

Searching for Prompt and Long-Lived Dark Photons in Electro-Produced e^+e^- Pairs with the Heavy Photon Search Experiment at JLab

P. H. Adrian,¹ N. A. Baltzell,² M. Battaglieri,³ M. Bondí,⁴ S. Boyarinov,² C. Bravo,^{1,*} S. Bueltmann,⁵ P. Butti,¹ V. D. Burkert,² D. Calvo,⁶ T. Cao,⁷ M. Carpinelli,^{8,9} A. Celentano,³ G. Charles,¹⁰ L. Colaneri,^{11,12} W. Cooper,¹³ C. Cuevas,² A. D'Angelo,^{11,12} N. Dashyan,¹⁴ M. De Napoli,⁴ R. De Vita,³ A. Deur,² M. Diamond,¹ R. Dupre,¹⁰ H. Eginyan,² L. Elouadrhiri,² R. Essig,¹⁵ V. Fadeyev,¹⁶ C. Field,¹ A. Filippi,⁶ A. Freyberger,² M. Garçon,¹⁷ N. Gevorgyan,¹⁴ F. X. Girod,² N. Graf,¹ M. Graham,^{1,†} K. A. Griffioen,¹⁸ A. Grillo,¹⁶ M. Guidal,¹⁰ R. Herbst,¹ M. Holtrop,⁷ J. Jaros,¹ R. P. Johnson,¹⁶ G. Kalicy,⁵ M. Khandaker,¹⁹ V. Kubarovsky,² E. Leonora,⁴ K. Livingston,²⁰ L. Marsicano,³ T. Maruyama,¹ S. McCarty,⁷ J. McCormick,¹ B. McKinnon,²⁰ K. Moffeit,²⁰ O. Moreno,^{1,16} C. Munoz Camacho,¹⁰ T. Nelson,¹ S. Niccolai,¹⁰ A. Odian,¹ M. Oriunno,¹ M. Osipenko,³ R. Paremuzyan,^{2,7} S. Paul,¹⁸ N. Randazzo,⁴ B. Raydo,² B. Reese,¹ A. Rizzo,^{11,12} P. Schuster,^{1,21} Y. G. Sharabian,² G. Simi,^{22,23} A. Simonyan,¹⁰ V. Sipala,^{24,8} A. Spellman,¹⁶ D. Sokhan,²⁰ M. Solt,¹ S. Stepanyan,² H. Szumila-Vance,^{2,5} N. Toro,^{1,21} S. Uemura,¹ M. Ungaro,² H. Voskanyan,¹⁴ L. B. Weinstein,⁵ B. Wojtsekhowski,² and B. Yale⁷

¹SLAC National Accelerator Laboratory, Stanford University, Stanford, CA 94309, USA

²Thomas Jefferson National Accelerator Facility, Newport News, Virginia 23606, USA

³INFN, Sezione di Genova, 16146 Genova, Italy

⁴INFN, Sezione di Catania, 95123 Catania, Italy

⁵Old Dominion University, Norfolk, Virginia 23529, USA

⁶INFN, Sezione di Torino, 10125 Torino, Italy

⁷University of New Hampshire, Durham, New Hampshire 03824, USA

⁸INFN, Laboratori Nazionali del Sud, 95123 Catania, Italy

⁹University di Milano Bicocca, 20126 Milano, Italy

¹⁰Université Paris-Saclay, CNRS, IJCLab, 91405, Orsay, France

¹¹Università di Roma Tor Vergata, 00133 Rome Italy

¹²INFN, Sezione di Roma Tor Vergata, 00133 Rome, Italy

¹³Fermi National Accelerator Laboratory, Batavia, IL 60510, USA

¹⁴Yerevan Physics Institute, 375036 Yerevan, Armenia

¹⁵C. N. Yang Institute for Theoretical Physics, Stony Brook University, Stony Brook, NY 11794, USA

¹⁶Santa Cruz Institute for Particle Physics, University of California, Santa Cruz, CA 95064, USA

¹⁷IRFU, CEA, Université Paris-Saclay, F-91190 Gif-sur-Yvette, France

¹⁸College of William & Mary, Williamsburg, Virginia 23187, USA

¹⁹Idaho State University, Pocatello, ID, 83209, USA

²⁰University of Glasgow, Glasgow G12 8QQ, United Kingdom

²¹Perimeter Institute, Ontario, Canada N2L 2Y5

²²Università di Padova, 35122 Padova, Italy

²³INFN, Sezione di Padova, 16146 Padova, Italy

²⁴Università di Sassari, 07100 Sassari, Italy

(Dated: January 12, 2023)

The Heavy Photon Search experiment (HPS) at the Thomas Jefferson National Accelerator Facility searches for electro-produced dark photons. We report results from the 2016 Engineering Run consisting of $10\,608\,\text{nb}^{-1}$ of data for both the prompt and displaced vertex searches. A search for a prompt resonance in the e^+e^- invariant mass distribution between 39 and 179 MeV showed no evidence of dark photons above the large QED background, limiting the coupling of $\epsilon^2 \gtrsim 10^{-5}$, in agreement with previous searches. The search for displaced vertices showed no evidence of excess signal over background in the masses between 60 and 150 MeV, but had insufficient luminosity to limit canonical heavy photon production. This is the first displaced vertex search result published by HPS. HPS has taken high-luminosity data runs in 2019 and 2021 that will explore new dark photon phase space.

PACS numbers: xxx

I. INTRODUCTION

Interest in searching for new, sub-GeV mediators with weak couplings to ordinary matter has grown exponentially in recent years, where such forces could play an

* Corresponding Author: bravo@slac.stanford.edu

† Corresponding Author: mgraham@slac.stanford.edu

essential role in production of sub-GeV dark matter in the early universe [1–4]. Additionally, and more generally, such experiments are a key complement to searches for new physics at high energies where new weakly coupled physics at low mass scales can be difficult to identify. Heavy photons, also known as “hidden-sector” or “dark” photons, are a benchmark example of such a mediator that also appears in many scenarios for physics beyond the Standard Model. Kinetic mixing of the heavy photon with the Standard Model photon through radiative loops of massive particles generates a weak coupling of the heavy photon to electrically charged particles [5–7]. As a result, heavy photons would be radiated by energetic electrons passing through a target in a process analogous to bremsstrahlung, but at parametrically lower rate, and can also decay to electron-positron pairs [8]. While our search is focused on heavy photons, it is also sensitive to dark forces with vector, axial-vector, scalar, or pseudo-scalar couplings to matter which will have similar signatures and could also be produced in our experiment.

The Heavy Photon Search Experiment (HPS) at the Thomas Jefferson National Accelerator Facility (JLab) in Newport News, Virginia, searches for heavy photons and other new force carriers that are produced via electro-production and decay to electron-positron pairs [8]. Note that if direct decays to dark matter (or other dark-sector particles) are kinematically allowed, those decays are expected to dominate over the decay to SM particles, so HPS is only sensitive to heavy photons with less than twice the mass of the dark matter. Experimental signatures are either a resonance in the invariant electron-positron mass distribution or displaced decay vertices with a particular invariant mass, depending on the heavy photon mass and the strength of its coupling to electrons. Over the past decade, searches for dark photons have been conducted over large regions of the dark photon mass/coupling parameter space, but much of that parameter space, including territory favored by thermal dark matter production in the early universe, remains unexplored and accessible to HPS [4]. Evidence for a dark force could be the first compelling evidence for a hidden sector and lead to identifying the dark matter.

For concreteness, we focus our discussion on the heavy photon, denoted A' . The A' is the mediator of a spontaneously broken “hidden” $U(1)'$ gauge symmetry. The A' interacts with SM particles through kinetic mixing with the SM $U(1)_Y$ (hypercharge) gauge boson, resulting at low energies in the effective Lagrangian density

$$\mathcal{L} \supset -\frac{\epsilon}{2} F'_{\mu\nu} F^{\mu\nu}, \quad (1)$$

where ϵ denotes the strength of the kinetic mixing, $F'_{\mu\nu} = \partial_\mu A'_\nu - \partial_\nu A'_\mu$ is the $U(1)'$ field strength tensor, and similarly $F^{\mu\nu}$ denotes the field strength of the SM photon. This A' -photon mixing allows heavy photons to be produced in interactions involving electromagnetically charged particles and, if sufficiently massive, to decay into pairs of charged particles like electron-positron pairs

or muon-antimuon pairs, or to hidden-sector states. The value of ϵ and the A' mass ($m_{A'}$) generated in the fundamental theory naturally fall into the sensitivity range of HPS in certain model scenarios [9–14].




FIG. 1. Diagram of the A' production off the tungsten target and decay to an e^+e^- pair.

The HPS experiment, which utilizes the Continuous Electron Beam Accelerator Facility (CEBAF) at JLab, can explore a wide range of heavy photon masses ($m_{A'} \sim 20 - 220$ MeV) and couplings ($\epsilon^2 \sim 10^{-10} - 10^{-6}$) using both resonance search and separated vertex strategies. In this paper, results from both strategies are reported, using the data from the 2016 Engineering Run which employed an electron beam with a current of 200 nA and an energy of $E_{\text{beam}} = 2.3$ GeV incident on a thin (4 μm) tungsten target, and integrating a luminosity of 10608 nb^{-1} . We have previously reported on the resonance search from our 2015 Engineering Run at 1.03 GeV [15]. In HPS, the A' s would be electro-produced on the target nuclei, and would subsequently decay to electron-positron pairs, shown in Figure 1. A charged particle spectrometer, triggered by an electromagnetic calorimeter, measures the momenta and trajectories of the pair, from which its invariant mass and decay position can be reconstructed. The A' decay length in the laboratory frame is given by

$$\ell_0 \simeq \frac{1.8 \text{ mm}}{N_{\text{eff}}} \left(\frac{E_{\text{beam}}}{2.3 \text{ GeV}} \right) \left(\frac{10^{-4}}{\epsilon} \right)^2 \left(\frac{100 \text{ MeV}}{m_{A'}} \right)^2, \quad (2)$$

where N_{eff} is the number of decay channels kinematically accessible ($N_{\text{eff}} = 1$ for HPS searches below the di-muon threshold) [8]. For larger couplings, the A' is essentially prompt and would appear as a narrow resonance, with a width set by the experimental resolution, on top of a broad distribution of background events from ordinary quantum electrodynamic (QED) processes. At smaller couplings the A' lifetime is long enough to give rise to secondary decay vertices, which can be distinguished from the prompt QED background, providing a second signature for heavy photon production.

The HPS experiment records copious QED trident production, as well as wide-angle bremsstrahlung production with subsequent conversion in the target or detector material, both of which produce the same final state. While these processes constitute physics backgrounds for the heavy photon search, they also enable important experimental checks and provide an experimental determination of our sensitivity, since the expected heavy photon production can be related to the measured trident production. The experimental mass resolution impacts the reach and is a critical input to the fits of the mass spectrum and to setting the width of the mass bins for the vertex search. It is calibrated directly from the data by measuring the invariant mass of Møller pairs, which have a unique invariant mass for any given incident electron energy. Similarly, the measured decay length distribution of the prompt trident signal provides a critical estimate of the decay length resolution.

The outline of the rest of the paper is as follows. Section II describes the beamline, target, and detector used by the HPS experiment. Section III gives an overview of the common elements of the data analysis described in the paper. Sections IV and V describe in detail the resonance search and displayed vertex search, respectively. Finally, Section VI gives a summary of the paper.

II. DETECTOR OVERVIEW

While the rejection of QED backgrounds motivates the best possible resolutions for e^+e^- mass and vertex position, the kinematic characteristics of the signal and beam backgrounds determine the overall layout of the HPS apparatus. Radiation of a mediator that is heavy compared to the incoming electron carries away most of the energy in the reaction, so $x = E_{A'}/E_{\text{beam}}$ is peaked strongly at 1 [8]. Since HPS operates at beam energies beyond 1 GeV, which are large compared to A' masses of interest, the A' is highly boosted with its momentum closely aligned with the beam direction. The A' subsequently decays to an e^+e^- pair, leaving that pair also boosted in the very forward direction and azimuthally back-to-back with respect to the beamline. Therefore, a detector with excellent forward acceptance immediately downstream of the target is required to detect the e^+e^- decay products and cleanly identify secondary vertices as close to the target and through-going beam as possible.

