ELSEVIER

Nuclear Instruments and Methods in Physics Research A 494 (2002) 520–525

NUCLEAR INSTRUMENTS & METHODS IN PHYSICS RESEARCH Section A

www.elsevier.com/locate/nima

Electronics for the $\mu \to e\gamma$ experiment

S. Ritt*

Paul Scherrer Institute, OLGA/021, CH-5232, Villigen PSI, Switzerland

For the MUEGAMMA Collaboration

Abstract

A new experiment at PSI searches for the forbidden decay $\mu \to e\gamma$ with a sensitivity down to 10^{-14}. This challenging experiment requires a series of new electronics in the trigger, the data acquisition and the slow control part. The described electronics not only satisfies the high demands in the sense of rate and resolution of the experiment, but is also significantly cheaper than conventional solutions.

PACS: 07.05.Hd; 07.05.Dz; 07.50.Ek; 84.70.+p

Keywords: MUEGAMMA; Trigger electronics; Fast waveform sampling; Slow control; High voltage supply

1. Introduction

A new experiment is planned at PSI to measure the lepton–flavour violating decay $\mu^+ \to e^+\gamma$ with a sensitivity down to 10^{-14} [1]. To achieve this sensitivity, a muon beam is stopped in a target at a rate of $\sim 10^8$ s^{-1}. The muon decay products are registered with drift chambers and a calorimeter consisting of 800 l of liquid xenon (LXe), the scintillation light of which is detected by 800 photomultipliers (PMTs) immersed in the xenon. The high particle rate and therefore the expected intense background makes it necessary both for the trigger electronics and the data acquisition (DAQ) electronics to show excellent pile-up rejection and timing resolution. Since the LXe calorimeter requires a large vacuum and pumping system, a complex slow control system is necessary to insure the long-term stability of the experiment. To fulfil these requirements, new approaches have been followed in electronics development which will be reported in the following sections.

The MUEGAMMA experiment is currently under construction and is expected to produce data starting from 2004.

2. Trigger electronics

A $\mu \to e\gamma$ event produces a 52.8 MeV gamma detected by the LXe calorimeter and a back-to-back 52.8 MeV positron which is registered in a set of radial drift chambers and a positron counter contained in a solenoidal magnetic field of 1.4 T. A trigger being sensitive to these events has in a first stage to discriminate these events from normal Michel events by a energy sum threshold in the LXe calorimeter. In a second stage, a back-to-back

*Tel.: +41-56-310-3728; fax: +41-56-310-2199.
E-mail address: stefan.ritt@psi.ch (S. Ritt).

0168-9002/02/$ - see front matter
PII: S0168-9002(02)01543-7

angular coincidence and a timing coincidence between the γ and the positron has to be made with an intrinsic trigger timing resolution in the order of 10 ns.

In a conventional way, the LXe energy sum would be done with analog fan-in units. Since the calorimeter is unsegmented, nearly every PMT sees part of the light produced by a γ shower, which makes it necessary to add the signals of all 800 PMTs. If each channel has a common noise of only a couple of millivolts, this would already result in an unacceptable noise level in the total analog sum. Furthermore, it is difficult to calculate the centre-of-gravity of all PMTs in an analog circuit, which is necessary to evaluate the angular correlation between the γ and the positron.

To overcome the problem of an analog summing, a digital trigger electronics has been designed. It utilizes 10-bit 100 MHz flash ADCs [2] and high-density FPGAs [3]. A "type-1" VME board houses 16 analog channels and calculates the sum of all channels in four stages after proper subtraction of the baseline for each channel, which reduces the contribution of common noise to a minimum. In addition, the channel with the maximal ADC value is determined in every 10 ns cycle. Monte-Carlo calculations have shown that the PMT with the maximal ADC hit is a good enough estimation of the shower centre and much easier to calculate than a real centre-of-gravity. The output of this board is transferred via LVDS lines to a "type-2" VME board, which contains only digital inputs and calculates the global ADC sum and the maximum PMT for all connected type-1 boards. A second layer of type-2 boards then completes a hierarchy that is able to accomodate all 800 channels from the calorimeter.

A key feature of the FPGA program is the baseline subtraction (Fig. 1). A register containing the current baseline level is maintained for every channel and subtracted from all ADC values. The baseline is derived from a running average over four ADC samples and "protected" against values from a real hit by a threshold comparator. Only if ADC values are below this threshold, they are attributed to the baseline value. After the baseline subtraction, the ADC values are run through a 10 bit $\times$ 10 bit lookup-table, which corrects PMT nonlinearities and applies an optional software gain. Since the FPGA can be re-programmed through the VME bus, this table together with the complete logic can be changed at any time.

