

HPS Physics Readiness, Revised

HPS Collaboration

(Dated: March 4, 2016)

I. Introduction

The Heavy Photon Search Experiment successfully installed, commissioned, and ran in Hall B at the Thomas Jefferson National Accelerator Facility during CEBAF's 2015 Spring Run. The ECal, trigger, and data acquisition systems had already been successfully commissioned during a brief run in December 2014. It remained to install the Silicon Tracker (SVT), which was completed at the end of February 2015, integrate its data acquisition system with that of the experiment as a whole, and commission the whole experiment with beams. This was accomplished during March 2015, while the Hall B beamline was also being commissioned with 2 GeV electrons. Progress was interrupted when a site wide power failure occurred March 25. The consequent loss of the CHLs delayed CEBAF operations until April 19, when operations were restored with a single CHL and 1 GeV beams were delivered to Hall B. Calibration of Hall B BPMs, careful optics matching with the accelerator, and implementation of beamline setup procedures allowed efficient HPS running on nights and weekends, with weekdays going to CLAS12 torus installation. HPS ran until May 18 thanks to a two week extension granted by the lab. This extension led to the first HPS data taking with a fully functioning detector, design spec beams, currents, and trigger and data rates; and with the SVT in design position, just 0.5 mm from the beams. All aspects of the experiment worked very well and in the end roughly 2 PAC days of 1 GeV data were taken at the 0.5 mm setting. The engineering run was a great success.

The HPS Experiment was originally proposed to PAC37 [1], which in 2011 recommended conditional approval C2 for 180 days, contingent upon a successful Test Run which it urged be carried out "before the 6 GeV shutdown so that the full experiment can be carried out in a timely manner." PAC39 [2], which met soon after the conclusion of the HPS Test Run Experiment in Spring 2012, awarded HPS a scientific rating of "A" and recommended HPS for C1 approval, leaving final approval in the hands of JLAB management. In Summer 2013, HPS proposed the full experiment to DOE, which reviewed and funded it in the fall of that year. This review also served as JLAB management's de facto review of the experiment. HPS addressed the reviewers' comments and recommendations in a report to DOE HEP and JLAB in Spring 2014. The same document re-

requested formal JLAB approval, leading to JLAB management awarding 25 PAC days to HPS for an Engineering Run in 2014-2015. Management asked to see performance demonstrations from the experiment before granting additional time. Later in Spring 2014, PAC 41 selected HPS's 25 day commissioning run plus 14 days for running at 4.4 GeV for High Impact Status, noting that it "was extremely timely." They concluded "The PAC believes this is very compelling physics and hopes it will be scheduled in a timely fashion." After taking into account the roughly 10 PAC days utilized during the Spring 2015 Engineering Run, HPS still has 15 PAC days remaining from Management's original 25 day award. HPS is presently running on weekends between February 4 and March 14, 2016.

The purpose of this document is twofold. First, in response to management's request for performance demonstrations, it will review the performance achieved during the 2015 Engineering Run and demonstrate that HPS is fully ready to take, process, and analyze high quality physics data. The status of each of the HPS subsystems is discussed in the following sections: Beamline, ECal and Trigger, Silicon Tracker, and Data Taking and Processing. HPS physics performance is then summarized. It has reached the level assumed in our proposals, so the HPS reach will be what was projected. Second, based on the performance it has demonstrated, HPS includes here its request for full and unconditional approval and asks it be granted the remainder of its 180 PAC days.

II. Beamline Performance

The beamline for the HPS Engineering Run was configured according to the design presented in the proposal [3]. The whole system worked as expected right from the start. A few changes were made after the commissioning run in November-December 2014 to address issues such as beam skewness and stability, which, in the end, were not present during the main part of the Spring run when Hall-B was the only hall receiving the beam. Beam for the run was made available during swing and owl shifts during the week and for all shifts over weekends, requiring that it be restored to the Hall and delivered to the HPS target almost daily. Once BPM calibrations were completed at the end of April 2015 and sound procedures for restoring beams were established, the time to set up a high quality beam was shortened to about 2 - 3 hours. During beam restoration, the Hall-B wire scans were used to monitor beam position and the profile at different locations. The beam profile at the harp closest to the target (actually 234 cm upstream), called 2H02A, was the best monitor of whether the beam was acceptable. In Fig.1, the beam position (top) and the beam width (bottom) are shown in both the horizontal (X) and vertical (Y) directions. The position

reproducibility was better than $50\text{ }\mu\text{m}$. The Y-width, the critical parameter for bringing the SVT close to the beam, was better than $60\text{ }\mu\text{m}$, within specifications.




