

The ATLAS Level-1 Calorimeter Trigger

R. Achenbach,^b P. Adragna,^d V. Andrei,^b P. Apostologlou,^e B. Åsman,^f C. Ay,^c
B.M. Barnett,^e B. Bauss,^c M. Bendel,^c C. Boehm,^f J.R.A. Booth,^{ae} I.P. Brawn,^e
P. Bright-Thomas,^a D.G. Charlton,^a N.J. Collins,^a C.J. Curtis,^d A. Dahloff,^c
A.O. Davis,^e S. Eckweiler,^c J.P. Edwards,^e E. Eisenhandler,^d P.J.W. Faulkner,^a
J. Fleckner,^c F. Föhlich,^b J. Garvey,^a C.N.P. Gee,^e A.R. Gillman,^e P. Hanke,^b
R.P. Hatley,^e S. Hellman,^f A. Hidvégi,^f S.J. Hillier,^a K. Jakobs,^c M. Johansen,^f
E.-E. Kluge,^b M. Landon,^d V. Lendermann,^b J.N. Lilley,^a K. Mahboubi,^b G. Mahout,^a
A. Mass,^b K. Meier,^b T. Moa,^f E. Moyse,^d F. Müller,^b A. Neusiedl,^c C. Nöding,^c
B. Oltmann,^c J.M. Pentney,^d V.J.O. Perera,^e U. Pfeiffer,^b D.P.F. Prieur,^e W. Qian,^e
D.L. Rees,^a S. Rieke,^c F. Rühr,^b D.P.C. Sankey,^e U. Schäfer,^c K. Schmitt,^b
H.-C. Schultz-Coulon,^b C. Schumacher,^b S. Silverstein,^f R.J. Staley,^a R. Stamen,^b
M.C. Stockton,^a S. Tapprogge,^c J.P. Thomas,^a T. Trefzger,^c P.M. Watkins,^a
A. Watson,^a P. Weber,^b and E.-E. Woehrling^a

^a School of Physics and Astronomy, University of Birmingham,
Birmingham B15 2TT, U.K.

^b Kirchhoff-Institut für Physik, University of Heidelberg,
D-69120 Heidelberg, Germany

^c Institut für Physik, University of Mainz,
D-55099 Mainz, Germany

^d Physics Department, Queen Mary, University of London,
London E1 4NS, U.K.

^e STFC Rutherford Appleton Laboratory,
Harwell Science and Innovation Campus, Didcot, Oxon OX11 0QX, U.K.

^f Fysikum, Stockholm University,
SE-106 91 Stockholm, Sweden
E-mail: e.eisenhandler@qmul.ac.uk

ABSTRACT: The ATLAS Level-1 Calorimeter Trigger uses reduced-granularity information from all the ATLAS calorimeters to search for high transverse-energy electrons, photons, τ leptons and jets, as well as high missing and total transverse energy. The calorimeter trigger electronics has a fixed latency of about $1\ \mu\text{s}$, using programmable custom-built digital electronics. This paper describes the Calorimeter Trigger hardware, as installed in the ATLAS electronics cavern.

KEYWORDS: Digital electronic circuits; Trigger concepts and systems (hardware and software).

* Corresponding author.

Contents

1. The ATLAS trigger system	2
1.1 Introduction	2
1.2 Outline of the paper	2
1.3 Requirements and trigger levels	3
1.4 The Level-1 Trigger	4
1.5 The Level-2 Trigger and the Event Filter	5
1.6 The data acquisition system	5
2. Level-1 Calorimeter Trigger introduction	5
3. The analogue front-end	7
4. The PreProcessor	8
4.1 Analogue signal-handling	9
4.2 The PreProcessor Multi-Chip Module	10
4.3 Output signals	11
4.4 Bunch-Crossing Identification (BCID)	12
4.5 Bunch-Crossing multiplexing (BC-mux)	14
5. The Cluster and Jet/Energy-sum Processors	14
5.1 The Cluster Processor Module (CPM)	15
5.1.1 The electron/photon and τ /hadron algorithms	15
5.1.2 Data input and fan-in/out	17
5.1.3 Processing	19
5.1.4 Readout to DAQ and Level-2	19
5.2 The Jet/Energy Module (JEM)	20
5.2.1 The jet algorithm	20
5.2.2 Energy sums	21
5.2.3 Data input and fan-in/out	22
5.2.4 Processing jets	22
5.2.5 Processing energy sums	24
5.2.6 Readout to DAQ and Level-2	24
5.3 The Common Merger Module (CMM)	24
5.3.1 Requirements and architecture	24
5.3.2 Hit counting	25
5.3.3 Total- E_T and missing- E_T summing	26
5.3.4 Total jet transverse energy	26
5.3.5 Outputs to CTP, DAQ and Level-2	27
5.4 The Processor Backplane (PB)	27

6. The Readout Driver (ROD) and data readout	29
6.1 Requirements and architecture	29
6.2 Data handling	30
6.2.1 Data-rate limitation	31
6.3 Structure of the ROD module	32
7. Infrastructure	34
7.1 Detector Control System (DCS)	34
7.2 The Timing Control Module (TCM)	35
7.2.1 TTC signal distribution	35
7.2.2 CANbus node	36
7.2.3 TCM variants, the auxiliary backplane and geographical addressing	36
7.3 Crate control and infrastructure	37
8. Status	38

1. The ATLAS trigger system

1.1 Introduction

The ATLAS detector at the CERN Large Hadron Collider (LHC) is one of the largest and most complex pieces of scientific apparatus ever built. The extraction of physics resulting from collisions of two 7 TeV beams of protons¹ at very high luminosity is a demanding procedure requiring deep understanding of the detector, and careful reduction and analysis of an enormous quantity of data. The ATLAS trigger system must make an irreversible, online selection of a tiny fraction of collisions within a very short time. The Level-1 Trigger (L1) provides the first and largest step of that selection, and delivers its decision within a fixed time of less than 2.5 μ s.

A comprehensive overview of the ATLAS detector is given in [1]. The present paper gives a fuller description of the Level-1 Calorimeter Trigger (L1Calo), which is a major component of L1. At the time of writing these papers, both the detector and the LHC machine were nearing completion. Future papers will describe the operation, software, commissioning and performance of the trigger.

1.2 Outline of the paper

In this section we discuss the basic triggering requirements for ATLAS, and give a brief overview of the three-level trigger that has been implemented. Section 2 presents general information regarding the overall design of the Level-1 Calorimeter Trigger. Section 3 describes the handling of the analogue signals used by the trigger, and in section 4 the PreProcessor which converts these signals into calibrated, correctly timed digital data is described. Section 5 concerns the two digital processors which implement the trigger algorithms, and section 6 the data readout to the data acquisition system and the Level-2 Trigger. Various aspects of the infrastructure, and the use of the Detector Control System, are discussed in section 7.

¹ The LHC will also run with heavy-ion beams, but that is not discussed in this paper.

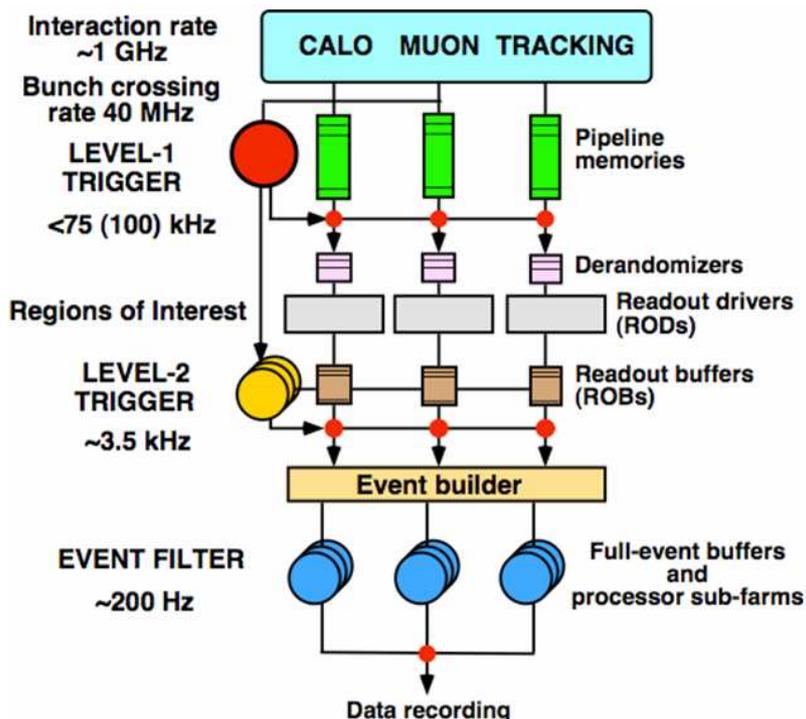


Figure 1. The three-level trigger and data acquisition systems of ATLAS.

1.3 Requirements and trigger levels

Many of the physics processes of interest at the LHC have very small cross-sections. The machine is therefore designed to achieve a luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ at 14 TeV centre-of-mass energy, producing a proton–proton interaction rate of about 1 GHz. Proton bunches collide in ATLAS every 25 ns (24.95 ns, to be more precise), and this means that at design luminosity every p – p collision of potential interest will be accompanied by an average of about 20 inelastic events per bunch-crossing in the detector. The recording rate for event data is limited to about 200 Hz, so the overall trigger rejection factor must be about 5×10^6 while at the same time achieving the maximum possible efficiency for the rare and exciting physics events.

ATLAS has adopted a three-level trigger system. A block diagram of the trigger and data acquisition systems is shown in figure 1. L1 is a synchronous system, using custom digital high-speed pipelined electronics to process a huge amount of reduced-granularity detector data in parallel. From the raw 40.08 MHz bunch-crossing rate (referred to later as 40 MHz, with multiples referred to as 80 MHz and 160 MHz) it must select candidates at a maximum rate of 75 kHz (and potentially 100 kHz), within a fixed time of less than $2.5 \mu\text{s}$.

The two following trigger levels, collectively called the Higher Level Trigger (HLT), use a high-capacity switched network of several thousand commercial computers which can access complete detector information to refine the selection. The Level-2 Trigger (L2) further reduces the rate to approximately 3.5 kHz within about 40 ms. The Event Filter (EF) can access fully built events, and uses offline analysis methods to achieve the final storage rate of about 200 Hz with event size of about 1.3 MB after about 4 s of processing. Both stages of the HLT use the full granularity and precision of calorimeter and muon chamber data, as well as Inner Detector tracking data. Better energy-deposition data improves threshold cuts, while track reconstruction makes particle identification, such as electrons versus photons, possible.

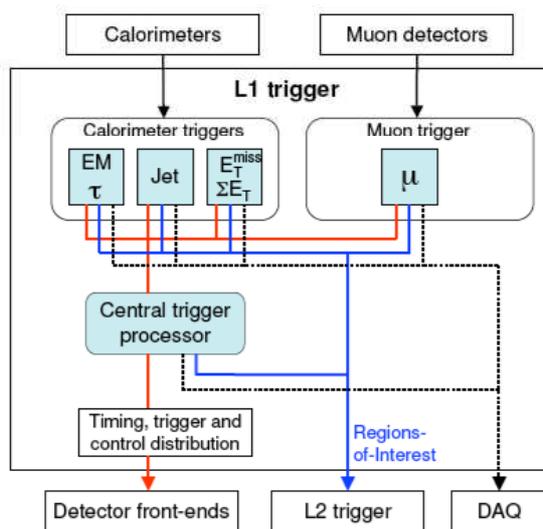


Figure 2. Block diagram of the Level-1 Trigger.

1.4 The Level-1 Trigger

In order to separate the desired rare processes from the predominant QCD jet production and other backgrounds, L1 searches mainly for exclusive signatures that could identify isolated high transverse-energy electrons, photons, muons, and τ 's, as well as missing transverse energy. Jets, high total transverse energy, and total jet transverse energy are also flagged.

In order to be able to do this quickly, L1 uses reduced-granularity data from the muon and calorimetric detectors: the resistive-plate chambers (barrel) and thin-gap chambers (endcap) for muons, and all of the calorimeters for electromagnetic clusters (i.e. electrons or photons), jets, τ 's decaying into one or more isolated hadrons (i.e. a τ /hadron trigger), and missing and total transverse energy. The maximum L1 accept rate that the detector readout systems can handle is 75 kHz, but they are required to be upgradeable to 100 kHz.

The 25 ns interval between successive bunch-crossings is far too short for processing and selecting events. In fact, given the size of the detectors and with much of the L1 electronics off the detector in a separate cavern, even the transmission delays for the signals are much longer than that. Therefore, the scheme adopted is for the detector data to be held in buffers while L1 makes its decision. If the bunch-crossing passes the L1 criteria a Level-1 Accept (L1A) signal is sent and the data are kept; if not they are deleted. The time allowed for the L1 stage depends on the size of the on-detector buffers, and that requires a compromise between a long enough processing time to allow effective trigger algorithms, and the cost and complexity of very large data buffers. In ATLAS the allowed decision time, or latency, for L1 was chosen to be 2.5 μ s. Since it is unacceptable to exceed this, L1 was designed to have a nominal latency of about 2.0 μ s in order to ensure an adequate safety margin. It is important to note that a large fraction of this time is consumed by the signal transmission delays from and back to the detector front-end electronics.

The Level-1 Trigger, shown in figure 2, uses three main components to make its decisions. The Level-1 Muon Trigger (L1Muon) uses track information from dedicated, fast muon chambers to identify high- p_T muon candidates. The Level-1 Calorimeter Trigger (L1Calo) uses calorimeter energy deposits to identify various types of high- E_T objects as well as energy sums of interest. Results from both of these systems are processed by the Central Trigger Processor

(CTP). The CTP implements a trigger ‘menu’ based on logical combinations of results from L1Calo and L1Muon. It can also pre-scale menu items, in order to make efficient use of the allowable rate bandwidth as the luminosity and background conditions change.

Events that pass the L1 selection conditions are transferred from the detector-specific front-end electronics to the data acquisition system. In addition, information from L1 itself, to indicate how it made its decision, is also read out. In parallel with this, L1 supplies information on so-called Regions-of-Interest (RoIs) to the Level-2 Trigger. For exclusive objects these RoIs are the geographical coordinates of the detector regions where they were found, as well as the criteria (e.g. thresholds) that they satisfied. This information is used by L2 to seed its selection process.

1.5 The Level-2 Trigger and the Event Filter

L2 uses full-granularity readout data from the data acquisition system and dedicated algorithms to refine the selections made by L1. L2 reduces the rate to a maximum of approximately 3.5 kHz, with an average latency of about 40 ms. In order to reduce the amount of data that must be transferred, L2 uses the RoI results from L1 to select a subset (~2%) of the total readout data to process. The RoIs from L1Muon, L1Calo and the CTP for a given event are assembled in a custom device known as the RoI Builder (RoIB). L2 consists of a large network of commercial CPUs linked by a high-capacity switched network.

The final stage of event selection is provided by the Event Filter. It works with fully built events, and therefore can use analysis procedures and algorithms similar to the offline data processing. Its processing time is about 4 s per event, and it further reduces the rate to about 200 Hz. The events passing the Event Filter selection are permanently stored for offline analysis, with event sizes of approximately 1.3 MB.

1.6 The data acquisition system

The data acquisition system (DAQ) receives event data from detector-specific Readout Driver (ROD) modules over 1600 point-to-point optical readout links. On request it sends data from within the Regions-of-Interest to L2, and for those events that satisfy the L2 criteria it carries out event building. The built events are then sent to the Event Filter, which selects events for permanent storage.

