TRI PP 95 9
March 1995

A Trigger for the Identification of Pions Stopped in an Active Target

K.J. Raywood*, A.R. Ambardar*, J.B. Lange* and M.E. Sevior**

*Dept. of Physics, University of British Columbia,
Vancouver B.C., Canada V6T 1Z1
**School of Physics, University of Melbourne,
Parkville Vic., Australia 3052

Abstract

The total cross sections of the $\pi^+ p \rightarrow \pi^+\pi^+ n$ and $\pi^- p \rightarrow \pi^+\pi^- n$ reactions near threshold can be used to obtain scattering lengths which are directly comparable to predictions of chiral perturbation theory. Data for these reactions were taken at TRIUMF using a segmented active scintillator target. The signature of a π^+ in an active target segment was a prompt pulse caused by the particle stopping, followed by a second pulse due to the π^+ to μ^+ decay. TRIUMF 500 Mhz transient digitisers were used to record the scintillator outputs in 2 ns steps so that the double pulses could be identified with high efficiency in off-line data analysis. A second level trigger able to reject events in less than 10 μs was necessary to avoid a prohibitively high deadtime due to the long read-out time of the digitisers. It was implemented with fast ECLine electronics and increased the useful event acquisition rate by a factor of more than forty.

(Submitted to Nuclear Instruments and Methods A)

1. Introduction

The $\pi - \pi$ interaction is well known to be a sensitive test of low energy QCD. For example, chiral perturbation theory has been used to make predictions of the $\pi - \pi$ scattering lengths with uncertainties of 5% [1]. It has also been known for many years that the threshold behaviour of $\pi p \rightarrow \pi\pi N$ reactions is sensitive to these scattering lengths [2]. This was the motivation for TRIUMF experiment 561 of which the first phase concentrated on the $\pi^+ p \rightarrow \pi^+\pi^+ n$ reaction. The results of the first phase and a full description of the experimental technique have been published [3,4]. The second phase of the program concentrated on the $\pi^- p \rightarrow \pi^+\pi^- n$ reaction although some data for $\pi^+ p \rightarrow \pi^+\pi^+ n$ were also accumulated. This second phase utilised a number of novel techniques which are reported here.

The experimental arrangement is shown in fig. 1. The pion beam was defined by three plastic scintillators S1, S2 and S3. The small scintillator S3 defined the beam spot well within the target dimensions while scintillators S1 and S2 were large enough to ensure that pions from preceding and following beam bursts were detected with high efficiency. The target consisted of five plastic scintillators placed in a stack along the beam direction. Another scintillator S4 was located behind the target and was used as a veto counter to define beam interactions within the target. Neutrons were detected in an array of large volume scintillator bars centered at zero degrees several metres downstream. The beam pions that do not interact were deflected away from the bars by a dipole magnet between the target and the array.

The first level of triggering for the system was a coincidence between a beam pion interacting in the target ($S1 \cdot S2 \cdot S3 \cdot \overline{S4}$) and a neutron detected in one of the bars. A neutron was defined by a signal from each end of one of the bars with the additional requirement that a charged particle was not detected in one of the veto scintillators in front of the neutron bars. This trigger was dominated by background events due to $^{12}C(\pi,n)$ reactions as well as random coincidences between a signal in the neutron bars and an interaction in the target. Thus additional methods of determining which events were due to $(\pi, 2\pi)$ reactions were needed.

The kinematics of the reactions near threshold constrained the neutrons to be emitted into a narrow cone around zero degrees and so the neutron bars had a large acceptance for the reaction neutrons. In addition, a large fraction of the pions from the reaction of interest were stopped in the active target since they were created just above threshold. Thus a coincidence between a positive pion stopped in the target and a neutron detected in the array was used to identify events from the reaction of interest. In this way the restrictive kinematics of the threshold reactions were used to suppress background reactions.

