Inclusive electron–nucleus scattering at $x > 1$ and high Q^2

J. Arrington,[1,2] C. S. Armstrong,[3] T. Averett,[2,3] O. K. Baker,[4,5] L. de Bever,[6] C. W. Bochna,[7] W. Boeglin,[8] B. Bray,[2] R. D. Carlini,[5] G. Collins,[9] C. Cothran,[10] D. Crabb,[10] D. Day,[10] J. A. Dunne,[5] D. Dutta,[11] R. Ent,[5] B. W. Filippone,[2] A. Honegger,[6] E. W. Hughes,[2] J. Jensen,[2] J. Jourdan,[6] C. E. Keppel,[4,5] D. M. Koltenuk,[12] R. Lindgren,[10] A. Lung,[9] D. J. Mack,[5] J. McCarthy,[10] R. D. McKeown,[2] D. Meekins,[3] J. H. Mitchell,[5] H. G. Mkrtchyan,[5] G. Niculescu,[4,13] I. Niculescu,[4,14] T. Petitjean,[6] O. Rondon,[10] I. Sick,[6] C. Smith,[10] B. Terburg,[7] W. F. Vulcan,[5] S. A. Wood,[5] C. Yan,[5] J. Zhao,[6] and B. Zihlmann[10]

[1] *Argonne National Laboratory, Argonne, Illinois, 60439*
[2] *Kellogg Radiation Laboratory, California Institute of Technology, Pasadena CA 91125*
[3] *College of William and Mary, Williamsburg, VA 23187*
[4] *Hampton University, Hampton VA 23668*
[5] *Thomas Jefferson National Accelerator Facility, Newport News VA 23606*
[6] *University of Basel, Basel Switzerland*
[7] *University of Illinois, Urbana-Champaign IL 61801*
[8] *Florida International University, University Park FL 33199*
[9] *University of Maryland, College Park MD 20742*
[10] *University of Virginia, Charlottesville VA 22901*
[11] *Northwestern University, Evanston IL 60201*
[12] *University of Pennsylvania, Philadelphia PA 19104*
[13] *Ohio University, Athens, Ohio, 45701*
[14] *George Washington University, Washington, D.C. 20052*

(Dated: March 12, 2005)

The abstract

PACS numbers: 25.30.Fj, 13.60.Hb

INTRODUCTION

EXPERIMENT

The experiment was one of the first to be carried out in Hall C at Jefferson Lab. An electron beam of 4.??? GeV impinged on a series of targets (liquid hydrogen and deuterim and C, Fe, and Au solid targets), and inclusive scattering of electrons was measured from 15 to 55 (74?) degrees.

Electron beam

The experiment used the electron beam provided by the Continuous Electron Beam Accelerator Facility at Jefferson Lab. A detailed description of the accelerator facility and electron beam parameters is given in [1]. E89-008 was one of the first experiments run at Jefferson Lab, and at the time the maximum electron beam energy was limited to 4.0 GeV. Data was taken with beam currents ranging from 10 to 80 μA, all at the maximum available energy.

The beam has a microstructure that consists of short (1.67 ps) bursts of beam coming at 1497 MHz. Each hall receives one third of these bursts, giving a pulse train of 499 MHz in each hall. The Beam Switch Yard takes the beam that has been extracted from the accelerator and sends the pulses to the individual halls. Beams of different energies can be simultaneously delivered into the three experimental halls. While CEBAF is capable of delivering simulataneous beams to each of the three experimental halls, only Hall C was operational in 1996.

After the electron beam has been accelerated to the desired energy in the main accelerator, it can be delivered into one or more of the three experimental halls. The beam is split at the end of the accelerator, and beam for Hall C is sent through the Hall C arc and into the end station. The arc is equipped with a variety of magnets used to focus and steer the beam, as well as several monitors to measure the energy, current, position, and profile of the beam. Figures 1 and 2 show the hardware in the Hall C Arc and Hall C beamline.

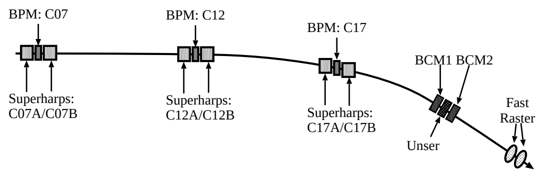

FIG. 1: Hardware in the Hall C Arc (not to scale).

Several harps and superharps are used to measure the beam profile. A harp consists of a frame with three wires, two vertical wires that measure the horizontal beam profile and one horizontal wire that measures the vertical beam profile. An Analog-to-Digital Converter (ADC) measures the signal on the wires and a position encoder

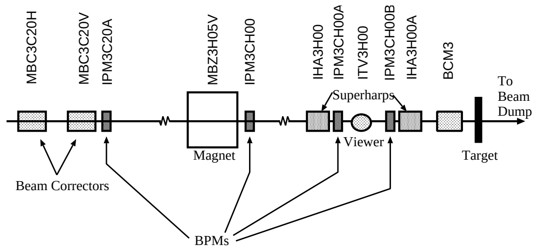

FIG. 2: Hardware in the Hall C beamline (not to scale).

measures the position of the ladder as they pass through the beam. Using the position information and the ADC measurements, the position and profile of the beam can be measured. Several harps are located throughout the accelerator for use in monitoring the position and shape of the beam. The superharps are essentially the same as the harps, but they have been more accurately fiducialized and surveyed for absolute position measurements. The superharps are primarily used for the beam energy measurement in the Hall C arc. Three superharps are located on aligned granite tables at the beginning, middle, and end of the Hall C Arc. Using the positions measured by the three superharps along with the field maps of the arc bending magnets, the beam energy and emittance can be determined. The absolute beam energy can be determined with a fractional uncertainty of $\sim$2x10^{-4} with this method and beam energy changes below the 10^{-4} level can be measured. During data taking, beam energy changes are monitored with the BPMs in the arc. Details of the superharp construction and operation can be found in [2].

The position of the beam in Hall C was monitored using four beam position monitors (BPMs). The BPMs are described in detail in [3]. Each BPM is a cavity with four antennae rotated $\pm 45°$ from the horizontal and vertical. Each antenna picks up a signal from the fundamental frequency of the beam which is proportional to the distance from the antenna. The beam position is then the difference over the sum of the properly normalized signals from two antennae on opposite sides of the beam. Because the position is determined by the ratio of signals in the antennae, the position measurement is independent of beam current. Non-linearity in the electronics can introduce a small current dependence in the BPM readout. For the range of currents used during E89-008, this led to an uncertainty of <0.5 mm. From these four antennae, the relative (X, Y) position of the beam can be determined once the signals from the four antennae have been properly calibrated. The beam position from the BPMs in the arc were compared to the Arc C superharps in order to calibrate the absolute position for the BPMs. The final accuracy of the beam position measurement was $\pm$1.0 mm, with a relative position uncertainty of 0.1-0.2 mm (neglecting the current dependence). The BPMs in the Hall C beamline were not calibrated against the superharps. The calibration of the BPMs was fixed at a nominal value, and the beam was steered so that x=1.8 mm, y=-1.0 mm at the final BPM. This was determined to be the correct position at the target based on requiring mid-plane symmetry in both spectrometers. This position was verified by placing a sheet of Plexiglas at the front of the scattering chamber and determining the beam position at the target from the position of the darkened spot on the Plexiglas.

The beam energy was determined by measuring the path and the settings of the magnets in the Hall C Arc. This is not done continuously, because the focusing elements in the arc are turned off for the measurement and the superharps are used to scan the beam, following the procedure of [4]. Using the superharps to measure the beam position at the beginning, middle, and end of the arc, the beam is steered to insure that it follows the central trajectory, with all corrector magnets turned off. One of the dipoles in the arc (the 'golden' magnet) has been precisely field mapped. The other dipoles are assumed to have the same field map, normalized to the central field value. With the precise knowledge of the field, and the absolute beam positions measured with the superharps, the field integral is well known, and the beam energy can be determined with an uncertainty of $\delta p/p \approx 2 \times 10^{-4}$. Details of the energy measurement and associated uncertainties can be found in Ref. [5]. However, after the analysis of the Arc measurements was completed, it was discovered that the degaussing procedure used for the Arc dipoles during the measurements was not the same as was used when the dipole fields were measured. The energy measurements assume that the dipole is run to 300 Amps, and then reduced to the desired current value. During data taking, the dipoles were only being ramped up to 225 Amps. This led to a difference in residual field which led to an overestimate of the beam energy. This effect was measured and a correction was applied to the Arc measurements based on the updated saturation corrections.

The BPMs can be used to monitor the beam energy when data taking is in progress. However, because the position is not measured as well with the BPMs as the superharps, and because the corrector magnets are energized, total integrated field ($\int B \cdot dl$) is only known to $\sim$0.2%. This limits the accuracy of the the absolute beam energy measurement to $0.2 - 0.3$% of the beam energy. However, relative beam energy changes can be detected at the $2 - 3 \times 10^{-4}$ level.

Final Arc energy measurement checked using elastic electron-proton scattering (see below??), measurements of differential recoil in elastic scattering from Berrillium and Oxygen, and measurements of the diffractive minimum in Carbon [6, 7]. It was found that the true beam energy was slightly ($\sim$0.1%) below the nominal energy from the arc measurement.

The beam current in the hall was measured with three microwave cavity beam current monitors (BCMs). The current is monitored by using the beam to excite resonant modes in cylindrical wave guides (the BCMs). The wave guides contain wire loop antennas which couple to resonant modes. The signal is proportional to the beam current for all resonant modes. For certain modes (e.g. the TM_{010} mode), the signal is relatively insensitive to beam position. By choosing the size of the cavity, one can choose the frequency of the TM_{010} mode to be identical to the accelerator RF frequency in order to make the cavity sensitive to this mode. The material and length can be varied to vary the quality factor, the ratio of stored energy to dissipated power, weighted by the resonant frequency, $Q = \omega_0 W/P_d$. The cavities and associated readout electronics as used during E89-008 are described in [8, 9]

Temperature changes can cause expansion or contraction of the cavity. This leads to a modification of the frequency of the TM_{010} mode and a detuning of the cavity away from the desired 1497 MHz. Therefore, as the temperature changes, the measured power decreases, giving an error in the current measurement. If the temperature is within 2 degrees of the tuning temperature, then the temperature dependence in the current measurement is proportional to $2Q\alpha\Delta T$ for small temperature variations (α is the thermal coefficient of expansion of the cavity, $\Delta R = \alpha R \Delta T$). This leads to a modest temperature dependence, $\approx 0.25\%$/degree C. However, if the operating temperature is several degrees away from the tuning temperature ($\sim$5 degrees), then the temperature dependence is greatly increased, and the error in the measured current is $\approx$1.5%/degree. Because of this large temperature dependence, Q was reduced by a factor of three from its initial value in order to minimize the temperature variation of the output. During E89-008, the temperature of the cavity was stable $\pm$0.2 C, and was less than 1 C from the tuning temperature, giving negligible ($< 0.05\%$) errors on the current measurement. In addition, the temperature of the readout electronics can lead to an error in the charge measurement. For BCM1, the temperature coefficient was $\sim 0.3\%/\Delta T$, and for BCM2 (the primary BCM for E89-008) it was somewhat better. However, the electronics room temperature was stabilized to $\pm$0.5 C, leading to uncertainties below the 0.2% level.

In addition to the microwave cavity BCMs, there is also a parametric DC current transformer (Unser monitor [10]) that measures the beam current. The Unser monitor has a very stable and well measured gain, but can have large drifts in its offset. Therefore, it is not used in the experiment to determine the accumulated charge. However, because the gain is stable, the Unser monitor is used to calibrate the gain of the microwave cavity BCMs. Calibration runs were taken about once a day in which the beam was alternately turned off and on over 2 minute intervals. During the beam off periods, the offsets of the Unser and cavity monitors were measured. During the beam on periods, the gains of the cavity monitors were calibrated using the known gain and measured offset of the Unser. The Unser gain was calibrated before the experiment by sending a precisely measured current through a wire running along the inside of the cavity. Analysis of all of the calibration runs indicated that the offsets and gains were stable during the experiment. A single gain (and offset) was determined for each BCM and that value was used throughout the run. The charge measurement was stable to within 0.5%, and the overall uncertainty on the absolute charge for each run was 1.0%.

The electron beam generated at CEBAF is a high current beam, with a very small transverse size ($\lesssim$200 μm FWHM). There are two rastering system+ to increase the effective beam size in order to prevent damage to the target or the beam dump. The fast raster system, 25 meters upstream of the target, is designed to prevent damage to the solid targets and to prevent local boiling in the cryogenic targets. The slow raster system is situated just upstream of the target, and is designed to protect the beam dump. During E89-008, the increase of the beam size caused by multiple scattering in the scattering chamber exit window and the Helium bag was enough to prevent damage to the beam dump without the slow raster, so it was not in use during data taking. Currents above 80μA would have required the slow raster.

The fast raster system consists of two sets of steering magnets. The first set rasters the beam vertically, and the second rasters the beam horizontally. The current driving the magnets was varied sinusoidally, at 17.0 kHz in the vertical direction, and 24.2 kHz in the horizontal direction. The frequencies were chosen to be different so that the beam motion would not form a stable figure at the target. Instead, it moves over a square area, $\sim$ 2.4mm across. The rastering was sinusoidal, and so the average intensity was greatest around the edges of the box, since this is where the beam is moving most slowly. Because the beam spends $\sim$ 40% of the time in the outermost 0.1-0.2mm of the box, the peak power density decreases more slowly than the inverse of the area of the raster pattern. However, the reduction of power density was sufficient to prevent any significant density fluctuations due to local boiling in the cryogenic targets for the currents used in this experiment.

The beam has an emittance of $\sim$2x10^{-9} mrad at 1 GeV (4σ value), and a somewhat lower value at higher energies. The fractional energy spread is $<10^{-4}$. The relative beam energy can be measured with a fractional uncertainty of 10^{-4} and is known absolutely to better than 10^{-3}. The nominal beam energy is determined from the magnet settings in the arcs in the accelerator or in the Hall C Arc. The beam energy can be measured by fixing the magnet settings in the Hall C Arc and measuring the beam position at the beginning, middle, and end of the arc in order to accurately measure the path length of the beam through the arc. By measuring the path of the electron beam and using precise field maps of the arc magnets, the field integral, $\int B \cdot dl$, through the arc is measured accurately, and this is used to determine the energy of the beam. For one and two pass beams, the energies measured in the arcs have been checked by measuring the differential recoil from a composite target, and by measuring the diffractive minimum in scattering from the Carbon ground state [6, 7].

Targets

The scattering chamber has room for two target ladders, one for cryogenic targets and one for solid targets. In order to use the solid targets, the cryotarget ladder must be lifted fully out of the beam and rotated 90° so that it is out of the beam path and does not interfere with the spectrometer acceptances. Then, the solid target ladder can be inserted.

The standard cryotarget ladder contains three pairs of target cells with one short cell (∼4 cm) and one long cell (∼15 cm) per pair. For this experiment, we had cryogenic Hydrogen and Deuterium targets, a pair of empty cells, and a pair of dummy cells used for measuring background from the aluminum target cell walls. The dummy cells consisted of two flat aluminum targets, placed at the same positions as the endcaps of the cryotarget cells, but with walls approximately 10 times thicker. This allows us to measure the background from the aluminum endcaps very rapidly, and makes the total thickness (in radiation lengths) of the dummy cells close to that of the full targets. Figure 3 shows the arrangement of the full cryotarget ladder. A complete description of the Hall C cryogenic target system can be found in ref. [11].

The cryotarget system has three separate loops (for Hydrogen, Deuterium, and Helium targets), with each loop linked to a short and long target cell. Figure 4 shows a side view of the two cells attached to a single target loop. Each loop consists of a circulation fan, a target cell, heat exchangers and high and low powered heaters. The target can dissipate in excess of 200 Watts of power deposited by the electron beam. In the loops, an axial fan inside a heat exchanger forces the target liquid to flow through two cells on an aluminum cell block, which is connected to the heat exchanger. Extending from each cell block are two target cells. The cells are thin aluminum cylinders made from beer can stock, 6.731 cm in diameter, with 0.0178 cm walls. The target liquid flows through these cells. Inside of the large cells are smaller aluminum flasks. The entrance and exit endcaps are both curved slightly, which gives a thickness variation with beam position. The maximum target length change for a 2mm beam offset is less than 0.5% for the 4cm cells, and ∼0.1% for the 15cm cells. During the cryotarget running, the beam position was typically kept within 1-2mm of the nominal central position, with an average offset of less than 1mm (better than 0.5 mm for all of the elastic runs). The heat exchanger has approximately 3.5 grams/second of 4 K liquid helium flowing through the refrigerant side, and provides the cooling for the target liquids. The cold helium is provided by the CEBAF End Station Refrigerator, and is returned at ∼21.5 K. High power heaters are used to maintain a constant heat load for the system, so that the cooling power stays constant as the beam current changes. There is sufficient cooling power to keep the heaters running on multiple cells. This meant that two cells (one hydrogen and one deuterium) could be kept ready for beam, eliminating delays caused when one loop needs to be powered down before another can be powered up. Low power heaters maintain the cry-

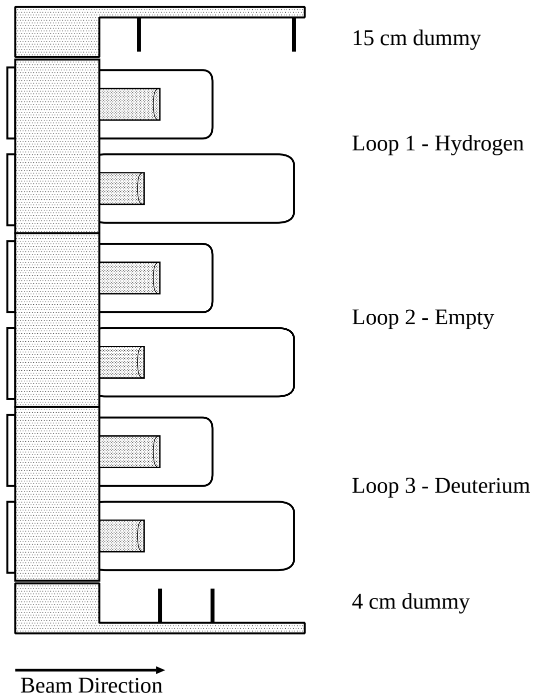

FIG. 3: Side view of the full cryotarget ladder.

