A MODULAR SOLID STATE DETECTOR FOR MEASURING HIGH ENERGY HEAVY ION FRAGMENTATION NEAR THE BEAM AXIS

C. J. ZEITLIN, K. A. FRANKEL, W. GONG, L. HEILBRONN, E. J. LAMPO, R. LERES, J. MILLER and W. SCHIMMERLING

Lawrence Berkeley Laboratory, Berkeley, CA 94720, U.S.A.

Abstract—A multi-element solid state detector has been designed to measure fluences of fragments produced near the beam axis by high energy heavy ion beams in thick targets. The detector is compact and modular, so as to be readily reconfigured according to the range of fragment charges and energies to be measured. Preamplifier gain settings and detector calibrations are adjustable remotely under computer control. We describe the central detector, its associated detectors and electronics, triggering scheme, data acquisition and particle identification techniques, illustrated by data taken with 600 MeV/u ^{56}Fe beams and thick polyethylene targets at the LBL Bevalac. The applications of this work to space radiation protection are discussed.

1. INTRODUCTION

ENERGETIC highly charged (HZE) nuclei ($Z \geqslant 2$) make up only a small fraction of the Galactic Cosmic Radiation (Simpson, 1983), but because they are densely ionizing, they are expected to contribute a substantial fraction of the dose equivalent to humans on expeditions outside the geomagnetosphere (Curtis *et al.*, 1991). In order to make informed decisions on radiation protection for humans in space, it is necessary to understand the biological effects of primary heavy ions and also to be able to accurately estimate the complex radiation field produced by interactions of heavy ions in shielding and tissue. To this end, for some time, ground-based experiments have been carried out to measure the physical and biological effects of HZE ions. These experiments have taken advantage of the fortuitous correspondence between the peak energies of the GCR heavy ion spectrum and the beam energies available at heavy ion accelerators such as the Bevalac at Lawrence Berkeley Laboratory. Iron is the heaviest ion present in significant numbers in the GCR, and is particularly prominent in the dose-weighted GCR flux.

Recently, experiments at the Bevalac have used high energy (up to 600 MeV/u) iron nuclei incident on a variety of biological targets. As a result of electromagnetic and nuclear interactions in beamline elements upstream of the target—including several 1/64″ of lead used to spread the beam—and in the target itself, the beam reaching the sample is inhomogeneous. It will be somewhat degraded and spread out in energy, and may undergo considerable fragmentation, as well. (For this reason, beam fragmentation has also been a source of concern in heavy ion radiotherapy.) Similarly, the primary GCR flux will be modified as it passes through spacecraft shielding and the astronauts' bodies. In both cases, it is important to know the characteristics of the resulting mixed radiation field. In the case of biological experiments, knowledge of the actual radiation field at the target is necessary for the accurate interpretation of results; for space radiation protection applications, fragmentation measurements will provide input data and benchmarks for computer simulations of different spacecraft shielding scenarios. Previous studies (Schimmerling *et al.*, 1989; Shavers *et al.*, 1990, 1993) used a simpler detector to measure fragmentation from ^{20}Ne ions in water, and compared the physics experiment results to both biology data and physics model predictions.

In this paper, we describe a compact detector system whose principal element is a 16-element stack of lithium-drifted silicon detectors, described in detail in Section 2. The intended purpose of the stack is to measure fragment charge over the entire range, i.e. from $Z = 1$ to $Z = Z_{beam}$ for moderately heavy beams (iron or similar) at a variety of energies and with a variety of target thicknesses. Isotopic resolution, while desirable, is not essential for this purpose. The stack detectors most important for particle identification subtended a solid angle of

approximately 2 millisteradians as seen from the target centers.

In Section 3, we compare techniques for identifying heavy particles, i.e. beam particles and heavy fragments. A χ^2 algorithm, which uses data from all hit detectors in the stack, is found to work well for heavy fragments. This algorithm has the significant advantage of providing identification of light fragments, down to $Z = 2$, unlike the other methods.

In Section 4, we discuss possible improvements to the system for future data taking.

2. MATERIALS AND METHODS

2.1. *Experimental layout*

A beam of ^{56}Fe ions was accelerated at the Lawrence Berkeley Laboratory BEVALAC with a nominal extraction energy of 600 MeV/u. We calculate from our data that the energy at the entrance to our experimental area was 578–579 MeV/u; the difference is attributable to ionization energy loss in upstream detectors used for beam control and monitoring.

Figure 1 is a schematic drawing of the experimental apparatus. The beam as delivered at the final dipole magnet (B2M1) is not large enough to irradiate typical samples uniformly. To spread the beam, lead foils of varying thickness were inserted into the beam upstream of B2M1. Beam particles were bent by B2M1 into the irradiation area, where the detectors were placed on two precisely aligned optical benches. Two scintillators, J1 and J2, were placed upstream of the target holder; a coincidence of hits in J1 and J2 in a 35 ns gate indicated the presence of a beam particle, and was required in all triggers.

In Table 1, we show calculated beam characteristics at various points along the beam line for the configuration without a polyethylene target. The probability for an iron nucleus to undergo a nuclear interaction is shown in the right-most column. These percentages are based on the energy-independent geometric cross section formula of Wilson *et al.* (1991). Approximately 3% of the incident beam particles interact in the lead foils; 2% interact in the aluminum exit window; and another 6% interact in the three scintillators used in triggering the experiment. In addition to the interactions upstream of the main detector systems, particles incident on the stack can interact in the detection medium. We calculate that the probability for iron to interact in silicon is approximately 1% per millimeter of silicon traversed. These interactions can adversely affect particle identification, and must be taken into account when calculating efficiencies. Also evident from Table 1 is the fact that a substantial amount of kinetic energy is lost in traversing the materials upstream of the detector box: for an incident energy of 578.5 MeV/u, even without a target, the energy of a primary at the first silicon detector in the box is degraded to 491 MeV/u. For the beam line elements up to and including PSD1Y, the total probability for primary iron to interact is 13.1% in the absence of a target. By comparison, we expect 16.1% of the incident iron particles to interact in 2 cm of polyethylene, 36% in 5 cm and 50% in 8 cm. Thus, in all cases, a significant amount of fragmentation of the beam occurs in materials other than the target.

The plastic scintillators used in this experiment are composed of CH, which is quite similar to polyethylene (CH_2). The elements J1, J2 and S have a combined thickness equivalent to about 7 mm of polyethylene, both in terms of energy loss and nuclear interaction cross section. In terms of energy loss, the silicon that comprises PSD1X and PSD1Y is roughly equivalent to 5 mm of polyethylene; the interaction length is equivalent to about 2.5 mm of polyethylene.

Table 2 shows the characteristics and locations of the detectors used in the experiment. The location of the elements is given relative to the center of the polyethylene target, which was constant from run to run. Position-sensitive silicon detectors are denoted by "PSD"; a detailed description can be found in (Wong *et al.*, 1990). These devices are lithium-drifted

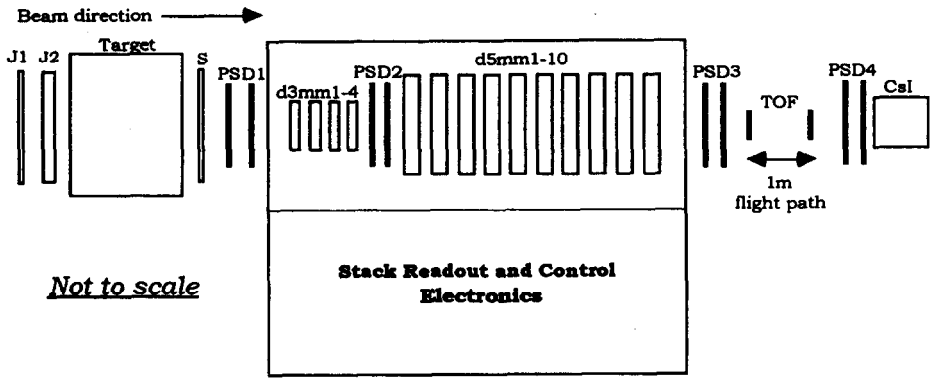

FIG. 1. Schematic representation of the detector elements used in the experiment, as seen from above. While not to scale, we have attempted to indicate the relative diameters and thicknesses of the detectors. See Table 2 for detailed information.

Table 1. Beam line elements; the incident energy, dE/dx in water, energy deposition and interacting fraction of the primary beam, with no polyethylene target, are shown

Element	E_{inc} (MeV/A)	dE/dx (H_2O) (keV/μm)	ΔE (MeV)	$\exp(-x/\lambda)$ (%)
Pb foils	578.5	180.5	2104.3	2.92
Al plate	539.9	182.1	614.8	2.04
MWPC	528.6	182.2	54.7	0.26
IC1	526.9	182.4	20.4	0.09
J1 scint	526.2	182.4	293.6	1.46
J2 scint	520.6	183.2	591.3	2.91
S scint	509.1	184.9	297.7	1.46
PSD1X	503.1	185.8	317.9	0.98
PSD1Y	497.4	186.7	319.4	0.98
d3mm1	491.0	187.7	1071.0	3.21

Targets were inserted between the J2 and S scintillators. Thus, while the values shown for elements up to and including J2 are correct for all target configurations, the values for elements S through d3mm1 change when a target is introduced.

silicon detectors with resistive equipotential stripes on the front surfaces. The deposited charge is divided in a position-dependent way into the two sides of the detector. Each PSD pair provided an (x, y) measurement. PSD1X and PSD1Y were used primarily for triggering (in coincidence with hits in the J1 and J2 scintillators). The active area of these detectors was considerably greater than that of the 3 mm thick detectors immediately downstream.

The preamplifier gain for the PSD1 detectors was set high to facilitate triggering on light particles. These channels saturated if the incident particle deposited more than about 130 MeV, corresponding to incident Z in the range 15–18.

