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

Abstract— The increased instantaneous luminosity of the Large Hadron Collider (LHC) in Run 3 brings the need for the upgrade of the A Toroidal LHC Apparatus (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 field programmable gate arrays (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 high-luminosity LHC (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 [∼2.5M look-up tables (LUTs)] of the Phase-I topological firmware is being adapted to a significantly higher latency budget (new data arriving every 1.2 μs) and a substantially tighter resource budget (∼100k 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— A Toroidal LHC Apparatus (ATLAS), field programmable gate array (FPGA), Large Hadron Collider (LHC), topological trigger.

I. INTRODUCTION

THE Phase-I upgrade of the A Toroidal LHC Apparatus (ATLAS) detector was installed during the Long Shutdown 2, taking place between 2019 and 2022. The current Run 3 has started in 2022 and will continue until 2025 with a peak luminosity of 2×10^{34} cm^{-2}s^{-1}.

The increase of the instantaneous luminosity of the large hadron collider (LHC) in Run 3 [1], brings the need for the upgrade of the ATLAS detector, including the trigger system.
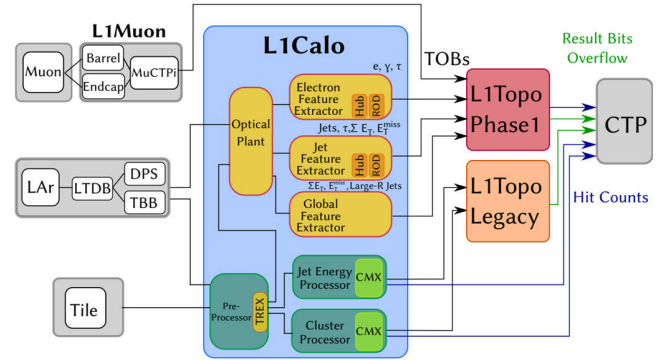
Fig. 1. Detailed block diagram of the Level-1 trigger system after the Phase-I upgrade [3].

Section II provides a detailed description of the Level-1 trigger system after the Phase-I upgrade as well as one of its new parts, the Phase-I L1Topo system [2]. Section III focuses on the hardware of the Phase-I L1Topo system. L1Topo firmware and performance in Phase-I are described in Section IV.

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 7.5×10^{34} cm^{-2}s^{-1}.

Section V provides a detailed description of the Level-0 trigger system after the Phase-II upgrade as well as one of its new parts, the global trigger system. Section VI focuses on the adaptation of the Phase-I L1Topo firmware to the global trigger operational environment. Behavioral simulation and implementation results are discussed in Sections VII and VIII, respectively.

II. LEVEL-1 TRIGGER SYSTEM AFTER THE PHASE-I UPGRADE

A detailed block diagram of the Level-1 trigger system after the Phase-I upgrade [3] is shown in Fig. 1. 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 μs.

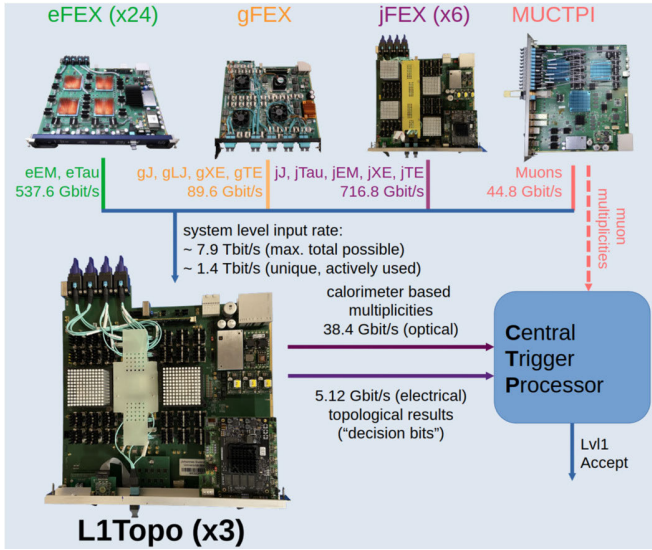As part of the Level-1 trigger system, the new Phase-I L1Topo system, which replaces its Phase-0 predecessor [3],

Fig. 2. L1Topo production module hardware overview, including the input sources and outputs [8].

