The Fast TracKer - A hardware track processor for the ATLAS trigger system

Karolos Potamianos[1,*]

[1]Deutsches Elektronen-Synchrotron (DESY), Hamburg, Germany

Abstract. The Fast Tracker (FTK) is a hardware upgrade to the ATLAS trigger and data acquisition system providing global track reconstruction to the High-Level Trigger (HLT) in a high pile-up environment. The FTK processes incoming data from the Pixel and SCT detectors (part of the Inner Detector, ID) at up to 100 kHz using custom electronic boards. ID hits are matched to predefined track patterns stored in associative memory (AM) on custom ASICs while data routing, reduction and parameter extraction is achieved with processing on FPGAs. With 8000 AM chips and 2000 FPGAs, the FTK provides enough resources to reconstruct tracks with transverse momenta greater than 1 GeV/c in the whole tracking volume with an average latency below 100 microseconds at collision intensities expected in Runs 2 and 3 of the Large Hadron Collider. The tracks will be available at the beginning of the trigger selection process, which allows the development of pile-up resilient triggering strategies to identify b-quarks and τ-leptons, mitigating pile-up dependent effects on jet and missing energy reconstruction, as well as providing the potential to devise new selections to look for particular signatures (e.g. displaced vertices) in the search for New Physics phenomena.

1 Introduction

The Large Hadron Collider (LHC) collides proton bunches every 25 ns at a center-of-mass energy of 13 TeV. With a luminosity of 2×10^{34} cm^{-2} s^{-1}, each collision produces on average 38 minimum-bias interactions (termed pile-up interactions, μ) with some runs having as high as $\mu = 70$ (see Figure 1). This results in a high detector occupancy and poses challenges to the event readout and reconstruction. Because of the limited data transfer and storage bandwidth available, a reduction of the event rate by 5 or 6 orders of magnitude is necessary. This conflicts with the requirement that interesting physics signatures to study the Standard Model (SM) and improve our reach to physics beyond the SM (BSM) be selected efficiently among the large background.

2 The ATLAS trigger system

The primary role of the trigger system of the ATLAS experiment [1] is to select which events are written to disk. In doing so, its goal is to keep a large fraction of events from interesting

*e-mail: karolos.potamianos@cern.ch

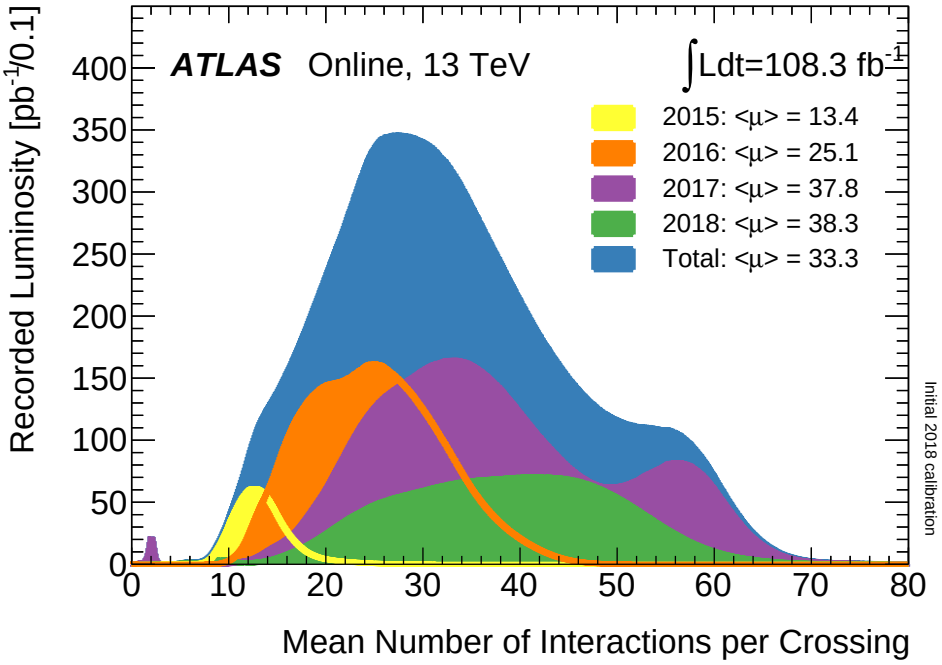

Figure 1: The luminosity-weighted distribution of the mean number of interactions per crossing μ for the 2018 pp collision data at 13 TeV centre-of-mass energy. All data recorded by ATLAS during stable beams through 12 June 2018 (LHC fill 6778) are shown. The luminosity is estimated using an initial calibration for the 2018 data. [4]

but rare processes in pp collisions. But the events produced by these processes are hidden in very large backgrounds and only a small fraction of the data collected by the experiment can be taken off the detector and stored on disk and subsequently on tape. This requires significant realtime data reduction which employs massive computational resources to parallelize and minimize the execution time of the complex algorithms used to select rare events. A multi-level trigger system is used to address this problem.