In addition to handling all of this data movement, the DAQ manages the configuration, control and monitoring of the entire ATLAS detector during data-taking. However, it does not supervise some of the functions needed to operate the detector hardware, such as power and gas systems — this type of functionality is provided by the Detector Control System (DCS).

2. Level-1 Calorimeter Trigger introduction

L1Calo is a fixed-latency, pipelined digital system using custom electronics. Its input data comes from about 7200 analogue trigger towers of reduced granularity, mostly 0.1×0.1 in $\Delta\eta \times \Delta\phi$, from all the ATLAS electromagnetic and hadronic calorimeters. (η is pseudo-rapidity and ϕ is azimuthal angle around the beam axis.) The L1Calo electronics has a latency of less than a microsecond, resulting in a total latency of about $2.1 \mu\text{s}$ for the L1Calo chain including cable transmission delays and the CTP processing time. This is well inside the required $2.5 \mu\text{s}$ envelope.

The L1Calo system is located entirely off the detector, in the large, separate electronics cavern known as USA15. A block diagram of L1Calo’s basic architecture is shown in figure 3. There are three main sub-systems. The PreProcessor first digitises the analogue calorimeter trigger-tower signals, then uses a digital filtering technique to associate the relatively wide

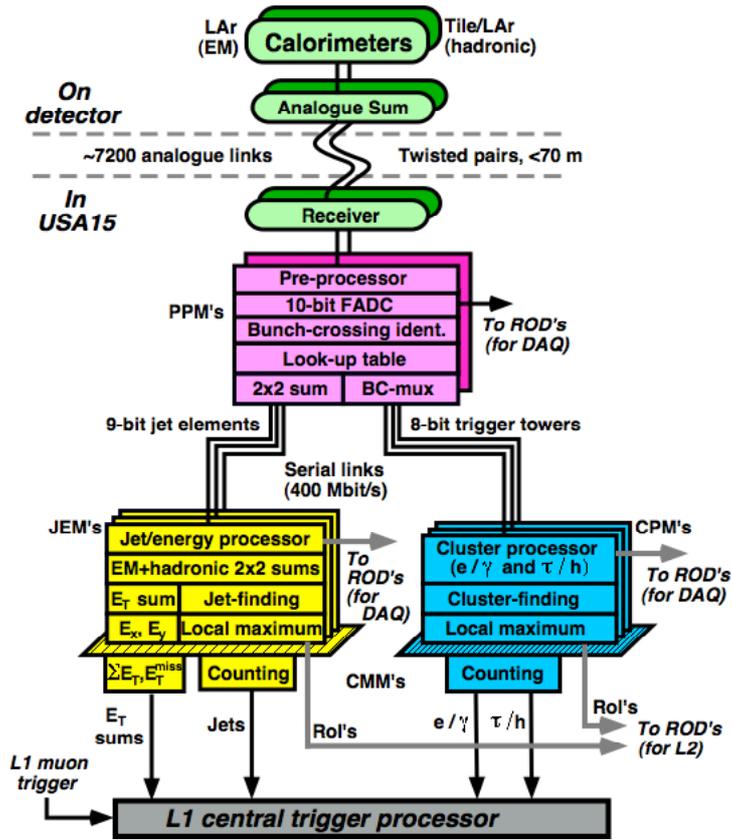


Figure 3. Architecture of the Level-1 Calorimeter Trigger. The real-time data path is in black, while the readout data paths are in grey.

analogue pulses with specific LHC bunch-crossings. A look-up table achieves several operations in one step: pedestal subtraction, final E_T calibration, noise suppression, and turning off problematic channels. The data are then transmitted in parallel to the two algorithmic processors: the Cluster Processor (CP) and the Jet/Energy-sum Processor (JEP). The CP identifies candidate electrons, photons and τ 's with high E_T above programmable thresholds and, if desired, passing isolation requirements. The JEP operates on so-called 'jet elements' at the somewhat coarser granularity of 0.2×0.2 in $\Delta\eta \times \Delta\phi$ to identify jets as well as produce global sums of total, missing, and jet-sum E_T . Both the CP and the JEP count 'hit' multiplicities of the different types of trigger objects, and send them, together with bits indicating which global E_T -sum thresholds were exceeded, to the CTP for use in the trigger menu. A detailed account of L1Calo trigger algorithms can be found in [2].

For all events that are selected by L1, a programmable selection of data from L1Calo is read out via Readout Driver modules to DAQ. At minimum, these data include trigger-tower energies and L1Calo results in order to allow calibration, monitoring and verification of the trigger. Digitised raw data, and intermediate results from points along the trigger logic chain, will also be read out during commissioning and early running in order to check correct functioning of the trigger, and later on whenever they are required to diagnose any problems. The trigger readout data may also provide useful diagnostic information for the LHC machine and ATLAS detectors.

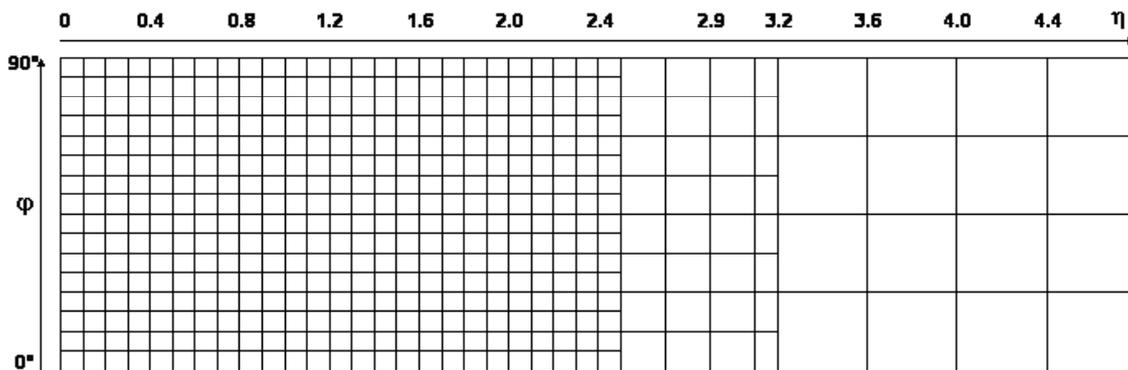


Figure 4. Trigger-tower granularity for $\eta > 0$ and one quadrant in ϕ .

In parallel with the information read out to DAQ, RoIs giving details of electron/photon, τ /hadron and jet cluster candidates are also read out by RODs and sent to the L2 RoI Builder. Missing, total, and total-jet transverse energy values are also sent.

All the main custom modules that comprise L1Calo are 9U (366 mm) in height and 400 mm deep. PreProcessor and ROD crates use standard VMEbus, while the Cluster and Jet/Energy-sum Processors use a custom backplane and a reduced VMEbus implementation. All modules include on-board monitoring of voltages and temperatures, interfaced via CANbus to the ATLAS Detector Control System.

The L1Calo hardware is designed to be relatively compact, with a high density of logic and interconnections. One of the reasons for this is to minimise the latency. Another feature is that some of the hardware modules are designed to carry out more than one role in the system, by using different firmware. This reduces the number of different module types, which in turn leads to a lower hardware cost and simplifies maintenance, at the cost of additional firmware complexity.

3. The analogue front-end

The ATLAS calorimetry [1] comprises the barrel, end-caps, and forward regions. In the barrel, a liquid-argon (LAr) electromagnetic (EM) calorimeter is surrounded by a scintillating-tile hadronic calorimeter (TileCal). In the end-caps liquid argon is used for both the EM and hadronic calorimeters. The forward calorimeters also utilise liquid argon.

Projective trigger towers are formed by analogue summation on the detector [3],[4]. They are 0.1×0.1 in $\Delta\eta \times \Delta\phi$ over most of the calorimetry, but larger in parts of the end-caps and in the Forward Calorimeters (where they are not projective in η), as shown in figure 4. Trigger towers cover the full depth of each of the electromagnetic or hadronic calorimeters. The number of calorimeter cells summed to form trigger towers depends on the granularity of the respective calorimeter, and ranges from a few in the end-caps up to 60 in the LAr EM barrel. In the Tile Calorimeter most towers are formed by summing five photomultiplier signals.

A distinction between EM and hadronic trigger towers is that the Tower Builder Boards used to sum EM trigger towers also convert the raw energy scale of the signals to transverse energy, but the hadronic trigger-tower signals from the Tile, LAr Hadronic End-Cap, and Forward Calorimeters are transmitted on the raw energy scale.

The analogue trigger-tower signals from the calorimeters are carried to L1Calo on 616 16-way twisted-pair cables. The twisted pairs are individually shielded, and there is also an outer global shield. The cables and connectors have been carefully selected to achieve less than

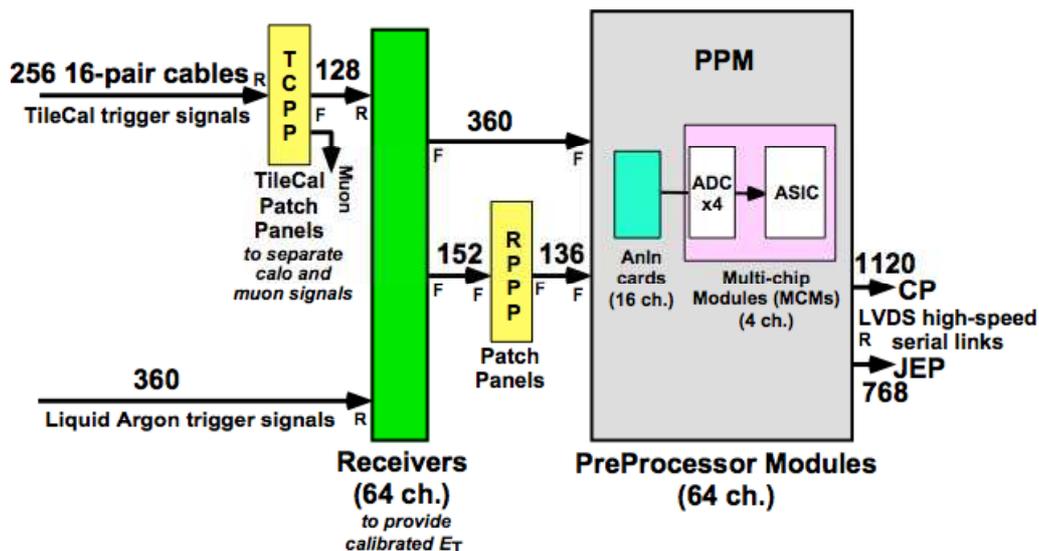


Figure 5. Front-end configuration of the Level-1 Calorimeter Trigger.

0.5% cross-talk between towers. The propagation velocity of signals is equivalent to 4.76 ns/m. The routing of the cables is specially optimised to reduce their length, and hence the delay, by penetrating the shielding between the main ATLAS cavern and USA15 through special holes leading directly to the trigger electronics racks. The lengths range from 30 m for the shortest LAr barrel cables to 70 m for the longest TileCal extended barrel cables.

Figure 5 is a block diagram showing how the signals are handled in USA15. The labels F and R on the diagram indicate the front and rear panels of modules, respectively, and the numbers indicate numbers of cables. The long cables from the TileCal also carry signals from the rear calorimeter sampling, which could be used if needed to help reduce backgrounds in the Level-1 Muon Trigger. The two kinds of TileCal signals are separated using patch-panels (TCPPs) upstream of the L1Calo Receivers. Figure 5 also shows the 776 short cables used in USA15 to route the signals to L1Calo. These are the same type of 16-pair cable as the long ones. They are individually trimmed to length in order to minimise latency, and also to achieve a tidy routing solution with these thick, stiff cables (see figure 6).

All trigger-tower signals pass through Receiver Modules [3] before being sent to the L1Calo PreProcessor. The Receivers include linear variable-gain amplifiers controlled by DACs. These are used to convert the hadronic trigger towers from energy to transverse energy, to compensate for attenuation in the different lengths of cable, and to set the E_T calibration of all signals. The Receivers also include a facility for monitoring a small, programmable selection of analogue signals; this is the only direct access to analogue calorimeter signals in ATLAS when the detector is closed.

4. The PreProcessor

The PreProcessor (PPr) digitises the trigger-tower signals, identifies the bunch-crossing they originate from, and does the final calibration and preparation of the signals for use in the algorithmic processors. It consists of 124 PreProcessor Modules (PPMs), each of which receives four analogue cables carrying a total of 64 trigger-tower signals on its front panel. The PPMs are housed in eight 9U VME64xP crates, six of which hold 16 PPMs and two of which hold 14 PPMs. Four of the crates process EM towers and four process hadronic towers. Figure 7 is a



Figure 6. Short analogue cables for $\eta > 0$, from Receivers (two racks at right) to PPMs (two racks at left).

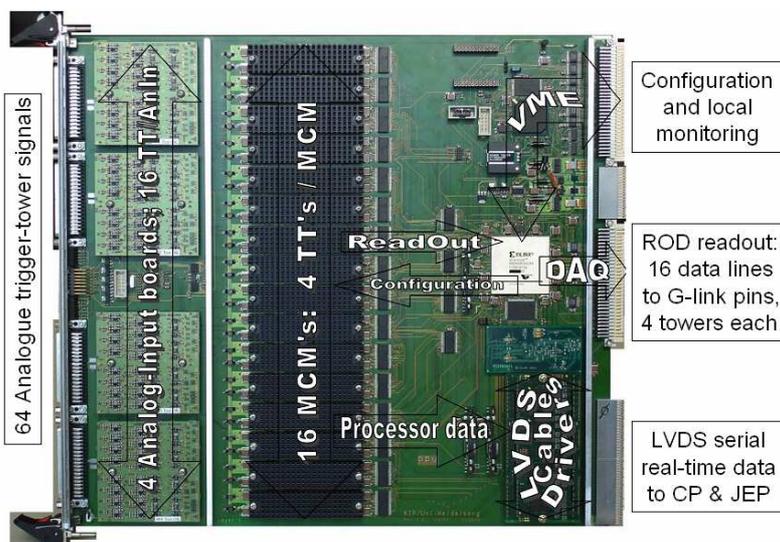


Figure 7. Photograph of a PreProcessor Module, with its main components labelled.

labelled photograph of a PPM, and figure 8 is a block diagram of its functionality. Each of the 9U crates also contains a Timing Control Module, to distribute the LHC clock and to monitor voltages and temperatures via CANbus, and a 6U VMEbus CPU module.

4.1 Analogue signal-handling

The differential signals first enter 16-channel analogue-input daughter cards, which convert them to single-ended signals with appropriate gain and bias for digitisation. A DAC provides a voltage offset used to set a suitable pedestal value. The pedestal allows small negative excursions

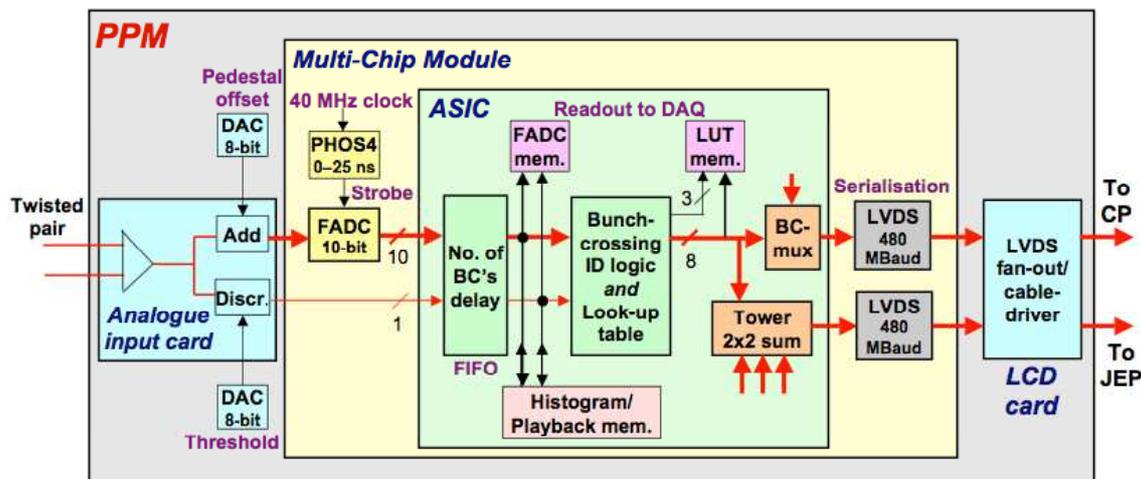


Figure 8. Block diagram of operations on one trigger tower by the PreProcessor Module.

