A Detector to search for $K^+ \to \pi^+\nu\overline{\nu}$

M.S. Atiya, I-H. Chiang, J.S. Frank, E. Garber, J.S. Haggerty, M. Ito[a],T.F. Kycia, K.K. Li, L.S. Littenberg, C. Ng, I. Polk, R. Rosenka, A.K. Sambamurti, A.J. Stevens, R.C. Strand and C. Zein

Brookhaven National Laboratory, Upton, N.Y. 11973

D.S. Akerib[b], H. Edwards, W.C. Louis[c], D.R. Marlow, P.D. Meyers, M.A. Selen[d], F.C. Shoemaker and A.J.S. Smith

Physics Department, Princeton University, Princeton, NJ 08544

S. Ahmad[e], E.W. Blackmore, D.A. Bryman, Chen Lejun, J.V. Cresswell, L. Felawka, N. Khan, P. Kitching, A. Konaka, Y. Kuno, J.A. Macdonald, T. Numao, P. Padley, J-M. Poutissou, R. Poutissou, J. Roy, R. Soluk and A.S. Turcot

TRIUMF, Vancouver, BC, Canada V6T 2A3

Abstract

The detector built for Experiment 787 to measure the decay $K^+{\to}\pi^+\nu\overline{\nu}$ at Brookhaven National Laboratory is described.

Published in Nuclear Instr. and Methods **A321** (1992) 129 – 151

I. INTRODUCTION

Higher-order predictions of the Standard Model (SM) [1] can be tested in an unambiguous way [2] by studying the decay $K^+{\to}\pi^+\nu\overline{\nu}$ which is suppressed in first order by the GIM [3] cancellation of flavor changing neutral currents. In the context of the SM assuming three generations, the branching ratio for $K^+ \to \pi^+ \nu\overline{\nu}$ is predicted to be in the range $(1-4) \times 10^{-10}$ [4]. Before Experiment 787 at Brookhaven National Laboratory was undertaken the experimental upper limit for $K^+ \to \pi^+ \nu\overline{\nu}$ was $B(K^+ \to \pi^+ \nu\overline{\nu}) < 1.4 \times 10^{-7}$ [5]. Based on early data from E787, an improved limit $B(K^+ \to \pi^+ \nu\overline{\nu}) < 3.4 \times 10^{-8}$ has been reported [6]. New limits for other rare decays including $K^+ \to \pi^+ \mu^+ \mu^-$ [7] , $K^+ \to \pi^+ \gamma\gamma$ [8] and $\pi^0 \to \nu\overline{\nu}$ [9] have also been obtained.

The main goal of Experiment 787 is to observe $K^+ \to \pi^+ \nu\overline{\nu}$ if it occurs within the range of predictions of the SM. A signal in the predicted range will place constraints on parameters including the top quark mass and certain of the Cabbibo-Kobayashi-Maskawa matrix elements which describe the weak mixing among the three generations of quark mass eigenstates. Observation of $K^+ \to \pi^+ \nu\overline{\nu}$ well above the prediction would indicate new physics beyond the SM such as the existence of new generations of leptons or quarks, or of one or more new light, weakly-interacting neutral particles.

The signature of $K^+ \to \pi^+ \nu\overline{\nu}$ is a single unaccompanied π^+ emitted by a stopped K^+. Fig. 1 shows the expected charged-particle spectra of the range and momentum for $K^+ \to \pi^+ \nu\overline{\nu}$ along with the other decay modes which could be sources of background.

FIGURES

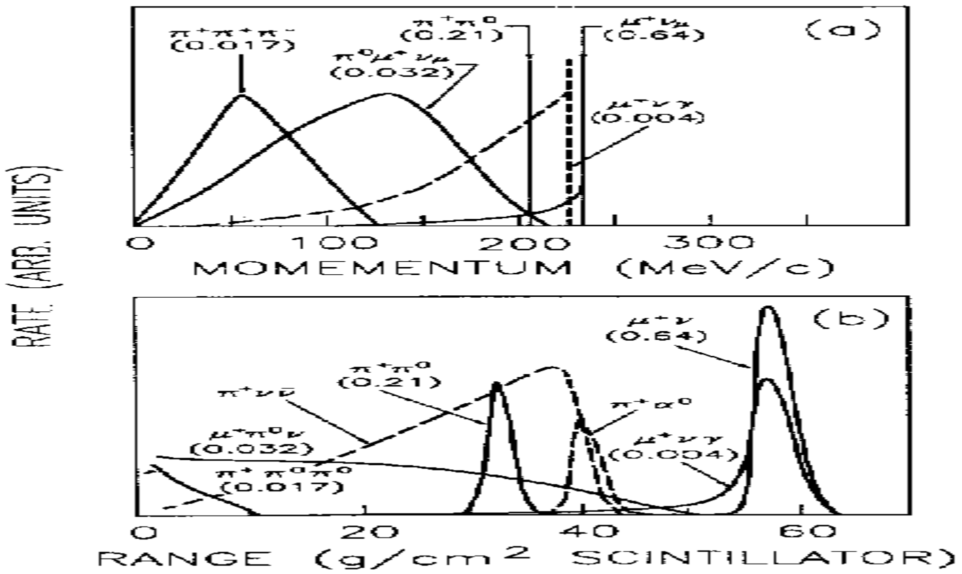

FIG. 1. Charged-particle spectra in (a) momentum and (b) range (including finite resolution) for $K^+ \to \pi^+ \nu\overline{\nu}$ and the major sources of backgrounds from Ref [5].

The decays $K^+ \to \mu^+ \nu$ (γ) $(K_{\mu 2})$ and $K^+ \to \pi^+ \pi^0$ $(K_{\pi 2})$ are major potential background sources with branching ratios of 0.635 and 0.212, respectively. It is possible to reduce the impact of $K_{\pi 2}$, at the cost of acceptance, by restricting the search to the 20% of the total spectrum in the kinematic region above the monoenergetic peak of $K_{\pi 2}$ (momentum $P_{\pi^+} > 205$ MeV/c) and having efficient photon detection. In this case only the effects of finite kinematic resolution could allow $K_{\pi 2}$ to contribute events in the signal region. Backgrounds with muons can be managed with a high level of π/μ discrimination.

The critical characteristics of the detector, in addition to high sensitivity, are unambiguous π/μ particle identification, efficient photon detection, and good kinematic resolution. These features are crucial also for rejection of the other potential background decay modes appearing in the kinematic region of interest, such as $K^+ \to \mu^+ \pi^0 \nu$ $(K_{\mu 3})$ and $K^+ \to \pi^+ \gamma\gamma$. The detector also must be capable of rejecting beam-related backgrounds from processes involving kaon charge exchange or pion scattering.

II. DETECTOR

A. Detector Overview

Schematic side and end views of the detector are shown in Fig 2.

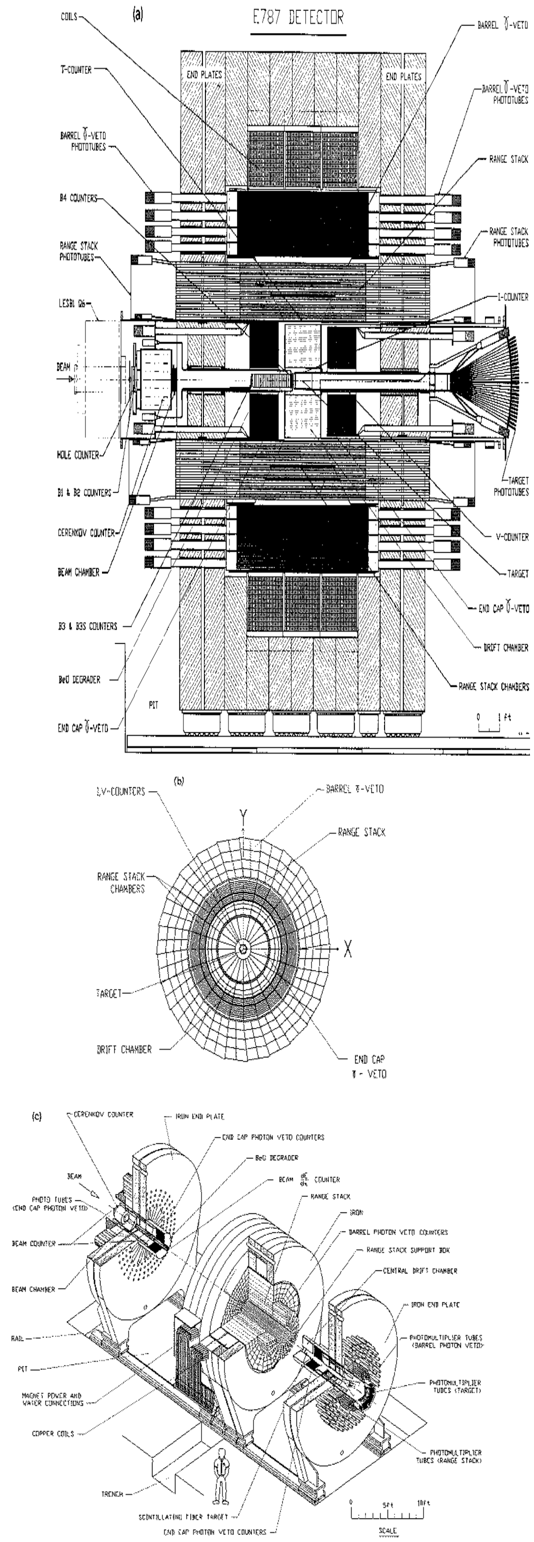

FIG. 2. Schematic views of the detector from the (a) side and (b) end in an expanded scale showing only the active elements. (c) Isometric view with the three magnet sections separated.

Incoming kaons are stopped in the target located at the center of the detector enclosed in a solenoid magnet which produces a uniform 1.0-T field along the beam direction. Pions from K^+ decay are momentum analyzed as they traverse the cylindrical drift chamber and are stopped in layers of scintillators (range stack) in order to measure the range and kinetic energy. The solid angle acceptance for π^+ detection is $\sim 2\pi$ sr. Surrounding the range stack and drift chamber is the photon detection system covering nearly 4π sr, comprised of a cylindrical barrel section and two endcap sections of electromagnetic shower detectors. The overall detector function is illustrated by the $K_{\pi 2}$ event shown in Fig. 3 in which the pion has been tracked out of the target, through the drift chamber, and stopped in the range stack, and the two photons from π^0 decay are detected in the range-stack and barrel-photon detectors.

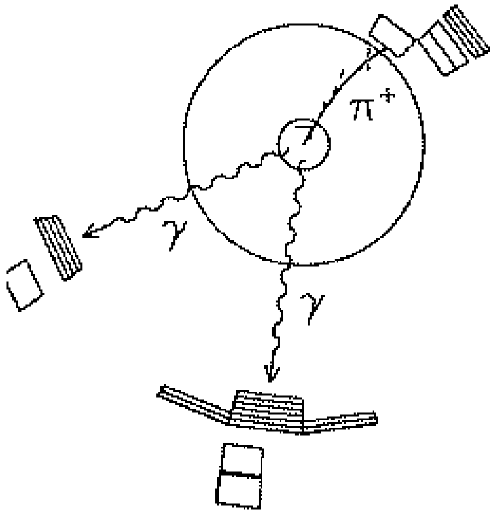

FIG. 3. Reconstructed $K_{\pi 2}$ event showing the detector elements hit by the π^+ and the two γs from the π^0.

Particle identification is obtained by exploiting the correlations among the range, momentum and kinetic energy of charged decay products in order to distinguish pions from muons. This kinematic technique is complemented by observing the characteristic decay sequence of $\pi^+ \to \mu^+\nu$, followed by $\mu^+ \to e^+\nu\overline{\nu}$ ($\pi^+ \to \mu^+ \to e^+$) in the range-stack photomultiplier signals from the pion stopping counter using transient digitizers.

There are also possible backgrounds related to the beam, such as beam pions scattering into the detector following K^+ stops, or K^+ charge-exchange reactions (K^+, K^0) resulting in decays such as $K^0 \to \pi^+ e^- \nu$, where the second charged particle is missed. In order to deal with these, the beam counters are capable of efficient K^+ identification and detection of secondary incident particles from the beam line. An active highly-segmented target system allows tracking of the stopping K^+s and of the decay π^+s near the decay vertex.

The individual subsystems are described in detail in the following sections.

B. Magnet

The solenoid magnet produces the axial field required for momentum determination from the central-drift-chamber tracks, and serves as the principal supporting structure for all detector components. To avoid inert material which would interfere with photon detection, and range and energy measurements, the magnet is made large enough to contain all the active components within the yoke and coils. The inner volume is 2.22-m long and 2.96-m in diameter.

The steel structure of the magnet is made from eleven circular plates 305-mm thick and 5.0-m in diameter. The magnet is divided into three sections (see Fig. 2c) supported independently on sets of rollers on tracks so that the sections can be moved apart axially for access to the interior volume. The center section carries the water-cooled copper coils made up of 18 sections of 22 turns and 2 cooling circuits each. It also carries the stainless-steel support frames for the barrel-photon-veto and range-stack assemblies through four 25.4-mm-thick aluminum plates spaced between the coil sections. The steel yoke is about 560-mm thick and the coil consumes 1.1 MW at 4400 A to produce the 1.0-T central field. The coil temperature rise is about 11°C at full field, and the detector interior temperature is maintained at ≤ 28°C. The interior volume is continuously flushed with dry N_2.

Each of the two end sections (endplates) is made of two 305-mm-thick steel plates from the Space Radiation Effects Laboratory (SREL) cyclotron separated by a 25.4-mm gap to minimize the fringe field at the upstream and downstream faces where the photomultiplier tubes (PMT) are mounted. Each endplate is penetrated by 192 holes for lightguides from the barrel photon veto, and has a 24-spoke structure in the radial region of the range stack also to give access for light collection. The amount of steel removed from the endplate is a compromise between the desired low fringe field and high light collection efficiency. The endplate retains about half its azimuthal steel for flux return. The center part of each endplate is a steel endplug to which the endcap sections of the photon veto are mounted. The upstream endplug also supports the beam instrumentation, and the downstream endplug supports the target. The central drift chamber is supported on a steel cylinder fitted between the downstream spoke structure and the endplug. The endplug with the endcap attached can be removed independently to give access to the ends of the drift chamber using a special fixture mounted on the same tracks as the magnet.

Before the detector elements were installed, the magnetic field was mapped using an array of Hall probes referenced to an NMR probe to establish the magnitude and uniformity in the central volume to be occupied by the drift chamber. The measurement showed the field to be uniform to $\pm 1\%$ overall, and $\pm 0.6\%$ excluding the regions not occupied by tracks of interest, leading to an estimated 0.2% contribution to the momentum resolution. The fringe fields at the PMTs are typically 2×10^{-3} T at full field and adequate shielding is provided by μ-metal shields.

The magnet is powered by two 720-KW, 4800-A power supplies connected in series with an LC network to filter the output. The field is monitored with a calibrated Hall probe mounted in the magnet and recorded during running by the data acquisition system.

