The RADPHI Detector

the RADPHI collaboration

R.T. Jones [1], T. Bogue [2], B.E. Evans, M. Kornicer,

Department of Physics, University of Connecticut, Storrs, CT 06268

A.R. Dzierba, R. Gardner [3], J.L. Gunter [4], D. Krop,
R. Lindenbusch, D.R. Rust, E. Scott, P. Smith, C. Steffen [5],
S. Teige,

Department of Physics, Indiana University, Bloomington, IN 47405

D.S. Armstrong, J.H.D. Clark [6], L.J. Kaufman [7], D.J. Steiner,

Department of Physics, College of William and Mary, Williamsburg, VA 23187

E. Frlez, D. Pocanic,

Department of Physics, University of Virginia, Charlottesville, VA 22903

J.J. Kolata, L.O. Lamm, G. Rogachev,

Department of Physics, University of Notre Dame, Notre Dame, IN 46556

C. Campbell, E. Collins, L. McGlinchey, P. Rubin [8] [9],
E. Walker,

Department of Physics, University of Richmond, Richmond, VA 23173

[1] Radphi Spokesman, 2001-present
[2] present address Department of Physics, New Mexico State University, Las Cruces, NM 88003
[3] present address Department of Physics, University of Chicago, Chicago, IL 60637
[4] present address Departments of Diagnostic Radiology and MR Research Laboratory, Mayo Clinic, Rochester, MN 55905
[5] present address National Center for Supercomputing Applications, Champaign, IL 61820
[6] present address American Physical Society, One Physics Ellipse, College Park, MD 20740
[7] present address U. Massachusetts Amherst, Amherst MA 01003
[8] present address Department of Physics and Astronomy, George Mason University, Fairfax, VA 22030
[9] Radphi Spokesman, 1998-2001

G.S. Adams, J. Napolitano,

Department of Physics, Rensselaer Polytechnic Institute, Troy, NY 12180-3590

H. Crannell, D.I. Sober,

Department of Physics, Catholic University, Washington D.C. 20064

R.R. Mammei [10] , and E.S. Smith

Thomas Jefferson National Accelerator Facility, Newport News, VA 23606

---

[10] present address Department of Physics, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061

Abstract

We describe the construction and performance of a detector system for measuring the all-photon decays of mesons photoproduced in a tagged photon beam with energies between 4.3 and 5.4 GeV and a flux of $5 \cdot 10^7$ tagged photons per second. The detector was optimized for the detection of the rare radiative decays of the $\phi(1020)$ meson. The primary detector element was a lead glass calorimeter. A single photon energy resolution of 11% at 1 GeV was achieved. Various veto and trigger components were also present. Final states with as many as seven photons were successfully detected and reconstructed.

1 Introduction

The all-neutral detector built for the RADPHI experiment [1] at the U.S. Department of Energy's Thomas Jefferson National Accelerator Facility (Jefferson Lab) was designed to detect and measure all-photon decays of ϕ mesons photoproduced in a 50 MHz tagged bremsstrahlung photon beam. The major component of this detector was a 620-element lead glass electromagnetic calorimeter. This paper describes the design, construction, calibration and operation of the detector. The performance of a detector with the capability of reconstructing all-neutral decays in a bremsstrahlung photon beam environment is of particular interest to future experiments with photon beams.

1.1 *Photoproduced Mesons with All-Neutral Decays*

The Continuous Electron Beam Accelerator Facility (CEBAF) at Jefferson Lab provides an opportunity to produce high quality photon beams via incoherent or coherent bremsstrahlung. These photon beams are characterized by high fluxes, high duty factor and superior emittance and make possible, for the first time, high statistics studies of peripherally produced mesons. Indeed, a large part of the planned physics program at an energy-upgraded CEBAF [2] involves the study of exotic hybrid mesons with quantum numbers not allowed within the conventional $q\bar{q}$ model of mesons. Incident photons are expected to be particularly effective in producing these states. In addition, the relatively large cross section for diffractive production of vector mesons, such as the ϕ meson, opens up the opportunity to study their rare radiative decays.

Decays of mesons into π^0's and η's are particularly interesting and result in all-photon final states. The π^0 nearly always decays into 2γ and the η has a large (43%) branching fraction into 2γ. A photon detector is sensitive to combinations such as $\pi^0\pi^0$, $\eta\eta$ and $\eta\pi^0$. The first two are constrained to have even spin and the third can have even or odd spin, with spin one having manifestly exotic $J^{PC} = 1^{-+}$ quantum numbers. Interesting rare radiative ϕ decays involve final states such as $\pi^0\pi^0\gamma$ and $\eta\pi^0\gamma$, these all-neutral modes being free of backgrounds from internal bremsstrahlung that are present when one or more of the decay products is charged.

1.2 Radiative ϕ Decays

The primary goal of the RADPHI experiment was to measure radiative decays of the ϕ meson, $\phi \rightarrow M\gamma$ where M is a pseudoscalar meson (P) such as the π^0 or η or a scalar meson (S) such as the $f_0(980)$ or $a_0(980)$. The decays involving the scalars are rare with branching ratios of order 10^{-4}. The ratio of their branching ratios provides information about the substructure of the $f_0(980)$ or $a_0(980)$ that could determine whether these states are conventional $q\bar{q}$ states, di-quark–anti-diquark states or $K\bar{K}$ molecules [3]. The primary decay modes of the f_0 and a_0 are $\pi\pi$ and $\pi\eta$ respectively. Some radiative meson decays of interest to the RADPHI experiment are listed in Table 1.

Experimental results for the radiative decays of the ϕ to two pseudoscalars have been recently published by groups working at storage rings. These experiments produced the ϕ mesons at rest in e^+e^- annihilations at an energy of $m_\phi c^2$ equal to 1020 MeV. These experiments at DAΦNE [4,5] in Frascati and the VEPP [6–8] in Novosibirsk provide a clean source of ϕ mesons but the maximum energy of the radiative photon, $E_\gamma^{max} = m_\phi - m_S$, is small, since the scalar mesons of interest have central masses of 980 MeV/c^2, making detection and measurement a challenge. In contrast, if the ϕ is photoproduced in a fixed target experiment, the radiative photon's energy is boosted by as much as a factor of 12 assuming a 6 GeV photon beam. Other differences in ϕ production using photons rather than e^+e^- annihilations make for an important complementarity between the two techniques.

1.3 Experimental Considerations

It is seen that all of the decays in Table 1 lead to a final state consisting solely of photons. At the energies sufficient to produce these mesons, all of the final-state photons are produced at forward angles in the laboratory frame. Thus a detector for these decays must have sufficient spatial resolution to distinguish several closely-spaced electromagnetic showers in the region close

to the forward direction. In the RADPHI experiment, this was accomplished using the Lead Glass Detector (LGD), an array of 620 lead-glass counters, described in Sect. 3.

In order to limit the rate of triggers to the data acquisition system, it was important to have additional information indicating that a ϕ or other vector meson was produced in the target. In most cases, the photoproduction of a vector meson from a proton target occurs at very forward angles, with a low-energy (typically less than 0.5 GeV kinetic energy) recoil proton emerging at a large angle (typically 40-60 degrees) to the beam direction. In order to detect these protons, a cylindrical scintillation counter system (the BSD) was placed around the production target.

These detectors, as well as various veto counters which served to reduce contamination from unwanted events, are described in Sect. 3.

2 Accelerator and Photon Beam Line

The CEBAF accelerator at Jefferson Lab is a five-pass recirculating linear accelerator, with superconducting radio-frequency cavities in its two straight sections. The accelerator delivered an electron beam of energy 5.65 GeV to experimental Hall B for this experiment.

A photon beam was generated by bremsstrahlung of electrons striking a thin gold foil. In the early portion of the run a foil of thickness $\approx 2 \cdot 10^{-4}$ radiation lengths was used. Later, due to concerns regarding electron beam-dump heating, the thickness was doubled to maintain the same photon flux at a lower electron beam current. The energy of individual photons was measured by energy analysis of the post-bremsstrahlung electrons using the Hall B tagged photon system [9]. For the RADPHI run only the portion of the tagging range between 4.38 and 5.38 GeV (counters 1 through 19) was used, that is between 77 and 95% of the electron beam energy. The electron current was selected such that the total tagged photon rate was $5 \cdot 10^7$ per second.

The photon beam remained in vacuum for about 10 meters beyond the tagger. It then emerged through a thin window into a 30-meter long polyethylene bag filled with helium. The bag terminated a few centimeters upstream of the RADPHI target. Monte Carlo studies and early beam tests indicated that this helium bag was a satisfactory solution to the problem of background generated by the beam passing through air. The distance from radiator to target was 40 meters. The photon beam dump was downstream of the RADPHI apparatus.

3 Experiment Hardware

3.1 Detector Overview

The RADPHI detector is shown schematically in Fig. 1 and its major components are listed in Table 2.

The beam was incident on a 2.66-cm diameter, 2.54-cm long beryllium target. These dimensions were chosen on the basis of GEANT Monte Carlo studies which took into consideration beam size, recoil energy loss, and detector rates. The target was suspended centered on the beam axis with three 28 AWG steel wires from a 50 cm diameter Plexiglas ring. A computer-controlled stepper motor could rotate the target into and out of the beam. Its single-step capability enabled beam-target scans to be performed remotely.

