

The ATLAS Level-1 Calorimeter Trigger: PreProcessor implementation and performance

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2012 JINST 7 P12008

(<http://iopscience.iop.org/1748-0221/7/12/P12008>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 137.138.93.221

The article was downloaded on 23/07/2013 at 14:20

Please note that [terms and conditions apply](#).

The ATLAS Level-1 Calorimeter Trigger: PreProcessor implementation and performance



B. Åsman,^p R. Achenbach,^g B.M.M. Allbrooke,^b G. Anders,^g V. Andrei,^g V. Büscher,^j H.S. Bansil,^b B.M. Barnett,ⁿ B. Bauss,^j K. Bendtz,^p C. Boehm,^p J. Bracinik,^b I.P. Brawn,^o R. Brock,^l W. Buttinger,^c R. Caputo,^j S. Caughron,^l L. Cerrito,ⁱ D.G. Charlton,^b J.T. Childers,^{d,1} C.J. Curtis,^b A.C. Daniells,^b A.O. Davis,^o Y. Davygora,^g M. Dorn,^g S. Eckweiler,^j D. Edmunds,^l J.P. Edwards,^o E. Eisenhandler,ⁱ K. Ellis,ⁱ Y. Ermoline,^l F. Föhlich,^g P.J.W. Faulkner,^b W. Fedorko,^l J. Fleckner,^j S.T. French,^c C.N.P. Gee,ⁿ A.R. Gillman,ⁿ C. Goeringer,^j T. Hülsing,^j D.R. Hadley,^b P. Hanke,^g R. Hauser,^l S. Heim,^l S. Hellman,^p R.S. Hickling,ⁱ A. Hidvégi,^p S.J. Hillier,^b J.I. Hofmann,^g I. Hristova,^a W. Ji,^j M. Johansen,^p M. Keller,^g A. Khomich,^g E.-E. Kluge,^g J. Koll,^l H. Laier,^g M.P.J. Landon,ⁱ V.S. Lang,^g P. Laurens,^l F. Lepold,^g J.N. Lilley,ⁱ J.T. Linnemann,^l F. Müller,^g T. Müller,^j K. Mahboubi,^e T.A. Martin,^b A. Mass,^g K. Meier,^g C. Meyer,^j R.P. Middleton,ⁿ T. Moe,^p S. Moritz,^j J.D. Morris,ⁱ R.D. Mudd,^b R. Narayan,^h M. zur Nedden,^a A. Neusiedl,^j P.R. Newman,^b A. Nikiforov,^a C.C. Ohm,^d V.J.O. Perera,^o U. Pfeiffer,^g P. Plucinski,^p S. Poddar,^g D.P.F. Prieur,^m W. Qian,ⁿ P. Rieck,^a E. Rizvi,ⁱ D.P.C. Sankey,ⁿ U. Schäfer,^j V. Scharf,^g K. Schmitt,^g C. Schröder,^j H.-C. Schultz-Coulon,^g C. Schumacher,^g R. Schwienhorst,^l S.B. Silverstein,^p E. Simioni,^j G. Snidero,ⁱ R.J. Staley,^b R. Stamen,^g P. Stock,^g M.C. Stockton,^k C.L.A. Tan,^b S. Tapprogge,^j J.P. Thomas,^b P.D. Thompson,^b M. Thomson,^c P. True,^l P.M. Watkins,^b A.T. Watson,^b M.F. Watson,^b P. Weber,^f M. Wessels,^g C. Wiglesworthⁱ and S.L. Williams^c

^aInstitut für Physik, Humboldt-Universität zu Berlin,
Berlin, Germany

^bSchool of Physics and Astronomy, University of Birmingham,
Birmingham, United Kingdom

¹Corresponding author.

^c*Cavendish Laboratory, University of Cambridge,
Cambridge, United Kingdom*

^d*CERN,
Geneva, Switzerland*

^e*Fakultät für Mathematik und Physik, Albert-Ludwigs-Universität,
Freiburg, Germany*

^f*II. Physikalisches Institut, Georg-August-Universität,
Göttingen, Germany*

^g*Kirchhoff Institut für Physik, Ruprecht-Karls-Universität Heidelberg,
Heidelberg, Germany*

^h*Physikalisches Institut, Ruprecht-Karls-Universität Heidelberg,
Heidelberg, Germany*

ⁱ*School of Physics and Astronomy, Queen Mary University of London,
London, United Kingdom*

^j*Institut für Physik, Universität Mainz,
Mainz, Germany*

^k*Department of Physics, McGill University,
Montreal QC, Canada*

^l*Department of Physics and Astronomy, Michigan State University,
East Lansing MI, United States of America*

^m*Department of Physics and Astronomy, University of Pittsburgh,
Pittsburgh PA, United States of America*

ⁿ*Particle Physics Department, Rutherford Appleton Laboratory,
Didcot, United Kingdom*

^o*Technology Department, Rutherford Appleton Laboratory,
Didcot, United Kingdom*

^p*Department of Physics, Stockholm University,
Stockholm, Sweden*

E-mail: jtchil0@yahoo.com

ABSTRACT: The PreProcessor system of the ATLAS Level-1 Calorimeter Trigger (L1Calo) receives about 7200 analogue signals from the electromagnetic and hadronic components of the calorimetric detector system. Lateral division results in cells which are pre-summed to so-called *Trigger Towers* of size 0.1×0.1 along azimuth (ϕ) and pseudorapidity (η). The received calorimeter signals represent deposits of transverse energy.

The system consists of 124 individual PreProcessor modules that digitise the input signals for each LHC collision, and provide energy and timing information to the digital processors of the L1Calo system, which identify physics objects forming much of the basis for the full ATLAS first level trigger decision. This paper describes the architecture of the PreProcessor, its hardware realisation, functionality, and performance.

KEYWORDS: Trigger concepts and systems (hardware and software); Trigger algorithms; Modular electronics; Front-end electronics for detector readout

Contents

1	Introduction	1
2	The L1Calo PreProcessor System	3
3	The PreProcessor Module	5
3.1	Analogue input boards	7
3.2	The multi-chip module	9
3.3	The PreProcessor ASIC	11
3.3.1	Bunch crossing identification and E_T measurement	12
3.3.2	Readout and diagnostic functions	13
3.4	The LVDS cable driver	14
3.5	The readout manager FPGA	15
3.6	The timing, trigger and control decoder module	15
3.7	Module protection, control and monitoring	15
3.8	Processing latency measurements	16
4	PreProcessor Module control and operation	17
4.1	The readout manager	17
4.2	Physics data taking	19
4.3	Calibration data taking	20
5	Calibration and performance of the PreProcessor System	21
5.1	Hardware validation	21
5.2	Timing calibration and performance	22
5.3	Bunch crossing identification calibration and performance	23
5.4	Energy calibration and performance	25
5.5	Performance during physics data taking	26
6	Conclusions	26

1 Introduction

ATLAS is one of four major experiments built at the Large Hadron Collider (LHC) accelerator ring. It is a general-purpose experiment for proton-proton collisions, designed to observe the wide spectrum of physics processes expected at the LHC. The innermost detector consists of a high resolution tracker which surrounds the interaction point and is immersed in a 2 T solenoid field. This is surrounded by Electromagnetic (EM) and Hadronic (HAD) Calorimeters, which are enclosed by a muon spectrometer. The EM Calorimeter uses lead absorbers immersed in Liquid Argon (LAr)

which acts as the ionising medium and is divided into three main sections, the barrel ($|\eta| < 1.5$),¹ the end-cap ($1.5 < |\eta| < 3.2$), and the forward calorimeter, or FCAL ($3.2 < |\eta| < 4.9$). The HAD Calorimeter is divided into similar sections covering the same η -regions and uses the same implementation for the end-cap and FCAL as the EM layer. The HAD barrel region is composed of steel and polystyrene tiles. Due to the two different implementations, these detectors are frequently referred to as the LAr and Tile Calorimeters, respectively. For further details about ATLAS see [1].

The LHC is designed to collide proton bunches at a rate of 40.08 MHz and, with an event size of 1.5 MB, ATLAS cannot record every event. The ATLAS Trigger reduces the event rate from this initial interaction rate by applying selections that limit recording to those events deemed interesting. The Trigger uses three rate reducing steps to accomplish this, which are called the Level-1 (L1), Level-2 (L2), and Event Filter (EF) Triggers. The L1 Trigger reduces the event rate from the LHC collision rate to 75 kHz, the L2 Trigger reduces this further to 5 kHz and, finally, the EF Trigger has an average output rate of 400 Hz with a maximum of ~ 1 kHz [1, 2].

The L1 Trigger is implemented using custom built electronics which is required to cope with the high collision rate. This is in contrast to the L2 and EF triggers which are implemented as software algorithms running on large computer farms. The L1 Trigger is mainly composed of the L1 Calorimeter (L1Calo) Trigger, L1 Muon (L1Muon) Trigger, and Central Trigger Processor (CTP). There are additional L1 triggers such as the beam pick-up triggers, the minimum bias scintillators, and the Zero Degree Calorimeter (ZDC), which uses the PreProcessor Module that will be described later. The CTP receives the data from all the L1 Systems and issues an L1 Accept (L1A) if predefined selections are met [1].

The L1Calo Trigger, diagrammed in figure 1, is designed to use spatially integrated calorimeter information (labeled *Analogue Sum* in the figure) to measure high energy electrons, photons, taus, well collimated jets of hadrons, and the total and missing E_T at the LHC bunch crossing (BC) rate. This task is divided among three main components, the PreProcessor (PP), the Cluster Processor (CP) and the Jet/Energy Processor (JEP). The 7168 summed analogue signals first arrive in the Receiver system where a linear gain may be applied. The signals then pass to the PP System which condition and digitise them at the LHC collision rate. The digitised signals are analysed to calculate the E_T and determine to which colliding bunch they belong. The results are sent in parallel to the CP and JEP Systems at the LHC BC frequency. Upon receiving an L1A from the CTP, the digitised signal data, E_T measurement, and BC identification (BCID) information are read out from the PP System and propagated to the ATLAS Data Acquisition (DAQ) System [3].

The CP performs a cluster finding algorithm and counts the number of electron/photon and tau/hadron clusters passing configurable thresholds. In this case, a cluster is defined as a well collimated energy deposit in the calorimeter where electrons would have most energy deposited in the EM calorimeter and taus would have energy in both the EM and HAD calorimeters. The JEP performs a similar algorithm and counts for jets which are defined as energy deposits in both

¹ATLAS uses a right-handed coordinate system with the origin at the nominal interaction point at the centre of the detector, z-axis along the beam pipe, x-axis pointing to the centre of the LHC ring, and y-axis pointing upward. Cylindrical coordinates (r, ϕ) are used in the transverse plane, ϕ being the azimuthal angle around the beam pipe. The pseudorapidity is defined in terms of the polar angle θ as $\eta = \ln \tan \theta / 2$. Transverse momentum and energy are defined as $p_T = p \sin \theta$ and $E_T = E \sin \theta$, respectively. The distance ΔR is defined as $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}$, where $\Delta\phi$ ($\Delta\eta$) is the separation in azimuthal angle (pseudorapidity).