The trigger latency is determined by the cycle time of 10 ns times the number of hierarchy stages. The flash ADC by itself has a latency of five cycles, which adds up to 150 ns for ten stages. If one takes into account the inter-board communication overhead plus the final correlation of the γ signal with the positron signal, one gets a total trigger latency of about 350 ns, which has to be accommodated by the pipeline of the front-end electronics. According to Monte-Carlo calculations, the energy cut together with the $\gamma - e^+$ correlation cuts down the trigger rate well below 100 Hz, which makes a second level trigger unnecessary.

To demonstrate the feasibility of such a trigger, a prototype board with the major components running at 100 MHz has been designed, built and successfully tested at PSI.

3. DAQ electronics

Most experiments use ADCs and TDCs for signal digitization. This technology reaches its limit when it comes to pile-up situations where two different events are separated in time by less than 30–40 ns. To disentangle these events, waveform sampling above 100 MHz becomes mandatory for all calorimeter channels. This is especially the case for the MUEGAMMA experiment, where a large background of 511 keV γ's originating from positron annihilations is superimposed on each signal. If the waveform sampling can be done with high accuracy for the PMT charge and timing, the usage of traditional ADCs and TDCs can be avoided and therefore the costs for the experiment can be reduced. To avoid expensive commercial flash ADCs in the GHz range, a previous development at PSI has been followed [4]. For a different experiment [5] an analog switched capacitor sampling chip has been designed. This chip consists of 128 capacitors, which sample the PMT input signal at a frequency of 500–1200 MHz. Instead of generating and distributing the sampling frequency directly, the capacitor switches are

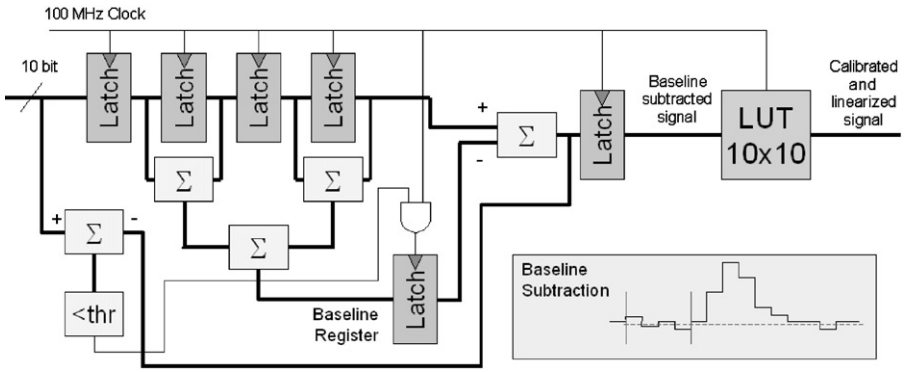

Fig. 1. Schematics of the baseline subtraction.

driven by an inverter chain, where a trigger signal propagates with high speed though a series of 128 double inverters. Since this is analog to a domino wave, the chip was named Domino Sampling Chip. When a waveform has been stored in the capacitors, they are read out at low speed (5 MHz) via a shift register and a commercial ADC (12 bit).

For the MUEGAMMA experiment, it has been concluded that a sampling speed of 2 GHz (500 ps bin width) is necessary to obtain a timing resolution of 50 ps via bin interpolation. If the domino wave runs in a circular fashion constantly and is only stopped by the trigger, the sampling capacitors can be used as an analog pipeline and delay cables can be avoided. This new chip, called Domino Ring Sampler is currently designed at PSI. The number of storage capacitors has been increased to 1024 and the readout speed via the shift register to 40 MHz. Eight data channels and one calibration channel are integrated on a single chip, which is housed and read out by a custom VME board using 12 bit flash ADCs and FPGAs. The sampling depth of 512 ns accommodates a typical LXe scintillation pulse of 100 ns width and a pipeline delay of 412 ns which is well above the expected level one trigger latency. The estimated costs per channel are USD 25 for the board and USD 6 for the chip. A first prototype using a 0.25μ radiation hard technology has been submitted recently.

