

The Phase-I Upgrade of the ATLAS First Level Calorimeter Trigger

Ivana Radoslavova HRISTOVA*, on behalf of the ATLAS Collaboration

Humboldt-Universität zu Berlin (DE)

E-mail: ivana.hristova@cern.ch

The level-1 calorimeter trigger (L1Calo) of the ATLAS experiment has been operating effectively since the start of LHC data taking, and has played a major role in the discovery of the Higgs boson. To face the new challenges posed by the upcoming increases of the LHC proton beam energy and luminosity, a series of upgrades is planned for L1Calo. An initial upgrade (Pre-Phase-I) is scheduled to be ready for the start of the second LHC run in 2015, and a further more substantial upgrade (Phase-I) is planned to be installed during the LHC shutdown expected in 2018. The calorimeter trigger aims to identify electrons, photons, taus and hadronic jets. It also determines total and missing transverse energy and can further analyse the event topology using a dedicated system incorporating information from both calorimeter and muon triggers. This paper also presents the Phase-I hardware trigger developments which exploit a tenfold increase in the available calorimeter data granularity when compared to that of the current system. The calorimeter signals will be received via optical fibres and distributed to two distinct processing systems. Those systems implement sliding window algorithms and quasi-offline algorithms to achieve object reconstruction and identification. The algorithms are implemented on high density electronics boards which make use of recent developments in high speed data transmission and FPGA technology. The presentation reviews the physics impact along with the current status of the hardware design, early prototypes and demonstrator boards.

*Technology and Instrumentation in Particle Physics 2014,
2-6 June, 2014
Amsterdam, the Netherlands*

*This research work was supported by the Bundesministerium für Bildung und Forschung (BMBF FKZ: 05H12KH2).

1. Upgrade Programme

1.1 LHC Upgrade

To reach the full potential of the LHC machine and thus to allow for exploring physics up to its frontier, luminosity production and data taking periods, *Runs*, are alternated with maintenance and upgrade installation periods, *Long Shutdowns (LS)*. At the time of writing these proceedings, the last quarter of LS1 (2013-2014) is leading to the completion of the hardware and software changes while the physics analyses of about 23 fb^{-1} of data recorded during Run 1 (2010-2012) are being finalised. During Run 2 starting in 2015 the instantaneous luminosity (L) and the average number of interactions per bunch crossing (i.e. the pile-up, $\langle\mu\rangle$) are expected to reach maximum values that are twice as large as those sustained in Run 1. Run 3, anticipated to commence in 2020 after LS2 (2018-2019), targets values of $L \sim 2.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and $\langle\mu\rangle \sim 60$ that are three times larger than those of Run 1. These projected numbers for Run 2 and Run 3 are tentative and may prove to be conservative, given the excellent performance, reliability and availability of LHC over the 3 years of Run 1 that surpassed design parameters and exceeded initial expectations. The preparation of the upgrade plans for Run 4 (2025-2028) is underway, however their review goes beyond the scope of this report.

The centre-of-mass energy (\sqrt{s}) increases gradually from 7-8 TeV in Run 1 to 13 TeV in Run 2 and reaches in Run 4 its design value of 14 TeV which is more favourable for Higgs production and searches for physics beyond the Standard Model. The initial Run 2 energy of 13 TeV, limited by the time needed to train the magnets up to full current, is possibly to be raised the ensuing year depending on experience during 2015. The bunch spacing is possibly reduced from 50 ns in Run 1 to 25 ns in the post-initial stages of Run 2; while the former remains a reserve option for the entire Run 2, the latter is preferred by experiments to minimise the event in-time pile-up. The shorter bunch spacing is achieved owing to the so-called batch compression and merging and splitting production scheme at the injector which reduces emittance and increases the brightness of the LHC proton-proton beams. A rise in the out-of-time pile-up is expected at 25 ns bunch spacing.

1.2 ATLAS TDAQ Upgrade

In response to the planned luminosity upgrades of the LHC, the ATLAS experiment has prepared a proposal for detector upgrades that require modifications and replacements of existing sub-systems as well as integration of new components into the detector. As part of the upgrade programme the trigger and data acquisition (TDAQ) system evolution has been considered and presented in [1]. The upgrade is a multi-stage programme that defines the design and preparatory activities for Run 2, Run 3 and Run 4 as Pre-Phase-I, Phase-I and Phase-II, respectively.