C. Beam Counters

An 800-MeV/c kaon beam from the Low Energy mass-Separated Beam (LESB–1) at the Alternating Gradient Synchrotron (AGS) is slowed by 535 mm of BeO degrader and stopped in the scintillator target. (BeO combines high density with low atomic number to minimize the effects of multiple scattering.) The intensity of the kaon beam incident on the detector has varied up to 2×10^6 K^+ per AGS spill (typically 1.6-s duration with 6×10^{12} 24-GeV protons striking the 89-mm-long Pt production target) producing up to 4.5×10^5 K^+ stopped in the target. The incident ratio of $K^+ : \pi^+ : p$ was about $1 : 2 : 1$. The beam counter system consists of a Čerenkov counter, several scintillation counters, and a 3-plane multi-wire proportional chamber, shown schematically in Fig. 2.

The kaons and pions in the beam are independently tagged by the acrylic-radiator Čerenkov counter, using the scheme devised by Fitch [10], and shown schematically in Fig. 4.

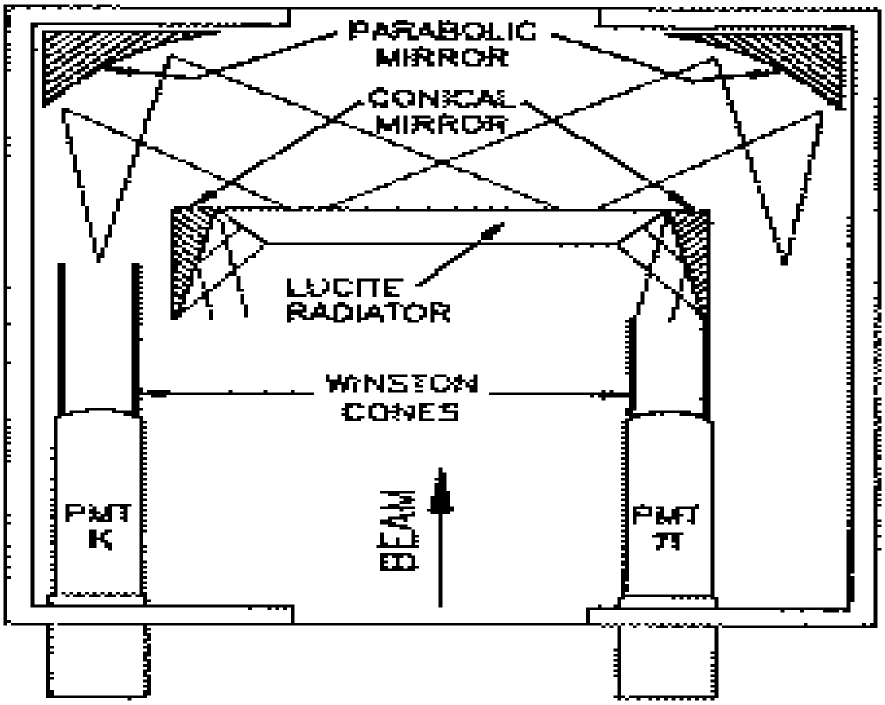

FIG. 4. Schematic showing the layout of the radiator, mirrors and phototubes of the beam Čerenkov counter.

The light from kaons is refracted out of the downstream face of the 150-mm-diameter, 25.4-mm-thick radiator with 97% efficiency, since at the beam momentum of 800 MeV/c the Čerenkov angle is close to Brewster's angle, and the light is polarized in the plane containing the beam. A mirror which is a paraboloid of revolution then focuses this light onto a ring of ten 50-mm-diameter EMI 9954KB PMTs each equipped with a Winston cone [11] cast of black epoxy resin and aluminized on the reflecting surfaces. The light from pions, having a larger Čerenkov angle, is totally reflected by the flat surfaces of the radiator, and exits nearly normally from the conical edge. It is then reflected by a conical mirror into the Winston cones of a second ring of ten PMTs. Both mirrors are made of acrylic with the active surfaces aluminized and the other surfaces painted black.

Each ring of PMTs is connected to a LRS 4413 CAMAC discriminator whose multiplicity output signal (100 mV per fired channel) is discriminated to obtain the tagging signal. The efficiency of the kaon channel is about 98% for a PMT multiplicity of seven. The pion channel inefficiency is 10^{-4} for a multiplicity of five PMTs. The multiplicity outputs of the discriminators are also fed to analog-digital converter (ADC), time-digital converter (TDC), and transient digitizer (TD) channels to be recorded with each event.

Upstream of the Čerenkov counter is a set of scintillator counters designed to monitor the incident beam. The $B1, B2$ counter assembly consists of two hodoscopes, both for the horizontal profile. $B1$ has elements 40-mm high and 35, 30, 20, 20, 20, 55-mm wide going from left to right as seen by the beam. $B2$ is the same with the order of elements reversed so that each counter intercepts the gaps between fingers of the other. In the vertical direction two counters, 12.7-mm high $\times$ 203-mm wide, are mounted 7.5 mm above and below the centre of the hodoscope and intercept only the small tails of the vertical beam distribution. Two L-shaped counters are configured to detect any halo of beam particles that miss the hodoscopes or the vertical-tail counters. Each scintillator counter element is viewed through an adiabatic ultra-violet-transmitting (UVT) acrylic lightguide by one or both halves of a 25.4-mm-square Hamamatsu R1548 dual PMT. All counters are 3-mm thick.

Immediately downstream of the Čerenkov counter is the beam multi-wire proportional chamber (BMPC) used to monitor the beam profile. It consists of 3 planes (X, U, V) of 12-μm-diameter Au-plated tungsten anode sense wires strung 1.27-mm apart and read out in pairs. High voltage is applied separately for each layer to the cathode plane consisting of a film of 25-μm-thick aluminized Mylar 3.18-mm away from the anode plane. The X-plane has 72 anode pairs running vertically and the U- and V-planes each have 60 pairs at $\pm 45°$

to the vertical. Each pair has an on-board preamplifier which drives the signal to a post-amplifier and discriminator circuit whose output feeds a LRS 1879 Fastbus pipeline TDC channel. The chamber operates with an 80 : 20 gas mixture of CF_4 and isobutane for fast response. At a high voltage of 3500 V each plane is typically >97% efficient.

A lead collimator downstream of the BMPC restricts the beam in the horizontal direction. It is 660-mm long and has a gap of 63.5 mm in the horizontal direction. Two counters, $B3$ and $B3S$, are located at the upstream face of the degrader, for beam tuning and diagnostic purposes. $B3$ is 76 mm $\times$ 64 mm and $B3S$ is 51 mm $\times$ 38 mm.

Three layers of $B4$ counters are sandwiched between the degrader and the upstream face of the target. The first two layers are 4-finger hodoscopes (each finger 25.4 mm $\times$ 101.6 mm) $B4Y$ and $B4X$. The third layer, $B4T$ is a single counter which has the same hexagonal shape as the target. It is 110-mm wide between parallel sides and extends 3 mm outside the target on all sides. The energy loss (dE/dx) in any of the 6-mm-thick B4 scintillators is used as part of the identification of incoming kaons.

Dispersion-cancelling circuits on all beam scintillator PMT signals are used to compensate for signal broadening due to the loss of high-frequency components on the $\sim$ 50-m-long coaxial cables from the PMTs to the counting room. There, the signals feed $\times 10$ amplifiers which in turn feed ADC and TD channels. A second output is discriminated for a TDC input. Normally, only $B4T$ and the kaon Čerenkov counter are included in the beam information sent to the trigger.

D. Target

For background rejection, it is desirable to locate the vertex of the K^+ decay to <5 mm transverse to the beam and to identify other charged particles in the target in time coincidence with either the incident K^+ or with the outgoing π^+. Inert material in the target was minimized to achieve good particle identification from energy-loss measurements of the stopping kaon and of charged decay products, and to maintain good photon detection efficiency. To meet these specifications plastic scintillating fibers were chosen [12]. The good time resolution necessary both for background rejection and for triggering purposes was obtained by using fast PMTs.

The 2-mm-diameter fibers were drawn in 200-m lengths by Optectron Corp. (Les Ulis, France) with a core of S-101 styrene scintillator (refractive index $n = 1.59$), and a vinyl-acetate ($n = 1.46$) cladding. In the core, butyl-PBD and POPOP shift the wavelength peak to 440 nm. The 25-μm cladding prevents damage to the clad-core interface where total internal reflection up to 6° captures 8% of the light. A 1000-Å Al coating was sputtered onto the fibers by Vacotec (Reconvilier, Switzerland) to provide a protective layer for the fragile vinyl-acetate cladding, and to prevent possible crosstalk of light from one fiber to another.

The fibers were cut into 3.12-m lengths and one end was polished for testing of light output using a ^{106}Ru source. Similar tests were repeated at each stage of subsequent fabrication to maintain quality. The typical fiber had an attenuation length of about 2 m but variations of more than a factor of two were observed. A bundle of six fibers was epoxied together in a triangular aluminum mold leaving 1 m free of epoxy at one end for routing to a PMT. Fibers were matched for grouping so that variation in light output from fiber to fiber in each bundle was kept to <10%. The sizes of the triangular bundles were kept identical to a tolerance of < 75 μm to enable precise interlocking to form the target. The resulting overall target is 75% active, the inert fraction being mostly epoxy.

An Al mirror was vacuum evaporated onto the polished upstream ends of the triangles in order to reflect scintillation light back towards the PMT mounted at the downstream end. This increased the response for light generated close to the mirror by a factor of 1.6. The Al mirror coating was about 1% transmitting in order to admit light from a light-emitting diode (LED) calibration system. The layout of the target triangles, looking downstream, is shown in Fig. 5.

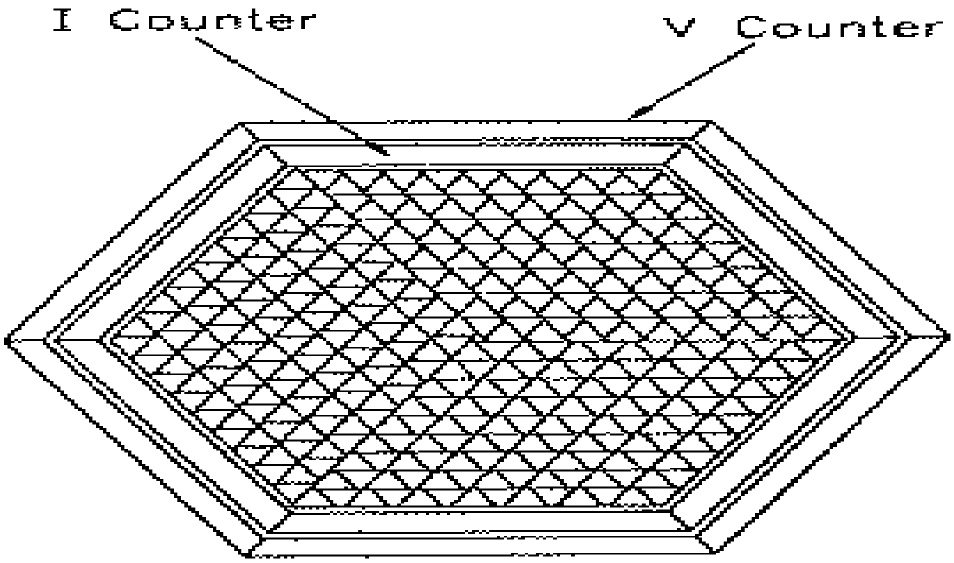

FIG. 5. Schematic showing the arrangement of the scintillator-fiber triangles and the surrounding I- and V-counters.

The unepoxied downstream region of each bundle was fitted with a black teflon sleeve, a brass mounting collar, and an UVT acrylic light-mixer block glued onto the six polished fiber ends for coupling to the PMT. At the target center is a single fiber coupled to its own PMT.

Each of the 379 bundles is coupled to a 10-mm-diameter Hamamatsu R1635-02 PMT in a miniature spring-loaded assembly. The PMTs have high-gain first dynodes for good one-photoelectron sensitivity. The relative quantum efficiency was determined for each PMT and matched to the brightness of each triangle. After the target was installed in the detector, average signals from muons in $K_{\mu 2}$ triggers were used to set gain values. Typically, the PMT readout yields one photoelectron per mm path for minimum ionizing particles.

Each PMT signal is connected to a $\times 10$ amplifier feeding a splitter-filter box. From there, signals are distributed to the ADC and TDC systems, and to the energy-sum units used in the fast trigger to require an energy-sum pulse of over 5 MeV coincident with the kaon beam logic. For each triangle, the TDC data are recorded for ± 500 ns with respect to the event. A signal for the number of triangles hit is fed to the trigger system as an estimator of the range of the decay particle in the target in order to reject $K_{\mu 2}$ decays. The threshold is set below the pulse height from a single photoelectron. Fig. 6a shows an event display of the target ADC information,

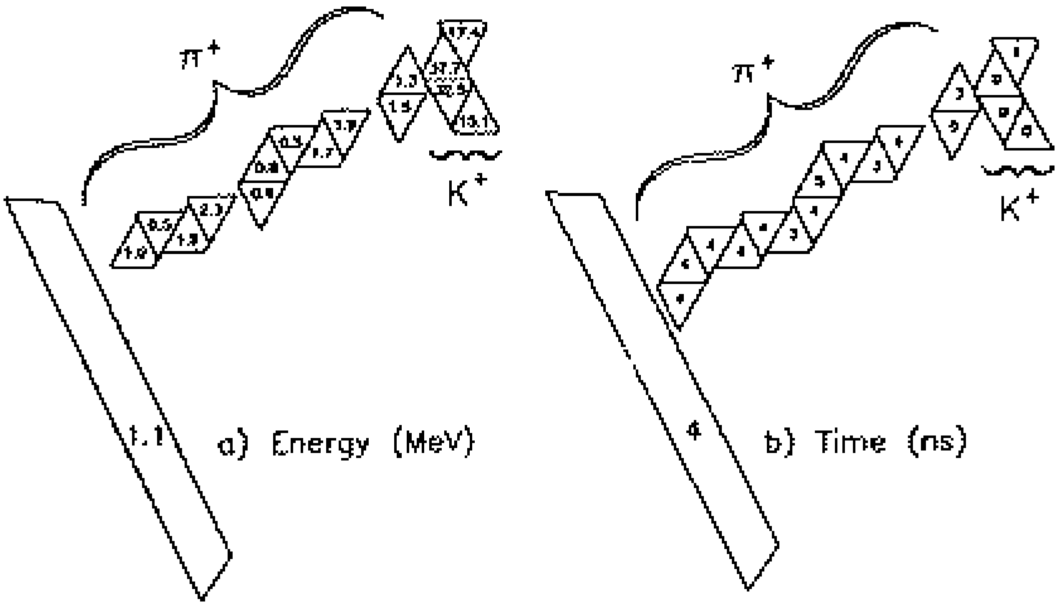

FIG. 6. Detail display showing the target elements struck by a stopping K^+ and the π^+ from its decay. (a) shows the calibrated energy in MeV deposited in each element and (b) shows the time of the hit in ns.

looking downstream, for a typical $K_{\pi 2}$ event. The TDC information for the first hit in the time window ± 50 ns for the same event is shown in Fig. 6b.

