The ATLAS Central Level-1 Trigger Logic and TTC System

Ask S.^{a,j}, Berge D.^a, Borrego-Amaral P.^a, Caracinha D.^e, Ellis N.^a, Farthouat P.^a,
Gällnö P.^a, Haas S.^a, Haller J.^{d,b}, Klover P.^a, Krasznahorkay A.^{a,c}, Messina A.^a,
Ohm C.^{a,i}, Pauly T.^a, Perantoni M.^g, Pessoa Lima Junior H.^g, Schuler G.^a, Sherman
D.^a, Spiwoks R.^a, Wengler T.^f, de Seixas J.M.^g, Torga Teixeira R.^{a,h}

^a Physics Department, CERN, CH-1211 Geneva 23, Switzerland

^b DESY, Notkestrasse 85, D-22603, Hamburg, Germany

^c ATOMKI, PO Box 51, HU-4010 Debrecen, Hungary

^d Universität Hamburg, Institut für Experimentalphysik, Luruper Chaussee 149, DE-22761, Hamburg, Germany

^e Universidade de Lisboa, Departamento de Física, Campo Grande ED C8, PT-1749-016, Lisbon, Portugal

^f School of Physics and Astronomy, University of Manchester, Manchester M13 9PL, UK

^g Universidade Federal do Rio de Janeiro, Instituto de Física, Caixa Postal 68528, Ilha do Fundão, BR-21945-970 Rio de Janeiro, Brazil

^h Now at Universidad Autónoma de Madrid, Carretera de Colmenar Canto Blanco, ES-28049 Madrid, Spain

ⁱ Now at Stockholm University, Vanadisvägen 9, SE-108-91, Stockholm, Sweden

^j Now at University of Manchester, School of Physics and Astronomy, M13 9PL Manchester, United Kingdom

E-mail: Philippe.Farthouat@cern.ch

ABSTRACT: The ATLAS central level-1 trigger logic consists in the Central Trigger Processor and the interface to the detector-specific muon level-1 trigger electronics. It is responsible for forming a level-1 trigger in the ATLAS experiment. The distribution of the timing, trigger and control information from the central trigger processor to the readout electronics of the ATLAS subdetectors is done with the TTC system. Both systems are presented.

KEYWORDS: ATLAS, LHC, Level-1 trigger, first-level trigger.

Contents

1. The Level-1 Trigger System	2
2. The MUON-TO-CTP Interface	3
2.1 MUCTPI Architecture	5
2.1.1 MIOCT Module	5
2.1.2 MIBAK Backplane	5
2.1.3 MICTP Module	6
2.1.4 MIROD Module	6
2.2 MUCTPI Implementation	6
2.2.1 Octant module (MIOCT)	6
2.2.2 The MICTP and MIROD Modules	10
2.3 The MUCTPI latency	11
3. The Central Trigger Processor	12
3.1 The Functionality of the CTP	12
3.1.1 The Trigger Formation	12
3.1.2 Additional Functionality	12
3.1.3 The Constraints	13
3.2 The Implementation of the CTP	14
3.2.1 Overview	14
3.2.2 The CTP_MI (Machine Interface) Module	15
3.2.3 The CTP_IN Module	15
3.2.4 The CTP_CORE Module	16
3.2.5 The CTP_CAL Module	18
3.2.6 The CTP_MON Module	18
3.2.7 The CTP_OUT Module	19
3.3 The CTP latency	19
4. Firmware management	19
5. The Timing, Trigger and Control System	20
5.1 Timing Signals from the LHC and synchronisation	20
5.2 CTP Links	22
5.3 TTC partition	22
5.3.1 TTC partition root	23
5.3.2 TTC optical tree and receivers	25
5.3.3 Timing adjustment	26
6. Modules fabrication	26
7. Performance of the system	26
8. Conclusion	28

1. The Level-1 Trigger System

The ATLAS Level-1 trigger [1] is based on information from the calorimeter and muon trigger processors. The trigger information consists of multiplicities for candidate electrons/photons, taus/hadrons, jets and muons, and of flags for total scalar transverse energy, total missing transverse energy, and total scalar jet transverse energy. The overview of the ATLAS Level-1 trigger is shown in Fig. 1.

Figure 1: The Overview of the ATLAS Level-1 Trigger

The calorimeter [2] and muon trigger [3], [4] processors provide trigger information to the Central Trigger Processor (CTP). The CTP forms the Level-1 Accept (L1A) decision and fans it out to the Timing, Trigger and Control (TTC) partitions of the experiment. Each TTC partition contains one Local Trigger Processor (LTP) [5], a TTC system proper [6], and a tree of ROD_BUSY modules [7]. The CTP provides trigger summary information to the Level-2 trigger system and to the data acquisition system [8]. It is configured, controlled and monitored by the ATLAS Online system [9]. The LTP provides the facility to run with the trigger and timing signals from the CTP, but it can also generate the signals locally. In order to allow several combinations of TTC partitions to run independently of the CTP, an LTP interface module [5] is used. It allows performing, for instance, concurrent calibration runs of the calorimeter detectors with the calorimeter trigger, and of the muon detectors with the muon trigger, without requiring any re-cabling.

The Central Trigger Logic described in this paper consists of the CTP, plus the Muon-to-CTP interface (MUCTPI) that sits between the detector-specific logic associated with the barrel

and endcap muon trigger systems and the CTP. Also described is the Timing, Trigger and Control system that passes trigger and timing signals to the detector front-end systems.

All the elements of the CTL are located in the ATLAS control room (USA15), in standard racks, and are subject to standard running conditions (i.e. they are not subject to radiations or to magnetic field). Most of them are implemented in 9U or 6U VME boards.

An important parameter for the CTL is the allowed latency of the system. Table 1 gives the latency envelope for each part of the level-1 trigger system, in time unit (ns) and in number of bunch crossing periods (BC). The overall allowed latency for the CTL (i.e. MUCTPI, CTP and TTC electronics) is 400ns maximum. As reported later in Sections 2.3 and 3.3, the measured values for the MUCTPI and CTP match the requirements.

Level-1 trigger system component	Latency [ns]	Latency [BC]
Detector specific MUON trigger	1293	52
CALORIMETER trigger	1465	59
MUCTPI	200	8
CTP	100	4
TTC electronics (LTP, TTCvi, TTCex) excluding cables and fibres	100	4

Table 1: Latency envelopes for each part of the level-1 trigger system.

2. The MUON-TO-CTP Interface

The ATLAS Level-1 trigger uses multiplicities of tracks found in the dedicated muon trigger detectors when forming the final trigger decision. An overview of the muon trigger system is shown in Fig. 2.