TDAQ is one of the largest and most complex systems of ATLAS with the trigger and data acquisition being closely integrated. The information from particle collisions in the form of input signals from the calorimeter and muon detectors is transmitted to their respective level-1 (L1) trigger electronics hardware systems where the embedded logic based on ASICs and FPGAs discriminates and counts candidate particles above threshold as well as identifies regions of interest (RoI). Upon the issue of a L1 accept (L1A) by the central trigger processor (CTP) all detector front-end buffers

	Rates		Data Flow	
	Run 1	Post-LS1	Run 1	Post-LS1
Input	20 MHz	40 MHz	1.6 MB/s	2.4 MB/s
Level-1 accept	70 kHz	100 kHz	100 GB/s	240 GB/s
Level-2 requests	25 kHz	40 kHz	8 GB/s	60 GB/s
Event building	6.5 kHz	12 kHz	10 GB/s	29 GB/s
Output	600 Hz	1 kHz	960 MB/s	2.4 GB/s

Table 1: Evolution of the TDAQ system in terms of the trigger rates and data flow throughputs. The Run 1 and the expected post-LS1 values are listed for the different trigger levels.

are read out and the RoI builder information is processed by the fast software algorithms of the level-2 (L2) trigger. Only events that pass the L2 accept criteria are fully built from the complete detector data and analysed by the event filter (EF) selection algorithms whose result finally triggers the data output to disk. The latencies of the three trigger levels are, respectively, $2.1 \mu\text{s}$, 60 ms, and 1 s. L2 and EF are jointly named the high-level trigger (HLT). Table 1 shows the evolution of the TDAQ system in terms of the trigger rates and data flow throughputs. The overall TDAQ capability is expected to increase by a factor of 2-3 in the post-LS1 operation compared to the Run 1 performance.

1.3 Level-1 Calorimeter (L1Calo) Trigger Upgrade

The L1Calo system successfully participated in Run 1 and provided high-quality information to the L1 trigger as well as to the HLT and readout systems. It is composed of pipelined custom electronics boards hosted in VME crates and cables that receive the input analogue signals from the liquid-argon (LAr) and Tile calorimeters and send the output results downstream for further treatment. The output signals correspond to electron/photon (e/γ) candidates, jets, single-hadron/tau (τ) candidates, missing transverse energy (E_T^{miss}), and total transverse energy (ΣE_T). In addition to triggering, this information is used to decide what caused the trigger, and to allow monitoring of the performance of the trigger system.

L1Calo will gradually evolve from the present-day architecture, with most of the old legacy components removed only when the new systems have been fully commissioned and validated. Smaller upgrade activities need to be accelerated and performed during LS1 in order to enhance the trigger selectivity in Run 2. The major changes take place during LS2. These proceedings describe the planned L1Calo system hardware changes with respect to the Run 1 architecture and report on the ongoing Pre-Phase-I and Phase-I upgrade activities. Further details can be found in [2].

2. L1Calo Pre-Phase-I Upgrade

As in the original design that was functional in Run 1, after LS1 the input trigger information from the upstream calorimeter electronics remains unchanged and consists of so-called trigger towers. The latter are analogue sums of energy depositions across the longitudinal layers of the calorimeters in areas of approximately $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$, 0.2×0.2 , and 0.4×0.4 in the central, endcap, and forward regions, respectively. L1Calo as configured in Run 1 will not be able to keep up with the

anticipated rise in instantaneous luminosity which necessitated flexible solutions to be sought and implemented into the existing system on a short time scale. Therefore the Pre-Phase-I upgrade items described below are mandatory for an optimal operation of L1Calo in Run 2.

2.1 PreProcessor (MCM Upgrade)

The digitisation of the analogue calorimeter trigger signals and the subsequent transverse energy (E_T) calibration, bunch-crossing identification (BCID) and the serial transmission of results to the cluster processor (CP) and jet/energy processor (JEP) sub-systems is performed in the multi-chip module (MCM) of the PreProcessor system which comprises 128 modules (PPMs) distributed over 8 crates.

As a pin-compatible replacement of the old MCM, the new MCM (nMCM) allows considerable improvements in performance without changes in the expensive hardware setup such as motherboards, crates and cabling. The production of 3,000 nMCM modules, their testing and installation in the ATLAS cavern is ongoing, to be followed by a period of integration into the L1Calo system by the end of 2014.

Lower electronics noise is inherent to the nMCMs. The replacement of the ASIC with an FPGA allows for implementing additional signal processing features and flexible algorithmic applications, such as autocorrelation FIR filter additionally to the matched filter, BCID-dependent dynamic average pedestal correction, and a dual look-up table optimised for both the electromagnetic and hadronic energy scales. Compared to the Run 1 system setup these modifications will significantly improve the E_T^{miss} and jet triggers at high pile-up as suggested by results from the analysis of simulation and reprocessed data. The Run 2 transmission speed of the output LVDS signals to the JEP system will double to 960 MBaud in order to be compatible with the Phase-I requirements.

