

Spring 5-2023

Design and Construction of a Longitudinally Polarized Solid Nuclear Target for CLAS12

Victoria Lagerquist

Old Dominion University, vlage001@odu.edu

Follow this and additional works at: https://digitalcommons.odu.edu/physics_etds



Part of the Nuclear Commons

Recommended Citation

Lagerquist, Victoria. "Design and Construction of a Longitudinally Polarized Solid Nuclear Target for CLAS12" (2023). Doctor of Philosophy (PhD), Dissertation, Physics, Old Dominion University, DOI:

10.25777/36yz-ft35

https://digitalcommons.odu.edu/physics_etds/180

This Dissertation is brought to you for free and open access by the Physics at ODU Digital Commons. It has been accepted for inclusion in Physics Theses & Dissertations by an authorized administrator of ODU Digital Commons. For more information, please contact digitalcommons@odu.edu.

**DESIGN AND CONSTRUCTION OF A LONGITUDINALLY POLARIZED
SOLID NUCLEAR TARGET FOR CLAS12**

by

Victoria Lagerquist
M.S. 2017, Old Dominion University
B.S. 2015, Old Dominion University

A Dissertation Submitted to the Faculty of
Old Dominion University in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

PHYSICS

OLD DOMINION UNIVERSITY
May 2023

Approved by:

Sebastian Kuhn (Director)

Stephen Bültmann (Member)

Balša Terzić (Member)

Ted Rogers (Member)

Linda Vahala (Member)

ABSTRACT

DESIGN AND CONSTRUCTION OF A LONGITUDINALLY POLARIZED SOLID NUCLEAR TARGET FOR CLAS12

Victoria Lagerquist
Old Dominion University, 2023
Director: Dr. Sebastian Kuhn

A new polarized nuclear target has been developed, constructed, and deployed at Jefferson Laboratory in Newport News, VA for use with the upgraded 12 GeV CEBAF (Continuous Electron Beam Accelerator Facility) accelerator and the Hall B CLAS12 (12 GeV CEBAF Large Acceptance Spectrometer) detector array. This ‘APOLLO’ (Ammonia PO-Larized LOngitudinally) target is a longitudinally polarized, solid ammonia, nuclear target which employs DNP (Dynamic Nuclear Polarization) to induce a net polarization in samples of protons (NH_3) and deuterons (ND_3) cooled to 1 K via helium evaporation, held in a 5 T polarizing field supplied by the CLAS12 spectrometer, and irradiated with 140 GHz microwave radiation. It was utilized in the RGC (Run Group C) experiment suite through a collaboration of the JLab Target Group, Old Dominion University, Christopher Newport University, the University of Virginia, and the CLAS Collaboration. RGC comprised six experiments which measured multiple spin-dependent observables across a wide kinematic phase space for use in nucleon spin studies. The dimensional constraints necessary for the incorporation of APOLLO into CLAS12, as well as the considerations necessary to utilize the CLAS12 solenoid, introduced unique challenges to the target design. This document presents the innovative solutions developed for these challenges including a novel material transport system, superconducting magnetic correction coils, and an all new bespoke NMR (Nuclear Magnetic Resonance) system. In addition to a detailed description of the complete target system and an initial report of the RGC experimental run, it will also present a study of Quark-Hadron Duality in the g_1 spin structure function based on Hall B EG1b data and pQCD fits from the JAM (Jefferson Lab Angular Momentum) Collaboration.

Copyright, 2023, by Victoria Lagerquist, All Rights Reserved.

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my advisor, Dr. Sebastian Kuhn. Words cannot express my sincere gratitude for the guidance, encouragement, and patience he has shown to me during my time as a graduate student. It is no exaggeration to say that I would not have completed this PhD had it not been for his relentless support. It has truly been a privilege to have him as a mentor. It seems unfair to summarize the magnitude of his kindness in such few sentences but I'm afraid words fail me.

I would also like to extend my gratitude to the Jefferson Lab Target Group. The chapters of this thesis are insufficient to convey the incredible work they have done in crafting the APOLLO target and I am grateful to have been allowed to contribute to it. Particularly, I would like to thank James Maxwell for his affable tutoring on all things academic and otherwise, James Brock for his candid life advice and invaluable technical instruction, and finally Chris Keith for his wealth of experience and expertise which is rivaled only by his generosity and willingness to share it with confused graduate students.

Of course, I also thank my committee members: Dr. Balša Terzić, Dr. Stephen Bültmann, Dr. Ted Rogers, and Dr. Linda Vahala for their valuable time and insight during the development of this thesis. I also wish to acknowledge the exceptional faculty and staff of the ODU Physics Department. There are too many to name but I must specifically thank Dr. Gail Dodge and Dr. Lepsha Vuskovic - not just for myself but for all of the women in the department who have been the beneficiaries of their continuous and tenacious support. Likewise, thank you to Lisa Okun for her tireless commitment to every student that has been fortunate enough to have crossed her path.

I would also like to thank my academic co-conspirator Pushpa Pandey for keeping me company on this project; I have every confidence her thesis will outshine this one by orders of magnitude. I as well thank my fellow first year classmates: Gabriel, Sajini, Dilini, Nilanga, Jiwan, Mark, and Wayne - without them I would not so much as made it past the starting line.

Finally, thank you to Dr. Palacios Serrano, Dr. Makita, and Dr. Sosa... We made it.

TABLE OF CONTENTS

	Page
LIST OF TABLES	vi
LIST OF FIGURES.....	vii
Chapter	
1. INTRODUCTION.....	1
2. PHYSICS BACKGROUND.....	5
2.1. ELECTRON SCATTERING	5
2.2. CROSS-SECTIONS	8
2.3. PARTON MODELS	14
2.4. DUALITY	18
2.5. EXISTING DATA.....	26
3. SOLID CRYOGENIC POLARIZED TARGETS.....	29
3.1. DYNAMIC NUCLEAR POLARIZATION	29
3.2. TARGET MATERIAL	32
3.3. MAGNETICS	35
3.4. MICROWAVES.....	37
3.5. NMR	38
4. THE APOLLO TARGET.....	41
4.1. JEFFERSON LAB FACILITIES	41
4.2. CRYOSTAT	58
4.3. INSERT ASSEMBLY, TROLLEY, AND BATH	68
4.4. OTHER SYSTEMS.....	82
4.5. INCORPORATION INTO CLAS12.....	94
5. PERFORMANCE AND CONCLUSION.....	104
5.1. PERFORMANCE DURING DEVELOPMENT	104
5.2. PERFORMANCE DURING EXPERIMENT	110
5.3. RGC OUTCOME	115
5.4. CONCLUSION.....	118
BIBLIOGRAPHY	119
APPENDIX	125
VITA	128

LIST OF TABLES

Table	Page
1. Selected W ranges, in GeV used for test of Quark-Hadron Duality.	22
2. Experimental Q^2 Ranges, in GeV^2	23
3. List of experiments within Run Group C.....	28
4. Table of CLAS12 Components	45
5. Specifications of the coils for two models of the solenoid.	50
6. List of target cells used for Run Group C.....	92
7. Timeline of RGC showing major events and their impact on run time.....	116

LIST OF FIGURES

Figure	Page
1. Feynman diagram of electron - nucleon single photon exchange.....	6
2. Illustration of resonant behavior.....	16
3. Example of Quark-Hadron Duality.....	18
4. Approach towards the scaling limit for the structure function $g_1(x, Q^2)$	20
5. Diagram showing the algorithm used to generate tests of quark-hadron duality.....	24
6. Results of the integrated duality study.....	25
7. Energy levels and transitions for electron-proton pairs.....	30
8. Example of diminishing polarization with dose.....	34
9. Polarization vs. magnetic field strength.....	36
10. Optimal microwave frequency varying with beam dose.....	37
11. Response of an NMR Q-curve to polarized material.....	40
12. Illustration of the Jefferson Laboratory CEBAF accelerator.....	41
13. Illustration of the Hall B CLAS12 spectrometer.....	43
14. Simplified OPERA3D rendering of the CLAS12 solenoid coils.....	47
15. Results of the CLAS12 mapping.....	48
16. Comparison of measured and modeled fields of the CLAS12 solenoid.....	49
17. An example of a four shim coil configuration model.....	52
18. Scans of initial of shim coil geometry.....	54
19. Shim coils creating two adjacent regions of opposite polarization direction.....	55
20. Field produced by the final shim coil arrangement.....	56
21. The APOLLO cryostat during construction.....	58
22. Diagram of the helium paths of the APOLLO cryogenic system.....	59

Figure	Page
23. Plot of vapor pressure vs. temperature for Helium-4	62
24. A simplified diagram of the APOLLO heat shielding.....	63
25. Successive layers of the APOLLO target's insulation system.....	65
26. Illustration of the capacitive level probe.....	67
27. A simplified diagram of the APOLLO insert assembly.....	68
28. The APOLLO insert assembly and cryostat.....	69
29. The APOLLO movable bath trolley.....	71
30. Movement mechanism of the insert assembly.....	72
31. Wiring of the movement spool and retraction tether.	73
32. Examples of failed tether configurations.....	74
33. The APOLLO trolley base.....	76
34. The downstream end of the APOLLO trolley.....	78
35. The APOLLO locating cage.....	79
36. Locator blocks and microwave exit horn.....	80
37. Final version of shim coils as they were wound on the shim mandrel.....	82
38. High temperature superconducting leads.....	83
39. The shim coil tab of the NMR software.	84
40. The main page of the of the new NMR software.....	85
41. Sample of NMR coil analyser.	87
42. Effect of rastering radius on the accuracy of NMR measurements.....	89
43. Effect of rastering offset on the accuracy of NMR measurements.....	90
44. Ammonia target cell being loaded into a transport container.....	91
45. Dummy target cell in bath as seen from the access hatch.....	93

Figure	Page
46. Target cell being lowered through the loading hatch.....	94
47. The complete Hall B beamline diagram for RGC.....	95
48. The upper half of the insertion cart.....	96
49. Cryostat when fully deployed into the spectrometer.	97
50. Coarse physical measurements being taken of the raster magnets.	99
51. Approximate model of a single raster magnet and a plot of its field values.	99
52. Plot of the modeled field of the raster magnet.....	100
53. Engineering diagram of the ELMO Møller shield.	102
54. Plot of cryostat helium usage vs. added heat load.	105
55. Effectiveness of the liquid He level probe PID system.....	106
56. Temperature changes experienced by the bath during tests.	107
57. Results from initial tests of the shim coils.	108
58. NMR measurement of TEMPO-doped deuterated-butanol.	109
59. Temperature change during a target swap in Hall B.	111
60. The raster calibration target.	113
61. An off-center beam spot compared to the raster calibration measurement.	114
62. Optical bullseye target.	115
63. Events gathered during the RGC experimental run.	117
64. On-Axis Field Map Full Length.....	125
65. On-Axis Field Map Side Regions	126
66. 1.25 cm Off-Axis Field Map Full Length	126
67. 30 cm Off-Axis Field Map Central Region	127
68. Off-Axis Field Map Radial Component	127

CHAPTER 1

INTRODUCTION

Since the discovery of the atom in the early 1800's, physicists have sought to understand the smallest components of matter. The early 1900's brought Rutherford's [1] discovery of the nucleus and by 1964 Gell-Mann [2] and Zweig [3] were proposing the fundamental constituents of nucleons: quarks. But with every advancement come new questions and the exact dynamic of quarks within the nucleon remains an active area of research more than fifty years later. When investigating the smallest units of our universe, the concept of observation becomes a slightly complicated matter. Any particles we might use to see an experimental interaction become themselves a part of it. The responsibility of extracting meaning from these complex events belongs to the field of collision physics. In collision experiments, a beam of particles is accelerated to high speeds before being directed into either another accelerated beam of particles or (as covered in this thesis) a stationary collection of them called a target. The particles then interact with each other and the aftermath is detected by sensors. The precise state of these ricocheted pieces reveals information about the nature of the particles before the collision. Through careful control of the initial parameters and precise measurement of the final ones, new insights on the nature of the target particles can be gleaned.

To understand the process of extracting those insights, we need to first understand what is actually happening during a collision. On the quantum scale, particles transfer momentum and energy through the exchange of virtual particles. The nature of the virtual particle depends on the force causing the exchange. For quantum interactions within an electromagnetic field (the field generated by charged particles - notably electrons), virtual photons act as the exchange mediators. These virtual photons have much the same properties as real photons but exist only within the confines of the collision event - being emitted by one particle and absorbed by another. For fundamental particles, such as the electron, this alteration of momentum and energy can be fully determined by measuring its final trajectory and speed. A composite target particle, however, can be altered in more complicated ways, sometimes undergoing complex changes of its internal structure. The nature of this process is decided by the amount of energy and momentum given to it by the virtual photon. We often consider the virtual photon as a sort of lens for viewing the target.

Take, for example, the proton. At the most zoomed out level it interacts with the photon as a single particle with measurable values for properties like mass and charge. However, a closer inspection reveals that these properties are in fact the combined effect of a collection of particles that make up the proton (quarks and gluons), each with their own observable values. Moreover, the degree to which one zooms in to examine this collection determines the observed complexity of the collection itself.

The mathematical framework used to explore these interactions is the scattering cross-section. A cross-section describes the probability that an incoming particle (electrons in the case of this thesis) will scatter in a particular direction with a particular energy. Since those observables are dependent on the exchanged photon, the function depends on the energy and momentum of the virtual photon and the structure of the target. The simplest target is a single, point-like particle whose cross-section is determined only by its charge and produces a Mott cross-section. However, if the target has internal structure like the proton, this cross-section gains additional terms (called structure functions) to account for the structure. Some of these additional terms rely only on the momenta of the internal components (unpolarized structure functions) but there are also terms that depend on a property called spin (polarized structure functions).

First measured in 1922 by Stern and Gerlach [4], spin is sometimes referred to as the intrinsic angular momentum of a particle. Indeed, as the name suggests, spin was originally assumed to arise from the rotational motion of a particle about its own axis. This explanation seems to be confirmed by certain properties associated with spin - such as its relationship to magnetic fields. It is understood that a charged object would produce electric currents if rotated - resulting in a magnetic moment pointing in a direction relative to the angular momentum pseudovector. However, particles such as electrons are not objects - they have no spacial extent and therefore *cannot* rotate. Even if they did, the rate of rotation that would be required to generate their observed magnet moment would exceed the speed of light. However, we still observe these particles to have rotational angular momenta and, furthermore, all particles of a given type always have the same quantized values of rotational momentum. We now understand spin as simply an intrinsic particle property the same way we do for mass and charge. As such, it plays a crucial roll in the mechanisms of particle collisions and, in turn, cross-sections.

Cross-sections (and by extension structure functions) are inherently statistical in nature. Since it is impossible to manipulate a single electron into emitting a single virtual photon to interact with an single specified target component, the technique is instead to take data from

millions of interactions and look at the cumulative results. This means that, to understand the role played by spin within a target, measurements must be comparative in nature so that extraneous influences are canceled out. First, cross-sections are measured with the target spin aligned in one direction, then the same measurements are made with the spins reversed in alignment. If all other parameters are held the same, the difference between the two cross-section measurements must depend on spin exclusively (see Chapter 2).

Which brings us to the question of how to convince a target particle (or, more aptly, a large collection of identical target particles) to be spin-aligned in a chosen direction. As mentioned, spin has an inherent relationship to magnetic fields. Their alignment can be controlled by the introduction of an external magnetic field. It is not sufficient, however, to simply apply that field. Though it may allow the alignment of a solitary particle, there are counter processes happening within collections of particles which limits the average alignment of the overall group (called polarization). To overcome this, a precisely controlled environment must be crafted which not only includes an exceptionally strong (5 T) and uniform (10^{-4}) magnetic field but also brings the target particles down to extremely low temperatures (1 K) and introduces high frequency (140 GHz) microwave radiation. Once met, these criteria allow the the process of dynamic nuclear polarization, or DNP (see Chapter 3).

DNP is a powerful and well-tested technique in the arsenal of nuclear physics. Proposed in 1953 by Overhauser [5], and first demonstrated in the 1950's [6] but it was SLAC in the 1970's that spearheaded its use for scattering experiments. In 1988, surprising results from European Muon Collaboration (CERN) brought even greater attention to the field of spin physics and triggered a cascade of similar experiments around the world including DESY/HERMES and Jefferson Lab. Each experiment provided new insights to the nature of spin within the nucleon, but there are still many unanswered questions (see Section 2.5).

One such question is how spin is distributed within the complex structure of a nucleon. Other intrinsic values, such as charge, are easily explained by simply summing together the individual values of the composite particles (quarks) making up the nucleon. And, at first glance, this would seem to hold true for spin as well. However, it was found that the combined behavior of these quarks contributed only a small fraction of the expected total spin, opening a mystery known as the spin crisis (or spin puzzle). Another, less concrete riddle is how to bridge the gap in how we formulate the description of partonic quarks (when they act individually) and how they behave as a combined hadronic group. In other words, how can we predict the overall nature of a proton from the first principles of its constituents

- a puzzle known as quark-hadron duality (see Section [2.4](#)).

It is in service to these pursuits and more that the topic of this thesis is centered. A brand new polarized nuclear target has been designed and constructed for use with the newly upgraded 12 GeV accelerator facilities at Jefferson Lab. The body of this text details its motivation and implementation, ending with current status and future expectations (see Chapters [4](#) and [5](#)).

CHAPTER 2

PHYSICS BACKGROUND

2.1 ELECTRON SCATTERING

The use of electrons as probes in scattering experiments is incredibly useful because electron dynamics and interactions are very well understood and the particles themselves are easily generated and controlled. They interact via the electroweak force which is well described via QED and, though the hadronic target they probe may undergo complicated changes, the final state of the deflected electron (being a fundamental particle) can be fully and easily measured and characterized.

In scattering experiments, only the initial and final states of particles are measurable. Any processes that occur during the interaction must be determined by comparison of the two. Moreover, some final states are never measured at all. Scattering processes in which only the deflected electron is characterized are called inclusive (as inferences about the interaction must encompass all possible final states for the unmeasured target particle). Events in which the final state of the target is measured explicitly are called exclusive and ones in which the target is broken apart and only some of its pieces are detected are called semi-inclusive.

Of course, quantum scattering does not occur through direct contact. Instead, particles can be described as interacting through the exchange of mediating force carriers (e.g., photons for the electroweak force and gluons for the strong force). In some cases, there are a complex series of interactions with multiple exchanges and self-interactions within the internal components of the target. These are sometimes called higher order processes. In the simplest case, however, one particle emits a single virtual¹ photon which is absorbed by the other in a process called single-photon exchange.

Variables

To analyze this process quantitatively, it is necessary to introduce some variables. The kinematic quantity of interest for any particle is its four-momentum. Four-momenta are

¹Virtual particles have the same properties as their real counterparts, except their energy can deviate from the relativistic requirement $E = \sqrt{m^2 + p^2}$. This deviation can occur only for a short time, hence virtual particles can travel only over short distances

spacetime vector quantities which combine a particle's standard three-dimensional momentum vector with its energy. They are extremely useful in that their squared values are Lorentz invariant, i.e., the same in all relativistic reference frames.

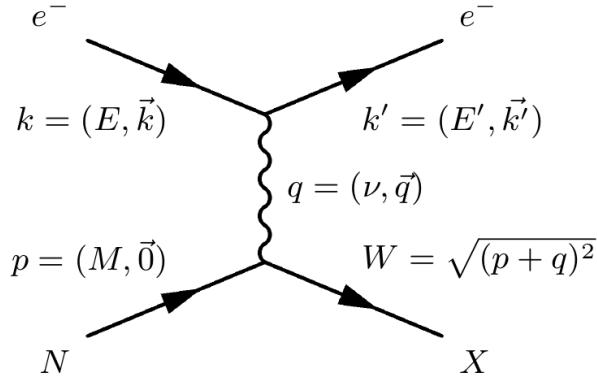


FIG. 1: Feynman diagram of electron - nucleon single photon exchange ($eN \rightarrow e'X$) where the electron has incoming (outgoing) four-momentum k (k'), the nucleon has initial four-momentum p and final mass W , and the exchanged virtual photon has four-momentum q . X denotes the target's unspecified final state particle (or collection of particles). Scanning the image from left to right gives a general timeline of events. On the right, an electron (top straight line) and nucleon (bottom straight line) exist independently. In the middle, they exchange a virtual photon (wavy line) which transfers energy and momentum between them. On the left, the particles again exist independent of each other in their final states.

In Fig. 1, we see an incoming electron with a four-momentum k (containing energy E and classical three-momentum \vec{k}) and a stationary nucleon with four-momentum p . As the nucleon is at rest² (with momentum $\vec{0}$), its energy is equal to its rest mass M .³ Connecting them, a virtual photon with a four-momentum of q , transfers energy ν and momentum \vec{q} , leaving final state four-momenta of k' and p' respectively. This transferred energy can be

²Another common reference frame is the Breit (or the infinite-momentum) frame in which $\vec{p} = (-)\vec{p}'$.

³Throughout, this text will be using natural units where $c=1$.

identified explicitly by the change in energy of the electron:

$$\nu = \frac{\mathbf{k} \cdot \mathbf{q}}{M} = E - E', \quad (1)$$

from which we can also express the fraction of the energy lost by the electron:

$$y = \frac{\mathbf{q} \cdot \mathbf{p}}{\mathbf{k} \cdot \mathbf{p}} = \frac{\nu}{E}. \quad (2)$$

Kinematic Domains

More salient though, is the transferred four-momentum or, more specifically, the negative of its square:

$$Q^2 \equiv -\mathbf{q}^2 = -(\nu^2 - \mathbf{q}^2), \quad (3)$$

which, as discussed, is constant across all inertial frames. It is this quantity that determines the resolution at which the electron is able to probe the hadron. For low Q^2 , the interaction unfolds as if the electron is interacting with the nucleon as a structureless particle. As the four-momentum transfer increases, the interaction transitions to that with a composite particle and the nucleon begins to display a structure of three bound quarks. Further increases reveal a so-called ‘sea’ of quarks⁴ in addition to the three valence ones.

Squaring the final state four-momentum of the target (\mathbf{p}') gives another useful quantity called the invariant final mass (squared) or sometimes simply the missing mass (W):

$$W^2 \equiv p'^2 = M^2 + 2M\nu - Q^2, \quad (4)$$

which is used to classify the end state of the target for instances in which it cannot be observed directly.

Scattering events in which the missing mass is measured to be that of the initial nucleon are identified as having undergone *elastic scattering* - the target remained unchanged in its ground state and the absorbed energy was simply transferred into kinetic energy. If that nucleon, however, was initially one of many within a compound nucleus, the absorbed photon can also provide sufficient energy to knock the target out of the energy well of that nucleus (while still leaving it unchanged) - in which case the missing mass will be that of a free nucleon plus the nuclear binding energy. These events are called *quasi-elastic scattering*.

Increasing the transferred energy further, we enter the domain in which the nucleon is fundamentally changed by the collision: *inelastic scattering*. Initially, these collisions cause

⁴In this sea, virtual gluons (the mediating particles for the fundamental strong force associated with quarks) split to form virtual quark-antiquark pairs which later recombine back into virtual gluons.

the nucleon to leave its ground state and enter an excited one - a short-lived, unstable particle known as a *nucleon resonance* (discussed later). With enough transferred energy, though, the target nucleon is altogether broken apart into an incoherent collection of many particles. In these events (known as *deep inelastic scattering*), the missing mass encompasses the entire group of final particles. In this region (*DIS*), the virtual photon is no longer described as interacting with the nucleon as a whole entity but, rather, the photon is described as interacting with a single (fundamental) constituent particle within the nucleon (a quark or gluon) sometimes called a parton.

2.2 CROSS-SECTIONS

After exchanging the virtual photon, the electron is deflected away from the nucleon by some angle which is measured (along with its energy) by detectors. The likelihood of finding the scattered electron in any given range of energy and deflection angle is given by the differential cross section:

$$\frac{d^2\sigma}{d\Omega dE'} = \frac{\alpha^2}{Q^4} \frac{E'}{E} L_{\mu\nu} W^{\mu\nu}, \quad (5)$$

where Ω is the solid angle of the scattered electron (a range in both the polar θ and azimuthal ϕ angles), $\alpha \approx \frac{1}{137}$ is the electromagnetic fine structure constant, and $L_{\mu\nu}$ and $W^{\mu\nu}$ are the leptonic (electron) and hadronic (nucleon) tensors. These tensors carry all of the information about the dynamics and internal structure of the particles.

As mentioned, the great benefit of using electrons as scattering probes is that the leptonic tensor is already well understood and fully calculable within QED using Dirac spinors $u(k, s)$ and the gamma matrix γ^μ (where s and s' are the initial and final spin states of the lepton). After summing over all final spin states, the leptonic tensor can be broken into two terms which are either symmetric (S) or anti-symmetric (A) with respect to the initial leptonic spin state:

$$\begin{aligned} L_{\mu\nu}(k, s; k', s') &= \sum_{s'} [\bar{u}(k', s') \gamma_\mu u(k, s)] * [\bar{u}(k', s') \gamma_\nu u(k, s)] \\ &= L_{\mu\nu}^{(S)}(k; k') + i L_{\mu\nu}^{(A)}(k, s; k') + L_{\mu\nu}'^{(S)}(k, s; k', s') + i L_{\mu\nu}'^{(A)}(k; k', s') \\ &= 2L_{\mu\nu}^{(S)}(k; k') + 2i L_{\mu\nu}^{(A)}(k, s; k'). \end{aligned} \quad (6)$$

Conversely, since the hadronic tensor describes a composite particle - subject to the strong force and QCD - it is not (as yet) exactly calculable. It can, however, still be parameterized

into symmetric and anti-symmetric terms similar to the leptonic tensor. Considering every transition from the ground state $|N\rangle$ to (all possible) excited states $|X\rangle$ (and utilizing the completeness of those states) while using the electromagnetic current operator $J_\mu(\zeta)$ (where ζ is the spacial four-vector) we find:

$$\begin{aligned} W_{\mu\nu}(q; p, S) &= \frac{1}{4\pi M} \int d^4\zeta e^{iq\cdot\zeta} \langle N_s(P) | J_\mu(\zeta) J_\nu(0) | N_s(P) \rangle \\ &= W_{\mu\nu}^{(S)}(q; p) + iW_{\mu\nu}^{(A)}(q; p, S), \end{aligned} \quad (7)$$

where S is the initial spin state of the nucleon.

Inserting these leptonic, Eq. (6), and hadronic, Eq. (7), tensor expressions into Eq. (5), we can separate the full cross-section formula into spin-dependent and spin-independent terms:

$$\frac{d^2\sigma}{d\Omega dE'} = \frac{\alpha^2}{2MQ^4} \frac{E'}{E} [L_{\mu\nu}^{(S)} W^{\mu\nu(S)} - L_{\mu\nu}^{(A)} W^{\mu\nu(A)}]. \quad (8)$$

2.2.1 UNPOLARIZED CROSS-SECTION

Adopting a convention for longitudinal (along the collision axis) lepton (\uparrow) and hadron ($\uparrow\downarrow$) spin orientations, we find that *summing* spin state combinations causes the anti-symmetric terms of Eq. (8) to vanish, producing an unpolarized cross-section:

$$\frac{d^2\sigma^{\uparrow\downarrow}}{d\Omega dE'} + \frac{d^2\sigma^{\uparrow\uparrow}}{d\Omega dE'} = \frac{d^2\sigma^{\text{unpol}}}{d\Omega dE'} = \frac{\alpha^2}{MQ^4} \frac{E'}{E} L_{\mu\nu}^{(S)} W^{\mu\nu(S)}. \quad (9)$$

Calculating $L_{\mu\nu}^{(S)}$ explicitly from Eq. (6), we find:

$$L_{\mu\nu}^{(S)}(k; k') = k_\mu k'_\nu + k'_\mu k_\nu - g_{\mu\nu}(k \cdot k' - m^2), \quad (10)$$

where $g_{\mu\nu}$ is the metric tensor and m is the lepton mass. While the leptonic contribution is completely described by its momenta and mass, the hadron requires additional factors to describe it. From Eq. (7), we find:

$$\begin{aligned} W_{\mu\nu}^{(S)}(q; p) &= (2M) \left(-g_{\mu\nu} + \frac{q_\mu q_\nu}{q^2} \right) W_1(p \cdot q, q^2) \\ &\quad + (2M) \left[\left(p_\mu - \frac{p \cdot q}{q^2} q_\mu \right) \left(p_\nu - \frac{p \cdot q}{q^2} q_\nu \right) \right] \frac{W_2(p \cdot q, q^2)}{M^2}, \end{aligned} \quad (11)$$

where W_1 and W_2 parameterize the internal structures and dynamics within the hadron. Finally, combining Eq. (9) - Eq. (11) gives:

$$\frac{d^2\sigma^{\text{unpol}}}{d\Omega dE'} = \sigma_{\text{Mott}} \left(2W_1(\nu, Q^2) \tan^2 \frac{\theta}{2} + W_2(\nu, Q^2) \right), \quad (12)$$

where:

$$\sigma_{\text{Mott}} = \frac{4\alpha^2 E'^2}{Q^4} \cos^2 \frac{\theta}{2}, \quad (13)$$

is the Mott cross-section (which is the cross-section found when scattering off a simple point-like particle).

These hadronic parameterizations, W_1 and W_2 , are dependent on both the transferred energy (ν) and transferred squared momentum (Q^2) from the virtual photon. As discussed, the resolution of the electron as a probe increases with Q^2 - moving from interactions with the overall hadron (in the elastic region) to interactions with individual partons (in the DIS region). In fact, once sufficiently within the DIS region, we find that W_1 and W_2 cease to depend on Q^2 - as the fundamental hadronic constituents have no internal structure and cannot be further resolved. The parameterizations instead depend only on the Bjorken scaling variable x :

$$x \equiv \frac{Q^2}{2M\nu} = \frac{Q^2}{W^2 - M^2 + Q^2}, \quad (14)$$

which represents the fraction of the total nucleon momentum carried by any particular parton struck in the collision. With this, we can define what are called the *unpolarized structure functions*:

$$\begin{aligned} F_1(x, Q^2) &\equiv MW_1(\nu, Q^2), \\ F_2(x, Q^2) &\equiv \nu W_2(\nu, Q^2), \end{aligned} \quad (15)$$

which, once sufficiently near the Bjorken scaling limit⁵ (where $Q^2 \rightarrow$ large and $\nu \rightarrow \infty$), become: $F_1(x, Q^2) \rightarrow F_1(x)$ and $F_2(x, Q^2) \rightarrow F_2(x)$, giving rise to the Callan-Gross relation:

$$F_2(x) = 2xF_1(x). \quad (16)$$

Being fully independent of spin, these unpolarized structure functions represent the momentum distribution within the hadron - with F_1 roughly representing the longitudinal contribution and F_2 depending on a mixture of longitudinal and transverse contributions. Substituting them into Eq. (9) gives us the completed unpolarized cross-section:

$$\frac{d^2\sigma^{\text{unpol}}}{d\Omega dE'} = \sigma_{\text{Mott}} \left(\frac{2}{M} F_1(x, Q^2) \tan^2 \frac{\theta}{2} + \frac{1}{\nu} F_2(x, Q^2) \right). \quad (17)$$

⁵This limit and its associated evolution behavior is dictated by the DGLAP equations[7–9].

2.2.2 POLARIZED CROSS-SECTION

Returning to Eq. (8) and the case that we do not sum over spin states (but instead take their *difference*), it is the symmetric contributions that cancel and we find:

$$\frac{d^2\sigma^{\uparrow\downarrow}}{d\Omega dE'} - \frac{d^2\sigma^{\uparrow\uparrow}}{d\Omega dE'} = \frac{d^2\sigma^{\text{pol}}}{d\Omega dE'} = \frac{2\alpha^2}{MQ^4} \frac{E'}{E} L_{\mu\nu}^{(A)} W^{\mu\nu(A)}. \quad (18)$$

In this case, from Eq. (6) we derive the asymmetric term of the leptonic tensor:

$$L_{\mu\nu}^A(k, s; k') = m\epsilon_{\mu\nu\alpha\beta} s^\alpha (k - k')^\beta, \quad (19)$$

and, from Eq. (7), that of the hadronic tensor:

$$\begin{aligned} W_{\mu\nu}^{(A)}(q; p, S) &= (2M)\epsilon_{\mu\nu\alpha\beta} q^\alpha \\ &\times \left[M S^\beta G_1(\nu, Q^2) + ((p \cdot q) S^\beta - (S \cdot q) p^\beta) \frac{G_2(\nu, Q^2)}{M} \right], \end{aligned} \quad (20)$$

where, just as in Eq. (11), G_1 and G_2 parameterize the internal structure of the hadron - in this case the spin dependent structure. Continuing as in Eq. (15), we can define the *polarized spin structure functions*:

$$\begin{aligned} g_1(x, Q^2) &\equiv M^2 \nu G_1(\nu, Q^2), \\ g_2(x, Q^2) &\equiv M \nu^2 G_2(\nu, Q^2), \end{aligned} \quad (21)$$

which produce the polarized cross-section:

$$\frac{d^2\sigma^{\text{pol}}}{d\Omega dE'} = \sigma_{\text{Mott}} \frac{1}{M\nu} 4 \tan^2 \frac{\theta}{2} \left((E + E' \cos \theta) g_1(x, Q^2) - 2Mx g_2(x, Q^2) \right). \quad (22)$$

Lastly, as this derivation only addresses the longitudinal spin alignments, we can also define the cross-section for transversely polarized collisions (where the target spin is aligned perpendicularly to that of the beam):

$$\frac{d^2\sigma^{\uparrow\Rightarrow}}{d\Omega dE'} - \frac{d^2\sigma^{\uparrow\Leftarrow}}{d\Omega dE'} = \frac{d^2\sigma^{\text{trans}}}{d\Omega dE'} = \sigma_{\text{Mott}} \frac{1}{M\nu} 4E' \tan^2 \frac{\theta}{2} \sin \theta \left(g_1(x, Q^2) + \frac{2E}{\nu} g_2(x, Q^2) \right). \quad (23)$$

Very roughly speaking, g_1 can be associated with the longitudinal contribution to the spin structure of the nucleon while g_2 correlates to a combination of longitudinal and transverse contributions - similarly to the unpolarized structure functions.

2.2.3 ASYMMETRIES

As stated, the four hadronic structure functions, F_1 , F_2 , g_1 , and g_2 , represent the complete internal dynamics of the nucleon that can be probed in inclusive lepton scattering via the electromagnetic interaction. However, the cross-sections presented thus far only ever produce them in coupled expressions. In principle, the longitudinal and transverse cross-sections could be combined to decouple g_1 and g_2 . In practice, however, the mechanical difficulty in constructing polarized (and particularly transversely polarized) experiments makes comparison to the unpolarized cross-section a more desirable route. To that end, we define the ratio of polarized Eq. (22), Eq. (23) to unpolarized Eq. (17) cross-sections,

$$\begin{aligned} A_{\parallel} &\equiv \frac{\sigma^{\uparrow\downarrow} - \sigma^{\uparrow\uparrow}}{\sigma^{\uparrow\downarrow} + \sigma^{\uparrow\uparrow}}, \\ A_{\perp} &\equiv \frac{\sigma^{\uparrow\Rightarrow} - \sigma^{\uparrow\Leftarrow}}{\sigma^{\uparrow\Rightarrow} + \sigma^{\uparrow\Leftarrow}}, \end{aligned} \quad (24)$$

as the longitudinal and transverse polarization asymmetries (where we have adopted the notation of $\sigma^{\uparrow\downarrow} \equiv \frac{d^2\sigma^{\uparrow\downarrow}}{d\Omega dE'}$). From this we find:

$$\begin{aligned} g_1(x, Q^2) &= \frac{F_1}{d'} \left[\tan\left(\frac{\theta}{2}\right) A_{\perp} + A_{\parallel} \right], \\ g_2(x, Q^2) &= \frac{F_1 y}{2d'} \left[\frac{E + E' \cos\theta}{E' \sin\theta} A_{\perp} - A_{\parallel} \right], \end{aligned} \quad (25)$$

where $d' = \frac{(1-\epsilon)(2-y)}{y(1+\epsilon R)}$ depends on:

$$R = \frac{F_2}{2xF_1} \left(1 + \frac{4M^2x^2}{Q^2} \right) - 1. \quad (26)$$

R is the ratio of longitudinal to transverse *photo absorption* cross-sections, and

$$\epsilon = \frac{1}{1 + 2 \tan^2 \frac{\theta}{2} \left(1 + \frac{y^2}{Q^2} \right)}, \quad (27)$$

is the ratio of longitudinal to transverse *virtual photon* polarizations. Thus far, we have only discussed spin orientations as they relate to the hadron and lepton in the overall scattering event. However, one can also independently consider the nucleon's interaction with the exchanged photon as a form of virtual Compton scattering.

Following a derivation similar to the one presented, virtual photon absorption cross-sections ($\sigma_{1/2}$ and $\sigma_{3/2}$) can be constructed in relation to the (anti-parallel and parallel) final state helicities (1/2 and 3/2) using the same hadronic structure functions as before.

The downside of such a formulation is that, as the photon exists only as an immeasurable mediating particle, these cross-sections are likewise immeasurable. They are, however, useful in constructing the *virtual photon absorption asymmetries*:

$$A_1 = \frac{\sigma_{1/2}^{(T)} - \sigma_{3/2}^{(T)}}{\sigma_{1/2}^{(T)} + \sigma_{3/2}^{(T)}} \quad \text{and} \quad A_2 = \frac{2\sigma_{1/2}^{(TL)}}{\sigma_{1/2}^{(T)} + \sigma_{3/2}^{(T)}}. \quad (28)$$

Here, A_1 correlates to transverse (T) photons while A_2 represents the interference between transverse and longitudinal (L). Relating these back to the polarization asymmetries (A_{\parallel} and A_{\perp}) helps to isolate and characterize individual structure functions. However, since the photon's propagation axis (along which its spin is aligned) is not guaranteed to match that of the electron, a photon depolarization factor,

$$D = \frac{1 - \epsilon \left(\frac{E'}{E} \right)}{1 + \epsilon R}, \quad (29)$$

must be incorporated into the final observable. With that, we can define A_1 and A_2 according to:

$$\begin{aligned} A_{\parallel} &= D(A_1 + \eta A_2), \\ A_{\perp} &= d(A_2 + \zeta A_1), \end{aligned} \quad (30)$$

where $\eta = \frac{\epsilon \sqrt{Q^2}}{(E - E' \epsilon)}$, $d = D \sqrt{\frac{2\epsilon}{(1+\epsilon)}}$, and $\zeta = \frac{\eta(1+\epsilon)}{2\epsilon}$. From there, we can express the spin structure functions as:

$$\begin{aligned} g_1 &= \frac{F_1}{1 + \gamma^2} (A_1 + \gamma A_2), \\ g_2 &= \frac{F_1}{1 + \gamma^2} \left(\frac{A_2}{\gamma} - A_1 \right), \end{aligned} \quad (31)$$

where $\gamma^2 = 4x^2 M^2 / Q^2$.

This derivation reveals some shape to the relative values between structure functions. Positivity constraints on the absorption cross-sections lead to the requirement that $|A_1| \leq 1$. Here, the case of $|A_1| = 1$ indicates a final state target with spin-1/2 - possible for elastic collisions with either the proton (low Q^2) or a constituent quark (high Q^2) polarized in the same direction as its host nucleon. These constraints also require that $|A_2| \leq \sqrt{R}$ (with the same elastic case of $|A_2| = \sqrt{R}$). Additionally, the Soffer limit [10] states: $|A_2| \leq \sqrt{R(1 + A_1)/2}$.

Finally, focusing on g_1 , we first find $A_1 = \frac{A_{\parallel}}{D} - \eta A_2$ from Eq. (30) which combines with Eq. (31) to yield:

$$g_1 = \frac{F_1}{1 + \gamma^2} \left(\frac{A_{\parallel}}{D} + (\gamma - \eta) A_2 \right). \quad (32)$$

Here, as Q^2 increases toward the Bjorken limit, η and γ both become small - allowing the second term to be ignored and approximations of g_1 to be made using only A_{\parallel} and F_1 .

2.3 PARTON MODELS

These structure functions are parameterizations of the hadronic tensor which is at present not exactly calculable. Although they characterize the internal dynamics of a compound particle, they *do* still mathematically treat the particle as a single object. The assumptions made regarding the nucleon's composite nature have, thus far, only been inferred through their behavior.

We can instead attempt an explicit description of the partonic composition of the nucleon. The first attempt at this was the Constituent Quark Model (CQM). This model proposed that the known baryons (p , n , Σ , etc.) were each composed of a different triplet of up, down, and strange quarks. Each of these quarks were assumed to contribute equally to the quantum numbers of the overall baryon - and the overall wavefunction was represented by an appropriate combination of its three constituents.

