Nuclear Inst. and Methods in Physics Research, A 1045 (2023) 167542

Contents lists available at ScienceDirect

Nuclear Inst. and Methods in Physics Research, A

journal homepage: www.elsevier.com/locate/nima

The WaveDAQ integrated Trigger and Data Acquisition System for the MEG II experiment

Marco Francesconi [a,b,*], Alessandro Baldini [b], Hicham Benmansour [a,b,c], Fabrizio Cei [a,b], Marco Chiappini [b], Gianluigi Chiarello [b], Luca Galli [b], Marco Grassi [b], Ueli Hartmann [c], Fabio Morsani [b], Donato Nicoló [a,b], Angela Papa [a,b,c], Stefan Ritt [c], Elmar Schmid [c], Giovanni Signorelli [b], Bastiano Vitali [b,d]

[a] Dipartimento di Fisica dell'Università di Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy
[b] Istituto Nazionale di Fisica Nucleare, Sezione Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy
[c] Paul Scherrer Institut, 5232 Villigen, Switzerland
[d] Dipartimento di Fisica dell'Università "Sapienza" di Roma, Piazzale A. Moro 2, 00185 Rome, Italy

ARTICLE INFO

Keywords:
Trigger
Data acquisition
MEG II
Lepton flavor violation

ABSTRACT

The complete MEG II Trigger and Data Acquisition System, named WaveDAQ, was installed and commissioned in Spring 2021 and successfully carried out the data taking campaign planned for the same year (Chiappini et al., 2021 [1]). It consists of 544 custom made 16-channel acquisition boards which contain the Domino Ring Sampler 4 chips for the analog sampling of the detector signals at > 1 GHz frequency. This paper presents how the MEG II physics objectives shaped the design and operation of the WaveDAQ, the result is a flexible and scalable trigger and data acquisition system.

1. The MEG II experiment

The MEG II experiment [2] at Paul Scherrer Institut (Switzerland) is the direct continuation of the MEG Experiment [3] and aims at improving the upper limit on the $\mu^+ \rightarrow e^+\gamma$ branching ratio of 4.2×10^{-13} (90% C.L.) [4].

The goal of increasing the sensitivity by an order of magnitude is accomplished by a deep redesign or complete replacement of the sub-detectors, so to obtain better resolutions on all observables by a factor 2. At the same time, the increased segmentation of the detector helps in maintaining the signal pile-up under control while operating the experiment at higher muon stopping rate on target with respect to MEG.

In particular in the Liquid Xenon detector [5], which measures the photon energy deposit, impact point and timing, the photomultipliers on the inner face were replaced by 4092 VUV-sensitive Silicon Photo-Multipliers (SiPM) by Hamamatsu Photonics to increase the light collection granularity and the photocathodic coverage. At the same time, the remaining faces are equipped with 668 photo-multipliers tubes whose positions were optimized to improve light collection. The positron trajectory is reconstructed by a newly designed single volume drift chamber [6] formed by 1728 drift cells read at both ends of the active volume. At the end of its path in the drift chamber, the positron hit some of the 512 plastic scintillator tiles which constitute the pixelated Timing Counter; each counter is equipped with two sets of SiPMs on opposite sides [7].

An excellent timing resolution is needed to suppress the accidental coincidence background and is a crucial request for the implementation of pileup rejection algorithms. In addition, Liquid Xenon detector channels must also provide charge extraction capabilities, in the same harsh pileup environment.

Finally, a huge fraction of the channels is connected to SiPMs and therefore need to be able to supply the bias voltage and provide the additional analog gain to match the available dynamic range.

2. Trigger and data acquisition requirements

The data acquisition requirements of the MEG II experiment were already largely addressed during the former MEG experiment, this led to the development and deployment of the Domino Ring Sampler 4 (DRS4) chip [8], an array of 1024 switched capacitor cells that provide full-waveform digitization above the 1 GSPS speed.

The storage of full analog waveforms make possible to analyze data offline optimizing feature extraction.

