

The ATLAS Muon to Central Trigger Processor Interface Upgrade for the Run 3 of the LHC

Aaron Armbruster, German Carrillo-Montoya, Magda Chelstowska, Patrick Czodrowski, Pier-Olivier Deviveiros, Till Eifert, Nick Ellis, Philippe Farthouat, Gorm Galster, Stefan Haas, Louis Helary, Orestis Lagkas Nikolos, Yusuf Leblebici, Antoine Marzin, Thilo Pauly, Vladimir Ryjov, Augusto Santiago Cerqueira, Kristof Schmieden, Marcos Silva Oliveira, Ralf Spiwoks, Joerg Stelzer, Alain Vachoux, Paschalis Vichoudis, and Thorsten Wengler

Abstract—To cope with the higher luminosity and physics cross-sections for the third run of the Large Hadron Collider (LHC) and beyond, the Trigger and Data Acquisition (TDAQ) system of the ATLAS experiment at CERN is being upgraded. Part of the TDAQ system, the Muon to Central Trigger Processor Interface (MUCTPI) receives muon candidates information from each of the 208 barrel and endcap muon trigger sectors, counts muon candidates for each transverse momentum threshold and sends the result to the Central Trigger Processor (CTP). The MUCTPI takes into account the possible overlap between trigger sectors in order to avoid double counting of muon candidates. A full redesign and replacement of the existing MUCTPI is required in order to provide full-granularity muon position information at the bunch crossing rate to the Topological Trigger processor (L1Topo) and to be able to interface with the new sector logic modules. State-of-the-art FPGA technology and high-density ribbon fibre-optic transmitters and receivers are being used to implement the MUCTPI in a single AdvancedTCA blade, compared to 18 9U VMEbus cards in the existing system. The upgraded MUCTPI features over 270 multi-gigabit optical inputs/outputs with an aggregate bandwidth of over 2 Tbit/sec. This work presents the hardware design, results from the validation of the first prototype, the testing of the optical interfaces, and integration tests with the muon sector logic.

I. INTRODUCTION

ATLAS [1] is a particle physics experiment at the Large Hadron Collider (LHC) at CERN. It observes proton-proton collisions at a center-of-mass energy of 13 TeV. With an average of 25 collisions per bunch crossing (BC) every 25 ns, there are 10^9 interactions per second potentially producing interesting physics. The online trigger of ATLAS is structured in a 2-level architecture in order to reduce the event rate from a bunch crossing rate of 40 MHz down to 1 kHz written to permanent storage. Figure 1 shows the Trigger and DAQ system block diagram. The first level (Level-1 trigger) is implemented with custom electronics, while the High-Level Trigger (HLT), is built from commercial computers, network switches, and custom software. The Data Acquisition (DAQ) system receives and buffers the event data from the detector-specific readout electronics in pipeline memories until the

A. Armbruster, G. Carrillo-Montoya, M. Chelstowska, P. Czodrowski, P.-O. Deviveiros, T. Eifert, N. Ellis, P. Farthouat, G. Galster, S. Haas, L. Helary, O. Lagkas Nikolos, A. Marzin, T. Pauly, V. Ryjov, K. Schmieden, M. Silva Oliveira, R. Spiwoks, J. Stelzer, P. Vichoudis, T. Wengler are with CERN, CH-1211 Geneva 23, Switzerland (corresponding author e-mail: marcos.oliveira@cern.ch)

A. Santiago Cerqueira is with UFJF, Juiz de Fora, MG, Brazil.

Y. Leblebici and A. Vachoux are with EPFL, STI LSM, Station 11, 1015 Lausanne, Switzerland.

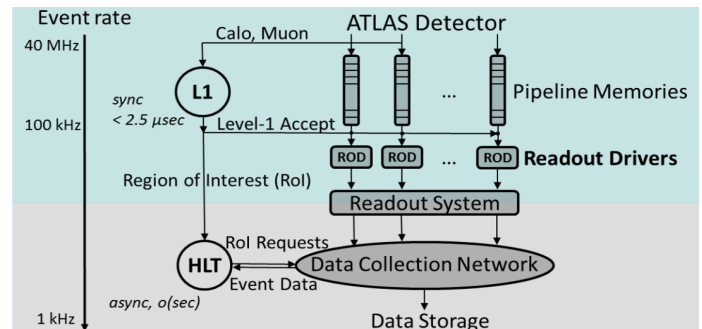


Fig. 1. Trigger and DAQ system

arrival of the Level-1 Accept (L1A) signal with a rate of up to 100 kHz.

