
Data Analysis for the Bo TPC

This report concerns the analysis of data from the Bo TPC instrument, which is under development as a test bed for electronics for liquid argon TPCs.

The Bo TPC contains three wire arrays, each made up of 48 wires. Array B is rotated by 60 degrees with respect to array A, while array C is rotated by $\frac{\pi}{3}$ in the opposite direction. 2048 readings, separated by intervals of 198 nanoseconds, are taken on each wire for each event. This allows us to produce plots of the signal for each array, as shown in figures 1a), 1b) and 1c). To create these plots from raw data (adc counts), we first subtract 400 (to compensate for the offset of 400 in the data) and divide by 10, then plot the data for each wire centred around the appropriate wire number. Notice that the signal shape varies between the arrays: for instance, in this particular counter orientation, the signal on array B is poor because the muon is travelling parallel to those wires.

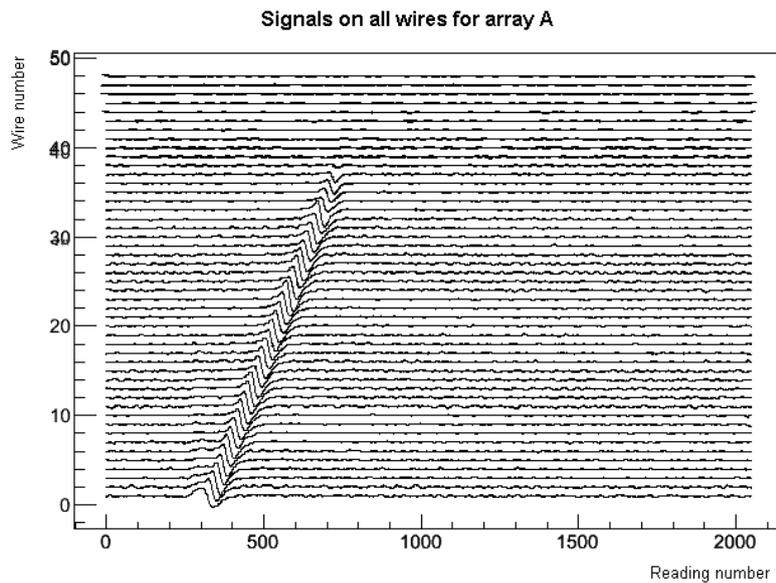


Fig. 1a)

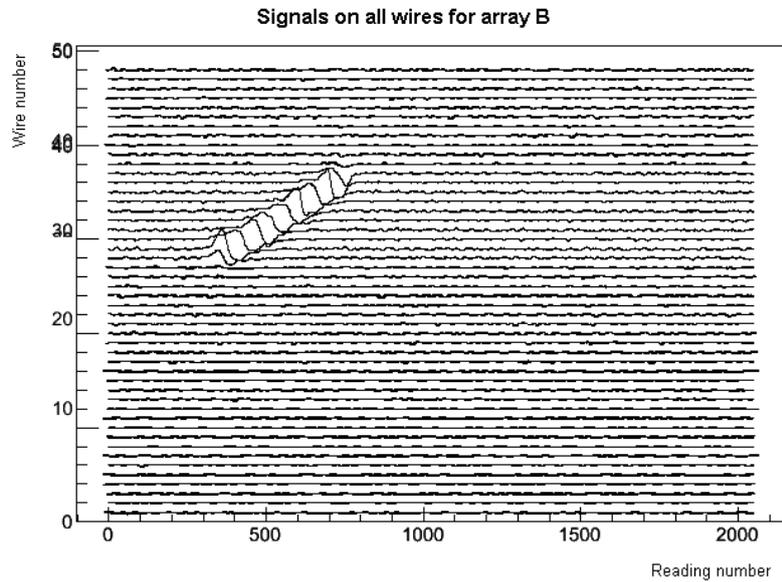


Fig. 1b)

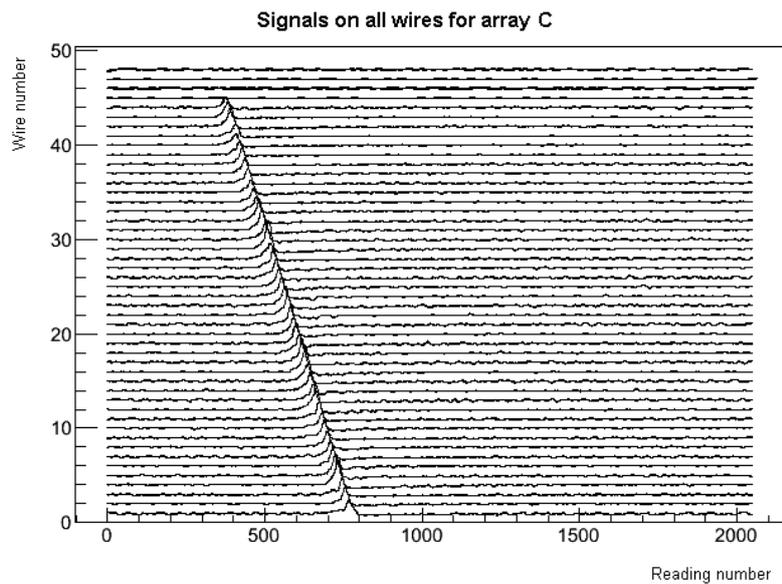


Fig. 1c)

We considered six different aspects of the data collection process:

- 1) Establishing a parameter for the signal/noise ratio.
- 2) Source of noise.
- 3) Normalising the data to compensate for the effect of muon path angle on signal strength.

-
- 4) Establishing a value for the electron drift speed in the liquid argon.
 - 5) Establishing a value for the purity of the liquid argon used in the instrument.
 - 6) Automating the process of selecting events which show a clear single muon path.

1 Signal/noise ratio

We used two different approaches to this calculation.

- 1) Integrating numerically over the signal, then integrating over a window of equal width containing only noise, and dividing the first value by the second. 400 is subtracted from the signal values before integration, to give a signal centred about zero, and we use absolute values for all data so the results are not affected by the negative sign on some readings.

- 2) Finding the peak height of the signal, then dividing this by the RMS value of the noise.

We found method 1) to be problematic because the window containing the signal normally contains less than 50 readings. Therefore if we select a window of equal size containing only noise and integrate over it, the analysis often produces values for the noise which are higher or lower than the true average simply because our randomly selected window contains a period of higher or lower noise than usual. We therefore preferred method 2.

1.1 RMS noise window

For this method it was necessary to decide how many readings to use to use in order to calculate the RMS noise. The following graphs show the value calculated for RMS noise for various different choices of the number of readings. Notice that the RMS noise is always calculated over a window that does not contain the signal.

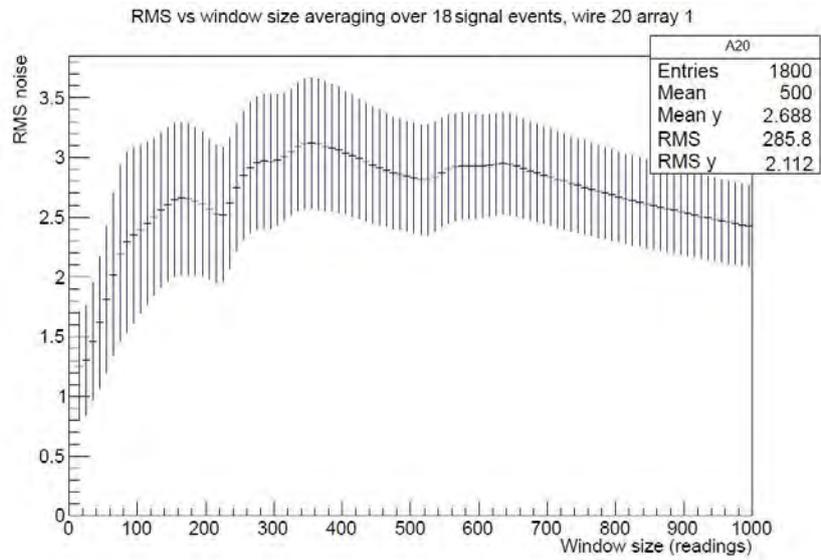


Fig. 2

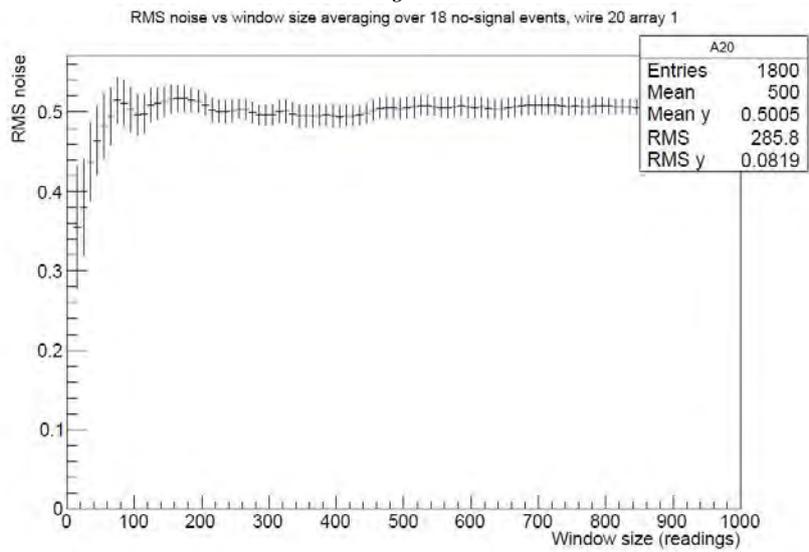


Fig. 3a)

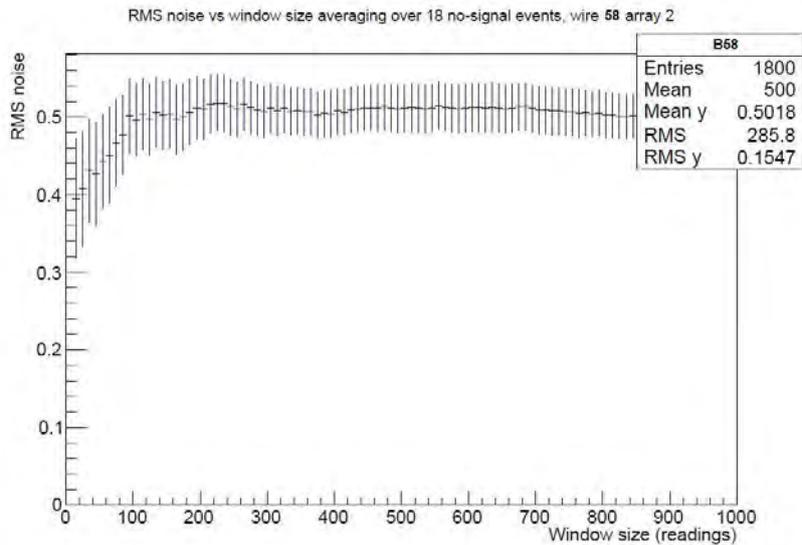


Fig. 3b)

The error bars are calculated using the formula:

$$e = \frac{S}{\sqrt{R}}$$

where S is the spread of values within the bin (the maximum value minus the minimum value), and R is the number of values in the bin (always 18 in this case).

The RMS is calculated using the formula:

$$RMS = \sqrt{\frac{\sum_{i=1}^N (x_i - 400)^2}{N}}$$

where x_i are the readings in the relevant interval, and N is the number of readings.

Notice that 400 is subtracted to compensate for the offset of 400 in the raw data.

A comparison of figure 2 with figures 3a) and 3b) demonstrates that the results for RMS noise have significantly wider spread, are considerably larger, and do not stabilise so consistently, when data is taken from outlying results for an event which contained a signal rather than from results from an event which did not contain a signal. It was therefore decided that the calculation of RMS noise should always be made using data from no-signal events. It can also be seen from figure 2 that the RMS noise result stabilises once the window is about 300 readings wide. This suggests that an appropriate window sampling size is about 400 readings, taken from a no-signal event.

1.2 Additional noise on signal events

In order to establish the reason for the increased noise on signal events, we compared the values for RMS noise before and after the signal. If the RMS noise were roughly equal in both cases, this would suggest that the extra noise is coming from the increased number of excitations due to extra particles (such as photons) resulting from the presence of the muon; but if the RMS noise were greater after the signal, this would suggest that the extra noise results from the fact that the electronics take some time to return to equilibrium after the signal.

