SLAC-155
UC-34
(TH) and (EXP)

CHARGE ASYMMETRY IN $K_L^0 \rightarrow \pi\mu\nu$ *

ROBERT L. PICCIONI

STANFORD LINEAR ACCELERATOR CENTER
STANFORD UNIVERSITY
Stanford, California 94305

PREPARED FOR THE U. S. ATOMIC ENERGY
COMMISSION UNDER CONTRACT NO. AT(04-3)-515

September 1972

Printed in the United States of America. Available from National Technical Information Service, U. S. Department of Commerce, 5285 Port Royal Road, Springfield, Virginia 22151.
Price: Printed Copy $3.00; Microfiche $0.95.

---

*Ph. D. dissertation.

ABSTRACT

For $\Gamma_{\pm} \equiv$ Rate $(K_L^o \rightarrow \pi^{\mp} \mu^{\pm} \nu)$, we have measured the charge asymmetry, $\delta \equiv (\Gamma_+ - \Gamma_-)/(\Gamma_+ + \Gamma_-)$, to be $(+ 0.28 \pm 0.05)\%$. Such an asymmetry is evidence of CP violation. We have compared our result with previous CP violation experiments in this and other decay modes of the neutral kaon. We find consistency among the experiments, and consistency with the predictions of the superweak theory.

ACKNOWLEDGEMENTS

This experiment was accomplished through the tireless efforts of many individuals.

The initiative and direction of Melvin Schwartz were ultimately responsible for the successful construction of the experimental facility. His imaginative hypotheses stimulated the search for new phenomena and potential biases. For this, and for his many acts of generosity, I offer him my warmest thanks.

The intelligence, experience, and dedication of Stanley Wojcicki were the heart and soul of our experimental effort. I thank him for his patience and his crucial assistance.

I wish to thank John Liu for his many contributions to the experiment, including constructing the scintillation counters, writing the Monte Carlo programs, and writing a major part of the analysis programs. Throughout the endless series of revisions, new formats, improved designs, and "why-don't-we-try-this?"'s, he remained as meticulous and patient as ever.

I wish to thank David Hitlin, David Fryberger, and David Dorfan for their work on the wire spark chambers and computer peripherals, and for their many important suggestions in regard to the analysis procedures.

I wish to thank Dan Porat for the tremendous job he did in creating the wire chamber readout and data acquisition logic. The great volume of data, the speed of processing, and the superb reliability of this

system are a tribute to his talents. I wish to thank Len Birkwood, Dale Ouimette, Karl Hense and Brian Brauer for their fine efforts in constructing and maintaining the readout and acquisition logic, and in maintaining the on-line computer system. I especially thank these gentlemen for being available around the clock during the experimental running.

I wish to thank Roger Coombes for his fine work in the mechanical design of the wire spark chambers, the spark chamber gas system, and the supports of the scintillation counters and the 60 ton lead wall. I thank Don Clark, Cliff Rasmussen and Felix Vargas for their contributions to these systems.

I gratefully acknowledge the aid of Uriel Nauenberg, Alan Franklin, Robert Morse, and Robert Messner of the University of Colorado for their assistance in the experimental running, and for providing the T counter array. I express my special thanks to Robert Messner for his many important contributions during the setup phase of the experiment.

I wish to thank Greg Donaldson, Allan Rothenberg and David Uggla for their work during the setup and running, and for their assistance and advice during the analysis of the data.

I thank Basil Meyer for his work on the on-line computer system.

I commend Ken Johnson and his colleagues in the Experimental Facilities Department for their excellent technical support.

I thank Joe Cobb and the staff of the Magnetic Measurements group for providing us with the field maps of our analyzing magnet.

I wish to thank Vern Price and the staff of the Accelerator

Operations group for their vigilance and efforts in maintaining the quality of our beam.

I gratefully acknowledge the staff of the SLAC Facility Stanford Computation Center for diligently running thousands of jobs on several hundred tapes and never losing one. My special thanks to Cheryl McGriff and Ralph Bowser for their friendly assistance.

My sincere thanks to Margaret Ledford for typing, and often re-typing, this thesis.

I wish to thank the members of the Technical Reports group for the excellent job they have done in preparing the figures for this thesis.

I thank my wife Joan for her encouragement and for her assistance in writing and proofreading this thesis.

I thank my father Oreste Piccioni for showing me the beauty of science, and for teaching me everything from 'rithmetic to regeneration. I always warm with pride when asked, "Are you the son of ...".

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

Introduction

Symmetry is one of nature's most prominent features. Symmetry reduces complexity and nature prefers simplicity. Throughout the range of observable phenomena, from galaxies to elementary particles, one sees the simplest, most symmetric shapes, configurations and motions. It is remarkable when one finds a phenomenon which violates a basic symmetry principle.

We present here a study of the symmetry between matter and antimatter. This symmetry, known as CP symmetry, is equivalent to the combination of two other symmetries C and P. C refers to charge conjugation, the symmetry between positive and negative charge. P refers to parity, the symmetry between objects and their mirror images. Prior to 1956, nature was thought to rigorously obey the symmetries P, C, and CP. At that time, violations of P and C were discovered in the weak interactions. These violations were due fundamentally to the existence of neutrinos which always have their spins aligned in a left-handed manner with respect to their velocities, and antineutrinos which always have their spins aligned in a right-handed manner. The nonexistence of right-handed neutrinos and left-handed antineutrinos violates P, and C but not the combined symmetry CP. Thus none of the observed violations of P and C were violations of CP, and the notion of CP invariance survived until 1964 and the experiment of Christensen et al.[1] which observed the decay $K_L^o \rightarrow \pi^+\pi^-$.

By now, CP violation has been seen in the two pion and the leptonic decays of the K_L^o. The relative degree of CP violation seen in the various modes has been the subject of considerable speculation, excitement, and investigation. We present here a measurement of the charge asymmetry in the decay $K_L^o \rightarrow \pi\mu\nu$. We examine the asymmetry across the Dalitz plot and compare our overall value to the charge asymmetry in $K_L^o \rightarrow \pi e \nu$ and to the CP violation seen in $K_L^o \rightarrow \pi\pi$.

We feel that our experiment has several important advantages. The use of wire spark chambers enables us to accumulate large numbers of events while obtaining precise measurements of the pion and muon trajectories. Our measurement of the K_L^o momentum, by the time-of-flight technique, adds a powerful tool to our analysis. Our multielement muon filter provides detailed information on each penetrating particle. The multitude and precision of our information enable us to detect, identify and measure contaminations which might otherwise have to be hypothesized and estimated. As the charge asymmetry is such a small effect, even minor systematic errors can seriously jeopardize any experimental measurement. We feel that the detailed knowledge of each event in our data is invaluable.

CHAPTER II

Theory

Our understanding of the neutral kaon system is based on a distinction between the eigenstates into which kaons are produced and the eigenstates from which kaons decay. Hadronic processes produce K^o's and $\bar{K}^o$'s. These states are the antiparticles of each other and differ only in their strangeness quantum number; the K^o has strangeness +1 and the $\bar{K}^o$ has strangeness -1. Second order weak transitions are therefore possible between the two states, and due to this mixing, the decay eigenstates are not the K^o and $\bar{K}^o$ but linear combinations thereof[2]. The decay states of definite lifetime are conventionally defined to be

$$|K_s^o\rangle = \left\{ (1+\epsilon)|K^o\rangle + (1-\epsilon)\ |\bar{K}^o\rangle \right\} / \sqrt{2(1+|\epsilon|^2)}$$

$$|K_L^o\rangle = \left\{ (1+\epsilon)|K^o\rangle - (1-\epsilon)\ |\bar{K}^o\rangle \right\} / \sqrt{2(1+|\epsilon|^2)}$$

where the K_s^o is the short-lived, and the K_L^o the long-lived kaon, and ϵ is the complex constant which measures the level of CP mixing in the decay states.

There are two possible sources of CP violation. The first alternative is that the weak interactions might contain small CP violating terms. The prospect of a force that almost conserves CP symmetry is unesthetic, however. The other alternative is that kaon states of definite CP are somehow mixed by an amount ϵ. The weak

interactions could then conserve CP, and the apparent violations[3] would have to be blamed on the process by which the kaon states were mixed.

The most widely held theory of CP violation is the superweak theory of Wolfenstein[4], which states that CP violation is due entirely to CP mixing in the decay states and is in no way intrinsic to the weak interaction itself. The theory further states that this mixing is caused by a hitherto unknown force weaker than the normal weak interaction. The superweak theory provides no further understanding of this force, but it does make definite predictions. As the violation alters the kaon states but not the decay process, the effect should be independent of the kaon's decay mode and kinematics. Recent results on the $K_L^o \rightarrow \pi\pi$ modes[5] are in excellent agreement with the superweak predictions.

We now proceed to compute an expression for the charge asymmetry assuming that CP violation is restricted to the mixing given by ϵ, as in the superweak theory. We define the K^o and $\bar{K}^o$ decay amplitudes:

$$f \equiv \langle \pi^- \mu^+ \nu \, | H_w | K^o \rangle \quad (\Delta S = + \Delta Q)$$

$$g \equiv \langle \pi^- \mu^+ \nu \, | H_w | \bar{K}^o \rangle \quad (\Delta S = - \Delta Q)$$

$$x \equiv g/f$$

by CPT invariance,

$$f^* = \langle \pi^+ \mu^- \bar{\nu} \, | H_w | \bar{K}^o \rangle$$

$$g^* = \langle \pi^+ \mu^- \bar{\nu} \, | H_w | K^o \rangle$$

where f, g and x can be functions of the Dalitz plot position. Combining these amplitudes appropriately we find the K_L^o decay amplitudes to be,

$$\langle \pi^- \mu^+ \nu \, | H_w | K_L^o \rangle \sim (1 + \epsilon) f - (1 - \epsilon) g$$

$$\langle \pi^+ \mu^- \nu | H_w | K_L^o \rangle \sim (1 + \epsilon) g^* - (1 - \epsilon) f^*$$

and for the decay rates,

$$\Gamma_{\pm} \equiv \text{Rate} \; (K_L^o \to \pi^{\mp} \mu^{\pm} \nu)$$

$$\Gamma_+ \sim (1 + |\epsilon|^2)(1 + |x|^2) + 2 \, \text{Re} \, \epsilon \, (1 - |x|^2) - 2 (1 - |\epsilon|^2) \, \text{Re} \, x + 4 \, \text{Im} \, \epsilon \, \text{Im} \, x$$

$$\Gamma_- \sim (1 + |\epsilon|^2)(1 + |x|^2) - 2 \, \text{Re} \, \epsilon \, (1 - |x|^2) - 2 (1 - |\epsilon|^2) \, \text{Re} \, x + 4 \, \text{Im} \, \epsilon \, \text{Im} \, x$$

to first order in ϵ,

$$\Gamma_+ - \Gamma_- \sim 4 \, \text{Re} \, \epsilon \, (1 - |x|^2)$$

$$\Gamma_+ + \Gamma_- \sim 2 \, (1 + |x|^2 - 2 \, \text{Re} \, x)$$

and finally for the charge asymmetry,

$$\delta \equiv \frac{\Gamma_+ - \Gamma_-}{\Gamma_+ + \Gamma_-}$$

$$\delta = 2 \, \mathrm{Re} \, \epsilon \; \frac{1 - |x|^2}{|1 - x|^2}$$

We emphasize the x, and hence δ, is in general a function of the Dalitz variables and could be different for $K_L^o \rightarrow \pi\mu\nu$ ($K_{\mu 3}^o$) and $K_L^o \rightarrow \pi e \nu$ (K_{e3}^o). x has been measured for both decays and is known to be small. The current world averages are[6]:

For K_{e3}^o $\quad$ $\mathrm{Re}\, x = 0.021 \pm 0.022$

$\mathrm{Im}\, x = 0.003 \pm 0.021$

For $K_{\mu 3}^o$ $\quad$ $\mathrm{Re}\, x = 0.09 \, {}^{+0.07}_{-0.09}$

$\mathrm{Im}\, x = 0.00 \, {}^{+0.11}_{-0.10}$

We expect, therefore, the effect of x should be small.

CHAPTER III

Experimental Setup

This study was performed at the Stanford Linear Accelerator Center on an experimental facility built by this group. Using this facility, we studied K_L^o decays from a neutral beam of well-defined direction and time structure. Charged decay products emerging from the beam were detected by arrays of scintillation counter banks and wire spark chambers. In order to measure the particle momenta, detectors were placed on both sides of a wide-aperature magnet. The decay products finally impinged on a lead wall which was followed by additional scintillation counter banks. Refer to Fig. 1.

The scintillation counters served to identify events of interest and to measure the K_L^o time-of-flight (TOF). The wire spark chambers provided an accurate determination of the particle trajectories. Particles which penetrated the lead wall were identified as muons.

Our data acquisition logic gathered and condensed the information from the detectors and transferred this data to our on-line computer. The computer monitored the data, carried out equipment performance tests, and recorded the events on magnetic tape for later analysis.

We now consider each aspect of the setup in detail.

A. Beam

Our K_L^o beam was produced by a primary electron beam incident on a beryllium target. The energy of the electron beam was generally 19 GeV.

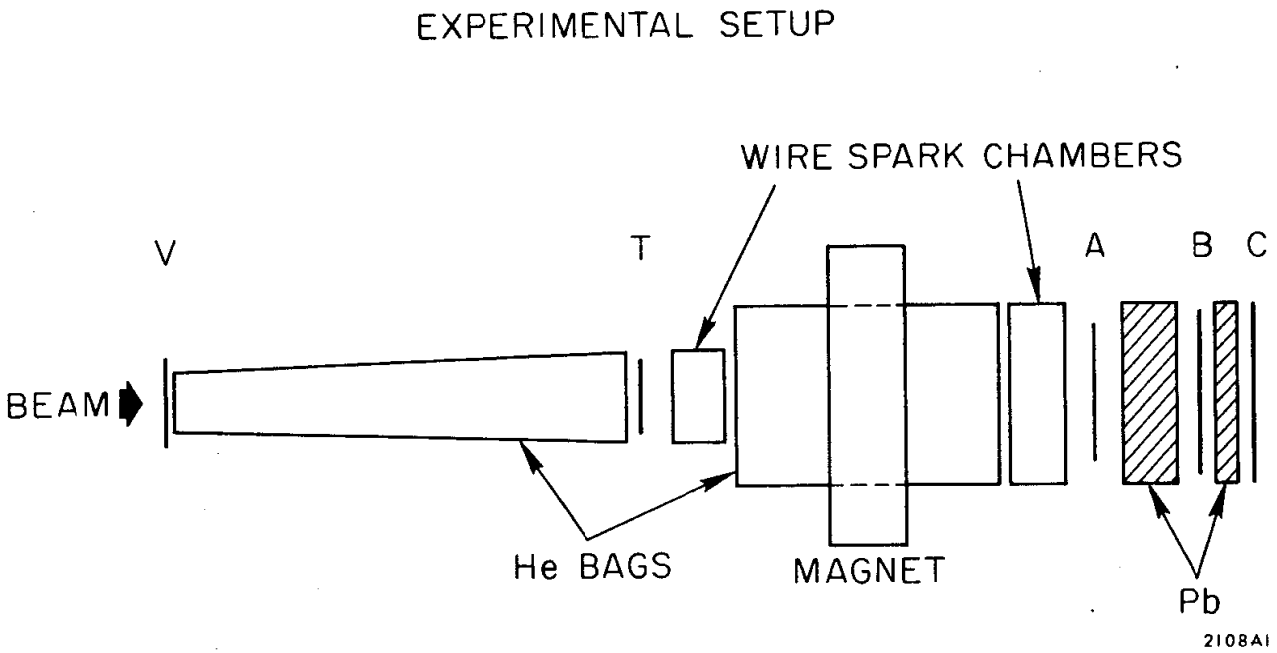

Fig. 1--Experimental setup.

The time structure of this beam was unusual. The beam came in pulses of 1600 nsec duration, typically 160 per second. The average current over one entire pulse was 3 ma. Each pulse was further divided into buckets of < 20 psec duration, with 12.5 nsec between buckets. Since the bucket separation was greater than the TOF difference between the fastest and the slowest K_L^o of interest, we could unambiguously and accurately determine the production time. By measuring the arrival time of its decay products, we determined the K_L^o TOF.

The electron beam struck a 1 radiation length beryllium target. Behind this target a coaxial cable was mounted in the beam line. As the electron beam passed through the target and hit the cable, it produced a signal which provided the time reference point for our TOF measurements. Our secondary K_L^o beam was taken at an angle of 50 mrad to the primary electron beam. The K_L^o beam passed through a hole in a 6 m thick iron and concrete shielding wall. The openings of this hole were blocked with a total of 20 cm of lead and 168 cm of polyethylene. The lead removed the gamma rays from the beam, and the polyethylene reduced the number of neutrons. At the exit hole of the shielding wall, a collimator reduced the beam to 5 by 12 cm. The beam then traveled through air for another 60 m before reaching the start of our detectors. Along this path, additional shielding was placed around the beam, and three sweeping magnets removed charged particles. A second collimator was placed 29 m upbeam of our momentum-analyzing magnet. This collimator removed scattered particles,

but did not intersect the beam itself. The total distance between the electron target and our analyzing magnet was 77 m. Care was taken in the design of this beam line so that no substantial material intersected the beam volume defined by the target and the shielding wall collimator. The beam's cross section at our detector was 64 cm wide by 28 cm tall.

The primary components of this beam were K_L^o's and neutrons. In the multi-GeV region, the K_L^o and neutron fluxes were comparable. The K_L^o momentum spectrum of our beam is shown in Fig. 2.

When the three sweeping magnets in the beam line were turned off, a substantial beam of fast muons entered our apparatus, providing a ready calibration source for our TOF measurements.

B. Decay Volume

Events of interest were restricted to K_L^o decays that occurred in our decay volume, a 6 m long helium bag preceded and followed by scintillation counters. Immediately downbeam of the final sweeping magnet was the V counter. Following the V counter was the upbeam side of the helium bag. This region was surrounded by iron and concrete shielding to prevent stray particles from entering the apparatus. Just downbeam of the helium bag was the T counter bank. Counts in the T bank coupled with an absence of counts in the V counter signaled decays (or possibly beam interactions) in the decay volume.