Surrounding the target and extending forward to 30° from the beam axis was a cylindrical barrel scintillator detector (BSD) which provided nearly full angular coverage for recoil protons. Surrounding the BSD was a cylindrical barrel gamma detector (BGD) composed of a lead-scintillating fiber matrix, which served to reject events (off-line) with large-angle photons.

The primary detector component was a 620-channel lead-glass wall (LGD) assembled to approximate a circle around the beam line with a 8×8 cm^2 central hole for the passage of the beam.

A 30-channel scintillator array (CPV), used to veto charged particles in the final state, was located upstream of the lead-glass array. The charged particle veto was applied during the off-line analysis because rates in the CPV were too high to permit its inclusion in the the on-line trigger.

Upstream of the target and all detectors was a 10-cm thick shield wall made of lead poured into a steel jacket. A hole through the wall permitted the photon beam to pass to the target. Mounted upstream of the wall was the lead and steel collimator shown in the second panel of Fig. 1, with inner diameter 6.5 cm, larger than the beam but smaller than the beam hole through the LGD. A scintillator hodoscope (UPV) was placed just downstream of this wall in order to veto beam halo interactions on the inner surface of the collimator. The hodoscope consisted of six horizontal and two vertical paddles, arranged around a square opening the size of the helium bag.

In the following sections each detector element is described in more detail. In these descriptions, a right-handed coordinate system is used with the origin at the center of the upstream face of the target with the z-axis along the beam direction and the y-axis up.

The active elements of the array were lead-glass bars recycled from the Brookhaven National Lab (BNL) E852 detector [10]. The dimensions of the bars were $4 \text{ cm} \times 4\text{cm} \times 45 \text{ cm}^3$. The long axis was oriented along the beam. The bars were wrapped in 0.0005 inch thick aluminized mylar and stacked in a 28×28 matrix with the corners removed so as to approximate a circular configuration (see Fig. 1). The four central blocks were removed to permit the unscattered photon beam to pass to the beam dump. A one-piece support structure held the phototubes in place relative to the lead-glass array, one tube per block, with an air gap for coupling. The entire assembly was enclosed in a light-tight box of Herculite.

The front face of the lead-glass detector was positioned 103 cm downstream of the target and subtended an angle of approximately 27° from the beam line. Many of its technologies were identical to those used in experiment E852 [11] at BNL in 1994 and 1995. The array was housed inside a steel frame on a transporter (see Fig. 2), which rolled on rails under the control of a drive system. The blocks rested on a steel platform which could be raised or lowered via an independent drive system. The dual systems permitted alignment of the array, while the horizontal system facilitated removing the lead-glass from the beam line between RADPHI running periods.

The phototubes were Russian-built FEU84-3 photomultiplier tubes with Cockcroft-Walton bases [12] designed and manufactured for RADPHI to reduce cabling and heating. Each base was powered externally by a low-voltage (15V) DC power supply, distributed via ribbon cable. Internally, an oscillator and a capacitor-diode chain multiplied the voltage to 175V DC, which drove a second chain that provided the voltages to the stages of the phototubes. The second stage was driven from both ends to eliminate voltage sag at both the photocathode of the tube and the anode.

The voltage for each base was individually controllable on a serial signal bus. Bases were assigned unique 16 bit addresses to identify row and column positions in the detector. This addressing permitted individualized voltage settings. A base could also be instructed to output a pulse with a total charge proportional to its voltage, for diagnostic purposes.

Monitoring of the lead-glass array was accomplished via a pulsed laser. A Laser Photonics LN-300C nitrogen laser illuminated a small block of scintillator which in turn illuminated a group of optical fibers. The fibers transported the light to a 1.27-cm thick sheet of polished Plexiglas which covered the upstream face of the LGD. The upstream face of the sheet was covered in aluminum foil. The fibers terminated at the edges of the sheet inside which

the light was distributed, a sufficient amount of it escaping into the lead-glass stack. Except around the three-inch hole in the center, where the intensity peaked, the illumination was quite uniform. Firing of the laser and subsequent phototube readout were coordinated as a special trigger initiated periodically during the run. The laser monitoring system was useful for setup and debugging, for initial gain adjustment and relative gain monitoring of the LGD.

3.3   The Barrel Detectors (BGD and BSD)

Two concentric barrel-shaped detectors surrounded the target, both originally built for the Jetset experiment [13,14]. These detected particles emerging from the target between 30° and 90° from the beam axis. The BSD was a hodoscope made up of three concentric cylinders of overlapping paddles. The innermost layer, of inner radius 35 cm, consisted of 12 paddles twisted to spiral around the barrel in a clockwise direction (looking downstream). The middle layer, of inner radius 36 cm, consisted of 12 paddles twisted to spiral counterclockwise. The outermost layer, of inner radius 37 cm, was composed of 24 paddles all parallel to the beam axis. The scintillators were 0.5 cm thick. The active area of the BSD extended from $z = -2$ cm to 68 cm and had full azimuthal coverage. A coincidence between counters in all three layers was used to define a triangular region, called a pixel, which enabled the recoil particle directions to be reconstructed in the off-line analysis. All 48 paddles were instrumented with Thorn-EMI 9954 phototubes.

Surrounding the scintillator array was the lead-scintillating fiber calorimeter. The BGD was installed in RADPHI to detect photons that emerge at large angles beyond the solid angle of the forward calorimeter. The fibers ran parallel to the z axis and extended from $z = -11$ cm to 75 cm. The detector was segmented azimuthally into 24 counters, each of which were read out on both ends. Upstream light readout was accomplished with 90°-bend light guides, the light being measured with Philips XP2020 phototubes. Constrained to fit into a small space just upstream of the forward detectors, the downstream light collection system consisted of a pad of Bicron (BCF-92), 1.5-mm square, multi-clad wave-shifting fibers positioned at the end of the scintillating fibers. Light from the wave-shifting fibers was detected by Thorn-EMI 9954 phototubes. The thickness of the BGD was 9.3 cm which amounts to about 5.8 radiation lengths at normal incidence. The inner radius of the counter was 39 cm.

3.4   The Charged Particle Veto (CPV)

The Charged Particle Veto (CPV) scintillator hodoscope was installed upstream of the lead-glass wall to tag charged particles. The hodoscope shad-

owed the upstream face of the lead-glass stack and so provided the capability of vetoing (off-line) events that contained charged particles in the final state. It was made up of 30 horizontal paddles, fifteen each side of center. The paddles were arranged so that neighbors overlapped both vertically and at the center as shown in Fig. 1 (except, of course, around the beam hole). The CPV scintillators were 0.4 cm thick and were staggered in z around a plane at $z = 90$ cm.

The paddles closest to the beam were narrower to approximately equalize rates in the counters. Rates were high enough to require that zener diode bases be installed on the outer sixteen paddles and transistorized bases be installed on the inner fourteen. To further reduce the rate-dependence of the gain in these counters, the phototube voltages were set at the low end of the efficiency plateau and signals were amplified by a $5\times$ amplifier before being split and analyzed. The phototubes were Thorn-EMI 9214B's.

3.5 Electronics

Photon tagger signals were discriminated with Phillips 715 constant fraction discriminator modules and the digital signals were brought to TDC's at the opposite end of the hall along 57.9 ± 0.3 m RG-58 cables. A logical OR of signals from the tagger hodoscope was sent through a 0.75c propagation-speed coaxial cable, type LM-240, via a shorter route to the trigger electronics.

Signals from all RADPHI detectors except the lead-glass array were split by impedance-matching splitters. One set of splitter outputs was sent through RG-174 cable to discriminators. The discriminator for the upstream hodoscope was the LeCroy 3412; for the barrel scintillator array, the LeCroy 3412A; for the barrel lead-scintillating fiber calorimeter, the LeCroy 3420 (constant-fraction); and for the charged-particle hodoscope, the LeCroy 4413. Discriminator outputs went to LeCroy 1877S FASTBUS TDC's operated in common stop mode, Struck 7200 VME scalers, and, in some cases, the trigger logic. The other set of outputs was sent through RG-58 cable to custom, 12-bit, integrating ADC's developed by Indiana University for E852 [11]. During the first 16 ns of the gate, the ADC sampled the signal baseline. It then integrated the signal relative to this baseline for the remaining $\sim$100 ns of the gate. The integrated charge was discriminated against two independent thresholds for use in higher level triggers, and digitized with a 12-bit successive approximation digitizer. The integration and discrimination circuitry for each channel was housed on a daughter card for easy replacement in case of failure. Digitization began 750 ns after the end of a gate and took 4 μs.

The two discriminators on each channel were arbitrarily designated the "high"

and "low" discriminators. The thresholds of all ADC discriminators were set during initialization of a data-taking period. The state of the two discriminators were represented as bits in the data word and also presented on the auxiliary FASTBUS backplane ∼100 ns after the end of the gate. The low discriminators were used in the second level trigger, and the high discriminators were used in the third level trigger.

The ADC's were read out in FASTBUS block-transfer mode. To reduce event size, only ADC channels with values at least 5 counts above pedestal were recorded. The typical pedestal width was 0.6 count.

3.6   Trigger

The RADPHI experiment employed a three-level trigger (see Fig. 3), described below.

3.6.1   Level 1 trigger

The first level was a logical coincidence of the signals from the barrel scintillator array and photon tagger, in anti-coincidence with the upstream veto array, during the LIVE condition from the hardware trigger supervisor. When the first level trigger condition was satisfied, the electronics generated gates for the ADC's and common-stop signals for the TDC's, initiated the second level trigger processing, and blocked further triggers until additional logic decided to reject or read out the event.