In addition to the scintillators and silicon detectors in the stack, several detectors were placed further downstream, with the goal of identifying very light particles, down to $Z = 2$. These systems consisted of a PSD pair (PSD3), a time of flight (TOF) telescope consisting of two thin silicon detectors separated by about 90 cm of vacuum, another PSD pair (PSD4) and a 10 cm thick cesium iodide (CsI) crystal. Due to difficulties in calibration and determination of geometric acceptance, these systems were not used in our particle identification scheme.

For each target, data were taken over a period of 20–30 min, with delivered beam intensities of about 10^3–10^4 particles/s during a spill. Beam intensity was nearly uniform across each 1 s spill, of which there were 15 per minute.

2.2. *Electronics*

As shown schematically in Fig. 1, the silicon detector stack is contained in a modular box. One side contains mounts for up to 16 silicon detectors in a gas-tight enclosure; a continuous flow of dry nitrogen was maintained over the detectors during data taking. The other side has mounts for the associated charge-sensitive preamplifiers. The proximity of the preamplifiers to the detectors minimizes signal losses. Other onboard electronics include ASCII interface units (a receiver–transmitter pair) and a controller board which allows remote control of the stack electronics by software running on a Gateway 2000/286 PC. Communications with the box are via the computer's serial port. The modular design allows the stack to be easily reconfigured according to the energy and charge of the fragments to be detected; the ASCII interface units make it possible to calibrate the detectors and set the preamplifier gains remotely, under computer control.

Charge is collected from three outputs on the PSDs: two are position-sensitive, the third is proportional to the total ΔE in the detector. The other silicon detectors require only a single channel to read out ΔE signals. Accordingly, two types of preamplifier boards (1- and 3-channel) were built for the stack. The design allows PSDs and ΔE detectors to be treated interchangeably by providing a dedicated slot for the preamplifier card for each detector; one need only insert the appropriate board type.

Appropriate voltages were applied to the detectors so that they were fully depleted. The 3 mm detectors were depleted by an applied voltage of 400 V, the 5 mm by 600 V and the PSDs by −200 V. Typical leakage current for a 3 mm detector was 3 μA, about the same as the PSDs; for the 5 mm detectors, the

Table 2. Detector elements, their areas, thicknesses and positions relative to the target center

Element	Type	Area (cm^2)	Thickness (μm)	Location (cm)
J1	Scint	25.81	1588	−27.5
J2	Scint	25.81	3175	−14.5
S	Scint	25.81	1588	+11.8
PSD1X, Y	Silicon	15.21	900	25.4, 26.7
d3mm1-4	Silicon	5.07	3035–3155	44.8–50.8
PSD2X, Y	Silicon	15.21	1016–1041	52.8, 54.8
d5mm1-10	Silicon	20.27	4924–5300	56.8–74.8
PSD3X, Y	Silicon	15.21	900	96.5, 97.8
TOF1, 2	Silicon	0.79	200	111.7, 199.0
PSD4X, Y	Silicon	15.21	900	262.8, 264.1
CsI	CsI	8.24	10 cm	268.9

The areas given here are nominal; the active areas, particularly of the silicon detectors, may be smaller. The PSD areas are 3 times larger than the d3mm areas; this strongly influenced trigger rates.

typical current was about 7 μA. Charge-sensitive preamplifiers were mounted in the box, in proximity to the detectors so that signal losses were minimized. On each preamplifier board is circuitry (controlled by the PC) to select the amplifier gain from among three settings, and to put the board into calibration mode. In calibration mode, a tuned capacitor of known value is charged by a precisely known voltage, which is pulsed at a selected frequency. The calibration voltage was set using the PC, through a single 12-bit digital-to-analog converter (DAC) on the controller board. Each discharging capacitor (one per channel) injects a known amount of charge into the preamplifier, providing an absolute scale for energy deposition. The uncertainty in the injected charge is dominated by the error in the value of the calibration capacitor, which we estimate to be 1–2% for all channels. This is an important factor in the χ^2 particle identification algorithm.

Three gain settings were available for each preamp channel. For optimal detection efficiency over the range of fragment charges and energies expected from 600 MeV/u iron running, we chose to leave all the stack preamplifiers at their middle setting, so that the preamps saturated at approximately 7–8 GeV deposited energy. (A charged particle traversing a silicon detector creates one electron–hole pair for every 3.62 eV of deposited energy.) With this gain selection, each channel employs a calibration capacitor with $C = 26.64$ pF. The dc voltage applied to the capacitors ranges from 0 to 12 V. A ΔE of 1 GeV results in 44.4 pC of liberated charge, so that the upper limit of calibration charge which could be injected into a preamp was equivalent to about 7 GeV of deposited energy.

The PSDs external to the stack (i.e. PSDs 1, 3 and 4) were read out by preamplifiers with circuitry essentially identical to that of the 3-channel stack preamps, except for the control logic. By far the most important of these was PSD1, used in triggering the experiment.

Output signals from the preamplifiers were carried over 100–120 ft cables to variable-gain shaping amplifiers. Shaper gains were adjusted for each run to optimize the dynamic range of the silicon detectors. These units feature a wide range of gain adjustment—more than an order of magnitude—which proved extremely useful. For example, in the run with 2 cm of polyethylene, the energy deposition from primary iron was expected to be roughly 1100 MeV in d3mm1 and about 6800 MeV in d5mm5. While both preamps were kept at the same gain, we adjusted the shapers so that both of these signals produced maximum output voltages around 7–7.5 V. The rise time of the shaper signals was roughly 4 μs. The shapers also provide a "fast out" for each channel with rise times on the order of 50 ns. Several of these were used in triggering.

To digitize the outputs of the shaper amplifiers, CAMAC-based Ortec AD811 Analog to Digital Converters were employed. The AD811 is peak-sensing, with 11 bits of dynamic range (maximum count 2047) in an 8-channel, single-width package. The digitization time of these units is 80 μs. Most of the AD811 units had been modified to accept a full-scale input of 8 V, compared to 2 V in unmodified units. The ADCs, scalers and TDC channels used to read out the TOF system were read out by a VAX 3500 computer and written to 8 mm tape.

2.3. *Triggers*

In Table 3, we show trigger definitions and corresponding detection thresholds. The table also shows the number of events for which each trigger fired for the runs with 2, 5 and 8 cm of polyethylene in the beam line. A coincidence of hits in J1 and J2 defined an incident beam particle. The L1 and L2 triggers formed a highly redundant pair, as did H1 and H2. The L4 trigger (referred to as the CsI trigger) was independent of both PSD1 and S.

The data with no target have been omitted from Table 3 due to a hardware problem with a scaler which resulted in incorrect counting at high beam intensity. As explained below, we have accordingly restricted our data analysis to only use events taken during low-intensity spills. This was only necessary for the no-target data; beam intensity for the other runs was much lower.

In the analysis, almost all the events which fired one or both of the H triggers were found to have fired the L1/2 triggers as well. This is not surprising given the geometry of the detector setup. As will be discussed in Section 3, for a particle to be identified by the χ^2 algorithm, three of the four 3 mm-thick detectors in the box were required to have ΔE measure-

Table 3. Trigger definitions and number of times fired with three different targets

Trigger	Elements	Threshold	#eV, 2 cm	#eV, 5 cm	#eV, 8 cm
L1	J1 · J2 · PSD1X	$Z \approx 10$	614,677	828,475	1,033,342
L2	J1 · J2 · PSD1Y	$Z \approx 10$	603,569	815,476	1,016,802
L3	J1 · J2 · TOF	—	931	765	736
L4	J1 · J2 · PSD3 · CsI	—	6276	6877	6433
H1	J1 · J2 · S · d3mm1	$Z \approx 14$	214,785	281,086	357,560
H2	J1 · J2 · S · d3mm2	$Z \approx 14$	207,252	271,766	349,299

The actual number of triggers written to tape for any trigger/run combination is much less than the numbers shown here, due to dead time. See Section 2.5 for a detailed discussion of dead time and related issues.

ments consistent with one another. This occurred only if the incident particle was within one detector radius of the nominal beam line and moving along a trajectory very nearly parallel to the beam line. All such particles also hit the PSD1 detectors, and since the H thresholds were higher than those of the L triggers, the L triggers fired with high efficiency on such events. Conversely, events triggered by the L triggers do not necessarily fire the H triggers, for two reasons: (1) since the PSD1 detectors covered a larger area than the 3 mm detectors, a particle could hit PSD1 without hitting the d3mm1 or d3mm2; (2) even if a particle hit d3mm1 and/or d3mm2, the energy deposition was sometimes below the H threshold, while the deposited energy in the PSD1 detector pair was above threshold for L1 and/or L2.

As can be seen from Table 3, the ratio of L1 to H1 triggers stayed nearly constant from run to run, ranging from 2.86 to 2.95. From Table 1, we see that the nominal area of PSD1X (used in the L1 trigger) was exactly 3.0 times that of d3mm1 (used in the H1 trigger). This strongly supports two straightforward conclusions: (1) the L1/H1 ratio was primarily determined by the geometry of the detectors; and (2) the beam was large (i.e. σ of several centimeters) at this distance downstream of the Pb spreading foils.

An unfortunate consequence of the trigger scheme as implemented was the fact that many L1/2 triggers were fired by events in which no particle hit any of the 3 mm detectors in the box. Since hits in those detectors were required for particle identification, such triggers were for the most part counterproductive, as their only effect was to cause deadtime. However, the scheme did allow us to trigger on an event topology we had not planned on—events in which a relatively heavy particle hit PSD1X and/or PSD1Y, depositing sufficient energy to fire the trigger, but missed the 3 mm detectors, while an accompanying light particle (which by itself would not have caused any trigger to fire) was seen in the 3 mm detectors.

