Technical Note
radphi-2004-301

Tagging Photons with the Radphi Detector

Richard Jones
University of Connecticut, Storrs CT

March 10, 2004

Abstract

The analysis of data taken with the Radphi experiment in summer 2000 has been hindered by the lack of a reliable method for associating an event in the detector with the energy of the photon in the beam that produced it. In principle the energy of the beam is tagged on an event-by-event basis by a coincidence between the detector and one of the channels of the tagging counters. This association is complicated by the fact that accidental coincidences with the tagger make event-by-event association ambiguous. Furthermore many of the interactions in the target that produce sufficient gamma energy to satisfy the trigger also include charged particles to which the calorimeters are largely insensitive. The result is that studies that look for association between the tagger channel and total energy in the detector have shown little if any correlation. This report shows how these problems can be overcome. An excellent agreement between the tagger channel and measured energy is demonstrated, in quantitative agreement with the resolution expected based upon prior studies of meson widths. The analysis procedures that incorporate the method are explained, and the total event yields obtained using photon tagging are provided.

The Radphi experiment took data during summer 2000 using the tagged photon beam line in Hall B at Jefferson Lab. The photon beam had an incoherent bremsstrahlung spectrum with an end-point of 5.65 GeV. The photon source was operated in two modes, the first with a thin radiator of approximately 10^{-4} radiation lengths and 130 nA of electron beam current and the second with a thicker radiator and 77 nA. Both configurations correspond to an inclusive rate of $5 \cdot 10^7$ /s in the tagging focal plane. The experimental trigger was based on a triple coincidence between three thin scintillator layers in the barrel detector referred to as the $bsdAND$. The range of interest in the beam spectrum was defined by a subset of the Hall B tagging counters subtending from 75% to 95% of the end-point energy. The electronic sum of the signals from these 19 counters called the $taggerOR$ was combined in coincidence with the $bsdAND$ to ensure that every trigger had a hit in at least one tagging counter within the tagging time window of 20 ns. Introduction of a veto from beam halo counters upstream of the detector, denoted the $upvOR$, completes the definition of the level-0 trigger as $bsdAND \cdot taggerOR \cdot \overline{upvOR}$.

The upv veto was used to suppress interactions coming from scraping of the beam on material upstream of Radphi, and just served to gate off the experiment for 40 ns whenever a halo track was seen. The $upvOR$ rate varied from 1.4 MHz near the beginning of the run down to around 400 KHz where it remained for the latter third of the run period, after the accelerator operators had improved the beam alignment. The effect of the $upvOR$ veto on the experiment was simply to reduce the live time of the experiment by a factor of 2-5%. Its effects are ignored in the following analysis.

1 Tagging of neutral triggers

1.1 charged particle veto

The analysis of tagging coincidences would be simple if a sample of fully-reconstructed events could be formed based on the information from the Radphi detector alone. Because the detector does not reconstruct charged tracks in the forward region, this would require a reliable method for extracting neutral events from the much larger background of events containing charged pions and kaons. For this purpose the design of the experiment incorporated a forward scintillator wall to veto charged particles, the cpv, but the rates in these counters were too high to permit an efficient cp veto in the online trigger. In order to pass the higher levels of the online trigger the total measured energy in the forward calorimeter had to be greater than about 3 GeV, which tends to favor reactions with several gammas in the final state, but the task of separating the mixed events from those with only gammas falls to the offline analysis.

Before a correspondence can be found between the measured energy and tagger channel a way must be found to effectively use the cpv to remove charged events. Fig. 1 shows the time difference between the $bsdAND$ and a hit in any one of the cpv paddles. There are 10^6 $bsdAND$ triggers in this sample. From the level of the continuum in Fig. 1 one deduces that the inclusive rate in the cpv was about 80 MHz. This rate was spread more or less evenly across the 30 scintillator paddles in the cpv so that the rate in a single counter was only a few MHz. A veto around the entire cpv coincidence peak in Fig. 1 would contain an average of 2 accidentals, which would imply a prohibitive loss of good all-neutral events if such a cut were applied.

A much tighter coincidence peak is obtained if the time correlation is taken between the tagger and the cpv. In Fig. 2 is shown the time difference between any hit in the cpv and any hit in the tagger. The tagger hits were selected outside the tagger coincidence window so that the cpv-tagger correlation in Fig. 2 is unbiased by the trigger. Part of the improvement in the coincidence timing in Fig. 2 over Fig. 1 comes from the use of times for each tagger channel that have been corrected for individual channel offsets, rather than using the electronic OR. It will be shown below that a careful selection of events with a single hit in the bsd and suitable corrections for time-of-flight in the barrel leads to a cpv-bsd timing peak that is nearly as well-defined as that shown in Fig. 2, but the fact remains that the tagger provides the single best time reference for this experiment.

1.2 choice of reference time

In discussions of timing spectra, the point is often raised that the complexities of the online trigger must be taken into account, in that sometimes the $taggerOR$ signal and sometimes the $bsdAND$ determined the common time reference for the TDC's. While this consideration does affect the time spectra for individual counters, it is not relevant to the spectrum of a time difference between hits in two counters. For time

