ATLAS Internal Note
LARG-NO-050
29 October 1996

A 16 BIT-40 MHZ READOUT SYSTEM BASED ON DUAL PORT ANALOG MEMORIES FOR LHC EXPERIMENTS.*

Julien Ardelean, Etienne Augé, Roger Bernier, Michel Bouchel,
Dominique Breton, Bernard Lavigne, Gisèle Martin-Chassard,
Irena Nikolic, Eric Plaige, Jean-Jacques Veillet.
LAL ORSAY, 91405 Orsay cedex, France

Eric Delagnes, Michel Huet.
CEA-DSM-DAPNIA Saclay, 91191 Gif sur Yvette, France

ABSTRACT

The very high collision rate and the scarcity of interesting events foreseen at the future Large Hadron Collider at CERN require the development of new types of readout electronics structures. Among those, the dual port analog memories provide the possibility to decrease the high rates of incoming analog data and thus to divide by an important factor the amount and the performances of the ADCs necessary for the analog to digital conversion. These memories allow indeed to perform this conversion only after the level 1 trigger, thus dividing by at least 100 the amount of samples to digitize, and to use much slower ADCs. Nevertheless, this induces that they are able to sample the incoming analog signal at high frequency (40 Mhz for the LHC) and to store it waiting for the level 1 trigger latency. Moreover, the very high dynamic range of the detector signal requires the analog memory to cover at least 12 bits of resolution, even with a multigain system.

The Laboratoire de l'Accélérateur Linéaire (LAL) from Orsay and the CEA-DAPNIA from Saclay have developed a complete readout solution to fit the requirements of the ATLAS liquid argon calorimeter system. The upstream part of the chain consists in warm preamps followed by bi-gain bipolar shapers, dual port analog memories (also named « analog pipelines » or « Switch Capacitor Arrays »), and 12bit-10MHz ADCs. These components are currently split on two different boards, but should be gathered on the same one in the future.

1. REQUIREMENTS

The basic requirements for the readout system are the following :
- sample at 40 Mhz the signals coming from the shapers.
- store data during the level 1 trigger latency (>= 2us).
- read 5 samples per event accepted by level 1 (the rate should be at most 75 kHz) and perform the analog to digital conversion.
- format and transmit the data to DSP boards which will provide on the fly feature and energy extractions before sending data to the level 2 event buffers.
- operate fully simultaneous write and read operations and deal with interleaved events.
- cover a dynamic range of 16 bits without degrading the calorimeter resolution (0.7% for the biggest signals).
- feed the level 1 trigger system with analog sums of the input signals.

The total number of channels to be equiped is very high (200,000) and the electronics has to fit in a limited volume with stringent constraints on power dissipation and accessibility.

Furthermore, the radiation is at the level of 1 to 2 krad per year.

2. SYSTEM DESIGN

As the detector front-end part of the system (preamp+shaper) was already existing and well caracterised in test beam [ref 1], the main work was to design and build the downstream part of the system, what led to the realisation of two different boards. The first one, called *Front-End board*, is intended to realize all the dataflow operations. The second one, called *Pipeline Controller*, has to provide the control of all the write and read sequences on the former.

These two boards have a 3*Europe x 340mm format with a VME 32bit interface. The crate is a standard 3*Europe VME crate with 21 slots, and the two non VME busses are private. The bottom one is used for transferring all the command lines from the *controller board* to the *Front-End board*. The middle one is used for sending the event data from the *Front-End board* to the board housing optical drivers. At the other end of the fibers, a third type of board housing DSPs provides feature and energy extractions [ref 2].

3. FRONT-END BOARD SPECIFICATIONS

* Paper presented at the Second workshop on Electronics for LHC experiments, Balatonfüred, September 1996

This board is able to receive and treat 32 channels coming from the detector. To fit the above requirements, we chose to design it in the following manner (Fig. 1) :
- the detector signals dynamic range of 16 bits going out of the preamplifiers is divided in two ranges of 12 bits. This operation, performed by bi-gain shapers, allows to split the 32 signals into two pairs of 12-bit analog memories without degrading the detector resolution.

