged particle[1] that reaches the DIRC plane, it hits a particular radiator bar inside one of the bar boxes. Cherenkov photons are produced and travel along the length of that bar and into one of the optical boxes. They then bounce off the mirrors inside that optical box and are eventually detected by the PMTs at the photodetection plane. For a typical event, about 30–40 photons are eventually detected, each by one of the 5760 PMT pixels, over a time span of $\mathcal{O}(100\,\text{ns})$. There is an $\mathcal{O}(100\,\text{ns})$ time span because each photon is produced at a different location along the Cherenkov cone of the charged particle while it traverses through the thin but finite thickness of the radiator bar. Depending on the angle at emission, some photons travel directly towards the optical box, which we call the *direct* photons, while some photons first travel away from the optical box and bounce back off the far-end mirrors before traveling towards the optical box, which we call the *reflected* photons. Therefore, for each *DIRC event*, that is when a charged particle hits a radiator bar producing *hits* in the PMT pixels, the PMT hit data can be a collection of points in 3D space (x, y, t), where (x, y) specifies the location of the hit pixel on the PMT plane and t is the hit time. Fig. 4-25 shows an example of such a 3-D hit pattern. The goal of DIRC reconstruction can therefore be thought of as a hypothesis testing problem of the particle mass, given the collection of 3D hits in (x, y, t) and the particle momentum measured by the tracking system, with the knowledge of how Cherenkov photons are produced, *i.e.,* its characteristic Cherenkov cone, and how they propagate in the detector geometry of the radiator bars and the optical box. The Cherenkov angle information, which depends on the charged particle speed, which in turn depends on the particle mass given its momentum, is in principle preserved all the way through the DIRC system from the production to the detection of the Cherenkov photons because the photons undergo either total internal reflection in the radiator bar or ordinary reflection off the mirrors in the optical box, which can both be modeled straightforwardly provided the detector geometry is known. A number of DIRC

[1]Note that *track* and *charged particle* are used interchangeably in this document.

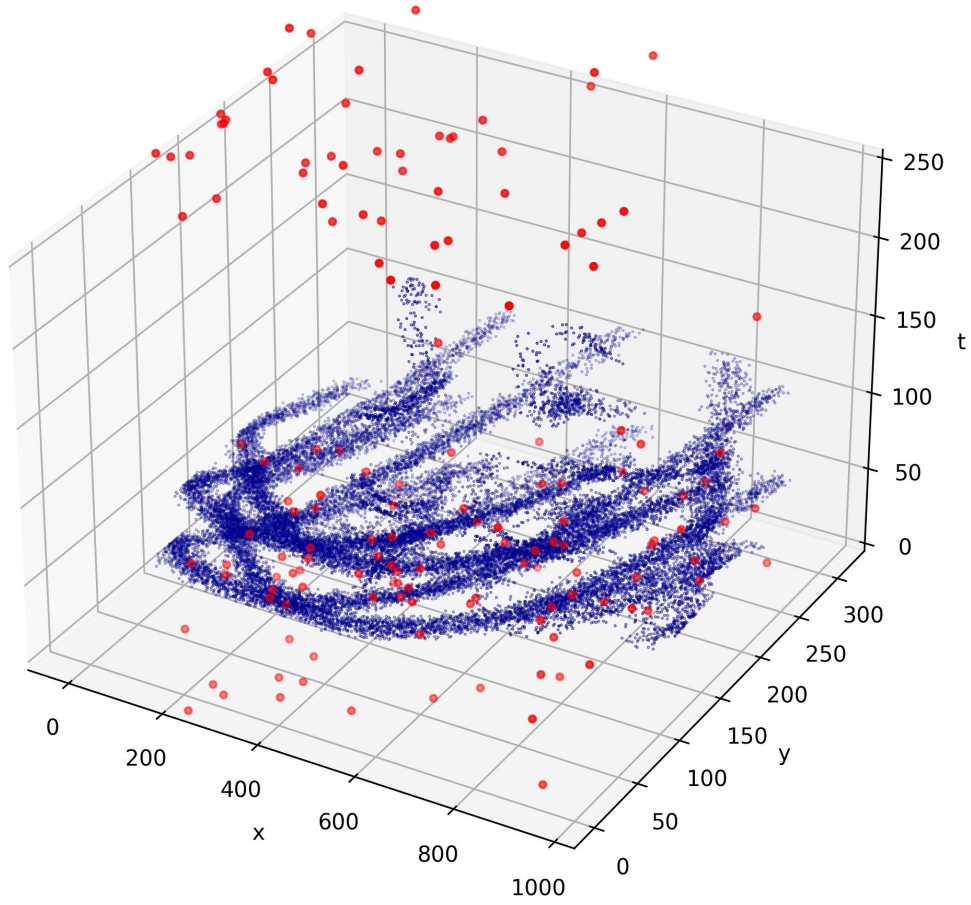

Figure 4-25: Example hit pattern in 3D from one π track. The blue points are $\mathcal{O}$(30k) support points in the KDE-based FastDIRC reconstruction (see Sec. 4.5.3), and the red points are data. The red data points that are far away from any blue points are various noises, such as electronic noise and Cherenkov photons from other tracks.

reconstruction strategies have been proposed over the past two decades since the first BaBar DIRC was proposed. They vary in their philosophy and possess different strengths and weaknesses depending on features of the detector to which they are being applied. Therefore, they are often complimentary and should be encouraged for pursuit in parallel when resources permit. We provide a brief overview of a few of the proposed DIRC reconstruction strategies here.

- geometrical reconstruction: it aims to reconstruct the Cherenkov angle from the detected hits on a per-photon basis using the detector geometry directly. It is used for the BaBar DIRC [61], PANDA DIRC prototype [79, 80], and now GlueX DIRC. The method is quite robust because it builds up physical observables step-by-step from the ground up. It makes certain simplifications that require the application of various correction factors. It is the default reconstruction algorithm for the GlueX DIRC. It will be discussed in more detail below. This method is also sometimes referred to as the look-up table (LUT) method.

- FastDIRC method: it performs a hypothesis testing directly on the detected 3D hit pattern. For each charged particle, it first performs fast simulation of the 3D hit patterns under different particle mass hypotheses. From the simulated 3D hit patterns under different hypotheses and the observed hits, kernel density estimation (KDE) can be used to calculate the likelihood of the observed hits against probability density functions (PDF) based on the simulated 3D hit pattern of each particle hypothesis. This method was proposed in [81] and proved very useful in the design of the GlueX DIRC optical box. The method is being actively investigated to be applied to a real data scenario with the GlueX DIRC. Preliminary results show comparable performance with the default geometrical reconstruction method, with details to be discussed in the following sections.

- time imaging reconstruction: the basic concept is to compare the measured time of each hit to the expected time per-pixel per-particle hypothesis. The expected time can be calculated analytically or from a full detector simulation. It was proposed and is being used by the Belle II time-of-propagation (TOP) counter [82, 83] as well as the PANDA DIRC design [80]. Application of this method to the GlueX DIRC was documented in [76].

- deep learning-based method: this method was recently proposed in [84] and studied in simulation. The idea is to train a deep neural network directly from training samples constructed from simulation or real data directly. Then a classification of the particle hypothesis is done using this network. The method showed promising results in simulation and is currently being explored for application in real data.

4.5.2 Geometrical Reconstruction

The geometrical reconstruction aims to measure the Cherenkov angle directly for each hit on the photodetection plane. It is based on a look-up table (LUT) in which it stores all the possible photon directions that might have hit a certain pixel from the end of a certain bar. This table is created with two steps. First, with the GEANT4 detector simulation, a large number of photons from the end face of a bar (all from the bar center or uniformly over the end surface) are generated using photon guns. These photons propagate through the wedge and the optical box before hitting the PMT plane. Next, for each pixel-bar combination, all possible photon directions emerging from the end of the bar are collected (with some averaging for paths that are very close to each other) and stored. Therefore, this LUT creates the following map

$$(\text{pixel}_i, \text{bar}_j) \rightarrow [\vec{v}_1 = (v_{x,1}, v_{y,1}, v_{z,1}), \vec{v}_2 = (v_{x,2}, v_{y,2}, v_{z,2}), ...]. \tag{4.3}$$

For every pair of pixel-bar combinations (since for every pixel hit, we also know which bar the charged particle hit), the LUT returns a list of possible photon directions at the end of the bar. However, since the Cherenkov photons internally reflect in the bars multiple times before emerging from the bar end, the exact direction from which they are emitted is unknown. For each direction $\vec{v}_k$ at the bar end, there are $2^3 = 8$ possible emission directions from combinations of the signs of each component, *i.e.,* $\vec{v}_k = (\pm v_{x,k}, \pm v_{y,k}, \pm v_{z,k})$ where each combination of the signs is one possible emission direction. All eight possible directions are used to calculate the angles they form with the particle direction. Only one of them is of course the correct direction and all the others are known as *ambiguities*. The ambiguities can be reduced in the following ways. One can discard the Cherenkov angles that are too large or too small because they are unphysical. One can also use the timing information. For every photon path from the end of a bar end to a pixel, the propagation time is also stored. This time can be combined with the average propagation time (calculated using some average wavelength hence some average index of refraction in fused silica) of the photon in the bars using the direction of the particle track and simple geometry. This time can be compared with the actual measured pixel hit time to form a $\Delta t = t_{\text{measured}} - t_{\text{calculated}}$. Only paths centered around 0 are selected. After ambiguity reduction, all possible Cherenkov angles are considered in a Cherenkov angle histogram. There should be a peak near the expected Cherenkov angle (calculated from the particle momentum and mass hypothesis), some smooth background, and possibily some ambiguity peaks. The peak is fitted, typically with a Gaussian distribution, and the width of the peak is known as the single photon resolution (SPR). A likelihood under a particle hypothesis is then calculated from the fitted Gaussian PDF and the differences of the log-likelihoods can be formed for a pair of particle hypotheses. Separation power, when the truth labels are known (either from simulation or relatively pure samples of tracks), can be then calculated from the delta log-likelihood distributions. More details on the geometrical reconstruction method can be found in *e.g.,* [61, 80]. It has been implemented in the GLUEX software for DIRC reconstruction and is currently used as the default method.

4.5.3 FastDIRC Method: Simulation and Reconstruction

Method Overview

The FastDIRC method for DIRC reconstruction was documented in detail in [64, 81]. Here, I summarize the important ingredients of this method.

The basic idea of FastDIRC is to perform hypothesis testing directly on the observed 3D hits for each track. It requires the PDF of the expected distribution of the hits in 3D for each particle hypothesis, given the track momentum (magnitude and direction) and track hit position on the DIRC wall as the parameters of the PDF. Due to the intrinsic statistical nature of the production, propagation, and detection of Cherenkov photons, Monte Carlo simulation is an effective and appropriate tool to sample the underlying PDF. With a large enough sample from the PDF, techniques such as kernel density estimation (KDE) can be used to get an accurate enough estimate of the underlying PDF. Once the per-track PDFs under different particle hypotheses are available, likelihoods of the observed hits against those PDFs can be computed and hypothesis testing performed.

It is evident that the key ingredient in this whole procedure is the accuracy of the per-track per-hypothesis PDF, that is, whether the estimated PDF represents the *true* underlying PDF from which the observed hits are sampled. The quality of the estimated PDF depends on many factors, which can be grouped into the following categories:

Track-related factors. This category primarily concerns the quality of the momentum measurement of the track, such as the magnitude, track direction, and timing. This information comes from the tracking system and is therefore out of the control of the DIRC system itself, although the DIRC can in turn provide useful feedback to the tracking system, as will be discussed in the following section.

Model-related factors. This category includes everything that goes into the modeling of the DIRC system, such as the physics modeling of Cherenkov photon production (*e.g.,* Cherenkov photon wavelength spectrum given all the material properties) and the geometry of the detector system (*e.g.,* if there is any misalignment of the optical components). This requires careful study of the observed data against the model.

Sample size. The PDF is estimated from the simulated hits, also called the *support points*, using KDE. Hence, the accuracy of this estimation depends on the number of support points used in the estimation, the larger the better. However, full detector Monte Carlo simulation routines, such as Geant4, are too computationally prohibitive for this task because it requires the simulation of $\mathcal{O}(1\,\mathrm{M})$ support points per track per hypothesis in order to obtain a sufficiently good estimation of the underlying PDF needed to perform the hypothesis testing. On the other hand, the level of detail in such Geant4-based simulations can also be an overkill because the Cherenkov photon propagation, which is the most time-consuming part in full simulations like Geant4, to a good approximation follows simple optics such as reflection and refraction. FastDIRC uses a novel fast ray-tracing technique suitable for geometries such as the DIRC bars and greatly speeds up the photon simulation, which in turn enables this reconstruction procedure.

FastDIRC Simulation: Cherenkov Photon Production

To define a simulated Cherenkov photon, there are several things that need to be specified: the location (x, y, z) and direction (dx, dy, dz) of the emission, and its wavelength. The Cherenkov angle, θ_C, which is part of the emission direction definition, also has wavelength dependence due to the wavelength dependence of the refractive index:[2]

$$\cos\theta_C = \frac{1}{n(\lambda)\beta}. \tag{4.4}$$

Emission location. The entrance position and direction of the track into a DIRC bar is given by the tracking system. The track then undergoes multiple scattering, *i.e.,* Molière scattering, as it traverses the bar.[3] FastDIRC allows the option to turn on Molière scattering and perform Monte Carlo sampling of the emission locations this way. This is more accurate but also slower and may not be necessary. Therefore, FastDIRC also allows sampling of the emission locations at regular intervals assuming the track follows a straight path inside the bar.

Cherenkov photon wavelength. The intrinsic wavelength dependence of the Cherenkov photons is defined by the Frank-Tamm formula in Eqn. (4.2). However, not all photons are equally likely to survive the DIRC system or to be detected due to the wavelength dependence of the materials present in the system (see Fig. 4-20a) and the quantum efficiency of the PMTs for detection [71]. To speed up the simulation, FastDIRC samples the expected wavelength spectrum at detection, taking into account the wavelength dependence of the materials and the quantum efficiency of the PMTs, and propagates only these photons. The additional loss of photons is only due to geometry, not material or detection efficiency. Once the wavelength is determined, one can obtain the corresponding indices of refraction of this photon in fused silica [85] and in water (see *e.g.,* [86]). The index of refraction determines the Cherenkov angle, θ_C, as well as the propagation speed, which is important for the timing modeling.

Emission direction. Due to the Cherenkov cone characteristics, it is most natural to consider the emission direction in the track coordinate system first (*e.g.,* track direction as the z-axis) and then transform back to the detector coordinate system. In the track coordinate system, the polar angle is the Cherenkov angle θ_C which is determined by the wavelength-dependent index of refraction obtained from above and the particle speed β. The azimuthal angle ϕ is sampled uniformly.

FastDIRC Simulation: Fast Ray-Tracing

One of the most important speed-ups of FastDIRC compared to the full simulations like in Geant4 is its fast ray-tracing through the bars. The algorithm and implementation are described in detail in [64, 81]. The key idea is to "unfold" the many bounces a photon experiences inside a bar into a straight path through "expanded"

[2] This is the origin of the chromatic effect.

[3] It is extremely unlikely that a track is stopped inside a bar because they are $\mathcal{O}$(GeV) particles. FastDIRC ignores this possibility in the simulation.

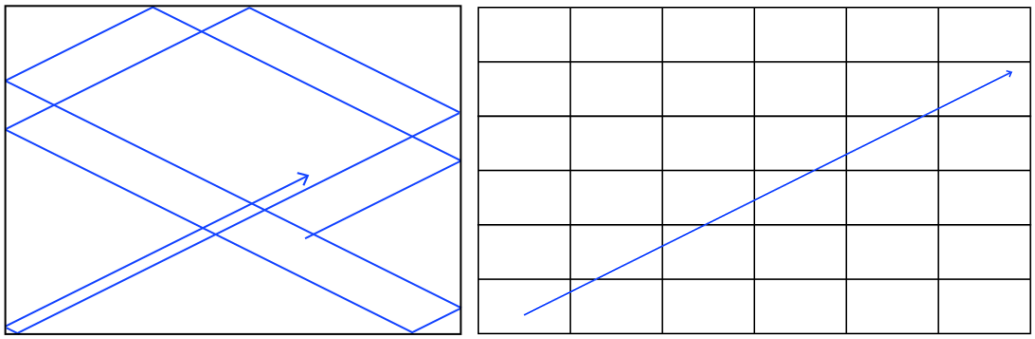

Figure 4-26: Graphical illustration of the fast ray-tracing algorithm in FastDIRC through the bars. Image source: [64].

bar volumes. A graphical illustration of this idea, taken from [64], is shown in Fig. 4-26. The propagation through the bar box wedge and inside the optical box is also done analytically.

KDE-based Reconstruction

Given the set of $\mathcal{O}(40)$ observed 3D hits and $\mathcal{O}(1\,\mathrm{M})$ simulated support points for a certain particle hypothesis (obtained from either the above described fast simulation or even full Geant4 simulation), one could perform a KDE-based likelihood calculation. The procedure in FastDIRC is shown as pseudocode in Algorithm 2.

Algorithm 2: Pseudocode for FastDIRC KDE-based reconstruction.

likelihood = 1;
for *every observed hit* O_i **do**
  for *every support point* S_j **do**
    $r_{i,j}^2 = \dfrac{(x_{O_i} - x_{S_j})^2}{s_x^2} + \dfrac{(y_{O_i} - y_{S_j})^2}{s_y^2} + \dfrac{(t_{O_i} - t_{S_j})^2}{s_t^2}$;
    likelihood *= $\exp\left(-\dfrac{r_{i,j}^2}{s_b^2}\right)$;

In the real implementation, the product of likelihoods is replaced with the sum of the logarithms of the likelihoods to avoid precision overflow. There is also an `if` statement before computing the exponential because most of the support points are far away from all observed hits under consideration. The exponentials in such cases would only yield very small numbers and slow the computation down significantly. The exponential term indicates that a Gaussian kernel is used, but other kernels were also explored and were found not to alter the results [62, 64]. s_x, s_y, and s_t are *not* resolutions on the hit position or timing, but are instead some scale parameters. However, the position and timing resolutions do set the natural scale for the scale parameters, so we do use the pixel size (6 mm) for s_x and s_y and the estimated timing

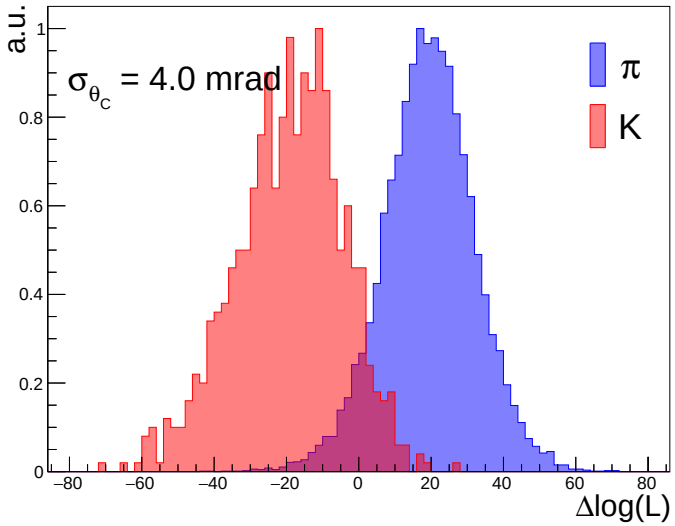

Figure 4-27: Example delta log-likelihood distribution for a selected sample of pions and kaons from GlueX data.

resolution (about 1 ns) for s_t. s_b is also a scale parameter that sets the scale of the hyper-distance $r_{i,j}$ and can be tuned (set to 1 as the default in FastDIRC).

