Available online at www.sciencedirect.com
ScienceDirect
ELSEVIER
Nuclear Instruments and Methods in Physics Research A 571 (2007) 527–561
NUCLEAR INSTRUMENTS & METHODS IN PHYSICS RESEARCH Section A
www.elsevier.com/locate/nima

The HARP detector at the CERN PS

M.G. Catanesi[a], M.T. Muciaccia[a], E. Radicioni[a], S. Simone[a], R. Edgecock[b], M. Ellis[b,1], S. Robbins[b,2], F.J.P. Soler[b,3], C. Gößling[c], M. Mass[c], S. Bunyatov[d], A. Chukanov[d], O. Klimov[d], I. Krasin[d], A. Krasnoperov[d], D. Kustov[d], B. Popov[d,4], V. Serdiouk[d], V. Tereshchenko[d], V. Carassiti[e], E. Di Capua[e], F. Evangelisti[e], G. Vidal-Sitjes[e,1,5], A. Artamonov[f,6], P. Arce[f,7], R. Brocard[f], G. Decreuse[f], B. Friend[f], S. Giani[f], S. Gilardoni[f,5], P. Gorbunov[f,6], A. Grant[f], A. Grossheim[f,5], P. Gruber[f,5], V. Ivanchenko[f,8], J.-C. Legrand[f], A. Kayis-Topaksu[f,9], J. Panman[f], I. Papadopoulos[f], J. Pasternak[f,5], E. Tcherniaev[f], I. Tsukerman[f,6], R. van der Vlugt[f], R. Veenhof[f], C. Wiebusch[f,10], P. Zucchelli[f,10], A. Blondel[g], S. Borghi[g], M. Campanelli[g], A. Cervera-Villanueva[g], M.C. Morone[g,11], G. Prior[g,5,12], R. Schroeter[g], I. Kato[h,13], U. Gastaldi[i], G.B. Mills[j,14], J.S. Graulich[k,15], G. Grégoire[k], M. Bonesini[l], F. Chignoli[l], F. Ferri[l], F. Paleari[l], M. Kirsanov[m], V. Postoev[m], A. Bagulya[n], V. Grichine[n], N. Polukhina[n], V. Palladino[o], L. Coney[p,14], D. Schmitz[p,14], G. Barr[q], A. De Santo[q,16], C. Pattison[q], K. Zuber[q,17], G. Barichello[r], F. Bobisut[r], D. Gibin[r], A. Guglielmi[r], M. Laveder[r], A. Menegolli[r], M. Mezzetto[r], A. Pepato[r], J. Dumarchez[s], S. Troquereau[s], F. Vannucci[s], U. Dore[t], A. Iaciofano[u], M. Lobello[u], F. Marinilli[u], D. Orestano[u], D. Panayotov[u,18], M. Pasquali[u], F. Pastore[u], A. Tonazzo[u], L. Tortora[u], C. Booth[v,*], C. Buttar[v,3], P. Hodgson[v], L. Howlett[v], R. Nicholson[v], M. Bogomilov[w], K. Burin[w], M. Chizhov[w], D. Kolev[w], P. Petev[w], I. Rusinov[w], R. Tsenov[w], S. Piperov[x], P. Temnikov[x], M. Apollonio[y,19], P. Chimenti[y], G. Giannini[y], G. Santin[y,20], J. Burguet-Castell[z], J.J. Gómez-Cadenas[z], P. Novella[z], M. Sorel[z], A. Tornero[z]

[a] *Università degli Studi e Sezione INFN, Bari, Italy*
[b] *Rutherford Appleton Laboratory, Chilton, Didcot, UK*
[c] *Institut für Physik, Universität Dortmund, Germany*
[d] *Joint Institute for Nuclear Research, JINR Dubna, Russia*
[e] *Università degli Studi e Sezione INFN, Ferrara, Italy*
[f] *CERN, Geneva, Switzerland*
[g] *Section de Physique, Université de Genève, Switzerland*
[h] *University of Kyoto, Japan*
[i] *Laboratori Nazionali di Legnaro dell' INFN, Legnaro, Italy*
[j] *Los Alamos National Laboratory, Los Alamos, USA*
[k] *Institut de Physique Nucléaire, UCL, Louvain-la-Neuve, Belgium*
[l] *Università degli Studi e Sezione INFN, Milano, Italy*
[m] *Institute for Nuclear Research, Moscow, Russia*
[n] *P. N. Lebedev Institute of Physics (FIAN), Russian Academy of Sciences, Moscow, Russia*
[o] *Università "Federico II" e Sezione INFN, Napoli, Italy*
[p] *Columbia University, New York, USA*
[q] *Nuclear and Astrophysics Laboratory, University of Oxford, UK*
[r] *Università degli Studi e Sezione INFN, Padova, Italy*
[s] *LPNHE, Universités de Paris VI et VII, Paris, France*
[t] *Università "La Sapienza" e Sezione INFN Roma I, Roma, Italy*

*Corresponding author. Tel.: +44 114 2223541; fax: +44 114 222 3555.
E-mail address: C.Booth@Sheffield.ac.uk (C. Booth).

0168-9002/$ - see front matter
doi:10.1016/j.nima.2006.08.132

[u]Università degli Studi e Sezione INFN Roma III, Roma, Italy
[v]Department of Physics, University of Sheffield, UK
[w]Faculty of Physics, St. Kliment Ohridski University of Sofia, Sofia, Bulgaria
[x]Institute for Nuclear Research and Nuclear Energy, Academy of Sciences, Sofia, Bulgaria
[y]Università degli Studi e Sezione INFN, Trieste, Italy
[z]Instituto de Física Corpuscular, IFIC, CSIC and Universidad de Valencia, Spain

HARP collaboration
Available online 10 November 2006

Abstract

HARP is a high-statistics, large solid angle experiment to measure hadron production using proton and pion beams with momenta between 1.5 and 15 GeV/c impinging on many different solid and liquid targets from low to high Z. The experiment, located in the T9 beam of the CERN PS, took data in 2001 and 2002. For the measurement of momenta of produced particles and for the identification of particle types, the experiment includes a large-angle spectrometer, based on a Time Projection Chamber and a system of Resistive Plate Chambers, and a forward spectrometer equipped with a set of large drift chambers, a threshold Cherenkov detector, a time-of-flight wall and an electromagnetic calorimeter. The large angle system uses a solenoidal magnet, while the forward spectrometer is based on a dipole magnet. Redundancy in particle identification has been sought, to enable the cross-calibration of efficiencies and to obtain a few percent overall accuracy in the cross-section measurements. Detector construction, operation and initial physics performances are reported. In addition, the full chain for data recording and analysis, from trigger to the software framework, is described.

PACS: 95.55.Vj

Keywords: Particle detectors

Contents

[1]Now at Imperial College, University of London, UK.
[2]Jointly appointed by Nuclear and Astrophysics Laboratory, University of Oxford, UK.
[3]Now at University of Glasgow, UK.
[4]Also supported by LPNHE, Universités de Paris VI et VII, Paris, France.
[5]Supported by the CERN Doctoral Student Programme.
[6]Permanently at ITEP, Moscow, Russian Federation.
[7]Permanently at Instituto de Física de Cantabria, Univ. de Cantabria, Santander, Spain.
[8]On leave of absence from the Budker Institute for Nuclear Physics, Novosibirsk, Russia.
[9]On leave of absence from Çukurova University, Adana, Turkey.
[10]On leave of absence from INFN, Sezione di Ferrara, Italy; now at SpinX Technologies, Geneva, Switzerland.
[11]Now at Università di Roma Tor Vergata, Italy.
[12]Now at Lawrence Berkeley National Laboratory, Berkeley, USA.
[13]K2K Collaboration.
[14]MiniBooNE Collaboration.
[15]Now at Section de Physique, Université de Genève, Switzerland, Switzerland.
[16]Now at Royal Holloway College, University of London, UK.
[17]Now at University of Sussex, Brighton, UK.
[18]On leave of absence from St. Kliment Ohridski University, Sofia, Bulgaria.
[19]Now at University of Oxford, UK.
[20]Now at ESA/ESTEC, Noordwijk, The Netherlands.

1. Introduction

Precise and comprehensive studies of secondaries from the interaction of protons and charged pions with nuclei, at incident particle momenta of a few GeV/c, are relevant in several areas of particle physics. Early studies exploring only small fractions of the phase space were hampered by large uncertainties in acceptance corrections, and were typically restricted to hydrogen and beryllium nuclei (see Refs. [1,2]). In recent years, interest in such studies was revived for three reasons.

The first arose in the studies required for an economic design of a Neutrino Factory [3–6]. The front stage of a Neutrino Factory is a high-intensity proton driver, which produces secondary charged pions from protons interacting with a high-Z nuclear target. In order to make optimum use of the unprecedented high beam power, a precise knowledge of pion yields is of great interest.

The second interest arose from the observation of substantial $\nu_\mu \to \nu_\tau$ oscillations in atmospheric neutrinos, while ν_e appeared not to oscillate into other neutrino flavours [7]. The initial proof of oscillations, and later the precise determination of the oscillation parameters, would have greatly benefited from a more precise knowledge of the atmospheric neutrino flux. Still today, absolute atmospheric neutrino fluxes are known with a $\sim 30\%$ uncertainty, and the ratio of ν_e and ν_μ of atmospheric neutrinos with a $\sim 7\%$ uncertainty. There are uncertainties from the incident flux of cosmic-ray particles and from the neutrino–nucleon cross-sections at low energy. However, the largest uncertainty stems from the lack of precise knowledge of the interactions of protons and pions with nitrogen and oxygen nuclei in the Earth's atmosphere [8–12].

In addition, strong interest was shown to determine the characteristics of secondary pion production from 8.9 GeV/c protons on beryllium (MiniBooNE experiment at FNAL [13]) and of 12.9 GeV/c protons on aluminium (K2K experiment at KEK [14]).

The third interest comes from the need for precise input data for the simulations of interactions of hadrons with all sorts of matter.

For these reasons, the PS214 (HARP) hadroproduction experiment [15] was constructed and took data at the CERN Proton Synchrotron (PS). HARP carried out a programme of measurements of secondary hadron production, over the full solid angle, produced on thin and thick nuclear targets by beams of protons and pions with momenta in the range 1.5–15 GeV/c .

HARP was designed to have a 4π acceptance. This was accomplished by the combination of a Time Projection Chamber (TPC) in which the target was embedded, with a forward spectrometer featuring a dipole magnet for momentum measurement. The identification of primary

particles in the beam was important, as was the identification of secondaries, both at large angles in the TPC, and at small angles in the forward region. Every effort was made to re-use existing equipment. However major parts of the experiment, notably the TPC, the Resistive Plate Chambers (RPC), the Cherenkov counter and the time-of-flight wall (TOFW), had to be designed and constructed specifically for HARP. The experiment was approved by CERN's Research Board in February 2000 and took data in the years 2001 and 2002. The experimental setup was decommissioned at the end of 2002.

2. Overall layout of the experiment

The experimental layout consists of four main functional parts:

(1) The beamline along with the detection and identification of incoming beam particles;
(2) The trigger detectors;
(3) The large-angle spectrometer housing the target;
(4) The forward spectrometer.

The beamline selects secondary particles with positive or negative charge emerging from a primary target in the extracted proton beam from the PS accelerator. It covers the momentum range between 1.5 and 15 GeV/c . Particle identification in the beamline is provided by two gas Cherenkov detectors (Beam Cherenkov A and B: BCA, BCB) and a pair of time-of-flight counters (TOF-A, TOF-B). Four multi-wire proportional chambers (MWPC) measure the position and direction of the beam particles upstream of the target. Several scintillation counters serve to trigger on single incoming beam particles. The beamline and its instrumentation are described in Section 3. The trigger detectors are described in Section 4.

The global layout of the HARP experiment is shown in Fig. 1. It covers a total length of 13.5 m along the beam direction. The large-angle spectrometer is housed in a solenoid magnet. It comprises the target and several concentric layers of cylindrical detectors. At the most inner radius the target is surrounded by the Inner Trigger Cylinder (ITC), six layers of scintillating fibres serving as an interaction trigger for large-angle particles. Most of the radial space of the solenoid magnet is occupied by the TPC. It provides track, momentum and vertex measurements for all outgoing charged particles in the angular range from 20° to 160° with respect to the beam axis. In addition, it provides particle identification by recording the particle's energy loss in the gas (dE/dx). The particle identification

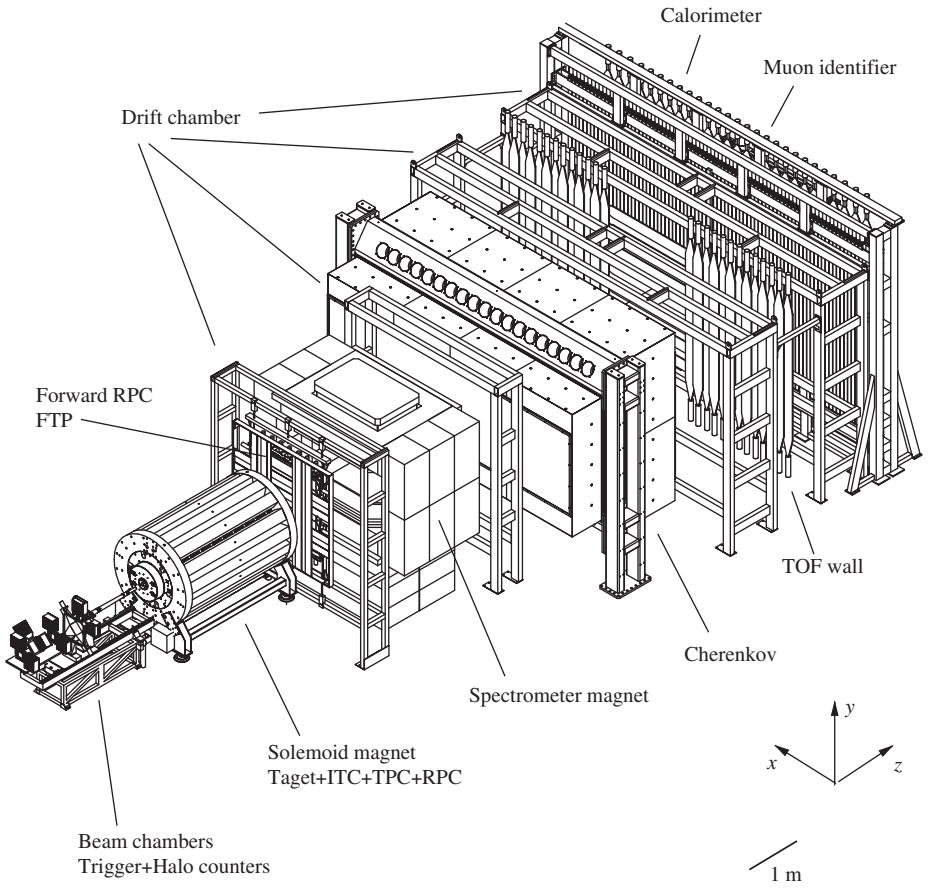

Fig. 1. Overall mechanical layout of the HARP detector. The different sub-detectors are shown. The target is inserted inside the TPC. The convention used for the coordinate system is drawn in the figure.

capabilities of the TPC detector are complemented by a set of multi-gap RPC serving as time-of-flight detectors and surrounding the TPC. The targets are described in Section 2.1, the solenoid magnet in Section 5.1.1, the TPC in Section 5.1 and the RPC detectors in Section 5.2.

The forward spectrometer measures particles produced in the forward direction at angles up to $\sim 14.3^\circ$. A double plane of scintillation counters (forward trigger plane, FTP) placed downstream of the solenoid magnet provides the forward interaction trigger. The tracks and momenta of the outgoing particles are measured in a dipole magnet and two sets of multi-wire drift chambers. Several additional sets of drift chambers are placed further downstream to follow the particle tracks between the various particle identification detectors. Discrimination between forward going protons and pions is provided by a threshold Cherenkov counter and a time-of-flight detector. The Cherenkov counter provides particle identification at high momenta and the (TOFW) detector at low momenta. The identification of photons is achieved through an iron absorber slab followed by a set of drift chambers and a lead-fibre calorimeter. Electrons are identified and their energy is measured in the calorimeter. The calorimeter also provides further particle identification for hadrons. The muon contamination of the beam is measured in a muon identifier consisting of thick iron absorbers and scintillation counters. The outgoing beam passes through the forward spectrometer. The TPC detector, the forward RPC time-of-flight detector and the FTP contain inactive areas for the outgoing beam. In the case of all other forward detectors, the outgoing beam passes through active detector areas. The dipole magnet is described in Section 6.1. The various detectors of the forward spectrometer are described in Section 6. The origin of the HARP coordinate system is placed in the centre of the target. The z-coordinate points in the direction of the beam, the y-coordinate upward, and the x-coordinate to the left when looking in the direction of the beam.

2.1. Targets

Both solid and liquid targets have been exposed to the beam. Target materials range from low to high atomic weight. In order to study the effects of secondary interactions, targets of different thicknesses have been used. In addition empty target runs were taken in order to measure background from interactions in materials surrounding the target.

2.1.1. Solid targets

Seven elements (Be, C, Al, Cu, Sn, Ta, Pb) were chosen to provide solid targets, covering the atomic number range from 4 to 82. The materials are all of very high purity, ranging from 99.95% to 99.999%, with details of the small amount of contaminants provided by the supplier. In order to observe the produced particles with the minimum amount of scattering or re-interaction, most measurements are made with targets which are thin discs. For the early runs, these are nominally 2% of a nuclear interaction length λ_{int} thick. To increase the fraction of triggers with beam interactions in the target, this thickness has been increased to 5% in 2002. Because of the size of the beam, the targets are all 30 mm in diameter. Particles produced at 90° with respect to the beam direction would therefore see a significant amount of material. To ensure that corrections for secondary interactions are satisfactorily modelled and corrected for, data have been also taken with thick targets, a full interaction length long (Be, C, Al, Cu, Ta, Pb).

In addition to the normal 'physics targets', special targets are used for calibration and alignment. These include a thin copper disc inclined at 45° to the beam, to provide further checks on particle production and re-interaction at large angles. Data have been also taken with replica targets from two accelerator neutrino experiments, MiniBooNE[21] [13] and K2K[22] [14], with the intention of reducing the systematic uncertainties in the calculated neutrino fluxes. In order to provide for tracking of the produced particles over as much as possible of the complete solid angle, the targets are mounted 500 mm inside the TPC, at the end of a long, hollow arm. The target support is constructed in two pieces. Each individual target is mounted at the upstream end of its own cylindrical support cap. This cap is made of aluminium, and the downstream 168 mm of this are only 0.3 mm thick. Thus, tracks produced in the forward hemisphere pass through no support material, while those with angle θ between 90° and 172°, measured with respect to the beam axis, pass through approximately $0.3\,\text{mm}/\sin(\theta)$ of aluminium. Targets less than 19 mm thick are mounted with their upstream face at the origin of the HARP coordinate system; thicker targets are centred on this point. The full interaction length targets also have a support collar, of 0.3 mm thick aluminium, near their upstream end. In addition, special data at $1.5\,\text{GeV}/c$ have been taken with water targets 60 and 600 mm long, with a 30 mm diameter, mounted in plexiglass tubes of 2 mm wall thickness ('water targets'). Measurements with these targets could improve the precision of the LSND background corrections to the observed $\overline{\nu}_\mu \to \overline{\nu}_e$ oscillation signal.

2.1.2. Cryogenic targets

Data have also been taken with cryogenic targets of H_2 ($L = 60$ and 180 mm), D_2 ($L = 60$ mm), N_2 ($L = 60$ mm) and O_2 ($L = 60$ mm). Data taking with H_2 and D_2 targets is fundamental to distinguish nucleon-induced contributions to the cross-sections from nuclear effects. In addition, H_2 and D_2 targets allow for comparison with previous hadron production data. Data collection with O_2 and N_2

[21]The MiniBooNE replica targets are made of Beryllium and have a three-wing shape, with the wings pointing at 120° of each other. They are 406 and 163 mm long, with a 29 mm diameter.

[22]The K2K replica targets are made of an alloy containing 98% Aluminium and are 650 and 200 mm long, with a 29.9 mm diameter.

targets yields direct information for the precise prediction of atmospheric neutrino fluxes.

The target envelope has a diameter of 35 mm and is made of 125 μm thick mylar foil. Its entrance window has a diameter of 20 mm, thus defining the effective diameter of the cryogenic target. The target is surrounded by aluminized mylar superinsulation layers and placed inside a cylindrical aluminium cryostat, 60 mm in diameter and 0.6 mm wall thickness, of similar shape to the solid target arm. The cryostat arm has mylar windows, each 250 μm thick, for the incoming and outgoing beams.

Because of differences in boiling temperatures, two cryogenic target setups were built, one for the H_2 and D_2 targets and one for the N_2 and O_2 targets. Buffer volumes and precise operating pressures are adapted to each of the different target setups. The operating pressure of the target defines the exact boiling temperature and target density and is monitored continuously through the Detector Control System (DCS). The typical target cool-down time is 6 h, whereas warming up to room temperature takes 2 h. Once cooled down, the targets remain operational for many days without any need for intervention.

Metrology measurements of the precise shape of the cryogenic target envelopes were performed at various gas pressures. As a result, the target lengths are known to better than 0.5%.

3. Particle beam

The experiment was installed in the T9 beam of the East Area of the CERN PS. This beamline is derived from a resonant extracted beam from the PS and sent towards a splitter which allows beams to be sent either to the North or South branches of the East Area. The North branch puts 24 GeV/c protons onto a primary target, usually made of aluminium, which is seen simultaneously by three beamlines: T9, T10 and T11.

The beam is extracted from the PS in a single 400 ms long spill of 2×10^{11} protons every 2.4 s onto the primary target. The PS accelerator operation is organized in 'supercycles', with a duration of a multiple of 1.2 s. During HARP data taking the length of the supercycle was either 16.8 or 19.2 s. Cycles of different length (in multiples of 1.2 s) perform the different functions of the accelerator. Up to three extractions during the 16.8 or 19.2 s long PS supercycle are sent to HARP.

The flux of protons and pions produced in the primary target is about 10^5 per spill, integrated over the full momentum range. A momentum selected fraction of secondaries is then transported towards the HARP experiment.

It should be noted that the T9 beamline [16] was originally built for test beams. It is limited to a maximum momentum of 15 GeV/c, because of the dipole bending power of the horizontal magnets. A major modification of the beamline was made to focus the beam on the experimental target in HARP. When taking data with the thick targets the total particle flux had to be kept below 2000 particles per 400 ms spill, to limit overlays in the TPC detector. During normal running with thin targets the particle flux was kept below 10 000–13 000 particles per spill for the same reason. The momentum resolution of the beam is about 0.24% in the momentum range used. At +3 GeV/c, a typical beam composition is ~70% of π, ~22% of p with the remaining split between minority components, such as e, K, and other particles.