The graph in figure 4) shows that the RMS noise is significantly greater immediately after the signal (which in this case occurs between reading 400 and reading 500), and only gradually decreases back to its original size. The table that follows indicates that this result is true for a range of different wires - values in this table are calculated from intervals of width 100 readings, with one interval directly before the signal and the other directly after.

We therefore conclude that the increased noise on signal events is a consequence of the time taken for the electronics to return to equilibrium after a signal.

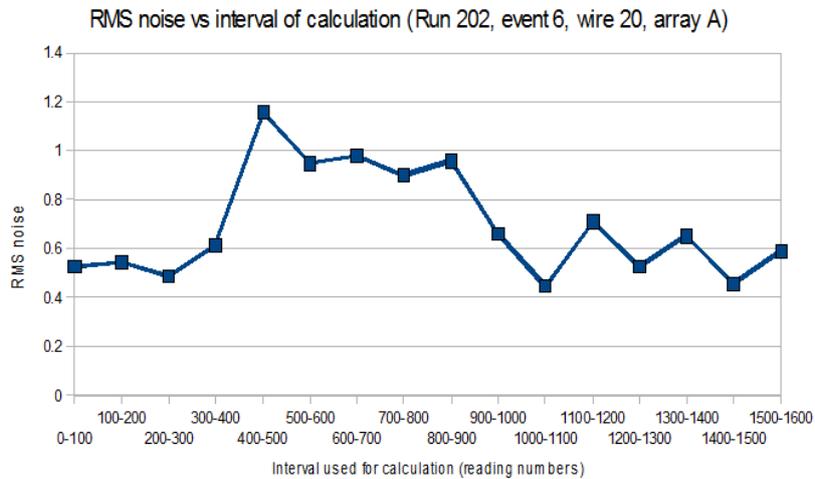


Fig. 4

Array	Wire	RMS noise before signal	RMS noise after signal
A	10	0.22	0.86
A	20	0.55	0.96
A	30	0.61	0.73
C	10	0.47	1.10
C	20	0.56	0.99
C	30	0.41	0.83

2 Source of noise

In order to further understand of the source of noise (including noise on non-signal events, which cannot be explained as the response of the electronics to the signal), we carried out analysis to determine whether the noise on different wires is coherent.

We first averaged the values in each time bin over the given range of wires, then took the RMS of these averages (A). We then calculated the RMS for each wire separately and averaged these results (B). If the variation on the wires is statistically independent, we expect

$$A = \frac{B}{\sqrt{W}}$$

where W is the number of wires used in the calculation.

The results of this analysis are summarised below. The values given are averages over eight different events, using wires 10-20 on each array, and taking the RMS over a window of size 300 readings.

Array	A	B	$\frac{B}{\sqrt{W}}$	Ratio $\frac{A\sqrt{W}}{B}$
A	0.54	0.2	0.16	1.24
B	0.55	0.19	0.17	1.15
C	0.58	0.23	0.17	1.34

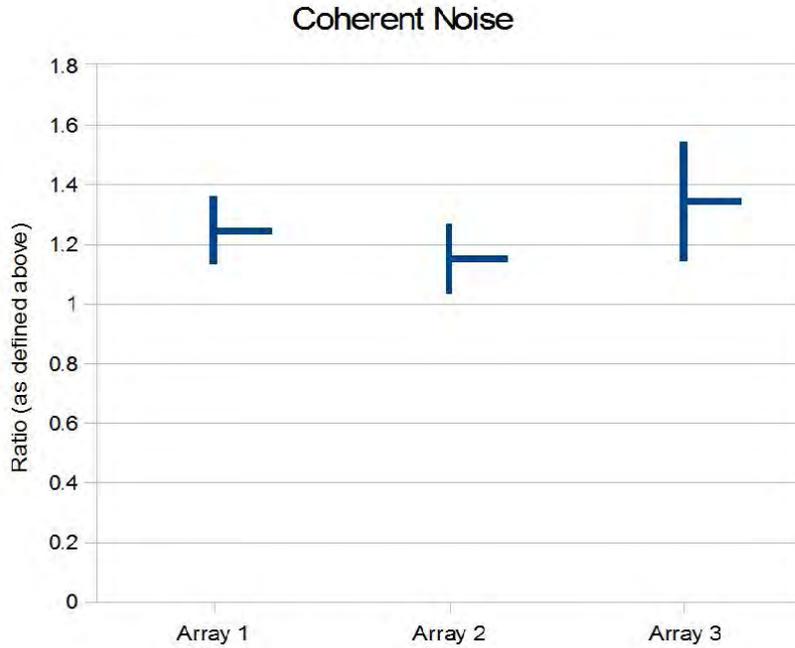


Fig. 5

Figure 5 shows the ratio for each of the three arrays. The error bars here simply range from the highest value obtained to the lowest value obtained.

It can be seen that in all cases, the final ratio is slightly greater than one. This indicates that there is some small correlation of the noise on the wires, which implies that there is a minor external source of noise such as the amplifier power source.

3 Normalisation

By considering the projection of the muon path on the wireplane, we see that the strength of the signal received from a given muon depends on the angle of the muon path in a way given by:

$$\frac{T}{\sin(\phi)\sin(\theta)}$$

where T is the true signal strength, ϕ is the azimuthal angle of the muon path (relative to the direction of the wires in the array under consideration) and θ is the polar angle of the muon path.

Therefore we would like to normalise the data by multiplying each result by the factor $\sin(\phi)\sin(\theta)$.

The angles for the muon path can be calculated by comparing the slopes of the muon tracks across plots similar to figures 1a), 1b) and 1c).

3.1 Slope calculation: first analysis

The value of the slope obtained for the muon path varies by about 10% depending on whether the best fit line approximating that path is fitted to maxima, minima or to the midpoints. We compared each of these three methods by calculating the chi-square value as well as the RMS value of the difference between the line and the actual data for each kind of line over a range of different events.

Notice that for the purpose of these tests, we defined the midpoint as the point when the data first passes through 400 after the maximum (recall that the minimum always occurs after the maximum).

For the line passing through the maxima, the RMS is calculated using the formula:

$$RMS = \sqrt{\frac{\sum_{i=1}^N (x_i - f^{-1}(y_i))^2}{N}}$$

The sum is taken over all the wires in the array. (x_i, y_i) are the coordinates of the maximum point for wire i on a plot similar to those shown in figure 1a)-c), so that x has units of number of readings, and y has units of wire number. (Note, however, that y need not be an integer, as it is taken from the data which has been centered around the appropriate wire number). $f(x)$ is the function representing the best fit line, so that the inverse $f^{-1}(y_i)$ gives the x value on that line corresponding to the y value of the maximum for wire i .

For the line passing through the minima, the calculation is similar, but we substitute the minimum point for the maximum point.

For the line passing through midpoint, the calculation is similar, but we substitute the midpoint as defined above for the maximum point.

We tested 25 different run-array combinations - notice however that tests involving array C left out the minimum line, because C is the collector plate and therefore does not show a minimum. The table below shows the average ratios between the RMS and Chi2 values obtained for the maximum, minimum and midpoint best fit lines. The graphs in Fig 6a) and 6b) show the RMS and Chi2 values for each of the tests, normalised to the midpoint value (in order to compensate for fluctuations arising from different signal strengths).

	RMS	Chi2
Maximum:midpoint	1.15	5.83
Maximum:minimum	0.4	1.02
Minimum:midpoint	3.17	8.77

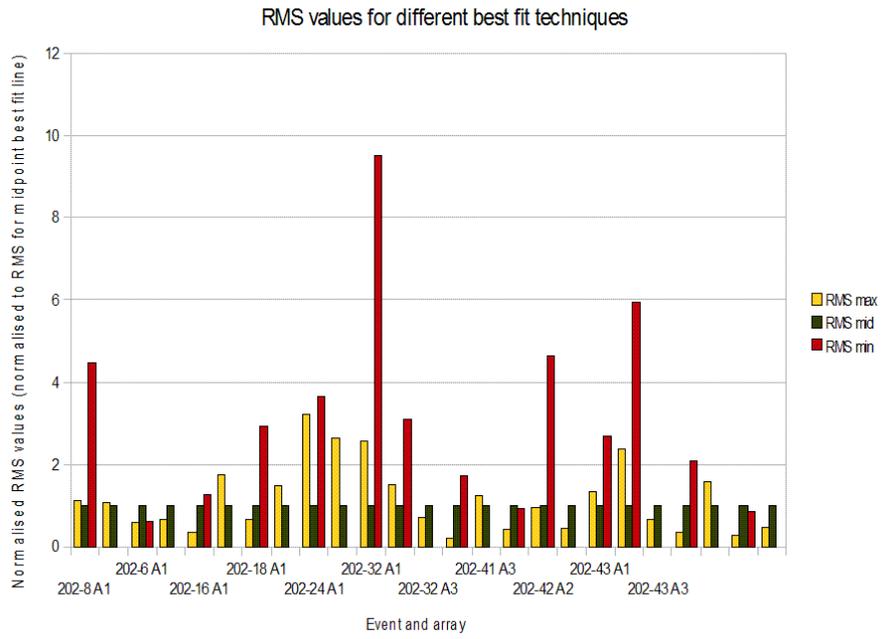


Fig. 6a)

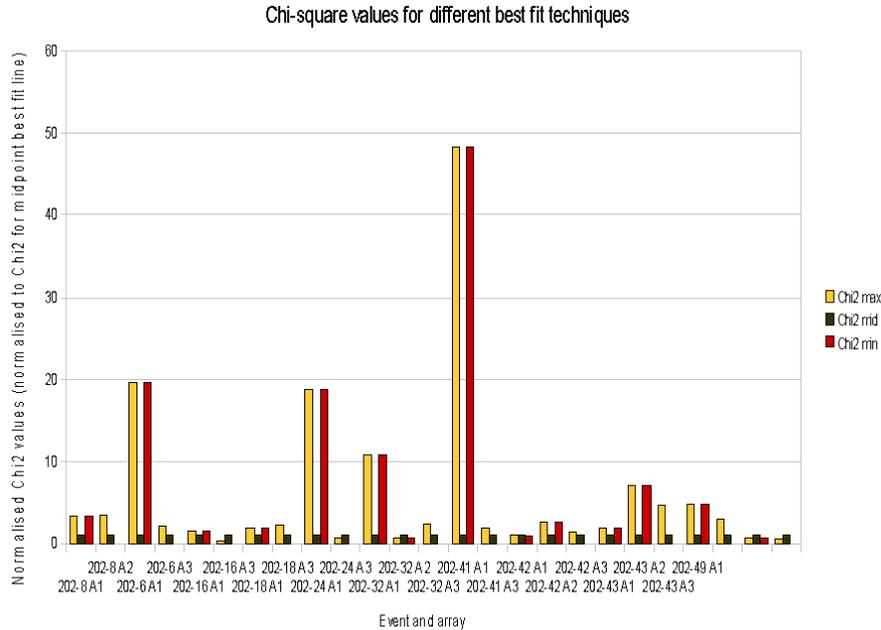


Fig. 6b)

These results suggest that a best fit line passing through the midpoint is most appropriate. The fit through the maxima appears slightly better than the fit through the minima, but this may be bias arising from the fact that the sampling size for lines through minima was smaller.

