



HAL
open science

ATLAS liquid argon calorimeter back end electronics

J. Bán, A. Bazan, F. Bellachia, A. Blondel, S. Böttcher, A. Clark, J. Colas,
M. Díaz-Gómez, B. Dinkespiler, I. Efthymiopoulos, et al.

► **To cite this version:**

J. Bán, A. Bazan, F. Bellachia, A. Blondel, S. Böttcher, et al.. ATLAS liquid argon calorimeter back end electronics. *Journal of Instrumentation*, 2007, 2, pp.06002. 10.1088/1748-0221/2/06/P06002 . in2p3-00142571

HAL Id: in2p3-00142571

<https://hal.in2p3.fr/in2p3-00142571>

Submitted on 10 May 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ATLAS liquid argon calorimeter back end electronics

The Liquid Argon Back End Electronics Collaboration

J. Ban^a, A. Bazan^b, F. Bellachia^b, A. Blondel^c, S. Boettcher^{a,1},
A. Clark^c, J. Colas^b, M. Diaz Gomez^c, B. Dinkespiler^{d,2},
I. Efthymiopoulos^{c,3}, M. Escalier^{e,4}, Lo. Fayard^f, A. Gara^{a,5},
Y. He^d, F. Henry-Couannier^g, F. Hubaut^{e,2}, G. Ionescu^{b,6},
A. Karev^{h,*}, L. Kurchaninov^{h,7}, R. Lafaye^b, B. Laforge^e, D. La
Marra^c, S. Laplace^b, O. Le Dortz^e, A. Léger^c, T. Liu^d,
D. Martin^e, P. Matricon^f, L. Moneta^{c,3}, E. Monnier^g,
H. Oberlack^h, J.A. Parsons^a, S. Pernecker^c, G. Perrot^b,
L. Poggioli^{b,8}, J. Prast^b, H. Przysiezniak^b, B. Repetti^{g,9},
L. Rosselet^c, I. Riu^{c,10}, P. Schwemling^e, S. Simion^{a,8},
W. Sippach^a, A. Straessner^c, R. Stroynowski^d, S. Tisserant^g,
G. Unal^{f,3}, H. Wilkensⁱ, I. Wingerter-Seez^b, J. Yang^d, J. Ye^d
and A. Xiang^d.

^a Columbia University, Nevis Laboratories, 136 South Broadway, P.O. Box 137,
Irvington, NY 10533, USA

^b LAPP - Laboratoire d'Annecy-le-Vieux de Physique des Particules, 9 Chemin de
Bellevue, BP 110, F-74941 Annecy-le-Vieux Cedex, France

^c Université de Genève, Département de physique nucléaire et corpusculaire, 24 Quai
Ernest-Ansermet, 1211 Genève 4, Suisse

^d SMU, Southern Methodist University, Department of Physics, Dallas, Texas
75275-0175, USA

^e LPNHE - Laboratoire de Physique Nucléaire et de Hautes Energies, Universités
Paris VI et Paris VII, 4 place Jussieu, Tour 33, Rez de chaussée, F-75252 Paris Cedex
05, France

^f LAL - Laboratoire de l'Accélérateur Linéaire, Bât. 200, F-91405 Orsay Cedex,
France

^g CPPM - Centre de Physique des Particules de Marseille, 163 avenue de Luminy,
Case 902, F-13288 Marseille Cedex 09

^h MPI, Max-Planck-Institut für Physik, Werner-Heisenberg-Institut, Foehringer Ring
6, 80805 Munich, Germany

ⁱ CERN - Centre Européen pour la Recherche Nucléaire, 1211 Genève 23, Suisse

¹ Now at Christian Albrechts University, Institute for Experimental and Applied
Physics, Kiel, Germany.

² Now at CPPM.

³ Now at CERN.



⁴ Now at CEA, Saclay.

⁵ Now at IBM T.J. Watson Research Center, USA.

⁶ Now at Laboratoire des Images et des Signaux (LIS-INPG), Grenoble, France.

⁷ Now at TRIUMF, Vancouver, Canada.

⁸ Now at LAL.

⁹ Now at Laboratoire d'Astrophysique de Marseille (LAM), France.

¹⁰ Now at U. Autònoma Barcelona, Spain.

* On leave of absence from JINR, Dubna, Russia.

E-mail: julie.prast@lapp.in2p3.fr, helenka.przysieznia@lapp.in2p3.fr

Abstract. The Liquid Argon calorimeters play a central role in the ATLAS (A Toroidal LHC Apparatus) experiment. The environment at the Large Hadron Collider (LHC) imposes strong constraints on the detectors readout systems. In order to achieve very high precision measurements, the detector signals are processed at various stages before reaching the Data Acquisition system (DAQ). Signals from the calorimeter cells are received by on-detector Front End Boards (FEB), which sample the incoming pulse every 25ns and digitize it at a trigger rate of up to 75 kHz. Off-detector Read Out Driver (ROD) boards further process the data and send reconstructed quantities to the DAQ while also monitoring the data quality. In this paper, the ATLAS Liquid Argon electronics chain is described first, followed by a detailed description of the off-detector readout system. Finally, the tests performed on the system are summarized.

Submitted to: *Journal of Instrumentation*

Keywords: LHC, ATLAS, Liquid Argon calorimeter, back end electronics.

Contents

1	Introduction	5
2	LArg electronics chain	6
2.1	Readout architecture	6
2.2	Front End electronics	8
2.3	Back End electronics	9
2.3.1	Optimal filtering	10
2.3.2	The Back End electronics system	11
2.3.3	Timing and trigger control	14
2.3.4	Level-1 receiver system	15
3	LArg Back End electronics decription	16
3.1	VME crate	16
3.1.1	Power supply	16
3.1.2	Air and water cooling of the electronics	17
3.1.3	Detector Control System (DCS)	18
3.1.4	Crate backplane and CP3 board	19
3.2	SPAC Master module	20
3.2.1	Description and functionalities of the module	20
3.2.2	SPAC protocol and Slave ASIC	21
3.2.3	SPAC interface	24
3.2.4	Power consumption, power and reset procedures	24
3.2.5	SPAC design and manufacturing	25
3.3	TTC and Busy module (TBM)	25
3.3.1	Timing, Trigger and Control distribution	25
3.3.2	Busy collection	26
3.3.3	Control outputs	27
3.3.4	VME interface	27
3.3.5	Power, reset and ESD considerations	27
3.3.6	TBM design and manufacturing	28
3.4	ROD module	28
3.4.1	FEB data input	29
3.4.2	TTC data input	32
3.4.3	Output flow	32
3.4.4	SDRAM	33
3.4.5	Serializer	35
3.4.6	Busy handling	35
3.4.7	VME interface and internal communication	35
3.4.8	ROD design, manufacturing and tests	36
3.5	Processing Unit	36
3.5.1	Input FPGA (InFPGA)	37

3.5.2	Digital Signal Processor (DSP)	39
3.5.3	Output FIFO	41
3.5.4	Control FPGA (ConFPGA)	42
3.5.5	Boot procedure	43
3.5.6	Power supply	43
3.5.7	JTAG chain	44
3.5.8	PU design and manufacturing	44
3.6	Transition Module	45
3.6.1	FIFO memory	45
3.6.2	Flow control	46
3.6.3	Power, reset, ESD and EMC considerations	46
3.6.4	TM design and manufacturing	47
3.7	ROD Injector	47
3.7.1	TTC receiver	49
3.7.2	Sequence generator	49
3.7.3	Data injector	50
3.7.4	VME interface	50
4	BE system tests	51
4.1	Injector tests	51
4.2	Configuration of the system	51
4.3	Speed of the DSP code	52
4.4	Commissioning in the cavern	53
5	Conclusion	55




Figure 1. Schematic 3D view showing the ATLAS calorimeter system with its EM barrel and end-cap LArg calorimeters, hadronic barrel and barrel-end tile calorimeters, hadronic LArg end-cap calorimeters and forward LArg calorimeters.

1. Introduction

The ATLAS experiment [1] at the LHC is a general purpose detector designed to exploit the full physics potential of the LHC at CERN. The collider will produce proton-proton collisions at a centre-of-mass energy of 14 TeV, as well as heavy ion collisions.

Liquid Argon (LArg) sampling calorimeters [2] are used in ATLAS for all electromagnetic calorimetry covering the pseudorapidity interval $|\eta| < 3.2$, as well as for hadronic calorimetry from $|\eta| = 1.4$ to the acceptance limit $|\eta| = 4.8$. The ATLAS calorimeters are shown in figure 1, with the LArg system as well as the barrel hadronic system.

The LArg calorimeters consist of four sub-detectors and are contained within three cryostats. The central cryostat houses the electromagnetic barrel calorimeter (EMB), while each end-cap cryostat contains an end-cap electromagnetic calorimeter (EMEC), a hadronic end-cap wheel (HEC) and one forward calorimeter (FCAL). In total, about 196000 calorimeter cells are to be read out.

A choice of common electronics for all calorimeters standardizes the hardware and will simplify the maintenance, The HEC nonetheless uses cold preamplifiers.

The main tasks for the readout electronics are :

- to measure, for triggered beam crossings, the energy deposit in each calorimeter

cell to better than 0.25% at high energy. The dynamic energy range will cover a maximum of 3 TeV down to a lower limit of 10 MeV set by thermal noise. Coherent noise over many cells used to measure the jet energy should be kept below 5% of the incoherent noise. The readout should proceed without any dead time up to a trigger rate of 75 kHz.

- to provide the trigger system with the energy deposited in trigger towers of size $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$. The trigger processor combines the information from all ATLAS sub-detectors to deliver at the 40 MHz bunch crossing rate a yes/no decision to read out the detectors.

An extensive phase of R&D led to the final design of the LArg readout electronics, which were then constructed, tested and installed in the ATLAS cavern, and finally underwent extensive commissioning. The whole chain is described in the following section, as an introduction to a detailed account of the back end electronics system.

2. LArg electronics chain

A general view of the readout architecture is given, followed by brief descriptions of the Front End (FE) and Back End (BE) electronics, as well as sections on the Timing and Trigger Control (TTC) and Level-1 receiver system.

2.1. Readout architecture

The high radiation levels in the vicinity of the detector on the one hand, and the necessity for a large signal dynamic range and for very low level coherent noise on the other hand, favour remote electronics combined with preamplification very near the detector. In addition, the large number of cables required imply data multiplexing in any off-detector solution. The final readout architecture takes into account all of these points and is sketched in figure 2.

When a charged particle traverses and ionizes the liquid Argon in the gap between a LArg copper electrode and a lead absorber, an ionization current is measured on the readout cells of the electrodes (see Reference [2] for more details). The pulse height is proportional to the energy deposit of the particle. The analog signal is received and digitized by the Front End Boards (FEBs) which are mounted directly on the detector and are part of the Front End electronics. The FEBs send the digitized pulse via optical links to the Read Out Drivers (RODs) which are installed in a radiation-free area next to the detector cavern and are part of the Back End electronics.

A careful summation of the analog signals is also formed, mostly in the FEBs or in the Tower Builder Boards (Tower Driver Boards for the HEC; see following section), but also in a few special cases in the receiver system (barrel-EC transition region and FCAL) to build primitives which are sent to the BE trigger electronics. They are received, reshaped and summed into trigger towers of size $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ before being transferred to the Level-1 calorimeter processor.

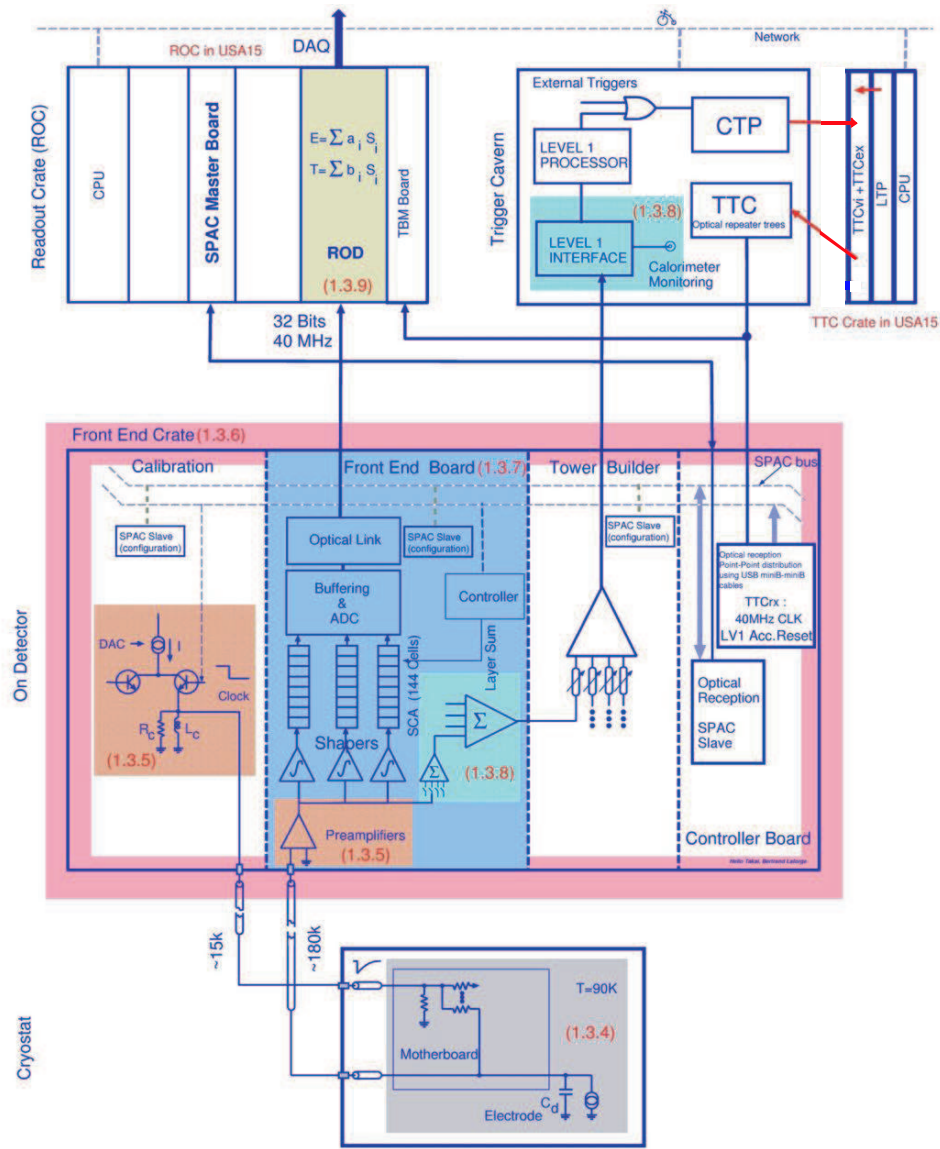


Figure 2. A block diagram of the readout electronics is shown, where the LArg calorimeter electrical circuit is schematized at the bottom of the figure, followed above by the on-detector Front End electronics crate (FEC) which contains amongst other modules the FEBs, while at the top left and right of the figure respectively are schematized a BE electronics crate containing the ROD boards (ROC: ROD crate) and a TTC crate.

Each ROD (192 in total) can be connected to up to 8 FEBs (1524 in total) and can process the signals of up to 1024 detector cells. The RODs send the data through optical links, the S-links, to the Read Out Buffers (ROBs) hosted on PCs. At this stage, the acquisition changes from a push mode driven by the Level-1 trigger to a pull mode driven by the event-filter farm, where event selection and building is performed, and the data are finally recorded.

2.2. Front End electronics

The sensitive analog electronics are housed on the detector. Inside the cryostat, the calorimeter electrodes are grouped to form readout cells and small coaxial cables bring the cell signals to the cold-to-warm feedthroughs (see Reference [2]). Front End electronic crates are mounted on the feedthroughs, in the crack between the barrel and end-cap calorimeters and at the rear of the end-caps. A Front End Crate contains several types of boards.