The time accuracy is limited by the small number of photoelectrons seen in each triangle and the intrinsic time response of the scintillation material. The leading edge time from each triangle, obtained from the first photoelectron, can be weighted by the energy to arrive at the best estimator for the kaon or pion time. The resulting time resolution from $K_{\mu 2}$ calibrations is 700 ps. The time information is crucial to efficient background rejection because only hits that are coincident in time with either the K or π need be considered. Tracks that occur at other times are indicative of accidentals from other beam particles and need not necessarily cause the event to be rejected. The time difference between the K and

π is used to eliminate in-flight decays.

In off-line analysis a pattern recognition program sorts the hit triangles into clusters of stopped kaons, charged decay products connected with a stopped kaon, separate photon conversions, and random backgrounds by using the target data for position, energy, and time. Candidate kaons or decay products are identified, to first order, using the energy-loss information which is corrected for the $\sim$25% inert material. K decay products within the detector tracking acceptance are close to minimum ionizing and travel within 30° of the perpendicular to the fibers and therefore deposit <2 MeV per triangle. The stopping kaons, however, travel parallel to the fibers with a large dE/dx and typically deposit >5 MeV per triangle. Each kaon excites up to 5 triangles, and the pion excites up to 25 triangles, depending upon the origin of the vertex in the target. Energy and path-length information from the target are used to correct the total energy and range determination for charged decay particles. Hit triangles that are disconnected from the kaon or pion tracks but still in time with the event can be used to identify photon conversions from π^0 decay in the target.

The spacial resolution in the xy plane is 1.8 mm (rms) transverse to a track in the target, consistent with the dimensions of the triangles. The resolution on the vertex position longitudinal to a track is 4.0 mm (rms). The longitudinal resolution is dominated by the size of the cluster of triangles hit by the stopping K, since the pion energy is not resolved within the stopping cluster. Using the vertex determination, a correction for the energy loss of the pion in the target is applied to the energy, range and momentum measurements obtained from the range stack and drift chamber.

In addition to the hexagonal array of scintillator fibers, the target assembly includes two sets of six conventional plastic scintillators which surround the target as shown in Fig. 5. The first set consists of the I-counters, 240-mm long $\times$ 6.4-mm thick, which define the fiducial stopping region by tagging charged decay products after a kaon stop and before they enter the drift chamber. Including the I-counters, the target assembly in the fiducial region is 127-mm across, from flat to flat. The second set, the V-counters (1960-mm long $\times$ 5-mm thick), overlaps the downstream end of the I-counters by 6 mm and helps to distinguish hits downstream of the fiducial region of the target fibers which serves to veto various background processes. The I- and V-counters are viewed through UVT lightguides by EMI 9954KB PMTs.

E. Drift Chamber

The primary functions of the cylindrical drift chamber [13], which subtends a solid angle of 2π sr for kaon decay products of momentum $P > 150$ MeV/c, are to provide a momentum measurement in the 1-T magnetic field and to achieve good tracking between the target and the range stack. Momentum resolution of $\sim 2.5\%$ (rms) is required to obtain the necessary level of background rejection related to $K_{\pi 2}$ and $K_{\mu 2}$ decays. The chamber is required to have a short memory time in order to operate successfully in the presence of high rates of kaons and pions in the beam.

The chamber wires are arranged in five superlayers of 36, 40, 50, 60, and 70 cells, respectively. The wires of layers 1, 3 and 5 are strung axially, and those of layers 2 and 4 are strung between feedthroughs offset by one cell, corresponding to stereo angles of approximately $\pm 3.5°$ to enable axial position measurements. The configuration of wires and wire diameters within a unit cell is shown in Fig. 7a.

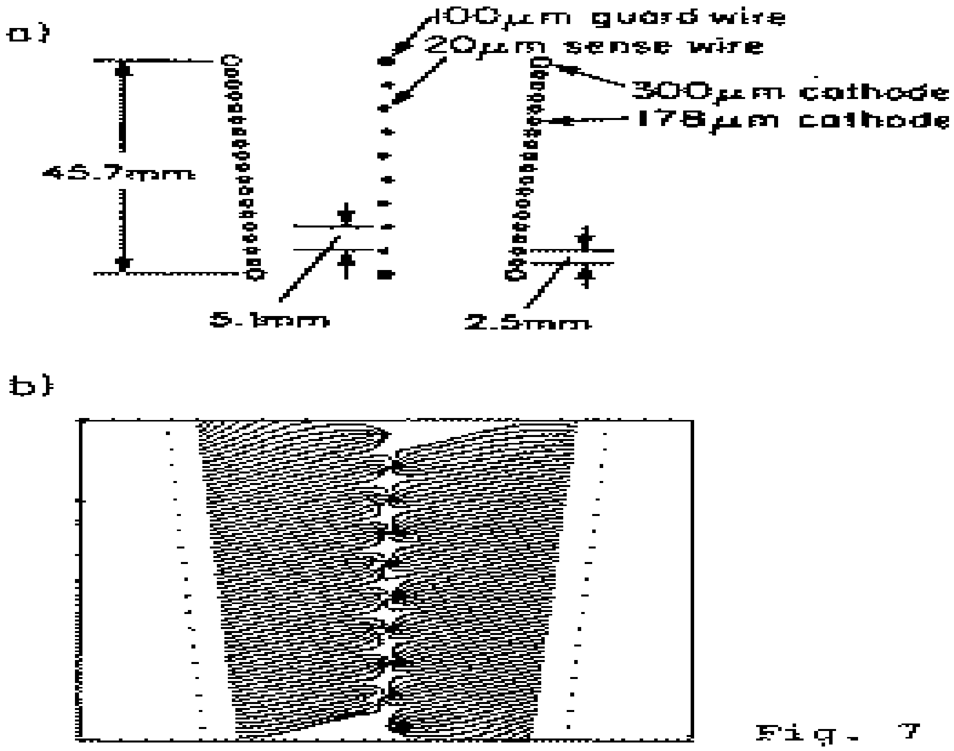

Fig. 7

FIG. 7. (a) Schematic showing the wire geometry in a drift cell. (b) Trajectories of ionization electrons to the sense wires in a 1.0-T magnetic field.

Cells within a layer are identical, but cell size varies from layer to layer between 12-mm and 17-mm half width. The cathode-wire plane, consisting of 19 Be-Cu wires strung 2.54-mm apart to a tension of 150 g , is shared by neighboring cells. The sense-wire plane has two Be-Cu end guard wires and eight 20-μm-diameter gold-plated tungsten anode wires spaced 5.08 mm apart at a tension of 50 g and staggered ± 254 μm from the midplane to resolve the left-right ambiguity locally within the cell. The end wires on both anode and cathode planes are of larger diameter as shown in Fig. 7a to optimize electric field uniformity and to maintain surface fields <20 kV cm^{-1} on the cathode wires when those on the sense wires are about 250 kV cm^{-1}. This arrangement without focussing or potential wires in the anode plane provided a high electric field in the drift region of $\sim$ 1.5 kV cm^{-1}, and reduced the Lorentz angle. To provide the correct potentials at the cylindrical boundaries of the chamber, thin copper-plated Mylar foils were attached to the graphite-epoxy support cylinders and set at the appropriate voltages. The trajectories calculated for a typical cell, of ionization electrons in a 1-T magnetic field are shown in Fig. 7b. Only the six middle anode wires are read out in order to avoid the cell-end distortions in the electric field and to reduce the effects of the Lorentz angle which is about 25° for a drift velocity of 5 cm μs^{-1}. The chamber is operated at atmospheric pressure with a gas mixture of Ar : $C_2H_6 = 50 : 50$. The argon fraction is bubbled through ethanol at 0° C.

The chamber occupies the radial region between 95 mm and 432 mm with a total length, including preamplifiers, connectors, and cables of 650 mm. The active volume, 508-mm long, is enclosed between 9.5-mm-thick precision-machined Al endplates carrying the slots to receive the feedthroughs which position and support the wires. Each feedthrough is made up of a machined comb of Vespel epoxied into an injection-molded Ryton insert which is epoxied into the slot. The overall mechanical tolerance for the location of each wire was measured to be $\sim$20 μm. The endplates are supported by inner (80-mg cm^{-2} thick) and outer (94-mg cm^{-2} thick) cylinders made of graphite-fiber epoxy, which was built up of four (inner) and five (outer) pre-impregnated layers with fiber orientation alternately axial and $\pm$45°. The endplates were prestressed prior to stringing with a system of temporary tie-rods which were adjusted and then removed as wires were installed, taking up the tension load. The total endplate loading due to wire tensions is 880 kg. The outer cylinder was installed after stringing was complete. During stringing, its compression load was taken up by a set of temporary posts. Subsequent replacement of individual broken or faulty wires has been done with the outer cylinder in place. The chamber is supported in the detector by a 3.2-mm-thick aluminum cylinder which extends from the magnet pole and is attached at the circumference of the downstream endplate.

Contact to each wire for high voltage (HV), ground, or onboard preamplifier, is via a printed edge connector card fitted into the feedthrough, to which the wire was soldered after tensioning. For the cathodes, HV is distributed to one third of a layer from one of 15 Bertan 1755N HV power supplies. For each of the 15 sections of HV distribution the voltage is graded in a resistor chain to compensate for the variation in cell width with radius in order to form a uniform drift field throughout the cell. The external regions of both endplates, on which the HV distribution and preamplifiers are mounted, are sealed and dry N_2 is circulated to control temperature and to eliminate HV breakdown due to atmospheric humidity.

Each sense-wire plane has six preamplifiers mounted on a card plugged onto the downstream feedthrough and grounded locally to the endplate. A typical preamplifier gain is 10 mV μA^{-1} with a risetime of 3 ns and power dissipation of 24 mW. The positive-ion-induced crosstalk of signals between adjacent wires of $\sim$15% is reduced to $< 5\%$ by a compensation resistor network which splits the input charge at the front ends of the preamplifiers. Each preamplifier drives 34 m of 50-Ω coaxial cable which degrades the signals to about half amplitude and risetime of about 7 ns when they reach the post-amplifier in the counting house. All cables, including signal, HV, preamplifier power, and test pulse, exit the magnet and go through a bulkhead connection at the downstream end of the detector.

The post-amplifier-discriminator circuits are housed, 24 per module, in a DIN 41494 standard sub-rack system (Euro-crates). The input signals are decoupled from ground by transformers and split for separate analog and discriminator processing. The gains of the amplifiers are 7–10 for the analog output and 100 for the discriminator path. Pole-zero cancellation is employed to minimize fall time and reduce the pulse tail, due to positive ions and to the 34-m cable, to a level of $< 2\%$. The discriminator produces a time-over-threshold output with a 10-ns minimum width and feeds an LRS 1879 Fastbus pipeline TDC for drift time (up to 350 ns) measurement.

The drift velocity v, Lorentz angle α, and time offset (pedestal) are determined iteratively by fitting tracks in the chamber. For v, zero-magnetic-field data were used. Position resolution obtained from $K_{\pi 2}$ data, shown by the residuals from track fits in Fig. 8,

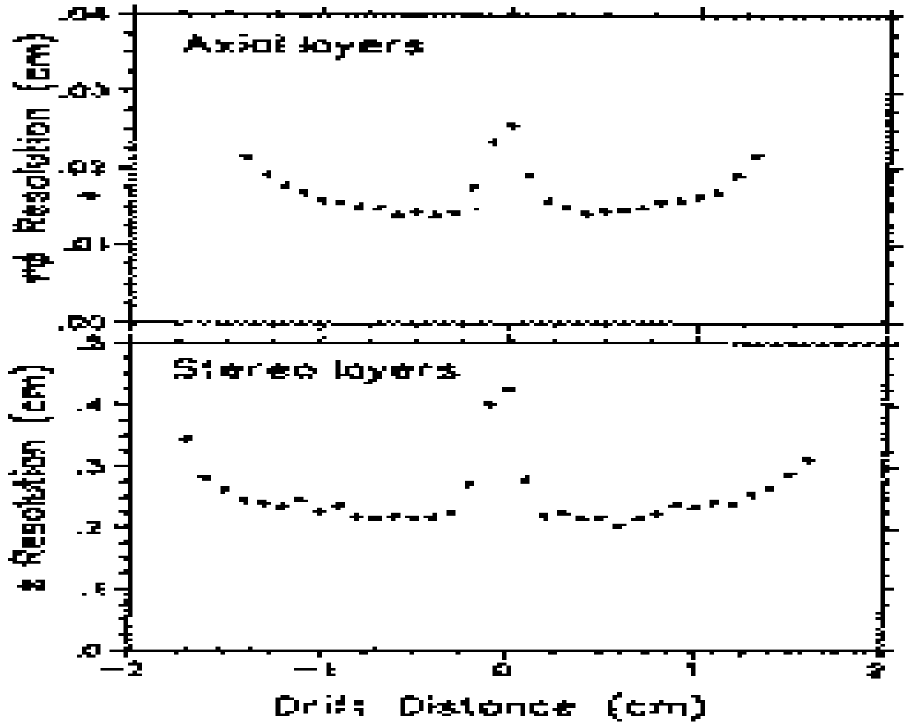

angle

FIG. 8. Typical local position resolution in the drift chamber as a function of drift distance for stereo and axial layers.

varies with drift distance and layer between 130 and 250 μm for the axial layers. The z (axial) resolution obtained from the stereo layers is between 2.2 and 4.2 mm. The momentum resolution, after correcting for the energy loss in the target, is shown in Fig. 9.

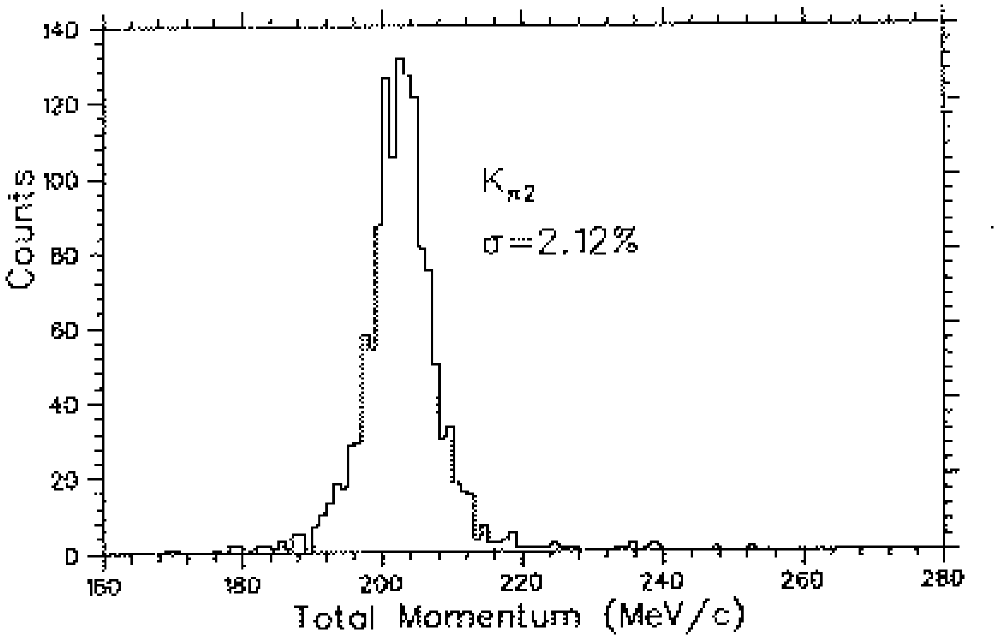

FIG. 9. Momentum for $K_{\pi 2}$ events corrected for the energy loss in the target system.

