

The ATLAS Level-1 Calorimeter Trigger Architecture

J. Garvey, S. Hillier, G. Mahout, T.H. Moyer, R.J. Staley, P.M. Watkins, A. Watson
School of Physics and Astronomy, University of Birmingham

R. Achenbach, P. Hanke, E-E. Kluge, K. Meier, P. Meshkov, O. Nix, K. Penno, K. Schmitt
Kirchoff-Institut für Physik, University of Heidelberg

C. Ay, B. Bauss, A. Dahloff, K. Jakobs, K. Mahboubi, U. Schäfer, T. Trefzger
Institut für Physik, Universität Mainz

E. Eisenhandler, M. Landon, E. Moyse, J. Thomas
Physics Department, Queen Mary, University of London

P. Apostoglou, B. M. Barnett, I. P. Brawn, A. O. Davis, J. Edwards, C. N. P. Gee,
 A. R. Gillman, V. J. O. Perera, W. Qian,
Rutherford Appleton Laboratory

C. Bohm, S. Hellman, A. Hidvégi, S. Silverstein
Department of Physics, University of Stockholm

Abstract—The architecture of the ATLAS Level-1 Calorimeter Trigger system (L1Calo) is presented. Common approaches have been adopted for data distribution, result merging, readout, and slow control across the three different subsystems. A significant amount of common hardware is utilized, yielding substantial savings in cost, spares, and development effort. A custom, high-density backplane has been developed with data paths suitable for both the em/τ cluster processor (CP) and jet/energy-summation processor (JEP) subsystems. Common modules also provide interfaces to VME, CANbus and the LHC Timing, Trigger and Control system (TTC). A common data merger module (CMM) uses FPGAs with multiple configurations for summing electron/photon and tau/hadron cluster multiplicities, jet multiplicities, or total and missing transverse energy. The CMM performs both crate- and system-level merging. A common, FPGA-based readout driver (ROD) is used by all of the subsystems to send input, intermediate and output data to the data acquisition system (DAQ), and region-of-interest (RoI) data to the level-2 triggers. Extensive use of FPGAs throughout the system makes the trigger flexible and upgradable, and several architectural choices have been made to reduce the number of inter-crate links and make the hardware more robust.

I. INTRODUCTION

DATA taking at the ATLAS experiment requires a fast, three-level trigger system to select physics processes of interest. The Level-1 trigger is designed to reduce the event

rate from the 40 MHz LHC bunch-crossing rate to below 75 kHz, based on information from the Level-1 Calorimeter and Muon trigger systems [1].

The Level-1 Calorimeter Trigger (L1Calo) is a pipelined processor designed to receive and analyze some 7200 summed analog trigger towers (granularity 0.1×0.1 in η and ϕ) from the electromagnetic and hadronic calorimeters, and send the results for each LHC bunch crossing to the Central Trigger Processor (CTP) less than 2 ns after the event took place. There, the calorimeter and muon trigger results are used together to produce a Level-1 Accept (L1A) decision.

The L1Calo system (Fig. 1) comprises three main subsystems. The PreProcessor (PPr) digitizes and processes the analog signals, assigning them transverse energy values and identifying them with specific bunch crossings. The digital data are then transmitted to the Cluster Processor (CP) and Jet/Energy-sum Processor (JEP) subsystems. The CP subsystem receives the digitized trigger tower sums and identifies isolated electron/photon and hadron/tau candidates. The JEP receives data in 0.2×0.2 sums of trigger towers (jet elements), and uses these sums to identify jets and produce global sums of transverse and missing transverse energy.

The CP and JEP send feature multiplicities and energy threshold information to the CTP for every bunch crossing. Upon a L1A decision from the CTP, the subsystems send stored input data as well as intermediate and output data to the data acquisition system (DAQ) for inclusion in the data stream. The types and positions of jet and cluster candidates are also collected and sent to the Region of Interest (RoI) builder for use by the Level-2 trigger.