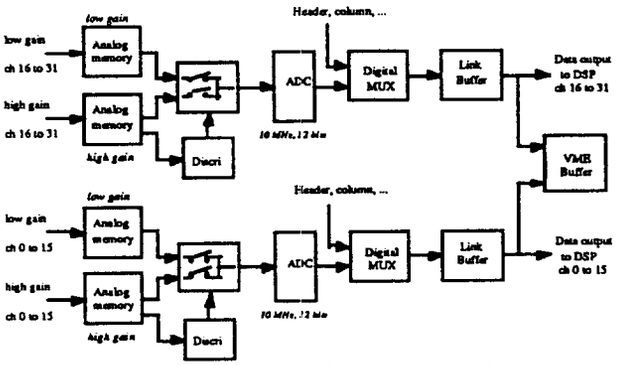

Fig. 1 : block diagram of the *Front-end board*

- inside each memory, the signal is sampled at 40 Mhz and stored until a level 1 trigger is received by the SCA controller. Then the five samples corresponding to the selected event are driven sequentially towards the output of the chip.
- at this point and for each sample, the internal 16 channels are multiplexed at a 10 Mhz rate. For each channel, the highest gain is compared to a saturation threshold (see Fig. 2). Depending on the result of the comparison, either the high or the low gain is sent to the ADC. The latter is a low cost-low power CMOS 12-bit 10 Mhz ADC.
- on the board, two sets of 16 bi-gain channels are implemented, each one with its own ADC. The data output by each of the latter is formatted and temporary stored in a FIFO (*link buffer*), waiting for being sent to the board providing the energy and timing extraction. The links between the two boards are provided by optical fibers and the data rates output by each ADC reach 200 Mbits/s on each fiber (including protocol).

At the output of the board, there is a spy memory to store on the fly the event data. This memory can be read through VME. This allow either to spy the data transfer or to realize the data acquisition if the DSP board is not connected downstream. In parallel, the *link buffer* can be loaded through VME to allow the checking of the data links.

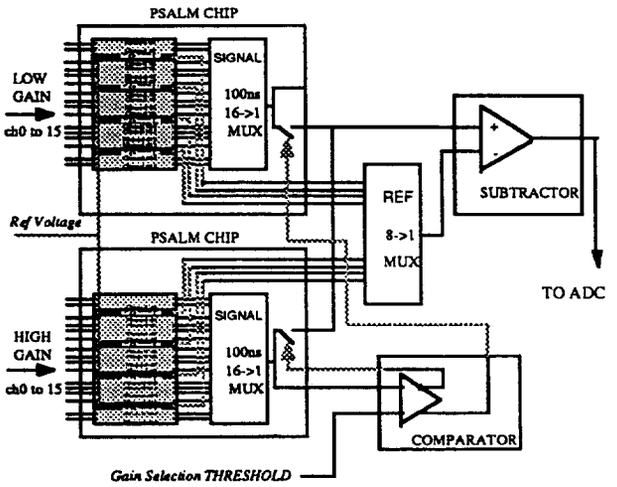

Fig. 2 : analog part of the *Front-end board*

On the *Front-end board*, a special care has been taken to avoid as much as possible the crosstalk between the digital part running either at 40 or at 10 Mhz and the analog signals. This implied the use of two common ground planes covering the complete board, special routing of critical signals, and decoupling and filtering of power supplies. Moreover, the most perturbative digital lines are transmitted in differential ECL.

4. PIPELINE CONTROLLER SPECIFICATIONS

The *Pipeline controller* board is intended to provide the complete control of the *Front-end board*. This is done through a private bus located on the back panel of the crate.

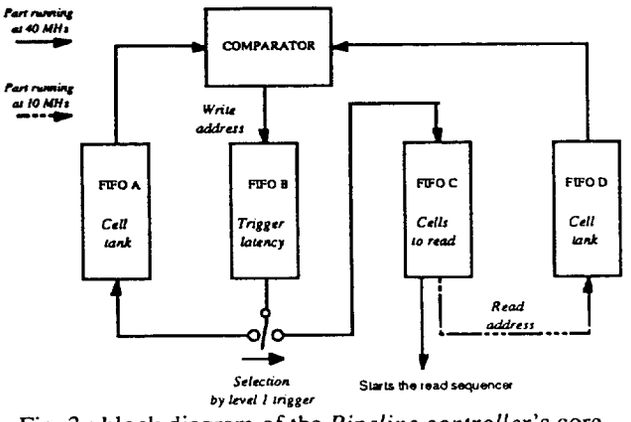

Fig. 3 : block diagram of the *Pipeline controller*'s core

The tasks to perform are the following :
1. concerning the analog memory :
- provide the 7-bit write and read address.
- during the read phase, provide the signals necessary for the sequence and the channel 4-bit address for the output multiplexor.

2. concerning the board environment :
- send the ADC 10 Mhz clock and tranfer the event data into the buffer FIFO (link buffer).
- drive the digital event formatters and the output buffers.