3.6.2   Level 2 trigger

The second level trigger requirement was a logical OR of the levels from the lead-glass detector ADC discriminators. This requirement selected events with at least one element of the lead-glass detector containing a minimum amount of energy, set to 1 GeV. In the case the level 2 trigger was not satisfied, the ADC's and TDC's had to be reset, requiring 250 ns, before another trigger could be received. The total incurred dead time for a level-1 trigger that was rejected at level 2 was about 1.2 μs. The largest single source of dead time in the experiment was from level-1 events that failed at level 2, so this reset time was a critical parameter.

3.6.3 Level 3 trigger

The third level trigger was a digital sum of the total energy in the lead-glass calorimeter. The sum was computed by a processor module designed and built by Indiana University for BNL E852. An 8-bit representation of the energy was presented to a LeCroy 2372 memory lookup where the value was compared to a threshold value equivalent to roughly 3 GeV.

In the event the level 3 trigger was not satisfied, a clear was sent to the TDC's. The ADC's had already digitized, so no resetting was necessary. The minimum dead time for an event that failed at level 3 was 8.5 μs. For every channel read into the energy sum module, this number increased by about 450 ns. The total dead time incurred for a level-1 trigger that failed level 3 had a most-probable value in the range 10 to 12 μs, and an average of 15 μs under standard beam and trigger conditions. All events which passed the level-3 trigger were saved on permanent storage for subsequent off-line analysis.

4 Calibration

4.1 LGD calibration

The LGD calibration procedure was based upon the fact that when the positions and energies of individual showers in the calorimeter are combined to form an invariant mass, the mass spectrum exhibits prominent peaks corresponding to known mesons decaying into multi-photon final states, eg. π^0, η, η', ω. These peaks are already visible in the reconstructed spectra before the gain calibration has been carried out, even when individual tubes vary in gain by a factor of 2 or more. The gains were set during the experimental run by adjusting the high voltage on individual tubes until their responses to an injected light pulse from the calibration laser were approximately equalized. The pulser equalization procedure was repeated periodically throughout the run to take into account changes in the response of individual blocks arising from radiation damage and other sources of long-term drift during the experiment. Inhomogeneities in the light distribution led to physical gains on particular channels that differed by more than a factor of two from the mean, with a 25% r.m.s. deviation. The goal of the off-line calibration was to measure these gain factors using experimental data, so that they could be used in turn to correct the data during reconstruction.

A calibration procedure based solely upon the reconstructed masses of known mesons appears to be problematic because it seeks to exploit one constraint (the known mass of the meson) to estimate multiple parameters (gain cor-

rection factors for every channel that contributes energy to the event). When many events containing a contribution from a given block are superimposed, however, the dependence of the average reconstructed mass on the gains of neighbor blocks tends to wash out, leaving the bias that comes from the gain of the block itself. Removing this bias by applying a gain correction to this channel leads to a narrower peak in the mass plot. Quantitatively, this is accomplished by adjusting individual gain factors to find the extremum of a single global function of the data,

$$F = \sum_{i=1}^{N}(m_i^2 - m_0^2)^2 + 2\lambda \sum_{i=1}^{N}(m_i^2 - m_0^2) , \qquad (1)$$

where N is the number of events in the calibration data sample and i denotes a single event in that sample. The masses m_0 and m_i are the physical mass of the meson being used for the calibration and the reconstructed mass in the LGD for event i. The first term in F measures the width of the reconstructed mass peak, while the second term is introduced with the Lagrange multiplier λ to embody the constraint $< m_i^2 > = m_0^2$.

For the purposes of RADPHI , the most convenient meson for calibration turned out to be the π^0 which appears as the dominant structure in the 2γ invariant mass plot for 2-cluster events. All events that were reconstructed with exactly two clusters and whose invariant mass lay within $\pm 30\%$ of the center of the π^0 peak were candidates for the calibration sample. The 2γ invariant mass is given by

$$m_i^2 = 2p_{1i}\, p_{2i}\, (1 - \cos\gamma_i) , \qquad (2)$$

where p_{1i}, p_{2i} are the reconstructed energies of the two showers, γ_i is the angle between the centers of the two showers as viewed from the target and the units are such that $c = 1$.

The reconstructed energy p_{ji} is approximately equal to the observed energy s_{ji} in shower j, but contains additional nonlinear corrections to account for angle-dependent attenuation and output coupling effects, written as

$$p_{ji} = (1 + g) \sum_{k \in S_j} E_{ki} = (1 + g)s_{ji} , \qquad (3)$$

where k labels an individual block contributing to shower j. The nonlinearity correction factor g is weakly dependent on the observed shower energy s_{ji} but not on the block energies E_{ki} individually. A calibration step consists of introducing a small channel-dependent gain correction factor ϵ_k such that

$E_{ki} \rightarrow E'_{ki} = (1 + \epsilon_k)E_{ki}$, where the prime is used to denote the corresponding quantity after the gain correction is applied.

Minimizing F in Eq. 1 directly with respect to the variables ϵ_k is made difficult by the nonlinear dependence of m_i on the block energies that appears in the factor g in Eq. 3 and also in γ_i. Progress can be made by observing that small shifts in the gain of a single channel will have little effect on the nonlinear correction factors or the shower centroids, but will rescale the p_{ji} of its shower:

$$\frac{\partial p'_{ji}}{\partial \epsilon_k} \simeq p_{ji} \frac{E_{ki}}{s_j} , \tag{4}$$

$$\frac{\partial m'^2_i}{\partial \epsilon_k} \simeq m_i^2 \frac{E_{ki}}{s_j} . \tag{5}$$

These approximations lead to a linear equation which is satisfied at the extremum of the penalty function F.

$$\begin{aligned}
\frac{\partial F'}{\partial \epsilon_k} &= 2 \sum_{i=1}^{N} (m'^2_i - m_0^2) \frac{\partial m'^2_i}{\partial \epsilon_k} + 2\lambda \sum_{i=1}^{N} \frac{\partial m'^2_i}{\partial \epsilon_k} \\
&= 2 \sum_{i=1}^{N} \left(m_i^2 - m_0^2 + \lambda + \sum_{k'} \epsilon_{k'} \frac{\partial m'^2_i}{\partial \epsilon_{k'}} \right) \frac{\partial m'^2_i}{\partial \epsilon_k} \\
&= 0 .
\end{aligned} \tag{6}$$

The solution is given by

$$\epsilon_k = [C^{-1}]_{kk'}(D - \lambda L)_{k'} , \tag{7}$$

where

$$C_{kk'} = \sum_{i=1}^{N} \left(\frac{\partial m'^2_i}{\partial \epsilon_k} \frac{\partial m'^2_i}{\partial \epsilon_{k'}} \right) ,$$

$$D_k = -\sum_{i=1}^{N} \left((m_i^2 - m_0^2) \frac{\partial m'^2_i}{\partial \epsilon_k} \right) ,$$

$$L_k = \sum_{i=1}^{N} \frac{\partial m'^2_i}{\partial \epsilon_k} ,$$

and the value of λ is fixed by the condition that the centroid of the reconstructed mass peak must lie at the physical mass:

$$\lambda = \frac{B + L^T C^{-1} D}{L^T C^{-1} L} , \tag{8}$$

13

where B is the mass bias $\sum_{i=1}^{N}(m_i^2 - m_0)^2$. Starting off with a uniform nominal gain factor for all channels, individual channel gain corrections are calculated using Eq. 7 and applied iteratively until the procedure converges to $\epsilon_k \to 0$ for all k. In practice, it was found that special care must be taken in the way that the matrix C is inverted. C is a square symmetric matrix with 620 rows and columns whose elements are determined statistically by sampling a finite sample of N calibration events. Even for very large samples there are instabilities that appear when taking the inverse C^{-1} which demand careful treatment. The nature of these instabilities can be best understood by expressing C^{-1} in terms of its spectral decomposition,

$$[C^{-1}]_{kk'} = \sum_{\alpha} \frac{1}{c(\alpha)} e_k(\alpha) e_{k'}(\alpha) \,, \qquad (9)$$

where $c(\alpha)$ are the eigenvalues and $e(\alpha)$ the corresponding orthonormalized eigenvectors of C. Of the 620 eigenvalues of C, there are generally a few whose values are very small and statistically consistent with zero. These terms tend to dominate the behavior of C^{-1} if it is calculated using exact methods. A better approach instead is to truncate Eq. 9 and include only eigenvectors in the sum whose eigenvalues are adequately determined by the data. This truncation implicitly recognizes that there are some linear combinations of the gains which cannot be determined from the given data sample, and simply leaves them unchanged from the initial conditions. Good convergence was obtained after 8-10 iterations.

5 Event reconstruction

5.1 Cluster pattern recognition

A photon incident on the LGD deposits energy in several LGD blocks. When these blocks are associated with each other, the position and energy of the incident photon can be measured. The algorithm to associate groups of blocks into "clusters" has three steps.

As a first step the algorithm finds the highest energy block in the LGD, called the "seed" block. Neighboring blocks are associated with the seed block into a cluster. In each step, one considers only blocks which are not already associated with clusters, or "active" blocks. Initially all blocks with ADC values over pedestal are active. Once one cluster is completed, the blocks in the cluster are removed from the list of active blocks and the process is repeated until the highest energy block remaining is below some minimum seed energy, typically

14

150 MeV. At this stage the found clusters are no larger than 3×3 blocks and not all blocks in the active list are used, i.e. associated with a cluster.