FIG. 1: The horizontal and vertical beam positions (top) and widths (bottom) as measured on 2H02A wire harp, mounted 234 cm upstream of the HPS target, after establishing the production beam during the HPS engineering run in May 2015.

It was possible to run the SVT Layer 1 at its proposed position, just 0.5 mm away from the beam plane, because of the excellent beam properties: remarkably reproducible beam profile (very close to the optics design), extremely low beam tails, and stable beam positions. In the end, using calibrated Beam Position Monitors (BPMs) and a feedback system that controls the beam motion at Hz level, beam position stability at the target was on the order of the beam width, $\sigma \sim 50\text{ }\mu\text{m}$. As a precaution, we inserted the SVT protection collimator, a 10 mm thick tungsten block with $4 \times 10\text{ mm}^2$ hole, to help protect against any accidental beam motion or irregularities in the beam setup. Prior to bringing the SVT close to the beam plane, we looked for short-term beam excursions during trips and recovery with a specially designed system. Studies were done by moving the harp wire or collimator edge close to the beam and recording a time history of count

rates in the downstream halo counters and the calorimeter within 15 μsec bins. Any beam motion towards the obstacle will markedly increase the count rates. No significant rate increases were observed in multiple beam trips and recoveries. In addition to the collimator and the feedback system based on BPMs, the beam Fast Shutdown (FSD) system was deployed using the beam halo counters. It could interrupt beam delivery within 5 ms if the rates in halo counters exceeded a fixed threshold. This level of protection proved to be adequate to run the experiment safely at its design conditions - 50 nA, 1 GeV beam on 4 μm tungsten target.

III. ECal and Trigger Performance

As described in our proposal to JLab management [4], several improvements and additions were made to the electromagnetic calorimeter (ECal) after the first test run in 2012. The key improvements that made significant impact on ECal performance were: 1) new large area (10x10 mm^2) APDs, Hamamatsu S8664-1010; 2) an LED based light monitoring system; 3) new, improved amplifiers and motherboards; 4) a flexible mounting system.

The ECal was installed, surveyed, and connected to electronics in the hall in October 2014. It was initially calibrated in situ using transversely penetrating cosmic muons. This allowed gains to be matched to the few % level. This was possible because the new APDs and low noise amplifiers allowed reliable measurements of the very low energy deposited by the cosmons, just ~ 18 MeV. Calibration at this level provided a reliable trigger on day one of our data taking. In both runs, November-December 2014 and April-May 2015, all 442 channels of the ECal worked as intended. Data at two beam energies, 2 GeV and 1 GeV, were taken at the proposed luminosities. Rates in individual counters were as expected from simulations. For the crystals closest to the beam this rate was ~ 1.3 MHz. No degradation of the ECal performance has been observed after exposure to radiation. Pedestals do show a dependence on rate, but this has been accounted for in the software.

A study of Coulomb scattered beam electrons using the cosmons calibration gives the energy resolution as $\sigma_E/E = 5\%/\sqrt{E}$. Further refinement of the gains using these same electrons improves the energy resolution. After a first pass, the resolution of the ECal in regions that exclude the calorimeter edges is 4%, , see Fig.2 (top), in comparison to 3.6% expected from simulation, and averages to $\sim 4.2\%$ for the entire region where Coulomb scattered electrons are detected. The entire pulse wave form is readout by the FADC for each crystal, allowing the start time of the signal to be extracted with high precision. This procedure gives ~ 0.3 ns time resolution for individual modules, as shown by the bottom plot in Fig.2. After corrections for small channel-to-channel

time offsets, this excellent resolution lets us apply tight timing cuts in the cluster reconstruction algorithm and in the cluster pair selection, thus avoiding out of time accidentals.




FIG. 2: Reconstructed energy of Coulomb scattered electrons in ECal (top). The energy resolution after calibration using Coulomb scattered electrons is $\sim 4\%$ in the regions that exclude the calorimeter edges. Single channel time resolution as a function of detected energy (bottom).