- The calibration board injects precisely known current pulses through high precision resistors located in the cold, to simulate energy deposits in the calorimeters.
- Front End Boards which:
 - amplify and shape the analog signals;
 - sum the calorimeter cells by trigger tower within each layer depth, and prepare the input signals for the tower builder board (see following item);
 - store the signals in an analog memory waiting for the decision by the Level-1 trigger;
 - digitize the selected pulses and transmit the multiplexed digital results on optical fibres.
- Tower Builder Boards (TBBs), for the EM barrel and endcap calorimeters, perform the final level of analog summation to form trigger tower signals and transmit the analog signals to the Level-1 cavern (USA15) for digitization and processing by the Level-1 trigger processor. As well, the board which builds the HEC trigger towers is called the Tower Driver Board (TDB) since no summation is implied here. The function of this board is to produce differential signals and to drive 70 m trigger cables.
- Front End Crate controller boards (FECcont) receive and distribute the 40 MHz clock, the Level-1 trigger accept signal (L1A), as well as other fast synchronous signals and information, to configure and control the FEBs and calibration boards.
- Various monitoring boards which ensure that the LArg system is operating correctly. They collect the information from the mechanical sensors which measure the position and stresses of the detectors, from the temperature gauges for the cryogenics and the LArg (LArg Temperature Monitor Board), and from the LArg purity probes (LArg Purity Monitoring Board).

The FEBs perform the first signal processing step. Each FEB receives the signals of 128 LArg detector cells. The triangular-shaped pulse of the ionization current of each detector cell is first amplified. In the case of the HEC only, the preamplifiers are installed directly on the detector and not on the FEBs. The preamplifier output is sent to shaper chips which amplify the pulse, split it into three gain scales in the ratio 1/9.3/93 and apply a bi-polar shaping function to each scale. The signal is then sampled at 40 MHz in Switched Capacitor Array (SCA) chips which store the samples in analog form during the Level-1 trigger latency ($< 2.5 \mu\text{s}$). Figure 3 shows the triangular pulse




Figure 3. Triangular pulse of the current in the LArg cell and the FEB output signal after bi-polar shaping. Also indicated are the sampling points every 25ns. Even though a maximum of 32 samples can be attained, only a few samples are typically used (5 to 7) during normal data taking.

current in the LArg cell as well as the FEB output signal after bi-polar shaping and sampling.

When the Level-1 trigger decision arrives, the optimal gain scale is selected on an event-by-event basis. The SCA samples are then digitized at a 5 MHz rate in a 12-bit Analog-To-Digital Converter (ADC) which, together with the gain selection procedure, fulfills the required 16-bit dynamic range over the whole energy interval. The FEB data are finally sent via the 1.6 Gbit/s G-link optical links to the RODs.

The time constant of the FEB shapers are chosen to minimize the overall noise level. The two main contributions are the electronic noise and pile-up noise, which respectively decrease and increase with the peak shaping time t_p , as can be seen in figure 4. The pile-up noise is due to background events occurring in the same beam crossing as a signal event, as well as in beam crossings occurring before and after. Digital signal processing in the BE electronics further optimizes the signal to noise ratio as a function of the actual luminosity.

2.3. Back End electronics

The BE electronics are contained in and make use of three systems situated in three different locations: the ROD which is the core of the BE, the TTC and the Level-1 receiver. The BE system must perform various tasks during physics and calibration runs. The physics mode is clearly the most demanding: the system reads the FE electronics, receives the TTC signals, processes the data and sends it to the acquisition system at a Level-1 rate of up to 75 kHz.

During calibration runs, signals of known intensity are generated and sent to the




Figure 4. Noise level as a function of the peak shaping time t_p . The electronic noise decreases with increasing t_p , while the pile-up noise increases. The optimal shaping time depends on the LHC luminosity L .

LArg detector cells by the calibration boards described previously. The trigger rate is lower than for physics (~ 10 kHz) due to the properties of the calibration pulser, and so there is more computation time available in the RODs.

The BE must also distribute the timing clock and trigger to the FE electronics and the RODs, as well as configure and control the FECs. More precisely, the FEB and TTC data readout and processing imply:

- buffering the events to absorb fluctuations in the time of arrival of triggers and to optimize the use of the available readout bandwidth;
- generating a Busy to pace the trigger processor;
- checking the FEB data consistency and the synchronization with the TTC stream;
- processing and monitoring of the data;
- reducing the data volume by a factor of ~ 2 ;
- sending out the processed data to the DAQ system.

All of these tasks must be performed fast enough to cope with the average Level-1 trigger rate of 75 kHz.

2.3.1. Optimal filtering The optimal filtering (OF) method [3] is at the core of the BE processing role. The OF algorithm, implemented in the ROD Digital Signal Processors (DSPs; see section 3.5), calculates the energy for each cell while minimizing the noise contribution. For cells with an energy above a certain threshold (approximately 5% to 10% of all cells), the precise timing of the signal as well as a data quality factor, a kind of χ^2 allowing to flag cells with large pile-up contribution, are determined.

The signal pulse is sampled every 25ns a programmable number of times, between 3 and 32. For physics, 5 samples are typically used, whereas for calibration, up to 32 samples can be taken. The energy E and time τ of the signal pulse are calculated in a weighted sum of the sample amplitudes, s_i :

$$E = \sum_{i=1}^N a_i (s_i - ped) \quad (1)$$

$$E\tau = \sum_{i=1}^N b_i (s_i - ped) , \quad (2)$$

where N is the number of samples and ped is the pedestal value of the corresponding readout channel. The optimal filtering weights, a_i and b_i , are evaluated while minimizing the dispersion in E arising from electronics and pile-up noise.

Since the error on τ is inversely proportional to E , the time is calculated only for channels above a given energy threshold, in the same way as the quality factor χ^2 . A simplified χ^2 , ignoring correlations between the sampling points, is computed to compare the measured pulse shape to the ideal pulse:

$$\chi^2 = \sum_{i=1}^N (s_i - ped - E(g_i - \tau g'_i))^2 , \quad (3)$$

where g_i and g'_i are respectively the normalized pulse shape and its derivative. These are estimated from the measured calibration pulse and from the electrical model parameters of the readout circuit (see Reference [4] for more details).

2.3.2. The Back End electronics system The BE electronics are situated in the Level-1 cavern (USA15), with approximately 70 m of optical fibres separating them from the detector. For each trigger, data from the TTC stream and from the FE electronics are pushed into the ROD modules where they are checked, processed, formatted and sent on a Read Out Link (ROL) to the ATLAS DAQ (see figure 2).

To increase modularity and allow for concurrent running of the various parts, the LArg calorimetry is split into 6 partitions: one for each half-barrel (EMB-A and EMB-C), one for each end-cap (EMEC-A and EMEC-C), one for the A side hadronic end-cap and forward calorimeter (HEC-A + FCAL-A) and one for the C side (HEC-C + FCAL-C). Partitioning is useful for debugging and calibration purposes, as well as for testing different parts of the detector independently of the others, during the installation phase and in case of problems during data taking. A photograph of the BE electronics racks can be seen in figure 5.

To each partition is associated a workstation (Partition Master), which controls and monitors the system, and a TTC sub-system (see section 2.3.3 for more details). A barrel partition e.g. EMB-A (see table 1 for the other partitions) is composed of 4 ROD crates and the associated FE electronics. Each ROD crate drives 8 half FECs and contains the following:

- A CPU connected to an Ethernet network controls the crate and the associated FECs.




Figure 5. The BE electronics racks in USA15. From left to right: the injectors rack with splitters above and the cooling system below, the FCAL and HEC rack, the EMEC-C rack, the EMEC-A rack, the four EMB racks, and finally 5 ROS racks.

- Two SPAC (Serial Protocol for the ATLAS Calorimeters) Master modules drive the 8 FECcont boards (one in each half FEC) to configure the various boards of the FEC and to read back status information.
- 14 ROD boards receive raw data from the 8x14 FEBs of the partition and calculate the energy, time and data quality factor.
- A TTC and Busy Module (TBM) receives the TTC information and delivers the TTC signals to the ROD boards in the crate via a dedicated CP3 backplane board.
- The TBM also collects the ROD Busy signals via the CP3 backplane.
- Transition modules (TM) (one per ROD) transmit the ROD result to the DAQ through the S-links (4 per TM).

A sketch of a barrel partition is shown in figure 6.

Such an architecture provides a completely autonomous partition which can run independently from the others. The composition of the other LArg partitions is very similar although the number of modules scales down with the number of channels to be read out as shown in table 1.

In addition, but in separate and dedicated crates, the ROD Injector module has been developed to simulate the flow of data coming from the FEBs and from the TTC system. It is used for testing the ROD boards and is described in detail in section 3.7. As an example, half a ROD crate (7 out of 14 RODs) can be driven using 6 Injectors.

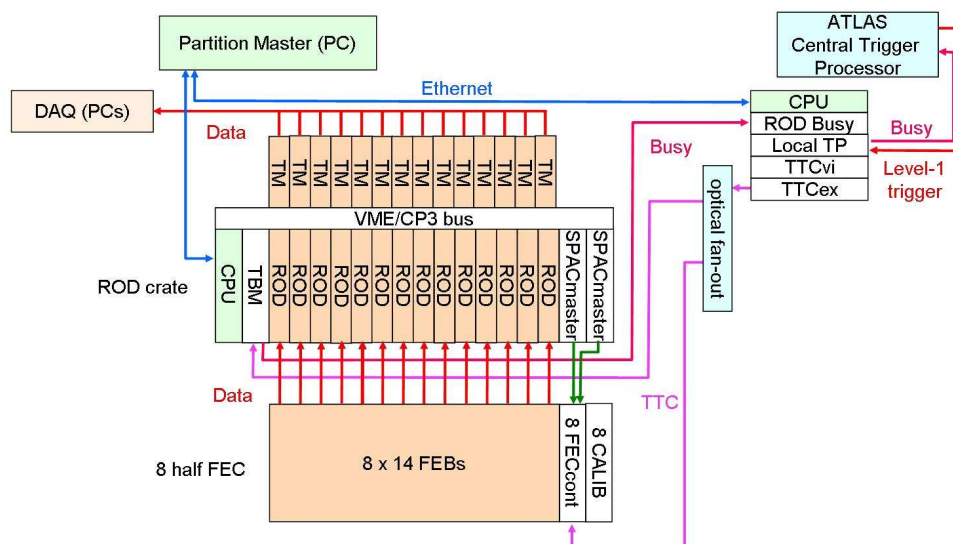


Figure 6. Sketch of a barrel partition. A FE crate (bottom) is schematized sending data to a BE crate (middle). The TTC crate (upper right except for the Central Trigger Processor) receives (Trigger) and sends (Busy) signals from and to the ATLAS Central Trigger Processor. The TTC signals are then sent to the ROCs and FECs. Also shown are the Partition Master PC as well as the DAQ PCs.

Table 1. Components of the LArg partitions.

Partition	EMB A + C	EMEC A + C	HEC-A + FCAL-A	HEC-C + FCAL-C	total
Work station	1 + 1	1 + 1	1	1	6
TTC crate	1	1	1	1	4
ROD crate	4 + 4	3 + 3	1	1	16
Crate CPU	4 + 4	3 + 3	2	2	16
SPAC Master	8 + 8	6 + 6	2	2	32
CP3 board	4 + 4	3 + 3	1	1	16
TBM	4 + 4	3 + 3	1	1	16
ROD board	56 + 56	35 + 35	5	5	192
TM board	56 + 56	35 + 35	5	5	192
FEB	448 + 448	276 + 276	38	38	1524
Calibration	32 + 32	28 + 28	5	5	130
FECcont	32 + 32	24 + 24	5	5	122
TBB	32 + 32	28 + 28	-	-	120
TDB	-	-	5	5	10

2.3.3. Timing and trigger control The TTC system uses common ATLAS components and relies in particular on the optical fan-out system developed within the framework of RD12 [5]. Four 6U VME crates are used to implement the system, where up to two partitions are implemented in a single crate. The signal flow between the FE, BE and TTC systems is illustrated in figure 6. Each TTC partition contains:

- a CPU,
- a local Trigger Processor (LTP), responsible for transmitting either the ATLAS global timing and trigger signals or locally generated timing and trigger signals (e.g. calibration is done in local mode) to the TTCvi and TTCex modules (see two items down).
- a ROD-Busy-Module. It receives from the TBM the combined result of all the ROD Busy signals of each ROD crate, and then delivers a partition wide Busy signal via the LTP to the Central Trigger Processor (CTP), which then throttles the trigger rate.
- TTCvi (TTC VME bus interface) and TTCex (TTC Encoder / Transmitter) modules [5]. The TTCvi encodes the TTC signals on two channels, one reserved for the L1A, the other for the LHC machine signals and various commands used by the system. These two signals are then sent to the TTCex, which merges them and performs the electrical to optical conversion. The optical signal is sent to the ROD crates and the FECs using optical fan-out modules, where it is electrically distributed to each relevant board.

One TTC link per ROD crate is needed, as well as two per half FEC for redundancy: in case of fibre or receiver failure, the second fibre is cabled replacing the previously connected one.

The LArg readout elements need the bunch-crossing signal (BC clock) from the LHC machine and the Level-1 accept signal (L1A) from the trigger system. The 40.08 MHz BC clock has to be precisely in phase with minimal jitter with the time of the crossings to benefit from the good timing characteristics of the calorimeters (better than 50 ps at high energy [1]). In addition, to synchronize all the readout elements, the LHC provides once per turn the Bunch Counter Reset signal (BCR) used to reset the Bunch Counter Identifier (BCID) in each readout element at a fixed time within the LHC cycle. The pre-defined LHC bunch structure will repeat itself every turn (every 88.924 μ s, which is equivalent to 3557 bunch crossings).

Processing prompt data from calorimeters and muon chambers, the CTP delivers a L1A signal to trigger the readout of events which fulfill a set of basic criteria, those which are most likely to contain interesting physics. Counting these L1A signals gives the Event Identifier (EVTID). This counter is synchronously reset by the Event Counter Reset (ECR) signal. Finally for each trigger, the CTP provides an 8-bit word to identify the type of trigger (physics, calibration, etc.). In addition, for testing and calibration purposes, a prepulse can be issued a pre-defined number of clock cycles before a calibration trigger is generated.

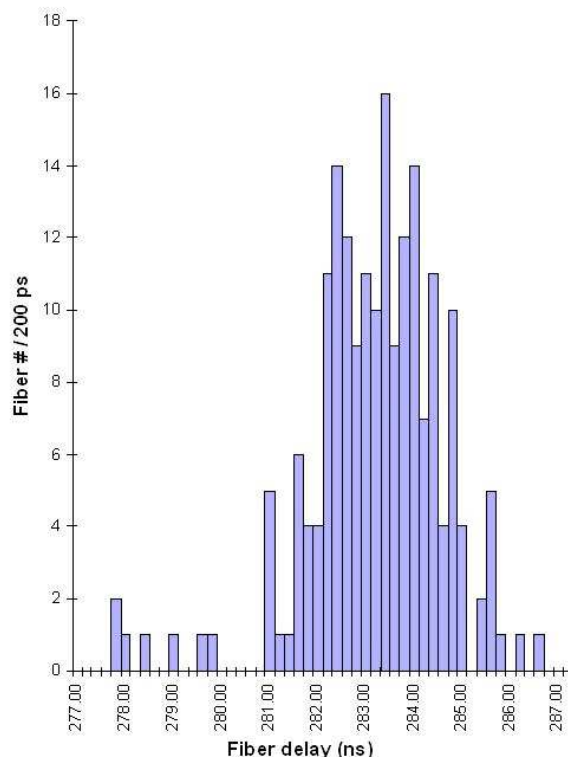


Figure 7. Measured time delays (in ns) of signal travelling in the TTC fibres which link the TTC rack to the FECs. The precision of each measurement is better than 100 ps. The r.m.s. of the distribution is 1.4 ns.

The trigger latency, which is the delay between the bunch crossing time and the time when the L1A signal arrives to the FE electronics, has been minimized to a value below $2.5 \mu\text{s}$. The TTC rack location has been chosen to minimize the length of the TTC fibres to the FECs and the associated contribution to the trigger latency. Figure 7 shows the measured delays in these fibres where an r.m.s. of 1.4 ns is observed. The precision of each measurement is better than 100 ps. The spread will be taken into account by programming a delay in a TTCrx [5] chip which receives the TTC signal in each of the FEBs. In addition, the programmable delay lines of the calibration boards are preset to reproduce the timing of signals generated by particles originating from the interaction point.

