

The ATLAS Hardware Track Trigger design towards first prototypes

Francesca Pastore*, on behalf of ATLAS Collaboration

Royal Holloway, University of London, United Kingdom

E-mail: francesca.pastore@cern.ch

In the High-Luminosity LHC, planned to start with Run 4 in 2026, the ATLAS experiment will be equipped with the Hardware Track Trigger (HTT) system, a dedicated hardware system able to reconstruct tracks in the silicon detectors with short latency. This HTT will be composed of about 700 ATCA boards, based on new technologies available on the market, like high-speed links and powerful FPGAs, as well as custom-designed Associative Memories (AM) ASICs, which are an evolution of those used extensively in previous experiments and in the ATLAS Fast Tracker (FTK). The HTT is designed to cope with the expected extreme high luminosity in the so called L0-only scenario, where HTT will operate at the L0 rate (1 MHz). It will provide good quality tracks to the software High-Level-Trigger (HLT), operating as co-processor, reducing the HLT farm size considerably, by lightening the load of the software tracking. All ATLAS upgrade projects are also designed for an evolved, so-called L0/L1 architecture, where part of the HTT is used in a low-latency mode (L1Track), providing tracks in regions of ATLAS at a rate of up to 4 MHz, with a latency of a few micro-seconds. This second phase poses very stringent requirements on the latency budget and to the dataflow rates. All the requirements and the specifications of this system have been assessed. The design of all the components has been reviewed and validated with preliminary simulation studies. After these validations are completed, the development of the first prototypes will start. In this paper we describe the status of the design review, showing challenges and assessed specifications, towards the preparation of the first slice tests with real prototypes.

*European Physical Society Conference on High Energy Physics - EPS-HEP2019 -
10-17 July, 2019
Ghent, Belgium*

*Speaker.



1. Introduction

The High-Luminosity LHC (HL-LHC) [1] will significantly extend the physics potential of experiments, by increasing the peak luminosity up to $7.5 \times 10^{34} \text{cm}^{-2} \text{s}^{-1}$ and providing 3000fb^{-1} of proton-proton collisions at a centre-of-mass energy of 14 TeV. Thanks to major upgrades of its components, the ATLAS [2] experiment will be able to collect this incredibly large data sample, mainly dedicated to electroweak and QCD precision measurements, exploration of Higgs boson properties, flavour physics and searches for new physics beyond the Standard Model.

To maintain full acceptance on processes at the electroweak energy scale, a major upgrade of the ATLAS trigger will be needed. The trigger strategy will remain focused on the identification of high momentum particles, in particular on single and double lepton signatures, with thresholds as low as in Run 2, as shown for some signals selected with single muons in Figure 1a. A completely new architecture of the trigger/DAQ system will exploit the full potential of the new LHC machine, accepting higher data throughputs: first-level trigger rates will increase from 0.1 to 1 MHz and the DAQ dataflow from 1 to 50 Tbps. In the high-luminosity scenario, the average number of interactions per bunch crossing (pile-up) will grow from the current 40 up to 200. In this harsh environment, reduced rejection power is expected for the trigger selection algorithms. To mitigate these effects, and in addition to the aforementioned extended readout capability, both upgraded detectors and new hardware components will be introduced.