The total trigger rate for running with 50 nA electron beam on 4 μm W target was 18 kHz with 12% dead time. About 90% of triggers were from our primary A' trigger. This rate is in good agreement with expectations. The trigger starts with a 3×3 clustering algorithm in the Crate Trigger Processor (CTP) which requires hits in good time coincidence, a cluster seed energy

above threshold, and computes the total cluster energy and time. The singles triggers allow cuts on the number of hits per cluster and cluster energy. The pairs triggers add cuts on the two-cluster energy sum and difference, geometric co-planarity, and a "transverse energy" requirement. The firmware allows four simultaneous physics triggers (two singles and two pairs) plus a random trigger. For production running, we use cuts optimized for A' selection without prescaling, which are based on real random triggers and trigger simulations. The other 4 triggers were configured to support trigger diagnostics and calibration reactions (e.g. elastics) with appropriate rates. The two singles triggers were prescaled by 2^{13} and 2^{11} , and an additional loose pair trigger by 2^{11} . Random triggers were taken at a rate of 100 Hz. Trigger efficiency was studied by including diagnostics information from the Sub-System Processor (SSP) in the data stream. This included the raw Ecal cluster data and a record of which clusters passed each trigger cut. Comparison of this data with the normal calorimeter readout for actual triggers, online and event-by-event, shows excellent agreement, better than 1%. We also measured the dead time online by counting the random trigger pulser in an ungated scaler, and one which was gated off when the DAQ was busy. This was repeated with the scaler readout of the Faraday cup in the Hall-B beam dump. The two methods agree very well and give a stable dead time of 12%. A screenshot from the online monitoring display in Fig.3 shows gated and -ungated trigger rates, and pre-scaled rates, for each trigger type. Livetime and beam current are also shown. It includes a strip chart to aid visual monitoring.




FIG. 3: Online monitoring tool for trigger and DAQ.

IV. SVT and SVT DAQ Status

HPS gets good low mass acceptance for heavy photons by positioning the layer one detectors of the Silicon Vertex Tracker (SVT) a scant 500 μm from the beam. Studies described above in Section II demonstrated that the beam spot was small enough, the beam position stable enough, and the beam tails low enough to make this possible. Positioning is accomplished with the help of precision linear shifts, which raise or lower the first three layers of the SVT to within 15 mrad of the beam direction. When beam conditions warrant, the detector is moved into place. When a beam trip with excessive halo counter noise occurs, the SVT bias voltage is turned off and the detector is retracted to a safe position. Once good beams are restored, the bias voltage is turned up and the detector moved back into place. Such trips are rare enough that little luminosity is lost with this procedure.

Data from the SVT in the running configuration showed maximum occupancies on the innermost strips of about 1%, just as simulation had predicted. The silicon microstrip detectors performed as expected, with S/N ~ 25 , timing resolution ~ 2 ns, essentially perfect efficiency, and only a few dead channels out of a total of 23,004 in the system. The SVT alignment, which a pre-run survey showed to be within ~ 100 μm of design targets, has been refined with beam-based alignment using the alignment program Millepede, and is already adequate to demonstrate excellent physics performance. Scattered 1 GeV beam electrons constitute most of the observed tracks; they can be extrapolated upstream to the target and downstream to the Ecal front face with about ~ 100 μm and 1 mm resolution respectively, as expected. As shown in Fig.4, using the latest round of alignment corrections, the SVT momentum scale is accurate to $\sim 1\%$, and the (multiple scattering limited) momentum resolution for full energy electrons using a Kalman-style fitter is $\sigma_p/p = 6.7\%$, which exceeds our original design and is essentially independent of momentum. The momentum resolution improves as $1/E_{beam}$ because the magnetic field of our analyzing magnet is in proportion to E_{beam} .

Several tracking refinements have already been incorporated in the reconstruction software: an accurate fringe field map (needed to extrapolate tracks accurately to the ECal and target); beam-based silicon detector alignment constants; and the Kalman Style track fitter (Generalized Broken Line or GBL) mentioned above. Tracking efficiency is presently about 95% per track, which is more than adequate for the physics. Elastic beam electron-target electron scatters (Moller events) provide good checks of SVT performance, shown below.

HPS relies on both good invariant mass resolution and excellent vertexing capability in its search




FIG. 4: Distribution of measured momenta for Coulomb scattered full energy electrons.

for heavy photons. The Moller invariant mass peak, shown in Fig.5, demonstrates that tracking resolution is already very good. The mass resolution is essentially at design and the mass scale offset within 3%.




FIG. 5: The invariant mass of e^-e^- from Moller scattering.

All in all, the SVT is operating reliably in its design location, efficiently taking data, and

demonstrating physics-quality performance.

