Nuclear Instruments and Methods in Physics Research A243 (1986) 323–347
North-Holland, Amsterdam

AN APPARATUS TO MEASURE THE STRUCTURE OF THE PION

C. BIINO, J.F. GREENHALGH, W.C. LOUIS, K.T. McDONALD, S. PALESTINI *, F.C. SHOEMAKER and A.J.S. SMITH

Joseph Henry Laboratories, Department of Physics, Princeton University, Princeton, New Jersey 08544, USA

C.E. ADOLPHSEN, J.P. ALEXANDER **, K.J. ANDERSON, J.S. CONWAY, J.G. HEINRICH, K.W. MERRITT + and J.E. PILCHER

Enrico Fermi Institute and Department of Physics, The University of Chicago, Chicago, Illinois 60637, USA

E.I. ROSENBERG and D.T. SIMPSON ++

Ames Laboratory and Department of Physics, Iowa State University, Ames, Iowa 50011, USA

Received 27 September 1985

We discuss the design and performance of a large-acceptance double-magnet spectrometer used in Fermilab Experiment 615. The apparatus was used to measure the properties of high-invariant-mass muon pairs produced in collisions of a pion beam with a tungsten target in order to infer the distribution of quarks within the pion and to study the production mechanism in detail. The emphasis was on interactions in which a single quark, and subsequently the muon pair, carried a large fraction of the pion's momentum. A hardware trigger processor enabled the detector to extract the relatively weak signal while operating in a high-rate environment.

1. Introduction

The prompt production of a pair of oppositely-charged leptons in hadron–hadron collisions has, by now, been well established to occur through quark–antiquark annihilation to a virtual photon, as originally suggested by Drell and Yan [1]. This is the dominant production mechanism for lepton pairs with invariant masses in the continuum region, for example between the J/ψ and Υ resonances [2,3]. Experiments such as the one described here have exploited this simple and intuitively appealing picture in order to extract information about the structure of hadrons and to gain insight into the nature of the strong interaction.

Fermilab experiment E-615 was dedicated to studying muon pairs produced in the reaction $\pi^{\pm}\mathrm{N} \rightarrow \mu^{+}\mu^{-}\mathrm{X}$ using 80 and 255 GeV/c pion beams incident upon a tungsten target. The muon pairs of interest carried a large fraction x_{F} (Feynman-x) of the available longitudinal momentum and had a large invariant mass M. As x_{F} of the muon pair tends toward unity, the valence antiquark in the pion carries a large share (x_{π}) of the momentum of the pion. Departures from the standard Drell–Yan parton model interpretation of the data are anticipated in this extreme of the kinematic range [4].

Evidence for such a departure was obtained in a previous experiment performed by this group [5]. We found that the muon-pair angular distribution changed from the form $1+\cos^2\theta$ to $\sin^2\theta$ as x_{π} approached unity, where θ is the angle between the μ^{+} and the beam pion in the muon-pair rest frame. This corresponds to a change of the virtual-photon polarization from transverse to longitudinal due to QCD corrections to the quark-antiquark annihilation process.

The present experiment was designed to confirm or refute this observation and at the same time to remeasure the structure function of the pion. By comparing data taken at two different center-of-mass energies $\sqrt{s}$ but at fixed values of the scaling parameter $\tau \equiv M^2/s$, we searched for scale-breaking effects which might be enhanced near $x_{\mathrm{F}} = 1$. Data collected with both π^{+} and π^{-} beams served to verify that the quark–antiquark annihilation picture is still valid at high x_{F}.

In this paper we describe the apparatus in detail, with some emphasis on the specialized trigger electronics used in the experiment. The acceptance and resolution characteristics of the detector are discussed as well.

We begin with an overview of the design of the experiment in sect. 2. We describe the principal ele-

* Now at INFN, Sez. Torino, Italy.
** Now at SLAC, Stanford, California 94305, USA.
+ Now at Fermilab, Batavia, Illinois 60510, USA.
++ Now at IBM, Endicott, New York 13760, USA.

0168-9002/86/$03.50 © Elsevier Science Publishers B.V.
(North-Holland Physics Publishing Division)

ments of the experiment in turn, beginning with a detailed discussion of the beam in sect. 3. The experimental target is described in sect. 4. The salient features of the selection magnet, the hadron absorber, and the spectrometer magnet are outlined in sects. 5–7. The particle detection devices, namely wire chambers and scintillation counters, are described in sects. 8 and 9. Details of the multilevel trigger are presented in sect. 10. Data acquisition is described in sect. 11. On-line and certain off-line tests of the performance of the trigger logic are discussed in sect. 12. Scintillation-counter efficiency studies are treated in sect. 13. The acceptance and resolution of the apparatus are discussed in sects. 14 and 15. In conclusion, a summary of the principal data sets is presented in sect. 16.

2. Overview of the design of the experiment

The apparatus described herein was a second-generation version of that used in our earlier work [6,7]. In both versions, muon pairs were produced in interactions of a hadron beam with a target and then analyzed in a magnetic spectrometer, after penetrating a beam dump immediately downstream of the target.

The present detector was designed to make accurate and complete measurements of the properties of muon pairs produced with large x_F. In particular, we sought good acceptance in the muon-pair angular-distribution variables. Hence good efficiency was required in the beam region, since that area was populated by the more energetic member of asymmetric pairs. This consideration ruled out a toroidal-spectrometer design [8] and the use of a heavier absorber in the beam [9]. We also sought a configuration of the apparatus that would allow the trigger to favor high-mass Drell–Yan events and discriminate against low-mass pairs and background associated with beam-pion decays. Because the cross section for the production of muon pairs with $M > 4$ GeV/c^2 is extremely small (on the order of 100 pb/nucleon), the detector was also required to operate in a high-intensity beam ($\sim 3 \times 10^8$ pion/s).

These requirements were met with a two-magnet design (see fig. 1). The first magnet (selection magnet) was located close to the target and contained the hadron absorber, which filled the gap between the magnet poles. The selection magnet imparted a transverse-momentum kick of 3.2 GeV/c which bent the tracks of high-mass $\mu^+\mu^-$ pairs until they were nearly parallel, thereby lending them an identifiable geometry. At the same time, low-momentum particles from low-mass pairs and pion decays were swept out of the fiducial volume of the detector downstream.

The central element of the detector itself was a

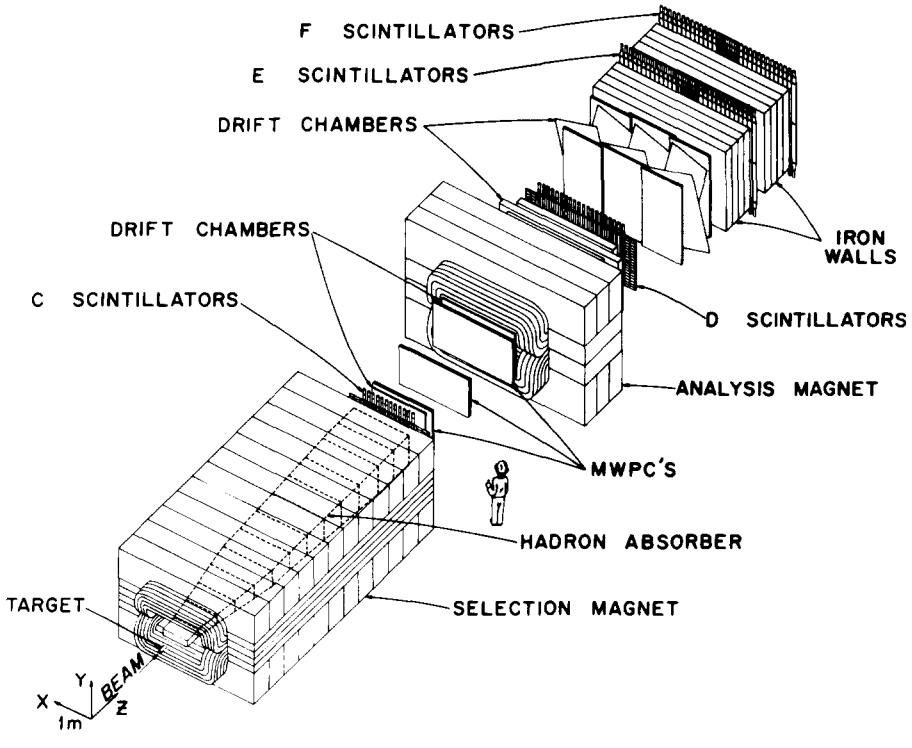

Fig. 1. View of the apparatus in the 255 GeV/c configuration. The A and B hodoscopes upstream of the target are not shown.

magnetic spectrometer consisting of a large-aperture analysis magnet and associated wire-chamber planes. A combination of multiwire proportional chambers and drift chambers afforded both good efficiency and position resolution in a high-rate environment. The wire chambers were sensitive through the beam region, yielding good acceptance over a wide range of cos θ at high x_F.

In order to operate in a beam of about 3×10^8 particles per second, the detector was equipped with a fast trigger based on plastic scintillation counters. A sophisticated hardware trigger processor [10] was built to reject muon pairs with low invariant mass and pairs that did not originate in the target.

Throughout this paper we shall use a right-handed coordinate system with the z-axis along the beam and the y-axis directed upward, as indicated in fig. 1.

3. Beam

The experiment was performed in the High Intensity Area of the Proton West Laboratory at Fermilab, where we were among the first users of the new superconducting accelerator, the Tevatron.

The data were collected at two beam momenta, 80 GeV/c [11] and 255 GeV/c [12], to permit a search for scaling violations in the production of muon pairs at high x_F. These beams were produced by 400 and 800 GeV/c protons, respectively. The yield of 255 GeV/c pions during the 800 GeV/c operation of the Tevatron was an order of magnitude larger per proton than the previously available yield (of 255 GeV/c pions) during the 400 GeV/c operation of the Main Ring. Even though the cycle time of the Tevatron was longer than that of the conventional-magnet 400 GeV accelerator, the 800 GeV superconducting accelerator still delivered about three times as many pions per hour as its predecessor.

3.1. Beam line

During each beam extraction (spill) of 10–20 s, about 5×10^{12} primary protons were delivered to a 40 cm long, 3.8 cm diameter beryllium target in the Proton West Area. Extraction occurred approximately once every minute. The secondary beam [13] was collected at 0°, and particles of the desired charge were selected by a dipole magnet and a collimator. The beam was then brought through a triplet of quadrupoles to an adjustable horizontal collimator used to select the momentum bite. The rest of the beam line consisted of a double focusing–defocusing transport and a final triplet for focusing on the experimental target. The beam line also had several horizontal bends to reduce the contamination of muons from beam-particle decays. The resulting secondary beam intensity was roughly 3×10^9 particles per spill at both 80 and 255 GeV/c.

The beam profile and position were monitored by several segmented-wire ionization chambers [14] that had wire spacings of 1 and 2 mm. The r.m.s. radius of the beam spot on the experimental target was about 1.0 cm. During the 80 GeV/c run, the beam traversed part of the fringe field of the selection magnet before reaching the target; by the average interaction point, the beam had departed from its nominal direction by +1.0 mrad in the x–z plane.

3.2. Beam flux

The intensity of the beam was monitored by four ionization chambers. One measured the flux of primary protons while the other three measured the pion flux. Two of the latter chambers were segmented into five concentric rings that were read out individually.

An absolute calibration of the ionization chambers was carried out on several occasions by comparing the activation of Cu foils placed in the beam with the number of ion-chamber counts recorded during a fixed number of beam spills. Measurements of the induced amounts of ^{24}Na and ^{52}Mn in the foils yielded calibration constants with a systematic uncertainty of about 20%.

3.3. Beam composition

The secondary beam on target had some contamination of kaons, protons, electrons, and muons. Estimates of the beam composition were based on previous measurements of charged-particle production at various secondary momenta by 400 GeV/c protons on beryllium targets [15]. Interpolating linearly between the reported values at 60, 120, 200, and 300 GeV/c, and assuming the particle production ratios are a function only of $x_{lab} \equiv P_{secondary}/P_{beam}$, we obtained the percentages for charged secondaries at the production target shown in table 1.

We were, of course, concerned only with the hadronic content of the beam at our target, since electrons could not produce an appreciable background. Taking account of the pion and kaon lifetimes, the beam momenta, and

Table 1
Secondary beam composition at the primary target.

Beam momentum [GeV/c]	Beam sign	$\pi^{\pm}$ [%]	$K^{\pm}$ [%]	$p^{\pm}$ [%]	$e^{\pm}$ [%]
80	−	74.3	6.1	2.2	17.4
255	−	85.6	5.7	1.5	7.2
255	+	52.6	4.5	40.5	2.4

Table 2
Hadronic composition of the beam at the experimental target.

Beam momentum [GeV/c]	Beam sign	$\pi^{\pm}$ [%]	$K^{\pm}$ [%]	$p^{\pm}$ [%]
80	−	91.9	5.2	2.9
255	−	92.8	5.5	1.7
255	+	53.8	4.1	42.1

the length of the beam line (250 m), we arrived at the hadronic composition of the beam at our target shown in table 2.

Muons from pions and kaons that decayed during the beam transport were a significant source of triggers, especially during the 80 GeV/c run, even though steps were taken to minimize the rate of accidental triggers. Only muons within ~ 5 cm of the beam axis at the experimental target could cause false triggers, since muons outside this radius were rejected by the "halo veto" described in sects. 9 and 10.1. Muons near the beam axis and with momentum close to that of the beam passed harmlessly through "beam holes" in the scintillator hodoscopes, described in sect. 9. The ratio of muons to pions in the beam was about 3% at 80 GeV/c and 1% at 255 GeV/c.

3.4. Beam momentum

Direct measurements of the beam-momentum spectra were made during special low-intensity runs using a magnetic spectrometer mounted just upstream of the target. Four stations of drift chambers (three planes each) were used together with a dipole magnet capable of producing field integrals up to 1.35 GeV/c. The two lever arms were 17.2 and 12.7 m. The momentum resolution obtained was 0.5% with a systematic error estimated to be 0.8%.

The results of these studies appear in fig. 2. The average beam momenta were 79.3 GeV/c and 254 GeV/c. The distributions shown have full widths at half-maximum (fwhm) of about 6 and 8%, respectively.

These measurements were confirmed by an analysis of the momentum spectra of off-beam-momentum muons, as observed in the detector proper. The momentum spectrum of negatively-charged muons associated with Level-1 triggers (sect. 10) in the 80 GeV/c data sample is given by the solid curve in fig. 3. The muon-momentum spectrum is observed to cut off sharply at 45 GeV/c, the kinematic lower limit expected from 79 GeV/c parent pions, corroborating the more direct measurement of the beam momentum discussed above.