Perhaps surprisingly, the CQM produces an accurate description of the majority of observable properties.⁶ However, it breaks down upon direct measurement of (for example) the quark contribution to the overall proton spin. Indeed, far from constituting the entirety of the spin of the proton, the combined valence quarks were measured [11] to contribute less than a quarter of it (at most). This prompted the infamous Spin Crisis (or, less dramatically, Spin Puzzle) which motivated many of the polarized scattering experiments through the late 20th century.

Since then, the partonic understanding of hadrons has been refined into the Quark-Parton Model (QPM). This model formulates the nucleon as an unspecified collection of quarks each contributing to the overall particle with various strengths. For example, under this model the unpolarized structure function F_1 can be written as:

$$F_1(x) = \frac{1}{2} \sum_i e_i^2 [q_i(x) + \bar{q}_i(x)] = \frac{F_2(x)}{2x}, \quad (33)$$

where e_i is the charge of 'ith' quark and q_i (\bar{q}_i) are the probability densities for a quark (anti-quark) to carry nucleon momentum fraction x . From this, $F_1(x)$ can be interpreted as the likelihood of striking a parton with momentum fraction x . Also, as this model inherently inhabits the Bjorken limit (with its description of free quarks), F_2 obeys the Callan-Gross relation of Eq. (16).

⁶In fact, it is still used in our classification of particles though their valence quarks.

Likewise, the QPM description of the polarized structure functions is:

$$g_1(x) = \frac{1}{2} \sum_i e_i^2 [\Delta q_i(x) + \Delta \bar{q}_i(x)], \quad (34)$$

where $\Delta q(x) = q^+(x) - q^-(x)$ is the difference between the number density, $q^{+(-)}$, of quarks with momentum fraction x and spin parallel (anti-parallel) to the nucleon. Here we sensibly see g_1 as being the distribution of longitudinal parton polarization while g_2 has no QPM counterpart.

Beyond the QPM, we now understand quarks to be strongly interacting through the strong force. The strength of this interactions is determined by the coupling value $\alpha_s(\mu^2)$, where μ is some scale. At low μ^2 (long distances), α_s is large enough to cause color confinement - a phenomenon in which quarks are never seen outside of hadrons or unbound. At high μ^2 (short distances), however, α_s falls off such that asymptotic freedom occurs and the quarks appear nearly as free particles.

Asymptotic freedom explains the near correctness of the QPM and means that, as Q^2 increases, it approaches a domain in which the QPM holds exactly.⁷ This gives a basis from which to address the strong force interactions through perturbation theory. Perturbative Quantum Chromodynamics (pQCD) presents a framework by which these small, residual strong force effects are treated as a series of corrections to (a field theory formulation of) free unbound quarks. This estimation holds well in the domain of high Q^2 DIS ($W > 2$ GeV).

For lower Q^2 , though, the effects become too significant to be approximated as mere perturbations. Moving down from the DIS regime (by decreasing W) one eventually enters the kinematic resonance region that produces excited baryons (Δ , N^* , etc.) as final state targets (see Fig. 2). In this region, scattering events are dictated by hadronic degrees of freedom (not the partonic degrees of freedom used in pQCD) and must be studied phenomenologically. Though these two regimes are *theoretically* incompatible, there does exist medium values of Q^2 (between DIS and resonance) in which the method of operator product expansion (OPE) can be used to bridge them.

⁷You may notice the definition of the Bjorken limit is expressed as just that: a limit.

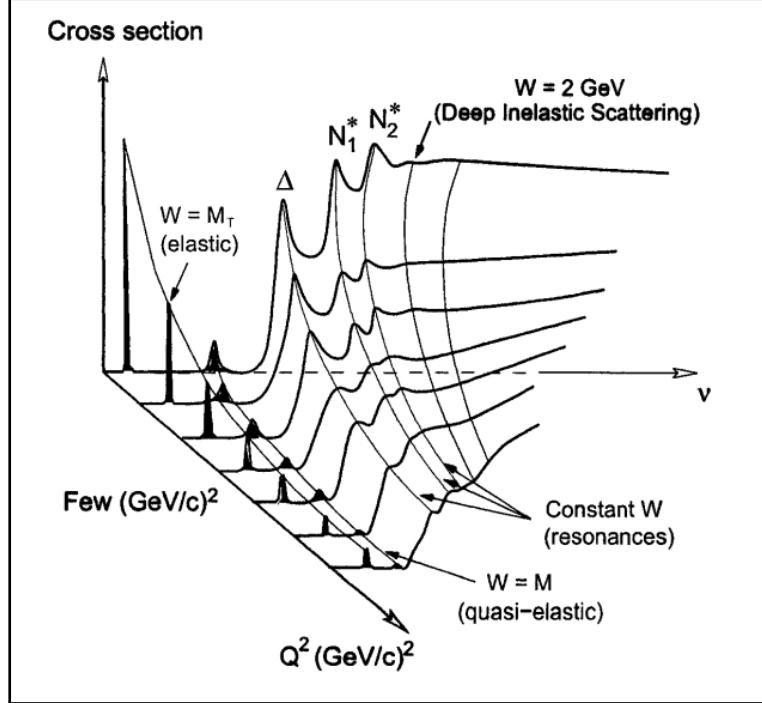


FIG. 2: Illustration of resonant behavior. Inclusive nuclear cross section plotted against Q^2 and ν . Reproduced from [12].

Operator Product Expansion

OPE was derived in 1968 by Kenneth Wilson[13]. We first consider two arbitrary local field operators: $\mathcal{O}_{a/b}$. The product of these operators, evaluated at two points whose spatial separation (δ) approaches zero, can be approximated as the expansion of a single operator (\mathcal{O}_k) with Wilson's coefficients (C_{abk}):

$$\lim_{\delta \rightarrow 0} \mathcal{O}_a(\delta) \mathcal{O}_b(0) = \sum_k C_{abk}(\delta) \mathcal{O}_k(0), \quad (35)$$

where these coefficients (but *not* the final operator) are exclusively dependent on the separation (δ). This relation requires that the overall evaluation occur at values large compared to δ .

Having already established that the coupling value α_s ‘runs’ (falls off at short distances) in the high Q^2 limit of DIS, we can manipulate this formula such that δ corresponds to α_s -

meaning calculations of the Wilson's coefficients (made within the domain of pQCD where α_s is definitely small) will remain valid - even outside the DIS region.

This first requires a Fourier transform to bring Eq. (35) into momentum space:

$$\lim_{q \rightarrow \infty} \int d^4 \delta e^{iq \cdot \delta} \mathcal{O}_a(\delta) \mathcal{O}_b(0) = \sum_k C_{abk}(q) \mathcal{O}_k(0), \quad (36)$$

which can then be applied to the electromagnetic current operators $J_{\mu/\nu}$ (introduced with the hadronic tensor in Eq. (7)) to expand them as functions of q .⁸

The resultant summed series contribution to $L_{\mu\nu} W^{\mu\nu}$ is found to consist of a terms on the order of:

$$\left(\frac{Q}{M}\right)^{2-\tau}, \quad (37)$$

where τ is known as twist. Note that the leading order contribution (the term with an overall power of 0) correlates to $\tau = 2$ (called twist-2 or leading twist) with higher twist contributions falling off in powers of M/Q .

From this we can express the spin structure functions within the relations:

$$\begin{aligned} \int_0^1 x^{n-1} g_1(x, Q^2) dx &= \frac{1}{2} a_{n-1} \\ \int_0^1 x^{n-2} g_2(x, Q^2) dx &= \frac{n-1}{2n} (d_{n-1} - a_{n-1}), \end{aligned} \quad (38)$$

where the left hand sides are the moments (x -weighted integrals) of $g_{1,2}$ [15] and a_{n-1} (d_{n-1}) are twist-2 (twist-3) operators calculated using the pQCD formulated Wilson's coefficients. In both cases, only the odd terms ($n = 1, 3, 5, \dots$) contribute to the asymmetric polarized structure functions.

In practice, we can extend the applicability of pQCD to lower μ^2 by adding terms of order $\frac{1}{Q^2}$. This method allows the theoretically explicit approximation of pQCD (with its partonic degrees of freedom) to extend down in Q^2 . Eventually, though, even this formulation fails and the effects (expressed via hadronic degrees of freedom) which dominate the resonance region cannot be reproduced via partonic descriptions. However, the boundary or, indeed, *overlap* between these two is explored in quark-hadron duality.

⁸For a more thorough handling of this, please refer to [14].

2.4 DUALITY

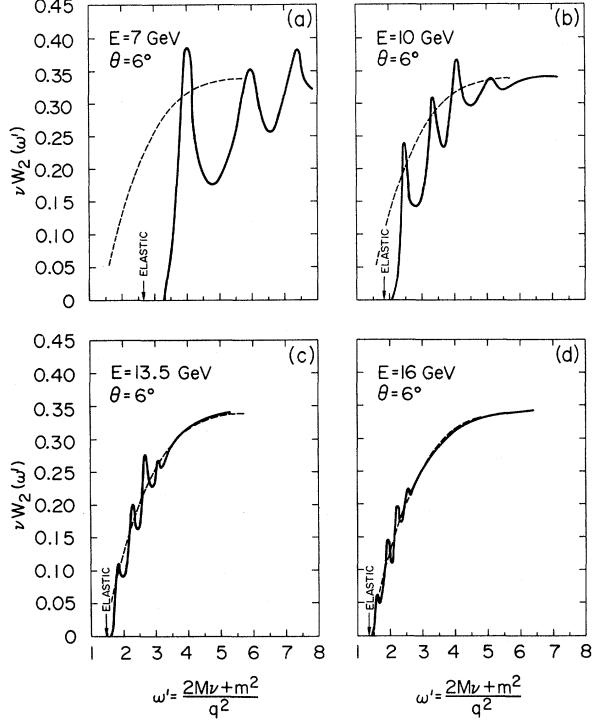


FIG. 3: Example of Quark-Hadron Duality. Values of F_2 (measured from inelastic electron-nucleon scattering) are shown (solid lines) plotted against the variable $\omega' = W^2/Q^2+1$, which approximates $1/x$ at large Q^2 . Panels are presented in order of increasing Q^2 . As shown, the prominent nucleon resonances at low Q^2 diminish with increasing Q^2 and approach the scaling limit (dashed lines) indicating a transition to quark-hadron duality. Reproduced from [16].

In 1970, Bloom and Gilman published the first investigations of quark-hadron duality [16, 17]. They discovered that partonic calculations of $F_2(x)$ which were made in the *valid* scaling limit and then extrapolated down into the *invalid* resonance region followed a surprisingly

equivalent curve to experimental data of $F_2(\nu, Q^2)$. Specifically, they found that the *average* of the resonance structures agreed well with of the extrapolations (see Fig. 3). This behavior was found to be true both locally (averaging over an individual resonance) as well as globally (averaging over the whole resonance region).

2.4.1 EG1B AND JAM

Following this, investigations were made to reproduce this effect with the other structure functions, including g_1 .⁹ Among these was the 2000-2001 EG1b experiment [18–22] in Hall B [23] of Jefferson Lab which produced *abundant* polarized data for both the proton and neutron throughout the entire resonance and DIS kinematic ranges (with: $0.06 \text{ GeV}^2 < Q^2 < 5 \text{ GeV}^2$, $W < 3 \text{ GeV}$, and electron beam energies of 1.6, 2.5, 4.2, and 5.7 GeV). EG1b represents the most comprehensive data set of its kind to date and provides an exceptional and unique opportunity to study quark hadron duality in g_1 .

Of course, robust experimental data is only half of the duality puzzle. For the partonic half, we look to the Jefferson Lab Angular Momentum Collaboration (JAM). The JAM Collaboration brings together theorists, experimentalists, and computer scientists from within the Jefferson Lab community in an effort to study the internal structure of nuclei through the theoretical lens of QCD. Specifically, for this thesis, I will be discussing their JAM15 [24] fit of inclusive, spin dependent, DIS, world data which comprises 2515 total data points from 35 experiments (including EG1b) across 4 facilities (CERN, SLAC, DESY, and JLab).

Unique to JAM15 is the use of a novel, iterative, Monte Carlo fitting method. This method utilizes data resampling and cross-validation techniques to generate bounded, PDF predicted observables - including g_1 for both the proton and neutron. This produces DIS region pQCD fits for not just leading twist, but also (optionally) including higher twist and target mass corrections (TMC) for both the proton and neutron. These were then extrapolated down into the resonance region for comparison with equivalent EG1b data (with the same x but lower Q^2). In this investigation, we studied both nucleon flavors each with the case of exclusively leading twist contribution and the case including target mass and higher twist contributions.¹⁰

⁹See the following section for a detailed history.

¹⁰For an more information on PDF's see [25] and for TMC's see [26–28]

2.4.2 RESULTS

Scaling Behavior

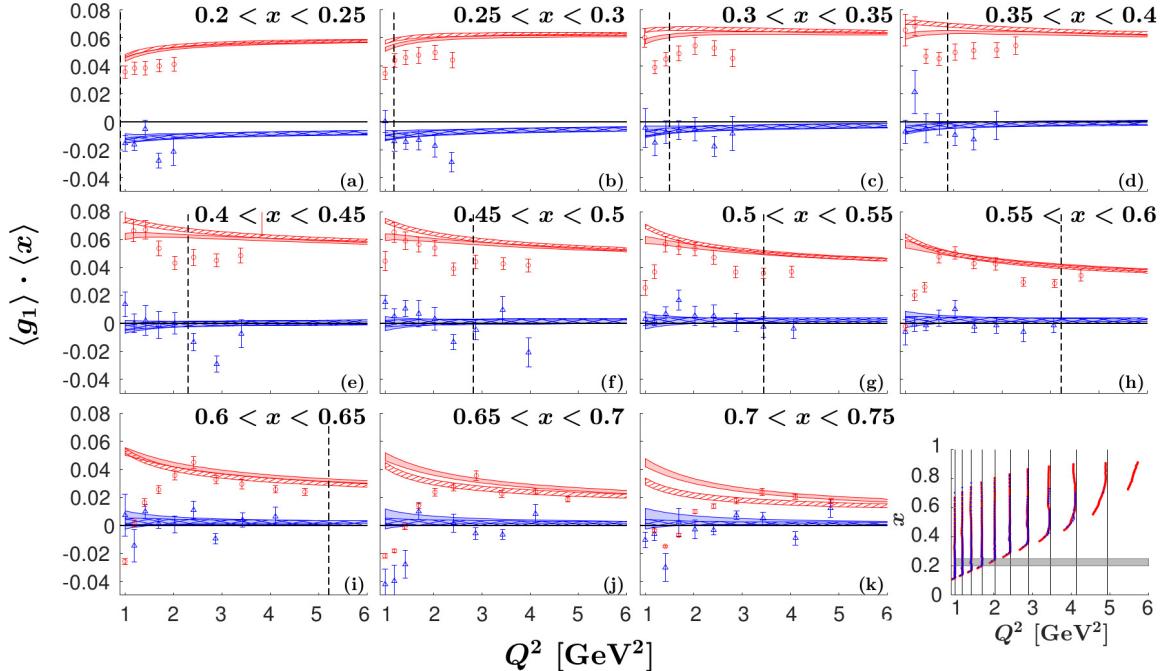


FIG. 4: Approach towards the scaling limit for the structure function $g_1(x, Q^2)$, averaged over 11 different x -bins of width $\Delta x = 0.05$, as a function of Q^2 . Data and JAM15 bands are shown multiplied with the average x for each bin for better clarity. The top (red) bands and data points (circles) are for the proton, and the bottom (blue) bands and data points (triangles) are for the neutron. The solid bands show the full prediction from the extrapolated JAM15 fit, including target mass and higher twist contributions. The striped band (proton) and the cross-hatched band (neutron) show the results including only the leading twist contribution. The vertical dashed line indicates the limit $W = 2$ GeV of the resonance region, which lies to the left. Last panel: kinematic location of all data points from EG1b in the x vs. Q^2 plane (red - proton, blue - neutron). The grey band indicates a sample interval in x over which the data are averaged (corresponding to panel (a)), and the vertical lines indicate the nominal central values of each Q^2 bin. Reproduced from [14].

As an initial comparison, we can examine the scaling behavior across all of the sets. Fig. 4 shows g_1 values from the EG1b experiment (vs. Q^2) averaged within bins of $\Delta x = 0.05$ (progressing from $x = 0.2$ to $x = 0.75$). Each bin shows the behavior of g_1 across the DIS and resonance regions (at least as much as can be shown for the given range of x and Q^2). As successive bins increase in momentum fraction x , they represent scattering events with successively larger momentum contributions - approaching the $x = 1$ limit of the overall nucleon itself. (A vertical dashed line is provided as reference to indicate the kinematic boundary between regions ($W = 2$ GeV) where applicable - with resonances laying to the left and DIS to the right.)

It would be expected for the experimental data (error bar points) to mostly agree with the JAM15 fits (shaded bands) in regions furthest into the DIS and for all sets to flatten as they approach that limit. Indeed, the first plot seems to show just that - albeit with limited (and offset) data points and no resonance region representation. However, moving further in x , significant structure begins to appear in the experimental data which deviates strongly from the JAM15 parameterization (but returns to agreement on either side of the deviation). This differs somewhat from (for example) the original F_2 findings in Fig. 3 which deviate *about* the theory curves but equal them *on average*. For our investigation, the data seem to consistently average to a value just *below* that of the extrapolated fits before eventually (in most cases) moving to agreement in high Q^2 .

Also of interest is the observation that this eventual agreement (and the resultant scaling behavior) occurs sooner (i.e., at lower Q^2) as the x bin increases. In fact, the plots in which the nominal x -value excludes the DIS region within EG1b kinematics show scaling onset as early as $Q^2 = 3$ GeV² - corresponding roughly to the end of the $\Delta(1232)$ resonance. This effect, if confirmed, may be useful in extracting spin structure function behavior at large x [29].

Integrated Duality

Moving on, we next conducted integrated tests of both local and global duality. We first selected six ranges of W corresponding to: four local ranges (covering discrete groups of resonance peaks) and two global ranges (one including the elastic peak and one excluding it) as shown in TABLE 1. For each range, we conducted a direct comparison (between experimental data and extrapolated fits) for truncated integrals of g_1 .

For the experimental data, elastic contributions were calculated as

$$g_1^{el} = \frac{1}{2} \frac{G_E G_M + \tau G_M^2}{1 + \tau} \delta(x - 1), \quad (39)$$

where $G_M = F_1 + F_2$ and $G_E = F_1 - \frac{Q^2}{4M^2} F_2$ are the magnetic and electric Sachs form factors [30]. This value was then added, by hand, to the truncated experimental integral (which had been calculated excluding the elastic limit). For the corresponding QCD fits, however, we simply extended the integral to $x = 1$ (without further compensation).

	Lower W limit	Upper W limit	Regions Covered by W Range
1(a)	1.072	1.38	$\Delta(1232)$
2(b)	1.38	1.58	$N(1440)1/2^+$, $N(1520)3/2^-$, and $N(1535)1/2^-$
3(c)	1.58	1.82	$N(1680)5/2^+$
4(d)	1.82	2	Remaining Resonances
5(e)	1.072	2	Global (excluding elastic peak)
6(f)	0.939	2	Global (including elastic peak)

TABLE 1: Selected W ranges, in GeV used for test of Quark-Hadron Duality. The listed resonance peaks are also visible in Fig. 3.

Each of the seven W ranges produced an independent test of duality. For a given test, experimental $g_1(x, Q^2)$ data were sorted into bins of Q^2 as listed in TABLE 2. Then, for each Q^2 bin, the data were integrated over x - with limits corresponding to the W limits for that particular test, such that:

$$\bar{\Gamma}_1(\Delta W, Q^2) = \int_{x_1(W_1, Q^2)}^{x_2(W_2, Q^2)} dx g_1(x, Q^2). \quad (40)$$

The statistical and experimental errors were finally added (in quadrature) into the integration and displayed as error bars.

The corresponding QCD fits were calculated by extrapolating the JAM15 results to the central Q^2 value of each bin - once each for the upper and lower limits. The theory integrals were then shown as bands corresponding to $\pm 1\sigma$ CL.

Lower Q^2	Upper Q^2	Central Q^2
0.92	1.10	1.00
1.10	1.31	1.20
1.31	1.56	1.43
1.56	1.87	1.71
1.87	2.23	2.04
2.23	2.66	2.43
2.66	3.17	2.91
3.17	3.79	3.47
3.79	4.52	4.14
4.52	5.40	4.94
5.40	6.45	5.90

TABLE 2: Experimental Q^2 Ranges, in GeV 2

Fig. 6 shows the results of these tests. As labeled in TABLE 1, the first four panels (a-d) show the tests of local duality while the last two panels (e and f) show the global duality tests.

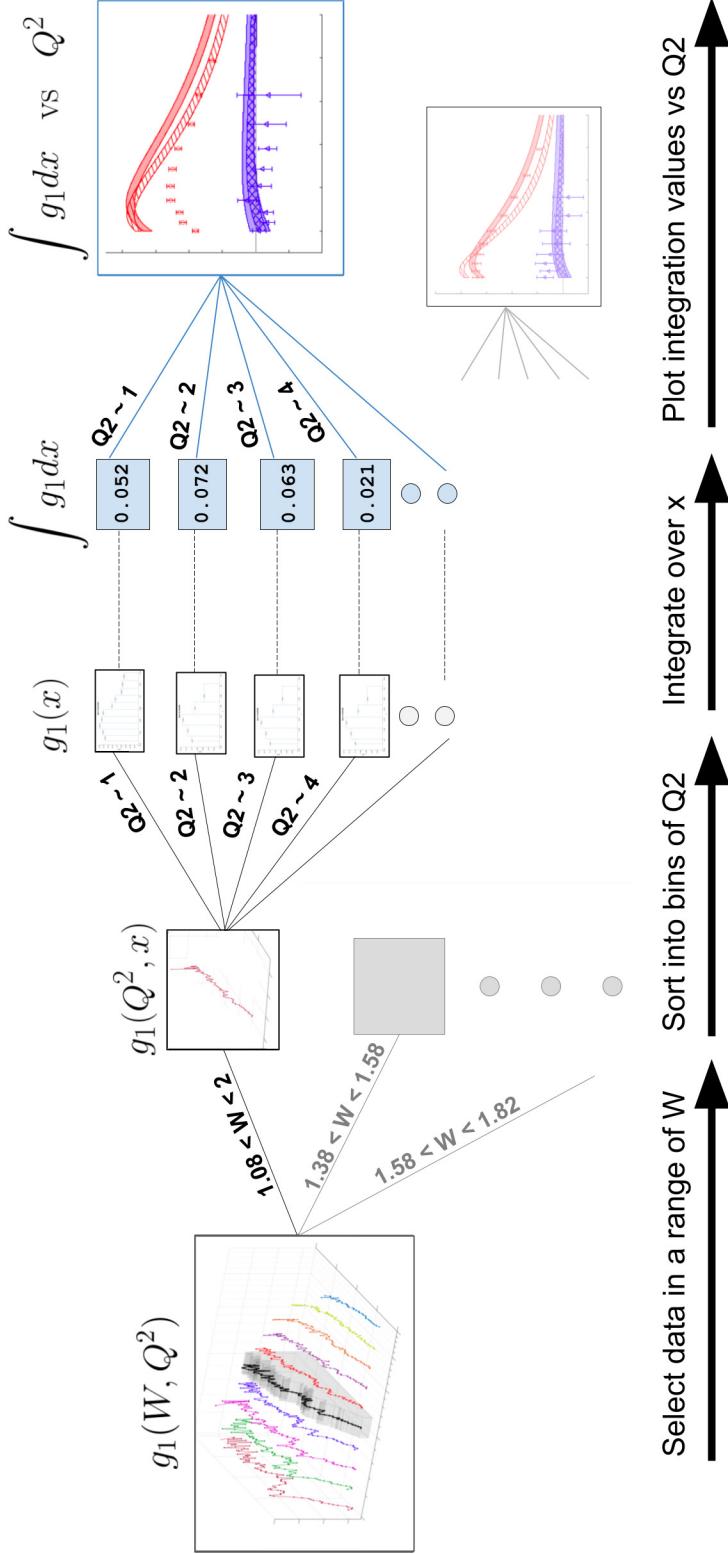


FIG. 5: Diagram showing the algorithm used to generate tests of quark-hadron duality. Data were first sorted into 6 bins of W as listed in TABLE 2. (Each of those bins generated a independent test of duality over a different kinematic band of the resonance region.) The W -binned data were then sorted into Q^2 bins - each of which were subsequently integrated by x . The result of those integrations were then plotted against the nominal Q^2 value of their Q^2 bin.

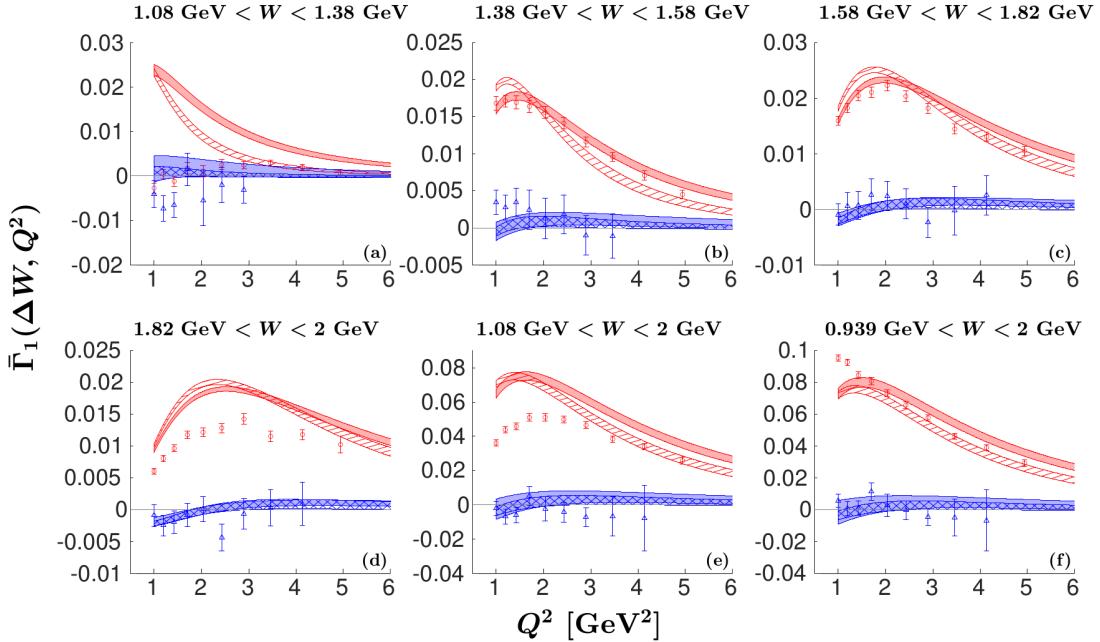


FIG. 6: Results of the integrated duality study. Shown are the $g_1(x, Q^2)$ integrals over regions of x ($\bar{\Gamma}_1(\Delta W, Q^2)$) corresponding to six kinematic regions of final state mass W , plotted as a function of Q^2 . Panel (a): The region of the first excited state of the nucleon, the $\Delta(1232)$ resonance. Panel (b): The region of the $N(1440)1/2^+$, $N(1520)3/2^-$, and $N(1535)1/2^-$ resonances. Panel (c): The region including the $N(1680)5/2^+$ resonance. Panel (d): The remainder of the customary resonance region, $1.82 \text{ GeV} < W < 2 \text{ GeV}$. Panel (e): The sum of regions (a) through (d), *i.e.*, the entire resonance region $1.07 \text{ GeV} < W < 2 \text{ GeV}$. Panel (f): Same as panel (e), with the elastic peak included: $0.938 \text{ GeV} < W < 2 \text{ GeV}$ (corresponding to a range in x extending all the way to $x = 1$). As in Fig. 4, the top (red) bands and data points (circles) are for the proton, and the bottom (blue) bands and data points (triangles) are for the neutron. The data points are shown with statistical and systematic uncertainties added in quadrature (error bars). The solid bands show the full prediction from the extrapolated JAM15 fit, including target mass and higher twist contributions. The striped band (proton) and the cross-hatched band (neutron) show the results including only the leading twist contribution. Reproduced from [14].

Beginning with panel (a), we observe that the resonance range closest to the elastic peak ($\Delta(1232)$) does not present compelling evidence for duality. Indeed, the experimental data all fall below or near zero while the extrapolations are all mostly (or significantly) positive. In fact, fits, as a rule, tend to predict rising positive asymmetries with large values of x (as is the case for low W). We see also that the effect is more pronounced in the higher twist fits (solid bands) than for the leading twist fits (striped bands).

As for the data, Δ resonances are known to be dominated by transitions in which $\sigma_{3/2}^{(T)}$ has a stronger effect than $\sigma_{1/2}^{(T)}$, which (as is apparent in Eq. (28)) produces negative A_1 (causing negative g_1). The disagreement is so strong that convergence of the data with the extrapolations only begins around $Q^2 > 3$ GeV 2 for the neutron, and even later (or not at all) for the proton indicating a clear failure of local duality in the Δ region.

In contrast to the ($\Delta(1232)$) resonance, the ($N(1440)1/2^+$, $N(1520)3/2^-$, and $N(1535)1/2^-$) resonance and the ($N(1680)5/2^+$) resonance of panels (b) and (c) display *compelling* evidence for local duality - particularly for the proton's higher twist fit. However, the final test of local duality, in panel (d), once again displays over-estimation in the theory curves compared to experiment. This is at least consistent with our assessment of panel (a) though - as this final range covers a collection of resonances that include Δ which carry the same helicity bias.

Lastly, we compare our global duality tests. Panel (e) displays the case where we exclude the elastic region (0.939 GeV $< W < 1.072$ GeV) whereas panel (f) shows that case of it being included. While the neutron results differ only minimally (with panel (f) having benefited from the addition of a positive valued elastic contribution and offering slightly better agreement), the proton plots present a *striking* comparison.

In the excluding case of panel (e), the data clearly suffer the same lowering influence from the Δ resonances that is observed in panels (a) and (d). However, the elastic region's positive contribution not only counteracts this discrepancy - it does so with a dependence on Q^2 which simultaneously corrects the low Q^2 data points *without* spoiling the agreement in the high Q^2 ones. As such, agreement occurs around $Q^2 = 1.4$ GeV 2 instead of $Q^2 > 3.5$ GeV 2 . This is perhaps particularly surprising given that the offending Δ resonances sit on both ends of the resonance region.

An overall assessment of these six tests suggests that duality may be a promising avenue by which to constrain future DIS fits. The requirement that extrapolations agree with either local resonances (which lack Δ) or with the entire resonance region (inclusive of the elastic range), may offer a route for greater data-informed constraints to QCD fits in the DIS region.

2.5 EXISTING DATA

Continuing the discussion of experimental data, the initial polarized resonance region experiments took place at SLAC in the late 1970's (for proton scattering [31, 32]) and early 1990's (for neutron scattering [33, 34]). The first of these, E80, utilized a polarized electron beam on polarized butanol with beam energies ranging from 6 to 13 GeV and a Q^2 of 2 GeV 2 . The second (E130) ran at a much higher 23 GeV with a Q^2 ranging from 3.5 to 10 GeV 2 . These early experiments seemed to confirm the naive parton model (for x greater than 0.6), indicating that nucleon spin could be accounted for entirely by its valence quarks. However, low luminosities due to the difficulty of producing polarized electrons made for large experimental uncertainties.

These were followed by the European Muon Collaboration (EMC) at CERN's Super Proton Synchrotron (SPS). These experiments replaced the electron beam with an automatically polarized muon one with higher beam energy providing greater kinematic reach - extending to a much lower x . This experiment reached energies of 200 GeV on an ammonia target. The results, published in 1988, showed violations of the Ellis-Jaffe sum rule in the high Q^2 , low x region - indicating valence quarks actually carry very little of the overall nucleon charge. These findings were what sparked the previously discussed spin crisis and the flurry of interest that followed.

A second generation of experiments at SLAC in the early 1990's utilized an improved electron beam on polarized ^3He (E142 [33] and E154 [35]) and polarized ammonia (E143 [34] and E155 [36]). These produced an incredibly high precision data set for $x > 0.02$ and remain unsurpassed in kinematic reach and statistical precision for high Q^2 .

Soon after, the HERMES Collaboration [37] (at DESY), the SMC and COMPASS experiments [38, 39] (at CERN), and Halls B [40] and A [41] (of Jefferson Lab) expanded the kinematic range of measured polarized data. Lastly, early in the 21st century (building on the previous experimental developments) high-precision experiments in Halls A [42–46], B [18–22, 47] and C [48] (of Jefferson Lab) provided a wealth of statistically precise polarized nucleon data that includes EG1b.

2.5.1 RGC EXPERIMENT SUITE

Even with this rich history of data, the field of polarized scattering experiments still has much to be explored. We have already discussed the current state of opacity surrounding the polarized structure functions, photon asymmetries, parton models, and duality - all of which

can only be clarified through continued polarized DIS investigations. In addition to these, however, are many other unanswered questions not yet touched on in this thesis. Transverse Momentum Distributions (TMDs), which are accessible through Semi-inclusive DIS (SIDIS) experiments, also require polarized data. Likewise, Generalized Parton Distributions (GPDs) require polarized Deeply Virtual Compton Scattering (DVCS) experiments to measure.

Fortunately, each of these pursuits require nearly the same experimental setup and can be carried out in tandem. To that end, a suite of polarized nuclear experiments were approved to complement the upgraded accelerator and detector capabilities at Jefferson Lab (see Section 4.1). The official list of accepted proposals are shown in Table 3. This suite (called Run Group C or RGC) represents an enormous investment of both physical resources (having been granted eight months of operation time) and effort (with collaborators from across the globe). The development of a polarized target which meets the needs of RGC is the subject of the remainder of this thesis.

Proposal ID	Title
E12-06-109	Longitudinal Spin Structure of the Nucleon
E12-06-109A	DVCS on the Neutron with Polarized Deuterium Target
E12-06-119(b)	DVCS on Longitudinally Polarized Proton Target
E12-07-107	Spin-Orbit Correlations with Longitudinally Polarized Target
E12-09-007(b)	Study of Partonic Distributions using SIDIS K Production
E12-09-009	Spin-Orbit Correlations in K Production with Polarized Targets

TABLE 3: List of experiments within Run Group C.

CHAPTER 3

SOLID CRYOGENIC POLARIZED TARGETS

3.1 DYNAMIC NUCLEAR POLARIZATION

Measuring the spin-dependent observables covered in the previous section require targets with a preferred orientation of their nuclear spin. This polarization can be expressed as a vector whose direction indicates the direction of alignment and whose magnitude indicates the excess of particles whose spins point in that direction (ranging from 0 to 1). For spin 1/2 particles, such as the proton and electron, spin can only be parallel (\uparrow) or anti-parallel (\downarrow) to the direction of measurement, producing a vector polarization of:

$$P = \frac{N_{\uparrow} - N_{\downarrow}}{N_{\uparrow} + N_{\downarrow}}, \quad (41)$$

where P is the vector polarization magnitude and $N_{\uparrow(\downarrow)}$ is the number of particles with spin pointing parallel (anti-parallel).

Under normal circumstances, a target sample has no preferential direction ($P=0$) as all spin directions have equal energy. However, if placed in an external magnetic field, the energy levels for the different spin orientations split via the Zeeman effect. The difference between these energy levels is determined by the the strength of the magnetic field (B_0) and the magnetic moment of the particle (μ), and is given by:

$$\Delta E = g\mu B_0, \quad (42)$$

where g is the particle's g-factor. This energy differential causes a distribution of energy level occupations which follows the Maxwell-Boltzmann law:

$$N_{\uparrow} = N_{\downarrow} e^{(-\Delta E/k_B T)}, \quad (43)$$

where T is temperature and k_B is the Boltzmann constant. This generates a polarization of:

$$P = \frac{e^{(\Delta E/2kT)} - e^{(-\Delta E/2kT)}}{e^{(\Delta E/2kT)} + e^{(-\Delta E/2kT)}} = \tanh\left(\frac{\Delta E}{2kT}\right), \quad (44)$$

known as as the thermal equilibrium polarization, or just TE.

The electron, with its large magnetic moment, reaches a TE polarization of 99% when placed in a 5 T magnetic field at 1 K. However, the proton with a magnetic moment of $\mu_p \approx \frac{1}{660}\mu_e$ placed in the same environment polarizes to only half of a percent. As the uncertainty of asymmetry measurements is proportional to the inverse of the polarization, it is necessary to obtain polarizations *much* higher than TE. To that end, dynamic nuclear polarization (DNP) must be employed.

Proposed in 1953 by Overhauser [5], it was first demonstrated by Carver and Slichter [6] the same year. The technique requires the same high magnetic field and low temperature that allows for high TE polarization of electrons, but also introduces free radical electrons (paramagnetic centers) and microwaves to transfer that polarization to the nucleon. Though there are several mechanisms at work depending on the setup, the simplest to explain is the solid state effect.¹

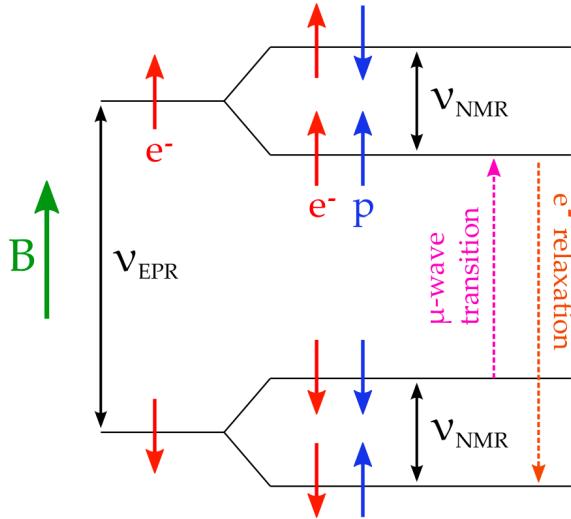


FIG. 7: Diagram of simplified energy levels and transitions for electron-proton pairs. ν_{EPR} indicates the frequency for electron transitions, and ν_{NMR} indicates the frequency for proton transitions, and \uparrow / \downarrow indicate spin orientation for the electron (red) and proton (blue). Reproduced from [51].

¹For a complete review of the of the various spin-transfer mechanisms, see [49] and, for a comprehensive discussion on the design and operation of polarized targets in general, see [50].

The solid state effect applies to the case of a single free electron paired with a nearby proton. The Zeeman split energy levels of the electron are then combined with the hyperfine splitting of the proton to produce four discrete energy levels corresponding to the four possible spin combinations of the particle pair (see Fig. 7).

Now, were these four levels independent, the dipole selection rules would exclude any transitions which simultaneously changed the spin of both particles. However, these energy levels (and the Hamiltonian which generates them) neglect the spin-spin interaction between the proton and electron. Such a term (H_{SS}), once added to the existing Zeeman interaction terms,

$$H = \vec{\mu}_e \cdot \vec{B} + \vec{\mu}_p \cdot \vec{B} + H_{SS}, \quad (45)$$

and solved perturbatively, results in the mixing of spin states and thus allows these forbidden transitions to occur. Functionally, however, these mixed levels are very close in value to the unperturbed ones and, as such, are labelled so in Fig. 7 for simplicity.

With this background, microwaves of frequency $\nu_{EPR} \pm \nu_{NMR}$ can be introduced to excite transitions from the lower electron level to the higher (where EPR and NMR indicate the electron paramagnetic resonance and nuclear magnetic resonance frequencies). This by itself would create an increased proton polarization, but only for individual proton-electron pairs. The final key to the solid state effect is the unmatched relaxation rates between the electrons and protons. The relaxation rate for electrons is much higher than that of protons (millisecond compared to minutes), meaning an electron can relax back into the lower energy state while maintaining the proton polarization thus quickly freeing the electron to pair with (and excite) another proton. In this way, a single free electron can generate and maintain a localized proton polarization much higher (and more quickly) than TE would allow.

This increased polarization is localized to the paramagnetic centers, though. Fortunately, the process of spin diffusion allows the proton polarization to propagate throughout the target sample. Nucleon dipole-dipole interactions cause adjacent protons to exchange spin orientations. As no energy is lost in the process, there is little delay between events (mere hundreds of microseconds between events).

Of course, this process is simplified. In reality, the paramagnetic centers are dense enough to have non-negligible effects on one another. In fact, for the materials and free radical densities typically used with DNP, the electrons form a reservoir with quasi-continuous bands of energy generated by the numerous dipole-dipole interactions with one another. These

bands have their own Boltzmann distributions characterized by a spin temperature T_s .² As the system absorbs microwave energy, polarization is exchanged between the electron spin reservoir and nucleons via thermal mixing. This process is referred to as the equal spin temperature theory and is actually the dominant process in our case. Fortunately, its requirements are identical to that of the solid state effect.

