ELSEVIER

Nuclear Instruments and Methods in Physics Research A 450 (2000) 391–398

NUCLEAR INSTRUMENTS & METHODS IN PHYSICS RESEARCH Section A

www.elsevier.nl/locate/nima

A quad 500 MHz waveform digitizer with differential trigger for use in the muon g-2 experiment

S. Dhawan[a], V.W. Hughes[a], D. Kawall[a], W. Liu[a,1], J. Pretz[a,*], R. Sumner[b]

[a]*Physics Department, Yale University, New Haven, CT 06511-8121, USA*
[b]*Cheesecote Mtn. Consultants, 24 Halley Drive, Pomona, NY 10970, USA*

Received 21 October 1999; received in revised form 5 January 2000; accepted 10 January 2000

Abstract

We have developed and tested a 4 channel, 500 MHz, 512 k sample, 8 bit waveform digitizer in a FASTCAMAC module. A novel feature is the implementation of a sparse scan during readout with a programmable differential trigger to reduce the data to a manageable level. The differential trigger allows reliable detection of pulses with low threshold in the presence of a variable baseline. © 2000 Elsevier Science B.V. All rights reserved.

Keywords: Waveform digitizer; Differential trigger

1. Introduction: need for a new waveform digitizer

The available waveform digitizers (WFD) on the market, either custom made [1,2] or commercial,[2,3] lack a differential trigger, i.e. the possibility to trigger on the time derivative of a signal instead of on the signal itself. This feature is important to detect small pulses on a varying baseline as explained in Section 2. Our WFD module was developed for the muon g-2 experiment E 821 [3][4] at Brookhaven National Laboratory (BNL) which is described in Section 3. In this section the requirements for the WFD determined by the experimental conditions are also discussed. Section 4 describes the module and the readout procedure. Tests of the functionality and the time stability are discussed in Section 5. In Section 6 data taken during the 1999 run period are presented.

2. Differential trigger

Instead of triggering on a input signal when its amplitude crosses a certain threshold, the differential trigger is obtained when the time derivative of the signal crosses a threshold. Fig. 1 illustrates this: two small negative pulses are superimposed on a large time-dependent baseline. The dots represent the sampled ADC values taken every 2 ns. Using a constant threshold as a trigger, the trigger occurs when the signal is below the dashed line in Fig. 1(a).

* Corresponding author. Tel.: + 1-203-432-3822; fax: + 1-203-432-3824.

E-mail address: pretz@sun2.bnl.gov (J. Pretz).

[1] Now at MICROSOFT, 1 Microsoft Way, Redmond, WA 98052, USA.

[2] STR7515, Bastian Technology GmbH & Co. KG, Backbarg 6, D-22889 Tangstedt, Germany.

[3] LRS2262, LeCroy Corp. 700 Chestnut Ridge Road, Chestnut Ridge, NY10877, USA.

[4] Muon g-2 Proposal BNL-AGS E821; *A new precision measurement of the muon g-2 value at the level of* 0.35 *ppm.*, 1995, Brookhaven National Laboratory., Muon g-2 Collaboration.

0168-9002/00/$ - see front matter © 2000 Elsevier Science B.V. All rights reserved.
PII: S0168-9002(00)00263-1

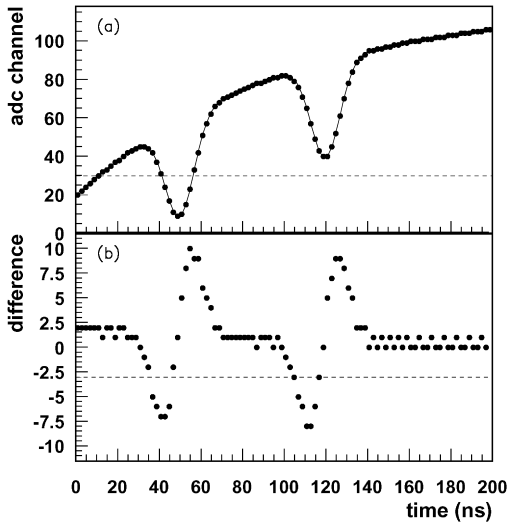

Fig. 1. Two small negative pulses are superimposed on a large time-dependent baseline. The dots in Fig. (a) indicate the ADC samples taken every 2 ns. In Fig. (b) the differences between two adjacent ADC samples are shown. The dashed lines indicate the trigger threshold for a constant trigger level (a) and a differential trigger (b). In (a) only the first pulse and the background at early times are read out, whereas in (b), using the differential trigger the two pulses are read out.