3.5. Beam spill structure

The Fermilab beam was delivered in 1.1 ns wide "buckets" every 18.8 ns, a period determined by the 53.1 MHz rf modulation of the accelerator. At an intensity of 4×10^9 pions in 20 s, the average number of particles per rf bucket was 3.8. The bucket population was subject to substantial variations during the spill.

We used a Cherenkov counter operated at atmospheric pressure to monitor the variations of spill inten-

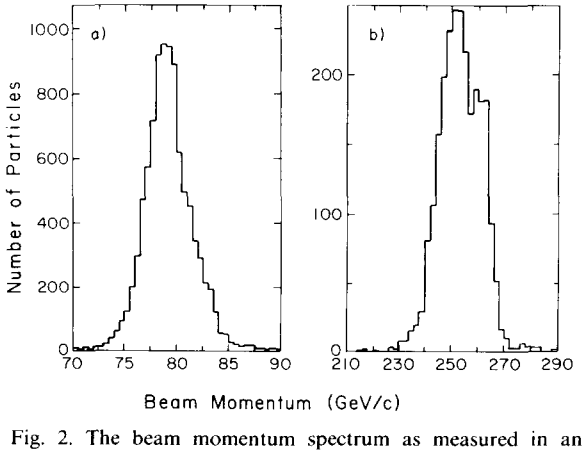

Fig. 2. The beam momentum spectrum as measured in an auxiliary spectrometer (a) during the 80 GeV/c run and (b) during the 255 GeV/c run.

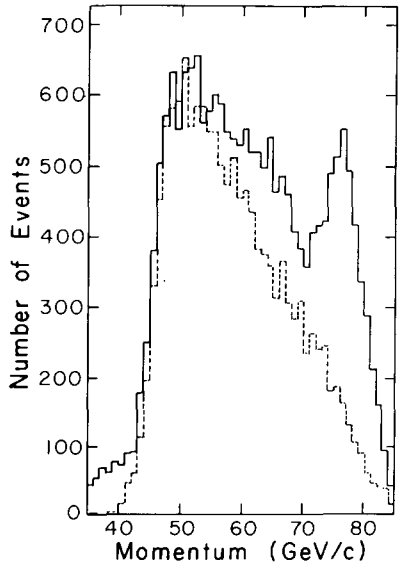

Fig. 3. The spectrum of muons in the beam as measured in the detector during the 80 GeV/c run. On-momentum muons were suppressed due to the "beam holes" in the scintillator banks. The dashed histogram is the result of a Monte Carlo simulation of the decay of 79 GeV/c beam pions.

sity. For the Cherenkov counter to serve as a beam-structure monitor, its discriminator threshold was set at about 10–15 simultaneous particles. Thus only the highly populated buckets were monitored. On the basis of studies of intensity correlations over time spans of the order of 100 ns, the effective spill length was found to be about 5–10 s. Therefore the peak instantaneous rates in the detector were typically 2–4 times the average rates stated in sect. 13.

The Cherenkov counter also provided a measurement of the detector dead-time. This was obtained by recording the number of Cherenkov counts in each spill as gated by logic levels describing the status of the triggering and data-acquisition systems. Another measure of the dead-time was furnished by a telescope of three 1.6 mm thick scintillation counters, gated in similar fashion, that viewed the target center at 90° from a distance of about 1.5 m during the 255 GeV/c run. The great majority of data were collected with a detector live-time between 60 and 70%. Virtually all of the dead-time was due, and in roughly equal measure, to the beam-muon veto logic (sect. 10.1) and to the time spent by the on-line computer in data acquisition (sect. 11).

4. Target

A tungsten target was chosen for two reasons. First, the use of a heavy target optimized the dimuon yield since the Drell–Yan cross section in nuclei depends on the atomic number A as A^1 [16] while the absorption cross section grows only as $A^{0.76}$ [17]. Second, because the interaction length for pions in tungsten is only 11.5 cm, the target could be relatively short, permitting useful vertex constraints to be applied in event reconstruction and fitting.

Our cylindrical tungsten * target measured 5.1 cm in diameter and 20.3 cm in length, corresponding to 1.7 inelastic interaction lengths for pions. Thus more than 80% of the beam interacted in the target. Because the target was also the equivalent of 55 radiation lengths, it contributed significantly to multiple Coulomb scattering of the muons produced in the target.

The longitudinal position of the target differed for the two beam momenta, as shown in fig. 4. For the 80 GeV/c run, the gap between the target and the upstream end of the beryllium absorber in the selection magnet was 2.4 cm. For the 255 GeV/c run, the separation was increased to 102 cm, and 56 cm of beryliumoxide (BeO) was placed in the gap (sect. 6).

* By weight, the target was 97.2% W, 2.8% Ni, Cu, and Fe. The pion-nuclear interaction length in this target was calculated to be 11.9 cm.

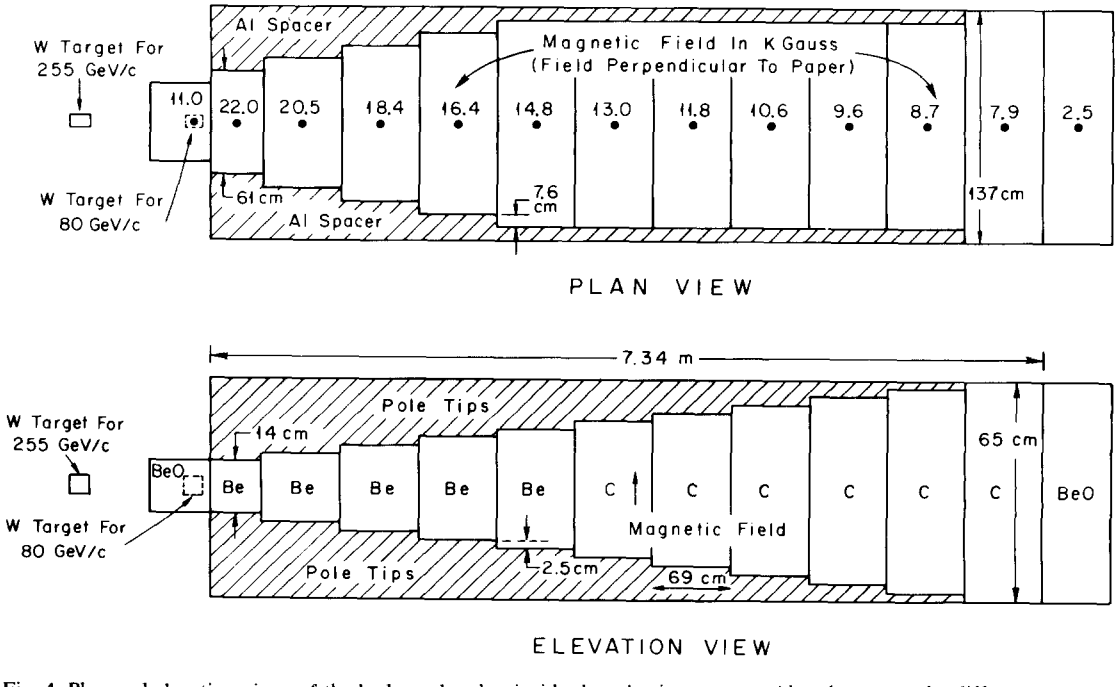

Fig. 4. Plan and elevation views of the hadron absorber inside the selection magnet. Also shown are the different target positions for the 80 and 255 GeV/c runs.

5. Selection magnet

The 400 t dipole selection magnet provided a momentum kick of 3.2 GeV/c, corresponding to a field integral of 10.7 T m. It was built by Fermilab for this experiment out of copper and iron from the main-ring magnets of the Argonne Zero Gradient Synchrotron.

Elevation and plan views of the selection-magnet aperture are shown in fig. 4. The aperture was 7.34 m long and 1.37 m wide, as measured along the pole pieces and between the coils. The vertical gap increased from 14 to 65 cm in eleven steps along the beam direction. The horizontal gap was tapered as well. The field was mapped using the Fermilab Ziptrack [18], a device employing three mutually-perpendicular search coils nested in a cart that rolled on an aluminum beam under computer control. The measurements showed that the horizontal (x) component of the field was everywhere less than 1% of the main (y) component. Furthermore, the integral along z of the x component was everywhere smaller than 0.1% of the integral of the y component. Assuming that B_x is zero, one can show that the field integral was the same along any line passing through the target and the gap of the magnet, even though the pole pieces were tapered in a step-wise fashion.

For the 80 GeV/c run, the detector upstream of the spectrometer magnet was compressed along the beam direction in order to improve the detector acceptance. This was achieved by rolling the selection magnet on rails 1.64 m toward the spectrometer magnet.

6. Hadron absorber

The hadron absorber that filled the selection magnet (fig. 4) was composed of low-Z materials in order to minimize multiple scattering. Good resolution in the initial track direction was important in the calculation of muon-pair kinematic quantities and the rejection of muons that did not originate in the target. The absorber consisted of 3.20 m of beryllium and 4.12 m of graphite inside the magnet, followed by 0.61 m of BeO located immediately downstream of the magnet. For the 255 GeV/c run, an additional 0.56 m of BeO was placed between the target and the upstream end of the magnet, making the absorber 15 absorption lengths deep for energetic pions and also 35 radiation lengths long.

7. Spectrometer magnet

The spectrometer magnet was a wide-aperture dipole with a magnetic-field strength of 13.8 kG at the center and a field integral of 2.9 T m (0.86 GeV/c). The magnetic-field volume was effectively 2.2 m long, 1.8 m wide, and 0.9 m high. (The pole pieces were 1.52 m along the beam.) The main component of the field was vertical and antiparallel to the field of the selection magnet during the 80 GeV/c run, but parallel instead during the 255 GeV/c run, the polarity chosen to optimize the performance of the trigger electronics. The field integral, based on measurements made with the Ziptrack, was found to be uniform over the cross section within 1%. The integral of the x component along the z-axis was less than 4 MeV/c.

8. Wire chambers

The muon trajectories were measured with 25 planes of wire chambers upstream and downstream of the spectrometer magnet.

8.1. Multiwire proportional chambers

Nine planes of multiwire proportional chambers (MWPC) were located between the selection magnet and the spectrometer magnet (see fig. 1). Successive planes in each of three separate chambers furnished x–u–v coordinate information. The oblique planes were inclined at $\pm 15.5°$ to the vertical direction. The anode planes consisted of 20 μm gold-plated tungsten wires, 2.1 mm apart, kept at zero voltage. The cathodes were 69 μm thick aluminized Mylar foils, raised to 4100 V. The anode–cathode separation was 6.4 mm. The chamber volume was filled with a gas consisting of argon (70%), isobutane (30%), and Freon 13B1 (0.1%). The mixture was bubbled through methylal at 10°C. In order to prevent the development of dark current, the cathode voltage was ramped down 20% between beam spills. During the spill, a chamber plane typically drew 100 μA. Each of the 7264 anode wires was connected to a preamplifier mounted on the chamber itself which drove the readout modules through 430 ns delay lines.

The detection efficiency averaged over all MWPC planes was 95%, and the position resolution achieved was about 0.75 mm. The transverse alignments of the MWPC's were determined to about 60 μm. The MWPCs are described in greater detail in ref. [12].

8.2. Drift chambers

Sixteen planes of drift chambers completed the tracking system. Four planes measuring the x coordinate were located upstream of the spectrometer magnet where they were used to improve the tracking accuracy. Two separate groups of six planes each that measured the coordinates x–x–v–u–x–x were located downstream of the spectrometer magnet. The final six planes were actually composed of 18 individual chambers that had been used in a previous experiment [19] and served as a model for the construction of the others. The cells

in the u–v chambers were inclined at $\pm 18°$ to the vertical. The drift cells were 1.9 cm wide and 6.4 mm thick, with potentials on the cathode wires graded from -1600 V to -100 V. The anode voltage, typically 1800 V, was also ramped down 20% between beam spills. The chamber gas was 70% argon bubbled through methylal at 10°C, and 30% isobutane. Signals from each of the 2128 channels were discriminated in amplifiers mounted on the chamber frames and transmitted to time-to-digital converters (TDC) through 76.2 m long (395 ns) ribbon cables. Each TDC module processed a group of 8 channels and had multihit capability. A given TDC took inputs from every 7th wire in a chamber so that the burden of handling high-rate regions could be spread over many digitizer boards. The drift time was measured in 31 bins, each ~ 8.3 ns wide. The drift velocity was about 45 μm/ns, except near anode wires and cell boundaries. The system is described more fully in ref. [19].

The measured average position resolution at high intensity was 290 μm in the drift chambers. The chamber alignments and time offsets were determined to an accuracy of about 20 μm and checked periodically. The time offsets were not measured for single channels but were averaged over groups of 56 cells. While it appears that a cell-by-cell calibration would have allowed us to reach about 200 μm position resolution, the momentum resolution of the spectrometer would not have been improved, since it was dominated by multiple scattering in the target and hadron absorber. The detection efficiency averaged over all drift-chamber planes was 90% [11,20].

In the reconstruction of particle trajectories the drift-chamber time information was used not only to measure position but to determine whether or not the tracks in the event belonged to the same rf bucket. For about 90% of the reconstructed tracks, a reasonably high number of hits were distributed on both the negative-x and positive-x sides of the drift cells, permitting the time offset to be determined from the fitted trajectory. A distribution of time offsets is shown in fig. 5 for a typical 80 GeV/c data tape. A resolution of about 2.5 ns was obtained, enabling us to distinguish tracks belonging to neighboring rf buckets, which were 7.5 standard deviations apart. A cut was placed at 5 standard deviations.

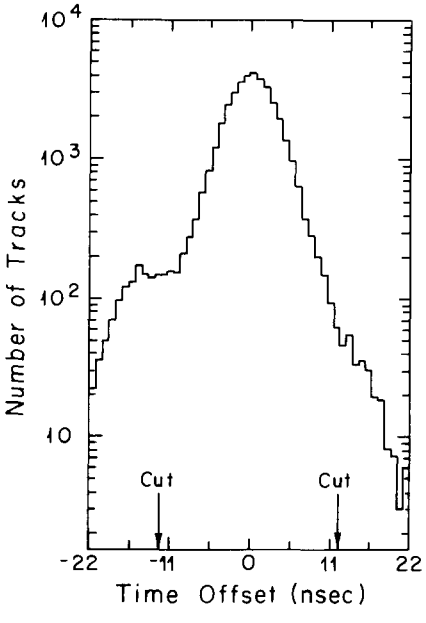

Fig. 5. The distribution of time offsets for tracks reconstructed in the drift chambers. Distance has been converted to time assuming a drift velocity of ~ 45 μm/ns (full scale corresponds to 1000 μm). The shoulder at -16 ns was due to muons in the bucket previous to that of the event trigger. (The out-of-time peak appears shifted from -19 ns because the track-reconstruction algorithm discarded hits which contributed significantly to the χ^2 of the fit.) In-time tracks were defined to be those lying between the indicated cuts.

