

The Readout System of the ATLAS Liquid Argon Calorimeters

A. Blondel^a, D. La Marra^a, A. Léger^a, G. Perrot^b, L. Poggioli^b, J. Prast^b, I. Riu^a, S. Simion^c

^aDépartement de Physique Nucléaire et Corpusculaire,
Université de Genève, 24 quai E. Ansermet, 1211 Genève 4, Switzerland

^bLaboratoire d'Annecy-le-Vieux de Physique des Particules (LAPP)
9 Chemin de Bellevue, BP 110 74941 Annecy-le-Vieux CEDEX, France

^cNevis Laboratories
136 South Broadway, P. O. Box 137, Irvington NY 10533, USA

Abstract

The ReadOut Driver (ROD) system is a key element of the ATLAS Liquid Argon Calorimeters readout system. It processes a predetermined number of samples of the bipolar output waveform from the calorimeter front-end electronics and precisely determines the energy deposited in each calorimeter cell and the timing of these signals at the Level one trigger output rate of 100 kHz. It applies an optimal filtering algorithm while minimizing the pileup and electronic noise and using coefficient constants determined from the calibration.

Around 190000 channel outputs are processed through the Liquid Argon ROD system. Only their energy, timing and a quality flag are sent to the data acquisition. The impossibility to recover the original data imposes severe reliability requirements to the ROD system.

The system consists of around 200 ROD modules, 200 transition modules and 16 custom-made backplanes. A ROD module receives data from 1024 calorimeter cells through eight 1.6 Gbit/s optical fibers and consists of one mother board with four daughter boards (called processing units) which contain two Digital Signal Processors (DSP) each. This modular design offers the possibility to use the latest development on DSP technology in the future. Two different DSPs have been tested and the results compared. These results together with the description of the Liquid Argon Calorimeters readout system are discussed.

I. INTRODUCTION

The ATLAS [1][2] experiment is a general purpose proton-proton detector designed to exploit the full discovery potential of the Large Hadron Collider at CERN. The overall

design is the result of the requirements of high precision muon momentum measurements, efficient tracking, large acceptance and very good electromagnetic calorimetry for electron and photon identification and measurements.

Four different detectors constitute the Liquid Argon calorimetry of ATLAS [3]: the electromagnetic barrel, the electromagnetic and hadronic endcaps (HEC) and the forward calorimeter. In total, around 190000 calorimeter cell outputs are to be read out. A high signal sampling frequency (40 MHz), a large energy dynamic range of the readout cells (from 50 MeV up to 3 TeV) and a good relative energy resolution are some of the main challenges of the Liquid Argon readout electronics.

II. THE LIQUID ARGON READOUT ARCHITECTURE

Signals from the detectors are processed by various stages before being delivered to the Data Acquisition system (DAQ). Figure 1 shows a simplified diagram of the different boards that process the data in its way to the DAQ. The calorimeter cell signals are received by the Front End Boards (FEB) housed on-detector in a radiation environment. Digitized data is processed in the ReadOut Driver boards (ROD), located in a radiation free environment, before being sent to the DAQ system.

III. THE FRONT END BOARDS

The front-end electronics for all Liquid Argon detectors is identical apart from the amplification stage which is done in the cryostat for the HEC and in the FEBs for the others. In the FEB, which treats 128 calorimeter channels, the signals are amplified, shaped and stored in analog levels in a switched capacitor array every 25 ns. Upon receipt of a Level one trigger, five (or more) samples are digitized using three gains in the ratio 1/10/100. By examining, event by event, the amplitude of the sample closest to the peak of the signal, the

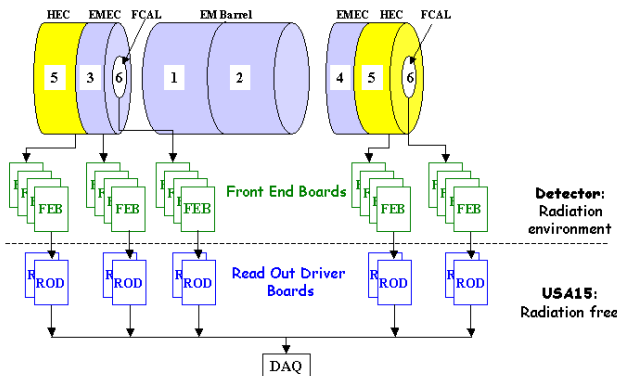


Figure 1: Readout architecture of the Liquid Argon Calorimeters.

gain scale is selected and the digitized samples transmitted to the ROD boards through optical fibers at 1.28 Gbit/s. The FEBs will be housed in 58 custom-made 9U crates.