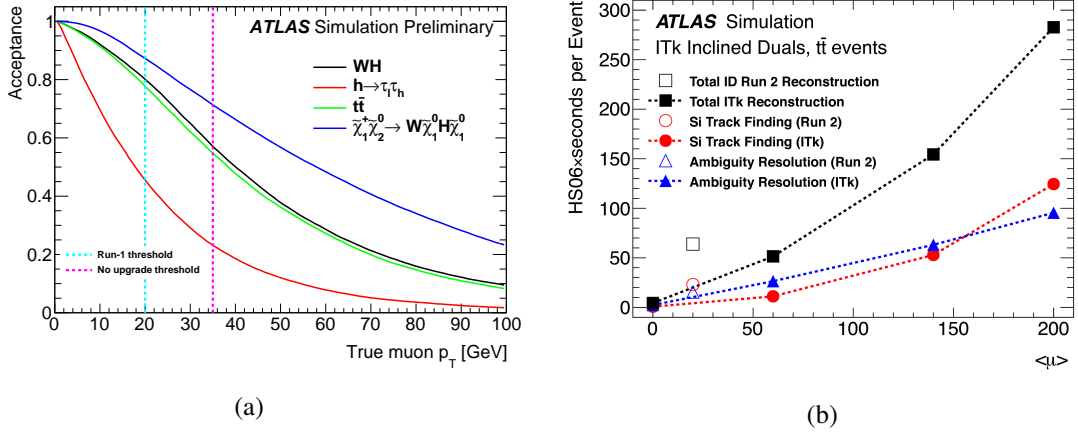


Figure 1: (a) Signal acceptance as a function of muon p_T for various signal predictions at 14 TeV with one muon in the final state [4], and (b) CPU required to reconstruct a $t\bar{t}$ event in the ITk as a function of the average pile-up, showing contributions from different algorithms [3].

The use of the new high precision full-silicon Inner Tracker (ITk [3]) as early as possible in the trigger selection is crucial. This detector, composed of inner layers of silicon pixels and outer larger-area strips, can provide optimal p_T -resolution, for rejecting low- p_T leptons, and the possibility of identifying the hard-scattering interaction, via the reconstruction of the primary vertex. This is a powerful tool to suppress the abundant low- p_T jets component from pile-up, helping in reducing the thresholds on hadronic signals (multi-jets and missing energy) and on any combined multi-object selection. It's worth noticing that data flow and event complexity do not scale linearly with

luminosity due to combinatorial effects. This has a particular impact on the software tracking algorithms, which require more and more CPU resources, as shown in Figure 1b.

2. The ATLAS Hardware Track Trigger project

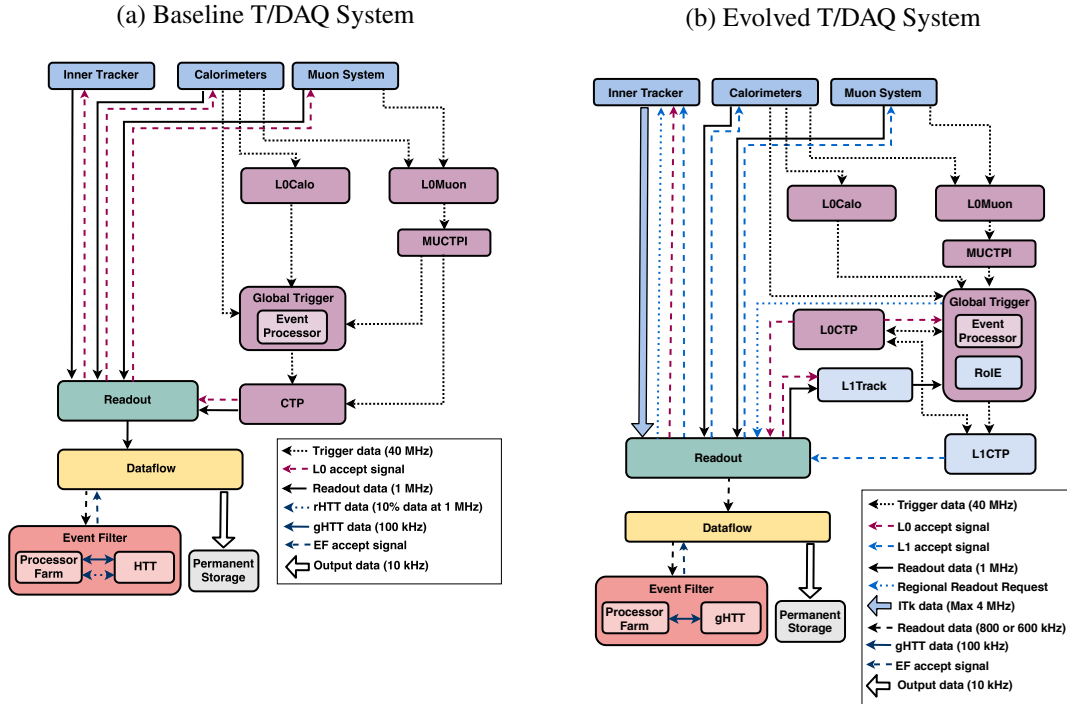


Figure 2: Diagram of the baseline single-level hardware trigger configuration (a), compared to the evolved T/DAQ System with a two-level hardware trigger (b) [4].