3.2 TARGET MATERIAL

Based on the requirements for DNP, the ideal target material would consist solely of the desired nucleons and free electrons. Obviously, for neutrons, that is impractical due to their instability and difficulty to produce in sufficient numbers. Whereas for protons the obvious candidate, H_2 (diatomic hydrogen gas), unfortunately enters a non-polarizable spin-zero state at low temperatures. So, compromise must be found in materials which have a high proportion of the desired nucleon (leading to high a dilution factor) and which can also obtain suitable degrees of polarization. The necessary inclusion of other particle species to the target introduces a background to the asymmetry measurements which must be mitigated by longer running times for a given statistical uncertainty:

$$t \propto \frac{1}{f^2 P^2 \mathcal{L} \Delta A^2}, \quad (46)$$

where t is the run time, ΔA is the asymmetry accuracy, \mathcal{L} is the beam luminosity, P is the polarization, and f is the dilution factor or, more precisely, the ratio of counts from the desired polarizable nucleons to total number of counts from *all* (free and bound) nucleons exposed to the beam. Additionally, to further minimize the experiment duration, materials should be able to obtain polarization in a reasonable amount of time. By the equal spin temperature theory (discussed earlier) the polarization of different spin 1/2 particles species are simultaneously dependent on their common spin temperature (T_{SS}) as given by:

$$P = \tanh \left(\frac{g_1 \mu_i B}{2k_B T_{SS}} \right). \quad (47)$$

So, provided the magnetic moment of the extraneous particles in the sample is sufficiently small compared to that of the desired nucleon of study, their effects are minimized as well as calculable even when not independently measured.

Common candidates for suitable target materials include lithium hydride (LiH), butanol (C_4H_9OH), ammonia (NH_3) and their deuterated counterparts. Lithium hydride, though it

²Incidentally, the effective proton polarization can be expressed by substituting T_s for T in Eq. (44).

polarizes well, only has one hydrogen per six nucleons. Butanol polarizes quickly, but cannot withstand the radiation damage caused by an electron beam. Ammonia, however, passes all the criteria, is well studied, and has been used in many polarized target experiments [52]. Depending on experimental criteria, the less common (but still stable) isotope $^{15}\text{NH}_3$ is sometimes used, as it is spin 1/2 as opposed to the naturally occurring spin 1 isotope $^{14}\text{NH}_3$.

For the introduction of paramagnetic centers, techniques vary with material. For butanol, and other materials which are liquid at room temperature, persistent radicals are introduced by dissolving in a chemical dopant such as porphyrexide or TEMPO. For materials like ammonia, though, radicals are created using high-intensity ionizing radiation which removes a proton from the molecule to form an NH_2^+ ion (with an unpaired electron). This process also produces other paramagnetic species such as f-centers and hydrogen atoms. These radicals are not believed to play a role in the polarization process but do, however, relax the nuclear spins towards their TE polarization - limiting the maximum achievable polarization. Fortunately, it seems these unwanted species recombine at a lower temperature than the desired NH_2^+ ion. This means that warm dose (near 90 k) irradiation can be used to produce highly polarizable ammonia samples.

Of course, as the experiment is conducted using a high-energy electron beam, cold dose (below 4 K) irradiation happens *in situ* continuously during the experiment. Once the number of free radicals increases beyond the optimal proportion, the polarizability of the material decreases as the oversaturation of paramagnetic centers presents an excess of spin relaxation paths. The solution is to carefully warm the target to temperatures slightly above the warm irradiation temperature (90-100 K) to allow the surplus radicals to recombine and restore polarizability. This procedure (called annealing) prolongs the lifetime of a target sample but, unfortunately, is an imperfect process. Over successive anneals, the material sustains irrecoverable damage and must eventually be retired.

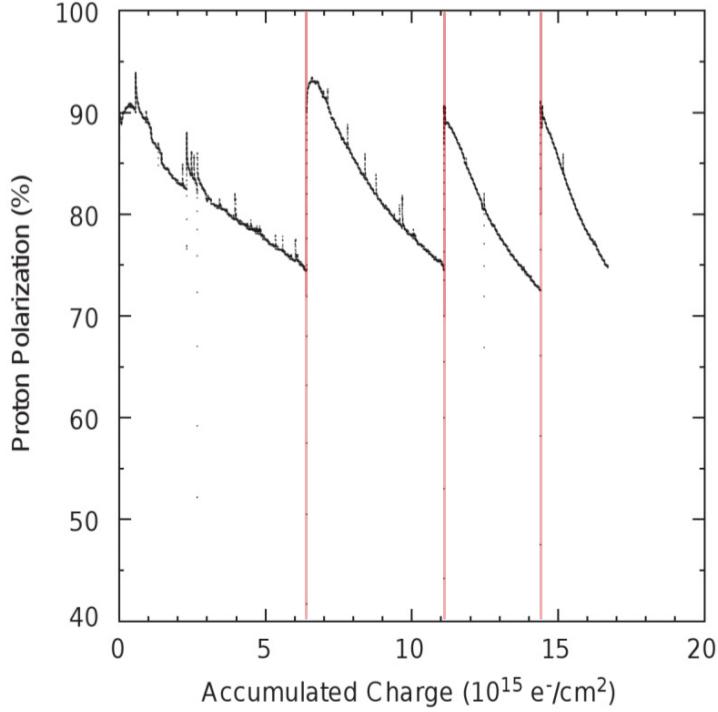


FIG. 8: Example from SLAC E155 of diminishing polarization with dose with multiple anneals (indicated in red). Modified from [53].

To delay this eventuality and best utilize the complete sample, it is important to distribute the beam dose as homogeneously as possible through the material. For liquid or gaseous targets, this is achieved naturally through the convective movement of the sample across the path of the beam. For solid targets, however, the beam itself must be moved (or *rastered*). The implementation of this is discussed in Sect 4.5.

Of course, all of the infrastructure surrounding and supporting the target cell that is traversed by the beam contributes to the event rate as well. When measuring asymmetries, the events for positive (n^+) and negative (n^-) polarizations are normalized to the total event count (n_{tot}):

$$A_{raw} = \frac{n^- - n^+}{n^- + n^+} = \frac{n^- - n^+}{n_{tot}}. \quad (48)$$

Ideally, however, we would normalize by only counts from the polarized particles of interest

- which would mean subtracting the irrelevant, diluting background counts (n_B) from the denominator:

$$A_{undiluted} = \frac{n^- - n^+}{n_{tot} - n_B}. \quad (49)$$

To that end, we can define the dilution factor as:

$$f = \frac{n_{tot} - n_B}{n_{tot}}, \quad (50)$$

to produce:

$$A_{undiluted} = \frac{A_{raw}}{f}. \quad (51)$$

Note also, that

$$f = \frac{n_{tot} - n_B}{n_{tot}} = 1 - \frac{n_B}{n_{tot}}, \quad (52)$$

returns us to our previous definition of “the ratio of desired nucleons to total nucleons” with the caveat that the “total nucleons” includes everything in the beam path. This leaves the non-trivial task of determining n_B . The most practical approach is to directly measure this value by removing the desired nucleons and repeating the experiment as closely as possible. This task requires the creation of special target cells with materials closely approximating the polarized ones using either pure nitrogen or carbon (to be discussed later).

Slightly complicating the issue is the nature of the polarizable compounds themselves. At room temperature, ammonia is a caustic, hazardous gas which melts at a temperature of 195 K. Liquid nitrogen is used to solidify it into a single block which is then crushed and sieved into irregular, but approximately equivalent pieces about a millimeter in width. These beads are then tightly packed into cells (leaving sufficient space between beads for the cooling via liquid helium) then weighed and characterized. This whole process must occur within a liquid nitrogen bath to protect the samples from evaporation and the sample preparers from corrosion.

3.3 MAGNETICS

As previously discussed, the polarization of any target material is strongly dependent on the strength of the external magnetic field (see Eq. (44)). In reality, though, magnetic fields are never perfectly uniform over an extended volume. So, some regions of target material

will see a different field than others. Fig. 9 shows, for a given microwave frequency, the polarization of ammonia and deuterated ammonia as a function of the field strength.

For the proton to keep above 80% polarization, the field can vary by up to 20 G around its optimal value (negative and positive peaks in Fig. 9), but the much more sensitive deuteron loses half its polarization with that degree of variance. So, for optimal polarization of all targets (and in all parts of the target cell), no part of the field can vary by more than a few G (at 5 T). In other words, the field uniformity $\frac{\Delta B}{B}$ must be no worse than 10^{-4} .

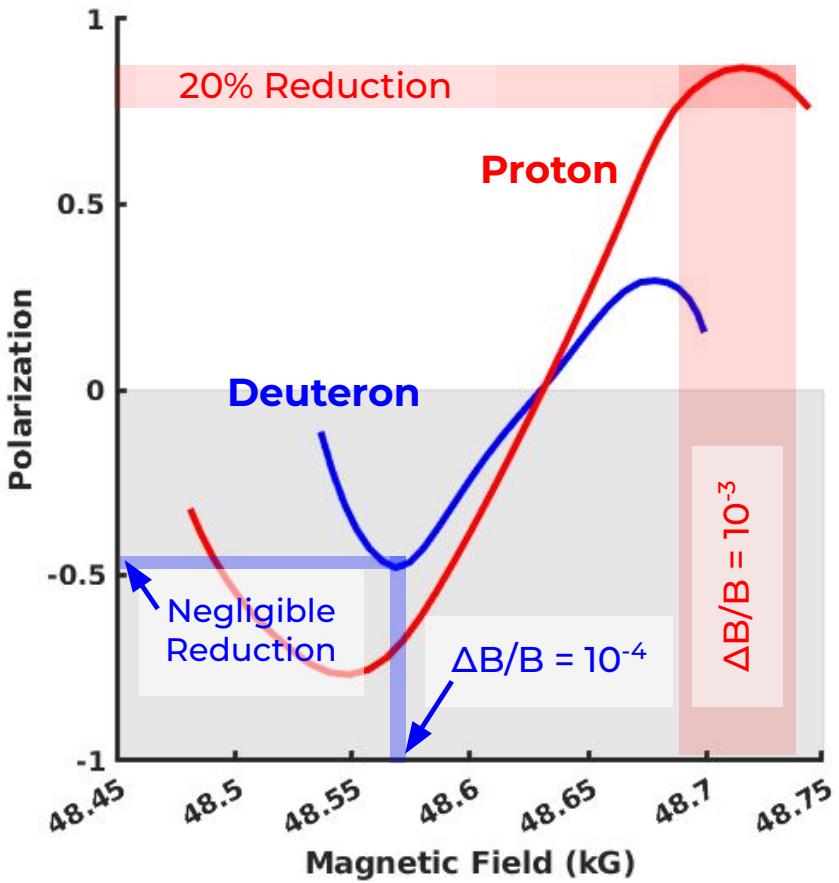


FIG. 9: Polarization of ammonia (red) and deuterated ammonia (blue) vs. magnetic field strength at a microwave frequency of 136 GHz with bars indicating the reduction of polarization for indicated field variance. Data from [54].

3.4 MICROWAVES

The final component needed for DNP is the microwave energy to induce spin transitions. These are generated using an Extended Interaction Oscillator (EIO) tube. Within the resonant cavity of EIO generator's klystron, accelerated electrons pass from a cathode filament to a collector, emitting coherent radiation at a frequency dependent on the size of the cavity. By mechanically adjusting this size, specific frequencies can be selected. For a 5 T field, this frequency is near 140 GHz (~ 2 mm wavelength).

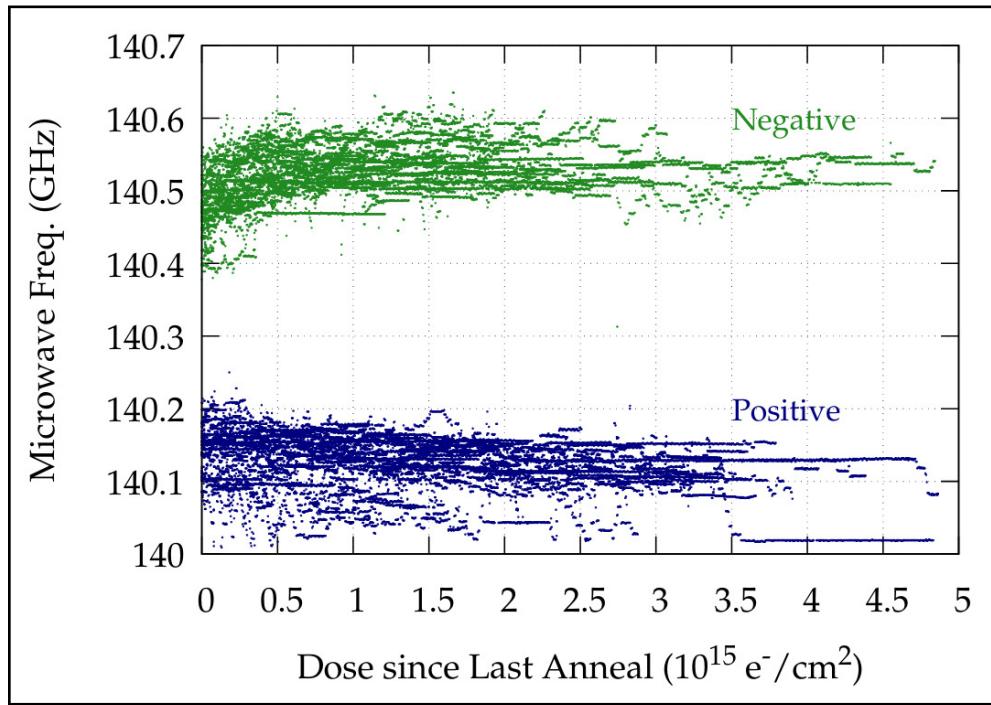


FIG. 10: Optimal microwave frequency varying with beam dose for positive and negative polarizations. Reproduced from [55].

This tuned energy is then directed down waveguides before exiting into a scattering chamber to irradiate the target material. While the overall power should be on the order

of a Watt, it is important that it be dispersed evenly throughout the cavity (and material) to prevent either unpolarized regions or regions of excess heat. To help achieve this, the exit of the waveguide is formed into a dispersive horn. Of course, as discussed in the DNP section, the energy transferred to the spin lattices actually forms a pseudo-continuous band of frequencies. So, to best utilize all possible polarization paths, the frequency of microwaves is modulated over a small range covering this band. Additionally, as the target material accumulates cold dose from the electron beam, the optimal polarization frequency changes as shown in Fig. 10.

3.5 NMR

Achieving maximum polarization while receiving beam is a delicate and time variable process. To ensure that all of the components of DNP are adjusted correctly, it is important to be able to monitor the target's performance in real time. For this, we use a CW-NMR (continuous-wave nuclear magnet resonance) system. Fundamentally, an NMR system is simply an RLC (resistor - inductor - capacitor) circuit. When alternating current flows through an RLC circuit, energy is periodically transferred between the magnetic field of the inductor and the electric potential field of the capacitor. As with any periodic process, there is an optimal frequency at which the least energy is lost. This is called the circuit's resonance frequency and it is given by $\omega_0 = \sqrt{\frac{1}{LC}}$. One measure of a circuit's efficiency is its quality factor $Q = \frac{\omega_0 L}{R}$. Driving the circuit at frequencies other than ω_0 increases its impedance, lowering the Q-factor. The closer the driving power of the circuit is to its resonant frequency, the higher its Q , and the less energy is lost during each cycle. Any increases in impedance, including interactions with the circuit's magnetic field, lower the quality factor.

As discussed previously, a spin possessing particle in a external magnetic field experiences Larmor precession. That is to say, its magnetic moment rotates perpendicular to the direction of the field and at a rate that is proportional to the field strength. This precession, if located within the magnetic field of an RLC inductor, induces a back EMF which affects the Q-factor of the circuit. The magnitude of this effect is directly proportional to the amount of polarization of the sample within the inductor. In a sense, the spin system responds to the RF field of the inductor in much the same way as it does to microwave irradiation. In both cases, the particles either emit or absorb energy to transition between spin states. For the circuit, this causes the inductance to become:

$$L_C(\omega) = L_0(1 + 4\pi\eta\chi(\omega)), \quad (53)$$

where L_0 is the unpolarized inductance, η is the filling factor, and $\chi(\omega)$ is the sample's magnetic susceptibility:

$$\chi(\omega) = \chi'(\omega) - i\chi''(\omega), \quad (54)$$

where χ' and χ'' are the dispersive and absorptive terms respectively. As the absorptive term indicates, negatively polarized particles absorbing energy to flip into a positive state, the overall polarization of the sample is proportional to the integral of this quantity over all frequencies (although, in practice, the only non-zero values are very near the Larmor frequency):

$$P \propto \int_0^\infty \chi''(\omega) d\omega. \quad (55)$$

This is apparent when plotting the circuit's Q-factor over a range of frequencies near ω_0 . In the absence of any polarization, the Q-factor steadily increases as it approaches the resonant frequency and smoothly decreases after passing it. However, in the presence of polarized material (and when tuned such that the ω_0 is near the Larmor frequency) this smooth curve is interrupted by a localized decrease (or increase for negatively polarized material) centered on the Larmor frequency. The integrated difference between this and the unperturbed plot is exactly proportional to the polarization of the sample (see Fig. 11).

Due to uncertainties in the filling factor (a measure of the coupling between the RF field and the sample) as well as the amount of polarizable material, the ratio between the integrated signal and polarization is not calculable from theory alone. However, as this ratio is dependent only on physical properties that do not change for a given arrangement, a single measurement of a known polarization allows for calibration of all subsequent measurements. Fortunately, such a known polarization value is readily available in the form of thermal equilibrium (TE) polarization (see Section 3.1 and Eq. (44)). Although small (and therefore difficult to measure), TE polarization is exactly calculable exclusively from temperature, particle species, and field strength. As this measurement determines the accuracy of all measurements that follow it, it is imperative that the NMR system is sensitive enough to reliably detect signals of such a small magnitude.

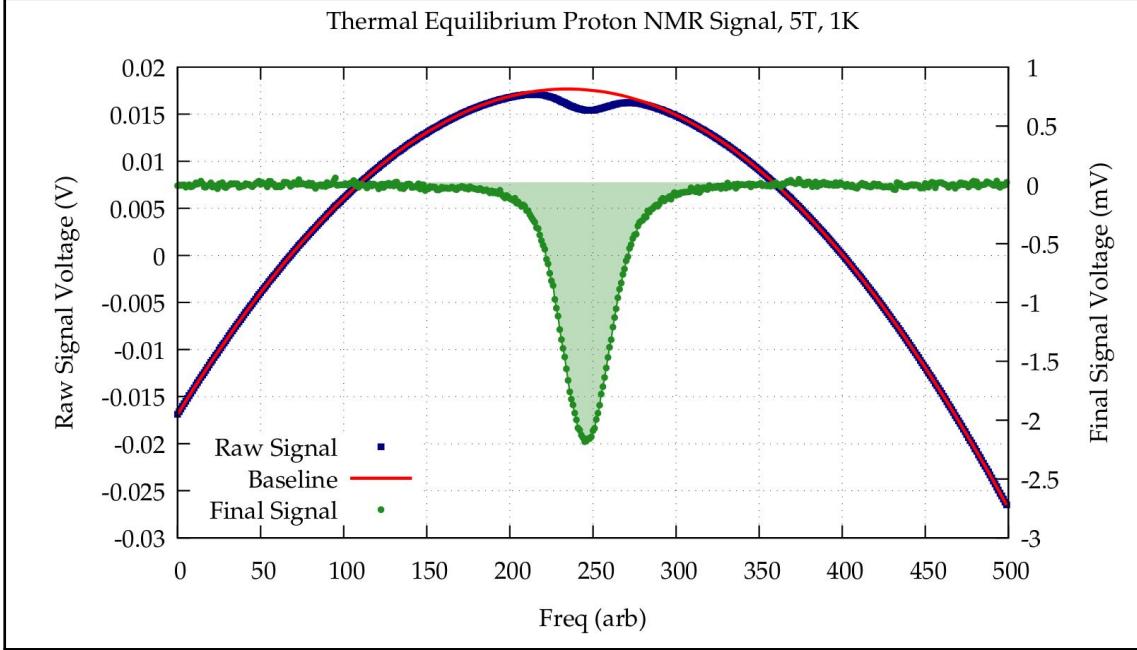


FIG. 11: Response of an NMR Q-curve to polarized material. Note that the voltage on the vertical axes refers to the voltage across the Q-meter's capacitor and is proportional to Q . The left axis shows the NMR circuit both without (red) and with (blue) polarized material within the magnetic field of the inductor. The integrated difference between the two (green) is proportional to the material's polarization (right vertical axis).

One limitation of this technique is that the magnitude of the back EMF caused by the sample is dependent not just on the sample's polarization, but also its location within the inductor's field - represented by the filling factor η . Particles located in strong regions of the field have a greater effect on the measured reduction in G than those in weaker parts of the field. So the accuracy of polarization measurements is dependent on inductor field uniformity. Care must be taken to optimize this uniformity and also to compensate for regions with smaller response so that strongly polarized particles in strong field locations do not mask depolarization in weak field regions.

CHAPTER 4

THE APOLLO TARGET

4.1 JEFFERSON LAB FACILITIES

CEBAF

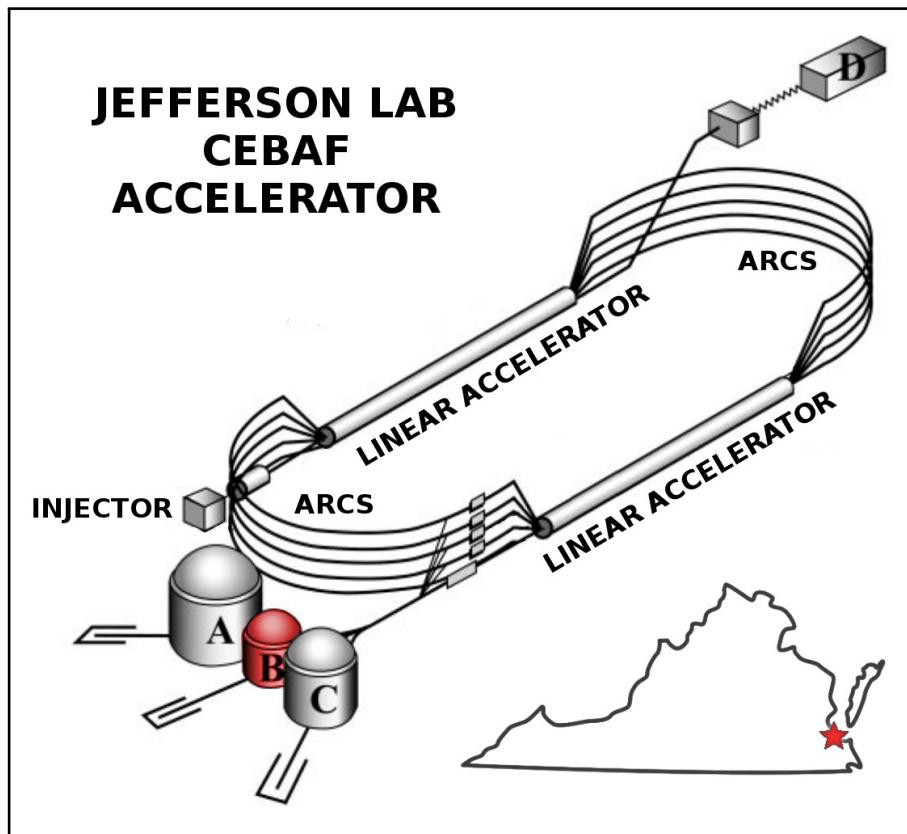


FIG. 12: Illustration of the Jefferson Laboratory CEBAF (Continuous Electron Beam Accelerator Facility) accelerator showing: north and south linear accelerators, compound recirculating arcs, beam injector, all four experimental halls (with Hall B highlighted in red), and the lab's location within the state of Virginia. Modified from [56].

The Continuous Electron Beam Accelerator Facility (CEBAF) is located at the Thomas Jefferson National Accelerator Facility (known colloquially as Jefferson Lab or simply JLab) in Newport News, Virginia. Established in 1984 as a one of the 17 US Department of Energy National Laboratories, the lab underwent an extensive upgrade project between 2010 and 2017 during which the available electron beam energy was doubled along with associated updates to multiple detector systems [56]. The accelerator can now produce a 12 GeV, high luminosity, high polarization, continuous-wave (CW) electron beam to four separate experimental end station halls simultaneously.

Beginning at the injector, four 250 MHz circularly polarized lasers (one for each hall) produce polarized electrons from a gallium arsenide cathode via the photoelectric effect. As these electrons move through the injector array, they are gathered into 499 MHz bunches and passed into one of two linear accelerators (linacs). These parallel linacs comprise a series of superconducting radio-frequency (SRF) cavities each of which transfer RF energy to the bunched electron beam.

At the end of either linac, the beam is curved through magnetic recirculation arcs which bend the electron path to connect it with the start of the opposite linac - reminiscent of a race track (see Fig. 12). As the magnetic fields required in these arcs vary with beam energy, each lap around the track requires a dedicated beam line and, as such, the total number of passes is limited (by infrastructure) to five.¹ From there, the electron bunches are directed to one of the four experimental halls (A, B, C or D), each of which can independently select the number of passes (beam energy) it receives as well as separately request beam currents (from 1 nA to 100 μ A).

The RGC experiments took place in Hall B with the maximum (five pass) beam energy of 10.6 GeV and beam currents of 4 nA and 10 nA. Once in the hall, beam position and quality is overseen by beamline monitoring components (see Fig. 47 and Section 4.5). Additionally, electron polarization is measured by a Møller polarimeter located upstream of the target (at the entrance of the hall).

¹Actually, due to its position at the opposite end of the track from the injector, the newest end station (Hall D) receives an extra half-lap for a total of 5.5 passes. This brings its beam energy up to the eponymous 12 GeV (as opposed to the 10.6 GeV available to the other three halls).

CLAS12 Spectrometer

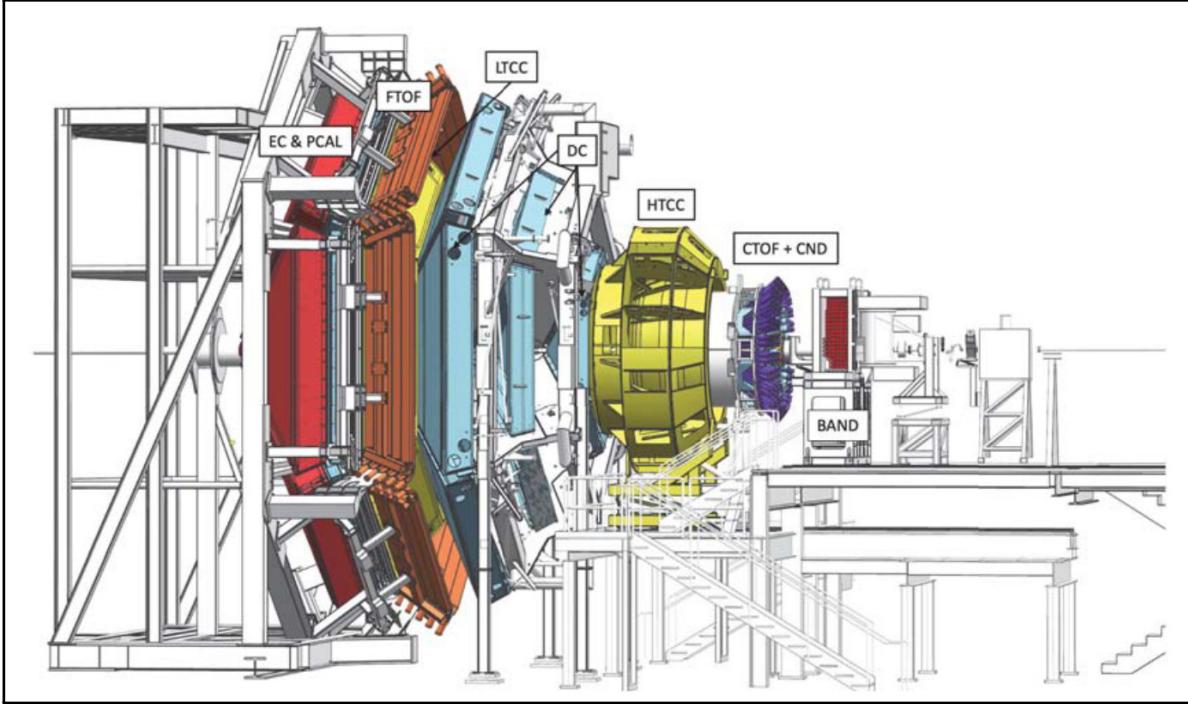


FIG. 13: Illustration of the Hall B CLAS12 spectrometer. The electron beam enters from the right, passes through the center of the spectrometer, and exits to the left. Note that, if comparing to Fig. 47 (or any other figure in the chapter), the orientation is mirrored and that the majority of the detector subsystems have not been rendered for that figure. Moving from downstream to upstream (left to right in this figure) the detector subsystems are: Electromagnetic & Preshower Calorimeters (leftmost red), Ring Imaging Cherenkov Detector (not pictured), Forward Time-Of-Flight Scintillation Counters (brown), Low Threshold Cherenkov Counters (leftmost yellow), Drift Chambers (light blue), Torus Magnet (white and interlaid with the drift chambers), Forward Tagger (obscured), High Threshold Cherenkov Counters (rightmost yellow), Solenoid Magnet (white), Silicon Vertex Trackers and Barrel Micromesh Trackers (obscured), and Central Time-of-Flight + Central Neutron Detectors (dark blue). Finally, the Back Angle Neutron Detector is shown (rightmost red) but **was not present** for the RGC experiment. Reproduced from [57].

Experiments in Hall B utilize the 12 GeV era CEBAF Large Acceptance Spectrometer (**CLAS12**) which replaced the previous 6 GeV era **CLAS** detector. In this context, large acceptance refers to the angular range over which the spectrometer can detect scattered particles (360 azimuthal degrees and 2.5 - 135 polar degrees). This is achieved through the near-cylindrical symmetry of the detector array (see Fig. 13).

The CLAS12 array is built out in concentric layers from its focal point (at the center of the solenoid) - with many detectors further divided azimuthally into six sectors. The array is conceptually split into two groups based on polar angle coverage (with 0° and 180° laying along the beam line). The Central Detectors (CD) surround the immediate vicinity of the target (35° - 135°) while the Forward Detectors (FD) extend downstream from it (2.5° - 40°). Combined, the system is able to characterize both charged and neutral particles across an exceptionally wide kinematic phase space.

As shown in Fig. 13, the FD array comprises: the Torus Magnet, Drift Chambers (DC), Electromagnetic and Pre-shower Calorimeters (EC/PCAL), Ring Imaging Cherenkov Detectors (RICH), Forward Time-Of-Flight Detectors (FTOF), Low Threshold Cherenkov Counters (LTCC), High Threshold Cherenkov Counters (HTCC), and the Forward Tagger (FT). Likewise, the CD array comprises: the Solenoid Magnet, Silicone Vertex Trackers (SVT), Barrel Micromesh Vertex Tracker MicroMegas (MVT/MM), Central Time-Of-Flight Detectors (CTOF), and the Central Neutron Detector (CND). Though a meaningful description of every component would exceed the scope of this text, Table 4 gives for a brief summary of each. Additionally, there are several which presented particular consideration to RGC.

Detector	Purpose
Central Detectors	
Solenoid [58]	Bends particles for use in CD tracking detectors and minimizes the radiation damage from high beam luminosity by bending low-energy background electrons inward to the Møller cone of the FT.
SVT [59]	Measures momentum and vertex of charged particles through track reconstruction via silicone strip sensors.
MVT/MM [60]	Improves particle track reconstructions, polar angle resolution of SVT, and vertex resolution of DC.
CTOF [61]	Identifies charged particles via plastic scintillators and photomultiplier tubes (PMTs).
CND [62]	Detects recoiled neutrons with momenta between .02 and 1 GeV for neutral particle identification.
Forward Detectors	
Torus [58]	Curves particle trajectories for identification / characterization via the FD array. Uses six large bore magnetic coils to provide an azimuthal field around the beamline.
DC [63]	Reconstructs charged particle tracks (via wire chambers) to extract momentum and vertex information. Sits within the Torus bores.
LTCC [64]	Differentiates between charged baryons with momenta between 3.5 and 9 GeV (refurbished from CLAS [65]).
HTCC [66]	Identifies particles with momenta greater than 5 GeV (via the production of Cherenkov radiation) and differentiates between charged particle species below 5 GeV.
FTOF [67]	Measures the time of flight of charged particles for particle ID and data triggering.
RICH [68]	Improves particle identification for charged baryons and for pions / kaons with momenta between 3 and 8 GeV.
EC/PCAL [69]	Identifies electrons, photons, neutrons, and $(\pi^0 \rightarrow \gamma\gamma)$ decay (refurbished from the CLAS EC [70]).
FT [71]	Detects electrons scattered between 2.5 and 5 degrees and identifies γ and π^0 with momenta up to 10 GeV. Contains the Møller cone which pairs with the solenoid to shield background radiation.

TABLE 4: Table of CLAS12 Components

Of the FD array, the Drift Chambers and Forward Tagger were most significant to the development of the RGC project. The FT detects very forward angle (small deflection) particles and, as such, sits immediately around the beamline. Incorporated into of the Forward Tagger is the Møller cone which shields the spectrometer from Møller background radiation generated by the beam. However, as will be discussed in Section 4.5, the RGC target requires the electron beam to be offset radially which reduces the effectiveness of the Møller cone and increases the rate of background events - potentially overwhelming the Drift Chambers closest to the beamline.

This can be compensated for by a reduction in beam current (which also lowers the overall data rate of the experiment) or, alternatively, the FT can be fully removed and replaced with a more robust Møller shield designed specifically for RGC. Removing the Forward Tagger, however, also removes a kinematic region from the spectrometer - one specifically important to DVCS studies. After extensive studies, a compromise was found in which both configurations were used during the experiment.

Another notable component in the FD array is the torus magnet which supplies the magnetic field needed to bend tracks for the particle identification. The geometry of the torus means that (depending on particle charge and torus current polarity) tracks can either bend inward toward the beamline (in-bending) or outward away from it (out-bending). This difference affects the acceptance and quality of data gathered in the experiment and had to be accounted for in the experimental run plan.

Finally, the solenoid magnet of the CD array was, of course, critically important to the RGC target in that is provided the overall polarizing field. As a spectrometer component, however, the solenoid's primary purpose is to provide the necessary magnetic field for the central tracking detectors and to deflect charged, low-energy background radiation away from the FD array. Since the CD detectors must occupy the interior of the solenoid to benefit from its field, the 78 cm warm bore diameter of the magnet is reduced to only \sim 10 cm. This places significant geometry constraints on any target - especially one needing to maintain the requirements of DNP.

CLAS12 Solenoid

The solenoid is composed of five separate superconducting cylindrical coils. The thin, outermost coil is counter-wound to the other four and runs the full length of the solenoid to minimize the fringe field outside the desired field region. [58] The other two sets of paired coils are thicker and concentrically nested within the outermost one (see Fig. 14).

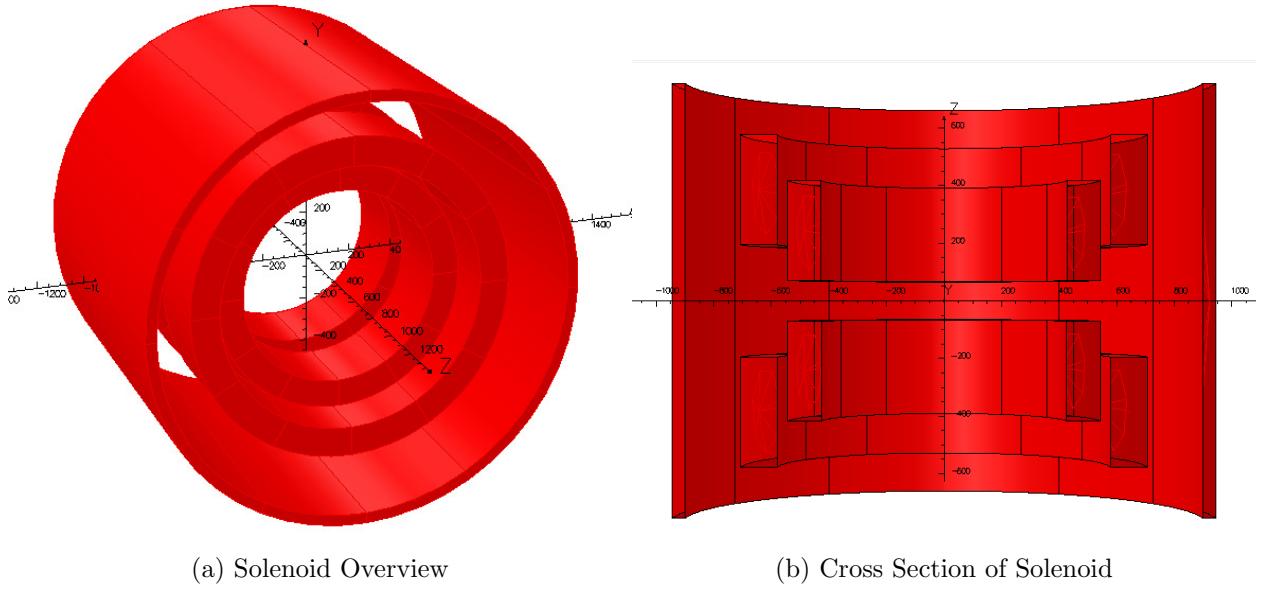


FIG. 14: Simplified OPERA3D rendering of the CLAS12 solenoid coils.

This assembly is wound on two independent bobbins and is contained within a sealed cryogenic housing. This housing prevents precision surveying of the coil geometry which, due to the process of shipping and installation, could potentially have shifted from its original position. Also, thermal contraction from cryogenic cooling and the expected shifts due to mutual forces of magnetic energization can only be estimated from calculation. As such, the precision of our knowledge of the location of the coils is not sufficient to precisely predict the field distribution.

Part of the commissioning of the solenoid was to conduct a mapping of its field. This task was carried out by running a set of three magnetic hall probes (one for each coordinate) down a tube that spanned the length of the solenoid. The canister containing the probes was driven forward in 1 cm steps, with each of the three probes logging their measurements at each step. The ends of the tube were affixed in place at either end of the solenoid by two large metal sheets that had been carefully and precisely placed and located. These sheets also allowed the tube to be moved away from the center for off-axis measurements at 1.25 cm and 30 cm (at various azimuthal angles). Additionally, to check self-consistency, the canister could be rotated within the tube in 90 degree steps and run that way. Most importantly, as far as DNP is concerned, special runs were conducted in 1 mm steps in just the center 10 cm region. The initial results of this central mapping are shown in Figure 15.

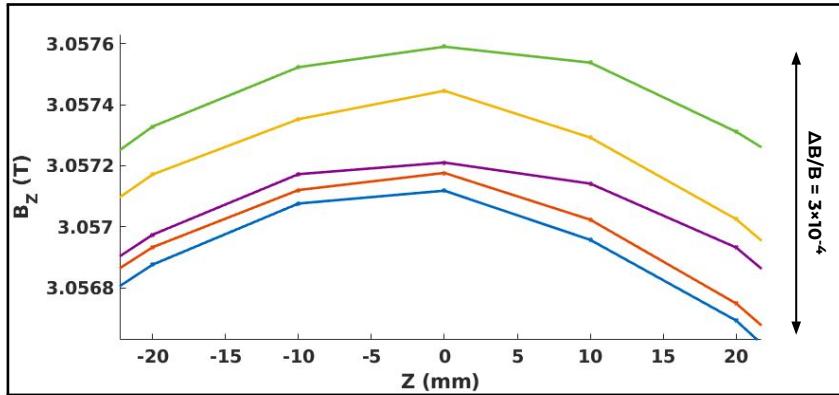


FIG. 15: Sample of results from the CLAS12 field mapping showing a uniformity of $\frac{\Delta B}{B} = 3 \times 10^{-4}$ in the central, on-axis region. Each line shows a set of measurements along the axis of the solenoid within a few cm of the center.

The measured homogeneity in the target region was initially measured to be 3×10^{-4} - three times worse compared to what is needed for optimal DNP. However, the data showed some discrepancies between measurements of the same type. To account for these, a closer analysis of the data was conducted which removed some, but not all, of the uncertainties,

and brought the homogeneity to 2×10^{-4} . Additionally, the measured field was found to deviate from that produced by the digital model of the solenoid. As such, efforts were taken to adjust the model to better match the measured field map. The outcome of that effort is shown in Figure 16.