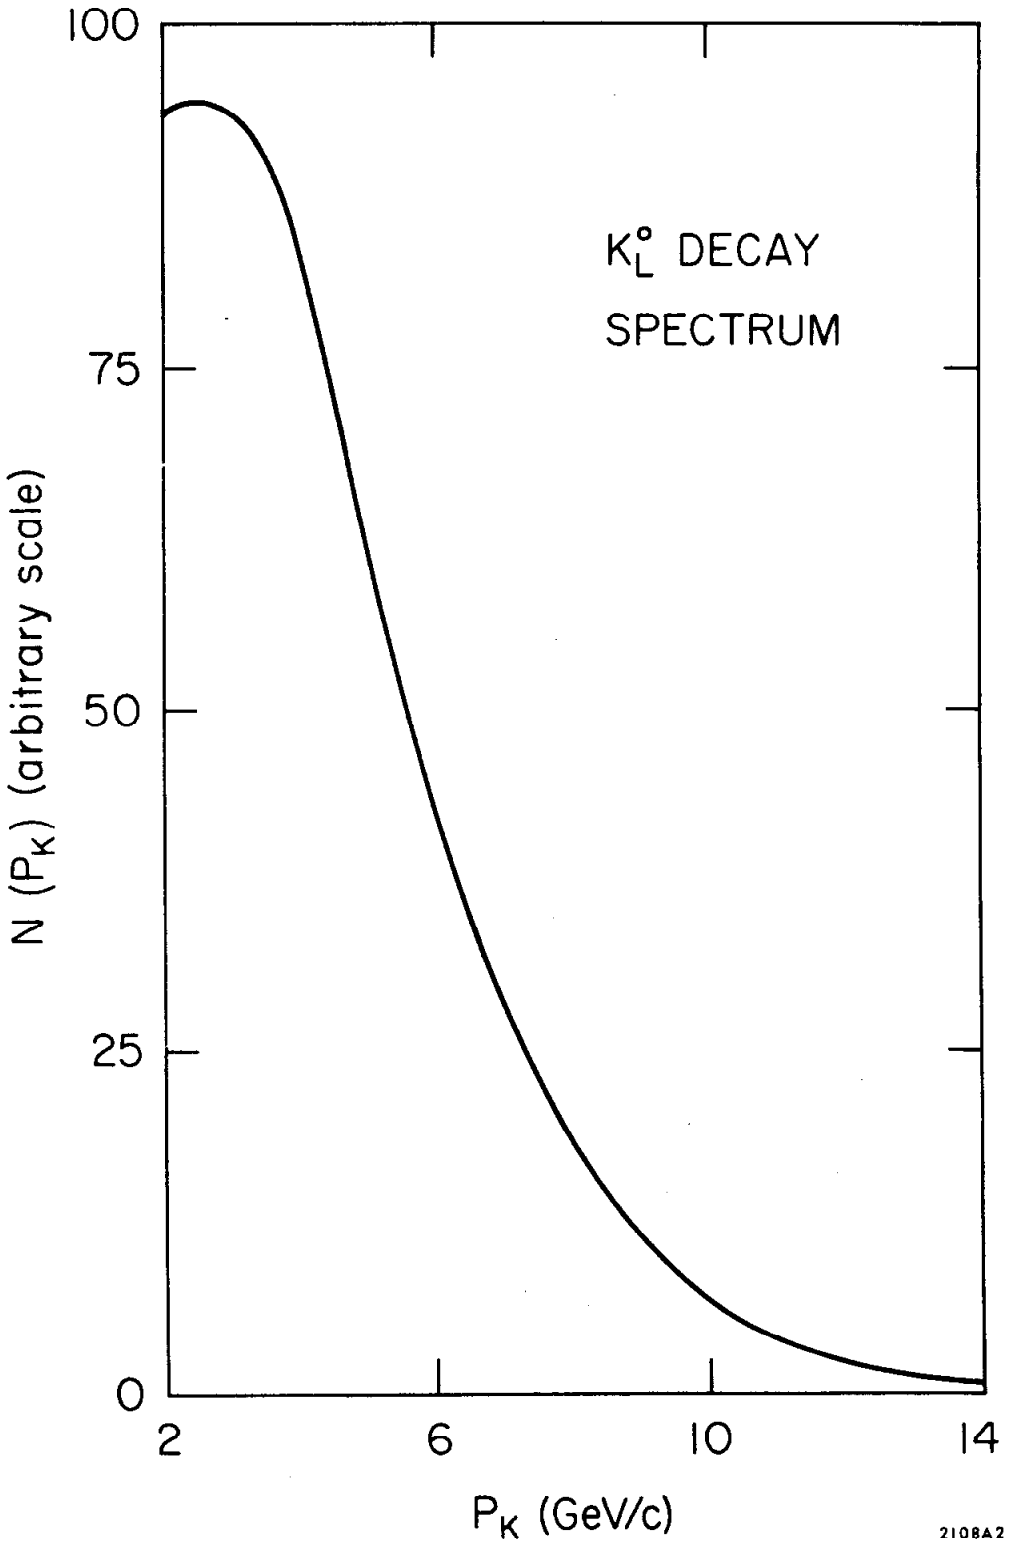

Fig. 2--K_L^o decay spectrum.

C. Scintillation Counters

Our five scintillation counter banks were of conventional design and construction. TOF measurements were made from the A, B and C counter banks. These counters were therefore viewed by one photomultiplier tube at each end. 56AVP tubes were used throughout except on the A counters which were equipped with XP 1021's for superior time resolution.

The V counter was located just upbeam of the decay volume. It was only 0.6 cm thick in order to minimize material in the beam line, and was viewed by four photomultiplier tubes in order to guarantee high efficiency. The presence of a signal in the V counter indicated a charged particle entering our decay volume.

The T counter bank was located just downbeam of the decay volume. This bank covered an area 100 cm wide by 120 cm tall, and consisted of 40 counters. To minimize material in the path of the decay products, the T counters were only 0.3 cm thick. Signals from the T bank indicated charged particles emerging from the decay volume.

Behind the downbeam wire chambers was a bank of 12 A counters covering an area 180 cm wide by 120 cm tall. Signals from the A bank indicated charged particles passing through the magnet.

Downstream of the A counters was a 75 cm thick lead wall. Behind this wall was a bank of 14 B counters covering an area 210 cm wide by 165 cm tall. Signals from the B bank indicated the presence of a penetrating particle.

Downstream of the B counters was a 30 cm thick lead wall.

Behind this wall was a bank of 16 C counters covering an area 240 cm wide by 165 cm tall. Signals from the C bank confirmed the presence of a penetrating particle.

The T and A counters were located 240 cm from the magnetic mirrors of our momentum-analyzing magnet. The magnetic field in this region was measured to be no higher than 15 gauss. Each photomultiplier tube was protected from this field by two concentric shields. The inner shield was molypermalloy; the outer shield was of iron. The B and C counters were 380 and 450 cm from the magnetic mirrors, respectively. As the field in this region was 2-3 gauss, the B and C shields were merely molypermalloy. The pulse heights in the A, B and C counters were digitized and recorded with each event. Some variations were found in the photomultiplier pulse heights between the magnet forward and magnet reverse data. These variations were typically a few parts per thousand, and resulted in an exceedingly small correction to the measured asymmetry. See VI.H for details.

The photomultiplier tube high voltages were set so that the pulse height from a minimum ionizing track was between 7 and 10 times the electronic trigger logic's discrimination threshold. These large pulse heights improved the time resolution, reduced pulse height time slewing, and minimized the importance of pulse height variations from magnetic field effects.

The computer checked the photomultiplier tube voltages every few hours. Voltages which had drifted were reported to the experimenter who could instruct the computer to adjust the voltages as desired.

Finally, every counter had a photodiode imbedded in it. These photodiodes could be activated under computer control, thus testing photomultiplier tubes and related electronics. This procedure, called a Pulser run, was performed every few hours.

D. Wire Spark Chambers

Our wire spark chambers provided accurate determinations of the particle trajectories. The details of the wire chamber construction, operation, and readout have been previously reported[7]. We will touch here only on some of the major points.

Two similar wire chamber arrays were constructed. One was placed upbeam of the magnet and one downbeam. Each array had 10 gaps, with each gap being composed of two parallel planes of wires. The planes were separated by 10 mm and the wires within each plane were separated by 1 mm. The wire diameter was 0.1 mm. One of the two planes was pulsed with high voltage, while the other, the ground plane, was used to read out the spark locations. Our wire chambers had several advantages. The wire planes put very little mass into the path of the beam and its decay products. By isolating the wires of the high voltage plane with resistors at the high voltage distribution bus, a spark on one wire could not rob charge from nearby wires. This isolation led to higher multitrack efficiency. The parallel wire scheme eliminated pulse propagation problems inherent in crossed wire chambers, thus providing high and uniform efficiency.

All the wire planes were perpendicular to the beam line (the z axis), and each gap measured spark locations with respect to one

coordinate direction. Within each array of 10 gaps there were 4 x-measuring gaps, 4 y-measuring gaps, 1 u-measuring gap, and 1 v-measuring gap. The x axis was horizontal and the y axis vertical. The u and v axes were rotated with respect to x and y. The upbeam chambers were 1.2 m square; the downbeam chambers were 1.2 m by 2.4 m. In each case, the u and v planes were somewhat larger. There were a total of 30,000 readout wires in the system.

The use of 10 gaps to determine 3-dimensional line segments provided a high degree of redundancy. This redundancy enabled us to isolate tracks of interest even in the presence of considerable background and occasional chamber inefficiency.

The chambers were typically pulsed with 4 KV. The voltage was delivered by discharging numerous charging cables through hydrogen thyratrons. The pulse length was 200 nsec, with a rise time of 40-50 nsec. A DC clearing field removed residual ionization from previous firings. It is interesting to note here, that without the clearing field, and at high trigger rates, the chambers would light up with many hundreds of sparks. The residual ionization from a preceding spark would reignite and form a cluster of sparks. These clusters would dance frantically over the surface of the chambers. Our computer was not amused by this mode of operation, but it does indicate the extremely high multitrack efficiency of our chambers.

The chambers were filled with a gas mixture of 90% neon, 10% helium, and a small amount of alcohol. The alcohol reduced spurious sparking and edge breakdown, and improved the multitrack efficiency.

The readout scheme was of unique design. It was unaffected by large magnetic fields, and was capable of handling an unlimited number of sparks. Each readout wire was connected to a capacitor. The current from a spark traveled down the readout wire and charged the capacitor. Later, all the capacitors were interrogated by the data acquisition logic.

Periodically, the computer would test all the wires and their electronic readouts. Each wire was required to have a nonzero counting rate. Each readout element was required to read ON when its capacitor was artificially charged, and read OFF after its capacitor had been cleared. Failures were reported to the experimenter. Due to the modular nature of the readout design, replacement of faulty electronic units was fast and straight-forward.

E. Momentum-analyzing Magnet

Our momentum-analyzing magnet had a gap of 2.5 m by 1 m by 1 m. It had magnetic mirrors to reduce fringe fields. At our normal running condition the magnet imparted a momentum transfer of 377 MeV/c to incident particles. The magnetic field was measured at 30,000 grid points. From these measurements, field integrals were computed and used in the analysis for momentum determinations. An NMR probe was placed just above the lower pole piece when the field was measured. The reading on this NMR was recorded throughout the experiment, and the measured field integrals were scaled appropriately. Over the course of the experiment, the NMR reading remained constant to 0.1%.

A shunt on the magnet power lines monitored the current and polarity of the magnet. We occasionally checked the magnet polarity by observing the deflection of a current-carrying wire in the field.

The volume inside the magnet gap was filled by a helium bag in order to reduce interactions and multiple scattering.

F. Lead Wall

Muons from $K^{o}_{\mu 3}$ decay were identified by their ability to penetrate our lead wall. This wall was divided into two sections. The first section was 75 cm thick and was followed by the B counters. The second section was 30 cm thick and was followed by the C counters. The minimum muon energy necessary to penetrate this assembly was 1.54 GeV. The lead wall was 7.7 interaction lengths thick.

The combination of the B and C information gave us powerful constraints on the presumed muon. The measured times, the horizontal positions, and the vertical positions in the B and C banks were required to concur with the A counter data and the measured trajectory in the downbeam chambers. Through these constraints spurious muons were eliminated.

After the lead wall was inserted, we noted a substantial background from beam interactions in the lead wall. In order to reduce a spray of slow neutrons from these interactions, we placed parafin on the upbeam and downbeam faces of the lead wall.

G. Electronic Trigger Logic

The function of the electronic trigger logic was to examine the

signals from the various scintillation counters and select events of interest.

Photomultiplier tube outputs were run through discriminators which converted them into standard logic level pulses. For each A, B, and C counter, the logic pulses from the upper and lower photomultipliers were combined to produce one logic pulse for the entire counter, whose time was the average of the upper and lower times.

The logic pulses were summed to determine the number of simultaneous signals in each bank. If the required number of counts in each bank was found in the proper time coincidence with the beam's cable pulse, an event was triggered, the data acquisition logic was initiated, and the trigger logic was gated off.

The requirements of the trigger logic could be changed under computer control. For $K^{o}_{\mu 3}$ runs the trigger was noV 2T A B C Cable; for muon runs it was V T A B C Cable. By the preceding notation, we mean the time coincidence of the logic signals defined below.

A - simultaneous firing of both ends of one or more A counters

2A - simultaneous firing of both ends of two or more A counters

B - simultaneous firing of both ends of one or more B counters

C - simultaneous firing of both ends of one or more C counters

T - firing of one or more T counters

2T - firing of two or more T counters

V - simultaneous firing of at least one tube on each side of the V counter

noV - absence of a logic V

Cable - occurrence of the beam cable pulse.

Every time an event occurred, the trigger logic produced logic gates which enabled latches and analog-to-digital converters (ADCs). The latches recorded every counter which fired during the event time. Pulse height ADCs integrated the photomultiplier tube outputs. The gate to the time ADCs was produced at a time determined by the beam's cable pulse. During the gate, each ADC integrated one of the photomultiplier tube discriminator outputs. Since each discriminator output was of constant amplitude, the integration measured the time overlap of the gate and the discriminator output. The gate was designed to begin before the start of any legitimate discriminator output, and to terminate before the end of any such output. Hence, the ADC measured the time of the photomultiplier signal relative to the beam's cable pulse.

H. Data Acquisition Logic

The data acquisition logic collected and condensed the data from the wire spark chambers and the scintillation counters, and fed this information into our computer. The internal workings of this logic have already been published[7].

After the trigger logic had signaled the occurrence of an event, the data acquisition system proceeded to interrogate all the wire chamber capacitors. Each spark would appear as one or more consecutive, charged capacitors. The sparks were isolated and one word of data, specifying the location and width, was produced for each spark. These data words were sent to the computer, and the readout capacitors were cleared. As a precaution against possible leakage, the capacitors

were also cleared 1 msec prior to each beam pulse.

After the wire chamber data had been processed, the logic began examining all the latches and ADCs. All nonzero data was sent to the computer.

The data transmitted to the computer entered its memory through a direct memory access channel, thereby reducing the load on the central processor and enhancing the data transfer rate. When the last data transfer was completed, the computer was interrupted.

Another feature of the data acquisition logic was a computer output device. This device enabled the computer to measure and control up to 200 voltage levels, and gave the computer access to a large number of relays, flip-flops, and pulse generators. These controls were used to change event triggers, pulse the photodiodes in various combinations, modify scope display modes, turn alarms on and off, enable and disable the trigger logic, and select data acquisition modes.

I. On-line Computer

Our on-line computer was a PDP-9 with a core size of 24K. Its main function was to accept, process, and record events. In addition, it performed tests on the data and the equipment to ensure proper operation of the experiment.

Upon receiving an interrupt from the data acquisition logic, the computer transferred the event from the direct memory access buffer to one of three processing buffers, and enabled the trigger logic. The use of multiple processing buffers enabled us to take events in

rapid bursts with a minimum of dead time. Once in the processing buffer, the event data was translated into a more convenient format. Then the numbers and positions of sparks, the spark widths, latches, times, pulse heights, and other data were histogrammed. The event was transferred to one of two magnetic tape buffers. One of these buffers was being filled while the other was being written onto the magnetic tape.

The computer displayed these histrograms on a storage scope. The experimenter could request any particular histogram, or he could request the computer to cycle through a sequence of displays. The latter mode was very useful. In less than one minute, the experimenter could see the entire gamut of distributions and readily spot malfunctions. The computer would also display individual events, showing all the sparks and scintillation counters which fired.

The computer saved histograms from old runs. By examining the experiment's history, we could look for such long term variations as the effects of increasing the alcohol content of the spark chamber gas. Periodic tests were made on the chamber readout electronics and the photomultiplier tube high voltages. After every Pulser run, the computer checked every latch and ADC to ensure the proper operation of the scintillation counters and their associated trigger and data acquisition logic.

CHAPTER IV

Data Taking Procedures

A. $K^{o}_{\mu 3}$ Runs

We divided our $K^{o}_{\mu 3}$ data taking into runs of two hours. Every four hours, we reversed the polarity of our analyzing magnet. In addition to the testing features provided on-line, data was periodically run through an off-line analysis program which provided further checks on the apparatus, such as the wire chamber plane efficiencies.

During the normal running, we took 10 events per second with an average of 8 sparks in each wire chamber gap.

B. Pulser Runs

Pulser runs were done every few hours in order to test the scintillation counters, photomultiplier tubes and associated electronics, and to calibrate the ADC scale factors. Because the events generated during the Pulser run were not beam associated, it was necessary to provide a timing signal from the pulse generator which took the place of the beam's cable pulse.

Through its control of external relays, the computer selected various combinations of scintillation counters and repeatedly pulsed their photodiodes, thus generating artificial events. A total of 2400 events were so generated with each counter appearing in at least 600 events. After all the events were taken, the computer checked to see that every latch, time ADC and pulse height ADC had registered appropriately.

After half of the events had been taken, the computer inserted an extra 7 nsec delay into the pulser timing signal. This had the effect of making all the counter times appear 7 nsec earlier. By examining the change in ADC values corresponding to this 7 nsec shift, we computed the time ADC scale factors.

C. Muon Runs

The purpose of the muon runs was to calibrate the time ADCs. By examining the arrival times of fast muons from the beam target, we determined one additive constant for each ADC that brought all the individual time measurements into one uniform scale of reference.