The ATLAS Run-2 trigger system, shown in Figure 2, consists of a hardware-based Level-1 (L1) trigger [2] and a software-based High-Level Trigger (HLT) [3]. The L1 trigger is used to determine *Regions-of-Interest* (RoI) in the detector using custom electronics which process coarse granularity information from the calorimeters and the muon detectors. The L1 trigger reduces the event rate from the LHC bunch crossing rate of 40 MHz to 100 kHz, with a decision time per event (L1 accept) of 2.5 μs.

The information from the RoIs identified at Level-1 are sent to the HLT which executes sophisticated software algorithms on commercially available computers using full granularity information in either the RoI or the whole event. The HLT reduces the event rate from 100 kHz to approximately 1 kHz, with an average processing time of about 200 ms per event, which is limited by technology and resource limitations.

The Level-1 trigger searches for signatures from high-p_T muons, electrons/photons, jets and τ-leptons decaying intro hadrons and topologies with large missing transverse energy (E_T^{miss}) or total transverse energy (E_T). It uses coarse-granularity information from the Re-

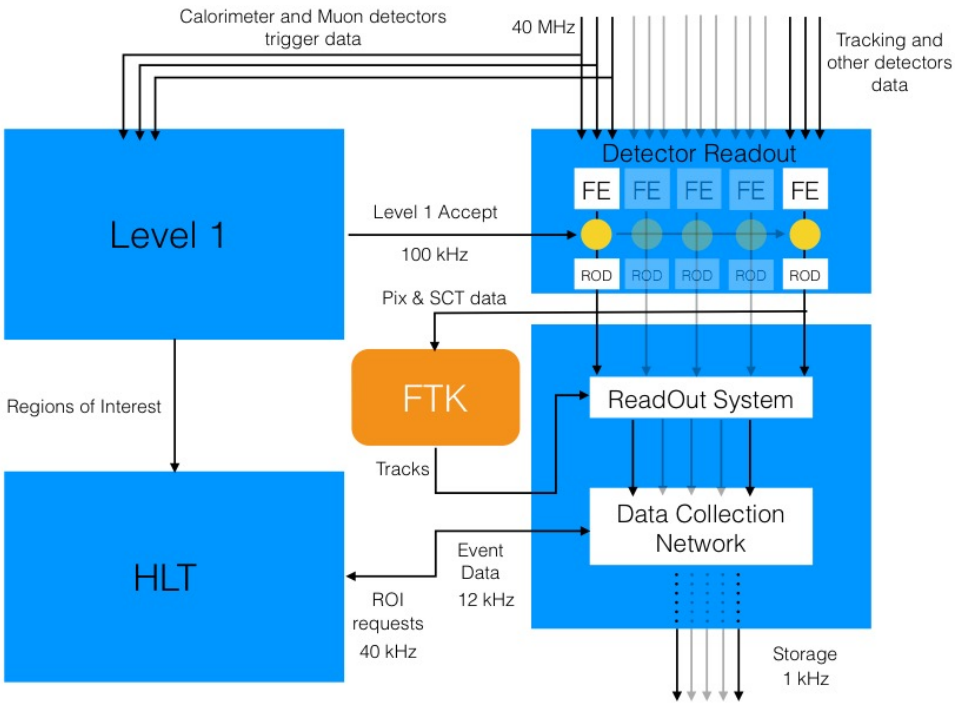

Figure 2: The ATLAS Trigger System. The Level-1 (L1) trigger provides Regions-of-Interest to the High Level Trigger (HLT) and produces L1 accepts (L1A) at 100 kHz, which trigger a full detector readout. The HLT farm reconstructs and filters the data, selecting interesting events at a rate of approximately 1 kHz, which are then written to disk. The Fast TracKer (FTK) will process data from the silicon detectors (Pixel and SCT) at the full L1A rate and provide track information within 100 μs to help the HLT make trigger decisions.

sistive Plate Chambers (RPC) and Thin-Gap Chambers (TGC) for high-p_T muons and all the calorimeter sub-systems for electromagnetic clusters, jets, τ-leptons E_T^{miss} and large E_T. Although the performance benefits would be large, the small timing window available to the L1 trigger system makes it impossible to use tracking information and data from the ATLAS silicon detectors, namely the Pixel [5], including the recently installed Insertable-B-Layer (IBL) [6], and the SemiConductor Tracker (SCT) [7] detectors. Moreover, global track reconstruction at the HLT cannot be performed within the HLT budget. Indeed, the offline reconstruction time (of which tracking accounts for over 50% of the total CPU time at $\langle\mu\rangle = 40$), shown in Figure 3 as a function of pile-up, has an exponential dependence on μ. Even using simplified algorithms, the HLT suffers from the same dependence on μ.