Figure 2: Overview of the ATLAS Level-1 muon trigger

The muon trigger detectors are resistive plate chambers (RPC) in the barrel region, and thin-gap chambers (TGC) in the end-cap and forward regions of ATLAS. Coincidences of hits in different detector stations are used to identify muon candidates. The muon trigger electronics

considers the transverse-momentum (p_T) of the muon candidates and classifies them according to six programmable p_T thresholds.

The muon trigger detectors are grouped into sectors, 64 for the barrel, 96 for the end-cap and 48 for the forward region (208 sectors in total). Each sector can provide information on up to two muon candidates (priority is given to the highest- p_T candidates). For each bunch crossing, the detector-specific sector logic sends to the MUCTPI information about the position and p_T (thresholds passed) of the muon candidates. A 32-bit word per sector and bunch crossing is received by the MUCTPI, representing more than 266Gbits/s in total.

The MUCTPI combines the information from all of the muon trigger sectors and calculates the total multiplicity for each of the six p_T thresholds, resolving possible double counting of muon candidate tracks that traverse more than one detector region. It forwards these multiplicity values to the CTP for use in forming the final Level-1 decision.

In addition to providing inputs to the CTP, the MUCTPI provides data to the Level-2 trigger and to the DAQ system for events selected at Level-1. The DAQ system receives a formatted copy of the information on candidate muon tracks, as well as the computed candidate multiplicity. The Level-2 trigger system receives a subset of the candidates (the 16 highest p_T candidates), ordered according to decreasing p_T . This information is used to define regions of interest (RoIs) that drive the Level-2 muon-trigger processing. While RoIs corresponding to a single bunch crossing are transmitted, one can read out several consecutive bunch crossings of data, allowing one to see activity in earlier and/or later bunch crossings compared to the triggered one.

The MUCTPI has to avoid double-counting muons which traverse more than one set of trigger chambers [10], since this would lead to an unacceptably high rate of fake di-muon triggers. This effect is particularly marked for low- p_T muons since their tracks are considerably deflected in the magnetic field of the toroid magnets. For example, muons may be seen both in the barrel and end-cap detectors.

Fig. 3 illustrates the overlap between the barrel and end-cap muon trigger chambers. Tracks for low and high- p_T muons are shown in the barrel and end-cap chambers. In addition an example of a muon track that would be detected in both barrel and end-cap systems is shown. The MUCTPI should detect this case and avoid double-counting the muon candidate in the multiplicity sum. It is worth noting that the magnetic field is non-uniform and can result in complex muon trajectories, particularly in the barrel–endcap transition region.

Figure 3: Barrel and end-cap muon trigger chambers

2.1 MUCTPI Architecture

Fig. 4 shows the architecture of the MUCTPI. The system consists of a 9U VME chassis with a special backplane and 18 custom-designed modules. It includes 16 MIOCT modules, each of which receives the data from one octant in azimuth angle, ϕ , and one half of the detector in pseudo-rapidity, η . The MIBAK active backplane sums the multiplicities of all MIOCT modules, provides readout data transfer, and distributes timing and trigger signals to all modules. The MICTP module receives the timing and trigger signals from the CTP (which are passed over the backplane to the other modules), and also forwards the multiplicities, calculated on the backplane for each of the p_T thresholds, to the CTP. The MIROD module collects information from the MICTP and the MIOCT modules, and sends the data after formatting to the Level-2 trigger and to the DAQ system.

Figure 4: MUCTPI architecture

The different components of the MUCTPI system are introduced below and described in more detail in the following sections.

2.1.1 MIOCT Module

Each of the 16 MIOCT modules receives information on the muon candidates from 13 trigger sectors and forms the local muon candidate multiplicities for each of the six p_T thresholds. It also resolves overlaps between the sectors of an octant in order to avoid double-counting of muon candidates. A more detailed description of the design and implementation of the MIOCT module is given in Section 2.2.1.

2.1.2 MIBAK Backplane

The MIBAK is a custom-built active backplane which receives the muon candidate multiplicities from the 16 MIOCT modules and forms the total-multiplicity sums. The multiplicity summing is performed by a binary adder tree implemented using Altera MAX7000 CPLDs for low latency. In addition, the backplane implements a bus for the transfer of readout

data from the MIOCT and the MICTP modules to the MIROD. This shared readout bus uses a token-passing protocol for arbitration and is based on Bus LVDS (BLVDS) signalling. Finally, the MIBAK distributes the trigger and timing signals from the MICTP module with low-skew to all the other modules in the crate. The MIBAK is mounted in the J3 position of the 9U VME crate.

2.1.3 MICTP Module

The MICTP module receives the total multiplicity sums for the six p_T thresholds from the MIBAK backplane and sends them to the CTP. It also writes the multiplicities into a pipeline for read-out by the MIROD module on reception of a Level-1 Accept signal from the CTP. In addition, the MICTP receives the timing and triggers signals, such as the 40 MHz bunch clock, the Level-1 Accept, the LHC orbit signal and the event counter reset, from the CTP and distributes them through the MIBAK backplane to the other modules in the crate. The MICTP provides features for monitoring the multiplicity sums over the VME bus. A more detailed description of the design and implementation of the MICTP module is given in Section 2.2.2.

2.1.4 MIROD Module

The MIROD module is the Read-Out Driver (ROD) of the MUCTPI. For every Level-1 Accept it collects information on the muon candidates from the 16 MIOCT modules and the multiplicities from the MICTP module and sends them after formatting to the data acquisition and Level-2 trigger systems. The system can be programmed to read out a time-frame of several time slices around the bunch crossing that caused the trigger. The standard S-LINK [11] optical-link mezzanine cards are used for this purpose. The MIROD also allows reading out the selected event-fragment data via the VME bus for monitoring purposes. A more detailed description of the design and implementation of the MIROD module is given in Section 2.2.2.

2.2 MUCTPI Implementation

2.2.1 Octant module (MIOCT)

The design and implementation of the MIOCT module is presented in the following sections.

MIOCT Architecture

The MIOCT is implemented as a 9Ux400 mm VME64x module. The 13 sector-logic inputs use 32-bit parallel LVDS signalling [12] at 40 MHz. Serial transmission was excluded because of the latency penalty of ~ 3 BC due to serialization and de-serialization. Using 2x68-pin high-density dual-stacked VHDCI connectors and low-skew SCSI-3 [13] twisted-pair cable, it is possible to fit all 13 sector-logic inputs on the front-panel of the module.

Fig. 5 shows a block diagram of the MIOCT module.

