Topical Workshop on Electronics for Particle Physics
Glasgow, Scotland, UK
30 September - 4 October 2024

Development of methodology and implementation of SoC-based compact single-board validation system for the ATLAS Phase-II level-0 muon trigger system

Yoshifumi Narukawa, on behalf of the ATLAS TDAQ Collaboration

The University of Tokyo,
7-3-1 Hongo, Bunkyo-ku, Tokyo, Japan

E-mail: yoshifumi.narukawa@cern.ch

Abstract: Firmware testing on actual hardware is an optimal way to validate large-scale FPGA-based trigger/DAQ systems. For the Phase-Ⅱ level-0 muon trigger system at the ATLAS experiment, a methodology using prototype ATCA-based Sector Logic (SL) boards was developed, featuring self-complete DAQ, high-statistics test patterns, and various nature of input test data. The design exploits Zynq SoC on the board for control, injection of BRAM-based test patterns data and flexible DAQ to probe the process. The success is owing to the flexibility of the SL board design around the SoC device. This methodology facilitates precise firmware validation and systematic debugging.

Keywords: Trigger concepts and systems (hardware and software); Data acquisition concepts; Simulation methods and programs; Performance of High Energy Physics Detectors

ATL-DAQ-PROC-2024-016
21 October 2024

Contents

1 Introduction

In high-energy particle physics experiments, the Trigger and DAQ (TDAQ) system plays a crucial role in maximizing the overall performance of the experiment. With advancements in integrated circuit technology, TDAQ systems have become more sophisticated and complex in recent years. To ensure accurate implementation of the system and achieve optimal TDAQ performance, developing a detailed and comprehensive hardware-based validation system is essential. In response to these challenges, we have designed and implemented an MPSoC[1]-based validation system for the **T**hin **G**ap **C**hamber (TGC) system in the ATLAS experiment [1, 2] at HL-LHC [3]. This system allows the demonstration of the dedicated trigger algorithm using large statistical "physical" datasets, such as Monte Carlo simulation data, collision data, and toy data of emulating infinite-momentum strainght tracks.

The operation of the ATLAS detector at HL-LHC will begin in 2030 for precision measurements of the Standard Model and to search for new physics with high statistics. The peak luminosity will be tripled, with the first stage (Level-0, L0) trigger rate expanded to 1 MHz and the trigger latency extended to 10 µs [4]. To handle the increased collision rate and L0 trigger rate, the readout and trigger electronics of the TGC will be upgraded [5].

The TGC trigger electronics in HL-LHC will mainly consist of two front-end boards and one back-end board, as shown in Figure 1a. The charge signal generated in the TGC chamber is amplified, shaped, and discriminated on the **A**mplifier-**S**haper-**D**iscriminator (ASD) boards, which are mounted directly on the TGC detectors. The **P**atch-Panel ASIC and **S**ender FPGA (PS) board collects these signals from ASD boards, assigns rising edges to a corresponding bunch crossing, and sends a 256-bit hit-bitmap to the Endcap **S**ector **L**ogic (Endcap SL) every 25 ns. The Endcap SL plays three roles: triggering, readout, and control. For triggering, it reconstructs muon tracks and estimates their transverse momentum (p_{T}) in realtime through layer coincidences of TGC hits. For readout, it retrieves hit data and the results of the trigger calculation and conduct local event building for triggered events, and transmits this information via an optical link to the central readout system. For control, it distributes timing signals to the front-end board, and conduct electronics configuration and monitoring.

[1] In this study, a type of SoC with multiple processors (MPSoC) is adopted.

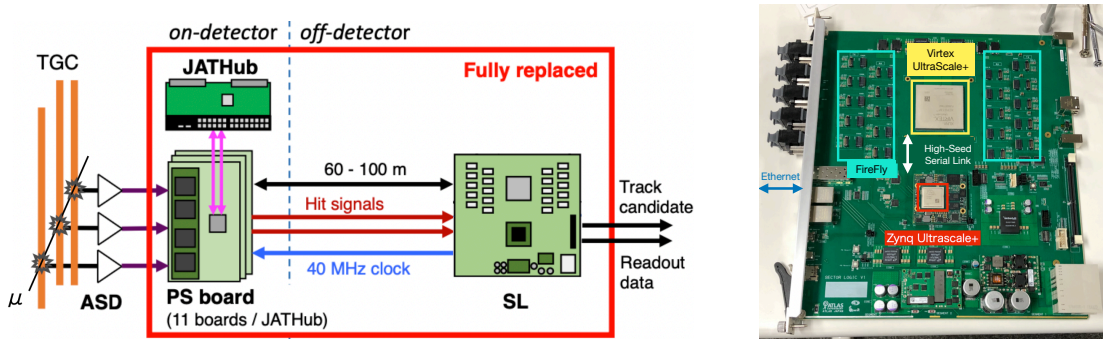

(a) (b)

Figure 1: Overview of the TGC Phase-II electronics system. 1a shows the schematics of the TGC electronics. 1b shows the picture of the Endcap SL board.

