

ATLAS Level-1 Calorimeter Trigger Upgrade for Phase-I

Weiming Qian, on behalf of the ATLAS Level-1 Calorimeter Trigger group

*STFC Rutherford Appleton Laboratory,
Harwell Oxford, Chilton, Didcot, Oxon, OX11 0QX, United Kingdom*

E-mail: Weiming.Qian@stfc.ac.uk

ABSTRACT: The ATLAS Level-1 Trigger requires several upgrades to maintain physics sensitivity as the LHC luminosity is raised. One of the most challenging is the electron/photon trigger, with a major development planned for installation in 2018. New on-detector electronics will be installed to digitize electromagnetic calorimetry signals, providing trigger access to shower profile information. The trigger processing will be ATCA-based, with each multi-FPGA module processing ~ 1 Tbit/s of calorimeter digits within the current 2.5 microseconds Level-1 Trigger latency limit. This paper will address the system architecture and design, and give the status of a current technology demonstrator.

KEYWORDS: Trigger algorithm, High-speed digital design, PCB simulation, ATCA.



Contents

1. Introduction	1
1.1 LAr calorimeter new digital TBB and R_{core} algorithm	2
1.2 ATLAS L1Calo Trigger Upgrade for Phase-I	2
1.3 eFEX architecture	3
1.4 eFEX design challenges	4
1.4.1 Multi-Gb/s data sharing	4
1.4.2 Density	5
1.4.3 Multi-Gb/s PCB design	5
2. High Speed Demonstrator (HSD)	5
2.1 High-speed PCB simulation	6
2.2 HSD initial test results	6
2.2.1 TDR test	6
2.2.2 TDT test	7
2.2.3 Differential skew test	8
3. Conclusion	8

1. Introduction

ATLAS [1] is one of the multi-purpose particle physics experiments at the Large Hadron Collider (LHC) at CERN. At its centre, bunches of protons are collided head-on at 40 MHz. A Trigger system is used to select only those bunch collisions with physics interest for further off-line analysis. The current ATLAS Trigger system has been designed to work up to the LHC design luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ and to produce an average output rate of 200 Hz. It consists of 3 levels. The Level-1 Trigger [2] is implemented in custom-built, pipelined, synchronous electronics. The Level-1 Trigger rate is limited by the readout bandwidth of the front-end detector electronics to 100 kHz or less, and the latency of the Level-1 Trigger is limited by the depth of the front-end detector pipeline memories to 2.5 μs .

A series of upgrades to the LHC have been planned to improve its physics potential. The present schedule foresees a major upgrade in 2018, which is called Phase-I, to increase the luminosity to $\sim 2 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$. However, the majority of the ATLAS front-end electronics will remain unchanged at Phase-I, so the above limits on the Level-1 Trigger rate and latency will still apply. The current ATLAS Level-1 Trigger system cannot cope with all of these requirements without increasing the trigger thresholds to a level that would undermine the sensitivity of ATLAS to physics processes of interest. This motivates the Phase-I Upgrade to the ATLAS Level-1 Trigger system.

The ATLAS Level-1 Trigger consists of three subsystems: Level-1 Calorimeter (L1Calo) Trigger, Level-1 Muon Trigger and Central Trigger Processor (CTP). The current L1Calo

Trigger uses trigger tower sums from Liquid Argon (LAr) electromagnetic calorimeter and Tile hadronic calorimeter, which are energy summations formed by on-detector Tower Builder Boards (TBB) with a typical granularity of 0.1×0.1 ($\eta \times \phi$). The strategy for the L1Calo Trigger upgrade for Phase-I is to use higher granularity information from the LAr electromagnetic calorimeter to run more effective algorithms, improving electron/photon isolation and energy cuts. Initial physics simulation studies have shown good results with this approach.

1.1 LAr calorimeter new digital TBB and R_{core} algorithm

Figure 1 shows the detailed structure of the LAr calorimeter and a logical trigger tower. The LAr calorimeter has four layers in depth: presampler, strip layer, middle layer and back layer. Currently, signals from all the calorimeter cells on all the layers within a trigger tower are summed together to give a single energy sum. For the Phase-I upgrade, a new digital TBB will be installed on the LAr calorimeter to digitize the calorimeter signals and provide trigger access to higher-granularity shower-profile information. The diagram on the right of **Figure 1** shows the baseline design of the new digital TBB, in which each trigger tower is divided into 10 subsums.




