

Fast waveform digitization with the DRS chip

S. Ritt

Abstract—The DRS chip was developed recently at PSI, Switzerland, using a 0.25 μm radiation hard CMOS technology. It implements a series of switched capacitor arrays (SCA), which allow the digitization of signals at speeds up to 5 GHz, at a power consumption and fabrication cost orders of magnitude lower than conventional flash ADCs. This allows a new generation of experiments with superior pile-up rejection and pulse shape discrimination, while simultaneously eliminating the need for traditional ADCs and TDCs. This paper explains the operating principle of the DRS chip and describes the deployment in the MEG experiment using 3000 channels in the MIDAS DAQ framework. Real time aspects of the data acquisition are covered and solutions are shown how to overcome the 880 MB/s raw data rate of the MEG experiment.

Index Terms—Switched capacitor array (SCA), Application-specific integrated circuit (ASIC), Waveform analysis, Data compression.

I. INTRODUCTION

THE MEG experiment [1] is currently being build at PSI, Switzerland. It searches for the forbidden decay $\mu^+ \rightarrow e^+ \gamma$ with a sensitivity of 10^{-13} . A beam with $10^8 \mu^+$ /sec is stopped and its decay products are registered with a drift chamber system, a timing counter array and a liquid xenon calorimeter, with a total amount of about 3000 read out channels. The high precision of the experiment requires amplitude resolutions of 12 bits and timing resolutions of 100 ps together with excellent pile-up rejection, which can only be accommodated with waveform digitization in the GHz range. Instead of using commercial flash analog-to-digital converters (FADC), a so-called switched capacitor array (SCA) chip has been developed at PSI, following an earlier development [2]. This solution is not only cheaper than commercial FADCs and takes much less power, it also has a higher channel density so 32 channels can be fit on a VME board, being able to digitize PMT and drift chamber signals at frequencies from 500 MHz to 5 GHz.

II. THEORY OF OPERATION

Since it is very hard to generate and distribute clock signals in the GHz range, the sampling frequency is generated on the chip with a series of daisy-chained inverters (Fig. 1). A wave propagates freely through these inverters and opens a series of storage capacitors, which sample an input signal from a com-

mon bus. The inverter chain is connected in a circular fashion, so the wave runs continuously until a trigger signal arrives, thus the name Domino Ring Sampler (DRS) was chosen for this ASIC. After the waveform is stored, it is read out via a shift register, which switches the capacitor contents to a common readout bus. The analog signal is then externally digitized with a commercial FADC with 33 MHz and 14 bits.

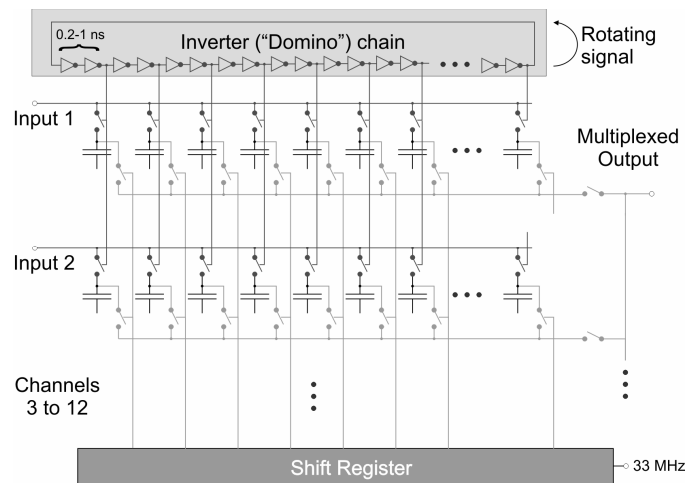


Fig. 1 Simplified schematics of one channel of the DRS chip.

The DRS3 chip is the third version of this design, designed with a radiation hard layout and fabricated in the UMC 0.25 μm 1P5M MMC process. The chip is 5 mm \times 5 mm large and contains 12288 sampling cells, which can be arranged as 1, 2, 3, 4, 6, or 12 channels then with 12288, 6144, 4096, 3072, 2048 or 1024 cells, respectively. The chip features a special readout mode which selects only a certain region-of-interest (ROI) for output and brings down the readout time to 3 μs for waveforms which contain 100 cells of interest for example. The DRS3 chip consumes 50 mW of power when operated at 2 GHz sampling speed.

The domino wave runs freely and its speed can be controlled by an analog voltage supplied by an external source. Since the domino speed depends on temperature and supply voltage, some stabilization is necessary. An external phase-locked loop (PLL) circuit regulates the domino speed via the analog voltage to match a reference clock in frequency and phase. The time jitter of this regulation is typical 200 ps. In addition the MEG experiment distributes a low-jitter 20 MHz reference clock to the 12th channel of each DRS3 chip. This clock signal is recorded together with 8 signal channels for calibration of each event.