3.1. Beam instrumentation

The beam instrumentation upstream of the target includes a pair of gas threshold Cherenkov counters (BCA and BCB), two scintillation halo counters (HALO-A and HALO-B), a beam time-of-flight system (TOF-A and TOF-B), a set of four MWPC and two special scintillation counters (BS and TDS). Their role is twofold: to provide the detection, tracking and identification of incoming beam particles (BC, TOF and MWPC) and to generate primary signals for the trigger decision (TOF-B, BS, TDS, HALO counters and optionally BCs).

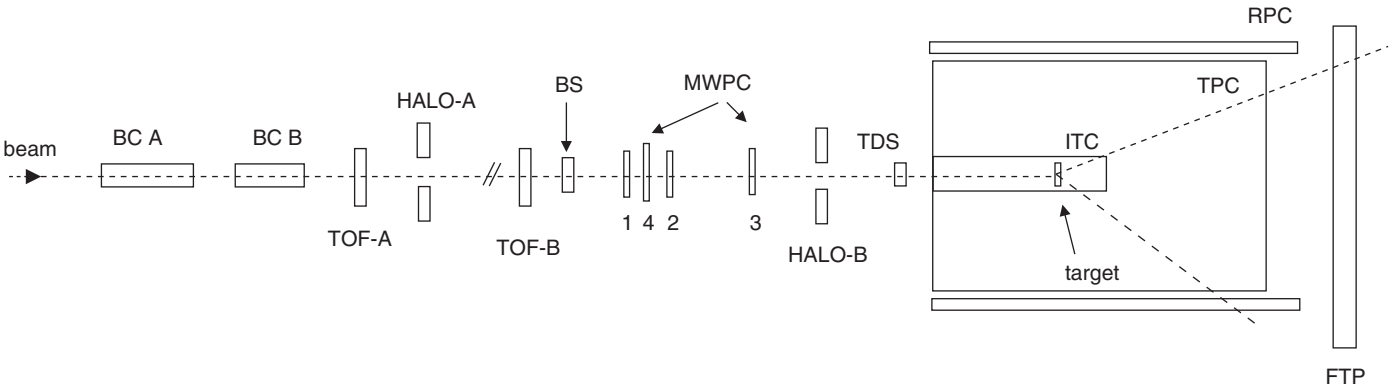

Fig. 2. Schematic view of the arrangement of all trigger and beam equipment: two gas threshold Cherenkov counters (BCA and BCB), two scintillation halo counters (HALO-A and HALO-B), a beam time-of-flight system (TOF-A and TOF-B), a set of four multi-wire proportional chambers and two special scintillation counters (BS and TDS). Detailed descriptions are given in the text. The MWPCs are numbered: 1, 4, 2, 3.

A schematic view of the relative positions of all beam and trigger detectors is shown in Fig. 2.

The task of beam particle identification is shared between the beam Cherenkov counters and the beam time-of-flight system. Below 3 GeV/c , the beam Cherenkovs are used to tag $e^{\pm}$. The TOFs are capable of resolving the more massive beam particles such as pions and protons. At 5 GeV/c the π/p-separation is done jointly by the beam TOF and one of the Cherenkovs (usually BCB), while the other Cherenkov (BCA) is used to tag $e^{\pm}$. At higher momenta the $e^{\pm}$-contamination drops to below 1% and at the same time the beam TOF system becomes unable to clearly separate pions and protons efficiently. The task of π(/K)/p-separation is performed by the beam Cherenkov detectors.

The tracking of beam particles is performed by the beam MWPCs located near the target. These chambers are accurately aligned with respect to the nominal HARP coordinate system. In practice, all other HARP detectors, primarily the drift chambers (Section 6.2), are aligned by reconstructing the non-interacting beam particles whose tracks are extrapolated from the beam MWPC measurements.

Table 1
Basic pressure settings (in bar absolute) for the beam Cherenkov counters

Beam momentum	BCA	BCB
$\leqslant$ 3 GeV/c	0.0 (-)	1.05 (e)
5 GeV/c	0.60 (e)	2.50 (π, e)
$\geqslant$ 8 GeV/c	1.25 (π)	1.50 (π)
$\geqslant$12 GeV/c (optional, for K-tagging)	3–3.5 (π, K)	3.5 (π, K)

The tagged particle types are indicated in parentheses. At 3 GeV/c , BCA is evacuated in order to reduce multiple scattering of the beam.

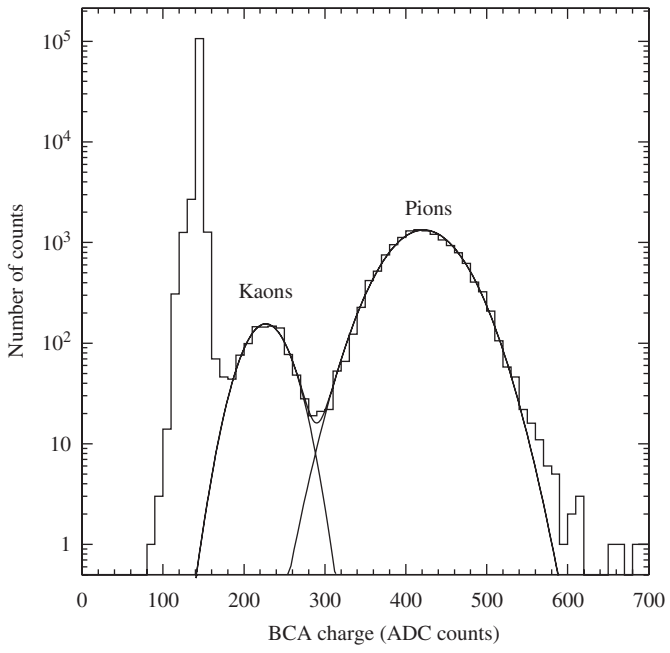

Fig. 3. The pulse-height spectrum from the BCA Cherenkov counter for the positive 12.9 GeV/c beam.

The gas threshold Cherenkov counters BCA (length 5000 mm) and BCB (length 3000 mm) belong to the standard equipment of the PS T9 beamline [16]. These counters were specially refurbished for HARP, with the installation of new mirrors and photomultipliers (PMT). The counters are filled with nitrogen at different pressures depending on the beam energy and the type of particles to be tagged (see Table 1). The measured tagging efficiency of both counters for pions and $e^{\pm}$ is close to 100%. The limited K-tagging capability at $\geqslant$12 GeV/c (Fig. 3) is sufficient given the relatively small fraction of kaons in the beam ($<$1%).[23]

The beam time-of-flight system with a 21.4 m-long base is used for beam particle identification up to 5 GeV/c . In addition, the beam momentum can be accurately determined by comparing the measured velocities of pions, protons and deuterons of the same beam, as illustrated in Fig. 4. Two identical scintillation hodoscopes, TOF-A and TOF-B [17] provide the average time-of-flight resolution of $\approx$ 100 ps. Each counter consists of eight 5 mm thick scintillator strips read out at both ends by Hamamatsu H1949 PMTs. The different strip widths are chosen to roughly equalize the count rates at different distances from the beam axis. The logical OR of all TOF-B channels is used in the trigger (see Section 4).

Four beam MWPCs are used for the off-line reconstruction of the beam particles, as well as for real-time beam tuning and beam-quality monitoring. They measure the beam particle position and the angle at the target with an accuracy of $<$1 mm and $<$0.2 mrad per projection. They are located close to the target in order to reduce multiple scattering effects between the track measurement and the target. The HARP MWPCs are refurbished standard CERN beam chambers of types XWCA (MWPC 1, 2 and 3) and XWCM (MWPC 4) [18]. Each chamber has two adjacent 10 mm-wide detection gaps with anode sense wire planes perpendicular to each other. The chambers 1–3 are used for particle tracking. They have a wire pitch of 1 mm and the sense wires are parallel to the x- and y-axis of the HARP coordinate system. The larger chamber 4 is rotated by 45° and has a 4 mm wire pitch. Its purpose is to resolve the tracking ambiguity in 2-track events and to detect beam halo particles. The MWPCs are equipped with a HARP-designed front-end read-out (a preamplifier and a discriminator/ECL driver) for each wire. The gas is a 1:1 mixture of Ar and CO_2 with an addition of 0.5% of water vapour to prevent ageing of the wires. The HV settings are chosen so as to maintain the efficiency above 99% for each chamber.

A beam muon identifier is placed at the downstream end of the HARP detector. Beam muons need to be identified

[23]The K-tagging option was used mainly during exposures with K2K replica targets at 12.9 GeV/c , for which a very pure proton beam was requested.

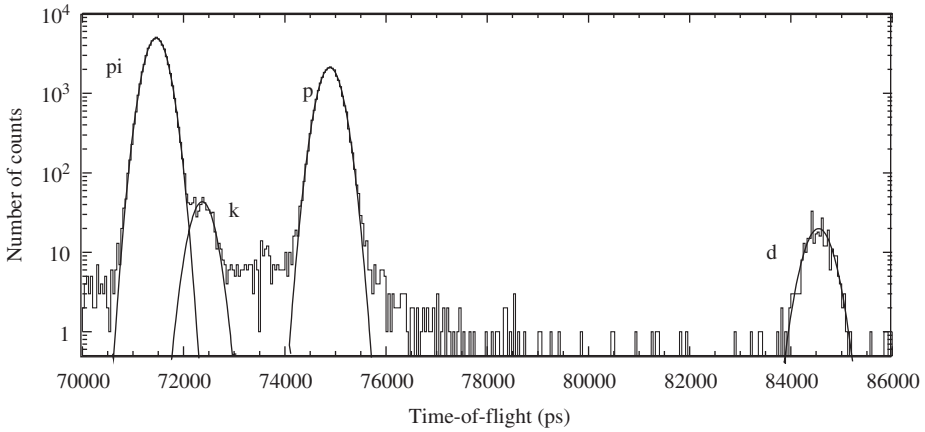

Fig. 4. Example for beam particle identification with time-of-flight at nominal beam momentum of 3 GeV/c . Electrons have been rejected with the Cherenkov counters.

since they would lead to a wrong interaction cross-section when accounted for as pions. At high momenta the most effective way to tag them is through non-showering in a sufficiently deep, longitudinally segmented calorimeter. At low beam momenta, beam muons, which generally result from kaon and pion decays along the beamline, are hardly discriminated from pions but are strongly bent by the magnetic dipole field. The beam muon identifier is a 1.40 m wide structure placed in the horizontal direction asymmetrically with respect to the beamline, in order to intercept all the beam muons which are horizontally deflected by the spectrometer magnet. It consist of a passive 0.40 m layer of iron followed by an iron-scintillator sandwich with five planes of six scintillator each, read out at both sides, giving a total of 6.4 λ_{int}.

4. Trigger detectors

The detectors used in the trigger can be divided into two groups: beam particles are detected upstream of the target, whereas the detection of secondaries emerging from the target relies on downstream detectors and detectors surrounding the target. The beam detectors have been described in Section 3.1. A schematic view of the relative position of all trigger and beam equipment is shown in Fig. 2.

4.1. Upstream trigger counters

Two HALO counters, made of one (HALO-A) and two (HALO-B) scintillator slabs each with a central hole, serve to veto events in which the beam particle is accompanied by a second particle in the halo of the beam. The primary purpose of HALO-A is to veto all particles which did not pass through the beam Cherenkov counters. It is read out by two Philips XP2020 PMT. The diameter of its hole is 9 cm. The final cleanup of the beam is done by HALO-B as near to the target as possible. Its hole diameter of 3 cm is consistent with the target diameter. To convert photons, a 11 mm ($\sim 2X_0$) thick lead plate is sandwiched between its two scintillator slabs. To achieve a maximal veto-efficiency, each halo counter enters the trigger with a logic OR of its channels.

The beam scintillator (BS) starts the decision logic of the trigger system. It is located a few centimeters downstream of TOF-B. In coincidence with a TOF-B hit it represents the lowest-level trigger (strobe). As a consequence, the timing of all trigger signals is given by the original BS signal, i.e. BS provides the timing reference for all ADC gates, TDC start and stop signals. BS consists of a single scintillator slab of 5 mm thickness and a sensitive area of 80 mm × 80 mm centred on the nominal beam axis. It is read out by a single Philips XP2020 PMT tube.

The Target Defining Scintillator (TDS) is a scintillator disc of 20 mm diameter and 5 mm thickness which is viewed by four PMT tubes (Hamamatsu R1635P 3/8 in). The PMTs are located at a distance of approximately 20 mm from the edge of the scintillator and optically connected via 'air lightguides', i.e., a tube formed with reflecting Mylar foil. They are housed in iron blocks in order to shield them from the magnetic stray field of the solenoid. The TDS is designed to have a very high efficiency and to define a subset of beam particles which are guaranteed to hit the target. Therefore, it is located as near as possible to the entrance of the TPC and with 20 mm its diameter is somewhat smaller than that of the target which is 30 mm. The TDS gives a signal if at least one PMT was hit. An efficiency of well above 99.9% is assured. The TDS has a sufficiently good time resolution (~130 ps) and stability to be used as an additional detector for the beam TOF system.

4.2. Downstream trigger counters

The ITC provides a trigger for large-angle secondaries emerging from the target. It is mounted inside the inner field cage (IFC) of the TPC (see the following section) and consists of an aluminium/carbon tube with a length of 1300 mm and an inner and outer diameter of 76 and

92 mm, respectively. Six layers of scintillating fibres (each of diameter 1 mm) are glued on the tube (see Fig. 5).

The readout is performed by 24 PMTs (Hamamatsu R1635P 3/8 in.) which are contained in a soft iron magnetic shield on the upstream end cap of the solenoid magnet. Triggering on a logical OR condition of all 24 channels provides a combined efficiency for a single track detection of well above 99%.

Downstream of the solenoid, the FTP covers the small angle region complementary to the ITC. It is made of two planes of scintillator slabs. Each plane consists of seven 1240 mm × 200 mm × 5 mm BC-408 Bicron scintillator, covering a total area of 1240 mm × 1400 mm. Each counter is read out at both sides by fine-mesh PMT tubes (Hamamatsu R2490), previously used in the trigger system of the NOMAD experiment [19]. The planes are at 90° with respect to each other in order to minimize the insensitive areas due to the spacing between the scintillators. The efficiency of the FTP for single tracks is >98%.

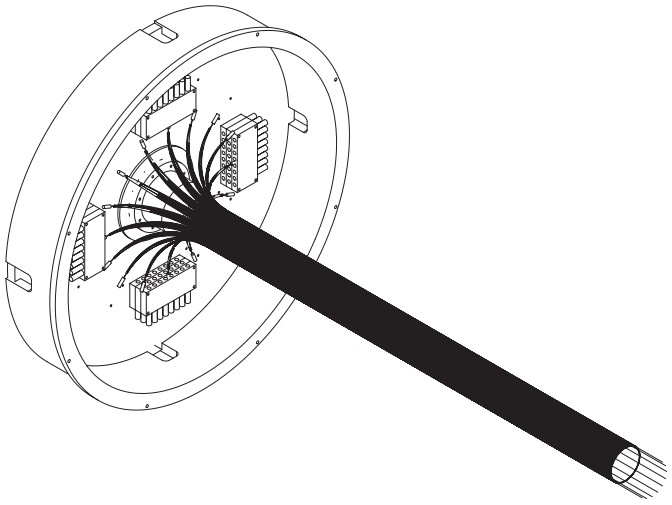

Fig. 5. Sketch of ITC. It is inserted into the inner field cage of the TPC. The compartment with the PMT housings is fixed outside the magnetic field on the upstream solenoid end-cap.

5. The large-angle spectrometer

The large-angle spectrometer includes a TPC and a system of RPC counters all inside a solenoidal magnet. The detector has to measure and identify large-angle tracks in the angular region $20^\circ \leqslant \theta \leqslant 160^\circ$. Particle identification is obtained mainly via the dE/dx measurement performed by the TPC. Additional particle identification is performed through a time-of-flight measurement in the RPCs.

5.1. Time projection chamber

One of the main physics requirements is an acceptance approaching the full solid angle coverage as closely as possible. Contrary to the usual approach with TPCs in fixed-target experiments, where the target is placed upstream of the TPC, HARP has chosen a geometry much closer to the one of collider experiments, with interactions located near the center of the active volume, where the IFC typically surrounds the beam pipe. The choice has been made to place the target in a truncated IFC, whose geometry is determined by the fact that, in the laboratory frame, the interaction products are boosted in the forward direction. As a consequence the TPC can track and identify particles in the large-angle region between 20° and 160°. Fig. 6 shows the schematic layout of the HARP TPC; the dimensions and parameters of the various elements are summarized in Table 2.

5.1.1. The solenoidal magnetic field

The TPC is inserted in the solenoid magnet operating at a field of 0.7 T. Its magnetic volume has a diameter of 0.9 m and a length of 2.25 m. It comprises 88 copper coils, 74 coils with 16 windings each and 14 coils with 20 windings

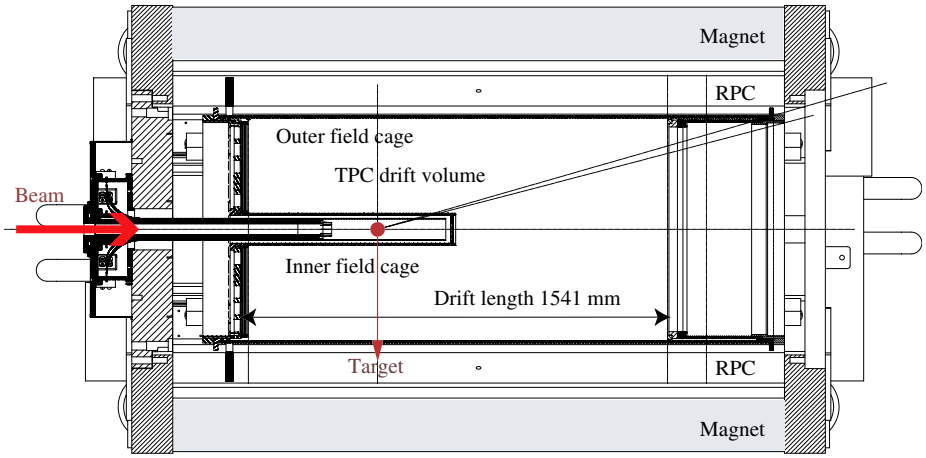

Fig. 6. Schematic layout of the TPC. The beam enters from the left. Starting from the outside, first the return yoke of the magnet is seen, closed with an end-cap at the upstream end, and open at the downstream end. Inside the yoke the cylindrical coils are drawn in grey. The field cage is positioned in the middle of this magnetic volume. The inner field cage is visible as a short cylinder entering from the left. Within it are the ITC trigger counter and target holder.

Table 2
Dimensions and parameters of the TPC

Item	Dimension/value
Active length (OFC drift length)	1541 mm
OFC inner/outer diameter	816 mm/832 mm
IFC length	776 mm
IFC inner/outer diameter	102 mm/106 mm
Field cage strip width/pitch	10 mm/11 mm
Field gradient	111 V/cm
Number of readout pads	3972
Pad dimensions	6.5 mm × 15 mm
Anode wire voltage	1820 V
Anode wire (W + Au) diameter/pitch	20 μm/4 mm
Distance between pad plane and anode wires	5 mm
Cathode wire (Cu + Be) diameter/pitch	70 μm/2 mm
Distance between anode wires and cathode plane	5 mm
Gating grid wire (Cu + Be) diameter	70 μm
Gating grid wire pitch (per plane)	4 mm
Gating grid effective wire pitch	2 mm
Distance between cathode plane and gating grid	6 mm
High-voltage end-plate thickness (aluminized mylar)	50 μm
Gas mixture	Ar (91%), CH_4 (9%)

each. The operating current is 889.4 A. The coils are surrounded by 16 iron beams (20 cm × 20 cm) running parallel to the field axis. The magnet is closed by a 15 cm thick cap on the upstream end, with a 15 cm diameter hole for the incoming beam and the target support. The cylindrical volume is left open at the downstream end for the passage of secondary particles to the forward spectrometer. The TPC detector is positioned in the most homogeneous part of the field volume, from 250 mm ($z = -500$ mm) to 1810 mm ($z = +1060$ mm) into the magnetic field volume.

The magnet was previously used for R&D for the TPC of the ALEPH experiment. For the HARP experiment, the magnet was extended by 50 cm, adding 20 copper coils. Its magnetic field was measured using a set of 48 Hall probes mounted on a rotating axis. Within the volume of the TPC the field along the z-axis is known accurately and differs by at most 3% from the nominal 0.7 T value (downstream and at large radii). The radial component of the field is typically less than 0.5% except at large radii where it locally represents up to 1% of the main field. There is no azimuthal field component. Outside the TPC volume, the field is non-homogeneous due to the open geometry of the magnet. The field measurements and field calculation based on OPERA software [20] were iterated until they agreed to better than 0.5%. The OPERA field calculation values are used in the track reconstruction programs.

5.1.2. The electric drift field

The HARP TPC field cage consists of an IFC, an outer field cage (OFC) and the high-voltage end plate. The amount of material in the radial direction has been kept to a minimum, compatible with the small dimensions of the detector and with the need to build an interaction trigger around the target region. Following simulation results [21] with the Maxwell software package [22] for the electric field geometry and the Garfield package [23] for the electron drift, the IFC and the OFC were both built with a double interleaved strip pattern. This pattern avoids electric field inhomogeneities and high field gradients, which would produce sparks or corona in the Ar gas. The IFC and the OFC have been built from 8 mm (OFC) and 2 mm (IFC) thick Stesalit® (65% glass fibres + epoxy) cylinders. A conductive pattern layer was applied by a galvanic process onto both sides of the IFC Stesalit® tube [24]. The electric field gradient is provided by means of rods with voltage degraders, glued on the field cage cylinders. For the IFC and the rods the coating consist of an initial 0.5–1 μm copper layer followed by a layer of 3–5 μm nickel with a final coating of 0.5 μm gold. The copper improves the adherence on the Stesalit®, the nickel provides good electrical conductivity, whereas the gold provides protection against oxidation. Since it was not possible to apply the same process for the OFC due to its large size, a foil of conductive strips has been glued inside the cylinder, and an aluminium foil has been glued onto the outside surface. The second field shaping layer is composed of individual aluminized Mylar strips. The OFC high-voltage membrane is set at 17.1 kV and the IFC one at 8.4 kV by two distinct HV power supplies. The IFC was made shorter than the OFC to minimize secondary interactions of forward going particles. For a more detailed description of the field cages layout see Ref. [25]. During TPC operation the HV at the downstream end of the IFC was different by about 2% from the HV value on OFC at the same z position. This difference causes electric field components which distort the nominal electric drift field configuration.