2.3.4. Level-1 receiver system The receiver system is part of the trigger sum chain and interfaces the TBBs/TDBs to the calorimeter Level-1 processor. One function of the trigger sum chain is to convert the signal from energy to transverse energy. The final gain adjustment is left to the receiver, which must also account for attenuation in the cable between the TBB/TDB and the USA15 cavern. Because of the need for a continuously variable gain over a relatively small range, a stage of programmable gain is included in the receiver module. Level-1 requires the trigger sums to be grouped in

$\Delta\eta \times \Delta\phi$ bins and this is done via a daughter remapping board located on each receiver module. Finally the system can select any raw trigger sum signals for diagnostic tests using special monitoring modules located in the receiver crates.

The receiver chain consists of four op-amps: the first two are located on a 16-channel variable gain amplifier daughterboard, the third performs an RC integration and the fourth is a driver circuit which sends signals to the Level-1 calorimeter processor. Each receiver module services 64 channels, along with circuitry to select channels for monitoring purposes. A detailed description of the functionality of the receiver system is given in Reference [6].

Each receiver crate contains 16 receiver modules, two monitoring modules, and one controller module. The crate is a standard 9U with a custom backplane, which is used to transport both the digital signals between the controller and the other modules, as well as the analog monitoring signals. The LAr system consists of 6 receiver crates: 2 for the EMB, 2 for the EMECs and 2 for the FCAL and HEC. Two additional crates are for the Tile calorimeters.

The TBB receiver system cabling is described in detail in Reference [7]. There are two receiver crates per rack, such that the entire receiver system occupies four racks.

3. LArg Back End electronics description

In the following sections, the ROD crate and its components are described in detail.

3.1. VME crate

The ROD crates are 21 slot Wiener 9U VME64x crates (Type UEV 6021-LHC9U-VME) for 9U VME cards with a remote water-cooled power supply [8]. Each crate measures 532.6 mm in height, 431.8 mm in width and is 720.0 mm deep. The crates are mounted in standard 52U LHC racks and contain: 1 CPU, between 6 and 14 RODs, 1 TBM, 2 SPAC Masters, a custom made CP3 backplane, and 6 to 14 TMs, according to the number of RODs.

The VME CPU, a Concurrent Technologies VP110 board [9], is installed in the first of the four 6U slots of the ROD crate. During system tests, the timing and trigger modules TTCvi and TTCex can be connected in the other 6U slots. Figure 8 illustrates a schematic top view of the ROD crate, where the locations of the various components are explicitly shown.

3.1.1. Power supply The remote power supply unit (PSU) is mounted on the back of the rack and is water cooled. A drop-less coupling system ensures a leak-less connection even if the system is working at over pressure (5-6 bar). The power is supplied to the crate via a low inductance 50 cm long cable, to achieve low noise conditions for the fast-switching electronic devices [10].




Figure 8. Schematic top view of the ROD crate showing the locations of the various crate components.

The crate is supplied with the following voltages: +5V, +12V, -12V, +3.3V and +48V. The main consumers are the RODs and TMs and if there are 14 of each in a crate, currents of 73A at 5V, 50A at 3.3V and 7.5A at 48V are reached. The total power delivered by the PSU is about 900W.

The temperature of the VME power supplies are monitored with probes mounted on the power modules. Over-heating of the PSU can be caused by an interruption of the cooling water system in the ATLAS cavern, in which case the crates are shut down.

3.1.2. Air and water cooling of the electronics The electronic boards installed in the VME crates are cooled by air which flows in a closed circuit inside the rack. The air is cooled by heat exchangers installed above each ROD crate. Temperature probes mounted inside the crates above the RODs measure the air temperature. In case of over-heating, the speed of the cooling fans below the boards is increased.

The ROD G-link HDMP-1024 chips require special water cooling to operate reliably at 80 MHz; the chip has to be kept below 35°C [11]. A cooling plate is mounted on the 8 G-link chips, as shown in figure 9. Flexible water pipes are connected on the front-plate of the ROD with a drop-less coupling system. The water is distributed to each board by a 1U distribution unit.

To prevent water leaking into the electronics system, the cooling circuit is operated at a pressure which is between 300 and 400 mbar lower than the air pressure in the counting room. The cooling unit uses demineralized water, cooled down by the primary water system installed in the ATLAS cavern. The power to be dissipated from the G-links of all 192 ROD boards is about 2300 W, well below the maximum

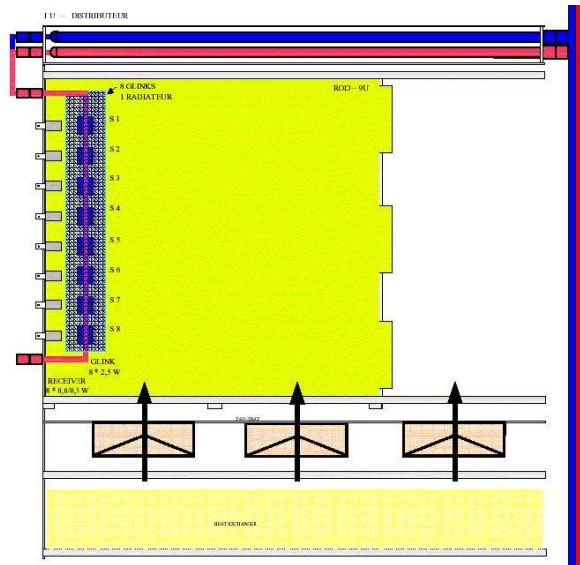


Figure 9. A side view of a ROD board inserted into a crate, where the water and air cooling can be seen. Air fans are shown below the boards, whereas the water cooling plate for the G-links and the associated pipes are shown on the left-hand side of the board. One can also see the P1, P2 and P3 connectors on the right-hand side of the board, respectively from top to bottom.

cooling power of 3860 W of the cooling unit. The plant is operated at a secondary circuit water temperature of $+20^{\circ}\text{C}$ and at a flow rate of $1.1\text{m}^3/\text{h}$ with a maximum differential pressure of 2 bar [12]. One main distribution line connects the seven ROD racks in parallel, and serves the 17 water distributors mounted above the ROD and the Injector crates. The G-link chip temperatures are approximately 28°C in a stable ROD processing mode.

3.1.3. Detector Control System (DCS) In addition to hardware control of the VME crates and the ROD cooling unit, the status of the BE electronics is continuously supervised by a software-based detector control system (DCS) [13]. The DCS checks in given time intervals the voltages, currents and temperatures measured inside the VME crate and its power supply. It also receives temperature, pressure and flow information from the ROD cooling station. The DCS software transmits warnings and alarms issued by the VME crate and cooling unit hardware controllers. It can also raise warnings and alarms by itself, if certain parameters exceed pre-defined limits. One example is the monitoring of the temperature on the VME power supply, which should not exceed a certain limit to prevent a hazardous automatic shut-down of the VME power. This temperature alarm may be raised e.g. in case of a failure of the primary cooling system.

The information about the BE system is further transmitted to DCS components which supervise the status of the complete ATLAS calorimeter hardware and eventually the status of the whole ATLAS detector.




Figure 10. Front side of the CP3 board and point-to-point connections between the TBM and the ROD modules.

3.1.4. Crate backplane and CP3 board The VME crate is equipped with the standard VME64 backplane supporting the 160-pin P1 and P2 connectors. This VME backplane is extended at the bottom by a custom made backplane, the CP3, supporting additional P3 connectors of the same type. The TM modules (see section 3.6) are situated at the rear of the crate while all the other electronic modules are inserted at the front.

The CP3 card is a passive printed circuit that links the TBM (see section 3.3) to the RODs (see section 3.4). It ensures the distribution of the TTC signals and the collection of the ROD Busy signals over point-to-point connections between the TBM slot number 5, and the ROD slots 6 to 21. It is capable of carrying the high speed digital TTC signals with minimal distortion, reflection, or cross talk. This requirement has been satisfied by the following implementations:

- Fast and low jitter Low Voltage Differential Signal (LVDS) chipset.
- A printed circuit with differential signal paths of $100 \Omega \pm 10\%$ characteristic impedance. All the paths have a controlled impedance of 50Ω . The two lines of a differential pair are routed on the same layer, are separated by the same distance all along, and have the same length.

For simplicity, and because they are less time critical, the ROD Busy signals run on 50Ω single ended transmission lines. Figure 10 shows the front side of the CP3 and the point-to-point connections.

Each CP3 has a unique identifier defined by two rotary switches. This identifier is readable by the TBM from its P3 connector. In addition, the CP3 includes a few service lines that allow to read the TM identifiers through the TBM. In this manner, the TMs situated at the rear of the crates, as well as the ROD crates with their associated CP3 cards, can be identified by VME through the TBM, simplifying maintenance.

The CP3 also distributes the electrical power to the TMs. The floating 48 V voltage




Figure 11. Rear side of the CP3 board and electrical connections.

used by the TMs is brought from the VME backplane by means of two cables attached to the two upper taps situated on the CP3. The third tap, at the bottom, is used to connect the rack and CP3 groundings. The CP3 is fastened onto the crate using 20 electrically grounded screws. Figure 11 shows the rear side of the CP3 and the electrical connections.

The CP3 measures 424.23 mm in width, 128.78 mm in height and is 4.64 mm thick. The VME64 connectors and the power taps have been inserted by force into the thick printed circuit (no soldering). The CP3 contains 3 signal layers and 5 equi-potential planes (see Reference [14] for a detailed description of the CP3).

3.2. SPAC Master module

The SPAC Master 9U VME module is dedicated to the configuration of the FE electronics. This VME module is situated in the ROD crates but is independent of the other crate boards, which are devoted to the data acquisition and are intensely active during physics or calibration runs. The SPAC Master module is mainly solicited during the configuration phase, to properly set up the various FE system boards (FEBS, Calibration boards, TBBs, FECcont boards).

3.2.1. Description and functionalities of the module The SPAC Master Module is optically connected to a FECcont board that redirects the requested READ/WRITE requests to the various boards populating the FECs. The data are then transmitted from their host board to the SPAC Master board through the FECcont that handles the electrical to optical conversion. Figure 12 shows this architecture. In the FEC, signals are transmitted to the FEBS using a dedicated copper bus. The MS1,2,SM1,2 (Master-Slave line, Slave-Master line) fibres denoted on the figure are the four optical connections shared by a SPAC Master board and a FECcont.

The data transmission between the SPAC Master board and a FECcont is done through a serial link using a protocol developed specifically for the ATLAS calorimeters (see Reference [15] for more details). The data are transmitted at a nominal rate of approximately 10 Mbits/s.

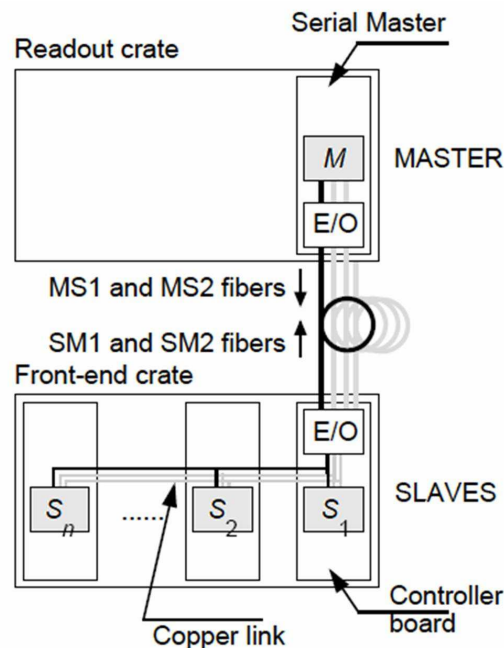


Figure 12. Functional layout of the SPAC configuration link showing the MS1,2,SM1,2 (Master-Slave, Slave-Master) optical connections shared by a SPAC Master board and a FECcont. S_1, \dots, S_n are the SPAC Slave circuits which equip each FEB to implement the SPAC link. The E/O indicates where the electrical to optical conversions are performed.

The SPAC Master - FECcont data exchange is a key issue of the whole system, which cannot be configured without it. Hence the reliability is increased by doubling the SPAC links between a SPAC Master and a FECcont. Consequently, the FECcont implements 2 Slave circuits in a special configuration allowing data transmission even if one of them fails. Nevertheless, only one pair of SM-MS links is used according to a configuration bit in the SPAC Master Control word and in each FECcont Slave ASIC. In addition the SPAC Slave has been designed in such a way that if it fails, it should still be able to transmit the signals to the SPAC copper bus delivering electrical SPAC signals to the FEBs.

The SPAC Master module is installed in a single slot of the VME crate. Each ROD crate contains 2 such modules. Each board is equipped with 4 independent SPAC links to a FECcont. It has 8 optical transmitters and 8 optical receivers: 4 transmitters or receivers are used, while the 4 others are spares for reliability purposes. Each optical connection is unidirectional. A SPAC Master board is shown in figure 13 and its functional layout in figure 14 showing that 2 Master interfaces are implemented in a single Altera ACEX EP1K50 Field Programmable Gate Array (FPGA).

3.2.2. SPAC protocol and Slave ASIC The SPAC protocol is a Master/Slave based protocol implementing a serial connection between one master and possibly multiple slaves (see Reference [15] for a detailed description). Each FEC board requiring an



Figure 13. Photograph of a SPAC Master board equipped with 4 SPAC links.

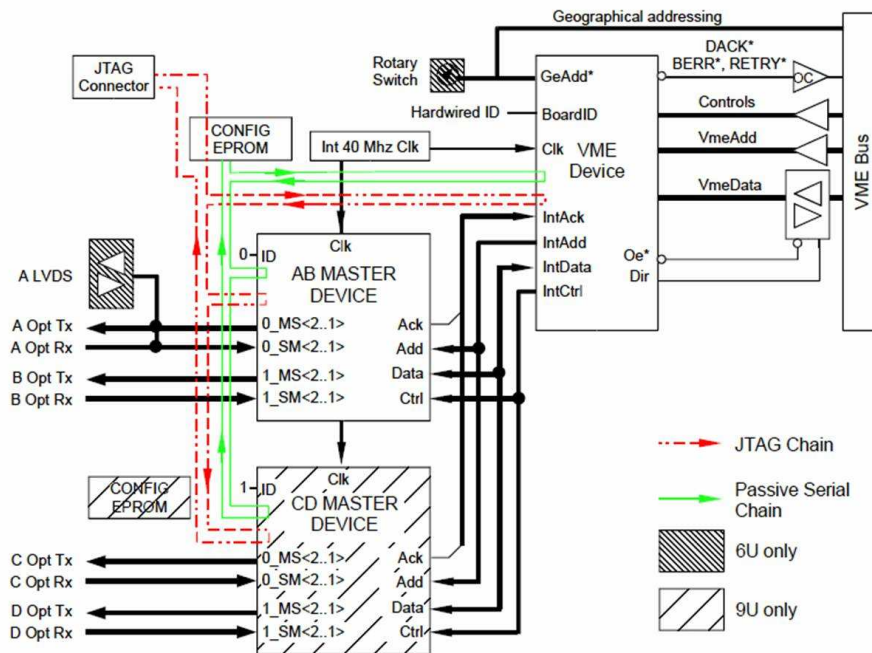


Figure 14. Functional layout of the SPAC Master board. Two Master interfaces are implemented in a single Altera ACEX EP1K50 Field Programmable Gate Array (FPGA).




Figure 15. Data format of a SPAC transaction between SPAC Master and SPAC Slave.

online configuration is equipped with a dedicated on-board SPAC Slave interface ASIC which communicates to the SPAC Master board through the FECcont board. Whereas this chip is used in a standard slave mode on all the FEBs, it is used on the FECcont with a different configuration (repeater mode) to allow SPAC frames transmission from/to the FECcont slave ASIC to/from the various FEBs slave ASICs.

The data transmission through the SPAC links is done using a Manchester encoding scheme [16] which increases the noise immunity of the transmission. The convention used is the following:

- a logical 1 is coded by a low to high transition (i.e. 50ns low, 50ns high at 10 Mbits/s);
- a logical 0 is coded by a high to low transition (i.e. 50ns high, 50ns low at 10 Mbits/s).

The system is formed of optical links operating at different wavelengths in order to benefit from existing qualified radiation hard components. The transmission through the MS links is done using 850nm transmitters on the SPAC Master whereas the SM links are implemented using 1310 nm transmitters on the FECcont. The SPAC Master module emitters and receivers are off-the-shelf components.