The readout of the sampling cells inside the DRS chip has

carefully been designed using operational amplifiers to deliver a linear input range of 1 V with a nonlinearity smaller than 0.5 mV and a temperature coefficient of 50 ppm/deg.C. The cells have an offset distribution (the so-called “fixed pattern noise”) of 6 mV (RMS), which gets corrected in the FPGA during readout, leaving a random noise of 0.25 mV (RMS). This is equivalent to a 12 bit signal-to-noise ratio. The bandwidth (-3db) of the DRS3 chip is 450 MHz. A small design change is planned to improve this even further.

III. DRS USAGE IN THE MEG EXPERIMENT

Two DRS chips are put on a PMC mezzanine board and two mezzanine boards are housed on a 6 HE VME carrier board resulting in 32 signal channels per VME board. The DRS chips are read out with 14 bit FADCs and a VIRTEX II PRO field-programmable gate array (FPGA). The FPGA contains two Power-PC CPU cores, which can be used for low level waveform analysis. Using mezzanine boards has the advantage that the MEG experiment can start with the existing DRS2 chips, which will later be replaced by DRS3 chips, without the need of replacing the expensive VME boards.




Fig. 2 VME board with one DRS2 mezzanine card installed (left). The second mezzanine card is removed and turned around to expose the two DRS2 chips (right bottom). The upper right mezzanine card contains two new DRS3 chips. It is connected to a USB interface board for test purposes.

Nine VME crates are connected to about 2000 channels from the anode wires and cathode strips of the MEG drift chamber system and to about 1000 photomultiplier channels. The drift chamber signals are digitized with 500 MHz, giving a time window of 2 μ s, while the PMT signals are digitized with 2 GHz. The 512 ns signal window accommodates for the 370 ns latency of the level-1 trigger in the MEG experiment and leaves 142 ns for the PMT signal.

The VME crates are read with STRUCK SIS3100 VME interfaces via fiber optical links. Each crate is connected to one front-end PC. Event building is done over Gigabit Ethernet by a central PC containing a large disk array for logging. In a first engineering beam time at the end of 2006, a few days of data taking have been performed which produced about one TB of data. The assembly of all detectors will continue in 2007 and physics data taking is planned starting in 2008 for several

years.

Reading out only waveforms from all detectors instead of ADC and TDC values is a paradigm shift. The amount of data increases by more than two orders of magnitude and needs new methods for data reduction and compression. On the other hand it gives many new ways of looking at the data and improving the signal quality significantly. First algorithms have been developed and used in the MEG experiment. A moving average algorithm has been applied to determine the baseline of a signal. By correlating neighbor channels, common low frequency noise can efficiently be suppressed. The residual baseline noise achieved by this method of 0.07 mV in the timing counter system of the MEG experiment is a huge improvement compared to previous experiments.

Good progress has also been made in pile-up recognition. As a rule of thumb, it can be concluded that pile-up can be identified if two hits are separated in time by about the rise-time of the signal. The liquid Xenon calorimeter signals have a rise time of about 8 ns, and Monte Carlo studies have shown that two hits of comparable size can be recognized if they are 10 – 15 ns separated. In the case of pile-up a “template fit” has shown good results. Many single waveforms are averaged for each PMT channel to obtain an average waveform which we call template. In case of n -fold pile-up, a sum of n templates is fitted to the signal, only leaving the position and amplitude of the templates as a fit parameter. This results in n position and amplitude values, which are comparable to TDC and ADC values.

IV. REAL-TIME ASPECTS

Each recorded signal contains 1024 bins in the MEG experiment. The amplitude and time of each bin are encoded in 12 bits, respectively. At the target event rate of 100 Hz, the 3000 channels produce a data rate of 880 MB/sec. Since this data rate is too high to be recorded, extensive data compression is required. Furthermore, in case of the DRS2 chip some complicated calibration using cubic splines is required to compensate for the nonlinearity of the chip. The data reduction algorithms are usually first developed in the front-end PCs in C++ and later moved upstream into the readout FPGA or then embedded PPC cores if applicable.




Fig. 3 The readout chain of the MEG experiment from the front-end electronics with the DRS chip to the back-end PC. The data transfer from the front-end PCs to the back-end PC is handled by the MIDAS DAQ system [3].

Following data reduction schemes have been successfully implemented or are currently under development:

- Huffman encoding for waveform compression. Different algorithms like one-dimensional JPEG or wavelet compression have been tried but it was concluded that non-lossless algorithms are can in-

roduce dangerous artifacts.