Figure 5: MIOCT block diagram

The main functionality of the MIOCT is implemented in one Altera Stratix II FPGA [14]. This chip features sufficient memory, logic and I/O resources to allow integration of the main MIOCT processing into a single device. There is an additional small FPGA for the VMEbus interface. The implementation is based on the VMEbus interface core [15] which was developed for the modules of the CTP.

Trigger Path

The signals received from the detector-specific sector logic are resynchronized with the system clock received from the MICTP. The incoming sector words are aligned in time to compensate for different latencies introduced by the RPC and TGC detector-specific electronics and cable lengths. The MIOCT checks the alignment to be sure that the data word of each sector corresponds to the same bunch crossing; this is achieved using a three-bit BCID field in the data received from the sector logic.

The multiplicity-summing logic counts the number of muon candidates for each of the 6 p_T thresholds, taking into account possible overlaps between sectors. The overlap-handling logic suppresses one of the muon candidates if there is a candidate in the corresponding overlap zone of an adjacent sector. The overlap candidates remain in the readout data block allowing for later checks. More details on the overlap handling are given later. The total number of muon candidates in each octant is encoded into six 3-bit words and sent to the MIBAK backplane for overall summing.

The internal logic is operated at four times the bunch clock (~ 160 MHz) in order to minimize the latency while maintaining a pipelined architecture. The latency of the trigger path is 2 BC (i.e. 50 ns).

Readout Path

All the data received from the muon sectors are stored in pipeline memories for the duration of the latency of the Level-1 trigger. In case of a Level-1 Accept, data of all input channels are read out from the pipelines and written into derandomizer FIFOs. The readout window is programmable up to ± 2 bunch crossings around the trigger. The data in the derandomizer FIFOs are zero-suppressed in order to reduce the data rate on the backplane, formatted and buffered

until they are read out via the MIBAK backplane. The data sent by the MIOCT module can also be read out for monitoring purposes using the VME bus.

Muon Candidate Overlap Handling

Each MIOCT module receives data from four barrel, six end-cap and three forward sectors of the muon trigger. Fig. 6 illustrates the “overlap regions” between these sectors, i.e. regions in which a muon may be seen also in an adjacent sector (taking into account both the detector geometry and the deflection of triggering muons in the magnetic field of the muon spectrometer).

The following types of overlaps between adjacent sectors can be handled by the MIOCT:

- Overlap between barrel sectors
- Overlap between barrel and endcap sectors
- Overlap between endcap sectors
- Overlap between forward sectors

Note that overlaps for chamber regions within sectors are handled by the detector-specific logic. There is no significant overlap between adjacent octants.

Figure 6: Overlap regions

Fig. 7 shows a block diagram of the implementation of the overlap-handling scheme. In a first stage, overlaps between muon candidates from different sectors are detected and flagged based on their location or sub-sector address and, in the case of the barrel/end-cap overlaps, their p_T threshold and charge sign as well. There is one overlap-detection unit for each of the overlap regions listed above. A second stage of the processing uses the overlap flags together with a comparison of the p_T values to determine which of the two candidates in overlap should be suppressed, priority being given to the higher- p_T candidate. Finally, the multiplicity is calculated from the p_T values of a maximum of 26 candidates in the octant, ignoring the candidates that are flagged to be suppressed.

Figure 7: Overlap handling architecture

All the functional blocks of the overlap handling are implemented using programmable look-up tables (LUT). The overlap handling policy can therefore easily be changed by reloading these LUT through the VMEbus interface. The contents of the LUTs can be determined from simulations of muons in the ATLAS detector and, in the future, using muons reconstructed in real data.

Snapshot/Test Memory

The MIOCT also features a snapshot and test data memory. The memory can be used to store the data from all 13 input sectors, as well as the calculated candidate multiplicity and the candidate suppression flags. This is useful during the timing-in of the system, as well as for diagnostics and monitoring purposes. The depth of the snapshot memory is 128k bunch crossings per sector, which corresponds to ~36 LHC turns. For module test purposes, the memory can also be used to replay test data. (Since the LVDS buffers on the sector-logic links are bi-directional this even allows performing a full functional test of another MIOCT module including the cable connections.)

The snapshot/test memory is implemented using two 1Mx36-bit Quad-Data Rate (QDR) SRAM devices. QDR SRAM features independent double data-rate (DDR) read and write ports. The memory devices on the MIOCT are clocked at 160 MHz, resulting in a bandwidth of ~23 Gbit/sec for read and write access. The memory interface is implemented using an IP core from Altera. The routing between the FPGA and the QDR SRAM is critical and requires length matching and 50Ω controlled impedance. A signal integrity analysis tool (Cadence SpectraQuest) was used to validate the termination scheme and the PCB layout.

Fig. 8 shows a picture of the MIOCT module.

Figure 8: Picture of a MIOCT module

2.2.2 The MICTP and MIROD Modules

Fig. 9 shows a block diagram of the MIROD/CTP module which uses a single hardware board design, with two sets of firmware for the different functions as discussed below.

Figure 9: MIROD/CTP module block diagram

The design is based on a single Altera Stratix-II FPGA, which has ample memory, logic and I/O resources, and therefore enables integration of the complete logic into a single device. Since all the necessary interfaces have been foreseen and are connected to the FPGA, the same PCB with different firmware programming can be used to perform the functions of the MICTP or of the MIROD. This requires mounting the front-panel connectors for the trigger and timing signals and the multiplicity for the MICTP, instead of the S-LINK mezzanine cards for the MIROD. Fig. 10 shows a picture of the MIROD/CTP module with the S-LINK mezzanine cards mounted for the MIROD function.

Figure 10: MIROD/CTP module picture

MIROD Functionality

When used to implement the MIROD, the FPGA receives the readout data via the BLVDS buffers from the MIBAK backplane and sends them, after processing, to the Level-2 and DAQ S-LINK mezzanine cards. It also receives the timing signals from the backplane.

MICTP Functionality

When using the board to implement the MICTP, the FPGA receives the trigger and timing signals from the front-panel through the appropriate level translators, and distributes them after reshaping to the other modules in the crate via the backplane. The FPGA also receives the outputs of the multiplicity adder trees on the backplane, latches them and sends them to the CTP via LVDS buffers. In case of an L1A the multiplicities are also written into a pipeline memory in order to be read out via the MIBAK bus.

The module features a programmable delay chip (DELAY25 [16]) which allows clocks with different phases to be generated in order to latch signals at the optimum time. The design also includes a TDC chip [17] to monitor the phase of the external timing signals and of the multiplicity words; this is useful during the timing-in of the system.