3 The ATLAS Fast TracKer

The Fast TracKer (FTK) [8] is a fast hardware-based track trigger system for ATLAS designed to perform a global track reconstruction receiving input from the silicon tracking detectors (part of the Inner Detector, ID) after each L1 trigger (at a rate of up to 100 kHz) and provide full-event track information to the HLT. It is designed to track all charged parti-

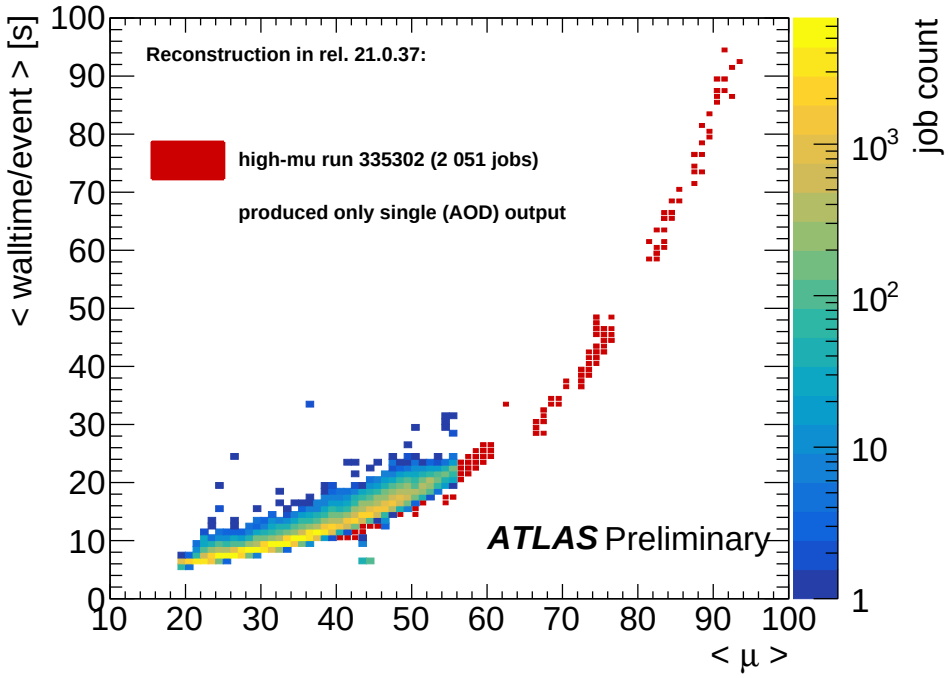

Figure 3: The dependency of reconstruction wall time per event as a function of the average number of interactions per bunch crossing for 13 TeV pp collisions. At $\langle\mu\rangle = 40$, the fraction of time spent on tracking tracking is over 50% of the total CPU time for recent software releases. [9]

cles with $p_T > 1$ GeV within $|\eta| < 2.5$[1] and provide the HLT with track parameters and hit coordinates within 100 μs.

The FTK aims at improving the efficiency of trigger selections that require tracking information, such as those to identify medium-p_T b's and τ's, with high background rejection. Both b-tagging and τ identification rely on track information, as the former are characterized by a displaced vertex that can be reconstructed from the tracks in the event, and the latter have significantly less tracks in a smaller cone than standard jets. This is especially important for measurements where b-jets and third generation leptons play a crucial role, such as Higgs coupling measurements or searches for particles predicted by BSM models. In addition to the physics gains, using FTK information in the HLT decision process can improve many trigger algorithms, in particular those relying on isolation variables (lepton triggers) and calorimeter information (jet and E_T^{miss} triggers), which are both affected by pile-up effects that can be mitigated using FTK tracks.