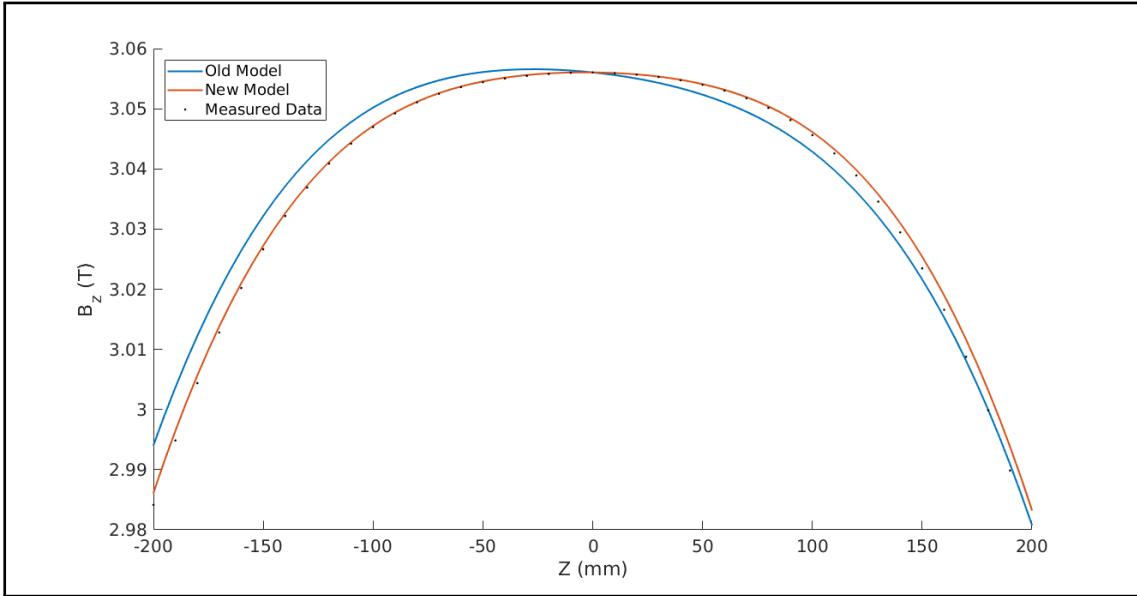


FIG. 16: Comparison of measured and modeled fields in the central region of the CLAS12 solenoid. See Appendix A for further comparisons and details on analysis.

To generate this updated model, small changes were made to the coils in the original model using finite element modeling software. The first (somewhat minor) change was to remove the small asymmetries present in the coil dimensions of the original model, which was done to ease the optimization process. From there, the innermost set were displaced asymmetrically along the z -axis (closer to the center by 5 and 7 mm, respectively) to match the asymmetry present in the measured data. They were also reduced symmetrically in radius by 25 mm.

The middle set of coils were not displaced along z or changed in radius (after the initial

removal of asymmetry), but they were extended outward along z , making them 20 mm longer than the original set. It is understood that adjustments of this magnitude represent non-physical changes to the solenoid model, however, they were necessary to produce any appreciable improvement to the agreement between the field map and the data. The final dimensions and placements for both models are shown in Table 5.

Coil	Inner Radius	Outer Radius	Z_1	Z_2
	New — Old	New — Old	New — Old	New — Old
1	402.36 — 427.36	519.68 — 544.68	63.81 — 70.81	410.65 — 417.65
2	402.36 — 426.63	519.68 — 544.50	-65.81 — -70.81	-412.65 — -417.65
3	707.71 — 707.71	580.06 — 580.06	194.89 — 194.89	596.37 — 576.37
4	707.71 — 707.66	580.06 — 578.74	-194.89 — -194.89	-596.37 — -576.37
5	-754.287 — -754.287	754.287 — 754.287	900.17 — 900.17	946.23 — 946.23

TABLE 5: Specifications of the coils for both models relative to the center of the solenoid (in mm) with differences between models shown in red.

Shim Coils

This new model is also not perfect, but it does adhere to the measured data much more closely. As such, it gives us a basis on which to design correction coils to bring the field within the needed uniformity for polarization. These coils would also enable localized adjustments of the field to be made without needing to change the overall field of the solenoid. This is useful for quick, temporary field changes, but also for adjusting the local field between polarization directions while maintaining the same overall solenoid field for particle tracking. Note that in Figure 9, the necessary field for positive and negative polarizations only differ by about 150 G.

As these correction coils will be incorporated into the target, the space they can exist in will be quite small. A preliminary study of coil performance indicated that the best off-axis

homogeneity was achieved with increasing coil radii so, as a compromise, the coil carrying mandrel (onto which the shims would be wound) was frozen in design at the largest possible radius of 3.3 cm and extended past the location of the target cells by only several cm. Also, since the uniformity in off-axis regions varied predictably and in conjunction with the on-axis uniformity (for coils of the same radius), it was found that different coil sets could be reliably compared using only the on-axis results.

Given the specific form that the variance in the solenoid field (hereafter referred to as the background field) took, it could have been possible to select a single coil, centrally located to counter the central peak of the background and render a mostly flat homogeneous field. However, this arrangement would leave very little room for discrepancies in coil placement and would only be useful in homogenizing the field to a single value (relative to the background). Conversely, any additional coils bring with them the necessity of additional leads which both require space in the target and present an added heat load. To leave room for potential asymmetries and to increase functionality, multi-coil designs were explored with a four coil arrangement being found to offer the best balance between utility and complexity.

The task then was determining the optimal configuration of those coils. For this, a Matlab script was written which attempted to minimize the difference between the field produced by a set of shims and some specified ideal field. The script calculated the field generated by the coils (along their shared axis) using a solenoidal application of the Biot–Savart law:

$$B_z = \frac{\mu_0 n I}{2} \left(\frac{z - z_1}{\sqrt{(z - z_1)^2 - R^2}} - \frac{z - z_2}{\sqrt{(z - z_2)^2 - R^2}} \right), \quad (56)$$

where μ_0 is the permeability of free space, n is the winding density, I is the current, R is the coil radius, z is the position of the field being calculated, and z_1 and z_2 are the beginning and end positions of the coil. As we assume the coils will be four layers thick, this calculation is performed four times per coil (once for each layer), then repeated for each of the four coils before being summed together. This occurs for every point along the length of the target region (in steps of 1 mm) to generate a close approximation of the total shim field. Several alternative techniques were tested, however, this method produced the best agreement with results from the 3D modeling software (OPERA3D) that was used to analyze any potential final coil configurations.

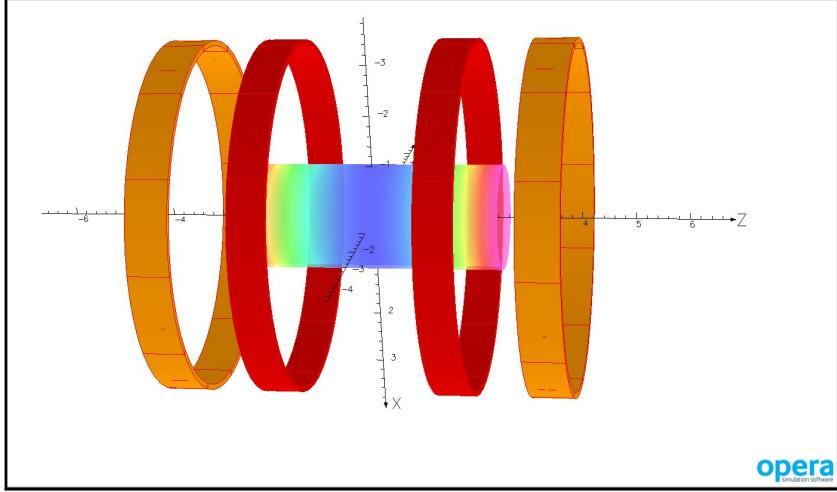


FIG. 17: An example of a four coil configuration as modelled with the OPERA3D software. The cylindrical center object indicates the magnetic field in the region of a target cell. The red and yellow hoops show the potential geometry of the shim coils which are symmetric about $z=0$.

The field generated from this calculation was then compared to an ideal, completely homogeneous field (B_0) to compute a merit of fit value: $\chi^2 = \int \frac{(B-B_0)^2}{B_0} dB$. This function was then run through one of Matlab's constrained minimization algorithms 'fmincon' to determine the best set of coil specifications. Each of the four coils has three degrees of freedom: position, width, and current.² This would make for 12 total free variables, however, the near symmetry of the background field allows for an assumption of symmetry in the coil configuration (about $z=0$) which reduces the problem to only 6 degrees of freedom.

Each degree of freedom also has limits on its allowable values. The available length on the mandrel to hold all four coils is 14 cm. This means the total combination of widths and positions cannot combine to exceed that length, nor can any adjacent coils be made to overlap one another. This could be programmed as a set of interdependencies or, more simply, the four dimensional variables (two coil widths and two coil positions mirrored across

²An alternative formulation could also have been used which varied the windings in each coil. But, for any change in winding density, the resultant total current could also be achieved by simply adjusting the current in the unchanged coil. So, an approximate winding density was selected based on the anticipated wire diameter.

the center) were reformulated as four widths - one for the width of the gap between the center ($z=0$) and the center-most coil, one for the width of the center-most coil, one for the width of the gap between the center-most and outer-most coil, and one for the width of the outermost coil. The only constraints then are that all widths must be positive (or zero) and that the sum of all widths must not exceed 7 cm. That leaves only the constraints for the two coil current variables which were bounded by ± 10 A due to the power supply limits.

A major difficulty in any minimization attempt is the avoidance of local minima. As the minimizer function moves towards a lower value, it can become trapped in a local minimum and conclude instead of finding a global minimum. The selection of the initial values used to begin the search greatly influences which minima the solver finds. Various schemes were employed to avoid this effect, but the most successful one ended up being a brute force, coarse scan of all initial geometric conditions which was only possible due to the low computational time of the solver (see Fig. 18).

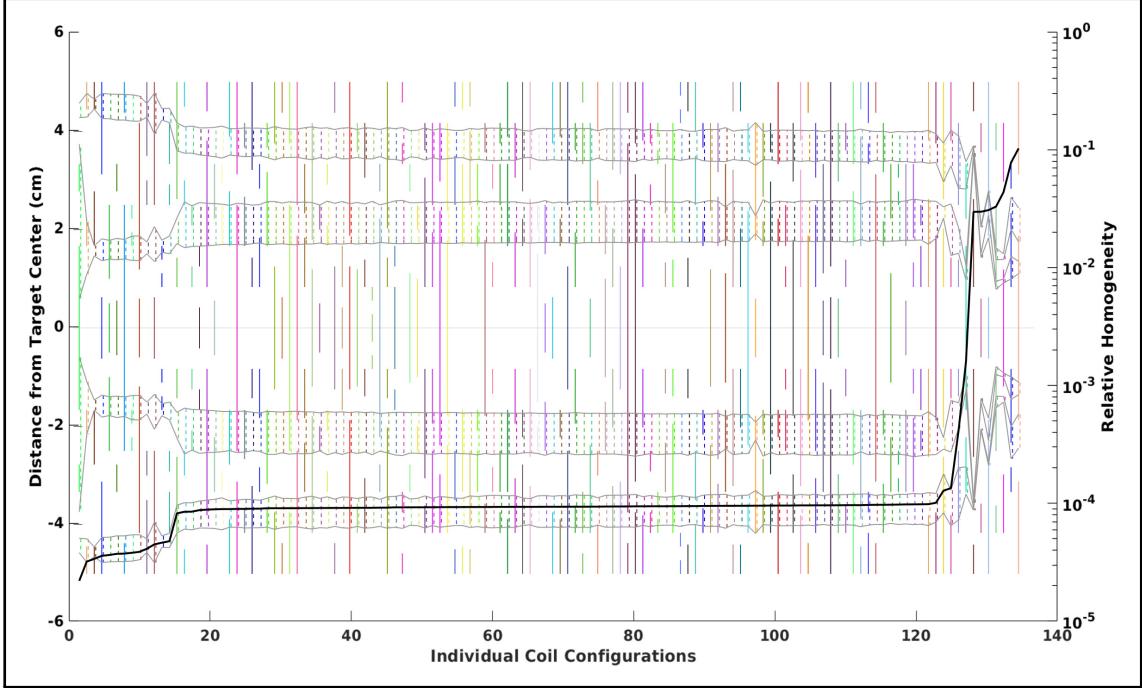


FIG. 18: The sample of results from a scan of initial conditions of shim coil geometry. Each minimization attempt is shown as a set of vertical lines with a shared color. The left vertical axis shows the the longitudinal length of the coils positions while the right vertical axis gives the homogeneity of each set (shown as a black line). The solid lines show the initial placement values given to the optimizer while the dotted lines show the output from that attempt (banded in light grey). Attempts are arranged along the horizontal axis from best to worst. The majority of initial values produce an extremely similar set of final coils which represent a large, easily found local minimum. However, a limited number of initial values manage to lead the solver to a significantly different coil geometry which preforms much better and may represent a small, difficult to find global minimum. Notably, the initial and final values appear completely uncoupled from one another. However, it is important to note that coil geometry constitutes just four of the six total degrees of freedom. The other two (coil currents) are not shown in this plot.

With all of that in place, the next parameter was the choice of ideal field. The most basic selection is simply a perfectly homogeneous field at the peak value of the background field. However, since a four-coil arrangement allows for movement of the shimmed field

above or below the background level, centering the background field half way between the two polarization directions offers the highest functionality for polarization. There is also an extreme case of potentially placing two target cells adjacent to each other and manipulating the field such that one cell sits in a region of positive polarization while the other sits in a region of negative polarization - a so-called dual cell configuration. Although this option was not pursued in the end, a test of that utility was undertaken [72].

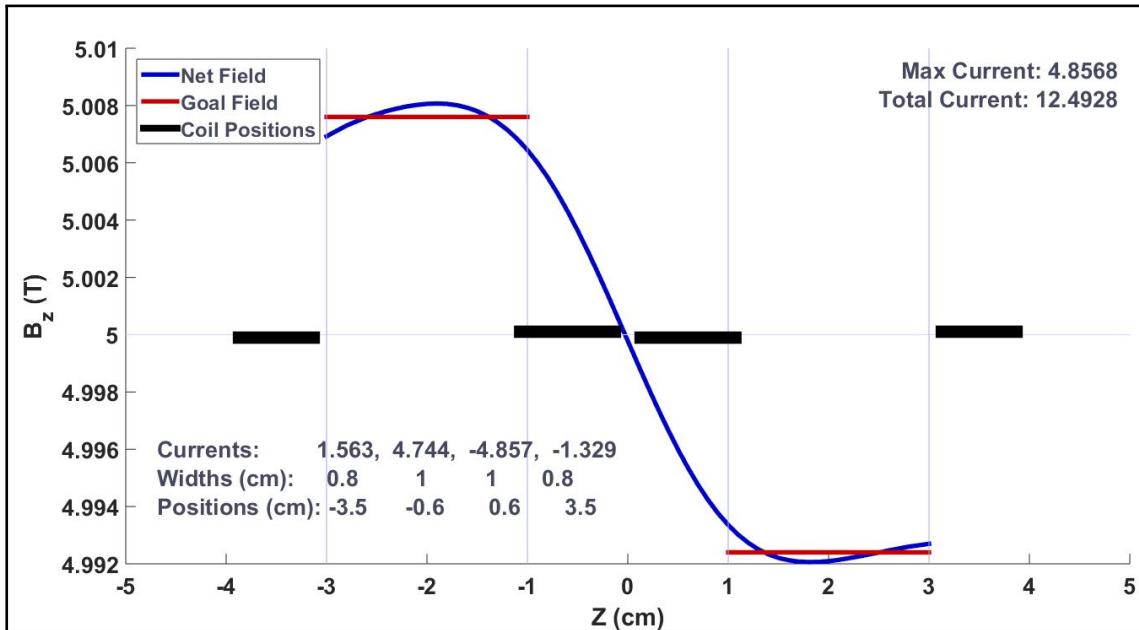


FIG. 19: Shim coils creating two adjacent regions of opposite polarization direction.

The final consideration is the χ^2 calculation. As is, the function optimizes homogeneity along the central axis for a single desired field exclusively. However, the function can be modified to include contributions from other sources as well, such as total current, maximum current, multiple goal fields, background asymmetries, or minimal total coil material. These calculations can be added in various weights to the minimization scheme to select a coil configuration that optimizes all aspects. With all of these considered, the final dimensions

were selected and finally the winding density was reduced as much as possible without removing functionality. The final configuration is given in Fig. 20. The shimmed field reaches a homogeneity of $\frac{\Delta B}{B} = 3 \times 10^{-6}$ (over 14 cm) and a maximum field displacement of 1% at 5 T.

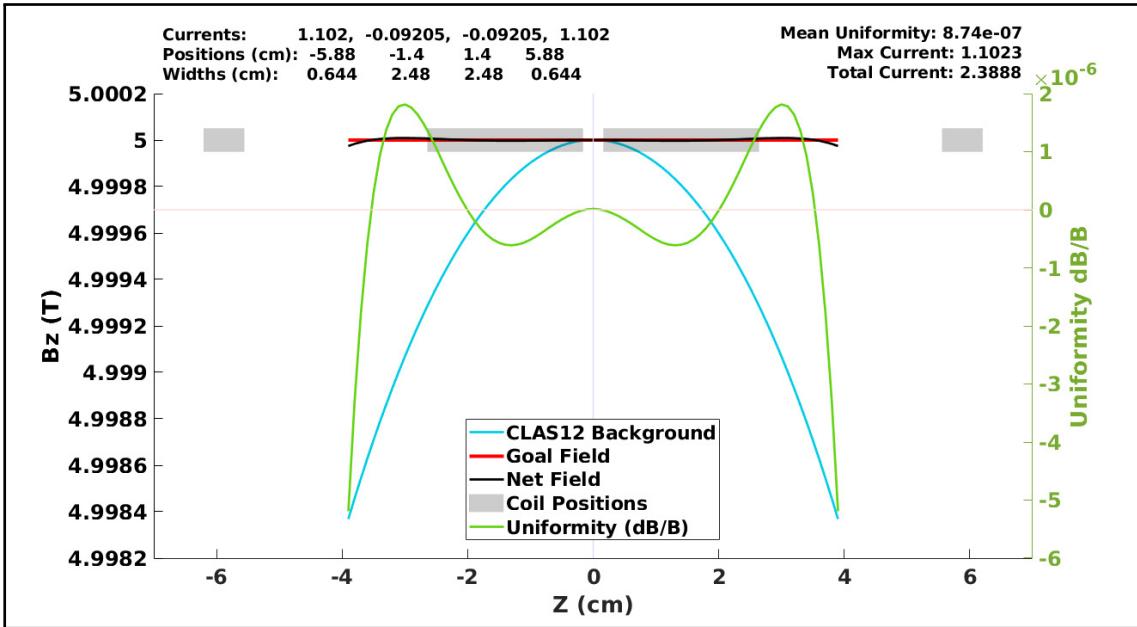


FIG. 20: Field produced by the final shim coil arrangement.

RGC Target

The CLAS12 polarized target for RGC (APOLLO - Ammonia POLarized LOngitudinally) is an evaporation refrigerated, dynamically polarized, longitudinal, solid ammonia nuclear target. The cooling is provided by liquid helium which is super-cooled via forced evaporation. High powered pumps drop the vapor pressure above the helium bath to reduce its temperature from around 4K to below 1K. Layers of shielding protect the pool from external heat. An outer vacuum chamber (OVC) insulates the shields while the evacuated gas from the cooling process flows over heat exchangers to remove stray external heat.

APOLLO sits along the center of the CLAS12 detector array with a radius of only a few cm at the tip, but a length of several meters. Within the refrigerator, target material (contained in modular cells) must be transported to the center of the detector array. This is achieved via a trolley mechanism that moves not only the cells, but the monitoring sensors for those cells as well. Also located at the tip of the target (center of detector) are NMR sensors and precision magnet correction controls (shims coils).

At the other end of the target are: an access port, electronics feedthroughs, exhausts, and valve controls. Thermometry and pressure sensors are located throughout the target as well as flow controllers and meters. There are also several strategically placed ohmic heaters in the cryostat as well to assist in homeostasis and rapid heat-up. All readings from the sensors are logged into a long-term data storage system (EPICS) for future analysis and diagnostics.

In following sections I will describe all of these systems in detail.

4.2 CRYOSTAT

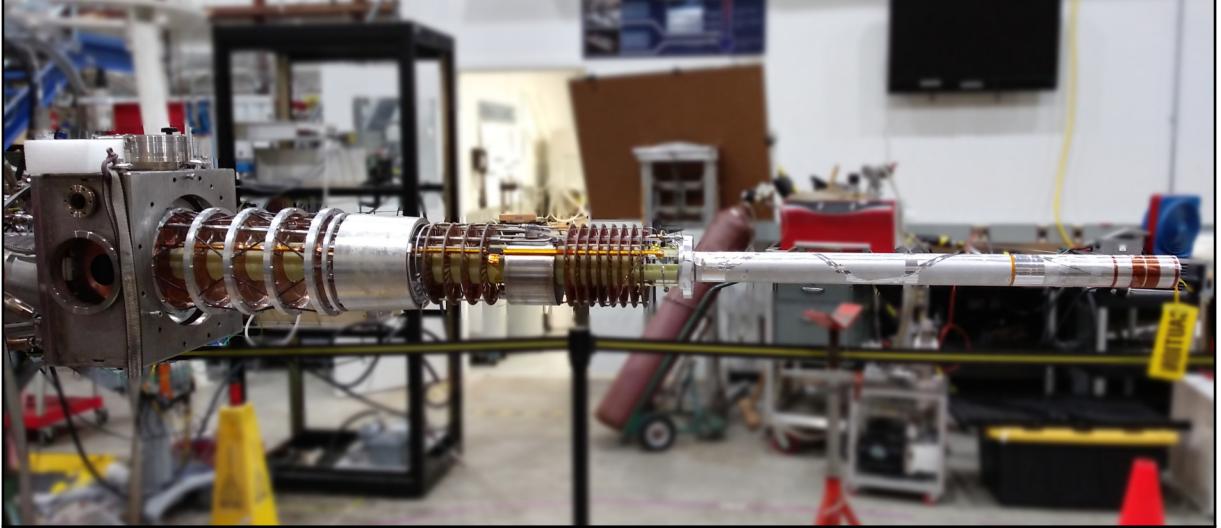


FIG. 21: The APOLLO cryostat during construction in the EEL building of Jefferson Lab.

The heart of the APOLLO target is its helium evaporation refrigerator. This cryostat serves to reduce the temperature of the target samples to the 1 K required for optimal polarization via DNP (see Section 3.1) as well as to cool the superconducting magnetic correction coils. The system produces a stable, thermostatic reservoir of liquid helium-4 in a chamber called the separator. This helium is then transferred to the far end of the target (called the bath) where it forms the pool in which the target cells are immersed. Gas from this pool is then pumped back over the refrigerating system, cooling components as it flows. The reduced pressure needed to lower the helium from its boiling point of 4.2 K (at 1 atm) down to 1 K is provided by strong vacuum pumps, maintained by efficient insulation, and monitored by a complex array of sensors.

Helium Path

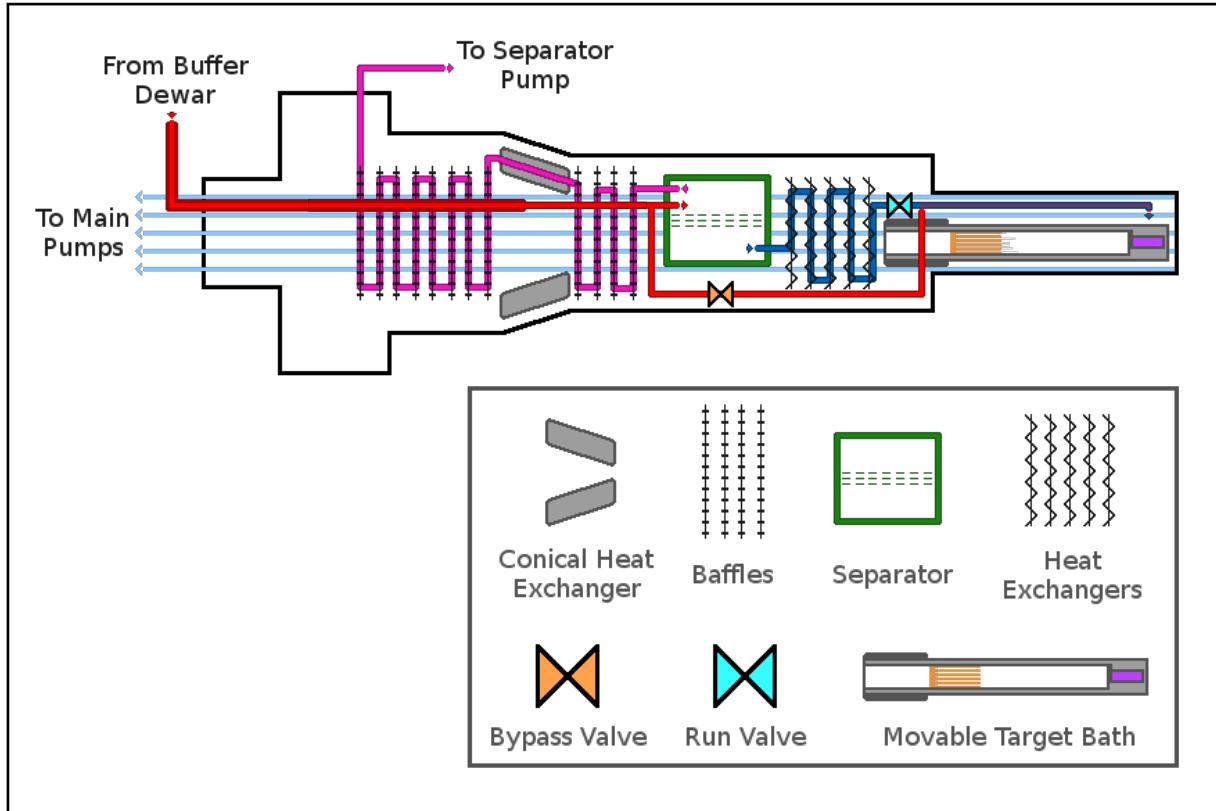


FIG. 22: A simplified diagram of the helium paths of the APOLLO cryogenic system. Liquid and gaseous helium is delivered from a buffer dewar via an insulated transfer line (red) to a phase separating chamber (green) or directly to the target bath via a bypass route (also red). From the separator, helium gas is removed by the separator pump - cooling several components along the way (magenta). Liquid helium from the separator is delivered to the target bath via the run valve route (blue). As the helium at the bath evaporates, the gas is removed by the main roots blower vacuum pumps (cyan) crossing the entire cryostat volume to provide cooling to heat exchangers and other components.

The liquid helium for APOLLO is supplied by the Jefferson Lab Cryogenics Department via its End Station Refrigerator (ESR). This on-site refrigeration plant recirculates and condenses the helium in a closed loop capable of dissipating collective heat loads of up to 4 kW. To ensure the most stable delivery, the helium is first dispensed into a buffer dewar immediately adjacent to the target. This dewar holds hundreds of liters of liquid helium which is capable of sustaining target cryogenics for many hours in the event of any supply disruption. Helium flows from the dewar through a custom vacuum insulated transfer line to the center of the target and is sprayed directly into a chamber called the separator. Along the way some of the helium is vaporized, however, the transfer line is specifically designed to reduce turbulence in the helium flow which can introduce periodic vapor lock and cause rhythmic fluctuations in the pressure (and therefore temperature) of the cryostat.

The lower section of the separator is divided from the top by a porous sheet through which the incoming 4K liquid He settles. The gas is meanwhile purged from the top, reducing the temperature and allowing for a highly stable and controlled delivery of liquid helium to the target bath. The gas expelled from the separator, passes through a series of baffles as well as the conical heat-exchanger (both discussed later) before returning to the ESR for re-condensing. On the way, it passes through a flow controller which monitors and regulates the separator flow.

From the bottom of the separator, the liquid helium flows through a transfer pipe which weaves through a system of heat exchangers - a series of consecutive, stamped, copper plates onto which the helium pipe is sintered. These heat exchangers serve the very important purpose of further cooling the liquid helium before it reaches the run valve. The run valve is a JT (Joule Thomson) needle valve which serves the dual purpose of regulating the rate of helium delivery to the final target region as well as contributing to the cooling via the Joule Thomson effect. Though more commonly seen in uses strictly involving gases, the Joule Thomson effect describes the temperature change experienced during an adiabatic expansion from a higher pressure region to a lower one. While the pressure in the separator is maintained around 1.5 atm (moderated by the flow controller) the volume past the valve is kept at a very low pressure ($\sim 10^{-4}$ atm), causing the liquid helium to expand and partially vaporize across the valve - greatly reducing in temperature as it does so.³

Past the run valve, the helium exits into the pooled liquid helium reservoir (called the target bath) which extends several feet back towards the separator. This bath consists of a

³Although helium typically experiences heating in such an expansion (given its negative Joule Thomson coefficient), below its inversion temperature of 40k the coefficient flips sign and cooling occurs instead.

two long, horizontal, concentric cylinders (described in section 4.3). Helium enters the outer cylinder via an opening on the top (just over the target cell which rests within it) while the inner cylinder is sealed to prevent helium entry (creating a void in the electron path to minimize the amount of material encountered by the electron beam before it reaches the target cell). The liquid helium occupies the space between these cylinders and its liquid level is measured by a capacitive level probe which also occupies that same space (discussed later). The relatively large size of this pool helps ensure stability in the liquid level surrounding the target cells.

Alternatively, the helium can bypass the separator (and run valve path) altogether (skipping the stabilizing and pre-cooling process) and be delivered directly to the bath at 4 K. This is typically only utilized in the initial stages of operation when the cryostat is well above operating temperature and benefits from the increased helium flow. This secondary path is controlled by the bypass valve. Both the run and bypass valves are controlled using a stepper motor and digital readback system (discussed later) which allows for extremely precise remote adjustments. Together with the flow controller, these three devices comprise the primary driving variables of the cryostat.

The helium pool itself is continuously pumped on by a paired set of high-powered Roots-blowers vacuum pumps which reduce the pressure at the bath to that needed for sub-Kelvin temperatures. In addition to providing the pressure drop needed for the Joule Thompson effect, they provide cooling in the form of reduced vapor pressure. The relationship between vapor pressure (P) and temperature (T) can be approximated by the Clausius-Clapeyron equation:

$$\ln\left(\frac{P_0}{P}\right) = \frac{L}{R}\left(\frac{1}{T_0} - \frac{1}{T}\right), \quad (57)$$

where P_0 and T_0 are a known vapor pressure and temperature (typically at 1 atm), L is the latent heat of vaporization, and R is the gas constant. However, for cryogenic practices this equation is not quite accurate and instead empirical look up tables are employed (see Fig. 23).

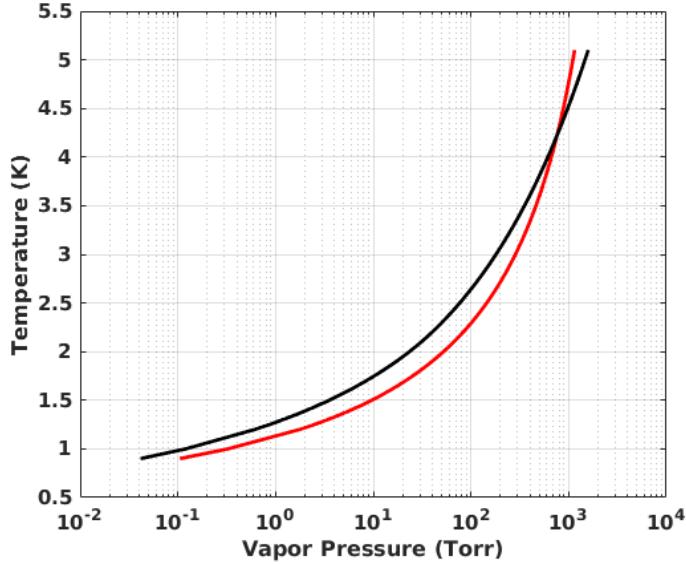


FIG. 23: Plot of vapor pressure vs. temperature for Helium-4. The red line represents the Clausius-Clapeyron approximation while the black line is interpolated from ITS-90 data [73].

So, to achieve the desired 1K temperature, a vapor pressure near 0.1 Torr ($\sim 10^{-4}$ atm) must not only be achieved, but maintained under heat load. During the experiment, we expect an incident heat load of 0.1 W from the 10 MeV energy loss of each electron at a 10 nA current and around 0.15 W from the microwaves. The latent heat of vaporization for helium is around 80 J/mol at 1 K which means 0.2 W of heat will vaporize about 2.5 mmol/s or 1300 m³/hr. The paired Roots-blowers are capable of purging 5700 m³/hr which is more than sufficient - provided the ambient heat load from all other sources does not exceed much more than half a Watt.

As the super-cooled helium vapor is removed from the bath reservoir, it flows back over the run-valve path heat exchangers - cooling the liquid that is on its way to the JT valve. Since a lower initial liquid temperature reduces the amount of gas that is produced during the JT expansion, this process is crucial to minimizing the incidental vapor load on the pumps and maximizing their cooling efficiency.

Exiting the cryostat on its way to the pumps, the gas first passes through a gate valve and flow sensor before also returning to the ESR for processing. As next discussed, the extremely

low operating temperatures of the cryostat demand diligent minimization of ambient heat loads. The cold vapor flow from the bath serves to remove incident external heat by contact with heat shields via the conical heat exchanger. It also flows across a series of flow-directing sheets, called baffles, which help to isolate and remove heat introduced at the back plate where the cryostat attaches to the rest of the target assembly (see Fig. 24).

Insulation

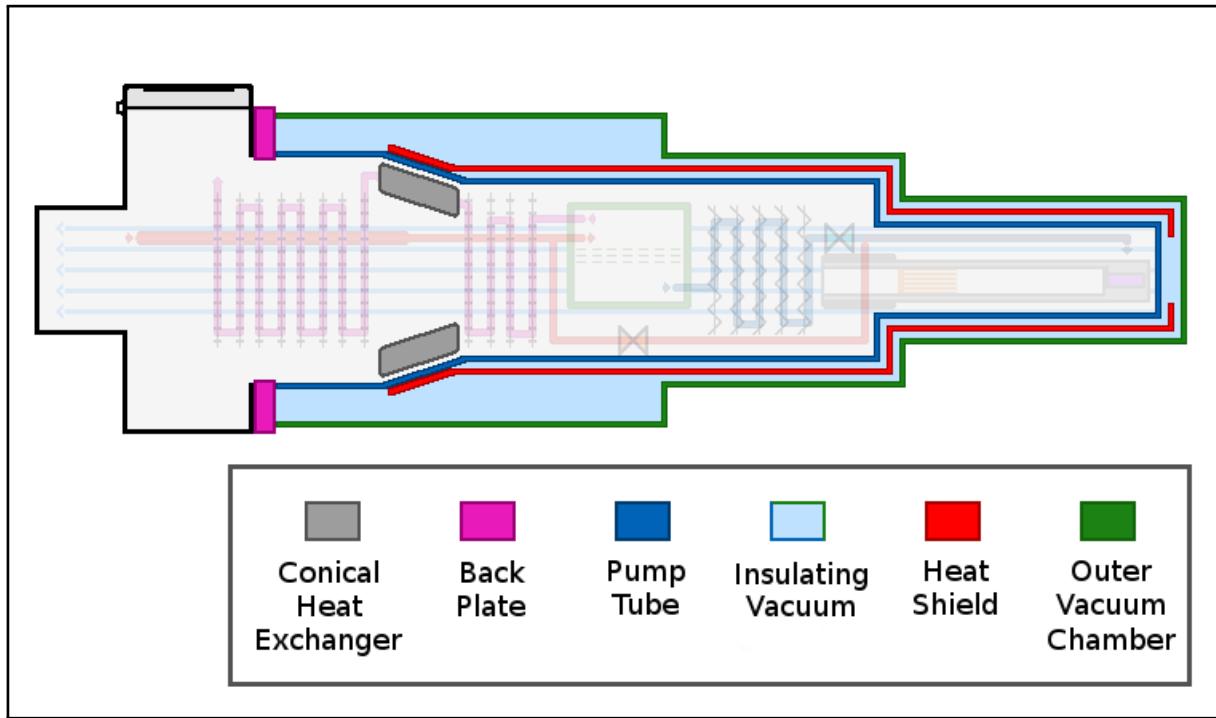


FIG. 24: A simplified diagram of the APOLLO heat shielding. An evacuated volume (Insulating Vacuum - cyan) is defined by its exterior wall (Outer Vacuum Chamber or OVC - green) and interior wall (Pump Tube - blue). Within this volume is a copper heat shield (red) as well as layers of Mylar wrap. Both walls attach to the cryostat at the Back Plate (magenta). Within the cryostat (interior to the pump tube) there is a Conical Heat Exchanger (grey) which transfers heat from the shielding to the exiting helium vapor.

The primary defense against direct environmental heat is the outer vacuum chamber or OVC (see Figs. 24 and 25). The OVC contains an evacuated volume which surrounds the target from the bath to the back plate. Its body is made of cylindrical aluminium with three successively narrower sections to accommodate the width restrictions of the CLAS12 detector. The terminal end of the OVC - which covers the bath section where the target material sits - is made of a thin, carbon fiber cylinder with a domed end. Carbon fiber reduces the radiation length of this portion while still maintaining a high degree of structural integrity. However, as carbon fiber is somewhat permeable under high vacuum, the section is also lined with a very thin aluminum foil. Also, as the tip of this domed section intersects the electron path, it has a thin aluminum window (without carbon fiber) through which the electron beam exits after passing through the target cells. This window is thin enough to sufficiently minimize events from electron interactions while still being resilient to the radiation damage from the beam.

The volume between these inner and outer walls is evacuated (via a turbo-pump) to a pressure of at least 10^{-5} Torr to prevent thermal heat conduction to the interior of the target. Within this chamber is a copper heat shield which protects against external radiative heat. This shield has an aluminum base which rests in contact with the inner wall of OVC. As this layer does not contribute to the structural integrity of the target or require a vacuum seal, it is free to be open-ended at the nose (not needing a solid exit window). Additionally, to further minimize the radiative thermal load, the vacuum space is wrapped in many layers of super insulation which are alternating layers of Mylar and low-conductive mesh.

The inner wall of the evacuated volume (called the pump tube) defines the boundary between the insulating vacuum and the complete helium volume of the cryostat. The body of the pump tube is made of cylindrical sections of G10 (a high-pressure fiberglass laminate) connected by transitional aluminum segments. While these sections are far enough away from the target cells to not influence the total radiation length, the G10 material is still useful in reducing the overall weight of the refrigerator. The pump tube terminates in a thin, domed aluminum nose piece. This piece is similar to the thin window of the vacuum can but, as the pressure differential is reversed compared to the can, it does not need to be as robust. This means the entire piece can be made from extremely thin aluminum (without any fiberglass), but also means the pressure within the helium volume cannot be allowed to drop below that of the OVC.

Notable on the pump tube is the aluminum section (on which the heat shield rests) corresponding to the conical heat exchanger. The conical heat exchanger is located on the

cryostat along the helium gas exit path. It serves as a heat sink for the pump tube and heat shield, utilizing the spent helium vapor from the bath as well as the extracted gas from the separator to further minimize the effects of black body radiation from those surfaces.