In the second step, the clusters are expanded by incorporating unused blocks contiguous with clusters into the original groups. If a block is near two step-1 clusters, the block is associated with both clusters, its energy shared in proportion to the energy contained in the central portion of the clusters.

The third step repeated step 1 but allows a seed block to have a lower minimum energy, typically 50 MeV. Once all possible seed blocks are exhausted, the photon positions and energies are determined.

5.2   Shower position and energy corrections

To a first approximation, the total energy of a reconstructed shower is equal to the sum of the observed energy in each of the blocks that belong to a cluster, associated using the algorithm described in Sect. 5.1. Improved resolution can usually be obtained by introducing a small nonlinear correction that takes into account a few-percent increase in the response of lead glass to showers above 1 GeV because of attenuation and light-collection efficiency effects in the blocks [15]. However the problem of shower reconstruction in the RADPHI experiment demands a more sophisticated approach than a simple nonlinearity correction because so many of the reconstructed showers are far from normal incidence.

To find the direction of the photon, a vector is constructed beginning at the center of the target and ending at the point (X_c, Y_c, Z_m) where X_c and Y_c are the measured coordinates (discussed below) of the shower centroid in the transverse plane of the LGD and Z_m is the longitudinal coordinate of the maximum of the shower profile inside the LGD. Determination of the the the dependence of the reconstructed momentum upon the unmeasured coordinate Z_m is especially important for showers far from normal incidence. The acceptance of RADPHI depends upon reconstructing showers as far as 25° from the normal, where the angle of incidence cannot be ignored. Furthermore at angles beyond 20° there are increasing effects from shower leakage out of the sides of the array, which introduces a bias in both the shower coordinates and the energy.

The RADPHI analysis relies on a detailed Monte Carlo shower simulation to map in three dimensions from unknowns (θ, ϕ, E), to measured values (X_c, Y_c, S) where E is the true energy of the photon and S is the observed energy summed over the cluster. Incident photons are generated in the target over a fine grid in both direction and energy, and the standard cluster algorithm is used to find the average centroid and observed energy for each grid

point. This map is parameterized using an expansion in a suitable set of basis functions to obtain algebraic forms for $X_c(\theta, \phi, E)$, $Y_c(\theta, \phi, E)$, and $S(\theta, \phi, E)$. This system of equations is then inverted using an iterative solution during analysis to provide corrected estimates for (θ, ϕ, E) for each reconstructed shower.

A slice through the function $S(\theta, \phi, E)$ at $\phi = 0$ and a range of values for E is shown in Fig. 4. The data points in the figure represent the average reconstructed shower energy for the given generated sample. The error bars show the r.m.s. spread in the reconstructed values that come from photoelectron statistics and shower fluctuations. The curve is the algebraic parameterization that emerged from a fit to a general form for $S(\theta, \phi, E)$. The depression in the response at normal incidence is an interesting effect that was discovered using this simulation, as discussed next.

The Monte Carlo shower simulation relies on the standard RADPHI Monte Carlo framework based upon GEANT3 [16]. Inside the LGD blocks the response of the lead glass is simulated in detail by the explicit creation and following of individual Cerenkov photons as they are produced in the development of the shower. The Cerenkov spectrum in the simulation is bounded at low frequencies by the response of the photocathode (cuts off below 1.8 eV) and at high frequencies by the attenuation curve of lead glass (cuts off above 4.1 eV). Between these two bounds, the photocathode efficiency, the attenuation of lead glass and the refractive index are all included with their known frequency dependence. Cerenkov light is polarized and the reflection and transmission of the polarized light is treated exactly in the simulation. The typical Cerenkov photon undergoes a number of reflections before reaching the photocathode or being absorbed. Wrapping the blocks with a thin layer of air between the glass surface and the surrounding aluminized mylar layer was an important factor in the simulation, as was the presence of an air gap at the interface between the block and the phototube. One consequence of the latter which was only realized after the simulation was carried out is that light emitted at the critical Cerenkov angle of 52° by a particle moving parallel to the block axis has zero probability of being detected at the phototube; these photons are entirely contained inside the block by total internal reflection.

Two important consequences follow from this observation. The first is that the overall shower response goes through a local minimum at normal incidence, where a large fraction of the shower particles are in the "blind spot" with momenta nearly parallel to the block axis. The simulation showed that at 20° the light output from a 1 GeV shower was about 20% greater than at normal incidence. The second consequence is that the observed lateral size of a shower must be larger than the size expected based upon the shower energy deposition profile. This follows from the fact that the most energetic particles in a shower are found near the center of the shower profile, and

16

these are the particles whose directions are most likely parallel to the block axis. The observed light yield coming from the core of the shower distribution is thus suppressed relative to light coming from the shower periphery where particle energies are lower and directions are more random. In fact, there was a discrepancy of nearly 50% between observed shower size in RADPHI data and Monte Carlo before this angle-dependent collection efficiency was discovered. After the effect was included in the simulation, the shapes of real and simulated showers were compared in terms of spatial moments up to the fourth moment, and they were found to be in agreement within errors.

6 Performance

6.1 Timing and rates

The RADPHI experiment ran for a total of about 1000 hr with photon beam on target. Data taking was divided into runs of 1-2 hr duration. Data were saved on local disk and then migrated off-line to permanent storage on the Jefferson Lab Mass Storage System. An average of 350 events were collected by the data acquisition system per second. The average event record size was 600 bytes, leading to a modest data rate of 200 kB/s. For approximately the first half of the run period, a radiator of nominal $2 \cdot 10^{-4}$ radiations lengths thickness was used with an electron beam current of 130 nA. For the second half the radiator thickness was increased to $3 \cdot 10^{-4}$ radiation lengths and the electron beam was turned down to 77 nA. At a distance of 40 m from the radiator, the RADPHI target of radius 12.8 mm intersected over 95% of the photon beam. When the RADPHI target was moved out of the beam, rates in the detector scintillator elements dropped by about one order of magnitude.

All of the signals that are relevant to the event trigger are listed in Table 3. The gate for the ADC's and TDC's was generated from the level-1 trigger (see Fig. 3). The high rate in the *cpvOR* prevented its use as a veto in the on-line trigger. The total proton photoproduction cross section integrated over the bremsstrahlung spectrum of the beam from pion threshold to the end-point leads to a total hadronic rate in the RADPHI target of 150 kHz. Increasing this rate by a factor of two to account for neutron photoreactions still does not approach the order of magnitude of the observed rates in the BSD and CPV.

GEANT-based Monte Carlo simulations of the RADPHI experiment, including the beam line, predict rates in agreement with those shown in Table 3 coming from electromagnetic backgrounds alone. Most of the rate of charged particles coming from conversions originating in the target is confined to angles a few degrees from the beam but the tails of this distribution extend out as far

17

as 60°, accompanied by a diffuse omni-directional background of low-energy deltas and gammas. The energy distribution of this background peaks in the few MeV region, except in the area immediately surrounding the forward beam hole where the typical energies are tens of MeV. The hard component coming from pair conversions of energetic beam gammas in the target escape through the forward beam hole and do not affect the experiment. Evidence that soft backgrounds of this sort are the dominant contribution to the rates in the trigger counters is seen in the marked decrease seen in Table 3 going from the innermost BSD layer *bsdOR-2* to the middle and outer layers. These three layers are in immediate contact with each other, with only the material of 5 mm of plastic plus two layers of tape shielding an outer layer from the flux seen by its inner neighbor. The fact that such a small amount of material led to a decrease in the observed rate by nearly a factor of two indicates that the source of the rate is primarily a few MeV-scale background.

The hit rate in the BGD was a strong function of threshold. During the early stages of the run, the BGD counters were operated at a low threshold corresponding to an electron-equivalent energy of 5 MeV. Under these conditions the total rate in the BGD was 800 kHz at standard operating intensity, as shown in the table. The BGD gains were lowered by a factor of 4 later in the run, effectively raising the thresholds to 20 MeV. This reduced the inclusive BGD rate to 120 kHz, in agreement with expectations based upon a total nuclear interaction rate of 300 kHz and 40% solid angle for the BGD acceptance. This observation is consistent with Monte Carlo simulations which show that the background rates in the barrel are dominated by hadronic sources for energies above 20 MeV. This is considered an effective lower bound on the energy of showers that may be reconstructed in the BGD.

There were no scalers or TDC's on the LGD signals so there is no direct measurement of the rates in the forward calorimeter. Instead these rates have been inferred from the increase with beam current in the average observed energy for a given block within the ADC gate. A minimum-bias trigger was formed by disabling levels 2 and 3 and accepting all level-1 events. By deriving the trigger from barrel and tagger elements only, an unbiased view of what is happening in the forward detectors is obtained. This view contains two components, one which is correlated to the trigger in the barrel (hadronic events are likely candidates) and the other which is uncorrelated with the barrel and consists of random forward hits that happen to fall within the 100 ns ADC gate triggered by the barrel. These two components are distinguished by running two minimum-biased runs under unchanged beam conditions, one at full beam intensity (77 nA) and the other at low intensity (2 nA).