2.2 Extended Merger Module (CMX)

The digital data from the PPMs are sent serially at 480 MBaud over LVDS cables to 56 cluster processor modules (CPMs) in 4 crates and 32 jet/energy modules (JEMs) in 2 crates. From there the results of the sliding-window algorithms that identify and count candidate e/γ , τ particles and jets as well as the computed partial E_T -sums are transmitted for summation and consolidation in the 8/4 common merger modules (CMMs) of the CP/JEP sub-systems before the counts above threshold are sent to the CTP. Each CMM receives 400 60- Ω single-ended input signals.

As a plug-compatible replacement for the CMM, the new extended CMM (CMX), both significantly expands and simultaneously provides full support of the Run 1 functionality. In addition to the multiplicity analysis of trigger objects (TOBs), the CMX is capable of accumulating and manipulating in FPGAs the TOB information which consists of its type, E_T , η and ϕ values before it is sent optically to the L1Topo processor. This increased payload requires 4-fold increase of the CPM/JEM backplane bandwidth to 160 Mbps. The CMX is able to handle the high-density high-speed signals owing to its 9 signal layers. Three working full specification prototypes and 20 production modules are available for assembly, tests and system integration.

2.3 Level-1 Topological Processor (L1Topo)

The level-1 topological processor (L1Topo) combines e/γ , μ , τ and jet candidates as well as E_T^{miss} information and performs real-time event selection based on their spacial information and geometrical relationships such as the angle, distance or invariant mass of pairs of TOBs. This is expected to improve the trigger performance by more than 30% by rejecting more background events at L1. The exact topological algorithms to be implemented in the processing FPGA firmware are still under study, but typically cuts will be applied on $\Delta\eta$, $\Delta\phi$, $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$, H_T , M_{eff} , or M_{inv} . L1Topo will be operational both in Run 2 and Run 3. The prototypes are undergoing final tests and the first batch of production modules is about to be ordered.

The L1Topo system consists of two or more L1Topo processor modules, each equipped with two processing FPGAs and one control FPGA, housed in a single ATCA shelf. Data arrive from the L1 muon and calorimeter trigger systems via four 48-way optical fibre ribbon bundles into four MTP-CPI connectors at 6.4 Gbps (link speed up to 13 Gbps supported), are then sent via octopus cables into four 12-fibre ribbons that are routed to twelve Avago miniPOD optical receivers (Rx) for opto-electrical conversion. The electrical signals are routed into the processing Xilinx Virtex-7 FPGA via its 80 on-chip multi-gigabit transceivers (MGT) and deserialised to 40 MHz. A low-latency, real-time communication path of 238 Gbps exists between the two processing FPGAs per module which provides 5.1 Gbps of payload for each of the 160 inputs assuming 8b/10b encoding, equivalent to 20480 bits per LHC bunch crossing. However, in Run 2 the input data to the processors are duplicated at source, and the two main FPGAs are supplied with the same data and operate independently and in parallel.

The algorithmic firmware is developed in a modular structure in order to decouple the core topological trigger algorithms, which require careful study and may change with time, from the common handling of FPGA I/O and the tools to provide generic functions such as sorting of TOBs, identification of overlapping TOBs, or setting cuts on angle or invariant masses of TOB pairs. The FPGA output consists of individual bits indicating the results of each specific algorithm. Each module transmits up to 64 bits per event to the CTP.

3. L1Calo Phase-I Upgrade

L1Calo will undergo a substantial transformation after Run 2 in order to adapt to the upgraded LAr calorimeter trigger electronics which will send digital fine granularity information consisting of so-called SuperCells [3]. The latter are digitised signals from all four LAr calorimeter layers (i.e. depth segmentation) and from varying $\Delta\eta \times \Delta\phi$ areas per layer. In total 10 SuperCells are sent per trigger tower, namely, 1+1 of size 0.1×0.1 in the front and rear layers, and 4+4 of size 0.025×0.1 in the two middle layers. This upgrade is necessary to accommodate the planned LHC luminosity during Run 3 and the required discriminatory power of the trigger while still operating at the disposal of the original calorimeter hardware. Detailed SuperCell simulation is being implemented in order to study the performance and optimise the system design parameters.