TOB type	Fibres	Copies	# TOBs	Name
Muon TOBs	6	4	8	MU
eEM TOBs (first fibre)	24	8 (4)	6	EM1
eEM TOBs (second fibre)	24	8 (4)	6	EM2
eTau TOBs (first fibre)	24	8 (4)	6	TAU1
eTau TOBs (second fibre)	24	8 (4)	6	TAU2
jJ, jLJ, jTAU, jEM TOBs (first fibre)	24	4	7	JET1
jJ, jLJ, jTAU, jEM TOBs (second fibre)	24	4	7	JET2
jFEX Missing/Total Energy (plus forward jJ, jLJ, jTAU, jEM)	12	4	XE+5	JFXE
Forward jJ, jLJ, jTAU, jEM TOBs from high-η jFEX	4	4	7	JF
gFEX jet TOBs (gJ, gLJ) and global quantities	8	4 (3)	6 (jets)	GFEX

"Copies" column shows the number of copies available at the FEX outputs, while the copies that are actually routed via the optical plant are shown in brackets.
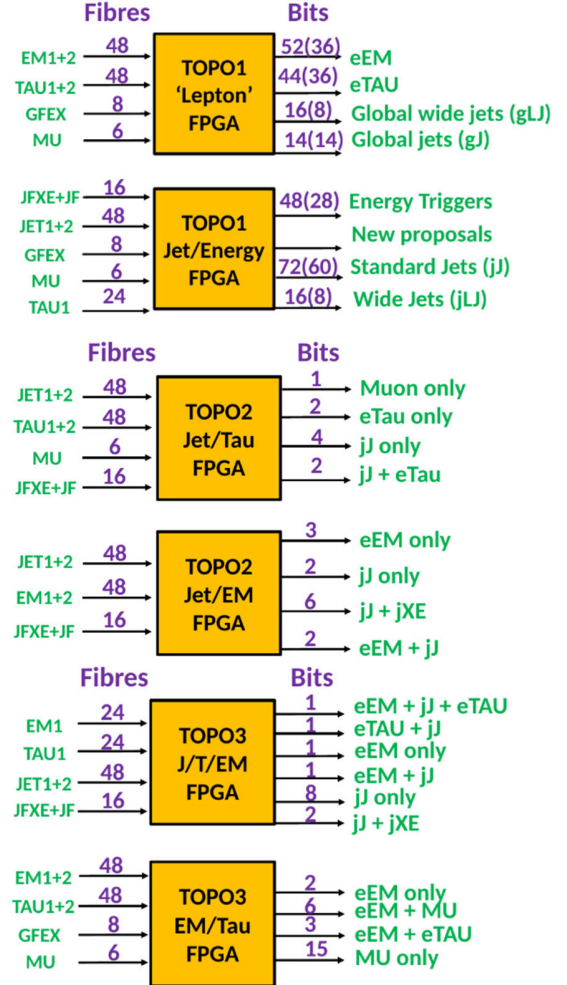

Fig. 3. Connectivity details of all L1Topo modules [11].

processes data from the new jet [4], electromagnetic [5], and global [6] feature extractors (FEXes) and the upgraded muon to central trigger processor interface (MUCTPI) [7] 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.

In order to guarantee a stable physics data taking during the commissioning period of the upgraded L1Topo system the legacy L1Topo system was kept operational. Currently, a full transition from the legacy to the upgraded system has taken place. All the Phase-I-related descriptions and results below correspond to the upgraded L1Topo system.

III. L1Topo HARDWARE OVERVIEW

The L1Topo system consists of three ATCA modules (Fig. 2) [8], each hosting two processor field programmable gate arrays (FPGAs) (Xilinx Ultrascale + 9P [9]).

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

High-speed optical transceiver modules (Avago Mini-POD [10]) are used for the modules' real-time data path to support data transmission at speeds up to 11.2 Gbit/s per link. The maximum possible total system level input rate is ~7.9 Tbit/s, while the unique, actively used, total data rate is ~1.4 Tbit/s.

The FEXes and the MUCTPI ATCA modules provide trigger objects (TOBs) to L1Topo. TOBs contain information on energy and position. 24 eFEX modules provide eTau and eEM objects that are sent on separate fibers, each type using two fibers. Six jFEX modules provide transverse energy (ET) sums and Jet objects. A single gFEX module provides energy information and complementary Jet objects. A MUCTPI module provides Muon TOBs.

Table I provides the summary of the type and multiplicity of the L1Topo inputs, while Fig. 3 shows the connectivity details of all three L1Topo modules [11].

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 printed circuit board (PCB) design routing techniques were used. Strict physical and spacing constraints were 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.

The PCB stack-up is also designed to support high-speed signals. A special PCB material (MEGTRON6 [12]) has good
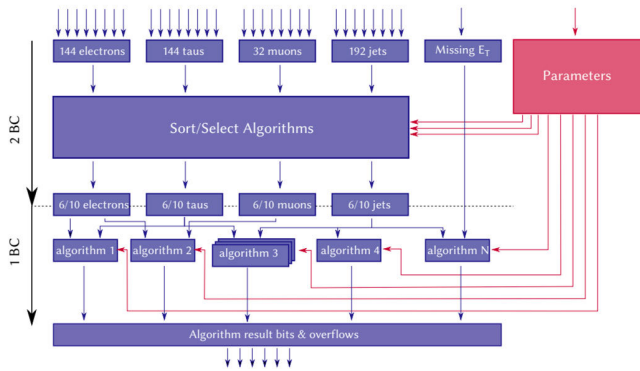
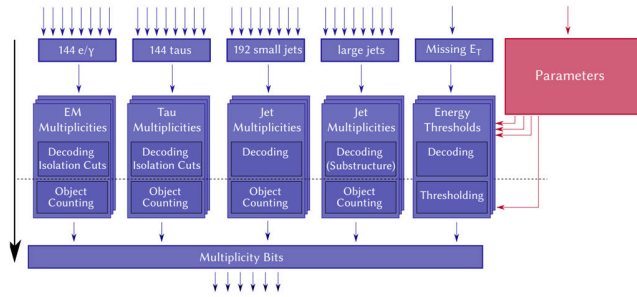Fig. 4. L1Topo "sort/select" and "decision" algorithms structure [15].