II. LEVEL-1 TRIGGER SYSTEM

Figure 2 shows the Level-1 trigger system block diagram. The Level-1 trigger system performs fast event selection based on reduced-granularity information from the calorimeters and muon detectors. Information from the calorimeter (L1Calo) and muon (L1Muon) trigger consist of multiplicities, energy and positions of trigger candidate objects. The Level-1 trigger data transfer is based on a system-synchronous clocking technique with the clock being distributed on a separate network, the so-called Timing, Trigger and Control (TTC) system [2]. The Muon to Central Trigger Processor Interface (MUCTPI) [3] combines the information from the muon trigger sector logic (SL) modules of the barrel and end-cap regions of the detector and then calculates the total multiplicity of muon candidates. The MUCTPI avoids double counting of muons candidates which transverse more than one detector region due to geometrical overlap of the chambers and the trajectory of the muon in the magnetic field. The MUCTPI sends the multiplicity for each of six energy thresholds to the Central Trigger Processor (CTP). It also sends muon position and energy information of selected muon candidates [4] to the Level-1 Topological Trigger Processor (L1Topo) [5]. The CTP combines information from L1Calo, MUCTPI, and L1Topo to make the final L1A decision.

III. MUCTPI UPGRADE

In an effort to improve the chances of seeing rare events, the luminosity of the LHC will reach twice its nominal

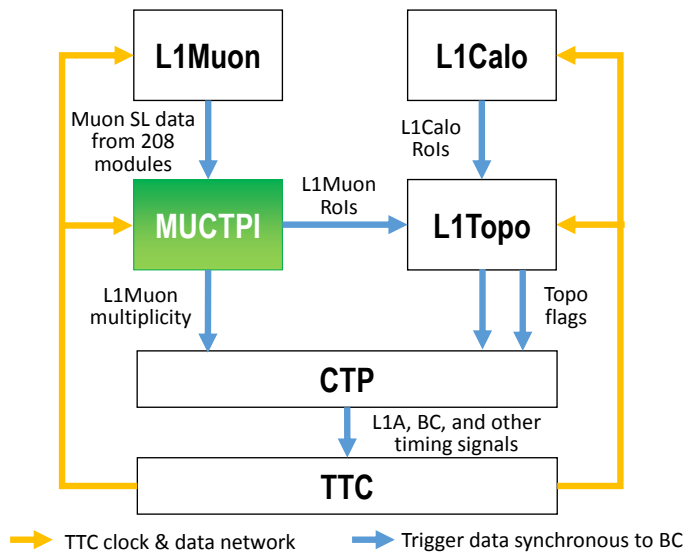


Fig. 2. Level-1 trigger system block diagram

luminosity value of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ in Run 3 (2021-2023). As the luminosity increases, the trigger system has to become more selective to keep output rates below 100 kHz in order to not overwhelm the DAQ and HLT systems. As part of the ATLAS Trigger and Data Acquisition (TDAQ) Phase-I Upgrade [6], a full redesign and replacement of the existing MUCTPI is required in order to provide full-granularity muon RoI information at the BC rate to L1Topo, allowing combined calorimeter/muon topological trigger algorithms, and to be able to interface to the new SL modules using high-speed optical links. The smaller physical dimensions of the optical modules combined with the higher FPGA densities available today enable the implementation of all the required MUCTPI functionality on a single AdvancedTCA [7] blade, compared to 18 9U VMEbus [8] cards in the existing system. The AdvancedTCA form factor offers improved hardware platform monitoring, cooling and power supply.

IV. MUCTPI ARCHITECTURE

Figure 3 shows the block diagram of the MUCTPI module. The new system [9] is based on a highly-integrated generation of FPGAs (Xilinx 20 nm Ultrascale devices), featuring a large number of on-chip multi-gigabit transceivers (MGTs) as well as 12-channel ribbon fiber optics receiver and transmitter modules (MiniPOD) for the data transfer. The higher bandwidth of the high-speed serial optical connections from the muon trigger sector logic modules to the MUCTPI system enables a higher number of muon candidates to be received. The upgraded MUCTPI will receive up to 4 candidates per trigger sector instead of 2. Figure 4 show a photo of the new MUCTPI prototype board.