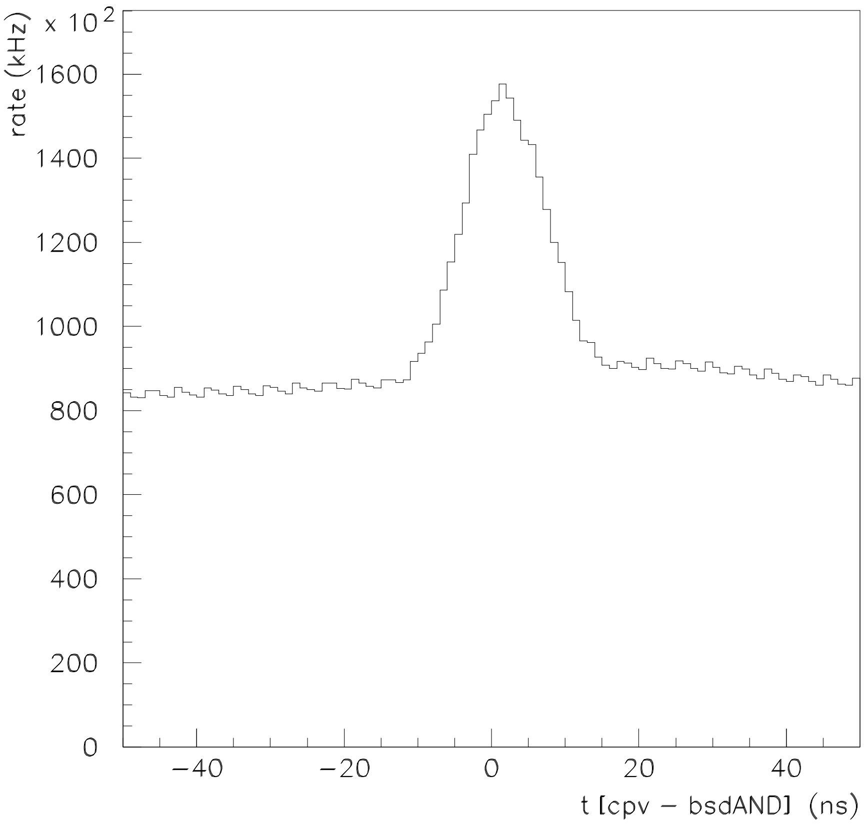

Figure 1: Time difference between any hit in the cpv and the bsdAND signal which generated the event trigger, normalized to 10^6 events. The continuum measures the background rate in the cpv, mainly coming from electromagnetic conversions in the target. Bin width is 1 ns.

differences only the logical condition of the trigger, not its precise timing, is relevant. Henceforth in this study it may be assumed that all time spectra use a single hit in a single tagger channel as a time reference.

The choice begs the question of *which* tagger time is suitable as a reference for any given event, since the rate in the tagger is sufficiently high that the typical event has 2 or more tagger hits within the region of the coincidence peak. The answer is proposed that *each* tagger hit is an equal candidate for a time reference for an event and all must be treated on exactly the same footing. While it may seem strange to be including a single event more than once in the analysis, the following conceptual argument serves to justify it.

Consider an analysis where the cross section for a well-defined photon energy is to be measured. In that analysis, one might well ignore the TDC information from all but one of the tagging counters, and look only at the coincidences with the relevant tagging channel. For such a case one need not worry about ignoring what other tagging channels would do with any given event, any more than the Radphi analysis needs to concerned with what was happening in regions of the tagger that were not turned on for the experiment. The single-channel tagging spectrum would be shaped very much like the one shown in Fig. 2, with true coincidences appearing as a peak on top of an accidental continuum. The single-channel tagging case is simple to handle tagger because the electronic double-pulse resolution on a single channel of 20 ns prevents more than one hit within a reasonable coincidence window and, even if that were not true, the rate in an individual counter is low enough to make the problem of double-coincidences irrelevant.

The analysis would proceed as follows. A timing spectrum would be formed in which the recoil time for the event is subtracted from the time for every hit in the tagging channel. Based on this timing spectrum, two windows would be assigned, the first containing the coincidence peak (the coincidences window) and the

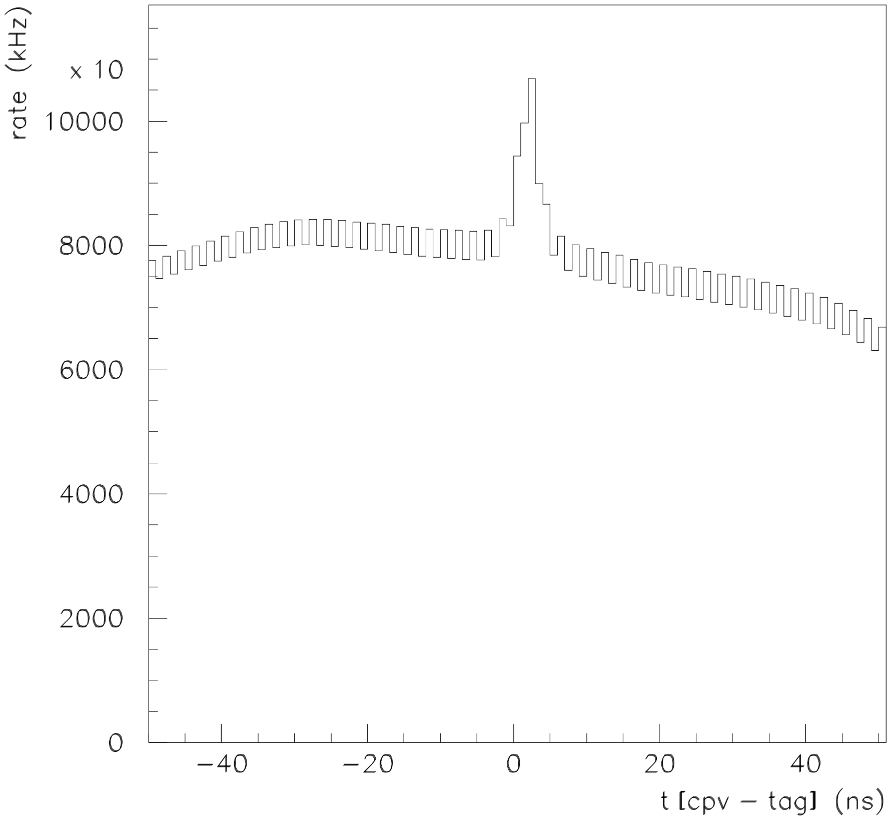

Figure 2: Time difference between any hit in the cpv and an arbitrary hit in the tagger, normalized to 10^6 tagger hits. The 2 ns microstructure of the CEBAF beam is visible in the spectrum, together with the main peak coming from tagged photons producing charged particles which strike the cpv. Bin width is 1 ns.

other outside the peak but nearby and of a similar width (the accidentals window). In the subsequent analysis of the information from the detector, two copies of every histogram would be formed: one which is filled on every event with a tagging hit in the coincidences window, and the other on every event in the accidentals. As a final step in the analysis, every coincidence spectrum would have its corresponding accidentals subtracted away. It is apparent that this procedure leads to an unbiased estimate in all experimental variables for the contribution coming only from beam photons within the range of the selected tagging counter.