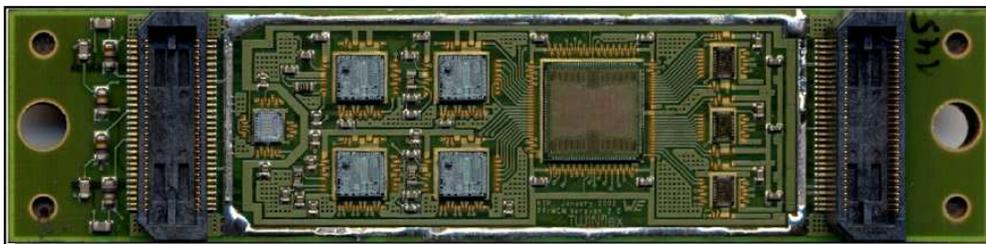


Figure 9. Photograph of a PreProcessor Multi-Chip Module without its cover. Connectors are at the left and right. The chip at left is the PHOS4 quad delay. The four chips left of centre are the FADCs, with the PPr-ASIC to their right then the three LVDS serial converters.

sions in the signal to be seen. The gain and pedestal are set such that the ADCs will saturate for signals corresponding to about 250 GeV, as this is also the level at which much of the trigger-tower summation logic on the calorimeters saturates. The implications of saturated signals for the trigger are discussed later in this paper.

4.2 The PreProcessor Multi-Chip Module

The main signal processing is performed on 16 Multi-Chip Modules (MCMs), each of which processes four trigger towers. The MCMs are easily replaceable plug-in components; figure 9 is a photograph. Four flash-ADCs (FADCs; Analog Devices AD-9042) digitise the signals to 10-bit precision at the bunch-crossing frequency of 40.08 MHz. (The FADCs actually digitise to 12 bits, but the two least significant bits are discarded.) The digitised values are sent to the PreProcessor ASIC (PPr-ASIC, designed at Heidelberg and built in a 0.6 μm process by AMS), which includes FIFOs for readout to DAQ. In order to be able to set up the timing and to understand the pulse shapes for the bunch-crossing identification logic, several FADC values (most often five) from bunch-crossings before and after the peaks of the pulses can be read out to DAQ.

The PreProcessor must synchronise the signals from a given bunch-crossing to compensate for differences in time-of-flight and signal-path lengths. Cables from the calorimeters to L1Calo are of many different lengths, so in general signals from the same bunch-crossing arrive at the PPMs at different times. To achieve synchronism both fine and coarse timing adjustments are provided. Fine adjustment of the FADC strobe so that it falls close to the peak position of each trigger-tower signal is performed by a four-channel ASIC (PHOS4, designed at CERN), which

provides delays programmable in 1 ns steps over a 25 ns range. The FADC output data for each trigger tower are re-timed to the main 40 MHz clock and then passed through a FIFO, which provides coarse timing correction in 25 ns steps over a wide range. Adjustments to the timing strobe and FIFO depth allow all trigger-tower data to be aligned in time.

The ASIC then assigns signals to the correct bunch-crossing, as described in section 4.4. The output of the bunch-crossing identification logic is a single 10-bit value correctly synchronised to the main clock. The 10-bit value is then used as the address for a look-up table. The contents of the look-up table are 8-bit E_T values that are used in the subsequent algorithmic processing with a nominal E_T scale of 1 GeV per count. Ten-bit values are retained up to the look-up table in order to ensure that full use can be made of the 8-bit range after the operation of the bunch-crossing logic and the transformations which are done in the look-up table. The use of a look-up table allows several operations to be carried out simultaneously, and in a very flexible way which allows non-linear transformations if necessary. The primary operation is a final tuning of the transverse-energy scale by using a linear transformation. At the same time, an offset that subtracts the pedestal is applied. Very small signals (~ 1 GeV) that are most likely due to noise can be set to zero. Saturated signals are set to full scale, i.e. 255 counts. Finally, problematic or dead trigger-tower signals can simply be set to zero. The output values from the look-up table are stored in a FIFO for readout to DAQ — this is important because they are the inputs used for algorithmic processing, and thus allow the trigger functionality to be monitored.

Following the look-up table, the four channels in each MCM are summed to form a 0.2×0.2 region in $\Delta\eta \times \Delta\phi$ for use in the Jet/Energy-sum Processor. This sum is truncated to nine bits, so in the rare case that the addition overflows it is set to full scale.

An additional feature of the PPr-ASIC is a playback memory, allowing test data to be introduced in order to verify operation of the digital part of the PreProcessor and the downstream trigger processors.

The same memories are also used in physics data-taking, to automatically build up tower-by-tower signal-rate histograms. For every tower there is a counter that monitors occurrences of energy above a configurable threshold. These data can be read out regularly via VMEbus to the crate-controller CPU, and from there to a dedicated monitoring computer where they are assembled in one data record covering all 124 PPMs. The rates can be displayed as histograms, thus providing a view of activity in the detector entirely free of trigger bias. This can help in finding and understanding malfunctions of the detector itself, as well as the LHC beams.

Downstream of the PPr-ASIC, the final stage on the MCM is to serialise the data to be transmitted to the CP and the JEP, since transmitting all 64 channels on a PPM in parallel would require an impossible number of connections and cables. The serialisation is done at 400 Mbit/s (480 MBaud including protocol bits) using two National Semiconductor LVDS serialisers for the four CP trigger towers (section 4.5 explains the additional ‘trick’ used to reduce the number of links from four to two), and one serialiser for the 0.2×0.2 jet data. In all cases, the LVDS 10-bit data includes an odd-parity bit to allow error-checking.

4.3 Output signals

The LVDS data streams need to be carefully handled in order to ensure that they reach the algorithmic processors with a negligible error rate. In addition, data near the edges of quadrants in azimuth have to be fanned out to more than one processor crate in order to allow the trigger logic to span these boundaries without any loss in efficiency (see section 5). This is done by sending the LVDS serial data streams from all 16 MCMs on a PPM to a small daughter card

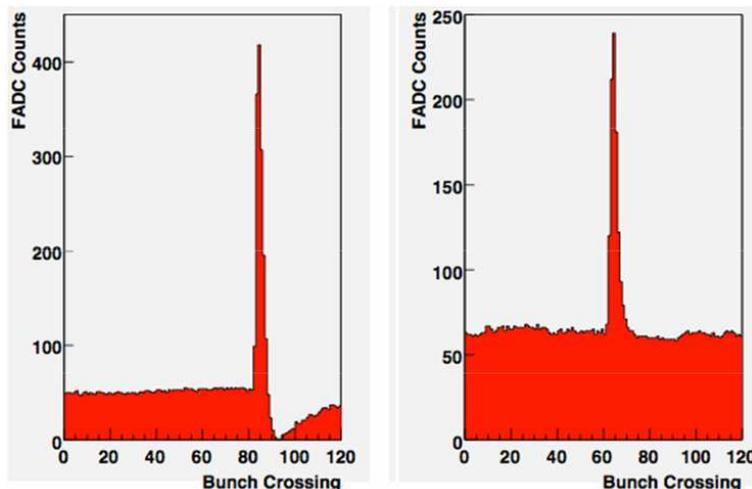


Figure 10. Calorimeter pulses digitised at 40 MHz by the PPM: LAr (left) and Tile (right).

holding four small FPGA drivers. In addition, to ensure reliable transmission down the 11 m cable links to the processors, some *RC* pre-compensation is applied to the signals. The LVDS cable links consist of 1888 assemblies of AMP-Tyco parallel-pair ‘Twinax’ cables, with four links per assembly. Tests have demonstrated bit-error rate limits on these links of $<10^{-14}$.

The data to be read out to DAQ from the FADC digitisations and look-up table outputs, for a programmable number of bunch-crossings surrounding the one of interest, are assembled and formatted by an FPGA (Xilinx XCV-1000E) for transmission to the Readout Driver modules. The readout data are sent to a small transition module mounted behind the crate backplane, which does parallel-to-serial conversion using an Agilent HDMP-1022 G-Link chip running at 640 Mbit/s followed by conversion from electrical to optical signal outputs. The same Xilinx FPGA also controls the PPM configuration and its VMEbus interface. Readout of PPM data via VMEbus is especially useful for commissioning and diagnostics.

4.4 Bunch-Crossing Identification (BCID)

The calorimeter trigger-tower pulses have rise times of the order of 50 ns, as seen in figure 10. This shows calibration signals from both the Liquid Argon Barrel Calorimeter and the TileCal, as digitised by a PPM but with many extra samplings read out. The sampling rate is 40 MHz, i.e. 25 ns. LAr signals are bipolar and have a long, negative undershoot after the peak, while the TileCal signals are unipolar.

It is of the utmost importance to associate trigger-tower signals with the correct LHC bunch-crossing, but since the pulses are several bunch-crossings in width a robust way to do this is required. Signals down to the lowest possible energies, pile-up at high luminosity, and very large saturated signals (above about 200–250 GeV) that might signal new and important physics must all be treated efficiently. The PreProcessor ASIC implements three separate methods for doing this: one for unsaturated pulses, a second for saturated pulses, and a third analogue method that provides useful redundancy especially when tuning the parameters of the first two methods.

The main method, used for normal unsaturated signals, is a digital pipelined finite-impulse-response (FIR) filter. A block diagram is shown in figure 11. The aim is to ‘sharpen’ the pulse before putting it through a peak finder. This is done by multiplying five consecutive samples by pre-defined coefficients and summing the resulting values. The coefficients are optimised for the pulse shape in each type of calorimeter. The peak-finder then compares the

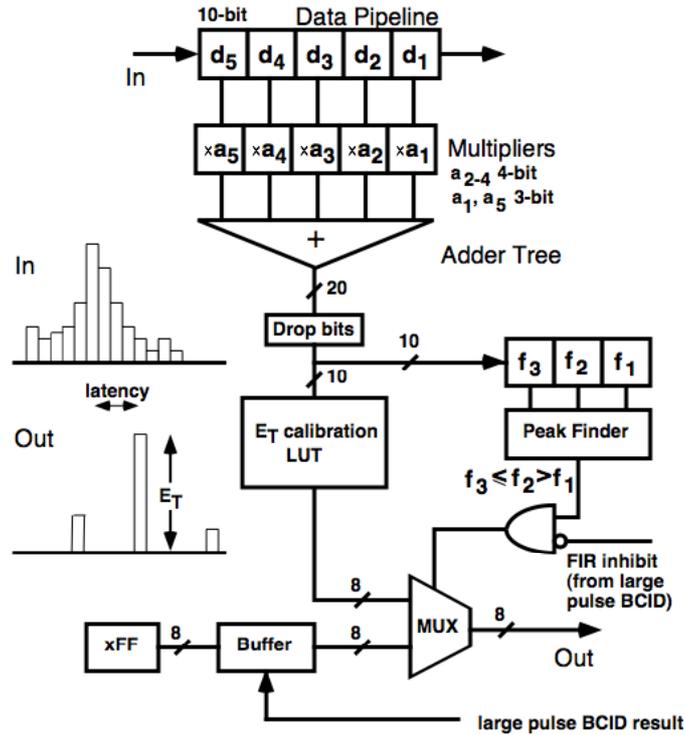


Figure 11. Finite-impulse-response filter for identifying the bunch-crossing. A typical input pulse and resulting output are shown.

sum with the values from the previous and following bunch-crossings, and looks for a maximum. Note that, since digital values can be equal, the comparison in one direction is ‘greater than’ while in the other it is ‘greater than or equal’. Simulation including pile-up and noise indicates that this method works efficiently over a signal range down to a few GeV and up to within a few GeV of the saturation level.

Saturated signals are handled by using two comparators on the leading edge of the pulse. A ‘low’ and a ‘high’ threshold are defined, based on the rise-time of the signals. This allows an estimate of when the peak would have occurred if the signal had not been saturated. Simulation indicates that this method works well from about 200 GeV up to the maximum energy range of the calorimeters.

The third method is primarily for checking the consistency of the first two methods, and to help in tuning the programmable parameters. It uses discriminators with programmable thresholds on the analogue signals, and is implemented on the analogue-input daughter cards. Since the peaking time for each type of calorimeter is known, bunch-crossing identification can be done using a programmable delay in the PPr-ASIC. The range of this method starts in the unsaturated region and extends to the maximum calorimeter energies.

A 10-bit FIR-filter result is computed on every clock cycle. For non-saturated pulses, the value is proportional to the analogue pulse size at the clock cycle identified by the peak-finder. On this clock cycle only, the FIR-filter result is sent to the look-up table to extract the final E_T value to use in the trigger algorithms. However, if the pulse is saturated then the tower is assigned the maximum 8-bit value of 255 GeV on the appropriate clock cycle. For clock cycles not corresponding to a pulse peak the tower E_T value is set to zero.

4.5 Bunch-Crossing multiplexing (BC-mux)

The data from the PPMs consist of 8-bit trigger-tower transverse energies to be sent to the Cluster Processor, and 9-bit 0.2×0.2 trigger-tower sums for the Jet/Energy-sum Processor. A ‘trick’ is used for the CP data in order to reduce the enormous number of data links required. Because the bunch-crossing identification uses a peak-finding scheme, any bunch-crossing with data in a given trigger tower must always be followed by one that is empty, i.e. zero. This can be used to allow two towers being sent to the CP to share a single serial link.

Trigger towers are paired up at the output stage of the PPr-ASIC. When a tower has a non-zero value, this is transmitted on the link along with a flag bit that indicates which of the two towers is being transmitted first. On the next bunch-crossing a value (or zero) for the other tower of the pair is sent to the link, again with a flag bit. This second flag bit is used to indicate whether the second tower’s value belongs to the same bunch-crossing as that of the first tower, or to the following bunch-crossing. This scheme is called bunch-crossing multiplexing, or ‘BC-mux’. By using it, data transmission to the CP is achieved with only two links per MCM instead of four. Note that in addition to the flag bit, an odd-parity bit is used for error detection.

For the JEP the sum of four towers is transmitted, so it is not the case that a non-zero value must be followed by zero, and this scheme cannot be used.

5. The Cluster and Jet/Energy-sum Processors

Many of the main functions of the Cluster Processor and the Jet/Energy-sum Processor are similar, and their designs take advantage of this by adopting similar architecture and utilising some common hardware modules. The electron/photon (e/γ) and τ /hadron algorithms in the CP and the jet algorithm in the JEP all search for features in overlapping, sliding windows. To avoid dips in efficiency at the boundaries of modules and crates a large amount of data duplication is required. Both of the processors divide the calorimeters into four quadrants in azimuth. As already mentioned in section 4.3, across the boundaries between quadrants efficiency is maintained by fanning out the trigger-tower and jet-element data sent from the PreProcessor. Within azimuthal quadrants, data are fanned out between neighbouring trigger modules by serial data transmission on crate backplanes; mapping of trigger towers to processor modules is arranged so that any module only needs to share data with the two adjacent modules, over short (~ 2 cm) links. This arrangement is such that no signal from the PreProcessor has to be sent to more than two crates. Thus the architecture minimises the number of cable links from the PreProcessor, and the backplane fan-out is also simplified.

