

L1Topo: The Level-1 Topological Processor for ATLAS Phase-I upgrade and its firmware evolution for use within the Phase-II Global Trigger

Viacheslav Filimonov, on behalf of the ATLAS TDAQ Collaboration

Abstract—The increased instantaneous luminosity of the LHC in Run 3 brings the need for the upgrade of the ATLAS trigger system. The newly commissioned Phase-I L1Topo system, which replaces its Phase-0 predecessor, processes data from the Feature Extractors (FEXes) and the upgraded Muon to Central Trigger Processor Interface (MUCTPI) to perform topological and multiplicity triggers. The L1Topo system consists of three ATCA modules, each hosting two processor FPGAs (Xilinx Ultrascale+ 9P). The L1Topo firmware is composed of a large number of sort/select, decision, and multiplicity algorithms, that are automatically assembled and configured based on the provided trigger menu. For the HL-LHC, the Phase-I L1Topo system will be replaced by a Global Trigger, a time-multiplexed system, which concentrates the data of a full event into a single FPGA. In order to match the new operational environment, the fully synchronous, very low latency (new data arriving every 25 ns), parallel implementation (~ 2.5 M LUTs) of the Phase-I Topological firmware is being adapted to a significantly higher latency budget (new data arriving every 1.2 μ s) and a substantially tighter resource budget (~ 100 k LUTs). The main challenge is to allow for multiple working points of the utilized resources and latency for each algorithm. A detailed overview of the Phase-I L1Topo hardware and firmware is provided. Preliminary performance results achieved by the Phase-I L1Topo together with a description of the challenges found during the commissioning process are included. Phase-II related firmware adaptations are also discussed.

Index Terms—LHC, ATLAS, FPGA, Topological trigger

I. INTRODUCTION

THE Phase-I upgrade of the ATLAS detector was installed during the Long Shutdown 2, taking place between 2019 and 2022. The following Run 3 has started in 2022 and will continue until 2025 with a peak luminosity of $2 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ (Fig. 1).

The increase of the instantaneous luminosity of the LHC in Run 3 [2], brings the need for the upgrade of the ATLAS detector, including the trigger system.




Fig. 1. Upgrade program of the LHC accelerator complex [1].

A detailed block diagram of the Level-1 trigger system after the Phase-I upgrade is shown in Fig. 2. The Phase-I Level-1 trigger system performs real time event selection. It reduces the event rate from 40 MHz down to 100 kHz, allowing to stay below the maximum readout rate of the ATLAS detector. The overall system latency budget is 2.5 μ s.

As part of the Level-1 trigger system, the new Phase-I L1Topo system [3], which replaces its Phase-0 predecessor [4], processes data from the new jet [5], electromagnetic [6], and global [7] Feature Extractors (FEXes) and the upgraded Muon to Central Trigger Processor Interface (MUCTPI) [8] to perform topological triggers as well as triggers, counting the number of objects (multiplicity triggers). The upgraded L1Topo system provides higher processing capabilities in order to make use of the input objects with increased granularity from the new FEXes and the MUCTPI.




Fig. 2. A detailed block diagram of the Level-1 trigger system after the Phase-I upgrade [9].

Manuscript received May 17, 2024.
Viacheslav Filimonov is with Institut für Physik, Johannes Gutenberg Universität Mainz, Germany. Corresponding author is Viacheslav Filimonov (email: viacheslav.filimonov@cern.ch).

The Phase-II upgrade of the ATLAS detector will be installed during the Long Shutdown 3, taking place between 2026 and 2029. The following Runs 4 and 5 are currently planned to take place between 2029 and early 2040s with a peak luminosity around $7.5 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ (Fig. 1).

In order to cope with the significant increase of the instantaneous luminosity of the LHC in Run 4, a further upgrade of the ATLAS trigger system is necessary.

The Phase-II Level-0 trigger system (Figure 3) will make use of an increased latency budget of 10 μs and have a Level-0 accept rate of 1 MHz.




Fig. 3. A block diagram of the Level-0 trigger system after the Phase-II upgrade [10].

As part of the Phase-II Level-0 Trigger System, the Global Trigger will replace the Phase-I Topological Processor. The Global Trigger system will absorb the functions of the Phase-I Topological Processor and significantly extend them by using full granularity calorimeter cells to perform offline-like algorithms, identifying topological signatures, processing the trigger information from the Run 3 hardware systems and transmitting the processed trigger information to CTP for final decision.