IV. THE ROD SYSTEM

A. The Requirements

Among the main requirements of the ROD system are the following:

- High channel density.
- Modular design. Basic components should be easily changed/upgraded.
- The maximum event processing time per event including histogramming is 10 μ s (for a Level one trigger rate of 100 kHz).
- Low power consumption.

B. The Functionality

A ROD module receives data from eight FEBs, that is (typically) five digitized samples from 1024 calorimeter cells. Compared with the ROD demonstrator board¹, which reads 256 channels, the present ROD module design has four times higher channel density. In total, around 200 ROD modules will readout the Liquid Argon Calorimeters.

The module is responsible of calculating the energy and the time relative to the peak of the signal for each of the channels. Since the raw data from the FEB is no longer available offline, it also performs monitoring of the calorimeters by building and updating histograms for selected

¹ The ROD Demonstrator board [4] was designed and built in 2000 as a first intermediate step towards the final ROD module in order to demonstrate that commercial Digital Signal Processors can perform the optimal filtering algorithm calculations fast enough. In addition, it has been used to readout and calibrate the calorimeter cells in several test beams.

channels. Calibration tasks for each channel are also performed.

C. The basic algorithm

The algorithm implemented in the ROD to extract the energy and time for each channel is a technique called optimal filtering [5]. The idea is to estimate these quantities in an accurate and computationally efficient way. The energy (E) and time (T) are expressed as a weighted sum of the samples S_i , as shown in the following expressions:

$$E = \sum_i a_i \cdot (S_i - Ped)$$

$$E \cdot T = \sum_i b_i \cdot (S_i - Ped)$$

where i extends over all samples, Ped is the pedestal value, and a_i , b_i are the optimal filtering weights. These weights are found by requiring a simultaneous minimization of the errors on the energy and time while satisfying a set of constraints. As T is inversely proportional to E , it is meaningful to calculate T only for channels with E above a given threshold (E_{th}). A quality parameter of the fit indicating how closely the samples follow the known waveform is also calculated for the same channels as for T . It is a simplified chi-square, i.e. ignoring the correlations between the different terms:

$$\chi^2 = \sum_i (S_i - Ped - E \cdot g_i)^2$$

where g_i is the expected normalized waveform for a given channel. In addition, the histograms of monitored quantities, such as the values of E , T and χ^2 , are updated.

V. THE ROD PHYSICAL DESCRIPTION

In order to be modular and decouple the different ATLAS detector readout systems from the DAQ, the Liquid Argon ROD system is divided into two different set of boards: ROD modules and Transition Modules (TM). A ROD module, installed at the front of a 9U VME crate, is dedicated to data processing while a Transition Module, installed at the back of the crate, interfaces the detector readout to the DAQ system. Additionally, a custom-made backplane is used, among other things, to transfer the signals between them. A total of 16 9U VME crates will house the Liquid Argon ROD system.

The ROD module is a 9U VME64x board housed in a 9U VME crate with 21 slots. It is the responsible of processing the data and transferring the result to the TM. As required, a modular design has been chosen to allow for an easy upgrade of the DSP components. It consists of a mother board [6] and four daughter boards called Processing Unit boards (PU) [7] mounted on top. Figure 2 shows a simplified scheme of the ROD module. Serial data (16 bits at 80 MHz) from the FEBs is received by the ROD mother board through eight optical receivers and de-serialized by the G-link chips [8]. Four Field Programmable Gate Array (FPGA) chips, called staging

