

Heavy Photon Search Experiment at Jefferson Laboratory: proposal for 2014-2015 run

P. Hansson Adrian, C. Field, N. Graf, M. Graham, G. Haller,
R. Herbst, J. Jaros*[†], T. Maruyama, J. McCormick, K. Moffeit,
T. Nelson, H. Neal, A. Odian, M. Oriunno, S. Uemura, D. Walz
SLAC National Accelerator Laboratory, Menlo Park, CA 94025

A. Grillo, V. Fadeyev, O. Moreno
University of California, Santa Cruz, CA 95064

W. Cooper
Fermi National Accelerator Laboratory, Batavia, IL 60510-5011

S. Boyarinov, V. Burkert, C. Cuevas, A. Deur, H. Egiyan, L. Elouadrhiri,
A. Freyberger, F.-X. Girod, S. Kaneta, V. Kubarovsky, N. Nganga, B.
Raydo, Y. Sharabian, S. Stepanyan[†], M. Ungaro, B. Wojtsekhowski
Thomas Jefferson National Accelerator Facility, Newport News, Virginia 23606

R. Essig
Stony Brook University, Stony Brook, NY 11794-3800

M. Holtrop[†], K. Slifer, S. K. Phillips
University of New Hampshire, Department of Physics, Durham, NH 03824

R. Dupre, M. Guidal, S. Niccolai, E. Raully, and P. Rosier
Institut de Physique Nucleaire d'Orsay, IN2P3, BP 1, 91406 Orsay, France

D. Sokhan
School of Physics & Astronomy, University of Glasgow, Glasgow, G12 8QQ, Scotland, UK

*Contact person

[†]Co-spokesperson

P. Schuster, N. Toro

Perimeter Institute, Ontario, Canada N2L 2Y5

N. Dashyan, N. Gevorgyan, R. Paremuzyan, H. Voskanyan

Yerevan Physics Institute, 375036 Yerevan, Armenia

M. Khandaker

Idaho State University, Pocatello, Idaho 83209

M. Battaglieri, R. De Vita

Istituto Nazionale di Fisica Nucleare, Sezione di Genova e

Dipartimento di Fisica dell'Università, 16146 Genova, Italy

S. Bueltmann, L. Weinstein

Old Dominion University, Norfolk, Virginia 23529

G. Ron

Hebrew University of Jerusalem, Jerusalem, Israel

A. Kubarovsky

University of Connecticut, Department of Physics, Storrs, CT 06269

K. Griffioen

The College of William and Mary, Department of Physics, Williamsburg, VA 23185

Y. Gershtein, J. Reichert

Rutgers University, Department of Physics and Astronomy, Piscataway, NJ 08854

(Dated: January 23, 2013)

ABSTRACT

The Heavy Photon Search (HPS) is an experiment proposed for Jefferson Laboratory to search for new heavy vector boson(s), aka "heavy photons" or "dark photons" or "hidden sector photons", in the mass range of 20 MeV/c² to 1000 MeV/c². Such particles will arise if there are additional U(1) gauge bosons in nature, and they will couple, albeit weakly, to electric charge through kinetic mixing. Many BSM theories predict the existence of additional U(1)'s, and recent observations of high energy electrons and positrons in the cosmic rays may be the result of primordial dark matter annihilating into heavy photons. HPS searches for electro-produced heavy photons using both invariant mass and separated decay vertex signatures using a compact, large acceptance forward spectrometer. The first stage of HPS, the HPS Test Run, was approved by the Jefferson Lab PAC37 on January 14, 2011, after which it was proposed to DOE HEP for funding and approved and funded by Summer 2011. The Test Run was built in 2011-2012, and installed and run at JLAB in Spring, 2012. PAC39 reviewed HPS in June, 2012, and on the basis of the successful run, granted it an "A" rating, a commissioning run with electron beams, and approval (C1) to proceed to the full experiment contingent only on final approval from JLAB management. This proposal describes the second stage of our program, the full HPS experiment, which is capable of searching for heavy photons over a wide and uncharted region in parameter space and discovering true muonium, the QED $\mu^+\mu^-$ atom. It reviews the scientific motivation of HPS; describes the proposed apparatus, data acquisition system, and trigger; reviews the results of the HPS Test Run which demonstrated that the experiment is technically feasible and that electron beam backgrounds are well-understood; summarizes expected performance and experimental reach; and concludes with budget and schedule information and a proposed run plan. Schedules at Jefferson Laboratory admit time for HPS to be commissioned in Hall B in the Fall of 2014 and to take data beginning in Spring 2015. With final approval from JLAB management, and with timely DOE approval and funding, the HPS Collaboration can design, build, test, and commission HPS in time to take advantage of these scheduling opportunities, and begin in earnest its search for spectacular new physics at the Intensity Frontier.

Contents

1. Introduction	7
2. Motivations for Searching for Heavy Photons	10
2.1. Theory Update	11
2.1.1. Heavy Photons and Dark Matter	12
2.1.2. Heavy Photons and Muon $g - 2$	16
2.2. Update on Experimental Status	16
2.3. HPS physics with True Muonium	17
2.4. HPS Searches for Hidden Sectors	18
3. Proposed Measurements	19
3.1. Search for the heavy photon	20
3.2. Search for true muonium	23
3.3. Other searches for hidden sector particles	24
4. Description of the HPS setup	26
4.1. Overview	26
4.2. HPS beamline	27
4.2.1. Layout of the HPS setup	28
4.2.2. Running Conditions	32
4.2.3. Beam Diagnostics	34
4.2.4. Vacuum chambers	35
4.2.5. Beam dumps and shieldings	37
4.2.6. Targets	38
4.3. Silicon vertex tracker	40
4.3.1. Layout	41
4.3.2. Module Design	42
4.3.3. Mechanical Support, Cooling and Services	46
4.4. Electromagnetic Calorimeter	48
4.4.1. Improvements to the existing calorimeter	50

4.4.2. Possibilities with new APDs	53
4.5. Muon system	55
4.5.1. Conceptual Design	57
4.6. Trigger and DAQ	61
4.6.1. SVT Data Acquisition	62
4.6.2. ECal and Muon Detector FADC Readout	65
4.6.3. Trigger System	68
4.6.4. Event Size and Data Rates	74
4.7. Offline Computing Model	75
5. May-2012 Test Run	77
5.1. HPS Test Run Apparatus	79
5.1.1. Test Run SVT	80
5.1.2. Test Run ECal	82
5.1.3. Test Run Data Acquisition	84
5.1.4. Test Run Trigger System	86
5.2. Multiple Coulomb Scattering Measurement	87
5.3. Test Run Apparatus Performance	92
5.3.1. SVT Performance	92
5.3.2. ECal performance	96
5.3.3. Trigger performance	97
6. HPS Performance Studies	100
6.1. Simulation of Backgrounds and Detector Occupancies	101
6.1.1. Simulation of Backgrounds	101
6.1.2. Simulated Tracker Occupancies	104
6.1.3. Simulated ECal Occupancies	107
6.2. ECal Trigger Rates	108
6.3. Muon system trigger rates	112
6.4. Track Reconstruction	115
6.5. Tracking Efficiency, Pattern Recognition and Fake Rates	116
6.6. Track Momentum and Spatial Resolution	116

7. Experimental Reach	123
7.1. Resonance Search	125
7.2. Displaced Vertex and Resonance Search	127
7.3. Reach in Mass-Coupling Parameter Space	128
8. Run Plan and Beam Time Request	129
9. Schedule and Cost Baseline	131
9.1. Cost	132
9.2. Schedule	135
9.3. Manpower	137
A. Simulation tools	144
B. Test Run SVT Performance	147
C. Test Run ECal Calibration	154
D. Production and Decay of the A'	155
References	157

1 Introduction

Access to higher and higher luminosities and ever faster detection and recording techniques enables searches for new physics at otherwise well-explored energies. This fundamental premise of Intensity Frontier physics has already seen dramatic demonstration at the e^+e^- B factories, where high luminosities and impressive data handling capacities have allowed extensive exploration of CP violation in the quark sector. The same principle is being exploited in new proposals to explore neutrino masses, mixings, and CP violation by directing ever more intense neutrino beams at massive detectors to push sensitivity well beyond present limits. At the Intensity Frontier, searches for new physics often rely on the study of rare processes and the search for subtle effects which would indirectly indicate physics beyond the Standard Model. But this is not the rule. New studies of otherwise commonplace phenomena at electron machines, like trident production off heavy nuclear targets, can, with sufficient sensitivity, explore whole new worlds and directly search for hidden sector particles and forces, those with very weak couplings to our Standard Model world. The Heavy Photon Search at Jefferson Laboratory does exactly this, utilizing the high duty factor CEBAF accelerator, intense beams, fast detectors, electronics and triggering, and state of the art data acquisition to explore a very common landscape in search of a most uncommon quarry.

Heavy photons, or dark or hidden sector photons, may well be part of our universe and related to the Dark Matter. Particles of dark matter, which interact very weakly with normal matter and account for a quarter of the universal mass-energy, are of course not yet detected. The Dark Matter can be thought of as constituting, or inhabiting, a hidden sector, since it interacts so weakly with normal, baryonic matter. This sector could include a complex of new forces and particles with which we barely interact. Stimulated by the observation of very high energy electrons and positrons in the cosmic rays and the difficulty of understanding their production in terms of tried and true SUSY dark matter annihilation, several authors realized that models in which massive dark matter particles annihilate to heavy photons, which in turn decay to high energy electron-positron pairs, could naturally account for the observations. These theories presume heavy photons couple to the dark matter, mediate its interactions, are produced in its annihilation, and weakly couple to electric charge. Heavy

photons in the mass range of 100 to 1000 MeV can reasonably account for the observed cosmic ray fluxes.

Many Beyond Standard Model theories generate extra $U(1)$ gauge groups, and the associated gauge bosons could have masses over a very wide range. As Holdum realized in the mid 80s, it is natural that such heavy photons kinetically mix with our own photon, leading to their induced coupling to electric charge. This mixing can be mediated by GUT level particles which carry both Standard Model hypercharge and its hidden sector analogue. Interestingly, the natural scale for this mixing results in heavy photons coupling to Standard model charged particles with couplings of order $10^{-3}e$. So heavy photons naturally couple to electrons, albeit with couplings much suppressed compared to those in standard QED. It follows that electrons will radiate heavy photons, and heavy photons will decay to electron-positron pairs or pairs of other kinematically accessible charged particles, but at rates significantly below QED trident production, and with lifetimes far longer than those expected from purely electromagnetic interactions.

HPS distinguishes heavy photons from the copious background of QED tridents by using both invariant mass and decay length signatures. With good mass resolution, heavy photons will appear as sharp resonances above the QED continuum. For suitable values of mass and coupling, heavy photons will have long lifetimes, resulting in discernible secondary decay vertices. The Heavy Photon Search employs a large acceptance forward magnetic spectrometer with precise momentum measurement and vertexing capability, followed by a highly segmented crystal Electromagnetic Calorimeter for fast triggering and electron identification. HPS depends on the 40 MHz readout capability of the silicon microstrip vertex tracker, 250 MHz FADC readout of the electromagnetic calorimeter, and very high rate triggering and data acquisition systems, to fully exploit CEBAF's essentially DC beams and high intensities. A muon identification system just downstream of the ECal significantly boosts the experimental reach for heavy photon masses above the dimuon threshold and provides an independent trigger. The beam is transported in vacuum through the entire apparatus to eliminate beam gas backgrounds; and the apparatus is split top-bottom, to avoid electrons which have multiple Coulomb scattered or radiated in the target.

HPS probes a unique region of the mass-coupling parameter space where the heavy photon signal would be lost in the trident background without the vertex signature, and it simultaneously accesses a region at higher coupling strength by relying on bump hunting

alone. HPS has sensitivity to a region of parameter space favored by accounting for the discrepancy between measured and calculated values for the muons $g-2$ with the existence of a heavy photon, and probes an extensive region suggested by parameters which could account for dark matter annihilations into heavy photons. In broader terms, HPS searches for heavy photons in a region suggested on very general theoretical grounds. As seen above, coupling strengths of order $10^{-3}e$ are theoretically natural; masses of order αm_W are also expected on general grounds. Interestingly, HPS is also sensitive to the production of true muonium, the QED atom comprised of $\mu^+\mu^-$, which is produced with a well-defined (and detectable) cross-section, and decays with a well-defined (and observable) lifetime to e^+e^- . HPS should discover true muonium, measure some of its properties, and find it a useful calibration signal.

This proposal seeks funding for the Heavy Photon Search (HPS) Experiment at Thomas Jefferson National Accelerator Facility. This experiment is the second stage of a program that was initiated with the Heavy Photon Search Test Run Proposal [1, 2], which was approved by the Jefferson Laboratory Program Advisory Committee PAC37 in January, 2011, and approved and funded by DOE HEP in the late Spring of 2011. PAC37 also conditionally approved the full experiment, contingent upon the Test Run results. During the remainder of FY2011 and the first half of FY2012, the Test Run apparatus, data acquisition system, and system software were designed, constructed, and tested. On April 19, 2012, the newly constituted HPS Collaboration installed the experiment in Jefferson Labs Hall B experimental area, and began commissioning the experiment parasitically, using the HDIce photon beam. Although the Jefferson Lab schedule did not accommodate the electron beam running which had been requested, the apparatus was fully commissioned by running parasitically in the photon beam. The trigger and data acquisition and storage systems worked well, and all systems performed as expected. Efficient track reconstruction in the Silicon Vertex Tracker was demonstrated, measurements of shower energies and positions were made in the Electromagnetic Calorimeter, and critical assumptions about background rates were tested. The critical test run goals were accomplished. A status report summarizing HPS progress and results was submitted to PAC39 [3] along with a request for unconditional approval for the full experiment. At its June, 2012 meeting, PAC39 graded HPS physics with an "A", approved a commissioning run with electrons, and granted us a so-called "C1" approval, which gave Jefferson Laboratory management the final say in granting HPS the

running time needed to search comprehensively for Heavy Photons. Since that approval, it has become clear that running time will become available in Hall B in late calendar 2014 for our commissioning run, when the upgraded CEBAF 12 will have been completed, commissioned, and operational. CLAS12 is the large general purpose apparatus being constructed to exploit CEBAF 12 in Hall B. Delays in funding will delay the construction of the CLAS12 magnets and will delay CLAS12 installation beyond 2015, thereby providing HPS the opportunity for a commissioning run in 2014 and the extended data collection run in 2015. To take full advantage of these scheduling windfalls, the HPS Collaboration has re-visited the original HPS design, and simplified and improved it. The resulting simplifications make it possible to construct and test HPS in time for installation in late 2014. The resulting improvements extend the reach far beyond that of the Test Run experiment, maximize the physics output during this time period, and let HPS begin searching for heavy photons in a large and hitherto unexplored region of parameter space.

In the following, this proposal motivates and describes the new HPS Experiment, documents the experience and performance obtained with the Test Run Apparatus, demonstrates that the backgrounds expected in electron running are understood and manageable, reviews the performance and physics reach of the new experiment, and outlines the budget, schedule, and milestones for constructing and deploying it. It concludes with a request for beam time.

2 Motivations for Searching for Heavy Photons

HPS will search for heavy photons, called A' 's, which are new hypothesized massive vector bosons that have a small coupling to electrically charged matter, including electrons. The existence of an A' is theoretically natural and could explain the discrepancy between the measured and observed anomalous magnetic moment of the muon and several intriguing dark matter-related anomalies. As discussed in the following section, HPS should also have the capability to make the first detection of *True Muonium*, a bound state of a $\mu^+ - \mu^-$ pair predicted by Quantum Electrodynamics (QED). The search for A' 's has generated enormous interest in the international physics community. This is evidenced, for example, by its inclusion in the recent Intensity Frontier Workshop [4, 5], many novel searches in colliding beam and fixed-target data (see [6] for a recent summary of results), and by numerous

new experiments (in addition to HPS) proposed to search for them, including APEX [7, 8], MAMI [9], and DarkLight [10]. We briefly review the theory and motivation for heavy photons and existing constraints on A' .

2.1 Theory Update

The A' is a new abelian $U(1)$ gauge boson with a weak coupling to electrically charged particles induced by “kinetic mixing” with the photon [11, 12]. Kinetic mixing produces an effective parity-conserving interaction $\epsilon e A'_\mu J_{EM}^\mu$ of the A' to the electromagnetic current J_{EM}^μ , suppressed relative to the electron charge e by the parameter ϵ , which can naturally be in the range $10^{-12} - 10^{-2}$ [13–16].

More broadly, “kinetic mixing” of the photon with new forces offers one of the few portals with which ordinary matter can be used to search for light new forces beyond the Standard Model consistent with known symmetries. An A' would also allow ordinary matter to have a small coupling to new particles in a “hidden sector” that do not interact with the Standard Model’s strong, weak, or electromagnetic forces. There has been intense speculation over the past three decades about the existence of hidden sectors. Theoretical models with dark matter, supersymmetry, and string theory constructions often employ hidden sectors with new particle content to resolve various phenomenological questions [17–21] (see [4] for a recent review). The photon mixing with an A' could provide the only non-gravitational window into their existence.

While loop level effects can naturally generate ϵ in an observable range, simple theory arguments offer less guidance for what range of A' mass to search for. Many mass generating mechanisms have been proposed – A' masses can arise, for example, via the Higgs mechanism as in the models of [22–25], or via a Stueckelberg mechanism, as often occurs in large volume string compactification models [4]. In models using a Higgs mechanism, a natural mass range for an A' is near (but beneath) the weak scale, in the MeV to GeV range. This mass range has received considerable attention in part because it may also allow A' s to resolve several anomalies (see below). Existing constraints are shown in Fig. 1. HPS will be sensitive to A' masses in between 20–1000 MeV.

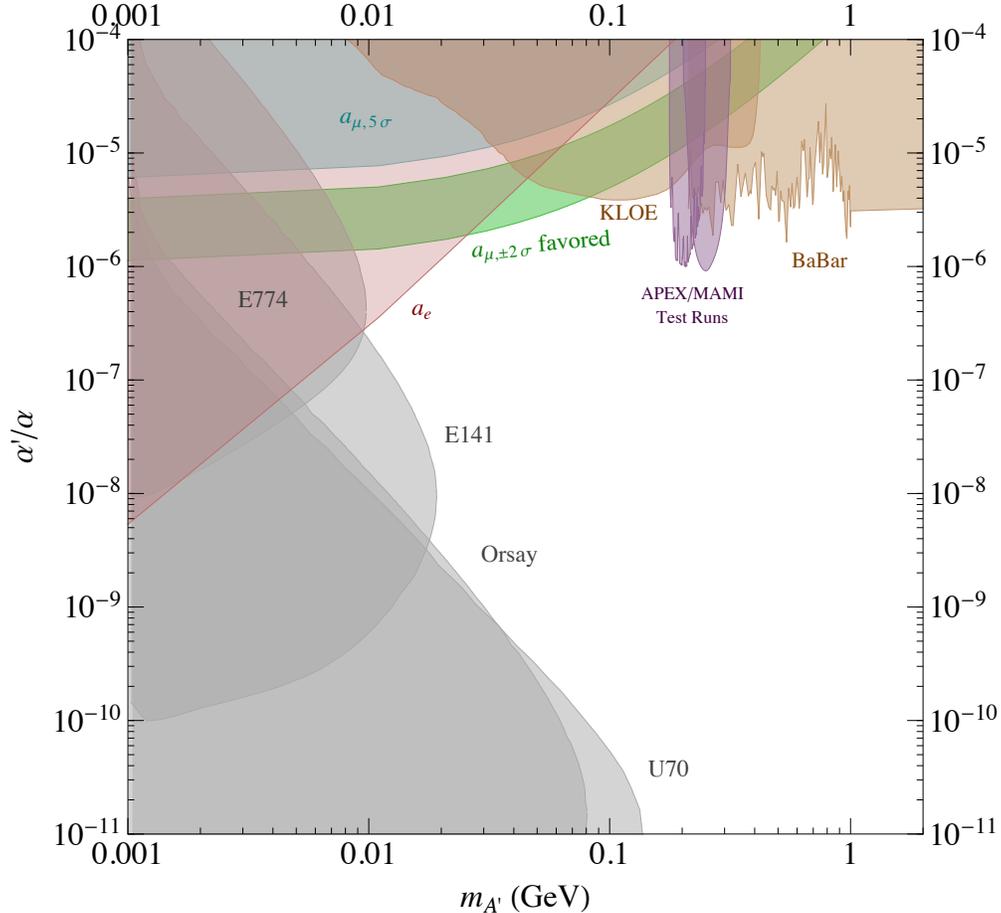


FIG. 1: Existing constraints on heavy photons (A'). Shown are existing 90% confidence level limits from the beam dump experiments E141, E774, Orsay, and U70 [26–29, 32, 34, 35], the muon anomalous magnetic moment a_μ [36], KLOE [37], the test run results reported by APEX [8] and MAMI [9], an updated estimate using a BaBar result [32, 38, 39], a constraint from supernova cooling [32] as updated in [40], and an updated constraint from the electron anomalous magnetic moment [30, 31]. In the green band, the A' can explain the observed discrepancy between the calculated and measured muon anomalous magnetic moment [36] at 90% confidence level.

2.1.1 Heavy Photons and Dark Matter

The possible role of heavy photons in the physics of dark matter [41, 42] has provided an urgent impetus to search directly for heavy photons. Results from two classes of dark matter searches — “indirect” searches for galactic dark matter annihilation and “direct” searches for dark matter scattering off nuclei — have both been interpreted as potential signals of dark matter interacting through a heavy photon. Both areas have developed considerably in recent years, but not decisively. Here we briefly summarize the status of dark matter, the case for its interactions with heavy photons, and pertinent recent developments in both

observation and theory. The motivation to test these theories of dark matter in a controlled laboratory experiment remains strong.

The concordance model of big bang cosmology — the Lambda Cold Dark Matter (Λ CDM) model — explains all observations of the cosmic microwave background, large-scale structure formation, and supernovae, see e.g. [43]. This model suggests that Standard Model particles make up only about 4% of the energy density in the Universe, while “dark energy” and “dark matter” make up 74% and 22%, respectively, of the Universe’s energy density. The concordance model does not require dark matter to have any new interactions beyond gravity with Standard Model particles. However, an intriguing theoretical observation, dubbed the “WIMP miracle”, suggests that dark matter does have new interactions. In particular, if dark matter consists of 10 GeV to 10 TeV particles interacting via an electroweak-strength force (weakly interacting massive particles or WIMPs), they would automatically have the right relic abundance consistent with the Λ CDM model.

If dark matter does interact with ordinary matter, such interactions could produce at least two observable consequences: dark matter particles in the Milky Way Galaxy (and other bound astrophysical systems) can annihilate into visible matter, which could be detectable as energetic cosmic rays and/or gamma rays at Earth (indirect detection). Dark matter passing through Earth can also scatter off nuclear targets, causing the target to recoil. This recoil is observable in radio-pure detectors with sufficiently low background rates of nuclear recoil (direct detection).

a. Indirect Detection

The satellites PAMELA [44] and Fermi [45], the balloon-borne detector ATIC [46], the ground-based Cherenkov telescope HESS [47, 48], and other experiments have all reported an excess in the cosmic-ray flux of electrons and/or positrons above backgrounds expected from normal astrophysical processes. The evidence for this excess has only grown, with new measurements of the cosmic-ray electron flux by PAMELA [49] and confirmation by Fermi of the positron excess [50]. It is expected that results from AMS-II will eventually shed more light on the spectrum of these excess cosmic-rays. However, the origin of these excess positrons and electrons remains unknown. It may plausibly arise from any of three possibilities: pair creation in nearby pulsars, acceleration in supernova shocks, or dark matter annihilation or decay.

If the excess arises from dark matter annihilation, two features are incompatible with

annihilation of “conventional” thermal WIMP dark matter charged under the Standard Model weak interactions, but compatible with an alternative explanation, namely that dark matter is charged under a new $U(1)'$ and annihilates into A' pairs, which decay directly into electrons and positrons, and/or into muons that decay into electrons and positrons (see e.g. [41, 42, 51–53]):

- The annihilation cross-section required to explain the electron signal is 50 – 1000 times larger than the cross-section favored for the “WIMP miracle”. This can be explained if dark matter interacts with an $\mathcal{O}(\text{GeV})$ -mass A' , which mediates a new moderate range force and enhances the annihilation rate at low velocities (the relative velocity of dark matter in the Galactic Halo, $v \sim 10^{-3}c$, is much lower than in the early universe, and the relative velocity in self-bound dark matter subhalos is lower still). We refer the reader to [54, 55] for a recent discussion.
- The PAMELA satellite did not see any anti-proton excess [56], which implies that, if dark matter annihilation is responsible for the positron/electron signals, it does not produce baryons. This contradicts expectations for dark matter annihilating through Standard Model interactions, but is expected if dark matter decays into light A' , which (for $m_{A'} \lesssim \text{GeV}$) are kinematically unable to decay into protons and anti-protons.

If dark matter annihilation produces the high-energy e^+e^- excess, correlated gamma-ray fluxes are expected from more distant astrophysical systems where dark matter can annihilate; such fluxes are not expected for the other possible explanations of the cosmic-ray excesses. Such gamma-ray fluxes have not been seen by satellite or ground-based gamma-ray telescopes, like the the Fermi Gamma-ray Telescope, MAGIC, HESS, or VERITAS. Bounds on the gamma ray flux from dwarf spheroidals [57], the outer Milky Way [58], the Galactic Center (e.g. [59, 60] and references therein), and distant galaxies [60, 61] and clusters [62] can thus be used to constrain dark matter interpretations of the Pamela/FERMI excess. In a similar spirit, dark matter annihilation in the epoch of atomic recombination would leave an imprint in the cosmic microwave background radiation, which is similarly constrained [63], and the self-interaction of dark matter via A' exchange could affect the shape of galactic halos [64, 65]. Each of these systems can be used to constrain models of the PAMELA/Fermi excess, albeit with large theoretical uncertainties.

The present situation can perhaps be summed up as follows: corroborating evidence for an explanation of the cosmic-ray excesses in terms of annihilating dark matter *could* have shown up, but have not. However, the size of the expected corroborating signals is very uncertain, so that the present situation is still inconclusive. Perhaps the best hope for a more definitive statement on a dark matter origin of the cosmic-ray excesses will arise from new CMB data expected from Planck, which will improve sensitivity to dark matter annihilation at the time of recombination by a factor of 10 over WMAP [63]. Planck should either find evidence for dark matter annihilation with a high cross-section (providing further support for dark matter interpretations of the e^+e^- excess), or more robustly constrain the minimal theories.

A very important caveat to the above discussion is that we assumed that dark matter *annihilations* to A' 's are the origin of the excesses. Instead, dark matter *decays* to an A' and other light hidden sector particles are also a viable possibility [66, 67]. In this case, an A' mass below ~ 1 GeV is again motivated by the absence of an antiproton signal, but the size of the e^+/e^- signal is set by the dark matter decay lifetime, and independent of the A' mass (recall that in the case of dark matter annihilations, the A' mass was an important ingredient in determining the size of the Sommerfeld enhancement and, thus, the annihilation cross section). Dark matter decays are less constrained than annihilations as a possible origin to the cosmic-ray excesses, as they produce a smaller corroborating gamma-ray signal (this signal is now proportional to the dark matter density ρ and not ρ^2). Also, no evidence is expected to show up in the CMB data, since the required dark matter lifetime to explain the cosmic-ray excesses is $\sim 10^{26}$ seconds, much larger than the time of the CMB formation ($\sim 10^{13}$ seconds).

b. Direct Detection

The search for dark-matter-nuclear scattering has also seen considerable developments recently, but remains equally ambiguous. Three experiments have reported excesses that *may* be attributable to dark matter, although more mundane explanations are certainly possible: DAMA/Libra [68], CoGeNT [69], which also reported an annual modulation signal [70], and CRESST [71]. If all or a subset of these signals have a dark matter origin, they are most readily attributed to light dark matter (~ 10 GeV). However, results from CDMS [72], XENON10 [73], and XENON 100 [74] appear to exclude the same parameter regions. Experimental and detector uncertainties remain large enough that perhaps some model

parameter space remains moderately consistent with all of these results [75]. Though the evidence for light dark matter is controversial, it does raise a puzzle: dark matter with such low masses and high couplings cannot easily interact through Standard Model forces (such as Z -boson exchange), without being excluded by measurements of the total Z width at LEP. If indeed dark matter is light, then it seems most likely to interact through a new mediator, a possibility that HPS will probe in the case of an A' .

We note that heavy inelastic dark matter ($\sim 100 - 1000$ GeV) interacting with nuclei through A' -exchange was a possible explanation for these direct detection anomalies a few years ago, and its annihilation to A' 's could also have explained the cosmic-ray excesses. However, this possibility is now highly constrained by results from XENON100 [76] and CRESST [71]. Light dark matter, as mentioned above, is still viable. In order to have a unified dark matter explanation of the cosmic-ray excesses and direct detection anomalies, one would now likely need two components of dark matter, one light and one heavy component. Theoretical examples of such a possibility have been discussed in the literature, see e.g. [67].

2.1.2 Heavy Photons and Muon $g - 2$

Besides being theoretically natural and having a possible connection to dark matter, an A' could explain the discrepancy between the measured and calculated value of the anomalous magnetic moment of the muon ($a_\mu = g - 2$) [36]. This long-standing puzzle has several possible resolutions, but among the simplest new physics explanations is the existence of a new force mediator that couples to muons, like the A' . The contribution to a_μ of the A' is like that of the photon, but suppressed by the mixing parameter ϵ^2 and dependent on the A' mass. The green region in Fig. 1 is the 2σ band in which the A' can explain the discrepancy. This is an intriguing region, which the HPS experiment will probe.

2.2 Update on Experimental Status

The most recent (as of October 2012) comprehensive update summarizing the experimental status of A' searches can be found in the presentations and summary talk of the Frascati “Dark 2012” workshop [6]. All relevant measurements and constraints, as of this workshop, are included in Fig. 1. One important change relative to a year ago is that an improved

measurement of $g - 2$ of the electron has slightly reduced the range of allowed parameter space (on the low mass range) for an A' to explain the $g - 2$ of the muon discrepancy [30, 31]. Additionally, searches for A' 's in rare ϕ decays at KLOE and rare π^0 decays at WASA have slightly reduced the allowed parameter space on the high mass range [77]. An ongoing search in the e^+e^- spectrum of data recently gathered by MAMI [9] should also be sensitive to the A' mass range explored by KLOE, but with sensitivity down to roughly $\epsilon \sim 10^{-3}$. Finally, improved theoretical calculations and modeling of the experimental acceptance have led to slightly revised constraints on the A' parameter space from past beam dump experiments sensitive to A' production and decay to e^+e^- pairs [78].

2.3 HPS physics with True Muonium

Positronium and muonium, bound states of (e^+e^-) and (μ^+e^-) pairs, respectively, have been produced and studied [79–81], but True Muonium has not yet been detected (see e.g. [82–92]). Together with tauonium $(\tau^+\tau^-)$ and tau-muonium $(\tau^\pm\mu^\mp)$, True Muonium is among the most compact pure QED systems. While $(\tau^+\tau^-)$ and $(\tau^\pm\mu^\mp)$ are difficult to detect since the τ has a weak decay that competes with the QED decay, the μ is very long lived so that the decay of True Muonium is purely a QED process.

The detection of True Muonium would be a significant discovery and would constitute a further important test of QED. A number of applications of True Muonium measurements have been highlighted in [84], designed to exploit True Muonium as a perturbative laboratory for QCD bound state physics. These include measuring dissociation cross-sections as a function of energy and lifetimes of the various states. More speculatively, the discrepancy between theory and experiment for $g - 2$ of the muon [93] and the discrepant measurement of the charge radius of the proton using muon bound states [94] suggest that further measurements of muon properties would be useful to resolve these puzzles.

Studies of the production and dissociation of True Muonium suggest that the yields in HPS should be sufficient for observation [95], and are discussed further in section 3. That HPS is uniquely suited for detecting True Muonium is straightforward to understand. The triplet True Muonium states 1^3S_1 , 2^3S_1 , and 2^3P_2 all eventually decay to e^+e^- final states, with lifetimes long enough to leave a detectable displaced vertex. In that important respect, triplet True Muonium states behave just like A' 's. True Muonium production kinematics

is a bit different. In HPS, True Muonium will be produced by electron scattering off the high- Z nuclear target. The so-called “single-photon” production mechanisms gives rise to True Muonium states with kinematics extremely similar to A 's – HPS will be most sensitive to these. The “three-photon” production mechanism, which is typically larger, gives rise to True Muonium with characteristically lower energy.

In addition to primary production mechanisms, there are a variety of secondary mechanisms that are important in targets thicker than $\sim 0.01\%$ radiation lengths. This was studied in some detail in [95], where it was shown that 1^3S_1 excitations (in the target) are the main source of 2^3S_1 and 2^3P_2 production. The $2S$ and $2P$ state are especially long-lived, so this finding suggests that HPS may first discover these states as they will comprise a sizable fraction of the e^+e^- decays with displaced vertex in the range of ~ 1 cm to several cm. The $1S$ state will be the main component of the decays in the region of ~ 1 cm and below.

2.4 HPS Searches for Hidden Sectors

As highlighted in the Intensity Frontier Workshop report [4], a well-motivated class of beyond the Standard Model scenarios include new particles that interact indirectly or very weakly with Standard Model matter (hidden sectors), possibly associated with dark matter. Low-energy and high-intensity experiments offer an excellent tool for exploring these possibilities, complementary to the ongoing efforts at high energy colliders.

HPS is primarily designed to look for new sub-GeV A 's that decay into lepton pairs. But if an A' is part of a larger hidden sector, as is often assumed in the literature, some fraction of the decays could be more intricate. For example, an A' might decay into hidden sector particles, which in turn may decay back into Standard Model lepton or photons. These decays would typically have displaced vertices and multiple leptons or photons. The phenomenology of a variety of such scenarios have been considered in [16, 96] (and references therein). Search strategies to look for more general decays of A' 's into hidden sector particles are actively being developed within HPS, and we comment in more detail on this physics opportunity in Section 3.3.

3 Proposed Measurements

The primary goal of the proposed experiment is to search for a heavy photon (dark photon) in the mass range from 20 MeV to 1000 MeV in at least two settings of beam energy 2.2 GeV and 6.6 GeV. HPS ultimately relies upon the precision measurement of two quantities: the invariant mass of the A' decay products and the position of the decay vertex. By placing a tracking and vertexing detector immediately downstream of the target inside an analyzing magnet, the complete kinematic information required for A' reconstruction can be obtained from a single system, whose proximity to the target naturally maximizes the acceptance of a relatively compact detector and provides excellent momentum and vertexing resolution. A finely segmented, fast electromagnetic calorimeter, just downstream of the tracker, provides a powerful high rate trigger, identifies electrons, and augments the electron energy measurement. Behind the ECal a muon system consisting of four planes of scintillator hodoscopes sandwiched between iron absorbers will be positioned. The muon system will provide a trigger for $(\mu^+\mu^-)$ detection and will be used for muon identification. It will extend the search for high mass A' in di-muon decay mode where electromagnetic backgrounds are much reduced. Very high rate data acquisition systems, for the tracker, Ecal and muon system, make it possible to trigger and transfer data at 10s of kHz, and run with negligible dead time.

The HPS experiment also has the potential to discover “true muonium”, a bound state of a $\mu^+\mu^-$ pair and to search for non-minimal hidden sector final states.

HPS plans to execute the full experiment in two phases. The first phase will start with a commissioning run in 2014 which will include data taking at 1.1 and 2.2 GeV beam energies. More extensive data taking will continue in 2015, with runs at 2.2 and 6.6 GeV. This first phase will use about one sixth of the total beam time that HPS has requested, roughly 4 weeks on the floor at each of 2.2 and 6.6 GeV, which, assuming 50% combined efficiency for the accelerator and detector, corresponds to one month total of (perfect) run time. The second phase of HPS running, which will occur in 2016 and beyond, will use the additional run time to extend the search for heavy photons to the largest possible region of parameter space, and study the properties of true muonium in detail.

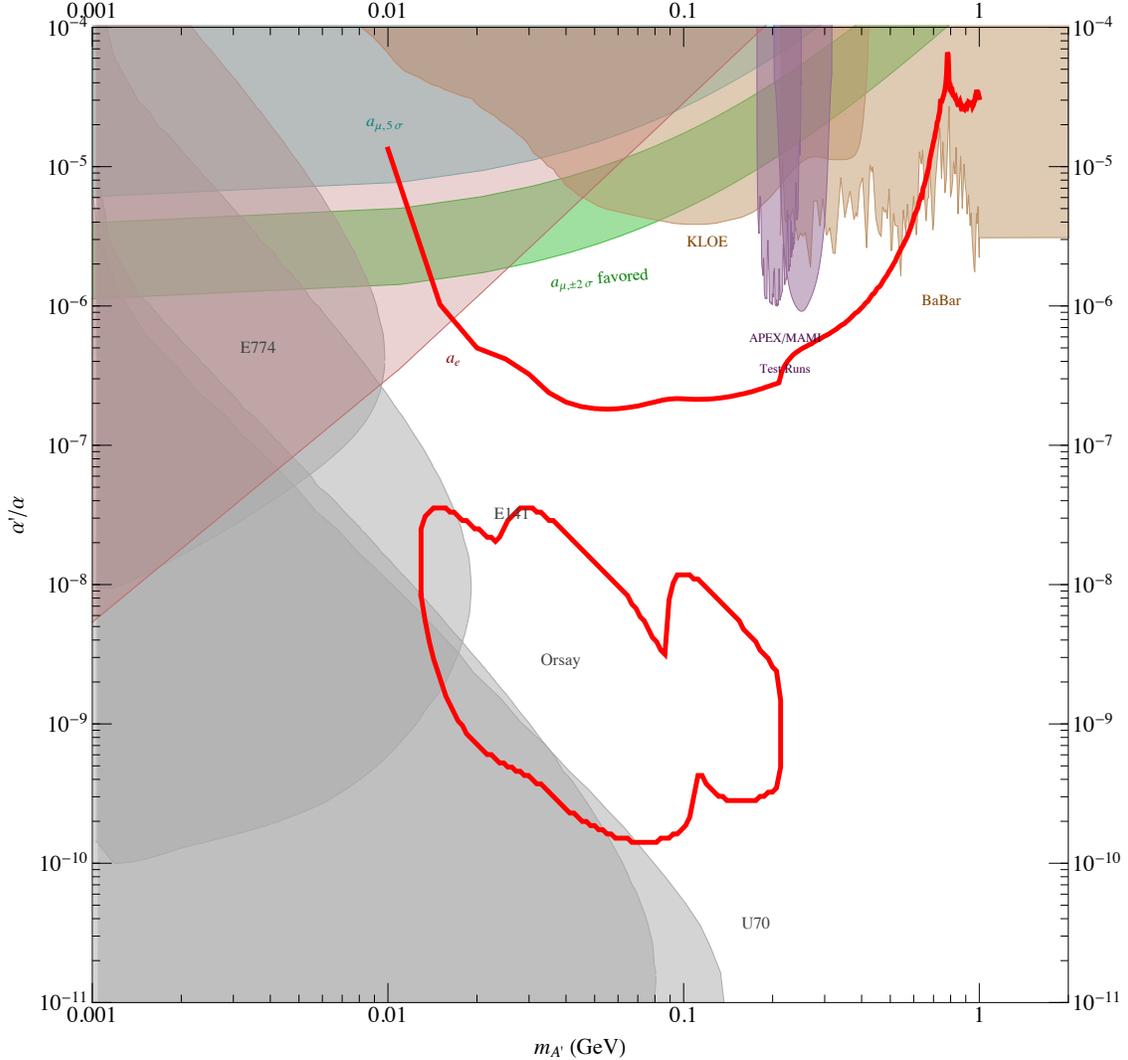


FIG. 2: Expected mass vs coupling parameter space reach full 2014-2015 running (solid red). Red line contour corresponds to 1 week of beam time at 1.1 GeV, and 3 weeks of beam time at 2.2 GeV and 6.6 GeV.

3.1 Search for the heavy photon

The expected in the first phase of the HPS is shown in Figure 2. The reach in mass-coupling parameter space is calculated using the simulated detector response as shown in Section 6. The plot shows two distinct regions, one at larger coupling corresponding to a purely bump-hunt region and another at lower coupling where the vertex of the A' decay is displaced.

A' particles are generated in electron collisions on a fixed target by a process analogous to ordinary photon bremsstrahlung, see Figure 3. The rate and kinematics of A' radiation

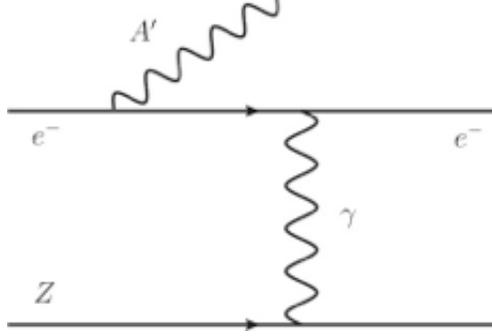


FIG. 3: Diagram of A' production by bremsstrahlung off of an incoming electron scattering with an atomic nucleus.

differ from massless bremsstrahlung in several important ways:

- **Rate:** The total A' production rate is controlled by $\alpha^3 \epsilon^2 / m_{A'}^2$. Therefore, it is suppressed relative to photon bremsstrahlung by $\sim \epsilon^2 m_e^2 / m_{A'}^2$. Additional suppression from small $\tilde{\chi}$ occurs for large $m_{A'}$ or small E_{beam} , see [32] for details.
- **Angle:** A' emission is dominated at angles $\theta_{A'}$ such that the virtuality of the intermediate electron, U (see definition in Eq.D2 in Appendix), is $U(x, \theta_{A'}) \lesssim 2U(x, 0)$ (beyond this point, wide-angle emission falls as $\theta_{A'}^4$). Near its median value, the cutoff emission angle is

$$\theta_{A',max} \sim \max \left(\frac{\sqrt{m_{A'} m_e}}{E_0}, (m_{A'}/E_0)^{3/2} \right), \quad (1)$$

which is parametrically smaller than the opening angle of the A' decay product, $\sim m_{A'}/E_0$. Although this opening angle is small, the backgrounds mimicking the signal dominate at even smaller angles.

- **Energy:** A' bremsstrahlung is sharply peaked at $E_{A'}/E_{beam} \approx 1$, where $U(x, 0)$ is minimized. When an A' is produced, it carries nearly the entire beam energy. In fact, the median value of $(1-x)$ is $\sim \max \left(\frac{m_e}{m_{A'}}, \frac{m_{A'}}{E_0} \right)$.
- **Lifetime** For the ranges of ϵ and $m_{A'}$ probed by this experiment, the mean decay length l_0 of the A' can be prompt or as large as tens of centimeters. All of the background decays promptly at the target.

The latter three properties are quite important in resolving signal events from the main backgrounds, as discussed below. More details of A' production and decay are given in Appendix D.

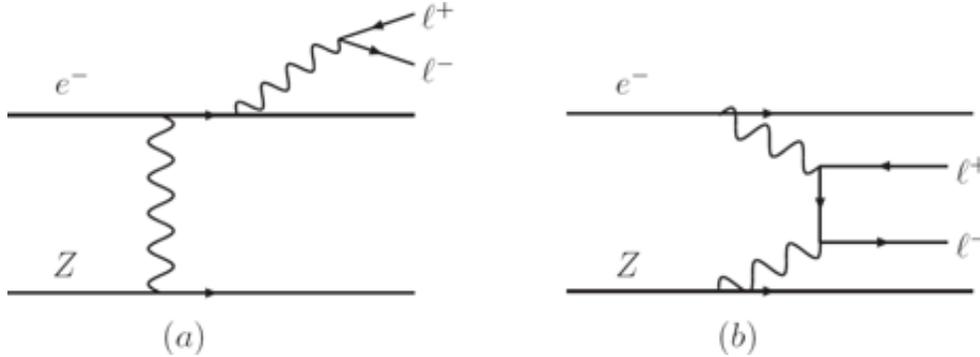


FIG. 4: Sample diagrams of (left) radiative trident (γ^*) and (right) Bethe-Heitler trident reactions that comprise the primary background to the $A' \rightarrow l^+l^-$ search.

The irreducible background rates are given by the diagrams shown in Figure 4. These trident events can be usefully separated into "radiative" diagrams (Figure 4 (a)), and "Bethe-Heitler" diagrams (Figure 4 (b)), that are separately gauge-invariant. These QED tridents dominate the final event sample, so we consider their properties in some detail here.

The contribution from the radiative diagrams (Figure 4 (a)) alone is also useful as a guide to the behavior of A' signals at various masses. Indeed, the kinematics of the A' signal events is identical to the distribution of radiative trident events restricted in an invariant mass window near the A' mass. Moreover, the rate of the A' signal is simply related to the radiative trident cross-section within the spectrometer acceptance and a mass window of width δm by

$$\frac{d\sigma(e^-Z \rightarrow e^-Z(A' \rightarrow l^+l^-))}{d\sigma(e^-Z \rightarrow e^-Z(\gamma^* \rightarrow l^+l^-))} = \frac{3\pi\epsilon^2}{2N_{eff}\alpha} \frac{m_{A'}}{\delta m} \quad (2)$$

This exact analytic formula was also checked with a MC simulation of both the A' signal and the radiative trident background restricted to a small mass window δm , and we find nearly perfect agreement. Thus, the radiative subsample can be used to analyze the signal, which simplifies the analysis considerably.

Although the Bethe-Heitler process has a much larger total cross-section than either the signal or the radiative trident background, it can be significantly reduced by exploiting its very different kinematics. In particular, the A' carries most of the beam energy while the recoiling electron is very soft and scatters to a wide angle. In contrast, the Bethe-Heitler process is not enhanced at high pair energies. Moreover, Bethe-Heitler processes have a forward singularity that strongly favors asymmetric configurations with one energetic, forward electron or positron and the other constituent of the pair much softer. These

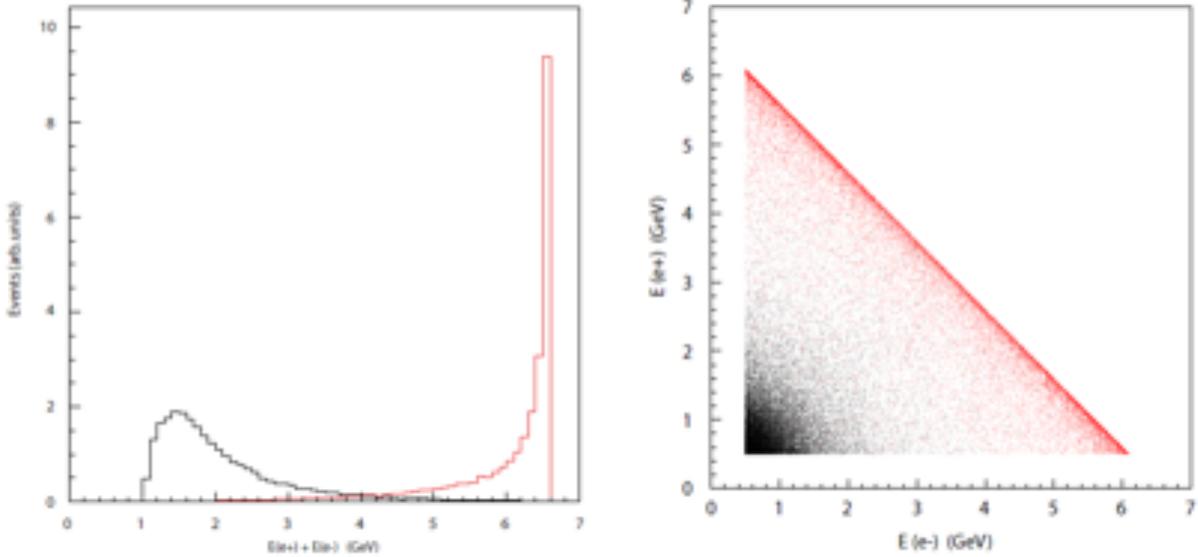


FIG. 5: Left: The distribution of Bethe-Heitler background events (black) and A' signal events (red) as a function of the sum of the electron and positron energy. Note that the signal is peaked at high energies, while the background is peaked at much lower energies. Right: The distribution of the positron versus electron energy for Bethe-Heitler background events (black dots) and A' signal events (red dots). Note that in both plots neither the signal nor background events have been normalized to the correct number. In reality, the number of background events is much larger than the number of signal events. Also, note that the electron energy here refers to the energy of the electron produced in the reaction, not the recoiling beam electron.

properties are discussed further in the Appendix of [4], and illustrated in Figure 5.

3.2 Search for true muonium

The proposed HPS experiment has the potential to discover “true muonium”, a bound state of a $\mu^+\mu^-$ pair, denoted here by $(\mu^+\mu^-)$. We expect that HPS will discover the 1S, 2S, and 2P true muonium bound states with its proposed run plan. The detection of these states should demonstrate the capability of the HPS experiment to identify rare separated vertex decays, and will provide a natural calibration tool for improving searches for heavy photons. The $(\mu^+\mu^-)$ atom is hydrogen-like, and so has a set of excited states characterized by a principal quantum number n . The binding energy of these states is $E = -1407 \text{ eV}/n^2$. The $(\mu^+\mu^-)$ “atom” can be produced by an electron beam incident on a target such as tungsten [82, 83].

With the existing proposal, HPS will search for true muonium just as it does for heavy

photons with separated vertices, requiring a vertex cut at about 1.5 cm to reject almost all QED background events, then searching for a resonance at $2 m_\mu$. An additional cut on the total energy of the e^+e^- pair of $E_{e^-} + E_{e^+} > 0.8 E_{beam}$ will also be required for triggering.

Based on [95], the total production yield for 1S, 2S, and 2P (including secondary production) leaving a target of thickness t_b (or larger) and satisfying the above requirements is,

$$N_{(\mu^+\mu^-)} = 200 \left(\frac{I}{450 \text{ nA}} \right) \left(\frac{t}{1 \text{ month}} \right) \quad (3)$$

where a beam energy $E_{beam} = 6.6$ GeV, and the nominal conditions of 450 nA beam current for 2 weeks of beam-time ($\sim 2.6 \times 10^6$ s) on a single foil has been assumed. The vertices near the cut of 1.5 cm will be dominated by the 1S state, while a tail of vertices extending out beyond a few cm is dominated by 2S and 2P.

Accounting for all the efficiencies associated with a separated vertex search, we would expect to see about 10–20 true muonium events in 2 weeks of 6.6 GeV run (we caution that the acceptance parameterization here is uncertain at the 50% level). The HPS experiment should be able to identify enough events to claim a discovery, and in addition, should be able to measure the mass of true muonium. There are certainly other properties of true muonium that would be interesting to measure. A measurement of the lifetimes would be interesting, as the lifetimes are sensitive to physics that couples to leptonic currents. With enough statistics, it should be possible to perform a measurement of the lifetimes of the 1S, 2S, and 2P states; work is ongoing to investigate this possibility.

3.3 Other searches for hidden sector particles

While the primary motivation for the HPS experiment is a search for A' decaying to lepton pairs, following Arkani-Hamed *et al.*[41], it is useful to explore the sensitivity of HPS is sensitive to other hidden sector particles, in particular those associated with Hidden Valley (HV) scenarios. As Strassler and Zurek[96] pointed out, HV scenarios provide many natural explanations of the nature of Dark Matter. The recent discovery of a boson that seems to have the properties of the Higgs boson of the Standard model also brings an old problem to the fore: why is the Higgs mass so light compare to the Planck scale? Searches at the LHC for supersymmetric partners and other particles with Standard Model couplings have

so far been unsuccessful, pushing the range of allowed masses for such particles higher and higher. HPS may be the first in a series of experiments that complement the energy frontier searches by looking for new, light particles that couple to the SM particles through new, very weak forces. Just as one should be ready for the unexpected when exploring the energy frontier at the LHC, so should we be in exploring the coupling frontier at with HPS.

In a general HV scenario, the new fermions may be lighter than the A' and have their own QCD-like forces, which results in them forming hidden mesons and baryons. Some of these hidden particles may be part of the Dark Matter. These hidden mesons would decay into SM fermion pairs either democratically or with mass-dependent branchings. One interesting

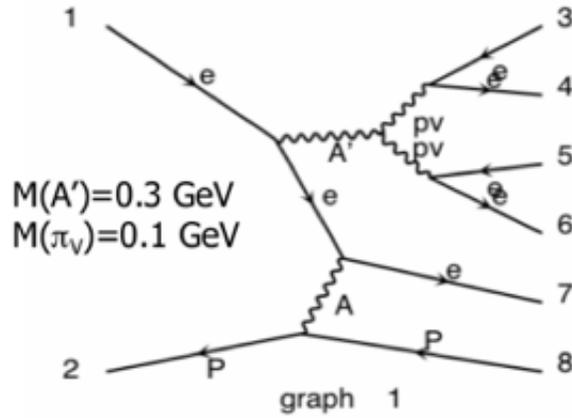


FIG. 6: Sample diagram of a non-Abelian hidden sector interaction.

case is where the A' decays into a pair of dark mesons (π_v), which in turn promptly decay into electron pairs (see Fig 6). The high multiplicity of the typical final states makes an exclusive search for events such as these difficult to trigger on, but it also reduces the background significantly by providing a number of new constraints and possible invariant mass bumps to search for and find. Simulation studies are ongoing towards estimating the reach HPS has for such hidden mesons.

4 Description of the HPS setup

4.1 Overview

HPS will utilize a setup located at the upstream end of experimental Hall-B at Jefferson Lab. The setup will be based on a three-magnet chicane, the second dipole magnet serving as the analyzing magnet for our forward spectrometer. The detector package will include a silicon tracker, an electromagnetic calorimeter, and a muon detector. High luminosities are needed to search for heavy photons with small couplings and masses in the 20 to 1000 MeV range. Utilizing CEBAFs essentially continuous duty cycle, the experiment can simultaneously maximize luminosity and minimize backgrounds by employing detectors with short live times and rapid readout. The HPS setup is designed to run with > 200 nA electron beams at energies from 1.1 GeV to 6.6 GeV impinging on a tungsten target of up to $0.0025 X_0$ located 10 cm upstream of the first layer of the tracker.

The HPS tracker consists of six double layer planes, 36 microstrip sensors in total. Placing the planes of the tracker in close proximity to the target means that the primary beam must pass directly through the middle of the tracking detector. This has necessitated that the sensors don't encroach on a dead zone, where multiple Coulomb scattered beam particles and radiative secondaries are bent into the horizontal plane, the so-called "wall of flame". However, since the energy released in the decay of a low mass A' is small relative to its boost, the opening angle between decay daughters can be quite small. To maximize the acceptance for low mass A' s, the vertical extent of the dead zone must be minimized and sensors placed as close as possible to the beam, so our design incorporates precision movers that can bring the silicon detectors to the required positions. Since interactions of the primary beam with air or even helium at atmospheric pressure gives rise to low-momentum secondaries that generate unacceptable occupancies, we have chosen to place the entire tracking and vertexing system in vacuum, in the Hall B pair spectrometer's magnet vacuum chamber. Silicon microstrip sensors are used in the tracker/vertexer because they collect ionization in 10s of nanoseconds and produce pulses as short as 50 – 100 ns. The sensors are read out continuously at 40 MHz using the APV25 chip, developed for the CMS experiment at the LHC. Running the high speed silicon module readout in vacuum further requires a vacuum

compatible cooling system, and data and power vacuum feedthroughs. All these features are incorporated in the design of the apparatus, as described below and have been tested in the May 2012 test run.