II. L1TOPO HARDWARE OVERVIEW

The L1Topo system consists of 3 ATCA modules (Fig. 4), each hosting 2 processor FPGAs (Xilinx Ultrascale+ 9P [11]).

High-speed optical transceiver modules (Avago MiniPOD [12]) are used for the modules' real-time data path to support data transmission at speeds up to 11.2 Gb/s per link. 118 input and 24 output fibers are accommodated per each FPGA.

Xilinx Ultrascale+ Zynq based control mezzanine provides configuration, monitoring and slow control functionality.

In order to optimize the signal integrity for the high-speed signals between the FPGA and the optical modules as well as other high speed components, dedicated high-speed PCB design routing techniques were used. Strict physical and spacing constraints are created and controlled for the high-speed differential pairs. Crosstalk is minimized by ensuring a sufficiently large pair to pair spacing. Phase tuning is used to stay within the tolerance limit. Trace width and spacing within each of the differential pairs is controlled to achieve the

necessary differential impedance.




Fig. 4. L1Topo production module hardware overview.

Lyr	Image
✓SMT	Solder Mask
✓1	Cu Final [Thk:35um]
	Panasonic M6 R-5670 [Thk:70um]
✓2	Cu Final [Thk:25um]
	Panasonic M6 R-5670 [Thk:70um]
✓3	Cu Base [Thk: 17um]
	Panasonic M6 R-5775 [Thk:100um]
✓4	Cu Base [Thk: 17um]
	Panasonic M6 R-5670 [Thk:100um]
	Panasonic M6 R-5670
✓5	Cu Base [Thk: 17um]
	Panasonic M6 R-5775 [Thk:100um]
✓6	Cu Base [Thk: 17um]
	Panasonic M6 R-5670 [Thk:100um]
	Panasonic M6 R-5670
✓7	Cu Base [Thk: 17um]
	Panasonic M6 R-5775 [Thk:100um]
✓8	Cu Base [Thk: 17um]
	Isola PCL-370HR [Thk:150um]
	Isola PCL-370HR
✓9	Cu Base [Thk: 70um]
	Isola PCL-370HR [Thk:100um]
✓10	Cu Base [Thk: 70um]
	Isola PCL-370HR [Thk:150um]
	Isola PCL-370HR
✓11	Cu Base [Thk: 70um]
	Isola PCL-370HR [Thk:100um]
✓12	Cu Base [Thk: 70um]
	Isola PCL-370HR [Thk:150um]
	Isola PCL-370HR
✓13	Cu Base [Thk: 17um]
	Panasonic M6 R-5775 [Thk:100um]
✓14	Cu Base [Thk: 17um]
	Panasonic M6 R-5670 [Thk:100um]
	Panasonic M6 R-5670
✓15	Cu Base [Thk: 17um]
	Panasonic M6 R-5775 [Thk:100um]
✓16	Cu Base [Thk: 17um]
	Panasonic M6 R-5670 [Thk:100um]
	Panasonic M6 R-5670
✓17	Cu Base [Thk: 17um]
	Panasonic M6 R-5775 [Thk:100um]
✓18	Cu Base [Thk: 17um]
	Panasonic M6 R-5670 [Thk:70um]
✓19	Cu Final [Thk:25um]
	Panasonic M6 R-5670 [Thk:70um]
✓20	Cu Final [Thk:35um]
✓SMB	Solder Mask

Fig. 5. L1Topo PCB stack-up.

The PCB stack-up (Fig. 5) is also designed to support high-speed signals. A special PCB material (MEGTRON6 [13]) has good dissipation factor and dielectric constant for high frequencies. It is highly heat resistant and provides ultra-low transmission loss. Microvias are used for high-speed signals, which are routed on the top and bottom inner layers in order to avoid stubs. Additionally, high-speed signal layers are shielded by the ground layers to minimize the crosstalk.

III. FIRMWARE AND PERFORMANCE IN PHASE-I

The L1Topo firmware is synchronous to the LHC Bunch Crossing (BC) with every new event data arriving every 25 ns. It is composed of a large number of sort/select, decision, and multiplicity algorithms, that are automatically assembled and configured based on the provided trigger menu.

Select algorithms select all Trigger Objects (TOBs) passing configurable parameter-based threshold. Sort algorithms output a list of the leading TOBs with the highest Transverse Energy (ET), that pass the configurable parameter-based threshold, and sort it by ET (Fig. 6).