This procedure can be generalized to cover the case where an inclusive measurement is desired of a photoproduction cross section averaged over a larger region of the focal plane. In such a case, the above single-channel analysis is repeated N times and then the resulting spectra are summed, leading to an average cross section weighted by the bremsstrahlung spectrum times tagging efficiency. The detailed shape of this weighting function is not of interest in the Radphi analysis; diffractive processes are weakly dependent on beam energy. The primary concern is to settle on a method that extracts the tagged component of the experimental spectra without bias and with known errors.

1.3 errors and weighted spectra

The method presented in the preceding section for subtracting accidentals from tagged spectra is equivalent to keeping a single copy of every analysis histogram and using a weight factor to keep track of the tagging information. The weight factor is calculated as follows. For every hit in any tagging counter that falls inside the coincidences window the weight factor is incremented by one, and for every tagger hit inside the accidentals window the weight factor is decremented by one. This leads to some events having a weight factor of zero, i.e. being ignored, but that is inevitable in an analysis where spectra are being subtracted. Viewed in this way, the method admits a straight-forward error analysis. Histogramming packages such as HBOOK have built-in facilities to track the statistical errors associated with weighted events.

1.4 combination of tagger and cpv

With the question of a reference time for Radphi resolved, the problem remains of how to efficiently remove the charged component from the sample. Noting the tight timing correlation between the cpv and the tagger, it is suggested that the cpv might be used to remove tagger hits associated with charged particles *before* they are considered for coincidence with the recoil detectors. Used offline in a similar way in which the upv was used online, a hit in the cpv simply gates off the tagger for a short period. Only tagger hits that do not appear near cpv hits can contribute toward the accidental-subtracted yields described above. In this way the coincidence with the tagger can be used to perform a double function, both to remove beam photons outside the tagging energy range and to remove all interactions with forward-going charged particles. In such an analysis one gives up the ability to decide on an event-by-event basis which are neutral/tagged and which are not and in return obtains an efficient and unbiased method for estimating the spectra that would be obtained if one were able to separate them.

The statistical errors are minimized when one choses as narrow a coincidence window as possible, which means in turn that a recoil time as precise as possible should be obtained for comparison with the tagger. Events with a single pixel in the bsd were selected and the average of the times for the three counters associated to the pixel was taken. When compared with the tagger, the coincidence peak shifted slightly later with increasing pixel index, which corresponds to the position of the pixel moving from upstream to downstream. For pixel rings 0-6 the slope was measured to be 0.5 ns/pixel. The dependence is linear within measurement errors. This dependence is taken out by forming a software “recoil time” from which the above z-dependent offset is subtracted from the pixel time average. The coincidence peak for pixel ring 7 was much weaker and did not reflect the trend of the others. It is left uncorrected.

The spectrum of the recoil time is shown in Fig. 3. The open histogram is obtained in reference to all tagger hits, while the shaded histogram refers to tagger hits with a cp veto. A 6 ns window around the cpv-tagger coincidence peak was used in forming the cp veto. The reduction factor of about 0.6 in the accidentals level between the two histograms in the figure reveals the cost of the veto. The coincidence peak is reduced by a much greater factor than 0.6 which confirms the prior assumption that most of the triggers in the Radphi data sample come from cp events. This statement will be quantified in the following sections.

The neutral-tagged histogram in Fig. 3 exhibits several features that are important to the understanding of how the gated tagger coincidence analysis works. The sawtooth structure seen across the spectrum is a consequence of the 2 ns pulse period of the accelerator and is expected in the difference spectrum between any two observables whose spread from time-of-flight is of order 1 ns or less. The prominence of the spikes in Fig. 3 indicate that resolution of the recoil time is of order 1 ns, which is at the level predicted on the basis of Monte Carlo for diffractive photoproduction from a nuclear target [1].

The tapering off of the accidental spectrum outside the interval [-40,40] ns is a consequence of the selection of tagging hits only within the window [-50,50] ns and the fact that the trigger defines the recoil time to lie within a 20 ns region around t=0. The region [-30,40] ns in Fig. 3 can be considered as a free-running correlation between the recoil and tagging counters with a given that *something* had to happen in a 20 ns coincidence window near t=0 in order for the event to pass the level-1 trigger. This would be expected to give rise to a small enhancement in the continuum over a 20 ns region containing the coincidence peak. This enhancement that comes from the level-1 trigger requirement is seen in Fig. 3 between the limits of approximately -5 and +15 ns. Unfortunately the boundaries of this region are very diffuse, evidence of the significant differences between the timing of individual bsd and tagging signals that formed the trigger. That the level of accidentals is roughly the same inside and outside the hardware trigger window is a consequence of running the tagger in at high rates.

This 20% enhancement is also seen in the shaded histogram in Fig. 3 but is partially obscured by the depression in the accidentals that appears around the coincidence peak. This depression occurs because, when the cp veto is applied to the tagger hits, it has a side effect of also suppressing a fraction of the recoils, even though whatever caused the recoil and cpv hits might be in random association with the tagger. This depression is a physical effect because it appears only when the cp veto is applied offline. The width of the depression corresponds to the 6 ns width of the cp veto cut.

Correct subtraction of accidentals in Fig. 3 requires the selection of non-overlapping windows for coin-