Over an extended period of research and development, the L1Calo architecture has undergone numerous revisions and

Manuscript received May 22, 2003.

Corresponding Author: S. Silverstein, Stockholm University Department of Physics, AlbaNova University Centre, 106 91 Stockholm, Sweden (telephone: (+46 8) 5537 8693, e-mail: silver@physto.se).

refinements. The final system design [2] has been chosen to be relatively “compact”, with a minimal number of crates and cable links, yielding benefits in overall system cost and latency. A number of common, multiple-use hardware modules have been developed for general use in the system, helping reduce hardware costs, numbers of spares, and design effort. And numerous refinements in the design help to make the trigger system more robust, reliable, and simpler to maintain.

This document gives an architectural overview of the L1Calo system, including the PPr, CP, JEP and read-out subsystems, with an emphasis on common hardware solutions and design features that make the system more compact, robust and reliable.

I. THE PREPROCESSOR

The PreProcessor receives and processes roughly 7200 analog signals from the ATLAS electromagnetic and hadronic calorimeters. The differential signals from the detector are brought to analog receiver stations where they are calibrated to an E_T scale and distributed to the front panels of the PreProcessor crates.

The PreProcessor is an eight crate system, of which four crates process electromagnetic trigger towers and four process hadronic towers. Each crate contains 16 PreProcessor Modules (PPMs), that can each receive and process 64 analog inputs.

The PPM is a 9U module residing in a standard VME64 crate. Differential signals entering the PPM first enter one of four analog input stage daughter modules. There they are converted to single-ended signals with appropriate bias for digitization.

The actual signal processing is performed by 16 Multi-Chip Modules (PPrMCM), each of which processes four trigger towers. Four 10-bit FADCs digitize the signals with a sampling frequency of 40 MHz. Fine adjustment of the digitisation strobes is performed by a special ASIC (PHOS4) that provides programmable delays in steps of 1 ns across the 25 ns LHC clock period. The digitized values are then sent to a custom PreProcessor ASIC (PPrASIC).

The PPrASIC synchronizes the four inputs for different path lengths and time-of-flight effects, assigns the signals to the correct bunch crossing, and uses a lookup table to produce calibrated 8-bit transverse energy values. It also sums the four values into 2×2 jet elements, performs “bunch-crossing multiplexing” (see below) and saves the input data in pipelines for eventual readout. Finally, two 10-bit LVDS serializers operating at 400 Mbit/s transmit the processed trigger tower data to the CP, while a third serializer sends the summed 9-bit jet elements to the JEP.

A. Bunch Crossing Identification (BCID)

The Preprocessor ASIC is capable of identifying a signal’s bunch crossing using three different methods that provide ample redundancy for consistency checks during system setup.

For normal, non-saturated signals (below about 250 GeV), a pipelined FIR filter processes five consecutive FADC-samples, by multiplying the samples with pre-defined coefficients and summing the resulting values. A subsequent “peak-finder” attributes the maximum value of this sum to