Scaling the previous VME-based system to the MEG II experiment was not sustainable, mainly because of the fourfold increase of the readout channels due to more segmented, detectors. From another

* Corresponding author at: Dipartimento di Fisica dell'Università di Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy.
E-mail address: marco.francesconi@pi.infn.it (M. Francesconi).

https://doi.org/10.1016/j.nima.2022.167542
Received 11 April 2022; Accepted 30 September 2022
Available online 10 October 2022

Fig. 1. WaveDREAM board in standalone configuration, used during detector commissioning and signal checks.

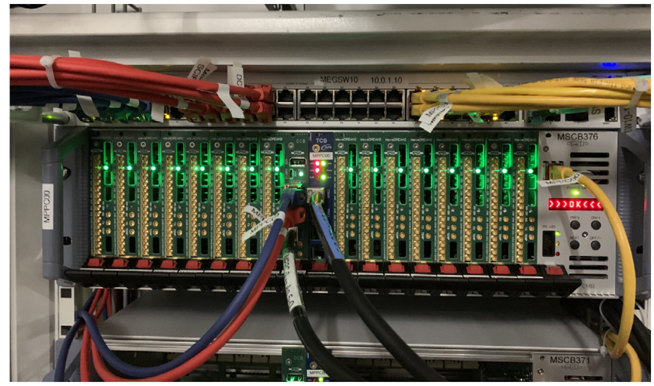

Fig. 2. WaveDAQ crate, from left to right: 8 WaveDREAM boards, the DCB boards, the TCB board, other 8 WaveDREAM boards, crate controller with power supply and fan circuitry.

point of view, the former data acquisition scheme had a long readout time required by the explicit interrogation of each card by the bus master.

These issues are solved by using the same DRS4 chip, which proved very successful in MEG, but in a different data acquisition system that fits in the same rack space, while also providing the additional features needed to operate SiPMs. The new system, called WaveDAQ, makes extensive use of Ethernet for the detectors readout, therefore directly pushing data from the hardware as soon as an event of interest is detected.

The main drawback of the DRS4 chip is a pretty long $\sim$ 620 μs dead time required to convert the charge stored in the capacitor bank through a slower speed external analog to digital converter. This limitation is well known and, in the former experiment, was limited by operating stringent online event selection in Field Programmable gate arrays (FPGA) [9].

Only a fraction of muon decays is indeed compatible with the two-body process the experiment is looking for; such candidates are selected by requiring the time and spatial coincidence of two child particles as well as a matching energy deposit by the photon [10]. In the WaveDAQ, the former trigger system is completely integrated with the data acquisition part sharing the same analog to digital converter used for the DRS4 readout. Newer FPGA technologies have been included in the trigger computation so that the information from a larger amount of channels can be combined leaving the latency untouched [11].

The trigger latency indeed is an important parameter in a DRS4-based system: the stop signal must reach all chips before the depth (1024 samples) of the analog memory in the chip is exceeded, otherwise the initial part of the waveform which generated the trigger is overwritten and therefore is lost.

From another point of view, a reduced trigger latency allows for faster sampling of the input signals by the DRS4 and therefore increases intrinsic time resolution.

3. The WaveDAQ system

The WaveDAQ system [12] is designed to be scalable from simple detector test setups to a complex system such as the MEG II one. This peculiarity proved very useful during the experiment commissioning when the DAQ had to be quickly adapted to perform various checks when needed.

The possibility of scaling the system is provided by the WaveDREAM (Waveform Drs4 REAdout Module), a custom self-contained DAQ board that contains two DRS4 chips for a total of 16 input channels.

The board comes in two types, having either a coaxial or a differential input connector. The first one is intended to be directly connected to the signal source and can provide up to 240 V to directly bias the SiPM without any external circuitry. The differential version, on the other hand, is designed to be used with detectors having front-end circuitry, such as the MEG II drift chamber.

The input stage of the WaveDREAM can be configured with analog gain from 0.25 to 100 to ease the matching with the 1 V DRS4 dynamic range. A pole-zero cancellation circuit and a tunable discriminator for triggering are also available on each channel.