In order to be able to determine the cross sections with the desired precision, it was necessary to be able to identify stopped positive pions with good efficiency and to be able to determine this efficiency with equal or better precision. Stopped positive pions were identified by the decay sequence $\pi^+ \rightarrow \mu^+\nu_\mu$ with the muons having a kinetic energy of 4.1 MeV. These have a range of 1.5 mm in the scintillator and Monte Carlo studies show that 80% of them are contained within the same segment as the stopped pion. The expected signal in a segment of the active target was a large prompt pulse produced by the beam pion and the stopped pion followed by a second pulse corresponding to the characteristic stopped 4.1 MeV muon. The time interval between the two pulses is of course random according to the radioactive decay law with the 26 ns

average lifetime of the charged pion. Thus the efficiency of stopped pion identification depends critically on how well two pulses can be resolved at short time intervals.

Several methods for identifying the stopped pion signature were developed. The most powerful technique was to use 500 MHz transient digitisers to record the scintillator outputs in 2 ns increments. The existence of a double pulse could then be determined in off-line data analysis using sophisticated curve fitting techniques. In order to achieve reasonable identification efficiency, the sampling time was long enough to include a large fraction of the decay muons. Thus the event size was large and therefore had a long read-out time. In addition the first level trigger rate was typically 3 kHz. The combination of large event size and high rate would have created a prohibitively high system deadtime when used with the first level of triggering alone. However, the TRIUMF-developed transient digitisers can be triggered and cleared at a rate of 2 MHz thus permitting the use of a second level trigger. To be effective, the second level trigger needed to be capable of rejecting, in a short time, a large fraction background events passed by the first level trigger.

2. The second level trigger

The second level trigger was designed to reject events in which no detectable second pulse occurs in any of the active target segments. In the context of this experiment, detectable refers to being able to identify a second pulse from the transient digitiser data. The algorithm chosen to achieve this task was the following. For every first level trigger, the signal from each target segment was integrated over a long period, 80 ns, corresponding to three pion lifetimes. The same signal was also integrated over a short period, 15 ns, corresponding to the duration of a prompt pion induced pulse. If the difference between these two integrations for a particular target segment was sufficiently small, then no detectable second pulse came from that segment. An event was rejected by the second level trigger when all target segments had been tried and none showed the presence of a second pulse. However, if the difference between the long and short integrations for one of the target segments was greater than some chosen critical value, then that event was passed by the second level trigger.

2.1. Overview of circuit operation

The second level trigger was constructed primarily from LeCroy ECLine trigger modules [5] . The integrations were performed by a Lecroy 4300B Fast Encoding and Readout ADC (FERA) module which contained sixteen channels of charge integrating ADC's. It is triggered by a single gate input which has a minimum duration of 50 ns and must precede the input signal by at least 20 ns. To make the long and short integrations with just one FERA module it was necessary to truncate each signal over the appropriate long and short periods. This operation was achieved with Philips 745 Linear Gates which passed the input signal unaltered while the gate was open and produced zero output when the gate was closed. The long and short gated signals from each target segment were integrated by the FERA which was operated with both pedestal subtraction and data compression enabled. This meant that integration values below the preset pedestals did not appear on the output. Thus, setting the pedestals on the six unused channels to the maximum value and the others to zero ensured that

the output stream always contained the five values of the long integrations and the five values of the short integrations.

The short integrated ADC channels were normalised to the long integration channels to correct for the different signal fractions seen by the two ADC's so that both had the same value when no second pulse was present. The cutoff value was added to the normalised short integration value at this point so that the difference between the long integration value and this sum was positive if a second pulse was present and negative otherwise. The normalisation and addition of the critical value was achieved through the use of a LeCroy 2372 Memory Lookup Unit (MLU) which was preloaded with the desired output for every possible input combination. The subtraction of the MLU output from the long integration value was performed by a LeCroy 2378 Arithmetic Logic Unit (ALU). A negative result, and hence no second pulse, produced an output carry (borrow) signal from the ALU. The full schematic diagram of the second level trigger is shown in fig. 2.