The CP is a four crate system. Each crate contains 14 Cluster Processor Modules (CPMs) and handles one calorimeter quadrant, as shown in figure 12. The JEP is a two-crate system, with each crate containing 16 Jet/Energy Modules (JEMs). Eight JEMs handle one calorimeter quadrant, while the other eight handle the quadrant opposite in ϕ . Within a quadrant, each CPM or JEM covers a relatively narrow slice in η and 90° in ϕ . By careful design, both CP and JEP crates use the same high-density custom 9U backplane. Results from the CPMs or JEMs are sent via the backplane to two Common Merger Modules (CMMs), one at each end of the main block of modules in each crate. The CMMs, which use identical hardware but different firmware loads, process the results from the CPMs or JEMs to produce sums over the entire crate, and send these results to a subset of the CMMs in order to produce system-wide results. These final results are sent on cables to the Central Trigger Processor.

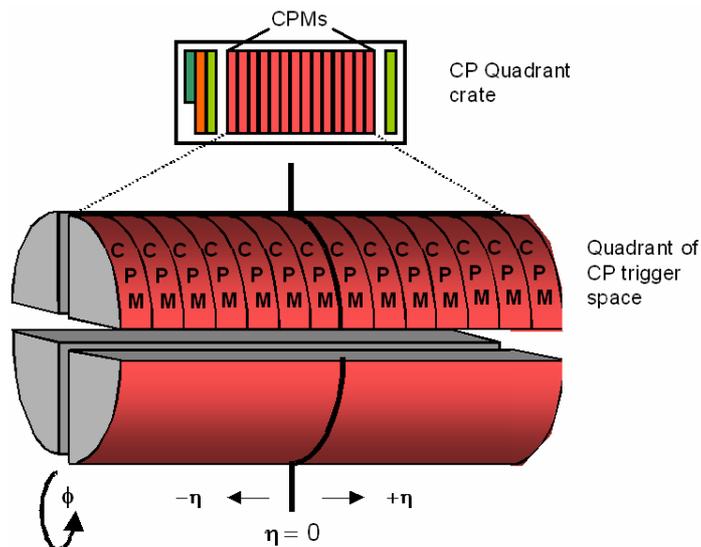


Figure 12. Layout of one crate of Cluster Processor Modules, covering one quadrant in ϕ . (In the Jet/Energy-sum Processor each quadrant is covered in a similar way, but by 8 JEMs.)

The calorimeter data handled by the CP extends out to $|\eta| < 2.5$, which is the limit of high-precision data from the Inner Detector and the EM calorimetry. The jet trigger extends further, to $|\eta| < 3.2$, which is the limit of end-cap acceptance. The E_T^{miss} and total- E_T triggers also include the forward calorimetry (FCAL), primarily to provide adequate E_T^{miss} performance. This extends the trigger to $|\eta| < 4.9$, and also allows the FCAL to be used for forward-jet triggers.

In addition to CPMs or JEMs and the CMMs, each of these 9U crates also contains a Timing Control Module to distribute the LHC clock and to monitor voltages and temperatures via CANbus, and a 6U VMEbus CPU module.

5.1 The Cluster Processor Module (CPM)

5.1.1 The electron/photon and τ /hadron algorithms

The function of the CPMs is to carry out the e/γ and τ algorithms and to count the multiplicity of successes, or hits, in the region covered by each module. The two algorithms use very similar logic, and are therefore executed together. Figure 13 illustrates the elements of the algorithms, which are run for all possible sets of overlapping 4×4 trigger-tower windows.

The e/γ algorithm searches for narrow, high- E_T showers in the EM calorimeters. The main background is an overwhelming rate of hadronic jets. Therefore, the characteristics used to enhance the selection at level-1 are to require transverse isolation, and that the showers should not penetrate to the hadronic calorimeter. The τ /hadron algorithm looks for τ decays into collimated clusters of hadrons, again permitting some level of isolation but in this case allowing the showers to penetrate into the hadronic calorimeters.

Consider the 2×2 trigger-tower region at the centre of the 4×4 trigger-tower window shown in figure 13. In the EM calorimeter E_T values are summed for the towers in each of the four possible 1×2 and 2×1 pairs within the region, in order to find relatively narrow showers while at the same time not losing efficiency for showers crossing tower boundaries. We do not worry about showers crossing into three or four of the towers because Monte Carlo studies [5] have shown that there is no significant loss in efficiency by summing only two towers, while the

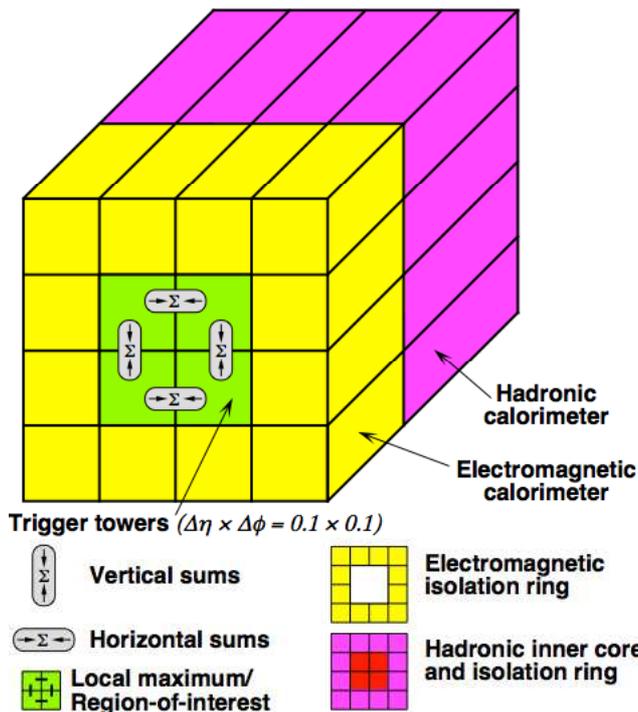


Figure 13. Elements used for the e/γ and τ /hadron algorithms.

improved selectivity in demanding narrow showers is useful. In the case of the e/γ algorithm at least one of the four sums is then required to pass a programmable ‘cluster’ threshold in E_T . For the τ /hadron algorithm each of the four EM 1×2 and 2×1 pairs is added to the sum of the 2×2 ‘core’ towers in the hadronic calorimeter, and at least one of the four sums is required to pass a threshold.

For the isolation requirements, the E_T values for the 12 EM towers surrounding the central 2×2 region are summed and required to be less than a programmable ‘EM isolation’ threshold. The E_T values for the 12 hadronic towers surrounding the central 2×2 core region are also summed and required to be less than a programmable ‘hadronic isolation’ threshold. For the e/γ algorithm only, to ensure that the shower is contained in the EM calorimeter, the sum of the core 2×2 hadronic region must be less than a programmable ‘hadronic veto’ threshold.

The ‘cluster’ threshold requirement is ‘greater than’, so that setting the threshold to its maximum value of 255 GeV makes it impossible to satisfy, and so turns it off. In a similar way, the various isolation thresholds require ‘less than or equal to’, so that setting a threshold to its full-scale value of 63 GeV effectively turns off that particular isolation requirement.

The isolation thresholds for both algorithms are fixed values, rather than ratios of isolation energy to cluster energy. Here too, physics studies [5] showed that using this approach, which is much simpler to implement, does not significantly decrease performance. In practice, the trigger menus will set much less demanding isolation thresholds (or none at all) for very energetic objects, while lower thresholds will need to have stricter isolation criteria in order to control the rates at the expense of some signal loss.

The CP provides 16 of these combinations, or sets, of cluster threshold and isolation conditions. Eight of the sets are for e/γ triggers, while the other eight can each be programmed to carry out either the e/γ or τ algorithms. For each set the cluster and isolation/veto thresholds

can all be chosen independently. It is important to note that because the sets are a combination of cluster and isolation thresholds, objects that pass one threshold set will not necessarily pass one with a lower cluster threshold.

If any trigger tower is saturated it could indicate new physics, and should produce a trigger. At least one threshold set should allow for this by not requiring isolation.

\geq	$>$	$>$
\geq	R	$>$
\geq	\geq	$>$

Figure 14. Local maximum test. The η axis is horizontal, the ϕ axis vertical.

An obvious problem is that it is possible for an object to satisfy the algorithm in more than one overlapping trigger-tower window. For example, a very clean electron or photon shower in just one trigger tower would satisfy the algorithm in four adjacent windows. This multiple counting, and also the ambiguity of the coordinates of the object to use for its region-of-interest, are avoided by requiring that the sum of the inner 2×2 region must be a local maximum compared to its eight overlapping nearest neighbours. The possibility of

comparing equal digital values must again be considered, so the method used for the overlapping windows uses four ‘greater than’ conditions (in the $+\eta$ and $+\phi$ directions) and four ‘greater than or equal’ conditions (in the $-\eta$ and $-\phi$ directions), as illustrated in figure 14. (R represents the region being tested, while the squares around it represent the overlapping adjacent 2×2 regions.) It should also be noted that for *both* the e/γ and τ /hadron algorithms the 2×2 sum includes both the EM and hadronic calorimeters; although it would be strictly correct to use only the EM sum for the e/γ algorithm, it has been shown in simulations that the simpler procedure of using the same sum as for the τ algorithm makes no real difference to performance. The coordinates of the local maximum defined in this way are used in the RoI.

5.1.2 Data input and fan-in/out

Each CPM processes 64 of the 4×4 overlapping windows described above. These are arranged in a 16×4 ($\phi \times \eta$) array (see figure 12), covering 90° in ϕ and 0.4 in η . The entire $|\eta| < 2.5$ range covered by the CP therefore requires a minimum of 13 CPMs per quadrant, but we have used 14 in order to match the PPMs, which are arranged symmetrically around $\eta = 0$. Thus the CPMs at the ends of each crate have a number of unused windows, and those that are used do not have all of their input towers populated. A block diagram of the CPM’s real-time data path is shown in figure 15.

If we regard one of the four inner towers as a reference (the lower-left one, see figure 16 upper left), then the reference towers for the 64 windows processed on the CPM are received directly from the PreProcessor, and are called ‘fully processed’. The remaining towers required to complete the windows are called the ‘environment’ (figure 16 lower left). Environment towers that are adjacent in ϕ to the fully processed ones are also received directly from the PreProcessor. The remainder are supplied by fan-in from the two adjoining CPMs in a serial format via the backplane.

In order to complete a 4×4 window around a reference tower, three additional rows and columns in both ϕ and η are required, for both the EM and hadronic layers. Thus, to process $M \times N$ windows requires $(M+3) \times (N+3) \times 2$ towers, which for a CPM handling 16×4 windows means $19 \times 7 \times 2$. However, the BC-mux scheme pairs towers in ϕ , thereby requiring $20 \times 7 \times 2 = 280$ towers and ignoring the row at the $-\phi$ end (figure 16 right). Four of the seven ϕ

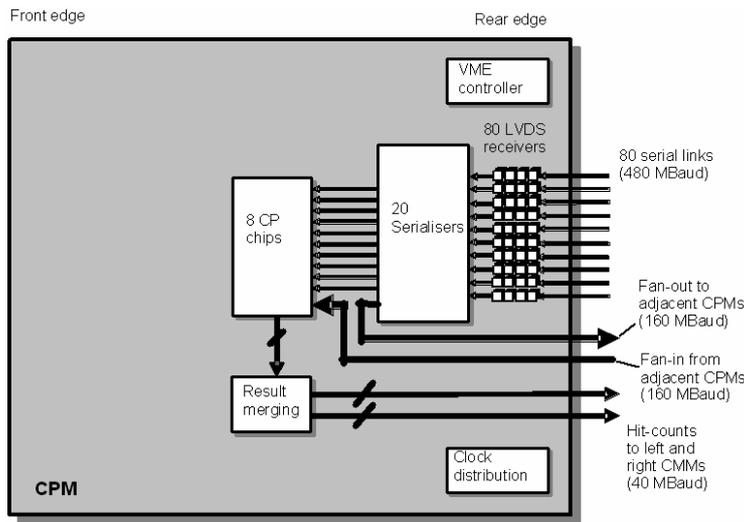


Figure 15. Block diagram showing the real-time data path of the CPM.

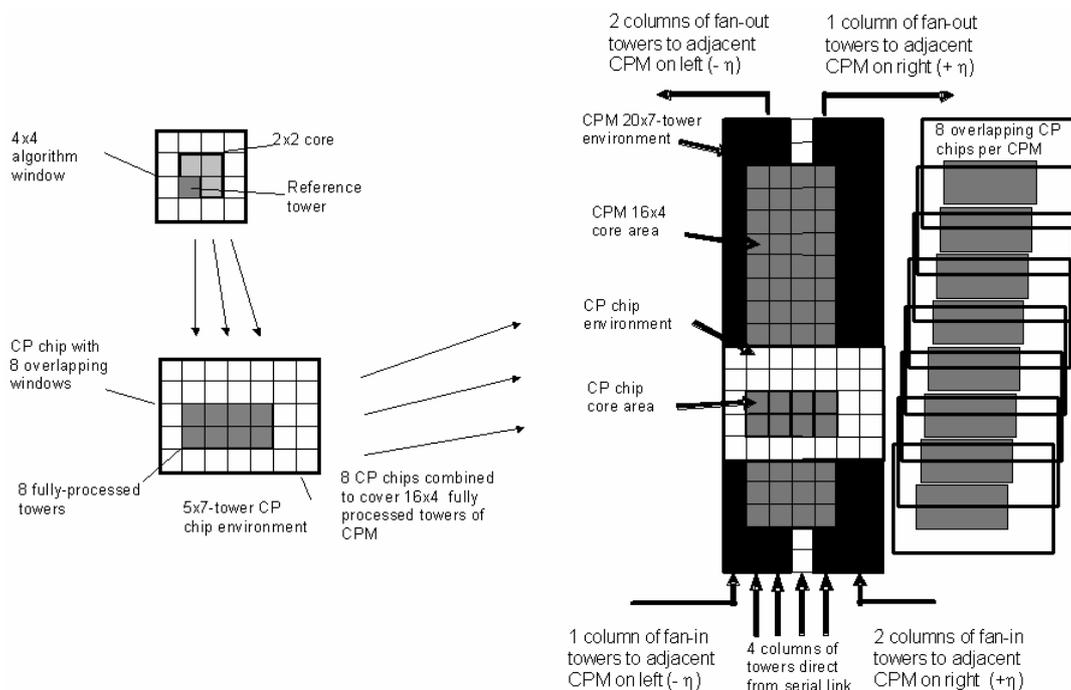


Figure 16. Illustration of algorithm windows and their handling on a CPM.

columns come directly from the PreProcessor, i.e. 80 EM and 80 hadronic trigger towers. The remaining 120 towers come as two columns from the neighbouring CPM at $+\eta$, and one column from the neighbouring CPM at $-\eta$. At the same time, one column is fanned out to the $+\eta$ neighbour and two to the $-\eta$ neighbour. Thus, each CPM shares three-quarters of its direct input data with its immediate neighbours.