HPS realizes this concept with a magnetic spectrometer, consisting of a multi-layer Silicon Vertex Tracker (SVT) situated within a large dipole magnet (0.24 T for the beam energy described in this note), to measure the momenta and trajectories of the e^+e^- pair. The field of the dipole is vertical, dispersing most of the beam electrons that have radiated in the target, as well as other electromagnetic backgrounds, into the horizontal plane (“sheet of flame”) containing the beam. As a result, the SVT is split into two segments, one above and one below the beam plane, which are positioned as close to it




FIG. 2. A cutaway view of the HPS detector showing the Silicon Vertex Tracker (SVT) in a vacuum chamber inside the bore of the spectrometer magnet and the Electromagnetic Calorimeter (ECal) downstream. The positions of the target and the front portions of the SVT are controlled by a set of linear positioning motors upstream of the detector.

as possible to maximize acceptance. The extent of the forward acceptance is limited by the background rate of single beam electrons that scatter in the target, which cannot mimic the signal but creates extreme occupancies ($\approx 4 \text{ MHz/mm}^2$) at the edge of the first layer of the SVT. As a consequence, the high repetition rate of the CEBAF beam (499 MHz), in tandem with a high-rate e^+e^- trigger with precision timing and similarly precise timing in the SVT, is required to select only in-time hits for readout and reconstruction. HPS uses a lead tungstate electromagnetic calorimeter (ECal) – also split above and below the beam plane – to provide this trigger and an off-line estimate of the precise hit time, along with particle identification for the tracks reconstructed in the SVT. These key components of the HPS apparatus are shown in Figure 2.

A. The JLab CEBAF

The HPS experiment utilizes beam from CEBAF at the JLab in Newport News, Virginia. CEBAF is oval shaped, consisting of two linacs connected by a pair of recirculating arcs, which enables injected beam to make multiple passes of the linacs — gaining 2.2 GeV per pass for up to 5.5 passes — before extracting the beam into one of four halls. Sub-harmonics of the 1.497 GHz beam may be simultaneously extracted into the different halls, allowing simultaneous operation of multiple experiments with high-rate (typically 499 MHz) beam. [16]. Operation at the JLab CEBAF is fundamental to the success of the HPS experiment. The experiment requires a very high repetition rate multi-GeV electron beam with low per-bunch charge, together with precision hit timing in all subsystems, in order to screen out the high rate of background hits from scattered single electrons. A higher per-bunch charge would spoil the clean tracking

and vertexing required for the displaced vertex search, while a lower current would require unacceptably long operations.

B. Hall B Beamline and Target

The HPS apparatus operates in the downstream alcove of experimental Hall B[17], as shown in Figure 3. The 2.3 GeV electron beam is transported ≈ 57 m from the upstream Hall B tunnel to HPS, passing through a number of quadrupole and dipole magnets that focus and steer the beam to the target. The extraordinary proximity of the SVT layers to the beam, as close as 500 microns between the edges of sensors and the center of the beam, places stringent requirements on the quality of the beam; a very small beamspot ($< 50 \mu\text{m}$ vertically) with vanishing low halo rate ($\lesssim 10^{-6}$ outside the Gaussian core) and excellent beam stability ($< 30 \mu\text{m}$ vertical variation).

Ensuring the safety of the SVT also requires multiple diagnostic and protection systems. During beam setup, the beam profile and position are measured by wire scanners (“harps”) located strategically along the beamline and used to tune the trajectory to produce the desired spot on the target. In addition, there are wires integrated into the movable structures of the SVT that are close to the target and precisely referenced to the positions of the silicon sensors that can be used to ensure the ideal profile and position of the beam. A typical scan of the beam with an SVT wire is shown in Figure 4. Beam position monitors (BPMs) are used to continuously monitor the transverse position of the beam at multiple locations during data-taking, and are tied to machine controls (orbit locks) to ensure the stability of the beam trajectory, as demonstrated by Figure 5. A set of halo counters around the apparatus monitors background levels to detect any scraping of the beam upstream of or inside the apparatus. In addition to providing the data for harp scans, the halo counters are tied to the Fast Shut Down (FSD) system of CEBAF, and can trigger beam shutdown within a few milliseconds of exceeding settable thresholds. Finally, a collimator with a choice of several apertures directly upstream of HPS is used to protect the detector from large beam excursions during tuning and operations.

The target for the experiment is chosen to be as thin as possible, given the upper limit on beam currents in Hall B to achieve the desired luminosity, in order to minimize occupancy in the detector, of which two-step processes in the target generate a small but significant component. The target system consists of a movable assembly with different thickness tungsten foils, in addition to carbon and polyethylene targets for calibration purposes. The data analyzed for this paper were taken with a $4 \mu\text{m}$ tungsten foil, equivalent to approximately 0.125% of a radiation length.

C. Silicon Vertex Tracker

The SVT is a six layer, high precision, silicon tracking and vertexing detector responsible for estimating both the mass and decay position of e^+e^- pairs by measuring the momenta and trajectories of charged particles. The design of the SVT, shown in Figure 6, is shaped by a few competing requirements. First, A' decay products have typical momenta $\lesssim E_{\text{beam}}/2$, so multiple scattering dominates mass and decay length errors for any feasible material budget. Second, the signal yields for long-lived A' s are very small, so the rejection of prompt vertices must be exceedingly strong, better than 10^{-6} , to reduce prompt background to the order of one event or less. Finally, as previously discussed, the passage of scattered and degraded primary beam through the apparatus creates a region of extreme occupancy and radiation in the same part of the detector that is critical for sensitivity to low-mass A' that have decay products nearly collinear with the beam. This puts low-mass acceptance at odds with tracking and vertexing purity and the material budget for the detector, requiring careful design to allow the largest usable acceptance. A prototype detector, with many of the same general features and utilizing the same sensor design, is described in more detail in [18].

The SVT employs radiation tolerant silicon microstrip sensors developed for the DØ RunIIb project[19], which allows the readout and cooling material to be placed outside the tracking volume. The sensors and their front-end readout electronics are cooled from the ends via their support structures to below -10°C to extend their lifetime at peak fluences exceeding 10^{16} electrons/ cm^2 (or 4×10^{14} [1 MeV neutron equivalent]). The SVT is split into mirror-symmetric halves, above and below the plane of the scattered and degraded beam. As a result, the regions of high occupancy are small spots along the sensor edges, so that only a very short length of the edge strips see high occupancy. As a result, long strips covering those regions have per-channel occupancies only a small factor larger than what pixels would experience.

Each layer of the SVT consists of sensors placed back-to-back 7.4 mm apart with a small stereo angle between them (100 mrad in front three layers; 50 mrad in back three), so that 3D space points can be determined. Each half of the detector – top and bottom – is further divided into two separate structures – front and back – with three layers each. The first three detector layers are a single sensor in width and spaced 10 cm apart along the beam direction, with the first layer just 10 cm downstream of the target. The next three layers are two sensors wide and spaced 20 cm apart beginning 50 cm downstream of the target, and consist of double width axial and stereo layers, which improves acceptance for low-momentum particles. All four detector segments, with 36 sensors and 23004 channels total, are placed as close to the beam as backgrounds allow, with acceptance down to 15 mrad above and below the beam plane with respect to the beam spot on the target. Since this places




FIG. 3. Engineering rendering of the downstream alcove of experimental Hall-B, where the HPS apparatus is located.




FIG. 4. An example beam vertical profile obtained with the SVT wire during final beam tuning.

the active (passive) edges of the sensors in the first layer 1.5 mm (0.5 mm) from the center of the beam, precision construction, alignment, and survey of the sensors are essential, and the structures holding the first three layers are movable, allowing them to be retracted from the beam during beam tuning. To eliminate displaced events and occupancy from beam-gas collisions, the SVT must operate inside the beam vacuum, and resides within a vacuum enclosure installed inside a dipole magnet with a downward-pointing central field of 0.24 T for this beam energy.

The sensors of the SVT are read out by APV25 ASICs [20] mounted on hybrid PCBs and wirebonded directly to the sensors. For each triggered event, the APV25 records six consecutive analog samples of the front end amplifier




FIG. 5. Distribution of the beam x and y positions reported by the 2H02 BPM, ≈ 2 m upstream of the target, over a period of roughly one hour.

output. An offline fit to the six samples allows a precise estimate of both the amplitude and the time of hits, as well as the discrimination of events with two hits that are very close in time. Power, control, and monitoring of the hybrids, and clocking, control and digitization of APV25 samples are performed by a set of Front End Boards (FEBs) also located inside the SVT vacuum enclosure, to minimize the length of the analog cables and reduce the number of signals that must penetrate the vacuum barrier. Being in vacuum, the FEBs also require liquid cooling, which uses a separate system from the sensor modules to allow the temperatures of the two systems to be set independently. Power and digital signals are passed from the FEBs via vacuum feedthroughs in a




FIG. 6. A diagram showing the SVT layout inside the vacuum box.




FIG. 7. The time resolution of SVT hits associated with tracks versus layer number.

pair of flanges to the power supplies and central data acquisition system (DAQ) outside of the vacuum chamber. The central DAQ for the SVT, based on the Reconfigurable Cluster Element (RCE) architecture[21], connects to the data flange via 50 m optical fibers, allowing it to be placed in a lower radiation environment. The SVT DAQ is capable of very high data rates, which is necessary to accommodate the torrent of irreducible trident backgrounds that must be accepted in order to search for rare A' events.

To further reduce occupancies for tracking, the CMS APV25 chip is used for readout in “multi-peak” mode, which records 6 samples of the signal development, allowing reconstruction of hit time with ≈ 2 ns resolution — near the level required to tag events in individual CEBAF bunches. Figure 7 shows the time resolution versus layer number, with the inner, high-occupancy layers having slightly worse resolution than the back layers. This is also reflected in the signal-to-noise of the sensors versus layer, shown in Figure 8.




FIG. 8. The mean signal-over-noise of SVT hits associated with tracks versus layer number.

D. Electromagnetic Calorimeter

The HPS Electromagnetic Calorimeter (ECal) [22] plays two critical roles. First, it provides a trigger for e^+e^- pairs with sufficient energy and time resolution to eliminate the overwhelming background of scattered single beam electrons. Second, it provides positive identification of electromagnetic energy deposits — from electrons, positrons, or photons — offline, with sufficient time resolution to tag them to a single CEBAF bunch, which can then be used to demand coincidence with tracks in the SVT. Like the SVT, the ECal must contend with extremely high rates and be relatively radiation tolerant in order to match the angular acceptance of the SVT as closely as possible.

The ECal meets these requirements through the use of 442 PbWO_4 crystals arranged in two identical arrays — placed symmetrically above and below the beam plane downstream of the SVT. The through-going degraded beam is transported between the two halves in a vacuum chamber to eliminate beam-gas backgrounds. Each half is a matrix of 5×46 PbWO_4 crystals. From the first row of each half, 9 crystals are removed nearest the through-going beam as the rate of scattered beam electrons is intolerably high in that region, well in excess of 1 MHz. The crystal layout and some mechanical elements of the ECal are shown in Figure 9. The ECal channels are read out via APD and 250 MHz Flash ADC (FADC) boards which record samples of the pulses every 4 ns. This provides a similar time window for triggers, whereas offline fitting of the FADC pulses provides a much better time estimate for ECal hits.

E. Trigger System

As outlined at the beginning of Section II, A' production is peaked at small angles with respect to the beam direction, so the e^+e^- decay daughters are typically back-to-back relative to the beam direction [8]. As a result,




FIG. 9. ECal crystal layout, as seen in the beam direction. For clarity, the top-half mechanical parts have been removed. For the bottom half, some mechanical elements such as the mother boards (in green) and the copper plates for heat shielding (in red) are visible. Between the two halves of ECal, the beam vacuum vessel is seen to be extended to the right to accommodate beam particles having lost energy through scattering or radiation.

when one daughter falls within the acceptance of the top half of the detector, the other will fall within the bottom acceptance. Meanwhile, the vertical magnetic field of the spectrometer magnet will bend the electron and positron in opposite horizontal directions. Therefore, the primary trigger for the experiment is a “pair trigger” in the ECal which requires energetic clusters in both halves (top and bottom) of the ECal, and with the two clusters displaced horizontally in opposite directions from the centerline according to their energies, since lower-energy particles will curve more in the magnetic field.

