Efficiency and resolution of 90-cell prototype

Trigger

We used a scintillator trigger for all the data in this report. In some runs a pair of scintillator bars set diagonally across the chamber picked out cosmic rays from a narrow Z range but covered the full width of the cells. In other runs the large plate scintillators were used to trigger on cosmics passing anywhere through one end of the prototype.

Operating Voltage

After the removal of a broken wire from the prototype in April 2009 all cells could be operated at high voltage. The voltages applied to cells were chosen so as to give the same field strength (same effective voltage) around all anode wires irrespective of their position in the 90-cell block. This was done by using a transfer matrix which describes the influence that a voltage applied to one anode wire has on the field around itself and the neighbouring anode wires. The elements of the matrix were determined experimentally by operating pairs of cells together and using the plasma propagation speed as the best measure of effective voltage. The transfer matrix is shown in Figure 1, and the result of using it to set all cells of the10x9 prototype to an effective voltage of 1500 V is shown in Figure 2. The voltages shown in Figure 2 are known as the standard values. When we wish to vary the operating voltage of the prototype we do so by adding a constant offset to the standard value for all cells. All of the runs used in this note were taken at 50V above standard. In principle it would be more correct to vary the voltage by multiplying the standard voltage values by a scale factor near 1, but the difference between this and adding offsets of up to 100 V is negligible.

	-0.00075	-0.00108	-0.00075	
-0.00075	-0.00227	-0.00833	-0.00227	-0.00075
-0.00108	-0.00833	0.873509	-0.00833	-0.00108
-0.00075	-0.00227	-0.00833	-0.00227	-0.00075
	-0.00075	-0.00108	-0.00075	

Figure 1. Transfer Matrix. The effective voltage on a given cell is obtained by multiplying the voltage applied to it by 0.873509 and adding the voltages of its nearest neighbours × -0.00833, its diagonal neighbours × -0.00227 , etc.

1765	1790	1790	1790	1790	1790	1790	1790	1790	1765
1790	1820	1820	1820	1820	1820	1820	1820	1820	1790
1790	1820	1830	1830	1830	1830	1830	1830	1820	1790
1790	1820	1830	1830	1830	1830	1830	1830	1820	1790
1790	1820	1830	1830	1830	1830	1830	1830	1820	1790
1790	1820	1830	1830	1830	1830	1830	1830	1820	1790
1790	1820	1830	1830	1830	1830	1830	1830	1820	1790
1790	1820	1820	1820	1820	1820	1820	1820	1820	1790
1765	1790	1790	1790	1790	1790	1790	1790	1790	1765

Figure 2. Standard Voltages. The applied voltages necessary to give an effective voltage of 1500 V on all cells. Groups of cells powered by the same CAEN HV channel share the same colour.

Geiger thresholds

The Geiger ASICs have thresholds that can be set to a 2-digit hexadecimal value. The same value is used for both the cathode and anode signals on the chip. Some Geiger cards were modified to amplify the cathode signals to a greater or lesser extent before sending them to the chip. We found that the efficiency for detecting the anode signal was fairly independent of threshold but some ASICs had to be set to a lower threshold in order to make them efficiency for cathode signals. All plots showing cathode performance were made from data taken with the threshold set to x60 and plots showing anode performance were taken at x90.

Pre-selections to remove junk events

Some events contain garbage data in which very large numbers of cells have hits. Another significant proportion of triggers contain few or no hits, maybe because the cosmic ray passed outside the block of cells, or the trigger was just a random coincidence between scintillator pulses. We remove these events by using cuts on the anode hits alone because the anodes generally suffer from fewer problems than the cathodes. First we apply a cut to the anode hit time, requiring it to come after the scintillator hit time. Cells with anode hits which fail this cut are ignored in all subsequent analysis. The distribution that we are cutting on is shown in Figure 3 for thresholds of x60 and x90. It can be seen that most anode hits are at the right time; between 0 and 5 us after the scintillator trigger. The number of early anode hits increases at low threshold, possibly because of noise.