9. Scintillation counters

The detector was equipped with 6 banks of plastic scintillators arranged in 14 planes, as illustrated in figs. 1 and 6. The A and B banks were designed to veto events associated with an incident muon outside the beam pipe. The C and D banks provided detailed information on the candidate muon-pair positions which was used by the hardware trigger processor. The E and F banks afforded a confirmation that the event indeed consisted of two muons.

The A and B banks were located upstream of the target. They covered an area approximately 2 m wide and 1m high centered on the beam pipe (which was 8.3 cm in diameter). The A hodoscope contained 56 counters, 36 Ax (vertical) and 20 Ay (horizontal), arranged in two planes; the B hodoscope contained 60 counters, 38 Bx and 22 By, also in two planes. The A and B scintillators were all 5.1 cm wide. To ensure that the particles being vetoed were muons, a 0.6 m thick iron absorber separated the A and B banks, and 1.4 m of iron and concrete shielded the B bank from the target region.

The C bank was located immediately downstream of the selection magnet and covered an area 1.42 m wide by 0.64 m high. It consisted of 28 Cx (vertical) counters, 48 Cy (horizontal) counters each spanning half of the aperture, and 31 Cu (oblique) counters arranged in two staggered planes. The D bank was placed downstream of the spectrometer magnet, between the two clusters of drift chambers. This bank was 2.24 m wide, 1.09 m high, and consisted of 44 full-length Dx counters, 48 Dy

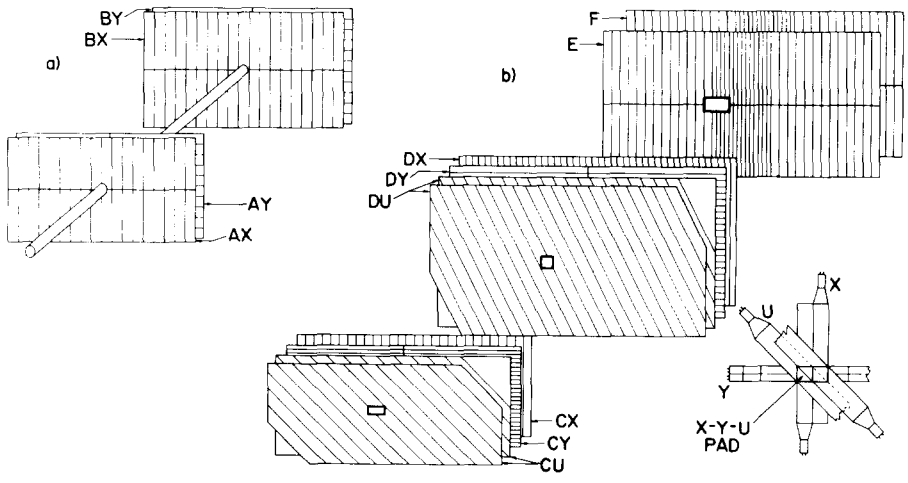

Fig. 6. (a) Layout (not drawn to scale) of the A and B scintillators used to form the "halo veto". The beam pipe passed through the banks as shown. (b) Layout (not drawn to scale) of the C, D, E, and F scintillators used to form the muon-pair trigger. The insert shows how a threefold coincidence of an x, either one of two y, and a u counter formed a "pad". The geometry required that the u bank be composed of two planes.

counters each spanning half the horizontal aperture, and 47 *Du* counters also arranged in two staggered planes. The insert in fig. 6b illustrates the way in which the counters overlapped.

Two 1 m thick iron walls were placed in succession downstream of the last plane of drift chambers. Their purpose was to absorb any nonpenetrating particles that might have "punched" through the absorber inside the selection magnet. The E and F banks were located between and behind the iron walls, respectively. Each bank covered an area 3.40 m wide and 1.77 m high and consisted of two rows of 40 vertical (x) counters that extended half the height of the bank. In each bank the 28 central counters were 5.7 cm wide, while the 52 outer ones were twice as wide. Adjacent counters were overlapped slightly to improve triggering and muon-identification efficiency.

To reduce the number of accidental triggers associated with muons (from pion decays) that remained in the beam, we created "beam holes" in the C and D banks along the line the beam would have followed through the apparatus. The hole was 5 cm high in the C bank, 9 cm high in the D bank, and 10 cm wide in both banks. The positions of the holes were changed according to the beam momentum. For the run at 255 GeV/c, beam holes were created in the E and F banks as well, each 30 cm wide and 15 cm high.

10. Trigger

The data-acquisition trigger, which was based entirely on scintillation-counter information, was designed to select events with two penetrating particles (muons) produced in the target and at the same time to discriminate against pairs with low invariant mass and pairs containing a "halo" muon from beam-pion decay.

Trigger requirements were imposed on three successive levels. During normal data-taking, events that satisfied all three levels of logic were written to tape. In addition, every 1000th Level-1 trigger was recorded for diagnostic and calibration purposes. The first level of logic was designed to select events having at least two distinct muons in coincidence with the beam spill structure (i.e. with a zero-crossing of the 53.1 MHz rf signal) and no halo particles. The first-level trigger initiated the wire-chamber readout and the latching of scintillator patterns. The second level of logic required that every event have at least two tracks pointing back to the target in the nonbending (elevation) view. The third level of the trigger demanded that the muon trajectories resemble those of high-mass pairs. The second and third levels could abort the data stream before it was recorded by the on-line computer and restore data-taking ability within a microsecond. A flow chart of the trigger logic is shown in fig. 7.

Most of the trigger electronics was designed and built specifically for this experiment. Circuit diagrams for some of the electronics developed at Princeton University may be found in ref. [11]. The discriminators, latches, and trigger processors used ECL integrated circuits while a few standard NIM modules were used in forming final coincidences at each trigger level.

A discussion of the overall effect of the higher-level triggers in sect. 10.5 follows a detailed description of the trigger logic in sects. 10.1 through 10.3 and a few

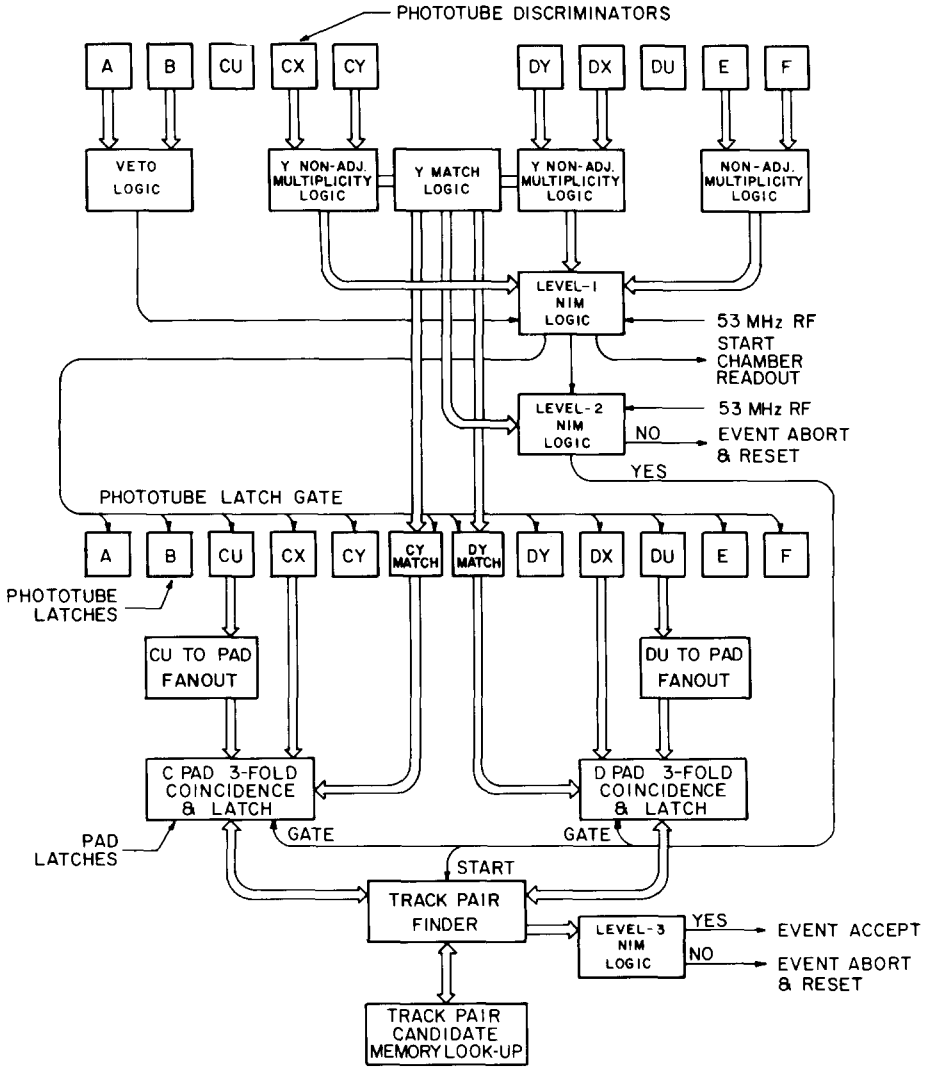

Fig. 7. Flow chart of the three levels of trigger logic.

remarks concerning trigger timing in sect. 10.1. Considerable care was taken to make certain that the trigger logic functioned properly. A summary of on-line and off-line diagnostics is deferred until sect. 12 so that an overview of the data acquisition may be given in sect. 11.

10.1. Level-1 trigger

The first-level trigger required a two-particle hit pattern in each of the C, D, E, and F banks. Events containing a halo muon in the A and B banks were vetoed. Information from the various scintillator banks was not correlated at this level except in the A- and B-hodoscope veto logic.

In order to reduce the likelihood of triggering on single particles accompanied by δ-rays, at least two non-adjacent counters were required to fire in each of the Cy, Dy, E, and F banks. If three or more counters fired in a plane, then the nonadjacency requirement was automatically satisfied for that plane. The scintillators in each of these four banks were physically split into two half-lengths but logically treated as full-length at the first and second levels of the trigger (but not at the third). This loss of information was partially offset for the C and D banks by ignoring at all trigger levels any y counter that was not overlapped by a struck x counter. (Individual counter hits were, however, latched for every bank for later off-line analysis.) The net effect of the nonadjacency requirements was to select events with

two tracks separated in the elevation (y) view at both the C and D hodoscopes and in the plan (x) view at both the E and F hodoscopes. The kernel of the Level-1 (and Level-2) ECL circuitry is sketched in fig. 8.

The A and B counters were used to veto rf buckets with incident particles traveling outside and nearly parallel to the beam. The veto requirement was a four-fold coincidence between a pair of x and y counters in one bank and a corresponding x–y pair in the other bank. To allow for small track inclinations, the counters in one bank were shifted by half of their width relative to the counters in the other bank. A and B counters that overlapped in a projection onto the x–y plane were considered corresponding counters. About 15% of the rf buckets were vetoed by the AB requirement.

In addition, we required that there be fewer than ~ 8 hits in each of the Cy, Dy, E, and F banks. Such high-multiplicity events would have accidentally satisfied the higher-level triggers but at the cost of increased detector dead-time. The resulting loss of good events was around 0.1% at high intensity, based on an extrapolation using the observed multiplicity distributions. A portion of the analog summing network that provided Level-1 and Level-2 hit-multiplicity outputs is shown in detail in fig. 9.

During the 80 GeV/c run, a further requirement was devised for the F bank in order to reduce the number of triggers due to off-momentum muons from beam-particle decays: at least one counter in the region $x <$ 0.36 m had to be struck. This requirement greatly reduced the number of triggers due to pairs of energetic negatively-charged muons (which tended to populate the positive-x side of the detector) at a small cost in acceptance. The loss of high-mass events was only 5%, while the gain in live-time was about 50%.

The Level-1 decision was reached about 40 ns after the signals emerged from the scintillation-counter discriminators. All the scintillator signals were recorded in latches and the readout of the wire-chamber planes was begun when a Level-1 trigger was issued. They latch modules (called phototube latches in fig. 7) could be read out, as well as set, via CAMAC. The also provided an output logic level for each channel which was used subsequently by the Level-3 circuitry.

10.2. Level-2 trigger

The second-level logic brought together information from the Cy and Dy counters to test the hypothesis that the tracks originated in the target.

