SEARCHING FOR HEAVY PHOTONS IN THE HPS
EXPERIMENT

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Sho Uemura
November 2016

© 2016 by Sho Uemura. All Rights Reserved.

Re-distributed by Stanford University under license with the author.

This work is licensed under a Creative Commons Attribution-Noncommercial 3.0 United States License.

<http://creativecommons.org/licenses/by-nc/3.0/us/>

This dissertation is online at: <http://purl.stanford.edu/hs558pd7550>

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

John Jaros, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Patricia Burchat

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Michael Peskin

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumport, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

The Heavy Photon Search (HPS) is a new experiment at Jefferson Lab that searches for a massive $U(1)$ vector boson (known as a heavy photon or A') in the MeV—GeV mass range and coupling weakly to ordinary matter through a kinetic mixing interaction. The HPS experiment seeks to produce heavy photons by electron bremsstrahlung on a fixed target, is sensitive to heavy photon decays to e^+e^- , and targets the range in heavy photon mass $m_{A'} \sim 20 - 600$ MeV, and kinetic mixing strength $\epsilon^2 \sim 10^{-5} - 10^{-10}$. HPS searches for heavy photons using two signatures: a narrow mass resonance and displaced vertices.

This dissertation presents the theoretical and experimental motivations for a heavy photon, the design and operation of the HPS experiment, and the displaced vertex search.

The data used in this dissertation is the unblinded fraction of the 2015 HPS run, for the period of operation where the HPS silicon vertex tracker (SVT) was operated at its nominal position. This data was recorded from May 13 to May 18, 2015, at a beam energy of 1.056 GeV and a nominal beam current of 50 nA. The integrated luminosity is 119 nb^{-1} , which is equivalent to 0.172 days of ideal running at the nominal beam current.

This dissertation presents results (signal significance and upper limits) from the displaced vertex search in the mass range $m_{A'} \sim 20 - 60$ MeV, and kinetic mixing strength $\epsilon^2 \sim 2 \times 10^{-8} - 10^{-10}$. This search does not have sufficient sensitivity to exclude a canonical heavy photon at any combination of $m_{A'}$ and ϵ^2 . The strictest limit achieved in this analysis on the production of a particle that decays like a heavy photon is 115 times the expected production cross-section for a heavy photon.

Factors limiting the sensitivity of this analysis are discussed. Projections of HPS performance with the full 2015 data set, and with planned improvements to the analysis, are presented. Comparisons are also made to earlier reach estimates.

Acknowledgments

I would like to thank everyone who made this dissertation, and the HPS experiment, possible.

First, thanks to John Jaros for being a great advisor, and letting me choose my direction while always asking the right questions. I would also like to thank my other reading committee members, Pat Burchat and Michael Peskin, and my other orals committee members, Tony Heinz and Tom Shutt. Thank you all for your help with this dissertation, and for making me a better physicist.

Thanks to the whole HPS collaboration for making this experiment work. Thanks to Matt Graham and Takashi Maruyama for really understanding how the experiment and analysis should work, and passing that knowledge on to me. Thanks to Tim Nelson for teaching me, by example, how to build a tracker, and Matt McCulloch for teaching me how to not break things. Thanks to Stepan Stepanyan for getting us beam and making sure we had a detector ready for beam. Finally, thanks to Pelle Hansson, Omar Moreno, and Matt Solt for putting in the hours with me in the trenches, both at SLAC and at JLab.

I would also like to thank my friends and family for their support throughout graduate school. This dissertation is dedicated to my loving father.

Contents

Abstract	iv
Acknowledgments	vi
1 Introduction	1
2 Motivation	6
2.1 Theory Summary	7
2.2 Observations Motivating a Heavy Photon	8
2.2.1 Dark Matter Annihilation	8
2.2.2 Halo Structure	10
2.2.3 Muon $g - 2$ Anomaly	10
2.3 Signatures	11
2.4 Overview of Searches	12
2.4.1 Beam Dumps	13
2.4.2 Thin Fixed Targets	13
2.4.3 Colliders	14
2.5 Overview of HPS	15
2.5.1 Signal Kinematics	15
2.5.2 Physics Backgrounds	16
3 The HPS Detector	21
3.1 Beamline	22
3.1.1 Hall B Beamline and Instrumentation	24

3.1.2	Beam Quality and Protection	26
3.1.3	HPS Beamline Elements	28
3.2	Silicon Vertex Tracker	29
3.2.1	Sensors and Readout	33
3.2.2	Mechanical Support	35
3.2.3	Power and Data Acquisition	40
3.2.4	Services	44
3.3	Electromagnetic Calorimeter and Trigger	46
3.3.1	Readout and Trigger	48
3.3.2	Trigger Cuts	50
4	Event Reconstruction and Selection	53
4.1	Reconstruction	53
4.1.1	Tracking	54
4.1.2	Vertexing	57
4.1.3	Clustering	58
4.1.4	Track-Cluster Matching	60
4.2	Tracker Alignment	61
4.2.1	Internal Alignment	61
4.2.2	Elastically Scattered Electrons	63
4.2.3	Møller Electrons	64
4.3	e^+e^- Selection Cuts	65
4.3.1	Tuning Cuts	68
4.4	Run Quality, Event Quality, and Data Normalization	68
4.5	Rate Comparison with Monte Carlo	71
5	Search for Displaced Vertices	78
5.1	Heavy Photon Monte Carlo	82
5.2	Vertex Cuts	83
5.2.1	Layer 1 Requirement	85
5.2.2	Track Quality	85
5.2.3	Beamspot Constraint	86

5.2.4	Isolation Cut	87
5.2.5	Momentum Asymmetry and Positron DOCA	88
5.2.6	Tuning Cuts	89
5.3	Fit Inputs	92
5.3.1	Estimating the Mass Resolution	93
5.3.2	Estimating the Signal Distribution	98
5.3.3	Fitting the Background Distribution	102
5.4	Finding Signal Significance	109
5.5	Setting Limits	113
5.5.1	Optimum Interval Method	113
5.5.2	Results	116
6	Conclusion	126

List of Tables

3.1	Layout of the HPS SVT. The angle of stereo sensors is relative to the bend plane.	32
3.2	Specifications of the SVT sensors. Breakdown voltage specification is value accepted for use in HPS; procurement specification was looser. .	34
3.3	Trigger cuts for the “pairs-1” trigger.	52
4.1	Base pair selection cuts for HPS. The effect of these cuts is shown in Figure 4.3.	67
5.1	The full set of pair selection cuts for the vertexing analysis. Cuts carried over from the base selection are on top; cuts specific (or tightened) for vertexing are on bottom.	83
5.2	Assumptions used for the reach estimate in the 2014 proposal.	122
5.3	Assumptions used for the “proposal-like” reach estimate.	124

List of Figures

1.1	The yellow-green contours outline the expected reach (for exclusion at 90% confidence level) for the HPS experiment with the run plan shown. The reach is split in two regions; the upper region corresponds to the bump-hunt and the lower region corresponds to the vertexing search. Existing limits from other experiments are plotted as shaded regions. The favored region from the muon $g - 2$ anomaly is plotted as a shaded green band (upper left).	4
1.2	Rough totals for integrated charge and event count during the 2015 HPS run. The numbers in this plot are not fully corrected for run and event quality.	5
2.1	One-loop diagram leading to kinetic mixing. Φ and Φ' are particles that couple to both the photon and the heavy photon. They could be massive (far above the SUSY-breaking scale), but the resulting kinetic mixing will persist down to low mass scales.	8
2.2	Dark matter annihilation to heavy photons, which can then decay to Standard Model particles.	9
2.3	Heavy photon correction to the muon magnetic moment.	11
2.4	Branching ratios for the heavy photon decay to Standard Model particles [35]. The bottom plot shows the range of $m_{A'}$ to which HPS is sensitive.	18

2.5	Heavy photon production and decay. The heavy photon (A') can be produced as final state radiation (right) or initial state radiation (left). In both cases, the heavy photon is on-shell and can travel some distance before decaying in free space.	19
2.6	Radiative trident diagrams. These diagrams are also known as timelike tridents, because the photon radiated by the incident electron (the upper photon in each diagram) is timelike.	19
2.7	Bethe-Heitler trident diagrams. These diagrams are also known as spacelike tridents, because the photon radiated by the incident electron (the upper photon in each diagram) is spacelike.	20
3.1	View of the HPS setup. The beam direction is left to right.	22
3.2	Schematic of the CEBAF accelerator, highlighting components added for the 12 GeV upgrade.	23
3.3	Location of the HPS setup, in the downstream alcove of Hall B. The beam direction is left to right.	24
3.4	A representative wire scan measurement of the beam size. This scan was taken using the 2H02 wire scanner just upstream of HPS, and reflects the beam size seen at the HPS target. This wire scan was taken early in the 2015 run; the beam size in X was reduced later in the run, and all physics data was recorded with a beam size of roughly $150 \times 50 \mu\text{m}$	25
3.5	A wire scan measurement of the beam halo. The black data points show the halo counter rates as a function of motor position. The two peaks correspond to the horizontal and vertical wires (each wire is iron, 1 mm diameter) crossing the beam. The red curve is a fit to the Gaussian core of the beam, and the blue curve is its convolution with the wire size. The data is consistent with the fits over five orders of magnitude, indicating that the beam halo is of that level.	27

3.6	A rendering of the ECal vacuum chamber. The large rectangular flange closes off the pair spectrometer vacuum chamber. The slot has a complex shape to pass different types of particles radiating from the target. From left to right: the round tube passes bremsstrahlung photons, the oval tube passes the beam (which is enlarged by small-angle scattering) and Møller-scattered electrons, and the wide slot passes bremsstrahlung electrons. The walls of the slot are quite thin and a strut of aluminum honeycomb (not visible) supports it against air pressure.	29
3.7	Beam’s-eye view of the HPS target and the front of the SVT after installation. The vertical rod from the target linear shift (foreground left), the ceramic support rod, and the target frame (center, offset to the right from the end of the support rod) are visible.	30
3.8	Schematic of the SVT and its support systems.	31
3.9	Rendering of the SVT as built, showing cooling lines and motion levers.	31
3.10	Schematic of the APV25. The APSP (analog pulse shape processor) is disabled in the “multi-peak” readout mode used by HPS.	34
3.11	One half-module for L4–6. The two hybrids (without readout chips, which would be mounted on the gold pads) are at the left and right ends. One sensor is in place, on the left. The carbon fiber support and Kapton passivation layer are visible on the right.	35
3.12	The L4–6 pairing fixture, with one half-module in place.	36
3.13	Renderings of the L1–3 (left) and L4–6 (right) module designs, with cutaways to show the spring pivots that hold the silicon under constant tension.	37
3.14	One of two U-channels for L1–3, fully assembled. The beam direction is left to right; the scan wires and motion lever are visible on the left.	38
3.15	One of two U-channels for L4–6, fully assembled. The beam direction is left to right.	38

3.16	Inside the SVT box, looking downstream. The L4–6 U-channels are installed. The rails for the L1–3 U-channels can be seen at the top and bottom. When the U-channels are installed, bearings on the U-channels roll along the horizontal slots until they drop down into the vertical slots, guiding the U-channel onto its kinematic mounts. The cooling lines (in braided jackets) and cables (ends wrapped in foil) for the L4–6 U-channels are visible to the left.	39
3.17	The SVT FEBs (frontend boards), mounted on their cooling plate. The red-and-white cables distribute low-voltage power to the FEBs; the red-and-black cables distribute high-voltage bias to the FEBs. The connectors for hybrid power, bias and data are covered with yellow Kapton tape. The mini-SAS connectors for high-speed data cables are at the bottom right of each FEB.	41
3.18	One HPS signal flange board. The left side of the board operates in vacuum, and has three mini-SAS electrical connectors for high-speed data cables, which connect to the FEBs. The right side of the board operates in air, and has two MPO multi-fiber connectors for data and control signals.	43
3.19	Test of the flange board potting process.	43
3.20	The RCE system. The COB (Cluster On Board) blade on the left hosts four RCE (Reconfigurable Cluster Element) daughterboards, which perform the data processing. The RTM (Rear Transition Module) on the right interfaces with fibers carrying data from the flange boards. The COB and RTM would normally be housed in an ATCA crate. . .	45
3.21	Beam’s-eye-view of the ECal. 18 crystals (the “electron gap”) in the innermost rows are missing, to make space for the oval bulge in the ECal vacuum chamber (Figure 3.6). The thermal enclosure that surrounds the ECal is not shown.	47
3.22	One ECal crystal with its readout electronics. The optical fiber was used in the CLAS IC for calibration and monitoring, but was removed for HPS. HPS uses LEDs mounted directly in front of the crystals. . .	48

4.1	Top: track slope and y at closest approach, for elastically scattered electrons in the top half of the detector. The stripes are the result of single-strip hits in the first layer of the tracker. Bottom: each point is the center of a Gaussian fit to a vertical slice of the top histogram. The linear fit to the points produces estimates of $y_{beamspot}$ and z_{target} .	64
4.2	Angles of Møller electrons (extrapolated to the target, and relative to the nominal beam direction) in the top half of the detector, in a narrow energy band centered at 400 MeV. Equation 4.8 predicts that Møller electrons at this energy should be emitted at 39.8 milliradians from the beam axis.	66
4.3	Cumulative effect of the different pair selection cuts on pair-1 events passing the top-bottom requirement and radiative cuts. The ratio of out-of-time to in-time events (outer peaks to central peak) decreases as cuts are applied (going from the black to the magenta distributions).	69
4.4	Momentum sum $p(e^+e^-)$ for data (black), Monte Carlo samples of tridents (red) and wide-angle bremsstrahlung conversions (blue), and the sum of the red and blue histograms (magenta). All histograms are normalized to integrated luminosity such that the sum of all bins for a histogram equals the total cross-section in nanobarns. If Monte Carlo completely describes data, black and magenta histograms should match.	73
4.5	Same histograms as Figure 4.4, except that data is multiplied by $1/0.65$ to account for the estimated level of detector inefficiency at high momentum. If Monte Carlo completely describes data, black and magenta histograms should match. Note the agreement at high momentum sum (not remarkable, since it is by construction).	74
4.6	Same histograms as Figure 4.5, but the left plot requires that the positron track have a layer 1 hit, which rejects most WAB conversions, and the right plot requires that the positron track not have a layer 1 hit, which rejects most tridents. Note the continued agreement between data and Monte Carlo at high momentum sum.	75

4.7	Distance of closest approach in the X-Z plane for the positron track, only plotting pairs with momentum sum above 0.8 times the beam energy (radiative cut). This is centered at 0 (where the beam spot is) if the positron comes from the target, but is usually positive for WAB positrons created downstream of the target.	76
4.8	$p(e^+e^-)_y \text{sign}(p(e^+)_y)$, only plotting pairs with momentum sum above 0.8 times the beam energy (radiative cut). This is positive if the pair momentum points in the same direction as the positron.	77
5.1	Dataset for the vertexing analysis. The core of the vertex distribution is the horizontal band at $z = -5$ mm; a heavy photon signal would appear as a vertical band extending upwards from the core.	79
5.2	Vertex z (units of mm) for a resolution-limited mass slice (in data) centered at $m_{A'} = 35.5$ MeV. A heavy photon signal would appear as a low-level tail extending to higher z (longer than the exponential tail of the blue curve, which is the prompt vertex distribution described in Section 5.3.3). For this mass slice, $z_{cut} = 26.6$ mm; there are four events with $z > z_{cut}$	80
5.3	The value of the cut $z > z_{cut}$, as a function of the center of the mass slice. The number of events per mass slice decreases and the z resolution improves as the mass increases; both effects tend to reduce z_{cut}	81
5.4	Cumulative effect of the different vertex cuts on events in the mass slice centered at 30 MeV, passing base event selection cuts and meeting the layer 1 hit requirement. The tails at high z are greatly reduced as cuts are applied (going from the black to the magenta distributions). . . .	84
5.5	Efficiency curves for $m_{A'} = 40$ MeV. Left: requiring layer 1 hits for both tracks. Right: full HPS acceptance.	86
5.6	Effect of the track quality cut on events in the mass slice centered at 30 MeV, and passing all other cuts.	87

5.7	Effect of the beamspot constraint cuts on events in the mass slice centered at 30 MeV, and passing all other cuts.	88
5.8	Illustration of a track with a mishit. This view shows the Z (beam direction, beam travels left to right) and Y (magnetic field direction) axes, and three SVT layers with hits (circles). The true particle trajectory (black) starts at the target, is scattered at layer 2, and continues on to layer 3. Because the true trajectory is not straight, a better track fit can be made using a hit (red circle) from a different particle. The resulting track fit (red line) misses the target and, if used in a vertex fit, will lead to a best-fit vertex downstream of the target.	89
5.9	Effect of the isolation cut on events in the mass slice centered at 30 MeV, and passing all other cuts. By design the isolation cut has little impact on vertices with $z \approx z_{target}$, since the tracks in those pairs pass close to the beamspot.	90
5.10	Effect of the WAB rejection cuts on events in the mass slice centered at 30 MeV, and passing all other cuts. The z distributions of events rejected by these cuts are wider than the events that pass, showing that while wide-angle bremsstrahlung conversions do not have systematic shifts in vertex z , they contribute to the tails of the z distribution. . .	91
5.11	Mass resolution vs. z for 40 MeV heavy photons in Monte Carlo. Top: mass residual (difference between the reconstructed mass and the true mass) vs. z for 40 MeV heavy photons. Bottom: the widths of Gaussian fits to vertical slices of the top distribution (blue points), and a linear fit to the points (red line). The slope of the linear fit is consistent with 0, showing that the mass resolution is constant with respect to z	95
5.12	Mass resolution vs. $m_{A'}$. The blue marker is the Møller mass resolution in data.	96
5.13	Distribution of reconstructed invariant mass of Møller pairs, with a Crystal Ball fit showing the mass resolution.	97

5.14	Left: rates of processes producing e^+e^- pairs, with the sum in black. Right: the radiative fraction, which is calculated by dividing the black histogram by the red histogram.	99
5.15	Top: distribution of decay z for heavy photons ($m_{A'} = 40$ MeV, $c\tau = 1$ mm) carrying at least 80% of the beam momentum. Bottom: the momentum distribution for the heavy photons.	100
5.16	Top: distribution of decay z for reconstructed heavy photons ($m_{A'} = 40$ MeV, $c\tau = 1$ mm) after all cuts. Bottom: efficiency curve $\frac{\epsilon_{reco}(m_{A'}, z)}{\epsilon_{reco}(m_{A'}, z_{target})}$, with fit.	101
5.17	Comparison of vertex Z distributions between data (black) and Monte Carlo (red).	102
5.18	Number of events found with $z > z_{cut}$ in each mass slice, as a function of the center mass of the mass slice. Since mass slices overlap, a single event appears many times in this plot. If the background model fully describes the data, the mean number per mass slice of events with $z > z_{cut}$ should be 0.5.	104
5.19	The set of events that appear with $z > z_{cut}$ in at least one mass slice. This is the same 2-D histogram as Figure 5.1, but only plotting events above $z = z_{cut}$ (the black curve). Some events are close to $z = z_{cut}$, but some events are substantially above the black line ($z \gg z_{cut}$). . .	105
5.20	The background CDF values (quantiles) for events in the mass slice centered at $m_{A'} = 35.6$ MeV that have $z > z_{cut}$ (this is the same mass slice as Figure 5.2). The two events with $z \approx 35$ have quantiles close to 1, suggesting that their values of z are larger than would be expected from the background model.	106
5.21	The background quantiles for events that have $z > z_{cut}$ in each mass slice. This is not a true histogram; each vertical column is a histogram similar to Figure 5.20, but a given event is plotted in every mass slice where it appears. The quantile for a given event changes with the mass slice because z_{cut} depends on the mass slice.	107

5.22	The set of events that appear with $z > z_{cut}$ in at least one mass slice (similar to Figure 5.19), for a trident Monte Carlo sample. One event has $z \gg z_{cut}$	108
5.23	Unbiased fit to estimate the background rate past z_{cut} in the mass slice centered at 35 MeV.	110
5.24	Local p -value for finding the observed number of events at each $m_{A'}$. The minimum value is 0.0661.	111
5.25	Left: distribution of the most significant p -values from 10000 runs of toy Monte Carlo. Right: mapping from local to global p -values. . . .	112
5.26	Comparison of the optimum interval method with cut-and-count using Feldman-Cousins limits. The y-axis is the limit on the total signal rate as a fraction of the background rate. The background (10000 events) has decay length 2, the signal has decay length 20, and the unexpected background (100 events) has decay length 5. Left plot is with only the expected background (no signal); right plot is with the unexpected background added (still no signal). The z_{cut} for 0.5 expected background events is 19.8.	115
5.27	90% CL upper limit on μ/μ_{exp} , the ratio of the true production rate to the expected production rate for a heavy photon. A value of 1 would mean exclusion; the lowest contour on this plot is 200, and the lowest point (at $m_{A'} = 49$ MeV, $\epsilon^2 = 2.1 \times 10^{-9}$) is 115. The vertical ridges in this plot correspond to the locations of events in mass space. . . .	117
5.28	μ_{exp} , the number of detectable heavy photon events expected with $z > z_{cut}$. The highest contour is at 0.03 events, and the highest point is 0.032.	118
5.29	Predicted 90% CL upper limit on μ/μ_{exp} , assuming no events with $z > z_{cut}$. A value of 1 would mean exclusion; the lowest contour on this plot is 90, and the lowest point is 86.2.	119

5.30	Predicted 90% CL upper limit on μ/μ_{exp} , projected for the full 2015 data set, and assuming no events with $z > z_{cut}$. A value of 1 would mean exclusion; the lowest contour on this plot is 20, and the lowest point is 14.7.	120
5.31	Predicted 90% CL upper limit on μ/μ_{exp} , projected for the full 2015 data set, using the full HPS acceptance, and assuming no events with $z > z_{cut}$. A value of 1 would mean exclusion; the lowest contour on this plot is 7, and the lowest point is 6.68.	121
5.32	The predicted values of μ_{exp} , the number of detectable heavy photon events expected with $z > z_{cut}$, that were used for the reach estimate in the 2014 proposal. The $m_{A'}$ and ϵ^2 ranges, and the binning, are different in this plot compared to other plots in this dissertation. The highest contour is at 10 events, and the highest point (at $m_{A'} = 22$ MeV, $\epsilon^2 = 5 \times 10^{-9}$) is 17.3 events.	123
5.33	Predicted μ_{exp} , the number of detectable heavy photon events expected with $z > z_{cut}$, with the assumptions listed in Table 5.3. The highest contour is at 7 events, and the highest point is 7.51 events.	125

Chapter 1

Introduction

Massive U(1) vector bosons, also known as heavy photons, are a natural consequence of many theories of physics beyond the Standard Model, and a basic component of many hidden sector models of dark matter. Such a particle would kinetically mix with the photon, giving it an effective coupling to electric charge much smaller than the photon’s direct coupling. The existence of such a heavy photon is a possible explanation for cosmic ray anomalies and the muon $g - 2$ anomaly, both of which are discussed in Section 2.2. Several current and proposed experiments are dedicated to the search for a heavy photon, and are summarized in Section 2.4.

The Heavy Photon Search experiment (HPS) is designed to produce heavy photons by sending the electron beam of Jefferson Lab’s CEBAF accelerator through a fixed target. The detector is a compact, large-acceptance forward spectrometer comprising a silicon microstrip tracker (silicon vertex tracker, or SVT) for momentum measurement and vertexing and an electromagnetic calorimeter (ECal) for triggering. The experiment has recorded physics data in 2015 and 2016, with future running planned.

A heavy photon is described by its mass $m_{A'}$ and the dimensionless kinetic mixing strength ϵ . Cross sections and decay rates scale with ϵ^2 and therefore it is common to quote results in terms of ϵ^2 . ϵ^2 also equals α'/α , the ratio between the heavy photon coupling to charge and the photon coupling to charge. Figure 1.1 shows current constraints in this parameter space from other experiments, and the projected HPS

reach as presented in the 2014 HPS proposal [1].

If produced at HPS, heavy photons will decay to e^+e^- pairs, possibly with some finite decay length; the HPS detector measures the momentum and decay vertex of the pairs. The invariant mass and vertex of the pairs are used in two searches for heavy photons, covering different regions of the parameter space: a “bump-hunt” and a vertexing search. The bump-hunt is a search for a narrow mass resonance above a smooth background, and is sensitive to heavy photons with relatively large couplings (and hence large production). The vertexing search is a search for e^+e^- pairs produced downstream of the target, and is sensitive to heavy photons with relatively small couplings (and hence long decay lengths). This dissertation presents the vertexing search.

The 2015 HPS run was at a beam energy of 1.056 GeV, and nominal current of 50 nA. The beam charge collected during physics data-taking is summarized in Figure 1.2. Because detector commissioning was in progress throughout the run, the SVT was not moved to its nominal position (0.5 mm from the beam) until late in the run. Roughly comparable amounts of data were recorded with the SVT at 1.5 mm and 0.5 mm from the beam.

This dissertation uses only the 2015 data from operation at 0.5 mm. A total of 1166 nb^{-1} of good data was recorded under these conditions, from May 13 through May 18. This value is corrected for trigger deadtime and run quality, as described in Section 4.4; it is equivalent to 1.69 days of ideal running at the nominal beam current. Approximately 90% of the data was blinded so detector performance studies and analysis development could be done on the other 10% without biasing the ultimate result. This dissertation uses only the unblinded fraction, which is a total of 119 nb^{-1} (0.172 days equivalent).

This dissertation presents the HPS experiment as a whole without specific reference to my personal contributions; those are summarized here. I was responsible for the survey of the SVT, and assembled the modules and U-channels (both described in Section 3.2.2). I installed and documented the SVT cooling systems described in Section 3.2.4. I developed the hit time reconstruction algorithm used in the SVT,

described in Section 4.1.1. I developed the detector readout and trigger simulations, which fully simulate the detector pileup, time evolution of pulses and readout pipelines, and trigger algorithms. I used this trigger simulation for the performance studies and tuning of trigger parameters presented in the 2014 HPS proposal; the final trigger parameters described in Section 3.3.2 are based on that work.

Finally, I have had primary responsibility for the displaced-vertex search. The analysis presented in Chapter 5 was outlined in the 2014 HPS proposal [1], but I implemented it. Section 5.5.2 discusses differences between the proposal outline of the analysis, and what I have done.