Simulations showed that the two clusters in signal events are nearly back-to-back azimuthally, so the trigger requires that the azimuthal coplanarity of the two clusters is close to 0, as shown in Figure 10. The trigger also places a cut on the minimum cluster energies as a function of their horizontal displacement from the center-line of the ECal, according to $E + F \times r > E_{\text{threshold}}$. Here E is the cluster energy, r is the distance of the cluster from the center of the calorimeter shown in Figure 10, and parameters F and $E_{\text{threshold}}$ are tuned using Monte Carlo simulation. This cut mostly eliminates the high rate of bremsstrahlung events with low-energy photons hitting close to the center of the ECal.

III. ANALYSIS OVERVIEW

Our search for heavy photons uses two different techniques, outlined in detail below. The first is a traditional resonance search, where we search for a resolution-dominated resonance shape superposed on the copious e^+e^- invariant mass distribution which arises primarily from QED tridents. Heavy photons with relatively large coupling strengths have very short decay lengths, so appear prompt and would be detected in this search. The second is a vertex search for e^+e^- decay vertices sig-

nificantly displaced from the target. The vertex search examines the observed decay length distribution mass-bin by mass-bin and looks for events beyond a cut where prompt backgrounds are expected to be small. Heavy photons with very small coupling strengths would have correspondingly large decay lengths, and would be detected in the vertex search. Thus HPS searches in two distinct regions of the heavy photon mass/coupling plane. Both searches are blind searches, in the sense that all analysis cuts are frozen after inspecting 10% of the data. The final analysis of the full data set incorporates those cuts.

Event selection and various data quality cuts are common to the two analyses, but not identical. The displaced vertex search, in particular, adopts special cuts to identify and eliminate long-lived backgrounds. The differences are detailed below. Both analyses calculate their sensitivity to heavy photon production using the observed flux of e^+e^- pairs, which is predominantly due to two QED processes, trident production and wide-angle bremsstrahlung conversion (cWAB), which have a known relationship to the heavy photon production rate. Trident production occurs via two processes, radiative (Figure 11) and Bethe-Heitler (Figure 12), and the interference between them. Both analyses require e^+e^- pairs with total energy near that of the incident electron, as expected for heavy photon production. For tridents, this means that the observed pair likely excludes the recoil electron. For cWABs which convert in the target or first detector layer, the observed pair is usually the conversion positron and the recoil electron. Monte Carlo simulation of trident and cWAB production, incorporating their calculated cross-sections, reasonably accounts for the observed rate and momentum spectrum of e^+e^- pairs, demonstrating a good understanding of the sample composition. This procedure reduces dependence on experimental efficiencies. Using this Monte Carlo estimation, the fraction of the observed events attributable to the purely radiative trident production diagram, which is proportional to heavy photon production, is determined. Hence, we calculate sensitivities to heavy photon production incorporating theoretical knowledge of trident and wide-angle bremsstrahlung cross-sections but normalized by the data.

Both analyses depend on knowing the experimental mass resolution and the invariant mass scale. For the resonance search, the mass resolution determines the width of the expected heavy photon resonance; the invariant mass scale, its exact position. For the displaced vertex search, the mass resolution determines what fraction of the signal appears in a given mass slice and how much background is included. Møller scattering results in e^-e^- pairs of fixed mass for a given beam energy. Measuring the position and width of the Møller peak enables calibration of mass scale and resolution.

The measured decay length distributions for background (prompt) e^+e^- pairs arising from tridents and cWABs can be characterized by a broad Gaussian cen-




FIG. 10. An illustration of an event satisfying trigger requirements. As described in the text, data for this analysis is collected using a pairs trigger that makes requirements on $|\phi_1 - \phi_2|$ and the relationship between $r_{1,2}$ and $E_{1,2}$.




FIG. 11. Diagram of radiative trident production off the tungsten target




FIG. 12. Diagram of Bethe-Heitler trident production off the tungsten target.

tered on the target location, with an exponential tail at large decay lengths. These features are the result of how the exiting e^+ and e^- multiple Coulomb scatter, where

resolution is dominated by scattering in the first detector layer. Good agreement between Monte Carlo and data decay length distributions confirms our understanding of decay length resolution.

The following subsections review the data samples, detector calibration and event reconstruction, event selection, sample composition, and mass resolution for the two analyses. The resonance search and displaced vertex search sections that follow discuss the specifics of each analysis in more detail.

A. Data Samples

The results presented here use data collected during the 2016 Engineering Run, which operated on weekends during February 20–April 25 of 2016. All data used for analysis were collected at a beam energy of 2.30 GeV with a current of 200 nA on a Tungsten foil target 4 μm ($\approx 0.125\%$ X_0 equivalent) thick. The total luminosity of this dataset is $10\,608\text{ nb}^{-1}$, comprising 7.2 billion triggered events from a total charge on target of 67.2 mC.

In addition to physics runs, a number of special runs were taken, such as field-off runs and runs with a trigger dedicated to collecting scattered single electrons over a wide range of scattering angles. Data from these runs were used to calibrate and align the ECal and SVT.

In addition to experimental data, the analysis presented here makes use of Monte Carlo simulation (MC) to understand some attributes of signal and background. MadGraph[23] is used to generate samples of A' signal at a range of masses, as well as tridents, which include both Bethe-Heitler and radiative tridents (which are kinematically identical to signal) and their interference term, and converted WAB events. Monte Carlo of Møller scattering events is also used to study the mass resolution. Beam backgrounds simulated using EGS5 [24], predominantly scattered single electrons, are overlaid on all samples, distributed according to the time structure of the beam. Simulation of generated samples uses GEANT4 [25] to model interactions with the detector, after which the de-




FIG. 13. Energy resolution of the ECal as a function of energy. The three points below 1.2 GeV were obtained from the 2015 Run, while the point at 2.3 GeV, which benefits from electronics upgrades, was obtained from the 2016 Run using elastically scattered electrons, and was not used in the fit.

tector response simulation and reconstruction are performed.

B. Detector Calibration and Event Reconstruction

Raw data from the detector and simulation are reconstructed to produce the physics objects used for analysis, which are reconstructed e^- and e^+ as well as A' candidates consisting of reconstructed e^+e^- pairs emanating from a common vertex, which we refer to as “V0 Candidates”. The reconstruction of e^- and e^+ is stepwise, taking place first separately in the ECal and the SVT, and then combining information from both subsystems.

1. ECal Calibration and Reconstruction

The crystals of the ECal are small compared to the Moliere radius in PbWO_4 , so in order to reconstruct and identify electrons and positrons in the calorimeter, the energy depositions in individual crystals must be calibrated and then combined, or “clustered”, to provide a good estimate of the energy of incident electrons and positrons.

Calibration uses both minimum ionizing particles (MIPs) from cosmic ray events, as well as samples of scattered beam electrons collected with a special trigger, to determine the conversion of pulse height to energy. The simple clustering algorithm, which begins with a high-energy seed and iteratively adds adjacent crystals above a threshold, results in good energy resolution, as shown in Figure 13. The pulse fit to the 250 MHz FADC readout stream also results in excellent time resolution, as shown




FIG. 14. Time resolution of hits as a function of hit energy. The time estimate comes from a fit to samples of the SiPM output at 4 ns intervals provided by the FADC readout used by the ECal.

in Figure 14. More details may be found in [22].

2. SVT Calibration and Reconstruction

The reconstruction of charged particle trajectories in the SVT detector starts with the formation of 3D space-points by combining the axial and stereo strip clusters on the two sides of each silicon module. In order to accept a 3D space-point, the two strip clusters’ reconstructed times are required to be in a time window of 12 ns from the trigger time and within 16 ns of each other. Three 3D space points in selected SVT layers are then grouped together to form a track seed and an initial estimation of the track parameters is obtained by performing a helical fit under the assumption of a uniform magnetic field. The track-seed finding efficiency is maximized by choosing multiple combinations of the 3D space point triplets with different layer combinations to start the pattern recognition. Track-seeds are then extended by iteratively adding 3D space-points located on the other SVT layers and performing a global helical track fit selecting the track candidate with minimum χ^2 during the procedure. At this stage, track candidates are required to have at least 5 associated 3D space points, momentum $p > 100$ MeV and track quality $\chi^2_{5\text{hits}} < 60$ and $\chi^2_{6\text{hits}} < 84$, for track candidates with 5 and 6 hits respectively. Track candidates are then refitted with the General Broken Lines (GBL) [26] algorithm to include the effects of multiple scattering and refine the initial estimate of the track parameters. The GBL-refitted trajectories are also used for calibration and alignment of the SVT using Millepede II [27]. The electron and positron particle candidates are then formed by requiring each reconstructed track to be associated with an ECAL cluster.

Using exactly two final state particles reconstructed in

the two detector volumes, vertices are then reconstructed using a global χ^2 minimization algorithm [28]. The final state particles used for vertex reconstruction are required to have an ECal cluster time difference within 2.5 ns and the electron momentum $p_{\text{ele}} < 2.18$ GeV. Successfully reconstructed vertices are required to have a total momentum $p_{\text{vtx}} < 2.8$ GeV.

The performance of the SVT can be characterized by its tracking efficiency, momentum resolution, and vertex position resolution. Tracking efficiency is measured by selectively dropping hits in a particular layer from the track finding code, extrapolating the track as measured by the other layers to that layer, and measuring the fraction of times hits are found within a predicted region. Efficiencies are $> 90\%$ in most of the SVT but are somewhat worse in the inner edges of the first two layers. Dead channels also have noticeable effects. In the analysis below, tracking efficiency effects are included in critical simulations. The hit-finding efficiency for two of the first layers is shown in Figure 15.

Momentum resolution is determined by measuring the momentum of elastically scattered beam electrons, which essentially have full beam energy (FEEs). Since the momentum resolution is dominated by multiple scattering effects in the SVT, determination of the momentum resolution of the highest momentum tracks suffices to characterize the resolution at all momenta. The Monte Carlo does not accurately account for the observed momentum resolution, with simulation being better than reality. Accounting for this discrepancy is important in order to understand the actual invariant mass resolution, a critical parameter in the analysis. Procedures for doing so are described in Section III E.

The vertex position resolution of the tracker is easily measured by vertexing the copious trident signal, which originates at the known target position. The vertex resolution is well described by Monte Carlo simulation and is detailed in Section V. The typical vertex resolution along the direction of the outgoing particle is 1 mm.

C. Event Selection

The event selection cuts were developed blindly, i.e. only 10% of the data was used for their optimization. Once determined, these cuts were used to select final events for both the resonance and vertex search analyses over the full data set.

The HPS experiment searches for A' 's through their decays to e^+e^- , so an event is required to contain at least one neutral, two-particle vertex (called a V0). Due to the kinematics of A' production, the electron and positron will almost always be in opposite halves of the HPS detector, so one track is required to be in the top half, the other in the bottom. One of the particles must be positively charged, the other negatively charged. Each of the particles must point to a cluster in the ECal. A V0 candidate is formed by fitting the two charged tracks




FIG. 15. The hit-finding efficiency versus the extrapolated SVT channel number for two SVT layers at the front of the detector, one each in the top (red) and bottom (blue) halves. The nominal center of the electron beam at these layers is ~ 1.5 mm from channel zero (which is the edge of the active sensor). The drop in efficiency is due to a combination of extrapolation error at the edge of the active volume and pile-up effects.

to a vertex, following the procedures described in [29]: The vector momentum sum of the electron and positron, P_{sum} , must meet the condition $P_{\text{sum}} < 1.2 \times P_{\text{beam}}$, where P_{beam} is the beam momentum (2.3 GeV in this run).

After the V0 candidates are formed, two V0 collections are created. These are the “unconstrained V0 candidates” (UC) and “target-constrained V0 candidates” (TC). In these collections, a V0 particle is created and defined as the parent of the corresponding e^+e^- pair. In the resonance search analysis, we use the TC V0s, where the z -coordinate of the vertex is constrained to be at the target position, and x - y coordinates are constrained at the beam spot coordinate. The displaced vertex search analysis specifically searches for long-lived particles that have traveled some distance before decaying. Therefore, in the displaced vertex search analysis, the UC collection is used.