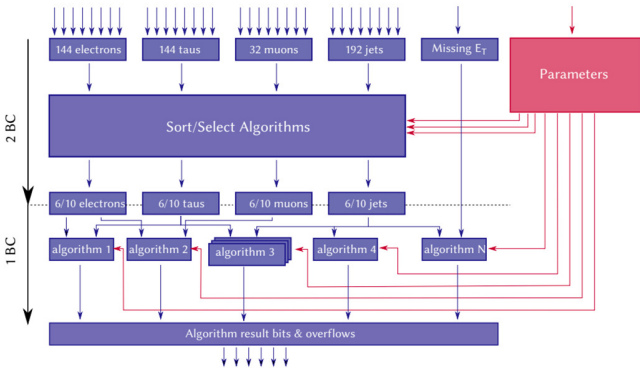


Fig. 6. L1Topo “Sort/Select” and “Decision” algorithms structure [4].

Decision algorithms perform calculations for one or more lists of TOBs, including angular differences, invariant masses, large jet reclustering, and missing transverse energy. Output decision bits, which indicate whether certain parameter-based trigger thresholds were passed, and overflow bits are both sent to the Central Trigger Processor.

Multiplicity algorithms perform non-trivial cuts on cross-dependent parameters and count the TOBs that pass (Fig. 7).

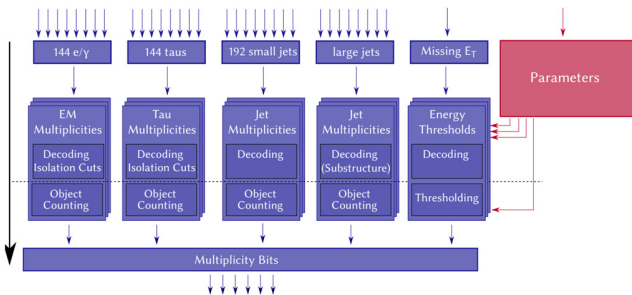


Fig. 7. L1Topo “Multiplicity” algorithms structure [4].

The latency budget for the algorithms is extremely tight. For example, decision algorithms only have 1 BC (25 ns) available.

Therefore, a full parallelization of the algorithms is required. Fig.8 shows an example implementation of one of the decision algorithms. The algorithm calculates the invariant mass and the $\Delta\phi$ for every combination of the input TOBs and applies corresponding thresholds. In case any combination satisfies the algorithm requirements, the trigger bit is fired. All the combinations are processed in parallel in just a single clock tick (25 ns). The price for the very low latency, however, is a very high resource usage: 2.5 million lookup tables (LUTs) across the 6 FPGAs of the L1Topo system.

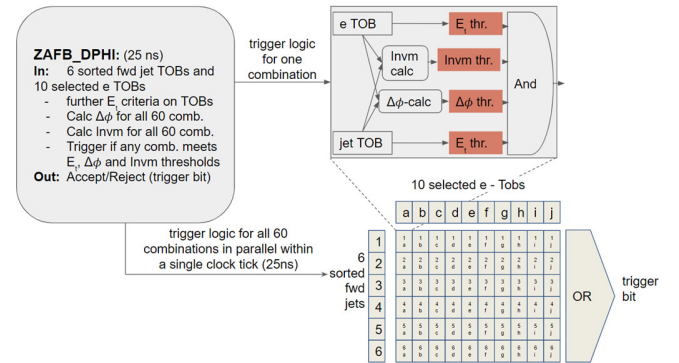


Fig. 8. Parallel implementation of the example decision algorithm [14].

As mentioned earlier, the algorithms are automatically assembled and configured based on the provided trigger menu. The algorithm parameters can be set and changed via the IPBus by the Online Software during a Run. As shown in Fig. 9, the topological trigger configuration is fully described in a single menu-driven json file, from which algorithm VHDL code, as well as IPbus address mapping, are automatically generated by a dedicated converter script. This ensures consistency between the firmware and the software.