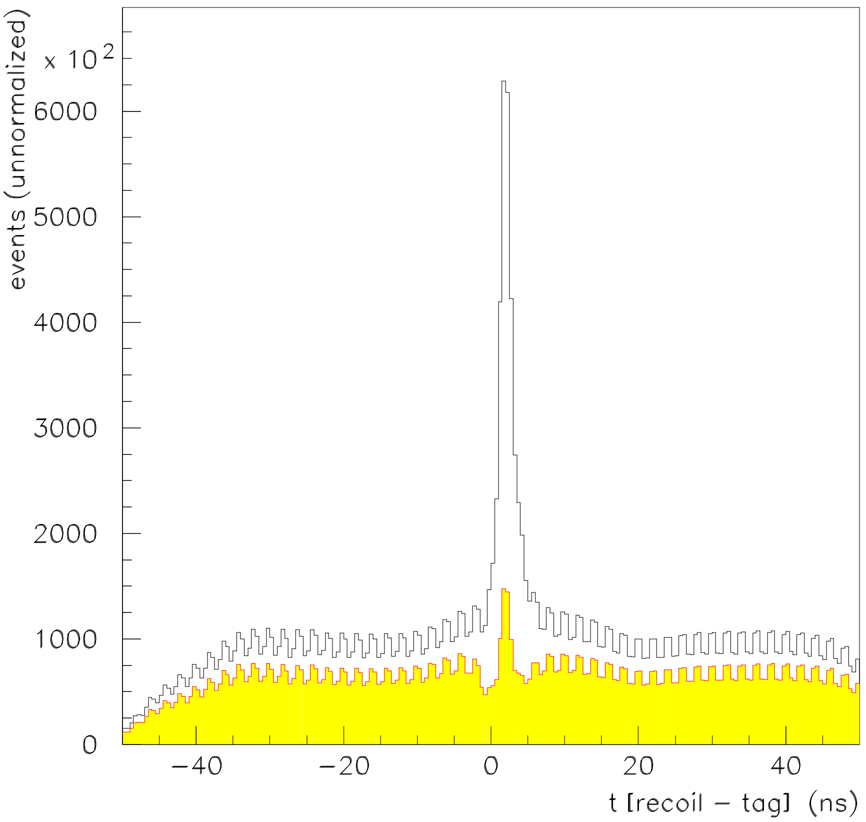

Figure 3: Spectrum of the recoil time relative to the tagger. The open histogram includes all tagger hits within a ±50 ns window around the event trigger, while the shaded histogram refers only to those tagger hits that have no cpv hit within the range [-1,5] ns in Fig. 2. Bin width is 0.5 ns.

cidences and accidentals, both within the range of the 20 ns hardware coincidence. As a compromise, the choice was made of the region [-1,5] ns for coincidences and [6,11] ns for accidentals. A small renormalization of the accidental weights can be used to account for the partial clipping of events in the accidentals window by the level-1 trigger. These two windows are shown in Fig. 4. The dashed histogram shows a similar depression in the accidentals window to the one discussed above that appears in the coincidences window when the cp veto is applied. Accidentals or coincidences, only tagger hits without nearby cpv hits are counted when forming a time reference for registering the recoil time. The two depressions are correctly canceled out in the subtraction.

2 Critical test: total neutral energy

2.1 total energy spectra

A test of above procedure would be to look at the total energy from all reconstructed showers in the forward calorimeter. This is a critical test, as can be seen from the total forward energy spectrum for all events shown in Fig. 5. The shape shows a threshold around 3 GeV coming from the online level-3 trigger, followed by a monotonically decreasing intensity with energy. Although the maximum energy is roughly consistent with the end-point energy of 5.65 GeV, there is very little evidence of an enhancement in the region of subtended by the tagger 4.39 - 5.38 GeV. Fig. 6 shows the same spectrum weighted by the tagger-cpv coincidences weight

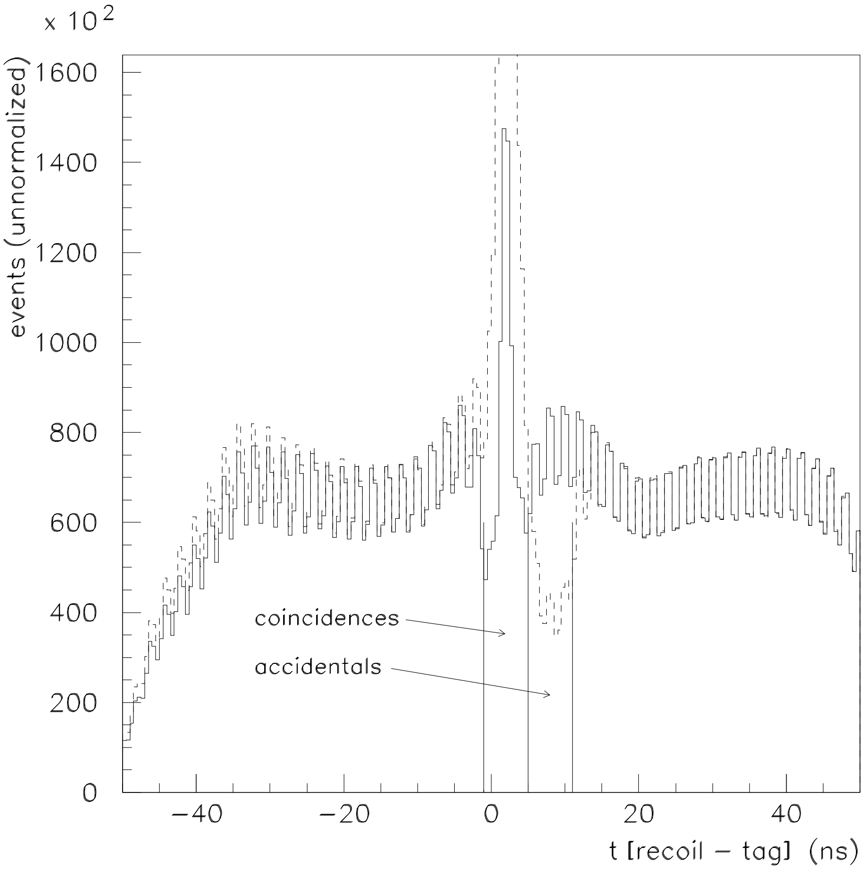

Figure 4: Recoil timing spectrum for coincidences [solid] and accidentals [dashed] with the tagger and cp veto. Bin width is 0.5 ns.

for tagging counter 10. The dashed histogram in the same figure indicates the corresponding spectrum for the accidentals for counter 10. The difference spectrum is shown in Fig. 7. The sum of two Gaussians gives an excellent empirical description of the spectral line-shape, with approximately equal numbers of events in the narrow spike at the maximum and the low-energy tail.

Fits like the one shown in Fig. 7 have been performed for all of the tagging channels. All channels exhibit a spectral shape similar to the one shown in the figure. Table 2 shows the fit parameters for the narrow Gaussian peak, and compares it with the energy calculated based upon the tagging channel. While not perfect, the agreement between the peak energy and the mid-energy for the given tagger channel is impressive when it is remembered that the energy calibration was based entirely on the invariant mass scale of two-cluster combinations. This is the first time that the total energy scale in the LGD has ever been verified.