The electromagnetic calorimeter (Ecal) consists of 442 PbWO_4 crystals (reconfigured from the CLAS Inner Calorimeter) that are read out with APDs and amplifiers. The short pulse widths allow the ECal to run at very high rates. The Ecal data is digitized in the JLAB FADC250, a 250 MHz flash ADC developed for the JLAB 12 GeV Upgrade program. The full analogue information from the FADCs coupled with the fine spatial information of the calorimeter is available to the trigger, which uses energy deposition, position, timing, and energy-position correlations to reduce the trigger rate to a manageable ~ 30 kHz. Like the tracker system, the electromagnetic calorimeter is split to avoid impinging on the dead zone. The beam and radiative secondaries pass between the upper and lower ECal modules, which are housed in temperature-controlled enclosures, needed to stabilize the energy calibration.

The muon detection system will be installed behind the ECal, which has absorbed most of the electromagnetic background produced in the target. The remaining backgrounds will be attenuated by the first absorber layer of the muon system. The muon system will consist of four double layers of scintillator hodoscopes sandwiched between iron absorbers. Light from scintillator strips will be transported to photo-multiplier tubes via wave-length shifting fibers embedded within the strips. As in case of the ECal, muon system will be divided into two parts, beam up and beam down. There is a vacuum chamber between the two parts to allow the beam and radiated secondaries to be transported in vacuum.

In the following, the various elements of the experiment are discussed in more detail, beginning with the beamline, continuing with the tracker/vertexer, electromagnetic calorimeter, muon system, data acquisition systems, trigger, and calibration system.

4.2 HPS beamline

The Hall-B beam line, magnetic elements, beam profile monitors, and beam position and current monitors upstream of the HPS setup will be used as is (after slight modifications for 12 GeV). The only modification needed for the upstream part of the beam line is the addition of a collimator upstream of the Hall-B tagger magnet. The role of the collimator is to prevent the beam from directly hitting the Si-tracker in an event in which the high intensity beam

can move up or down from its nominal position. The collimators, which can consist of a 1 cm thick tungsten plates with different size oval holes (to match the beam profile), can be incorporated into the Hall-B photon tagger radiator ladder to provide vertical alignment. Horizontal alignment of the whole system will also be needed for fine tuning of the collimator position relative to the beam. Downstream of the HPS setup, there will be two beamlines, the electron beam line that will transport electron beam to the Hall-B electron beam dump, and a photon beam line that will transport the photon beam generated in the target to a photon beam dump mounted on the space frame. The electron beam will be transported in vacuum all the way through to the beam dump. Following the vacuum chamber in the last chicane dipole, the photon beam will go to the dump in Helium bag. There will be an H-corrector installed on the electron line after the HPS chicane to compensate any possible mis-steering of the beam in the chicane and to make sure that the electron beam stays on the original beamline to the dump. A YAG viewer will be used to monitor beam position just before the dump. The beam position on the dump monitor must stay unchanged before and after energizing the chicane.

4.2.1 Layout of the HPS setup

The HPS experiment will use the same three magnet chicane that was used for the CLAS Two Photon Exchange experiment (TPE). The layout of the beam line and the chicane is shown in Figure 7. The Hall B pair spectrometer magnet, an 18D36 (pole length 91.44 cm, gap $45.72 \times 15.24 \text{ cm}^2$, max-field 1.5 T), will serve as the analyzing magnet. The dipole field direction (Y) is perpendicular to the horizontal (XZ) plane. The Hall B Frascati H magnets (pole length 50 cm, gap $21 \times 8.25 \text{ cm}^2$, max-field 1.2 T) will be used as the first and the last dipoles of the chicane. The analyzing magnet will be operated at a 0.25 T-m field for 1.1 GeV, a 0.5 T-m field for 2.2 GeV, and a 1.5 T-m field for 6.6 GeV running. The two bending magnets will be set to 0.125 T-m, 0.25 T-m, and 0.75 T-m fields, respectively. The distance between the centers of the magnets will be about 50 cm bigger than it was for the TPE run, to accommodate the muon system. The location of the analyzing magnet along the beam will be exactly the same as it was for the TPE run.

The analyzing magnet will be displaced to beam left by ~ 11 cm in order to optimize the detector acceptance for e^+ and e^- , see Figure 8.

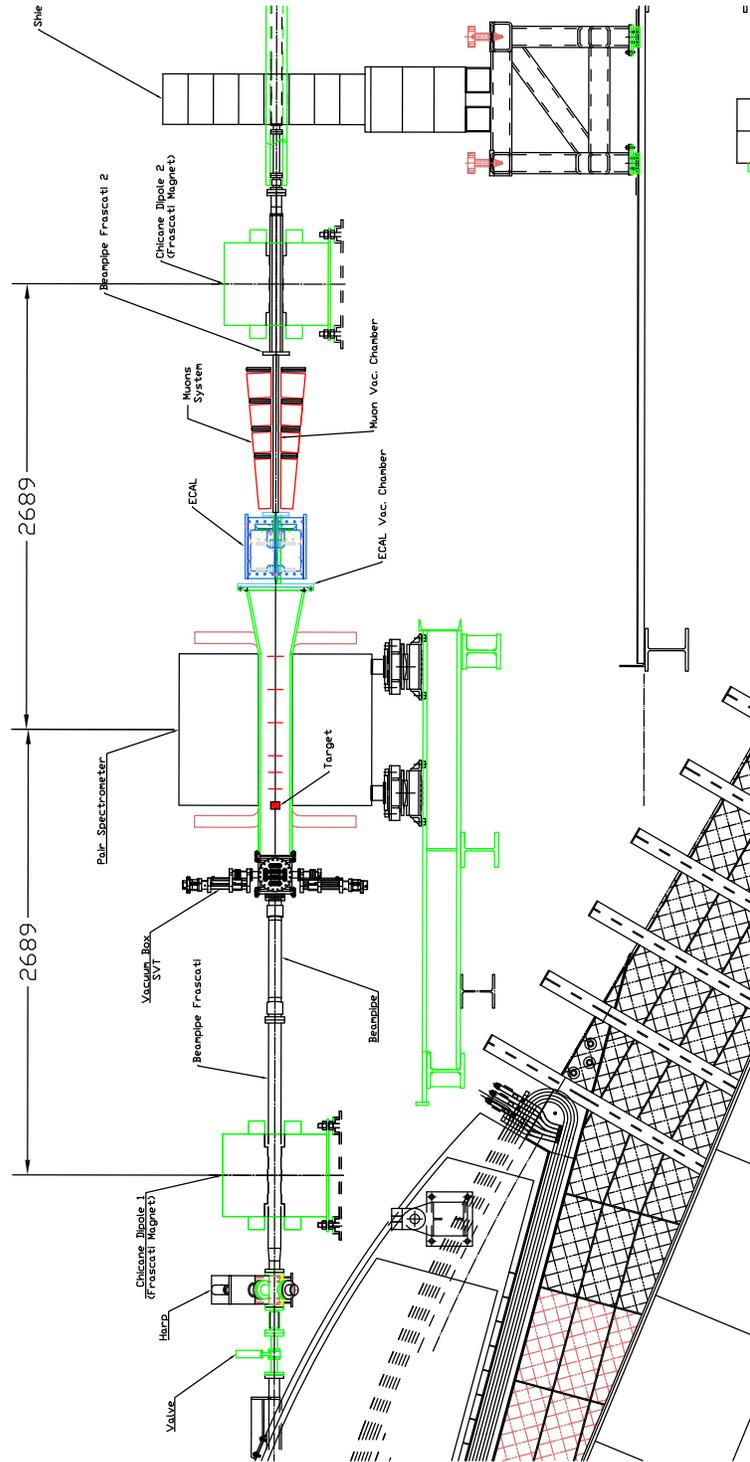


FIG. 7: Beam line configuration for the HPS test run with electron beams. The chicane configuration is similar to a previously run CLAS experiment.

The HPS target is positioned at the upstream edge of the analyzing magnet's pole. The distance from the target to the first layer of the silicon tracker is 10 cm, and to the face of the

electromagnetic calorimeter ~ 137 cm. There is continuous vacuum for the electron beam throughout the entire setup ending in the Hall B electron beam dump. The Si-tracker and the target will be located inside the Hall-B pair spectrometer vacuum chamber. The SVT vacuum box is mounted on the upstream end of the analyzing magnet vacuum chamber to provide connections for the SVT motion system, the cooling system, power and signal cables, and the target motion system. The Ecal vacuum chamber is attached to the downstream end of the analyzing magnet vacuum chamber, above and below which are placed the Ecal modules. Downstream of the Ecal vacuum chamber, another vacuum chamber is attached, leading through the muon system and the downstream chicane magnet.

The analyzing magnet, the Hall B pair spectrometer dipole, has its own power supply. The Frascati H magnets will use one common power supply and will be powered by the Hall B "mini-torus" power supply. There will be a shunt installed between the two Frascati magnets to allow independent small changes in currents on those two magnets if necessary (as it was done during the TPE experiment, although never used). Both power supplies are bipolar, so the magnets can be degaussed when needed. From available field map data at 900 A, the $\int Bdl$ of Frascati H magnets along the path of the electron beam is 0.663 T-m. The specified max current for these magnets is 950 A. In order to get 0.75 T-m an additional 10% increase in field value will be needed. From initial evaluation of the magnet design and characteristics, it should not be a problem to run them at 10% higher currents over their specified max current. If this is not possible, reducing the gap by 1/3 inch can yield the desired $\int Bdl$.

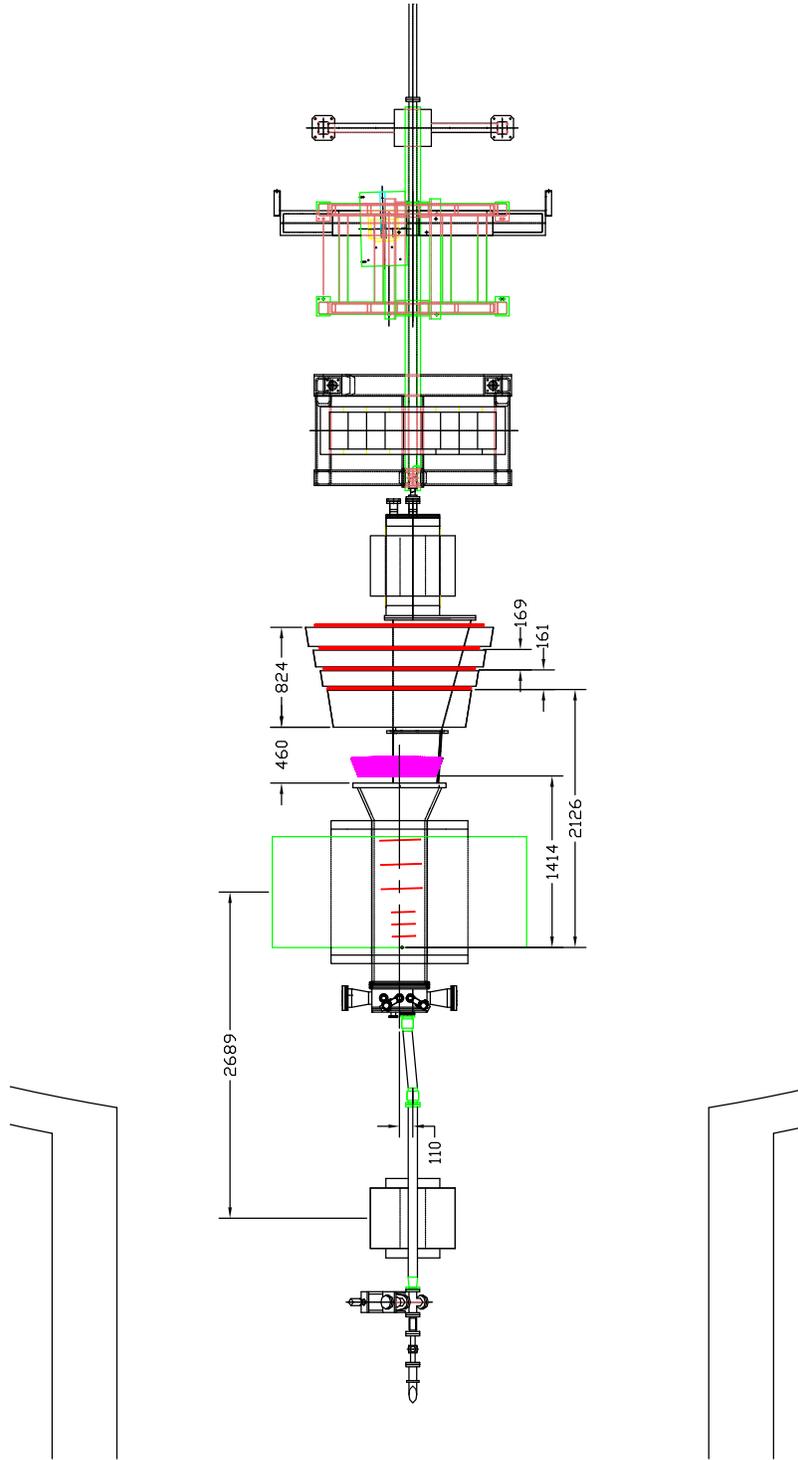


FIG. 8: Top view of the beam line configuration for the HPS test run with electron beams. The analyzing magnet is shifted by 4 inches (110 mm) to beam's left to get optimal acceptance for both e^+ s and e^- s.

4.2.2 Running Conditions

The HPS will use ~ 1.1 GeV, ~ 2.2 GeV, and ~ 6.6 GeV electron beams of up to 500 nA incident on a thin tungsten (W) target. Operational experience (with 6 GeV) showed that the CEBAF beam is very clean, and is contained within ± 0.5 mm with halo at the level of less than 10^{-5} . It is expected that the beams from the 12 GeV machine will be of comparable quality, at least for up to 3-pass beams (up ~ 6.6 GeV), so the primary electron beam should cleanly pass through the dead zone gap of the HPS setup.

For optimizing the vertexing performance and acquiring physics data, an asymmetric beam profile is desirable. Since the vertex resolution in the non-bend plane will be high, beam sizes of $< 50 \mu\text{m}$ in the Y direction are preferable. The momentum measurement will not benefit from small beam sizes in the X direction, and if the beam sizes in both dimensions are too small, the target foil will overheat. For these reasons the required beam sizes for HPS will be $\sigma_X \sim 250\mu\text{m}$ and $\sigma_Y < 50\mu\text{m}$. The HPS beam parameter requirements are presented in Table I.

Parameter	Requirement			Unit
E	1100	2200	6600	MeV
$\delta E/E$	$< 10^{-4}$			
Current	< 200	< 400	< 500	nA
Current Instability	< 5			%
σ_x	< 300			μm
σ_y	< 50			μm
Position Stability	< 30			μm
Divergence	< 100			μrad
Beam Halo ($> 5\sigma_Y$)	$< 10^{-5}$			

TABLE I: Required beam parameters.

The B-line optics in the 6 GeV era was checked using simulation and a beam test of the system. The optics program ELEGANT [97] was used to determine the optimized B-line parameters needed to achieve an asymmetric beam size, $\sigma_X \approx 250\mu\text{m}$ and $\sigma_Y \approx 20\mu\text{m}$, at the HPS test run target location. Beam tests were conducted in Hall B to validate these optics simulations during the Two Photon Exchange experiment when 2.2 GeV beam was available (February of 2011). Parameters were set for a beam profile of $\sigma_X \approx 300 \mu\text{m}$ and $\sigma_Y \approx 10\mu\text{m}$ at the Hall B tagger beam profiler (~ 8 meters upstream of the proposed HPS

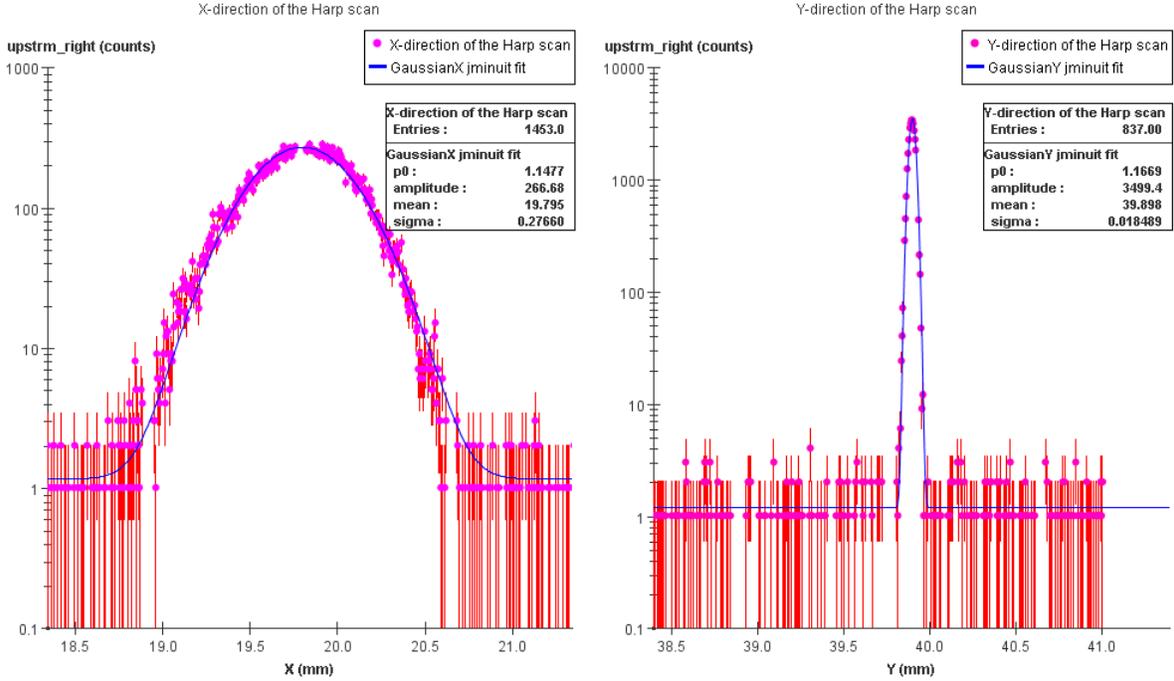


FIG. 9: Wire harp scan after loading optics parameters from the ELEGANT program. The wire scan speed was 0.1mm/s, readout speed is 15Hz. Based on the width of the Y-profile, the beam position stability is $< 20\mu\text{m}$. Note: any beam motion with more than $10\mu\text{m}$ amplitude and faster than 1Hz is included in the scan.

target location). Several beam profile scans with different scanner and data readout speeds were performed to check the beam stability and the systematics in the measurements. One of the scans is shown in Figure 9. As can be seen from the figure, the required profile can be reliably achieved. Several scans performed over two hours resulted in a consistent and stable beam profile. Based on the width of the Y-profile, beam position stability is $< 20\mu\text{m}$. Note that any beam motion with more than $10\mu\text{m}$ amplitude and faster than 1Hz is included in the scan.

The beamline optimizations have been performed for the 12 GeV CEBAF machine including the proposed changes for Hall-B/CLAS12 operations. Using the program ELEGANT and inputting the new locations of magnetic elements and their field maps, the beam profile was optimized at the HPS target location. In Figure 10 the beam sizes and the beam angles are shown for 6.7 GeV setup. The required beam size is achievable within the operating specifications of all the quadrupoles. Since HPS will run at beam energies < 6.7 GeV it is straight forward to scale (linearly) the magnets down to the other energies. The beam size/angle (beam transport) remains the same for 1.1, 2.2, and 4.4 GeV energies with the

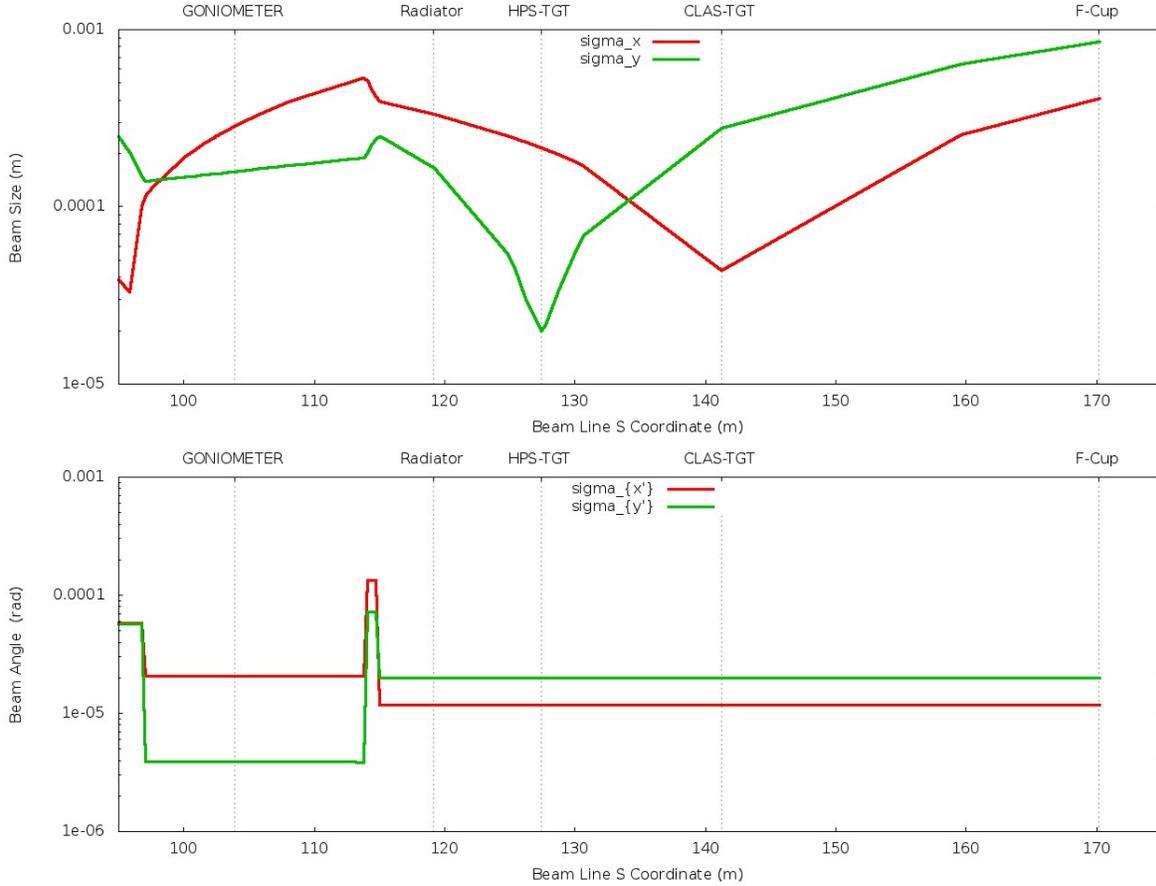


FIG. 10: Beam sizes in X and Y along the B-line in the upstream tunnel and in the region of the HPS test run setup. At the HPS target an asymmetric beam profile $\sigma_X = 300\mu\text{m}$ and $\sigma_Y = 20\mu\text{m}$ can be achieved with existing B-line optics.

exception of the small emittance increase at 6.7 GeV.

4.2.3 Beam Diagnostics

Beam position and current will be controlled using inputs from two sets of cavity beam position monitors (BPMs), that are located in the upstream tunnel. Sets of corrector dipoles and quadrupoles are routinely used to tune the beam for Hall B (2C21 to 2C24). A pair of BPMs, 2C21 and 2C24, will define the incoming trajectory of the beam and are included in the fast feedback loop. Readings from these BPMs will be used to maintain stable beam positions and currents. The stability of beam positions at two different locations also ensures the stability of the beam inclination.

The beam profile will be measured using three wire scanners. Two are installed in the

tunnel, the first one at 2C23, and the second one before the Hall-B tagger magnet, (2c24 harp, called "tagger harp"), about 8 meters upstream of the HPS target. The third wire harp, 2H00 harp, will be located just before the first chicane dipole. The first two profilers will be used to establish the required beam parameters during the initial setup. The Hall-B tagger magnet will be energized when beam tune is in progress, diverting the beam away from the HPS apparatus. After an acceptable beam profile is achieved, the tagger magnet will be degaussed and turned off, and the beam will be put through the HPS system and the beam profile will be checked using the 2H00 wire harp. The backgrounds in the HPS silicon tracker from beam profiling using the 2H00 harp have been simulated. At 5 nA beam current, the radiation damage is equivalent to about 10 sec. of production beam current on the HPS target, so is not a concern.

An insertable YAG screen beam viewer will be installed in the downstream alcove of Hall-B, before the Faraday cup, ~ 40 meters downstream of the HPS target. Both the position and profile of the beam will be used to setup the chicane magnets and to monitor beam quality during the run. A set of beam halo counters mounted along the beam line provides continuous and fast monitoring of the beam conditions. These counters are like those used for beam profile measurements. Excess noise in the beam halo counters triggers the machine fast shutdown system (FSD) in order to terminate beam in the event of beam excursions which could damage the HPS detectors. The FSD will occur in less than $40 \mu\text{s}$. In addition to halo counters, a beam offset monitor (BOM) will be installed upstream of the 2H00 wire harp. It is similar to the BOM used in CLAS. A short quartz cylinder, about 1 cm in diameter, with optical fibers attached around the edge will be centered on the beam. Even a few electrons in the beam tail will generate light in the cylinder that will be detected in a multi-anode PMT attached to the readout fibers. Errant beam motion towards the collimator located upstream of the tagger magnet will generate more light and increase the counts in the quartz cylinder, signalling a potential problem. The BOM will be wired to FSD as part of the equipment protection system.

4.2.4 Vacuum chambers

The SVT vacuum box will be attached to the existing magnet vacuum chamber as shown in Figure 11. Power, high voltage, and data signals to and from the hybrids are connected

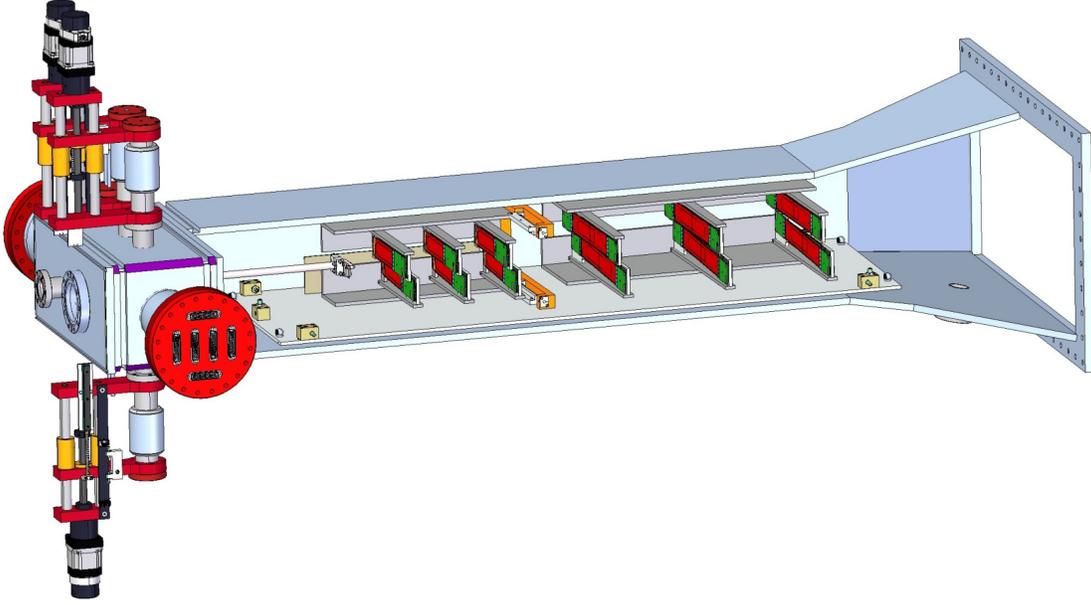


FIG. 11: Rendering of the SVT inside the Hall-B pair spectrometer vacuum chamber and the upstream vacuum box with SVT and target connections.

through two 8" flanges on the sides of the vacuum box. Two vertical linear motion mechanisms driven by stepper motors are used to position the SVT upper and lower modules with a precision of $1.25 \mu\text{m}/\text{step}$. A third linear motion mechanism is used to position the target on or off the beam. All the stepper motors are placed at a large enough distance from the magnet to avoid any ill effect from the magnetic fringe field. An existing stepper motor driver and EPICS-based control software will be used.

The scattering chamber between the top and bottom parts of the ECal is a critical beamline element. In order to keep the calorimeter as close as possible to the beam plane, include sufficient thermal insulation for the ECal, and maintain as wide a vacuum gap as possible, the top and bottom plates of the scattering chamber must be quite thin. At the location where the primary beams (e^- and γ) exit, the openings in the chamber have been enlarged. In Figure 12 a rendering of the scattering chamber in between the two halves of the ECAL is shown. The front flange of the chamber connects directly to the magnet vacuum chamber. Vacuum is maintained only on the electron side (beam right) since the backgrounds on the positron side are negligible. This design is based on detailed GEANT4 simulations of the background rates and acceptance of the ECal. It places crystals within 20 mm from the beam plane to maximize low-mass acceptance. In order to avoid excessive

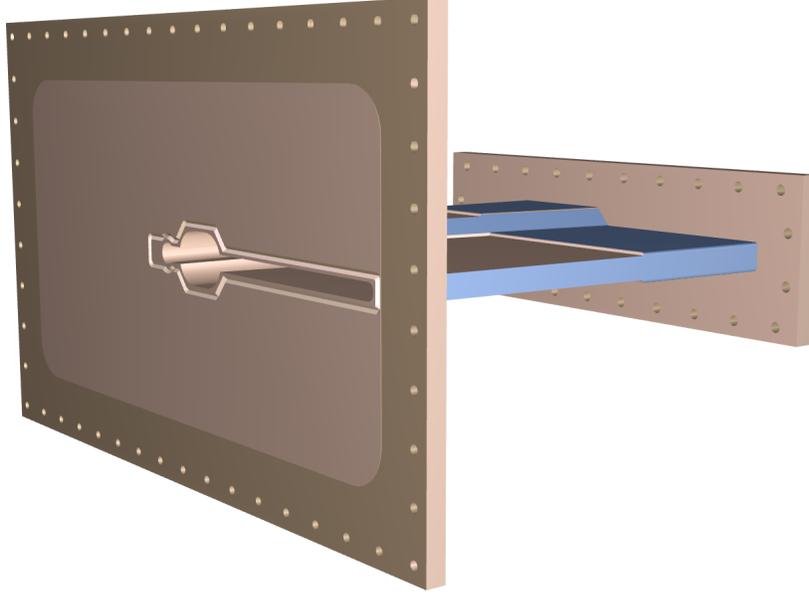


FIG. 12: Rendering of the ECal and the ECal vacuum chamber.

deformation of the thin walls of the vacuum chamber, an aluminum honeycomb support is inserted between the upper and lower walls, to beam's right.

The ECal vacuum chamber will be connected to the muon system vacuum chamber, located between the two halves of the muon system. Special openings for the photon and electron beams are not needed in the muon system vacuum chamber. The gaps for the radiated secondary electrons are essentially projections of those in the ECal vacuum chamber. At its upstream end, the muon vacuum chamber will have a gap of ~ 5 cm. At the downstream end that gap will be ~ 7.5 cm.

The last vacuum chamber, which passes through the third dipole, does not need to have a narrow opening. It will have size of the Frascati H magnet gap. At the downstream end of this chamber, there will be flange with two exit windows, a Kapton window for the photon beam to exit the chamber and go to the photon beam dump through a Helium bag, and a vacuum continuation to the standard beam line for the electron beam to go to the Hall-B electron beam dump.

4.2.5 Beam dumps and shieldings

The Hall-B electron beam dump will be used to terminate the electron beam. Due to its high intensity, the beam will not be dumped in the Faraday cup (FC); instead, the existing

beam blocker before the FC will be used to terminate the beam. The photon beam will be dumped in a photon beam dump, which will be a hut made of lead bricks located on the space frame. There will be a shielding wall after the last chicane magnet to prevent radiation from reaching the detector systems on the downstream side of the Hall.

4.2.6 Targets

A thin tungsten foil is used as the target. High Z material is chosen for its short radiation length, to minimize the hadronic production relative to the electromagnetic trident and A' production. The target is located 10 cm in front of the first plane of silicon strip detectors.

The primary target, 10 mm square, is 0.00125 radiation lengths (approximately $4 \mu\text{m}$ tungsten). Mounted immediately above it is a similar area of 0.0025 radiation lengths, available for some of the data taking, adjusting the beam current as appropriate. The foil can be fully retracted from the beam, and is inserted on to the beam line from above, using a stepping motor linear actuator. The bottom edge of the foil is free-standing so there is no thick support frame to trip the beam when the target is inserted. Its position is adjustable vertically allowing either thickness to be selected, and different sections of the tungsten can be used in the event of beam damage. The support frame on the beam-right side of the target is made thin enough to prevent radiation damage to the silicon in the event of an errant beam caused, for example, by an upstream chicane magnet trip.

The target is intended to operate with beam currents up to 450 nA, which produce strong local heating. The strength of tungsten drops by an order of magnitude with temperature increases in the range of 1000 C. In addition, the material re-crystallizes above this range, which increases the tendency for cracking where thermal expansion has caused temporary dimpling. For these reasons, it was decided to keep the temperature rise less than about 1000 degrees, which is accomplished by selecting an adequately large beam spot area. For example at 200 nA the rms beam radii will be held above 20 by 250 μm , or 40 by 250 μm for 400 nA. Simulations have shown that these beam spot sizes do not diminish the pair reconstruction resolution of the experiment.

In addition to the target, a set of tungsten beam-fiducial wires will be installed immediately in front of the silicon detectors. One horizontal wire, 20 μm diameter, and one 30 μm wire at 9 degrees to the horizontal, will be mounted on a frame attached to the upper

movable silicon support plate, and similarly for the bottom plate. The frames for the wires are wide enough that they do not occlude the silicon active area. The wires can be used to locate the position of the beam relative to the silicon. To accomplish this safely, the vertical separation between the front silicon sensor and its nearest wire is 8 mm. This separation, and their small diameters, also means that, when the sensors are positioned for data taking, the wires have a negligible effect on acceptance. The wires are also available for use as a fairly conventional wire scanner. In particular they can provide information about the minor and major axes, and the tip angle (roll), of a strongly elliptical beam.

4.3 Silicon vertex tracker

Achieving the best possible acceptance at low A' mass requires positioning the silicon as close as possible above and below the primary beam, where radiation and occupancy are both limiting factors. As a result, the silicon must be actively cooled to retard radiation damage, hits must be read out quickly and tagged with the best possible time resolution to reduce effective occupancies, and the tracker must operate in a vacuum to eliminate beam-gas secondaries. At the limit of feasibility from these considerations, the silicon in the first layer is only 0.5 mm from the center of the beam, so prudence dictates that the tracker be retractable from the beam. Meanwhile, achieving the best possible acceptance at high mass means that the active volume of the tracker fill as much of the magnet bore as possible. Finally, both the mass and vertex resolution that determine the experimental sensitivity are dominated by multiple scattering effects, so minimizing material in the tracker is the principal design goal. A more complete discussion of the key requirements and design principles of the HPS Silicon Vertex Tracker (SVT) may be found in the initial proposal to the JLab PAC [1].

The HPS Test Run SVT, described and discussed in Chapter 5, achieved these goals with the minimal possible apparatus capable of delivering A' physics during a short run. Unlike the initial proposal for the full experiment, this design used a single type of silicon microstrip module with small stereo angle arranged into five layers and compromised redundancy, precision, and longevity in order to compress the project timeline and reduce the budget. In the process of developing this design, it was found that this simple system was capable of delivering a surprising fraction of the physics potential anticipated for the full experiment. With this in mind, we now propose a design for the HPS SVT that builds upon the HPS Test SVT, principally by addressing the compromises made for HPS Test to ensure the best possible performance for A' physics within the envelope of the existing beam line layout and analyzing magnet. This design uses the exact same sensors and readout chips, retaining the most successful elements of the HPS Test design and addressing the weaknesses identified during assembly and operation to ensure the success of the experiment.

4.3.1 Layout

The layout of the HPS SVT is summarized in Table II and shown in Figure 13. There are six measurement stations, or “layers,” placed immediately downstream of the target. Each layer comprises a pair of closely-spaced planes and each plane is responsible for measuring a single coordinate, or “view”. Introduction of a stereo angle between the two planes of each layer enables three-dimensional tracking and vertexing.

Layer number	1	2	3	4	5	6
nominal z , from target (cm)	10	20	30	50	70	90
Stereo Angle (mrad)	100	100	100	50	50	50
Bend-plane resolution (μm)	≈ 60	≈ 60	≈ 60	≈ 120	≈ 120	≈ 120
Non-bend resolution (μm)	≈ 6	≈ 6	≈ 6	≈ 6	≈ 6	≈ 6
Number of sensors	4	4	4	8	8	8
Nominal dead zone in y (mm)	± 1.5	± 3.0	± 4.5	± 7.5	± 10.5	± 13.5
Module power consumption (W)	6.9	6.9	6.9	13.8	13.8	13.8

TABLE II: Layout of the HPS SVT. The angle of stereo sensors is relative to the axial sensors that have their readout strips parallel to the edge of the dead zone.

The layout of the first three layers is the same as in the HPS Test SVT, with a single sensor of coverage both above and below the beam and 100 milliradian stereo angle to balance acceptance against the resolution required for vertexing. The last three layers are two sensors wide in the bend direction to better match the ECal acceptance and use 50 milliradian stereo, as in HPS Test, to break the tracking degeneracy that creates fake tracks from ghost hits in layers with the same stereo angle. The first five layers cover the full acceptance of the ECal with one redundant layer. The sixth layer has only slightly less acceptance than the ECal and results in an extra safety factor in tracking purity and improved momentum resolution for the vast majority of tracks. Altogether, this layout comprises 36 sensors and 180 readout chips for a total of 23040 readout channels.

Acceptance for larger A' masses is limited by the size of the magnet but low-mass sensitivity depends on reconstructing tracks as close to the primary beam as possible; minimizing the so-called “dead zone” surrounding the degraded primary beam in the mid-plane of the detector. For the tracker, there are a number of considerations including proximity to beam halo and radiation damage in the first layer, ability to resolve time-overlapping hits, and

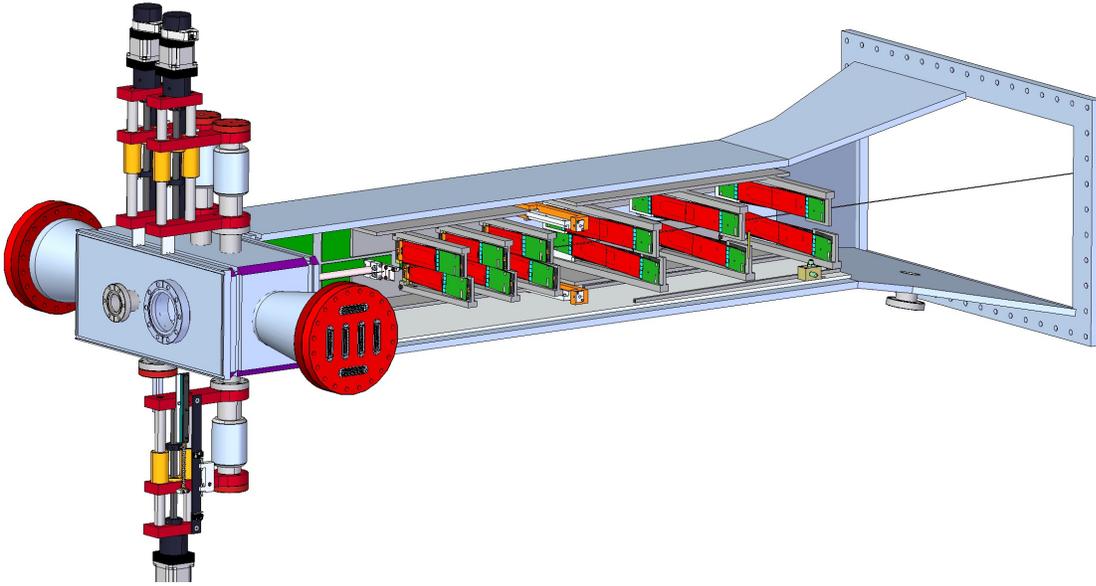


FIG. 13: A rendering of the concept for the HPS SVT. The beam enters from the left through the vacuum box providing services to the SVT. The silicon is shown in red and the hybrid readout boards in green.

the wall of pattern recognition errors at very high occupancies. With sensors capable of operation to 1.5×10^{15} 1 MeV neq., readout with single-hit resolution of ~ 2 ns and two-hit resolution of ~ 50 ns, and pattern recognition robust to 2% occupancy; the closest tolerable position of Layer 1 results in tracking acceptance outside of a 15 mrad dead zone around the beam plane. In this configuration, the edge of the silicon in Layer 1 is 0.5 mm away from the center of the beam where occupancy from beam-gas curlers would be unacceptable. Therefore, along with the drive to minimize multiple scattering errors, the desire to maximize low-mass acceptance creates the principal design challenges for the SVT: a lightweight, movable, liquid-cooled tracker with superior time resolution and operated in vacuum.

4.3.2 Module Design

One strength of the HPS Test SVT is exceptionally low material budget, an average of 0.7% X_0 per double-sided layer in the tracking volume with only 10% of this from support material. The HPS Test sensor modules achieve this figure by compromising the straightness, mechanical stability and cooling of the sensors. The module design for HPS aims to maintain

this material budget but eliminate these compromises by retaining the design of the carbon fiber half-modules but mounting them in a more robust way. Furthermore, this design allows the existing half-modules built for HPS Test to be reused for the first three layers of HPS, enabling the development of this module concept and the assembly of HPS to commence immediately and with a very small initial investment.

A half module for the HPS Test SVT consists of a single microstrip sensor and a hybrid electronic readout board glued to a polyimide-laminated carbon fiber composite backing, as shown in Figure 14. A window is machined in the carbon fiber leaving only a frame around

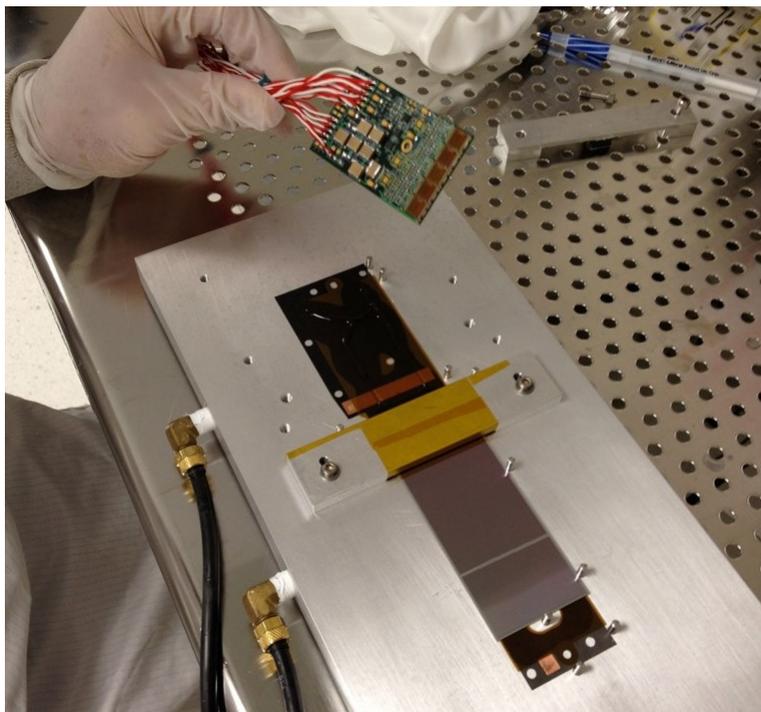


FIG. 14: A half module being assembled for Layers 1-3 of the HPS SVT. A silicon sensor and an APV25 readout hybrid are being glued to the Kapton laminated carbon fiber support frame.

the periphery of the silicon to minimize material. A $50\ \mu\text{m}$ sheet of polyimide is laminated to the surface of the carbon fiber with 1 mm overhang at all openings to ensure good isolation between the back side of the sensor, carrying high-voltage bias, and the carbon fiber which is held near ground. The sensors are single-sided, radiation tolerant, p^+ in n bulk, AC coupled, poly-biased sensors fabricated on $\langle 100 \rangle$ silicon manufactured by the Hamamatsu Photonics Corporation for the RunIIb upgrade of the $D\emptyset$ silicon detector. The sensors have $60(30)\ \mu\text{m}$ readout(sense) pitch and are read out by APV25 chips operating in multi-peak mode, allowing single hit position resolution of $\sim 6\ \mu\text{m}$ and tagging of individual hit times with a

precision of approximately 2 ns. Although specified to achieve 350 V bias, approximately 90% of sensors break down above 1000 V, sufficient to operate layer one of the SVT with full efficiency for six months at full beam intensity.

For HPS Test, the half-modules were sandwiched back-to-back around an aluminum cooling block at the hybrid end and a similar PEEK spacer block at the other. To allow for module rework, the modules were assembled with hardware and thermal compound instead of adhesive and have no stiffening material between the two sensors. For simplicity, only the cooling block at the hybrid end is supported, resulting in deviations in the planarity of the sensors as large as a few hundred microns at the cantilevered end. Furthermore, the lack of cooling at the unsupported end where there is no heat source from readout electronics limits the temperature achievable at the most highly irradiated portion of the sensor, a very small spot along the edge of the sensor adjacent to the dead zone. Improved cooling is necessary to achieve the longevity required for the longer running time envisioned for the HPS detector.

For layers 1-3 of HPS, these same half-module structures will be tensioned, like bicycle spokes, between a pair of cooling blocks held by a grooved aluminum base, as shown in Figure 15. Rather than building manifolds to provide cooling to these blocks and attempt

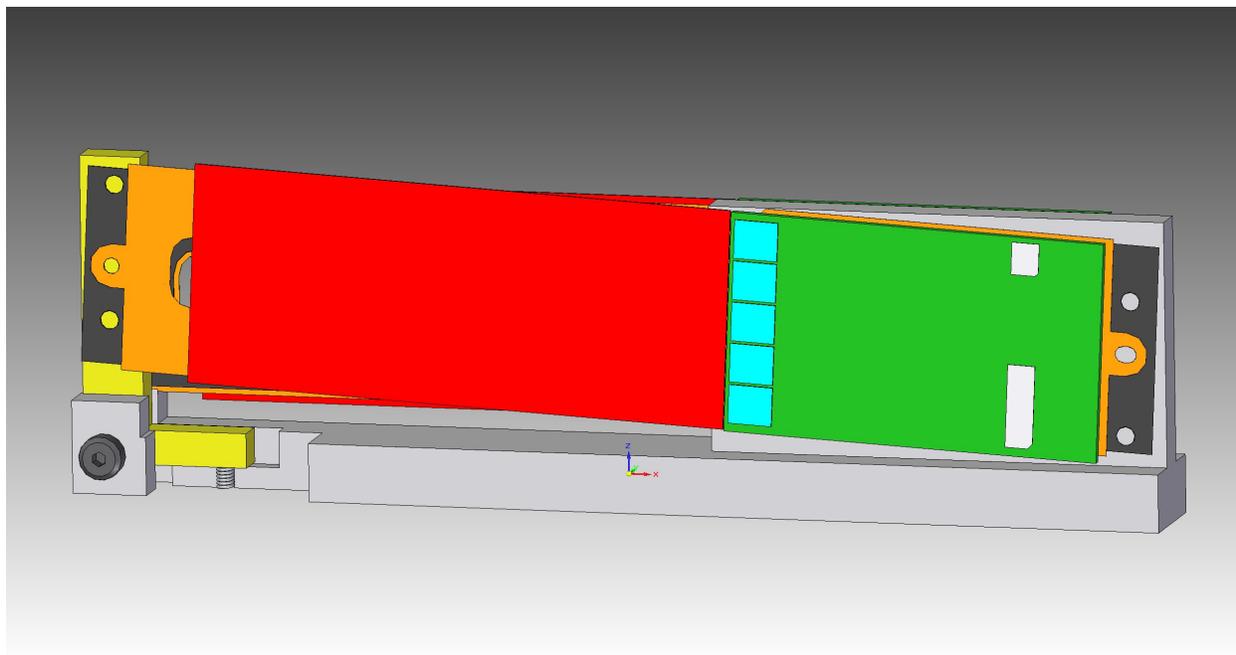


FIG. 15: A rendering of the concept for the new modules for Layers 1-3 of the HPS SVT. A cutaway at the left shows the spring and lever mechanism that maintains tension on the carbon fiber of the half modules.

to isolate them thermally from the underlying support structure as in HPS Test, the entire aluminum support structure will be cooled. The block at the hybrid end of the module is fixed, while the other pivots on a shoulder screw with a small stainless compression spring providing tension of approximately 40 N and taking up the 60 micron differential contraction between the half module and the support structure during a 30 °C cool down. The same low-viscosity thermal compound used in HPS Test will provide the thermal contact in the pivot joint between the grooved base and the pivoting block and generates a negligible temperature drop across the gap. This arrangement results in much flatter silicon, much better mechanical stability and much better cooling that provided by the previous design for layers 1-3. With temperature stability during running better than 1 °C, dimensional stability of the tracker will be better than intrinsic measurement resolutions and more than an order of magnitude better than the resolution limitation from multiple scattering effects.

More importantly, this scheme allows the ultra-thin design to be extended to the longer, double-sensor half-modules of layers 4-6 that have a pair of silicon sensors in the middle and a readout hybrid at each end, as shown in Figure 16. With a larger heat load to transmit,

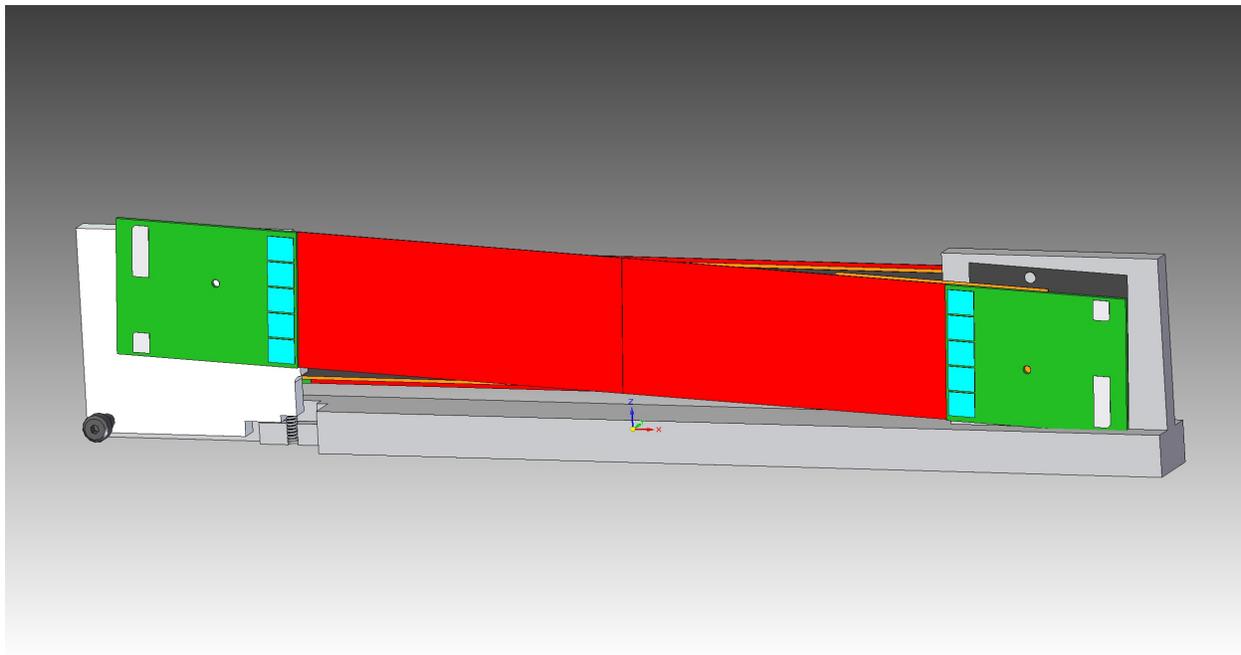


FIG. 16: A rendering of the concept for the new modules for Layers 4-6 of the HPS SVT. A cutaway at the left shows the mechanism responsible for keeping the half modules under tension.

the temperature drop through the pivot joint of the moving block will be approximately 0.5 °C. To accommodate the length of these double modules across the width of the vacuum

chamber, the hybrids will be shortened by approximately 20%. Through the use of polyimide flex cables instead of twisted pair and the elimination of superfluous circuitry, this footprint can be achieved with little effort. Aside from this, only minor changes to the conceptual design of the half module concept of HPS Test are envisioned to reduce assembly effort.

4.3.3 Mechanical Support, Cooling and Services

Sag of the aluminum support plates, when subjected to the bending load of the long motion levers, was the largest source of mechanical imprecision in the HPS Test SVT. For HPS, this motion system will be reused but with changes to eliminate this weakness. First, only layers 1-3 will retract, reducing the length of the support plate by a factor of two. Replacing the twisted pair readout with flex cables eliminates the largest external load on the plate. Finally, the 0.5 inch plate will be replaced by a 0.25 inch plate with 0.25 inch sides, forming a u-channel for increased stiffness, as shown in Figure 17. The walls of this

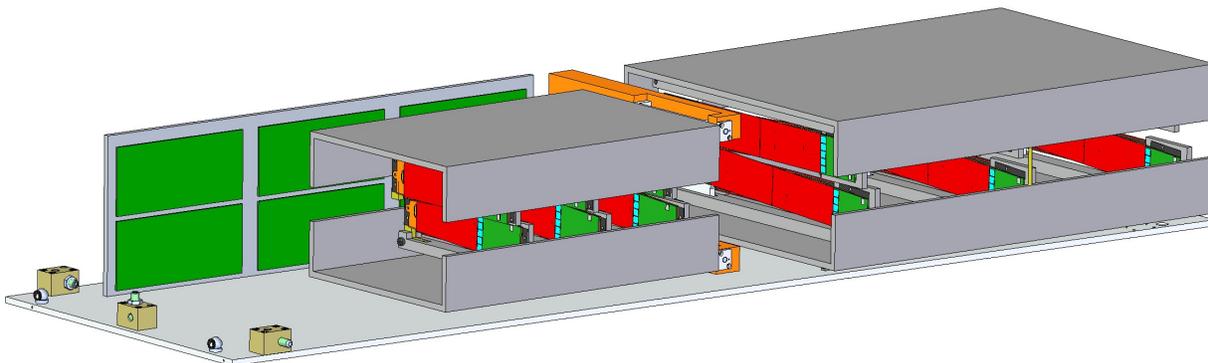


FIG. 17: A rendering of the support concept the HPS SVT. The modules are mounted in cooled channels. The channels for Layers 1-3 pivot on a downstream "C-support" and are moved by lever extending upstream to linear shifts on the vacuum box, as in HPS Test. Layers 4-6 are fixed in place. The DAQ boards, shown in green, are mounted to a separate cooling plate located in a low-neutron region upstream on the positron side of the detector.

support channel will extend almost to the dead zone and the entire structure will be cooled by large, embedded copper tubes. Surrounding the modules over most of the solid angle, these support channels will nearly eliminate thermal radiation from the walls of the vacuum chamber, the primary heat load on the sensors. Layers 4-6 will be mounted inside similar cooled channels but will be fixed since they are already far enough from the dead zone

for safety during beam tuning. These two support structures will be mounted to a single baseplate as before, with complete adjustability relative to the vacuum chamber.