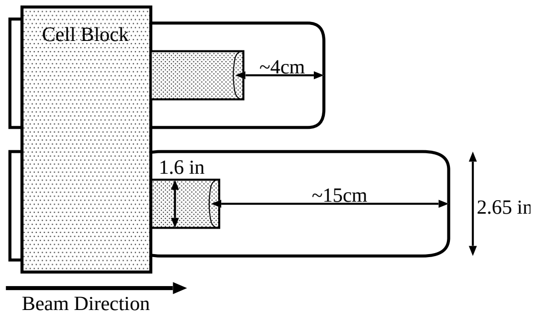

FIG. 4: Side view of one cryotarget loop.

otargets at their operating temperatures, and correct for small fluctuations in the beam current. The hydrogen target is operated at ∼0.2MPa (29 PSIA) pressure, and a temperature of 19K. In this state, the boiling temperature of hydrogen is 22.8K. The deuterium target is also operated in a subcooled fashion, at 22K. Table (??) lists the targets available in the cryotarget ladder for E89-008.

The loops are connected to a vertical lifting mechanism, which lifts the target ladder in order to place the desired cell in the beam. In addition, if the ladder is lifted to its highest position, the entire assembly can be rotated out of the beam by 90°. This allows the insertion of the solid target ladder and keeps the cryotarget cells and lifting mechanism clear of the spectrometer acceptances.

The temperature of the target cryogen is determined by a resistance measurement of two Lakeshore Cernox resistors for each loop, and the absolute temperature is measured to an accuracy of ∼100mK. Changes in the temperature are measured with 50mK accuracy. The density dependence on temperature is $\frac{1}{\rho}\frac{d\rho}{dT} = -1.25\%/K$, leading to an an uncertainty in density of less than 0.2%. Pressure changes have a much smaller effect on the density, $\frac{1}{\rho}\frac{d\rho}{dP} = 0.01\%/PSIA$, and were negligible in the final density uncertainty. The overall uncertainty in the calculation of the density (without beam) is ∼0.4%, mainly due to the uncertainty in the relative amounts of ortho and para hydrogen and the uncertainty in the equation of state. The length of the target cells has been corrected for thermal contraction (∼0.4% at the operating temperatures, and a 0.2% uncertainty is assumed for this correction. The uncertainties in the target thicknesses are summarized in table (??).

The density of the hydrogen is 0.07230(36) g/cm^3 at the operating temperature of 19K. The deuterium has a density of 0.1670(8) g/cm^3 at 22K. There is an additional current-dependent uncertainty in the density due to local target boiling. The analysis of the density dependence for runs up to August 1996 is described in [12]. Figure ?? shows the normalized yield (events per charge) for the 15 cm cryogenic deuterium target taken at the end of the experiment. During E89-008 data taking, the cryogenic targets were run at or below 55 μA, with a ±1.2 mm beam raster. For this current and raster size, there is no significant loss of target density. However, it was discovered after E89-008 that the beam tune into Hall C was not perfect, and that the unrastered beam size was larger than the desired 80-100μm [13]. In later runs, the tune was improved and the spot size reduced. Because the raster motion is sinusoidal in x and y, the beam spends a large fraction of the total time near the edges of the raster pattern. Therefore, the intrinsic size of the beam is still important when determining localized boiling. For the runs where the beam tune was improved, there was a density loss of ∼0.04%/mm/μA. This would correspond to a density loss of 1.8% at 55 μA with a 1.2mm raster. Our typical beam cross section was ∼3 times larger then for the improved tune, and was always $\gtrsim$2 times larger. While the beam spot may not have been small enough during E89-008 to have as large of an effect as seen with the improved beam tune, we cannot be sure that the spot size was completely stable during the run. This means that the effect of localized target boiling during data taking could have been larger or smaller than the effect measured during our test run. Therefore, we apply no correction to the density for target boiling, but assign an uncertainty of $0.013\%/mm/\mu$A (one third of the measured effect for the improved tune) to our target density, corresponding to a 0.6% uncertainty at 55 μA.

Target	Thickness (radiation lengths)	Thickness (mg/cm^2)	$\delta t/t$
C	2.09%	.8915(12)	0.5%
C	5.88%	2.510(10)	0.5%
Fe	1.54%	.2129(3)	1.0%
Fe	5.84%	.8034(11)	2.0%
Au	5.83%	.3768(6)	1.0%

TABLE I: Solid target thicknesses. All targets contained natural isotopic abundances.

Samples of the gases used to fill the targets were taken in order to measure the purity of the cryotargets. For the hydrogen gas used during E89-008, the target was 99.8% Hydrogen, and this was corrected for in the elastic analysis. The quantity of impurities (Nitrogen and Oxygen) was small enough that the background to the elastic measurement is negligible. For the deuterium, the gas was 99.6% Deuterium by number of nuclei, 99.2% by mass.

The solid target ladder is water cooled and has space for three thin targets and two thick targets (see figure 5). Two Carbon, two Iron, and one Gold target were used during the experiment (see table I). The target was cooled by flowing water through a copper tube that was attached to the back of the target. The tube was shaped so that water flowed past each target on all four sides. In addition to the physics targets, a Beryllium-Oxide (BeO) target was attached to the bottom of the ladder. It did not need to be water cooled because it was only used for beam tuning. At low currents, the beam spot is visible on the BeO target, and the the spot can be used to determine the position of the beam at the target. At higher current, the spot is visible on all of the targets. The ladder can be rotated so that the spectrometers can have a clear view of the target, without interference from the sides of the target frame. The targets were approximately 3.0 cm high and 4.2 cm wide, but when clamped into the frame, the area visible to the beam was 2.0 cm by 3.3 cm.

The beam has a roughly gaussian distribution, with a width of about 200μm, and so the size of the beam spot

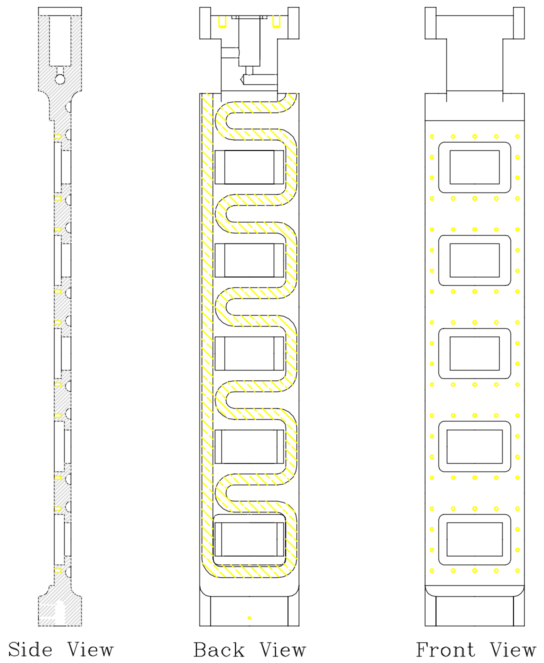

FIG. 5: The E89-008 solid target ladder. The bottom two slots are deep enough to hold thick (Carbon) targets. The BeO target (not shown) was hung from the bottom of the ladder. The shaded region on the back view shows the copper tubes that carry the cooling water.

on the target is determined by the raster size ($\pm$1.2mm horizontally and vertically for E89-008). The maximum beam position deviations were less than 4 mm, so there was always at least 5 mm clearance from the frame of the target ladder. This was sufficient to insure that there was no problem with background from the halo of the beam striking the frame. Since the beam profile monitors can only measure the profile of the beam where the intensity is relatively large, we took some test runs with the beam 1mm to 4mm away from the BeO target in order to look for non-gaussian tails to the beam profile. The test gave a crude measurement of the beam width which was consistent with the 200 μm measured by the harps. Any non-gaussian tail was below the 10^{-7} level at 1.5 mm.

The position of the target ladder was not fully surveyed after it was installed because it was replaced at the beginning of the run due to a vacuum leak. We know the position of the targets transverse to the beam to $\pm$2 mm, which is sufficient to insure that the beam was always well clear of the target frame. However, we do not know its exact location upstream or downstream of the central position. In addition to the overall uncertainty in position along the beam direction, there was some tilt to the ladder that caused this position to vary between different targets. From looking at the reconstructed target position (along the beam direction) for each target at identical kinematics, we estimate the offset to be $\approx$4.6 mm over the length of the target ladder, with the central target within 1mm of the nominal target position. Since almost all of the data was taken on the central three targets, we assume a position uncertainty of $\pm$1.3mm. In addition, if the beam is not on the exact center of the target, the angle of the target ladder will give a z-position offset. For a 20° target rotation (the maximum angle) and a 2 mm beam offset, this corresponds to a 0.7mm position offset. Combining the two effects, we assign an uncertainty of $\pm$1.5 mm in the z-position of the target.

For very forward angle data taking, this position uncertainty causes an uncertainty in distance from the target to the solid angle defining slit, which causes an error in the solid angle assumed in the analysis. The target-slit distance was $\sim$127 cm in both spectrometers, so a $\pm$1.5 mm position error gives a $0.12\%/\sin\theta$ error in the theta and phi acceptance, and a $0.25\%/\sin\theta$ error in the total solid angle and extracted cross section. Because the position of the beam varies on a similar scale ($\sim$1-2 mm), the large angle data will have a similar uncertainty in the target-slit distance, and we assign an uncertainty of 0.25% to the measured cross section, independent of target angle.

Because of the uncertainty in target position, and the fact that some of the data was taken with extended targets, we reconstructed events from the focal plane to the target with reconstruction matrix elements that were optimized for an extended target. Since this reconstruction set does not assume that you are at the central position, it will be insensitive to small position variations.

HMS Spectrometer (skip SOS?)

The standard detector package in Hall C at CEBAF consists of two magnetic spectrometers with highly flexible detector packages. The High Momentum Spectrometer has a large solid angle and momentum acceptance and is capable of analyzing high-momentum particles (up to 7.4 GeV/c). The HMS is a 25° vertical bend spectrometer, with superconducting magnets in a QQQD configuration. The magnets are supported on a common carriage that rotates around a rigidly mounted central bearing. The detector support frame is mounted on the same carriage as the magnets, thus fixing the detector frame with respect to the optical axis. The shielding hut surrounding the detector package is supported on a separate carriage. Figure 6 shows a side view of the HMS spectrometer and detector hut.

The magnets are cooled with 4K Liquid Helium provided by the CEBAF End Station Refrigerator (ESR). Under standard operating conditions, the HMS magnets require a flow of approximately 4 grams/second, running in parallel to the four magnets, to keep the magnet reservoir full and provide cooling for the current leads. The quadrupoles are cold Iron superconducting magnets. Soft

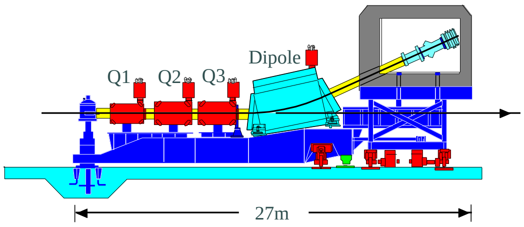

FIG. 6: Side view of the HMS.

	CDR	Final Design
Max. Central momentum (GeV/c)	6.0	7.4
Momentum bite	20%	20%
Momentum resolution [$\delta p/p$]	0.1%	0.04%
Solid angle (no collimator)	10 msr	8.1 msr
Acceptance - scattering angle		$\sim \pm 32mr$
Acceptance - out-of-plane angle		$\sim \pm 85mr$
Scattering angle resolution	0.1 mr	0.8 mr
Out-of-plane angle resolution	1.0 mr	1.0 mr
Extended target acceptance	20 cm	$\sim$ 10 cm
Vertex reconstruction accuracy	$\sim$ 1 mm	3 mm

TABLE II: HMS design goals and final performance. Resolutions include the effects of $\sim$200μm resolution per plane in the drift chambers and multiple scattering for 2.5 GeV/c electrons.

Iron around the superconducting coil enhances the central field and reduces stray fields. The quadrupoles are 'degaussed' by running the currents up to 120% of their 4 GeV/c values, and then lowering the currents to the desired values. The quadrupole current is provided by three Danfysik System 8000 power supplies. These supplies are water cooled and can provide up to 1250 Amps at 5 Volts. In addition to the quadrupole coils, each magnet has multipole windings. The correction coils are powered by three HP power supplies, capable of providing up to 100 Amps at 5 Volts. The multipole corrections to the quadrupoles were measured to be small when the magnet was mapped, and it was decided not to use the multipole correction coils for the standard point-to-point tune.

The HMS dipole is a superconducting magnet with a 25° bending angle for the central ray. The dipole has a bend radius of 12.06 m and a gap width of 42 cm. Its effective field length is 5.26 m (calculated assuming a perfect dipole, with a 25° bend and 12.06 m radius). It has been operated at up to 1350 Amps, corresponding to a central momentum of just over 4.4 GeV/c. The current is provided by a Danfysik System 8000 power supply capable of providing up to 3000 Amps at 10 Volts.

The HMS was operated in its standard tune: point-to-point in both the dispersive and non-dispersive direction. This tune provides a large momentum acceptance, solid angle, and extended target acceptance (see table II). In this tune, Q1 and Q3 focus in the dispersive direction and Q2 focuses in the transverse direction. The optical axis of each quadrupole was determined using the Cotton-Mouton method [14]. The optical axes were found to be different from the mechanical axes by up to 2mm, and all magnets were aligned with respect to the optical axis. When installed, the magnets were aligned to 0.2 mm, but move slightly when the spectrometer is rotated. The magnets move up to 1.0 mm, but the positions are reproducible up to 0.5 mm. The dipole field is monitored and regulated with an NMR probe. The quadrupole fields are regulated by monitoring the current in the magnets. The fields of dipole and quadrupoles are stable at the 10^{-4} level. Table II summarizes the design goals from the CEBAF Conceptual Design Report [15] and final performance of the HMS. The determination of the final optical tune and the analysis of the optics data to generate the reconstruction matrix elements is described in Ref. [7].

A slit system was installed in front of the first quadrupole, allowing remote insertion of various collimators. There are three HEAVYMET (machinable Tungsten with 10% CuNi; density=17 g/cm^3) collimators and one blank space in the slit box. The first collimator is a 3.175 cm thick sieve slit used for optics testing. It is an array of small holes (0.508 cm diameter) used to compare focal plane distributions to data with known angular distributions in order to study the optics of the spectrometer. Two holes are missing in the sieve slit in order to verify proper left-right and top-bottom reconstruction. The central hole is smaller than the others in order to measure the resolution of the angular reconstruction. The other two collimators are octagonal apertures designed to limit the solid angle acceptance of the HMS. Both are 6.35 cm thick and have flared holes that match the acceptance of the spectrometer. The large slit has a solid angle of $\sim$ 6.8 msr and was designed to keep losses within the spectrometer low for a point target (no loss in the magnetic elements for a $\pm$5% momentum bite, <2% for a momentum bite of $\pm$10%). The small slit was designed to give small losses in the spectrometer for an extended target ($\lesssim$0.1% for $\pm$10% with a 4 cm target, $\lesssim$0.1% for $\pm$5% with a 10cm target). For E89-008, all data was taken using the large octagonal collimator.

Detectors

The standard HMS and SOS detector packages are very similar. Each spectrometer contains two drift chambers, two sets of x-y hodoscopes, a gas Cerenkov detector, and a lead glass shower counter. The drift chambers provide tracking information, the hodoscopes are used to form the primary trigger, and the calorimeter

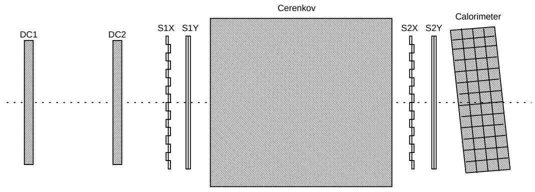

FIG. 7: Schematic diagram of the HMS detector hut.

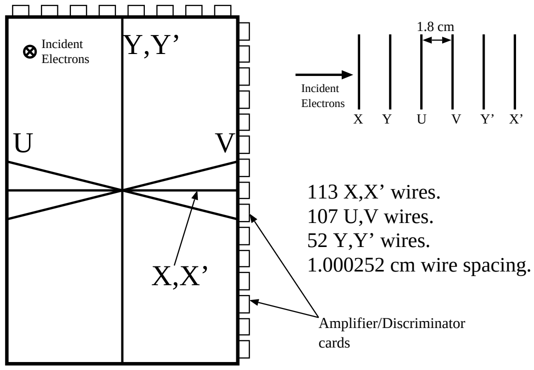

FIG. 8: Front view of the HMS drift chambers. The lines shown within the chamber indicate the region of coverage for the x, y, and v wire planes. The position of the readout cards is shown on the outside of the chamber.

and Cerenkov signals are used for particle identification (pion rejection) in the trigger and in the offline analysis. A schematic of the HMS detector package is shown in figure 7.