Once the log-likelihood for a given particle hypothesis is computed, differences of the log-likelihoods for different hypotheses can be formed and the so-called delta log-likelihoods (DLLs), defined *e.g.*, $\Delta LL = LL_\pi - LL_K$ for pion hypothesis versus kaon hypothesis, can be analyzed to evaluate the separation power of the DIRC for particle identification.

To evaluate the DIRC performance in data, we first select relatively pure samples of pions and kaons from data without using the DIRC information. This can be done by analyzing exclusive final states such as ρ (ϕ) meson photoproduction as in $\gamma p \to p\rho$ ($\gamma p \to p\phi$) with subsequent decay of $\rho \to \pi^+\pi^-$ ($\phi \to K^+K^-$) for pion (kaon) samples. One could then apply further event selection, such as kinematic fitting and invariant mass cuts, to improve the purity of the pion/kaon samples. After such selections have been applied, the selected pion or kaon tracks are treated as having the truth labels. Then, we apply either the geometrical reconstruction or the FastDIRC KDE-based reconstruction to those tracks and analyze the DLLs distributions from these relatively pure pion and kaon samples. If the DIRC reconstruction works as expected, then the $\Delta LL = LL_\pi - LL_K$ distribution should have a positive mean with some spread for the relatively pure pion sample and should have a negative mean with some spread for the relatively pure kaon sample. An example DLL distribution following such a procedure for a selected set of pions and kaons of 3 GeV/c momentum hitting a selected region of the DIRC wall is shown in Fig. 4-27.

4.5.4 Application and Development of FastDIRC with Experimental Data at GlueX

Mapping the Geometry

FastDIRC was developed and used in the design of the GlueX DIRC optical box. As a result, the code itself (hosted in [87]) inherits features of early iterations of the detector design. In addition, some considerations were not important during the design stage (*e.g.,* matching its internal coordinate system with the Hall-D global coordinate system) but became so for the application of it to experimental data. The first step in the application of FastDIRC to the data of the GlueX DIRC was to establish a proper mapping to adapt the internal FastDIRC geometry to the actual detector realization. Here, we discuss a few notable considerations.

In the early iterations of the GlueX DIRC design, the long axis of the radiator bars was vertical and the optical box was situated at the top of the bar boxes, as evident in Figure 2-1 of [64]. In the actual realization of the detector, the long axis of the radiator bars is horizontal and the optical boxes are on the sides of the bar boxes as shown in Fig. 2-1 and Fig. 4-6. The early design and reconstruction studies using FastDIRC did not include the side mirrors because the shift of the hit pattern in the x(y)-direction in the FastDIRC (Hall-D) coordinate system was not very important for reconstruction, as explicitly stated in the caption of Figure 2-1 of [64]. The internal coordinate system in FastDIRC is left-handed and the azimuthal ϕ angle starts with $\phi = 0$ at the positive y-axis direction and increases towards the positive x-axis direction , whereas the the Hall D coordinate system is right-handed and the ϕ angle starts with $\phi = 0$ at the positive x-axis and increases towards the positive y-axis. Fig. 4-28, taken as an example from [88] which contains more details, shows a graphical representation of the mapping between the Hall D coordinate system and the internal FastDIRC coordinate system for the two bar boxes below the beamline and the South/Lower optical box. Finally, the optical box geometry, such as mirror positions, was slightly different in FastDIRC than in the final design, and needed to be matched. It is also worth noting that FastDIRC internally only accommodates two bar boxes and one optical box. For the GlueX DIRC, there are four bar boxes, two of which are paired with one optical box.

With all the considerations above in mind, we decided that the most sensible strategy for adapting FastDIRC for the GlueX DIRC at the current stage was to keep the internals of FastDIRC intact as much as possible to maintain its internal consistency and to carefully map the Hall D coordinate system convention onto the FastDIRC internal convention. The implementation of this adaptation is hosted at [89] with instructions for usage with GlueX data. Future integration into the GlueX system could consider a more ambitious overhaul of even the internals of FastDIRC to make it match more seamlessly with GlueX.

Track Timing

After the initial adaptation and validation of FastDIRC to reconstruct experimental data, we began more in depth investigation of the FastDIRC modeling and DIRC

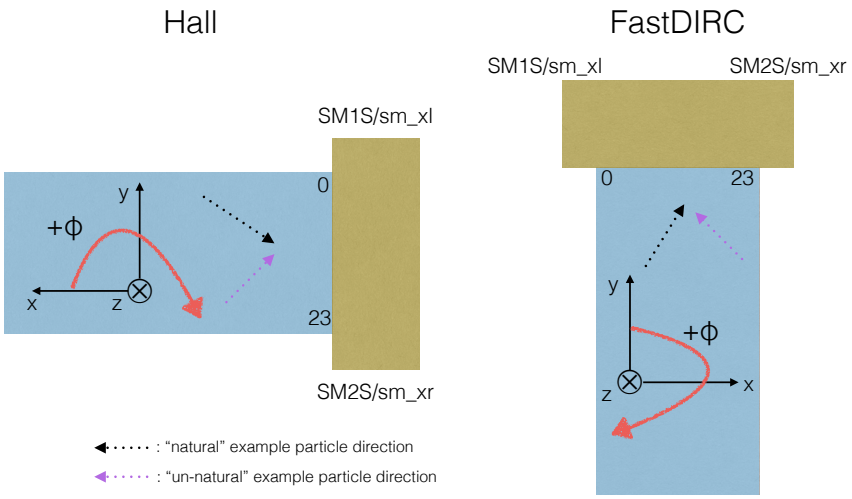

Figure 4-28: A graphical representation of the mapping between the Hall D coordinate system (left) and the internal FastDIRC coordinate system and convention (right) for the two bar boxes below the beamline and the South/Lower optical box looking from downstream towards upstream.

performance. One of the first things we did was to look at the raw 2D hit pattern on the PMT plane and the timing distributions of the observed hits in kinematic bins (binned by the location along the x-axis and by bar number) of the tracks hitting the DIRC wall with the selected pure samples of pions and kaons. We compared the observed distributions with the simulated ones from the FastDIRC model, which are the support points for constructing the particle hypothesis PDFs. Fig. 4-29 shows such a comparison for pion and kaon tracks in a particular x-bin of one radiator bar. The data distribution (blue) seemed like a smeared version of the FastDIRC Monte Carlo model (red) and it seemed to affect kaons more than pions. To match the observed data distributions, we considered applying additional time smearing of the PDF support points. Fig. 4-30 shows a comparison of the DLL distributions with and without the additional smearing where we applied a 1 (1.5) ns additional smearing to the pion (kaon) PDF support points. As expected, the additional applied time smearing made the DLLs wider, but the separation power was improved by about 30% due to the further separation of the means.

The observation that kaon tracks seemed to have worse timing characteristics than the pion tracks pointed to a potential difference in how pions and kaons were treated in the reconstruction. This observation prompted the tracking group to further investigate and they later found out the cause, which had to do with details of how the timing information was taken from different subdetectors. Specifically, when the track fit under a certain hypothesis succeeded, the timing information of the track was always taken from the subdetector with the best timing resolution, usually the time-of-flight detector. When the track fit failed, the timing information from the drift chambers, which had inferior timing resolution, was used. On average, kaon hypothesis fits failed more often than the pion ones (whose cause was still under

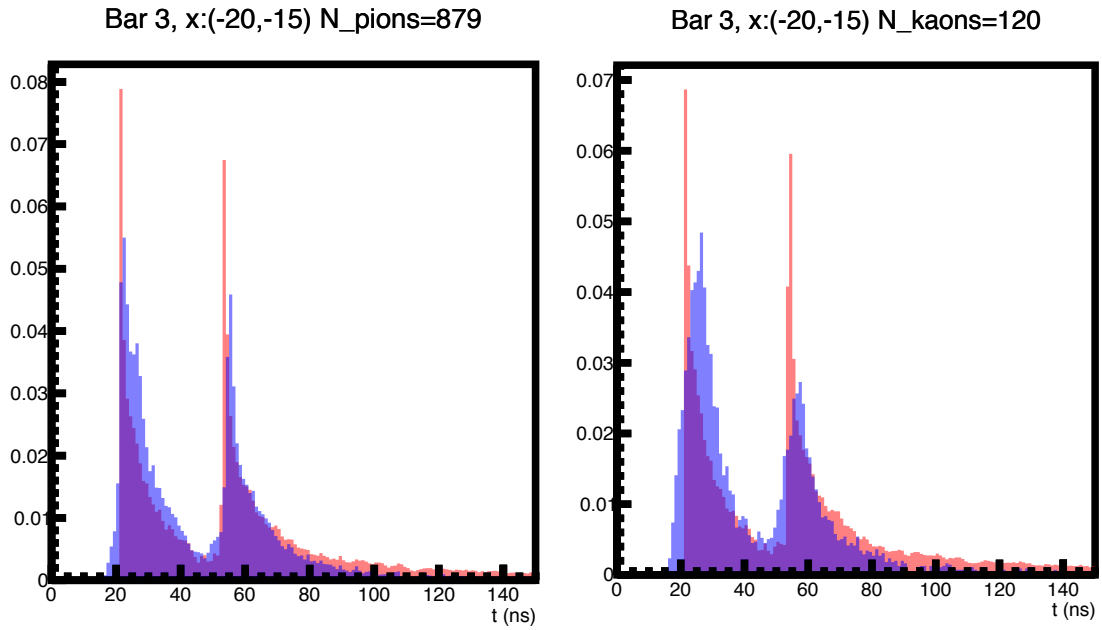

Figure 4-29: Comparison of hit timing for the selected 3 GeV/c pion tracks (left) and kaon tracks (right) in a (x, bar) bin. Blue is the observed data and red is from the FastDIRC PDF model. These plots were made before the track timing improvement.

investigation), resulting in worse kaon track timing. An improvement to the track timing was later implemented.

This observation happened during the analysis of the first 10-day commissioning run of the DIRC, and it demonstrated that the DIRC could already provide information not available using the baseline GlueX detector because the time-of-flight detector could not provide effective pion/kaon separation for tracks with momentum beyond 2 GeV/c. In addition, it demonstrated the complementarity of the FastDIRC method and the default geometrical reconstruction. This difference did not impact the geometrical reconstruction as much because the timing information was used only as a selection criterion for the observed hits and was not used directly in the likelihood calculation in the geometrical reconstruction.

Photon Yield

The yield of Cherenkov photons per track is an important performance characteristic of the DIRC system. However, there was no existing definition within the FastDIRC framework. I developed a working definition of photon yield within the FastDIRC framework based on the *neighborhood* of an observed hit. For any given observed hit, it is effectively a point in the 3D (x, y, t) space. We define the neighborhood of this hit to be a "cylinder" within a "height" in time, *e.g.,* ± 5 ns, and within a circular or rectangle "side" on the 2D PMT plane, *e.g.,* a circle of 8.5 mm in diameter. The PDF support points under a given particle hypothesis are also just 3D points in the

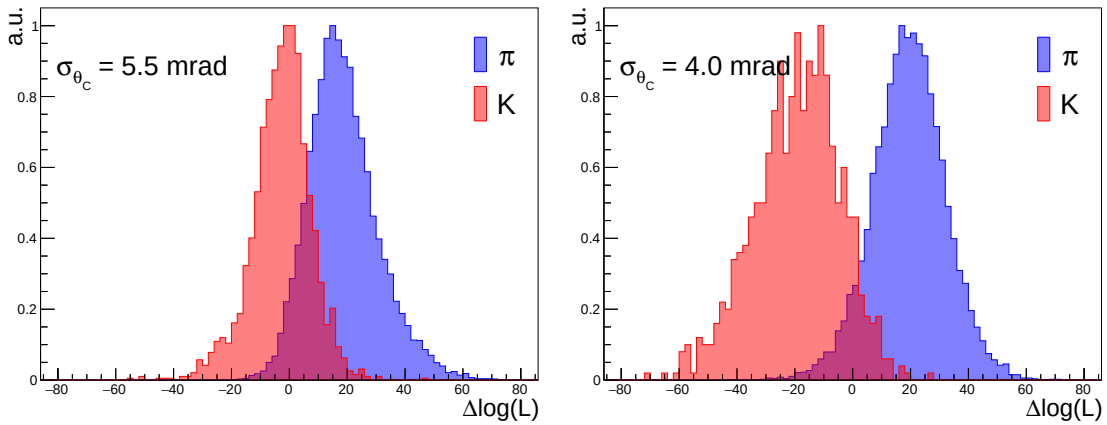

Figure 4-30: Comparison of the delta log-likelihood distributions for the selected pion and kaon tracks without (left) and with (right) the additional time smearing on the PDF support points.

(x, y, t) space. We could then define photon yield as the number of observed hits whose neighborhood contains at least one support point. Intuitively, the support points form "strips" in the 3D (x, y, t) space, as illustrated in Fig. 4-25, where the observed photons are likely to be. A real Cherenkov photon will land "near" those strips, whereas background hits in the PMT, such as Cherenkov photons from other particles in the event or simply electronic noise, should be randomly distributed in the 3D (x, y, t) space. Therefore, counting the observed hits that are "near" a "strip" should provide a good measure of the photon yield. Fig. 4-31 shows an example photon yield distribution using the above definition for tracks with momentum greater than 4 GeV/c integrated over a bar box. The data was taken from the Fall 2019 DIRC commissioning run. This showed good agreement with the photon yield distributions obtained from the default geometrical reconstruction.

Delta Observables

The performance of a KDE-based method, such as FastDIRC, depends critically on the accuracy of the modeling of the detector. When applied to data from a real detector where the underlying truth is not known, it is never the case that the modeling of the detector is perfect. Efforts are needed to investigate the discrepancies between the model and the observed data within the framework.

Unlike the geometrical reconstruction approach which tries to reconstruct the Cherenkov angle and performs hypothesis testing on this physically meaningful observable, the KDE-based reconstruction takes a "one-stop" approach where it attempts to perform hypothesis testing directly on the 3D hit information with no intermediate step. This "one-stop" approach allows naturally incorporating all the known effects in the modeling (such as the chromatic effect) without making the many simplifications in the geometrical reconstruction. It also has the potential of leveraging the full statistical power by using the low level hit information directly, provided

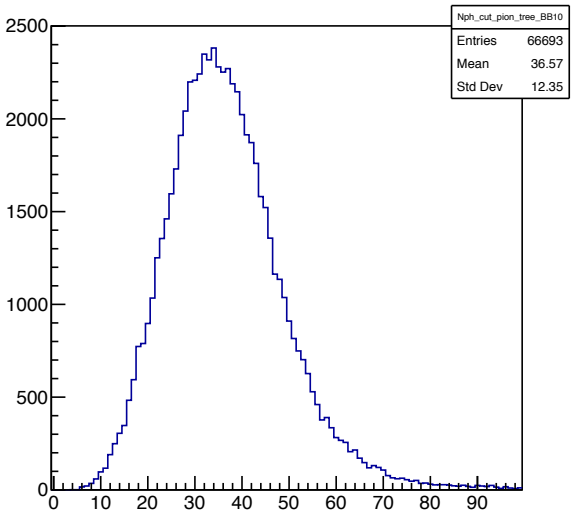

Figure 4-31: Preliminary photon yield distribution using the FastDIRC neighborhood definition for tracks with momentum greater than 4 GeV/c integrated over a bar box. Data from the Fall 2019 DIRC commissioning run.

the PDFs are known to a good enough precision. However, the lack of intermediate observables does make the diagnosis of the cause of the imperfect modeling more challenging, as is often the case during the early stages of understanding a detector. In addition, since KDE-based methods had not been used to reconstruct experimental data before and the application of FastDIRC, a KDE-based method, to reconstruct GlueX DIRC data was the first time, no diagnostic methods were readily available. I attempted some exploratory work along this direction.

Inspired by observables used in the geometrical reconstruction, I proposed the "delta observables" based on the neighborhood concept in the definition of photon yield within the FastDIRC framework. The idea is to simply look at the difference in x, y, and t between the observed hits and the support points within their neighborhoods. For each observed hit, we loop over the support points in its neighborhood and construct the following quantities: $\Delta x = x_O - x_{S,i}$, $\Delta y = y_O - y_{S,i}$, and $\Delta t = t_O - t_{S,i}$ for the i-th support point. We then construct them for all the observed hits. We can also compare the ones obtained from the data hits with the simulated hits that are generated directly from the PDF itself. Fig. 4-32 shows example distributions of such "delta observables" from a selected sample of kaon tracks using data hits and simulated hits. The data and simulated hits clearly show different behaviors in their peak locations and in their shapes, especially for Δy and Δt. The cause of the difference and the interpretation of such observables should be further investigated.

Reconstruction Performance and Future Improvements

The design goal of the per-track resolution of the Cherenkov angle for the GlueX DIRC was 2.5 mrad [62]. This corresponds to a π/K separation of about 3σ at 3.7 GeV/c. Calibration and performance analysis of the GlueX DIRC are still underway

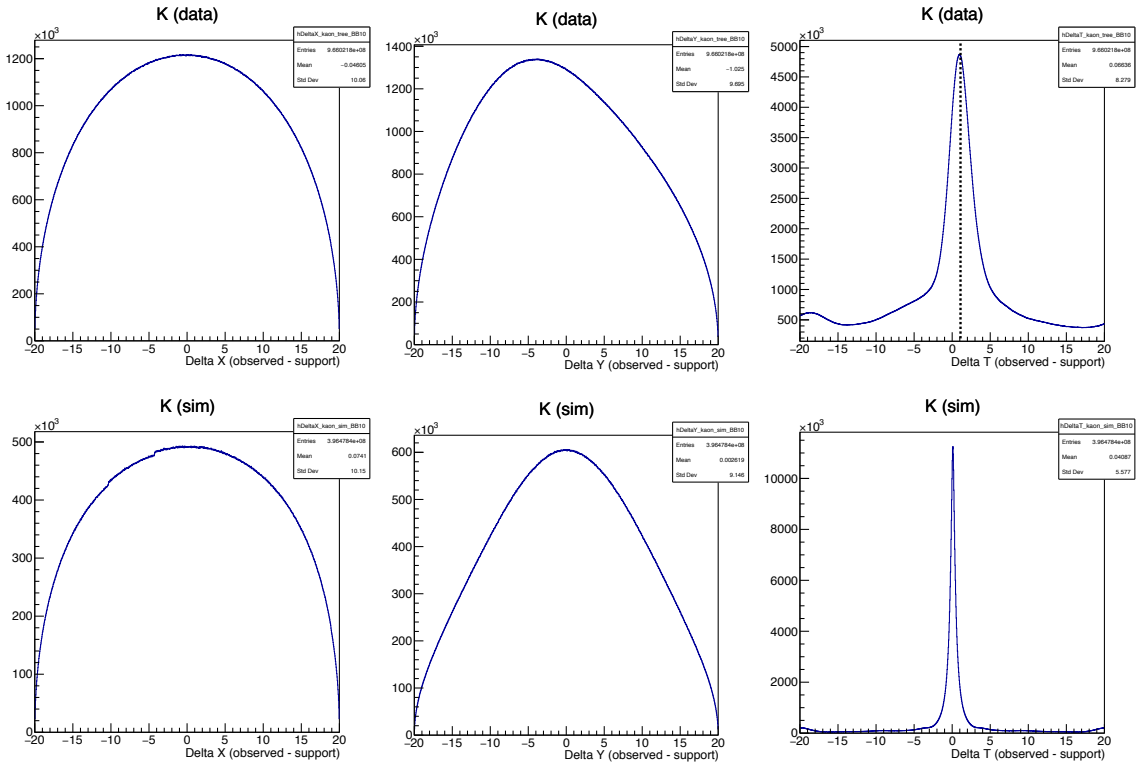

Figure 4-32: Delta observables with FastDIRC reconstruction from a selected sample of kaon tracks using data hits (top row) and simulated hits (bottom row).

as of the writing of this thesis. Preliminary analyses from both the geometrical reconstruction and FastDIRC have demonstrated 3σ separation in regions of the phase space around 3 GeV/c. The per-track Cherenkov angle resolution can be written in the following form

$$\sigma_{\theta_C}^{\text{track}} = \sqrt{\left(\frac{\sigma_{\theta_C}^{\text{photon}}}{\sqrt{N_{\text{photon}}}}\right)^2 + \left(\sigma^{\text{correlated}}\right)^2}, \tag{4.5}$$

where $\sigma_{\theta_C}^{\text{photon}}$ is the single photon resolution (SPR), N_{photon} is the number of Cherenkov photons and $\sigma^{\text{correlated}}$ is the correlated term. From preliminary analysis using geometrical reconstruction [90], the achieved single photon resolution is about 8 mrad, comparable to the expected values from Geant4 simulation. The analysis of the dependence of the per-track resolution on photon yield suggested that the correlated term was large [91]. The correlated term refers to factors that affect all photons and it includes effects such as the tracking system resolution, the alignment between the DIRC and the tracking system, and the alignment of the DIRC system itself (*e.g.*, the mirrors in the optical box and the alignment between the optical box and the bar boxes). Active studies are underway to understand the cause of the large correlated term.