2.2. Detailed description of data flow

The REQ (REQuest readout) output of the FERA went true upon completion of conversion and indicated that the module contained valid data. This was fed to the REN (Readout ENable) input which began the data transfer. The first word to appear on the ECL port output of the FERA was the header word which was identified by having the most significant bit (bit 16) set. Following that were the data which had the channel number in bits 12 to 15 and the value of the integration in bits 1 to 10 (the FERA used in this experiment had ten bit resolution). The particular arrangement used here ensured that the long integration data were transferred first. The ECL port fed a bus connected to two LeCroy 2375 Data Stacks. These were configured to accept only those data words that had a preselected combination of the three most significant bits. This feature enabled the header word to be ignored (bit 16 clear) and the data corresponding to the long integrations (bit 15 clear) to be stored in one data stack while those corresponding to the short integrations (bit 15 set) were stored in the other. The hand-shaking signals that control that data flow between the FERA and the data stacks are also shown in the diagram. When the data transfer was complete, REQ went false and the REN input was routed to the PASS output. The REQ signal was delayed on route to the REN input so that a pulse was produced on the PASS output at the end of the data transfer. This was connected to the ADI (All Data In) inputs of the data stacks and signified the end of the data stream.

The first word written into a Data Stack initiated a read cycle. Thus the ten data bits of the first long integration value were fed to the register A input of the ALU. The nine most significant bits of the first short integration value (bits 2 to 10) and the three bits of the segment identifier (bits 12 to 14) were fed to the lowest twelve input bits of the MLU. Only nine data bits were used because the available MLU's had a maximum of twelve input bits for sixteen output bits. The MLU output was fed to the register B input of the ALU and the RDY (data ReaDY) signal of the MLU was connected to the Strobe B input. The ALU was operated in subtraction mode (A – B) with both input registers latched on a leading edge to Strobe B (B strobes both). If the result of the subtraction was negative the COUT (Carry OUT) went true and was false otherwise. The DR (Data Ready) signal from the ALU was normally true but went false during processing and then returned to true to indicate valid data on the

output. In order to obtain a pulse, this signal was inverted and delayed by the ALU processing time. It was then used in combination with the COUT to produce either a second level ACCEPT when COUT was false, in which case processing terminated, or a pulse to the RE (Read Enable) of each Data Stack when COUT was true. The latter initiated the next read cycle causing the long and short integrations of the next target segment to be tested. Note that for the first four read cycles the ROF (Read OverFlow) from each Data Stack was false but went true when the last data words were issued. If a second level ACCEPT did not result from testing these, then a REJECT pulse was produced and processing terminated.

The time for the second level trigger to make a reject decision was dominated by the FERA conversion time. For the ten bit model used in this experiment this was 4.8 μs. With the data compression mode enabled there was an additional 2.3 μs before readout began. The five long integration values were then transferred to the first data stack with each cycle taking 100 ns. The five short integration values were transferred next but, as previously mentioned, main trigger processing was initiated by the first word written to the second data stack. The processing time for each loop was the propagation time through the data stack, the MLU, the ALU, and the logic gates. This was just under 200 ns and so all five loops were executed in 1 μs. Thus the total time taken for a reject decision was 8.6 μs and so the second level trigger could operate at a maximum rate of about 110 kHz. The first level trigger rate was typically 3 kHz and so the second level trigger introduced a deadtime of about 3%.

2.3. Calibration

For calibration purposes, data were collected with the MLU loaded with zeros for every possible input value. In this way there was never a negative result from the subtraction and a second level ACCEPT was issued for every first level trigger. The FERA data were read out through CAMAC and a scatter plot of the long integration value versus the short integration value was constructed for each target segment. An example of these is shown in fig. 3 for target segment *b*. The signals at large pulse heights were caused by charged particles from nuclear interactions initiated by the pion beam. The first level trigger preselected such events, the vast majority of which had no second pulse. The straight line displayed in the plot is a least squares fit to the points in the main band. This line was used to define the normalisation of the short integration values to the long ones for the target segment.

The differences between the long integration values and the normalised short integration values were calculated for all target segments. The presence of a second pulse was indicated by an excess in the long integration value. The critical cutoff value was chosen to exclude most of the events from the main band. The normalisation constants and cutoff values for each target segment were then used as input to a program which loaded the MLU in readiness for the main data-taking runs. Each possible input address to the MLU corresponded to a target segment number in the three most significant bits followed by the top nine bits of a short integration value. For every possible input address, the program calculated the sum of normalised short value and the cutoff value for the appropriate target segment and loaded the result into the corresponding address of the MLU.