For the HL-LHC scenario, ATLAS has approved the Hardware Track Trigger project (HTT) [4], planning to build a pure hardware system to perform tracking in a low latency environment. The baseline T/DAQ architecture is shown in Figure 2a. The first level trigger (Level-0 or L0) will reduce the 40 MHz collision rate to 1 MHz, with a latency of $10 \mu s$, using signals from the calorimeters and the muon system. In a Region-of-Interest approach, the L0A signal indicates the type of selection and the region in which the L0 succeeded. The Event Filter is the successive trigger level, able to execute software selection algorithms, depending on the different L0 requests, and reducing the rate down to 100 kHz. The Event Filter is based on a farm of commodity processors, with the HTT included as co-processor to handle the tracking needs and reduce the demanding CPU load. This architecture allows to reduce the farm size by almost a factor 10 compared to the non-HTT scenario.

Uncertainties in the projected trigger rates for hadronic signals and in the occupancy in the inner pixel detector layers under ultimate HL-LHC conditions, with the subsequent increase in event size, motivate the design of an evolved architecture schema. As shown in Figure 2b, it will accept a 4 MHz input L0 rate and add a second hardware-level trigger (Level-1 or L1) with more

extended latency (below $35 \mu\text{s}$). Its main purpose is to allow HTT regional tracking and provide a rudimentary primary vertex reconstruction. The new schema can mitigate both of the above risks, reducing the readout rate down to $600 - 800 \text{ kHz}$, allowing opportunities of reducing thresholds and improving acceptance for many important physics signatures. In this scenario HTT is the key player of the L1 selection, and needs to process tracks within the allocated latency.

parameter	rHTT	L1Track	gHTT
maximum input rate	1 MHz	2 – 4 MHz	100 kHz
fraction of ITk data processed	5 – 10%	5 – 10%	100%
number of ITk logical layers	8	8	13
maximum latency (from request to return of data)	10 ms	$6 \mu\text{s}$	10 ms
minimum track p_T	2 GeV	4 GeV	1 GeV
η coverage	$ \eta < 4$	fixed by Pixel readout	$ \eta < 4$
minimum efficiency for muons $p_T > 10 \text{ GeV}$	98% in $ \eta < 2.5$		
minimum efficiency for electron $p_T > 10 \text{ GeV}$	95% in $ \eta < 2.5$		
minimum efficiency averaged over all η	95%		
minimum efficiency in each η region	85%		

Table 1: Key performance requirements on the three HTT configurations

The evolution scenario is not a key driver of the HTT specifications, but the possibility of low latency processing (below $6 \mu\text{s}$) must be included in the design. For this reason, the HTT design does not impact the final tracking capabilities of the ITk tracker, but only requires a fast readout on its front-end electronics. Flexibility is one of the design principles: HTT can operate in different conditions by means of different configurations. It can execute either regional or global tracking, depending on the L0 selection type, with different performance summarised in Table 1.

Regional HTT (rHTT) provides fast coarse resolution tracks with $p_T > 2 \text{ GeV}$. Only a reduced number of ITk layers (8 among a maximum of 13) are read out and only in regions around specific L0 selections, which correspond to less than 10% of the full readout. The rHTT is designed to operate on single lepton selections but an extended usage for multi-object selections is considered, mainly for the evolved scenario. In this latter case, the system is called L1Track, since it has the additional requirement of short latency processing, which motivates the increase of the minimum track p_T , set at 4 GeV. Coarse resolution and limited η coverage is expected since only few ITk inner pixel layers can have high bandwidth instrumentation. Global tracking (gHTT) provides full-resolution tracks with $p_T > 1 \text{ GeV}$, in the full tracker acceptance $|\eta| < 4$. For this reason it operates at reduced rate (100 kHz) in both scenarios. Quasi-offline algorithms can be used with these tracks, making gHTT crucial for b -tagging, τ selection, missing- E_T and pile-up mitigation purposes.