Figure 1.1: The yellow-green contours outline the expected reach (for exclusion at 90% confidence level) for the HPS experiment with the run plan shown. The reach is split in two regions; the upper region corresponds to the bump-hunt and the lower region corresponds to the vertexing search. Existing limits from other experiments are plotted as shaded regions. The favored region from the muon $g - 2$ anomaly is plotted as a shaded green band (upper left).

Figure 1.2: Rough totals for integrated charge and event count during the 2015 HPS run. The numbers in this plot are not fully corrected for run and event quality.

Chapter 2

Motivation

There is current theoretical and experimental interest in a massive $U(1)$ vector boson that does not couple directly to particles in the Standard Model, but gains a weak effective coupling to charged particles through kinetic mixing. This particle is commonly called a heavy photon, dark photon, or A' , and is characterized by its mass $m_{A'}$ and a dimensionless coupling constant ϵ .

The heavy photon is primarily motivated as part of a larger “dark sector” of particles that are not charged directly under the Standard Model forces. Some sort of “portal” is necessary to create an interaction between the dark sector and the Standard Model. The possible portals are restricted by the symmetries of the Standard Model, and the dominant candidates are commonly referred to as the vector, Higgs, neutrino, and axion portals; the heavy photon is the mediator for the vector portal [2].

Dark sector particles are natural candidates for dark matter, and a heavy photon can be part of a mechanism that produces the observed dark matter abundance. Dark sector models for dark matter are motivated by astrophysical observations supporting the existence of non-Standard Model interactions between dark matter particles, or between dark matter and the Standard Model. Conversely, the detection of a heavy photon would strongly support the existence of a dark sector. Knowledge of the mass and couplings of the heavy photon would inform models of dark matter interactions.

2.1 Theory Summary

The basic assumption behind the heavy photon is that there exists a second (broken) $U(1)$ symmetry, and that it interacts with Standard Model hypercharge via kinetic mixing [3]. At low energies, this leads to the following gauge field Lagrangian, where $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the electromagnetic field strength, $F'_{\mu\nu} = \partial_\mu A'_\nu - \partial_\nu A'_\mu$ is the heavy photon field strength, and ϵ is the dimensionless coupling constant:

$$\mathcal{L}_{\text{gauge}} = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} - \frac{1}{4}F'^{\mu\nu}F'_{\mu\nu} + \frac{1}{2}\epsilon F^{\mu\nu}F'_{\mu\nu} \quad (2.1)$$

The kinetic mixing means that the fields are not orthogonal. Orthogonality can be restored by redefining the electromagnetic field according to $A^\mu \rightarrow A^\mu + \epsilon A'^\mu$. This modifies the interaction Lagrangian as follows:

$$A^\mu J_\mu^{EM} \rightarrow A^\mu J_\mu^{EM} + \epsilon A'^\mu J_\mu^{EM} \quad (2.2)$$

This implies that particles with electric charge acquire a coupling, proportional to ϵ , to the heavy photon.

If particles exist that are charged under both fields, kinetic mixing may arise from a one-loop diagram similar to Figure 2.1, with a natural scale of $\epsilon \sim 10^{-2} - 10^{-4}$; on the other hand, GUT models require that one-loop contributions to ϵ vanish, and instead motivate two-loop contributions at $\epsilon \sim 10^{-3} - 10^{-6}$ [4]. Because kinetic mixing is a renormalizable interaction, ϵ is independent of the masses of the particles that give rise to it. String theory models can motivate much smaller ϵ , as low as 10^{-12} [5, 6].

There is a wide range of reasonable values for the mass $m_{A'}$. Models where supersymmetry breaking is communicated by the kinetic mixing lead to natural mass scales of MeV-GeV [7, 8, 9]. String theory models typically tie the mass scale to ϵ , and can motivate masses down to the meV scale [5, 6].

Figure 2.1: One-loop diagram leading to kinetic mixing. Φ and Φ' are particles that couple to both the photon and the heavy photon. They could be massive (far above the SUSY-breaking scale), but the resulting kinetic mixing will persist down to low mass scales.

2.2 Observations Motivating a Heavy Photon

Two broad classes of observations motivate searches for the heavy photon. First, some astrophysical observations suggest that dark matter is part of a dark sector, and interacts through a heavy photon. Second, precision tests of the Standard Model find anomalies that could be explained by a heavy photon.

2.2.1 Dark Matter Annihilation

Dark matter annihilation mediated by a heavy photon has been proposed as an explanation for several anomalies that have been reported in cosmic ray and X-ray observations.

The PAMELA satellite first observed that the positron fraction $\phi(e^+)/(\phi(e^+) + \phi(e^-))$ increases above 10 GeV. This is inconsistent with the assumption that high-energy positrons originate in secondary production processes, from cosmic-ray nuclei interactions with interstellar gas. Meanwhile, the antiproton fraction is consistent with secondary production. Measurements by the Fermi Large Area Telescope and the Alpha Magnetic Spectrometer (AMS) confirmed the observation and extended it

to higher energies [10, 11].

The annihilation cross-section implied by the positron flux is large compared to the cross-section expected for a thermal relic, and also much larger than any observed annihilation to hadrons [12]. This is consistent with annihilation to heavy photons which then decay to electrons and positrons, as shown in Figure 2.2 [13]. A heavy photon with $m_{A'} < 2m_p$ cannot decay to $p\bar{p}$, and $m_{A'}$ in the MeV–GeV range creates Sommerfeld enhancement that boosts the annihilation cross-section at low velocities. Models with $m_{A'}$ in the 200–900 MeV range were tested in [14]; if annihilation is dominated by local subhalos, lower mass ranges are allowed [15].

Figure 2.2: Dark matter annihilation to heavy photons, which can then decay to Standard Model particles.

The positron fraction has more recently been disfavored as a motivation for the heavy photon. A larger AMS data set shows the full shape of the positron spectrum is softer than would be expected from a heavy photon decaying directly to e^+e^- [16], but consistent with a pulsar origin for cosmic ray positrons [17]. Meanwhile, CMB observations put limits on the DM annihilation rate at recombination; measurements by Planck severely limit the DM annihilation explanation for the positron fraction [18]. The positron spectrum may still be consistent with heavy photons decaying to intermediate states that then decay to e^+e^- [17].

Gamma ray and X-ray excesses also motivate models of dark matter that include heavy photons. Dark matter annihilation has been proposed as a source for an excess

of high-energy gamma rays from the Galactic Center in Fermi data, and hidden sectors can play a role [19]. A 3.5 keV line seen in galaxy cluster X-ray spectra [20, 21] has been suggested to come from collisional excitation and de-excitation of dark matter, in a model named eXciting Dark Matter (XDM) [22].

2.2.2 Halo Structure

Observations of dark matter halos of galaxies and dwarf galaxies have historically disagreed with simulations of halos based on the assumption of collisionless cold dark matter [23]. These discrepancies have motivated models (“self-interacting dark matter,” or SIDM) in which dark matter is self-interacting with a large but velocity-dependent cross-section, consistent with a heavy photon mediator. This allows momentum diffusion in halos without disrupting high-velocity events such as the Bullet Cluster [24, 25].

Rotational velocities of observed Milky Way dwarf satellite galaxies are lower than the rotational velocities that simulations predict for the Milky Way’s largest dark matter subhaloes. This suggests that either the most massive dark matter subhalos fail to form dwarf galaxies, the massive subhalos do not exist, or the rotational velocities for a given subhalo mass are lower than predicted; this is known as the “too big to fail” problem [26]. Self-interacting dark matter reduces the central densities of subhalos and thus the rotational velocities [27]. Similarly, the observed density profiles of galaxies are better fit with a constant density core than the cuspy models from collisionless dark matter: this is the “cusp-core” problem [28]. Both the observational evidence and the simulations continue to develop, and it may be possible to resolve the conflicts without abandoning collisionless dark matter [29]. For example, baryonic outflows in galaxies may transfer energy to the dark matter and create the same “cored” distributions produced by self-interacting dark matter [30].

2.2.3 Muon $g - 2$ Anomaly

The measured value of the muon $g - 2$ (anomalous magnetic moment, also known as $a_\mu = (g - 2)/2$) is more than 3 standard deviations away from the value predicted

by the Standard Model [31, 32]. A heavy photon can contribute to the magnetic moment through the diagram shown in Figure 2.3; there is a specific band in the $\epsilon - m_{A'}$ parameter space (shown in green in Figure 1.1) where the heavy photon correction accounts for the anomaly [33]. The good agreement of the electron $g - 2$ with the Standard Model excludes a different region (drawn in red); similarly, there is an excluded region (drawn in green, above the favored band) where the heavy photon correction to muon $g - 2$ exceeds the observed anomaly [34]. This favored region has been targeted by many experiments, and is now excluded for the case of heavy photon decays to visible particles.

Figure 2.3: Heavy photon correction to the muon magnetic moment.

2.3 Signatures

The heavy photon's only coupling to the Standard Model is through kinetic mixing with the photon. So assuming that heavy photon decays to the dark sector are kinematically forbidden, the production and decay of a heavy photon are closely related to the production and decay of a virtual photon with the same mass.

The branching fractions for decay, as a function of $m_{A'}$, can be derived from the ratios of cross sections for different final states of e^+e^- collisions, as a function of center-of-mass energy: see Figure 2.4. The decay width and proper lifetime are given by the following formulas, where N_{eff} is the effective number of possible available

decay channels: neglecting phase-space corrections, $N_{eff} = 1$ for $m_{A'} < 2m_\mu$, and $N_{eff} = 2 + R(m_{A'})$ for $m_{A'} \geq 2m_\mu$, where $R(Q) = \frac{\sigma(e^+e^- \rightarrow \text{hadrons}, Q)}{\sigma(e^+e^- \rightarrow \mu^+\mu^-, Q)}$.

$$\Gamma = \frac{N_{eff} m_{A'} \alpha \epsilon^2}{3} \quad (2.3)$$

$$\tau = \frac{\hbar}{\Gamma} = \frac{3\hbar}{N_{eff} m_{A'} \alpha \epsilon^2} \quad (2.4)$$

Since ϵ^2 is small, the decay width is very narrow ($\Gamma/m_{A'} \propto \alpha \epsilon^2$) and the heavy photon appears as a sharp resonance.

In general, the rate of production of a heavy photon with mass $m_{A'}$ is proportional to the rate of the corresponding process for virtual photons. The kinematics are the same, so the virtual photon process is an irreducible background except that the heavy photon is on-shell and can travel a measurable distance. A measurement of the virtual photon process can be used to normalize the sensitivity of a search in a data-driven way. The ratio of differential cross-sections is given as follows, where the cross-section of the virtual photon process is integrated over the mass range bounded by $m_{A'} \pm \frac{\delta m}{2}$.

$$\frac{d\sigma(X \rightarrow A'Y \rightarrow ZY)}{d\sigma(X \rightarrow \gamma^*Y \rightarrow ZY)} = \left(\frac{3\pi\epsilon^2}{2N_{eff}\alpha} \right) \left(\frac{m_{A'}}{\delta m} \right) \quad (2.5)$$

2.4 Overview of Searches

Recent and planned searches for heavy photons can be categorized by the production mechanism, the signature for detection, and the type of facility.

Detection signatures include a mass resonance in the decay products (e^+e^- , $\mu^+\mu^-$, or $\pi^+\pi^-$ pairs), missing mass (which allows a search for heavy photons decaying to the dark sector), or displaced vertices.

The common production mechanisms are bremsstrahlung ($e^-Z \rightarrow e^-ZA'$), Drell-Yan ($q\bar{q} \rightarrow \gamma A'$), e^+e^- annihilation ($e^+e^- \rightarrow \gamma A'$), and meson decay ($\pi^0 \rightarrow \gamma A'$, $\eta \rightarrow \gamma A'$, $\phi \rightarrow \eta A'$, and others). Some experiments look for mass resonances by making all possible pairs (for example, all $\mu^+\mu^-$ pairs) without identifying the production

process; these are called “inclusive” searches.

Fixed-target experiments and colliders are complementary. The higher luminosities at fixed-target experiments allow sensitivity at smaller ϵ , particularly for displaced-vertex searches. Colliders, with their higher center-of-mass energy, can reach larger $m_{A'}$.

2.4.1 Beam Dumps

An electron or proton beam dump will generate heavy photons through bremsstrahlung, Drell-Yan, or meson decays. If the heavy photon decay length is comparable to the dump length, heavy photons can travel through the dump and decay to visible particles after the dump. A detector downstream of the dump detects the decay products.

Beam dump experiments have very high luminosities, due to the high currents and thick targets. Several electron beam dump experiments were run as searches for axion-like or Higgs-like particles, but were reinterpreted to give limits on a heavy photon: E137 [36] and E141 [37] at SLAC, E774 [38] at Fermilab, and an experiment at Orsay [39]. Similarly, a proton beam dump experiment at the U70 accelerator was reinterpreted to give limits for heavy photons, both from π^0 decay and from proton bremsstrahlung [40].

Planned proton beam dump experiments include SHiP at CERN and SeaQuest E1067 at Fermilab. SHiP is a general-purpose search for dark sector particles [41]; E1067 is an upgrade to an existing dimuon spectrometer to allow parasitic searches for heavy photons [42]. NA64 at CERN is an electron beam experiment that uses an electromagnetic calorimeter as an active dump [43].

2.4.2 Thin Fixed Targets

An electron beam incident on a thin target will generate heavy photons through bremsstrahlung, $e^-Z \rightarrow e^-ZA'$. The signature is a mass resonance in the decay products (typically e^+e^- ; $\mu^+\mu^-$ and $\pi^+\pi^-$ are accessible with beam energies above a couple of GeV). In contrast to beam dump experiments, these experiments are sensitive to heavy photons that decay promptly, without a resolvable decay length.

The mass resolution can be improved by detecting the recoiling beam electron and/or nucleus.

HPS is the first experiment of this type to attempt a displaced-vertex search.

Two experiments have been run with electron beams and existing fixed-target spectrometers: one experiment using the A1 spectrometer at the Mainz Microtron [44], and the APEX experiment using the HRS spectrometer at JLab Hall A (a test run was completed, a full run is scheduled) [45, 46]. Two experiments are planned: DarkLight at the JLab Low-Energy Recirculator Facility [47], and MAGIX at the Mainz Energy-Recovering Superconducting Accelerator [48].

A positron beam incident on a thin target could generate heavy photons through annihilation. Experiments of this type can detect monophoton final states ($e^+e^- \rightarrow \gamma A' \rightarrow \gamma, \text{dark}$) and identify the heavy photon in the missing mass spectrum. Experiments are proposed at INFN Frascati (PADME [49]) and VEPP-3 at the Budker Institute at Novosibirsk [50].

Proton or ion beams can also be used: the HADES spectrometer was used to measure inclusive e^+e^- pairs from p -on- p , p -on-Nb, and Ar-on-KCl interactions. The assumed signal was from meson decays [51].

2.4.3 Colliders

Most collider searches for heavy photons have come from flavor factories, where meson decays are a natural source of heavy photon production. e^+e^- annihilation or Drell-Yan can also be a significant source of heavy photons. Displaced-vertex searches are often possible.

Among e^+e^- colliders, BaBar has set limits using $\Upsilon \rightarrow \gamma A'$, $A' \rightarrow \mu^+\mu^-$ [52, 53]; KLOE has set limits using $\phi \rightarrow \eta A'$ and $e^+e^- \rightarrow \gamma A'$ production modes and dielectron, dimuon, and dipion decays [54, 55, 56, 57, 58]. Belle-II is expected to improve on the BaBar result [59].

$p-p$ colliders have also been used. WASA at COSY looked for $\pi^0 \rightarrow \gamma A'$, $A' \rightarrow e^+e^-$ [60]. LHCb plans to make a significant impact with a search in $m_{A'} < 2m_\mu$ using $D^* \rightarrow D^0 A'$, $A' \rightarrow e^+e^-$, a search in $m_{A'} > 2m_\mu$ using inclusive production and

$A' \rightarrow \mu^+\mu^-$, and a displaced-vertex search using $\eta \rightarrow \gamma A'$, $A' \rightarrow \mu^+\mu^-$ [61, 62].

2.5 Overview of HPS

In HPS, the heavy photon is generated through electron bremsstrahlung on a tungsten target, and detected in its decay to e^+e^- . The process is shown in Figure 2.5. This process is the analogue of the radiative trident process shown in Figure 2.6, with a heavy photon substituted for the radiated photon. Only the e^+e^- pair is triggered on and reconstructed; the recoil electron is typically outside the detector acceptance. Both the pair invariant mass and vertex position are reconstructed.

Two analyses are performed on the same data set. A pure bump-hunt search at large ϵ assumes the heavy photon decays at the target, which improves the mass resolution. This is a search for a mass resonance, using only the invariant mass of the pair; it is described fully in [63]. A displaced-vertex search at small ϵ rejects pairs originating from the target. Because the heavy photon must travel a measurable distance, this search is not sensitive to large ϵ .

This dissertation presents the displaced-vertex search.

2.5.1 Signal Kinematics

The kinematics of heavy photon production can be calculated using the Weizsäcker-Williams approximation (where the nucleus is replaced by an effective photon flux); this is done in [64], based on work in [65, 66]. An exact (to leading order, with radiative corrections) calculation shows similar results [67, 68]. The differential cross-section for an electron with energy E_0 to produce a heavy photon with energy $E_{A'} = xE_0$, at angle $\theta_{A'}$ from the beam (in the lab frame), is

$$\frac{d\sigma}{dx d\cos\theta_{A'}} \approx \frac{8Z^2\alpha^3\epsilon^2 E_0^2}{U^2} \frac{\chi}{Z^2} \left(1 - x + \frac{x^2}{2} - \frac{x(1-x)m_{A'}^2(E_0^2 x \theta_{A'}^2)}{U^2} \right) \quad (2.6)$$

where χ/Z^2 ($\sim 5 - 10$ for values of $m_{A'} < 100$ MeV, where HPS is sensitive) depends on atomic form factors and kinematics, and $U(x, \theta_{A'}) = E_0^2 x \theta_{A'}^2 + m_{A'}^2 \frac{1-x}{x} + m_e^2 x$ is

the virtuality of the electron in the intermediate state. The approximation assumes $m_e \ll m_{A'} \ll E_0$ and $x\theta_{A'}^2 \ll 1$. The characteristic angle of emission is set by $U(x, \theta_{A'}) - U(x, 0) \sim U(x, 0)$, which occurs at $\theta_{A'} \sim \frac{m_{A'}\sqrt{1-x}}{xE_0}$.

Integrating out $\theta_{A'}$ and ignoring m_e , the differential cross-section is

$$\frac{d\sigma}{dx} \approx \frac{8Z^2\alpha^3\epsilon^2x}{m_{A'}^2} \frac{\chi}{Z^2} \left(1 + \frac{x^2}{3(1-x)}\right) \quad (2.7)$$

There is a cutoff at $1 - x = \max\left(\frac{m_e^2}{m_{A'}^2}, \frac{m_{A'}^2}{E_0^2}\right)$, and the median value of $1 - x$ is $\max\left(\frac{m_e}{m_{A'}}, \frac{m_{A'}}{E_0}\right)$.

The overall rate is controlled by $\frac{\alpha^3\epsilon^2}{m_{A'}^2}$, so heavy photon production is increasingly less likely at larger $m_{A'}$. From Equation 2.7, the heavy photon typically takes most of the incident electron's momentum: $E_{A'} \approx E_0$. This means that the typical opening angle of the decay products is $2m_{A'}/E_0$, and the typical vertex displacement is

$$\gamma c\tau \approx \frac{3\hbar E_0}{N_{eff}m_{A'}^2\alpha\epsilon^2c} \quad (2.8)$$

From Equation 2.6, $\theta_{A'}$ is dominated by small values: the typical value of $\theta_{A'}$ is $\left(\frac{m_{A'}}{E_0}\right)^{3/2}$, so $\theta_{A'}$ is smaller than the opening angle of the decay products and it is a good approximation to assume the heavy photon is emitted in the beam direction.

The recoiling incident electron is soft and also largely balances the $p_\perp$ of the heavy photon, and therefore is at relatively wide angle $\theta_R \sim \sqrt{\frac{m_{A'}}{E_0}}$. It rarely makes it through the full length of the detector, and is not easily tracked.

2.5.2 Physics Backgrounds

The major sources of e^+e^- backgrounds are radiative tridents, Bethe-Heitler tridents, and pair conversions of wide-angle bremsstrahlung photons. The first two of these are $e^-Z \rightarrow e^+e^-e^-Z$ processes, and are called “tridents” because of the three-lepton final state. The third is $e^-Z \rightarrow e^-\gamma Z$ (bremsstrahlung) followed by $\gamma Z \rightarrow e^+e^-Z$ (pair conversion), which can mimic a trident event.

As discussed in Section 2.3, the radiative trident and heavy photon processes

have the same kinematics. The radiative trident process is therefore an irreducible background, except that heavy photons can be distinguished by the mass resonance and measurable decay length.

The Bethe-Heitler trident process (Figure 2.7) is another background. The interference between the radiative and Bethe-Heitler diagrams is significant; since radiative tridents are of special interest to HPS, the interference terms and the Bethe-Heitler diagrams are lumped together as “non-radiative” tridents. In contrast to heavy photons and radiative tridents, pairs from non-radiative tridents are not peaked at high x . In general, the two electrons are not distinguishable and either e^+e^- pair is a potential background.

Wide-angle bremsstrahlung (where the photon and electron are both at large angles from the beam axis), followed by pair conversion of the photon (either in the target or in the first parts of the detector), is another significant background. Since the photon is on-shell, the e^+e^- pair from the pair conversion is generally not in the HPS acceptance. However, the scattered incident electron and the positron from the pair conversion are often in the acceptance.

Figure 2.4: Branching ratios for the heavy photon decay to Standard Model particles [35]. The bottom plot shows the range of $m_{A'}$ to which HPS is sensitive.

Figure 2.5: Heavy photon production and decay. The heavy photon (A') can be produced as final state radiation (right) or initial state radiation (left). In both cases, the heavy photon is on-shell and can travel some distance before decaying in free space.

Figure 2.6: Radiative trident diagrams. These diagrams are also known as timelike tridents, because the photon radiated by the incident electron (the upper photon in each diagram) is timelike.

Figure 2.7: Bethe-Heitler trident diagrams. These diagrams are also known as space-like tridents, because the photon radiated by the incident electron (the upper photon in each diagram) is spacelike.

Chapter 3

The HPS Detector

The HPS detector is a spectrometer with a silicon vertex tracker (SVT) for momentum measurement and an electromagnetic calorimeter (ECal) for energy measurement and trigger. The detector is installed in the middle dipole of a three-magnet chicane (see Figure 3.1), with the field extending from the target to the end of the tracker.

The detector design is determined by physics needs. To capture low-mass A' , the detector must have acceptance at small angles from the beam. To get the best possible vertex resolution, the detector must operate as close to the target as possible. Because multiple scattering dominates tracking resolution at HPS energies, the material in the tracking volume must be kept as low as possible.

Elastic scatters in the target send large numbers of electrons into the detector acceptance, so it needs to tolerate high rates and have a selective trigger. Beam-gas interactions would create large detector backgrounds and fake A' decays downstream of the target, so the beam must travel in vacuum all the way through HPS. Bremsstrahlung energy losses in the target cause beam electrons to bend in the dipole field, forming a “sheet of flame.” To avoid this, no detector material is placed in the beam plane.

All parts of the HPS detector have the same minimum vertical angle, or “dead zone,” at 15 milliradians. This is set by the maximum rate tolerable in the first layer of the SVT: the occupancy within a resolution-limited hit time window must be less than 1% for clean track reconstruction, and radiation damage must be kept low

enough that the SVT remains fully functional after six months of operation.

Figure 3.1: View of the HPS setup. The beam direction is left to right.

3.1 Beamline

HPS uses the CEBAF (Continuous Electron Beam Accelerator Facility) accelerator at Jefferson Lab (Figure 3.2). CEBAF is a recirculating linac, where the electron beam can take multiple passes through the same set of accelerating cavities. The superconducting RF cavities used at CEBAF allow a continuous duty cycle, where beam bunches pass through the accelerator at 1500 MHz without interruption. A system of RF separators delivers beam to each of four experimental halls at 250 or 500 MHz, and allows each hall to select its own beam energy. HPS is installed in Hall B, the hall that typically operates at the lowest beam current.

The main detector in Hall B is CLAS (CEBAF Large-Acceptance Spectrometer), which is used for precision nuclear physics experiments with low-current (sub- μA) beams of electrons or photons. CEBAF recently underwent a major upgrade to increase the maximum beam energy from 6 GeV to 12 GeV, and add a new experimental hall (Hall D). CLAS is undergoing a major rebuild (to become the CLAS12 detector) for 12 GeV operation, and this work is still in progress. The 2015 and 2016 HPS runs were conducted after CEBAF began 12 GeV operations but before much of CLAS12 was complete; HPS is the first Hall B experiment of the 12 GeV era, albeit not operating at 12 GeV (the maximum field of the HPS magnets limits HPS to 6.6 GeV

beam).

The injector energy is 100 MeV and one pass through the linacs adds 2.2 GeV to the beam energy, so in normal operation, the available beam energies at Hall B are $100 \text{ MeV} + n * 2.2 \text{ GeV}$ where n is 1 through 5. During the 2015 engineering run, a mechanical problem disabled one of the two CEBAF helium liquifiers. With half the cooling power, the superconducting cavities could only be run at half the nominal gradient. HPS took the opportunity to run at 1.056 GeV, an energy that is not normally available.

HPS relies on the continuous beam structure at CEBAF to reduce pileup. A beam bunch arrives at HPS every 2 ns, which is comparable to the time resolution of the detectors. This means that beam backgrounds are spread in time as uniformly as possible. A larger bunch spacing or lower duty cycle would increase the amount of beam background that overlaps an event of interest.

Figure 3.2: Schematic of the CEBAF accelerator, highlighting components added for the 12 GeV upgrade.

Figure 3.3: Location of the HPS setup, in the downstream alcove of Hall B. The beam direction is left to right.

3.1.1 Hall B Beamline and Instrumentation

HPS is installed in the downstream alcove of Hall B: see Figure 3.3. Since CLAS12 construction is in progress, a plain beam pipe was installed to pass the beam through CLAS, and HPS operations were limited to nights and weekends to allow CLAS construction to continue.