I successfully adapted FastDIRC to reconstruct data from the GlueX DIRC

detector. I also began further development of the KDE-based reconstruction algorithm to provide further insight to improve its internal modeling to better describe the observed data and to provide insight into the DIRC performance. More efforts are needed to understand the new observables and develop other useful quantities to disentangle different effects in the FastDIRC simulation modeling. We should continue to compare the FastDIRC model with experimental data. The existing full Geant4 simulation of the DIRC should also be a useful tool to understand and improve the FastDIRC modeling.

4.6 Calibration and Alignment of the DIRC

With the successful installation and collection of commissioning data, the focus has shifted to the understanding of the commissioning data and the characterization and improvement of the reconstruction performance. The performance of the tracking system impacts the DIRC reconstruction regardless of the reconstruction method used. However, the exact strategies for calibration and alignment[4] of the DIRC optical system may vary depending on the reconstruction method. I discuss the efforts on those fronts in this section.

4.6.1 Tracking Improvement

The charged particle information is a critical input to the DIRC system. It not only defines the magnitude and direction of the momentum vector but also the impinging location on the DIRC wall. For charged particles that hit the DIRC, the last tracking point is the last FDC package, which is located at $z \cong 340\,\text{cm}$ in the Hall D global coordinates, while the DIRC wall is located at $z \cong 590\,\text{cm}$. Extrapolation has to be done over a distance of about 2.5 meters over which there also exists fringe magnetic field from the main 2T solenoid field. Performance of the tracking contributes to the correlated term in Eqn. (4.5) in the DIRC per-track Cherenkov angle resolution. Preliminary analysis of the February 2019 commissioning data suggested a correlated error that was larger than expected from simulation. Coincidentally, some GlueX collaborators were developing a GEM/TRD prototype for a future experiment. In coordination with them for the December 2019 commissioning run period, we installed the GEM/TRD prototype immediately in front of the DIRC wall, which provided another tracking point that can be used to evaluate the quality of track extrapolation. Analyzing this dataset led to another improvement in the track reconstruction. Previously, the Kalman filter algorithm was only run in the direction from the outermost to innermost hits because the focus was on the momentum at the interaction

[4]The terms "calibration" and "alignment" are used loosely in the context of this document. They broadly refer to efforts that attempt to improve the modeling of the DIRC system to better represent the actual detector. Calibration is often a more general term because it can refer to tasks that implement per-pixel timing correction as discussed previously. However, since the DIRC is essentially an optical system, its calibration mostly involves aligning the various components and correcting for the misalignment. Hence, the term alignment is also often used.

vertex not on extrapolation. The quality of the extrapolation was sufficient for the time-of-flight system because the TOF scintillator paddle size was large enough that the inaccuracy in the exact track hit position and direction was irrelevant. However, this was not case for the DIRC system because of its sensitivity to the track extrapolation. A "forward" Kalman filter was implemented by the tracking group to re-fit the hits from the innermost to the outermost, resulting in a better measurement of the track momentum exiting the last FDC tracking point. Studies found that for the few locations where GEM/TRD data were available, the track hit position achieved the expectation from the DIRC Technical Design Report [91, 62]. In the 2020 run, another layer was added to the GEM/TRD package that should give a measurement of the track direction just in front of the DIRC wall. Analysis of this dataset is still ongoing as of the writing of this thesis, as is the further improvement of the tracking code.

4.6.2 Calibration and Alignment in Geometrical Reconstruction

The key quantity used in the geometrical reconstruction is the Cherenkov angle. It is calculated with ambiguities given the PMT pixel and track information. The Cherenkov angles are then plotted and used in the hypothesis testing. Hence, the calibration and alignment strategy in the geometrical reconstruction framework mostly involves applying various correction factors to the Cherenkov angle. The corrections developed for the GlueX DIRC are described in [90] and I summarize a few of them here.

Per-PMT/pixel correction. This is a factor calculated per-PMT, or even per-pixel if a large enough data sample is available. For each PMT/pixel, the $\Delta\theta_C = \theta_C^{\text{reconstructed}} - \theta_C^{\text{expected}}$ can be calculated. Ideally, this should be centered at 0 with some spread. However, in a real detector, it can have a non-zero offset and also be different for different PMTs/pixels. This offset factor is calculated and applied on a per-PMT/pixel basis. The aim of this factor is to account for the misalignment within the optical box and between the optical box and the bar box.

Cherenkov ring fit correction. If one transforms the Cherenkov photon direction obtained from the LUT calculated in the detector/Hall D coordinate system to the charged particle coordinate system, those photon directions should form a ring, reflecting the characteristic Cherenkov cone. By performing a fit to the circle, offsets can be determined and applied to the charged particle direction. This factor aims to account for rotation offsets of the bars and the misalignment between the tracking system and the DIRC bars.

Chromatic correction. Cherenkov photons are emitted according to some wavelength spectrum, not monochromatically. Combined with the wavelength dependence of the index of refraction of the fused silica, it means those photons are emitted at slightly different Cherenkov angles and propagate at slightly different phase velocities. The Cherenkov angle and the propagation speed are correlated. Therefore, one can use the propagation time to apply a small factor to the Cherenkov angle.

The impact of those corrections on the SPR was visible but it did not seem to translate to as large of an impact on the separation power. Active work is ongoing.

4.6.3 DIRC Calibration and Alignment as an Optimization Problem

Problem Set-up

The FastDIRC method performs fast simulation of Cherenkov photons through the DIRC optical system in order to create the support points needed for the KDE-based reconstruction. It naturally incorporates information about the optical components involved, which are the subject matter for calibration and alignment. There are many such components (individual bars, bar box wedge, mirrors in the optical box, and the PMTs) and even more combinations of possible ways of misalignment. They all affect with varying importance the 3D information (x, y, t) of the support points, which in turn impacts the reconstruction. In addition, due to the complex geometry, the impacts on the support points from the rotations and offsets of the various optical components are often correlated. Therefore, it is natural to cast the DIRC calibration and alignment problem as a multi-dimensional optimization problem. The parameters are the various ways of misalignment, *i.e.,* the rotation and offsets of the various components. The objective is to make the simulated support points represent the underlying PDF of the observed points as best as it can.

Proof-of-Principle Study

As a proof-of-principle study, we set up a closure test with FastDIRC in the following way. First, we pick a set of charged particle tracks of one species, *e.g.,* pions, and simulate the "observed" hits with some injected "truth" offsets and rotations for some selected set of components, such as the angle of the 3-segmented mirrors. We also simulate the support points for those tracks under the "default" condition, *i.e.,* the nominal geometry with no offsets or rotations. Next, in the optimization loop, we treat those offsets and rotations as unknown free parameters. For each set of parameters requested by the optimization engine, we can simulate the "calibration" set of support points. We can then compute the log-likelihoods and form the delta log-likelihood for the "default" condition and the "calibration" condition and use the delta log-likelihood as the objective function to be optimized.[5] The closure is achieved if the optimization engine is able to recover the injected truth offsets and rotations.

The characteristics of this optimization problem are that the function is a black-box, non-parametric function in the sense that it is a Monte Carlo simulation with statistical nature and there is no functional form from the parameters to the objective function. It is also quite expensive to evaluate because even though it leverages the fast ray-tracing capability of FastDIRC, it still requires the simulation of a large

[5]The choice of using a "default" set and forming delta log-likelihood between the "default" and "calibration" sets is to circumvent the fact that the likelihood function as it is currently implemented is not properly normalized.

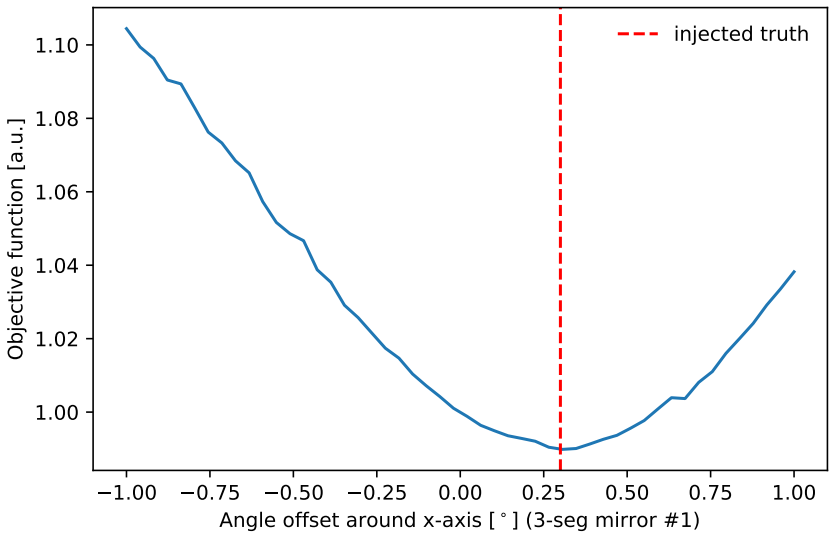

Figure 4-33: One-dimensional scan of the objective function as a function of an offset parameter in a DIRC calibration closure study.

amount of support points for a large number of tracks at each optimization iteration. Bayesian optimization described in Chapter 3 was used as the optimization strategy, using the implementation of `scikit-optimize` package [92].

Fig. 4-33 shows a scan of the objective function as a function of a rotation parameter. The red line indicates the injected truth value. This verifies that the chosen objective behaves as expected. Fig. 4-34 shows preliminary results of a simple three-dimensional problem (there are only three injected offset and rotation parameters). The 2D profiles in Fig. 4-34 are the posterior distributions of the parameters from the internal Bayesian model. The black points are the sampled sets of offset and rotation parameters. The 1D curves show the partial dependence on each parameter. The red dashed lines and stars show the injected truth values. We see that black dots are converging towards the red as desired. This is a relatively simple problem because it is only three-dimensional and the three parameters are expected to have a relatively large impact on 3D hit pattern because they are the angle offsets (along the three axes) between the bar box and the optical box. Systematic studies of more challenging and realistic cases are underway.

Towards Application to Real Data

With promising results from the closure tests in Monte Carlo, we applied this approach with GlueX data. A first challenge was to evaluate the effectiveness of the objective function used in the Monte Carlo closure study. Fig. 4-35 shows a comparison of the log-likelihood distributions with data (left) and simulation (right). We see that the data distribution has a much longer tail presumably because of a combination of the large amount of noise hits that are far way from the support points and the modeling was unsatisfactory. Due to this difference, we investigated other choices of objective

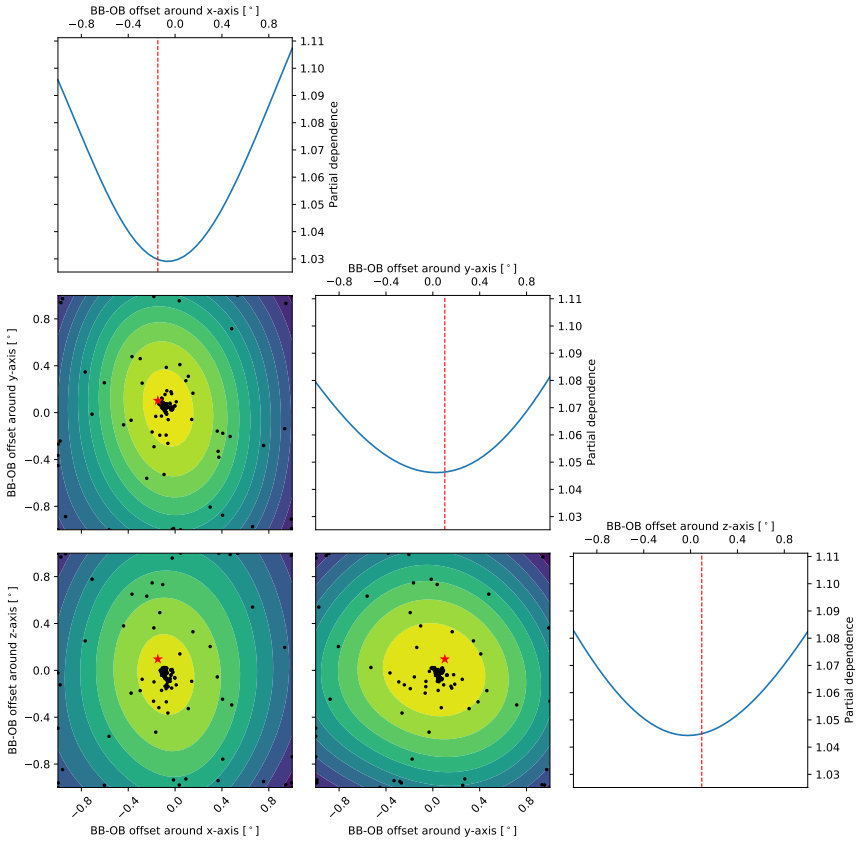

Figure 4-34: Preliminary results from a three-dimensional problem in a DIRC calibration closure study. See text for the explanation and interpretation of the figure.

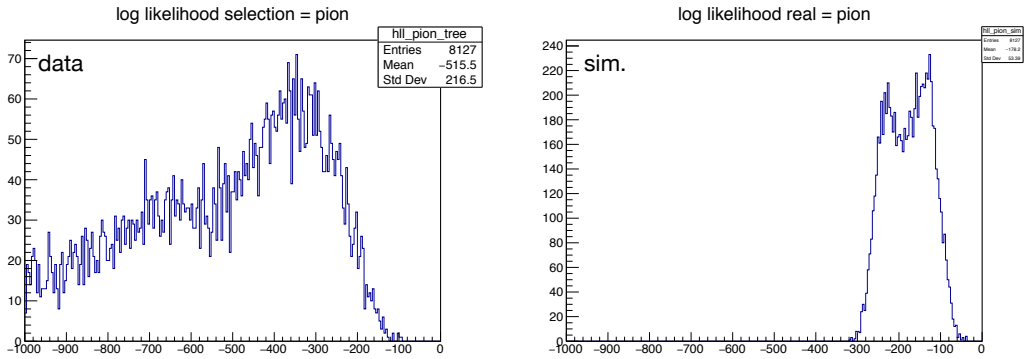

Figure 4-35: Comparison of the log-likelihood distribution in the FastDIRC method on data (left) and simulation (right).

functions by performing scans of the objective functions as a function of the offset parameter. Fig. 4-36 shows an example of such scans of different objective functions as a function of an offset parameter (angle offset of the middle mirror of the 3-segmented mirror). The blue and orange curves are variants of the objective function used above (with the "calibration" set and "default" set) while the red and purple curves are variants of using the separation power directly. In simulation, they all seem to agree with each other regardless of the choice of objective function. In data, however, the objective function used in the closure study does show a strong minimum, but it does not translate to the separation power, which is the ultimate objective of the calibration and alignment procedure. This suggests that more systematic studies are necessary to understand the choice of objective functions and how they impact the performance. One future direction is to properly normalize the likelihood function in the FastDIRC likelihood calculation and use the likelihood function directly.

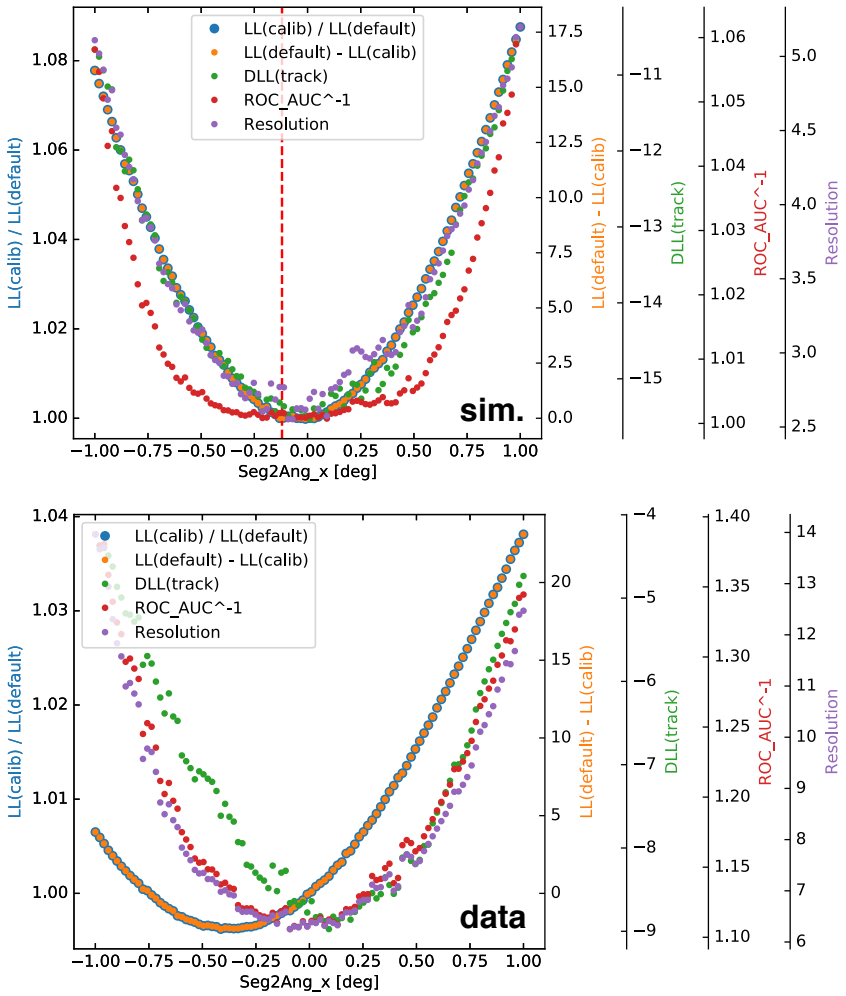

Figure 4-36: Scans of objective functions as a function of an offset parameter for simulation (top) and data (bottom).

Chapter 5

Search for Photoproduction of Axion-like Particles at GlueX

Originally proposed to solve the strong CP problem, axions and axion-like particles (ALPs) are hypothetical pseudoscalar particles found in many proposed extensions to the Standard Model of particle physics. Recently, ALPs with masses in the MeV-to-GeV scale received considerable interest as discussed in Sec. 1.3. This section presents a search for photoproduction of axion-like particles using data in photon-proton interactions collected by the GlueX experiment at Jefferson Laboratory in the $\gamma\gamma$ and $\pi^+\pi^-\pi^0$ final states of the ALPs. A general discussion and overview of the analysis strategy is discussed in Sec. 5.1. The dataset and event selection are discussed in Sec. 5.2.1 and Sec. 5.2.2, respectively. The normalization fits and the extraction of the mass resolution are discussed in Sec. 5.3. The Monte Carlo simulation samples and the extraction of the acceptance and efficiency are presented in Sec. 5.4. The method and validation of the signal searches are presented in Sec. 5.5. The systematic uncertainties are discussed in Sec. 5.6. The results of the searches are presented in Sec. 5.7. The search is currently blinded and the analysis is under internal GlueX collaboration review. We expect to set world-leading limits for most of the ALP masses over which the search is conducted.