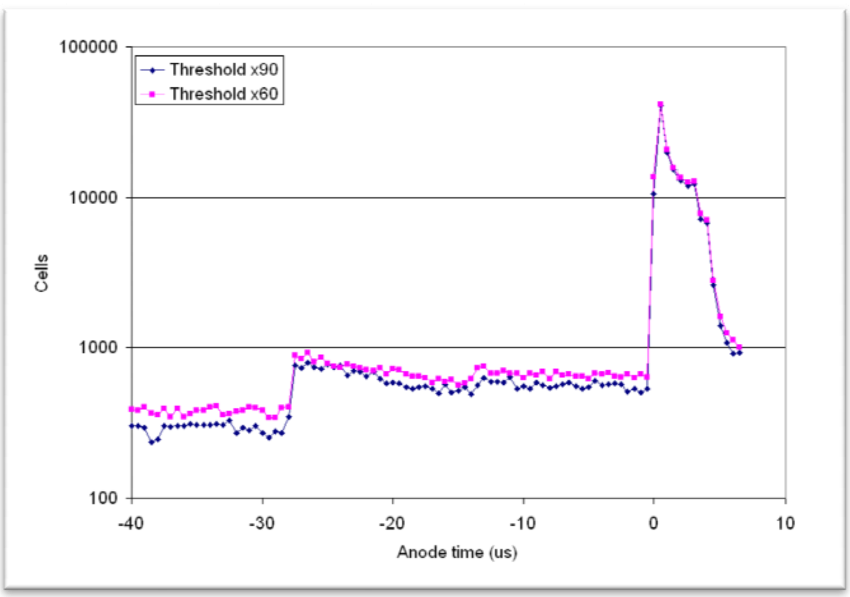

Figure 3. Anode hit time distribution at x60 and xA0 thresholds.

Having excluded bad anodes as describes above, we require that the event has between 7 and 19 good anode hits in order to proceed with the analysis. The lower limit allows for tracks which cross just 9 slightly inefficient cells, while the upper limit allows for tracks at 45º which can hit two cells in every layer. The distribution that we are cutting on is shown in Figure 4.

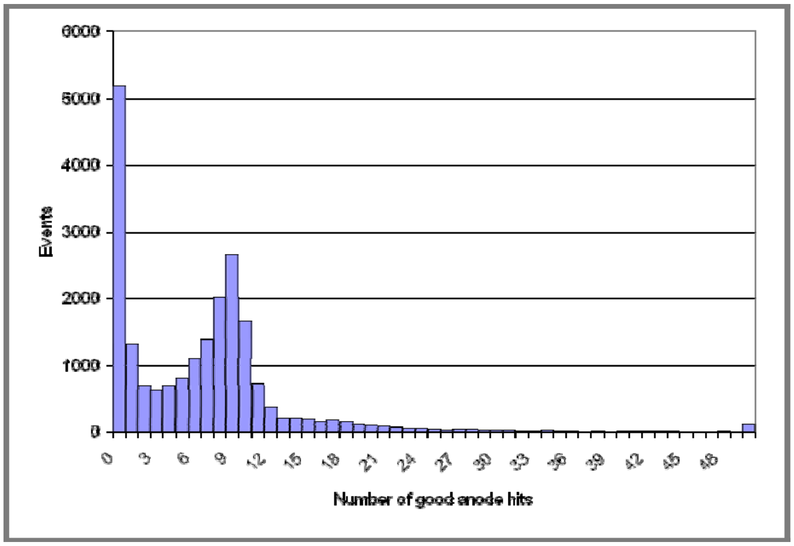

Figure 4. Number of good anode hits. Events with 7 to 19 hits are accepted.

Anode-only fit 1: hit pattern

First we do a simple χ^2 fit to the pattern of hit cells, allowing up to two of the cells to be excluded from the fit because they might be noise or cross-talk.

Each good anode hit is treated as a data point located in the centre of its cell and having uncertainties in both the X and Y directions of 0.4 of a cell width. A straight line χ^2 fit is made to all the points and accepted as the track if the value of χ^2/dof is less than 2.2. If the χ^2/dof is above this limit then one data point at a time is taken out of the fit and the choice which gives the lowest χ^2 is tested. If this choice has χ^2/dof < 2.2 then the fit is accepted as a track and the excluded cell is flagged as noise and takes no part in the remainder of the analysis. Otherwise the procedure continues with all combinations of two cells being excluded and the remainder fitted. If at this stage no good fit is found the fitting is abandoned and the event is ignored.

The number of good anode hits included in the track versus the number excluded from the track is shown in Table 1. Figure 5 shows the distribution of these hits across the 90 cells.