2.4. Performance

The value of the second level trigger was indicated by the increase in the rate of acquisition of useful events compared with the same system without the second level trigger. A useful event was one that could be identified as being due to a stopped pion. There was a substantial fraction of good stopped pion events that could not be identified as such. Most of these were from pions that decayed so quickly that the muon pulse could not be resolved from the pion pulse. Thus, though it was necessary to know the ratio of useful events to all stopped pion events, there was no loss of information by rejecting non-useful events with the second level trigger.

A histogram of the difference between the long and normalised short integration values was constructed for each target segment during the analysis of the data. An example of these for target segment *e* and an incident π^+ energy of 200 MeV is shown in fig. 4. All recorded events must have been accepted by the second level trigger so those with an insufficient difference from this target segment must have had an excess in another segment. Events with a muon-induced second pulse from the decay of a stopped pion were identified in the analysis of the transient digitiser data as described in section 3.2. The distribution of long minus normalised short integration values of events that were identified as a stopped pion in this target segment is also shown in fig. 4 as the shaded histogram. These are the useful events and clearly an insignificant fraction of them were rejected by the second level trigger.

The increase in useful event acquisition rate is a consequence of the reduction in system deadtime due to the large number of first level triggers that were rejected. With a first level trigger rate of 2.8 kHz and no second level trigger, the system deadtime was a crippling 98%. With the inclusion of the second level trigger in the system, 96% of the first level triggers were rejected and the the deadtime was reduced to 14% indicating that the useful event acquisition rate was increased by a factor of about forty-three. The first level trigger rates, second level trigger rejection fractions and the corresponding deadtimes for various data-taking runs are shown in table 1.

3. The transient digitisers

The 500 MHz transient digitisers developed at TRIUMF have been described in detail in ref. [6]. Each was based on a 128-cell gallium arsenide charged coupled device (CCD). A 1 GHz signal generated by an external radio frequency synthesiser was prescaled internally to 500 MHz and used as the clock signal to drive the acquisition of the input signal into the CCD. Upon receipt of an external trigger pulse the prescaler switches to division by 256, thus providing a low frequency (approximately 4 MHz) clock signal for output of the captured signal. This was passed to an eight bit ADC which provided the digital output stream. An external abort signal could also be applied to the prescaler to asynchronously reset it to the high frequency and thus resumption of input signal accumulation.

The transient digitisers were developed primarily for use in experiment 787 at Brookhaven National Laboratory. The binary output stream is read out through FASTBUS in that experiment. At the time of running the experiment described here, such a read-out system was not available. However, the transient digitisers also provided another form of output. The binary stream was sent internally to a digital to analogue converter (DAC) giving an analogue output signal at low impedance that is stretched in

time, relative to the original input signal, by a factor of 128. This signal was redigitised with commercially available high-frequency waveform digitisers whose output streams were read out via CAMAC.

3.1. Operation

A block diagram of the trigger scheme for the transient digitiser readout is shown in fig. 5. A first level trigger stopped the CCD input accumulation and triggered the low frequency output. The first level trigger signal was also delayed by 18 μs before being sent to the waveform digitisers to stop their input accumulation and initiate their readout. This delay time ensured that at least 140 ns of the original signal was sent to the waveform digitisers. The second level trigger processing was also initiated by the first level trigger. A second level accept was used to trigger the complete event readout by the data acquisition computer. However, a second level reject blocked the stop signal to the waveform digitisers and aborted the transient digitiser output. Recall that the second level reject pulse arrived at a fixed time of about 8.6 μs after the first level trigger thus making it possible to block the waveform digitiser stop signal as shown in the diagram.

The stretched signals from each target segment were redigitised with LeCroy quad waveform digitisers and the resulting data stream was read out through CAMAC. Two such modules, (Lecroy models 2262 and 2261), were needed to accommodate the five target segments and the extra three channels were used to record the signals from scintillators S1, S2 and S3. They are also based on CCD's and can operate at frequencies from 8 MHz to 40 MHz. To avoid pulse shape inconsistencies arising from the redigitisation, it was necessary to ensure that the clock signal to the waveform digitisers was synchronised with the low frequency clock signal used by the transient digitisers for output. The 1 GHz RF signal was fed to an external divide by 128 prescaler to provide an 8 MHz clock signal that was phase locked to the 4 MHz clock signal of the transient digitisers.