Readout and power for HPS Test required 30 conductors per hybrid, or a total of 600 lines, with even that count requiring elimination of sense for DVDD on the hybrids. It does not appear feasible to scale this solution by nearly a factor of two for HPS. Instead, we plan to provide digitization and optical readout of the APV25 data as well as shared power for the hybrids on boards placed inside the vacuum chamber, as described in Section 4.6.1. In consideration of the DAQ requirements and the environment inside of the vacuum chamber, it appears that there is a natural location for support and cooling of the necessary boards adjacent to layers 1-3 on the positron side. This structure consists of a single vertically-oriented plate with an embedded cooling loop, shown in Figure 17. By separating the readout boards on the outside of this plate by the same 20 cm separation of layers 4-6, a single cable solution can be employed to connect the hybrids of layers 4-6 to these boards. This leaves layers 1-3 equidistant from the pair of readout boards on the inner side, where a second cable type can connect to the existing pigtail connectors for those hybrids. The feedthroughs required for power and data in this design fit easily on one of the two flanges in the existing vacuum box, leaving the other flange for additional cooling, eliminating the need for a cooling manifold inside of the vacuum chamber.

4.4 Electromagnetic Calorimeter

The Ecal, depicted in Figure 18, consists of 442 lead-tungstate PbWO_4 crystals with avalanche photodiode (APD) readout and amplifiers. The short output pulse widths permit operation at very high rates. Indeed, (PbWO_4) crystals with APD readout are ideally suited to deal with the expected high radiation and high rate environment and they can operate in the fringe field of the HPS magnetic field as well. The lead-tungstate modules, see Figure 19, are taken from the Inner Calorimeter (IC) of the JLab CLAS detector, which was built by IPN Orsay (France) and other Hall B collaborators and was used in a series of high energy electroproduction experiments. Orsay played a key role in the design and fabrication of the support frames, thermal enclosure, amplifiers, and motherboards of the IC, and is playing a similar role with the HPS ECal. The PbWO_4 crystals are 16 cm long and tapered. The cross section of the front face is $1.3 \times 1.3 \text{ cm}^2$, and $1.6 \times 1.6 \text{ cm}^2$ at the back. Modules in the ECal are arranged in two formations, as shown in Figure 18. There are 5 layers in each formation; four layers have 46 crystals and one has 37. The ECal is mounted downstream of the analyzing dipole magnet at the distance of about 137 cm from the upstream edge of the magnet. The two ECal modules are positioned just above and below the ECal vacuum chamber, through which the beam, radiated photons, and the wall of flame will pass unimpeded. The innermost edge of the crystals is just 2 cm from the beam. In order to stabilize the calorimeter's performance, the crystals, APDs, and amplifiers are enclosed within a temperature controlled environment, held constant at the level of 1° F . The energy resolution of the system, expected from operational experience with the IC, is $\sigma_E/E \sim 4.5\%/\sqrt{E}$ (GeV). As in the IC, PbWO_4 modules are connected to a motherboard that provides power and transmits signals from individual APDs and amplifier boards to the data acquisition system. The ECal data is digitized using the JLAB FADC250, a 250 MHz flash ADC developed for the 12 GeV Upgrade. Pulse height information and spatial and timing information from each crystal are available for the trigger decision, which uses this information to reduce the trigger rate to a manageable $\sim 30 \text{ kHz}$ (see Section 4.6.3 for details).

The HPS calorimeter described above was built and used during the HPS test run in April-May of 2012. This was the first time that a readout and trigger system utilizing the JLAB FADC250s had been used in a real experiment. The trigger algorithm was designed to

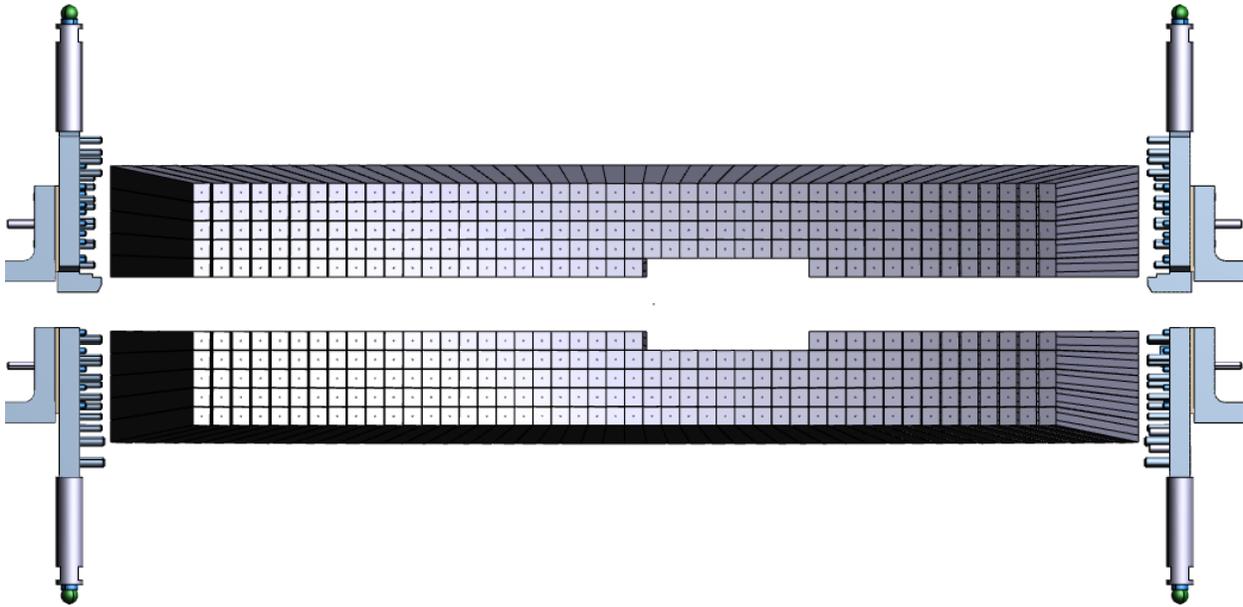


FIG. 18: Arrangement of ECal crystals. The two modules are positioned above and below the beam plane. Each module has 5 layers. There are 46 crystals in each layer, with the exception of the layers closest to the beam plane in which 9 crystals are removed to allow a larger opening for the outgoing electron and photon beams.

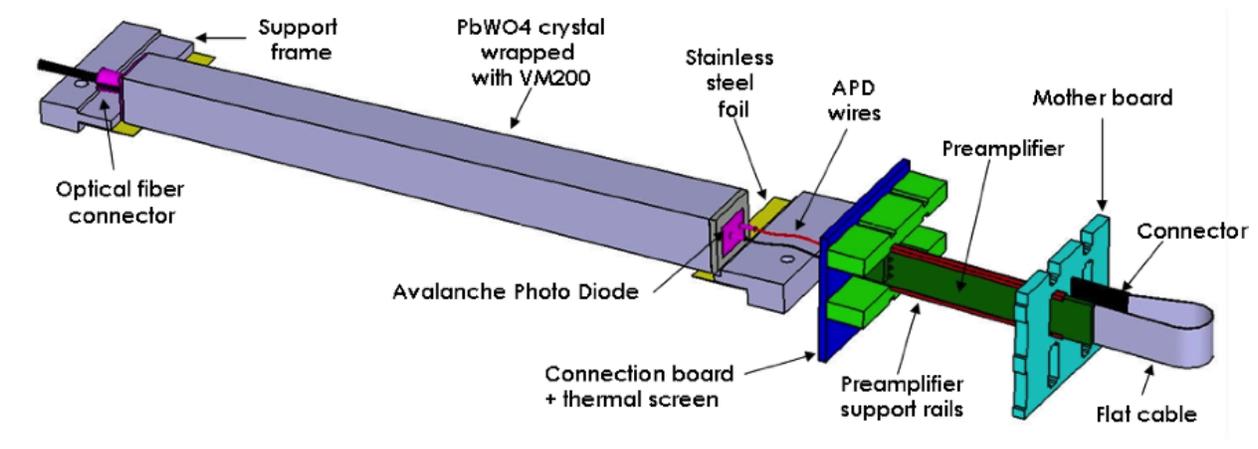


FIG. 19: The ECal module is composed of a 16 cm long lead-tungstate crystal, Avalanche Photo Diode, and a amplifier board.

satisfy the HPS event selection criteria and was implemented with newly developed FPGA-based trigger processors. Not all aspects of the trigger system were tested in the HPS Test Run because of the low interaction and background rates associated with photon running. But the ECal and its readout performed well and the critical goals for the Test Run run were achieved (see Section 5 for details). While the ECal performance during the test run

was satisfactory, several aspects are in need of improvement, as described below.

4.4.1 Improvements to the existing calorimeter

There are no plans for major mechanical changes to the ECal proper. The crystals, support frames, and thermal enclosure operated as designed and will stay unchanged. Most of the needed changes are related to the signal readout chain and problems encountered in the Test Run with the ECal and FADC250. Plans for modifications and/or improvements are as follows:

1. **Replacing the ECal mounting system** - During the test run the ECal was supported by the Hall-B pair spectrometer mount rails which also support with the pair spectrometer hodoscopes. Photon running prevented the installation of the ECal vacuum chamber and relaxed the requirements for precise ECal alignment. Consequently, the ECal was simply hung from the mount rails using long threaded rods. This system must be replaced with a more robust and finely adjustable (both horizontally and vertically) support mechanism to align the ECal correctly with the ECal vacuum chamber.
2. **Modifications to the side brackets to accommodate fiber bundles for the light monitoring system** - A light monitoring system was not used during the test run. While the design of the ECal enclosure was done in such a way that it can accommodate optical fibers attached to the front face of crystals, the side plates that hold the crystal frames do not have inlets for the accompanying fiber bundles. Space is available on the side plates for a straight-forward modification which will allow the addition of a light monitoring system.
3. **Modification of motherboards** - One of the issues faced during the test run was excess noise on some channels of the motherboards. Most missing channels, those which are absent on the performance figures in Section 5, were in fact switched off because they were very noisy and there was no time to debug them. New motherboards will be designed and built to resolve these noise issues. One of options under discussion is to replace long motherboards with shorter ones with power and signal connectors

located on the top (for the top module) and the bottom (for the bottom module) of the thermal enclosure.

4. **Signal splitting** - From the experience gained with the JLAB FADC250 by the HPS group and others, it is evident that the FADCs can be used both for precision time measurements and as real-time scalers simply by developing the appropriate firmware. Precise pulse timing will allow tighter coincidence windows and lower backgrounds, and real time scalers will provide good online monitoring of detector performance. The present ECal readout configuration uses signal splitters to divide the preamplifier signal in a 2:1 ratio, sending 2/3 of the signal to a discriminator that has built-in scalers and feeds the TDC channel. The other 1/3 of the signal goes to the FADC for the energy measurement. Removing the splitter will increase the signal on the FADC input by $\times 3$. This will allow lowering the amplifier gain and thereby lowering the noise level. This effort will make use of developments already planned for other detectors at JLAB.

5. **New preamplifiers** - The noise will be further reduced with the addition of new preamplifiers, allowing an additional reduction in effective energy threshold. The impact of the lower noise/threshold system is twofold: first it will improve the ECal's energy resolution; and second it will make the ECal sensitive to minimum ionizing particles which pass through the crystals transversely. With sensitivity to cosmic ray muons, which will pass through the ECal transversely when it is installed in HPS, the Ecal crystals can be calibrated for MIPS, and their effective gains balanced. HPS collaborators from INFN Genova have shown that with such a low noise, low threshold system, the ECal can distinguish the MIP energy deposition from noise, see Figure 20 and explanations below.

6. **Light monitoring system**

For an experiment like HPS, where backgrounds must be well understood and need to be strongly suppressed, the trigger bias must be minimized. In particular, having stable, known, and uniform thresholds at the trigger readout is necessary in order to avoid bias in the event selection. Such uniformity and stability can be achieved with the installation of an online gain monitoring system. This system will introduce

short light pulses into the front face of the crystals. The crystals already have fiber holders attached, allowing implementation of this system without having to modify the crystals or wrapping.

Optical fibers will be used to transmit light from a calibration source to the crystals to test the response of the APDs. The response of the system could change in time because of losses in crystal transparency due to radiation damage or because of gain variations of the APDs. Such a calibration system has been developed for several experiments (CMS at CERN for instance) with various light sources. The system for the ECal will be developed at IPN Orsay during 2013 and in the first half of 2014, and will be ready for installation at JLAB for the commissioning run in the fall of 2014. Each module will have a red and blue mono-color LED light source for monitoring purposes. Blue light transmission, corresponding to the domain of the crystal's emission spectrum, is very sensitive to the presence of color centers which are produced by radiation damage. So the blue light source will monitor variations of the response in the main domain of the spectrum. The response to red light is less sensitive to the color centers, and so permits monitoring the APD gains more directly. Thus the use of two colors separates gain variations due to the APD and electronics from those due to changes in the crystal transparency, and provides clear information on the state of the electronics. The main challenge for the system is to guarantee stability at a level of 2%. The test of the system will be carried-out at IPN Orsay, in order to study its efficacy and also to test the radiation resistance of the various materials.

7. **New low-voltage power supply** - The existing low voltage power supply is a manually controlled, single output power supply that feeds the four ECal motherboards through a custom-made patch panel. The present system cannot control the voltage supplied to preamplifiers at different parts of the ECal, and controlling or resetting them remotely has proved to be very inconvenient, requiring frequent access to the Hall, especially during commissioning. Newly available low voltage power supplies are much more flexible. The one that is the most suitable for the ECal APD preamplifiers is the WIENER MPV 8008LD. This power supply is being used at JLAB and the control software exists, so it will be easy to incorporate it into HPS.

4.4.2 Possibilities with new APDs

Installing new APDs on the current crystals will significantly improve the ECal performance, but doing so is expensive, so replacement is only being considered if a funding source beyond DOE HEP is secured. Replacing the old $5 \times 5 \text{ mm}^2$ Hamamatsu S8664-55 APDs with $10 \times 10 \text{ mm}^2$, Hamamatsu S8664-1010 will improve two critical characteristics of the present calorimeter modules. First, the new APDs from Hamamatsu have much better performance than the ones which are currently installed. Data from Hamamatsu shows that APDs made from the same wafer have excellent gain uniformity. With $\pm 10\%$ known uniformity at the gain of 100, the required variations in bias voltage are only $\sim 4.5 \text{ V}$. Even for large samples of APDs (1300), the required bias voltage differences are $\sim 50 \text{ V}$, which is half that of the current APDs. The ECal supplies bias voltage to groups of APDs, so with new APDs, with their smaller voltage-gain variations, it will be possible to achieve much better uniformity in the response of the calorimeter modules, and consequently better uniformity in trigger thresholds.

Secondly, the new APDs have a 4 times larger readout area, ensuring 4 times more light collection and therefore 4 times larger signals. This will allow the use of lower gain amplifier modules which in turn will decrease electronic noise. Tests performed for another calorimeter, now in production at INFN Genova for JLAB Hall-B, showed that amplifier boards with a factor 2 lower gains have a noise level $< 5 \text{ MeV}$. The energy deposition in the HPS PbWO_4 crystals from minimum ionizing cosmic muons passing transversely to the crystal axis is $\sim 15 \text{ MeV}$. Moving the energy thresholds close to 5 MeV will allow the MIP peak to be clearly distinguished, and will let the calorimeter be calibrated and monitored with cosmic muons. The HPS collaborators from the INFN group have performed the first tests of the Hamamatsu $10 \times 10 \text{ mm}^2$ APDs and a new amplifier board on HPS crystals. In Figure 20, the charge distribution of a single crystal system is shown for $5 \times 5 \text{ mm}^2$ (left) and $10 \times 10 \text{ mm}^2$ (right) APDs. A coincidence signal from scintillator telescopes positioned above and below the module provides the trigger. The crystal was positioned horizontally as it would be in HPS, so the cosmic muons would pass through it perpendicular to the crystal axis. The red line histogram is for all events triggered by the scintillation telescope and corresponds to the noise. The black line histogram corresponds to the charge detected within 100 ns of the trigger time. The MIP peak is clearly visible and well isolated from the

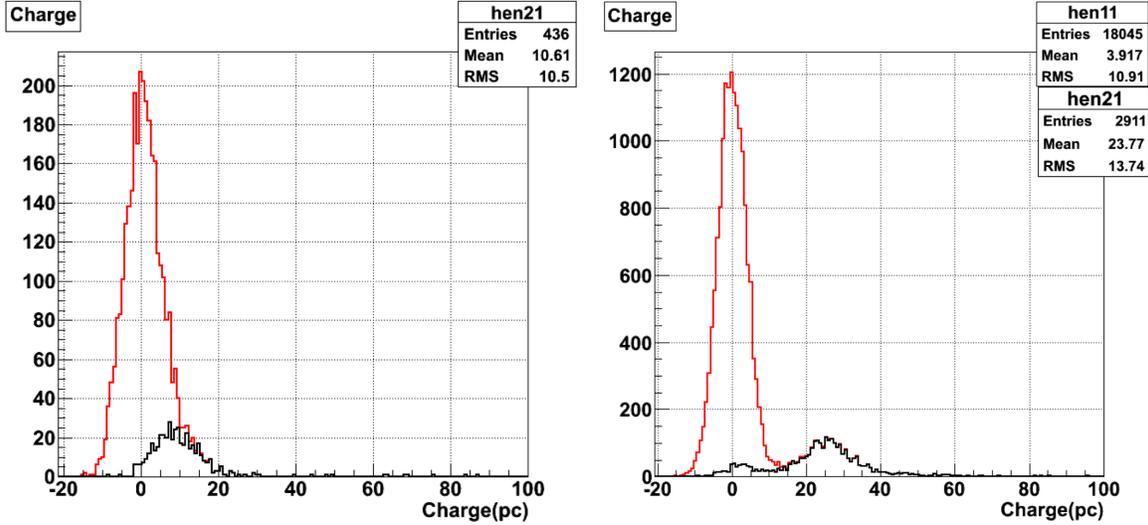


FIG. 20: Charge distribution from readout of the HPS calorimeter crystal with Hamamatsu S8664-55 (left) and S8664-1010 (right) APDs, and the new low noise amplifier board. The red line histogram corresponds to the charge distribution for all triggers coming from the scintillators positioned above and below the crystal. The black line shows the distribution for hits in the crystal within 100 ns of the trigger signal.

noise for the S8664-1010 APD readout. For the S8664-55 APD, the MIP signal is also seen, but its charge distribution is under the noise peak and it does not have well defined peak position. Using MIP calibration in conjunction with the light monitoring system will ensure stable and reliable performance of the ECal and the trigger system. As a bonus, the lower noise will allow the use of lower thresholds and improve the calorimeter's energy resolution. The new amplifier boards have to be designed to work with new APDs.

The total cost of replacing all ECal APDs is about 500K\$. The HPS collaborators from IPN Orsay have applied for a European Research Council (ERC) Advanced Grant 2013 to purchase APDs, for manpower costs to replace the old ones, to design and build new preamplifier boards, and to assemble and test the ECal with the new modules. This grant also includes light monitoring system. If successful the grant will cover most of the ECal modifications, and will significantly reduce the support needed from DOE for the more modest upgrades we have proposed above.

4.5 Muon system

The muon detector will match the geometrical acceptances of the tracker and ECal, and will be about a cubic meter in size. With such geometrical coverage, the efficiency of detecting high mass A' 's in the $\mu^+\mu^-$ decay channel will be higher than for e^+e^- decays, see Figure 21. The di-muon decay channel of the A' has the advantage of having greatly reduced electromagnetic backgrounds. In this case, the only particle background in a muon counter would come from the photoproduction of π^+ and π^- pairs that are not fully stopped in the ECal or absorber. The expected low background and the high detection efficiency make the di-muon final state an attractive complement to the e^+e^- final state. It will add substantial territory in the mass and coupling parameter space as show in Figure 22 and will offer valuable cross-checks. With the addition of a muon system, HPS will be the only fixed target experiment proposed to date to search for heavy photons in an alternative to the e^+e^- decay mode.

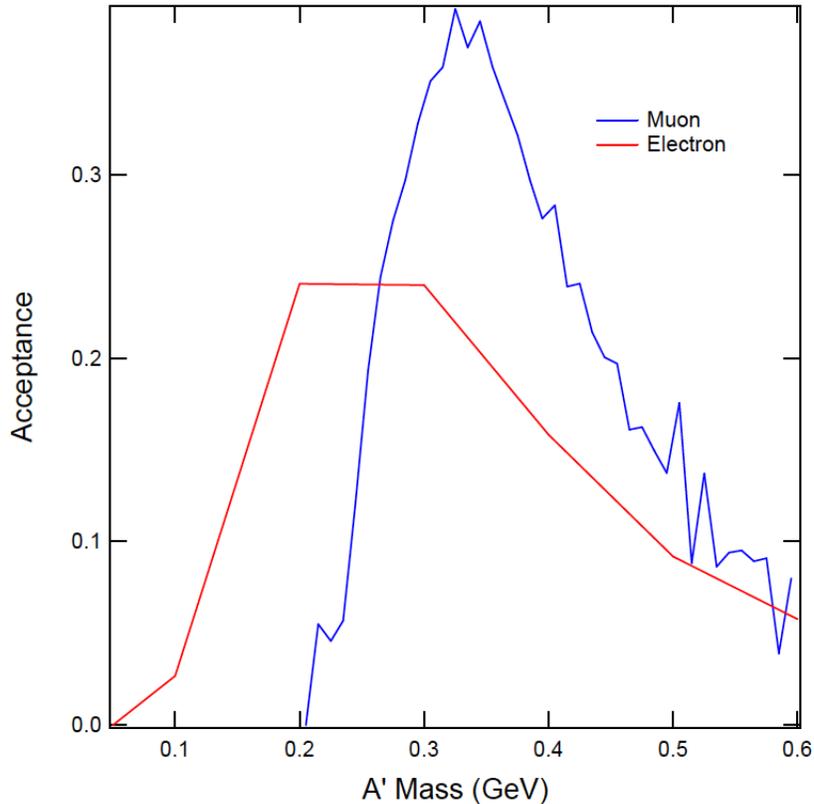


FIG. 21: A' detection efficiency through $\mu^+\mu^-$ (blue) and e^+e^- (red) decay channels as a function of mass for 6.6 GeV beam energy.

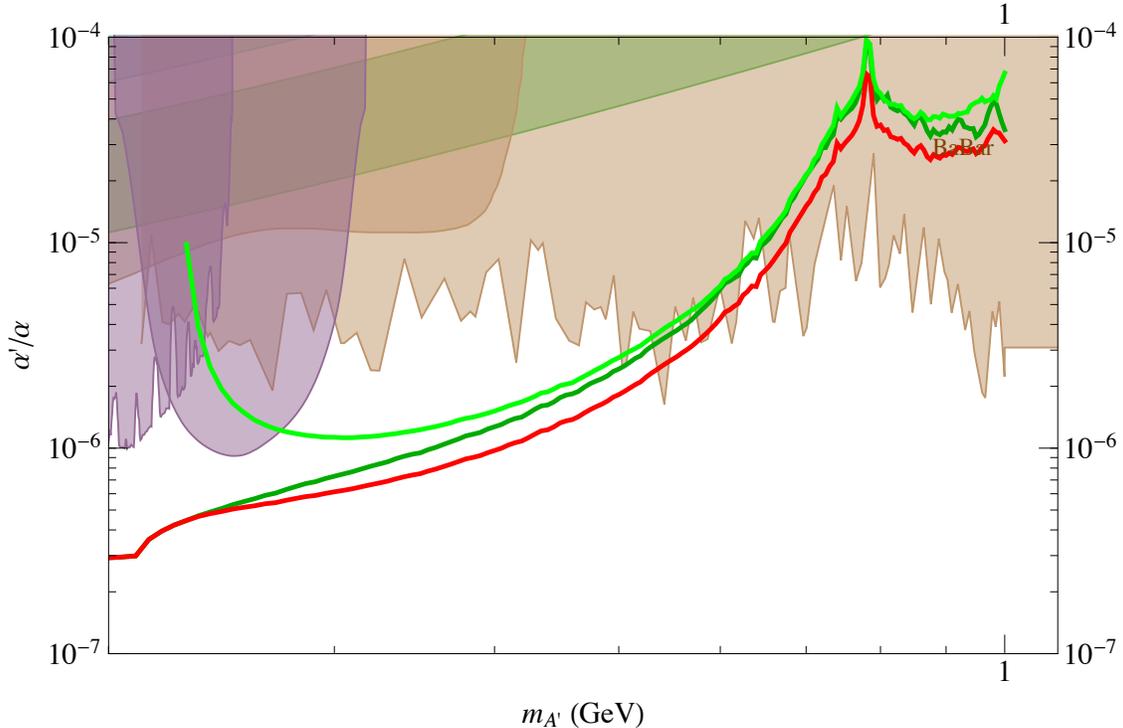


FIG. 22: Experimental reach for 2 weeks of beam time at 6.6 GeV for $\mu^+\mu^-$ (light green) and e^+e^- (dark green) decay channels. The red line is the combined reach.

The muon system can easily be constructed with layers of scintillator hodoscopes sandwiched between iron absorbers, and can easily be added downstream of the rest of the HPS apparatus. The number of layers and the thickness of absorbers is defined by the π/μ rejection factor. The schematic design of the muon detector was optimized using the GEANT-3 model for the ECal with added layers of iron and scintillators. In the simulation, muons and pions in the momentum range of 1 to 4 GeV/c first passed through the 16 cm of lead tungstate in the ECal and then entered a muon counter with various total absorber thicknesses (see [1] for details). Detection efficiencies for pions (ϵ_π) and muons (ϵ_μ) were then calculated as a function of absorber thickness and particle momentum. For low-energy particles (< 1.7 GeV) detection in all four layers of scintillator hodoscopes was not considered. Depending on the momentum, particles were not traced behind the third, fourth or fifth absorber. Figure 23 shows the resulting rejection factor $\epsilon_\pi/\epsilon_\mu$. The right-hand plot shows the dependence of $\epsilon_\pi/\epsilon_\mu$ on the total thickness of the iron absorber, with the best rejection at about 75 cm. The right-hand plot shows $\epsilon_\pi/\epsilon_\mu$ for a 75 cm absorber as a function of muon momentum. The suppression of individual pions by two orders of magnitude will suppress

pion pairs by 4 orders of magnitude.

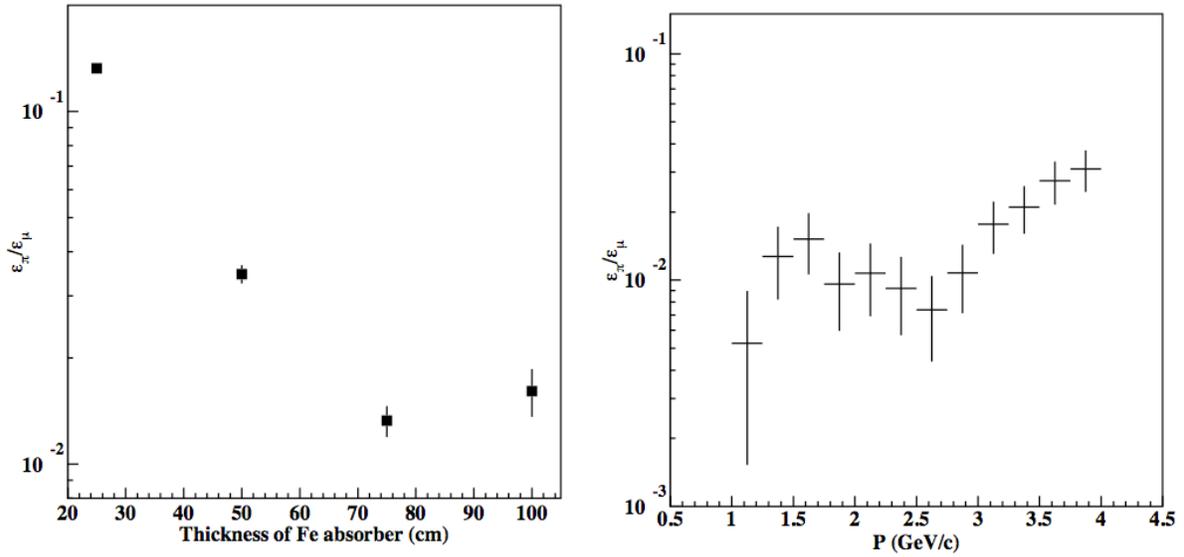


FIG. 23: Pion-muon rejection factor $\epsilon_\pi/\epsilon_\mu$ versus total iron absorber thickness (left) and versus particle momentum for a 75 cm absorber (right).

4.5.1 Conceptual Design

On the basis of these simulations, we have designed a muon detector composed of four iron absorbers (total length of $30 + 15 + 15 + 15 = 75$ cm) with a double-layer scintillator hodoscope positioned after each absorber. The muon detector will be mounted behind the ECal. The front face of the first absorber will be at ~ 180 cm from the target. Similar to the Ecal, the muon detector will consist of two halves, one above and one below the beam plane. This segmentation is necessary in order to minimize the effects of the “sheet-of-flame”, the multitude of low-energy particles in the horizontal plane, swept into the detector acceptance by the dipole analyzing magnet.

The dimensions of the hodoscopes and absorbers are defined using simulations of $A' \rightarrow \mu^+\mu^-$. In Figure 24 the points where muon pairs from A' decays of different masses intercept the (XY) plane located at 210 cm from the target are shown. Both muons are required to be detected in the ECal. Overall dimensions of the hodoscopes and absorbers are shown in Table III. Figure 25 shows a CAD drawing of the HPS detector with the muon system on the right, which includes the 4 absorbers (gray), the vacuum box (light gray) between the upper and lower sections, and the final set of scintillator paddles (red). The ECal is

directly upstream from the muon detector, with its crystals shown in yellow. In front of the ECal is a large gray vacuum flange. The silicon tracker is represented by red and gray rectangles and the red point on the left is the target position. The vertical gap between the first hodoscope layers of the two halves is about 5 cm and increases to about 7 cm for the last layer of hodoscopes.

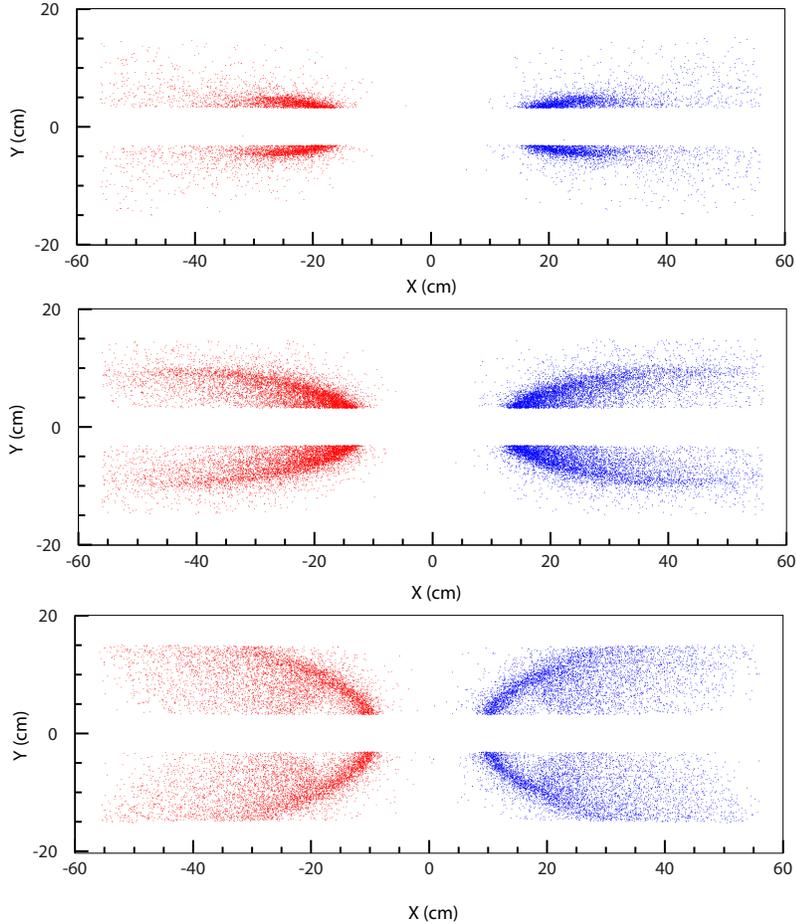


FIG. 24: Points where μ^+ (red) and μ^- (blue) pairs from A' decays intercept the (XY) plane located 210 cm from the target for A' masses 250 MeV (top), 300 MeV (middle), and 400 MeV (bottom).

For the hodoscopes we plan to use the same extruded scintillator strips with embedded wavelength-shifting fiber and phototube readout that was developed for the CLAS12 Preshower Calorimeter. These scintillator strips are 45 mm x 10 mm in cross section, and can be cut to any length. Widths can be reduced as needed for the muon counter. Each strip contains two long tunnels, created in the original extrusion process, into which wavelength shifting fibers can be inserted. Each hodoscope will consist of one X and one Y plane. The

TABLE III: Dimensions (in cm) of the muon system scintillation hodoscopes (H) and iron absorbers (A).

	H1	H2	H3	H4
Distance from target	212	232	252	272
Width	112	125	138.5	152
Height	10.5	11.5	12.5	13.5
	A1	A2	A3	A4
Distance from target	207	227	247	267
Width	108.5	122	135	148.5
Height	10	11	12	13
Thickness	30	15	15	15

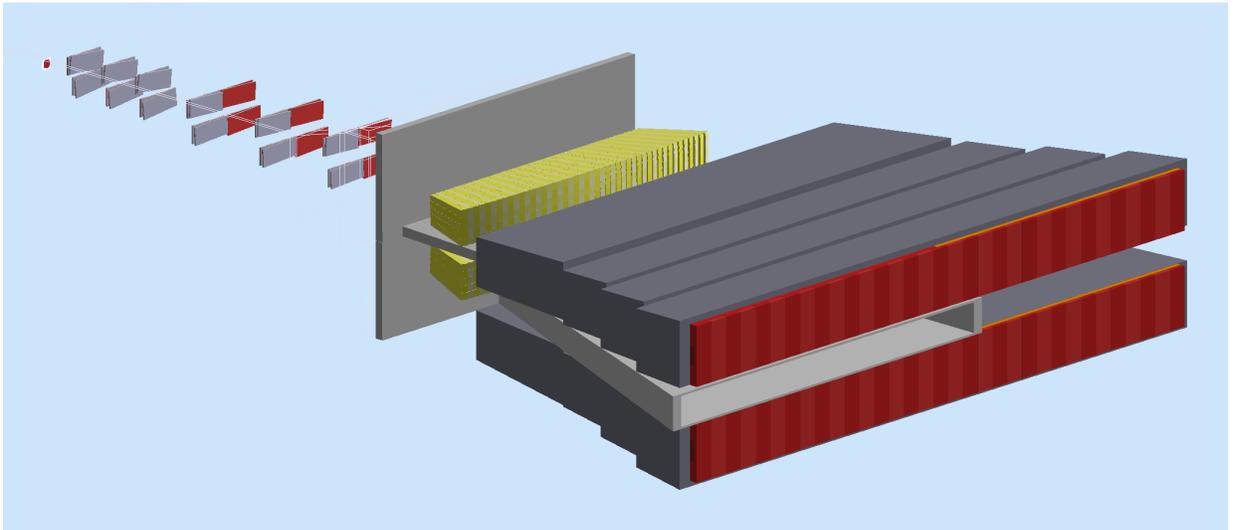


FIG. 25: CAD drawing of the HPS detector setup. From left to right this consists of the target (red dot), the silicon tracker (gray and red rectangles), the large shielding wall (gray), the ECal lead tungstate crystals (yellow, two shades), the muon counter absorbers (gray), and the final muon counter scintillators (red, two shades).

vertical strips of the last hodoscope are shown in two colors in Figure 25. The horizontal counters of the last hodoscope plane is shown in Figure 26. The horizontally aligned strips will extend over half the width of the detector and will be read out from the outer ends. The upper and lower hodoscopes in each plane will have their own vertically aligned strips, which will be read out only on their outer ends. The inner end is inaccessible because of the vacuum box, and there is no particular advantage to having a double readout on these short (135 mm) strips.

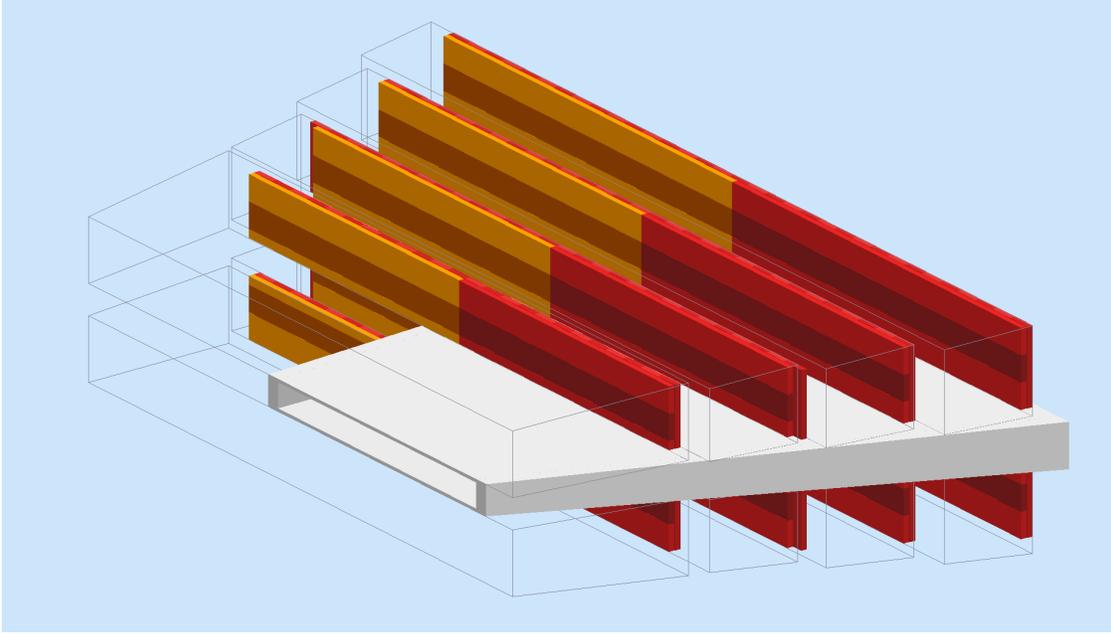


FIG. 26: The horizontal scintillator configuration for the muon counter. Scintillators are shown in red and yellow/brown. The white/gray structure is the vacuum box. Each hodoscope layer (top and bottom) contains three long strips, read out on both ends.

The system will be instrumented with less than 256 readout channels so that the requisite electronics will fit into a single VME/VXS crate. The signals from each channel (PMT) will be sent to a FADC. We intend to borrow the CLAS12 Preshower Calorimeter electronics and HV system. Just as with the ECal, the FADCs will be used to construct a muon trigger for the experiment. In the current design there will be 3 pairs (left-right) of horizontal strips in each of 8 hodoscope planes (48 total) and a total of 208 vertical strips in 8 hodoscope planes. The number of vertical strips per plane increases slightly with distance from the target to keep a constant angular coverage. The maximum is 33 per hodoscope in the back plane.

Full Monte Carlo simulations with realistic event rates are currently under study in order to finalize the design details of the muon counter. The crucial issues are the event rates in the scintillators near the beamline (which already has initiated a redesign of the vacuum chamber to reduce background), the target-to-muon-counter tracking resolution and the detection efficiency, and the achievable trigger rates. Any changes to the detector as a result of these studies are expected to be minor and will not alter the conceptual design presented here.

4.6 Trigger and DAQ

The HPS experiment data acquisition (DAQ) handles the acquisition of data for the three sub-detectors: the SVT, ECal and the Muon System. HPS employs two DAQ architectures: the SVT is readout with Advanced Telecom Communications Architecture (ATCA) hardware while the ECal and Muon system use VXS based hardware. The trigger system receives input from the ECal and Muon System, and distributes a trigger signal to all three detector sub-systems to readout a selected event. Figure 27 gives a schematic block diagram of the DAQ system.

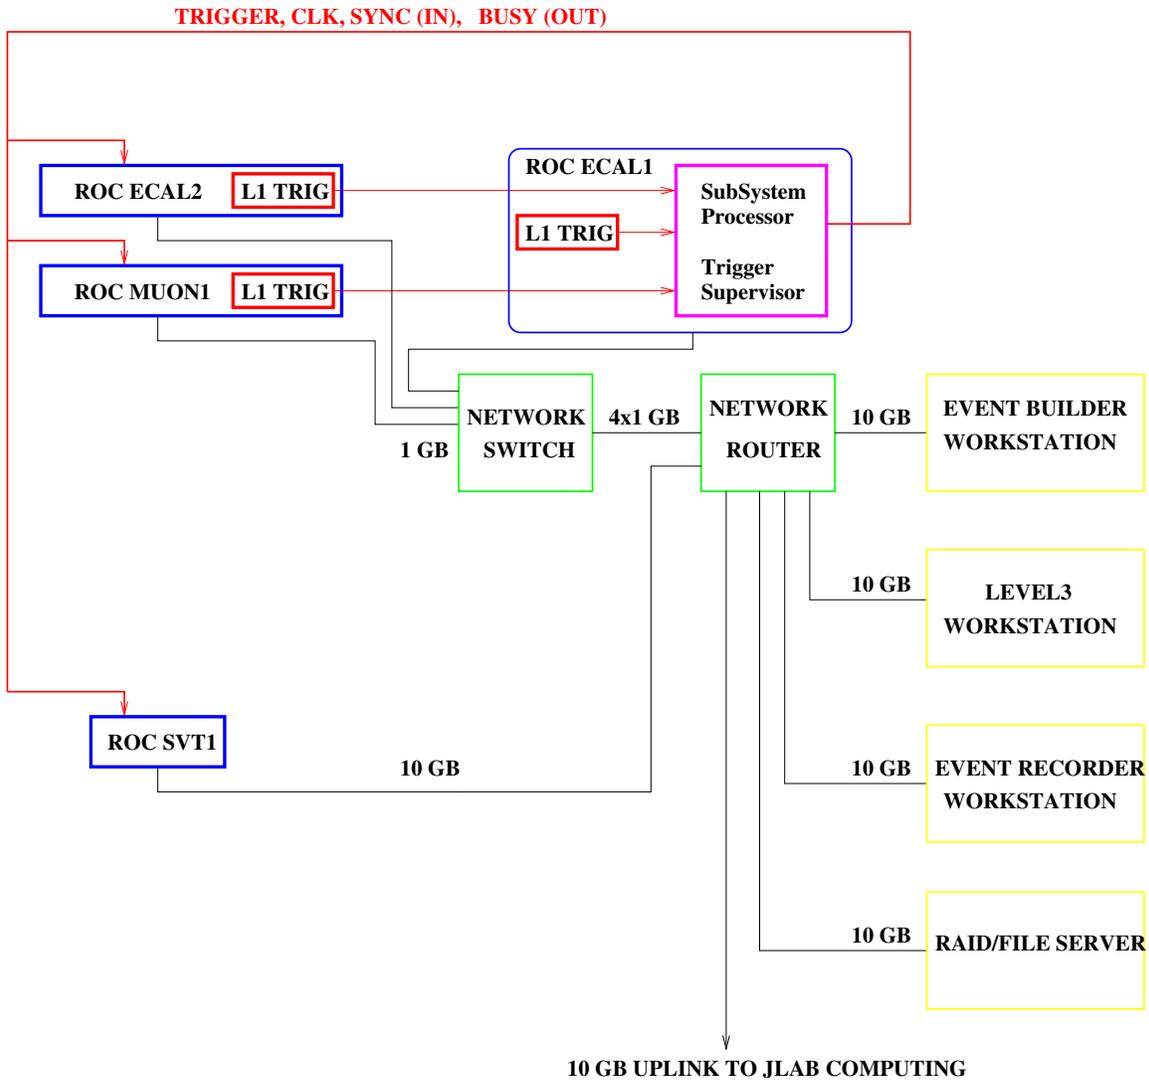


FIG. 27: Schematic block diagram of the data acquisition system.

For the ECal and Muon System, every VXS crate contains a Readout Controller (ROC)

that collects digitized information, processes it, and sends it on to the Event Builder. The ROC is a single blade Intel-based CPU module running DAQ software under CentOS Linux OS. For the SVT ATCA system, the ROC application runs on an embedded processor situated on the ATCA main board. The Event Builder assembles information from the SVT, ECal and Muon System ROCs into a single event which is passed to the Event Recorder that writes it to a data storage system capable of handling up to 100 MB/s. The Event Builder and other critical components run on multicore Intel-based multi-CPU servers. The DAQ network system is a Foundry router providing high-speed connections between the DAQ components and to the JLab computing facility. The SVT ROC, which must handle large data volumes, has a 10 Gbit/s link to the Foundry router, while a 1 Gbit/s link is adequate for the ECal and Muon System. A 10 Gbit/s uplink to the JLab computing facility is used for long-term storage.

Section 4.6.1 describes the SVT DAQ in more detail while the VXS based readout for the ECal and muon detector is described in Sec. 4.6.2. The trigger system is explained in more detail in Sec. 4.6.3.

4.6.1 SVT Data Acquisition

The goal of the SVT DAQ is to support the continuous 40 MHz readout and processing of signals from the 36 silicon strip sensors of the SVT. It also selects and transfers those events that were identified by the trigger system to the JLab DAQ for further event processing at rates close to 50 kHz. High data volumes result from high occupancy in the detector, pile-up from multiple bunches, and sampling pulse heights in six consecutive time buckets for each hit in order to facilitate reconstruction of the hit time to high accuracy.

The system adopted is an evolution of the SVT DAQ used for the HPS Test Run, described in Sec. 5.1.3. Several features changed in response to the new SVT design and the evolution of SLAC's ATCA system. The new SVT has nearly twice the number of sensors as the Test Run detector, necessitating a more compact way to transfer data and power to the individual sensor modules through the vacuum flanges. Accordingly, the new system incorporates a front end board within the vacuum volume for power distribution and signal digitization, allowing many fewer vacuum connections per sensor and less interference within the vacuum volume. Problems encountered with reflections on long twisted pair data lines,

although ultimately overcome, have been avoided altogether by incorporating a flange board just outside of the vacuum and adopting an optical link. The ATCA system has evolved to using optical input, so this change lets HPS optimally piggyback on SLAC's ATCA system development.

Each of the 36 silicon strip sensors is connected to a hybrid board incorporating five 128-channel APV25 front-end ASICs [98, 99], see Figure 42. The APV25 ASIC, initially developed for the Compact Muon Solenoid silicon tracker at the Large Hadron Collider at CERN, was chosen because it provides excellent signal to noise, analog output for optimal spatial resolution, and multi-bucket output for good time resolution. Each hybrid board has five analog output lines (one for each of the APV25 ASICs) which are sent to the front-end readout board using low power LVDS differential current signals over about 1 m of flex cable. At the readout board, a preamplifier scales the APV25 differential current output to match the range of a 14-bit Analog to Digital Converter (ADC). Each front-end board services four hybrids. The ADC operates at the system clock frequency of 41.667 MHz. The digitized output from the front-end board is sent through compact 8-pair mini-SAS cables to the vacuum flanges to connect to the external DAQ which resides outside the vacuum chamber. The front-end readout board houses a FPGA and buffers to allow for the control of the distribution of clock, trigger and I²C communication with the APV25 ASICs. To further simplify the services and minimize cabling that enter through the vacuum flanges, it contains linear regulators to distribute and regulate three low voltage power lines to the APV25 ASICs and the high voltage bias. Figure 28 shows a schematics layout of the downstream readout chain of the SVT.

The digitized signals are converted to optical signals just outside the vacuum flange on custom built flange boards. Each flange board houses optical drivers to handle the electrical-optical conversion and to transmit the optical signals over ~ 10 m fibers to the SVT DAQ. The flange board also interfaces the low- and high voltage power transmission from the CAEN power supplies to the vacuum chamber.

The main SVT DAQ uses the ATCA system for high speed data transfer. The optical signals from four hybrids, one half flange board, are received at one of four sections of the Rear Transition Module (RTM) boards of the ATCA crate, see Fig. 29. Each section of the RTM connects to one of four FPGA units on the main ATCA board, the COB (Cluster On Board). Each FPGA is housed on a separate daughter board called a Data Processing

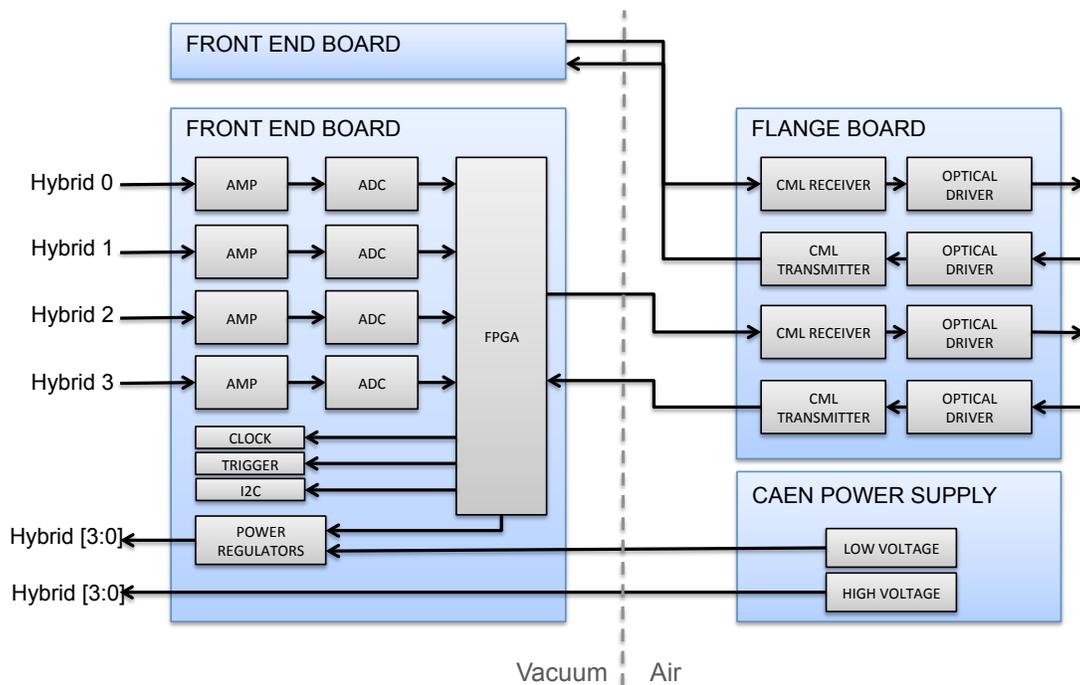


FIG. 28: Schematic overview of the front end and flange boards of the downstream part of SVT DAQ.

Module (DPM). The modular ATCA design permits the HPS DAQ to re-use architecture and functionality from other DAQ systems such as the ATLAS muon system whose components are similar to that used by HPS. Figure 43 shows the boards designed and used for the HPS Test run. In order to minimize the complexity of the system inside the vacuum chamber, all signal processing is done at the DPM. Each DPM receives the digitized signals from the RTM, applies thresholds for data reduction and organizes the sample data into UDP datagrams. One of the DPMs functions as the trigger interface which receives trigger signals from the optical fiber module on the RTM, distributes clock and trigger signals, and handles communication with the JLab trigger supervisor and the RCE. The RCE (Reconfigurable Cluster Element) is a generic computational building block on the trigger interface DPM running a 450 MHz PPC processor with 2 GB of DDR3 memory. Four COBs housed in two ATCA crates are sufficient to handle the 36 hybrids of the SVT.

The RCE receives and buffers UDP datagrams from the data and trigger DPMs and assembles them into full event frames. The RCE also runs an implementation of the JLab ROC application that integrates the SVT event frames into the JLab DAQ system described above. The RCE node transfers data to the JLab DAQ through a 10 Gbit/s Ethernet

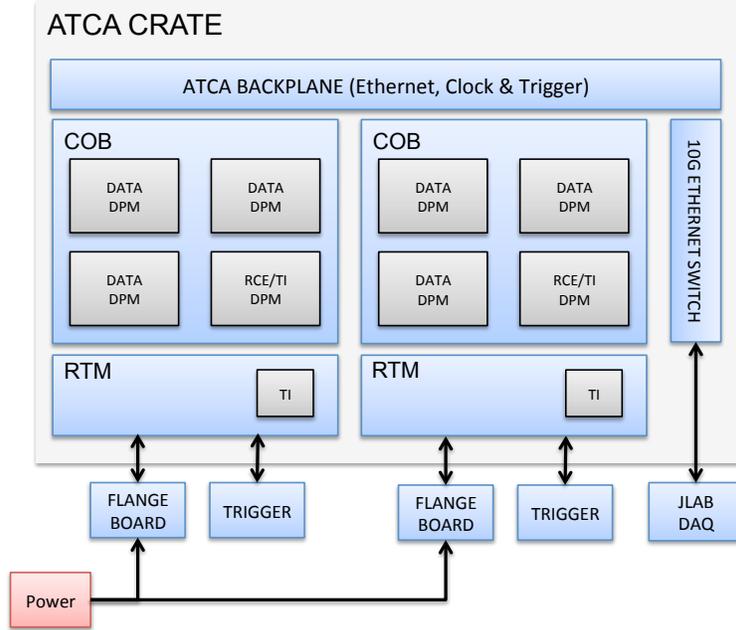


FIG. 29: Schematic block diagrams of the SVT data acquisition system.

backend interface. The maximum readout rate of the SVT is approximately 50 kHz, limited by the APV25 readout rate.

4.6.2 ECal and Muon Detector FADC Readout

The readout of the ECal and Muon System are essentially identical. The analog signals from the individual APD's of the ECal (shaped and amplified as described in Sec. 4.4) and phototubes of the Muon System are input to a single channel on the 16-channel JLab FADC250 VXS module (FADC), shown in Fig. 30.

Three 20-slot VXS crates are needed to accommodate the system: one for each half of the ECal with 221 channels and one for the muon detector with a total of 232 channels.

The FADCs store 12-bit digitized samples at 250 MHz in $8 \mu\text{s}$ deep pipelines. When a trigger is received, the appropriate part of the pipeline is accessed. If a FADC signal exceeds a predefined threshold within that time window, the integrated amplitude of a pre-defined number of samples before (NSB) and after (NSA) the signal passed threshold in addition to the time are recorded as explained in Fig. 31. This scheme significantly compresses the data input to the FADC. During data analysis, a pedestal value is subtracted to obtain the actual summed energy.



FIG. 30: A Jefferson Lab FADC250 VXS module.

The main characteristics of the FADC are:

- 12-bit digitizer with sampling rate of 250 Msps,
- 50Ω termination input,
- front-end input range: -0.5V , -1V or -2V (sufficient to avoid signal clipping for large pulse heights),
- nominal charge resolution between 10-39 fC per ADC (see Tab. IV).

Input range (V)	Nominal charge resolution (fC per ADC count)
-0.5	9.76
-1.0	19.53
-2.0	39.06

TABLE IV: Nominal FADC charge resolution for different front-end input ranges.

The FADC has two parallel data paths, shown in Fig. 31: the readout and trigger paths.