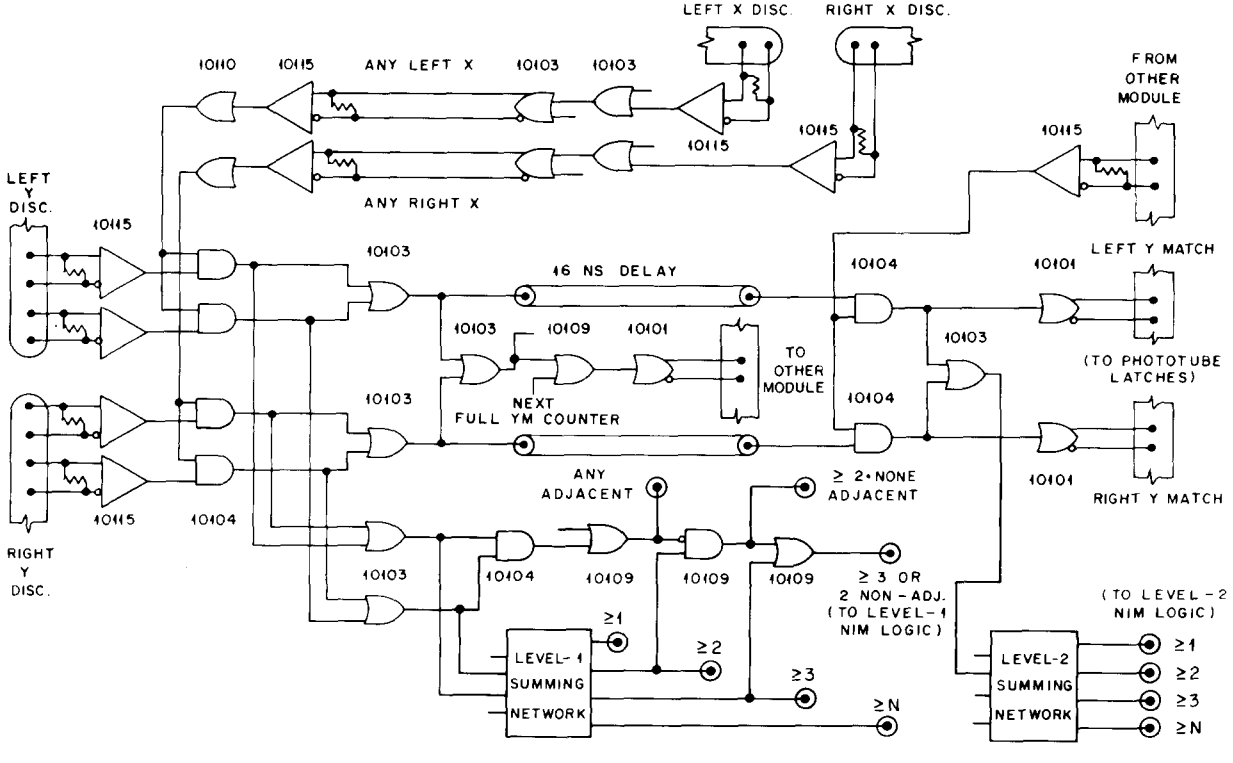

Fig. 8. Scheme of the ECL circuitry for the Level-1 and -2 C- and D-hodoscope logic. The "top" and "bottom" E- and F-counter discriminator outputs passed through Level-1 nonadjacency logic similar to that shown for the "left" and "right" Cy- and Dy-counter signals.

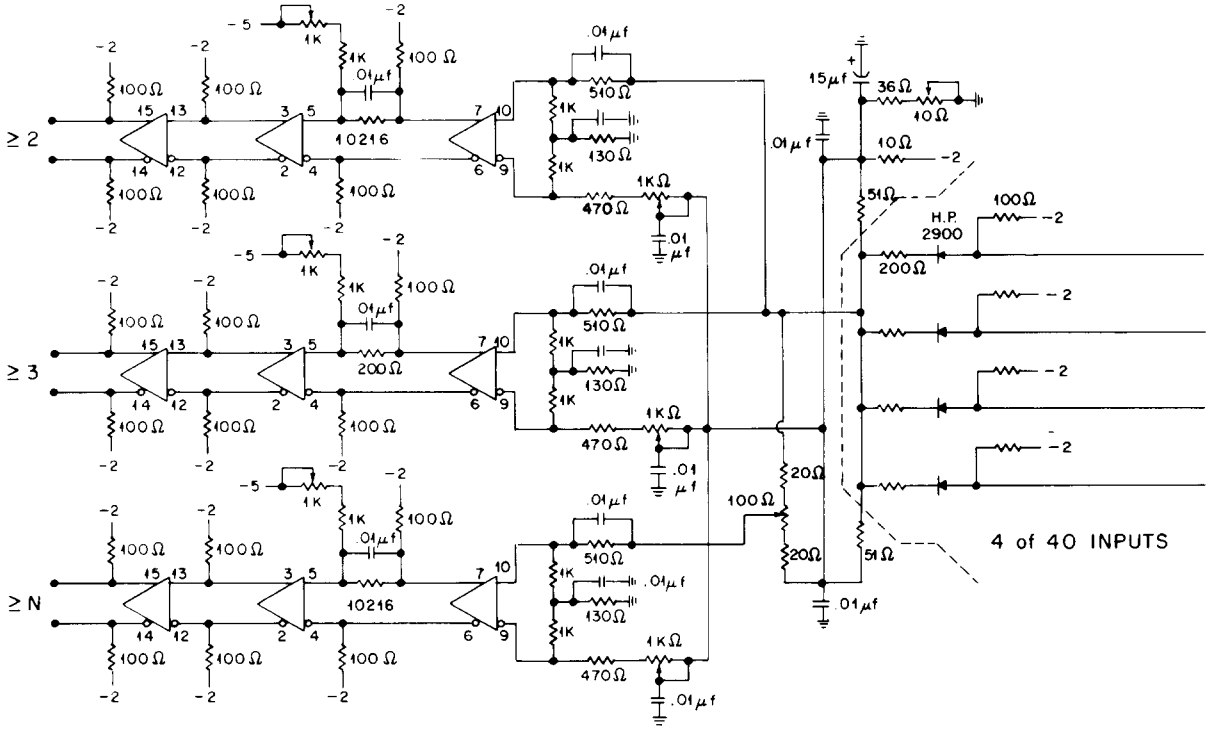

Fig. 9. A detail of the analog summing network used in the Level-1 and -2 circuitry. Outputs are shown at the left.

In the Level-2 circuitry, the 48 *y* scintillators in each of the *C* and *D* hodoscopes were clustered logically into 12 horizontal *ym* "counters" each of which spanned the full horizontal aperture. That is, if any one of the four scintillators making up a *ym* counter was hit, the *ym* counter itself was considered struck.

The widths of the *Cy* and *Dy* counters were such that a *Dy* counter subtended the same angle as a *Cy* counter when viewed from the target. Furthermore, the entire *C* hodoscope was shifted down by one *Cy* counter's width so that a muon from the target which struck *Cym* counter *i* would strike either *Dym* counter *i* or $i-1$, allowing for multiple scattering and some vertical focusing in the two magnets. *Cym* counter *i* was said to have a "match" if either *Dym* counter *i* or $i-1$ was struck. Similarly, *Dym* counter *j* had a "match" if either *Cym* counter *j* or $j+1$ was struck. This "*y*-match" logic is sketched in fig. 8.

The Level-2 trigger logic required that at least two struck *Cym* counters be found with matching *Dym* counters struck and also that at least two struck *Dym* counters be found with matching *Cym* counters struck. The two requirements are not necessarily equivalent. The Level-2 NIM circuit placed the *Cym* ≥ 2 and *Dym* ≥ 2 outputs of the *y*-match logic in coincidence with the Level-1 trigger and the accelerator rf pulse.

As noted above, all of the scintillator discriminator outputs were latched if an event passed the Level-1 trigger. In addition, the left and right halves of the *ym* counters were latched independently for subsequent use at the third trigger level. These latched bits could be non-zero only if the *ym* counter as a whole had a match in the other hodoscope (see fig. 8).

The Level-2 decision was available 10 ns after the Level-1. Including the Level-2 requirement in the Level-1 trigger would have been desirable, but the fixed-length signal cables for the drift chambers did not permit even an additional 10 ns delay in the start of the readout.

The second-level logic reduced the number of accepted triggers by a factor of about 1.5 during the 80 GeV/*c* run and 2.0 during the 255 GeV/*c* run. If the Level-2 trigger was not satisfied, a reset pulse was issued and the Level-1 logic was re-enabled.

10.3. Level-3 trigger

The third-level logic compared the information from the *C*- and *D*-scintillator banks to analyze track candidates in the magnetic bending plane. This permitted, in effect, a trigger requirement that a track pair have high invariant mass.

No mathematical computation of an invariant mass was actually made in our trigger processor, unlike some others [7,21]. Instead, a 15-bit word was formed for each track containing the "addresses" of the struck scintillators along the track's path through the two

counter banks. In principle, the 30 bits characterizing a *pair* of tracks could have furnished an address into a 2^{30}-bit memory, each location of which had been programmed to contain a 1 or 0 depending on whether the candidate pair was to be accepted or rejected, based on prior off-line calculations of such a pair's invariant mass. In practice, a fast lookup into only $\sim 2^{18}$ bits of memory sufficed because the selection magnet and, to a lesser extent, the analyzing magnet bent the tracks into configurations that could be well described by 18 bits of information. This reduction was achieved by taking various differences among the addresses of the struck counters in the C and D banks, as described below.

The "raw data" used at the third trigger level consisted entirely of the output logic levels from the C- and D-hodoscope phototube latches. A threefold coincidence of x, ym, and u counters created rectangular "pads" in each hodoscope * as illustrated in fig. 6b and in fig. 10, which depicts an actual event from the point of view of the third-level processor. These coincidences were latched in the pad-latch modules (fig. 7), which also could be read out or written into via CAMAC.

Given a set of struck pads, a hardware track-pair finder assigned a pair of pads, one in the C hodoscope and one in the D hodoscope, to a candidate particle trajectory. For the muon-pair trigger, such track candidates were constructed two at a time and labelled α and β. The track candidates were then presented to memory-lookup (MLU) modules which rendered a trigger decision.

10.3.1. Track-pair-finder

The track-pair finder encoded the parameters of each track candidate into a 15-bit word. The lowest order 5 bits specified the x coordinate of the struck pad in the C hodoscope, the next 6 bits gave the x coordinate of the struck pad in the D hodoscope, and the remaining 4 bits identified the y coordinate of the C-hodoscope pad. It was not necessary to encode the y coordinate of the struck pad in the D hodoscope because the track-pair finder associated a given C-hodoscope pad with only those pads in the D hodoscope belonging to ym counters that were roughly on a line with the target, in keeping with the Level-2 logic. Furthermore, the track-pair finder did not permit the α- and β-track pads to share a y address in either hodoscope, in conformity with the Level-1 nonadjacency requirements.

As can be inferred from fig. 10, quite often many candidate track pairs could be constructed for a given event (though only the tracks later reconstructed using the wire chambers are shown in the figure). Accordingly, the track-pair finder not only performed an encoding function but looped over all possible pairs of candidate tracks, subject to the Level-1 and -2 trigger requirements.

* In the C hodoscope, the left and right halves of only the upper 11 ym counters entered the third-level logic.

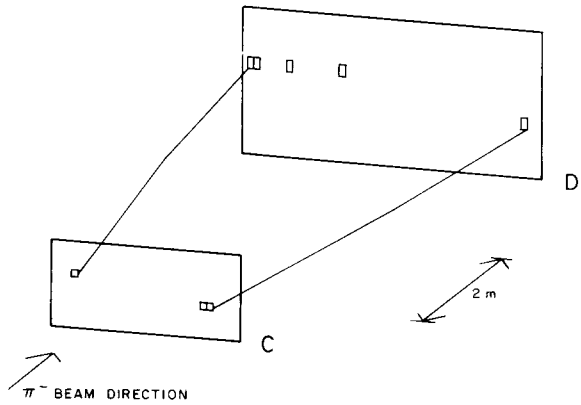

Fig. 10. View of the struck scintillator "pads" presented to the Level-3 track-pair finder in a typical event.

Fig. 11 illustrates the flow of data between the track-pair finder (in the center of the figure) and the pad-latch modules (on either side) and between the track-pair finder and the memory-lookup modules (beneath). The track pair finder consisted of eight functionally-similar circuits, labelled by hodoscope (C or D), pad projection (x or y), and track (α or β). Six of the circuits provided the MLUs with the requisite binary addresses of the projections of the candidate-tracks' "struck" pads. The actual encoding was done with ECL 10165 integrated circuits. "New data" and "carry" lines between the circuits provided a means of stepping through the struck pads in an orderly and efficient way. The first four circuits (shown within a dashed line) belonged physically to one module, called the y module. The remaining four circuits were each housed in individual x modules. A representative x-module circuit ($Dx\alpha$) is diagrammed in fig. 13, while the more complicated circuit for the y module is sketched in fig. 12. A photograph of the five modules is displayed in fig. 14.

The construction of a candidate track pair began with the elevation (y) view. The inputs to the y module were single-ended ECL output levels from the pad-latch modules indicating which horizontal (y) rows of pads had a hit. A Level-2 trigger caused the four circuits of the y module to assign struck rows (labelled $Cy\alpha$, $Dy\alpha$, $Cy\beta$, and $Dy\beta$) to the tracks, beginning with the lowest valid rows. The addresses of two of the assigned rows, $Cy\alpha$ and $Cy\beta$, were encoded for use by the MLUs. As each of the four rows was found, an x-enable line was asserted causing the pad-latch bits for that row to be placed on one of the four busses (fig. 11) that provided single-ended ECL inputs to the x modules. The x

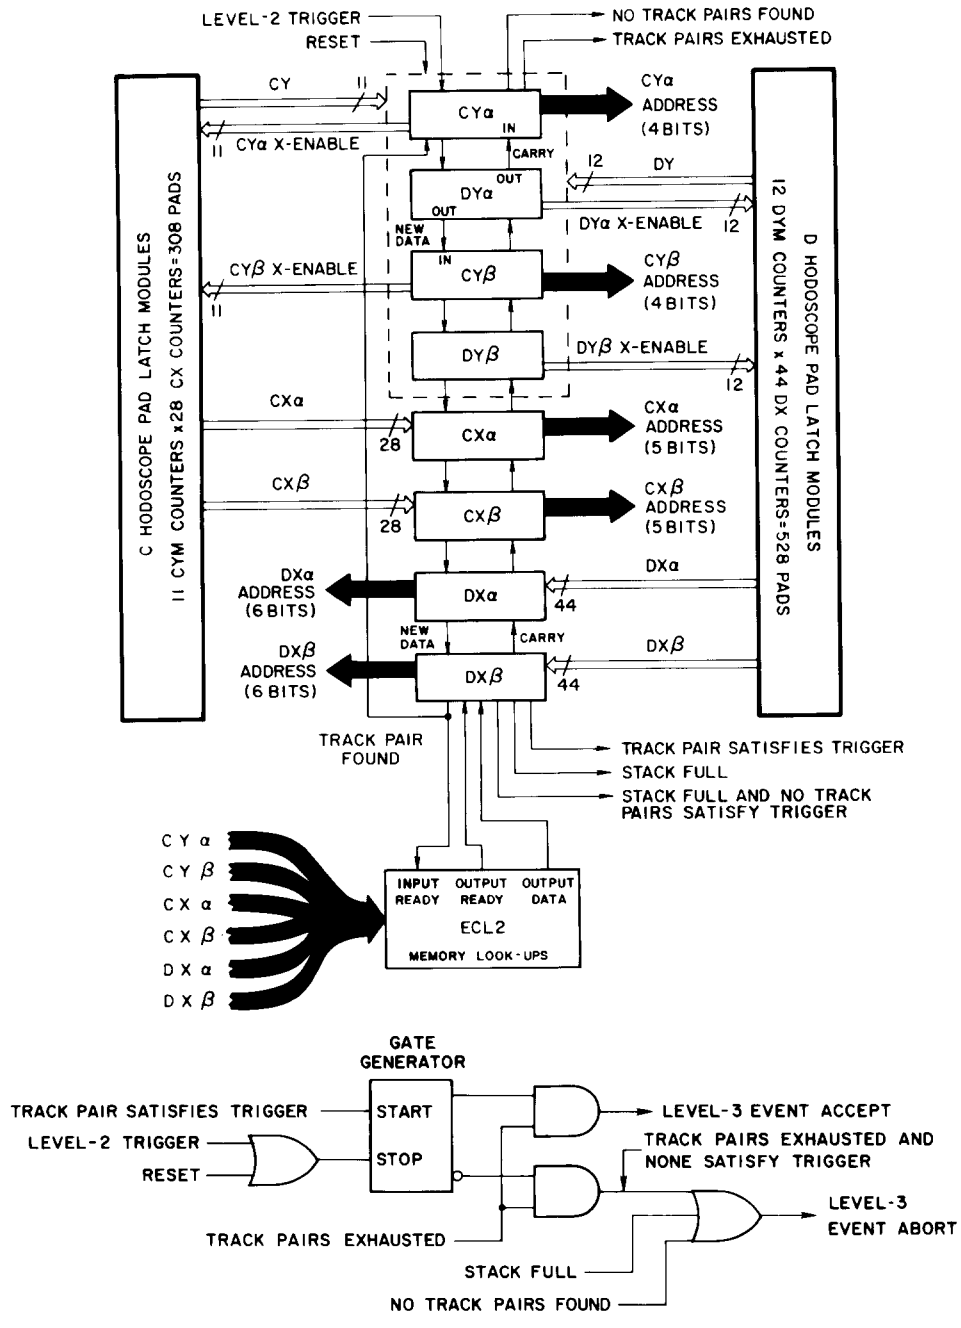

Fig. 11. Flow chart of the track-pair finder. Also shown are data and control lines to and from the pad-latch modules and the memory-lookup modules. The track-pair finder sent four of the five outputs, shown at the right, to the Level-3 NIM logic, shown at the bottom of the figure.

modules in turn assigned the first available hit in each selected row to a track and encoded the column addresses of these hits, thereby completing the specification of a candidate track pair.