Drift chambers

The HMS drift chambers consists of six planes, two measuring x (the dispersive direction), two measuring y (the non-dispersive direction), and two that were rotated $\pm 15°$ from the x planes (the u and v planes). The planes were ordered x, y, u, v, y', x' as seen by incoming particles. The chambers had an active area of approximately 113 cm (x) by 52 cm (y) with a sense wire spacing of 1 cm. Figure 8 shows a front view of the HMS chambers. The planes were spaced 1.8 cm apart and the two drift chambers were separated by 81.2 cm. Each active plane contained alternating field and sense wires. The sense wires (anodes) are 25 μm diameter Gold-plated tungsten wire, and the field wires (cathodes) are 150 μm Gold-plated copper-beryllium wires. In between these planes were planes of guard wires. The sense wires detect the ionization from passing charge particles, and the field and guard wires are maintained at negative high voltage in order to isolate the sense wires and provide the electric field that attracts the ionized electrons to the sense wires. The voltage for the guard wires varied depending on its distance from the nearest sense wire, from -1800 V to -2500 V. This provided equipotential contours that were roughly circular. The distance between the wire and the track is determined by the drift time of the electrons.

When a charged particle passes through the chamber, the gas is ionized, and the liberated electrons are attracted to the nearest sense wire by the voltage differential maintained by the chamber. By detecting which wire sensed the particle, the position is measured with a 0.5 cm accuracy (half the wire spacing). The time required for the electrons to drift to the wire is measured by taking the time difference between the passage of the charged particle and the signal on the wire. This allows a much better determination of the position of the particle. By measuring the position with 6 planes, the x and y position of the particle and it's trajectory through the chamber can be measured. A complete description of the HMS drift chambers can be found in Ref. [16].

The HMS chambers are filled with an argon/ethane mixture (equal amounts by weight) along with $\sim$1% Isopropyl alcohol. The gas mixing system is located in a shed above the experimental hall and provides parallel gas streams to the two chambers. An MKS 647 menu driven 4-channel controller operates the system. The gas flow is controlled with MKS 1259c proportional mass flow control valves. The flow is monitored by temperature controlled alcohol bubblers on the gas lines going to the chambers.

The sense wires are read out in groups of 16, each connected to a LeCroy 2735DC or Nanometric N-277-L amplifier/discriminator card. The discriminator thresholds for all of the cards is provided by single external Acopian low voltage supply which was controlled remotely from the counting house. The threshold voltage supply in the counting house was set between 5.0 and 5.5 Volts during the experiment, but there is a 1-2 Volt drop between the source and the chambers downstairs. The signals from the discriminator cards are carried on twisted pair ribbon cable and go to LeCroy 1877 multi-hit Time-to-Digital Converters (TDCs) located in the back of the detector hut. The trigger is formed in the counting house and a TDC stop signal is sent back to the hall. The TDCs can store all hits (up to 16 per wire) that came within the last 32 μs. Because the total time between a particle in the spectrometer and the trigger arriving at the TDC is less than 2 μs, we programmed the TDCs to read out events within a window of $\sim$4μs. The drift chamber TDCs measure the time that the wire detected the electrons created by the ionization of the chamber gas, relative to the time of the trigger.

Using the hodoscope TDCs to dettermine the time that the particle passed through the focal plane (again, relative to the trigger), we can determine the time it took for the electrons created by the ionizing particle to 'drift' to the wire. This drift time is converted into a drift distance which is then added to the wire position in order to get the position of the event. The conversion from drift time to drift distance is determined by comparing the distribution of drift times in the chamber with expected position distributions of events within a cell. Combining the hits in all six planes allows us to determine on which side of each wire the particle passed. We make a small angle approximation and assume that for planes that measure the same coordinate, but which are offset by 1/2 cell, the particle passed between the two wires that fired. For events where only one of the two matching planes fired and for unmatched planes (u and v), we look through all left-right combinations and take the track with the minimum χ^2. The final position resolution is approximately 280 μm per plane.

Hodoscopes

The HMS and SOS each have two pairs of (x-y) hodoscopes, identical except for size of the elements. Each hodoscope plane is constructed of 9 to 16 elements. The hodoscope elements are long narrow strips of BC404 scintillator with UVT lucite light guides and Philips XP2282B phototubes on both ends. When charged particles pass through the paddles, they ionize the atoms in the material. The liberated electrons excite molecular levels in the scintillator, which emit light when they decay. The light is detected by Photomultiplier Tubes (PMTs) at the ends of the paddles. The light emitted along the length of the paddles will be detected by PMTs after the light has had time to traverse the length of the paddle. Light that is not emitted along the length of the paddle, but which hits the surface of the scintillator at greater than the critical angle, will be completely reflected and will also reach the PMTs. The HMS scintillators are wrapped with one layer of aluminum foil and two layers of Tedlar to make them light tight with a minimum amount of additional material. The SOS scintillators are wrapped with 1 layer of Aluminized mylar and 1 layer of Tedlar. The scintillators have approximately 0.5 cm of overlap between the paddles in order to avoid gaps between the elements. In the HMS, all of the scintillators are 1.0 cm thick and 8 cm wide. The x elements are 75.5 cm long, and the y elements are 120.5 cm long. The x planes have 16 elements each and the y planes have 10 elements each, giving each x-y pair an active area of 120.5 cm by 75.5 cm. The front and back planes are separated by approximately 220 cm. In the SOS, the front hodoscope pair is smaller than the back. The front x plane (S1X) has 9 elements, 36.5 cm x 7.5 cm x 1.0 cm and the front y plane (S1Y) has 9 elements that are 63.5 cm x 4.5 cm x 1.0 cm. The total active area of the front hodoscope is 63.5 cm x 36.5 cm. The rear hodoscope planes are larger versions of the front planes. The S2X plane is made up of 16 elements, each 36.5 cm x 7.5 cm x 1.0 cm and S2Y has 9 elements, 112.5 cm x 4.5 cm x 1.0 cm. Once again, the widths and lengths of the planes were matched so that the full area (112.5 cm x 36.5 cm) is active. The front and back planes in the SOS are separated by approximately 180 cm.

Each scintillator element is read out by PMTs at both ends. The 8-stage PMTs are connected to bases with zener stabilization in the first and last two stages. The anode output from the bases is sent to a patch panel in the detector hut through $\sim$30 feet of RG58 cable, and then goes upstairs to the counting house through $\sim$450 feet of RG8 cable. The signals are run through a splitter, giving two signals with 1/3 and 2/3 of the amplitude of the original input signal. The smaller signal is put through $\sim$400 ns of RG58 cable delay and then goes to the Analog-to-Digital Converters (ADCs) that measure the integral of the signal. The larger signal goes to PS7106 leading edge discriminators. One set of outputs from the discriminators goes to custom logic delay modules and then to Fastbus TDCs and VME scalers. The other set of outputs is sent to a LeCroy 4654 logic module. This module generates the OR of all tubes on one side of a given plane (e.g. S1X+). The outputs we use for the trigger logic are the AND of the sets of tubes on each side of a plane (e.g. S1X $\equiv$ [S1X+] & [S1X-]) as well as the OR of the front (and back) pairs of planes (e.g. S1 $\equiv$ [S1X] + [S1Y]).

The hodoscope PMTs were gain matched using a ^{60}Co gamma ray source at the center of each element. The tube voltages are set such that the Compton edge from the gamma rays gives a pulse height of 175 mV at the discriminator inputs in the electronics room. Timing calibrations of the scintillators was done using data taken during running. Corrections for the 'time walk' due to variations in pulse height and offsets between the individual elements are determined using an offline fitting procedure. The procedure used to determine the timing calibrations is described in detail in Ref. [6]. The final timing resolution achieved was $\sim$100 ps per plane for the HMS, and 80-100 ps per plane for the SOS. The increased timing resolution in the SOS is offset by the reduced lever arm for the time-of-flight measurement, due to the smaller separation of the hodoscope planes. This gives a measurement of the particle velocity, $\beta = v/c$, with an RMS resolution σ_β=0.018 at $\beta = 1$ for both spectrometers. The resolution improves as β decreases because the uncertainty in the time at the hodoscope planes is constant, but the flight time is larger. Therefore, the relative uncertainty is proportional to the inverse of the time of flight, which is proportional to β.

Gas Cerenkov

The HMS Čerenkov tank is cylindrical, with an inner diameter of ∼150 cm and a length of ∼165 cm. The effective length (before the mirrors) is approximately 120 cm. The tank is designed to run at gas pressures of up to 3 atmospheres, as well as running below atmospheric pressure. This allows the Čerenkov to be set up for e/π separation using nitrogen at ∼ 1 atmosphere of pressure, or π/p separation using 2-3 atmospheres of Freon-12. For this experiment, the tank was filled with 0.42 atmospheres of Perfluorobutane (C_4F_{10}, n=1.00143 at 1 atmosphere, 300K) giving an index of refraction of 1.0006. This gives a pion threshold of just over 4 GeV/c and electron threshold of ∼15 MeV/c). The expected yield was ∼11 photoelectrons, and the average measured signal from an electron was ∼ 10 photoelectrons. There were two mirrors at the back of the tank which reflected and focussed the Čerenkov light into two 5-inch Burle 8854 PMTs. In addition, the PMT front surfaces were coated with a wavelength shifting coating in order to improve the PMT quantum efficiency in the Ultraviolet wavelengths. The PMT has a UV window, but UV light is cut off below 200nm. The coating (paraterphenyl, 2400nm thick) fluoresces at 380nm when struck by light below 200nm. This allows some fraction of the 200nm light to be detected by the PMT. The tank has circular entrance and exit windows of 0.1016 cm Al (.27 g/cm^2). The combined thickness of the entrance window and C_4F_{10} gas is ∼0.7 g/cm^2. However, the main source of δ-ray production is the two hodoscope planes ∼20cm in front of the Čerenkov detector (∼2.3 g/cm^2 total thickness).

In both spectrometers, signals from the PMTs came up from the detector hut to the counting house through ∼10m of RG58 cable and ∼150m of RG8 cable. The signals are run through a 50-50 splitter and one set of outputs goes through 360ns of RG58 cable delay to a LeCroy 1881M ADC. The second set of outputs was summed in an Philips 740 linear fan-in module and put through a discriminator to give signals for the trigger logic as well as outputs for TDCs and scalers.

Because the signal from the Čerenkov was used in the trigger, the high voltages were adjusted so that the height of the signal from each tube was identical to within about 10% in the HMS and 20-30% in the SOS. Then a single threshold was applied to the sum of the analog signals from the PMTs. The final voltages varied between 2550 and 2750 Volts in the HMS, and 2650-2800 in the SOS. In the HMS, the mean number of photoelectrons is ≈10, and the trigger threshold corresponds to ∼1.5 photoelectrons. This means that the Čerenkov trigger signal is ≳99.9% efficient. While the mean signal in the SOS is larger than the HMS (∼12 photoelectrons), the difference in gain between the SOS PMTs means that the mean signal can be as low as 9 photoelectrons. The SOS trigger threshold corresponded to ∼1.7 photoelectrons, making the Čerenkov trigger signal ≳99.8% efficient.

The normalization of the signals from the gas Čerenkov counters were determined by measuring electrons in the spectrometer, and converting the ADC signal to the number of photoelectrons detected. A clean, high-statistics sample of detected electrons is chosen using the calorimeter to reject pions, and tracking to insure that the event points to the center of one of the mirrors. The number of photoelectrons detected should have a Poisson distribution. For each mirror-PMT combination, the mean and standard deviation of the ADC spectrum are determined, and the conversion from ADC channels to photoelectrons is determined by requiring that the mean value is equal to the square of the standard deviation.

The HMS Čerenkov detector has a larger active area than the calorimeter, and so all events within the acceptance of the calorimeter were far enough from the outer edges of the mirror that all of the Čerenkov light was captured. The mean HMS signal was 10 photoelectrons, but was reduced 10-20% at the edges of the mirrors. However, this was still a large enough signal to provide very efficient electron detection (≳ 99.2% everywhere for a 2 photoelectron cut) with better than 500:1 pion rejection for a cut at two photoelectrons. The majority of pions that have a signal above 2 photoelectrons are pions that produce a knock-on electron of high enough energy to emit Čerenkov light. At high momentum, the pion rejection is limited by the production of knock-on electrons above the electron Čerenkov threshold. This limits the gas Čerenkov pion rejection to ∼500:1. Figure 9 shows the HMS Čerenkov spectrum for runs with high and low pion to electron ratios, taken without the particle identification in the trigger. The final cut was placed at 2 photoelectrons in order to reject pions with a single photoelectron signal and maintain a high efficiency.

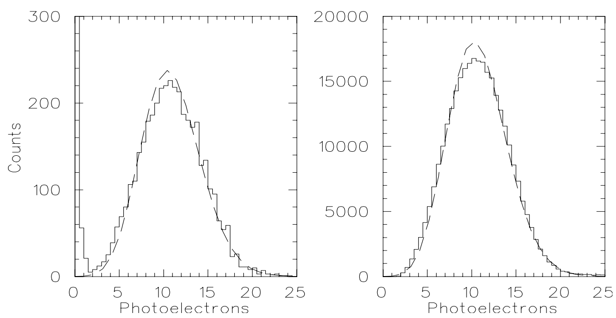

FIG. 9: HMS Čerenkov spectrum for runs with high (∼100:1) and low ratio of pions to electrons. Most of the pions appear at zero photoelectrons. The counts at ∼1 photoelectron are pions with single photoelectron noise.

Lead glass calorimeter

Both the HMS and SOS had lead glass calorimeters used to identify electrons and reject pions. The lead-glass counter is an electromagnetic calorimeter that detects the energy deposited when an electron enters the lead-glass. A high energy electron will radiate photons through Bremsstrahlung in the calorimeter, which will in turn generate positron-electron pairs. These pairs will also radiate photons, and a shower of particles (photons, electrons, and positrons) will be generated. The PMTs on the lead-glass blocks detect the Čerenkov light given off by the charged particles. This signal is proportional to the total track length of charged particles in the calorimeter (for particles above the Čerenkov threshold) which is in turn proportional to the energy of the initial electron. Electrons, positrons, and photons will deposit their entire energy in the calorimeter giving a detected energy fraction of one. The energy fraction is the ratio of energy detected in the calorimeter to particle momentum (determined from the tracking for charged particles). Hadrons (mostly negative pions for E89-008) usually deposit a constant energy per layer, due to ionization and direct Čerenkov light. The pions typically deposit ~300 MeV in the calorimeter. Therefore, pions will show up as a peak in the energy fraction distribution at $E_{cal}/p = 0.3GeV/p$. A negative pion can have a charge-exchange reaction in the calorimeter and produce a neutral pion with a significant fraction of the initial pion's momentum. In this case, the pion will decay into two photons, and the full energy of the neutral pion can be deposited in the calorimeter. This leads to a high-energy tail for pions that goes up to an energy fraction of one. However, the neutral pion will not have the full momentum of the initial charged pion, and unless the charge-exchange reaction and pion decay occur in the front of the calorimeter, some of the particles in the shower will leak out the back of the calorimeter, and their energy will not be measured. At momenta significantly above 300 MeV/c, this high energy tail is the dominant contribution to pion misidentification.

The calorimeters were of identical design and construction except for their total size. Each calorimeter is a stack of 10 cm x 10 cm x 70 cm blocks of TF1 lead glass, with a PMT on one end. The blocks are stacked transversely to the incoming particles, four layers deep and 13 blocks high in the HMS (11 in the SOS), for a total of 52 (44) modules and an active area of 130(110) cm x 70 cm. The calorimeters are rotated 5° from the optical axis in order to avoid loss through the cracks between the modules (see figure 7). TF1 lead glass has a density of 3.86 g/cm^3 and a radiation length of 2.54 cm, making the entire calorimeter ~16 radiation lengths total thickness. Each block is wrapped with one layer of Aluminized mylar (25 μm) and 2 layers of Tedlar PVF film (38 μm each) to increase reflection and make the modules light tight. Each module was read out from one end by an 8-stage Philips XP3462B 3-inch phototube. The gains of the phototubes and attenuation of the blocks were measured and the best blocks were paired up with the worst phototubes to minimize the signal variation over the calorimeter. The attenuation length varied between 50 and 100cm (at $\lambda = 400$nm). The operating voltages were set to match the gain of the individual modules. The outputs were gain matched to within 20%, and the final differences were corrected in software. A detailed description of the calorimeter design and performance will be published elsewhere [17]. In addition, each block had a light guide input for use with a laser gain monitoring system. The gain monitoring system was in place for the calorimeter at the time of the run, but was not used because it had not been sufficiently tested at that time.

The signals from the phototubes are taken from the detector hut to the electronics room through ~30 feet of RG58 and ~450 feet of RG8 coaxial cable. The signal is then run through a 50-50 splitter. One set of outputs is sent through 400 ns of RG58 delay cable to a LeCroy 1881M ADC and the other set is sent to Philips 740 linear fan-in modules to be summed. The sum in the first layer (PRSUM) and the sum in the entire calorimeter (SHSUM) are discriminated to give three logic signals for the trigger. PRHI and PRLO are high and low thresholds on the energy in the first layer, and SHLO is a cut on the total energy in the calorimeter. Also, groups of four modules are summed, sent through discriminators, and scaled in order to look for dead or noisy tubes.