The SL is an ATCA blade equipped with high-bandwidth optical I/O (Samtec FireFly), a Virtex UltraScale+ FPGA (XCVU13P-1FLGA2577E), and a Zynq UltraScale+ MPSoC (XCZU5EV-2SFVC784I), as illustrated in Figure 1b. The FPGA plays a central role by receiving signals from the PS boards, calculating muon trigger candidates, and reading out hits and trigger information for events selected by the L0 trigger. This large-scale FPGA consists of four silicon chips and is equipped with high-performance multi-gigabit SERDES transceivers (GTY) distributed across 32 quad banks for efficient data transfer. The MPSoC interfaces with the TDAQ systems via Ethernet and serves as the control master for the Virtex UltraScale+ FPGA and the PS boards. It also includes high-performance multi-gigabit SERDES transceivers (GTH) for communication with the Virtex UltraScale+ FPGA. It provides flexible functionality through its processor system, which is leveraged for control and debugging, making it an essential component in the validation system.

2 L0 endcap muon trigger and its validation system

Trigger algorithm The TGC chambers consist of seven layers, which are grouped into three super-layers (stations): M1 (three layers), M2 (two layers), and M3 (two layers), located at $z = 13, 14, 14.5$ m, respectively (Figure 2a). Muons originating from the interaction point (IP) are deflected in the toroidal magnetic field, and their momentum is estimated using point-angle measurements. These measurements rely on the observables ($\Delta\phi$, $\Delta\theta$), defined as the difference between the reconstructed trajectories and the straight-line tracks connecting space points on the pivot station (M3) and the interaction point, as illustrated in Figure 2a. Look-up tables (LUTs) then map the ($\Delta\phi$, $\Delta\theta$) observables to an estimate of the transverse momentum.

The logic is implemented in the VU13P FPGA and structured as a sequence of three consecutive modules. The first module handles hit coincidences among two or three layers within each station. The output of this stage has a granularity of 1/2 or 1/3 of the channel size of the individual layers, taking advantage of the staggered layer structure within the stations. The second module performs segment reconstruction, which takes coincidences across the three stations and calculates the ($\Delta\phi$, $\Delta\theta$). The third module estimates p_{T} using LUTs.

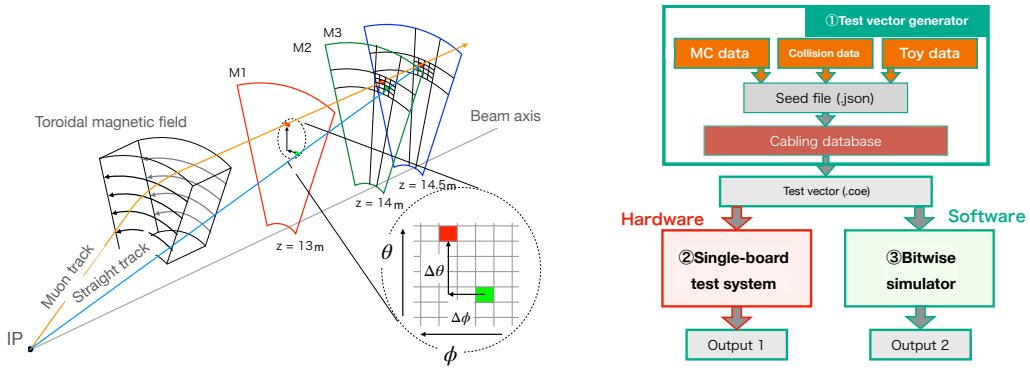

(a) (b)

Figure 2: Overview of the trigger algorithm and validation system. 2a is a conceptual illustration of the muon trigger algorithm. 2b shows the overall design of the trigger validation system.

Integrated validation framework To ensure that the trigger logic is accurately implemented in the FPGA, we have designed and implemented an integrated trigger validation system (Figure 2b). This system consists of three main tools: 1. test vector generator, 2. single-board test system, and 3. bitwise simulator.

