

Prototype Cluster Processor Module for the ATLAS Level-1 Calorimeter Trigger.

G. Anagnostou, J. Garvey, S. Hillier, G. Mahout*, R.J. Staley, P.J. Watkins, A. Watson

School of Physics and Astronomy, University of Birmingham, Birmingham B15 2TT, UK

R. Achenbach, P. Hanke, W. Hinderer, D. Kaiser, E-E. Kluge, K. Meier,
O. Nix, K. Penno, K. Schmitt

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

B. Bauss, A. Dahlhoff, K. Jakobs, K. Mahboubi, U. Schäfer, J. Thomas, T. Trefzger

Institut für Physik, Universität Mainz, D-55099 Mainz, Germany

E. Eisenhandler, M. Landon, D. Mills, E. Moyse

Physics Department, Queen Mary, University of London, London E1 4NS, UK

P. Apostologlou, B.M. Barnett, I.P. Brawn, A.O. Davis, J. Edwards, C. N. P. Gee,
A.R. Gillman, R. Hatley, V.J.O. Perera

Rutherford Appleton Laboratory, Chilton, Oxon OX11 0QX, UK

C. Bohm, S. Hellman, S. Silverstein

Fysikum, University of Stockholm, SE-106 Stockholm, Sweden

Abstract

The Level-1 Calorimeter Trigger consists of a Preprocessor, a Cluster Processor (CP), and a Jet/Energy-sum Processor (JEP). The CP and JEP receive digitised trigger-tower data from the Preprocessor and produce trigger multiplicity and region-of-interest (RoI) information. The CP Modules (CPM) are designed to find isolated electron/photon and hadron/tau clusters in overlapping windows of trigger towers. Each pipelined CPM processes a total of 280 trigger towers of 8bit length at a clock speed of 40 MHz. This huge I/O rate is achieved by serialising and multiplexing the input data. Large FPGA devices have been used to retrieve data and perform the cluster-finding algorithm. A full-specification prototype module has been built and tested, and first results will be presented.

I. INTRODUCTION

At the full LHC design luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, there will be approximately 23 proton-proton interactions per bunch crossing. The ATLAS Level-1 trigger [1] is to reduce the 1 GHz interaction rate to a trigger rate of 75 kHz for input to

the Level-2 trigger. The reduction is performed by processing reduced-granularity data from the calorimeters and muon spectrometer. Potentially interesting events are selected by identifying electron/photon candidates, jets, single-hadron/tau candidates, missing transverse energy, and total transverse energy. Muon candidates are identified in a separate trigger. The Level-1 decision will be based upon a combination of these signals, and is made by the Central Trigger Processor (CTP). The total Level-1 latency is 2 μs during which data from the entire ATLAS detector will be stored in analogue or digital pipeline memories.

This paper describes the Cluster Processor Module of the Level-1 calorimeter trigger system, designed to identify transverse energy clusters associated with electrons/photons and isolated hadrons/taus. After a more detailed description of the Level-1 calorimeter trigger system, algorithms used to identify electromagnetic and hadronic clusters will be presented. Its implementation in hardware requires a high bandwidth of input data, and design choice and the use of technology such as FPGAs will be discussed. Finally, an existing full specification CPM will be shown and early results from stand-alone tests presented.

* Corresponding author: gm@hep.ph.bham.ac.uk




Figure 1: The ATLAS Level-1 Calorimeter Trigger System

II. THE ATLAS LEVEL-1 CALORIMETER TRIGGER SYSTEM

As shown in Figure 1, the ATLAS Level-1 Calorimeter Trigger system consists of three subsystems, namely the Preprocessor, electron/photon and tau/hadron Cluster Processor (CP), and Jet/Energy-sum Processor (JEP). The CP and JEP will receive digitised calorimeter trigger-tower (TT) data from the Preprocessor, and send trigger multiplicity information to the Central Trigger Processor via Common Merger Modules (CMM). Using Readout Driver (ROD) modules, the CP and JEP will also provide region-of-interest (RoI) information for the Level-2 trigger, and intermediate and final results to the data acquisition (DAQ) system for monitoring and diagnostic purposes.

The calorimeter trigger covers the region $|\eta| < 5$ and $\phi=0$ to 2π . On the detector, cells are combined to form trigger towers, with a reduced granularity of $(\Delta\eta \times \Delta\phi) = 0.1 \times 0.1$ over the region $|\eta| < 2.5$ and a variable granularity elsewhere. Analogue pulses enter the Preprocessor, where they are digitised to 10-bit precision at a frequency of 40 MHz. After Bunch-Crossing Identification (BCID), the precise value of transverse energy for each trigger tower is produced in a look-up table. The transverse energy is an 8-bit word, giving a transverse energy scale linear up to 255 GeV. The total number of trigger towers is 7200 of which 6400 are processed by the Cluster Processor, corresponding to an area of $|\eta| < 2.5$. Each Cluster Processor Module sends the number of e/γ and τ /hadron clusters it has found, up to a maximum of 7 per set of thresholds, to two merger modules located in the same crate. All of the four crates in the CP are merged and sent to the CTP to build the Level-1 decision. Similar architecture [2] is used for the Jet Processor where jets are identified and their multiplicities are also sent via CMM to the CTP, as well as the total and

missing transverse energy. On receipt of a Level-1 request, both subsystems, Jet and Cluster Processor, send RoI information, Trigger Tower Data and multiplicity information to the ROD.