The serial data transmitted directly from the PPM to the CPM pass through the backplane and are converted to parallel format by an LVDS de-serialiser stage, then go to 20 so-called ‘serialiser’ FPGAs. These unpack the data and then re-serialise them at 160 Mbit/s for transmission single-ended to both the on-board algorithmic ‘CP chips’ and, for the fanned-out

towers, to the neighbouring modules via the backplane. Each of these 160 Mbit/s streams carries a 4-bit nibble from the 10-bit BC-mux trigger-tower words, and so five data streams are needed for four towers and three data streams for two towers. This means that the fan-in and fan-out require a total of 320 connections between each CPM and the backplane. These data streams finally arrive at the CP chips, where they are de-serialised and unpacked from the BC-mux format for use in the algorithms. The use of a 160 MHz clock demands that the signals, both on-board and from neighbouring modules, must be carefully timed to nanosecond precision.

5.1.3 Processing

The eight CP chips on each CPM are large FPGAs (Xilinx XCV-1000E). Each chip processes eight algorithm windows, arranged as 2×4 in $\phi \times \eta$ and thus requiring towers from a 5×7 -tower region. The eight CP chips therefore form a 1-dimensional array in ϕ (see figure 16 far right), with their core towers adjacent in ϕ and their environments overlapping.

The algorithms run at the LHC clock frequency of 40 MHz. Because of the local maximum requirement that is used to avoid double-counting of hits and also to define the RoIs, only one of the four windows in each half of an FPGA can produce a hit for a given threshold set, so there is a maximum of two hits per FPGA. The results from each FPGA therefore consist of two 16-bit words, to indicate whether each of the 16 threshold sets was satisfied for each half of the FPGA.

The overall result for an entire CPM is the sum of the eight CP-chip results, in the form of multiplicity counts of hits for each of the 16 threshold sets. These are limited to three bits each, so each multiplicity count saturates at seven. These results are sent over the backplane to the two Common Merger Modules in the crate, each of which counts the overall crate results for eight of the threshold sets. The final results are sent to each CMM as a 25-bit word, the 25th bit being odd parity.

5.1.4 Readout to DAQ and Level-2

In order to be able to diagnose and monitor the performance of the CPMs, it is possible to read out their input data and their results to the DAQ for recording, and for online or offline analysis. In addition, RoI information must be sent to L2. These data are available for any bunch-crossing that has led to an L1A. The readout occurs after L1A is received, and is therefore not part of the synchronous real-time data path.

The input data consists of the trigger-tower values as received in the serialiser FPGAs, while the results consist of the hit data from the CP chips. The readout functions are controlled separately for DAQ and RoI data by Readout Controller (ROC) logic on two FPGAs. The data are held in FIFOs, and on receiving an L1A the data corresponding to the relevant bunch-crossing are transferred to high-speed serial readout links using Agilent HDMP-1022 G-Link chips running in 20-bit mode at 800 Mbit/s. This is followed by conversion from electrical to optical signal outputs, and transmission on optical fibres from the front panel to the RODs, as illustrated in figure 17.

The readout to DAQ can be programmed to include or omit both types of data. In addition to the bunch-crossing of interest the number of adjacent bunch-crossings can also be selected — this is useful for setting up timing, and for diagnostics. The RoI readout for L2 has a somewhat different format, as it must specify the coordinates of RoIs, and is limited to the bunch-crossing of interest. It also includes a bit indicating that saturation occurred.

The FIFOs used for recording input data can also be loaded with data patterns that can be ‘played back’ to the module in place of real, external input data. This playback facility is extremely useful for testing the operation of individual modules and the downstream CMMs.

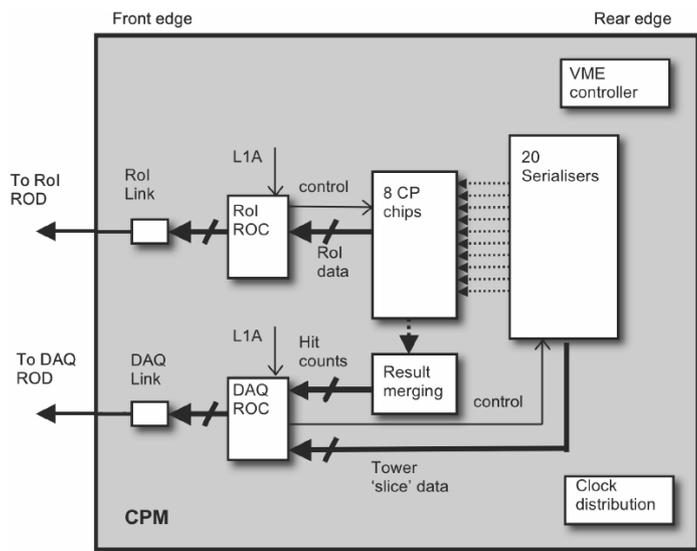


Figure 17. Block diagram of the CPM read-out (real-time data is shown as dotted lines).



Figure 18. Photograph of a CPM. From right to left, de-serialiser and serialiser chips, 8 CP chips, optical readout to DAQ and L2.

Figure 18 is a photograph of a CPM. Real-time signals enter and leave from the rear of the module (right), and readout to DAQ and L2 are from the front panel (left).

5.2 The Jet/Energy Module (JEM)

5.2.1 The jet algorithm

The JEMs carry out the jet algorithm and count the multiplicity of hits in the region covered by each module. They also serve as the first stage of the missing- E_T and total- E_T triggers, by summing the E_T components E_x and E_y , and the total E_T , over the region covered. For these purposes the granularity need not be as fine as for the e/γ and τ /hadron algorithms, and there is no need to keep EM and hadronic calorimeters separate. Therefore, the JEMs work with ‘jet elements’ that are the sum of 2×2 trigger towers in the EM calorimeters added to 2×2 trigger towers in the hadronic calorimeters, giving a basic granularity of 0.2 in $\Delta\eta$ and $\Delta\phi$. (Some jet

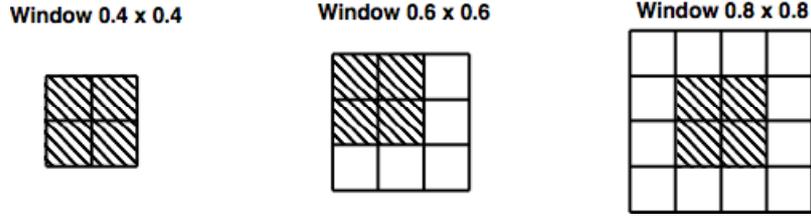


Figure 19. Jet trigger algorithm windows, based on jet elements of 0.2×0.2 and showing RoIs (shaded). In the 0.6×0.6 case there are four possible windows containing a given RoI.

elements in the outer regions of η are 0.3×0.2 .) Jet element E_T values are 10-bit words with, nominally, 1 GeV of E_T per count.

The jet algorithm has similarities to the e/γ and τ /hadron algorithms. The sums of E_T in windows consisting of 2×2 , 3×3 or 4×4 jet elements, i.e. window sizes of 0.4, 0.6 or 0.8 in $\Delta\eta$ and $\Delta\phi$, are compared to jet thresholds. The choice of window size will depend on the desired jet multiplicity: the largest window size includes more of the jet energy and therefore has the highest efficiency, while the smaller window sizes are better for resolving multiple jets. These windows are illustrated in figure 19. They overlap and slide by one element, i.e. 0.2 in $\Delta\eta$ and $\Delta\phi$. The E_T in the jet window must be greater than the 10-bit threshold; this allows the threshold to be turned off by setting it to the maximum value of 1023.

As with the e/γ and τ /hadron algorithms, it is possible for a jet to exceed the threshold in more than one window, so again it is necessary to require that a 2×2 region (i.e. summed over 0.4×0.4 in $\Delta\eta \times \Delta\phi$) must be a local maximum (figure 14) — this is used to identify jet RoI locations as well as resolve any ambiguous hits. Note that for the 0.4×0.4 and 0.8×0.8 windows there is only one possible window for each RoI, but for 0.6×0.6 there are four possible windows surrounding each RoI (figure 19), so the one with the highest E_T sum is used.

Eight independent sets of jet thresholds are available, with a nominal resolution of one count per GeV. Each threshold set is a combination of a threshold for jet E_T and a choice of jet window size. Note that because the requirements are a combination of threshold and window size, objects that pass one threshold set will not necessarily pass one with a lower jet threshold.

In addition to trigger-tower saturation, the 2×2 sums done by the PreProcessor can also overflow since these sums are limited to nine bits. In either case, the PreProcessor sets the 2×2 sum to its full-scale value of 511 GeV before sending it to the JEM. When the JEM sums the 2×2 EM and hadronic regions to form a jet element it sets the jet element to its full scale value of 1023 GeV if either the EM or hadronic regions are set to full scale. All windows in which this happens are flagged, and will automatically produce a jet trigger if any jet threshold is not set to full scale.

In the FCAL, jet elements of 0.4 in $\Delta\phi$ and summed over the whole of the FCAL η -range cover the region out to $|\eta| < 4.9$. Provision has been made for a trigger on forward jets, either in the FCAL alone or in the FCAL and end-caps, but the exact algorithm is not yet decided.

5.2.2 Energy sums

The initial summing for the missing- E_T and total- E_T triggers is also done on the JEM, but the final summing and comparison with thresholds are done on the system-level Common Merger Module, as described in section 5.3. For missing- E_T the JEM has to multiply each jet element by the appropriate geometrical constants to obtain its transverse-energy components E_x and E_y . A quad-

linear compressed format, described below, is used to send the results to the CMM. Each JEM processes data from a single quadrant, so no signed arithmetic is needed on the JEMs themselves.

The sums of E_x , E_y and E_T on JEMs are done to 12 bits, so any value above 4095 GeV produces an overflow and the quad-linear code transmitted to the CMM is set to full scale. The same thing is done if any of the input jet elements in the sum are saturated, due to either trigger-tower saturation or upstream arithmetic overflow.

5.2.3 Data input and fan-in/out

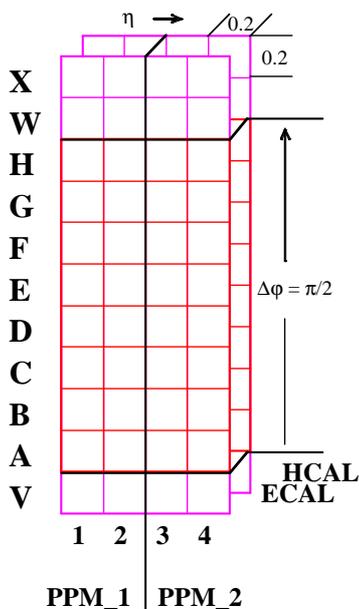


Figure 20. JEM direct-input data. Overlapping ϕ data are needed for rows X, W and V.

The 32 JEMs used by L1Calo are housed in two 9U crates. In each crate, eight JEMs handle one quadrant in ϕ and the other eight JEMs handle the opposing quadrant. As shown in figure 20, each JEM receives most of its serial direct-input data from two EM and two hadronic PPMs, covering a jet-element space of 8ϕ -bins \times 4η -bins. Overlap data from an additional four PPMs in each of the two neighbouring quadrants are required for the jet algorithm. The total number of direct-input signals per JEM is 88.

Figure 21 is a block diagram of the JEM, and figure 22 is a photograph of the module. The direct input data arrive on cables connected to the backplane, and are de-serialised to 10-bit words (9-bit E_T with an odd-parity bit) at the LHC clock rate of 40 MHz. These data are presented to a bank of input processor FPGAs, located on daughter modules (R, S, T, U). After a check for errors, the EM and hadronic values for each 2×2 element are summed into 10-bit values. The resulting jet-element values are multiplexed to 80 MHz and sent to the jet processor FPGA.

Similarly to the CPM, the jet algorithm examines 4×4 jet-element windows around each reference jet element. In order to process the JEM's core of 4×8 potential jet positions, an environment of 7×11 jet elements is required. To accomplish this, the input FPGAs send fanned-out copies of shared jet elements to the neighbouring modules via the backplane; for the JEMs this is done with a clock speed of 80 MHz. Three-quarters of the jet elements are duplicated in this manner.

5.2.4 Processing jets

The jet processor FPGA (Xilinx XC2V-3000) identifies and counts clusters of 2×2 , 3×3 or 4×4 jet elements that exceed programmable E_T thresholds and are centred around a local maximum. There are eight independent jet definitions, each consisting of a programmable threshold associated with a selectable window size. The results are sent via the backplane to one of the CMMs in the crate, designated as the jet CMM, in the form of eight 3-bit multiplicities, plus one bit of odd parity. In more detail, the steps carried out are as follows.

Jet elements are first summed to produce 60 (6×10 in $\eta \times \phi$) 2×2 clusters, 45 (5×9) 3×3 clusters, and 32 (4×8) 4×4 clusters. Cluster sums containing saturated jet elements are flagged. The central 32 (4×8) 2×2 clusters are compared with their nearest neighbours to determine whether they are local maxima, and therefore possible jet candidates.

automatically pass all enabled thresholds. For JEMs covering only the barrel and end-cap calorimeters, 3-bit multiplicities of jet clusters satisfying each of the eight jet definitions are produced. JEMs that also include FCAL elements produce eight 2-bit central jet multiplicities, and four 2-bit FCAL jet multiplicities.

5.2.5 Processing energy sums

Each JEM sums the total scalar E_T of all jet elements that arrive directly from the PreProcessor. It also multiplies each jet element by the appropriate geometrical constants to obtain its transverse-energy components E_x and E_y . The algorithms operate on the 32 core jet elements, and are implemented partially in the input processors and partially in the sum processor FPGA (Xilinx XC2V-2000). A low threshold is applied to the jet elements entering the E_x and E_y adder trees, and a separate threshold is applied to data entering the total- E_T adder tree; both of these are to reduce the effects of noise. The jet elements are converted to x and y components by multiplication with cosine and sine of ϕ , respectively. The resulting sums of E_T , E_x and E_y are transmitted via the backplane to the second CMM in the crate, designated as the energy summation CMM.

Energy summation is performed with a precision of 1 GeV for the total E_T . For E_x and E_y summation is done using 12-bit multipliers and products, working to 0.25 GeV precision before the result is rounded to the nearest 1 GeV. The energy scale goes up to 4095 GeV. Signals exceeding full scale are saturated to 4095 GeV.

A data compression technique is used in order to be able to use the same number of energy-sum connections via the backplane to the CMM as for the cluster and jet multiplicities. The three energy-sum words are transmitted as eight bits each, plus one odd-parity bit. To achieve this a quad-linear compression scheme is employed, using a 6-bit mantissa and a 2-bit exponent that multiplies the mantissa by 1, 4, 16 or 64 to yield the full value transmitted. In cases of saturation in the jet elements or along the summing chain the 8-bit code is set to full scale (0xFF). This produces 4032 GeV when decoded.

5.2.6 Readout to DAQ and Level-2

As with the other types of modules in the trigger, provision is made in the JEM design for extensive monitoring via the data acquisition system. All jet-element input data arriving directly on cable links from the PreProcessor, and all jet and energy-sum results sent to the CMMs, are captured for each accepted event. The buffers can hold data for up to five bunch-crossings around the one that caused the event, an option that is primarily useful in setting up the timing or diagnosing problems. In addition, RoI data on jet hits are captured for just the bunch-crossing of interest.

On receiving an L1A, the data corresponding to the bunch-crossing of interest are transferred to high-speed serial readout links using 16-bit Agilent HDMP-1032 G-Link chips running at 640 Mbit/s. This is followed by conversion from electrical to optical, and the signals are sent from the front panels on optical fibres to the DAQ and RoI Readout Drivers.

For diagnostic purposes, additional playback memories upstream of the main JEM logic and spy memories for testing are provided. The playback memories can be filled with test patterns under VME control.

