ELSEVIER

Nuclear Instruments and Methods in Physics Research A 433 (1999) 257–261

NUCLEAR INSTRUMENTS & METHODS IN PHYSICS RESEARCH Section A

www.elsevier.nl/locate/nima

RICH pattern recognition for LHCb

for the LHCb Collaboration

Roger Forty

CERN EP Division, CH-1211, Geneva 23, Switzerland

Abstract

An algorithm for performing pattern recognition in RICH detectors is described. It uses a maximum-likelihood technique, comparing the predicted and observed distributions of detected photoelectrons. The algorithm is able to handle regions of high track density, for which LHCb provides a good example.

1. Introduction

The LHCb experiment has recently been approved, as a dedicated b-physics experiment for the LHC. As the production of b-hadrons at high energy is peaked towards the beam axis, the experiment will study the pp collisions with a forward spectrometer, covering polar angle out to about 300 mrad. Details of the experimental setup can be found elsewhere [1,2]. Particle identification is crucial for the b-physics programme of the experiment, over a wide range of momentum, so it is equipped with two RICH detectors, combining three radiators. The first, RICH-1, has aerogel at its entrance window, and is filled with C_4F_{10} gas. It covers the full acceptance of the spectrometer, identifying particles with low or intermediate momentum ($\sim$ 1–60 GeV/c). The second, RICH-2, covers the small-angle region up to 120 mrad, and as the forward tracks have a harder momentum spectrum it is filled with CF_4 gas, giving π–K separation up to 150 GeV/c. The photodetectors require a spatial granularity of $\sim 2\ \text{mm} \times 2\ \text{mm}$ and should cover a total detection area of about 3 m^2 with the highest possible efficiency; possible technologies are described elsewhere [2,3].

A typical b event is shown in Fig. 1, where the photodetector planes of each RICH detector are displayed side by side, and the Cherenkov rings are superimposed. In RICH-1 the clear rings from the C_4F_{10} gas are visible, along with the less easily identified, large diameter, rings from the aerogel. Rayleigh scattering in the aerogel is also taken into account, giving a background of a few scattered hits per track. Clearly, the pattern-recognition algorithm must be able to handle regions of high track density.

2. Local pattern recognition

The focussing mirrors of the RICH detectors are tilted, to bring the images out of the spectrometer acceptance. The ring images are therefore not perfect circles. Instead of searching for ring-like

E-mail address: roger.forty@cern.ch (R. Forty)

0168-9002/99/$ - see front matter
PII: S0168-9002(99)00310-1

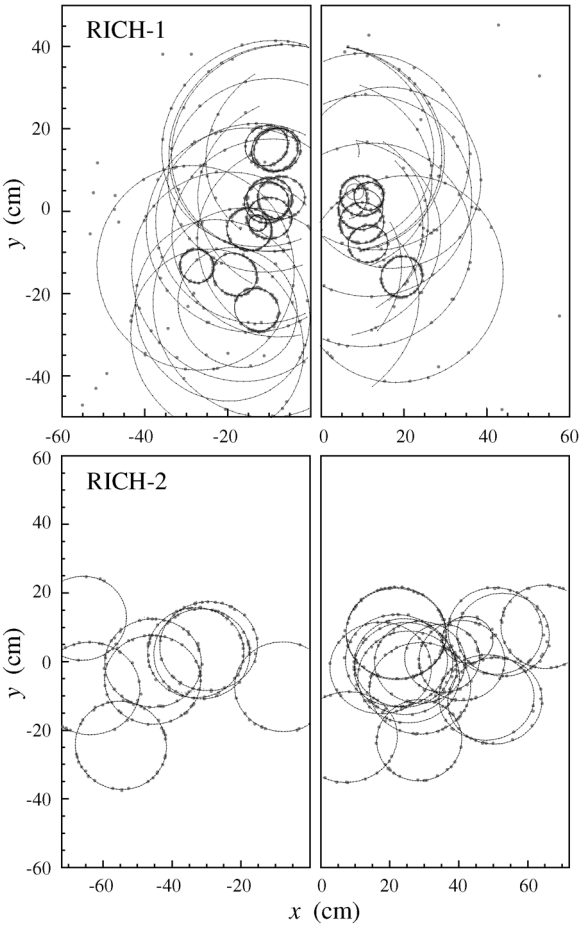

Fig. 1. Event display of a simulated $B^0 \to \pi^+\pi^-$ event. The event generator is PYTHIA, and a full GEANT description of the experiment is used.