	7	8	9	10	11	12	13	14	15	16	17	18
0	1301	1887	2463	1386	497	156	38	26	16	7	6	0
1	59	94	136	182	112	66	36	22	11	10	5	0
2	10	15	22	50	43	47	26	22	18	11	0	0

Table 1. Number of hits associated with track in fit 1 (horizontal) versus hits flagged a noise (vertical)

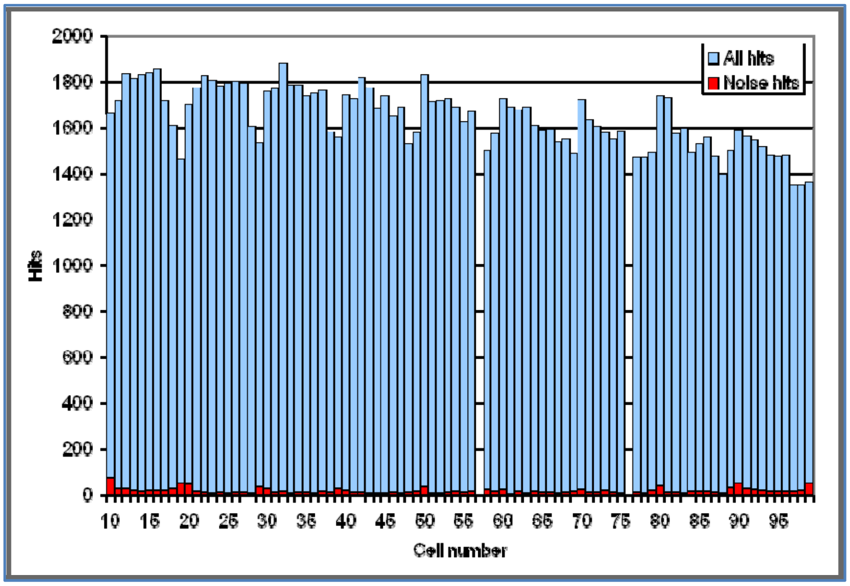

Figure 5. The map of all hits and those classified as noise by fit 1. The number of noise hits is low and they tend to be at the corners and edges of the 10×90 block.

The resulting track at level 1 is described by X0, the x coordinate of its intersection with the mid-plane of the prototype and by its angle to the vertical. The distribution of these parameters is shown in Figure 6.

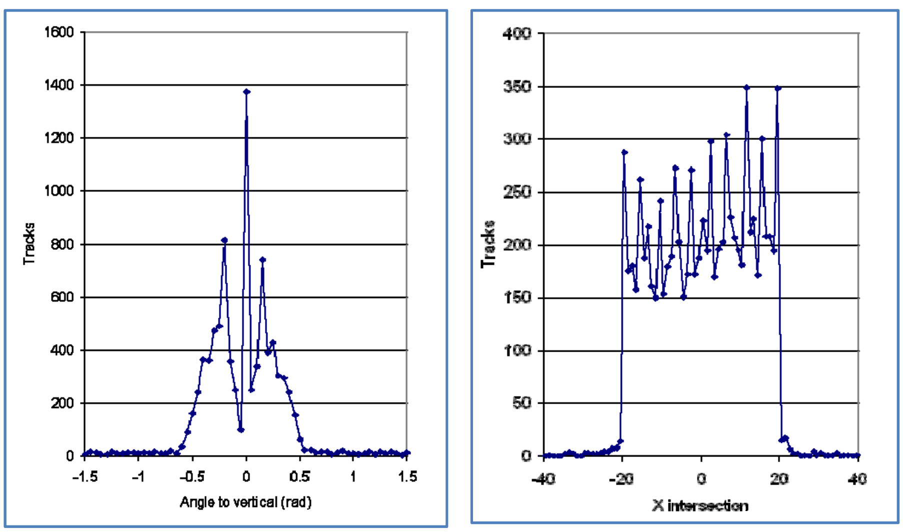

Figure 6. Distribution of the track angle to the vertical (radians) and intercept with the mid-plane (cm).

Anode-only fit 2: drift times

Events where a fitted track was found at level 1, and where its angle to the vertical was less than 45 degrees are passed to level 2 for a fit that makes use of the anode time information. The level 2 fit searches for a minimum in χ^2 where the variables are track angle, track intercept and (optionally) the time of the track relative to the trigger, *t0*. For every cell, *i*, the drift radius, r_i, is calculated from the measured anode time t_i using $r_i = D(t_i - t0)$ where *D* is a drift time-distance relation. The function *D* was at first taken from a gas simulation and later tuned to the data as described below. For every drift circle there are two possible sides, L or R, which could be tangents to the track hypothesis and the closest one is used. The error on the drift circle radius is also a function of $t_i - t0$ taken from the simulation but it is never far from ±1 mm. When minimising with respect to t0 it is possible that for some cells $t_i - t0$ will be negative. In this case the function D returns a radius of zero and a penalty of $((t_i - t0)/0.05\mu s)^2$ is added to the χ^2 of this hypothesis. The choice of the L or R hypothesis for each drift circle means that the χ^2 function can have local minima so a gradient descent minimisation does not work. Instead we use a dumb search over a grid of X0 values within ±2 cm of the value found in fit 1, angles within ±0.15 radians of the value from fit 1 and t0 within ±0.5 μs of the expected value.