F. Range stack

The energy, range, and decay sequence of charged particles are measured in the range stack. Charged particles which emerge from the drift chamber outer circumference are stopped in a cylindrical array of layers of plastic scintillators segmented azimuthally into 24 sectors. Each range-stack sector consists of 20 layers of 1.82-m-long $\times$ 19.05-mm-thick plastic scintillation counters which were fabricated of BC408 by Bicron Corp. (Newbury OH, USA). An additional innermost layer, 520-mm long $\times$ 6.35-mm thick (T counter), is used in the trigger to define the $\sim 2\pi$ acceptance of the range stack as viewed from the target, coinciding roughly with that of the drift chamber. The scintillator material was chosen for its good light output, intrinsic short pulse width ($\sigma \sim 1$ ns), and reasonable attenuation length. Two layers of range-stack multi-wire proportional chambers (RSPC) [14] are located following layers 10 and 14, as shown in Fig. 10 in order to obtain tracking information within the range stack.

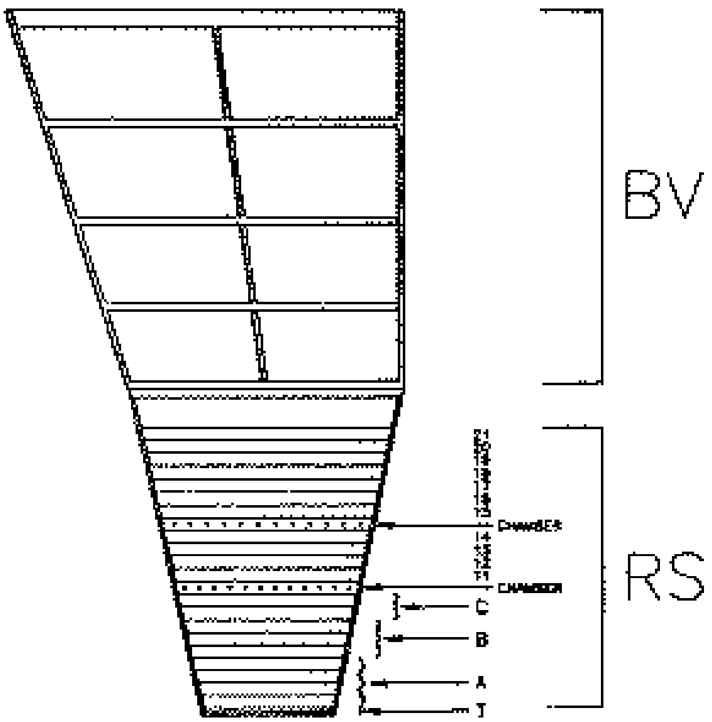

FIG. 10. Schematic end view of one range-stack sector with the corresponding barrel photon veto modules at the outer radius (top). The A, B, C superlayers and the positions of the proportional chambers are indicated.

Overall, the range-stack sectors fill the radial region between 451 mm and 896 mm.

The range-stack system is supported in the magnet by a pair of welded stainless-steel web-like frames at the upstream and downstream ends of the magnet volume, fastened to the center section of magnet yoke indirectly through the support frame for the barrel photon veto. Each frame extends towards the center plane of the detector up to the axial extent of the drift chamber, approximately, at the inner diameter, and tapers back in a conical shape towards the outer diameter. This design provides adequate support without introducing inert material into the fiducial region.

The twenty 19-mm-thick layers form the stopping region of the range stack. The nine innermost counters are grouped into superlayers A (4-fold), B (3-fold), and C (2-fold), as shown in Fig. 10, each viewed by a single PMT at each end. This arrangement simplified the problem of fitting the PMTs into the limited space, but at the cost of range resolution for charged particles that stop in this region. In the strategy to restrict the search for $K^+\to\pi^+\nu\overline{\nu}$ to the kinematic region above the $K_{\pi 2}$ peak, where the π^+s of interest have a range in scintillator of > 25 g cm^{-2}, range resolution is not critical in the inner ~ 18 cm. The remaining eleven layers are each viewed by a PMT at each end, for a total of 720 PMTs on the entire range stack, including the T-layer.

Scintillation light is brought out of the magnet at both ends of the detector by UVT acrylic lightguides made in two sections. The 900-mm-long inner straight sections were glued to the scintillator layers before they were stacked, with 25-μm-thick Al foil between layers, in radial groups to provide for the gaps to be occupied by the RSPCs. Each radial group was wrapped tightly with Kapton tape to support it as a unit to be slid into the frame. The lightguide sections were wrapped loosely with black polyester tape and fitted with shims for alignment and support. The inner sections pass through the gaps between the flux-return spokes of the magnet endplates and each is joined to the outer section through an air gap. The outer section of each lightguide is curved to join with the PMT array where it is glued to the PMT. EMI 9954KB PMTs are used for the stopping layers and 28-mm-diameter Hamamatsu R1398 for the T layer. The light output at each end produces approximately 15 photoelectrons in the PMT per MeV deposited in the scintillator.

PMT signals are passively split and recorded by ADCs, TDCs, and by the TD system. The latter provides a pulse-shape history in each tube as well as precision timing ($\sigma \sim$ 500 ps). The signals fed to the TDs from the A, B, C counters are processed through amplifiers with logarithmic transfer functions in order to compress the larger dynamic range of analog pulses from these thick layers.

Signals from the range stack form the basis for most low-level trigger decisions. Individually discriminated range-stack PMT signals provide information for fast pattern recognition. For example, the stopping layer determined from the deepest layer struck provides a coarse measure of the range of a charged track. About 80% of $K_{\mu 2}$ events are rejected by vetoing on the long range muon which tends to penetrate deeper than pions from the $K^+\to\pi^+\nu\overline{\nu}$ endpoint. In addition, the pulse heights from the range stack are summed in hextants; discriminated sums from hextants not involved in a charged-particle track are used to derive veto signals to augment the photon detection systems.

The range-stack energy calibration was done for each end separately, using muons from $K_{\mu 2}$ triggers, by comparison to a Monte Carlo simulation. Attenuation effects were determined from cosmic-ray data taken with zero magnetic field. Drift-chamber tracks were extrapolated into the range stack to get the z position and the attenuation length was adjusted so that the z positions determined from them agree with the DC extrapolation. The average attenuation length determined this way was 2.07 m. However, the effects of the exponential attenuation can be eliminated by taking the geometric mean of the pulse heights from both ends of the counter. The range-stack calibration was monitored using a prescaled sample of $K_{\mu 2}$ triggers taken during running, and high voltages were adjusted as necessary to compensate for changes.

From the calibrated ADC information, energy and range measurements for π^+ from $K_{\pi 2}$ and μ^+ from $K_{\mu 2}$ decays are shown in Fig. 11.

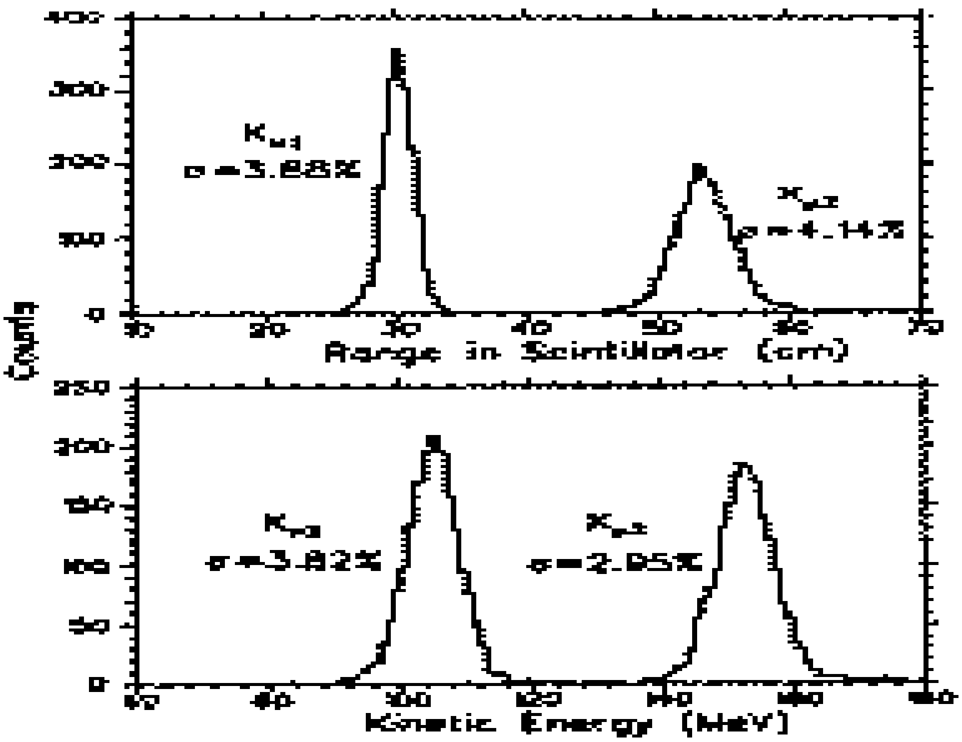

FIG. 11. The (a) range and (b) energy for $K_{\pi 2}$ and $K_{\mu 2}$ events from the range stack, corrected for losses in the target and other materials in the trajectory.

The range determination uses polar-angle information from the drift chamber and a correction based on the energy measurement in the last counter hit to establish the range contribution from the stopping layer. Both energy and range are corrected for the target contribution.

The two layers of multi-wire proportional chambers (RSPC) in each sector are located radially from the detector axis at approximately 635 mm and 720 mm. These chambers measure the axial (z) and azimuthal ($r\phi$) positions of the track in order to refine the range measurement for pions in the off-line analysis. In addition, the z information from the inner RSPC layer is used to give the slope of the track in order to improve the range determination in the second trigger level for online event selection.

Fig. 12

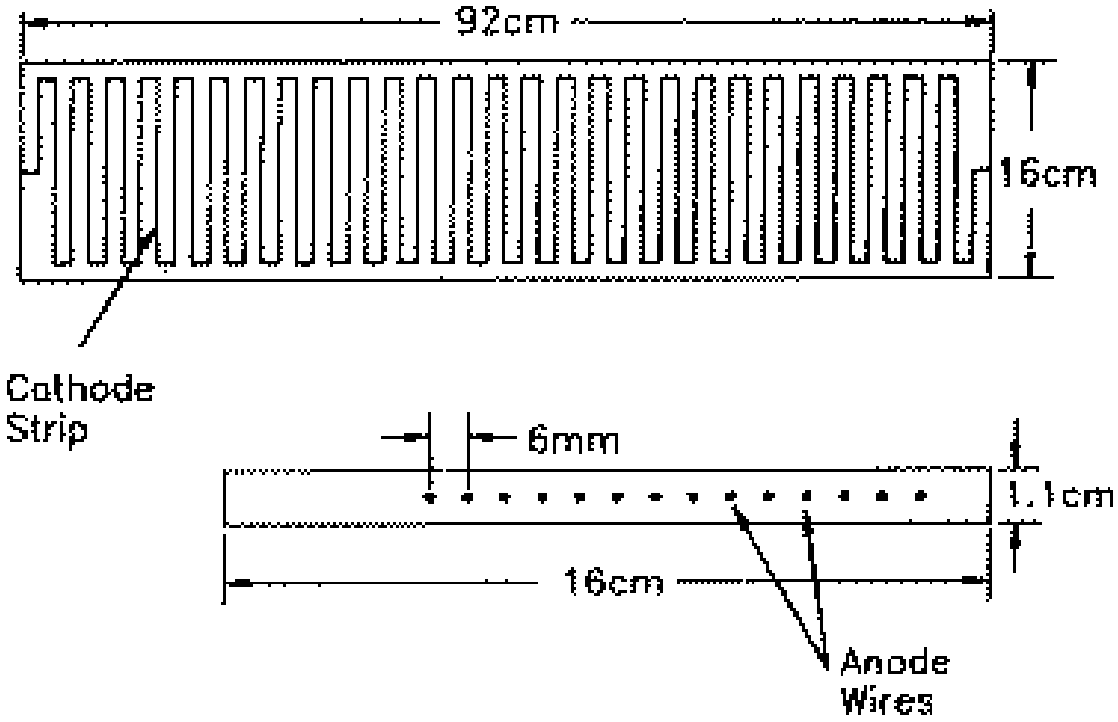

FIG. 12. Schematic showing the serpentine cathode and anode wire layout of the RSPCs. The dimensions shown are approximate for the inner layer.

is a schematic of the design of one of the chambers. To minimize the introduction of inert material and maintain full azimuthal coverage, the RSPC frames were constructed of very light, thin boxes molded from Kevlar cloth in epoxy. Each chamber is approximately 11-mm thick with a total material thickness of 0.315 g cm^{-2}. The active area of the inner (outer) chambers is 926 (1067)-mm long by 160 (180)-mm wide. Electrical components are epoxied and soldered in place. The 25-μm-diameter Au-plated tungsten anode wires are strung 6-mm apart along the z direction. The $r\phi$ measurement is obtained from the position of the struck wire using a coupled delay-line readout scheme, similar to that of ref. [15]. The maximum drift time in the 9-mm gap is $\sim$ 60 ns using a gas mixture of argon : isobutane : methylal of 65.2 : 26.4 : 8.4 and high voltage of 2.5 kV. The z measurement is derived from the end-to-end time difference of signals from the serpetine-patterned copper cathode etched onto G10, making up a 150-ns-long transmission line. The system of 48 chambers is read out by two 96-channel LRS 1879 Fastbus TDC modules in addition to the fast digitizing second-level trigger. The position resolutions (σ) are roughly 10 mm in z and 2 mm in $r\phi$.

G. Transient Digitizers

Unambiguous pion identification is essential to assure adequate rejection of muons from $K_{\mu 2}$ even after kinematic constraints are applied. The information required to certify each pion in the range stack is available in the history of pulse shapes from the PMTs which view the scintillator layers in the region of the pion stop. The pion will decay into a monoenergetic (4 MeV) muon and a neutrino, with the muon in turn decaying into a positron (up to 53 MeV) and two neutrinos. In contrast, a stopping muon will lack the 4-MeV signal and decay directly into a positron and two neutrinos.

In order to observe the decay sequence with high acceptance in the scintillator with a typical signal width of 30–40 ns and with the pion lifetime being 26 ns, a sampling rate of $> 3 \times 10^8$ s^{-1} is needed to be able to detect early decays (< 10 ns). Since the muon lifetime is 2.2 μs, a memory depth of $\sim$ 10 μs is needed to insure high efficiency for detection of the muon decay. The required dynamic range of > 7 bits is determined by the energy loss in each scintillator of the stopping pion (1 – 30 MeV), the decay muon (4 MeV), and the decay positron (3 – 10 MeV). Other necessary features are zero suppression to keep the number of data words reasonable, fast data readout in a standard format to keep deadtime low at high rates, and access to front-end high-speed computing for data compaction and on-line trigger decisions.