Common Features

A VME interface is required for both applications and is based on the same design as the MIOCT, however without using a separate small FPGA. The module also features a snapshot and test data memory which can be used to capture data from the MIBAK backplane for monitoring and diagnostics, or to replay data for module test purposes. The memory is implemented as a single 1Mx36-bit QDR SRAM device.

2.3 The MUCTPI latency

The measured latency of each stage of the MUCTPI is given in Table 2.

MUCTPI stage	Remark	Latency
MIOCT	LVDS input to multiplicity out on MIBACK	50 ns
MIBAK	MIOCT output to MICTP input, including the adder tree	22 ns
MICTP	MIBACK adder tree output to LVDS output to CTP	11 ns
Overall MUCTPI latency	83 ns	

Table 2: MUCTPI latency

The overall latency is 83ns, i.e. 3.3 bunch crossing periods. It is much lower than the allowed 200ns latency. This is mainly due to the fact that the MIOCT logic has been fully integrated in a large FPGA which is internally running at four times the LHC lock speed (160 MHz).

3. The Central Trigger Processor

The Central Trigger Processor (CTP) combines information from calorimeter and muon trigger processors and makes the final Level-1 Accept (L1A) decision on the basis of lists of selection criteria (trigger menus). In addition to the event-selection decision, the CTP provides trigger summary information to the Level-2 trigger and the data-acquisition system. It further provides accumulated and bunch-by-bunch scaler data for monitoring of the trigger, detector and beam conditions.

3.1 The Functionality of the CTP

3.1.1 The Trigger Formation

The CTP receives trigger information from the calorimeter and muon trigger processors. This information consists of multiplicities for candidate electrons/photons, taus/hadrons, jets and muons, and of flags for total scalar transverse energy, total missing transverse energy, and total jet scalar transverse energy. In addition, it receives inputs from various other sources, including luminosity detectors and beam pick-ups. The CTP also provides internal triggers from two random generators, eight bunch-crossing groups, and two pre-scaled clocks.

All the calorimeter and muon triggers are programmable, and several p_T thresholds are used concurrently for each type of trigger object. Up to 160 input bits can be taken into account by the CTP at any given time. The total number of input bits available to the CTP can be higher than this because selection of which will be used at any given time is performed by the modules at the input to the CTP.

The CTP generates a L1A signal derived from the trigger inputs according to the Level-1 trigger menu. The menu consists of up to 256 trigger items, each of which is a combination of one or more conditions on trigger inputs. For example, if “EM10” represents the multiplicity for the trigger input for electrons/photons with a transverse energy threshold of 10 GeV, “1EM10” symbolizes the condition of there being at least one electron/photon above that threshold. Several such conditions can be combined to make a trigger item. Each trigger item also has a mask, a priority and a pre-scaling factor. The priority is used in the algorithm for the preventive dead-time generated by the CTP, see Section 3.1.2. High-priority items see as little deadtime as possible, while low-priority items see a comparatively higher deadtime. The L1A is the logical OR of all trigger items. An example of an excerpt of a Level-1 trigger menu is shown in Table 3.

1MU6	mask = ON, priority = LOW, pre-scaling = 1000
2MU6	mask = ON, priority = HIGH, pre-scaling = 1
1EM10 AND XE20	mask = ON, priority = LOW, pre-scaling = 1

Table 3: An Example of an excerpt of a Level-1 Trigger Menu

3.1.2 Additional Functionality

The CTP not only generates the L1A, but it also provides with each L1A signal an 8-bit trigger type word which is sent promptly via the TTC system to the detector front-end electronics and

which summarises the type of trigger; this can be used to steer event data processing, e.g. in the detector Read-Out Driver (ROD) modules.

For each accepted event, the CTP sends information to the Region-of-Interest Builder (RoIB) of the Level-2 trigger system for guidance of the Level-2 trigger algorithms. It sends more detailed information to a Read-Out System (ROS) that can be read out by the data acquisition system for events retained after the Level-2 selection. This information is a superset of the RoI information and can contain data for several bunches before and after the one that triggered; this can be used for debugging and monitoring purposes. The data sent to the RoIB and the ROS provide much more information than the trigger-type word, and allows detailed understanding of why the event was selected.

The CTP provides an absolute GPS-based UTC time stamp that is included in the trigger information that is sent to the RoIB and ROS. This allows one to correlate the event to observations from other sources, e.g. data taken in other particle or astrophysics experiments (see also Section 3.2.4).

In order to prevent front-end buffers from becoming full and to simplify the readout electronics, the CTP generates deadtime with three mechanisms: a fixed, but programmable number of untriggered bunch crossings after each L1A, and two leaky-bucket algorithms which limit the number of L1As to be generated in a given period of time (the number of triggers and the time period are programmable parameters). The two leaky-bucket algorithms are used to define two priorities of trigger items as discussed above. The CTP also introduces deadtime when a BUSY signal is asserted by one of the multiple readout modules of ATLAS.

The leaky-bucket algorithm is the following: a bucket, which leaks at a constant rate, R , has its contents increased by a given value, X , every time a L1A occurs. When the bucket is full deadtime is introduced. The values of the programmable parameters X and R directly determine the values of the two parameters defining the deadtime, i.e. the time period and the maximum number of L1A issued during this time period.

The CTP also provides mechanisms for generating signals used in event fragment synchronisation. The event fragments are identified at all levels of the readout with a L1ID (event number) and a BCID (bunch crossing number within one LHC turn of 3564 BCs). The BCID values are reset by a bunch counter reset signal issued by the TTC system (see Section 4) while the L1ID are periodically reset by an event-counter reset (ECR) signal issued by the CTP. The ECR can be issued periodically (with a programmable period) or on software request.

The CTP provides extensive facilities for monitoring: snapshots of incoming data; scalars for bunch-by-bunch monitoring of the inputs (see Section 3.2.6); and scalars of trigger inputs and trigger items, before and after pre-scaling, integrated over all bunches. It also provides the capability of issuing calibration triggers following several mechanisms (see Section 3.2.5).

3.1.3 The Constraints

The formation of the trigger is required to be performed within four bunch crossings from the last input arriving at the CTP until the output of the L1A from the CTP; this corresponds to a latency of 100 ns.

The Level-1 trigger menu used for the trigger formation is expected to change frequently depending on the physics, beam and detector conditions. High flexibility therefore has to be provided for the trigger formation.