Fig. 5. L1Topo "multiplicity" algorithms structure [15].

Fig. 6. Parallel implementation of the example decision algorithm [16].

Fig. 7. Algorithm assembly and configuration.

dissipation factor and dielectric constant for high frequencies. It is highly heat resistant and provides ultralow 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.

The L1Topo modules have been successfully produced and their functionality has been fully tested [13].

IV. FIRMWARE AND PERFORMANCE IN PHASE-I

The L1Topo firmware is synchronous to the LHC bunch crossing (BC) with new event data arriving every 25 ns [14]. 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 [15]. Most of the algorithms run at a 40 MHz clock, apart from a few select algorithms that use a 160 MHz clock due to a high number of input TOBs.

Select algorithms select all TOBs passing a configurable parameter-based threshold. Sort algorithms output a list of the leading TOBs with the highest ET that pass the configurable parameter-based threshold and sort it by ET (Fig. 4).

Decision algorithms perform calculations for one or more lists of TOBs, including angular differences, invariant masses, large jet reclustering, and missing ET. Output decision bits, which indicate whether certain parameter-based trigger thresholds were passed, and overflow bits are both sent to the central trigger processor (CTP).

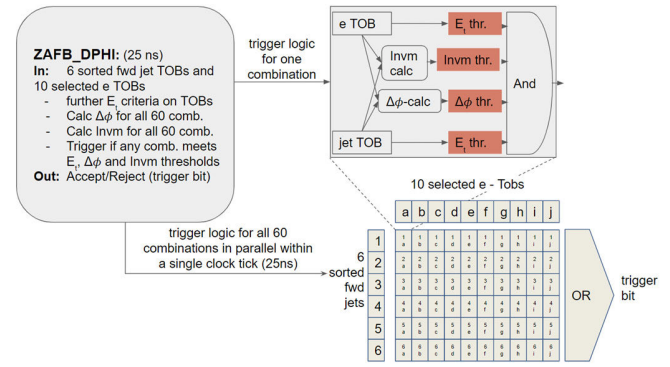Multiplicity algorithms perform nontrivial cuts on cross-dependent parameters and count the TOBs that pass (Fig. 5).
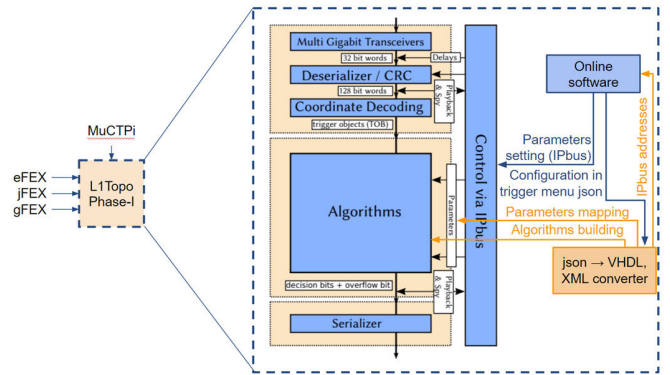
Multiplicity algorithms are implemented on two FPGAs of the L1Topo system (TOPO1), while the rest of the FPGAs perform the sort/select and decision algorithms (TOPO2 and TOPO3).

The latency budget for the algorithms is extremely tight. For example, decision algorithms only have one BC (25 ns) available. Therefore, a full parallelization of the algorithms is required. Fig. 6 shows an example implementation of one of the decision algorithms [16]. 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 six FPGAs of the L1Topo system.

As mentioned earlier, the algorithms are automatically assembled and configured based on the provided trigger menu. The inputs distribution is fixed and defined by the trigger menu. A new trigger menu configuration requires a new firmware implementation. The algorithm parameters, however, can be set and changed via the IPBus by the Online Software during a Run, without the need to rebuild the firmware. As shown in Fig. 7, 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.
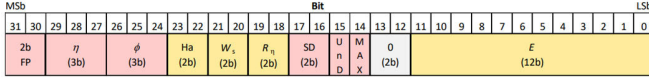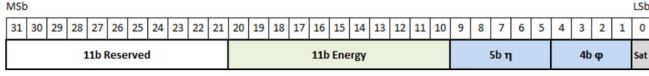
Fig. 8.   eFEX EM TOB format [17].