A system of 5×10^8-sample s^{-1}, 8-bit transient digitizers (TD) [16] was constructed to meet these specifications. The range-stack signals are multiplexed by summing corresponding layers from four adjacent sectors into one TD channel. At this level, random background rates and signal degradation are not significant problems. Since the detailed information this system offers is also valuable for other detector systems (e.g. photon veto and target) in improving acceptance and accidental rejection, the system is expandable and other channels have been added as they have become available.

The system is housed in Fastbus, supporting up to eight 4-channel modules per crate. Each crate contains a control board (TDMaster) used to read out the data from the 32 channels over a special high-bandwidth bus on the auxiliary backplane, and transfer the data on the main Fastbus backplane to a SLAC Scanner Processor (SSP) [17] crate master. In addition, each crate houses two 500-MHz clock fanout boards. Cooling of each 1.6-kW crate is by forced air through water-cooled heat exchangers.

Fig. 13

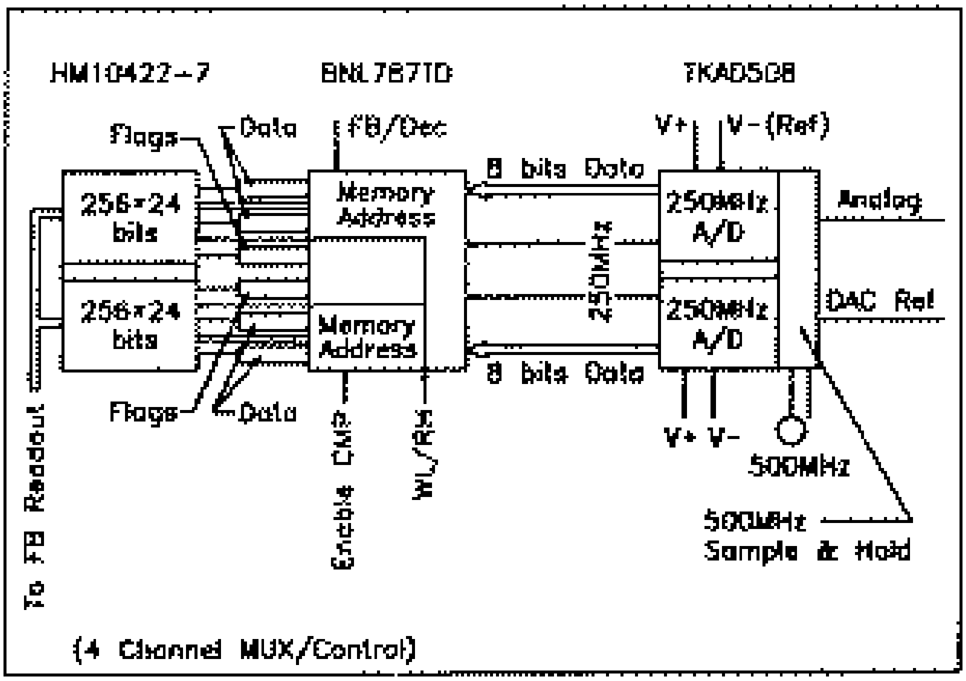

FIG. 13. Block diagram of one channel of the 500-MHz digitizer.

shows a block diagram of the basic digitizing channel. It consists of four parts: the flash ADC, the custom data handling chip, fast memories and data acquisition/readout protocol electronics. Flash ADCs offer the important features of ability to suppress zeros prior to writing and of fast access to the data, as well as the possibility of a large memory depth. The effective dynamic range can be increased, for example, by using bilinear conversion or logarithmic signal processing on the input.

Each channel in the system uses a Tektronix ADC hybrid TKAD508 consisting of a 500-MHz sample-and-hold circuit and two interleaved 250-MHz ADCs shown schematically in Fig. 13. The device requires two differential analog inputs to the Sample-and-Hold, a 500-MHz clock signal, and two reference voltages to the ADCs. The sample-and-hold drives the two ADCs and supplies the necessary clocking and buffering. The output of the hybrid is a pair of interleaved bytes of data at a rate of 250 MHz. One byte is output on the rising edge of the clock, while the other is output on the falling edge of the clock. The device is capable of 7.2 effective bits performance at 100 MHz, and 6.8 effective bits at 250 MHz.

The negative polarity PMT signal is fed to one analog input, and the other input is set at a DC level generated by one of the 32 12-bit Digital-to-Analog converters (DAC) contained on the TDMaster board. This allows fine software control on the effect of the digital threshold to provide zero suppression by compensating small DC voltage shifts in the input signals. The settable DAC levels also provide an efficient means of on-line testing and calibration. This method of using the device provides an input dynamic range of up to 1 V. In addition, a fiducial time marker derived from the trigger is induced on the DC level.

A custom IC (BNL787TD) called the macrocell was designed to handle the high data rate (2 bytes every 4 ns), to provide zero suppression of flash ADC data prior to writing into memory, and to carry out the various housekeeping and digital readout duties required in such a device. The macrocell is a time demultiplexer capable of handling input data rates of up to 400 MHz (but normally runs at 250 MHz), implemented in Signetics ACE2200 current-mode-logic technology and has roughly 2150 gate equivalents. It takes as input 2 bytes of data and a synchronizing clock signal and outputs 4 bytes of data at half the input frequency. In addition the IC generates two identical 8-bit addresses for the external memory and an 8-bit timing marker generated by an internal timer. The output of the timer labels the data words with their relative timing, necessary because of the zero-suppression feature of the system.

The external memory is 4-ns Hitachi HM10422-7 256 x 4 ECL. The data word width is 48 bits: 32 bits of data, 8 bits of time and 8 bits of flags to identify the 4 multiplexed PMT signals in each TD channel. To reduce the load on the address lines from the macrocell, the twelve memories per channel are organized into two banks of 6 memories (hence the two identical address-line sets generated by the macrocell).

The memory for each channel forms a circular buffer, with zero-suppressed data over-writing old data until a trigger is received. Typically, data from -3 μs to +6 μs with respect to the K decay is saved. Zero suppression is achieved by incrementing the memory address only for data bytes exceeding 7 counts as well as for one adjacent byte before and after. The memory depth is extended beyond that afforded by the 8-bit timer (2 μs) by forcing a write to memory whenever the timer counter overflows. Other conditions which synchronize the data to a trigger from the detector, and special diagnostic states also cause the memory address to be incremented to store data. The extension of the timer data to greater than 2 μs and the rearrangement of the data is done subsequent to readout in software at the crate level.

Crate-level intelligence is used to form high-level triggers on-line based on the TD data. The crate controller (SSP) is a special purpose integer-arithmetic computer with roughly 5 – 10 Mips of computing power. In addition, a special front-end processor was developed to handle the raw data from each channel prior to being read out by the SSP. This module, called the Transient Digitizer Smart Controller (TDSC), uses an AMD 29K pipeline RISC processor with separate data, address and I/O busses and a compatible floating point co-processor, and is capable of about twice the effective computing power of the SSP.

The transient digitizer system is used for both on-line trigger and off-line π/μ separation. Fig. 14

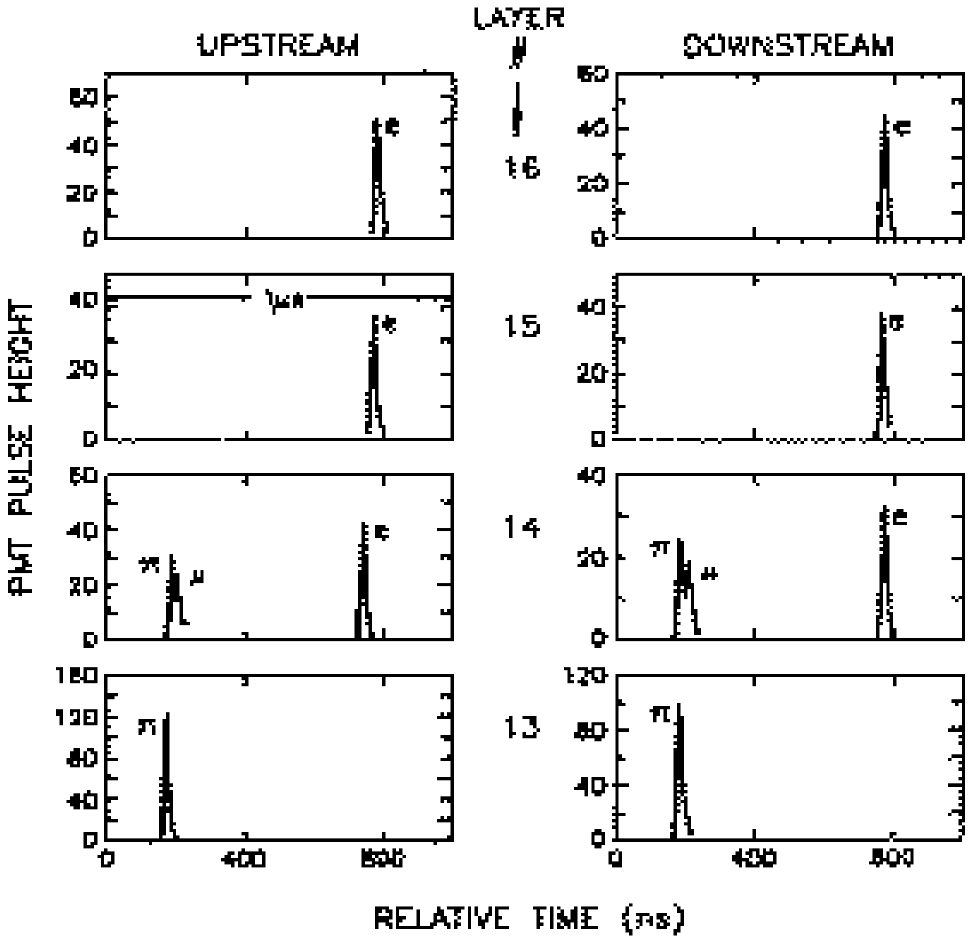

FIG. 14. TD data from four adjacent range-stack layers, showing the pulses for a pion entering from the bottom, stopping in layer 14, decaying to a muon, and finally the muon decaying to a positron.

shows the power of correlated time/pulse-height measurements in a typical $K_{\pi 2}$ event. The π^+ penetrated through several range-stack layers, stopped in layer 14, decayed into a muon, and the muon subsequently decayed into a positron which penetrated several more layers.

Fig. 15

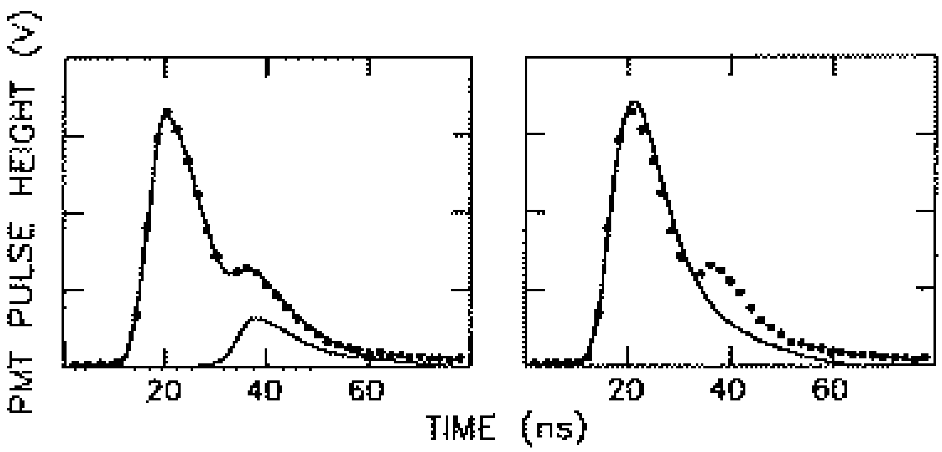

FIG. 15. TD data from a stopping pion showing the second pulse from the decay to a muon. Analytical fits to (a) two pulses and (b) one pulse, from which π/μ particle identification is done.

shows the data from the stopping layer. One can clearly see the primary pulse with the secondary muon decay. In the simplest algorithm to separate pions from muons on-line, a test on the ratio of pulse-area to pulse-peak-value is done and a minimum value is required which tends to reject single pulses (muons) and accept double pulses (pions). This technique typically accepts 72% of pions and rejects 86% of muons.

In off-line analysis, the pulse-shape data are tested for goodness-of-fit to a single-pulse (muon) or a two-pulse (pion) hypothesis as illustrated in Fig. 15. From a successful two-pulse fit one can extract the energy and time of the secondary muon. The curve in Fig. 16a

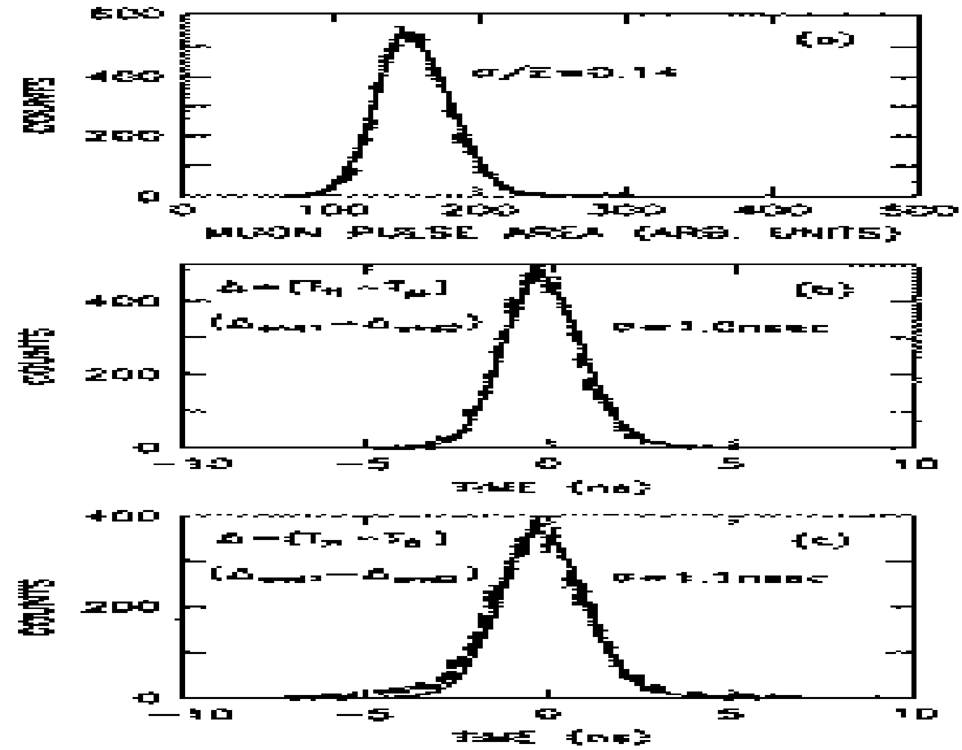

FIG. 16. (a) Energy spectrum for the 4-MeV μ from π decay from TD data. (b) End-to-end time difference for the $\pi \to \mu$ decay and (c) $\mu \to e$ decay from TD data.