The Hall B beamline can be configured for delivery of both electron and photon beams. Soon after it enters the hall, the CEBAF electron beam passes through a large “tagger” magnet. If a target is positioned upstream of the tagger magnet, a beam of bremsstrahlung photons can be sent to experiments while the tagger magnet sends the beam to a dump (the “tagger dump”) and the bremsstrahlung electrons to wire chambers (“tagging” the energy of each photon). For HPS operation the beamline is run in the electron configuration (tagger magnet de-energized, electrons to HPS and main dump), but the photon configuration (tagger magnet energized, electrons to tagger dump) is still used for beam setup since it allows for tuning the upstream part of the beamline without impacting the beam at HPS.

The Hall B beam instrumentation consists of beam position monitors (BPMs),

halo monitors, wire scanners, viewers, and a Faraday cup.

BPMs detect the current induced by passing beam bunches in order to measure the current and/or position of the beam. Since this is a noncontact measurement, BPMs operate continuously during operations and provide a record of beam changes during data taking.

Halo monitors are small particle detectors (mostly scintillation counters) positioned around the beam pipe. If the beam is obstructed, defocused or missteered, it will scatter into the halo monitors.

Wire scanners consist of a set of wires mounted on a motorized frame. The beam position and size is measured by scanning the wires across the beam; the halo counter rates increase in proportion to the amount of beam subtended by a wire. An example wire scan is shown in Figure 3.4. All harps have at least one vertical and one horizontal wire; some have a diagonal (45°) wire for measuring the beam tilt and/or thicker wires for measuring the tails of the beam.

Figure 3.4: A representative wire scan measurement of the beam size. This scan was taken using the 2H02 wire scanner just upstream of HPS, and reflects the beam size seen at the HPS target. This wire scan was taken early in the 2015 run; the beam size in X was reduced later in the run, and all physics data was recorded with a beam size of roughly $150 \times 50 \mu\text{m}$.

Beam viewers consist of a thin screen that emits light when the beam passes

through it, and a video camera for viewing the shape and position of the beam. Fluorescent screens are sensitive but are susceptible to saturation and blooming effects, and are most useful for viewing the beam position. A screen using optical transition radiation (OTR) is dimmer, but does not saturate and accurately shows the beam shape. One fluorescent screen is mounted in front of the Hall B tagger dump. Two fluorescent screens of different types, and one OTR screen, are mounted on a motorized carousel in front of the Hall B main dump.

The Hall B beamline terminates in a Faraday cup, which directly measures the beam charge. This is the most accurate measurement of the beam current.

3.1.2 Beam Quality and Protection

HPS has two primary concerns regarding the beam. First, since the SVT is so close to the beam, the SVT must be protected both from beam being directed into the silicon and causing damage, and from stray beam electrons hitting the inner regions of the active silicon and adding to the detector pileup. This implies strong beam protection controls and low beam halo. Second, a small beam spot is important for event quality cuts; if the beam spot is small (ideally, smaller than detector resolution), poorly reconstructed tracks or vertices that are not consistent with an origin at the beam spot can be rejected.

Active beam protection is provided by the CEBAF fast shutdown (FSD) system. The FSD system is an interlock which, when triggered, shuts off the electron gun in the injector. In addition to the large amount of beam instrumentation that is normally connected to the FSD, the set of halo counters closest to HPS were used to provide an FSD signal. This halo counter FSD was configured so that an abnormally high rate on the halo counters (the threshold being set as low as possible without causing spurious beam trips) would trip the beam in 1 ms. If the beam were to move into the SVT sensors, it would hit the inactive edge of the silicon (0.5 mm from the nominal beam position) first. The beam scattering from the silicon would trip the halo counter FSD, hopefully before the beam could reach the active silicon (1.5 mm from the nominal beam position).

A collimator provides passive beam protection. The collimator is a tungsten plate 1 cm thick with machined slots of different widths, mounted on a linear shift so the appropriate slot can be selected and positioned precisely. For the 2015 run, the 4 mm slot was used. If the beam were to move into active silicon, the collimator would absorb and spread the beam enough that the silicon would not be damaged in the time it would take for the FSD to stop the beam.

Measurements of the beam halo, such as shown in Figure 3.5, show that the beam profile does not deviate from a Gaussian for five orders of magnitude; the rate of beam electrons outside of ± 0.5 mm is 10^{-5} of the beam current. At this level, the rate of beam halo electrons hitting the innermost strips of the SVT is less than the rate of scattered electrons from the target.

Figure 3.5: A wire scan measurement of the beam halo. The black data points show the halo counter rates as a function of motor position. The two peaks correspond to the horizontal and vertical wires (each wire is iron, 1 mm diameter) crossing the beam. The red curve is a fit to the Gaussian core of the beam, and the blue curve is its convolution with the wire size. The data is consistent with the fits over five orders of magnitude, indicating that the beam halo is of that level.

HPS originally requested a beam size that was smaller in Y than in X. Since the

vertex resolution is better in Y ($\sim 150 \mu\text{m}$) than in X ($\sim 300 \mu\text{m}$), and target heating was incorrectly believed to limit the minimum safe beam size, the desired beam shape was narrow ($\sim 50 \mu\text{m}$) in Y and wider ($\sim 300 \mu\text{m}$) in X. One challenge in making such an asymmetric beamspot is that the beam tilt must be precisely controlled. As shown in Figure 3.4, this was achieved. It was later realized that target heating was not a serious concern, and all 2015 physics data was recorded with a beam size of $\sim 50 \mu\text{m}$ in Y and $\sim 150 \mu\text{m}$ in X.

3.1.3 HPS Beamline Elements

The HPS chicane, shown in Figure 3.1, consists of three dipoles with fields in the vertical direction. The three magnets of the chicane are repurposed from a past Hall B experiment. The detector is installed in the central magnet (the “pair spectrometer” magnet), which is a 18D36 magnet (pole length 91.44 cm, width 45.72 cm), and was operated at a field strength of 0.24 T for the 2015 run. The outer magnets are identical “Frascati”-type magnets, equidistant from the analyzing magnet and operated at equal field strengths, such that the $\int Bdl$ of each Frascati magnet is half that of the analyzing magnet (and opposite in sign). This ensures that the beam trajectory downstream of the chicane is the same whether or not the chicane is energized.

A series of connected vacuum chambers constitute the HPS beam path. A short rectangular chamber upstream of the pair spectrometer magnet provides feedthroughs for SVT services (motion, cooling, power, data), further described in Section 3.2.4. A long rectangular chamber (inherited with the pair spectrometer magnet) fills the magnet bore and extends downstream; the downstream segment flares outward to allow charged particles to bend. The next chamber (shown in Figure 3.6) closes off the vacuum chamber, except for a slot-shaped channel that allows the beam and the sheet of flame to pass through to the dump; the ECal (which does not operate in vacuum) closely surrounds the slot on top and bottom.

The HPS target is a set of foils mounted on a common frame [69]. There are two tungsten foils, one graphite foil, and one polyethylene foil. The tungsten foils have

Figure 3.6: A rendering of the ECal vacuum chamber. The large rectangular flange closes off the pair spectrometer vacuum chamber. The slot has a complex shape to pass different types of particles radiating from the target. From left to right: the round tube passes bremsstrahlung photons, the oval tube passes the beam (which is enlarged by small-angle scattering) and Møller-scattered electrons, and the wide slot passes bremsstrahlung electrons. The walls of the slot are quite thin and a strut of aluminum honeycomb (not visible) supports it against air pressure.

design thicknesses of 0.125% and 0.25% radiation lengths; from measurements made during target assembly, the true thicknesses are 0.116% and 0.223%. The 0.125% X_0 foil is for runs at 1.1 and 2.2 GeV, and the 0.25% X_0 foil is for runs at 4.4 and 6.6 GeV. A linear shift moves the frame up and down so the foils can be moved in and out of the beam; since the target needs to be at the face of the magnet, and the linear shift is on one of the flanges of the upstream vacuum chamber, the target frame is cantilevered on a ceramic support rod.

3.2 Silicon Vertex Tracker

The silicon vertex tracker (SVT) measures track momentum and vertex position.

HPS uses silicon microstrips because they provide good time resolution, low mass, and high rate. Because a microstrip sensor can only make a 1-D measurement (it identifies the strip that was hit, but not the hit position along the strip), each measurement station (called a “layer”) uses two sensors at an angle relative to each other. For HPS

Figure 3.7: Beam's-eye view of the HPS target and the front of the SVT after installation. The vertical rod from the target linear shift (foreground left), the ceramic support rod, and the target frame (center, offset to the right from the end of the support rod) are visible.

Figure 3.8: Schematic of the SVT and its support systems.

Figure 3.9: Rendering of the SVT as built, showing cooling lines and motion levers.

this stereo angle is kept small, so all sensors have their strips pointing roughly in the bend direction. A small stereo angle is a compromise: it sacrifices measurement resolution in the strip direction, but it is less likely that two particles will create “ghost hits” when the strips they hit intersect each other. The stereo angle is not uniform throughout the SVT (it is 100 milliradians in the upstream half, 50 milliradians in the downstream half), which prevents ghost hits from creating ghost tracks.

The SVT is made of six layers at different distances from the target: this number allows for 5-hit tracking even if a particle misses one layer. Layer 1 is 10 cm from the target; this is the closest we can safely operate while allowing 500 μm distance from the beam to the edge of the sensors, which have a 1 mm border of inactive silicon. Layer 6 is 90 cm from the target, just at the end of the uniform field region of the analyzing magnet; maximizing the length of the tracker maximizes the momentum resolution. Layers 1–3 use single sensors; layers 4–6, where tracks have bent out to the sides, use two sensors joined end to end. The layout of the six layers is summarized in Table 3.1.

Table 3.1: Layout of the HPS SVT. The angle of stereo sensors is relative to the bend plane.

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

Because the SVT operates in a high magnetic field and in vacuum, materials must be compatible. All materials are checked for vacuum compatibility by pumping test samples to high vacuum in a test chamber. Only nonmagnetic materials are used; inductors used on the frontend boards are air-core (instead of ferrite, which would saturate).

3.2.1 Sensors and Readout

HPS uses silicon microstrip sensors that were originally produced for the run IIb upgrade of the DØ detector at Fermilab. This upgrade would have replaced the entire SMT (silicon microstrip tracker). The full SMT upgrade was cancelled in favor of an insertable Layer 0, but not before sensors had already been procured. The sensors used for HPS are those that would have been used for layers 2–5 of the new DØ tracker.

The HPS sensors are single-sided p+n with AC-coupled readout: the bulk is lightly doped n-type silicon with $\langle 100 \rangle$ crystal orientation, and the strip implants are strongly p-type doped. The strips are biased through polysilicon resistors at the ends of the strips, and capacitively coupled to aluminum readout strips that run on top of the strips. Only every other strip is read out; the “readout” strips capacitively couple to the intermediate “sense” strips, so a hit in a sense strip splits its charge between the neighboring readout strips.

Radiation damage limits the useful lifetime of silicon sensors. Incident particles can displace silicon atoms from their places in the crystal lattice, which effectively converts the n-type bulk of the sensor to p-type (type inversion). This increases the depletion voltage, so the sensor bias must be increased to keep the same charge collection efficiency; the sensor lifetime is therefore limited by the breakdown voltage of the sensor. The defects also increase the leakage current, which leads to increased sensor heating. The HPS sensors are specified to have a breakdown voltage of greater than 350 V, and a further selection was made to only use sensors with a breakdown voltage in excess of 1000 V.

The sensors are read out by the APV25 readout chip [70]. This chip was developed for silicon microstrip readout in the CMS tracker. Because the APV25 can read out multiple consecutive samples of its shaper waveform, it can be used for pileup rejection and high-precision hit time reconstruction. This is an essential feature for HPS and other experiments (notably the Belle II SVD [71]) with CW beam and high pileup.

Each APV25 chip has 128 input channels. One channel consists of a charge-sensitive preamp with an optional inverter, CR-RC shaper, and 192-cell analog pipeline. We run the chip with a 24 ns clock: the design clock period is 25 ns to match the

Table 3.2: Specifications of the SVT sensors. Breakdown voltage specification is value accepted for use in HPS; procurement specification was looser.

Thickness	320 μm
Overall area (L $\times$ W)	100 mm $\times$ 40.34 mm
Active area (L $\times$ W)	98.33 mm $\times$ 38.34 mm
Strip pitch (count)	30 μm (1277)
Readout pitch (count)	60 μm (639)
Depletion voltage	110-130 V (typical)
Breakdown voltage	>1000 V

LHC bunch crossing period, but 24 ns is an even multiple of the standard JLab clock. On each clock, each channel samples its shaper output and stores it in a cell of its pipeline. On a trigger, each channel reads out the appropriate pipeline cell, and the chip multiplexes the 128 signals onto a single differential current output. A configurable “latency” setting determines which pipeline cells are read out; since the cells of interest are those that were written at the time of the triggering event, the latency should approximately equal the trigger delay.

Figure 3.10: Schematic of the APV25. The APSP (analog pulse shape processor) is disabled in the “multi-peak” readout mode used by HPS.

HPS runs the APV25 in its “multi-peak” readout mode, which allows us to read out six consecutive pipeline cells for each trigger. Six samples per trigger, combined with knowledge of the shaper pulse shape (from an offline calibration), allows us to

fit the shaper output to the sum of one or two pulses. The result is a reconstructed hit time and amplitude that is unaffected by a pileup hit that comes before or after the hit of interest.

3.2.2 Mechanical Support

The base unit of the SVT is a “half-module” comprising a low-mass support structure, sensors, and hybrid readout circuit boards. Half-modules are the only unit of the SVT that cannot be disassembled and reworked. Two types of half-modules are used: layers 1–3 use single-ended half-modules with one sensor and hybrid each, and layers 4–6 use double-ended half-modules with two sensors and hybrids each. A single half-module provides a single measurement (axial or stereo) for one half (top or bottom) of a layer.

Figure 3.11: One half-module for L4–6. The two hybrids (without readout chips, which would be mounted on the gold pads) are at the left and right ends. One sensor is in place, on the left. The carbon fiber support and Kapton passivation layer are visible on the right.

The hybrid circuit board carries the APV25 readout chips, and connects the sensor to the rest of the DAQ. The input channels of the APV25 chips are wirebonded directly to the sensor; the APV25 power, output channels, and control lines are wirebonded to the hybrid. The hybrid also carries filter capacitors for the sensor bias, and temperature sensors to monitor the sensor temperature.

The carbon fiber support structure provides structural support for the silicon, and acts as a ground plane for the half-module. A layer of Kapton insulation isolates the carbon fiber from the back surface of the sensor, which is held at high voltage. The carbon fiber and Kapton are thinner than the silicon and contribute negligibly to the

material seen by particles; cutouts further reduce any effect. The sensors and hybrids are glued to the support structure with epoxy.

Two half-modules are paired back-to-back to form a “module.” The axial half-module is oriented with its strips pointing in the bend direction; the stereo half-module is rotated so it dips into the beam plane on the positron side (where beam backgrounds are less intense). The modules are assembled using pairing fixtures, which are machined to set the edges of the sensors at precisely the correct height and angle.

Figure 3.12: The L4-6 pairing fixture, with one half-module in place.

The aluminum module support holds the half-modules at both ends. Heat generated by the hybrids is pulled out through the module support, which is in close thermal contact with the APV25 chips and the sensors through parallel paths so that

the sensors can be kept colder than the APV25 chips. The module support and the half-modules contract at different rates when the SVT is cooled to its operating temperature, so the module support must apply constant tension to keep the half-modules flat. This is done with a spring pivot on one side of the module support.

Figure 3.13: Renderings of the L1–3 (left) and L4–6 (right) module designs, with cutaways to show the spring pivots that hold the silicon under constant tension.

Three modules are mounted on a common aluminum support structure to form a “U-channel.” The sidewalls add to the rigidity of the U-channel and shield the sensors from thermal radiation. The module mounting surfaces are recessed by the correct amounts to put the layers at the correct distances from the beam. The SVT is divided into four U-channels: top and bottom L1–3, top and bottom L4–6.

Each U-channel is supported at three points using kinematic mounts, which guarantee repeatable positioning when the U-channels are installed. The L4–6 U-channels rest on three kinematic mounts. The L1–3 U-channels rest on two kinematic mounts, which serve as a hinge at the downstream end of the U-channels, and are supported on the upstream end by motion levers which tilt the U-channels up and down. In addition to modules, the L1–3 U-channels carry scan wires so that the beam position can be measured relative to the silicon.

Survey

The SVT is surveyed to determine the positions of the sensors in the detector volume. This serves two purposes. First, surveying the SVT checks that it was assembled as designed, and allows adjustable components to be brought to their design positions.

Figure 3.14: One of two U-channels for L1–3, fully assembled. The beam direction is left to right; the scan wires and motion lever are visible on the left.

Figure 3.15: One of two U-channels for L4–6, fully assembled. The beam direction is left to right.

Figure 3.16: Inside the SVT box, looking downstream. The L4–6 U-channels are installed. The rails for the L1–3 U-channels can be seen at the top and bottom. When the U-channels are installed, bearings on the U-channels roll along the horizontal slots until they drop down into the vertical slots, guiding the U-channel onto its kinematic mounts. The cooling lines (in braided jackets) and cables (ends wrapped in foil) for the L4–6 U-channels are visible to the left.

Second, the precision of tracks reconstructed with the SVT is limited by the precision with which the sensor positions are known; the survey provides the initial knowledge of the detector geometry, which must be good enough to allow efficient track reconstruction, and close enough to the true geometry for track-based alignment (see Section 4.2.1) to work.

The basic tool for the SVT survey is a coordinate-measuring machine (CMM). A CMM uses optical and/or touch probe measurements to locate target points in three dimensions.

Because the SVT is assembled in a modular way, with repeatable positioning at each stage, the survey can be done in stages as well. Each module is surveyed to find the positions of the silicon relative to the module mounting points. Each U-channel is surveyed to find the positions of the module mounting points relative to the U-channel. After the U-channels are installed in the SVT box, the SVT box is surveyed to find the positions of the U-channels relative to the SVT box; the U-channel kinematic mounts are adjusted during the survey to bring the U-channels to their nominal positions. Finally, after the SVT box is installed in the pair spectrometer vacuum chamber, the SVT box is surveyed to find the position of the SVT box relative to the rest of the detector.

3.2.3 Power and Data Acquisition

The power and data paths for the SVT are constrained. All signals must pass through a pair of 8-inch vacuum flanges at the upstream side of the analyzing magnet, so the number of signals has to be reduced. The closest available rack for the SVT power supplies and DAQ is 20 meters from the alcove where HPS is installed, so the analog APV25 output signals have to be converted to digital optical signals. Therefore HPS digitizes the signals and regulates the low-voltage power supplies inside the vacuum chamber, on frontend boards (FEBs) that are mounted on a cooling plate next to layers 1–3.

Each FEB can service four hybrids: a pair of L1–3 modules or a single L4–6 module. A single bundle of impedance-controlled twisted pair magnet wire connects

Figure 3.17: The SVT FEBS (frontend boards), mounted on their cooling plate. The red-and-white cables distribute low-voltage power to the FEBS; the red-and-black cables distribute high-voltage bias to the FEBS. The connectors for hybrid power, bias and data are covered with yellow Kapton tape. The mini-SAS connectors for high-speed data cables are at the bottom right of each FEB.

each FEB to its set of hybrids, carrying low-voltage power, high-voltage sensor bias, analog APV25 output signals, and digital control and trigger signals.

Each FEB digitizes the output signals from 20 APV25 chips. Each differential current signal is converted to a voltage by a preamp and digitized by a 14-bit ADC. Each FEB carries a single Xilinx Artix-7 FPGA, which packs the ADC data to be streamed off the FEB. The FPGA also monitors the hybrid state and configuration. All data and control signals are on a single high-speed data link, carried by a standard mini-SAS cable.

The FEBs also distribute low-voltage power to the hybrids. A single voltage supplied to the FEB is split into four independent voltages (one per hybrid) using a combination of switching and linear voltage regulators. This improves noise performance and reduces the number of voltages that must be passed into the vacuum chamber. High-voltage sensor bias is also routed through the FEBs, but is passed through directly.

Two sets of cables connect the FEBs to the vacuum chamber flanges: mini-SAS cables carrying digital signals, and twisted pair cables carrying low-voltage power and high-voltage bias. In both cases, the number of connections is too high for conventional vacuum feedthroughs. Instead, HPS uses “flange boards.” Each board has a vacuum side and an air side, to which connections are made using solder or standard connectors. The middle section of the board carries signal traces but is kept smooth; the board is then passed through a machined gap in the vacuum flange, and epoxy is poured to fill the space around the board: see Figure 3.19. The flange on beam-right carries two flange boards: one for low voltage and one for high voltage. The flange on beam-left carries four signal flange boards (one of which is shown in Figure 3.18), which use fiber transceivers on the air side to convert the electrical signals to optical signals.

The low and high voltages for the SVT are supplied by Wiener MPOD power supplies. The low-voltage supplies use sense lines to regulate the voltage actually supplied to the FEBs, and compensate for voltage drop in the cables.

The core of the SVT DAQ is the RCE platform. This is a general-purpose DAQ

Figure 3.18: One HPS signal flange board. The left side of the board operates in vacuum, and has three mini-SAS electrical connectors for high-speed data cables, which connect to the FEBs. The right side of the board operates in air, and has two MPO multi-fiber connectors for data and control signals.

Figure 3.19: Test of the flange board potting process.

system developed at SLAC. The system is built on the ATCA (Advanced Telecommunications Computing Architecture) industry standard and is housed in a standard ATCA crate.

The data processing, trigger handling, and event building is done on a COB (Cluster On Board) blade, which carries a set of daughterboards: four DPMs (Data Processing Modules) and one DTM (Data Transport Module). All of these are generic hardware that can be used for any experiment using the RCE platform. The only HPS-specific hardware is the (RTM) Rear Transition Module, which interfaces the COB to the fiber bundles connected to the signal flange boards. The SVT DAQ uses two fully loaded COBs and two RTMs. Each DPM contains two data processing nodes, each of which runs a Xilinx Zynq system-on-a-chip which integrates an ARM processor (running the Linux operating system) and an FPGA. The nodes can therefore process and reduce data on the FPGA at high speed, and perform high-level functions on the processor. The DTM contains a single node, which handles timing and trigger distribution. An implementation of the JLab TI (Trigger Interface) module is integrated in the DTM firmware for HPS.

On each trigger, the SVT DAQ reads out six 14-bit samples for every one of 23040 APV25 channels. This is far too much data to store in full, so the DAQ applies a data reduction threshold to select the channels to be recorded in the event. The threshold rule used in the 2015 run was for three samples to exceed the channel pedestal by at least three times the channel noise. The channel pedestal and noise are taken from offline calibrations, and are the mean and standard deviation of samples from the channel in the absence of an input signal.

3.2.4 Services

All the services needed to operate the SVT — motion, cooling, and power — must be supplied from outside the vacuum. These enter the vacuum through flanges on the upstream vacuum chamber.

Three linear shifts provide independent control of the upper and lower U-channels for layers 1–3, and the target frame. Each linear shift is driven by a stepper motor;

Figure 3.20: The RCE system. The COB (Cluster On Board) blade on the left hosts four RCE (Reconfigurable Cluster Element) daughterboards, which perform the data processing. The RTM (Rear Transition Module) on the right interfaces with fibers carrying data from the flange boards. The COB and RTM would normally be housed in an ATCA crate.

the three stepper motors are powered and controlled by a Newport XPS controller. The motor turns a lead screw to create precise linear motion, which is transmitted inside the vacuum chamber through a bellows.

The SVT uses two independent cooling loops. The silicon sensors need to be kept cold (below 0° C) to prevent reverse annealing. Reverse annealing is an effect where the performance of radiation-damaged silicon worsens after exposure to high temperatures. The FEBs only need cooling to remove the heat they generate, and their cooling loop can run near room temperature.

The SVT is cooled by circulating cooling fluid through copper lines that are pressed into the U-channels. As discussed in Section 3.2.2, the modules and half-modules are engineered to provide short and parallel cooling paths from the sensors and readout chips to the U-channels. A low-viscosity fluid is required to keep flow rates high at low temperature; HPS uses Novec 7000, a hydrofluoroether compound. The SVT chiller operates at a setpoint of -20° C. The cooling flow is split between the top and bottom halves of the SVT before it enters the vacuum chamber through ceramic feedthroughs; the L1–3 and L4–6 U-channels are connected in series.

The FEBs are cooled by circulating distilled water through copper lines pressed into the FEB cooling plate. The FEB components are kept in direct thermal contact with the cooling plate: the heat-generating components are on the side of the board that faces the cooling plate, and the plate has machined pockets lined with thermally conducting pads. The FEB chiller operates at a setpoint of 25° C.

Power provision to the SVT and FEBs is explained in Section 3.2.3.

3.3 Electromagnetic Calorimeter and Trigger

The HPS ECal is a homogeneous crystal calorimeter, containing 442 lead tungstate (PbWO_4) scintillating crystals [72]. The ECal is based on the CLAS Inner Calorimeter (IC); the crystals are reused from the CLAS IC, and the basic design of each crystal module is unchanged. The mechanical and electronic design of the HPS ECal was led by the same IPN Orsay group that built the CLAS IC.

The crystals have a trapezoidal shape, 16 cm long, with front faces $1.3 \times 1.3 \text{ cm}^2$

and back faces $1.6 \times 1.6 \text{ cm}^2$. Lead tungstate has a fast time response, which allows for good time resolution and high pileup tolerance. An avalanche photodiode (APD) is glued to the back face of each crystal for readout. Two LEDs (one red, one blue) are mounted on the front face of each crystal, and are used to monitor the stability of the readout gain and radiation damage to the crystals.

The ECal is split into top and bottom halves; each half contains 5 rows of 46 crystals, except for the innermost row which has 9 crystals removed to avoid the region of highest beam background (this gap in row 1 is known as the “electron gap”). The crystals are spaced as closely as possible. Since the crystals are tapered, they fan out to the sides and an incident particle will typically not hit a crystal head-on. Scintillator response is sensitive to the crystal temperature, so the ECal is surrounded by a thermal enclosure. The innermost rows of crystals are 2 cm from the beam plane, and the front plane of the ECal is 139.3 cm from the nominal target position.