Further cuts on the V0 properties were imposed to minimize accidental backgrounds, maximize the signal-to-background ratio of the radiative signal, and reduce physics backgrounds. Accidental backgrounds can be minimized by optimizing the cut on the time difference between the two ECal clusters. Figure 16 shows this cluster time difference, which is sharply peaked at 0. The bottom panel shows the same data, but with the vertical scale magnified to show the structure in the tails, displaying peaks that occur at multiples of 2 ns, the spacing between CEBAF’s electron bunches. It shows that accidental coincidences between bunches occur at a low level. This distribution is fit with the function given in Equation (3) as a sum of peaks where each subpeak is parameterized as the sum of two Gaussian functions, one describing its core and another, wider and of lower amplitude, its tail. The ratio of the amplitudes of these two

Gaussians is constrained to be the same for all peaks. The optimum time interval is chosen to maximize the ratio $S/\sqrt{S+B}$ where S is the integral of the central peak in the given $\pm\Delta t$ cut range, and $S+B$ is the integral of signal plus background.

$$F = \sum_{i=0}^{N_{\text{peak}}} a_i \cdot (\text{Gauss}(x - \mu_i, \sigma_{1,i}) + b \cdot \text{Gauss}(x - \mu_i, \sigma_{2,i})) \quad (3)$$

For the resonance (vertex) search, the absolute value of the cluster time difference must be less than 1.43 ns (1.45 ns).




FIG. 16. Top and bottom cluster time difference, when the cluster energy sum is in the range 1.9 GeV to 2.4 GeV. Bottom figure is the same as the top with vertical axis adjusted to show peaks in tails.

Figure 17 shows the differential cross sections for the various physics processes that contribute to the event sample as a function of the V0 momentum. Radiative tridents are labeled in red, the WAB sample in orange, the trident sample in cyan, and the sum of WAB and tridents in blue. Radiative tridents are peaked at high momenta, whereas the full trident sample (which includes radiative tridents and Bethe-Heitlers and their interference) and WABs are broadly enhanced at lower momenta. The sensitivity of the resonance search is proportional to the radiative fraction, so a cut in the minimum V0 momentum that maximizes the ratio $N_{\text{rad}}/\sqrt{N_{\text{tot}}}$ is optimal. For the resonance (vertex) search, this occurs at 1.9 GeV (1.85 GeV).




FIG. 17. Differential cross-section as a function of V0 momentum for different MC samples. Radiative tridents are presented by red markers, the WAB sample is represented by the orange-colored histogram, the trident sample is shown by the cyan histogram, and the blue histogram is the sum of WAB and tridents.

Finally, a cut on the maximum V0 momentum reduces background from the cWABs, which extends beyond the beam energy. For both the resonance and vertex searches, the maximum V0 momentum must be less than 2.4 GeV.

Figure 18 compares the data with the Monte Carlo after all the above cuts. The data and MC are in broad agreement, giving evidence that the sample composition is understood. At lower V0 momentum, the data fall below the MC, primarily because the trigger efficiency for low-energy clusters is not perfectly accounted for in the Monte Carlo. Momentum resolution effects, also not perfectly accounted for, explain small data/Monte Carlo discrepancies at the high edge of the distribution. The invariant e^+e^- mass distribution for events passing these final cuts is shown in Figure 19. The highlighted region in green is the mass range where the resonance search was performed. The mass range for the displaced vertex search is discussed in Section V.




FIG. 18. Number of events as a function of V0 Momentum for a subset of the data, shown in yellow, and different MC samples. Radiative tridents are presented by red markers, the WAB sample is represented by the orange-colored histogram, the trident sample is shown by the cyan histogram, and the blue histogram is the sum of WAB and tridents.




FIG. 19. The mass distribution after all event selection cuts (described in the text). The highlighted area in green represents the range where the resonance search is performed.

D. Sample Composition and Fraction of Radiative Rate

While the nominal A' events are primarily at high x , trident events cover the entire x range and have a higher rate at low x . The HPS detector accepts $\sim 0.5 < x < 1$, and although events with $x < 0.8$ are not useful for A' searches, they provide a high statistics sample for calibrations and sample composition studies.

It can be shown that the expected signal cross-section is [8]:

$$\left. \frac{d\sigma_{A'}}{dm} \right|_{m=m_{A'}} = \frac{3\pi m_{A'} \epsilon^2}{2N_{\text{eff}} \alpha} \left. \frac{d\sigma_{\gamma^*}}{dm} \right|_{m=m_{A'}} \quad (4)$$

The luminosity, detector acceptance, and efficiency are factored out from this equation and it is rearranged to give an equation to calculate the upper limit on ϵ^2 via an upper limit on the signal rate. N_{eff} is the ratio of the

sum of all branching ratios to the branching ratio of the electron-positron decay channel, and is one for all masses in this search. The differential cross section $d\sigma_{\gamma^*}$ is taken at specifically the A' mass, and the notation indicating this will be dropped. This gives:

$$\epsilon^2 = \frac{2\alpha N_{\text{sig}}^{\text{up}}}{3\pi m_{A'} \frac{dN_{\gamma^*}}{dm}} \quad (5)$$

where $N_{\text{sig}}^{\text{up}}$ is the upper limit on the number of signal events observed in the data. Section IV A 1 discusses in detail how this upper limit is set. The focus of this section is to present how the differential rate dN_{γ^*}/dm is evaluated in the analysis.

The differential γ^* rate is only defined theoretically and is not something that can be directly extracted from the data. We start by defining the radiative fraction as:

$$f_{\text{rad}} = \frac{\frac{dN_{\gamma^*}}{dm}}{\frac{dN_{\text{bkg}}}{dm}} = \frac{\frac{dN_{\gamma^*}}{dm}}{\frac{dN_{\text{tri}}}{dm} + \frac{dN_{\text{wab}}}{dm}} \quad (6)$$

where N_{tri} and N_{wab} are the number of trident and WAB events, respectively.

Using this definition the equation for ϵ^2 can be rewritten as:

$$\epsilon^2 = \frac{2\alpha N_{\text{sig}}^{\text{up}}}{3\pi m_{A'} f_{\text{rad}} \frac{dN_{\text{bkg}}}{dm}} \quad (7)$$

It is important to note that the differential background rate measured with the data is a function of the reconstructed mass, while the differential rate of γ^* is a function of the true γ^* mass. This means Equation (7) is interpreted to have the true mass in the numerator and the reconstructed mass in the denominator to minimize systematic uncertainty from events migrating into other mass bins due to resolution effects. The differential background rate is extracted directly via the fit described in Section IV A 1. The radiative fraction is constructed entirely via Monte Carlo simulations, and since the signal model is only sensitive to the peaking part of the signal, only the contribution from reconstructed events that correctly match the radiative electrons, and not the recoil electron, is used.

Figure 20 shows the radiative fraction versus invariant mass for the resonance search selection. The corresponding plot for the displaced vertex search is shown in Section V.

E. Invariant Mass Resolution

Searching for a resonance peak on top of a large background requires accurate knowledge of its width. The width of the expected A' signal is dominated by the experimental resolution, so it is critical that the mass resolution is well understood. The mass resolution for ob-




FIG. 20. The differential rates in units of MeV^{-1} for all the MC samples that go into the radiative fraction for the target-constrained vertex fit used in the resonance search analysis. The radiative component is a function of the true mass of the MC generated γ^* for the event. The WAB and Trident components are a function of the invariant mass of the selected V0 candidate. The bottom panel shows the radiative fraction. The red line is the fifth order polynomial fit, and the dotted line is the zeroth order polynomial fit.

served Møller events is compared to Monte Carlo simulations, which are then tuned to get agreement. The adjusted Monte Carlo is then used to derive the expected mass resolution for all masses of interest to the analyses. These steps are detailed in this section.

The Monte Carlo is used to evaluate the mass resolution using simulations of A' signal at several fixed mass points. These generated signal events are processed through the GEANT4 simulation chain with a full detector model. Since the natural width of the A' is significantly smaller than the detector resolution (by more than a factor of 1000), the observed width of the signal shape is determined solely by the mass resolution.

1. Using the Møller Resonance to Calibrate the MC Mass Resolution

The Møller process $e^-e^- \rightarrow e^-e^-$ provides a direct measurement of the mass resolution since the center of mass energy of a beam electron and an electron at rest is equal to the invariant mass of the final state electrons (called the Møller mass). A beam energy of 2.3 GeV will have a Møller mass of 48.5 MeV. Just like the A' process, the observed width of the Møller invariant mass is dominated by detector resolution. Furthermore, the Møller and A' final states both have particles of equal mass which will multiple scatter in the detector material essentially identically. The mass resolution for e^+e^- and e^-e^- final states is expected to be equivalent at the same invariant mass. Figure 21 shows the e^-e^- invariant mass distributions of Møller events in the data (cyan) and MC (blue). These histograms have been scaled to have the




FIG. 21. Mass distribution of Møller events from data and MC overlaid. Histograms are scaled to have the same maximum value. The cyan line represents data, while the blue line represents MC.

same maximum bin value. Note the mass resolution in data is about a factor of two worse than in MC.

The Møller mass can be written in terms of the momenta of the final state particles (P_1 and P_2) and the angle theta between them (θ), neglecting the mass squared terms:

$$M = 2\sqrt{P_1 P_2} \cdot \sin \frac{\theta}{2}. \quad (8)$$

This formula demonstrates the source of the discrepancy in the Møller mass resolutions in data and MC can be modeled as discrepancies in the momentum resolution and/or angular resolution.

2. Momentum Resolutions With Full Energy Electrons

Elastically scattered full beam energy electrons (FEEs) provide an experimental check of the momentum scale and resolution. Since the electron is so light compared to a tungsten nucleus, it loses nearly zero energy in elastic interactions. Consequently, elastically scattered beam electrons are expected to appear as a single peak in the electron momentum distribution. The width of this peak is a measurement of the momentum resolution at the beam energy. As it is natural to expect better momentum resolution for 6 hit tracks compared to 5 hit tracks, these resolutions are measured separately. The top and bottom tracks are also separated because the two detector halves are not expected to have systematically identical misalignments. Figure 22 shows FEE peaks for 6 hit negative tracks in the bottom half of the tracker, where the cyan line is from data and the blue is MC. In this particular case, the data resolution is a factor of 1.6 times worse than the MC resolution. Over all the categories, the momentum resolution in data is worse than that in MC by factors ranging from 1.3 to 1.6.

Adding additional momentum smearing can bring the MC and data mass into agreement. The smearing coefficients for each MC category (bot/top/5-hit/6-hit) are parameterized by:

$$\Sigma_{\text{smear}} \equiv \frac{\sigma_{\text{smear}}}{P_{\text{MC}}} = \sqrt{\left(\frac{\sigma_{\text{data}}}{\mu_{\text{data}}}\right)^2 - \left(\frac{\sigma_{\text{MC}}}{\mu_{\text{MC}}}\right)^2}. \quad (9)$$

where σ_{smear} is the factor by which an MC electron with a given momentum (P_{MC}) is smeared. The data and MC FEE momentum resolutions are σ_{data} and σ_{MC} , respectively. Finally, μ_{data} and μ_{MC} are the mean values of the FEE momentum peaks.

The MC tracks are then smeared with the appropriate Σ , depending on the category. Figure 23 compares the smeared MC momentum distribution in blue with data in cyan. The mean of the MC distribution has been shifted slightly so that the peaks overlap for ease of comparison. The matching between MC and data for other categories is comparable. In all cases, there is good agreement between data and the smeared MC distributions. Accordingly, smearing is applied to all the tracks from the Møller and A' MC samples.




FIG. 22. FEE momentum distributions for 6 hit tracks in the bottom half of the tracker. The cyan line represents data and the blue line represents the un-smeared MC.

3. Recalculated Mass after MC Momentum Smearing

The Møller mass is recalculated using the smeared electron momenta. The mass taking into account the smeared momenta can be expressed in terms of the un-smeared mass, using Equation (10).

$$M(\text{ee})^{\text{smeared}} = \sqrt{\frac{P_{1,\text{smeared}}}{P_{1,\text{rec}}} \frac{P_{2,\text{smeared}}}{P_{2,\text{rec}}}} \cdot M(\text{ee}) \quad (10)$$

Here, $M(\text{ee})^{\text{smeared}}$ is the smeared mass, $P_{1,\text{smeared}}$ ($P_{2,\text{smeared}}$) is the smeared momentum of 1st (2nd) particle,




FIG. 23. FEE momentum distributions for 6 hit tracks in the bottom half of the tracker. The cyan line represents data and the blue line represents the smeared MC momentum.