The above spectra are all taken for the case of two-cluster events. Similar plots are obtained for events with higher multiplicities, albeit with slightly larger r.m.s. values as expected based on the known calorimeter resolution. Fig. 8 compares the subtracted total-energy spectra for the highest and lowest counters in the tagging focal plane. The correspondence with the energy computed based upon the tagger channel is similar to the two-cluster case, systematically low by about 100 MeV. An offset of this order is expected in order to account for the energy carried away by the recoil proton.

A comparison of the shapes of the two spectra in Fig. 8 shows that there is a significant tail on the low side of the resolution function, perhaps indicating some loss of energy out the sides of the detector combined with the effects of dead channels that appeared throughout the run, but mainly explained by the distribution of the kinetic energy transfered to the recoil proton. The tail is less prominent in the spectrum from the

Table 1: Parameters of two-Gaussian fit shown in Fig. 7. Statistical errors are at the level of the least significant digit.

parameter	Gaussian 1 value	Gaussian 2 value
height	6378	4191
mean (GeV)	4.732	4.355
r.m.s. (GeV)	0.266	0.495

Table 2: Parameters of a Gaussian fit to the narrow peak in the total forward photon energy spectrum for subtracted spectra corresponding to each individual tagging channel. Statistical errors are at the level of the least significant digit.

tagging channel	nominal mean (GeV)	mean (GeV)	r.m.s (GeV)
0	5.360	5.200	0.280
1	5.310	5.160	0.288
2	5.210	5.120	0.278
3	5.190	5.060	0.276
4	5.140	5.030	0.270
5	5.090	4.960	0.273
6	5.040	4.910	0.263
7	5.020	4.880	0.254
8	4.940	4.820	0.261
9	4.880	4.770	0.264
10	4.830	4.730	0.267
11	4.770	4.670	0.259
12	4.710	4.600	0.250
13	4.640	4.550	0.251
14	4.590	4.510	0.253
15	4.550	4.480	0.253
16	4.510	4.430	0.248
17	4.460	4.380	0.255
18	4.410	4.350	0.246

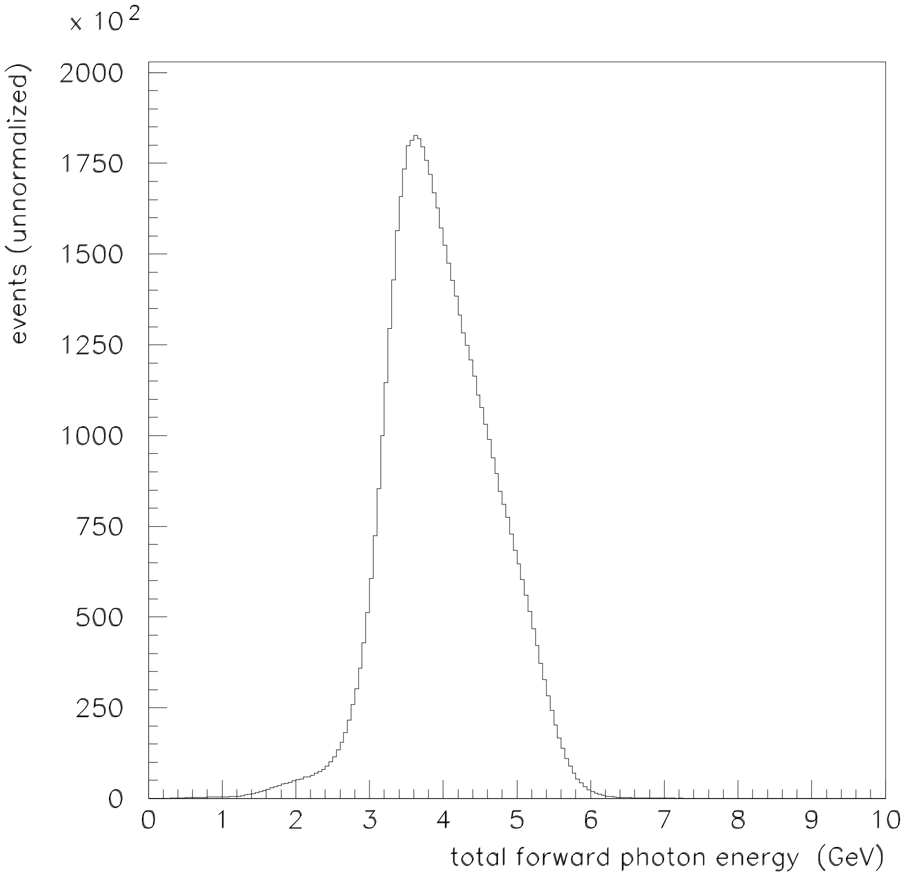

Figure 5: Summed energy from all reconstructed forward clusters in an event. The sample includes all events with between 2 and 7 clusters, a single recoil and no observed neutral energy in the barrel.

low energy end of the tagger only because it is truncated by the online level-3 threshold around 3 GeV. The widths of the narrow peak listed in Table 2 are in good agreement with the r.m.s. expected on the basis of the empirical resolution function [2] for single showers

$$\frac{\sigma_E}{E} = 0.036 + \frac{0.073}{\sqrt{E}} \tag{1}$$

whose value varies from 265 MeV at 4.41 GeV to 310 MeV at 5.36 GeV. An interesting study that has not yet been performed would be to look for a correlation between the energy transfered to the recoil proton and the missing energy in the forward that gives rise to the low-energy tail in Figs. 7-8.

2.2 event selection

The above spectra were taken on a small subset of the total Radphi event sample that was selected to enhance fully-contained forward events. The selection entails three cuts: rejection of events with at least one cluster near the boundaries of the calorimeter (fiducial cut), rejection of events with neutral energy in the barrel (barrel gamma veto) and rejection of events with more than one pixel cluster in the barrel scintillator (single recoil). The above analysis was repeated with each of the above cuts suspended in order to examine their effects on the ability of the analysis algorithm to reconstruct an exclusive final state.