shows a 15% energy resolution on the 4-MeV muon, due to a combination of the basic resolution of the device and the fitting-scheme resolution. Reduction of random backgrounds is achieved by comparing the timing of the secondary and tertiary pulses relative to the primary pulse on both ends of the range-stack counter. The end-to-end timing resolution from $K_{\pi 2}$ data for the secondary pulse (muon) is shown in Fig. 16b, and for the tertiary pulse (positron) in Fig. 16c. Fig. 17

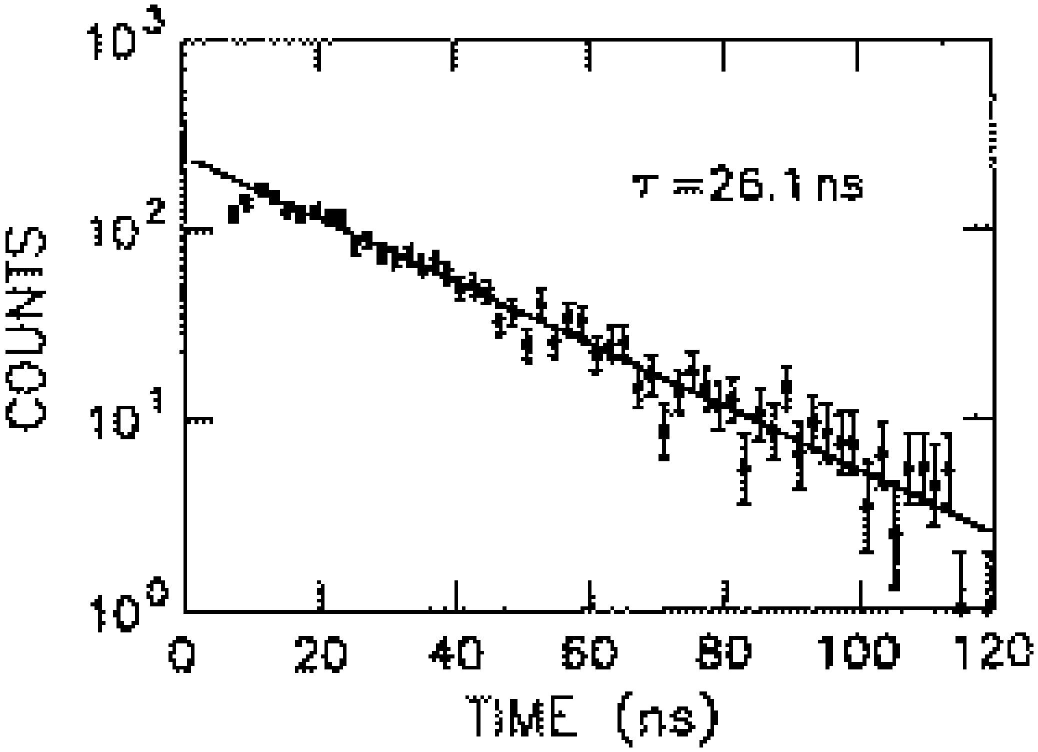

FIG. 17. Pion decay-time spectrum from TD data showing the π lifetime fit.

shows a distribution of the secondary pulse time relative to the primary pulse compared to the pion lifetime of 26.1 ns. The dropoff at times < 8 ns costs about 25% in decay detection efficiency. The final μ/π rejection achieved using both the secondary and tertiary pulse detection is about 5.5×10^{-6}.

H. Photon Veto

Photon detection capability with an inefficiency of $< 10^{-5}$ is a key strategy in suppressing backgrounds from $K_{\pi 2}$ decays and other radiative processes. Photons accompanying kaon decays are detected in the $\sim 4\pi$ sr array of electromagnetic shower counters consisting of alternating layers of 1-mm-thick lead and 5-mm-thick scintillator. Roughly two thirds of the coverage is obtained from the barrel assembly, which surrounds the range stack, 2π in azimuth and 1.90 m axially. Most of the remaining solid angle is subtended by the two endcaps located upstream and downstream of the drift chamber. These systems are augmented by the capability to identify photon conversion elsewhere in the detector, particularly in the target and range stack, and to some extent, in the degrader, using the beam counters.

The barrel is constructed in 24 eight-module sectors as shown in Fig. 2 and 10, supported in a frame with a web-like cross-section made of 1.5-mm-thick stainless steel and fastened to the four aluminum plates supported on the magnet yoke. The azimuthal boundaries of each sector are tilted so that the inter-sector gaps of inert material (air, wrapping, and stainless steel) do not project back to any part of the target. The center divider in each sector splits each layer into halves so that the azimuthal segmentation is 48-fold. The four radial modules add up to a total of 14.3 radiation lengths, and are made up of 16, 18, 20, and 21 lead-scintillator layers, respectively with increasing radius. The 3600 scintillator panels in 75 sizes of increasing width were fabricated of BC408 by Bicron Corp.

Each module is epoxied to a 150-mm-long acrylic mixer block on each end. The light collected from the scintillator layers in each of the 384 mixer blocks is coupled via a flexible wafer made of silicone dielectric gel (Dow Corning SYLGARD 527), to a 67-mm-diameter, 850-mm-long UVT acrylic lightpipe fed through a 73-mm-diameter hole in the magnet end-plate steel to a 77-mm EMI 9821KB PMT mounted on the upstream or downstream end of the magnet. Geometrically, the fraction of scintillator light reaching the photocathode varies from 30% for the inner modules to 17% for the outer modules, and produces $\sim$10 photoelectrons per visible MeV energy loss in the scintillator. The PMTs are individually shielded from the $<$ 15-G fringe field with μ-metal sleeves. Each PMT signal is connected to a passive splitter with outputs to LRS CAMAC 4300B FERA ADC and LRS 1879 Fastbus TDC systems, and to the energy-sum units used in the fast trigger.

In the regions upstream and downstream of the drift chamber covered by the two endcap photon-veto assemblies [18], the lead-scintillator layers are oriented perpendicular to the detector axis so that all photons originating from the target will convert with high probability, leaving acceptable visible energy. Both endcaps surround a 200-mm-diameter hole on the beam axis allowing for the beam upstream, and the target assembly downstream. Each endcap is segmented into 24 azimuthal modules with 66 petal-shaped lead-scintillator layers making up 12.4 radiation lengths. The 3168 scintillator pieces were machined from NE104 (Nuclear Enterprises, Scotland) using a diamond-tipped cutter which left a surface finish that required no polishing to obtain adequate internal light reflection. The light output of each sheet of NE104 was tested using a ^{106}Ru radioactive source, and the material in each module was selected to insure reasonably uniform response from module to module.

In order to transport the light axially with a minimum of inert material in the crucial volume, light from the scintillator layers in each module is absorbed and re-emitted by a fluorescent wavelength-shifter bar which covers most of the outer edge [19]. The wavelength shifter is viewed by a 52-mm EMI 9954KB PMT via an adiabatic lightguide, coupled with a flexible wafer made of silicone dielectric gel to a 1.3-m-long, 41-mm-diameter UVT acrylic lightpipe through a hole in the magnet endplug steel. Each PMT is locally shielded against the small fringe magnetic field. The wavelength shifters were made of 6.5-mm-thick acrylic doped with BBOT (Bicron Corp.), chosen for its very fast light output decay time of 1.6 ns. The emission and absorption properties of NE104 and BBOT are well matched for our geometry and give $\sim$8–10 photoelectrons per visible MeV in the modules. The attenuation lengths of the NE104 and BBOT are approximately 400 and 800 mm, respectively. Other tests showed that the output from NE104 would not degrade by more than $\sim$20% due to radiation damage over the expected life of the experiment.

The 24 modules in each endcap are enclosed in a welded 700-μm-thick stainless steel frame. A detailed view of a module is shown in Fig. 18.

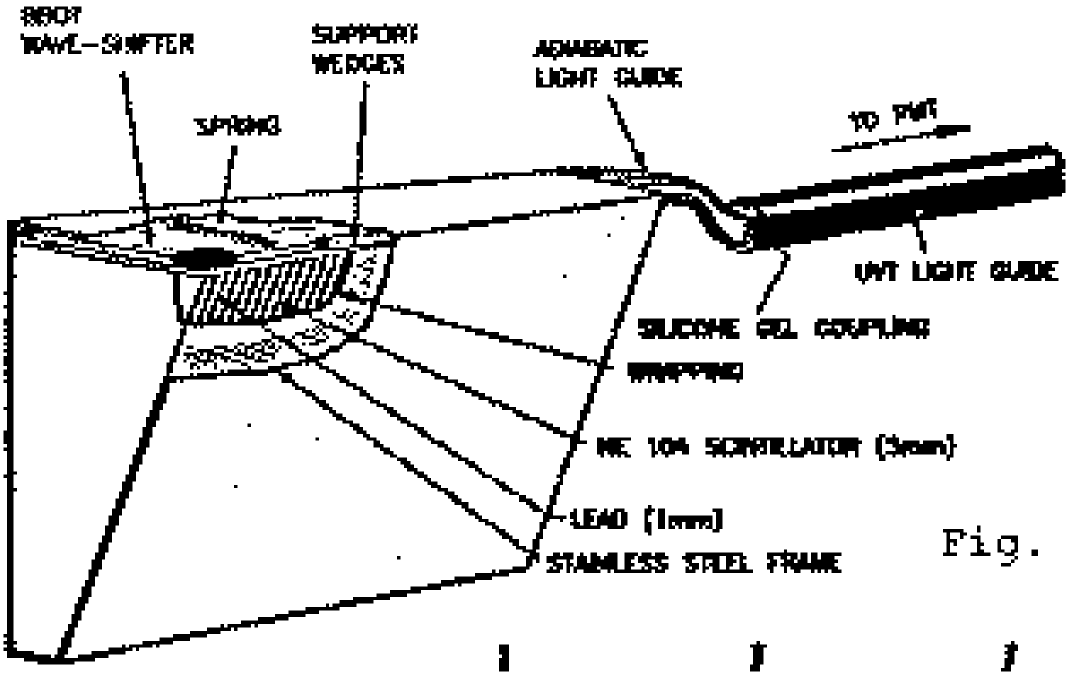

Fig. 18

FIG. 18. Isometric schematic showing the structure of a photon veto endcap module.

An air gap between the scintillator and the wavelength shifter is maintained by two 300-μm-diameter copper wires extending the full length of the module. The air gap minimizes reabsorption by the primary scintillator layers of light reflected internally in the wavelength shifter. The modules are supported in the frame in such a way that uniform pressure applied by small springs holds the wavelength shifters in place. Each endcap assembly ($\sim$750 kg) is supported in a 4.8-mm-thick aluminum cylinder bolted to the magnet steel endplug.

The relative quantum efficiencies for the PMTs were measured and tubes were selected to further balance the output. The PMT signal cables feed $\times 10$ amplifiers in the counting room. One amplifier output is connected to a passive splitter which distributes signals to the ADC system, the discriminators feeding the TDC system, and the summing units used for the fast trigger. The second output is summed in groups of 4 separated modules and the sum is fed to a TD channel.

Calibration of the barrel is performed by comparing the energy depositions from cosmic ray data with the calculated energy loss of minimum-ionizing radiation in the known scintillator thickness. Even though the trigger and analysis constraints eliminate events passing through the scintillator at shallow angles, cosmic muons have a broad asymmetric spectrum due to large Landau fluctuations. In addition, there is a large end-to-end variation in light output (factor 2.3–7.3) because the ~1.2-m average attenuation length is short compared to the 1.9-m active length. (The attenuation length was determined for the barrel in a similar manner to the range stack.) The position dependence of this effect can be eliminated by considering the geometric mean of the light output from both ends.

Photon vetoing is based on visible energy, *i.e.*, energy deposited in the scintillator. However, total energy calibration is useful as a check (see below), and for other physics measurements involving photons. To turn calibrated visible energy in the barrel into actual photon energy, a correction is made for the invisible energy deposited mostly in the lead but also in the support structure and wrapping materials. This correction was calculated using a Monte Carlo simulation of the passage of electrons and photons through matter using the program EGS [20]. Three energies, 20, 100, and 200 MeV, were run under various conditions, and gave a typical visible energy fraction $f_{vis} = 0.29$. The tilted support structures show up as ~15% dips in f_{vis} as a function of azimuth of entry, as shown in Fig. 19.

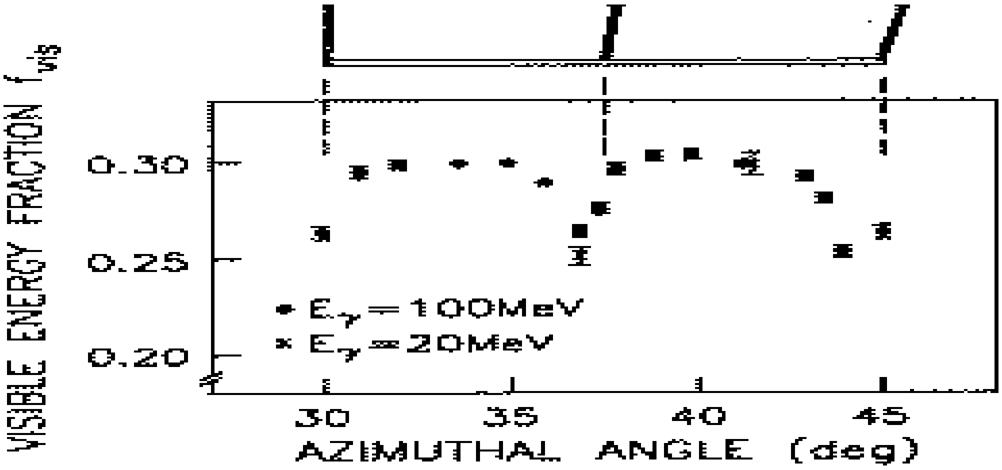

FIG. 19. Azimuthal dependence of the visible energy fraction in the barrel photon veto, correlated with the support structure shown at the top.

For the endcaps an initial crude calibration was obtained using cosmic ray muons, although, unlike the barrel, the geometry and orientation of the endcaps in the detector makes a clean trigger, for which the energy deposition is known, impossible. Proper calibration is done in beam with a special trigger to use the 152-MeV monoenergetic muons from $K_{\mu 2}$ decay which stop in the first third of the endcap, after losing energy in the target and other parts of the detector. The trigger and analysis select events for which the muon energy is contained mostly in one module in order to provide an easily recognizable peak. The expected visible energy is determined by a Monte Carlo calculation. The HV was adjusted to balance the fit peaks in all modules and establish the required dynamic range, normalized to the fit Monte Carlo spectrum. The systematic uncertainties on the gains set in this manner are estimated to be about 10%.

Since this calibration relies heavily on a Monte Carlo calculation, an independent check was carried out by reconstructing the energy of the neutral pion in $K_{\pi 2}$ decay. The kinematics are strongly biased against both photons hitting the endcaps when the π^+ is accepted in the range stack. However, by using the more reliably calibrated barrel, it is possible to use events where one photon is in the barrel and one in an endcap, and add their measured energies to reconstruct the π^0 energy. Taking into account the visible energy fraction f_{vis}, which should be similar for barrel and endcap systems, the consistency of the endcap energy calibration could be checked. As a test of the overall photon energy calibration and the determination of f_{vis}, the summed energy from the barrel and the endcaps of the two photons from π^0 in clean $K_{\pi 2}$ triggers was measured and is shown in Fig. 20.