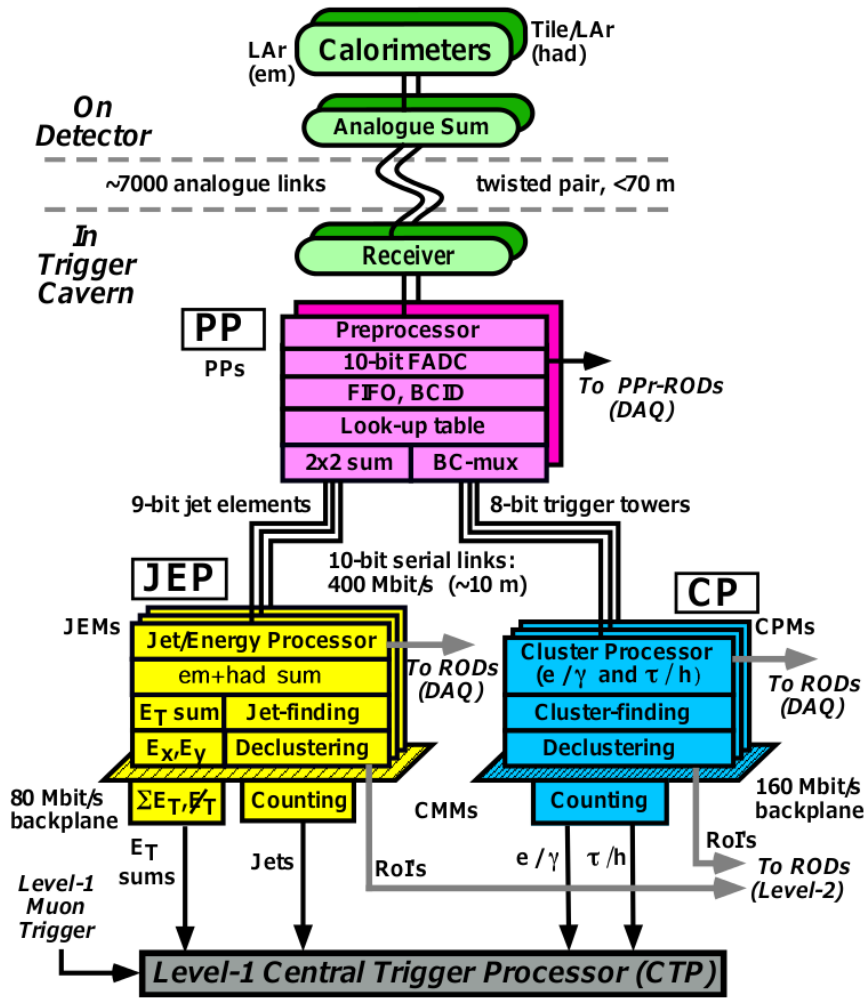


Figure 1. A diagram of the Level-1 Calorimeter Trigger architecture with the real time data flowing from top to bottom beginning in the Calorimeter detector front-ends and traveling via long signal cables to the L1Calo system.

calorimeters with a broader profile than those used by the CP. In addition, the JEP calculates the total and missing E_T , and missing E_T significance for each event. This information is sent to the CTP where it is combined with data from the L1Muon Trigger and other L1 subsystems in order to determine an L1 decision. The η/ϕ coordinate of each cluster or jet, called a Region of Interest (RoI), is transmitted to the L2 Trigger and used as a seed for the L2 reconstruction algorithms [4].

This paper presents the implementation, operation and performance of the PP System. First the hardware will be discussed, followed by the command and control mechanisms. Finally, the performance of the PP System will be presented using physics data provided by the LHC.

2 The L1Calo PreProcessor System

Within the L1Calo Trigger, the PP System is responsible for processing the LAr and Tile Calorimeter dedicated trigger output to produce a measurement of E_T , up to 250 GeV, which is then trans-




Figure 2. η - ϕ coverage of the PreProcessor System in one quadrant of the Electromagnetic Calorimeter. The coverage of a single PPM, Analogue Input (AnIn) board, Multi-Chip Module (MCM), and Trigger Tower (TT) are shown. Groups of 4×4 TTs with the same colour are contained within a single MCM and are summed to form a single Jet Element. For $2.9 < \eta < 3.2$ ($3.2 < \eta < 4.9$) only half (quarter) of a PPM is needed per quadrant due to the reduced granularity.

mited to the CP and JEP Systems for object identification. The PP must also identify the LHC bunch in which the signal was belongs, referred to as the BCID. These tasks must be performed at the LHC BC frequency of 40.08 MHz. In addition, the PP System must respond to an L1A, issued by the CTP, by reading out the event that triggered the L1A and transmitting the data to the ATLAS DAQ System.

The calorimeters contain over 200,000 channels, all of which cannot be processed at the LHC BC frequency, due to design limitations such as cost, physical space, and the required decision latency. Therefore, the calorimeter signals are summed within the calorimeter front-end electronics to form 7168 analogue signals, called Trigger Towers (TT). One TT can be the sum of up to 60 calorimeter cells. The typical dimensions of a TT are $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ with coarser resolution above $|\eta| > 2.5$. There are independent TTs for the EM and HAD layers.

L1Calo must provide triggers over the entire calorimeter coverage, $-4.9 < \eta < +4.9$ and 2π radians in ϕ . Figures 2 and 3 show the η - ϕ coverage for a single EM and HAD ϕ -quadrant for positive η . In the region $|\eta| < 2.4$, each PPM processes 64 TTs with the coverage of $4\text{TT} \times 16\text{TT}$ in $\Delta\eta \times \Delta\phi$ with the standard TT size of 0.1×0.1 . For $2.4 < |\eta| < 2.9$, only half of the channels on the PPM are used due to the varying TT sizes shown in the figures. Within $2.9 < |\eta| < 3.2$, one PPM covers two quadrants.

The remaining region, $3.2 < |\eta| < 4.9$, only requires one PPM to cover all four ϕ -quadrants. In this region, which receives signals from three FCAL layers, the TTs are different in size. The first layer (FCAL1) is nearer to the interaction point and is more granular than the two following layers (FCAL2 and FCAL3). In FCAL1, each TT is approximately 0.4×0.4 in $\Delta\eta \times \Delta\phi$, see figure 2. In the FCAL2 and FCAL3, the TTs are similar to those in FCAL1 except they are roughly twice the size in η , see figure 3.

The L1Calo PP System is made of 124 PreProcessor Modules (PPM) installed in eight standard VME crates with a 21-slot backplane. The VME bus follows the VME64x standard meaning that the top two thirds of the crate backplane are occupied by a VME64xP-VIPA backplane. The lower




Figure 3. η - ϕ coverage of the PreProcessor System in one quadrant of the Hadronic Calorimeter. Similar to figure 2 except the Forward Calorimeter has an additional layer and larger bins in η .

third is occupied by a custom backplane that a PPM uses to output the real-time data to the CP and JEP Systems. Each PP crate contains one network connected Single Board Computer (SBC) (slot 1), followed by three empty slots, 16 PPMs, and one Timing Control Module (TCM). Two crates only contain 14 PPMs because of the coarser granularity at higher $|\eta|$.

The SBC acts as a local controller for the VME bus and is used to communicate with and configure the PPMs via the VME J1 backplane connector. The SBC can also access all PPM data locally during data taking. The VME J2 backplane connector is used to output the PPM data to the ATLAS DAQ System via a dedicated optical data link which is described in section 3.5.

The TCM receives protocol signals (e.g. LHC clock, L1A, and more) from the ATLAS Timing, Trigger and Control (TTC) System via an optical link [4]. The VME J0 backplane distributes these signals using point-to-point routing on an auxiliary printed circuit board. The TCM is also master of the Control Area Network (CAN) bus in each crate, which is used to monitor environmental parameters important to the safe operation of the hardware. The CAN bus requires two serial lines across the slots on the backplane which are also implemented on the VME J0 backplane. The unoccupied pins of this backplane are used for servicing the PPM, e.g. reloading or reprogramming of the CAN micro-controller. The TCM front panel has a VME bus-display for visual diagnostics.

The power consumption was measured using a single, fully populated PP crate in a test lab setup that was operating all PPMs in the typical data processing mode, resulting in 175 A on the +3.3 V supply and 150 A on the +5.0 V supply. This corresponds to 84 W per PPM, which is well below the worst case estimate of 100 W [5]. In order to dissipate the heat created, fans at the bottom of the crate force a high flux air stream upwards.

3 The PreProcessor Module

The PPM is the primary component of the PP System and is shown in figure 4. The PPM is a printed circuit board with a height of 366 mm (9 NIM units) and a depth of 400 mm. It is made of eight layers (2.0 mm thick), five layers for routing and three for full or partial power planes. The




Figure 4. The PreProcessor Module with a view of the front panel at the left and the side on the right. Data flows, in general, from left to right in the side-view. Daughter boards and other components are outlined and described in the text.

daughter boards are easily seen and perform many of the important signal-processing tasks. In the event of a hardware failure, the problematic daughter board can be replaced instead of the full PPM.

A PPM can receive up to 64 TT signals, as described in section 2, which corresponds to an η - ϕ coverage of 0.4×1.6 . For the PPMs processing signals from the range $2.4 < |\eta| < 2.9$ and $2.9 < |\eta| < 3.2$, fewer signals are received and, therefore, some channels are unused.

In figure 4, the flow of data through the PPM is from left (front-panel) to right (backplane). Analogue signals arrive at the front panel where four Analogue Input Boards (AnIn), each processing 16 TTs, apply voltage offset and scaling to the input signals. The signals are routed to 16 Multi-Chip Modules (MCM), each processing four TTs. The MCM contains four Analogue-to-Digital Converters (ADC) to digitise the signals and a four channel Application Specific Integrated Circuit (ASIC). The PreProcessor ASIC (PPrASIC) is responsible for the bulk of the signal processing in the PPM, such as measuring E_T , performing BCID, and buffering data for readout. These data are transmitted to the CP and JEP via a Low-Voltage Differential Signalling (LVDS) Cable Driver card (LCD) which completes the real-time data flow through the PPM.

The Readout Manager (ReM) FPGA controls the readout of the data upon the arrival of an L1A provided by the CTP. The data from the 16 PPrASICs is read out and packed for serial transmission to the ReadOut Driver (ROD), via the Rear G-Link Transition Module (RGTM), the first step in the ATLAS DAQ System. The ReM is controlled over VME, using the SBC, to configure the PPM for proper data processing. The TTC Decoder card receives and decodes protocol signals from the ATLAS TTC System for use within PPM components. The CAN module collects and transmits environmental parameters to the ATLAS Detector Control System (DCS) [6].

There are two Complex Programmable Logic Devices (CPLD) on the PPM main board, the VME CPLD and the Flash CPLD. The VME CPLD is connected to the VME bus and forwards all traffic to the appropriate chips depending on the addresses requested. It also provides access to a readable unique module number and VME status register. The Flash CPLD controls loading of the configuration code, stored in an onboard Flash-RAM, into the ReM FPGA and four FPGAs located on the LCD. The Flash-RAM is large enough to store six versions of ReM FPGA configuration bit files and two versions for each LCD FPGA. The Flash CPLD can also be used to load configuration code to the FPGAs directly from the SBC, bypassing the flash memories.

Space on the front panel is very limited due to the large number of input signals. Nevertheless, the free space in the centre of the panel is allocated to visual status indicators (LEDs) and a matrix display element, which shows more complex information, i.e. module status coded as a single character.