This controller already fits the full ATLAS performances. This implies that it is able to run at 40 MHz for driving the complete sequences with a trigger rate up to 100kHz. Its core is based on a block of 4 FIFOs and one comparator (see Fig. 3). This system deals with the 128 cell addresses which are continuously running through the FIFOs. It allows to provide simultaneous read and write, and to deal with interleaved events. Moreover, it ensures a high level of reliability for debugging the potentially doubled or missing addresses. Thus, in any case of error, flag bits are set up in the dataflow to inform the central Data Acquisition system.

This controller is presently being integrated in an ASIC in order to be mounted directly on each *Front-end board*. This will allow to avoid the back pannel interconnexions and to improve the flexibility for the tuning of the differents boards.

5. ANALOG MEMORY SPECIFICATIONS

The analog memory used on the *Front-end board* has been developped both in the AMS CMOS 1.2um and DMILL BICMOS rad-hard technologies [ref 3].

The Tab. 1 summarizes the main features of the two chips.

	AMS Version	DMILL Version
Number of cells	128	128
Number of signal channels	16	8
Number of reference channels	4	2
Storage capacitors value	0.45pF	0.75pF
Memory cell size.	187um*54um	140um*45um
Reference multiplexor	external	on-chip
Amplifier input transistors type	NMOS	PMOS
Chip size	6.4*9.6 mm2	3.1*7.9 mm2

Tab. 1 : Main features of AMS and DMILL chips.

The memory consists in a switched capacitor array where the analog signal is sampled, stored, and read as a voltage. The global architecture of the chip is given on Fig. 4.

The AMS version contains 20 analog channels. 16 of them are used to store the signal coming out from the shaper. The remaining 4, equally spaced in the chip (Fig. 2), store a reference level. During the read-out operation, an off-chip amplifier subtracts the closest reference channel to each signal channel. This pseudo-differential working has shown a common mode noise rejection with a factor higher than 4 when the chip is used in simultaneous read and write operation.

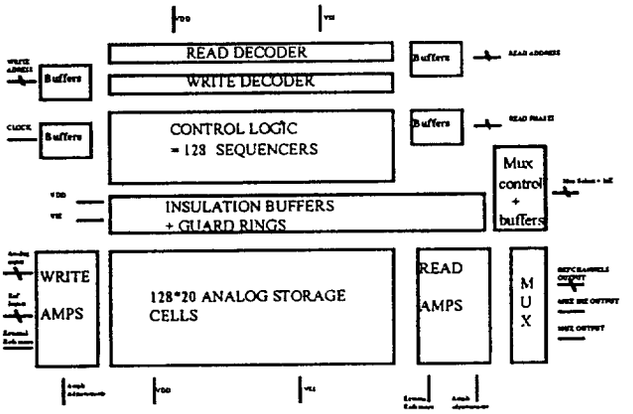

Fig. 4 : analog memory architecture.

Each of the 20 channels (Fig. 5) consists in 2 write amplifiers, 128 storage cells, a read-out amplifier and a pedestal analog subtraction system not used in our application.

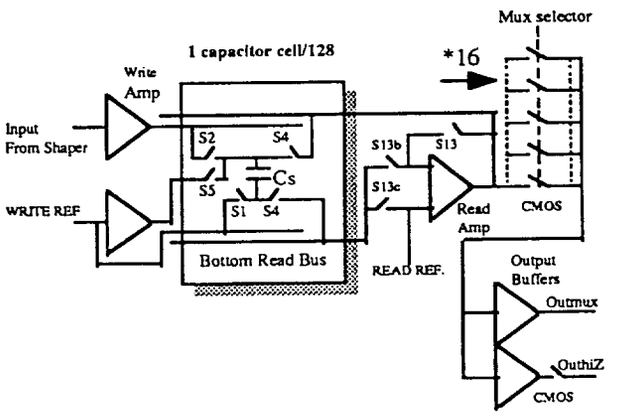

Fig. 5: channel block diagram

The write buffer is an O.T.A. type amplifier with enhanced slew-rate capabilities (Fig. 6) used as a voltage follower. The same structure with different transistors sizes is used for all the other amplifiers of the chip. The purpose of the write amplifier is to present a constant high impedance to the shaper output and to avoid signal distortion caused by high dI/dt in the bondings. Its characteristics are the following :

- 100Mhz Gain-Bandwidth product in order to minimize the contribution of the write amplifier to the filtering.

- Very good unity-gain stability, to avoid any signal distortion.

- Less than 1% integral non linearity on a 4.5V range with (VDD-VSS)=6V.

- 100 V/us slew-rate to limit distortion for high energy pulses.