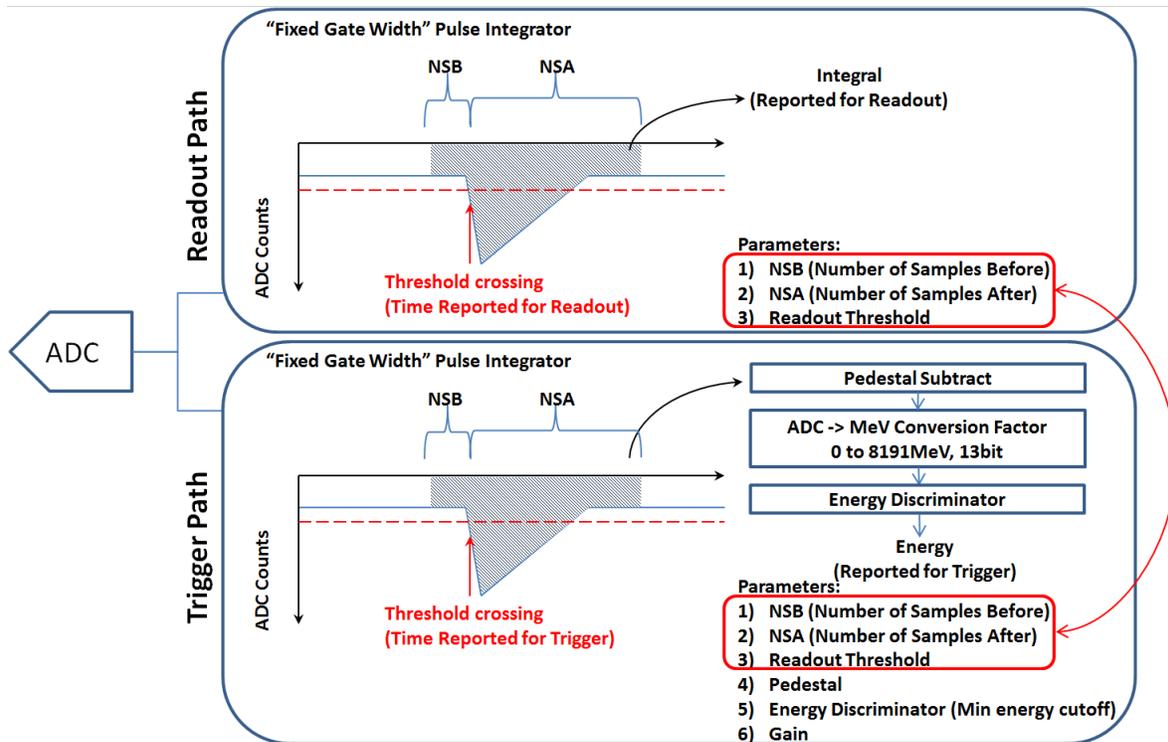


FIG. 31: FADC data paths

The trigger path runs continuously to report hits to the trigger system. The readout path only reports hits to the DAQ when the FADC receives a trigger.

For the readout path, every FADC has the following parameters:

- the number of samples integrated before the signal crossed threshold (NSB),
- the number of samples integrated after the signal crossed threshold (NSA),
- the readout threshold, measured in ADC counts.

The number of samples for a given channel integration is the sum of NSB+NSA samples. It is a fixed gate width pulse integration with no pedestal subtraction where the sum is stored in a 17-bit register for readout (pedestal subtraction happens offline).

For the trigger path, every channel has, in addition to NSB and NSA:

- trigger threshold, measured in ADC counts,
- a pedestal,
- a conversion factor (gain) that converts the ADC counts to energy in MeV (with 13 bits: from 0 to 8191 MeV),

- an energy discriminator threshold (minimum energy cutoff).

Note that the threshold for the trigger path can be set independently from the readout threshold. The pedestal value is subtracted from the integrated sum over NSB+NSA samples and converted to MeV units using a supplied gain conversion factor. The energy discriminator can be used to cut off low energy pulses before reporting to the Crate Trigger Processor (CTP). The values reported to the CTP are the 13-bit pulse energy and the time at which the pulse crossed the threshold. Data for every channel is sent to the CTP every 32 ns (if there is no hit a 0 energy pulse is sent) which sets a worst case double pulse resolution of 32 ns for individual channels, but less if pulses occur in adjacent 32ns windows.

4.6.3 Trigger System

The trigger system is designed to efficiently select e^+e^- and $\mu^+\mu^-$ events by using information from the ECal and Muon System. For e^+e^- events, the trigger looks for time coincidences of clusters in the top and bottom half of the ECal. The clusters also have to satisfy loose topological selections optimized on A' events to further reduce the rate. For $\mu^+\mu^-$ events, signals from at least the first two layers of the muon hodoscopes are combined with an ECal signal consistent with a minimum ionizing particle (MIP).

As described above in Sec. 4.6.2, the first stage of the trigger logic is incorporated into the FPGA's on the FADC boards which sends crystal energy and time information to the CTP. With the available 3-bit time information, the CTP can in principle look for time coincidence of crystal signals with 4 ns resolution (HPS will use a 8 ns time coincidence interval). The final trigger decision is made in the CTP and Sub-System Processor (SSP). The Trigger Supervisor generates all necessary signals and controls the entire DAQ system readout through the Trigger Interface (TI) units. The TI units are installed in every crate that participate in the readout process.

The trigger system is free-running and driven by the 250 MHz global clock and has essentially zero dead time at occupancies expected by HPS. The Trigger Supervisor can apply dead time if necessary, for example on a 'busy' or 'full' condition from front-end electronics. The system is designed to handle trigger rates above 50 kHz and a latency set to $\approx 3 \mu\text{s}$ to match that required by the SVT APV25 chip.

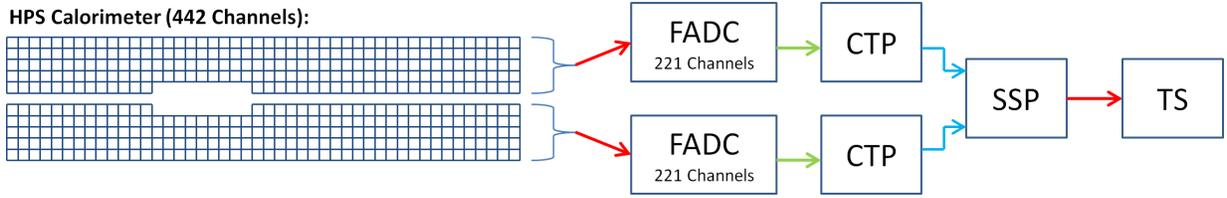


FIG. 32: Block diagram of the ECAL trigger system consisting of the FADC that samples and digitizes the detector channel signals and sends them for cluster finding in the CTP. The CTP clusters are sent to the SSP where the final trigger decision is taken based on pairs of clusters in both halves of the ECal. The decision is sent to the Trigger Supervisor (TS) that generates the necessary signals to readout the sub-detectors.

e^+e^- Trigger

The trigger system for e^+e^- events can be broken down into three sections (see Fig. 32):

- FADC (pulse finding): Samples and digitizes the signal pulses from each detector channel. Sends the measured pulse energy and arrival time to the CTP.
- CTP (cluster finding): Groups FADC pulses from each half of the ECal into clusters. The cluster energy, arrival time, and hit pattern are sent to the SSP.
- SSP (cluster pair finding): Searches for time coincidence of pairs of clusters from the top and bottom half of the ECal and applies topological selections.

The time coincidence window of pairs of clusters in the top and bottom half of the ECal are programmable with 4 ns resolution. The cluster finding algorithm is very fast and makes use of the parallel processing nature of FPGA's by simultaneously searching for 125 clusters, up to 3x3 in size, across the calorimeter crystal array, see Fig. 33. It performs the following

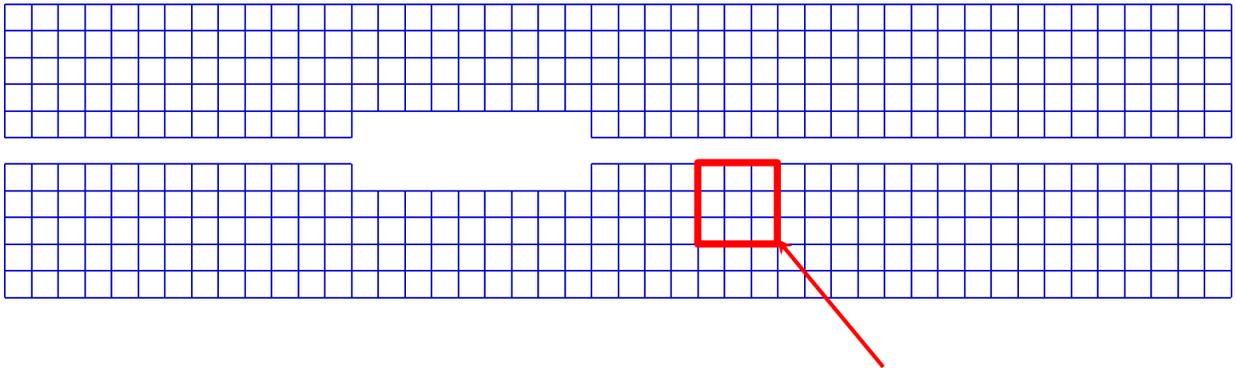


FIG. 33: Cluster finding algorithm.

tasks:

- Adds energy from hits together for every 3x3 square of channels in ECal.
- Hits are added together if they occur (leading edge) within a programmable number of 4 ns clock cycles of each other (HPS will use 8 ns time coincidence time interval).
- If the 3x3 energy sum is larger than the programmable cluster energy threshold and the sum is greater than any neighboring 3x3 windows, the CTP reports the cluster parameters (time, energy, position and 3x3 hit pattern) to the SSP.

The CTP evaluates all hits in its half of the calorimeter every 4 ns. A programmable time window is used to allow hits that are out of time with each other to be considered as part of a cluster sum. This is done by reporting hits when they occur and then reporting them again for the next N number of 4 ns clock cycles, where $N \in [0, 7]$. This is useful to deal with skew and jitter that develop from the detector, cabling, and electronics. As described above, the CTP only selects the 3x3 window with the highest energy sum of its neighbors. This filtering is applied to deal with overlapping clusters and cases where the cluster is larger than a 3x3 window.

The CTP sends the following information about the clusters to the SSP:

- 13-bit cluster energy (MeV).
- Cluster position (crystal index: x,y).
- Cluster time (with 4 ns resolution).
- Cluster 3x3 hit pattern (the detector channels reporting a hit in the cluster).

The cluster position is the coordinate of the peak crystal energy as seen from a 3x3 view. The 3x3 cluster hit pattern can be used by the SSP to help filter strange cluster patterns and/or make a low resolution cluster centroid computation. The SSP collects the cluster information from the two halves of the calorimeter and further applies topological selections optimized to further reduce background rates with small impact on the e^+e^- signal:

- Energy sum, $E_{min} \leq E_{top} + E_{bottom} \leq E_{max}$
- Pair time coincidence, $|t_{top} - t_{bottom}| \leq \Delta t_{max}$

- Energy difference, $|E_{top} - E_{bottom}| \leq \Delta E_{max}$
- Energy slope, $E_{cluster_with_min_energy} + R_{cluster_with_min_energy} \times F_{energy} \leq Threshold_{slope}$
- Co-planarity, $|\tan^{-1}(\frac{X_{top}}{Y_{top}}) - \tan^{-1}(\frac{X_{bottom}}{Y_{bottom}})| \leq Coplanarity_{angle}$
- Number of hits in 3x3 window, $\#hits_{3 \times 3} \geq HitThreshold$

where E_{max} , Δt_{max} , ΔE_{max} , $Threshold_{slope}$, F_{energy} , $Coplanarity_{angle}$ and $HitThreshold$ are programable parameters.

The SSP can also create a trigger decision based on the existence of a single cluster in the ECal exceeding the energy threshold which is useful for commissioning and calibration runs.

Online event analysis will be provided to be compared against trigger event data for immediate verification (on each trigger cut: cluster energies, positions, timing, energy slope, coplanarity and hit threshold). With identical trigger and data readout paths and high energy resolution, very precise agreement can be expected between trigger and readout.

$\mu^+\mu^-$ Trigger

The muon trigger will look for $\mu^+\mu^-$ pairs by finding energy depositions consistent with those expected from minimum ionizing particles in the layers of the muon system. The trigger algorithm in the CTP of the muon system VXS crate will produce a muon pair trigger in four steps:

- search for MIP hits using energy selections on the hits reported by the FADCs which satisfy $E_{MIP}^{thr} < E_{\mu_hodo}$
- use the time information of the reported MIP hits to select coincidences between the two planes of each hodoscope layer, quadrant by quadrant.
- look for coincidences in successive quadrants of at least the first three layers of the muon hodoscopes
- select pairs of triple coincidences in opposite quadrants of the detector and report the times and positions of coincident triple hits to the SSP

If it is necessary to reduce the rate further, the SSP can in addition look for time and position coincidences of MIP hits in the ECal, defined in the ECal crate CTPs as hits with 1 or 2 crystals and energy within predefined thresholds: $E_{MIP}^{min} < E_{ECal_channel} < E_{MIP}^{max}$. The SSP will send the final decision regarding the $\mu^+\mu^-$ trigger to the Trigger Supervisor board.

Diagnostic Tools

The previous experience with the similar (but much simpler) trigger system showed that diagnostic tools are necessary to make sure that the calorimeter and trigger electronics work as expected.

Scalers will be implemented for every ECal channel. The example of this diagnostic tool is presented in Fig. 34 from the previous version of the calorimeter. Hot or dead channels are easily identified and disabled online. Diagnostic scope permits to analyze on-line the

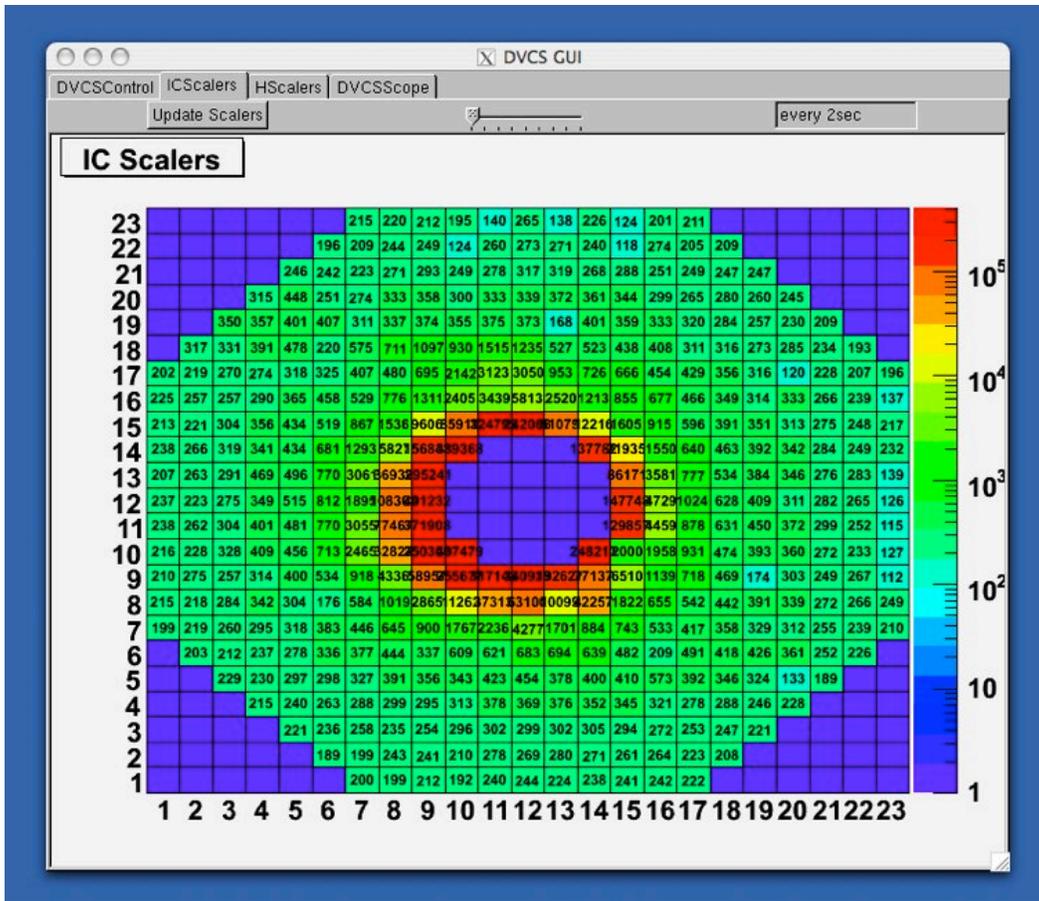


FIG. 34: Scalers (example from the previous version of the calorimeter).

trigger logic. The goal to have the Two-Dimensional Analyzer is to provide a remote debug

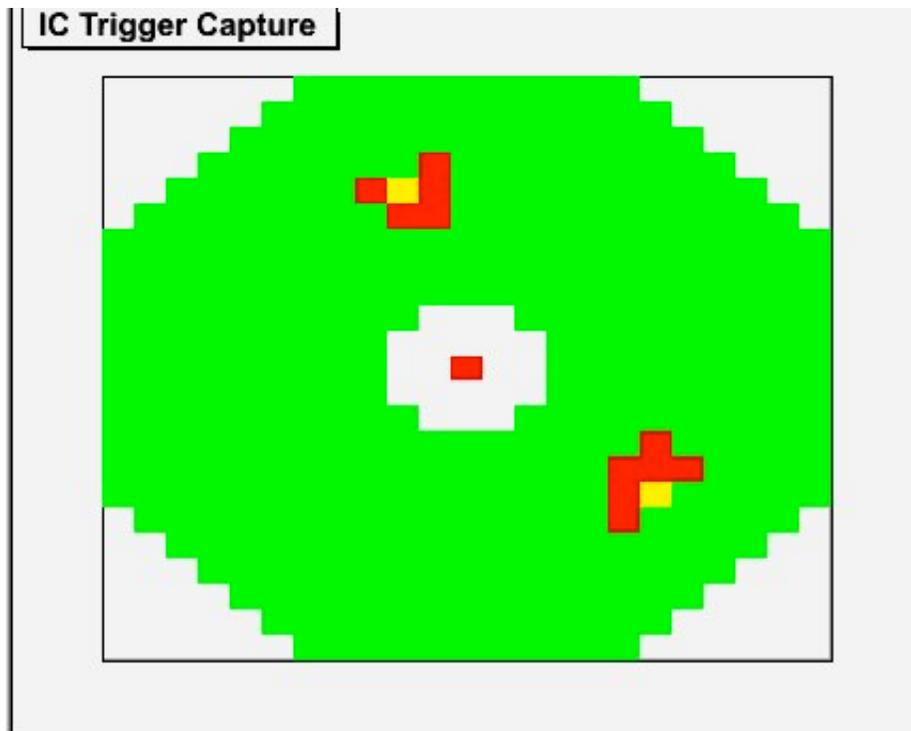


FIG. 35: Diagnostic scope (example with two clusters found from the previous version of the calorimeter). Green - no hits, red - tower with hit, yellow - cluster found.

interface to identify bad channels, verify cluster finding algorithms and check timing. The details of this analyzer are as following:

- Logic analyzer runs in parallel, non-intrusive, to the calorimeter trigger
- Can setup trigger on any ECal pixel arrangement and/or cluster count
- Can move forward/backward in time by 250 ns to see timing details
- Will be customized for HPS geometry and hardware

The example of the 2D analyzer is presented in Fig. 35. Two clusters are displayed in the picture. The red color displays the hits in the calorimeter and the center of clusters is displayed in yellow.

In addition to the scalers, the distributions on individual ADC channel pulse energy will be provided. The cluster hits positions and energy from SSP processor will be histogrammed as well. Two histograms (accepted and rejected) will be provided for each trigger cut used in the trigger logic.

	Occupancy(%)			Event size (kB)			Data rate (MB/s)		
Beam energy (GeV)	1.1	2.2	6.6	1.1	2.2	6.6	1.1	2.2	6.6
SVT	0.5	0.3	0.3	2.5	1.7	1.5	43.1	27.2	18.9
ECal	3.0	4.2	4.7	0.3	0.3	0.3	4.9	4.8	3.9
Muon	10.0	10.0	10.0	0.2	0.2	0.2	3.8	3.4	2.7
Total	-			3.0	2.2	2.0	51.9	35.4	25.6

TABLE V: Summary of the occupancy, event size and data rate expected for the runs at the three beam energies in the run plan.

4.6.4 Event Size and Data Rates

The high occupancies in the detector require a high readout bandwidth to be able to transfer hits from the detectors to disk. The event sizes and rates are based on estimates from full Geant4-based simulations including all known backgrounds. As expected the SVT dominates the expected rates. The noise hit occupancy in the SVT is kept low by requiring that three of the six samples are above an effective threshold of three times the noise level. The dominant contribution to the occupancy is instead the estimated high rate of beam background hits. This is estimated from detailed full simulation resulting in an occupancy of around 0.3% or an average of 61 channels above threshold. Each SVT channel has, in addition to the six digitized samples, header information that identifies the the channel number and it's chip address. The complete SVT event size also include the overhead from each FPGA and the JLab data stream bank header. The maximum average event size increased with decreasing beam energy since a larger fraction of backgrounds get larger opening angles and thus potentially higher than the 15 mrad vertical dead zone angle. For a beam energy of 1.1 GeV, the average SVT event size is 2.5 kB and the rate is 43 MB/s, well within the SVT DAQ capabilities. The ECal and Muon System each contribute with an event size of approximately 0.3 kB and maximum rates of about 5 MB/s. Further studies will determine the exact Muon System occupancy. For the data rates a conservative estimate of 10% is used. The rates expected are well within the 100 MB/s limit for each VXS crate used in the ECal and muon DAQ system. Table V summarizes the event size and data rates. The highest overall rate, for a 1.1 GeV run, and that needs to be written to disk is 52 MB/s which is within the current DAQ system design limit of 100 MB/s.

Run	E_{beam} (GeV)	Time (days)	Events ($\times 10^9$)	Raw data (TB)	Processed data (TB)
2014	1.1	21	33	100	445
2014	2.2	21	29	63	282
Total	-	42	62	163	727
2015	2.2	35	48	105	470
2015	6.6	35	38	76	341
Total	-	70	86	181	810

TABLE VI: Summary of the raw and processed data expected from the HPS runs.

Beam energy	Raw (EVIO) event size (kB)	Reconstructed (LCIO) event size (kB)
1.1 GeV	3.0	13.4
2.2 GeV	2.2	9.8
6.6 GeV	2.0	8.9

TABLE VII: Data event sizes.

4.7 Offline Computing Model

The following is an outline of the offline computing model envisioned for satisfying the analysis needs of the HPS experiment. The raw data collected over the running periods must be processed through calibration passes, reconstructed into calibrated objects useful for analysis and separated into analysis streams. Corresponding Monte Carlo will also need to be produced and separated into the same analysis streams.

Table VI shows the expected number of triggered events and the total amount of data expected over the different runs. For the six-week run in 2014 we expect to collect approximately 163 TB of raw data. We assume that the experiment collects data for all of its available beam time and the time allocated for detector commissioning, even though the experiment reach only assumes 50% availability; this gives a conservative estimate of computing requirements. For trigger rate estimates, we use the ECal trigger rate from Section 6.2; based on 6.3, the muon system trigger rate is expected to be negligible.

The raw data must be processed to produce physics data objects that can be analyzed. This reconstruction process will also include filters to select events of physics interest. We use the event size estimates in Table VII, which are based on Table V from the previous section.

For modeling signals, estimating backgrounds and confirming the understanding of the detector performance, extensive Monte Carlo simulation is needed. Table VIII summarizes

Event type	Sim. stage	Size/triggered event (kB)	Mass points
Beam bkg.	evgen	37.0	1
A' signal	evgen	0.5	10
A'+beam bkg	evgen	37.4	10
Beam bkg.	MC output	79.5	1
A' signal	MC output	2.5	10
A'+beam bkg	MC output	82.0	10

TABLE VIII: Event sizes in kB per triggered event, including pileup events for beam background.

the typical event size at the various stages of the simulation.

In total approximately 1/10th of the number of data events collected need to be simulated. For the 2014 run this amounts to 6.2 billion events. In addition, 100 million A' events must be simulated at each of 10 mass points at each beam energy. This totals 2 billion A' events in each of 2014 and 2015, assuming that changes in running conditions prevent us from using the same simulated data set for both 2.2 GeV run periods.

In total 1855 (2236) Tb of storage for data (raw, reconstructed and simulated) is needed for the 2014 (2015) run.

Tape is currently the only economical storage solution for storing all of the raw, simulated and processed data. The processing of the raw data is foreseen to occur at JLab. Given a typical bandwidth between sites of 3 to 4 Tb/day, analyses needing access to the hit level information will need to be run at JLab or run on small samples of exported data unless they can take advantage of the analysis streams. Analysis streams of events satisfying pre-selections criteria for targeted analyses will be exported to remote sites. Likewise, the size of the simulated data samples suggests that the simulation should be processed and stored at JLAB and that only analysis streams or small samples of the full data will be exported.

Disk space at JLAB and SLAC will be needed for ntuples, code releases, and scratch areas. Additionally, disk space at SLAC will be needed for the imported analysis streams. It is estimated that 100 TBytes will accommodate the needs for for the 2014 data analysis work and an additional 100 TBytes would be needed for 2015.

The HPS storage requirements are summarized in Tab. IX.

Approximately 0.1 CPU seconds are needed to reconstruct a data event on a typical 2.4 GHz core. This would require a total of 2.2 million CPU hours of processing for the entire 2014 dataset at the JLab processing center. To simulate events, approximately 0.02

Storage category	2014 (TB)	2015 (TB)
Raw data	163	181
Processed raw data	727	810
Simulated data	965	1244
Total tape space	1855	2236
Disk space	100	100

TABLE IX: Data storage summary.

Computing category	2014	2015
Raw data processing ($\times 10^6$ CPUh)	1.7	2.4
Simulation production ($\times 10^6$ CPUh)	8.8	10.1
Total ($\times 10^6$ CPUh)	10.5	12.5

TABLE X: Computing needs summary.

CPU seconds are needed for a beam event and approximately 0.7 seconds for an A' event. In total 8.8 million CPU hours are needed for Monte Carlo simulation for the 2014 run. Based on experience with previous experiments, it is reasonable to estimate that the net CPU needed for analysis work (batch and interactive) will be comparable to that needed for production. The HPS offline computing requirements are summarized in Tab. X.

The Jefferson Lab Computing Center provides computing and storage for experiments at JLab. A request will be submitted for data storage (tape and disk), computing resources (CPU hours for simulation and production), and data transfers to/from JLab.

5 May-2012 Test Run

The HPS Test run was proposed to DOE early in 2011 as the first stage of the HPS experiment. Its purposes included demonstrating that the apparatus and data acquisition systems are feasible and that the trigger rates and occupancies encountered in electron-beam running are as simulated. It also provided valuable experience to the HPS Collaboration in all aspects of designing, building, installing, and running the experiment at JLab. Furthermore, in the case that HPS Test detector met all of the performance goals and was given dedicated running time with an electron beam, the HPS Test Run could provide new sensitivity to heavy photons. The HPS Test apparatus was installed on April 19, 2012, and ran parasitically with the HDice experiment, using a photon beam, until May 18. The JLab

run schedule precluded any dedicated electron beam running, but the HPS Test Run was allowed a short and valuable dedicated run with the photon beam.

This section briefly reviews the HPS Test Run apparatus, a simplified version of that planned for the full HPS experiment, and demonstrates the feasibility of the detector technologies proposed for silicon tracker, ECal, and data acquisition systems. It documents the performance of the trigger, data acquisition, silicon tracker, and ECal and shows that the performance assumed in calculating the physics reach of the experiment is realistic. Of particular importance, data from dedicated photon beam running has been used to compare the observed trigger rates with that expected in simulation. The trigger rate is almost entirely due to photons which have converted to e^+e^- upstream of HPS and is sensitive to the multiple Coulomb scattering of electrons and positrons in the conversion target. Since scattered primary beam is the dominant source of occupancy in running HPS in an electron beam, good agreement between data and simulation confirm the background simulation used to benchmark the physics reach of the HPS experiment.

In addition to this important test of our background simulation, the test run accomplished the following goals which are explained below in Sec. 5.3:

1. More than 97% of SVT channels functioned properly
2. SVT readout signal to noise of 25.5
3. SVT hit time resolution of 2.6 ns; this proves hit time reconstruction will work for HPS
4. SVT hit efficiency greater than 98%
5. Survey-based SVT alignment performs as expected and will allow track-based alignment
6. 87% of ECal crystals functioned properly, and defects will be corrected by planned ECal upgrades
7. Calibrated ECal using SVT tracks
8. Fully integrated SVT and JLab DAQ
9. Trigger functions as designed; FADC trigger rate tested to greater than 100 kHz

5.1 HPS Test Run Apparatus

In Figure 36, the layout of the parasitic run is shown. The silicon vertex tracker was installed inside the Hall B pair spectrometer magnet vacuum chamber with the electromagnetic calorimeter mounted downstream. Both the tracker and the ECal were retracted off the beam plane to allow clean passage of the photon beam through the system.

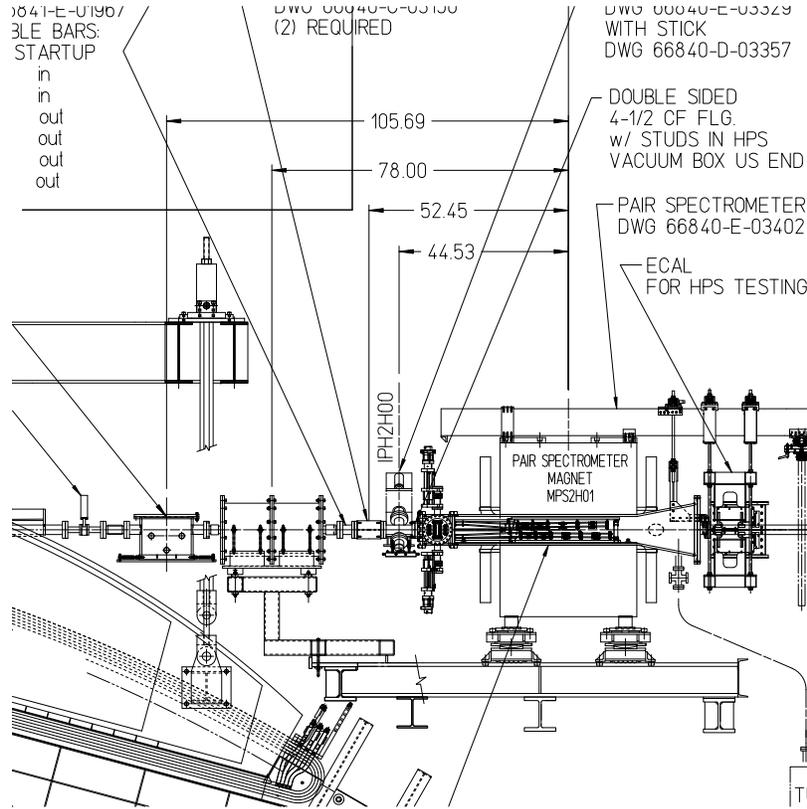


FIG. 36: Layout of the HPS parasitic run.

For dedicated HPS running the photon beam was generated in the interaction of the 5.5 GeV electrons with a gold radiator of 10^{-4} r.l., located ≈ 8 meters upstream of the PS pair converter. After collimation ($D = 6.4$ mm), the photon beam passes through the pair converter and through the HPS system. The converter was located ≈ 77 cm upstream of the first layer of the silicon vertex tracker. Data were taken on different converters (empty, 1.8×10^{-3} r.l., 4.5×10^{-3} r.l., and 1.6×10^{-2} r.l.). These measurements were repeated for the reverse field setting of the pair spectrometer dipole.

5.1.1 Test Run SVT

The silicon tracking and vertexing detector for HPS Test, or SVT, operates in an existing vacuum chamber inside the pair spectrometer analyzing magnet in Hall B at JLab. The design principles of the SVT are described in further detail in the HPS Test Run Proposal [2]. There are five measurement stations, or “layers,” placed immediately downstream of the target. Each layer comprises a pair of closely-spaced planes and each plane is responsible for measuring a single coordinate, or “view.” Introduction of a stereo angle between the two planes of each layer provides three-dimensional tracking and vertexing throughout the acceptance of the detector with one redundant layer.

In order to accommodate the dead zone, the SVT is built in two halves that are mirror reflections of one another about the plane of the nominal electron beam. Each half consists of five double-sided modules mounted on a support plate that provides services to the modules and allows them to be moved as a group relative to the dead zone. Each module places a pair of silicon microstrip sensors back-to-back at a specified stereo angle with independent cooling and support.

Modules with 100 milliradian stereo are used in the first three layers to provide higher-resolution 3-d space points for vertexing. The 50 milliradian stereo of the last two layers breaks the tracking degeneracy of having five identical layers and minimizes fakes from ghost hits, improving pattern recognition while still providing sufficient pointing resolution into Layer 3 for robust hit association in the denser environment there. These stereo angles are the same as those proposed in Section 4.3 for the new SVT. The details of the five layers are shown in Table XI and a rendering of the detector layout is shown in Figure 37. Figure 38 shows a photograph of both completed detector halves prior to final assembly. Altogether, this layout comprises 20 sensors and hybrids and 100 APV25 chips for a total of 12780 readout channels.

Power is provided to each hybrid using CAEN power supplies on loan from Fermilab. Three low voltages are supplied for the APV25 and one high voltage to reverse bias the sensor. The supplies that provide sensor bias are capable of 500V operation and can be used to test operation at high voltage in close proximity to an electron beam. The total power consumption of each hybrid during normal operation is approximately 1.7 W, which is removed by the cooling system. Care was exercised in selecting power and data cables

Layer	1	2	3	4	5
Nominal z , from target (cm)	10	20	30	50	70
Stereo angle (mrad)	100	100	100	50	50
Bend plane resolution (μm)	≈ 60	≈ 60	≈ 60	≈ 120	≈ 120
Non-bend plane resolution (μm)	≈ 6	≈ 6	≈ 6	≈ 6	≈ 6
# sensors	4	4	4	4	4
Nominal dead zone (mm)	± 1.5	± 3.0	± 4.5	± 7.5	± 10.5
Power consumption (W)	6.9	6.9	6.9	6.9	6.9

TABLE XI: Layout of the HPS Test SVT.

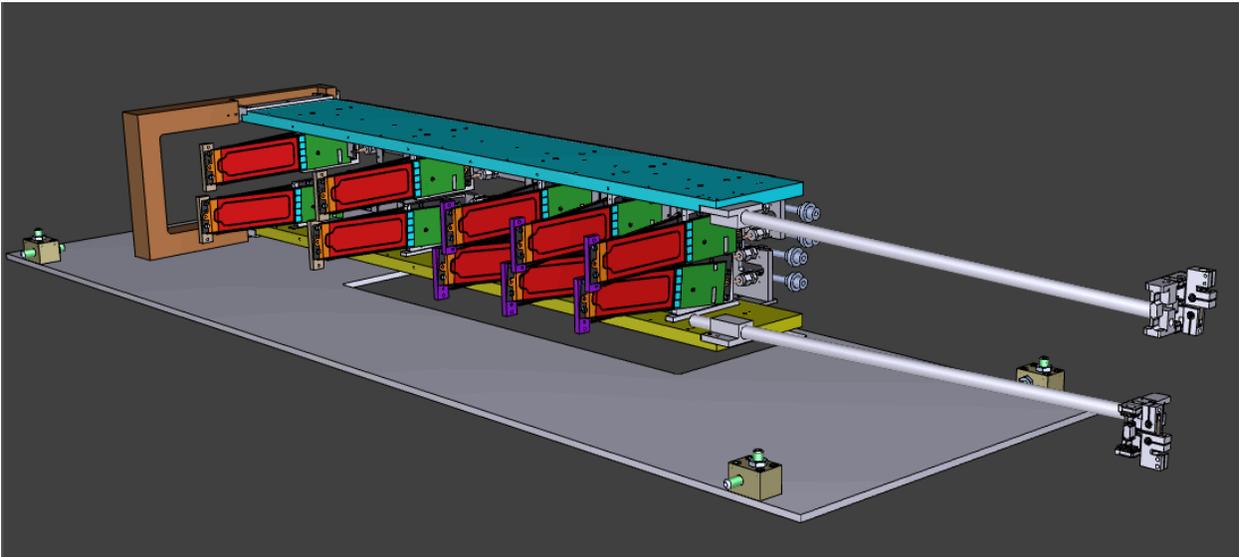


FIG. 37: A partial rendering of the HPS Test SVT solid model showing the modules of the upper and lower half-detectors on their support plates, the hinged C-support, the motion levers, the cooling manifolds on their strain relief plate and the baseplate with its adjusters. The sensors are shown in red and the hybrids in green. The beam enters from the right.

to ensure vacuum compatibility and sufficient radiation hardness. A custom junction box interfaces the CAEN power supply output channels to the SVT hybrids. Control of the supplies is provided via an EPICS graphical user interface, which allows monitoring of the detector and interlock protection.

The linear shifts that define the opening of the SVT are controlled by a pair of stepper motors located in low field regions at the ends of the analyzing magnet. For photon running, these are locked in the open position, but for electron running they will be connected and controlled through EPICS so that the distance between the beam and the sensors can be adjusted to balance detector occupancies and acceptance.



FIG. 38: Both halves of the HPS Test SVT fully assembled at SLAC.

5.1.2 Test Run ECal

The electromagnetic calorimeter (ECal) for HPS, as described in Section 4.4, was built and tested in the test run. The only differences between the test run ECal and what is proposed here for HPS are in the position and the vacuum chamber. The vacuum chamber between the two ECal modules was not used for the photon test run; instead a 2" beam pipe was used to transport photon beam from the pair spectrometer vacuum chamber to the HDIce target. The ECal was mounted downstream of the analyzing dipole magnet at the distance of ~ 148 cm from the upstream edge of the magnet. The two ECal modules were positioned just above and below the beam pipe such that the edge of the crystal closest to the beam was 3.7 cm from it.

For the test run, the ECal made use of the existing low and high voltage systems from the CLAS IC, as well as signal cables and splitters. Connectors on the existing signal cables were rearranged to accommodate the new layout of the channels.

Assembly of the bottom half to the ECal is shown in Fig. 39. Figure 40 shows the ECal



FIG. 39: Assembly of the ECal bottom module.

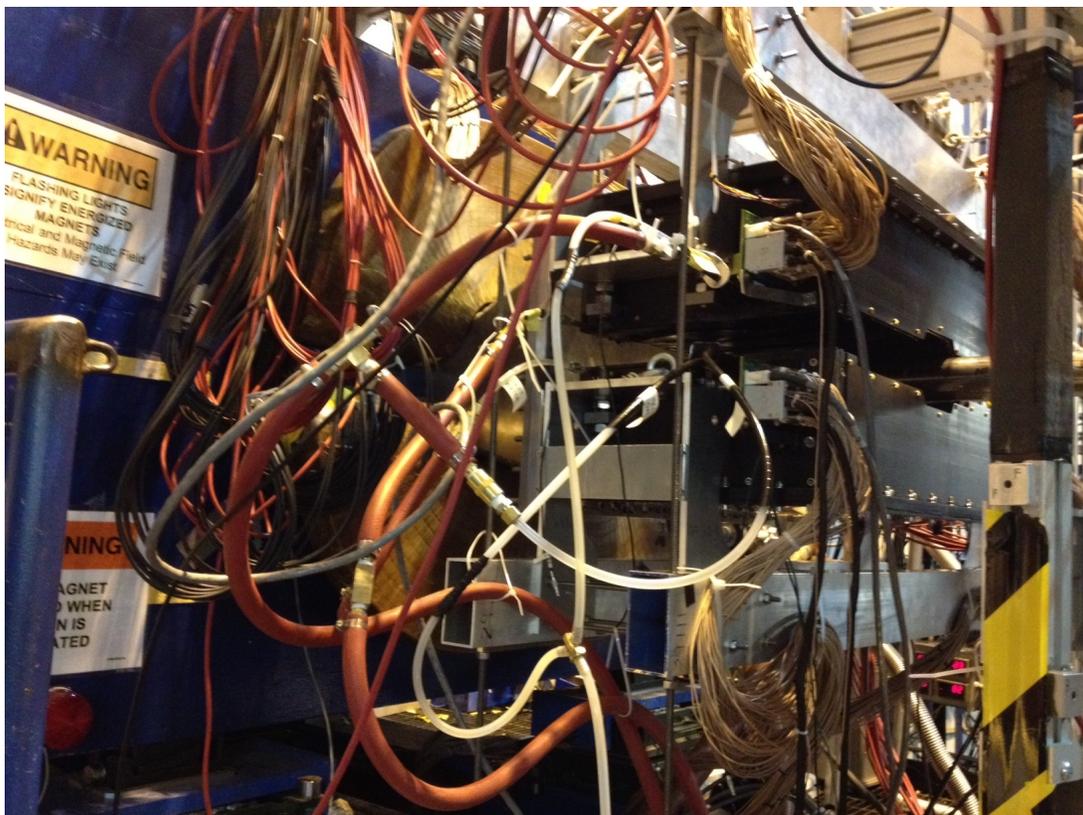


FIG. 40: ECal mounted downstream of the Hall-B pair spectrometer for the parasitic run with photon beams. Hoses for the cooling system, and the power and signal cables for beam-right side of both modules are visible.

in its installed position for the parasitic run with photon beam.

5.1.3 Test Run Data Acquisition

The data acquisition system (DAQ) for the HPS Test run was a somewhat simplified version of the DAQ proposed for HPS in Sec. 4.6. In addition to simpler trigger logic, the primary difference for the ECal is a different signal processing for the trigger which results in slightly worse single-crystal energy resolution and no possibility to calibrate individual crystal gains at the trigger level. For the SVT, the smaller number of channels eliminated the need for optical conversion and aggregation of data and detector power inside the vacuum chamber. Finally, most data links had 1 Gbit bandwidth, sufficient for the purposes of the test run.

In other respects, the HPS Test run DAQ was essentially identical to that proposed for HPS and used to verify the overall technical approach of the system. The ECal provides data to the FADC-based Level 1 trigger system. Accepted events are read out from the ECal and SVT DAQ and processed by an event builder before output to disk. For the ECal, the Readout Crate Controllers (ROCs) were the same as those proposed for HPS and are installed in VME, VME64X and VXS crates running mvme6100 controllers, with prpmc880 or pmc280 co-processor modules. For simplicity, a hybrid approach was used for the SVT DAQ in the test run where the ROC ran on an external PC connected to the ATCA crate. Similar to that proposed for HPS, a Foundry Router was used as the backbone of the DAQ system, providing 1 Gbit and 10 Gbit connections between components and to the JLAB Computer Center. The Event Builder, Event Recorder, and other critical DAQ components ran on 4-CPU Opteron-based servers, which was sufficient for the test run. The RAID5 storage system had 100 MB/s capability, sufficient to handle the anticipated data rates for electron running.

The SVT DAQ system was based on the same overall architecture as that described in Sec. 4.6.1, see Fig. 41 for a schematic layout. Being only half the size of the HPS SVT, the largest difference is the provision for individual power and data for each sensor and readout chip from the power supplies and DAQ outside of the vacuum chamber. This simplification reduced development time and cost at the expense of a large number of connections and vacuum feedthroughs. With a total of 20 silicon microstrip sensors, each one connected to an onboard hybrid readout card hosting five 128-channel APV25 ASICs (shown in Figure 42), 600 lines for power and data are required. Without optical conversion inside the chamber

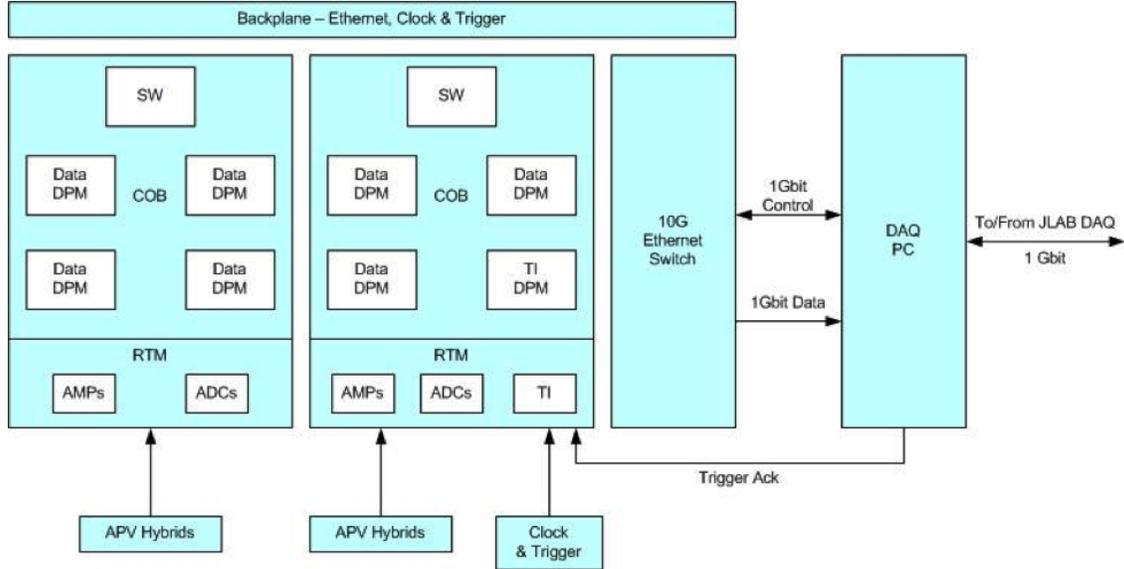


FIG. 41: Schematic of the SVT DAQ used for the test run. Note that the hybrids are connected directly to the RTM and that an external DAQ PC is used for control and transfer of data to the JLab DAQ.

as proposed for HPS, analog signals from the APV25 chips are carried directly via multi-twisted-pair cable to the Rear Transition Module (RTM), see Fig. 43, in the ATCA crate located outside the vacuum chamber. On the RTM, a pre-amplifier converts the APV25 differential current output to a voltage output scaled to the sensitive range of a 14-bit ADC. The RTM is organized into four sections with each section supporting 3 hybrids (15 channels). The ADC is operated at the system clock of 41.667 MHz. The ATCA main board, the Cluster on Board or COB (see Fig. 43), is similar to the one described for the HPS DAQ with the important exception that one of the Data Processing Modules (DPMs) functions as the trigger interface and there is no Reconfigurable Cluster Element (RCE) module. Instead, the DPMs package and send the data from the hybrids through a 1 Gbit ethernet connection to an external PC which serves the same purpose as the RCE module in the HPS DAQ. The ATCA crate hosts two COB cards, one supporting four data processing DPMs and the other supporting three data processing DPMs and one trigger DPM for a total capacity of 21 hybrids, one more than required. The external PC supports three network interfaces; two standard 1 Gbit ethernet interfaces used for slow control, monitoring and communication with JLab DAQ and one low latency custom reception card supporting TTL trigger acknowledge interface to the trigger DPM in the ATCA crate. This PC hosts the SVT control and monitoring software as well as the ROC application described above.

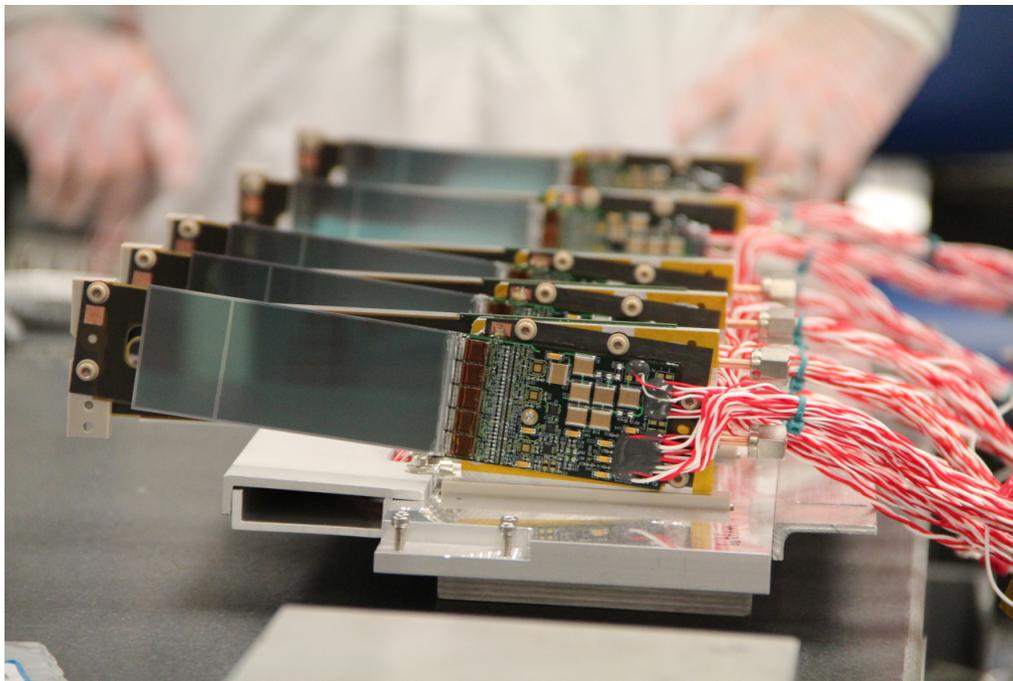


FIG. 42: View from upstream of one half of the SVT modules mounted on the support plate. Signal, power and control are soldered, and potted, to pads at one end while the five APV25 chips are wirebonded to the silicon sensor at the other.

The ECal DAQ system used in the test run is very similar to that described for HPS in Sec. 4.6.2. The only significant difference is that in the test run, the signals from the ECal modules were sent to a signal splitter. One of the outputs of the splitter is fed to a discriminator that also has an internal scaler, and then to a TDC channel. The other output is sent to the JLab FADC250 VXS module, shown in Fig. 30. There are differences between the test run trigger system and the trigger system proposed in Section 4.6.3; these differences are described in further detail in Sec. 5.1.4.

5.1.4 Test Run Trigger System

The trigger system used in the test run is described in the HPS status update to PAC39 [3] in full detail; only the chief differences will be described here. It is generally similar to that described in Section 4.6.3; the same hardware (FADC, CTP, and SSP) was used.

The key differences were in the FADC integration algorithm, the energy resolution reported by the FADC to the CTP, and the SSP trigger decision algorithm.

In the test run, the FADC used a time-over-threshold algorithm to integrate hit energy for

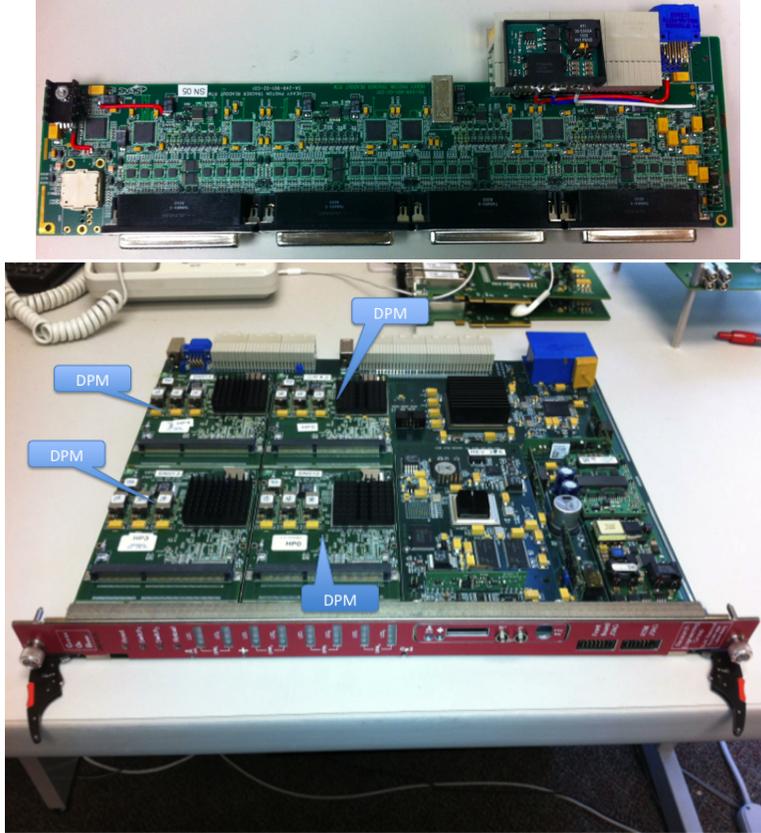


FIG. 43: Picture of a RTM (top) and COB board (bottom) used in the HPS test run 2012.

use in trigger decision, where only samples above threshold were summed. This algorithm is shown in Figure 44.

The FADC reported hits to the CTP in an 8-bit format consisting of a 5-bit channel sum and a 3-bit timestamp. This meant that the pulse integral had to be truncated to fit in 5 bits; consequently some energy resolution was lost. A block diagram of the FADC internal processing and reporting to the CTP is shown in Figure 45.

For the test run, the simplified trigger logic in the SSP was a simple threshold: the trigger was configured to fire on a single cluster with energy exceeding a given threshold. However, the full trigger described in Ref. [2] was implemented and tested.

5.2 Multiple Coulomb Scattering Measurement

Occupancies close to the beam create many of the key challenges in the HPS experiment and determine the limits of sensitivity to low A' masses. These occupancies are dominated by

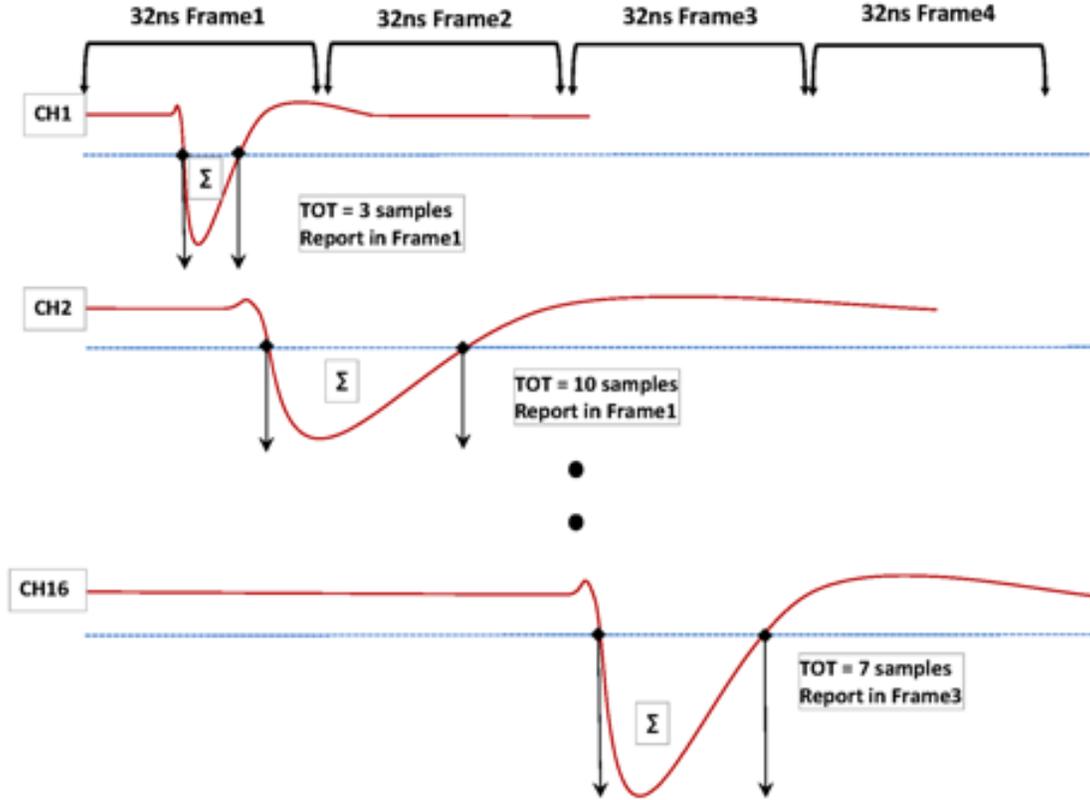


FIG. 44: Example of input signals, and how they are integrated and digitized for the test run trigger.

electrons which have multiple Coulomb scattered to relatively large angles in the converter. Because HPS is sensitive to scattering angles far out on the tail of the multiple Coulomb scattering distribution, well beyond the angles important in other experiments, care must be taken to ensure our simulations are correct in this regime. In particular, Geant4 overestimates the multiple Coulomb scattering rate by a factor of two at large angles as explained in detail in the appendix (see Fig. 88). One of the main goals of the test run in 2012 was to evaluate the description of the tails of the multiple Coulomb scattering in order to gain further confidence in our expected detector occupancy. As will be shown below, data from the test run can be used to confirm our model of multiple Coulomb scattering despite the fact that all data was taken with a photon beam.