The digitised scintillator output of a sample event is shown in fig. 6. The data have been subtracted from a baseline that was obtained by intermittently forcing the data acquisition system to record the transient digitiser output when no signal was observed in any of the beam defining scintillators. Thus even though the photomultipliers produce negative pulses, they are shown as positive here. The presence of a second pulse in target segment e is clearly seen in these data.

3.2. Analysis of the digitised signals

The data from the transient digitisers were analysed by fitting a linear superposition of average pulse shapes with the height and time of each pulse shape as free parameters. To obtain starting values for the fit, the digitised scintillator output was scanned for local maxima (with a limit of five). The first of these was identified with the incoming pion. Of the remaining local maxima, the one with the greatest pulse height was assigned the role of a muon candidate. If no second local maximum of sufficient height was found then a typical height and a small time offset were used for the initial values of the second peak. The non-linear least squares fitting routine MRQMIN [7] was employed for the fits to the transient digitiser data. The shape of the peaks was assumed to be some average pulse shape whose magnitude was proportional

to the energy deposited in the scintillator. It was assumed to be independent of height and represents the response of the phototube to a burst of scintillation light due to the passage of a charge particle. The average pulse shape in each target segment was determined from the transient digitiser data of events in which a beam pion interacted in the target without the production of a stopped π^+. Two examples of fitting the sum of two peaks to the transient digitiser data are shown in fig. 7.

The pulse heights were calibrated [8] to energy in units of MeV by normalising the average pulse height with the average energy lost in the scintillator by passing 200 MeV pions. The height of the second peak in the digitised scintillator output should fall within a narrow range of values as it was due to the 4.1 MeV muon from the decay of a stopped pion. A histogram of the heights of the second pulse in target segment e is shown in fig. 8. The concentration of pulse heights due to pion decay muons is clearly seen. Also shown are the upper and lower limits on this quantity that were used to identify the stopped pion signature in the subsequent analysis. A histogram of the time between the first and second pulse is shown in fig. 9. The best fit of an exponential decay curve to the data between the shown limits resulted in a half-life parameter of 26.4 $\pm$ 0.5 ns, in agreement with the charged pion half life of 26.030 ns. The lower limit of 16 ns in the decay time was due to the unreliability of the fitting procedure for peak separations less than this. The upper limit of 64 ns was due to the inability of the procedure to fit pulses that were not entirely digitised.

3.3. Stopped pion detection efficiency

In the measurement of the near threshold cross section for $\pi^- p \rightarrow \pi^+ \pi^- n$ both a π^- and π^+ could stop in the same target segment. The π^- is rapidly absorbed into a carbon nucleus where it liberates its rest mass energy via any one of a number of reactions thus creating a range of prompt pulse heights. Thus the ability of the system to detect the positive pion with known efficiency that is independent of the prompt pulse height is a critical requirement of the experiment.

Stopped pions were identified as those events where a second pulse, with decay time and peak height within the specified limits, was resolved in fitting the transient digitiser data. As mentioned earlier, the number of such events that were rejected by the second level trigger was insignificant. Therefore to determine the stopped pion detection efficiency for each target segment, it was sufficient to determine the ratio of the number of identified events to the actual number of stopped pions. These were obtained in separate runs where a low energy (30 MeV) π^+ beam was stopped in each target segment by means of an energy degrading system consisting of a selection of aluminium sheets of various thickness placed upstream of the beam scintillators. Data were collected with the second level trigger disabled and with five different aluminium sheet thicknesses, corresponding to pions stopped in each target segment. The number of pions that stopped in a particular target segment was obtained by counting only those events where a signal was present in all of the preceding scintillators but in none of the subsequent scintillators. The digitised target scintillator output of each of these events was fitted with a linear combination of two peaks and those whose second-peak height and decay time were within the assigned limits were counted as being identified. The resulting efficiencies for each target segment are shown in the last column of table 2. Also shown are the minimum and maximum decay times used to define an identified stopped pion and the expected fraction of pions with a decay

time between these limits.