The raw ADC values are corrected in two ways. First, the signal is corrected for attenuation through the block to remove the signal dependence on distance from the PMT. Then, each channel has a gain correction factor applied, determined by fitting a value for each block in order to match the sum of the blocks to the energy as determined from the momentum reconstruction. Figure 10 shows the calorimeter spectrum for two runs (low and high pion to electron ratio), after a Čerenkov cut has been applied. For the SOS, the calorimeter is identical, and the resolution and pion rejection are nearly identical to the HMS. Figure 11 shows the resolution as a function of momentum for both calorimeters. The curves shown are fits to the resolution, giving a $6.5\%/\sqrt{E}$ for the HMS, and $5.6\%/\sqrt{E}$ for the SOS. While the calorimeters and readout electronics are identical in the two spectrometers, the HMS had additional noise at the ADC which worsened the average resolution.

Trigger and data acquisition

The HMS and SOS have separate trigger electronics, which provide independent triggers for events in each spectrometer. There are two different types of single spectrometer triggers when running in electron detection

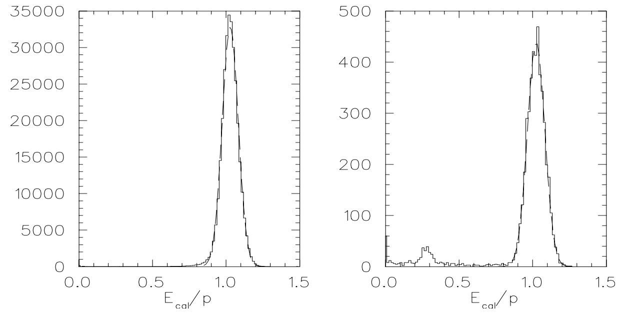

FIG. 10: HMS shower counter spectrum (Energy measured by calorimeter divided by the particle momentum) after a cut on the Čerenkov signal has been applied. The dashed lines are gaussian fits. The left figure is for a run with a low pion to electron ratio (Fe data at 30°, p=2.06 GeV/c). The right is for a high pion to electron run (Fe at 30°, p=1.11 GeV/c) and shows a clear pion peak, even after the Čerenkov cut. The pions deposit approximately 250 to 300 MeV of energy in the calorimeter, so the pion signal appears at ∼.3 GeV/p_π (∼0.27). The pion peak is wider than the electron peak because the energy deposition is roughly constant, so the energy fraction is widened by the size of the momentum acceptance (∼20%).

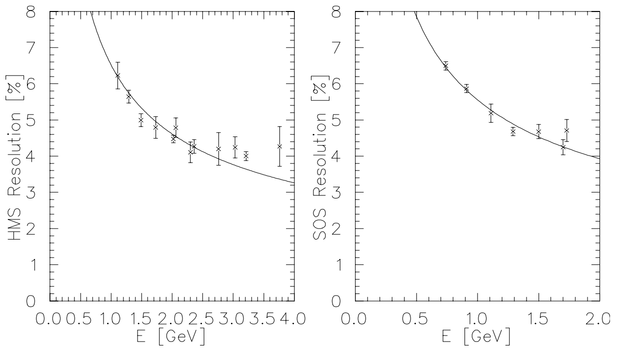

FIG. 11: HMS and SOS shower counter resolution vs energy. The fits are $6.5\%/\sqrt{E}$ for the HMS and $5.5\%/\sqrt{E}$ for the SOS.

mode. ELREAL is the electron trigger, and requires scintillator hits plus user defined particle identification signals. There is also a pion trigger (PION), which requires just scintillators (and can be vetoed by the Čerenkov if desired), and can be dynamically prescaled independently of the electron triggers. The trigger electronics in Hall C provide single spectrometer triggers and coincidence triggers. The Trigger Supervisor (TS) is programmed to accept, reject, or prescale each of the different trigger types, depending on the needs of the experiment. For E89-008, only singles electron triggers were taken. Pion singles triggers were blocked, and coincidence triggers were prescaled away. However, a coincidence trigger means that there was a singles trigger in each spectrometer. While the COIN triggers, generated in the 8LM, are prescaled away at the TS, if the HMS and SOS triggers come within 7 ns (the TS trigger latching time) of each other, the TS will treat the event as if it were a coincidence trigger, even though the COIN trigger was ignored. Even though the coincidence event contains good HMS and SOS data, the timing for the ADC gates and TDC stops is sometimes incorrect for the coincidences triggers, since the timing was not set for taking coincidence data. The rate of coincidences was low enough that the inefficiency caused by missing these triggers was between 10^{-7} to 10^{-3}, except for a handful of runs. For these runs with extremely high SOS trigger rates, the SOS triggers were prescaled at 100:1 or greater. Because the prescaling occurs before the triggers are latched, the rate of SOS triggers that can cause a false coincidence is also reduced by a factor of 100 or more. After taking the prescaling of the SOS triggers into account, the inefficiency caused by this accidental identification of singles triggers as coincidence events is always negligible ($<0.1\%$).

The HMS and SOS have separate trigger electronics, which provide independent triggers for events in each spectrometer. There are two different types of single spectrometer triggers when running in electron detection mode. ELREAL is the electron trigger, and requires scintillator hits plus user defined particle identification signals. There is also a pion trigger (PION), which requires just scintillators (and can be vetoed by the Čerenkov if desired), and can be dynamically prescaled independently of the electron triggers. The trigger electronics in Hall C provide single spectrometer triggers and coincidence triggers. The Trigger Supervisor (TS) is programmed to accept, reject, or prescale each of the different trigger types, depending on the needs of the experiment. For E89-008, only singles electron triggers were taken. Pion singles triggers were blocked, and coincidence triggers were prescaled away. However, a coincidence trigger means that there was a singles trigger in each spectrometer. While the COIN triggers, generated in the 8LM, are prescaled away at the TS, if the HMS and SOS triggers come within 7 ns (the TS trigger latching time) of each other, the TS will treat the event as if it were a coincidence trigger, even though the COIN trigger was ignored. Even though the coincidence event contains good HMS and SOS data, the timing for the ADC gates and TDC stops is sometimes incorrect for the coincidences triggers, since the timing was not set for taking coincidence data. The rate of coincidences was low enough that the inefficiency caused by missing these triggers was between 10^{-7} to 10^{-3}, except for a handful of runs. For these runs with extremely high SOS trigger rates, the SOS triggers were prescaled at 100:1 or greater. Because the prescaling occurs before the triggers are latched, the rate of SOS triggers that can cause a false coincidence is also reduced by a factor of 100 or more. After taking the prescaling of the SOS triggers

into account, the inefficiency caused by this accidental identification of singles triggers as coincidence events is always negligible (<0.1%).

The first part of the trigger comes from the hodoscope signals which fire when a charged particle passes through the spectrometer. The gas Čerenkov counter and calorimeter signals are used to determine if the event is an electron or a pion. Triggers with no Čerenkov signal were labeled as pions. If an event had either a Čerenkov signal or a large shower counter signal, it was counted as an electron. This was highly efficient for electrons, since either detector can identify the event as an electron, but limited the hardware pion rejection. Since the Čerenkov has a large pion rejection ($\sim$500:1 HMS, $\sim$150:1 SOS), the pion rejection in hardware was usually limited by the rejection of the shower counter. Because of this, the thresholds for the calorimeter were set as high as possible, while still having a high ($>90\%$) electron efficiency. This gave a final online pion rejection of $\sim$20:1 for the HMS at the lowest momentum, and better than 100:1 as the momentum increased. Because the SOS was operated at lower momenta, the online rejection was as low as 10:1 for some kinematics. In order to improve the pion rejection, the 4th layer of the calorimeter was removed from the hardware sum. For momenta below $\sim$1.5 GeV, the energy from electrons is contained almost entirely in the first three layers and only pions deposit energy in the last layer. By removing this layer, we reduce the pion signal without losing any signal from the electrons. After the raw spectrometer trigger was formed (the 'pretrigger'), additional logic provided the final trigger for the Trigger Supervisor (TS) which generates the necessary ADC gates and TDC stop and start signals for the event.

Each hodoscope plane consisted of 9-16 individual elements, each of which was read out on both sides (the 'positive' and 'negative' ends). The signals from the tubes were discriminated and the tubes from the positive (and negative) ends were ORed together to give the signals S1X+, S1X-, ... A hit in a given plane was defined as a coincidence of a hit in one of the positive tubes and a hit in one of the negative tubes, (e.g. S1X $\equiv$ [S1X+] + [S1X-]). This definition does not require both tubes to be on the same scintillator, but requires much less electronics and does not cause any significant amount of random signals. Two scintillator triggers are then formed. 'STOF' was defined as the coincidence of one of the front planes and one of the back planes, which is the minimum hit requirement for a good time of flight measurement in the scintillators. 'SCIN' required that 3 of the 4 planes fired, and provided a tighter scintillator trigger. Since any combination of three planes will include one S1 plane and one S2 plane, every time SCIN was true, STOF was also true. Occasionally, a PMT would be lost due to a failure in the high-voltage for that channel or a problem with the PMT base. Each plane has an average efficiency of $\sim$99.9%, and so even when a PMT is lost, the trigger efficiency for events passing through that scintillator is still $\sim$99.7%. There were no cases where multiple PMT signals were missing in the trigger electronics. The only cases in which two or more PMT signals were lost involved either a problem in the electronics (after the trigger signal is formed) or else a failure in the PMT base that affected the ADC, but not the discriminated signals used in the trigger. This problem occurred when an anode solder connection broke in such a way that it become AC-coupled. This meant that no charge could flow across the connection, but that the signal could still be large. This gave a distorted pulse shape with a very narrow negative voltage spike followed by a narrow positive spike. However, the signal is still able to fire the discriminator, which generates the trigger and TDC signals. Therefore, the ADC signals were lost, but there was no significant inefficiency in the hodoscope trigger. Finally, for an event with both a good Čerenkov and shower counter signal, the ELLO trigger (see next section) will fire on the STOF hodoscope condition. Therefore, even if a plane was missing completely, the trigger could still fire with one front and one back hodoscope hit as long as both particle identification signals were present. The necessary Čerenkov and shower counter signal were both $\gtrsim 98\%$ efficient for all except the lowest momentum settings in the SOS, so in general the STOF signal (two hodoscope planes) was sufficient to generate a trigger.

Because of the high pion to electron ratio for some of the kinematic settings, we require the event to pass some particle identification cuts before generating a trigger. In order to have a high efficiency for detecting electrons, we accepted a trigger as an electron if either the Čerenkov fired or if the calorimeter had a large enough signal. This allows an extremely high electron efficiency even if one of the two detectors has a low efficiency, but limits our hardware pion rejection. The Čerenkov signal used in the trigger (CER) was true if the Čerenkov sum fired the discriminator, set at between one and two photoelectrons. The shower counter signals had discriminators on the total (hardware) sum of all blocks (SHSUM) and the sum of all blocks in the first layer (PRSUM). The total energy had one discriminator threshold (SHLO) and the pre-radiator had one discriminator with a high threshold (PRHI) and one with a low threshold (PRLO). The final electron trigger (ELREAL) was the OR of the two conditions. ELHI required a high calorimeter signal, but no Čerenkov signal, while ELLO required a Čerenkov signal, but not a calorimeter signal. ELHI was defined as the coincidence of SCIN, PRHI, and SHLO (a tight scintillator cut and both a high pre-radiator sum and total energy sum from the calorimeter). ELLO required the Čerenkov signal (by vetoing with the $\overline{CER}$ signal) as well as two of the following: a tight hodoscope condition (SCIN), a loose hodoscope condition (STOF), and a shower counter signal (PRLO). If the SCIN signal (3/4 hodoscope planes) is present for an event, there must also be a STOF signal (which requires one front plane and one back plane). This means that ELLO requires the Čerenkov and either the 'good' scintillator trigger (SCIN), or the minimum scintillator trigger (STOF) and the lower shower counter signal (PRLO).

In addition to providing the information used in the trigger, all of the intermediate signals are sent to scalers and TDCs. The TDCs are mainly used as latches, and tell which signals were present when the trigger was taken. This allows us to determine what kind of event formed the trigger. The scalers allow us to look at raw rates and look for certain types of electronics problems in the intermediate steps of trigger formation. We also use the scalers to measure computer and electronics dead time by comparing the number of triggers that were formed with the number that were accepted.

The maximum data taking rate is limited by the fastbus conversion and data readout time. In basic data acquisition mode, the total time to process an event is just under 1 ms. The time is broken up as follows: $\sim 95\mu$s for fastbus data conversion, $\sim$150 μs for the fastbus crate controller (FSCC) to read the data from the ADC modules into it's FIFO, and $\sim$650 μs for the FSCC to take the data from it's FIFO into memory and send it out over ethernet. This limits data acquisition to $\sim$1.1 kHz, but gives large computer dead times even at lower rates. Several improvements have been made to improve the data rate and decrease dead time. First, because the FSCC is inefficient at sending data over the ethernet, the readout of the fastbus data was modified so that when running in 'parallel' mode, the data was read out from the FSCC FIFO through a VME CPU. This reduced the processing time to $\sim$95 μs for fastbus conversion, and $\sim$400 μs for the data readout. In addition, optimization of the fastbus readout of the TDCs and ADCs reduced the fastbus readout time to 300 μs, giving a total time to process the event of $\sim$400 μs and a trigger rate limit of $>$ 2 kHz when running in parallel mode. However, the dead time is still large for rates well below this limit. The fraction of events missed is equal to the fraction of the time the computer is busy which equals the rate of events taken over the maximum rate (2-2.5 kHz), so even at 500 Hz the computer dead time is $\sim$20-25%. In addition to the improvements gained by running in parallel mode, the fastbus modules we use allow buffering of 8 events. This allows the trigger supervisor to accept new triggers as soon as the fastbus conversion is done, rather than waiting for the full conversion and readout time. This means that the dead time is roughly one quarter of what it is in non-buffered mode. The total processing time for an event is still $\sim 400\mu$s, so the total event rate limit does not improve, but fewer events will be missed for rates lower than the maximum.

For a handful of runs, there was a problem in the synchronization between the drift chamber TDCs and the hodoscope, Čerenkov, and calorimeter ADCs and TDCs. This happened when excessive 'noise' caused extra triggers to appear at the fastbus crate containing the drift chambers. In buffered mode, each crate digitizes and stores up to 8 events. If an extra trigger comes to the crate, it will perform an extra read. Because the individual TDC stop signals are not present, it will tag the data for this read as being incomplete. However, it is stored in the buffer and not read out until a real trigger causes the event builder to read the data from each crate. The TDC readout caused by the bad trigger will take the place of the TDC readout caused by the current trigger. After this point, the drift chamber events are always being read out with data from the previous event, or data from earlier events, if the noise caused multiple false triggers. Because this affected only a very small part of the data, the runs where there was a synchronization problem were discarded.

DATA REDUCTION

The analysis of the raw data files was done using the standard Hall C event reconstruction software. The event reconstruction code reads the raw events, decodes the detector hits, and generates tracks and particle identification information for each event. In addition, it keeps track of the hardware scalars and generates software scalers for the run. The detector calibrations were done using separate code and the results were taken as input to the event reconstruction software. The data is output in three forms. Report files contain the hardware and software scalars, as well as calculated detector efficiencies. PAW [18] HBOOK files contain the standard set of histograms which can be used to check detector performance and monitor the hardware during a run. PAW Ntuple files contain the event by event information, and are the main output used in the final physics analysis. Histograms and Ntuples are generated using the CERN HBOOK libraries. Input parameters, software scalars, histograms and tests are handled using the CEBAF Test Package (CTP) [19]. After the tracking, efficiency, and particle identification information is generated by the analysis package, The physics analysis is done using separate stand-alone Fortran and PHYSICA [20] code.

Before the events are reconstructed, the TDCs that record the intermediate trigger signals are examined, and events are rejected unless they contain both a Čerenkov signal (CER) and a shower counter signal (PRLO, SHLO, or PRHI). See section ?? for the definition of the trigger signals. This effectively modifies the online trigger from an OR of the two detectors to an AND. The shower counter signal required in the calorimeter based trigger (ELHI) sometimes has an electron efficiency as low as 90% (at the lowest momentum settings). However, it requires that the total energy be above a fixed threshold (SHLO) and that the energy in the first layer be above a fixed threshold (PRHI). It is this 'high' threshold on the first layer energy that causes most of the inefficiency for electrons in the ELHI trigger. By requiring only one of the three signals (SHLO, PRHI, or PRLO, which is a lower threshold on the pre-radiator energy), the efficiency becomes very high (>99%).

This offline 'trigger modification' is done for two reasons. First, in order to insure that the trigger efficiency would be high even if one of the detectors was not working well, the thresholds were set relatively low. This limited the online pion rejection. By modifying the trigger requirements before reconstructing the event, we can reduce the size of our data set by a factor of two. This significantly reduces the time required to analyze the data set.

In addition to reducing the data set, this cut has an additional benefit in the SOS. In the SOS Čerenkov signal, there was significant noise in the ADC readout which limits the offline pion rejection. Because the noise was in the ADC, the trigger signal was not affected, and the pion rejection is not reduced. Therefore, we use a combination of the trigger signal (a ∼1.7 photoelectron on the clean signal) and a cut on the Čerenkov ADC (3.3 photoelectrons on the noisy signal). The online cut rejects pions at ∼250:1, and the offline cut rejects pions at ∼170:1. The combined efficiency is estimated to be between 300:1 and 380:1, and we assume 300:1 when estimating the pion contamination. The worst case pion contamination after the final particle identification cuts is ∼3%, and only occurs for the largest angle data, where the statistical uncertainties and systematic uncertainties due to other backgrounds are their largest (>10%).