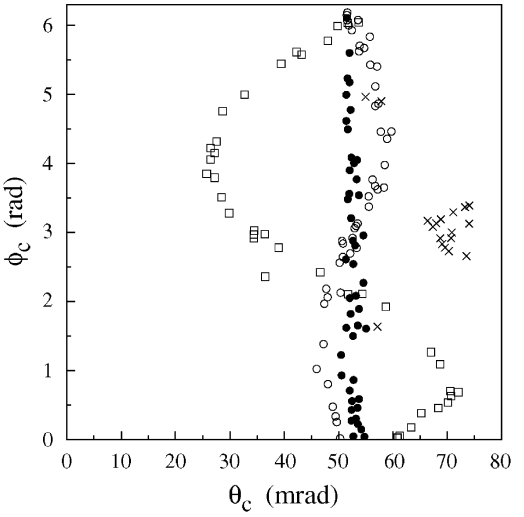

Fig. 2. Reconstructed Cherenkov angles (θ_c,ϕ_c) for the hits in the event of Fig. 1, when calculated assuming that the photons were emitted from the gas radiator of RICH-1, from one of the tracks that gave the overlapping rings at $y \sim 20$ cm. The symbol indicates which track the hit truly originated from, with solid points for the track that has been used for the angle reconstruction.

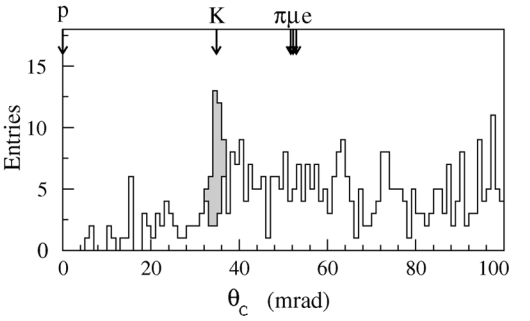

Fig. 3. Distribution of Cherenkov angle for the hits in the event of Fig. 1, calculated relative to a track near $y = 0$. The hits that are truly from that track are shaded, and the expected values of θ_c for the different particle types are indicated.

patterns, the task of pattern recognition is simplified by first reconstructing the Cherenkov angles at emission that correspond to a given association of photon hit and track. For spherical focussing mirrors this can be achieved by the solution of a quartic equation [4], assuming that the track direction has been provided by the dedicated tracking chambers of the experiment. Then the signal hits from a track will be at constant polar angle θ_c and randomly distributed in the azimuthal angle ϕ_c. By performing this reconstruction, distortions of the rings due to the optics are automatically corrected for.

In the case of overlapping rings, the hits from the other tracks give background to the signal of the track under study. This is illustrated in Fig. 2. Projecting such a distribution onto the θ_c-axis, a peak is seen at the true value of the Cherenkov angle for the signal track, but with background from the other tracks. The essence of "local" pattern recognition is to search for such peaks in the θ_c distribution. It is known as local, since each track is treated independently. This also makes it fast,

taking typically 240 ms to analyse an LHCb event [5] on a nominal 1000 MIPS processor, fast enough that it may even be used in the later stages of the trigger.

Instead of assuming continuous θ_c, one can use the fact that there are only five common charged particles in LHCb events, (e, μ, π, K, p). Thus one can instead compare the likelihood for each of these mass hypotheses. This is illustrated in Fig. 3, for a track that was truly a kaon. One is no longer performing pattern recognition and *then* fitting for the Cherenkov angle, but rather doing both in a single step.

3. Global pattern recognition

In regions of high track density, the local approach becomes difficult due to the high background. However, the dominant source of this "background" is the signals from other tracks, so a simultaneous fit to all tracks' hypotheses can take account of it. One could try to fit for θ_c of all tracks, but this gives a complex multi-parameter fit, and it is difficult to incorporate backgrounds. Instead, a transformation is made back to the detector plane. Then, for a given set of particle mass-hypotheses the probability is calculated that a photon signal would be seen in each pixel of the detector, from each track. Comparing the observed number of photoelectrons in the pixels, a likelihood can be calculated. Finally, the set of mass-hypotheses that maximize the likelihood is searched for.

