Operation of the upgraded ATLAS Central Trigger Processor during the LHC Run 2

H. Bertelsen[a]**, G. Carrillo Montoya** [b]**, P.-O. Deviveiros** [b]**, T. Eifert** [b]**, G. Galster**[a,b]**,**
J. Glatzer[b]**, S. Haas**[b]**, A. Marzin**[b]**, M. V. Silva Oliveira**[b,c]**, T. Pauly**[b]**, K. Schmieden**[b]**,**
R. Spiwoks[b]**, J. Stelzer**[b]

[a]*University of Copenhagen,*
 Nøorregade 10, PO Box 2177 1017 Copenhagen
[b]*CERN European European Organization for Nuclear Research,*
 CERN, CH-1211 Geneva 23, Switzerland
[c]*Universidade Federal de Juiz de Fora,*
 José Lourenço Kelmer - Martelos, Juiz de Fora - MG, 36036-330, Brazil

 E-mail: `antoine.marzin@cern.ch`

ABSTRACT: The ATLAS Central Trigger Processor (CTP) is responsible for forming the Level-1 trigger decision based on the information from the calorimeter and muon trigger processors. In order to cope with the increase of luminosity and physics cross-sections in Run 2, several components of this system have been upgraded. In particular, the number of usable trigger inputs and trigger items have been increased from 160 to 512 and from 256 to 512, respectively. The upgraded CTP also provides extended monitoring capabilities and allows to operate simultaneously up to three independent combinations of sub-detectors with full trigger functionality, which is particularly useful for commissioning, calibration and test runs. The software has also undergone a major upgrade to take advantage of all these new functionalities. An overview of the commissioning and the operation of the upgraded CTP during the LHC Run 2 is given.

KEYWORDS: ATLAS, Central Trigger Processor, Trigger upgrade.

1. Introduction

The ATLAS experiment [**?**] is a particle detector located at the Large Hadron Collider (LHC) [**?**] at CERN. The LHC is a proton-proton and heavy ion collider with a circumference of 27 km. It is designed to operate with a collision rate of 40 MHz and a center-of-mass energy of 14 TeV. During Run 1 in 2012, the LHC delivered an integrated luminosity of 20 fb^{-1} at $\sqrt{s} = 8$ TeV and with a bunch-spacing of 50 ns. For the Run 2 data-taking period which has started in 2015, the center-of-mass energy has been increased to 13 TeV and the bunch-spacing has been reduced to the nominal value of 25 ns. The higher cross-sections, in particular for strong interaction dominated processes, and instantaneous luminosity increase the total interaction rate per bunch-crossing by up to a factor of six compared to Run 1. In order to cope with these harsher conditions, the ATLAS trigger system has undergone a major upgrade during the shutdown of 2013/2014.

The ATLAS trigger system is implemented as two stages to select collision events relevant for physics analyses. The first stage, called the Level-1 trigger (L1), uses custom-built hardware to reduce the 40 MHz bunch-crossing rate to a trigger rate of up to 100 kHz. The trigger decision is formed by the Central Trigger Processor (CTP) [**?**] based on coarse-granularity information from the calorimeters and dedicated muon chambers. The second stage, called the High-Level trigger (HLT), is software-based and further reduces the trigger rate to 1 kHz using fast reconstruction algorithms in detector regions of interest provided by the L1 trigger. Unlike L1, the HLT can use information from all sub-detectors with their full granularity.

The trigger menu consists of trigger items, each of which is a logical function of trigger inputs. The CTP was already used to the limits of its capacity during the Run 1, and therefore it has been upgraded to increase its available resources. For instance the number of trigger inputs was increased from 160 to 512 and the number of trigger items from 256 to 512 to have more flexibility in the selection of events. The installation of the new L1 Topological Processor (L1Topo) [**?**] also required an upgrade of the interface between the muon trigger and CTP. Furthermore, up to three independent combinations of sub-detectors can now use the CTP simultaneously for parallel commissioning and calibration.