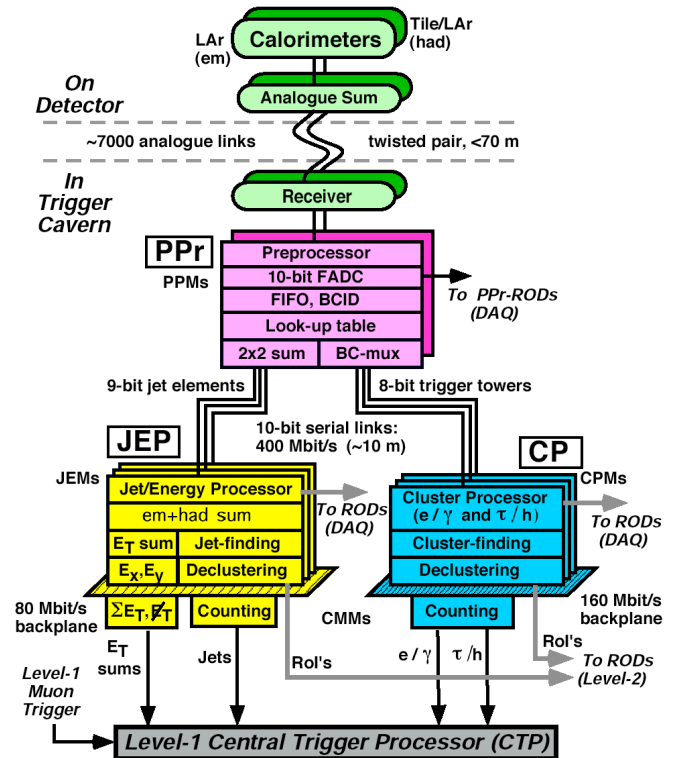


Fig. 1. Architectural overview of the ATLAS Level-1 Calorimeter trigger, showing the real-time data path from the calorimeters to the Central Trigger Processor, as well as the readout to the DAQ and Level-2 systems.

the corresponding bunch crossing. The working range of the method spans from small signals (few GeV equivalent) up to the near saturation level of around 250 GeV.

For saturated signals, two consecutive samples are compared to a “low” and a “high” threshold, making use of the finite “peaking-time” (50 ns) of an analog input signal. Thus, detection of a “leading edge” allows attribution of the “virtual peak” to a specific bunch crossing. This method is valid from around 200 GeV through the maximum energy range of the calorimeter.

A third method, provided for consistency checking, uses comparators with programmable thresholds on the analog input daughter modules to present a “rising edge” signal to the MCM. Given the known peaking time, BCID can be performed using an appropriate programmed delay in the PPrASIC. The validity of this method begins well above the comparator threshold, and extends up to the full energy range. There is thus a large overlap with the previous two methods, allowing consistency checks between the methods to be performed during set-up.

The FIR filter output is presented to a lookup table to extract a calibrated E_T value for the tower. If the BCID criteria are met, this value is sent to the CP and JEP outputs. In the case of saturation, the tower is assigned the maximum 8-bit value (255 GeV). A 2×2 jet element with a sum that overflows, or which contains a saturated tower, is also assigned the maximum 9-bit value (512 GeV). All data words are accompanied by an odd-parity bit for error detection.

A. Bunch-Crossing Multiplexing (BC-mux)

The PPrMCM sends its outputs to a daughter card on the PPM subsystem via 10-bit LVDS serial links for electrical

fanout to the CP and JEP. The four trigger towers serviced by each MCM are sent to the CP, and one 9-bit, 2×2 sum of these towers (jet element) is sent to the JEP. To economise on the number of links needed, it is important to note that the BCID algorithm, being essentially a “peak finding” scheme, will always follow an occupied BC with an empty (zero) one. This allows two trigger towers to share a single serial link.

Towers are paired in the PPrASIC output stage. When a tower has a non-zero value, its value is sent to the link, along with a flag bit that indicates which of the two towers is being transmitted first. In the next LHC clock cycle, an E_T value (or zero) for the other tower is sent to the link, again with a flag bit. For this clock cycle, the flag bit is used to indicate whether the second tower’s value belongs to the same BC as that of the first tower, or to the following BC.

By using BC-mux, data for the CP subsystem can be achieved with only two links per MCM instead of four. It is not possible to use this scheme on the summed 4-tower jet element, since there is no assurance that such a sum will be followed by a zero in the next BC.

A. Output Signal Fanout and Pre-compensation

A daughter card on the PPM distributes data to the CP and JEP subsystems via LVDS serial links operating at 400 Mbit/s over at least 10 m of shielded “parallel pair” cables. To provide the trigger algorithms with the needed overlapping data between detector quadrants, some links must be fanned out 1×2 . Furthermore, measures must be taken to improve signal driving capabilities, since observed signal attenuation and distortion from such long cables can compromise data integrity.