The test vector generator is a software tool designed to generate pseudo-test input bitstreams. These test vectors are based on Monte Carlo simulation data of realistic muon trajectories or toy data that emulates infinite straight-line tracks. A relational database stores cabling information and the SL input bitstream data format, allowing the tool to map individual TGC hits to bit assertions in the pseudo input data. Regardless of the source of the test vectors, the tool can consistently generate input data patterns coherently for both hardware and software test benches. The single-board test system is an MPSoC-based validation platform that runs the trigger logic using test vectors as input on actual hardware. The input and output data are packaged into event data and recorded by a local readout system implemented within the MPSoC device. The bitwise simulator is a software tool that fully emulates the trigger logic with bitwise precision. The input and output data formats are designed to match exactly those of the hardware. By processing identical input data through the hardware and the simulator and comparing the consistency of their outputs, comprehensive, efficient, and detailed validation of the trigger system is achieved.

Implementation of single-board test system Figure 3 shows the schematic of the single-board test system. In addition to the main trigger logic, the system includes a timing controller to manage test pulse trigger and Level-0 Accept (L0A) signal, as well as test vector injectors that provide pseudo data synchronizteed with the test pulse trigger. The readout path has been modified to enable a single-board readout system via the MPSoC, which runs Linux and is directly connected to the programmable logic. In this setup, the MPSoC also acts as the master for control. Tests are conducted through the following steps.

1. The MPSoC writes the test vector into the BRAM equipped with the test vector injector.
2. The timing controller generates the test pulse trigger as a single clock tick pulse synchronized with the 40 MHz clock.

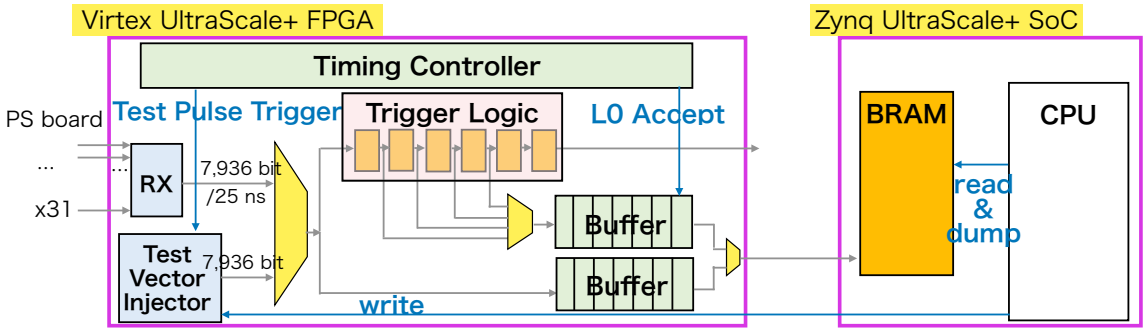

Figure 3: Block diagram of single-board test system

3. The test vector injector feeds test vectors, consisting of 7,936 bits, into the trigger logic.

4. The trigger output is stored in the buffer (L0 buffer). The timing controller generates the L0A signal, a single clock tick pulse, with a defined delay relative to the test pulse trigger, emulating the fixed latency of the L0 logic. Upon receiving the L0A, the corresponding data — including the output from the test vector injection – is transferred to the MPSoC and stored in the BRAM on the programmable logic. The software running on the MPSoC's processing system can then access the stored data in the BRAM for readout.

This test system is designed not only to process large datasets quickly but also to enable detailed debugging of the firmware. The timing controller emulates the role of the Central Trigger Processor and Trigger Timing Controller systems [1], generating timing signals such as the test pulse trigger and L0A signals as described above. The L0A signals are produced after a set latency period following the test pulse trigger, reflecting the L0 latency. The depth of the L0 buffer is fixed and configured to correctly read out the appropriate bunch crossing data, taking into account the latency of the trigger logic and the interval between the test pulse trigger and the L0A signals. The test vector injector stores 7,936-bit test vectors in the BRAM and injects them simultaneously as a single bunch crossing data, synchronized with the test pulse trigger signal. The BRAM has a depth of 60, corresponding to 60 events, and can be repeatedly updated by the MPSoC CPU, enabling high-statistics testing. Trigger data readout allows for verifying both final and intermediate results, primarily for debugging purposes. The trigger readout selector can switch between outputs from various trigger modules. The L0 buffer depth is configured with specific values to read data at each stage of the trigger calculation, enabling confirmation of correct data retrieval for each L0A. This setup also ensures the trigger logic functions with the expected latency at every step. For inter-chip communication, two bi-directional serial links (16.08 Gbps) are established: one for the control and the other for the readout functionalities.