We have not endeavored to determine the thresholds for L3 or L4. Events which fired one or both of these triggers but no others were not used in the analysis. The CsI trigger proved valuable for measuring the efficiencies of the other triggers. The CsI crystal had a fairly small area and was placed far downstream, so relatively few particles reached it, as can be seen in Table 3. The low TOF trigger rate was due to the very small area of those detectors and the long lever arm between them.

Figure 2 is a simplified schematic diagram of the trigger logic, excluding the TOF trigger which is not of interest here. Note that a hit in S was not required for the L triggers, since it was expected that events with only light particles would not cause a hit above threshold in S. Various combinations of coincidences were read out by scalers, e.g. J1 · J2, which counted the total number of incident beam particles. The ratio of Masters to PreMasters gives a direct measurement

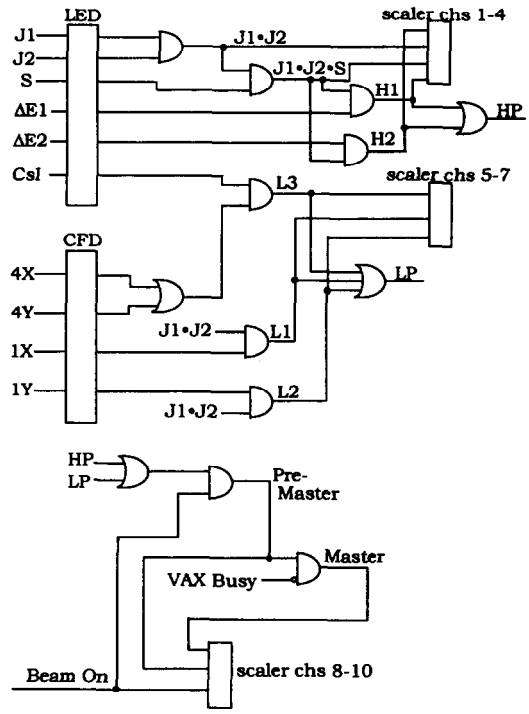

FIG. 2. Simplified schematic of the trigger logic.

of the experimental livetime, which varied substantially from run to run. A reliable livetime measurement is an essential part of the normalization procedure.

Figure 3 shows histograms of the energy deposition in d3mm1 for events which fired only the L1/2 triggers (solid) and for events which fired the H1 trigger (dashed). Events in the H1 plot could also have fired L1/2, and usually did; events in the L1/2 plot did *not* fire H1. The efficiency of H1 drops sharply at its threshold, near 400 MeV, while L1/2 continue to be efficient at lower ΔE_1. The few H1 events below this threshold are probably attributable to noise. In the region where H1 is efficient, there is a nearly identical (but suppressed by almost two orders of magnitude) distribution of events which only fired L1/2. These events arise from the few percent inefficiency of H1. A bump in the vicinity of 200 MeV is visible in the solid histogram; careful analysis shows that most of these are due to heavy particles at the edge (or perhaps beyond the edge) of d3mm1. Cuts on event geometry remove this artifact. An analogous plot of ΔE_2 for H2-triggered events is essentially identical. The data from the 5 cm run were used here; data from the other runs are quite similar. Note the presence of pile-up events in Fig. 3 above about 1500 MeV.

2.4. *Calibration*

Since the shaper amplifier gains were changed from run to run, calibration data were taken for each run

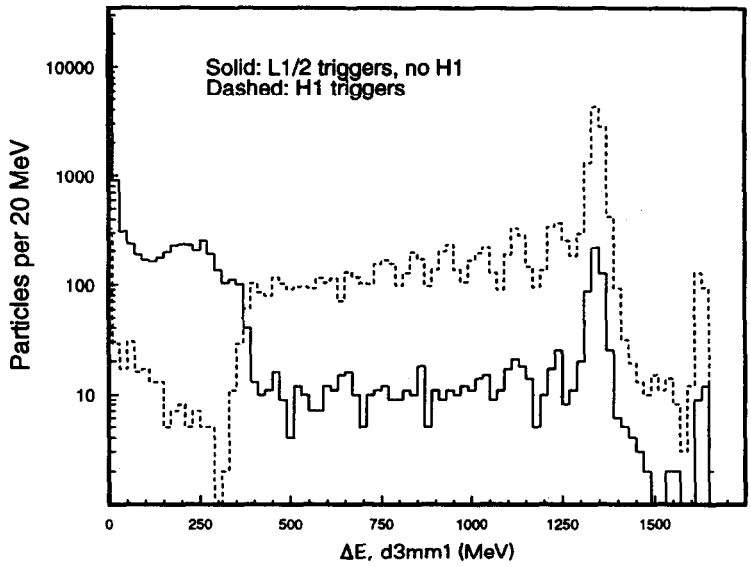

FIG. 3. ΔE spectrum seen in the first stack detector, d3mm1, for events in which the H1 trigger fired (dashed) and in which L1 and/or L2 fired without H1 firing (solid).

at points corresponding to ΔE of 0.2, 0.5, 1.0, 1.5, 2.5, 4.0 and 6.0 GeV. Depending on the channel and target, this meant that between two and seven points were usable. Figure 4 shows a plot of ADC counts versus the ΔE equivalent to the injected calibration charge for channel d5mm1, for the run with the 5 cm target; several features of this plot are characteristic of all our calibration data. Of primary importance is the high degree of linearity over a wide dynamic range which is clearly demonstrated here. In addition to the calibration points, each channel has a non-zero pedestal, determined from the data runs. In principle, a linear fit of the calibration points should go through the pedestal, i.e. the point at $\Delta E = 0$; such a fit (excluding the pedestal) yields an excellent χ^2. The fit line, extrapolated to $\Delta E = 0$, predicts a higher pedestal than is observed in the data, by about 10 counts as can be seen in the figure. If we add the pedestal point to the calibration points in the fit, the χ^2 becomes poor, although the slope changes by only about 3%. The cause of the discrepancy between predicted and observed pedestals is unknown; it does

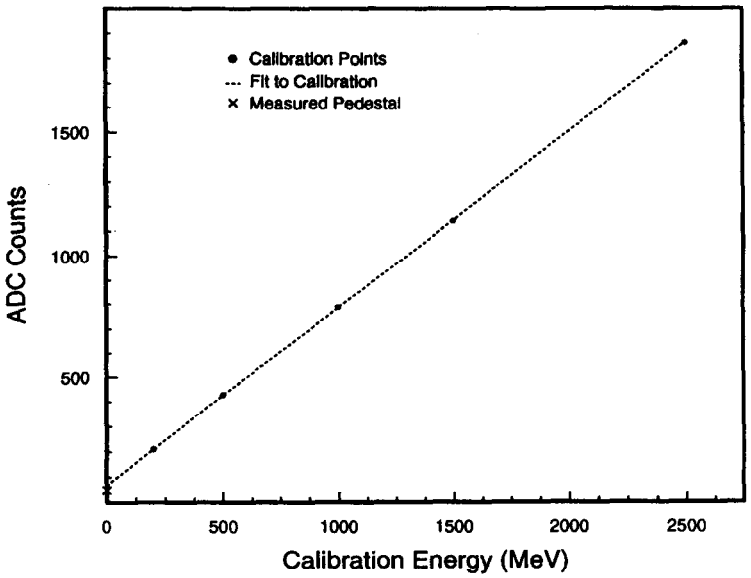

FIG. 4. Typical plot of ADC counts versus input charge. The input charge is expressed in equivalent deposited energy in silicon. While the calibration points are almost perfectly fit by a straight line, the extrapolated line does not match the pedestal seen in data runs.

not present a major problem in particle identification. It is not a product of, nor is it affected by, the 1–2% uncertainty in calibration capacitor values.

One possible explanation of the pedestal discrepancy is a small shift of the applied dc calibration voltage relative to ground. Another possibility is that the preamps and/or ADCs behave non-linearly at the lowest ΔEs. The integral non-linearity of the Ortec AD811 is not specified for the lowest 1% its range, according to the manufacturer's technical data sheet. It thus appears quite likely that ADC non-linearity is a contributor (possibly dominant) to the observed pedestal discrepancy. If the former explanation is correct, then a constant should be added to the calibration ΔE values in the fits; if the latter explanation is correct, one way to address it is to perform two fits per channel, one between the pedestal and the lowest calibration point, and one using only the calibration points (ignoring the pedestal). The data analysis was performed both ways, with negligible differences in the results. The results presented here make use of the latter method.

2.5. *Data acquisition, timing and pile-up*

Data were read from the CAMAC controller using a Micro-programmed Branch Driver (MBD) and a VAXserver 3500. The VAX 3500 was used to run the data acquisition software, and the MBD was used as an interface between the 3500 and the CAMAC controller. The main source of computer dead time came from data reads by the MBD. The MBD uses a little less than 8 μs to read each data word (2 bytes/word), and each event read out 79 words. Thus, the total amount of dead time per event due to a data event read was about 600 μs. Approximately 90 μs were added to dead time at the beginning of each event to allow the CAMAC modules to process the incoming signals. Additional time was needed to clear the old data from the CAMAC modules and prepare the modules for the next read, so the total amount of dead time per event was about 700 μs. Each event was stored in buffers located in the MBD. Once a beam spill was over, the data stored in the buffers was transferred to the 3500, thereby eliminating the contribution to the dead time from transfer times between the MBD and the 3500.

The 80 μs needed to process each event in the CAMAC ADCs created a small window in which a signal due to an event from a different beam particle could be processed along with the original signal, thus giving a finite probability of "pile-up" in the event signals. Beam currents were held low in order to minimize the amount of pile-up in the ADCs. We estimate that the percentage of events contaminated by pile-up was 24% for the no-target run, 14% for the 2 cm target, 9% for the 5 cm target and 4% for the 8 cm target. Pile-up events were explicitly expunged in the data analysis.