5.3 The Common Merger Module (CMM)

5.3.1 Requirements and architecture

The CMMs ‘merge’ results from entire crates of CPMs or JEMs by counting the total number of e/γ , τ /hadron and jet hits, and by summing the total E_T , E_x and E_y . A second stage of logic

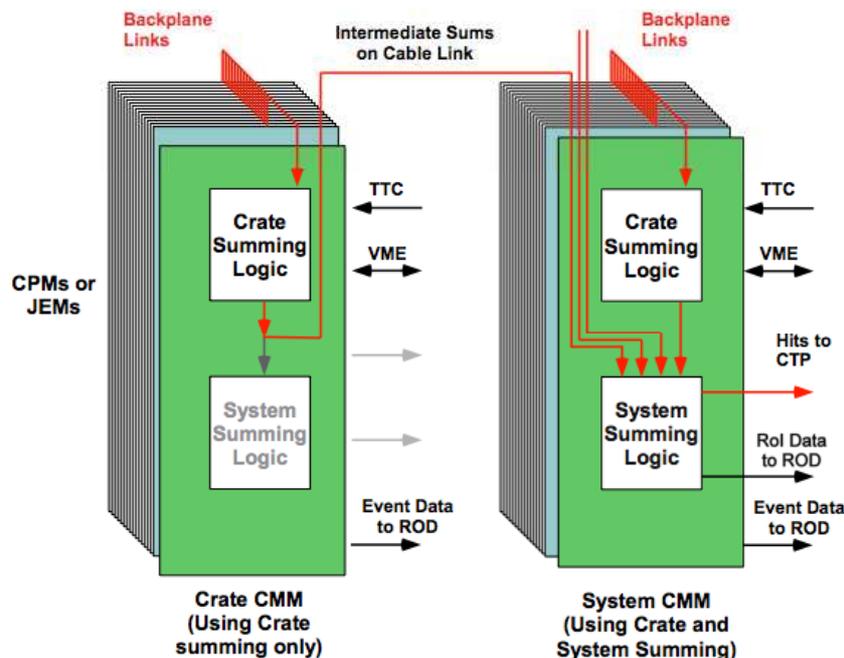


Figure 23. CMM connectivity, including crate-level and system-level functionality.

allows a subset of the CMMs to perform system-wide counts of the various types of hits, and to do the system-wide energy sums. At the system level the total- E_T and missing- E_T are compared to thresholds, along with an approximation to the total E_T in jets. All this is achieved with a single hardware module design that runs different versions of firmware in its FPGAs. This is facilitated by the use of a common, custom backplane for both CP and JEP crates (section 5.4).

The CPM and JEM results used to produce the crate results are transmitted to the two CMMs in each CP and JEP crate via the backplane as parallel, single-ended point-to-point CMOS signals running at 40 MHz. Each CMM receives up to 400 bits of data per bunch-crossing from the 14 CPMs or 16 JEMs in each crate. A large FPGA performs crate-level merging. For each of the four types of CMM (the two groups of eight threshold sets per crate in the CP, and the jets and energy sums in the JEP), the crate-level results are transmitted to a single CMM which performs the system-level merging on a second FPGA. Both FPGAs are Xilinx XCV-1000E. Transmission of the crate results is via short, parallel-data LVDS cables connected to small cards behind the backplane, except for the crate results from that CMM itself which are transmitted internally on the module. A block diagram of the basic functionality of the crate-level and system-level CMMs is shown in figure 23. Figure 24 is a photograph of a CMM.

5.3.2 Hit counting

The CP has 16 threshold sets (i.e. combinations of cluster and isolation thresholds), eight reserved for e/γ triggers and eight individually programmable to function either as e/γ or as τ /hadron triggers. For each threshold set, the number of hits above threshold is summed to a maximum of seven (i.e. three bits) over the entire detector. One CMM per crate handles the first eight threshold sets, and the second CMM per crate handles the other eight. The two CMMs in one of the four crates act as ‘system’ CMMs, one for the eight designated e/γ triggers and the other for the eight e/γ or τ /hadron triggers. The two ‘system’ CMMs first form their own crate sums, then add in the other three crate sums (received on cables), and send the results to the Central Trigger Processor. Each system CMM sends eight 3-bit multiplicities to the CTP, plus a parity-check bit.

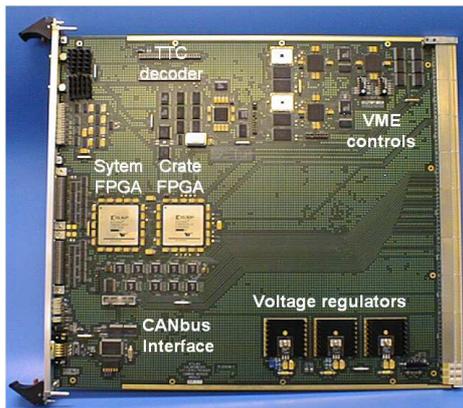


Figure 24. Photograph of a CMM.

Similarly in the JEP there are eight jet threshold sets, each a combination of jet threshold and jet-window size. The number of hits for each threshold set is summed up to a maximum of seven over the entire detector excluding the forward calorimeters. (A programmable option to include the forward calorimeters in the overall jet trigger is available.) These crate sums are done by one CMM in each of the two JEP crates, and one of these acts as ‘system’ CMM. The system CMM sends eight 3-bit multiplicities to the CTP, plus a parity-check bit.

A more limited counting capability is required for jets in the forward calorimeters, and this is done in the same CMMs. At present, four forward-jet thresholds are foreseen, with multiplicities up to three counted separately at each end of the detector for a total of 16 bits.

5.3.3 Total- E_T and missing- E_T summing

The three energy sums, E_x , E_y and E_T , from each JEM are sent to the second CMM in each of the two jet crates, as 8-bit words using a quad-linear code (see section 5.2.5), thus using the same number of input-data bits as the jet multiplicities. On arrival at the CMM the linear form of the summed energy is recovered, and the data from all the JEMs in the crate are summed. The system-wide sums are done by one of the CMMs, again designated the ‘system’ CMM. These transverse-energy sums cover the entire detector, including the forward calorimeters.

The total- E_T sum is compared to four thresholds, with values up to 2 TeV in 4 GeV steps. Four bits, indicating which thresholds were passed, are sent to the CTP.

The ϕ -quadrant architecture of the trigger means that the E_x or E_y components from any one JEM always have the same sign values, which are therefore not transmitted. Since each JEP crate handles two opposing quadrants in ϕ , the component sums from each quadrant (i.e. 8 JEMs) simply have to be subtracted.

A look-up table is used to perform the final quadrature addition of E_x and E_y , as well as compare with eight thresholds in a single step. In order to cover as wide a dynamic range as possible while maintaining precision at the low-energy end, four different energy ranges are provided in the look-up table, with the choice based on the value of the larger of E_x and E_y . The thresholds have step sizes from 1 GeV to 8 GeV, depending on the range. Eight bits, indicating which thresholds were passed, are sent to the CTP.

5.3.4 Total jet transverse energy

Extra logic in the jet ‘system’ CMM calculates an approximation to the total transverse energy in jets, E_{Tj} , by multiplying the multiplicity of jets exceeding each threshold by an energy value

Table 1. L1Calo results sent to the Central Trigger Processor.

Cable origin	Number of bits	Type of information
Cluster Processor CMM 1	8 thresholds \times 3 bits	e/γ multiplicities
Cluster Processor CMM 2	8 thresholds \times 3 bits	τ /hadron or e/γ multiplicities
Jet/Energy-sum Processor CMM 1	8 thresholds \times 3 bits	Jet multiplicities
Jet/Energy-sum Processor CMM 1	4 thresholds \times 2 bits \times 2 ends	Forward-jet multiplicities
Jet/Energy-sum Processor CMM 1	4 thresholds \times 1 bit	Total jet E_T
Jet/Energy-sum Processor CMM 2	8 thresholds \times 1 bit	Missing- E_T threshold bits
Jet/Energy-sum Processor CMM 2	4 thresholds \times 1 bit	Total- E_T threshold bits

close to that of the threshold. This estimated value is compared to four thresholds, and the results sent to the CTP.

To understand the estimator, consider an event containing m jets passing the threshold $E_T > x$ GeV and n jets passing $E_T > y$ GeV, with $m > n$ and $y > x$. There are $m - n$ jets in the E_T range x GeV $\leq E_T \leq y$ GeV. These jets have a total E_T of at minimum $(m - n)x$ GeV and at maximum $(m - n)y$ GeV. A value closer to x gives a more accurate estimate because the jet E_T spectrum falls steeply with E_T ; the value actually used is programmable.

The hardware implementation of the E_{TJ} calculation uses look-up tables to convert groups of jet counts to transverse energies. The energies are summed and thresholds applied to yield the four-bit E_{TJ} hit result.

5.3.5 Outputs to CTP, DAQ and Level-2

The CMMs send 104 bits of L1Calo results, via cables from the front panels of the four ‘system’ CMMs, to the Central Trigger Processor. These are summarised in table 1.

In common with other module types, both CMM input data and results are available for readout to the data acquisition, with an option for how many bunch-crossings to read out. In the case of the CMMs both the ‘crate’ inputs and results, and the ‘system’ inputs and results, can be read out to facilitate testing and diagnostics. Once more, the data are sent to the RODs using high-speed serial optical links running at 800 Mbit/s.

The global RoI data sent to the Level-2 trigger from the energy-summing chain consist of the total values of E_x , E_y , and E_T as well as the threshold results on missing- E_T and total- E_T . The jet chain sends the threshold results from the total jet E_{TJ} estimation.

5.4 The Processor Backplane (PB)

Both the Cluster and Jet/Energy-sum Processors have very high numbers of signals entering and leaving their CPMs, JEMs and CMMs. The dominant contributions to this are the many input signals coming directly from the PreProcessor, and the need to share the majority of these signals between neighbouring modules because of the sliding-window nature of the EM, τ /hadron and jet algorithms. Even with 9U-high modules, these requirements cannot be met using standard VME and thus demand a very high density custom-built backplane. On the other hand, the input and shared-signal requirements of the JEMs are very similar to those of the CPMs. CPM and JEM output data do differ, but a scheme was developed to handle the missing- E_T and total- E_T data going to the CMMs using the same number of bits as the other algorithms. With careful design, we could then use a common backplane for both the CP and the JEP.

The most obvious feature of the resulting backplane (see photographs in figure 25) is that this large, monolithic printed-circuit board is almost completely covered by high-density con-



Figure 25. Photographs of the Processor Backplane. Left: module side, with CPU connections in slots 1 and 2 and TCM in slot 21. Right: from rear of crate, showing cable supports for LVDS serial links into CPMs and JEMs.

Table 2. Processor Backplane slot organisation.

Position(s)	Module type
1 (and/or 2)	VME 6U CPU (controller), mounted in 9U VME Mount Module (VMM)
3	Common Merger Module (CMM)
4–19 or 5–18	Jet/Energy Module (JEM) or Cluster Processor Module (CPM), respectively
20	Common Merger Module (CMM)
21	Timing and Control Module (TCM)

nectors, with a total of about 22,000 pins. Each position in the crate is designed to service a specific type of module, as specified in table 2. The 1148 pins per CPM or JEM slot receive the serial cable links from the PreProcessor to the CPMs or JEMs, and provide point-to-point connections for data sharing between neighbouring CPMs or JEMs, point-to-point connections from each CPM or JEM to both CMMs for results merging, as well as other functions requiring fewer pins as discussed below.

The primary connector choice for the backplane is the Hard Metric (HM) family, specified by international standard IEC 1076-4-101. In addition to high density, the large number of high-speed input and output signals requires connectors with good signal characteristics and high reliability. This connector style provides connections with 2 mm pitch, five columns wide and with a choice of connector heights. Ground-return shields are added on both sides of the connectors. A signal density of 20 signals per centimetre of card edge can be attained while maintaining a signal:ground ratio as low as 4:3. This connector range was also chosen due to the availability of a cable assembly that is suitable for carrying the high-speed LVDS serial cable links from the PreProcessor system to the CPMs and JEMs. These links are brought to the modules through the backplane using long through-pins. The male connectors on the backplane have three mating levels on the front side to reduce the maximum insertion force. A guide-pin assembly helps ensure alignment during board insertion, thus reducing insertion force and minimising the risk of damage to the pins from improper insertion.

The serial-link inputs to the CPMs and JEMs are brought to the backplane via untwisted shielded-pair cable assemblies. These are commercially assembled Z-Pack HM assemblies from AMP/Tyco. Each assembly contains four differential input pairs, plus an additional ground shield per two differential cables. To isolate the high-speed LVDS signals from potential

sources of noise, the cable inputs are arranged in blocks of four cable assemblies (16 pairs), with ground pins between the LVDS signals and any other signal pins.

Signals are shared between neighbouring CPMs or JEMs by using short point-to-point links between modules. Each CPM has 320 links running at 160 Mbit/s, while each JEM has 330 links running at 80 Mbit/s. These pins occupy the majority of the processor card edge. They have been placed to preserve signal integrity and to minimise cross-talk at 160 Mbits/s by having at least one adjacent ground pin, grouping pins to have equal numbers of inputs and outputs in each block, and keeping routing short by never taking signals more than one row up or down.

Each CPM or JEM sends its real-time outputs to the two CMMs via single-ended point-to-point backplane links. Each CMM therefore receives results from up to 16 processor modules, each comprising 24 data bits and one parity bit. Since some of these signals must travel almost the full width of the crate, great care has been taken with their routing and impedance.

A commercial VME single-board processor (CPU) in slot 1 is used for configuration and control of the CP and JEP subsystems. To economise on signal pins, a highly reduced custom subset of the VMEbus protocol, called ‘VME--’, is used. This takes up just 43 pins of a high-density backplane connector, including 16 data bits and allowing 24-bit addressing.

Dedicated address pins on the backplane provide every module in the CP and JEP with information to determine their crate and slot position in the system, and thus select appropriate VME, CANbus and TTC addresses. Some modules, notably CMMs and those JEMs covering the forward regions, also load different FPGA firmware based on their geographic address.

A Timing and Control Module (TCM; see section 7.2) in slot 21 of each crate receives the Timing, Trigger and Control (TTC) information optically, and transmits it via the backplane in differential PECL format over point-to-point links to all the processor modules.

The TCM also provides the interface between the crate and the ATLAS Detector Control System (DCS). The backplane includes a differential CANbus to all modules in the crate, allowing readout of temperature and voltage-supply information through the TCM to the DCS.

6. The Readout Driver (ROD) and data readout

6.1 Requirements and architecture

All of the main modules in L1Calo have extensive and flexible capabilities to read out their input data and results, both for online monitoring and for subsequent storage and offline analysis. The DAQ readout data are intended for trigger monitoring, diagnostics and calibration, and include information collected from the inputs and outputs of the processing modules. In addition, RoI data must be collected and transmitted to the Level-2 Trigger RoI Builder (RoIB). The data inputs and outputs, and the processing requirements, are significantly different in the two cases.

All processing modules can sample a programmable number of consecutive time-slices around a particular bunch-crossing. There is a minimum effective dead-time of four bunch-crossings between L1As; this means that a maximum of five time-slices can be read out in normal physics running, although more are needed for some of the calibration with PPMs. The time offset of the readout relative to the L1A may be programmed via registers in the modules. When all modules are properly configured, the data derived from a specified LHC bunch crossing can be tracked through the complete trigger system. Readout of multiple time-slices is useful in setting up and checking the internal digital timing of the trigger, and is essential to verify the correct operation of bunch-crossing identification of the long calorimeter pulses. RoI data are read out only for the single time-slice corresponding to the L1A.