Two new sub-systems will be added: the electromagnetic (eFEX) and jet (jFEX) feature extractors which will gradually overtake the functions of the old CP and JEP sub-systems. Optionally a global

feature extractor (gFEX) tailored to large-area jets will be added. Common to the three FEXs are the input data sources as well as the system infrastructure design (ATCA module control, read-out, monitoring and timing interface), however the processor module requirements are sufficiently different in each case to require separate hardware designs. Therefore the standard functionality is decoupled from the specific algorithmic trigger processing and implemented in common hardware modules, however the dedicated FEX modules still share a large portion of common design standards, principles and techniques, and where feasible, firmware (e.g. IPbus firmware and software, control and configuration FPGAs). A demonstrator FEX module is being built for the various tests. The digitised E_T values are transmitted from the LAr calorimeter over optical fibres at a baseline speed of 6.4 Gbps, while above 9 Gbps links will be used if the ongoing tests prove successful. However, the trigger analogue signals from the Tile calorimeter are still digitised in the L1Calo PPMs and delivered to the JEMs by new link daughter cards. Data is transmitted with standard 8b/10b encoding which corresponds to 128 bits per bunch crossing per fibre.

The high density of high-speed PCB signal tracks is of particular concern in the design of the FEX modules. The following sections briefly describe the hardware and the functionality that are specific to each FEX sub-system. A comparison of Run 1 and Run 3 parameters is listed in Table 2.

3.1 eFEX System

Each optical fibre from the electromagnetic and hadronic calorimeters carries the data from a trigger tower area of 0.2×0.1 and 0.4×0.2 , respectively. The connection to the eFEX is made via four 48-way MTP connectors mounted on the ATCA backplane, and to the Processor FPGA via twelve 12-channel Avago miniPOD transceivers mounted in-board. Of the total received 136 multi-Gbps signals 52 are directed to a single FPGA, 72 to two FPGAs and 12 to three FPGAs. The signal mapping has been carefully optimised taking into account the system layout, cable paths and the needs of the processing algorithms. The three-way fanout is implemented using discrete, high-speed electrical buffers. For the two-way fan out a less demanding method in terms of layout components is considered, namely, PMA loopback of signals through the multi-gigabit receiver/transmitter pairs in the FPGAs. A signal delay of 25 ns and some degradation of signal quality is associated with this technique.

Four Processor FPGAs per eFEX are responsible for identifying TOBs by evaluating the object type (e, γ or τ), the measured energy, E_T , and the η and ϕ coordinates. The input and output data are also recorded in scrolling memories, and on receipt of L1A signal, transferred to a derandomising buffer where a data packet is built, including the bunch crossing number, and transmitted to the Read-out Interface FPGA. One of the Processor FPGAs collects all the TOBs in a single output packet and transmits copies serially at 6.4 Gbps via electrical-optical transmitters to up to six L1Topo modules. The eFEX power consumption is estimated to be 190 W per module.

3.2 jFEX and gFEX Systems

Each optical fibre from the electromagnetic and hadronic calorimeters carries the data from an area of trigger towers of 0.4×0.2 . The connection to the jFEX is made via four 72-way MTP/MPO connectors mounted on the ATCA backplane, and to the Processor FPGA via 24 12-channel MicroPOD

Run 1	Run 3
CPM	eFEX
identify e/γ and τ candidates	
56 CPMs in 4 VME crates	24 eFEX modules in 2 ATCA shelves
$\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ trigger towers	up to 10 (1+4+4+1) LAr SuperCells per trigger tower; layers 1 and 2: 0.025×0.1 , layers 0 and 3: 0.1×0.1
0.4×0.4 sliding-window algorithm with no electromagnetic calorimeter depth segmentation seeded by max. E_T -sum of 2 adjacent trigger towers in 0.2×0.2 window	sliding window-algorithm seeded by max. E_T SuperCell in layer 2, add max. E_T adjacent in ϕ and both neighbour SuperCells in η ; cluster of 3×2 SuperCells; add corresponding cells in layer 1; add pair of cells in layers 0 and 3 around the seed and its max. E_T neighbour in ϕ
isolation: E_T -sums in 2×2 hadronic towers behind and 12 electromagnetic towers surrounding the cluster	flexible jet rejection algorithms: R_η , R_{core} (adapted from offline lepton identification)
1 GeV baseline E_T scale	250 MeV baseline E_T scale
core area 0.4×1.6 , environment area 0.7×1.9	core area up to 1.7×0.8 , environment area 1.8×1.0
LAr analogue inputs to PreProcessor	LAr digital inputs via 6.4 Gbps optical links (> 9 Gbps under investigation)
280 electromagnetic and hadronic trigger tower signals, 160 directly transferred	> 230 electromagnetic and hadronic signals, 90/25 (up to 100/36) electromagnetic/hadronic fibres directly transferred
8 CP ASICs, 0.4×0.2 region per ASIC	4 Xilinx FPGAs, 0.6×1.0 region per FPGA
JEM	jFEX
identify jet, large-area τ , $\sum E_T$ and E_T^{miss} trigger candidates	
32 JEMs in 2 VME crates	8 jFEXs in 1 ATCA shelf
$\Delta\eta \times \Delta\phi = 0.2 \times 0.2$ jet elements	0.1×0.1 electromagnetic and hadronic trigger towers
0.8×0.8 sliding window algorithm seeded by 0.4×0.4 jet cluster	0.9×0.9 (1.7×1.7) sliding window algorithm at 6.4 (10) Gbps, flexible jet algorithms using shape variables and radial weights, jet-area pile-up correction
1 GeV baseline E_T scale	250 MeV baseline E_T scale
core area 0.8×1.6 , environment area 1.4×2.2 , per JEM	core area 1.2×0.8 , environment area 2.0×1.6 , per FPGA
LAr analogue inputs to PreProcessor	LAr digital inputs via 6.4 Gbps optical links (> 9 Gbps under investigation)
77 jet element signals, 44 directly transferred	416 electromagnetic and hadronic signals per module, 208 directly transferred
1/1 Jet/Sum Altera FPGA	4 Xilinx FPGAs

