21st International Conference on Computing in High Energy and Nuclear Physics (CHEP2015) IOP Publishing
Journal of Physics: Conference Series **664** (2015) 082013 doi:10.1088/1742-6596/664/8/082013

Operation of the Upgraded ATLAS Level-1 Central Trigger System

Julian Glatzer on behalf of the ATLAS Collaboration

CERN European Organization for Nuclear Research, 1211 Geneva 23, Switzerland

E-mail: julian.glatzer@cern.ch

Abstract. The ATLAS Level-1 Central Trigger (L1CT) system is a central part of ATLAS data-taking and has undergone a major upgrade for Run 2 of the LHC, in order to cope with the expected increase of instantaneous luminosity of a factor of two with respect to Run 1. The upgraded hardware offers more flexibility in the trigger decisions due to the factor of two increase in the number of trigger inputs and usable trigger channels. It also provides an interface to the new topological trigger system. Operationally - particularly useful for commissioning, calibration and test runs - it allows concurrent running of up to three different subdetector combinations. An overview of the operational software framework of the L1CT system with particular emphasis on the configuration, controls and monitoring aspects is given. The software framework allows a consistent configuration with respect to the ATLAS experiment and the LHC machine, upstream and downstream trigger processors, and the data acquisition system. Trigger and dead-time rates are monitored coherently at all stages of processing and are logged by the online computing system for physics analysis, data quality assurance and operational debugging. In addition, the synchronisation of trigger inputs is watched based on bunch-by-bunch trigger information. Several software tools allow for efficient display of the relevant information in the control room in a way useful for shifters and experts. The design of the framework aims at reliability, flexibility, and robustness of the system and takes into account the operational experience gained during Run 1. The Level-1 Central Trigger was successfully operated with high efficiency during the cosmic-ray, beam-splash and first Run 2 data taking with the full ATLAS detector.

1. Introduction

The Large Hadron Collider (LHC) [1] at CERN is a 27 km proton–proton (proton–Pb ion, Pb ion–Pb ion) collider with a maximum design collision energy of $\sqrt{s} = 14$ TeV. The LHC is designed for a collision rate of 40.08 MHz at a bunch spacing of approximately 25 ns and an instantaneous luminosity of approximately 10^{34} cm^{-2} s^{-1}.

The ATLAS experiment [2] is one of the two general-purpose detectors at the Large Hadron Collider. At the ATLAS detector, particle collisions are recorded using an inner detector tracking system, electromagnetic and hadronic calorimeters and a muon spectrometer. A multi-stage trigger system [3] is used to identify interesting collision events. The first stage, the Level-1 trigger, reduces the 40.08 MHz bunch-crossing rate to a rate of up to 100 kHz using information from the calorimeters, dedicated muon detectors, and the forward detectors close to the beam line. Using the information from all subdetectors in regions where high-p_{T} activity has been previously detected by the Level-1 Trigger ("Regions of Interest") or the full event information, the High-Level Trigger (HLT) [4] reduces the rate to a permanently recorded rate of up to 1 kHz.

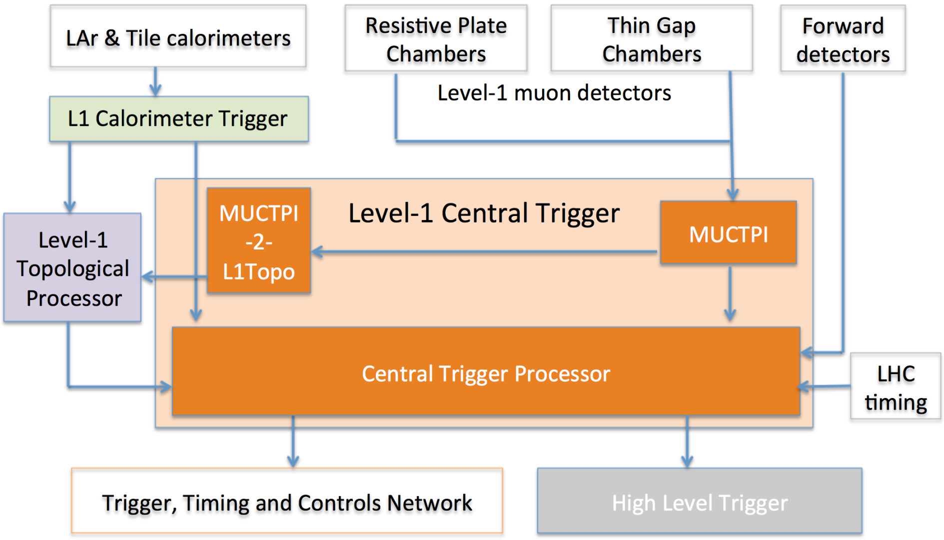

Figure 1. Schematic overview of the ATLAS Level-1 trigger system.