The measured efficiencies are about 10% lower than expected. Part of this must be due to an intrinsic inefficiency in the curve fitting procedure as a result of the inability to obtain a satisfactory fit to pulses that are not similar enough to the average shape. There is also a problem with the technique of determining the true number of stopped pions as it was necessary to count only events with a sufficiently large prompt pulse to eliminate pions that stopped near the edge of a target segment with the decay muon leaking into an adjacent segment. Thus an uncertainty of about 10% must be assigned to the measured stopped pion detection efficiencies. Nevertheless, the results indicate that the system is capable of identifying most stopped positive pions with decay times between about 16 ns and 64 ns regardless of the height of the prompt pulse.

4. Conclusions

The second level trigger developed for TRIUMF experiment 561 utilised the high speed ECLine trigger modules developed by LeCroy. It was well matched to the performance of 500 MHz CCD transient digitisers and was capable of selecting candidate stopped π^+ events with well known efficiency and high speed. The measurement of the cross section of $\pi^-p \rightarrow \pi^+\pi^-n$ near threshold required this trigger.

Acknowledgements

The authors thank all the collaborators of TRIUMF experiment 561 for their contributions to the experiment. We are also grateful to D. Bryman, J.V. Cresswell and M. LeNoble for allowing us to use the transient digitisers. We express our gratitude to N. Shibaoka for designing the frequency divider. Finally the authors gratefully acknowledge financial support from NSERC which made this project possible.

References

[1] J. Gasser and H. Leutwyler, *Phys. Lett.* **125B** (1982) 312.

[2] M.G. Olsson and L. Turner, *Phys. Rev.* **181** (1969) 2141.

[3] M.E. Sevior, A. Ambardar, J.T. Brack, P. Camerini, F. Duncan, J. Ernst, A. Feltham, N. Grion, R.R. Johnson, G. Koch, O. Meirav, D.F. Ottewell, R. Rui, G.R. Smith, V. Sossi, D. Theis and D. Vetterli, *Phys. Rev.* **D48** (1993) 3987.

[4] M.E. Sevior, A. Ambardar, J.T. Brack, P. Camerini, F. Duncan, J. Ernst, A. Feltham, N. Grion, R.R. Johnson, G. Koch, O. Meirav, D.F. Ottewell, R. Rui, G.R. Smith, V. Sossi, D. Theis and D. Vetterli, *Phys. Rev. Lett.* **66** (1991) 2569.

[5] L.B. Levit, G.J. Blanar and M.L. Vincelli, *IEEE Trans. Nucl. Sci.* **33** (1986) 925.

[6] D. Bryman, J.V. Cresswell, M. LeNoble, R. Poutissou, *IEEE Trans. Nucl. Sci.* **38** (1991) 295.

[7] W.H. Press, *Numerical Recipes: The Art of Scientific Computing*, Cambridge University Press, Ch. 14 (1988).

[8] J.B. Lange, *M.Sc. Thesis*, (1992), University of British Columbia.

Table 1. The first level trigger rates, second level trigger rejection fractions and the corresponding deadtimes for various data-taking runs.

Beam Polarity	Energy (MeV)	1st Level Rate (kHz)	2nd Level Rejection (%)	Deadtime (%)
+	200	1.3	95	15
+	184	1.2	94	14
+	172	1.2	94	15
−	200	2.6	96	16
−	184	2.6	97	13
−	172	2.8	96	14

Table 2. The software limits on decay time used to identify stopped pions, the expected fraction of pions with decay times between these limits and the measured stopped pion detection efficiencies for each target segment.