3.2 Slope calculation: second analysis

We considered a more precise fitting technique in order to improve the fit of the lines passing through the maxima and minima. We selected a window around each maximum, extending on each side of it until the point where the values decreased back down to one third of the maximum value. We then fitted a third degree polynomial to the data points (given by $\frac{adccounts-400}{10}$) falling within this window. The maximum value of this function was then substituted for our original maximum value. A similar procedure was applied to the the minima. A plot of one such polynomial appears below:

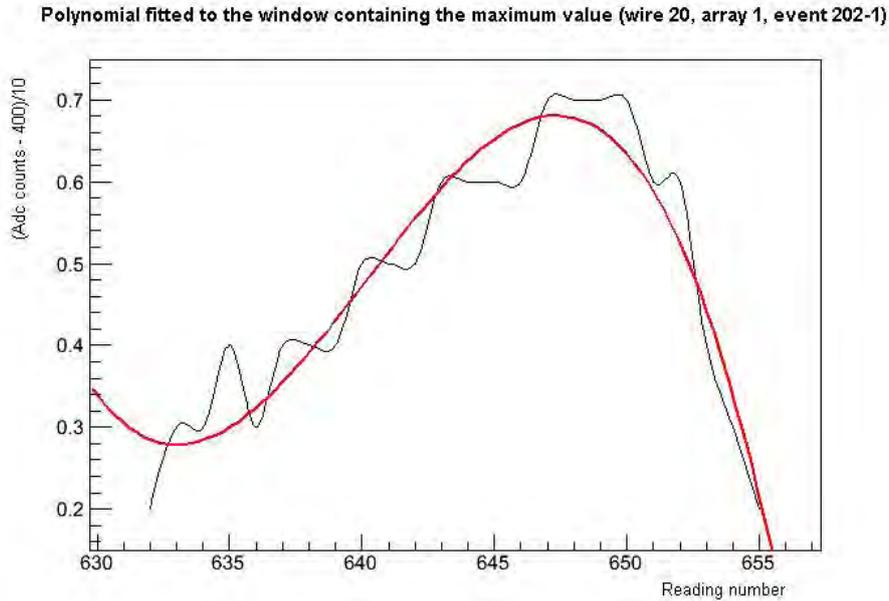


Fig. 7)

The following table shows the average ratios between the RMS and Chi2 values obtained for the original lines and the new lines.

	RMS	Chi2
Original maximum line: new maximum line	1.09	2.04
Original minimum line : new minimum line	4.51	1.46
Original midpoint line: new maximum line	0.75	0.13
Original midpoint line: new minimum line	0.57	0.08

These ratios show that this new technique has indeed given a better fit than the original versions of the lines passing through maxima and minima. However, even with this improved technique, the line passing through the midpoints still gives the best overall fit to the data.

3.3 Angle calculation

The azimuthal angle can be calculated using the slopes from any two planes. The slopes are used to find the time between readings on two consecutive readings: for slope m , the time between readings is $t_i = \frac{198 \times 10^{-9}}{m_i}$ where the numerical factor simply converts to correct units.

As an example, the derivation of the formula for the azimuthal angle using data from arrays A and B is as follows.

The rate of procession of the signal in array A is given by $v = \frac{w}{t_A \cos(\phi)}$ and since array B is rotated by $\frac{\pi}{3}$ with respect to array A, the rate in array B is $v = \frac{w}{t_B \cos(\phi + \frac{\pi}{3})}$, where w is the spacing between the wires (0.004695), t_A is the time between readings on adjacent wires in array A, and t_B is the time between readings on adjacent wires in array B.

Assuming that the velocity of the electron cloud between arrays A and B is always the same, these two rates are equal, so we can equate the right hand sides of our equations and rearrange to obtain $\sin(\phi)(2t_A - t_B) = \sqrt{3}t_B \cos(\phi)$.

Further rearrangement yields

$$\phi = \tan^{-1}\left(\frac{\sqrt{3}t_2}{2t_1 - t_2}\right)$$

where ϕ is measured relative to the direction parallel to the wires.

Alternatively, we can leave the expression in terms of slopes (the numerical factor 198×10^{-9} will cancel):

$$\phi = \tan^{-1}\left(\frac{\frac{\sqrt{3}}{2} \frac{m_2}{m_1} - \frac{1}{m_2}}{\frac{1}{m_1} - \frac{1}{m_2}}\right)$$

Similar calculations yield formulae for other combinations of arrays:

$$\phi = \tan^{-1}\left(\frac{\sqrt{3}t_3}{2t_1 + t_3}\right)$$

$$\phi = \tan^{-1}\left(\frac{\sqrt{3}t_3}{2t_2 + t_3}\right) - \frac{\pi}{3}$$

Notice that this produces the azimuthal angle relative to the A plane: we add $\frac{\pi}{3}$ to get the correct angle for the B plane, and subtract $\frac{\pi}{3}$ to get the correct angle for the C plane.

The slopes can also be used to obtain the polar angle. The value v calculated above should equal $v_d \tan(\theta)$ where v_d is the electron drift speed (1500 m s^{-1}). Substituting one of the formulae for v and rearranging, we obtain the result:

$$\theta = \tan^{-1}\left(\frac{w}{\sin(\phi)v_d t_1}\right)$$

3.4 Evaluating the normalisation

The graphs that follow give a qualitative picture of how the data are affected by the normalisation. They are histograms showing the relative frequency of different values for the maximum reading on the signal (obtained from the polynomial fitted to the maximum, as in section 3.2) before and after normalisation, for an arbitrarily selected set of wires. Notice that wires near the start and end of the range are excluded, because they tend to show more erratic results.

Maximum reading (on fitted polynomial) for unnormalised data on array A

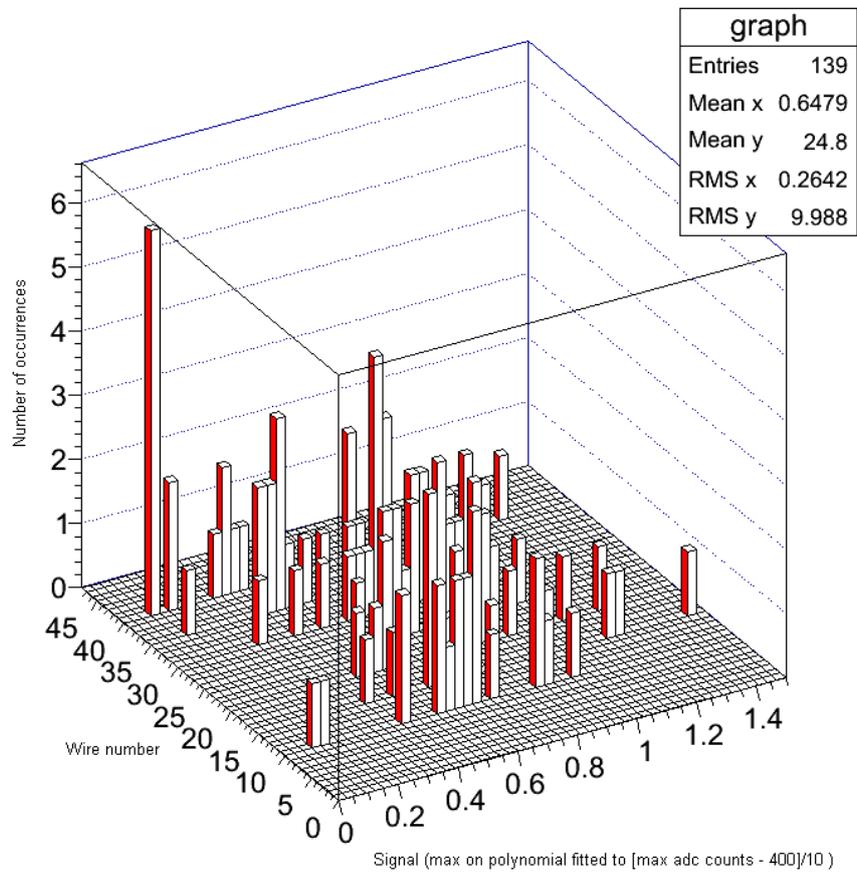


Fig. 8a)

Maximum reading (on fitted polynomial) for normalised data on array A

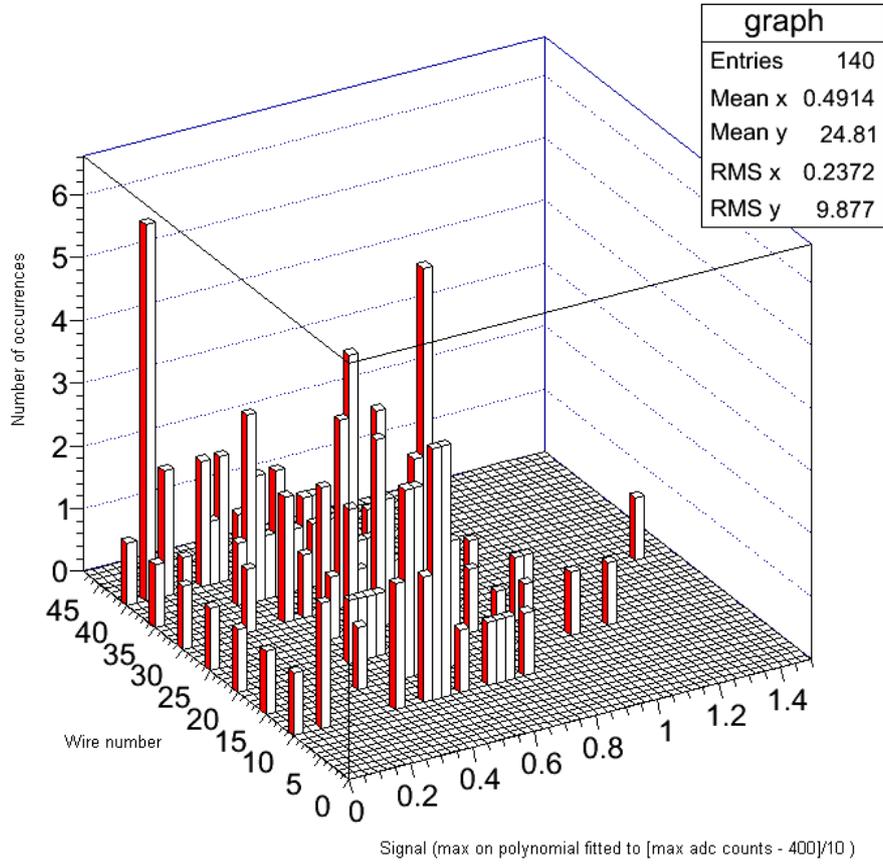


Fig. 8b)

Maximum reading (on fitted polynomial) for unnormalised data on array C

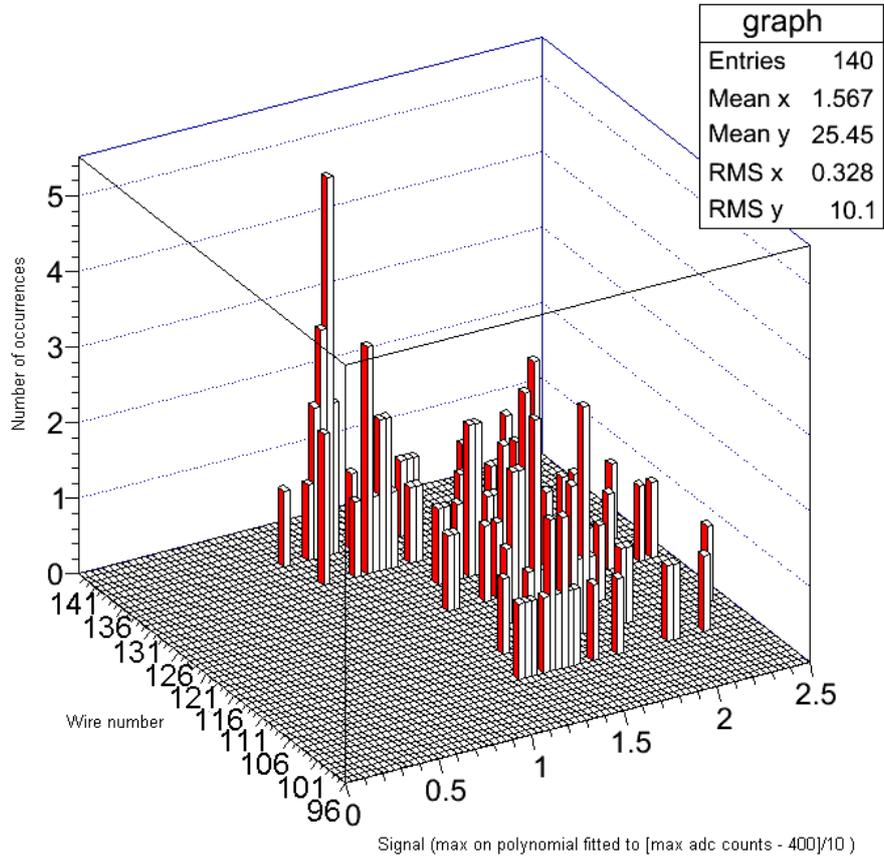


Fig. 8c)