The re-enforcement of lower-level trigger requirements took place in the y module. Because the α and β tracks were not allowed to coincide in the y projection (unlike the x) at either scintillator plane, the two x-enable lines asserted by the $Cy\alpha$ and $Dy\alpha$ circuits were also used internally to veto the corresponding inputs to

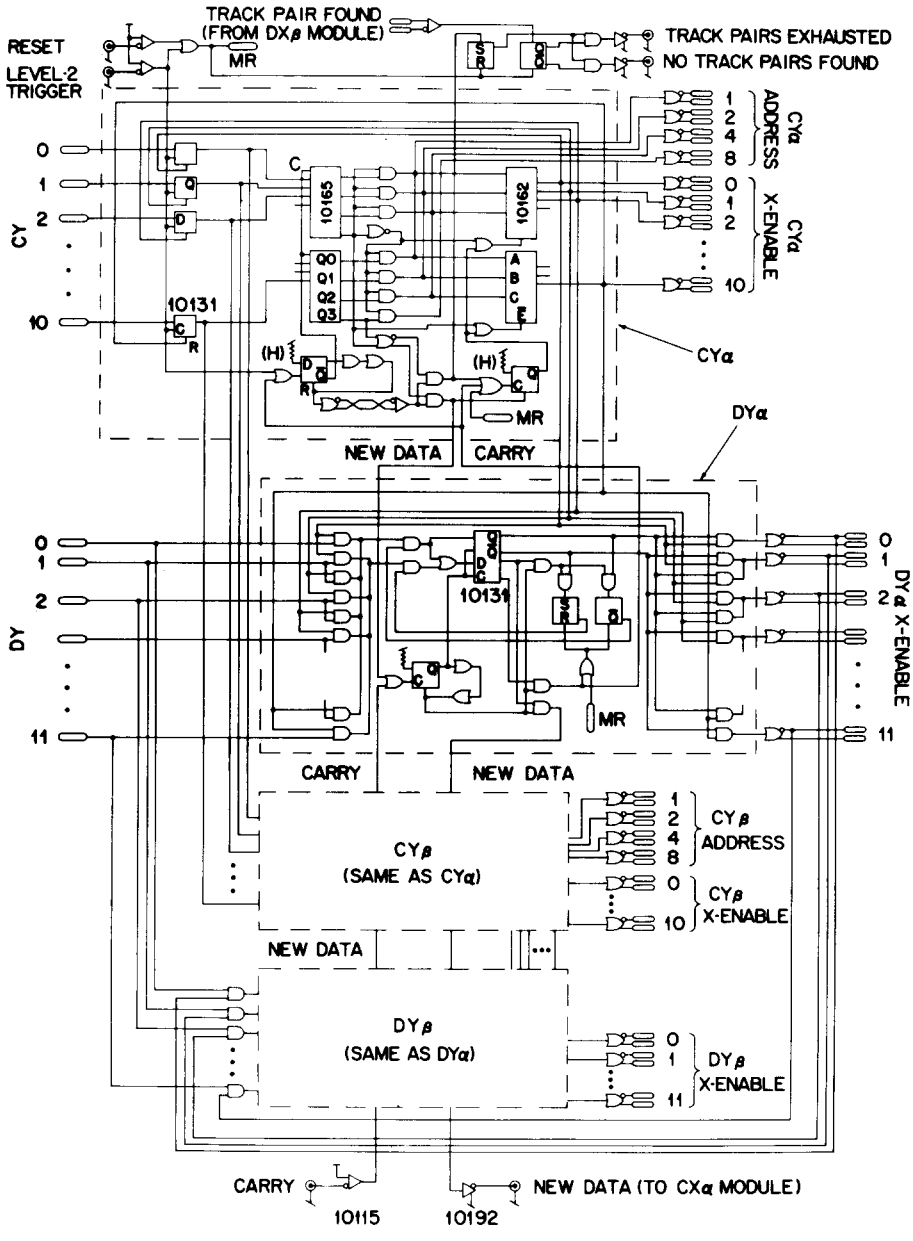

Fig. 12. Scheme of the ECL circuitry of the y module of the track-pair finder.

the $Cy\beta$ and $Dy\beta$ circuits, as shown in fig. 12. In addition, the $Cy\alpha$ and $Cy\beta$ x-enable levels were used internally to apply the Level-2 y-match requirement in the selection of y rows by the $Dy\alpha$ and $Dy\beta$ circuits. These two constraints necessitated wiring the four y circuits on a single board.

Loop control was maintained throughout the track-pair finder as illustrated in the following example. If circuit $Cy\alpha$ found valid inputs, it sent a "new data" pulse to the next circuit down the line ($Dy\alpha$). That circuit in turn passed a "new data" signal to its successor ($Cy\beta$) if it found valid inputs; otherwise, it sent a "carry" pulse back to the previous circuit ($Cy\alpha$). In the latter case, circuit $Cy\alpha$, having already removed the lowest-order input from further consideration with a reset of the input flip-flop, would then encode the current lowest-order input and reissue "new data". A "new data" signal from the last circuit in the chain ($Dx\beta$) signified that a track-pair candidate had been found.

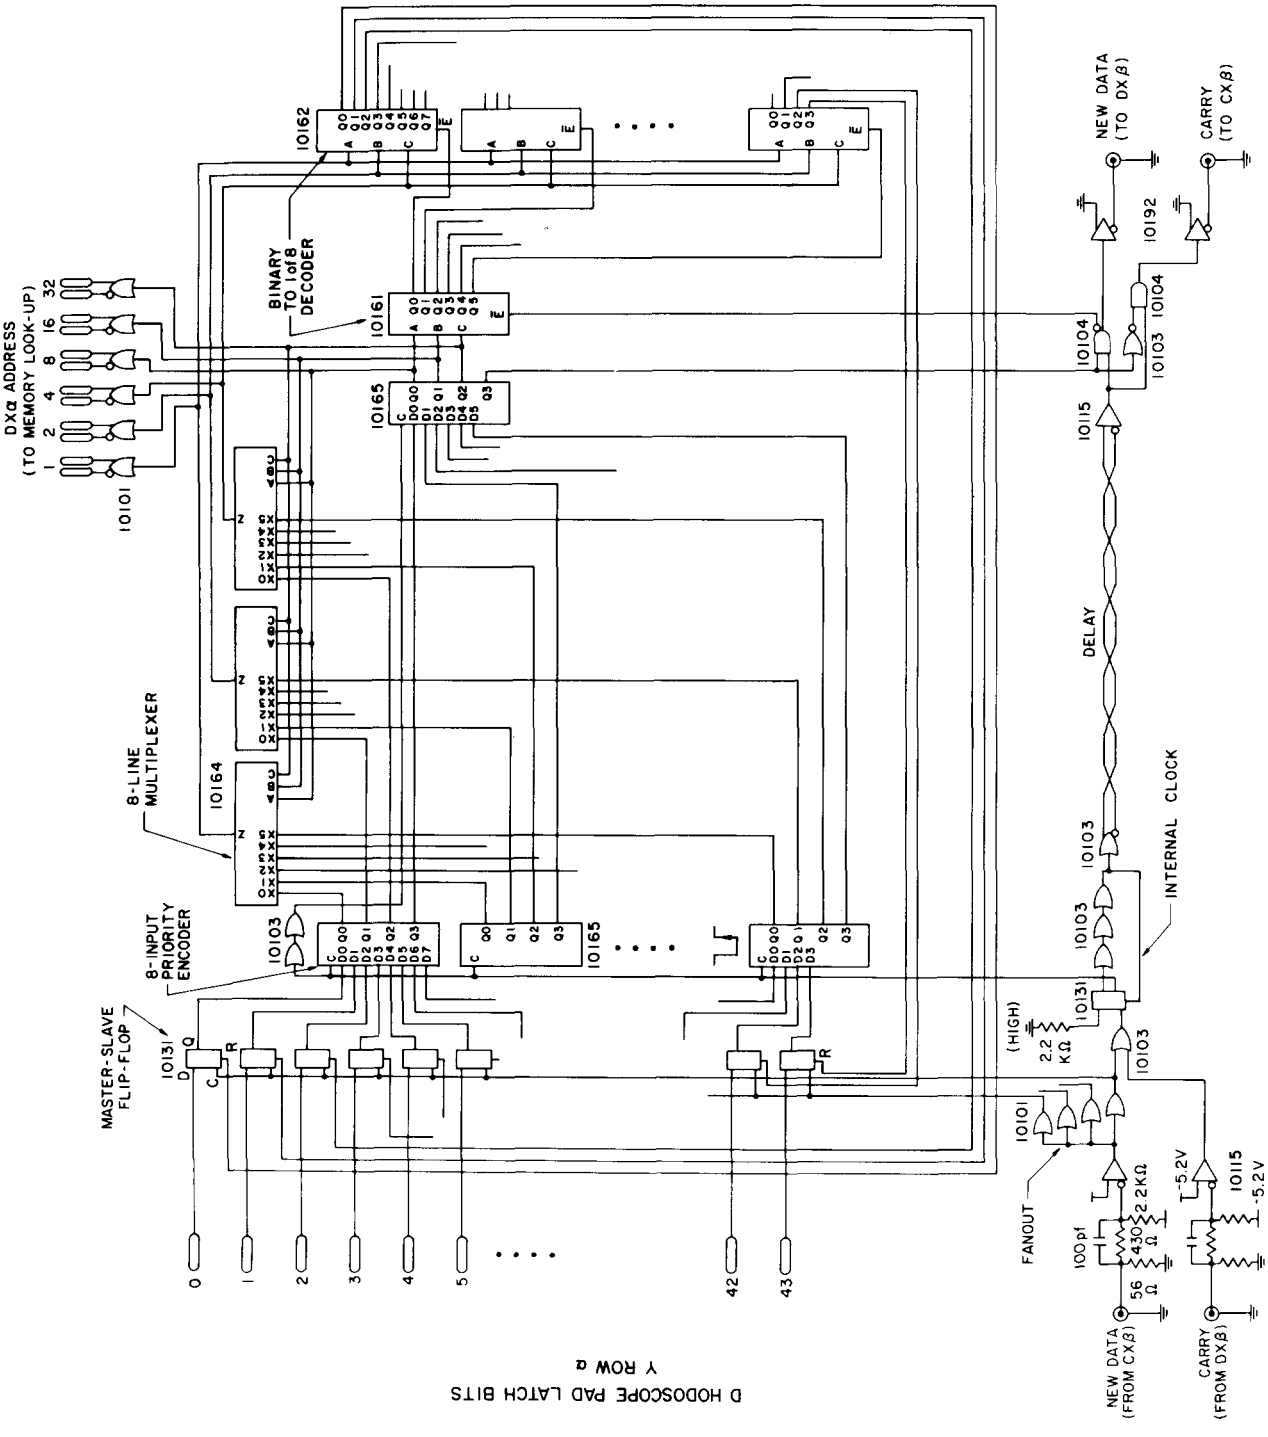

Fig. 13. Scheme of the ECL circuitry of one of the four x modules of the track-pair finder.

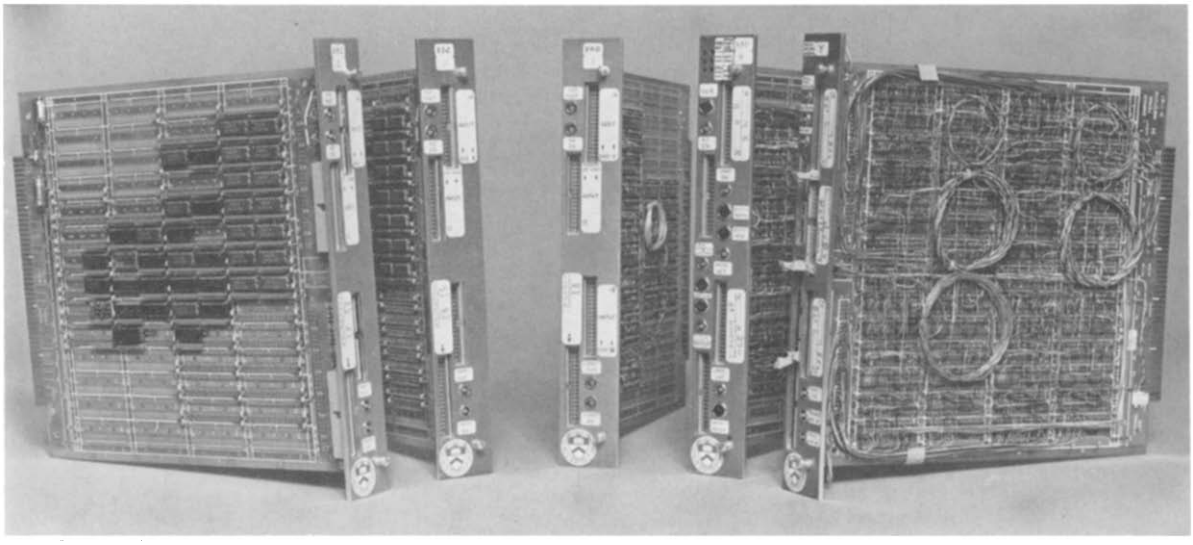

Fig. 14. The five modules of the track-pair finder photographed prior to shipment from Princeton to Fermilab. Wiring was done using an insulation-displacement technique.