Tracking

The trajectory of the event at the focal plane is measured with two drift chambers, each with six planes. The position of the track as it passes through a plane is determined by starting at the position of the wire that detects the particle, and adding the distance of closest approach between the track and the wire. This distance is determined by measuring the time difference between the time that the particle passed through the focal plane (as determined by the scintillators) and the time at which the wire detected the particle passing. It is assumed that the particle is moving nearly perpendicular to the plane, and that the point of closest approach is in the plane of the drift chamber. In addition, small corrections are applied for the time required for the signal to propagate along the wire and differences in cable lengths from between the chamber and the Time-to-Digital Converters (TDCs).

The drift chamber hits are used to identify clusters of hits (space points) in the front chamber. The drift time is determined from the drift chamber TDC values and the hodoscope start time. For each space point, a 'stub' is fit. This is a track determined using just the hits in the first chamber. For each wire in a space point, the particle could have gone past the wire on the left or the right. The left-right determination can be made by fitting a stub through the space point for each left-right combination (2^6 stubs per space point) and choosing the stub with the lowest χ^2. However, in order to improve the speed of the tracking algorithm, we use a small angle approximation for the y and y' planes in the HMS (High Momentum Spectrometer), and all of the planes in the SOS (Short Orbit Spectrometer). In the y and $y\prime$ planes (or any two parallel planes), the wires within each plane are separated by 1 cm, but the parallel planes are offset 0.5 cm. If you have a hit in both planes, you can choose the left-right combination that makes the particle go between the wires. For planes that are close together and incoming particles that are nearly perpendicular to the drift chambers, this is a very good approximation. Therefore, a space point with one hit in each of the six planes has only 2^4 possible left-right combinations in the HMS (since the left-right determination for the y and y' planes is made using the small angle assumption), and no left-right ambiguity in the SOS. Approximately 3% of the time, a plane is missing and the left-right determination for its partner plane is made by looping through all 2^5 possible left-right choices that are not determined by the small angle assumption and choosing the stub with the lowest χ^2. After all space points have been found in the front chamber, and stubs fit for each one, the code finds space points and stubs for the second chamber. Finally, for each combination of stubs in the front and back chambers, a full track is fit if the two stubs were consistent (*i.e.* the slopes of the stubs must be consistent, and they must point to each other). Each of these tracks is recorded along with the χ^2 of the fit.

In bench tests, the HMS and SOS chambers had resolutions of $\lesssim$150 μm per plane. However, in the final two-chamber tracking, there are additional resolution effects coming from the resolution of the start time from the hodoscopes, wire position offsets or wire sagging, and errors in the drift chamber position or angles. By comparing the position measurements of the individual planes and comparing them to the final fitted track, we obtain a tracking resolution of ∼280 μm per plane in the HMS, and ∼180 μm in the SOS. For the HMS, each chamber has two planes that measure y, and four planes that primarily measure x. This gives a position resolution in $x(y)$ of $\sim 140\mu m$ $(200\mu m)$ and an angular resolution of 0.24 mr for $\frac{dx}{dz}$ and 0.34 mr for $\frac{dy}{dz}$. The resolution on the momentum and reconstructed angles is given in table II and is a combination of the drift chamber resolution and the error in the track reconstruction. The resolution on the reconstructed quantities is worse at lower electron energy as multiple scattering in the target, scattering chamber, and magnet entrance window. At low momentum spectrometer momentum settings, the multiple scattering dominates the resolution. For the SOS, there are six measurements per plane, with equal x and y information, giving a position resolution of $\sim$ 105 μm and an angular resolution of ~0.30 mr. Note that while the position resolution is better in the SOS, the angular resolution is comparable in the two spectrometer because the SOS chambers are separated by 49.5 cm, while the HMS chambers are separated by 81.2 cm.

Before a fitted track is accepted as a good track, cuts are applied to reject bad fits caused by space points with missing wires or with noise hits. The track is used to determine which hodoscope elements and which calorimeter blocks the particle passed through, and cuts are applied on the particle velocity, the signal in the calorimeter, and the measured dE/dx in the hodoscope, all as measured using the detector elements that lie on the fitted track. In addition, a hard cut is placed on the χ^2 of the fit for the track. If multiple tracks pass these cuts, then the track with the best χ^2 is selected as the final track. In our analysis, the hard cuts were opened up, allowing all good tracks to pass, and the best track was selected using χ^2. Typically, multiple tracks are found in 1-2% of events (3% worst case). Most of these tracks come from finding space points with slightly different sets of wires. Typically, 5 of the wires occur on both tracks, and only the sixth differs (or is missing). In these cases, the tracks are nearly identical, and the choice of the lower χ^2 is effective in selecting the appropriate track when one of the hits is a 'random' hit. The fraction of events with true multiple particles in the spectrometer is typically less than 0.1%, and is always less than 1%.

The event reconstruction code generates information for the tracks at the focal plane, and reconstructed tracks at the target. The focal plane quantities are the x and y positions and slopes of the track at the focal plane (x_{fp}, y_{fp}, x'_{fp},and y'_{fp}), in the coordinate system where $\hat{z}$ is parallel to the central ray, $\hat{x}$ points downwards, and $\hat{y}$ points left when viewing the spectrometer from the target. The reconstructed values are δ, y_{tar}, x'_{tar}, and y'_{tar}, where $\delta = (p_{recon} - p_0)/p_0$, with p_0 equal to the spectrometer central momentum, y_{tar} is the horizontal position at the target plane (perpendicular to the spectrometer central ray), and y'_{tar} and x'_{tar} are the tangents of the in-plane and out-of-plane scattering angles, with $\hat{x}$ pointing downwards, $\hat{y}$ pointing left, and $\hat{z}$ pointing towards the spectrometer. Note that while x'_{tar} and y'_{tar} are the slopes of the tracks ($x'_{tar} = \frac{dx_{tar}}{dz_{tar}}$), they are often referred to as the out-of-plane and in-plane scattering angles, and given the units of radians (or milliradians).

Cuts are applied to the reconstructed target quantities in order to eliminate events that are outside of the spectrometer acceptance but which end up in the detectors after multiple scattering in the magnets or shielding.

HMS	SOS
$\lvert x'_{tar}\rvert < 90mr$	$\lvert x'_{tar}\rvert < 40mr$
$\lvert y'_{tar}\rvert < 55mr$	$\lvert y'_{tar}\rvert < 80mr$
$\lvert y_{tar}\rvert < 7cm +$ (target length)/2	$\lvert y_{tar}\rvert < 2cm +$ (target length)/2
$\lvert\delta\rvert < 14\%$	$-16\% < \delta < 12\%$

TABLE III: Cuts on reconstructed tracks.

The cuts are kept loose enough to avoid losing any real events due to the finite tracking resolution caused by the drift chamber position resolution and by multiple scattering in the target and the entrance and exit windows in the spectrometer. In addition, we apply a cut on the reconstructed momentum. This cut is applied so that we analyze data in the momentum region where we have good matrix elements for reconstructing the track to the target. The tracking cuts applied are listed in table III.

In the HMS, the x'_{tar},y'_{tar}, and y_{tar} cuts typically rejected ~1.0% of the total tracked events, and never more than 2%. Of these events, $80 - 90\%$ come from events that are outside of the acceptance, but scatter back into the detectors at the dipole exit or in the vacuum pipe afterwards. Therefore, the cuts are $> 99.5\%$ efficient for good events. Of the events that scatter inside of the spectrometer and end up in the detector stack, $\gtrsim 90\%$ are rejected in the tracking cuts or with the background cuts (described later). More than half are rejected by the tracking cuts, and therefore the worst case loss to tracking cuts of 2% indicates a worst case of scraping events of 4%. With $\gtrsim 90\%$ rejection, this leaves a possible contamination of 0.4%. No correction is made to the cross section, but a $\pm 0.5\%$ uncertainty is assumed due to possible inefficiency in the cuts or contamination due to scraping events.

Particle identification

For many of the e89-008 kinematics, there was a large pion background, sometimes up to 100 times the electron rate. Loose cuts on the gas Čerenkov detector and lead-glass shower counter were used to reject pions in the trigger, and tighter cuts were applied in the offline analysis.

The Čerenkov consisted of four mirrors and PMTs in the SOS, and two in the HMS. In both cases, the Analog-to-Digital Converter (ADC) output from each PMT was converted into the number of detected photoelectrons. The Čerenkov signal for the event is just the sum of the signals from the phototubes. No corrections were applied for position dependence of the signal, but the cuts were chosen to give high efficiency over the entire acceptance of the spectrometer.

For the calorimeter, one ADC value is measured for each module. The ADC value is converted to energy deposited in the block in GeV. Clusters of hits are located, and the energy per layer and total energy is calculated for each cluster. For each track found by the drift chambers, the energy associated with that track is the energy of the cluster that the track points to, if there is one. The energy is corrected for attenuation in the blocks based on the distance of the hit from the PMT, as determined by the tracking.

In addition to electrons, the spectrometer detects negative hadrons (mostly pions). The gas Čerenkov detector and lead-glass shower counter can separate the electrons from the hadrons. The trigger electronics require a signal from either one of these detectors before the event is accepted. Over the full range of the data, the ratio of pions to electrons varies between 10^{-3} and 10^{3}. In order to have a clean sample of electrons, a cut is applied requiring a good signal from both the Čerenkov and the shower counter.

Figure 9 shows the HMS Čerenkov spectrum for runs with high and low pion to electron ratios, taken without the particle identification in the trigger. The threshold on the Čerenkov signal in the trigger electronics corresponds to a cut at $\sim$1.5 photoelectrons, while the average signal was 10 photoelectrons. In order to improve pion rejection in software, the event was required to have more than 2 photoelectrons for the HMS. On average, this cut is 99.8% efficient, but at the edges of the mirrors in the HMS, the signal drops as low as $\sim$8-9 photoelectrons, which causes the inefficiency to increase by up to 0.8%. Figure 12 shows the measured number of photoelectrons as a function of the vertical position of the track at the HMS Čerenkov mirrors. The data is corrected for the average efficiency (99.8%), and a systematic uncertainty of 0.5% is assigned to the Čerenkov cut. The pion rejection for this cut is $\sim$550:1, with the main source of pion contamination coming from pions which produce knock-on electrons in the material immediately in front of the Čerenkov tank. If the knock-on electron is above the Čerenkov threshold ($\sim$15 MeV/c), it can emit Čerenkov light and cause the pion to be misidentified as an electron.

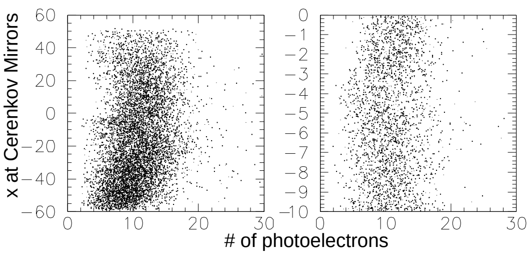

FIG. 12: HMS Čerenkov signal versus horizontal position at the mirrors. On the right is a blowup of the overlap region. Note that even at the lowest point in the dip, the mean signal is still 8-9 photoelectrons.

The lead-glass shower counter was also used to reduce the pion contamination. Because the calorimeter does not cover the complete acceptance of the spectrometer (some tracks miss the calorimeter for extreme values of δ), the reconstructed focal plane track was projected to the calorimeter and a fiducial cut was applied requiring that the track was at least 3 cm inside of the edge of the calorimeter.

In the HMS, the intrinsic calorimeter energy resolution is $\sim 6\%/\sqrt{E}$, but during the first half of the running, the ADC pedestals had small fluctuations, and the overall resolution was somewhat worse. The ADC offsets make discrete jumps, leading to offsets in the measured energy for pions and electrons. In most cases, the separation between the pions and electrons (pions should appear at $\sim$0.3 GeV) is large enough that the pion rejection is unaffected. In addition, because the calorimeter energy fraction cut was lowered as the momentum increased (see below), the calorimeter cut is efficient enough that there is no significant inefficiency for electron detection for this run. The fluctuations only occured during the first half of the run (after which the bad ADC was replaced), and only affected $\sim$1/3 of the runs during that period. For the majority of the runs, the electron energies were large and the fluctuations were small. For these cases, the pion rejection and electron efficiency were not significantly affected. For runs where the electron energies were smaller or the fluctuations large, the energy cut was lowered if the Čerenkov cut and reduced pion rejection were sufficient to remove the pions. Runs where this was not possible due to the large pion background were removed from the data set. For some of these runs it would have been possible to measure the pedestal shifts using the values from blocks that had no signal from the electron. However, all of the data that was rejected was taken at kinematics where there were other runs which were unaffected by the pedestal jumps. Therefore it was decided to eliminate the bad runs entirely and take the reduced statistics, rather than trying to correct these runs and have larger systematic uncertainties due to reduced electron efficiency or a non-negligible pion background.

The HMS detected particles with momenta between 0.995 GeV/c and 4 GeV/c. For the lowest momentum, where the resolution is the worst and the pion-electron separation is the smallest, the electron was required to have an energy fraction, E_{cal}/p, greater than 0.7. This cut is always 3σ or greater, ($\gtrsim$ 99.9% efficient) even for runs where the resolution is worse than usual due to pedestal drift. As the momentum increases, the energy fraction measured for electrons is still one, and the pion peak shifts to lower energy fraction ($\sim$ 0.3 GeV/p). During a portion of the running, all at higher momenta, the calorimeter ADC signals made discrete jumps during the course of a run. Therefore, while the resolution of the electron peak improves as the energy increases, there were some runs where the effective width was significantly larger then the normal $6\%/\sqrt{E}$. Therefore, the energy fraction cut was varied with energy, so that it was always highly efficient ($> 99.8\%$) for all energies, including runs where the pedestals varied during the run. The final cut used was:

$$E_{cal}/p > 0.7 - 0.07 * (p - 0.995) \qquad (1)$$

which corresponds to an energy cut of $0.7315p - 0.07p^2$. As the momentum increases, the energy resolution improves and the energy fraction cut decreases, increasing the electron efficiency of the cut. In addition, the absolute energy cut increases with momentum (for momentum values below 5 GeV/c), while the energy of the main pion signal remains constant. Therefore, the pion rejection is also improved as the momentum increases. However, even at very high energies there is still a small probability that a pion will deposit enough energy and be misidentified as an electron. While the majority of pions deposit roughly 0.3 GeV in the calorimeter, there is a small tail in the calorimeter energy distribution for pions that extends out to the full pion energy. The tail comes from pions which undergo a charge exchange interactions and become neutral pions. The neutral pions can decay into photons in the calorimeter, and their full energy can be deposited in the calorimeter. For the kinematics measured in e89-008, it is the lower momentum values where the pion rejection is most important, and in this region it is the resolution of the pion energy deposition that limits the pion rejection, rather than the tail. The HMS calorimeter pion rejection is $\sim$ 25:1 at 1 GeV, 50:1 at 1.3 GeV, and 150:1 at 1.5 GeV. For the HMS, the combination of Čerenkov and Calorimeter cuts reduces the pion contamination in the final data to $< 1.0\%$ for all kinematics. Figure 13 shows calorimeter signal (E_{cal}/p) versus the Čerenkov for the HMS at a central momentum of 1.11 GeV/c, with a pion to electron ratio of $\sim$70:1. For some higher momentum runs, the ratio of pions to electrons is much higher, but the calorimeter pion rejection improves as the energy increases, making this one of the worst cases for pion contamination. Figure ?? shows the pion to electron ratio (as calculated from the hardware scalers) versus the momentum for all of the data runs. The line shows the π/e ratio at which there is a 1% contamination after the particle identification cuts. The π/e ratio for the run is determined by taking the ratio of the PION and ELLO hardware scalers. At very high π/e ratios, the ELLO scaler will have a significant contribution from pions which produce a knock-on electron of sufficient energy to give a signal in the Čerenkov. The ELLO scaler was corrected for the expected pion contamination, based on the pion rejection of the Čerenkov trigger signal. Therefore, the calculated π/e ratio is accurate for π/e $\geq$ one. However, for π/e $\ll 1$, the calculated π/e is too high, due to electrons which do not fire the Čerenkov discriminator and are identified as pions.

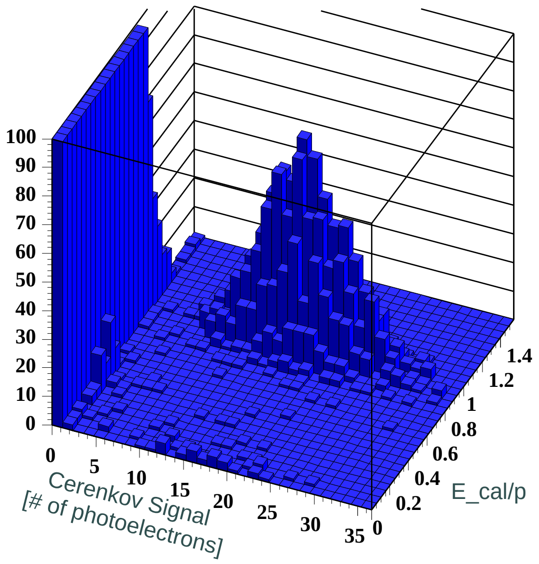

FIG. 13: Calorimeter (E_{cal}/p) versus Čerenkov for HMS run with a pion to electron ratio of approximately 70:1. The majority of the pions occur at 0 photoelectrons, though approximately 1% have a single photoelectron signal from noise.