2.6. *Position measurements*

The silicon position-sensitive detectors provide three ΔE signals each, two of which are sensitive to position, while the third signal is proportional to total ΔE as in the non-position sensitive silicon detectors. Each PSD is capable of measuring x or y, depending on its orientation. For an x orientation, we designate the position signals E and W; for y, T and B. The measured positions u and v (closely related to x and y) are given by:

$$u = r \sin\left(\frac{\pi}{2}\frac{q^E - q^W}{Q_u}\right)$$

and

$$v = r \sin\left(\frac{\pi}{2}\frac{q^T - q^B}{Q_v}\right), \tag{1}$$

where r is the radius of the detector and Q is the total charge (directly proportional to total ΔE). In principle, $Q_u = q^E + q^W$ and similarly $Q_v = q^T + q^B$. The position signals are processed in an identical manner to all the other ΔE signals, and we make the correspondence $q \propto \Delta E$. Thus, in practice

$$u = r \sin\left(\frac{\pi}{2}\frac{\Delta E^E - \Delta E^W}{\Delta E^E + \Delta E^W}\right)$$

and

$$v = r \sin\left(\frac{\pi}{2}\frac{\Delta E^T - \Delta E^B}{\Delta E^T + \Delta E^B}\right). \tag{2}$$

In Wong *et al.* (1990), u and v were related to x and y in a way which eliminated distortions, which were ascribed to nonuniformity in the thickness of the resistive layers across the detector surfaces. The functions used to define the conformal transformations were fitted to calibration data taken with a special grid collimator. We did not perform a similar calibration, and thus are unable to make the distortion corrections, which appear to be most significant near detector boundaries. Hence, we make the approximations $x = u$ and $y = v$. We define $R = \sqrt{x^2 + y^2}$.

Figure 5 shows the x and y distributions as measured in PSD2 for events in the 2 cm run which had mutually consistent hits in at least three of the first four detectors. Data from the other runs were very similar. Both distributions are quite broad and centered approximately on 0. Because the beam was large—with a full-width at half maximum in both x and y as large or larger than the diameter of the detector—the distributions do not represent measurements of the beam size.

Figure 6 shows a scatter plot of y vs x as measured at PSD2 for the same data set. Structure in the form of alternating light and dark bands in y is visible. This is a non-critical detector artifact, which implies that the resolution in y is no better than 1–2 mm. From this figure, it is evident that the beam was distributed quite uniformly over the d3mm detectors. The

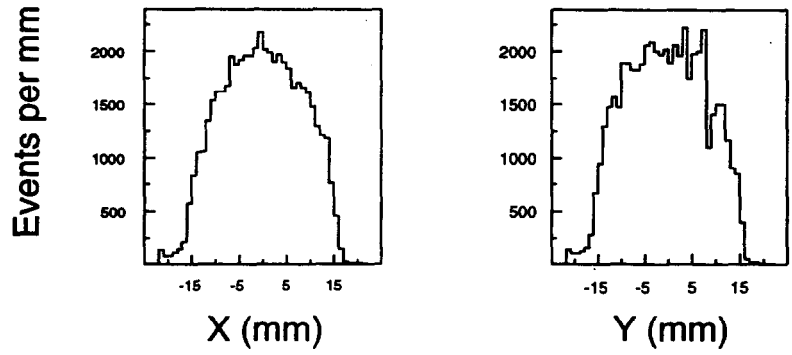

FIG. 5. Distributions of x and y as measured in PSD2 for the 2 cm run.

densely populated region of this plot is contained in a circle with an approximate radius of 16 mm, larger than the 12.7 mm radius of the d3mm1–4 detectors. There is also some distortion (as expected), which causes the populated area to deviate from being perfectly round. The most important conclusion we draw from this plot is that simply requiring good correlations in ΔE among the first four detectors is virtually equivalent to imposing a geometrical cut. Quantitatively, we find that 94% of the events in this plot are contained within a measured radius of 16 mm.

3. PARTICLE IDENTIFICATION

3.1. *Introduction: correlation methods*

Each detector in the stack which is traversed by a charged particle yields a measurement of the energy deposited, ΔE_i for the ith detector. Data were taken using three thicknesses of polyethylene—2, 5 and 8 cm—and also with no polyethylene in the beamline, allowing us to measure the fragmentation of iron in the beamline materials other than the target. Heavy fragments were identified by three techniques, which proved to be mutually consistent to

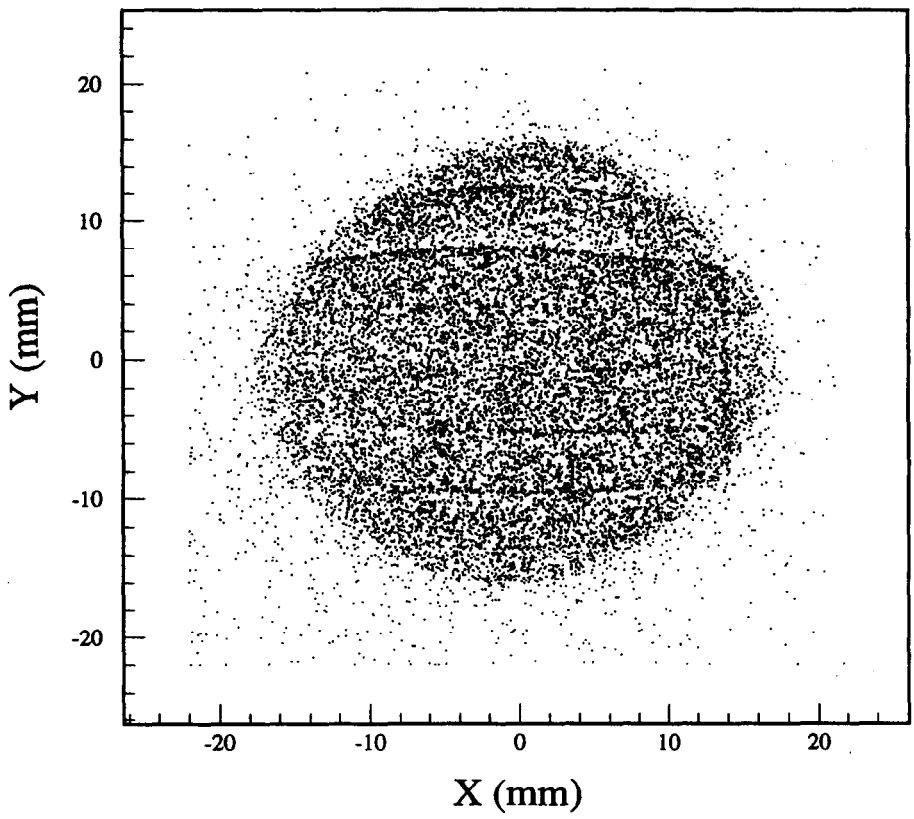

FIG. 6. Scatter plot of y vs x for the 2 cm run. By simply requiring good correlations in d3mm1–4, a densely populated disk is visible.

a high degree: (1) ΔE_1 vs E_{tot}, the sum of the ΔE_i; (2) ΔE_1 vs ΔE_2; and (3) a χ^2 algorithm which made use of the energy deposited in each hit silicon detector. For lighter fragments ($Z \leqslant 16$ or so, depending on the target thickness), the graphical techniques (methods 1 and 2) were impractical, and thus we relied solely on the χ^2 determination of fragment Z. The excellent agreement of the χ^2 algorithm results with those from the other methods gives us confidence that light particles are correctly identified. In addition to finding Z, the χ^2 algorithm also measures the kinetic energy, T, of incident particles.

3.1.1. ΔE *vs total* E. Particles incident on the detector stack can be identified in a variety of ways. In Fig. 7, we show ΔE_1 vs E_{tot} for the run with a 2 cm polyethylene target. Events were required to have well-correlated hits in at least three of the first four detectors in the stack. Obvious clusters corresponding to particles of a single Z, with varying total energies, are apparent. Clusters are well-defined down to about $Z = 16$; each cluster has a tail at more or less fixed ΔE_1, extending to lower E_{tot} values. These tails arise from nuclear interactions in the detector box. Data taken in the other target configurations show qualitatively similar clusters in ΔE_1 vs E_{tot} plots. Clusters of events at lower Z become better defined as target thickness increases. Figure 7 shows a "kink" around $Z = 17$, owing to the fact that most particles with this charge or less are not stopped in the stack, and hence the true total energy is not measured.

Table 4 shows the range in silicon of the primary iron for the four runs. The probability for an iron ion to undergo a nuclear interaction in silicon is roughly 1% per mm. For all but the 8 cm run, incident iron particles have significant probability for undergoing an interaction in the detector box. The interaction probabilities for heavy particles near iron (e.g. manganese) are only slightly lower. This explains the prominence of the tails seen in the ΔE_1–E_{tot} plots.