All hits which were used at level 1 are included in the fit at level 2. The χ^2/dof of the level 2 fit is shown in Figure 6. A cut is made at 2. Events which pass this cut are used in the following analysis and their track parameters are shown in Figure 6.

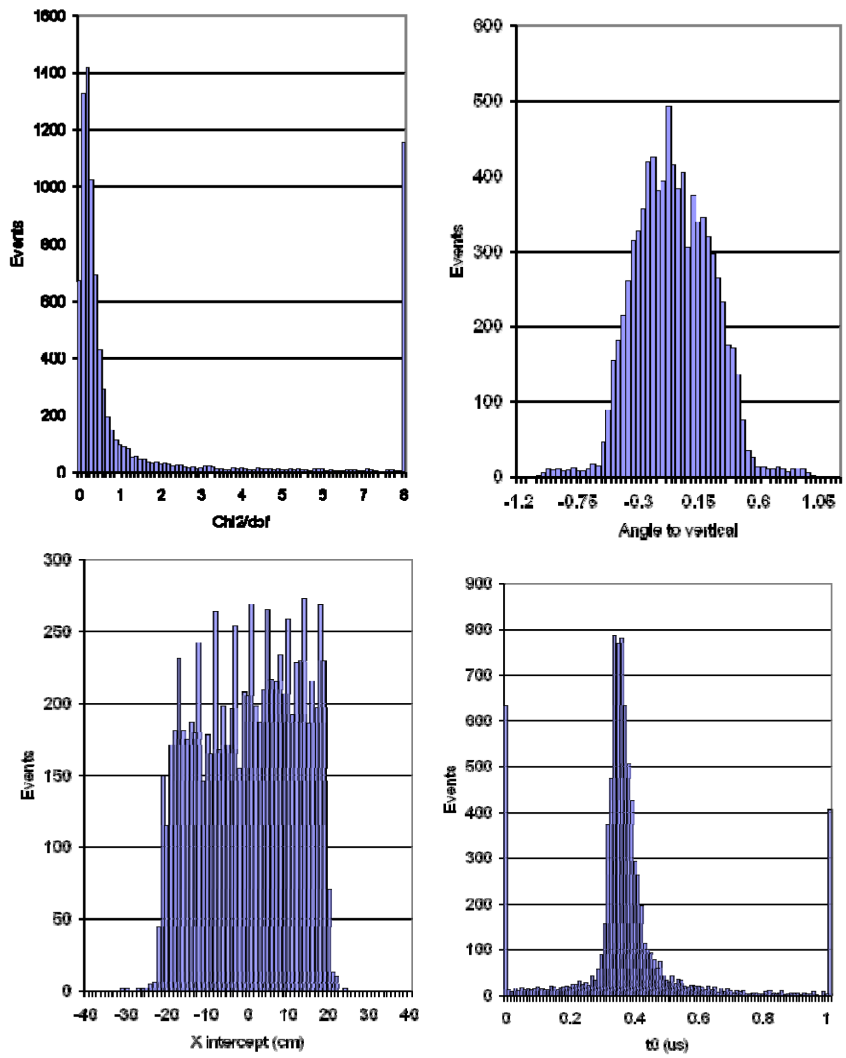

Figure 6. Features of the level 2 fit. Track angle and intercept are smoother than at level 1. The t0 distribution has a sharp peak with width of 0.05 μs.

Drift time-distance relation

A time-distance relation derived from the magboltz gas simulation combined with a FlexPDE electrostatic simulation was used initially. In order to refine this model and make an unbiased test of anode efficiency we analysed all events by masking out all hits on one layer at a time, then doing the entire fit (levels 1 and 2) on this reduced set of hits and using the reconstructed track to test the cells in the masked layer. In these fits with masked layers we kept t0 fixed to a value which had previously been found to be appropriate for this data run.

The distance between the track extrapolated to the masked cell and the anode wire is plotted in Figure 7 against the anode drift time. The agreement between the simulated and measured r-t relations is quite good, but in order to improve further we put the measured r-t relation back into the reconstruction and repeated the procedure.