Only the first of the two pulses causes a trigger, the second pulse stays above the threshold, because the baseline has changed. In addition, a low threshold leads to a permanent trigger at early times due to the baseline, as can be seen in Fig. 1(a), which results in a high data rate written to tape. Fig. 1(b) shows the differential signal, i.e. the difference of two consecutive ADC samples. The dashed line indicates the differential threshold. Here only the two pulses cause triggers.

The differential trigger of the WFD is of general utility whenever the signal to be digitized is coincident with a time-dependent background and this occurs generally in any experiment using a pulsed beam. It is also extremely useful when the incoming signal is AC coupled.

3. The g-2 experiment: requirements for the WFD

The experimental conditions of the muon g-2 experiment E821 determine the requirements for the development of this new WFD.

The muon g-2 experiment measures the anomalous magnetic moment of the muon by observing the spin precession in a storage ring through detection of the arrival time and energy of the positrons from the decay $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$. After injection of 3.1 GeV muons into the storage ring, positrons are detected in calorimeters for 600 μs, which corresponds to approximately 10 muon lifetimes in the laboratory system. Over the 600 μs interval the positron rate varies from a few MHz to a few Hz. The average arrival time of the positrons must be determined with an accuracy of better than 60 ps over the 600 μs to achieve a sufficiently low systematic error. This necessitates good rejection of pile-up pulses, which requires a high sample rate of approximately 500 MHz.

Because of the decaying background, the pedestal varies from early to late times. A differential trigger is required to detect low-energy positrons with the same efficiency over the entire 600 μs period.

Finally, each calorimeter consists of 4 blocks of lead/scintillating fibers each read out separately by a PMT. Hence the WFD was designed with 4 input channels.

To digitize the signal at 500 MHz, 300 k samples have to be stored. The time between two fills of the storage ring magnet is 33 ms, during which time the 300 k samples must be analyzed to identify positron pulses. Six fills are obtained during one accelerator cycle. The time between two cycles is 2.5 s, which is the maximum time available to read out the data from the WFD module.

At present, the g-2 experiment uses a 400 MHz WFD, developed at Boston University [3], which fulfills all the requirements except for the differential trigger.

4. Hardware and firmware description

The CAMAC module is designed for high performance and maximum flexibility. It contains 4 independent channels, a multi-event buffer memory, a 500 MHz clock and the clock control, and provides for FASTCAMAC and FERA (Fast Encoding and Readout ADC) readout. The circuit diagram is shown in Fig 2. Fig. 3 shows a photograph of the 4 channel digitizer module. It is

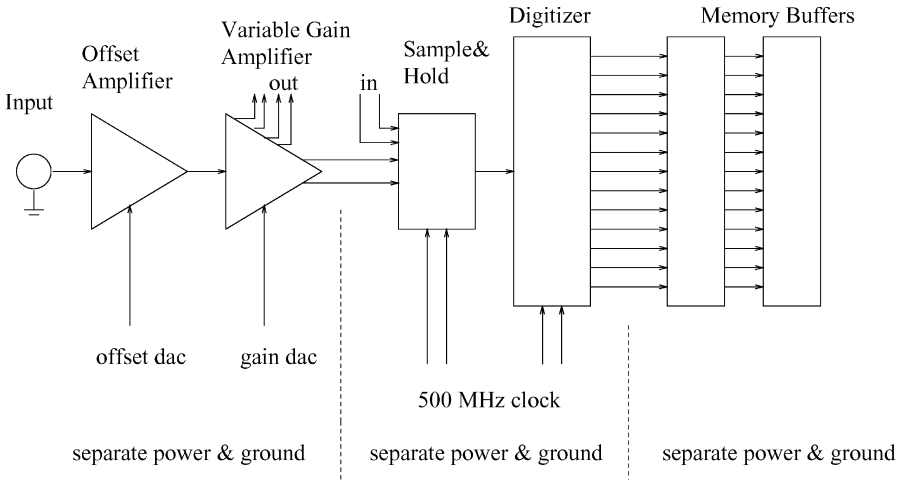

Fig. 2. The multi-chip module (one for each of the four channels) consists of a front-end amplifier (with offset and gain control), a digitizer (500 MHz sample/hold and ADC), and a circular memory buffer (2 chips, 256 k samples each). The three sections have separate power and ground (isolated on the multi-chip module) and the signal is fully differential between sections. This eliminates for example any possible coupling of correlated noise into the signal (address noise from the memory chips).

a double width CAMAC module. The module consumes about 30 W (each channel requires about 6 W) and supplementary ± 6 V power.