3.1.2. ΔE_1–ΔE_2 *correlations*. In Fig. 8, we show ΔE_1 vs ΔE_2 for the 2 cm run. The requirement of good correlations in the first few detectors comes into play in this plot, in that the tails of events with poor correlations have been excluded. Again, clusters of events can be seen down to about $Z = 16$. The no-target data is quite similar. As target thickness increases, so too does the spread of possible energies for fragments of a given Z. This causes the clusters in the ΔE_1–ΔE_2 plots to spread out and eventually

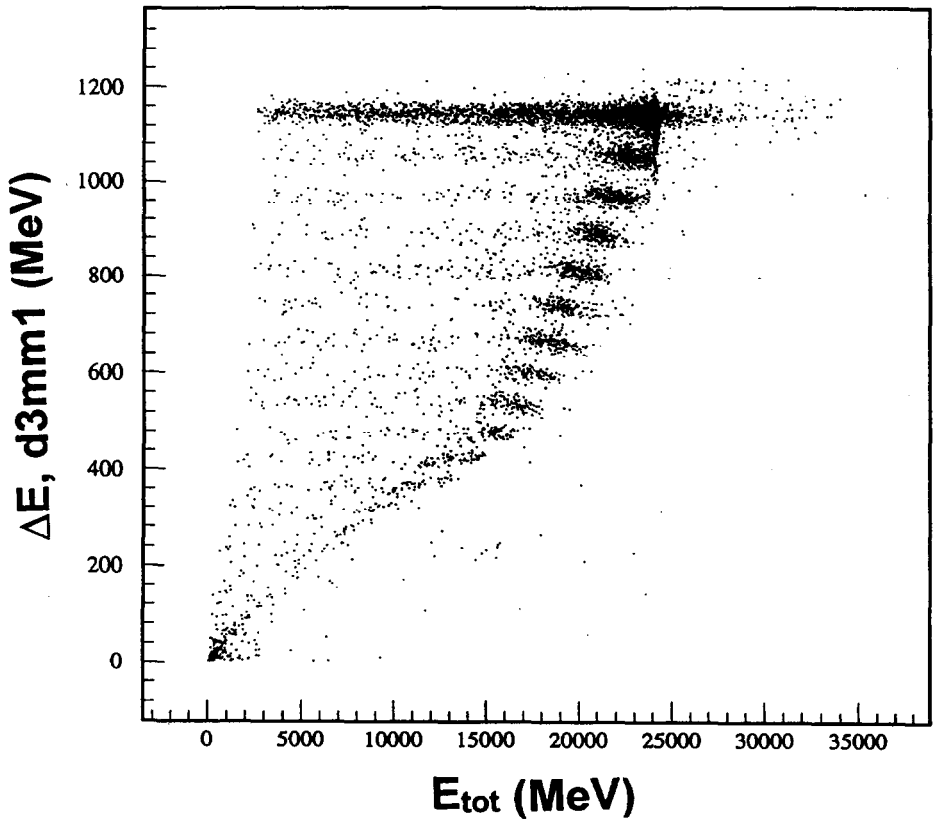

FIG. 7. ΔE_1 (vertical axis) vs E_{tot} (horizontal axis) for the data taken with 2 cm of polyethylene. Particle species show up as distinct clusters of events, which are well-defined down to about $Z = 16$. The only cut on these data required good correlations in three of the first four stack detectors.

Table 4. Incident energy of primary iron on the first detector in the box for the four target configurations

	Polyethylene target thickness			
	None	2 cm	5 cm	8 cm
E_{inc} (MeV/A)	492	424	314	173
Range (mm)	46.9	39.6	24.4	9.0

The incident energy determines the residual range in silicon, shown in the bottom row of the table.

overlap. Thus, in the 5 cm data, clusters can only be seen down to about $Z = 20$; and in the 8 cm data, no obvious clusters are visible. The events in Fig. 8 are the same sample as in Fig. 7.

3.2. *The χ^2 algorithm*

3.2.1. *Calculations.* To take advantage of the fact that many detectors are hit on each event, we have created an algorithm for identifying particles which uses each valid ΔE_i in an event. Unlike the correlation methods described above, this algorithm works (at least in principle) over the full range of Z of the incident particles. We used the Bethe–Bloch formula (Particle Data Group, 1992) to calculate the energy deposited in each detector for a wide range of isotopes and energies incident on d3mm1. Ignoring the δ term (due to the density effect), and important only at higher energies) the Bethe-Bloch equation gives:

$$-\frac{dE}{dx} = 0.3071\left(\rho\frac{z}{a}\right)\frac{Z^2}{\beta^2}\left[\log\left(\frac{2m_e\beta^2\gamma^2}{16z^{0.9}}\right) - \beta^2\right], \tag{3}$$

for a particle of charge Z and velocity β traversing a material of density ρ, atomic number z and mass number a. This gives dE/dx in units of MeV/cm when ρ is given in g/cm^3; m_e is in units of eV. Strictly speaking, the equation is valid only for an infinitesimal thickness dx. What is actually measured is the integral of dE/dx over finite thicknesses. As equation (3) is not readily amenable to analytic integration, we must integrate numerically with a finite step size. In each step, the energy loss results in a decrease in β (i.e. β is not constant across the step), making the integral sensitive to the choice of step size. A step size which is too large will underestimate the integral. We found that a 1 μm step size was adequate for these calculations, in that the results were virtually unchanged compared to calculations with a 0.5 μm step size, but systematically higher than those using a 2 μm step.

We used equation (3) to calculate a table of ΔE_i for 117 isotopes with Z from 1 to 26 and kinetic energy in the range 150–550 MeV/u at 1 MeV/u

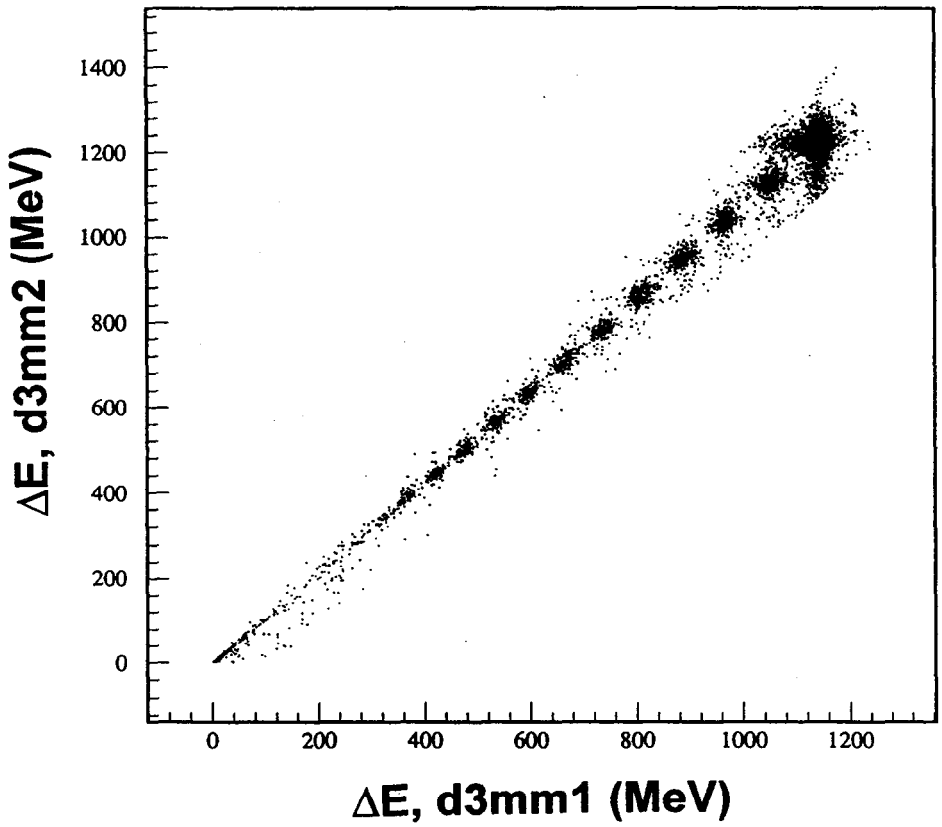

FIG. 8. ΔE_2 vs ΔE_1 for the same data sample as in Fig. 7. As in that plot, clusters of events correspond to particle species.

intervals. For a given Z, A and T (the kinetic energy of the particle as it enters d3mm1), we calculated dE/dx in silicon according to equation (3). The incident energy was used to determine the starting β. In succeeding steps, β was decreased according to the relation $d\beta/dx = (1/M\beta\gamma^3)(dE/dx)$. The gaps between detectors were taken to be filled with N_2 at 1 atm. The result of these calculations is a data file stored on disk for use by the particle identification routines. The calculations use the measured detector thicknesses, and we assume that the first 99% of the thickness is active (i.e. the detectors are taken to have dead layers equal to 1% of their total thickness, on the downstream side).

3.2.2. *Chi-squared definition.* For each event, we have a set of measurements, denoted by ΔE_i(meas). We also have the calculated values of ΔE described above, denoted by $\Delta E_i(Z, A, T)$ to indicate that the energy deposited in the ith detector is a function of the charge, mass and kinetic energy of the incident particle. We then define:

$$\chi^2(Z, A, T) = \sum_i \frac{(\Delta E_i(\text{meas}) - \Delta E_i(Z, A, T))^2}{\sigma_i^2} \quad (4)$$

where the sum runs over all detectors with usable hits. (Methods for classifying hits as usable are discussed below.) Roughly speaking, the algorithm identifies particles by finding the combination of (Z, A, T) which yields the lowest χ^2. There are several uncertainties which contribute to σ_i; these are described in Section 3.2.8.

3.2.3. *Interactions in the detector box.* Before the χ^2 algorithm proceeds, a crude search for interactions in the detector box is conducted. The search is applied only to events where at least the first eight box detectors have hits above threshold. It begins with the seventh detector (the first of the 5 mm thick units), and compares $\Delta E_7/\Delta x_7$ (i.e. the average dE/dx in the detector) to $\Delta E_9/\Delta x_9$. In the absence of a charge-changing interaction, the value in the latter detector should be higher. If it is measured to be less than 95% of the value in the earlier detector, the detector in between (#8) is examined. If it, too, has an average dE/dx less than 95% of that in detector 7, an interaction is considered to have occurred. If no interaction is found in this comparison, we move one detector downstream and iterate the procedure. Iteration continues until an interaction is found or the stopping detector is reached. Events with interactions need not be rejected; it is only necessary to discard detectors downstream of the interaction when calculating χ^2. The search described here is complemented by checks on the hits in the first four detectors, described in the next section.