During the deployment and commissioning of the FTK, which started in 2016 and continues until the end of Run-2 and during the LHC Long Shutdown 2 (LS2), a partially installed

[1] ATLAS uses a right-handed coordinate system with its origin at the nominal interaction point (IP) in the centre of the detector and the z-axis along the beam direction. The x-axis points from the IP to the centre of the LHC ring, and the y-axis points upward. Cylindrical coordinates (r, ϕ) are used in the transverse (x, y) plane, ϕ being the azimuthal angle around the beam direction. The pseudorapidity is defined in terms of the polar angle θ as $\eta = -\ln[\tan(\theta/2)]$.

FTK system is processing ATLAS data in parasitic mode, i.e. writing track information on disk but without it being used by the HLT to make a trigger decision.

3.1 FTK processing

To cope with event rates of up to 100 kHz, the FTK needs to perform tracking faster than the HLT by several orders of magnitude. This can only be achieved with highly parallel processing. The ATLAS tracking volume is split into 64 $\eta - \phi$ regions (towers) which are processed independently (with some overlap due to the finite size of the beam's luminous region in z and the finite curvature of charged particles in the magnetic field). In addition, the data volume is decreased using a hit clustering algorithm, which yields cluster information which is then processed by the rest of the FTK rather than the full information from the Pixel or SCT detectors. During some of the processing steps, the cluster information is also re-binned in coarse-resolution Super-Strips (SS).

Tracking inside the FTK system is performed in two stages. First, track candidates are identified comparing resonstructed SS to predefined *patterns* stored in memory. These patterns are determined from MC simulation [11] and encode the trajectory in terms of SS of particles as they traverse the silicon detector volume. These coarse-resolution track candidates, called roads, seed a high-resolution tracking performed by FPGAs. The pattern matching procedure is performed using a custom associative memory (AM) ASIC designed for very high speed which allows a comparison of the incoming data to all stored patterns simultaneously. By considering only hits from the roads, the combinatorial problem of pattern recognition and parameter extraction is significantly reduced, as well as the fitting time.

The parameters of the pattern matching can be adjusted. Narrow roads permit fast track fitting but require many patterns to be stored and searched in the AM. Wide roads allow storing fewer patterns but slow down the track fitting due to increased combinatorics. A variable resolution feature using ternary bits [12] and a programmable number of matching layers allows optimizing between wide and narrow roads. Furthermore, the track parameter evaluation is reduced to a set of scalar products, as the fit method is based on a linearized model, using dedicated sets of coefficients. Each set is valid for a small region of the ID.

With the IBL, there are 12 silicon detector layers in the ATLAS ID: four Pixel and eight SCT layers. The first-stage 8-layer fit uses 3 Pixel hits (excluding those from the IBL) as well as five SCT hits (excluding 3 of the stereo hits). For the second stage, the IBL and SCT stereo hits are added to form 12-layer tracks. The linearized model is used for the coarse 8-layer pattern matching and fitting, as well as for the second stage 12-layer fit (but with different coefficients and number of coordinates). One (two) missing layer is allowed at the first (second) stage.

4 The FTK hardware

The functionalities mentioned in the previous section are implemented in custom electronic boards, hosted in Versa Module Europa (VME) crates and Advanced Telecommunications Computing Architecture (ATCA) shelves. Six types of custom-design electronic boards make up the FTK system: the Input Mezzanines (IM), the Data Formatters (DF), the Auxiliary Cards (AUX), the Associative Memory Boards (AMB), the Second Stage Boards (SSB) and the FTK to Level-2 Interface Cards (FLIC), which are described below. An AUX and AMB pair forms a Processor Unit (PU). The complete system has about 8000 AM chips and 2000 FPGAs from different vendors and models. This computing power is distributed on 32 DF

Table 1: Number of boards, crates, FPGAs and 2 Gbps links of the FTK system for Run-2 (Run-3). The links from the AUX to the DF and SSB (from the SSB to the HLT) are operated at 6.4 Gbps (3 Gbps).

Board Name	Units	Crates/Shelves	FPGA	Connectivity
IM	128	4x 14U ATCA	2x Xilinx Artix-7 or 2x Sparatan-6	4 SPF+
DF	32		1x Xilinx Virtex-7	160 (288) QSFP to AUX and SSB
AUX	64 (128)	6x 9U VME	6x Intel/Altera Arria V	2x QSFP from DF, 1x SFP to SSB
AMB	64 (128)		2x Xilinx Artix-7 + 2x Xilinx Spartan-6	16 AM06 x 4 LAMBs
SSB	32	8x 9U VME	5x Xilinx Kintex-7	QSFP from DF, SFP from AUX, and SFP to FLIC
FLIC	2	1x 6U ATCA	4x Xilinx Virtex-6	SFP from SSB, and SFP to HLT

boards, 64 AMB and AUX cards[2], 32 SSBs and 2 FLIC. Table 1 summarizes the number of FTK boards, FPGAs, and links.