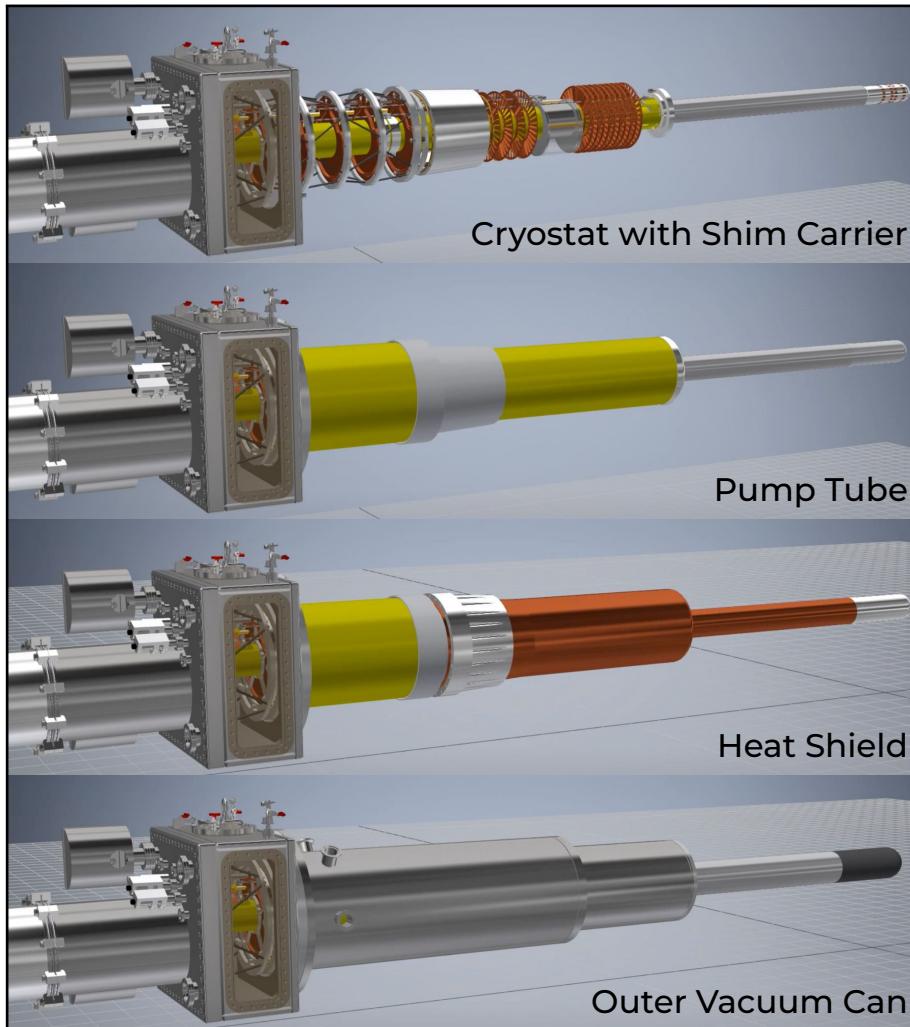


FIG. 25: Digital renderings of successive layers of the APOLLO target's insulation system. In these models, components constructed from G10 are colored yellow, copper components are copper colored, black indicates carbon fiber, and various grades of aluminium are rendered in silver or grey.

Monitoring Sensors and Controls

Both the OVC and pump tube (as well as the refrigerator itself) connect at the back plate which is the large aluminum box structure that enables access to the target material during target swaps (see Section 4.3). As this structure is the terminus of the insulated region but still part of the exit path of the helium gas, it exchanges a significant amount of heat from the environment to the exiting vapor. For high enough flows, such as those during experimental production, the back plate can actually become overly cool causing water to condense on its exterior. As this poses a small risk to nearby electronics, the back plate is equipped with an ohmic heater and thermal probe which are automated to maintain a thermal homeostasis.

There is also a small heater in the target bath which allows for rapid heating of the nose section when needed (e.g.; for target swaps) and to simulate heat loads for testing. Additionally, there are many temperature and pressure sensors throughout the cryostat. The readouts from these probes, along with that of every control, are sent through the site-wide monitoring and control system known as EPICS (Experimental Physics and Industrial Control System). This system allows for remote control of the target over a secure network as well as logging a complete history of the target's operation. It also allows for integration of the target into the existing accelerator controls enabling automatic accelerator fail-safes to be triggered by the target status.

Another functionality of the EPICS system is the ability to automatically adjust any incorporated *control* variable based on the reading from any incorporated *probe* variable. For example, the back plate thermal probe could be used to control the current to the back plate heater in order to maintain a user designated temperature (though, in practice, this particular function is accomplished with hardware). The process of adjusting a control variable to maintain a value on a monitoring variable is achieved via a PID loop which is a well-established, standard feedback algorithm for controls systems. For the APOLLO target, the PID loop functionality is primarily used to control the run and bypass valves.

During initial cool down, which takes many hours, the PID monitors the helium flow to balance cooling rate and helium consumption. When the fridge is near operating temperature, but not in active use, the PID is used to maintain a set temperature in the cryostat, again to conserve helium usage while still allowing for expedited target readiness once needed. During operation, however, neither the helium flow nor temperature probe reading are sufficient to precisely maintain the liquid helium levels and ensure the target cells are immersed. To that end, a liquid helium level probe has been incorporated into the bath (see Fig. 26).

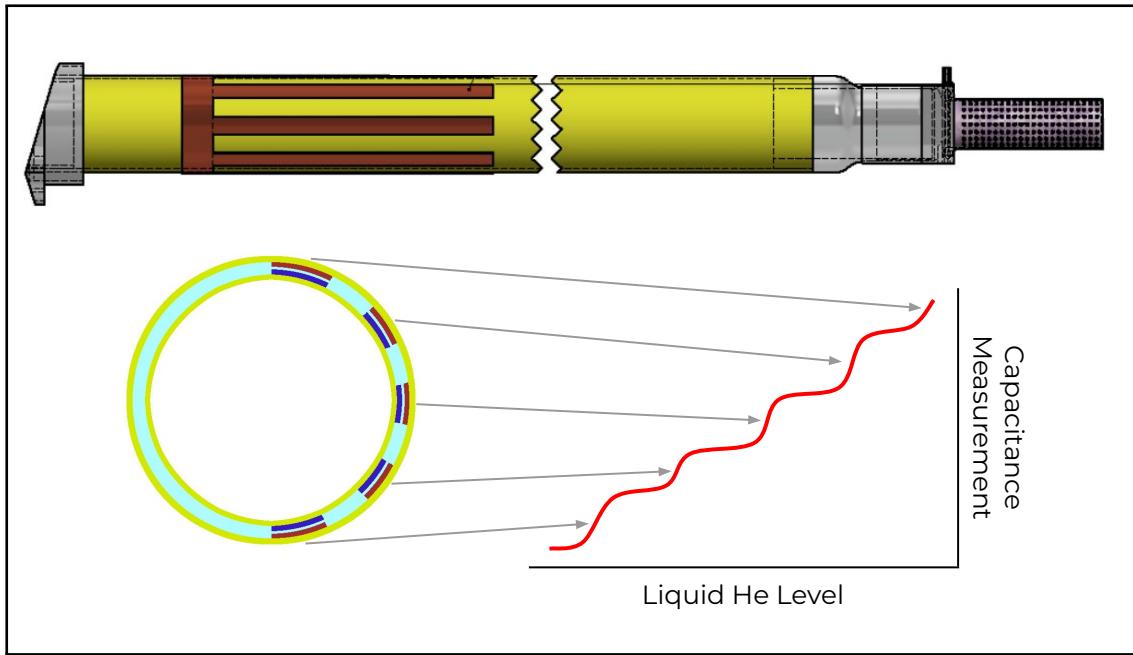


FIG. 26: Illustration of the capacitive level probe. The top figure shows a side-on view of the interior tube of the movable bath with the location of the capacitive strips in maroon. The lower left figure shows a face-on view of the strips as they are positioned within the two concentric cylinders (yellow) of the bath. As liquid helium fills the space (cyan) between the two parallel plates (maroon and blue), the capacitance increases as shown on the lower right plot.

This level probe is located in the liquid reservoir of the movable bath, between its concentric tubes (see Section 4.3). The probe consists of two thin copper foils each cut into five, 20 cm long strips that run lengthwise down the bath and are separated by thin Kapton spacers. Each foil is independently wired back to the electronics rack such that they form the two sides of a parallel plate capacitor. Any liquid helium that fills the gap between the foils acts as a dielectric - changing the total capacitance of the circuit. The foils are positioned such that each strip sits at a different elevation in the liquid reservoir. As the liquid level rises, it periodically encounters and fills successive strips. These discrete capacitive increases are monitored - with the fifth increase indicating a completely full bath. The PID loop then adjusts the run valve as needed to maintain a capacitance value between appropriate jumps.

4.3 INSERT ASSEMBLY, TROLLEY, AND BATH

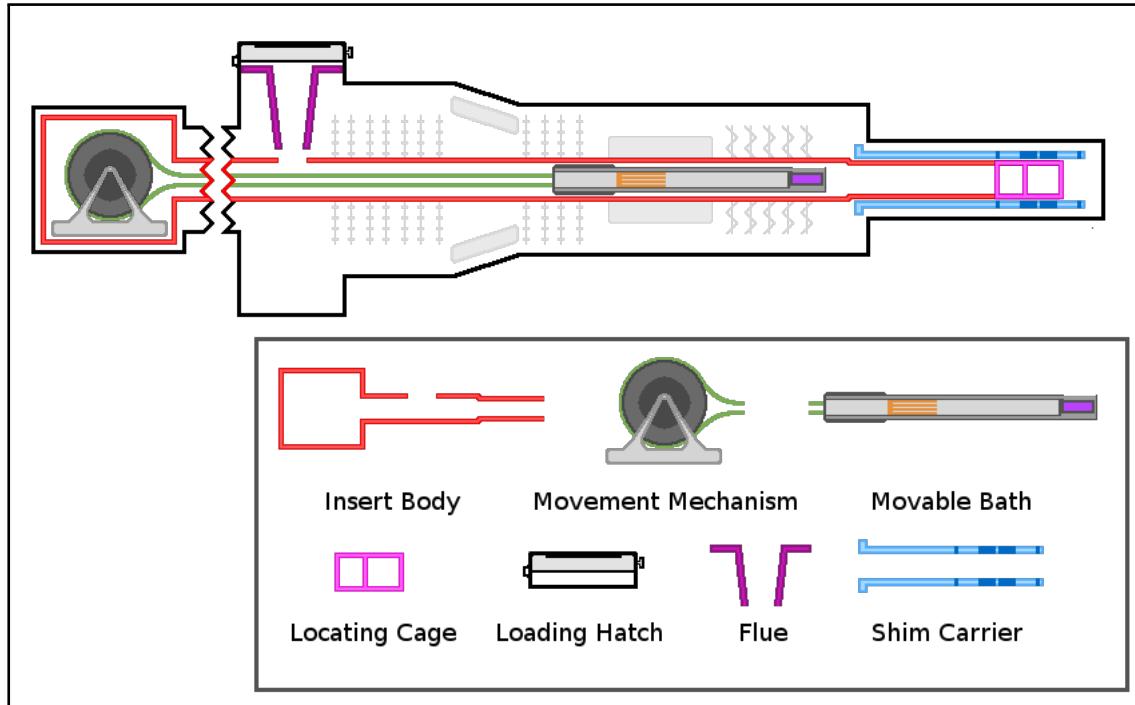


FIG. 27: A simplified diagram of the APOLLO insert assembly. The insert body runs through the center of the target. It includes a back chamber box (at the far upstream end of the target) which houses the movement mechanism spool for the movable target trolley. This mechanism transports the trolley from one end of the insert to the other via tethers (green). At a retracted position the trolley is positioned such that the target bath is directly under the loading hatch (adjacent to the back plate of the cryostat) - allowing access to the target cells (via a loading flue). At the inserted position the trolley is positioned such that the target bath fits firmly into the locating cage which is attached to the far downstream end of the insert body (interior to the shim carrier). In this diagram the trolley is half way between these two locations.

Within the central axis of the refrigerator sits the insert assembly. This assembly carries many of the necessary components of the target downstream to the target nose where the cell is located during data taking. To ensure access to its complicated systems, this assembly is independently removable from the rest of the refrigerator. It extends the full length of the target and contains: monitoring sensors, the NMR system, the microwave delivery system, and the movable bath trolley.



FIG. 28: The insert assembly (foreground) next to the target cryostat (background). The insert fits thought the center of the cryostat as shown in Fig. 27.

Trolley

The most notable feature of the insert is its trolley system. In previous polarized target experiments at Jefferson lab, target material had to be accessed by fully removing a cell-holding structure from within the cryostat. The APOLLO target, however, must sit

horizontally in-line with the electron beam. As such, any structure lengthy enough to reach the center of the CLAS12 detector would require disassembling a large portion of the beam-line to extract.⁴ Additionally, removing and replacing large components from a cryostat requires considerable time to cool back down. In the past, this was mitigated by placing several target cells into the target at a time which could be alternately moved into place as needed. However, once again the space restrictions of CLAS12 prevent such a system. The solution was a novel mechanism which transports individual cells horizontally within the cryostat between a positions at either end of the fridge.

In this system, a movable bath runs along a track within the insert assembly similar to a trolley car. The movement is driven by a pulley and paired set of tethers which pull the bath from the far downstream end of the insert where data is taken (within the detector array) to an upstream loading hatch (outside of the detectors) where samples can be easily accessed and exchanged. As this entire path is contained within the cryostat (and the only components being swapped are small target cells) the refrigerator is able to recover its temperature very quickly - on the order of an hour for the entire operation.

While some systems stay in place at the far end of the insert during this process, it is necessary for the trolley to carry several components with it. The entire liquid helium bath is contained in the trolley (since otherwise the system would need to be able to remove a cell *horizontally* out of the bath and yet keep the bath superfluid tight once the cell was replaced). Likewise, as the liquid helium level probe, bath heater, and bath thermal probes⁵ are all contained within the bath reservoir, they must also move with the trolley.

⁴As is the case for the insert itself. Although the insert is much more accessible than the target as a whole, removing it still represents a significant investment of time and effort.

⁵Additionally, the thermal probes need to move with the trolley in order to monitor the temperature of target materials during cell exchanges.



FIG. 29: The movable bath trolley which runs within the insert to transport target cells in and out of position within the detector. The left-hand side of the image is the upstream end of the trolley terminating with the trolley base (see Fig. 33), which attaches to the movement tethers and houses the ball-bearings by which the trolley glides. The right-hand side of the image is the downstream end of the trolley terminating with the Kel-F bath (see Fig. 34), which holds the target cells. Connecting the two ends is the trolley body made up of two concentric cylinders containing the liquid level probes and thermal sensors. Liquid helium fills the space between the cylinders as well as the Kel-F bath.

Movement Mechanism and Tether

The trolley is pulled by two tethers. The *retraction* tether attaches from a spool located in the insert box at the far upstream end of the insert assembly (see Fig. 30). From the spool, it passes via pulley into the bottom of the insert pipe and attaches directly to the upstream base of the movable trolley (see Fig. 33). As the spool rotates counter-clockwise, this tether retracts and drags the trolley upstream towards the spool.

The *insertion* tether is counter-wound on the same spool and runs fully to the far downstream end of the insert tube (near the cage) where it passes through a second pulley before running back upstream to attach to the trolley. When the spool is wound clockwise, this tether retracts, which drags the trolley downstream towards the cage. This system allows both forward and reverse movement of the trolley to be coupled and driven by a single spool.

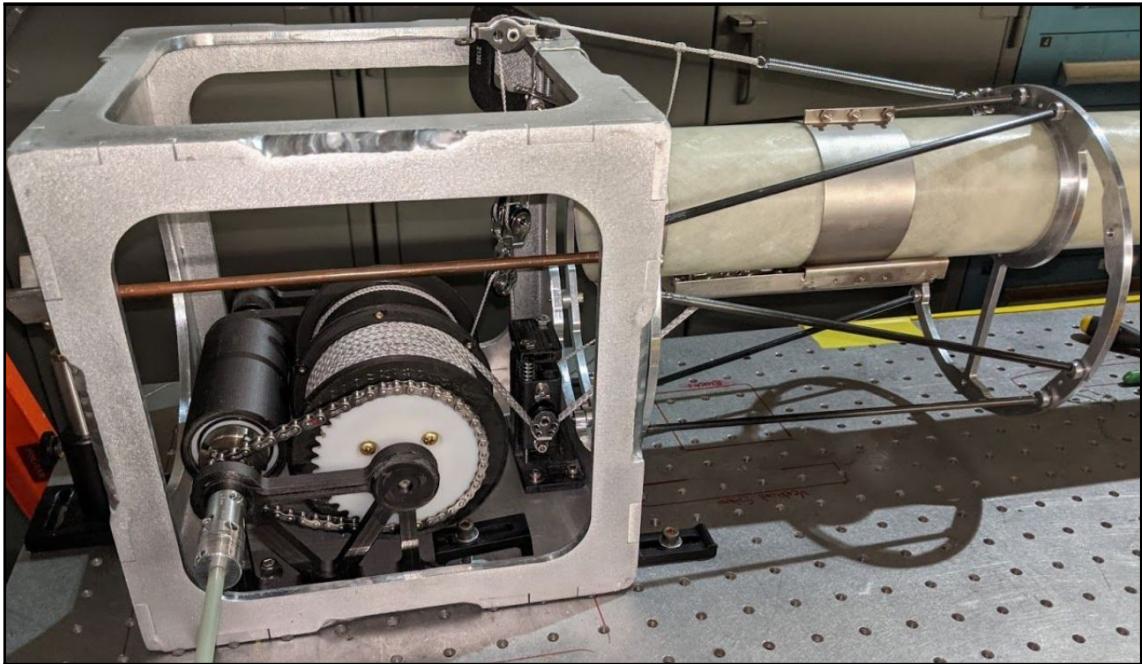


FIG. 30: Movement mechanism located within the upstream box of the insert assembly. This spool drives the trolley via a paired tethers and a tensioning pulleys.

The 3D printed spool and spool frame are mounted in the back box of the insert which in turn fits into the the upstream-most end of the overall target body. Instead of a simple axle, the spool rotates on a threaded rod - the pitch of which matches the diameter of the insertion tether such that, while rotating, the spool shifts transverse position at the exact rate necessary to lay the retraction tether into a perfectly neat single layer. This prevents the risk of entanglement which might freeze the trolley movement and endanger the cryogenic safety of the target material. The movement of the spool is driven by a rotating feedthrough on the side of the insert box which connects to a hand wheel outside of the target body by a threaded rod and articulated junction. This connection also includes a torque-limiting slip clutch to prevent damage from over-extension of the spool.

The retraction tether also serves the purpose of carrying the necessary electrical wiring to the trolley. This consists of two cernox thermal probes (one near the top of the bath reservoir and one near the bottom), a power line for the ohmic heater inside of the bath, and the coaxial wires for the capacitive level probe. As each cernox has both voltage and current

taps, this makes for 12 separate leads. These travel through the center of the retraction tether and pass with it into the spool body.

There, the wires exit from the tether and connect to a breakout circuit board which, in turn, connects to a central slip ring. This ring (or electrical rotary joint) accepts signals from its rotating outer shell and transmits them to a stationary set of leads at its axis.⁶ From there, the leads connect to electrical feedthroughs which enable the signals to exit the target. This system allows for continuous communication with the trolley electronics - most importantly the cernoxes that monitor the temperature of the target material during transit to and from the cryostat.

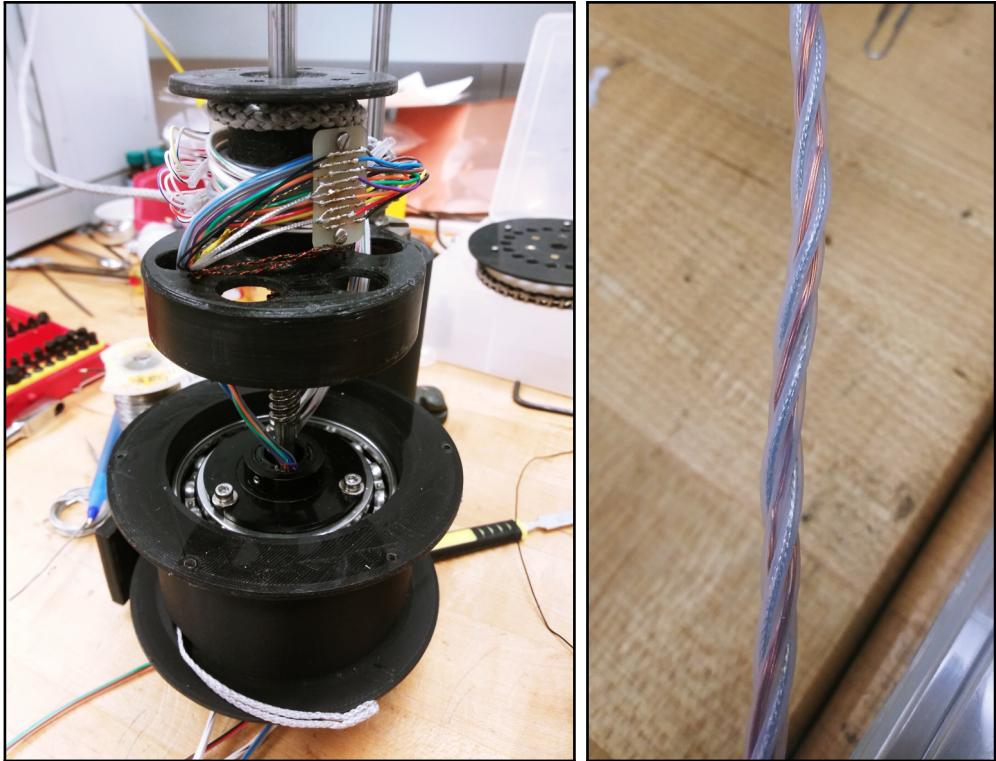


FIG. 31: Left: Internal wiring of the movement spool where the retraction tether connects to the slip ring. Right: Final configuration of retraction tether internal wiring. This braid of wires is covered by a nylon sheath before assembly.

⁶In other words, it allows one end of the wires to rotate around the spool without twisting the other end.

These circuits must not only make consistent connection with the trolley but also reliably maintain it after numerous transit cycles during which the conductors transition rapidly from the cryogenic temperature of the cryostat to the room temperature of insert box while *also* being under tension and bending around the circumference of the spool. Therefore, the composition and structure of wires and tether body needed to be carefully selected and rigorously tested.

The heater, when activated, is intended to introduce heat to the bath, however, it is not desirable for it to conduct heat there at other times. Likewise, the two coaxial cables of the liquid level probe must be able to transmit minute changes in capacitance without introducing excessive noise while also being small enough to minimize both heat conduction and tether radius. The eight leads for the cernoxes, too, need to minimize heat load without risking fragility. Also, as these wires wrap around the spool, they need to be simultaneously loose enough that they do not experience material strain and taut enough that they do not entangle with each other (and have enough force to move the trolley).

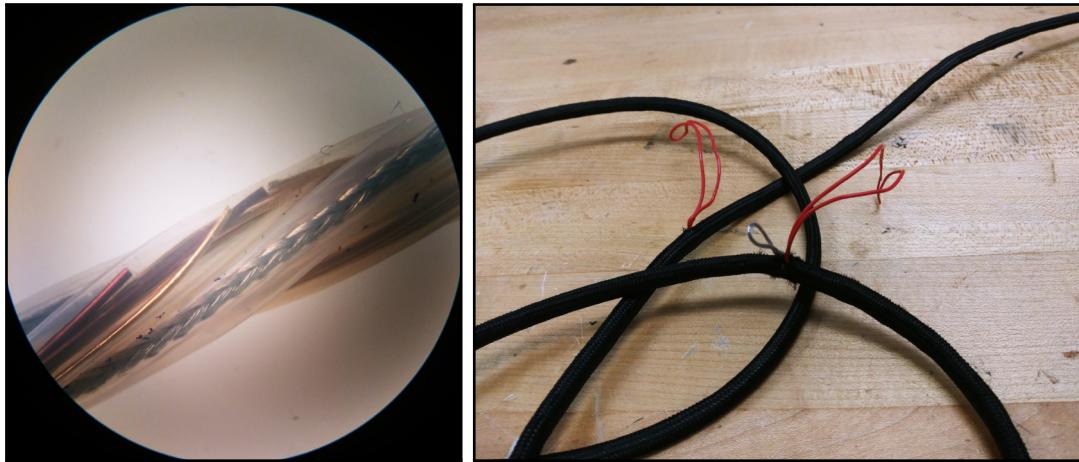


FIG. 32: Examples of failed tether configurations. Left: Microscope view of broken wires within the protective tubing. Right: Break-through of wires from within the exterior polyester wrap.

The solution was to construct the tether out of four twisted PTFE tubes. Two of the tubes contain the coaxial leads for the level probe, one contains the four leads for the bottom cernox, and the last contains the four leads for the top cernox as well as both leads for the heater. The twist ensures that the average travel length of all tubes are approximately equal while spooling. This prevents relative motion between tubes and discourages tangles. As the electrical wires run freely within the tubes, the external strain is minimally transferred to the more delicate electrical conductors. Finally, an external sheath wraps around the four tubes and carries the majority of the tension for trolley movement. This configuration was determined after extensive prototyping and testing.

Trolley Testing

During operation, target cells will need to be exchanged every 2-3 days on average. This means that, over the course of the scheduled experiments in run group C, the trolley is expected to travel to and from the cage around 100 times before any significant target maintenance might allow for convenient replacement of the tether. With a safety factor included, we required any tether configuration to survive at least 1000 movement cycles without failure. As each round-trip takes nearly a minute to complete, manually testing each tether prototype for the necessary number of cycles would take an excessive amount of time so, instead, an automatic testing rig was created.

This tabletop rig connected the spool to a stepper motor which was, in turn, digitally connected to two physical switches attached to the insert. As the trolley reached either end of its travel, it made contact with one of the switches which signaled the motor to reverse direction and logged a completed journey. Meanwhile, electrical continuity through the tether was continuously monitored. (For testing purposes, the leads to the capacitive level probe were shorted to one another at the trolley to allow for a simple circuit.) These signals were fed into a custom-made LabView program which processed them, controlled the motion (via a National Instruments DAQ box), and generated a text log of the tether's status and performance. This setup could be left running autonomously overnight and allowed multiple tether configurations to be tested with ease. The final configuration achieved 1500 cycles before succumbing to a known material flaw that was introduced during assembly and removed from the final build.

Another potential source of failure is in the trolley's movement itself. The trolley rolls down the insert tube on plastic bearings contained within two mirrored loops on either side of the trolley base (see Fig. 33). These bearings make contact with the insert tube's interior

wall and allow the trolley to glide with minimal friction. To prevent the trolley from rotating during transit, the trolley base runs over the microwave waveguides and locator blocks as well as special rails that run the length of the insert pipe (also Fig. 33). These correspond to cutout grooves on either end of the trolley base and serve as steering tracks. Additionally, every structure within the insert pipe is beveled to remove any possibility of collisions that may interrupt the trolley movement. Finally, to help direct helium gas flow, the back of the trolley base is fitted with a thin kapton film wiper which gently contacts the interior of the pipe to form a loose seal. The multiple tests conducted on the retraction tethers also served to thoroughly test the overall trolley movement - ensuring a high degree of confidence in this novel mechanism.

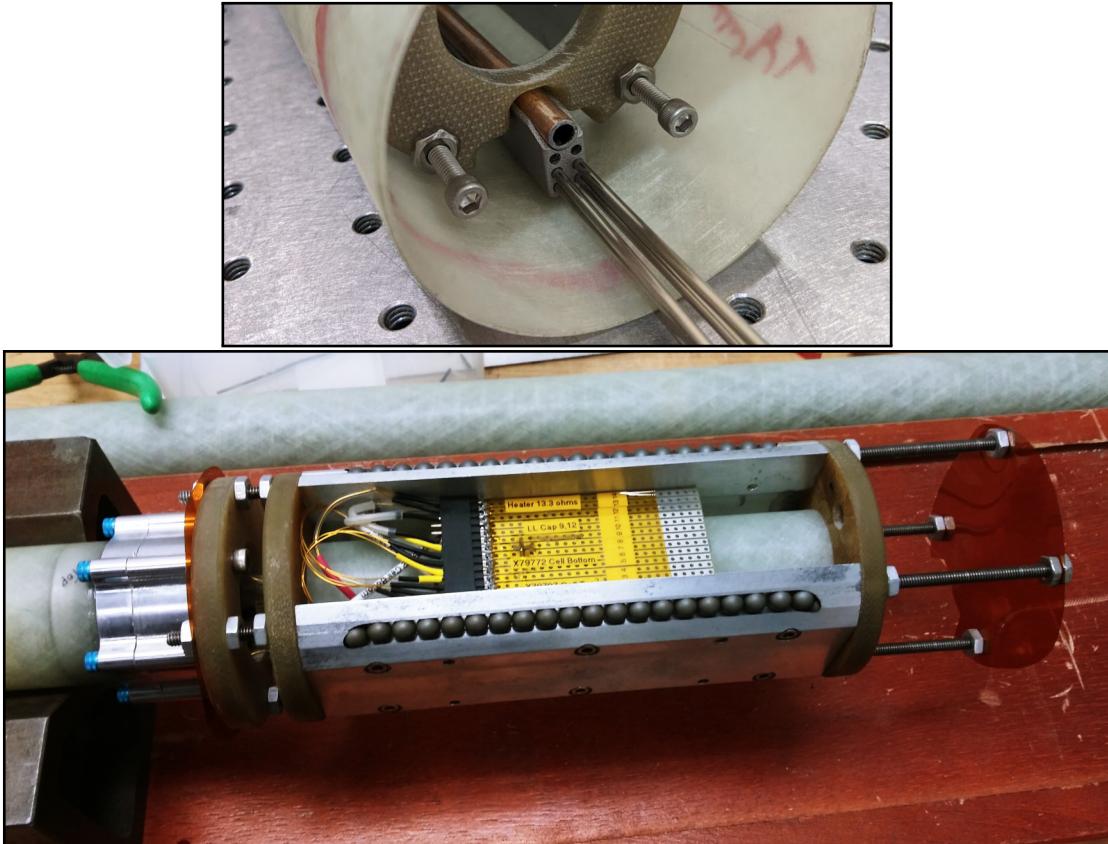


FIG. 33: Top: The trolley base above the microwave waveguide locating block. Bottom: The upstream base of the trolley showing its ball bearing wheels, connection points for the tether wiring, and back kapton film wiper.

The back of the base is where the tethers connect to the trolley. The base's hollow interior is also where the communications wires exit the polyester sheath of the retraction tether and connect to a breakout circuit board. From there, the circuits enter the tube of the bath reservoir that extends out from the trolley base.

Bath

As already described, this is a compound reservoir with two concentric tubes - an inner one which the level probe wraps around (see Fig. 26) and an outer one which protects the level probe and connects to the final bath cell reservoir (see Fig. 34). The space between these tubes defines the liquid helium volume is also where the thermal probes are located.

Although the distance between the probes and the target cells introduces some offset in the temperatures measured while the bath is empty, when supercooled liquid helium is present its superfluid thermal properties ensure the probes are representative of the cell temperature.⁷ Importantly, since the bath reservoir extends fully to the trolley base, the vertical offset of the probes allows for verification of the presence of liquid helium - as first one, then the other, detect a temperature change as the bath fills. This helps in the initial calibration of the level probe.

⁷This is actually only true in the absence of a polarizing field - cernox thermal probes have a known property of slight (~10%) offset when in strong magnetic fields. As such, bath temperature is determined via vapor pressure during data taking.



FIG. 34: The downstream end of the trolley - the target cell helium reservoir (with an empty target cell in place). The oval cutout on the right is the window through which liquid helium drops into the bath and also reveals the smaller, interior bath tube around which the level probe (not pictured) wraps. This inner tube terminates at the target cell holder where a thin aluminum window prevents liquid helium intrusion. The dark stains on the bath are a result of the etching process needed for effective bonding with cryogenic epoxy.

Attached to this compound tube and terminating the trolley unit, is the bath proper (see Fig. 34). This bath is carved out of Kel-F which is hardy against both radiation and cryogenic temperature cycling. Its thin walls hold liquid helium around the target cells while minimally impeding scattered particle. Meanwhile its thicker base (upstream of the cell) attaches securely to the outer G10 bath pipe. Running through this base and extending into the Kel-F bath, is the inner G10 bath pipe which terminates with a thin aluminum window as well as the aluminum slot which holds the target cells.

Insert Cage

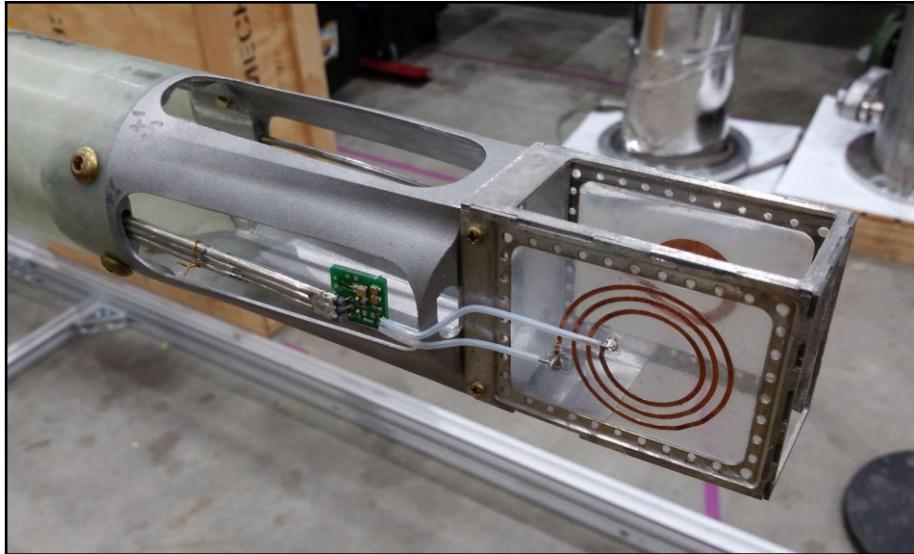


FIG. 35: The locating cage which hold the target cell bath in place next to the NMR coils. The near side of the cage (beam-right) has a spiral deuteron NMR coil and the far side (beam-left) has a simpler, circular proton coil. Also visible is the deuteron cold tank circuit board (just upstream of the deuteron coil).

When returning from the retracted position, the Kel-F bath is guided into the locator cage at the end of the insert pipe (see Fig. 35) by the ramped entrance of the cage. Once in place, cells sit directly above the microwave horn, between the NMR coils, and beneath (though slightly forward of) the liquid helium outlet. The locating cage itself is made out of thin, 3D printed aluminium. The distal half forms an interlocking rectangular cube which can be removed via screws for NMR coil servicing. The proximal end is a thicker, single piece which attaches to the insert pipe (also with screws) and is further supported by the ridged NMR coaxial cables and microwave waveguide. The microwave horn that terminates the wave guide forms the bottom of the cage (see Fig. 36).

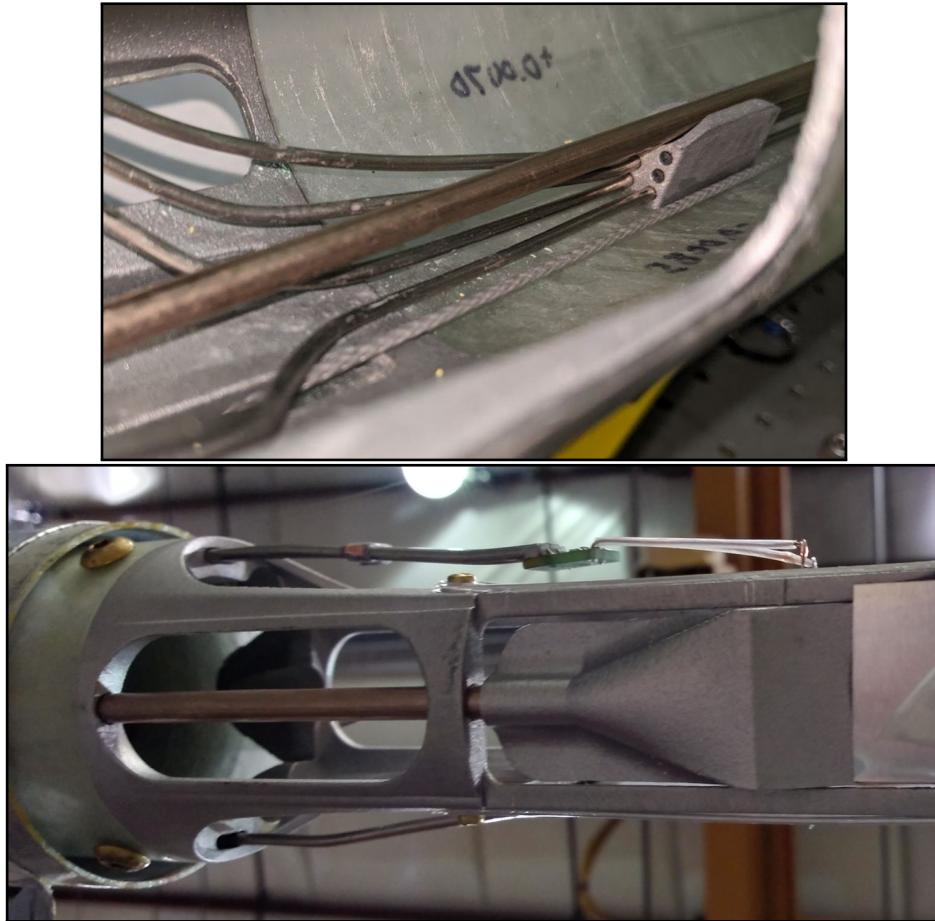


FIG. 36: Top: Locator blocks which are periodically spaced along the lower inner surface of the insert pipe to secure the microwave waveguide and NMR cables. Bottom: Microwave exit horn located on the bottom of the bath cage to evenly distribute the microwave radiation preventing hot (or cold) spots.

The microwave horn is a reflector which disperses the microwaves as they exit the waveguide. The waveguide runs along the bottom interior of the insert tube and is held securely in place by periodic locator blocks on which it rests. The latter half of the waveguide (the section attached to the horn) is made of copper which is an excellent transmitter of microwaves. This first half of the waveguide (nearest the microwave generator) is made of a copper-nickel alloy (commonly called German silver). This alloy has a slight disadvantage at transmitting microwaves, but also has a diminished thermal conductivity which allows it to serve as a

thermal insulator for the latter half. The microwaves run the full length on the insert in a straight line bending only slightly once they exit the insert box.

Running through the same locating blocks are the cables serving the two NMR coils. The circuits for these cables connect through individual SMA feedthroughs on the back wall of insert box (behind the spool). Flexible coaxial cables snake around the outside of the insert box before entering through the bottom of the insert pipe a short distance downstream. From there, they transition to ridged coaxial cables which run securely within the insert pipe though to the same locating blocks as the microwave waveguide (see Fig. 36). Finally, at the locating cage, they split to terminate at the proton coil on the beam-left side and the deuteron coil tank circuit (and subsequently coil) on the beam-right side. These coils are laminated between thin FEP (fluorinated ethylene propylene) sheets onto slim aluminium frames which attach directly onto the locating cage - allowing them to float securely next to the target bath (see Fig. 35).

4.4 OTHER SYSTEMS

Shim Coil Implementation

Attached to the cryostat - exterior to and surrounding the locating cage - is the shim carrier (mandrel) on which the superconducting magnetic correction (shim) coils are wound (see Figs. 25, 27, and 37). These coils use 34 gauge (54 filament NbTi) magnet wire with a critical temperature around 15 K and maximum current of 15 Amps.⁸

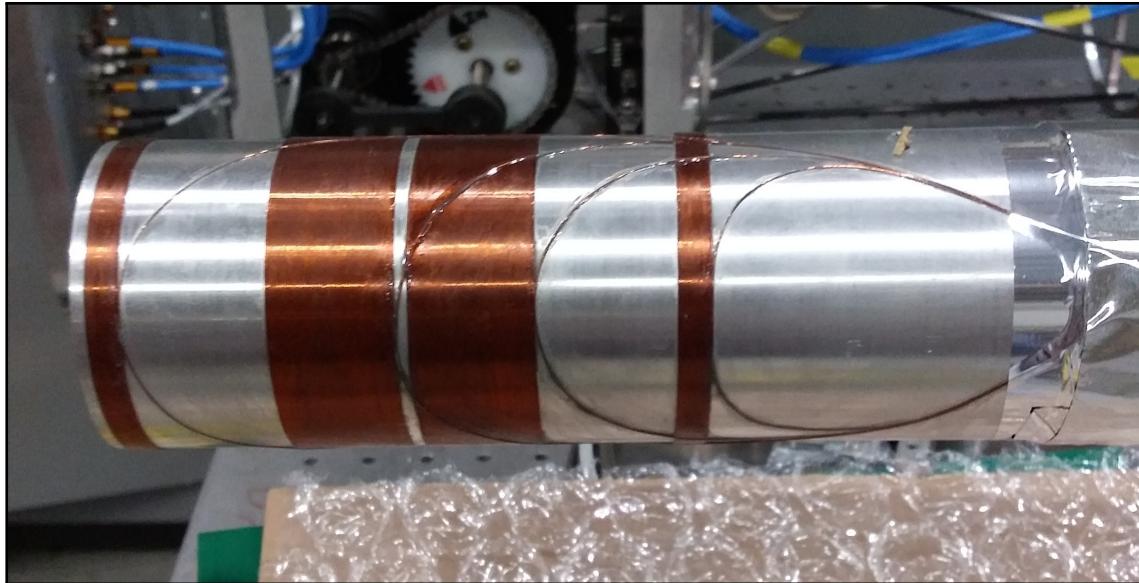


FIG. 37: Final version of shim coils as they were wound on the shim mandrel. These coils were optimized as discussed in Section 4.1 resulting in four coils with two layers of windings each. The center two coils each have a width of 24.5 mm (268 total windings) and are separated from each other by a 2.5 mm gap. The outer two each have a width of 6.5 mm (69 total winding) and are separated from the inner two by a gap of 28.5 mm.