Each channel consists of a coaxial input (DC-coupled and terminated in 50 Ω), a programmable offset amplifier, a variable gain amplifier with a bandwidth limiting filter, a sample and hold, an 8 bit, 500 mega-sample per second digitizer, and a 512 k sample circular buffer memory. The input and offset amplifier are made with commercial surface-mounted components. The variable gain amplifier, sample and hold and digitizer are all high-speed custom bipolar chips. The memory consists of two custom CMOS chips (256 k samples each), operating with a 250 MHz clock (this chip stores data on both edges of the clock). The five custom monolithic chips in each channel were provided by LeCroy Corporation, and the detailed specifications are proprietary. The signal bandwidth is greater than 200 MHz, and the full scale input signal range is continuously variable from ± 25 to ± 250 mV. The least significant bit (LSB) ranges from 200 μV to 2 mV. The input offset can be adjusted over a range of ± 250 mV.

Each channel is constructed as an 80 pin plug-in multi-chip module, with the custom semiconductor dice wire bonded to a multi-layer substrate containing the remaining components for one channel. The input signal is a separate coaxial cable. The 500 MHz sampling clock is differential ECL, on terminated, controlled impedance lines. The remaining signals (the control signals and the readout data path) are relatively slow (<20 MHz), and are a mixture of ECL and CMOS. The board is divided into three sections: the input and amplifiers, the sample and hold and digitizer, and the memory. These three sections have completely separate power and ground. The signals between the sections are differential, which completely eliminates any feedback of memory noise to the front end. Indeed, the noise level (with no signal) varies from 0.5 to 1 LSB RMS, depending on the gain setting, and shows no correlation with memory address. There is no detectable cross-talk between the 4 channels in the CAMAC module.

Four of the digitizer channels are installed in a CAMAC board containing the control logic, data path, event memory and the 500 MHz clock. The control logic is provided by a large FPGA (Field Programmable Gate Array, Xilinx 4013). The data path from the digitizers goes through the FPGA to the event memory (512 K samples, 2 Mbytes). The

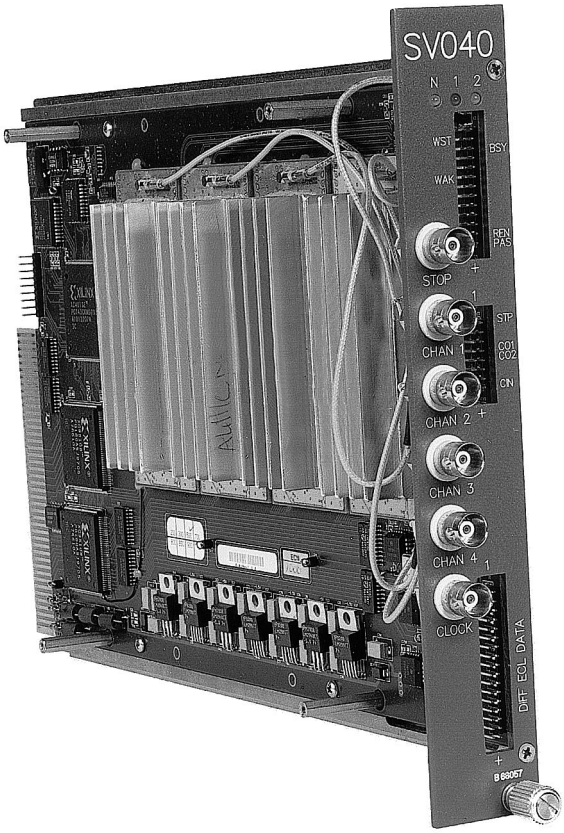

Fig. 3. A photograph of the 4 channel digitizer module. Each channel has its own heat sink mounted on the MCM. Analog signals, external clock and external stop signals are brought in via BNC connectors.