- Low noise (80uV rms in a 100Mhz bandwidth).

- Very low power consumption (6 mW).

All this characteristics are maintained for a capacitive load in the range of 4pF to 10pF.

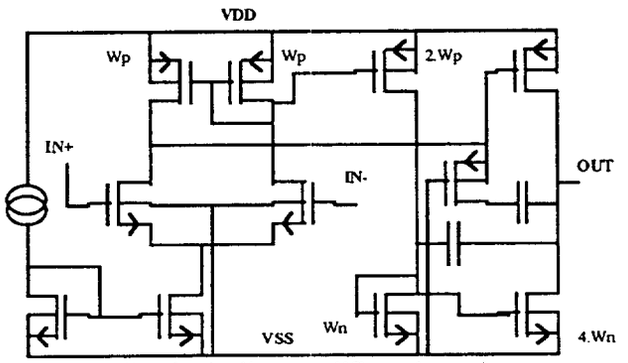

Fig. 6: pipeline amplifier schematics.

At each leading edge of the 40MHz clock, the write address is decoded and one of the 128 write sequencers is enabled and starts the following actions :

- The storage capacitor is reset to a reference voltage (S1 and S5 closed) to ensure that the voltage stored on the capacitor will be absolutely independent of those present on this capacitor at the beginning of the write operation, and to decrease the voltage swing necessary to track the signal.

- Then, the S1 and S2 switches connect the capacitor between the output of the write amplifier and the reference write bus. The duration of the loading of the capacitor is two clock cycles, longer than 12 ln(2) x (Rs1+Rs2) = 30ns (where Rs1 and Rs2 are the on-state resistors of the S1 and S2 switches), necessary to achieve a 12-bit precision.

- The S1 switch is turned off a few nanoseconds before S2, so that the injected charge and sampling time are made independent of the input voltage. As S1 is a minimum size switch, the value and the cell-to-cell spreading of the injected charge are kept small. A special care has been taken in the layout of the clock and signal distribution to limit both the sampling time jitter and the cell-to-cell sampling time spreading.

As the write operation lasts five clock cycles, 6 of the 128 sequencers are always active.

Before each read-out operation, a residual charge due to the limited open loop gain of the read amplifier is stored on the parasitic capacitor Cb of the 'bottom' read bus. So, when a read address is selected, this bus is immediately reset (S13 and S13c closed, S13b opened). S13c must be a quite big switch to allow a fast (50ns) reset. Its high injected charge during its turn-off is compensated by the correct sizing of S13b. This new reset operation presents important advantages compared to the conventional one [ref 4] :

- There is no risk of oscillations during reset.

- The noise sampled on Cb during the reset is low (only kT/Cb due to S13c), compared to the noise of the read-out amplifier in a classical reset.

After this operation, S4 switches are turned on, connecting the storage capacitor across the read amplifier. After 350ns, corresponding to the amplifier settling time, the 16 channels are sequentially multiplexed toward the 2 output buffers at the rate of 10 Mhz. One of the two outputs is sent to the gain discriminator and the other, switchable, is sent to the ADC. The total 16-channel readout time including multiplexing and gain selection is less than 2us.

All the digital inputs use ECL levels. The 40Mhz clock and the write address inputs are differential. To limit on-chip D/A couplings, the digital and the analog parts are separated by guard rings and have different power supplies. The switches command busses may be vectors of digital to analog pollution, via parasitic capacitors with analog busses or via couplings through the analog switches themselves. Thus to keep these signals noiseless, a third set of power supplies is provided especially for the command buffers.

6. SYSTEM PERFORMANCES

The performances summarized on Tab. 2 have been measured on the VME *Front-end boards* working in a simultaneous 40Mhz write/read operation, and with the ADCs continuously running at 10 MHz. 5 samples are picked up for each event. Fig. 7 shows a typical high signal sampled by the pipeline and digitized by the 12bit ADC. The input of the preamplifier is a calibration pulse. The absissa is 25nS per cell.

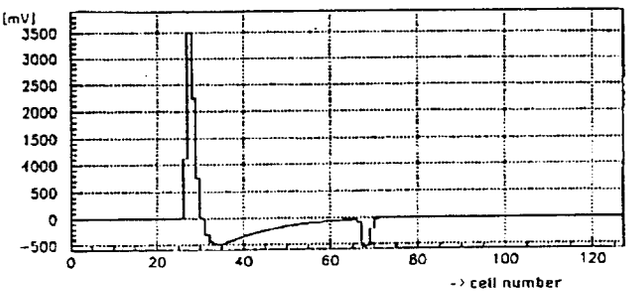

Fig. 7 : typical response for a large calibration signal

The performances of the AMS and DMILL chips are quite similar, i-e a more than 12bit dynamic range and a better than 100ps time resolution.