⁸Which is functionally limited to 10 Amps due to both the maximum rating for the electrical feedthrough that delivers current into the fridge and the maximum output of the power supply itself.

During operation, the temperature of the shim mandrel averages 5 K - safely below the wire's critical temperature. However, the path supplying current to the coils is well above this. To minimize the ohmic heat load, the magnet wire connects to high temperature super conducting (HTC) leads. These leads have a critical temperature around 70 K and allow the transition to normal conducting leads to occur much further back in the cryostat - before even the separator. However, since HTC is made of delicate sputtered ceramics, they require considerably more care in construction and handling.

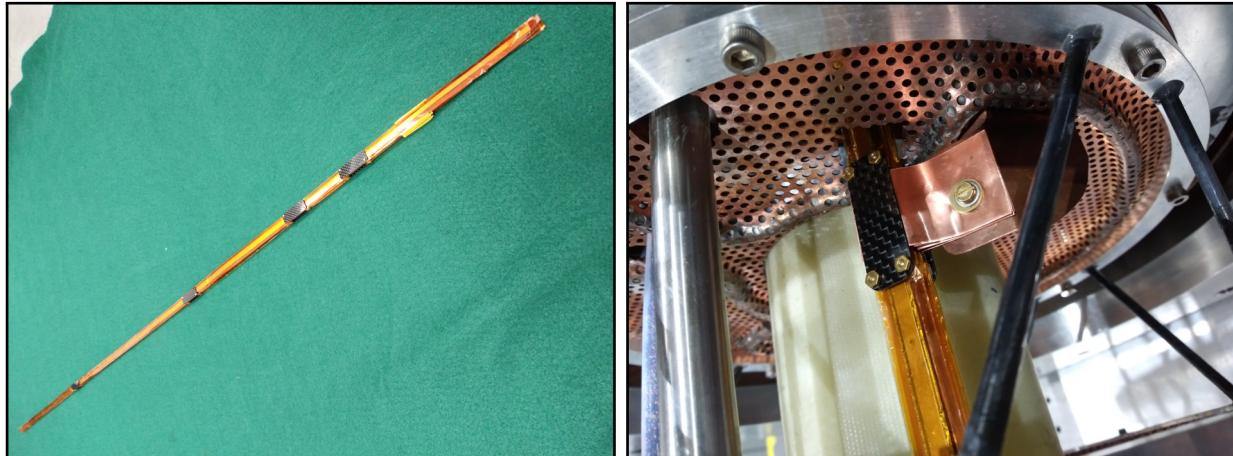


FIG. 38: Left: Completed stack of 8 high temperature superconducting leads. Right: Thermal conduction strips (heat sinks) connecting the leads to the cryostat's heat exchangers.

To prevent damage to these ribbon-shaped conductors, they were soldered onto stacked G10 frames using a proprietary, low-temperature solder. These stacks are periodically heat-sunk to the cryostat by copper strips layered with insulating sapphire slides. In the region in which the cryostat reaches around 50 K, the leads transition to normal conducting brass strips and, finally, flexible brass wires. These terminate at a high current (10 Amps) feedthrough from which they are connected to the power supply.

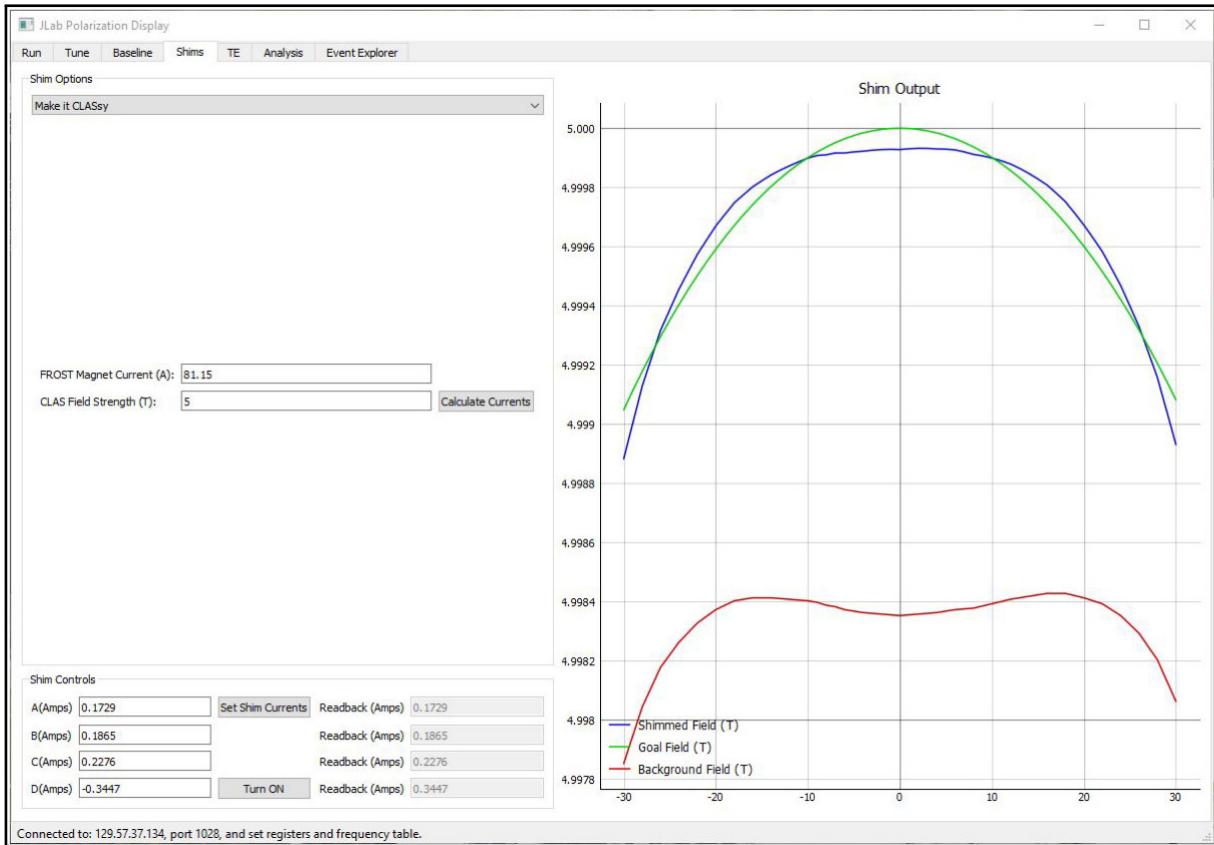


FIG. 39: The shim coil tab of the NMR software. The left panel shows the input options for desired and background fields as well as power supply controls. The right panel shows the expected field results.

The power supply interfaces remotely to the NMR controls which allows for remote manipulation of the magnetic field in response to polarization conditions. The system enables the field to be homogenized and shifted to a range of field values including those needed for NMR baseline measurements.

NMR Controls

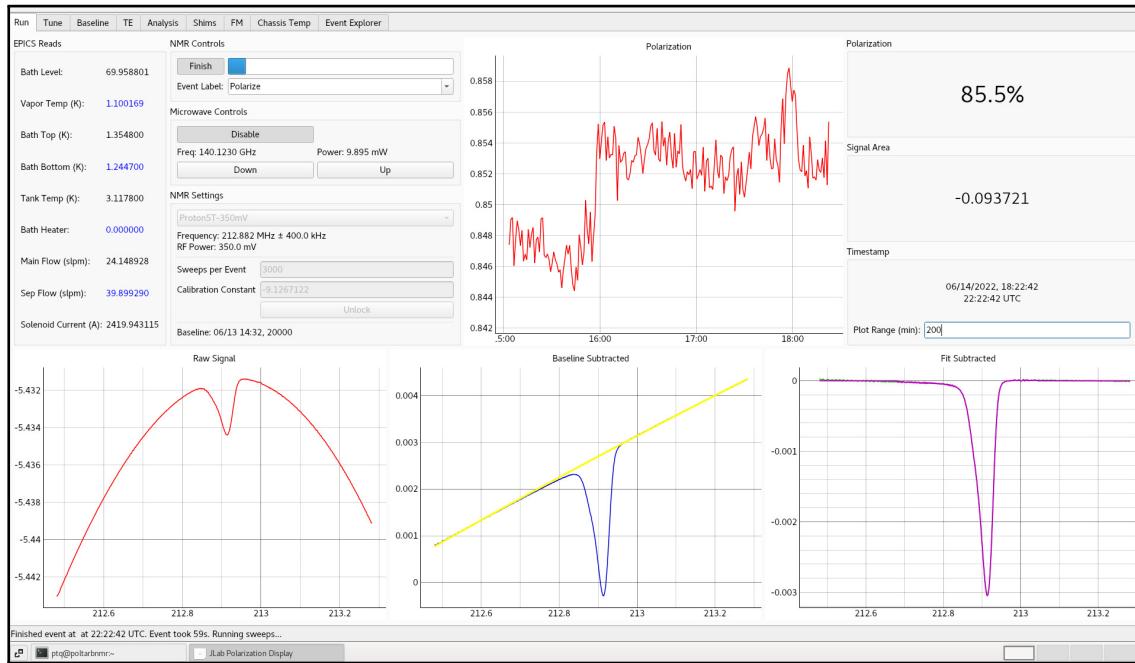


FIG. 40: The main page of the new NMR software. The bottom three charts show the most recent single measurement (raw, background subtracted, and fit subtracted) with analysis values in the top right. The middle top chart shows a history of recent measurements and the top left shows target parameters and NMR settings.

The NMR data acquisition software was created explicitly for the APOLLO target along with the entire NMR measurement system. Historically, polarized target NMR systems have used Liverpool Q-Meters developed in the late 1970's. Since then, individual components have fallen out of production causing difficulty in maintaining legacy hardware. There are currently efforts to reproduce these systems with equivalent modern components; however, advances in the field have also introduced opportunities to improve both the hardware and software.

Among those improvements is the creation of a so-called cold tank board. This board moves both the capacitor and resistors of the RLC circuit down into the cryostat (immediately adjacent to the inductor) and greatly reduces the overall noise of the measurement. The reduced noise is especially critical for the measurement of deuteron polarization since the accuracy of the calibration constant is entirely dependent on the initial TE measurement which, with the former system, approached the magnitude of the background noise for the less polarizable deuteron. As the Larmor frequency (and therefore required inductance) varies between proton and deuteron, the system has two distinct circuits to service them separately (the deuteron with a cold board and the proton without).

Tuning of the circuit is achieved via cold board varactor-diodes whose capacitance varies with an applied bias voltage. These diodes are slightly temperature sensitive, so the board also features a small thermal probe and heater. Additionally, to remove the need to match cable lengths, a remotely adjustable phase shifter was added to the system as well. Both the varactor and shifter controls are integrated into the new Python software which replaces the traditional DAQ system. The software interfaces with the hardware and Rohde & Schwarz RF generator via custom FPGA controllers. This system also interfaces with the microwave controls allowing operators to adjust and modulate the frequency as needed to optimize polarization (which is monitored and logged continuously during operation).

NMR Coils

As discussed in Section 3.5, a significant factor in NMR performance is the shape of the inductance probe. The signal strength coming from a specific portion of material is proportional to the square of the transverse field⁹ from the coil at the location of that material, for a given current in the coil. It would seem prudent to make as large a coil (and therefore as large a homogeneous field region) as possible; however, larger coils also represent larger inductances. To balance the RLC circuit, that also means a smaller capacitance. In theory, there is no problem with this but, in reality, the noise induced from stray environmental capacitance quickly drowns out the signal if the capacitance is allowed to become too small. As is it, the APOLLO NMR system best supports coils around 70 nH in inductance. With such a constraint, it is beneficial to explore coil geometries that make the most use of their fields. To that end a suite of potential coil geometries were modeled and analysed.

⁹The signal depends on the transverse component of field due to Larmor precession and is squared because the proportionality between coil current and magnetic field is the same as between the induced EMF in the coil and the time variation of the magnetic moment of the precessing species.

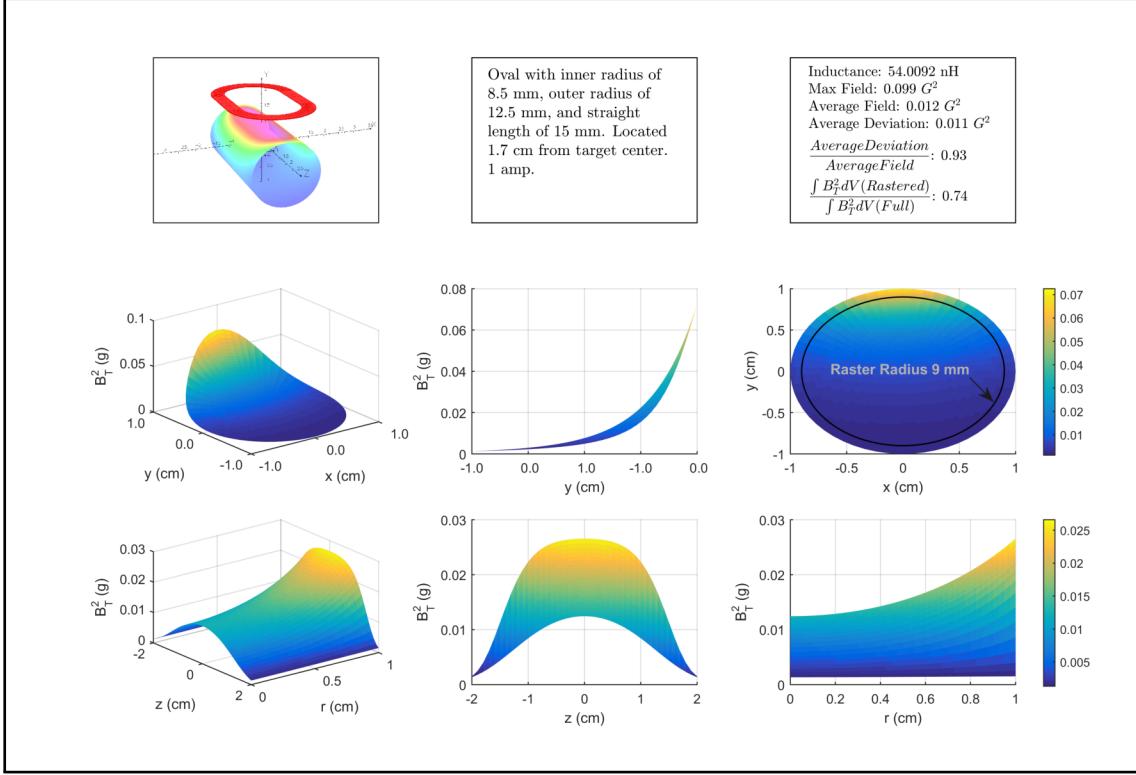


FIG. 41: Sample of NMR coil analyser. The top line has a visual of each coil taken from OPERA3D (left) and a brief description of the model (center). It also has a box of calculated values (right) including: Inductance, Maximum/Average Field Values, and several measures of Uniformity. Average Deviation refers to the average difference between the field at any point in the target region and the field at the center of the target. The ratio of the volume of rastered and unrastered fields has a maximum value of 0.81 due to geometry. The middle row shows the field compressed along z . The left plot gives a 3D view of the z -averaged field (vs. x and y). The center and right plots show the same field in 2D - along y and top-down, respectively. Note that the center plot does not show the field averaged in x , only a 2D view of a 3D plot. The right plot includes a reference circle indicating the extent of the rastered area. The rightmost (middle) color bar applies to all middle row plots. The bottom row shows the field compressed axially. Similar to the middle row, the left plot gives a 3D view of the field (vs. z and r), while the center and right plots show the 2D equivalents -along z and along r . Again the 2D plots do not show the field averaged in z or r , only 2D views of the left plot. The rightmost (lower) color bar also applies to all three bottom plots.

This analysis calculated each coil's inductance, maximum and average field strengths, field uniformity, and percentage of field which falls within the rastered radius of a target cell. An example of such an analysis is shown in Figure 41. The study helped inform decisions about NMR coil shape, but it also gave insight into how we can expect the beam rastering to effect the NMR signal.

During the experiment, the electron beam travels across the cross-sectional area of the target cell (see Section 4.5). This extends the life of the cell by distributing the radiation damage evenly throughout the material volume. However, to ensure the beam never intersects the body of the target cell, this region does not extend fully to the edge of the material. This means that the outermost shell of material receives no radiation dose and therefore does not age and depolarize with the inner volume. This outer shell is also physically closer to the NMR probe coil. So, as the experiment runs, a significant portion of the contribution to the NMR signal comes from an increasingly misrepresentative sample of material.

Polarization in the rastered area (P_R) decreases with dose as

$$P_R = P_0 e^{-D/K}, \quad (58)$$

where D is radiation dose, P_0 is initial polarization, and $K = 25.6 \times 10^{15} \frac{\text{e}}{\text{cm}^2} \times \text{raster area}$. If we assume the polarization in the unrastered area (P_U) stays at P_0 , the measured polarization (P_M) is:

$$P_M = C_R P_R + C_U P_U, \quad (59)$$

where

$$C_R = B_{T(Rastered)}^2 / B_{T(Total)}^2, \quad (60)$$

and

$$C_U = B_{T(UnRastered)}^2 / B_{T(Total)}^2, \quad (61)$$

and $B_T^2 = \int B_X^2 + B_Y^2 dV$. A demonstration of this effect is shown in Fig. ?? and 43.

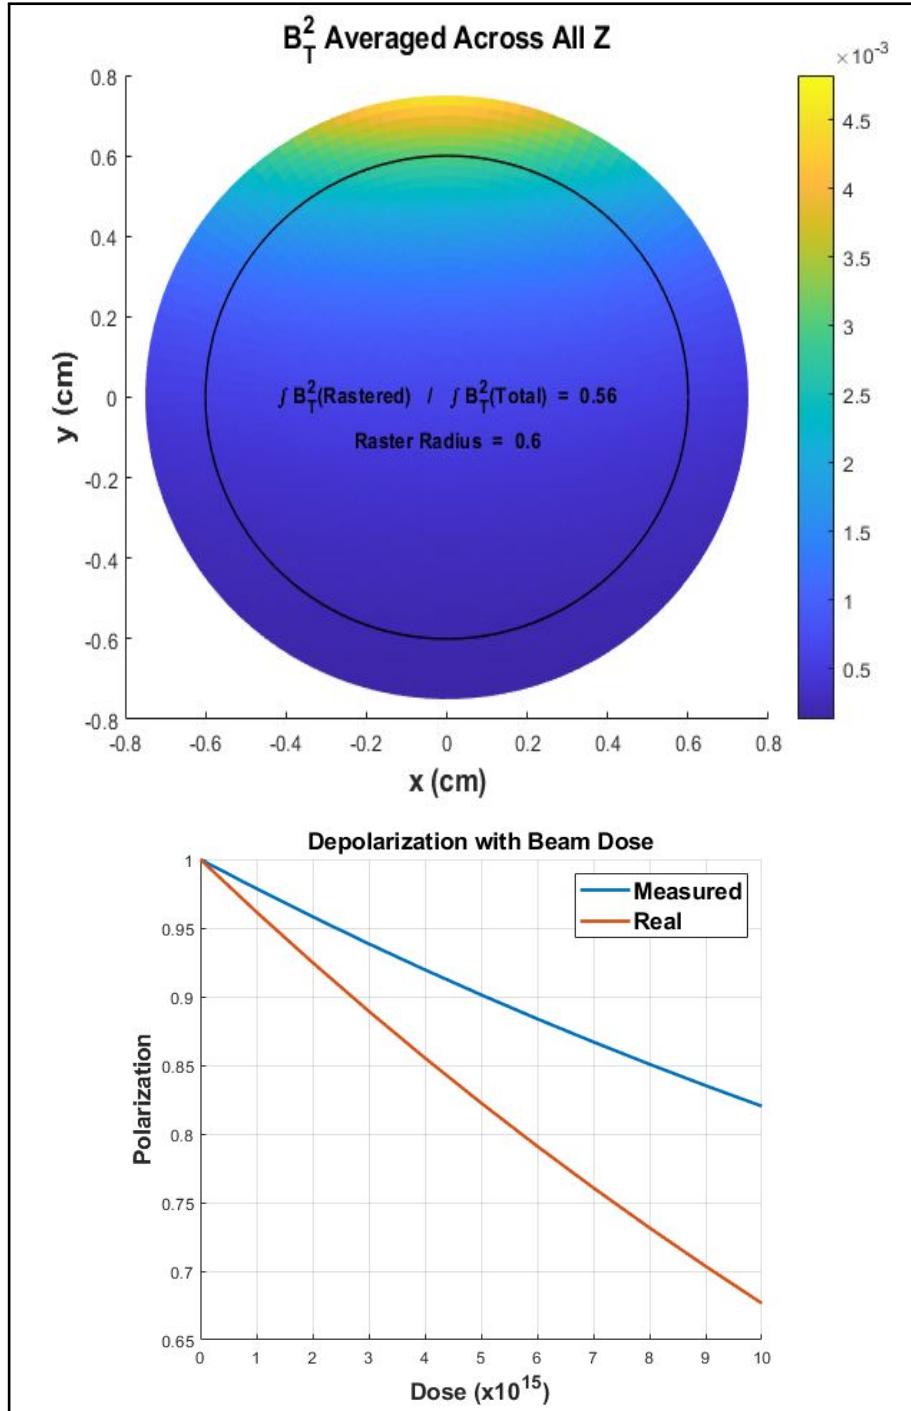


FIG. 42: Analysis showing the effect of rastering *radius* on the accuracy of NMR measurements as a function of beam dose. The top plot shows a cross section of the target material with color indicating field strength from a nominal NMR coil. The black circle indicates the extent of the rastered beam. The bottom plot shows how the polarization measured by the NMR becomes increasingly inaccurate to the real polarization within the raster radius.

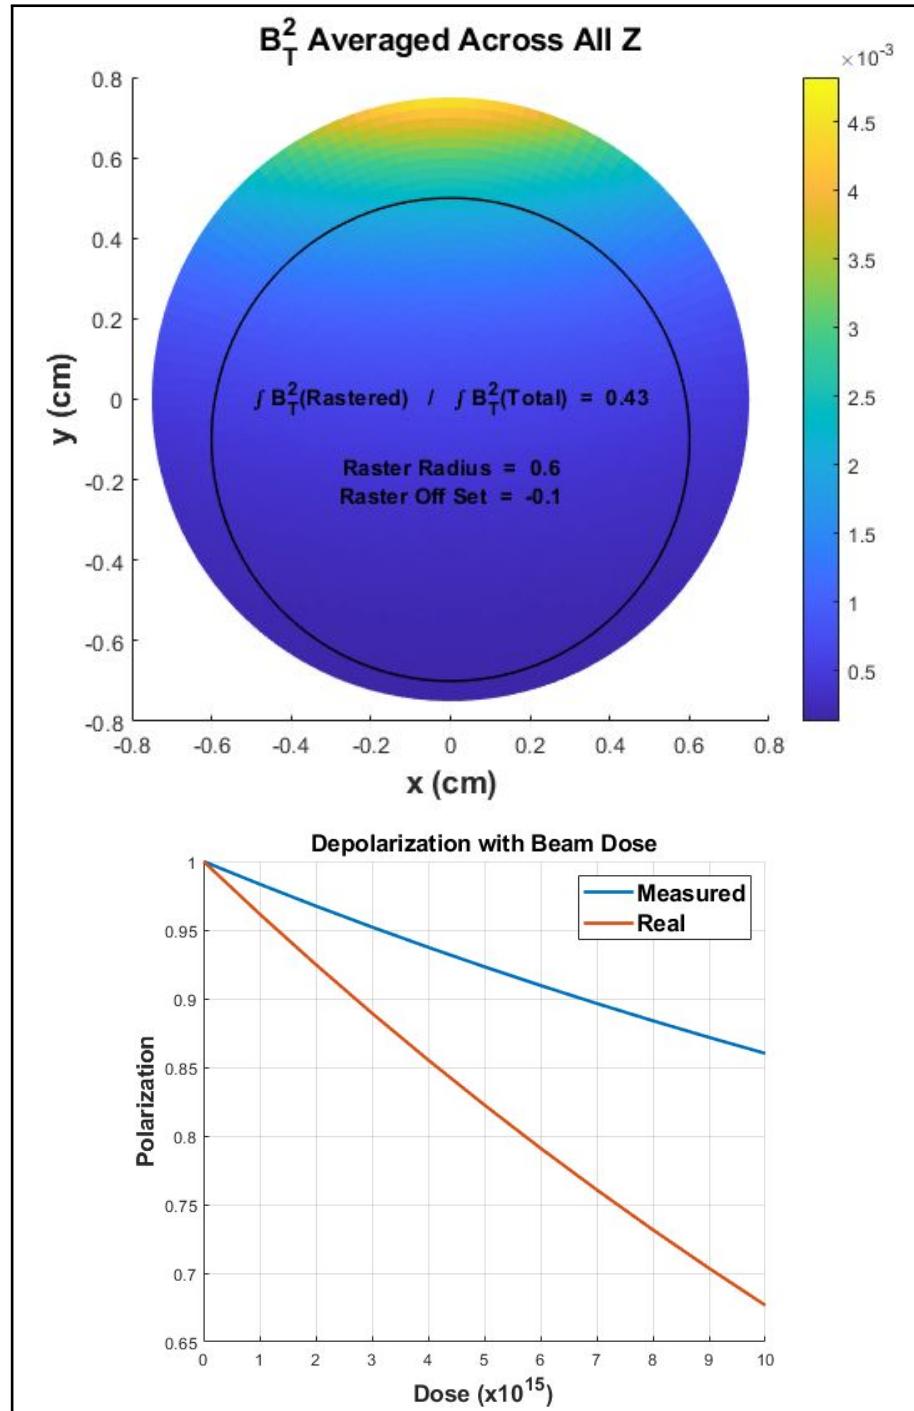


FIG. 43: Analysis showing the effect of rastering *offset* on the accuracy of NMR measurements as a function of beam dose. Plots as described in Fig ??.

Target Cells

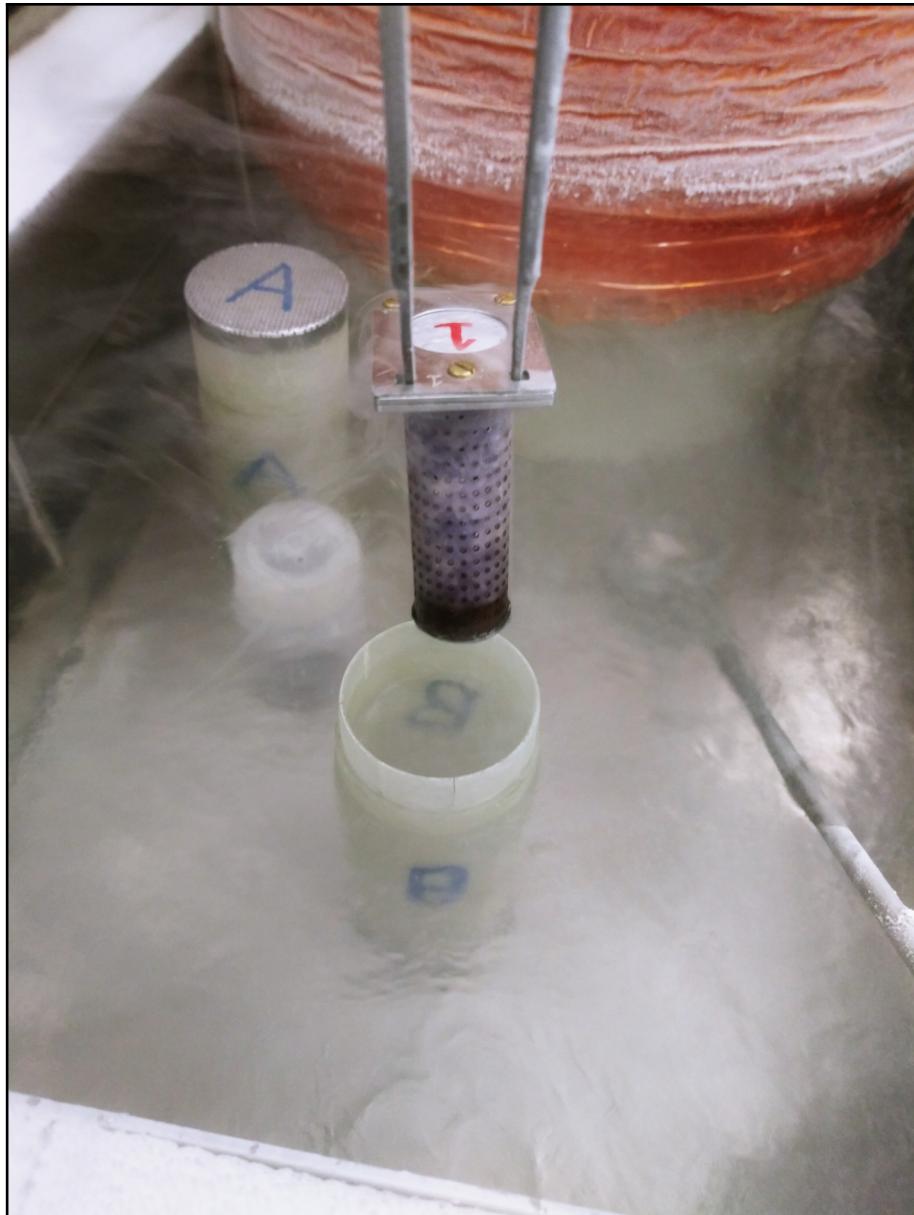


FIG. 44: Ammonia target cell being loaded into a transport container.

As discussed, Run Group C took data on both the proton and deuteron via ammonia and deuterated ammonia. However, to analyse the data gained from these materials, other measurements must be made as well. For instance, taking data with no ammonia in place allows for contributions from the surrounding structures to be measured and subtracted. Additionally, the target dilution factor (from the unpolarizable composition of the ammonia), can be approximated with the use of calibration targets made from equivalently sized carbon and polyethylene cells.

These cells attempt to replicate the same density and distribution as the ammonia cells by spacing three disks of material along the cell's length. There are also cells to help facilitate operation and detector characterization. A special raster target has cross-hair wires to detect the beam position (via the spike in events measured by the detector) for use in commissioning the beam rastering system. A similar optics target uses the same technique with a staggered bullseye pattern to determine cell alignment. There is also an empty cell which varies from the others by only the position of a single thin window. This helps to quantify the vertex reconstruction and DC alignment. Table 6 contains a list of all cells used for the experimental run.

Contents	Diameter (mm)	Length (cm)	Purpose
NH ₃ / ND ₃	20 / 15	5	Production Data
C / CH ₂ / CD ₂	20 / 15	5	Effective Dilution Factor Determination
Empty	20 / 15	5	Background Subtraction
Empty	20	2.5	DC Alignment / Vertex Reconstruction
Wire Cross Hair	20	5	Raster Calibration
Optics Target	20	5	Target Cell Alignment

TABLE 6: List of target cells used for Run Group C. In addition to these, cells containing tempo doped epoxy, butanol, and deuterated butanol were used during development and testing. The epoxy cells were beneficial due to their room temperature tolerance. The butanol cells, while requiring cryogenics, posed far less risk to personnel and equipment than ammonia and represented far less investment risk.

Cell Exchange Procedure

Each of these cells are anchored to a front plate of aluminium¹⁰ which features a double-holed tab used to safely transport the cell to and from the bath (see Fig. 44). A special loading (and unloading) tool fits securely around the tab and two hooks latch into place through the holes. The cells can then be moved to or from a slotted receptacle within the bath which couples to the bezel of the front plate, locking it in place (shown in Fig. 34).

As the ammonia targets must stay below 100 K, this operation is very time sensitive. Additionally, opening the cryostat to access the bath presents the risk of introducing air and other contaminants to the system. To ameliorate this, the access port, under which the trolley retracts the bath, is fitted with a flue which constricts and directs vapor (see Figs. 27 and 45). During cell swaps, the cryostat is placed into a state of helium purge, the outward flow of which sufficiently minimizes any impinging gas.



FIG. 45: Dummy target cell in bath as seen from the access hatch. The reflective surfaces are the angled walls of the loading flue.

¹⁰This front plate also includes a thin circular aluminium film window (matching the thin window on the other end of the target cell) which minimizes unwanted particle scattering from the cells.

The cells themselves also present a potential source of contamination. As they cannot practically be stored in liquid helium, the cells must be reliably cleaned of their storage cryogen (argon or nitrogen). This is achieved by holding the cell in a drying apparatus for a short period before loading. This holding platform flows 80-90 K gaseous helium over the cells allowing the storage cryogen to sublimate while still maintaining the necessary temperature. Once the cells are within the bath, the cryostat is quickly returned to its operating state to protect the material from excess heating.

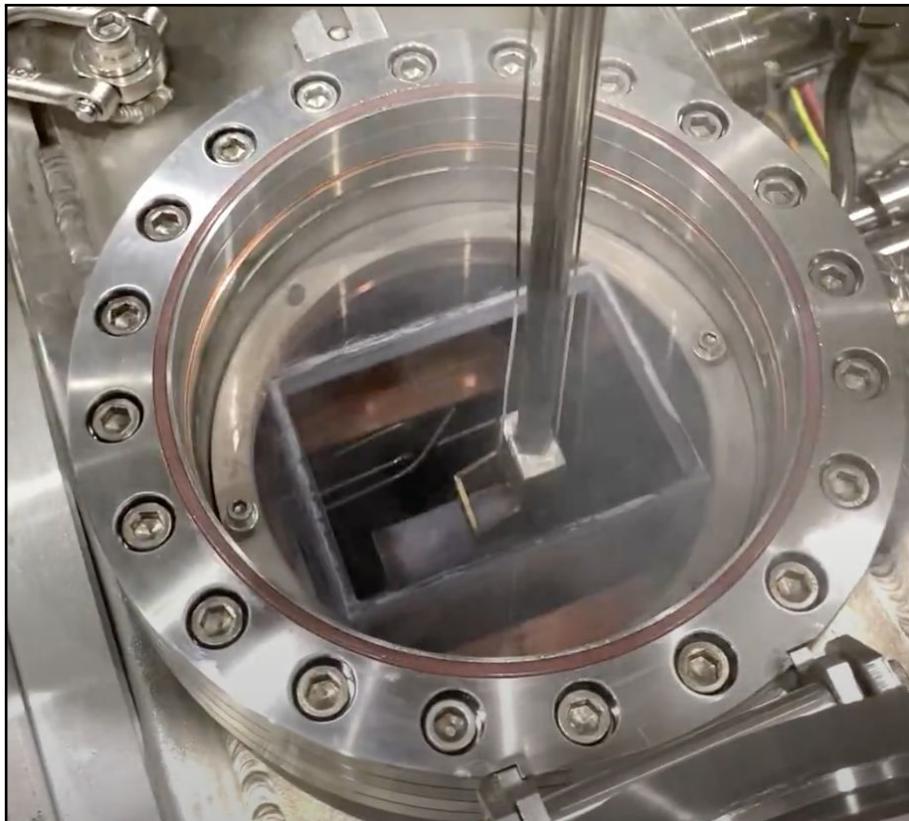


FIG. 46: Target cell being lower through the loading hatch and flue into position within the trolley bath using the cell loading tool.

4.5 INCORPORATION INTO CLAS12

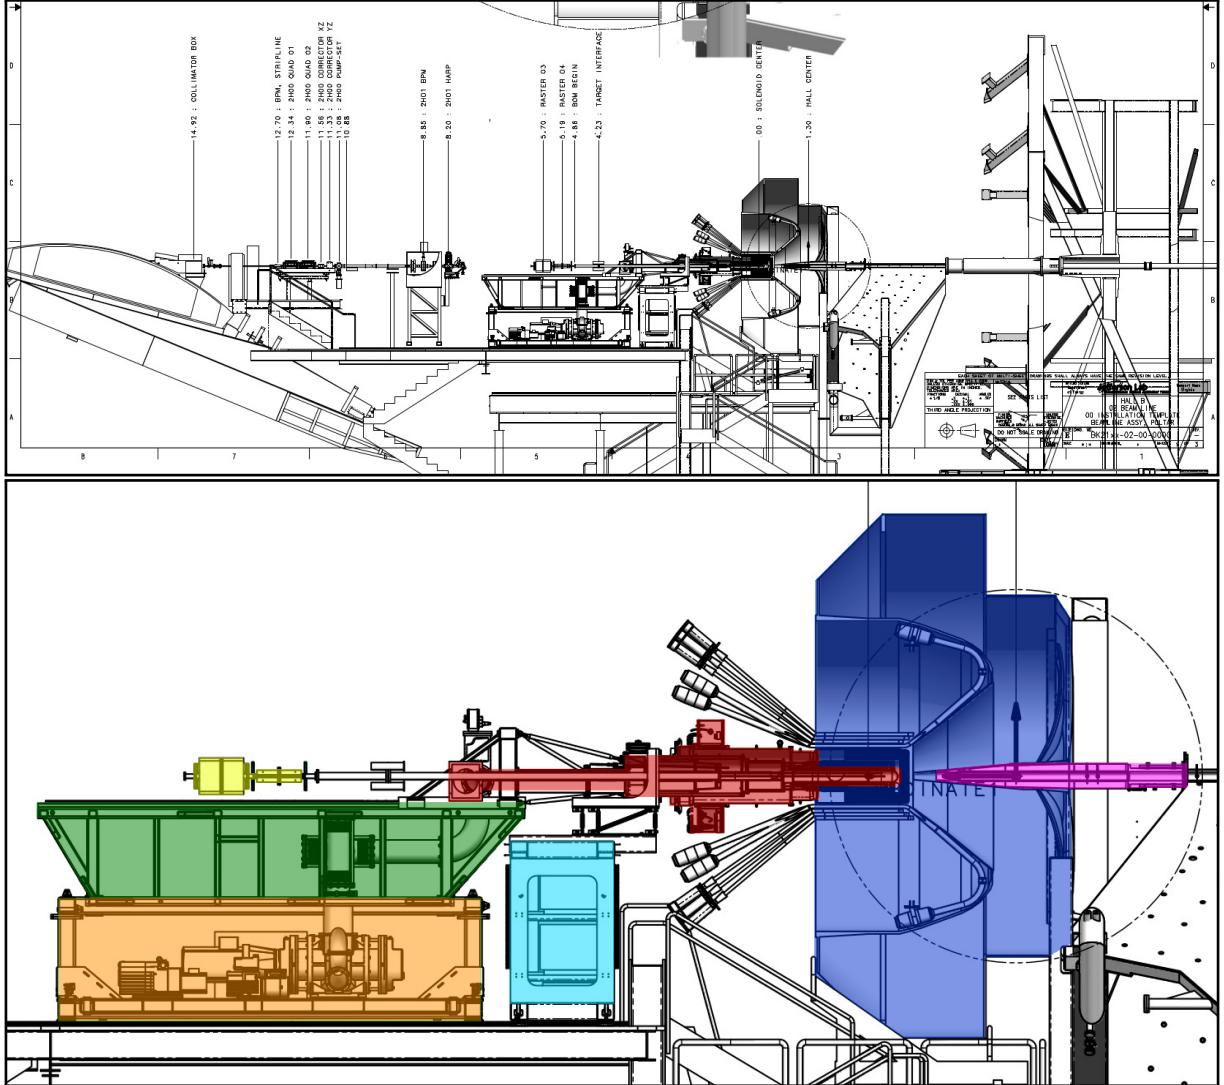


FIG. 47: Top: The complete Hall B beamline diagram for RGC featuring the target (center), spectrometer (right half), raster magnet and other beamline elements (left half). Bottom: A magnified view of the center of the same diagram with colors highlighting: red for the target body, blue for the spectrometer entrance, magenta for the ELMO shield, yellow for the downstream raster magnet, green for the target cart electronics rack, orange for the pump cart, and cyan for the detector platform used for accessing the loading hatch.

Insertion Cart

As numerous stated, the APOLLO target operates at the center of the CLAS12 spectrometer array. To bring the body of the target into its nominal in-beam position within the detector, it is mounted onto an injector cart (see Fig. 47). The upper half of this structure (called the electronics rack) is topped with a sliding mounting cage into which the target is secured. This cage rests on rails which allow the refrigerator to glide forward into a cantilevered position in front of the cart and within CLAS12. The fridge is secured within this mounting cage via adjustable rods which fine tune its position and attitude.