Figure 1. Left: LAr calorimeter internal structure. Middle: Trigger tower logical structure. Right: Higher granularity trigger tower subsums for Phase-I upgrade

Figure 2 shows a new algorithm, R_{core} , which will operate on the middle layer of the higher granularity trigger tower subsums. Compared to the current algorithm, which can only set static isolation thresholds, this new algorithm can in addition use shower shape information to select narrow showers such as those arising from electrons or photons. One quantity of R_{core} that can be exploited is the measurement of the width of a shower in transverse direction. R_{core} is defined as $\sum 2 \times 3 / \sum 2 \times 7 > a$.




Figure 2. R_{core} algorithm

1.2 ATLAS L1Calo Trigger Upgrade for Phase-I

The overall structure of the upgraded ATLAS Level-1 Trigger is shown in **Figure 3**. The green box represents L1Calo Trigger at Phase-I. The current L1Calo system, consisting of the Preprocessor, Jet/Engery Processor and Electron/Tau Processor, will continue to run at least in part throughout the Phase-I period. Areas in bright read represent changes from the current system prior to Phase-I: a new Multi-Chip Module (nMCM) will replace the current Multi-Chip

Module on the Preprocessor to improve digital filtering, and a new merger module (CMX) will replace the current Common Merger Module to provide energy and position information needed by the Level-1 topology trigger processor.

For the Phase-I upgrade, L1Calo will include 2 new subsystems: a Digital Processing System (DPS) and a Feature Extractor subsystem with electron/photon processor (eFEX) and jet processor (jFEX). The DPS will perform digital filtering and energy calibration on the higher granularity calorimeter information received from the new digital TBB on the LAr calorimeter. The eFEX/jFEX will run new trigger algorithms (e.g. R_{core}) to identify calorimeter trigger signatures, which will then be sent to the Level-1 Topology trigger processor (L1Topo) and CTP. The Tile calorimeter front-end electronics will remain unchanged at Phase-I, and thus a new JEM Daughter Board (JemDboard) will be developed to provide hadronic calorimetry information to the eFEX/jFEX processors. The data links from the DPS/JemDboard to the eFEX/jFEX will run in the range 6–10 Gb/s over optical fibres. A complex Optical Patch Panel (OPP) will be designed to map from detector to trigger geometry.