	AMS Version	DMILL Version
total noise (mV rms)	0.8	0.65
fixed sequence noise (mV rms)	0.3	0.2
maximum signal (V)	4.5	4.5
dynamic range (bits)	12.3	12.6
relative non-linearity on a 3V range	<1%	To be measured
cell-to-cell gain non uniformity (% rms)	<0.015	To be measured
cell-to-cell sampling time dispersion	<40ps rms	250ps pp
sampling time jitter	<40ps rms	<40ps rms
crosstalk with upper channel	1.5 %	<0.3%
crosstalk with other channels	<0.05%	<0.05%

Tab 2. : performances of the boards with two different chips.

The fixed sequence noise is due to the fixed part of the sequence environment at the moment where the signal is sampled in the capacitors. There are two major contributions :

- the dispersion of the switch charge injection along the cells,
- the position of the write pointer, but also other signals that could have the same state every time you sample on a given capacitor.

Nevertheless, this part of the noise remains very small (0.2 to 0.3 mV) and has to be added in quadrature to other sources.

The 1.5% crosstalk measured between a channel and its upper neighbour on the AMS chip is due to a 7fF parasitic capacitor between the output of the read amplifier of one channel and the 'bottom' read bus of the upper one (layout error). It has been improved on the DMILL chip as shown in Tab. 2. All the other crosstalks are very small.

The difference in sampling time dispersion between the two chips is mainly due to a difference in the routing of the clock signals. Thus a tree distribution doesn't seem to be the best solution.

As shown on Fig. 8, the relative non linearity of the total chain (calibration, preamplifier, shaper and pipeline) is less than 1% on a 3.5V range.

All those meaurements have been done on the *Front-end board.* Thus the noise characteristics for example are not due only to the chip but also to all the possible sources of coupling between it and the digital environment. The results for the chip alone on a dedicated test bench would be of course still better.

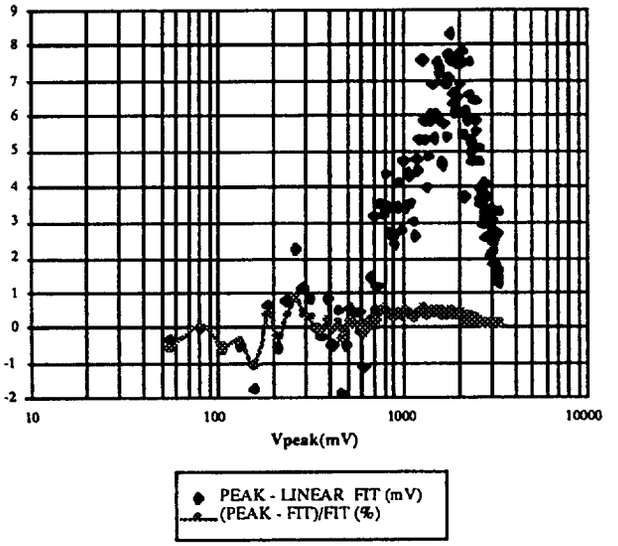

Fig. 8 : global chain linearity measurement.

7. TEST BEAM

7.1 . Test beam setup

A system consisting of 5 Front-end boards equiped with AMS chips and a common Controller has been used to read the prototype ATLAS liquid argon calorimeter in a test beam at CERN in june 1996. 120 calorimeter cells with warm preamplifiers and bi-gain shapers have been equipped. The complete setup is shown on Figure 9. Electrons of energies 20, 50, 100, 150, 200 and 300 GeV have been recorded with the beam pointing at a « reference point ». A few other electron runs have been taken with the beam in another position. Calibration data (Pedestal and Test-Pulse runs) were also recorded. We present some preliminary results below.

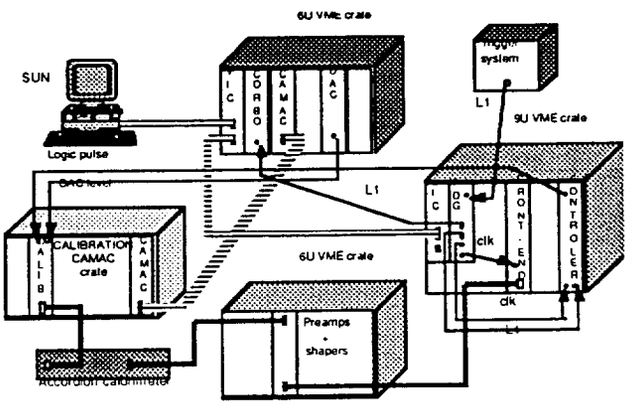

Fig. 9 : setup for testbeam at CERN in june 1996.