Muon runs were taken with the three beam sweeping magnets turned off. The trigger logic requirement was set to V T A B C Cable. This trigger favored single charged particles completely traversing our apparatus. Due to the high flux of muons, wire chamber data was not recorded during muon runs.

D. Special Calibration Runs

In addition to the normal running, we took three small data samples under special conditions. One sample contained regenerated $K^0_{\pi 2}$ decays, and was used to study our experimental resolutions. Another sample was designed to collect $K^0_{\pi 3}$ decays and was used to determine our momentum spectrum. The final sample measured the effects of beam interactions.

We collected a sample of $K^0_{\pi 2}$ decays by placing a 35 cm thick

copper block in the beam. The K_L^o beam entering the copper produced a flux of K_s^o's through the phenomenon of regeneration[8]. For this running, the V counter was moved to a position 120 cm upbeam of the T bank, and the regenerator was placed just upbeam of the V. The trigger logic was set to noV 2T 2A Cable.

We readily isolated the charged $K_{\pi 2}^o$ decays of the K_s^o downbeam of the copper block. As this mode is so heavily constrained, we could, for example, compute the pion momenta from the measured angles and known masses. Then, by comparing the measured and computed pion momenta, we determine the precision of our momentum measurements. By similar procedures, we determine the precision of our TOF and beam direction measurements.

In order to determine the K_L^o spectrum, we took a special data sample to collect charged $K_{\pi 3}^o$ decays. As the amount of energy released in the $K_{\pi 3}^o$ decay is smaller than in the leptonic decays, the kinematically computed kaon momenta were subject to less error. During this running, the trigger logic was set to noV 2T 2A cable.

We studied the effects of beam interactions by removing the helium bag from our decay volume and placing a 2.5 cm thick carbon slab in the beam. Carbon was chosen because of its nuclear similarity to helium. The carbon slab completely covered the beam's cross section and contained 42 times as much mass as the entire helium in the decay volume. By selecting those events which verticized near the slab, we isolated an enriched sample of beam interactions. We took data at seven different positions along the beam line in order to obtain a representative measure of the interaction contamination in our main data.

CHAPTER V

Data Analysis Procedure

The output of the on-line computer comprised 470 magnetic tapes of raw data containing 75 billion bits of uncorrelated spark positions and uncalibrated ADC readings. The data was analyzed on an IBM 360/91, and required 300 hours of central processor time. The first stage of the analysis was to convert the spark positions into reconstructed particle tracks, and convert the ADC time data into times measured relative to the beam. The first stage condensed the data to 24 billion bits written on 73 magnetic tapes. The second stage of the analysis was to match front and rear track segments through the magnet, compute particle momenta, identify muons, and correlate tracks coming from common vertices. The second analysis stage condensed the data to 6.3 billion bits written on 17 magnetic tapes. In the third, and final analysis stage, events were selected and categorized according to a large number of criteria. The number of events and their asymmetry in each category finally determined our result.

A. Wire Chamber Reconstruction

The most time consuming phase of the analysis was the correlation of sparks to form particle tracks. Consider the extraction of line segments from one set of 4 x gaps.

Our object was to find all straight lines formed by 3 or 4 sparks. We took every combination of one spark in gap 1 and one spark in gap 3,

and projected the line joining these two sparks into gaps 2 and 4. If we found sparks in the vicinity of the projections, a suitable line had been found. Next, we proceeded in the same manner starting with spark pairs from gaps 2 and 4, and projecting into gaps 1 and 3. During this procedure, it was possible for a given spark to be part of several different lines. For each line so isolated, a best-fit slope, intercept, and quality-of-fit were computed.

We found that the set of lines, extracted by the above procedure, contained a number of spark combinations which accidentally happened to lie on straight lines and did not represent true particle tracks. The expected number of such accidental coincidences can be readily computed. For our conditions, we could expect 4.5 accidental 3 spark lines in each set of 4 gaps. In order to reduce this number, we compared the quality-of-fits of lines with common sparks, and discriminated against the poorer lines.

The above procedure was implemented so that the computer time required per event varied linearly with the number of sparks per gap.

Having found all the lines in one set of x gaps, we applied the same procedure to find all the lines in the corresponding set of y gaps. Next, we determined which x line went with which y line. We projected each x and y line into the u gap. Since the u direction was a linear combination of the x and y directions, we could compute, for each pair of x and y lines, the expected spark position in the u gap. If a spark had actually occurred in the vicinity of this position, a complete track had been found. This sequence was repeated for the v gap, and all possible tracks were collected.

Discrimination against accidental coincidences was handled in a manner similar to the 4 x gap case.

B. Time Calibration

In order to convert the time ADC data into usable measurements, we had to calibrate each ADC. Care was taken to operate the ADCs in their linear range. We needed to determine, for each ADC, the number of ADC counts per nsec and the TOF corresponding to a given ADC value. The pulser runs determined the ADC scale factors, as discussed above.

The muon runs determined one additive constant for each ADC that brought the time measurements from the various counters into one uniform system of reference. Typical muon runs had at least 75 acceptable events in each counter. The times for these events were averaged, and one constant was produced for each ADC. By subtracting these constants from all subsequent ADC data, the times in all counters were measured with respect to a common reference point. All that remained was to find one overall constant, TSHIFT, which would give the time of the average selected muon with respect to the beam production time.

The final time calibration constant, TSHIFT, was determined by fitting the data to the $K^o_{\mu 3}$ hypothesis. Without the knowledge of the K^o_L momentum, all of the kinematic parameters of the $K^o_{\mu 3}$ decay may be computed from the pion and muon directions and momenta, the particle masses, and the K^o_L direction. Unfortunately, this computation results in two solutions due to a quadratic ambiguity. There were, however, a substantial number of events in which the two solutions were almost equal. Using this sample of events, we computed the required value of

TSHIFT needed to bring the measured times into agreement with the times computed from kinematic fitting.

We found that the precision of our TOF measurements was 0.35 nsec (rms deviation). A substantial part of this was apparently due to jitter in the ADC time gate.

C. Track Matching Through the Magnet

In the ideal case, a particle's trajectory has the same y slope and intercept on both sides of the magnet. The upbeam and downbeam x line segments intersect near the magnet's midplane, and the difference in the x slopes determine the particle's momentum. In actual practice, these ideal conditions are somewhat degraded. Each potential pairing must satisfy four equations relating the front and rear track segments. As there is only one unknown, the momentum, there are three constraints.

The track matching proceeded by considering every pair of one upbeam and one downbeam track. An approximate momentum was computed using the average transverse momentum transfer of the magnet and the difference of the upbeam and downbeam x slopes. Using this approximate momentum, vertical focusing corrections were made. Pairings with substantial differences in the corrected y slopes were discarded. Next, the point of closest approach of the two tracks was found. Pairings with a large distance of closest approach, or with a point of closest approach outside the magnet volume, were discarded.

If all the above criteria were met, the particle's trajectory was integrated through the field. Using the upbeam track data, the

approximate momentum, and the magnetic field measurements, we integrated the trajectory through the magnet to the downbeam chambers. The integrated trajectory was required to be consistent with the downbeam track. Having satisfied all these criteria, the track segments were deemed matched. The momentum was corrected to compensate for the difference in the x-direction cosine between the integrated trajectory and the downbeam track. The upbeam track direction was corrected for bending in the magnetic field inside the upbeam chambers.

D. Muon Identification

Muons were identified by their ability to penetrate the lead wall. All particle tracks which had been successfully matched through the magnet and correlated with an A counter, were considered as potential muons. Each of these tracks were projected through the lead wall onto the planes of the B and C counters. For the track to remain a muon candidate, counters were required in the vicinity of the projections in each bank. Due to multiple Coulomb scattering, the particle's actual trajectory could deviate substantially from the projection, especially for low momentum muons. For each set of one track and one A, B, and C counter, a quality-of-fit was computed which reflected the differences in the measured times and the horizontal and vertical position deviations relative to the expected scattering.

The particle's time at each counter bank was determined by the average of the times of the upper and lower photomultiplier tubes. The particle's vertical position was determined by the difference between the upper and lower times. Occasionally, the ADC time data of

a given counter would be lost due to electronic dead time. In order to reduce loss of events, muon candidates were permitted to have inconsistent ADC data in one of the A, B, or C counter banks; the quality-of-fit of such events were heavily penalized, however.

Muons with exceedingly poor quality-of-fits were discarded. Marginally acceptable muons were discarded if a substantially superior muon was found.

E. Selection of $K^{o}_{\mu 3}$ Events

Having matched particle tracks through the magnet and identified the muon, the next step was to find the pion. All tracks which were not identified as muons were potential pions. Each of these tracks was examined to determine if it verticized with the muon. The events were then divided in groups according to the results of their muon identification and pion search.

The prime events fell into two groups. In both groups, there was one unambiguous muon which had a suitable vertex with an upbeam track, assumed to be the pion. In the first group, the pion track was seen in both the upbeam and downbeam chambers, and hence its momentum was measured. For this group we had a one constraint fit to the $K^{o}_{\mu 3}$ hypothesis. In the second group, the pion track was seen only in the upbeam chambers, and we had a zero constraint fit.

The other event groups were not included in the charge asymmetry sample, but were used to study systematic errors. One such group included the events with two distinct muons. In another group, one

penetrating particle was seen, but it was not clear which track correlated with the counters behind the lead wall. And in yet another group, no pion could be found.

By this point, all the involved analysis had been done on the individual events, and all that remained was to count and categorize. We present in Table 1 the numbers of events that survived the various stages of the analysis as discussed above.

TABLE 1

Events Remaining Through the Various Analysis Stages	Number of Events (in thousands)
Events taken during acceptable running conditions	18,200
Events with reconstructable wire chamber data	17,300
Events with sufficient tracks and counters	16,500
Events with identified muons	15,500
Events with a two-prong vertex in the decay volume	14,300

Events With Identifiable Muons are Subdivided Into Event Groups as Described Below		Number of Events (in thousands)
Name	Description	
2 TRACK	Muon identification unambiguous Pion and muon seen in both chambers Vertex in decay volume	8,452
1.5 TRACK	Muon identification unambiguous Pion not seen in downbeam chambers Vertex in decay volume	4,701
NO PION	Muon identification unambiguous No track verticizes with muon	1,235
AMBIGUOUS	Muon identification ambiguous, either track could correlate with counters behind lead wall. Vertex in decay volume	312
2 MUON	Two distinct muons present. Muons come from common vertex in decay volume.	193
MULTIPRONG	Muon identification unambiguous More than two tracks share same vertex.	244

CHAPTER VI

Systematic Corrections

The object of our effort is to determine the intrinsic charge asymmetry in $K^{o}_{\mu 3}$ decay. We are constrained, however, to perform our experiment in the presence of extraneous phenomena which affect the particles we study and shift the value we measure away from the true charge asymmetry. These shifts, or systematic errors, must be individually measured and subtracted out. The identification and computation of systematic corrections is the trickiest aspect of the analysis.

One approach to the computation of these corrections is to try to isolate a subsample of the data in which the systematic error can be experimentally measured. Another approach is to isolate a subsample in which the bias has been enhanced in some known way. Using the asymmetry of the subsample and the enhancement factor, we can determine the error in the total data.

A. Pion Interactions

Pions, being hadronic particles, have a substantial probability of suffering a nuclear interaction in traversing even a small amount of matter. If a pion is absorbed or significantly scattered before it can be detected by our apparatus, the involved event would be lost. If, furthermore, pions of one charge are preferentially lost, a systematic bias will be introduced into the data.

In order to correct for this systematic bias, we measured the pion losses in our apparatus. The material that the pion was required to

traverse in order to be detected was divided into two identical sections. The upbeam section contained the decay volume helium, the helium bag wall, the T counter bank and 10 wire chamber gaps. The downbeam section contained the magnet volume helium, the helium bag wall, a plastic sheet 3 mm thick and 10 wire chamber gaps. The plastic sheet served only to duplicate the mass of the T counter bank. By measuring the pion losses in the downbeam section, we determined the pion losses throughout the apparatus.

We tried, therefore, to isolate events in which the pion should have been present in the downbeam section. By counting the number of such events with and without downbeam pions, we determined the pion losses. The isolation of these special events was done kinematically. Using the upbeam pion direction and the measured muon and kaon data, we could compute the pion momentum as a function of the kaon momentum. The locus of solutions is a closed curve as shown in Fig. 3. For each event, we computed from the upbeam geometry the minimum pion momentum necessary for the pion to traverse the magnet and the downbeam wire chambers; call this lower limit PMIN. In a large number of events, the line P_π = PMIN lay below the closed curve. These events were guaranteed to have pion momenta greater than PMIN, and should have had a detected downbeam pion. In another large number of events, the line P_π=PMIN intersected the curve at two kaon momentum values PK1 and PK2, but the kaon momentum determined by the TOF was substantially greater than either PK1 or PK2. These events also should have had a detected downbeam pion.

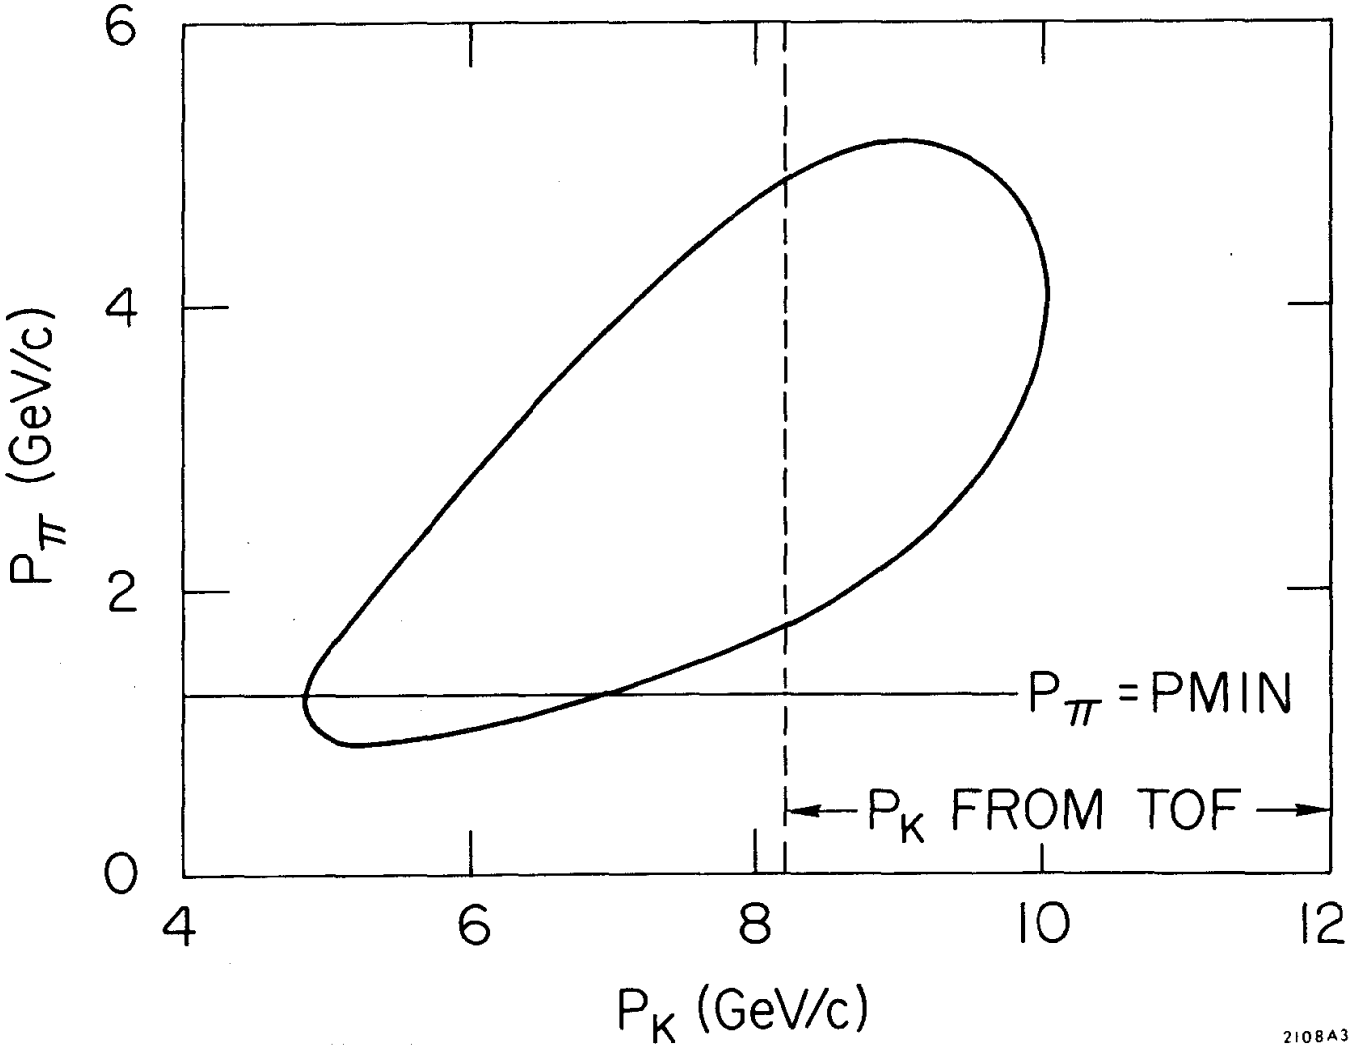

Fig. 3--Locus of solutions for a selected event in the P_π, P_K plane.

The pion loss analysis was carried out using 5.3 million events from the final charge asymmetry data sample. Of these, 2 million passed one of the above criteria and should have had a downbeam pion. Of these 2 million, 7.13% did not have the expected downbeam pion, and were hence classified as pion losses. The muon charge asymmetry of these events was $\delta_{abs} = -3390 \pm 2620$ ppm. This would lead to a correction of $\Delta_1 = 7.13\% \times (\delta_{abs} - \delta_{raw})$, where δ_{raw} is the measured overall asymmetry before corrections.

The measured level of pion losses of 7.13% was substantially larger than the expected hadronic absorption. We found that the kinematic selection procedure described above failed to exclude all events in which the pion could not have traversed the magnet. This led to an apparent loss of 3.5%. Furthermore, spark chamber inefficiencies and program reconstruction failures led to an apparent loss of 1.5-2.0%. The remaining 1.5-2.0% was due to the hadronic absorption in the 0.8 gm/cm^2 of material in the downbeam half of the apparatus. Since our analysis depended only on the product of the fraction lost and asymmetry of the loss, the inclusion of these extraneous charge symmetric losses does not alter our result.