5.1.3. Gas choice

A physics requirement is the capability of the experiment to collect large data samples. The design goal of 1000 events per 400 ms spills required a very careful study, since a rate of ≈ 2.5 kHz is a challenging issue for a TPC of this size. The sources of dead time in a TPC are the drift time, determined by the drift velocity in the gas volume, and the readout time of the front-end electronics. Both of these have to be reduced as much as possible.

The gas mixture and the applied drift field have been chosen to ensure that the dead time due to the drift time is ⩽10% of the total dead time. The gas mixture eventually retained for the HARP TPC is Ar/CH_4 91%/9%. At the chosen drift electric field $E = 111$ V/cm, the drift velocity is 5.17 cm/μs, which keeps the total drift time within the required limits. The transverse diffusion, under these conditions and with a magnetic field $B = 0.7$ T, is $\sigma_T \approx 200\,\mu m/\sqrt{cm}$. The anode wires provide a gas amplification of approximately 2×10^4.

5.1.4. Amplification region and pad plane geometry

Charge amplification in front of the pad plane is provided by a set of wire planes. When the ionization electrons reach the readout chamber, at the upstream end of the TPC, they first encounter the gating grid consisting of two interleaved single-wire planes at a base voltage of −67 V and a sweep voltage of ±35 V. To reduce the ion feed-back into the drift volume the gating grid is normally kept closed, and only opened during the drift time of a triggered event. Electrons passing the gating grid are multiplied by the anode wire plane (set at 1820 V) located in between the cathode plane and the pad plane (both kept at null potential). The anode plane and cathode plane are each made of a single 150 m long wire strung onto frames which feature six spokes positioned at the boundaries of the six pad planes boards. Pins inserted into the spokes hold the wire in position. The use of a single wire provides equilibrium of wire-tension forces on both sides of the wire-support spokes, which have a width of 6 mm only. This construction minimizes the dead space created by the supports of the wires. Further parameters and dimensions of the wire planes are given in Table 2.

Given the limited dimension of the TPC, a compromise had to be made between the conflicting requirement of providing many measurement points along high-p_T tracks and the need not to spoil the point resolution by reducing too much the ionization statistics per pad. The choice has been made to subdivide the pad plane, made of six independent sectors, radially into 20 rows of pads. The pad dimensions are $6.5 \times 15\,\text{mm}^2$ and there are from 11 (at the inner radius) to 55 (at the outer radius) such pads per row per sector. The transverse pad dimension and the anode-to-pad distance have been optimized for a FWHM of the pad response function corresponding to the size of the pad, thus ensuring adequate charge sharing between adjacent pads. A mechanical drawing of the TPC pad layout is shown in Fig. 7.

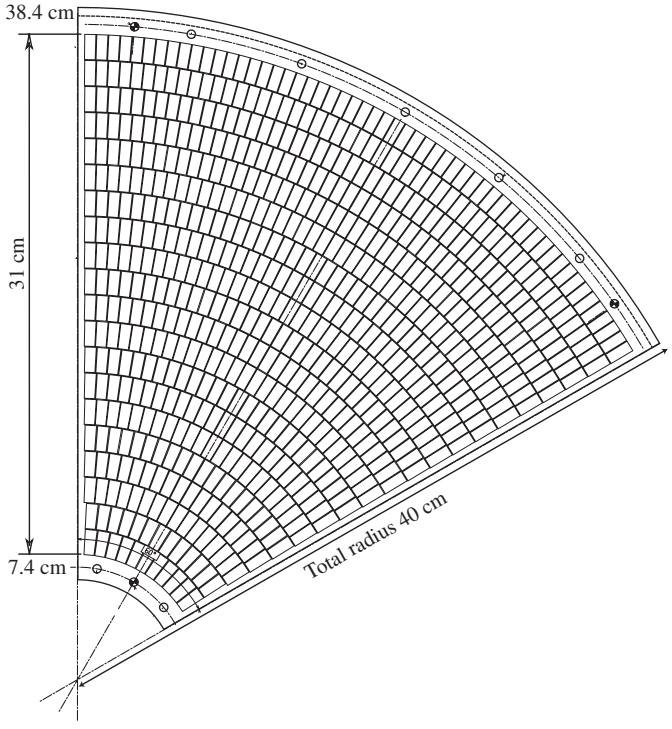

Fig. 7. Mechanical drawing of a sector of the TPC, the layout of the pads is indicated.

5.1.5. TPC readout

The pad plane is made in standard PCB (printed circuit board) technology. The sectors are six-layer PCBs, carrying 662 pads on one side and the front-end preamplifiers on the other. The pads are connected to ALCATEL SMB302 preamplifiers/shapers, whose typical signal rise-time is ≈ 150 ns. The preamplified signals are fed, via bundles of pico-coaxial cables, to 48-channel flash-ADC (FADC) cards, located in VME crates. Fourteen such cards were needed to read out one sector. The FADC cards had been developed as a prototype [26] of the ALICE TPC [27] front-end system; due to their attractive performances they were adopted by the HARP and CERES [28] experiments. Each FADC card provided 10 bit, 10 MHz sampling, giving several samples during the typical duration of the signals, and allowing a complete waveform analysis during the off-line reconstruction.

The total drift time of the TPC corresponds to 300 samples; additional samples (30 before the start of the triggered event's drift time and 30 after the end) were saved to check for piled-up events. Advanced processing capabilities such as on-the-fly zero suppression, pre- and post-sampling, and data compression were available to reduce the data volume. The data compression was implemented in the following way: the 10-bit words were packed in blocks of 32 and organized as a time series, pad by pad. The event rate goal could generally be met, with a total dead-time usually ranging from 500 to 600 μs; different target thicknesses and beam energies made this figure significantly lower or higher, primarily depending on particle multiplicity in the chamber.

5.1.6. TPC operation

Difficulties were encountered in the process of soldering the flexi-striplines onto the backplanes of the six pad boards. This problem remained the largest source of dead channels during operation. Since the pad boards were to be installed inside the TPC enclosure to ensure gas tightness this caused delays in the assembly schedule, shortening the commissioning time.

The rather large variation in the electronic response of the channels created difficulties in determining the operational parameters. The gas gain was set by ensuring efficiency for minimum-ionizing particles (m.i.p.) and avoiding saturation for the energy deposition of ^{83}Kr decays, which deposit the equivalent of about 15 m.i.p. During the 2 years of data taking in total three periods of ^{83}Kr exposure were performed. As shown in the following

section this method is not sufficient to track the performance of the chamber.

After the data taking in the beam, calibration data were taken using a pulser to characterize each single electronic channel. In addition, an exposure to ^{55}Fe sources and ^{83}Kr gas were performed, completed with an extended cosmic-ray exposure. These exposures are described below.

5.1.7. TPC calibration

The relative pad response, the absolute time offset and the drift velocity need to be calibrated. After the calibration of the response, the resolutions on the transverse momentum, $p_\perp$, and the energy-loss, dE/dx, are determined [29–32].

The TPC suffered from a number of shortfalls discovered during and after the data taking:

(1) large excursions of the gains of the pad preamplifiers;
(2) a relatively large number of dead or noisy pads;
(3) large pad gain variations with time;
(4) static distortions caused by the inhomogeneity of the magnetic field, the accidental HV mismatch between the IFC and OFC and a partial 'transparency' of the cathode wire grid;
(5) cross-talk between pads caused by capacitive coupling between signal lines in the multilayer printed boards;
(6) dynamic distortions caused by build up of ion-charge density in the drift volume during the 400 ms long beam spill.

The first four effects are corrected and the latter two are taken into account in the ongoing analysis. A detailed discussion of these effects and their corrections will be presented in a separate paper. Here, only a short description is given of the effective measures taken.

5.1.7.1. Characterization of electronics channels. The response of the pads and their electronics to induced charges is not constant in time, nor equal between pads. Therefore, methods need to be employed which can characterize the behaviour of the electronics, track the variations of the response in time and equalize the response between pads. The pad preamplifier gains show variations of typically 25% and need to be calibrated. The quality of the pad equalization enters directly into the accuracy with which the φ position of the track points can be determined.

In order to characterize the individual electronic channels, in particular their pulse shape, the response to a short pulse applied to each pad separately was measured. Known amounts of charge were injected in each pad and the response of the preamplifiers recorded by the data acquisition system. Three points were taken along the gain curve of each preamplifier, so that a complete channel-by-channel characterization in terms of gain and linearity could be made.

Another method [32] was developed to study the variations of the overall gain of each pad, including the gas gain. Accumulating all the data taken during a longer period in time (the duration of a 'run', i.e. a few hours) for each pad a 'super-event' is constructed as described below. With the observation that most tracks have their origin in the target the energy-flow through each ring of pads is equal if one can restrict the summation to tracks which traverse all pad rings. The latter condition can be met by making a selection of the time-of-arrival of the charge, constraining the ADC readings to those corresponding to tracks which are consistent with originating at the target and traversing the full radius of the sensitive volume of the TPC. With the appropriate geometrical corrections the average gain of each individual pad can be obtained. The spread in the overall gains is quite large. They can range from 0.5 up to 2.0 and even beyond. This method is sufficiently simple that the time-dependence can be followed. Fig. 8 shows the variation observed of the pad gain as a function of time. We use this method in the software chain reducing the fluctuation in the response between pads down to a 3% level. In addition, this method allows dead and noisy channels to be identified.

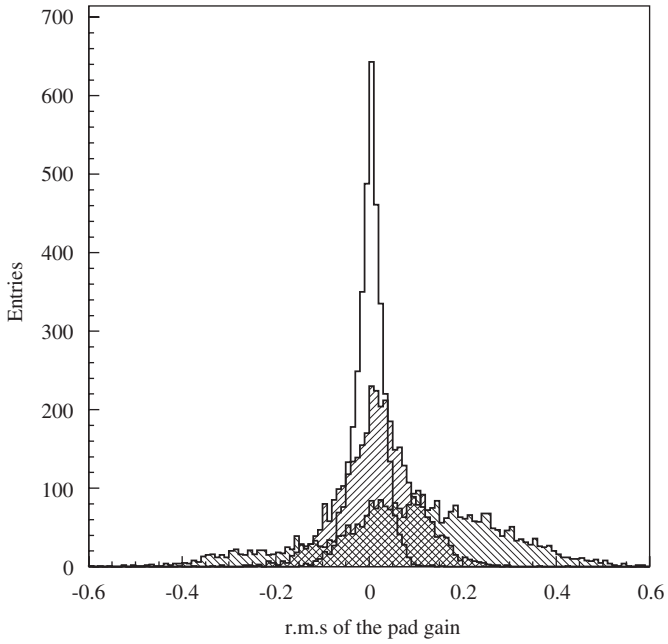

Fig. 8. Variation of the pad gains as a function of time. The unfilled histogram shows the relative pad-by-pad difference for two runs taken with 36 h difference, it has an r.m.s. of $\approx 3\%$. The narrower hatched histogram is for two runs taken with a time difference of a week ($\approx 8\%$ r.m.s.). The broader hatched histogram is for two runs taken with a time difference of 2 months ($\approx 18\%$ r.m.s.).

5.1.7.2. Absolute equalization and energy scale with ^{55}Fe and ^{83}Kr. In order to study the absolute gain of the pad response and to cross-check the method based on 'super-events', radioactive isotopes mixed into the TPC gas (^{83}Kr) or positioned just outside the HV membrane (^{55}Fe) were employed in dedicated measurement periods. The main observable peaks in the TPC gas are 5.9 and 3.0 keV for Fe sources, and 41.6 and 32.2 keV for Kr. The measured

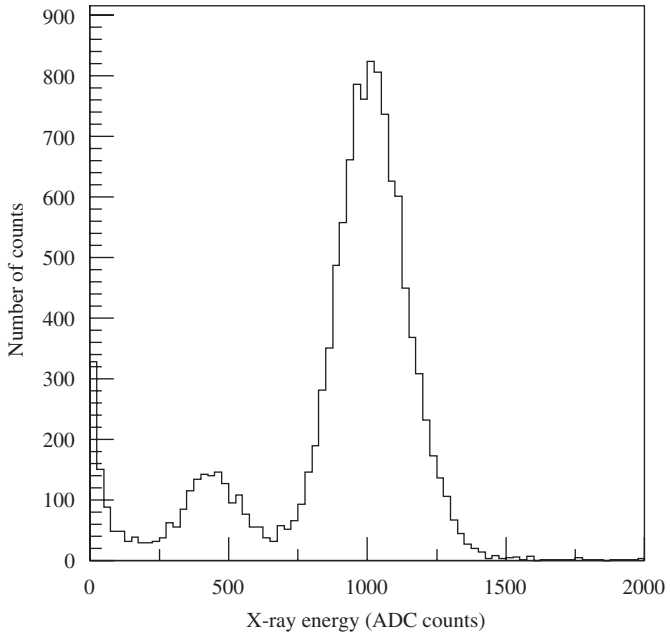

Fig. 9. ^{55}Fe spectrum as measured in the HARP TPC for a set of contiguous pads that happened to have identical electronic gains and no cross-talk. The 5.9 keV and Ar-escape 3.0 keV peaks are clearly visible.

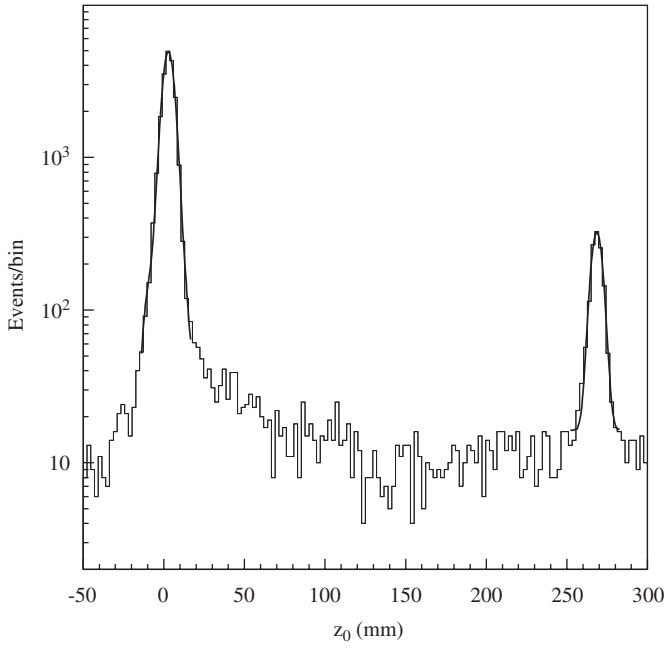

Fig. 10. Longitudinal position distribution of reconstructed primary vertices. The high peak centred at zero corresponds to the thin (5%λ_{int}) Ta target while the lower peak corresponds to the 2 mm stesalit end-plate of the inner field-cage. The data are taken with a 3 GeV/c beam.

^{55}Fe spectrum is shown in Fig. 9. The energy resolution for the 5.9 keV X-rays is about 12%.

The energy range provided by these sources is useful as an absolute energy calibration tool. The lowest energy peak of the ^{55}Fe decay spectrum is close to the typical ionization of a minimum-ionizing particle, the ^{83}Kr one is similar to the behaviour of the nuclear fragments. Thus the full energy range provided by these sources can be used to calibrate the response over a large interval, from m.i.p. to heavily ionizing tracks. The calibration with radioactive isotopes required special data taking periods and was performed only at three occasions during the run using the ^{83}Kr source and at the end of the data taking using the ^{55}Fe source. Since the total ionization corresponding to the energy release of the γ from the source has to be collected a clustering algorithm is applied summing the charge measured in neighbouring pads. This procedure therefore restricts the calibration to pads which are not adjacent to boundaries of the readout sectors.

The data collected with ^{83}Kr source, that generates a wide illumination of the pad plane, is also used to evaluate the overall gain constants allowing a consistency check with the results obtained with the 'super-event' method.

This cross check is, of course, limited to only three time periods and cannot be applied to boundaries of the readout sectors. The values obtained are in agreement at the 10% level.

5.1.7.3. Calibration of drift velocity. Two techniques were foreseen to monitor the drift velocity. The first is based on the usage of a KrF excimer laser to produce photo-electrons at known positions at the end of the drift volume. This allows the measurement of the arrival times of those electrons after having drifted to the readout pads and to use the large drift distance to compute the drift velocity. Using this method an initial evaluation of 5.17 cm/μs was obtained for the drift velocity.

The other method is based on a *tomography* of the experimental setup. It uses reconstructed interaction points of beam particles interacting in the target and in the material at the end of the IFC. Locating known mechanical structures of the detector, the velocity can be deduced by comparison of the known positions to the drift time. An example for a data taking period with a 5% λ_{int} Ta target is given in Fig. 10, where the target and the end of the IFC are clearly seen. From the precisely known distance between the target and the flange of the IFC the drift velocity has been determined to be 5.11 cm/μs with a 0.2% error. This method also calibrates the absolute time-offset of the TPC electronics, providing a reference between measured drift time and track position.

Under the influence of atmospheric pressure and temperature variations the drift velocity shows variations (below 4%) as a function of time. These can be followed using the method described above.

Given all the advantages described above the method based on tomography was chosen to calibrate the drift velocity, the maximal difference with respect to the other method being 1.7%. The drift velocity is evaluated on a day-by-day basis.

5.1.7.4. Cross-talk correction. The cross-talk problem is caused by capacitive couplings between the outputs of some preamplifiers and the inputs of other preamplifiers. It can affect the tracking and energy resolution of the TPC, since it tends to artificially reduce the pulse-height (on average by about 20%) and delay the signals.

The detailed characterization of the electronics described in Section 5.1.7.1 also provided a complete map of the cross-talk effects. With this information appropriate deconvolution algorithms were developed in the reconstruction software. In addition, the effect of the cross-talk could be simulated in detail, so that the remaining degradation of performance could be modelled. The measurements make it possible to determine all parameters needed to fully describe the problem, namely the transfer functions of each pad, a map of cross-talk relations and the capacitive couplings between such pads.

It was found that the induced cross-talk signals could be predicted using the Fourier transforms of the signal from the pad inducing the cross-talk signal and the transfer function of the pad on which the cross-talk was induced, along with the capacitive coupling between the two pads.

A Monte Carlo program has been used to generate a correction based on the understanding of the capacitive coupling. The signal shapes are corrected by predicting the cross-talk signals using a fast Fourier transform (FFT model) and subtracting these predicted cross-talk signals from the measured ones.

The effect of the correction has been studied using both Monte Carlo and data. Studies using cosmic-ray data taken during 2003 show a decrease in $r\varphi$ residuals of 24%. For z residuals the decrease is 16% [33].

5.1.7.5. Static distortion correction. After the end of the beam data taking period, during the year 2003, a special calibration run was performed. By placing a long scintillator rod in the IFC, cosmic-ray tracks which pass through the centre of the TPC could be triggered. These cosmic rays appear in the reconstruction as two tracks emerging from the IFC as if they were back-to-back tracks with the same momenta but with opposite charges. The scintillator was 610 mm long and 25 mm × 25 mm in cross-section and placed in such a way that its centre was located approximately at the position of the target. Under ideal conditions, tracks extrapolated to the axis of the IFC should show a perfect geometrical match. Any deviation from the ideal behaviour is an indication of static distortion: a priori, one cannot exclude a voltage misalignment between the IFC and OFCs, since they are powered independently and the voltage dividers cannot be guaranteed to be identical.

Careful studies of this sample of calibration events indicated a voltage misalignment of the order of 100 V out of a total of about 8×10^3 V, corresponding to $\approx 1\%$. The results of a study using tracks recorded in the beam show an effective mismatch of the order of 150 V. This value was confirmed by looking at the cosmic-ray tracks sample collected in 2002. Additional large HV excursions (+500 and −500 V) were applied in 2003 in order to magnify the effects for more precise study. Appropriate distortion correction algorithms in the TPC reconstruction software allow the effective correction of the voltage offset and of the inhomogeneities of the magnetic field ($E, B, E \times B$ effect).

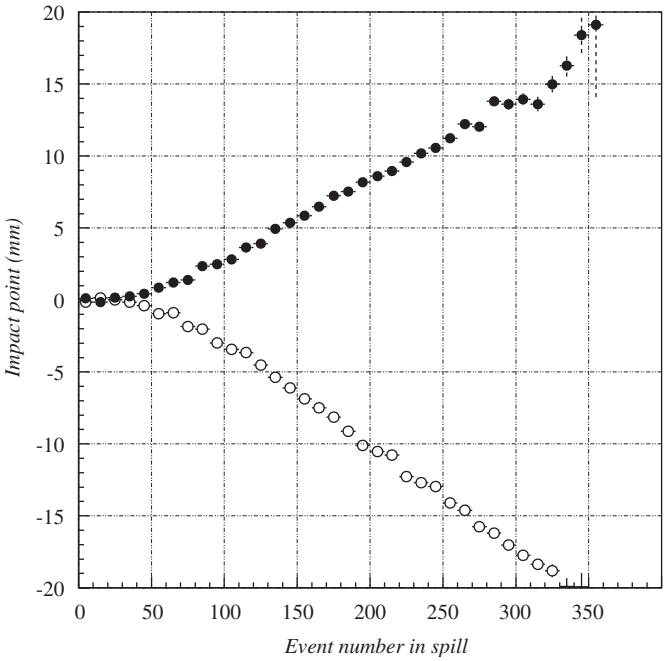

Fig. 11. Demonstration of effect of dynamic distortions as a function of event number in the spill. The filled (open) circles show the average extrapolated distance of positively (negatively) charged tracks from the incoming beam particle trajectory.

5.1.7.6. Dynamic distortions. Unlike the static distortions described above, dynamic distortions change with time and show a dependence on the beam tuning. This is clearly a charging effect of the chamber during the spill. One possible explanation for this effect is the accumulation of positive ions generated by the beam crossing the detector, so increasing the perpendicular component of the electric field. Another possibility is a malfunctioning of the gating grid. Changes in the rate of the trigger or the opening of beam collimators cause an increase or decrease in the dynamic distortions. Different runs or settings are affected in different ways, so that the net effect is sometimes large and at other times negligible.

Methods to correct the dynamic distortions are under development. However, for the first analysis using the large-angle spectrometer, the strategy will be not to introduce corrections, but, given the large statistic available (a few millions of events per setting), to restrict the analysis to the subsample of events not affected by the dynamic distortions.[24] An example is given in Fig. 11, where the

[24]In the worst case (only 10% of the sample used) the statistical error will be $\leqslant 1\%$ to be compared to a systematic error of a few %.

average extrapolated distance of TPC tracks (corrected first for static distortion) from the incoming beam particle is plotted separately for positively and negatively charged tracks as a function of the event number within the spill. It can clearly be seen that this distance diverges from zero after a plateau. The plateau range is used to select events which do not suffer from dynamic distortions.