7.2 . *Conditions of the data taking*

Because of the shaper boards modularity, only 12 channels in each analog memory chip are receiving shaper signals. The input of the remaining 4 channels is connected to ground. The readout is performed via VME by a Sun workstation, sitting 35 meters away from the front-end crate. The trigger rate is small (at most 50 Hz), no new trigger being allowed before the readout of the previous one is completed. As a result, the write addresses are always running in the same order going from 0 to 127.

The setup is such that no VME activity takes place when an event is written into the analog memory, and when it is read from it and digitized. On the contrary, the 40 MHz clock ,the writing in the analog memory and the 10 MHz ADC clock never stop. For each event, 5 time samples are read-out to tape.

The pedestal trigger is synchronous with the 40 MHz clock. This is also the case for the calibration pulses, the third sampling falling at the maximum of the pulses. The physics trigger selects electrons such that the third sample falls within +- 1.5 ns of the maximum. Up to now, we have used only the third sample to estimate the signals amplitude, with a small correction for time jitter in the case of electrons (at most +- 0.5 %). No attempt has been done yet to correct for cross-talk.

Up to an electron energy of 200 GeV, only the high gain channels are used. On the contrary, both gains are used for 300 GeV electrons. The data at 300 GeV are not yet analysed.

The shaper signals being bi-polar, an offset of 0.6 V was chosen, leaving at most 3.4 V to measure the positive lobe of the signals.

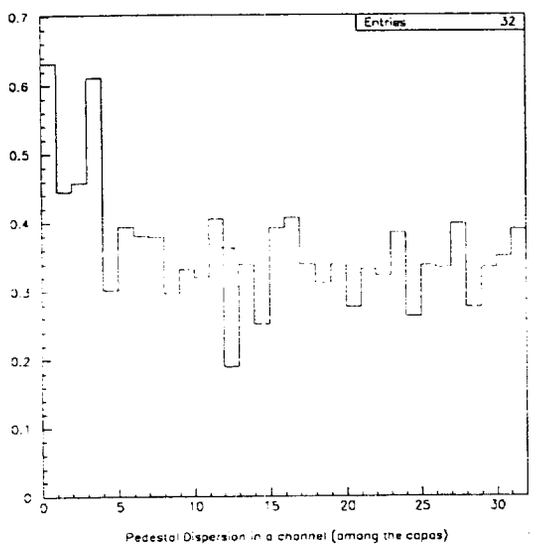

Fig. 10 : fixed sequence noise per channel

Two high gain channels and one low gain one are not working because of a bad cable between shapers and front-end boards. One of these channels affects slightly the measurement done at the reference point.

7.3 . *Fixed Sequence Noise*

Figure 10 shows for each channel of board 1 the dispersion of the 128 pedestal values (one pedestal value per capacitor). This gives a measurement of the fixed sequence noise, equal to 0.3 mV (11 MeV for the high gains). This quantity should be added in quadrature with the noise if the write address increases randomly.

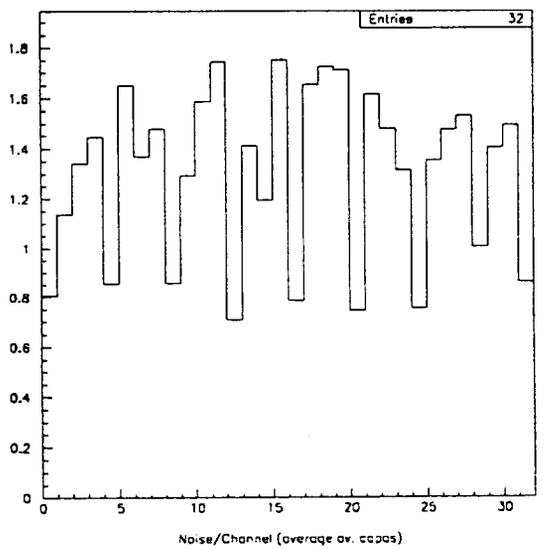

Fig. 11 : noise per channel in board 1

7.4 . *Noise*

Figure 11 shows the noise per sample for each high gain channel of board 1. For the channels connected to shapers, the noise is of the order of 1.5 mV (equivalent to 56 MeV) while it is only 0,85 mV (32 MeV) for the unconnected ones. For the low gains, the noise if 0.95 mV (507 MeV) per sample.