This analysis suffered somewhat because the pion momentum spectrum downbeam of the magnet differed from that upbeam. Low momentum pions were unlikely to traverse the magnet and unlikely to contribute to our pion loss analysis. As the pion absorption cross sections are the largest and the most charge asymmetric at low momenta, we had to refine our analysis. To this end, we endeavored to eliminate the events with the lowest pion momenta by imposing a kinematic cut, and we computed a

second order pion absorption correction from the measured asymmetry in the 2 TRACK events as a function of the measured pion momentum.

As described above, using the upbeam pion direction, and the measured kaon and muon data, we were able to kinematically isolate events in which the pion was almost certainly above any given cutoff. We therefore excluded from our final data sample all events in which the pion momentum might have been below 500 MeV/c. Note that this selection was done without reference to the downbeam pion track, if any. Although this procedure resulted in the loss of 20% of our data, Monte Carlo calculations showed that it reduced the number of pions below 500 MeV/c by a factor of 10.

The final stage of the pion absorption analysis was the computation of a second order correction, Δ_2. In our computation of Δ_1, we selected events which should have had a downbeam pion; let f(P) be the selection probability for an event with pion momentum P. Had f(P) been constant we would have obtained a uniform measure of the pion absorption over all momenta. Δ_2 corrected for the variations in f(P) from its mean value F.

To obtain the differential pion absorption bias as a function of pion momentum, we examined the measured asymmetry in the 2 TRACK events, and assumed that variations in this asymmetry were due entirely to pion absorption. In the 2 TRACK events, the pion was required to traverse twice the amount of material required of all prime events, so that the bias introduced at given momentum P was $(\delta_{2TRK}(P) - \delta_{raw})/2$, where δ_{raw} is as given above and δ_{2TRK} is the measured asymmetry in the 2 TRACK events. If $\epsilon(P)$ is the fraction of all prime events with

pion momentum P, then we have

$$\Delta_2 = \sum_P \epsilon(P) \left[F - f(P)\right] \left[\delta_{2TRK}(P) - \delta_{raw}\right]/2$$

$\epsilon(P)$ and $f(P)$ were determined from the Monte Carlo data, and are shown in Fig. 4. We found that $\Delta_2 = +167 \pm 119$ ppm. Combining Δ_1 and Δ_2, the final pion absorption correction to the charge asymmetry was $-7.13\% \times \delta_{raw} - 75 \pm 221$ ppm.

B. Pion Decay and Pion Interaction

Any process which causes a muon signature, other than a muon from $K^o_{\mu 3}$ decay, could result in a systematic error in the measured asymmetry. As pions can decay to muons $\pi \rightarrow \mu\nu$, these secondary muons could cause spurious events. In addition, should a pion manage to penetrate the lead wall and create a muon signature, it could also cause a spurious event.

The pion decay contamination was a readily calculable correction. A series of Monte Carlo runs generated samples of normal $K^o_{\mu 3}$ decays and $K^o_{\mu 3}$, $K^o_{\pi 3}$ and K^o_{e3} decays with subsequent pion decays. The generated events were then run through the same analysis programs as the real data. The ratios of the number of events entering the final data sample from each of the initial modes determined the level of pion decay events in the real data. As the pion decay process and the questions relating to event acceptance probabilities are well-defined, we are confident that our calculation accurately duplicated the true level of pion decays. The asymmetry of these events were determined from known values[9]. The

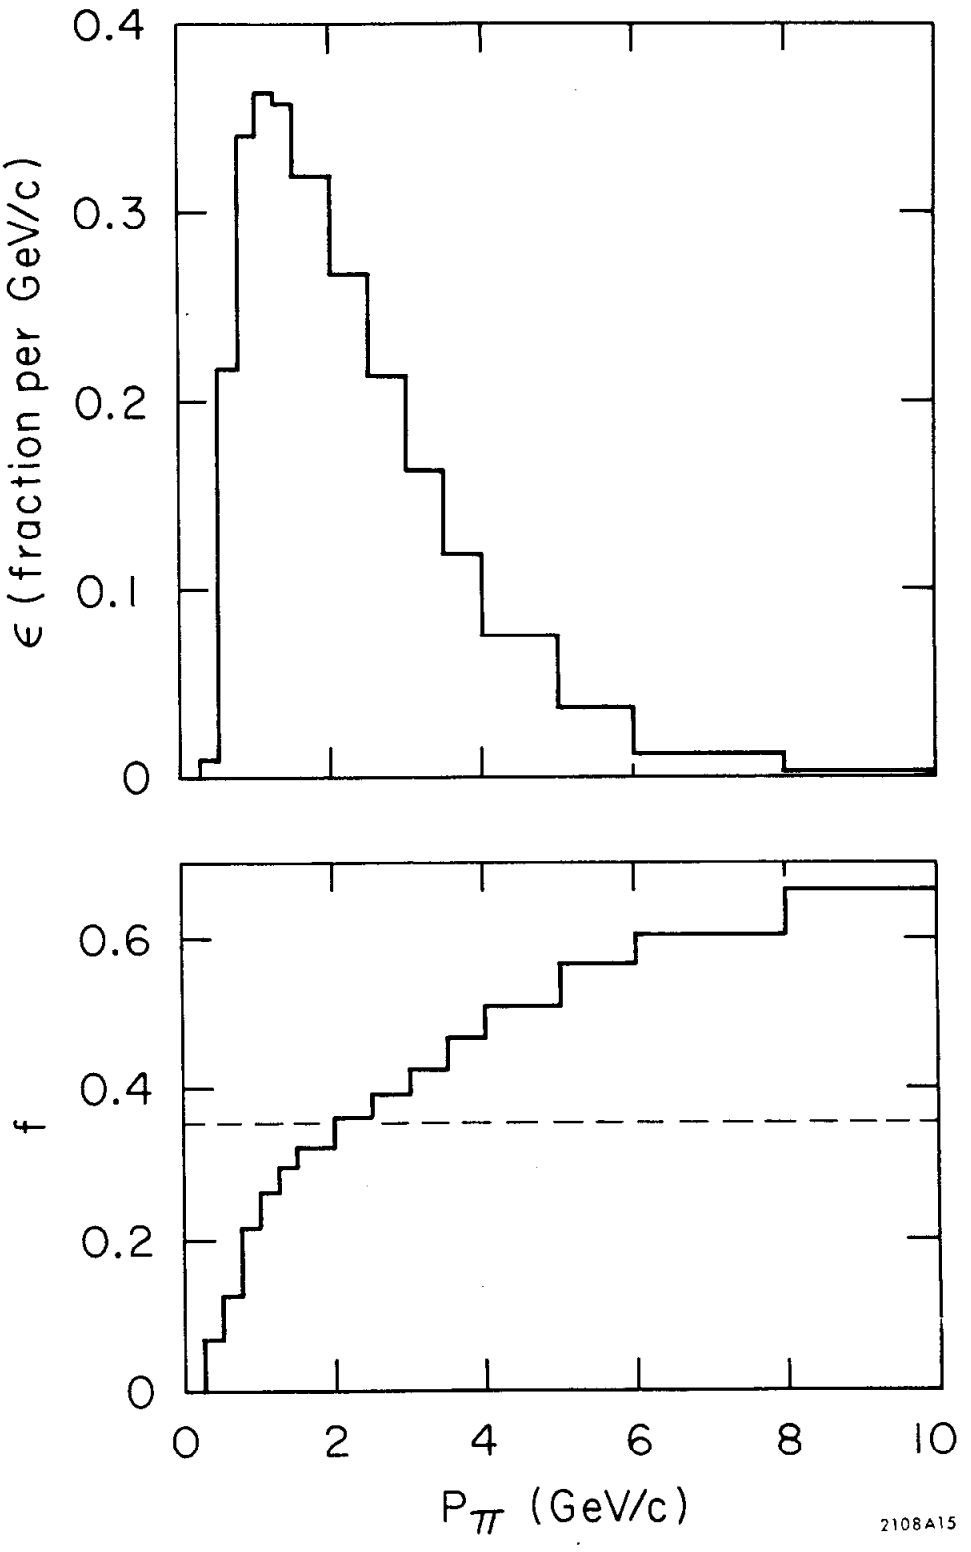

Fig. 4--$\epsilon(P_{\pi})$ and $f(P_{\pi})$ as determined by the Monte Carlo. See the text for their definitions.

contributions from the various modes, and the total correction to the charge asymmetry are given below:

Event Trigger by Pion from:	Asymmetry of Pions	Fraction of Total Data	Correction
$K_L^o \rightarrow \pi^+\pi^-\pi^o$	0. %	3.10 ± 0.09%	0 ppm
$K_L^o \rightarrow \pi e \nu$	-0.32 ± 0.04%	3.87 ± 0.11%	+124 ± 15 ppm
$K_L^o \rightarrow \pi \mu \nu$	-0.29 ± 0.13%	2.95 ± 0.09%	+ 86 ± 38 ppm

Our approach to the pion penetration problem was to isolate the $K_{\pi 3}^o$ decays which crept into our final event sample. Since the intrinsic charge asymmetry in $K_{\pi 3}^o$ decay was identically zero, any asymmetry observed in this event sample must have been due to the pion penetration process. We utilized a remarkable kinematic variable called $P_o'^2$, which effectively separated $K_{\pi 3}^o$ decays from the leptonic decays solely on the basis of the charged decay products[10]. We have computed this variable for all 2 TRACK events in which the transverse momentum of the two charged particles and the missing neutral were less than the kinematic limit for $K_{\pi 3}^o$ decay. In Fig. 5 we have shown the experimental distribution and the distribution of the Monte Carlo data with pion decays. In this distribution, the leptonic modes peak at negative values, and the $K_{\pi 3}^o$ mode peaks near zero. The excess in the experimental distribution near $P_o'^2=0$ is most likely due to pion penetration.

The quantity we required was the product of the number of pion penetrations and their asymmetry. We obtained that product from the experimental asymmetry in the region $0 \leq P_o'^2 \leq 0.01$ $(GeV/c)^2$, and the number and asymmetry of the $K_{\mu 3}^o$ events in that region. Since the number of events here was small, the uncertainty in the experimental

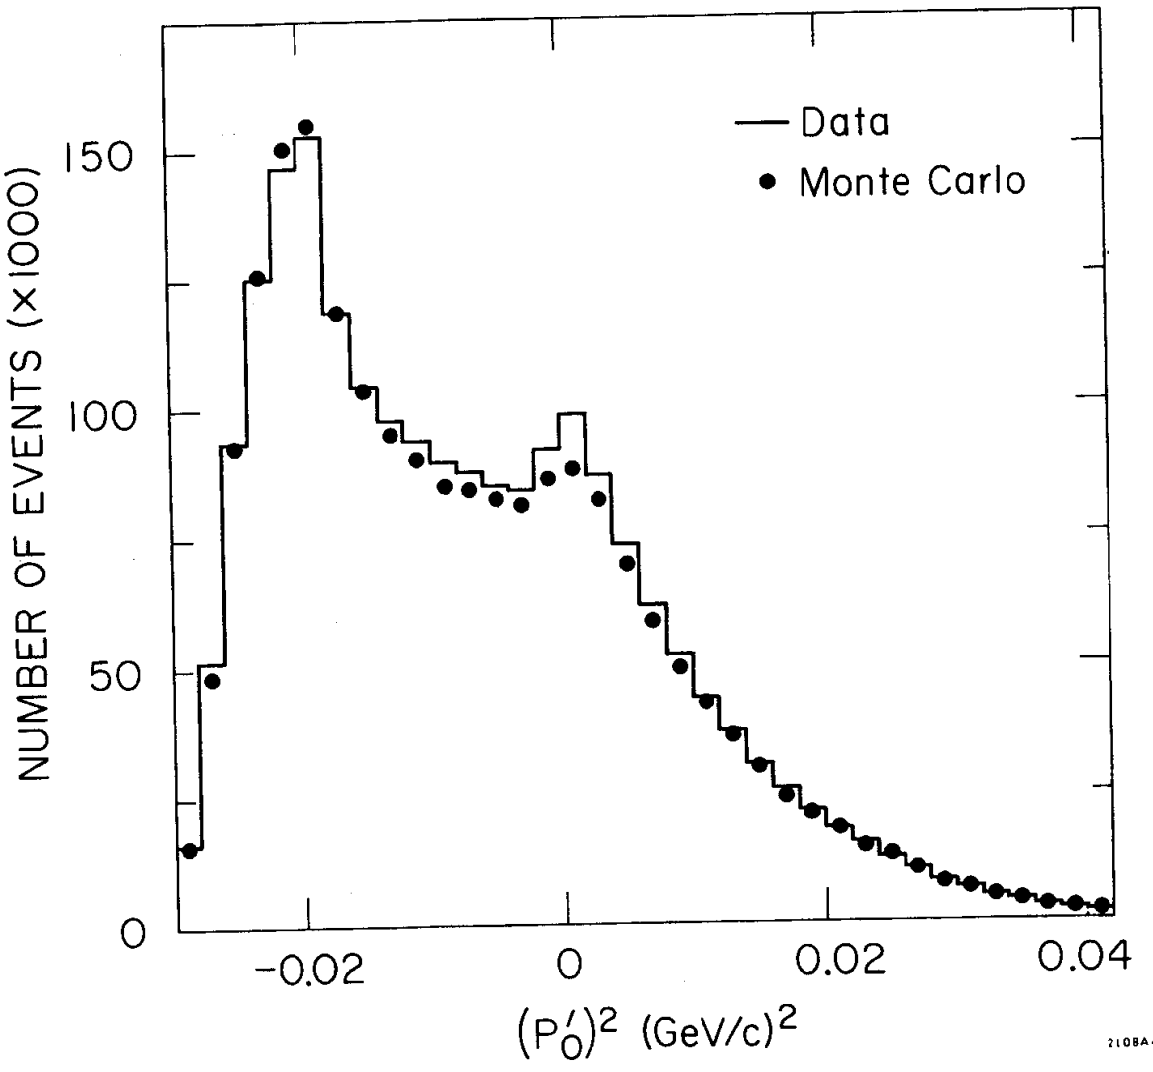

Fig. 5--P'^2_o distribution for selected events for the Monte Carlo and the experimental data.

asymmetry was 50% of its value. Thus, we were insensitive to the uncertainties in the Monte Carlo-determined value of the number of $K^o_{\mu 3}$ events under the pion peak. The required product was determined for each of several momentum bins. Since penetrating pions came from all modes, $K^o_{\mu 3}$, K^o_{e3}, and $K^o_{\pi 3}$, and since our determinations were based solely on the $K^o_{\pi 3}$ sample which lay in selected region, the penetration product and its error were multiplied by an enhancement factor computed from the Monte Carlo data. Again, uncertainties in the Monte Carlo determinations were completely dominated by the large uncertainties in the experimental asymmetry. Combining the results for all momentum bins, we found that the pion penetration correction to the final asymmetry was -394 ± 271 ppm. The results for the individual momentum bins are shown below,

P_μ (GeV/c)	No.Events (10^3)	No.Events in Region (10^3)	$K^o_{\mu 3}$ in Region (%)	Raw Asym. - Asym.due to $K^o_{\mu 3}$ (10^{-4})	Enhancement Factor	Correction (ppm)
< 4	6361	231.5	47	$+14 \pm 21$	3.21	-128 ± 192
4-6	1726	47.9	63	$+79 \pm 46$	6.83	-318 ± 185
> 6	92	1.9	75	-243 ± 227	9.21	$+52 \pm 49$

C. Muon Interactions

The muon interactions which occurred in the lead wall might have introduced systematic errors into our data. Possible asymmetries in Coulomb scattering or in end-of-range behavior required investigation.

The Coulomb scattering proceeding through one photon exchange is

charge symmetric, but higher order terms could be asymmetric. Camilleri et al.[11] have studied muon-proton scattering at high energies, and have found no asymmetry between positive and negative muons to the level of 1%. We have shown in Fig. 6 the charge asymmetry as a function of the observed scattering of the muon in traversing the lead wall. We found no variation and hence concluded that no correction was needed for possible Coulomb scattering asymmetries.

An end-of-range asymmetry resulted from the capture of negative muons by the lead nuclei. Whereas positive muons merely decayed $\mu^+ \rightarrow e^+ \nu\nu$ when they reached the end of their range, negative muon capture and subsequent nuclear disintegration could lead to neutron or photon emission. As these neutrals could traverse more lead and then convert in the C bank, the effective range of negative muons was enhanced.

Another end-of-range asymmetry could result from differences in the ionization losses of positive and negative muons. These losses should be charged symmetric to first order in the electromagnetic interaction, but might be asymmetric in higher orders. Such asymmetries should not exceed 1% of the total ionization loss; thus in our case, the difference should be < 20 MeV[12].

We have examined our data for these end-of-range effects and show in Fig. 7, the charge asymmetry as a function of the computed muon energy at the C bank. A large negative asymmetry is apparent at the lowest muon energy, but the effect does not extend to higher energies. We have therefore excluded from our final data sample all events in which the computed muon energy at the C bank was below 220 MeV.