Figure 3.21: Beam’s-eye-view of the ECal. 18 crystals (the “electron gap”) in the innermost rows are missing, to make space for the oval bulge in the ECal vacuum chamber (Figure 3.6). The thermal enclosure that surrounds the ECal is not shown.

Figure 3.22: One ECal crystal with its readout electronics. The optical fiber was used in the CLAS IC for calibration and monitoring, but was removed for HPS. HPS uses LEDs mounted directly in front of the crystals.

3.3.1 Readout and Trigger

Each ECal crystal has its own preamplifier. The preamp output signals are routed through a motherboard and out to cables, which connect to FADC250 digitizer boards in VXS crates.

The FADC250 (FADC for short) is a general-purpose readout board developed at Jefferson Lab. Each FADC board has 16 input channels, which are digitized continuously to 12-bit precision at 250 MHz. The digitized samples are stored in pipelines so they can be read out in response to a trigger. Several readout modes exist, and the FADC has the capability to perform pulse integration or emulate a constant-fraction TDC, but HPS uses the readout mode that outputs a window of 100 samples around the trigger time. This allows pulse fitting, for the best possible time resolution.

The FADC is also the first step in the trigger chain; the trigger chain is implemented entirely using VXS boards developed by Jefferson Lab as a general-purpose trigger framework for experiments at the lab. The digitized samples are passed to an algorithm that continuously looks for threshold crossings, and integrates the signal amplitude within a fixed window around each threshold crossing. The integrated amplitude is then converted to an energy using calibrated values of pedestal and

gain. Hits found in this way (crystal position, energy, time of threshold crossing) are reported to the Global Trigger Processor (GTP) every 32 ns.

The two GTP boards (one per ECal half) cluster the hits according to a simplified algorithm. According to this algorithm, a cluster is a set of hits in a 3×3 block of crystals, where the center hit is at least 50 MeV and is a local maximum in energy, and the other hits in the cluster are within 16 ns of the center hit. Both GTP boards report clusters (center crystal, center hit time, number of hits, total energy) to the single Subsystem Processor (SSP) board.

The Subsystem Processor (SSP) makes the trigger decision. For the HPS physics trigger, the clusters from the two halves of the ECal are paired up, and each pair is tested against the defined trigger cuts. If a pair meets the trigger requirements, a trigger is sent to the Trigger Supervisor board (TS), which then distributes the trigger to the Trigger Interface (TI) boards installed in all detector readout subsystems. If any subsystem is not ready to accept a trigger, or if the trigger follows too closely after another trigger (as defined by rules configured in the TS), the TS will reject the trigger.

A total of six trigger types are defined in the SSP. The HPS physics trigger is nicknamed “pairs-1” and is detailed in the next section. A looser variant of the same trigger is nicknamed “pairs-0” and is used for diagnostics. A single-cluster trigger tuned to trigger on electrons near the beam energy (elastic scatters) is called “singles-1,” and a looser variant is “singles-0.” To keep the rates of “pairs-0,” “singles-1” and “singles-0” triggers low as a fraction of the total trigger rate, they are prescaled so only one in 2^N (N in the range of 10 to 13, depending on the trigger) of each trigger is accepted. A pulser trigger generates a constant rate of triggers at 100 Hz regardless of detector conditions. A “calibration” trigger is used for cosmic ray calibration of the ECal when the beam is off, triggering on a pair of scintillator paddles mounted under the ECal.

The total trigger rate during normal operations in the 2015 run was roughly 19 kHz, of which 16.6 kHz was pairs-1.

Measurements of livetime (the fraction of time during which the DAQ is sensitive to triggers) are important for normalizing the data. Two measurements of livetime

are available. The first measurement uses the Faraday cup (see Section 3.1.1) and a gating signal output by the trigger system which indicates whether the DAQ is live. In addition to a scaler that records the Faraday cup measurement of integrated beam charge, a “gated Faraday cup scaler” records integrated beam charge with the DAQ live. The gated Faraday cup scaler is the standard measurement of integrated charge used by CLAS. The ratio of the gated to ungated scalers measures the fraction of the integrated beam charge which was accumulated with the DAQ live. The second measurement uses the pulser trigger: the pulser trigger fires at a constant rate of 100 Hz, so the number of pulser triggers actually recorded is a direct measurement of the livetime. Section 4.4 discusses how these measurements compare, and why the pulser livetime is used for this analysis.

3.3.2 Trigger Cuts

The HPS physics trigger is tuned to maximize efficiency for heavy photon decays: e^+e^- coincidences with total energy near the beam energy. For reasons explained below, there is no attempt to impose a minimum total energy requirement. The main single-particle backgrounds are elastic scatters (electrons with energy near the beam energy) and bremsstrahlung photons. The main two-particle background is wide-angle bremsstrahlung, where both the electron and photon hit the ECal.

A cluster reported to the SSP will not have energy equal to the incident particle, or even to the energy the full HPS reconstruction would compute for that cluster. First, the FADCs are calibrated such that the reported energy of a hit equals the energy deposited in the crystal; whenever a particle showers in the ECal, some energy is absorbed in the vacuum flange in front of the ECal, penetrates through the back of the ECal, or escapes through the gaps between crystals. These effects are compensated for in the full reconstruction but are ignored in the trigger. Second, the GTP imposes a 3×3 limit on the dimensions of a cluster, even though a single shower can deposit energy in a larger area. For these reasons, the energy thresholds used in the trigger cuts are rather lower (by roughly a factor of 0.8) than would be expected based on true particle energies.

Furthermore, many particles hit the ECal in or near the innermost row of crystals. For these particles, much of the shower is lost in the beam gap and the energy of the cluster can be much less than the particle energy. Since these may still be desirable events (many decays from lower-mass heavy photons land in the innermost row of crystals), the trigger cuts mostly avoid imposing strong minimum requirements on cluster energy or total energy.

The time coincidence between the top and bottom clusters is required to be within 12 ns. This allows for time walk; since the hit time reported by the FADC is a simple threshold crossing with a low threshold, time walk can be substantial.

The SSP applies an initial requirement on the energy and size of each cluster; it only considers clusters with energy between 54 and 630 MeV, and at least 1 hit. For a pair, the sum of the cluster energies is required to be between 180 and 860 MeV, and the difference is required to be less than 540 MeV. The minimum hit requirement is obviously trivial; the minimum cluster energy requirement is largely redundant with the GTP requirement on the center hit energy. Of the other requirements, the most important is the maximum energy sum requirement: most of the events rejected by this cut are pairs of two beam-energy electrons.

The coplanarity cut is intended to select e^+e^- events. On average, e^+e^- events should be symmetric around the beam axis (the line pointing along the beam direction at the target, starting at the beam spot). Events with two electrons will usually have both clusters on beam-right, and will fail this cut. The cut requires that the two clusters be on opposite sides of the beam axis: the azimuthal angle of each cluster is computed relative to the beam axis, and the difference between the two angles is required to be within $\pm 30^\circ$ of 180° . (In practice, this cut is implemented as a mask of pairs of cluster positions that satisfy this cut.)

The energy-distance cut acts on the lower-energy cluster of the pair, and imposes the requirement $E_{low} + (5.5 \text{ MeV/mm})r_{low} > 600 \text{ MeV}$, where E_{low} is the cluster energy and r_{low} is the cluster distance from the beam axis. Restated as $E_{low}/(600 \text{ MeV}) + r_{low}/(600/5.5 \text{ mm}) > 1$, the effect of this cut is to reject the pair if the lower-energy cluster is low-energy and close to the beam axis. This primarily rejects wide-angle bremsstrahlung events, where the photon is typically the lower-energy

particle, and is usually close to the beam axis. The cut also rejects beam-energy electrons that hit close enough to the ECal edge that most of the energy is lost; these would pass the cluster energy cuts, but are closer to the beam axis than a genuine low-energy charged particle would be.

Table 3.3: Trigger cuts for the “pairs-1” trigger.

Time difference	$ t_{top} - t_{bot} \leq 12 \text{ ns}$
Cluster energy	$54 < E < 630 \text{ MeV}$
Cluster size	$N_{hits} \geq 1$
Energy sum	$180 < E_{top} + E_{bot} < 860 \text{ MeV}$
Energy difference	$ E_{top} - E_{bot} < 540 \text{ MeV}$
Coplanarity	$ \phi_{top} - \phi_{bot} - 180^\circ < 30^\circ$
Energy-distance cut	$E_{low} + (5.5 \text{ MeV/mm})r_{low} > 600 \text{ MeV}$

Chapter 4

Event Reconstruction and Selection

The data from the HPS detector must be reconstructed and calibrated before it can be used in an analysis. Furthermore, the contents of the data must be understood and shown to be consistent with known physics before we can trust it in a search for new physics. This chapter describes that process.

4.1 Reconstruction

The HPS event reconstruction is based on the `org.lcsim` software toolkit [73]. The basic outputs of the event reconstruction are tracks in the SVT and clusters in the ECal. The reconstruction also matches clusters to tracks, and fits pairs of charged tracks to vertices.

The HPS coordinate system used in reconstruction is defined by the Hall B beam-line. The Z-axis points downstream along the Hall B axis, in the direction the beam takes before it reaches the HPS chicane. The Y-axis points up, and coincides with the magnetic field direction. The X-axis points to beam-left, in the bend plane. The origin is set at the intersection of the beam with the nominal target position.

Because the target is not in the center of the chicane, the HPS beam axis (the direction of the beam when it intersects the target) is rotated by 30.5 milliradians

relative to the Z-axis. The reconstructed momenta of particles at the target are distributed symmetrically around the HPS beam axis, not the Hall B axis. Since the HPS beam axis is the natural reference frame for reconstructed particles, reconstructed momenta and vertex positions are typically rotated by 30.5 milliradians to put them in the HPS beam frame.

4.1.1 Tracking

A charged particle track in a uniform field follows a helix. HPS uses the “perigee” parametrization for tracks, where the track position and direction are defined at the point of closest approach to the origin (the perigee) in the X-Z frame. The standard definition of the perigee parametrization fixes the magnetic field along the Z-axis; this is not the case for HPS, but we retain the standard parameter names (e.g. z_0). The curvature is defined as $C = -q/R$, where q is the charge and R is the helix radius; it is proportional to $1/p_{\perp}$, where $p_{\perp}$ is the component of momentum that is perpendicular to the magnetic field. The track position is defined by the distance of closest approach in the X-Z frame d_0 , and the height at closest approach z_0 . The track direction is defined by the angle from the Z-axis at closest approach ϕ_0 , and the slope at closest approach $\tan \lambda$.

For HPS tracks, the point of closest approach to the origin is always near the origin, and the track direction there is always close to the Z-axis. This means that d_0 and z_0 can be interpreted (approximately) as the track x and y at $z = 0$, and ϕ_0 and $\tan \lambda$ are approximately equal to θ_x (the angle of elevation from the Y-Z plane) and θ_y (the angle of elevation from the X-Z plane).

Multiple scattering in the silicon sensors means that different segments of the track may be described by different helices. A full HPS track fit (after GBL refit, as described in Section 4.1.1) has a different helix before and after every sensor: a track with hits on 12 sensors has 13 track segments. The track segment of interest for the analysis is the one at the target (before the first sensor), and this is what is used unless otherwise specified.

Hit Reconstruction

For each strip above threshold, the SVT readout stores a waveform of six samples at 24 ns intervals. The readout is aligned with the trigger such that typically, two or three of the samples are taken before the trigger, and three or four samples are taken after the trigger. This is enough points to both fit a pulse that is in time with the trigger, and a pileup pulse that may arrive before or after the trigger.

The APV25 response is modeled as a filter with four poles, three of them coincident. This is the transfer function, where $\tau_1 \approx 72$ ns and $\tau_2 \approx 12$ ns are the two time constants:

$$\tilde{F}_{4pole}(\omega) = \frac{1}{(1 + j\omega\tau_1)(1 + j\omega\tau_2)^3} \quad (4.1)$$

The pulse shape is the impulse response, which is the inverse Fourier transform of the transfer function. With t defined as the time after the start of the pulse, the impulse response is as follows:

$$F_{4pole}(t) = \frac{\tau_1^2}{(\tau_1 - \tau_2)^3} \left(e^{-\frac{t}{\tau_1}} - e^{-\frac{t}{\tau_2}} \sum_{k=0}^3 \frac{\left(\frac{\tau_1 - \tau_2}{\tau_1 \tau_2} t\right)^k}{k!} \right) \quad (4.2)$$

The pulse fitting algorithm attempts two different fits: a fit to a single pulse with unknown time and amplitude, and a fit to the sum of two pulses with different times and amplitudes. The χ^2 of the two fits are compared to decide whether the waveform is better fit by a single pulse or two pulses. The pedestal and the shaping times τ_1 and τ_2 are calibrated offline. It is possible to fit the pedestal on an event-by-event basis to better handle an effect where the pedestal shifts in high-occupancy channels; this part of the algorithm was not active for this analysis.

After pulse fitting, each strip above threshold in the SVT has one or two hits, with times (relative to the trigger) and amplitudes (in units of ADC counts). It is not crucial to express the hit amplitudes in terms of energy deposition, and gains in different channels do not vary by enough to affect the reconstruction, so amplitudes are left in terms of ADC counts even though gain information is available from calibrations.

Fitted hits are merged to form clusters if they appear to be from the same

particle. The position of a cluster is the amplitude-weighted centroid ($x_{cluster} = (\sum x_i a_i) / (\sum a_i)$) of the hits. The time of a cluster is instead weighted by the square of the amplitude ($t_{cluster} = (\sum t_i a_i^2) / (\sum a_i^2)$), because the time resolution is significantly worse for low-amplitude hits.

The clustering algorithm applies thresholds that are based on σ_{noise} , the noise level for each channel. A cluster starts with a “seed” hit with amplitude at least $4\sigma_{noise}$. Hits in adjacent strips are added to the cluster if they have amplitude at least $3\sigma_{noise}$ and are within 8 ns of the cluster time.

Track Finding and Refit

Before track finding, the 1-D strip clusters are paired up to form 3-D hits. For every cluster on every sensor, the clusters on the paired sensor are tested to see if the two clusters are close enough in space (the strip clusters must cross) and in time (the cluster times must be within 16 ns of each other) for them to correspond to a single particle. In addition, both clusters must be within 12 ns of the trigger time. If the two clusters pass these cuts, a 3-D hit is created at their intersection. Since the two sensors in a stereo pair are at a small distance from each other, there is a parallax effect such that the position of the intersection depends on the track angle. Therefore, the 3-D hit position is recalculated each time the hit is used in a track fit.

Tracks are found using a simple track finding algorithm called SeedTracker, which was developed for use in design studies for the SiD detector [74]. SeedTracker tests different combinations of hits as possible track candidates. The order in which it tests hits is defined by “strategies,” where each strategy defines the layers of the SVT as “seed,” “confirm,” or “extend” layers. First, one 3-D hit is chosen from each of the three seed layers. The initial track candidate is found by fitting a helix to the three hits. Second, the track candidate is tested against every 3-D hit in the single confirm layer; if any hit in the confirm layer is consistent with the track candidate, the hit is added to the track candidate and the track is refit (if no hit is found, the track candidate is discarded). Third, the track candidate is tested against every 3-D hit in the two extend layers; again, if a hit is consistent with the track, the hit is added to the track. If at least one hit is added to the track candidate in the extend stage

(for a total of five hits), the track candidate is accepted as a track. A single strategy should test every possible combination of hits from its seed and confirm layers. Four strategies are used because a single strategy will not find tracks for particles that miss one of its seed or confirm layers.

Multiple scattering deflects particles randomly at every silicon sensor, so a single helix cannot fit all the hits perfectly. A helix that describes the particle trajectory well in one part of the SVT will not describe hits in other parts of the SVT, and the position residuals for those hits will be large. In SeedTracker, a single helix is used to fit all hits, but the cumulative multiple scattering error is added to the position resolutions in later layers, reducing their importance in the fit. The resulting helix is a good fit to the particle trajectory at the target, and a poor fit downstream, but the fit χ^2 accounts for this. However, this fit does not make full use of the hit information; in particular, the momentum resolution is poor because the hit resolution in the downstream half of the SVT (layers 4–6) has been artificially worsened by multiple scattering.

The tracks found by SeedTracker are refit using the General Broken Lines (GBL) method [75, 76]. The GBL track model accounts fully for multiple scattering by treating a track as a collection of track segments connecting a series of scatterers; for HPS, each silicon sensor is a scatterer. A GBL fit attempts to minimize both the hit residuals and the kinks (scattering angles) between consecutive track segments. Since the full hit resolution is used at every layer, the momentum resolution is improved relative to the SeedTracker fit; in this regard GBL provides performance equivalent to a Kálmán filter. Since scatters in downstream layers of the SVT are recognized as scatters and do not perturb the track fit at the target, the track angle resolutions at the target are also improved relative to the SeedTracker fit.

4.1.2 Vertexing

Pairs of tracks are vertexed using a fast vertex fit that finds the best-fit vertex position and track parameters based on the track parameters and covariance matrices, and optional additional vertex constraints [77]. The momentum and mass of the vertex

are calculated using the fitted track parameters, at the fitted vertex position.

The HPS vertex reconstruction uses constraints on the x , y , and z location of the vertex. All constraints are limited by the vertex resolutions in those directions; the x and y constraints are limited by the beamspot size, and the z constraint is limited by knowledge of the target position, but these are all smaller than the vertex resolutions. The “ z -constrained” fit requires that the vertex be consistent with the z location of the target. The “target-constrained” fit requires that the vertex be consistent with the z location of the target, and with the x and y locations and sizes of the beam spot. The “beamspot-constrained” fit requires that the vertex position and momentum are such that the vertex momentum points back to the beamspot at the target z .

4.1.3 Clustering

For each ECal crystal above threshold, the ECal readout stores a waveform of 100 samples (at a fixed time offset relative to the trigger) at 4 ns intervals. The waveform is fit to the ECal pulse shape, which is modeled by the following function, where t is the time after the start of the pulse, and τ is the shaping time (roughly 2.4 ns for most channels) [78]:

$$F_{3pole}(t) = \frac{t^2}{2\tau^3} e^{-\frac{t}{\tau}} \quad (4.3)$$

HPS calls this pulse model a “three-pole” function because it is the inverse Fourier transform (impulse response) of a filter with three coincident poles, which has this transfer function:

$$\tilde{F}_{3pole}(\omega) = \frac{1}{(1 + j\omega\tau)^3} \quad (4.4)$$

The ECal pulse fitter does not consider the possibility of pileup: it always attempts to fit the waveform as the sum of a constant pedestal and a single pulse. Since the pulse fitter uses a χ^2 fit, it implicitly assumes Gaussian white noise (each sample has a Gaussian distribution, and there is no correlation between samples). The shaping time τ is held constant in the fit (an offline calibration finds τ for every crystal), but the pedestal, the start time of the pulse, and the amplitude of the pulse are fitted independently for every event.

The pulse amplitudes found by the pulse fitter are converted to energy using gain constants, which are calibrated for every channel. Two main methods are used for gain calibration: cosmic rays and elastic scatters [79].

Cosmic rays passing downwards through the ECal are recorded when the beam is off, using two scintillator paddles as a trigger source. The gains are calibrated to the known rate of energy loss for minimum ionizing particles. Since no beam is needed, this method is used for the initial gain calibration.

The cosmic ray signal is small compared to typical ECal hits, so data from the run is used to refine the gain calibration. Clusters from elastically scattered electrons are identified in the data, and the gain constants for every crystal are adjusted until the cluster energies match the energies seen in Monte Carlo simulations of elastic scatters. Since elastically scattered electrons have energy near the beam energy, the two calibrations cover the full range of energies of interest.

Hits in ECal crystals are combined to form clusters using an algorithm adapted from the CLAS Inner Calorimeter [80]. Hits that are local maxima in energy (no neighboring crystal has more energy) are identified as “seed” hits for clusters. Other hits are added to a cluster if a neighboring crystal contains a seed hit or a hit that has already been added to a cluster; this continues until all hits have been added to clusters. Finally, a time cut is applied, which removes hits from clusters if they are separated from the seed hit by more than 8 ns.

The cluster time is defined as the seed hit time: since the seed hit has the highest energy, it has the best time resolution. The cluster energy is initially set to be the sum of the component hit energies, and the position is initially set to be the logarithmically weighted centroid, according to this formula where $\omega_0 = 3.1$ is an empirical constant and x_i are the centers of the front faces of the crystals with hits:

$$x_{cluster} = \frac{\sum_{i=1}^N \omega_i x_i}{\sum_{i=1}^N \omega_i}, \quad \omega_i = \max \left(0, \omega_0 + \ln \frac{E_i}{E_{tot}} \right) \quad (4.5)$$

However, both the cluster energy and position must be corrected to make them reflect the actual energy and impact position of the incident particle. The sum of the hit energies does not account for the portion of the electromagnetic shower that

escapes off the side of the ECal, is absorbed in the vacuum flange in front of the ECal, or is lost through the back of the ECal. The incident particle generally does not enter the ECal parallel to the crystal axis because the crystals fan out to the sides, and particles can impact the ECal at a range of angles. Since the electromagnetic shower deposits energy at all depths in the ECal, the maximum energy deposition may not occur in the same crystal whose front face the particle hit.

Both the energy and position corrections are based on Monte Carlo studies: electrons, positrons, and photons of varying energies and angles are simulated, and the true energy and position are compared to the reconstructed values to find the corrections as a function of particle type, cluster energy, and cluster position [81]. Because the shower properties and direction both depend on particle type, the energy and position corrections depend on the particle type and are only applied after track-cluster matching.

4.1.4 Track-Cluster Matching

Tracks are matched to clusters by extrapolating the track to the ECal and comparing the cluster position to the extrapolated track position. Since the ECal is outside the magnetic field, this requires extrapolating tracks through a non-uniform field. The tracks are swum through the three-dimensional fieldmap for the pair spectrometer magnet, and the intersection with the front face of the ECal is compared with the positions of clusters in the ECal. Since the resolution of both the track extrapolation and the cluster position are energy-dependent, the distance between the extrapolated track position and the cluster position is converted to a χ^2 value that is scaled relative to the distance distribution that is seen in Monte Carlo samples. Each track is matched to the cluster with the best match χ^2 , provided that the χ^2 is less than 30; this threshold is arbitrary and high, and analyses are expected to impose a stricter matching requirement.

4.2 Tracker Alignment

The positions of the silicon sensors in the SVT are not perfectly known. The survey, described in Section 3.2.2, defines the sensor positions to within $100\ \mu\text{m}$, but this level of uncertainty is still enough to add significantly to the track fit χ^2 , create systematic shifts in the track parameters, and dilute resolution.

Alignment is the process of using the data to find the true sensor positions, and is done in two steps. Internal alignment finds the positions of the sensors relative to each other, with the objective of minimizing the track χ^2 . The track χ^2 is not sensitive to certain transformations (known as weak modes) of the entire tracker: for example, a uniform translation of all sensors would translate all tracks, but leave the track χ^2 unchanged. Global alignment uses tracks to measure the beam energy, beam angles and the beamspot position, compares them to their nominal values, and adjusts the alignment constants to bring the SVT to the correct position and orientation relative to the beamline.

4.2.1 Internal Alignment

The internal alignment is done using Millepede-II, a software package developed for fast alignment of large tracking detectors, and used by CMS [82, 83]. Millepede takes a set of tracks as input, and finds corrections for the sensor positions and orientations that minimize the χ^2 for all tracks. For each track, the input to Millepede contains all the contributions to the GBL track χ^2 (the hit residuals for each sensor and the track kinks), and the derivative of each contribution with respect to the track parameters and the alignment corrections; this information is computed as part of the GBL fit. Finding the optimal alignment corrections is then reduced to a system of linear equations, which Millepede solves by matrix inversion. The alignment corrections are applied relative to the sensor positions found in the SVT survey (Section 3.2.2).

Each sensor has a total of six possible alignment corrections: three translations and three rotations. Not all are important: for example, a translation along the strip direction has no effect unless a track is near the end of a sensor. The three possible translations, in order of decreasing importance, are along the measurement direction,

along the sensor normal, and along the strip direction. The three possible rotations, in order of decreasing importance, are around the sensor normal, around the strip direction axis, and around the measurement direction axis.

For this analysis, Millepede was only allowed to change the translations along the measurement direction. An improved alignment, which will add the rotations around the sensor normal, is in progress.

Millepede can be run multiple times, each run improving on the previous run. For each run, the data is reconstructed, Millepede is run on the reconstructed tracks, and the new alignment corrections are added to the corrections used for the reconstruction. Often different sets of alignment corrections are “floated” in each run. For example, it may not be desirable to let Millepede change all of the alignment corrections at once; instead, several layers of the SVT are floated in each run.

Weak modes occur when it is possible to shift all the sensors and all the tracks in a such a way that they remain consistent with each other. The simplest weak modes are those that shift one of the track parameters by a constant, for all tracks. Since there are five track parameters, there are five of these weak modes. Shifting tracks (horizontally or vertically) corresponds to shifting every sensor by the same amount. Rotating tracks (horizontally or vertically) corresponds to a linear shear of the SVT: every sensor shifts by a distance proportional to its distance from the target. Finally, a shift in the track curvature corresponds to a quadratic shear of the SVT in the horizontal direction: every sensor shifts by a distance proportional to the square of its distance from the target.

The curvature correction is found using elastically scattered electrons, which have a known momentum. The horizontal and vertical shift corrections are also set using elastically scattered electrons, by requiring that they extrapolate back to $x = 0$, $y = 0$ at the target z . The horizontal and vertical tilt corrections are set using electrons from Møller scatters, by requiring that their angles are consistent with the nominal beam direction.

4.2.2 Elastically Scattered Electrons

Elastic electron scatters from target nuclei ($e^-Z \rightarrow e^-Z$) are the dominant source of charged particles in the HPS acceptance. Since the electron mass is much less than the mass of a tungsten nucleus, the energy of these electrons is indistinguishable from the beam energy. The rate of elastically scattered electrons is high enough that they populate the entire HPS angular acceptance.

Because the energy of elastically scattered electrons is known, their curvature is known and they are used to calibrate the curvature correction.