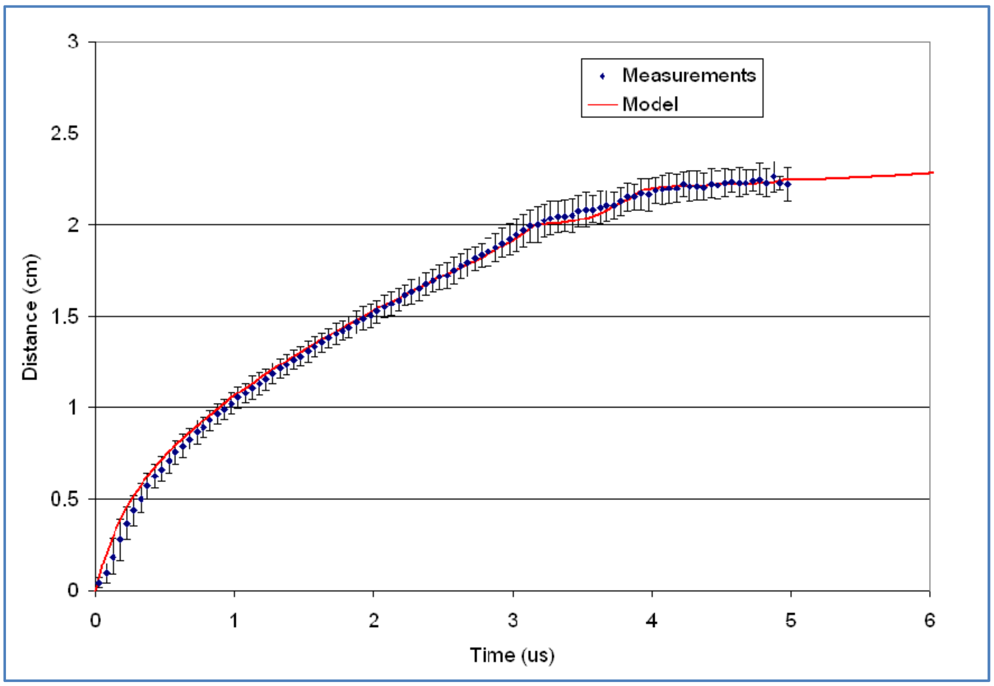

Figure 7. The red line is from the Magboltz/FlexPDE simulation. The points are from the data, with error bars showing the r.m.s. spread of the data that was averaged to make each point.

Drift distance resolution

With a drift time relation that now matches the data quite well we can look at the drift distance resolution. Figure 8 shows the distance between the drift circle and the track for various radius ranges. The best resolution is ±0.7 mm for radius between 5 and 20 mm. Shorter distances suffer from low primary ionisation statistics and longer distances suffer from low and non-uniform electric field at the edges of the cell.

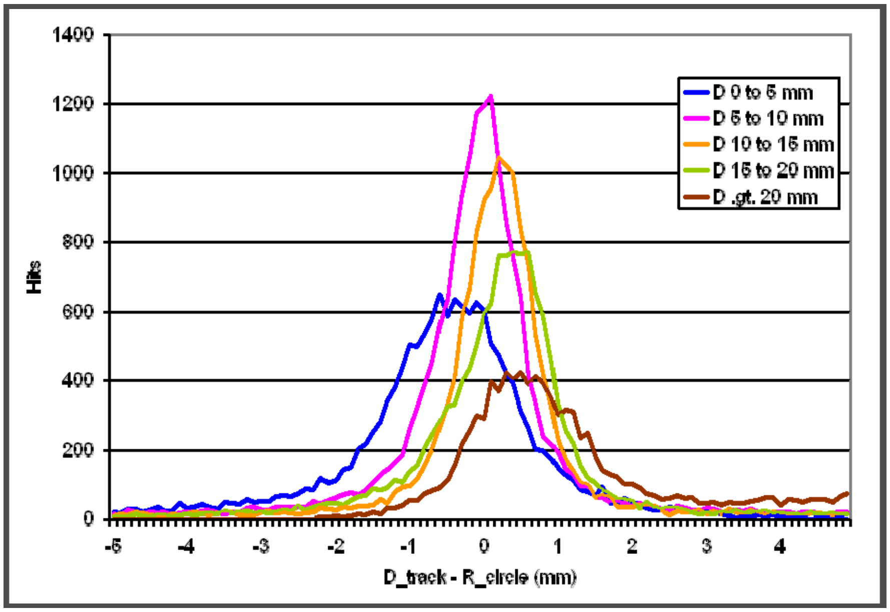

Figure 8. Resolution of the drift distance measurement at various distance ranges.