5.1 Analysis Overview

This chapter presents a dedicated search for photoproduction of axion-like particles in the $\gamma\gamma$ and $\pi^+\pi^-\pi^0$ channels using the GlueX Phase-I dataset. The overall analysis strategy largely follows what was outlined in [23] and is summarized here:

1. define the dataset and perform event selection without examining the evidence for the existence of ALPs;
2. perform a bump hunt of the $m_{\gamma\gamma}$ ($m_{\pi^+\pi^-\pi^0}$) invariant mass spectrum in the mass range between the π^0 and η (η and ω) mesons to obtain local p-values and extract the observed upper limits on the ALP yields, if no evidence for ALPs is found;
3. compare the observed upper limits on the ALP yields and the expected ALP yields,

obtained via the normalization equation Eq. (5.1), and place upper limits on the a-gluon coupling, c_g/Λ.

Following Ref. [22], we denote the m_a-dependent mixing terms between the ALP and pseudoscalar mesons as $\langle \boldsymbol{a\pi^0} \rangle$ and $\langle \boldsymbol{a\eta} \rangle$. The expected ALP yield in final state $\mathcal{F}(=\gamma\gamma$ or $\pi^+\pi^-\pi^0)$ in a small bin of $[s,t]$, where $t \equiv (p_a - p_{\text{beam}})^2$, is related to the observed $\pi^0 \to \mathcal{F}$ and $\eta \to \mathcal{F}$ yields in the bin, $n_{\pi^0}(s,t)$ and $n_{\pi^0}(s,t)$, by [23]

$$n_a(s,t) \approx \left(\frac{f_\pi}{f_a}\right)^2 \left[|\langle \boldsymbol{a\pi^0} \rangle|^2 \frac{n_{\pi^0}(s,t)\epsilon(m_a,s,t)}{\mathcal{B}(\pi^0 \to \mathcal{F})\epsilon(m_\pi,s,t)} + |\langle \boldsymbol{a\eta} \rangle|^2 \frac{n_\eta(s,t)\epsilon(m_a,s,t)}{\mathcal{B}(\eta \to \mathcal{F})\epsilon(m_\eta,s,t)} \right] \mathcal{B}(a \to \mathcal{F}), \tag{5.1}$$

where f_π and $f_a \equiv -\Lambda/32\pi^2 c_g$ are the pion and ALP decay constants, $\mathcal{B}(\pi^0, \eta \to \mathcal{F})$ are the known meson-decay branching fractions [1], ϵ denotes the m_a-dependent product of the detector acceptance and efficiency, and $\mathcal{B}(a \to \mathcal{F})$ is the ALP-decay branching fraction [22]. Equation (5.1) assumes that t-channel processes are dominant, which is known to be true at GlueX energies for $-t \lesssim 1\,\text{GeV}^2$ [26]. For the $\pi^+\pi^-\pi^0$ decay, the ALP-pion mixing term is negligible and can be ignored.

5.2 Dataset and Event Selection

5.2.1 Dataset

This analysis considers the entirety of the GlueX Phase-I production dataset, including Spring 2017, Spring 2018 and Fall 2018 datasets. The process of going from raw electronic signals to reconstructed physically meaningful quantities is described in some detail in Chapter 2 and summarized in Fig. 2-24. The data reconstruction is coordinated centrally by the GlueX collaboration through *reconstruction launches*. The next stage is up to the analyst to write reaction channel-specific programs, known as *plugins*, to perform crude selection of the data so that the selected events match the reaction topology under consideration, *e.g.,* having the correct number of charged particles and photons. Since it is often the case that many groups have shared interest in the same reaction channels, this stage is now also centrally coordinated by the GlueX collaboration through *analysis launches* using the `ReactionFilter` plugin. The selection requirements applied at this stage will be described in Sec. 5.2.2. The output of these analysis launches are analysis trees in ROOT format [93] corresponding to various reaction channels.

The two channels of this analysis are (i) $\gamma p \to p\gamma\gamma$ and (ii) $\gamma p \to p\pi^+\pi^-\pi^0$, followed by $\pi^0 \to \gamma\gamma$. Because these channels are always produced in the standard analysis launches due to the shared interest by other collaborators for different physics, we use the officially produced analysis trees as the starting point of our analysis. The trees are labeled as `tree_gg__B4` for the $\gamma\gamma$ channel and `tree_pi0pippim__B4` for the $\pi^+\pi^-\pi^0$ channel. Details about the reconstruction and analysis launches for each run period are listed in Table 5.1.

Table 5.1: Reconstruction (REST) and analysis launch versions used in this analysis.

	$\gamma\gamma$		$\pi^+\pi^-\pi^0$	
	REST	analysis	REST	analysis
Spring 2017	ver03	ver20	ver03	ver27
Spring 2018	ver02	ver02	ver02	ver02
Fall 2018	ver02	ver02	ver02	ver02

5.2.2 Event Selection

In this analysis, we start with the output trees produced by the `ReactionFilter` plugin used in the official GlueX analysis launches described in Sec. 5.2.1. We then apply further selection requirements on those trees, using the GlueX `DSelector` library which is adapted from `TSelector` library in ROOT. Some selection cuts are fiducial and are straightforward to apply, while others, such as the kinematic fit cut, are optimized based on their effects on the expected sensitivity of the search. Sensitivity here refers to the expected sensitivity of the search on c_g/Λ and it can be calculated without performing the bump hunt using the following procedure at each mass point m_a:

1. sum over the t bins on the RHS of Eq. (5.1) to obtain the expected ALP yield at each m_a as a function of f_a;
2. set the LHS equal to $2\sqrt{B}$, where B is the number of background candidates within $\pm 2\sigma(m_a)$ of m_a ($\sigma(m_a)$ is the mass resolution at m_a, see Sec. 5.4) determined from the fit to the $m_{\mathcal{F}}$ spectrum described in Sec. 5.3;
3. and finally, solve for f_a and hence c_g/Λ.

The expected sensitivity can then be used to optimize the event selection in the following steps:

1. choose a baseline set of selection cuts and compute the baseline sensitivity;
2. for each cut under consideration, calculate the sensitivities for each cut value;
3. decide on the action about each cut (apply or not, at what cut value).

Plugin Requirements

Before describing our event selection in detail, we first summarize the selection applied in the standard `ReactionFilter` plugin in the official GlueX analysis launches. We adopt all the default selection criteria in the `ReactionFilter` plugin library that produce `tree_gg__B4` and `tree_pi0pippim__B4` analysis trees. The details of this plugin and its analysis cuts are described in Ref. [94]. Since the basic analysis cuts are on the particle level and the same for both the $\gamma\gamma$ and $\pi^+\pi^-\pi^0$ channels, we list the relevant ones together in Table 5.2. The energy threshold for photon reconstruction is set as 100 MeV for both the BCAL and FCAL to ensure good reconstruction efficiency. There are a number of Δt timing cuts, where Δt refers to the difference

between the observed particle hit time and the expected particle hit time for a number of subdetectors and particles. The particle hit time is the time between some start time, denoted here as the RF, and some stop time, detected by the subdetector of interest. The start time is referred to as RF because the CEBAF accelerator RF time is the most precise among all the timing system involved. Once the photon beam bucket containing the photon that initiated a certain reaction in the main GlueX detector is identified, often via the start counter, the timing of this photon bucket is determined by propagating the corresponding electron beam bucket time using the CEBAF RF clock. These Δt variables provide loose particle identification selection. Additional PID selection is provided by cutting on the energy loss variable dE/dx in the CDC for pion and proton candidates.

The missing energy selection and kinematic fit cut both leverage the exclusivity of the reaction enabled by the hermetic angular coverage of the GlueX detector. The so-called *missing* four-momentum is the difference between the initial and final state four-momenta, $p_{\text{missing}} = p_{\text{init}} - p_{\text{final}}$. The missing energy is the energy component of the missing four-momentum and the missing mass squared is defined as $m^2_{\text{missing}} = p^2_{\text{missing}}$. Kinematic fitting is a regression process in which the four-momenta of the final state particles are fitted, taking as inputs the measured quantities and their error estimations, under kinematic constraints such as the conservation of total four-momentum of the reaction (see *e.g.,* Ref. [95] for details of the implementation of kinematic fitting in GlueX). The objective function is often defined as a χ^2 function and the optimizer aims is minimize the χ^2. The results of the kinematic fitting include the goodness-of-fit figures-of-merit, *e.g.,* χ^2/nDoF or the corresponding p values, and the fitted four-momenta of final state particles. At the plugin level, convergence of the kinematic fit is required, and we explore tighter requirements on the quality of the kinematic fit in our event selection procedure described in the next sections.

Sensitivity-based Selection

Following the plugin selection, we begin the sensitivity-based selection procedure. First, we choose the set of baseline selection criteria which are either fiducial cuts or loose selection cuts. They are summarized in the middle columns of Tables 5.3 and 5.4 for the $\gamma\gamma$ and $\pi^+\pi^-\pi^0$ channels, respectively. Next, we look at each of the variables to see if there is anything more to gain by altering the selection criteria. The results of the selection optimization are summarized in the right columns of Tables 5.3 and 5.4.

Below, we provide a brief summary of the meaning and rationale of the selection criteria.

Beam energy and Mandelstam t

The photon beam energy fixes s. For simplicity, we focus only on the coherent peak region as it has the highest yields and a small beam energy window also gives us less beam-energy-dependent corrections to consider. We require the beam energy to be in the range (8.0, 9.0) GeV, which corresponds to an s bin of (15.9, 17.8) GeV^2. For Mandelstam t, we follow the same analysis strategy as done in Ref. [23]. We discard the region $|t| \lesssim 0.1\,\text{GeV}^2$ because the efficiency is small and sharply varying

Variable Name	**Selection**
E_γ	> 100 MeV
γ Δt (BCAL/RF)	$\pm$ 1.5 ns
γ Δt (FCAL/RF)	$\pm$ 2.5 ns
$\pi^\pm$ Δt (BCAL/RF)	$\pm$ 1.0 ns
$\pi^\pm$ Δt (TOF/RF)	$\pm$ 0.5 ns
$\pi^\pm$ Δt (FCAL/RF)	$\pm$ 2.0 ns
$\pi^\pm$ Δt (SC/RF)	$\pm$ 2.5 ns
p Δt (BCAL/RF)	$\pm$ 1.0 ns
p Δt (TOF/RF)	$\pm$ 0.6 ns
p Δt (FCAL/RF)	$\pm$ 2.0 ns
p Δt (SC/RF)	$\pm$ 2.5 ns
$\pi^\pm$ dE/dx (CDC)	$< \exp(-7\,\lvert\mathbf{p}\rvert + 3) + 6.2$ [keV/cm]
p dE/dx (CDC)	$> \exp(-4\,\lvert\mathbf{p}\rvert + 2.25) + 1$ [keV/cm]
Missing energy	$-3.0 <$ Missing energy < 3.0 GeV
Kinematic fit	convergence

Table 5.2: Summary of selection cuts in the `ReactionFilter` plugin.

Variable Name	**Baseline Selection**	**Optimized Selection**
Beam energy	(8, 9) GeV	–
Mandelstam $-t$	(0.1, 1) GeV^2	–
Missing mass squared	(-0.05, 0.05) GeV^2	–
Vertex z position	(50, 80) cm	–
Vertex radial position	< 1 cm	–
Proton momentum	> 350 MeV	–
FCAL shower radial position	(25, 100) cm	–
BCAL shower z position	(150, 380) cm	–
Number of unused tracks	0	–
Unused energy	< 100 MeV	–
dist($x_4(\gamma_1)$, $x_4(\gamma_2)$)	> 0 cm	> 12 cm
Photon energy	(0.1, 10) GeV	(0.5, 10) GeV
Kinematic fit confidence level	$> 10^{-7}$	> 0.02

Table 5.3: Summary of selection criteria for the $\gamma\gamma$ channel. "–" in the optimal selection column indicates that there is no change from the baseline values.

Variable Name	Baseline Selection	Optimized Selection
Beam energy	(8, 9) GeV	–
Mandelstam $-t$	(0.1, 1) GeV^2	–
Missing mass squared	(-0.05, 0.05) GeV^2	–
Vertex z position	(50, 80) cm	–
Vertex radial position	< 1 cm	–
Proton momentum	> 350 MeV	–
FCAL shower radial position	(25, 100) cm	–
BCAL shower z position	(150, 380) cm	–
Number of unused tracks	0	–
Unused energy	< 100 MeV	–
Photon energy	(0.1, 10) GeV	–
dist($x_4(\gamma_1)$, $x_4(\gamma_2)$)	> 0 cm	> 12 cm
$m(\pi^0)$ (measured quantities)	(100, 170) MeV	(110, 155) MeV
Kinematic fit confidence level	$> 10^{-7}$	$> 10^{-3}$

Table 5.4: Summary of selection criteria for $\pi^+\pi^-\pi^0$ channel.

with $|t|$. We discard the region $|t| > 1\,\mathrm{GeV}^2$ because the approximation of t channel domination in Eq. (5.1) begins to break down.

Vertex position

The reaction vertex is defined as the point of closest approach (POCA) of the final state reconstructed tracks. For the $\gamma\gamma$ channel, there is only one charged particle in the final state, the recoil proton, so the reaction vertex is defined as the POCA between the proton track and the beamline. We place requirements on both the z position and the radial position r of the vertex to ensure that the reaction is consistent with having occurred within the liquid hydrogen target.

Photon shower position

The FCAL is azimuthally symmetric with respect to the beamline. There is a large electromagnetic background from the photon beam near the beamline, *i.e.,* at low radial shower distances. Both the energy resolution and the reconstruction efficiency degrade in the transition region between the BCAL and FCAL; therefore, we exclude the large radial distance photon showers in the FCAL and large z photon showers in the BCAL. We also place a cut on the z position of the BCAL photon showers in the backward direction to ensure full containment of the shower for better energy resolution and reconstruction efficiency. We also require the FCAL photons to be well separated spatially with the photon distance cut.

Minimum energy and momentum

Energy thresholds are required to ensure good reconstruction efficiency. Such cuts include the requirement on minimum proton momentum and the photon energy. Furthermore, we found in the selection optimization procedure that increasing the minimum photon energy requirement improves the expected sensitivity for the $\gamma\gamma$ channel. Fig. 5-1 shows that cutting harder than the default minimum shower energy of 100 MeV improves the sensitivity. This is because decay photons of π^0 and η

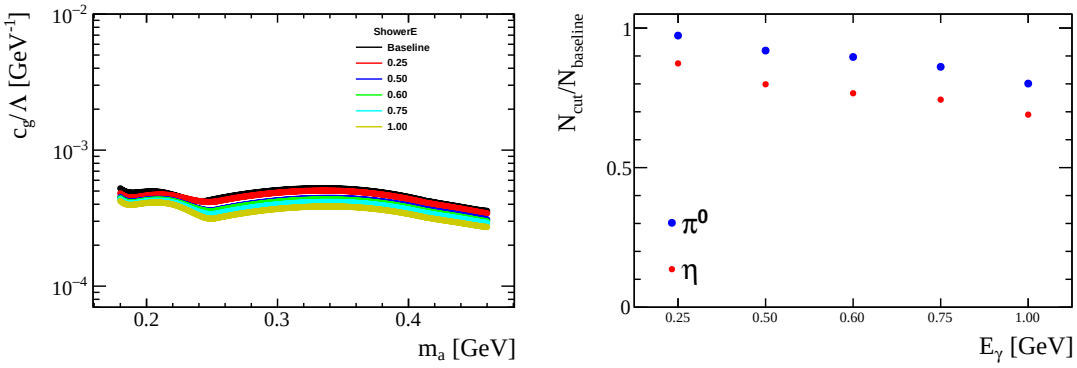

Figure 5-1: The (left) expected sensitivity and (right) efficiency for the π^0 and η mesons for various photon minimum energy cuts.

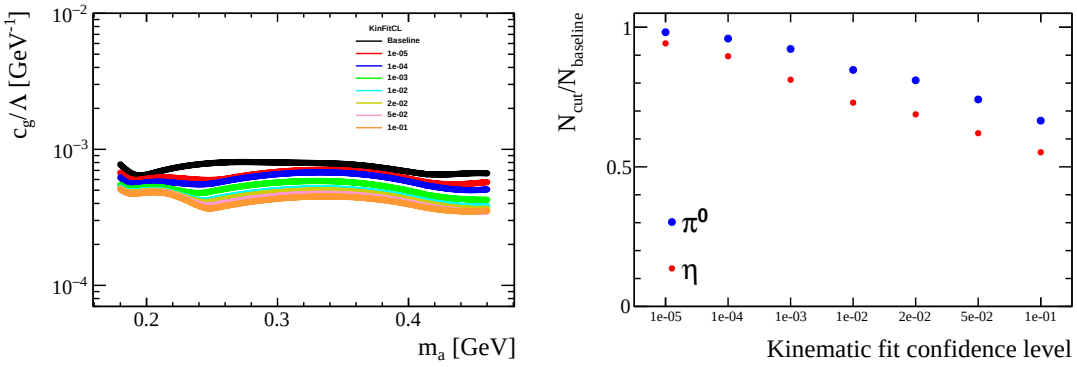

Figure 5-2: The (left) expected sensitivity and (right) efficiency for the π^0 and η mesons for various kinematic fit cuts.

mesons, as well as ALPs if they exist, tend to be more energetic compared to photons from other processes. We choose to require at least 500 MeV. In principle, cutting harder could improve the sensitivity at the $\mathcal{O}(10\%)$ level; however, this induces additional mass dependence in the efficiency which increases the normalization systematic uncertainty.

Reaction exclusivity

Thanks to the tagged photon beam and the nearly hermetic coverage of the GlueX detector, we can apply selection cuts on the exclusivity of the reaction. We require no extra charged particles other than the required ones in both channels and a maximum of 100 MeV of unassociated energy in the detector. The missing mass describes the exclusivity of the reaction and is correlated with the kinematic fit confidence level. We place a loose requirement on the missing mass squared and optimize the kinematic fit confidence level. Fig. 5-2 shows the expected sensitivity and efficiency for the π^0 and η mesons for various kinematic fit confidence level cuts for the $\gamma\gamma$ channel. A cut of confidence level greater than 0.02 is chosen as this is the loosest cut with near-optimal sensitivity; *n.b.*, the sensitivity estimates like Fig. 5-2 do not consider systematic effects, and therefore, we choose to use looser cuts if the estimated performance is nearly the same.

Table 5.5: Fiducial regions of the searches for both $a \to \gamma\gamma$ and $a \to \pi^+\pi^-\pi^0$ decays.

All searches	$8 < E_{\text{beam}} < 9\,\text{GeV}$
	$p_p > 0.35\,\text{GeV}$
	$-t < 1\,\text{GeV}^2$
	$2.45 < \theta_\gamma < 9.7°$, $11.66 < \theta_\gamma < 37.4°$
	$\alpha(\gamma_1, \gamma_2) > 1.15°$
$\gamma\gamma$ channel	$0.5 < E_\gamma < 10\,\text{GeV}$
	$-t > 0.2\,\text{GeV}^2$
$\pi^+\pi^-\pi^0$ channel	$0.1 < E_\gamma < 10\,\text{GeV}$
	$-t > 0.15\,\text{GeV}^2$

More detailed discussion on each of the selection cuts in both channels can be found in [96].