Figure 46 gives a schematic view of the main differences between the photon and electron beam setup. In particular, the angular distribution of the pair produced electron and positrons emerging from the converter has comparable contributions from *i*) the pair production

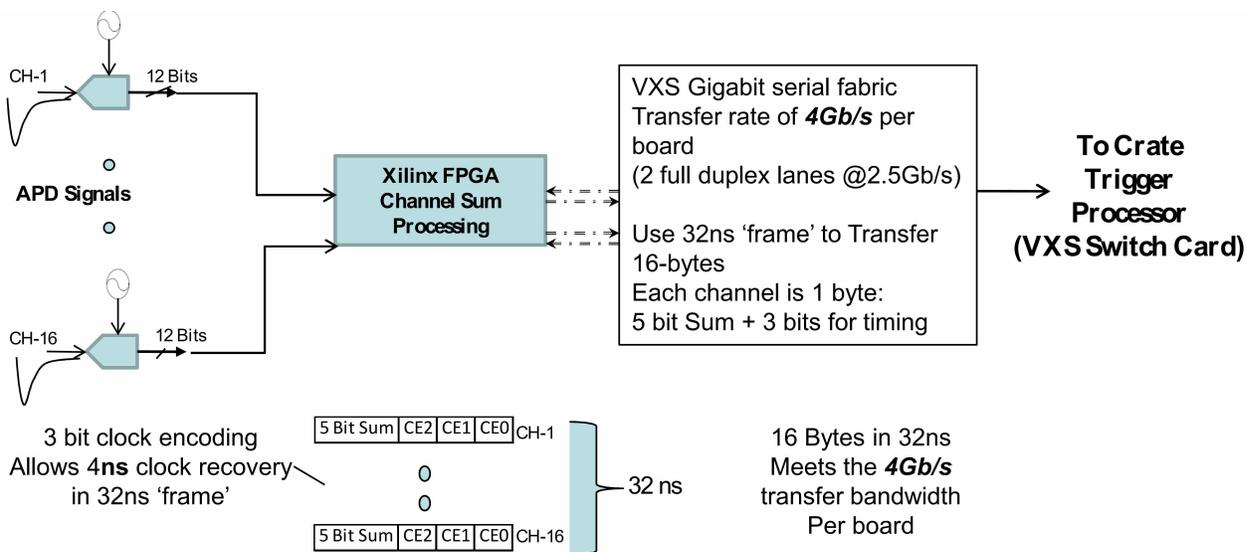


FIG. 45: Block diagram for the trigger system.

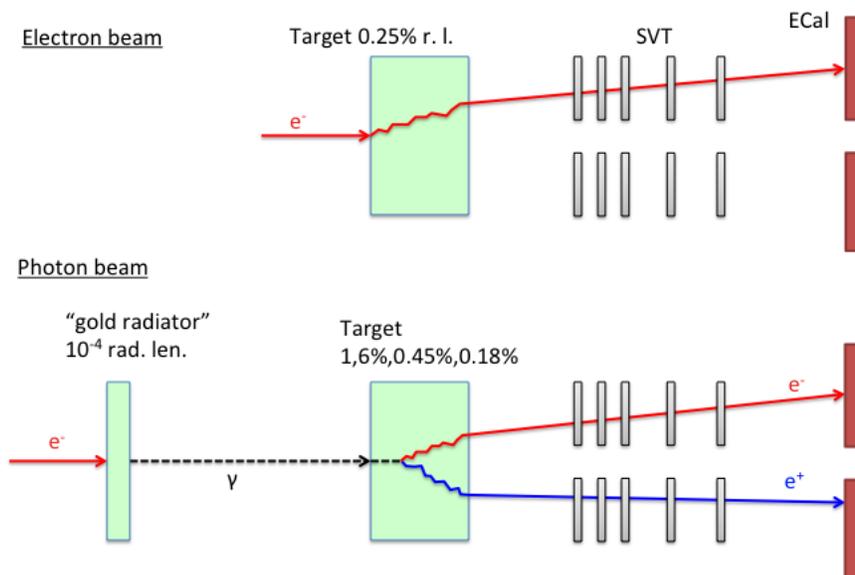


FIG. 46: Schematic comparison of the the setup in the test run photon beam compared to the HPS electron beam.

angle and *ii*) the multiple Coulomb scattering of the electron and position in the converter after production, see Fig. 47.

The measured angular distribution in the ECal for the three converter thicknesses are shown in Fig. 48 (left). The photon beam line during the test run produced a relatively

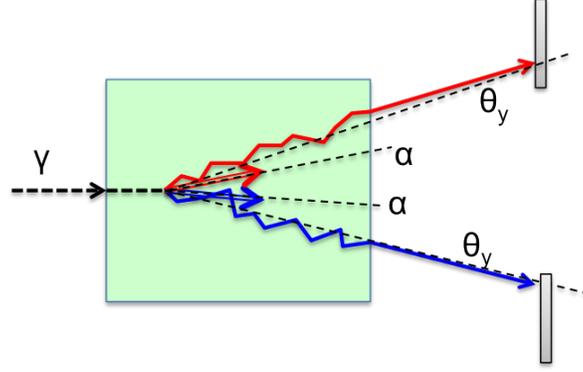


FIG. 47: Schematic description of the relevant angles for pair production in the test run.

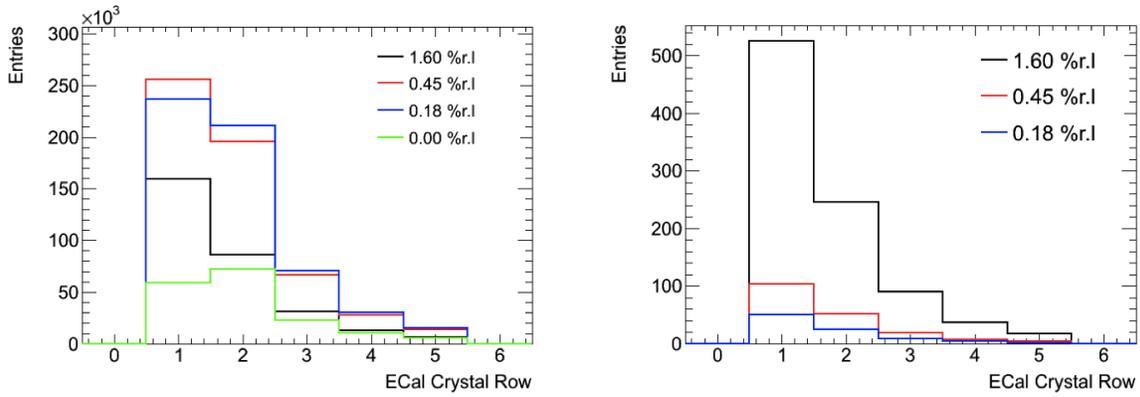


FIG. 48: Measured raw vertical angular distributions before (left) and after (right) normalization and background subtraction.

large fraction of pairs originating upstream of the converter. This contribution was measured during data taking with "empty" converter runs i.e. removing the converter but with all other conditions the same. The upstream background measured in the "empty" converter runs was subtracted from the other runs, properly normalized using the measured integrated currents detailed in Tab. XII. The background fraction for the three converter thicknesses was 16%, 52% and 71% for converter thicknesses of 1.6%, 0.45% and 0.18%, respectively.

Converter thickness (%r.l.)	Duration (s)	e^- on converter (nC)
1.6	911	24385.9
0.18	2640	193508.9
0.45	2149	140709.9
0	1279	88079.6

TABLE XII: Measured integrated currents for the runs used to measure the angular distributions.

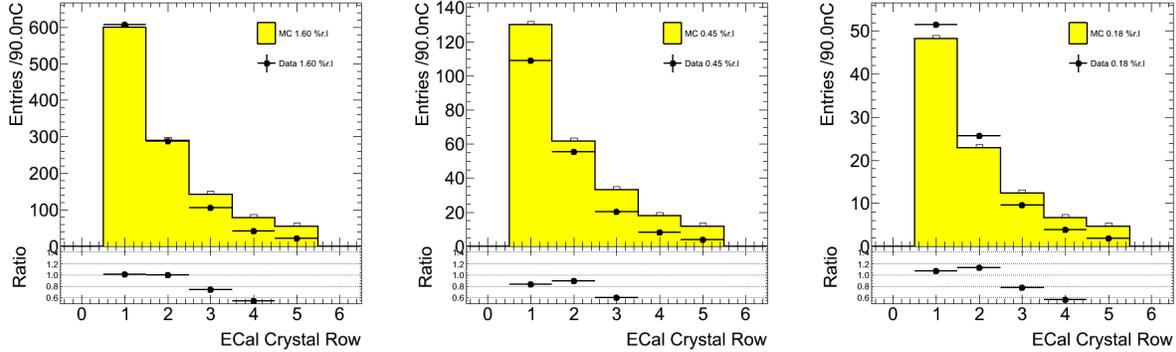


FIG. 49: Comparison between the observed and predicted angular distribution using EGS5 for converter thickness of 1.6% (left), 0.45% (middle) and 0.18% (right). Only statistical uncertainties are included.

Converter (% r.l.)	1.60	0.45	0.18
EGS5	1162 ± 112	255 ± 28	94 ± 17
GEANT4	2633 ± 250	371 ± 38	114 ± 18
Observed	1064 ± 2	196 ± 1	92 ± 1

TABLE XIII: Observed and predicted number of events for 1 s of beam at 90 nA for three different converter thicknesses. The uncertainty on the prediction includes systematic uncertainties.

The background fraction was also cross-checked by pointing back tracks reconstructed in the SVT to identify the fraction of tracks not emanating from the converter. This can be seen in Fig. 56 (bottom) where small satellite peaks at ± 10 mm can be identified as tracks from the upstream background. The angular distribution, after normalization and subtracting the upstream background, are shown in Fig. 48 (right). We also checked that the contribution from photons to our triggered sample was less than 2% (without angular selections which would further reduce the contribution).

These measured angular distributions are compared to simulation to validate the modeling of the multiple Coulomb scattering. As described in more detail in Appendix A, EGS5 is used to generate the electromagnetic interactions in the converter while GEANT4 is used to simulate the particles after the converter. Figure 49 shows the angular distribution comparing data and EGS5 normalized to 1 s of beam at 90 nA beam current. The total rate measurements are in Fig. 50 and summarized in Tab. XIII. The total systematic uncertainty was estimated to be between 10-18% depending on the run including: a 5% uncertainty on the integrated current normalization, alignment of the ECal, uncertainty from the background normalization, and limited Monte Carlo statistics.

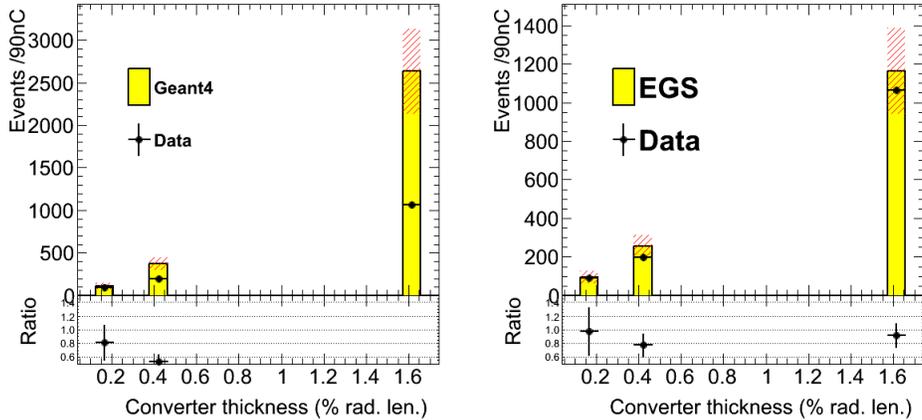


FIG. 50: The measured rate as a function of converter thickness comparing GEANT4 (left) and EGS5 (right).

In summary, the accurate modeling of the multiple Coulomb scattering is fundamental to estimate occupancies and trigger rates for HPS. EGS5 predicts the correct angular distribution across all converter thicknesses while GEANT4 overestimates the rates; with the disagreement increasing with larger converter thickness. This preliminary result verifies our modeling of the multiple Coulomb scattering using EGS5 for HPS.

5.3 Test Run Apparatus Performance

As previously noted, all running of the HPS Test apparatus was with photon beams, using the Hall-B pair spectrometer (PS) pair converter as a target. This target, located ~ 67 cm from the first layer of the tracker, produced a sufficient flux of electrons and positrons to test the principles of running the HPS experiment. This section will report on a few selected preliminary results that demonstrate our understanding of the system.

5.3.1 SVT Performance

During the duration of the dedicated photon run all SVT hybrids and APV25 readout chips were configured to their nominal operating points [98] while all sensors were reverse-bias at 180 V. The sensors were operated within a temperature range of 20 to 24°C throughout the test run. Multiple calibration runs established a noise level of 60-68 ADC counts (≈ 750 - 850 electrons) which was stable across all hybrids. With a linear gain up to ≈ 3 MIPs, the

cluster charge for hits on a track follow the characteristic Landau shape (see Fig. 51) with a mean of about $25,500 e^-$ as expected.

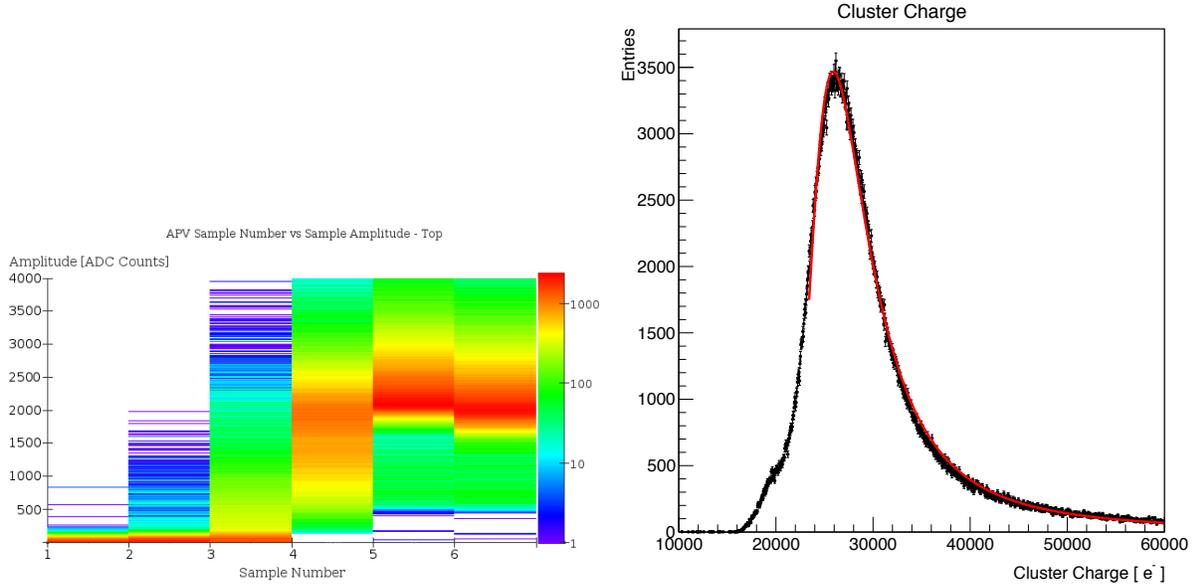


FIG. 51: The six pedestal subtracted samples associated with a hit on a track are shown on the left plot along with a distribution of the cluster charge exhibiting the characteristic Landau shape on the right.

One of the important tests of the SVT was the operation of the APV25 chips in multi-peak readout mode discussed in Sec. 4.3. The six samples of the APV25 pulse shaper output are fitted to an ideal $CR - RC$ function to extract the amplitude and the t_0 of the hit. The typical cluster shape obtained is shown in Figure 51 also demonstrates that the SVT was well timed in to the trigger with the rise of the pulse at the 3rd sampling point. After clustering hits on a sensor, the hit time for each cluster is computed as the amplitude-weighted average of the fitted t_0 channel times. The t_0 -resolution is studied by comparing the cluster hit time with the average of all cluster hit times, the "track time". Figure 52 shows the track time, with the expected jitter due to clock phase and trigger, and the residual to the individual cluster times. After correcting for offsets from each sensor (time-of-flight, clock phase) the extracted t_0 resolution is 2.6 ns (larger than the 2.4 ns due to correlations between the cluster and "track" times). This is somewhat worse than the ≈ 2 ns resolution expected in Section 6 which we attribute to our pulse shape fit procedure function (see Appendix B for more details).

Throughout the duration of the test run, approximately 97% of the 12,780 SVT channels

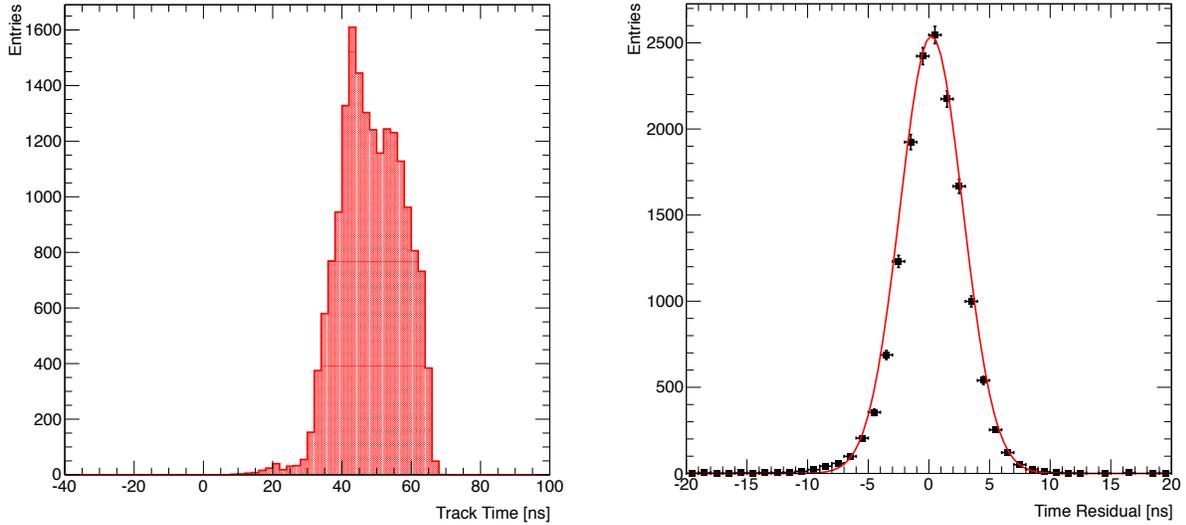


FIG. 52: Track time distribution (left) and cluster time residual (right). The track time is measured relative to the APV25 clock. The width of the distribution is due to trigger jitter (24 ns jitter in the tracker readout clock, plus 16 ns jitter in the trigger system). The cluster time residual is for a representative sensor relative to the track time.

were found to be functional. The identified dead or noisy channels were largely attributed to a known noisy sensor and a few noisy or misconfigured chips which will be resolved for future running. The increased noise resulted in occupancies and data rates that were higher than what were expected from simulation. However, after masking out all known noisy channels found during the commissioning of the SVT, good agreement between simulation and test run data was achieved as shown on Figure 53. Similarly, the hit efficiency was measured to be above 98% for known good layers, see Figure 54. Further details can be found in the Appendix B.

The SVT was aligned using a combination of optical, laser and touch probe surveys at SLAC and JLab. The optical survey of individual modules with precision of a few microns are combined with a touch-probe survey of the overall SVT support structure, with 25-100 microns precision, to locate the silicon sensor layers with respect to the support plates and the mechanical survey balls on the base plate. After full assembly and installation of the SVT at JLab, a mechanical survey of the SVT base plate position inside the pair spectrometer vacuum chamber is used to determine the global position of the SVT with respect to CEBAF beam line. The resulting survey-based alignment has the position of the silicon sensors correct to within a few hundred microns as shown in the reconstructed track

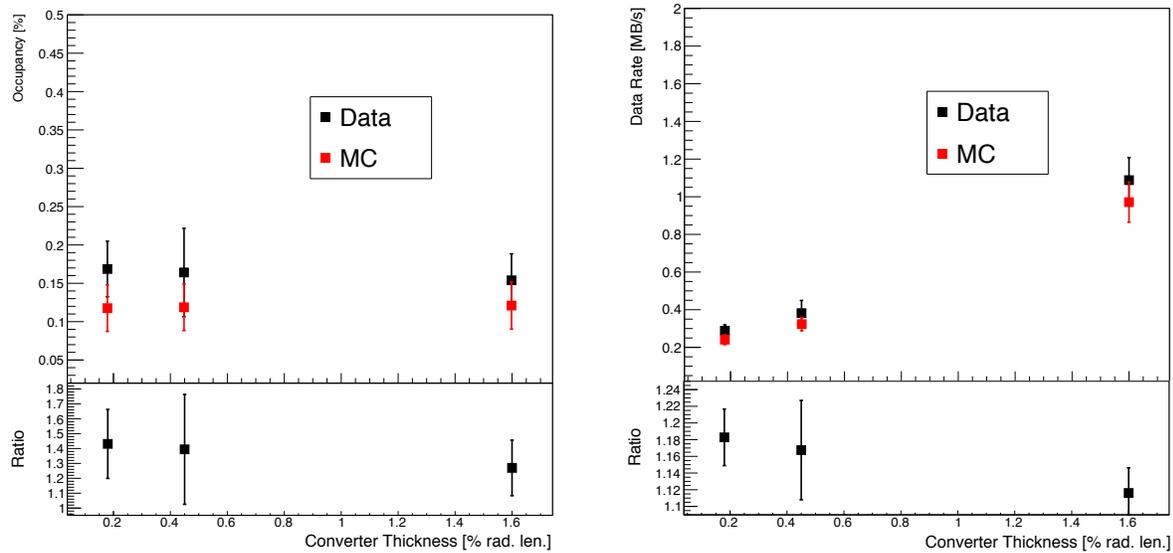


FIG. 53: Comparison of SVT occupancy and data rates between test run data and that expected from simulation after noisy channels are masked.

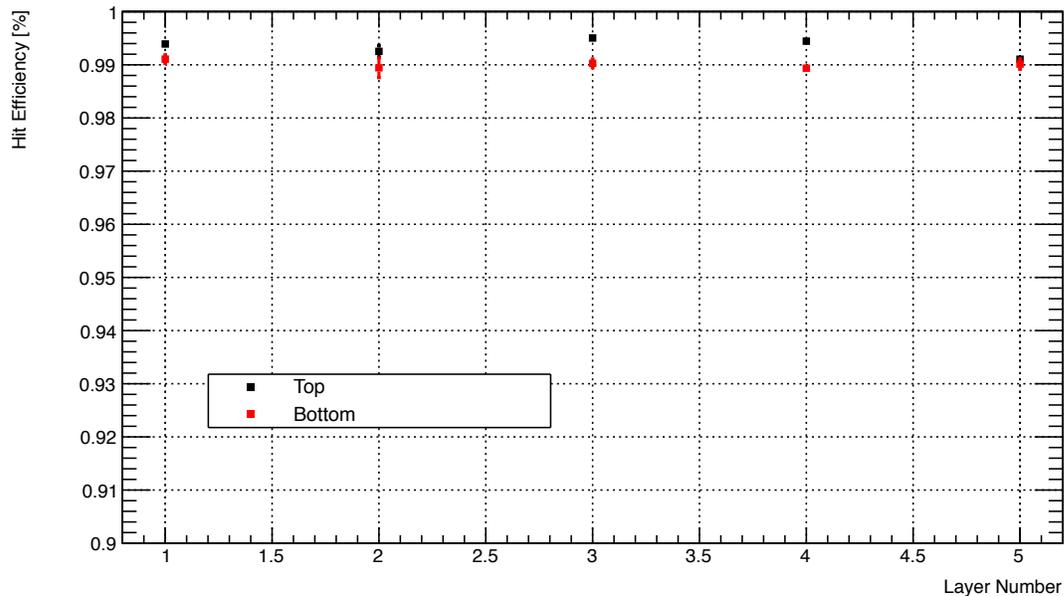


FIG. 54: The hit reconstruction efficiency as a function of detector layer. The variation across the layers can be explained by known DAQ issues.

residuals in Fig. 55.

We also extrapolate the reconstructed tracks back to the converter located ≈ 77 cm from our first silicon layer to understand the tracker alignment w.r.t. to the other components on

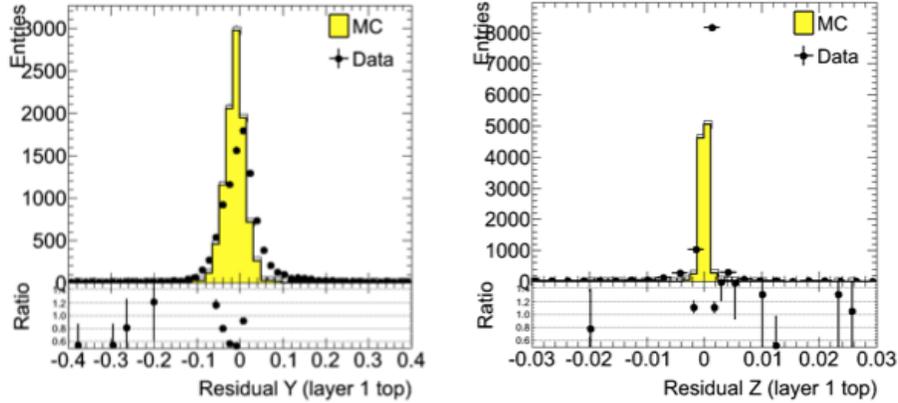


FIG. 55: Residual between the actual hit position and the predicted position in layer 1 from the reconstructed tracks in the bend (left) and non-bend (right) plane in the top half of the SVT after mechanical survey. The filled histogram show the predicted residual for simulated events with an ideal geometry. The X-scale is in mm.

the beam line. Figure 56 shows good agreement of the reconstructed track position at the converter with that predicted from simulation. The offset of the horizontal position simply reflects the fact that the positions are reconstructed in an SVT-centered coordinate system, which is tilted with respect to the beam coordinate system.

With initial residuals less than $\sim 500 \mu\text{m}$ across all layers of the tracker and a reconstructed beam profile similar to that expected from simulation, it appears these survey techniques are adequate to bootstrap the SVT alignment. For HPS, we are developing a more sophisticated global track-based alignment technique to reach the final alignment precision. This framework will also enable us to explore and understand important details such as weak modes and how dedicated alignment runs (e.g. with magnetic field off or with different targets) may shape operational procedures during HPS running.

5.3.2 ECal performance

Of 442 crystals/channels, 39 were disabled or disconnected and were not read out by the DAQ. 13 of these were not read out because of a shortage of FADC readout boards. The remainder either had no HV bias on the APD, or were disabled in the FADC software due to noise. In the data, we identified two types of abnormal channels. One FADC was not sending trigger signals correctly, resulting in low efficiency. This affected the 13 channels read out by that FADC. 5 channels were diagnosed as noisy because they had a high incidence of hits

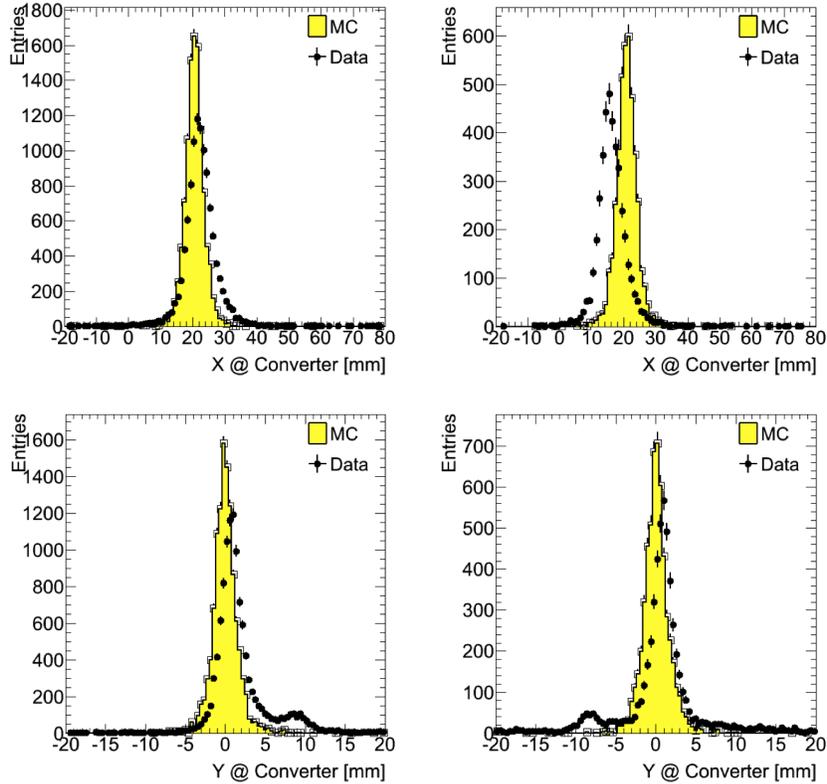


FIG. 56: Extrapolated track position in the bend (top) and non-bend (bottom) direction for tracks reconstructed in the top (left) and bottom (right) half of the SVT. The filled histograms show the prediction from simulation using an ideal geometry. A shift in the bend-plane coordinate for tracks in the bottom half (top right) is likely due to alignment or incomplete description of the magnetic field at the edge of the magnet. The extra bumps in the data at ± 10 mm arise from backgrounds originating upstream of the converter.

out of coincidence with the trigger. A large number of channels were originally misidentified as noisy because they had much higher hit occupancy than neighboring channels. Gain calibration shows that these channels have high gain (and thus lower energy threshold) but are otherwise normal. The abnormal channels were ignored in analysis in order to simplify comparison with Monte Carlo. This leaves 385 useful channels—87% of the ECal.

Each ECal crystal was calibrated for pedestal, noise and gain. This is described in Appendix C.

5.3.3 Trigger performance

As described in Section 5.1.4, the trigger and DAQ integrate pulses differently to measure hit energy. The trigger integrates using a time-over-threshold window, and the DAQ readout

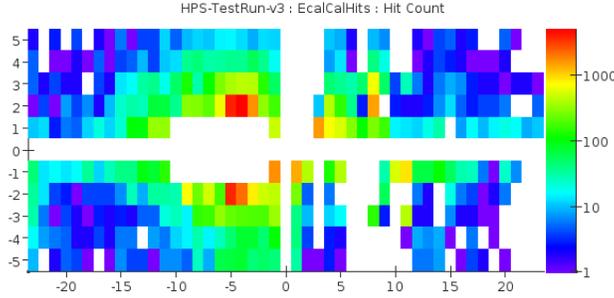


FIG. 57: Hit rates (above an energy threshold, to cut out effect of gain variations) vary smoothly with position.

integrates using a constant window (5 samples before and 30 samples after a threshold crossing). For every event, the trigger reports as a bitmask the trigger decision (top trigger, bottom trigger, or both) and the time the trigger fires.

We study trigger performance by simulating the trigger for each event and comparing to how events were actually triggered. First, we simulate the FADC to convert from readout hits (constant integration window) to trigger hits (time-over-threshold integration). We then simulate the CTP clustering algorithm and the trigger decision (described in Section 5.1.4), and compare the trigger decision and trigger time reported by the simulation to what was reported by the real trigger.

To eliminate trigger bias in checking the trigger decision, we use a tag and probe method: to check trigger performance in one half of the ECal, we tag events where there was a trigger in the other half, and exactly one probe cluster in the ECal half under test. We then measure trigger efficiency (proportion of tagged events where there was a trigger) as a function of ADC counts and energy of the probe cluster. These turn-on curves are shown for the top half of the ECal in Figure 58. The trigger threshold is seen to be 1280 ADC counts as expected. The threshold is not perfectly sharp in this analysis because of uncertainties in the conversion from readout to trigger hits but based on comparisons with Monte Carlo simulation we believe the trigger worked exactly as specified. The trigger threshold in terms of cluster energy is very uneven for two reasons; gain variations between different ECal crystals lead to threshold variations and the nonlinearity of the time-over-threshold integral means that the effective threshold is higher for clusters that span multiple crystals. Overall the trigger appears to have functioned exactly as intended. Changes planned for the next run (constant integration window and per-crystal gain calibration constants for the trigger)

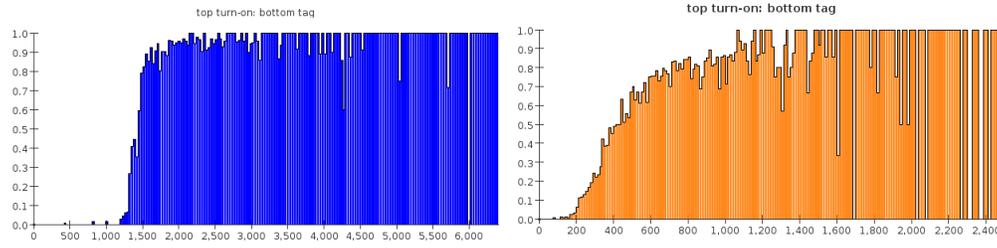


FIG. 58: Trigger turn-on as a function of probe cluster ADC counts (left) and probe cluster energy in MeV (right). Both plots are for the top half of the ECal; bottom is similar. Energy is not corrected for sampling fraction.

will solve both of the issues that led to threshold variations in the test run.

6 HPS Performance Studies

We use the HPS detector simulation system based on SLAC's org.lcsim infrastructure for full GEANT4 simulation of the passage and interaction of charged and neutral particles through the SVT and the ECal to the muon detector. In the SVT, the simulation creates realistic energy deposits in the silicon microstrip detectors, accounts for dead material, simulates APV25 signal sampling every 25 ns, creates clusters, and performs track finding and reconstruction. In the ECal, the geometry for the flange and vacuum chamber is based on a tessellated representation imported directly from the CAD drawings. It creates energy deposits in individual trapezoidal-shape $PbWO_4$ crystals, simulates FADC signal time evolution and sampling every 4 ns, and generates triggers based on the FPGA trigger algorithm implementation. To maintain the chicane beamline configuration, the field strength of the chicane magnets must scale with the beam energy. The performance studies were made using the field strength of the analyzing magnet of 0.25 Tesla at 1.1 GeV, 0.5 Tesla at 2.2 GeV, and 1.5 Tesla at 6.6 GeV. Figure 59 shows a lcsim rendering of the HPS detector.

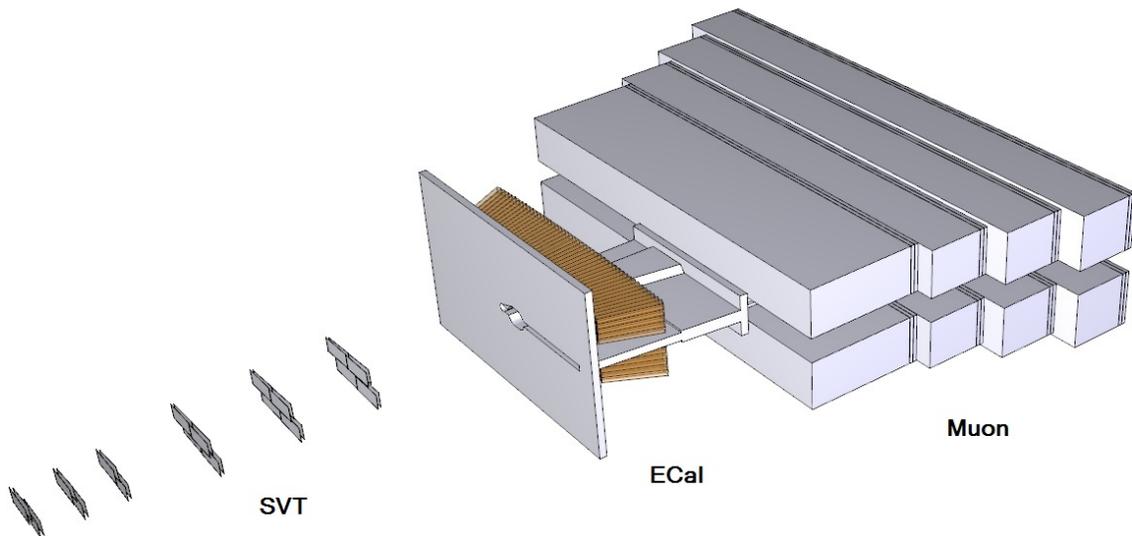


FIG. 59: Rendering of the HPS detector simulation.

6.1 Simulation of Backgrounds and Detector Occupancies

6.1.1 Simulation of Backgrounds

The multiple Coulomb scattering and bremsstrahlung processes in the target will generate high intensity fluxes of electrons and photons in the very forward direction, while the large Møller interaction cross section with atomic electrons will generate high intensity low energy electrons. We use the high energy interaction simulation tools GEANT4 and EGS5 to simulate these backgrounds. In the original HPS proposal to JLab PAC37 [1], we described a significant disagreement between these tools. GEANT4 predicted a broader angular distribution of multiple scattered electrons than EGS5, resulting in twice the occupancy in the tracker near the dead zone and much higher ECal trigger rates. The HPS Test run was motivated in part by the need to resolve this discrepancy, and the outcome of the test run is described in the previous section. The algorithms used in the codes to simulate the multiple scattering have been studied, and the findings are summarized in Appendix A. The test run result and the algorithm studies have confirmed that EGS5 can describe the multiple scattering tails more accurately than GEANT4. All the electromagnetic interactions in the target are simulated with EGS5.

When bound electrons in the target are ionized by incoming electrons or secondary photons, outer shell electrons will fill the vacancy and characteristic X-rays are emitted. These X-rays can contribute background hits in the SVT when a conversion takes place in the silicon sensors via the photoelectric effect or Compton scattering. Since X-ray production from electrons interactions is not fully simulated in EGS5, we estimate the X-ray intensity at SVT Layer 1 using the impact ionization cross section, σ_I , [100], the fluorescence yield, ω , [101], the photoabsorption length in Tungsten, λ_W , to account for the self-absorption, and the solid angle of the SVT Layer 1. Table XIV summarizes these parameters and the expected X-ray fluxes at SVT Layer 1 for 0.25% X_0 Tungsten and 100 nA beam current in 8 ns time window. X-ray backgrounds are not a concern.

Hadrons are also produced in the target. Hadron production is at least three orders of magnitude smaller than the electromagnetic interaction. The polar angle for hadron production is predominantly larger than 100 mrad, whereas the HPS detector acceptance is limited to less than 100 mrad. Furthermore, the hadron energy spectrum is soft as they are produced

	Energy (keV)	σ_I (barns)	ω	λ_W (μm)	N_γ at Layer 1 in 8 ns
K-shell	60	40	0.95	100	0.5
L-shell	10	1000	0.30	5	2
M-shell	2	20000	0.02	0.2	0.1

TABLE XIV: X-ray intensities.

from the $1/k$ bremsstrahlung spectrum and more than 90% of the hadrons are swept away by the analysing magnet before reaching the ECal. The hadron production is simulated using GEANT4 and FLUKA. In a target thinner than about 5% X_0 , the “virtual” photon interaction is dominant [102]. The inclusive hadron production $\sigma(eA \rightarrow X)$ is simulated from the photonuclear process $\sigma(\gamma A \rightarrow X)$ using the equivalent photon approximation,

$$\sigma(eA \rightarrow X) = \int \sigma_k(\gamma A \rightarrow X) dn(k),$$

where $dn(k)$ is the number of equivalent photons with energy k [103] and there are approximately 8×10^{10} photons/sec in 6.6 GeV 100 nA beam. Table XV summarizes the pion single rates from 1% X_0 Tungsten target and 6.6 GeV 100 nA beam. While pion production is larger in GEANT4, the energy spectrum is softer and consequently the single rate of pions reaching the ECal is lower in GEANT4. While pions look like a minimum ionizing particle in the ECal most of the time, they can deposit significant energy when π^0 are produced in the ECal crystals, and together with the beam background, they contribute accidental coincident triggers.

	Total production rate (kHz)	Single rate reaching the ECal (kHz)
GEANT4	410	8
FLUKA	240	15

TABLE XV: Pion single rates from 1% X_0 Tungsten target at 6.6 GeV 100 nA.

Other beam induced background we considered are:

- Beam halo

Beam halo was measured using a large dynamic range halo monitor during the 6 GeV era. The beam halo that extends to 2 mm was found at the level of 10^{-7} . At this level, the halo contribution in the SVT occupancy is negligible. It is expected that the behavior of the 12 GeV machine will be understood at the same level.

- Synchrotron radiation

Synchrotron radiation is produced from the last dipole magnet in the beam line in the vertical plane, and from the chicane magnets in the horizontal plane. Since the characteristic energy is proportional to E_{beam}^2 , synchrotron radiation is of concern only at 6.6 GeV. The characteristic energy (k_c), the average energy (k_{ave}), and the power of the synchrotron radiation is summarized in Table XVI. None of the radiation from the last dipole will enter the HPS detector as the radiation will be intersected by the beamline collimator. The radiation from the chicane magnets is in the dead zone, and none of the detector components are designed to intersect the beam plane.

Source	k_c (keV)	k_{ave} (keV)	N_γ per e-	Power (mW) at 100 nA
Vertical bend	19	5.9	4.0	2.4
Frascati Magnet	52	16	4.6	7.4
PS magnet	44	14	9.3	13

TABLE XVI: Synchrotron radiations at 6.6 GeV.

- EM induced backgrounds

Electromagnetic fields induced by the high intensity beam could interfere with the SVT and its electronics as the detector is located as close as 0.5 mm from the beam. We have evaluated the direct beam field and its wake field, the diffraction radiations from the beamline apertures, and the transition radiations from the target. The intensities of these EM induced backgrounds are small and no interference with the SVT is expected. The beam charge per 2ns CEBAF bunch is only few thousand electrons, many orders of magnitude lower than that in other experiments which have chosen to shield against it.

6.1.2 Simulated Tracker Occupancies

Figure 60 shows the distribution of charged particle hits in the SVT Layer 1 which is located 10 cm from the target. The beam energy is 6.6 GeV, and the target thickness is 0.25% X_0 . Multiple Coulomb scattered beam electrons are confined to within 0.5 cm of the beam axis ($x=y=0$), while the low energy Møller electrons are distributed in a parabolic shape. There are very few positrons. From these distributions, the detector occupancy in the horizontal silicon strip in the 8 ns time window is calculated for a 400 nA beam current and five different target thicknesses, 1.0% X_0 , 0.5% X_0 , 0.25% X_0 , 0.1% X_0 , and 0.05% X_0 , as shown in Fig. 61. For a 0.25% X_0 target and 430 nA beam, the occupancy is 1% at a distance of 1.5mm from the beam in Layer 1, which corresponds to a dead zone of ± 15 mrad. As long as the product of target thickness (T) and beam current (I) is constant, the same A' production rate is maintained. Since multiple scattering and hence the effective beam size is reduced in a thinner target, it is advantageous to use a thinner target and a higher current. Using the constraint that the occupancy is 1% at 15 mrad, we find the beam current I which gives this occupancy for each of several potential target thicknesses T . The quantity $(I \cdot T)^{1/2}$, which is approximately proportional to the sensitivity S/\sqrt{B} , is given in Table XVII, showing how the sensitivity improves as the target thickness decreases.

Target thickness (% X_0)	Beam Current (nA)	$\propto S/\sqrt{B}$
1.0	60	7.7
0.5	170	9.1
0.25	430	10.4
0.10	1330	11.6
0.05	2860	11.9

TABLE XVII: Beam current yielding 1% occupancy in SVT layer 1 for various target thicknesses at 6.6 GeV, and the relative experimental sensitivities which result.

The run conditions for other possible beam energies are studied using the same criterion that the maximum occupancy in SVT Layer 1 does not exceed 1%. Table XVIII summarizes the target thickness and proposed beam current.

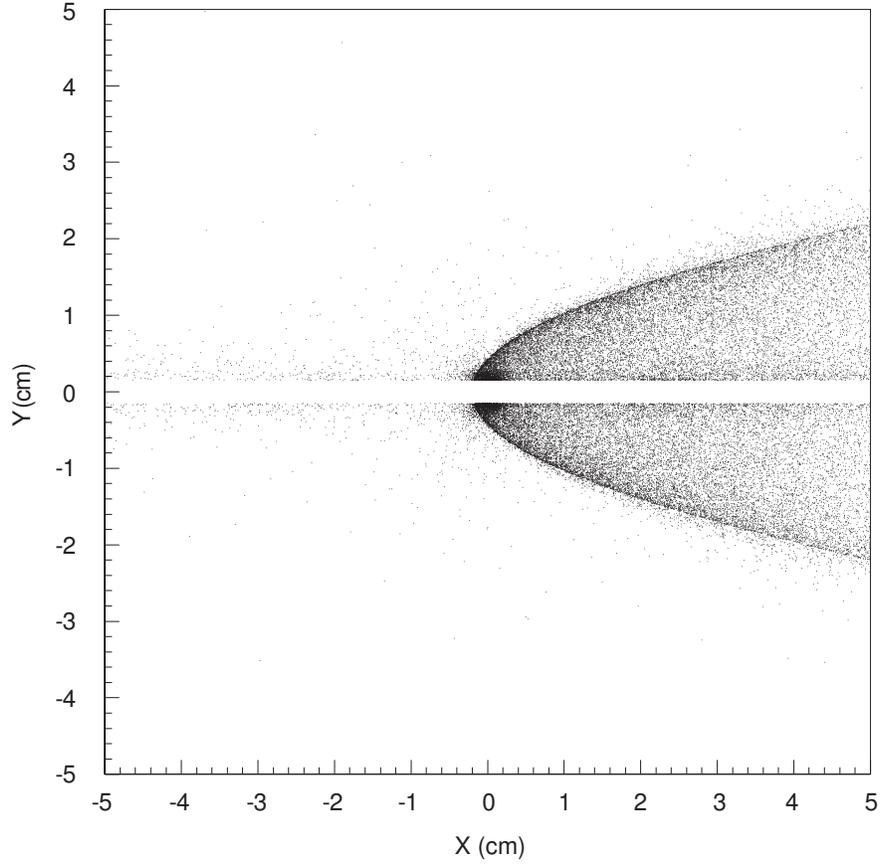


FIG. 60: Charged particle distribution in SVT Layer 1.

Beam Energy	Target thickness (% X_0)	Beam Current (nA)
1.1	0.125	50
2.2	0.125	200
4.4	0.25	300
6.6	0.25	450

TABLE XVIII: Run conditions.

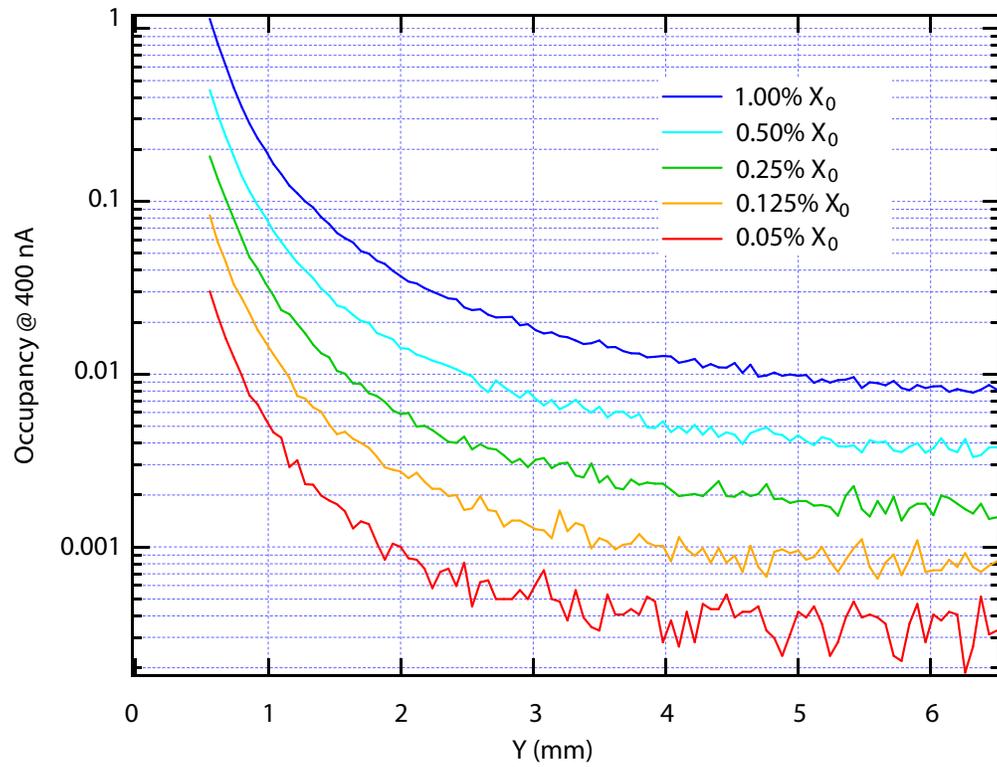


FIG. 61: Layer 1 silicon strip occupancy at 400 nA vs. distance from the beam in mm.

6.1.3 Simulated ECal Occupancies

There are two factors limiting the allowable ECal occupancy. First, the ECal readout algorithm uses a window of fixed size to integrate hit energy. This window was set to 140 ns (35×4 ns) for the test run, and so the number of hits above readout threshold in a 140-ns time window should be well below 1. Figure 62 shows that the maximum rate in any crystal is 500 kHz, which translates to 0.07 hits in 140 ns. Second, because the FADC only reads out on a rising threshold crossing, each hit above threshold causes dead time for that crystal until the pre-amplifier output falls back below threshold. Figure 63 shows the fraction of time each crystal spends above threshold. The maximum dead time is 0.03, meaning that even the hottest crystal is sensitive to new hits 97% of the time.

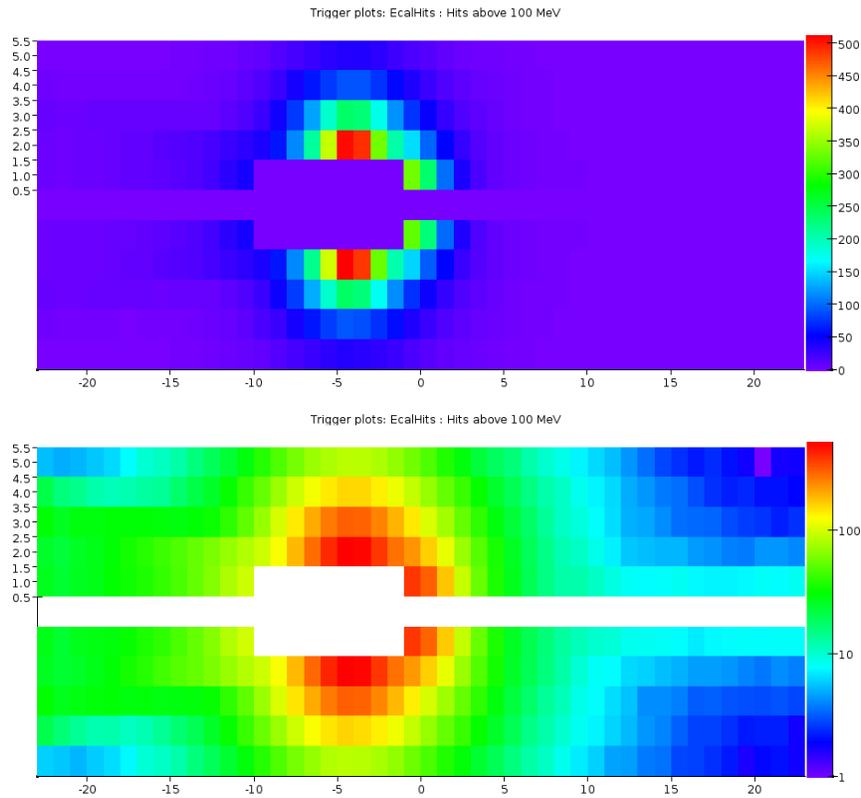


FIG. 62: Rate of hits over 100 MeV (units of kHz) per crystal (X and Y axes are the crystal index), for 2.2 GeV beam at 200 nA. Top plot uses linear scale for the Z-axis; bottom plot is log scale.

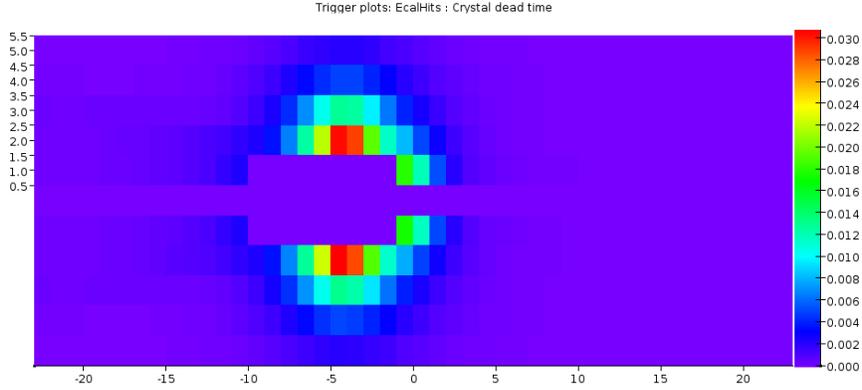


FIG. 63: ECal readout deadtime fraction for 2.2 GeV beam at 200 nA, with a threshold of 75 MeV for each crystal.

6.2 ECal Trigger Rates

The proposed ECal trigger was simulated to test trigger selections, verify that the trigger has acceptable efficiency for A' events, and verify that the trigger rate is compatible with the HPS DAQ in all running conditions.

The CEBAF beam bunch structure was simulated by sending one bunch equivalent of electrons, 625 (1.1 GeV), 2,500 (2.2 GeV) and 5,625 electrons (6.6 GeV), through the target. A total of 50 million bunches of beam background (equivalent to 100 ms of beam) were generated at each beam energy. The details of the target interactions are given in Sec. 6.1.1. Since the trident production process could not be done with EGS5, trident events were generated with MadGraph/MadEvent and overlaid on the beam background bunches with average rate expected from the trident cross section. For the trigger acceptance studies, A' events were generated with MadGraph/MadEvent at beam energies of 1.1, 2.2, and 6.6 GeV.