Anode efficiency

The probability of finding a hit on the anode (with any positive time) versus distance of the track from the anode is shown in Figure 9a. If we take only tracks which pass within 20 mm from the anode and plot the efficiency against cell number we get the values shown in Figure 9b. We know that a cell is dead for of order 1 ms after a hit because this is the time taken for the positive ions to clear away from the wire and the field to be restored to its normal value. Given that the rate of cosmics through each cell is about 100 Hz we are not surprised by a10% inefficiency.

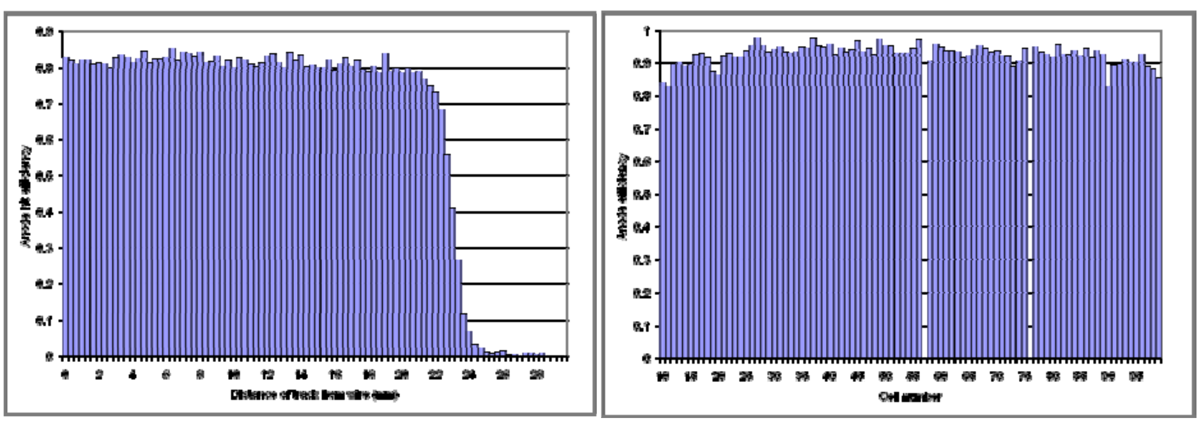

Figure 9. Anode efficiency. The left plot shows efficiency versus distance from the wire. Efficiency is flat up to the edge of the cell at 22mm. The right plot shows efficiency versus cell number.

Cathode analysis

The cathode analysis uses all cells associated with tracks that pass the anode analysis cuts. There are known to be a significant number of cells that will have missing cathode hits because the plasma is blocked by something on the wire. Cathode hits may also be missed because of problems in the Geiger readout cards or because the signal was below threshold.

First we look at the propagation time in cells where a signal was found on both cathodes. The propagation time is $t_p = t_{c1} + t_{c2} - 2t_a$. A distribution of all such propagation times is shown in Figure 10a. There is a typical propagation time of about 52 μs and a tail on the high side. The tail is due to hits in cells that had not fully recovered from their previous hit. Their effective voltage is lowered by the presence of positive ions near the wire and this lowers the plasma propagation speed. Figure 10b shows a map of the propagation time versus cell number, showing that all working cells have similar propagation speed, meaning that our choice of operating voltages has set all cells to approximately the same effective voltage.