The backplane is equipped with four connectors each having a high pin count, seen in figure 4. The connectors J0, J1, and J2 are standard VME backplane connections used for inter-board communication, and command-and-control of the PPMs using the SBC. A custom 2 mm CompactPCI connector is mounted in the remaining space and is a pass-through for the real-time signals output by the LCD. Any bending or misalignment at insertion time would damage the crate's backplane, so a mechanical stiffening is added for insertion and extraction. As seen in figure 4, vertical bars and horizontal rods give the board good rigidity, when force is applied at insertion into the backplane. Electromechanics complies with IEEE standard 1101.10 [7]. Reliable long-term operation is ensured by screw-fixation of removable parts (e.g. connectors, daughterboards).

3.1 Analogue input boards

A single PPM receives up to 64 differential analogue signals from the calorimeters on four 16-way twisted-pair cables each carrying 16 TTs. Example signals from the LAr and Tile Calorimeters are shown in figure 5. The LAr and Tile front-end electronics, which produce these signals, construct a signal whose peak is linearly proportional to the total E_T contained in the TT. The signal voltage to E_T conversion is calibrated such that $1 \text{ V} \approx 100 \text{ GeV}$. The LAr signals, which account for about 75% of all signals processed by the PPM, are bi-polar. This means the signal contains a 'fast' half-wave, used for triggering purposes, and a 'slow' half-wave of opposite polarity but equal area [8]. Signals from Tile are uni-polar as is typical for Photo-Multiplier tubes [9]. The input connectors are specially designed to minimise cross-talk. Each connector has a pair of fixation screws to ensure proper electrical contact and protect against accidental disconnection.

The incoming analogue signals are routed directly to the AnIn daughter boards. There are four such boards on the module matching up to the connectors on the front panel. The daughter board is interfaced to the PPM using the Common Mezzanine Card (CMC) standard [10].




Figure 5. Example Trigger Tower signals arriving at the PreProcessor Module from the (a) Liquid Argon EM calorimeter layer and (b) Tile Hadronic calorimeter layer.

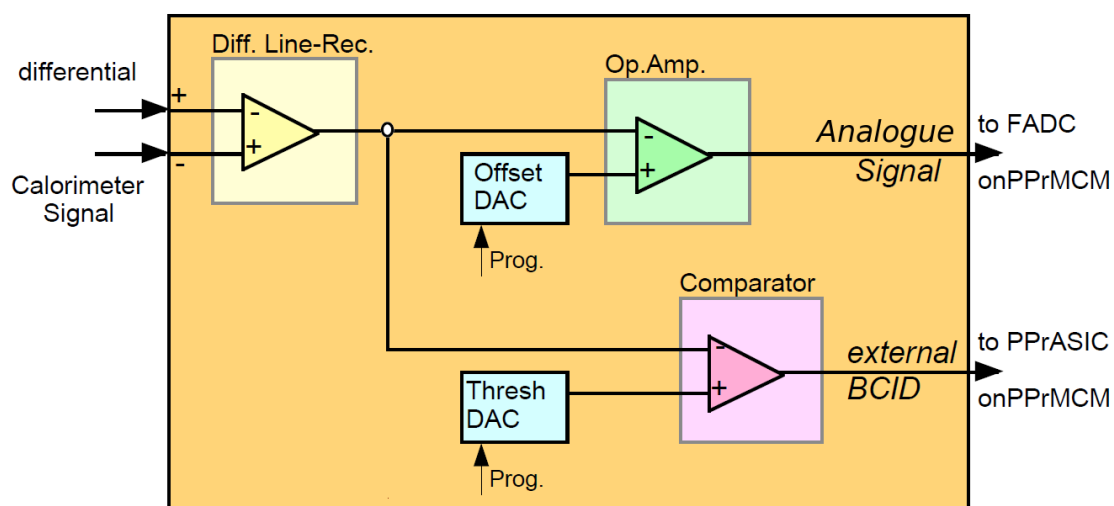


Figure 6. Block diagram of the signal conditioning and external BCID performed on the Analogue Input daughter board for a single TT signal.

A diagram of the AnIn functionality is shown in figure 6 for a single channel. The differential analogue signal arrives from the PPM front panel connector at the input of a differential line receiver which forms a single ended signal. An SPI-programmable, 8-bit Digital-to-Analogue Converter (DAC) produces a voltage level which is added to the input signal. This voltage offset is necessary to ensure that the signal is within the linear digitisation window of the ADC located on the MCM. The analogue signal is also rescaled by an operational amplifier to match this linear digitisation range of the ADC which has a size of ~ 1 V.

Using a LAr signal, three examples of the effects of signal conditioning are illustrated in figure 7. The voltage offset, referred to as the DAC-offset, is represented at the bottom (in green). The digital offset can be set from 0-255 with DAC=0 corresponding to 1.76 V and DAC=255 to 2.28 V. The digitisation window of the ADC is also shown with a range of 1.9-2.9 V (in blue). The

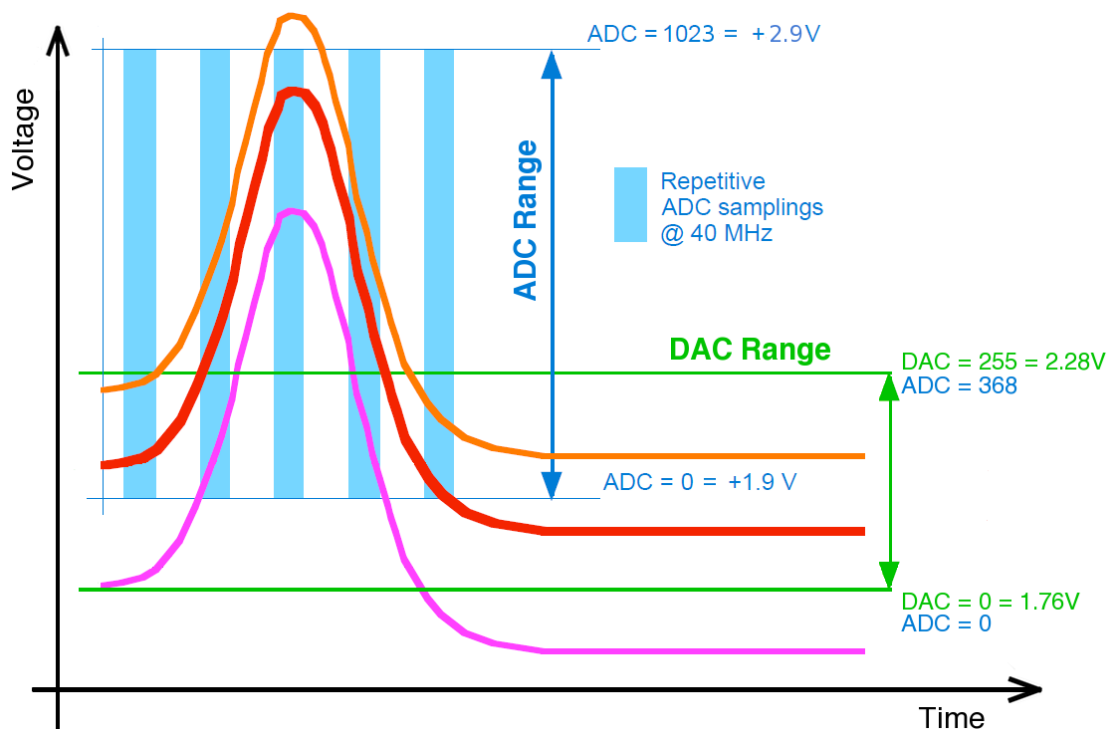


Figure 7. Three example Liquid Argon signals are shown, each with different voltage offsets applied using the DAC onboard the AnIn. The pedestal must be adjusted to ensure the signal is within the 1.0 V window of the ADC located on the MCM.

lowest signal example (purple) shows the case where the signal pedestal is below this window and therefore acts as an effective threshold, eliminating excessive noise if needed. The central example (red) places the signal baseline just above the lower limit of the ADC window, which allows a measurement of the noise distribution before and after the signal. This is the DAC-offset model used during physics data taking because it allows for the largest digitisation range. The highest signal (orange) gives an example where the LAr signal undershoot is fully contained within the digitisation window, allowing it to be measured if necessary.

The analogue signal is propagated in parallel to a discriminator, which sets a digital output to the logic of ‘1’ when the signal crosses a programmable voltage threshold. The threshold is adjustable for each channel independently using a second DAC. The digital output is routed to the MCM and used by the PPrASIC to mark the corresponding LHC clock cycle. This constitutes one of the three implemented methods to perform BCID, called the *External BCID*, which will be discussed in more detail in section 3.2.

3.2 The multi-chip module

The MCM contains most of the signal processing components of the PPM. Each MCM processes four TTs. The main MCM components are the PHOS4 chip, the four ADCs, the PPrASIC, and the three LVDS Serialisers as shown in figure 8.

The signal from the AnIn boards is propagated to the MCM where it is first digitised in the ADC. The LHC clock (40.08 MHz) is used as the strobe in the ADC which therefore provides a




Figure 8. The bonded components on the MCM (brass-lid removed), showing the (A) PHOS4, (B) ADCs, (C) PPrASIC, and (D) LVDS Serialisers.

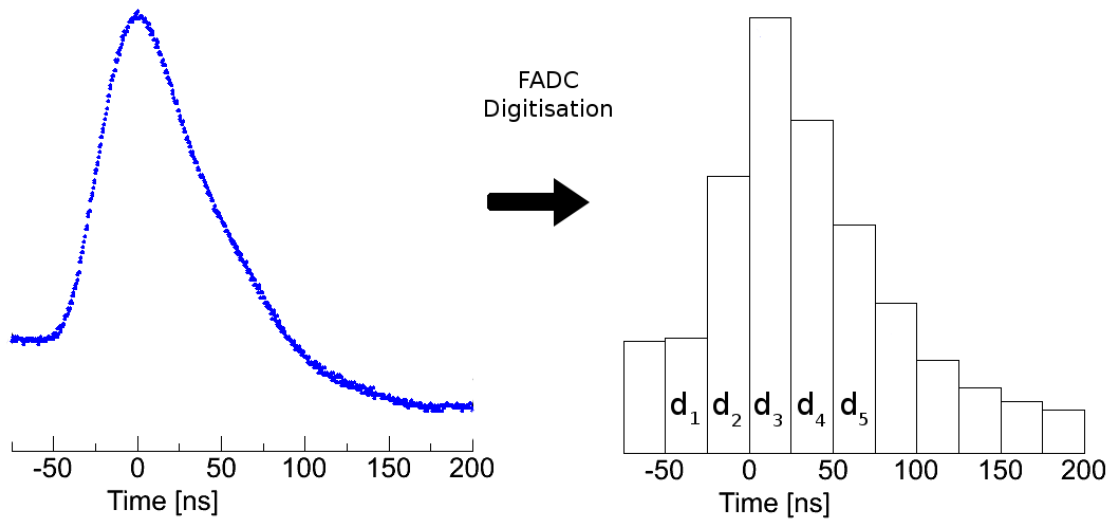


Figure 9. Example of the ADC digitisation. Right: analogue pulse; left: digitised pulse.

sample of the signal every ~ 25 ns as shown in figure 9. The strobe can be delayed by up to 24 ns, in 1 ns steps, using the PHOS4 chip [11]. The input signal spans multiple BCs such that the peak can fall between two digitised samples which would negatively affect the E_T measurement. Therefore, this nanosecond delay must be calibrated to ensure the signal peak is sampled by the ADC. This is called the *fine-timing* calibration and is discussed in section 5.2. One PHOS4 chip can provide four independent delays, hence, only one is needed on each MCM.