III. CLUSTER PROCESSOR MODULE REQUIREMENTS

The CPM receives digitised data from the Preprocessor Module and sends in real time cluster information to the CMM. It also provides information to the readout chain. Its requirements are as follows:

- Identify possible isolated electrons, photons and semi-hadronic τ decays
- Calculate multiplicities of e/γ candidates and τ candidates for different threshold conditions on E_T
- Transmit these multiplicities as input to the Level-1 trigger decision
- Transmit Trigger Tower data, multiplicities and RoI coordinates to ReadOut Driver Modules.

IV. THE CLUSTER PROCESSOR ALGORITHM

The calorimeter trigger is designed to find electromagnetic clusters and isolated hadrons. A large fraction of the isolated hadron triggers will probably come from taus.

A. Trigger Windows

Overlapping windows of $(\Delta\eta \times \Delta\phi) = 0.4 \times 0.4$ in both electromagnetic and hadronic calorimeters are examined (Figure 2:a). These windows scan in steps of 0.1 in η and ϕ over the trigger phase space of the calorimeter. For each of these windows, we define a set of trigger clusters:




Figure 2: Cluster Trigger definitions

- Electromagnetic clusters: four possible sums of two adjacent central electromagnetic trigger towers (b)
- Electromagnetic isolation sum: sum of the 12 electromagnetic calorimeter trigger towers surrounding the central 2x2 towers (c)
- Hadronic isolation sum: sum of the 12 hadronic calorimeter trigger towers surrounding the central 2x2 towers (c)
- For e/γ trigger, an additional central hadronic isolation sum: sum of the 2x2 hadronic calorimeter central trigger towers
- For tau/hadron trigger, clusters are four possible sums of two adjacent central e.m. towers with the four central hadronic towers
- Central Region of Interest: sums of the 2x2 central electromagnetic and hadronic trigger towers (d)

B. Trigger Algorithm

The 4x4 window is declared an e.m. candidate object if:

- The Central RoI is a local maximum
- The electromagnetic isolation sum, the hadronic isolation sum and the central hadronic isolation sum are all below their specified thresholds
- At least one of the 4 electromagnetic clusters is above a trigger threshold.

It is declared an isolated hadronic trigger object if:

- The Central RoI is a local maximum
- The electromagnetic isolation sum and the hadronic isolation sum are both below their specified thresholds
- At least one of the 4 hadronic clusters is above a trigger threshold.

A total of 16 sets of threshold values, consisting of the trigger threshold and the relevant isolation thresholds, can be flagged. Of these, eight are allocated to electromagnetic clusters, and eight can be allocated to electromagnetic clusters or isolated hadron triggers.

V. HARDWARE IMPLEMENTATION OF THE CLUSTER PROCESSOR ALGORITHM

Due to the high density of signals to process to perform the cluster finding algorithm, early design of the CPM was done with an ASIC. Today, Programmable Logical Devices,

such as FPGAs [3], are getting bigger, faster and the increase of I/O makes them very appealing for our purpose. They also provide the option to redesign the algorithm at later stage during the running of LHC, FPGAs have now been chosen to perform the CP algorithm. We are still limited by the number of inputs per chip, and per board. The choice has been to design a chip capable of processing a total of 8 4x4 windows. A CPM is populated with 8 of them, and data flow needs to be reduced by:

- Multiplexing and serialising data before reaching the CPM
- Serialising data at 160 MHz at the input of the CP chip
- Sharing TT data across the board but also between adjacent modules

This latter requirement has involved building a custom backplane, where 160 MHz signals are exchanged between adjacent boards.

A. Cluster Processor Data Flow

The Cluster Processor has to process a total of 6400 TT of 8 bits each. This amount of data is reduced by a factor 2 by multiplexing data of 2 TT adjacent in phi. Due to the nature of the calorimeter pulses and the BCID process, a TT containing energy is immediately followed on the next clock tick by a null value. A lot of null data would have to be transferred and so a multiplexing scheme, called BC-mux, fills this instead with the data value of the TT adjacent in phi.

An additional reduction factor of 8 is achieved by serialising the 8-bit information at 400 Mbit/s using LVDS links. Only 4 crates, with 14 CPMs each, are required to cover the full calorimeter phase space.

B. Cluster Processor Module Data Flow

Data have to be shared between chips on board but also with chips belonging to adjacent modules. One CPM processes 64 4x4 windows or a total of 2048 TT. By sharing data between chips, and between chips from adjacent modules, the CPM only needs to receive 280 TT. For each chip, data come:

- Directly from the input of the board
- From adjacent modules in the same crate, fanned out through the backplane
- From its adjacent neighbours in phi

In total and in real-time, the CPM was designed to:

- Receive 80 LVDS signals
- Fan in/out 120 TT from/to its adjacent modules
- Transmit multiplicity information to the CMM

A custom built backplane has been designed to handle all those signals, between CPMs and from CPM to CMM. It requires a total number of 1150 pins per slot.