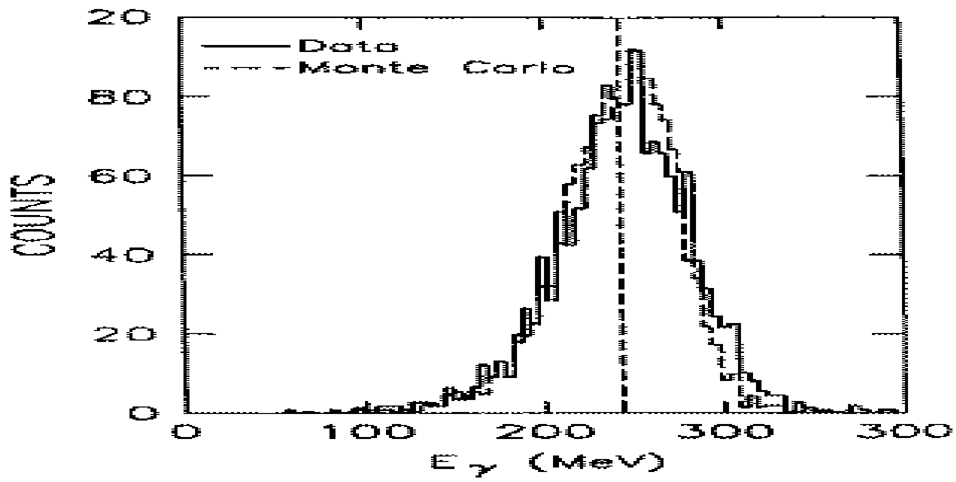

FIG. 20. Summed photon energy corrected for the visible energy fraction from $K_{\pi 2}$ events. The π^0 energy in such events is shown by the vertical dashed line.

The peak is close to the expected value of 245.6 MeV, the π^0 energy.

$K_{\pi 2}$ events are also used to determine the time calibration of the barrel and endcap in the TDCs, since the π^+ in the range stack defines the event time reference. The time resolution is limited by the 2-ns time-bucket width in the pipeline TDCs.

The photon-veto system performance is optimized by simultaneously minimizing two characteristics — the accidental rate and the photon inefficiency. The accidental rate causes good events to be rejected due to random uncorrelated hits, resulting in a loss in acceptance. The photon inefficiency limits the primary function of the system to reject backgrounds. In our photon-veto system, the parameters to optimize are the low-energy threshold and the width of the coincidence-resolving-time window.

For a given energy threshold, the accidental rate is reduced by narrowing the vetoing time window. This is determined by discriminator pulse widths in the trigger and by the time resolution of the detector and electronics in the off-line analysis, which uses leading-edge timing, rather than overlap. To study these effects, we use $K_{\pi 2}$ data and plot the leading-edge time of all individual energy depositions in the barrel and endcaps with respect to the measured time of the π^+. These plots show a peak corresponding to photons from π^0s and a background due to accidentals. Fig. 21

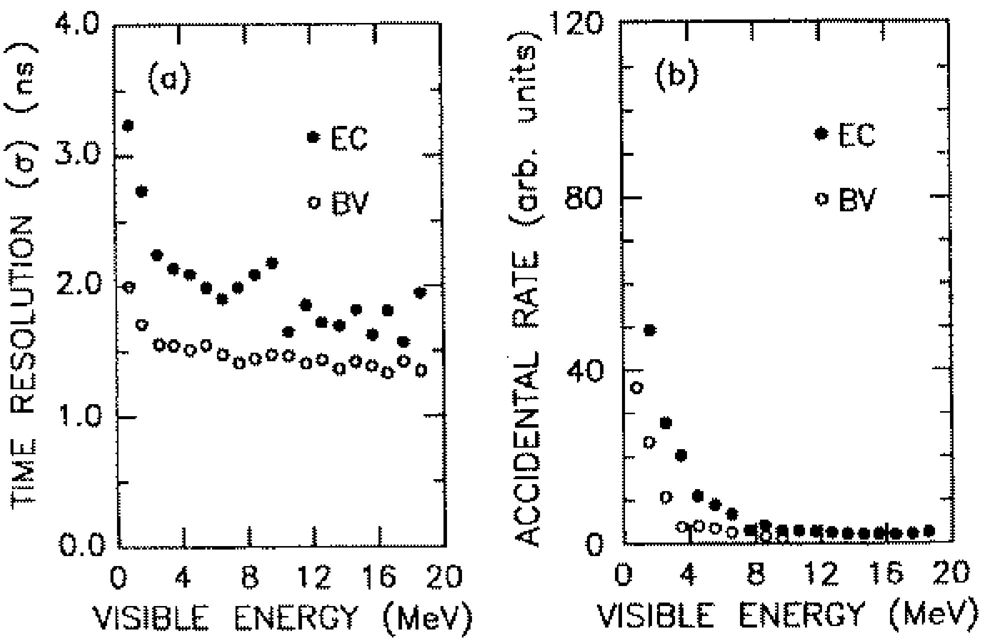

FIG. 21. Barrel and endcap (a) time resolution and (b) accidental rate as a function of visible energy.

shows the time resolution and accidental rate inferred from such plots as a function of visible energy. The barrel benefits in resolution from being read out at both ends, and in accidental rate from being farther from the beam than the endcaps.

In the trigger, pulse widths make the effective time window for vetoing tens of nanoseconds. Visible-energy vetoing thresholds are 5 MeV for the sum of barrel energy and 10 MeV for the sum in each endcap. Off-line, time windows are reduced to as low as 12 ns and thresholds are typically 1 MeV for the entire detector.

The inefficiency of the photon-veto system can be measured directly in terms of π^0 detection in $K_{\pi 2}$ decay (see below). However, in order to guide the design of the detector and to study its performance, Monte Carlo simulations were performed. Inefficiency comes from three sources: escape of photons before conversion, sampling fluctuations in which little or no energy is deposited in scintillator, and photonuclear interactions in which no charged products are detected. The Monte Carlo program uses the EGS package [20] for electromagnetic interactions, and photonuclear reactions are simulated using PICA [21].

Other routines propagate the reaction products, including neutrons, through the detector, and describe the final-state nuclear de-excitation by a simple evaporation model. The effects of the last few MeV of nuclear excitation are studied separately using various assumptions about the cascade of low-energy photons. For our low detection thresholds, the inefficiency is quite sensitive to these assumptions, giving estimates with an uncertainty of about an order of magnitude.

Nonetheless, the simulations indicate that detection inefficiencies for single photons in $K_{\pi 2}$ events vary from 10^{-2} to 10^{-4} for photon energies between 20 and 225 MeV. Sampling fluctuations dominate the loss of low-energy photons (<50 MeV), while photonuclear interactions dominate the loss of high-energy photons. For the latter, photonuclear losses are larger than escape losses (the next-largest effect) by about a factor of five.

The primary function of the photon detection system — rejecting photons from π^0s produced by kaons decaying at rest — benefits from several factors beyond simple detector efficiency: There are two photons in each event; both have energies greater than 20 MeV and one always has energy greater than 122 MeV; the high-energy photon is constrained by momentum conservation and our π^+ selection to penetrate the most efficient sections of the photon detector, including the range stack which is one radiation length thick with little inert material. The π^0 detection inefficiency measured by observing the magnitude of the 205-MeV/c momentum peak surviving the $K^+ \to \pi^+ \nu \overline{\nu}$ analysis [9] is $1 - 3 \times 10^{-6}$, depending on analysis cuts, and is consistent with Monte Carlo simulations.

I. Trigger

An event trigger is refined in four stages of increasing sophistication and decision time. The first two stages, Level 0 and Level 1, are executed in a custom-built system of ECL logic boards [22] designed around the configuration of the range stack, but including the essential input from the beam, target and photon-veto systems. The K-stop logic is formed externally in fast NIM logic from the beam scintillators and Čerenkov counters and a discriminated analog sum of the target-fiber signals. Level 0 is entirely composed of fast combinational logic, and introduces about 40 ns of deadtime per event. Level 1 is implemented in memory lookup units (MLU), and takes several microseconds to make a decision. Trigger decisions beyond Level 1 involve arithmetic processing of ADC and TD data. In Level 2 the ADC data from the range-stack track are selected and summed in CAMAC in a few tens of microseconds and arithmetic processing in Fastbus of the TD data requires about 150 microseconds. Level 3 applies loose cuts and compacts the data in about 500 microseconds. The latter will be described further with the Data Acquisition system in the next section.

Low-level triggers are classified into groups depending on which of two timing strobes are used to issue gate signals to the various detector systems. The beam strobe, related to a stopping kaon, is formed in NIM logic and determined by the kaon Čerenkov time. The detector strobe, related to a charged particle entering the fiducial volume of the range stack, is formed in the trigger system using the discriminated PMT signals from the T and A layers from all 24 sectors. For each sector, the mean time from the two ends of A defines the time for a $T \cdot A$ coincidence, which is then ORed from all sectors to form the detector strobe. Cable lengths were trimmed to keep sector-to-sector jitter to < 1 ns.

The physics triggers, *i.e.* kaon decays, as well as range-stack calibration and beam-background studies, use both strobes, but are initiated by the detector strobe. A schematic diagram of the Level 0 system is shown in Fig. 22.

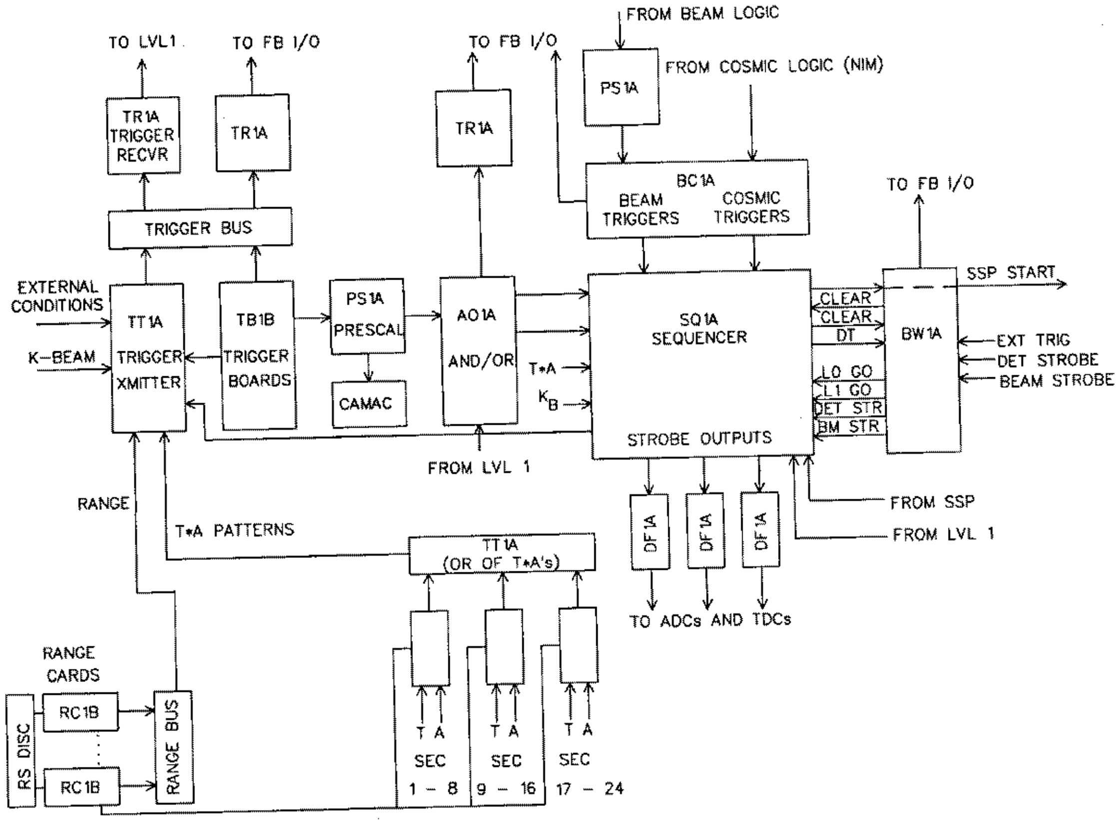

Fig. 22. Block diagram of the Level 0 trigger system.

FIG. 22. Block diagram of the Level 0 trigger system.

A K^+ stopping in the target decays into a π^+ that produces a $T \cdot A$ when it enters the range stack. The range (stopping layer) of the π^+ is deduced by the Range Cards (RC1B), each of which takes the discriminator signals from one complete range-stack sector, and determines layers belonging to a positive-charged track. Pulse heights from the range stack segmented into hextants, barrel veto and endcaps are summed and discriminated to produce veto signals, which are latched and placed on the 48-line Trigger Bus by the Trigger Transmitter (TT1A). The stopping layer information from the Range Cards is also latched onto the Trigger Bus, along with external conditions from the beam and target systems, to establish a minimum decay time with respect to a K^+ stop, and topology information from a fast MLU which recognizes valid patterns of $T \cdot A$ hits.

Up to 16 Level 0 triggers can be implemented using ECL logic chips on wire-wrap Trigger Boards (TB1B) that make decisions based on the state of the Trigger Bus. For example, the Level 0 trigger for $K^+ \to \pi^+ \nu \overline{\nu}$ requires, in addition to a $T \cdot A$ coincidence, an incident K^+ (identified by the Čerenkov) to have stopped in the fiducial volume of the target and decayed after the minimum-time condition (determined from the $B4$, target, I-, and V- counters), a charged-particle track in the range stack that penetrated into layer B, but not to layer 19 or beyond, and no energy above threshold in the photon-veto system.

The information on the Trigger Bus is updated with every $T \cdot A$ until any of the Trigger Boards find a match and issues a 50-ns-long interim inhibit (DT) signal to the Trigger Transmitter which freezes the Trigger Bus until the Level 0 signal has had time to propagate through its Prescaler (PS1A) channel, the AND/OR board (AO1A) , and the Sequencer (SQ1A). The prescaling factor is adjustable from 1 to 2^{20}. If the prescaler output is *true*, the Sequencer will issue beam and detector strobes via Delay-Width-Fanout boards (DF1A) to the appropriate ADCs and TDCs, send a signal to initiate Level 1, and set a flip-flop which further extends the DT signal.

The information transmitted by the AND/OR board verifies that the Level 0 and Level 1 triggers are of the same kind, and that the relative timing of the Level 0 and 1 signals are correct. Both the Trigger Bus and the AND/OR Board are read by receiver boards (TR1A), which feed the data to I/O boards in Fastbus to be accessed by the SSP crate master as seven 32-bit trigger words which summarize the Level 0 and Level 1 trigger results, as well as bus and MLU status for every defined trigger type. These data eventually form part of an event written to tape.

Cosmic ray triggers (for barrel-veto calibration) require only the detector strobe, while beam triggers require only the beam strobe. These triggers are much simpler, since they essentially bypass the Range Card and Trigger Bus structures, and do not involve Level 1. A special board is used to latch the prescaled beam trigger and output the proper strobes and levels to the Sequencer and I/O boards. The cosmic triggers, based on $T \cdot A$ and I-counter hits recognized in a fast MLU board, are also prescaled and fed through a special board to the Sequencer and I/O boards.