The digitised results from the four ADCs are clocked in to the PPrASIC at the LHC bunch frequency where the data is used for the E_T measurement and BCID (for details see the next section). While the CP uses E_T and BCID information from each TT, the JEP uses reduced granularity Jet Elements of 0.2×0.2 in $\eta \times \phi$ (see figures 2 and 3) which corresponds to the four channels on a single MCM. The data are output at the bunch frequency to three commercial LVDS transmitters which serialise the three 10-bit outputs from the PPrASIC at a rate of 480 Mbits/second (including protocol bits). The three serial streams are routed to the LCD via the motherboard (details in section 3.4).

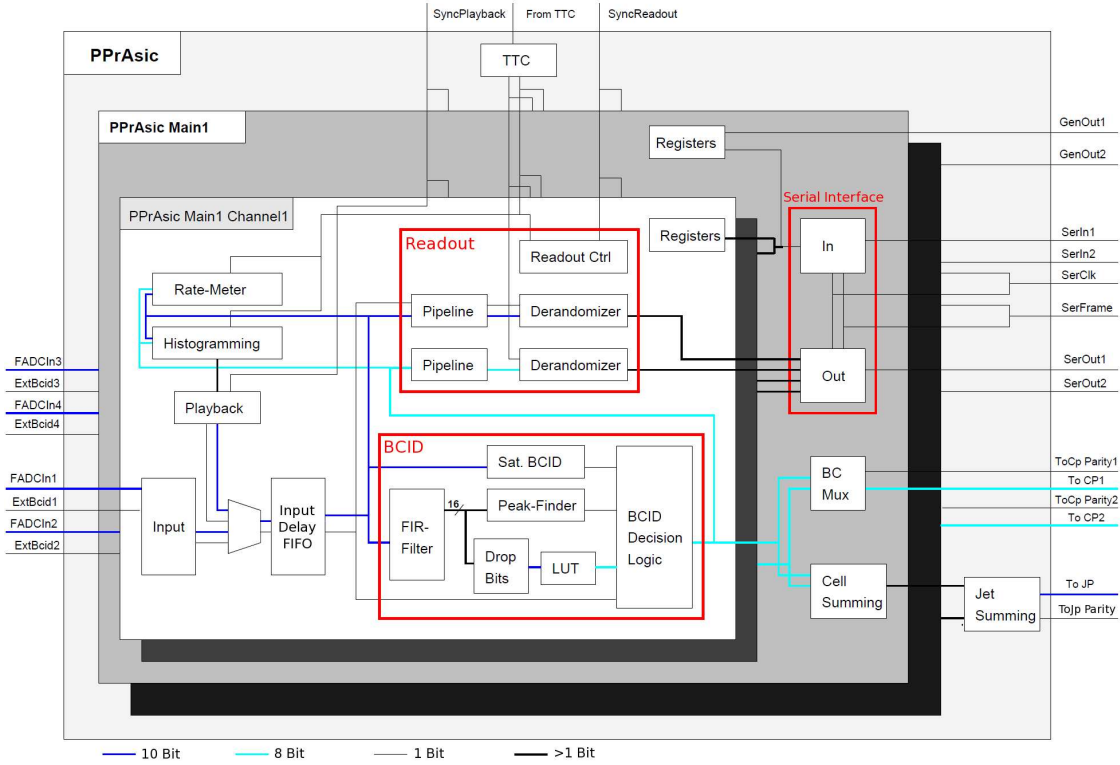


Figure 10. Diagram of the PPrASIC logic for one channel.

The data, read out upon receiving an L1A, are passed directly from the PPrASIC to the ReM FPGA on the motherboard via a serial communication bus, which is discussed in section 3.5.

3.3 The PreProcessor ASIC

The PPrASIC processes the output from all four ADCs in parallel. Figure 10 shows the processing details for one channel inside the ASIC with input signals on the left and output on the right. The 10-bit ADC values enter an Input First-In, First-Out (FIFO) memory with a maximum depth of 16 locations. The FIFO depth is programmable for each channel and acts effectively as a delay to allow for the alignment of calorimeter signals arriving at different times due to differences in transmission latency, caused by variable cable lengths and detector geometry effects. This setting is referred to as the *coarse-timing* calibration. The *External BCID* arrives in a similar 1-bit FIFO which acts to delay its arrival in the PPrASIC to ensure the logic bit is relayed in the correct BC.

The BCID and E_T measurement (real-time processing) and data readout (data acquisition) are the two most important functions of the PPrASIC, as highlighted in figure 10. The buffered ADC data is used to perform a BCID and E_T measurement and is also sent to a secondary memory for later readout. The results of the BCID and E_T measurement are also stored for readout. In the event of a positive BCID, the E_T and four-channel sum of the E_T are propagated to the CP and JEP Systems, respectively.

There are also monitoring and stand-alone testing capabilities built in to the PPrASIC, such as Rate Metering, Energy Histogram, and Playback memory. Communication with the PPrASIC

is performed via two serial interfaces, each of which controls two of the four channels. Details are discussed in the following sections.

3.3.1 Bunch crossing identification and E_T measurement

The first BCID method, the *External BCID*, was discussed in section 3.1 and raises a digital logic of ‘1’ if the analogue signal passes a programmable voltage threshold. The 1-bit result is synchronised to the proper BC using the Input FIFO, and an additional delay (not shown in figure 10) can be added before it arrives in the BCID logic block (center-bottom of the same figure) to ensure it is in time with the *Peak-Finder* and *Saturated BCID* results.

The *Peak-Finder BCID* is the second BCID algorithm. It begins by routing the 10-bit values from the FIFO to a Finite-Impulse-Response (FIR) Filter. The FIR Filter buffers five ADC values, d_i , one from the current BC, and two preceding and two following BCs. These are used with five programmable coefficients, a_i , to produce an output value of 16-bits using equation (3.1).

$$f = \sum_{i=1}^5 d_i a_i \quad (3.1)$$

The five coefficients are tuned to the input signal shape on a per TT basis which acts to suppress noise. This tuning results in values of a_i that are optimised to the signal-to-noise ratio contained to each bin, i , for well behaved physics pulses.

The FIR output value, f , is transmitted in parallel to the Peak-Finder logic as well as the Drop-Bits and Look-Up-Table (LUT). The Peak-Finder performs the logic $f_{-1} \leq f_0 < f_{+1}$ where f_0 is the current FIR output, f_{+1} the one from the preceding, and f_{-1} from the following. The logic can also be programmed to use $f_{-1} < f_0 < f_{+1}$, however, the former relation is used in standard running conditions. This is the main method for BCID used for signals that are not saturated, i.e. having $d_i < 1023$. The output of this logic is transmitted to the BCID Decision Logic block and combined with the two complimentary BCID algorithms, the *External BCID* and the *Saturated BCID* (discussed below).

Following the parallel path in figure 10, the 16-bit FIR output is first reduced to a 10-bit value by dropping a programmable number of consecutive bits. Depending on the FIR Filter Coefficients, these ten bits are chosen to maximise the E_T resolution. The LUT contains a programmable table that transforms a 10-bit input to an 8-bit output representing the TT E_T measurement. This 8-bit LUT-value is sent to the BCID logic block and is propagated if the BCID logic is satisfied.

The third, and final, method for BCID is the *Saturated BCID* algorithm which operates directly on the 10-bit ADC output. This algorithm uses three programmable values, the Saturation Level, D_{sat} , the Saturation High Threshold, D_{hi} , and the Saturation Low Threshold, D_{lo} . The algorithm is evaluated if the ADC output, d_i , satisfies $d_i \geq D_{sat}$, thus, D_{sat} defines a saturated signal, as shown in figure 11. The algorithm is only performed on the first saturated ADC output and rearms after $d_i < D_{sat}$ to ensure the identification is only performed once per signal. The *Saturated BCID* can assign the identification to the first, d_s , or the second, d_{s+1} , saturated ADC value based on the following logic: if $d_{s-1} > D_{hi}$ and $d_{s-2} > D_{lo}$ then d_s is the identified BC, otherwise, it is assigned to d_{s+1} . The algorithm is making an approximate measurement of the rising edge of the trigger signal to determine the correct BC, and the thresholds, D_{hi} and D_{lo} , must be calibrated to perform the identification properly.




Figure 11. Diagram of a saturated, digitised signal with three thresholds (D_{sat} , D_{hi} , D_{lo}) that define the *Saturated BCID* algorithm.

The three BCID algorithms are evaluated on every BC, producing a logical output of ‘1’ for success or ‘0’ for failure, and arrive in the BCID logic block where they are combined to form an overall decision. In addition, the E_T range (output from the LUT) of $0 \text{ GeV} < E_T < 255 \text{ GeV}$ is divided into three regions and each region can be programmed to require any combination of the three BCID algorithms. For standard data taking, the *Peak-Finder* BCID is required for the low and middle energy regions, while the *Peak-Finder* and *Saturated* BCIDs are required for the high energy region. If the BCID logic satisfies the required combination for a given pulse, the E_T is transmitted to the CP and JEP Systems.

For transmission to the CP, which uses the full 0.1×0.1 TT granularity, the LUT outputs from two channels are multiplexed which halves the number of cables needed. This is possible since, after a successful BCID, the following BC cannot contain another local maximum, so any non-zero value is always flanked by two zeros. To allow for proper unpacking of the multiplexed data upon reception in the CP, a *mux-bit* is added to indicate both the channel and BC to which the 8-bit data belongs. An odd parity bit is also added, making 10-bits in total, to allow for transmission error detection. For transmission to the JEP, which uses a reduced 0.2×0.2 granularity, the four 8-bit LUT values are summed to form a 9-bit Jet Element. An odd parity bit is also included to form a 10-bit word. The three data streams are then sent to the serializers on the MCM for transmission to the CP and JEP.

3.3.2 Readout and diagnostic functions

The readout of the ADC, E_T and BCID results is another primary task of the PPrASIC, also highlighted in figure 10. These data are not only important for validating the BCID logic and E_T measurement, but are required to properly calibrate the PPM hardware. The readout of the ADC data is also important because the analogue trigger signals are built in the front-end electronics of the calorimeters and are accessible for readout only in the PP System. The number of ADC samples

and BCID results (meaning LUT output and the logic bits from all three BCID algorithms) that are read out are independently programmable. During standard data taking, five ADC samples (centred on the peak similar to d_1-d_5 in figure 9), and one BCID result are read out, which constitutes the minimum amount of data needed to validate the processing logic in the PPrASIC. This is referred to as 5+1 readout mode and yields 276 bits of data from a single TT per L1A. Other readout modes exist, such as 15+1 readout which is used for calibration data taking. In the lab, the data is read out locally over the VME backplane in 127+7 mode for stand-alone testing and diagnostics.