Maximum reading (on fitted polynomial) for normalised data on array C

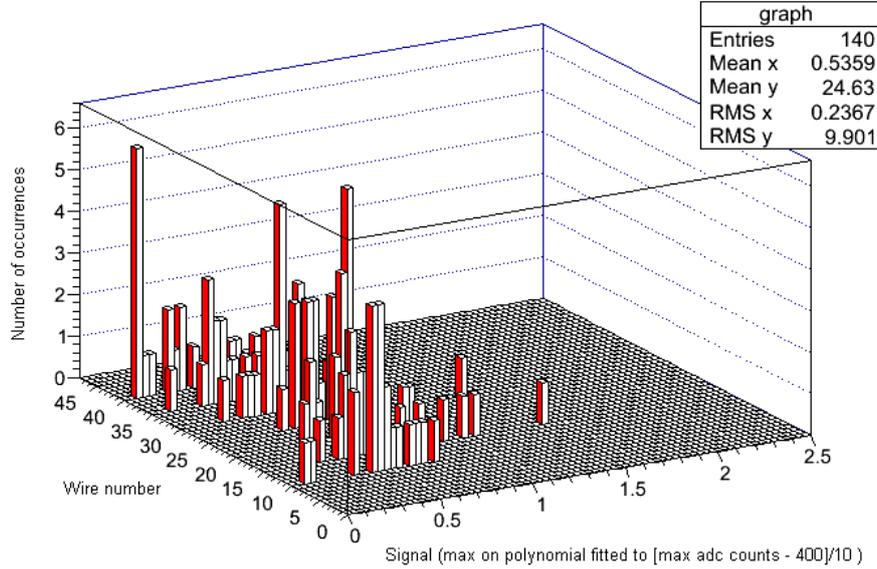


Fig. 8d)

We calculated the uncertainty on the normalisation factor over a range of events in order to determine the precision of our normalisation process. These uncertainties were calculated assuming that the uncertainty on the wire spacing is 0.0000005 m and the uncertainty on the electron drift velocity is 50 ms^{-1} . However, an uncertainty of 0.0000005 on the wire spacing is for an ideal case; in fact the uncertainty may be significantly greater, meaning that these uncertainties may be underestimates.

Event	Uncertainty on normalisation factor
Run 202, event 6	0.0000113
Run 202, event 8	0.0005090
Run 202, event 16	0.0000283
Run 202, event 24	0.0000767
Run 202, event 29	0.0004480
Run 202, event 32	0.0000053
Run 202, event 36	0.0014000
Run 202, event 41	0.0000460
Run 202, event 43	0.0180000
Run 202, event 49	0.0001840
Average	0.0020711

The fact that the uncertainties on the normalisation are small indicates that we are achieving reasonable precision in this process.

If the normalisation is working correctly, we would expect to see that the strength of the correlation between signal and azimuthal or polar angle decreases after the data has been normalised. As a preliminary check on whether this might be the case, compare the graphs in figures 9a) to 9d). At least for this randomly selected data, we see that the upward/downward trends are less apparent after normalisation. It is a positive sign that in both cases the magnitude of the slope for the best fit line is smaller after normalisation, since we would expect zero slope if there is no correlation between signal and angle.

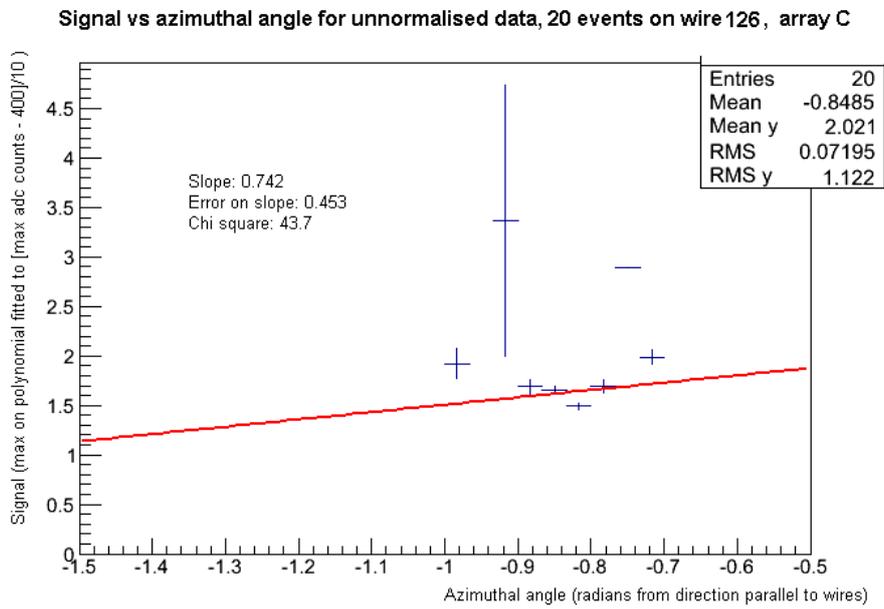


Fig. 9a)

Signal vs. azimuthal angle for normalised data, 20 events on wire 126, array C

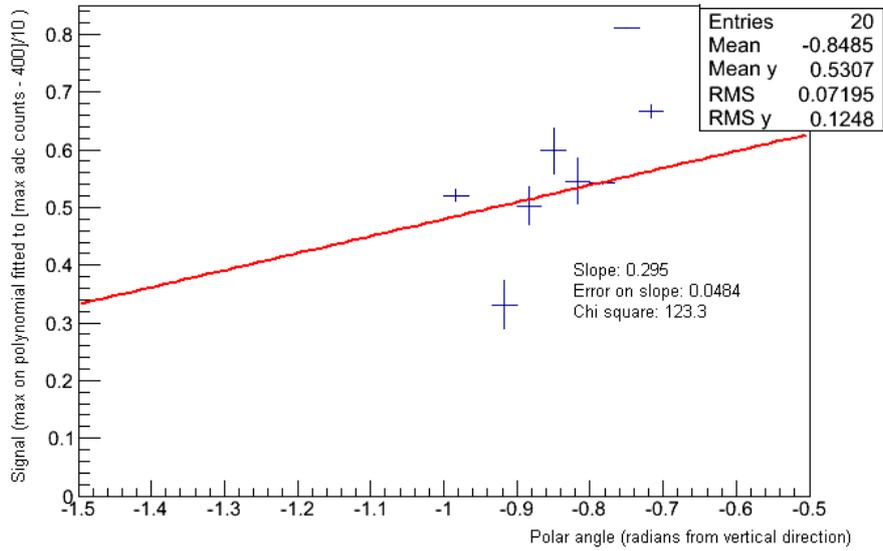


Fig. 9b)

Signal vs polar angle for unnormalised data, 20 readings on wire 126, array C

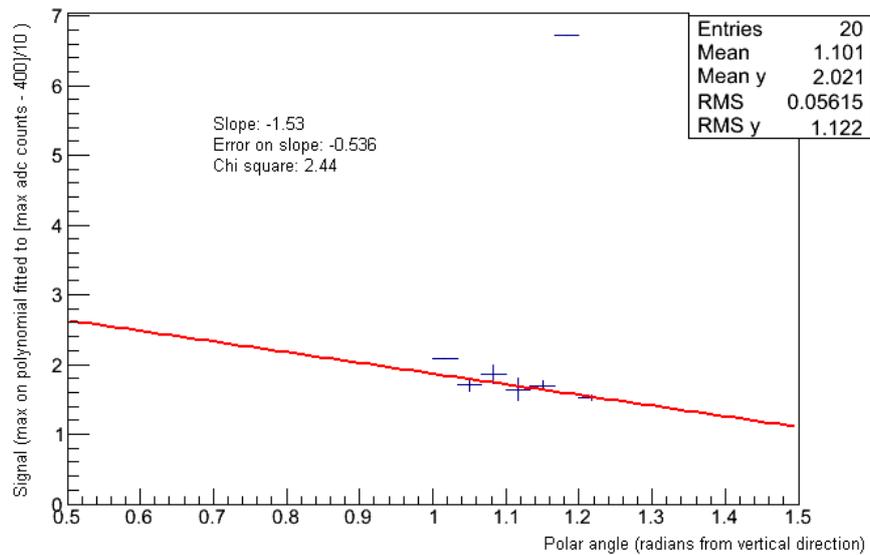


Fig. 9c)

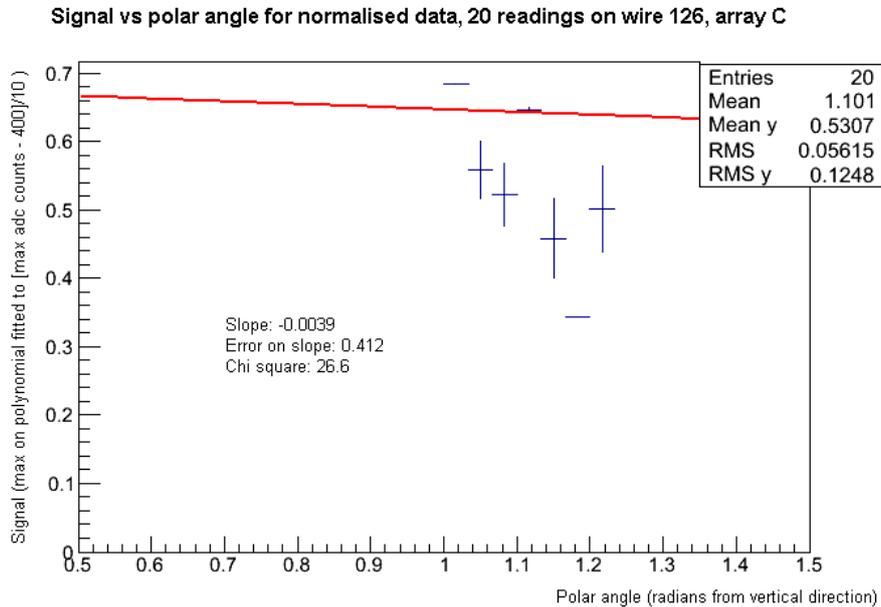


Fig. 9d)

4 Drift Velocity

As muons pass through liquid argon, they ionise particles and produce electrons, which are detected on arrival at the wireplanes. It is essential useful to have an estimate for the drift velocity of these electrons, because the drift time is used to establish the position and transverse geometry of events. The drift velocity is also useful to give an indication of the strength of the electric field being applied. In order to obtain such an estimate, we carried out two consecutive runs with the counters at different vertical heights. When the counters are moved down, the electrons take a longer time to get to the array, and therefore for any given wire, the average time for the maximum reading will be shifted to the right. We assume that the angular distribution is the same for all vertical heights; this assumption is valid provided that there is nothing in the immediate surroundings which blocks a significant proportion of muons.

4.1 Number of readings: first analysis

To establish how many readings we would need to perform this calculation to the appropriate degree of accuracy, we carried out a preliminary analysis on two short runs, each having eight readings, with the counters at different vertical heights. The table below shows the results on these two runs (results are in multiples of 198 nanoseconds, which is the time between

two consecutive readings). Notice that there is a clear shift in the average position of the maximum reading between the two runs.

	207 A	207 B	207 C	208 A	208 B	208 C
Average max position	995.63	707.75	1158.63	424.13	592.25	476.75
RMS	130.05	543.76	118.94	118.65	253.96	167.89
Uncertainty	45.98	192.25	42.05	41.95	89.76	59.36

The average change in time from the results for the A and C planes is 1.24×10^{-4} seconds (data from array B has been omitted, because the results for the B plane are poor for this orientation). The counters were moved through a vertical height of 16.5 cm between these two runs. Thus we calculate a drift velocity of approximately 1330 ms^{-1} . However this result is very approximate, given that it is based on so few readings.

Figure 8 shows the uncertainty versus the number of readings for three sets of data - two from array A and one from array C. Polynomials have been fitted to this data to give an idea of the trend.