2. Overview of the Level-1 trigger system

A schematic of the upgraded L1 trigger system is shown in figure **??**. The L1 Central Trigger (L1CT) systems which have been upgraded for Run 2 are indicated by dashed lines. Information from the electromagnetic and hadronic calorimeters are processed by the L1Calo system. This uses coarse granularity calorimeter towers to identify electrons, photons, taus and jets, and then provides a count of the number of objects above various energy thresholds as well as the transverse energy sum and missing transverse energy. The L1Muon system uses dedicated Resistive Plate Chambers (RPC) and Thin Gap Chambers (TGC) with fast readout to identify muons in the barrel and end-caps, respectively. The L1Topo module allows selection of events based on their topology, such as the angular separation between trigger objects. The Muon-to-CTP Interface (MUCTPI) [**?**] is composed of 16 Muon Interface Octant (MIOCT) modules. It receives trigger information from 208 L1Muon sectors, calculates the multiplicity of muon candidates for up to six momentum thresholds, and sends them to the CTP. In order to transmit the η and ϕ coordinates of the muons to L1Topo,
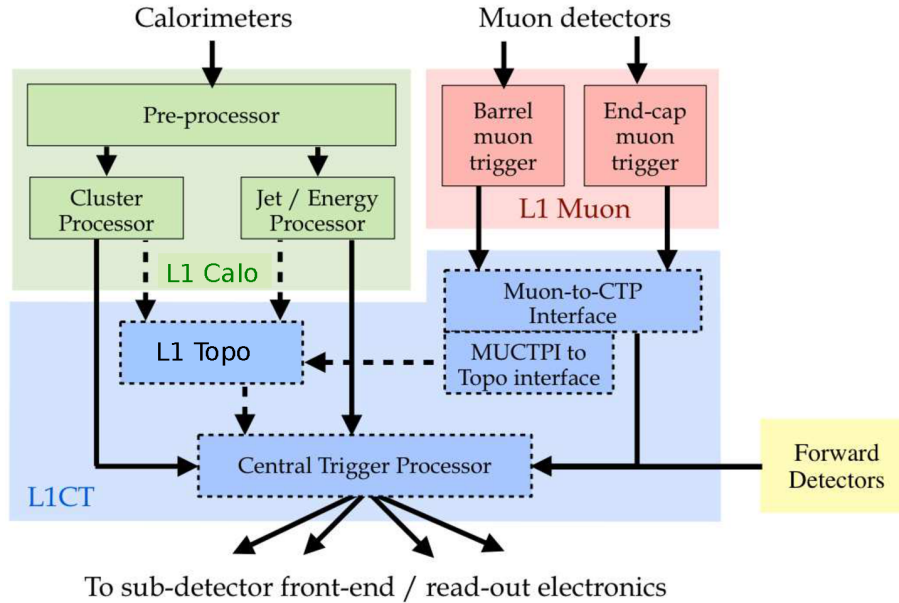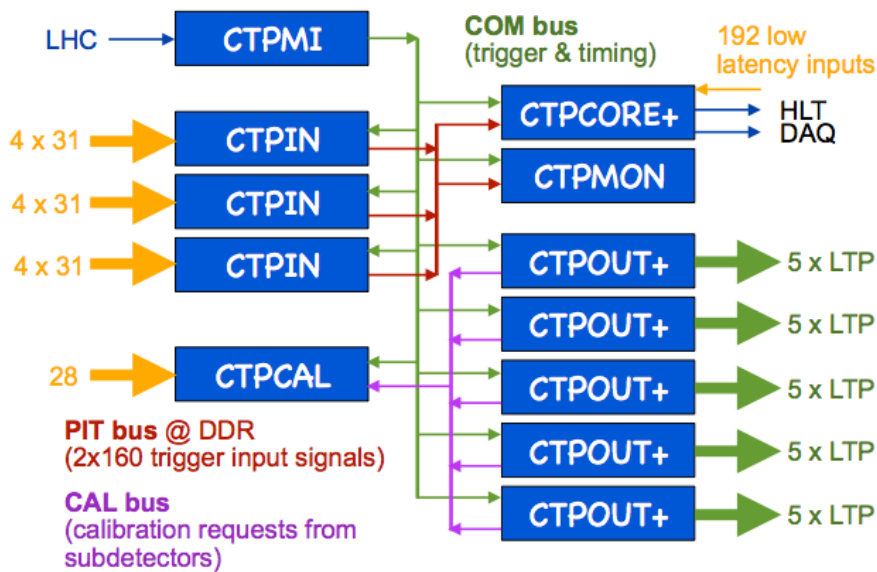
Figure 1. Schematic overview of the Level-1 trigger system. The dashed lines indicate the upgraded L1CT systems for Run 2.