The rate in a block is defined as the fraction of events for which the block's ADC is over threshold divided by the gate width. This rate is the sum of the barrel-correlated component which does not depend on beam current (a con-

stant probability divided by a constant gate width) and the barrel-accidental part which is linear in beam current and disappears in the limit of low rate. When the beam current was reduced from 77 nA to 2 nA the rates across the LGD decreased by only a factor of about 8 instead of the factor of 38 expected if only accidentals were present. On the other hand, if only barrel-correlated hits were present in the LGD then this rate should have been independent of beam current. This shows that at 77 nA the LGD inclusive rates are accidentals-dominated and at 2 nA they are dominated by the barrel-correlated component.

Within errors, the inclusive LGD rate that would be measured on a free-running scaler connected to each block is simply the difference in the above-defined block rates between the high and low-intensity runs. These rates are plotted as the data points in Fig. 5a as a function of the distance of the block from the beam axis. The histogram in the figure is the Monte Carlo estimate for the LGD rates arising only from electromagnetic background. Note that the expected hadronic rate for individual blocks is negligible on this scale. The excess of the data over Monte Carlo at large radius suggests that there are sources of background in the experimental hall that are not included in the simulation. The simulation includes the principal components of the Hall B photon beam line starting at the radiator and including the (empty) CLAS target and downstream yoke aperture. It was on the basis of this simulation that the helium bag and lead shielding wall upstream of the RADPHI experiment were introduced. An excellent agreement between observed and predicted rates is seen across most of the face of the LGD. The ADC threshold used in the Monte Carlo for this comparison is 15 MeV, which corresponds to the on-line threshold that was applied to the LGD data by the data acquisition.

Note that in Fig. 5 a marked depression appears at small radius, relative to the observed rates. These blocks are in the vicinity of the beam hole. In addition to suffering from the highest rates, these blocks also sustained the greatest radiation damage. The eight blocks closest to the beam axis (first data point) are the most affected, but some effects can be seen at neighboring points. These data which were taken toward the end of the RADPHI run period provide a quantitative measure of the effects of radiation damage on the response of the LGD. More insight can be provided by the pulse-height spectrum of these background hits in the minimum-bias sample. All of the spectra show a maximum intensity at threshold and an exponential tail that extends to the GeV region. The mean of this distribution is plotted in Fig. 5b as a function of block distance from the beam axis. As in Fig. 5a the points are the data and the histogram is the Monte Carlo prediction. The comparison is sensitive to the exact threshold used, which was 15 MeV for Monte Carlo but varied between 10 MeV and 20 MeV for real data, depending on the channel. Even with this caveat, the general trends are very similar between data and Monte Carlo, including the forward rise that is expected based upon the kinematics

of electromagnetic showers. The depression of the response of the innermost blocks is due in part to radiation damage, but also to the fact that these blocks contained so much background that the calibration procedure tended to artificially suppress their gains relative to their neighbors in the interest of optimizing the total shower energy resolution.

6.2 *Photon tagging at high rates*

Nominal operating conditions called for an inclusive rate of $5 \cdot 10^7$ γ/s in the tagged range from 77% to 95% of the end-point energy. The electronic sum of the signals from the 19 tagging counters, called the *taggerOR*, was included in the level-1 trigger to ensure that every trigger had a hit in at least one tagging counter within the tagging coincidence window of 20 ns.

By conventional standards for tagged photon experiments, $5 \cdot 10^7$ Hz is a high tagging rate that requires special care in the treatment of accidental coincidences. On average, within the on-line coincidence window of 20 ns one accidental tagging coincidence is expected in addition to any true coincidence that might be present for a given event. After the timing differences between channels have been eliminated in the off-line analysis and the timing resolution optimized by correcting for signal propagation delays, the coincidence window can be reduced to a few ns, at which point the accidentals probability is reduced well below unity. Even then, it is a mistake to suppose that accidentals can be ignored in the analysis. The reason for this is that, with a level-1 trigger dominated by background as in the case of RADPHI, the number of accidental tags can outnumber the true tags even if the absolute probability for accidentals on a given event is small compared to unity. As an example, consider an analysis in which the final accidental probability is 30%. If only 10% of the experimental triggers correspond to actual tagged photons interacting in the target then a 30% accidentals probability leads to a ratio of 3 in accidental/true coincidences in the tagged event sample. These are the approximate conditions for the RADPHI analysis.

The correct treatment of these accidentals is to repeat the entire analysis for any observable twice, the first time with all spectra weighted by the number of tagging hits inside the (off-line) coincidence window and the second time weighted by the number of tagging hits inside a window of similar width but shifted to exclude the coincidence peak. This shifted window is known as the "accidentals window" and spectra weighted with the tagged accidentals weight are known as "accidental spectra." Spectra weighted by the number of tagging hits inside the coincidence window are called "coincidence spectra." Coincidence and accidental spectra can be formed for an arbitrary observable in the experiment. Their difference leads to accidentals-subtracted spectra

or "tagged spectra" which reflect both in shape and magnitude what would have been measured if only the tagged component of the beam had been present. Apart from the inflation of the statistical errors which follows from the weighting and subtraction, the tagged spectra are identical to what would be obtained if the bremsstrahlung beam outside the tagger momentum window were absent and only the beam photons coming from electrons counted by the tagger were present in the beam.

The way that photon tagging is carried out in the RADPHI analysis is modified by the fact that the rates in the CPV are even higher than the those in tagger. Rates in the individual CPV counters are only a few MHz but the inclusive rate is over 80 MHz, too high to permit an efficient on-line veto and high enough to demand special care in the timing optimization and the treatment of accidental vetoes in the off-line analysis.

The most accurate time reference in the experiment comes from the tagger. The time resolution of this reference is limited by the 500 ps bin width of the TDC, but this did not limit its use in the analysis because the the light-collection dispersion of hits in the CPV scintillators is of order 1 ns. With the tagger time as a reference, the time spectra of the CPV counters exhibit a small coincidence peak of about 2.5 ns f.w.h.m. over a large continuum of accidentals. Even though the peak itself is a small bump on top of a large continuum, it actually represents 80% of all tagged events. Hence it is essential that a charged particle veto be applied before events are analyzed in terms of all-neutral final states.

In the off-line analysis, the charged particle (CP) veto is applied by taking each hit in the CPV counters and eliminating all hits in the tagging counters that fall within a certain time window of the CPV hit. The way to minimize accidental vetoes is to use as narrow a veto window as possible, without allowing charged final states to leak into the sample. A veto window of 6 ns is sufficient to contain the entire CP peak at a cost of only 30% in accidental vetos. The CP veto is applied in the off-line analysis simply by eliminating from the tagger hit list any hits which within 3 ns of any CPV hit. This is done before the tagging weights are computed, so that only the beam particles that cannot be associated with forward-charged final states are considered "tagged." Done in this way, the suppression of events with charged tracks in the forward region is carried out statistically, as a part of the photon tagging analysis, rather than as a cut on an event-by-event basis.

The event sample that remains after charged final states have been eliminated derives its event time from the barrel scintillators where a hit from a charged track, presumably the recoil proton, was required by the trigger. The overlap of counters in the three layers of the BSD forms a "pixel" which marks the point of impact of the track on the surface of the barrel. Knowledge of the position of

the impact allows the light-propagation time to be subtracted from the times of the hits in each of the three counters forming the pixel. The average of these three corrected times for a given pixel is called the "pixel time." Taken relative to the tagger, the pixel time measures the time-of-flight for the recoil particle. An average time-of-flight of 5 ns/m is subtracted from the pixel time to form the "recoil time." Defined in this way, the recoil time gives an optimum time resolution of the BSD relative to the tagger.

Fig. 6 shows the recoil time from the barrel scintillators referenced to the tagger for all level-3 triggers. All tagger hits are treated on the same footing, leading on average to more than one entry in the spectrum per event. In the case where more than one barrel pixel is present in an event, the earliest pixel defines the recoil time. The regularly spaced spikes in the spectrum are significant; they reflect the 2 ns period of the CEBAF beam.

Including all tagger hits results in the black histogram in the figure. The shaded histogram results when only the tagger hits that survive the CP veto are used. Accidental vetoes cause the overall scale to be suppressed by a factor of 0.7 after the CP veto is applied. The number of counts in the coincidence peak is reduced by much more than this factor, indicating that most of the recoils in the original sample are associated with charged tracks in the forward region. This shows the importance of the CP veto in isolating a clean sample of neutral decays.

The third histogram outlined in gray in the figure illustrates a subtle point regarding the correct treatment of accidentals in the case of CP-vetoed tags. Note that after the CP veto is applied, a depression appears in the level of the accidentals continuum in the vicinity of the coincidence peak. This is because a large fraction of the recoils occur in association with one or more hits in the CPV, even when the event that caused them did not come from a tagged photon. Thus the suppression of tagger hits with CP coincidences not only has the desired effect of removing the charged component from the recoils-tagger coincidence peak but as a by-product it depresses the accidental baseline under the peak. A similar baseline depression must also be present in the accidentals window or an over-subtraction will result when the coincidence-accidentals difference is taken.

The gray histogram in the figure illustrates what results when the correct procedure is applied as follows. The coincidences window is a region in the recoils-tagger timing spectrum of width Δt_c which contains the entire coincidence peak. The accidentals window is an image of the coincidences window shifted by an amount $\Delta t_a \geq \Delta t_c$ so that it contains none of the coincidence peak but samples a common accidentals baseline with what lies within the coincidence window. All tags within a CP veto window of Δt_v relative to any CPV hit in the event are eliminated, and the remaining tags are used relative

to the event recoil time to count the coincidence weight for the event. A similar procedure is followed to compute the accidentals weight, except that tags are eliminated which fall in a shifted CP veto window, still of width Δt_v but shifted by Δt_a. This has the effect of removing from the accidentals sample those events which have associated recoil and CPV hits that are not in association with the tagger, in exactly the same way that such events were eliminated from the coincidences sample. The dark shaded region containing the peak in Fig. 6 indicates the contribution to the signal from the coincidences window, while the dark shaded region just to the right of it indicates the accidentals window contribution. Subtraction of the two produces the neutral-tagged yield.