Target Segment	Minimum Decay Time (ns)	Maximum Decay Time (ns)	Expected Fraction	Stopped π^+ Detection Efficiency
a	16.4	63.4	0.44	0.40
b	15.8	65.6	0.46	0.42
c	15.8	66.8	0.47	0.43
d	17.8	63.6	0.42	0.39
e	16.6	63.6	0.44	0.43

Figure captions

Fig. 1. The experimental arrangement used to measure the cross sections of the $\pi^+p \rightarrow \pi^+\pi^+n$ and $\pi^-p \rightarrow \pi^+\pi^-n$ reactions near threshold. The aluminium degrader was only used to stop pions in a specific target segment during calibration runs.

Fig. 2. A schematic diagram of the second level trigger.

Fig. 3. A scatter plot of the long integration value versus the short integration value for target segment *b*. Also shown is the straight-line fit to the main band that was used to define the normalisation of the short integration values to the long ones. The second level trigger was disabled to collect these data.

Fig. 4. The distribution of the difference between the long integration value and the normalised short integration value for target segment e and an incident π^+ energy of 200 MeV. The upper histogram is due to all events that passed the second level trigger and the shaded area indicates the events that were identified as containing a stopped pion in target segment e.

Fig. 5. A block diagram of the the transient digitiser read-out scheme utilising the first and second level triggers.

Fig. 6. The digitised scintillator output of a sample event.

Fig. 7. Two examples of fitting the sum of two peaks to the transient digitiser data.

Fig. 8. A histogram of the heights of the second pulse in target segment e.

Fig. 9. A histogram of the time between the first and second pulse.