The d_o (impact point in the xy plane with respect to the incoming beam track) is indeed a very sensitive parameter (and in this sense an estimator) of distortions (both static and dynamic) of the track trajectory. As a consequence, when the d_o distribution is no longer distorted, we see a clear improvement in the detector performances. For example, when the static distortion correction is applied and only the first part of the spill is used, in the dE/dx plot the proton curve becomes clearly defined and the pion spectrum is less broad while at the same time the $p_\perp$ resolution is improved by a factor of 2.

5.1.8. Track reconstruction

The reconstruction of tracks in the TPC proceeds through a number of distinct steps.

First, clusters are searched for in individual pad rows. Pre- and post-sampled 100 ns time slices are selected in individual pads above a defined threshold. These time clusters are given a reference time. Several time estimators can be defined using the time series. The reference time is defined on the rising edge of the signal when the pulse goes over threshold. Hits in neighbouring pads with time stamps that differ by less than 600 ns are included in the cluster formed within the pad row. Each cluster gets a weighted position in the $r\varphi$ direction along the pad row using the pad positions and in the z direction using the time information.

The study of the cluster resolution has shown dependencies on the number of hits used to calculate the cluster position and on the vicinity to a dead region. Values can range from 1 mm (3 hits or more, no dead region) to 3–4 mm (1 hit, close to a dead region). Due to the limited number of measured points per track, in order not to spoil the momentum resolution the "lower" quality clusters were not rejected but enter in the reconstruction with a lower weight.

In a second step, a pattern reconstruction algorithm assigns clusters to tracks. The algorithm uses a general framework for track finding in multi-dimensional space [34], in this case applied to a three-dimensional situation. The algorithm first builds a network of all possible links between the clusters. Then it builds a tree of connected clusters, starting from 'seeds'. Despite the magnetic field, the track model approximates tracks locally as straight lines. The branch of the tree which is retained as the best continuation of the track is determined by examining pairs of fully grown branches and selecting the better one.

Once clusters are assigned to a track, the track is fitted to a helix. The fitting procedure is based on the algorithm developed by the ALEPH Collaboration [35] with slight modifications [36]. The fit consists of two consecutive steps: a circle-fit in the xy plane based on a least square method [37], and a straight line fit in zs_{xy} plane.[25] As in the original ALEPH method, weights are applied to take into account the differences in cluster quality. Tracks which are emerging from the target are refitted using the constraint of the extrapolated beam particle. In the analysis, the error propagation matrix can be evaluated using a Kalman filter method and energy loss inside the target region can be corrected for.

5.1.9. Performances

5.1.9.1. Momentum resolution and dE/dx analysis with cosmic rays. A complete characterization of performance can be obtained by the analysis of the cosmic-ray calibration datasets mentioned in Section 5.1.7. Three evaluations of the transverse momentum $p_\perp$, or the inverse of the radius of curvature ρ, can be obtained. The two track arms, emerging from the IFC as if they were produced in the target, can be fitted and compared independently. In addition, the full extent of the cosmic-ray track traversing the centre of the field-cage can be used as a high-precision measurement of $p_\perp$: since its lever arm is more than twice that of a normal track, the intrinsic $p_\perp$ resolution is more than four times better.

To estimate the momentum resolution one uses the following equation:

$$\frac{\Delta p_\perp}{p_\perp} = \frac{\Delta\rho}{\rho} = \frac{\rho_1 + \rho_2}{\sqrt{2}\cdot\rho_{\mathrm{tot}}} \tag{1}$$

where ρ_{tot} refers to the full-length cosmic track. The value of ρ is defined positive or negative depending on the apparent charge-sign of the track, so that $\Delta\rho$ is centred at zero and has a continuous distribution. $\Delta\rho$ is a better estimator of the resolution than $\Delta p_\perp$ which has a discontinuity at infinity (corresponding to straight tracks).

The performances of the present analysis are summarized in Figs. 12–14. Fig. 12 gives the present status of the calibration (with an average resolution of $\Delta p_\perp/p_\perp \approx 23\%$ in the relevant $p_\perp$ region), adequate for the measurement of the π production yield using the TPC. Further advances in the analysis and calibration are likely to lead to improvements of the resolution. Figs. 13 and 14 obtained, respectively, during the cosmic ray and the physics data taking (with a thin $5\%\lambda_{\mathrm{int}}$ Ta target) clearly show a good separation of p, π and e based on dE/dx measurements.

5.1.9.2. Analysis of selected event topologies. Special topologies (such as elastic p and π scattering) are selected to check dE/dx and momentum resolution. The purpose is to verify the calibrations and assess the correctness of the whole procedure. The simplest and most interesting analysis to be performed as a cross-check of the calibration procedure is the search for elastic scattering events.

[25] The s_{xy} coordinate is defined as the arc length along the circle in the xy plane between a point and the impact point.

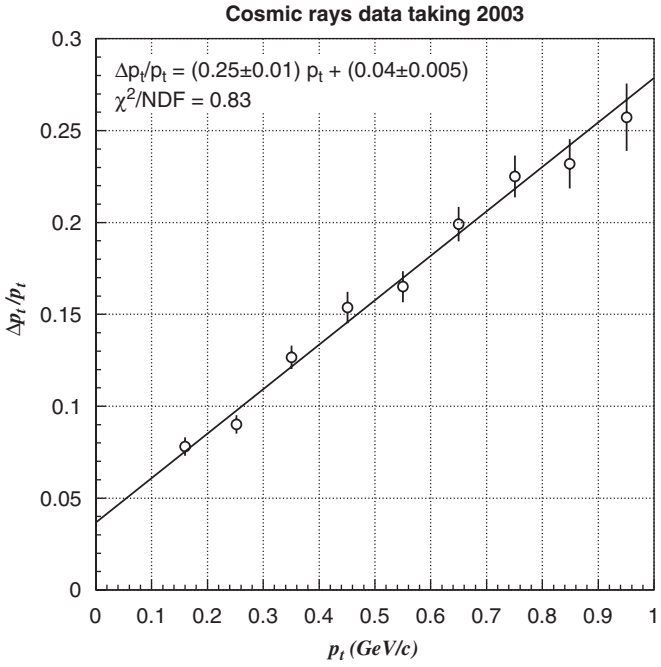

Fig. 12. Transverse momentum resolution of the TPC determined using cosmic-ray tracks as explained in the text.

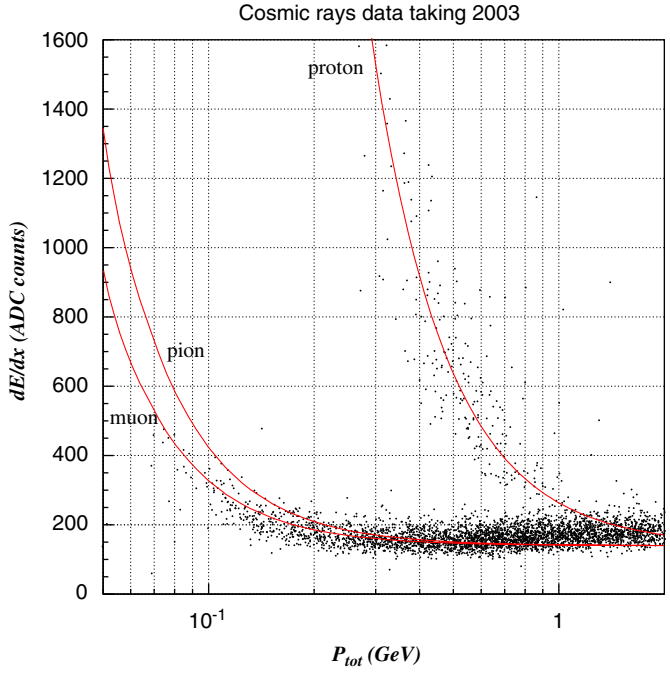

Fig. 13. Particle identification by dE/dx using cosmic rays (which are mainly μ). The different populations are clearly visible.

Elastic scattering processes of the type pp → pp and πp → πp in interactions of π and p on a liquid-H_2 target at 3 GeV/c were studied. The incoming particle is identified by the beam TOF detectors. One of the two emerging tracks escapes in the forward direction sometimes without being measured in the TPC, and the second one is detected at

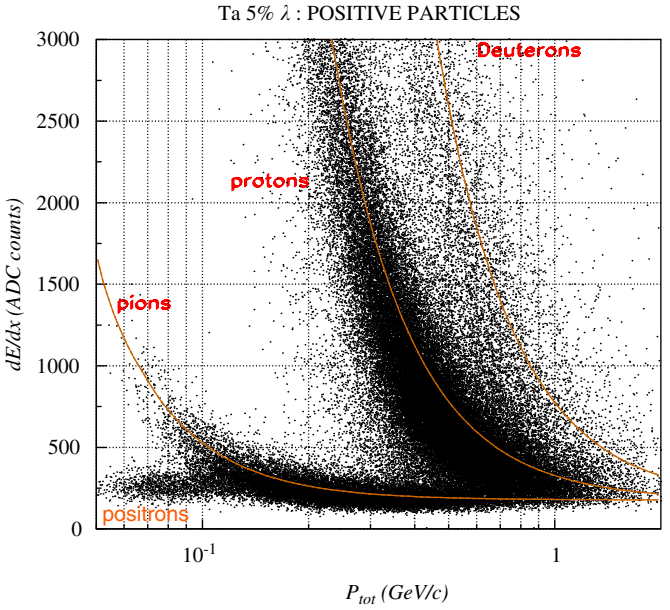

Fig. 14. Particle identification by dE/dx using tracks produced by the beam interactions. The different populations are clearly visible including the deuteron production.

large angle in the TPC. The event selection is performed by requiring at most two tracks in the TPC, allowing both for the large angle and forward particle to be seen.

The elastic reactions at 3 GeV/c can be grouped as follows:

- About 47% of the elastic reactions have a forward proton with angle less than 100 mrad and a large-angle track with momentum less than 300 MeV/c. The large-angle track is stopped inside the target region and the forward track is not detected.
- About 45% of the elastic reactions have a forward proton between 100 and 200 mrad and a large-angle track with momentum in the range from 300 to 620 MeV/c. The large-angle track is well reconstructed in the TPC, the forward one produces only a few hits which is not enough for a momentum estimation.
- About 8% of the elastic reactions have a forward proton between 200 and 300 mrad and a large-angle track with momentum in the range from 620 to 900 MeV/c. Both large angle and forward tracks can be seen by the TPC.

The large-angle tracks from the last two items are used in this analysis [32].

A missing mass distribution can be constructed from the quantity:

$$MM^2 = (p_{\text{beam}} + p_{\text{target}} - p_{\text{TPC}})^2 \quad (2)$$

where p_{beam}, p_{target} and p_{TPC} are the 4-momenta of the incoming beam particle, target particle and the particle scattered at large angle and measured in the TPC, respectively.

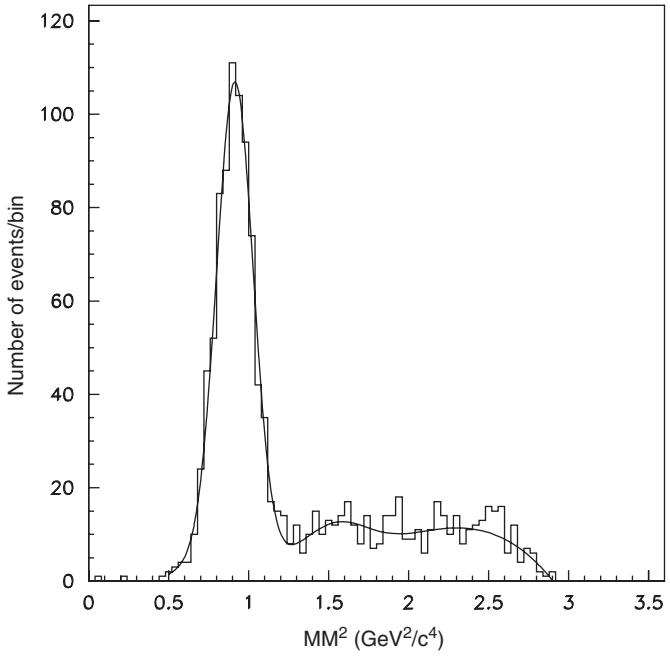

Fig. 15. Missing mass plot for pp → pp elastic scattering events. The peak position corresponds to the squared mass of the proton. The histogram is obtained by requiring proton identification by dE/dx in the TPC.

Fig. 15 shows the distribution of MM^2 for pp → pp after an additional selection using the particle identification by dE/dx in the TPC. The peak corresponding to the missing mass of the proton is clearly visible. The continuous line also includes a background parametrization in which the Δ (1232) production is present. The capabilities of event selection and the correct position of the missing mass peak are a valuable cross-check of the correctness of the calibration procedure. Moreover, after subtraction of the background continuum, the number of events counted in the region of the elastic peak is corrected for the efficiency, and then used to directly evaluate the elastic cross-section. The results obtained are in agreement with previous experiments [38].

5.2. Resistive plate chamber system

RPC were chosen as a time-of-flight system in the large and medium angle acceptance region for particle identification at momenta in which dE/dx measurements in the TPC cannot distinguish between electrons and pions. The basic requirements for the RPC design were an efficiency close to 100% and a time resolution of ~200 ps. Based on these requirements, multi-gap RPCs have been chosen [39].

The RPC system consists of 46 identical chambers, 10 mm thick, 150 mm wide, and 2 m long. Thirty of them are arranged as a barrel around the TPC providing full coverage in azimuth and covering polar angles from 17° to 142° with respect to the beam axis. Chamber dimensions and shape were chosen such that two overlapping layers of RPCs fit into the ~25 mm radial free space between the TPC and the coils of the solenoidal magnet, see Fig. 16(a). The other 16 RPCs are installed downstream of the TPC, perpendicular to the beam. They are arranged in four groups of four overlapping RPCs each, see Fig. 16(b) and cover the angular region 6°–16°.

5.2.1. Mechanical chamber design

The active element of the RPC is a four-gap glass stack with 0.3 mm gas gaps. It consists of two identical sets of three glass plates each, arranged symmetrically on both sides of a central readout electrode. The glass plates are 0.7 mm thick and made of standard float glass[26] with a specific resistivity of $\sim 10^{13}\,\Omega$ cm. The gaps are created by four fishing lines per gap, stretched over the full length of the detector. The glass stack is 106 mm wide, 1930 mm long, and 7.8 mm thick. A cross-section through its short side is shown in Fig. 17.

The two outer glass plates of each set are coated on the outside with graphite tape[27] with a resistivity of 200 kΩ/square. A negative high voltage (−6000 V) is applied to the two outer graphite coatings of the two sets. The graphite coatings next to the central readout electrode are connected to detector ground. The middle glass plate of each set is left electrically floating. In addition to the central readout electrode reference electrodes, connected to ground and electrically isolated from the HV layer by 175 μm of mylar, are placed on the outside of the two glass sets. All read-out electrodes (5 μm Cu on a 75 μm Kapton layer) are segmented into 64 strips of dimensions 29 mm × 104 mm with 1 mm distance between them and connected at one end to the preamplifier inputs.

5.2.2. Front-end electronics

The front-end electronics consists of on-chamber preamplifiers and a discriminator/splitter stage at about 5 m distance from the chamber.

Each RPC has eight preamplifiers leading to 368 readout channels for the full system. The preamplifiers are plug-in modules to the chamber housing. They consist of a fast summing amplifier (Analog Devices AD8009), which combines the signals of the eight strips of a pad. In order to decouple the strips from each other and match the strip impedance, each strip is first connected to a low noise fast transistor (PHILIPS BFR92A); the transistor outputs are connected through printed-board delay lines of equal geometric length to the input of the fast summing amplifier. The delay lines are designed to render the timing of all eight strips within one pad identical (this goal was only met within ~100 ps). Each preamplifier has a calibration input through which test pulses can be injected. Output pulses from the fast summing amplifier have a signal rise time of ~1 ns and the FWHM of the pulse is ~15 ns.

The amplified signals are transmitted through mini coaxial cables over a distance of 0.8–2.5 m to a passive

[26]GLAVERBEL S.A., 116, Chaussée de la Hulpe, B-1170-Bruxelles.
[27]ESD EMI Engineering Corp., Toyo 5-chome, Tokyo 135–0016.

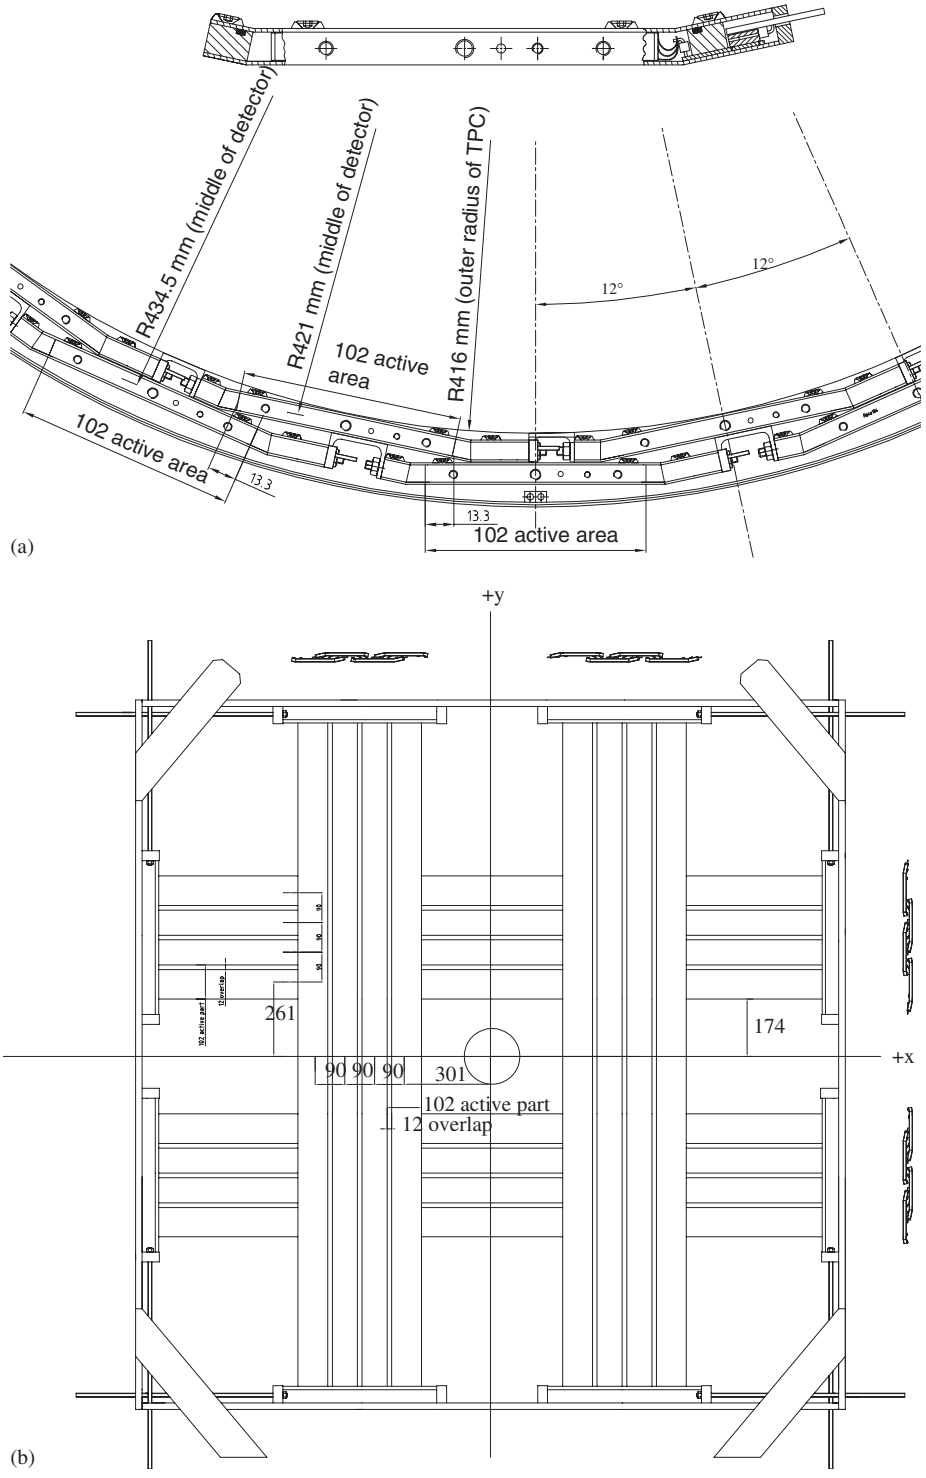

Fig. 16. (a) Cross-section through single RPC and arrangement of barrel RPCs around the TPC; (b) layout of forward RPCs.

connector board and from there through 5 m long 50 Ω cables to a discriminator and splitter unit. Each splitter module has 16 inputs and serves two RPCs. In the splitter module the signals are again amplified (by a factor ~5) using the same AD8009 chip as in the preamplifier and split into two signal streams. One of the two streams is discriminated (with a threshold of about 5 mV) and the discriminator output is sent via 80 m of twisted pair cable to a TDC.[28] The other signal part is sent as an analogue signal over 80 m twisted pair cable and an impedance adapter to a QDC.[29]

[28]CAEN V775 with ~35 ps per count.
[29]CAEN V792 with 0.1 pC per count.

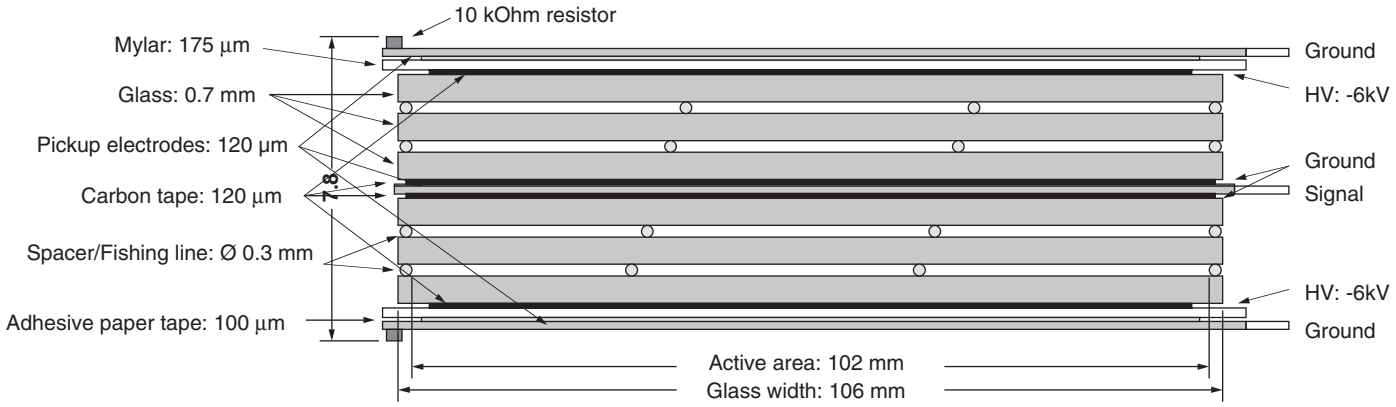

Fig. 17. Cross-section through the glass stack, shown is the arrangement of the components used in the RPC assembly. Note the different horizontal and vertical scales.