The power of this joint CPV-tagging analysis to isolate a clean sample of neutral γ,p events is demonstrated by the quality of the total energy signal seen in the forward calorimeter in association with a single tagging counter. The first panel of Fig. 7 shows the summed energy of all reconstructed showers in the forward calorimeter, for events with a single recoil and no extra charged or neutral energy in the barrel. These spectra are shown in order to demonstrate that a simple event-by-event tagging analysis is ineffectual at these high rates. The dashed histogram is filled for all events that contain at least one coincidence with a selected tagging counter at the high-energy end of the tagger around 5.36 GeV, while the solid histogram is taken in coincidence with a tagging counter at the low end of the tagged range around 4.41 GeV. Little difference can be seen between the two spectra apart from a bulge on the high side of the peak in the case of the higher-energy counter. The second panel shows the same two spectra after the application of the CPV-tagging analysis described above, applied to the same two tagging counters. The vertical lines indicate the expected median energy of the beam spectrum subtended by the respective tagger channels, based upon the known electron beam energy and the field in the spectrometer. A downward shift of about 100 MeV and a tail to lower energies is expected in these spectra because of the energy from the incident photon that is carried away by the recoiling target. This agreement between expectation and measurement in the total energy scale was obtained without fine-tuning, purely on the basis of the LGD calibration that adjusted the observed 2γ mass peaks to align with the physical masses of the π^0 and η.

The counting rates in the two tagging counters shown in Fig. 7 are roughly equal. The difference in the tagged yields for the two counters that is seen in the second panel reflects the decrease in the cross section times acceptance for 1-prong reactions across the photon energy range of the tagger.

The importance of the accidentals subtraction in the tagging analysis is illustrated in Fig. 8. Only events reconstructed with two clusters in the LGD are included in this sample. The upper open histogram is the total forward energy spectrum for all events in coincidence with a selected tagging counter

near the middle of the tagged range. The anti-coincidence with the CPV has been applied. The shaded histogram shows the same spectrum after tagger accidentals have been subtracted, giving the tagged spectrum for that tagging counter. The curve through the tagged spectrum is a fit to the sum of two Gaussians, one representing the central peak and the other associated with the low-energy tail. The r.m.s. width of the central peak of 280 MeV is in good agreement with expectations for the summed energy of two-cluster events as parameterized in Eq. 10 of Sect. 6.5.

6.3   Radiation Damage

Online monitoring of the LGD during the experiment indicated that the gain 8 blocks immediately adjacent to the beam hole decreased as the run progressed. This observation was based upon the laser monitor system, the raw pulse-height distributions, and the channel gains which were periodically recalibrated during the run, as outlined in Sect. 4. A similar effect was seen, but to a lesser degree, for the next layer of blocks once removed from the beam hole. During a pause in the experiment, visual inspection of the blocks indicated that the glass was darkening, a well-known effect of radiation damage on lead glass. Fig. 8 illustrates the gain reduction with beam time (roughly proportional to integrated radiation dose) for a typical block adjacent to the beam hole. It is apparent that the gain change is a gradual, cumulative effect rather than a sudden change which might be characteristic of a beam mis-steering event.

The magnitude of the gain loss (order 40%) was such that it could be compensated by adjustments of the PMT high voltages. This was done periodically during the experiment. The last datum in Fig. 8 shows the result of one adjustment. However, this is only a partial solution, since a radiation damaged block produces fewer photoelectrons in the PMT, resulting in a degraded energy resolution which cannot be compensated by increasing the PMT gain. Thus it was desirable to "heal" the radiation damage as much as possible.

Radiation damage in lead glass is known to be temporary, and to largely heal itself on the time scale of a few months. The healing can be accelerated by the use of ultraviolet (UV) light. This approach was adopted for the most affected blocks. During an extended down-time in the run, the PMT and base for selected blocks were removed and a UV light guide attached to a quartz-envelope Mercury vapor lamp was inserted. The output of the lamp was 5 W/cm^2 in the range 300 to 480 nm, with a peak intensity at 365 nm. The affected blocks were each illuminated for periods of 6-8 hours. These blocks showed a gain increase of 30% following this treatment, nearly recovering their initial performance.

24

It should be noted that the change in the response of the blocks due to radiation damage and the gain recovery following UV treatment was even more dramatic in the data from the laser monitoring system (typically a factor of two change). The difference between the shift in the pulser response and that seen in the gain constants from the calibration can be qualitatively understood by noting that the laser illuminates the front of the block and thus measures the transmission of the entire block, while the showers seen in the calibration data create Cerenkov light throughout the volume of the block, and are therefore less sensitive to attenuation effects in the upstream region of the block. The radiation damage is expected to be concentrated within one or two radiation lengths of the front surface of the detector, and this was confirmed by visual inspection, in qualitative agreement with the difference between the laser monitor data and the calibration data.

6.4 Yields

The integrated live-time of the experiment under standard operating intensity and trigger was $1.56 \cdot 10^6$ s corresponding to an integrated luminosity of 27.1 pb^{-1} in $\gamma\,^9$Be interactions within the tagged region of the beam photon spectrum. If nuclear corrections may be neglected, this corresponds to a little over 100 pb^{-1} in γp interactions. The trigger acceptance for γp$\rightarrow$pX reactions where X decays to all-neutral final states is on the order of 50%, where losses come mainly from absorption of recoil protons in the target and escape of final-state gammas through the forward hole in the calorimeter.

The total numbers of events collected and reconstructed at successive stages in the analysis are shown in Table 4, broken down according to the shower multiplicity in the forward calorimeter. The sample shown in a given row consists of all events shown on the previous row that satisfy the one additional requirement listed in column 1. None of the cuts shown actually modify the way a given event is reconstructed, so there is no cross-over between the columns. Each column represents an independent and distinct set of initial triggers. Which multiplicity is assigned to a given event depends upon the tuning of the cluster-finding algorithm. The results shown in the table were obtained with settings which give an acceptable trade-off between under-counting and over-counting of clusters.

6.5 Energy and position resolution

Lacking a source of electron or photon showers of a well-defined energy, the RADPHI experiment had to rely on the observed width of known mesons to deduce the energy resolution of the lead-glass calorimeter. The observed width

of narrow mesons such as the π^0 and η that undergo 2γ decay is determined by the single-shower energy and position resolutions of the LGD. In cases where the spatial contribution can be neglected, the single shower energy resolution was extracted by selecting pairs with one of the two showers in a narrow energy window and examining the energy spectrum of the other, for a given cluster-separation angle. This energy spectrum shows peaks that correspond to the masses of the π^0 and η and whose line-shapes are convolutions of the energy response functions for the two showers plus contributions from uncertainties on the shower centroid positions. The contribution from spatial resolution to the width of the peaks was minimized by focusing first on the η, which is associated with pairs of showers that are well separated on the face of the calorimeter. By analyzing the dependence of the peak width on the energies of the individual showers, the convolution was inverted to obtain the r.m.s. resolution for single showers as a function of shower energy without introducing a model for the energy dependence. Once the energy resolution had been determined in this way, the spatial resolution was then examined by looking at the excess width of the π^0 peak over what was expected based upon energy resolution alone. In the end, a unified analysis including both energy and spatial resolution effects was able to reproduce both the π^0 and η profiles.

Assuming that the spatial resolution is not important for $\eta(2\gamma)$ decays, the r.m.s. shower energy resolution can be extracted without assuming any functional form for its dependence on shower energy. This model-independent solution was then compared with the standard parametrization [17] of the lead glass energy resolution

$$\frac{\sigma_E}{E} = \frac{B}{\sqrt{E}} + A \,. \tag{10}$$

The first term on the right contains the effects of photoelectron statistics, while the second term wraps up all of the systematic block-to-block differences and calibration errors which prevent the width of the response function from approaching a delta function in the high-energy limit. In order to describe the π^0 peak width, it is necessary to introduce a model for the shower spatial resolution, which itself depends upon shower energy. The energy dependence is proportional to $1/\sqrt{E}$ with a proportionality constant that depends on the size of the LGD block [18]. Eq. 11 is adopted for showers at normal incidence, with the constant C expected to be around 7 mm·GeV$^{-\frac{1}{2}}$:

$$\sigma_x = \frac{C}{\sqrt{E}} \,. \tag{11}$$

In the RADPHI geometry many of the showers are far from normal incidence and so shower depth fluctuations of roughly one radiation length also contribute to the centroid position resolution along one of the spatial axes. This

was taken into account in the analysis by projecting one radiation length along the shower axis onto the transverse plane and adding it in quadrature to the base term in Eq. 11 to obtain the major axis of the error ellipse on the shower centroid. The final values for all parameters were determined by simultaneous analysis of the η and π^0 data where all of the above effects are included for both.