The readout path uses two 11-bit FIFOs, labeled *Pipeline* in figure 10, with a depth of 128 locations in order to store the latest 10-bit ADC values and 8-bit LUT values. The extra bit in the ADC FIFO is used to store the *External BCID* value and the three extra bits in the LUT FIFO are used to store the three BCID algorithm bits. The FIFOs provide a window of $128 \times 25\text{ns} = 3.2\ \mu\text{s}$ from the time the values enter the FIFO to when they are lost. The L1A must arrive from the CTP within this window, or the data is lost. The design latency limit for the L1 Trigger is $2.0\ \mu\text{s}$ with an additional $0.5\ \mu\text{s}$ contingency which is well within the latency of the readout FIFOs. The L1A causes event-related data in both of the *Pipeline* memories to be copied to their respective *Derandomizer* memories where it is stored for readout. The data from two channels are serialised to form one data stream in the Serial Interface, making two streams per ASIC that are sent to the ReM FPGA for further processing.

Stand-alone testing and statistical monitoring are also foreseen in the PPrASIC design by inclusion of an additional memory per TT (labeled *Histogramming* in figure 10) with 256 11-bit memory locations. During ATLAS data taking this memory is used to histogram either the ADC or LUT output above a programmable threshold, which can be read out between L1As in order to monitor single channel noise and provide an unbiased signal spectrum. For testing, it can be loaded with a data pattern that is then propagated to the BCID logic block and further to the CP and JEP in order to test the full L1Calo System with known and programmable data when calorimeter input signals are not available.

Another important monitoring capability of the PPrASIC is the Rate Metering which is used during physics data taking to measure the signal rate at the L1Calo inputs. This functionality can be programmed to use the ADC or LUT as input. A 20-bit data counter is used to record the number of samples greater than a programmable threshold. The counting stops when the 20-bit value is saturated or a programmable, 16-bit time limit is reached. Then a status bit is set to alert the user that the counters are ready to be read out.

3.4 The LVDS cable driver

As discussed in section 3.3.1, each MCM produces three high frequency signals, i.e 16 MCMs produce 48 in total, which are routed via the motherboard to the LCD daughter board. Special routing techniques are applied for these signals to ensure error-free transmission. An impedance-matched strip-line together with pre-emphasis of the digital signal at the source is used to fulfil the task.

The architecture of the L1Calo System requires TT signals on the azimuthal boundaries of each module to be duplicated and sent on separate cable links [5]. Furthermore, the PPM must have the capability to drive the cable link a significant length ($\sim 11\ \text{m}$) to the downstream processor crates. The LCD performs this task before the signals are taken to a backplane connector.

The duplication of signals is done inside four FPGAs (X2CV250), where the output is sent in parallel to two output drivers. It is necessary to pre-compensate amplitude losses due to the integration properties of the parallel pair cable transporting data to distant processor crates. The required passive components (R, C) are placed close to the FPGA outputs. The dimensioning of the pre-compensation network (R, C values) is optimised for cables with a length of 10 to 15 meters.

The signals then leave the PPM via a custom backplane which is just a feed-through to the rear of the crate where LVDS cables are connected and secured.

3.5 The readout manager FPGA

The ReM FPGA, seen in figure 4, is a Xilinx XCV1000-E. It is responsible for many important functions within a PPM, including the primary function for which it is named, managing the readout of event data from all 64 TTs on the PPM.

As described in section 3.3.2, there are $2 \times 16 = 32$ serial links output by the PPrASICs. The ReM FPGA reformats this data in preparation for transmission to the ATLAS DAQ. There are 16 input streams to the RGTM module which are connected via the VME J2 backplane connector. The data from four channels are serialised by the ReM FPGA such that each input stream carries the data from one PPrASIC. These input streams are serialised on the RGTM and transmitted via optical G-Link (960 Mbits/second) to the ROD, the first step in the ATLAS DAQ System.

The ReM FPGA is also the primary control and communication hub for the PPM and its functionality will be further discussed in section 4.

3.6 The timing, trigger and control decoder module

The PPM, as a pipelined device, is driven by the LHC clock. The readout facilities require protocol signals such as the L1A and an Bunch Counter Reset (BCR). Synchronisation of internal counters, such as the Bunch-Crossing counter or the L1 Event counter, is given by periodic RESET signals. Furthermore, a synchronous START/STOP signal is required for stand-alone testing using the playback memories.

These signals are provided by the TTC System [12] via optical links to each of the eight PP crates. A TCM takes in the optical signal stream and provides electrical output for point-to-point distribution to the 16 PPMs. The signals arrive on the auxiliary backplane in a daughter board called the TTC Decoder seen in figure 4. The protocol stream is unpacked and individual signals are distributed to their destinations on the motherboard via the ReM FPGA. Access to parameter settings on the TTC Decoder is given through a separate Inter-Integrated Circuit (I²C) bus controlled by the ReM FPGA, which maps the registers to VME.

3.7 Module protection, control and monitoring

The operating and environmental conditions of the electronics are continuously monitored by the ATLAS DCS to ensure the safety of the hardware. If anomalous conditions arise, the DCS will take actions to protect the hardware, such as cutting power to hardware due to temperatures leaving the acceptable range. The DCS uses an industry standardised “slow-control” system CAN-Bus, which is an automated, multi-master, broadcast serial bus for connecting electronic control units [13]. Voltages, currents, and temperatures are monitored to determine the safety of the system.

From input connector to input of MCM	2
Input to output of MCM	13
Input of LCD to the PPM backplane connector	1
Input to output of PPM	16

Table 1. PPM signal latency across PPM components in units of LHC BCs (~ 25 ns).

The infrastructure of a crate (supply voltages, supply currents, cooling fan speed) can be controlled via the system CAN-Bus. The TCM acts as CAN master in a PP crate, interfacing the DCS with the internal bus in the crate, where the individual modules are identified by their slot number. Reliable system operation requires also monitoring of quantities on the module level, i.e. temperatures, operating voltages, supply currents, and supply voltages produced on-board for special components. This is achieved on the CAN daughter module with a CAN-Bus interface, which transmits the module information to the TCM.

The CAN daughter board consists mainly of a chip-size controller running CAN-Bus embedded software. It receives some quantities directly in digital form (e.g. MCM temperatures). Others are digitised on the CAN-board itself (e.g. supply voltages). All data are transmitted to the CAN master over bus lines on the auxiliary backplane installed on the VME-J0 connectors (see figure 4).

Each PPM is fused on board for each supply voltage from the crate to prevent damage to the backplane in case of a fault on the module itself. Passive fuses would produce a high voltage drop at the required currents, therefore an active solution is chosen. Power ON/OFF is handled by a hot-swap controller connected to a switch in the handle of the front panel, and also connected to the CAN-bus controller. The controller monitors the ramping of each voltage and unforeseen conditions lead to the module being switched OFF.

Many critical quantities monitored via the CAN-bus are digitised by an additional micro-controller on the PPM, an ATMEGA16. These include power supply voltages, the voltage across a temperature-measuring diode on the ReM FPGA, and the voltages across temperature-measuring diodes on the 16 MCMs. In addition, this micro-controller drives the 7*5 matrix LED display which is used to indicate the current state of the PPM such as 'C' meaning the PPM has been configured for data taking, or 'L' meaning the ReM FPGA has been loaded.

3.8 Processing latency measurements

Since the PP is a pipelined system, the trigger decision latency is very important to consider in the design. The latency includes the processing time in the detector front-end electronics, analogue signal transmission to the L1 Trigger Systems, processing time within the L1 Trigger Systems, and L1A signal propagation time to detectors and trigger systems in order to initiate readout. The detailed results of detector measurements are held in the electronics of the detector systems until the trigger decision arrives causing readout of the data. The maximum time for data storage on detectors is $\sim 2 \mu\text{s}$ (plus $\sim 0.5 \mu\text{s}$ contingency). Hence, a trigger decision must be propagated back within this time interval or detector data are lost.

The propagation time of a given analogue input signal (counted from the signal's peak) has been measured within the components of the PPM and is listed in table 1 in units of LHC clock ticks (~ 25 ns). The total latency contribution of the PPM to the Level-1 pipelined system is 16 LHC

clock cycles ($16 \times 25 \text{ ns} = 400 \text{ ns}$) as compared to 17 estimated in the TDR [5]. The data to the CP are not subject to the 4-cell summing in the ASIC hardware. Hence, the CP data depart one clock tick earlier. The latency measurements have been verified in 2008 using final production modules.

4 PreProcessor Module control and operation

The PPM is a standard VME slave module, and the SBC controls and communicates with it over the VME backplane. The VME bus provides full read/write access to the PPM, which is completely mapped to an 8 MB VME address space. All control and status registers, and readout data from the PPrASIC, are accessible to the SBC.

On board the PPM, the VME bus protocol is maintained by permanent code residing in the VME CPLD (see section 3), and all VME communication is routed through this chip. At power up, the PPM is not fully operational as the ReM FPGA and LCD FPGAs require loading of the configuration bit files, which are stored in the Flash-RAM chip installed on the PPM motherboard. The Flash CPLD controls the loading of these FPGAs and these control registers are mapped to the VME address space. Once the FPGAs are loaded the components of the PPM are ready to be configured.

4.1 The readout manager

The top priority of the ReM FPGA is receiving, formatting, and merging the event data from the 16 PPrASICs for transmission to the ROD. Otherwise, the ReM FPGA acts as the central control and operation hub of the installed daughter modules. There are many important components of the ReM FPGA, as shown in figure 12, which will be described below.

The VME Manager (top in figure 12) connects the ReM FPGA to the VME bus. Software applications, running on the SBC, send configuration data via the VME bus to the VME Manager, which then distributes the data to on-board destinations. Command and control registers are provided to initiate certain tasks, such as readout of configuration data, PPrASIC monitoring data or PPrASIC event related data over VME. The ReM FPGA gathers requested information and stores it in VME accessible registers. The VME Manager also collects bitwise status and error data, received from the interfaced devices or generated by internal algorithms, necessary to monitor the PPM operation.

The AnIn Manager (top-left in figure 12) provides four unidirectional Serial Peripheral Interface (SPI) data buses for transferring configuration data to the DACs mounted on the four AnIn boards. The ASIC Manager uses 32 bidirectional serial data links to interface the ReM FPGA to the 16 PPrASICs. The ReM FPGA writes trigger configuration data to the PPrASICs and receives readout, status, and monitoring data via these serial streams. During data taking, event data are reconstituted into 13-bit words and stored in dual ported memories. The I²C Manager uses two I²C data buses to communicate with the 16 PHOS4 chips and the TTC Receiver Chip (TTCrx), located on the TTC Decoder daughter board. The ReM FPGA transfers configuration data to these devices, and reads back configuration data only from the TTCrx (the PHOS4 is not readable). The MCM Control Signals module transmits control signals to components of the MCM, for instance the START/STOP signals used to synchronise calibration runs (see section 3.6).