The 30 bits describing a newly constructed candidate track-pair were presented to Fermilab ECL2 memory-lookup modules [22] where a comparison was made with stored patterns (sect. 10.3.2). Upon completion of the comparison, an "output ready" pulse was sent to the "carry" input of module $Dx\beta$ (see fig. 11), initiating the construction of another track-pair candidate differing from the previous one, perhaps, only in the $Dx\beta$ address.

The conditions under which the loop was terminated depended upon the setting of switches inside the $Dx\beta$ module. In one mode of operation, the track-pair finder quit as soon as a track-pair candidate satisfied the memory lookup. In another, the track-pair finder continued to construct track pairs regardless of the outcome of the memory-lookup comparison for each candidate track pair. In either mode, a switch-selectable limit could be placed on the number of track-pair candidates constructed.

Except during special calibration runs, we allowed the trigger processor to construct up to 16 track-pair candidates per event, independent of the outcome of each memory lookup. This latter mode of operation enabled us to reject an event whenever the "stack full" condition of ≥ 16 pairs was reached. The cut at 15 track-pair candidates eliminated noisy events that would eventually have found their way through the third-level trigger. We estimated that the resulting loss of efficiency was in the range 2–5%, depending on the beam intensity. Fig. 15 shows a typical distribution of the multiplicity of track pairs presented to the MLUs.

A Level-3 trigger was issued only if at least one track-pair candidate satisfied the memory-lookup comparison and the number of track-pair candidates was less than 16. The state of the track-pair finder after quitting the search for track-pair candidates was signaled by one of three mutually-exclusive NIM outputs, labelled "no track pairs found", "track pairs exhausted", and "stack full" in fig. 11. The first and third always led to event aborts while the second led to an abort only if none of the track-pair candidates satisfied the memory lookup. These three outputs were combined with a fourth, labelled "track pair satisfies trigger", in the Level-3 NIM logic depicted at the bottom of fig. 11. While the first three outputs were asserted only upon

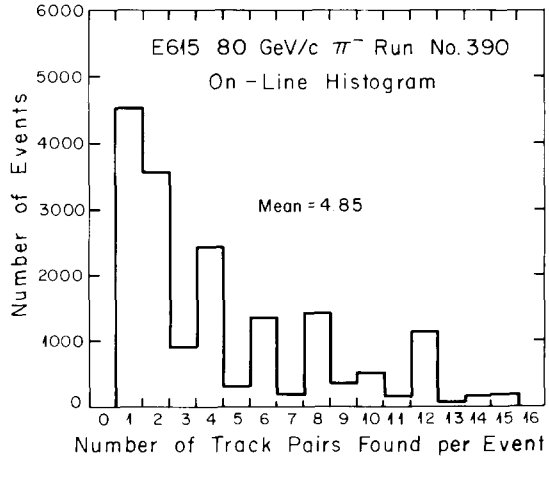

Fig. 15. Histogram of the number of candidate track pairs found by the Level-3 logic.

completion of the loop, the fourth one was set either "true" or "false" for each track-pair combination; thus a flip-flop (NIM gate generator) was required to produce a signal indicating whether or not *some* track-pair candidate had satisfied the memory lookup. A fifth NIM output, labelled "stack full and no track pairs satisfy trigger", was essential to trigger-logic diagnostics. Whenever one of the five track-pair-finder NIM outputs was asserted, a corresponding bit was set in an input register whose contents were included in the event record written to tape. These trigger-status bits (among others) enabled us to monitor the performance of the trigger on-line as well as off-line.

Typical performance of the Level-3 logic was as follows. "Stack full" aborts occurred roughly 16% of the time. About 50% of the Level-2 triggers were aborted because of "no track pairs found". This was attributable to the tighter timing requirements enforced at Level-3, as discussed in sect. 10.4. Finally, some 28% of the Level-2 triggers were aborted because none of the fewer-than-16 track-pair candidates satisfied the memory lookup ("track pairs exhausted and none satisfy trigger"). If one or more of the track pairs tested *did* fit an allowable pattern as stored in the MLUs *and* no more than 15 track-pair candidates could be constructed for the event, then, of course, a Level-3 trigger was issued and the event was recorded by the on-line computer.

Adding the above three figures, we see that 94% of the second-level triggers were discarded at the third level, corresponding to a rejection factor of 16.7. In fact, the number of triggers was reduced by a factor of about 12 during the 80 GeV/c run and a factor of about 28 during the 255 GeV/c run. The *combined* Level-2 and Level-3 rejection factors ranged between 17 and 20 during the 80 GeV/c run and between 50 and 65 during the 255 GeV/c run.

10.3.2. Memory-lookup procedure

The MLUs were interconnected as shown in fig. 16. They resided in a nonstandard CAMAC crate [22] and could be loaded and tested by the on-line computer.

The memory lookup took place in three logical stages, permitting a 30-bit lookup with only about 2^{18} bits of memory. This was possible because the 30 bits were not all equally significant. The large momentum kick of the

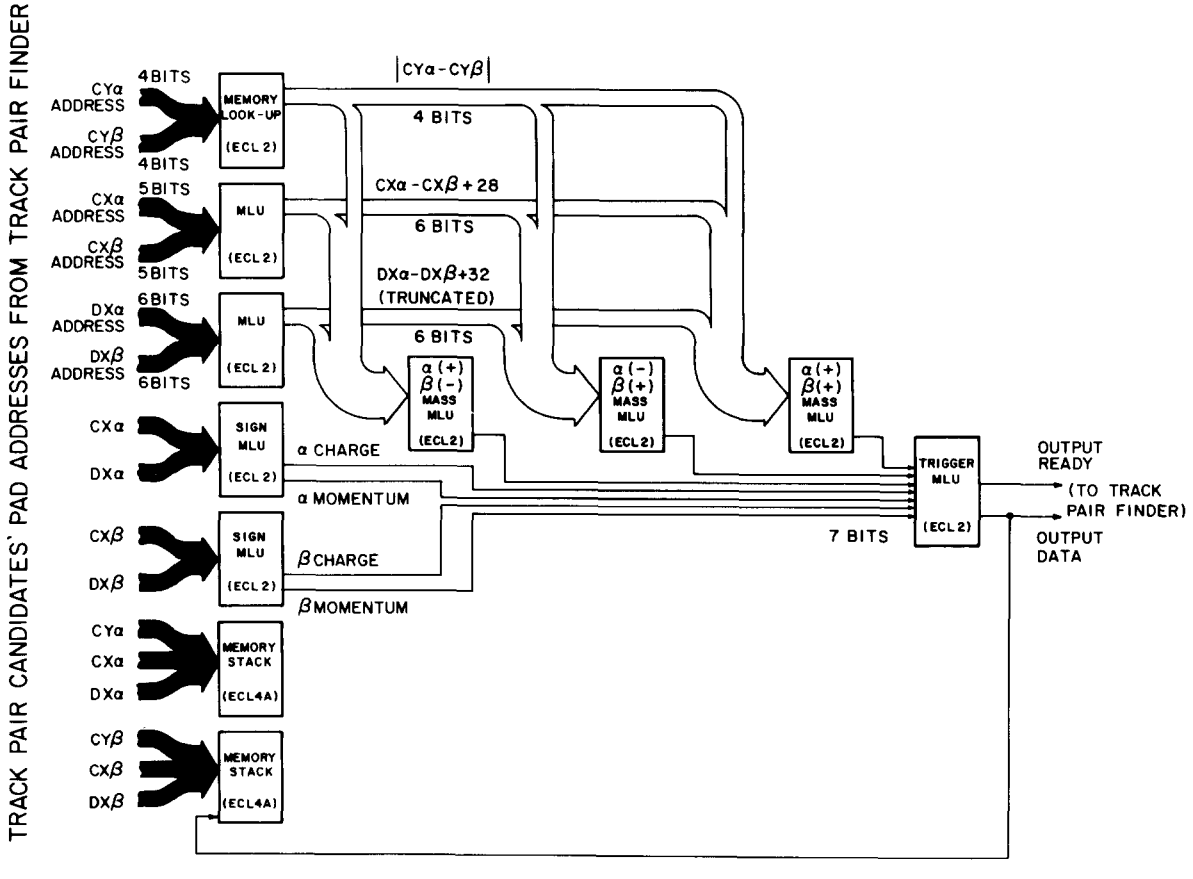

Fig. 16. Flow chart of the memory-lookup phase of the Level-3 logic.

selection magnet gave muon pairs a topology in which the pair invariant mass was, to a good approximation, a function only of the track separations in the C and D hodoscopes.

Three MLUs were programmed with the results of calculations of all possible track separations. The y-view track separation $|Cy\alpha - Cy\beta|$, which required a 4-bit output, was calculated in one MLU. (The Level-2 trigger requirement made the corresponding calculation for the D hodoscope superfluous.) For the x view, 6-bit results were calculated and loaded in two other MLUs for $Cx\alpha - Cx\beta$ and for $Dx\alpha - Dx\beta$. Both positive and negative differences were kept by adding 28 and 32 to the respective results. Because there were 44 Dx counters, the latter subtraction actually required 7 bits. However, if the counter difference was greater than 31 (or less than -31), we truncated the result to 31 (or -31) with very little loss in the effectiveness of the trigger. The three calculations proceeded in parallel and reduced the initial 30 bits to 16.

In the second memory-lookup stage, the reduced 16-bit pattern was presented to a 2^{16}-bit memory in which was stored a 1-bit "yes" or "no" decision for all possible patterns.

In practice, a decision on the 16-bit input pattern depended on 2 additional bits of information, the signs of the charges of the α and β tracks. These were obtained from two more MLUs ("Sign" MLUs in fig. 16) belonging to the first stage. Because the selection magnet deflected the muons in the x view, the sign of the charge could be obtained from the 11 bits in the Cx and Dx addresses of each track. Though the granularity of the hodoscope banks was too coarse to make an accurate momentum measurement, the sign of the charge was determined correctly for over 99.5% of all tracks.

Because the 16-bit pattern presented at the second memory-lookup stage did not include the signs of the charges, four different MLUs were needed to store the decisions appropriate to the 4 charge combinations $++$, $+-$, $-+$, and $--$. In fact, we decided never to trigger on a pair in which both tracks had the same charge as the beam particles. Thus only three 2^{16}-bit "mass" MLUs were used.

The three mass-MLU decision bits were combined with the two track-sign bits in a single MLU at the third stage (called the "trigger" MLU). The mass-MLU decision bit corresponding to the actual track charges became the output trigger bit, provided each track satisfied one further criterion.

We took advantage of the power of the MLU scheme to include two other "quality" bits in the final decision. Many (Cx, Dx) combinations could not be generated by tracks from high-mass muon pairs. The MLUs that determined the signs of the charges were easily programmed to assign an additional bit to each track, called "momentum" in fig. 16, signifying whether or not the muon followed a kinematically allowed trajectory.

Setting the contents of the sign MLUs and the second-stage mass MLUs required a detailed understanding of the apparatus. This was obtained in the test run of the experiment in 1982, which did not use the Level-3 trigger. By the start of the main data collection in 1983, full analysis programs were available and were used to verify the programming of the MLUs through the analysis of data taken with only the Level-2 trigger. In this way the boundaries around kinematically allowed regions for + and − tracks and between "yes" and "no" mass-MLU decisions were altered slightly from initial estimates based on a Monte Carlo simulation.

Though the primary purpose of the mass MLUs was to exclude certain track-pair topologies on the basis of kinematics, we found it expedient to exclude one class of track-pair topologies simply on the basis of geometry. Because the u counters were arranged in two completely-overlapping planes (fig. 6b), two u counters were struck for every particle passing through the bank. As a result, when the projections of a pair of hits were close together, the u counters could not always resolve the twofold ambiguity as to the actual track intercepts in the x- and y-counter banks. Accordingly, we programmed the mass MLUs to reject track-pair candidates whenever the vertical track separation at the C hodoscope was a minimum, that is, whenever the 4-bit input $|Cy\alpha - Cy\beta|$ was equal to 1. Although this cut reduced the geometric acceptance of the detector by ~ 22%, the number of track-pair candidates constructed using "ghost" hits was also reduced, thereby improving trigger rates and the quality of data written to tape.

For diagnostic purposes, all of the 30-bit patterns describing the candidate track pairs constructed for an event were stored in ECL4A memory-stack modules [22]. Each pair's trigger-MLU output bit was also stored. These 31-bit words were written on magnetic tape by the on-line computer as part of the event data record.

10.3.3. Level-3 trigger dead-time

Following a Level-2 trigger, the track-pair finder took a minimum of 300 ns to produce a candidate track pair (less than 40 ns per circuit). The memory-lookup sequence required approximately 250 ns per track pair. After receiving a "carry" signal from the MLUs, the track-pair finder took another 330 ns to declare that all track pairs had been exhausted. Thus the Level-3 dead-time for an event that produced only one track-pair candidate was 940 ns (including NIM-circuit delays). An event with 5 track-pair candidates therefore introduced a dead-time of about 2.0 μs whether or not it satisfied the memory lookup. Events aborted on the "stack full" condition cost the experiment about 4.9 μs each, while events aborted because no track pairs could be found took as little as 90 ns.

Weighting the dead-times for the three categories of aborted events by the percentages given above (sect. 10.3.1) and adding 6% of 2 μs for the accepted events, one finds that the dead-time introduced by the Level-3 circuitry was only about 1–2 μs per Level-2 trigger and was less than 1% of the total dead-time of the experiment.