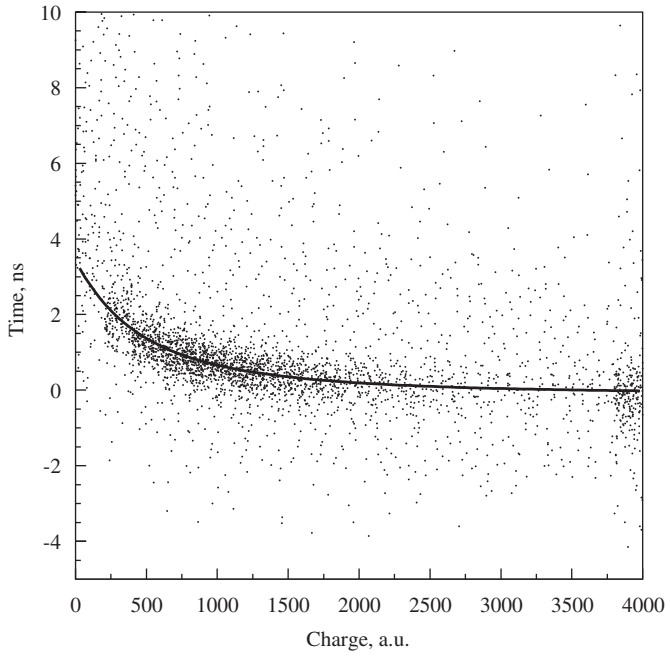

Fig. 18. Time-charge response to negative pion tracks as measured in runs with 8 GeV/c beam hitting 5% λ_{int} Cu, Ta and Pb targets; superimposed is the functional form used to describe the time-charge dependence.

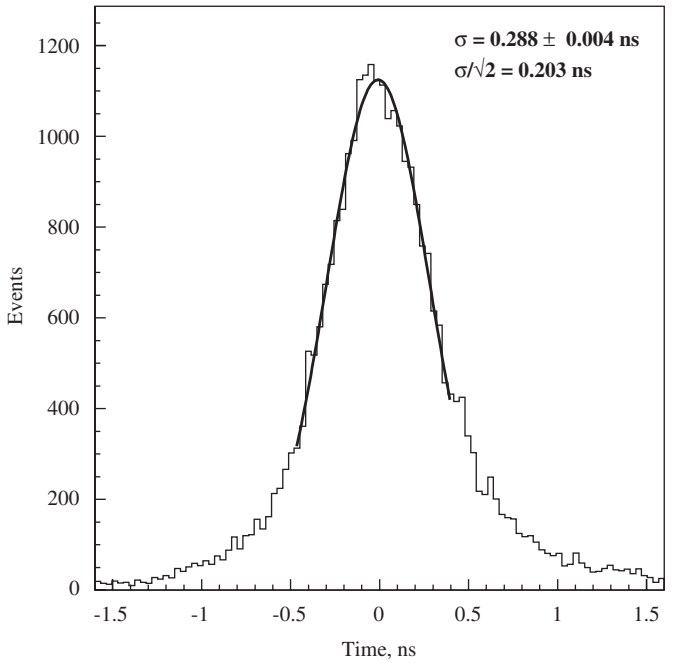

Fig. 19. Difference in time-of-flight measured in two RPCs for tracks passing through pad overlap regions. The sigma of the Gaussian fit gives the convoluted resolution of the two contributing pads, averaged over all pads.

5.2.3. Operation and performance

The RPCs were operated in avalanche mode with a gas mixture of 90% tetrafluorethane $C_2F_4H_2$, 5% sulfur hexafluoride SF_6, and 5% isobutane C_4H_{10}. The HV was −6 kV over two gas gaps. The gas flow was ∼1 volume change per hour with two neighbouring RPCs connected in series to the same gas line. The operating temperatures were, depending on the location, 27–30 °C in the barrel (stable) and 20–30 °C in the forward chambers (varying with hall temperature).

Typical random noise rates (above discriminator threshold) were $\leqslant$200 Hz per RPC, i.e., $\leqslant$1 kHz/m^2, with a few chambers showing 2–10 times higher noise rates[30]. During the 2 years of operation the random noise rates were stable.

The time-charge response of readout channels to pion tracks is shown in Fig. 18 together with the functional form[31] used for all tracks to correct for the time-charge dependence. The measured time-charge dependence for

[30]When some of these chambers were opened we found in all cases that the gas gaps were locally not maintained at 0.3 mm. The fishing lines, serving as outer spacers, had slipped out from between the glass plates over some distance.

[31]Correction $= a + b/Q + c/Q^2 + d/Q^3$, where Q is the measured charge.

protons is shifted typically by about 500 ps towards shorter times due to their higher ionization rate, hence steeper rise of the pulse leading edge.

The efficiency and time resolution of the RPCs have been measured for a few chambers in a dedicated beam exposure; intrinsic time resolutions of ~150 ps and ~98% efficiency over a HV plateau of a few hundred volts were found. The efficiency has been determined using two stacked RPCs as 'trigger' and then checking for the presence of a hit in the third intermediate RPC. It should be noted that particle fluxes in this run were 3 kHz/cm^2 in the centre of the beam (compared to rates below 1 Hz/cm^2 in normal data taking).

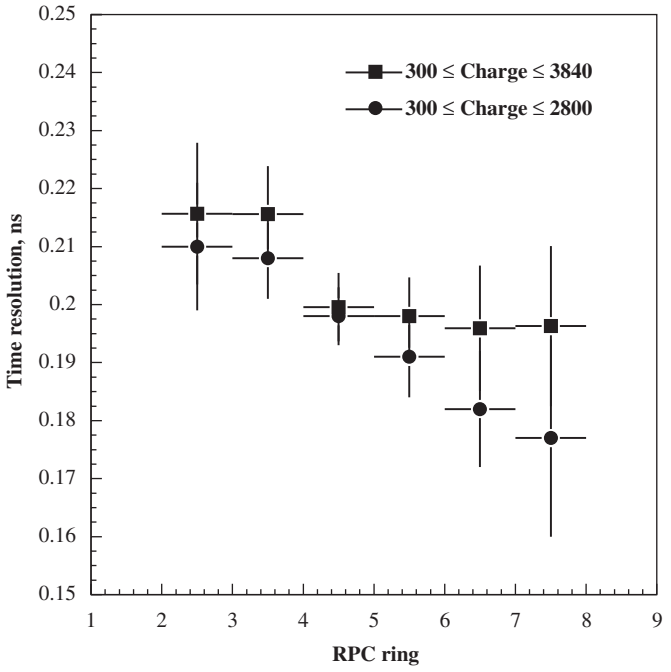

Fig. 20. Time resolution averaged over rings of pads with the same z-position as a function of the pad number for two ranges of signal charge. The ring number 1 is the most upstream, while the ring number 8 is the most downstream one. z-position of the target is within the ring number 4.

An estimate of the time resolutions for the RPCs installed in the detector has been obtained by studying negative pion tracks passing through the overlap region between two neighbouring chambers. Fig. 19 shows the spectrum (accumulated over all pads) of the time difference between two measurements of the same track in two overlapping RPC pads. The width of the distribution corresponds to the convoluted time resolutions of the two RPCs, leading to a resolution of $\sigma = 203$ ps , if both pads are assumed to contribute equally. The resolution measured in this way includes uncertainties related to the time-slewing correction, the impact point corrections [40,41] and correction for a temperature drift of RPC time response. Alignment of time response of all 240 pads was achieved by subtracting a constant, specific for each pad, from the measured time. This correction absorbs the signal propagation times from preamplifier to the discriminator, then from the discriminator to TDC module, processing time of the signal in the electronic modules, etc. These constants have been obtained by comparing the measured time for relativistic pions and their calculated time-of-flight based on momentum and track length measured in the

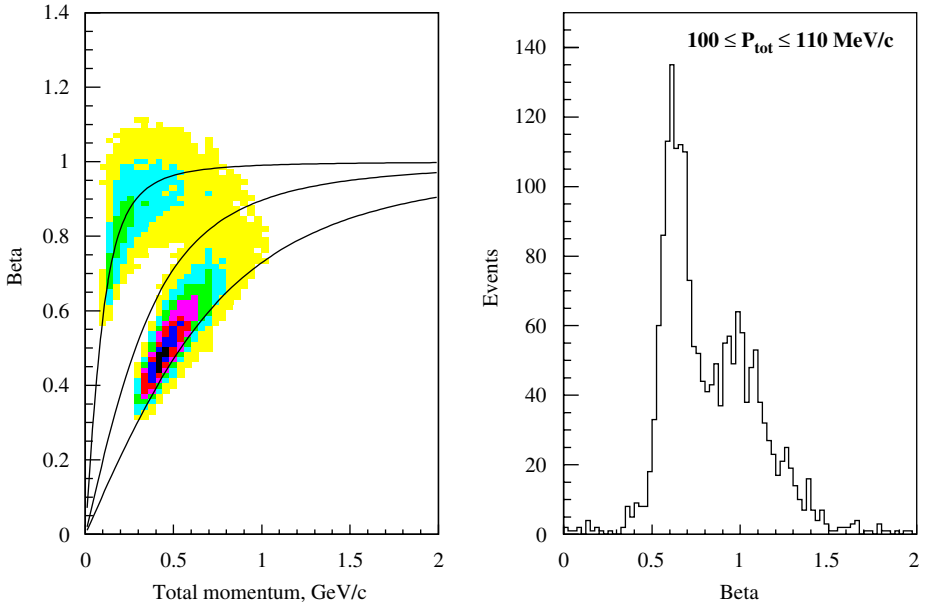

Fig. 21. $\beta = v/c$ versus momentum plot for positive tracks (left). β is calculated from the time-of-flight measured with RPCs (pion time-charge correction applied for all tracks) and the track length reconstructed in the TPC. The nominal β versus p lines for π, kaons, and protons are shown by solid lines. The proton line passes below the proton "island" because of a systematic time shift of proton signals when pion time-charge correction is applied to them. One-dimensional projection of β in the indicated momentum range (right).

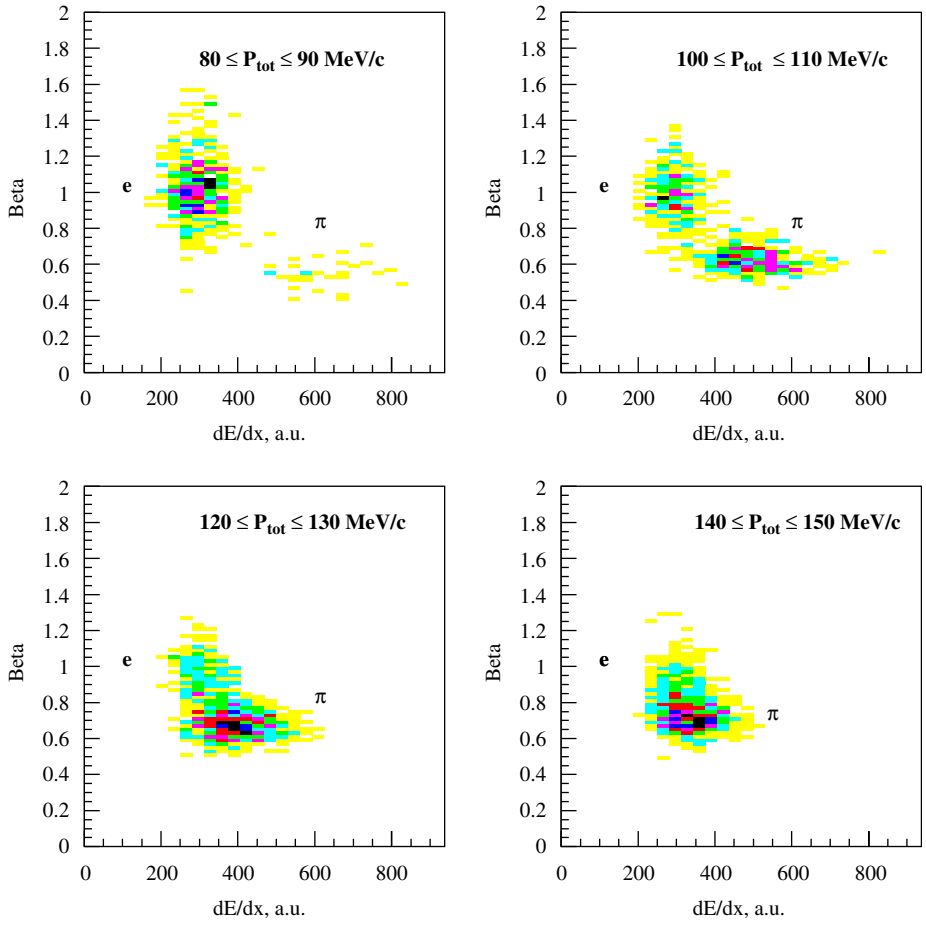

Fig. 22. β versus dE/dx for the same tracks as in Fig. 21 in the indicated momentum region. "Islands" of pions and electrons are clearly discernible.

TPC. A linear dependence of the RPC time response on ambient temperature was found, with a slope as big as 49±5 ps/deg, and was corrected for on a run-by-run basis individually for each pad. It has been observed that the time resolution is independent of particle momentum but varies with pad position along the beam and measured charge of the signal (see Fig. 20). A better time resolution can be obtained for hits caused by neutral particles (e.g. photons from π^0 decays, converting in RPC material). We do not quote it here because the goal for the RPCs is to measure the time-of-flight of charged tracks.

Fig. 21 shows the capability of the RPCs to distinguish between different particles as extracted from runs with 8 GeV/c beam on 5% λ_{int} Cu, Ta and Pb targets for positive secondaries. The left plot shows tracks of all momenta, while the right one is the one-dimensional projection of the indicated momentum slice.

Fig. 22 demonstrates the particle identification power of combined measurement of time-of-flight by RPC and energy losses by TPC. "Islands" of electrons and pions can be identified and separated up to momenta of 150 MeV/c .

6. The forward spectrometer

The forward spectrometer measures the momentum and trajectory of tracks produced at small angles ($\theta \leqslant 15^\circ$) via a set of drift chambers. Particle identification is performed by combining informations from a threshold Cherenkov counter (at higher momenta), a TOFW (at lower momenta) and an electromagnetic calorimeter.

The currently used particle identification algorithms for the forward spectrometer are described in a separate publication [42]. First physics results of the analysis performed with the HARP forward spectrometer on positive pion production in 12.9 GeV/c p-Al collisions are reported in Ref. [43].

6.1. Dipole magnet

The dipole magnet used for momentum measurement in the forward spectrometer has an aperture of 2.41 m × 0.88 m × 1.72 m ($x \times y \times z$). Large copper coils, two above and two below the 1.40 m × 0.90 m ($x \times z$) pole pieces are powered at a current of 2910 A. Its field is non-homogeneous: the vertical component B_y amounts to 0.5 T

in the central region and rapidly approaches zero outside the aperture. Secondary particles leaving the TPC pass through an integral field of $\int B_y \, dL$ of 0.66 Tm. The field of the spectrometer magnet was mapped with a set of Hall probes and the results of these measurements were fitted using the procedure described in Ref. [44]. The agreement between the measured and fitted values of the field is about 1%. The latter are used in the reconstruction program.

6.2. Drift chambers

A set of large drift chambers[32] is placed upstream and downstream of the spectrometer magnet (see Fig. 1) to act as a tracking device for the forward going particles. Five modules of the NOMAD drift chambers (NDC) are used: one directly upstream of the dipole magnet, one directly downstream of the magnet, and three modules further downstream, covering a large lateral surface between the Cherenkov detector and the TOFW. Each NDC module consists of four chambers. The properties of these chambers have been described in Ref. [46] and only their main characteristics and the modifications needed for HARP will be reported here.

The mechanical suspension system has been redesigned to allow hanging the modules rotated by 90° with respect to the NOMAD situation. The dimensions of the NDC modules (3 m ×3 m) allow for full coverage of the forward solid angle as defined by the aperture of the spectrometer magnet. A charged particle can leave up to 12 hits in each traversed module, which allows a fairly accurate segment reconstruction, though each chamber represents 2% of a radiation length. Three more drift chambers were used as preshower in front of the electromagnetic calorimeter.

Each chamber is made of four panels of lightweight composite materials delimiting three 8 mm gas gaps. The very efficient NOMAD gas mixture—Ar (40%)–C_2H_6 (60%)—could not be used because of the stricter CERN safety rules, and a classical, non-flammable mixture: Ar (90%)–CO_2 (9%)–CH_4 (1%) was chosen instead. The central gap is equipped with sense wires at 0° with respect to the vertical axis, the direction of the dipole magnetic field; in the outer gaps the wires are at +5° and −5°. This small stereo angle allows for three-dimensional reconstruction with full precision in the horizontal direction, essential for the momentum measurement via track bending in the spectrometer magnet. The sense wires are equally spaced and provide drift cells of ±32 mm around each sense wire, the cells being separated by potential wires. The sense wires are held at +1300 V and the potential wires at −2900 V. The field shaping is obtained through aluminium strips glued onto the panels and held at decreasing potentials, between −2900 and 0 V, providing a drift field of 900 V/cm.

[32]These chambers have been reused from the NOMAD experiment [45], where they served both as a target for neutrino interactions and as a tracker for the produced charged particles.

The avalanche signals from the sense wires are amplified through a transimpedance amplifier, discriminated by a LeCroy MVL 4075 comparator with a threshold corresponding to an input current of 10 μA and sent to a multihit 128-channels CAEN V767 TDC, running in an emulated 'common stop' mode. The TDC bin width is 0.78 ns, which translates into less than a 40 μm equivalent drift distance, well below the expected spatial resolution of the chambers (a few hundred microns, see below). The trigger signal of the experiment is sent to all TDCs, both as a start ('trigger') signal on each TDC board and also on a special TDC channel to allow correcting for any jitter in the trigger time.

On-line monitoring of the low and high voltages as well as noisy and dead wires was available throughout the experiment. Two out of 69 planes were disconnected after experiencing a short circuit, and only about 1% of the remaining wires were dead, due to problems with the wire itself, its preamplifier, connectors or TDC channel.

To extract the performance of the chambers (spatial resolution and hit efficiency) an alignment of all wires was done first. We used the samples of cosmic events recorded via a special trigger requiring a coincidence between the FTP, the TOFW and the cosmic wall (CW). It selects roughly horizontal muons covering almost the entire chamber cross-section. A further selection on the timing reduces the sample to tracks going in the same direction as the beam: the decoding and reconstruction programs make this assumption to deconvolute the drift time, the particle's time-of-flight and the signal propagation time along the cables from the total time measured by the TDC. The alignment process then consists in reconstructing these tracks and using the distribution of hit residuals to correct

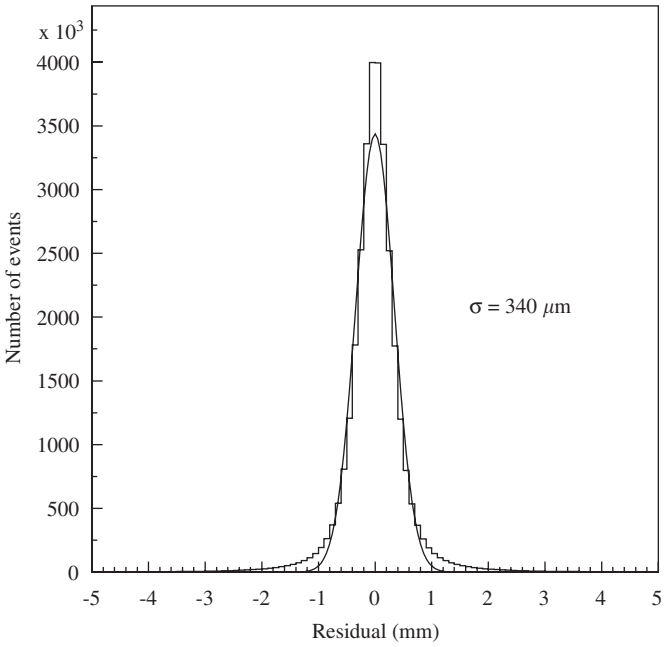

Fig. 23. Residuals (in mm) for a sample of cosmic tracks in the drift chambers.

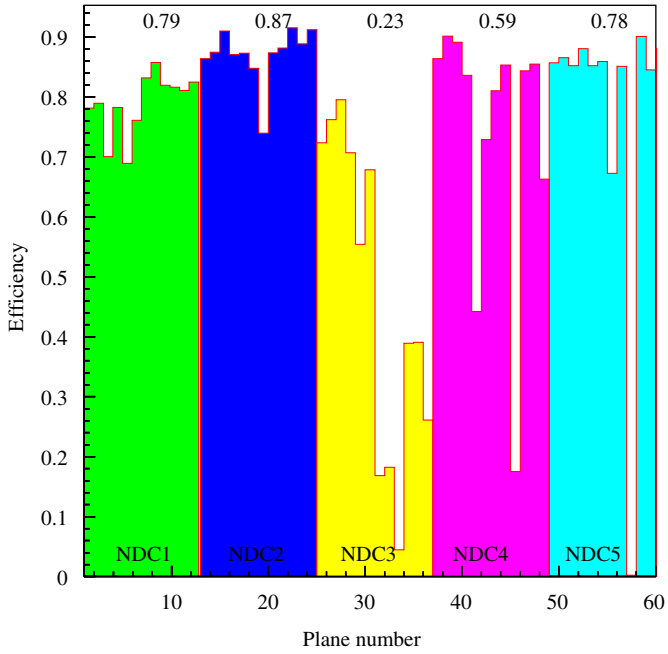

Fig. 24. Hit efficiencies of the drift chamber planes. Different modules of drift chambers (from NDC1 to NDC5, 12 consecutive planes each) are shown by different shades. The average efficiency for each NDC module is also given.