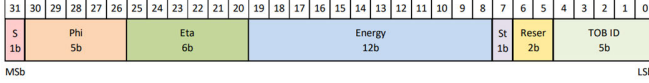Fig. 9.   jFEX Small-R Jet TOB format [17].

Fig. 10.   gFEX Jet TOB format [17].

TABLE II
FIELDS OF THE eFEX EM TOB (ADAPTED FROM [17])

Label	No. bits	Description
E	12	12-bit cluster energy in units of 100 MeV.
0	Various	An unused field, in which the bits are set to zero.
MAX	1	Set if the seed supercell is a local maximum (i.e. there is no more energetic supercell in the tower).
UnD	1	UpnotDown: set if the seed includes the supercell above (higher ϕ) rather than below the central supercell of the seed. e/γ & Baseline τ TOBs only.
SD	2	Seed η position within TOB trigger tower, in units of 0.025. e/γ & Baseline τ TOBs only.
R_η	2	Count of the number of R_η thresholds satisfied. e/γ TOB only.
W_s	2	Count of the number of W_s thresholds satisfied. e/γ TOB only.
Ha	2	Count of the number of Hadronic thresholds satisfied. Thresholds are assumed to be in increasing order. e/γ TOB only.
ϕ	3	ϕ position of the TOB trigger tower relative to the algorithm origin, in units of $\pi/32$ (≈ 0.1)
η	3	η position of the TOB trigger tower relative to the FPGA algorithm origin, in units of 0.1. Value = 0–5.
FP	2	FPGA Number 0-3 within the eFEX module. FPGAs are ordered in increasing η. Each FPGA covers an η region of 0.4

TABLE III
FIELDS OF THE jFEX SMALL-R JET TOB [17]

Label	No. bits	Description
Sat	1	Saturation bit on the jFEX input
φ	4	Local φ coordinate
η	5	Local η coordinate
Energy	11	Transverse energy, 200 MeV per bit
Reserved	11	Reserved bits

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 four different input sources included different input format of TOBs, different granularity of TOB coordinates and complicated detectors' geometry. Figs. 8–10 show examples of the TOB formats, received by L1Topo. Tables II–IV provide additional details regarding the fields of the TOBs.

The eFEX EM TOB format has energy information transmitted in units of 100 MeV. The local coordinates are in units of 0.1 for η and ϕ, indicating the position of the TOB trigger tower relative to the algorithm origin [17].

The jFEX Small-R Jet TOB format has energy information transmitted with an LSB of 200 MeV, which allows for energy of up to 409.6 GeV. The 5-bit η and 4-bit ϕ are local bin numbers [17].

TABLE IV
FIELDS OF THE gFEX JET TOB [17]

Label	No. Bits	Description
TOB ID	5	TOB identifier, according to the following: 00001: leading gBlock in η bins (0–5). 00010: leading gBlock in η bins (6–11). 00011: second leading gBlock in η bins (0–5). 00100: second leading gBlock in η bins (6–11). 00101: leading large-R jet in η bins (0–5). 00110: leading large-R jet in η bins (6–11).
Reser	2	Reserved
St	1	Status of jet. For gBlocks and gJets, set to 1 if E_T exceeds threshold. For ρ TOBs, set to 1 if calculation is valid.
Energy	12	E_T (12-b): 0–819 GeV in 200 MeV steps. This is the same for gBlock, gJet, and ρ TOBs.
Eta	6	η position of the TOB. Set to zero for the ρ TOB.
Phi	5	ϕ position of the TOB. Set to zero for the ρ TOB.
S	1	Saturated: indicates energy saturation in the TOB energy (gBlock, gJet). Always 0 for ρ since if an individual gTower saturates, it is not used for calculating ρ.

The gFEX Jet TOB format has energy information transmitted with an LSB of 200 MeV, which allows for energy of up to 819.2 GeV. In the gFEX case ϕ granularity is η dependent, with larger ϕ intervals in the forward region [17].

Within the L1Topo algorithm firmware, however, it is essential to have all the objects in the same format in order to be able to perform their matching and comparisons. Therefore, a detailed understanding of local η/ϕ coordinates with different granularities is needed in order to translate them into global coordinates. More details can be found in [15].

Despite all the challenges, 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. 11, showing the mass spectra of $J/\psi \rightarrow \mu^+\mu^-$ and $\Upsilon \rightarrow \mu^+\mu^-$ candidates reconstructed using the dedicated stream for B-physics and light states from 2023 dataset. As can be seen from the results, L1Topo chains provide about 70% of unique rate for J/Ψ and Υ candidates. Additional requirements made by L1Topo algorithms include restrictions on invariant mass and opening angle. In addition, the bottom plot includes a requirement on Barrel-Only L1 muon objects [18].

V. LEVEL-0 TRIGGER SYSTEM AFTER THE PHASE-II UPGRADE

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 (Fig. 12) will make use of an increased latency budget of 10 μs, compared to 2.5 μs in Phase-I, and have a Level-0 accept rate of 1 MHz, compared to 100 kHz in Phase-I [19].