3.2 The Implementation of the CTP

3.2.1 Overview

The CTP consists of the following modules:

- Up to three input (CTP_IN) modules
- A core (CTP_CORE) module for trigger formation, read-out, monitoring and time stamping
- A bunch-by-bunch monitoring (CTP_MON) module
- Up to four output (CTP_OUT) modules connecting to the LTPs in the TTC partitions of the detectors
- A machine interface (CTP_MI) module for timing signals
- A calibration (CTP_CAL) module for receiving calibration requests from the detectors

The CTP modules are housed in a 9U VME64x crate. In addition to the standard VMEbus, the CTP modules use dedicated buses for transmission of synchronized and aligned trigger inputs (PITbus, PIT = pattern in time), for the common timing and trigger signals (COMbus), and for the calibration requests from the detectors (CALbus). An overview of the implementation of the CTP is shown in Fig. 11.

Figure 11: The CTP design with its modules, backplanes, and signals

3.2.2 The CTP_MI (Machine Interface) Module

The CTP_MI is the timing module of the CTP. It receives the timing signals (the 40 MHz bunch crossing clock BC, the ORBIT signal defining the 3564-BC LHC turn) from the LHC machine or can generate them locally when the machine is not running. It generates the Event Counter Reset (ECR) signal on a regular basis or under software control in order to maintain the event synchronisation in all subdetectors. Both the timing and control signals are sent to the COMbus, to be distributed to the other modules of the CTP. The CTP_MI also receives, monitors and masks the BUSY signals that may be received from the subdetectors or generated internally. An overall BUSY signal is sent to the CTP_CORE to generate deadtime.

All the signals are synchronised with the bunch crossing clock (BC). In order to avoid meta-stability, the phase of all incoming signals with respect to BC is measured in order to decide to strobe them either with the rising edge or the falling edge of BC.

A block diagram of the CTP_MI module is shown in Fig. 12. It is based on FPGAs (Altera Cyclone) and uses the CERN high-performance TDC (HPTDC) [17] for phase measurement. In the shaping stage the width of the signals are adjusted.

Figure 12: Block diagram of the CTP_MI module

3.2.3 The CTP_IN Module

A block diagram of the trigger path of the CTP input module is shown in Fig. 13.

Figure 13: Trigger path diagram of the CTP_IN module

The CTP_IN module receives trigger inputs from the Level-1 calorimeter or muon trigger processors and other sources. It measures the phase of the trigger inputs using a CERN High-Performance TDC (HPTDC) [17]. The clock of the CTP_IN module is adjusted using a CERN DELAY25 chip [16] to sample the input data with the optimal phase.

The inputs are aligned with respect to each other in time using pipelines in an FPGA to delay inputs that arrive comparatively early to match the one that arrives last. The CTP_IN module selects and routes trigger inputs to be sent to the PITbus using a switching matrix implemented in a CPLD. The synchronized trigger inputs can be stored in diagnostic memory (dual-port RAM) for debugging and monitoring purposes.

Monitoring scalars are included in the CTP_IN. These counters can be incremented either when a single input is present or when a given pattern is present in a group of inputs (used in the case the inputs are encoded multiplicity values for instance).

3.2.4 The CTP_CORE Module

A block diagram of the trigger path of the CTP_CORE module is shown in Fig. 14.

Figure 14: Trigger path of the CTP_CORE module

The CTP_CORE module receives the pattern-in-time (PIT) data from the PITbus and combines them using Look-up Tables (LUT) and Content-Addressable Memories (CAM) to form up to 256 trigger items. Each CAM contains a 256-bit word and is ternary, i.e. allows bitwise matching of “0”, “1” or “don’t care”. The 256 trigger items are prescaled using 24-bit prescalers (PSC). The items can be masked (MSK). The mask is the logical OR of a programmable general mask, the local deadtime generation and the general BUSY of the experiment (readout deadtime). After this mask stage, the items are Ored to generate the Level-1 Accept (L1A). When a L1A is issued, a local deadtime time of n BC ($n=5$ in normal running) is generated. The trigger-type word is formed in encoding the trigger items after masking (TYP).

The CTP_CORE module sends information to the RoIB of the Level-2 trigger system and to the ROS of the data acquisition system. The readout and monitoring of the CTP_CORE module is shown in Fig. 15.

Figure 15: Readout and monitoring of the CTP_CORE module

For every L1A, data in a programmable window around the triggered bunch crossing are copied from pipeline memories into FIFOs embedded in two FPGAs. The data captured include:

- 160 PIT signals and 12 internal triggers (8 bunch crossing groups, 2 random triggers and 2 pre-scaled bunch crossings)
- 256 trigger items before prescaling (TBP)
- 256 trigger items after prescaling (TAP)
- 256 trigger items after veto (TAV)
- 64-bit UTC time-stamp

The time stamping in the CTP_CORE module is based on the LHC General Machine Timing (GMT) system [18] which provides a precise GPS-based UTC time reference. The GMT signal is received using a CTRP card on the commercial single-board computer in the CTP crate. The CTRP card produces a 40.000 MHz and a 1 Hz PPS signal (“pulse-per-second”) which are sent to a daughter card on the CTP_CORE module. The PPS signal is phase locked to the GPS absolute time with a precision of 20 ns and is initially obtained from a commercial GPS satellite receiver handled by the machine team. The daughter card contains a counter which counts this signal, giving an absolute time in seconds. The 40 MHz clock, phase locked to the PPS signal is multiplied by 5 in the daughter card. The resulting 200-MHz clock feeds a counter which allows to measure a relative time with a 5 ns precision. For every L1A, a 64-bit time-stamp is produced with a relative precision of 5 ns and a 20 ns absolute precision (given by the PPS signal). The time stamping system is shown in Fig. 16.

Figure 16: Time stamping of the CTP_CORE module

3.2.5 The CTP_CAL Module

Sub-detectors have the possibility of requesting calibration triggers during physics data taking. In order to avoid collisions of requests from different sub-detectors, a time-sharing mechanism has been implemented. Each LHC turn (a $88.9 \mu\text{s}$ period between two Orbit signals) is assigned a number modulo 16. Sub-detectors may be assigned, exclusively, turn numbers during which they are allowed to request calibration triggers. The CTP_CAL module time-multiplexes the calibration requests on the CALbus, accepting inputs from the different detectors according to the turn number, and sends accepted requests to the CTP_IN. The CTP_CAL module also provides for up to 28 other external inputs for beam pick-up monitors and test triggers which are also transmitted to the CTP_IN module.

3.2.6 The CTP_MON Module

A block diagram of the CTP monitoring module (CTP_MON), is shown in Fig. 17.