The complete chain of signal evolution in the ECal crystals and signal processing through the trigger system was simulated by following closely the ECal trigger description in Sec. 4.6.3. Starting from the energy deposits in the ECal crystals, signals were generated using the CR-RC shaper function with a time constant of 15 ns measured with the ECal crystals, amplitudes were sampled and pulse data evaluated every 4 ns (simulating FADC), and the cluster finding algorithm and trigger logic were applied (simulating CTP and SSP). The simulation has been tested against the actual performance of the test run detector and DAQ: see Sec. 5.3.3.

The trigger parameters described in Sec. 4.6.3 are chosen by running the simulation and plotting the relevant variables for beam background and A' events. This is done for each beam energy and a set of A' masses for each beam energy. Figure 64 shows the coplanarity angle vs. the azimuthal angle of the lower-energy cluster, indicating that A' events tend to have small coplanarity angles. Figure 65 shows the distance from beam axis vs. energy of the lower-energy cluster, indicating that the energy-distance cut can reduce the beam background effectively. Figure 66 shows the cluster energy difference vs. energy sum, indicating that the energy sum cut can retain A' events effectively.

These cuts are chosen to lie between the loosest reasonable values (accept as many A' events as possible) and the tightest (reject as many background events as possible). In some cases this leads to different cut values at different beam energies—for example, the coplanarity cut is looser at 1.1 GeV because the background events are clustered at large deviations from coplanarity and a relatively loose cut rejects most of them, but the cut is tighter at higher beam energies.

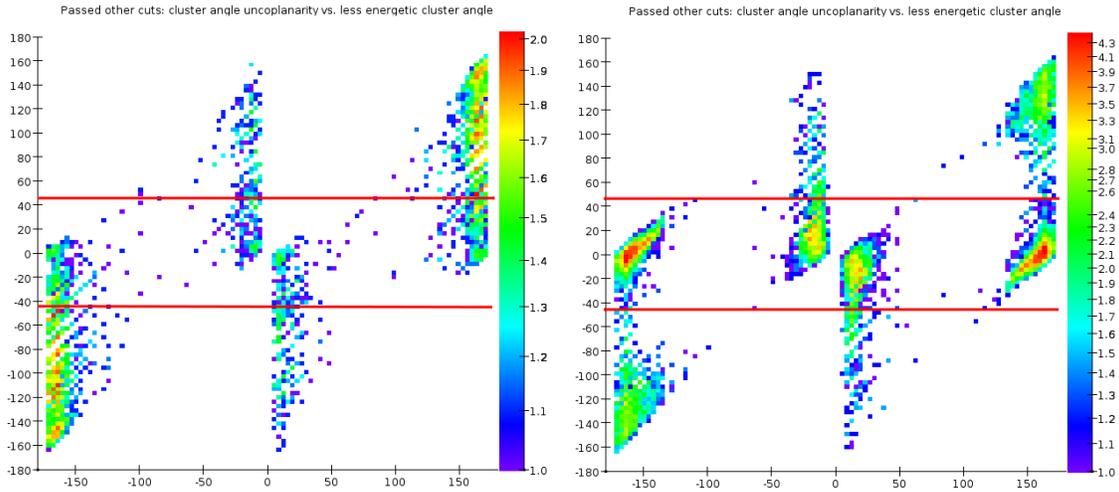


FIG. 64: Deviation of cluster pairs from coplanarity (units of degrees) for 2.2 GeV beam; background and 75 MeV A' trident events are shown. The X-axis is the azimuth around the beam axis (ϕ_1) of the lower-energy cluster, such that 0 degrees is the positron side of the detector and 180 degrees is the electron side; the Y-axis is the difference between the azimuth angles ($\phi_1 - \phi_2 - 180$) of the two clusters. The coplanarity acceptance region is the space between the red lines.

The following trigger parameters were determined to be independent of beam energy:

- Minimum cluster energy (E_{min}): 0.1 GeV
- Distance (r_{edist}) in the energy-distance cut: 200 mm

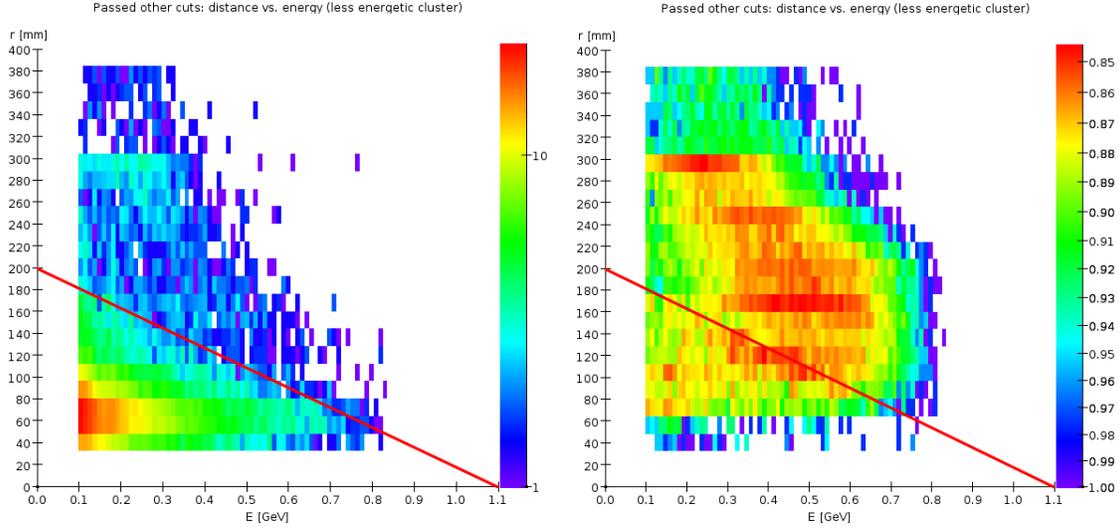


FIG. 65: Energy and distance from beam axis of the lower-energy cluster, for 2.2 GeV beam; background and 75 MeV A' tridents are shown. The energy-distance acceptance region is above the red line.

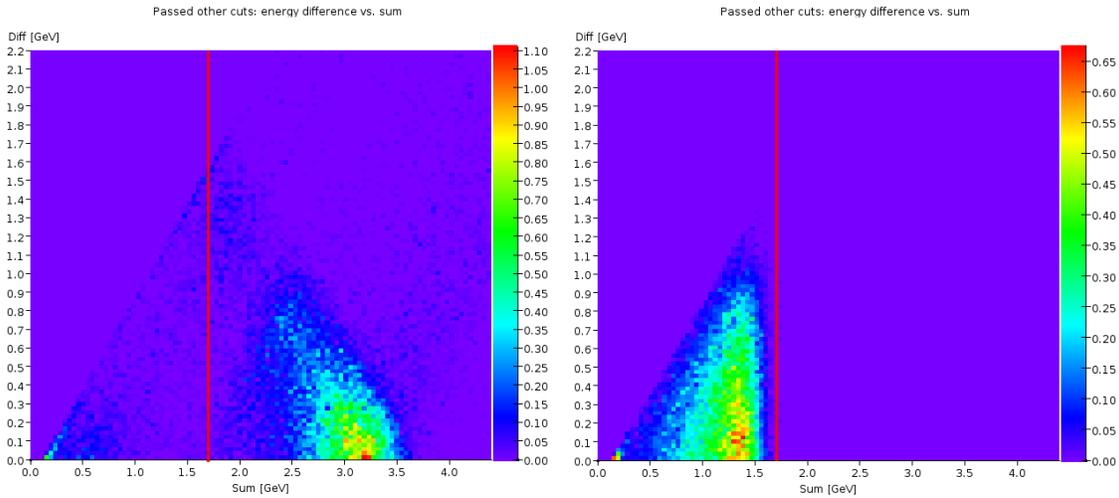


FIG. 66: Energy sum and difference of cluster pairs, for 2.2 GeV beam; background and 75 MeV A' tridents are shown. The energy difference acceptance region is left of the red line.

- Energy (E_{dist}) in the energy-distance cut: $0.5 \times E_{beam}$

Table XIX summarizes the trigger parameters that were found dependent on the beam energy. The remaining trigger parameters given in Sec. 4.6.3 did not have a significant effect on specificity of the trigger—neither the background trigger rate nor the trigger efficiency is sensitive to those parameters..

Trigger rates are shown in Table XX. These rates are safely under the maximum readout

Beam energy [GeV]	E_{max} [GeV]	$E_{sum_{max}}$ [GeV]	$\Delta\phi_{max}$ [°]
1.1	0.7	0.8	90
2.2	1.6	1.7	45
6.6	5.0	5.5	60

TABLE XIX: Trigger parameters optimized for different beam energies.

Sample	Rate (kHz)
1.1 GeV beam background	15.7 ± 0.4
1.1 GeV beam background+tridents	18.3 ± 0.4
2.2 GeV beam background	11.2 ± 0.3
2.2 GeV beam background+tridents	15.8 ± 0.4
6.6 GeV beam background	10.2 ± 0.3
6.6 GeV beam background+tridents	12.6 ± 0.4
6.6 GeV beam background+tridents+pions (FLUKA)	13.4 ± 0.4
6.6 GeV beam background+tridents+pions (G4)	13.5 ± 0.4

TABLE XX: Trigger rates using various background samples, with statistical uncertainties.

rate of 43 kHz set by the SVT DAQ. Furthermore, tightening the coplanarity and energy-distance cuts lowers trigger rates to ≈ 10 kHz at 1.1 and 2.2 GeV and ≈ 5 kHz at 6.6 GeV, while reducing the A' efficiency by no more than 2 percentage points; this provides further safety margin in case trigger or data rates are higher than expected. The addition of pions to the 6.6 GeV background sample has only a small effect on the trigger rate.

Trigger efficiency for A' events is defined as the fraction of A' tridents (generated without fiducial cuts) that produce a trigger.

To evaluate the performance of the experiment, we are interested in the combined efficiency of the trigger and tracker: the fraction of A' tridents that produce a trigger and leave enough hits in the tracker for a pair of tracks to be reconstructed. We simulate charge deposition and readout of the tracker (turning off the generation of noise hits), and check each sensor for hits. If the DAQ reads out hits in four stereo pairs in each half of the tracker, the event is considered to be in the combined acceptance.

Figure 67 shows the ECal trigger efficiency and the ECal/SVT-combined efficiencies for A' events at 1.1, 2.2, and 6.6 GeV. Both trigger and tracker acceptances are dominated by the geometric acceptances of the ECal and tracker. A' prompt decays are assumed.

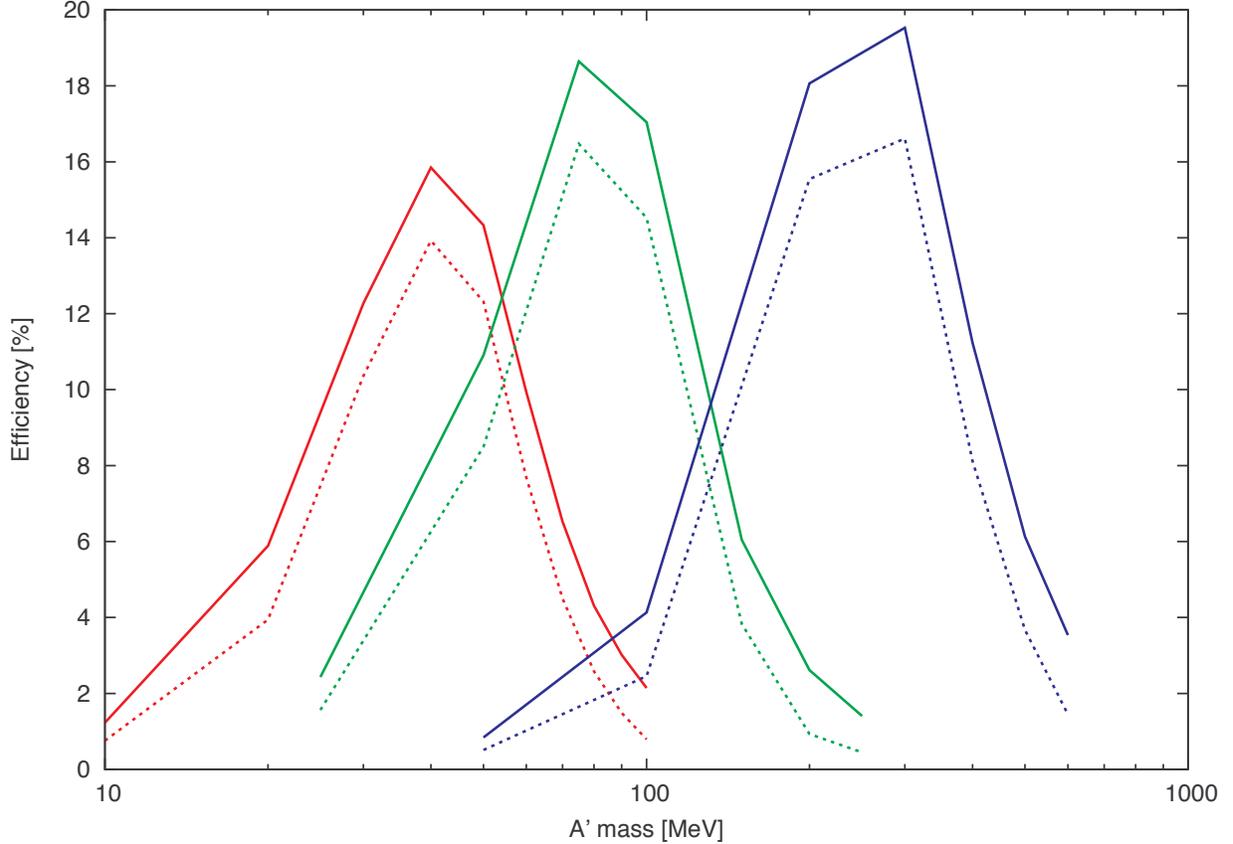


FIG. 67: Trigger efficiency (solid lines) and combined efficiency (dashed lines) as a function of A' mass, at beam energies of 1.1, 2.2 and 6.6 GeV (red, green and blue respectively).

6.3 Muon system trigger rates

Like the ECal, the muon system trigger rates are dominated by beam backgrounds. A GEANT4 model of the HPS detector was used to estimate the rates, following the conceptual design for the muon system presented in Section 4.5. Figure 25 shows the layout of the system. Each of eight hodoscope layers (four layers in the top part of the detector and four in the bottom) consists of two planes of scintillator strips. One plane, called the Y-plane, has strips oriented horizontally. The other, called the X-plane, has them oriented vertically. The Y-strips are segmented exactly in the middle and outside ends, left or right, are read out. Six Y-strips make up each half plane, so the total number of Y-strips is 48. X-planes are divided into 4.5 cm wide segments, with 240 strips in total. It should be noted that the number of vertical strips in the conceptual design is only 208 to limit the total number of readout channels to 256, one crate's worth. Since the rates in the vertical strips at the

edges of the hodoscope are very low, eight strips in each plane can be paired to make four readout channels without negatively impacting the detector occupancy. In Table XXI, the lengths and widths of the detector readout segments (counters) used in the simulation are presented. There is a 14 cm wide and 3.5 cm high gap introduced into the model, centered on the point where the electron beam passes through the muon vacuum chamber, to avoid this high rate region. As shown in Figure 24 this gap has a negligible effect on the detection efficiency of muon pairs from A' decays.

TABLE XXI: Lengths and widths of the hodoscope strips. Dimensions are centimeters.

Readout plane	Layer 1	Layer 2	Layer 3	Layer 4
X-plane width	4.5	4.5	4.5	4.5
X-plane length	10.5	11.5	12.5	13.5
Y-plane width	3.5 all three	3.5, 4, 4	3.5, 4.5, 4.5	4.5 all three
Y-plane length	56	62.5	70	76

Events generated in EGS5 were used as an input to the GEANT4 simulation. The CEBAF beam bunch structure was simulated by sending one bunch equivalent of electrons, 5,625 e^- 's (6.6 GeV), through the target to generate secondaries and scattered beam particles. The secondaries were followed through the apparatus to simulate the detector response. As expected, the highest background rates are seen in the Layer 1 hodoscope and are ~ 0.7 MHz in both the X-strips near the electron beam location and the beam-left Y-strip closest to the beam plane, see Figure 68. Rates in the vertical strips far from the beam position are very low, allowing multiple strips to be combined into a single readout channel in order to reduce the number of PMTs and electronic channels.

The coincidence rates between hodoscope planes in a given layer and between different layers have been studied using a 16 ns coincidence time window. On the left of Figure 69, the coincidence rates between X- and Y-quadrants, top-left (TL), top-right (TR), bottom-left (BL), and bottom-right (BR) of the Layer 1 hodoscope are shown. On the right, the figure shows coincidence rates of respective quadrants of Layers 1 and 2. The fact that there is a significant reduction of the rates from 2-plane (1.2 MHz) to 2-layers (0.07MHz) coincidences indicates that hits are mostly from uncorrelated background. For the muon trigger, a coincidence of two opposite quadrants (TLxBR) or (TRxBL) is required along with triple coincidences of the first three layers of hodoscopes in each quadrant. The rates

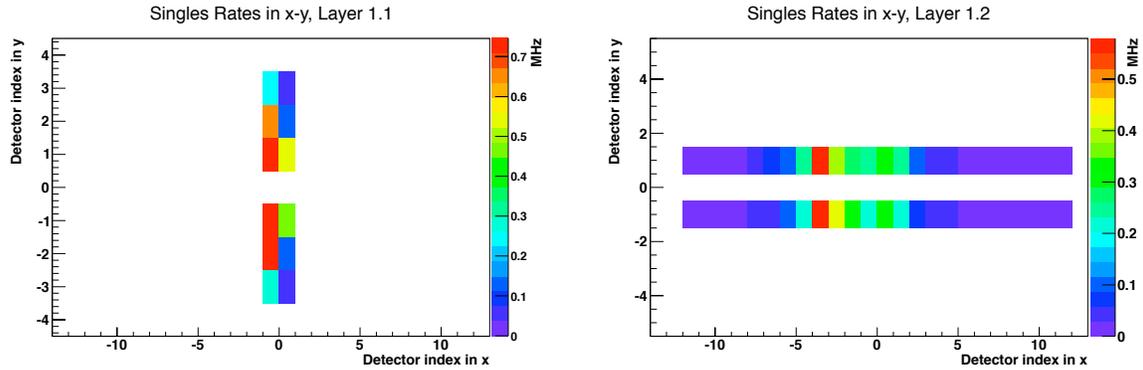


FIG. 68: Rates in Y- (left) and X-planes (right) of the Layer 1 hodoscopes for hits with energy deposition > 0.5 MeV.

of the triple quadrant coincidences within 16 ns are shown in Figure 70. The maximum trigger rates are in the beam-right (electron side) quadrants and are on order of 7 kHz. In the beam-left quadrants (positron side), the tripple coincidence rates are < 1 kHz. Since an overall trigger requires hits in two opposite quadrants, the maximum rate will be < 1 kHz. While a further reduction of rates will be possible with the inclusion of MIP hits in the ECal, the < 1 kHz is a small addition to the total trigger rate from the Ecal trigger (see discussions in Section 6.2) and will keep overall trigger rate well within the limit of allowed rates for the HPS DAQ.

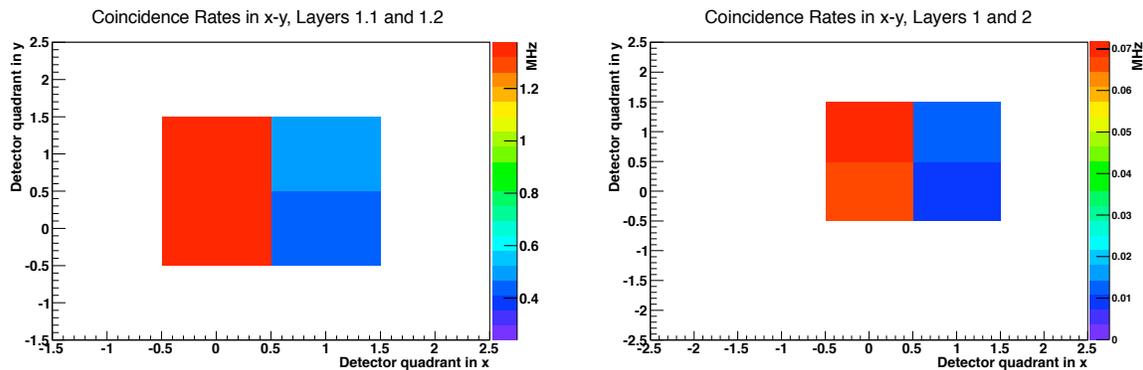


FIG. 69: Coincidence rates between X- and Y-quadrants of the Layer 1 hodoscope (left graph) and coincidence rates between Layer 1 and 2 (right graph).

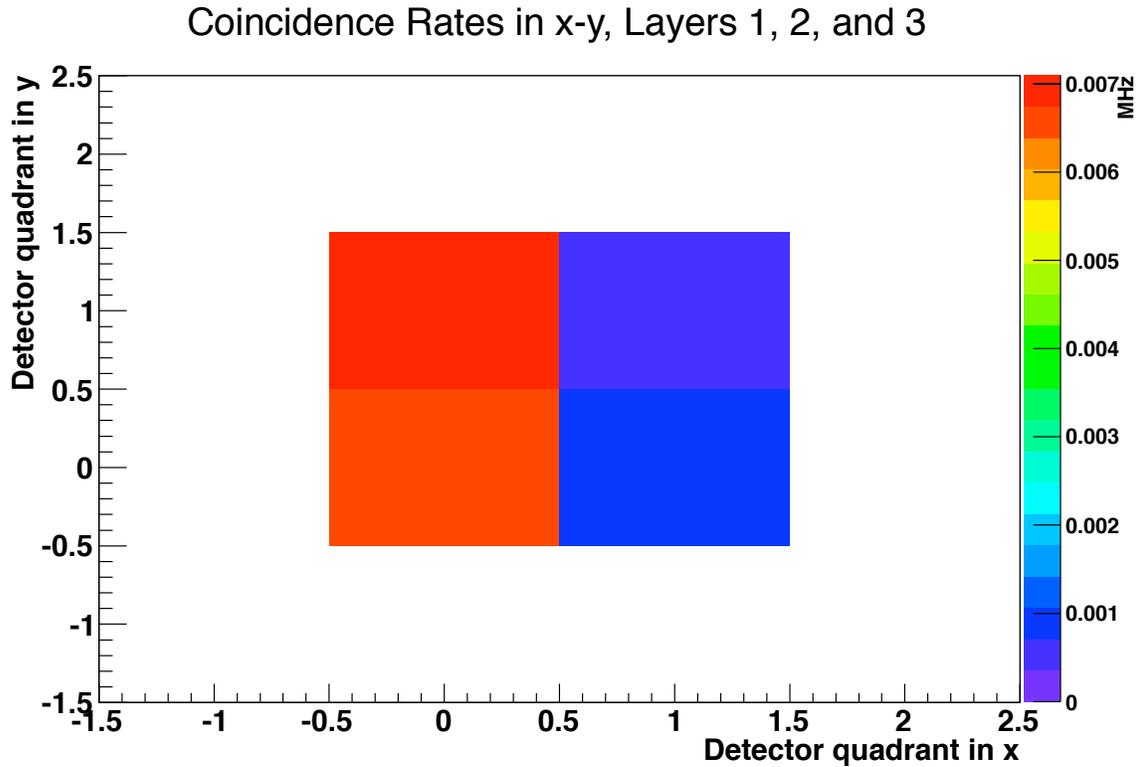


FIG. 70: Coincidence rates in first three layers of muon hodoscopes. The coincidence time window is set to 16 ns.

6.4 Track Reconstruction

In order to study the tracking performance of the detector, we use samples of A' events at a variety of energies and decay lengths. On top of each event, we overlay backgrounds produced by the passage of beam electrons equivalent to our optimized run conditions at different beam energies and with a W target and a beamspot with a Gaussian sigma of $40\mu\text{m}$ in the vertical direction and $200\mu\text{m}$ in the horizontal. The beam energies, currents, target thickness and analyzing magnetic field used for these simulations are:

- 50nA at 1.1 GeV with $X_0 = 0.125\%$ and $B=0.25$ T
- 200nA at 2.2 GeV with $X_0 = 0.125\%$ and $B=0.5$ T

- 300nA at 4.4 GeV with $X_0 = 0.25\%$ and $B=1.0$ T
- 450nA at 6.6 GeV with $X_0 = 0.25\%$ and $B=1.5$ T

At each energy, we evaluate momentum, invariant mass, and vertex resolution. The plots shown in the following section typically use the 4.4 GeV beam as an example.

6.5 Tracking Efficiency, Pattern Recognition and Fake Rates

Due to the requirements imposed on the tracks, the efficiency for finding tracks in the geometric acceptance is not 1. The average track reconstruction efficiency is 98% (Fig. 71) and the bulk of the inefficiency comes from the cut on the total χ^2 of the track. Of the reconstructed tracks, a small percentage include a hit that is not from the correct electron. These “bad” hits may be from one of the high energy beam electrons scattered from the target into the detector or from a lower energy secondary. The left plot of Fig. 72 shows the number of bad hits/track for both the electron and positron from the A' decay. The number of tracks with 0 bad hits is $> 98\%$. The right plot of Fig. 72 shows the layer number of the bad hit. The rate of mishits are slightly higher in the downstream 3 modules due to the larger stereo angle We'll show how these bad hits affect the track parameters in the next section.

6.6 Track Momentum and Spatial Resolution

The momentum resolution is shown in Fig. 73 as a function of momentum for tracks with 0 bad hits and for tracks with one or more. The momentum resolution for well-reconstructed tracks is $\delta p/p = 3.5\%$ for $B=1.0$ T (appropriate for a beam energy of 4.4 GeV) and is roughly inversely proportional to B .

One quantity we use to determine track quality is the distance of closest approach (DOCA) to the beam axis. We use this instead of the DOCA to the target beam spot since we are interested in long-lived decays and tracks from those will not point back to the target. We separate the distance into the bend plane (XOCA) and non-bend plane (YOCA) distances. Below, in Fig. 74, is the resolution of these quantities as function of momentum. The resolution is, on average, about $100\mu\text{m}$ ($300\mu\text{m}$) in the non-bend (bend) direction but

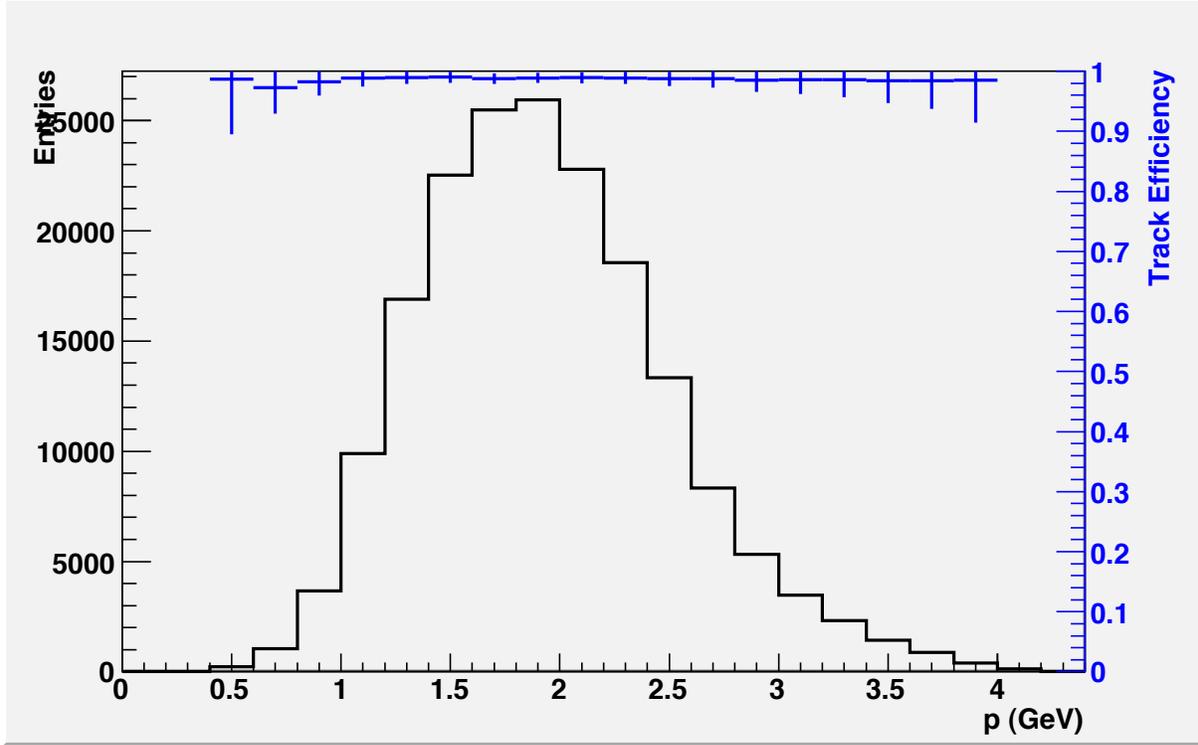


FIG. 71: Track reconstruction efficiency versus track momentum (right axis). The black histogram (left axis) show the track momentum distribution.

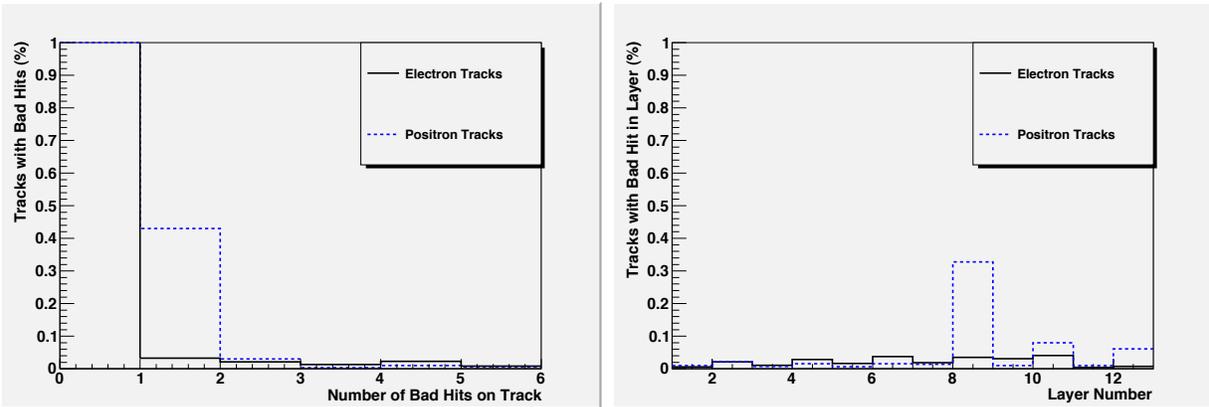


FIG. 72: The number of bad hits (left) and the layer number of the bad hit (right) for electron (black) and positron (blue) tracks.

increases significantly at low momentum. The position resolution for tracks with one or more bad hits is somewhat worse, depending on which layer the bad hit is. Tracks with bad hits in layers 1 or 2 are a major contribution to the tail of the vertex position distribution.

For long lived A' decays, the position of the decay vertex is an important discriminating variable. The dominant background to A' production is radiative events which originate in

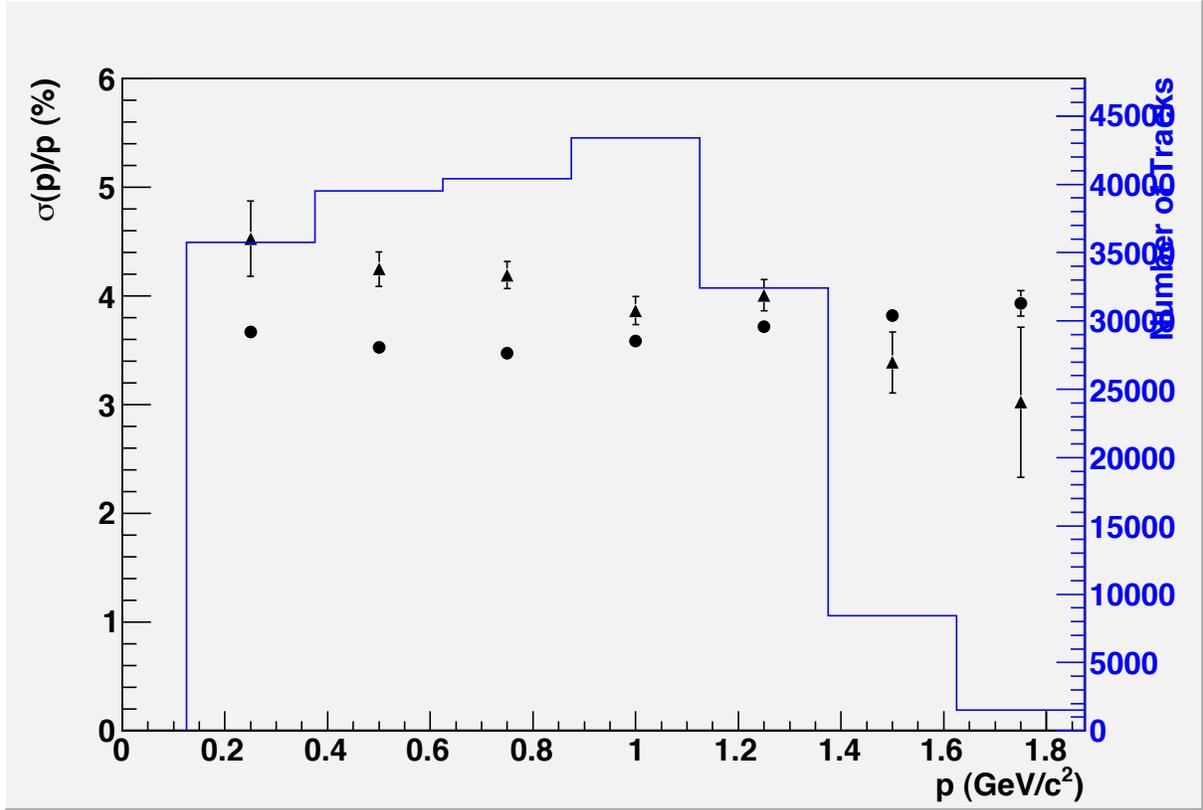


FIG. 73: Fractional momentum resolution versus momentum for a beam energy of 4.4 GeV and a magnetic field of 1.0 T.

the target. Distinguishing A' decays from the background therefore depends on the vertex resolution and in particular on the tails of the vertex distribution. In order to study the tails, we use large samples of A' events decaying promptly overlaid on top of the simulated beam background events.

Each pair of oppositely charged tracks is fit to a common vertex using a Kalman filtering method first suggested by Billoir [104], [105] and used in many experiments. The method uses the measured helix parameters and their correlations to determine the most likely decay position of the A' and also returns fitted momenta for each particle. We actually fit each pair twice with different hypotheses of their origin. We constrain either the vertex to be consistent with an A' :

- which originates in the $200\mu\text{m} \times 40\mu\text{m}$ beamspot at the target, and moves off in the direction given by the measured A' momentum. This fit will be used for the vertexing search.

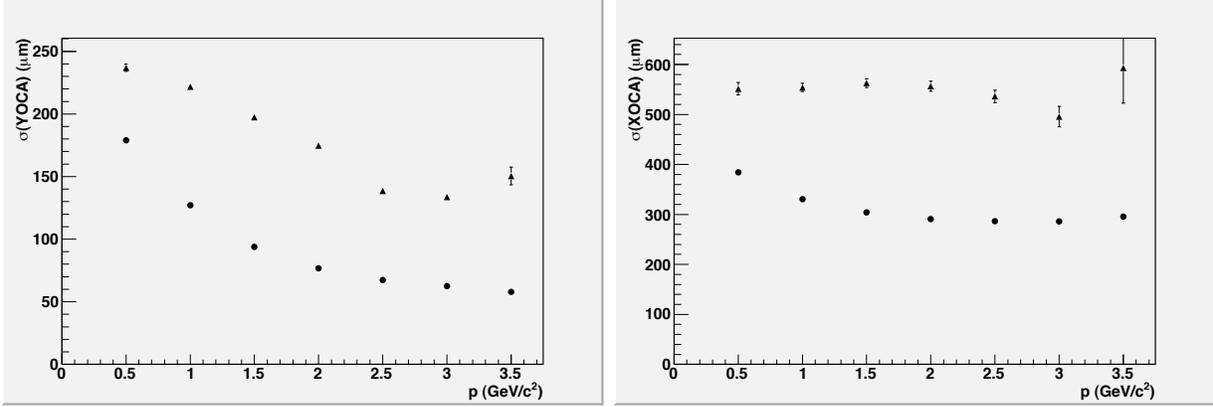


FIG. 74: The resolution of the position of closest approach to the beam axis versus track momentum in the (left) non-bend direction and (right) bend direction. The dots represent tracks with 0 bad hits and triangles with one or more.

- which originates and decays at the target within the $200\mu\text{m} \times 40\mu\text{m}$ beamspot. This fit will be used for the bump-hunt only search.

For each electron/positron pair reconstructed in the tracker, we compute the invariant mass based on the fitted momenta of the tracks. The mass resolution depends on the invariant mass of the pair and is shown in Fig. 75. The closed circles in Fig. 75 shows the improvement in the resolution for the second fit, where the decay is assumed to occur in the target.

Even for prompt decays, the z vertex position (V_z) distribution of all reconstructed e^+e^- pairs (solid black histogram, Fig. 77) shows a long tail, still significant beyond 5cm. This tail is primarily comprised of events where one or both of the tracks use one or more bad hits. Fortunately there are a number of quantities we can use to minimize the tails. Namely, for purposes of this proposal, we make the following cuts:

- The χ^2 of each track is less than 20
- The total momentum of the A' candidate is less than the beam energy
- A very loose cut on the reconstructed vertex position $|V_x| < 400\mu\text{m}$ and $|V_y| < 400\mu\text{m}$
- The clusters in layer 1 of each track must be isolated from the next closest cluster by at least $500\mu\text{m}$
- A χ^2 cut on the vertex fit of less than 5.

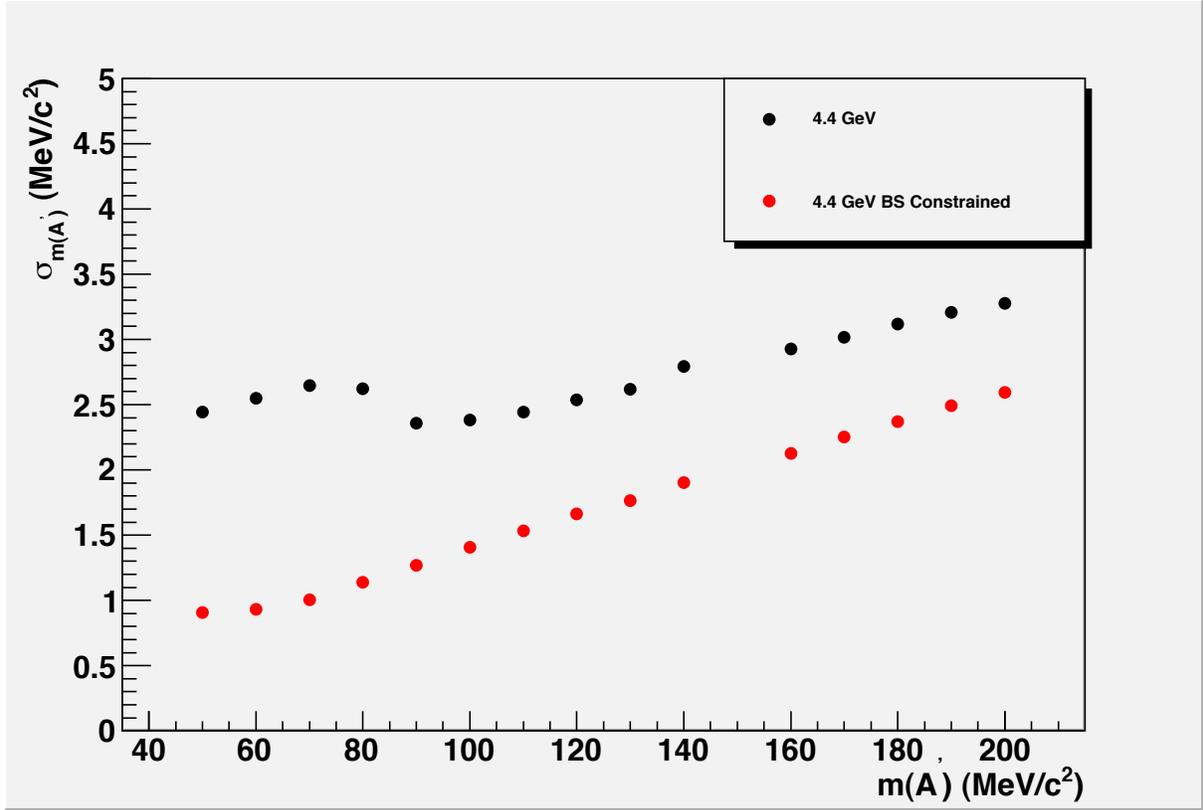


FIG. 75: The gaussian width of the mass distributions (MeV/c²) vs generated A' mass (MeV/c²). The open circles are the resolutions when the decay is constrained to the target beamspot and the closed circles are without this constraint.

The vertex resolution depends on the invariant mass of the particles being vertexed. Lower masses have worse Gaussian resolutions as shown in Fig. 76. This is expected since the error on the opening angle (θ), due to multiple scattering, scales like: $\sigma(\theta)/\theta \sim (1/E)/(m/E) \sim 1/m$.

Figure 77 shows the vertex resolution for samples of 80 MeV and 160 MeV A' events. The cuts above remove almost all of the tail past 1.5cm (points with errors in Fig. 77) while retaining 50% of the e^+e^- pairs from the A' candidate. The events on the tail are enhanced with vertices where there are one or more bad hits on the track (represented by the blue histogram in Fig. 77), although there is still a contribution from well-reconstructed tracks. The rejection of tracks with bad hits depends strongly on the precision of the virtual A' trajectory, which in turn depends on the size of the beamspot. Having a beamspot significantly smaller than the intrinsic tracker resolution, $100\mu\text{m}$ in the non-bend and $300\mu\text{m}$ in the bend directions, is important.

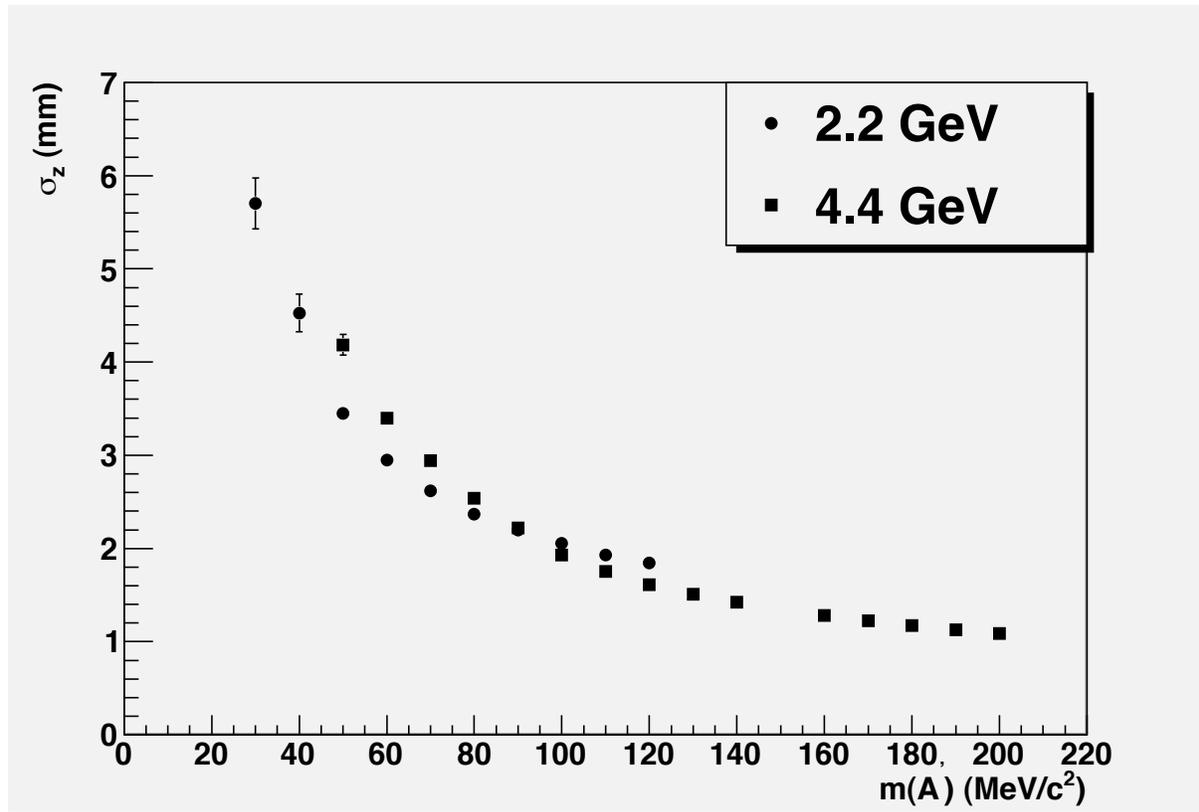


FIG. 76: The Gaussian resolution dependence versus A' mass for signal-only events.

In practice, there is much more we can do to clean up the vertex and mass resolution both at the track level (e.g. remove hits that are clearly from scattered beam electrons) and at the vertex level. These will be pursued in the near future.

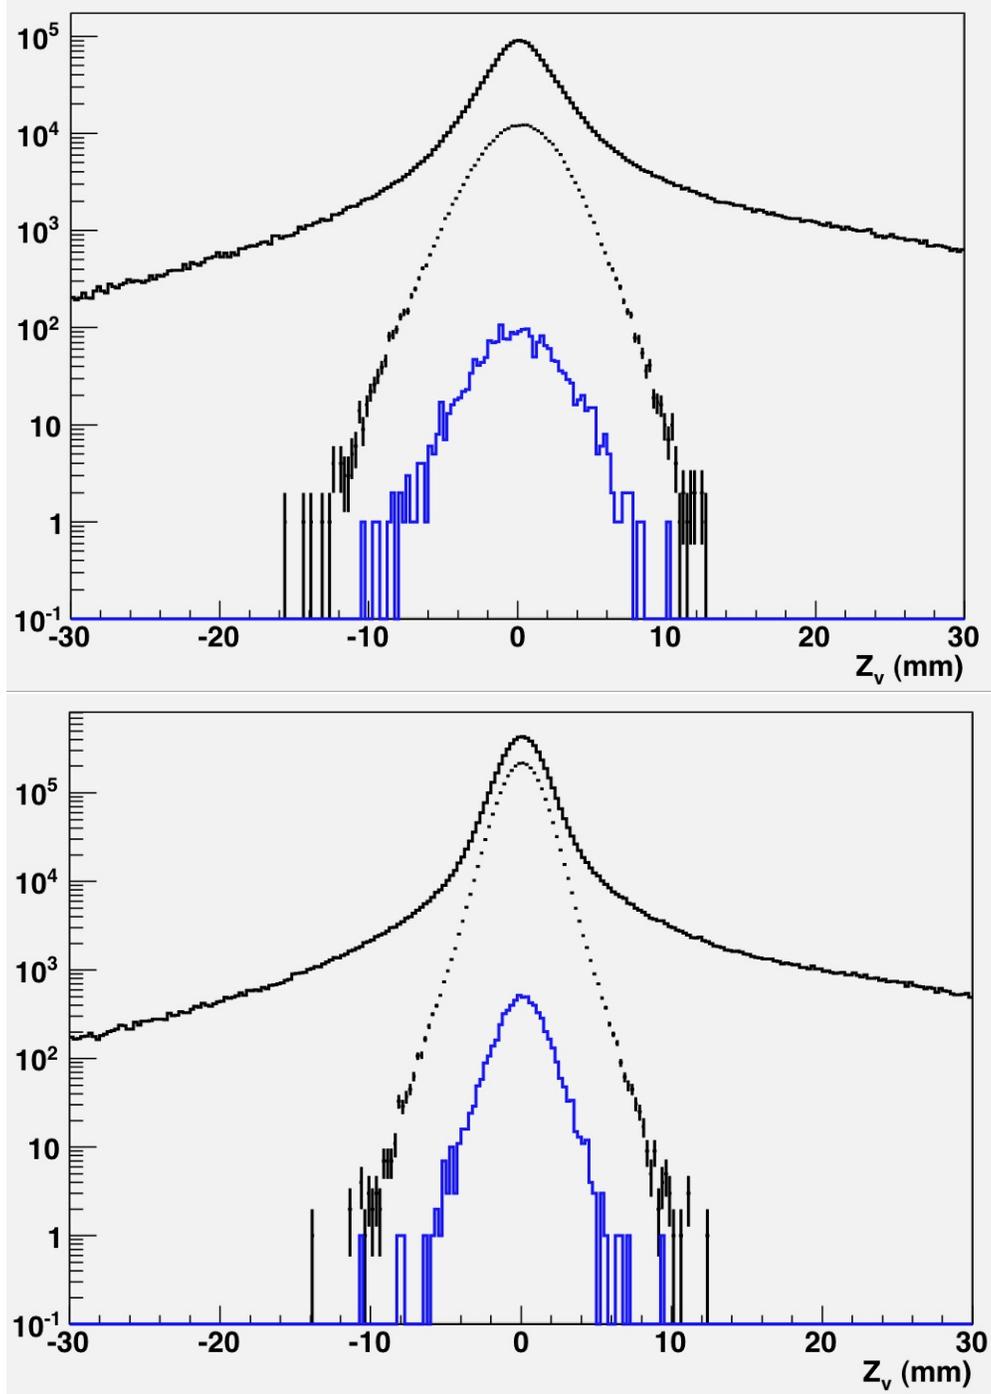


FIG. 77: Distribution of the reconstructed vertex position along the beam axis for 4.4GeV 80MeV (top) and 160MeV (bottom) A' events before (solid black) and after (points with errors) selection. The blue histogram shows the distribution for pairs that have at least one bad hit after selection.

7 Experimental Reach

The primary search channel for this experiment is $A' \rightarrow e^+e^-$, with or without a displaced vertex, depending on the magnitude of the coupling α' . As such, the primary irreducible background is QED trident production, with rate given by the diagrams shown in Figure 3. Trident events can be usefully separated into “radiative” diagrams (Figure 4(a)), and “Bethe-Heitler” diagrams (Figure 4(b)), that are separately gauge-invariant. The expected parameter reach of the experiment is shown in Figure 2. Below, we discuss how this was calculated.

The contribution from the radiative diagrams (Figure 4(a)) alone is a useful guide to the behavior of A' signals at various masses. In particular, the kinematics of A' signal events is identical to that of radiative trident events restricted to an invariant mass window near the A' mass. Moreover, the rate of the A' signal is simply related to the radiative trident cross-section within a small mass window of width $\delta m_{A'}$ by [1],

$$\frac{d\sigma(e^-Z \rightarrow e^-Z(A' \rightarrow e^+e^-))}{d\sigma(e^-Z \rightarrow e^-Z(\gamma^* \rightarrow e^+e^-))} = \left(\frac{3\pi\epsilon^2}{2N_{eff}\alpha}\right)\left(\frac{m_{A'}}{\delta m_{A'}}\right) \quad (4)$$

where N_{eff} counts the number of available decay states. A fraction ϵ_{bin} of signal events will have reconstructed masses within the mass window, because of the finite mass resolution (for a $2.5 \times \sigma$ mass resolution window, $\epsilon_{bin} = 0.8$). Equation 4 corrected for ϵ_{bin} allows us to conveniently express the sensitivity to A' signals in terms of the radiative portion of the total QED trident statistics, which we will do shortly.

The Bethe-Heitler process has a much larger total cross-section than either the signal or the radiative trident backgrounds, but exploiting its different kinematics can significantly reduce it. In particular, the A' carries most of the beam energy (see the discussion in Section 3.1) while the recoiling electron is very soft and scatters to a wide angle. In contrast, the Bethe-Heitler process is not enhanced at high pair energy. Moreover, Bethe-Heitler processes have a forward singularity that strongly favors asymmetric configurations with one energetic, forward electron or positron and the other constituent of the pair much softer. The geometric acceptance and trigger requirements select the region of phase space where signal is dominant, and the Bethe-Heitler background is smallest, as illustrated by Figure 5 (it should be emphasized, however, that even in this region the Bethe-Heitler background

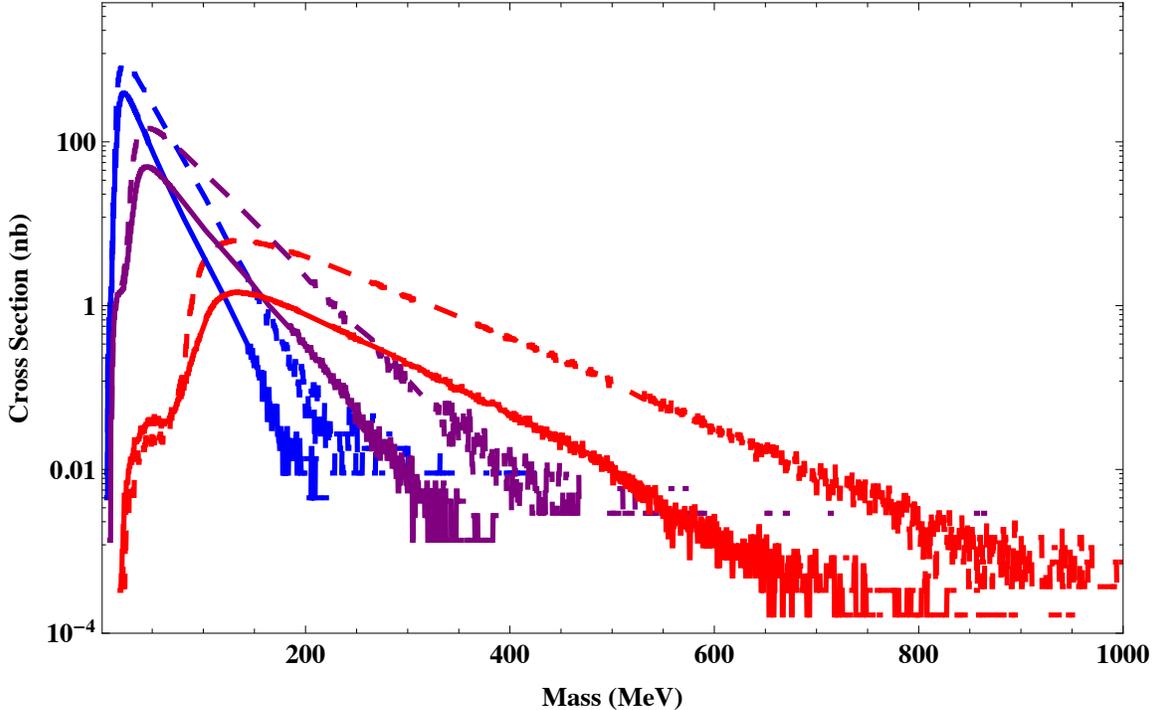


FIG. 78: The Bethe-Heitler (dashed) and radiative (solid) cross-section/MeV after acceptance for 1.1 (blue), 2.2 (purple) and 6.6 GeV (red).

rate exceeds that of radiative tridents by roughly a factor of 5). The radiative and Bethe-Heitler cross-sections, after accounting for acceptance in the HPS detector, are shown in Figure 78.