The SVT DAQ also performed well. Full readout and integration of the SVT with the rest of HPS were achieved just before the power outage in late March. Once beams were restored, the SVT could be timed in and first tracks seen. Initial running of the SVT was in a "safe" position, with layer 1 a full millimeter beyond the nominal position $500\text{ }\mu\text{m}$ from the beams. This was still adequate to see tracks and even get some physics acceptance. Once the final beam stability studies had been completed, the SVT was moved into its nominal 0.5 mm running position. It first began taking "real" data at trigger rates of 10-15 kHz. The SVT DAQ handled these rates, but incurred rather large dead times, in the 15-20% range, owing to fixed dead times associated with readout of the SVT front end chips. The event sizes of roughly 6 kB/event were larger than had been anticipated, further stressing data transfers and exacerbating storage requirements, and resulting in general DAQ trips from time to time. Two changes improved the SVT DAQ performance. First, the SVT timing was slightly changed to correct a small inefficiency coming from late arriving hits. Second, and more significantly, event buffering in the SVT front end readout chip was implemented, which allowed up to 4 triggers to be sent in close succession before a dead time penalty was imposed at the trigger supervisor level. With this new buffering scheme, the SVT DAQ handled trigger rates up to the 20 kHz level with only 10-15% dead time. The higher trigger rate capability let us relax some of the trigger cuts, thereby increasing trident acceptance. While the present data transfer rates are already adequate for running HPS, additional improvements have been made which provide another factor of two head room, and should bring the trigger rate capability of the SVT up to its design value of about 50 kHz without incurring additional deadtime.

V. Data Taking and Processing

As has already been stated in this document, the DAQ worked very well, allowing for a full system event rate of 18 kHz and a transfer rate to disk of 200 MB/s at about 90% livetime. This data was then transferred to the tape silo during the day while the DAQ was not operating and stored on tape at a rate of 150 MB/s.

The CODA data acquisition software performed adequately for the experiment. Some improvements to the DAQ system, which will also be very important for CLAS12, have already been implemented for the present 2016 run. The DAQ computers have been moved to 64-bit Linux, with more CPU cores, more memory and better networking. This will allow for better performance and a higher transfer rate of data, without the data transfer interfering with the DAQ performance.

The setup in the counting house now has nine operational workstations for monitoring the experiment. We have detailed monitoring systems for the beamline, ECAL, the SVT and the trigger system, which are monitored continuously during the run. In addition to the individual monitoring of the HPS subsystems, we have a system in place to monitor the overall data quality. The collaboration has built up a considerable amount of data taking expertise, including shift taking expertise and individual system expertise. This, together with the expertise of the CEBAF operators, made it possible to recover fairly quickly from the daytime shutdowns and allowed us to take production quality data during our engineering run.

We are presently making a final and complete reconstruction pass through the data. Our data blinding policy dictates that only 10% of the data will be available for study until a given analysis is ready to be frozen, at which point the full data set will be made available to that analysis. Performance studies and physics analyses have already shown that we have taken very high quality data. Several physics analyses are now underway.

VI. HPS Performance Summary

The excellent subsystem performance HPS has achieved translates into excellent physics performance. The use of a vertex constraint allows a significant improvement in the HPS momentum resolution, and consequently in mass resolution appropriate for the bump hunt. Moller scatters provide a perfect check of the invariant mass resolution. Fig.6 shows the mass resolution as a function of mass from simulation and from data at the Moller peak. HPS mass resolution has reached the level used in the proposal for determining the experimental reach. In Table I the list of key experimental parameters as proposed and measured during the 2015 run is given, demonstrating the excellent performance of the HPS apparatus.

TABLE I: HPS Key Performance Parameters

Parameter	Proposal value	Measured value
Beam Current	50 nA	50 nA
SVT Occupancy	$\leq 1\%$	$\sim 1\%$
ECAL Rates	≤ 0.5 MHz	≤ 1.3 MHz
DAQ/trigger Rate	18 kHz	19 kHz
Pair Mass Resolution	1.5 MeV	1.6 MeV
Pair Vertex Resolution	4.4 mm	4.6 mm




FIG. 6: HPS invariant mass resolution as a function of mass at 1.056 GeV. The red dots show the simulated A' mass resolution, the green triangle the simulation of the Moller mass at 1.056, and the blue square the observed resolution for Mollers.

Bump hunt and vertex analyses are making good progress. We hope to submit our first results for publication in summer 2016. To whet the appetite, Fig.7 shows the bump hunt invariant mass distribution for the 10% of the 2015 Engineering Run data which has been unblinded.

Trident production normalizes the absolute sensitivity of HPS. The reach projections in our proposal and subsequent documents depend on the experiment recording the number of fast, forward e^+e^- pairs expected from QED trident production. In fact, the observed yield of tridents including detector cuts and efficiencies, is in reasonable agreement with our earlier estimates and in good agreement with present Monte Carlo expectations, see Fig.8 below.