Detector calibration

Calibrations had to be performed in order to match the timing of the individual scintillator elements, to calibrate the gains of the calorimeter and Čerenkov PMTs, and to convert the drift chamber TDC values to drift distances. For the gas Čerenkov, the final gains were calculated by hand. For each PMT, one gain parameter was needed; the number of ADC channels per photoelectron. The pedestal values were subtracted from the ADCs, and the gains were determined by finding the one photoelectron peak or by comparing the mean and widths of the signal in a central region. The drift chambers, hodoscopes, and calorimeter had a more complicated calibration procedure that involved running the tracking code and saving information for many events, and then fitting for the corrections using stand-alone code.

Not including any sections from my thesis here. Look for papers with details, or else just refer to thesis

data quality/elimination criteria?

CORRECTIONS

Overview. Yields corrected for dead time, efficiency, etc... Acceptance, radiative corrections, bin centering, etc... determined by applying all of these corrections to a model, comparing the predicted yield to the data, correcting the starting model and iterating.

Computer and electronic dead time

The main corrections to the measured number of counts come from data acquisition dead times and inefficiencies in the trigger hardware and the drift chambers. Electronic deadtime is caused when triggers are missed because the hardware is busy when an event that should generate a trigger comes in. When a logic gate in the trigger is activated, the output signal stays high for a fixed time. If another event tries to activate the gate in that time, it is ignored. If the mean event rate is R, then the probability of finding n counts in a time t is given by the Poisson distribution:

$$P(n) = \frac{(Rt)^n e^{-Rt}}{n!}, \tag{2}$$

and the probability distribution for the time between events is

$$P(t) = Re^{-Rt}. \tag{3}$$

An event will be missed if it comes within a time τ of an event accepted by the gate, where τ is the gate width of the logic signal. If the probability for this to occur is small enough, then this is nearly identical to the probability of an event coming within time τ of the previous event (whether or not the previous event triggered the logic gate). Therefore, for small dead times the fraction of measured events is equal to the probability that the time between events will be greater than τ:

$$\frac{N_{measured}}{N_{total}} = \int_{\tau}^{\infty} Re^{-Rt} dt = e^{-R\tau}. \tag{4}$$

In the trigger, all of the logic gates have a width of 30 ns, except for the hodoscope discriminators. The hodoscope discriminators have a very low threshold, and so their gate width was set to 50 ns in order to eliminate double pulsing of the discriminators caused by ringing of the signal. However, the hodoscope discriminators are not dead when their outputs are active. If a new signal comes in while the discriminator output is high, the output signal is extended to 60ns after the latest hit. Therefore, $\tau = 30$ns for the electronic dead time. For the trigger rates measured in this experiment, the live time was very close to 100%, and could be approximated by $e^{-R\tau} \approx 1 - R\tau$. To correct for the dead time, we generated four versions of the final electron trigger, each with a different gate width (τ = 30,60,90, and 120 ns). We then made a linear extrapolation to zero dead time in order to determine how many events were lost in the real electron trigger (τ = 30 ns). For each run we measured the electronic dead time and corrected the final cross section for the number of triggers lost. For the HMS, the maximum correction was $\approx 0.1\%$, and for the SOS it was $\lesssim 0.02\%$.

DO WE CARE ABOUT COINCIDENCE BLOCKING? ARE WE DOUBLE CORRECTING FOR IT HERE AND IN THE COMPUTER DEADTIME?????

There is another source of electronic deadtime, coming from singles triggers which were generated properly, but which were interpreted as coincidence triggers due to a random coincidence with an SOS trigger. The data acquisition system included HMS and SOS singles triggers, as well as coincidence triggers. Coincidence triggers only came as the result of random electron coincidences in the spectrometers. While the COIN triggers formed in the 8LM were prescaled away at the trigger supervisor (TS), if the HMS and SOS singles triggers come within the latching time of the TS ($\sim$7 ns), then the event will be treated as a coincidence. While each coincidence trigger indicates a trigger for both the HMS and SOS, they are not analyzed because the timing was not set up properly for coincidences, and there could be mistiming in the ADC gates and TDC stops. Because an event with HMS and SOS events coming within the TS latching time will be treated as a coincidence event, an SOS trigger coming between 7 ns before and 7 ns after an HMS trigger will cause the event to be tagged as a coincidence. If the rate of triggers in the SOS is R, and the time window for a coincidence trigger is τ (15 ns in this case), then the probability of an SOS trigger causing a random coincidence with an HMS trigger is:

$$\int_{0}^{\tau} Re^{-Rt} dt = 1 - e^{-R\tau}. \tag{5}$$

For $R\tau \ll 1$, the coincidence blocking deadtime can be approximated as $1 - e^{-R\tau} \approx 1 - (1 - R\tau) = R\tau$. For the most part, the coincidence blocking caused an inefficiency between 10^{-7} and 10^{-4} of the events. However, there were a few runs where the SOS singles rate was high enough to cause $\gtrsim$0.2% of the HMS events to be taken as coincidence triggers. However, for all of the runs where the SOS rate was high enough to cause a noticeable dead time, the SOS triggers were prescaled by a factor of 100 or more. This reduced the number of SOS triggers available to make a false coincidence with the HMS in the TS, and made the dead time negligible for these runs as well.

A more significant source of dead time for this experiment was the computer dead time. In this case, events are lost because a hardware trigger is formed when the data acquisition system is busy processing the previous event. The total processing time for an event is $\sim$300-400μs. However, when running in buffered mode the data acquisition can accept a new trigger before the old trigger is fully processed. It is only dead for $\sim$100μs, while the fastbus conversion of the data is in progress (see section ?? for more details). The computer dead time is measured by counting the number of triggers that were formed and the number of triggers that were processed by the Trigger Supervisor. The number processed over the number generated is the live time of the data acquisition system. The dead time is calculated for each run, and the cross section is corrected for the lost triggers. Figure 14 shows the computer deadtime for all runs. A few runs were taken in non-buffered mode, and have a processing time of 300-400μs, depending on the average size of the event. The average event size is dependent on the ratio of HMS to SOS events and the pion to electron ratio, since electrons will usually have extra ADC and TDC values for the calorimeter and Čerenkov signals. For some early runs, the parallel readout of multiple crates was not enabled and the event processing time was roughly 800μs. Note that at very high rates ($\gtrsim$2kHz) the deadtime is larger than expected for a 100μs processing time. This is because the minimum time between events is 100μs in buffered mode, but each event still requires $\sim$400 μs to process fully. Therefore, the maximum rate is $\sim$2500 Hz, and the effective processing time increases from 100 to 400μs as the incoming event rate goes beyond 2500Hz.

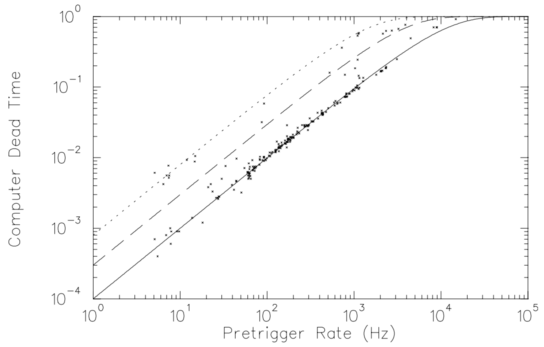

FIG. 14: Measured computer dead time vs. pretrigger rate. The solid line is the expected value for a processing time, τ, of 100μs, the dashed line is for τ=350μs and the dotted line is τ=800μs. Note that there is some uncertainty in the calculated pretrigger rate. The value plotted is the average rate over the entire run. Therefore, if the beam is off for part of the run, or if the current changes as a function of time, the pretrigger rate shown will not exactly represent the instantaneous rate that determines the deadtime.

Tracking and trigger inefficiency

Events are also lost due to detector inefficiencies that cause triggers to be missed, or inefficiency in the drift chambers or tracking algorithm that cause real events to be lost in the event reconstruction. Inefficiencies in the hodoscopes can cause a plane not to fire. The efficiency of each scintillator is determined by taking tracks that point to the center of the paddle (excluding the outer 2 cm of each paddle in the HMS, 1.25 cm in the SOS) and determining how often each paddle fires. Using the measured efficiency of the scintillators, we calculate the probability of missing a trigger due to hodoscope inefficiency and correct the counts for this loss. Because the trigger requires only three of the four planes, the scintillator trigger efficiency is always high, >99.4% for all HMS runs and $\gtrsim$99.8% for the SOS. In the HMS, the data is corrected run by run for the scintillator inefficiency, as determined by the measured hodoscope efficiencies for the run, and a 0.05% systematic uncertainty is assumed in the correction.

Additional trigger inefficiency can come if the particle identification signals in the trigger do not fire. The thresholds in the trigger are $\gtrsim$99.5% efficient for the Čerenkov, and >90% efficient for the Calorimeter (better than 99% efficient for higher energies). Since the trigger requires only one of the calorimeter signal or the Čerenkov signal, the PID is greater than 99.95% efficient in the trigger. Because the PID cuts in the analysis are tighter than the cuts in the trigger, we do not apply a correction for inefficiency in the trigger PID, we apply a single correction to take into account the total inefficiency of all PID cuts. The electron efficiency and pion rejection of the cuts was determined by taking runs with the particle identification signals removed from the trigger. In addition, the pion rejection is checked for each run by examining the calorimeter energy distribution after the final Čerenkov cut has been applied to insure that there is a clean separation of the pion and electron peaks, and that the pion contamination is at or below the level expected from the Čerenkov and calorimeter pion rejection.

Even if a trigger is formed, there will be some events where there is not enough information to reconstruct a track. The main sources of inefficiency of this kind are events where too many or too few wires fire in the drift chambers. If too few wires fire, the left-right ambiguity cannot be well determined, and a track is not fit. If too many wires fire, then the tracking takes a large amount of CPU time (finding all pairs and combinations of pairs of hits), and the chance of having a 'noise' hit included in the track increases.

The tracking efficiency is defined as the number of events for which a track is found, divided by the number of 'good' events (*i.e.* the number which we expect to have a real track). A trigger is defined as being a 'good' event if there was a trigger for the spectrometer, the time of flight determined before tracking determines it was a forward-going particle (rather than a cosmic ray), and one of the two drift chambers had less than 15 hits. We assume that events where both chambers have more than 15 hits are caused by electrons (or pions) which scrape the edge of one of the magnets and cause a shower of particles. Therefore, while there was a real particle, it was not within the acceptance of the spectrometer, and we should not correct for losing it due to tracking inefficiency. An event in which only one drift chamber had 15 hits is assumed to be a good event with additional hits due to noise in the chamber (which sometimes causes all 16 wires on a single discriminator card to fire) or the production of a knock-on electron which produces another short track and therefore another cluster of hits in one of the chambers. Since both of these conditions occur for good events within the acceptance of the spectrometer, we correct for these losses in the tracking efficiency. Once we require that one chamber was clean (<15 hits), then the number of tracks is corrected for the fraction lost to a single noisy chamber, a chamber with less than 5 planes hit, or events in which a consistent track cannot be made from the hits in the two chambers.

The tracking efficiency is calculated for all events, events passing a particle identification cut, events within a fiducial region of the hodoscopes, and events passing both the fiducial and PID cuts. This is because the efficiency calculated for all events includes the tracking efficiency for pions and background events as well as the real electrons. For runs where the electron cross section is low, the majority of events are pions or background electrons. By applying a PID cut, we reject the majority of the pions. By applying the fiducial cut, we look at the central and low momentum region, where the electron cross section is largest, and the signal to background ratio is larger. The data is corrected for the efficiency calculated using events passing the PID and fiducial cuts.

The HMS tracking efficiency is typically 93-97%. Roughly 1% of the loss comes from the drift chamber inefficiency causing too few hits, and the rest comes primarily from noise in a single chamber giving more than 15 hits in a plane. The tracking efficiency has large variations, but it was checked for several low and high tracking efficiency runs that the majority of event lost came from random noise in the amplifier/discriminator cards or the TDC.

The main uncertainty in this correction comes from the assumption that all events with one noisy chamber correspond to real events, and events with two noisy chambers correspond to bad (scraping) events. By looking at patterns of drift chamber and hodoscope hits for events where both chambers have >15 hits, we determined that $\geq$90% of these events come from showers in the detector. Since the maximum fraction of these events is <5% of the total events (after the PID and fiducial cuts), the maximum loss of good events is $\lesssim$0.5%. Similarly, $\geq$90% of the events where one chamber has >15 hits correspond to events where there is a single good track in the chambers and hodoscopes, but additional hits in one chamber, usually for a set of wires on a single amplifier/discriminator card. Usually 3-4% of the events have one noisy chamber, leading to a typical correction for 'junk' events of $\leq$0.4%. For a handful of runs, the number of events lost due to one chamber with >15 hits was as high as 8%, leading to a possible error of $\leq$0.8%. We correct the data for the measured efficiency (after PID and fiducial cuts) and assign an uncertainty of $\pm$0.5% to the correction.

PID efficiency

Target density fluctuations

Background electrons

In addition to rejecting pions, it is also necessary to reject background electrons. These are electrons that are not coming from the scattering of beam electrons in the target. There are two main sources of background electrons. First, there are events where particles coming from upstream or downstream of the target (beam halo scattering off of the beam pipe or background from the beam dump) enter the spectrometer after the magnets and create low energy electrons that reach the detectors. There are also 'secondary' high energy electrons that are produced in the target rather than being scattered from the beam.

In the HMS, background events come from low energy electrons from the beam dump scattering into the detector hut near the exit of the dipole. There is a vacuum pipe that runs through the magnets and into the detector hut. Particles in the hall that pass through the vacuum pipe after the magnets can be scattered into the detector hut (or produce knock-on electrons that make it into the hut). When the focal plane tracks are projected backwards to a point just before the entrance to the shielding hut, the events that come from scattering in the vacuum pipe can be seen as a 'ring' in the x-y plane, while real events are seen in the center. Prior to the experiment, shielding was added to decrease the background from particles entering the spectrometer after the magnets. In the analysis, a cut is applied to remove events that come from outside of the vacuum pipe. In addition, because these are low energy electrons, most are rejected in the calorimeter cut. The combination of the cut at the entrance to the hut and the calorimeter is sufficient to eliminate this source of background. Figure ?? shows a run with a very low rate of real events as well as a high rate run. In the low rate run, the events coming from the vacuum pipe are clearly visible. Because most of the background particles in the hall come from the beamline or the beam dump, they are traveling nearly horizontally when they pass through the vacuum pipe. This means that they pass through significantly more material if they strike the top or bottom of the pipe, and so have a greater chance of being scattered into the hut than particles which pass through the sides of the pipe.

There were also a significant number of events in which particles above the spectrometer momentum would hit the bottom of the dipole and be scattered into the spectrometer, or produce lower energy electrons which made it through the last part of the dipole and into the hut.

Before e89-008, shielding was added at the back of the dipole, in between the vacuum pipe and the magnet in order to reduce the background. In the analysis, the combination of the calorimeter cut, the cut at the hut entrance, and the cuts on reconstructed target quantities eliminated these events.

In addition to background coming from the low energy electrons, there are secondary electrons produced in the target. Since they are secondary electrons, rather than scattered electrons, they are a background for the measurement. The main background of secondary electrons most likely comes from electro-production and photo-production of neutral pions. These pions then decay into photons which can produce positron-electron pairs. This background is charge-symmetric, and can be measured by running with the spectrometers in positive polarity, and detecting the produced positrons. For the largest angles (55° and 74°), this background was significant. In this case, the positron production cross section was fit from our measurements and subtracted from the electron data. For the smaller angles, this background was negligible (<1%).

Positive polarity data was typically only taken for one or two targets for each kinematics. We parameterize the ratio of positron to electron production in terms of the target thickness (in radiation lengths), and extrapolate the measured positron cross sections to the thickness of the other targets. The e^+/e^- ratio can vary by up to a factor of four between the different targets, but the positron rate differs from the parameterization by only $\sim$10% over this range. Most of the positive polarity data were taken with the thick targets in order to maximize the positron statistics. Therefore, the extrapolation of the measured e^+/e^- ratio between the different thick target had only a small uncertainty $\approx$1-2%, while the extrapolation to thin targets was uncertain at the $\sim$10% level. However, the ratio of positrons to electrons was near unity for the thick targets, but only $\sim$30% for the thin targets. Therefore, the uncertainty due to the target thickness extrapolation is $\lesssim$3% of the total electron cross section. Rather than making a point by point subtraction of the measured positron cross section, all positron data at 55deg and 74deg was fit in order to obtain the cross section to be subtracted due to the charge-symmetric background. The uncertainty in the positron fit was a combination of the uncertainty due to target thickness differences, and due to the statistics of the measurements.

Figure 15 shows the background subtracted electron and raw positron cross sections for scattering from the thick Gold at 55°, and from the thick Iron and thin Carbon targets at 74°. At 55°, the charge symmetric background is $\sim$10% of the electron cross section for the thick targets, and $\sim$3-4% for the thin targets. At 74°, the background can be equal to or larger the electron cross section for the thick targets, and $\sim$20% for the thin targets.