To compute the reach of the HPS experiment, we simulate the production of irreducible trident reactions in the detector. We additionally apply a mock-up of the geometric acceptance for the tracking and of the trigger requirements. In addition, high-statistics Monte Carlo samples at particular invariant masses have been used to estimate the background rejection efficiency for a vertex-based search. We produce generator-level events using MadGraph and MadEvent [106] to compute the full matrix elements for $e^-Z \rightarrow e^-(e^+e^-)Z$ in leading order QED, but neglecting the effect of nuclear excitations on the kinematics in inelastic processes. We use the QED nuclear elastic and inelastic electric form-factors in [107]. The MadEvent code was modified to properly account for the masses of the incoming nucleus and electron in event kinematics.

We use a “reduced-interference” approximation that simplifies our analysis and is much less computationally intensive. In this approximation, we treat the recoiling e^- and the e^- from the produced pair as distinguishable. Furthermore, we separate trident processes into

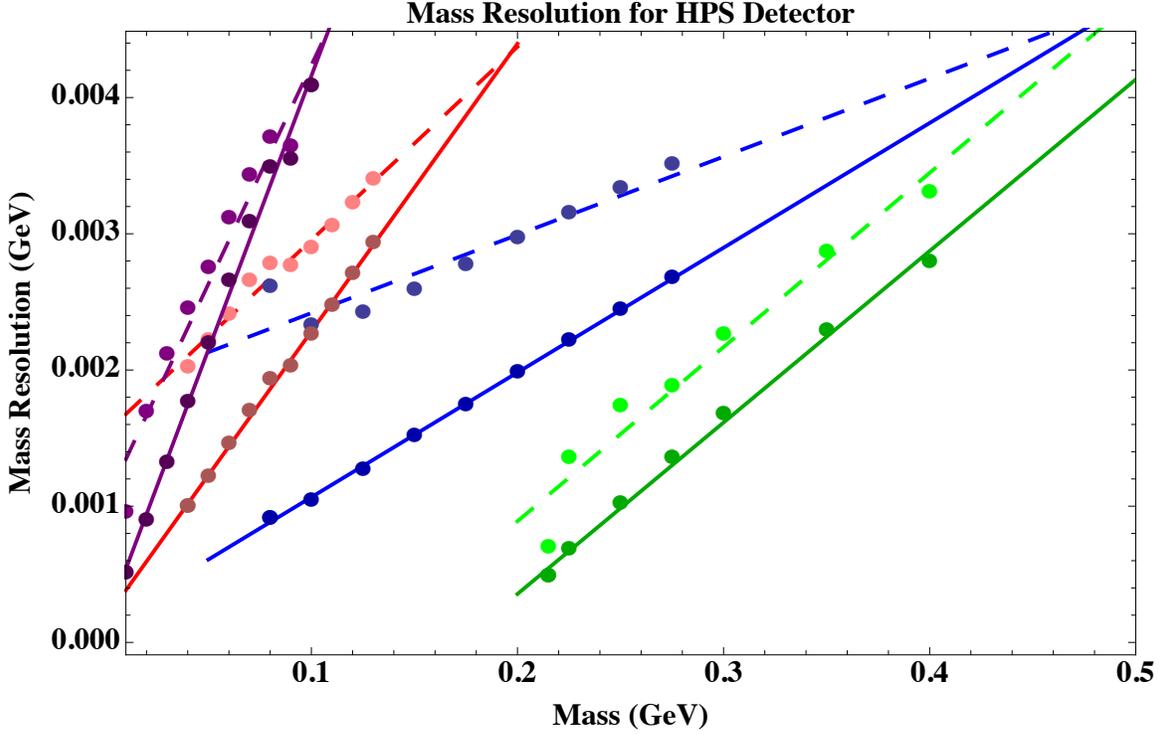


FIG. 79: The invariant mass resolution versus mass for 1.1 (purple), 2.2 (red), 6.6 GeV electron decays (blue) and 6.6 GeV muon decays (green). The points are from simulated data at various masses while the lines are linear fits to the points. For each energy, the points with worse resolution (dashed line) are from the vertex fit without constraining the decay to the target while the better resolution (solid line) require the decay to be prompt. The fitted curves are used in the reach calculation; the resolution from the dashed lines are used for the displaced vertex search while the bump-hunt search uses the resolution from the solid lines.

the radiative diagrams and the Bethe-Heitler diagrams, and we calculate the cross-section for both of these diagrams separately. Within the acceptance and signal region for the HPS experiment, the Bethe-Heitler reactions dominate the trident rate by 4:1. We have checked that the “reduced-interference” approximation does not correct the trident cross-section by more than 10% in a representative kinematic region [4].

7.1 Resonance Search

Equation 4 is used to compute the reach for a resonance search in the e^+e^- or $\mu^+\mu^-$ final state. We start by simulating radiative and Bethe-Heitler trident events and require that e^+e^- or $\mu^+\mu^-$ pairs pass the detector acceptance cuts. We additionally require that the total energy exceed 80% of the beam energy and that each track have at least 0.5 GeV

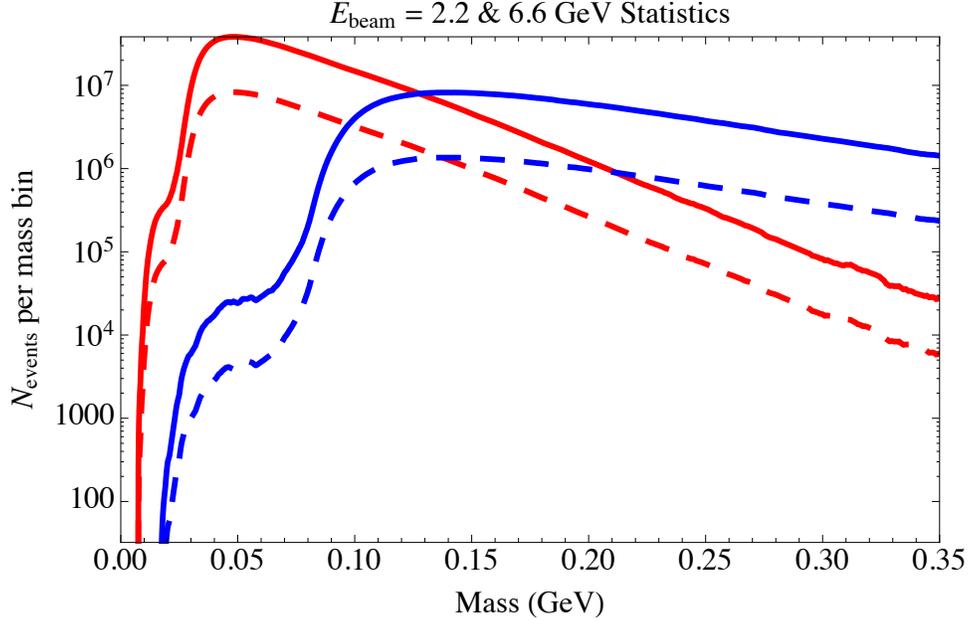


FIG. 80: Right: Distribution of statistics in full resonance search. The solid (dashed) curves indicate the number of background QED trident (radiative) events expected in a resolution-limited mass window of width $\delta m_{(A')} = 2.5\sigma(m_{A'})$. The red curve corresponds to the distribution of statistics for a one month run at 6.6 GeV beam energy (blue), with a current of 450 nA on a 0.25% X_0 target, while the dashed maroon curve corresponds to six weeks at 2.2 GeV beam energy (red) and a 0.125% X_0 target.

of energy. We will refer to these cuts collectively as the “detector/trigger mock-up”. We compute the total differential cross section, as a function of invariant mass, for radiative and Bethe-Heitler trident events to pass the detector/trigger mock-up cuts, and from this the final statistics are computed assuming a run duration according to the run plan laid out in Section 3, and conditions in Section 6.4. The mass resolution and the background statistics expected in each resolution-limited mass window are shown in Figures 79 and 80.

To quantify statistical sensitivity, we assume that the trident background in the resonance search can be modeled by a smoothly varying function and subtracted off. The significance is then determined by the ratio of the signal within an invariant mass window to $\sqrt{N_{bin}}$, where N_{bin} is the total background statistics in the same window. Using equation 4, the sensitivity for a resonance search is determined by

$$\left(\frac{S}{\sqrt{B}}\right)_{bin} = \left(\frac{N_{radiative}}{N_{total}}\right) \sqrt{N_{bin}} \left(\frac{3\pi\epsilon^2}{2N_{eff}\alpha}\right) \left(\frac{m_{A'}}{\delta m_{A'}}\right) \epsilon_{bin} \quad (5)$$

Here, $\left(\frac{N_{radiative}}{N_{total}}\right)$ is the fraction of radiative reactions among all QED trident events in the

search region. This quantity is determined by simulation as described below. N_{bin} is the total number of QED trident events residing in a given invariant mass search bin, and is determined by

$$N_{bin} \equiv \epsilon_{reco}^2 \times \epsilon_{stat}(m_{A'}) \times \sigma_{trigger} \times L. \quad (6)$$

Here L is the integrated luminosity, $\sigma_{trigger}$ is the trigger cross section, $\epsilon_{stat}(m_{A'})$ is the fraction of the total statistics in an invariant mass window centered on $m_{A'}$ of size $\delta m_{A'} = 2.5\sigma(m_{A'})$, and $\epsilon_{reco} \cong 0.85$ is the efficiency for reconstructing each track that is within the geometric acceptance of the detector.

7.2 Displaced Vertex and Resonance Search

A search for resonances that decay with cm-scale displaced vertices opens up sensitivity to much smaller couplings than can be observed through a resonance search alone. The vertex reconstruction and quality selection is discussed in Section 6.4. For the purpose of computing reach, the vertex quality requirements reduce the signal efficiency by a factor $\sigma_{sig} \sim 0.5$. We use the high-statistics Monte Carlo studies described in Section 6.4 to model the tails of the vertex distribution for decays at the target. These vertex distributions have been generated at a few different masses for each beam energy. Away from these masses, we parameterize the background rejection factor $\epsilon_{rejection}(zcut)$, the fraction of events with a fake vertex beyond a beam line distance of $zcut$, by a smooth interpolation.

Because the fake vertex distribution falls quite rapidly, the greatest sensitivity is achieved far on the vertex tail, where less than one background event is expected. For the purpose of computing reach, we have determined a mass-dependent choice of $zcut(m_{A'})$ such that the expected background in each resolution-limited mass window $\delta m_{A'}$, with reconstructed vertices beyond this cut, does not exceed 0.5 events in the run period. This requires rejection $\epsilon_{rejection}(zcut)$ of background events from the target at the level of 10^{-6} to 10^{-7} , achieved for $zcut \sim 5\text{-}30$ mm (see Figure 81).

The geometric acceptance falls off at decay lengths greater than 10 cm. For simplicity we compute reach using the geometric acceptance for $z=0$, but only considering decays with

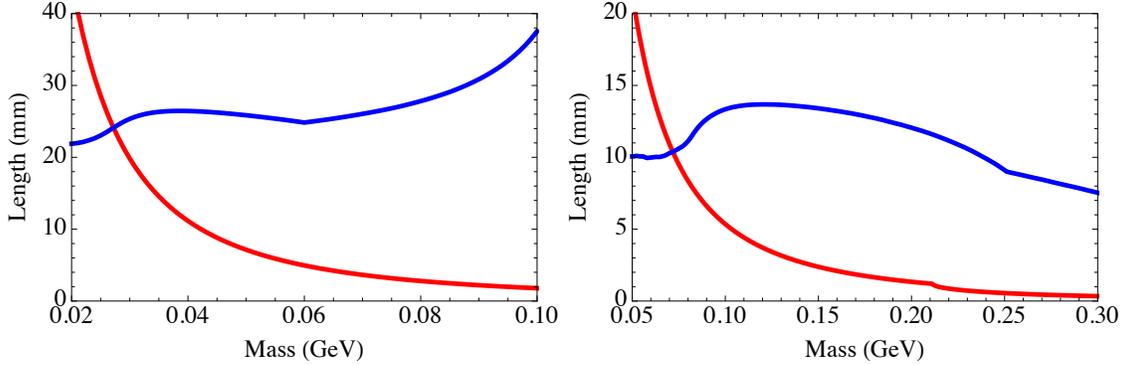


FIG. 81: In blue: the minimum vertex displacement z_{min} (in mm) along the beamline, required for the vertex-based resonance search at 2.2 GeV (left) and 6.6 GeV (right). These are chosen to bring the expected background to 0.5 events in each resolution-limited mass window. In red: the $\gamma c\tau$ assuming $\epsilon = 10^{-4}$.

$z < z_{max}=10$ cm, so that the fraction of signal events included in the vertex search is

$$\epsilon_{sig}(zcut) \cong \epsilon_{vtx} \times \left(e^{-\frac{zcut}{\gamma c\tau}} - e^{-\frac{zmax}{\gamma c\tau}} \right). \quad (7)$$

Accounting for the reduced acceptance of both signal and background events, the statistical significance expected can be computed from that of the pure resonance search as:

$$\left(\frac{S}{\sqrt{B}} \right)_{bin,zcut} = \left(\frac{S}{\sqrt{B}} \right)_{bin} \frac{\epsilon_{sig}(zcut)}{\sqrt{\epsilon_{rejection}(zcut)}} \quad (8)$$

where $(\frac{S}{\sqrt{B}})_{bin}$ is given by (5). For the small expected background rate (0.5 events/bin), however, this formula becomes irrelevant, as the exclusion sensitivity of the experiment is limited by the probability of a downward fluctuation in the signal. Thus, for the vertex reach contours in Figure 2, we additionally require an expected signal

$$S_{bin,zcut} = \left(\frac{N_{radiative}}{N_{total}} \right) N_{bin} \left(\frac{3\pi\epsilon^2}{2N_{eff}\alpha} \right) \left(\frac{m_{A'}}{\delta m_{A'}} \right) \epsilon_{sig}(zcut) > 2.4 \text{ events} \quad (9)$$

7.3 Reach in Mass-Coupling Parameter Space

Using the $S/\sqrt{(B)}$ for the bump-hunt and displaced vertex searches as described above, we estimate to cover the regions of coupling vs mass parameter space shown in Figure 82. The contours in the plot show the the two-sigma exclusion regions for:

- purple, dashed: 1 week of 50nA, 1.1 GeV beam on a 0.125% target
- blue, dashed: 1 week of 200nA, 2.2 GeV beam on a 0.125% target
- blue, solid: 3 weeks of 200nA, 2.2 GeV beam on a 0.125% target
- dark green: 2 weeks of 450nA, 6.6 GeV beam on a 0.25% target, detecting $A' \rightarrow e^+e^-$
- light green: 2 weeks of 450nA, 6.6 GeV beam on a 0.25% target, detecting $A' \rightarrow \mu^+\mu^-$
- red: the statistical combination of all of the above
- green shaded: 3 months each of 2.2 GeV and 6.6 GeV (same currents and thicknesses as above)

8 Run Plan and Beam Time Request

The Jefferson Laboratory PAC39 graded HPS physics with an "A", approved a commissioning run with electrons, and granted so-called C1 approval for the full HPS experiment. The total requested beam time for the experiment is 180 days. For the upcoming commissioning run we have requested 3 weeks. Anticipating early running in Hall-B, we propose to conduct the experiment in two phases. The first phase, expected to run in 2014-2015, will complete the commissioning run and begin the production running at 2.2 GeV and 6.6 GeV beam energies as detailed below. The second phase, which will use remaining beam time, can be scheduled later in 2015 or in 2016 and beyond, and will continue runs at 2.2 and 6.6 GeV and possibly other energies.

We plan to execute the first phase of the experiment in two run periods using the apparatus described above. First, we will perform a commissioning run which should produce our first physics output, and then, after a month or two down time, continue with a longer run at multiple beam energies to cover as much parameter space as possible. The experimental apparatus, if funded on time, will be ready to be commissioned and take physics data in the fall of 2014 when the first physics quality beams should be available in Hall-B. The proposed run plan for phase one is as follows:

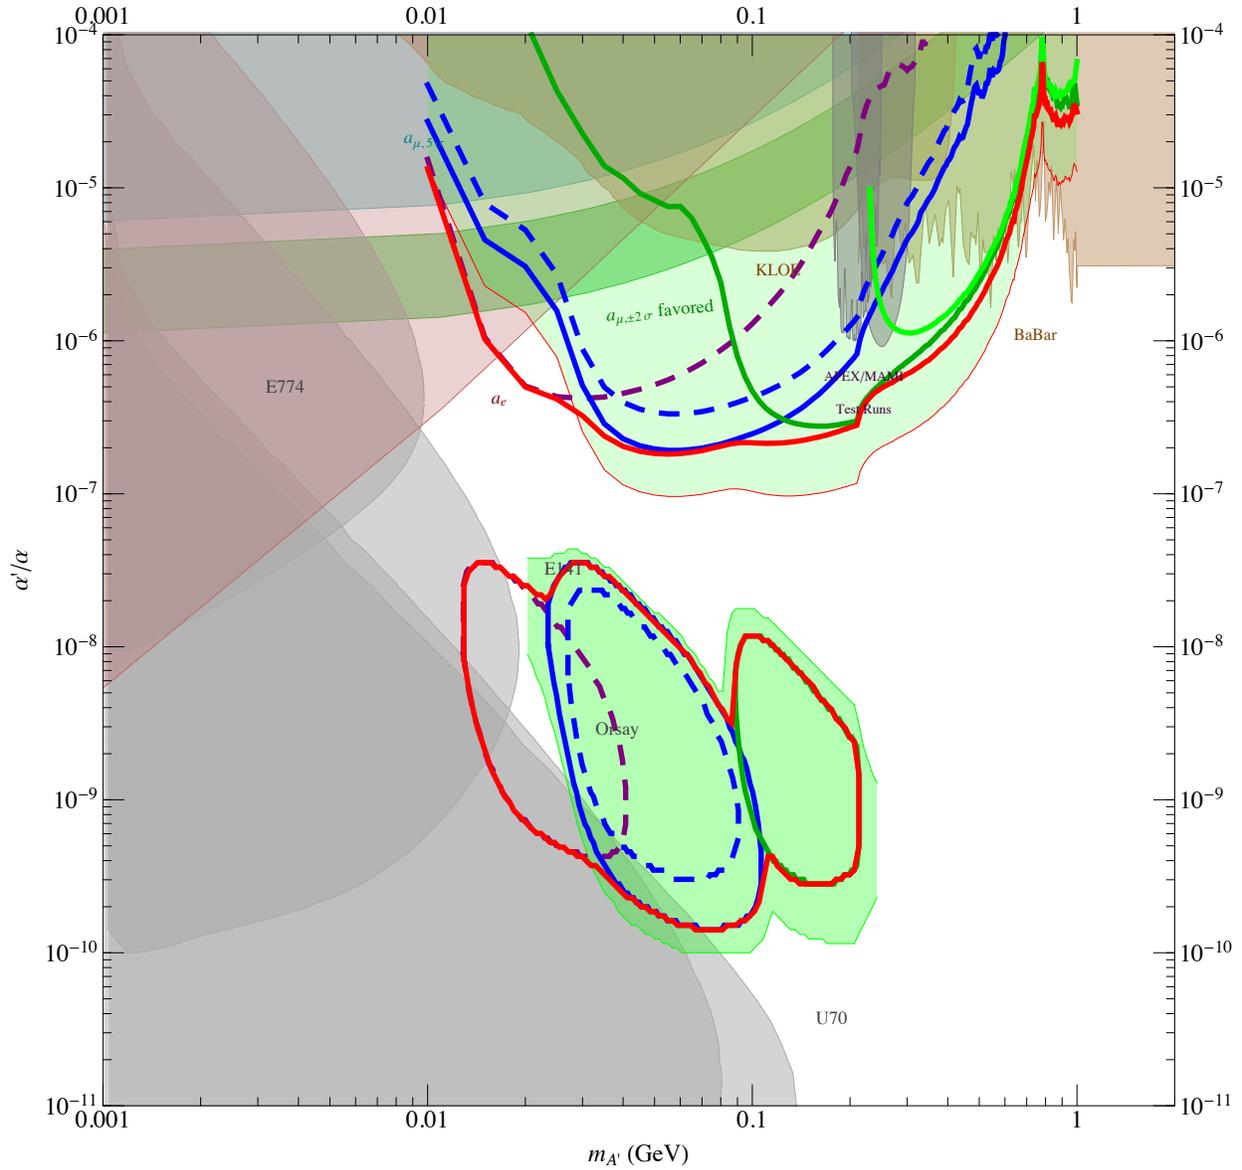


FIG. 82: The expected reach in mass-coupling parameter space. See text for details about different features.

- **Commissioning run in 2014, total of 3 weeks of beam time (6 weeks on the floor assuming 50% for combined efficiency of the accelerator and the detector):**
 - 1 week of detector commissioning
 - 1 week of physics run at 2.2 GeV
 - 1 week of physics run at 1.1 GeV
- **Physics run in 2015, total of 5 weeks of time beam (10 weeks on the floor**

assuming 50% for combined efficiency of the accelerator and the detector):

- 1 week of detector commissioning
- 2 weeks of physics run at 2.2 GeV
- 2 weeks of physics run at 6.6 GeV

If more beam time is available in 2015, HPS will continue data taking at these or other energies.

The proposed run plan will cover the remaining region of parameter space favored by the muon g-2 anomaly, and will explore a significant region of parameter space, not only at large couplings ($\alpha'/\alpha > 10^{-7}$), but also in the regions of small couplings, down to $\alpha'/\alpha \sim 10^{-10}$. This small coupling region is not accessible to any other proposed experiment. The excellent vertexing capability of the Si-tracker uniquely enables HPS to cover the small coupling region.

In the proposed run plan we have assumed running at a non-standard energy for 12 GeV CEBAF, $E = 1.1$ GeV, in 2014. In case this will not be possible due to scheduling conflicts, we will continue to run at 2.2 GeV instead. We would then expect to complete the 1.1 GeV running at some time in 2015 by reducing the time at 2.2 GeV. The gap between the run periods in 2014 and 2015, on the order of two months, will be used to improve, correct or fix apparatus if necessary.

In summary, we request time for the first phase of HPS experiment, first in the fall of 2014 for a total of 3 weeks of beam time (6 weeks on the floor), with beam energies 1.1 GeV, two weeks, and 2.2 GeV, two weeks. Second, we request 5 weeks of beam (10 weeks on the floor) that will be equally shared between beam energies of 2.2 GeV and 6.6 GeV. We expect that the second part of the run will be in 2015.

9 Schedule and Cost Baseline

Cost estimates for engineering, designing, fabricating, assembling, testing, and installing the Heavy Photon Search detector are given below. The costs assumes considerable savings from the reuse of many parts of the Test Run, which was already assembled made from the donated silicon microstrip sensors from Fermilab, the use of some DAQ crates and equipment

from SLAC, and many contributions from JLab, including PbWO_4 calorimeter crystals, the chicane and analyzing magnets, magnet power supplies, and beam diagnostic apparatus. Much of the calorimeter readout electronics utilizes designs which are already in place for the Hall B 12 GeV upgrade, eliminating engineering and design expense. The SVT DAQ benefits from SLAC's development of an ATCA readout system, and incorporates many of its existing designs. Very significant cost savings come from utilizing the FADCs and data acquisition system being developed for the upgraded CLAS12 detector, which will be available free of charge to HPS. The Orsay group is contributing engineering and design efforts for the ECal and its vacuum chamber, affording additional savings.

The costs are given in an accompanying WBS summary table, below, which itemizes the major items subsystem by subsystem, and indicates whether JLab (J) or SLAC (S) takes responsibility for construction. Engineering, design, and technician labor rates are fully loaded, including benefits and lab overheads, which differ between the two laboratories. Our DAQ and beamline cost estimates have been made by engineering groups at SLAC and JLab which are experienced in cost estimation and actively involved in many related projects. The SVT estimates came from physicists and engineers on the project, with experience in designing and fabricating silicon detector systems. The Ecal estimates come from physicists and engineers at JLab and Orsay who have constructed a similar system, the CLAS IC, in the recent past.

The schedule for the overall project is included in a Project Summary table below. A brief description of the schedule for the different subsystems is also given. The overall schedule contingency is about 20%, and depends critically on the assumption that funding is available by mid CY2013. The HPS construction project has been organized into a Work Breakdown Structure (WBS) for purposes of planning, managing and reporting project activities. Work elements are defined to be consistent with discrete increments of project work. Project Management efforts are distributed throughout the project, including conceptual design and R&D. The HPS has 12 WBS Level-2 elements, see Table XXII.

9.1 Cost

The costs include Labor and M&S. The labor includes only engineering or technician manpower in professional centers at SLAC or JLAB. It does not include labor provided

TABLE XXII: Project WBS structure.

WBS	NAME
1.1	Beamline
1.2	SVT
1.3	SVT DAQ
1.4	ECAL
1.5	Muon
1.6	TDAQ
1.7	Slow Controls
1.8	Installation & Commissioning
1.9	Electron Running
1.10	SLAC Travel Meetings
1.11	SLAC Travel Running
1.12	Project Management

by physicists, which is the dominant contribution to the project. Labor rates have been applied following the official shop rates at SLAC and JLAB, which include already 31% or 57% fringe benefits, respectively. M&S have been determined from a best estimation of the commercially available parts, benefiting from our experience with the actual costs of the HPS Test Run. The overheads have been added to both labor and M&S, being respectively 53% and 7.65% at SLAC, 49% for both labor and M&S at JLAB. SLAC travel includes 53% overheads. Contingencies have been set at 10% for catalogue items, 20-25% for items similar to previous design, 30-50% if new design. Since the project is staged over three years, an annual inflation rate of 2.5% is included in the FY14 and FY15 costs.

Beamline expenses for HPS are held to a minimum by using the 18D36 magnet currently installed in Hall B as the analyzing magnet, the two existing JLab Frascati chicane magnets and the existing Test Run vacuum chamber with the SVT vacuum box. Some overall engineering and design will be required, beam pipes fabricated, the Muon System vacuum chamber designed and built, a vacuum chamber built for the downstream Frascati magnet, and a photon dump and shielding inserted behind the second chicane magnet. Total beamline expenses are about \$324k.

Three out of the five planes of the SVT Test Run will be reused after modification of their supports which will provide improved mechanical stability and better cooling. Three

new planes with double sensors and their supports will be designed and built from scratch. Fermilab will donate the needed silicon microstrip detectors, as it had for the HPS Test Run. The tracker/vertexer will cost about \$639k.

The SVT DAQ requires small modifications to the existing hybrid; new readout and flange board engineering design, prototyping, and production; APV25 and chip procurement; and fabrication and testing. The SVT DAQ also requires designing and prototyping the Trigger Interrupt ACTA card and new firmware for the APV25 to provide event buffering to accommodate higher trigger rates. ATCA crates, and standard RCE cards are also required. The expenses are dominated by engineering development, and total \$607k.

JLab will donate the PbWO_4 crystals used in the ECal. Orsay will donate engineering and design for a new enclosure for the crystals, but Jlab will need to fabricate the enclosure, the crystal support structure, the readout motherboard and connection board, and support fixtures. It will also make repairs to the existing motherboards and acquire new power supplies. The total expense will be roughly \$312k, including fabrication, assembly, and testing.

Trigger and DAQ electronics for the ECAL are being developed for the CLAS upgrade, so relatively little engineering and technician time will be needed in preparation of the HPS Test Run except for providing special purpose firmware. Many components, including the 250 MHz FADC boards and crates will be provided at no cost since they can be borrowed from the CLAS upgrade. The system test expenses will also be borne by JLab Hall B. The total cost is \$184k.

The Muon system costs are in purchasing photomultipliers and scintillator, designing and fabricating the absorbers and Muon System vacuum chamber, and building support stands and providing cables. The total is \$419k.

The Slow Controls are needed to monitor the operations of the three sub-detectors. In addition, they will control and interlock the movements of the SVT with respect the beamline and provide beam protection interlocks. The total cost is \$221K, which is essentially the labor required to integrate the HPS with the existing Slow Control system in the Hall-B.

The offline computing resources will be provided by JLAB with the exception of the local storage at SLAC for analysis streams, ntuples, code releases and scratch areas.

Travel and lodging expenses for SLAC trips to JLab are also included in this proposal. During design and construction, there will be a small number of trips to solidify and review

designs, and to work together to begin DAQ integration of the SLAC and JLab systems. Funds are reserved for collaboration meetings to be held during calendar 2013 and 2014 at JLab. Funds are also reserved to staff installation, commissioning, and data taking runs. The total is \$172K.

The total cost for HPS is \$3.035 M.

HPS is seeking funding from other sources for the Muon System and the Ecal. William&Mary will submit an MRI proposal to NSF for the Muon System, requesting \sim \$200k. IPN ORSAY (France) has submitted a proposal to a French funding agency for the ECal Light Monitoring System (\$100k) and for a new, high performance APDs to improve crystal readout (\$500k) and other expenses related to ECal fabrication and test. Note that the new APDs are not part of this proposal. If these requests are approved the corresponding funding will be subtracted from the total cost of the HPS.

9.2 Schedule

Our goal is to be ready to install the HPS in Hall-B at JLAB by September 2014, and proceed with commissioning on beam with the CEBAF early physics beam window opportunity in October 2014. The data taking will until summer 2015. Meeting this schedule will require approval and funding as soon as possible, preferably by June 2013. Schedules for each of the major subsystems of the experiment are attached below, and summarized here. The total construction schedule extends over 16 months, assuming the funding available mid-2013. The schedule contingency is about 20%.

The conceptual design of the beamline will be done during 2013. Formal beamline engineering will start when funding is secured. A Beamline Engineering Design Review will be held in February 2014 to validate the concept before the start of major spending. Final Engineering and Construction will start in Spring 2014 and be completed well before the installation time in October 2014, providing substantial float.

Using keep-alive funds, the Test Run SVT will be shipped back to SLAC by early February 2013 to rework the modules for the first three layers of the HPS and commission the motion control systems. The conceptual design of the Layers 1-2-3 and Layers 4-5-6 will start already in spring 2013 using keep alive funding. An Engineering Design Review of the SVT will be held in November 2013, after the funding has been released and before the final production.

Preliminary engineering for the SVT DAQ is already underway, using the same keep alive funds.

The SVT DAQ will formally begin work in the second half 2013, after funding and an Engineering Review. The assembly and integration test at SLAC are expected in spring 2014 and the SVT will be ready for shipping on July 2014. The SVT will be ready for installation in mid-August 2014. SVT installation in the analyzing magnet vacuum chamber will occur in September, depending on the schedule of the Hall-B. The SVT schedule has 1 month of float between the shipping and the test at JLAB, which will be eventually used as contingency for construction work at SLAC.

The Ecal work will start in the second half 2013 and run through June 2014. The ECAL will be ready for installation by June 2014.

The Muon System work will start in July 2013 with a design phase which will be validated in an Engineering Design Review in September 2013. It will be ready for installation on August 2014, with one month of float with respect to the installation date.

The schedule includes seven milestones to track the progress of each subsystem. They will monitor the subdetector readiness after testing at the respective assembly sites, and the readiness for installation at JLAB. Also, ad-hoc Engineering Design Reviews will be conducted by the PM for each subsystem before major costs are incurred.

TABLE XXIII: Project Milestones.

WBS	Milestones	Date
1.3.2.8	Flange Board ready	25-Oct-13
1.3.1.11	FE Board Ready	25-Nov-13
1.3.4.6	Hybrid Ready	25-Nov-13
1.2.15	Layer 1-3 Ready	14-Mar-14
1.3.5.6	Flex Cable Ready	17-Mar-14
1.2.16	Layer 4-6 Ready	18-Apr-14
1.3.3.15	DAQ Ready	25-Apr-14
1.2.12	SVT Ready to Ship	16-Jun-14
1.2.14	SVT Ready For Installation	15-Aug-14
1.5.9	Muon Ready for installation	11-July-14
1.4.12	ECAL Ready for the installation	8-Aug-14
1.1.21	HPS Installed on Beamline	26-Sep-14
1.10.7	HPS ready for the beam	27-Sep-14

TABLE XXIV: Planned Review.

WBS	Engineering Reviews	Data
1.2.2	SVT Design Review	14-Oct-13
1.5.1	Muon Design Review	4-Nov-13
1.10.1	Installation Review	25-Aug-14

9.3 Manpower

The manpower needed to design, fabricate, assemble, test, install, and commission the HPS is captured in the WBS tables. The HPS Collaboration has the personnel needed to realize this project.

Beamline design work will be done at JLab by Arne Freyberger, F-X Girod and Stepan Stepanyan and at SLAC by Ken Moffeit; engineering at SLAC by Marco Oriunno, Dieter Walz, and Clive Field; fabrication in the JLab shops; and installation by the Hall B crew.

Engineering for the Ecal is being done by Philippe Rosier at Orsay in consultation with Marco Oriunno.

Beam diagnostics and slow control will be supported by Nerses Gevorgyan (Yerevan) and Hovanes Egiyan.

The Tracker/Vertexer will be engineered and designed by Marco Oriunno, Tim Nelson and Per Hansson, with additional help from Vitaly Fadeyev, Alex Grillo, and Bill Cooper, all experienced with silicon detector systems. Others at SLAC and UCSC will help with test and assembly, including Matt Graham, Takashi Maruyama, John Jaros, and graduate students Sho Uemura and Omar Moreno. Matt McCulloch will serve as the technician at SLAC.

The Ecal is being designed by the Orsay Group, especially Philippe Rosier, Emmanuel Rindel, Emmanuel Raully, Raphael Dupre, and Michel Guidal, with participation by the Jlab group, especially Stepan Stepanyan, and F.-X. Girod. Others at JLab and in the collaboration will help in assembly and test of the ECal, especially the group from INFN Genova(Italy).

The SVT DAQ is being done by Haller's group at SLAC, including Gunther Haller, Ryan

Herbst, Tung Phan, and Raghuveer Ausoori. SVT Physicists Per Hansson, Alex Grillo, Vitaliy Fadeyev, and Tim Nelson will collaborate closely. Postdocs and students will help debug, test, and certify DAQ electronics.

The Ecal Trigger/DAQ work is done in Sergey Boyarinov’s group, which supports Hall B activities at JLAB, and with Chris Cuevas’s group, which has designed the FADC250. R. Dupre and V. Kubarovsky will collaborate with this group in assembling and testing the electronics, programming the trigger, and integrating it with the Ecal hardware.

The HPS collaboration is nearly 60 strong, so has adequate manpower for overall installation, commissioning, and data taking. Simulation work is supported by Maurik Holtrop, Matt Graham, M.Ungaro, and Takashi Maruyama, along with help from students Sho Uemura and Omar Moreno, and Norman Graf and Jeremy McCormick at SLAC. Data management and storage and computing infrastructure will be overseen by Sergey Boyarinov and Maurik Holtrop and Homer Neal, all very experienced professionals. Analysis and simulation studies have been initiated by Maurik Holtrop, Matt Graham, Sho Uemura, and Takashi Maruyama. Students are actively being engaged.

The HPS collaboration is managed by its three spokespersons, Maurik Holtrop, John Jaros, and Stepan Stepanyan and its Executive Committee, which consists of the spokespeople along with Takashi Maruyama, Matt Graham, Tim Nelson, and F-X Girod. Ten working groups supervise the progress of each sub-system. The Project Manager is Marco Oriunno.

TABLE XXV: Working groups.

HPS working Groups	Chair (Deputy)
Beamline	K. Moffeit (FX Girod)
SVT	T.Nelson (V.Fedayev)
ECAL	R. Dupre (S.Stepanyan)
DAQ	S. Boiarinov (P.Hansson)
Trigger	V. Kubarovsky (T.Maruyama)
Slow Control	H. Egiyan (N. Gevorgyan)
Muon	K.Griffioen (Y.Gershtein)
Software	M.Holtrop (S. Uemura)
Analysis	M. Graham (O. Moreno)
Project Management	M. Oriunno (S. Stepanyan, J.Jaros)

TABLE XXVI: Total Labor (FTE).

SLAC					JLAB			
FTE	FY13	FY14	FY15	TOT	FTE	FY13	FY14	TOT
ME	0.14	0.67	0.10	0.91	ME	0.05	0.23	0.28
MD	0.11	0.35	0.00	0.46	MD	0.17	0.48	0.65
MT	0.09	0.63	0.00	0.72	MT	0.00	0.31	0.31
EE	0.21	1.35	0.03	1.59	EE	0.16	1.7	1.86
ET	0.00	0.02	0.00	0.03	ET	0.03	0.12	0.15

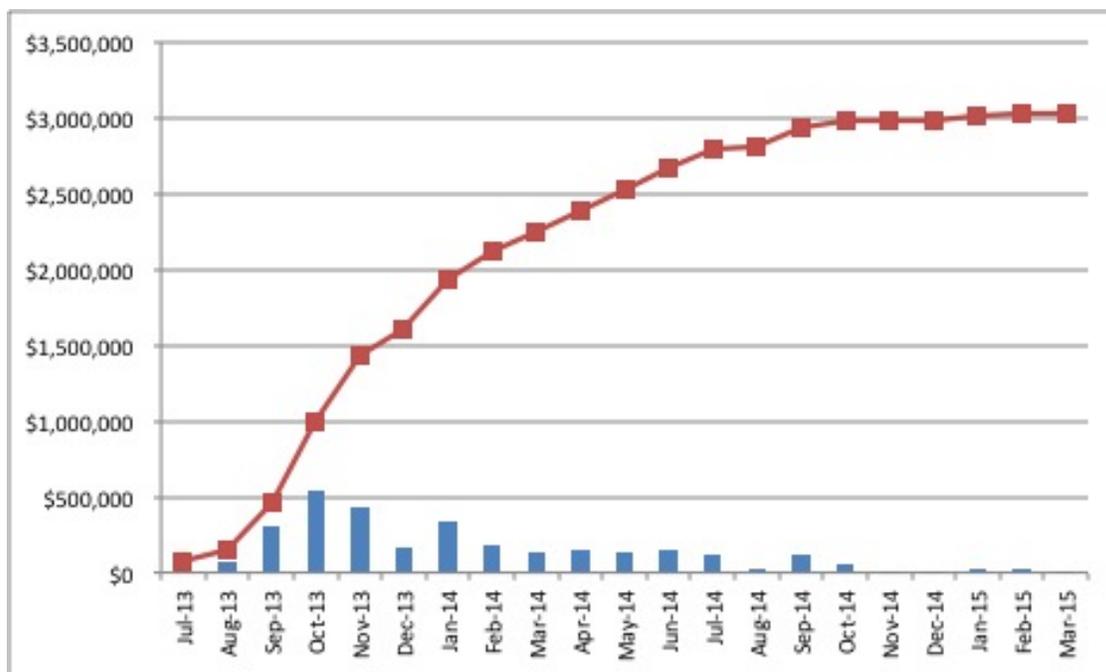


FIG. 83: Spending profile (costs after overheads and contingency).

TABLE XXVII: Summary of HPS Budget.

WBS	Name	Labor	L Cont	Material	M cont	Labor Tot	Material Tot	Total
1	HPS	\$1,414kK	\$402K	\$930K	\$288K	\$1,816K	\$1,219K	\$3,036K
1.1	Beamline (J)	\$100K	\$30K	\$149K	\$45K	\$130K	\$193K	\$324K
1.2	SVT (S)	\$363K	\$113K	\$124K	\$40K	\$475K	\$163K	\$639K
1.3	SVT DAQ (S)	\$327K	\$87K	\$142K	\$51K	\$414K	\$193K	\$607K
1.4	ECAL (J)	\$32K	\$10K	\$208K	\$62K	\$42K	\$270K	\$312K
1.5	Muon (J)	\$74K	\$25K	\$244K	\$75K	\$99K	\$319K	\$419K
1.6	TDAQ (J)	\$141K	\$43K	\$0	\$0	\$184	\$0	\$184K
1.7	Slow Control (J)	\$113K	\$28K	\$64K	\$16K	\$141K	\$80K	\$22K
1.8	Installation & Commissioning (S)	\$59K	\$18K	\$0	\$0	\$76K	\$0	\$76K
1.10	SLAC Travel Meetings (S)	\$64K	\$9K	\$0	\$0	\$73K	\$0	\$73K
1.11	SLAC Travel for Setup and Running (S)	\$78K	\$21K	\$0	\$0	\$99K	\$0	\$99K
1.12	Project management (S)	\$63K	\$19K	\$0	\$0	\$82K	\$0	\$82K

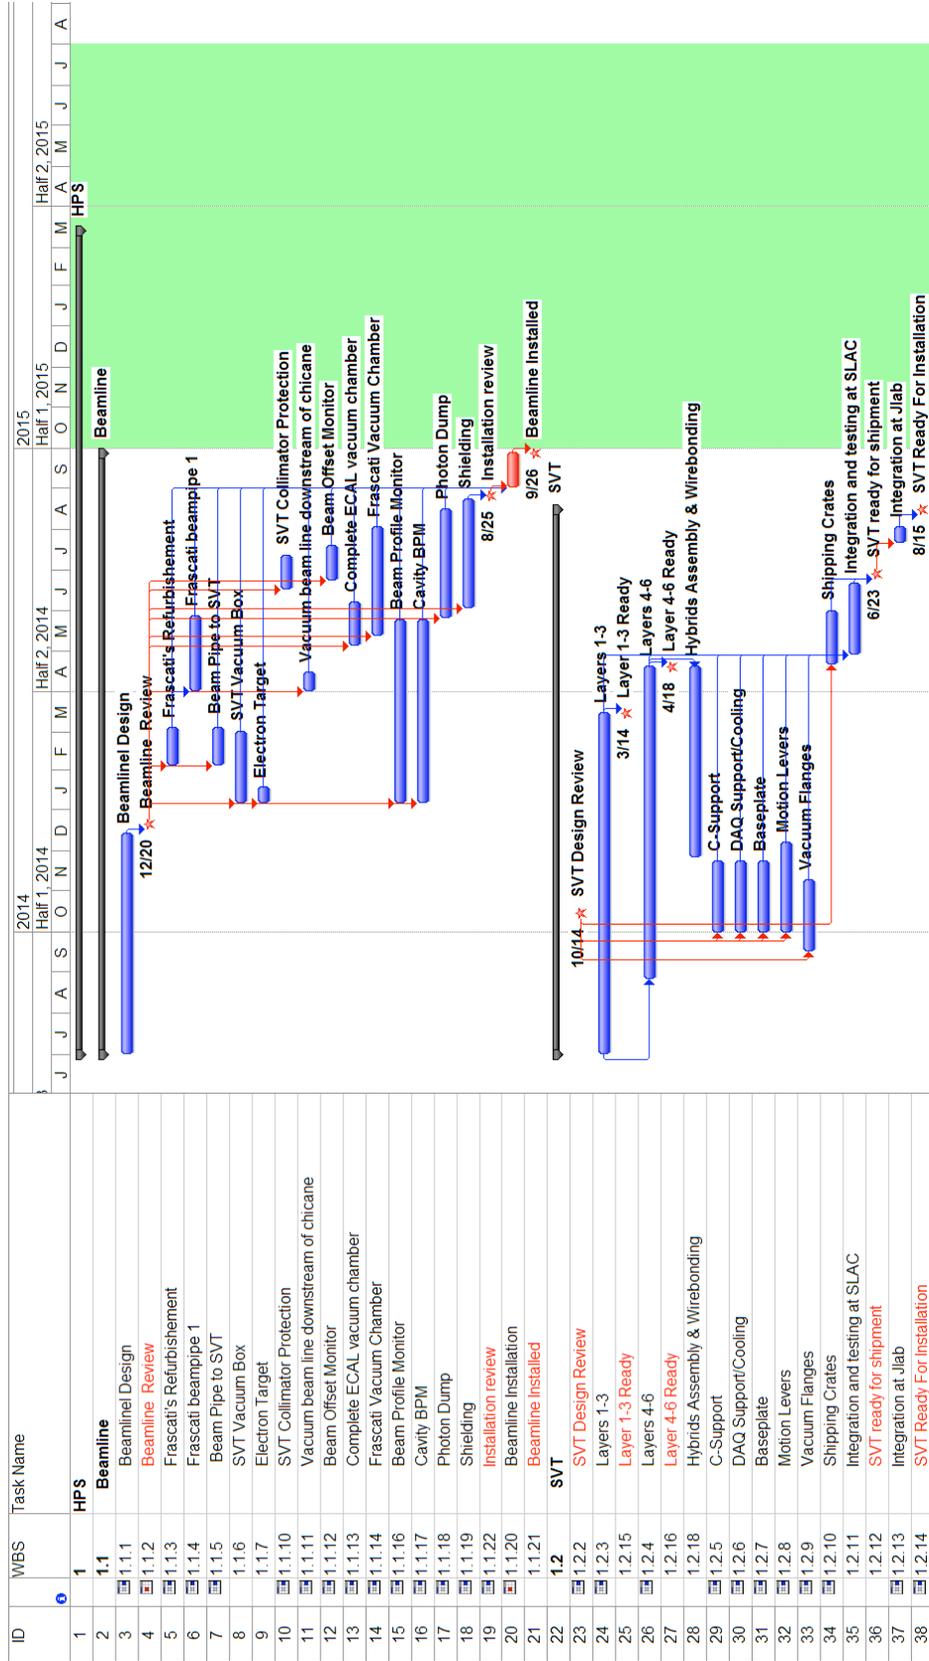


FIG. 84: HPS schedule.

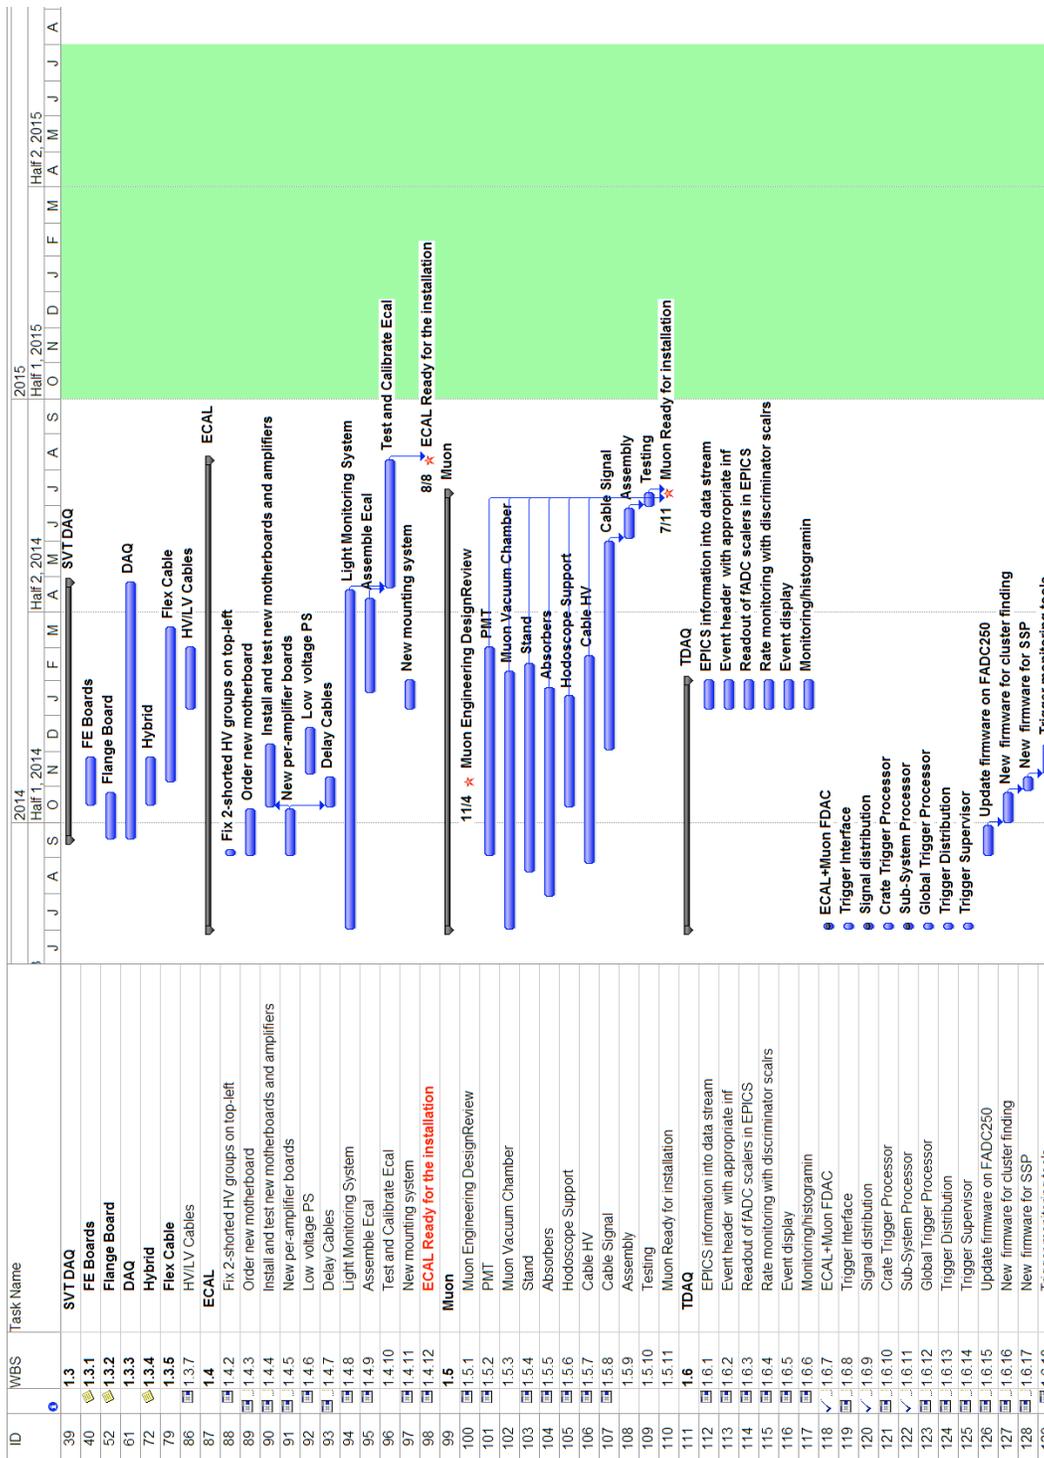


FIG. 85: HPS schedule.

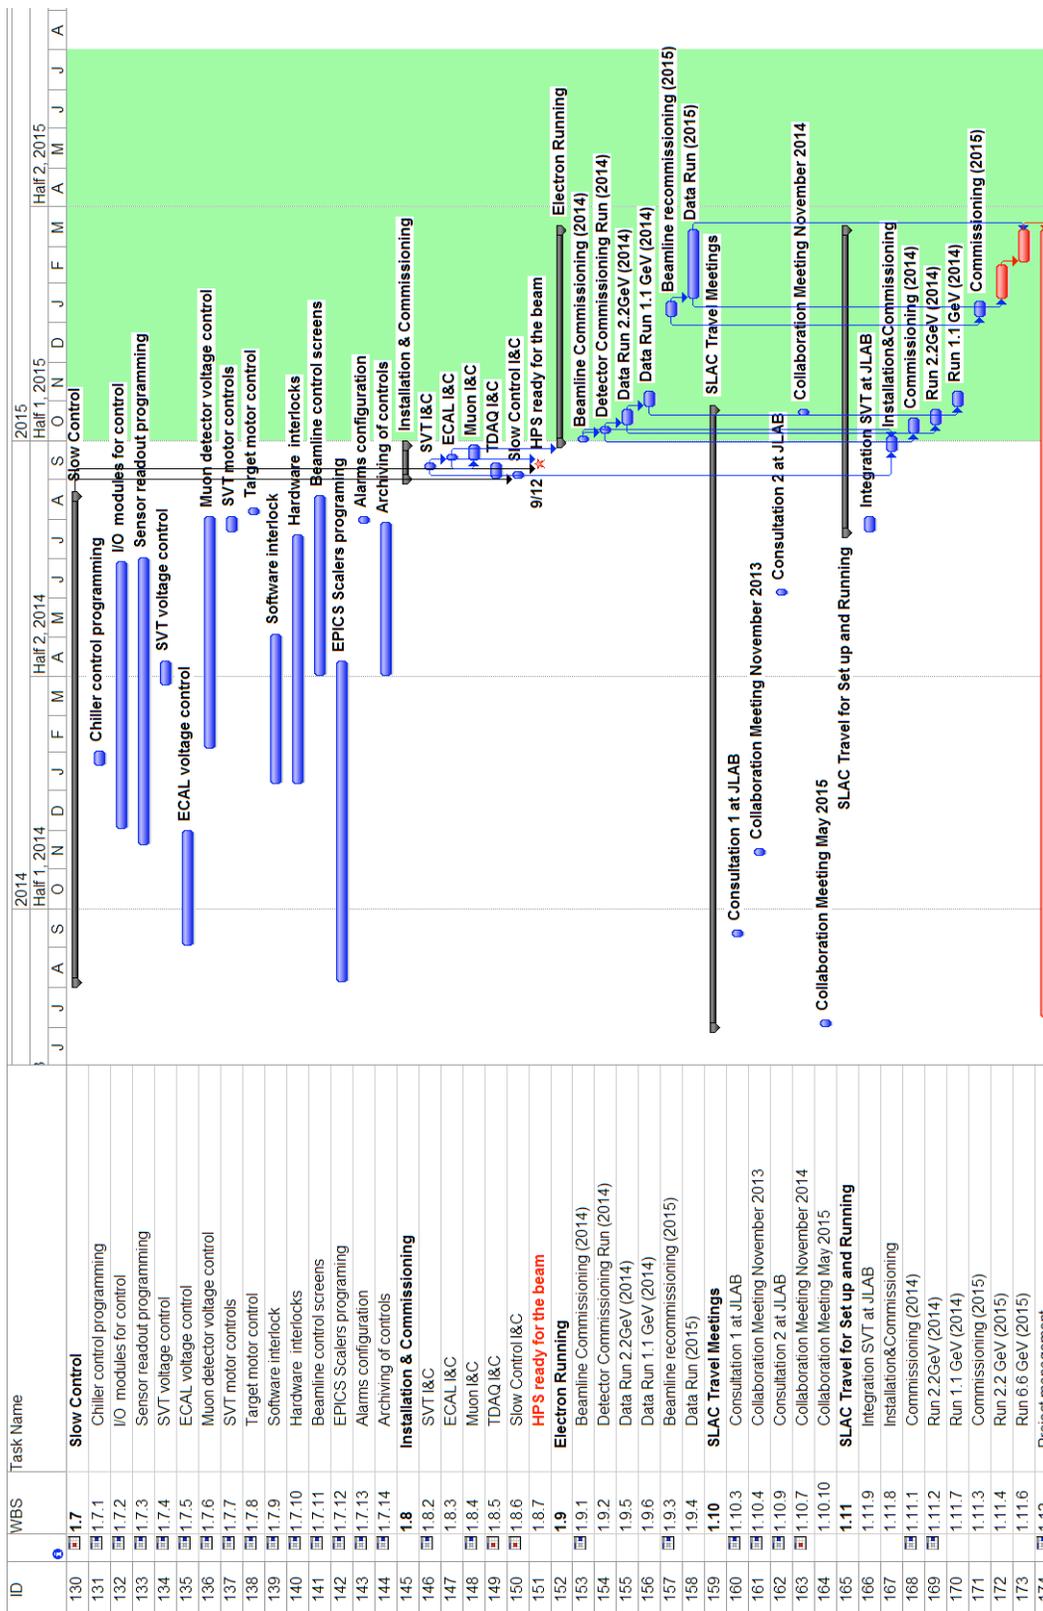


FIG. 86: HPS schedule.

A Simulation tools

The simulation tools play a critical role in simulating the background environment, optimizing the detector setup, and developing the trigger and reconstruction strategies. We use GEANT4 and EGS5 to simulate electromagnetic interactions. There is generally good agreement between these two codes. In particular, no inconsistencies have been found on secondary particle yields or energy spectra. However, we have found significant disagreements on the angular distributions in the multiple scattering, bremsstrahlung and pair production processes.