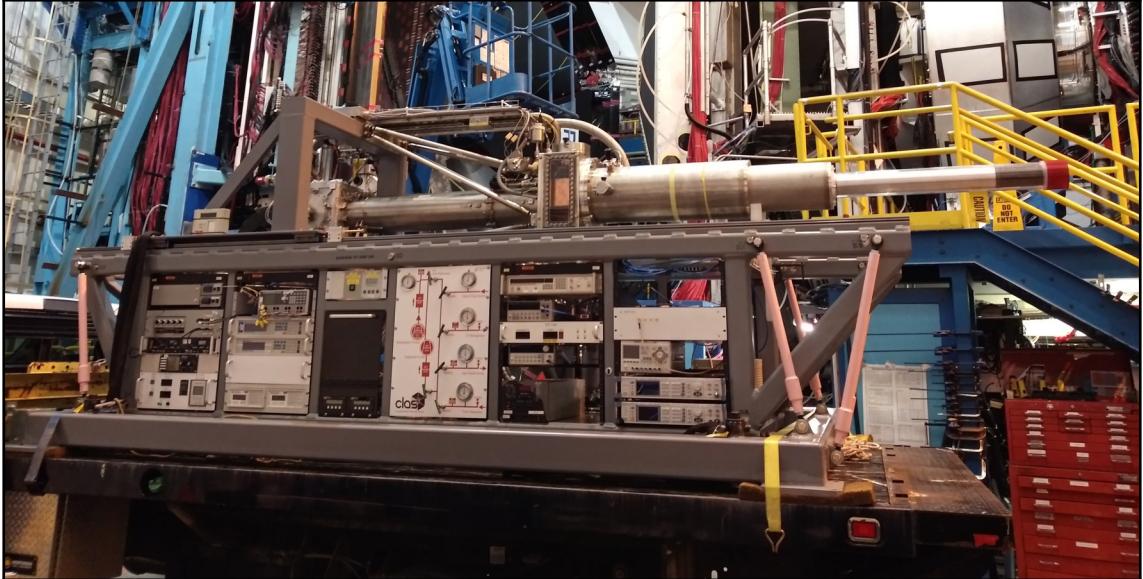


FIG. 48: The upper half of the insertion cart - electronics rack, target cage, and target (in the retracted position) upon delivery to the experimental hall.

Beneath the cage, the electronics rack houses the equipment for the monitoring sensors and power supplies along with the gas panel valves that regulate portions of the helium path into and out of the cryostat (see Fig. 48). This section of the cart is insulated from the lowest section (pump cart) by vibration absorbing pneumatic cushions. The pump cart

is then itself connected to rails on the floor of Hall B which allow the whole assembly to smoothly move into and out of place in front of the spectrometer.

This compact design allowed the target to be nearly fully assembled before entering the experimental hall and greatly improves the efficiency of installation. Once in the hall, the freed space on the topmost portion of the cart (created by rolling the cryostat into the spectrometer) is occupied by beamline components and the microwave delivery system.

Also, as shown in Figs. 27 and 49, the loading hatch for accessing target cells is located at the back of the cryostat above the back plate - which is the point at which the target, when deployed, enters the CLAS12 spectrometer. To reach the hatch, an existing detector platform was modified to include railing and a helium vapor shield (to protect sensitive spectrometer components).

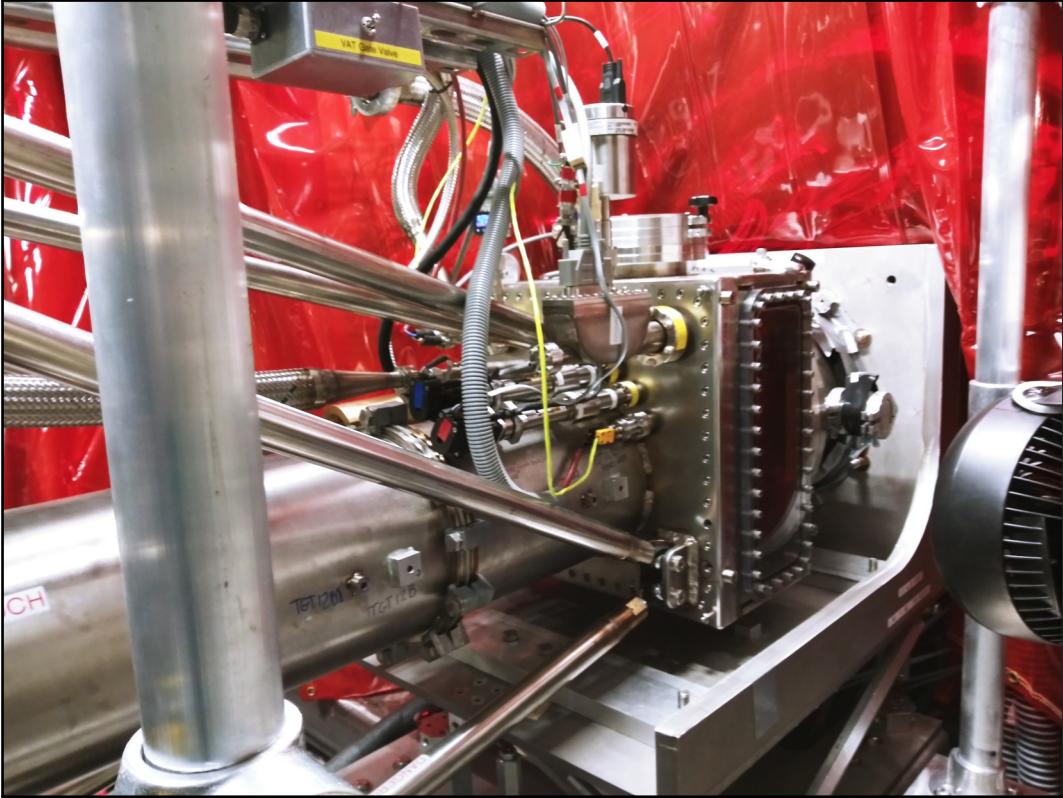


FIG. 49: The back of the cryostat (with loading hatch on top and target cage adjustment rods attached to the back) as it sits when fully deployed into the spectrometer. The red tarp acts as a helium shield for the detector components immediately behind it.

Raster System

A crucial component of the beamline is a paired set of raster magnets. As discussed previously, the electron beam must constantly travel across the cross section of the target material to evenly distribute the radiation damage and heat throughout the sample. This is achieved by generating a precisely controlled magnetic field far upstream of the target which bends the electron path very slightly away from its original trajectory down the center line of the target. The beam then travels in this new direction (diverging from the center line) until it passes through a second magnetic field which bends its trajectory back parallel to (but offset from) its original path. There is a set of these paired fields for both the x and y directions. The degree and direction of offset is determined by the combined magnitudes of those fields as well as the distance between magnet pairs.

This system was successfully employed during the 6 GeV era of CLAS; however, with the increased beam energy of the upgraded accelerator (and resultant stronger field requirement to bend particle paths) it was unclear how much of an offset the extant raster magnets would be able to provide. Unfortunately, in the decades since their construction, the details of their composition were lost - save the value of their inductance ($425 \mu\text{H}$) and maximum allowable current (240 A). As all that was needed was a conservative initial estimate, an attempt was made to reconstruct the missing information.

To determine the rastering limits for a 12 GeV beam, the maximum field strength of the raster magnets, integrated along the beam line, needs to be known. The number of windings in each magnet (and thus total field strength) were unknown. We concluded that the best avenue would be to infer the number of windings from the known inductance values of the magnets. From there, we could calculate the field produced at maximum current and, subsequently, the maximum displacement possible for a '12 GeV' (actually 10.5 GeV) beam.

To achieve this, physical measurements of the raster magnets were taken. Since the original specifications of the magnets no longer exist, estimates were made using the external dimensions of the magnets. After compensating slightly for the insulating wrapper, the magnets were assumed to have a cross sectional area of 3.5 inches by 0.8 inches and a straight length of 15 inches.



FIG. 50: Coarse physical measurements being taken of the raster magnets.

These estimates were used to create approximate models of the magnets using OPERA 3D. This software models conductors as solid blocks of uniform current density, without regard to windings. A dummy value of 100 A/m^2 was used for the current density - the equivalent of $\sim 0.18 \text{ A}$ total current if the magnets had just a single winding. The resultant field was exported for analysis in Matlab.

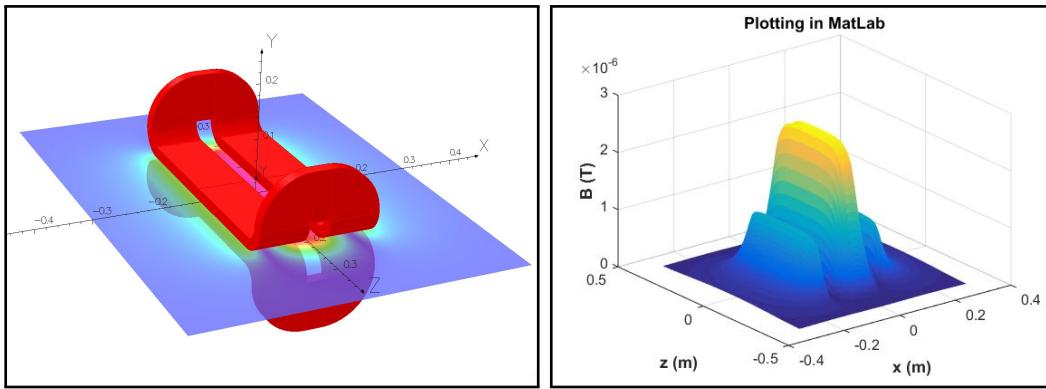


FIG. 51: Left: Approximate model of a single raster magnet in OPERA 3D with a single plane of (B_y) field values shown. Right: Plot of those field values in Matlab to be integrated and combined with others to calculate the total stored energy of the field.

Using Matlab to calculate the total integrated squared field (together with the reported inductance value of $425 \mu\text{H}$ and dummy current) we can calculate the number of windings.

Inductance is related to the energy stored in a magnetic field by:

$$U = \frac{1}{2}LI^2, \quad (62)$$

where L is the inductance and I is the current (per winding). We also have:

$$U = \int \frac{1}{2} \frac{B^2}{\mu_0} dV, \quad (63)$$

where B^2 is the squared magnitude of the field and μ_0 is the vacuum permeability. We combine those to get:

$$\int B^2 dV = \mu_0 LI^2 = \mu_0 (425 \mu\text{H}) (0.18 \text{A/windings})^2, \quad (64)$$

which was found to be approximately 20 windings. From there, the total maximum current was assumed to be 4.8 kA (maximum current times 20 windings).

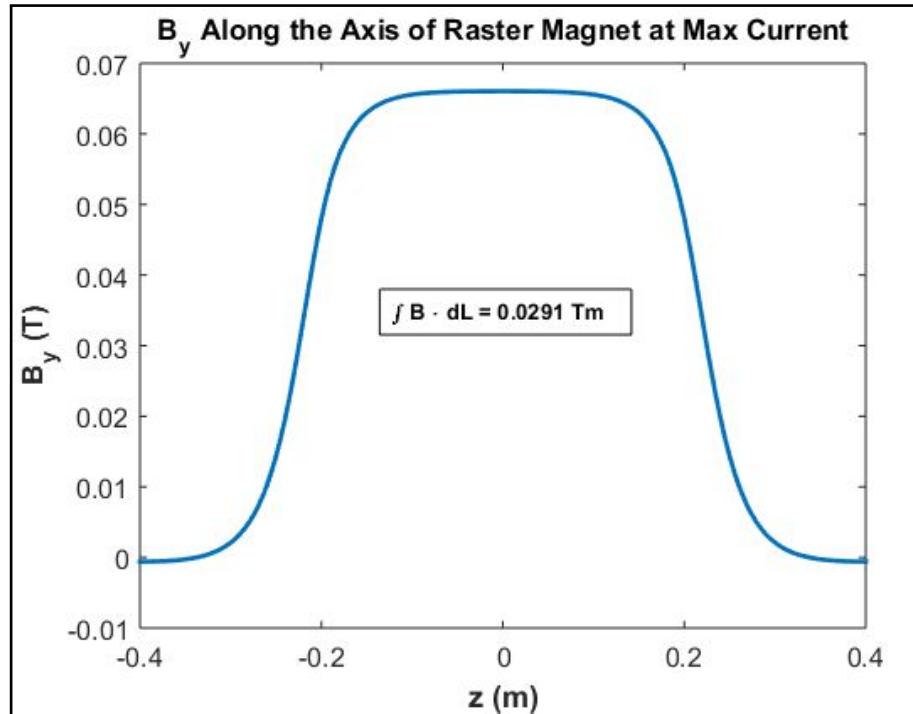


FIG. 52: Plot of the modeled vertical field component along the axis of a single raster magnet at maximum current.

Scaling the modeled field to this current, the vertical field component (B_y) was plotted (Fig. 52) and integrated along the horizontal length (z) of the magnet to find the angle of deflection via:

$$\theta_{\text{deflection}} = \int B_y dl (0.3/10.5) \approx 8.3 \times 10^{-4} \text{ rad.} \quad (65)$$

The final step in determining the maximum raster amplitude was to multiply the tangent of the deflection angle by the distance between the two sets of raster magnets. After consulting with beamline experts, initial placement positions were decided upon. These positions were integrated into a preliminary beamline diagram. This positioning allows for approximately 46 feet (14 m) of distance between magnets which generates a maximum raster amplitude of just over 1 cm - sufficient for Run Group C.

As this was just an initial estimate, part of the commissioning of the raster system included testing and calibrating. The primary concern of the raster is to avoid moving the electron beam into any beamline components while maximizing its coverage of the cell interior. The previously mentioned raster calibration cell (see Table 6) was designed to allow the system to steer a safe, low current beam across its cross hairs while recording detector rates at every position. With the resultant spikes of collision events giving a clear indication of raster extant and allowed definitive calibrations of the voltage to raster radius to be made (see Section 5.2).

This calibration confirmed that the magnets provide more than enough field for the desired raster size - indicating that our initial analysis underestimated the integrated field. This is possibly due to the presence of ferromagnetic material in the magnets which would increase the field beyond that which was calculated from coarse dimensional estimates.

ELMO Møller Shield

As discussed in section 4.1, the forward tagger (FT) detects small angle forward scattered electrons and photons. It also includes a tungsten cone which shields the detector from background electromagnetic radiation generated by the electron beam - particularly Møller electrons. However, this Møller shield is designed for a centrally located electron beam. The rastering needed for the APOLLO target offsets the beam, potentially defeating the shield and overwhelming the CLAS12 drift chambers (DC). One solution was to replace the FT with more robust shielding which allows for beam rastering at the expense of a kinematic region. Studies were conducted to develop this bespoke Møller shield. Additionally, as DVCS experiments greatly benefit from the inclusion of the FT, studies were also conducted to determine the maximum raster allowable with the existing FT configuration.

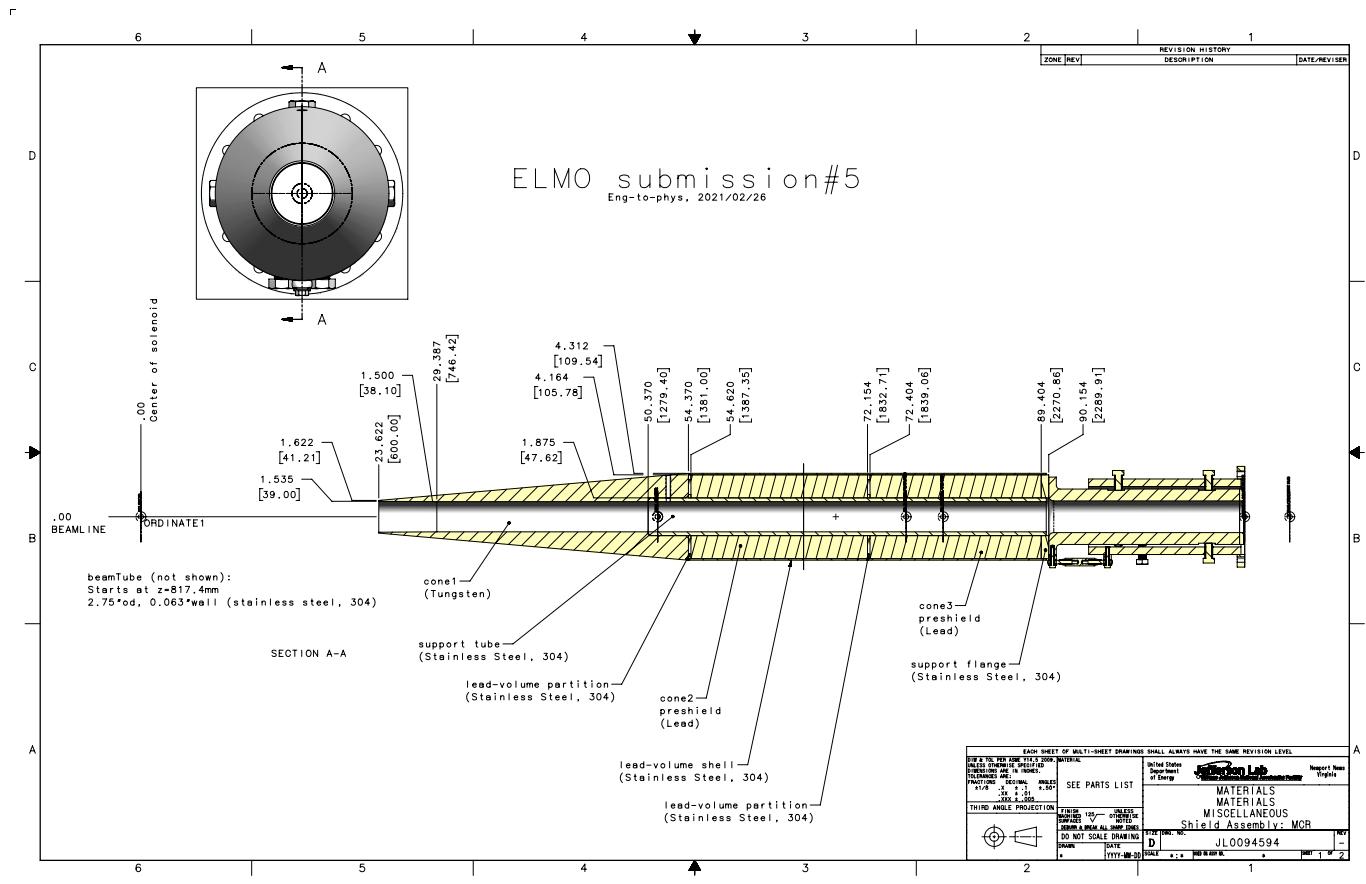


FIG. 53: Engineering diagram of the ELMO Møller shield for use in the FT-Off configuration (when the CLAS12 Forward Tracker is removed from the detector array). This shield enables an increased beam raster radius (and therefore target cell radius) of 1 cm compared to 0.75 cm for the FT-On configuration.

These studies were conducted using GEant Monte-Carlo (GEMC) software with the aim of maximizing the rastered offset of the beam while keeping the DC event rates within an acceptable threshold. The culmination of these studies produced the ELMO (ELECTRON Møller) shield which replaces the FT and minimizes background rates from beams rastered

up to a radius of 1 cm. As shown in Fig. 53, the ELMO shield is composed of tungsten, stainless steel and lead. These studies also concluded that, without the ELMO shield, the FT could be used with beams rastered up to 0.75 cm.

This means that for FT-Off configurations (using ELMO without the FT), target cells could have a diameter of 2 cm while FT-on configurations (ones including the the FT with its standard shielding) needed target cells of no more than 1.5 cm diameters. While this may not seem like a significant difference in size, the decrease in cells dimension effects the frequency of target swaps as well as the effectiveness of the NMR system. As previously described, the NMR coils are held statically in place at the nose of the target - meaning smaller target cells are further from the coils and therefore produce a smaller overall signal (less easily discernible above noise).

Finally, the reduced raster radius leads to a smaller rastered-over area (by a factor of 2). This is mitigated by reducing the beam current by the same factor of 2, leading to similar time intervals between necessary target anneals. This reduction was also necessary because the less efficient shielding of the standard FT Moller cone led to too high background in the drift chambers.

CHAPTER 5

PERFORMANCE AND CONCLUSION

5.1 PERFORMANCE DURING DEVELOPMENT

The process of integrating all of the components described so far into a fully operational target required years of development and testing. While the general principles at play are well-established, implementing them for the specifics of Run Group C offered unique challenges and necessitated extensive trials. As the components were designed and tested in the EEL building of Jefferson Lab, away from the helium recirculation plant; cryogenic testing could only occur on an incidental basis (or at the higher temperatures offered by other cryogens). It was therefore of critical importance to have an operational cryostat early in the process to enable unhindered development of the other components.

Cryostat

The immediate success of the refrigerator on its first cool-down was a pivotal achievement for APOLLO's development. Not only was it able to successfully achieve sub-Kelvin temperature, it also maintained very low temperature while under significant heat load. Using the bath heater to simulate microwave, beam, and external heat sources, the cryostat demonstrated a cooling power well beyond any anticipated need (see Fig. 54). Furthermore, it did so efficiently, with minimal helium consumption. This meant that, in addition to reducing the overall heat load on the ESR once in the hall, the subsequent cool down tests in the EEL could run with less overhead - since the helium had to be supplied via limited standalone dewars (without helium recapture). This success was followed by others as the level probe and automated PID controls were implemented with high effectiveness (see Fig. 55).

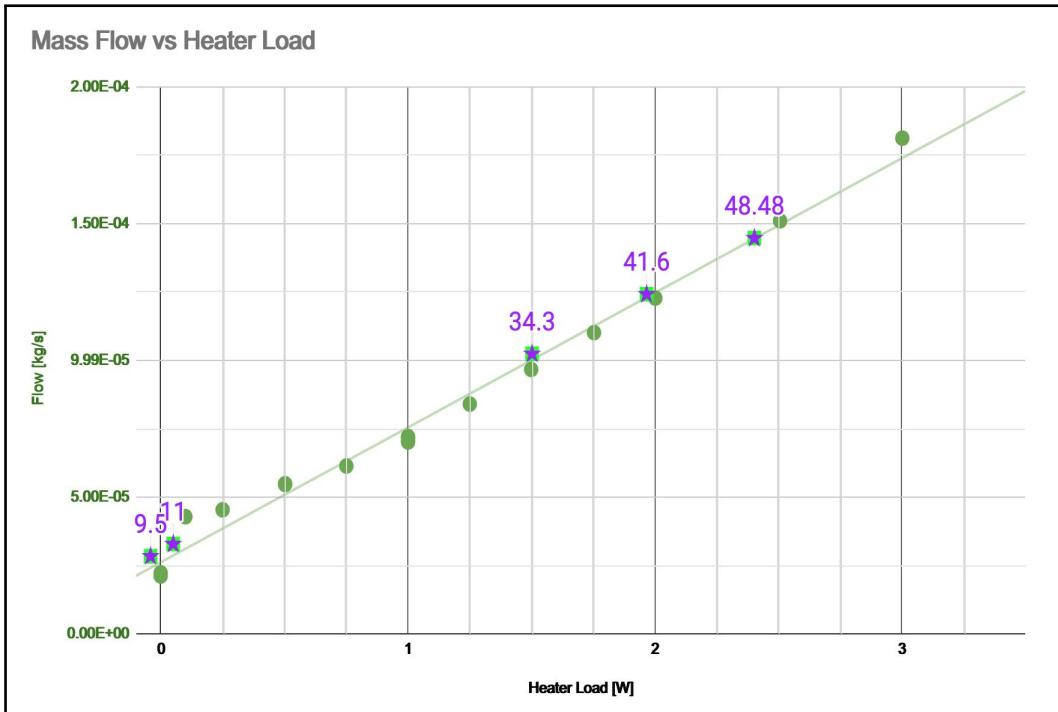


FIG. 54: Plot of cryostat helium usage vs. added heat load. The green points show measurements of heat delivered via an ohmic heater located in the target bath during an early (2019) cool-down test. The purple points show subsequent (2021) measurements of heat delivered via microwaves (with the numbers indicating He flow in slpm). For all measurements, the cryostat was operating at its running temperature near 1 K with lower flow rates performing at the lowest temperature (~ 0.9 K) and temperatures increasing with flow rate until a cutoff temperature of 1.2 K.

Microwaves

Adding to this efficiency, it was found that the microwave system actually over-performed in its initial configuration of ~ 1 Watt. By comparing target polarizations across a range of delivered microwave powers, it was determined that the cells were fully saturated using only a fraction of the maximum available wattage. Furthermore, frequency modulation was incorporated into the microwave system which broadened the band of frequencies delivered to the target. This both improves overall polarization and also reduces the sensitivity to microwave tuning.

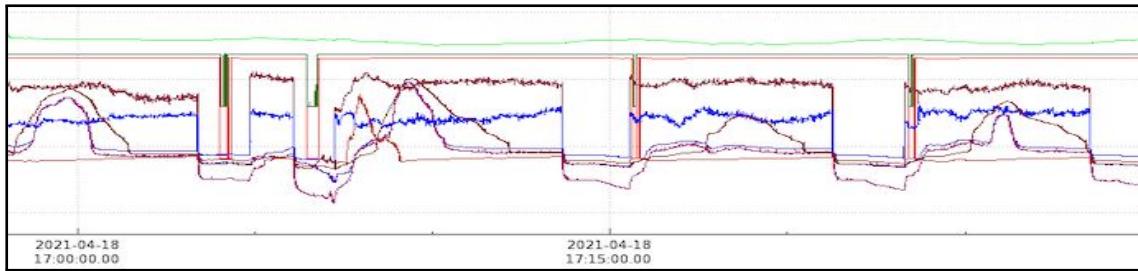


FIG. 55: Plot demonstrating the effectiveness of the liquid He level probe PID system. The different colored lines indicate various thermal probes along the cryostat with the exception of the topmost pale green one which indicates the liquid He level (all plotted against time with arbitrary vertical units). The erratic jumps in the lower lines display the sudden temperature changes induced by the introduction or removal of microwaves energy. Note that the upper green line, however, stays remarkably static throughout indicating the effectiveness of the automated PID system.

Trolley

The final milestone was the ability to maintain appropriate temperatures during material changes and rapidly recover to operating temperature afterward. This depended not only on the cooling rate of the cryostat, but also on the efficiency of the trolley system. As shown in Fig. 56, the cell exchange protocol was extremely successful. Not only did target cells stay well within their safe handling temperatures, but the refrigerator returned to an operational state very quickly. This is in addition to the already discussed success of the movement mechanisms reliability (see Section 4.3).

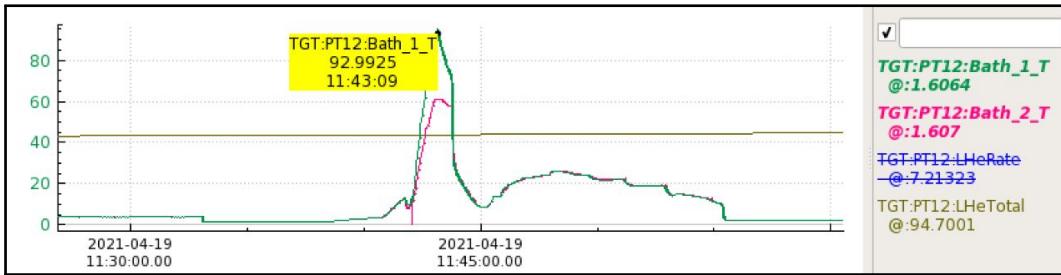


FIG. 56: An example of the temperature changes experienced by the bath during tests of the target cell exchange procedure. The green and pink lines display the temperature reading of thermal probes located at the top and bottom of the bath. This plot demonstrates the efficiency of the cell swapping technique developed for APOLLO with cells staying below 100 K and the cryostat recovering to operational temperature in less than an hour. It is worth noting, as well, that the displayed temperatures are measured somewhat upstream of the target cells (in a warmer section of the cryostat) meaning they represent the *absolute maximum* potential temperature reached by the bath. The cells themselves are actually substantially cooler.

Shims

Obviously, as the APOLLO development was happening outside of Hall B, the CLAS12 solenoid was not available for testing. Fortunately, a preexisting 5 T solenoid (which was used during the FROST experiment [74]) was available with a surprisingly comparable warm bore diameter. This much smaller solenoid more than met the testing needs of APOLLO. However, its central polarizing field was considerably more uniform than that expected from CLAS12. So uniform, in fact, that there was no anticipated benefit to correcting it via shims coils - which limited the type of testing that could be performed on them. However, the question of whether or not the CLAS12 solenoid would be need to be corrected was yet unanswered. To help with that endeavor later, a test was also conducted to see how sensitive the APOLLO system was to field inhomogeneity. After an initial test to confirm that the shims were able to manipulate the local field as desired, that field was then intentionally spoiled by introducing inhomogeneity to observe its effect on measured polarization. As shown in Fig. 57, the APOLLO NMR system was able to clearly reflect the intended field effects induced by the shims, giving confidence to their effectiveness.

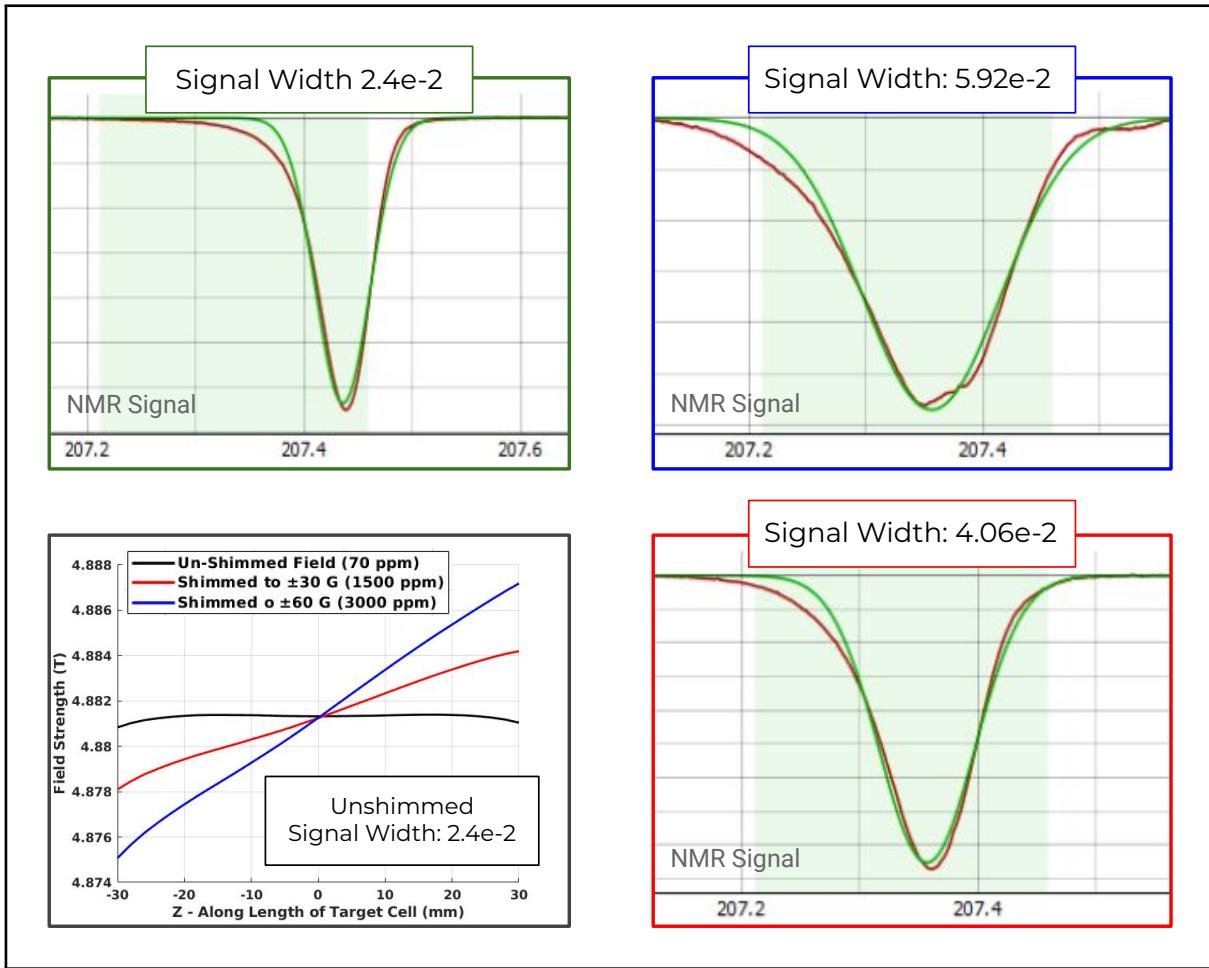


FIG. 57: Results from initial tests of the shim coils. The bottom left plot shows the local polarizing magnetic field across the length of a target cell. The *black* line represents the unmodified background field (produced by the FROST solenoid). As indicated by the *black* text box, this field produced an NMR signal width of 2.4×10^{-2} MHz. The shim coils were then energized such that the overall field was uniformly increased by 95 G across the full cell length. This produced the NMR signal shown on the top left. As expected, this signal was shifted laterally by an appropriate amount (indicating a change in field magnitude) but retained its prior width (indicating no change in field uniformity). After returning to the nominal field, the shim coils were then energized such that they intentionally introduced inhomogeneity. This is represented by the *blue* and *red* lines on the bottom left plot which correspond to the *blue* and *red* plots on the top and bottom right. As anticipated, these signals were *not* shifted but *did* increase in signal width to 5.9×10^{-2} MHz and 4.1×10^{-2} MHz respectively.

NMR and Polarization

After an extensive and intensive period of development, the NMR system was deployed with outstanding success. As evidenced by the multiple NMR figures presented already, the bespoke Python software suite produces exceptionally clear, real-time polarization measurements reliably and intuitively. Likewise, the novel advancements made to the DAQ, Q-meter, and tank circuits represent an achievement fully independent of the APOLLO project - a point made clear by the fact that it received a patent during the period of target development. In that time, it allowed the target to be tested against various polarizable materials. In the final test before deployment, polarizations of tempo-doped epoxy reached 70%, irradiated NH₃ reached +94% and -78%, and (excitingly) tempo-doped Deuterated-butanol reached a polarization of 28% (see Fig. 58). The final test, however, was of irradiated ND₃ which only achieved a polarization of 0.3%. This measurement might have raised serious concern about the efficacy of the target had not the NMR already thoroughly confirmed APOLLO's ability to polarize deuterated material. The rigorousness of the NMR system allowed for a reasonable assumption that the low polarization was due to the specific material sample.

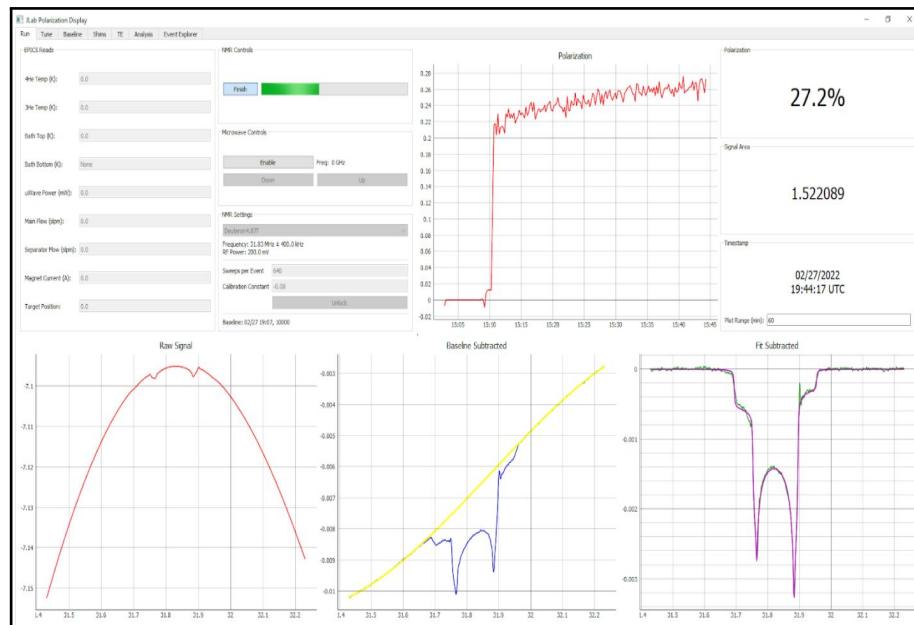


FIG. 58: NMR measurement of TEMPO-doped deuterated-butanol in a 2 cm diameter target cell and irradiated with 136.7 GHz microwaves.

Run Plan

With the APOLLO target fully developed, the RGC experiment suite could begin. At the beginning of the run period, the planned experiment schedule was an eight-month run spanning from June 2022 to March 2023 totaling 120 PAC days (equivalent to 240 calendar days). Of those 120, 30 would utilize the FT-on configuration (discussed in Sections 4.1 and 4.5) with the associated reduced beam current (4 nA), raster radius (6 mm), and target cell diameter (15 mm). Those 30 days were to be split evenly between NH3 and ND3 with a beam energy of 10.5 GeV and exclusively in-bending torus polarity.

Following that, the remaining 90 PAC days would use the FT-off (or ELMO) configuration (with a one week allowance for removing the forward tagger and installing the ELMO Möller cone). This would enable the larger (20 mm diameter) target cells to be used (with a 9 mm raster radius) as well as the higher beam current of 10 nA. Once again, the PAC days were to be split equally between NH3 and ND3 with 10.5 GeV beam energy for both. This portion of the run, however, would dedicate 1/4 of its time to out-bending torus polarity.

All of this, however, would follow the initial set up phase during which the target would be installed in the hall and commissioning at 2.2 GeV. Additionally, the raster system would be calibrated, luminosity studies conducted, and polarization measurements made of both the beam and target.

5.2 PERFORMANCE DURING EXPERIMENT

As covered in Section 4.5, the target was designed to optimize the installation process with the contained insertion cart being able to be moved into position as a single unit. Although this plan had to be slightly modified due to geometric limitations in the hall, the installation still worked as intended to minimize time overhead. After only minor delays, the first beam was successfully delivered to APOLLO on June 11, 2022 marking the commencement of RGC.

Cryostat, Trolley, and Shims

The outstanding cryostat performance that was established during development continued once in the hall. When idle, the target maintained a nominal temperature of ~ 0.94 K (with 0.145 g/s LHe consumption) which increased to only ~ 1.19 K (and 0.305 g/s LHe) when under significant heat load. Additionally, the target swapping procedure was perfected such that the cell exchange could take as few as ~ 15 minutes and be completed using only

a couple well trained experts. Over the course of the experiment, cells were changed 75 times with no sign of fatigue in the trolley mechanism or communication tether.

The shim coils, however, proved to be unnecessary as the CLAS12 solenoid outperformed its expected homogeneity. Some small studies were done to investigate potential impacts of modifying the local field using the coils. However, the effects were on the order of the NMR signal noise and would have required substantial experiment downtime to pursue. Unfortunately, before such further investigations could be made, an error occurred which caused one of the coils to overheat and lose continuity. However, even with that loss, the remaining coils were able to easily compensate to produce the desired homogeneity (albeit with reduced functionality).

The more significant magnetic event of the experiment was produced by the solenoid which, during regular target maintenance, suffered a firmware error. This error caused the solenoid to attempt a polarity switch while at full field. The result was a catastrophic failure of the magnet power supply; however, the damage was fully contained. This event did lead to a major experiment downtime, while the power supply was repaired and the incident investigated. Once repaired, though, the experiment was able to resume.

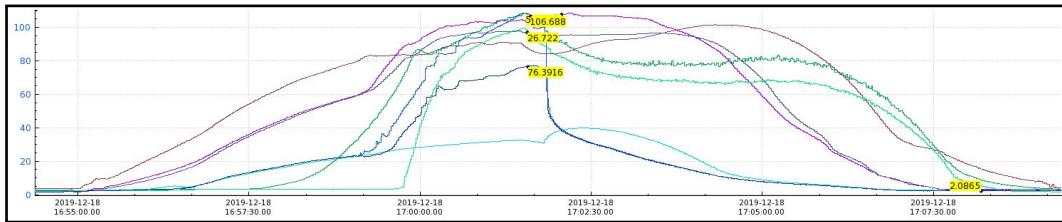


FIG. 59: Temperature change during a target swap in Hall B.

NMR

The NMR system was also a great success. The comprehensive software package ran continuously and reliably to provide polarization data both to on-site operators and remote users. The innovation of digital circuit tuning allows NMR adjustments to be made at-will,

without the need to pause operation for physical access to the experimental hall - with proton signal strength and measurement fidelity commensurate with that expected from external probe coils. The deuteron signal, however, represented a considerable improvement over past systems. The inclusion of the cold tank allowed deuteron polarization to be discernible directly from the raw signal - before even background subtraction.