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
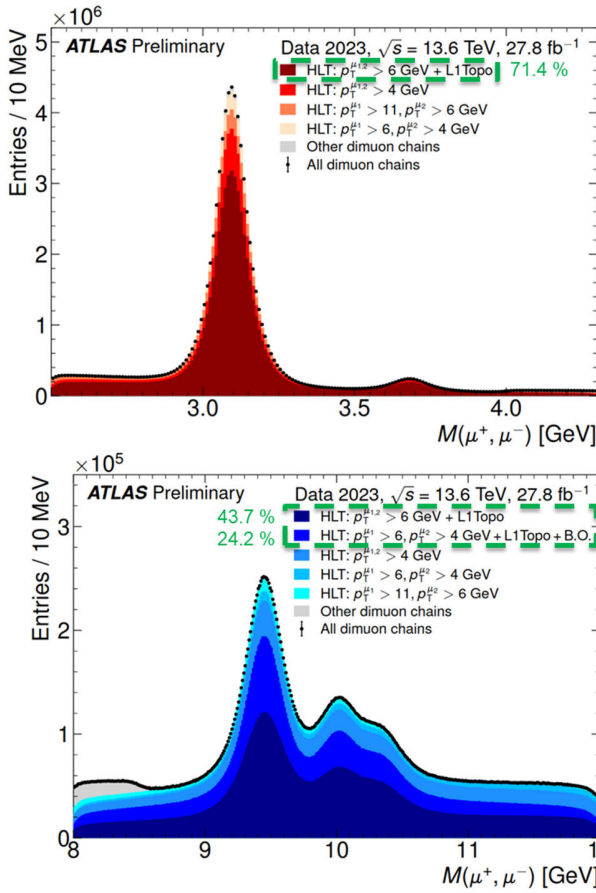
Fig. 11. Mass spectra of (top) and (bottom) candidates reconstructed using the dedicated stream for B-physics and light states from 2023 dataset [18].
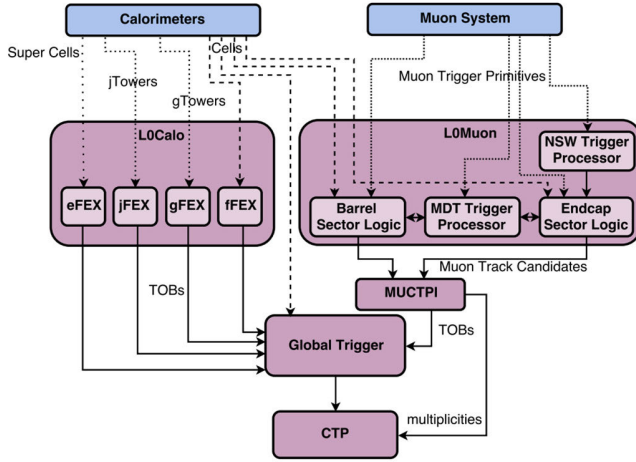

Fig. 12. Block diagram of the Level-0 trigger system after the Phase-II upgrade [19].

transmitting the processed trigger information to CTP for final decision.

VI. New Hardware and Firmware Adaptation for Phase-II

Global trigger is 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
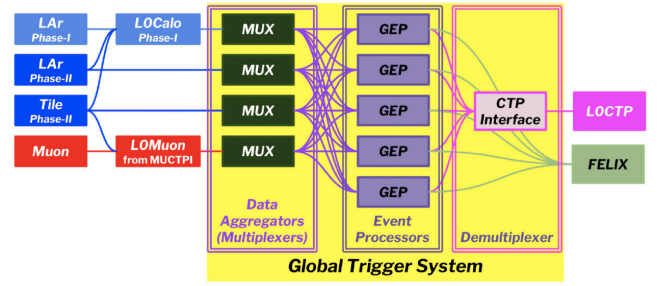

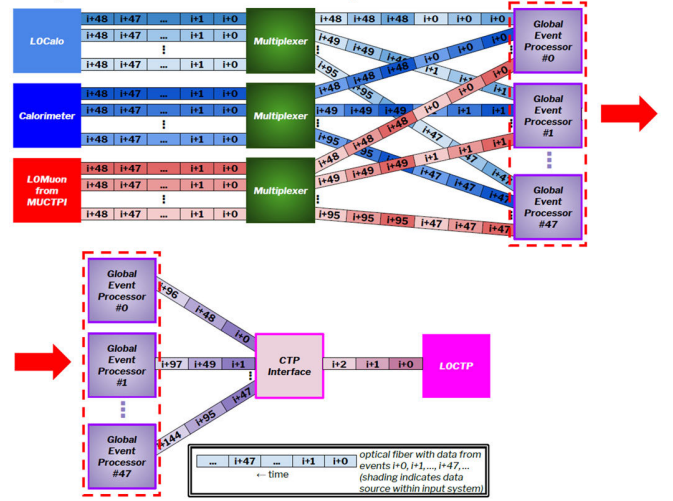Fig. 13. Layers of the global trigger system [19].