The WaveDREAM can operate in standalone mode (Fig. 1) using the on-board ethernet socket or can be inserted in a 16 slot crate (Fig. 2) for larger setups, up to 256 channels. In this configuration two additional boards are needed: the DCB (Data Concentrator Board) provides a common Ethernet connection for the configuration of the whole crate and sequences the transmission of packets originating from the WaveDREAMs. The second Board is the TCB (Trigger Concentrator Board) which receives up to 5.12 Gbps worth of information from each WaveDREAM so that a trigger decision can be based on any combination of amplitudes and times from every input channel.

For larger systems, the TCB forwards the trigger information to a central trigger system also made of TCBs. When an event of interest is identified by this central system, the stop signal is sent to all DRS4 chips and a trigger number is distributed to all boards to tag the piece of information associated with that particular event.

4. The MEG II trigger and data acquisition system

The full MEG II Trigger and Data Acquisition system consists of 524 WaveDREAM boards distributed in 34 crates and is expected to operate at a trigger rate of $\sim$ 15 Hz.

The readout Ethernet infrastructure was also deployed at the same time. It is constituted by five *"Top-Of-Rack"* switches, which aggregates the crates in the same rack, and a second aggregation switch that connects to the DAQ computed utilizing 10 Gbps interfaces. For both roles, we selected Mikrotik CRS series switches.

We opted to consolidate as much as possible into a single compute node which proved, as described later, capable enough to process the full DAQ throughput expected for MEG II.

5. Event building and data reduction

A single-thread program cannot digest the throughput generated by the WaveDAQ, therefore a multi-thread approach was adopted using concurrency primitives provided by modern C++11.

At network throughputs above 1 Gbps, a non-negligible part of the current CPU resources is currently taken by retrieving each network packet from the Linux Kernel. Such congestion was easily solved by spreading the packets across multiple sockets, by distributing the load across multiple threads and by using system calls which allows retrieving multiple packets at once.

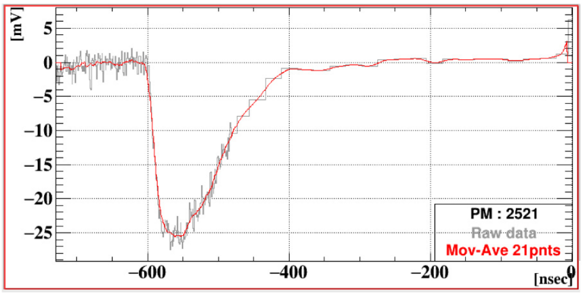

Fig. 3. Example of a waveform rebinning outside the signal time region. Depending on peak amplitude (in this case ∼ −25 mV), a rebinning value of 8 was chosen to be applied after the peak (in the example, from ∼ −480 ns onward).

The next step is the event building according to the trigger event number which was distributed together with the stop signal. This stage accounts for packet order scrambling and potential packet losses, which can happen during the transmission.

Then each waveform is calibrated to compensate for the variability in the characteristics of different DRS4 chips. This step is strictly needed because it is followed by a data reduction stage that must operate on well-calibrated signals so to effectively reduce the event size to be written on disk.

We currently have three data reduction strategies. The first one allows a selection of only a specific time window in the waveform to be recorded. The second can apply a threshold in the peak-to-peak amplitude to discard waveforms without any associated signal. Finally, the third algorithm averages up to 16 consecutive analog samples (what we call "rebin") outside the region selected by the trigger. In this way, the overall charge associated with the pulse is fully preserved and the timing of the pulse can be extracted with the same resolution of the original waveform.

An example of the effect of the rebinning algorithm is clearly shown in Fig. 3 for a SiPM channel in the Liquid Xenon Detector inner face.

6. System installation and commissioning

The full MEG II Trigger and Data Acquisition system was completely installed in March 2021 [1] and is depicted in Fig. 4.