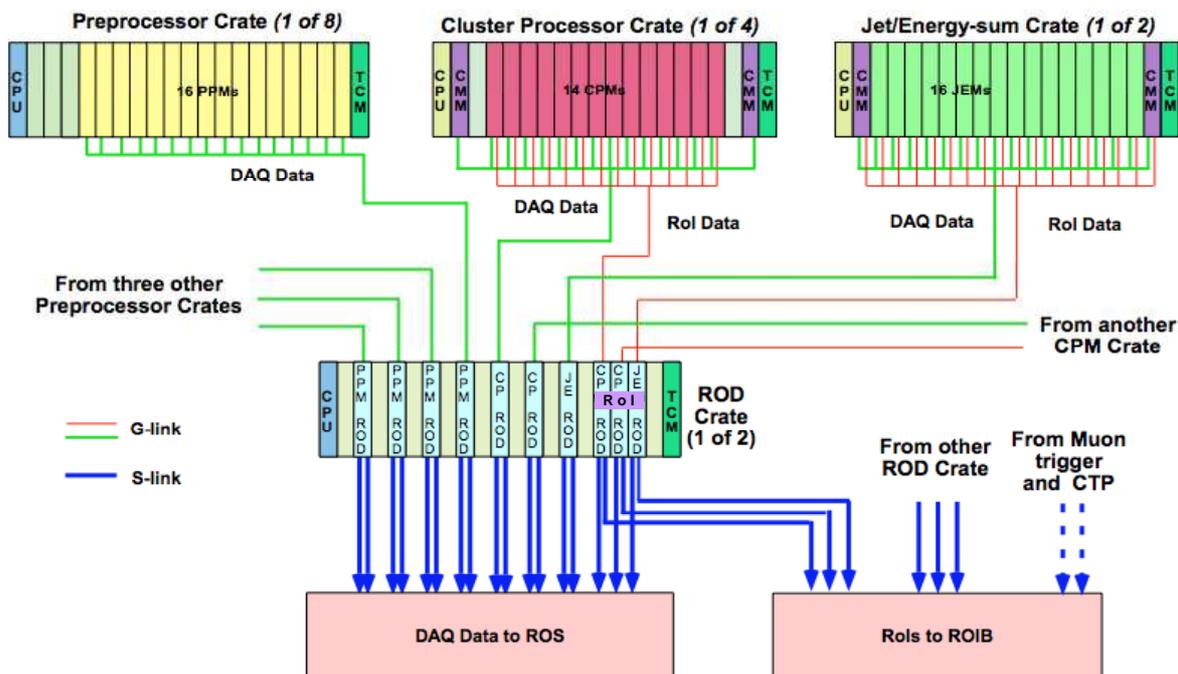


Figure 26. Diagram of the overall crate layout, with lines indicating the readout data paths. Note that CMM RoIs are only for energy sums.

A common system of readout from all the trigger modules, using high-speed serial links over optical fibres, is used. The data are collected by Readout Driver modules, which apply formatting to ATLAS standards before transmission to the DAQ and the Level-2 Trigger. A single ROD-module implementation has been made possible by careful design and the use of many different firmware versions running on the same type of hardware. This can handle the different types of data coming from the various trigger modules, for both DAQ and RoI data.

There are 124 PPMs, 56 CPMs, 32 JEMs and 12 CMMs to read out to DAQ, and there is RoI data from all but the PPMs and most of the CMMs. RODs could be organised in various ways, by readout type or by processor crate for example. We chose to have one ROD per processor crate for DAQ data and a second ROD per crate for RoI data. This allows all of the data processing within one processor crate to be monitored directly on the same ROD. The overall layout is shown in figure 26. Up to 16 optical-fibre inputs are required for each ROD handling PreProcessor or CP data, and 18 for each ROD handling JEP data. Therefore, the common ROD design has 18 inputs. Twenty RODs are needed for the entire calorimeter trigger, 14 for DAQ and 6 for RoIs. This requires two VME crates since two other modules per crate are also needed.

The data from the ROD inputs must be combined and formatted according to the ATLAS standard, and sent to the Readout System (ROS) and the Level-2 RoI Builder on standard Readout Links. These links mark the boundary of L1Calo responsibility.

6.2 Data handling

Data arrive at the RODs on optical links using a serial format generated by the G-Link chipset from Agilent. The transmitting G-Link encodes 16 or 20 bits of user data into a 20- or 24-bit frame, which it transmits serially at 800 or 960 MBaud respectively. The effective data transfer rate is 640 or 800 Mbit/s. The G-Link receiver recovers the clock and user data from the serial data stream, and also checks the framing bits to verify link stability.

Table 3. Numbers of RODs and S-Links.

Subsystem	No. of RODs	S-Links per ROD		Total S-Links	
		to DAQ	to RoIB	to DAQ	to RoIB
PreProcessor	8	4	0	32	0
CP DAQ	4	2	0	8	0
CP RoI	4	1	1	4	4
JEP DAQ	2	4	0	8	0
JEP RoI	2	1	1	2	2
Totals	20			54	6

The readout process is initiated by the L1A signal generated by the CTP and distributed by the TTC system. All processing modules copy data continuously from their 40 MHz pipelines into dual-port scrolling memories. On receipt of an L1A signal, each module extracts data from its scrolling memories into FIFOs to await readout. Each FIFO is connected via a shift register to one of the G-Link user-data pins. Logic on the module moves data from the FIFOs into the shift registers, asserts the Data Available (DAV) signal to the G-Link, and transmits the shift register contents for the event concerned. An odd-parity bit is appended to each active G-Link pin when the shift register contents have been sent. During quiescent periods the transmitting G-Link sends fill-frames to maintain the lock between the transmitter and receiver.

In the ROD, the G-Link receiver presents the recovered 40 MHz clock and parallel 16- or 20-bit user data. Firmware in the ROD collects, processes and stores the data in 32-bit-wide FIFOs with the format of the output event record. The conversion process includes checks of parity and Bunch-Crossing Number (BCN), as well as zero-suppression or data compression. The ROD also receives the L1A signal via the TTC subsystem, shortly followed by the event trigger-type needed for the ATLAS ROD fragment header. When all required information has been received from the TTC and G-Links, the ROD assembles a complete ATLAS event fragment with a header, a payload from the G-Links, and a trailer. The ROD module is equipped with four S-Links to transmit the output data to the ROS and the RoIB. The S-Links [6] provide data transport at up to 160 Mbyte/s per link.

For a ROD handling DAQ data, the data volume depends on the L1A rate, the occupancy of the detector (hence the LHC luminosity and background), the number of time-slices being read out, and the type and effectiveness of the zero-suppression or data compression algorithms. Some of these parameters are hard to predict, so the ROD must include configuration options to use from one to four S-Links to despatch the output data. At high data rates, the input G-Link data volume may temporarily exceed the bandwidth of the output S-Links, and buffers are provided to smooth the data flow.

The volume of RoI data handled by a ROD typically occupies only a few percent of the S-Link bandwidth. A single S-Link carries the RoI data from an RoI ROD to the RoIB, and a second S-Link carries a copy to the ROS. The remaining two links are unused on RoI RODs. Table 3 summarises the number of RODs, and links to the ROS and RoIB.

6.2.1 Data-rate limitation

Limits on readout data rates are imposed by the G-Links between the trigger modules and RODs, and by the S-Links between the RODs and the ROS. It is also desirable to reduce the volume of data read out to DAQ in order to reduce the mass storage required for ATLAS events.

The number of G-Links handling DAQ data is kept to one per trigger module in order to minimise both the power consumption per module and the number of RODs. (The number of RoI G-Links never needs to be more than one per module.) Once the full DAQ data rate can be handled by one G-Link per module, further data reduction can be carried out by the RODs. The amount of data read out from each trigger module can be changed only by varying the number of time-slices of DAQ data that are read out, and this is controlled on the modules. Except for the PPMs, the number of time-slices to be read out from a particular type of module is the same for all the different types of data read out from the module. The number of time-slices of the two types of PPM data can be separately controlled, as the PPM cannot read out five time-slices of both FADC and look-up table output data at the full L1A rate of 100 kHz. At 100 kHz L1A rate, up to a total of eight time-slices (e.g. 5 FADC and 3 look-up table) can be read out.

Several tools are available to keep data rates within acceptable limits. First, the number of time-slices read out from the trigger modules can be controlled. Second, the ROD can reduce independently the number of time-slices of each type of data. This may include entirely eliminating intermediate data that are not needed, especially when nominally equal to other data, e.g. at the two ends of data-links in the trigger. Third, data containing a large number of zeroes can be reduced in volume without loss of information by zero suppression. Finally, for FADC data where zero suppression is ineffective but very low data values are much more frequent than others, lossless data compression is available to be used to reduce data volume.

To prevent data loss, each ROD module in ATLAS is required to provide a ROD BUSY signal when its buffers are nearly full. This is used to generate a veto signal for the CTP. The L1Calo ROD implements this by continuously comparing the depth of data in its buffers to a programmable ‘nearly full’ threshold. The ROD asserts the BUSY front-panel output whenever this threshold is exceeded, and removes the BUSY when all buffer levels fall to the threshold or below.

When everything is running smoothly and the calibration is stable, the absolute minimum amount of data readout required consists of one time-slice of look-up table data from the PreProcessor, together with the final results sent to the CTP. This allows the nature of each level-1 trigger to be understood, and the operation of the digital part of L1Calo itself to be verified. It is useful to add the RoI readout data to this bare minimum. Reading out three or five time-slices of FADC data adds useful checks of calorimeter signal quality and bunch-crossing identification.

6.3 Structure of the ROD module

A block diagram of the ROD module is shown in figure 27, and a photograph in figure 28. The input signals enter via the front panel. Eighteen optical receivers are followed by G-Link receivers to de-serialise the data, which then go to five Input FPGAs. Four of these handle four inputs each, and the fifth the remaining two inputs. After a parity check, data compression is applied where selected, either zero-suppression or, for FADCs, lossless data compression (see above). The data are buffered, and then formatted into S-Link packets. A total of 13 different firmware versions had to be developed to run in these FPGAs, to handle both DAQ and RoI data from the different types of trigger modules, and various compression options and test options. It is possible to run different firmware versions independently for each input channel, and this is required since the ROD must be capable of handling the mixture of CPMs or JEMs and CMMs present in CP and JEP crates. The devices used are Xilinx Virtex-II Pro XC2VP20.

The Input FPGAs feed their results to a Switch FPGA. This device merges the data, and feeds it to between one and four of the output S-Links depending on the data volume and

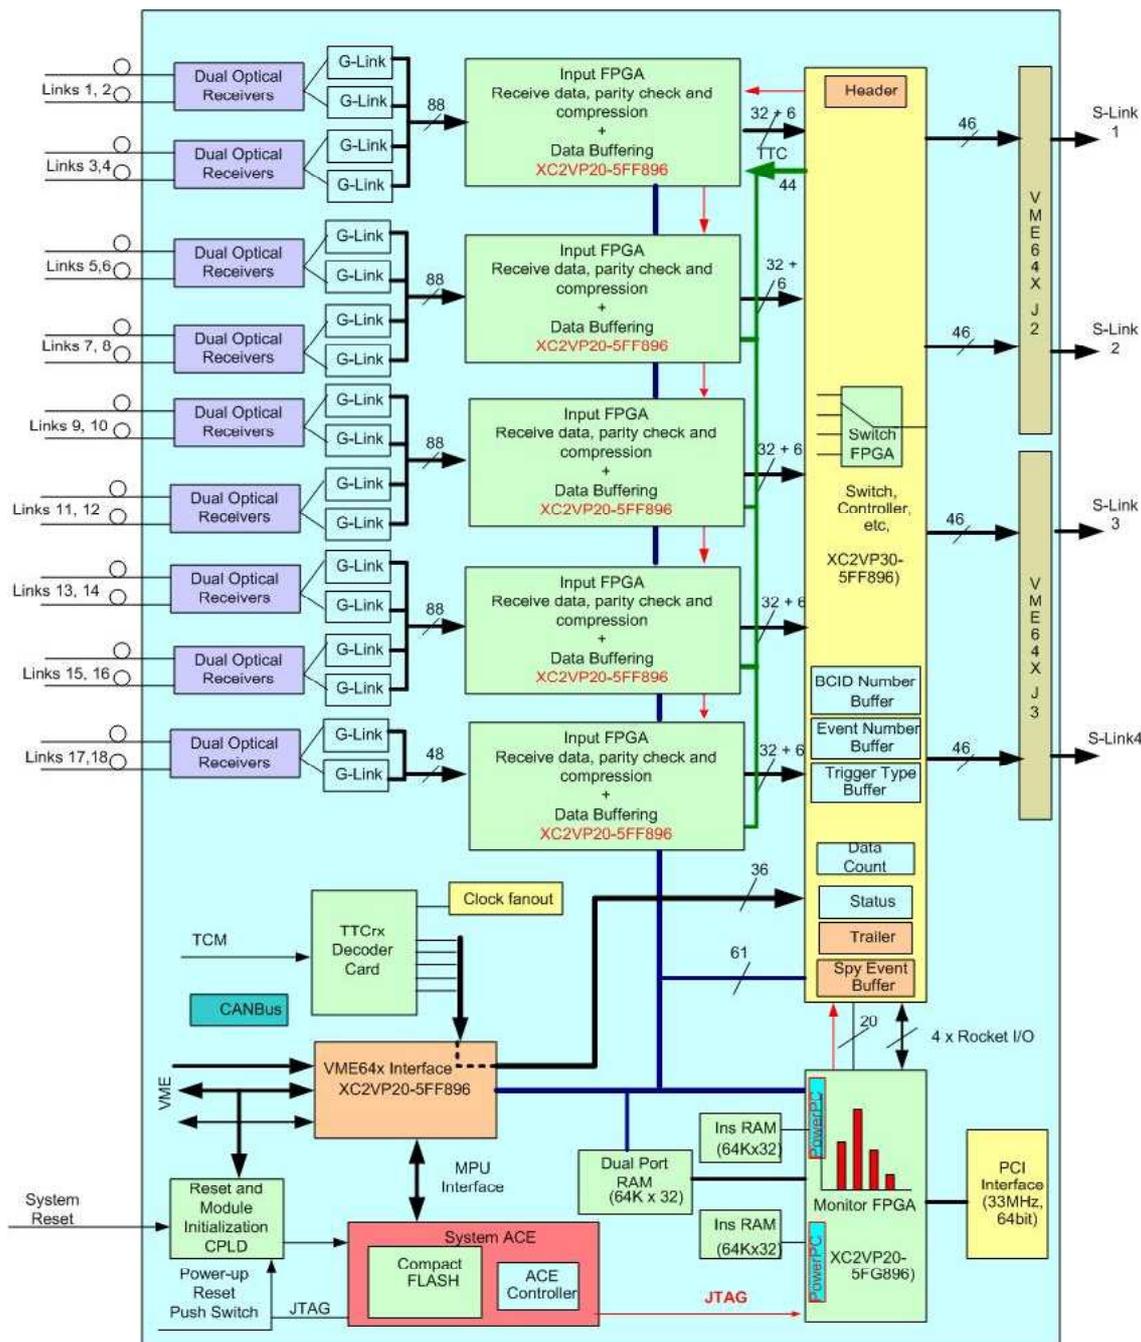


Figure 27. Block diagram of the Readout Driver module.

application. It can also sample the data and send a ‘spy’ stream off to the Monitor FPGA via one of two possible high-speed routes: 20 links that can handle a total of 3.2 Gbit/s, or Xilinx Rocket I/O. The Switch FPGA is a Xilinx Virtex-II Pro XC2VP30.

The ROD outputs to the S-Links pass through the J2 and J3 connectors on the VME backplane. The four S-Link drivers themselves are on a rear-transition module mounted behind the backplane, and are standard HOLA (High-speed Optical Link for ATLAS) daughter cards.