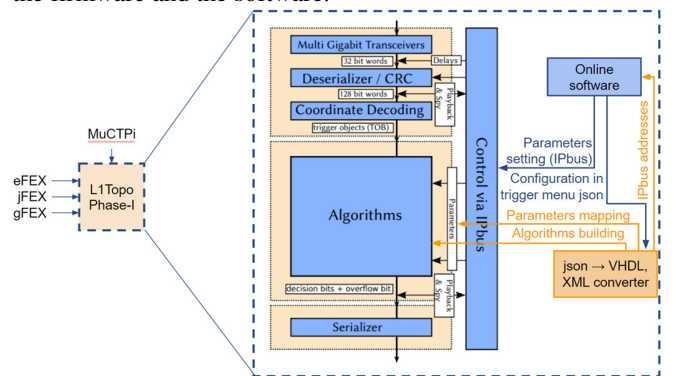


Fig. 9. Algorithm assembly and configuration.

The Phase-I L1Topo system has been fully commissioned with the rest of the new L1 trigger systems in ATLAS. The main commissioning challenges due to 4 different input sources included different input format of TOBs, different granularity of TOB coordinates, complicated detectors’ geometry and different time of TOBs’ readiness. Nevertheless, the Phase-I L1Topo system has come into routine operation taking data in 2024. All L1Calo and L1Muon triggers are going through L1Topo, making it a crucial element of the Level-1 trigger system. The L1Topo contribution can be clearly seen in the

example performance results in Fig. 10, where L1Topo chains provide about 70 % of unique rate for J/Ψ and Υ candidates.




Fig. 10. The mass spectra of $J/\Psi \rightarrow \mu^+\mu^-$ (left) and $\Upsilon \rightarrow \mu^+\mu^-$ (right) candidates reconstructed using the dedicated stream for B-physics and Light States from 2023 dataset [15].

IV. FIRMWARE ADAPTATION FOR PHASE-II

For the HL-LHC, the Phase-I L1Topo system will be replaced by a Global Trigger, a time-multiplexed system, which concentrates the data of a full event into a single FPGA. It is composed of three main layers: a Multiplexing (MUX) layer, a Global Event Processor (GEP) layer and a Demultiplexing layer (Global-to-CTP Interface (gCTPi)), which implements an interface to the Central Trigger Processor (CTP) (Fig. 11). This architecture provides a synchronous interface to the rest of the ATLAS detector.

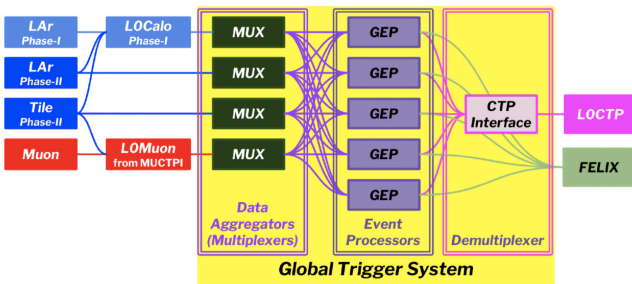


Fig. 11. Layers of the Global Trigger System [10].

MUX nodes within the Global Trigger receive data from LOCalo, Calorimeter and MuCTPI every BC and transmit a full event to a single GEP node every 49 BC (Fig. 12). As a result of this round-robin scheme, the per-event processing time on each GEP, i.e. the time until the next event arrives, is 1.2 μ s.




Fig. 12. Schematic view of the time multiplexing within the Global Trigger System [16].

In order to process the full event data, each GEP node will host an abundance of algorithms (Figure 13). To fit all these algorithms onto a single FPGA, the resource budget of the Hypothesis block (topological algorithms firmware block) in Phase-II is extremely tight with 100k LUTs.

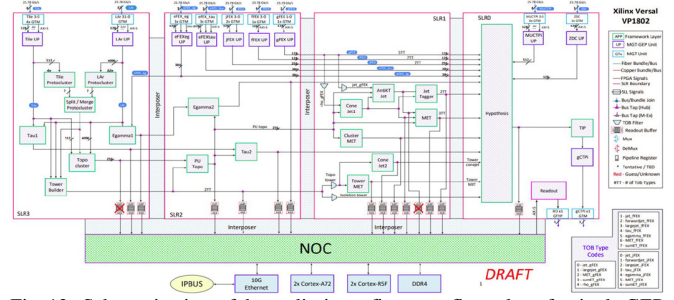


Fig. 13. Schematic view of the preliminary firmware floor plan of a single GEP node [16].

The main strategy within the Global Trigger operational environment is to fit within the tight resource budget at a cost of higher latency. Namely, instead of parallel processing of all TOB combinations in a single clock tick, as it is done in Phase-I L1Topo operational environment, sequential processing is implemented. This leads to a significant resource reduction.