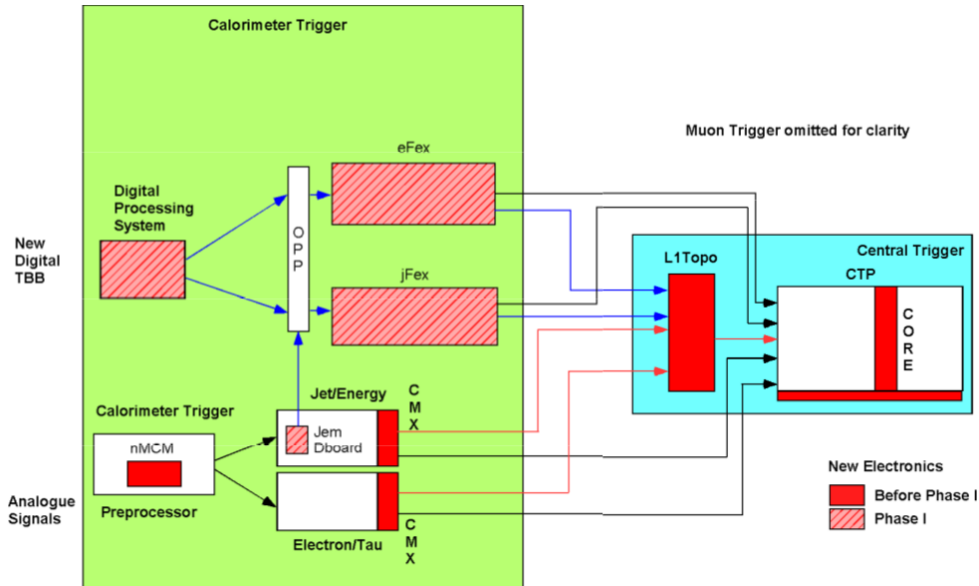


Figure 3. L1Calo view of ATLAS Level-1 Trigger at Phase-I upgrade

1.3 eFEX architecture

The eFEX processor will search for electron/photon signatures in the region $-2.5 \leq \eta \leq 2.5$ and $0 \leq \phi \leq 2\pi$. It will be designed as a modular subsystem using the ATCA [3] standard, with each module responsible for processing a core area that is a subset of the whole trigger region. As the new algorithms are still based on sliding windows, an eFEX module searching electron/photon signatures in its core area also needs environment information from its surrounding area. Hence, a substantial volume of calorimeter information will be shared between neighboring modules. The partitioning of the calorimeter data into eFEX modules needs to balance the total number of modules, the fibre count per module, the complexity of fibre mapping between DPS and eFEX, and the difficulty of data sharing between adjacent modules. Many partitioning scenarios have been explored, with the following two of particular interest because of their simple 1-to-2 data sharing requirement on the borders:

- a) An eFEX module processes a core area of $0 \leq \phi \leq 2\pi$ and $\Delta\eta = 0.4$, i.e. whole slice in ϕ .

- b) An eFEX module processes a core area of $-2.5 \leq \eta \leq 2.5$ and $\Delta\phi = 0.3$, i.e. whole slice in η .

Partitions larger than these require a super-dense eFEX module that is too difficult to achieve with current technology. Partitions that do not use a whole slice in either ϕ or η require larger numbers of eFEX modules to implement, resulting in a less efficient system. They also require both 1-to-2 data sharing on the module borders and also 1-to-4 data sharing on the module corners.

Based on the above partitioning analysis, an eFEX processor is foreseen which consists of about 20 eFEX modules housed in two ATCA crates. Each eFEX module will receive data from the DPS over a maximum of 200 optical fibres running at 6–10 Gb/s. These will be received on the modules by 12-channel optical receivers (Avago MiniPOD or MicroPOD [4]), which will be mounted in-board to minimize the length of the high-speed PCB tracks. To make testing and module service easier, all the input optical fibres will be routed through ATCA zone 3 via optical backplane connectors. There will be at least 4 Xilinx Virtex-7 high-end FPGAs on each eFEX module for algorithm processing, with additional FPGA resources dedicated to implementing Readout Driver (ROD) and Region-of-Interest (coordinates in η/ϕ plane of trigger objects) functions.

1.4 eFEX design challenges

The eFEX processor will be a very dense system, running at very high-speed and based on an ATCA standard of which there is limited experience within the particle physics community. There are many challenges in designing such a system.

1.4.1 Multi-Gb/s data sharing

As previously discussed, data sharing is essential for the eFEX algorithm, based as it is on sliding windows. In the eFEX processor, data sharing is needed at three levels: between modules in different crates, between modules within the same crate, and between FPGAs on the same module. Data are transported from the DPS to the eFEX over multi-Gb/s serial optical links. Electrical fan-out on a PCB at multi-Gb/s speeds is very challenging. At present, the maximum usable PCB trace length from an optical transceiver on a host board is about 20cm [5], which means that electrical fan-out of multi-Gb/s links, if it works, cannot be used between eFEX modules directly. A possible way to bypass this signal integrity problem associated with optical links would be to de-serialize and re-serialize data inside the FPGAs on the eFEX modules, before resending the data to neighbouring modules at multi-Gb/s over the backplane or electrical cables. Unfortunately, deserilization/serialization requires 3~4 clock ticks (of the 40 MHz LHC clock). It must therefore be discounted as the requirement that the L1Calo upgrade stays within the current 2.5 μ s Level-1 latency envelope imposes a very tight latency budget. This leaves two options for data sharing between eFEX modules:

- a) Passive optical splitting. This is limited to 1-to-2 splitting by the optical power budget of MiniPOD/MicroPOD transceivers and it is very marginal, hence requiring a very careful optical patch panel design to minimize optical insertion loss.
- b) Serial link duplication at the DPS. This is the preferred solution as it solves the signal integrity issues related to optical links and does not increase the overall latency.