The fiducial cut is implemented using the software library function `lgd_cluster_cleanup()`. It removes events with any clusters whose center lies near the outer or inner boundary of the detector. It has essentially no effect on the sample of events with two clusters, but has increasingly visible effects on higher multiplicity

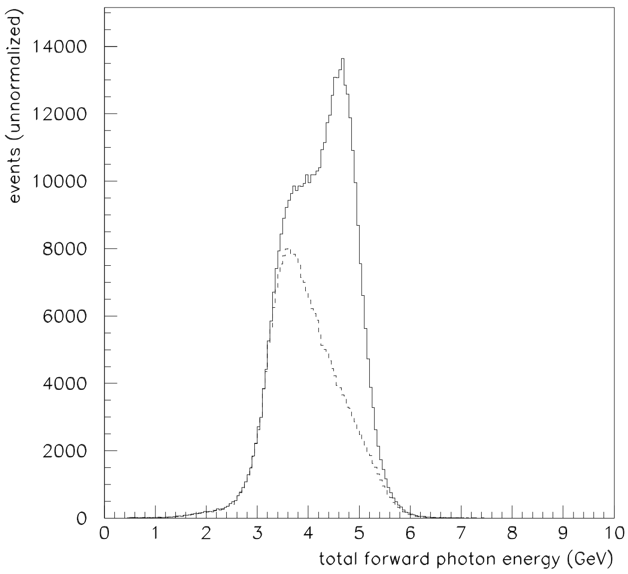

Figure 6: Summed energy from all reconstructed forward clusters in an event, weighted by the the tagger-cpv coincidence condition (solid histogram) and accidental condition (dashed histogram). The sample includes all events with between 2 clusters, a single recoil and no observed neutral energy in the barrel.

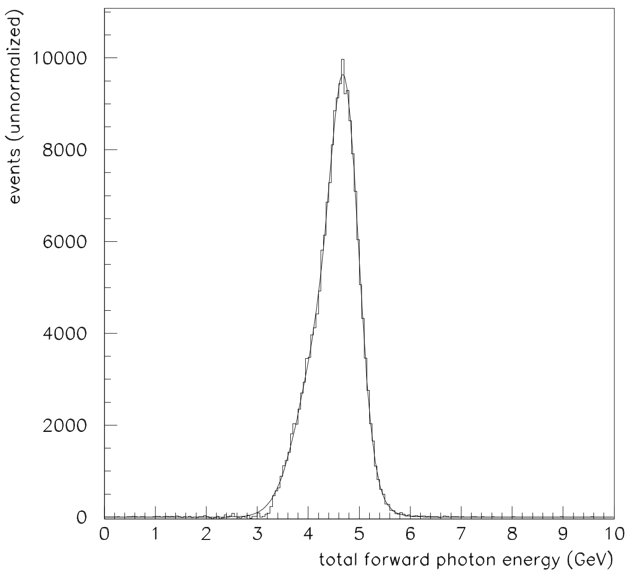

Figure 7: Difference spectrum from Fig. 6 which estimates the spectrum that would be obtained from neutral tagged events free from accidental effects. The fit is to a sum of two Gaussians with parameters given in Table 1.

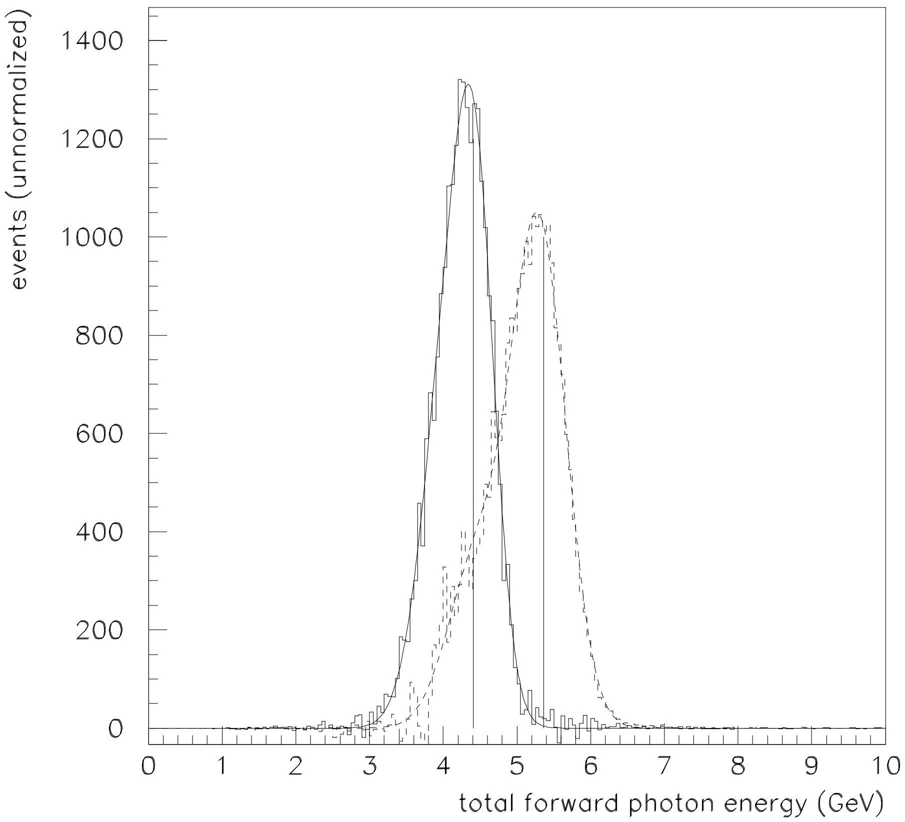

Figure 8: Total energy spectrum from the forward calorimeter for five-cluster events with one recoil and no neutral energy in the barrel. The spectra shown correspond to the lowest (solid) and highest (dashed) tagging channels in the experiment. The calculated mid-energies for these two tagging channels are indicated with the vertical lines.

samples. For example, only 2/3 of the five-cluster events survive this cut. When the cut is suspended the maximum of the total energy peaks are shifted downward relative to the above figures by about 4% and the r.m.s. of the peaks are increased by about 10%. The conclusion is that the fiducial cut is valuable for analyses involving higher multiplicity final states, and involves relatively modest cost in terms of statistics.

The barrel gamma veto is implemented by looking for any bgd hit over ADC threshold and vetoing the event if no corresponding hit is found in one of the two straight scintillators overlapping the bgd counter. In contrast to the fiducial cut above, the barrel gamma veto is both very powerful and costly in terms of statistics. Fig. 9 shows the total forward photon energy spectrum corresponding to tagger channel 10 for all two-cluster events without (open histogram) and with (shaded histogram) the barrel gamma veto enabled. The figure shows that without this cut the total forward energy would bear little relation to the tagger energy; without the cut there is no real peak in the spectrum and the maximum occurs well below the energy of the tagging channel. The cut eliminates all but about 15% of the events, a fact which is more or less the same for all cluster multiplicities. Without it there is very little chance of knowing that the event was contained in the forward calorimeter. The conclusion is that either this cut or a refinement thereof must be strictly applied for all subsequent analysis or a method must be developed to include the barrel neutral clusters in the reconstructed event.