Table 5 shows data from the first eleven detectors for an event in which an ion with $Z = 20$ enters the box and interacts in the ninth or tenth detector. The χ^2 method found the theory values which best agreed with the measured values, with $Z = 20$, $A = 40$ and $T = 424$ MeV/u. Theory and measurement agree well up to the ninth detector, the likely site of a nuclear interaction. χ^2 for this event was 2.99 for the first eight detectors. Note too that the average dE/dx steadily increases (within errors) as the particle traverses the detector material, as expected.

3.2.4. *Candidate selection.* Given the large number of (Z, A, T) combinations, it is impractical to test every event against every possibility. Instead, we find a limited number of candidates and calculate only those χ^2 values. To find candidates, we look for the first detector which has a ΔE compatible with other detectors. We require that at least three of the first four detectors in the stack must be at least somewhat mutually consistent. Events are rejected if both the first and second detectors fail these cuts. This removes events which lie outside our fiducial region, as well as events in which the particle undergoes a nuclear interaction in one of the first few detectors. Obviously, more heavy particles will fail this cut than will light particles; it is approximately a 10% effect for iron.

Given an event where at least three of the first four detectors are usable, we search at intervals of 10 MeV/u in T, with the criterion that $\Delta E_{\text{first}}(Z, A, T)$ is within $\pm 5\%$ of ΔE_{first}(meas). (In all cases, first = 1 or 2.) If ΔE_{first}(meas) < 100 MeV, the $\pm 5\%$ cut is relaxed to ± 5 MeV. If no suitable candidate (Z, A, T) is found, the event is rejected. Typically, 20–40 candidates are found. Only physically reasonable values of T (depending on target thickness and Z) are considered.

After one or more candidates have been selected, χ^2 is calculated for the candidate (Z, A, T) and the next nine values of T (i.e. at intervals of 1 MeV/u instead of the 10 MeV/u used in the candidate search). After all relevant values of χ^2 have been calculated, we pick the lowest, and calculate the probability $P(\chi^2, N)$ of finding a better χ^2 with N degrees of freedom. If P is less than 90% (a rather

Table 5. Values of ΔE measured (m), theory (th), calculated error, and average dE/dx, versus detector number for an event with $Z = 20$

Det #	1	2	3	4	5	6	7	8	9	10	11
ΔE_i(m)	678	701	698	748	266	251	1380	1377	1577	446	391
ΔE_i(th)	669	704	715	756	254	250	1348	1383	1498	1706	1985
σ_i	16	16	16	17	15	14	37	37	43	15	14
$\Delta E_i/\Delta x_i$	2233	2278	2283	2429	2617	2410	2706	2692	3078	867	753

All quantities are in units of MeV, except for dE/dx, which is in MeV/cm. The theoretical values were determined by the χ^2 algorithm as described in the text.

arbitrary choice), we consider this a "good" (Z, A, T).

3.2.5. *Unusable detectors.* After all the χ^2 values have been calculated, we search—starting with the last hit detector—for detectors which contribute either (1) 60% or more of the total χ^2, or (2) an individual χ^2 of 3.85 or more, which corresponds to a probability of less than 5%. If such a detector is found, the χ^2 analysis is repeated from the start, with the offending detector ignored. If no good χ^2 is found in the first iteration, the procedure is repeated. The algorithm continues iterating until either a good χ^2 is found or only three detectors remain usable. Three of the first four detectors are required to be used in all cases. Note that this procedure skews the distribution of χ^2 towards low values.

Table 6 shows an example of an event where two detectors—in this case, #4 and #10—were removed from the analysis. With these detectors excluded, the analysis found $Z = 22$, $A = 50$ and $T = 312$ MeV/u, with an excellent χ^2 of 0.8 for 8 degrees of freedom. Clearly, ΔE_4 is inconsistent with its neighbors, and with the value expected by the analysis, by several standard deviations. If detector 4 were included in the calculation of χ^2, the event would fail any reasonable χ^2 cut and would be rejected. Given the excellent χ^2 found with the other detectors, it seems most prudent to simply exclude the inconsistent detector and keep the event. At present, we do not have a definitive explanation for this type of failure, which is observed in roughly 10% of our events. The problem is not confined to any particular channel or channels, and we have found no correlating factors.

The inconsistency of ΔE_{10} with the calculated value in Table 6 may be attributable to the difficulties associated with modeling the energy loss at the extreme end of the particle range. For example, we have not incorporated into our model the very thin plating on the detectors, nor have we modeled the (also very thin) dead layers on the upstream sides of the detectors. For an event such as this one where, evidently, the particle barely penetrates the last detector, those omissions may result in a significant difference between the measurement and the calculation. Fortunately, as demonstrated by this event, this complication is easily managed by discarding the information from the last detector.

3.2.6. *Ambiguity in* Z. Our highest priority in particle identification is to make the correct determination of Z. For the great majority of events, even if there are several good (Z, A, T) combinations found by the χ^2 algorithm, all of them have the same Z. In those cases, we simply chose the values of A and T which yielded the lowest χ^2. If there were two or more good Zs found (recall that by "good", we mean $P(\chi^2, N) < 0.9$), it becomes necessary to choose one Z. In Figures 9 and 10, we show plots of ΔE_1 (Fig. 9) and ΔE_2 (Fig. 10) for each of the four runs of interest (no target, 2, 5 and 8 cm targets). In all but the 8 cm target runs, both ΔE_1 and ΔE_2 show several clear peaks, each of which corresponds to a single element. Based on these plots, we make cuts to determine Z_1 (i.e. Z based solely on ΔE_1) and Z_2 (based on ΔE_2). If an event falls into one of the valleys between the peaks, we do not assign Z for that detector. We then apply the following criteria to resolve events with one or more good Zs from χ^2:

(1) If $Z_1 = 0$, then $Z_d = Z_2$ (if Z_2 was assigned).
(2) If $Z_2 = 0$, then $Z_d = Z_1$ (if Z_1 was assigned).
(3) If $Z_1 = Z_2$, then $Z_d = Z_1 = Z_2$.
(4) If $Z_1 \neq Z_2$, then $Z_d = \max(Z_1, Z_2)$.
(5) If one of the good Zs from χ^2 equals Z_d, then $Z = Z_d$, if not, choose the highest Z.

In all cases, we are able to assign Z with a high level of confidence. In the following section, we will show that in general, we get excellent agreement between Z as determined by χ^2 and as determined by the correlation methods (ΔE_2 vs ΔE_1 and ΔE_1 vs E_{tot}), in the limited range where those methods are applicable.

3.2.7. *Features of the* $\Delta E(d3mm1)$ *plots.* Here, we briefly digress to discuss the ΔE_1 plots shown in Fig. 9. For all runs, the spectra have minima around 100–150 MeV, the approximate vicinity of the L1/2 trigger threshold. All runs show peaks for iron and—except for the 8 cm data—several other elements.

In the no-target run, we note the lack of a distinct peak for $Z = 25$. Because of the relatively small number of interactions in this run, the tail of the iron distribution obscures the manganese peak. For the 2 and 5 cm runs, distinct Mn peaks appear. As the target thickness increases, the elemental peaks spread out, until they completely disappear in the 8 cm data.

The data in these plots are those events which were identified by χ^2, using all trigger types. For

Table 6. Values of ΔE measured (m), theory (th), calculated error, and average dE/dx, versus detector number for an event with $Z = 22$

Det#	1	2	3	4	5	6	7	8	9	10	11
ΔE_i(m)	946	1024	1060	886	398	396	2249	2602	4625	527	0
ΔE_i(th)	949	1016	1055	1147	394	393	2232	2638	4674	789	0
σ_i	21	23	24	20	18	18	60	69	122	23	0
$\Delta E_i/\Delta x_i$	3116	3327	3467	2878	3918	3808	4412	5086	9024	1023	0

Note the excellent agreement between measured and theoretical ΔE values for all detectors except 4 and 10.

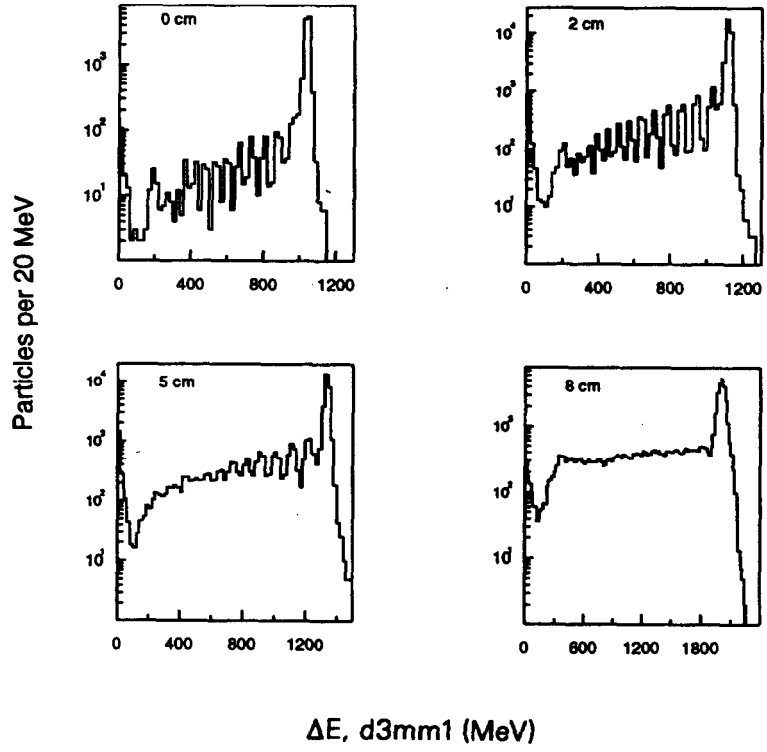

FIG. 9. ΔE_1 spectra for the four runs. The data consist of events for which the H1/2 triggers fired and an acceptable χ^2 was found.