1.4.2 Density

Zone 3 of the ATCA form factor can accommodate only 4 MTP-CPI [6] optical backplane connectors and, due to the larger optical attenuation found in devices of higher multiplicity, it is desirable to use connectors that utilize a maximum of 48 fibres each. In total, then, a maximum of 192 fibres can be routed through ATCA zone 3 per module, which therefore cannot support the two eFEX partition scenarios proposed above without further optimization.

1.4.3 Multi-Gb/s PCB design

Multi-Gb/s PCB design is probably the biggest challenge, where the signal integrity issue is ubiquitous in the whole system. To make a multi-Gb/s PCB successfully, a very careful design is needed in all the following areas:

- Impedance control
- High frequency attenuation
- Crosstalk
- Clock jitter
- PCB differential skew
- Power distribution system

In the multi-Gb/s speed range, the problems in the above areas are entangled together, making system testing/debugging extremely difficult.

2. High Speed Demonstrator (HSD)

Given such big challenges and limited experience in these areas, a relatively simple ATCA module, High Speed Demonstrator, has been designed to explore these new technologies.

The main purpose of the HSD is to explore multi-Gb/s PCB simulation and the correlation of simulation versus hardware measurement, so that PCB simulation can be used to guide the future eFEX design. A systematic methodology (**Figure 4**) has been demonstrated in the HSD design process.

The HSD (**Figure 5**) uses a Xilinx Virtex-6 FPGA (XC6VHX255T) [7] as the multi-Gb/s data sink/source. It has 24 GTX (gigabit transceivers) running at 5 Gb/s and 24 GTH (high-speed gigabit transceivers) running at 10 Gb/s. Many new technologies are implemented, including:

- Clock jitter-cleaning circuitry
- Multi-Gb/s fan-out circuitry
- 12-way parallel optical transmitters and receivers
- ATCA Intelligent Platform Management Controller (IPMC)
- Blind via PCB technology

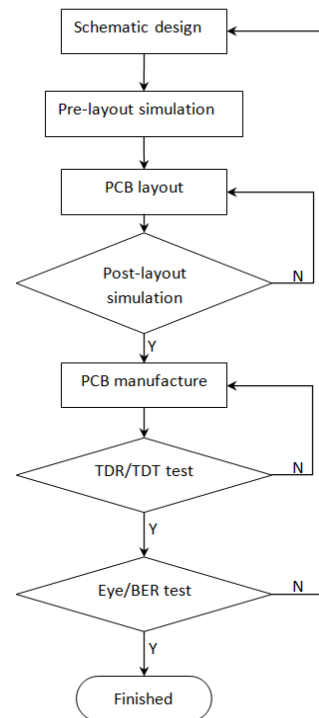


Figure 4. HSD design procedure




Figure 5. HSD

A variety of serial links on the HSD are terminated at SMA connectors, thus enabling various link topologies to be easily formed and measured.

2.1 High-speed PCB simulation

PCB simulation has been used at various stages during the HSD design flow.

The pre-layout simulation is used to derive PCB layout rules on PCB stack-up, trace impedance control, PCB material selection, crosstalk control, PCB via design and maximum achievable PCB trace lengths. An example is shown in **Figure 6**, in which the signal loss is simulated for two PCB materials. For the standard PCB material FR4, with loss tangent of 0.035 (left graph), the loss is dominated by dielectric loss above 1 GHz. Using a modified FR4 with a loss tangent of 0.01 (right graph), the PCB dielectric loss is brought down under the PCB resistive loss up to 10 GHz. From this simulation, it was concluded that a PCB material with loss tangent of 0.01 is adequate for the HSD. Any PCB material with even smaller loss tangent would increase the cost significantly without much improvement in the overall PCB performance.




Figure 6. PCB material selection for HSD

The post-layout simulation is done when the PCB layout is finished, to verify the quality of layout, and employs models of the serial-link channels extracted directly from the finished PCB layout. In PCB Time Domain Reflectometry (TDR)/Time Domain Transmission (TDT) tests, the performance of the serial link channels are measured both in time and the frequency domains, and then correlated to the post-layout simulation. Any unexpected difference should be understood and accounted for (an example is shown in next section). In the Eye/Bit-Error-Rate (BER) test, the measured Eye diagram and BER are compared to the simulation, and the latter is also used to tune the equalization setting of multi-Gb/s transceivers.

2.2 HSD initial test results

Some initial tests on the HSD are presented in this section.

