

Development of the Barrel Sector Logic firmware of the ATLAS Muon Spectrometer for the High-Luminosity LHC^(*)

F. MORODEI on behalf of the ATLAS COLLABORATION

*Dipartimento di Fisica, Università La Sapienza & INFN Sezione di Roma
Piazzale Aldo Moro 5, 00185, Roma, Italy*

received 13 February 2024

Summary. — The Muon Spectrometer of the ATLAS detector at the Large Hadron Collider will undergo a substantial improvement during the Long Shutdown 3 to cope with the instantaneous luminosity and radiation levels foreseen for the High-Luminosity LHC. The entire trigger and readout electronics of the Resistive Plates Chambers will be replaced. All data coming from these detectors will be collected by the on-detector Data Collector and Transmitter boards and then sent to the off-detector Sector Logic boards, which will execute the muon Level-0 trigger algorithm and perform the readout logic. This contribution describes the development of the barrel Sector Logic firmware and the related challenges.

1. – Introduction

During Long Shutdown 3 (2026-2028) the ATLAS detector [1] will undergo a major upgrade (Phase-II Upgrade) in preparation to the High-Luminosity LHC (HL-LHC) operations. Many detector components will be improved or completely replaced in order to cope with the much harsher conditions due to the increase in the instantaneous luminosity, which may reach 7.5 times the nominal LHC value of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The higher event pile-up and radiation levels of HL-LHC impose to completely replace the Trigger and Data Acquisition (TDAQ) system of the Resistive Plate Chambers (RPCs), except for the current front-end boards. The RPCs are gaseous detectors arranged in three concentric rings in the barrel region of the Muon Spectrometer and are used to trigger muons. Currently the RPC system is composed of three stations of doublet RPC chambers, two of them placed in the middle region of the barrel (BM1 and BM2) and one in the outer region (BO). With the Phase-II Upgrade an RPC station of triplet chambers will be installed in the barrel inner region (BI) to increase the redundancy of the trigger system and to cover the current acceptance holes [2].

^(*) IFAE 2023 - “Poster” session




Fig. 1. – Scheme of the barrel RPC TDAQ system for the HL-LHC.

2. – The Phase-II RPC TDAQ system

During the Phase-II Upgrade the current Pad and Splitter boxes will be replaced by 1546 Data Collector and Transmitter (DCT) boards, installed on-detector to collect the front-end RPC data. Two different kinds of DCT board are foreseen: in the BI region the digitisation and time measurement of the RPC signals will be performed by the new front-end boards and therefore the DCTs will receive digital signals, while in the BM and BO regions the DCTs will receive analog pulses from the BM/BO legacy front-end electronics and will digitise them. Through the Low Power Gigabit Transceiver (IpGBT), a radiation tolerant ASIC, the DCT will communicate with the off-detector Sector Logic (SL) boards, sending the RPC hits and receiving commands and control data. The barrel SL boards (32 in total), connected with up to 50 DCTs each, will have to generate muon trigger candidates for the ATLAS Level-0 (L0) trigger and perform the readout logic. The SL board consists of a Xilinx Virtex Ultrascale+ XCVU13P FPGA which implements the SL firmware, a Xilinx Zynq Ultrascale+ SoC module to communicate with the Detector Control Systems (DCS) and the L0 barrel TDAQ server and an IPMC module to communicate with the ATCA shelf manager. The muon trigger candidates produced by the SL firmware will be sent to the MDT Trigger Processor (MDT-TP) for confirmation and then to the Muon Central Trigger Processor Interface (MUCTPI). The RPC data will be sent to the ATLAS readout system through the Front-End Link Interface eXchange (FELIX) modules. The requirements of the L0 RPC-based trigger for HL-LHC will be 100 kHz trigger rate and 390 ns latency, while the global requirements for the ATLAS L0 trigger will be 1 MHz trigger rate and 10 μ s latency [3]. The barrel RPC TDAQ system for the HL-LHC is schematically shown in fig. 1.

3. – The SL FPGA firmware

The development and the implementation of the SL firmware have been particularly challenging due to very strict timing and spatial constraints. The XCVU13P FPGA is divided into four die slices called Super Logic Regions (SLRs), thus a careful floorplanning has been conceived to divide the logic among them. Each SLR is provided with 32 high speed transceivers (GTYs), having a transmitting (TX) and a receiving (RX) fibre each. GTYs are organised in groups of four, called Quads, which share the same reference clock. Figure 2 presents a scheme of the four SLRs of the SL FPGA, showing the GTY usage, bandwidth and connections with the external systems adopted for the SL firmware.

SLR0 receives the BI RPC data and the Tile Calorimeter data through 10 and 6 GTYs respectively, decodes and reorders them by Bunch Crossing (BC), since the DCT hits