2.1 HTT processing

A detailed description of HTT can be found in [4]. HTT processing is an evolution of the current ATLAS Fast Tracker (FTK) [5], with the same tracking strategy but evolved hardware within a more modular system. The tracking processing is divided in three steps. The first step performs a track finding based on pattern recognition, aimed at reducing the huge combinatorics. It

compares combinations of eight input hits with all pre-stored patterns within Associative Memory (AM) ASICs, in one clock cycle. Before matching, silicon hits from outer ITk layers (including at least one pixel layer, with acceptable occupancy) are clustered into low granularity super-strips to reduce the possible number of patterns stored in memory. For every matched combination, a first-stage track-fit is executed, using eight full-resolution hits extracted from the patterns. Linearised fits based on principal component analysis [6] can be easily implemented in FPGAs, coupling hit positions with pre-stored track parameters. The algorithm calculates helix parameters and the resulting χ^2 allows the selection of good track candidates. A dedicated algorithm further reduces the number of output tracks, removing duplicates that share the majority of hits. These two steps are the core processing of rHTT. Resolution of track parameters is limited due to the reduced lever arm and distance to impact point, but a rejection factor of five can be reached on single leptons with more than 95% efficiency (see an example ROC curve in Figure 3a). Global tracking gHTT foresees an additional second-stage track-fit, again executed in FPGAs. It extrapolates first-stage tracks into the remaining ITk layers and finds nearby hits, to execute similar fit algorithm. Track candidates from the second-stage have resolution close to what is achieved in software algorithms. A comparison of the different resolution on z_0 (the distance from the interaction point) along η is shown in Figure 3b.

The baseline HTT system and the L1Track systems are separately dimensioned to meet their goals. System design parameters, needed to assess the requirements in Table 1, are extracted with a detailed simulation of the full processing chain.

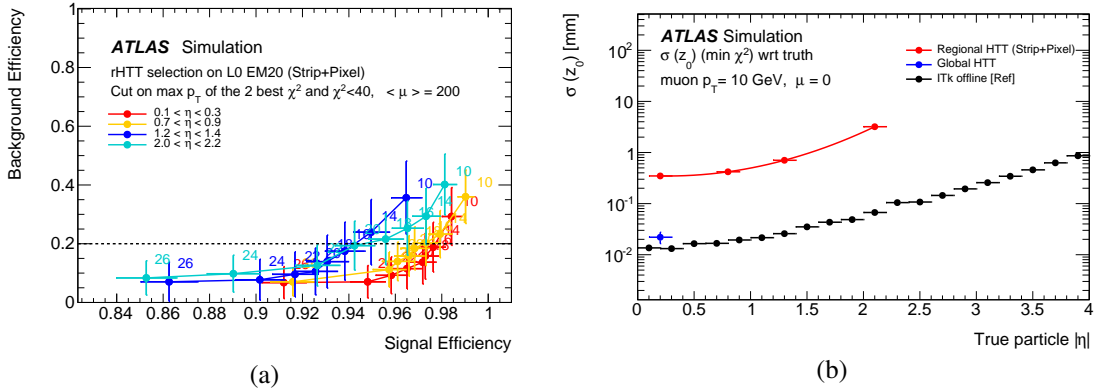


Figure 3: (a) Rejection as a function of efficiency in different η regions, when requiring a first-stage track for 20 GeV electrons (the numbers next to the lines show the cut on track p_T), and (b) comparison of the z_0 resolution for 1st-, second-stage fitting and offline reconstruction as a function of η [4].

3. System description

The HTT is a massively parallel system, made of custom-made boards organized in HTT-units. An overview is shown in Figure 4. Each unit is interfaced to the Event Filter farm via dedicated servers called HTT-Interfaces (HTTIF). Each HTT-unit contains two kinds of boards, one for the first-stage (Associative Memories Tracking Processor, AMTP) and one for the second-stage




Figure 4: Overview diagram of the HTT system for the baseline scenario [4].

(Second Stage Tracking Processor, SSTP). All boards are based on the same Tracking Processor (TP) motherboard, an ATCA blade, which only differ by mounting different mezzanines.

For the baseline scenario, inter-board/intra-board data rates and power consumption are key drivers of the system specification. The number of AMTP boards is given by the maximum input rate that each board, covering a fixed detector region, can afford. The size of the region is also linked to the minimum number of AM chips that can be placed on-board. On the other hand, AM power consumption is proportional to the input processing rate, and power limit per board fixes the maximum area density of chips. A larger number of chips allows more patterns to be stored, and this gives the ability to go to lower p_T or have finer granularity patterns, thus extending the system capabilities while lightening the load on the track-fit stage. Given the strong link between the number of AM chips per board and the FPGA resources, a good compromise must be found to limit the overall power consumption. Second-stage fitting has reduced inputs, so power consumption and dataflow are not critical.