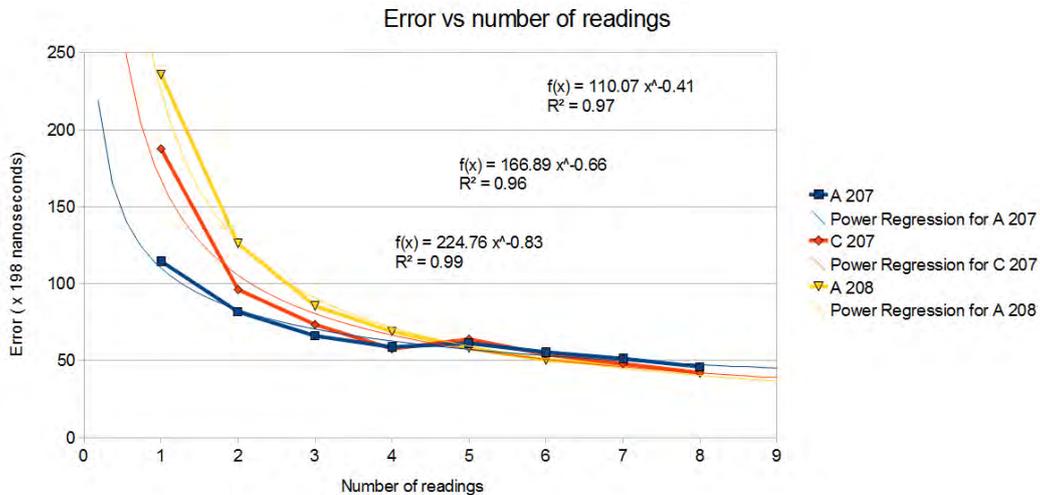


Fig. 10)

Notice that the equations for the best fit lines suggest that the uncertainty is approximately proportional to the number of readings to the power of negative half, which is what we expect from standard statistical results. Using these equations, we can extrapolate to predict the number of events that would be required to achieve an uncertainty for each average maximum position of less than 10 counts (1.98 microseconds):

	Predicted number of events
Run 207, Array A	64
Run 207, Array C	71
Run 208, Array A	43
Average	59.06

In general, roughly half of the events recorded are good events which can be used in a calculation of this kind. Therefore we will need to record at least 120 events in order to produce results with an error of less than 1.98 microseconds.

The change in height of the counter has an estimated error of 0.4 cm. Using standard error propagation methods, we find that if this error were combined with an error of 1.98 microseconds on each time measurement, the overall error would be about 3%, or 40 ms^{-1} for a result of 1330 ms^{-1} .

4.2 Number of readings: second analysis

To improve accuracy, we considered a second method of establishing the change in time: plot graphs of signal time vs polar angle for a single wire, fit a line to the data, and compare the y intercepts for such graphs over the two different runs. We hoped this method would help eliminate some of the variation that arises from different tracks having different polar angles. An example of such a plot is shown in figure 11.

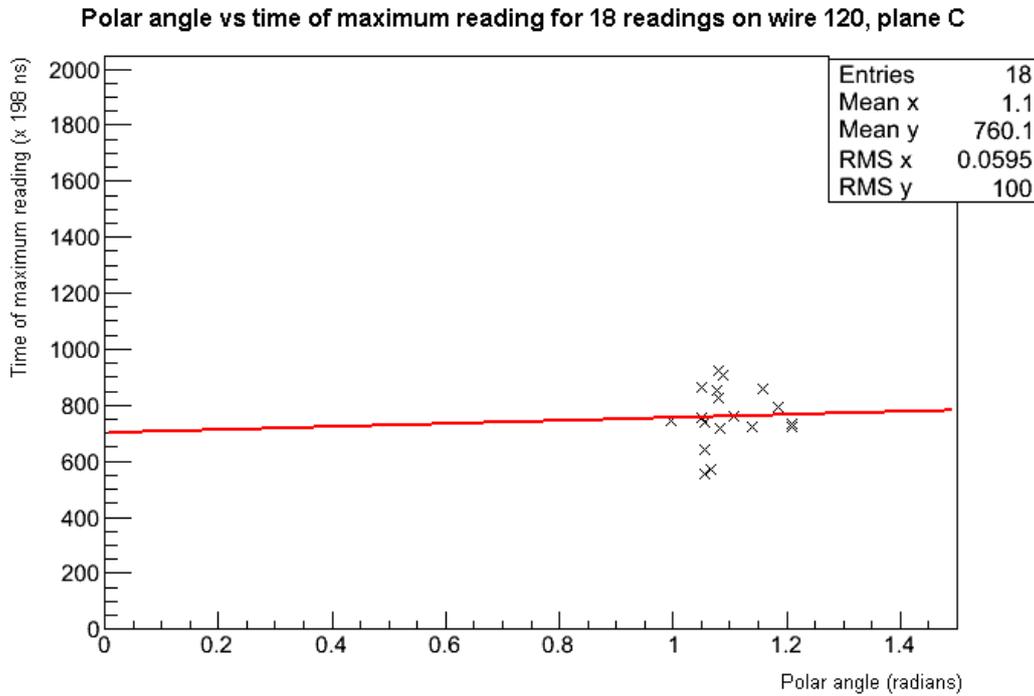


Fig. 11)

We initially found that the uncertainty on the y intercepts was too large for these results to be useful; moreover, because the uncertainty tends asymptotically towards a value of about 400 counts (80 microseconds) as shown in figure 11, this method would not produce useful results even were we to take many more readings.

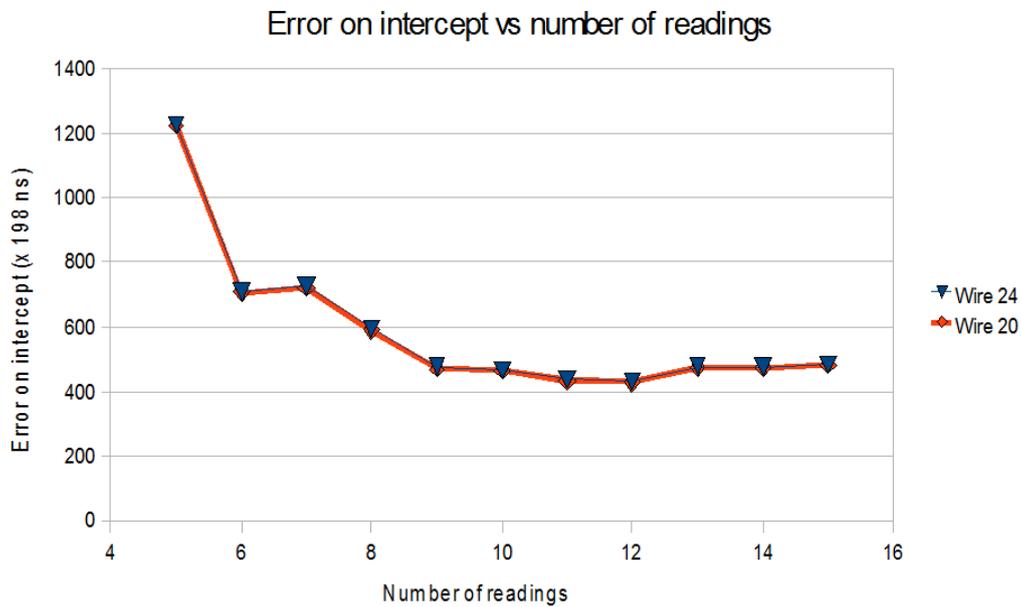


Fig. 12)

However, this result can be significantly improved if, instead of fitting the line directly to the data, we fit to x values with the average angle (in this case 1.103 radians) subtracted. As shown in figure 12, the uncertainty on the y intercept now tends asymptotically towards a value of about 25 counts (4 microseconds), which would give us an overall uncertainty of about 9%, or 117 ms^{-1} for a result of 1330 ms^{-1} .

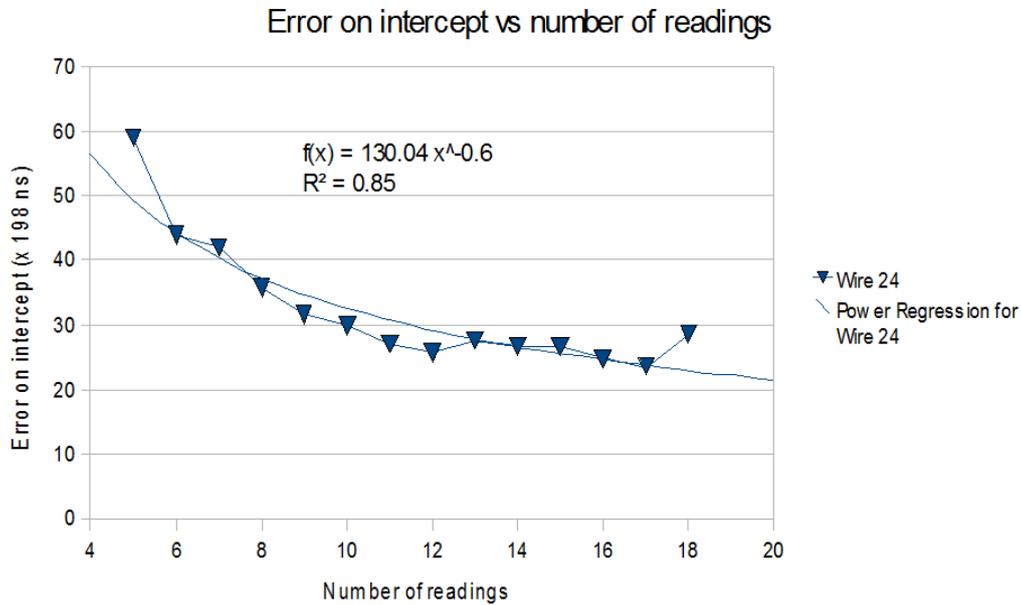


Fig. 13)

This figure suggests we need about twelve good events in order to minimise the uncertainty. The preliminary runs available to us had about eight good events. When this method is applied to this data, we obtain a result of $1400 \pm 300 \text{ ms}^{-1}$ - the uncertainty is still a little too large for the result to be useful. In order to minimise the uncertainty, we carried out this procedure over a number of different wires and averaged the results. Applying this strategy to seven wires produced an average drift velocity value of $1400 \pm 100 \text{ ms}^{-1}$, as shown in the table below.

Wire	Drift velocity	Error
10	1451.0	289.8
16	1454.6	301.1
20	1443.0	298.1
24	1410.0	276.3
28	1416.0	287.4
30	1403.4	283.9
35	1382.0	276.1
Average	1422.87	108.7

However, it might at first seem that there could be a circularity in this reasoning: we need to assume a value of drift velocity in order to calculate the polar angle, which is then used to calculate a value of the drift velocity.

In order to ascertain the influence of the initial assumption on the final value, we plotted the output drift velocity against the input drift velocity.

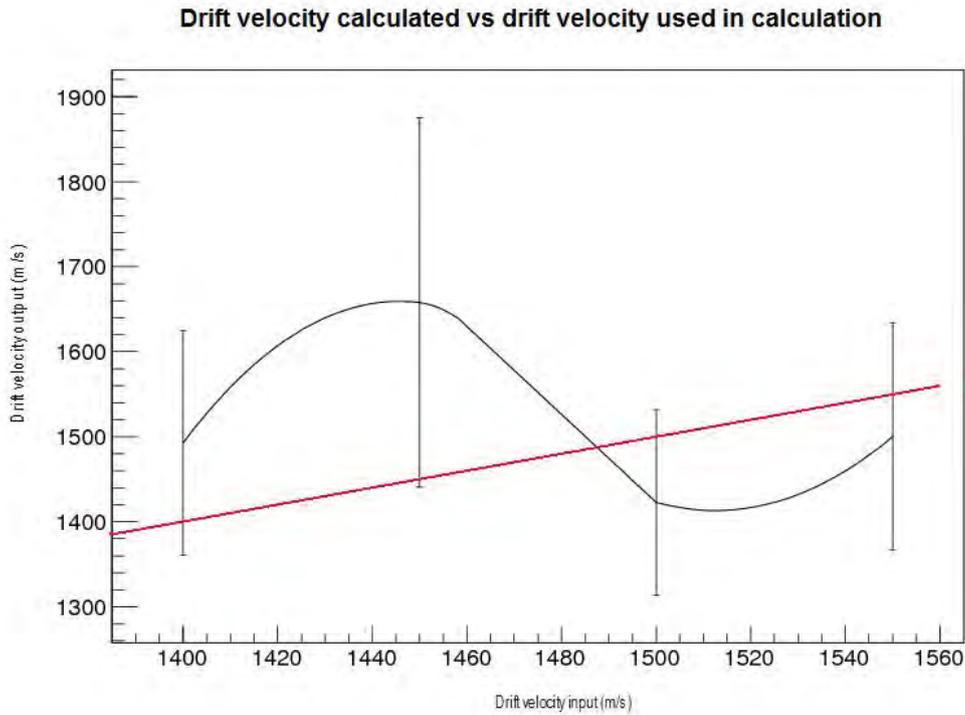


Fig. 14)

This plot was made by averaging the drift velocity obtained from a given input value of drift velocity over eight different events. It shows that the input and output values for drift velocity are approximately equal at a point between 1400 and 1500 ms^{-1} , which suggests that the true value for drift velocity lies somewhere in that range.