The data are transmitted in 9-bit words composed of a byte of useful data and a Continue/Last bit (logical 0 for the value Last). The data transmission format is described in figure 15. A SPAC frame consists of the following:

- A preamble string sent at the beginning of each data frame to restore a correct electrical level on the differential electrical or optical receivers on the receiving side (see STR1-STR2 in figure 15) which leads to very good transmission.
- The address of the SPAC Slave to which the SPAC frame is issued (ADD).
- The SPAC address of the on-board resources to which the SPAC frame is issued (SAD).
- A first checksum (CHK1) word which allows detection of badly transmitted addresses.
- The data words to be transmitted to the recipient (DAT1-n).
- A second checksum (CHK2) which tests the data integrity and can correct badly transmitted data with 1 or 2 bits flipped during transmission.

If a slave does not answer to a READ request, a configurable software timeout is generated by the online software. If a slave receives a bad frame (either a bad checksum




Figure 16. Transmission error rate in a SPAC link as a function of the SPAC Master FPGA frequency when the connected SPAC Slave ASIC is clocked with a 40 MHz signal. This shows that the system can easily tolerate a 2 MHz frequency misalignment.

or a frame inconsistency), it emits a SPAC interrupt to the Master and a flag is set in its internal Slave Status Register. The SPAC Master which receives the interrupt sets a flag in its status register.

The clocks on the SPAC Master board and on the Slave ASICs can be asynchronous and even have a small discrepancy in frequency. Figure 16 shows the error rate as a function of the frequency of the SPAC Master FPGA when the connected SPAC Slave ASIC is clocked with a 40 MHz signal. This shows that the system easily tolerates a 2 MHz frequency misalignment. The actual system will use very similar frequencies at both ends: 40 MHz from a crystal on the SPAC Master Board and the LHC 40.08 MHz clock on the FEBs.

3.2.3. SPAC interface The SPAC Master allows to use all the possible interfaces implemented on the SPAC Slave ASIC, which can access on-board resources (memory or registers [17]) through either a parallel interface or two Inter-Integrated-Circuit (I2C) links. In the first case, the command is directly interpreted and executed on the slave parallel interface. In the second case, the I2C command is built and put in a dedicated ASIC memory of the SPAC Slave; the SPAC Master then orders the SPAC Slave to issue this I2C command. Hence an I2C command requires at least one or more extra SPAC Master VME transactions with respect to writing/reading through the parallel interface. The SPAC Master Board is equipped with emission/reception First-In-First-Out's (FIFOs) which are used either to build the SPAC commands or to read back the expected information using VME transactions to the SPAC Master module.

3.2.4. Power consumption, power and reset procedures The SPAC Master module has a power consumption of 15 W when SPAC transactions occur at the maximum SPAC link bandwidth. The electrical consumption is detailed in table 2. Four fuse carriers are mounted on the board and host two spare fuses of each type (5 A and 500 mA fuses).

At initialization, the VME interface reads the geographical address of the slot where

Table 2. Power consumption of the SPAC module.

Power supply	Consumption	Fuse value
5V	<3A	5A
3.3V	250mA	500mA

the board sits and uses it to define the base address of the module. FIFOs are emptied and optical transmitters are idle. The RESET procedure has the same characteristics as the POWER-ON except that the VME interface is not reinitialized.

3.2.5. SPAC design and manufacturing The SPAC Master module has 2 male right angle VME64x connectors. The Printed Circuit Board (PCB) measures 400 mm in height, 366 mm in depth and is 1.6 mm thick. The board is composed of eight layers: four signal layers and four ground and power layers. Thirty five modules were produced and tested: 31 used by the experiment and 4 spares. The full production of the boards was performed by a unique vendor and included: PCB manufacturing, sub-contracting, procurement of the components, assembly, burn-in sequences, flying probe tests, front panel production and mounting. Functional tests were then performed on reception of all the boards: the four SPAC links of each module were connected to FECcont boards to check for errorless communication in any optical fibre configuration, and the optical power of the transmitters on the SPAC master modules was also measured and stored into a production database.

3.3. TTC and Busy module (TBM)

The TBM module distributes the TTC signals and collects the Busy signals inside each ROD crate (see Reference [18] for more details), where it sits in slot number 5. Figure 17 shows a photograph of a module.

3.3.1. Timing, Trigger and Control distribution The TBM receives signals from the TTC system via an optical fibre, with a TRR-1B43 optical receiver from TrueLight. The receiver converts the TTC optical signal into an electrical differential signal (LVDS). To set with a minimum amount of distortion a known inactive state on the TTC lines, a slight bias is introduced between the positive and the negative inputs of the first electrical buffer. The TBM then duplicates the LVDS TTC signal and distributes it to each ROD module through the P3 connector, using point-to-point connections on the CP3 board.

At the receiving end on a ROD module, a pair of lines of 100 Ω differential characteristic impedance is ended by a 100 Ω resistor. If the slot is empty, the corresponding TTC signal is absent, due to the absence of the terminating resistor which is necessary to close the driver current loop.




Figure 17. Photograph of the TBM module.

While the TTC signal is electrically transmitted to the RODs, a copy of the TTC signal enters a TTCM mezzanine card [18] situated on the TBM, where a TTCrx chip decodes it. If the TTCrx ready-flag is false i.e. missing TTC signal, an interrupt is generated to warn the VME crate controller. The state of this flag is also visible on the TBM front panel. The TBM can be programmed to count the Level-1 triggers and to generate an interrupt after a pre-defined number of triggers. During calibration runs, the Busy signal is used in conjunction with this trigger count, in order to synchronize the data taking of all the sub-detector channels.

3.3.2. Busy collection The TBM, which receives the ROD Busy signals from the P3 connector, enables or disables each of them according to a dedicated mask register, then makes a logical OR out of them to create the Crate-Busy signal. This signal is sent through a coaxial cable to the ROD-Busy-Module, which is situated in the TTC crate [19]. This module has 16 coaxial inputs and outputs their logical OR. The TBM also implements control and monitoring functions on the input signals.

The logic associated in the TBM with the Busy signals is contained in an FPGA, in conjunction with another FPGA dedicated to the VME interface (both FPGAs are ACEX EP1K50QC208-2 from Altera). The TBM can be programmed to generate a VME interrupt when the first ROD Busy signal arrives, as well as when all the ROD Busy signals disappear. A schematic view of Busy signal logic is shown in figure 18.

The Low Voltage Transistor Transistor Logic (LVTTL) Busy signals are active in the high logic state. Pull-down resistors (10 k Ω) at the input of the TBM cause the empty slot lines to be inactive. The Busy lines on the CP3 board are long and consequently sensitive to reflections, and hence have a controlled 50 Ω characteristic impedance. Reflections are avoided by using a small resistor placed in series at the ROD side, thus absorbing the reflected wave.

After they have entered the TBM, the ROD Busy signals are individually masked

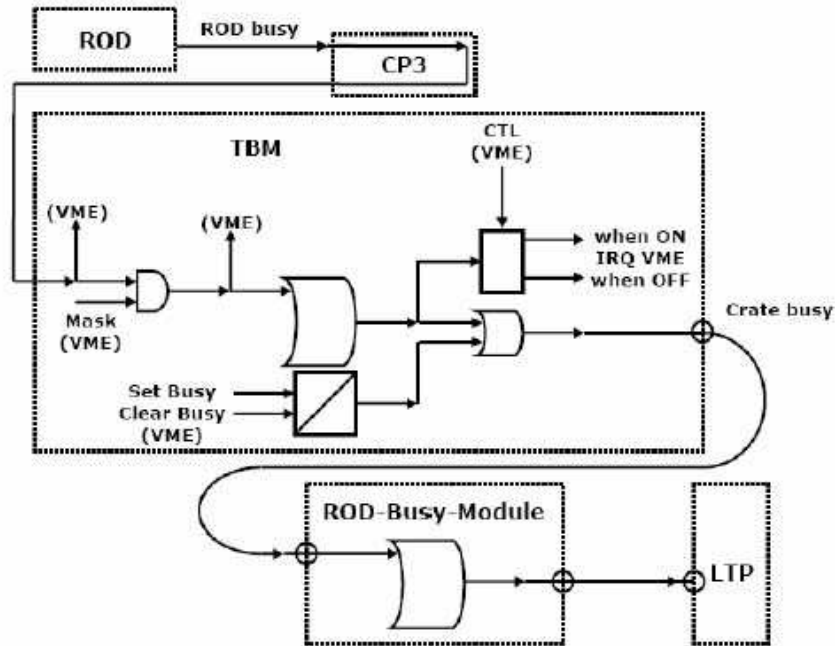


Figure 18. The ROD Busy signals are collected by the TBMs and the ROD-Busy-Modules before being transmitted to the Trigger Processor.

or enabled. They can be read as they are at their inputs on the P3 connector or after the mask. In order to prevent unwanted triggers, the Crate-Busy signal is set ON at power-up or when the TBM is reset.

3.3.3. Control outputs The front panel Crate-Busy signal is available from a coaxial LEMO type connector. A second identical LEMO output is available for scope monitoring and a third one is under the control of Set and Clear VME commands. The first LEMO cable is used to connect the TBM to one of the 16 inputs of the ROD-Busy-Module. The 3 LEMO Control outputs are driven by 3 separate outputs of the Busy FPGA with 50 Ω resistors in series to prevent reflections.

3.3.4. VME interface The TBM is connected to the VME-bus and has four interrupt sources: an end of calibration sequence (when a pre-defined number of Level-1 triggers has been reached), the first enabled ROD Busy ON, the last enabled ROD Busy OFF, and a lost TTC clock. The core of the VME interface is common to all LArg ROD boards. The registers that are specific to the TBM are implemented in its Busy FPGA. The low part of the VME address is transmitted by the VME FPGA to the Busy FPGA, allowing the commands that are dedicated to the TBM to be interpreted and executed in the Busy FPGA.

3.3.5. Power, reset and ESD considerations The TBM current consumption is 0.5 A on the +5 V and 0.6 A on the +3.3 V. Re-triggerable variable resistors are used as fuses

to protect the crate backplanes from hazardous currents that could occur in the TBM. When the power is set on, the on-board logic is reset and the FPGAs automatically configure themselves in less than 1 second. To protect against hazardous electrostatic discharges (ESD), two copper strips establish a connection to ground through 2 M Ω resistors when the board is inserted into the crate.

3.3.6. TBM design and manufacturing The TBM module has 3 male right angle VME64x connectors. It measures 400.05 mm in height, 360 mm in depth and is 2mm thick. The top and bottom parts which slide into the VME crate rails are only 1.6 mm thick.

Many paths between integrated circuits on the board are longer than ten centimeters, and care has been taken to avoid parasitic reflections. The PCB of the TBM contains 3 signal layers and 5 equi-potential planes, and the same care has been taken for the routing of signals as for the CP3 (see section 3.1.4).

For the long unipolar paths, either a parallel or a series impedance matching has been implemented. The parallel matching is made of a Thévenin network placed at the end. The series matching is made of a 50 Ω resistor placed at the emitter side. These matchings are used between the VME FPGA and the Busy FPGA. It was verified that pulses as short as 15ns are correctly transmitted, while no matching is necessary for pulses beyond 100ns.

Due to its large dimensions, the board is subject to mechanical distortions. When it is plugged into the VME crate, the board is maintained horizontally straight by the top and bottom rails of the cage. To avoid vertical distortions, a metal bar is screwed onto the board, near the VME connectors.

3.4. ROD module

The Read Out Driver is at the core of the off-detector data processing. Each ROD module receives the digitized raw data from 8 FEBs through 8 optical fibres. This corresponds to 1024 (8×128 channels/FEB) detector cells being read out. The module is composed of a ROD motherboard, equipped with 4 Processing Units (PUs) placed on separate daughterboards, each housing 2 Digital Signal Processors (DSPs). Highly performant plug-in DSPs are chosen to achieve an average processing time per event of less than $\sim 12 \mu\text{s}$, as required by the 75 kHz Level-1 trigger rate.

The ROD processes the data and sends the results of the DSP optimal filtering calculations (see section 2.3.1) to the TM installed at the back of the VME crate. When the ROD cannot output data towards the TM, it sends a Busy signal which paces the Level-1 trigger. For monitoring purposes, output data can also be sent to a Synchronous Dynamic Random Access Memory (SDRAM), where they are read out through the VME interface.

The LArg detectors are read out using 192 ROD modules. A photograph of a ROD board is shown in figure 19. RODs are also used for the ATLAS Tile calorimeter




Figure 19. Photograph of a ROD motherboard equipped with four processing units. The 8 optical receivers (see section 3.4.1) and the G-link cooling plate are clearly visible on the left-hand side of the board.

readout.

3.4.1. FEB data input Data from one FEB are sent to a ROD through the 1.6 Gbit/s G-link optical link which encodes data from the 16 FE ADCs (128 calorimeter channels).

As shown in figure 20, the ADC event fragment format consists of the following 16-bit words: a frame-start word (all ones), two event header words, data words carrying the calorimeter signal data, an event trailer word, and at least one frame-end word to separate the next event (all zeros). For all words except the start and end of event, bit 14 is used as a parity bit to impose odd parity. More details about the FEB output data format are given in Reference [20].

Data from each ADC are encoded on the FEB on a two-bit 40.08 MHz serial data bus. The even (odd) bits are transmitted on the 0 (1) output bit line with the most significant bit first. The resulting 32 bits at 40.08 MHz corresponding to the 16 ADCs are multiplexed in a custom made multiplexer component, the SMUX [21]. It was developed according to a 2:1 multiplexing in order to generate a stream of 16 bits at 80.16 MHz at the output of the FEB. In addition, the SMUX generates a flag which is set to 1 (0) for the multiplexer cycle when data from channels 0-63 (64-127) are sent. These 16 bits are then serialized by the HDMP 1022 G-link chip on the FEB.

3.4.1.1. Optical receiver and link Serial data from the FEBs are received by the ROD motherboard through eight optical receivers, the ORx. These are custom made devices and provide the optical to electrical conversion of the data. To reach the desired sensitivity (-17 dBm), special care was taken for the packaging and electrical

1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0	P	0	0	ADCID				PHASE			EVENTN					
0	P	0	0	BCID												
0	P	0	0	F	L	B	A	CELLN								
0	P	Gain		ADC sample												
0	P	Gain		ADC sample												
0	P	Gain		ADC sample												
0	P	Gain		ADC sample												
0	P	Gain		ADC sample												
0	P	Gain		ADC sample												
0	P	Gain		ADC sample												
0	P	Gain		ADC sample												
0	P	0	0	1	S	E	SCAController status							1		
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Figure 20. Event format at the input of the PU.

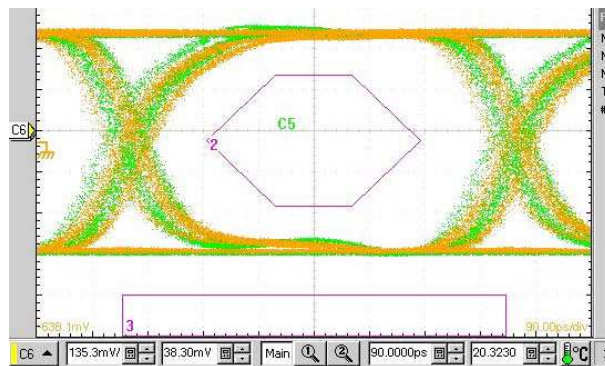


Figure 21. Eye pattern diagram of an ORx chip. Entries inside the marked fiducial region provide a measure of the bit error rate, while the widths of the rise and fall lines represent effects of the timing jitter.

filtering. The receivers can be seen in figure 19 and details of the ORx chips are given in Reference [22], figure 21 shows the ORx eye pattern diagram, illustrating the high quality of the 1.6 Gbit/s transmission.

The incoming data are deserialized by Agilent HDMP-1024 G-link chips and transmitted in 16-bit words to the receiver FPGAs. The HDMP-1024 chip allows to reconstruct parallel FEB data as they were loaded into the HDMP-1022 transmitter chip located on the FEB. The chip set hides all the complexity of encoding, multiplexing, clock extraction, demultiplexing and decoding. An external quartz provides an 80.16 MHz clock reference to the G-link chip, allowing the component to lock on the incoming data.