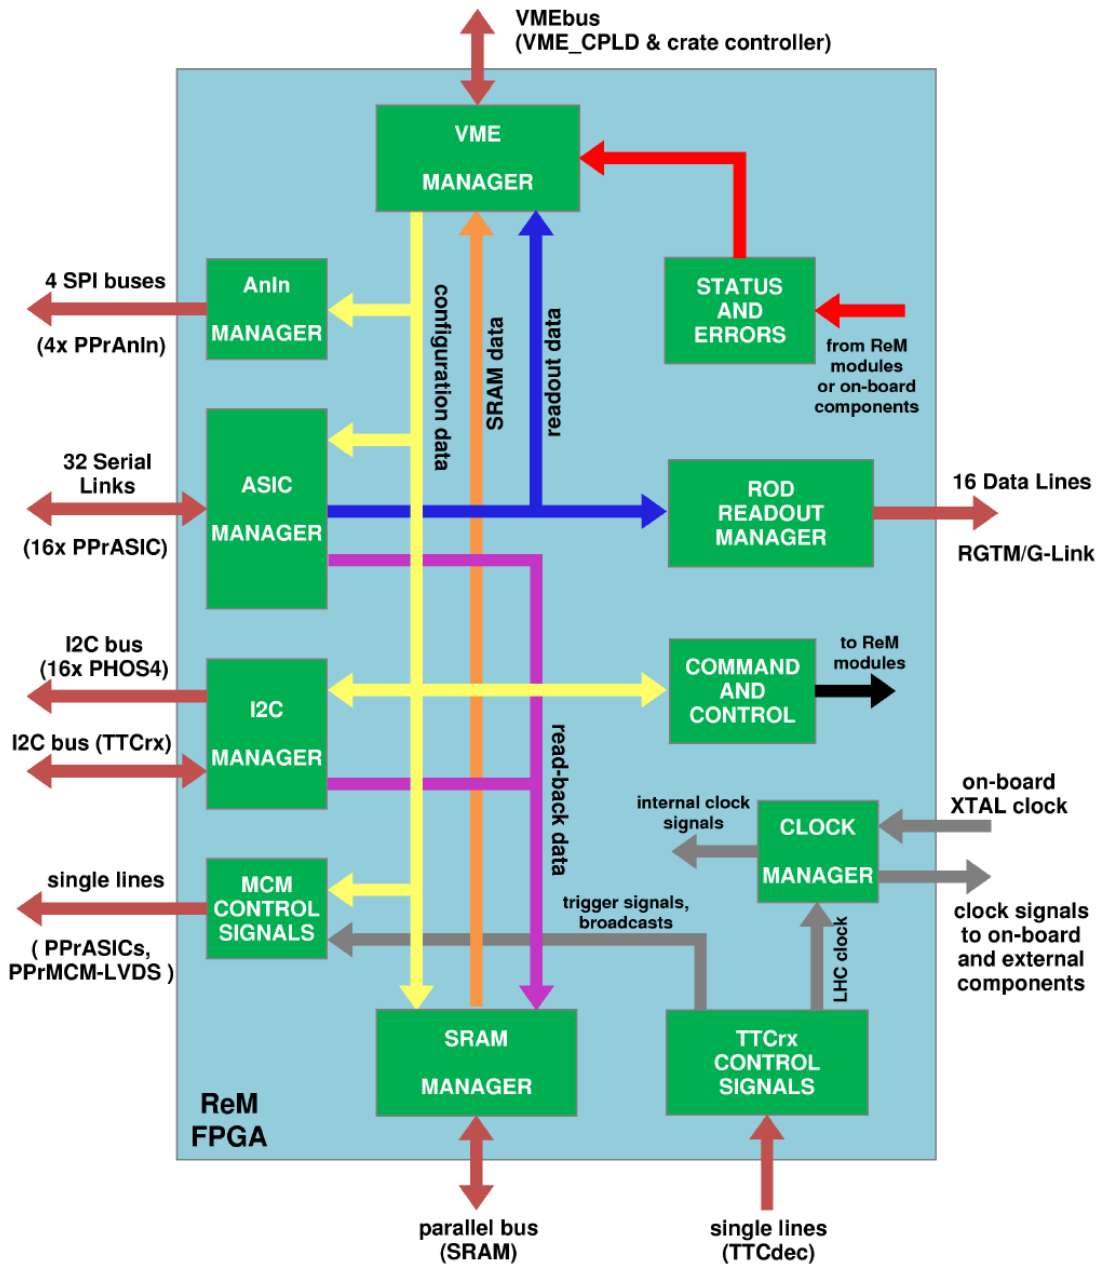


Figure 12. Diagram of the Readout Manager FPGA design implementation.

The ROD Readout Manager (top-right in figure 12) retrieves the event data from the dual ported memories and merges four channels into one serial data stream. 16 unidirectional data lines transfer the processed PPrASIC data to the ROD, via the RGTM. The Command and Control module handles commands addressed to the ReM FPGA and implements the requested actions.

The TTCrx Control Signals module (bottom-right in figure 12) receives protocol signals (i.e. LHC clock, bunch counter reset, event counter reset, L1A) from the TTC Decoder card. Most of these signals are fanned out to other modules in the ReM FPGA and to the PPrASICs. The

Clock Manager can use the LHC clock, when available, or a clock that runs at 40.00 MHz and is located on the TTC Decoder. This clock is mainly used for stand-alone testing. The Clock Manager provides the input clock to the many PPM components and derives additional clocks for the I²C and SPI buses which require 2 MHz and 100 kHz frequencies, respectively.

One bidirectional parallel data bus connects the ReM FPGA with an on-board Static RAM (SRAM) device. The 4 MB SRAM acts as a physical extension to the ReM FPGAs internal memory resources. The SRAM is used to verify proper loading of configuration data and to provide VME access to monitoring data from the PPrASICs. Configuration data received by the VME Manager is written to SRAM by the SRAM Manager and used as a *reference*. When the SBC requests configuration data to be read from the hardware, the requested information is copied to a second location in the SRAM referred to as the *read-back* block. Configuration data can then be validated by comparing the *reference* with the *read-back*, and if they are not equal an error is reported. For monitoring histograms, each of which consist of 256 data words, the user sets a bit in a ReM FPGA command register to initiate the readout from all 64 channels. The monitoring data is stored in the SRAM and when retrieval is complete a flag is set to indicate reading over VME can proceed.

4.2 Physics data taking

The standard operating mode of the PPM is the data taking mode, where signal processing and data acquisition are the priority tasks of the PPM. To review, this involves the following steps:

- analogue TT signals, composed of energy deposits in multiple calorimeter cells, are conditioned in the AnIn board,
- these signals are then digitised in the ADC at the LHC clock frequency of 40 MHz,
- using the digitised signals, the PPrASIC performs the BCID and E_T measurements and the results are transmitted in real-time to the CP and JEP Systems,
- and, finally, when an L1A is received, event data is sent from the PPrASIC to the ReM FPGA, which packs the data and transmits them to the ATLAS DAQ System.

To reach this mode of operation, after powering a crate of PPMs, the controlling software must initiate the loading of the configuration bit files into the ReM and LCD FPGAs. Then the components of the PPM can be configured by accessing the many registers in the ReM FPGA. The ReM FPGA can be set to *DAQ Mode*, which means, as long as this mode is active, all configuration requests related to the real time or readout data paths are denied to ensure stable operating conditions. The Rate Metering and Histograms are exempt from this restriction since their readout is decoupled from the event readout.

When an L1A arrives in the PPrASIC, the ADC, BCID, and E_T results are copied from the *Pipeline* memories to the *Derandomiser* memories. Once the data is available in the *Derandomiser*, the PPrASIC transfers the data to ReM FPGA via the serialiser interfaces. The readout of physics data is given priority while other processes, such as readout of rate metering and histogram data, are lower priority and are initiated by requests sent via the VME backplane. The PPrASIC sends a data bit to the ReM FPGA every LHC clock cycle with a data word size of 13 bits. In the absence of event data and diagnostic data, a status word is transmitted with information about the PPrASIC

Calibration runs without calorimeter signals	
DAC Scan	Fixes the signal pedestal to the same value for all channels (default is 32 ADC units).
Pedestal Scan	Measure the mean value and width of the signal pedestal distribution for each TT.
Playback Test	Test PPrASIC real-time processing and readout using simulated signal patterns. Stand-alone testing of upstream processors.
Calibration runs with calorimeter signals	
Energy Scan	Compare the E_T measured in the PPM with the reconstructed E_T using calorimeter data at multiple energies.
PHOS4 Scan	Determine the proper PHOS4 setting to ensure the signal peak is digitised optimally.

Table 2. List of typical calibration runs.

status, such as if the histograms and rate metering are available for readout. In the event of receiving an L1A or a request for diagnostic data, the PPrASIC will transmit the data, and upon completion, return to transmitting status words. If an L1A is received while diagnostic data is being readout, the event data is immediately transmitted, and any remaining diagnostic data follows the last word of event data.

The readout process in the ReM FPGA is driven by the data received from the PPrASIC. The data arrives in the ReM FPGA via 32 serial links where the data from each link is buffered in dual ported memories. Physics data is formatted according to the ATLAS DAQ standard and transmitted to the ROD via the RGTM installed on the backplane whereas diagnostic data is removed and made available for readout over the VME backplane. There are six readout modes available: 5+1 (5 ADC samples and 1 E_T), which is used in standard physics data taking; 3+1 may be used for high luminosity data taking; and 7+1, 9+3, 11+5, and 15+1 which are used for special calibration data taking to evaluate the input signal shapes and BCID performance.

4.3 Calibration data taking

Some of the PPM settings, such as those related to signal pedestal and timing, require special runs. The calorimeters can provide test signals at fixed energies which are used in some of these calibration runs. Table 2 lists the standard calibration runs that are taken on a regular basis to check the stability of the PPM.

The DAC and Pedestal Scans are only relevant to the PPM and do not require calorimeter signals. In fact, typically the inputs to the PPM are disabled to suppress any noise coming from the electronics upstream. The DAC Scan takes a fixed number of events at each DAC (this refers to the 'offset DAC' in figure 6) setting in order to extract the linear dependence between the DAC and digitised ADC values. This is used to fix the DAC such that the zero voltage level is the same for all channels. In addition to calibration, single bit errors, due to problems in the ADC, can be identified in the DAC Scan as well as problems with the DAC itself.

The Pedestal Scan is typically run after the DAC Scan, and simply takes a fixed number of events in order to measure the pedestal mean and sigma. Problematic channels can also be identified

in this test via broad pedestal or multi-peaked distributions. Any problematic MCM identified by these tests can be easily replaced due to the modularity of the PPM.

The Playback Test uses the 256-words long memories in the PPrASIC which can be loaded with any pattern necessary to test the logic in the PPrASIC or even the CP and JEP Systems. These 10-bit values are then fed directly into the PPrASIC instead of using the ADC output. For the playback of data to begin synchronously across multiple crates, START and STOP signals are transmitted via the TTC System.

The Energy Scan ensures that the energy measured by L1Calo agrees with the reconstructed energy from the calorimeter data. This is done using calibration pulses produced in the calorimeters with multiple fixed energies. A fixed number of events are taken at each energy in order to measure the linear response.

The PHOS4 Scan, also called the fine timing calibration, uses calibration pulses from the calorimeters, but at a single energy. A fixed number of events are taken at each time delay step of the PHOS4 chip from 0 to 24 ns. The data can be reconstructed to form a single pulse with a sample at every nanosecond instead of the typical sample every 25 ns. Using this reconstructed pulse, the optimal PHOS4 setting can be determined in order to sample the peak of the signal in the calibration regime.

5 Calibration and performance of the PreProcessor System

The PP System has been running within the ATLAS detector since Spring 2008 and is now precisely calibrated using LHC collision data. The initial validation of the PP System required testing every PPM before being shipped to CERN for installation in the full system. Once installed, long term testing began with the full L1Calo Trigger System. This includes performing regular calibration and analysing physics data to verify system performance. The tests and performance results will be presented in the following sections.