In the evolved scenario, the size of the system is driven by latency requirements. The slowest components - the AM blocks - are dominant, so enough link bandwidth is allocated to these components. Since system size scales approximately with $1/\min-p_T$, one can also exploit higher threshold to cope with unexpected larger latency. In addition, in this scenario, two copies of patterns are allowed per board, to have higher probability of freeing AM ASICs chips before the next L0 regional request. Given these tight links between components and the constraints, the specifications of the system have been set, and the prototype development has started.

3.1 Components specification and status

The TP is a high-bandwidth motherboard, that can support either two small mezzanines (SSTP) or one large mezzanine (AMTP). It contains a Rear Transition Module (RTM), 10 Gbps links, one large FPGA and a SoC to support full system monitoring. In addition to data communication, it also executes specific algorithms, like pixel clustering and duplicate track removal. Critical aspects under evaluation in this phase are the choice of the mezzanine connectors (candidate Z-ray [7]), and of the high-speed links, and the evaluation of the thermal and mechanical models.

The AMTP mounts one large pattern recognition mezzanine (PRM), for the first-stage processing. The core component is the AM ASIC, which is a low-power CMOS device that can execute

about 30 peta bit-wise comparisons per second. Its low power and small size allow high density of chips on board. AM09 has been chosen as the chip for HTT, which is the evolution of FTK AM06, improving by a factor of three the number of patterns ($3 \times 128k$), the I/O bandwidths (from 100 to 250 MHz clock), and the power schema (from 1.11 to 0.42 fJ/comparison/bit), thanks to a new technology cell, based on 28 nm technology. Power is driven by bit-comparison (1 W plus 0.05 W/MHz), so that with 16 bits at 50% bit-flip and data at 50 MHz on 8 buses, one can expect 2.5 W consumption. Each bus is equipped with 1 Gbps LVDS links. Current understanding of AM07 tests and of AM08 simulations shows that AM09 will be within the allocated budget in terms of both area and power consumption. Four groups of 5 AM ASICs (about 7 M patterns per PRM) will be placed on the mezzanine, and connected to a large FPGA for data sharing and first-stage track fitting (~ 1 GHz fits/board). The maximum fit rate and output bandwidth are driven by the evolution scenario, while the FPGA choice is driven mostly by commercial aspects. The requirement to support high request rates of constants for the fit, for example, makes Intel Stratix 10 [8], with internal High Bandwidth Memory (HBM), a good candidate. For the second-stage tracking the SSTP will mount two small mezzanines, the track fitting mezzanine (TFM), which will follow the same connector and power dissipation schema as the PRM.

Current estimates foresee 24 HTTIF servers, 48 ATCA shelves, each containing 14 blades, 12 AMTPs and 2 SSTPs. It will provide more than 1400 cards for 5.3 billion patterns, for a total of 11520 ASICs and 1728 FPGAs. The overall power-budget is 385 kW (including contingency), with about 1 Tbps output bandwidth. Power budget and dataflow remain challenging and are key drivers of the system design, which just passed the specification review process.

4. Conclusions

The HTT is under design and is a crucial component of the ATLAS trigger upgrades in the HL-LHC. It is designed with good flexibility and modularity, to run as both a regional and a global tracking co-processor in the HLT (and, if needed, the system can evolve to run as an L1-track-trigger). The required performance, in terms of efficiency and resolution, has driven the design choices, and the experience in the chosen technologies allows for detailed estimates of resources and project planning.

References

- [1] G. Apollinari et al., [CERN-2017-007-M](#)
- [2] ATLAS Collaboration, [2008 JINST 3 S08003](#)
- [3] ATLAS Collaboration, [CERN-LHCC-2017-021](#)
- [4] ATLAS Collaboration, [CERN-LHCC-2017-020](#)
- [5] ATLAS Collaboration, [CERN-LHCC-2013-007](#)
- [6] H. Wind, [CERN-EP-INT-81-12-REV, 1982](#)
- [7] Samtec, [Z-ray ULTRA-LOW PROFILE HIGH-DENSITY ARRAYS](#)
- [8] Intel, [Stratix 10 documentation](#)