The readout of all nine channels per chip with a single 40 MHz flash ADC consumes 230 μs, which is an acceptable dead time for an experiment running at $\sim$100 Hz. After the readout, the waveform is processed in the onboard FPGAs. At a first step a baseline subtraction is applied in the same way as in the trigger boards. Then a zero suppression logic discards channels which have no hit to reduce the amount of readout data. The FPGA algorithm checks if a channel contains a single hit or if a pile-up occurred by calculating the first derivative and searching for zero crossings. In case of a pile-up, the full waveform is transferred to the backend computers for further analysis. If no pile-up occurred, the signal charge is calculated by numerical integration and the time is evaluated by fitting the leading edge. These two numbers are equivalent to traditional ADC and TDC values and sent to the backend computers instead of the full waveform which then reduces the amount of data dramatically.

Since the FPGAs are re-programmable through VME, the algorithm can be changed and optimized during the set-up of the experiment. If additional devices like scalers or constant fraction discriminators are necessary, they can be easily emulated in the FPGAs. It is planned to use VME boards with the DRS chip for the calorimeter readout and the trigger boards with 100 MHz digitizing for the readout of the anodes and cathodes of the drift chambers, thus eliminating the need of any other device except pre-amplifiers and waveform digitizers. This reduces the hardware costs for the experiment by nearly an order of magnitude.

4. Slow control electronics

Most experiments use an inhomogeneous set of buses and devices for slow control like GBIP, RS232, PLCs, PCs, multimeters and terminal servers. It then becomes problematic to feed all this data into the central DAQ. For the MUEGAMMA experiment, it has been concluded that a field bus system with controller nodes directly connected to the sensors and actors would be beneficial. The field bus would then be connected directly to the backend computer, so that all software components of the data acquisition system have direct access to the slow control components of the experiment. While there exist many different commercial field bus systems, they are usually not optimized for this type of application, which results in an protocol overhead and the problem of interfacing these systems to experiment devices like high voltage power supplies or vacuum equipment.

To avoid these problems, a new field-bus-like system has been developed. It utilizes a new generation of microcontrollers [6,7] which contains eight channels of 12-bit ADCs, two channels of 12-bit DACs, 24 channels of digital IOs, a serial interface, a flash EPROM and a microcontroller on a single chip. The communication over the bus is done asynchronously via the serial interface and a RS-485 transceiver running at 340 kBaud. By using low load transceivers, up to 256 nodes can be connected to a single bus segment. The physical bus connection consists of 10-wire flat cable, which also carries supply voltages for the nodes. This makes it possible to design a minimal node with a footprint of only 40×40 mm^2 (Fig. 2). This node can be attached either directly to sensors, or used as a piggyback board on other electronics boards. Several 3 HE boards have been designed using this node for analog, digital, temperature and 220 V input–output, fulfilling the interface requirements for the complete experiment. Two special interface nodes have been designed to attach either to a standard RS-232 PC port or to the parallel PC port on one side and to the field bus on the other side, a node for the USB standard is planned. These nodes are used to interface the backend PC to the field bus, and to connect third pary RS232 devices to the bus via a protocol translator software running on the node.

This field bus together with the protocol definition is called Midas Slow Control Bus (MSCB), which expresses the tight integration into the Midas DAQ system [8]. The protocol is

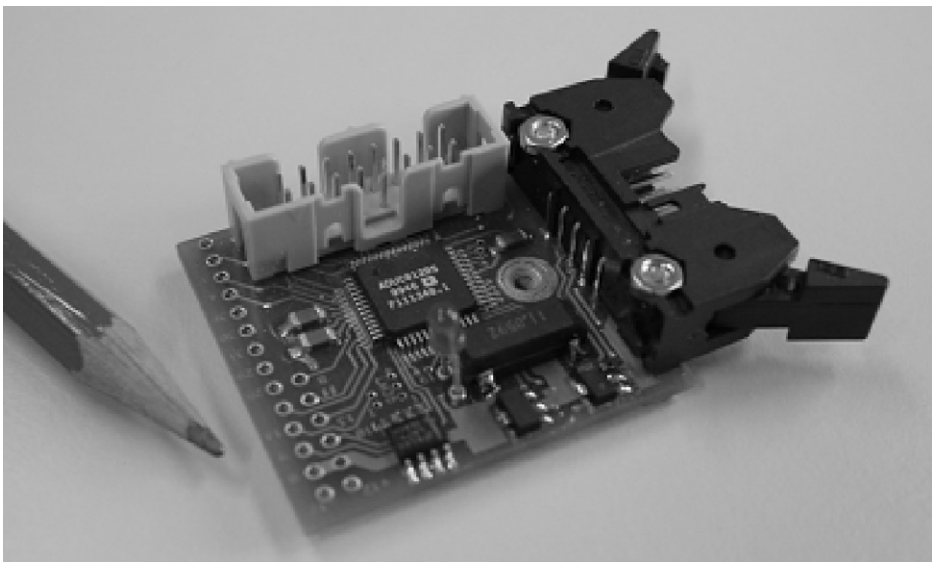

Fig. 2. Slow control node.