Accidental Subtraction

The GlueX photon beam is tagged by measuring the post-Bremsstrahlung electrons in the Tagger, described in Sec. 2.2. For each beam bucket, there often exists more than one tagged beam photon entering the liquid hydrogen target that could have initiated the reaction detected in the main GlueX detector. The experimenter has no way of knowing which photon in the beam bucket initiated the reaction on an event-by-event basis; therefore, each tagged photon in the correct beam bucket is considered with the final state particles to form a combination, or a *combo*, of an exclusive reaction. This inevitably results in the inclusion of the wrong combos consisting of the incorrect beam photon in the analysis. Such wrong combos can sometimes pass all the selection criteria described above if the wrong beam photons have energies close to that of the correct beam photon, contributing to the *accidental background*. Although the accidental background cannot be eliminated on an event-by-event basis, it can be subtracted *statistically*. The accidental background can be estimated by deliberately choosing the out-of-time photons, *i.e.,* those tagged photons from the adjacent beam buckets of the correct bucket, to form combos with the final state particles. The magnitude of such accidental background can be estimated by the number of out-of-time buckets included in the analysis. This statistical subtraction of the accidental background is known as *accidental subtraction*. We include the accidental subtraction in our analysis whenever applicable.

5.2.3 Fiducial Region

The fiducial region refers to the kinematic phase space that our analysis considers. It is defined by the selection criteria on the kinematic quantities of both the initial and final state particles, such as the beam photon energy, Mandelstam t, final state photon directions and so on. We summarize them in Table 5.5.

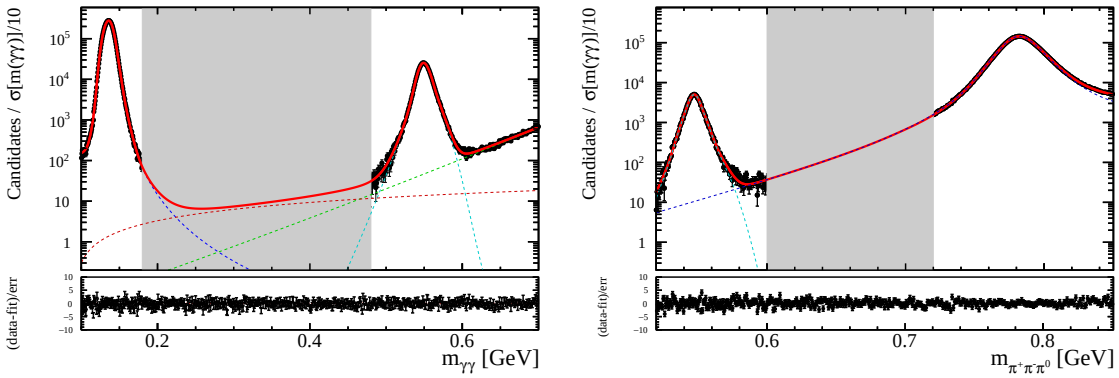

Figure 5-3: Fits to the (left) $\gamma\gamma$ and (right) $\pi^+\pi^-\pi^0$ invariant mass spectra after applying the event selection, including accidental subtraction, used to determine the π^0 and η yields (shown here integrated over t). The residuals account for both the statistical and modeling uncertainties.

5.3 Invariant Mass Spectra and Normalization Fits

After applying the event selection criteria described in Sec. 5.2.2, we proceed to examine the resultant invariant mass spectra of $m_{\gamma\gamma}$ and $m_{\pi^+\pi^-\pi^0}$, performing normalization fits to extract the π^0 and η yields, and the mass resolution.

5.3.1 Normalization Fits

As of the writing of this thesis, the search is blinded, so we need a way to validate the fits in the blinded region without unblinding the data. To do that, we allow one signal anywhere in the search region in the fit and look at the pulls. No information regarding the signal is reported to ensure blinding. Fig. 5-3 shows the $\gamma\gamma$ and $\pi^+\pi^-\pi^0$ invariant mass spectra obtained after applying the full selection and accidental subtraction. The increase in the uncertainty of the background fluctuation compared to using non-accidental subtracted spectrum is only about 2% from the 8 available out-of-time beam buckets in the sample. Binned maximum likelihood fits are performed using the RooFit library in ROOT to these spectra in bins of t to obtain the observed π^0 and η yields needed in Eq. (5.1); Fig. 5-3 shows the fit results integrated over t. It is worth noting that the bump-hunt procedure we use to search for ALP signals, described in detail in the next section, introduces additional complexity to the background models beyond what is used here. The fits shown in Fig. 5-3 are only used to determine the meson yields and the mass resolution.

The $\gamma\gamma$ fit model consists of the following components:

- the π^0 component is modeled by a sum of six Crystal Ball functions with three having power law tails on the low side of the peak and the other three having the power law tails on the high side of the peak;
- the η is modeled by a sum of one Gaussian function and three Crystal Ball functions with two having power law tails on the low side of the peak and the other on

the high side of the peak;

- the ω component is modeled by a sum of two Crystal Ball functions with one power law tail on the low side of the peak and the other on the high side of the peak;
- a linear background;
- and a signal component with a mass anywhere in the search region, but no information is reported about this component to ensure the blinding is maintained.

The $\pi^+\pi^-\pi^0$ fit model consists of the following components:

- the η meson is modeled by a double Gaussian;
- the true ω lineshape is taken from Ref. [97] and then convolved with a resolution function which is modeled by a sum of four Crystal Ball functions with two having power law tails on the low side of the peak and the other two having the power law tails on the high side of the peak;
- the remaining smooth background is modeled by a power law in Q;
- and a signal component with a mass allowed to be anywhere in the search region, but no information is reported on this component.

5.3.2 Mass Resolution

We define the mass resolution as half of the width of the region containing 68% of the probability of the signal shape. To assess the uncertainty on the mass resolutions for the π^0 and η mesons, we apply the bootstrapping method.[1]

For the $\gamma\gamma$ channel, the π^0 and η yields are 4.4 ± 0.1 million and 0.62 ± 0.02 million, respectively. Here, the systematic uncertainties, which are obtained by varying both the pseudoscalar and background models, are dominant. In addition, the $\gamma\gamma$ mass resolution is determined to be about 5.9 ± 0.0045 MeV and 9.1 ± 0.0352 at m_{π^0} and m_η, respectively (note that these are only statistical uncertainties obtained from this bootstrapping procedure, and their values demonstrate that the statistical uncertainties here are negligible). Monte Carlo simulation, described in Sec. 5.4, is used to interpolate between these values to obtain the resolution in the $m_{\gamma\gamma}$ region considered in the ALP search with a precision of 2%.

For the $\pi^+\pi^-\pi^0$ channel, the η yield is 70 ± 1 thousand, where again the systematic uncertainty is dominant. The $\pi^+\pi^-\pi^0$ mass resolution is determined to be about 5.7 ± 0.03 MeV and 10.5 ± 0.01 MeV at m_η and m_ω, respectively (note that these are only statistical uncertainties obtained from this bootstrapping procedure, and their values demonstrate that the statistical uncertainties here are negligible). Monte Carlo simulation, described in Sec. 5.4, is again used to interpolate between these values

[1]This involves re-sampling the $\gamma\gamma$ or $\pi^+\pi^-\pi^0$ sample with replacement $N = 100$ times to obtain N mass spectra. The same fitting procedure is performed on each of the N mass spectra, and histograms are filled with the π^0 and η resolutions from each sample. These histograms are fitted with a Gaussian model, the width of which is assigned to be the uncertainties.

to obtain the resolution in the $m_{\pi^+\pi^-\pi^0}$ region considered in the ALP search with a precision of 2%.

5.4 Monte Carlo Simulation

Monte Carlo simulation samples are used in this analysis to obtain the mass resolution function and the ratios of the product of the acceptance and efficiency of the SM pseudoscalars and the ALPs.

5.4.1 Sample Generation

We use the GENR8 event generator [98] to simulate ALPs with various masses. The GENR8 event generator assumes purely t-channel production, the same assumption under which the approximations made to obtain Eq. (5.1) are valid. Since the normalization is done in narrow $[s, t]$ bins as shown in Eq. (5.1), the production mechanisms do not need to be well understood. Once the $\gamma p \to pa$ photoproduction events are generated, they are processed through the standard GLUEX simulation pipeline described in detail in Sec. 2.4.3. The reconstruction and analysis software versions were chosen to exactly match the ones used to process the corresponding data. In addition, the effects of the interactions of other beam photons are included in the simulation via the inclusion of random trigger data into the simulation with the correct proportion as described in Sec. 2.4.1. The simulation pipeline was handled by the `MCwrapper` tool [28], which streamlines the input specifications, implements consistency with corresponding data reconstruction, seamlessly accesses computer offsite resources, and produces Monte Carlo samples in proportion to the actual data taken.

5.4.2 Mass Resolution Function

$\gamma\gamma$ channel

The simulated $\gamma\gamma$ mass spectra are fitted with a model which consists of a sum of two Crystal Ball functions with one having the power law tail on the low side and the other on the high side. The resolution is taken to be half the width of the region containing 68% of the probability density of the model. The left panel in Fig. 5-4 shows a comparison of the resolution in data and MC. We see that the MC slightly underestimates the resolution across the mass range of interest. Therefore, we shift the MC resolution curve by a function linear in mass to minimize the discrepancy at the π^0 and η masses, shown at right in Fig. 5-4. After the shift, the relative errors at the π^0 and η masses are 0.93% and 0.45%, respectively. We assign a conservative 2% error on the resolution across the mass range in the search to account for possible discrepancies in the mass dependence of the resolution in MC.

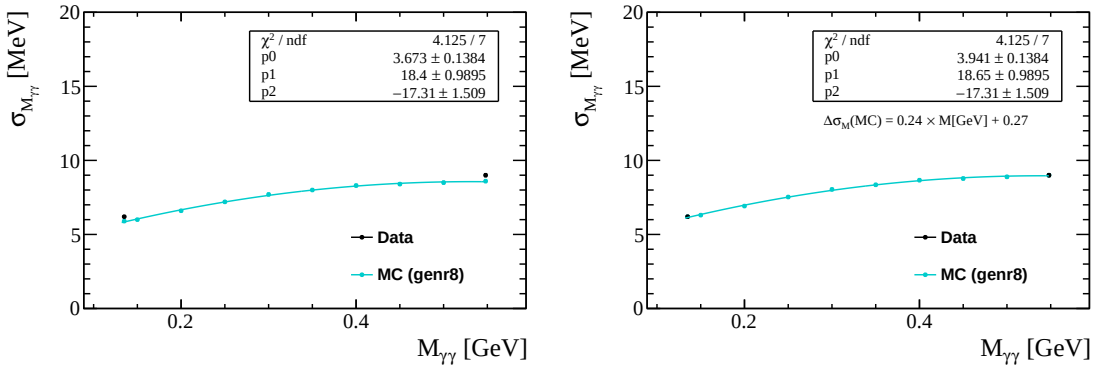

Figure 5-4: Mass resolution in MC versus data for the $\gamma\gamma$ channel (left) without and (right) with shifting all MC points down by the linear function shown on the figure.

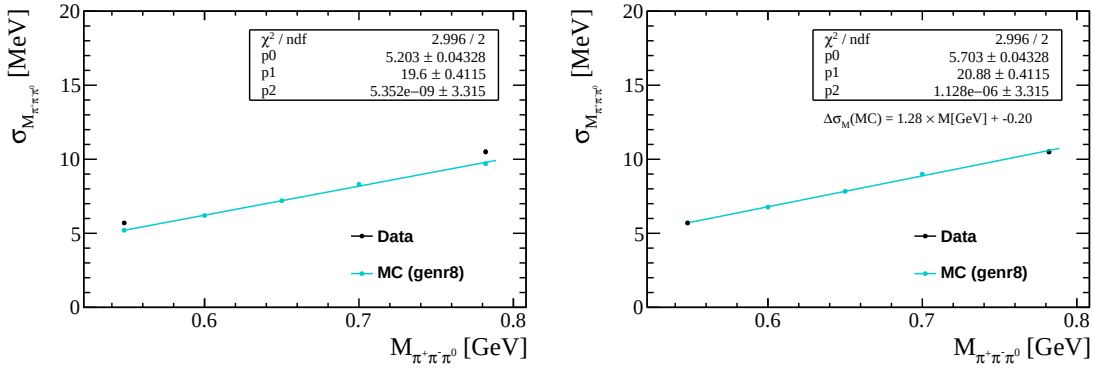

Figure 5-5: Mass resolution in MC versus data for the $\pi^+\pi^-\pi^0$ channel (left) without and (right) with shifting all MC points down by the linear function shown on the figure.

$\pi^+\pi^-\pi^0$ channel

The same procedure is used for the $\pi^+\pi^-\pi^0$ channel, see Fig. 5-5. The resultant relative errors at the η and ω masses are 0.1% and 0.8%, respectively. We again assign a conservative 2% error on the resolution across the mass range in the search to account for possible discrepancies in the mass dependence of the resolution in MC.

5.4.3 Acceptance and Efficiency

One of the advantages of the search strategy in this analysis is that only the relative efficiency to reconstruct the ALP and pseudoscalar-meson decays to the same final state is needed; knowledge of the absolute efficiencies is not required. The acceptance is defined here as the probability that a reaction producing an ALP in a $[s,t]$ bin will have all final-state particles in the fiducial region defined in Table 5.5. This acceptance is strongly dependent on m_a and requires careful treatment as described below. Reconstruction efficiency is defined as the probability of reconstructing a particle if it is in the fiducial region. Our choice of fiducial region is designed to

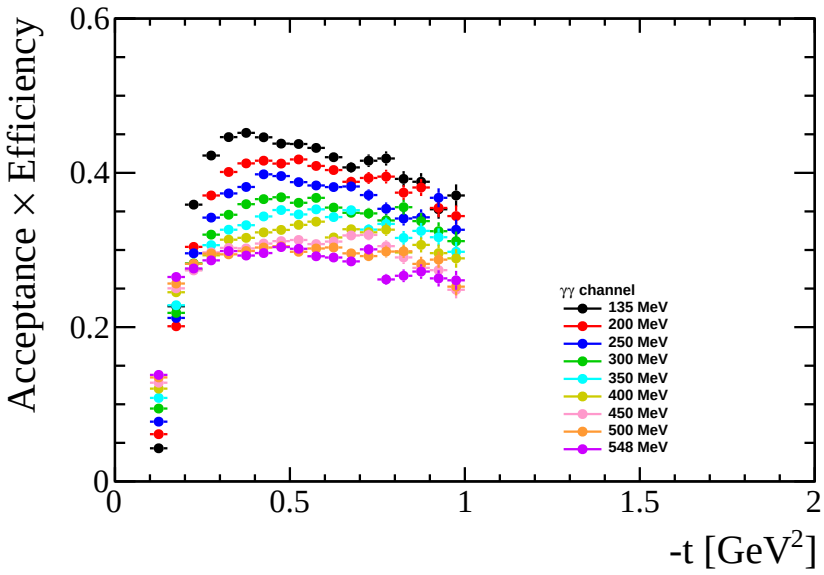

Figure 5-6: Product of the acceptance and efficiency in simulated $\gamma\gamma$ events.

minimize the m_a and t dependence, since only this dependence enters into Eq. (5.1). We denote ϵ, as in Eq. (5.1), to be the product of the acceptance and reconstruction efficiency, and we obtain this quantity and evaluate its uncertainties from simulation.

$\gamma\gamma$ channel

Fig. 5-6 shows the product of the acceptance and reconstruction efficiency for $\gamma\gamma$ events in bins of t. The error bars are binomial.[2] Fig. 5-6 indicates a rather strong m_a dependence in the t distribution due to the mass dependence of the acceptance. Kinematically, for low mass ALPs at small t, the decay photons from a have small opening angles and are traveling close to the beamline. We placed a minimal radial distance cut on the shower position in the forward calorimeter in order to reduce photon beam-induced electromagnetic background. This requirement is the major contributor to the low acceptance of those low mass ALPs at small t. In order to better control this effect, we perform a phase space Monte Carlo study to investigate this mass dependence.

The phase space Monte Carlo is generated using ROOT's `TGenPhaseSpace` class. It consists of two steps. The first step generates $\gamma + p \to p + a$ for a given E_γ, m_a, and t, fixing the outgoing polar angle θ of the a while the azimuthal angle ϕ is generated uniformly. The second step generates the $a \to \gamma\gamma$ decay. The 4-momentum of the ALP is given by the first step and its decay is purely phase space, handled by `TGenPhaseSpace`. We then apply the photon fiducial cuts for the FCAL, BCAL, minimum photon energy, and minimum momentum cut for protons. To get the acceptance value and its uncertainty in a t bin (we use 0.05 GeV^2 wide bins), we generate phase space MC events and evaluate the acceptance for each of the following

[2]For t bin i, $\sigma_i = \sqrt{p(1-p)/N(\text{generated})}$, where $p = N(\text{passed})/N(\text{generated})$.

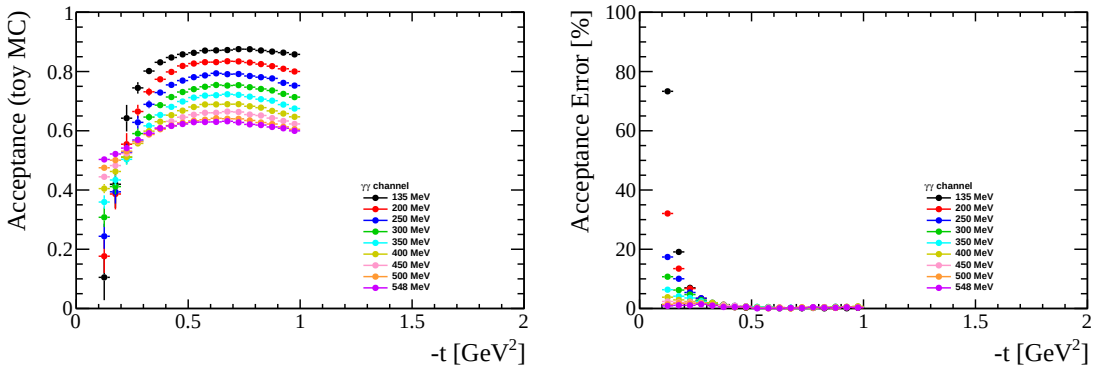

Figure 5-7: The (left) acceptance of $\gamma\gamma$ events from toy MC in bins of t and (right) the relative uncertainty obtained from the variation across the t bin as described in the text.

three cases:

(1) generate with t fixed to be the lower bin edge;

(2) generate with t fixed to be the upper bin edge;

(3) generate t uniformly in the t bin.

We then compute two differences: (i) the acceptance difference between case (1) and case (3) and (ii) the acceptance difference between case (2) and case (3). We take the result from case (3) as the acceptance value for each bin and we assign the uncertainty on the acceptance as half of the larger difference among (i) and (ii). We repeat this procedure for all t bins and all ALP masses. Figure 5-7 shows the acceptance in bins of t for three representative mass points and their relative uncertainties. The error bars are large for small t bins and low ALP masses because the detector acceptance varies the most rapidly at low masses near the beamline. In the search, we exclude the first two t bins due to low efficiency and large uncertainty and start with $t = 0.2$ GeV^2. If we divide out the detector acceptance effect in Fig. 5-7 from the product of the acceptance and efficiency in Fig. 5-6, we obtain the efficiency in bins of t in Fig. 5-8. The efficiency shows weak mass dependence; therefore, we can reliably interpolate this factor between generated MC mass points. We perform spline interpolation between the mass points for acceptance and linear interpolation for efficiency. We then obtain the product of acceptance and efficiency for all mass points of interest in bins of t by multiplying the interpolated acceptance and efficiency values as shown in Fig. 5-9

$\pi^+\pi^-\pi^0$ channel

A similar procedure described above for the $\gamma\gamma$ channel is applied to study the $\pi^+\pi^-\pi^0$ channel. Fig. 5-10 shows the product of the acceptance and reconstruction efficiency of $\pi^+\pi^-\pi^0$ events in bins of t at the generated mass points. Even though Fig. 5-10 does not indicate as strong an m_a dependence as that for the $\gamma\gamma$ channel, we perform a similar phase space MC study to evaluate the systematic uncertainty as done in the $\gamma\gamma$ channel. We also apply the polar angle acceptance cuts on the $\pi^{\pm}$. Figure 5-11

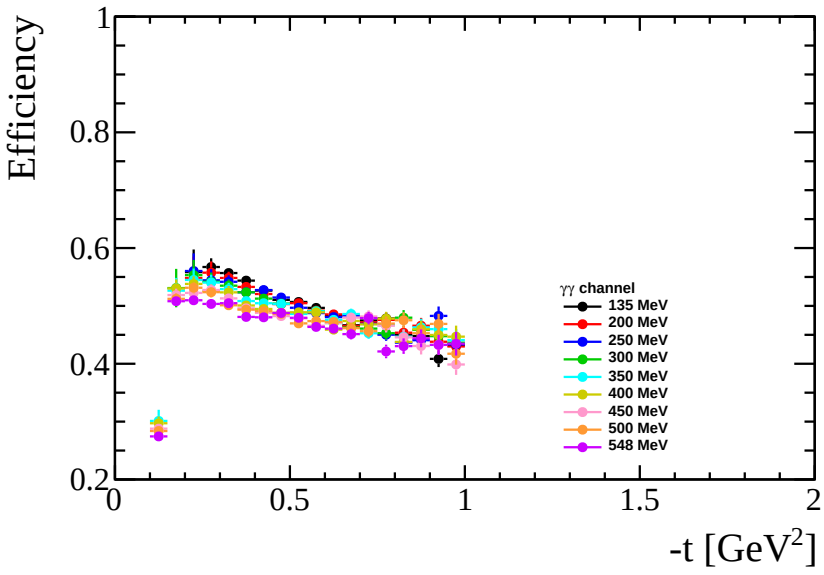

Figure 5-8: Efficiency of $\gamma\gamma$ events in bins of t.