Multiple Scattering

EGS5 simulates the electron elastic scattering using the Molière theory [108] as formulated by Bethe. [109] It is based on a small angle approximation ($\theta \ll 1$ radian), and the angular distribution approaches asymptotically to Gaussian at small angles, and to Rutherford's Coulomb scattering function at large angles given by,

$$F(\theta) \sim \frac{1}{(1 - \cos\theta + \frac{\chi^2}{2})^2}. \quad (1)$$

Instead of using the complex and time consuming Molière's formula, GEANT4 uses two functions explicitly, Gaussian at small angles and the Rutherford function Eq. (1) at large angles with a requirement that these two functions and their first derivatives are joined continuously. GEANT4, however, uses a different power in the denominator in Eq. (1) which is close to 2 but not exactly equal to 2 and is dependent on the target material and thickness.

Several comparisons have been made in the angular distribution $F(\theta)$ in the differential cross section $d\sigma = F(\theta)d(\cos\theta)d\phi$ for 2.2 GeV electron scattering from 0.125% X_0 Tungsten target. The EGS5 simulation is compared with Molière's analytical formula in Figure 87(a), demonstrating a good agreement between EGS5 and the Molière theory. While the Molière theory is based on a small angle approximation, the multiple scattering theory developed by Gaudsmit and Saunderson is valid for any angle by means of an expansion in Legendre polynomials. [110] The validity of the small angle approximation is checked by compar-

ing the Molière integral with the Goudsmit-Saunderson theory as shown in Figure 87(b), demonstrating that the Molière theory is accurate in the angular region of the HPS detector.

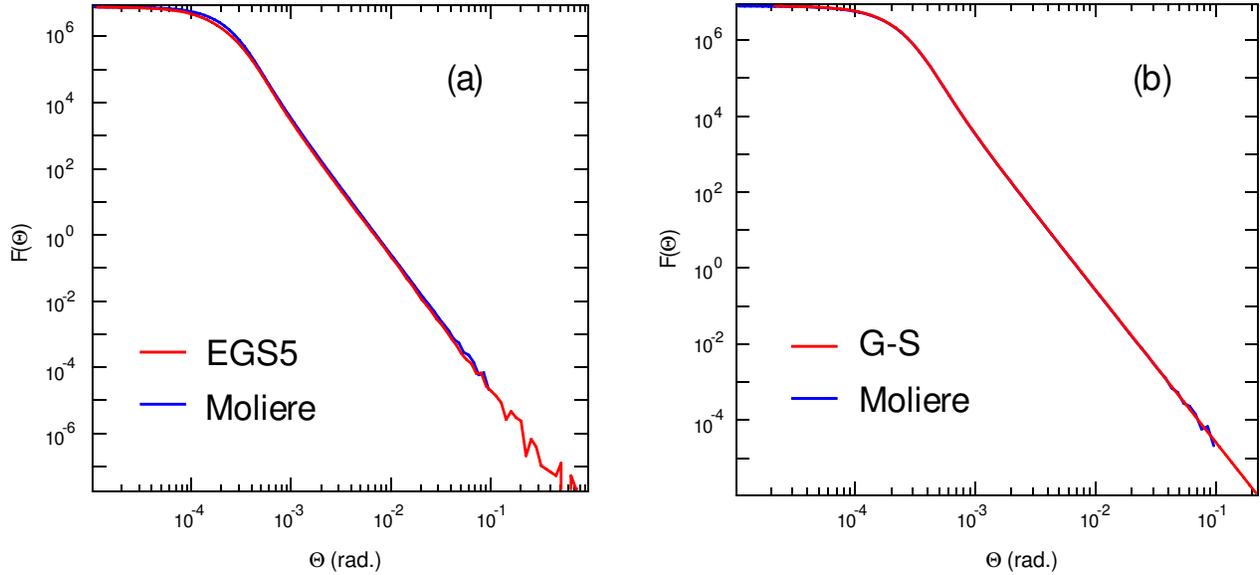


FIG. 87: (a) Molière vs. EGS5 (b) Molière vs. Goudsmit-Saunderson

Figure 88 shows the angular distribution comparison between the GEANT4 simulation and the Molière integral. GEANT4 is in good agreement with the Molière integral up to about 1 mrad, then it deviates at larger angles, predicting roughly twice the cross section at 15 mrad, where the HPS tracker sensor edge is located.

D. Attwood et al. measured 170 MeV muon angular distributions and compared with GEANT4 simulations and the Molière theory. [111] They concluded that GEANT4 simulation over-estimated the scattering tail by about a factor of two, and the data were consistent with the Molière theory. G. Shen et al. [112] and B. Gottschalk et al. [113] also showed that the Molière theory was consistent with the measurements on a wide variety of target materials.

Angular distributions in the bremsstrahlung and pair production processes

While GEANT4 and EGS5 are in good agreement in the production rates and the secondary particle energy spectra, there are significant differences in the angular distribution in the secondary particles. In EGS5, the angular distributions are sampled from the following differential cross section for the bremsstrahlung process, [114]

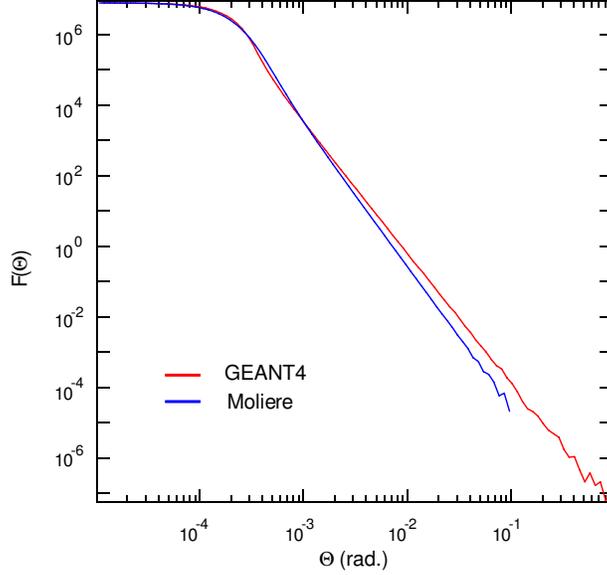


FIG. 88: Molière vs. GEANT4

$$d\sigma(k, \theta_\gamma) = \frac{4Z^2 r_0^2 dk}{137 k} y dy \left\{ \frac{16y^2 E}{(y^2 + 1)^4 E_0} - \frac{(E_0 + E)^2}{(y^2 + 1)^2 E_0^2} + \left\{ \frac{E_0^2 + E^2}{(y^2 + 1)^2 E_0^2} - \frac{4y^2 E}{(y^2 + 1)^4 E_0} \right\} \ln M(y) \right\},$$

where k photon energy, θ_γ photon polar angle, E_0 and E are initial and final electron energy, and

$$y = E_0 \theta_\gamma; \frac{1}{M(y)} = \left(\frac{k}{2E_0 E} \right)^2 + \left(\frac{Z^{1/3}}{111(y^2 + 1)} \right)^2,$$

and for the pair production process, [115]

$$\frac{d\sigma}{dE_\pm d\Omega_\pm} = \frac{2\alpha Z^2 r_0^2 E_\pm^2}{\pi k^3} \left\{ -\frac{(E_+ - E_-)^2}{(u^2 + 1)^2} - \frac{16u^2 E_+ E_-}{(u^2 + 1)^4} + \left\{ \frac{E_+^2 + E_-^2}{(u^2 + 1)^2} + \frac{4u^2 E_+ E_-}{(u^2 + 1)^4} \right\} \ln M(u) \right\},$$

where k photon energy, E_\pm e^\pm energy, θ_\pm e^\pm polar angle, and

$$u = E_\pm \theta_\pm; \frac{1}{M(u)} = \left(\frac{k}{2E_+ E_-} \right)^2 + \left(\frac{Z^{1/3}}{111(u^2 + 1)} \right)^2.$$

GEANT4 uses an approximate function to simulate the angular distributions in the bremsstrahlung and pair production processes given by

$$f(u) = C[ue^{-au} + due^{-3au}],$$

with $u = E_0\theta_\gamma$ for incident electron energy E_0 and the polar angle θ_γ of the bremsstrahlung photon, and $u = E_\pm\theta_\pm$ for the pair energy E_\pm and polar angle θ_\pm in the pair production. Since the production angle is typically $1/\gamma$, GEANT4's approximations are acceptable for most of the high energy detector simulations. However, GEANT4 simulations are inconsistent with the data in the following two cases in the HPS Test Run,

- 1) GEANT4 prediction on the bremsstrahlung photon angular distribution is too narrow, resulting in too few collimator scattering, and
- 2) GEANT4 prediction on the pair angular distribution is too narrow, resulting in too few Ecal trigger rates.

Conclusions

Because of these inaccuracies in GEANT4 the electromagnetic interactions in the target are simulated by EGS5, and all the particles that come out of the target are passed on to the HPS detector simulation system based on GEANT4.

B Test Run SVT Performance

SVT Calibration

In order to prepare the SVT for real physics data-taking, the SVT was calibrated. This involved the extraction of the mean baseline (pedestal), noise and gain for each of the 12,780 SVT channels. All measurements were made with the APV25 readout chips configured to their nominal operating points [98] and all sensors biased to 180 V. The APV25s were operated in “multi-peak” mode with six samples being readout per trigger. This, in turn, allowed for the extraction of the t_0 and amplitude of the signals being read out.

Figure 89 shows an intensity plot of the pedestals along with the readout noise as a function of channel number for a single hybrid. The noise was computed by taking the RMS of the Gaussian distributed pedestals for each of the channels and was observed to be

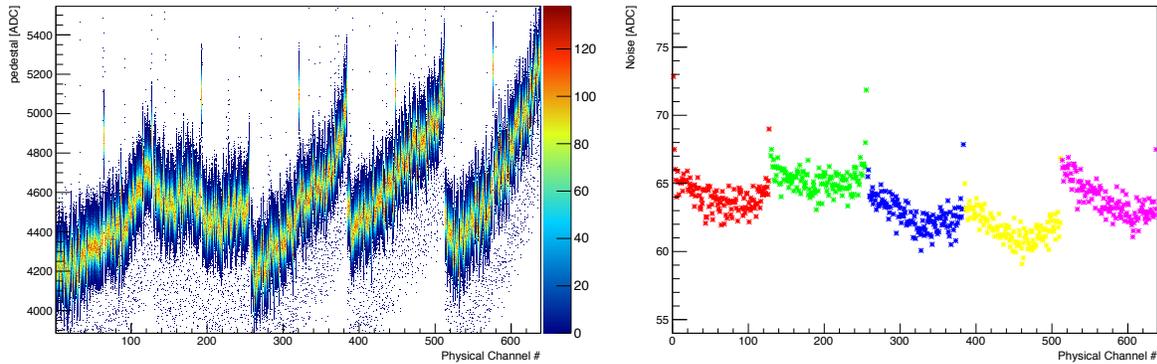


FIG. 89: The baseline distributions for a single hybrid are shown on the left plot. The overall baseline shifts were calibrated out. The measured noise is shown on the right with the different colors denoting a different APV25 readout chip. The large noise values at the chip edges have also been observed by the CMS collaboration and the cause is still under investigation.

consistently within 60 - 68 ADC counts ($\approx 750 - 850$ electrons). Figure 89 also demonstrates a feature common amongst all APV25 readout chips, mainly the large values of noise for channels lying near the chip edges. This was also observed by the CMS collaboration and the cause is still under investigation.

Another aspect which is crucial to full characterization of the SVT is an understanding of the signal response and the associated gain. Using the APV25 internal calibration circuitry, a known fixed charge was injected into all channels. This allowed for the accurate determination of the signal response along with its scaling with input charge. As shown on

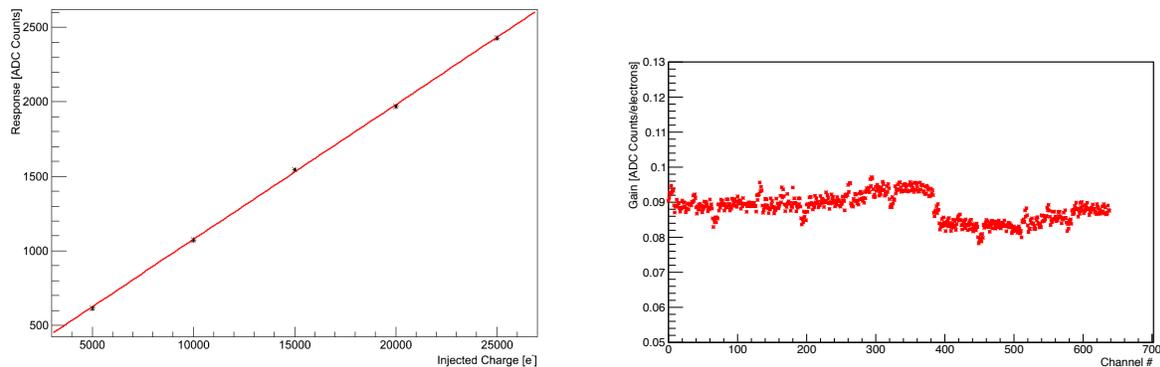


FIG. 90: Fit of the signal response as a function of charge over a range of 0 to 25,000 electrons is shown on the left plot. The response shows good linearity up to ≈ 3 MIPs. The gain uniformity, shown on the right, was within the expected range across all chips.

Figure 90, the response exhibits good linearity up to ≈ 3 MIPs while the gain uniformity was within the expected range across all chips on a hybrid.

SVT Hit Time Resolution

As discussed in Sec. 4.3, the APV25 multi-peak readout is crucial to the time stamping on hits in the SVT. This, in turn, allows the effective occupancies to be lowered for pattern recognition during electron running. Six samples of the APV25 shaper output for each trigger are fitted to an ideal $CR - RC$ function to extract the amplitude and hit time. The χ^2 distribution of these fits from test run data is as expected for four degrees of freedom. After clustering hits on a sensor, the hit time for each cluster is computed as the amplitude-weighted average of the channel hit times. Since we have no measurement of the “true” hit time, we study the overall SVT hit timing performance using the average of all cluster times in a track as the “track time,” and take the residual of the cluster time relative to that. The observed track time, shown in Fig. 52, has the expected amount of trigger jitter due to the readout clock and trigger system jitter. After correcting for offsets for each sensor (time-of-flight, clock phase) the RMS of the final residual distribution is roughly 2.4 ns for each sensor. Because the track time is calculated using the individual hit times, the hit time is positively correlated with the track time; thus the RMS of the residual is slightly smaller than the true time resolution. The standard deviation of this residual for n -hit tracks where all hits have the same time resolution is reduced by a factor of $\sqrt{(n-1)/n}$; since most of our tracks have 8 clusters, the true time resolution is 2.6 ns. This is somewhat worse than the ≈ 2 ns resolution expected in Sec. 6, but we believe this discrepancy is due to our fit function. Our pulse shape fit assumes an ideal CR-RC pulse shape; since the actual pulse shape has a slower rise time, there is a systematic pull on the hit time when a hit comes immediately before the APV clock time. This is visible in Figure 91 as a shift in the residual at certain values of track time. Work is in progress to use the actual pulse shape in time reconstruction; this should improve time resolution to that expected from previous studies. Reducing the APV25 pulse shaping time will also improve time resolution. In short, these results show that we can achieve time resolution adequate for pileup rejection during electron running.

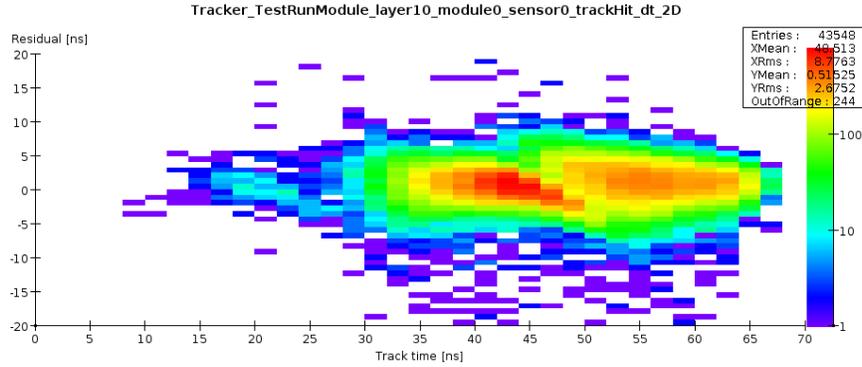


FIG. 91: Plot of the time residual for a representative sensor vs. the track time. The kinks in the horizontal band are caused by the fitter; without them the time resolution (measured by taking the projection of this histogram) would be better.

SVT Track Reconstruction

Clustered hits in Si planes within each layer are combined to form 2-dimensional “stereo hits” which are used by the tracking algorithm to form tracks. The determination of the probability that a stereo hit is formed, or hit efficiency, provides insight as to the performance of each of the SVT layers.

In order to obtain the hit efficiencies, tracks were fitted using only 4 of SVT layers. The resulting track was then extrapolated to the layer omitted from the fit. If the track was found to lie within the sensitive volume of the layer, a search for a stereo hit within the layer acceptance was conducted. The hit efficiency was then determined as

$$\epsilon_{\text{hit efficiency}} = \frac{\text{Tracks with hit on missing layer}}{\text{Tracks within layer acceptance}}.$$

The hit efficiencies per layer were calculated using all dedicated runs. It must be noted that those tracks found to intersect bad channels as well as those that lie on the sensor edges were excluded from the calculation. As can be seen from Figure 92, the average hit efficiency per layer was found to be $\approx 99\%$.

All events containing pairs of oppositely charged tracks were fit to a common vertex using a simple vertexing algorithm which searches for the distance of closest approach between the two tracks. The reconstructed vertex position along the beam axis for both data and Monte Carlo is shown on Figure 93. The positioning of the test run SVT layers was such that observation of pairs produced by the incident photon required both electrons to experience

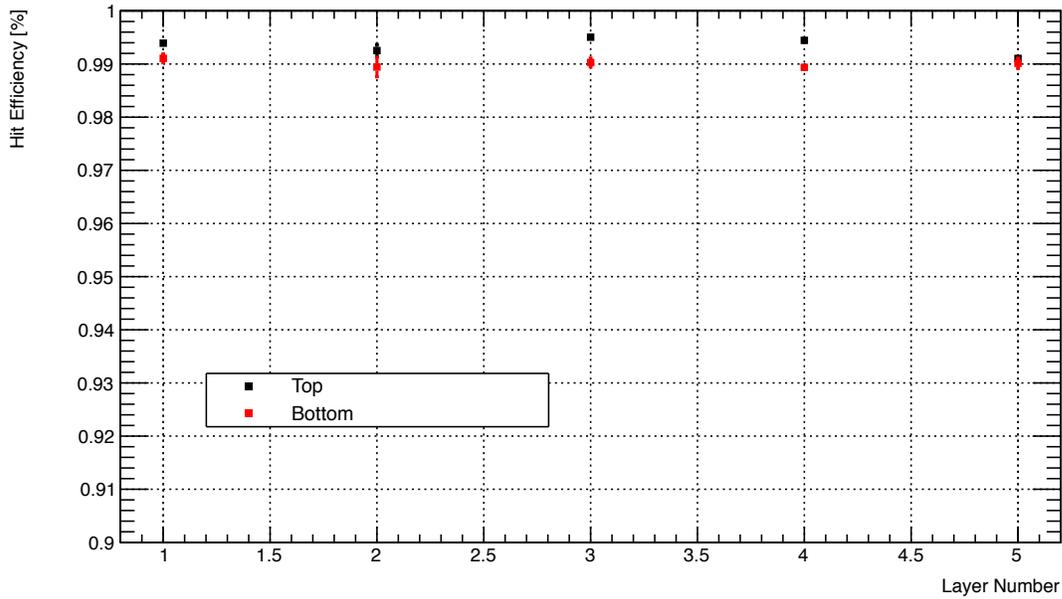


FIG. 92: Average hit efficiency, excluding known bad channels, from all dedicated photon runs as a function of layer number.

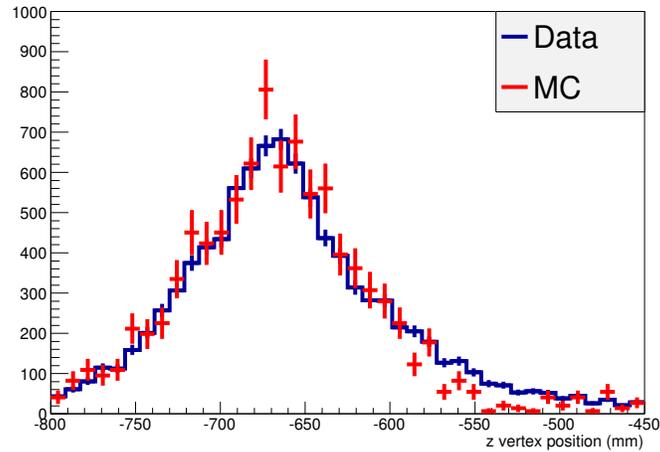


FIG. 93: The reconstructed vertex position along the beam axis for both data (blue) and Monte Carlo (red).

a hard scatter within the target. This results in a broadening of the Gaussian distributed reconstructed vertex position.

SVT Test Run DAQ Performance

The expected data rates and event sizes for each of the dedicated photon runs were estimated using a full simulation of the SVT DAQ and compared to observed values. As discussed in Section 5.1.3, the digitized samples from three hybrids were received by a single DPM. The DPM then required that at least three of the six samples exceeded a threshold of two times the noise level for that channel. An additional “pile-up” cut requiring that (sample 2 > sample 1) or (sample 3 > sample 2) was also applied. This was meant to eliminate hits arising from the falling edge of previous hits expected to occur when running at the highest occupancies.

All samples were placed into their own container along with the channel number, hybrid identifier, chip address and DPM identifier. An additional layer of encapsulation or bank was used to store all samples emerging from a single DPM along with the DPM identifier, the event number an error bit and hybrid temperatures. A diagram of the container along with the sizes of each of the elements is shown on Figure 94. Overall, the container overhead

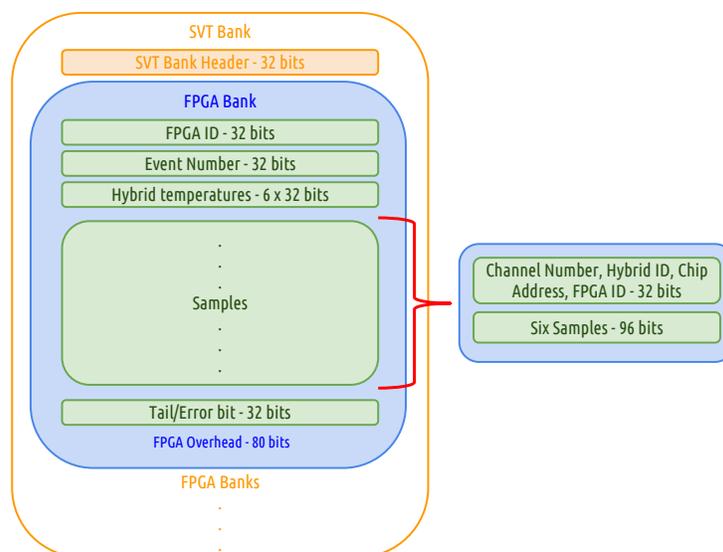


FIG. 94: SVT data format. Samples readout from three hybrids are processed by a single FPGA and are placed within a single container or FPGA bank. An additional layer of encapsulation is used to store all of the FPGA banks.

contributed a total of 326 bytes to an event with an additional 16 bytes per hit.

The observed occupancy expected for each of the converter thicknesses along with the

corresponding data rate are shown on Figure 95. The data rates observed during the test

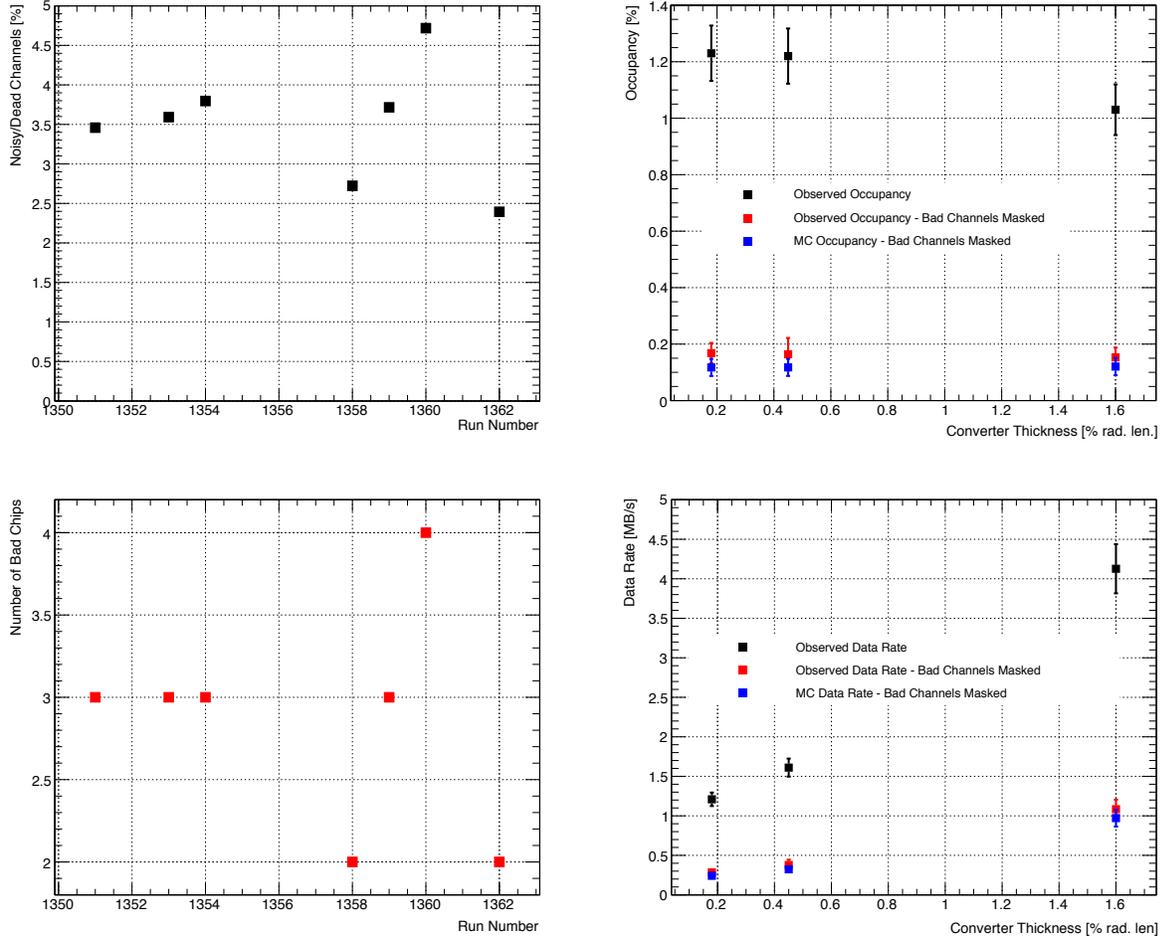


FIG. 95: The plots on the left shows the percentage of bad/noisy channels observed during each of the runs (top left) along with the number of noisy/misconfigured readout chips (bottom left). Most of the noisy channels present during runs were due to misconfigured chips. The plots on the right shows the occupancies (top right) and data rates (bottom right) observed for each of the targets used before and after noisy channels were masked. Once all noisy channels were masked, the observed data rates agree well with those predicted by the SVT readout simulation.

run were much higher than expected. This can be attributed to a known noisy sensor and a few noisy chips which appeared during certain runs as can be seen on Figure 95. The causes of both these issues are now well understood and will be resolved for future running.

A better comparison between simulated and observed data rates can be obtained by masking out all known noisy channels found during the commissioning of the SVT along with the noisy chips. The resulting occupancies and data rates are also shown on Figure 95. The simulated occupancies shown include a contribution of 0.02% (3 hits) due to noise and

the data rates are estimated using the trigger rates observed during each of the dedicated photon runs. As can be seen from the figure, the occupancies and data rates after most bad channels are masked are well in agreement with those predicted by Monte Carlo.

C Test Run ECal Calibration

The noise and pedestal of the readout chain are calibrated by running the ECal readout in a mode where the preamplifier output is sampled every 4 ns in a time window of 100 samples: by looking at a part of the window before the hit, we calibrate the readout channel.

For ECal analysis, cluster reconstruction was done using the algorithm described in [1]: build clusters around seed hits (hits above a “seed” energy threshold and with greater energy than any neighboring hits), and add all neighboring hits above an “add” energy threshold.

We calibrate gain of the individual ECal channels using the SVT measurement of track momentum. The ratio of cluster energy to track momentum is calculated both for Monte Carlo simulation and test run data at each point in the ECal, and we find the value of gain for each channel that brings the two into agreement. We use a formula to compute the “weighted E/p” for a crystal, representing the average E/p for clusters that include the crystal: $\frac{\sum_j w_{j,i}}{\sum_j \frac{E_j}{E_j} w_{j,i}}$. We disable all SVT and ECal channels in the simulation that were inoperable or noisy in the test run, so any efficiency or bias effects that affect the real data should be reflected in the simulation as well.

The calibrated gains are corrected by the ratio between the weighted E/p values from Monte Carlo and real data. The E/p in Monte Carlo data is also affected by the gain because the trigger thresholds change, so both Monte Carlo and data reconstruction are rerun with each iteration of gain calibration. It takes up to 4 iterations for the gains to stabilize; the final values are shown in Figure 96.

These gains can then be used to convert from ADC counts in a channel to energy deposited into that ECal crystal. The other information needed to find the energy of an incident particle is the sampling fraction—the ratio of energy read out from crystals to energy of an incident particle. The conventional sampling fraction—the fraction of incident energy that is deposited in crystals—is approximately 0.9 for our ECal, and less at edges. For our readout, there is additional energy lost because crystals under the readout threshold are not read

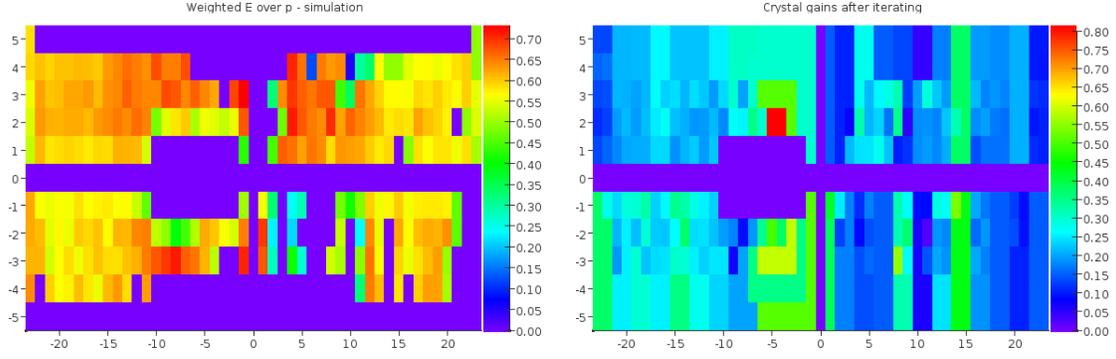


FIG. 96: Weighted E/p from Monte Carlo simulation (left), calibrated values of gain in units of MeV per ADC count (right).

out. The weighted E/p used in calibration (see Figure 96) is an approximate measurement of sampling fraction, but the sampling fraction is energy-dependent because of the effect of readout threshold. A full computation of sampling fraction can be done using simulation.

D Production and Decay of the A'

A' particles are generated in electron collisions on a fixed target by a process analogous to ordinary photon bremsstrahlung, see Figure 3. This can be reliably estimated in the Weizscker-Williams approximation (see [1-4]). When the incoming electron has energy E_0 , the differential cross-section to produce an A' of mass $m_{A'}$ with energy $E_{A'} \equiv xE_0$ is

$$\frac{d\sigma}{dx d\cos\theta_{A'}} \approx \frac{8Z^2\alpha^3\epsilon^2 E_0^2 x}{U^2} \tilde{\chi} \times \left[\left(1 - x + \frac{x^2}{2}\right) - \frac{x(1-x)m_{A'}^2 E_0^2 x \theta_{A'}^2}{U^2} \right] \quad (\text{D1})$$

where Z is the atomic number of the target atoms, $\alpha = 1/137$, is the angle in the lab frame between the emitted A' and the incoming electron,

$$U(x, \theta_{A'}) = E_0^2 x \theta_{A'}^2 + m_{A'}^2 \frac{1-x}{x} + m_e^2 x \quad (\text{D2})$$

is the virtuality of the intermediate electron in initial-state bremsstrahlung, and is the Weizsacker-Williams effective photon flux, with an overall factor of removed. The form of and its dependence on the A' mass, beam energy, and target nucleus are discussed in [107]. For HPS with $E_0 = 6.6$ GeV, we find $\tilde{\chi} \sim 7(4, 1)$ for $m_{A'} = 100$ (200, 500) MeV/ c^2 .

The above results are valid for

$$m_e \ll m_{A'} \ll E_0, \quad x\theta_{A'}^2 \ll 1. \quad (\text{D3})$$

For $m_e \ll m_{A'}$, the angular integration gives

$$\frac{d\sigma}{dx} \approx \frac{8Z^2\alpha^3\epsilon^2x}{m_{A'}^2} \left(1 + \frac{x^2}{3(1-x)}\right) \tilde{\chi}. \quad (\text{D4})$$

Assuming the A' decays into Standard Model particles rather than exotic, its boosted lifetime is

$$l_0 \equiv \gamma c\tau \approx \frac{0.8\text{cm}}{N_{eff}} \left(\frac{E_0}{10\text{GeV}}\right) \left(\frac{10^{-4}}{\epsilon}\right)^2 \left(\frac{100\text{MeV}}{m_{A'}}\right)^2, \quad (\text{D5})$$

where we have neglected phase-space corrections, and N_{eff} counts the number of available decay channels. If the A' couples only to electrons, then $N_{eff} = 1$. If the A' mixes kinetically with the photon, the $N_{eff} = 1$ for $m_{A'} < 2m_\mu$ and $2 + R(m_{A'})$ for $m_{A'} \geq 2m_\mu$, where [116]

$$R = \left. \frac{\sigma(e^+e^- \rightarrow \text{hadrons})}{\sigma(e^+e^- \rightarrow \mu^+\mu^-)} \right|_{E=m_{A'}}. \quad (\text{D6})$$

For the ranges of ϵ and $m_{A'}$ probed by this experiment, the mean decay length l_0 can be prompt or as large as tens of centimeters.

The total number of A' produced when N_e electrons scatter in a target of $T \ll 1$ radiation lengths is

$$N \sim N_e \frac{N_0 X_0}{A} T \frac{Z^2 \alpha^2 \epsilon^2}{m_{A'}^2} \tilde{\chi} \sim N_e C T \frac{\epsilon m_e^2}{m_{A'}^2}, \quad (\text{D7})$$

where X_0 is the radiation length of the target in g/cm², $N_0 \approx 6 \times 10^{23} \text{mole}^{-1}$ is Avogadro's number, and A is the target atomic mass in g/mole. The numerical factor $C \approx 5$ is logarithmically dependent on the choice of nucleus (at least in the range of masses where the form-factor is only slowly varying) and on $m_{A'}$, because, roughly, $X_0 \propto \frac{A}{Z^2}$ (see [7, 32, 107]).

For a Coulomb of incident electrons, the total number of A' s produced is given by

$$N \sim 10^5 \left(\frac{N_e}{1C}\right) \tilde{\chi} \left(\frac{T}{0.1}\right) \left(\frac{\epsilon}{10^{-4}}\right)^2 \left(\frac{100\text{MeV}}{m_{A'}}\right)^2. \quad (\text{D8})$$

-
- [1] A. Grillo *et al.* [HPS Collaboration], HPS Proposal to JLab PAC37 PR-11-006, http://www.jlab.org/exp_prog/PACpage/PAC37/proposals/Proposals/
- [2] A. Grillo *et al.* [HPS Collaboration], HPS Test Run Proposal to DOE, <https://confluence.slac.stanford.edu/display/hpsg/Project+Overview>
- [3] P. Hansson *et al.* [HPS Collaboration], HPS Update PAC 39, <https://confluence.slac.stanford.edu/display/hpsg/Project+Overview>
- [4] J. L. Hewett, H. Weerts, R. Brock, J. N. Butler, B. C. K. Casey, J. Collar, A. de Gouvea and R. Essig *et al.*, arXiv:1205.2671 [hep-ex].
- [5] M. Kamionkowski, S. M. Koushiappas and M. Kuhlen, Phys. Rev. D **81**, 043532 (2010) [arXiv:1001.3144 [astro-ph.GA]].
- [6] “Dark2012: Dark Forces at Accelerators”, <http://www.lnf.infn.it/conference/dark/index.php>
- [7] R. Essig, P. Schuster, N. Toro and B. Wojtsekhowski, JHEP **1102**, 009 (2011) [arXiv:1001.2557 [hep-ph]].
- [8] S. Abrahamyan *et al.* [APEX Collaboration], Phys. Rev. Lett. **107**, 191804 (2011) [arXiv:1108.2750 [hep-ex]].
- [9] H. Merkel *et al.* [A1 Collaboration], Phys. Rev. Lett. **106**, 251802 (2011) [arXiv:1101.4091 [nucl-ex]].
- [10] M. Freytsis, G. Ovanesyan and J. Thaler, JHEP **1001** (2010) 111 [arXiv:0909.2862 [hep-ph]].
- [11] B. Holdom, Phys. Lett. B **166** (1986) 196.
- [12] P. Galison and A. Manohar, Phys. Lett. B **136** (1984) 279.
- [13] M. Goodsell, J. Jaeckel, J. Redondo and A. Ringwald, JHEP **0911**, 027 (2009) [arXiv:0909.0515 [hep-ph]].
- [14] M. Cicoli, M. Goodsell, J. Jaeckel and A. Ringwald, JHEP **1107**, 114 (2011) [arXiv:1103.3705 [hep-th]].
- [15] M. Goodsell, S. Ramos-Sanchez and A. Ringwald, JHEP **1201**, 021 (2012) [arXiv:1110.6901 [hep-th]].
- [16] R. Essig, P. Schuster and N. Toro, Phys. Rev. D **80** (2009) 015003 [
- [17] M. Goodsell and A. Ringwald, Fortsch. Phys. **58**, 716 (2010) [arXiv:1002.1840 [hep-th]].
- [18] P. Candelas, G. T. Horowitz, A. Strominger and E. Witten, Nucl. Phys. B **258**, 46 (1985).

- [19] E. Witten, Nucl. Phys. B **268**, 79 (1986).
- [20] S. Andreas, M. D. Goodsell and A. Ringwald, arXiv:1109.2869 [hep-ph].
- [21] J. Jaeckel and A. Ringwald, Ann. Rev. Nucl. Part. Sci. **60**, 405 (2010) [arXiv:1002.0329 [hep-ph]].
- [22] P. Fayet, Phys. Rev. D **75** (2007) 115017 [arXiv:hep-ph/0702176].
- [23] C. Cheung, J. T. Ruderman, L. T. Wang and I. Yavin, Phys. Rev. D **80** (2009) 035008 [arXiv:0902.3246 [hep-ph]].
- [24] N. Arkani-Hamed and N. Weiner, JHEP **0812**, 104 (2008) [arXiv:0810.0714 [hep-ph]].
- [25] D. E. Morrissey, D. Poland and K. M. Zurek, JHEP **0907** (2009) 050 [arXiv:0904.2567 [hep-ph]].
- [26] S. Andreas, C. Niebuhr and A. Ringwald, Phys. Rev. D **86** (2012) 095019 [arXiv:1209.6083 [hep-ph]].
- [27] M. Davier and H. Nguyen Ngoc, Phys. Lett. B **229** (1989) 150.
- [28] A. Konaka, K. Imai, H. Kobayashi, A. Masaike, K. Miyake, T. Nakamura, N. Nagamine and N. Sasao *et al.*, Phys. Rev. Lett. **57** (1986) 659.
- [29] J. Blumlein and J. Brunner, “New Exclusion Limits for Dark Gauge Forces from Beam-Dump Data,” Phys. Lett. B **701** (2011) 155 [arXiv:1104.2747 [hep-ex]].
- [30] H. Davoudiasl, H. -S. Lee and W. J. Marciano, “Dark Side of Higgs Diphoton Decays and Muon G-2,” Phys. Rev. D **86** (2012) 095009 [arXiv:1208.2973 [hep-ph]].
- [31] M. Endo, K. Hamaguchi, and G. Mishima, [arXiv:1209.2558 [hep-ph]].
- [32] J. D. Bjorken, R. Essig, P. Schuster and N. Toro, Phys. Rev. D **80**, 075018 (2009) [arXiv:0906.0580 [hep-ph]].
- [33] J. D. Bjorken *et al.*, Phys. Rev. D **38** (1988) 3375.
- [34] E. M. Riordan *et al.*, Phys. Rev. Lett. **59** (1987) 755.
- [35] A. Bross, M. Crisler, S. H. Pordes, J. Volk, S. Errede and J. Wrbanek, Phys. Rev. Lett. **67** (1991) 2942.
- [36] M. Pospelov, Phys. Rev. D **80** (2009) 095002 [arXiv:0811.1030 [hep-ph]].
- [37] KLOE-2 Collaboration, Phys. Lett. B **706** (2012) 251 [arXiv:1110.0411 [hep-ex]].
- [38] M. Reece and L. T. Wang, JHEP **0907** (2009) 051 [arXiv:0904.1743 [hep-ph]].
- [39] B. Aubert *et al.* [BABAR Collaboration], Phys. Rev. Lett. **103**, 081803 (2009) [arXiv:0905.4539 [hep-ex]].

- [40] J. B. Dent, F. Ferrer and L. M. Krauss, arXiv:1201.2683 [astro-ph.CO].
- [41] N. Arkani-Hamed, D. P. Finkbeiner, T. R. Slatyer and N. Weiner, Phys. Rev. D **79** (2009) 015014 [arXiv:0810.0713 [hep-ph]].
- [42] M. Pospelov and A. Ritz, Phys. Lett. B **671** (2009) 391 [arXiv:0810.1502 [hep-ph]].
- [43] E. Komatsu *et al.* [WMAP Collaboration], Astrophys. J. Suppl. **192**, 18 (2011) [arXiv:1001.4538 [astro-ph.CO]]; D. J. Eisenstein *et al.* [SDSS Collaboration], Astrophys. J. **633**, 560 (2005) [astro-ph/0501171]; S. Perlmutter *et al.* [Supernova Cosmology Project Collaboration], Astrophys. J. **517**, 565 (1999) [astro-ph/9812133]; A. G. Riess *et al.* [Supernova Search Team Collaboration], Astron. J. **116**, 1009 (1998) [astro-ph/9805201]; M. Kowalski *et al.* [Supernova Cosmology Project Collaboration], Astrophys. J. **686**, 749 (2008) [arXiv:0804.4142 [astro-ph]].
- [44] O. Adriani *et al.* [PAMELA Collaboration], Nature **458**, 607 (2009) [arXiv:0810.4995 [astro-ph]].
- [45] M. Ackermann *et al.* [Fermi LAT Collaboration], Phys. Rev. D **82**, 092004 (2010) [arXiv:1008.3999 [astro-ph.HE]].
- [46] J. Chang, J. H. Adams, H. S. Ahn, G. L. Bashindzhagyan, M. Christl, O. Ganel, T. G. Guzik and J. Isbert *et al.*, Nature **456**, 362 (2008).
- [47] F. Aharonian *et al.* [H.E.S.S. Collaboration], Phys. Rev. Lett. **101**, 261104 (2008) [arXiv:0811.3894 [astro-ph]].
- [48] F. Aharonian *et al.* [H.E.S.S. Collaboration], Astron. Astrophys. **508**, 561 (2009) [arXiv:0905.0105 [astro-ph.HE]].
- [49] O. Adriani *et al.* [PAMELA Collaboration], Phys. Rev. Lett. **106**, 201101 (2011) [arXiv:1103.2880 [astro-ph.HE]].
- [50] M. Ackermann *et al.* [The Fermi LAT Collaboration], Phys. Rev. Lett. **108**, 011103 (2012) [arXiv:1109.0521 [astro-ph.HE]].
- [51] M. Cirelli, M. Kadastik, M. Raidal and A. Strumia, Nucl. Phys. B **813**, 1 (2009) [arXiv:0809.2409 [hep-ph]].
- [52] I. Cholis, D. P. Finkbeiner, L. Goodenough and N. Weiner, JCAP **0912**, 007 (2009) [arXiv:0810.5344 [astro-ph]].
- [53] I. Cholis, G. Dobler, D. P. Finkbeiner, L. Goodenough and N. Weiner, Phys. Rev. D **80**, 123518 (2009) [arXiv:0811.3641 [astro-ph]].

- [54] D. P. Finkbeiner, L. Goodenough, T. R. Slatyer, M. Vogelsberger and N. Weiner, *JCAP* **1105**, 002 (2011) [arXiv:1011.3082 [hep-ph]].
- [55] T. R. Slatyer, N. Toro and N. Weiner, arXiv:1107.3546 [hep-ph].
- [56] O. Adriani, G. C. Barbarino, G. A. Bazilevskaya, R. Bellotti, M. Boezio, E. A. Bogomolov, L. Bonechi and M. Bongi *et al.*, *Phys. Rev. Lett.* **102**, 051101 (2009) [arXiv:0810.4994 [astro-ph]].
- [57] M. Ackermann *et al.* [Fermi-LAT Collaboration], *Phys. Rev. Lett.* **107**, 241302 (2011) [arXiv:1108.3546 [astro-ph.HE]].
- [58] M. D. Kistler and J. M. Siegal-Gaskins, *Phys. Rev. D* **81**, 103521 (2010) [arXiv:0909.0519 [astro-ph.HE]]. K. N. Abazajian, S. Blanchet and J. P. Harding, *Phys. Rev. D* **85**, 043509 (2012) [arXiv:1011.5090 [hep-ph]].
- [59] M. Papucci and A. Strumia, *JCAP* **1003**, 014 (2010) [arXiv:0912.0742 [hep-ph]].
- [60] G. Hutsi, A. Hektor and M. Raidal, *JCAP* **1007**, 008 (2010) [arXiv:1004.2036 [astro-ph.HE]].
- [61] J. Zavala, M. Vogelsberger, T. R. Slatyer, A. Loeb and V. Springel, *Phys. Rev. D* **83**, 123513 (2011) [arXiv:1103.0776 [astro-ph.CO]].
- [62] X. Huang, G. Vertongen and C. Weniger, *JCAP* **1201**, 042 (2012) [arXiv:1110.1529 [hep-ph]].
- [63] S. Galli, F. Iocco, G. Bertone and A. Melchiorri, *Phys. Rev. D* **80**, 023505 (2009) [arXiv:0905.0003 [astro-ph.CO]]; T. R. Slatyer, N. Padmanabhan and D. P. Finkbeiner, *Phys. Rev. D* **80**, 043526 (2009) [arXiv:0906.1197 [astro-ph.CO]]; S. Galli, F. Iocco, G. Bertone and A. Melchiorri, *Phys. Rev. D* **84**, 027302 (2011) [arXiv:1106.1528 [astro-ph.CO]]; D. P. Finkbeiner, S. Galli, T. Lin and T. R. Slatyer, *Phys. Rev. D* **85**, 043522 (2012) [arXiv:1109.6322 [astro-ph.CO]].
- [64] J. L. Feng, M. Kaplinghat and H. -B. Yu, *Phys. Rev. Lett.* **104**, 151301 (2010) [arXiv:0911.0422 [hep-ph]].
- [65] M. R. Buckley and P. J. Fox, *Phys. Rev. D* **81**, 083522 (2010) [arXiv:0911.3898 [hep-ph]].
- [66] J. T. Ruderman and T. Volansky, “Decaying into the Hidden Sector,” *JHEP* **1002** (2010) 024 [arXiv:0908.1570 [hep-ph]].
- [67] R. Essig, J. Kaplan, P. Schuster and N. Toro, “On the Origin of Light Dark Matter Species,” Submitted to: Submitted to Physical Review D [arXiv:1004.0691 [hep-ph]].
- [68] R. Bernabei *et al.* [DAMA and LIBRA Collaborations], *Eur. Phys. J. C* **67**, 39 (2010) [arXiv:1002.1028 [astro-ph.GA]].

- [69] C. E. Aalseth *et al.* [CoGeNT Collaboration], Phys. Rev. Lett. **106**, 131301 (2011) [arXiv:1002.4703 [astro-ph.CO]].
- [70] C. E. Aalseth, P. S. Barbeau, J. Colaresi, J. I. Collar, J. Diaz Leon, J. E. Fast, N. Fields and T. W. Hossbach *et al.*, Phys. Rev. Lett. **107**, 141301 (2011) [arXiv:1106.0650 [astro-ph.CO]].
- [71] G. Angloher, M. Bauer, I. Bavykina, A. Bento, C. Bucci, C. Ciemniak, G. Deuter and F. von Feilitzsch *et al.*, arXiv:1109.0702 [astro-ph.CO].
- [72] Z. Ahmed *et al.* [CDMS-II Collaboration], Phys. Rev. Lett. **106**, 131302 (2011) [arXiv:1011.2482 [astro-ph.CO]]. Z. Ahmed *et al.* [CDMS Collaboration], arXiv:1203.1309 [astro-ph.CO].
- [73] J. Angle *et al.* [XENON10 Collaboration], Phys. Rev. Lett. **107**, 051301 (2011) [arXiv:1104.3088 [astro-ph.CO]].
- [74] E. Aprile *et al.* [XENON100 Collaboration], Phys. Rev. Lett. **107**, 131302 (2011) [arXiv:1104.2549 [astro-ph.CO]].
- [75] C. Kelso, D. Hooper and M. R. Buckley, Phys. Rev. D **85**, 043515 (2012) [arXiv:1110.5338 [astro-ph.CO]].
- [76] E. Aprile *et al.* [XENON100 Collaboration], “Implications on Inelastic Dark Matter from 100 Live Days of Xenon100 Data,” Phys. Rev. D **84** (2011) 061101 [arXiv:1104.3121 [astro-ph.CO]].
- [77] T. Beranek, and M. Vanderhaeghen, [arXiv:1209.4561 [hep-ph]].
- [78] S. Andreas, C. Niebuhr, and A. Ringwald, [arXiv:1209.6083 [hep-ph]].
- [79] M. Deutsch, Phys. Rev. **82**, 455 (1951).
- [80] J. I. Friedman and V. L. Telegdi, Phys. Rev. **105**, 1681 (1957).
- [81] V. W. Hughes, D. W. McColm, K. Ziock and R. Prepost, Phys. Rev. Lett. **5**, 63 (1960).
- [82] E. Holvik and H. A. Olsen, Phys. Rev. D **35**, 2124 (1987).
- [83] N. Arteaga-Romero, C. Carimalo and V. G. Serbo, Phys. Rev. A **62**, 032501 (2000) [hep-ph/0001278].
- [84] S. J. Brodsky and R. F. Lebed, Phys. Rev. Lett. **102**, 213401 (2009) [arXiv:0904.2225 [hep-ph]].
- [85] S. M. Bilenky, V. H. Nguyen, L. L. Nemenov and F. G. Tkebuchava, Yad. Fiz. **10**, 812 (1969).
- [86] V. W. Hughes and B. Maglic, Bull. Am. Phys. Soc. **16**, 65 (1971).
- [87] J. Malenfant, Phys. Rev. D **36**, 863 (1987).

- [88] S. G. Karshenboim, U. D. Jentschura, V. G. Ivanov and G. Soff, Phys. Lett. B **424**, 397 (1998).
- [89] D. A. Owen and W. W. Repko, Phys. Rev. A **5**, 1570 (1972).
- [90] U. D. Jentschura, G. Soff, V. G. Ivanov and S. G. Karshenboim, hep-ph/9706401.
- [91] U. D. Jentschura, G. Soff, V. G. Ivanov and S. G. Karshenboim, Phys. Rev. A **56**, 4483 (1997) [physics/9706026].
- [92] S. G. Karshenboim, V. G. Ivanov, U. D. Jentschura and G. Soff, J. Exp. Theor. Phys. **86**, 226 (1998) [Zh. Eksp. Teor. Fiz. **113**, 409 (1998)].
- [93] G. W. Bennett *et al.* [Muon G-2 Collaboration], Phys. Rev. D **73**, 072003 (2006) [hep-ex/0602035].
- [94] R. Pohl, A. Antognini, F. Nez, F. D. Amaro, F. Biraben, J. M. R. Cardoso, D. S. Covita and A. Dax *et al.*, Nature **466**, 213 (2010).
- [95] A. Banburski and P. Schuster, Phys. Rev. D **86**, 093007 (2012) [arXiv:1206.3961 [hep-ph]].
- [96] M. J. Strassler and K. M. Zurek, Phys. Lett. B **651**, 374 (2007) [hep-ph/0604261].
- [97] M. Borland, A Flexible SDDS-Compliant Code for Accelerator Simulation, ANL, Argonne, IL 60439, USA
- [98] L. Jones, APV25-S1: User guide version 2.2, RAL Microelectronics Design Group, 2011.
- [99] M. Raymond *et al.*, APV25 production testing and quality assurance, 8th Workshop on Electronics for LHC Experiments, Colmar, France, 9-13 Sep 2002
- [100] D.H.H. Hoffmann *et al.*, Z. Physik **A293**, 187 (1979)
- [101] J.H. Hubbell *et al.*, J. Phys. Chem. Ref. Data, Vol. **23**, 339 (1994)
- [102] H.-J. Möhring, Nucl. Instrum. Methods **A292**, 482 (1990)
- [103] Y.M. Budnev *et al.*, Phys. Rep. **15C**, 181 (1975)
- [104] P. Billoir, R. Fruhwirth, and M. Regler, Nucl. Instr. And Meth. **A241**, 115 (1985).
- [105] P. Billoir and S. Qian, Nucl. Instr. And Meth. **A311**, 139 (1991).
- [106] J. Alwall, P. Demin, S. de Visscher, R. Frederix, M. Herquet, F. Maltoni, T. Plehn and D. L. Rainwater *et al.*, JHEP **0709**, 028 (2007) [arXiv:0706.2334 [hep-ph]].
- [107] K. J. Kim and Y. -S. Tsai, Phys. Rev. D **8**, 3109 (1973).
- [108] G. Molière, Z. Naturforsch. **3a**, 78 (1948)
- [109] H. A. Bethe, Phys. Rev. **89**, 1256 (1953)
- [110] S. A. Goudsmit and J. L. Saunderson, Phys. Rev. **57**, 24 and **58**, 36 (1940); K. Okei and T.

Nakatsuka, Proceedings of the 17th EGS User's Meeting.

- [111] D. Attwood et al., Nucl. Instrum. Methods **B251**, 41 (2006)
- [112] G. Shen et al., Phys. Rev. **20**, 1584 (1979)
- [113] B. Gottschalk et al., Nucl. Instrum. Meth. **B74**, 467 (1993)
- [114] H.W. Koch and J.W. Motz, Rev. Mod. Phys. **31**, 920 (1959)
- [115] J.W. Motz, H. A. Olsen, and H.W. Koch, Rev. Mod. Phys. **41**, 581 (1969)
- [116] J. Beringer et al. (Particle Data Group), Phys. Rev. D86, 010001 (2012)