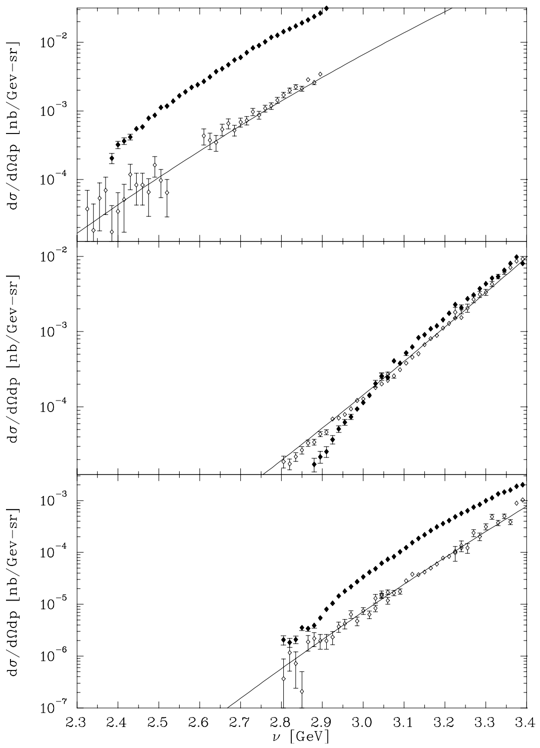

FIG. 15: Electron and positron cross sections. The filled diamonds are the measured electron cross section after subtraction of the charge symmetric background. The hollow diamonds are the measured positron production cross section. The top plot is 55deg data measured with the thick Gold target (5.8% of a radiation length), the middle is 74deg with the thick Iron target (5.8% r.l.), and the bottom is 74deg data with the thin Carbon target (2.1% r.l.). The solid line is the fit to all positron data that is subtracted from the electron cross section.

Angle and energy determination

Spectrometer acceptance

For a fixed angle and momentum setting, the HMS (and SOS) will measure data in a limited range of angles and momenta around the central values. As we move away from the central kinematics, some fraction of the events will be lost if they hit the collimator, scrape the walls of the magnets, or miss detector elements required for the trigger or in the data analysis. For scattering with a cross section σ, the number of events detected in the spectrometer will be a function of the point where the scattering occurs in the target, and the kinematics of the spectrometer:

$$N = \int d\delta dx' dy' dx dy dz \cdot \sigma(\delta, x', y', x, y, z) \cdot A^6(\delta, x', y', x, y, z), \quad (6)$$

where $A^6(\delta, x', y', x, y, z)$ is the acceptance function of the spectrometer which represents the probability that a scattering event coming from the point (x, y, z), with kinematics defined by δ, x', and y' will be detected. We can use a model of the spectrometer to perform a Monte Carlo calculation of the acceptance function of the spectrometer. However, it is not feasible to generate enough statistics in the Monte Carlo to have a high precision calculation of acceptance as a function of all 6 variables. Therefore, we would like to define a simplified acceptance function, which averages over the behavior of several of the variables.

As long as the target is thin enough that there is no significant loss of beam intensity as a function of position along the target, the cross section is independent of x, y, and z. The cross section is then just a function of δ, x', and y'. This means that we can now integrate over x, y, and z over the region of interest (as defined by the position and size of the beam and target), and come up with an acceptance function in terms of just δ, x', and y' which takes into account the acceptance of the spectrometer in x, y, z, and which is independent of the scattering kinematics:

$$N = \int d\delta dx' dy' \cdot \sigma(\delta, x', y') \int dx dy dz \cdot A^6 \equiv \int d\delta dx' dy' \cdot \sigma(\delta, x', y') \cdot A^3(\delta, x', y') \tag{7}$$

In order to further simplify the acceptance function, we can fix the central angle of the spectrometer, and convert from x' and y' to the in-plane and out-of-plane scattering angles θ and ϕ. Because the inclusive cross section is independent of ϕ, we can integrate over ϕ and define a two-variable acceptance function, $A^2(\delta, \theta) = \int A^3(\delta, \theta, \phi) d\phi$, such that

$$N = \int d\delta d\theta \cdot \sigma(\delta, \theta) \cdot A^2(\delta, \theta). \tag{8}$$

We can generate events in x, y, z, δ, θ, and ϕ in the Monte Carlo, and bin the results as a function of just δ and θ in order to determine the acceptance of the spectrometer. The Monte Carlo model has three main elements: the event generator, the transportation of the particle through the magnets, and the list of materials and apertures that cause multiple scattering or stop the particles. The event generator creates a large set of initial particles distributed uniformly in δ, θ, ϕ, x, y, and z. The particles are then run forward through the model of the spectrometer, and focal plane tracks are recorded for all particles which make it all of the way through the detector stack. These tracks are reconstructed to the target in the same way as the measured events.

The magnetic portion of the spectrometer is modeled using the COSY INFINITY program from MSU[21]. COSY takes a list of positions, fields, and lengths for the quadrupoles and dipoles in the spectrometer and generates a forward matrix that converts from rays at the target to rays at the focal point (or any other point in the spectrometer). The transport matrix calculates the focal plane quantities (x_{fp}, x'_{fp}, y_{fp}, and y'_{fp}) based on the target quantities $x_{tar}, x'_{tar}, y_{tar}, y'_{tar}$, and $\delta = (p - p_0)/p_0$, where p_0 is the central momentum setting of the spectrometer. The expansion for each of the focal plane quantities is of the following form:

$$x_{fp} = \sum_{i,j,k,l,m} F^x_{ijklm} \cdot x^i_{tar} y^j_{tar} (x'_{tar})^k (y'_{tar})^l \delta^m \qquad (1 \leq i+j+k+l+m \tag{9}$$

where N is the order of the expansion, F^x_{ijlkm} is one column of the forward transport matrix (one column for each of the four focal plane quantities), and i, j, k, l, and m are integers between 0 and N. For the HMS, the forward transport matrix is calculated to 5th order, and for the SOS it is calculated to 6th order. In both cases, a significant fraction of the matrix elements are zero. For example, because of mid-plane symmetry, all terms contributing to y_{fp} and y'_{fp} are zero if the combined power of the y_{tar} and y'_{tar} terms is even (*i.e.* if j+l is even).

COSY is used to generate forward matrices that take an event from the target to several points in the magnetic system, not just the focal plane. The events are transported to the beginning and end of each magnet in order to reject events that are outside of the acceptance of the magnets. In addition, the position for the event is determined 2/3 of the way through Q1 and Q2 in order to reject events that hit the inside of the magnet. COSY also generates reconstruction matrices, used to determine the target quantities y_{fp}, x'_{fp}, y'_{fp}, and δ from the focal plane tracks. Because δ is not directly measured at the focal plane, only four quantities can be reconstructed. For purposes of calculating the reconstruction matrix elements, the events are assumed to come from x_{fp}=0, where x_{fp} is the vertical position at the target. Thus, the reconstruction of the target quantities is of the form:

$$y_{tar} = \sum_{i,j,k,l} R^y_{ijkl} \cdot x^i_{fp} y^j_{fp} (x'_{fp})^k (y'_{fp})^l \qquad (1 \leq i+j+k+l \leq N) \tag{10}$$

where R^y_{ijkl} is one column of the reconstruction transport matrix. For the HMS, the COSY generated reconstruction matrix elements were used to reconstruct the target quantities from the measured focal plane quantities in the real data. For the SOS, the reconstruction matrix elements were fitted from data. The fitting procedure is described in [22] and involved fitting sieve slit data in order to reconstruct the angles, elastic data (with a known p-θ correlation) to reconstruct momentum, and sieve slit data from targets at different positions along the beam to reconstruct y_{tar}. For the HMS, the COSY reconstruction matrix elements were used because elastic data was not available over the entire range of momenta needed for the analysis of the e89-008 data. However, comparison of the data to the Monte Carlo and the reconstruction of the sieve slit data indicate that the COSY matrix elements give a good reconstruction of the data.

Finally, multiple scattering effects are applied to the events, and cuts representing physical apertures or software cuts applied to the real data are applied to the events. The most significant multiple scattering occurs in the target material and scattering chamber exit window. While there is greater multiple scattering in the detector material itself, the scattering that occurs before the particle passes through the magnets has the most significant effect on the resolution. Gaussian multiple scattering was applied to the events for scattering in the target and the scattering chamber exit window and spectrometer entrance window. The particles were projected forward to the slit box, and particles outside of the octagonal collimator were rejected. The events were transported through the magnetic field to various points in the spectrometer using the COSY generated forward matrix elements. Cuts were applied at the entrance and exit of each magnet, at a point 2/3 of the way through Q1 and Q2, and at the beamline apertures between the dipole exit and the entrance to the detector hut. Events that hit the magnets or apertures in the spectrometer are rejected. Particles that reached the detector hut were projected through each of the detector systems, with multiple scattering applied for the detectors and the air in the hut. Events which missed detector elements that are required in the trigger or in the data analysis were thrown out. The position at the wire chamber planes were smeared out with the wire chamber resolution and recorded, and tracks were fit through the 'measured' positions. This track was reconstructed to the target using the COSY reconstruction matrices. The COSY matrix elements were used for reconstruction for both the HMS and SOS Monte Carlos. Even though we fit the reconstruction matrix elements for the SOS data analysis, we use the COSY values in the Monte Carlo so that we have a consistent model for both forwards and backwards reconstruction. Then, the cuts that were applied to the reconstructed data were applied to the Monte Carlo events. The events that passed through the spectrometer and were reconstructed to the target were binned in δ and θ. The acceptance for a given δ, θ bin is defined as the number of events that pass all cuts and are reconstructed into that bin divided by the expected number of events generated in that bin (*i.e.* the total number of generated events divided by the number of (equally sized) δ, θ bins).

The Monte Carlo distributions of events at the focal plane were compared to the distributions from the data. From this, offsets between the detectors in the Monte Carlo and in the spectrometer were determined, and these offsets were applied to the Monte Carlo. It was noted that the Monte Carlo events were being cut off by the vacuum pipe between the HMS dipole and the detector hut, while in the real data, events were not being lost. Because the vacuum pipe was not precisely surveyed in the spectrometer, it was shifted down 2.0 cm in the model in order to match the cuts seen in the data.

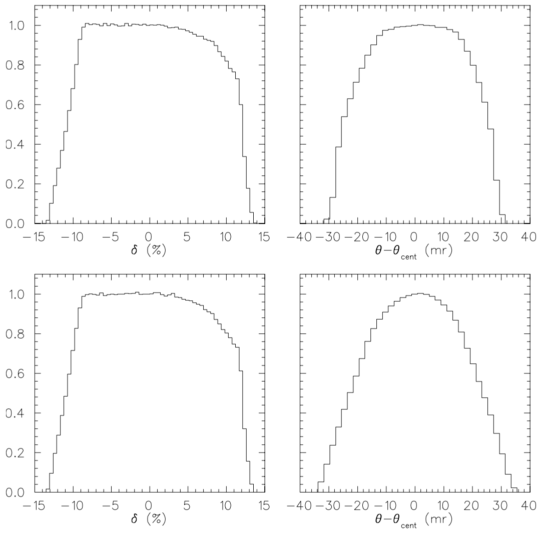

FIG. 16: HMS δ and θ acceptance for 55°. The top figures are for a point target, the bottom for a 4cm target. The curves are arbitrarily normalized to one at the peak value.

Figure 16 shows the HMS δ and θ acceptance at 55° for a point target, and for the short (4.2cm) target. Note that at 55°, the target length as seen by the spectrometer is 3.4cm. The acceptance is normalized to one for the central δ or θ value. For the SOS, the extended target causes a significant loss of events as $|\delta|$ increases. Section ?? shows comparisons of the data to Monte Carlo for a version of the Monte Carlo which has the elastic cross section. This allows us to compare the data to the Monte Carlo directly, without having to divide the cross section out of the data.

Rather than dividing the acceptance out of the data for each δ, θ bin, the acceptance correction was applied at the same time as the bin centering corrections in order to reduce the systematic uncertainties and model dependence of that correction. The procedure is described in detail in the following section.

Bin centering corrections

In order to measure the cross section at fixed values of p and θ, we must bin the data and make a correction to convert from binned counts (which represent the integral of the cross section over the bin) to the value of the cross section at the center of the bin. The goal of the analysis was to extract the cross section for a range of p values at a fixed angle. Therefore, the initial procedure involved binning the data into small p,θ bins, corresponding to the δ, θ bins used in determining the spectrometer acceptance. Each bin then was corrected by the Monte Carlo acceptance for that bin. The acceptance corrected counts were then rebinned into 15 MeV momentum bins and summed over the full θ acceptance of the spectrometer ($\sim\pm 25$ mr for the HMS, $\sim \pm 60$ mr for the SOS). The cross section variation over the 15 MeV p bin was generally small, and the correction was determined by taking a model cross section and calculating the ratio of the central cross section to the average cross section over the momentum bin:

$$p \text{ Binning Correction} = \frac{\sigma^*(p_0, \theta) \cdot \Delta p}{\int_{p_0-\Delta p/2}^{p_0+\Delta p/2} \sigma^*(p, \theta) dp}, \qquad (11)$$

where σ^* is the model differential cross section, and Δp is the momentum bin size. Since the number of counts in a p bin measures the integral of the cross section over that bin (the denominator in the above expression), multiplying the measured counts by this bin correction factor yields the central value of the cross section. Because this correction is small (usually $<1\%$, and always $\lesssim 5\%$) and the model has been adjusted to reproduce the data, the uncertainty on this correction is quite small.

This procedure can be extended to take into account both the p bin and the θ binning:

$$(p, \theta) \text{ Binning Correction} = \frac{\sigma^*(p_0, \theta_0) \cdot \Delta p \cdot \Delta\theta}{\int_{\theta_0-\Delta\theta/2}^{\theta_0+\Delta\theta/2} \int_{p_0-\Delta p/2}^{p_0+\Delta p/2} \sigma^*(p, \theta) dp d\theta}. \qquad (12)$$

However, as noted before, the θ bin size is the entire θ acceptance of the spectrometer. Over this range, the cross section variations can be very large (more than an order of magnitude). In this case the correction is often large, and the model dependence in this correction can be the dominant systematic uncertainty in the analysis.

There were two changes made to the above procedure in order to reduce the size and the uncertainty of this correction. Note that a linear variation to the cross section over the acceptance will have no bin centering correction, and only higher order variations will produce a correction. Therefore, the bin centering correction, coming from higher order variations of the cross section, will grow rapidly with the size of the θ bin. This means that one could reduce the size of the correction by applying a tight θ cut. This would reduce the correction, but would also throw out a large part of the data. However, the θ range is already limited by the acceptance of the spectrometer. When we apply the acceptance correction, we increase the weight of the counts at the edges in θ, where the acceptance is falling off. This is done so that the measured counts represent the incoming counts, before they are cut out by the collimator. We then are measuring the counts over the full θ range of the spectrometer, and so in the bin centering correction we compare the central value of the cross section to the integral over the full θ range. If we do not correct for the θ acceptance, then we are measuring the cross section times the acceptance, and therefore reduce the weight of the measurement when θ is far from the central angle. We can modify our procedure to take advantage of the fact that the data has reduced acceptance at large angles by rewriting equation 12 with the *acceptance weighted* cross section in the denominator:

$$(p, \theta) \text{ Binning Correction} = \frac{\sigma^*(p_0, \theta_0) \cdot \Delta p \cdot \Delta\theta}{\int_{\theta_0-\Delta\theta/2}^{\theta_0+\Delta\theta/2} \int_{p_0-\Delta p/2}^{p_0+\Delta p/2} A(p, \theta) \cdot \sigma^*(p, \theta) dp d\theta} \qquad (13)$$

The denominator now represents the acceptance weighted counts, which gives less weight to the values of θ far from the central angle, thus reducing the correction. This means that by applying the acceptance correction at the same time as the bin centering, we can reduce the size of the binning correction, and therefore the associated uncertainty.

The other improvement involved binning the data in different variables. Once we have applied the bin centering correction, we are looking at the cross section at a fixed valued of p and θ. At that point, we can freely translate to any other desired variables that specify the kinematics. This means that if we start with variables other than p and θ, bin the data and apply acceptance and bin centering corrections, we can convert back to the desired p and θ values. Thus, if we can replace p with some other variable, over which the θ variation of the cross section is smaller, we can bin the data over θ and have a significantly smaller bin centering correction than when we use p and θ. For fixed p, the cross section varies by a factor between 5 and 200 over the theta acceptance of the HMS ($\sim 3^\circ$, or roughly 1/2 to 1/3 of the spacing for the angles shown). This is what causes the large correction using the method of equation 12. The correction is especially large at the higher values of p, corresponding to the large Q^2 values which are of the most interest, and where the model cross section is least well known. For fixed values of x, the cross section variation over the HMS θ acceptance is typically a factor of 1.5 to 3, and is always $\lesssim 10$. The θ variation for fixed ξ is even smaller, usually less than a factor of 2, and is smallest at the high Q^2 values (corresponding to large scattering angles). Therefore, by binning in ξ and θ, and including the acceptance in the correction, rather than directly to the binned counts, we have a significantly smaller bin centering correction of the form:

$$(\xi, \theta) \text{ Binning Correction} = \frac{\sigma^*(\xi_0, \theta_0) \cdot \Delta\xi \cdot \Delta\theta}{\int_{\theta_0-\Delta\theta/2}^{\theta_0+\Delta\theta/2} \int_{\xi_0-\Delta\xi/2}^{\xi_0+\Delta\xi/2} A(\xi, \theta) \cdot \sigma^*(\xi, \theta) d\xi d\theta} \qquad (14)$$

where σ^* is now the differential cross section $\frac{d\sigma}{d\xi d\Omega}$, rather than $\frac{d\sigma}{dpd\Omega}$.

Figure 17 shows the size of the bin centering correction for 30deg, taking fixed p, x, or ξ and binning over a $\pm$1.4deg bin. For each variable, the correction was calculated using two models in order to estimate the model dependence. The top line is using our final model of the cross section. The bottom line comes from adding an additional Q^2 dependence to the model. The standard model is typically within 10% of the data (and always within 30%), and has small (<10%) variations in the ratio of data to model when comparing different angles. The modified model ($\sigma^* = \sigma \cdot \frac{Q^2}{\langle Q^2 \rangle}$) introduces large discrepancies between the model and data (up to a factor of 5), and introduces a large angular variation in the ratio of data to model. While this severely overestimates the uncertainty in the θ dependence of the model, it still leads to a small uncertainty in the correction when taking fixed ξ.