Since the number of observed photoelectrons in a given pixel obeys Poisson statistics, the likelihood is simply given by

$$\mathscr{L} = \prod_{\text{pixels } i} \frac{e^{-v_i} v_i^{n_i}}{n_i!}$$

where n_i is the observed number of photoelectrons, and v_i is the expected signal, in each pixel. $v_i = \sum a_{ij}$, the sum over tracks of the contribution from track j to pixel i, given in turn by

$$a_{ij} = \varepsilon_i N_j^h \iint_{\text{pixel } i} f_h(\theta,\phi)\, d\theta\, d\phi.$$

Here ε_i is the detection efficiency of the pixel, N_j^h is the total number of photons emitted by track j under mass-hypothesis h, and $f_h(\theta,\phi)$ is the probability-density distribution of photons emitted by the track, relative to its direction. For the signal Cherenkov photons this last term is assumed to be a Gaussian distribution in θ_c, and constant in ϕ_c:

$$f_h(\theta,\phi) = \frac{1}{(2\pi)^{3/2}\sigma(\theta)} \exp\left[-\frac{1}{2}\left(\frac{\theta-\theta_h}{\sigma(\theta)}\right)^2 \right]$$

where θ_h is the expected Cherenkov angle for hypothesis h, and $\sigma(\theta)$ is the single-photon resolution. The resolution has contributions from the emission-point uncertainty (independent of θ) and chromatic dispersion in the radiator (proportional to $1/\tan\theta$); the θ dependence is taken into account in the fit. On the detector plane this probability density corresponds to an approximately circular ring, with a radial cross-section that is roughly Gaussian. When rings overlap, the probability of seeing a photoelectron is enhanced. This can be seen in Fig. 4, where the calculated probability density function is shown for a zoomed region of the event of Fig. 1, for a given choice of particle mass-hypotheses. This can be compared with the detected

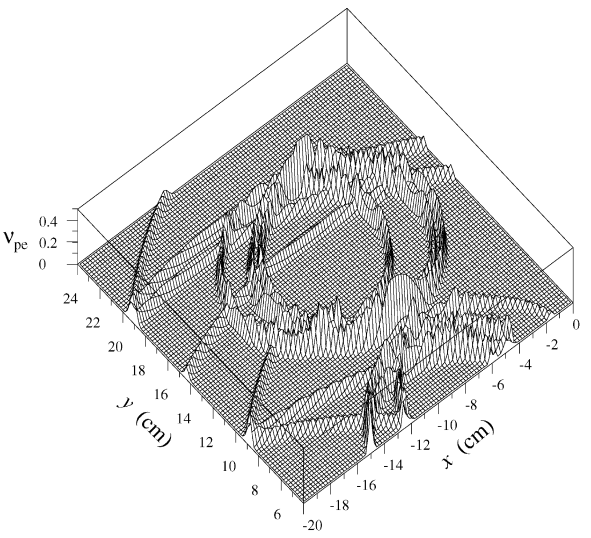

Fig. 4. Expected number of photoelectrons in each pixel, for a region of the RICH-1 detector in the event of Fig. 1.

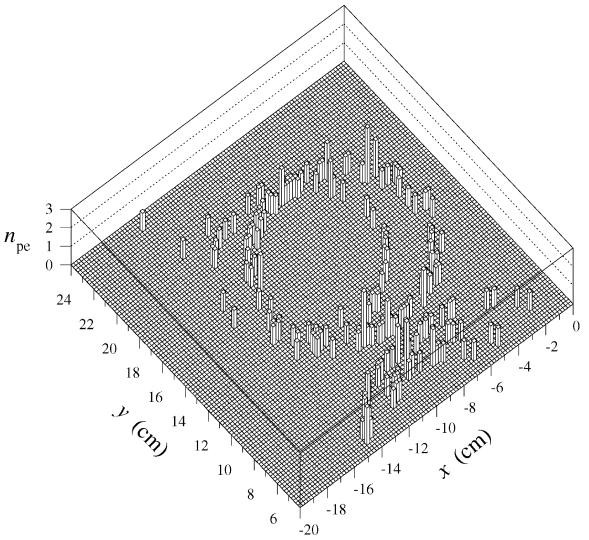

Fig. 5. Observed number of photoelectrons in each pixel, for the same region of the detector as Fig. 4.

photoelectrons over that same region, shown in Fig. 5, where each bin represents a single pixel. The likelihood that the observed hits agree with the expected distribution is calculated from a comparison of these two distributions.

The set of mass-hypotheses that maximise that likelihood is then searched for. The search strategy starts by assuming that all tracks are pions (the most numerous particle type in LHCb events). Then the likelihood is recomputed changing each track hypothesis in turn (to e, μ, K, p), leaving the other tracks unchanged. The change that gives the largest increase in the likelihood is then selected, and the procedure is iterated until no further improvement is found. Typically 25 iterations are required per LHCb event.