Figure 17: Block Diagram of the CTP_MON module

The CTP_MON module receives and resynchronizes the trigger inputs from the PITbus. It decodes and selects trigger inputs that are to be monitored. It monitors (counts) trigger information on a bunch-by-bunch basis. One can build a histogram of the number of events per BC for each individual input. It may happen that a single PIT signal has no meaning, for instance when it is part of a multiplicity word. In these cases, several PITbus signals can be grouped together and the CTP_MON module histograms occurrences of patterns. One can for instance histogram a given muon multiplicity on a bunch-per-bunch basis.

The CTP_MON module uses numerous segmented memories in four Altera Stratix FPGAs for the counters. A total of 2 MByte is required, i.e. up to 160 trigger inputs times 3564 bunch crossings times 30 bits.

3.2.7 The CTP_OUT Module

The CTP_OUT module receives timing and trigger signals from the COMbus and fans them out to the LTPs of the different TTC partitions. It also receives the busy signals and calibration requests from the sub-detectors. The module masks and monitors the busy signals, and makes them available on the COMbus. It forwards the calibration requests to the CALbus.

Five links (CTP Links) are available per CTP_OUT module, giving 20 links for the complete CTP. One link is used per subdetector, and one link can drive several TTC partitions as seen in Section 4.2.

3.3 The CTP latency

The latency of the CTP is defined as the time between the arrival of the last input in the CTP_IN module and the L1A availability at the output of the CTP_OUT module. The minimum latency is obtained when the clock which strobes the CTP_IN input signals is just a setup-time after the arrival of the last input signal. In this condition, the measured latency of the different stages of the CTP is given in Table 4. The overall latency can be as low as 101ns and is very close to the target value of 100ns.

CTP stage	Remark	Latency
CTP_IN	LVDS input to PITbus	23 ns
CTP_CORE	PITbus to COMbus	74 ns
CTP_OUT	COMbus to LVDS output	4 ns
Overall CTP latency	101 ns	

Table 4: CTP latency

4. Firmware management

FPGA's are extensively used in the CTP and the MUCTPI and two points merit attention: the maintainability of the firmware during the lifetime of the experiment and the remote configuration of the FPGA's.

It is important to be able to maintain the firmware during the lifetime of the experiment, i.e. to keep track of the version of the firmware and to be able to implement necessary modifications. For that purpose, the VHDL sources and the files generated by the vendors' firmware development suites are stored in a repository database centrally maintained by the CERN Information Technologies department (and also used to store the control and data acquisition code used in the system).

The FPGAs' firmware can be downloaded through the VME interface in case a change is required. For changing the trigger menu a change of the firmware is usually not required. A software program translates the logical equations of the trigger menu into LUT and CAM values. These are downloaded into the LUTs and CAMs of the CTP_CORE module at the beginning of a run. The exceptions to this rule are the switching matrices of the CTP_IN modules. In case the selected input signals are changed the same software program generates automatically the VHDL code of the switch matrices and scripts take care of the compilation and of the in-situ programming. However, this is not expected to happen very often.

5. The Timing, Trigger and Control System

The TTC system is used to distribute and fan-out the timing signals (BC, Orbit), the trigger signal (L1A), the eight-bit trigger type summary word, and some commands (Bunch Counter Reset [BCR], Event Counter Reset [ECR]). These signals and data are sent to the read-out and front-end electronics of the subdetector systems. In order to be able to operate sub-detectors or parts of sub-detectors concurrently and independently, e.g. in calibration periods, the TTC system is partitioned. A total of about 40 partitions will be used as shown in Table 5.

Some of the TTC elements that are used have been developed as a common project between LHC experiments and are maintained centrally by the electronics group of the CERN physics department (PH-ESE). A complete description of the TTC project and of these common modules can be found in Ref. [6].

Associated with the TTC system are modules that receive Busy signals from the partitions, as well as calibration requests from the different sub-detector systems.

Sub-detector	Number of partitions
Inner Detector Pixels	3
Inner Detector SCT	4
Inner Detector TRT	3
EM liquid-argon calorimeter	4
Hadronic end-cap calorimeter	1
Forward calorimeter	1
Tile calorimeter	4
Muon TGC	4
Muon RPC	2
Muon MDT	4
Muon CSC	2
Forward detectors	5
Level-1 calorimeter trigger	1
Level-1 muon trigger	1
Level-2/DAQ	1

Table 5: ATLAS TTC partitions

This section details the different parts of the TTC system from the LHC machine interface through to the sub-detectors.

5.1 Timing Signals from the LHC and synchronisation

The LHC bunch structure is shown in Fig. 18. The machine makes available a clock signal per beam (BC) as well as an Orbit signal which allows synchronising within a full LHC turn of 3564 bunch crossings.

Figure 18: LHC bunch structure

To maintain synchronisation within ATLAS, each event is associated with two identifiers: the bunch crossing identifier (BCID) and the event identifier (L1ID). The BCID gives the BC number within an LHC turn at which the event was stored. To maintain synchronisation a bunch counter reset (BCR) signal is issued at each LHC turn. This signal is formed in one of the TTC modules of each partition (TTCvi) and its timing is adjusted with respect to the Orbit signal. The L1ID is formed in counting the L1A. A 24-bit counter is embedded in the readout electronics and an event counter reset (ECR) signal is regularly issued to reset these counters and make sure that the synchronisation is maintained. In addition an 8-bit counter in the readout electronics counts the number of ECR. The concatenation of this 8-bit counter with the previously mentioned 24-bit counter forms a unique 32-bit L1ID. The ECR is issued by the CTP and distributed by the TTC system.

Optical fibres are used to transmit the different LHC clocks and Orbit signals from the machine to the experiments. One clock per beam (BC1 and BC2), and one reference clock (BC) to which BC1 and BC2 are locked during physics runs, are made available; in addition, there are two Orbit signals (one per beam). Furthermore, information concerning the status of the machine (mode of operation, energy, etc.) is received via the Beam Synchronous Timing (BST) network [22].

A so-called LHC-TTC [19] crate contains optical-to-electrical converters, an interface (RF2TTC) board that decodes the BST packets and selects and adjusts the timing of the BCx and Orbitx to be used, and fan-out modules. The selected BCx and Orbitx are sent to the CTP from where they are distributed to all sub-detectors.

The clock signals are transmitted over several kilometres of optical fibres. The propagation delay on these fibres is sensitive to the ambient temperature and it is expected to have timing variations of a few nanoseconds over a year. It is therefore important to be able to monitor these variations and to compensate for them. Beam pick-up electrodes located on both sides of the experiment are used to monitor the clock phase with respect to the beam arrival times. A global timing correction can then be applied in the RF2TTC module, making this timing adjustment transparent for the different sub-detectors.