The G-link indicates the status of the optical link to the receiver FPGA by a link-

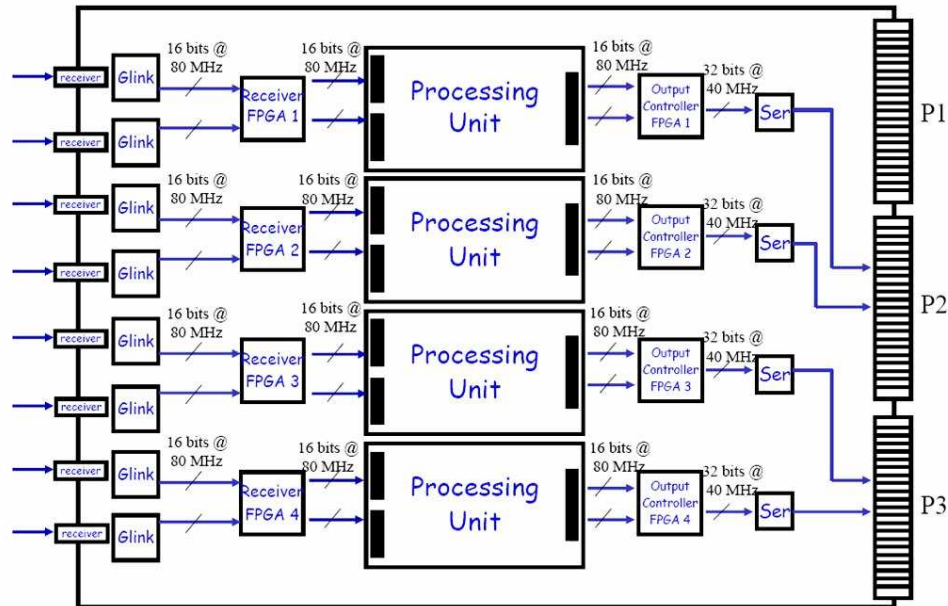


Figure 22. Data distribution on the ROD board in normal mode, using all PUs and S-links.

ready bit. If the G-link loses the frame lock in the incoming data stream, it is set to 0 (false) and the G-link output data are composed of words with all bits set to 1.

3.4.1.2. Receiver FPGA Four Altera EP1K100 receiver FPGAs are used on the ROD motherboard. The incoming data of two different G-link chips are sent to one receiver FPGA, which routes the data further to the PUs. In normal working mode, the data of one receiver FPGA are transferred to one PU. This is shown in figure 22.

In the so called staging mode, a PU receives the data from two FPGAs which are linked by a special data bus. The staging mode, with half the processors installed, allows for a ROD system with a reduced number of input links to the DAQ. This mode is used by the Tile calorimeter readout which has a less demanding processing power, since there are only 32 to 45 readout channels per FEB as compared to the 128 channels for the LArg calorimeters.

Data from each receiver FPGA are sent to the PU through a data bus, running at 80.16 MHz. FEB data are synchronized in the receiver FPGA using an 80.16 MHz clock made by an on-board Phase Lock Loop (PLL) chip and the 40.08 MHz TTC clock signal. In this way, the ROD data are synchronized with the LHC clock and not with the external G-link reference quartz clock.

In total, 18 bits are transmitted to the PU per FEB: 16 bits of data, the flag signal to distinguish data from channels 0-63 or 64-127, and the link-ready signal indicating if the G-link has lost the frame lock. In this case, the receiver FPGA does not transfer the incoming data to the PU.

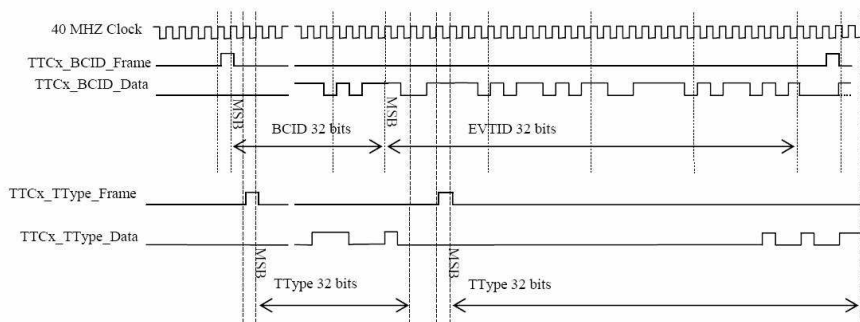


Figure 23. Schematic view of the communication protocol between the TTC FPGA and the PUs.

In order to test the ROD board in stand-alone mode i.e. without optical data injection, an additional functionality has been added in the receiver FPGA. The embedded memory can be used to store a maximum of 1024 32-bit words (corresponding to one 5 sample FEB event). The simulated event can be sent either a finite number of times or in a continuous loop to the PU, with the possibility to adjust the delay between events. The same event is duplicated and sent to both inputs of the PU, to simulate the arrival of data from two FEBs.

3.4.2. TTC data input The TTC signal from the CP3 is received by the TTCrx chip which decodes it and sends the resulting signals to the TTC FPGA. The FPGA receives the Trigger Type information and the BC, BCR and ECR signals. It calculates the BCID and EVTID and then distributes all information to the PU through point-to-point connections. Figure 23 shows the communication protocol between the TTC FPGA and the PUs.

The TTC FPGA sends the BCID, the EVTID and the Trigger Type using 5 lines: one 40 MHz synchronization clock, two frame lines which define the beginning of an element transfer, and two serial lines for the data transmission. One frame line, with 2 phases of 32 bits, and one data line are used to transmit the 12-bit long BCID and the 32-bit long EVTID. The other frame line, with only one 32-bit phase, and data line are used to transmit the 8-bit Trigger Type.

Depending on the mode of operation, the clock is generated differently. During LHC running, the ROD board will receive the 40.08 MHz LHC clock from the CP3 backplane. In local or VME mode, the clock comes either from a local clock oscillator or from the Injector module (see section 3.7). The clock is distributed to each sub-module through point-to-point connections.

3.4.3. Output flow The DSP 16-bit output data are stored in two FIFOs which are read by an Output Controller (OC). There are four OCs on a ROD module, one per PU. The OC reads the event fragments from the two DSP FIFOs and according to its




Figure 24. Functional diagram of the Output Controller interfaces.

configuration, outputs the data to the SDRAM and/or to the serializer chip. Figure 24 shows a functional diagram of the OC interfaces. The OC formats the event, including a header and a trailer according to the specifications given by the ATLAS DAQ group [23]. Table 3 shows the ROD event fragment format.

The OC reads the DSP FIFOs (16 bit wide) at 80 MHz and outputs 32-bit wide data at 40 MHz. FIFO data are stored according to the big-endian protocol, most significant word first. The protocol between the FIFOs and the OC is not based on the notion of event, but on data bunch, which can include a partial event or several events. The transfer from the FIFO is initiated by the OC as soon as the empty flag is false i.e. when the FIFO is no longer empty. The OC reads a 16-bit wide FIFO word at each rising edge of the 80 MHz clock, such that two clock periods are needed to obtain a 32-bit pattern.

The first 32-bit pattern contains the data bunch size, which can include a partial event or several events. The OC stores this value and decrements it at each new 32-bit word transferred. When the OC reaches the last word, it gives a 50ns pulse on the Event End line, indicating to the PU that the whole data bunch is transferred. The first 32-bit pattern of an event contains its size: the number of 32-bit words in the current event, excluding itself but including the last word. The last word of the event must be 0xe0e000000. The OC verifies this and in case of an error, will set a flag in the status register and stop event transfers.

The OC always starts reading the FIFO associated to the first DSP of a PU (FIFO1), and then reads the second FIFO2. The event notion is kept: if the data bunch contains several events, they are read in one at a time, to let the OC format the ROD event fragment. A complete OC description is given in Reference [24].

3.4.4. SDRAM The OC outputs the data to the SDRAM, and/or to the TM to send to the DAQ, or to both in the so called spy mode. The SDRAM mode allows to read data through the VME bus and is mainly used for debugging purposes e.g. when the DAQ system is not available or for data consistency checks in the local CPU. The spy mode is mainly used during data acquisition, to check the quality of the incoming FEB data, without disturbing the DAQ data flow. In this mode, the data flow is slowed down a little, since the SDRAM must be refreshed by the OC and synchronized with the serializer.

The spy mode is activated through VME, which also defines the number of words

Table 3. ROD event fragment format at the output of an output controller. The data words marked in *italic* are generated by the DSP but not sent in the event fragment.

		Beginning of fragment (0xb0f00000)
Header	OC	Start of header marker (0xee1234ee)
		Header size (0x00000009)
		<i>Number of Words of DSP1</i>
	DSP1	Format version (0x02040000)
		Source identifier
		Run Number
		Level 1 ID
		Bunch crossing ID
		Level 1 trigger type
		Detector event type
		Header event DSP1
Data and		Data
status		:
elements		:
		<i>End of event (0xe0e00000)</i>
		<i>Number of Words of DSP2</i>
	DSP2	Format version (0x02040000)
		Source identifier
		Run Number
		Level 1 ID
		Bunch crossing ID
		Level 1 trigger type
		Detector event type
		Header event DSP2
		Data
		:
		:
		<i>End of event (0xe0e00000)</i>
		Number of status elements (0x00000000)
Trailer	OC	Number of data elements
		(words DSP1 + words DSP2-9)
		Status block position (0x00000000)
		End of fragment (0xe0f00000)

stored in the SDRAM. When the spy mode is enabled, the events are sent to the DAQ through the TM, and they are also stored in the SDRAM. When the SDRAM threshold is reached, the OC finishes copying the current event to the SDRAM and stops sending it data. From then on, the OC sends the events only to the DAQ and the SDRAM data can be read via VME without disturbing the TM event transfer. Once the SDRAM is emptied, the spy mode can be re-enabled.

The SDRAM is a 128 Mbit high speed DRAM. It is internally configured as a quad-bank DRAM with synchronous interface. Each bank is organized in 4096 rows times

256 columns times 32 bits. Read and write accesses to the SDRAM are burst oriented: accesses start at a selected address and continue for a programmed number of locations. The SDRAM is designed to operate at 3.3V and requires 4096 auto refresh cycles every 64 ms (at least one command every 15625 μ s).

The SDRAM status and control flags (enable, empty flag, full flag, etc.), as well as its READs, WRITEs and refreshing are handled by the OC.

3.4.5. Serializer The output data are sent to the TM through the VME backplane (P2 and P3 connectors). Each ROD to TM channel uses one DS90CR483/484 serializer-deserializer chipset and is implemented in differential LVDS. This mechanism reduces the noise associated with the transmission. The chipset converts 48 bits of TTL data into 8 LVDS data streams running at 240 MHz. The output of one OC consists in a 48-bit bus: 32 bits of data, 6 bits of control signals generated for the S-link, and 10 spare bits. The TM receives, deserializes and sends the data to the corresponding S-link mezzanine board.

A total of 14 LVDS pairs (28 pins) are required per OC in the P2 and P3 connectors: 8 pairs for data transmission to the TM, 5 pairs for return information and one pair for the ROD 40 MHz clock signal which is transmitted in parallel with the data streams. The 5 return bits from the S-link are the following: the Link Full flag indicating the Read Out System (ROS) is busy, the Link Down flag indicating the communication with the DAQ is no longer ensured, and the 3 Return Lines LR0, 1 and 2 from the ROS. These 5 signals are converted from LVDS to LVTTTL and sent to the corresponding OC FPGA for analysis.

A complete OC description is given in Reference [24].

3.4.6. Busy handling The ROD module generates the Busy signal and sends it to the TBM through the CP3 board. There are 8 sources of Busy per ROD Module, one per DSP. Each source is individually maskable to isolate faulty or disconnected PUs. Busy signals are then logically *ORed* and a LED on the ROD front panel gives the status of the Busy. The status and duration of the Busy signals from each source are also available through VME. After a POWER-ON, the Busy output is forced to avoid incoming triggers during the initialization phase. For more details about the Busy module of the ROD board, see Reference [25].

3.4.7. VME interface and internal communication The VME bus system consists of 4 sub-buses: the Data Transfer bus, the Arbitration bus, the Priority Interrupt bus and the Utility bus. Data transfer is asynchronous, supporting modules with a broad variety of response times. See Reference [25] for more details on the VME bus system.

An adjustable timeout is implemented in the ROD module and if it is reached, the ROD sends a Bus Error line (BERR) signal to the VME bus. This timeout is adjustable up to 800 μ s. The VME Configuration ROM / Control and Status Register Requirements (CR/CSR) are also implemented in the ROD module. This includes a

ROM containing the serial number of the board, as well as the possibility to change the base address of the board in the crate. This function also allows enabling or disabling the ROD module.

3.4.7.1. Protocol between the VME FPGA and other ROD module FPGAs The ROD VME interface allows data transfers between the ROD crate CPU and the ROD VME FPGA, but also with the PUs, the OCs, the TTC FPGA and the receiver FPGAs. While data transfers between the CPU and the VME FPGA follow the VME standard protocol, a custom made serial protocol is used between the VME FPGA and the other chips of the ROD module. This protocol allows read and write cycles to all sub-modules, as well as broadcast writes to the PUs. For more details on these protocols, see Reference [25].

3.4.7.2. On line firmware reconfiguration through Boundary Scan All ROD FPGAs and associated Electrically Erasable Programmable Read-Only Memories (EEPROMs; motherboard and PUs) are reprogrammable through boundary scan, either using the JTAG connector, or from the VME bus, which allows firmware modifications without removing the ROD module from the crate. In this last case, the VME FPGA drives the JTAG signals according to the content of specific VME registers.

3.4.8. ROD design, manufacturing and tests The ROD module measures 366.7 mm in height, 400 mm in depth and is 2.4 mm thick. The motherboard has 3 right angle VME64x connectors. A metal bar mounted close to the VME connectors ensures a good mechanical rigidity of the module.

The ROD motherboard printed circuit has 6 signal layers and 4 equi-potential planes. Care has been taken to avoid parasitic reflections on long lines. All lines are 125 μm wide to maintain an impedance of $56 \Omega \pm 10\%$. The impedances of the motherboard and the daughterboards are adapted with a serial resistor of 47Ω at the source of the signal on the emitter side.

The connections to the 9 Ball Grid Array (BGA) devices on the ROD motherboard have been controlled after board assembly by X-ray inspection and functional tests of the board.

3.5. Processing Unit

The ROD Processing Unit (PU) is a 120 mm \times 85 mm mezzanine board which is plugged into the ROD motherboard. The connection with the motherboard is performed through 3 AMP board-to-board connectors: two 64-pin AMP-120525-1 input connectors and one 84-pin AMP-120525-2 output connector. A photograph of a PU board can be seen in figure 25.

The PUs are composed of two DSP blocks, which process the data of up to 128 calorimeter cells coming from one FEB in normal mode, and up to 256 cells from two




Figure 25. Photograph of the ROD PU with its two clearly visible DSPs. The PU measures 120mm by 85mm.

FEBs in staging mode. The DSP calculates the energy deposited in each cell, the time of the signal peak as well as a quality factor (see section 2.3.1). In addition, it performs error detection and provides histograms for monitoring.

Each DSP block is composed of one Input FPGA (InFPGA), a TMS320C6414 720 MHz DSP from Texas Instruments and one Output FIFO. The PU also contains a control FPGA (ConFPGA) used for the VME and TTC interfaces. Input and output calorimeter data streams are synchronized by the motherboard with the TTC clock signal multiplied by two (80.16 MHz), while VME and TTC exchanges are performed at the TTC clock speed. The Bus driven by the DSP and the DSP core itself are cadenced with clocks derived from a 60 MHz local oscillator. Figure 26 shows the architecture of the PU board.

3.5.1. Input FPGA (InFPGA) The InFPGA detects start of events and parallelizes incoming FEB data sent by the receiver FPGA (see 3.4.1 for the data encoding description). The InFPGA verifies data consistency, to check for potential data corruption e.g. single event upsets (SEU) due to radiation effects. Several checks are performed : verifying that parity is odd and that all samples of a given channel have the same gain values, testing the consistency of header and trailer words of the incoming events, etc. When an error is detected, the InFPGA fills the event status word correspondingly. This word is interpreted by the DSP which will take the appropriate action, such as asking for a FEB module reset. For more details about the InFPGA checks and more generally about the PU board, one can see Reference [26].