CAMAC interface is provided by three CPLDs (one Xilinx 9536 and two 95108 s). An interface to the LeCroy FERA bus is also provided. The control logic and interfaces are completely programmable. The FPGA is programmed automatically on power up from an EEPROM memory. The FPGA (and the EEPROM) can also be reprogrammed from CAMAC. The CAMAC interface is designed to be FASTCAMAC compatible. The board also allows interleaving the 4 channels in pairs, to create 2 channels, each with 1 GHz sampling rate.

The 4 channels operate in unison, with a common clock, each channel digitizing and storing the data for one of the 4 signal inputs. The digitizers are started and stopped by synchronously gating the clock. The clock can be the internal 500 MHz crystal oscillator, or an external source, either sine wave or differential ECL. The external clock can be any frequency from DC to 500 MHz. Differential ECL outputs of the gated clock are provided to allow synchronized operation of multiple digitizer modules. The clock and clock control are implemented with high-speed ECLIPS logic.

This programmability has allowed a custom program to be designed for the muon g-2 experiment. To allow maximum resources for this application, only a normal CAMAC readout is implemented. FASTCAMAC and FERA interfaces are disabled. The final program uses more than 75% of the resources in the FPGA (which limits the maximum processing speed).

During a muon g-2 decay cycle, the analog signals coming from the four segments of the calorimeter are digitized for about 600 μs (300 k samples), after which the clock is stopped and the digitizers are read out. During the readout of the digitizer memories, the data are examined by the logic programmed in the FPGA, and only those data satisfying the g-2 differential trigger criteria are stored in the CAMAC event buffer memory.

To obtain acceptable processing speed, the trigger algorithm is pipelined. Each stage described below is a registered pipeline stage. The readout logic and trigger pipeline are clocked at 15 MHz. The readout and processing of the 300 k samples take about 20 ms. The processing must be complete, and the digitizer reset before the next fill, 33 ms later.

Stage 1: Read 4 bytes (one sample) from the digitizer buffer memories and store in the first pipeline register. These are the data for the same time slot for the 4 channels. These 4 bytes also enter a 16 word deep by 32 bit wide data delay FIFO memory which is programmed inside the FPGA, and is not an external FIFO chip.

Stage 2: Add the 4 ADC values in two stages. First, sum in pairs to make two 9 bit words, and add these to make one 10 bit word which is stored in the second pipeline register.

Stage 3: Add the differential threshold value (programmable) to the sum in pipeline register 2, thus creating the current trigger threshold, and store this in the pipeline register 3.

Stage 4: Compare the current threshold from pipeline register 3 to the sum in register 2, which is now the sum for the next (later) time slot. Store the result of the comparison in pipeline register 4. If the sum is greater than the current threshold, i.e. the sum has increased by more than the differential threshold, the result is a trigger.

Stage 5: The trigger causes the header to be written to the event buffer memory. The header contains the time of the trigger (number of samples) relative to the beginning of the readout.

Stage 6: A sufficient number of pipeline stages are inserted to adjust the number of pre-trigger data samples as desired. The trigger then loads the event size (programmable, the sum of the pre-trigger and post trigger samples) into a counter. When the counter is greater than zero, the output of the data delay FIFO (delayed by 16 clock cycles) is written to the event buffer memory on each clock cycle. Each time a word is written, the counter is decremented. When the counter has decremented to zero, the record is complete, and the writing stops.

The trigger algorithm continues until all 300 k samples have been examined. If a new trigger is found during the current trigger, the counter is reloaded with the event size, and the record is simply extended to include the new trigger. There is no dead-time associated with this trigger and no potential events are missed.

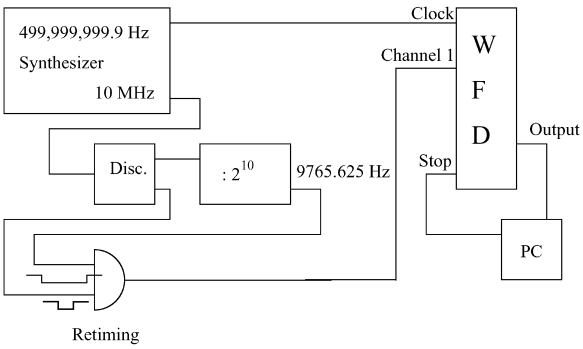

Fig. 4. The test setup: A 499,999,999.9 Hz synthesizer signal is used as a clock signal. A 10 MHz signal from the the same synthesizer is divided by 2^{10} which results in a 9765.625 Hz signal. This signal is re-timed with the original 499,999,999.9 Hz signal (logical AND in lower left corner). The output signal of the AND unit which was about 20 ns wide was then put in one of the input channels of the WFD. The WFD was read out and stopped by a PC.