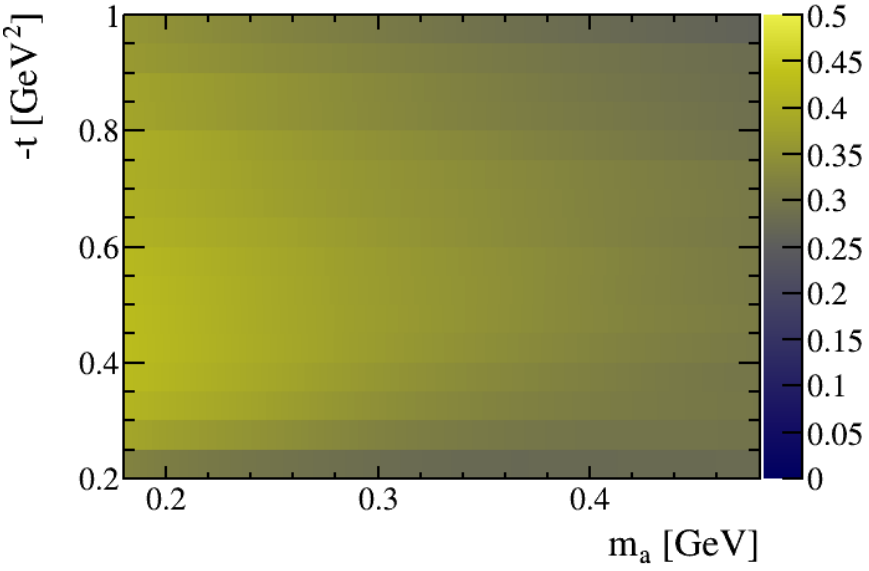

Figure 5-9: Acceptance $\times$ efficiency map in bins of t and ALP mass m_a for $\gamma\gamma$ channel.

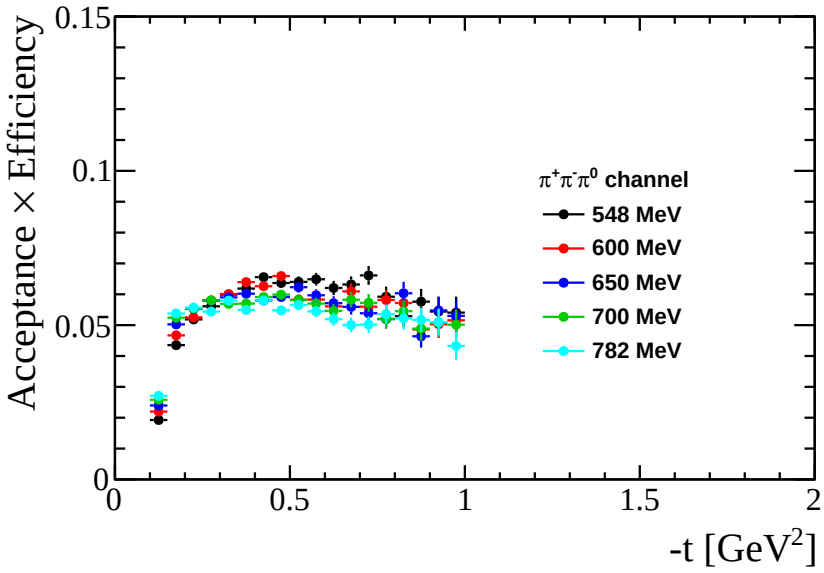

Figure 5-10: Product of the acceptance and efficiency of $\pi^+\pi^-\pi^0$ events in bins of t.

shows the acceptance in bins of t and their relative uncertainties. Again, if we divide out the detector acceptance effect in Fig. 5-11 from the product of the acceptance and efficiency in Fig. 5-10, we obtain the efficiency in bins of t in Fig. 5-12. The efficiency again shows weak mass dependence. The interpolated acceptance times efficiency map is shown in Fig. 5-13.

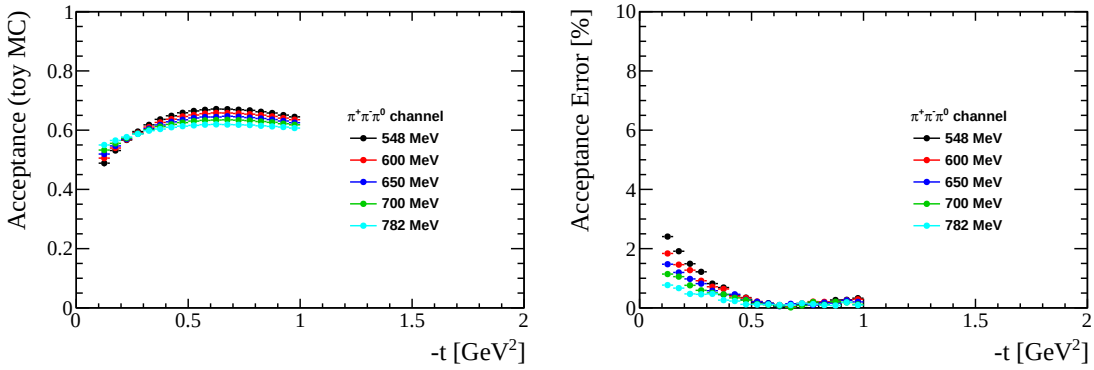

Figure 5-11: The (left) acceptance of $\pi^+\pi^-\pi^0$ events from toy MC in bins of t and (right) the relative uncertainty obtained from the variation across the t bin as described in the text.

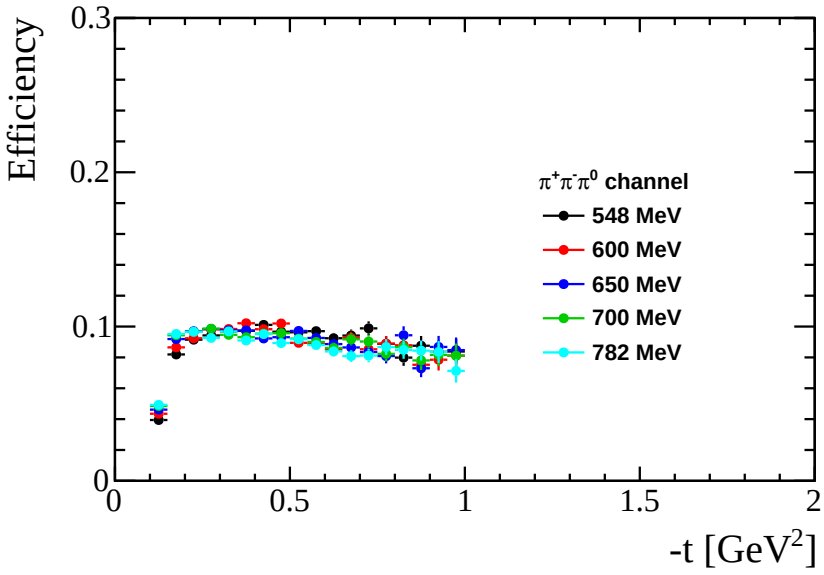

Figure 5-12: Efficiency of $\pi^+\pi^-\pi^0$ events in bins of t.

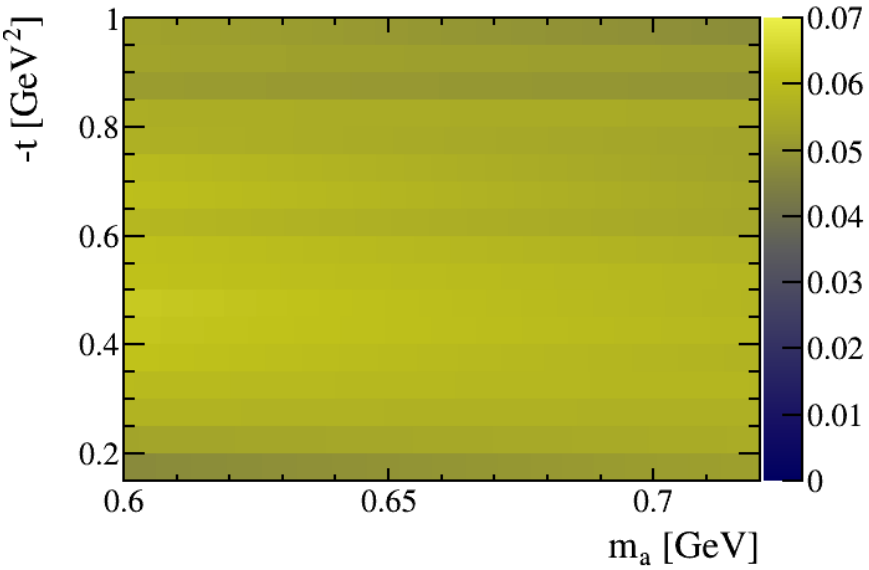

Figure 5-13: Acceptance × efficiency map in bins of t and ALP mass m_a for $\pi^+\pi^-\pi^0$ channel.

5.5 Signal Searches

5.5.1 Method

The signal-search, *i.e.,* bump hunt, strategy and method will follow that proposed in [99] and are similar to those used in [100, 101, 102]. We summarize the general principles of the strategy here and refer the interested reader to Ref. [99] for details:

- the mass spectra for ALP final states $\mathcal{F} = \gamma\gamma$ and $\pi^+\pi^-\pi^0$ are scanned in steps of about half the mass resolution, $\sigma(m_{\mathcal{F}})/2$, seaching for ALP contributions;
- around each mass hypothesis, m_a, a binned extended maximum likelihood fit is performed in a $\pm 12.5\sigma(m_{\gamma\gamma})$ or $\pm 7.5\sigma(m_{\pi^+\pi^-\pi^0})$ window around the m_a; a narrower window is used in the $\pi^+\pi^-\pi^0$ final state due to the small distance between the η and ω peaks compared to $\sigma(m_{\pi^+\pi^-\pi^0})$;
- the profile likelihood method is used to determine the local p-values and the ALP signal-yield confidence intervals (CIs); the trial factors are obtained using pseudo-experiments for each final state;
- the *bounded likelihood* approach [103] is used when determining the confidence intervals, which defines $\Delta \log \mathcal{L}$ relative to zero signal, instead of the best-fit value, if the best-fit signal value is negative.

The fit model at each mass hypothesis consists of the ALP signal component and a background component. The ALP signal mass distributions are well modeled by a Gaussian function, whose center is at the mass point m_a and width determined as described in Sec. 5.4.2.

The background models include the meson components described in the Sec. 5.3 and shown in Fig. 5-3. Specifically, they include[3] the π^0 and η resonances and the ω tail in the case of the $\gamma\gamma$ channel, and the η and ω resonances with a smooth background in the case of $\pi^+\pi^-\pi^0$ channel. In addition, the background models also include Legendre polynomial terms up to $\ell = 4\,(2)$ for $\mathcal{F} = \gamma\gamma\,(\pi^+\pi^-\pi^0)$. The data-driven model-selection process of Ref. [99] is performed. Specifically, the *aic-o* method in Ref. [99] is used, which penalizes the log-likelihood of each background model according to its complexity (number of parameters). This is done by adding a penalty term to the likelihood function

$$\Lambda = -2 \log \hat{\mathscr{L}} + c \cdot n(\text{par}), \tag{5.2}$$

where $\hat{\mathscr{L}}$ is the maximum likelihood value, $n(\text{par})$ is the number of unknown parameters in the model, and c is chosen to be 2 using the Akaike Information Criterion (AIC). The "o" in the *aic-o* method means that all odd Legendre polynomial terms up

[3]It also includes the out-of-time invariant mass spectrum as histogram PDF. As a result, even though the fit is technically to the in-time spectrum to preserve the Poisson statistics, it effectively is performed on the accidental-subtracted spectrum. We also put a model uncertainty for the large resonance PDFs in the peak region, which has negligible impact on the fit because they are far away from the fit windows.

to ℓ are always included because they are orthogonal to the signal component, which is an even function on [-1,1], onto which the invariant mass spectra are mapped before performing the fits. The uncertainty of this model-selection process is included in the profile likelihood following Ref. [104]. Specifically, Λ is calculated for all models, then the confidence interval is obtained from a profile of the penalized likelihood Λ, where the model index m is treated as a discrete nuisance parameter. The CIs are then defined as usual, *e.g.*, a 68.3% interval is defined as the region where $\Lambda - \Lambda(\hat{S}) \equiv \Delta\Lambda < 1$, see Refs. [99, 104] for more details. The *bounded likelihood* approach [103] is utilized to avoid unphysical (negative) limits. This involves taking $\Delta\Lambda$ starting from $S = 0$ rather than the best-fit signal yield value $\hat{S}$ if $\hat{S} < 0$ (when $\hat{S} > 0$ no modification is made). While this produces conservative limits, this approach has two benefits: it enforces that only physical (nonnegative) upper limits are placed on the ALP yields, and it prevents these limits from being much better than the experimental sensitivity if a large deficit in the background yield is observed.

5.5.2 Ensemble Validation

We perform a toy Monte Carlo ensemble study first to validate that the fitter produces desired properties in bias and coverage of the signal estimator. The study is done in the following way:

1. the normalization fit results, described in Sec. 5.3 using the nominal background model with and without additional injected signals with strengths of $S/\sqrt{B} = 0$, 1, 2, and 5, is used to generate an ensemble of N ($N = 1000$ in this study) toy datasets for each injected signal strength;

2. for each toy dataset, a fit to the full mass spectrum is performed using the nominal background model but excluding the $\pm 2.5\sigma$ signal region around the test mass, the fit result is used as the nominal background model at this mass;

3. the bump-hunt fit is performed and the signal estimator, its 68.3% and 90% CIs, and local significance, *etc.*, are reported;

4. we investigate the distribution of the pull which is defined as

$$\frac{\hat{S} - S(\text{injected})}{\sigma(\hat{S})},$$

   where $\sigma(\hat{S})$ is taken to be half of the full 68.3% CI;

5. we also investigate the coverage given at 90% CL, which is defined as the fraction of toy datasets of the ensemble where the corresponding CIs contain the injected signal (this is what matters for our limit setting).

We also run the search procedure on the background-only ensemble to obtain the following information:

- Expected limits and their uncertainty bands. After the search procedure is run on the N toy datasets, there is a distribution of the 90% upper limit (UL) on the signal yield at each mass point. The median is taken as the expected limit and the intervals which contain (15.9%, 84.1%) and (2.3%,97.7%) of the instances are taken as the $\pm 1\sigma$ and $\pm 2\sigma$ regions, respectively.

- Distribution of signed local significances. This is the distribution of all the reported local significances at all mass point of the ensemble. It will be normalized by the ensemble size and compared against the same distribution from the search on data.

- Distribution of the largest observed local significance. This will be used to account for the look-elsewhere effect in order to convert a local significance to a global one in data.

$\gamma\gamma$ channel

Pull distributions for an example test mass point with various injected signal strengths are shown in Fig. 5-14. They all roughly follow standard normal distributions as expected. We show the mean bias and 90% CIs coverage plots in Fig. 5-15 for additional mass values. The means are sufficiently unbiased and the 90% CIs show good coverage, which is what matters in the search. Finally, we return to the choice of $\ell = 4$ in the Legendre background. Given our use of 25σ-wide fit windows, this background can describe any background structures not accounted for in our nominal background model with a characteristic width of $\approx 5\sigma \approx 50\,\text{MeV}$. We do not expect any structures this narrow and consider this choice to be conservative; however, the impact on the sensitivity is small, so we prefer to be conservative here given the blinded nature of our search.

$\pi^+\pi^-\pi^0$ channel

A similar ensemble study is carried out for the $\pi^+\pi^-\pi^0$ channel. Due to the small distance between the η and ω peaks relative to the mass resolution in the $\pi^+\pi^-\pi^0$ channel, we used a narrower window of 15σ around each test mass point ($\pm 7.5\sigma$ on each side). Accordingly, we set the maximum allowed Legendre polynomial order in the background model to be 2, instead of 4. Example pull distributions for one test mass with various injected signal strengths are shown in Fig. 5-16. They all roughly follow standard normal distributions as expected. Figure 5-17 shows that the mean bias of the pull distributions and the 90% CIs look as expected across the mass range of interest.

5.6 Systematic Uncertainties

There are two main categories of systematics on the limit: one for the observed ALP yield and the other for the normalization of the expected ALP yield.

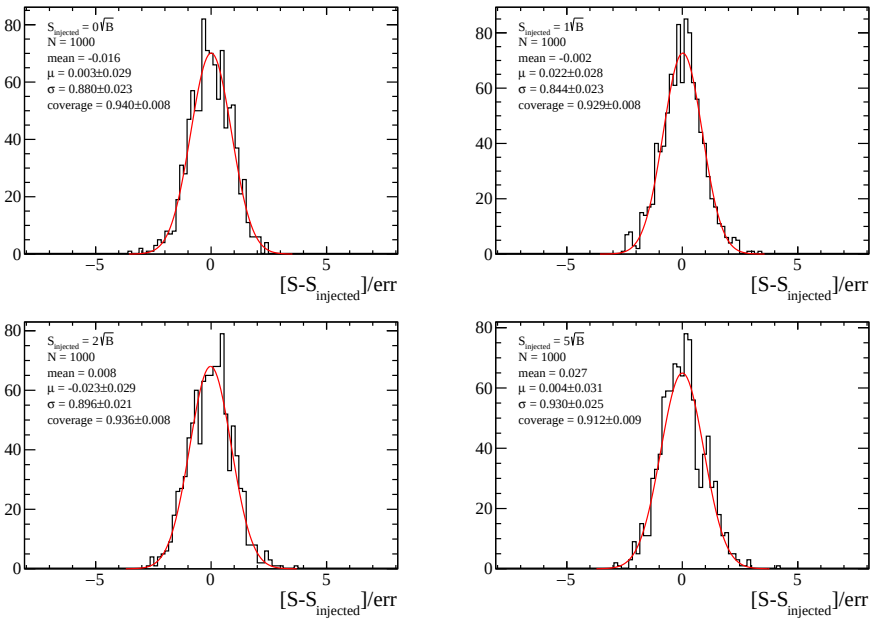

Figure 5-14: Pull distributions of the ensemble for test mass point at 350 MeV for various injected signal strengths.