2.2.1 TDR test

Figure 7 shows the TDR test on one PCB channel of the HSD. The measured differential impedance is 110 Ω , just on the edge of the PCB impedance specification ($100 \Omega \pm 10\%$). However, there is a huge negative reflection at the channel entrance point. This problem has been traced to the SMA launch pad, as shown

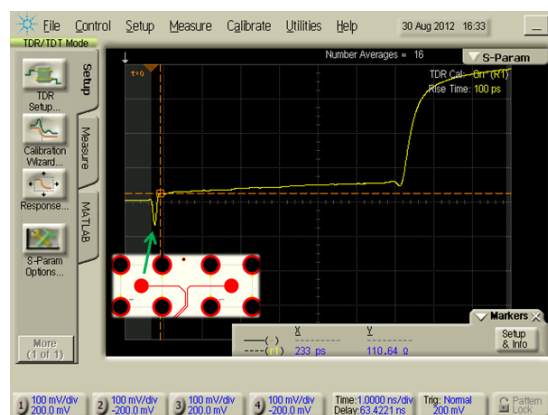


Figure 7. HSD TDR test

in **Figure 7**. The SMA launch pad is the mounting point for the centre signal pin of the SMA connector, and hence matches its size. This relatively large pad (1.8mm in diameter) creates an excessive capacitance of ~ 1.2 pf, causing a negative reflection in the TDR test.

2.2.2 TDT test

Figure 8 shows the insertion loss extracted from the TDT test on the same PCB channel, together with PCB simulations. The blue curve is the original PCB channel simulation, which deviates from the measurement (red curve) significantly. The green curve is the PCB channel simulation taking into account the excessive SMA launch capacitance. It agrees with the measurement well to 6.5 GHz. Beyond this the TDT measurement reaches the noise floor of the scope, but the simulation still traces the trend reasonably well.

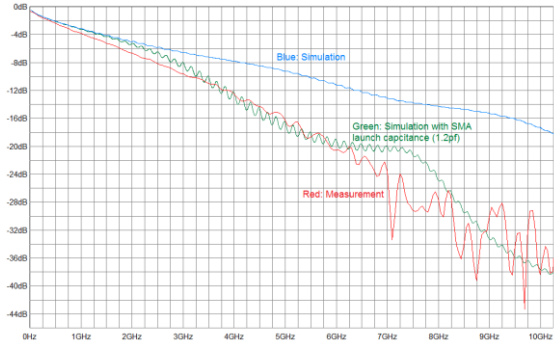


Figure 8. HSD TDT test

In order to optimize the SMA launch performance, the SMA launch structure is modeled in 3D EM solver (ANSYS HFSS) as shown in **Figure 9**. The top model is the current SMA launch on the HSD. The bottom model is the optimized SMA launch. The ground planes are highlighted in pink in both 3D models. The graphs to the right are the corresponding S-parameters with insertion losses in brown and return losses in red. It becomes clear from these 3D simulations that a circular ring cut on the ground plane underneath the SMA signal pin pad would greatly improve the SMA launch performance.