While the Level-1 trigger is implemented in custom-built hardware and operated synchronously to the LHC bunch crossing, the High-Level Trigger uses software-based trigger algorithms which are run on a dedicated farm of computers.

An overview of the ATLAS Level-1 trigger system is shown in figure 1. For the Level-1 trigger decision coarse granularity information from the LAr and Tile calorimeters is obtained by the Level-1 Calorimeter Trigger (L1Calo) [5]. L1Calo provides clusters above a given energy threshold, jet multiplicities for various energy thresholds, a transverse energy sum and missing transverse energy trigger, as well as dedicated tau triggers. Muon triggers are based on hits in dedicated muon detectors with a fast readout, the Resistive Plate Chambers (RPC) and Thin Gap Chambers (TGC). The information is combined in the Muon-to-CTP Interface (MUCTPI), which is discussed in more detail in section 2.2. In the Level-1 Topological Trigger (L1Topo) [6], which was installed for the LHC Run 2, trigger decisions are made based on topological information, such as angles between trigger objects or the invariant mass of two or more trigger objects. These trigger objects are provided to L1Topo by L1Calo in the case of calorimeter trigger objects or the MUCTPI in the case of muon trigger objects. The final Level-1 Accept decision is made in the Central Trigger Processor (CTP), described in section 2.1. In addition the CTP is responsible for the distribution of trigger, timing and control (TTC) signals via a dedicated network.

For the Run 2 data-taking period of the LHC the collision energy will be increased from 8 TeV to 13 TeV, leading to a 20% higher proton–proton cross section and a doubled hard interaction cross section. Furthermore, the instantaneous luminosity will be increased by at least a factor of two with respect to Run 1. In order to account for the higher interaction rates the Level-1 Accept rate was increased by 25%, the event storage rate of the ATLAS data acquisition system was doubled and more selective trigger algorithms have been made possible by the addition of the Level-1 Topological Trigger.

During the two year long shutdown in 2013/2014, the central trigger has undergone major upgrades which were necessary to prepare for the new Run 2 trigger requirements. In these proceedings an overview of the Level-1 Central Trigger and its hardware implementation is given in section 2, while in section 3 the online software infrastructure is described.

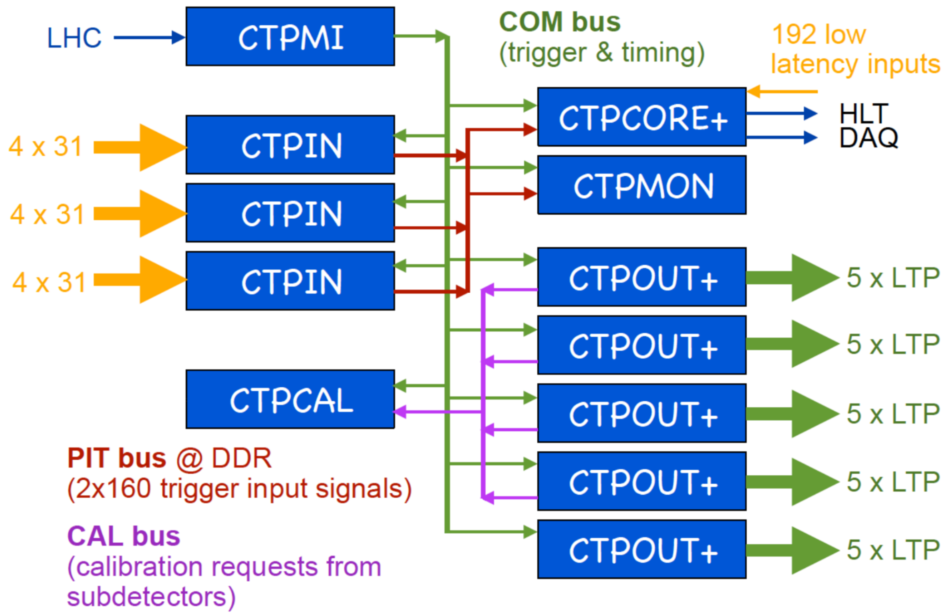

Figure 2. Schematic overview of the ATLAS Central Trigger Processor and the use hardware modules and backplane connections.

2. The Level-1 Central Trigger

The Level-1 Central Trigger consists of the CTP [7], which is discussed in section 2.1, and the MUCTPI, which is discussed in section 2.2.