A. Muon Sector Processor

The functionality of 8 existing modules, i.e. one half of the detector, will be merged into a single large Virtex UltraScale FPGA (VU160), the Muon Sector Processor (MSP). The

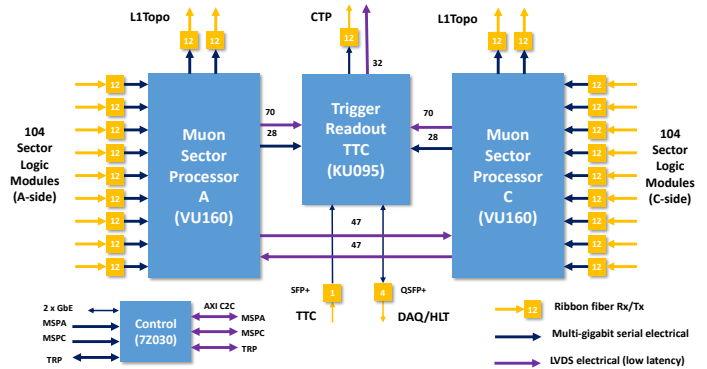



Fig. 3. MUCTPI architecture



- Avago MiniPODs
- Point-of-load DC/DC converters
- MSP FPGAs (VU160)
- 12/24 MPO connectors
- TRP FPGA (KU095)
- TTC SFP module
- SoC FPGA (7Z030)
- JTAG/UART ports
- DAQ/HLT QSFP
- 48 V to 12 V DC/DC
- DDR3L SDRAM
- IPMC mezzanine



Fig. 4. MUCTPI prototype

two FPGAs together receive and process muon trigger data from 208 sector logic modules connected through high-speed serial optical links using MiniPOD receiver modules. The MSP FPGAs also copy information on selected muon trigger objects to several L1Topo modules using MiniPOD transmitter modules.

B. Trigger, Readout, and TTC processor

The Trigger and Readout Processor (TRP) FPGA is a Kintex UltraScale device (KU095) that merges the information received from the two MSP FPGAs through LVDS and MGT links, and sends the results to the CTP. In addition, it will be used to implement muon topological trigger algorithms. This is possible because all the trigger information is available in a

single module with low latency. The same FPGA also receives, decodes and distributes the TTC information. Finally, it sends the muon trigger information to the DAQ and HLT systems when it receives an L1A decision.

C. System on chip

A Xilinx Zynq-7000 System on Chip (SoC) (7Z030) is used for configuration, control and monitoring of the module. The device integrates a programmable logic part with a dual-core ARM processor subsystem. The processor subsystem will run the required software to interface the MUCTPI to the ATLAS run control system through a Gigabit Ethernet (GbE) interface. It will also be used for environmental monitoring of components of the board via I2C, such as the power supply, optical modules, and FPGAs. The values read include voltages, currents, temperatures, optical input power, clock status, etc. Finally, the SoC is used to load the configuration bitstreams into the three FPGAs.

D. On-board connectivity

Connections using both general purpose I/O pins and dedicated MGT pins are used to exchange information between the FPGAs on the board. Each MSP FPGA can share data with the other MSP FPGA using 47 LVDS pairs. Operating each LVDS pair at a bit rate of 1.28 Gb/s, results in a total bandwidth of ≈ 60 Gb/s each way. This is sufficient to share all the trigger information from the barrel region between the two MSP FPGAs. In addition, 70 LVDS pairs are connected from each MSP FPGA to the TRP FPGA, resulting in a total bandwidth of ≈ 90 Gb/s. This connection will be used to send pre-processed trigger information to the TRP FPGA with low-latency.

All the on-board LVDS links were tested at 1.28 Gb/s. No error was detected and an eye-diagram horizontal opening of 70 % was measured.

Twenty-eight MGT links are connected from each MSP FPGA to the TRP FPGA. Operating each of these MGT links at a bit rate of 12.8 Gb/s with 8b10b encoding results in a total bandwidth of ≈ 570 Gb/s. Up to 4 links will be required for the transfer of the readout data. The remaining links could be used to transfer muon candidate information for muon topological trigger processing.