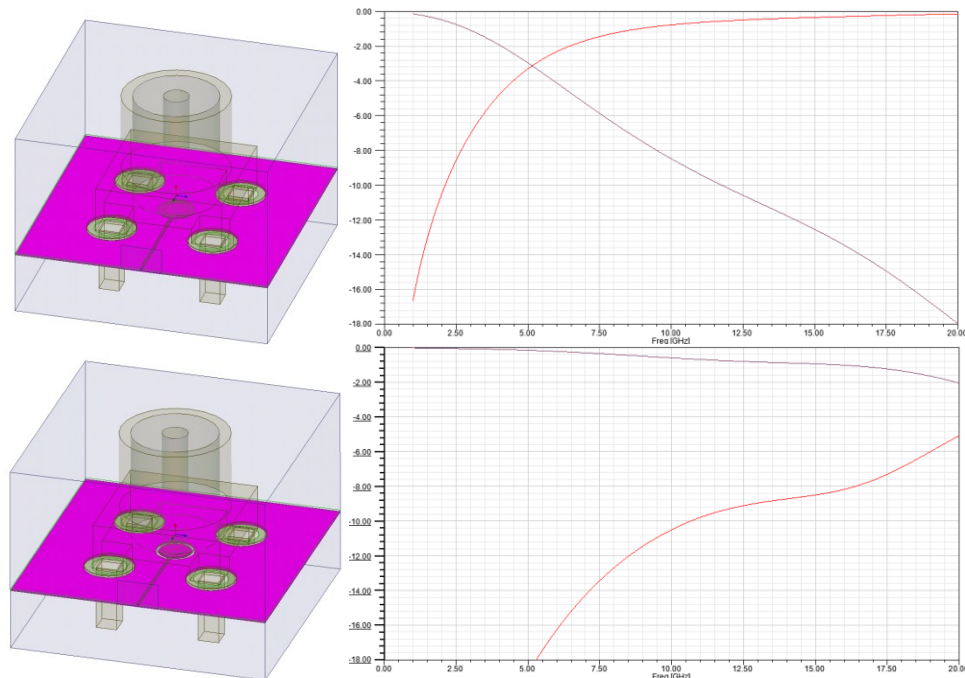


Figure 9. 3D modeling of SMA launch structure

2.2.3 Differential skew test

Differential skew control is very important for multi-Gb/s PCB design. The design goal is that the differential skew should be less than 0.1 Unit of Interval (UI) of the serial data stream, e.g. 10 ps at 10 Gb/s. The differential skew on a PCB is caused by the weave pattern of the fibre glass in the PCB dielectric material. There are many ways to reduce the PCB differential skew, including routing signal traces at small angle zigzags or choosing a specific weave pattern for the PCB material. In the HSD manufacture, we specified a 22° rotation of the module relative to the PCB panel. **Figure 10** shows a resulting ~1 ps differential skew, an excellent result on a 50cm channel.

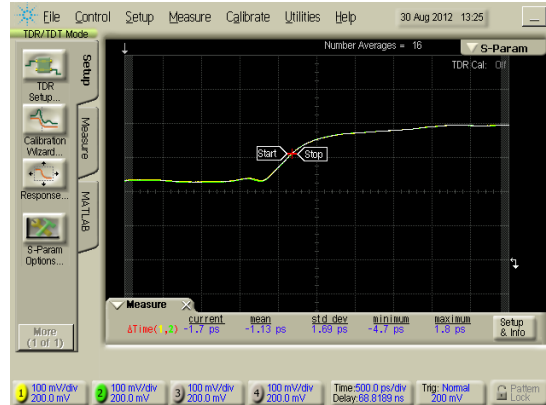


Figure 10. HSD differential skew test

3. Conclusion

The ATLAS L1Calo Trigger will be upgraded as part of the ATLAS Phase-I upgrades for 2018. The conceptual design has been explored with many problems/challenges well understood. The detailed system/module specifications will depend on further discussions between the ATLAS L1Calo and LAr groups. A High-Speed Demonstrator has been designed to test new technologies enabling the trigger upgrade. A new PCB design methodology centred around PCB simulation and measurement has been demonstrated in the HSD design flow. The HSD initial tests have shown some good results and uncovered some problems at the same time. Simulation is an essential tool in diagnosing PCB problems.

References

- [1] ATLAS collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, 2008 *JINST* **3** S08003. <http://iopscience.iop.org/1748-0221/3/08/S08003>
- [2] ATLAS collaboration, *ATLAS Level-1 Trigger: TDR*, CERN/LHCC/98-14. <http://atlas.web.cern.ch/Atlas/GROUPS/DAQTRIG/TDR/tdr.html>
- [3] AdvancedTCA PCIMG 3.0 Short Form Specification http://www.picmg.org/pdf/PICMG_3_0_Shortform.pdf
- [4] MicroPOD™ and MiniPOD™ 120G Transmitters/Receivers http://www.avagotech.com/pages/minipod_micropod
- [5] SFF-8431 Specifications for Enhanced Small Form Factor Pluggable Module SFP+ <ftp://ftp.seagate.com/sff/SFF-8431.PDF>
- [6] MTP-CPI Coplanar Optical Backplane Connector System <http://www.molex.com/molex/products/family?channel=products&chanName=family&key=mtpcpi>
- [7] <http://www.xilinx.com/products/silicon-devices/fpga/virtex-6/index.htm>