The MEG II experiment takes full advantage of the WaveDAQ trigger capabilities: the muon beam with an intensity of above 10^7 μ^+/s is stopped in the target at the experiment center producing a huge number of uninteresting decay products. A $\mu^+ \rightarrow e^+\gamma$ candidate is singled out from this crowded environment through the following steps:

- The photon energy is reconstructed by the weighted sum of all photosensors in the Liquid Xenon detector. The individual photosensor weight takes into account the variability in sensor gains and detection efficiency.
- The pulse time is extracted by the analog discriminator of the WaveDREAM whose output is oversampled at 640 MHz so to obtain a 1.56 ns time bin on the single channel.
- The impact point position of the positron in the pixelated Timing Counter is identified and its location and time are compared with the ones obtained from the reconstructed shower position in the Liquid Xenon detector inner face.

The complete MEG II Trigger was first used in 2021. It was operated with loosened cuts, and therefore at an increased trigger rate, to assure selection efficiency since the detector response was not yet fully calibrated. Further improvements on the trigger side are expected for next year. The MEG II Trigger system is also used to select events from detector calibration sources, such as the α particle emitters submerged in liquid xenon [13] or the laser system used by the pixelated Timing Counter [14].

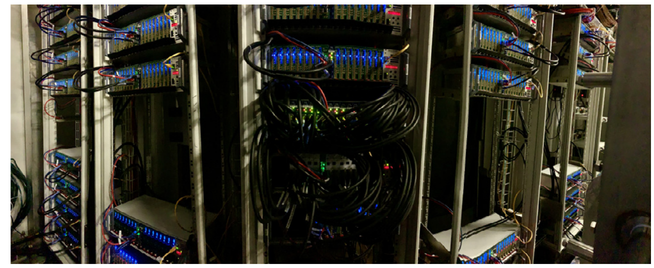

Fig. 4. MEG II WaveDAQ system, central crates with black cables contains the central trigger system of the experiment and the corresponding fan-out logic.

On the DAQ side, it was quickly realized that a reduction of the number of generated Ethernet packets, by increasing the maximum allowed size (also known as Maximum Transfer Unit), would have been very helpful in saturating the network capabilities. Unfortunately, the current implementation of DCB firmware cannot cope with MTU size above the standard 1528 bytes, and, given the size of a DRS4 waveform, two packets have to be sent for each channel.

The maximum aggregated throughput that the online machine can handle is ∼ 8 Gbps with more than 1.5 Million packets per second being collected.

The MEG II experiment raw event size, after waveform calibration, is ∼ 20 MB and was successfully reduced by a factor ∼ 4 with the previously mentioned algorithms. In this way, the data writing on disk, at MEG II trigger rate of 15 Hz, stays below 130 MB/s which is the current limit caused by the data compression software we are using.

7. Conclusions

The MEG II experiment started its physics data collection in 2021 with a ∼ 1 month data acquisition period which provided the opportunity to develop and gain experience with the trigger and the data reduction algorithms described in this paper.

The WaveDAQ system scaled as expected from a few crates to the final system and proved to be stable over this first data taking campaign.

The final experimental sensitivity of $6{\times}10^{-14}$ will require three years of data acquisition [15].

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

The authors are thankful to the former collaborator Ryu Sawada who developed the data reduction code in MEG I.

References

[1] M. Chiappini, M. Francesconi, S. Kobayashi, M. Meucci, R. Onda, P. Schwendimann, Towards a new $\mu \rightarrow e\gamma$ search with the MEG II Experiment: From design to commissioning, Universe 7 (12) (2021) 466.
[2] A.M. Baldini, E. Baracchini, C. Bemporad, F. Berg, M. Biasotti, G. Boca, P.W. Cattaneo, G. Cavoto, F. Cei, M. Chiappini, et al., The design of the MEG II experiment, Eur. Phys. J. C 78 (5) (2018) 1–60.
[3] J. Adam, X. Bai, A.M. Baldini, E. Baracchini, C. Bemporad, G. Boca, P.W. Cattaneo, G. Cavoto, F. Cei, C. Cerri, et al., The MEG detector for $\mu^+ \rightarrow e^+\gamma$ decay search, Eur. Phys. J. C 73 (4) (2013) 1–59.