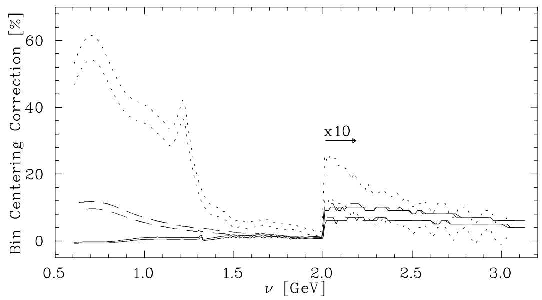

FIG. 17: Bin centering corrections at 30deg for a $\pm$1.4deg bin. The dotted line is the correction at fixed p, dashed is for fixed x, and solid is for fixed ξ. The top line in each case represents the correction calculated using the standard cross section model. The bottom line is for the model with a large Q^2 dependence, used to estimate the uncertainty in the correction.

In the real data, the acceptance does not always include a symmetric region in θ about the central value in ξ. The acceptance of the spectrometer is a roughly rectangular region in δ and θ. A fixed ξ bin is a roughly straight line through the δ,θ acceptance region. For a value of ξ corresponding to $\delta = 0$, $\theta = \theta_0$, the entire θ range is included in the bin. For ξ bins corresponding to high or low values of δ (at the central angle), only part of the θ acceptance lies within the spectrometer acceptance. Therefore, the bin centering corrections are largest at the edge of the momentum acceptance, where a bin of fixed ξ only includes half of the θ acceptance. Instead of comparing the average cross section to the central value, we are comparing the average to the extreme value, and so the maximum bin centering corrections occur at the edge of the acceptance.

We apply an overall 1% systematic uncertainty in the cross section due to the bin centering correction. In addition, we apply an additional systematic uncertainty equal to 10% of the correction made. The maximum bin centering correction (for 15deg, very low ν) is 20%, leading to a 2% uncertainty in the correction (in addition to the 1% overall uncertainty).

Because the correction for the cross section variation over the ξ bin is small, it is a good approximation to separate the binning centering correction into two pieces. By separating the ξ and θ bin centering corrections, the corrections involve one dimensional integrals over the model cross section, rather than a two-dimensional integral. This significantly reduces the time required to calculate the correction.

In order to check the acceptance and bin centering correction, runs with significant overlap in momentum were taken. This allows us to have multiple measurements of the same cross section, taken in different regions of the spectrometer. Figure ?? shows the cross sections (in arbitrary units) from three runs with central momentum settings of 2.06, 2.20, and 2.36 GeV/c. It also shows the difference between the fit and the individual points as a function of δ. The typical deviations from the fit are consistent with statistical uncertainties of the individual points ($\chi^2_\nu = 1.10$ for 72 degrees of freedom), and a systematic uncertainty of 1% is applied to the acceptance at the peak value.

The data is cut when the acceptance for a ξ bin falls below 50% of the maximum acceptance. The uncertainty associated with the acceptance is 1% (1.3% in the SOS) combined in quadrature with 4% of the difference between the acceptance for the bin and the maximum acceptance. Therefore, for a bin with an acceptance of 0.5, the systematic uncertainty is $(.01^2 + (.04 * (1 - 0.5))^2) =$ 2.2%. In δ and θ, the acceptance is roughly rectangular, and falls from 1 to 0 very quickly. Where the acceptance drops very rapidly, the Monte Carlo is very sensitive to small offsets or differences in resolution. Therefore, the uncertainty is large for a δ bin at the edge of the acceptance. However, when the data is taken as a function of ξ, the decrease in the acceptance comes mainly from the fact that the kinematic transformation between ξ and δ means that only a certain portion of the ξ bin has acceptance. Because the fraction that is populated comes from the mapping between ξ and δ rather than losses at the edges of the spectrometer, it is less sensitive to any small offsets or resolution differences. Therefore, the uncertainty in acceptance correction is relatively insensitive to the size of the correction, and even for an acceptance of 0.5 (which leads to a 100% correction in the cross section), the uncertainty is small.

Radiative corrections

This section can probably be pretty short. Main points are:

-Presription is described in (DHP, JRA thesis)

-Model used is (y-scaling+convolution), with ad-hoc adjustments to reproduce global behavior of measured cross sections.

We Apply rad. corr (and everything else) to a model of the cross section, compare to data, adjust model, iterate... (seperate section for the details of this). Mention tests (thick vs. think, different starting models, etc...)

Coulomb distortion

After the incoming electron passes through the atomic electrons of the target atom, it sees a bare nucleus, and is accelerated by the electric field of the nucleus. This acceleration leads to an increase in the energy of the incoming electron, and a decrease in the energy of the scattered electron. This means that the energy of the initial and scattered electron at the scattering vertex is not the same as the energies determined by measurements of the beam energy and the scattered electron momentum. This change in kinematics can have a significant effect on the measured cross section. In addition, the electric field of the nucleus can lead to a deflection of the electron when the scattering occurs at the edge of the nucleus. This deflection of the electron means that at fixed spectrometer angle, we are measuring over a range of scattering angles.

We estimate the effect of the Coulomb energy correction by calculating the cross section from our model with and without the energy shift due to the Coulomb acceleration. In order to estimate the energy shift, we treat the nucleus as a uniform sphere of radius R_0. Then, the electric potential for a point r inside of the nucleus ($r < R_0$) is given by:

$$V(r) = -\frac{Ze}{8\pi\epsilon_0 R_0}\left(3 - \frac{r^2}{R_0^2}\right). \quad (15)$$

with $V(\infty)$ defined to be zero. Outside of the electron cloud, the potential from the nucleus is canceled by the potential from the electrons. However, at typical electron distances, the potential is $\sim 10^{-4}$ of the potential at the surface of the nucleus. We thus neglect the shielding by the atomic electrons, and the energy change for the electron at the surface of the nucleus is:

$$\Delta E(R_0) = eV(R_0) = \frac{Ze^2}{4\pi\epsilon_0 R_0} = 1.44 MeV \frac{Z}{R_0}. \quad (16)$$

Assuming that the scattering occurs uniformly throughout the nucleus, we calculate the average energy shift for the scattering:

$$\langle \Delta E \rangle = \frac{\int_0^{R_0} V(r) r^2 dr}{\int_0^{R_0} r^2 dr} = \frac{6}{5}\Delta E(R_0). \quad (17)$$

Nucleus	R_0 [fm]	$\Delta E(R_0)$ [MeV]	$\langle \Delta E \rangle$ [MeV]	RMS $p_\perp$ [MeV/c]
^{12}C	3.23	2.67	3.2	1.5
^{56}Fe	4.85	7.72	9.3	4.4
^{197}Au	6.88	16.53	19.8	9.8

TABLE IV: Effective radius, Coulomb energy correction (at surface and averaged over the nucleus), and RMS transverse momentum kick for the target nuclei. The radius is taken from [23], and is the effective radius for the nucleus, assuming a spherical nucleus with uniform charge density.

Table IV gives the values for R_0, $\Delta E(R_0)$, and $\langle \Delta E \rangle$ used in the correction. Using this average energy correction, we estimate the correction to the cross section by calculating the cross section for our model at the nominal kinematics, and with the Coulomb energy correction applied ($E \rightarrow E + \langle \Delta E \rangle$, $E' \rightarrow E' + \langle \Delta E \rangle$, and ν remains constant at the point of interaction). We take the modification of the cross section model as our correction to the data for the Coulomb energy correction. The correction is roughly proportional to $\langle \Delta E \rangle$, and averages 2% for Carbon, 5.5% for Iron, and 9.8% for Gold. The largest corrections to the data occur at 74deg, and are at most 6% for Carbon, 15% for Iron, and 24% for Gold.

In addition to the energy change for the initial and scattered electron, the Coulomb field of the nucleus will lead to a deflection of the electron. The maximum deflection occurs when the electron grazes the nucleus. In this case, the incoming electron can be approximated by integrating the component of the force transverse to the electron direction, neglecting the change in the trajectory. In this case, the transverse 'kick' received by the electron is:

$$\Delta p_\perp = \int_{-\infty}^{0} F_\perp dt = \frac{1}{c}\int_{-\infty}^{0} F_\perp dr_\parallel = \Delta E(R_0)/c \quad (18)$$

The worst case is for gold, where $\Delta p_\perp$=16.5 MeV/c for an electron at $r_\perp = R_0$. This leads to an angular deflection of $\Delta\theta = \Delta p_\perp / p_{beam} = 4.1$ mr, which is much larger than the uncertainty in the θ reconstruction. In addition, there will be a transverse kick of similar magnitude to the scattered electron. Because the scattered electron energy can be much lower than the beam energy (as low as $\sim$600 MeV), the deflection can be much larger. A Monte Carlo calculation was used to determine the distribution of $\Delta p_\perp$ for events generated uniformly within the nucleus. The distribution of $\Delta p_\perp$ for Carbon, Iron, and Gold is relatively flat, and was approximated by a flat distribution with a width chosen to match the RMS value of the calculated distribution. The correction to the cross section was determined by comparing the model cross section at the measured angle to the average value over the θ range determined by combining the angular range of the incoming electron ($\Delta\theta = \Delta p_\perp/p_{beam}$) with the angular range of the scattered electron ($\Delta\theta' = \Delta p_\perp/p'$). The angular range can be large for high ν (low E'), but the cross section has the greatest θ variation at low ν, and the correction is never very large. While the angular deflection range is proportional to the $\Delta p_\perp$ kick, the correction grows at least as fast as the square of the angular range. The correction is $\lesssim$5% for Gold, $\lesssim$2% for Iron, and $\lesssim$0.5% for Carbon, and has the opposite sign as the correction for the energy change of the electrons (except when the correction is very small).

Figure 18 shows the correction for Iron, as a function of angle. The crosses show the correction to the model when the coulomb energy correction is applied, the diamonds show the correction to the model coming from the deflection of the electrons, and the circles show the combined effect. For Gold the correction is roughly twice as large, and for Carbon, the total correction is roughly one third of the correction for Iron. The Coulomb correction for the Hydrogen elastic scattering data has a negligible effect on the cross section, and a small effect effect on the measured position of the W^2 peak. However, the effect was small enough that it does not significantly affect the conclusions of the spectrometer momentum and beam energy calibration.

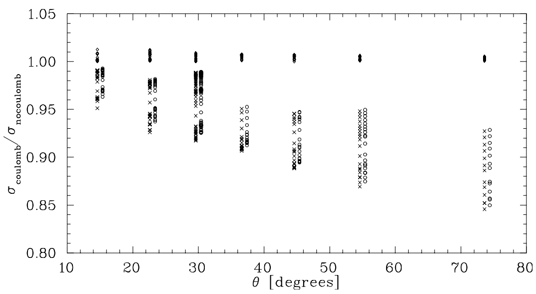

FIG. 18: Coulomb corrections for the Iron data. The crosses represent the change in the model cross section when ΔE is applied. the diamonds are the correction when the angular deflection is applied, and the circles are the combined effect. The correction is roughly twice as large for Gold, and one third of the size for Carbon. The multiple points at each angle represent different values of ν. The corrections are largest for the lowest ν values.

The main source of uncertainty in the correction comes from the assumption that the nucleus can be modeled as a sphere with uniform charge distribution, and the uncertainty in the radius chosen for the sphere. In addition, it is assumed that the electron scattering occurs uniformly throughout the volume of the nucleus. However, the boost in electron energy will modify the cross section as a function of position from the center of the nucleus, leading to a slightly non-uniform distribution of events. We estimate that the uncertainty associated with modeling the nuclei as uniform spheres, and the choice of radius (given in table IV) is less than 8% of the correction. The model dependence in calculating the correction is less than 5% of the correction. Finally, the maximum difference in cross section between the center of the nucleus and the edge of the nucleus is ~10% in Gold, ~7% in Iron, and ~2-3% in Carbon. We assume that the average effect of including the cross section weighting is always less than half of the maximum cross section variation, and use half of this value as the overall uncertainty. In the current analysis, we use the maximum correction to determine the overall systematic uncertainty for each target, giving an upper limit for the uncertainty in the extracted cross section of 0.6% for Carbon, 1.7% for Iron, and 2.8% for Gold. This uncertainty is fairly small relative to the other systematic uncertainties (typically 3.5-4.0%). With a more careful comparison of different models for the charge distribution and the effects of neglecting the cross section weighting, these uncertainties should be significantly reduced from their present values, and should have a negligible effect on the total systematic uncertainties for Carbon and Iron, and a small effect for Gold.

As part of the radiative correction procedure, the model cross section is corrected for radiative effects, and the scattering kinematics are corrected for energy loss in the target and in the spectrometer vacuum window. However, while the coulomb correction could also be applied as part of the radiative correction procedures, there are two advantages to making a separate correction. First, we need to apply the same correction to the data from previous measurements [24–26] in order to compare results (the analysis of the NE3 data and the inclusive analysis of the NE18 data did not include coulomb corrections except for the extrapolation to nuclear matter [27]). Only the Iron data is compared to the SLAC results, and the average coulomb correction is ~4%, and the maximum correction is 10%. In addition, while the energy change due to the coulomb correction is applied as a shift in energy, the deflection of the electron due to the coulomb field leads to an averaging of the cross section over a range in θ. Including this in the radiative correction procedure would significantly increase the CPU time required to determine the radiative corrections.

EXTRACTION OF THE CROSS SECTION

Some details on starting model. Explanation of procedure (as mentioned in notes in the Radiative corrections section).

NOTE: Make sure to mention the proton and neutron form factors used in the model, and in the extraction of the y-scaling. Based on Rosenbluth form factors, which appear to suffer from large two-photon exchange contribution [28–31] (as well as corrections due to Coulomb distortion [32]). However, since what we want is the CROSS SECTION for e-p scattering, they provide the *correct parameterization for the cross section in the one-photon exchange approximation that we use. [7, 33].*

This work was supported in part by DOE Grants W-31-109-ENG-38 and DE-FG02-95ER40901, NSF Grant 0099540, and the Schweizerische Nationalfonds.

[1] C. W. Leemann, D. R. Douglas, and G. A. Krafft, Ann. Rev. Nucl. Part. Sci. **51**, 413 (2001).
[2] C. Yan et al., Nucl. Inst. & Meth. **A365**, 261 (1995).
[3] P. Gueye, Tech. Rep., CEBAF Internal Report (1995).
[4] C. Yan, Tech. Rep., CEBAF Internal Report (1996).
[5] P. Gueye, M. Tiefenback, and C. Yan, Tech. Rep., CEBAF Internal Report (1996).
[6] J. Arrington, Ph.D. thesis, California Institute of Technology (1998).
[7] D. Dutta et al., Phys. Rev. C **68**, 064603 (2003).
[8] G. Niculescu, Tech. Rep., CEBAF Internal Report (1995).
[9] C. Armstrong, Tech. Rep., CEBAF Internal Report (1996).
[10] K. B. Unser, Tech. Rep., CERN SL/91-42 (1991).
[11] F. Duncan, Tech. Rep., CEBAF Internal Report (1997).
[12] K. Gustafsson, Tech. Rep., CEBAF Internal Report (1996).
[13] R. Ent, private communication.
[14] J. Cobb and J. Murray, Nucl. Inst. & Meth. **46**, 99 (1967).
[15] *CEBAF Conceptual Design Report* (CEBAF press, Newport News, 1995).
[16] O.K.Baker and other, Nucl. Inst. & Meth. **A367**, 92 (1995).
[17] H. Mkrtchyan, what is this reference????????? ???????).
[18] C. Group, *Paw (physics analysis workstation) users guide, program library q121, cern.*
[19] S. A. Wood and W. A. W. III, in *Conference Record of the Eigth Conference on Real-Time Computer Applications in Nuclear, Particla and Plasma Physica*, edited by D. Axen and R. Poutissou (TRIUMF, Vancouver B.C. Canada, 1993), p. 60.
[20] J. L. Chuma, *Physica analysis package, triumf.*
[21] M. Berz, *Cosy infinity version 7 reference manual, nscl technical report msucl-977, michigan state university* (1995).
[22] K. A. Assamagan, D. Dutta, and P. Welch, Tech. Rep., CEBAF Internal Report (1997).
[23] H. Überall, *Electron Scattering from Complex Nucli* (Academic Press, New York, 1971).
[24] D. H. Potterveld, Ph.D. thesis, California Institute of Technology (1989).
[25] B. W. Filippone et al., Phys. Rev. C **45**, 1582 (1992).
[26] J. Arrington et al., Phys. Rev. C **53**, 2248 (1996).
[27] D. B. Day et al., Phys. Rev. C **40**, 1011 (1989).
[28] P. G. Blunden, W. Melnitchouk, and J. A. Tjon, Phys. Rev. Lett. **91**, 142304 (2003).
[29] J. Arrington, Phys. Rev. C **69**, 032201(R) (2004).
[30] J. Arrington, Phys. Rev. C **68**, 034325 (2003).
[31] Y. C. Chen, A. Afanasev, S. J. Brodsky, C. E. Carlson, and M. Vanderhaeghen, Phys. Rev. Lett. **93**, 122301 (2004).
[32] J. Arrington and I. Sick, Phys. Rev. C **70**, 028203 (2004).
[33] J. Arrington, Phys. Rev. C **69**, 022201(R) (2004).