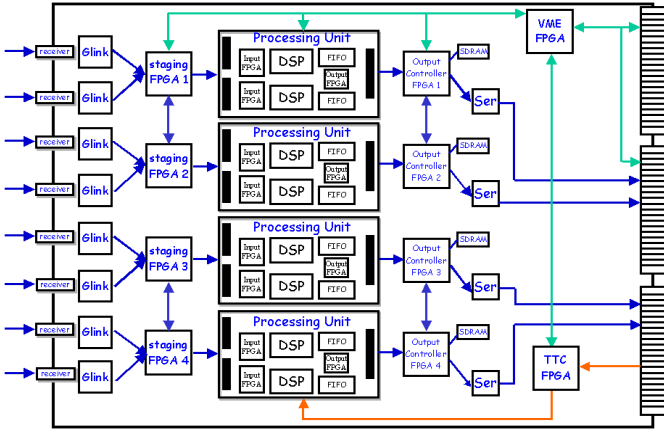


Figure 2: The ROD module scheme.

FPGAs, route the data from the G-link chips to the PU boards. Two DSPs are mounted in each PU and perform the optimal filtering calculations. The input FPGAs of the PU are used to convert the data serial to parallel, to check for data transmission errors and to appropriately rearrange the data for use in the DSP. The output FPGA is used to interface the VME and the Trigger Timing and Control (TTC) information² to the DSPs. The output data with the DSP calculations is stored in two FIFOs in the PU. Four FPGAs in the ROD mother board, called Output Controller FPGAs, get the data from the FIFOs and send it to the Synchronous Dynamic Random Access Memory (SDRAM) for monitoring purposes and to the serializer chips. These latter serialize and send the data in LVDS signals at 280 MHz to the TM. The VME FPGA interfaces the ROD with the VME and deals with the busy signal (signal generated by the ROD to stop the Level one trigger; for example, in case the DSP is busy with data processing). The TTC FPGA gets and distributes the TTC information to the ROD.

The Transition Module is a 9U VME64x board that has four de-serializer chips, four FIFOs and four S-link interface cards [9]. The de-serializer chips de-serialize and send the data to the FIFOs. These FIFOs are needed in order not to lose data when the signals that require data sending be stopped (“link down” and “link full” of the S-link protocol) come back from the DAQ. The S-link interface cards are the responsible of sending the output data to the DAQ (32 bits at 40 MHz).

A. Staging mode

Due to contingency, the ROD mother board will be equipped with only half of the PUs (the so-called staging mode) at the beginning of LHC. This is the reason why a data bus between staging FPGAs (32 bits at 80 MHz) has been introduced. Data from four G-link chips is routed through one staging FPGA to one PU board. Therefore, in staging mode

² The TTC information contain the trigger type, the event number and the trigger number.

the DSP processes double the number of channels (256 channels) than in normal mode with all PUs (128 channels).

VI. THE DSP EVALUATION RESULTS

Two Processing Unit boards equipped, respectively, with one Texas Instruments (TI) DSP C6203 and one TI DSP C6414 have been built by Nevis Laboratories and LAPP Ancey. Figures 2 and 3 show both boards.




Figure 3: PU with Texas Instruments DSP C6203.




Figure 4: PU with Texas Instruments DSP C6414.

The characteristics of the two DSPs are described in table 1. Both have fixed-point arithmetic and the main differences are the internal clock frequency, the memory and its organization.

Table 1: Main characteristics of the TI DSPs C6203 and C6414.

TI DSP C6203	TI DSP C6414
<ul style="list-style-type: none"> • 300 MHz; 3.3 ns core cycle • Fixed-point arithmetic • VelociTI™ architecture • Eight 32-bit instructions/cycle • 7 Mbit internal memory: 3 Mbit Program 4 Mbit Data RAM. • 32-bit EMIF • 4 DMA channels • 384-pin BGA package • 3.3V I/O, 1.5V core 	<ul style="list-style-type: none"> • 600 MHz; 1.67 ns core cycle • Fixed-point arithmetic • VelociTI.2™ architecture • Eight 32-bit instructions/cycle • L1/L2 Memory architecture: 128 kbit L1 Program Cache, 128 kbit L1 Data Cache, 8 Mbit L2 RAM/Cache. • 2 EMIF: 16-bit and 64-bit • 64 EDMA channels • 532-pin BGA package • 3.3V I/O, 1.4V core