$P_{1,\text{rec}}$ ($P_{2,\text{rec}}$) is the reconstructed (unsmeared) momentum of the 1st (2nd) particle, $M(\text{ee})$ is the unsmeared target constrained mass. The momentum resolution discrepancy between data and MC is assumed to be independent of momentum. This is expected since $\sigma(p)/p$ is nearly constant over all momentum, being multiple scattering dominated.

After smearing the mass with Equation (10), the smeared mass of Møller events shown in Figure 24 (blue) is obtained. Incorporating smearing, the mass resolution discrepancy is reduced from about a factor of 2 to about 6%, which indicates that momentum resolution accounts for nearly all the mass resolution discrepancy between MC and data.




FIG. 24. Smeared mass distribution of Møller MC events (blue), and Møller events in data (cyan)

4. Parametrizing the A' Mass Resolution

We study the expected mass resolution for A's of various masses using a collection of A' samples with masses




FIG. 25. Mass resolutions. Filled (empty) star markers represent the unsmeared (smeared) target constrained MC Møller mass resolution. Diamond-shaped markers show the target constrained Møller mass resolution from data. Filled squares show the smeared target constrained A' mass resolutions, while the filled triangles represent UC A' mass resolutions. The solid curve in red is the fit over target constrained A' mass resolutions.

ranging from 40 MeV to 175 MeV, with the momenta of the e^- and e^+ tracks smeared with the procedure described above. The smeared mass distributions of all the A' MC samples are fit with a Gaussian function to obtain the mass resolutions. These smeared A' mass resolutions

and the Møller mass resolutions are shown in Figure 25. The vertical axis shows the mass resolution, and the horizontal axis represents the mean value, the mass, of the Gaussian fit.

The filled (empty) star markers represent unsmeared (smeared) TC MC Møller mass resolution. The diamond-shaped marker shows the TC Møller mass resolution from data. Blue filled squares represent smeared TC A' mass resolutions for different masses, while black triangles represent UC A' mass resolutions. These smeared points are fit with an $\mathcal{O}(4)$ polynomial function, which is used to parameterize the mass resolution as a function of mass. Only the fit of the TC A' smeared mass resolutions is shown, the solid red curve.

IV. RESONANCE SEARCH

This section describes the resonance search analysis. Event selection has been described above in Section III C. There the invariant e^+e^- mass distribution for events passing these final cuts was shown in Figure 19. The region highlighted in green is the mass range where the resonance search was performed. Below we describe the resonance search technique, systematic uncertainties, and final physics results. All Monte Carlo momenta and masses used in this section are smeared according to the procedure described in Section III E 2.

A. Statistical Analysis

If an A' exists within the acceptance of HPS, it will manifest itself as an excess in the e^+e^- invariant mass spectrum (a “bump”). The excess is expected to take the form of a Gaussian centered at the mass of the A' ($m_{A'}$) with a width equal to the mass resolution for that point as discussed in Section III E.

However, since the mass of the A' is not known, it is necessary to search for it at all possible masses. To do this, HPS employs a resonance search over a mass range of 39 MeV to 179 MeV, in steps of 1 MeV, using a maximum likelihood fit ratio to test the background-only hypothesis at each mass hypothesis. The full methodology of this process is discussed in detail in this section.

1. Resonance Search Methodology

First, a fit window is selected centered on each mass hypothesis. The width of this window is chosen carefully so as not to introduce a bias in the signal yield and to minimize the signal yield uncertainty due to the background shape uncertainty. An exception occurs when the mass hypothesis is near the edge of the invariant mass

distribution, and the fit window extends into a region where there are no reconstructed events. In these cases, the window is shifted such that the lower (upper) edge is at the lowest (highest) mass event, which results in the window no longer being centered on the mass hypothesis.

The probability density function for this window is defined by Equation (11).

$$P(m_{e^+e^-}) = \mu \cdot \phi(m_{e^+e^-} | m_{A'}, \sigma_{m_{A'}}) + 10^{L_N(m_{e^+e^-} | \vec{t})} \quad (11)$$

Where $m_{e^+e^-}$ is the e^+e^- invariant mass, μ is the signal yield, $\phi(m_{e^+e^-} | m_{A'}, \sigma_{m_{A'}})$ is a Gaussian probability distribution describing the signal shape, and $L_N(m_{e^+e^-} | \vec{t})$ is a Legendre polynomial of the first kind of order N with coefficients (also the nuisance parameters) $\vec{t} = \langle t_0, t_1, \dots, t_N \rangle$ used as the background model. We used order 5 polynomials at low mass and order 3 above 66 MeV. The fit window width is an integer multiple of the mass resolution varying from 6 to 10 depending on the mass.

A hypothesis test is constructed with a background-only null hypothesis, $\mathcal{H}_0$. This assumes no signal, and thus $\mu = 0$ in Equation (11). The certainty to which the null hypothesis may be rejected is determined by calculating a likelihood ratio test statistic.

We define:

$$\lambda(\mu) = \frac{\mathcal{L}(\mu, \hat{\theta})}{\mathcal{L}(\hat{\mu}, \hat{\theta})} \quad (12)$$

where $\mathcal{L}(\mu, \theta)$ represents the Poisson likelihood function for the data, where μ is the signal yield and θ represents the nuisance parameters. More specifically, $\hat{\theta}$ represents the set of nuisance parameters which maximizes the likelihood estimate for the case where $\mu = 0$. The denominator represents the combination of signal yield $\hat{\mu}$ and nuisance parameters $\hat{\theta}$ which maximizes the likelihood. A maximum likelihood fit of Equation (11) is performed, once with μ restricted to zero and once with μ as a free parameter, to determine the maximum value of the likelihood function for each model.

We construct a test statistic, similar to that used in [30], based on the likelihood ratio of the background plus signal model and the background-only model:

$$\tilde{q}_0 = \begin{cases} -2 \ln \lambda(0) & \hat{\mu} > 0 \\ +2 \ln \lambda(0) & \hat{\mu} \leq 0. \end{cases} \quad (13)$$

This creates a mapping from the test statistic to the likelihood ratio for all possible $\hat{\mu}$, which is sufficient to employ the Neymann-Pearson lemma. This test statistic will be Gaussian distributed in the asymptotic limit of large sample size [31]. The probability of the null hypothesis being consistent with the data given this test statistic is:

$$p = \int_{\tilde{q}_{0,\text{obs}}}^{\infty} f(\tilde{q}_0 | 0) d\tilde{q}_0 \quad (14)$$

where $f(q_0 | 0)$ is the expected Normal distribution of q_0 .

Putting these definitions together we can write the p-value simply as:

$$p = \begin{cases} 1 - \Phi(\sqrt{\tilde{q}_0}) & \tilde{q}_0 \geq 0 \\ 1 - \Phi(-\sqrt{-\tilde{q}_0}) & \tilde{q}_0 < 0 \end{cases} \quad (15)$$

where $\Phi(Z)$ is the Gaussian cumulative distribution function (CDF).

The p-value calculated using Equation (15) can then be compared to some threshold α to establish whether $\mathcal{H}_0$ is sufficiently unlikely to model the data. Conceptually, the p-value indicates the probability of the data being consistent with the background-only hypothesis. In particle physics, α is by convention chosen to be approximately 3×10^{-7} to make a discovery claim, or “ 5σ ”. However, since we consider more than one mass hypothesis, expected random fluctuations will produce a lower p-value somewhere in the search space as we search more masses. To account for this, it is necessary to estimate a correction due to this effect, commonly known as the “look-elsewhere effect”.

2. The Look-Elsewhere Effect

In this analysis, the search is performed independently for multiple mass hypotheses. Each individual search is summarized by the local p-value, but this only accounts for local statistical fluctuations for that individual search at that mass. When searching multiple mass hypotheses, only the most significant p-value obtained over all masses searched is considered. Properly interpreting this p-value must then include a correction accounting for the additional statistical fluctuations expected from searching more than a single mass hypothesis. If all search regions were independent of each other this correction could simply be approximated for p-values much less than 1 by:

$$p_{\text{global}} = N_{\text{reg}} p_{\text{local}}$$

where N_{reg} is the number of independent search regions searched [32]. In this case, since the raster size of the search is 1 MeV and this is less than the mass resolution, the mass hypotheses are not independent; therefore, N_{reg} is not simply the number of mass points in the search. The number of search regions is approximated via

$$N_{\text{reg}} \approx \frac{W}{\sigma_{\text{ave}}}$$

where W is the width of the full search window in mass and σ_{ave} is the average mass resolution in the window. It is found for this search that $N_{\text{reg}} \approx 30$.

3. Calculating an Upper Limit on the Signal Yield

The upper limit is calculated via an asymptotic approximation using methodology described in [31]. We

compare a hypothesis at some fixed signal rate to the maximum likelihood estimate of the signal rate or the background-only hypothesis, so we define:

$$\tilde{\lambda}(\mu) = \begin{cases} \frac{\mathcal{L}(\mu, \hat{\theta}_\mu)}{\mathcal{L}(\hat{\mu}, \hat{\theta}_\mu)} & \hat{\mu} \geq 0 \\ \frac{\mathcal{L}(\mu, \hat{\theta}_\mu)}{\mathcal{L}(0, \hat{\theta}_0)} & \hat{\mu} < 0 \end{cases} \quad (16)$$

We then construct a maximally powerful test statistic:

$$\tilde{q}_\mu = \begin{cases} -2 \ln \tilde{\lambda}(\mu) & \hat{\mu} \leq \mu \\ +2 \ln \tilde{\lambda}(\mu) & \hat{\mu} > \mu \end{cases} \quad (17)$$

We expect this test statistic to be Gaussian distributed in the asymptotic limit as is discussed in [31]. We can then write the probability that the data agree with a signal hypothesis of μ as:

$$p_\mu = \begin{cases} 1 - \Phi(-\sqrt{-\tilde{q}_\mu}) & \tilde{q}_\mu < 0 \\ 1 - \Phi(\sqrt{\tilde{q}_\mu}) & 0 \leq \tilde{q}_\mu \leq \mu^2/\sigma^2 \\ 1 - \Phi(\frac{\tilde{q}_\mu + \mu^2/\sigma^2}{2\mu/\sigma}) & \mu^2/\sigma^2 < \tilde{q}_\mu \end{cases} \quad (18)$$

where Φ is again the Gaussian CDF. We then write the probability that the data agree with the hypothesis of a signal rate of μ , as is prescribed in [33]:

$$p_b = \begin{cases} \Phi(-\sqrt{-\tilde{q}_\mu} - \mu/\sigma) & \tilde{q}_\mu < 0 \\ \Phi(\sqrt{\tilde{q}_\mu} - \mu/\sigma) & 0 \leq \tilde{q}_\mu \leq \mu^2/\sigma^2 \\ \Phi(\frac{\tilde{q}_\mu - \mu^2/\sigma^2}{2\mu/\sigma}) & \mu^2/\sigma^2 < \tilde{q}_\mu \end{cases} \quad (19)$$

We then continue to follow the prescription from [33] and define:

$$\text{CL}_s(\mu) = \frac{p_\mu}{1 - p_b} \quad (20)$$

The upper limit on the signal rate is chosen to be the particular μ such that $\text{CL}_s(N_{\text{sig}}^{\text{up}}) = 0.05$.

B. Systematic Uncertainties

There are two categories of uncertainties in this analysis: the uncertainty of our estimate of the mass resolution and that in estimating the radiative fraction. The two main contributors to the mass resolution uncertainty are our understanding of the target position and the momentum resolution of the apparatus. We estimate the uncertainty due to the mass resolution by varying the smearing coefficients extracted to replicate the mass resolution observed at the Møller mass according to their statistical uncertainties. We simulate the experiment with the target position at $\pm 0.5\text{mm}$ and compare the resulting mass distributions. We then add the two uncertainties in quadrature at each mass independently and choose the largest uncertainty across the entire spectrum, which is 3.4%. We account for this uncertainty by performing the

final fit to the data 10,000 times while varying the signal shape width by this amount and selecting the 84% quantile of the results.