4.1 Input processing and data formatting

The IM and DF cards are responsible for clustering the data and distributing it to the rest of the FTK, respectively. Each IM contains two FPGAs (either two Xilinx Artix-7 or two Spartan-6), each connected to two SFP+ transceivers[3] each handling 1 ID link. Each DF can host up to 4 IM daughter cards and can thus process the data of up to 16 ID links, contains a Xilinx Virtex-5 FPGA and is connected to a Rear Transition Module (RTM) supporting up to eight Quad-SFP (QSFP) and six SFP optical transceivers.

The DF system geometrically organizes the incoming clustered data in $\eta - \phi$ projective towers of dimension $\Delta\phi \times \Delta\eta \sim 32^\circ \times 1.2$ and in logical layers. Each DF board receives the data from 2 towers and feeds data to 4 core processors and 1 SSB board, and in addition can send ID data not belonging to the towers it serves to other DF boards. Optical links placed on the RTM are used to communicate with the PUs. Communication with DFs in the same crate uses the full-mesh backplane while inter-shelf communication uses an extra link in the RTM.

The IM process about 750 Gbps of data from O(400) ID links. Each DF receives data at up to 32 Gbps of clustered ID data, of which about 30 Gbps are routed to the core processors, while 40 Gbps (25 Gbps) of data can be sent to other DFs in the same shelf (other shelves).

4.2 Pattern matching and track fitting

At the core of the FTK are the PUs, composed by an AMB and an AUX card. Each PU performs the pattern matching and a first-stage (8-layer) fit, the most computationally-intensive steps of the pipeline. Data reaches the AUX card via QSFP connectors, already organized by layers. The card implements the Data Organizer (DO), a small database in which all clusters are organized according to a coarse-resolution position identifier (SS), the Track Fitter (TF), which constructs 8-layer tracks, and the Hit Warrior (HW), which removes duplicate tracks based on a number of common hits and χ^2.

The SS information is sent to the AM Board Serial Link Processor (AMBSLP) for pattern matching. The TF subsequently receives the list of roads found by the AMBLSP as well as the cluster information associated to them (from the DF). Based on the SS content of each road, the cluster information is received by the DO and the packet of clusters belonging to each road is sent to the TF. The TF builds all combinations of clusters in a road (with 1

[2] 128 card pairs are foreseen for Run-3

[3] The small form-factor pluggable (SFP) is a compact, hot-pluggable optical module transceiver used for both telecommunication and data communications applications.

cluster per layer) and evaluates the χ^2, sending all good candidate tracks to the SSB. The computations on the AUX are distributed among 6 Intel/Altera Arria V FPGAs, of which 2 control the input and the output, and 4 are devoted to the TF. The AMB receives the SSs from the AUX and sends them to the AM chips. The pattern matching is complete as soon as all hits reach the chips. Match roads are sent back to the AUX as soon as they are matched. The AMB comprises 2 Xilinx Artix-7 FPGAs controlling the input and output logic, and 2 Xilinx Spartan-6 FPGAs controlling the VME interface and the internal state of the board. Pattern matching is performed by 64 AM06 [13] ASICs installed on 4 LAMBs (Local Associative Memory Boards) daughter cards.

Each AUX receives data on 8 6.4 Gbps links from the DF and sends data on 1 6.4 Gbps link to the SSB. High speed serial links in the VME P3 connector guarantee 12 Gbps input from the AUX to the AMBSLP and 16 Gbps from the AMBSLP to the AUX.

4.3 Second stage fitting and interface to the HLT

The AUX does not exploit the full precision of the ID because it doesn't use the information from all tracking layers. The SSB rejects fake track combinations and improves the helix parameter resolution starting from 8-layer track fits and using information from the 12 layers. It also removes duplicate tracks. Each SSB receives the output of 4 AUX cards via a RTM, as well as the IBL and stereo SCT hits associated to the 2 corresponding towers from the DF system.

The SSB has 3 main functions: the Extrapolator, which uses 8-layer track information to compute the likely position of hits in the other 4 layers, the Track Fitter (TF) which determines the best helix parameters from hits in roads using 12 silicon layers, and the HW, which is similar to that of the AMB. These functions are implemented in 5 Xilinx Kintex-7 FPGAs (4 for Extrapolation and TF, and 1 for the HW functionality).