In fact, the power of this cut for cleaning up fully-contained forward all-neutral events suggests that a piece of what appears in the low-energy tail of the peaks in Figs. 7-8 corresponds to events with neutral barrel energy that "leaks" through the barrel gamma veto cut by appearing in the two bgd counters that overlap

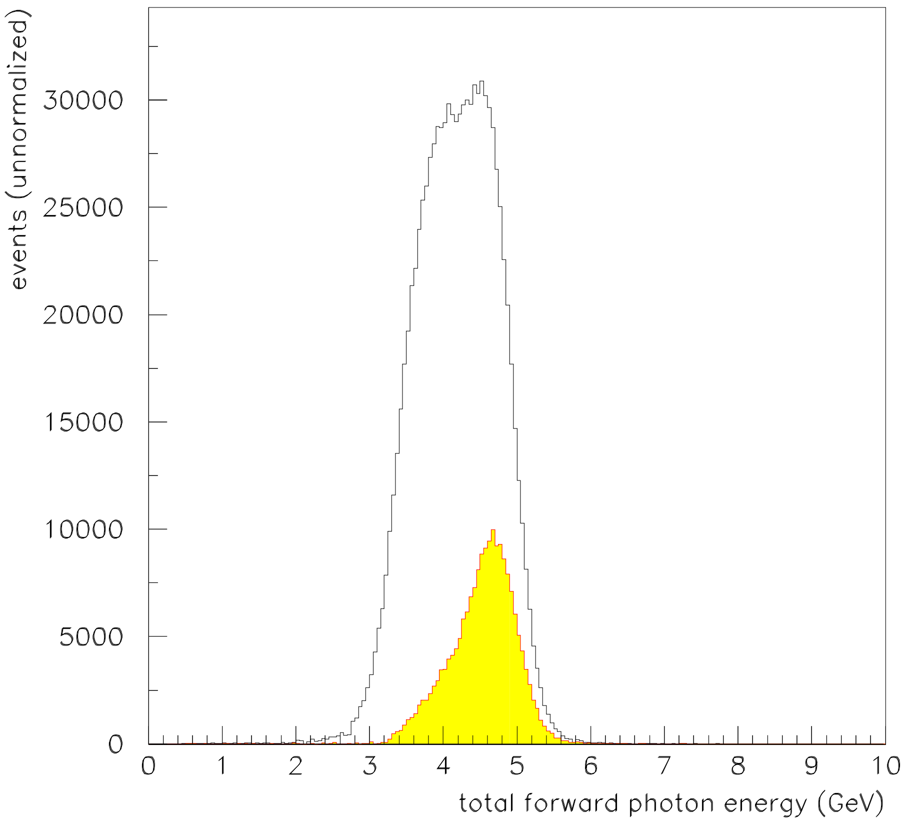

Figure 9: Total energy spectrum from the forward calorimeter for two-cluster events without (open histogram) and with (shaded histogram) the barrel gamma veto enabled.

the recoil pixel. These two counters are not included in the barrel neutral energy sum because they normally contain some energy from the stopping recoil track. A more sophisticated barrel gamma veto cut might make use of z-coordinate information in the bgd to separate the recoil track from additional neutral clusters in these counters. One way to test this is to repeat the above analysis with the zero neutral energy cut in place and in addition to require that among the two bgd counters behind the recoil pixel, only one of them may see a hit. This should reduce by a factor of two the leaks mentioned above. The result is shown in Fig. 10, which confirms the hypothesis. The r.m.s. of the peak is decreased by 10% by the additional requirement of at most one bgd hit, and the spectrum of the eliminated events shows no peak at the tagger energy. The cost of the additional requirement in terms of statistics is 20% in good agreement with expectations based upon the leaky veto hypothesis, leading to the conclusion that the barrel gamma veto should include the additional requirement of no more than one counter hit.

The single-recoil cut, by contrast, has a relatively weak effect on either the quality of the total energy signal or the statistics. It cuts roughly 20% of the statistics from the sample, regardless of multiplicity, and adds these events mainly to low-energy tail of the total energy peak. No other explanation for this effect would be required than simply that the additional bsd hits in the event allow more neutral energy in the bgd to "hide" from the barrel gamma veto cut. Since this cut is relatively cheap in terms of statistics and enables a cleaner identification of neutral clusters in the barrel, the conclusion is that this cut should be retained in subsequent analysis.