There is also provision for external triggers, carrying their own strobe timing, *e.g.* the endcap-veto $K_{\mu 2}$ calibration trigger. External triggers are input via a "Bells & Whistles" (BW1A) board which generates the required timing strobes and handles all communication with the Sequencer and I/O boards. The other function performed by the BW1A board is to verify correct ADC gate timing, and SSP starts, and it is able to generate a *clear* signal to the Sequencer, for trigger-system testing without the need for computer intervention.

In the case of $K^+ \to \pi^+ \nu \overline{\nu}$, the Level 1 decision is based on a refined range measurement taking into account the polar angle of the π^+, and the range in the target. An MLU address is created out of the z position of a hit in an RSPC by converting the end-to-end timing difference to an analog level and digitizing it to a 4-bit word with a flash ADC. This takes about 1 μs and gives a position resolution of $\sim$ 10 mm. The range of the π^+ or μ^+ in the target, can be deduced from the number of struck target elements which is also flash digitized to a 4-bit word. These, together with another 4 bits of stopping layer information derived from the Level 0 Trigger Bus, make up the 12-bit address to a second MLU whose contents determine the Level 1 trigger decision. If a Level 1 *false* is returned, a *reset* is issued to restart the TDCs and clear the ADCs without interrupting the SSPs.

If the Level 1 decision is *true*, the Sequencer initiates Level 2 by sending a *start* signal to the Fastbus trigger SSP which instructs the secondary SSPs to read out the LRS CAMAC 4300B FERA ADCs and LRS Fastbus 1879 TDCs, and to the TDSCs in the TD system. Level $2A$ is performed during the readout of the ADCs. A charged-track energy sum is formed by digitally adding the ADC outputs of the target and range-stack counters as they are transferred from the CAMAC modules to a Fastbus buffer memory via a front panel ECL readout port. LRS ECLine modules (LRS 2372 MLU and LRS 2378 ALU) are used to screen out information from non-track-associated ADCs and to perform the sum to which a maximum energy cut is applied.

In parallel with the Level $2A$ process in CAMAC, the raw TD data are processed in Level $2B$ by the TDSCs resident in the TD Fastbus crates. The TD pulse-shape information from the range-stack stopping counter is analyzed for evidence of a second pulse indicative of decay of the stopped pion. Only events with a separated second pulse or with a pulse-height to pulse-area ratio consistent with a single pulse augmented by $\pi^+ \to \mu^+\nu$ decay are retained. Prior to the implementation of TDSCs, the pulse-analysis algorithm was executed in the SSPs after the TD data were read out.

For each event that passes both Level $2A$ and Level $2B$, data are collected in memory by the SSPs in each Fastbus crate for further processing by the data acquisition system. The trigger Sequencer receives a *clear* signal from the Fastbus Master SSP and then issues the *reset* to the ADCs and TDCs. After waiting 7 μs for the resetting to finish, the DT signal to the Trigger Transmitter is cleared and the trigger is restarted.

Table I

TABLES

TABLE I. Typical Trigger and Data Acquistion for $K^+ \to \pi^+ \nu\overline{\nu}$

Level	Rejection	Execution Time per event	Number Surviving	Deadtime
K_{STOP}/1.6-s spill	—	—	5×10^5	—
Level 0	36:1	50 ns	1.4×10^4	25 ms
Level 1	14:1	3 μs	10^3	42 ms
Level 2A	2.3:1	50 μs	430	50 ms
Level 2B	17:1	150 μs	25	65 ms
Readout	1:1	8 ms	25	200 ms

summarizes typical trigger and data acquisition performance for $K^+ \to \pi^+ \nu \overline{\nu}$ triggers only at a rate of 5×10^5 K_{stop} per 1.6-s AGS beam spill. In practice, additional physics and monitor triggers are suitably prescaled and added up to a maximum deadtime fraction of 30 – 35%.

J. Data Acquisition

The logic of the trigger system initiates signal digitizing in the data acquisition system. After Level 1, the higher levels of trigger decision involve close links with the data acquisition in order to minimize the large overhead invested in reading out and logging the data. The process up to the point of reading the front-end data into SSP memory in each Fastbus crate was described in the previous section. Once the decision is made (after Level 2) to collect and format the data in the SSPs, the data path continues in the data acquisition system in four stages as indicated in Fig. 23.

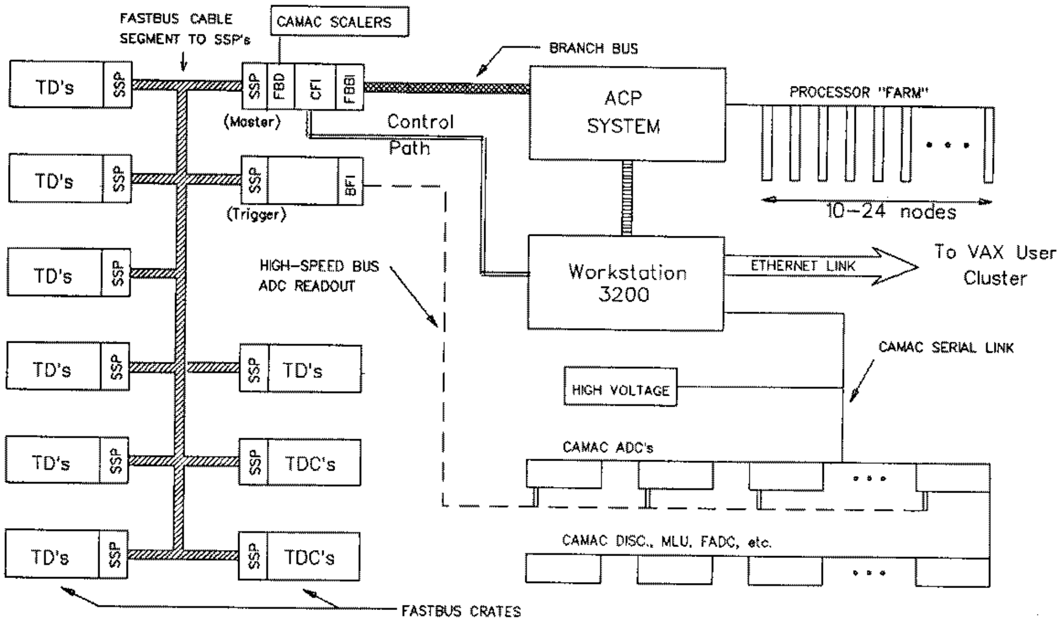

Fig. 23. Block diagram of the data acquisition system.

FIG. 23. Block diagram of the data acquisition system.

Each stage is performed by a different microprocessor and runs asynchronously from the others. As long as no stage saturates its allotted time, only the first stage contributes to the dead time.

Stage 1 is the event-by-event readout. During each beam spill from the AGS, acquisition of the data is done independently in each secondary Fastbus crate for all events that pass the trigger. The data are collected by an SSP and stored in its data memory (512 Kbytes). The only communication between Fastbus crates is done via the Fastbus cable segment to synchronize the event number, to acknowledge completion of the readout in each crate and to report possible errors.

The limitations at this stage come from the conversion time of the front-end modules, the size of the data memory of each SSP, the speed of the SSPs, and the amount of data formatting required. For the moderate data rates possible in the present configuration of the experiment, the 8-ms time for this step has not been critical because the trigger decision was fast enough that most events were rejected before it was necessary to format them.

The second stage occurs during the period (typically 1.6 s) between beam spills. Data are collected from each crate over the Fastbus cable segment and events are assembled in memory by a Master SSP. Events are then dispatched to the memory of an available node in the VME ACP [23] system via the Branch Bus, using a Fastbus Branch Coupler. The data transfer rate on the Fastbus cable segment is about 18 Mbyte s^{-1} and on the Branch Bus about 15 Mbyte s^{-1}. The typical event length at this stage is 20 kbytes.

At the third stage an array of up to 24 ACP microprocessor nodes, located in two VME crates controlled by a host VAX [24], is used to reject events (functioning as the Level 3 trigger) and to compress the data. Each node uses a Motorola 68020 processor and 2 Mbytes of memory. Processing in the nodes has been limited primarily to compression of TD data resulting in a factor of two reduction in the event size, and to some analysis of special test and calibration triggers. For example, TD channels which contained only the fiducial time marker were cut. The system has the potential to perform analysis and apply conservative cuts equivalent to a first pass off-line analysis.

In the final stage, the data are collected from the nodes by the ACP host Q-bus [24] computer, stored in memory and copied to tape by independent processes. Extra records containing calibration and control information are added. At this stage the bandwidth limitations come mostly from the Q-bus which must handle all transfers from VME to memory, from memory to tape, all disk activity and all DECNET/ETHERNET activity. An effective throughput of 180 Kbyte s^{-1} to tape is presently achieved.

On-line monitoring of detector performance and data quality was done by transferring samples of events from the host VAX over DECNET [24] to other VAX nodes running analysis and display software similar to the off-line packages. For example, this on-line package produces occupation plots of detector subsystems, and monitors parameters such as rates, gains, event lengths, and persistant missing channels, with the purpose of quickly identifying failures in the detector, trigger, and data acquisition systems.

III. SUMMARY

The detector for AGS Experiment 787 at Brookhaven National Laboratory has been built to begin the search to observe the decay $K^+\to\pi^+\nu\overline{\nu}$ if it occurs at the level predicted by the Standard Model. New limits for $K^+\to\pi^+\nu\overline{\nu}$, $K^+\to\pi^+\mu^+\mu^-$, $K^+\to\pi^+\gamma\gamma$, and $\pi^0 \to \nu\overline{\nu}$ decays have been reported from the analysis of the first data run which corresponds to $\sim 10^{10}$ stopped K^+s. Subsequent data corresponding to an order of magnitude more statistical sensitivity have been obtained and analysis is in progress. Several improvements to the detector and beamline are under way or under study to ensure that the full sensitivity required to test the Standard Model predictions will be achieved.

ACKNOWLEDGMENTS

This project was carried out thanks to the talented contributions of the following people and we are pleased to acknowledge them here: G. Azuelos, D.G. Bartolomeo, P. Bichoneau, J.C. Biggs, C. Bopp, D.R. Brady, S. Chidzik, M. Constable, H. Coombes, J. Cracco, J.R. Cuccia, S. Daviel, W. Groom, R. Henderson, W.E. Lenz, C. Lim, S. Moreale, G.J. Munoz, R. Openshaw, R. Rabberman, H. Ratzke, W. Sands, H.W. Sauter, R.J. Strzelinski, and M.R. Van Lith.

We thank D.I. Lowenstein for use of the steel for the detector magnet. We appreciate the help from W.J. McGahern and the AGS Experimental Areas Group, and from R.J. Spellman and the BNL Central Shop staff.

We wish also to acknowledge the funding support of this project by the National Research Council of Canada, the Natural Sciences and Engineering Research Council of Canada, and the United States Department of Energy under Contracts No.DE-AC02-76CH00016, No.W-7405-ENG-36 and No.DE-AC02-76ER03072.

IV. REFERENCES

a) Present address: Physics Dept., Princeton University, Princeton, NJ 08544.

b) Present address: California Institute of Technology, Pasadena, CA 91125.

c) Present address: Los Alamos National Laboratory, Los Alamos, NM 87545.

d) Present address: Physics Dept., Cornell University, Ithaca, NY 14853.

e) Present address: Physics Dept., Rice University, Houston, TX 77251.

[1] S.L. Glashow, Nucl. Phys. **22** (1961) 579; S. Weinberg, Phys. Rev. Lett. **19** (1967) 1264; A. Salam, in *Elementary Particle Theory* Proc. 8th Nobel Symposium, Aspenäsgarden, 1968, ed. N.Svartholm (Almqvist and Wiksell, Stockholm, 1968)

[2] for reviews of rare kaon decays, G.L. Kane and R.E. Shrock, in "*Intense Medium Energy Sources of Strangeness*", Proc. of the Santa Cruz Workshop, 1983, ed. T. Goldman, H.E. Haber and H.F.-W. Sadrozinski (AIP conf. Proc. No.102) p.123; J.S. Hagelin and L.S. Littenberg, Progress in Particle and Nuclear Physics **23** (1989) 1; D.A. Bryman, Int'l. Journal of Modern Physics **A4** (1989) 79.

[3] S.L. Glashow, J. Iliopoulos, and L. Maiani, Phys. Rev. **D2** (1970) 1285.

[4] J. Ellis, J.S. Hagelin and S. Rudaz, Phys. Lett. **B192** (1987) 201; J. Ellis, J.S. Hagelin,

S. Rudaz and D.-D. Wu, Nucl. Phys. **B304** (1988) 205; C.O. Dib, I. Duneitz and F.J. Gilman, Phys. Rev. **D39** (1989) 2639.

[5] Y. Asano *et al.*, Phys. Lett. **107B** (1981) 159.

[6] M. Atiya *et al.*, Phys. Rev. Lett. **64** (1990) 21.

[7] M. Atiya *et al.*, Phys. Rev. Lett. **63** (1989) 2177.

[8] M. Atiya *et al.*, Phys. Rev. Lett. **65** (1990) 1188.

[9] M. Atiya *et al.*, Phys. Rev. Lett. **66** (1991) 2189.

[10] V. Fitch and R. Motley, Phys. Rev. **101** (1956) 496.

[11] H. Hinterberger and R. Winston, Rev. Sci. Inst. **37** (1966) 1094.

[12] J.S. Frank and R.C. Strand, *Proc. Workshop on Scintillating Fiber Detector Develöpement for the SSC*, vol. 1, p. 359, (Fermilab, Batavia, 1988).

[13] J.V. Cresswell, S. Ahmad, E.W. Blackmore, D.A. Bryman, N. Khan, Y. Kuno and T. Numao, IEEE Trans. on Nucl. Science **35** (1988) 460.

[14] D. Akerib, Princeton University Ph.D. Thesis (unpublished), 1990.

[15] D. Bryman, J. Cresswell and R. Skegg, Nucl. Instr. Methods **141** (1977) 573.

[16] M. Atiya, M. Ito, J. Haggerty, C. Ng, F.W. Sippach, Nucl. Instr. Methods **A279** (1989) 180.

[17] T. Barklow, T. Glanzman, A.J. Lankford and K. Riles, IEEE Trans. on Nucl. Science **33** (1986) 775.

[18] J. Roy, University of British Columbia M.Sc. Thesis (unpublished), 1988.

[19] G. Keil, Nucl. Instr. Methods **83** (1970) 145, and G. Keil, Nucl. Instr. Methods **83** (1970) 111.

[20] R.L. Ford and W.L. Nelson, SLAC Report 210, June, 1978.

[21] T.A. Gabriel and R.G. Alsmiller, Jr., Nucl. Sci. and Eng. **49** (1972) 243.

[22] M. Selen, Princeton University Ph.D. Thesis (unpublished), 1989.

[23] I. Gaines *et. al.*, *The ACP Multiprocessor System at Fermilab* in Proceedings of Computing in High Energy Physics, Asilomar State Beach, California, 1987, or Fermilab Report Fermilab-Conf-87/21.

[24] VAX, Q-bus and DECNET are trademarks of Digital Equipment Corp.