Elastically scattered electrons are also a convenient population of events for other global alignment studies, such as finding the target position in z . All tracks should converge at the beamspot when projected to z_{target} , the z -location of the target. When viewed in the Y-Z plane and projected to any z , the extrapolated y of a track should follow Equation 4.6; if projected to z_{target} , all tracks should extrapolate to $y_{beamspot}$, the y -location of the beamspot. If evaluated at $z = 0$, Equation 4.6 reduces to Equation 4.7, which relates the two unknowns ($y_{beamspot}$ and z_{target}) to two observables (track slope and extrapolated y at $z = 0$). Figure 4.1 histograms (for tracks in the top half of the detector) the extrapolated y at the closest approach to the origin (roughly the extrapolated y at $z = 0$) against the track slope, and fits the histogram to Equation 4.7 to find $y_{beamspot}$ and z_{target} relative to the top half of the SVT.

$$y(z) = y_{beamspot} + slope \times (z - z_{target}) \quad (4.6)$$

$$y(z = 0) = y_{beamspot} - slope \times z_{target} \quad (4.7)$$

Global alignment corrections are chosen to bring the fitted values of $y_{beamspot}$ close to 0 for both halves of the detector: essentially, the two halves are moved vertically to bring their beamspot measurements into agreement. The fitted values of z_{target} are consistent with $z_{target} = -5$ mm on both top and bottom. This could mean that the whole SVT is shifted by +5 mm relative to its nominal position (highly unlikely) or that the target is shifted by -5 mm (possible). We believe the target is at -5 mm relative to its nominal position, and therefore choose not to apply global alignment

Figure 4.1: Top: track slope and y at closest approach, for elastically scattered electrons in the top half of the detector. The stripes are the result of single-strip hits in the first layer of the tracker. Bottom: each point is the center of a Gaussian fit to a vertical slice of the top histogram. The linear fit to the points produces estimates of y_{beamspot} and z_{target} .

corrections to move the SVT by +5 mm and bring z_{target} to 0.

The track extrapolation in the X-Z plane is used similarly to find global alignment corrections that bring x_{beamspot} to 0.

4.2.3 Møller Electrons

Møller scattering events can be identified in the data as e^-e^- pairs with momentum sum equal to the beam momentum.

Because the final state is fully reconstructed, the momentum sum should point in the beam direction. This is useful for monitoring the beam direction, but it is difficult to use this as a constraint on the global alignment because both halves of the SVT contribute to the measurement of the momentum sum. Fortunately, the two-body

kinematics of Møller scattering is sufficiently constrained that each electron on its own can provide a constraint, which applies independently to each half of the SVT.

For each electron, there is a strict relation between the energy E and the angle from the beam axis θ . To lowest order in the electron mass m_e , this relation is identical to Compton scattering:

$$m_e c^2 \left(\frac{1}{E} - \frac{1}{E_{beam}} \right) = (1 - \cos \theta) \quad (4.8)$$

This means that all Møller electrons with a known energy will have the same angle from the beam axis. The electron energy is well-measured, since the track curvature has been calibrated using elastic scatters. Selecting Møller electrons with energies in a narrow band, as in Figure 4.2, and extrapolating these tracks back to the target, results in a narrow ring of track angles centered at the beam direction.

Global alignment corrections are used to shift the track angles and center the ring around the nominal beam direction. Even though only a section of the full ring is captured in the HPS acceptance, there is enough of an arc to determine the horizontal position of the center, and the vertical position of the center can be determined using the radius of the ring (known from Equation 4.8).

4.3 e^+e^- Selection Cuts

After reconstruction, all possible e^+e^- pairs in the event are tested against a set of cuts. There is no explicit requirement that the electron and positron be the pair of particles that caused the trigger. There is also no fiducial requirement, since the inner edge of the detector acceptance is key for sensitivity to low-mass heavy photons. The base selection is intended to remove accidental coincidences from the pair sample; the pair sample should contain only events where the electron and positron originate in the same interaction.

The “pairs-1” trigger is the HPS physics trigger, described in Section 3.3.2. It is tuned to accept e^+e^- pairs, and is the overwhelming majority of the event rate (16.6 kHz out of 19 kHz).

Figure 4.2: Angles of Møller electrons (extrapolated to the target, and relative to the nominal beam direction) in the top half of the detector, in a narrow energy band centered at 400 MeV. Equation 4.8 predicts that Møller electrons at this energy should be emitted at 39.8 milliradians from the beam axis.

The electron and positron are required to be in opposite halves of the detector: this cut is implemented as a requirement that the y -coordinates of the two clusters have opposite signs. The trigger requires a top-bottom coincidence, so repeating the requirement as an event selection cut does not reduce the efficiency. This cut eliminates any possibility of confusion in the track or cluster reconstruction, since the hits from the electron and positron are guaranteed to be well separated.

Track-cluster matching is important for two reasons: the cluster time resolution is better than the track time resolution, and track-cluster matching eliminates many misreconstructed tracks. Two checks are done on the quality of the track-cluster matching for both particles. First, a cut is made on the χ^2 of the track-cluster match; this is a requirement on the distance between the cluster position and the track extrapolation to the ECal. Second, a cut is made on the track-cluster time difference.

Since the track and cluster times are referenced differently (the track time is relative to the trigger time, and the cluster time is relative to the start of the ECal readout window), a constant offset of 43 ns is subtracted.

Three more simple cleanup cuts are applied. A loose track quality cut is applied on the χ^2 of each GBL fit; this is only meant to reject very poor track fits. Elastically scattered electrons with $p(e^-) \approx E_{beam}$ are the main pileup background in the tracker, and are rejected with a maximum momentum requirement on electrons. A momentum sum cut rejects pairs with a momentum sum too far in excess of E_{beam} ; this further reduces the rate of random coincidences with elastically scattered electrons.

The last cleanup cut is a cut on the cluster time difference. This selects time coincidences. The track time difference could be used similarly, but the cluster time resolution is better.

Finally, a “radiative cut” is applied for heavy photon analyses. This is a minimum requirement on the momentum sum, at $0.8E_{beam}$. As shown in Section 2.5.1, most heavy photons and radiative tridents are produced with energy near E_{beam} ; the radiative cut keeps most of these and rejects the Bethe-Heitler tridents that dominate at low momentum.

Trigger type	“pairs-1” trigger
Run and event quality	see Section 4.4
Top-bottom requirement	$\text{sign}(y_{cl}(e^-)) \neq \text{sign}(y_{cl}(e^-))$
Track-cluster matching (position)	$\chi_{match}^2 < 10$
Track-cluster matching (time)	$ t_{cl} - t_{trk} - 43 < 4 \text{ ns}$
Track quality	$\chi_{trk}^2 < 50$
Elastics cut	$p(e^-) < 0.75E_{beam}$
Momentum sum cut	$p_{tot}(e^+e^-) < 1.15E_{beam}$
Cluster time coincidence	$ t_{cl}(e^-) - t_{cl}(e^+) < 2 \text{ ns}$
Radiative cut	$p_{tot}(e^+e^-) > 0.8E_{beam}$

Table 4.1: Base pair selection cuts for HPS. The effect of these cuts is shown in Figure 4.3.

4.3.1 Tuning Cuts

The cuts are tuned on the data, using the cluster time difference to separate “good” and “bad” events. Pairs with large cluster time difference ($|t_{cl}(e^-) - t_{cl}(e^+)| > 3$ ns) are accidental coincidences; pairs with small cluster time difference ($|t_{cl}(e^-) - t_{cl}(e^+)| < 1$ ns) are dominated by true time coincidences. An effective cut should reject pairs with large cluster time difference, and not pairs with small cluster time difference.

Figure 4.3 shows the effect of the cuts on the distribution of cluster time differences. This distribution is the sum of the distributions of true coincidences and random coincidences. The distribution of true coincidences is a single peak with shape determined by the time resolution, and the distribution of random coincidences is the sum of multiple peaks spaced by the 2 ns beam period, with a slowly varying envelope shaped by the efficiencies of the trigger and track reconstruction. If (as is the case for all cuts shown in the figure) a cut is not directly sensitive to the cluster time difference, and yet suppresses the outer peaks more than the central peak, it must be rejecting random coincidences.

The cluster time coincidence cut is applied after the other cuts, and selects only the central peak. The rate of random coincidences contaminating the central peak can be estimated from the outer peaks: the fraction of random coincidences after the event selection cuts is roughly 1.5%.

4.4 Run Quality, Event Quality, and Data Normalization

All the data used in the analysis should have the same detector conditions and efficiencies: this is ensured by appropriate selections of runs and events. The HPS data is divided into runs, which are roughly two hours each unless interrupted by DAQ problems. All “golden runs” included in this analysis have the same configuration for the HPS setup: the tungsten target with design thickness of $0.125\% X_0$, nominal beam current of 50 nA, the trigger cuts listed in Section 3.3.2, and the SVT position

Figure 4.3: Cumulative effect of the different pair selection cuts on pair-1 events passing the top-bottom requirement and radiative cuts. The ratio of out-of-time to in-time events (outer peaks to central peak) decreases as cuts are applied (going from the black to the magenta distributions).

at nominal (layer 1 silicon at 0.5 mm from the beam).

A more fine-grained run quality selection is necessary because even a golden run can include substantial amounts of data that were recorded without the detector being in the desired configuration: a good-quality run can have periods of poor-quality data. The most common situation is a beam trip. According to the procedures that were followed by shift workers in the 2015 run, the high-voltage bias to the SVT sensors was lowered when the beam was lost, and bias was only restored after the beam returned. Furthermore, if the trip was determined to be caused by the halo counter FSD (see Section 3.1.2), the SVT would be moved out to 1.5 mm from the beam, and only moved back in after the beam returned. The result is that the SVT would not be in its normal configuration after beam restoration: the bias would be low, and the SVT position might be away from nominal. Between beam restoration and the

restoration of SVT bias and position, data would be recorded but it would not be reconstructable, since at low bias the SVT would not see hits, and at a position away from nominal the SVT alignment would not be correct.

The SVT bias and position history are recorded in a database, and the database is used to find the time intervals when the SVT was in its nominal configuration. These time ranges are then used to select events with good run quality.

There are additional effects in the SVT that can affect event quality, and can be identified on an event-by-event basis. The SVT DAQ can enter an error state during a run, which affects all data for the rest of the run for a single sensor; the error state is flagged in each event, so these events are easy to reject. Roughly 40% of the data used in this analysis was recorded before the correct setting for the APV25 latency was found; with the incorrect setting, if an event occurs too late relative to the APV25 clock, the SVT hits will not pass the data reduction threshold. This affects 1/3 of the events in these runs, and the events are identified and rejected based on their trigger timestamps. Finally, a small fraction (3.5%) of events in every run have elevated noise in the SVT because the APV25 pipelines were being filled as the APV25 was outputting the digital header for a previous event. This effect is called “burst-mode noise” after the APV25 feature that allows a trigger to be received while a previous event is being read out. Events with burst-mode noise are identified by counting the number of isolated (no neighboring strips with hits) low-amplitude hits: most normal events have no such hits, and the noisy events have many (tens to hundreds), so this analysis rejects events with 3 or more isolated low-amplitude hits.

Since the expected rate of heavy photons can be normalized to the data, a precise measurement of the integrated luminosity is not critical to the analysis. Normalizing the data to the integrated luminosity is still essential for understanding the detector efficiencies and comparing the data to the cross sections of known processes. In order for the integrated luminosity to be useful, it should be corrected for the run and event quality selections.

The luminosity is the product of the beam current, target thickness, and experiment livetime. The beam current is measured by a Faraday cup, as described in Section 3.1.1. The target thickness is taken from measurements made during target

assembly. The experiment livetime is the product of the trigger livetime and the SVT DAQ efficiencies (latency and burst-mode noise) explained above.

Section 3.3.1 explains the two measurements of trigger livetime. One, from the Faraday cup, measures the fraction of the integrated beam charge which was accumulated with the DAQ live. The other, from the pulser trigger, measures the fraction of time for which the DAQ was live. The Faraday cup is more precise, both because it accounts correctly for fluctuations in beam current, and because it has less statistical uncertainty (the Faraday cup scaler rate is roughly 45 kHz at 50 nA, much higher than the 100 Hz frequency of the pulser trigger). However, the pulser livetime is recorded more frequently (every second, as compared to every 4-5 seconds), and this makes it easier to integrate the luminosity over the run quality time ranges. A comparison of the Faraday cup and pulser livetimes shows agreement at the 1% level or better, so the pulser livetime is acceptable.

The luminosity is integrated over the time ranges identified (as explained above) for good SVT bias, good SVT position, and SVT DAQ error status. The resulting integrated luminosity is appropriate for comparisons with Monte Carlo samples: it assumes an always-on beam, a detector always in its nominal state, and a DAQ that always takes good data.

According to this procedure, the total integrated luminosity for data recorded in the 2015 run with the SVT at 0.5 mm is 1166 nb^{-1} . The unblinded fraction of this data is a total of 119 nb^{-1} .

4.5 Rate Comparison with Monte Carlo

Simulated samples of all types of e^+e^- pair events of interest (tridents, heavy photons, and wide-angle bremsstrahlung) are generated using MadGraph/MadEvent version 4 [84]. Beam backgrounds, such as elastically scattered electrons, that are created in the target and create pileup in the detector, are simulated using EGS5 [85] and Geant4 [86]. Particle interactions in the HPS detector are fully simulated using SLIC [87]. The HPS readout and trigger are simulated using software based on the org.lcsim software toolkit [73].

The trident processes (radiative and Bethe-Heitler tridents, as described in Section 2.5.2) are simulated using MadGraph/MadEvent. One generator, called the “full-diagram” generator, incorporates all of the trident diagrams and their interference terms, and thus simulates (to lowest order) the combined rate of all $e^-Z \rightarrow e^+e^-e^-Z$. Since radiative tridents are important for normalizing the heavy photon rate, a separate generator is used to generate radiative tridents on their own. The heavy photon Monte Carlo generator, and the use of the radiative trident Monte Carlo sample, is explained in Chapter 5.

Wide-angle bremsstrahlung conversions (WABs) are the other important source of e^+e^- pairs; as explained in Section 2.5.2, the primary electron and the positron from the pair conversion can fake an e^+e^- pair. MadGraph/MadEvent is used to produce $e^-\gamma$ events because it handles the angle correlations and the momentum transfer to the nucleus correctly. The photon pair conversion is simulated in EGS5 (if in the target) or Geant4 (if in the SVT).

Ignoring random coincidences (as shown in Section 4.3.1, they are a small contribution), we believe the pair events in the data are fully simulated by the sum of the full-diagram trident and WAB generators. Figure 4.4 shows the momentum sum (the magnitude of the total pair momentum) for data and the two Monte Carlo samples. If the Monte Carlo samples fully describe the data, the data histogram should match the sum of the Monte Carlo histograms, but it does not. Instead the rate in data undershoots the rate in Monte Carlo at every value of the momentum sum.

Preliminary studies suggest that there is a significant inefficiency in the track reconstruction that is not reflected in the Monte Carlo. Since the heavy photon analysis only uses pairs with momentum sum above $0.8 \times E_{beam}$, it is most important to characterize the problem in that region of phase space.

We make the assumption that the inefficiency uniformly affects all high-momentum pairs. Scaling the data histogram by $1/0.65$, as in Figure 4.5, forces the data and Monte Carlo rates to agree at high momentum sum; this is equivalent to assuming that there is an 65% efficiency factor that is present in data but not in Monte Carlo. Since the rates at lower momentum sums do not agree, either the efficiency effect is momentum-dependent, or there is another explanation (such as a problem with the

Figure 4.4: Momentum sum $p(e^+e^-)$ for data (black), Monte Carlo samples of tridents (red) and wide-angle bremsstrahlung conversions (blue), and the sum of the red and blue histograms (magenta). All histograms are normalized to integrated luminosity such that the sum of all bins for a histogram equals the total cross-section in nanobarns. If Monte Carlo completely describes data, black and magenta histograms should match.

MadGraph/MadEvent generator) for the discrepancy at low momentum sums. In either case we must justify assuming uniform inefficiency at high momentum sum, when the behavior at lower momentum sum is clearly different. We must show that after scaling the data, the data and Monte Carlo agree at high momentum sum — the rates agree by construction, but other comparisons should also agree.

One such comparison is to check the relative rates of tridents and WABs. Tridents and WABs have significantly different kinematics, and might be expected to be affected differently by a momentum-dependent inefficiency. Since the MadGraph/MadEvent generators are different for the two Monte Carlo samples, a problem affecting one generator would change the relative rates of the two samples. Several

Figure 4.5: Same histograms as Figure 4.4, except that data is multiplied by $1/0.65$ to account for the estimated level of detector inefficiency at high momentum. If Monte Carlo completely describes data, black and magenta histograms should match. Note the agreement at high momentum sum (not remarkable, since it is by construction).

variables can be used to separate tridents from WABs.

One variable identifying WABs is the absence of a hit in layer 1 of the SVT for the positron track. A trident positron should always (aside from detector inefficiencies) have a layer 1 hit, since the positron is created in the target and the angular acceptances are roughly the same for all layers. WAB conversions often occur in layer 1 (or later), and the positron track will not have a layer 1 hit unless the conversion occurs in the upstream sensor of the layer 1 stereo pair. Figure 4.6 shows the momentum sum for pairs with and without layer 1 hits on the positron track. In both, data and Monte Carlo show good agreement at high momentum sum. This indicates that data and Monte Carlo agree on the relative rates of tridents and WABs; the same efficiency applies to both.

WABs can also be separated by looking at the positron track's distance of closest

Figure 4.6: Same histograms as Figure 4.5, but the left plot requires that the positron track have a layer 1 hit, which rejects most WAB conversions, and the right plot requires that the positron track not have a layer 1 hit, which rejects most tridents. Note the continued agreement between data and Monte Carlo at high momentum sum.

approach. If the pair conversion happens in silicon and not the target, the positron track will not extrapolate to the target, but will instead miss to the side in the X-Z plane. This can be seen in Figure 4.7, which plots the positron track's distance of closest approach for pairs with high momentum sum: the distribution for trident Monte Carlo events is symmetric around the beam spot, but the distribution for WAB Monte Carlo events is strongly skewed. Data agrees with the sum of the two Monte Carlo samples, which again supports the conclusion that Monte Carlo describes the data well at high momentum.

Finally, WABs can be separated based on the direction of the pair momentum. Since the WAB positron is created in the pair conversion of a photon, the electron from the pair conversion is created at a similar angle in the Y-Z plane; both are on the same side of the beam plane. This means the electron from the pair conversion is never part of the e^+e^- pair that passes the selection cuts in Section 4.3, and therefore that the positron in a WAB event is always on the same side of the beam plane as the missing electron. Since the momenta of the two electrons and the positron approximately sum to the beam momentum (the momentum transfer to the nucleus is typically not large), this means that the momentum of the positron is on the opposite side of the beam

Figure 4.7: Distance of closest approach in the X-Z plane for the positron track, only plotting pairs with momentum sum above 0.8 times the beam energy (radiative cut). This is centered at 0 (where the beam spot is) if the positron comes from the target, but is usually positive for WAB positrons created downstream of the target.

plane from the total momentum of the e^+e^- pair. In contrast, radiative tridents have no such correlation, and Bethe-Heitler tridents have the opposite correlation. Figure 4.8 plots (for pairs with high momentum sum) the y -component of the momentum sum times the sign of the y -component of the positron momentum, and shows that this is typically negative for the Monte Carlo WAB sample and positive for the Monte Carlo trident sample. Again, data agrees with the sum of the two Monte Carlo samples, which again supports the conclusion that Monte Carlo describes the data well at high momentum.

We conclude that after the radiative cut, the Monte Carlo samples describe the data well, except that the rate in data is 0.65 times the rate expected from Monte Carlo.

Figure 4.8: $p(e^+e^-)_y \text{ sign}(p(e^+)_y)$, only plotting pairs with momentum sum above 0.8 times the beam energy (radiative cut). This is positive if the pair momentum points in the same direction as the positron.

Chapter 5

Search for Displaced Vertices

The HPS vertexing search is a low-background search for heavy photons with displaced vertices. This search is sensitive to heavy photons with low coupling, and therefore low production but long decay lengths.

The pair selection for the vertexing search requires cuts in addition to the basic cuts described in Section 4.3. These vertex quality cuts are tuned to suppress the tails of the vertex distribution, and are described in Section 5.2.

The pairs passing cuts are reduced to a 2-D dataset of points (m, z) , where each point is the mass m and vertex location z of a pair: see Figure 5.1. The vertex z is the distance along the HPS beam axis, and is defined relative to the nominal target position; as explained in Section 4.2.2, the actual target position is at $z_{target} \approx -5$ mm. To test for a heavy photon with $m_{A'} = m$, a resolution-limited mass slice ($|m - m_{A'}| < 1.4\sigma_m$) is taken from the 2-D data set, and the vertex z of the pairs in this mass slice are plotted as in Figure 5.2.

A cut is made at a minimum $z > z_{cut}$ to remove the large background from prompt e^+e^- pairs. z_{cut} (plotted in Figure 5.3) is set so the expected number of events past z_{cut} is 0.5 according to a model of the prompt pairs distribution, described in Section 5.3.3. This minimizes the prompt backgrounds and maximizes the heavy photon signal. The rectangular region formed by the mass and z cuts is the signal box for the heavy photon with mass $m_{A'}$.

If a heavy photon exists at some $m_{A'}$, heavy photon decays will appear in the data

Figure 5.1: Dataset for the vertexing analysis. The core of the vertex distribution is the horizontal band at $z = -5$ mm; a heavy photon signal would appear as a vertical band extending upwards from the core.

with a Gaussian mass distribution centered at $m_{A'}$, and a distribution of locations that extend from z_{target} to positive z . These events will appear in those mass slices that are close enough to $m_{A'}$ that they substantially overlap the Gaussian.

Two types of analyses are performed on the data in the signal box. The first is a cut-and-count analysis, where the number of pairs found in the signal box is compared to the number expected; this analysis gives the significance of any observed excess over background, and is described in Section 5.4. The second analysis uses the optimum interval method to set an upper limit on the heavy photon production rate, and is described in Section 5.5.

Figure 5.2: Vertex z (units of mm) for a resolution-limited mass slice (in data) centered at $m_{A'} = 35.5$ MeV. A heavy photon signal would appear as a low-level tail extending to higher z (longer than the exponential tail of the blue curve, which is the prompt vertex distribution described in Section 5.3.3). For this mass slice, $z_{cut} = 26.6$ mm; there are four events with $z > z_{cut}$.

Figure 5.3: The value of the cut $z > z_{cut}$, as a function of the center of the mass slice. The number of events per mass slice decreases and the z resolution improves as the mass increases; both effects tend to reduce z_{cut} .

5.1 Heavy Photon Monte Carlo

Heavy photon Monte Carlo samples are generated for values of $m_{A'}$ spaced out across the region of interest: 15, 16, 17, 18, 19, 20, 22, 24, 26, 28, 30, 35, 40, 50, 60, 70, 80, and 90 MeV. The generator is similar to those described in Section 4.5. It uses MadGraph/MadEvent, and fully simulates the momentum and angle spectra of the produced heavy photons. In order to get complete coverage of $z > z_{target}$ for the purposes of the vertexing analysis, the decay vertices were displaced (in the direction of the heavy photon momentum, and accounting for variations in γ) according to an arbitrary decay length of $c\tau = 1$ mm.

5.2 Vertex Cuts

Table 5.1: The full set of pair selection cuts for the vertexing analysis. Cuts carried over from the base selection are on top; cuts specific (or tightened) for vertexing are on bottom.

Trigger type	“pairs-1” trigger
Track-cluster matching (position)	$\chi^2_{match} < 10$
Track-cluster matching (time)	$ t_{cl} - t_{trk} - 43 < 4$ ns
Cluster time coincidence	$ t_{cl}(e^-) - t_{cl}(e^+) < 2$ ns
Top-bottom requirement	$\text{sign}(y_{cl}(e^-)) \neq \text{sign}(y_{cl}(e^+))$
Elastics cut	$p(e^-) < 0.75E_{beam}$
Momentum sum cut	$p_{tot}(e^+e^-) < 1.15E_{beam}$
Radiative cut	$p_{tot}(e^+e^-) > 0.8E_{beam}$
Layer 1 requirement	layer 1 hits for both tracks
Track quality	$\chi^2_{trk} < 30$
Beamspot constraint	$\chi^2_{bsc} < 10, \chi^2_{bsc} - \chi^2_{unc} < 5$
Layer 1 isolation	see text
Momentum asymmetry	$ p(e^-) - p(e^+) / (p(e^-) + p(e^+)) < 0.4$
Positron DOCA	$d_0(e^+) < 1.5$ mm

Broadly speaking, there are three types of events that the vertexing cuts are meant to eliminate: layer 1 scatters, mishits, and wide-angle bremsstrahlung pairs.

Multiple scattering in layer 1 is the source of the Gaussian core of the vertex distribution, and also plays a role in the tails. If one of the particles scatters in the first layer of the SVT, the track parameters at the vertex will be shifted. The distribution of scattering angles is Gaussian at small angles where the scattering process is approximated by a random walk, but at large angles, the distribution is described by the Molière distribution, which approaches the power-law Rutherford scattering distribution at larger angles.

Mishits happen when a particle scatters in the second (or later) layer of the SVT. The scatter can shift the track so that a layer-1 hit from a different particle is more in line with the hits from layers 2 on; the track fit will then add the wrong hit to the track, and the track parameters at the vertex will be shifted.

The electron and positron of a wide-angle bremsstrahlung pair do not have the

same origin: the electron comes from the bremsstrahlung interaction in the target, and the positron comes from a pair conversion that may happen in the target or in the SVT. If the pair conversion happens in the SVT, the positron will not extrapolate to the target, and the reconstructed vertex may be pulled away from the target (the z distribution is wider than the inherent vertex resolution). Because the positron is typically collinear with the photon, the effect on vertex z is typically small, but these events are an identifiable component of the vertex distribution tails.

Figure 5.4 shows the effect of the vertex cuts on a mass slice in data. 60% of vertices near z_{target} (presumed to be well-reconstructed) pass the vertex cuts, and this cut efficiency is similar for different masses. The cut efficiency for displaced vertices (in Monte Carlo samples of heavy photons) is similar and does not degrade significantly with increasing z . The individual cuts are described below.

Figure 5.4: Cumulative effect of the different vertex cuts on events in the mass slice centered at 30 MeV, passing base event selection cuts and meeting the layer 1 hit requirement. The tails at high z are greatly reduced as cuts are applied (going from the black to the magenta distributions).