5 Purity Measurement

For any liquid argon TPC it is important to monitor the purity of the liquid argon. Electronegative impurities will reduce the number of electrons which arrive at the wire arrays in the signal, thus reducing and potentially obliterating the signal.

In order to measure the purity of the liquid argon used in these experiments, two sets of data were taken, with the positions of the counters inverted between them. The inversion changes the average distance travelled by the

electrons to arrive at any particular wire, and this is reflected by a change in the average time of the signal on each wire. The further the electrons have to travel through the argon, the greater the proportion of electrons that will be lost through interactions with impurities. Therefore a comparison of the change in the signal strength with the change in the time of the signal should yield some indication of the purity of the argon: the more contaminated the argon, the higher the degree of correlation between these quantities.

Difference in average max signal vs difference in average arrival time, on counter inversion

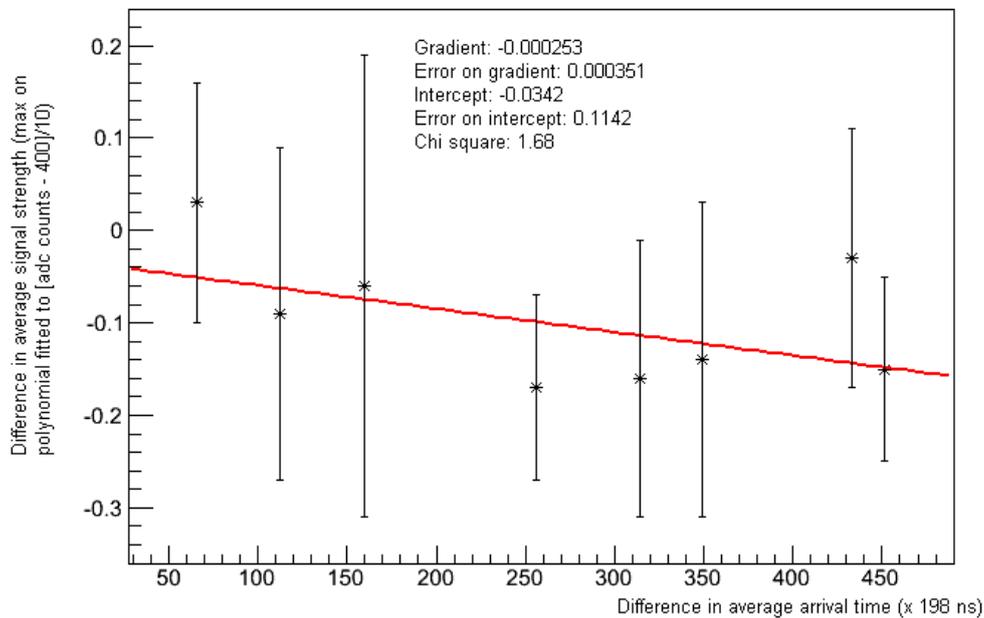


Fig. 15)

Figure 15) shows a graph of the change in average signal strength versus the change in average signal time, averaging over six events on eight different wires (each point represents a different wire). The error bars are calculated using the formula:

$$\sigma = \frac{RMS}{\sqrt{N}}$$

where the RMS is the characteristic RMS for the signal strength on that particular wire, and N is the number of readings used to find the average (6 in this case).

The gradient of the best fit line on this graph represents the change in the signal strength (in units of [adc counts]/10) per distance travelled through

the argon (in units of m). Observe that the uncertainty on the gradient is large, so that zero falls within its range. Therefore this data does not provide good statistical evidence in support of any diminution in the signal strength, or correspondingly any impurity in the argon.

If the argon were completely pure, we would expect 95% of experiments to measure a change in average signal strength which is within 2σ of zero, where σ is the standard deviation, given by the formula above. Therefore we can conclude only that the true change in signal strength on each wire is less than 2σ . The table below shows, for each wire tested, the value of 2σ , the change in average time (in units of 198 ns), and the upper limit on the change in signal per unit distance travelled through the argon (in units of [adc counts]/10 per m) calculated from those values of change in time and 2σ . This gives an indication of the degree of sensitivity of this test - any impurity in the argon must be such as to produce a change in signal per unit distance which is of the order of the values given in this table, or smaller.

Wire	Change in average time (x 198 ns)	2σ	Upper limit on attenuation constant
106	65.72	0.27	14.7
109	112.24	0.35	11.34
111	159.65	0.49	11.14
116	256.21	0.20	2.85
119	314.57	0.29	3.37
119	314.57	0.29	3.37
121	349.3	0.32	3.54
125	432.95	0.28	2.30
126	451.37	0.21	1.67

We could increase the sensitivity of this analysis by obtaining more data so that we could take averages over a larger number of events. This would decrease 2σ and thus decrease the upper limit for the change in signal per unit distance. To give an idea of the changes in sensitivity, figures 16a), 16b) and 16c) show how the values of 2σ vary with the number of events used to calculate the average.

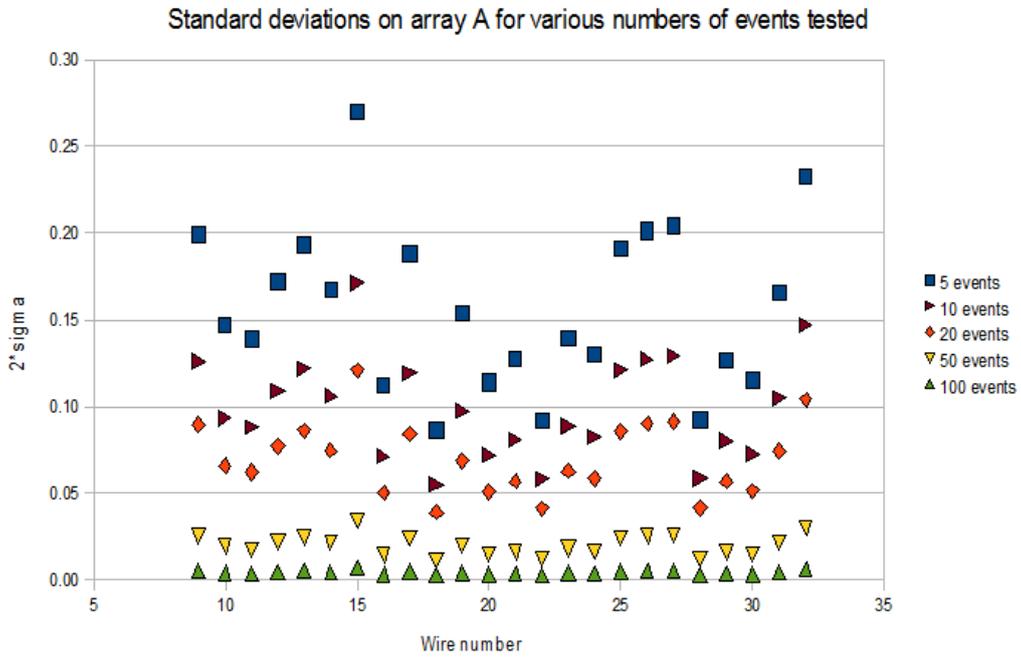


Fig. 16a)

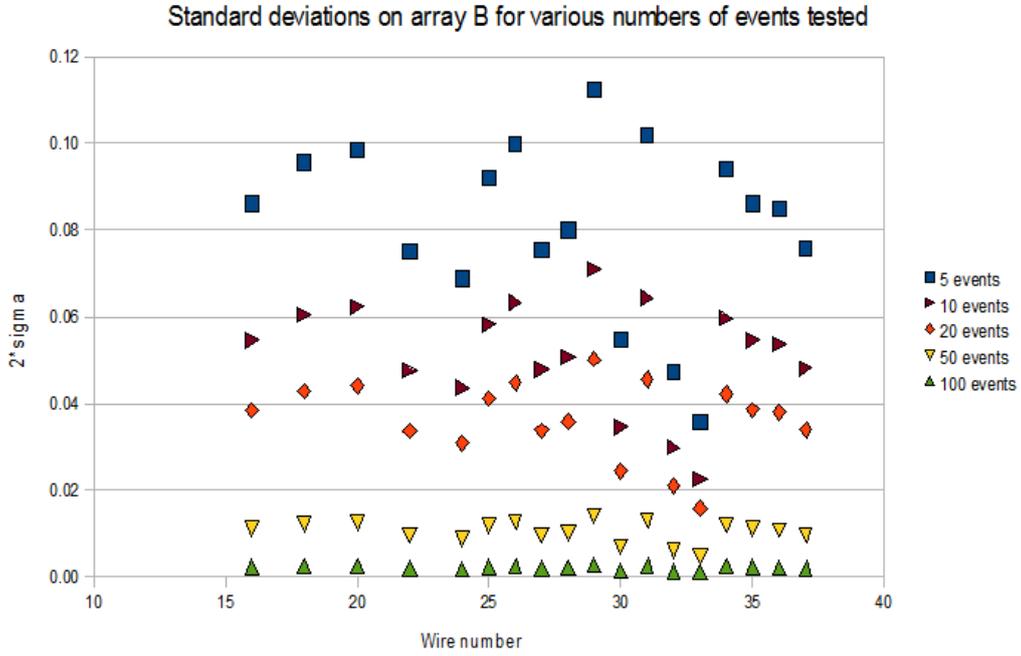


Fig. 16b)

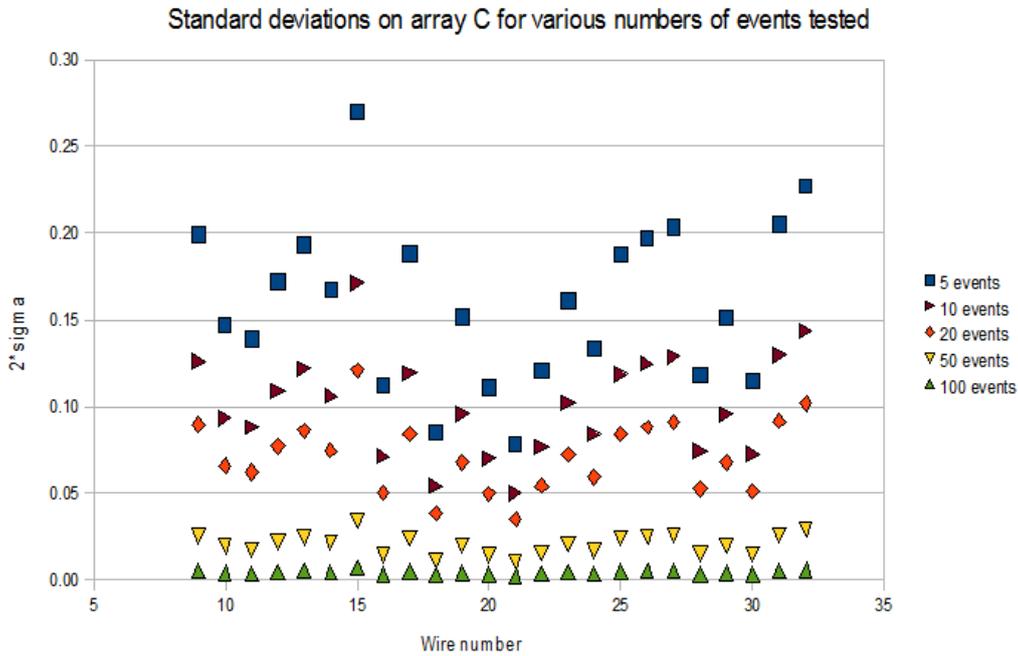


Fig. 16c)

Notice that not all wires are included on these plots. This is because wires at the outer edges of the arrays often fail to show any signal, and therefore we cannot calculate the RMS for these wires.

6 Automating Event Selection

About half of the events recorded by the Bo TPC are unsuitable for purposes such as finding the angle of a muon's path. The figures below show several events which cannot be used in this way. Figure 17a) shows an event where no muon passes through, figure 17b) shows an event where the muon gives off a second particle, and figure 17c) shows an event with a shower of particles rather than a single muon.