An event size of 50 words with 7 pre-trigger samples is used for the muon g-2 readout. The trigger threshold (the delta used in pipeline register 3) is programmed as 7. The 4 channel event record for one trigger event is shown in Fig. 7. The sum of the 4 channels (the total energy in the calorimeter) is given in Fig. 8(a), and the derivative is shown in Fig. 8(b). Since there are typically only 100 triggers during the 600 μs decay period, the total event data is only 5000 words (20 k bytes), a data reduction of 60 : 1 compared to reading and storing the complete digitized data sample (1.2 Mbytes). The module event memory is able to store the data from 6 fills of the storage ring. The data are read out over CAMAC during the 2.5 s interval between accelerator cycles.

The differential trigger algorithm implemented here was designed for the particular experimental

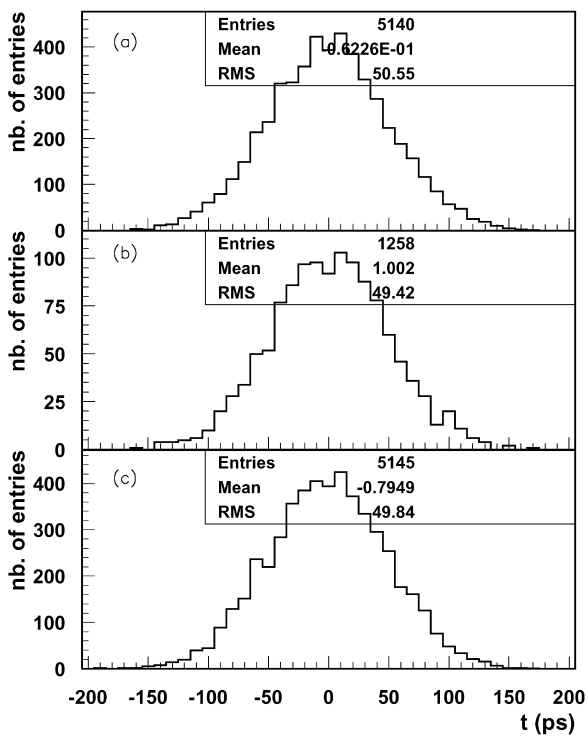

Fig. 5. The distribution of time differences between two 20 ns pulses (minus an offset of 1 024 000 ns) for three different cases: both pulses are stored on the first chip (a), the first pulse is on the first chip the second pulse is on the second chip (b), and both pulses are stored on the second chip (c). The mean values are consistent with 0 at a ps level.

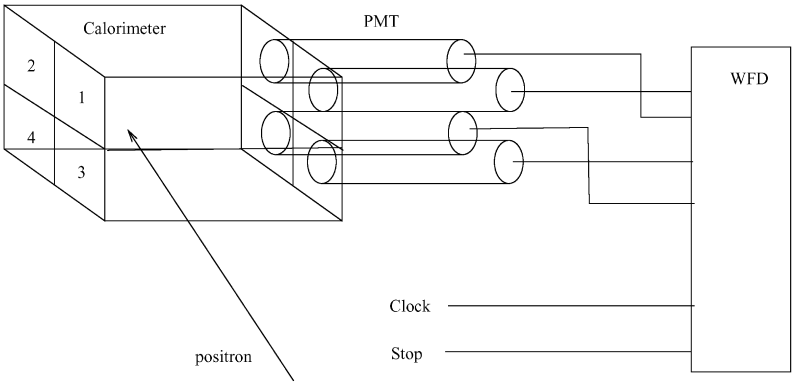

Fig. 6. In the muon g-2 experiment the calorimeters are segmented. Each of the four segments of a detector was read out by a PMT. The analog signals of the four PMT tubes are the input signals to the WFD. The clock signal comes from a 500 MHz frequency generator.[5] The stop signal which triggered the readout of the module was sent 600 μs after the injection of the beam into the storage ring.

situation of the g-2 experiment. The programmability of the module allows other algorithms to be designed for other applications. The module can be reprogrammed over CAMAC in a few seconds, or a second program can be stored in the EEPROM program memory, allowing a complete change in the function of the module in about 300 ms. For example, multilevel triggering, or coincidences using the three channels, resulting in a (possibly delayed) acquisition on the fourth channel. The possibilities are limited only by the users imagination (and the size of the FPGA).