the MUCTPI firmware has been upgraded to overclock the two electrical outputs of each MIOCT to a rate of 320MHz. In addition to the MUCTPI firmware upgrade, a new MuCTPiToTopo interface module has been installed to convert the electrical MIOCT output to optical signals for the L1Topo input. Finally, the CTP forms the L1-Accept (L1A) signal and transmit it to the sub-detectors.

3. Upgrade of the Central Trigger Processor

The CTP forms the L1 trigger decision based on the information received from L1Topo, L1Calo and L1Muon, and distributes the L1A and LHC timing signals to the sub-detector readout systems via the Timing, Trigger and Control (TTC) network. An overview of the CTP modules with their connections is shown in figure **??**. It is composed of the following custom-built VME modules which are housed in a single 9U VME crate:

- CTP Machine Interface (CTPMI): receives timing signals from the LHC and distributes them to the other modules. No upgrade was needed for Run 2.

- CTP Input Modules (CTPIN): receive and align the trigger inputs coming from L1Calo and L1Muon. Each of the three CTPIN modules receives up to 124 trigger inputs on four cables, yielding a total of 372 trigger inputs. A switch matrix selects and routes the trigger inputs which are sent to the CTPCORE+ and CTPMON modules through the Pattern In Time (PIT) backplane. The firmware has been upgraded to use the PIT bus at double the data rate (80 MHz) in order to use 320 trigger inputs instead of 160 as in Run 1.

- Upgraded CTP Core Module (CTPCORE+): this module has been completely re-designed and re-built for Run 2 to support more trigger inputs and double the number of trigger items
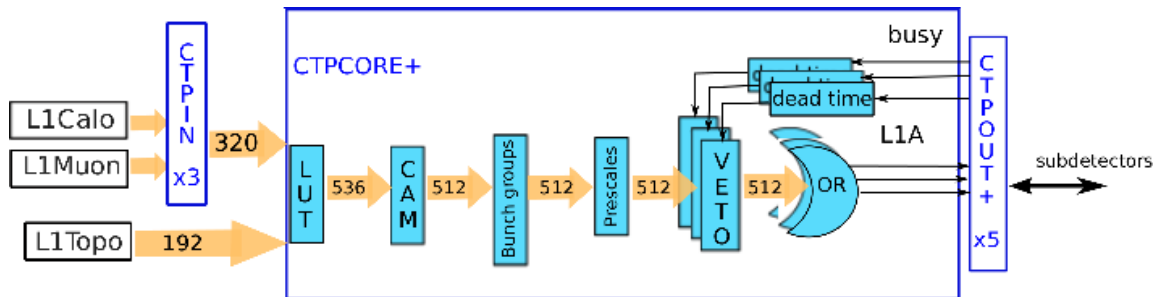
Figure 2. Schematic overview of the CTP modules described in the text with their connections.

from 256 to 512. It also supports multiple-partitions and provides better monitoring. The CTPCORE+ receives 320 trigger signals from the CTPIN modules via the PIT bus and 192 trigger signals from three direct low-latency input cables. The direct inputs are currently used by ALFA [**?**], a detector located far from interaction point, and L1Topo which both require reduced latency to compensate for additional processing time. Extra connectors for optical inputs are also available to be used with future detector upgrades.