The InFPGA allocates 32 kbits of embedded memory per FEB. This memory is configured as a dual port memory, used to store data before its transmission to the DSP. It is separated into two equal banks, one of which is written to, while the other is read

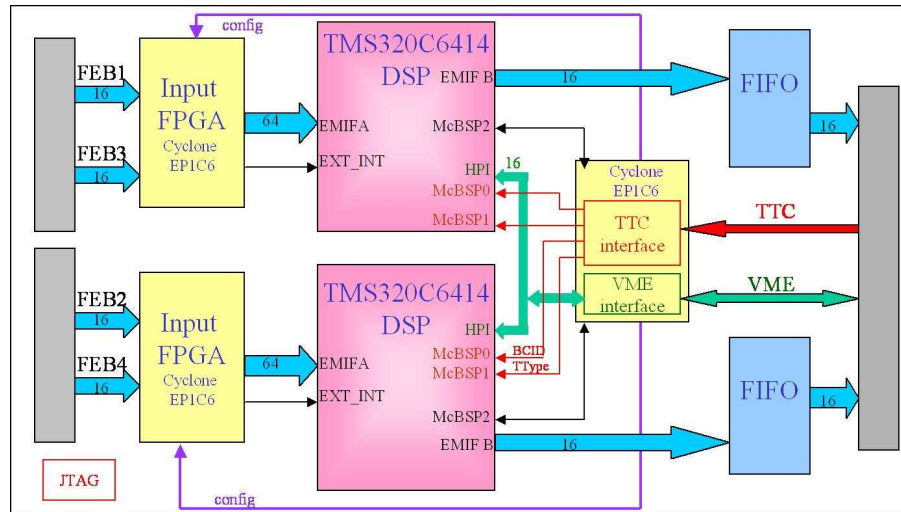


Figure 26. Block diagram of the PU board with two TMS320C6414 Texas Instruments DSPs.

out by the DSP. The configuration of the memory as RAM type allows writing data to non-consecutive addresses, and reorganizing them to optimize the event format for the DSP algorithms. However, if the event is bigger than 16 kbits, which is the size of one memory bank, the event cannot be completely rearranged.

The InFPGA allows two main modes of FEB data organization. The first organizes data by channel to optimize the DSP calculations. The internal memory size of the InFPGA allows this for FEB data with up to 7 samples (16 kbit event, for all 16 FEB ADCs). The second is a transparent format for any number of samples and gains, which was used during the ATLAS CBT (Combined Beam Test) [27] and for a general test of the readout. Note that the maximum sized event composed of 32 samples and 3 gains reaches almost 223 kbits per FEB i.e. about 14 times the size of one InFPGA memory bank.

When a complete FEB event is stored in the memory, or if the number of words written in the bank reaches the data block size set by the DSP, the InFPGA sends an interrupt to the DSP which launches a Direct Memory Access (DMA) to read the event. The readout is clocked by the 120 MHz DSP clock. From the readout side, the internal dual port memory of the Input FPGA is seen by the DSP like a FIFO, such that the data are read in consecutive memory accesses.

Despite the fact that one PU receives the data from 2 (4) FEBs in normal (staging) mode, the FEBs are treated separately by the InFPGAs. Indeed, the FEBs operate independently, and due to the varying lengths of the optical fibres linking the FEBs to the RODs, data from different FEBs arrive at different times. Hence the two DSP blocks are completely independent one from each other, and the logic in the InFPGA allowing the data treatment of 2 FEBs in staging mode is doubled. Nonetheless, the FEB treatment is strictly identical. The synchronization is performed in the DSP, using TTC data.

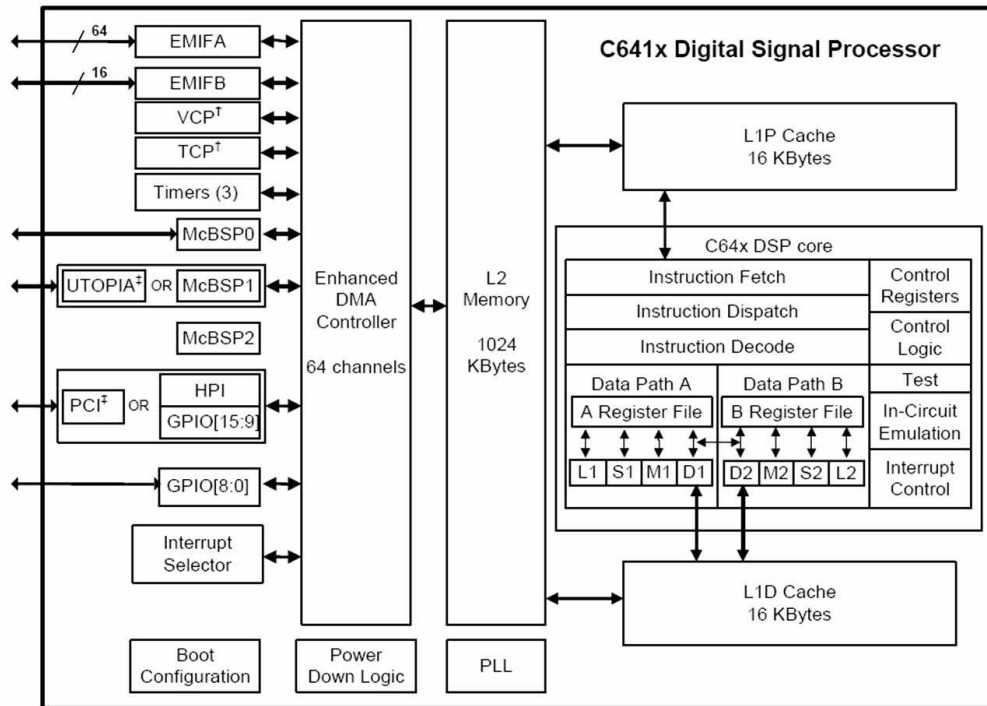


Figure 27. The TMS320C6414 DSP architecture.

3.5.2. Digital Signal Processor (DSP) The 720 MHz TMS320C6414GLZ is a high performance fixed point DSP, executing up to 5.7×10^9 instructions per second. The core processor has 64 general purpose registers of 32-bit word length and eight independent functional units: two multipliers for a 32-bit result and six arithmetic logic units. It is based on an advanced Very Long Instruction Word (VLIW) architecture, allowing up to eight functional units every clock cycle. The DSP uses a set of peripherals, adapted to the ROD needs:

- Three multichannel buffered serial ports (McBSPs) allowing full duplex 32-bit word transmission. McBSP0 and McBSP1 are used to transfer incoming TTC information, whereas McBSP2 is used to transfer VME commands to the DSP and read back status information as interrupt vectors.
- A user-configurable Host Port Interface (HPI), which is a convenient way of accessing the DSP internal memory. It is used to download code at boot time and to read data during normal processing.
- Two glueless external memory interfaces (EMIF), configured as a synchronous memory interface. The 64-bit wide EMIFA is used to input FEB data, whereas the 16-bit wide EMIFB is used to output data. Both EMIFs run at 120 MHz.
- 16 General Purpose Input/Output ports (GPIO) used to interface the ConFPGA or the InFPGA.

Figure 27 shows the DSP architecture.

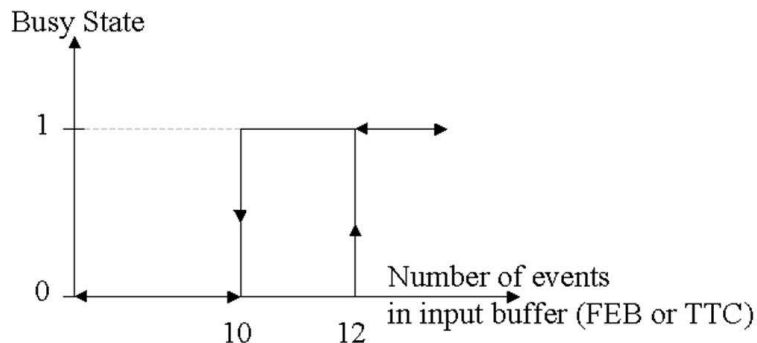


Figure 28. Hysteresis Busy generation.

Data goes into the DSP through two circular buffers (FEB and TTC), and comes out through a third one. All buffers are 16 events deep and allow to smooth out fluctuations in the average incoming data rate.

Every external transfer is handled by the Enhanced Direct Memory Access controller (EDMA). The EDMA interrupt subroutine is aroused every time a DMA transfer is finished (FEB data, TTC data or output). It is used to increment or decrement counters, allowing the DSP to know how many events are stored in each circular buffer. These counters are used for the Busy generation to control the data flow.

To generate the Busy signal, a hysteresis level type security is used on the input buffers (see figure 28). The Busy signal is generated when the input buffer (either FEB or TTC) is filled with at least 12 events (adjustable number). It is removed when the number of events in the input buffers decreases below 10 (also adjustable).

The DSP tasks are described in the following paragraphs:

- The first task is the synchronization, which checks for consistency between FEB and TTC data. The TTC information always arrives before the FEB data. It is taken as the reference and is assumed uncorrupted. If de-synchronization occurs (missing TTC or FEB event), a synchronism is searched for in all available data using a double loop scheme. If FEB data are missing, NULL events are sent out for each lost FEB event. If TTC information is missing, the corresponding FEB event is discarded and nothing is output.
- The second task is the processing, which involves physics calculations (see section 2.3.1), input to output copy mode for test purposes, calibration algorithms and generating monitoring histograms. This last function is essential since the FEB raw data are generally not transferred beyond the ROD. It includes monitoring physics information e.g. baseline, and logging of errors e.g. parity, SEU. All this information is transferred via VME to the local CPU.

During calibration runs, charges of various amplitudes are injected in the electronics chain. The DSP computes first and second moments of the pulse shape for each channel of the calorimeter.

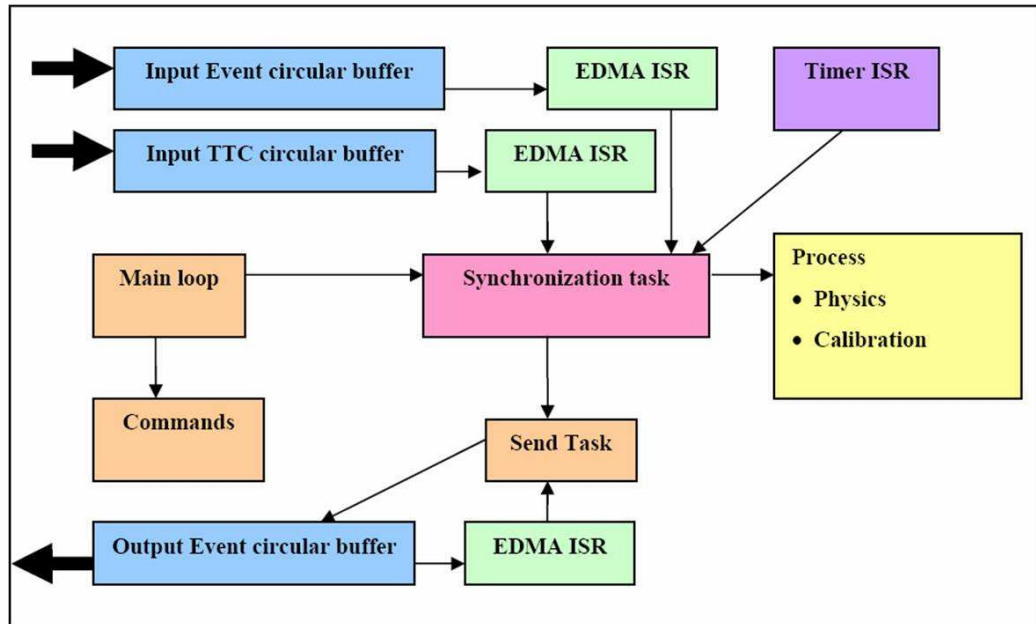


Figure 29. DSP code structure.

- After synchronization and processing, the DSP's last task is to send the data: they are prepared in the output buffer and sent to the ROD motherboard.

The DSP uses a two-level cache based architecture. The Level-1 cache consists of a 128 kbit program and a 128 kbit data memory. The Level 2 consists of an 8 Mbit memory space which is shared between program and data. This memory is used to store the DSP software, the input buffers (FEB and TTC data), the output buffers, the histograms and the calibration constants. The Level-1 cache is relatively small and causes a 30 % stall-cycle time (cycles lost because the instructions or the data are not in the Level-1 cache memory). The CPU stalls until the data or instructions are copied from the L2 memory into the cache. A dedicated memory map, as well as pipeline memory accesses are used to ensure minimal stalling when accessing data.

The DSP code is written in C and is then optimized using the Code Composer Studio program. This decreases the code complexity and simplifies its legibility and maintenance. The implementation of monitoring histograms is presently being investigated. The code structure is illustrated in figure 29.

3.5.3. Output FIFO The PUs contain IDT72V273L7-5BC 3.3V high density super synchronous Output FIFOs. Each one is organized in 16k-words of 18 bits. In most cases, this allows the DSP to send entire events or even several events at a time in single output DMA, thus optimizing data transfers. For output events bigger than 16k-words, which can be the case for raw data or 3 gains-32 sample events, the DSP slices up the event.

The FIFO has separate write and read clocks allowing DSP writes at 120 MHz

and motherboard reads at 80 MHz. The PU is thus completely independent of the motherboard output stage.

The FIFO has five flag pins: empty, full, half-full, programmable almost-empty and programmable almost-full. The programmable almost-empty and almost-full flags are configured by the DSP during the initialization phase. The offset default value is 1024 words. These five flag pins allow the DSP to control the status of the FIFO.

3.5.4. Control FPGA (ConFPGA) The ROD PU also contains a third FPGA, the ConFPGA, which has 3 main purposes:

- The DSP interface with the Output FIFOs and the output stage of the motherboard. The ConFPGA provides the status of the FIFO flags to the DSP and interfaces the DSP with the motherboard OC. The ConFPGA can also handle the DSP and master FIFO reset upon reception of a VME request, as well as test points and LEDs of the PU for debugging purposes, in particular of the DSP code.
- The TTC interface between the motherboard and both DSPs. The ConFPGA buffers the TTC data coming from the motherboard TTC FPGA. It distributes TTC signals to both DSPs through point-to-point connections, using two McBSP serial ports per DSP: McBSP0 is reserved for the BCID and EVTID transmission, while McBSP1 is for the Trigger Type. The communication protocol used between the ConFPGA and the DSPs is the same as the one used between the TTC FPGA and the PUs (see section 3.4.2).
- The VME interface of the PU, which allows to control the PU by VME. A custom made protocol is used between the motherboard and the PU, and allows point-to-point communication with a reduced number of pins. In this protocol the PU always operates as a slave (see section 3.4.7).

The VME interface of the PU is 32-register wide and allows:

- The configuration of the PU, as well as readout of the status through dedicated registers.
- The DSP boot and configuration as well as histogram readout through the 16-bit HPI. The connection between the DSPs and the ConFPGA enables the VME to access both DSPs simultaneously, allowing boot sequences which are twice as fast. When reading the internal DSP memory through the HPI bus, an average VME bandwidth of 4 Mbytes/s is observed.
- The communication with each DSP through the full duplex serial port McBSP2, allowing independent and simultaneous data transfers, both in transmission and in reception. In transmission mode, McBSP2 allows sending DSP commands or configuration words such as the run number. In reception mode, this port is used to receive interrupt vectors which describe the cause of the issuance of the DSP interrupt to the VME processor. For this specific purpose, there is one internal FIFO (32 words of 32 bits) per DSP in the ConFPGA. The status of these FIFOs can be read through the status register.

- The InFPGA boot. Due to its technology, based on SRAM memory, the InFPGA must be booted after each POWER-ON. This is ensured by the VME CPU. This procedure provides flexibility since the InFPGA firmware can be changed at any time. Both InFPGAs can be configured at the same time, in a so called broadcast mode.
- The InFPGA configuration and status read through a 3-bit custom made protocol. This allows to load the number of samples and gains per event, as well as the pre-defined block size for DSP input DMA transfer, and can also provide information about the InFPGA, such as the firmware version.