2.1. Central Trigger Processor

Based on the trigger information provided by L1Calo, MUCTPI, and forward detectors, the CTP forms the Level-1 Accept decision and distributes it together with the timing information to the subdetectors via the TTC network.

In order to allow integration of the Level-1 topological processor and more flexibility in the operation and monitoring, the CTP was substantially upgraded during the long shutdown in 2013/2014. The upgraded hardware offers more flexibility in the trigger decisions due to the double number of trigger inputs and usable trigger channels. As the Level-1 topological processor adds considerable latency to the trigger path, dedicated low-latency inputs have been added to the CTP and are now used by the L1Topo and ALFA subdetectors. Operationally the ability to run up to three different trigger sessions was added for more efficient use of the time for testing and commissioning. The trigger sessions are separated logically, but share the same trigger path up to the pre-scaling operation. Only one trigger session is interfaced to the High-Level Trigger and data acquisition system and can be used for physics data taking. For Run 2 the CTP is operated with a higher output rate of 100 kHz compared to 75 kHz in Run 1 due to the increased capabilities of the High-Level Trigger.

An overview of the CTP is shown in figure 2. It is composed of custom-built VME modules, which are connected to each other using the PIT bus (trigger input signals), the COM bus (trigger and timing) and the CAL bus (calibration requests). The VME modules are the three CTP Input Modules (CTPIN), the CTP Machine Interface Module (CTPMI), the CTP Core Module (CTPCORE+), the five CTP Output Modules (CTPOUT+), the CTP Calibration Module (CTPCAL) and the CTP Monitoring Module, which are described in more detail in the following:

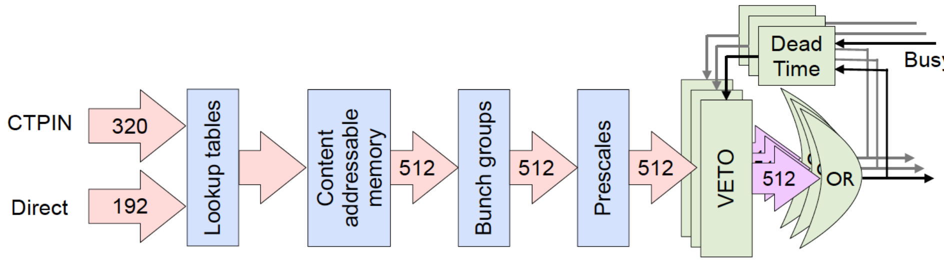

Figure 3. Schematic view of the trigger path in the Central Trigger Processor.

- The CTP Input Modules (CTPIN) receive up to 124 trigger inputs from subdetectors on four cables each. The input modules synchronise and align the trigger signals to the LHC bunch crossing and transmit selected trigger signals through the Pattern In Time (PIT) bus to the CTPCORE+ and CTPMON modules. For Run 2 a firmware upgrade enabled the use of the PIT bus at double the rate allowing 320 trigger items to be transferred to the CTPCORE+ and CTPMON modules.
- The CTP Machine Interface Module (CTPMI) receives the LHC clock and orbit signals and distributes them to the other modules.
- The CTP Core Module (CTPCORE+) is the module that obtains the Level-1 Accept decision based on the trigger menu with pre-scale factors and beam patterns (bunch groups). 320 input signals are received from the CTP Input modules via the PIT bus and 192 trigger signals from three direct low-latency input cables to the module itself. With two Virtex 7 FPGAs for the readout and monitoring of the Level-1 trigger decision and the trigger path, the computing power of the CTPCORE+ was significantly increased with respect to Run 1. For the trigger path, the number of logical trigger items was increased from 256 to 512 and the number of bunch groups was increased from 8 to 16. In addition the monitoring capabilities have been improved significantly; for Run 2 it is possible to monitor 256 instead of 12 items with bunch-by-bunch counters.
- The CTP Output Modules (CTPOUT+) distribute the Level-1 Accept and timing signals via five cables each. The CTPOUT+ modules receive calibration requests and signals from the subdetectors if the subdetector readout is busy and the ATLAS data taking should be inhibited to prevent readout data corruption. For Run 2 the output modules were upgraded to support the additional Level-1 Accept, timing and readout busy signals necessary for the parallel running of three partitions and a fifth output module was added.
- The CTP Calibration (CTPCAL) and CTP Monitoring Modules (CTPMON) handle calibration requests and monitor trigger signals bunch-by-bunch.

For Run 2 new hardware for the CTP Core Module and the CTP Output Modules was designed and commissioned. In addition the PIT bus data rate was increased and the CTP Input Module and CTP Monitoring Module firmware was adapted correspondingly.