The plots above demonstrate that the observed reach for the HPS bump hunt will agree with projections made in the proposal. For the vertexing search, we see in Fig.9 that the all-important tails of the vertex distribution, implemented with cuts to insure high track purity so as to reduce




FIG. 7: The distribution of the e^+e^- invariant mass measured in the unblinded data (10%) from the Spring 2015 Engineering Run used in the bump hunt analysis.




FIG. 8: The observed cross-section for e^+e^- pairs (from tridents) as a function of the sum of the energies of the e^+ and e^- for the data (in black) and the Monte Carlo (in red) shows reasonable agreement between the two. The observed cross section includes acceptance and efficiency effects, and is normalized by the integrated beam current.

spurious large decay length vertices, are in excellent agreement with Monte Carlo simulation. The same cuts were used as those which had been used in the proposal, and they have the desired effect. Remember that the analysis proceeds by counting the number of vertices beyond a cut where the background is at the level of 0.5 events. Long-lived A' events will survive the cut. For tridents, the efficiency of the tracking and vertexing cuts in the data matches that in the Monte Carlo, and this efficiency for A 's (in full Monte Carlo) and hence the vertexing reach per PAC day in the 2015 data is consistent with that of the proposal.




FIG. 9: A sample of the vertex distribution for events with invariant mass between 38.5 MeV and 42.9 MeV from the unblinded sample of the 2015 engineering Run data. Other mass bins are similar. The observed vertex distribution (black) is in good agreement with the Monte Carlo prediction (red), even in the extreme tails. The slight offset of the data with respect to the Monte Carlo reflects the need for further tuning of the actual target location in the data, but does not affect the conclusion.

These vertex distributions can be fit with a Gaussian core and exponential tail. The fits are used to predict the cut in vertex z position beyond which we expect just 0.5 background event for one PAC week of data. This so-called z -cut distribution is shown in Fig.10. The distribution is in very good agreement with that used in the proposal to estimate the reach. So, modulo an efficiency factor, the reach of the present data will coincide with that projected from the proposal.

To reiterate, the trident yield and invariant mass resolution observed give HPS the reach that was projected in the proposals. Present vertex cuts reduce the far tails of the vertex distributions




FIG. 10: The flight distance in millimeters beyond which the expected background from a PAC week of data at 1.056 GeV is 0.5 events versus the e^+e^- mass in GeV. The data is given in black, the value used in the proposal to estimate the reach in blue.

as needed for the A' search, but with slightly reduced efficiency compared to the proposal. We expect further cut optimization to minimize the resultant small impact on the expected vertex reach.

VII. Requests for Approval and Scheduling

In its 2015 Engineering Run, HPS proved that it is a working experiment ready to conduct a meaningful search for heavy photons. It took enough data to begin the search at low masses. The Hall B beamline delivered the needed small spots, low beam halo, and beam position stability at the $< 60 \mu\text{m}$ level, as needed for the experiment. This beam allowed the SVT to be positioned as per design just $500 \mu\text{m}$ from the beam and operate there efficiently and reliably. The ECal pre-run calibration with cosmic rays was more than adequate to determine the ECal's energy response, set the needed trigger thresholds, and record events with low noise and good positional and energy resolution. A sophisticated, high-rate trigger which exploited both energy and position information of clusters in the ECal performed perfectly; online diagnostics proved all the algorithms fully efficient; and tridents were recorded at the expected levels in the data. Data taking worked

well, with high rates of data routinely transferred and stored.

Operationally, the Collaboration maintained the HPS subsystems, monitored performance of the detector and trigger, worked effectively with MCC to monitor beams, and took good data. Offline, reconstruction has proceeded efficiently, delivering a Pass1 for initial studies in less than a month after data taking. The analysis crew has successfully generated the needed final calibrations, alignment, and tracking improvements. Physics performance, as shown above, is at the level assumed in our proposals, and our reach per PAC week is essentially as proposed. HPS is fully ready to search for Heavy Photons and only needs adequate running time and good beam quality to deliver very topical and exciting results.

Accordingly, the HPS Experiment requests that its conditional status (C1) be removed, that it be granted full and unconditional approval and classified as an approved experiment, and that it be allotted the remainder of its PAC approved running time.

-
- [1] HPS Collaboration,
https://www.jlab.org/exp_prog/PACpage/PAC37/proposals/Proposals/New%20Proposals/PR-11-006.pdf
 - [2] HPS Collaboration, https://www.jlab.org/exp_prog/proposals/12/C12-11-006.pdf
 - [3] HPS Proposal for 2014-2015 Run,
<https://confluence.slac.stanford.edu/display/hpsg/HPS+Proposals>.
 - [4] HPS Plans for 2014 and beyond, HPS-NOTE 2015-013 (2015).