Fig. 14 shows an example implementation of the same decision algorithm as used in the example in Fig. 8. However, in this case only a single logic block, which performs the required calculations and applies thresholds, is implemented. All the input TOB combinations are sequentially fed into the logic block. The functionality remains the same. In case any of the TOB combinations satisfies the algorithm requirements the trigger will be fired. This implementation leads to a significant resource reduction. 31732 LUTs required to implement the algorithm in parallel are reduced down to 636 LUTs with the sequential implementation. The price to pay, however, is a higher latency. Namely, more than 60 clock sub-ticks are required in this case. Nevertheless, this implementation fits nicely within the Global Trigger resource and latency budget. It should be also mentioned that the clock frequency can be increased in order to reduce the algorithm latency, as shown in the example.

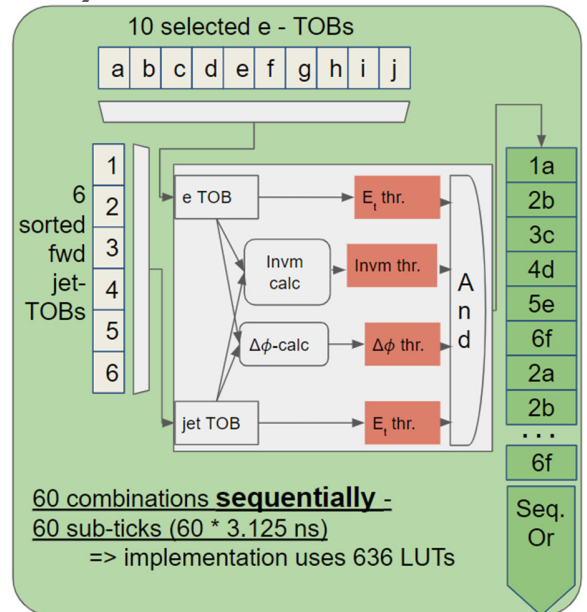


Fig. 14. Sequential implementation of the example decision algorithm [14].

A. Select algorithms

Similar to the Phase-I L1Topo firmware implementation, in order to stay within the acceptable latency and resource budget, the number of input TOBs per type needs to be limited for decision algorithms. Therefore, select algorithms select all TOBs, passing configurable parameter-based thresholds. Select algorithms within the Global Trigger environment are implemented sequentially. As shown in Fig. 15, the core of the select algorithm is a so-called selector block, which receives an input TOB, converted to a GenericTOB format (an internal common format for the Hypothesis block). The thresholds can be configured by a dedicated parameter record. The output of the selector block indicates with a single bit whether the input TOB has passed the selection. In case the input TOB passes the selection, it is reduced down to a ReducedTOB format (keeping only the bits required by the downstream algorithms) and forwarded further down the trigger chain. Otherwise, an EmptyTOB will be forwarded.

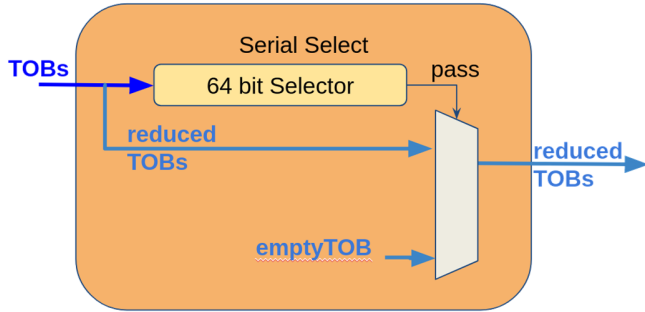


Fig. 15. Sequential implementation of the select algorithm [17].

B. Multiplicity algorithms

Adding a dedicated counter after the selector block creates the logic for a multiplicity algorithm (Fig. 16). Counter width can be flexibly configured.




Fig. 16. Sequential implementation of the multiplicity algorithm [17].

C. Sort algorithms

Sequential implementation of the sort algorithms consists of sorting stages, the number of which corresponds to the number of input ReducedTOBs to be sorted (Fig. 17). Within the each sorting stage the ET of the input ReducedTOB is compared to the ET of the currently stored ReducedTOB. In case the ET of the input ReducedTOB is higher it will be stored in the current stage and the previously stored ReducedTOB will be forwarded to the next stage and compared there. After the last ReducedTOB has left the last sorting stage, the stored leading ReducedTOBs are piped out.