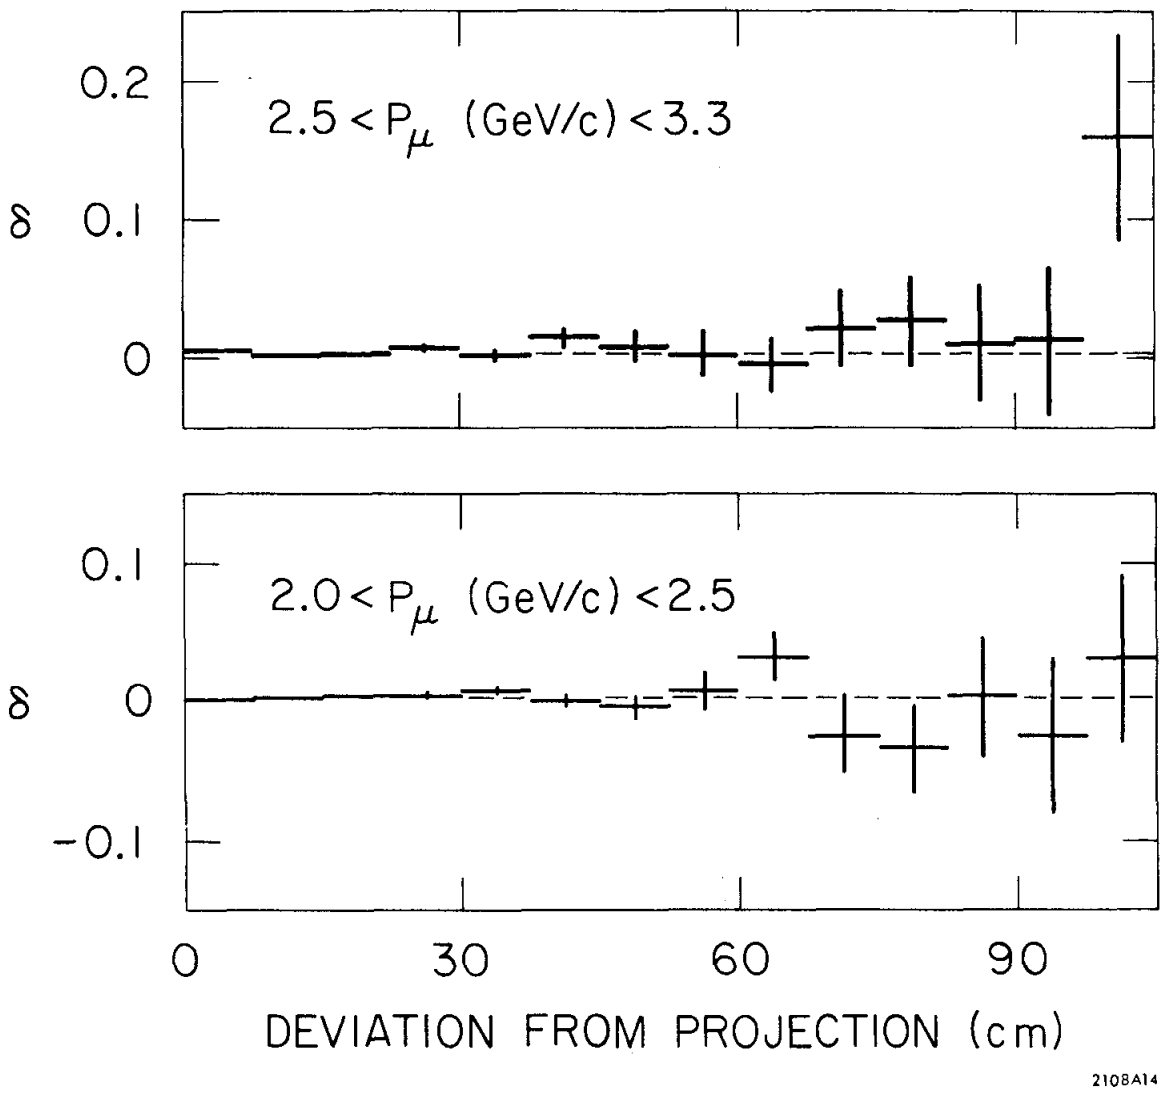

Fig. 6--Charge asymmetry vs. the scattering of the muon in the lead wall. Abscissa measures the horizontal deviation of the muon C counter from the projected wire chamber trajectory.

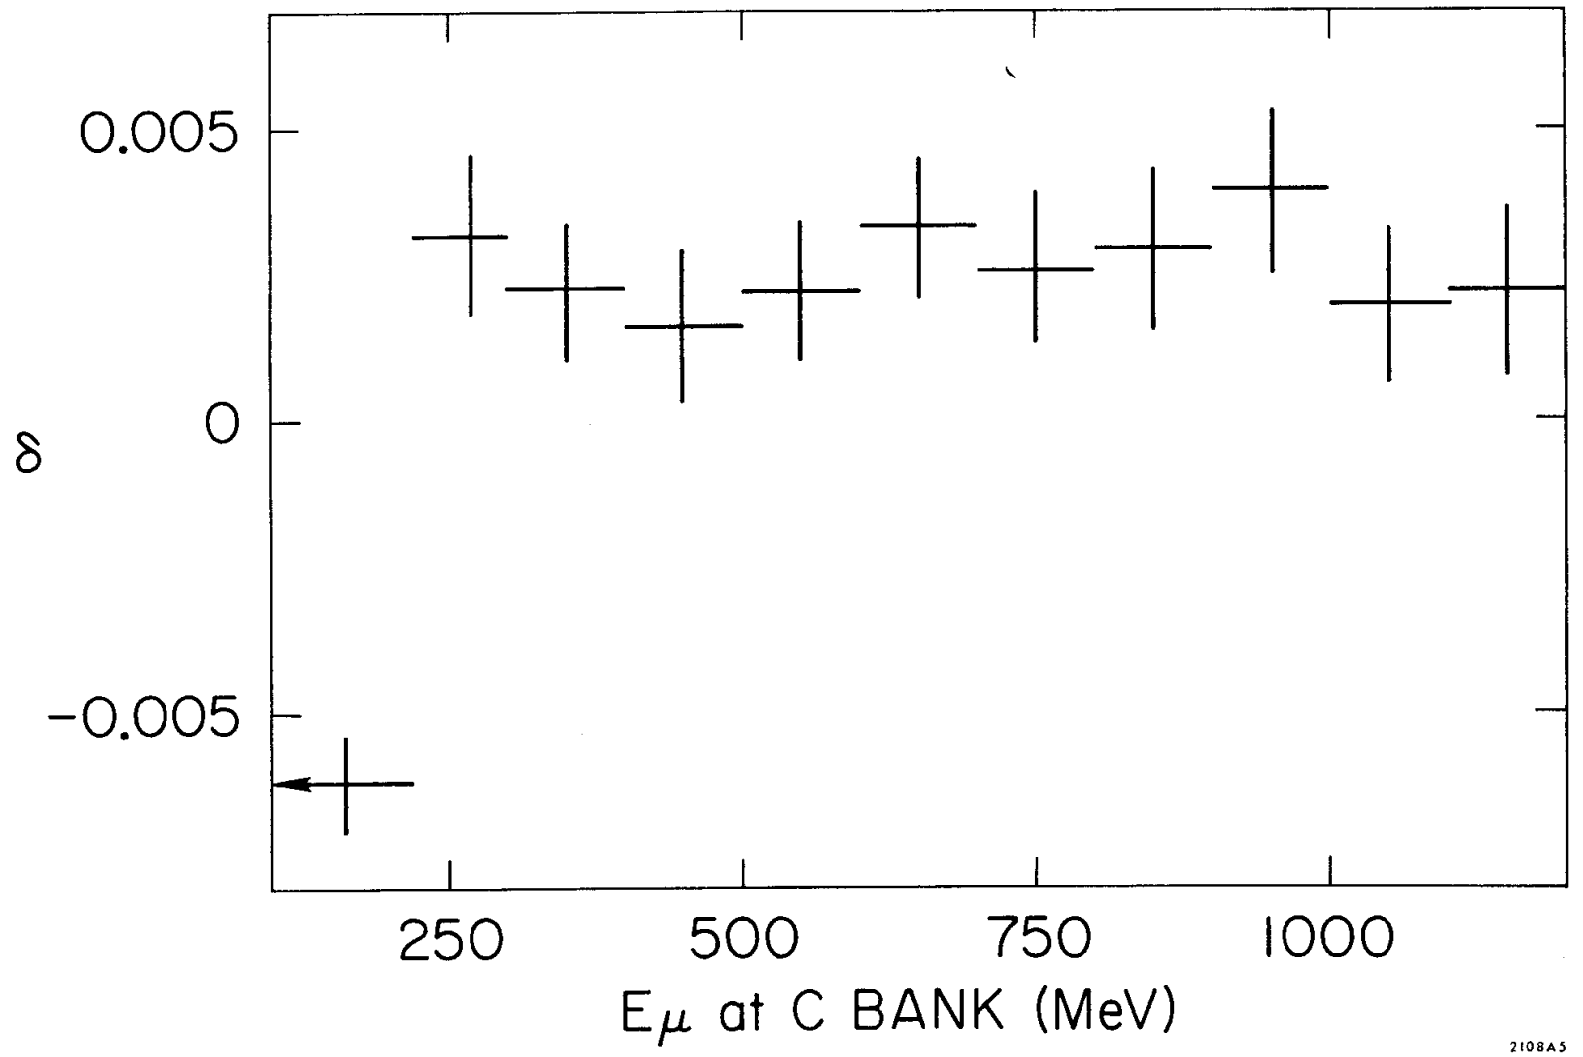

Fig. 7--Charge asymmetry vs. the computed muon energy at the C bank. Events in which this energy was < 220 MeV were placed in the lowest bin, including events in which the muon was apparently below threshold.

D. Beam Interactions

The kaons and neutrons in the beam occasionally underwent nuclear interactions with the helium nuclei in the decay volume. To some extent, these interactions simulated $K^{o}_{\mu 3}$ decays. We had to determine the number of interactions which crept into our final data sample and their asymmetry.

Our approach was to place a 2.5 cm thick carbon slab in the beam to enhance the level of interactions. The carbon slab was successively placed at 7 different positions along the beam line. At each position, data was taken in each magnet polarity. We chose carbon because of its nuclear similarity to helium. Assuming that the production cross sections scale as the two-thirds power of the atomic number, we expected the number of interactions in the slab to be 29 times the number of interactions in all of the decay volume helium. And since the carbon slab interactions were localized, we could isolate event subsamples with 2000 times the normal level of interaction contamination.

In the analysis of the carbon interactions, we restricted our attention to events originating from the slab itself; hence the effects of regeneration were not considered by this analysis. As the level of regeneration was so low, only the coherent transmission regeneration had any significant effect on the charge asymmetry. See VI.G for a discussion of regeneration.

We have shown in Fig. 8 the vertex distributions of the events from two carbon runs, both with and without the final cuts. We have presented in Fig. 9 the transverse momentum distribution of the muon, for the events originating from the slab. Note that these interaction

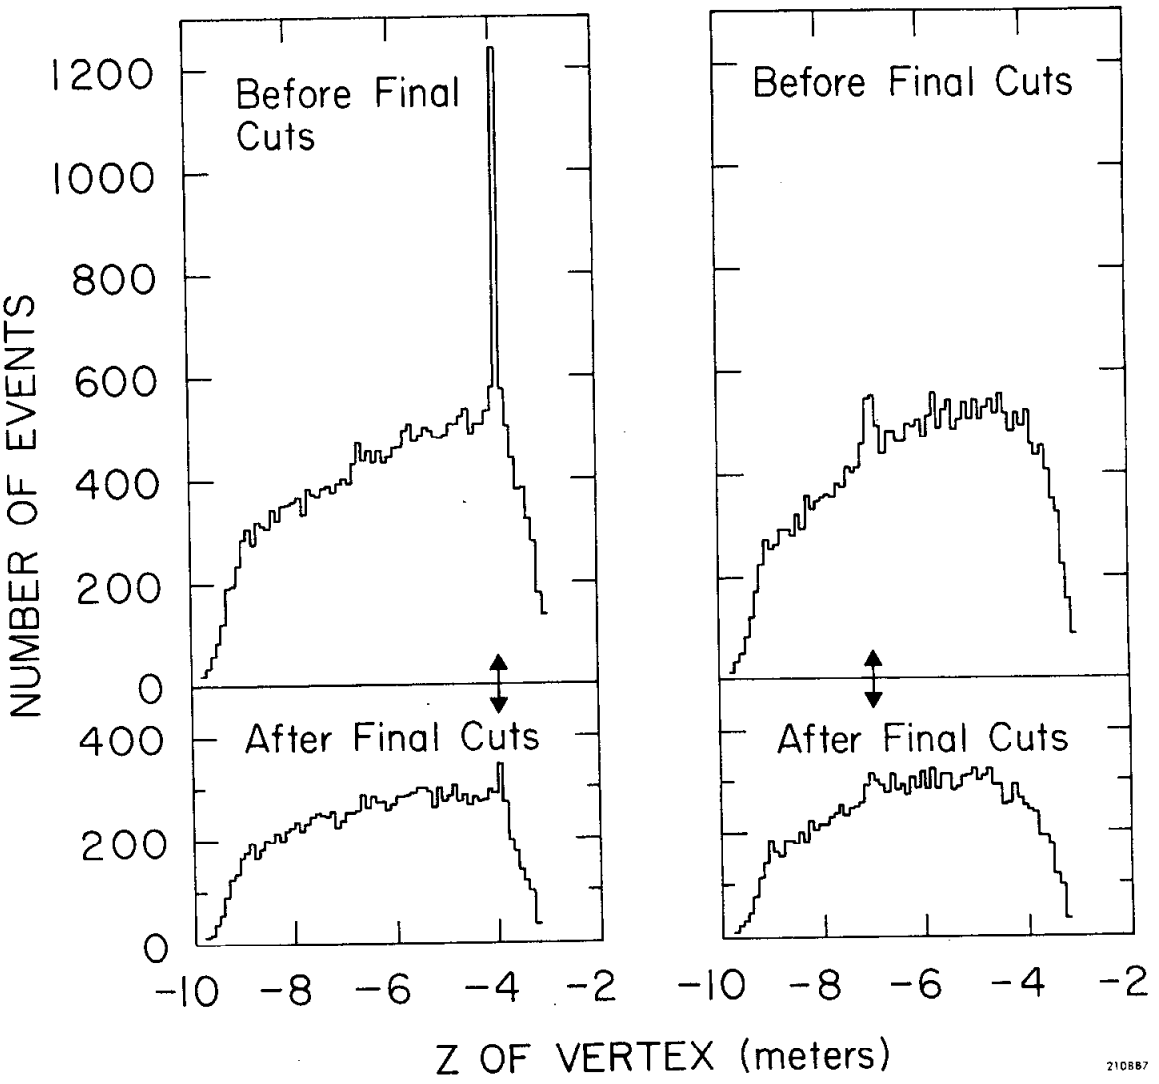

Fig. 8--Z position of vertex for carbon runs. Arrows indicate the position of the carbon slab.

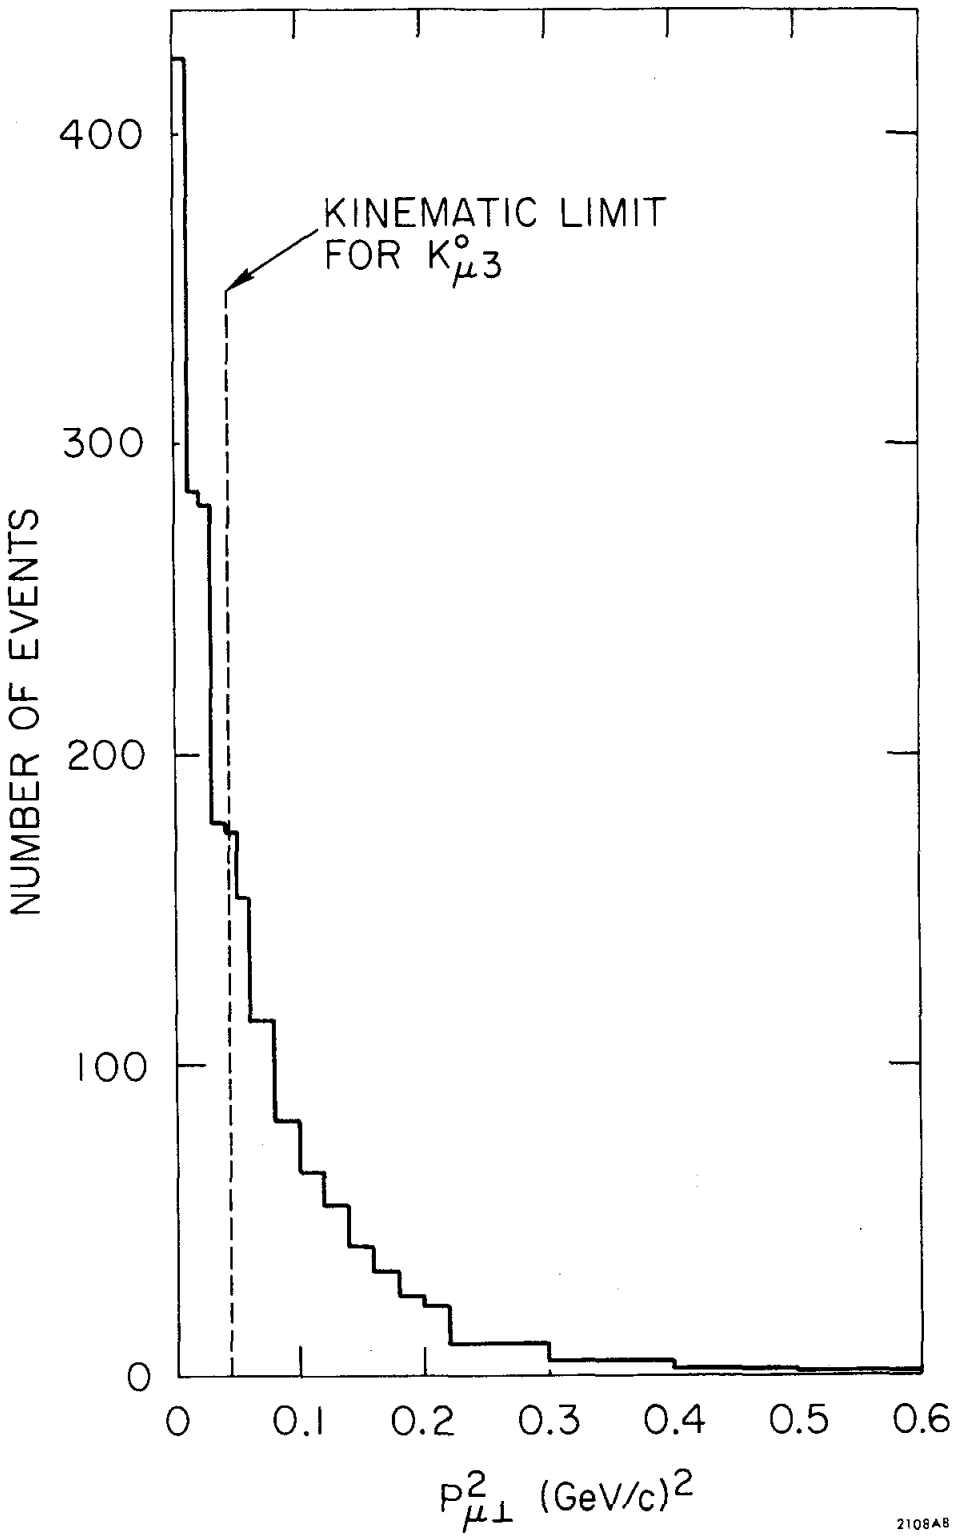

Fig. 9--Transverse momentum squared of the muon for carbon runs. Data is shown before final cuts. K_L decays have been subtracted.

events extend to very high transverse momenta as compared to the $K^{o}_{\mu 3}$ events. To reduce the level of these interactions, we imposed a cut on the muon transverse momentum. See VII.A for details.