The invariant mass-squared of two photons is given by $2E_1 E_2(1 - \cos\theta_{12})$, where E_1 and E_2 are the corresponding photon energies and θ_{12} is their angular separation. The expression for the variance of m^2 in terms of the variances V_E on the shower energies and V_x, V_y on the shower positions is given by

$$V_{m^2} = \sum_{i=1}^{2} \left[\left(\frac{\partial m^2}{\partial E_i}\right)^2 V_{E,i} + \left(\frac{\partial m^2}{\partial x_i}\right)^2 V_{x,i} + \left(\frac{\partial m^2}{\partial y_i}\right)^2 V_{y,i} \right. $$
$$\left. + 2\left(\frac{\partial m^2}{\partial x_i}\right)\left(\frac{\partial m^2}{\partial y_i}\right) V_{xy,i} \right] , \tag{12}$$

where index i counts photons, and $V_{xy,i}$ is the covariance between the x and y coordinates of the shower centroid for shower i. Eq. 12 can be rewritten in the form

$$\frac{V_{m^2}}{m^4} = \frac{V_E(E_1)}{E_1^2} + \frac{V_E(E_2)}{E_2^2} + \frac{V_s(m^2)}{m^4} , \tag{13}$$

where spatial derivatives and variances are grouped into the term $V_s(m^2)$. A lengthy but straight-forward calculation gives V_s in terms of the underlying position uncertainties of the two showers. For example, the contribution from the x coordinate of shower 1 is given in by

$$\frac{\partial m^2}{\partial x_1} = -\frac{2p_{z,1}}{z_0 E_1^2} \left[p_{x,2} (p_{y,1}^2 + p_{z,1}^2) - p_{x,1} (p_{y,1} p_{y,2} + p_{z,1} p_{z,2}) \right] , \tag{14}$$

where component $k = 1, 2, 3$ of reconstructed shower $j = 1, 2$ is written $p_{k,j}$. The other spatial derivatives have a similar form and they can be obtained by the proper variable substitution in Eq. 14. The position z_0 of the shower maximum is not directly measured, but from Monte Carlo it is estimated to be about 20 cm into the glass and weakly dependent on energy. The results of the resolution analysis are not sensitive to changes in z_0 on the order of one radiation length. The fact that departure from normal incidence cannot be ignored requires a treatment of centroid errors that couples the uncertainties in the x and y coordinates. The results are summarized in Eqs. 15-17.

27

$$V_x = \frac{C^2}{E} + (X_0 \sin \theta \cos \phi)^2 \tag{15}$$

$$V_y = \frac{C^2}{E} + (X_0 \sin \theta \sin \phi)^2 \tag{16}$$

$$V_{xy} = X_0^2 \sin^2 \theta \cos \phi \sin \phi \tag{17}$$

The geometry and material properties of the LGD lead to nominal values $C = 7.1$ mm·GeV$^{\frac{1}{2}}$ and radiation length $X_0 = 31.6$ mm [19].

The η sample used in the resolution study contained 8M events. Pairs were selected with shower separation $D_{\gamma\gamma} \geq 20$ cm to enhance the η peak, as shown in the second panel of Fig. 10. The π^0 sample consisted of 15M events selected by limiting the invariant mass $m < 0.5$ GeV/c^2. The corresponding invariant mass distribution is shown on the first panel of Fig. 10 contains 8M events. The π^0 and η peaks are fitted with a Gaussian over a polynomial background. The width of the η peak receives only a few percent contribution from the V_s term in Eq. 13. Neglecting this contribution permitted a model-independent extraction of the energy dependence of the single-shower energy resolution in the LGD, as described above. The results are shown by the data points in Fig. 11. They are in good agreement with a fit based upon the form of Eq. 10, as indicated by the solid curve in the figure. Including the spatial contribution in Eq. 13 in the fit and allowing C to vary as a free parameter leads to the dashed curve in Fig. 11 which lies very close to the solid curve, showing that spatial resolution plays an insignificant role in determining the observed width of the η in this sample.

If the results of the analysis of the η width are applied to the π^0 sample, the predicted widths come out smaller than the measured widths by about 40% when the spatial contribution is neglected. This indicates that the spatial and energy resolution contributions are comparable in the case of the π^0. A global analysis was applied to both the η and the π^0 samples in which both the energy and spatial contributions to the error were included in a uniform way. A total of 50 measured widths taken from different regions in photon energies and separation angles were used as inputs, and the constants A, B and C were treated from Eqs. 10 and 11 as free parameters in the fit. The results of the fit for the energy resolution are shown by the dotted curve in Fig. 11. The best-fit values from the fit were $A = 0.035 \pm 0.008$, $B = 0.073 \pm 0.006$ GeV$^{\frac{1}{2}}$ and $C = 6.4 \pm 0.1$ mm·GeV$^{\frac{1}{2}}$. The χ^2 returned by the fit is 1.5 per degree of freedom.

7　Conclusions

The RADPHI detector was able to detect exclusive final states consisting of up to seven photons. The calibration, trigger and reconstruction techniques developed worked well in the challenging high-rate environment of a bremsstrahlung photon beam. Acceptable resolution was achieved down to energies lower than 200 MeV. The detector was able to survive the high rate environment with minor and correctable radiation damage.

8　Acknowledgments

References

[1] "A Measurement of Rare Radiative Decays of the ϕ Meson", Jlab Experiment E-94-106.

[2] A. R. Dzierba, "QCD Confinement and the Hall D Project at Jefferson Lab", eConf (2001) C010430.

[3] F. Close, *et al.*, Nucl. Phys. B389 (1993) 513.

[4] A. Aloisio, *et al.*, Phys. Lett. B536 (2002) 209.

[5] A. Aloisio, *et al.*, Phys. Lett. B537 (2002) 21.

[6] M. N. Achasov, *et al.*, Phys. Lett. B479 (2000) 53.

[7] M. N. Achasov, *et al.*, Phys. Lett. B485 (2000) 349.

[8] R. R. Akhmetshin, *et al.*, Phys. Lett. B462 (1999) 380.

[9] D. I. Sober, *et al.*, Nucl. Inst. Meth. A440 (2000) 263.

[10] S. Teige, *et al.*, Phys. Rev. D 59 (1999) 1.

[11] R. R. Crittenden, *et al.*, Nucl. Inst. Meth. A 387 (1997) 377.

[12] P. Smith, *et al.*, Nucl. Inst. Meth. A414 (1998) 466.

[13] R. T. Jones, *et al.*, Nucl. Inst. Meth. A343 (1994) 208.

[14] S. Passaggio, *et al.*, Nucl. Inst. Meth. A371 (1996) 188.

[15] B. B. Brabson, *et al.*, Nucl. Inst. Meth. A332 (1993) 419.

[16] R. Brun, *et al.*, GEANT User's Guide, Program Library W5013, CERN.

[17] R. Fruhwirth, *et al.*, Ďata Analysis Techniques for High-Energy Physics, Cambridge University Press, 2000.

[18] R. Wingman, International Series of Monographs on Physics, Vol. 107, Oxford University Press, 2000.

[19] C. P. Steffen, Ph.D. thesis, Indiana University, 2001.

Table 1
Some meson decays to all-neutral final states.

Decay	Photon multiplicity
$\phi \to f_0(980)\gamma \to \pi^0\pi^0\gamma$	5γ
$\phi \to a_0(980)\gamma \to \eta\pi^0\gamma$	5γ
$\phi \to \eta\gamma$	3γ
$b_1(1235) \to \omega\pi^0$	5γ
$a_0(980) \to \omega\gamma$	4γ
$\omega \to \pi^0\gamma$	3γ
$\omega \to \eta\gamma$	3γ

Table 2
Glossary of detector subsystems.

Symbol	Full name	Description
LGD	Lead Glass Detector	a circular array of 620 $4\times4\times45$ cm^3 lead glass blocks
BSD	Barrel Scintillator Detector	a 3-layer cylindrical scintillator array surrounding of the target
BGD	Barrel Gamma Detector	a cylindrical lead-scintillating fiber array surrounding the BSD
CPV	Charged Particle Veto	a plane scintillator hodoscope covering the face of the LGD
UPV	Upstream Pair Veto	a scintillator array surrounding the beamline upstream of the target

Table 3

Signal rates, widths and dead-time factors for the major components of the RADPHI detector. Rates shown are those obtained at full operating intensity.

signal	rate (Hz)	duration	dead-time fraction
$taggerOR$	$5.0 \cdot 10^7$	5 ns	0.07 [1]
$cpvOR$	$8.2 \cdot 10^7$		0.04 [2] [4]
$bgdOR$	$8.0 \cdot 10^5$		0.03 [3] [4]
$upvOR$	$0.9 \cdot 10^6$	40 ns	0.04
$bsdOR$-2	$6.4 \cdot 10^6$	10 ns	0.005 [5]
$bsdOR$-1	$2.7 \cdot 10^6$	10 ns	0.002 [5]
$bsdOR$-3	$1.8 \cdot 10^6$	10 ns	0.001 [5]
$bsdAND$	$7.9 \cdot 10^5$	20 ns	0.016
Level-1 [6]	$5.8 \cdot 10^5$	10 ns	0.006
Level-2	$2.7 \cdot 10^5$	1.2 μs [7]	0.31
Level-3	$7.9 \cdot 10^3$	15 μs [8]	0.11
readout	350	650 μs [9]	0.23

[1]based on 25 ns gate from individual channel discriminators
[2]based on 10 ns gate from individual channel discriminators
[3]based on 40 ns gate from individual channel discriminators
[4]this is a veto inefficiency, no effect on experimental live-time
[5]from 5 ns gate on individual channel discriminators
[6]the level-1 logic signal before it is gated by the busy signal
[7]minimum dead time from receipt of level-1 gate to end of fast-clear when event fails level-2
[8]average dead time from receipt of level-1 gate to end of fast-clear when event passes level-2 and fails level-3
[9]average event readout time