E. Off-board connectivity

Each MSP FPGA receives the muon candidate information from 104 SL modules through 9 MiniPODs. It also transmits muon candidate information to L1Topo through 24 MGT links using 2 MiniPODs. The TRP FPGA receives TTC clock and data through a SFP+ module, and sends information to the DAQ and HLT systems using a QSFP+ module. In addition, one MiniPOD is available for sending trigger bits (muon multiplicities, trigger flags, etc.) to the CTP using up to 12 MGT links. For backwards compatibility and to minimize latency, a parallel electrical LVDS signal connection through a 68-pin SCSI VHDCI connector is also foreseen for sending trigger information to the CTP.




Fig. 5. Serial link test automation

V. TESTING OF MGT SERIAL LINKS

In order to simplify the board layout, swapping of high-speed serial links and their polarities were allowed. Due to the very high number of high-speed serial links (334) in the MUCTPI, a collection of Python scripts were written in order to extract interconnectivity information from the back-annotated board design netlist and to automate the test running and compilation of eye diagrams for all the links running at 6.4, 9.6, and 12.8 Gb/s.

The integrated bit error ratio test firmware from Xilinx (IBERT IP core) was used to test the MGT links in all the FPGAs. Figure 5 shows the serial link test automation flow diagram. TCL scripts are generated to configure and read the results from the Xilinx Vivado software based on the board design netlist and FPGA package pin files. The Xilinx Vivado software communicates with the MUCTPI through Ethernet using a hardware server (“virtual cable”) running on the SoC which is connected to the FPGAs via the JTAG chain.

All of the MUCTPI serial connections including on-board and off-board MGT links were checked for errors by transmitting and receiving PRBS-31 pattern data. In addition, a long-term BER run measurement with 116 serial links running at 12.8 Gb/s was done. No error was detected in 10 days which corresponds to a BER (Bit Error Rate) lower than 1×10^{-15} with a confidence level of 95%.

Figure 6 shows the eye diagrams of 12 MGT optical links between the two MSP FPGAs through MiniPOD transmitter and receiver modules. All the links are running at 6.4 Gb/s with an excellent eye-opening of 100% vertical and 80% horizontal.

VI. TTC RECOVERY

The MUCTPI features a Clock and Data Recovery (CDR) chip to recover the TTC clock and data which was tested and worked. As an exploratory work, we decided to investigate an alternative way to recover the TTC clock and data using only an on-chip MGT. Figure 7 shows the block diagram of this alternative TTC recovery and clock distribution circuitry. An off-the-shelf OC-3 SFP transceiver module converts the TTC signal from optical to electrical. An on-chip MGT recovers the TTC clock and encoded data. This MGT oversamples the TTC signal because the TTC bit rate of 160 Mb/s is lower than the minimum MGT operation bit rate of 500 Mb/s. The phase of the recovered clock is unknown at this stage. A custom alignment block guarantees a fixed phase for the recovered clock by aligning the recovered clock to the edges of the TTC signal. An external clock generator and a crystal are used to generate the reference clock required by this MGT channel.




Fig. 6. Eye diagrams of 12 MGT optical links between the two MSP FPGAs through MiniPOD transmitter and receiver modules

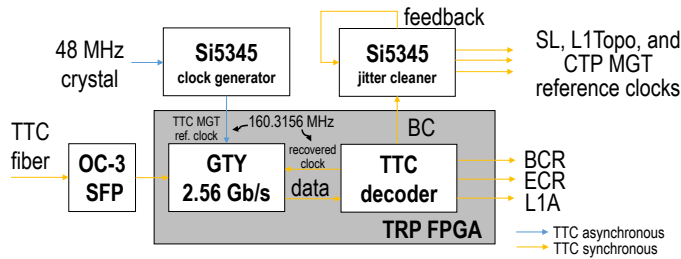


Fig. 7. TTC recovery and clock distribution

As the recovered TTC clock is not sufficiently clean to drive the reference clocks of the MGT channels directly, a jitter cleaner is used to generate the MGT reference clocks with their required frequencies. The jitter cleaner uses an external feedback connection to ensure fixed input-to-output delay. A custom logic block decodes the TTC data which is distributed to the other FPGAs through on-board electrical connections.