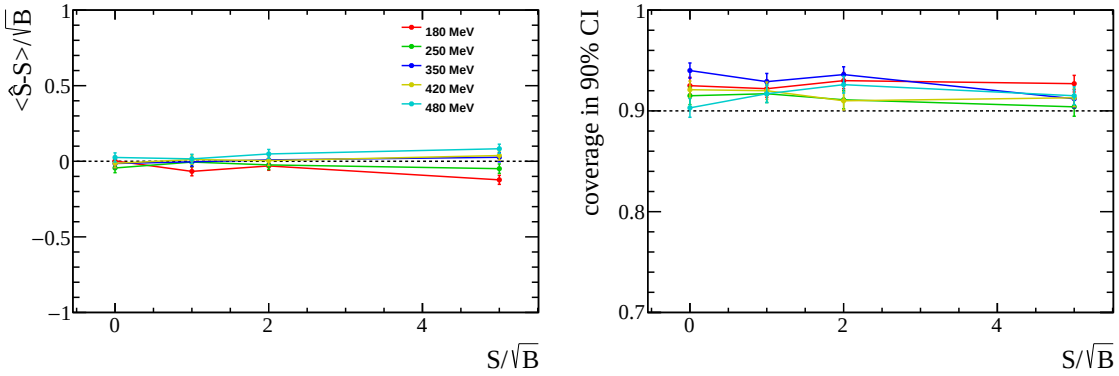

Figure 5-15: Mean bias of pull distributions (left) and 90% CI coverage for $\gamma\gamma$ channel ensemble study.

5.6.1 Systematic uncertainties on the observed ALP yield

There are two sources of systematics that can contribute to the bump-hunting fits: signal model and background model.

Signal model

As described in detail in Sec. 5.5, we use a simple Gaussian model for the signal model. At each test mass, the shape is fixed by the mass resolution at that mass point, which is extracted from Monte Carlo with data-driven corrections applied, described in detail in Sec. 5.4. In both $\gamma\gamma$ and $\pi^+\pi^-\pi^0$ channels, we assign a conservative 2% error for all mass points. As shown in Fig. 5-18, this translates roughly into a 1%

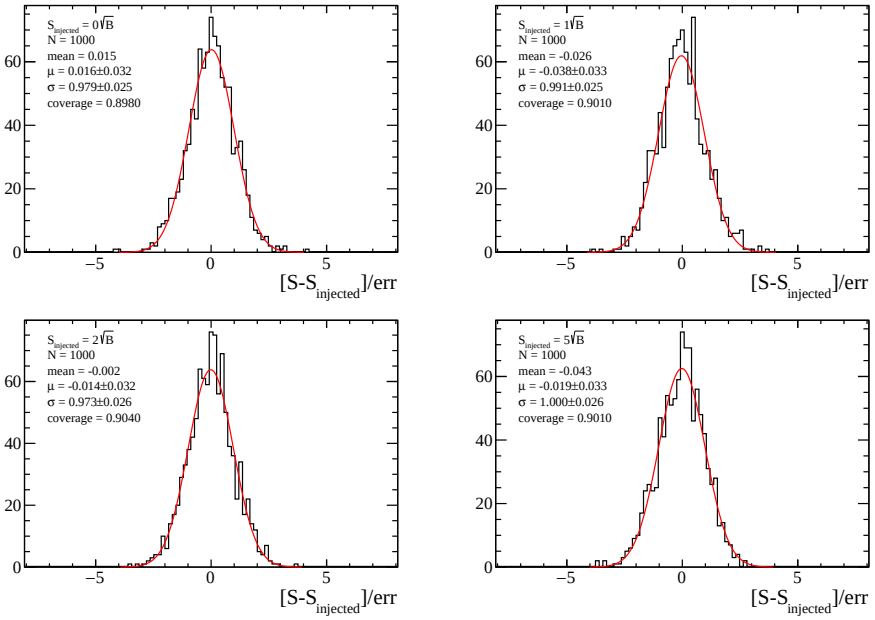

Figure 5-16: Pull distributions of the ensemble for test mass point at 650 MeV for various injected signal strengths.

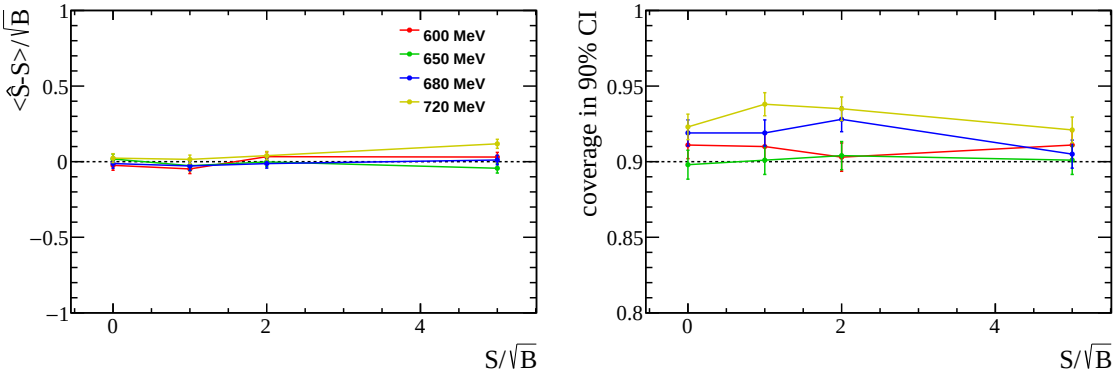

Figure 5-17: Mean bias of pull distributions (left) and 90% CI coverage for $\pi^+\pi^-\pi^0$ channel ensemble study.

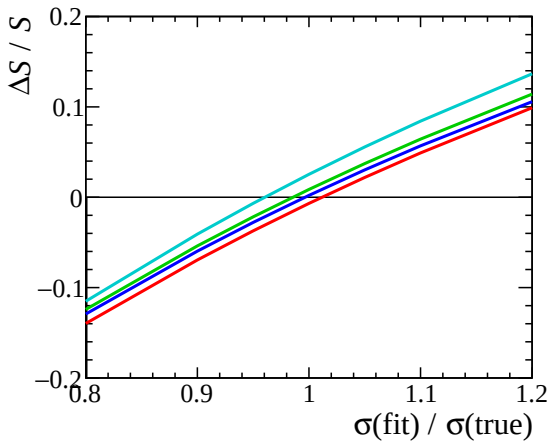

Figure 5-18: Relative bias in signal yield versus relative shift in resolution used in the bump-hunting fits for true signal rates of (cyan) 1, (red) 2, (green) 5, and (blue) 10 times $\sqrt{B}$, where B is the background yield in a $\pm 2\sigma$ window around the signal mass. Roughly, the relative error on the signal is about half the shift in the resolution.

error on the signal yield.

Background model

The handling of the background model in the bump-hunting fits is described in detail in Sec. 5.5. The procedure performs profiling of the likelihood function, treating model index as a nuisance parameter. Therefore, the background model uncertainty is built into this procedure. Here, we estimate the size of this uncertainty by comparing the full lengths of the 90% CI in the case where the profiling over different models is done versus the case where only the best model is considered using the same ensemble described in Sec. 5.5. Fig. 5-19 shows the results for the two channels for different mass points for a few signal strengths. The uncertainty is obviously mass dependent, and we note a range of 2–10% for the $\gamma\gamma$ channel and 2–8% for the $\pi^+\pi^-\pi^0$ channel.

5.6.2 Systematic uncertainties on the expected ALP yield

The expected ALP yields are shown in Eq. (5.1). There are three main sources that can cause systematic uncertainties on the expected ALP yield: acceptance times efficiency ratio, π^0 and η yields, and branching fractions.

Acceptance × efficiency ratio

We described in detail in Sec. 5.4 our procedure to obtain the acceptance × efficiency ratios from Monte Carlo. The largest uncertainties happen at low t bins for low masses due to errors on the steep gradient of the acceptance. We propagate the errors on the acceptance and efficiency in the product and show their impact on the

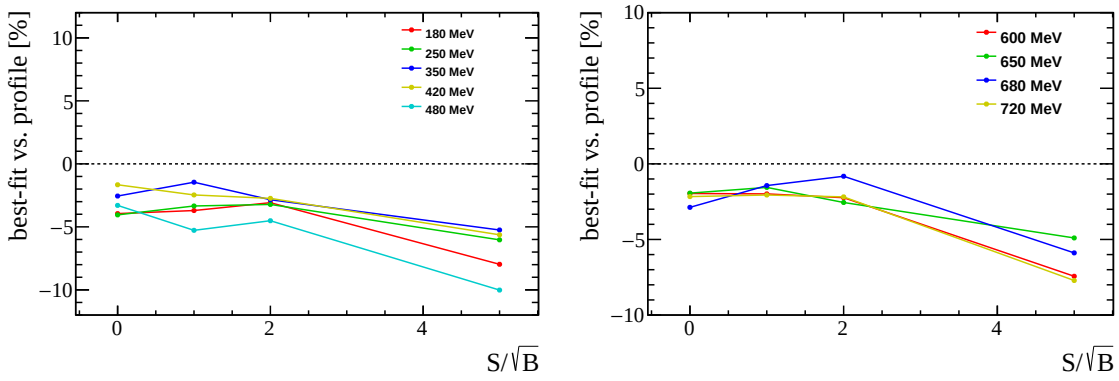

Figure 5-19: Relative difference in the 90% CI lengths in the nominal procedure where the profiling over different background models versus only using the best model for (left) the $\gamma\gamma$ channel and (right) the $\pi^+\pi^-\pi^0$ channel.

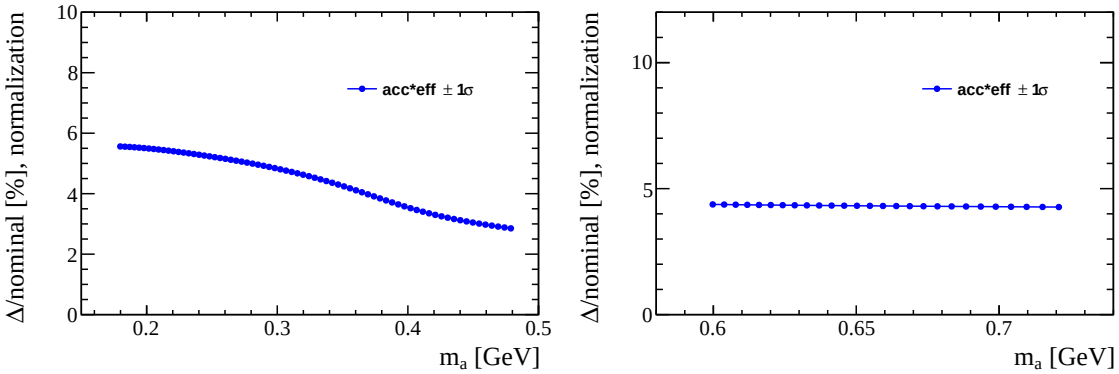

Figure 5-20: Relative uncertainty in expected yield from acceptance × efficiency for (left) the $\gamma\gamma$ and (right) the $\pi^+\pi^-\pi^0$ channel.

expected ALP yield in Fig. 5-20. This effect results in a mass dependent 2–6% effect for the $\gamma\gamma$ channel and a 5% effect for the $\pi^+\pi^-\pi^0$ channel.

π^0 and η yields

We estimate the systematics on the π^0 and η yields in the two channels by running the data fits with different π^0 and η shapes and background models.

The models used in the fits are described in Sec. 5.3 in detail. For the $\gamma\gamma$ channel, the nominal model used a linear background. We estimated this uncertainty by running the fits with second-order and third-order polynomials. For the $\pi^+\pi^-\pi^0$ channel, there is no explicit background model under the η peak other than the tail from the ω, which is modeled by a sum of four Crystal Ball functions for the resolution function. We estimated this uncertainty by running with two and six Crystal Ball functions. The π^0 and η shapes are modeled by a sum of Crystal Ball functions and/or Gaussian functions. We estimated this uncertainty by adding more or using fewer Crystal Ball functions.

Table 5.6: Summary of relative systematic uncertainties. Those specified as a range are mass dependent.

Source	$a \to \gamma\gamma$	$a \to \pi^+\pi^-\pi^0$
Signal model	1%	1%
Background model	2–10%	2–8%
Acceptance×efficiency	3–6%	5%
π^0 and η yields	3%	1%
Branching fractions	0.1–0.5%	1.5%
Total	5–12%	5–9%

The peaks are reasonably clean and we observe 3% variation on π^0 and η yields in the $\gamma\gamma$ channel and a 1% variation on η yield in the $\pi^+\pi^-\pi^0$ channel.

Branching fractions

There is some uncertainty on the branching fractions used in the calculation of expected ALP yields: $\pi^0/\eta \to \gamma\gamma$ and $\eta \to \pi^+\pi^-\pi^0$. We use the PDG values and they are small: less than 0.5% for $\gamma\gamma$ channel and about 1.5% for the $\pi^+\pi^-\pi^0$ channel.

We summarize the systematic uncertainties for the two channels in Table 5.6.

5.7 Results

As of the writing of this thesis, the search is currently blinded. We show the expected sensitivity from the search and the blinded results in this section.

5.7.1 Expected Sensitivity

Figure 5-21 shows our expected sensitivity compared to existing limits on the ALP-gluon coupling ($c_\gamma = 0$, $c_g = 1$) including the systematic uncertainties.

5.7.2 Blinded Results

We adopt a blinding procedure as follows:

- mass regions of interest are masked and residuals are used to investigate the fit quality as shown in Sec. 5.3;
- any local significance greater than $+2\sigma$ is reported as $+2\sigma$; we also adjust the expected signed local significance distribution to account for this until unblinding.

Figures 5-22 and 5-23 show the blinded search results.

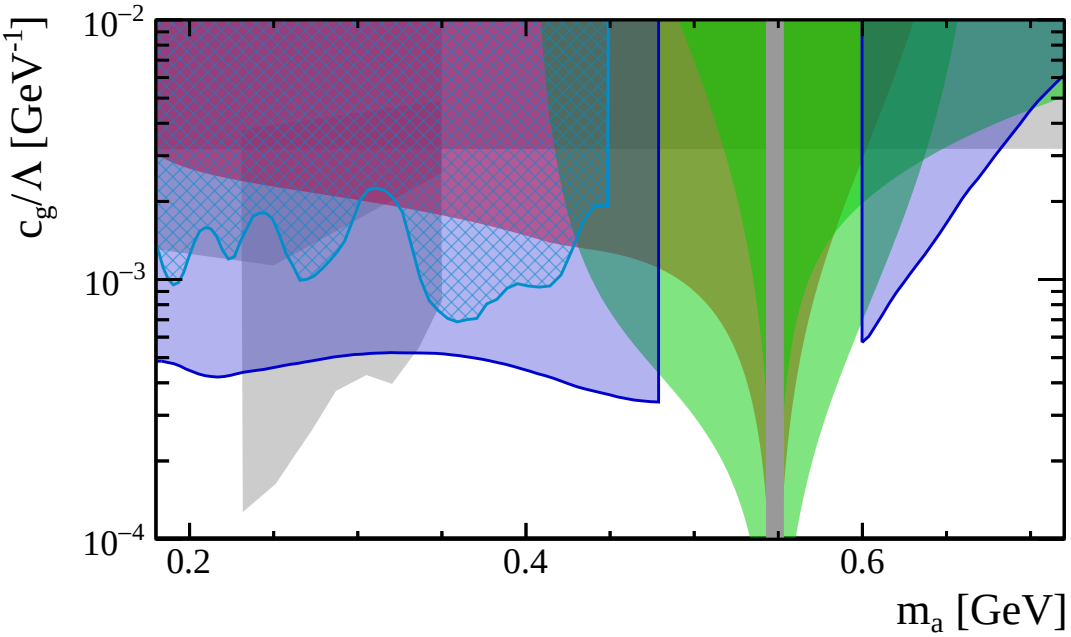

Figure 5-21: Expected sensitivity for the ALP-gluon coupling ($c_\gamma = 0$, $c_g = 1$) from this search (blue) compared to the current bounds [22] from (red) LEP [105, 106], (gray) kaon decays [107, 108, 109, 110], (hashed cyan) existing GlueX limits [26], and (green) ϕ and η' decays [1]. *N.b.*, the gray kaon limits depend on the unknown UV completion of the full theory; therefore, these have $O(1)$ uncertainty and should be taken *cum grano salis*.

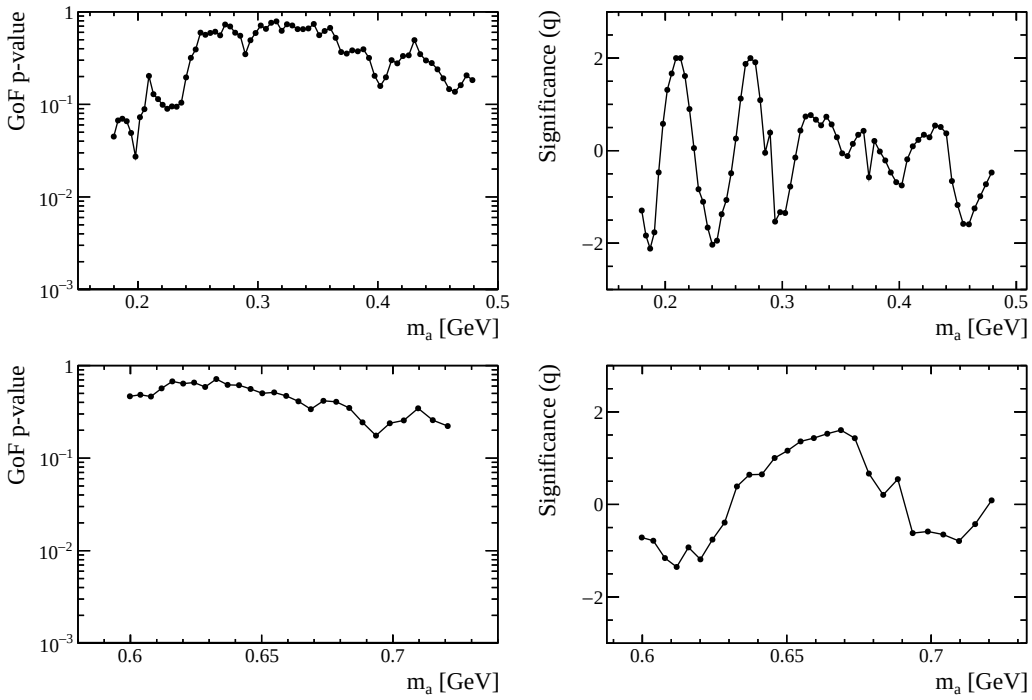

Figure 5-22: Goodness-of-fit (left column) and signed local significance (right column) as a function of test mass for $\gamma\gamma$ (top row) and $\pi^+\pi^-\pi^0$ (bottom row).

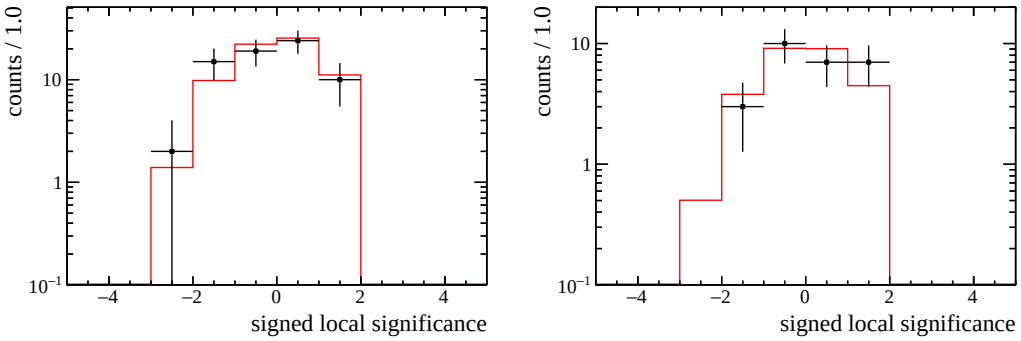

Figure 5-23: Signed local significances at each scan mass from (black) all fits and (red) the expected distribution for the (left) $\gamma\gamma$ channel and (right) $\pi^+\pi^-\pi^0$ channel; if the best-fit signal-yield estimator is negative, the signed significance is negative and *vice versa*. The error bars account for the correlation between nearest-neighbor fit results, which often produces outliers in pairs due to the $\sigma/2$ step size.