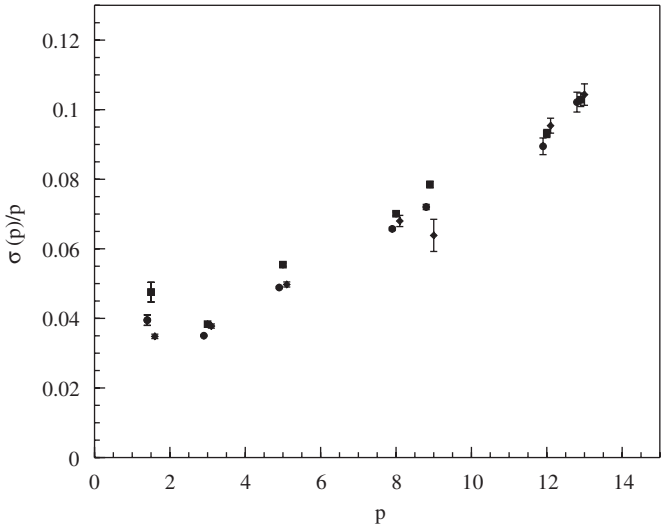

Fig. 25. Momentum resolution $\sigma(p)/p$ (Gaussian fit) as a function of momentum (in GeV/c) for the drift chambers: the data taken using several well-defined discrete beam momenta and no target for pions (filled circles) and protons (filled squares). Also shown (filled diamonds) the corresponding resolution found using the Monte Carlo simulation.

Table 3
Physical properties of the perfluorobutane C_4F_{10} radiator

Boiling point (°C)	−2
$(n-1)\times 10^6$ at atm. pressure	1415
Lorentz factor, γ_t	18.8
π threshold (GeV/c)	2.6
K threshold (GeV/c)	9.3
p threshold (GeV/c)	17.6
Θ_{CKV} max (mrad)	53
$\frac{d^2N}{dx\,dE}\left(\frac{\text{photons}}{\text{cm eV}}\right)$	1.04

the wire positions. The alignment procedure also determines wire time offsets and drift velocities per plane. In total more than 15 000 parameters have to be adjusted in an iterative procedure, reconstructing one million tracks over about 10 iterations. The alignment parameters are validated also using beam data taken with a thick target. The resulting residual distribution is shown in Fig. 23 and gives a spatial resolution of about 340 μm, sufficient for the requirements of the HARP experiment. This value is about two times worse than that measured in NOMAD. This can be explained by the new gas mixture: the longitudinal dispersion of the drifting electrons and the much reduced size of the plateau in the drift velocity distribution as a function of the electric field both contribute to degrading the resolution.

Reconstructing long tracks allows us to compute hit efficiencies per chamber plane. Fig. 24 shows their distribution for the five tracking modules (NDC1 to NDC5, 12 consecutive planes each). The average efficiencies for the central modules (NDC1, NDC2 and NDC5) lie between 80% and 85%, much less than in NOMAD (>95%): this loss of performance can again be attributed to the change of gas mixture and electronics. The reduced efficiencies of the two side modules (NDC3 and NDC4) however, reflect more a specific degradation with time of those chambers, which have a different gas tightness system (o-rings and clamps instead of polymerized silicon joints, see Ref. [46]) and which were not used for 2 years (between the NOMAD and HARP data-taking periods).

Tracks are reconstructed in two steps. In the first step a segment reconstruction algorithm builds track segments in each module. A special software implementation of the Kalman filter technique [47] is used in the second step for matching of track segments and hit collection, taking into account multiple scattering and the exact field map of the spectrometer magnet. To increase track reconstruction efficiency not only tracks with good upstream and downstream segments, with respect to the dipole magnet, are considered, but also tracks with short upstream segments or even with only some hits upstream imposing a vertex constraint. The resulting efficiency for forward track reconstruction amounts to ~80% in the region of interest for physics analysis. The achieved momentum resolution is shown in Fig. 25.

6.3. Cherenkov detector

In the forward direction, the particle identification capabilities of the TOFW are supplemented by a threshold gas Cherenkov detector located immediately downstream of the NDC module after the dipole bending magnet. The radiator gas chosen is perfluorobutane C_4F_{10}. In addition to its environmentally safe properties, its high refractive

index allows the detector to be operated at atmospheric pressure in threshold mode to separate secondary pions from protons (see Table 3 for more details).

The structure of the Cherenkov detector consists of a large steel frame with a 6 m × 3 m opening. This frame supports two horizontal rows of 19 photomultipliers (PMT) each, one row on the top and one on the bottom. The suspensions of the two sets of mirrors are also fixed on this frame. Two large boxes made of thin stainless steel sheets on an aluminium structure define the Cherenkov vessel with a total volume of about 31 m^3. These boxes are welded directly on the steel frame. They are equipped with large honeycomb panels on the entrance and exit paths of the particles. Despite its large volume, every effort was made to keep the Cherenkov vessel as leak tight as possible to avoid losses of the expensive C_4F_{10} radiator gas.

The 8-in diameter EMI 9356-KA PMT were chosen for their very low noise and high gain. In order to increase their useful light collection area to a diameter of 340 mm, the PMT were matched to aluminized Winston cones machined from 150 mm thick PMMA plates. To allow for possible replacements of PMT during operation, optical windows made of 10 mm thick Schott B270 glass separate the PMT from the gas volume. The shielding of the photodetectors against the stray magnetic field of the neighbouring dipole is performed with individual double walls of low carbon steel (5 mm) and Armco alloy (1 mm).

The particles traverse about 2 m of the radiating medium and generate photons that are deflected by about 135° upwards or downwards by two large cylindrical mirrors 6 m long and with a radius of curvature of 2.4 m. The reflecting layers of aluminium with a protective coating of magnesium fluoride were evaporated on appropriately bent 3 mm thick polycarbonate panels. The average reflectivity of the mirrors was about 90%. A modular structure of assembled honeycomb panels supports the mirrors. The goal of the mechanical design was to obtain a large rigid structure with the minimal mass budget along the path of the dominant flux of particles. The light collection system was optimized using the Zemax-EE v.10 optical design software. A light collection efficiency of about 80% was reached.

A permanent flow of about 6 m^3/h of the gas radiator is maintained by an elaborate gas purification and circulation system. The gas is first compressed to 10 bar at 60 °C before being pushed through dessication columns. It is then expanded down to 3 bar in a cold box maintained at −40 °C where it liquefies. The residual nitrogen and oxygen gas are then purged off and the resulting pure liquid returns to the supply bottle. The operation is completely autonomous under the control of a computer. The purity of the gas is continuously monitored with a dedicated ultrasonic device.

In order to identify a particle crossing the active volume a calibration must be performed which assigns the total number of photo-electrons to an individual cluster. Each PMT was equipped with a diode emitting a tiny pulse of light during the inter-spills of the normal data taking. In this way the single photo-electron peaks and the pedestals of each channel can be determined.

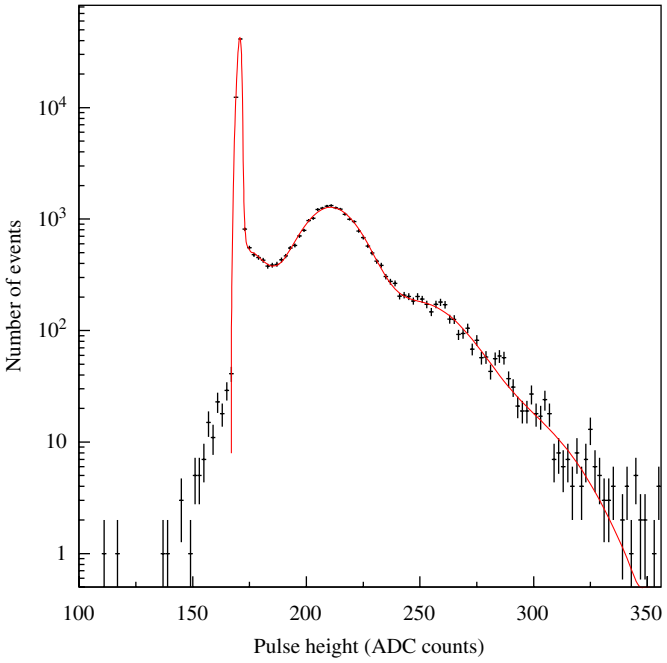

Fig. 26. Photo-electron spectrum from the LED flashing system. The pedestal peak is clearly visible as is the first photo-electron peak. The fit (full line) uses a function which takes into account the Poisson distribution of photo-electron counts and smearing from the electronics.

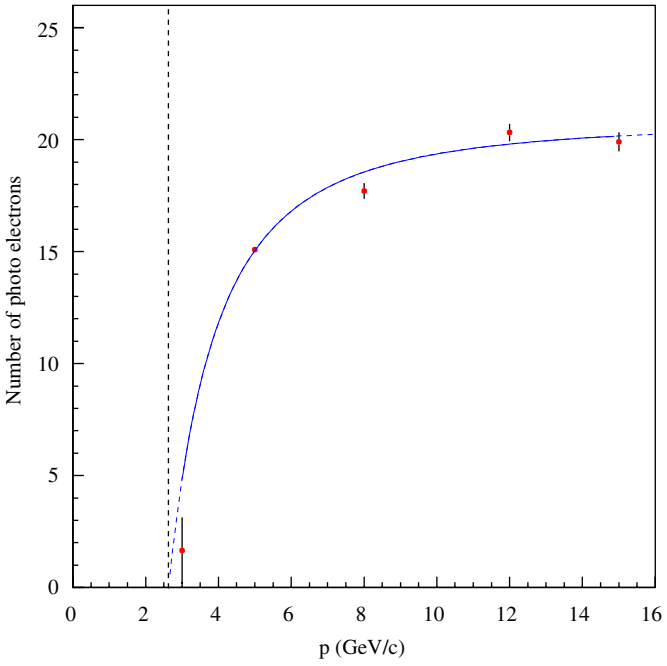

Fig. 27. Pion light yield as a function of particle momentum: the dots are the experimental points taken at five different beam momenta (3, 5, 8, 12 and 15 GeV/c). The solid curve is a fit to the data with the function described in the text. The dashed vertical line marks the threshold for light production from pions in a gas with refractive index $n = 1.0014$.

Fig. 26 shows a typical histogram used for the Cherenkov calibration; the pedestal and the first photo-electron peak are clearly visible. Every 2 days and after any change in the working conditions a calibration of the 38 channels is performed.

Signals in the detector are grouped into clusters, each giving a total number of photo-electrons and a position in the x–y coordinate; since the position is computed as a centre of gravity of the hit PMTs, the y coordinate of a cluster is something abstract, showing how much the signal is shared between the upper and lower rows of PMT, while the x coordinate is directly related to the passage of the charged particle through the gas volume. Using runs at varying energies and selecting the beam particles in a proper way a plot for the light yield versus momentum can be obtained, as shown in Fig. 27.

The curve adjusted to the data is:

$$N_{\text{phel}} = N_0 \cdot (1 - 1/n^2 \cdot (1 + (m_\pi/p)^2))$$

where the parameter N_0 is the number of photo-electrons in the regime of high momentum ($p \to \infty$), n is the refractive index of the gas mixture (constant) and m_π is another parameter of the fit (related to the threshold value) which should be close to the mass of the selected species (pions in this particular case). For the selected sample the fit gives:

- $N_0 = 20.9 \pm 0.2$,
- $m_\pi = 139 \pm 3\,\text{MeV}/c^2$.

This last result is in agreement with the value for the pion mass. The performance of the apparatus can be inferred from the data themselves by exploiting the redundancy of the HARP spectrometer. Information from the NDC reconstruction is used to infer the position of the Cherenkov light cone axis on the PMT plane. Combining the information from the NDC tracks, the Cherenkov hits and the energy measured in the calorimeter, one can discriminate between electrons and pions. An example is shown in Fig. 28 for a 3 GeV/c run. The scatter plot of the number of photo-electrons from Cherenkov hits versus the energy seen by the electromagnetic calorimeter exhibits two well-defined populations. The pion-like sample is characterized by a low number of photo-electrons (the pions being just above the threshold of 2.6 GeV/c) together with

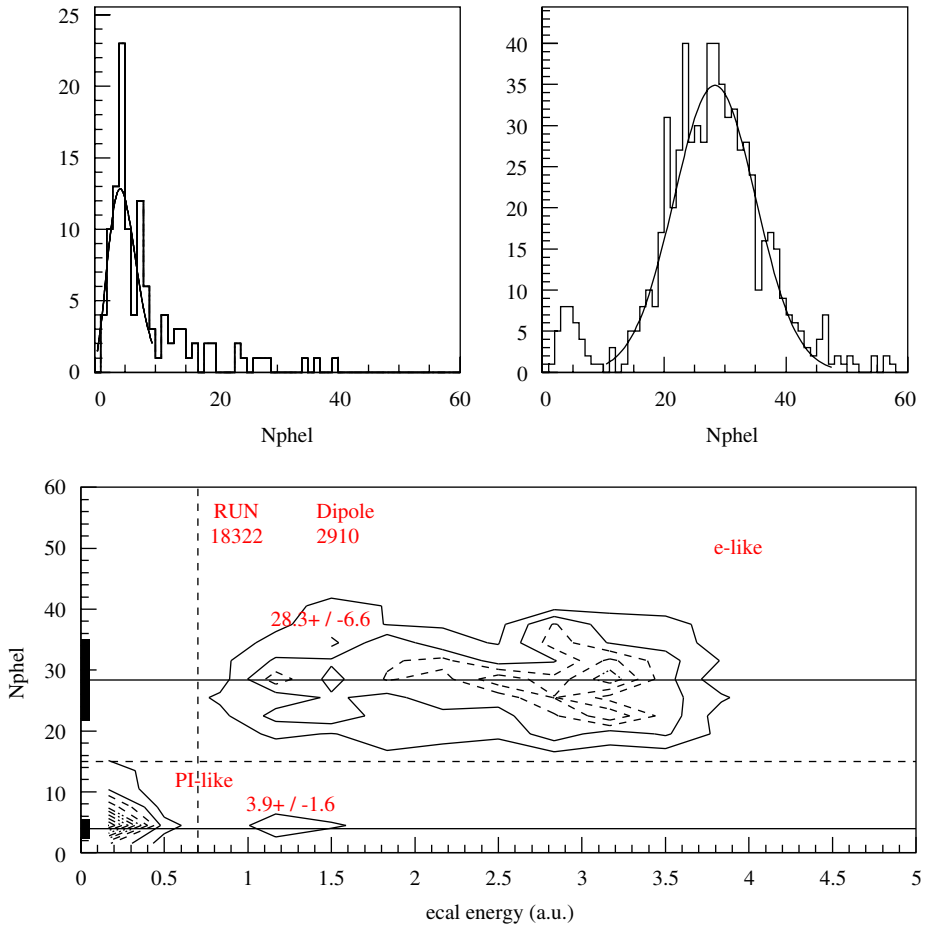

Fig. 28. Scatter plot of the number of photo-electrons per Cherenkov cluster versus collected calorimeter energy, showing a clear difference in two populations (namely π-like and electron-like samples). Notice that the energy scale, here shown as arbitrary units, is consistent with GeV within few percent (see the calorimeter section for further details). The dark bands on the left side represent the average number of photo-electrons with their uncertainty, resulting from a fit which is shown on the top of the picture.

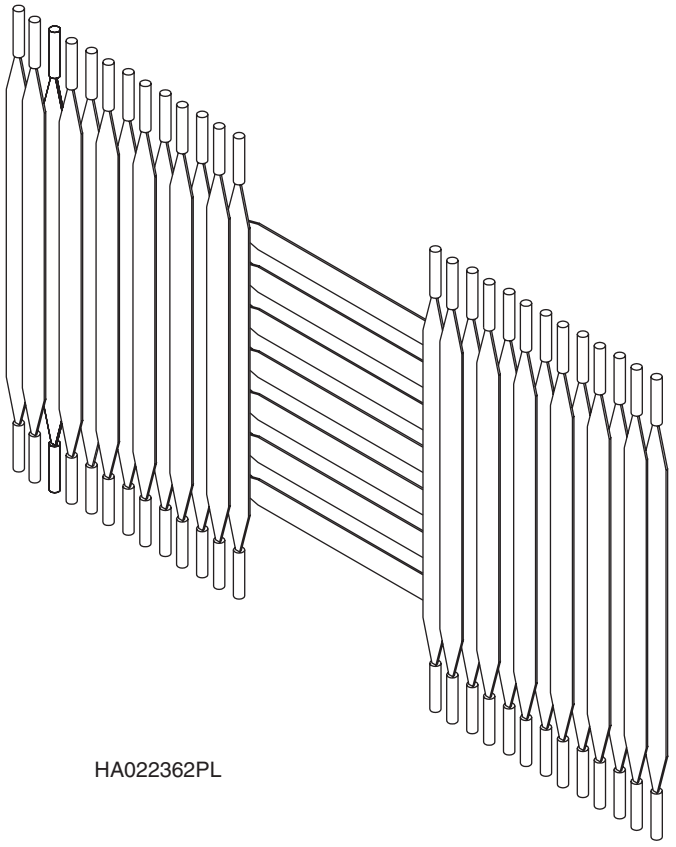

Fig. 29. Layout of the TOF wall of the HARP experiment. The beam enters the wall from the front-left in the figure. The auxiliary detectors used for the calibration are not shown.

a small energy release in the electromagnetic calorimeter. The electron-like sample is instead characterized by a high number of collected photo-electrons per Cherenkov hit and a typical energy close to the nominal $3\,\mathrm{GeV}/c$ measured with the e.m. calorimeter.

6.4. TOF wall detector

The requirements for the forward TOFW called for a time resolution of $\sigma \simeq 250\,\mathrm{ps}$ to separate pions from protons at 4σ up to $3.5\,\mathrm{GeV}/c$, on the basis of a 10 m flight path, and a good transverse segmentation to avoid particle pile-up on single counters. Particle identification is achieved by combining leading-edge time measurements (from TDC) with pulse-height informations for time-walk corrections (from ADC).

The layout of the HARP TOFW, covering an active area of $\sim 6.5 \times 2.5\,\mathrm{m}^2$ at a distance of 10 m from the production target, is shown in Fig. 29; for more details see Ref. [49]. In the left/right palisades, scintillators are 250 cm long and lie vertically, while in the central palisade scintillators are 180 cm long and lie horizontally. The scintillator counters are made of BC-408 bars from Bicron, 2.5 cm thick and 21 cm wide. The chosen scintillator material has a good bulk attenuation length (~3.8 m), a high scintillation efficiency ($\sim 10^4$ photons/MeV) and a high speed (2.1 ns decay time), that make it suitable for fast timing applications over large areas. In each wall the counters are partially overlapped by 2.5 cm to ensure hermetic coverage. Moreover, the particles passing through the overlapping areas can be used to cross-calibrate the counters within a given wall. Two additional sets of scintillation counters ('calibration counters') are placed upstream and downstream of the TOFW to trigger on cosmic events.

At the two ends of each scintillator slab, plexiglass fish-tail light guides have been glued, tapered from the dimensions of the ends of the slabs to those of a plexiglass cylinder ('collar') coupled to the photocathode of a PMT tube (Philips XP2020) through a disc of silicone elastomer, matching the optical refractive indices and leaving no air gap in between.

The scintillators are wrapped in aluminium foil and sealed with black plastic. The quantum efficiency of XP2020 PMT (~26% at 401 nm) is well matched with the spectrum emitted by BC-408. The XP2020 is characterized by a fast rise time (1.5 ns) with a minimal transit time spread (~250 ps), good linearity of response and can achieve a gain $\simeq 3 \times 10^7$ at a typical voltage of 2200 V with a low background noise. The phototube is shielded from the residual magnetic field by a cylindrical screen of mumetal.

The analog signal from the PMTs is fed, after a 40 m long RG-213 cable,[33] to an active splitter chain in the counting room, that divides the signal 25% to the ADC line and 100% to the TDC line, after a leading edge discriminator,[34] with a −38 mV threshold. The discriminator signals are first delayed by ~220 ns using single high-quality cables, then regenerated by a fast discriminator (Line Receiver) and finally processed by the TDCs (model CAEN V775, 32 channels, ~35 ps/ch). The second output of the splitter is sent to a charge-integrating ADC (model CAEN V792, 12 bit, 0.1 pC/ch) after being delayed by ~240 ns.

When a particle crosses a scintillator bar j ($j = 1, \ldots, 39$), equipped with two PMT i ($i = 1, 2$) at a time t_0, at a distance x from its centre, the time difference Δt_i^j between the STOP from the PMT i and the START from a reference counter t_s is given by

$$\Delta t_i^j = t_0 + \frac{L/2 \pm x}{v_\mathrm{eff}} - t_\mathrm{s} + \delta_i^j, \quad i = 1, 2 \tag{3}$$

where L is the scintillator length, v_eff the effective light velocity in the scintillator bar ($v_\mathrm{eff}^{-1} \simeq 5.9\,\mathrm{ns/m}$) and δ_i^j include all delays (cables, PMT transit time, ...). The quantity

$$\Delta t_+^j = \frac{\Delta t_1^j + \Delta t_2^j}{2} = t_0 + \frac{L}{2 \cdot v_\mathrm{eff}} - t_\mathrm{s} \tag{4}$$

[33] Used instead of standard RG-58 coaxial cables, to minimize signal distortion.

[34] Model Lecroy 4413, 16 channel.

is independent of the impact point x along the counter j and allows a measurement of the time-of-flight to be made. The intrinsic crossing time (t_0) resolution for counter j is given by [49]:

$$\sigma_0^j = \frac{1}{2}\sqrt{\sigma_1^{j\,2} + \sigma_2^{j\,2}} \tag{5}$$

where $\sigma_{1,2}$ are the corresponding PMT j intrinsic resolution. The stability of the digital electronics is monitored by pulsing simultaneously all the channels of each discriminator with an electronic signal generated by a pulser circuit. In addition the whole time-of-flight system (counters and electronics) is globally monitored through laser pulses which are directly sent to the centre of each scintillator by an optical fibre injection system several times a day. The laser flash could be triggered externally, in the inter-spill, by a NIM level sent by a calibration trigger and so data could be accumulated in dedicated runs or during the standard data-taking. The laser light is beam split to a fast photodiode, that gives the start for the TDC system and is injected into a bundle of fibres that transmit the pulse to the different scintillator channels (for details see Ref. [48]).

The intrinsic counter time resolution of the HARP TOFW has been measured as $\sigma_0 \sim 160$ ps on average (see Ref. [49] for details). The hit position along a single counter can be obtained by the quantity Δt_-, that is the time difference between the signal from the two counter's PMTs, with a resolution given by $\sigma_0 \times v_{\text{eff}} \simeq 3.2$ cm. The resolution on the TOF measurement for particle identification in the experiment can be expressed as

$$\sigma_{\text{TOF}} = \sqrt{\sigma_0^2 + \sigma_{t_s}^2 + \sigma_{\text{calibr}}^2} \tag{6}$$

where σ_{t_s} is the resolution on the start signal (t_s) and σ_{calibr} is the resolution of the calibration method used, throughout the data taking period: this includes initial calibration of delays and their time monitoring. To avoid a deterioration of σ_{TOF}, the term σ_{calibr} must be kept below 100 ps.