The results of our analysis showed that the level of interactions was very low. The event trigger rate was increased by only 3%, and only 0.49% of the final event sample could be attributed to interactions in the carbon slab. Examining the distribution of vertex positions, we computed the level of interactions from the excess in the number of events with vertices near the slab. By noting the change in the asymmetry near the slab, and knowing the fraction of interactions near the slab, we computed the asymmetry of the interaction events. The systematic correction to the charge asymmetry was obtained by averaging the effects from each slab position and dividing by the production enhancement factor. The results are,

Net interaction fraction in carbon data	$0.49 \pm 0.05\%$
Asymmetry of interaction events	$-5.0 \pm 11.0\ \%$
Production enhancement factor	29
Correction to the final charge asymmetry	$+8 \pm 19$ ppm

E. Charge Resolution

If there were any ambiguity in the determination of the charge of the muon, our measured charge asymmetry would be diluted. Any legitimate particle would have a bend angle of at least 30 mrad in traversing our magnet. As our angular resolution was approximately 1 mrad, we could not misidentify a particle's charge. However, there

were a number of events in which we could not positively determine which of the particles was the muon. Events were placed in the AMBIGUOUS group if the best muon candidate was not at least three standard deviations superior to all other candidates. This event group was not included in the final charge asymmetry.

The only remaining concern was that more events of one charge than the other may have fallen into the AMBIGUOUS group. This might have happened if, for example, more positive pions penetrated to the B bank. We have studied the charge asymmetry of the pions which reach the B counters but not the C counters, and we found that their asymmetry was $+2.07 \pm 0.37\%$. Furthermore, we found that 0.23% of our data fell into the AMBIGUOUS group with two B counters and one C counter. This entire sample was most likely due to accidental coincidences between the muon C counter and a pion which penetrated to the B bank. We therefore applied a correction of -25 ± 25 ppm.

F. Geometric Asymmetries and Magnet Reversals

If the detection efficiency was different for positive and negative muon events, a systematic error would have been introduced into the measured charge asymmetry. If, for example, the positive muons always hit the left side of our apparatus, and if that side were slightly smaller, we would have missed a number of positive muon events, and our measured asymmetry would have been reduced. The standard solution to this problem has been to take half the data in each of the two magnet polarities. If positive muons were lost in one polarity, negative muons would be lost in the other polarity. Then, we computed

the charge asymmetry as,

$$\delta = (r - 1)/(r + 1)$$

where,

$$r^2 = \frac{N_{+\uparrow}}{N_{-\uparrow}} \times \frac{N_{+\downarrow}}{N_{-\downarrow}}$$

and, $N_{+\uparrow}$ = the number of μ^+ events with the magnet forward, etc.

Constant geometric asymmetries were completely eliminated by the above method. We have verified this by discarding all events which involve counter C1. The remaining events have a large geometric asymmetry, but have the same charge asymmetry as the original sample, within statistics. By successively discarding events with C1, C2, etc., we generated a series of data groups with widely varying geometric asymmetries. As seen in Fig. 10, the charge asymmetries of these data groups show no variation with the geometric asymmetry.

Finally, we tested for time-dependent geometric asymmetries, such as gradual changes in spark chamber efficiencies. Consecutive data runs with opposite magnet polarities were paired and the charge asymmetry was computed for each of 169 such pairs. We found that these asymmetries were consistent with a constant value, and that the χ^2 for this hypothesis was 178 for 168 degrees of freedom, corresponding to a confidence level of 27%.

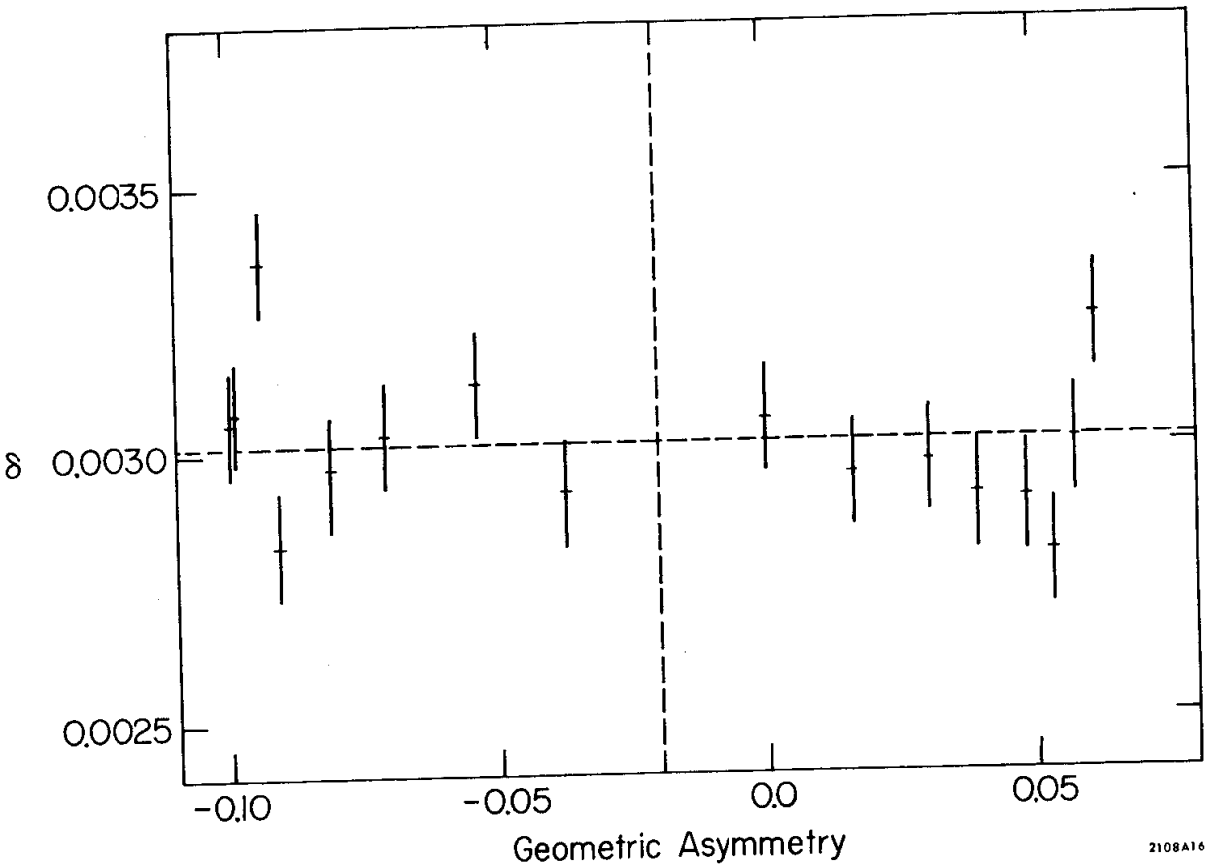

Fig. 10--Charge asymmetry vs. geometric asymmetry for data groups defined in the text. Error bars indicate the mean deviation from the overall average. Dotted lines indicate the charge asymmetry and geometric asymmetry for the total data.

G. Regeneration

In addition to the beam interactions discussed above, we considered the effects of coherent transmission regeneration in the V counter and in the helium of the decay volume. Owing to the difference in K^o, $\bar{K}^o$ cross sections, when a K_L^o beam traverses material the K^o, $\bar{K}^o$ mixture is altered. As this is equivalent to changing the CP-mixing parameter ϵ, coherent regeneration introduces a systematic bias into the measured charge asymmetry. Consider, first, the regeneration in the V counter. The coherent regeneration amplitude following a thin regenerator is known to be[13],

$$A(Z) = if_{21}\ \lambda\ N\ L\ \exp\left[(i\delta - \tfrac{1}{2})\ Z/\Lambda\right]$$

where,

f_{21} is half the difference of the K^o and $\bar{K}^o$ forward scattering amplitudes

λ is the K_L^o wave length

N is the density of nuclei

L is the thickness of the regenerator

δ is the K_L^o, K_s^o mass difference in units of the K_s^o lifetime

Λ is the K_s^o mean decay length

Z is the distance downbeam of the regenerator

Using the charged kaon total cross sections measurements of Abrams et al.[14], and the phase determinations of Bott-Bodenhausen et al.[15], we found that $|A(o)| = 3 \times 10^{-4}$. The effect of this bias was to increase the charge asymmetry by 2 Re A(Z) in the vicinity of the

regenerator. This led to a correction of -12 ± 6 ppm.

Consider now the regeneration in the helium of the decay volume. Within a difuse regenerator, the coherent amplitude attains a constant value after the first few K_S^o mean decay lengths. This value is,

$$A = if_{21} N \lambda \Lambda/(\tfrac{1}{2} - i\delta)$$

Extrapolating the helium parameters from the carbon data, we found that helium regeneration led to a correction of -4 ± 2 ppm.

H. Pulse Height Variation

As noted in III.C, we observed slight shifts in the photomultiplier tube pulse height distributions between the magnet-forward (↿) and magnet-reverse (⇂) data. We have shown in Fig. 11 the pulse height spectra for C15, the counter in which this effect was the greatest. The mean ↿ pulse height was 1.25 counts larger than the mean ⇂ pulse height. As the overall mean was 233 counts, the spectra were shifted by 0.54%. As C15 was near the left edge of the C bank, only μ^+'s hit it for ↿, and only μ^-'s hit it for ⇂. Thus the μ^- pulse heights were effectively reduced by 0.54% in this counter, and more μ^-'s would fail to trigger the C15 discriminator. Fortunately, the discrimination thresholds were typically 30 counts, and very few events were affected. We have estimated the number of events per count at threshold, and the fractional shift between the μ^+ and μ^- spectra for each counter.

The results, averaged within each counter bank, are

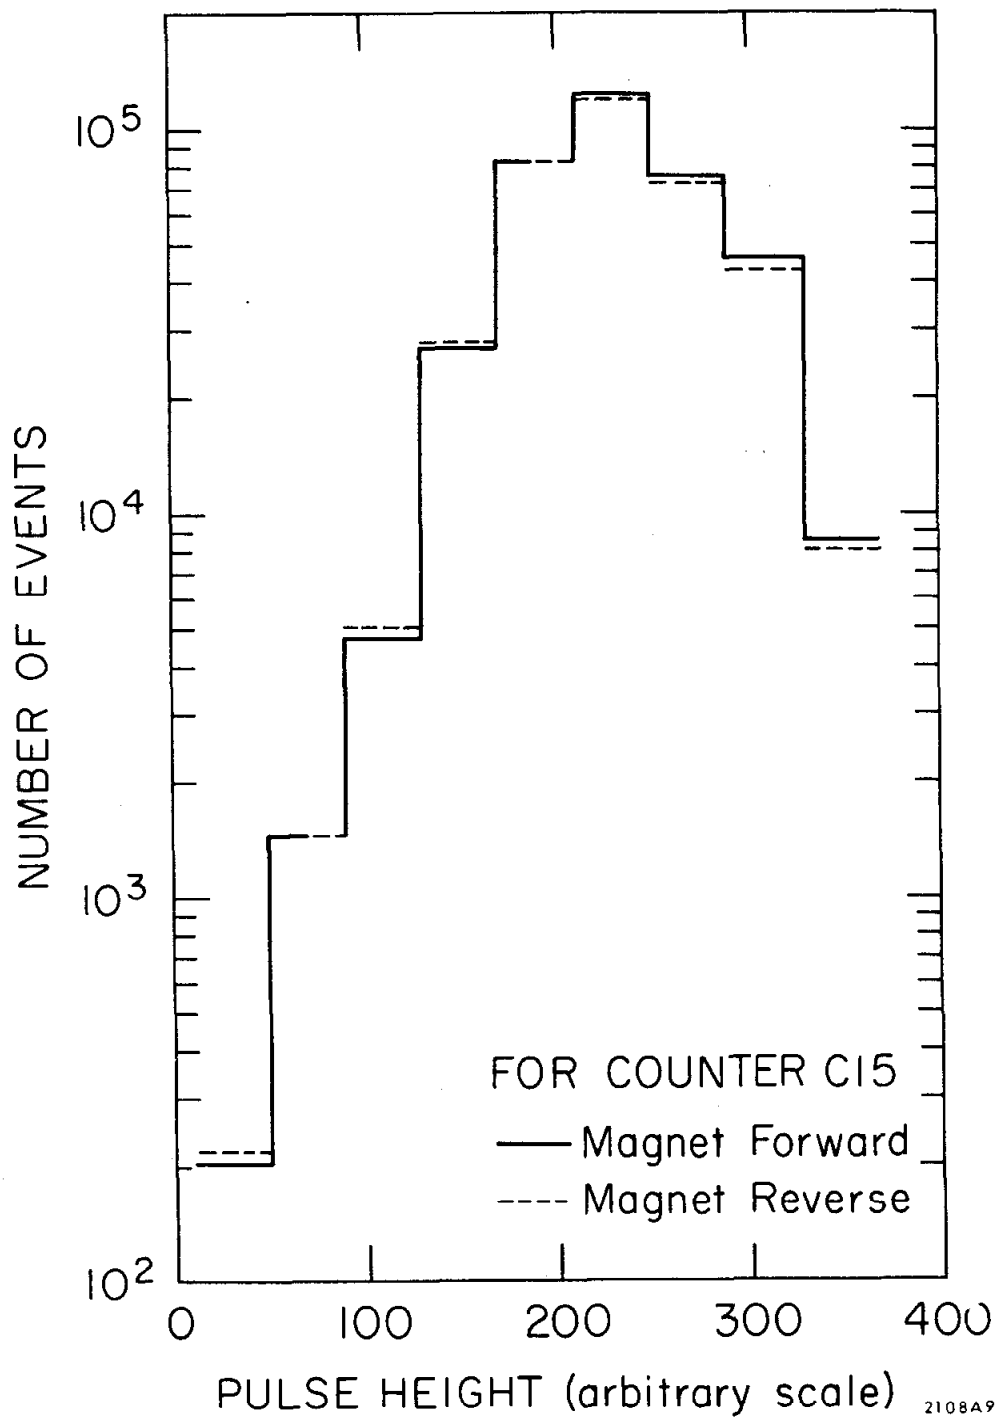

Fig. 11--Pulse height distribution for counter C15 for magnet forward and reverse.

Counter Bank	Fraction of Events per count at Threshold (ppm)	μ^+/μ^- Spectrum Shift (counts)	Correction (ppm)
A	20 ± 20	-0.05±0.02	+0.5±0.5
B	32 ± 32	+0.07±0.02	-1.1±1.1
C	10 ± 10	+0.30±0.02	-1.5±1.5
Total			-2.1±1.9

I. Accidentals

Our K_L^o beam had an especially poor duty cycle. In order to obtain a high event rate, we were forced to run at very high instantaneous beam currents. In addition to the tracks of interest in each event, we recorded any uncorrelated particle tracks which happened to occur within the resolving time of our spark chambers. As the beam current was high, the number of such accidental tracks was high. Our reconstruction programs had to select and correlate the tracks which belonged to the true event, and discard the accidental background. Some fraction of these accidental tracks would, however, get through our selection criteria, and become erroneously correlated into an "event". We had to determine the bias introduced by this process.

There were three main sources of accidentals: accidental matching of upbeam and downbeam particle tracks through the magnet, accidental correlation of muon counters and downbeam tracks, and accidental verticizing of upbeam tracks.

As discussed above, there were three constraints in the matching of tracks through the magnet. Accidental matches were, therefore, extremely unlikely. Furthermore, such matches should not have been charge asymmetric.

The correlation of muon counters and downbeam tracks was also well-constrained, and accidentals were unlikely. For each muon, we computed a χ^2 which measured how well the data fit the constraints. We expected the χ^2 distribution of the accidental "muons" to extend to much higher values than that of the real muons. Hence, if these accidentals altered the measured charge asymmetry, they would have done so to a much greater extent at high χ^2 values where they were in relative abundance. Upon examining the charge asymmetry as a function of the muon χ^2, we found that the asymmetry was constant to a confidence level of 97.5%.

The accidental verticizing of upbeam tracks was more worrisome. There was only one constraint in the verticizing procedure, and we found that there was a background of charge asymmetric, lone muons. To determine the effect of these accidentals, we took each of the muons which did not verticize with another track in the same event, and tried to verticize it with a pion track from a genuine $K^o_{\mu 3}$ event. If the tracks formed an acceptable vertex, we subjected this "event" to all the event selection cuts. We found that 788 $\pm$ 10 ppm of our events were due to such accidental verticizing, and that the asymmetry of these events was + 1.0 $\pm$ 1.3%. This led to a correction of -8 $\pm$ 10 ppm.

J. General Checks

We have applied corrections for all the systematic biases which we have been able to isolate in the data. What about other, as yet unknown, systematic biases; how can we insure that our measured asymmetry is free of such unknown effects? The answer is, of course, that

we can not be completely certain that we have thought of all possible biases, but we can reduce the chance of substantial error by a thorough search of our data. If there were a sizable systematic bias in our data, we could reasonably expect that it would affect certain subsamples of our data more than others. Then it would readily be spotted in the appropriate distribution.

In this spirit we have examined the asymmetry as a function of many variables. The intrinsic asymmetry should be constant, and independent of all the variables we will discuss. For each such distribution, therefore, we will show in the figure the χ^2 per degree of freedom obtained in attempting to fit the charge asymmetry to a constant value. We emphasize that in these distributions, the systematic corrections already discussed have not been subtracted, as we are unable to quantitatively discuss the various functional dependences of the corrections. We can, however, make some qualitative comments about the major corrections, those being the pion absorption, the pion penetration and the pion decay corrections.

Pion decay is inhibited at high energies by relativistic factors, so that we expect relatively more contamination from pion decays at low muon energy, where it will reduce the measured asymmetry.