Studies have shown that an FPGA (e.g. Xilinx Virtex XCV50E) [3] can provide reliable signal distribution and fanout at 400 Mbit/s. Combined with an RC “pre-compensation” network (Fig. 2) to compensate cable signal losses, bit-error rates of less than 10^{-14} have been achieved over 15 meters of cable.




Fig. 2. An RC pre-compensation network for driving high speed serial LVDS signals over long cables.

I. THE CLUSTER AND JET/ENERGY-SUM PROCESSORS

The CP and JEP subsystems share many architectural features and common hardware. The jet algorithm in the JEP and the em/hadron cluster algorithms in the CP both perform feature searches in overlapping windows. Therefore, a large amount of data duplication between processor units is required. Both subsystems divide the calorimeters into four quadrants, with processor modules in each quadrant sharing input data with nearest neighbors over short point-to-point links. Overlapping data from neighboring quadrants are

provided by duplicated serial links from the PreProcessor. This architecture minimizes the number of cable links from the PreProcessor, thus allowing the system to be built more compactly. Because backplane fanout is only between nearest neighbors (~ 2 cm), signal timing is also simplified.

The CP is a four-crate system, with 14 Cluster Processor Modules (CPMs) in each crate covering one calorimeter quadrant. The JEP is contained in two crates, with 8 Jet/Energy Modules (JEMs) from each of two opposing quadrants (16 total) in each crate. Results from the processor modules are brought to two Common Merger Modules (CMMs) in each crate, which sum the data to produce crate-level results. The CMMs also perform the system-level summation of data from the different crates, and report those results to the CTP.

A. The Cluster Processor Module (CPM)

The e/γ cluster algorithm identifies 2×2 clusters of trigger towers with at least one two-tower sum (1×2 or 2×1) of nearest neighboring electromagnetic towers exceeding a predefined threshold. Isolation thresholds are set for the 4×4 surrounding ring in the EM calorimeter, as well as the 2×2 hadronic tower sum behind the cluster, and the 12-tower hadronic ring around it. The hadron/ γ algorithm is similar: nearest-neighbor sums of EM plus hadronic trigger towers are compared with an energy threshold, and separate isolation thresholds are set for the surrounding 12-tower rings in the EM and hadron calorimeters. Multiple counting of clusters is avoided by requiring the central 2×2 region of the cluster to be a local maximum. The CPM identifies and report clusters satisfying 8–16 e/γ and 0–8 γ /hadron criteria, for a total of 16 independent cluster types.

Each CPM receives and deserialises data from 80 LVDS links, corresponding to 4×20 trigger towers in \square and \square , both EM and hadronic. Twenty serialiser (“SRL”) FPGAs (XCV100E) receive the data and serialize them to 160 MHz. The data are then shared between neighboring modules via the backplane, and finally fanned out to eight “CP” FPGAs, which perform the cluster algorithms. The SRL FPGAs also store input in pipelines for eventual readout to DAQ upon a L1A.

The 8 “CP” FPGAs are implemented in large devices (XCV1000E), each of which services 8 overlapping 4×4 windows. Pipelines implemented in each CP FPGA save output data for readout to the DAQ system, and cluster types and coordinates for readout as RoIs to Level-2.

Two “hit multiplicity” FPGAs collect and sum the cluster multiplicities from the CP FPGAs for reporting to the crate-level merging of Level-1 results. The multiplicities are reported on two 25-bit data streams, each containing 3-bit multiplicities for 8 of the 16 cluster types, plus one bit of odd parity. If more than 7 instances of a cluster type are identified, the multiplicity is reported as 7.

Two additional “Read-Out Controller” (ROC) FPGAs collect input data from the SRL FPGAs, RoI data from the CP FPGAs and output data from the hit multiplicity FPGAs upon a L1A, and transmit them to DAQ and Level-2.