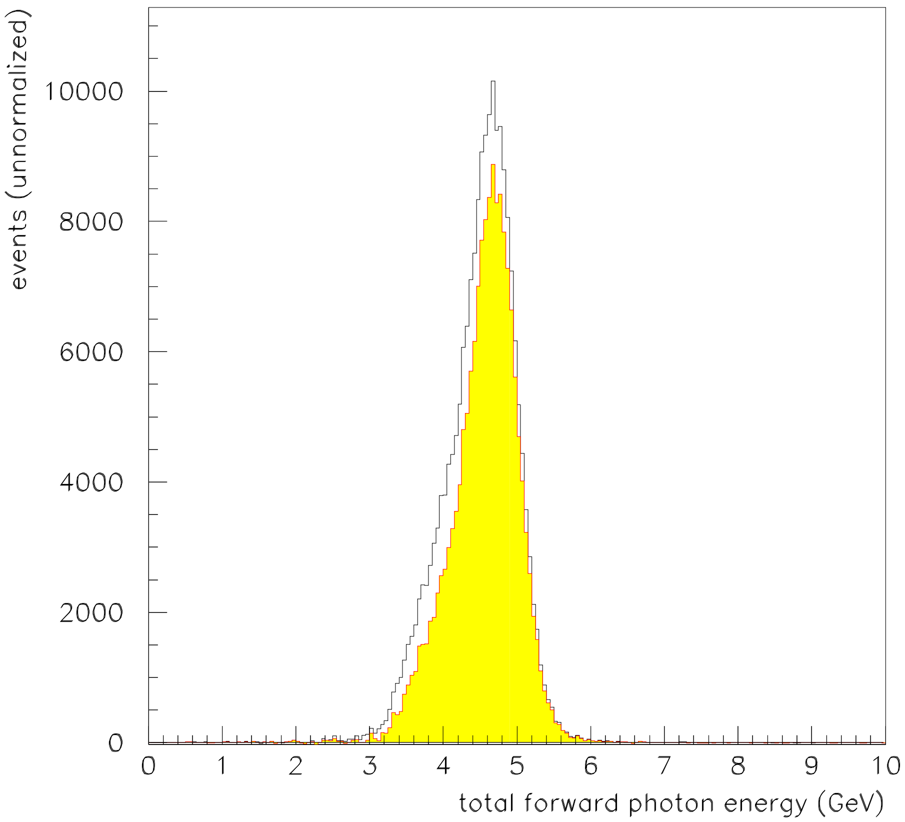

Figure 10: Difference spectrum for tagging channel 10 from Fig. 7 (open histogram) compared with the same spectrum with the additional requirement that only one bgd counter saw any energy during the event.

3 Implementation in code

The software tools for carrying out the above analysis have been verified and imported into the Radphi cvs area. Three existing functions in the `makehits` package have been mostly rewritten, together with their associated data structures. The `make_recoil()` function now groups together neighboring pixels into clusters described by the structure `recoil_t` which contains suitably averaged pixel information, including the corrected *recoil time* described in a previous section and an average dE/dx for the passing charged particle. In addition to a list of all pixel clusters found in the event, there is also an index to the pixel cluster to be used as the interaction time for the event. Most analyses will accept only events with a single unique recoil, in which case the recoil time is unambiguous.

The pair of functions `make_tagger_clusters()` and `make_tagger_photons()` together prepare the information needed by the user code to implement the tagging algorithm. In `make_tagger_clusters()` hits in neighboring channels are compared and, in the case where to hits occur within a coincidence window, the two hits are combined and assigned to the higher tagger channel. The higher tagger channel has a larger bend angle and is likely to be hit first in the case that a single electron passes through more than one counter in the tagger focal plane. The same coincidence window is applied between right and left phototubes on the tagger to form the hit list for a single channel as is subsequently used between neighboring channels to associate tagger clusters. Once an irreducible set of tagger hits has been formed, `make_tagger_photons()` may be called to count coincidences and accidentals. Coincidence and accidental repeat counts are available in the `tagger_photons_t` structure both for the individual tagger channels and the totals across the focal plane. The following code snippet shows how a user might use the information in the `tagger_photons_t`

Table 3: Total number of events of each multiplicity in the entire Radphi data sample after cuts have been applied to select fully-contained forward events and tagging subtraction has been performed.

multiplicity	count
2	3,613,000
3	2,006,000
4	1,277,000
5	494,000
6	177,000
7	45,000

structure to obtain a neutral-tagged spectrum.

```
float weight;
tagger_photons_t *tagging;
tagging = data_getGroup(event,GROUP_TAGGER_PHOTONS,0);
weight = tagging->neutral_coincidences - tagging->neutral_accidentals;
/*
 * calculate some event variable x here
 */
hfill(id,x,0.,weight);
```

4 Total tagged event yields

The subtraction of accidental tags has a large effect on the statistical errors in all measured spectra. In addition to reducing the total number of events via the coincidence requirement, the errors no longer follow the $\sqrt{N}$ formula after subtraction. Precise error estimates may be obtained by turning on explicit computation of errors in the histograms when they are created and then using the weighted filling technique described in the previous section. It turns out that the ratio of true/accidental coincidences is on the order of 1, so the size of a sample can still give an indication of the expected statistical error. The statistical errors of a subtracted spectrum with N events and a signal/background ratio of 1 are roughly equivalent to those of a sample of $N/3$ pure signal events, provided that the spectrum of signal and background are not too different in shape. For the actual analysis the correct error treatment must be applied, but in advance of that some idea of the statistical quality of the Radphi sample can be obtained with this rule of thumb. In Table 3 is shown the total count of events of each multiplicity are shown that remains after the cuts described above are applied to select fully-contained forward neutral triggers, and the tagging subtraction has been performed.

5 Conclusions

A method has been proposed which employs a statistical technique to exploit the photon tagger and forward scintillator information in the Radphi data sample. The method enables the extraction of spectra corresponding to all-neutral triggers in true association with tagged photons in the presence of an arbitrary background from accidental coincidences, at the cost of some statistical precision. The method has been tested by selecting a sample of fully-contained forward events and comparing the total energy of reconstructed showers in the forward calorimeter with the energy of the incident photon as given by the tagger. Excellent agreement has been observed between expected and measured forward energy, once one takes into account the typical 100 MeV that is transfered to the recoil proton. A tail on the low-energy side of the total energy spectrum is attributed primarily to the spread in the energy of the recoil proton (and nucleus).

Examination of the cuts used to select a sample of fully-contained forward events reveal the importance and power of the barrel gamma veto to remove events with significant neutral energy outside the forward calorimeter. In fact, without the barrel gamma veto the peak in the total forward energy spectrum at the tagged energy was completely washed out. Without the bgd information, reconstructing a clean sample of exclusive neutral events in the forward calorimeter would have proved impossible with Radphi. On the other hand, the barrel gamma veto also imposes a steep cost of nearly an order of magnitude in terms of statistics.

It is therefore highly desirable to find and perfect a method for recovering these events by reconstructing barrel photons.

The statistics of a clean sample of all-neutral tagged fully-contained forward events of a given multiplicity have been computed. Sample sizes of order a million are available up to multiplicity five. Software tools for applying the method presented have been checked into the Radphi cvs repository, and verified. Instructions for their use have been provided.

References

[1] R.T. Jones, "A second look at rates in Radphi", presented at Radphi collaboration meeting, Richmond, VA (Sept. 1997) published online at http://zeus.phys.uconn.edu/radphi/radphi-9-97/.

[2] M. Kornicer and R.T. Jones, "LGD shower resolution", Technical Note *radphi-2003-701* (July 2003) published online at http://zeus.phys.uconn.edu/radphi/lgdres-7-2003/.