Fig. 14. Schematic view of the time MUX within the global trigger system [21].

layer [global-to-CTP interface (gCTPi)], which implements an interface to the CTP (Fig. 13). This architecture provides a synchronous interface to the rest of the ATLAS detector, while decoupling the GEP layer from the LHC BC rate, allowing complex asynchronous algorithms that can emulate more closely the event filter trigger processing.

The main hardware element of the global trigger is the global common module (GCM) [20]. GCM modules compose every layer of the global trigger. A single GCM module has a standard ATCA form factor and hosts two Xilinx Versal Premium XCVP1802 devices along with 20 Samtec FireFly optical modules. Due to hardware design limitations (thermal and space), only a fixed configuration of the two FPGAs on the GCM is supported. Thus, one FPGA is configured as a MUX node and the other as either a GEP or a gCTPi node.

All the MUX nodes within the global trigger receive the event data 0 from L0Calo, calorimeter, and MuCTPi in an LHC-synchronous manner and transmit it to GEP node 0. Similarly, in the next BC ($\sim$25 ns), the event 1 is transmitted to the GEP node 1 and so on until GEP node 47 and afterward back to the GEP node 0 (Fig. 14). As a result of this round-robin scheme, the per-event processing time on each GEP node, i.e., the time until the next event arrives on a particular GEP node, is 48 $\times$ 25 ns = 1.2 μs [21]. The input TOB format for Phase-II is still being defined.

Bringing event filter-like capability to the Level-0 trigger system is the driving design goal of the global trigger. In contrast to previous hardware trigger upgrades, the main emphasis
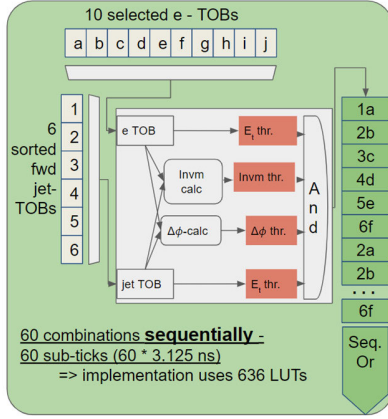
Fig. 15. Sequential implementation of the example decision algorithm [16].
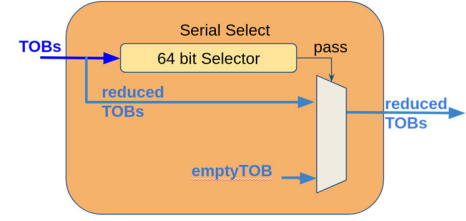

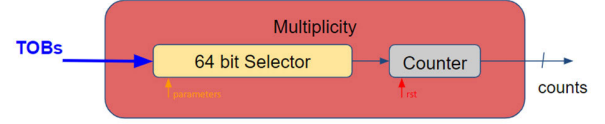Fig. 16. Sequential implementation of the select algorithm [22].

Fig. 17. Sequential implementation of the multiplicity algorithm [22].

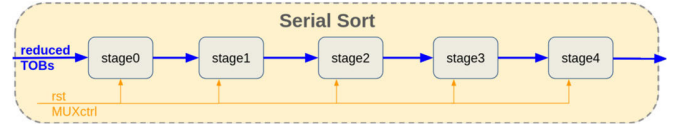Fig. 18. Sequential implementation of the sort algorithm [22].

is on the firmware rather than hardware. The system design attempts to minimize the distribution of data within the system in order to maximize flexibility by concentrating the complete BC event data on a single GEP node.

In order to process the full event data, each GEP node will host an abundance of algorithms. Currently, about 20 different algorithm blocks are planned for the target device: VP1802 with a total of ∼3 million LUTs. The logic resources are allocated based on the expected algorithm needs. According to the proposed allocation [21], the programmable logic resource budget of the hypothesis block (topological algorithms firmware block) in Phase-II is extremely tight: 100k LUTs.

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. 15 shows an example implementation of the same decision algorithm as used in the example in Fig. 6. 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. 31 732 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 subticks 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, where algorithm uses a 320 MHz clock. Within a GEP node, different algorithm blocks can run at different frequencies with 320 MHz being the maximum allowed frequency due to the thermal restrictions of the hardware.

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 a configurable parameter-based threshold. Each select algorithm has its own adjustable threshold. Select algorithms within the global trigger environment are implemented sequentially. As shown in Fig. 16, 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 [22].

B. Multiplicity Algorithms

Adding a dedicated counter after the selector block creates the logic for a multiplicity algorithm (Fig. 17). Counter width can be flexibly configured. The counter value is the output of a multiplicity algorithm.