The initial timing calibration (determination of the delay constants δ_i^j at time $T = T_0$) has been done with a cosmic ray run in July 2001. Afterwards periodic recalibrations have been performed with cosmic rays every 2–3 months during HARP data taking, to re-determine the values of the calibration constants δ_i^j at fixed times T_k. In between, the drift of these delays, as a function of the running time T, was traced by the laser calibration system. The precision on the time calibration constants, as determined by the cosmic rays run and traced by the laser system, was estimated to be $\sim$70 ps.

In the experiment, the time-of-flight of particles produced at the target is obtained from the difference between the times measured in the TOFW and in the TOF-B counter, which has an intrinsic time resolution of about 100 ps, or a combination of TOF-A, TOF-B and TDS, with an intrinsic resolution of $\sim$70 ps. Therefore, the final time resolution on the time-of-flight measurement is $\sim$200 ps, considerably better than the design value of 250 ps. Fig. 30 shows in a $1/\beta$ versus $p * Q$ (where Q is the sign of the

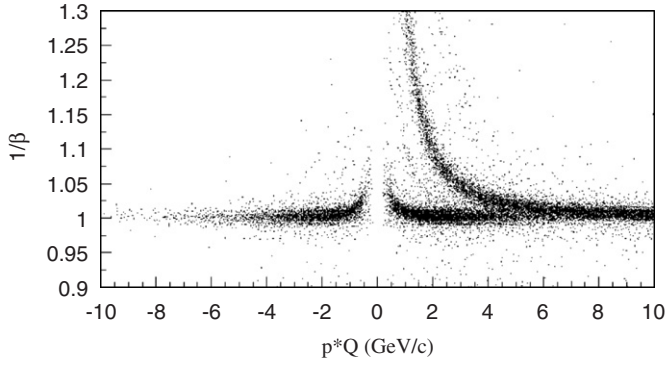

Fig. 30. $1/\beta$ versus $p * Q$ for secondary particles in 12.9 GeV/c p-Al interactions (thin 5%λ_{int} Al target).

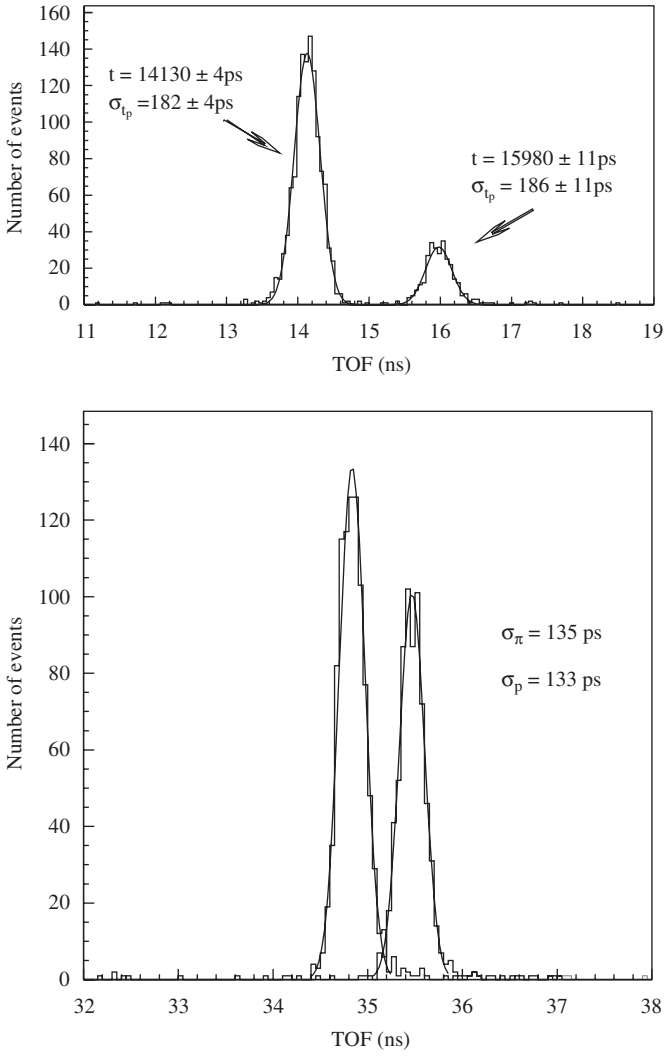

Fig. 31. Particle identification with the TOFW detector, with 3 GeV/c (top: all outgoing tracks) and 5 GeV/c (bottom: tracks in the overlap region of two scintillators) incident unseparated hadron beam. The pion and proton peaks are clearly visible. The time-of-flight is computed between TOF-B and the TOFW system.

electric charge) plot clear bands corresponding to positive and negative pions, kaons and protons.

Particle identification in the HARP TOFW relies on the combination of particle momenta, as measured from the drift chambers, and the time-of-flight between a start signal and a stop signal from the TOFW itself. The previous calibration issues are essential for the quality of the extracted TOF PID and thus the determination of particle masses. After the calibration procedure, π and p are separated at better than 5σ at $3\,\mathrm{GeV}/c$ incident momentum, as shown in Fig. 31.

About 12% of the tracks hit the overlap region of the TOFW detector, between two counters. For these tracks two independent time-measurements are available, improving the resolution. For tracks in this overlap region, kaons and pions are separated at 3σ level up to $3\,\mathrm{GeV}/c$, the full relevant region for MiniBooNE and K2K, while the K/p separation extends up to $5\,\mathrm{GeV}/c$.

6.5. Electron identifier

The electron identifier was designed to provide electron–pion separation when charged pions, accompanied by knock-on electrons, are occasionally identified as electrons by the Cherenkov counter.

The additional capability of detecting π^0's was achieved by simply placing a passive converter of 20 mm thick iron in front of the electron identifier to convert a good fraction of photons from π^0 decays. The π^0's are then identified through the two-photon invariant mass reconstructed under the assumption that both photons originate from the target centre.

The sequence of components of the Electron Identifier in the downstream direction is the following:

- an iron photon converter (2 cm thick);
- three NOMAD drift chambers;
- two planes of existing calorimeter modules from the CHORUS experiment [50];
- a plane of scintillators, the CW,[35] to generate a cosmic-muon trigger (in coincidence with the TOFW) to monitor the response of the calorimeter modules.

The two calorimeter planes (EM1 and EM2) consist of 62 and 80 modules, covering a total active width of 4.96 and 6.4 m, respectively. Each module is composed of scintillating fibres (1 mm diameter) embedded in extruded lead sheets with a volume ratio $\frac{1}{4}$. EM1-type modules, 2.62 m long, are 40 mm thick corresponding to $5.37X_0$ and $0.19\lambda_{\mathrm{int}}$. Scintillating fibres are grouped in two bundles thus providing two independent counters. Each bundle is read at both ends by a 1 inch. PMT. EM2-type modules, 3.35 m long, are 80 mm thick corresponding to $11X_0$ and $0.4\lambda_{\mathrm{int}}$. In this plane scintillating fibres are collected in one bundle only, read by a 2 inch. PMT at each end. The signals from the

[35]The 'cosmics wall' is an array of 32 scintillators of 3.20 m × 0.20 m.

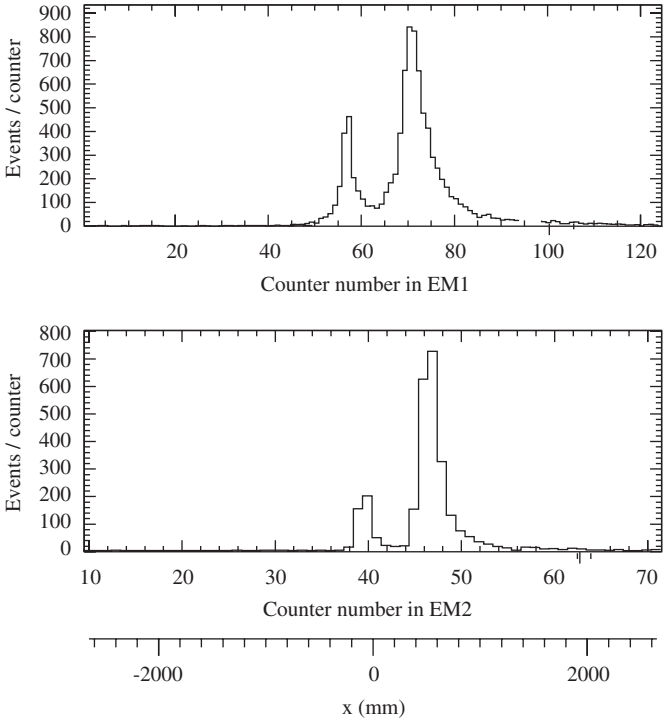

Fig. 32. Occupancy distribution in EM1 and EM2 for incoming electrons at $3\,\mathrm{GeV}/c$ momentum.

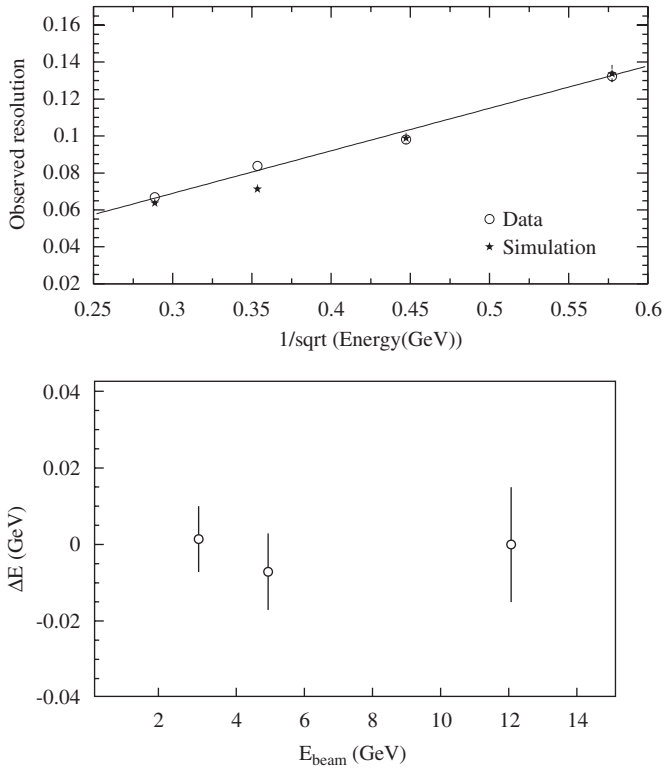

Fig. 33. Top: overall resolution as a function of energy. The open dots represent the measured values at four different energies, the stars indicate the values computed by the simulation. The continuous line is the parametrization $\sigma_E^{\mathrm{obs}}/E = 23\%/\sqrt{E(\mathrm{GeV})}$. Bottom: deviation from linearity as a function of energy.

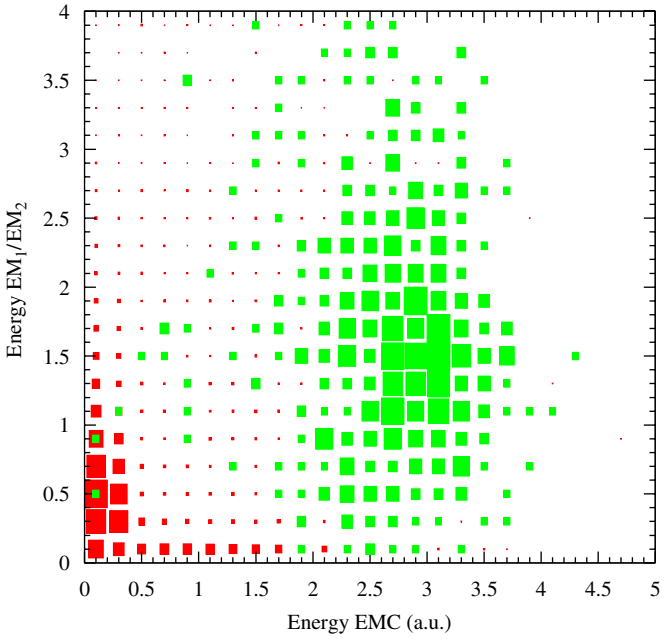

Fig. 34. Two-dimensional distribution of total calorimeter energy and ratio of the energy in the first and second plane for 3 GeV particles. Electrons, identified with the beam Cherenkov, are clustered with a higher value in both parameters and are shown with the lighter shade, pions are concentrated in the lower left corner and are indicated with the darker shade.

calorimeter modules are read out by CAEN V792 QDCs. All the assembled modules were individually tested using cosmic-ray data before and immediately after installation in the experiment and the HV of the PMTs were set to equalize their response within 5%.

In Fig. 32 the occupancy distribution in the EM1 and EM2 calorimeter arrays for an electron-enriched sample at a 3 GeV/c momentum is presented. The distributions show, as expected, two well separated peaks: the first peak, placed along the beam direction, is due to the bremsstrahlung photons emitted by electrons travelling through the magnetic field and the second peak is due to the electron itself.

Dedicated cosmic-ray data samples, taken using the muon trigger mentioned above, were used to provide the calorimeter calibration for equalization of the response in different modules and the measurement of the attenuation length.

The measured energy resolution, reported in Fig. 33, is about $23\%/\sqrt{E(\mathrm{GeV})}$ at all energies, well above the nominal value of

$$(13.8 \pm 0.9)\%/\sqrt{E(\mathrm{GeV})} + (-0.02 \pm 0.04)\%,$$

obtained by the CHORUS collaboration using testbeam data [50]. Simulation studies have shown that the resolution observed in the data can be explained as a convolution of the detector resolution with the spread of the electron beam energy at the entrance of the spectrometer, as shown in Fig. 33. The calorimeter resolution can therefore be assumed to be consistent with the nominal value. In addition, results for the measured attenuation length show that the refurbished modules give results compatible with the ones obtained in the CHORUS experiment.

Table 4
Main datasets taken by HARP at CERN PS in 2001–2002

Target	Momentum (GeV/c)	Length (λ_{int})	Events (10^6)
Be	3, 5, 8, 12, 15	2%, 5%, 100%	37.4
C	3, 5, 8, 12, 15	2%, 5%, 100%	30.7
Al	3, 5, 8, 12, 15	2%, 5%, 100%	34.5
Cu	3, 5, 8, 12, 15	2%, 5%, 100%	36.6
Sn	3, 5, 8, 12, 15	2%, 5%	23.7
Ta	3, 5, 8, 12, 15	2%, 5%, 100%	38.2
Pb	3, 5, 8, 12, 15	2%, 5%, 100%	44.9
N	3, 5, 8, 12, 15	6 cm	13.0
O	3, 5, 8, 12, 15	6 cm	15.5
H	3, 5, 8, 12, 15	6, 18 cm	32.0
D	3, 5, 8, 12, 15	6 cm	21.0
MiniBooNE	+8.9	5%, 50%, 100%, Replica target	22.6
K2K	+12.9	5%, 50%, 100%, Replica target	15.3
H_2O	+1.5	10%, 100%	6.4
Pb, Ta, Cu	+1.5	5%	3.2

Data were all taken with both positive and negatively charged beams, except where explicitly indicated. Some sets include dedicated empty target runs.

The calorimeter is capable of providing stand-alone particle identification on the basis of the total energy deposition and of the longitudinal shower development, evaluated from the ratio of the energy deposition in the two calorimeter planes. The two-dimensional distribution of total calorimeter energy and ratio of the energy in the first and second plane is shown in Fig. 34 for 3 GeV particles: the electrons are well separated from the pions.

7. Detector operation

HARP has taken data in 2001 and 2002, collecting ~420 M triggers with different target and beam settings. A summary of the taken data is shown in Table 4.

To increase trigger purity, data taking in 2002 used thin targets of 5% of an interaction length instead of 2%. In addition to the foreseen settings, additional data were taken with a water target at the lowest possible beam momentum (1.5 GeV/c) at the end of the 2002 period. About 30 TB of raw-data were recorded in about 150 days of running. In the following, all steps of the data acquisition and recording are outlined from the trigger to the software analysis framework.

7.1. Trigger

The central trigger generates several gates in order to ensure that the spill and inter-spill period are clearly separated. The nominal duration of a T9 spill is 400 ms; the minimum delay between the beginning of two consecutive

spills is about 2 s. In addition, all trigger crate controllers, called Local Data Collectors (LDCs), are informed about the beginning of a spill (B-SPL) and its end (E-SPL). B-SPL and E-SPL are stored as special events.

The trigger logic defines a 'physics trigger' to enrich the sample with beam particles giving an interaction in the target. The central part of the physics trigger decision is a programmable logic unit (PLU) [51], which was designed for the trigger of the CHORUS experiment. When a strobe is received, 16 input bits are latched and according to an internal lookup table a pattern of 16 output bits is generated. Input signals are logic signals of all detectors which are of importance for the trigger: TOF-A, TOF-B, BS, TDS, HALO-A, HALO-B, BCA, BCB as beam detectors and FTP, ITC, Cherenkov as interaction detectors. The PLU allows for maximum flexibility for the trigger configuration: the lookup table is loaded at the beginning of each run and may therefore be adapted to different conditions such as beam momentum, beam composition and target. The timing of the trigger signals does not depend on the complexity of the decision as it would be for binary logic realized with standard NIM electronics only.

A standard physics trigger condition for thin targets is $\mathrm{BS} \times \mathrm{TOF\text{-}A} \times \mathrm{TOF\text{-}B} \times \overline{\mathrm{HALO\text{-}A}} \times \overline{\mathrm{HALO\text{-}B}} \times \mathrm{TDS} \times (\mathrm{ITC} + \mathrm{FTP})$. For thick targets a signal in FTP or ITC is not required since most beam particles interact in the target. For beam configurations with extremely high electron content, one of the beam Cherenkov counters was used to veto these events at the trigger level.

For the precision requirements on the normalization in HARP it is essential that a down-scaled sample of minimum-biased data is recorded. The least biased trigger condition is the BS×TOF-B coincidence. Thus, a certain fraction of all strobes reaching the PLU is read out, no matter whether they caused an interaction or not. The value of the down-scale factor[36] i is adjusted in such a way that the dead time due to the down-scaled strobe events does not substantially reduce the number of physics triggers.

Apart from physics and normalization triggers provided by the PLU, calibration events also have to be taken: pedestal events, pulser-related events and events induced by cosmic muons. Common to these events is that they are taken separately by a single subsystem or a set of subsystems. These triggers are acquired during an inter-spill gate, during which the central trigger releases control over the local trigger systems.

Pedestal events are taken once at the beginning of a run. These data are used to monitor the pedestal of all ADC channels as well as to calculate its value and feed it back into the module for an on-line zero-suppression.

Cosmic muon events are used for calibration and alignment. For these, the HARP detector is divided into two subsystems which take data independently:

- TPC cosmics trigger (also referred to as 'upstream' cosmics trigger): TPC and RPCs are read out. The trigger signal is provided by a coincidence of signals in two opposite RPC chambers.
- NDC cosmics trigger ('downstream' cosmics trigger): NDCs and calorimeter are read out. The trigger signal is a coincidence of FTP, TOFW and CW.

7.2. Data acquisition

The HARP Data Acquisition System (DAQ) is made of a switched 100 MBit Ethernet network linking readout processors, event-builders, recording facilities, monitoring and control workstations. The decision to base the readout hardware on VME, the readout and processing nodes on 32-bits Intel processors, and the operating system on Linux has very much simplified the development and the deployment of the system.

As a software framework for the HARP DAQ, DATE [52], the data acquisition prototype of the ALICE experiment at CERN [53], has been chosen. It allowed the implementation of the needed functionalities with minimal add-ons to the DAQ core, and includes a complete set of logging and monitoring tools.

Each sub-detector is connected to an independent set of VME crates, called subsystems. The only exception is the Cherenkov, read out directly by the trigger subsystem, due to its small number of channels. The highest data rate is produced by the TPC, read out by seven crates, six of which are located in the experimental area on both sides of the TPC itself.

The VME trigger crates (LDCs) are arranged in three logical levels as shown in Fig. 35. In the running mode with the complete detector, the distribution of physics trigger signals, the veto/busy logic and spill gates/signals are exclusively handled by the central trigger. The only exception to this rule is formed by calibration triggers which can be taken by each system between spills.

Three processes run on the crate controller LDC: the readout, the recorder and a run-control agent. While the readout process interacts with hardware to arm, readout, reset or disarm several VME modules, the recorder only pulls data from the shared memory and pipes it through the network interface. The shared memory segment is the buffer decoupling the hardware-related operations from the recording and the event building. The readout process does not know if running in global or local mode but acts depending on the kind of trigger and creates sub-event fragments.

Data produced by the LDCs are sent to one or more recording pipelines, each one consisting of an event-building stage, optionally an on-line data objectification, and a storage system. As a rule in the HARP DAQ, one computer performs one stage.

An event-builder (EVB), also called GDC (Global Data Collector), is a computer in charge of collecting and assembling event fragments into full events, checking for

[36]Typical values are $i = 32$ and 64.

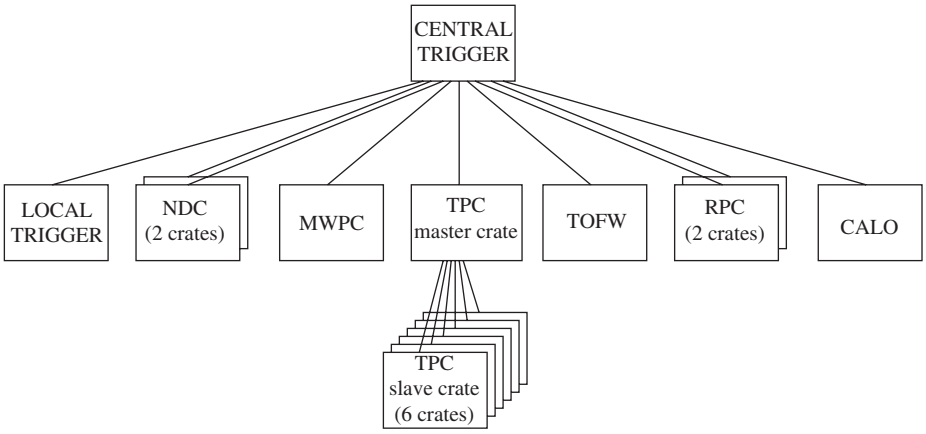

Fig. 35. The central trigger (LDC 16) performs the global trigger decision and distributes the information to the CIRQ trigger modules [51] of the LDC of the second level (LDCs 1–9). The TPC master crate serves the six slave crates (LDCs 10–15) with trigger signals and communicates their busy signals back to the central trigger.

consistency and for compliance with the foreseen trigger structure. Several event-building rules allowed separating the DAQ in many subsystems. Partial event building was possible, for example to allow for building only fragments belonging to a certain subsystem in the case of between-spill calibration events.