5. Tests of the module

Fig. 4 shows the setup used to test the functionality and the time stability of the WFD. A 499,999,999.9 Hz synthesizer signal is used as a clock signal. A 10 MHz signal from the same synthesizer is divided by 2^{10} which results in a 9765.625 Hz signal. This signal is re-timed with the original 499,999,999.9 Hz signal (logical AND in lower left corner of Fig. 4). The output signal of the AND unit which was about 20 ns wide was then put in one of the input channels of the WFD. The WFD was read out and stopped by a PC.

[5] PTS (Program Test Sources Inc.), 9 Beaver Brook Road, P.O.Box 517, Littleton, MA 01460, USA.

The time difference of two consecutive 20 ns wide pulses was determined. Fig. 5 shows this distribution. The time of the pulse was determined from the center of gravity: $t_{\mathrm{pulse}} = \sum y_i t_i / \sum y_i$ where y_i is the pedestal subtracted ADC sample at time t_i. The RMS of 50 ps is mainly due to the algorithm of finding the pulse time, t_{pulse}, which was kept intentionally simple here.

As the data on the WFD are internally stored on two different chips ($0 < t < 500$ μs and 500 μs $< t < 1000$ μs, the two memory buffers in Fig. 2), the three plots show the three different cases where:

- both pulses are stored on the first chip,
- the first pulse is on the first and the second pulse, is on the second chip,
- both pulses are stored on the second chip.

The mean values of the three distributions are consistent with zero on a picosecond level. The tests were repeated at different times of the day, i.e. at different temperature conditions. No dependence was found.

6. Results from the muon g-2 data taking

This WFD was used during the 1999 data-taking of the muon g-2 experiment in one of the 24 detectors in parallel with the 400 MHz WFD mentioned in Section 3. Each of the four PMT signals

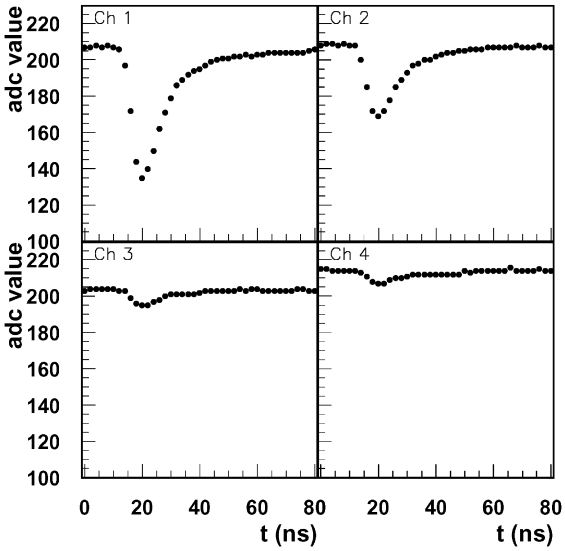

Fig. 7. Digitized analog signals of the four PMT for a positron entering the upper part of the detector.

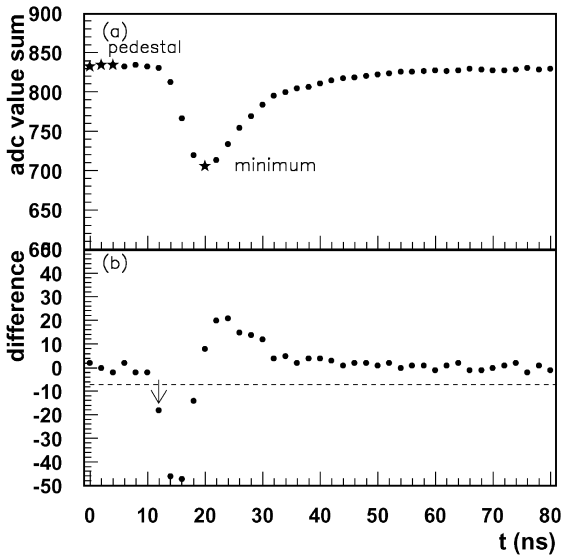

Fig. 8. The upper plot (a) shows the sum of the four signals in Fig. 7. In the lower plot (b) the differential signal, i.e. the difference of two consecutive channels from plot (a) is shown. The arrow points to the channel which caused the differential trigger indicated by the dashed line.