5.2 CTP Links

As mentioned in Section 3.2.7, the timing and trigger signals are fanned out by the CTP_OUT modules on up to 20 CTP links. One link is used per sub-detector. In general, there are several TTC partitions for a given sub-detector, and the corresponding CTP link is daisy chained from one partition to another (see Section 4.3.1).

The CTP delivers on the CTP_OUT links the BC, ORBIT, L1A and ECR signals, plus the eight-bit Trigger-Type summary word. It receives a BUSY signal and a three-bit calibration trigger request. Note that if a sub-detector contains more than one partition, the BUSY signal that is transmitted to the CTP is the logical “OR” of the BUSY signals of these sub-detector partitions (see Fig. 21). Table 6 and Fig. 19 detail the distribution of signals between the CTP and the TTC partitions.

The CTP links are made using low skew 25-pair cables and all the signals are transmitted using the LVDS standard.

Name	Function
BC	LHC clock
ORBIT	LHC ORBIT signal used for instance to issue the BCR signal
L1A	L1 Accept signal
Trigger-Type	Eight-bit trigger-type word issued by the CTP along with each L1A
ECR	Event Counter Reset signal. Signal used to reset the 24 low-order bits and to increment the eight high-order bits of the L1ID
Pre-pulse	A signal generated by the CTP a defined number of BC before a L1A is issued. This would allow sub-system to generate tests signals (e.g laser or generator shots) in time for being readout when this L1A occurs.
BUSY	The BUSY signal generated by the RODs of the sub-detector when their buffers are almost full. This is used by the CTP to introduce dead-time.
Calibration	Three-bit word issued by the sub-detector and used by the CTP to generate calibration triggers

Table 6: Signals exchanged between the CTP and a TTC partition

Figure 19: Interface between the CTP and a TTC partition

5.3 TTC partition

Each TTC partition contains one Local Trigger Processor (LTP) and, optionally an LTP Interface module (LTPIM). It also contains a TTC system proper, consisting of the VME to TTC interface (TTCvi) and the electrical to optical conversion module (TTCex or TTCtx), and a busy tree which allows the subdetector to throttle the generation of L1As using ROD_BUSY modules.

Each partition has its own TTC network controlled by a TTCvi module, and can be run in two modes:

- Global mode, where the TTC signals are taken from the CTP and where the sub-detector data are part of the main ATLAS run and read out through the main ATLAS central DAQ system.
- Local mode, where the partition is run stand-alone with “private” TTC signals and where the data are kept separate from the main ATLAS run.

5.3.1 TTC partition root

An ATLAS TTC partition is driven by a set of four or five modules as shown in Fig. 20:

- the optional LTPIM (not shown on the picture);
- the LTP;
- the TTCvi [20];
- the TTCxx (ex, vx or tx);
- the ROD-Busy module.

The LTP interfaces the partition to the CTP when running in Global mode (i.e. under control of the CTP) and to the local trigger logic when running in Local mode.

Figure 20: A TTC partition root

The interface to the CTP is done through a differential link (CTP Link) that was already discussed above. As the number of TTC partitions in ATLAS is high, it is not affordable to have as many CTP Links as partitions. Only 20 links will be available, one link per sub-detector. Each CTP Link connects the CTP to all of the LTPs of a given sub-detector using a daisy chain as shown in Fig. 21.

The TTCvi provides the A- and B-Channel signals to the TTCxx module which contains the encoder and electrical-to-optical converter, and feeds the TTC optical network.

The ROD-Busy module gathers the Busy signals from the RODs attached to the partition to form an overall BUSY signal which will throttle the trigger source (either the CTP or some local trigger logic). The BUSY signal sent back to the CTP is the logical “OR” of the BUSY signals issued from the LTPs that are daisy-chained.

Figure 21: Daisy chain of the Busy signals of the different partitions of a sub-detector

When running in local mode, all signals can be generated and handled locally (i.e. a connection with the CTP is not needed); however, the BC and Orbit signals are always available on the CTP links and may be used. Local trigger inputs are made available and output towards the trigger inputs of the TTCvi. The readout deadline is in that case locally handled; the local BUSY signal is not transmitted to the CTP. Similarly, if needed, the Trigger-Type has to be generated locally.

The test triggers and the local commands can be issued either externally (with a generator for instance) or internally, using a pattern-generator of one megasample depth.

Fig. 22 shows a view of the TTC root modules as used in local mode.

Figure 22: View of the TTC root modules when in local mode

It is possible to run several partitions in local mode with the same timing and trigger signals (e.g. a run involving all the partitions of a given sub-detector). In this case, one of the LTPs acts as a master. It will transmit on its output CTP link all the trigger and timing signals used locally and will receive the BUSY from the other LTPs.

In order to be capable of using together several sub-detectors without the CTP (for instance to run the muon chambers together with the muon trigger systems, and, at the same time, the calorimeters together with the calorimeter trigger) an additional LTP Interface Module (LTPIM) can be introduced between the CTP and the first LTP of a CTP link. This module acts as a switch and allows one to connect several sub-detectors without the CTP. One of the LTPs involved is, in this case, playing the role of the CTP. Fig. 23 shows a configuration with three sub-detectors connected.

Figure 23: Three sub-detectors, each of them having four TTC partitions. The top left LTP (marked Master LTP) acts as a CTP in this configuration.

As the distance between two LTPIM can be large, special care has been taken for signal recovery and active equalisers are used. In addition, programmable delay chips have been introduced in order to be able to adjust the timing of each signal.

5.3.2 TTC optical tree and receivers

The encoded optical signal generated in the TTCex or TTCtx module is fanned-out in optical passive filters. Three types are used in ATLAS, all housed in a 1U 19" sub-rack: single 1-to-32 splitter, dual 1-to-16 splitters and quad 1-to-8 splitters. Single-mode fibres are used from the TTCex or TTCtx up to the splitter inputs; multi-mode fibres are used after the splitters. At the receiving end, a PIN diode converts the optical signal to an electrical one and an ASIC, the TTCrx, recovers the BC clock and the encoded signals.

5.3.3 Timing adjustment

Programmable delays are available at many points in the system in order to be able to do fine timing adjustments of the clock on the front-end electronics, or coarse delay adjustments (in step of one clock period, 25 ns) to adjust the L1A timing on the front-end (to match the level-1 pipeline length) and the BCR (to make sure that all the BCID [bunch crossing identifier] values are properly set in the entire read-out system). The timing adjustment is done for each TTC partition using programmable delays in the LTP, the TTCvi and the TTCrx.