We apply the electron reconstruction algorithm on pedestal events, and on events taken randomly during the physics runs. A typical energy distribution is shown in Figure 12, showing a noise contribution of 250 MeV in the high gains, including 100 MeV which are coherent over the front, middle and back layers of the detector. For the low gains, the noise contribution to electrons is 1.77 GeV, with a coherent contribution of 365 MeV. Given the ratio between high and low gains, equal to 14.15, one concludes that the noise contribution from before the shapers is 210 MeV (out of which 100 MeV coherent)

whereas the noise coming from the shapers and the front-end boards is 125 MeV (out of which 25 MeV coherent) in the case of high gain. For comparison, the electronics used previously gives a contribution of the order of 300 MeV.

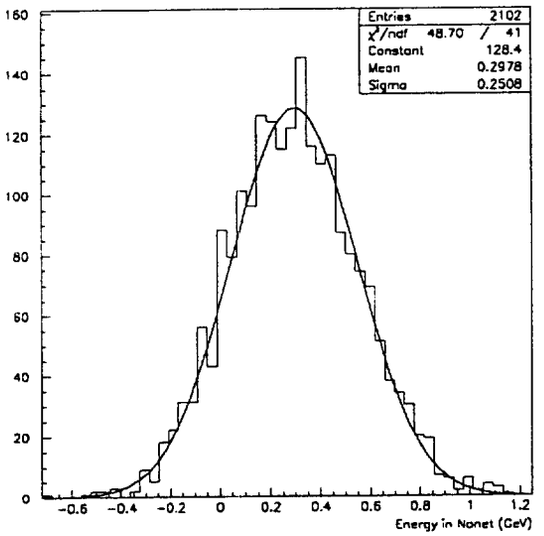

Fig. 12 : energy spectrum for random events

On the contrary, the noise is dominated by the shapers and front-end boards contributions in the case of low gains.

7.5 . *Calibration Pulses*

The variation of the measured signals (after pedestal subtraction) as a function of the DAC driving the calibration pulse generator is fitted for each channel with a polynomial of degree 3. The gains are uniform between channels within better than 10%. The non-linear terms are small (respectively -6% and +3% of the linear term at full scale). The fit residuals are all within +- 0.1% of the full scale.

Figure 13 shows the resolution on the amplitude measurements in each channel for both gains as a function of the amplitude. To a good approximation, it is given by the noise terms of 56 MeV for the high gain and 507 MeV for the low gain, as expected from the pedestal runs, down to resolutions below 0.1% : we have no evidence of a constant term above 0.1% in the resolution.

Saturation is observed around 3V above pedestal, both for high and low gains. It comes from the amplifier sitting between the shaper and the analog memory. The ratio between saturation and noise is therefore equal to 2000 for the high gains (11 bits), and to 28300 for the bi-gain system (14.8 bits). The dynamic range of the front-end board alone is 49950 (15.6 bits).

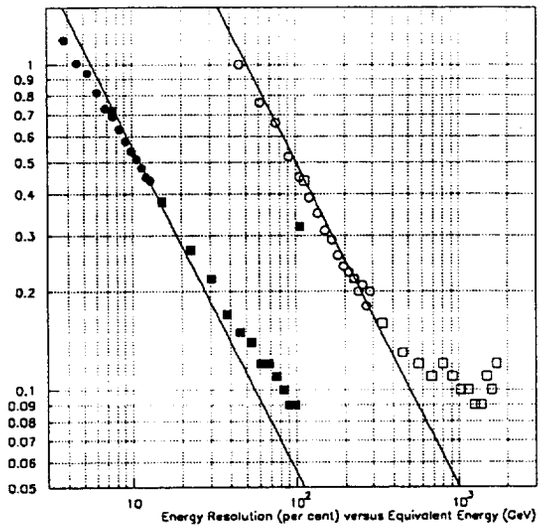

Fig. 13 : energy resolution for calibration signals

7.6 . *Electron Data*

As for the usual analysis, corrections have been applied as a function of the impact point coordinates within the impact cell, and as a function of the sharing of the energy between the front, middle and back layers of the calorimeter.