VII. SL SYNCHRONIZATION & ALIGNMENT

The MUCTPI receives and sends data synchronously to the TTC clock. Although transmitters and receivers use the same clock frequency, the phase offset has to be extracted from the clock embedded in the received data for each of the 208 inputs. The phase offset for each of the 208 MUCTPI sector data inputs is different due to the length mismatch among the clock and data optical fibers, as well as the part-to-part skew of each of the sector logic module components. Therefore, for synchronizing the data from each of the muon SL modules, the MUCTPI has to compensate the input phase-skew with respect to the TTC clock and align the signals in multiples of the bunch-crossing period of 25 ns. Figure 8 shows the SL synchronization and alignment firmware. The write control logic block is used to detect the BC frame boundaries, and dual port memories are used to transfer each

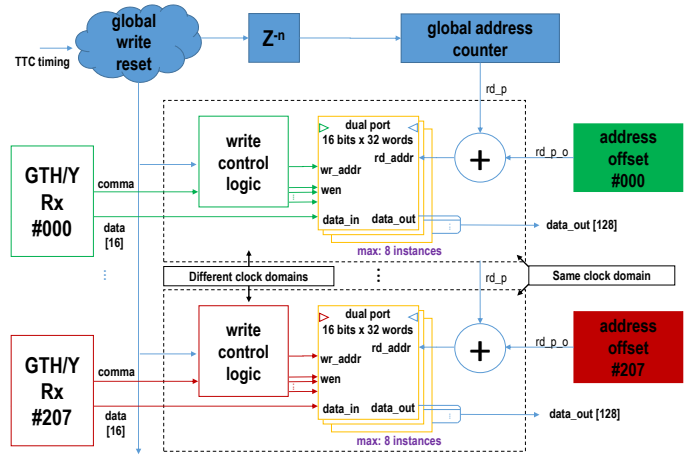


Fig. 8. MUCTPI sector data input synchronization and alignment circuit

input from their respective clock domains into a single clock domain for combined data processing. This circuit can cope with phase variation of the received data and non-deterministic data transfer latency of FPGA transceivers by monitoring received data timing and setting delay parameters in the write control logic. A reduced version of this firmware was already implemented in a pre-prototype module and tested successfully with the muon SL prototypes. The complete version of this firmware has been tested on the MUCTPI with 72 optical inputs coming from two different sources and will also be used during the upcoming integration tests with the muon SL modules.

VIII. CONCLUSION

The prototype of the upgraded MUCTPI has been tested successfully. All the features including power supply, clock circuitry, FPGA configuration, hardware monitoring, inter-FPGA communication, external optical connections, and the SoC were validated. Another two partially assembled prototypes are available for firmware and software development.

The MUCTPI serial link connections were also tested and excellent eye openings were measured for the serial links running at 6.4 Gb/s. Compatibility for future operation at 9.6 & 12.8 Gb/s was also confirmed. The TTC clock recovery has been tested and fixed-latency was measured. The muon SL data synchronization & alignment was tested successfully using 72 optical inputs from two different sources. Integration tests with the muon trigger sector logic modules will take place in the coming months. We plan to upgrade the MSP FPGAs to a pin compatible Ultrascale+ device (VU9P) for the next version of the MUCTPI in order to benefit from increased memory and logic resources.

REFERENCES

- [1] ATLAS Collaboration, "The ATLAS Experiment at the CERN Large Hadron Collider," *Journal of Instrumentation*, *JINST 3 S08003*, 2008.
- [2] S. Ask *et al.*, "The ATLAS Central Level-1 Trigger Logic and TTC System," *Journal of Instrumentation*, *JINST 3 P08002*, 2008.
- [3] S. Haas *et al.*, "The ATLAS Level-1 Muon to Central Trigger Processor Interface," in *Electronics for particle physics. Proceedings, Topical Workshop, TWEPP-07, Prague, Czech Republic, September 3-7, 2007*.

- [4] M. V. Silva Oliveira *et al.*, “The ATLAS Level-1 Muon Topological Trigger Information for Run 2 of the LHC,” *Journal of Instrumentation*, vol. 10, no. 02, p. C02027, 2015.
- [5] R. Caputo *et al.*, “Upgrade of the ATLAS Level-1 Trigger with an FPGA based Topological Processor,” in *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, 2013.
- [6] G. Aad *et al.*, *Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System*, CERN-LHCC-2013-018, 2013.
- [7] “Advanced TCA base specification,” *PICMG 3.0 Revision 3.0*, 2008.
- [8] “IEEE Standard for a Versatile Backplane Bus: VMEbus,” *ANSI/IEEE Std 1014-1987*, 1987.
- [9] P. Vichoudis *et al.*, “The ATLAS Muon-to-Central-Trigger Processor Interface,” in *Topical Workshop on Electronics for Particle Physics*, 2017, in preparation.