Because it is more effective to concentrate the HW functionality in the SSBs that send final tracks to the FLIC, there are two types of SSB, with identical Extrapolator and TF functions but different HW functions: the preliminary SSB (pSSB) which sends its tracks to $+\phi$-neighboring final SSB (fSSB) and the fSSB which receives tracks from its own TF and the -ϕ-neighboring pSSB, performs overlap removal, and outputs the track list to the FLIC. In the pSSB, the HW functionality is essentially a fan-out.

The SSB receives data at up to 25 Gbps (6.4 Gbps from each AUX) and shares data with other SSBs over optical links to perform overlap removal and merges FTK data within a core crate for output to the FLIC via 2 optical connections at 3 Gbps.

The FLIC collects the reconstructed track information from the SSB, reducing the data volume and converting it to a format compatible with the HLT software. The FLIC system has a total maximum bandwidth of 32 Gbps, provided by 8 SFP links from the SSB and 8 SPF+ links to the HLT. Each FLIC comprises 4 Xilinx Virtex-6 FPGAs (2 for processing and 2 for monitoring). The 2 FLIC boards are hosted in an ATCA shelf.

5 Commissioning and performance

The FTK system is being deployed in the ATLAS underground counting room (USA15). Currently, several slices[4] are operated in parallel, processing ID data either within an ATLAS data-taking run or standalone. While the system is being commissioned, several milestones have already been achieved. One of these milestones is producing 8- and 12-layer tracks in a slice with 4 IMs, 1 DF and 1 PU (AUX+AMB), outputting the data in the ATLAS

[4] A slice corresponds to a subset of the full FTK system, connected to parts of the ID.

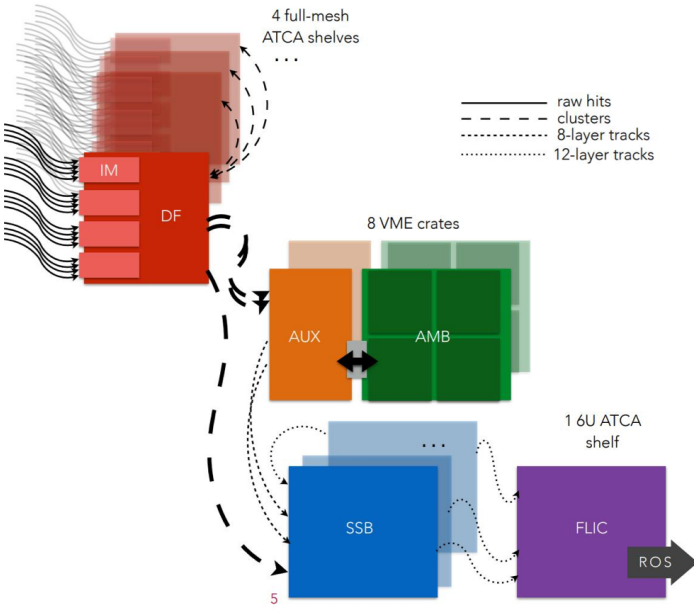

Figure 4: The data flow in the ATLAS Fast Tracker (FTK). The FTK Input Mezzanine (IM) cards, hosted on the Data Formatters (DF) receive raw hit data from the silicon detectors and perform hit clustering. The DF is responsible for re-distributing the clustered data (long dashes) into 64 $\eta - \phi$ towers and sends it to the Processing Units (PU), which consist of an Auxiliary (AUX) and an Associative Memory (AMB) board. The AUX performs 8-layer clustering using roads matching predefined patterns provided by the AMB. The Second Stage Boards (SSB) combine 8-layer track data (small dashes) from the AUX and the DF to form 12-layer tracks (dotted line), which are then formatted for the HLT by the FTK to Level-2 Interface Cards (FLIC).

format. Additional milestones are planned until Run-3 when the full FTK system is expected to provide track information which the HLT will use to perform trigger decisions.

In parallel to the commissioning of the FTK boards and the online infrastructure, improvements to the offline software are underway. In particular, the determination of the patterns to be loaded in the AM chips is constantly being improved. Figure 5(a) shows the distribution of the number of 8-layer tracks produced per event from an FTK slice. Figure 5(b) shows the p_T distribution of tracks from the FTK functional simulation (FTKSim) [14] matched to tracks output by the FTK slice, showing very good agreement. Figure 5(c) shows the fraction of FTK slice output tracks that are matched to tracks produced by running FTKSim on ATLAS RAW data. The efficiency, currently around 90-95%, is constantly being improved with improvements to the firmware and the FTK configuration (patterns and constants), and FTKSim. Figure 5(d) shows the fraction of 8-layer tracks with a minimum of 7 hits output by the FTK slice with missing hits in a Pixel or SCT layer, and those with no missing hits, as a function of the recorded event number.