- Re-binning of the waveforms. Some section of the waveform like the “tails” of a PMT hit are less important so several bins can be averaged and put into one bin.
- Zero suppression for “empty” channels. For each event the baseline noise distribution is calculated and a channel is considered to contain no hit if no bin deviates from the baseline average by more than $5 \times \sigma$.
- ADC/TDC evaluation using the template fit. Since the waveform is discarded after the fit, this technique is only applied to less important event types such as calibration events.
- Region-of-interest (ROI) readout of the waveform. For each subdetector a different waveform region in respect to the trigger is defined. The signal outside this region is discarded.
- Anode-cathode correlations in the drift chamber system. If an anode wire does not contain any hit, the corresponding cathode signals are discarded even if they contain some signal.
- Third level trigger. The back-end PC calculates some global quantities for each event and applies appropriate filtering.

After some of these algorithms have been implemented on the front-end PCs, it became obvious that a multi-threaded architecture is necessary. The MEG experiment uses nine front-end computers with dual Xenon CPUs, each connected to one VME crate. This distributes already the computing load to the PCs, but no usage of the two CPUs inside each PC is made, which limits the overall data rate to 7 Hz, far below the desired rate of 100 Hz. Studies have shown that the bottleneck is the calibration algorithm for the DRS2 chip running on the front-end PC. Only one CPU is occupied with a load of 70%, since it waits the remaining 30% for the VME data transfer which is done in PCI master mode and therefore does not require the CPU.

The CPU clock speed currently hits its physical limit around 3 GHz and won't improve in the future as it has done in the past. What changes however is the number of cores per CPU. Soon we will have quad core CPUs and more are likely to come. To make use of this technique, a multi-threaded architecture has been implemented in the framework of the MIDAS DAQ system on the front-end PCs. One thread transfers the data from VME and distributes it to n calibration threads via a set of ring buffers in a round-robin fashion. This distribution is done on a “zero-copy” basis, meaning that the VME data is already written into the ring buffer and no additional memory copy is required. A careful optimization of the ring buffer routines resulted in a very low overhead of this event distribution which can handle 10,000's of events per second with a CPU load of only a few percent. Each calibration thread picks up an event from the input ring buffer, calibrates and analyzes it and puts it into an output ring buffer, from where it is retrieved by

a collector thread and sent over the network to the back-end PC as shown in Fig. 4.

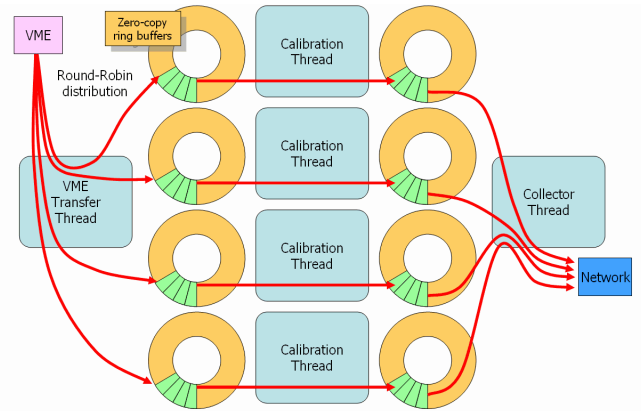


Fig. 4 Multi-thread architecture implemented on each front-end PC.

The Xenon CPUs support hyper-threading, which allows the partitioning of one CPU into two virtual CPUs giving an average performance improvement of about 30% compared to the single CPU. With hyper-threading each front-end PC contains four virtual CPUs and it was concluded that four calibration threads give the best performance. Each of the four calibration threads occupy about 90% of one CPU, while the VME readout thread and the collector thread each use 20%. This adds up to 400%, indicating that all four CPUs are optimally used. Using this technique, the event rate could be improved from 7 Hz to 30 Hz without any zero suppression. With zero suppression and the ROI readout mode, the data rate can be improved even above 100 Hz.

V. CONCLUSION

Waveform digitizing in the GHz range with the DRS chip opens new exciting possibilities for pile-up recognition, noise reduction and pulse-shape discrimination. It eliminates completely the need of traditional ADCs and TDCs at an attractive cost. The downside is however that waveform recording tremendously increases the demand for computing power in online environments, which can only be accommodated by new strategies such as effective multithreading. Given the current development in the CPU market with its multi-core technology, this will be an important trend in the future of real time computing.

REFERENCES

- [1] T. Mori *et al.*, PSI R-99-5 Experiment Proposal, Paul Scherrer Institute, Villigen, 1999
- [2] C. Brönnimann *et al.*, NIM A420, 264 (1999)
- [3] S. Ritt, P.A. Amaudruz, <http://midas.psi.ch>