The uncertainty of the radiative fraction has two contributions, from mismodeling the detector in MC and from uncertainties in the cross-sections used to scale the rate of each of the components of the radiative fraction. Efficiencies, momentum resolution, and acceptance of the final selection were varied in MC simulations to study the detector mismodeling uncertainty contribution. It was found that these effects introduce an uncertainty less than 1% on the radiative fraction. The first component of uncertainty from cross-section scaling is from the uncertainty in their evaluation by MadGraph, which we evaluated to be roughly 7% in total. The last component of uncertainty comes from our modeling of the rate of accidental track coincidences. After adding this in quadrature with the uncertainty from our evaluation of cross-sections the total uncertainty on the radiative fraction is determined to be at most 7.4%. This is accounted for by simply scaling the radiative fraction down by this amount.

C. Results

We calculated p-values in the mass range $m(e^+e^-) \in (39\text{ MeV} - 179\text{ MeV})$ with the method described in Section IV A 1. Our search for a resonance failed to reject the null hypothesis at every searched mass point. The smallest local p-value = 6.38×10^{-3} is observed for the mass $m(e^+e^-) = 94\text{ MeV}$. After accounting for the look-elsewhere effect [32], the global p-value (Figure 26) corresponds to about 1σ . Figure 26 shows the p-values for




FIG. 26. The local p-values produced by a resonance search analysis of 100% of the HPS 2016 data set. Here, the green line represents the global 1σ threshold, the orange line represents the global 2σ threshold, and the red line represents the global 3σ threshold. Mass resolution systematics are included in this plot.

the searched mass hypotheses. For each mass hypothesis,

we then calculate upper limits on ϵ^2 (Figure 27) with the method described in Section IV A 3. Figure 27 shows the




FIG. 27. The ϵ^2 upper limit produced by the resonance search analysis of 100% of the HPS 2016 data set, including all systematic uncertainty effects. The green band represents the 68% quantile range while the orange band represents the 95% quantile range.

upper limit results from the HPS 2016 data set and includes all systematic uncertainty effects discussed in Section IV B. The green band represents the 68% quantile range while the orange band represents the 95% quantile range of limits set on an ensemble of background-only simulation. This analysis of the e^+e^- mass distribution in the range 39 MeV to 179 MeV did not yield any statistically significant variations from the background-only hypothesis; therefore we report upper limits on ϵ^2 in the searched mass range. The integrated luminosity of the reported run is insufficient to cover new territory in the “ ϵ^2 vs mass” A' parameter space; however, the results are consistent with the results reported by previous experiments sensitive to this range in mass and coupling of dark photons.

V. DISPLACED VERTEX SEARCH

The goal of the displaced vertex analysis is to search for long-lived A' s produced in the target that decay to e^+e^- pairs in the range 1–10 cm downstream. These rare signal processes must be distinguished from a large number of prompt QED tridents which can occasionally appear to originate downstream of the target because of detector resolution, scattering effects, or tracking errors. Consequently, this search is limited by the vertex position resolution of HPS, anomalous scatters, and the quality of the tracking. For incident electron energies of a few GeV, the vertex position resolution is dominated by multiple scattering in the tracker, particularly in the first layers.

The basic principle of the analysis is illustrated in Figure 28, which shows the vertex distribution (in the coordinate along the beam axis, z) for reconstructed e^+e^-




FIG. 28. The vertex distribution along the beam direction for the full data set (black) and a simulated 105 MeV A' with a uniform decay in z (blue) in a 105 ± 4.7 MeV mass slice. See text for details.

pairs in the invariant mass slice of 105 ± 4.7 MeV. The black distribution shows data, which are composed entirely of prompt backgrounds. The blue distribution shows the shape of the acceptance from a simulated 105 MeV A' , assuming a decay uniform in z . The actual normalization and decay distribution of the A' distribution is dependent on ϵ^2 , and is, in general, very small compared to the background. Note that the background is well characterized by a Gaussian peak centered on the target location, with an exponential tail on its high side. The search is conducted at values of z beyond which 0.5 background events are expected from an exponential fit to the tail, which we call z_{cut} . Since a near-zero background region is necessary to search for a very low signal rate, every decay downstream of z_{cut} (the yellow region) is considered as a signal candidate. This search is performed using mass slices over the entire mass range considered in this analysis.

The following sections describe the event selection, analysis technique, and results of the displaced vertex search.

A. Event Selection

In addition to the cuts that select V0s described in Section III C above, the displaced vertex search, which depends critically on tracking, imposes several additional cuts on track and vertex quality. For both the electron and positron, the difference between the track time and the associated cluster time is required to be less than 4 ns, to reduce accidental backgrounds. Tracks are required to have a $\chi^2/\text{d.o.f.} < 6$ along with a minimum momentum of 0.4 GeV to eliminate those that arise from particles that suffer very large hard scattering in the tracker. Electron tracks are required to have a momentum magnitude less than 1.75 GeV in order to remove contamination from

full energy electrons (i.e. electrons that scatter elastically off the nucleus of the tungsten target) whereas positron tracks have no such requirement. The unconstrained vertex fit is required to have $\chi^2 < 10$ to reduce e^+e^- pairs that are inconsistent with originating from a single decay vertex.

The final set of cuts, described in the next section, is imposed to separate the prompt background that falsely reconstructs downstream of the target from true long-lived particles. These cuts are aimed at eliminating nearly all backgrounds arising from prompt sources, leaving a clean signal region beyond the z_{cut} .

B. Reducing High z Backgrounds

To reduce prompt backgrounds that reconstruct at large z , the so-called “high z background”, additional cuts beyond the event selection cuts described above must be employed. Most of the high z background results from a prompt track scattering in layers 1 or 2 of the tracker (both the active and inactive detector material) or from mis-reconstructed tracks. There are several handles that can be used to distinguish between a true displaced vertex and a high z background. In general, a true displaced vertex will have a good vertex χ^2 ; will project back to the beam spot; and will be comprised of tracks that each have large vertical impact parameters. These conditions are rarely true for high z backgrounds. In addition, to guard against high z backgrounds due to mis-tracking, the so-called “isolation cut”, described below, is implemented. All these cuts have been designed to eliminate most high z background events while having minimal impact on the efficiency to detect the A' signal. They were tuned using a 10% sample of the data.

An A' with a relatively short decay length will have L1 hits for both daughter particles, whereas an A' with a longer decay length may have one or both of these particles miss L1 due to geometrical acceptance effects as shown in Figure 29. For prompt processes, two effects may cause particles to “miss” the first layer. First, hit detection inefficiencies in L1 may cause particles to be undetected even though the particle traverses the active sensor plane. Second, particles from the target can interact with or convert in the inactive material in L1, resulting in no L1 hit, but scatter into the acceptance of the downstream layers and be detected. These effects are illustrated in Figure 30. Consequently, the analysis is divided into several mutually exclusive categories based on which layer has the first hit for each of the two daughter particles. If both particles have an L1 hit, the event is placed in the so-called “L1L1” category. If exactly one particle hits L1 and the other particle misses L1 but hits L2, the event is placed in the “L1L2” category. If both particles miss L1 but have their first hits in layer 2, the event is placed in the “L2L2” category. These are the only three possible categories since the tracking algorithm requires at least 5 hits on a track in the 6 layer




FIG. 29. Top: A schematic of a relatively short A' decay length in which both daughter particles have a layer 1 hit. This is referred to as the “L1L1” category. Bottom: A schematic of a relatively long A' decay length in which one of the daughter particles misses layer 1 (but hits L2) and the other daughter particle hits layer 1. This is referred to as the “L1L2” category.

SVT. For the purposes of this analysis, only the L1L1 and L1L2 categories are used. The probability of A' s populating the L2L2 category requires such long lifetimes and correspondingly low rates that much more luminosity is required to see them. The L2L2 category will add significance to future analyses that incorporate detector upgrades and have improved luminosity.

The L1L1 and L1L2 categories are analyzed separately for several reasons. First, the vertex position resolution is highly dependent on which layer is hit first. The closer the first hit is to the target, the better the vertex position resolution. Second, the nature of the backgrounds varies in the different categories. In the L1L1 category high z backgrounds are typically due to mis-tracking and multiple scattering in the active region of L1 sensors, whereas backgrounds in the L1L2 and L2L2 categories are typically due to hit inefficiency effects, multiple scattering in both active and inactive regions of L1, converted WABs, mis-tracking, and even trident production in L1.

The following cuts are implemented for both L1L1 and L1L2 to reduce the high z backgrounds.




FIG. 30. Top: A schematic of a prompt background process that has a hit inefficiency in layer 1 and is placed in the L1L2 category. Bottom: A schematic of a prompt background process in which one of the daughter particles scatters away from the beam in the inactive silicon of layer 1 and into the acceptance of the tracker. This process is placed in the L1L2 category and also reconstructs a false vertex downstream of the target.

1. V0 Projection to the Target

The V0 position is projected back to the target location at $z = -4.3$ mm using the V0 momentum vector direction. There its x - y coordinates are compared to those of the beam spot. The position of the beam center is corrected for run-by-run variations and then the projected vertex position is required to be within a 2σ elliptical region of the mean beam position in x - y space.

2. Isolation Cut

Mis-reconstructed tracks are tracks that contain at least one hit that is not created by the particle responsible for the majority of the other hits on the track. For instance, a track can reconstruct a real particle trajectory but include a spurious hit from a beam electron, recoil electron, converted photon, or noise hit. When this mis-reconstructed hit is in L1 and is closer to the beam than

the true hit, it can result in a vertex that appears downstream of the target, often significantly downstream, i.e., one that appears signal-like.

Mis-reconstructed hits in L1 often occur as a result of scattering in L2; those in L2, if L1 is missing, from scattering in L3. Such scattering can cause the track to extrapolate to the incorrect hit and occurs at a significant enough rate that it needs to be mitigated. The isolation cut provides a simple test to see if substituting a nearby hit in the innermost layer would give a track that is more consistent with coming from the beam spot, and is thus more likely the correct hit. If such a hit is found, the event is eliminated. The isolation cut compares the distance between the hit associated with the track and that closest to it in the direction away from the beam, called the isolation value δ_{iso} , to the track vertical impact parameter y_0 , as shown in Figure 31 for the L1L1 case. Multiple scattering and beam size effects complicate this cut. Reconstruction errors on the impact parameter Δy_0 are comparable to the beam size (both $\sim 100 \mu\text{m}$) so both must be accounted for. The isolation cut for the L1L1 category and for L1L2 tracks that pass through L1 is as follows:

$$\delta_{iso} + \frac{1}{2} (y_0 - n_\sigma \Delta y_0) > 0. \quad (21)$$

For the L1L2 tracks that miss L1, it is:

$$\delta_{iso} + \frac{1}{3} (y_0 - n_\sigma \Delta y_0) > 0. \quad (22)$$

The factor of $1/2$ in Equation (21) is the ratio of the distance between the first two layers and the distance between L2 and the target. The factor of $1/3$ in Equation (22) is the ratio of the distance between L2 and layer 3 and the distance from layer 3 to the target, appropriate when the first hit is in L2. Multiple scattering is taken into account with the error term n_σ , where n_σ is selected to be 3 and Δy_0 is the combination of the projected impact parameter resolution and the beam size. A Monte Carlo study shows that the cut eliminates most high z background due to mis-tracking but has minimal impact on signal efficiency.

3. Impact Parameter Cut

For the A' signal, a true displaced vertex will have large vertical impact parameters (y_0) for both its electron and positron tracks. Furthermore, these impact parameters are correlated with z , increasing with increasing z . For prompt background that reconstructs at large z , this is usually not the case. Instead, it is likely that just one particle has a large scatter away from the beam plane (and thus a large impact parameter) and the other particle is either consistent with coming from the beam spot or has an impact parameter smaller than is expected from signal. With a cut on the impact parameters of both e^+e^-




FIG. 31. A geometric picture of the isolation cut comparing the distance between the nearest hit away from the beam δ_{iso} and the longitudinal impact parameter of the track y_0 . The correct track is in blue and the incorrect track found by the tracking algorithm is in dashed purple. This can result in a falsely reconstructed vertex downstream of the target.




FIG. 32. Prompt background that falsely reconstructs at a large z due to an e^- particle with a large scatter away from the beam plane in L1 of the SVT. The corresponding e^+ does not have a large scatter and the track points back near the primary. A cut on the impact parameter can eliminate such backgrounds.

tracks, such backgrounds can be eliminated. This concept is illustrated in Figure 32. The impact parameters for signal display correlated bands in the y_0 - z space, average y_0 increasing as z does. This correlation is approximately linear for the masses of interest in this analysis, so the cut depends linearly on z . Since larger mass A' s have larger decay angles on average, they will also have larger impact parameters. Thus, the cut is also parameterized as a function of mass. Both the electron and positron are required to satisfy the impact parameter condition. Before imposing the cut, the y position of the beam is corrected for changing beam conditions.