Pion absorption is largest, and most asymmetric, at low pion energies. The primary effect of pion absorption in our data is to reduce the measured asymmetry.

The probability of pion penetration, either through direct transmission or through hadronic showers, increases at high energies. This tends to affect the distributions at high "muon" energy.

Finally, there are natural correlations among the variables of interest. For example, effects that arise due to phenomena at low muon energy will reappear in many other distributions. We must restrict our attention to the search for large anomalous fluctuations until such time as the detailed dependences of the systematic corrections already applied can be determined.

We expect possible biases to vary with physical parameters such as the momenta of particles, the position of the decay vertex, and the counters.

Fig. 12 shows the dependence of the charge asymmetry on the momentum of the muon, and of the pion. The muon distribution is consistent with no dependence. The pion distribution suffers from a substantial deviation (χ^2=11.0) in the region 0.5-1.0 GeV/c. This is precisely the region of maximum differential pion absorption where we expect a reduced asymmetry, although the magnitude of the effect is considerably larger than expected.

Fig. 13 shows the dependence of the charge asymmetry on the muon counters. To compute this distribution, we have combined the events from symmetrically located counters, and computed the asymmetry for each pair. All the distributions agree adequately with the hypothesis of no dependence on the scintillation counter involved. Among other things, this supports the contention of VI.H that the corrections due to pulse height variation are indeed small.

Fig. 14 shows the dependence of the charge asymmetry on the Z position of the vertex. Again the asymmetry appears flat, and reinforces the contention of VI.G that the regeneration in the V counter is small.

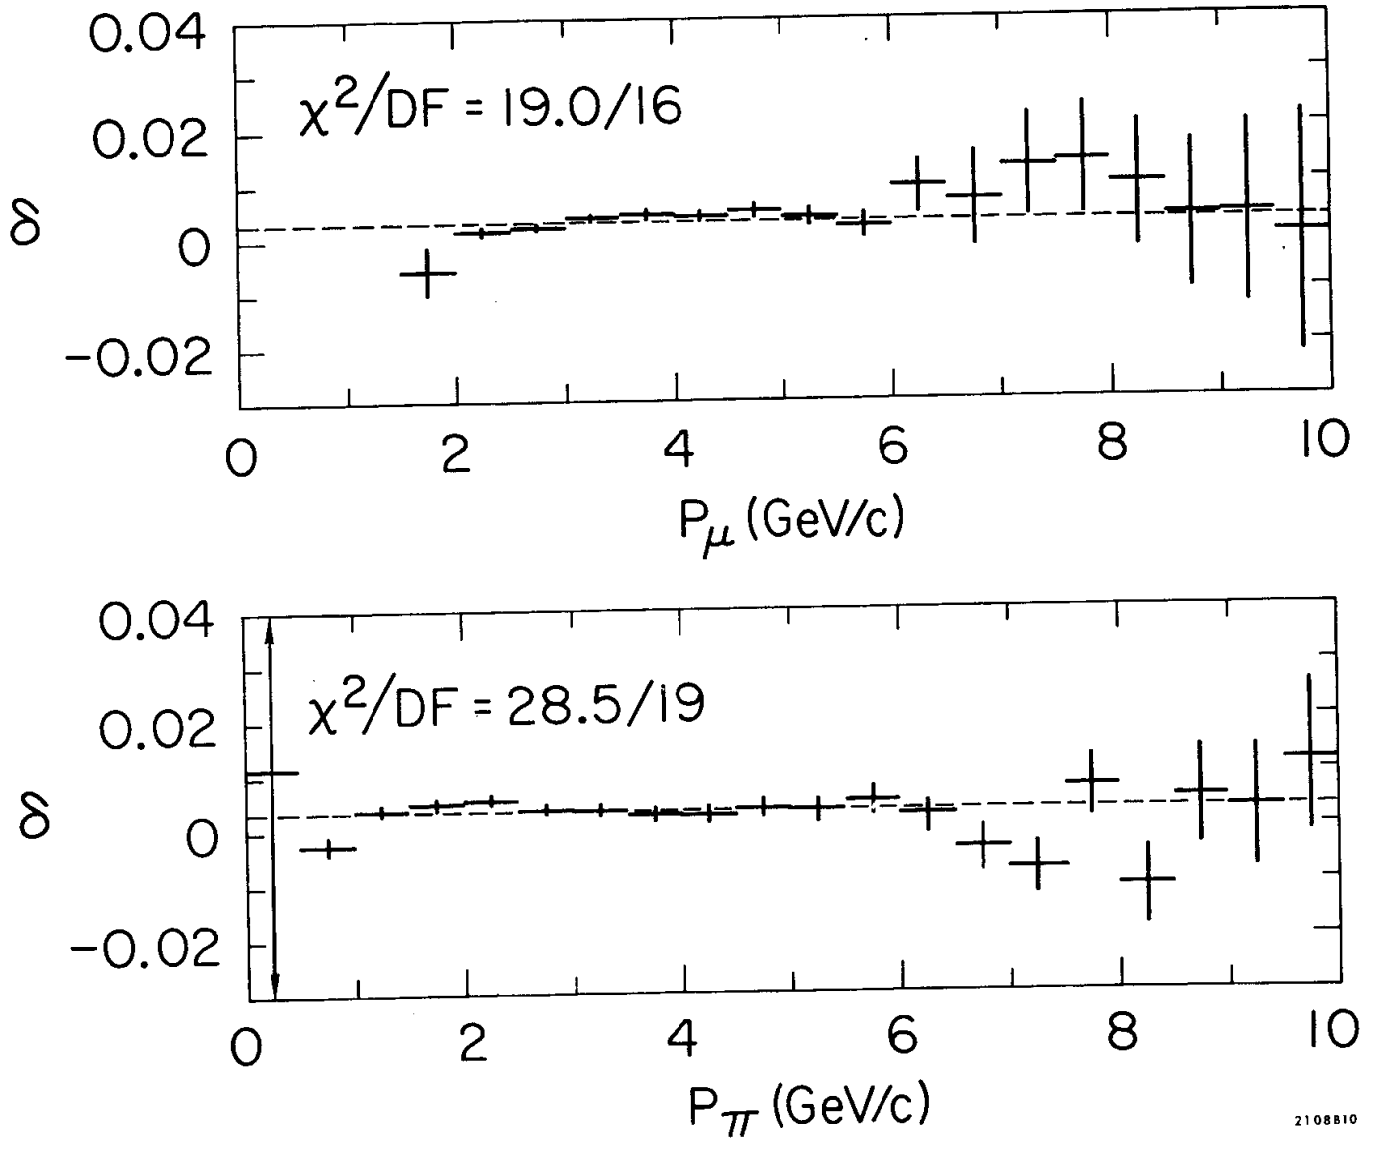

Fig. 12--Charge asymmetry vs. momentum of the muon and the pion.

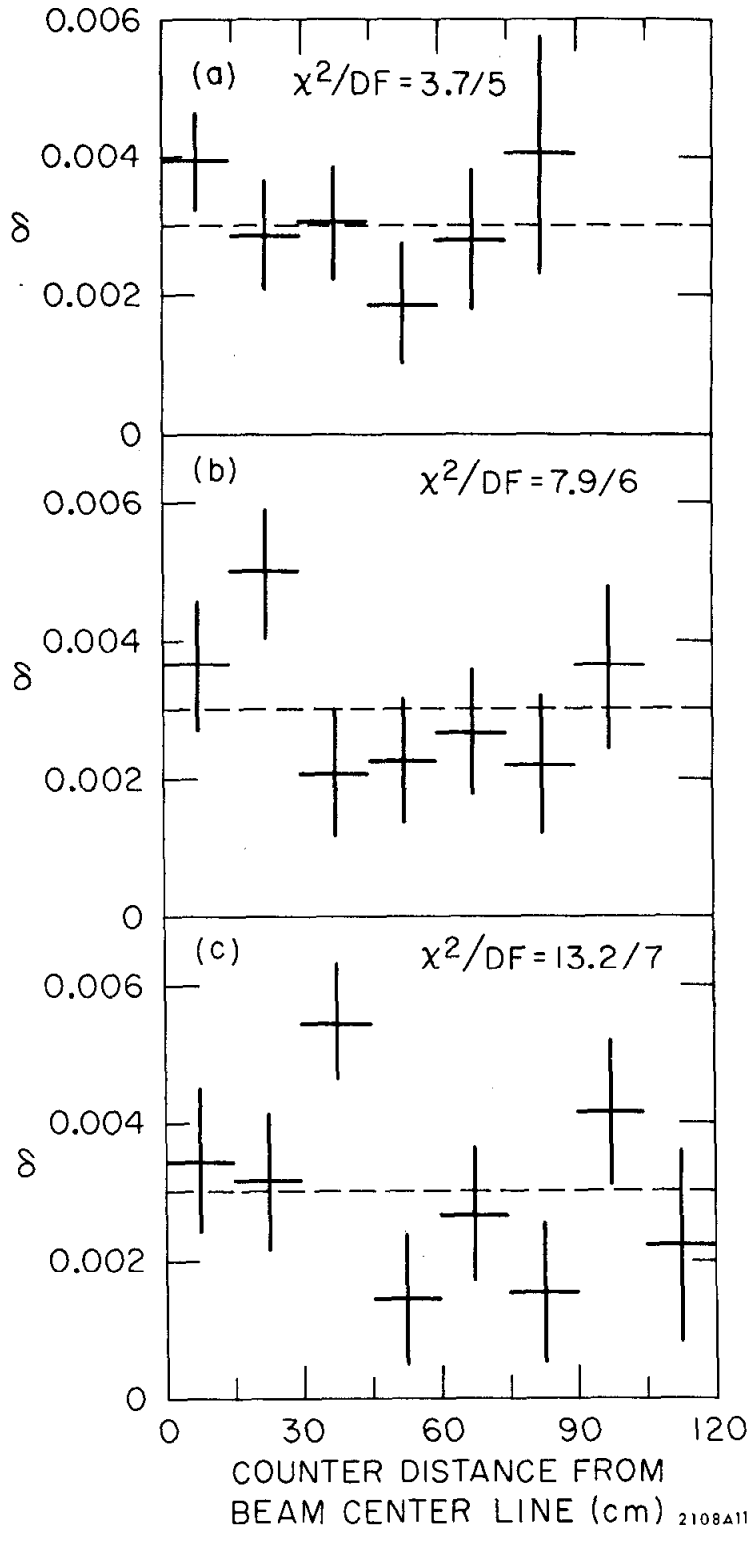

Fig. 13--Charge asymmetry vs. position in the A, B and C banks.

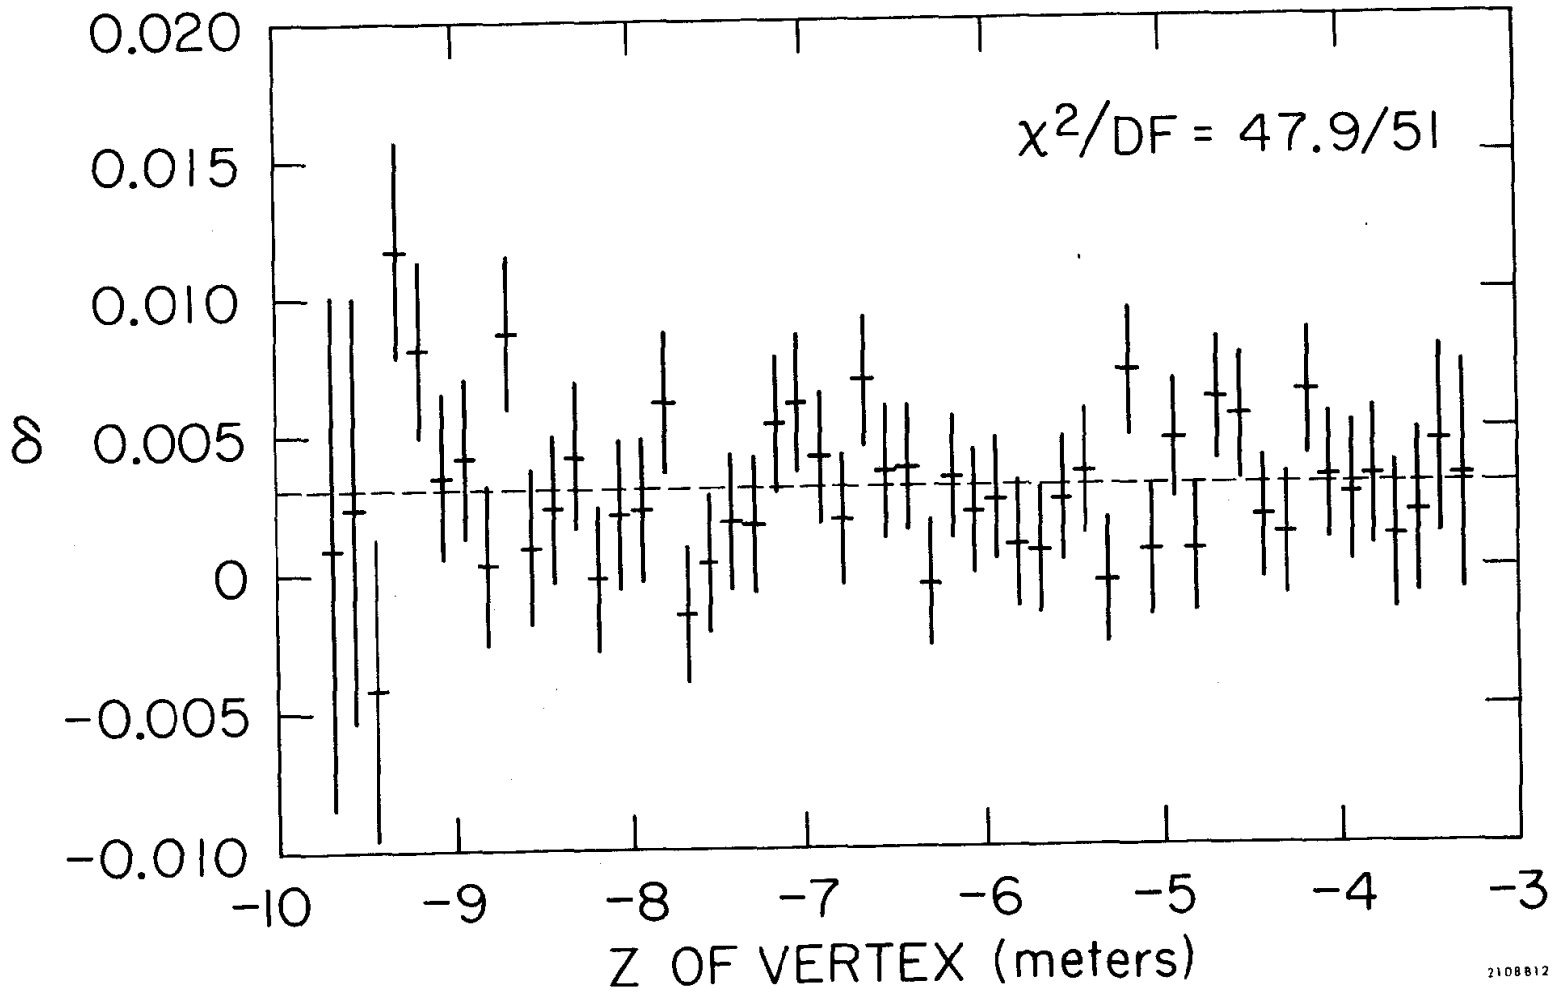

Fig. 14--Charge asymmetry vs. Z position of vertex.

Since the helium bag around the decay volume is much larger than the beam's cross section, we have no reason to suspect contamination at the X,Y edges of the decay volume. We have looked for such effects, however, and find that the asymmetry is independent of X and Y position of the decay vertex.

CHAPTER VII

Conclusions

A. Final Cuts

In order to reduce the level of certain contaminations, we imposed a set of final cuts on the data. These cuts eliminated much good data, but greatly purified the final sample.

To avoid end-of-range effects, events were eliminated if the computed muon energy at the C bank was less than 220 MeV. See IV.C for details.

To reduce beam interaction contamination, events were eliminated if the transverse momentum squared of the muon was greater than 0.035 $(\text{GeV/c})^2$. The kinematic limit for $K^o_{\mu 3}$ is 0.047 $(\text{GeV/c})^2$.

To reduce pion absorbtion difficulties at low pion momenta, events were eliminated if the kinematically fitted pion momentum could have been less than 500 MeV/c. In making this cut, no reference was made to the measured pion momentum, if any. See IV.A for details.

Events in which the pion and muon fired the same T counter were eliminated since they were inconsistent with the trigger requirement. These events may have resulted through the occurence of an accidental T counter, or by erroneous assignment of both tracks to the same counter in cases of adjacent T counters. This class of events had a normal asymmetry, but was excluded in order to improve the consistency between the Monte Carlo and the real data.

To exclude beam interactions in the T counters, events with vertices within 25 cm of the T bank were eliminated.

To reduce the magnitude of the correction required for pion decay and pion penetration, we endeavored to eliminate all $K^o_{\pi 3}$ events from our data. These events were isolated kinematically as discussed in VI.B. We eliminated all events in which the value of P'^2_o was positive. This cut reduced the number of $K^o_{\pi 3}$ events in the data by a factor of 6 as determined from the Monte Carlo events.