The signal path to form the Level-1 Accept decision in the CTP is shown schematically in figure 3. It is implemented in an FPGA located on the CTPCORE+ module. 320 trigger signals are received from the CTP input modules via the PIT bus. The 192 trigger inputs from the L1Topo and ALFA subdetectors are received directly in the core module to reduce the latency of these signals. Using lookup tables and content addressable memory the multiplicities of the trigger objects are decoded and combined into 512 trigger items using logical operations. The trigger items are put into coincidence with up to 16 pre-defined patterns which allow or disallow triggers for a given LHC bunch crossing number ("bunch groups"). The physics bunch group is filled with the beam pattern for a coincidence of the trigger signals with the filled bunches directly, however also bunch groups with empty bunches for calibration triggers are used. For every trigger a configurable pre-scale factor can be applied to randomly select a fraction of events. Finally trigger items can be vetoed to protect the subdetector front-ends from overflow or data corruption. The veto is applied if one of the detector readout systems is over-occupied ("busy") or based on a configurable dead-time. The configurable dead-time uses a fixed dead-time applied after each trigger ("simple dead-time") and up to four leaky bucket algorithms which model a derandomizer with a fixed size and readout rate ("complex dead-time"). A Level-1 Accept signal is forwarded to the subdetectors to trigger the readout of the front-end electronics if any of the 512 items was triggered and not vetoed by the bunch groups, pre-scale factors and dead-time vetoes.

2.2. The Muon-to-CTP Interface

The MUCTPI receives muon candidates from 208 muon trigger sectors in 16 so-called Muon Interface Octant (MIOCT) modules. A correction is applied to avoid double-counting of muons which cross overlapping muon detectors. The total muon multiplicity for six momentum thresholds is calculated and sent to the CTP for every bunch crossing. Region of Interest information is supplied to the High-Level Trigger for events which are triggered at Level-1.

After the Phase-I upgrade, scheduled to start in 2018, the MUCTPI system will be upgraded to provide full resolution muon inputs to the Level-1 Topological Trigger. For Run 2 it was possible to profit from the flexible Run 1 hardware design and to overclock two electrical cables per MIOCT module to a rate of 320 MHz. For every octant the pseudorapidities, azimuthal angles and transverse momenta of up to two muon candidates are encoded into 8 bits each and sent to the Level-1 topological processor [8]. Altogether 256 bits of muon information per event are available to form a topological trigger decision. For this setup the firmware needed to be adapted and only an additional interface, which converts the electrical signal from the MIOCT modules to optical signals, had to be newly developed.

3. Design and Implementation of the Online Software

As the CTP is at the core of the ATLAS trigger system, a stable suite of online software is necessary. Any failure might result in an interruption of the data taking of the whole ATLAS detector, so that a time-consuming restart would need to be triggered.

During Run 1 the online software of the CTP was designed as a monolithic application using several threads. This main application provided the logic for trigger configuration, control of the board functions, and monitoring for each of the hardware boards. While the software worked very well in Run 1, the disadvantage of this design was that problems in non-essential monitoring tasks were sometimes able to stop ATLAS data taking. Furthermore the combination of configuration and control features in one application required a complicated build process and was hard to maintain.

For the LHC Run 2 the online software needs to support running three trigger sessions in parallel and avoid accessing the same hardware registers at the same time. In view of the mentioned disadvantages and the added complexity due to the multi-partition run model it was decided that for LHC Run 2 the online software design would be moved from a hardware- to a function-oriented design with the following design goals:

- Factorisation of the configuration, control and monitoring tasks.
- Access to trigger configuration limited to one application.
- Strict partitioning: every session can only access its private registers. If access to non-private registers is necessary, it is provided via a server application.

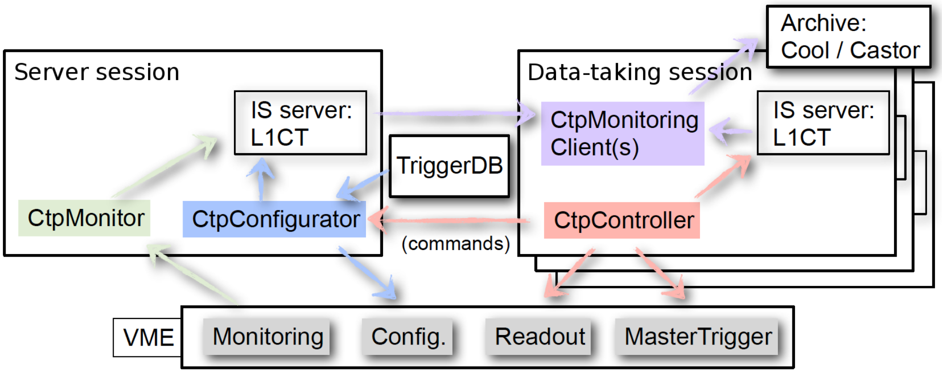

Figure 4. Schematic model of the Level-1 Central Trigger online software. It is implemented in a client–server architecture which ensures that three data-taking sessions can use the Central Trigger Processor in parallel. On the left-hand side the applications in the server session, on the right-hand side the applications in up to three data-taking sessions and on the bottom the function commands sent to the low-level software are shown.