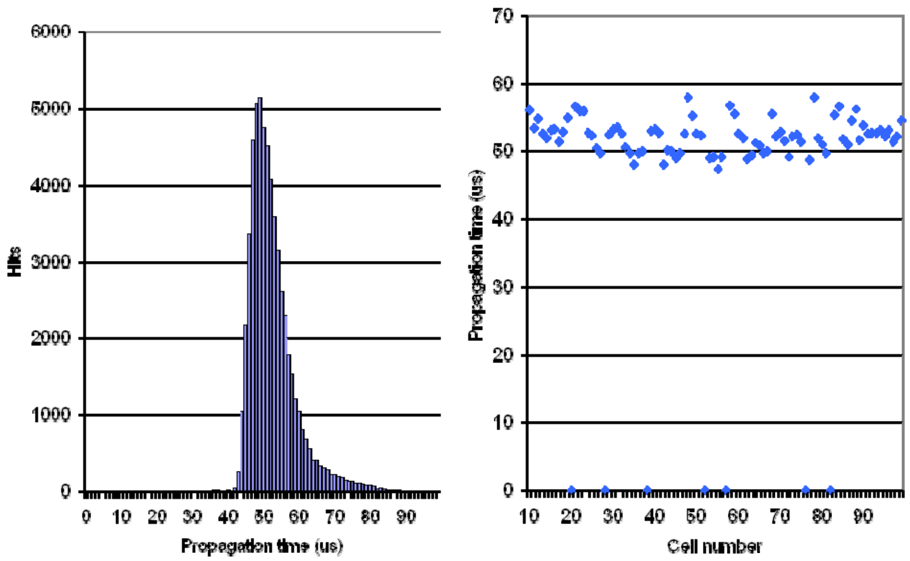

Figure 10. Plasma propagation time. On the left a distribution from all cells, on the right the mean propagation time versus cell number.

Hit cells can be divided into those where the plasma propagates to both cathodes, to cathode C1 only, to C2 only or to neither end. Figure 11 shows a map of the probabilities of these four outcomes for the 90cells. There are propagation problems in 24 cells out of the 90. About half of these are thought to be dues to Geiger card faults and half to blockages on the wire.

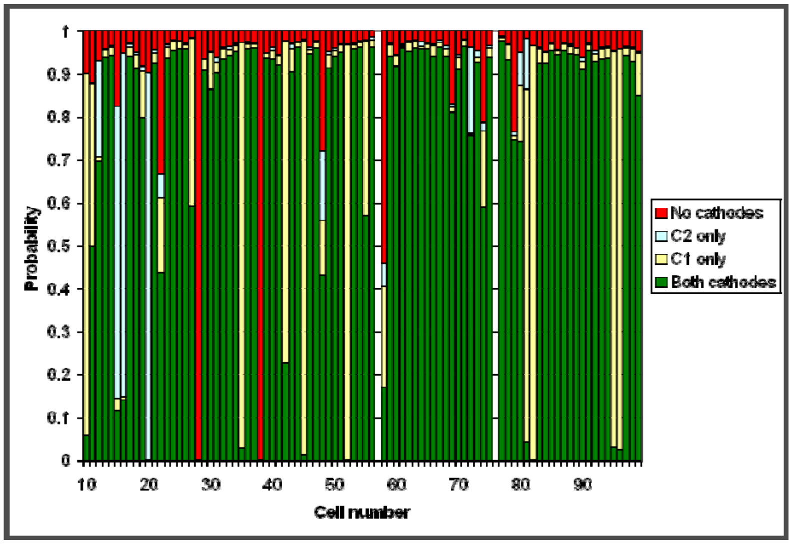

Figure 11. Plasma propagation probability versus cell number.

If a signal is found on both cathodes then the Z position of the hit is reconstructed as $Z_b = L(t_{c1} - t_{c2})/2t_p$. If only C1 has a signal then the Z position is $Z_{c1} = L[(t_{c1} - t_a)/t_{pdef} - ½]$, while for C2-only hits we use $Z_{c2} = L[½ - (t_{c2} - t_a)/t_{pdef}]$, where t_{pdef} is a default propagation time typical for the run.

We select events having Z_b measurements in at least 5 layers and then do a simple χ^2 straight line fit to all of the Z_b hits in the ZY plane. The Z_b measurement error is assumed to be 1.2 cm and the Y measurement error is negligible in comparison. If the χ^2/dof of this fit is below 5 it is accepted for the analysis below. The quality of this fit is indicated in Figure 12. The χ^2 is usually good and the scatter plot shows that the great majority of tracks are reconstructed in the diagonal band between the scintillator bars. This plot uses data taken with the narrow scintillator bars, whereas all other plots in this note use data with the large scintillator panel which covers a much larger region of the prototype.