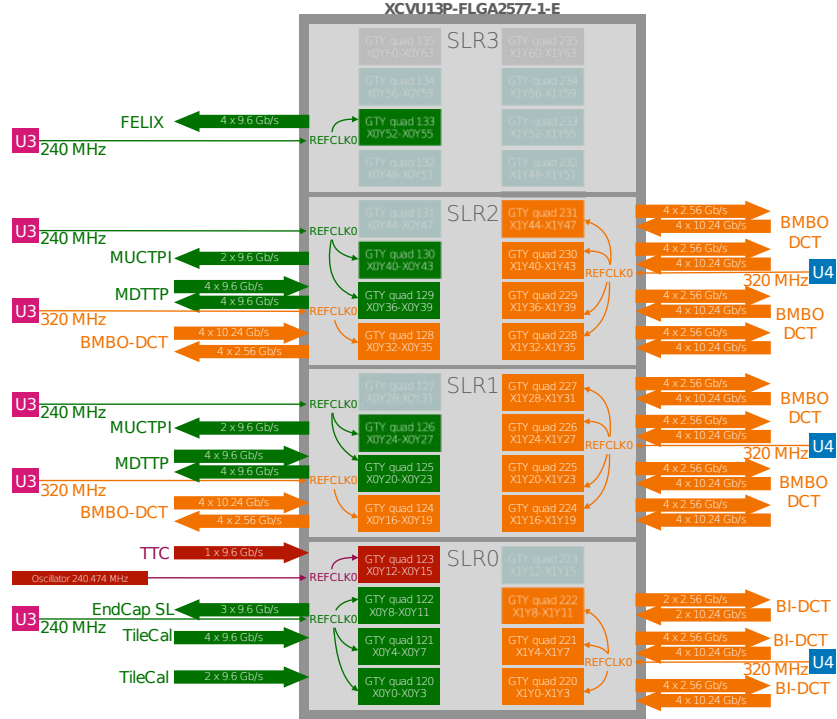


Fig. 2. – SL FPGA transceiver usage, bandwidth and connections with the external systems.

arrive to the SL with non-fixed latency, and passes them to SLR1 and SLR2. SLR1 and SLR2 receive the BM/BO RPC data through 20 GTYs each, decode and reorder them and use them together with the BI RPC data from SLR0 to perform the L0 RPC-based trigger algorithm, independently for each half sector. Each trigger algorithm instance generates up to two candidates per BC requiring a hit coincidence in at least three RPC stations. SLR1 and SLR2 then send the trigger candidates to the MDT-TP via 2 TX fibres and receive back confirmation via 2 RX fibres each. The trigger candidates are finally sent to the MUCTPI through 2 fibres for each SLR. Moreover, SLR0, SLR1 and SLR2 pass all the RPC data to SLR3, where the readout logic stores and reorders them by BC using 50 RAM memories. When a L0-Accept signal is received for a certain BC, the corresponding RPC data are sent to the FELIX modules through 3 GTYs. Old data are discarded if no L0-Accept signal is received. All the DCT GTYs use the lpGBT encoding, work with a 320 MHz reference clock and have TX and RX line rates of 2.56 and 10.24 Gb/s respectively, while all the other GTYs use the 8b/10b encoding, work with a 240 MHz reference clock and use a line rate of 9.6 Gb/s both for TX and RX. The GTY reference clocks are generated by two jitter cleaners (U3 and U4) using the 40 MHz LHC TTC reconstructed clock. The 40 MHz TTC clock is used also to generate, through clock multipliers embedded in the FPGA, the 240 MHz and 320 MHz clocks needed by the firmware logic.

The VHDL simulation and implementation of the SL firmware have been performed with the Xilinx-AMD Vivado software. The multi-die structure of the XCVU13P device and the high frequency clocks used by the internal logic made a careful planning of

Resource	Utilization	Available	Utilization %
LUT	605645	1728000	35.05
LUTRAM	1341	791040	0.17
FF	471106	3456000	13.63
BRAM	497.50	2688	18.51
IO	227	448	50.67
GT	78	128	60.94
BUFG	18	1344	1.34
MMCM	1	16	6.25

Fig. 3. – Overall post-implementation resource utilisation of the SL FPGA.

the various logic blocks necessary to fulfill the spatial and timing constraints. Indeed the firmware logic and the I/O ports had to be divided among the four SLRs so as to minimise the number of signal crossings, since signals can be transmitted between different SLRs only through dedicated routes which have a slower transmission speed with respect to internal routes and are placed mostly in the central part of the die. After many trials, the design that optimises the FPGA resource utilisation has been found to be the one described above, which separates the logic blocks related to the BI DCTs and to the two halves of the BM and BO DCTs in different SLRs. This division is possible since the trigger algorithm can be run independently for the two sides of the barrel, with the only common signals being the BI data. The remaining SLR has been dedicated entirely to the readout logic. To solve the timing closure issues due to SLR signal crossings, pipeline registers have been added according to the number of SLRs crossed by the signals. The overall post-implementation resource utilisation, reported in fig. 3, is 35% of the lookup tables (LUT) and 14% of the flip-flops (FF), so there are still enough resources available for the last missing pieces of logic needed to complete the SL firmware (mostly monitoring logic).

4. – Conclusion

The RPC trigger and readout electronics of the ATLAS barrel will be completely replaced during LS3 in preparation to the HL-LHC, with the installation of on-detector DCT boards, to collect RPC data, and off-detector SL boards, to perform trigger and readout logic. The SL firmware development has been challenging because of the large amount of resources needed to process all RPC data and the very strict timing constraints. A careful floorplanning has allowed to achieve a successful implementation of the SL firmware. Few pieces of logic still need to be added to complete the firmware, but their addition is not expected to cause issues, since there are still a lot of FPGA resources available. The next step is the validation of the SL firmware using the SL prototype.

REFERENCES

- [1] THE ATLAS COLLABORATION, *JINST*, **3** (2008) S08003.
- [2] THE ATLAS COLLABORATION, *Technical Design Report for the Phase-II Upgrade of the ATLAS Muon Spectrometer*, CERN-LHCC-2017-017.
- [3] THE ATLAS COLLABORATION, *Technical Design Report for the Phase-II Upgrade of the ATLAS Trigger and Data Acquisition System*, CERN-LHCC-2017-020.