Chapter 6

Summary

The overarching theme of this thesis centers around problems in the study of the strong nuclear force. The GlueX experiment, of which I have been a member, aims to perform quantitative tests of Quantum Chromodynamics in the nonperturbative regime by searching for and studying hybrid mesons. I have presented efforts in the construction, commissioning, and reconstruction of the GlueX DIRC detector, that will upgrade the particle identification capability of the GlueX experiment in order to fully exploit its physics potential. The detector has been successfully constructed and commissioned. Preliminary analysis of the commissioning data showed promising results. Further reconstruction and calibration efforts are underway to better understand and improve the detector performance.

Originally proposed to solve the strong CP problem, axions and axion-like particles are hypothetical pseudoscalar particles found in many proposed extensions to the Standard Model of particle physics. I have presented a search for photoproduction of axion-like particles using data in photon-proton interactions collected by the GlueX experiment at Jefferson Laboratory in the $\gamma\gamma$ and $\pi^+\pi^-\pi^0$ final states of the ALPs. We expect to set world-leading limits across most of the ALP masses over which the search is conducted. The analysis is currently under internal review of the GlueX collaboration.

In addition, the Monte Carlo modeling of the strong interaction at low energies leads to challenges known as the event generator tuning problem. I have presented a novel approach to the Monte Carlo event generator tuning problem using Bayesian optimization.

Bibliography

[1] Particle Data Group, P. A. Zyla *et al.*, *Review of Particle Physics*, Progress of Theoretical and Experimental Physics **2020** (2020), 083C01.

[2] B. Ketzer, B. Grube, and D. Ryabchikov, *Light-meson spectroscopy with COMPASS*, Progress in Particle and Nuclear Physics **113**, 103755 (2020).

[3] Hadron Spectrum Collaboration, J. J. Dudek, R. G. Edwards, P. Guo, and C. E. Thomas, *Toward the excited isoscalar meson spectrum from lattice QCD*, Phys. Rev. D **88**, 094505 (2013).

[4] Joint Physics Analysis Center, A. Rodas *et al.*, *Determination of the Pole Position of the Lightest Hybrid Meson Candidate*, Phys. Rev. Lett. **122**, 042002 (2019).

[5] GlueX Collaboration, M. Dugger *et al.*, *A study of decays to strange final states with GlueX in Hall D using components of the BaBar DIRC*, Jefferson Lab PAC 42 Proposal (2014) arXiv:1408.0215 [physics.ins-det].

[6] R. Peccei, *The Strong CP Problem and Axions* (Springer, Berlin, Heidelberg, 2008), Lecture Notes in Physics, vol 741.

[7] A. Hook, *TASI Lectures on the Strong CP Problem and Axions*, arXiv:1812.02669 [hep-ph], 2018.

[8] C. A. Baker *et al.*, *Improved Experimental Limit on the Electric Dipole Moment of the Neutron*, Phys. Rev. Lett. **97**, 131801 (2006).

[9] J. M. Pendlebury *et al.*, *Revised experimental upper limit on the electric dipole moment of the neutron*, Phys. Rev. D **92**, 092003 (2015).

[10] A. Hook, S. Kumar, Z. Liu, and R. Sundrum, *High Quality QCD Axion and the LHC*, Phys. Rev. Lett. **124**, 221801 (2020).

[11] P. Di Vecchia and G. Veneziano, *Chiral dynamics in the large N limit*, Nuclear Physics B **171**, 253 (1980).

[12] R. D. Peccei and H. R. Quinn, CP *Conservation in the Presence of Pseudoparticles*, Phys. Rev. Lett. **38**, 1440 (1977).

[13] R. D. Peccei and H. R. Quinn, *Constraints imposed by* CP *conservation in the presence of pseudoparticles*, Phys. Rev. D **16**, 1791 (1977).

[14] F. Wilczek, *Problem of Strong P and T Invariance in the Presence of Instantons*, Phys. Rev. Lett. **40**, 279 (1978).

[15] S. Weinberg, *A New Light Boson?*, Phys. Rev. Lett. **40**, 223 (1978).

[16] R. Kallosh, A. Linde, D. Linde, and L. Susskind, *Gravity and global symmetries*, Phys. Rev. D **52**, 912 (1995).

[17] Y. Nomura and J. Thaler, *Dark matter through the axion portal*, Phys. Rev. D **79**, 075008 (2009).

[18] M. Freytsis and Z. Ligeti, *On dark matter models with uniquely spin-dependent detection possibilities*, Phys. Rev. **D83**, 115009 (2011).

[19] M. J. Dolan, F. Kahlhoefer, C. McCabe, and K. Schmidt-Hoberg, *A taste of dark matter: Flavour constraints on pseudoscalar mediators*, JHEP **03**, 171 (2015), [Erratum: JHEP07,103(2015)].

[20] Y. Hochberg, E. Kuflik, R. Mcgehee, H. Murayama, and K. Schutz, *SIMPs through the axion portal*, (2018).

[21] P. Svrcek and E. Witten, *Axions in string theory*, Journal of High Energy Physics **2006**, 051 (2006).

[22] D. Aloni, Y. Soreq, and M. Williams, *Coupling QCD-Scale Axionlike Particles to Gluons*, Phys. Rev. Lett. **123**, 031803 (2019).

[23] D. Aloni, C. Fanelli, Y. Soreq, and M. Williams, *Photoproduction of Axionlike Particles*, Phys. Rev. Lett. **123**, 071801 (2019).

[24] K. J. Kelly, S. Kumar, and Z. Liu, *Heavy Axion Opportunities at the DUNE Near Detector*, (2020).

[25] T. Fujiwara, T. Kugo, H. Terao, S. Uehara, and K. Yamawaki, *Non-Abelian Anomaly and Vector Mesons as Dynamical Gauge Bosons of Hidden Local Symmetries*, Progress of Theoretical Physics **73**, 926 (1985).

[26] GlueX Collaboration, H. Al Ghoul *et al.*, *Measurement of the beam asymmetry* Σ *for* π^0 *and* η *photoproduction on the proton at* $E_\gamma = 9$ *GeV*, Phys. Rev. C **95**, 042201 (2017).

[27] GlueX Collaboration, S. Adhikari *et al.*, *Beam asymmetry* Σ *for the photoproduction of* η *and* η' *mesons at* $E_\gamma = 8.8$ GeV, Phys. Rev. C **100**, 052201 (2019).

[28] S. Adhikari *et al.*, *The GlueX beamline and detector*, Nucl. Instrum. and Methods Sect. A **987**, 164807 (2021).

[29] C. W. Leemann, D. R. Douglas, and G. A. Krafft, *The Continuous Electron Beam Accelerator Facility: CEBAF at the Jefferson Laboratory*, Annual Review of Nuclear and Particle Science **51**, 413 (2001).

[30] C. Meyer, *The GlueX Experiment: (part 1) A Short History of GlueX*, GlueX-doc-4901.

[31] U. Timm, *Coherent Bremsstrahlung of Electrons in Crystals*, Fortschritte der Physik **17**, 765 (1969).

[32] K. Livingston, *The Stonehenge technique. A method for aligning coherent bremsstrahlung radiators*, Nucl. Instrum. and Methods Sect. A **603**, 205 (2009).

[33] M. Dugger *et al.*, *Design and construction of a high-energy photon polarimeter*, Nucl. Instrum. and Methods Sect. A **867**, 115 (2017).

[34] F. Barbosa *et al.*, *Pair spectrometer hodoscope for Hall D at Jefferson Lab*, Nucl. Instrum. and Methods Sect. A **795**, 376 (2015).

[35] N. Jarvis *et al.*, *The Central Drift Chamber for GlueX*, Nucl. Instrum. and Methods Sect. A **962**, 163727 (2020).

[36] L. Pentchev *et al.*, *Studies with cathode drift chambers for the GlueX experiment at Jefferson Lab*, Nucl. Instrum. and Methods Sect. A **845**, 281 (2017), Proceedings of the Vienna Conference on Instrumentation 2016.

[37] T. Beattie *et al.*, *Construction and performance of the barrel electromagnetic calorimeter for the GlueX experiment*, Nucl. Instrum. and Methods Sect. A **896**, 24 (2018).

[38] K. Moriya *et al.*, *A measurement of the energy and timing resolution of the GlueX Forward Calorimeter using an electron beam*, Nucl. Instrum. and Methods Sect. A **726**, 60 (2013).

[39] F. C. Pilat, *The 12 GeV Energy Upgrade at Jefferson Laboratory*, Proceedings of LINAC2012 (2012).

[40] E. Pooser *et al.*, *The GlueX Start Counter Detector*, Nucl. Instrum. and Methods Sect. A **927**, 330 (2019).

[41] B. Mecking *et al.*, *The CEBAF large acceptance spectrometer (CLAS)*, Nucl. Instrum. and Methods Sect. A **503**, 513 (2003).

[42] E. Smith, GlueX-doc-4915.

[43] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand, *Parton fragmentation and string dynamics*, Physics Reports **97**, 31 (1983).

[44] T. Sjöstrand *et al.*, *An Introduction to PYTHIA 8.2*, Comput. Phys. Commun. **191**, 159 (2015).

[45] Pythia home page: http://home.thep.lu.se/Pythia/.

[46] P. Skands, S. Carrazza, and J. Rojo, *Tuning PYTHIA 8.1: the Monash 2013 Tune*, Eur. Phys. J. **C74**, 3024 (2014).

[47] A. Buckley, H. Hoeth, H. Lacker, H. Schulz, and J. E. von Seggern, *Systematic event generator tuning for the LHC*, Eur. Phys. J. **C65**, 331 (2010).

[48] P. Ilten, M. Williams, and Y. Yang, *Event generator tuning using Bayesian optimization*, Journal of Instrumentation **12**, P04028 (2017).

[49] Spearmint GitHub repository: https://github.com/HIPS/Spearmint.

[50] J. Snoek, H. Larochelle, and R. P. Adams, *Practical Bayesian Optimization of Machine Learning Algorithms*, in *Advances in Neural Information Processing Systems 25*, edited by F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, pp. 2951–2959, Curran Associates, Inc., 2012.

[51] Y. Yang, TuneMC package, https://github.com/yunjie-yang/TuneMC.

[52] P. I. Frazier, *A Tutorial on Bayesian Optimization*, arXiv: 1807.02811.

[53] J. Mockus, *Bayesian Approach to Global Optimization: Theory and Applications* (Kluwer Academic Publishers, 1989).

[54] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas, *Taking the Human Out of the Loop: A Review of Bayesian Optimization*, Proceedings of the IEEE **104**, 148 (2016).

[55] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (MIT Press, 2006).

[56] P. A. Cherenkov, *Visible Emission of Clean Liquids by Action of γ Radiation*, Dokl. Akad. Nauk SSSR **2**, 451 (1934).

[57] I. Y. Tamm and I. M. Frank, *Coherent Radiation of Fast Electrons in a Medium*, Dokl. Akad. Nauk SSSR **14**, 107 (1937).

[58] The Nobel Prize in Physics 1958, `https://www.nobelprize.org/prizes/physics/1958/summary/`.

[59] C. Grupen and B. Shwartz, *Particle Detectors*, 2 ed. (Cambridge University Press, 2008).

[60] L. A. Jarvis *et al.*, *Cherenkov Video Imaging Allows for the First Visualization of Radiation Therapy in Real Time*, International Journal of Radiation Oncology*Biology*Physics **89**, 615 (2014).

[61] I. Adam *et al.*, *The DIRC particle identification system for the BaBar experiment*, Nucl. Instrum. and Methods Sect. A **538**, 281 (2005).

[62] GlueX Collaboration, *FDIRC Technical Design Report*, GlueX-doc-2809, 2015.

[63] B. Dey *et al.*, *Design and performance of the focusing DIRC detector*, Nucl. Instrum. and Methods Sect. A **775**, 112 (2015).

[64] J. Hardin, *Upgrading Particle Identification and Searching for Leptophobic Bosons at GlueX*, Ph.D. thesis, Massachusetts Institute of Technology, 2018.

[65] LORD MicroStrain® Sensing Systems, G-LINK-200, `https://www.microstrain.com/wireless/g-link-200`.

[66] Omega™, PX429-SGV, `https://www.omega.com/en-us/pressure-measurement/pressure-transducers/p/PX429-SGV`.

[67] LORD MicroStrain® Sensing Systems, V-LINK-200, `https://www.microstrain.com/wireless/v-link-200`.

[68] Omega™, KMQXL-062U-48, `https://www.omega.com/en-us/temperature-measurement/temperature-probes/probes-with-integral-connectors/p/JMQSS`.

[69] LORD MicroStrain® Sensing Systems, IEPE-LINK-LXRS, `https://www.microstrain.com/wireless/IEPE-LINK`.

[70] M. Shepherd, *Notes from the road: DIRC bar box shipping*, https://mailman.jlab.org/pipermail/halld-pid-upgrade/2018-June/000359.html.

[71] Hamamatsu H12700 Multi-Anode Photomultiplier Tube series, `https://www.hamamatsu.com/resources/pdf/etd/H12700_H14220_TPMH1379E.pdf`.

[72] A. Ali *et al.*, *Installation and Commissioning of the GlueX DIRC*, Journal of Instrumentation **15**, C09010 (2020).

[73] M. Patsyuk, *Optical Interface Using RTV-615 Silicone Technical Writeup*, GlueX-doc-3983, 2019.

[74] A. Ali *et al.*, *The GlueX DIRC Program*, Journal of Instrumentation **15**, C04054 (2020).

[75] U. Tamponi, *The TOP counter of Belle II: Status and first results*, Nucl. Instrum. and Methods Sect. A **952**, 162208 (2020), 10th International Workshop on Ring Imaging Cherenkov Detectors (RICH 2018).

[76] M. Patsyuk, R. Dzhydaglo, J. M. Hardin, and Y. Yang, *GlueX DIRC Writeup*, GlueX-doc-3974, 2019.

[77] M. Shepherd, *DIRC bar dimensions*, GlueX-doc-3733, 2018.

[78] Y. Yang, *DIRC As-Built Geometry in HDDS*, GlueX-doc-3876, 2019.

[79] R. Dzhygadlo *et al.*, *Simulation and reconstruction of the PANDA Barrel DIRC*, Nucl. Instrum. and Methods Sect. A **766**, 263 (2014), RICH2013 Proceedings of the Eighth International Workshop on Ring Imaging Cherenkov Detectors Shonan, Kanagawa, Japan, December 2-6, 2013.

[80] A. Ali *et al.*, *Particle identification algorithms for the PANDA Barrel DIRC*, Journal of Instrumentation **15**, C09057 (2020).

[81] J. Hardin and M. Williams, *FastDIRC: a fast Monte Carlo and reconstruction algorithm for DIRC detectors*, Journal of Instrumentation **11**, P10007 (2016).

[82] M. Starič, K. Inami, P. Križan, and T. Iijima, *Likelihood analysis of patterns in a time-of-propagation (TOP) counter*, Nucl. Instrum. and Methods Sect. A **595**, 252 (2008), RICH 2007.

[83] M. Starič, *Pattern recognition for the time-of-propagation counter*, Nucl. Instrum. and Methods Sect. A **639**, 252 (2011), Proceedings of the Seventh International Workshop on Ring Imaging Cherenkov Detectors.

[84] C. Fanelli and J. Pomponi, *DeepRICH: learning deeply Cherenkov detectors*, Machine Learning: Science and Technology **1**, 015010 (2020).

[85] J. Cohen-Tanugi *et al.*, *Optical properties of the DIRC fused silica Cherenkov radiator*, Nucl. Instrum. and Methods Sect. A **515**, 680 (2003).

[86] I. Thormählen, J. Straub, and U. Grigull, *Refractive Index of Water and Its Dependence on Wavelength, Temperature, and Density*, Journal of Physical and Chemical Reference Data **14**, 933 (1985).

[87] J. Hardin, FastDIRC package, `https://github.com/jmhardin/FastDIRC`.

[88] Y. Yang, *Geometry mapping: FastDIRC and hall coordinates*, GlueX-doc-4922.

[89] Y. Yang, FastDIRC4GlueX package, `https://github.com/yunjie-yang/FastDIRC4GlueX`.

[90] R. Dzhygadlo, *Status of Geometrical Reconstruction*, GlueX-doc-4435.

[91] J. Stevens, *Status of DIRC and Track Extrapolations*, GlueX-doc-4463.

[92] scikit-optimize package: https://github.com/scikit-optimize/scikit-optimize.

[93] R. Brun and F. Rademakers, *ROOT — An object oriented data analysis framework*, Nucl. Instrum. and Methods Sect. A **389**, 81 (1997), See also root.cern.ch.

[94] *GlueX Analysis Launch Cuts*, wikipage.

[95] P. Mattione, *Least Squares Kinematic Fitting of Physics Reactions*, GlueX-doc-2112 (2016).

[96] Y. Yang and M. Williams, *Analysis note: search for photoproduction of axion-like particles*, GlueX-doc-4866.

[97] M. N. Achasov *et al.*, *Study of the process $e^+e^- \to \pi^+\pi^-\pi^0$ in the energy region $\sqrt{s}$ below 0.98 GeV*, Phys. Rev. D **68**, 052006 (2003).

[98] GENR8 Monte Carlo Event Generator, GlueX-doc-11 (1998).

[99] M. Williams, *A novel approach to the bias-variance problem in bump hunting*, Journal of Instrumentation **12**, P09034 (2017).

[100] LHCb Collaboration, R. Aaij *et al.*, *Search for Dark Photons Produced in 13 TeV pp Collisions*, Phys. Rev. Lett. **120**, 061801 (2018).

[101] LHCb Collaboration, R. Aaij *et al.*, *Search for $A' \to \mu^+\mu^-$ Decays*, Phys. Rev. Lett. **124**, 041801 (2020).

[102] LHCb Collaboration, R. Aaij *et al.*, *Searches for low-mass dimuon resonances*, JHEP **10**, 156 (2020).

[103] W. A. Rolke, A. M. Lopez, and J. Conrad, *Limits and confidence intervals in the presence of nuisance parameters*, Nucl. Instrum. Meth. **A551**, 493 (2005).

[104] P. D. Dauncey, M. Kenzie, N. Wardle, and G. J. Davies, *Handling uncertainties in background shapes*, JINST **10**, P04015 (2015).

[105] S. Knapen, T. Lin, H. K. Lou, and T. Melia, *Searching for Axionlike Particles with Ultraperipheral Heavy-Ion Collisions*, Phys. Rev. Lett. **118**, 171801 (2017).

[106] G. Abbiendi et al. and The OPAL Collaboration, *Multi-photon production in ee collisions at $\sqrt{s} = 181-209 GeV*, **26**, 331.

[107] E. Izaguirre, T. Lin, and B. Shuve, *Searching for Axionlike Particles in Flavor-Changing Neutral Current Processes*, Phys. Rev. Lett. **118**, 111802 (2017).

[108] C. Lazzeroni *et al.*, *Study of the $K^\pm \to \pi^\pm\gamma\gamma$ decay by the NA62 experiment*, **732**, 65 .

[109] E. Abouzaid *et al.*, *Final results from the KTeV experiment on the decay $K_L \to \pi^0\gamma\gamma$*, Phys. Rev. D **77**, 112004 (2008).

[110] G. P. S. Gori, K. T. noted that an additional bound can be derived from $K^+ \to \pi^+\pi^0$ followed by a-π^0 mixing. We have derived this bound taking $\langle \boldsymbol{a\pi^0} \rangle$ from Ref. [22], and using kaon-decay data as in Refs. [107, 108].