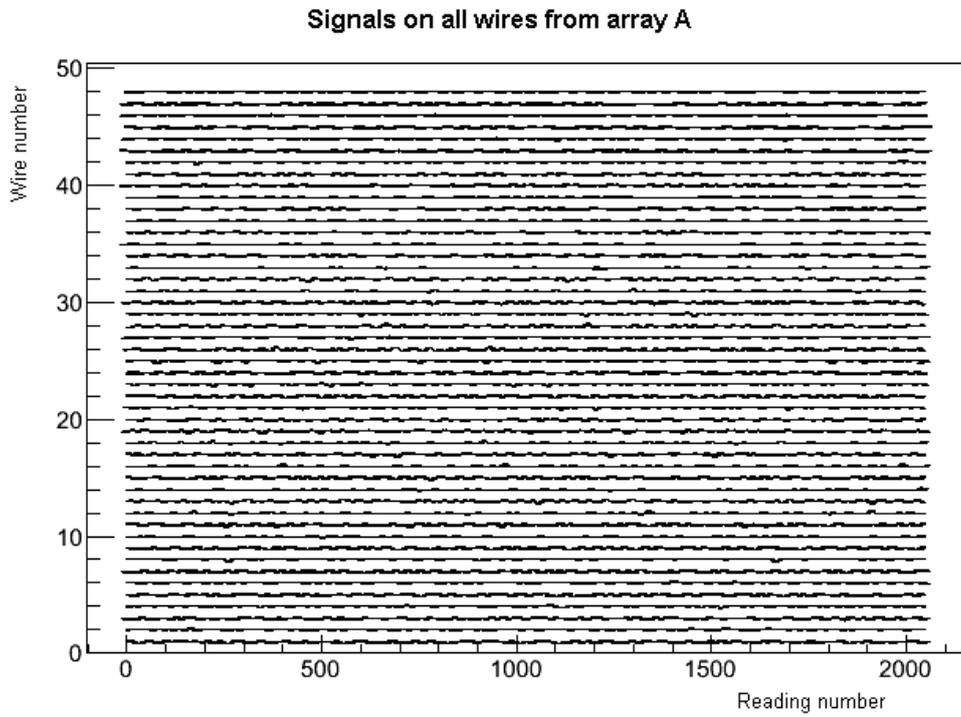


Fig. 17a)

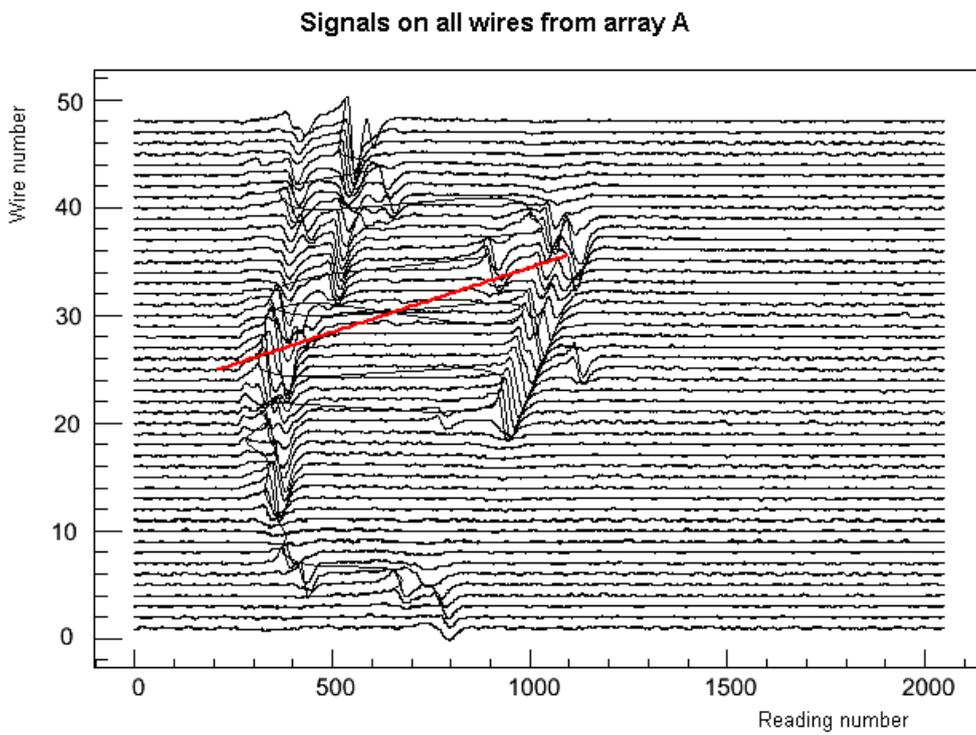


Fig. 17b)

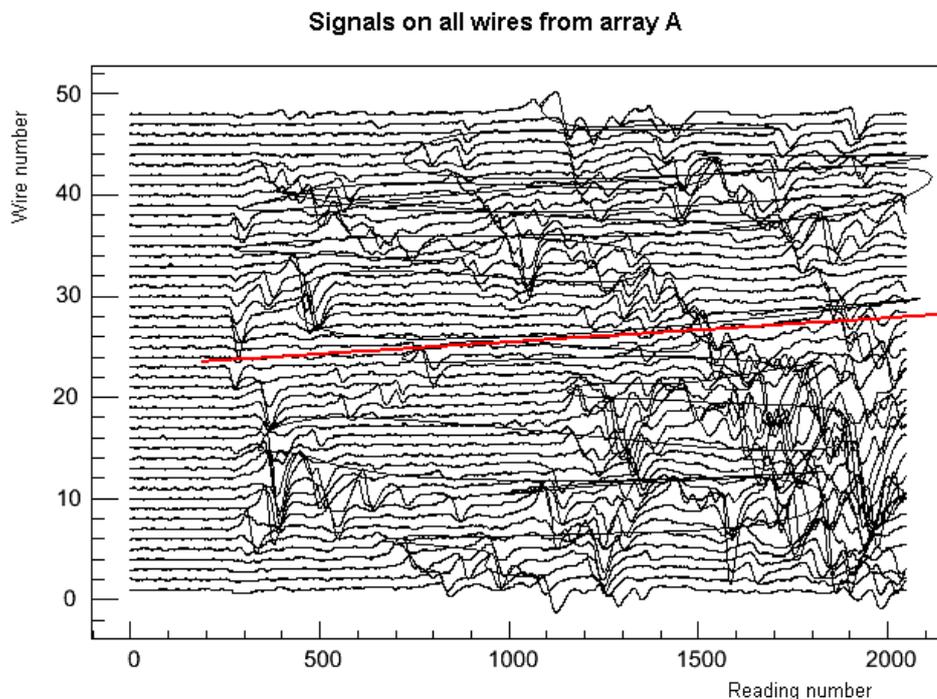


Fig. 17c)

It is possible to decide which events can be used by inspecting them by eye, but the process is lengthy. Therefore we developed a program which selects events automatically. The program carries out two checks on the data:

1) Test whether there is more than one maximum (i.e. is there more than one distinct string of data containing values greater than 406). Events having two or more such maxima are labelled bad.

2) Find the distance between the midpoint (the point where the data crosses the 400 line, between the maximum and minimum reading) and the corresponding position on the best fit line. Events where this distance is greater than 20 are labelled bad.

Events which pass both tests are labelled good.

6.1 Evaluation of program

In order to establish that the program was working properly, we compared the decisions made by the program with the comments made by a human inspecting the relevant events by eye. The results are in agreement in most

clear cases: here, 9 out of 46 events are labelled bad by both the program and human, and 19 out of 46 are labelled good by both program and human. In the remaining borderline cases, viewing the plot by eye does not yield a definite good or bad verdict, and in such cases the program gives variable results; this is understandable because in such cases there is no definite answer as to whether the event is usable or not. With more time, the program could be further developed to identify such borderline events as a separate category.

The table below shows the results for the first fifty events on run 202 (barring three events where the data cannot be loaded on the programs).

Event number (on run 202)	Visual inspection	Program verdict
1	Good	Good
2	Bad	Bad
3	Good	Good
4	Clear slope, track not clean	Good
5	Bad	Bad
6	Good	Good
8	Good	Good
9	Bad	Bad
10	Good	Good
11	Messy but perhaps usable	Good
12	Noisy but perhaps usable	Bad
13	No event	Bad
14	Good	Good
15	Bad	Bad
16	Good	Bad
17	Good	Good
18	Good	Good
19	Messy but perhaps usable	Good
20	Bad	Bad
21	Messy but perhaps usable	Good
23	Messy parts, perhaps usable	Good
24	Good	Good
26	Messy but perhaps usable	Good
28	Bad	Bad
29	Good	Good
30	Messy middle, but usable	Good
31	Bad	Bad

32	Good	Good
33	Messy but perhaps usable	Good
34	Bad	Bad
35	Messy in the middle	Bad
36	Good but a few bumps	Good
37	Bad	Bad
38	Good	Good
39	Messy but perhaps usable	Good
40	Messy but usable	Good
41	Good	Good
42	Good	Good
43	Good	Good
45	Good	Good
46	Good	Good
47	Good	Good
48	Messy but perhaps usable	Good
49	Good	Good

As an example of a borderline event, figure 18) shows event number 11 on this run, which a human observer classified as messy but perhaps usable, but the program labelled as Good. It can be seen that despite the messiness, there is a clear slope and the best fit line remains close to the data, so it is reasonable for this to be considered a good event.

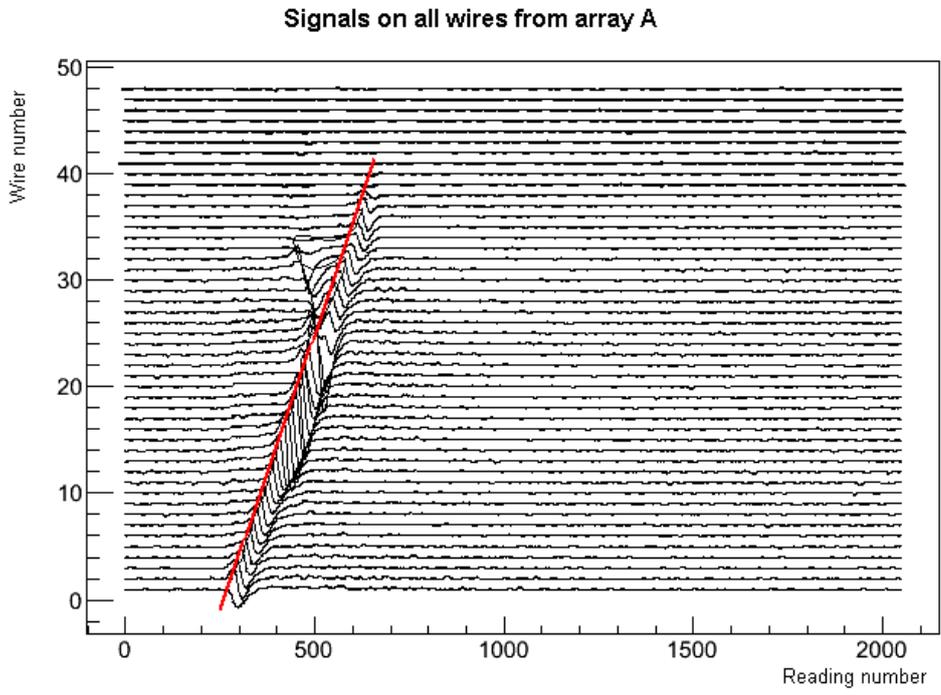
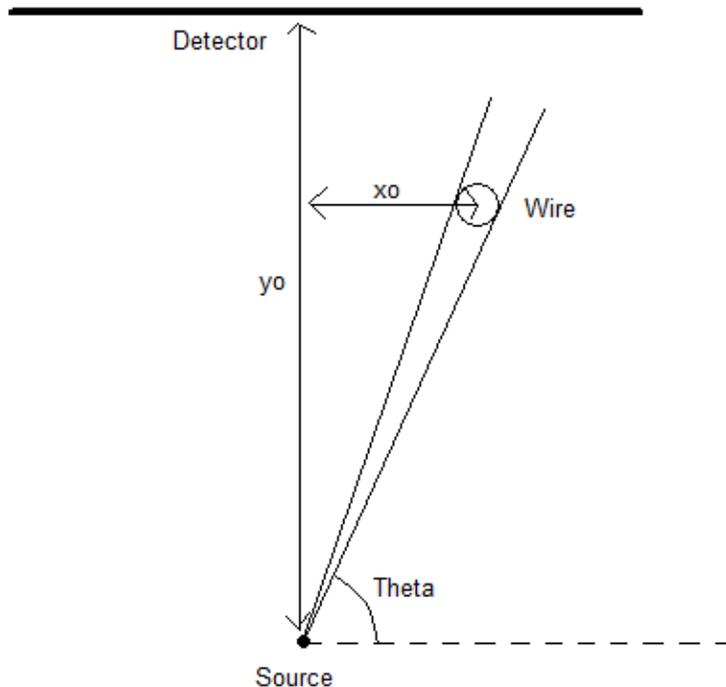


Fig. 18)

Light Blocked by Wire Arrays

Consider a photon detector of width 20 cm. Five planes containing arrays of wires are placed between the counter and the source. The wires have diameter 0.015 cm, and the distance in the plane between their centres is 0.5 cm. The spacing between the planes is 0.5 cm and the distance between the detector and the closest plane is also 0.5 cm. We would like to know what fraction of the light which would otherwise have reached the counter will be blocked by these wireplanes.