5.1 Hardware validation

All PPMs used in the ATLAS Trigger are validated in the lab, at the Kirchhoff-Institut für Physik, before being transferred to CERN. There are two test setups used, the first validates the real-time and readout data paths of all 64 channels of a single PPM using an analogue signal source to populate all 64 channels. The real-time data is captured by a custom unit, built in-house, receiving the full quantity of LVDS cables. However, recording of all channels in parallel at full speed is complicated and not necessary in the laboratory. Therefore, a group of four channels is selected, routed to deserialisers, and captured in a fast memory for a limited duration that is long enough to perform data integrity checks. Test results for all components are stored in a dedicated database, which holds tables showing:

- the installation status of the system in ATLAS-USA15,
- the status and test results of PPM motherboards fully equipped with daughter boards,
- the status and test results of AnIn daughter boards,
- the status and test results of MCMs,

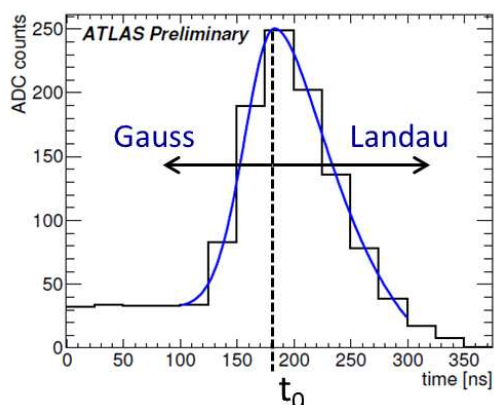


Figure 13. An example LAr calibration pulse as read out by the L1Calo PreProcessor System using the extended readout mode with 15 digitised ADC samples. The signal is fitted by an optimised Gauss/Landau function.

- the status and test results of LCD daughter boards,
- and status and test results of other peripheral modules (TTC, CAN, RGTM).

This database is also used to record test information and feedback from physics running at CERN in the case of hardware failures and provides a way to track component history.

The second test uses a fully operational and fully populated PP crate, as described in section 3. Here the environmental and operational stability of a system crate can be tested over a longer period of time. No input signals are used, but the playback memories of the PPrASIC are programmed to provide data test patterns that stress the processors and test for logic, single-bit, and memory errors.

These two test facilities are also used to study any hardware that is returned from CERN due to errors reported during data taking.

5.2 Timing calibration and performance

The analogue signals must be precisely aligned in time at the PPM input because sampling at the peak position is essential to ensure a correct BCID and precision E_T measurement. The fine-timing calibration was first established with the calorimeter pulser systems, as described in section 4.3, and then refined using the first LHC beam delivered to the detector as splash events in November 2009 [14]. Improved timing delays were applied early after the first 7 TeV collisions were delivered by the LHC at the beginning of 2010, based on the analysis of the recorded collisions data. Since then, the timing has been incrementally improved, such that, for the majority of recorded data, the timing of most towers is better than ± 2 ns which provides close to ideal performance.

Figure 13 shows an example of a digitised LAr calibration pulse as read out by the PPM using the extended readout mode with 15 digitised ADC samples. This signal is fitted using an optimised function. Depending on the calorimeter region, a hybrid function composed of either a Gauss or a Landau function applied on the rising edge combined with a Landau function on the falling edge was found to give the best fits. These Gauss/Landau or Landau/Landau fit functions are used to reconstruct the original pulses in order to extract the fine-timing information beyond the 25 ns

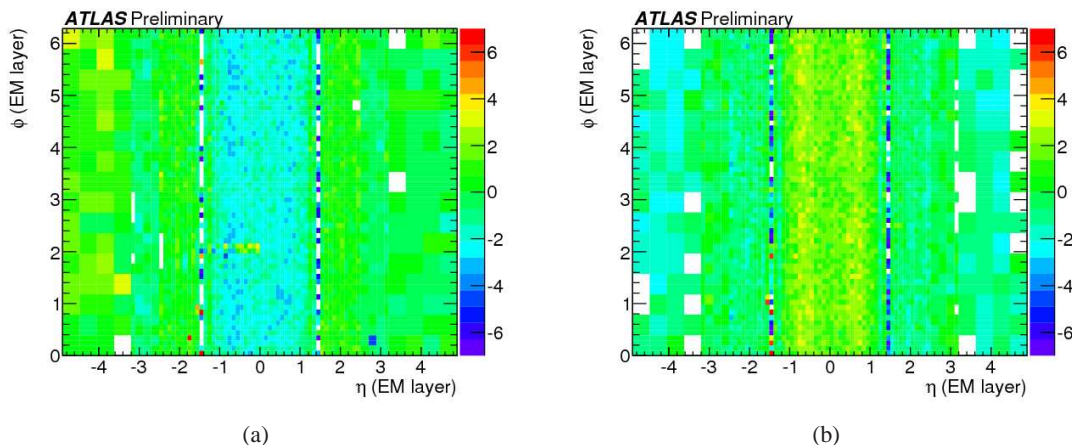


Figure 14. L1Calo trigger tower timing offsets in ns as function of η and ϕ for the electromagnetic (EM) calorimeter layer. (a) shows the results using collision data from the initial 2011 running period while (b) shows the results after applying corrections to the timing delays. White bins have no measurement due to lack of statistics.

sampling resolution. In order to avoid the large parameter space of these fitting functions, some fit parameters are constrained. In particular the widths of the Gauss and Landau sub-functions are derived from special calibration runs and are fixed for the analysis of the proton-proton collision data. As it is known that the pulses provided by the calorimeter pulser systems are slightly broader than those created by particles from collisions, the impact on the fit method and on the timing results needs further study.

The status of the TT timing as achieved at the start of the 2011 data taking period is shown in figure 14. The distributions show the offsets from the ideal timing, defined as the mean difference between the fitted maximum position t_0 and the middle of the central bin, in units of nanoseconds (cf. figure 13). The $\eta - \phi$ maps compare the distributions at the beginning of the 2011 running period with those after having applied resulting correction factors to the hardware timing delays. While the timing for the majority of the TTs is already within ± 2 ns for the early measurement, there exist some larger offsets, mainly due to modification and repair of calorimeter electronics during the 2010/11 winter shutdown which were compensated for by the corrections applied later.

5.3 Bunch crossing identification calibration and performance

Identifying the correct LHC BC to which an L1A belongs is very important for efficiently operating the L1Calo trigger. Since the TT signals span several BCs, a robust method is used to assign the pulse to the correct BC. It must operate correctly from very low energy signals up to saturation above the maximum energy of approximately 250 GeV. Three methods were described in section 3.3.1, *External BCID*, *Peak Finder BCID*, and *Saturated BCID*. The main method used with unsaturated signals is the *Peak Finder BCID* which sharpens the pulse using the FIR Filter before running the Peak Finder algorithm.

The FIR Filter coefficients being used at the beginning of the data taking period in 2010 were derived from the analysis of calorimeter calibration pulses [15]. After sufficient collision data were recorded, an improved set of coefficients were produced using signal shapes, each normalised to the

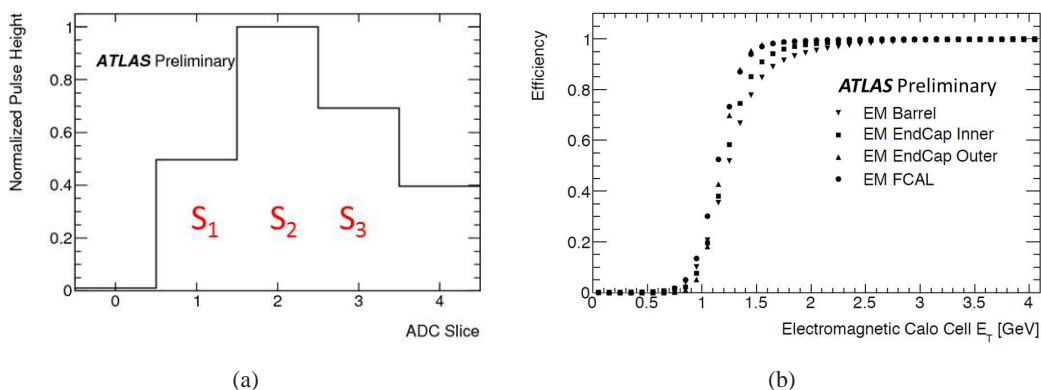


Figure 15. (a) The pedestal subtracted and normalised ADC pulse shape for an example trigger tower as derived from the analysis of 2010 collision data. (b) Efficiency for an EM trigger tower energy to be associated with the correct bunch crossing as a function of the sum of the raw cell E_T within that tower. Note the improving efficiency at higher η which is related to the $\sin(\theta)$ suppression of noise.

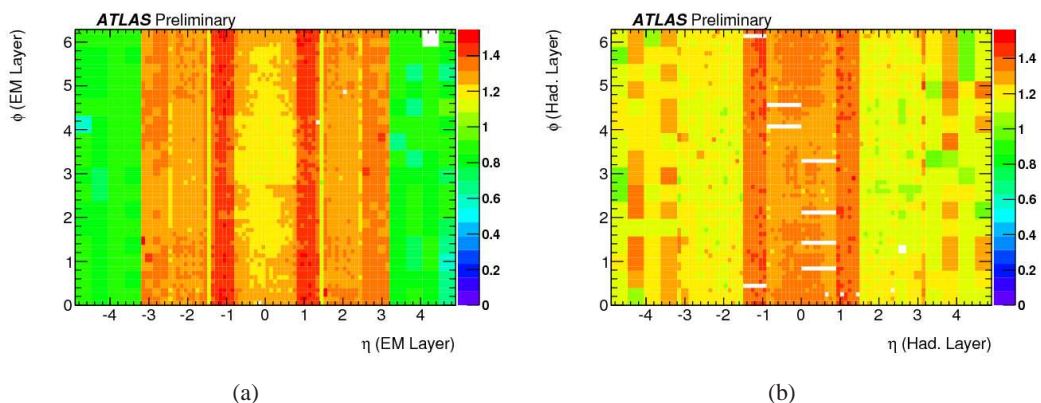


Figure 16. The sum $S_1 + S_3$ as a function of η and ϕ for the (a) electromagnetic and (b) hadronic calorimeter layer. The value S_i is the normalised pulse height of the i -th ADC sample as illustrated in figure 15(a). The plots show the results from the analysis of 2010 collision data. White bins indicate a channel that is not in use due to disabled hardware.

signal peak, as measured from proton-proton collisions. Figure 15(a) shows the pedestal subtracted and normalised ADC pulse shape for an example TT. The sum of the two samples surrounding the peak ($S_1 + S_3$), where S_i is the i -th ADC sample normalised by the signal peak, provides a measure for identifying regions in the calorimeter with similar pulse shape. As shown in figure 16, the pulse shape varies mainly along η , reflecting the transition between different calorimeter regions and varying detector geometry and implementation. A final set of region specific FIR coefficients is derived which follow the shapes of the normalised signals. The overall normalisation and drop-bits range is chosen such that the 8-bit LUT coverage is maximised.