3 Trigger performance evaluation using single-board test system

We have successfully utilized the single-board test system to evaluate the performance of the trigger logic, running a chain of algorithms on Monte Carlo simulation data for realistic muon trajectories. Among these, the trigger efficiency derived from the high-statistics Monte Carlo simulation data

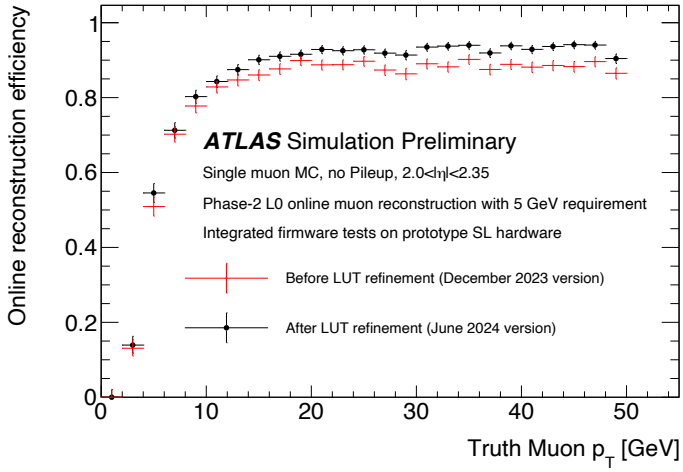

Figure 4: The trigger efficiency of the L0 endcap muon trigger [6]. This was measured using a single-board test system and a single-muon Monte Carlo data.

has been found to be a reliable indicator for assessing the quality of the latest firmware, as well as for validating and debugging it. Figure 4 (black points) shows the trigger efficiency as a function of the truth-level p_{T} of the muons in the pseudorapidity range of $2.0 < |\eta| < 2.35$, with a 5 GeV p_{T} threshold for 500,000 simulated muons. The plot indicates a high efficiency, as expected, with the efficiency about 94% at the plateau of the efficiency curve, demonstrating a good level of maturity in the trigger algorithm development.

Additionally, the figure shows performance improvements by comparing the online muon reconstruction efficiency between two different firmware versions used for refining the LUTs. Initially, it was observed that the efficiency was lower than the expected value from the simulation, which was approximately 94%. A detailed investigation of the test vector and comparison between expected outputs and observed outputs allowed us to identify the issues in the logic and LUTs. As this example shows, this test system is useful not only for checking the performance but also for debugging. We are fully utilizing this test system to realize the optimal L0 endcap muon trigger system.

4 Conclusion

To achieve the optimal TDAQ system, it is essential to establish a detailed and comprehensive validation system using actual hardware. One solution we have developed is the SoC-based single-board test system. This system allows us to process a sufficient amount and variety of datasets. By monitoring the trigger outputs or intermediate outputs and comparing them with input data or simulations, we can verify that the logic circuits are correctly implemented. We have completed the design and implementation of this test system for the Phase-II TGC Endcap SL and are now fully utilizing it in studies aimed at maximizing trigger performance. Moreover, the validation methods and implementation techniques used in this system are widely applicable to other electronic systems incorporating SoC devices.

References

[1] ATLAS Collaboration, The ATLAS Experiment at the CERN Large Hadron Collider, 2008 JINST 3 S08003, https://doi.org/10.1088/1748-0221/3/08/s08003

[2] ATLAS Collaboration, ATLAS muon spectrometer: Technical Design Report, CERN-LHCC-97-022, https://cds.cern.ch/record/331068

[3] I.B. Alonso, O. Brüning, P. Fessia, M. Lamont, L. Rossi, L. Tavian, M. Zerlauth, High-luminosity large hadron collider (HL-LHC): Technical design report, 2020, CERN Yellow Reports: Monographs, 10, http://dx.doi.org/10.23731/CYRM-2020-0010

[4] ATLAS Collaboration, Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System, CERN-LHCC-2017-020, https://cds.cern.ch/record/2285584

[5] ATLAS Collaboration, Technical Design Report for the Phase-II Upgrade of the ATLAS Muon Spectrometer, CERN-LHCC-2017-017, https://cds.cern.ch/record/2285580

[6] ATLAS Collaboration, ATLAS Experiment public result, July 2024, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/L0MuonTriggerPublicResults