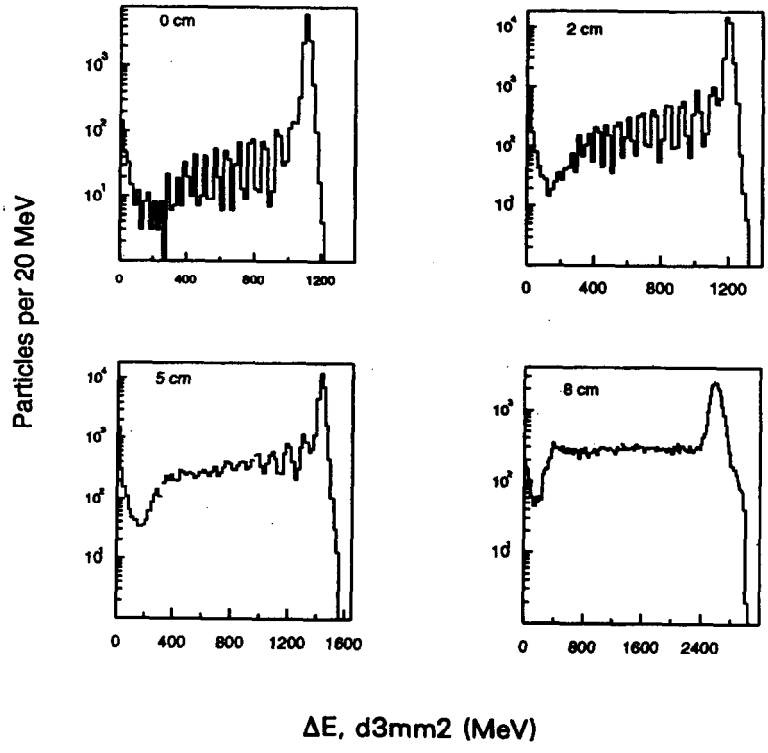

FIG. 10. ΔE_2 spectra for the four runs. The data are the same as in Fig. 9.

comparison, consider Fig. 3 and the bottom left plot in Fig. 6. Both show ΔE_1 in the 5 cm data. The most striking difference between them is the suppression of events at ΔE_1 below about 250 MeV. The count of events in that region drops by roughly an order of magnitude when we require the event to be identified by χ^2. As mentioned in Section 2.3, most of these events were heavy particles outside the fiducial area; as shown in Section 2.6, cuts requiring good correlations among the first four detectors—made in Fig. 9 but not in Fig. 3—removes those, and explains much of the difference. Additionally, most of the events in the lowest bin of ΔE are within a few counts of pedestal; these are likely to have triggered on particles which hit PSD1 (firing L1/2) but missed d3mm1, etc. These too are outside the fiducial region, and are legitimately removed when χ^2 fails to identify them.

3.2.8. *Contributions to σ.* Several factors contribute to σ_i; we take the RMS sum of these factors, as follows:

$$\sigma_i = \sqrt{(\sigma_{\rm cal})^2 + (\sigma_{\rm thk})^2 + (\sigma_{\rm BB})^2 + (\sigma_{\rm ped})^2}.$$

All the terms except the pedestal error are taken to scale with the measured ΔE. Specifically, we take the calibration error to be $\sigma_{\rm cal} = 2.0\% \times \Delta E$; the uncertainty in detector thickness and dead layer depth is given by $\sigma_{\rm thk} = (1.0\%,\ 1.5\%) \times \Delta E$, where the 1% error applies to the 3 mm detectors, and the 1.5% to all others; and we take the uncertainty in the Bethe–Bloch calculation to be $\sigma_{\rm BB} = 1.0\% \times \Delta E$. The pedestal errors are taken to be 10 MeV for the 3 mm detectors, and 16 MeV for the PSD2 elements and the 5 mm units. These numbers reflect the discrepancy between fit and measured pedestals, as described in Section 2.4.

The pedestal error ultimately limits our ability to separate species at low Z. At the energies typical for this experiment, a helium ion will deposit about 7 MeV in a 3 mm thick detector, and lithium will deposit about 16 MeV. The 10 MeV uncertainty in the pedestal dominates the error term at these ΔEs, and allows about 1σ separation in each detector between helium and lithium. With several detectors hit, this should be sufficient to separate them, and extensive visual event scanning shows that $Z = 2$ and $Z = 3$ events are correctly identified for the most part. Signals from protons are not sufficiently above the noise to be reliably identified.

Overall, we have chosen to err on the side of making the uncertainties too large rather than too small. This means that our χ^2 distributions will be skewed towards lower values, and more events will be categorized as "ambiguous" by the algorithm than would be the case if the errors were set smaller. The beneficial aspect is that more events will pass the $P(\chi^2, N)$ cut than would otherwise be the case. The fact that the analysis removes detectors which contribute large values of χ^2 to the total (as explained in Section 3.2.5) also skews the distributions to lower values.

Table 7. Number of events on tape, passing the first pass filter and the fraction passing these minimal cuts

	No target	2 cm	5 cm	8 cm
Events on tape	65,970	160,429	173,981	200,550
Passed filter	32,954	74,687	83,332	88,866
% passed	50.0	46.6	47.9	44.3

3.3. *Comparison between PID methods*

3.3.1. *First pass filter.* The data for each run were passed through a simple filter program which performed three functions:

(1) events in which there was no hit above pedestal in both d3mm1 and d3mm2 were rejected;

(2) run-dependent calibration constants were applied;

(3) data were written to disk in a compact format.

Table 7 shows the effect of the filter for the four runs of interest. Despite the very minimal requirements, fewer than half the events pass. We attribute this to the large number of L1/2-triggered events in which no particle hit the smaller d3mm detectors. For the no-target run, we limited the selection to spills in which there were fewer than 2^{14} (= 16,384) pre-masters, as a hardware problem with the scaler rendered higher count rates unreliable.

3.3.2. *Second pass analysis.* In the next analysis pass, the data were analyzed by the χ^2 algorithm. Table 8 shows the number of events which passed at this stage. For an event to pass, the algorithm had to find at least one χ^2 such that $P(\chi^2, N) < 0.9$. No explicit cuts were made on position. Likely pile-up events were rejected by cuts on ΔE_1, ΔE_5, and ΔE_7 which required these to be less than carefully-chosen cut values. (The cut values were target-dependent. These detectors—d3mm1, PSD2X and d5mm1—were chosen because each is the first of a set of detectors with the same radii. See Table 2 for details.) At least three of the four d3mm detectors were required to have been deemed usable in the calculation of χ^2. In addition, only events which fired one or more of the H1/2, L1/2 triggers were accepted.

The fraction of passing events in the no-target run is much smaller than that for the other runs. This can, for the most part, be explained by the high fraction of pile-up events found in that run. Given that the beam current was much higher in that run than any other, this is not a surprising result.

The main analysis program which performed the χ^2 analysis linked in the HBOOK routines (Brun and Lienart, 1987) written at CERN. Using

Table 8. Number and fraction of events passing the χ^2 analysis

	No target	2 cm	5 cm	8 cm
Passed filter	32,954	74,687	83,332	88,866
Passed χ^2	13,864	46,785	51,210	54,721
% passed	42.1	62.6	61.5	61.6

HBOOK, we created files with "ntuples" of the events which passed cuts. The ntuples are readable by the data analysis package PAW (Brun *et al.*, 1989). Among its many useful features, PAW allows the user to make graphical cuts on the data. In each ntuple, we stored the ΔE readings from all the silicon detectors, ADC readings from the CsI and trigger information. Also stored in the ntuple were the most germane results from the χ^2 analysis: Z, A, T $P(\chi^2, N)$ and the ΔE_i (theory) values.

3.3.3. *Comparison of PID results.* Figure 11 shows ΔE_1 vs E_{tot} for H1/2 triggers in the 5 cm run. Also shown in this plot are graphical cuts ("gates") used for determining Z. The precise locations of the gates are somewhat arbitrary; we tried to put them midway between dense clusters of events.

We have also used PAW to make gates in the ΔE_1–ΔE_2 plots. In the no-target and 2 cm runs, this method has about the same range of applicability as the ΔE–E_{tot} method. However, in the 5 cm data, only a few clusters are visible in the ΔE_1–ΔE_2 plot, and the method is useless in the 8 cm data.

Tables 9–12 show comparisons of event totals using the three different methods on the various data sets. The comparisons were restricted to events which fired the H1 and/or H2 triggers, since those events were most likely to have good geometry. The agreement is generally excellent, within a few percent. By far the most significant disagreements are to be found in the counts of $Z = 25$ events between methods: in cases where it is applicable, the ΔE_1–ΔE_2 method finds significantly fewer events than the other two methods. A similar trend is apparent for $Z = 24$ as well. We believe that ΔE_1–ΔE_2 is most accurate in these cases, since the other two methods appear to occasionally misidentify Fe as Mn. Since there is so much more Fe than Mn, a small percentage of misidentified events can result in a significant, and incorrect, increase in the number of identified Mn. The same logic applies to events misidentified as $Z = 24$.