10.4. Trigger timing

The 18.8 ns rf cycle determined the time scale in which the various elements of the trigger had to operate. The distribution of arrival times of scintillator pulses had a fwhm of about 5 ns, though the tails extended the timing spread to 15–20 ns. The scintillator pulses were discriminated to produce logic signals with widths between 25 and 35 ns, since the Level-1 and -2 logic was found to work at full efficiency when tested with 15 ns wide signals that arrived simultaneously.

Because the logic signals were so wide, the first two trigger levels were not always able to reject pairs formed by muons belonging to different, consecutive rf buckets. However, the third-level trigger was very effective in eliminating such background pairs because it was based on scintillator signals that were latched on the trailing edge of a 5 ns wide gate. The lower-level logic, as noted above, required its input signals to overlap for 15 ns.

10.5. Effectiveness of the trigger

The trigger rates for two representative runs presented in table 3 summarize the overall effectiveness of the trigger logic in reducing the number of events written to tape.

An indication of the effect of the second and third levels of trigger processing on the acceptance of the detector is provided by fig. 17. The solid curve is the observed $\mu^+\mu^-$ invariant-mass spectrum for prescaled Level-1 triggers recorded during the 255 GeV/c π^- run. Level-1 triggers that also passed the Level-3 trigger, as signified by the Level-3 trigger-status bits, are shown by a dashed curve. The higher-level trigger processors strongly suppressed low-mass muon pairs without affecting the shape of the high-mass distribution. The overall acceptance at low mass (below 2.0 GeV/c^2) was reduced by a factor of 32. At high mass (above 4.0 GeV/c^2), the detection efficiency was reduced by about 40% due to several effects: among them, somewhat inefficient *Cu* and *Du* scintillation counters (sect. 13) which entered the third-level but not the lower-level trigger logic, the Level-3 memory-lookup requirements themselves (in particular the $|Cy\alpha - Cy\beta| \geq 1$ cut), and the Level-3 track-pair multiplicity cut.

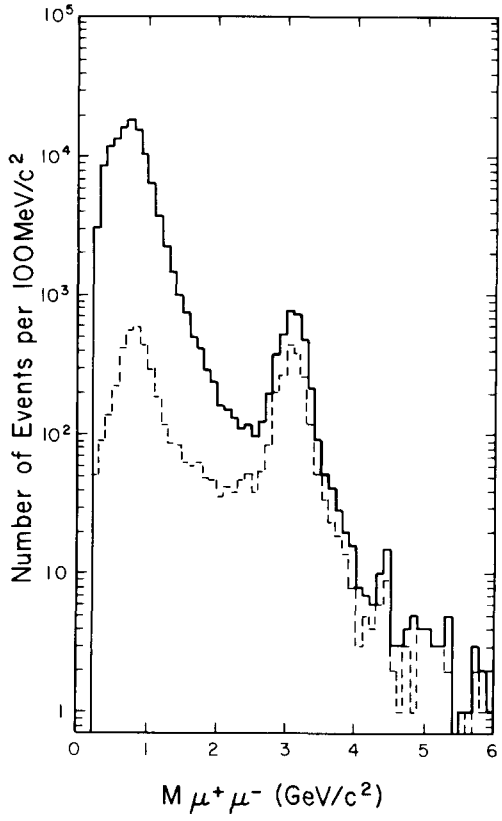

Fig. 17. The $\mu^+\mu^-$ pair invariant-mass spectrum for the entire sample of Level-1 triggers collected during the 255 GeV/c π^- run (solid curve) compared to the subsample of Level-1 triggers that also satisfied the second- and third-level trigger requirements (dashed curve).

Table 3
Typical trigger rates (per beam spill).

	80 GeV/c run no. 390	255 GeV/c run no. 490
Protons	4.9×10^{12}	3.6×10^{12}
π^-	2.7×10^{9}	2.0×10^{9}
Level-1	2.4×10^{4}	6.3×10^{4}
Level-2	1.6×10^{4}	3.3×10^{4}
Level-3	1390	1260
Live-time	66%	59%

11. Data acquisition

Data acquisition was performed by a PDP 11/45 computer equipped with a bipolar memory unit for fast program execution, as well as Digital Pathways, Inc. bank-switchable memory modules for temporary event storage. The computer was interfaced to the data-acquisition electronics via two CAMAC branches with a total of 14 crates.

The data-acquisition program was a version of MULTI [23] supported by Fermilab. Continuous moni-

toring of the experiment was provided by event displays of scintillator hits, chamber hits, and trigger-status bits, as well as by numerous diagnostic histograms and scatter plots. By rewriting the MULTI histogramming package in assembly language, we achieved a factor of 40 improvement in speed over the standard version.

The events were written onto high-density tapes (6250 bpi), each containing about 180 000 triggers. For each event all wire-chamber hits, phototube-latch words, pad-latch threefold coincidences, detector-magnet currents, and trigger-logic status bits were recorded on tape. Valid trigger-processor track-pair candidate data from the memory-stack modules were also recorded for each Level-3 trigger and for every prescaled Level-1 trigger that passed the Level-2 requirements (and thereby advanced to the Level-3 logic). During normal running conditions, a typical event record consisted of 450 16-bit words. End-of-spill records contained scaler and ionization-chamber sums and Hall-probe measurements of the field strengths of the analysis and selection magnets. The contents of the MLUs were written to tape at the beginning and at the end of each run so that the integrity of the stored patterns could be verified and the patterns themselves known precisely for every run.

The dead-time associated with data acquisition by the computer was roughly 2 ms per event (0.75 ms minimum).

12. Trigger logic diagnostics

The trigger logic was tested with event-simulation programs executed (automatically) by the on-line computer at the beginning of each data run. A thorough check could be performed because every discriminator and latch could be loaded by the computer with patterns simulating the full complexity of signals to be processed by all three trigger levels. Prior to executing the event-simulation routines, the computer verified the contents of the MLUs.

The Level-3 logic was subjected to lengthier tests during the weekly shut-down periods and found to process as many as 100 000 simulated events without an error. The performance of the chamber readout was also checked during shut-downs.

So that electronic failures could be localized quickly and easily, subsystems at each trigger level could be tested individually. For instance, the three subsystems within the third-level logic (the pad-latch modules, the track pair finder modules, and the memory-lookup and memory-stack modules) could be tested independently or in series. This flexibility of the hardware and software was indispensable during the debugging of the trigger.

The scintillation-counter latch bits written on tape enabled us to monitor the performance of the trigger by off-line analysis within hours of the data collection. The off-line tests were similar to the on-line diagnostics but had the advantage that the recorded scintillator information reflected actual signal-timing variations. The off-line tests using Level-3 triggers had the disadvantage, of course, of testing only events that actually fulfilled all the hardware-trigger requirements. Tests using prescaled Level-1 triggers had no such bias, but were statistically less significant when applied on a run-by-run basis.

Typically 98.5% of the Level-3 triggers written on tape satisfied the Level-1 and Level-2 logic, based upon an application of the Level-1 requirements to the latched scintillator hits and an application of the Level-2 y-match multiplicity requirements to the latched matches. The E- and F-bank trigger requirements accounted for about 90% of the failures. Because the latched E- and F-counter signals were not used in the trigger, unlike the latched C- and D-counter signals, small timing differences between the individual counter-bank triggers and latches were expected to result in failures more often associated with the E and F counters than with the C or D.

Because the *Cym* and *Dym* counters that had matches found in the Level-2 logic were latched in addition to the scintillators themselves, a test of the consistency of the latches could also be performed. Discrepancies between the latched matches and those computed from the individual scintillator hits were present at a rate of 2% during the 80 GeV/c run and 6% during the 255 GeV/c run. The errors were equally divided between the cases of too many and too few matches. Since the trigger processors always worked correctly with simulated patterns, the problem was associated with the time of arrival of the pulses. Triggers on muons belonging to different rf buckets accounted for both effects, since the latches tended to record delayed pulses, while the trigger logic was more sensitive to signals from early buckets.

The pad latches were found consistent with the x-, ym-, and u-counter phototube latches for all but 0.1% of the events tested. This small discrepancy was completely understood in terms of the latching and gating circuitry in the phototube- and pad-latch modules.

The Level-3 trigger processor was extremely reliable. Starting from the recorded pad latch data, the off-line program calculated the patterns that should have been read out of the memory-stack modules for each event. Discrepancies with the actual patterns were found in less than 0.001% of the triggers, consistent with observations made using the on-line testing programs. This was to be expected, since the Level-3 trigger was based on latched pad information to begin with. In this respect, the off-line tests of the third-level logic differed qualitatively from the tests of the lower-level logic, which were also tests of latching efficiency.

13. Scintillator efficiency studies

The voltages for the scintillation-counter photomultipliers were set so that the signals from the conversion of gamma rays emitted by a ^{60}Co source had uniform and ample pulse heights. Exploiting the fact that ^{60}Co emits two gammas simultaneously, we timed the counters relative to one another by placing the same radioactive source between each scintillator and a movable test counter. This procedure was repeated between the 80 and 255 GeV/c runs.

The scintillation counters suffered some inefficiency during high-intensity running. Typical counting rates during the 80 GeV/c run were as follows:
- 15 MHz in the C bank and in the MWPCs,
- 12 MHz in the D bank and in the downstream drift chambers.

Some counters were exposed to a much higher flux than others. The highest rate recorded in an individual counter in each of the planes of the C bank was as follows:
- 2.5 MHz in a Cx counter,
- 2.3 MHz in a Cy counter,
- 3.6 MHz in a Cu counter.

The maximum rates during the 255 GeV/c run were about two-thirds of the above.

To measure the scintillator efficiency at high intensity, we made a set of special runs with simplified versions of the trigger. We took data with a Level-3 requirement but without the E or F banks in the Level-1 trigger. In other runs, we took data with only the Level-1 requirement but without the C or D banks. During part of each special trigger run, the polarity of the selection magnet was reversed in order to illuminate the whole detector evenly with tracks. Reconstructed tracks were used to measure the efficiencies of counters in the banks not used in the trigger.

While the efficiencies for most banks were quite good (better than 99% on the average), we found that the efficiency was lower for the Cu and Du counters: their average efficiencies were 95 and 97%, respectively. The inefficiencies were not localized in single counters. Though they were fairly uniform along the vertical direction, they were greater on the (positive-x) side of the detector where the particle flux was higher due to muons from beam-particle decays. Fig. 18 shows, for example, the efficiency of the D bank along the x projection. The inefficiency was observed to increase linearly with the intensity, almost disappearing in the runs taken at low intensity, as shown by fig. 19.

This behavior was consistent with the known limitations of the Zener-diode-stabilized phototube bases. Because the areas of the Cu and Du counters were almost twice those of counters in the corresponding x and y banks, the u counters experienced the highest rates. They were subject to average rates as large as 3–4 MHz per counter, and instantaneous rates which were more than twice as high during short time intervals because of the spill structure.

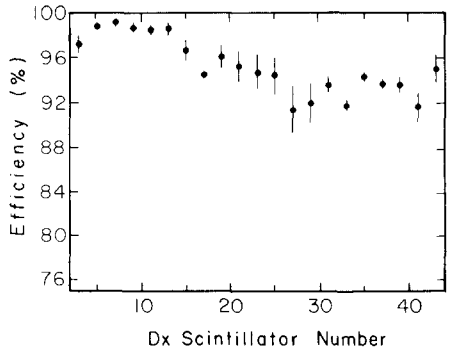

Fig. 18. The x projection of the D scintillation-counter hodoscope efficiency as measured during the 80 GeV/c run, in bins two Dx counters wide. The dip in the central region was due to particularly high beam rates in certain Du counters.

The Monte Carlo program used to compute the acceptance and the resolution of the spectrometer (sects. 14 and 15) also simulated the inefficiency of the apparatus. Fig. 20 compares, for example, the y projection of the C scintillator-bank illumination for simulated and real data, using events with invariant mass larger than 4 GeV/c^2.

Track reconstruction in the wire chambers was little affected by the high particle intensity because of the large number of chamber planes. The efficiency for the reconstruction of a particle pair was almost 99%. Fig. 21 exhibits a typical MWPC efficiency plot. A dip is seen where the intensity was highest. Fig. 22 gives an example of the efficiency across a drift-chamber plane.

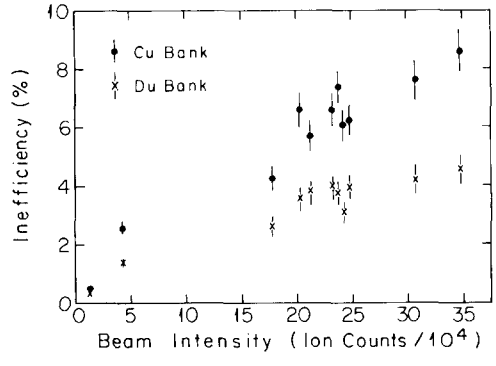

Fig. 19. The observed average inefficiency in the Cu and Du scintillator banks as a function of the beam intensity recorded by an ionization chamber during the 80 GeV/c run. Normal operating beam intensity was $\sim 20\times10^4$ ion counts/spill, corresponding to $\sim 2.5\times10^9$ pions/spill.

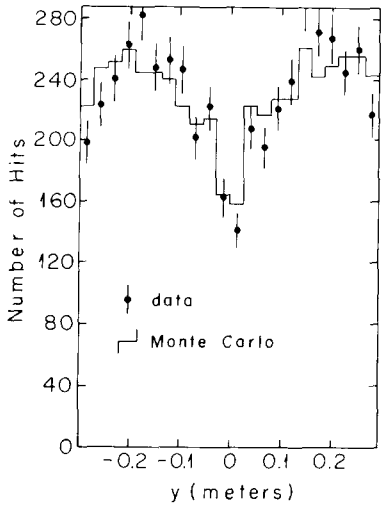

Fig. 20. Comparison of the observed hit distribution in the C scintillator bank (projected on the y axis) with that predicted by a Monte Carlo simulation including the effect of scintillator inefficiency.

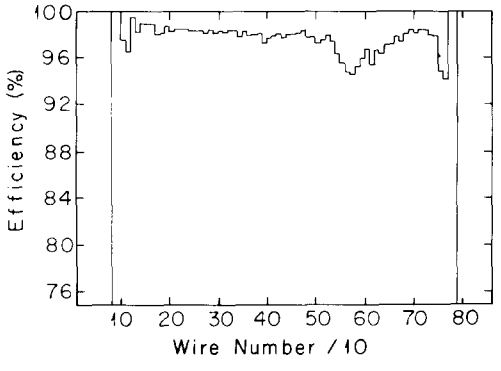

Fig. 21. The efficiency of a typical MWPC plane as a function of wire number. A dip occurred in the region of beam muons.