Table 2: Comparison between the Run 1 (legacy CPM, JEM) and Run 3 (upgrade eFEX, jFEX) systems in terms of their design and functionality features.

devices supported in-board by custom-built mechanics to provide stability and heat dissipation. The input data are partitioned between four Processor FPGAs by making use of the PMA loopback scheme where needed and received as differential signals to the GTH high-speed receivers of each FPGA. A large amount of data needs to be duplicated between neighbouring FPGAs to allow the algorithms to access the environment data from four towers around the core tower. Instead of transmission between jFEXs, data sharing is achieved by sending two copies of all signals from the LAr calorimeter and JEM daughter cards to a jFEX module.

The four Processor FPGAs per jFEX identify TOBs as jet and τ candidates and calculate E_T^{miss} separately in segments of η . This allows for pile-up corrections to be performed by algorithms implemented in downstream devices, such as an in-board FPGA or the L1Topo processor. One of the Processor FPGAs collects all the TOBs and transmits copies serially at 6.4 Gbps via two MicroPOD devices (i.e. 24 optical fibres) to up to six L1Topo modules.

The design described above corresponds to an environment area of 0.9×0.9 , while possible use of higher speed links would allow the support of larger jet sizes, up to 1.7×1.7 corresponding to a jet radius parameter of 0.85. However, the feasibility for coping with higher data rates, higher signal counts and data duplication schemes in dense electro-optical systems is still under study. Therefore the addition of a separate dedicated gFEX module has been proposed for identifying large-area electro-weak jets in regions of up to 1.8×1.8 as well as boosted physics objects at high \sqrt{s} .

The gFEX design shares many features and functional components with the eFEX and jFEX modules. Data from the electromagnetic and hadronic calorimeter from an area of 0.2×0.2 are sent via 32 optical links per 0.8η ring, for six rings, to large FPGAs with up to 96 transceivers. The input fibres are connected to 16 12-channel parallel-optic receivers and distributed among four Processor FPGAs directly or via two-way or three-way fan out, depending on the η range. Two gFEX modules in a single ATCA shelf each cover half the calorimeter η range. The architecture is suitable for other trigger algorithms which require the same granularity including E_T^{miss} and $\sum E_T$.

4. Summary

The upgrade of the ATLAS level-1 calorimeter trigger is necessary in order to increase the rejection power and preserve selectivity of the trigger at the high luminosity and pile-up during the LHC Run 2 and Run 3. During Run 3 the fine granularity digital information from the calorimeter is handled by offline-like algorithms running in new feature extractor modules, eFEX, jFEX and gFEX, featuring modern FPGAs, high-density circuit boards and high-speed optical transmission.

References

- [1] ATLAS Collaboration, *ATLAS TDAQ System Phase-I Upgrade Technical Design Report*. CERN-LHCC-2013-018, ATLAS-TDR-023, 2013. <https://cds.cern.ch/record/1602235>.
- [2] M. Wessels, on behalf of the ATLAS Collaboration, *Upgrade of the ATLAS Level-1 Calorimeter*. ATL-DAQ-PROC-2014-003, 2013. <https://cds.cern.ch/record/1703259>.
- [3] ATLAS Collaboration, *ATLAS Liquid Argon Calorimeter Phase-I Upgrade Technical Design Report*. CERN-LHCC-2013-017, ATLAS-TDR-022, 2013. <https://cds.cern.ch/record/1602230>.