C. Sort Algorithms

Sequential implementation of the sort algorithms consists of sorting stages, the number of which corresponds to the number of input TOBs to be sorted and is equal to the pipeline depth (Fig. 18). The number of the sorting stages is fixed for a particular firmware build, but can be adjusted by setting a corresponding parameter prior to a firmware rebuild. 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.
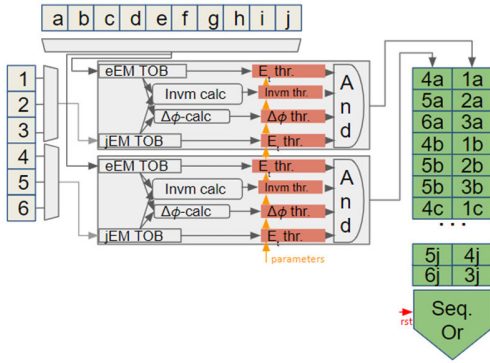
Fig. 19. Semi-sequential implementation of the example decision algorithm with double amount of logic resources in order to halve the latency [22].
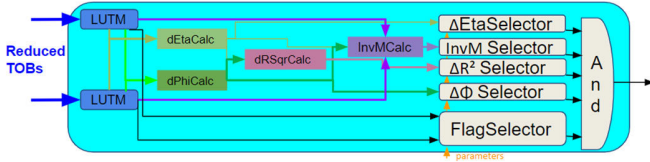

Fig. 20. Generic decision algorithm block diagram [22].

Fig. 21. Hypothesis algorithms implementation result [22].

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. 19). This feature allows to flexibly tune the utilized resources versus latency and helps to meet the overall latency budget. In order to use this feature, a corresponding parameter has to be set and the firmware has to be rebuilt.

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. 20). 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.

VII. 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. A built-in Vivado Simulator [23] was used. This is a hardware description language (HDL) simulator. It supports timing and functional simulations for VHDL, Verilog as well as SystemVerilog.

As an example, a behavioral simulation of the serialized decision algorithm (InvmDrSqrIncl2) is described. The TOBs
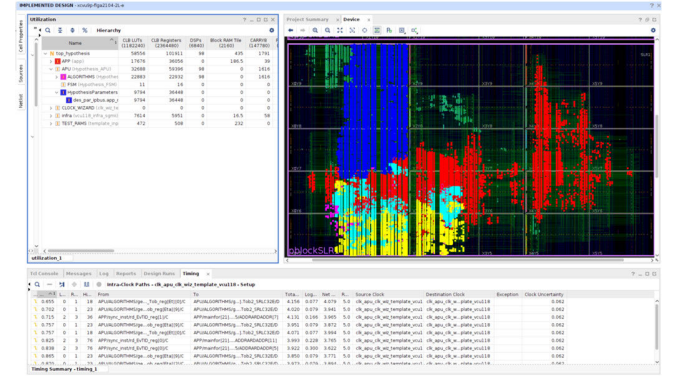
from the two input lists are serially provided to the functional algorithm blocks that perform calculations and apply thresholds for one combination at a time. The output is a trigger bit, which is either accept or reject. In this case, the simulation showed that event numbers 2 and 5 have triggered. The input TOBs are generated and the decision parameters are set by the simulation environment. The algorithm requires one clock tick to process one combination of input TOBs. Since in this example, we have two input lists, n^2 clock ticks are required. Running at 320 MHz clock frequency, we have 3.125 ns per clock tick. In case of ten input TOBs per list, the processing time takes 312.5 ns, which is the dominant latency. In addition to that a small pipeline latency of the algorithm itself and the pipeline latency of the Finite State Machine, controlling the data flow has to be accounted for.

At the next step, the original decision algorithm (InvmDrSqrIncl2), implemented in parallel is simulated. Using the same input TOBs and decision parameters as for the serialized implementation, event numbers 2 and 5 have also triggered as expected.

The same simulation has been performed for all the serialized algorithms. Correct triggering on the same events in both the serialized version and the original parallel implementation has been verified. A full list of the simulated algorithms as well as example simulation waveforms can be found in [22].

Further development plans include implementing a simulation functionality for the hypothesis algorithms to support a complete trigger menu validation.

VIII. IMPLEMENTATION

The first implementation of the hypothesis algorithms has been performed with serialized components assembled according to a certain trigger menu (v24.0.23), which is the menu in use at the time of writing. Most of this menu is covered except for a few algorithms that still need to be serialized. The menu coverage includes 21 out of 25 Select and Sort algorithms, 150 out of 151 multiplicity algorithms, and 41 out of 41 Decision algorithms.

Fig. 21 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. Tests at higher clock frequencies are foreseen as well in order to ensure a robust design. 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.

IX. 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.

Following the Phase-I upgrade, the Phase-II upgrade activities are currently ongoing. 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 μs) and a substantially tighter resource budget (∼100k 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.

Firmware and hardware developments for the global trigger are currently ongoing. The next steps include development of the final firmware for all the node types of the global trigger along with the production GCM hardware.

REFERENCES

[1] S. Fartoukh et al., *LHC Configuration and Operational Scenario for Run 3*, Standard CERN-ACC-2021-0007, 2021.