A good indicator of a successful BCID, as well as timing calibration, is the efficiency of associating small energy deposits to the correct BC. Figure 15(b) shows the efficiency of an EM TT energy being associated with the correct BC, as a function of the raw calorimeter cell E_T within that tower for different regions of the EM calorimeter. In order to remove the majority of fake

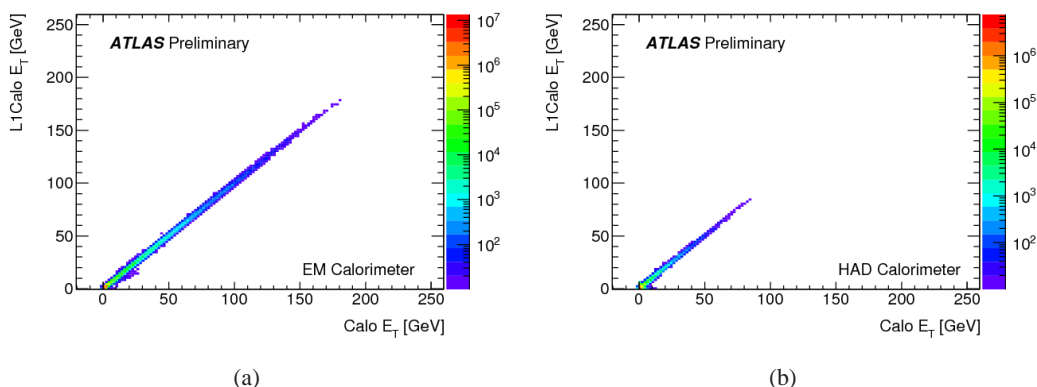


Figure 17. L1Calo trigger tower transverse energy versus offline transverse energy. The offline transverse energy is derived by summing the individual calorimeter cells associated to a tower. The plots show the results for the (a) electromagnetic and (b) hadronic calorimeters using 2011 collision data.

triggers due to small energy deposits, a noise cut is applied to the energy in the LUT. The effect of this cut at around 1.2 GeV is reflected in the turn-on curve which is in line with the optimal performance as expected from simulations.

5.4 Energy calibration and performance

Another critical aspect of operating the L1Calo trigger is the energy calibration of the input signal which translates ADC counts in the PPM to E_T delivered to the CP and JEP systems for further processing. Currently all calibration coefficients are implemented in an intermediate board, called the Receiver [16], through which all analogue signals pass while traveling from the detector to the PPM inputs. The Receiver contains an operational amplifier which is used to scale the analogue signals to ensure the E_T measured in the PPM and calorimeters is equal. It is planned to use the LUT for future corrections of dead material, crack losses and non-linearities.

The present calibration is derived from the Energy Scans, described in section 4.3, using dedicated calibration pulser runs which are regularly taken between LHC luminosity fills. Based on these energy scans, the analogue gain factors are derived for every TT by comparing the energy measured in L1Calo to the more precise calorimeter measurement. The status of the energy calibration is regularly verified in the analysis of collision data. Figure 17 shows the energy correlation plots between trigger and offline calorimeter E_T and reflects the good agreement between the L1Calo and calorimeter measurements.

By the end of the 2010 running period, sufficient data had been collected in order to perform detailed studies of the energy calibration on a channel-by-channel basis and as a function of relevant observables. Figure 18 shows the derived fractional difference between L1Calo and offline transverse energy as a function of the offline transverse energy. The L1Calo energy is calculated using two different methods: the energy based on the ADC peak sample and the energy based on the result of the LUT. Disregarding a minor overall offset, the 2010 calibration reveals small LUT deviation at low energies. This effect was found to be due to a rounding bias in the way the LUT was derived which was successfully corrected for the 2011 data taking period such that the agreement is now better than $\sim 1\%$ for $E_T > 10$ GeV.




Figure 18. Fractional difference between L1Calo transverse energy and offline transverse energy as a function of the offline transverse energy. The L1Calo energy is calculated using two different methods: the energy based on the ADC peak sample and the energy based on the result of the LUT. The plots show the results for the electromagnetic barrel (EMB) derived from (a) 2010 data in comparison with the results using (b) 2011 data which include improvements in the LUT calculation.

5.5 Performance during physics data taking

The calibration activities described above result in a very well calibrated PPM and L1Calo Trigger as depicted in figure 19. Figure (a) shows the rate for three primary triggers produced by the L1Calo Trigger versus the instantaneous luminosity. The trigger rates themselves are stable and mostly scale well over a wide range of luminosities and time. As expected, pile-up effects mainly affect the missing and total E_T triggers as well as the jet trigger items based on forward calorimetry. The stability of the system can also be seen in the sharp, rising edge of the efficiency curves which saturate at high values as presented in figure (b) for an inclusive electron trigger. Here the trigger efficiency versus E_T is shown for a L1 trigger threshold and the corresponding L2 and EF triggers which are seeded by the L1 trigger.

6 Conclusions

The PreProcessor System is a central component of the Level-1 Calorimeter Trigger. It uses a pipelined design to process 7168 analogue Trigger Tower signals provided by the calorimeters in order to measure the signal E_T and assign the signal to the proper LHC BC in under a microsecond. The integrity of these measurements has a direct impact on the efficiency of the ATLAS Trigger System.

The PreProcessor Module, in particular, has been shown to be operating successfully and essentially error free after the initial validation in the lab and at CERN, followed by the 2010 and 2011 LHC physics data taking periods. During the 2010 physics data taking, incremental improvements of the timing, the BCID performance and the energy calibration established the L1Calo System with close to ideal performance such that in 2011 only minor adjustments were necessary. With even larger data samples being recorded in the near future, further optimisation of the calibration will be possible such as a tower-by-tower energy calibration based on identified physics objects with precisely known energies, for example electrons from Z decays.

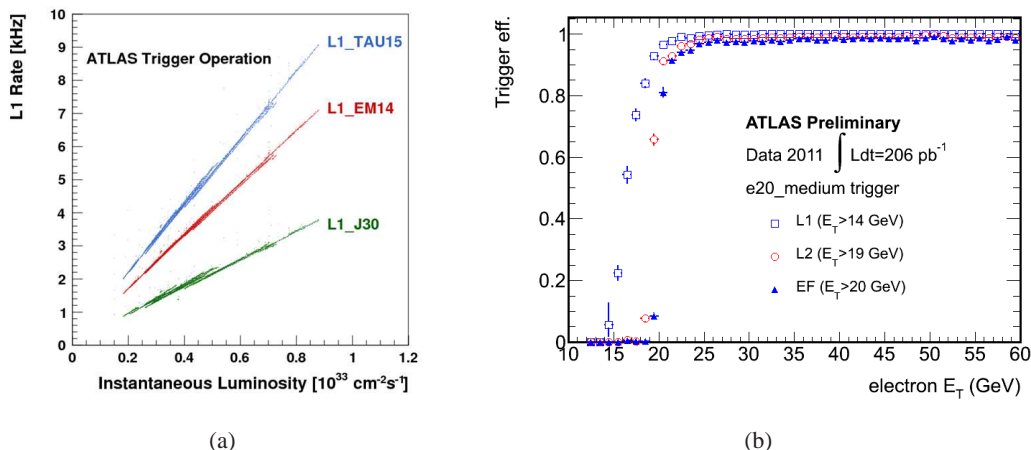


Figure 19. (a) Unprescaled L1 rates from the initial 2011 data taking period as a function of the instantaneous luminosity for an electromagnetic trigger with a threshold of 14 GeV, a tau trigger with a threshold of 15 GeV and a jet trigger with a threshold of 30 GeV. The instantaneous luminosity used is the online measurement. (b) Efficiencies for e20_medium at each trigger level (L1, L2 and EF) measured with $Z \rightarrow ee$ events using the tag-and-probe method. Efficiencies are measured as a function of the offline electron E_T for candidates satisfying tight identification requirements. Opposite sign electron pairs with $80 < M_{ee} < 100$ GeV are used for the $Z \rightarrow ee$ selection.

Acknowledgments

We are deeply indebted to the skills and patient work of Simon Pyatt, Alexander Leonhardt and Xenakis Serghi.

We are very grateful to the large number of people who worked on this project and made numerous important contributions during the 16 years it has taken to come to fruition, but who now work elsewhere or on other projects. We cannot list everyone, but specifically mention Oliver Stelzer-Chilton, Werner Hinderer, Dan Husmann, Volker Schatz, and Bernd Stelzer.

The Level-1 Calorimeter Trigger has been supported by BMBF and the Helmholtz Alliance, Germany; The Swedish Research Council and the Knut and Alice Wallenberg Foundation, Sweden; and the Particle Physics and Astronomy Research Council (up to 2007) and the Science and Technology Facilities Council (from 2007), United Kingdom.

References

- [1] ATLAS collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, 2008 *JINST* **3** S08003 [[INSPIRE](#)].
- [2] ATLAS collaboration, *Performance of the ATLAS Trigger System in 2010*, *Eur. Phys. J. C* **72** (2012) 1849 [[arXiv:1110.1530](#)].
- [3] ATLAS collaboration, *ATLAS high-level trigger, data-acquisition and controls: Technical Design Report*, CERN-LHCC-2003-022 (2003).
- [4] R. Achenbach et al., *The ATLAS level-1 calorimeter trigger*, 2008 *JINST* **3** P03001 [[INSPIRE](#)].
- [5] ATLAS collaboration, *ATLAS level-1 trigger: Technical Design Report*, CERN-LHCC-98-14 (1998).

- [6] H. Boterenbrood et al., *Design and Implementation of the ATLAS Detector Control System*, [ATL-DAQ-2003-043](#) (2003).
- [7] *IEEE Standard for Additional Mechanical Specifications for Microcomputers Using the IEEE Std 1101.1-1991 Equipment Practice* (1997), <http://ieeexplore.ieee.org/servlet/opac?punumber=4595>.
- [8] N. Buchanan et al., *ATLAS liquid argon calorimeter front end electronics*, [2008 JINST 3 P09003](#) [[INSPIRE](#)].
- [9] A. Cerqueira, J. Seixas and L. Calôba, *Analog system for building the first-level triggering signal provided by the hadronic calorimeter of ATLAS detector*, *Nucl. Instrum. Meth. A* **570** (2007) 117 [[INSPIRE](#)].
- [10] *IEEE Standard for a Common Mezzanine Card (CMC) Family* (2001), <http://ieeexplore.ieee.org/servlet/opac?punumber=7510>.
- [11] T. Toifl, P. Moreira, A. Marchioro and R. Vari, *4-channel rad hard delay generation ASIC with Ins timing resolution for LHC*, *IEEE Trans. Nucl. Sci.* **46** (1999) 139 [[INSPIRE](#)].
- [12] RD12 PROJECT collaboration, B. Taylor, *TTC distribution for LHC detectors*, *IEEE Trans. Nucl. Sci.* **45** (1998) 821 [[INSPIRE](#)].
- [13] International Organization for Standardization, *ISO 11898-1:2003 - Road vehicles — Controller area network (CAN) — Part 1: Data link layer and physical signalling*, (2003).
- [14] ATLAS collaboration, *ATLAS Level-1 Calorimeter Trigger: Timing Calibration with 2009 LHC Beam Splashes*, [DAQ-PUB-2010-001](#) (2010).
- [15] ATLAS collaboration, *Digital Filtering Performance in the ATLAS Level-1 Calorimeter Trigger*, [DAQ-PROC-2010-012](#) (2010).
- [16] W.E. Cleland, B. Liu, J. Rabel and G. Zuk, *Receiver/Monitor System for the ATLAS Liquid Argon Calorimeter*, Technical Report [ATL-AL-EN-0043 v.4](#), CERN, Geneva Switzerland (2006).