4. Removing Tracks with Shared Hits and Selecting Single $V0$ s

The last step is to remove both tracks with shared hits and events with multiple $V0$ particles. In the reconstruction, tracks are allowed to share hits with other tracks, and these shared hits can be from hits from another particle. There is evidence in both data and MC that tracks with shared hits may produce high z background events. To eliminate this possibility, tracks that share any hits with any other track are eliminated. The final requirement is that each event must have exactly one $V0$ candidate that passes all previous cuts. This will prevent there being multiple candidate vertices in an event, which is extremely unlikely *a priori*.

The complete cut flow for all background reduction cuts for the L1L1 (L1L2) category is shown in Figure 33 (Figure 35). The resulting reconstructed z vs. mass for events in the L1L1 (L1L2) category is shown in Figure 34 (Figure 36). Note that no mass bins in either plot show significant concentrations of events beyond the z_{cut} . Further note the nearly complete absence of any events at large decay lengths, beyond 50 mm for L1L1 and out to 75 mm for L1L2, although there is acceptance in these regions.




FIG. 33. Top: The cut flow for data in the L1L1 category. Bottom: The cut flow for 80 MeV displaced A' MC in the L1L1 category.




FIG. 34. The final selection for the L1L1 category is plotted as reconstructed z vs reconstructed e^+e^- mass for the full data set. The black line shows the value of z_{cut} versus mass and the yellow shaded region is, roughly, the region of sensitivity to A' events.

C. Defining the Signal Region

Because of the low rate of the expected signal, a signal region must be defined such that very little background is expected. Both the background and signal fall exponentially in the z -direction; however, the background falls at a much faster rate. Thus a nearly zero background region can be found downstream of a sufficiently large z value. Specifically, this is done as a function of mass since a signal is expected at a specific invariant mass and the vertex position resolution is dependent on the opening angle, and hence mass-dependent. With this in mind, the z vs mass distribution is sliced into overlapping bins of width equal to $\pm 1.9\sigma_m(m)$ for a mass m in the bin center. Each mass slice is fitted in z using the following continuous and differentiable empirical function consisting of Gaussian core and exponential tail.

$$F(z) = \begin{cases} Ae^{-\frac{(z-\mu_z)^2}{2\sigma_z^2}} & \text{if } \frac{z-\mu_z}{\sigma_z} < b \\ Ae^{-\frac{b^2}{2}-b\frac{z-\mu_z}{\sigma_z}} & \text{if } \frac{z-\mu_z}{\sigma_z} \geq b. \end{cases} \quad (23)$$

The parameter b is the number of standard deviations from the mean that the fit function changes from a Gaus-

sian to an exponential tail. All of A , μ_z , σ_z , and b are determined by the fit for each mass slice.

Using the results of the fit function, the z value beyond which the background fit function predicts 0.5 background events defines the z_{cut} . Or more precisely:

$$0.5 = \int_{z_{\text{cut}}}^{\infty} F(z) dz. \quad (24)$$

After this fit is performed in every mass slice and the z_{cut} is found, the final z_{cut} in a given mass slice is found by fitting the z_{cut} distribution as a function of mass without the points in the mass bin of interest in order to be unbiased (i.e. using the mass sidebands). An example background fit to a mass slice in the full data set is shown in Figure 28 and the z_{cut} for both the L1L1 and L1L2 categories, along with the signal region in yellow, is displayed on Figure 34 and Figure 36, respectively.

D. Computing the Expected Signal Yield

Computing the expected rate of A' s in the signal region takes careful consideration of several z -dependent factors including decay length distributions (as a function of the




FIG. 35. Top: The cut flow for data events in the L1L2 category. Bottom: The cut flow for 80 MeV displaced A' events in the L1L2 category.

model parameters), detector acceptance effects, and efficiency effects. The first step is to compute the truth signal distribution for long-lived A' s, which is exponential in z . The normalized truth signal shape as a function of $c\tau$ is an exponential given by:

$$S_{\text{truth}}(z, m_{A'}, \epsilon) = \frac{e^{-(z_{\text{targ}} - z)/\gamma c\tau}}{\gamma c\tau} \quad (25)$$

This function is normalized such that the integral from z_{targ} , the z position of the target, to infinity is unity so that it gives the expected signal density distribution (i.e. $\int_{z_{\text{targ}}}^{\infty} S_{\text{truth}}(z, m_{A'}, \epsilon) dz = 1$). In this equation, $\gamma = \frac{E}{m_{A'}}$ is the relativistic constant where the A' energy is computed to be $E = 0.965 E_{\text{beam}}$ (which is the mean of the x distributions across all relevant masses).

After computing the truth distributions, detector acceptance must be taken into account. The SVT is designed to have a θ_y acceptance beyond 15 mrad for prompt decays. However, downstream decays must have a larger opening angle to remain in the acceptance of the SVT, and the farther downstream the decay, the more likely the daughter particles will miss the SVT. The geometrical acceptance drops dramatically with increasing decay length as is shown in Figure 37. The geometrical acceptance cannot be measured in data and must be

derived from simulation.

Finally, putting this all together and integrating the signal shape across the z range of interest gives the formula for the expected signal past z_{cut} as a function of mass and ϵ denoted as $S_{\text{bin}, z_{\text{cut}}}(m_{A'}, \epsilon)$.

$$S_{\text{bin}, z_{\text{cut}}}(m_{A'}, \epsilon) = S_{\text{bin}}(m_{A'}, \epsilon) \times \int_{z_{\text{targ}}}^{z_{\text{max}}} S_{\text{truth}}(z, m_{A'}, \epsilon) \epsilon_{\text{vtx}}(z, m_{A'}) dz. \quad (26)$$

In this equation, $S_{\text{bin}}(m_{A'}, \epsilon)$ is the expected signal yield within prompt acceptance and within a finite mass bin computed from the number of e^+e^- pairs and the radiative fraction, shown in Figure 38, calculated for the specific event selection used in the displaced vertex search (see Section III C). Additionally, $\epsilon_{\text{vtx}}(z, m_{A'})$ is the normalized efficiency as a function of z including acceptance and all efficiency effects (including analysis cuts and the z_{cut}). The value z_{max} is the minimum z beyond which signal is not expected.¹ The expected A' rates computed with this equation are used as an input to set the final limit in Section V F.

The expected signal yields for the data set in the L1L1 and L1L2 categories are shown in Figure 39. For the L1L1 category, a peak of 0.32 A' events is expected at 75 MeV A' mass and $\epsilon^2 = 2.4 \times 10^{-9}$ while for the L1L2 category, a peak of 0.22 A' events is expected at 75 MeV A' and $\epsilon^2 = 1.7 \times 10^{-9}$. Adding the two categories, the yield peaks at 75 MeV and $\epsilon^2 = 2.1 \times 10^{-9}$ with 0.52 expected A' events as shown in Figure 40. This shows that for suitable parameters, the HPS sensitivity is closely approaching that needed to exclude some parameter space of canonical A' production. In the following section, we will discuss the systematic uncertainties and the procedure used to set upper limits.

¹ Note that the integral is taken starting from the target and not the z_{cut} since the z_{cut} is already applied as an analysis cut at this point. This is done because the z in the integral is a truth value while the z_{cut} is a reconstructed value derived from data.




FIG. 36. The final selection for the L1L2 category is plotted as reconstructed z vs reconstructed e^+e^- mass for the data set. The black line shows the value of z_{cut} versus mass and the yellow shaded region is, roughly, the region of sensitivity to A' events.




FIG. 37. The product of geometrical acceptance and efficiency for displaced A' MC for the L1L1 and L1L2 categories as well as their sum for 80 MeV displaced A' 's.




FIG. 38. Estimate of the fraction of radiative events as a function of invariant mass. This distribution is obtained from MC.




FIG. 39. Top: The expected detected A' yield in the L1L1 category versus mass and epsilon. Bottom: The expected detected A' yield in the L1L2 category versus mass and epsilon.




FIG. 40. The expected detected A' yield for the combination of both L1L1 and L1L2 categories versus mass and epsilon.

E. Systematic Uncertainties

Systematic Description	L1L1 Value	L1L2 Value
e^+e^- Composition	$\sim 7\%$	
Mass Resolution	$\sim 3\%$	
Analysis Cuts	$\sim 8\%$	$\sim 13\%$
A' Efficiency	$\sim 5\%$	
Total in Quadrature	12%	16%
Target position	$\sim 5\text{-}10\%$ (m/ ϵ dep)	

TABLE I. A summary of systematic uncertainties that impact the final result of the displaced vertex search. Where there is a single number the systematic effect is the same for L1L1 and L1L2.

The systematic uncertainties from the experiment and the displaced vertex analysis have been quantified for both the L1L1 and L1L2 samples and are summarized in Table I. These sources of uncertainties are described below.

A source of systematic uncertainty that is shared with the resonance search is the uncertainty in the e^+e^- composition that is expressed in the error of the radiative fraction. See Section IV B for details.

An underestimate of the mass resolution would result in signal leaking out of a mass bin. Thus, uncertainty in the mass resolution is a source of systematic uncertainty. As described in Section III E, we obtain the mass resolution as a function of A' mass using A' MC which has the e^+ and e^- momenta smeared by the data/MC ratio of FEE resolutions. As a cross-check, we do the same for Møller MC and compare that to the resolution seen in Møller data. The Møller comparison gives very good agreement between data and MC, with the data having only a 5% higher resolution compared to MC. We use this 5% seen in the Møller samples to estimate a systematic on the number of signal events due to the mass cut and find that it is $\sim 3\%$.

There are systematic uncertainties associated with the analysis cuts, particularly the cuts to reduce high- z background (see Section V B). Recall that we use the radiative fraction to normalize the rates at event selection level while the relative efficiency from going from event selection to the final selection is accounted for using A' MC. There are small differences in the MC and data efficiencies of the final cuts and these have to be accounted for as systematic scaling errors.

To do this, we calculate the efficiencies of each cut, with all other cuts applied, for data and trident MC events and take the ratio as the relative scaling that must be applied to the final limits. There are only four categories of cuts to consider: the V0 projection to the target, the isolation cuts, the impact parameter cuts, and the shared hits cuts. The results of this study give the product of the efficiency ratios (data/MC) for L1L1 (~ 0.92) and L1L2 (~ 0.88). The inverse of these ratios is applied

to the final limits.

From mechanical measurements, the target position is estimated to be known within ± 0.5 mm from the nominal position. Any change in the assumed target position will result in an overall shift in truth z distributions of displaced A' 's, and thus is a source of systematic uncertainty. For example, if the target is 0.5 mm more upstream than assumed, the entire displaced A' truth distribution will also shift upstream by 0.5 mm (without changing z_{cut}) resulting in the actually expected signal yield that is less than the calculated signal yield. For a given A' mass, this discrepancy will depend significantly on ϵ because of varying decay length, and can be calculated by simply recomputing both the signal yield and the limit at a different target position (± 0.5 mm). The ratio of the limit from a target at 0.5 mm upstream of the nominal position to the target at the nominal position is shown in Figure 41. This mass and ϵ dependence are used in the final estimate of systematic uncertainties.




FIG. 41. The ratio of the limit for the L1L1 category from the target 0.5 mm upstream of the nominal position to the target at the nominal position using the Optimum Interval Method.

We combine the ϵ /mass independent systematic uncertainties in quadrature and then combine those in quadrature at each combination of ϵ and mass to obtain a map of the uncertainty vs ϵ /mass. This uncertainty is then used to scale the upper limits we obtain from the data.

F. Upper Limit on A' Rate

The Optimum Interval Method (OIM) [34] is used to set a limit on the cross-section of the canonical A' model. OIM was originally developed for direct detection dark matter experiments in which one expects a small signal where the signal shape in one variable is known and there is a small, but not necessarily understood, background. The OIM is an extension of the Maximum Gap Method, which searches for the largest gap in signal space that has no background events in order to set a limit. The OIM generalizes this method to an arbitrary number of background events between any two events in signal space and

sets a limit based on the optimum interval and automatically selects the interval to avoid experimenter bias. In addition, the absolute cross-section of the signal does not need to be known. Instead the OIM finds the optimum interval and sets a limit at the smallest cross-section at a specified confidence interval C_0 , 90% for this analysis.