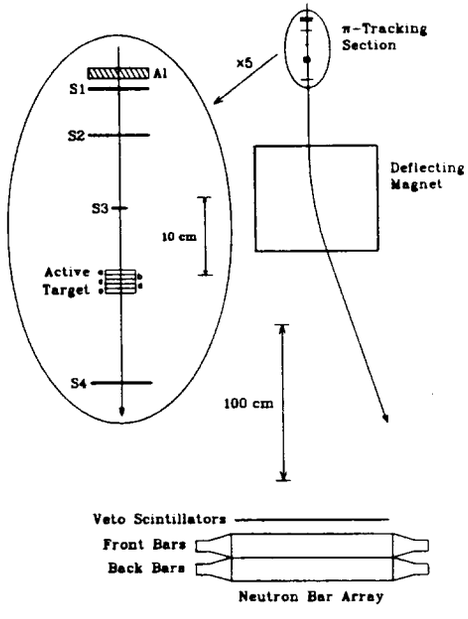

Fig. 1

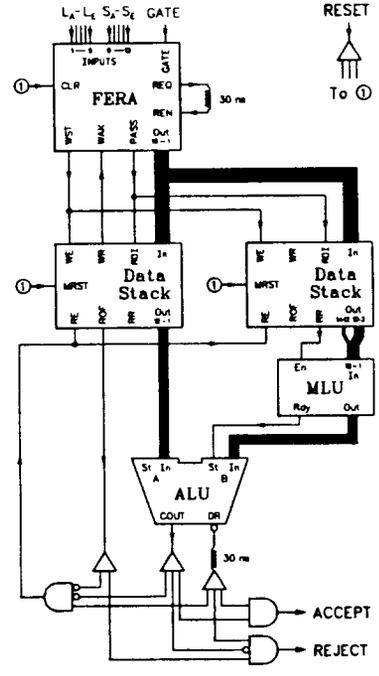

Fig. 2

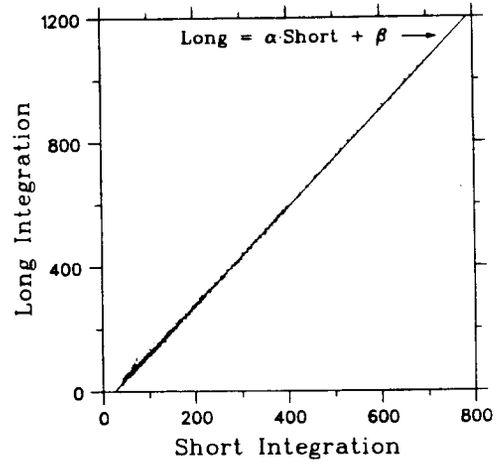

Fig. 3

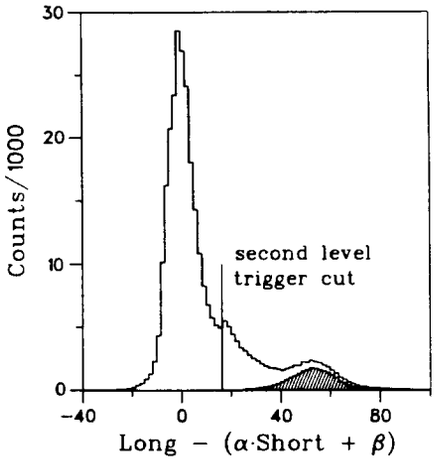

Fig. 4

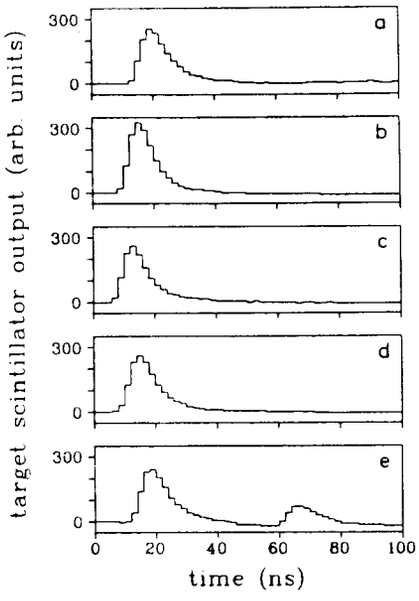

Fig. 6

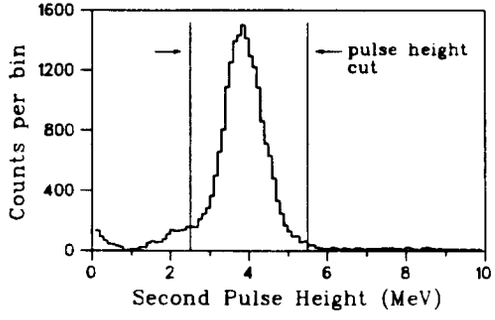

Fig. 8

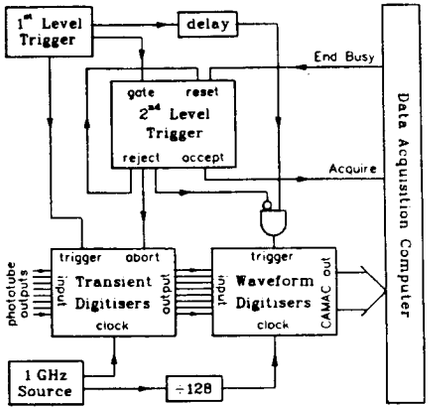

Fig. 5

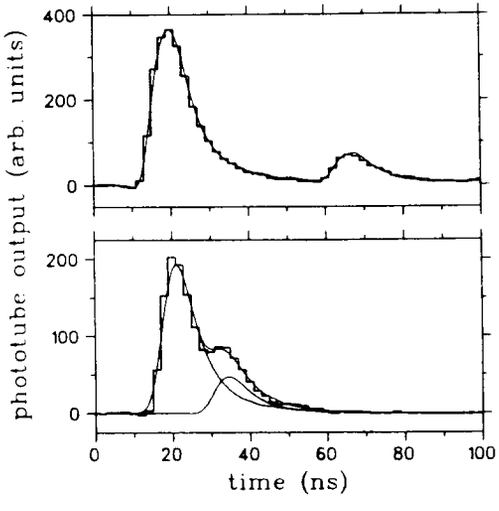

Fig. 7

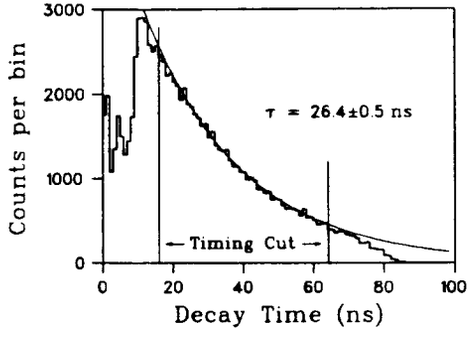

Fig. 9