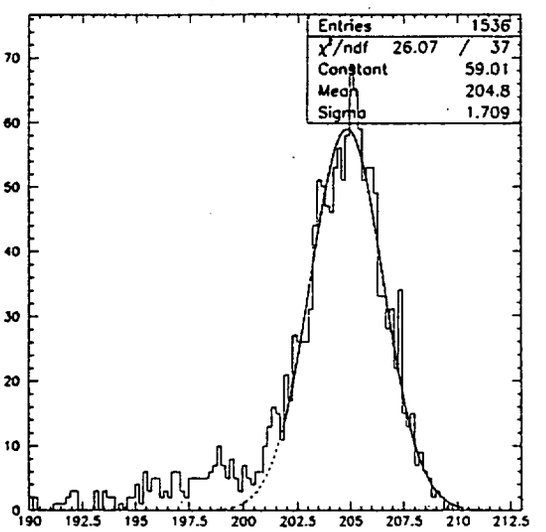

Fig. 14 : energy spectrum for 200 GeV electrons

Figure 14 shows the energy distribution for 200 GeV electrons. The resolution is 0.83 %, comparable to the one obtained with usual electronics. However, 10 % of the events appear in a low energy tail which is not explained at the moment, and which does not seem to come from the front-end boards. At 50 GeV, the resolution is 1.64 % and a much smaller tail is observed.

Figure 15 shows the energy resolution as a function of energy. As previously mentionned, the data used here are affected by a dead channel (cell in a corner of the nonet in the front layer), and the energy resolution is slightly degraded with respect to the numbers quoted just above.

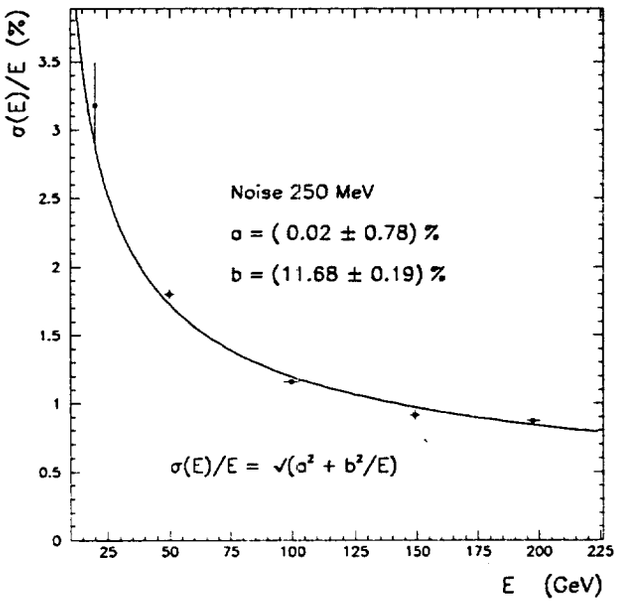

Fig. 15 : energy resolution versus beam energy

A beam energy spread equal to 0.15% at 200 GeV and 0.29 % at the other energies has been subtracted.When the noise term is fixed to 250 MeV in the resolution, the fit gives a sampling term of 11.7%, slightly higher than usual, and a constant term compatible with 0. The quoted errors include systematic and statistical contributions. If the sampling term is forced to be as low as 11%, then the fit is poor and the constant term remains below 0.3%, which is still smaller than what is observed with usual electronics.

8. CONCLUSION

The preliminary results show performances in real beam conditions at the level expected from the test bench measurements. One can therefore anticipate that , with shapers integrated on the front-end board, and with the DMILL version of the chip, the front-end boards fullfill the requirements and may be used in the readout of the ATLAS accordion calorimeters. Furthermore, the common work recently started with teams involved in similar developments at Nevis labs and at the University of Alberta has already brought a few ideas to further improve the level of performances of the analog memory chip.

In conclusion, the 4 year R&D effort naturally converts now into the design of the real Front-end board for the ATLAS accordion calorimeters. The first aim is now to equip several thousands channels of «module 0 », both for the barrel and for the end-caps parts early in 1998.

REFERENCES

[ref 1] : RD3 collaboration, "Performance of a large scale prototype of the ATLAS electromagnetic calorimeter", Nuclear Instruments and Methods, A364 (1995) 290-306.

[ref 2] : J.Colas & al, "MUSIP, A 40 MHZ FILTER CARD FOR THE ATLAS ELECTROMAGNETIC CALORIMETER SIGNAL", Second Workshop on Electronics for LHC Experiments, September 1996.

[ref 3] : M.Dentan & al, "DMILL, A MIXED ANALOG-DIGITAL RADIATION-HARD BICMOS TECHNOLOGY FOR HIGH ENERGY PHYSICS ELECTRONICS", I.E.E.E T.N.S vol. 43, N° 3, part 2, June 1996.

[ref 4] : S. Kleinfelder, " A 4096 CELL SWITCHED CAPACITOR ANALOG WAVEFORM STORAGE INTEGRATED CIRCUIT", I.E.E.E T.N.S vol. 37, N° 3, June 1990.