- Ease of access to the monitoring information for remote central trigger experts. As it is often important for remote central trigger experts to get a quick and good overview of the ATLAS and central trigger running conditions, web pages allow remote access to the readout busy information, trigger rates at different steps in the Level-1 Accept formation procedure and other information.

Following the design goals, a client–server model with four main applications was developed. A schematic overview of the model is shown in figure 4. The main applications are divided into a server session that coordinates access to shared registers and other resources or functions of the hardware and up to three parallel data-taking sessions which provide the interface to the shift crew and services specific to one data-taking session. During the collision data taking only one such data-taking session is operated.

In the following, the four main applications of the client–server model are described.

The Configurator is a server application that reads the trigger menu from the database and writes its contents, such as trigger logic, items, and pre-scale factors, to the hardware and configures general shared parameters prior to starting a run. It serves as a proxy for the data-taking sessions and provides a reservation system so that shared resources can be accessed by data-taking sessions running in parallel without interference. A copy of the configuration is shared with the monitoring clients for monitoring and archiving purposes. During the run the configurator updates pre-scale factors and bunch group masks. Due to the important tasks executed by the Configurator, it is not possible to restart it during the run.

The Monitor periodically reads out and publishes raw data from the hardware. This raw data includes trigger rates, event information and the information when subdetector readouts are busy. The Monitor can be restarted during a data-taking run.

The Controller is the interface to the CTP hardware used by the shift crew in the ATLAS control room. It is embedded in the ATLAS TDAQ run control software framework and with it the trigger can be held and resumed, the pre-scale factors can be changed before and during a run and the CTP readout and busy propagation is configured. In addition the Controller configures and defines periodic signals such as luminosity blocks (the basic time unit for storing luminosity information including events with similar conditions) and event counter resets. Due to the important tasks executed by the Controller, it is not possible to restart it during the run.

Monitoring Clients are small single-purpose applications that republish, combine and archive monitoring information together with the configuration data for display in a human-readable and persistent way. While the Monitor's task is to publish data in a raw format, monitoring clients may perform complex data analysis and error detection. Monitoring clients are restartable during a data-taking run, so that less stringent requirements on the stability can be posed.

The CTP software framework is on the one hand characterised by software components whose stable operation is essential for the ATLAS data taking and whose development is conservative, with only absolutely necessary updates applied during a running period. On the other hand it contains monitoring software where flexibility is needed to react to current data-taking conditions. The modular design of the CTP software framework allows for fast development cycles for monitoring software, while guaranteeing long term stable operation for the configuration and control software. The compartmentalisation of tasks and hardware access allows for writing software that guarantees safe and stable operation in the complex multi-partition environment and easy maintainability. In addition it enables the use of load balancing as it is not necessary to run all applications on hosts that provide access to the hardware VME interface.

4. Conclusions

In view of the new trigger requirements needed for LHC Run 2 conditions the ATLAS central trigger system was upgraded. The upgraded central trigger system allows for a higher number of input signals from subdetector systems, features low-latency inputs and provides more trigger items. The full system was already successfully operated in many hours of testing with and without beam. The first events from proton–proton and proton–collimator collisions ("beam splashes") have been successfully recorded with high efficiency.

References

[1] Evans L and Bryant P 2008 *JINST* **3** S08001
[2] ATLAS Collaboration 2008 *JINST* **3** S08003
[3] Schoerner-Sadenius T 2003 *Mod.Phys.Lett.* **A18** 2149–2168 (*Preprint* `hep-ex/0307078`)
[4] Jenni P, Nessi M, Nordberg M and Smith K (ATLAS Collaboration) 2003 *CERN-LHCC-2003-022* URL `http://cds.cern.ch/record/616089`
[5] Achenbach R *et al.* 2008 *JINST* **3** P03001
[6] Simioni E 2014 (*Preprint* `1406.4316`)
[7] Ask S *et al.* 2008 *JINST* **3** P08002
[8] Artz S *et al.* 2015 *JINST* **10** C02027