B. The Jet/Energy Module (JEM)

The jet algorithm identifies transverse energy sums within overlapping windows consisting of 2×2 , 3×3 or 4×4 jet elements (corresponding to sizes 0.4, 0.6 or 0.8 in \square and \square), and comparing them against predefined energy thresholds.

Multiple counting of jet candidates is avoided by requiring the window to surround a 2×2 cluster whose sum is a local maximum. The location of this 2×2 cluster also defines the coordinates of the jet RoI. All three jet window sizes are calculated by default, and eight independent combinations of energy threshold and window size are available for trigger menus.

The energy-summation algorithm produces unsigned sums of E_T , E_X and E_Y and uses the system-level sums of these to report on total and vector E_T thresholds to the CTP.

Each JEM receives and deserializes data from 88 LVDS links, corresponding to 4×11 jet elements in \square and \square , both EM and hadronic. The 40 MHz data are then received by four “input” FPGAs (XC2V1500), which sum the EM and hadronic parts of each jet element to 10-bit values and multiplex these sums to 80 MHz for fan-out. Both backplane and on-board transmission of data to the main processor units are carried out at 80 MHz, with the five least significant bits in the sum transmitted first. Pipelines in each input FPGA save input data for readout to the DAQ system upon a L1A.

The jet and energy-summation algorithms are implemented in one large “main processor” FPGA (XC2V2000) per JEM. The main processor is also responsible for reporting results to the crate-level mergers, as well as pipelining of DAQ and RoI information for readout. The jet output of each JEM is a 25-bit data stream consisting of 8 3-bit jet multiplicities and one odd-parity bit. The energy output is also a 25-bit data stream containing the unsigned values of E_T , E_X and E_Y , each compressed from 12 bits to an 8-bit “quad-linear” scale (6-bit mantissa plus two multiplier bits). The energy output is also accompanied by an odd parity bit.

A single ROC FPGA collects input data from the input FPGAs and output and ROI data from the main processor FPGA for readout to DAQ and Level-2.

C. The Common Merger Module (CMM)

Two modules in each CP and JEP crate carry out crate-level merging of results received from the crate’s processor modules. In the CP crates, each merger module is responsible for calculating 3-bit cluster multiplicities for 8 of the 16 e/\square and \square hadron cluster definitions. In the JEP crates, one merger module produces 3-bit multiplicities for the 8 jet definitions, while the other produces sums of E_T , E_X and E_Y . A common merger module (CMM) has been developed for all merging functions, using a large FPGA with multiple firmware configurations to do crate-level merging on up to 400 bits of data per BC from the crate’s processor modules [4].

In addition to crate-level summation, the CMM also performs system-level merging. Parallel LVDS cable links between the subsystem crates bring the crate-level results to one CMM of each type, which is designated the system-merger. A second large FPGA on the CMM carries out the system level merging and reports final results to the CTP.

The FPGAs on the CMM also store input and output data into pipelines for readout to DAQ upon a L1A.

D. The Processor Backplane (PB)

The CP and JEP subsystems use a common, custom processor backplane (PB). Each position in the PB is meant for a single kind of module. In addition to 16 CPM/JEM

positions flanked by two CMM positions, there are two positions for inserting a VME CPU (mounted in a 9U adapter module), and one position for the Timing and Control Module (TCM), which serves as an interface for the TTC and slow control systems.

Physically, the PB is a monolithic backplane of 9U height, and hosts up to 21 single-width modules. It is populated almost entirely with 2mm HM connectors, with 1148 signal and ground pins in each JEM/CPM and CMM position. Three high-current DIN pins (3.3 V, 5 V, GND) provide DC power to each module near the bottom of the crate. The PB has 8 signal layers sandwiched between 10 ground planes, allowing excellent impedance control and low cross-talk [5]. Two signal layers provide point-to-point links between neighboring processor modules for input data fan-in/out. A CANbus on one of these layers provides an interface to the detector control system (DCS). Another four layers provide diagonal connections from outputs at the top and bottom of each CPM or JEM to the two merger modules at the right and left of the processor modules. The remaining two layers distribute TTC signals from the TCM. To conserve pins, the VME bus is reduced to only the 43 signal lines needed to provide A24D16 slave access to the modules.