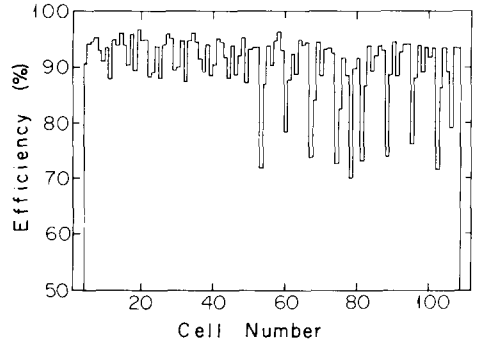

Fig. 22. The efficiency of a drift-chamber plane as a function of cell number. The pattern of inefficiency with a period of seven cells was due to a faulty time-digitizer board.

14. Acceptance

The acceptance of the detector was calculated with a Monte Carlo program. A data string completely equivalent to that written on tape by the on-line computer was formed and subsequently analyzed by the same programs that processed real data.

The calculation simulated the production of muon pairs both in the target and in the hadron absorber, the propagation of the muons through the apparatus, and the detection of tracks in scintillators and wire chambers, including detector efficiency and resolution. The momentum of the beam particles was generated according to the measured spectrum. Fermi motion in nuclei was simulated according to the model of ref. [24]. Multiple scattering and energy loss due to ionization and bremsstrahlung, including Landau–Vavilov fluctuations, were accounted for in the transport of each particle. The magnetic-field variation within the selection magnet was also simulated.

The acceptance for the 80 and 255 GeV/c runs, including effects of detector efficiency, is shown as a function of five relevant parameters in fig. 23. The acceptance depended essentially linearly on longitudinal momentum, resulting in a larger acceptance for a given value of x_F at the higher beam momentum. Apart from the enhancement of the acceptance due to the focusing effect of the selection magnet, the limitations on the acceptance were largely due to the finite transverse apertures of the detector. Because the wire chambers were not deadened in the central region, the acceptance remained high at $x_F = 1$.

When calculating differential cross sections, we used values of the detector acceptance averaged over each bin for each reconstructed parameter. Inasmuch as the cross section is a function of each parameter, this procedure would appear to require prior knowledge of the cross section. We achieved accuracy in simulating the dependence of the data on relevant physical parameters by iterating the computation of the acceptance. The first approximation was based on Monte Carlo events generated according to an extrapolation of previously measured cross sections. The computed acceptance was used to correct the real data, which were then used as a better approximation to the cross section in performing a second acceptance calculation. Only small corrections were necessary in the subsequent iterations. The consistency of the method was demonstrated by the rapidity of the convergence and the stability of the result.

15. Resolution

The Monte Carlo program was also used to estimate the measurement resolution for various parameters describing a lepton pair (see table 4).

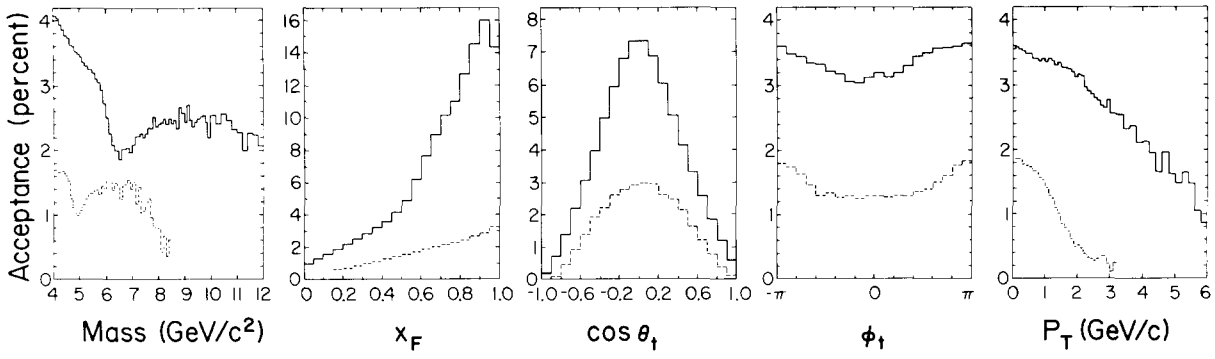

Fig. 23. The calculated acceptance in five kinematic variables for the 255 GeV/c run (solid curves) and the 80 GeV/c run (dashed curves). Effects of detector efficiency and resolution are included. The acceptance in each variable reflects an integration over the other variables, weighted by an assumed production cross section.

Table 4
R.m.s. resolutions for various physical parameters.

Parameter	σ during 80 GeV/c run	σ during 255 GeV/c run
Mass	0.14 GeV/c^2	0.20 GeV/c^2
x_F	0.04	0.03
x_π	0.03	0.03
x_N	0.03–0.05	0.01
P_T	0.13 GeV/c	0.19 GeV/c
cos θ_t	0.05–0.07	0.08

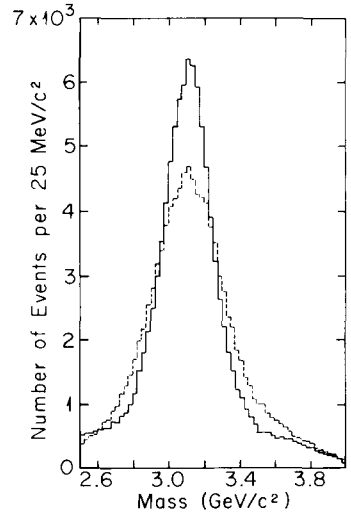

Fig. 24. The reconstructed $\mu^+\mu^-$ pair invariant-mass spectrum in the J/ψ resonance region. The solid curve is based on optimal use of position and angle measurements in the detector, while the dashed curve is based on use of angle measurements only. The resolution in mass is dominated by the effect of multiple scattering in the tungsten target.

The dominant factor in determining the resolution was multiple scattering of the muons in the target and hadron absorber, except at large values of x_π where the uncertainty in the beam momentum was important. We minimized the degradation of the resolution due to multiple scattering by reconstructing the track parameters at the target, taking into account the correlation between scattering angle and cumulative transverse displacement and assuming that the two muons in an event originated at a common point in the target. A χ^2 fitting procedure was devised which included these constraints on the observed positions and angles of the tracks. The resulting improvement in mass resolution is illustrated in fig. 24 using J/ψ events observed in the experiment. The broader curve resulted from a mass calculation using only the measured slopes and momenta of the muons, while the narrower curve was obtained using all available information according to the χ^2 method.

16. Data collection

Data were collected from December, 1983 through July, 1984 in three exposures, summarized in table 5. For each data set, the total number of pions on target was estimated to 20% by correcting the integrated number of beam particles on target (as recorded by the ion chambers) for detector live-time and hadron contamination.

A three-week test run in 1982 yielded a sample of 1322 $\mu^+\mu^-$ pairs with $M > 4.0$ GeV/c^2 produced by a π^- beam of momentum 263 GeV/c [12]. The principal improvements to the apparatus made on the basis of this test included the installation of the Level-3 trigger processor, drift chambers upstream of the analyzing magnet, and an additional plane of scintillators (the *Du* counters).

Fig. 25 shows the invariant-mass spectrum for $\mu^+\mu^-$

Table 5
Summary of the data collection.

Beam momentum [GeV/c]	Beam particle	Spill length [s]	Total number of "live"pions on target	$\mu^+\mu^-$ pairs with $M > 4.0$ GeV/c^2
80	π^-	10–15	7.2×10^{13}	4060
255	π^-	15–20	5.5×10^{13}	~ 50000
255	π^+	20	3.8×10^{13}	~ 9000

pairs from the 255 GeV/c π^- run (uncorrected for acceptance). In addition to the continuum of Drell–Yan events above the J/ψ resonances at 3.1 GeV/c^2, evidence for the upsilon family of resonances may be seen in the mass range 9–10.5 GeV/c^2.

First results from the analysis of these data are being published elsewhere [25]. The results confirm and extend our previous findings [3,5] that there are striking features to muon-pair production by pions whose momentum is carried largely by a single quark.

Acknowledgments

We wish to acknowledge the vital contributions to the success of the experiment from the support staffs at our universities and at Fermilab.

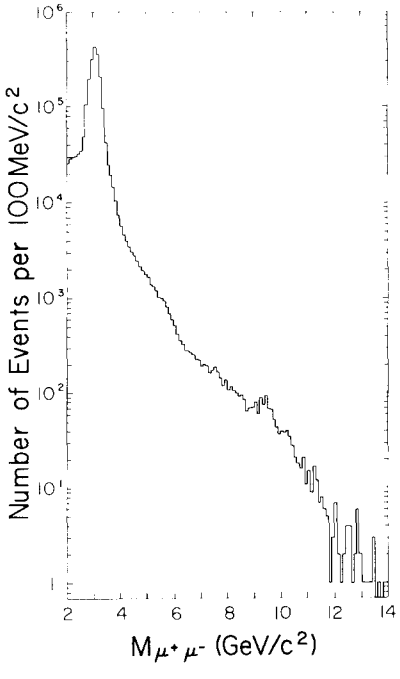

Fig. 25. The invariant-mass spectrum of $\mu^+\mu^-$ pairs from the 255 GeV/c π^- data set, uncorrected for acceptance.

The first- and second-level trigger electronics, together with the phototube- and pad-latch modules, was designed by Carl Bopp at Princeton University, with support from Ann David and Marius Isaila. The expert construction of the drift chambers and scintillation-counter hodoscopes C through E was under the supervision of Howard Edwards.

The AB veto logic and MWPC readout electronics were designed by Harold Sanders and constructed by the Electronics Development Group at The University of Chicago. The veto-hodoscope arrays and the proportional chambers were designed and constructed with the assistance of the Engineering Services Group under Richard Armstrong.

We gratefully acknowledge technical support in the construction of the F-scintillator hodoscope from Norbert Cent and George Holland at Iowa State University.

The Accelerator Division of Fermilab deserves special commendation for the excellent operation of the Tevatron immediately after it was commissioned. The selection magnet was designed at Fermilab by Ed LaVallie who also supervised its assembly. The coils were constructed in the magnet shop headed by Jack Jagger. Strong support was provided by the electrical, mechanical, and beam-line groups of the Experimental Areas Division headed by Roger Dixon with assistance from Joel Butler in the Research Division. PREP furnished the on-line computer and the standard fast-electronics modules used in the experiment. We enjoyed excellent support from the on-line and off-line computer groups under the direction of Jeff Appel and Al Brenner. The ECL/CAMAC trigger modules were developed by a group under Tom Nash and supported by Merle Haldeman. The Ziptrack facility was made available by the Research Division Facilities Department and was supervised by Willie Yang and Al Ito. Jon Hawkins was our liason physicist.

This experiment was supported in part by the US Department of Energy under contracts DE-AC02-76ER03072 and W-7405-ENG-82-KA-01–01 and by the National Science Foundation.

References

[1] S.D. Drell and T.-M. Yan, Phys. Rev. Lett. 25 (1970) 316, 902(E); S.D. Drell and T.-M. Yan, Ann. Phys. 66 (1971) 578.
[2] D. Antreasyan et al., Phys. Rev. Lett. 39 (1977) 906; D.M. Kaplan et al., Phys. Rev. Lett. 40 (1978) 435.
[3] G.E. Hogan et al., Phys. Rev. Lett. 42 (1979) 948; C.B. Newman et al., Phys. Rev. Lett. 42 (1979) 951;
[4] E.L. Berger and S.J. Brodsky, Phys. Rev. Lett. 42 (1979) 940;
E.L. Berger, Z. Phys. C4 (1980) 289;
S. Matsuda, Phys. Lett. 119B (1982) 207.
[5] K.J. Anderson et al., Phys. Rev. Lett. 43 (1979) 1219.

[6] K.J. Anderson et al., Phys. Rev. Lett. 36 (1976) 237.
[7] G.E. Hogan, Nucl. Instr. and Meth. 165 (1979) 7.
[8] H.J. Frisch et al., IEEE Trans. Nucl. Sci. NS-27 (1980) 150;
L. Anderson et al., Nucl. Instr. and Meth. 223 (1984) 26.
[9] J. Badier et al., Nucl. Instr. and Meth. 175 (1980) 319.
[10] A brief description of the trigger processor appears in: J.F. Greenhalgh, Proc. Symp. on Recent Developments in Computing, Processor and Software Research for High-Energy Physics, Guanajuato, Mexico, eds., R. Donaldson and M.N. Kreisler (1984) pp. 63–68.
[11] S. Palestini, Ph.D. Dissertation, Princeton U. (1984) unpublished.
[12] J.P. Alexander, Ph.D. Dissertation, U. of Chicago (1985) unpublished;
C.E. Adolphsen, Ph.D. dissertation, U. of Chicago (1985) unpublished. These theses present results from a run at 263 GeV/c prior to the commissioning of the Tevatron. The 255 GeV/c data are presently being analyzed.
[13] B. Cox, The High Intensity Laboratory, in Fermilab Report (January, 1979) unpublished.
[14] F. Hornstra, Fermilab TM-380 (July, 1972) unpublished.
[15] H.W. Atherton et al., CERN/SPS Yellow Report 80–07 (1980) unpublished.
[16] J. Badier et al., Phys. Lett. 104B (1981) 335;
H.J. Frisch et al., Phys. Rev. D25 (1982) 2000.
[17] A.S. Carroll et al., Phys. Lett. 80B (1979) 319.
[18] R. Yamada et al., Nucl. Instr. and Meth. 138 (1976) 567;
A. Ito et al., Fermilab TM-1200 (July, 1983) unpublished.
[19] G. Gollin, M.V. Isaila, F.C. Shoemaker and P. Surko, IEEE Trans Nucl. Sci. NS-26 (1979) 59.
[20] D.T. Simpson, M.S. thesis, Iowa State U. (1984) unpublished.
[21] H. Areti et al., Nucl. Instr. and Meth. 212 (1983) 135.
[22] E. Barsotti et al., IEEE Trans. Nucl. Sci. NS-26 (1979) 686;
E. Barsotti et al., Fermilab TM-821, 2nd ed. (July, 1982) unpublished.
[23] J.F. Barlett et al., IEEE Trans. Nucl. Sci. NS-26 (1979) 4427;
L.M. Taff et al., Fermilab PN-97.5 (October, 1979) unpublished.
[24] A. Bodek and J.L. Ritchie, Phys. Rev. D23 (1981) 1070.
[25] S. Palestini et al., Phys. Rev. Lett. 55 (1985) 2649;
J.P. Alexander et al., to be submitted to Phys. Rev. D.