The Monitor FPGA provides several possibilities for spying on the data. The FPGA itself has two embedded PowerPC cores that could be used. In addition, a PCI interface has been

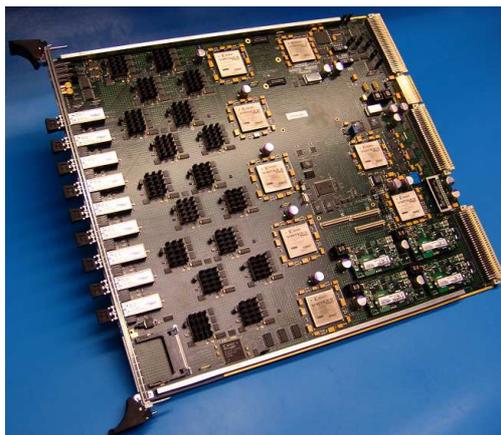


Figure 28. Photograph of a ROD. Left to right: optical receivers, G-Link chips, five Input FPGAs. At right (top to bottom): VME, Monitor, and Switch FPGAs.

added, via a dual-port RAM, to allow a PMC daughter card with its own processor should that prove useful. Finally, the data are accessible via VME. This allows the CPU controller in the ROD crate to access data from all the RODs in the crate and use these data to monitor performance. The Monitor FPGA is a Xilinx Virtex-II Pro XC2VP20.

The different firmware versions needed for the RODs are stored on Compact Flash cards, utilising the Xilinx System Advanced Configuration Environment (System ACE).

7. Infrastructure

7.1 Detector Control System (DCS)

The ATLAS Detector Control System monitors and controls the power and cooling in the electronics racks and crates in a standard way [1]. However, in the case of L1Calo we also want to be able to monitor voltage-supply levels and temperatures on the individual modules. This is because the modules are custom-built, are very densely populated in some cases (especially PPMs, CPMs, and RODs), include many large and expensive FPGA devices, and utilise a number of local voltage regulators in order to make best use of the crate power-supply arrangements. To do this monitoring, each crate utilises a facility on its Timing Control Module (TCM; see section 7.2). The TCM connects to an external CANbus within the ATLAS DCS, and it manages and monitors an in-crate CANbus that connects to all the trigger modules. On each trigger module a microcontroller connected to the CANbus monitors a collection of voltages and temperatures appropriate for that type of module, and reports results to the TCM. Alarms indicating parameters outside programmable limits are passed up through the DCS and are visible remotely.

The overall scheme is shown in figure 29. The ATLAS central DCS controls and monitors environmental conditions in all the racks. Within L1Calo, a separate CANbus monitors each crate's power supplies, temperatures and fan speeds via standard logic in the crate fan-trays. What we have added is that every Calorimeter Trigger module includes a CANbus node (shown as N in the figure), connected to on-board sensors measuring temperatures and voltages appropriate for that type of module. The nodes are standardised on all L1Calo modules, and are Fujitsu MB90F594 microcontrollers that include a 10-bit ADC. They are connected via the backplane to another microcontroller B, on the TCM, acting as a CANbus bridge. The second port of the bridge is connected via an external CANbus to a Local Control Workstation (LCS)

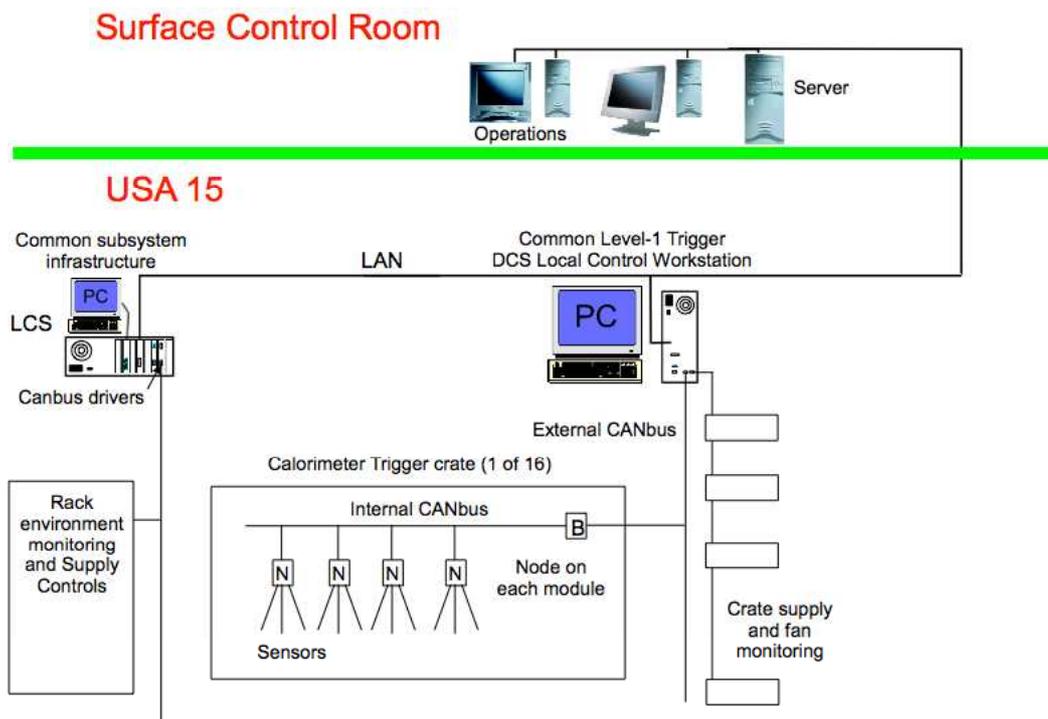


Figure 29. ATLAS DCS architecture and L1Calo monitoring.

computer, where information is gathered from all L1Calo crates using an implementation of the CANOpen protocol. The LCS acts as the interface to the DCS system for the whole of the ATLAS experiment, sending error messages to the main control room when problems occur.

A watchdog timer provides a hardware confirmation that periodic checking of data values within the CANbus system is being done.

7.2 The Timing Control Module (TCM)

All L1Calo crates contain a Timing Control Module at their right-hand end. This carries out two main functions: it distributes the LHC 40.08 MHz clock signals and other information from the ATLAS TTC system, and it acts as a gateway between the internal CANbus, which monitors voltages and temperatures on all the trigger modules in the crate, and an external CANbus connected to the ATLAS DCS. The TCM also provides a visual diagnostic display of the VME status on its front panel.

The TCM is a relatively simple 9U module. However, there are two versions: a standard VME version for use in the PPM and ROD crates, and another version for use in the CP and JEP crates with their custom backplane and reduced VMEbus. The two versions are functionally very similar, as shown in figure 30.

7.2.1 TTC signal distribution

The timing information from the TTC system includes the 40.08 MHz LHC clock, the LHC Orbit signal, the L1A signal, the Trigger Type, and the Event Counter Reset signal. The signals are received at L1Calo in a 6U VME crate by a CTP Interface module. After passing through a Local Trigger Processor (LTP) module, the information is encoded by a TTCvi module into a single serial signal, and then distributed optically by a TTCex module. The final stage of distribution is provided by the TCM, which receives the TTC optical signal, converts it to

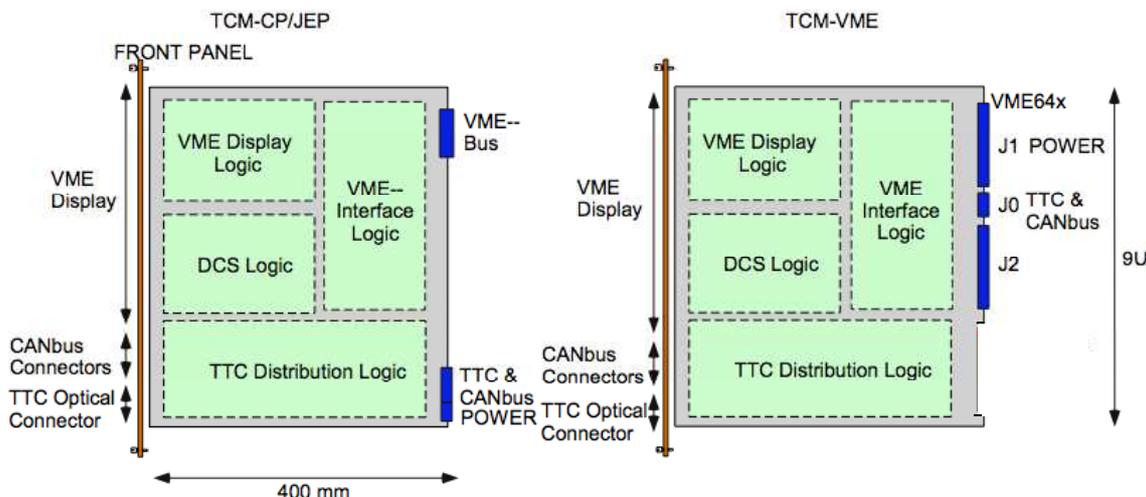


Figure 30. Schematic layout of the TCM: CP/JEP version (left) and VME64 version (right).

differential electrical format (PECL), and sends it to individual trigger modules over the Processor Backplane in CP/JEP crates, or via the auxiliary backplane (section 7.2.3) in VME64x crates. On receipt in the trigger modules, the TTC signal goes to a TTCdec decoder daughter card, a standard item on all L1Calo modules, where a TTCrx chip recovers the clock and other signals for use on the module.

7.2.2 CANbus node

The system for monitoring voltages and temperatures on all the trigger modules was described in section 7.1. Within each crate, the internal CANbus is managed by the CANbus controller on the TCM. The Fujitsu MB90F594 microcontroller has two CANbus ports, and communicates via its second port with an external CANbus network coming from the DCS Local Control Workstation. This is connected via the TCM's front panel.

The microcontroller on the TCM reads measured values from each module in the crate at roughly 16-second intervals, and creates a table of values which is sent to the Local Control Station over the external CANbus. If a voltage or temperature on a module is outside of programmable limits an error message is sent to the TCM, which passes it to central DCS. Error messages are asynchronous and have a higher priority than normal data packets — they are received by the central DCS within roughly 0.1 seconds.

7.2.3 TCM variants, the auxiliary backplane and geographical addressing

Two versions of the TCM are necessary, due to the incompatible connector layouts in standard VME crates (PreProcessor and ROD) and custom-backplane crates (CP and JEP). Figure 30 shows the layouts of these two versions. The CP/JEP version obtains power from the bottom of the board, VME signals from the VME-- bus at the top, and sends TTC and CANbus signals through the backplane connector at the bottom. In the VME version, power, TTC and CANbus signals are connected to the upper part of the board. For this version the J3 area is not connected because incompatible connectors are used in the PreProcessor and ROD crates, and the module must not interfere with the connector in either type of crate.

There are no suitable traces in the VME64x and VME64xP backplane to carry the crate geographical address, the TTC and the CANbus signals. An auxiliary backplane has therefore been added at the rear of PreProcessor and ROD crates, connecting to the rear of the J0

connectors on the VME64 backplane. This distributes the necessary signals from the TCM's J0 connector to slot numbers 2 to 20. The geographical addresses are bussed, while the TTC signals are point-to-point.

VME64x and VME64xP provide a set of five geographical addressing pins, which a module may use to determine its slot number within a crate. In ROD and PreProcessor crates, L1Calo defines four further custom pins on the J0 connector to provide crate-number addressing. The crate number is distributed over the auxiliary backplane to each module position. In CP/JEP crates, three pins are provided by the custom backplane and the crate number is determined by setting a switch on the backplane.

7.3 Crate control and infrastructure

There are three types of 9U crates in the L1Calo system. The PreProcessor uses VME64xP crates with air-cooled power supplies mounted above the crate. The RODs use VME64x crates with remote water-cooled power supplies mounted at the rear of the rack. The CP and JEP use the same basic crates, but with a custom 9U-high backplane (see section 5.4). These crates also have remote water-cooled power supplies at the rear of the rack. The custom backplane carries a minimal implementation of the VMEbus protocol on high-density connectors.

In all crates, the control, configuration, and low-level monitoring carried out via the VMEbus is managed by a Concurrent CPU. This is a 6U VME module in slot 1, and operates under Linux. In the PreProcessor and ROD crates the first two slots in the crate are provided with a card cage to support the module. In the CP and JEP crates an essentially passive, custom 9U adapter module, called the VME Mount Module (VMM), is used to hold the CPU at the front of the crate, and to interface it mechanically and electrically to the reduced VMEbus implementation.



Figure 31. Photograph of the rear of the JEP, showing LVDS cable links, cable supports, and other infrastructure.

In many cases it is necessary to mount various types of rear-transition modules behind the backplane, and the mechanics for these must provide support and some protection. These modules include: small G-Link driver cards to serialise the PPM readout data and convert them for optical transmission, passive cards to provide cable connectors between crate-level and system-level CMMs, and S-Link modules for the RODs to transmit the readout and RoI data.

All PreProcessor and Processor crates have a very large number of LVDS serial-transmission links connected to their backplanes. These relatively fragile cables and connectors occupy a large volume in the crate, coexisting with a number of optical fibres, rear-transition modules, crate power-supply cables and water-cooling pipes. Their weight is supported by custom metal support structures. The general situation is illustrated in figure 31.

8. Status

At the time of writing (late 2007–early 2008) the entire calorimeter trigger system has been installed in ATLAS. The trigger itself is being commissioned, and systematic testing and calibration of input signals from the calorimeters is underway using both cosmic-ray triggers and the calorimeter calibration systems. L1Calo has participated in a series of increasingly demanding integration runs with the rest of the detector; these used cosmic-ray triggers from the Tile Calorimeter, and more recently cosmic-ray triggers from L1Calo itself. Procedures for operating L1Calo and monitoring the trigger performance are evolving steadily. We expect the trigger to be operating well by the time the LHC starts up.

Acknowledgements

The Level-1 Calorimeter Trigger has complex interfaces with other ATLAS systems, and we have benefited from the help and advice of many people, far too many to list. We are very grateful in particular to colleagues working on the Liquid Argon Calorimeters, the Tile Calorimeter, the Central Trigger Processor, the Level-2 Trigger, the Data Acquisition, the Detector Control System, and ATLAS Technical Coordination. We are especially indebted to the Pittsburgh group for close cooperation on their analogue Receivers. Special thanks go to Nick Ellis, one of the originators of this project, who has given us much support and advice over many years.

Installation of the trigger has proved to be a very demanding and complex task that included cable and fibre installation, and the solution of a number of electrical and mechanical problems. 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 — physicists, engineers, technicians and students — 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 Georgios Anagnostou, Darren Ballard, Dan Beckett, Richard Carney, Adam Connors, Chris Day, Magnus Engström, Dan Husmann, Kithsiri Jayananda, Michael Keller, Richard Matson, Adrian Mirea, Tamsin Moye, Volker Schatz, Bernd Stelzer, Oliver Stelzer-Chilton, Peter Stock, Bill Stokes, Scott Talbot, Dimitrios Typaldos and Gernot Weber.

The Level-1 Calorimeter Trigger has been supported by BMBF, 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*, to be published in JINST (2008).
- [2] E. Eisenhandler, ATLAS Level-1 Calorimeter Trigger Algorithms, ATLAS note ATL-DAQ-2004-011 (2004).
- [3] N.J. Buchanan et al., *ATLAS Liquid Argon Calorimeter Front End Electronics*, submitted to JINST.

- [4] A.S. Cerqueira, J.M. Seixas and L.P. 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.
- [5] ATLAS Level-1 Trigger Group, *ATLAS First-Level Trigger Technical Design Report*, CERN/LHCC-98-14 (1998).
- [6] S-Link specification at <http://www.cern.ch/hsi/s-link/>.