The experimental setup

My first step was to find the total angular range blocked by each of the planes. To do this, I first considered the light blocked by a single wire. I observed that all light rays falling between the two rays which are tangent to the outside of the wire at the point of their intersection with the wire will be blocked. Thus I wrote the following simultaneous equations:

(1)

$$y = \tan(\theta)x$$

(2)

$$dy/dx = \tan(\theta)$$

(3)

$$(x - x_o)^2 + (y - y_o)^2 = r^2$$

(4)

$$dy/dx = \frac{x_o - x}{y - y_o}$$

The origin of the coordinate system is at the source, with the x axis parallel to the detector and the y axis perpendicular to it.

(1) is the equation representing the ray travelling outwards from the source at an angle θ to x axis. (2) gives the gradient of this line.

(3) is the equation of the outside of the wire at a distance y_o from the source and x_o from the centre of the detector (r is the radius of the wire, 0.0075 cm). (4) gives the gradient of the outside of the wire for particular values of x and y .

Solving these equations for θ , we find

$$\theta = \sin^{-1}\left(\frac{\pm r}{\sqrt{x_o^2 + y_o^2}}\right) + \tan^{-1}\left(\frac{y_o}{x_o}\right)$$

(using $+r$ gives the angle further to the left, while $-r$ gives the angle further to the right).

I therefore wrote a spreadsheet giving a full list of the angular ranges excluded by the wires, for two different values for the distance between the source and the detector (20 cm and 100 cm). From this information I could easily calculate the total angular range blocked by any wireplane individually, as shown below.

Distance between plane and source	Sum of angular ranges excluded
17.5	0.03339
18.0	0.03225
18.5	0.03174
19.0	0.03097
19.5	0.03024
97.5	0.00630
98.0	0.00626
98.5	0.00623
99.0	0.00620
99.5	0.00617

However, I could not simply sum these results to find the total proportion of light blocked, as there is the possibility that the planes closer to the source might partially shade some of the planes further off, which would decrease the total proportion. Nor was it practicable to manually compare the angles excluded by each plane to determine the overlap. Therefore I wrote a program using C++ which took angles at intervals of 0.000001 across the full range of interest and determined the proportion of these which fell inside at least one of the angular ranges excluded by the wires.

For the range of interest I used all angles between $\frac{\pi}{2} + \tan^{-1}(\frac{10}{d})$ and $\frac{\pi}{2} - \tan^{-1}(\frac{10}{d})$ where d is the distance between the detector and source. Any light ray travelling at an angle outside this ray will miss the detector in any case, and so it does not matter whether it is excluded by a wireplane or not.

The results of this analysis are displayed below. The second column gives the percentage of light that would be blocked by the given plane individually, while the third column gives the percentage of light blocked by all the planes up to and including the given plane. It can be seen that the shading of later planes by earlier ones is quite minimal.

For a distance of 20 cm between detector and source:

Plane number	Source-plane distance	Light excluded by plane	Cumulative total
Plane 1	17.5 cm	3.113 %	Plane 1: 3.113 %
Plane 2	18.0 cm	3.098 %	Planes 1-2: 6.094 %
Plane 3	18.5 cm	3.114 %	Planes 1-3: 9.074 %
Plane 4	19.0 cm	3.114 %	Planes 1-4: 11.984 %
Plane 5	19.5 cm	3.114 %	Planes 1-5: 14.781 %

For a distance of 100 cm between detector and source:

Plane number	Source-plane distance	Light excluded by plane	Cumulative total
Plane 1	97.5	3.005 %	Plane 1: 3.005 %
Plane 2	98.0	2.990 %	Planes 1-2: 5.543 %
Plane 3	98.5	2.975 %	Planes 1-3: 8.066 %
Plane 4	99.0	2.959 %	Planes 1-4: 10.572 %
Plane 5	99.5	2.945 %	Planes 1-5: 13.065 %

Notice that the results in the second table, if converted to angular ranges, do not entirely agree with the corresponding results in the first table. This is because the first table gives the sum of all angles excluded by the given plane, whereas the second includes only angles between $\frac{\pi}{2} + \tan^{-1}(\frac{10}{d})$ and $\frac{\pi}{2} - \tan^{-1}(\frac{10}{d})$.

Positions for Photon Counters

Consider a scintillator sheet of radius 6" and a detector of radius 4". Assume that the angular distribution of the photons released from the scintillator is uniform. What is the optimum positioning of the scintillator above the detector in order to maximise the proportion of the photons released that will reach the detector? I used C++ programming to apply Monte Carlo methods to this integration. I selected a value for h which I intended to test. I chose a point R along the radius of the scintillator, then selected an azimuthal angle and a polar angle to characterise the direction of travel of the photon emitted, and calculated the coordinates of the point at which such a photon would intersect the plane containing the detector.

When the scintillator is at a height h above the detector, and a photon is released from coordinates (x,y,h) relative to the centre of the detector, with azimuthal angle ϕ and polar angle θ , it will intersect the plane of the detector at coordinates:

$$x_2 = R + h \times \tan(\theta)\cos(\phi)$$

$$y_2 = h \times \tan(\theta)\sin(\phi)$$

The photon enters the detector if the condition $x_2^2 + y_2^2 < 16$ is met.

I then carried out this test multiple times, varying the values of θ and ϕ separately in order to cover the entire range of photons emitted. I assumed that the angular distribution of photons was uniform. This meant that to generate values of ϕ I merely had to start from 0 and increment by my chosen interval (0.01) until I reached $\phi = \pi$. However, the frequency of θ must be weighted by a factor of $\sin(\theta)$ in order to cover a spherical surface uniformly. Therefore in order to obtain values for θ I defined a variable x such that $\theta = \cos^{-1}(1 - x)$; I then started from $x = 0$ and incremented x by the chosen interval (0.0001) until I reached $x = 1$. I used the results of these tests to obtain the proportion of photons emitted at radius R which would be expected to arrive at the detector.

Finally, I repeated this procedure for different values of R. The values of R ranged from 0 to 6, at intervals of 0.1. However, I also had to make sure to weight the proportion obtained at each different R, because there is a larger area corresponding to a given increment in radius at the edge of the scintillator than at the centre. To compensate for this effect, I multiplied the proportion obtained for each value of R by $2\pi R \times 0.1$ (i.e. the area corresponding to $R + dR$, with $dR = 0.1$ "). To obtain the final result, I then

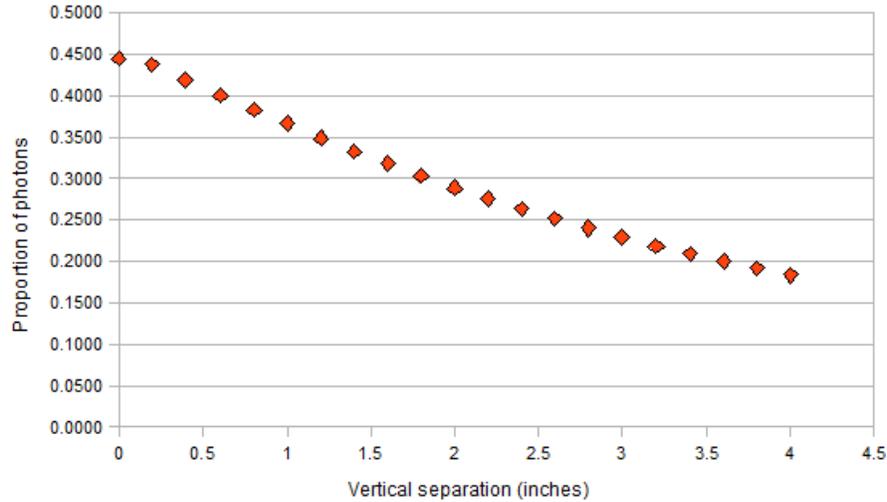
divided the sum of all these weighted proportions by the total area. Thus I used the formula:

$$proportion = \sum_{R=0}^6 \frac{p(R) \times \pi R \times 0.1}{\pi \times 6^2}.$$

I then applied the resulting program to a range of values of h, in order to ascertain the optimum positioning of the equipment. My results are displayed below.

Height	Proportion of photons detected
0	0.4444
0.2	0.4369
0.4	0.4184
0.6	0.4002
0.8	0.3825
1.0	0.3652
1.2	0.3485
1.4	0.3325
1.6	0.3171
1.8	0.3025
2.0	0.2885
2.2	0.2752
2.4	0.2626
2.6	0.2506
2.8	0.2393
3.0	0.2285
3.2	0.2183
3.4	0.2087
3.6	0.1996
3.8	0.1909
4.0	0.1828

Proportion of photons detected vs vertical distance between scintillator and detector



I conclude that the optimal experimental arrangement has the scintillator placed directly on top of the detector.

Discussion

I went through several stages in this calculation. My original intention was to integrate over the entire scintillator by incrementing x and y (rather than using R and weighting the results). I also intended to use a random number generators to produce appropriate values of ϕ and θ .

However, a number of tests on the accuracy of my results revealed significant problems. My first test was to see if the program produced the correct result for a vertical separation of zero between the scintillator and detector - in that case, the proportion of photons obtained should just equal the ratio of areas: $\frac{\text{detector}}{\text{scintillator}}$. I obtained 0.44, which is the expected result.

I next ran the program for a vertical separation of 100". At such a great separation, we would expect the proportion of photons obtained to be approximately equal to the area of the detector divided by the surface area of the half-sphere with radius 100 subtended by the scintillator. This did not prove to be the case. I then carried out a similar test but reduced the area of the scintillator to the single point, in case its finite area were making the approximation untenable, but the results still did not match.

After some investigation, I discovered that the cause of the problem was

the quantisation of the θ and ϕ values arising from the fact that the random number generator can only produce integer values. In order to obtain decimal values between 0 and π for ϕ , I had been generating random integers between 0 and 31420 and dividing the result by 10000; similarly, to obtain decimal values between 0 and 1 for x , I had been generating random integers between 0 and 10000 and dividing the result by 10000. This strategy limits the fine-graining of the values obtained for ϕ and θ , leading to inaccuracies.

Thus an obvious solution was to increase the range of numbers generated and the divisor by the same factor in both cases. However, this was not an option, because the random number generator can only produce values up to 32767. Therefore I decided to avoid using a random number generator, and instead simply took values of ϕ and θ and incremented them as described. However, this strategy significantly increased the run time of the programme; in order to decrease the time I therefore changed from integrating over the entire scintillator to integrating over R only and weighting the results appropriately.

The same checks of accuracy were carried out on this programme as on the previous one. It gives the expected result (0.44) for a height of zero. It also gives the correct result for a single point at a vertical separation of 100", but only provided that x is changed by increments of 0.00001. Such small increments are not practicable for the full program, because the run times become unmanageable; therefore I chose to use increments of 0.0001, meaning that the results obtained are only approximate. At a height of 100", the program predicts that a proportion 0.00085 of photons will strike the target, as compared to the proportion 0.0008 obtained from considering an area on the surface of the sphere, as described above. Thus the results should be taken as approximate (accurate to within a factor of about 7%).

I do not think this error should be a source of serious concern, because the trend - a steady decrease in the proportion of photons as the scintillator is raised - is clear, and indeed, the same trend was evidenced by all versions of this program. Therefore this error probably does not affect the reliability of my conclusion.