The maximum data rate was determined, to first approximation, by the number of available recording pipelines, two in normal running conditions.

One workstation was dedicated to running the main run-control process and its user interface. The main run-control task was monitoring and modifying the status of all the DAQ instances, i.e., all the LDCs and GDCs, taking appropriate actions on user commands or on requests from the DAQ processes. Every computer running an instance of the DAQ was also running a local run-control process sending status information and acting upon request of the main process.

Two database servers, accounting for a total of 1 TB fail-safe disk space, and capable of very fast I/O (about 40 MB/s concurrent read and write speed), made the local disk buffer. The total capacity corresponds to about two days of data taking at the maximum trigger rate.

HARP's main I/O was based on Objectivity/DB [55]. The DAQ system provided two slightly different approaches for writing objectified data to the so-called off-line federation in Objectivity's jargon. In the first approach, data were pipelined from the EVB to an objectifier process taking care of writing Objectivity files to the on-line federation, physically resident in the local disk buffers. Those files were then transferred and committed to the off-line federation, located in the CERN computing centre, before being removed from the on-line one. The second approach relied on the storage of temporary raw-data files from the EVB to the local disk buffer. Those files were then asynchronously processed by the objectifier, and data were committed directly to the off-line federation. The temporary files were deleted only upon successful storage of the raw-data to the computing centre. Both approaches proved to be valid. The second, accounting for one less step, and de-coupling the event-builder from the objectifier, is simpler and proved to be slightly more stable.

Depending on the beam energy, the target configuration and the running conditions, the actual event rate has been between 200 and 400 Hz. The DAQ was designed for a maximum event rate of 2.5 kHz, corresponding to 1000 events per 400 ms-long PS spill. This figure was unprecedented for a TPC, which is in our case the most demanding detector from the point of view of the data rate. Such a trigger rate has never been used in practice, the main reasons being the need to avoid piled-up events in the TPC and to ensure the necessary trigger purity.

The total data rate strongly depended on the number of spills per PS cycle assigned to HARP. The maximum rate in normal data-taking mode has been about 6 MB/s, comfortably split between the two available recording pipelines. Special calibration runs, when the trigger was continuously active, were able to saturate the pipelines with a rate of about 20 MB/s.

Overall, the DAQ and recording system showed excellent stability, with an overall efficiency of 88% during the 1st year of data taking, and about 95% during the 2nd year.

7.3. Detector control

The HARP experiment is controlled and monitored by the DCS. The system consisted of two independent subsystems, one (HVC) for the control and monitoring of the CAEN High Voltage power supplies (920 channels) and the other (ADSM) for the monitoring of analog and discrete signals (160 channels).

The High Voltage Control uses the PVSS2 SCADA in the first supervisory PC and DIM/SLiC client-servers on two PCs for the CAEN HV power supplies of types SY 127, 527, and 403. The Front-End PCs run Linux software and communicate with the first supervisory PC, where the SCADA software PVSS2 is running under Linux. The data

from the HV units is accessible via a DIM (Distributed Information Management) Client/Server application. The Client is a DIM API manager which is part of the PVSS2 software and located in the supervisory computer. The Server consists of the SLiC software located in the front-end HV PCs.

The ADSM is based on the National Instruments FieldPoint system running under LabView on the supervisory PC with an OPC client-server on the processing PC. The FieldPoint network PC runs the Windows NT operating system and accesses the networks via two OPC server applications. The PC communicates with an OPC client in the second supervisory computer, where LabView runs under Windows 2000.

7.4. Software framework

In HARP, standard software engineering procedures were adopted. User requirements were defined and then transformed into a set of software requirements. The architectural Object Oriented design was worked out. The detailed software design and code production in C++ language were implemented. Unit testing, system testing, and release procedures were defined and implemented.

The HARP software components described in Fig. 36 have been developed and used for detector calibration and performance studies, trigger and background studies, beam particle identification, on-line applications, data quality, and productions for data analysis.

A short description of the functional specifications of these software components is given below.

DAQ is the data acquisition software based on the DATE package. HarpEvent is the component implementing the HARP transient event model, including a structured description of settings, reconstruction objects, MC information, etc. It is based on the Gaudi framework developed by LHCb [57]. HarpDD is the component implementing the HARP detector geometry and materials data in a neutral-representation format. It also contains the alignment and calibration data description. DetRep is the component creating the various geometrical representations of the detector objects most suitable for use by the physics applications. It is based on the GEANT4 solid modelling. ObjyHarp is the component implementing the HARP persistent event model. It is based on ObjectivityDB database, and mirrors the transient event model. In 2003 HARP migrated its data to Oracle, thus an equivalent component implementing the HARP persistent event model in Oracle exists. ObjectCnv is the component implementing the unpacking of the raw data and the construction of the transient C++ objects used by the physics applications. It can use transparently both on-line data and stored off-line data, as well as Monte Carlo output. ObjyPersistency is the component implementing the adapter to use the Objectivity or Oracle databases, while allowing the physics applications not to depend at compile time on the I/O solution. EventSelector is the component implementing the event selection and data navigation functionality. Simulation is based on GEANT4. This has also been used for the T9 beam simulation, and for understanding and resolving trigger rate problems. Reconstruction is the component implementing the computation of reconstructed objects at various levels. It uses a dedicated package for Kalman filter operations. HarpUI is the component implementing the event display. It is also used on-line, and is based on ROOT [54]. An example of

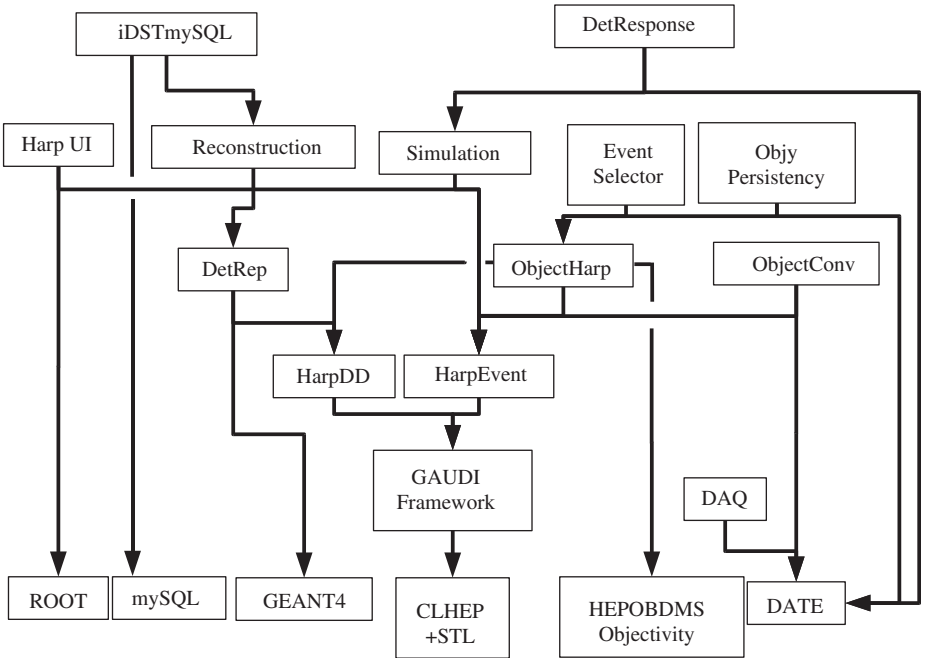

Fig. 36. HARP Software architecture. The various components are explained in the text.

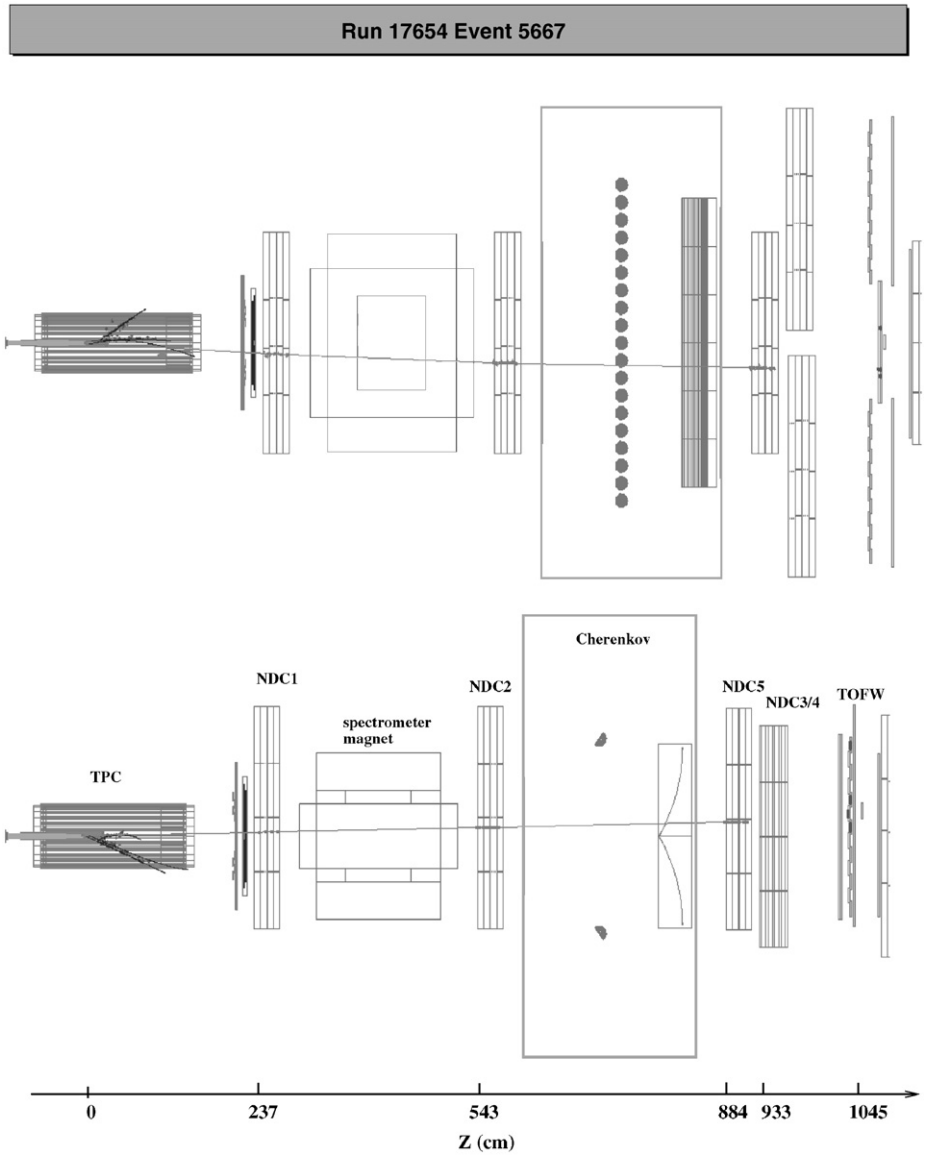

Fig. 37. HARP event display: an example of a reconstructed real data event (top and side views). Main HARP sub-detectors are shown. Short tracks are reconstructed in the TPC, while the longer track is reconstructed in the forward spectrometer.

reconstructed real data event is shown in Fig. 37. DetResponse is the component implementing the digitization of the main detectors. iDSTMySQL is the component implementing the DST concept for distribution in the collaboration. It contains the persistent-capable physics objects (including reconstruction, simulation, geometry, and event model objects). It supports both a neutral file format and Linux MySQL.

The entire software chain can be rerun on the data produced/retrieved in iDSTmySQL without needing to access the CERN mass storage system (CASTOR) and central databases (Objectivity, Oracle). This allows full distribution of the analysis work within the HARP institutes.

HARP data are reconstructed in production at a rate varying from 0.3 s/evt/GHz (large angle) to 2.1 s/evt/GHz (all forward detectors). HARP data are simulated in production at a rate of 1.7 s/evt/GHz for both large angle and forward detectors simulation. These rates allow in all cases more than one million events per day per 10 GHz (e.g., three standard processors clocked at 3 GHz).

7.5. Monte Carlo simulation

One of the main external packages in the HARP software, see Section 7.4, is GEANT4 [56]. It is used as a toolkit based on the Gaudi framework [57]. All GEANT4 components including UI interface are initialized and executed via Gaudi algorithms and services, which allows the same user interface both at reconstruction and simulation level.

A complete set of GEANT4 physics processes is available for simulation of particle transport and interaction inside the HARP detector. For electromagnetic interactions the standard package is used [56]. For performance study in the TPC the low-energy package was applied which allows us to see effects of gaseous detectors. For simulation of hadron nucleus interactions the set of physics processes provided for LHC detectors are used [58].

The simulation of HARP sub-detectors response is subdivided into three main steps:

- GEANT4 simulation of particle transport, giving energy deposition in sub-detectors and creation of Monte Carlo hits;
- Simulation of electronic signals;
- Application of efficiency functions, resolution functions, and dead channel data and creation of Monte Carlo digits in experimental data format.

The true Monte Carlo information is also recorded. Only beam particles and energetic particles produced inside the target are stored. The complete history of the event is available during event simulation, which allows all Monte Carlo hits to be associated with these initial particles. The true Monte Carlo information is propagated to simulated detector response (digits) and to reconstructed objects.

8. Conclusions

The HARP experiment has successfully operated in 2001 and 2002. It has collected data for the measurement of hadron production cross-sections in the range 1.5–15 GeV/c for incoming protons and pions. A total of about 420 millions of triggers were taken. The main components of the detector have been built in less than 1 year and have proven to give satisfactory performances. The redundancy of the particle identification capabilities of the apparatus is of special importance to reduce systematic errors in the cross-section measurement.

The currently used particle identification algorithms for the forward spectrometer are described in a separate publication [42]. First physics results of the analysis performed with the HARP forward spectrometer on positive pion production in 12.9 GeV/c p-Al collisions have been published recently [43].

Acknowledgements

We gratefully acknowledge the help and support of the PS beam staff and of the numerous technical collaborators who contributed to the detector design, construction, commissioning and operation. In particular, we would like to thank D. Conventi, M. Delattre, C. Detraz, A. Domeniconi, M. Dwuznik, D. Lacroix, M. Lollo, J. Loquet, J. Mulon, L. Musa, X. Pons and M. Scandurra, for their support in the construction of the detector. The collaboration is indebted to M. Baldo Ceolin, A. De Min, M. Doucet, L. Linssen, A. Pullia, A. Semak, E. Usenko and V. Zaets for their contributions to the experiment.

We acknowledge the contributions of V. Ammosov, G. Chelkov, D. Dedovich, F. Dydak, M. Gostkin, A. Guskov, D. Khartchenko, V. Koreshev, Z. Kroumchtein, I. Nefedov, J. Wotschack, and A. Zhemchugov to the work described in this paper.

The experiment was made possible by grants from the Institut Interuniversitaire des Sciences Nucléaires (Belgium), Ministerio de Educacion y Ciencia, Grant FPA2003-06921-c02-02 and Generalitat Valenciana, grant GV00-054-1, CERN (Geneva, Switzerland), the German Bundesministerium für Bildung und Forschung (Germany), the Istituto Nazionale di Fisica Nucleare (Italy), Joint Institute for Nuclear Research (Dubna), INR RAS (Moscow) and the Particle Physics and Astronomy Research Council (UK). We gratefully acknowledge their support.

References

[1] J.V. Allaby, et al., CERN report 70–12, 1970.
[2] T. Eichten, et al., Nucl. Phys. B 44 (1972) 333.
[3] B. Autin, A. Blondel, (Eds.), CERN 99-02, ECFA 99-197.
[4] S. Ozaki, et al., (Eds.), Feasibility Study-II of a Muon-based Neutrino Source, Report BNL-52623, June 2001.
[5] A. Blondel, et al., CERN-2004-02, ECFA/04/230.
[6] M.M. Alsharoa, et al., Phys. Rev. ST. Accel. Beams 6 (2003) 081001.
[7] Y. Fukuda, et al., SuperKamiokande Collaboration, Phys. Rev. Lett. 81 (1998) 1562;
Y. Fukuda, et al., SuperKamiokande Collaboration, Phys. Rev. Lett. 82 (1999) 2644.
[8] G. Battistoni, Nucl. Phys. Proc. Suppl. B 100 (2001) 101.
[9] T. Stanev, Rapporteur's talk at the 26th International Cosmic Ray Conference, Salt Lake City, Utah, USA;
B.L. Dingus, et al. (Eds.), in: AIP Conference Proceedings, vol. 516, 2000, p. 247.
[10] T.K. Gaisser, Nucl. Phys. Proc. Suppl. B 87 (2000) 145.
[11] R. Engel, T.K. Gaisser, T. Stanev, Phys. Lett. B 472 (2000) 113.
[12] M. Honda, Nucl. Phys. B 77 (1999) 140.
[13] E. Church, et al., MiniBooNE Collaboration, nucl-ex/9706011; LAUR-97-2120, 1997.
[14] M.H. Ahn, K2K Collaboration, et al., Phys. Rev. Lett. 90 (2003) 041801.
[15] M.G. Catanesi, et al., Proposal to study hadron production for the neutrino factory and for the atmospheric neutrino flux, CERN-SPSC/99-35, 15 November 1999.
[16] L. Durieu, A. Mueller, M. Martini, Optics studies for the T9 beam Line in the CERN PS East area secondary beam facility, PAC-2001-TPAH142 Presented at IEEE Particle Accelerator Conference (PAC2001), Chicago, Illinois, 18–22 June 2001;
L. Durieu, et al., The CERN PS East area in the LHC era, in: Proceedings of PAC'97, Vancouver, 1997;
L. Durieu, O. Fernando, Design of the T9 (ATLAS/CMS) for EHNL, CERN PS/PA Note 96-38.
[17] K. Pretzl, et al., Invited talk at the International Symposium on Strangeness and Quark Matter, Crete, 1999, p. 230.
[18] C. Beugnet, XWCA/M: chambre proportionnelle multifils, CERN SL/BI-EA Note, ⟨http://sl-div-bi-ea.web.cern.ch/sl-div-bi-ea/areas/EQUIPS/xwca.htm⟩.
[19] B. Schmidt, et al., Nucl. Instr. and Meth. A 428 (1999) 299.
[20] OPERA, 3D version, Vector Fields Limited, Oxford, England.
[21] A. Lundborg, Diploma Thesis, Uppsala, 2002.

[22] ⟨www.ansoft.com⟩.
[23] Garfield, CERN Program Library W5050.
[24] C. Manaranche, et al., Photostructured coating on a voltage degrader for a Time Projection Chamber (TPC), CERN-OPEN-2002-008, 11 January 2002.
[25] G. Prior, HARP Collaboration, Nucl. Phys. Proc. Suppl. 125 (2003) 37.
[26] J. Baechler, et al., Front-End electronics for the ALICE TPC detector, CERN-ALI-98-022, 16 October 1998.
[27] ALICE Collaboration, Technical design report of the Time Projection Chamber, CERN-LHCC-2000-001, 2000.
[28] NA45/CERES Collaboration, Status of the NA45/CERES experiment and plans for 2000, CERN-SPSC-2000-009, 26 March 2000.
[29] E. Radicioni, Characterisation, calibration and performances of the TPC of the HARP experiment at the CERN PS, IEEE Trans. Nucl. Sci. NS-52 (6) (2005) 2986.
[30] S. Borghi, HARP Collaboration, Nucl. Phys. Proc. Suppl. 150 (2006) 223.
[31] P. Chimenti, Ph.D. Thesis, University of Trieste, 2006.
[32] P. Temnikov, et al., pp elastic scattering at 3 GeV/c with liquid H2 target in HARP, INFN-LNL 209-06.
[33] L. Howlett, Simulation and correction of cross talk in the HARP time projection chamber, Ph.D. Thesis, Sheffield, 2004.
[34] J.W.E. Uiterwijk, J. Panman, B. Van de Vyver, Nucl. Instr. and Meth. A 560 (2006) 317.
[35] J. Knobloch, et al., Status of the Reconstruction Algorithms for Aleph, ALEPH-Note 88-46.
[36] M.C. Morone, Evaluation of silicon sensors for the ATLAS silicon tracker, and TPC reconstruction in the HARP experiment, Ph.D. Thesis, University of Geneva, 2003.
[37] N.I. Chernov, G.A. Ososkov, Comput. Phys. Commun. 33 (1984) 329.
[38] I. Ambats, et al., Phys. Rev. D 9 (1974) 1179.
[39] V. Ammosov, et al., arXiv:hep-ex/0204022;
V. Ammosov, et al., arXiv:hep-ex/0205061.
[40] M. Bogomilov, et al., Nucl. Instr. and Meth. A 508 (2003) 152.
[41] G. Barr, et al., Nucl. Instr. and Meth. A 533 (2004) 214.
[42] HARP Collaboration, Particle identification algorithms for the HARP forward spectrometer, Nucl. Instr. and Meth. A, in press.
[43] M.G. Catanesi, et al., HARP Collaboration, Nucl. Phys. B 732 (2006) 1 (arXiv:hep-ex/0510039).
[44] H. Wind, J. Comput. Phys. 2 (1968) 274.
[45] J. Altegoer, et al., NOMAD Collaboration, Nucl. Instr. and Meth. A 404 (1998) 96.
[46] M. Anfreville, et al., Nucl. Instr. and Meth. A 481 (2002) 339.
[47] A. Cervera, J.J. Gomez-Cadenas, J.A. Hernando, Nucl. Instr. and Meth. A 534 (2004) 180.
[48] M. Bonesini, et al., IEEE Trans. Nucl. Sci. NS-50 (2003) 1053.
[49] M. Baldo-Ceolin, et al., Nucl. Instr. and Meth. A 532 (2004) 548.
[50] S. Buontempo, et al., Nucl. Instr. and Meth. A 349 (1994) 70;
E. Di Capua, et al., Nucl. Instr. and Meth. A 378 (1996) 221.
[51] M.G. Van Beuzekom, et al., CHORUS Collaboration, Nucl. Instr. and Meth. A 427 (1999) 587.
[52] H. Beker, et al., The ALICE DATE: the benefits of using hardware and software industry standards in a real-time environment, CERN-ALICE-PUB-97-12.
[53] ALICE Technical Proposal for a Large Ion Collider Experiment at the CERN LHC, CERN/LHCC/95-71.
[54] R. Brun, F. Rademakers, Nucl. Instr. and Meth. A 389 (1997) 81.
[55] ⟨www.objectivity.com⟩.
[56] A. Dell'Acqua, et al., GEANT4 Collaboration, Nucl. Instr. and Meth. A 506 (2003) 250.
[57] G. Barrand, et al., Comput. Phys. Commun. 140 (2001) 45.
[58] ⟨http://geant4.web.cern.ch/geant4/geant4.html⟩.