5.2.1 Layer 1 Requirement

The track reconstruction will find tracks with hits in any 5 out of the 6 layers. This means that tracks can be reconstructed without layer 1 hits (as long as they have hits in all other layers). Tracks without layer 1 hits have degraded mass and vertex resolution. Furthermore, the tails of the vertex distribution extend to larger values of z . For these reasons, it is not possible to use tracks with and without layer 1 hits as part of the same data set. As discussed in Section 5.5.2, a full analysis of the 2015 data will include pairs without layer 1 hits as a separate data set. For simplicity, this analysis is limited to pairs with layer 1 hits on both tracks.

The background from wide-angle bremsstrahlung conversions is also significantly reduced by this cut. In order to create charged tracks, the bremsstrahlung photon must convert in the target, either layer 1 sensor, or early enough in the upstream layer 2 sensor to make a hit there. But in order to create charged tracks with layer 1 hits, the photon must convert in the target or the upstream layer 1 sensor. The silicon sensors ($0.35\% X_0$) are significantly thicker than the portion of the target traversed by the average photon (half of $0.125\% X_0$), so requiring that the track have a layer 1 hit cuts this background by roughly a factor of three.

The layer 1 requirement has a significant effect on the reconstruction efficiency for long-lived, low-mass heavy photons. As seen from the target, all layers of the SVT have their inner edges at ± 15 mrad vertical angle from the beam plane. As seen by a heavy photon decaying downstream of the target, layer 1 is at a significantly larger vertical angle than the others, and so the minimum $m_{A'}$ needed to hit layer 1 is larger. Put another way, this means that the maximum z for detecting a heavy photon of given $m_{A'}$ is smaller if layer 1 is required; this is shown in Figure 5.5. The impact of this effect is discussed in Section 5.5.2.

5.2.2 Track Quality

As explained in Section 4.3, the track χ^2 is not a very selective cut when studied as part of the base event selection (where the objective is to reject accidental coincidences). However, poor track fits can lead to poor vertex fits, and so events with

Figure 5.5: Efficiency curves for $m_{A'} = 40$ MeV. Left: requiring layer 1 hits for both tracks. Right: full HPS acceptance.

poor track χ^2 do contribute to the tails of the vertex z distribution. This is shown in Figure 5.6.

5.2.3 Beamspot Constraint

The beamspot constraint uses the vertex fit to identify the common situation where one track points back to the beamspot at the target as it should, but the other does not. If the second track (due to scattering, mishits or any other effect) intersects the beam axis downstream of the target, the reconstructed vertex will be pulled downstream along the first track. The effect is that the reconstructed z will be pulled downstream of the target, but the reconstructed vertex y is also pulled up or down from the beam axis in a way that is inconsistent with a true displaced decay. The beamspot constraint is meant to identify these inconsistencies.

As explained in Section 4.1.2, the vertex reconstruction can use a “beamspot-constrained” fit where the vertex momentum is required to point back to the beamspot at z_{target} . The χ^2 of this vertex fit is the sum of two factors: the consistency of the vertex (how close the tracks are to intersecting each other) and the consistency with the beamspot constraint (how close the vertex is to pointing back to the beamspot). The “unconstrained” vertex fit χ^2 contains only the first factor.

Cuts are applied on both the beamspot-constrained vertex χ^2 and the difference

Figure 5.6: Effect of the track quality cut on events in the mass slice centered at 30 MeV, and passing all other cuts.

between the beamspot-constrained and unconstrained vertex χ^2 . The effect of these cuts is shown in Figure 5.7.

5.2.4 Isolation Cut

The isolation cut rejects possible mishits by looking at the other hits in layer 1 of the SVT. If a mishit on a track pulls the vertex downstream, the real hit from the particle will be close to the hit that was mistakenly associated with the track, but further away from the beam plane. Since the track parameters at the vertex are most strongly determined by the layer 1 and layer 2 hits, and the lever arm from the target to layer 2 is twice that from layer 1 to layer 2, typically the track y at the target will shift by double the distance between the real hit and the mishit. This geometry is illustrated in Figure 5.8.

Therefore an “isolation” value is calculated for each of the layer 1 hits, as the

Figure 5.7: Effect of the beamspot constraint cuts on events in the mass slice centered at 30 MeV, and passing all other cuts.

distance to the closest hit in the outwards direction. If the isolation is less than 0.5 times the track's distance of closest approach to the beamspot in the $y - z$ plane, the hit on the track may be a mishit and the closest hit might be the real hit. If this is the case for any of the four isolations (one for each layer 1 sensor), the pair is rejected. The effect of this cut is shown in Figure 5.9.

5.2.5 Momentum Asymmetry and Positron DOCA

Two cuts are specifically designed to reject wide-angle bremsstrahlung pairs (the layer 1 requirement also significantly reduces the rate of WABs).

Because the electron from a bremsstrahlung interaction typically carries more momentum than the photon, the electron in a WAB pair usually has higher momentum than the positron. Heavy photons and radiative tridents have a symmetric distribution of electron and positron momentum, so a cut is used to reject pairs where the

Figure 5.8: Illustration of a track with a mishit. This view shows the Z (beam direction, beam travels left to right) and Y (magnetic field direction) axes, and three SVT layers with hits (circles). The true particle trajectory (black) starts at the target, is scattered at layer 2, and continues on to layer 3. Because the true trajectory is not straight, a better track fit can be made using a hit (red circle) from a different particle. The resulting track fit (red line) misses the target and, if used in a vertex fit, will lead to a best-fit vertex downstream of the target.

electron has much higher momentum than the positron.

If the pair conversion happens in the SVT, the positron track will curve wide of the target since the positron is roughly collinear with the photon, which does not bend in the magnetic field. This appears in the track parameters as a positive DOCA (distance of closest approach) in the X-Z plane. Therefore, pairs with large positive positron DOCA are rejected.

The effect of both cuts is shown in Figure 5.10.

5.2.6 Tuning Cuts

Data is used to tune the cuts to keep pairs with z close to z_{target} , and reject pairs with large positive z . The intent is to maximize the cut efficiency for pairs with well-reconstructed z , and minimize the tails of the z distribution for prompt pairs (thus minimizing z_{cut} , and maximizing the usable signal region for heavy photons).

In general, tuning cuts on the data can introduce bias. In this case, only the unblinded 10% of the full dataset is being used (and will be used, even after the data is unblinded) for tuning cuts; even if a heavy photon is present in the tuning data, the number of expected events is negligible. Also, since the cuts are not tuned for

Figure 5.9: Effect of the isolation cut on events in the mass slice centered at 30 MeV, and passing all other cuts. By design the isolation cut has little impact on vertices with $z \approx z_{target}$, since the tracks in those pairs pass close to the beamspot.

individual mass slices, any possible heavy photon signal will not be as prominent in the tuning process as it would be in the analysis.

Heavy photon Monte Carlo is used to check that none of the cuts have an unexpected adverse effect on efficiency for displaced vertices.

Figure 5.10: Effect of the WAB rejection cuts on events in the mass slice centered at 30 MeV, and passing all other cuts. The z distributions of events rejected by these cuts are wider than the events that pass, showing that while wide-angle bremsstrahlung conversions do not have systematic shifts in vertex z , they contribute to the tails of the z distribution.

5.3 Fit Inputs

To test for a heavy photon at mass $m_{A'}$ and coupling ϵ^2 , the signal, background, and resolutions must be modeled as inputs to the analyses.

The mass cut is set to keep only events with $|m - m_{A'}| < 1.4\sigma_m(m_{A'}, z)$, where σ_m is the mass resolution expected for heavy photons with mass $m_{A'}$. The mass resolution depends on the mass and the vertex position, and is estimated using Monte Carlo: this is explained in Section 5.3.1. This mass window is chosen to accept a large fraction of the signal events, without accepting too many background events. A window of $\pm 1.4\sigma$ (which accepts 83.8% of the signal events) optimizes significance for high-statistics, high-background experiments where significance is proportional to $S/\sqrt{B}$. The optimality of this window is not exact for low statistics (where the Poisson distribution cannot be safely approximated as a Gaussian), but is still approximately true.

After cutting on m , the distribution of events in z is studied. An additional cut is made to keep only events with $z > z_{cut}$, rejecting the region where the background strongly dominates and there is no sensitivity to a signal. z_{cut} is chosen such that only 0.5 events are expected past $z = z_{cut}$, based on the fitted shape of the background distribution. The amplitude of the background distribution is taken from the peak of the vertex distribution; the shape is taken from Monte Carlo as described in Section 5.3.3.

Setting limits requires knowledge of the expected signal distribution, $S(z; m_{A'}, \epsilon^2)$. This is described in Section 5.3.2.

5.3.1 Estimating the Mass Resolution

The mass resolution σ_m for a e^+e^- pair depends on the momentum resolutions $\sigma_{p_{e^+}}, \sigma_{p_{e^-}}$ for the two particles, the resolution σ_θ of the opening angle, and their covariances. Since the opening angle and momentum measurements predominantly rely on different parts of the SVT (opening angle on the upstream half of the SVT, momentum on the downstream half), their covariance is negligible. Neglecting the electron mass, and using the small-angle approximation for the opening angle,

$$m = \sqrt{(1 - \cos \theta)p_{e^+}p_{e^-}} \approx \frac{1}{\sqrt{2}}\theta\sqrt{p_{e^+}p_{e^-}} \quad (5.1)$$

$$\sigma_m \approx \frac{1}{\sqrt{2}} \left(\theta \frac{\sqrt{p_{e^+}p_{e^-}}}{2} \left(\frac{\sigma_{p_{e^+}}}{p_{e^+}} + \frac{\sigma_{p_{e^-}}}{p_{e^-}} \right) + \sigma_\theta \sqrt{p_{e^+}p_{e^-}} \right) \quad (5.2)$$

θ is the only term in this expression with a strong dependence on m or z : it is proportional to m . Because the momentum resolution is dominated by multiple scattering, the fractional momentum resolutions $\frac{\sigma_{p_{e^+}}}{p_{e^+}}$ and $\frac{\sigma_{p_{e^-}}}{p_{e^-}}$ do not depend strongly on the momentum; nor do they depend on the track angles. The opening angle resolution is determined by the resolutions for the two track slopes, which do not depend strongly on the track slopes, so σ_θ is roughly constant. The conclusion is that σ_m is expected to depend linearly on m , and not at all on z .

Mass resolution is determined using the Monte Carlo samples of reconstructed heavy photons described in Section 5.1. For each $m_{A'}$, the residual between the reconstructed mass and true mass is plotted against the true z . The width of the residual distribution at each z is fitted with a Gaussian to get the mass resolution at that z , $\sigma_m(m_{A'}, z)$. The mass resolution is fitted with a first-order polynomial in z : $\sigma_m(m_{A'}, z) = a_0(m_{A'}) + a_1(m_{A'})z$. Then the polynomial coefficients are themselves fitted with first-degree polynomials in $m_{A'}$: $a_0(m_{A'}) = a_{00} + a_{01}m_{A'}$, $a_1(m_{A'}) = a_{10} + a_{11}m_{A'}$. The result of this procedure is a model for the mass resolution: $\sigma_m(m_{A'}, z) = (a_{00} + a_{01}m_{A'}) + (a_{10} + a_{11}m_{A'})z$.

The fitted values of a_{10} and a_{11} are consistent with 0, as expected. After this procedure, this is the mass resolution model (including statistical uncertainties in the

last digits of the coefficients):

$$\sigma_m(m_{A'}, z) = 0.00072(2)\text{GeV} + 0.0382(4)m_{A'} \quad (5.3)$$

Møller scattering is one check of the mass resolution. As explained in Section 4.2.3, pairs of electrons from Møller scattering have a fixed invariant mass equal to the center-of-mass energy. The width of the Møller mass distribution is therefore a useful check. Figure 5.13 shows the Møller mass distribution in data, which has $\sigma_m = 2.168$ MeV. As shown in Figure 5.12, this is within 10% of the heavy photon mass resolution from Monte Carlo (1.973 MeV). The difference between data and Monte Carlo resolutions is believed to be due to the incomplete SVT alignment, for which work is still in progress.

Figure 5.11: Mass resolution vs. z for 40 MeV heavy photons in Monte Carlo. Top: mass residual (difference between the reconstructed mass and the true mass) vs. z for 40 MeV heavy photons. Bottom: the widths of Gaussian fits to vertical slices of the top distribution (blue points), and a linear fit to the points (red line). The slope of the linear fit is consistent with 0, showing that the mass resolution is constant with respect to z .

Figure 5.12: Mass resolution vs. $m_{A'}$. The blue marker is the Møller mass resolution in data.

Figure 5.13: Distribution of reconstructed invariant mass of Møller pairs, with a Crystal Ball fit showing the mass resolution.

5.3.2 Estimating the Signal Distribution

The expected distribution of signal events can be expressed as $\mu_{exp}s(z)$, where μ_{exp} is the expected number of reconstructed heavy photon events after the mass and z cuts, and $s(z)$ is the probability density function (normalized to unit integral) of the z locations. The distribution is estimated as the following product of terms, where each term can depend on $m_{A'}$ and ϵ^2 :

$$\mu_{exp}s(z) = (N_{A'}\epsilon_{reco}(z_{target})) \frac{e^{\frac{z_{target}-z}{\gamma c\tau}}}{\gamma C\tau} \frac{\epsilon_{reco}(z)}{\epsilon_{reco}(z_{target})} \epsilon_{cut}(z) \quad (5.4)$$

$N_{A'}$ is the number of heavy photons produced in the target. The exponential function is the distribution of decays along z , and is normalized to 1. $\epsilon_{reco}(z)$ is the efficiency to detect and reconstruct an e^+e^- pair produced at a given z . $\epsilon_{cut}(z)$ is the efficiency of the mass and z cuts for a heavy photon with mass $m_{A'}$ decaying at a given z . It equals 0 for $z < z_{cut}$, and 0.838 for $z \geq z_{cut}$.

In principle, this distribution should be smeared by σ_z , the resolution of the vertex position: this is not done since the signal distribution varies slowly on the scale of σ_z (which is 3-6 mm, depending on m).

Production Rate and Radiative Fraction

$N_{A'}\epsilon_{reco}(z_{target})$ is estimated using data and Equation 2.5, which shows that $N_{A'}$ is linked to $\frac{dN_{rad}}{dm}$, the number of radiative tridents produced with masses around $m_{A'}$. The data gives $\frac{dN_{e^+e^-}}{dm}\epsilon_{reco}(z_{target})$, the number of e^+e^- pairs produced and detected in a mass window around $m_{A'}$; some fraction of these are radiative tridents. The fraction is estimated using Monte Carlo.

The Monte Carlo samples described in Section 4.5 are used to calculate the normalized cross-sections (after detector and reconstruction efficiencies, and all cuts) for radiative tridents only, for all tridents, and for wide-angle bremsstrahlung conversions. The radiative trident fraction is the ratio of the cross-section for radiative tridents to the total cross-section for e^+e^- pairs.

Figure 5.14: Left: rates of processes producing e^+e^- pairs, with the sum in black. Right: the radiative fraction, which is calculated by dividing the black histogram by the red histogram.

Decay Length

The decay length is calculated using Equation 2.4, which gives the lifetime τ for a given $m_{A'}$ and ϵ^2 . The boost γ equals $E_{beam}/m_{A'}$ if the heavy photon takes the full beam energy, but this is not completely accurate; in reality the average boost is slightly less than the maximum.

Monte Carlo samples are used to find the correct distribution of decay lengths. As shown in Figure 5.15, the decay z is plotted for an MC sample of heavy photons with mass $m_{A'}$ and an arbitrary lifetime, with the requirement that the heavy photon momentum be at least $0.8E_{beam}$ (matching the analysis cut), and fit with an exponential. The decay constant of the exponential is compared to the $\gamma c\tau$ that would be expected from $\gamma = E_{beam}/m_{A'}$; this shows that the typical γ is roughly $0.95E_{beam}/m_{A'}$.

Efficiency

The efficiency ϵ_{reco} for reconstructing a heavy photon decay depends on $m_{A'}$ and the decay z . The measurement of the radiative trident rate implicitly includes a factor of $\epsilon_{reco}(m_{A'}, z_{target})$, so Monte Carlo is only needed to estimate $\frac{\epsilon_{reco}(m_{A'}, z)}{\epsilon_{reco}(m_{A'}, z_{target})}$, the efficiency falloff as a function of vertex displacement. This assumes that any detector-based inefficiencies not represented in the Monte Carlo are independent of

Figure 5.15: Top: distribution of decay z for heavy photons ($m_{A'} = 40$ MeV, $c\tau = 1$ mm) carrying at least 80% of the beam momentum. Bottom: the momentum distribution for the heavy photons.

z .

The efficiency falls off with z because the further downstream the decay, the larger the opening angle in the Y-Z plane necessary to hit layer 1 of the SVT. At some cutoff value of z it is no longer possible for both the electron and the positron to hit layer 1; the efficiency starts to fall off well before that cutoff. This effect is more severe for lower $m_{A'}$, where the opening angle is smaller.

The efficiency is determined using the Monte Carlo samples of generated and reconstructed heavy photons described in Section 5.1. For each $m_{A'}$, the distribution of decay z is filled both for generated and reconstructed heavy photons (examples in the top halves of Figures 5.15 and 5.16 respectively). The ratio of the two distributions

gives $\epsilon_{reco}(m_{A'}, z)$; for $m_{A'} = 40$ MeV, this is shown in the bottom of Figure 5.16. This is scaled so it equals 1 at $z = z_{target}$, and is fitted with a function of the form $\frac{\epsilon_{reco}(m_{A'}, z)}{\epsilon_{reco}(m_{A'}, z_{target})} \approx \exp(p_3 z^3 + p_2 z^2 + p_1 z + p_0)$. The parameters p_3 , p_2 , p_1 , and p_0 are fit with polynomials in $m_{A'}$: p_0 and p_1 are fit with first-order polynomials, and p_2 and p_3 are fit with third-order polynomials.

Figure 5.16: Top: distribution of decay z for reconstructed heavy photons ($m_{A'} = 40$ MeV, $c\tau = 1$ mm) after all cuts. Bottom: efficiency curve $\frac{\epsilon_{reco}(m_{A'}, z)}{\epsilon_{reco}(m_{A'}, z_{target})}$, with fit.

5.3.3 Fitting the Background Distribution

The background distribution consists of a Gaussian core and a non-Gaussian tail. The width of the Gaussian core is set by multiple scattering and is well understood, but the tails extend much farther than the Gaussian. For $z > z_{cut}$, the background distribution is dominated by the tails: z_{cut} is typically at least $5\sigma_z$.

The background distribution is estimated from Monte Carlo. This is necessary because if the background distribution were fit from the data, and there is actually a heavy photon signal in the data, the fit would be pulled so as to understate the size of the signal. As in Figure 5.17, comparisons of the background distribution in Monte Carlo and data show that the Monte Carlo accurately simulates the processes that control the vertex resolution: the resolution in the Gaussian core and tails found in Monte Carlo are good fits to the data.

Figure 5.17: Comparison of vertex Z distributions between data (black) and Monte Carlo (red).

The 2-dimensional vertex distribution from Monte Carlo is scanned in m , using the same mass cut that is used for the analysis.

The background distribution is fit with a 4-parameter piecewise function defined

from a Gaussian and an exponential. The Gaussian is defined with the usual parameters of mean $\bar{z}$ and standard deviation σ . The parameter b defines the distance from the Gaussian's mean where the distribution transitions to the exponential. The exponential is defined in terms of a tail length l , and its amplitude is fixed by the requirement that the function be continuous. This function is similar to the ‘‘Gauss-Exp’’ function described in [88], except GaussExp fixes the exponential's decay length by requiring that the function be continuously differentiable.

$$B(z; \bar{z}, \sigma, b, l) = \begin{cases} e^{-\frac{(z-\bar{z})^2}{2\sigma^2}} & \text{if } z \leq \bar{z} + b, \\ e^{-\frac{b^2}{2\sigma^2} - (z-\bar{z}-b)/l} & \text{if } z \geq \bar{z} + b \end{cases} \quad (5.5)$$

The values of b and l found at each $m_{A'}$ are fitted to cubic polynomials in $m_{A'}$. When estimating $B(z)$ for data, the values of b and l are fixed to the values found for Monte Carlo, but the Gaussian parameters are allowed to float.

As shown in Figure 5.2 for a representative mass slice, this procedure fits the data well for several orders of magnitude. However, it is not immediately obvious whether the background model completely describes the extreme tail of the distribution.

Characterizing Excess Backgrounds

If the tails of the vertex distribution were completely described by the exponential tail of the background model of Equation 5.5, each mass slice should by construction (on average, and in the absence of a possible heavy photon signal) contain 0.5 events with $z > z_{cut}$. But as seen in Figure 5.18, most mass slices contain more than one event with $z > z_{cut}$, and some contain as many as 4 events; this is a significant excess. This indicates either that there is a heavy photon in the data, or that there is a background not described by the background model.

The events with $z > z_{cut}$ can be characterized by their distributions in z and in mass. Events consistent with the background model should have z close to z_{cut} (the median $z - z_{cut}$ is proportional to l), but be uniformly distributed in mass, with an average of 0.5 events per mass slice. Any other component of the background might have a different distribution in z , but will be broadly distributed in mass. Events from

Figure 5.18: Number of events found with $z > z_{cut}$ in each mass slice, as a function of the center mass of the mass slice. Since mass slices overlap, a single event appears many times in this plot. If the background model fully describes the data, the mean number per mass slice of events with $z > z_{cut}$ should be 0.5.

heavy photon decays should have large z and a tight distribution in mass (limited by the mass resolution).

If the excess events are described by the background model, the events with $z > z_{cut}$ in a mass slice should be distributed according to the exponential distribution. Transforming by the cumulative distribution function (CDF) of the exponential, $z \rightarrow 1 - e^{(z_{cut}-z)/l}$, gives the quantile of each event in the background model; if the background model describes the events, the quantiles should be distributed uniformly. Figure 5.20 shows the quantile histogram for one mass slice; Figure 5.21 shows the quantiles for all mass slices. Some of the events with $z > z_{cut}$ do appear to follow this pattern, but some have large quantiles, meaning that they have z larger than would

Figure 5.19: The set of events that appear with $z > z_{cut}$ in at least one mass slice. This is the same 2-D histogram as Figure 5.1, but only plotting events above $z = z_{cut}$ (the black curve). Some events are close to $z = z_{cut}$, but some events are substantially above the black line ($z \gg z_{cut}$).

be typical from sampling the background model.

If we (somewhat arbitrarily) separate the two sets of events by cutting at a quantile of 0.9, the events with low quantiles contribute 0.52 events per mass slice; this suggests that these events are consistent with the background model and occur with the correct frequency, but the high- z events with quantiles above 0.9 are in excess of the expected background. These are the five events on the upper right of Figure 5.19, with masses between 30 and 50 MeV and $z > 30$ mm.

The five events isolated by the quantile cut are spread broadly in mass. As a group, their RMS in mass is 4.2 MeV, much larger than the estimated mass resolution (2.2 MeV) at the mean mass (38.2 MeV). No mass slice captures more than three of these

Figure 5.20: The background CDF values (quantiles) for events in the mass slice centered at $m_{A'} = 35.6$ MeV that have $z > z_{cut}$ (this is the same mass slice as Figure 5.2). The two events with $z \approx 35$ have quantiles close to 1, suggesting that their values of z are larger than would be expected from the background model.

events; if all five events had the same true mass, the lowest- and highest-mass events would be deviations of more than $2\sigma_m$. In other words, it is unlikely that these events originate from a process with a fixed mass. These events do not match the signature of a heavy photon.

The rate of these events is also not consistent with a heavy photon, or with beam-gas interactions. The expected number of heavy photon decays in this data is much less than 1, as discussed in 5.5.2. Beam-gas interactions, where beam electrons interact with residual gas in the beamline to create e^+e^- pairs downstream of the target, are another potential source of displaced vertices, but cannot explain these events. The beamline vacuum was measured to be in the 10^{-5} torr range, which would (under generous assumptions) contribute less than 0.01 event to this data set.

There is other evidence that these events are a rare detector background and not

Figure 5.21: The background quantiles for events that have $z > z_{cut}$ in each mass slice. This is not a true histogram; each vertical column is a histogram similar to Figure 5.20, but a given event is plotted in every mass slice where it appears. The quantile for a given event changes with the mass slice because z_{cut} depends on the mass slice.

a heavy photon. A sample of trident Monte Carlo (with 1/3 as much integrated luminosity as the data used for this analysis) contains one similar event: see Figure 5.22. This suggests that these events are a background that can be reproduced and understood in Monte Carlo.

Since the current rate of the excess background is greater than the expected rate of heavy photon decays, rejecting them is essential to the full HPS analysis: the number of excess background events will scale with luminosity, and will always overpower any heavy photon signal. The current vertex cuts have mostly been tuned to reject the exponential tail of the vertex z distribution, or to reject specific detector backgrounds.

Studying the excess background in Monte Carlo is likely to identify a specific process that produces these events, and may lead to a specific cut that can reject them.

Figure 5.22: The set of events that appear with $z > z_{cut}$ in at least one mass slice (similar to Figure 5.19), for a trident Monte Carlo sample. One event has $z \gg z_{cut}$.

5.4 Finding Signal Significance

The signal significance is a measure of the inconsistency of the data with the background-only assumption. It is expressed in terms of a p -value, which is the probability under the background-only hypothesis of observing an apparent signal that is at least as significant as what was seen in the data. The convention in particle physics is to convert the p -value to an equivalent significance Z , defined such that p equals the p -value of a Z -sigma upward fluctuation in a Gaussian variable.

For this analysis, the significance at a given $m_{A'}$ is calculated using n , the number of pairs in the mass slice with $z > z_{cut}$, and b , the number of background events expected. p equals the probability $P(n; b) = \sum_{k=n}^{\infty} \frac{b^k}{k!} e^{-b}$ of drawing at least n events from a Poisson distribution with mean b .

According to the expected background distribution and the definition of z_{cut} , b should equal 0.5. However, as shown in Figure 5.18 and explained in 5.3.3, there appears to be an additional background that is above the 0.5 level, and varies smoothly with m . Because this cannot be a heavy photon, we can treat it as part of the background.