A comparison of the average event processing time has been performed by using a code optimized for each DSP which processes 128 channels per event and calculates E , T and χ^2 for all channels by using the optimal filtering algorithm. The code also performs histogramming of E and T for all channels and all three gains. These histograms have 128 bins, 16-bit bins for low and medium gains and 32-bit bins for high gain and are updated event by event only for channels with $E > E_{th}$ (1.9 ADC counts). The DSP computation time is obtained real time by using one of the DSP counters and stored per event in the output data. Table 2 shows the average event processing time depending on the number of channels histogrammed (channels with $E > E_{th}$): 30%, 50% and 80% for each DSP. As expected, the PU with the 600 MHz DSP takes less time than the PU with the 300 MHz DSP. However, it does not take half of the time. The main reason is the different way the memory is organised in the DSPs.

Table 2: Average DSP processing time versus percentage of channels histogrammed.

	30 %	50 %	80 %
300 MHz DSP C6203	~7.2 μ s	~7.6 μ s	~8.4 μ s
600 MHz DSP C6404	~5.5 μ s	~6.2 μ s	~7.8 μ s

The TI C6414 DSP has two level memories: L1 (cache) and L2 (RAM). The cache memory is distributed between Data (L1D) and Program. The histogramming procedure involves three steps: read the bin content, update it and write it back into memory. When reading, the bin content might not be allocated in the L1D cache (that is called a L1D miss) and consequently, it takes more time since it has to be fetched from the L2 memory and a space allocated in the L1D cache. When writing the bin content back to memory, the fastest is to write it directly into the L2 memory. However, when the data is present in the L1D cache, it needs to be written there first (that implies a L1D dirty line) and only later, updated in the L2 memory. That takes longer also. These effects explain the fact that the more the number of channels histogrammed, the less powerful the DSP C6404 is compared to the DSP C6203. This effect can be seen in the results of table 2 when comparing the processing time between DSPs.

In summary, the requirement of a computation time not exceeding 10 μ s is met by both DSPs. Since the TI C6414 DSP has more memory and gives larger margin in the computation time requirement it has finally been chosen for the final ROD system.

VII. STATUS AND SUMMARY

The readout of the ATLAS Liquid Argon Calorimeters has been described. In particular, a physical description of the ReadOut Driver module and Transition Module has been given. A processing time evaluation of two different DSPs

have been done by using two processing unit boards built with two different Texas Instruments DSPs. The TI DSP C6414 has been chosen for the final ROD since it has more memory and is faster. The requirements to the readout which include a modular design and a maximum latency of 10 μ s are met. The power consumption has been estimated to be around 80 Watts by using some measurements of the PU consumption. First prototypes of the final ROD system are expected in 2003. A series production is expected to start in 2004.

VIII. REFERENCES

- [1] ATLAS collaboration, Technical Proposal for a General Purpose pp Experiment at the Large Hadron Collider at CERN, CERN/LHCC/93-43, LHCC/P2, 15 December 1994.
- [2] ATLAS collaboration, ATLAS Detector and physics performance Technical Design Report (TDR), CERN/LHCC/99-14, 25 May 1999.
- [3] ATLAS collaboration, ATLAS Liquid Argon Calorimeter TDR, CERN/LHCC/96-41, 15 Dec. 1996.
- [4] I.Efthymiopoulos et al. The ROD Demonstrator for the LArgon Calorimeter board description.
http://atlas.web.cern.ch/Atlas/GROUPS/LIQARGON/Electronics/Back_End/Rod/Archives/docs/lardemorod_v3r1.doc
- [5] W.E. Cleland and E.G. Stern, Nuclear Instruments and Methods A338 (1994) 467.
- [6] A.Blondel et al, The ROD Mother Board for the ATLAS Liquid Argon Calorimeters. Board description.
<http://www.cern.ch/Imma.Riu/>
- [7] J.Prast, The ATLAS Liquid Argon Calorimeters ROD, the TMS320C6414 DSP Mezzanine board PU documentation.
<http://dpnc.unige.ch/LArgROD/>
- [8] The G-link chip. Low Cost Gigabit Rate Transmit/Receive Chip Set with TTL I/Os.
<http://hsi.web.cern.ch/HSI/components/serialisers/hp/hdmp1022.html>
- [9] CERN S-link homepage.
<http://hsi.web.cern.ch/HSI/s-link/>