Table 4

Event yields at successive stages in the RADPHI analysis, broken down according to shower multiplicity in the forward calorimeter.

sample	2γ	3γ	4γ	5γ	6γ	7γ
on tape	$2.29 \cdot 10^8$	$2.47 \cdot 10^8$	$2.12 \cdot 10^8$	$1.34 \cdot 10^8$	$7.16 \cdot 10^7$	$3.50 \cdot 10^7$
tagged [1]	$9.05 \cdot 10^7$	$1.06 \cdot 10^8$	$9.35 \cdot 10^7$	$5.89 \cdot 10^7$	$3.06 \cdot 10^7$	$1.42 \cdot 10^7$
cp veto [2]	$3.27 \cdot 10^7$	$2.34 \cdot 10^7$	$1.72 \cdot 10^7$	$8.09 \cdot 10^6$	$3.50 \cdot 10^6$	$1.33 \cdot 10^6$
1-prong [3]	$2.50 \cdot 10^7$	$1.81 \cdot 10^7$	$1.33 \cdot 10^7$	$6.27 \cdot 10^6$	$2.71 \cdot 10^6$	$1.02 \cdot 10^6$
fiducial cut [4]	$2.28 \cdot 10^7$	$1.36 \cdot 10^7$	$8.93 \cdot 10^6$	$3.55 \cdot 10^6$	$1.36 \cdot 10^6$	$4.45 \cdot 10^5$
barrel γ veto [5]	$3.83 \cdot 10^6$	$1.89 \cdot 10^6$	$1.29 \cdot 10^6$	$4.59 \cdot 10^5$	$1.65 \cdot 10^5$	$4.33 \cdot 10^4$

[1]after subtraction of accidental tags

[2]after elimination of tags with coincidences in CPV

[3]single isolated charged hit in barrel scintillators

[4]no showers in LGD located on inner and outer boundaries

[5]no energy in BGD apart from that associated with the recoil proton

Fig. 1. Cut-away views of the RADPHI detector. In the 3-d view (panel 1) the CPV detector has been omitted to reveal the lead-glass array. All counters are shown in the 2-d view (panel 2).

Fig. 2. The LGD transporter. The glass stack is shown, with the uninstrumented corner blocks removed for clarity.
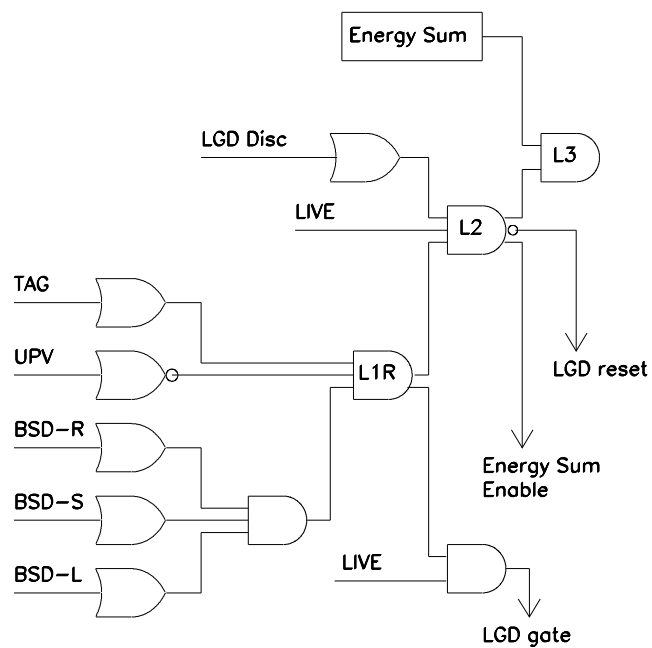
Fig. 3. The RADPHI trigger logic. L1R denotes Level-1 raw, Level 1 included L1R in coincidence with LIVE.
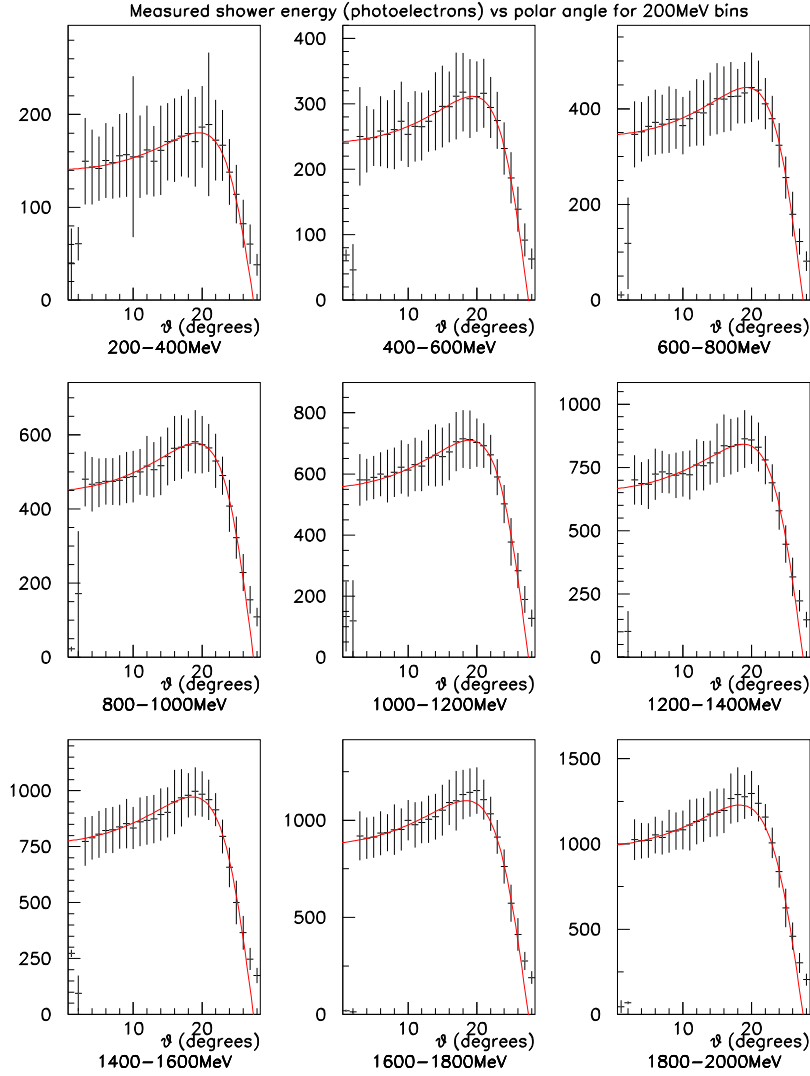
Fig. 4. Measured photoelectron yield in LGD as a function of angle for gammas generated in the target at $\phi = 0$. The generated shower energies are distributed uniformly within the respective energy intervals. The data points are the average and r.m.s. yields for a sample of several hundred simulated showers. The curve is a fit to the data.

Fig. 5. Characteristics of unbiased flux observed in individual blocks in the LGD as a function of distance from the beam. The points are derived from data and the histograms from a Monte Carlo simulation of the electromagnetic background coming from the beam and target. Rates (first panel) and average energy (second panel) include all hits over 15 MeV.

38

Fig. 6. Spectrum of the recoil time relative to the tagger. The black histogram includes all tagger hits, while the shaded histogram refers only to those tagger hits that have no CPV hit within the veto window. The darker shaded regions indicate the coincidences and accidentals windows. Bin width is 0.5 ns.

Fig. 7. Summed energy from all reconstructed forward clusters in an event. The sample includes all events with between 2 and 7 clusters, a single recoil and no extra energy in the barrel. The first panel shows the spectrum corresponding to events in coincidence with tagging counter 19 (solid histogram) at the low-energy end of the tagged photon range and counter 1 (dashed histogram) at the high-energy end. The second panel shows the same two spectra after the CPV-tagging analysis described in the text has been carried out for the same two tagging counters.

40

Fig. 8. Summed energy from all reconstructed forward clusters in an event before (open histogram) and after (shaded histogram) tagging accidentals have been subtracted. Only tagging channel 10 was included in the analysis, corresponding to a 50 MeV bin in the beam energy spectrum centered at the position of the vertical line. The sample includes all events with exactly 2 reconstructed clusters, a single recoil and no extra energy in the barrel. The curve is a fit to the sum of two Gaussians.

Fig. 9. The effect of radiation damage on the central part of the detector. The last point shows the gain after an adjustment of the phototube high voltage.

Fig. 10. Invariant mass distribution of $\pi^0(2\gamma)$ (first panel) and $\eta(2\gamma)$ (second panel). The parameters shown are the height (P1), mean (P2) and sigma (P3) of the Gaussian peak fitted to the data over a polynomial background described by parameters P4-P6.

Fig. 11. Energy resolution of showers in the LGD obtained from analysis of the 2γ sample. Points represent the free solution to the η squared-mass resolution measurements when the contribution from the spatial resolution has been neglected. The solid line represents the fit to the η data with the standard energy resolution model. The dashed line represents the fit to the η data when the spatial contribution is taken into account by Eq. 13. The dashed and solid lines are nearly indistinguishable on this scale. The dotted line corresponds to the simultaneous fit to the η and π^0 data with the same function.