highly optimized for the requirements in a typical experiment. Only a few bytes, protected against transmission errors by a 8-bit CRC code, are necessary to read out a channel in a master–slave configuration, where the central backend PC acts as the master. The whole readout cycle lasts less than 300 μs which is superior to most existing field bus systems. Since the protocol uses 16-bit addressing, up to 65 536 nodes can be connected to the same bus. Each node can contain up to 256 internal "channels", which can be addressed individually and correspond to either an analog or digital I/O point. In addition, each node has up to 256 "configuration parameters", which are stored in the node's EEPROM and can be used by the application program running on the node. A typical usage are offset and gain constants. The microcontroller nodes can be programmed in the C language and re-programmed through the field bus. Each node contains the description of its channels and parameters in the EEPROM, which makes the network self-documenting and a central configuration database needless. On the PC side, a C library and LabView drivers have been written.

5. High voltage system

A new high voltage supply based on the MSCB system has been designed as well. Each channel contains a microcontroller node and a series of opto-couplers to regulate an external applied high voltage between zero volts and the external voltage (Fig. 3). Since each opto-coupler can sustain about 200 V, a series circuit of 12 opto-couplers are good for 2400 V regulation range and a current of 1 mA. The non-linearity of the opto-couplers is compensated by a regulation loop, where the output voltage is measured via a voltage divider and compared to a demand value from the microcontroller DAC. By using over-sampling, the output voltage is measured by the microcontroller ADC with a resolution of 16 bit. To ensure high stability and a low-temperature coefficient, two high accuracy reference voltages are distributed to all channels. Each microcontroller continuously performs self-calibrations using these voltages, which yields in a absolute voltage

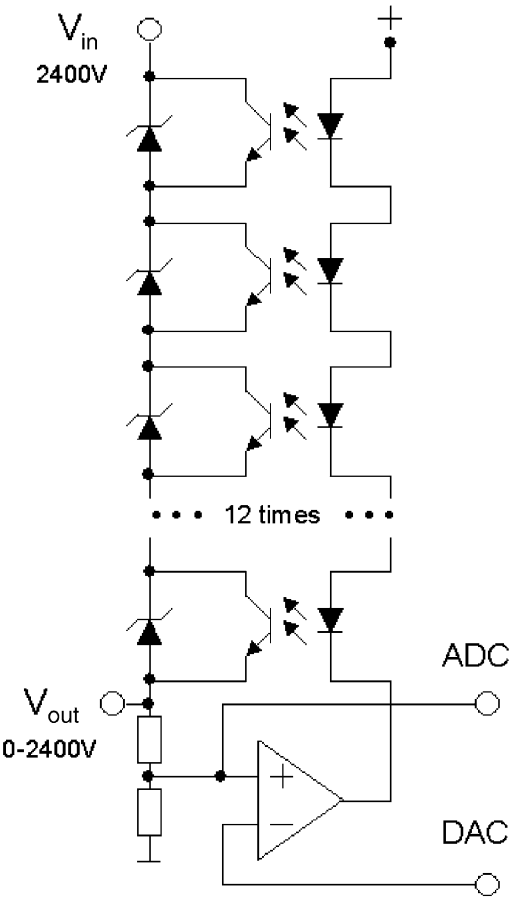

Fig. 3. Simplified schematics of a high voltage channel.

accuracy below 0.3 V over the full temperature range in a typical experiment. An additional current measurement not shown in the schematics can be used for diagnostics and to implement a current trip with a response time of about 10 μs. A motherboard houses 12 high voltages channels, and a 19 in. crate contains 192 channels at a cost of about USD 30 per channel excluding the external high voltage supply. Since the channels are directly connected to the MSCB system, a full crate can be read out in less than 60 ms from the central DAQ system.

References

[1] T. Mori, et al., PSI R-99-05 Experiment Proposal, Paul Scherrer Institute, Villigen, 1999.

[2] Analog Devices AD9218 Datasheet, http://products.analog.com/products/info.asp?product=AD9218.
[3] Xilinx Virtex E Series, http://www.xilinx.com.
[4] C. Brönnimann, et al., Nucl. Instr. and Meth. A 420 (1999) 264.
[5] http://pibeta.psi.ch.
[6] http://www.cygnal.com.
[7] Analog Devices ADuC812 Datasheet, http://products.analog.com/products/info.asp?product=ADuC812.
[8] http://midas.psi.ch.