3.5.5. Boot procedure A precise procedure is followed when booting the PU:

- When the power is turned on, a DSP reset is automatically provided by a micro-monitor chip.
- An EEPROM of type Altera EPC2 downloads the firmware in the ConFPGA.
- The ConFPGA is reset by the general reset of the motherboard. VME accesses to the PU are from then on possible.
- The InFPGA firmware is downloaded from the VME, through the ConFPGA, according to the procedure described above. Once booted, a bit is set in the status register of the PU and a LED is switched on.
- The DSP code and calibration constants are downloaded from the VME CPU, through the ConFPGA using the HPI interface.
- The InFPGA is reset by the VME, through the configuration register of the ConFPGA. It can then be accessed and configured.
- The Output FIFO is reset from the DSP.
- A JTAG reset is also performed on both DSPs, so that they are properly initialized. While the DSP reset initializes the DSP core, the JTAG reset initializes the DSP's emulation logic. Both resets are required for proper operation and must be asserted upon power up. This can be performed from the motherboard, through the JTAG registers.

3.5.6. Power supply The PU is supplied with a 3.3V voltage, coming from the mother board and generated by the VME 48V backplane, through a DC-DC regulator. The PU power supply is provided through the three PU connectors, via 31 ground pins and 34 3.3V pins. The power consumption is less than 1.5A for the whole PU. On the PU, the following voltages are derived from this 3.3V supply:

- 1.5V for the InFPGA and ConFPGA cores with two TPS76815Q regulators from Texas Instruments.
- 1.4V for the DSP core with a TPS54610 adjustable regulator from Texas Instruments.

- 3.3V for the DSP Input/Output, which is set up after the DSP core voltage and with specific setup requirements.

3.5.7. JTAG chain The PU includes a boundary scan chain (JTAG chain) which has 2 purposes on the board:

- Boundary scan analysis, to detect faults coming from the production. In particular, it allows identifying stuck at, short and open circuits. The boundary scan is the main test to validate the electrical functionality of the PU, since about 90 % of nets are testable.
- Altera components programming, allowing FPGAs and EPC2 configuration.

The JTAG chain is controlled from the motherboard, either from a dedicated connector or directly through the VME bus.

3.5.8. PU design and manufacturing The PU is a 10 layer board with 6 signal layers and 4 power planes, with approximately 2600 etches, and 2300 through hole vias of 0.5 mm in diameter and 0.25 drill size under the DSP (0.8 mm BGA), and 0.6 mm in diameter and 0.3 drill size elsewhere. The minimal line width and spacing is of 120 μm . The planes are organized such that the signal layers have almost the same impedance ($\sim 50 \Omega$ for a 0.120 mm wide net) in order to preserve signal integrity.

The PU layout has fast Input/Output pins with fall times between 1ns and 3ns, which can contribute to noise generation, signal reflection, cross talk and ground bounce, such that special care has been taken to:

- filter noise coming from the power distribution using decoupling capacitors and power planes;
- respect sensitive signals including clocks using:
 - point-to-point connections, such as a zero delay PLL for clocks when necessary;
 - traces as straight as possible with the clock signals routed first;
 - a minimized number of vias.
- match impedances and terminate lines using a special layout cross section and 50 Ω serial resistors at the emitter side when necessary;
- minimize cross talk between parallel traces using a minimum of 0.36 mm spacing at each side of the clocks.

The PCB is a dielectric of FR4 type, with double face varnish, chemical Ni Au finishing and layer impedances within a 10% margin. After the PCB production, a complete electrical test was performed on each PCB. The ROD PU has about 285 Surface Mount components, cabled on both sides of the PCB, amongst which are 7 BGAs: two 532-pin 0.8mm-pad, two 100-pin 1mm-pad and three 256-pin 1mm-pad BGAs. Prior to assembly, the PCBs were stored under controlled temperature, to prevent oxidation. A double refusion process was performed with nitrogen gas, to ensure a better cabling quality between the PCB and the components, and for a better board aspect. After



Figure 30. Photograph of the TM module with 4 HOLA mezzanine cards.

assembly, a visual check was performed, as well as a BGA check using X rays, and a JTAG interconnection test [28]. The rate of success for these tests was of the order of 95%. Various reception tests were also performed with a 99% success rate.

3.6. Transition Module

To each ROD module is associated a TM. Situated at the rear of the ROD crate, it serves as the interface between the RODs and the ROS. The ROD output data are sent to the TM through the VME backplane. One TM supports four S-Link mezzanine cards with optical drivers, called HOLA's. Each HOLA is associated to one PU. A TM module is shown in figure 30. See Reference [29] for a detailed description of the TM module.

3.6.1. FIFO memory Each ROD to TM channel uses a serializer-deserializer chipset. The use of these chips introduces a latency of 5 to 6 cycles in the data flow between the input of the serializer on the ROD and the output of the deserializer on the TM. There are always several data words being transferred from the ROD to the TM, even when the OC has stopped its transmission. A FIFO (IDT72V205L-15), which memorizes the




Figure 31. Duplex optical connector which connects onto the HOLA board.

32 data bits and the CTRL bit, is used as a decoupling buffer between the deserializer and the HOLA card.

3.6.2. Flow control Each of the four ROD-TM channels ends with an S-Link that connects the ROD system to the ROS. The transmitting end of the S-Link is implemented on the HOLA card. Figure 31 shows the duplex optical connector which connects onto the HOLA board. Each data word transmitted is accompanied by an additional CTRL bit which identifies it (block address, event header, end of block, etc.).

Data can only be transferred to the HOLA when the S-Link Link Full Flag is inactive. This flag indicates that the ROS has reached a state in which it cannot process new data words coming from the link. In this case, the HOLA sets its Full Flag which forces the TM FIFO memory to stop transmitting data to the HOLA. However, the FIFO continues to accept data from the ROD. When the FIFO is half full (more than 128 words have been stored and still not read), it activates the FIFO Full line to tell the OC to stop transmitting data. The S-Link checks for errors in the data and control words, but does not correct for them. These can be reported on a block-by-block basis or on a word-by-word basis.

3.6.3. Power, reset, ESD and EMC considerations The typical TM current consumption with 4 HOLAs is 0.255 A from 48 V. Table 4 shows the TM current consumption at 40 MHz. A re-triggerable variable resistor (PolySwitch from RAYCHEM) is used to protect the CP3 board from hazardous currents that could occur in the TM. When the power is turned on or when the ROD generates a TM-reset, the on-board logic is reset, the FIFOs are cleared and the HOLAs reset themselves. In case of possible hazardous electrostatic discharges (ESD), two copper strips establish a connection to ground through 2 M Ω resistors when the board is inserted into the crate. For electromagnetic compatibility (EMC), the front panel is equipped with an EMC

Table 4. TM current consumption from 48 Volts.

TM with 4 HOLAs transmitting	TM with 4 HOLAs 2 transmitting 2 not transmitting	TM with 2 HOLAs transmitting
255mA (13W)	219mA (11W)	163mA (8W)

gasket and is connected to ground through four $0\ \Omega$ conductors.

3.6.4. TM design and manufacturing The two TM module connectors are female right angle VME64x connectors. The module itself measures 220 mm in height, 400 mm in depth and is 1.54 mm thick. It is designed to be mounted at the rear of a 9U crate. The printed circuit contains 4 signal layers and 4 equi-potential planes. Care has been taken to avoid parasitic reflections on the paths, in the same way as for the CP3 (see section 3.1.4) and for the TBM (see section 3.3.6).

3.7. ROD Injector

The ROD Injector simulates the flow of data coming from the FEBs and from the TTC system, and is used for testing the ROD boards. Three test benches made use of the Injector during the ROD production phase and will also be used for maintenance.

The Injector is interfaced with the RODs and reproduces the FEB behaviour in terms of its digital functionalities without the complex analog part of the FEB. With its 5 independent data injectors (organized in 10 half-FEB data injectors) and three 1-to-2 optical splitters, it is able to drive one full ROD (8 FEB inputs). Many Injectors can be run simultaneously to test a full ROD crate. Half a ROD crate (7 out of 14 RODs) is driven using 6 Injectors and 30 1-to-2 optical splitters. The ROD crate electronics were commissioned in this way.

An integrated sequence generator reproduces the LHC machine timing signals and the experiment trigger decision logic signals. The Injector can drive the ROD timing logic using a cable connected directly to the ROD without the need of any TTC modules. It can also drive the TTC system using its own TTCvi interface without the need of an LTP module. The Injector can generate normal sequences and events, as well as all kinds of errors in order to test the error recovery procedures of the ROD. It can also generate errors on a single FEB.

The Injector is a 9U VME64x module and is fully integrated in the ROD software environment. It has a standard VME64x interface, which integrates A32-D32 and A32-D16 data access in single and block transfer modes. One or more modules can operate while being connected in the same or in a different crate as the RODs. Figures 32 and 33 show respectively the architecture and a photograph of an Injector board.

The functionalities are implemented in 7 on-board FPGAs: one is dedicated to the VME interface, five to the data injectors which all have the same code, and finally one

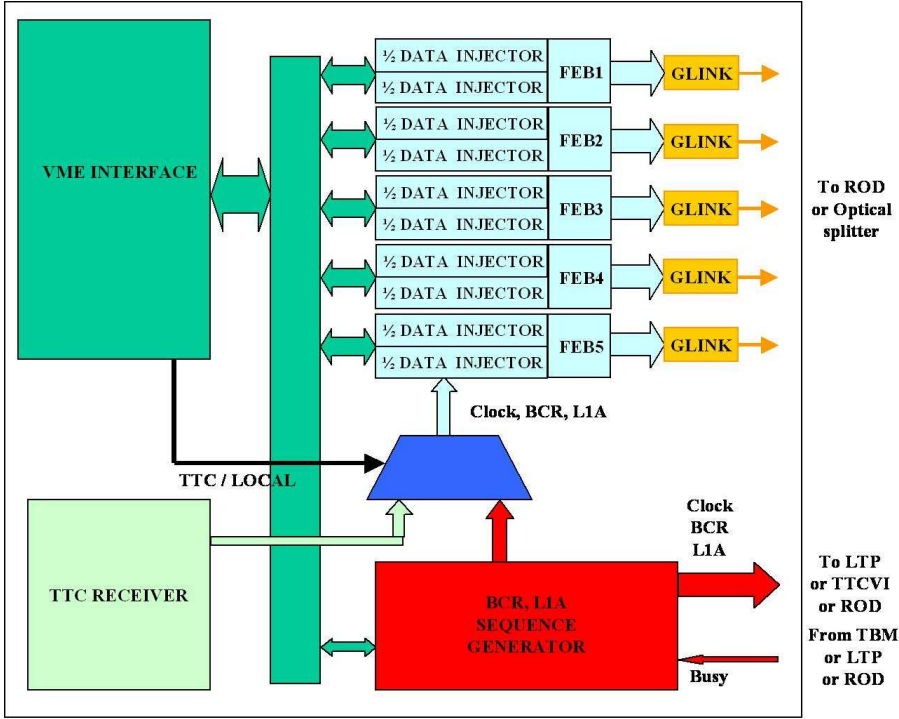


Figure 32. Architecture of the ROD Injector.




Figure 33. Photograph of the ROD Injector board.

for the sequence generator.

3.7.1. TTC receiver A TTC receiver is implemented in order to synchronize many Injectors, the first of which provides the TTC signals to the TTC system. Another scheme is that all Injectors are driven by an independent external source driving the TTC system e.g. the LTP internal generators. The TTC receiver consists of a TTCrx with switches to set its geographical address. The signal coming from the CP3 backplane goes through an LVDS receiver before going to the TTCrx. The module can be plugged into a ROD crate and uses the same TTC P3 signal distribution as the ROD. A Quartz Phase-Locked Loop (QPLL) is incorporated to reduce the amount of jitter produced by the TTC system

3.7.2. Sequence generator The sequence generator provides the experiment with timing signals and trigger decision logic signals either directly to the ROD module or to the TTC system via the TTCvi or the LTP. These control and timing signals are:

- A 40.08 MHz clock generator to emulate the machine BC (beam crossing) signal.
- A 40.08 MHz clock down-counter controlled by VME to generate the machine BCR (bunch counter reset).
- A L1A generator controlled by VME to emulate the experiment trigger decision logic. It generates a ROD L1A and an independent L1A for each half-FEB data injector. A sequence of L1As is stored in a memory and this memory is read at the LHC clock frequency either once or in continuous loop mode. This allows all possible L1A patterns and sequences, legal or illegal. The maximum duration of a single sequence is 800 μ s. The number of L1As sent is programmable, from 1 to continuous.
- A Busy signal in order to control the flow of L1A from outside.

All of these signals are available on the front panel to interconnect to the TTCvi and TBM modules. They are also available on a dedicated LVDS cable to interconnect directly with a ROD.

The sequence generator gets its L1A data from a 16-bit 32 K deep memory. Only 13 bits are used: 10 for the L1A corresponding to each half-FEB data injector, 1 for the ROD, 1 to trigger an oscilloscope and the last one to generate a calibration signal. When the L1A generator is started, and from the next occurrence of a BCR, this memory is read every 25ns and its content is sent. The size of the memory used can be defined in a register. A loop mode allows to loop on the memory when a sequence needs to be repeated indefinitely.

The sequence of L1A can be suspended by the presence of different Busy signals. These Busy signals can be enabled with a mask register. In fact, when a Busy signal is present, the memory continues to be read, but the signals coming out of it are suppressed. This is done in order to respect the position of the signals relative to the BCR for a given sequence.

A status register gives information about the status of the different Busy signals. A register can be set with the number of L1A to be transmitted to the ROD. It operates through the generation of a Busy signal. Another register counts the number of L1A which have effectively been transmitted to the ROD. A register controlling a delay line is used to set the position of the clock relative to the data transmitted to the G-link.

3.7.3. Data injector There are 10 half-FEB data injectors on the board organized in groups of two. Each half-FEB injector gets its data from a 16-bit 32 K deep memory. Data are stored in lots of 8 consecutive 16-bit words corresponding to the 8 FEB ADCs it must emulate. The lots are organized following the FEB format with a start of event, the control words where the BCID and the EVTID are stored, the data words, and finally an end of event. The BCID and EVTID can either be set within the memory or replaced on the fly with values from the TTC system. The event size and of the memory size used to store all events are defined in registers.

On reception of a dedicated L1A, the data injector will send, after a programmable time, an event in the FEB format to mimic the FEB response. Up to 81 five sample events can be stored into the memory, where parity is also stored. When the EVTID and BCID are generated on the fly, parity for these two words is calculated locally. A flag can also be set to disable this parity generation. The L1A comes either from the sequence generator or from the TTC system. The stored events can be read once or in continuous loop mode. The size of the memory used for storing the events is programmable.

The incoming L1A are stored in a FIFO. In normal working mode, they can come every 125ns or more, with a mean maximal rate of 100 kHz. For each L1A in the FIFO, a complete event is read from the memory after a latency time pre-defined in a register. If this FIFO overflows e.g. because too many L1As have been sent in the time necessary to transfer the corresponding events, a FIFO_FULL bit is set in the status register and the data injector does not accept a L1A until the system is reset.

Each data injector is connected to a G-link. In the same way as the FEB data, data coming out of two half-FEB injectors are multiplexed into the SMUX [21], the SMUX output then enters the HDMP-1022 transmitter, and the signal coming out of the transmitter is sent to the OTx optical transmitter.

3.7.4. VME interface The VME FPGA communicates with the other FPGAs through a dedicated 16-bit address bus and a bi-directional 16-bit data bus. It implements the registers necessary to control the other FPGAs as well as the general setup of the board, but it is mostly a VME interface between the VME bus and the internal communication bus. All VME address decoding is done inside this FPGA. This interface allows different types of resets on the board through a specific reset register. It can start and stop the BCR generator and the L1A Generator. A control register defines the mode of operation of the board and a status register gives information about the different components on the board.

4. BE system tests

Many quality checks and tests have been performed on the various modules or sub-components of the BE electronics system. These checks took place at the different collaborating institutes, as well as at CERN, at the EMF (Electronics Maintenance Facilities) or during the CBT (Combined Beam Test), when a slice of the ATLAS detector was set in the H8 beamline of the SPS (Super Proton Synchrotron) [4]. The detector, electronics and software have also undergone extensive testing during installation and commissioning in the USA15 cavern.