The results for the OIM for the L1L1 and L1L2 categories on the full dataset are shown in Figure 42. For the full dataset in the L1L1 category, the best limit is set at $m_{A'} = 80.2$ MeV and $\epsilon^2 = 2.1 \times 10^{-9}$ with a factor of 9.1 times the canonical A' cross-section. The interpretation of this value is for an A' -like model with 9.1 times the cross-section. The model is excluded at that mass and ϵ^2 with 90% confidence. For the L1L2 category the best limit is at $m_{A'} = 69.2$ MeV and $\epsilon^2 = 1.9 \times 10^{-9}$ with a factor of 13.9 times the canonical A' cross-section. These results include the systematic uncertainties described in Section V E.

The limits derived when the L1L1 and L1L2 categories are combined are shown in Figure 43. Combining the L1L1 and L1L2 categories gives the best limit at $m_{A'} = 82.0$ MeV and $\epsilon^2 = 1.7 \times 10^{-9}$ with a factor of 7.9 times the canonical A' cross-section. With the current luminosity it is not possible to set upper limits on canonical A' production in the parameter plane.




FIG. 42. Top: The limit from the Optimum Interval Method for the L1L1 category. Bottom: The limit from the Optimum Interval Method for the L1L2 category.




FIG. 43. The limit from the Optimum Interval Method for the combination of both L1L1 and L1L2 categories for the full dataset.

VI. SUMMARY

This paper has presented the HPS results from its 2.3 GeV 2016 engineering run. Evidence for heavy photons was searched for with both resonance search and vertex search techniques. Our previous resonance search results, from the 1.06 GeV 2015 engineering run, have been updated to use a more modern statistical approach. The 2016 data have extended the coverage in heavy photon mass to 180 MeV in the resonance search, and exclude canonical A' production over the mass range 40–180 MeV down to the level of $\epsilon^2 = 10^{-5}$ as shown in Figure 44.

The resonance search result confirms the results of previous searches but does not extend their sensitivity. The vertex search, reported here for the first time, explores A' masses in the range 60–150 MeV/ c^2 for ϵ^2 in the region 10^{-8} to 10^{-10} . This is virgin parameter space, so far unexplored by other experiments, which is preferred territory for models assuming thermal production of hidden sector dark matter during the Big Bang. Being statistically limited, the present search does not reach the sensitivity needed to see canonical A' production in this region, but it does, at its point of optimal sensitivity, exclude production of long-lived e^+e^- pairs with 7.9 times the expected heavy photon cross-section and has afforded a first sensitive search for e^+e^- secondary vertices in electro-production at low energy. At its peak sensitivity in mass and ϵ^2 , the experiment would have expected to see 0.5 A' events, so it is approaching the sensitivity needed for a canonical A' search. Over much of the range in A' mass, backgrounds were controlled to a

level that should allow future vertex searches, with significantly greater luminosity, to explore interesting regions of parameter space. HPS data runs in 2019 and 2021 have acquired this needed luminosity, and we project sensitivity to canonical A' production over a range of mass and ϵ^2 parameters when those data are fully analyzed.

VII. ACKNOWLEDGEMENTS

The authors are grateful for the outstanding efforts of the Jefferson Laboratory Accelerator Division, the Hall B engineering group, and Forest McKinney of UC Santa Cruz in support of HPS. The research reported here is supported by the U.S. Department of Energy Office of Science, Office of Nuclear Physics, Office of High Energy Physics, the French Centre National de la Recherche Scientifique, United Kingdom’s Science and Technology Facilities Council (STFC), the Sesame project HPS@JLab funded by the French region Ile-de-France and the Italian Istituto Nazionale di Fisica Nucleare. Jefferson Science Associates, LLC, operates the Thomas Jefferson National Accelerator Facility for the United States Department of Energy under Contract No. DE-AC05-06OR23177.




FIG. 44. The exclusion of dark photon parameter space by this analysis and statistical recasting of our 2015 result. Existing limits from beam dump [35–43], collider [44–49] and fixed target experiments [50–53] are also shown. The region labeled “ a_e ” is an exclusion based on the electron $g - 2$ [54–57]. The green band labeled “ $a_\mu \pm 2\sigma$ ” represents the region that an A' can be used to explain the discrepancy between the measured and calculated muon anomalous magnetic moment [58, 59].

- [1] *Fundamental Physics at the Intensity Frontier* (2012) arXiv:1205.2671 [hep-ex].
- [2] R. Essig *et al.*, in *Community Summer Study 2013: Snowmass on the Mississippi* (2013) arXiv:1311.0029

- [hep-ph].
- [3] J. Alexander *et al.* (2016) arXiv:1608.08632 [hep-ph].
- [4] M. Battaglieri *et al.*, in *U.S. Cosmic Visions: New Ideas in Dark Matter* (2017) arXiv:1707.04591 [hep-ph].

- [5] B. Holdom, Phys. Lett. B **166**, 196 (1986).
- [6] P. Galison and A. Manohar, Phys. Lett. B **136**, 279 (1984).
- [7] P. Fayet, Nucl. Phys. B **347**, 743 (1990).
- [8] J. D. Bjorken, R. Essig, P. Schuster, and N. Toro, Phys. Rev. D **80**, 075018 (2009).
- [9] N. Arkani-Hamed, D. P. Finkbeiner, T. R. Slatyer, and N. Weiner, Phys. Rev. D **79**, 015014 (2009), arXiv:0810.0713 [hep-ph].
- [10] N. Arkani-Hamed and N. Weiner, JHEP **12**, 104 (2008), arXiv:0810.0714 [hep-ph].
- [11] M. Baumgart, C. Cheung, J. T. Ruderman, L.-T. Wang, and I. Yavin, JHEP **04**, 014 (2009), arXiv:0901.0283 [hep-ph].
- [12] R. Essig, P. Schuster, and N. Toro, Phys. Rev. D **80**, 015003 (2009), arXiv:0903.3941 [hep-ph].
- [13] C. Cheung, J. T. Ruderman, L.-T. Wang, and I. Yavin, Phys. Rev. D **80**, 035008 (2009), arXiv:0902.3246 [hep-ph].
- [14] D. E. Morrissey, D. Poland, and K. M. Zurek, JHEP **07**, 050 (2009), arXiv:0904.2567 [hep-ph].
- [15] P. H. Adrian *et al.* (HPS), Phys. Rev. D **98**, 091101 (2018), arXiv:1807.11530 [hep-ex].
- [16] C. Leemann, D. Douglas, and G. Krafft, Ann. Rev. Nucl. Part. Sci. **51**, 413 (2001).
- [17] N. Baltzell *et al.* (HPS), Nucl. Instrum. Meth. A **859**, 69 (2017), arXiv:1612.07821 [physics.ins-det].
- [18] M. Battaglieri *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **777**, 91 (2015).
- [19] V. M. Abazov *et al.* (D0), FERMILAB-PUB-02-327-E, FERMILAB-DESIGN-2002-01 (2002).
- [20] M. J. French *et al.*, Nucl. Instrum. Meth. A **466**, 359 (2001).
- [21] R. Herbst *et al.*, in *Proceedings, 21st Symposium on Room-Temperature Semiconductor X-ray and Gamma-ray Detectors (RTSD 2014): Seattle, WA, 2014* (2016) p. 7431254.
- [22] I. Balossino *et al.* (HPS), Nucl. Instrum. Meth. A **854**, 89 (2017), arXiv:1610.04319 [physics.ins-det].
- [23] J. Alwall *et al.*, Journal of High Energy Physics **2014**, 79 (2014).
- [24] H. Hirayama, Y. Namito, A. F. Bielajew, S. J. Wilderman, and W. R. Nelson, (2005).
- [25] S. Agostinelli *et al.* (GEANT4), Nucl. Instrum. Meth. A **506**, 250 (2003).
- [26] C. Kleinwort, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **673** (2012), 10.1016/j.nima.2012.01.024.
- [27] V. Blobel, NIM A, 5 (2006).
- [28] P. Billoir and S. Qian, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **311**, 139 (1992).
- [29] W. D. Hulsbergen, Nucl. Instrum. Meth. A **552**, 566 (2005), arXiv:physics/0503191.
- [30] G. Aad *et al.* (ATLAS), Phys. Rev. D **86**, 032003 (2012), arXiv:1207.0319 [hep-ex].
- [31] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, The European Physical Journal C **71**, 1554 (2011).
- [32] E. Gross and O. Vitells, The European Physical Journal C **70**, 525–530 (2010).
- [33] T. Junk, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **434**, 435 (1999).
- [34] S. Yellin, Phys. Rev. D **66**, 032005 (2002).
- [35] J. D. Bjorken, S. Ecklund, W. R. Nelson, A. Abashian, C. Church, B. Lu, L. W. Mo, T. A. Nunamaker, and P. Rassmann, Phys. Rev. **D38**, 3375 (1988).
- [36] E. M. Riordan *et al.*, Phys. Rev. Lett. **59**, 755 (1987).
- [37] A. Bross, M. Crisler, S. Pordes, J. Volk, S. Errede, and J. Wrbanek, Phys. Rev. Lett. **67**, 2942 (1991).
- [38] A. Konaka *et al.*, Phys. Rev. Lett. **57**, 659 (1986).
- [39] M. Davier and H. Nguyen Ngoc, Phys. Lett. **B229**, 150 (1989).
- [40] J. D. Bjorken, R. Essig, P. Schuster, and N. Toro, Phys. Rev. **D80**, 075018 (2009).
- [41] S. Andreas, C. Niebuhr, and A. Ringwald, Phys. Rev. D **86**, 095019 (2012).
- [42] J. Blumlein *et al.*, Z. Phys. **C51**, 341 (1991).
- [43] J. Blumlein *et al.*, Int. J. Mod. Phys. **A7**, 3835 (1992).
- [44] M. Reece and L.-T. Wang, JHEP **07**, 051 (2009), arXiv:0904.1743 [hep-ph].
- [45] B. Aubert *et al.* (BaBar), Phys. Rev. Lett. **103**, 081803 (2009), arXiv:0905.4539 [hep-ex].
- [46] D. Babusci *et al.* (KLOE-2), Phys. Lett. **B720**, 111 (2013), arXiv:1210.3927 [hep-ex].
- [47] F. Archilli *et al.* (KLOE-2), Phys. Lett. **B706**, 251 (2012), arXiv:1110.0411 [hep-ex].
- [48] R. Aaij *et al.* (LHCb Collaboration), Phys. Rev. Lett. **120**, 061801 (2018), arXiv:1710.02867 [hep-ex].
- [49] R. Aaij *et al.* (LHCb Collaboration), Phys. Rev. Lett. **124**, 041801 (2020).
- [50] S. Abrahamyan *et al.* (APEX), Phys. Rev. Lett. **107**, 191804 (2011), arXiv:1108.2750 [hep-ex].
- [51] H. Merkel *et al.*, Phys. Rev. Lett. **112**, 221802 (2014), arXiv:1404.5502 [hep-ex].
- [52] G. Agakishiev *et al.* (HADES), Phys. Lett. **B731**, 265 (2014), arXiv:1311.0216 [hep-ex].
- [53] J. R. Batley *et al.* (NA48/2), Phys. Lett. **B746**, 178 (2015), arXiv:1504.00607 [hep-ex].
- [54] R. Bouchendira, P. Cladé, S. Guellati-Khélifa, F. m. c. Nez, and F. m. c. Biraben, Phys. Rev. Lett. **106**, 080801 (2011).
- [55] T. Aoyama, M. Hayakawa, T. Kinoshita, and M. Nio, Phys. Rev. Lett. **109**, 111807 (2012), arXiv:1205.5368 [hep-ph].
- [56] D. Hanneke, S. Fogwell, and G. Gabrielse, Phys. Rev. Lett. **100**, 120801 (2008).
- [57] H. Davoudiasl, H.-S. Lee, and W. J. Marciano, Phys. Rev. **D86**, 095009 (2012), arXiv:1208.2973 [hep-ph].
- [58] M. Pospelov, Phys. Rev. **D80**, 095002 (2009), arXiv:0811.1030 [hep-ph].
- [59] G. Bennett *et al.* (Muon G-2 Collaboration), Phys. Rev. **D73**, 072003 (2006), arXiv:hep-ex/0602035 [hep-ex].