[4] A.M. Baldini, Y. Bao, E. Baracchini, C. Bemporad, F. Berg, M. Biasotti, G. Boca, M. Cascella, P.W. Cattaneo, G. Cavoto, et al., Search for the lepton flavour violating decay $\mu^+ \to e^+\gamma$ with the full dataset of the MEG experiment, Eur. Phys. J. C 76 (8) (2016) 434.

[5] S. Ogawa, MEG II Collaboration, et al., Liquid xenon calorimeter for MEG II experiment with VUV-sensitive MPPCs, Nucl. Instrum. Methods Phys. Res. A 845 (2017) 528–532.

[6] A.M. Baldini, G. Cavoto, F. Cei, M. Chiappini, G. Chiarello, A. Corvaglia, M. Francesconi, L. Galli, F. Grancagnolo, M. Grassi, et al., The ultra light drift chamber of the MEG II experiment, Nucl. Instrum. Methods Phys. Res. A 958 (2020) 162152.

[7] M. De Gerone, A. Bevilacqua, M. Biasotti, G. Boca, P.W. Cattaneo, F. Gatti, M. Nishimura, W. Ootani, G. Pizzigoni, M. Rossella, et al., A high resolution Timing Counter for the MEG II experiment, Nucl. Instrum. Methods Phys. Res. A 824 (2016) 92–95.

[8] S. Ritt, Design and performance of the 6 GHz waveform digitizing chip DRS4, in: 2008 IEEE Nuclear Science Symposium Conference Record, IEEE, 2008 pp. 1512–1515.

[9] L. Galli, F. Cei, S. Galeotti, C. Magazzù, F. Morsani, D. Nicolo, G. Signorelli, M. Grassi, An FPGA-based trigger system for the search of $\mu \to e\gamma$ decay in the MEG experiment, J. Instrum. 8 (01) (2013) P01008.

[10] L. Galli, A. Baldini, P. Cattaneo, F. Cei, M. De Gerone, S. Dussoni, F. Gatti, M. Grassi, F. Morsani, D. Nicolo, et al., Operation and performance of the trigger system of the MEG experiment, J. Instrum. 9 (04) (2014) P04022.

[11] M. Francesconi, A. Baldini, F. Cei, M. Chiappini, L. Galli, M. Grassi, U. Hartmann, F. Morsani, D. Nicolò, A. Papa, et al., Low latency serial communication for MEG II trigger system, Nucl. Instrum. Methods Phys. Res. A 936 (2019) 331–332.

[12] L. Galli, A.M. Baldini, F. Cei, M. Chiappini, M. Francesconi, M. Grassi, U. Hartmann, M. Meucci, F. Morsani, D. Nicolò, et al., WaveDAQ: An highly integrated trigger and data acquisition system, Nucl. Instrum. Methods Phys. Res. A 936 (2019) 399–400.

[13] D. Nicolò, A. Baldini, C. Bemporad, F. Cei, M. Chiappini, M. Francesconi, L. Galli, M. Grassi, T. Iwamoto, F. Morsani, et al., Real-time particle identification in liquid xenon, IEEE Trans. Nucl. Sci. 68 (11) (2021) 2630–2636.

[14] G. Boca, P.W. Cattaneo, M. De Gerone, M. Francesconi, L. Galli, F. Gatti, J. Koga, M. Nakao, M. Nishimura, W. Ootani, et al., The laser-based time calibration system for the MEG II pixelated timing counter, Nucl. Instrum. Methods Phys. Res. A 947 (2019) 162672.

[15] A.M. Baldini, V. Baranov, M. Biasotti, G. Boca, P.W. Cattaneo, G. Cavoto, F. Cei, M. Chiappini, G. Chiarello, A. Corvaglia, et al., The search for $\mu \to e\gamma$ with 10^{-14} sensitivity: The upgrade of the MEG experiment, Symmetry 13 (9) (2021) 1591.