Figure 6 shows the effect of disabled Pixel and SCT modules[5] on FTK track finding efficiency and the recovery using a WildCards algorithm, in which disabled modules are assumed

[5]These modules were disabled during the 2017 running.

to always have a hit in the pattern recognition stage. The WildCards algorithm recovers the track finding inefficiency due to disabled modules but also other sources of inefficiencies.

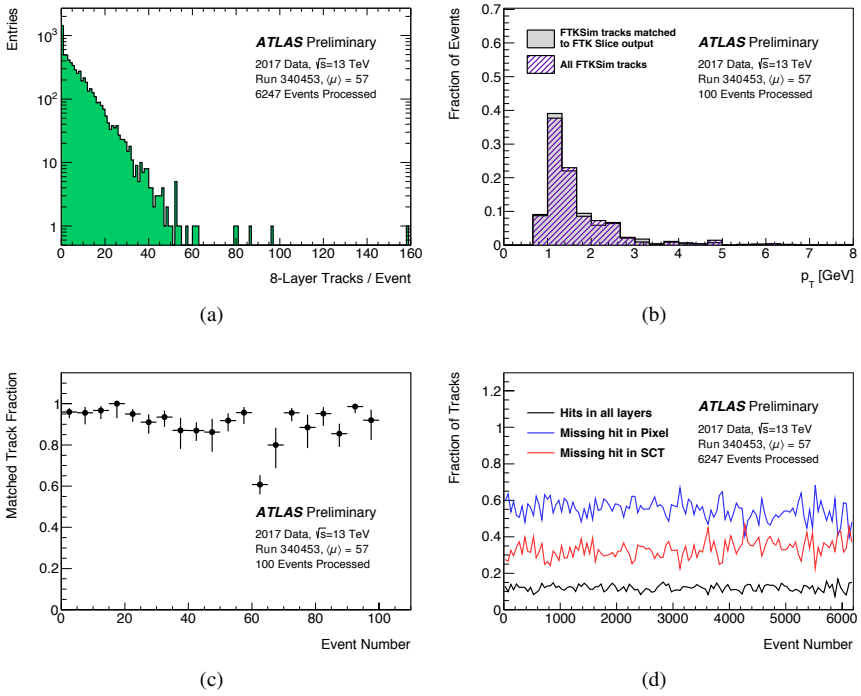

Figure 5: (a) Number of 8-layer tracks produced per event from an FTK slice. (b) The p_T distribution of FTKSim 8-layer tracks matched to tracks output by the FTK slice. The distribution of matched simulated tracks (grey) is compared to all simulated tracks for the same events (purple). (c) Fraction of FTK slice output tracks that are matched to tracks produced by running FTK functional simulation (FTKSim) [14] on ATLAS RAW data. (d) The fraction of 8-layer tracks with a minimum of 7 hits output by the FTK slice with missing hits in a Pixel (blue), a SCT layer (red), and those with no missing hits, as a function of recorded event number. The slice contains 4 IMs, one DF, one AUX and one AMB. In this slice, instead of sending 8-layer track information to the SSB, the AUX outputs directly to the ATLAS readout. The 8-layer track information consists of hit coordinates, pattern identification information, and which layers are included in the 8-layer track fit. The slice spans $-0.2 < \eta < 1.24$ and $2.4 < \phi < 2.8$. These plots cover a subset of ATLAS Run 340453, a 13 TeV pp run which began on November 9, 2017, in which 6247 events were processed by FTK at a rate of 35 kHz. [10]

6 New trigger capabilities

In addition to enhancing the capabilities of the current ATLAS trigger with respect to identifying b-jet and τ-lepton (central to probing the Standard Model and to searching for new physics signatures) and mitigating the effect of multiple interactions, which are the main priority, the FTK is able to provide new trigger capabilities. In particular, it can be used to look