Though the true polarization of each target sample will not be conclusively known until final analysis is completed, the NMR very effectively served its intended purpose of providing real-time relative polarization results - thereby ensuring that all target parameters were optimized to produce the maximum polarization possible for each run. The incorporated microwave controls meant that target experts could easily adjust for the shifting ideal frequency that occurs with accumulated radiation dose and monitor each cell's performance to determine when an anneal or retirement became necessary. An issue with the deuteron circuit did force the target to undergo a repair in the run. However, the modularity of the target insert enabled the repair to be conducted rapidly and the run continued soon thereafter.

Material and Polarization

The commissioning beam received on June 11 was met with an initial NH₃ polarization of \sim 65%. Unfortunately, it was discovered that not all of the material available for RGC was immediately highly polarizable. As suspected from the final polarization test of ND3, the deuterated samples in particular required a fair amount of beam time to receive the cold dose radiation needed for high quality polarization (see Section 3.1). However, once sufficiently irradiated, most samples performed well - with NH₃ and ND3 reaching NMR measured polarizations of 80% and 50% respectively. Likewise, the beam polarization was excellent - averaging \sim 80%.

Fortunately, the modular design of the target cells made it extremely easy to track each sample's dose and performance. This is compared to previous experiment designs in which material holding cups had to be emptied and refilled to change samples. The APOLLO design also enabled higher scrutiny of beam dose distribution within each cell. Solid ammonia has a well-known phenomenon of irradiation induced color change - shifting from clear white to bright purple. Though this color fades faster than the polarizability effects of dose (meaning it cannot be used as a sample rejection parameter), it does offer a visual confirmation of beam distribution.

Early in the run, it was discovered that target cells were not receiving perfectly uniform

beam exposure. As shown in Fig. 61, they appeared to be at least partly situated lower than the nominal electron beam axis. This was somewhat of a surprise as the beam position had already been assessed during the raster calibration process. During the raster calibration, a cross-hair target cell was placed in the beam line and a low-current beam was then used to correlate detector occupancy to beam position (see Fig. 60).

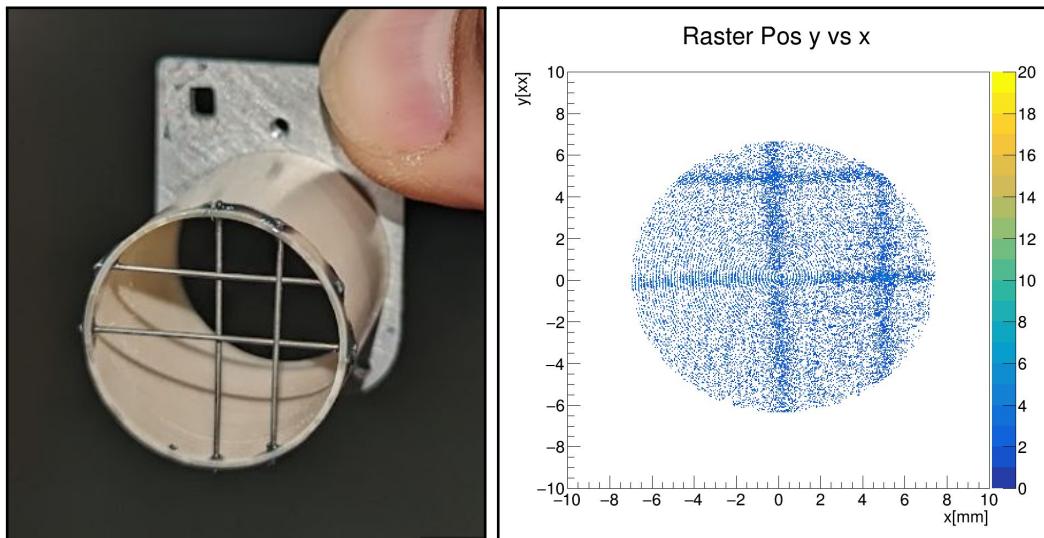


FIG. 60: The raster calibration target (left) and resultant detector response (right) used to commission the rastering system.

This system was implemented to calibrate magnet currents to raster radius, but also it also had the effect of seemingly confirming target cell positioning. The raster calibration cell, however, only provided positioning information for the downstream end of the cell (where the cross-hairs were located). The removable cell window though by which the beam spot could be observed is on the upstream end of the cell. This means that a tilted misalignment of the cells could explain the discrepancy.

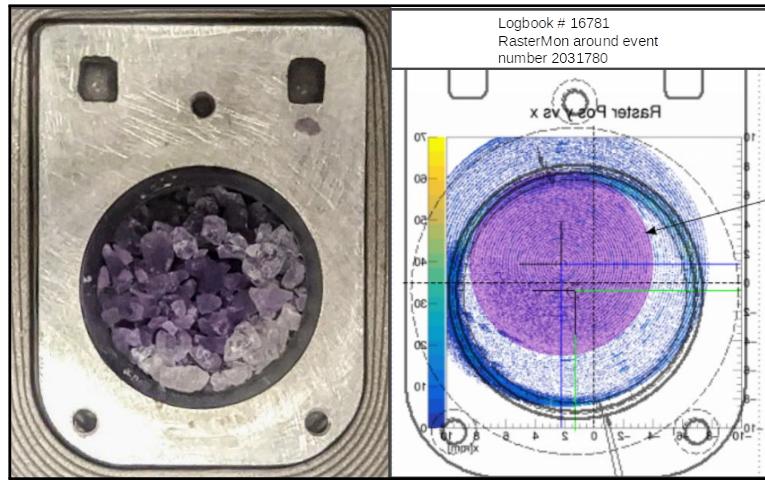


FIG. 61: An off-center purple beam spot on a target sample compared to the raster calibration system measurement.

To investigate this potentiality, an additional optical target cell was created. It utilized the same mapping technique as the raster calibration cell but with an axially offset brass bullseye. As shown in Fig. 62, this study confirmed the misalignment and allowed corrective adjustments to be made. Although this improved the beam distribution, there still remained a small offset. As discussed in Section 4.4 (see Fig. 43), the presence of unrastered material near the NMR coils creates a dose varying discrepancy between the NMR-measured and real target polarization. This problem compounds over the material's lifetime as successive anneals increase the disparity between rastered and unrastered sections. To minimize this effect, it became standard practice to periodically redistribute material within each cell.

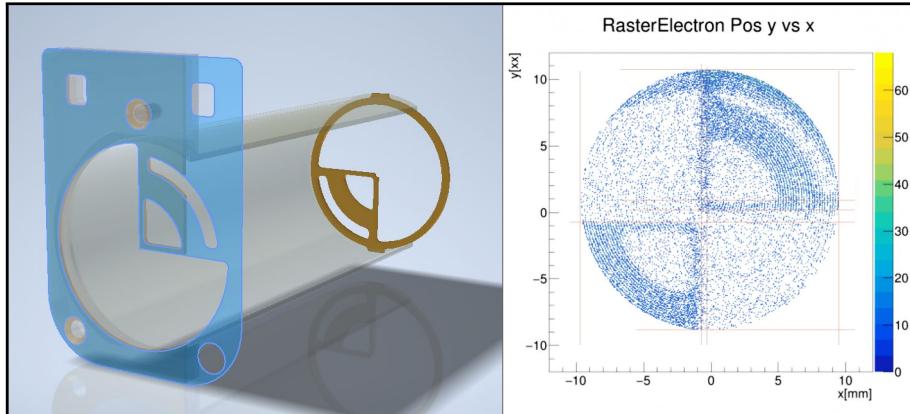


FIG. 62: Optical bullseye target used to assess the roll pitch and yaw of the target cells.

5.3 RGC OUTCOME

RGC ran for a total of 10 calendar months and 95 PAC days representing a full 80% of the originally allotted beam time (see Table 7). Of the time lost, ~ 88 hrs (3.7 PAC days or 2.8%) can be attributed to target malfunction. However, considering the standard method of material exchange in previous experiments would take the greater part of a day to accomplish, the APOLLO target (with its rapid trolley system) provided exceptional beam usage - enabling RGC to collect data on over 45×10^9 scattering events: 20×10^9 in the FT-off configuration and 25×10^9 in the FT-on configuration. Fig. 63 shows the complete data collection history of the RGC experiment with each color change in the histograms represents a change of target type.

Date	Event	Scheduled (Days)	Received (Days)	Downtime (hr)
June 11, 2022	2.2 GeV Commissioning	-	-	-
June 15, 2022	Run Valve Maintenance	-	-	2 [0]
June 15, 2022	Beginning of First FT-on	30	35 [28]	-
June 26, 2022	1st NMR Repair	-	-	86 [86]
Aug 7, 2022	Run Valve Icing	-	-	2 [2]
Aug 26, 2022	Changeover to FT-off	-	-	-
Sept 1, 2022	Beginning of FT-off	53	35 [25.4]	-
Oct 14, 2022	Reversal of Solenoid Field	-	-	-
Nov 2, 2022	Run Valve Icing	-	-	2 [0]
Nov 3, 2022	Run Valve Icing	-	-	3.5 [0]
Nov 11, 2022	Solenoid Failure	-	-	-
Jan 30, 2023	Beginning of Second FT-on	37	25 [25]	-
Jan 31, 2023	RGC Back BEAM ON	-	-	3.5 [0]
Mar 20, 2023	End of Experiment	-	-	-
-	Totals	120	95 [78.4]	99 [88]

TABLE 7: Timeline of RGC showing major events and their impact on run time. Numbers in brackets indicate the effective impact due coincidence with accelerator down time.

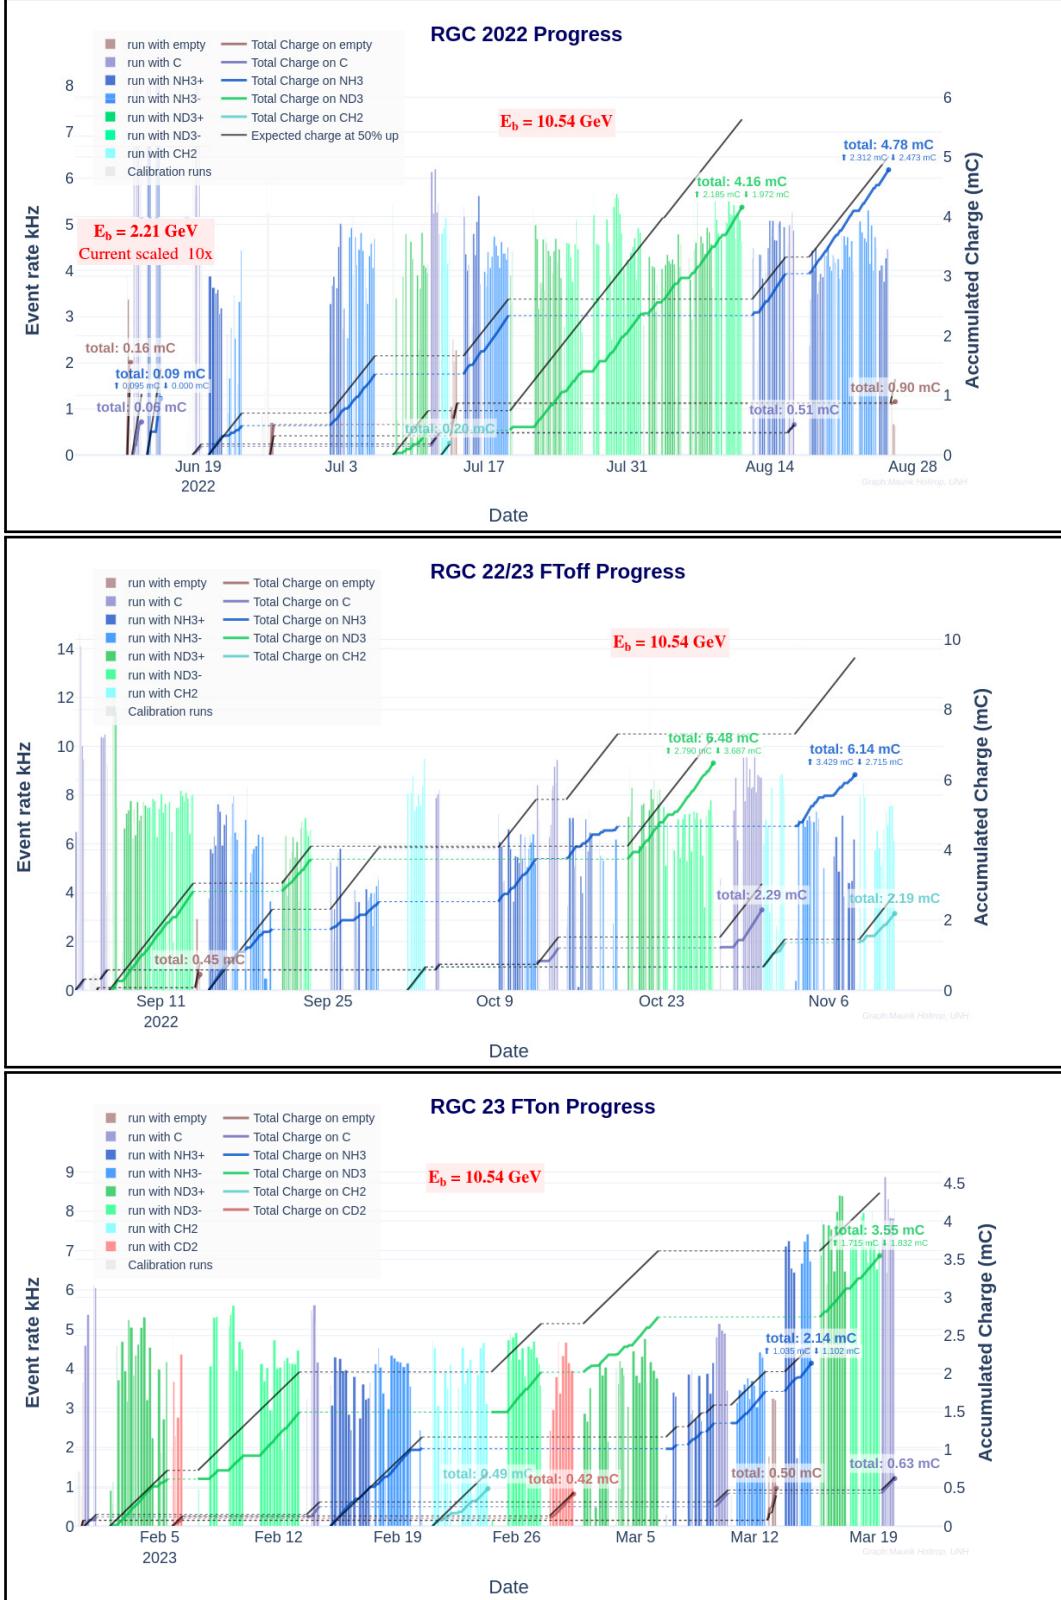


FIG. 63: Events gathered during the RGC experimental run.

5.4 CONCLUSION

Despite some setbacks, the experiment run was a great success. It achieved its goal of gathering high quality polarized nuclear data thanks to the diligent work of the many RGC collaborators. The APOLLO target met its mandate of providing all the elements of DNP: a consistent cryogenic refrigerator, a homogeneous magnetic field, tunable microwave irradiation, and a real time polarization measurement system. The innovative material transport system ensured this all was achieved with minimal downtime while still allowing for frequent material changes. Furthermore all of this was accomplished despite the difficulties introduced by a global pandemic. Although the Run Group C experiment has concluded, the APOLLO target is far from retired. Even while the final NMR characterization studies are being conducted and preliminary data analysis is being presented, the future of APOLLO is already being planned. Thanks to the modularity of the APOLLO system, each component and subsystem can be easily modified while still maintaining the core infrastructure, allowing it to be seamlessly adapted to accommodate future polarized experiments.

For example, one proposal calls for *dual polarization* of both lithium deuteride and ammonia which would require holding the two independently polarized materials in proximity for simultaneously data taking. This experiment would require new shim coil geometry, a new LHe bath and target cells, and additional NMR coils with their corresponding hardware. Fortunately, the shim coil mandrel was designed to be separable from its carrier and the software which calculated the original coil geometry can trivially output new results for different requirements. Likewise, the HTC magnet leads were also implemented such that access and modification does not require altering the cryostat in any way. Alterations to the LHe bath also do not necessitate changes to the overall trolley design or movement mechanism. As for the additions to the NMR, the system already accommodates multiple NMR circuits and (although the coils would need to be developed and the hardware updated) such changes are compatible with the existing technique. Whether or not this specific proposal is pursued, it seems clear that APOLLO will be able to meet the requirements for a wide array of future physics proposals.

Beyond even the potential for scattering experiments, there are also discussions of establishing a polarized material development program at Jefferson Lab which could investigate the mechanisms behind warm and cold dosing, quantify material degradation and annealing processes, characterize materials, and provide high performing polarizable samples for future experiments. This development program will benefit from the experience collected during the APOLLO project.

BIBLIOGRAPHY

- [1] E. Rutherford, The Scattering of Alpha and Beta Particles by Matter and the Structure of the Atom, *Phil. Mag. Ser. 6* **21**, 669–688 (1911).
- [2] M. Gell-Mann, A Schematic Model of Baryons and Mesons, *Phys. Lett.* **8**, 214–215 (1964).
- [3] G. Zweig, An SU(3) Model for Strong Interaction Symmetry and its Breaking. Version 2, in *Developments In The Quark Theory Of Hadrons. Vol. 1. 1964 - 1978*, edited by D. B. Lichtenberg and S. P. Rosen (Feb. 1964), pp. 22–101.
- [4] W. Gerlach and O. Stern, Experimental Proof of the Magnetic Moment of the Silver Atom, *Z. Phys.* **8**, 110–111 (1922).
- [5] A. W. Overhauser, Polarization of Nuclei in Metals, *Phys. Rev.* **92**, 411–415 (1953).
- [6] T. R. Carver and C. P. Slichter, Experimental Verification of the Overhauser Nuclear Polarization Effect, *Phys. Rev.* **102**, 975–980 (1956).
- [7] Y. L. Dokshitzer, Calculation of the Structure Functions for Deep Inelastic Scattering and e+ e- Annihilation by Perturbation Theory in Quantum Chromodynamics, Sov. Phys. JETP **46**, 641–653 (1977).
- [8] V. N. Gribov and L. N. Lipatov, Deep Inelastic e-p Scattering in Perturbation Theory, Sov. J. Nucl. Phys. **15**, 438–450 (1972).
- [9] G. Altarelli and G. Parisi, Asymptotic Freedom in Parton Language, *Nucl. Phys. B* **126**, 298–318 (1977).
- [10] J. Soffer and O. V. Teryaev, Improved Positivity Bound for Deep Inelastic Scattering on Transversely Polarized Nucleon, in Workshop on Polarized Protons at High Energies - Accelerator Challenges and Physics Opportunities (May 1999), pp. 419–421.
- [11] J. Ashman et al. (European Muon), A Measurement of the Spin Asymmetry and Determination of the Structure Function $g(1)$ in Deep Inelastic Muon-Proton Scattering, *Phys. Lett. B* **206**, edited by V. W. Hughes and C. Cavata, 364 (1988).
- [12] V. A. Sulkosky, “The Spin Structure of ${}^3\text{He}$ and the Neutron at Low Q^2 : A Measurement of the Generalized GDH Integrand”, PhD thesis (William-Mary Coll., 2007).
- [13] K. G. Wilson, Nonlagrangian Models of Current Algebra, *Phys. Rev.* **179**, 1499–1512 (1969).

- [14] V. Lagerquist, S. E. Kuhn, and N. Sato, Detailed study of quark-hadron duality in spin structure functions of the proton and neutron, *Phys. Rev. C* **107**, 045201 (2023).
- [15] J. M. Cornwall and R. E. Norton, Current Commutators and Electron Scattering at High Momentum Transfer, *Phys. Rev.* **177**, 2584–2586 (1969).
- [16] E. D. Bloom and F. J. Gilman, Scaling, Duality, and the Behavior of Resonances in Inelastic Electron-Proton Scattering, *Phys. Rev. Lett.* **25**, 1140 (1970).
- [17] E. D. Bloom and F. J. Gilman, Scaling and the Behavior of Nucleon Resonances in Inelastic Electron-Nucleon Scattering, *Phys. Rev. D* **4**, 2901 (1971).
- [18] K. V. Dharmawardane et al. (CLAS), Measurement of the x- and Q^2 -Dependence of the Asymmetry A_1 on the Nucleon, *Phys. Lett. B* **641**, 11–17 (2006).
- [19] Y. Prok et al. (CLAS), Moments of the Spin Structure Functions g_1^p and g_1^d for $0.05 < Q^2 < 3.0 \text{ GeV}^2$, *Phys. Lett. B* **672**, 12–16 (2009).
- [20] P. E. Bosted et al. (CLAS), Quark-Hadron Duality in Spin Structure Functions g_1^p and g_1^d , *Phys. Rev. C* **75**, 035203 (2007).
- [21] N. Guler et al. (CLAS), Precise Determination of the Deuteron Spin Structure at Low to Moderate Q^2 with CLAS and Extraction of the Neutron Contribution, *Phys. Rev. C* **92**, 055201 (2015).
- [22] R. Fersch et al. (CLAS), Determination of the Proton Spin Structure Functions for $0.05 < Q^2 < 5 \text{ GeV}^2$ using CLAS, *Phys. Rev. C* **96**, 065208 (2017).
- [23] B. A. Mecking et al. (CLAS), The CEBAF Large Acceptance Spectrometer (CLAS), *Nucl. Instrum. Meth. A* **503**, 513–553 (2003).
- [24] N. Sato, W. Melnitchouk, S. E. Kuhn, J. J. Ethier, and A. Accardi (Jefferson Lab Angular Momentum), Iterative Monte Carlo Analysis of Spin-Dependent Parton Distributions, *Phys. Rev. D* **93**, 074005 (2016).
- [25] F. E. Close and A. W. Thomas, The Spin and Flavor Dependence of Parton Distribution Functions, *Phys. Lett. B* **212**, 227–230 (1988).
- [26] O. Nachtmann, Positivity Constraints for Anomalous Dimensions, *Nucl. Phys. B* **63**, 237–247 (1973).
- [27] H. Georgi and H. D. Politzer, Freedom at Moderate Energies: Masses in Color Dynamics, *Phys. Rev. D* **14**, 1829 (1976).

- [28] A. De Rujula, H. Georgi, and H. D. Politzer, Trouble with Xi Scaling?, *Phys. Rev. D* **15**, 2495 (1977).
- [29] T. Liu, R. S. Sufian, G. F. de Téramond, H. G. Dosch, S. J. Brodsky, and A. Deur, Unified Description of Polarized and Unpolarized Quark Distributions in the Proton, *Phys. Rev. Lett.* **124**, 082003 (2020).
- [30] R. C. Walker et al., Measurements of the Proton Elastic Form-Factors for $1 \leq Q^2 \leq 3$ GeV/c² at SLAC, *Phys. Rev. D* **49**, 5671–5689 (1994).
- [31] G. Baum et al., Measurement of Asymmetry in Spin Dependent e-p Resonance Region Scattering, *Phys. Rev. Lett.* **45**, edited by V. W. Hughes and C. Cavata, 2000 (1980).
- [32] G. Baum et al., A New Measurement of Deep Inelastic e-p Asymmetries, *Phys. Rev. Lett.* **51**, edited by V. W. Hughes and C. Cavata, 1135 (1983).
- [33] P. L. Anthony et al. (E142), Determination of the Neutron Spin Structure Function, *Phys. Rev. Lett.* **71**, edited by V. W. Hughes and C. Cavata, 959–962 (1993).
- [34] K. Abe et al. (E143), Measurements of the Proton and Deuteron Spin Structure Functions g_1 and g_2 , *Phys. Rev. D* **58**, 112003 (1998).
- [35] K. Abe et al. (E154), Precision Determination of the Neutron Spin Structure Function g_1^n , *Phys. Rev. Lett.* **79**, edited by B. Frois, V. W. Hughes, and N. De Groot, 26–30 (1997).
- [36] P. L. Anthony et al. (E155), Measurements of the Q^2 Dependence of the Proton and Neutron Spin Structure Functions g_1^p and g_1^n , *Phys. Lett. B* **493**, 19–28 (2000).
- [37] A. Airapetian et al. (HERMES), Evidence for Quark Hadron Duality in the Proton Spin Asymmetry A_1 , *Phys. Rev. Lett.* **90**, 092002 (2003).
- [38] B. Adeva et al. (Spin Muon), Measurement of the Spin Dependent Structure Function $g_1(x)$ of the Deuteron, *Phys. Lett. B* **302**, edited by V. W. Hughes and C. Cavata, 533–539 (1993).
- [39] V. Y. Alexakhin et al. (COMPASS), The Deuteron Spin-dependent Structure Function g_1^d and its First Moment, *Phys. Lett. B* **647**, 8–17 (2007).
- [40] J. Yun et al. (CLAS), Measurement of Inclusive Spin Structure Functions of the Deuteron, *Phys. Rev. C* **67**, 055204 (2003).
- [41] Z. E. Meziani et al., Higher Twists and Color Polarizabilities in the Neutron, *Phys. Lett. B* **613**, 148–153 (2005).

- [42] X. Zheng et al. (Jefferson Lab Hall A), Precision Measurement of the Neutron Spin Asymmetries and Spin-Dependent Structure Functions in the Valence Quark Region, *Phys. Rev. C* **70**, 065207 (2004).
- [43] X. Zheng et al. (Jefferson Lab Hall A), Precision Measurement of the Neutron Spin Asymmetry A_1^n and Spin Flavor Decomposition in the Valence Quark Region, *Phys. Rev. Lett.* **92**, 012004 (2004).
- [44] P. Solvignon et al. (Jefferson Lab E01-012), Quark-Hadron Duality in Neutron (He-3) Spin Structure, *Phys. Rev. Lett.* **101**, 182502 (2008).
- [45] D. S. Parno et al. (Jefferson Lab Hall A), Precision Measurements of A_1^n in the Deep Inelastic Regime, *Phys. Lett. B* **744**, 309–314 (2015).
- [46] D. Flay et al. (Jefferson Lab Hall A), Measurements of d_2^n and A_1^n : Probing the Neutron Spin Structure, *Phys. Rev. D* **94**, 052003 (2016).
- [47] R. G. Fersch (CLAS), Quark-Hadron Duality of Spin Structure Functions in CLAS EG1b Data, *Few Body Syst.* **59**, edited by R. Gothe, Y. Ilieva, V. Mokeev, E. Santopinto, and S. Strauch, 108 (2018).
- [48] F. R. Wesselmann et al. (RSS), Proton Spin Structure in the Resonance Region, *Phys. Rev. Lett.* **98**, 132003 (2007).
- [49] M. Thorsten et al., Dynamic Nuclear Polarization at High Magnetic Fields, eng, *The Journal of Chemical Physics* **128**, 052211–052211 (2008).
- [50] Crabb, D. G. and Meyer, W., Solid Polarized Targets for Nuclear and Particle Physics Experiments, *Annual Review of Nuclear and Particle Science* **47**, 67–109 (1997).
- [51] J. D. Maxwell, “Probing Proton Spin Structure: A Measurement of g_2 at Four-momentum Transfer of 2 to 6 GeV²”, PhD thesis (University of Virginia, 2017).
- [52] W. Meyer, Ammonia as a Polarized Solid Target Material: A Review, *Nucl. Instrum. Meth. A* **526**, edited by S. Goertz, W. Meyer, and G. Reicherz, 12–21 (2004).
- [53] P. M. McKee, Observations of radiation damage and recovery in ammonia targets, *Nucl. Instrum. Meth. A* **526**, 60–64 (2004).
- [54] M. L. Seely, “Dynamic Nuclear Polarization of Irradiated Target Materials”, PhD thesis (Yale, Dec. 1982).
- [55] J. D. Maxwell et al., Design and Performance of the Spin Asymmetries of the Nucleon Experiment, *Nucl. Instrum. Meth. A* **885**, 145–159 (2018).

- [56] C. W. Leemann, D. R. Douglas, and G. A. Krafft, The Continuous Electron Beam Accelerator Facility: CEBAF at the Jefferson Laboratory, *Ann. Rev. Nucl. Part. Sci.* **51**, 413–450 (2001).
- [57] V. D. Burkert et al., The CLAS12 Spectrometer at Jefferson Laboratory, *Nucl. Instrum. Meth. A* **959**, 163419 (2020).
- [58] R. Fair et al., The CLAS12 Superconducting Magnets, *Nucl. Instrum. Meth. A* **962**, 163578 (2020).
- [59] M. A. Antonioli et al., The CLAS12 Silicon Vertex Tracker, *Nucl. Instrum. Meth. A* **962**, 163701 (2020).
- [60] A. Acker et al., The CLAS12 Micromegas Vertex Tracker, *Nucl. Instrum. Meth. A* **957**, 163423 (2020).
- [61] D. S. Carman, G. Asryan, V. Baturin, L. Clark, R. De Vita, W. Kim, B. Miller, and C. Wiggins, The CLAS12 Central Time-of-Flight System, *Nucl. Instrum. Meth. A* **960**, 163626 (2020).
- [62] P. Chatagnon et al., The CLAS12 Central Neutron Detector, *Nucl. Instrum. Meth. A* **959**, 163441 (2020).
- [63] M. D. Mestayer et al., The CLAS12 Drift Chamber System, *Nucl. Instrum. Meth. A* **959**, 163518 (2020).
- [64] M. Ungaro et al., The CLAS12 Low Threshold Cherenkov Detector, *Nucl. Instrum. Meth. A* **957**, 163420 (2020).
- [65] G. Adams et al., The CLAS Cherenkov Detector, *Nucl. Instrum. Meth. A* **465**, 414–427 (2001).
- [66] Y. G. Sharabian et al., The CLAS12 High Threshold Cherenkov Counter, *Nucl. Instrum. Meth. A* **968**, 163824 (2020).
- [67] D. S. Carman et al., The CLAS12 Forward Time-of-Flight System, *Nucl. Instrum. Meth. A* **960**, 163629 (2020).
- [68] M. Contalbrigo et al., The CLAS12 Ring Imaging Cherenkov Detector, *Nucl. Instrum. Meth. A* **964**, 163791 (2020).
- [69] G. Asryan et al., The CLAS12 Forward Electromagnetic Calorimeter, *Nucl. Instrum. Meth. A* **959**, 163425 (2020).

- [70] G. Asryan et al., The CLAS12 Forward Electromagnetic Calorimeter, *Nucl. Instrum. Meth. A* **959**, 163425 (2020).
- [71] A. Acker et al., The CLAS12 Forward Tagger, *Nucl. Instrum. Meth. A* **959**, 163475 (2020).
- [72] V. Lagerquist (CLAS12 Polarized Target), Magnetic Field Requirements for the CLAS12 Polarized Target, *PoS PSTP2019*, edited by J. Pierce, L. Crow, C. Jiang, F. Li, L. Broussar, C. Keith, D. Gaskel, and M. Poelke, 055 (2020).
- [73] F. Pobell, *Matter and methods at low temperatures*, en, 3rd ed. (Springer, Berlin, Germany, 2006).
- [74] C. Keith et al., The Jefferson Lab Frozen Spin Target, *Nucl. Instrum. Meth. A* **684**, 27–35 (2012).

APPENDIX A

SOLENOID MAPPING

The measured data used for this optimization were taken along the full length of the solenoid both on the axis and at a radial position of 30 cm. The assumed cylindrical symmetry of the coils negates any potential azimuthal field components (B_ϕ) as well as any on-axis radial field components (B_r). This leaves only the B_z (at $r = 0$ cm), B_z (at $r = 30$ cm), and B_r (at $r = 30$ cm) data to be compared. Additionally, the data were taken at a reduced field strength (due to the limitations of the probes). To compensate, both models have been scaled to match the central maximum field.

On-Axis

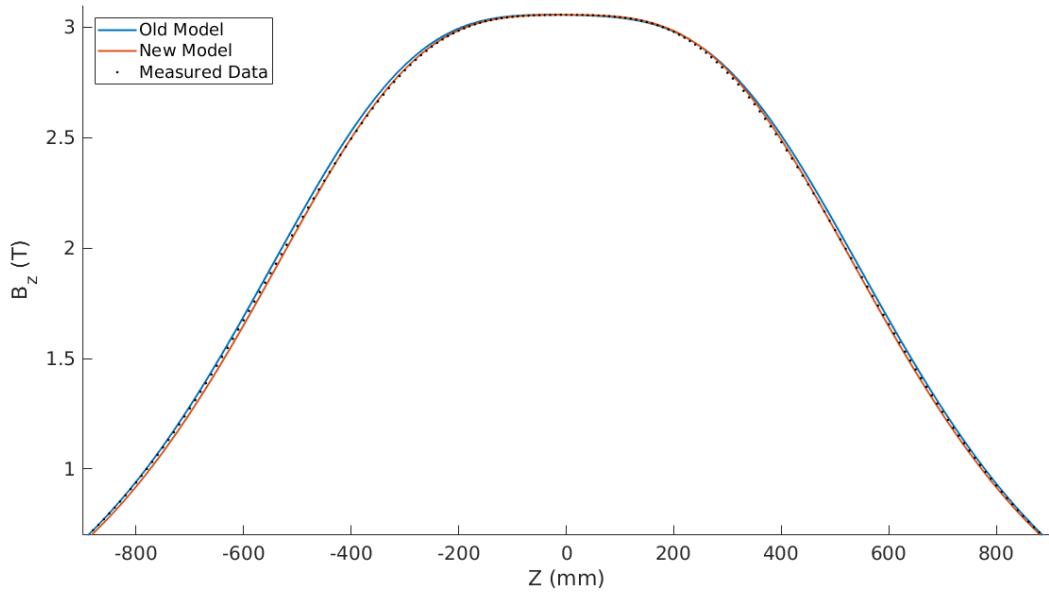


FIG. 64: On-Axis Field Map Full Length

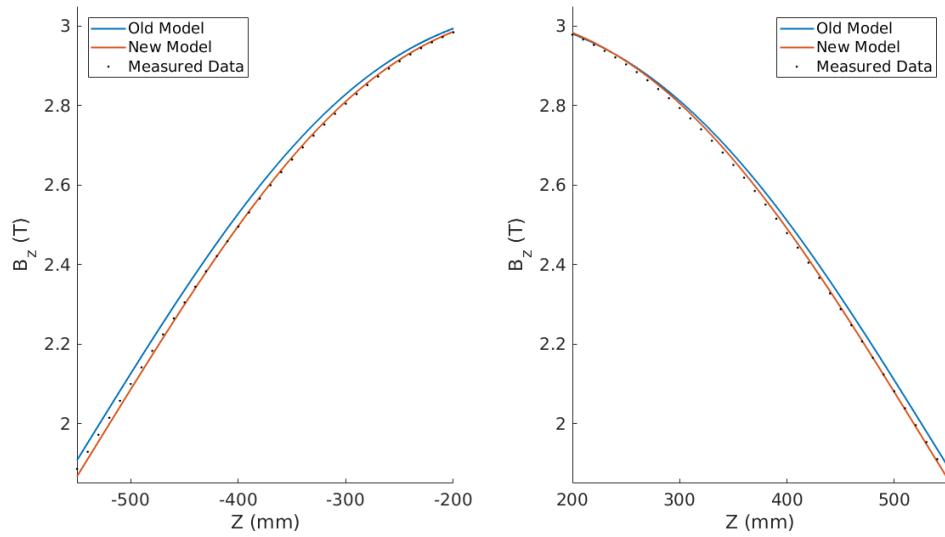


FIG. 65: On-Axis Field Map Side Regions

Off-Axis

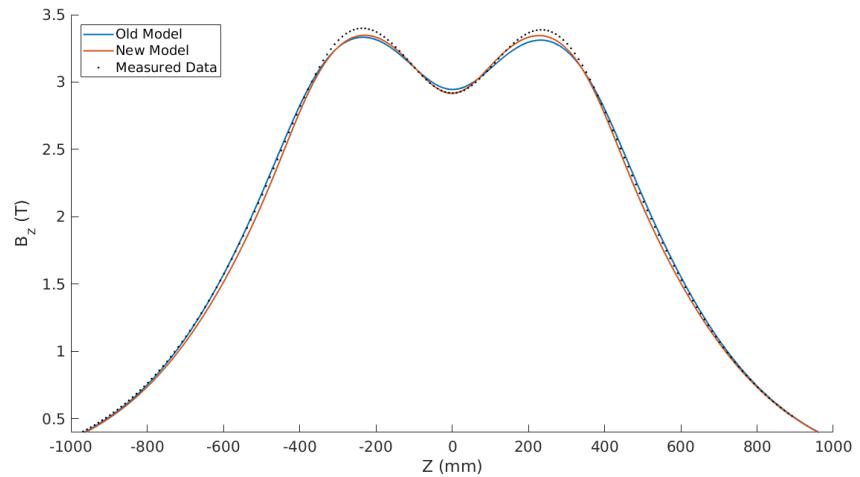


FIG. 66: 1.25 cm Off-Axis Field Map Full Length

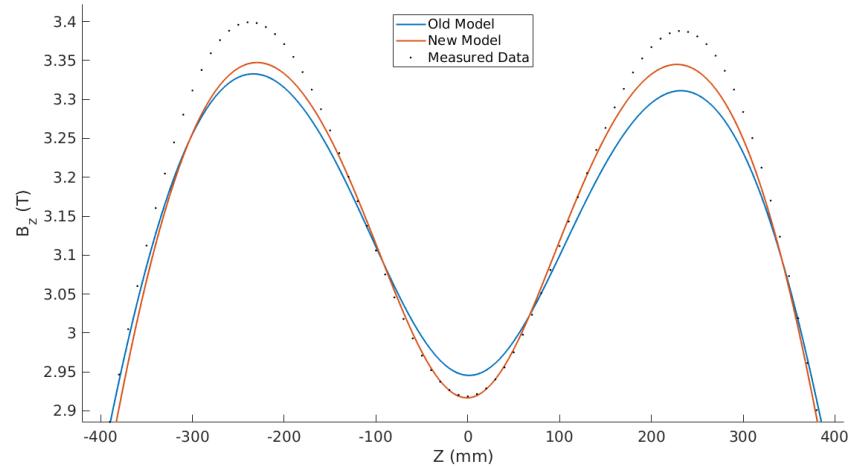


FIG. 67: 30 cm Off-Axis Field Map Central Region

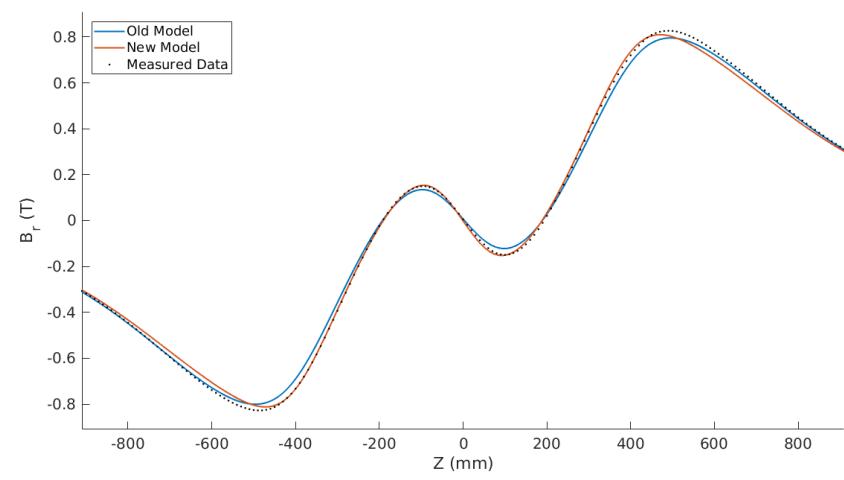


FIG. 68: Off-Axis Field Map Radial Component

VITA

Victoria Lagerquist
 Department of Physics
 Old Dominion University
 Norfolk, VA 23529

Education:

M.S. Physics, Old Dominion University, Norfolk, VA. December 2017
 B.S. Physics, Old Dominion University, Norfolk, VA. May 2015

Selected Publications:

- V. Lagerquist, S. E. Kuhn, and N. Sato. Detailed Study of Quark-Hadron Duality in Spin Structure Functions of the Proton and Neutron. *Phys. Rev. C* **107**, 045201 (2023).
- Lagerquist, V. Magnetic Field Requirements for the CLAS12 Polarized Target. *PoS PSTP2019* (eds Pierce, J. *et al.*) 055 (2020).
- Fair, R. *et al.* The CLAS12 superconducting magnets. *Nucl. Instrum. Meth. A* **962**, 163578 (2020).
- Fersch, R. *et al.* Determination of the Proton Spin Structure Functions for $0.05 < Q^2 < 5 \text{ GeV}^2$ using CLAS. *Phys. Rev. C* **96**, 065208. arXiv: 1706.10289 [nucl-ex] (2017).