The L1A signal is formed in the CTPCORE+ following the processing stages shown in figure **??**. The 512 trigger input signals are logically combined to form 512 trigger items using Look-Up-Tables (LUT), which select object multiplicities and can perform OR operations, and Content Addressable Memory (CAM) which performs AND and NAND operations. Each trigger item is checked to coincide with an LHC bunch-crossing according to a programmable list. Up to 16 LHC bunch patterns are configurable, such as colliding or empty bunches, instead of 8 for Run 1. Trigger items can be individually prescaled using a random prescaling algorithm. The prescale factor of each item is configurable and can be changed during data-taking to adjust the trigger rate. A trigger veto is produced by busy signals from sub-detectors or by internal dead-time generation. Five different dead-time contributions can be applied: a simple dead-time with a fixed but configurable number of vetoed bunch-crossing after a L1A, and 4 leaky bucket algorithms [**?**] which prevent sub-detector front-end buffers from becoming full. The veto gate as well as the enable/disable mask is first applied to each trigger item. The L1A signal is then formed as the OR of all enabled trigger items, and sent to sub-detectors via the CTPCORE+ modules.

The CTPCORE+ module block diagram is shown in figure **??**. The upgraded module contains two Virtex 7 FPGAs to increase computing capacity. One is dedicated to the processing of trigger signals while the second is a dedicated monitoring processor, through which CT-

Figure 3. Schematic overview of the processing stages in the CTP.

PCORE+ brings significantly improved capability in that area. Trigger rate monitoring is provided for each trigger item before prescale (TBP), after prescale (TAP) and after veto (TAV). Per-bunch monitoring of up to 192 counters (64 TBP, 64 TAP, 64 TAV) is available, as well as dead-time monitoring.

- Upgraded CTP Out Modules (CTPOUT+): distribute the L1A and timing signals to the sub-detectors, and receive their calibration and busy signals. These modules have been re-designed and re-built for Run 2 to improve diagnostics, test features, monitoring and support the additional signals needed for multiple-partition configuration. The CTPOUT+ module block diagram is shown in figure **??**.

- CTP Monitoring module (CTPMON): performs per-bunch monitoring of trigger signals on the PIT backplane. No upgrade was needed for Run 2.

- CTP Calibration module (CTPCAL): receives and transmits sub-detector calibration requests to CTPCORE+. No upgrade was needed for Run 2.

The distribution of trigger and timing signals between the CTP modules is done by a common backplane (COM bus). The backplane has been upgraded to house a fifth CTPOUT+ module and support the additional signals needed for multiple-partition configuration.

The partitioning of the CTP is useful for parallel commissioning and calibration of different sub-detectors. In this configuration, all data-taking sessions are logically separated. Each trigger item is assigned to only one data-taking session which has its own configuration of prescales and bunch-group patterns. Furthermore, each session can use a different dead-time configuration to better address sub-detector requirements. This configuration has been successfully tested hardware-wise and its commissioning in the ATLAS control room is envisaged during the 2015/2016 winter shutdown.

4. Central Trigger Processor Software

The CTP is essential for triggering events and hence a software failure could potentially stop data-taking for the whole detector. The software architecture has been completely re-designed to avoid any conflict when several data-taking sessions simultaneously use the CTP. The software is based on four main applications:
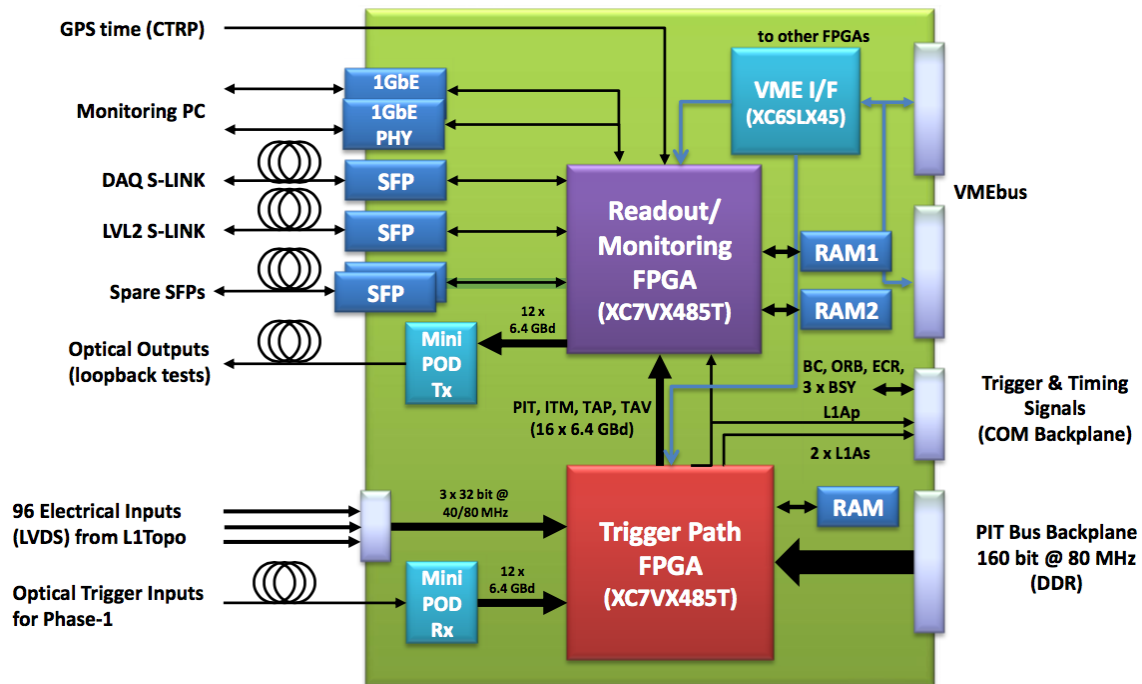
Figure 4. CTPCORE+ module block diagram.