Fig. 17. Sequential implementation of the sort algorithm [17].

D. Multiple working points

Since a significant difference in arrival times of the various input objects is expected it is essential to allow for multiple working points of the utilized resources and latency for each algorithm. For example, the implemented algorithms are able to double the amount of logic resources in order to halve the latency, if needed (Fig. 18). This feature allows to flexibly tune the utilized resources versus latency and helps to meet the overall latency budget.

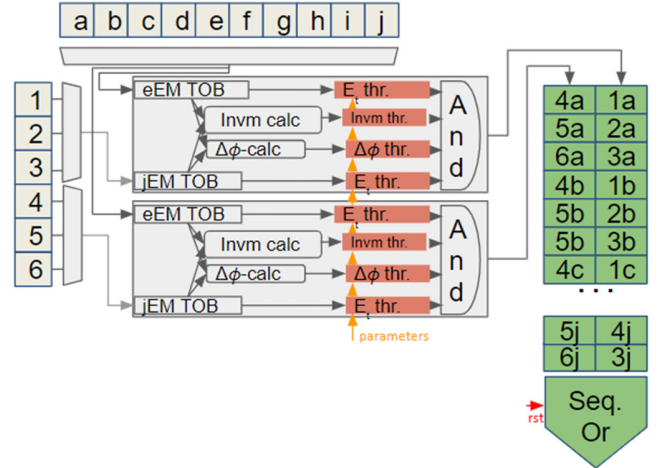


Fig. 18. Semi-sequential implementation of the example decision algorithm with double amount of logic resources in order to halve the latency [17].

E. Generic decision algorithms

Since most of the decision algorithms are composed of the same building blocks, it is possible to implement a so-called generic decision algorithm (Fig. 19). This algorithm includes all the necessary calculations, such as $\Delta\eta$, $\Delta\phi$, ΔR^2 and invariant mass, and can be configured to accept one or two input lists of TOBs. What selectors are used is also configurable.




Fig. 19. Semi-sequential implementation of the example decision algorithm with double amount of logic resources in order to halve the latency [17].

V. BEHAVIORAL SIMULATION

As mentioned earlier, the Phase-II Global Trigger Hypothesis algorithm structure is extremely similar to the algorithm structure within the Phase-I L1Topo. Therefore, the well tested Phase-I L1Topo algorithm blocks are serialized and reused.

In order to verify the functionality of the serialized algorithm blocks versus the original parallel implementation, a dedicated behavioral simulation has been performed using a Phase-I L1Topo simulation environment adapted for the Hypothesis verification.

As an example, Fig. 20 shows the waveform of the serialized decision algorithm (InvMDrSqrIncl2) simulation. The TOBs from the two input lists are serially provided to the functional algorithm blocks that perform calculations and apply thresholds

for one combination at a time. The output is a trigger bit, which is either accept or reject. In this case events number 2 and 5 have triggered. The input TOBs are generated and the decision parameters are set by the simulation environment.

Fig. 21 shows the simulation waveform of the original decision algorithm (InvMDrSqrIncl2), implemented in parallel. The same input TOBs and the decision parameters as for simulation of the serialized implementation are used. Events number 2 and 5 have triggered as expected.




Fig. 20. The waveform of the serialized decision algorithm simulation [17].



Fig. 21. The simulation waveform of the original decision algorithm, implemented in parallel [17].

VI. IMPLEMENTATION

The first implementation of the Hypothesis algorithms has been performed with serialized components assembled according to a certain trigger menu. Fig. 22 shows the implementation result. This test implementation requires 56k LUTs, which comfortably fits within the 100k LUTs budget. With 300 MHz clock timing closure is met. Xilinx VU9P device is the target device in this case. Decision algorithms are marked with yellow, Select and sort algorithms with cyan and multiplicity algorithms with purple.