Since the background is larger than 0.5, b must be approximated using the data. b for a given $m_{A'}$ hypothesis must be approximated in a way that is not biased by the data in the $m_{A'}$ mass slice. This is done for every mass slice as shown in Figure 5.23 by taking the mass slices that do not overlap with the $m_{A'}$ mass slice, and fitting the trend in n as a quadratic function of $m_{A'}$. Assuming there is only one heavy photon, this procedure yields an estimate of b for a $m_{A'}$ hypothesis that is not perturbed by the possible presence of a heavy photon.

The p -values from this procedure are shown in Figure 5.24. The smallest p -value is 0.0661, but it is not correct to use this as the significance of the data set. These are “local” p -values: the p -value for a mass slice represents the consistency of that slice of the data with the background-only hypothesis. But the objective of this analysis is to find the significance of the full data set, not individual slices: how consistent is the data with the background-only hypothesis? This is expressed by a “global” p -value.

The “global” p -value is found by taking the minimum of all local p -values, and

Figure 5.23: Unbiased fit to estimate the background rate past z_{cut} in the mass slice centered at 35 MeV.

finding the probability of finding a p -value at least as significant as this one. The global p -value is always larger (less significant) than the local p -value because there are multiple local p -values. This is known as the “look-elsewhere effect.”

A simple brute-force method of finding the global p -value is to run a Monte Carlo simulation of the background-only hypothesis and tabulate the probability of finding a given minimum local p -value. The cumulative distribution of minimum local p -values, scaled to 1, gives the mapping to global p -value.

The Monte Carlo model of the background-only hypothesis must reflect the background seen in the data. The distribution of events past z_{cut} is fit to a quadratic function, and data sets are drawn from this distribution: there are 12 events past

Figure 5.24: Local p -value for finding the observed number of events at each $m_{A'}$. The minimum value is 0.0661.

z_{cut} , and so the number of events in each Monte Carlo data set is drawn from a Poisson distribution with a mean of 12. The distribution of minimum local p -values is shown on the left side of Figure 5.25, and the cumulative distribution is on the right.

After this procedure, the global p -value is found to be 0.72 for a local p -value of 0.0661. This is completely consistent with the background-only hypothesis; it means that in fact the local p -value is slightly less significant than the median p -value. Expressed as a equivalent one-sided significance for a Gaussian, the significance of this result is -0.59σ .

Figure 5.25: Left: distribution of the most significant p -values from 10000 runs of toy Monte Carlo. Right: mapping from local to global p -values.

5.5 Setting Limits

An upper limit on the heavy photon production at a given $m_{A'}$ and ϵ^2 is the maximum rate at which heavy photons could be produced, and still be consistent with the data. The confidence level used for this analysis is 90%: in other words, if a heavy photon signal does exist at a given rate, the limit set by this analysis will (incorrectly) exclude that signal rate only 10% of the time. The meaningful target for this analysis is the heavy photon production rate given by Equation 2.5; if the upper limit at a given $m_{A'}$ and ϵ^2 is below that rate, the analysis has (at 90% CL) excluded the possibility of a heavy photon at that $m_{A'}$ and ϵ^2 .

Upper limits do not distinguish between a lack of sensitivity (insufficient data to say anything meaningful about the presence or absence of a signal) and the presence of a signal: the upper limit will be high in either case.

5.5.1 Optimum Interval Method

The method chosen for setting limits is the “optimum interval” method by Yellin [89]. This method was developed for dark matter direct detection experiments, and is intended for experiments where the signal shape is known, but the backgrounds are not fully understood and there is the possibility of an unexpected background. A particular strength of the method is that it minimizes the influence of a background that is concentrated in one part of the data distribution. This analysis uses Yellin’s implementation of the optimum interval method, which is publicly available [90].

The optimum interval method sets a one-sided upper limit (with confidence level C) on the number of signal events μ in a one-dimensional data set, where the shape of the signal distribution is known. For HPS, the data set is the distribution of vertex z locations, after applying the mass and z_{cut} cuts; the signal distribution is the $s(z)$ found in Section 5.3.2 for the $m_{A'}$ and ϵ^2 being tested.

The method works by testing a proposed signal rate μ against the data with a confidence level C . The cumulative distribution function of the signal, $S(z)$, is known. A change of variables is made from the measured variable z to a new variable $x = \mu S(z)$. Under the signal assumption, the expected distribution of the data is uniform

in x with unit density, and has total width μ . An interval (x_1, x_2) , with $x_1, x_2 \in [0, \mu]$, contains a number of expected signal events equal to its width $\Delta x = x_2 - x_1$. If an unexpected background is present and is distributed differently from the signal, the data will not be distributed uniformly in x , and events will be spaced more widely where the background is not present.

The next step is to search for the “optimum interval,” the interval that most strongly rejects the proposed signal rate. This is the interval (x_1, x_2) that contains the smallest number of actual events n relative to its width Δx . Put another way, if the function $C_n(\Delta x, \mu)$ is the probability that all intervals of width Δx contain more than n events, the optimum interval is the (x_1, x_2) that maximizes $C_n(\Delta x, \mu)$.

For the optimum interval, x_1 and x_2 always coincide with 0, μ , or events in the data (otherwise the interval can be widened to increase Δx without changing n). Thus the program only needs to loop over every interval between two events, of width x expected events and containing n actual events. The value of $C_n(\Delta x, \mu)$ for the optimum interval is called C_{Max} , and if it exceeds a threshold $\bar{C}_{Max}(C, \mu)$, μ is rejected with confidence level C . The upper limit on μ is the value for which $C_{Max} = \bar{C}_{Max}(C, \mu)$.

The functions $C_n(x, \mu)$ and $\bar{C}_{Max}(C, \mu)$ pay the statistical penalties for using the data to pick the best interval. Since they are not specific to the signal distribution, they are calculated using Monte Carlo and stored in lookup tables that are distributed with the software.

The optimum interval method can be used with a known background; in this case, the known background density is added to the signal density. Since the known background for HPS falls off rapidly, relatively little is to be gained from this: after the cut in z , the remaining known background is tightly clustered at the edge of the range of z , so the optimum interval method effectively ignores it even without subtraction. Therefore the known background is not used as an input to the optimum interval calculation.

A toy model was used to assess the optimum interval method for use in HPS, and the results are shown in Figure 5.26. The toy signal and background are both exponential distributions, but the background has a short decay length of 2, and the

signal has a long decay length of 20; in units of mm, these are typical values for HPS. 10000 background events are generated; there is no signal. A nuisance background of 100 events, with a decay length of 10, is present in the plot on the right. The different limits (as a number of signal events) are plotted as a function of z_{cut}

The optimum interval method was compared against cut-and-count limits using the Feldman-Cousins method (with no background subtraction) [91]. As expected, the upper limits from the optimum interval method are always as strong as, or stronger than, the upper limits from the cut-and-count method given the same information.

Optimum interval limits are insensitive to background events at the edge of the search range. This can be seen around $z = 22$ and $z = 26$ in the right-hand plot, where there are events from the nuisance background: the cut-and-count limits jump up discontinuously when z_{cut} moves past the event, but the optimum interval limits vary smoothly. It is still clear that it makes sense to avoid as much background as possible, so setting z_{cut} for expected 0.5 background events is still appropriate.

Figure 5.26: Comparison of the optimum interval method with cut-and-count using Feldman-Cousins limits. The y-axis is the limit on the total signal rate as a fraction of the background rate. The background (10000 events) has decay length 2, the signal has decay length 20, and the unexpected background (100 events) has decay length 5. Left plot is with only the expected background (no signal); right plot is with the unexpected background added (still no signal). The z_{cut} for 0.5 expected background events is 19.8.

5.5.2 Results

The optimum interval method gives limits at 90% confidence level on μ , the number of signal events past z_{cut} and after acceptance and efficiency effects. Dividing the limit on μ by μ_{exp} , the expected number of signal events for a heavy photon, gives a limit on the dimensionless ratio of the true production rate to the expected production rate for a heavy photon. A limit on μ/μ_{exp} of less than 1 means the limit on μ is less than μ_{exp} ; in other words, the heavy photon is excluded (at 90% CL) at that $m_{A'}$ and ϵ^2 .

The value μ_{exp} is itself of some interest: it relates directly to the strongest possible limit on μ/μ_{exp} . The strongest limit on μ is obtained when no events are observed in the mass slice with $z > z_{cut}$, in which case the 90% CL limit on μ is 2.303 (this is true not only for the optimum interval method, but in general): if $\mu = 2.303$, there is exactly a 10% probability of seeing no events. Therefore, exclusion of the heavy photon (a limit on μ/μ_{exp} of 1 or less) is only possible if $\mu_{exp} > 2.303$.

The limits from this analysis are shown in Figure 5.27, which shows that this analysis, on this data, is a factor of 115 from excluding any portion of the parameter space. As shown in Figure 5.28, $\mu_{exp} = 0.032$ at most, which is too small for exclusion even in the absence of excess background events.

Reach Projections for the 2015 Run

Several factors will improve this reach in future iterations of this analysis. First, the number of events with $z > z_{cut}$ was larger than predicted by the background fit. Better background rejection cuts may be able to reduce this background, and bring the reach closer to the bound set by μ_{exp} . The reach with zero events with $z > z_{cut}$ is shown in Figure 5.29, where the strongest predicted limit is down to 86.2. If, after refining cuts, the background shape comes to match the exponential form used in this analysis, the reach will be very close to this zero-events reach.

Second, this analysis only used the unblinded fraction of the 2015 data. The full set increases integrated luminosity by a factor of 9.77 (from 199 to 1166 nb^{-1}), and therefore the number of detectable heavy photons. Since z_{cut} must be increased to keep the level of background at 0.5 events per mass slice, μ_{exp} does not scale linearly

Figure 5.27: 90% CL upper limit on μ/μ_{exp} , the ratio of the true production rate to the expected production rate for a heavy photon. A value of 1 would mean exclusion; the lowest contour on this plot is 200, and the lowest point (at $m_{A'} = 49$ MeV, $\epsilon^2 = 2.1 \times 10^{-9}$) is 115. The vertical ridges in this plot correspond to the locations of events in mass space.

with the amount of data. As shown in Figure 5.30, this improves the exclusion to 14.7.

Finally, this analysis only used the events where both tracks made hits in layer 1 of the SVT. As explained in Section 5.2.1 Figure 5.5 shows the difference in efficiency between the current analysis and the full acceptance. Doing this will require tuning cuts separately for the sets of events that miss layer 1, and combining the data sets (the optimum interval method can be used for this [92]). As shown in Figure 5.31, this improves the exclusion to 6.68. The mass range of the search is also substantially improved, because the lower values of $m_{A'}$ are more strongly affected by the layer 1

Figure 5.28: μ_{exp} , the number of detectable heavy photon events expected with $z > z_{cut}$. The highest contour is at 0.03 events, and the highest point is 0.032.

requirement.

Some data was taken in 2015 with the SVT opened wider than its nominal position (silicon at 1.5 mm from the beam, instead of 0.5 mm). The integrated luminosity is similar to the data that was used in this analysis, but the acceptance is worse. At best, this can only double μ_{exp} , so even using this data there will be no excluded region.

Reach projections for the 2016 data and the planned 2018 run are in progress.

Reach Comparison with 2014 Proposal

The reach estimates in the 2014 proposal (see Figure 1.1) are substantially different from the projected reach shown above. The proposal estimates were made assuming

Figure 5.29: Predicted 90% CL upper limit on μ/μ_{exp} , assuming no events with $z > z_{cut}$. A value of 1 would mean exclusion; the lowest contour on this plot is 90, and the lowest point is 86.2.

1 week of beam time, and used some assumptions that are now understood to be inaccurate. In addition, the trident rates in data are different from the rates in Monte Carlo trident samples. Even accounting for these differences, there are significant differences between the μ_{exp} estimated using this analysis, and the μ_{exp} that was used for reach estimates in the proposal.

The reach estimates for the proposal were made by calculating μ_{exp} assuming the same procedure presented here: apply cuts to reduce the backgrounds at large z , make a resolution-limited mass slice and apply a z_{cut} that reduces the background to 0.5 expected events.

These are assumptions that were made for the proposal, and how each differs from

Figure 5.30: Predicted 90% CL upper limit on μ/μ_{exp} , projected for the full 2015 data set, and assuming no events with $z > z_{cut}$. A value of 1 would mean exclusion; the lowest contour on this plot is 20, and the lowest point is 14.7.

our current understanding of the experiment:

- The integrated luminosity was expected to be 1 week of 50 nA beam, on a target with the nominal thickness of 0.125% of a radiation length; the actual target was somewhat thinner, at 0.116% X_0 .
- The beam energy was expected to be 1.1 GeV; in reality it was 1.056 GeV. The beam energy affects the vertex resolution, mass scale and production rates, but the difference in beam energy is small.
- Efficiency for heavy photons was assumed to be uniform to $z = 100$ ($\epsilon_{reco}(m_{A'}, z) = \epsilon_{reco}(m_{A'}, z_{target})$). This is of course not true, as explained in Section 5.3.2. The

Figure 5.31: Predicted 90% CL upper limit on μ/μ_{exp} , projected for the full 2015 data set, using the full HPS acceptance, and assuming no events with $z > z_{cut}$. A value of 1 would mean exclusion; the lowest contour on this plot is 7, and the lowest point is 6.68.

impact depends on mass (just like the efficiency) but is roughly a factor of 2 to 3.

- Rates for heavy photons and trident backgrounds were estimated using a MadGraph/MadEvent generator and a rough mock-up of the geometric acceptance (not a full Monte Carlo simulation of the detector, trigger, and reconstruction). The rate estimate from the mocked-up acceptance is high by roughly a factor of 2, compared to the rate of fully reconstructed e^+e^- pairs in trident Monte Carlo samples. This suggests the mocked-up geometric acceptance may be too generous.

- Efficiencies were estimated at 85% reconstruction efficiency per track, and 50% cut efficiency per pair. In short, efficiency is $0.85^2 \times 0.5 = 0.36$ if both particles are within the geometric acceptance, 0 otherwise. The reconstruction and cut efficiency estimates are in rough agreement with what is seen in Monte Carlo samples, though as shown in Section 4.5, the efficiency in data appears to be 65% that in Monte Carlo.
- The width of a mass slice was set at $m_{A'} \pm 1.25\sigma_m$, instead of $m_{A'} \pm 1.4\sigma_m$. The effect is a slight reduction in the mass cut efficiency for a heavy photon: $\epsilon_{cut} = 0.789$ instead of 0.838.

These assumptions are summarized in Table 5.2, and the μ_{exp} used for the proposal is shown in Figure 5.32.

Table 5.2: Assumptions used for the reach estimate in the 2014 proposal.

Beam energy	1.1 GeV
Integrated luminosity	50 nA, 1 week, target thickness 0.125% X_0
Efficiency dependence on z	Uniform to $z = 100$ ($\epsilon_{reco}(m_{A'}, z) = \epsilon_{reco}(m_{A'}, z_{target})$)
Efficiency at z_{target}	$0.85^2 \times 0.5 = 0.36$ if both particles are within the geometric acceptance, 0 otherwise

The reach projections in the current analysis can be rerun using the trident Monte Carlo sample and some of the proposal assumptions to produce a “proposal-like” reach estimate. A comparison between the proposal and proposal-like estimates is a first step towards understanding any problems with either procedure. The assumptions for the proposal-like estimate are shown in Table 5.3, and the resulting μ_{exp} is shown in Figure 5.33.

The proposal (Figure 5.32) and proposal-like (Figure 5.33) estimates show large differences in μ_{exp} . The proposal μ_{exp} has its peak at a lower $m_{A'}$ (22 MeV, as opposed to 28 MeV), and the peak is higher (17.3 events, as opposed to 7 events). Some differences between Figures 5.32 and 5.33 are to be expected, since the beam

Figure 5.32: The predicted values of μ_{exp} , the number of detectable heavy photon events expected with $z > z_{cut}$, that were used for the reach estimate in the 2014 proposal. The $m_{A'}$ and ϵ^2 ranges, and the binning, are different in this plot compared to other plots in this dissertation. The highest contour is at 10 events, and the highest point (at $m_{A'} = 22$ MeV, $\epsilon^2 = 5 \times 10^{-9}$) is 17.3 events.

energy and target thickness are different. However, these differences are of order 10% and should not produce large changes, or shift the peak $m_{A'}$ by more than the change in the beam energy. The event selection cuts (and thus the z_{cut}) have also changed from the proposal, but not enough to explain the changes seen.

In addition, the low-mass reach in the proposal does not seem plausible: the proposal shows reach at $m_{A'}$ below the lower limit of the HPS acceptance. The proposal shows exclusion at masses as low as $m_{A'} = 14$ MeV, but Monte Carlo samples of low-mass heavy photons show that the efficiency at $z = z_{target}$ falls rapidly in this mass range, from 1% at $m_{A'} = 20$ to 0.06% at $m_{A'} = 16$ MeV, and 0.01%

Table 5.3: Assumptions used for the “proposal-like” reach estimate.

Beam energy	1.056 GeV
Integrated luminosity	50 nA, 1 week, target thickness 0.116% X_0
Efficiency dependence on z	Uniform to $z = 100$ ($\epsilon_{reco}(m_{A'}, z) = \epsilon_{reco}(m_{A'}, z_{target})$)
Efficiency at z_{target}	From Monte Carlo

at $m_{A'} = 15$ MeV. This is expected since the decay products from heavy photons in this mass range are near the minimum opening angle detectable by the SVT; the rate of pairs with reconstructed masses below 15 MeV is still appreciable, but these are mostly pairs with higher masses whose masses were shifted downwards by the mass resolution.

This suggests that the simplified model of the HPS geometric acceptance used for the proposal may have omitted important features of the detector, or that the method used to estimate the detector efficiency may have been flawed. Further work is needed to identify the problem, and confirm that there is no mistake in the current analysis.

Figure 5.33: Predicted μ_{exp} , the number of detectable heavy photon events expected with $z > z_{cut}$, with the assumptions listed in Table 5.3. The highest contour is at 7 events, and the highest point is 7.51 events.

Chapter 6

Conclusion

This dissertation covers the initial stages of the HPS experiment. The detector has been built, and brought into operation. We are approaching a full understanding of the event reconstruction and the data. The analyses are approaching their final state.

The displaced-vertex search is fully developed. The event selection has been tuned, and the backgrounds after event selection have been characterized. The steps of the analysis have been tested, and perform as expected. Data and Monte Carlo are each used where appropriate, and the analysis is data-driven to the extent possible.

The remaining steps to a complete result are clear. The excess background at large z (Section 5.3.3) must be understood and rejected. Pairs with missing layer 1 hits are an additional data set, which require separate tuning of cuts and modeling of signals and backgrounds, and the data sets must then be combined. Similarly, the data taken with the tracker at 1.5 mm must be combined with the 0.5 mm data used in this dissertation.

The signal significance analysis detailed in Section 5.4 can be refined considerably. A profile likelihood analysis should make better use of the mass and z information, will directly compute the likelihood of an observed excess being a background fluctuation or a heavy photon, and will identify preferred parameters ($m_{A'}$, ϵ^2 , and production rate) for an excess.

This dissertation is written at the time that HPS is becoming a mature experiment. We now understand how the experiment works, and what we can do to improve future

runs. Based on these results, HPS is planning to reconfigure the SVT to get better acceptance at large z , and an upgrade is being considered to add a “layer 0” to improve the vertex resolution. Reach estimates for future runs will also be much better informed.

Bibliography

- [1] HPS Collaboration. *Heavy Photon Search Experiment at Jefferson Laboratory: Proposal for 2014-2015 Run*. May 10, 2013. URL: https://confluence.slac.stanford.edu/download/attachments/86676777/hps_2014.pdf?version=1&modificationDate=1368227154000 (visited on 09/17/2016).
- [2] R. Essig et al. “Dark Sectors and New, Light, Weakly-Coupled Particles”. In: (Oct. 31, 2013). arXiv: 1311.0029 [astro-ph, physics:hep-ex, physics:hep-ph]. URL: <http://arxiv.org/abs/1311.0029> (visited on 07/27/2016).
- [3] Bob Holdom. “Two U(1)’s and Charge Shifts”. In: *Physics Letters B* 166.2 (Jan. 9, 1986), pp. 196–198. ISSN: 0370-2693. DOI: 10.1016/0370-2693(86)91377-8. URL: <http://www.sciencedirect.com/science/article/pii/0370269386913778> (visited on 07/27/2016).
- [4] Nima Arkani-Hamed and Neal Weiner. “LHC Signals for a SuperUnified Theory of Dark Matter”. In: *Journal of High Energy Physics* 2008.12 (Dec. 29, 2008), pp. 104–104. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2008/12/104. arXiv: 0810.0714. URL: <http://arxiv.org/abs/0810.0714> (visited on 07/28/2016).
- [5] Mark Goodsell et al. “Naturally Light Hidden Photons in LARGE Volume String Compactifications”. In: *Journal of High Energy Physics* 2009.11 (Nov. 6, 2009), pp. 027–027. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2009/11/027. arXiv: 0909.0515. URL: <http://arxiv.org/abs/0909.0515> (visited on 08/24/2016).
- [6] Michele Cicoli et al. “Testing String Vacua in the Lab: From a Hidden CMB to Dark Forces in Flux Compactifications”. In: *Journal of High Energy Physics*

- 2011.7 (July 2011). ISSN: 1029-8479. DOI: 10.1007/JHEP07(2011)114. arXiv: 1103.3705. URL: <http://arxiv.org/abs/1103.3705> (visited on 08/24/2016).
- [7] Matthew Baumgart et al. “Non-Abelian Dark Sectors and Their Collider Signatures”. In: *Journal of High Energy Physics* 2009 (04 Apr. 3, 2009), pp. 014–014. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2009/04/014. arXiv: 0901.0283. URL: <http://arxiv.org/abs/0901.0283> (visited on 08/24/2016).
- [8] David E. Morrissey, David Poland, and Kathryn M. Zurek. “Abelian Hidden Sectors at a GeV”. In: *Journal of High Energy Physics* 2009 (07 July 15, 2009), pp. 050–050. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2009/07/050. arXiv: 0904.2567. URL: <http://arxiv.org/abs/0904.2567> (visited on 08/24/2016).
- [9] Clifford Cheung et al. “Kinetic Mixing as the Origin of Light Dark Scales”. In: *Physical Review D* 80.3 (Aug. 12, 2009). ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.80.035008. arXiv: 0902.3246. URL: <http://arxiv.org/abs/0902.3246> (visited on 08/24/2016).
- [10] The Fermi LAT Collaboration et al. “Measurement of Separate Cosmic-Ray Electron and Positron Spectra with the Fermi Large Area Telescope”. In: *Physical Review Letters* 108.1 (Jan. 5, 2012). ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.108.011103. arXiv: 1109.0521. URL: <http://arxiv.org/abs/1109.0521> (visited on 08/24/2016).
- [11] AMS Collaboration et al. “First Result from the Alpha Magnetic Spectrometer on the International Space Station: Precision Measurement of the Positron Fraction in Primary Cosmic Rays of 0.5-350 GeV”. In: *Physical Review Letters* 110.14 (Apr. 3, 2013), p. 141102. DOI: 10.1103/PhysRevLett.110.141102. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.110.141102> (visited on 08/24/2016).
- [12] Ilias Cholis et al. “High Energy Positrons From Annihilating Dark Matter”. In: *Physical Review D* 80.12 (Dec. 10, 2009). ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.80.123511. arXiv: 0809.1683. URL: <http://arxiv.org/abs/0809.1683> (visited on 08/24/2016).

- [13] Nima Arkani-Hamed et al. “A Theory of Dark Matter”. In: *Physical Review D* 79.1 (Jan. 27, 2009). ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.79.015014. arXiv: 0810.0713. URL: <http://arxiv.org/abs/0810.0713> (visited on 07/29/2016).
- [14] Douglas P. Finkbeiner et al. “Consistent Scenarios for Cosmic-Ray Excesses from Sommerfeld-Enhanced Dark Matter Annihilation”. In: *Journal of Cosmology and Astroparticle Physics* 2011 (05 May 3, 2011), pp. 002–002. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2011/05/002. arXiv: 1011.3082. URL: <http://arxiv.org/abs/1011.3082> (visited on 08/24/2016).
- [15] Tracy R. Slatyer, Natalia Toro, and Neal Weiner. “Sommerfeld-Enhanced Annihilation in Dark Matter Substructure: Consequences for Constraints on Cosmic-Ray Excesses”. In: *Physical Review D* 86.8 (Oct. 22, 2012). ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.86.083534. arXiv: 1107.3546. URL: <http://arxiv.org/abs/1107.3546> (visited on 08/24/2016).
- [16] AMS Collaboration et al. “High Statistics Measurement of the Positron Fraction in Primary Cosmic Rays of 0.5-500 GeV with the Alpha Magnetic Spectrometer on the International Space Station”. In: *Physical Review Letters* 113.12 (Sept. 18, 2014), p. 121101. DOI: 10.1103/PhysRevLett.113.121101. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.113.121101> (visited on 08/24/2016).
- [17] Ilias Cholis and Dan Hooper. “Dark Matter and Pulsar Origins of the Rising Cosmic Ray Positron Fraction in Light of New Data From AMS”. In: *Physical Review D* 88.2 (July 29, 2013). ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.88.023013. arXiv: 1304.1840. URL: <http://arxiv.org/abs/1304.1840> (visited on 08/25/2016).
- [18] Mathew S. Madhavacheril, Neelima Sehgal, and Tracy R. Slatyer. “Current Dark Matter Annihilation Constraints from CMB and Low-Redshift Data”. In: *Physical Review D* 89.10 (May 7, 2014). ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.89.103508. arXiv: 1310.3815. URL: <http://arxiv.org/abs/1310.3815> (visited on 08/24/2016).