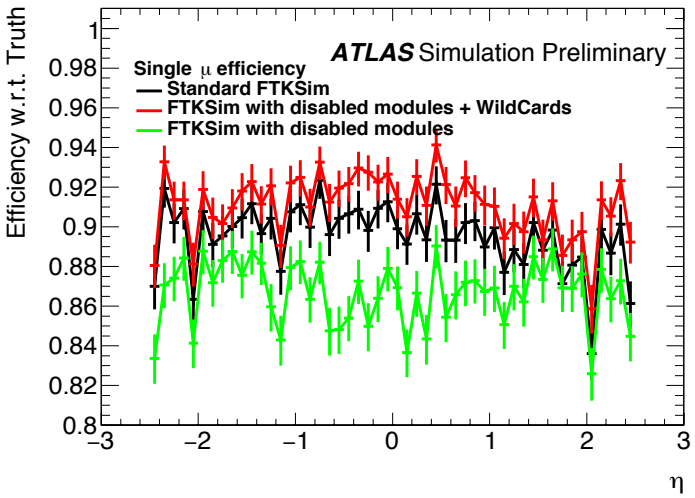

Figure 6: The FTK track finding efficiency from FTK functional simulation [14] of Monte Carlo single muon events as a function of the muon pseudorapidity η . The black line shows the simulation output assuming perfect detector operation. The green line shows the simulation output when hits from the list of disabled Pixel and SCT modules from 2017 running are excluded from the pattern recognition and track finding, emulating typical detector operation. A drop in the efficiency can be seen when disabled modules are taken into account. The red line shows the efficiency using the WildCards algorithm, in which disabled modules are assumed to always have hits in the pattern recognition stage. The track fitting stage is unchanged by the algorithm. Using this algorithm, the efficiency improves significantly. The lower efficiency at $\eta > 0$, due to the pattern banks, has been corrected in newer versions. [10]

for signatures that are hard to trigger on at Level-1 and have a distinctive tracker activity. With full reconstruction of tracks with $p_T > 1$ GeV, the FTK gives access to b-quarks with lower momentum than previously available. It also allows reconstructing "delayed particles" looking for displaced vertices, which is a distinct signature of long-lived BSM particles. This requires special "delayed particle" patterns to be loaded in the AM chips (more precisely, in a fraction of the available pattern space). These differ from those looking for the primary vertices within the region where protons are colliding and the (slightly) displaced secondary vertices of b-tags.

7 Conclusion

The ATLAS FTK is expected to provide high quality tracks to the HLT algorithms at the full Level-1 rate of 100 kHz. This will allow a more efficient collection of pp collision events from the LHC with b-jets or τ-leptons, mitigating pile-up dependent effects on jet and missing energy reconstruction, and will also provide the potential to devise new selections to look for particular signatures (e.g. displaced vertices) in the search for New Physics phenomena. The hardware has been produced and is being installed in the ATLAS counting room. The commissioning of the full system has started, and will continue until the end of LS2, processing tracks in a parasitic mode until it is ready to assist the HLT in collecting ATLAS data for interesting physics events.

References

[1] ATLAS Collaboration, "The ATLAS Experiment at the CERN Large Hadron Collider", JINST 3 (2008) S08003.

[2] ATLAS Collaboration, "ATLAS level-1 trigger : Technical Design Report", CERN-LHCC-98-014, ATLAS-TDR-12

[3] ATLAS Collaboration, "ATLAS high-level trigger, data-acquisition and controls : Technical Design Report", CERN-LHCC-2003-022, ATLAS-TDR-16

[4] https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2

[5] G. Aad *et al.*, "ATLAS pixel detector electronics and sensors", JINST 3 (2017) P07007

[6] ATLAS IBL Collaboration, "Production and Integration of the ATLAS Insertable B-Layer", JINST 13 (2008) T05008

[7] A. Abdesselam *et al.*, "The integration and engineering of the ATLAS SemiConductor Tracker Barrel", JINST 3 (2008) P10006

[8] ATLAS Collaboration, "Fast TracKer (FTK) Technical Design Report", CERN-LHCC-2013-007, ATLAS-TDR-025

[9] https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ComputingandSoftwarePublicResults

[10] https://twiki.cern.ch/twiki/bin/view/AtlasPublic/FTKPublicResults

[11] A. Annovi *et al.*, "Variable resolution Associative Memory optimization and simulation for the ATLAS FastTracker project", PoS RD13 (2013) 014.

[12] ATLAS Collaboration, "A new "Variable Resolution Associative Memory" for High Energy Physics", ATL-UPGRADE-PROC-2011-004

[13] A. Annovi *et al.*, "AM06: the Associative Memory chip for the Fast TracKer in the upgraded ATLAS detector", JINST 12 (2017) C04013

[14] J. Adelman *et al.*, "ATLAS FTK challenge: simulation of a billion-fold hardware parallelism", ATL-DAQ-PROC-2014-030

[15] C.-L. Sotiropoulou, *et al.*, "A Multi-Core FPGA-based 2D-Clustering Implementation for Real-Time Image Processing", IEEE Trans. Nucl. Sci. 61 (2014) 3599