was digitized by a separate WFD channel at a 500 MHz rate (see Fig. 6). Fig. 7 shows a signal in the four channels of the calorimeter. The positron enters the upper part of the calorimeter where most of its energy is deposited in the front segment. Only a small amount of energy is deposited in the lower part. Fig. 8(a) shows the sum of the four signals in Fig. 7. In Fig. 8(b) the differential signal, i.e. the difference of two consecutive time bins is shown. The arrow points to the channel which caused the differential trigger indicated by the dashed line. It was set here to be greater than 7 ADC channels in the summed signal.

Fig. 9 shows an energy spectrum of the decay positrons given by the amplitude of the signal. The amplitude was calculated from the pedestal which was taken to be the average of the first three samples in the signal as indicated by the stars in Fig. 8(a) minus the minimal ADC value. Channel 200 in Fig. 9 corresponds to approximately 3 GeV. The lowest entries are around the threshold. The trigger becomes 100% efficient at approximately channel no. 20 (i.e. 300 MeV). It can be seen from Fig. 8(a) that the rise time of the signal is approximately two time bins or 4 ns. This means that a signal with an amplitude of 20 ADC channels which corresponds to roughly 2.5 times the threshold will cause a trigger whenever it arrives within a 2 ns bin. Fig. 10 shows the pedestal of the summed signal versus time. It varies by about 30 ADC

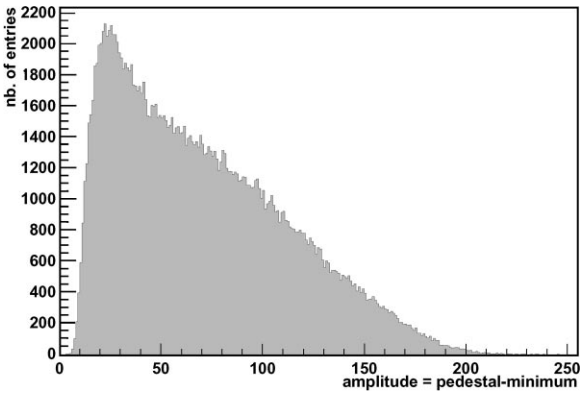

Fig. 9. The energy spectrum from the sum of all four channels. The energy was calculated from the pedestal which was taken to be the average of the first three samples in the signal as indicated by the stars in Fig. 8(a) minus the minimal ADC value (also indicated by a ★). An ADC value of 200 corresponds to approximately 3 GeV. The trigger becomes 100% efficient at approximately channel no. 20 (i.e. 300 MeV).

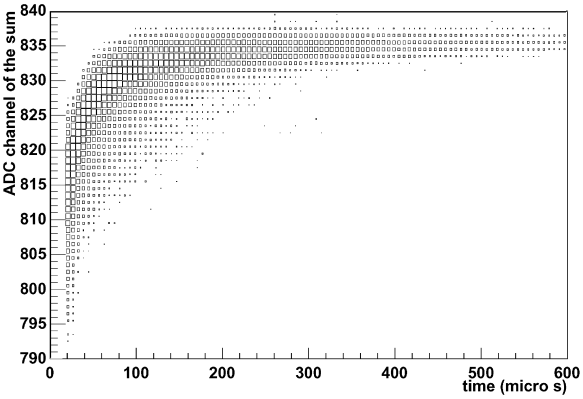

Fig. 10. The pedestal of the summed signal versus time. It varies by about 30 ADC channels from early to late times, emphasizing the need for a differential trigger to detect low-energy pulses.

channels from early to late times, emphasizing the need for a differential trigger to detect low-energy pulses.

7. Conclusions

We have successfully developed and tested a new 4 channel, 500 MHz, 512 k sample, 8 bit waveform digitizer using a differential trigger and it has been used in the muon g-2 experiment at BNL. The module is especially useful when signals are superposed on a time dependent baseline.

Acknowledgements

This work was supported in part by the US Department of Energy.

References

[1] D.A. Bryman et al., IEEE Trans. Nucl. Sci. NS38 (1991) 295.
[2] M. Atiya et al., Nucl. Instr. and Meth. A 279 (1989) 180.
[3] R.M. Carey et al., Phys. Rev. Lett. 82 (1999) 1632.