- [19] Dan Hooper and Lisa Goodenough. “Dark Matter Annihilation in The Galactic Center As Seen by the Fermi Gamma Ray Space Telescope”. In: *Physics Letters B* 697.5 (Mar. 2011), pp. 412–428. ISSN: 03702693. DOI: 10.1016/j.physletb.2011.02.029. arXiv: 1010.2752. URL: <http://arxiv.org/abs/1010.2752> (visited on 08/25/2016).
- [20] Esra Bulbul et al. “Detection of An Unidentified Emission Line in the Stacked X-Ray Spectrum of Galaxy Clusters”. In: *The Astrophysical Journal* 789.1 (July 1, 2014), p. 13. ISSN: 0004-637X, 1538-4357. DOI: 10.1088/0004-637X/789/1/13. arXiv: 1402.2301. URL: <http://arxiv.org/abs/1402.2301> (visited on 08/25/2016).
- [21] Alexey Boyarsky et al. “An Unidentified Line in X-Ray Spectra of the Andromeda Galaxy and Perseus Galaxy Cluster”. In: *Physical Review Letters* 113.25 (Dec. 15, 2014). ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.113.251301. arXiv: 1402.4119. URL: <http://arxiv.org/abs/1402.4119> (visited on 08/25/2016).
- [22] Douglas P. Finkbeiner and Neal Weiner. “An X-Ray Line from eXciting Dark Matter”. In: (Feb. 26, 2014). arXiv: 1402.6671 [astro-ph, physics:hep-ph]. URL: <http://arxiv.org/abs/1402.6671> (visited on 08/25/2016).
- [23] David H. Weinberg et al. “Cold Dark Matter: Controversies on Small Scales”. In: (June 4, 2013). arXiv: 1306.0913 [astro-ph]. URL: <http://arxiv.org/abs/1306.0913> (visited on 08/25/2016).
- [24] David N. Spergel and Paul J. Steinhardt. “Observational Evidence for Self-Interacting Cold Dark Matter”. In: *Physical Review Letters* 84.17 (Apr. 24, 2000), pp. 3760–3763. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.84.3760. arXiv: astro-ph/9909386. URL: <http://arxiv.org/abs/astro-ph/9909386> (visited on 08/26/2016).
- [25] Sean Tulin, Hai-Bo Yu, and Kathryn M. Zurek. “Beyond Collisionless Dark Matter: Particle Physics Dynamics for Dark Matter Halo Structure”. In: *Physical Review D* 87.11 (June 7, 2013). ISSN: 1550-7998, 1550-2368. DOI: 10.1103/

- PhysRevD.87.115007. arXiv: 1302.3898. URL: <http://arxiv.org/abs/1302.3898> (visited on 08/25/2016).
- [26] Michael Boylan-Kolchin, James S. Bullock, and Manoj Kaplinghat. “Too Big to Fail? The Puzzling Darkness of Massive Milky Way Subhaloes”. In: *Monthly Notices of the Royal Astronomical Society: Letters* 415.1 (July 21, 2011), pp. L40–L44. ISSN: 17453925. DOI: 10.1111/j.1745-3933.2011.01074.x. arXiv: 1103.0007. URL: <http://arxiv.org/abs/1103.0007> (visited on 08/25/2016).
- [27] Mark Vogelsberger, Jesus Zavala, and Abraham Loeb. “Subhaloes in Self-Interacting Galactic Dark Matter Haloes”. In: *Monthly Notices of the Royal Astronomical Society* 423.4 (July 11, 2012), pp. 3740–3752. ISSN: 00358711. DOI: 10.1111/j.1365-2966.2012.21182.x. arXiv: 1201.5892. URL: <http://arxiv.org/abs/1201.5892> (visited on 08/25/2016).
- [28] Rachel Kuzio de Naray and Kristine Spekkens. “Do Baryons Alter the Halos of Low Surface Brightness Galaxies?” In: *The Astrophysical Journal* 741.2 (Nov. 10, 2011), p. L29. ISSN: 2041-8205, 2041-8213. DOI: 10.1088/2041-8205/741/2/L29. arXiv: 1109.1288. URL: <http://arxiv.org/abs/1109.1288> (visited on 08/26/2016).
- [29] Kyle A. Oman et al. “The Unexpected Diversity of Dwarf Galaxy Rotation Curves”. In: *Monthly Notices of the Royal Astronomical Society* 452.4 (Oct. 1, 2015), pp. 3650–3665. ISSN: 0035-8711, 1365-2966. DOI: 10.1093/mnras/stv1504. arXiv: 1504.01437. URL: <http://arxiv.org/abs/1504.01437> (visited on 08/25/2016).
- [30] F. Governato et al. “Cuspy No More: How Outflows Affect the Central Dark Matter and Baryon Distribution in Lambda CDM Galaxies”. In: *Monthly Notices of the Royal Astronomical Society* 422.2 (May 11, 2012), pp. 1231–1240. ISSN: 00358711. DOI: 10.1111/j.1365-2966.2012.20696.x. arXiv: 1202.0554. URL: <http://arxiv.org/abs/1202.0554> (visited on 08/25/2016).
- [31] Muon, Collaboration, and G. W. Bennett. “Final Report of the Muon E821 Anomalous Magnetic Moment Measurement at BNL”. In: *Physical Review D* 73.7 (Apr. 7, 2006). ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.

- 73.072003. arXiv: hep-ex/0602035. URL: <http://arxiv.org/abs/hep-ex/0602035> (visited on 08/26/2016).
- [32] Thomas Blum et al. “The Muon (G-2) Theory Value: Present and Future”. In: (Nov. 9, 2013). arXiv: 1311.2198 [hep-ex, physics:hep-lat, physics:hep-ph]. URL: <http://arxiv.org/abs/1311.2198> (visited on 08/26/2016).
- [33] Maxim Pospelov. “Secluded U(1) below the Weak Scale”. In: *Physical Review D* 80.9 (Nov. 2, 2009). ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.80.095002. arXiv: 0811.1030. URL: <http://arxiv.org/abs/0811.1030> (visited on 08/10/2016).
- [34] Motoi Endo, Koichi Hamaguchi, and Go Mishima. “Constraints on Hidden Photon Models from Electron G-2 and Hydrogen Spectroscopy”. In: *Physical Review D* 86.9 (Nov. 27, 2012). ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.86.095029. arXiv: 1209.2558. URL: <http://arxiv.org/abs/1209.2558> (visited on 08/26/2016).
- [35] Jia Liu, Neal Weiner, and Wei Xue. “Signals of a Light Dark Force in the Galactic Center”. In: *Journal of High Energy Physics* 2015.8 (Aug. 2015). ISSN: 1029-8479. DOI: 10.1007/JHEP08(2015)050. arXiv: 1412.1485. URL: <http://arxiv.org/abs/1412.1485> (visited on 08/26/2016).
- [36] J. D. Bjorken et al. “Search for Neutral Metastable Penetrating Particles Produced in the SLAC Beam Dump”. In: *Physical Review D* 38.11 (Dec. 1, 1988), pp. 3375–3386. DOI: 10.1103/PhysRevD.38.3375. URL: <http://link.aps.org/doi/10.1103/PhysRevD.38.3375> (visited on 07/29/2016).
- [37] E. M. Riordan et al. “Search for Short-Lived Axions in an Electron-Beam-Dump Experiment”. In: *Physical Review Letters* 59.7 (Aug. 17, 1987), pp. 755–758. DOI: 10.1103/PhysRevLett.59.755. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.59.755> (visited on 07/29/2016).
- [38] A. Bross et al. “Search for Short-Lived Particles Produced in an Electron Beam Dump”. In: *Physical Review Letters* 67.21 (Nov. 18, 1991), pp. 2942–2945. DOI:

- 10.1103/PhysRevLett.67.2942. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.67.2942> (visited on 07/29/2016).
- [39] M. Davier and H. Nguyen Ngoc. “An Unambiguous Search for a Light Higgs Boson”. In: *Physics Letters B* 229.1 (Oct. 5, 1989), pp. 150–155. ISSN: 0370-2693. DOI: 10.1016/0370-2693(89)90174-3. URL: <http://www.sciencedirect.com/science/article/pii/0370269389901743> (visited on 08/27/2016).
- [40] J. Blmlein and J. Brunner. “New Exclusion Limits on Dark Gauge Forces from Proton Bremsstrahlung in Beam-Dump Data”. In: *Physics Letters B* 731 (Apr. 2014), pp. 320–326. ISSN: 03702693. DOI: 10.1016/j.physletb.2014.02.029. arXiv: 1311.3870. URL: <http://arxiv.org/abs/1311.3870> (visited on 08/27/2016).
- [41] SHiP Collaboration et al. “A Facility to Search for Hidden Particles (SHiP) at the CERN SPS”. In: (Apr. 20, 2015). arXiv: 1504.04956 [hep-ex, physics:physics]. URL: <http://arxiv.org/abs/1504.04956> (visited on 08/27/2016).
- [42] S. Gardner, R. J. Holt, and A. S. Tadepalli. “New Prospects in Fixed Target Searches for Dark Forces with the SeaQuest Experiment at Fermilab”. In: *Physical Review D* 93.11 (June 10, 2016). ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.93.115015. arXiv: 1509.00050. URL: <http://arxiv.org/abs/1509.00050> (visited on 08/27/2016).
- [43] S. N. Gninenko. “Search for MeV Dark Photons in a Light-Shining-through-Walls Experiment at CERN”. In: *Physical Review D* 89.7 (Apr. 8, 2014). ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.89.075008. arXiv: 1308.6521. URL: <http://arxiv.org/abs/1308.6521> (visited on 08/27/2016).
- [44] H. Merkel et al. “Search for Light Massive Gauge Bosons as an Explanation of the $(g-2)_{\mu}$ Anomaly at MAMI”. In: *Physical Review Letters* 112.22 (June 4, 2014). ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.112.221802. arXiv: 1404.5502. URL: <http://arxiv.org/abs/1404.5502> (visited on 08/28/2016).

- [45] Rouven Essig et al. “An Electron Fixed Target Experiment to Search for a New Vector Boson A' Decaying to $E+e^-$ ”. In: *Journal of High Energy Physics* 2011.2 (Feb. 2011), pp. 1–35. ISSN: 1029-8479. DOI: 10.1007/JHEP02(2011)009. arXiv: 1001.2557. URL: <http://arxiv.org/abs/1001.2557> (visited on 08/28/2016).
- [46] S. Abrahamyan et al. “Search for a New Gauge Boson in the A' Experiment (APEX)”. In: *Physical Review Letters* 107.19 (Nov. 4, 2011). ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.107.191804. arXiv: 1108.2750. URL: <http://arxiv.org/abs/1108.2750> (visited on 08/28/2016).
- [47] J. Balewski et al. “The DarkLight Experiment: A Precision Search for New Physics at Low Energies”. In: (Dec. 15, 2014). arXiv: 1412.4717 [hep-ex, physics:physics]. URL: <http://arxiv.org/abs/1412.4717> (visited on 08/28/2016).
- [48] Achim Denig. “Recent Results from the Mainz Microtron MAMI and an Outlook for the Future”. In: *AIP Conference Proceedings*. XVITH INTERNATIONAL CONFERENCE ON HADRON SPECTROSCOPY: Hadron2015. Vol. 1735. AIP Publishing, May 25, 2016, p. 020006. DOI: 10.1063/1.4949374. URL: <http://scitation.aip.org/content/aip/proceeding/aipcp/10.1063/1.4949374> (visited on 08/29/2016).
- [49] Mauro Raggi, Venelin Kozhuharov, and P. Valente. “The PADME Experiment at LNF”. In: (Jan. 8, 2015). arXiv: 1501.01867 [hep-ex, physics:hep-ph]. URL: <http://arxiv.org/abs/1501.01867> (visited on 08/28/2016).
- [50] B. Wojtsekhowski, D. Nikolenko, and I. Rachek. “Searching for a New Force at VEPP-3”. In: (July 20, 2012). arXiv: 1207.5089 [hep-ex, physics:nucl-ex]. URL: <http://arxiv.org/abs/1207.5089> (visited on 08/28/2016).
- [51] HADES Collaboration et al. “Searching a Dark Photon with HADES”. In: (Nov. 1, 2013). arXiv: 1311.0216 [hep-ex, physics:hep-ph]. URL: <http://arxiv.org/abs/1311.0216> (visited on 08/29/2016).

- [52] BABAR Collaboration et al. “Search for Dimuon Decays of a Light Scalar Boson in Radiative Transitions $\Upsilon \rightarrow \gamma$ ”. In: *Physical Review Letters* 103.8 (Aug. 18, 2009), p. 081803. DOI: 10.1103/PhysRevLett.103.081803. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.103.081803> (visited on 08/29/2016).
- [53] Abner Soffer. “Constraints on Dark Forces from the B Factories and Low-Energy Experiments”. In: (Sept. 18, 2014). arXiv: 1409.5263 [hep-ex, physics:hep-ph]. URL: <http://arxiv.org/abs/1409.5263> (visited on 08/29/2016).
- [54] KLOE-2 Collaboration. “Search for a Vector Gauge Boson in Phi Meson Decays with the KLOE Detector”. In: *Physics Letters B* 706 (4-5 Jan. 2012), pp. 251–255. ISSN: 03702693. DOI: 10.1016/j.physletb.2011.11.033. arXiv: 1110.0411. URL: <http://arxiv.org/abs/1110.0411> (visited on 08/28/2016).
- [55] KLOE-2 Collaboration et al. “Limit on the Production of a Light Vector Gauge Boson in Phi Meson Decays with the KLOE Detector”. In: *Physics Letters B* 720 (1-3 Mar. 2013), pp. 111–115. ISSN: 03702693. DOI: 10.1016/j.physletb.2013.01.067. arXiv: 1210.3927. URL: <http://arxiv.org/abs/1210.3927> (visited on 08/28/2016).
- [56] D. Babusci et al. “Search for Light Vector Boson Production in $e^+e^- \rightarrow \mu^+\mu^-\gamma$ Interactions with the KLOE Experiment”. In: *Physics Letters B* 736 (Sept. 2014), pp. 459–464. ISSN: 03702693. DOI: 10.1016/j.physletb.2014.08.005. arXiv: 1404.7772. URL: <http://arxiv.org/abs/1404.7772> (visited on 08/28/2016).
- [57] A. Anastasi et al. “Limit on the Production of a Low-Mass Vector Boson in $e^+e^- \rightarrow U\gamma$, $U \rightarrow e^+e^-$ with the KLOE Experiment”. In: *Physics Letters B* 750 (Nov. 2015), pp. 633–637. ISSN: 03702693. DOI: 10.1016/j.physletb.2015.10.003. arXiv: 1509.00740. URL: <http://arxiv.org/abs/1509.00740> (visited on 08/28/2016).

- [58] KLOE-2 Collaboration et al. “Limit on the Production of a New Vector Boson in $\mathrm{e^+e^-} \rightarrow \mathrm{U}\gamma$, $\mathrm{U} \rightarrow \pi^+\pi^-$ with the KLOE Experiment”. In: *Physics Letters B* 757 (June 2016), pp. 356–361. ISSN: 03702693. DOI: 10.1016/j.physletb.2016.04.019. arXiv: 1603.06086. URL: <http://arxiv.org/abs/1603.06086> (visited on 08/28/2016).
- [59] Gianluca Inguglia. “Belle II Studies of Missing Energy Decays and Searches for Dark Photon Production”. In: (July 7, 2016). arXiv: 1607.02089 [hep-ex, physics:hep-ph]. URL: <http://arxiv.org/abs/1607.02089> (visited on 08/29/2016).
- [60] WASA-at-COSY Collaboration et al. “Search for a Dark Photon in the $\pi^0 \rightarrow \mathrm{e^+e^-}\gamma$ Decay”. In: *Physics Letters B* 726 (1-3 Oct. 2013), pp. 187–193. ISSN: 03702693. DOI: 10.1016/j.physletb.2013.08.055. arXiv: 1304.0671. URL: <http://arxiv.org/abs/1304.0671> (visited on 08/28/2016).
- [61] Philip Ilten et al. “Dark Photons from Charm Mesons at LHCb”. In: *Physical Review D* 92.11 (Dec. 21, 2015). ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.92.115017. arXiv: 1509.06765. URL: <http://arxiv.org/abs/1509.06765> (visited on 08/29/2016).
- [62] Philip Ilten et al. “Inclusive Dark Photon Search at LHCb”. In: *Physical Review Letters* 116.25 (June 23, 2016). ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.116.251803. arXiv: 1603.08926. URL: <http://arxiv.org/abs/1603.08926> (visited on 08/29/2016).
- [63] Omar Moreno. “Search for a Heavy Photon in the 2015 Engineering Run Data of the Heavy Photon Search Experiment”. Jan. 1, 2016. URL: <http://escholarship.org/uc/item/44n7c8qc> (visited on 08/31/2016).
- [64] James D. Bjorken et al. “New Fixed-Target Experiments to Search for Dark Gauge Forces”. In: *Physical Review D* 80.7 (Oct. 28, 2009). ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.80.075018. arXiv: 0906.0580. URL: <http://arxiv.org/abs/0906.0580> (visited on 07/26/2016).

- [65] Yung-Su Tsai. “Pair Production and Bremsstrahlung of Charged Leptons”. In: *Reviews of Modern Physics* 46.4 (Oct. 1, 1974), pp. 815–851. DOI: 10.1103/RevModPhys.46.815. URL: <http://link.aps.org/doi/10.1103/RevModPhys.46.815> (visited on 08/30/2016).
- [66] Yung Su Tsai. “Axion Bremsstrahlung by an Electron Beam”. In: *Physical Review D* 34.5 (Sept. 1, 1986), pp. 1326–1331. DOI: 10.1103/PhysRevD.34.1326. URL: <http://link.aps.org/doi/10.1103/PhysRevD.34.1326> (visited on 08/30/2016).
- [67] T. Beranek, H. Merkel, and M. Vanderhaeghen. “Theoretical Framework to Analyze Searches for Hidden Light Gauge Bosons in Electron Scattering Fixed Target Experiments”. In: *Physical Review D* 88.1 (July 29, 2013), p. 015032. DOI: 10.1103/PhysRevD.88.015032. URL: <http://link.aps.org/doi/10.1103/PhysRevD.88.015032> (visited on 08/30/2016).
- [68] T. Beranek and M. Vanderhaeghen. “Study of the Discovery Potential for Hidden Photon Emission at Future Electron Scattering Fixed Target Experiments”. In: *Physical Review D* 89.5 (Mar. 10, 2014). ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.89.055006. arXiv: 1311.5104. URL: <http://arxiv.org/abs/1311.5104> (visited on 08/15/2016).
- [69] Clive Field. *HPS Target*. HPS Note 2015-008. 9/8/15. URL: https://misportal.jlab.org/mis/physics/hps_notes/viewFile.cfm/2015-008.pdf?documentId=11 (visited on 10/19/2016).
- [70] M. J. French et al. “Design and Results from the APV25, a Deep Sub-Micron CMOS Front-End Chip for the CMS Tracker”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*. 4th Int. Symp. on Development and Application of Semiconductor Tracking Detectors 466.2 (July 1, 2001), pp. 359–365. ISSN: 0168-9002. DOI: 10.1016/S0168-9002(01)00589-7. URL: <http://www.sciencedirect.com/science/article/pii/S0168900201005897> (visited on 08/04/2016).

- [71] Ted Liu et al. “The Belle II Silicon Vertex Detector”. In: *Physics Procedia* 37 (Jan. 1, 2012), pp. 867–873. ISSN: 1875-3892. DOI: 10.1016/j.phpro.2012.02.428. URL: <http://www.sciencedirect.com/science/article/pii/S1875389212017804> (visited on 08/05/2016).
- [72] Ilaria Balossino et al. “The HPS Electromagnetic Calorimeter”. In: (Oct. 14, 2016). arXiv: 1610.04319 [hep-ex, physics:nucl-ex, physics:physics]. URL: <http://arxiv.org/abs/1610.04319> (visited on 10/24/2016).
- [73] Norman A. Graf. “Org.lcsim: Event Reconstruction in Java”. In: *Journal of Physics: Conference Series* 331.3 (2011), p. 032012. ISSN: 1742-6596. DOI: 10.1088/1742-6596/331/3/032012. URL: <http://stacks.iop.org/1742-6596/331/i=3/a=032012> (visited on 09/30/2016).
- [74] Richard Partridge. “Seed-Based Tracking Algorithm for SiD”. Vancouver Linear Collider Workshop/GDE Meeting. Vancouver, Canada, July 19, 2006. URL: <https://agenda.linearcollider.org/event/316/contributions/11519/> (visited on 10/17/2016).
- [75] V. Blobel. “A New Fast Track-Fit Algorithm Based on Broken Lines”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*. TIME 2005 Proceedings of the 1st Workshop on Tracking in High Multiplicity Environments 1st Workshop on Tracking in High Multiplicity Environments 566.1 (Oct. 1, 2006), pp. 14–17. ISSN: 0168-9002. DOI: 10.1016/j.nima.2006.05.156. URL: <http://www.sciencedirect.com/science/article/pii/S0168900206007996> (visited on 10/17/2016).
- [76] Claus Kleinwort. “General Broken Lines as Advanced Track Fitting Method”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 673 (May 2012), pp. 107–110. ISSN: 01689002. DOI: 10.1016/j.nima.2012.01.024. arXiv: 1201.4320. URL: <http://arxiv.org/abs/1201.4320> (visited on 09/30/2016).

- [77] P. Billoir and S. Qian. “Fast Vertex Fitting with a Local Parametrization of Tracks”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 311.1 (Jan. 1, 1992), pp. 139–150. ISSN: 0168-9002. DOI: 10.1016/0168-9002(92)90859-3. URL: <http://www.sciencedirect.com/science/article/pii/0168900292908593> (visited on 09/09/2016).
- [78] Nathan Baltzell. *ECal Pulse Fitting*. HPS Note 2015-010. 9/9/15. URL: https://misportal.jlab.org/mis/physics/hps_notes/viewFile.cfm/2015-010.pdf?documentId=12 (visited on 10/19/2016).
- [79] Holly Szumila-Vance. *HPS Energy Calibration for the Spring 2015 Engineering Run*. HPS Note 2016-007. 6/8/16. URL: https://misportal.jlab.org/mis/physics/hps_notes/viewFile.cfm/2016-002.pdf?documentId=18 (visited on 10/19/2016).
- [80] R Niyazov and Stepan Stepanyan. *DVCS Calorimeter Prototype Reconstruction and Calibration Procedures*. CLAS Note 2005-007. 7/20/05. URL: <https://misportal.jlab.org/ul/Physics/Hall-B/clas/viewFile.cfm/2005-007.pdf?documentId=148> (visited on 10/19/2016).
- [81] Holly Szumila-Vance and Michel Garcon. *HPS/ECal Simulations: Energy and Position Reconstruction for Electrons, Positrons and Photons*. HPS Note 2014-001. 8/8/14. URL: https://misportal.jlab.org/mis/physics/hps_notes/viewFile.cfm/2014-001.pdf?documentId=1 (visited on 10/19/2016).
- [82] V. Blobel. “Software Alignment for Tracking Detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*. TIME 2005Proceedings of the 1st Workshop on Tracking in High Multiplicity Environments1st Workshop on Tracking in High Multiplicity Environments 566.1 (Oct. 1, 2006), pp. 5–13. ISSN: 0168-9002. DOI: 10.1016/j.nima.2006.05.157. URL: <http://www.sciencedirect.com/science/article/pii/S0168900206007984> (visited on 10/17/2016).

- [83] Volker Blobel, Claus Kleinwort, and Frank Meier. “Fast Alignment of a Complex Tracking Detector Using Advanced Track Models”. In: *Computer Physics Communications*. Computer Physics Communications Special Edition for Conference on Computational Physics Trondheim, Norway, June 23-26, 2010 182.9 (Sept. 2011), pp. 1760–1763. ISSN: 0010-4655. DOI: 10.1016/j.cpc.2011.03.017. URL: <http://www.sciencedirect.com/science/article/pii/S0010465511001093> (visited on 10/17/2016).
- [84] Johan Alwall et al. “MadGraph/MadEvent v4: The New Web Generation”. In: *Journal of High Energy Physics* 2007 (09 Sept. 6, 2007), pp. 028–028. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2007/09/028. arXiv: 0706.2334. URL: <http://arxiv.org/abs/0706.2334> (visited on 09/06/2016).
- [85] Hideo Hirayama et al. *The EGS5 Code System*. SLAC-R-730. SLAC National Accelerator Laboratory, Dec. 19, 2005. URL: <http://rcwww.kek.jp/research/egs/egs5.html> (visited on 09/09/2016).
- [86] S. Agostinelli et al. “Geant4a Simulation Toolkit”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (July 1, 2003), pp. 250–303. ISSN: 0168-9002. DOI: 10.1016/S0168-9002(03)01368-8. URL: <http://www.sciencedirect.com/science/article/pii/S0168900203013688> (visited on 09/09/2016).
- [87] Norman Graf and Jeremy McCormick. “Simulator For The Linear Collider (SLIC): A Tool For ILC Detector Simulations”. In: *AIP Conference Proceedings*. CALORIMETRY IN HIGH ENERGY PHYSICS: XII International Conference. Vol. 867. AIP Publishing, Oct. 27, 2006, pp. 503–512. DOI: 10.1063/1.2396991. URL: <http://scitation.aip.org/content/aip/proceeding/aipcp/10.1063/1.2396991> (visited on 10/10/2016).
- [88] CMS Collaboration. “Search for Resonant Pair Production of Higgs Bosons Decaying to Two Bottom Quark-Antiquark Pairs in Proton-Proton Collisions at 8 TeV”. In: *Physics Letters B* 749 (Oct. 2015), pp. 560–582. ISSN: 03702693. DOI: 10.1016/j.physletb.2015.08.047. arXiv: 1503.04114. URL: <http://arxiv.org/abs/1503.04114> (visited on 09/06/2016).

- [89] S. Yellin. “Finding an Upper Limit in the Presence of an Unknown Background”. In: *Physical Review D* 66.3 (Aug. 20, 2002), p. 032005. DOI: 10.1103/PhysRevD.66.032005. URL: <http://link.aps.org/doi/10.1103/PhysRevD.66.032005> (visited on 05/17/2016).
- [90] Steven J. Yellin. *Optimum Interval Software*. 2011. URL: <http://cdms.stanford.edu/Upper/> (visited on 07/18/2016).
- [91] Gary J. Feldman and Robert D. Cousins. “Unified Approach to the Classical Statistical Analysis of Small Signals”. In: *Physical Review D* 57.7 (Apr. 1, 1998), pp. 3873–3889. ISSN: 0556-2821, 1089-4918. DOI: 10.1103/PhysRevD.57.3873. URL: <http://link.aps.org/doi/10.1103/PhysRevD.57.3873> (visited on 07/18/2016).
- [92] S. Yellin. “Some Ways of Combining Optimum Interval Upper Limits”. In: (May 15, 2011). URL: <http://arxiv.org/abs/1105.2928> (visited on 07/18/2016).