One type of event which could easily be misidentified (especially by χ^2) as Mn or Cr is an incident Fe which undergoes a peripheral interaction in d3mm1,

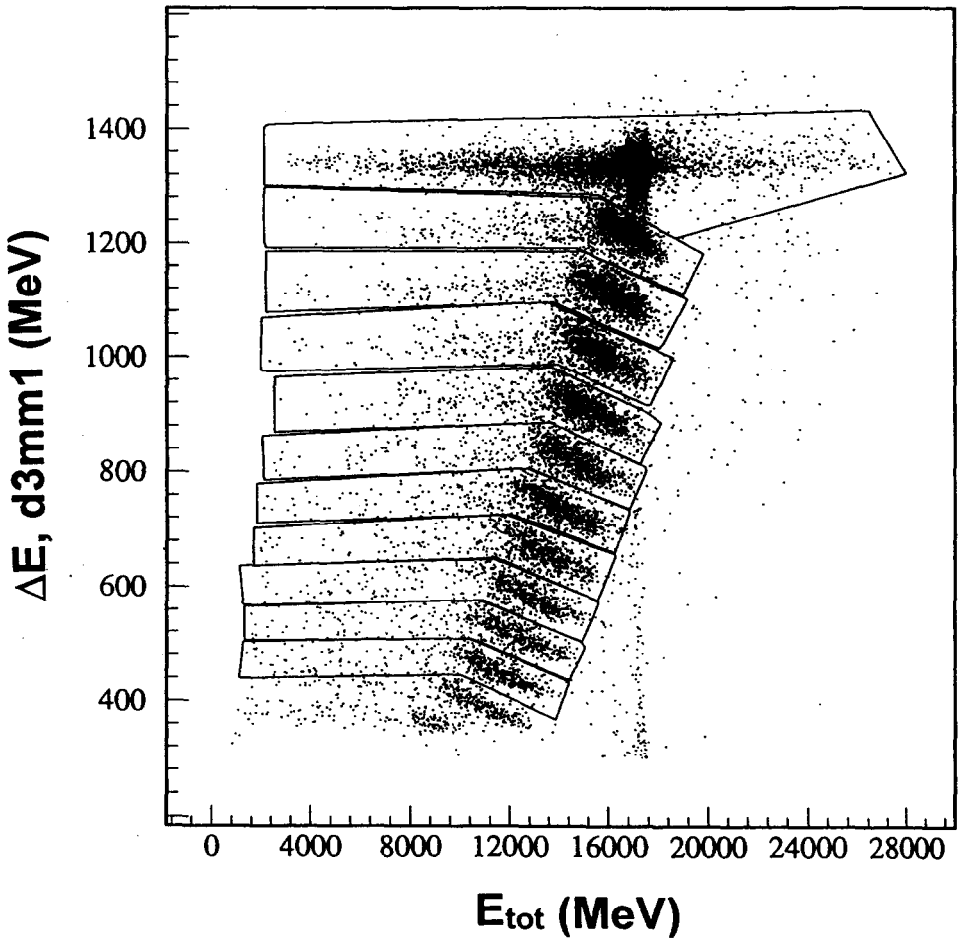

FIG. 11. ΔE_1 vs ΔE_{tot} for the 5 cm run (events passing χ^2). Particle identification cuts are shown. The other runs were handled similarly.

Table 9. No-target results

Z	#events $-\chi^2$	#events $-\Delta E_1$–E_{tot}	#events $-\Delta E_1$–ΔE_2
26	11,192	11,054	10,874
25	331	318	239
24	207	194	183
23	125	132	131
22	139	130	129
21	101	96	97
20	68	66	70
19	54	55	58
18	45	50	50
17	48	43	46

Table 11. Results with the 5 cm target. The ΔE–ΔE method is not applicable for $Z < 22$

Z	#events $-\chi^2$	#events $-\Delta E_1$–E_{tot}	#events $-\Delta E_1$–ΔE_2
26	27,728	27,019	26,680
25	3380	3171	3045
24	2750	2675	2569
23	2196	2176	2016
22	2026	2041	2110
21	1530	1441	—
20	1419	1404	—
19	1018	967	—
18	965	866	—
17	680	710	—

losing 1 or 2 units of charge. The event analysis may well discard d3mm1 in such a case for being inconsistent with the rest of the event; the detectors downstream of d3mm1 then (correctly) identify the event as Mn or Cr, as the case may be. The ΔE_1–E_{tot} method is not as vulnerable to this problem as is χ^2, and for ΔE_1–ΔE_2 it is not a problem at all. ΔE_1–E_{tot} *is* vulnerable to a different problem: on a small fraction of the events, ΔE_1 reads low. Most of these events occur when the particle is near the edge of the d3mm1. Similar effects are seen in d3mm2 and d3mm3. These problems are probably related to imperfect alignment of the detectors in the stack. For all but the 8 cm data, the gate in the ΔE_1–E_{tot} plot is essentially equivalent to a simple cut on ΔE_1; hence, a low reading from d3mm1 results in misidentification. By requiring consistency between the first two detectors, the ΔE_1–ΔE_2 method avoids these problems and is probably the cleanest technique. One of the consequences of requiring good correlation in the first two detectors is the elimination of a small number of events which the other methods are able to identify (note the smaller numbers for Fe in the last columns of Tables 9–12).

It should be noted that there is some inherent uncertainty in the event counts given by the graphical methods. This is due to the fact that the gates represent a "best guess". We estimate that the event counts from these methods are good to no better than about 5%.

In principle, the χ^2 algorithm can separate isotopes of the same species. In practice, the relatively large uncertainties in calibration and detector thickness (including dead layer thickness) preclude isotopic resolution with the present system. However, given the excellent agreement between χ^2 and the other methods (especially for $Z < 25$), we conclude that the method does an adequate job of identifying heavy particles, and we expect that lower-Z particles have also been correctly identified. Thus, the design goal of the stack—particle identification across the entire range of fragment Z—has largely been achieved. Future upgrades, discussed in the following section, may provide the final refinements needed to identify all fragment species.

4. CONCLUSIONS

To summarize, we have developed and tested a solid state detector system for measuring heavy ion fragmentation near the beam axis. The detector is compact and can be quickly configured to optimize its performance over a specified range in fragment charge and energy. Onboard electronics allow preamplifier calibration and gain control to be handled remotely under computer control. The well-understood properties of silicon detectors and the well-developed theory of dE/dx energy deposition enable us to use a χ^2 technique which makes maximum use of the data. Analysis of data taken with ^{56}Fe

Table 10. 2 cm target results

Z	#events $-\chi^2$	#events $-\Delta E_1$–E_{tot}	#events $-\Delta E_1$–ΔE_2
26	33,255	32,497	32,254
25	2028	2110	1863
24	1653	1629	1588
23	1226	1186	1188
22	1105	1107	1067
21	814	799	796
20	772	783	766
19	517	505	505
18	471	474	460
17	340	342	338

Table 12. 8 cm target results. The ΔE–ΔE method is not applicable for this data set

Z	#events $-\chi^2$	#events $-\Delta E_1$–E_{tot}
26	22,200	22,589
25	5190	4568
24	4213	3955
23	3206	3041
22	3201	3062
21	2333	2299
20	2221	2172
19	1507	1498
18	1530	1540
17	1145	1157

beams incident on thick CH_2 targets shows that the detector can resolve charges ranging from $Z = 26$ down to $Z = 2$. At the high end of the Z range, visual methods allow us to confirm that the χ^2 method works well.

Our experience with this system suggests that several possible improvements are desirable. To improve calibration, it would be useful to decrease the uncertainty in the values of the calibration capacitors, and to take more calibration points at low ΔE in order to map out possible non-linearities in the preamplifier/shaper/ADC chain; lower preamplifier noise is necessary to identify protons. Several aspects of the layout of the detector elements and the triggering could also be improved. In particular, PSD1 should not be used in the trigger, since its area is so much greater than d3mm1. The gain of the PSD1 electronics should also be set so that it can be used to measure heavy particle ΔE and position, thus providing trajectory information when combined with PSD2 position measurements. The H1/2 thresholds could be set much lower. Amplification and digitization of the scintillator signals, particularly S, should be handled with more care. Finally, a trigger which will fire on single light particles (down to protons) should be implemented.

Acknowledgements—This work was funded by the National Aeronautics and Space Administration under NASA Grant L14230C through the U.S. Department of Energy under Contract No. DE-AC03076SF00098.

REFERENCES

Brun R. and Lienart (1987) HBOOK User Guide: CERN Computer Center Program Library Long Writeup: Version 4, CERN-Y250.

Brun R., Couet O., Vandoni C. E. and Zanarini P. (1989) PAW: A General Purpose Portable Software Tool for Data Analysis and Presentation, CERN-DD/89/17.

Curtis S. B., Townsend L. W. and Wilson J. W. (1991) Effects of HZE Particles on Astronauts, LBL-31717.

Particle Data Group (1992) Review of Particle Properties. *Phys. Rev.* **D45**.

Schimmerling W., Miller J., Wong M., Rapkin M., Howard J., Speiler H. G. and Jarret B. V. (1989). The Fragmentation of 670 AMeV Neon-20 as a function of depth in water, I. Experiment. *Radiat. Res.* **120,** 36.

Shavers M. R., Curtis S. B., Miller J. and Schimmerling W. (1990) The Fragmentation of 670 AMeV Neon-20 as a function of depth in water, II. One-generation transport theory. *Radiat. Res.* **124,** 117.

Shavers M. R., Frankel K., Miller, J., Schimmerling W., Townsend L. W. and Wilson J. W. (1993) The Fragmentation of 670 AMeV Neon-20 as a function of depth in water, III. Multi-generation transport theory. *Radiat. Res.* **136,** 1.

Simpson J. A. (1983) Introduction to the Galactic Cosmic Radiation. In *Composition and Origin of Cosmic Rays* (ed. Shapiro M. M.). Reidel Publishing, Dordrecht, Netherlands.

Wilson J. W., Townsend L. W., Schimmerling W., Khandelwal G. S., Khan F., Nealy J. E., Cucinotta F. A., Simonsen L. C., Shinn J. L. and Norbury J. W. (1991) *Transport Methods and Interactions for Space Radiations*, NASA Reference Publication 1257, 227.

Wong M., Schimmerling W., Phillips M. H., Ludewigt B. A., Landis D. A., Walton J. T. and Curtis S. B. (1990) The multiple Coulomb scattering of very heavy charged particles. *Med. Phys.* **17,** 163.