VI. PROTOTYPE CLUSTER PROCESSOR MODULE LAYOUT

To perform all the requirements listed in III, the CPM has been designed with the following layout:

- 80 400 Mbit/s LVDS links, to collect data from Preprocessor Modules
- 80 LVDS deserialisers, to convert data to 40 MHz 10-bit parallel words
- 20 Serialiser (SRL) chips, to distribute data at 160 MHz:
 - Onboard to perform cluster finding algorithm
 - Via backplane to adjacent modules
- Receive data at 160 MHz from neighbouring modules
- 8 CP chips, to perform the Cluster Finding Algorithm
- 2 Hit Merger chips, to calculate and transmit multiplicities to Level-1 Trigger decision via CMM
- 2 readout controller (ROC) chips, to pipeline Trigger Tower Data, multiplicities and RoI co-ordinates
- On Level-1 request, to send previous information to ROD module to help build Level-2 decision and read out to DAQ

A. Serialiser Chips (SRL)

The serialiser chip performs two tasks:

- Multiplexes and reserialises data at 160 MHz to perform the cluster finding algorithm, in real time.
- Stores data in a pipeline, waiting for a Level-1 Readout request to output them toward the readout controller

The serialiser chip has been implemented inside a FPGA Xilinx VirtexE XCV100E. There are 20 per CPM, half dedicated to the electromagnetic data, the other half to the hadronic.

B. Cluster Processor Chip (CP)

The Cluster Processor Chip requires a bigger device than the SRL chip. The algorithm has been implemented in a XCV1000E with a total of 660 pins and 1.3 million gates. The outputs of the chip are:

- RoI co-ordinates
- Sets of thresholds, among 16, which have been passed.




Figure 3: The Cluster Processor Chip latency

So far, only simulation results can give an estimation of the time required to perform the algorithm. The total latency is 6 ticks but, as seen on Figure 3, only 1 tick is required to perform the algorithm. An extra half tick has to be added to merge and calculate the multiplicity of each threshold for all 8 CP chips. It is done with two extra FPGAs called Hit Merger.

C. Readout Controller Chip (ROC)

The readout controller has been implemented in a Xilinx XCV100E FPGA. Two devices are used to drive on one path the RoI information to the Level-2 system, and, on another path, TT and hit multiplicity data for the DAQ system. Each ROC empties its local pipeline and other available pipelines of the SRL and CP chips, tags the event and sends the results via G-link to the ROD system.

VII. CLUSTER PROCESSOR MODULE PROTOTYPE

Figure 4 shows a picture of a full specification prototype CPM. This is a 9U board, 16 layers and it contains a total of 32 FPGAs: 20 SRL Chips, 8 CP chips, 2 Hit Mergers and 2 ROCs for the Level-2 and DAQ. FPGA configurations are stored in 3 FlashRams. Eighty LVDS deserialisers are positioned close to the backplane input, immediately followed by the SRL chips. The bigger devices are the 8 CP chips, and near the top front panel, the two ROC prepare data to be output via 2 G-link connectors.




Figure 4: The Cluster Processor Module

Early tests have been performed with the custom-built backplane. A 6U Concurrent CPU was mounted on an adaptor card to match the backplane layout, and a Timing Control Module was used as a spy on the VME bus. The OS system was Linux.

First tests have been to download successfully FPGA configuration inside the FlashRam, and from FlashRam to FPGA. Figure 5 shows a calibration pattern generated by the SRL chip in order to calibrate the different inputs of the CP chip.

The CP algorithm has not yet been tested as no LVDS signals were available from adjacent modules. But the SRL chip has a playback memory, also used as pipeline memory, which was used to send data to the inputs of the CP chip. Data were correctly recovered inside the CP chip, showing that the real time data path on board was working.




Figure 5: The 160 MHz calibration pattern from the SRL chip to calibrate all data inputs of the CP chip

VIII. CONCLUSION

The CPM is a complex board, designed to overcome the I/O constraint of FPGAs available on the market. It requires a high bandwidth of data across the board and between boards.

In the end, the choice of FPGAs will pay off in the future by giving a more flexible system to the trigger. Early tests have shown encouraging results but are limited as just stand-alone tests. The next step will have to integrate this board with other modules of the ATLAS Level-1 trigger system, modules at a similar stage of development. But successful results have shown that such a design is feasible, with the additional benefit to have a compact system for ATLAS.

IX. REFERENCES

- [1] ATLAS First Level-1 Trigger Technical Design Report, CERN/LHCC/98-14 and ATLAS TDR-12, 30 June 1998: <http://atlasinfo.cern.ch/GROUPS/DAQTRIG/TDR/t dr.html>
- [2] "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
- [3] <http://www.xilinx.com>