The integral over the pixels for the contribution from each track, a_{ij}, simplifies when the pixel is small compared to the resolution, as then the probability density does not vary strongly over the pixel and can be factorised:

$$\iint_{\text{pixel } i} f_h(\theta,\phi)\,\mathrm{d}\theta\,\mathrm{d}\phi \approx f_h(\theta_i,\phi_i)\iint_{\text{pixel } i}\mathrm{d}\theta\,\mathrm{d}\phi$$

$$\approx f_h(\theta_i,\phi_i)\frac{4A}{R^2\theta_i}$$

Table 1
Results from the global pattern recognition applied to 500 events

Rec	True particle type						$\mathscr{P}$
	e	μ	π	K	p	X	
e	6233	7	328				0.95
μ	8	224	554			31	0.27
π	5	10	13114	1		8	0.99
K		1	39	1083		11	0.96
p	1		4	1	427	1	0.98
X	3	8	197	27		3990	0.94
ε	0.99	0.90	0.92	0.97	1.00	0.99	

where A is the pixel area and R is the radius of curvature of the focussing mirror. The log-likelihood can be rearranged in the following form:

$$\ln\mathscr{L} = -\sum_{\text{tracks } j}\mu_j + \sum_{\text{hits } i}\left(n_i \ln\sum_{\text{tracks } j} a_{ij}\right) + C$$

where $C = \sum \ln n_i!$ is a constant and can therefore be ignored when comparing the likelihood of different hypotheses, and $\mu_j = \sum a_{ij}$ is the total number of expected photoelectrons from track j (under a given mass hypothesis), which needs only be calculated once per event. The remaining summation, which must be performed repeatedly (for each combination of track mass-hypotheses) is only over the *hit* pixels, and since the pixel occupancy is rather low (typically 1% or less) the algorithm is reasonably fast: it takes about 1.2 s per LHCb event (on the nominal 1000 MIPS processor).

This formalism can be straightforwardly generalised to include background sources of photodetector hits, either pixel-related (such as electronic noise) or track-related (such as Rayleigh-scattered photons from the aerogel). The performance has been studied by applying the algorithm to all tracks in 500 simulated b events. The result is shown in Table 1, where there is one entry for each track: the column indicates the true particle type (X if below threshold in all radiators) and the row indicates the reconstructed particle type. Correctly identified tracks lie on the diagonal of the table, and as can be seen the performance is good. It can be characterised in terms of the purity $\mathscr{P}$, the fraction of tracks

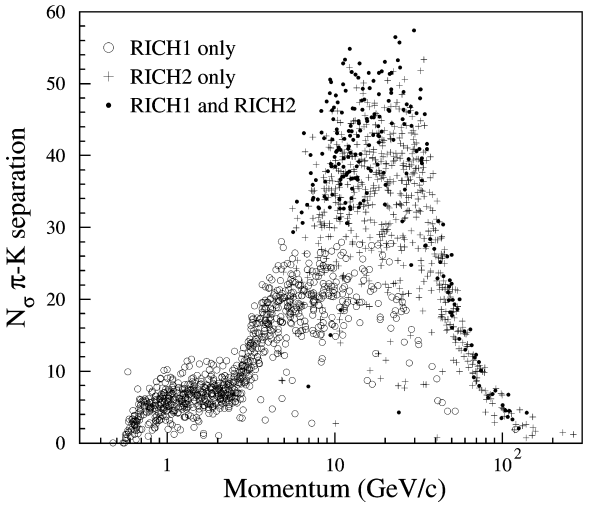

Fig. 6. π–K separation (in σ) as a function of momentum, for all true pions in 100 b events. The symbol indicates which detector(s) the pion passed through.

identified as a given particle type that are correctly identified, and the efficiency ε, the fraction of tracks truly of a given particle type that are correctly identified. The purities and efficiencies are all in the region 90–100% (except the muon purity, which suffers from the much more numerous pions, close in mass, but for which LHCb has a dedicated muon identification system).

The method is found to be relatively insensitive to random noise in the photodetectors, until the probability of a pixel firing is increased beyond a few percent: this follows from the use of the tracking information to predict where the signal photoelectrons are expected. Instead of simply choosing the maximum-likelihood solution, one can vary the cut on the separation between different hypotheses, expressed in terms of Gaussian sigma, $N_\sigma = \sqrt{2\Delta \ln \mathscr{L}}$. The π–K separation is shown in Fig. 6. As can be seen, good separation is achieved over the region 1–150 GeV/c in momentum, and thus the particle-identification requirements of LHCb can be satisfied.

References

[1] LHCb Technical Proposal, CERN/LHCC 98-4.
[2] A. Go, Nucl. Instr. and Meth. A 433 (1999) 153.
[3] S. Easo, Nucl. Instr. and Meth. A 433 (1999) 159.
[4] R. Forty, O. Schneider, LHCb/98-040.
[5] A. Schöning, LHCb/97-018.