[2] *L1Calo Phase-I L1Topo Specifications*. Accessed: May 1, 2024. [Online]. Available: [Online.] Available: https://twiki.cern.ch/twiki/bin/viewauth/Atlas/LevelOneCaloUpgradeModules

[3] ATLAS Collaboration, *Technical Design Report for the Phase—I Upgrade of the ATLAS TDAQ System*, Standard CERN-LHCC-2013-018, 2013.

[4] M. Weirich, "Development of new ATLAS trigger algorithms in search for new physics at the LHC," Ph.D. thesis, Institut für Physik, Johannes Gutenberg-Universität Mainz, Mainz, Germany, 2021.

[5] *L1Calo Phase-I eFEX Specifications*. Accessed: May 2, 2024. [Online]. Available: [Online.] Available: https://twiki.cern.ch/twiki/bin/viewauth/Atlas/LevelOneCaloUpgradeModules

[6] The ATLAS L1Calo Group, *Global Feature Extractor of the Level-1 Calorimeter Trigger: ATLAS TDAQ Phase—I Upgrade gFEX Final Design Report*, Standard ATL-COM-DAQ-2016-184, 2016.

[7] R. Spiwoks et al., *The ATLAS Muon-to-Central Trigger Processor Interface (MUCTPI) Upgrade*, Standard ATL-DAQ-PROC-2017-013, 2017.

[8] R. Gugel. *L1Topo: The Catalyst in the Level-1 Trigger System for Run 3*. Accessed: Jul. 10, 2024. [Online]. Available: https://indico.cern.ch/event/1392489/

[9] Xilinx. *UltraScale Architecture and Product Data Sheet: Overview*. Accessed: May 5, 2024. [Online]. Available: https://www.xilinx.com/support/documentation/data_sheets/ds890-ultrascale-overview.pdf

[10] *MiniPOD AFBR-814xyZ, AFR-824VxyZ, 14 Gbps/Channel Twelve Channel, Parallel Fiber Optics Modules*, Standard AV02-4039EN, Avago Technologies, San Jose, CA, USA, 2013.

[11] S. Hillier et al., *Overview of Phase-1 Topological Trigger Connectivity and Commissioning Stages*, Standard ATL-COM-DAQ-2017-106, 2017.

[12] *Megtron 6, Ultra-low Loss, Highly Heat Resistant Circuit Board Material*. Accessed: May 10, 2024. [Online]. Available: https://industrial.panasonic.com/ww/products/pt/megtron/megtron6

[13] B. Bauss et al., *Level-1 Topological Processor (L1Topo) ATLAS TDAQ Phase—I Upgrade Hardware Tests*, ATL-COM-DAQ-2019-136, 2019.

[14] J. Damp, "Search for dijet resonances with the level-1 topological processor at ATLAS," Ph.D. thesis, Johannes Gutenberg-Universität Mainz, Mainz, Germany, 2020.

[15] J. Damp, *Level-1 Topological Processor (L1Topo) ATLAS TDAQ Phase—I Upgrade Firmware Algorithms*, Standard ATL-COM-DAQ-2019-010.

[16] E. Meuser. *The ATLAS Level-1 Topological Processor*. Accessed: May 15, 2024. [Online]. Available: https://cds.cern.ch/record/2869237

[17] The ATLAS L1Calo Group. *L1Calo Phase—I Feature Extractor Trigger Object Formats*. Accessed: Jul. 12, 2024. [Online]. Available: https://edms.cern.ch/ui/file/1492098/1/L1CaloTOBFormats_v015_docx_cpdf.pdf

[18] *ATLAS Experiment—Public Results*. Accessed: May 10, 2024. [Online]. Available: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/BPhysicsTriggerPublicResults

[19] ATLAS Collaboration, *Technical Design Report for the Phase II Upgrade of the ATLAS TDAQ System*, Standard CERN-LHCC-2017-020, 2017.

[20] *ATLAS TDAQ Phase—II Upgrade: Hardware Specifications for the Global Trigger*. Accessed: Jul. 12, 2024. [Online]. Available: https://edms.cern.ch/ui/file/2677534/1/Hardware_v0.7.pdf

[21] *ATLAS TDAQ Phase-II Upgrade: Firmware Specifications for the Global Trigger*. Accessed: Jul. 12, 2024. [Online]. Available: https://edms.cern.ch/document/2677532/1

[22] *ATLAS TDAQ Phase—II Upgrade: Global Event Processor Hypothesis Algorithms Specification*. Accessed: Jul. 15, 2024. [Online]. Available: https://edms.cern.ch/ui/file/2856352/1/Hypothesis_0.5.pdf

[23] *Vivado Design Suite User Guide: Logic Simulation (UG900)*. Accessed: Jul. 12, 2024. [Online]. Available: https://docs.amd.com/r/en-US/ug900-vivado-logic-simulation/Simulating-with-Vivado-Simulator