- Configurator: in charge of the hardware configuration of the CTP which is common to all data-taking sessions. This application runs on a dedicated CTP software infrastructure and can receive commands from the different data-taking sessions. It handles the requests from concurrent sessions via a finite state machine which prevents any conflict during VME access.

- Controller: this application is the interface between each data-taking session and the CTP. It therefore belongs to the session software infrastructure. The Controller is in charge of the configuration specific to each session, such as the dead-time, and interacts with the Configurator via commands sent on the network.

- Monitoring: this application periodically reads all monitoring information from the hardware, such as the trigger rates, and publishes them on an information server for further use. This application is common to all data-taking sessions and therefore runs on the same dedicated infrastructure as the Configurator to minimise the amount of VME access.

- Monitoring clients: this is a suite of software applications which read the monitoring information published by the Monitoring, Configurator and Controller applications and re-publish them in a more human readable way, such as on a webpage. They are also in charge of archiving monitoring information.

With the exception of the Controller, all software applications can be safely restarted at any time without any interruption of data-taking. This modular approach has proven to be very stable and robust over long periods of operation.
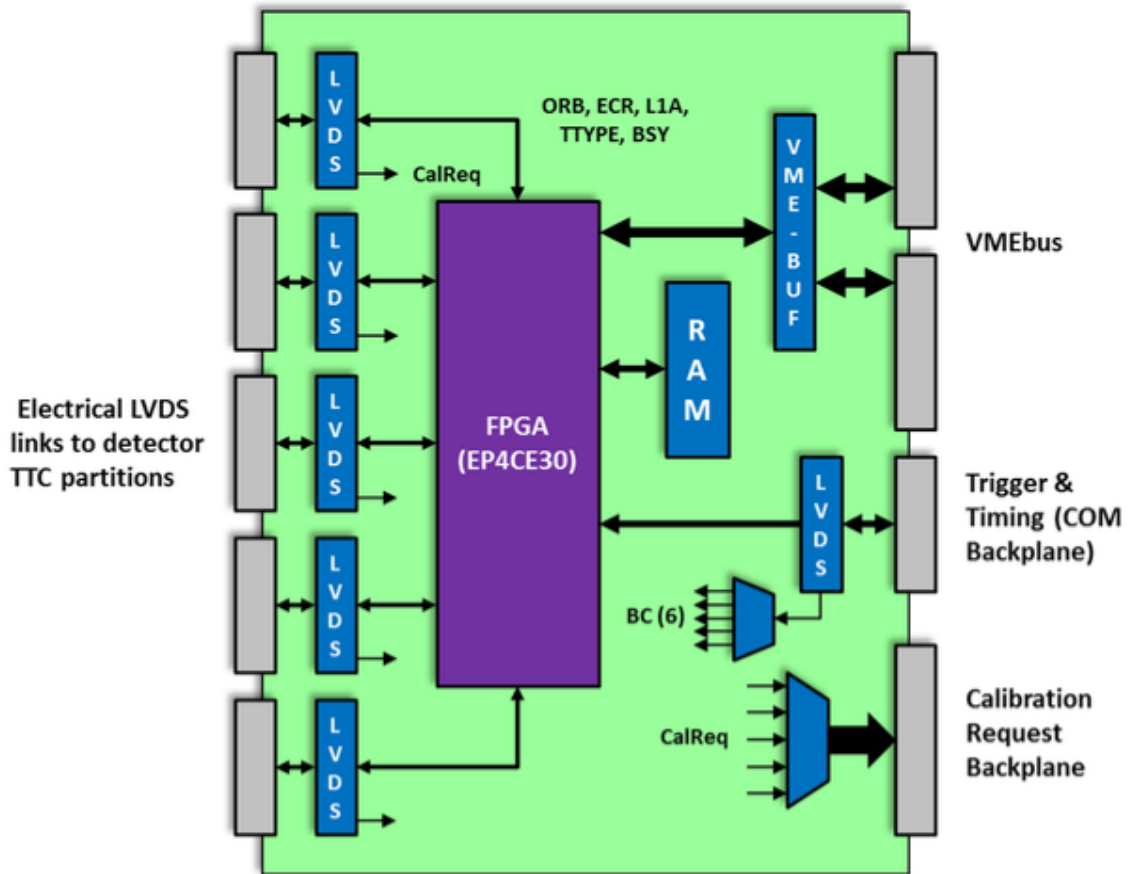
Figure 5. CTPOUT+ module block diagram.