The LVDS serial link and merger interconnect cables are connected to the rear side of the PB and passed through it to the modules in front. This allows a system with few cables on the front panels of the modules, allowing fewer re-cabling errors and cable damage over the lifetime of the experiment. The crate also provides extended geographic addressing. Upon power-up, address pins tell each module its location in the crate, as well as the crate’s location in the system. The modules use this information to set unique VME, CAN and TTC addresses. Modules with multiple FPGA configurations can also use this information to automatically load the appropriate configuration.

II. THE READOUT DRIVER (ROD)

The L1Calo has two separate readout systems. Input, output and some intermediate data from each module are read out to the DAQ system, and the CP and JEP subsystems report feature types and coordinates as RoI data to the Level-2 trigger. The readout system has been designed to report 1 BC of RoI data and up to 5 BC of DAQ data per event at a L1A rate of up to 75 kHz. A common approach has been adopted over all L1Calo subsystems for DAQ and RoI readout.

For each module to be read out, readout FIFOs on each processor FPGA or ASIC are read out as 40 MHz serial streams to a Read Out Controller (ROC) for timing alignment. The ROC passes the serial streams in parallel to the inputs of a 20-bit G-link transmitter (Agilent HDMP 1022), [7] which transmits them at 800 Mbit/s to a Readout Driver (ROD).

A common ROD module is used by both the DAQ and RoI readout subsystems to gather and report data from the PPr, CP and JEP subsystems, using a different firmware configuration in a large FPGA for each readout task. The ROD is a 9U module residing in a standard VME64 crate. It has 18 G-link receivers (Agilent HDMP 1024) which pass the parallel outputs to the FPGA for data compression, zero suppression, and some data monitoring. The ROD also contains four S-link [8] transmitters for passing compressed event data to the DAQ and RoI readout buffers. The number

of G-link inputs and S-link outputs used, and the mapping between them, depends on the type and source of data being read out.

III. CONCLUSIONS

We have described a trigger system with a compact architecture that minimizes cost and latency without sacrificing performance or flexibility. Several common hardware solutions have been adopted that also reduce cost, design effort, and the numbers of spare boards. And numerous refinements, such as LVDS precompensation, eliminating most cabling to the front of the CP and JEP subsystems, and self-configuration of modules using extended geographic addressing, make the trigger more robust, reliable, and easy to maintain.

IV. REFERENCES

- [1] ATLAS Level-1 Trigger Group, *ATLAS First-Level Trigger Technical Design Report*, ATLAS TDR-12, CERN/LHCC/98-14, CERN, Geneva, 1998.
- [2] Module information available on ATLAS Level-1 Calorimeter Trigger home page: <http://hepwww.pp.rl.ac.uk/Atlas-L1>
- [3] Xilinx home page: <http://www.xilinx.com/>
- [4] I Brawn, *One Size Fits All: Multiple Uses of Common Modules in the ATLAS Level-1 Calorimeter Trigger*, 7th Workshop on Electronics for LHC experiments, CERN/LHCC/2001-034, 22 October 2001, p. 253
- [5] M. K. Jayananda, *ATLAS Level-1 Calorimeter Trigger – Simulation of the backplane for Common Merger Module*, ATLAS ATL-DAQ-2000-004, CERN, Geneva, 2000.
- [6] Timing, trigger and control system (TTC) information may be found at: <http://www.cern.ch/TTC/intro.html>
- [7] Agilent G-link information available at: <http://www.semiconductor.agilent.com/>
- [8] CERN S-link Specification: <http://www.cern.ch/HIS/s-link>