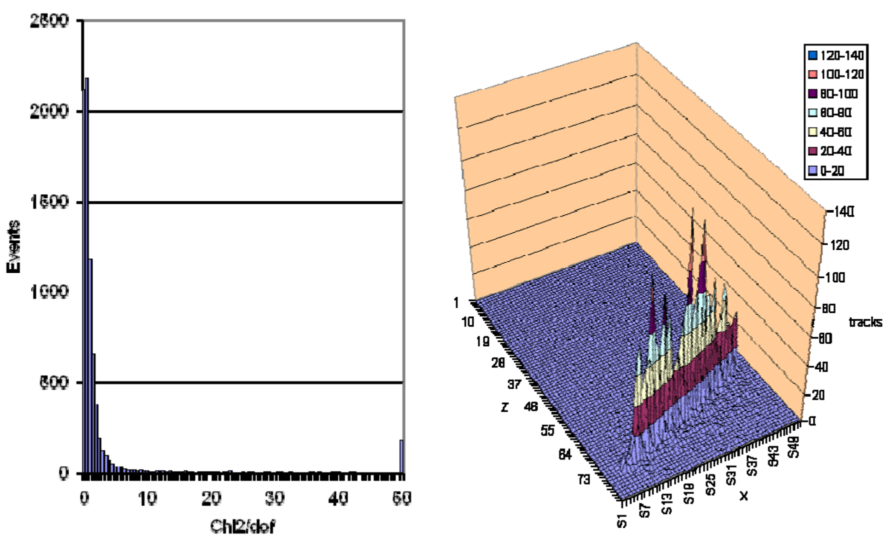

Figure 12. Quality of the ZY fit. Left is the chi2/dof distribution. Right is the impact point of the fully reconstructed track on the horizontal mid-plane of the 90-cell block.

The distance in Z between the reconstructed track and the hit Z measurement is shown in Figure 13 for the various types of hits. In the case where both cathodes are present the r.m.s. width is 1.1 cm, but note that the real accuracy is slightly worse, around 1.3 cm, because the hit itself was used in the track fit. For cases where only one cathode was present the Z reconstruction accuracy is 9.5 cm for C1-only hits and 25 cm for C2-only hits. The difference is because the scintillator panels are close to the C1 end, so that C1-only plasmas have less far to travel and the error on their propagation velocity has less effect. In a low rate environment where the cell is always fully recovered before the next hit the accuracy of single cathode hits will be better.

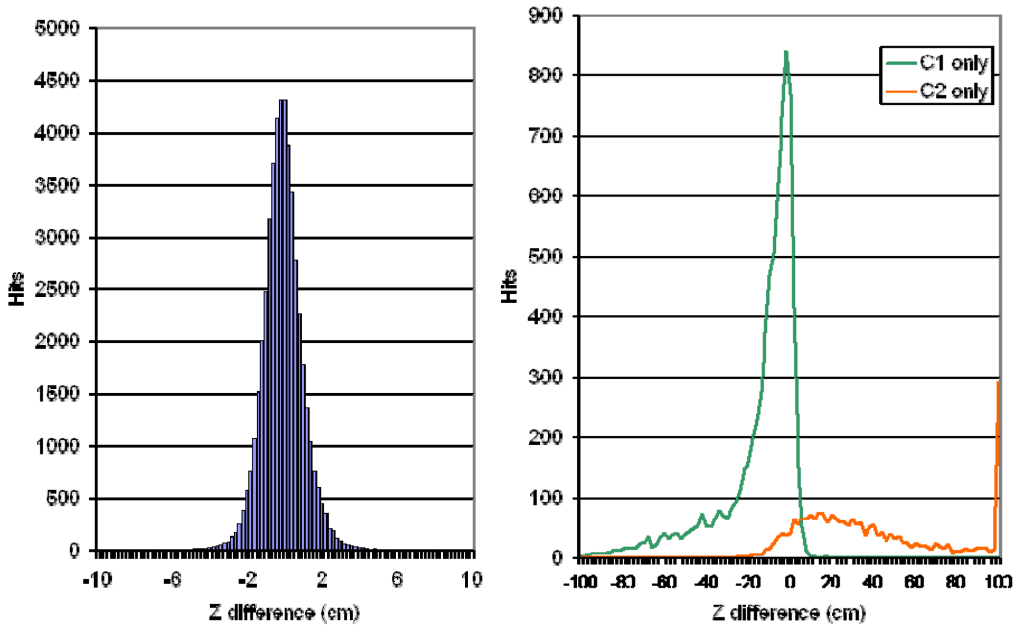

Figure 13. Z resolution for cells where both cathode signals are present (left) and cells where one cathode signal is missing (right).

Conclusions

For cells that are working well the radial resolution is about 0.7 mm and the Z resolution is about 1.3 cm. These values are compatible with the performance of NEMO-3 and with our expectations.

Two cells have dead anodes in this analysis but we know that the problem is in the Geiger cards and the cells themselves are functioning normally. The anode hit efficiency in the other 88 cells is about 90% and we expect this to be higher in a low rate environment. The anode efficiency is flat up to 22 mm from the wire, indicating that there is not a problem with electron attachment to electronegative impurities in the gas.

A large number of cells have problems with plasma propagation. Diagnosing these problems is best done with the scope readout and is not the subject of this note. For the cells which have no plasma blockages and no faults in the Geiger cards, the plasma propagation efficiency is typically 95%.