5. Commissioning of the Central Trigger Processor

Prototypes of all upgraded CTP modules were successfully tested in the laboratory and the final boards installed in the ATLAS cavern at the end of 2014. The overall latency of the trigger path from the CTPIN to the CTPOUT+ modules with the upgraded CTP has been measured to be six bunch-crossings, corresponding to an increase of 1 bunch-crossing with respect to Run 1 but still within the latency budget of eight bunch-crossings. The phase adjustment of timing signals to synchronise all sub-detectors was implemented at the beginning of 2015. Dead-time parameters have been tuned separately for each sub-detector in order to minimise total dead-time, which is crucial to reach a L1 rate of 100 kHz. High rate tests have demonstrated that the full ATLAS data-acquisition chain can run successfully at this L1 rate. The first beam splashes were delivered by the LHC on April 5th, 2015. Dedicated triggers with a well known timing were used in the CTP to trigger these events. Almost all beam splashes were successfully selected, and the recorded events then used to compute the timing of all other triggers. The first collisions at $\sqrt{s} = 900$ GeV happened on May 15th and were also successfully triggered. The timing of the CTP was further adjusted according to information from these collision events and the CTP is now fully operational in the ongoing Run 2 data-taking.

6. Conclusion

The Central Trigger Processor is responsible for forming the Level-1 trigger decision in the ATLAS experiment. It has undergone a major upgrade during the long shutdown of 2013/2014 in order to cope with the increasing luminosity delivered by the LHC during the second run period starting in 2015. In particular, the number of trigger inputs and trigger items have been increased to have more flexibility in the selection of events relevant for physics analyses. Moreover, up to three sub-detector combinations can now simultaneously use the CTP for parallel commissioning and calibration. These improvements required the upgrade of the hardware or firmware of most CTP modules, as well as a complete re-design of the software architecture. The upgraded system has been successfully commissioned before and during the first LHC beams, and is now being used in a stable way for Run 2 data-taking.

References

[1] The ATLAS collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider.*, 2008 *JINST* **3** S08003.

[2] L. Evans and P. Bryant, *LHC Machine.*, 2008 *JINST* **3** S08001.

[3] S. Ask et al., *The ATLAS central level-1 trigger logic and TTC system*, 2008 *JINST* **3** P08002.

[4] E. Simioni, *The Topological Processor for the future ATLAS Level-1 Trigger: from design to commissioning*, ATL-COM-DAQ-2014-019, arXiv:1406.4316 (2014).

[5] S. Artz et al., *The ATLAS Level-1 Muon Topological Trigger Information for Run 2 of the LHC*, 2015 *JINST* **10** C02027.