The various tests confirmed that the following BE system requirements were fulfilled: FEB and TTC data readout and processing, local processing of data during calibration, distribution of the timing clock and trigger to the FE and to the RODs, and configuration and control of the FE crates. In what concerns the 75 kHz physics mode rate, partial system tests were performed e.g. speed of the DSP data processing, high incoming L1A rates (up to 100 kHz), confirming the robustness and reliability of the system, and that the LArg BE system can tolerate the required Level-1 rate.

In the following paragraphs, the essential tests are described.

4.1. Injector tests

Tests were performed using the ROD Injectors. One Injector is capable of testing the 8 FEB inputs of a ROD module: three of its five data injectors are associated to *one to two* optical splitters while the other two are not, giving the 8 G-links. The ROD data are read out through 4 S-links connected to a custom interface board set in a readout PC. Machine signals and the L1A Trigger decision signal are generated by the Injector and sent to the TTCvi which distributes them through the TTCex, TBM, CP3 backplane and TTCrx to the ROD and to the 1/2 data injectors.

This setup was very useful in the debugging phase of the ROD prototypes. Its sequence generator allowed to test the TTC system and ROD board operation e.g. the BCID and EVTID generation. The well known relationship between BCR and L1A allowed to understand all TTC synchronization problems between the TTC system and the events received within the DSP. It also allowed to generate near the limit conditions in terms of L1A position and frequency.

With its data generators, it was easy to generate known events to verify the proper operation of the ROD input chain up to the DSP. This same setup has been used for testing the whole of the ROD production. Figure 34 shows an example of a test using the Injector board.

4.2. Configuration of the system

The various components of the LArg electronics are initialized and configured before any data is processed. The Partition Master (see section 2.3.2) can run the ATLAS data acquisition software (TDAQ) and thus every partition can operate in stand-alone mode.




Figure 34. Schematic view of a BE electronics test setup using an Injector board, both for FEB and TTC data input.

The Partition Master synchronizes the operations of the ROD and TTC crates, sending various commands to their crate controllers : BOOT, CONFIGURE, START and STOP. In a more global mode where several partitions are run from a global trigger, the control is transferred to a higher level. The Partition Masters then receive the commands from above and just transmit them.

The ROD Crate Controller configures and controls each board in the ROC and in the associated FECs. First the system is booted. The DSP code and the code necessary to configure the Input FPGA are then loaded from source files. The components are configured : the number of samples to read out is sent to the DSPs and to all FEBs through SPAC links, various delays are set either through the TTC system or using SPAC links; for calibration runs, calibration board Digital-to-Analog Converter (DAC) values, delay values and pulsing pattern are loaded. The run can then be started.

The time it takes to configure the FE and BE systems is presently dominated by the ROD access to the database (from 30 seconds up to three minutes). Nonetheless, the configuration of the BE electronics lasts less than one second per board, and since all crates are performed in parallel, the whole BE system configuration takes less than 15 seconds. For the FE system, it takes approximately 30 seconds.

The system configuration is performed regularly, and is robust and efficient.

4.3. Speed of the DSP code

The DSP code (written in C) calculates the average energy (E), the energy sum projected along the ATLAS x , y and z axis (E_x , E_y , E_z), the average time and the quality factor. The DSP calculations have been tested to give the same results as those performed offline. The speed of the system is limited by the downstream DAQ readout system.

Events with up to 32 samples can be processed.

Once the data arrive from the TTC or from the FEBs, an interrupt is lifted in the DSP and the data are stored (see section 3.5.2). If the buffer is full, a Busy signal is generated. The time of arrival of the TTC data is recorded and these are henceforth used as the reference. A synchronization code then searches for FEB data which correspond to the reference TTC data. If there are no FEB data found within a pre-defined time interval, a null event is sent to the ROS. If FEB data are found, they are compared with those of the TTC. If the EVTID and BCID do not match, the data are thrown away and the next stream of FEB data is taken. When a match is found, the DSP calculations are performed.

The data input and DSP calculations are performed in parallel. The input time for data with 5 samplings and one gain is approximately $10 \mu\text{s}$ per event. Once the DSP calculations are performed, the data are output through the OC in less than $10 \mu\text{s}$. This step is also performed in parallel and does not lengthen the processing time of events.

The DSP code was tested and timed on a full BE setup. In the example shown in figure 35, the energy ($2.7 \mu\text{s}$), time and quality factor ($3.7 \mu\text{s}$) computation contribute $6.4 \mu\text{s}$ to the time it takes to process the data of one FEB (128 cells, 5 samplings). In addition, (Ex,Ey,Ez) and the maximum energy cell per layer in each trigger tower are computed, adding $1.6 \mu\text{s}$. The computation of the checksum, which is checked offline to verify the data integrity after transfer, takes up another $0.7 \mu\text{s}$. The overall time, taking into account the FEB to Level-1 TTC synchronization, is $9.3 \mu\text{s}$.

However, since the time and the quality factor will be computed only for the cells above 5σ of the noise (10 to 20% of the cells), the (Ex,Ey,Ez), time and quality, and checksum contributions will consequently be reduced. For example, in a measurement where 25% of the cells were above threshold, the reduction factors were of 1.6/0.6 (Ex,Ey,Ez), 3.7/1.1 (time and quality) and 0.7/0.5 (checksum). The overall computing time was $5.5 \mu\text{s}$.

These processing time measurements are safely below the upper limit of $13.3 \mu\text{s}$ set by the maximum Level-1 trigger rate of 75 kHz, leaving sufficient space to implement monitoring histograms. Nonetheless, a full system test has yet to be performed.

4.4. Commissioning in the cavern

As soon as the FEBs were installed in the cavern and connected to the BE electronics, calibration pulses were injected on the detector and read back with the whole FE and BE readout chain. Regular calibration runs were recorded: pedestal runs, delay runs (signal as a function of the time for pulsed channels for a fixed pulse height) and ramp runs (signal as a function of the DAC value with fixed timing of the calibration pulse). The integrity of all the connections, the whole readout chain, the calibration system as well as the status of the detector cells were efficiently checked.

A dedicated DSP code was developed to allow for faster analysis and data size reduction in calibration runs. This code computes the mean values and r.m.s. of the

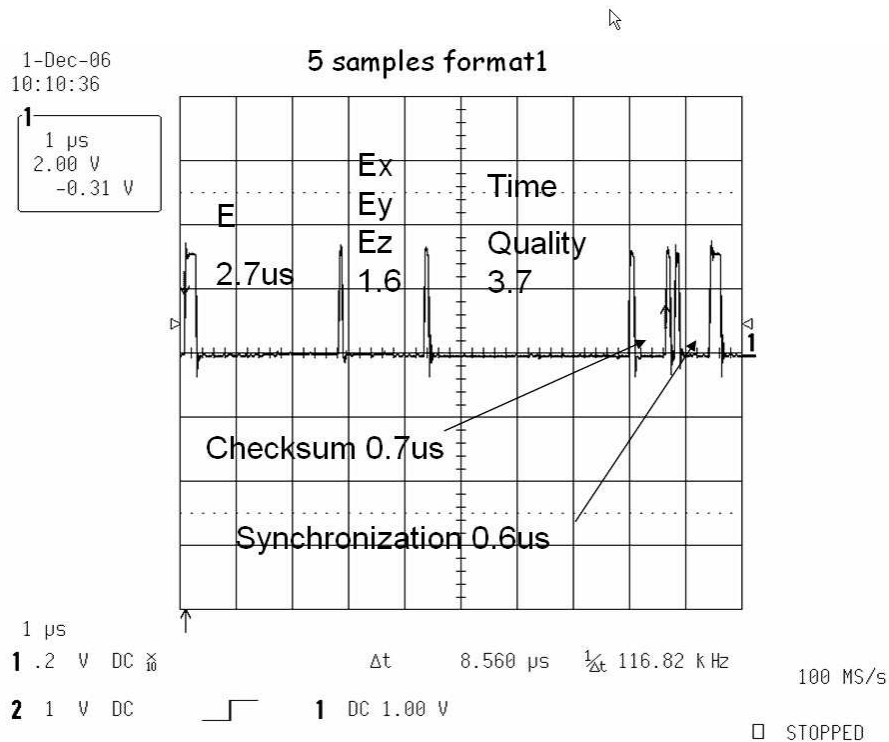


Figure 35. DSP pulses signaling the processing stages as seen on an oscilloscope, for Injector generated events. The times to compute the energy, the energy sum projected along the ATLAS x, y and z axis (Ex, Ey, Ez), the time and the quality factor, and the times to perform the checksum (to verify data integrity) and the FEB to Level-1 TTC synchronization, are measured for all 128 cells of a FEB.

signal in each cell for a given number of triggers with identical calibration parameters.

The system loops over many sets of DAC values, delay values, pulsing pattern and gain, sending typically 100 calibration pulses per point and per calorimeter cell. Null events are sent out to the ROS (to keep synchronization with this system), and when the last event of a given configuration is reached, the results of the DSP computations are also transmitted, and another set of calibration parameters are used.

As an example, when the system loops around the delay values, each cell is pulsed 100 times with a given DAC value at a given time delay between the calibration pulser and the LHC clock. For each trigger, the ADC value of each of the samples is read, the mean and r.m.s. of these ADC values over all the triggers for each sample are computed in the DSPs. The delay is sequentially increased in steps until the sampling period is covered. Eventually one obtains an averaged profile of the response of the cell to the input DAC. The readout pedestal is finally subtracted to restore the proper baseline. A delay run pulse shape is shown in figure 36.

In addition, the commissioning of the BE electronics with cosmics has allowed to exercise the readout chain and to check and adjust the detector timing at the ns level. A pulse shape generated by a cosmic muon traversing the calorimeter is shown in figure 37.




Figure 36. A delay run pulse shape for a DAC value of 500 and with a delay step of 1.04ns.

The pulse was sampled every 25ns for 30 samples. The horizontal line corresponds to the value of the channel pedestal.

5. Conclusion

The ATLAS calorimeters electronics, which were chosen after an R&D phase and were constructed, installed and extensively tested, have been described in this paper. They will measure the energy deposit in each calorimeter cell to better than 0.25% at high energy, and will provide the trigger system with the energy deposited in each layer of all the trigger towers active in an event.

The whole LArg electronics chain has been described, followed by an overview of the ROD, TTC and Level-1 receiver systems, with the tasks they must perform during the physics and calibration runs. The ROD crate components have then been described in detail.

It has been shown that the following BE system design requirements have been fulfilled: the local processing of data during calibration, the distribution of the timing clock and trigger to the FE and to the RODs, the configuration and control of the FE crates, and of course the FEB and TTC readout and processing.

In what concerns the physics mode at 75 kHz, partial tests have shown that the BE system can safely support this rate. Strengthened by the experience acquired, the collaboration is presently in a phase of global commissioning, namely the LArg




Figure 37. A pulse shape generated by a cosmic muon traversing the electromagnetic calorimeter with a sampling every 25ns, for 30 samples. The horizontal line corresponds to the value of the channel pedestal. An ADC count of 2500 above the pedestal is approximately equivalent to 25 GeV deposited in this readout cell.

electronics as well as all sub-detector electronics combined, to prepare for the first LHC data.

Acknowledgments

We wish to thank D. Cuisy, B. Debennerot, H. Rosenzweig, P. Rusquart, R. Sliwa from LAL, N. Chevillot, F. Corageoud, J.-M. Nappa, J.-L. Panazol from LAPP, and M.-C. Cloarec, M. Dhellot, C. Goffin, A. Guimard, J.-M. Parraud, A. Sefri from LPNHE, for their precious and essential engineering contributions; M. Citterio from Universita degli Studi di Milano and B. Cleland from the University of Pittsburgh for their reviewer skills; M. Delmastro from CERN and L. Serin from LAL for some figures.

References

- [1] ATLAS Technical Proposal, CERN/LHCC/94-43.
- [2] Aubert B *et al.* 2006 *Nucl. Inst. Meth. A* **558** 388
- [3] Cleland W E and Stern E G 1994 *Nucl. Inst. Meth. A* **338** 467
- [4] Neukermans L *et al.* 2001 ATLAS Publication ATL-LARG-2001-008; Prieur D 2005 ATLAS Publication ATL-LARG-PUB-2005-001; Banfi D *et al.* 2006 ATLAS Publication SN-ATLAS-2005-054, *J. Instrum.* **1** 08001; Aleksa M *et al.* 2006 ATLAS Publication ATL-LARG-PUB-2006-003

- [5] See <http://ttc.web.cern.ch/TTC/intro.html>
- [6] Cleland W E *et al.* 2006 *Receiver/Monitor System for the ATLAS Liquid Argon Calorimeters* CERN EDMS Document <https://edms.cern.ch/document/347184/>
- [7] Cleland W E 2006 *Cabling of the ATLAS Liquid Argon Receiver System* CERN EDMS Document <https://edms.cern.ch/document/347182/>
- [8] Wiener Type UEP 6021-LHC9U, OB06.011J, OF00.070B, OP07.1160. See <http://www.wiener-d.com> and <http://ess.web.cern.ch/ESS/EquipmentSelectorGuide/Crates/crates.htm>
- [9] Datasheet of the VP110 VME CPU by Concurrent Technologies. See <http://www.gocct.com/sheets/vp11001x.htm>
- [10] Baron S 2003 *Remote Power Supply Project* CERN Publication CERN-EPESS-2003-005
- [11] Hubaut F *et al.* 2004 *The ATLAS LArg ROD G-links Cooling System* ATLAS Publication ATL-ELEC-2004-002
- [12] Corbaz and Houd C 2005 *ATLAS LAr ROD G-links Racks Cooling System* CERN EDMS Document <https://edms.cern.ch/document/488953/>
- [13] ATLAS LAr DCS Group 2005 *ATLAS Liquid Argon DCS* CERN EDMS Document <https://edms.cern.ch/document/597764/>
- [14] Matricon P *et al.* 2004 *The CP3 Board for the ATLAS LARG ROD System* CERN EDMS Document <https://edms.cern.ch/document/478376/>; Matricon P 2006 *CP3 Manual* CERN EDMS Document <https://edms.cern.ch/document/754062/>
- [15] Hubaut F *et al.* 2006 *Specifications for the Serial Protocol for the Atlas Liquid Argon Calorimeters* CERN EDMS Document <https://edms.cern.ch/document/110351/>
- [16] See <http://www.erg.abdn.ac.uk/users/gorry/course/phy-pages/man.html>
- [17] Dhellot M *et al.* 2004 *SPAC Master VME 64 boards User Guide*, ATLAS Publication ATL-ELEC-2004-004
- [18] Matricon P *et al.* 2005 *TBM (and TTCM) : The Trigger and Busy Module for the ATLAS LARG ROD System* CERN EDMS Document <https://edms.cern.ch/document/653609/>
- [19] Gällnö P 2001 *ATLAS ROD Busy Module, Technical description and users manual* CERN EDMS Document <https://edms.cern.ch/document/319209/>
- [20] Parsons J and Simion S 2006 *Output Data Format for the ATLAS LAr Front End Board* CERN EDMS Document <https://edms.cern.ch/document/785268/>
- [21] Dzahini D *et al.* 2007 *The SMUX chip Production Readiness Review* CERN EDMS Document <https://edms.cern.ch/document/824318/>
- [22] ORx components, see e.g. <http://www-hep.phys.sinica.edu.tw/~atlas/lar.html>
- [23] Bee C *et al.* 2006 *The raw event format in the ATLAS Trigger and DAQ* CERN EDMS Document <https://edms.cern.ch/document/445840/>
- [24] La Marra D 2004 *Output Controller FPGA of the LArgROD board, Description* CERN EDMS Document <https://edms.cern.ch/document/460775/>
- [25] La Marra D *et al.* *VME and BUSY FPGA for the ROD motherboard* CERN EDMS Document <https://edms.cern.ch/document/466898/>
- [26] Prast J 2005 *The ATLAS Liquid Argon Calorimeters Read Out Driver (ROD) : The TMS320C6414 DSP Mezzanine board* CERN EDMS Document <https://edms.cern.ch/document/598713/>
- [27] Combined Beam Test, see http://atlas.web.cern.ch/Atlas/GROUPS/-LIQARGON/Comb_TB/index.html
- [28] Boundary-Scan Testing in Altera Devices, AN 39: IEEE 1149.1 (JTAG)
- [29] Matricon P 2005 *TM Manual* CERN EDMS Document <https://edms.cern.ch/document/653614/>