6. Modules fabrication

All the modules have been built, tested and deployed in the ATLAS experimental zone. For both the CTP and the MUCTPI, it has been decided to build three systems. One to be used in the experiment, one used as a running spare and the third used in the laboratory for software development.

Three complete CTP systems were built and are now fully working. In the case of the MUCTPI, there exist three crates with three MICTP and MIROD boards, but only 35 MIOCT boards were built, which is sufficient for having running spares and a development system.

In all, 80 ROD-BUSY, 66 LTP and 34 LTPIM modules have been produced and made available to the different ATLAS sub-detectors. The TTC modules (TTCvi, TTCex, TTCtx, ...) have been centrally produced by a CERN support group [21] which also takes care of the maintenance and support.

7. Performance of the system

The CTP and the MUCTPI have been installed in the ATLAS experimental zone together with the TTC distribution system. Figs. 24, 25 and 26 show the CTP and the MUCTPI in situ as well as a typical TTC crate and its optical fanout.

Figure 24: Photograph of the CTP in the ATLAS underground counting room

Figure 25: Photograph of the MUCTPI in the ATLAS underground counting room. Eight MIOCT modules (on the right side) still have to receive their input cables.

Figure 26: Photograph of one TTC crate and its optical fan-out

Extensive tests have been performed of all functions of the CTP and the MUCTPI, in particular the real-time trigger path, but also the readout, monitoring and busy mechanism. The

latency of the system (the time from input to output) was measured to be 4 BC for the MUCTPI and 4 BC for the CTP. Both values are within the latency budget foreseen for these components.

The DAQ and Level-2 readout of the final, full CTP crate has been tested at the maximum L1A rate of 100 kHz without introducing deadtime. The same has been tested for a demonstrator prototype version of the MUCTPI using two MIOCT boards and an average of three muon candidates per event.

The system has been used extensively operating the ATLAS detector in situ with a trigger on cosmic-ray muons. Fig. 27 shows a sample display of a cosmic-ray event triggered by the end-cap muon trigger with hits in the muon precision chambers.

Figure 27: Event display of a cosmic-ray muon triggered via the central trigger logic

The MUCTPI and CTP have been in routine operation in the ATLAS underground counting room since 2006 providing cosmic-ray and technical triggers through the final trigger chain. Their readout is integrated with the final data-acquisition chain as well as with the Level-2 trigger system and has been shown to withstand the maximum L1A rate of 100 kHz. The configuration, control and monitoring of the MUCTPI and CTP is integrated in the ATLAS online operation framework.

Connection tests have been performed with the muon trigger electronics, as well as with the calorimeter trigger system. Currently, more than 30 TTC partitions are being used, connected to the CTP using 16 output links. Each of these partitions delivers the TTC signals from the LTP to the sub-detector front-ends in the cavern and the RODs in the underground counting room. These TTC partitions are routinely driven by an LTP in stand-alone runs and by the CTP in combined runs.

8. Conclusion

The ATLAS Level-1 Central Trigger Logic, consisting of the MUCTPI, the CTP and the TTC system, has been successfully built and installed in the experiment and the complete trigger and readout chain is being operated while triggering on cosmic rays. The overall system provides a large degree of flexibility and programmability while meeting very stringent requirements in terms of latency.

Acknowledgements

We would like to thank Sophie Baron, Stephane Detraz and Markus Joos from the CERN PH-ESE group for their support of the TTC systems.

We would like to thank Gary Beetham, Quentin King and Javier Serrano from the CERN AB department for their invaluable help in integrating the GMT hardware in the CTP.

References

- [1] The ATLAS Collaboration, *First-level Trigger Technical Design Report*, CERN/LHCC/98-14, June 1998.
- [2] R. Achenbach et al, *The ATLAS Level-1 Calorimeter Trigger*, 2008_JINST_3_P03001
- [3] *The ATLAS Level-1 Barrel Muon Trigger*, <http://sunset.roma1.infn.it/muonl1/docs/publications>
- [4] *The ATLAS Level-1 End-cap Muon Trigger*,
http://atlas.web.cern.ch/Atlas/project/TGC/www/doc/MuonEndcap_rev01.pdf
- [5] P. Gallno, *The ATLAS Local Trigger Processor (LTP)*, <http://edms.cern.ch/document/374560>
- [6] S. Baron, *The TTC System*, <http://ttc.web.cern.ch/TTC/intro.html>
- [7] P. Gallno, *The ATLAS ROD_BUSY Module*, <https://edms.cern.ch/document/319209/1>
- [8] *The ATLAS Dataflow System*, <http://atlas.web.cern.ch/Atlas/GROUPS/DAQTRIG/dataflow.html>
- [9] *The ATLAS Online System*, <http://atlas-onlsw.web.cern.ch/Atlas-onlsw>
- [10] Ph. Farthouat, *Interfaces and Overlaps in the Level-1 Muon Trigger System*,
<https://edms.cern.ch/document/114604>.
- [11] O. Boyle et al., *The S-LINK Interface Specification*, <http://cern.ch/s-link>
- [12] National Semiconductor, *LVDS Owners Manual*, <http://lvds.national.com>.
- [13] ANSI, *SCSI Parallel Interface 4 (SPI-4)*, May 2002.
- [14] Altera Corporation, *Stratix II Device Handbook*, April 2006, <http://www.altera.com>.
- [15] R. Spiwoks, *The VMEbus Interface of the Central Trigger Processor*,
<http://edms.cern.ch/document/428910>
- [16] A. Marchioro, *The CERN Four-Channel 1/2-ns Programmable Delay Line (DELAY25)*, <http://proj-delay25.web.cern.ch/>
- [17] http://tdc.web.cern.ch/tdc/hptdc/docs/hptdc_manual_ver2.2.pdf
- [18] G. Beetham and Q. King, *Interface Specification for the CTRA LHC General Machine Timing Receiver Card*, <https://edms.cern.ch/document/338833/1.2>
- [19] S. Baron et al., *Status of the TTC upgrade*, in *Proc. 2006 LECC workshop*, CERN-2007-001, http://documents.cern.ch/cgi-bin/setlink?base=cernrep&categ=Yellow_Report&id=2007-001, p.p. 236-240
- [20] P. Gallno, *The TTCvi Module – Technical Description and User Manual*,
<http://www.cern.ch/TTC/TTCviSpec.pdf>

- [21] *Support for the TTC system*, <http://ess.web.cern.ch/ESS/TTCsupport/index.htm>
- [22] R. Jones, *The Interface and Structure of the Machine Beam Synchronous Timing Message for the LHC Experiments*, LHC EDMS note LHC-BOB-ES-0001.
<https://edms.cern.ch/document/638899/0.1>