To reduce hadronic penetration (primarily by pions), events with more than one C counter were eliminated. This cut required some consideration to insure that it did not asymmetrically eliminate genuine events. Genuine events might fire multiple C counters through three mechanisms. A $K^o_{\mu 3}$ event in which both the muon and the pion penetrated the lead wall could have two C counters. However, such events would be classified as AMBIGUOUS, and would not be included in the final data sample. A single muon could physically traverse two counters if it hit near a counter boundary with a large incident angle; this effect would be charge symmetric. A muon could fire multiple C counters by delta ray emission. Delta ray emission is charge symmetric to first order, but could be charge asymmetric in higher orders, just as Coulomb scattering. The charge asymmetry as a function of the number of C counters is shown in Fig. 15. We see that the asymmetry of the events with two C counters is 3300 ppm above the normal value. This asymmetry may be due to hadronic penetration, in which case we wish to eliminate these events, or it may be due to delta ray asymmetries, in which case we would introduce a bias by eliminating these events. Fortunately, the total effect of these events is only 130 ppm so we can afford to straddle the issue by eliminating these events and applying a

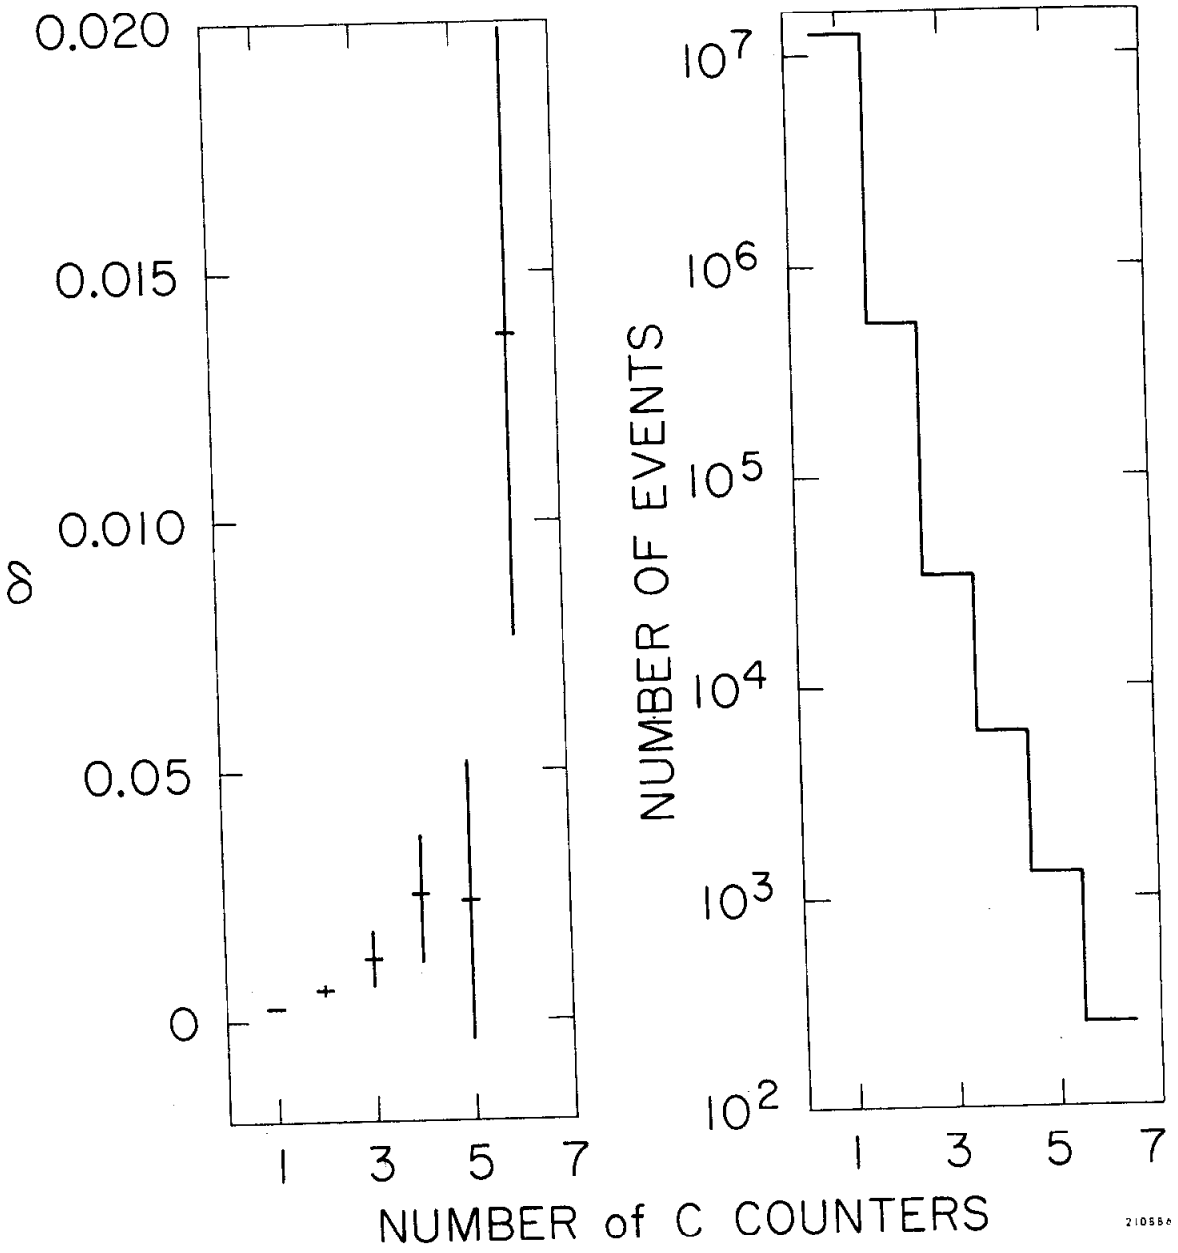

Fig. 15--Charge asymmetry vs. number of C counters.

correction of +65 ± 65 ppm.

We present in Table 2 the results of these cuts.

B. Numerical Results

Certain of the systematic corrections given below are stated as an absolute bias. The relative biases indicate that the measured asymmetry has been shifted by a multiplicative constant rather than an additive constant. For example, if the charge of the muon were reversed in 1% of our events, then the magnitude of the measured asymmetry would be 2% too low. We would then want to apply a relative correction of +2%. The tabulation of all corrections and the final asymmetry are given in Table 3.

Assuming the absence of $\Delta S = -\Delta Q$ amplitudes, we find that,

$$\mathrm{Re}\ \epsilon = (1.39 \pm 0.25) \times 10^{-3}$$

Alternatively, we may assume the current world values for Re ϵ[5], in which case we find that,

$$\frac{1 - |x|^2}{|1-x|^2} = 0.92 \pm 0.17$$

Assuming that x is small, as shown by the data of reference 6, we may refine our limit to,

$$|1-x| = 1.04 \pm 0.10$$

This limit provides a sensitive measure of the real part of x only, which we find to be

$$\mathrm{Re}\ x = -0.04 \pm 0.10 \qquad (K^0_{\mu 3})$$

TABLE 2

Asymmetries and Number of Events for each Cut
(See VII.A for description of cuts)

Event Sample	Asymmetry in ppm	No.Events Millions
2 TRACK and 1.5 TRACK events before cuts	+ 3060 ± 280	13.15
Events cut for low muon energy at C bank	- 6120 ± 860	1.35
Events cut for high muon transverse momentum	+14320 ± 1560	0.41
Events cut for multiple C counters	+ 7030 ± 1400	0.51
Events cut for same pion and muon T counter	+ 1810 ± 1260	0.63
Events cut for vertex near T counter bank	- 1080 ± 4770	0.04
Events cut for fitted pion momentum	+ 6370 ± 690	2.07
Events cut as possible $K_{\pi 3}$ candidates	+ 3690 ± 1520	0.43
Events surviving all cuts		
2 TRACK	+ 3360 ± 440	5.07
1.5 TRACK	+ 2330 ± 620	2.63
Total	+ 3010 ± 360	7.70

TABLE 3

Tabulation of Systematic Corrections and Final Charge Asymmetry

Description	Absolute Corrections ppm	Relative Corrections %
Events passing final cuts	+3010 ± 360	
Pion absorption	- 75 ± 221	-7.13 ± 0.02
Pion decay	+ 210 ± 41	+7.33 ± 0.15
Pion penetration	- 394 ± 271	
Beam Interactions	+ 8 ± 19	
Pulse height variation	- 2 ± 2	
Asymmetric contributions to AMBIGUOUS event group	- 25 ± 25	
Asymmetric elimination of multiple C counter events	+ 65 ± 65	
Regeneration	- 16 ± 6	
Accidentals	- 8 ± 10	
Sum	+2773 ± 509	+0.20 ± 0.15
Evaluation of relative corrections	+ 6 ± 5	
Final corrected Charge Asymmetry	+2779 ± 509	

Furthermore, we have determined the charge asymmetry as a function of the Dalitz variables. We present in Fig. 16 the charge asymmetry vs. the kinetic energy of the muon, the pion, and the neutrino in the K_L^o center of mass system. Note that systematic corrections have not been applied to these distributions. We see no strong correlations in these distributions. In fitting them to the hypothesis of a constant asymmetry, we find

$$T_\mu : \chi^2 = 24.5 \text{ for } 26 \text{ degrees of freedom}$$
$$T_\pi : \chi^2 = 36.7 \text{ for } 23 \text{ degrees of freedom}$$
$$T_\nu : \chi^2 = 46.5 \text{ for } 49 \text{ degrees of freedom}$$

None of the fits are significantly improved by the inclusion of linear or quadratic dependences. The T_π fit has a confidence level of only 4%, but no systematic effect is apparent. We now compare our result with previous determinations[9].

$K_{\mu 3}^o$ Charge Asymmetry

Dorfan et.al.	$(+4.03 \pm 1.34) \times 10^{-3}$
McCarthy et.al.	$(+2.14 \pm 1.01) \times 10^{-3}$
This experiment	$(+2.78 \pm 0.51) \times 10^{-3}$
Average	$(+2.80 \pm 0.43) \times 10^{-3}$

K_{e3}^o Charge Asymmetry

Saal et.al.	$(+2.46 \pm 0.59) \times 10^{-3}$
Marx et.al.	$(+3.46 \pm 0.33) \times 10^{-3}$
Ashford et.al.	$(+3.6 \pm 1.8) \times 10^{-3}$
Average	$(+3.23 \pm 0.28) \times 10^{-3}$

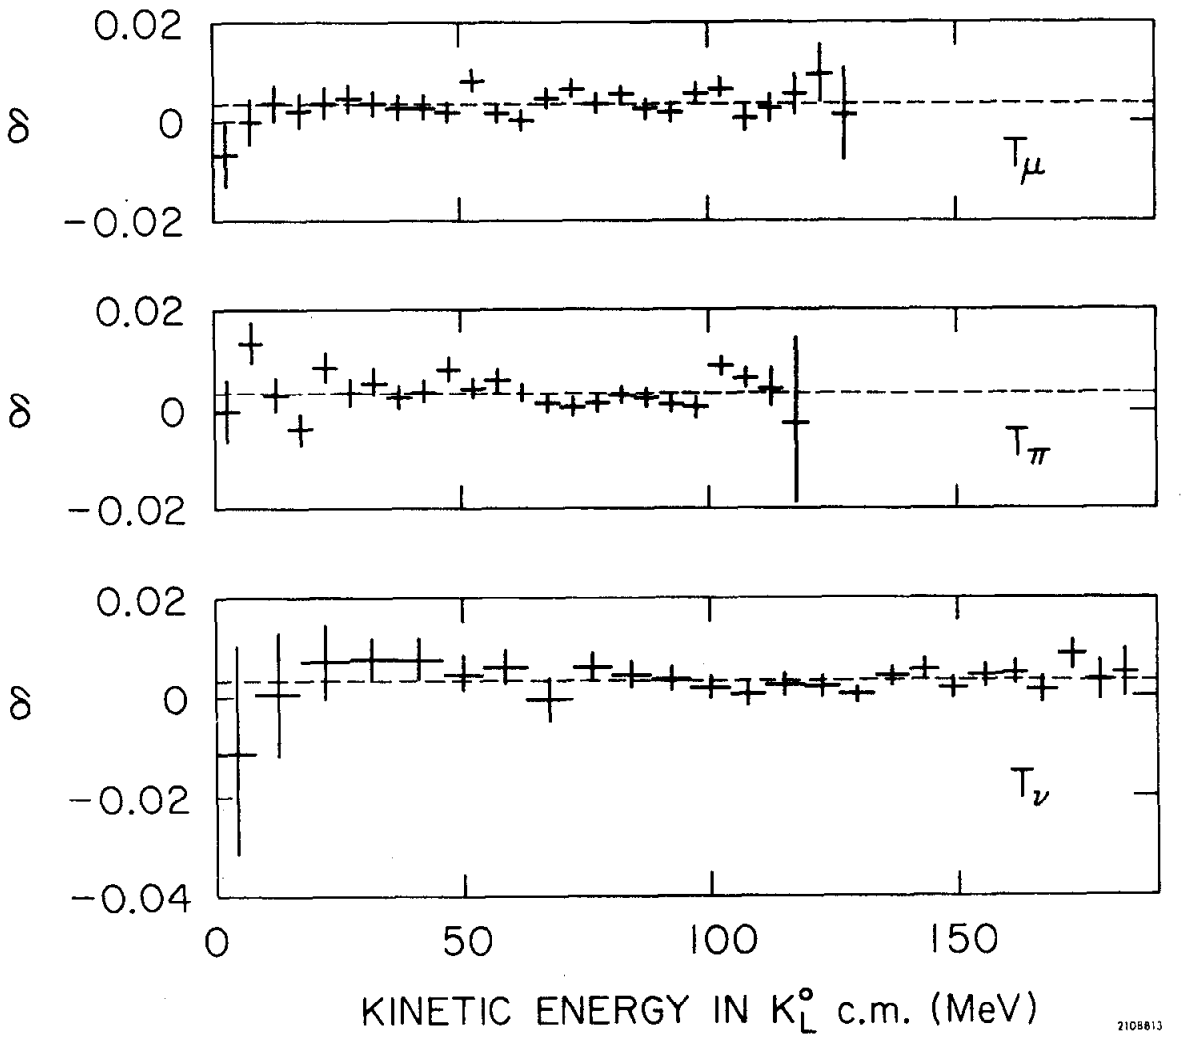

Fig. 16--Charge asymmetry vs. kinetic energies of the muon, pion, and neutrino in the K_L° center of mass.

$K^0_{\mu 3}$ and K^0_{e3} Average	$(+3.10 \pm 0.24) \times 10^{-3}$
World Average of 2 Re ϵ =	$(+3.02 \pm 0.14) \times 10^{-3}$

The last line gives the current value of 2 Re ϵ as derived from the experiments on the magnitude and phase of η_{+-} and η_{oo}.

C. Interpretation

We have measured the charge asymmetry in $K^0_{\mu 3}$ decay and find it to be $(+2.78 \pm 0.51) \times 10^{-3}$. This result is consistent with the previous measurements in $K^0_{\mu 3}$ and K^0_{e3}. Considering the difficulty of the experiments and their extreme sensitivity to systematic bias, the degree of consistency among these experiments is remarkable.

Assuming that CP violation is due solely to CP mixing in the decay states, as in the superweak theory, and assuming the absence of $\Delta S = -\Delta Q$ amplitudes, the results of the $K^0_L \rightarrow \pi\pi$ experiments predict a value of $(+3.02 \pm 0.14) \times 10^{-3}$ for charge asymmetry in both $K^0_{\mu 3}$ and K^0_{e3}, and further predict that the asymmetry should not vary across the Dalitz plot. The measured charge asymmetries are consistent with this value. The constancy of the asymmetry vs. the center of mass energies of the decay products, as discussed VII.B, is further confirmation of these assumptions.

We feel confident in our result for several reasons. The multitude and precision of our information greatly reduced our susceptibility to systematic biases. Due to the design of the apparatus, the biases were small; none exceeded our final uncertainty. Almost all of

these corrections were directly measured rather than theoretically derived. Monte Carlo methods were only used in the determination of certain well-defined geometric and kinematic ratios. Finally, we have checked many sensitive distributions for the existence of unknown biases. We see no evidence thereof, and feel that should such unknown biases exist, their effect on our result would be small.

REFERENCES

1. J. H. Christensen, J. W. Cronin, V. L. Fitch, and R. Turlay, Phys. Rev. Letters 13, 138 (1964).

2. M. Gell-Mann and A. Pais, Phys. Rev. 97 1387 (1955).

3. We say "apparent violation" because the possibility remains that the known forces, which all conserve CP, might conspire in some elusive manner which, because we experiment on an earth and not an antiearth, interacts preferentially with the K^o (or $\bar{K}^o$) component of the long-lived kaon. The unseen interaction could mix the CP states without violating CP invariance. Many such suggestions have come and gone, however, and it seems unlikely that CP violation is not a fundamental phenomenom of nature.

4. L. Wolfenstein, Phys. Rev. Letters 13, 562 (1964).

5. K. Winter, Rapporteur's talk, Amsterdam International Conference on Elementary Particles (1971). Winter's best fit to the world data yields;

$$\eta_{+-} = (1.956 \pm 0.062) \times 10^{-3} \times \exp\,[i(41.6 \pm 2.9)^o]$$
$$\eta_{oo} = (2.15 \pm 0.09) \times 10^{-3} \times \exp\,[i(42.3 \pm 2.9)^o]$$
$$\epsilon = (2.02 \pm 0.05) \times 10^{-3} \times \exp\,[i(41.8 \pm 2.8)^o]$$
$$\mathrm{Re} = (1.51 \pm 0.07) \times 10^{-3}$$

6. K. Winter, Rapporteur's talk, Amsterdam International Conference on Elementary Particles (1971), and G. Burgun et al., Lett. Nuovo Cimento 2, 1169 (1971).

7. R. Coombes et al., Nucl. Instr. Methods 98, 317 (1972).

8. A.Pais and O. Piccioni, Phys. Rev. 100, 1487 (1955).

9. Particle Data Group, Physics Letters 39B, (1972). References to individual experiments are as follows:

D. Dorfan et al., Phys. Rev. Letters 19, 987 (1967).

H. Saal, Thesis, Nevis-169 (1969).

J. Marx, Physics Letters 32B, 219 (1970).

R. McCarthy, Thesis, LBL-550 (1971).

V. Ashford, Phys. Letters 38B, 47 (1972).

10. $P_o'^2$ is defined to be the square of the kaon momentum in the longitudinal rest frame of the charged pair, assuming them to be pions. The original use of this variable is due to A. Astier et al., Proceedings of the Aix-en-Provence International Conference on Elementary Particles (1961). This work is now available in English as SLAC Translation No. 143.

11. L. Camilleri et al., Phys. Rev. Letters 23, 149 (1969).

12. Recent work by J. D. Jackson indicates that the range difference of μ^+ and μ^- should be 2-3 x 10^{-3} and should favor longer ranged μ^-'s. An effect of this small magnitude will be completely eliminated by the cut we impose.

13. W. A. W. Mehlhop et al., Phys. Rev. 172, 1613 (1968).

14. R. J. Abrams et al., Phys. Rev. 4D, 3235 (1971).

15. Bott-Bodenhausen et al., Phys. Letters 24B, 438 (1967).