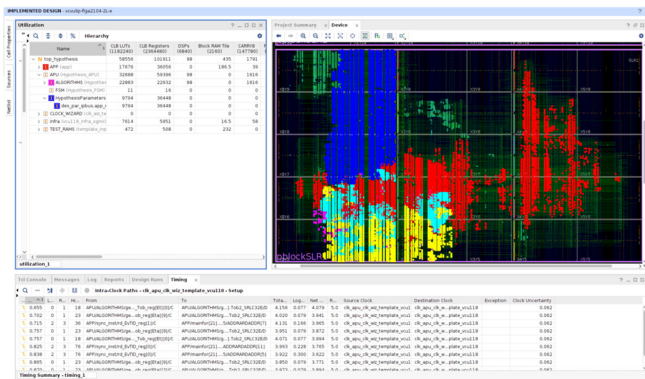


Fig. 22. Hypothesis algorithms implementation result [17].

VII. CONCLUSION

The new Phase-I L1Topo system hardware and firmware has been developed in order to process the data from the FEXes and the MUCTPI. Performing topological and multiplicity triggers, L1Topo is an essential component within the Level-1 trigger system. Preliminary performance results achieved by the Phase-

I Level-1 Topological trigger system together with a description of the challenges found during the commissioning process are included.

The fully synchronous, very low latency (new data arriving every 25 ns), parallel implementation (~ 2.5 m LUTs) of the Phase-I Topological firmware is being adapted to the Global Trigger operational environment with a significantly higher latency budget (new data arriving every 1.2 μ s) and a substantially tighter resource budget (~ 100 k LUTs). Namely, serialization of the algorithms, currently implemented in parallel, is ongoing. The key feature allows for multiple working points of the utilized resources and latency for each algorithm – an essential requirement, since a significant difference in arrival times of the various input objects is expected. The functionality of the serialized algorithms is verified with dedicated behavioral simulations. The first implementation fits comfortably within the allocated resource budget and meets timing.

REFERENCES

- [1] LHC / HL-LHC Plan. Last update February 2022. URL: <https://hilumilhc.web.cern.ch/content/hl-lhc-project>
- [2] S. Fartoukh et al., “LHC Configuration and Operational Scenario for Run 3”, CERN-ACC-2021-0007
- [3] L1Calo Phase-I L1Topo specifications, tech. rep., url: <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/LevelOneCaloUpgradeModules>
- [4] J. Damp, “Search for Dijet Resonances with the Level-1 Topological Processor at ATLAS”, PhD thesis: Johannes Gutenberg-Universität Mainz, 2020
- [5] M. Weirich, “Development of New ATLAS Trigger Algorithms in Search for New Physics at the LHC”, PhD thesis: Johannes Gutenberg-Universität Mainz, 2021
- [6] L1Calo Phase-I eFEX specifications, tech. rep., url: <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/LevelOneCaloUpgradeModules>
- [7] M. Begel et al., “Global Feature Extractor of the Level-1 Calorimeter Trigger: ATLAS TDAQ Phase-I Upgrade gFEX Final Design Report”, ATL-COM-DAQ-2016-184
- [8] R. Spiwojs et al., “The ATLAS Muon-to-Central Trigger Processor Interface (MUCTPI) Upgrade”, ATL-DAQ-PROC-2017-013
- [9] ATLAS Collaboration, “Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System”, CERN-LHCC-2013-018
- [10] ATLAS Collaboration, “Technical Design Report for the Phase II Upgrade of the ATLAS TDAQ System”, CERN-LHCC-2017-020
- [11] Xilinx, UltraScale Architecture and Product Data Sheet: Overview, url: https://www.xilinx.com/support/documentation/data_sheets/ds890-ultrascale-overview.pdf
- [12] “MiniPOD AFBR-814xyZ, AFR-824VxyZ, 14 Gbps/Channel Twelve Channel, Parallel Fiber Optics Modules”, Tech. Rep. AV02-4039EN, Avago Technologies, 2013
- [13] Megtron 6, Ultra-low Loss, Highly Heat Resistant Circuit Board Material, url: <https://industrial.panasonic.com/ww/products/pt/megtron6>
- [14] E. Meuser, “The ATLAS Level-1 Topological Processor”, url: <https://cds.cern.ch/record/2869237>
- [15] ATLAS Experiment – Public Results, url: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/BPhysicsTriggerPublicResults>
- [16] ATLAS TDAQ Phase-II Upgrade: Firmware Specifications for the Global Trigger, tech. rep., url: <https://edms.cern.ch/document/2677532/1>
- [17] ATLAS TDAQ Phase-II Upgrade: Global Event Processor Hypothesis Algorithms Specification, tech rep., url: https://edms.cern.ch/ui/file/2856352/1/Hypothesis_0.5.pdf