Performance Studies of the Associative Memory System of the ATLAS Fast Tracker

Tomoya Iizawa and Xin Wu

Université de Genève,

24 rue du General-Dufour, 1211, Geneva, 4, Switzerland

Email: Tomoya.Iizawa@cern.ch

Chiara Roda, Milene Calvetti,

Marco Piendibene, and Paolo Mastrandrea

Università di Pisa and INFN

Edificio C, Largo Bruno Pontecorvo, 3, 56127 Pisa, Italy

Alberto Annovi, Calliope-Louisa Sotiropoulou, *Member IEEE*

Nicolò Vladi Biesuz, and Paola Giannetti

INFN Pisa,

Edificio C, Largo Bruno Pontecorvo, 3, 56127 Pisa, Italy

Agostino Lanza, Simone Sottocornola,

and Andrea Negri

Università di Pavia and INFN,

Via Agostino Bassi, 6, 27100 Pavia, Italy

Abstract—The Associative Memory (AM) system of the Fast Tracker (FTK) processor has been designed to perform pattern matching using the hit information of the ATLAS experiment silicon tracker. The system is one of the two main processing elements of FTK and is mainly based on the use of Application Specific Integrated Circuits, the AM chips, specifically designed to execute pattern matching with a high degree of parallelism. It finds track candidates at low resolution that are seeds for a full resolution track fitting. The AM system implementation is based on a collection of “AM boards”, the “Serial Link Processors” (AMBSLP). The AMBSLP is based on a network of high speed serial links to sustain very high data traffic. It has a high power consumption (~ 250 W) because of its high performance requirements and, therefore, the AM system needs custom power and cooling.

This proceedings reports on the performance studies of the system made with the first production of 64 AMBSLPs integrated in FTK and results from the first ATLAS data during 2018.

I. INTRODUCTION

To obtain further knowledge on the particle physics, precise measurements of the Standard Model (SM) and searches for physics beyond the SM are performed with the Large Hadron Collider (LHC) at European Organization of Nuclear Research (CERN) built in Switzerland and France. Protons are accelerated and collide at a center of mass energy of 13 TeV. Particles which are generated by interactions of the protons are detected by the ATLAS detector [1] which is a combination of tracking detectors, calorimeters, and muon detectors.

At the LHC, proton bunches collide at a rate of 40 MHz, and it is impossible to store all the collision information due to finite size of data storage. The ATLAS experiment implements trigger system which performs event selection to reduce data rate before recording data to the storage. Figure 1 shows a schematic overview of the ATLAS trigger system. The trigger system consists of two stages which are Level-1 (L1) trigger and High Level Trigger (HLT). The L1 trigger is

hardware based system which does fast event selection, while the HLT is software based system performing more precise event selection than the L1 trigger. Since events discarded at the trigger stage cannot be recovered anymore, it is critical to implement efficient selection which keeps as many signal events as possible while rejecting background events.

Fig. 1. An overview of the ATLAS trigger system.

One of the major challenges for the trigger system in high luminosity environment is to cope with the multiple proton-proton interactions per bunch crossing (pile-up) which makes it more difficult to distinguish signals from backgrounds, and deteriorates resolution of observed quantities. The number of simultaneous interactions would be 70 in LHC Run3, starting from 2021. An event display of simulation of the bunch crossing with pile-up as expected in 2021 is shown in Fig. 2. A straight forward way to reduce event rate in high pile-up environment is applying higher threshold on transverse momentum or energy of each physics object, but this also reduces signal events and phase space for later analysis.

Fig. 2. An event display of the bunch crossing with pile-ups from CMS simulation [2].

In this situation, efficient use of track information has large potential to improve the trigger decision whilst maintaining low thresholds for analyses. Since the track information has high resolution and granularity, and it is measured by the detector placed at the most inner part of the ATLAS detector complex, it provides precise information near the collision point. Currently track information is not used at the L1 trigger, and it is used at the HLT but only inside limited regions since it takes much time and the processing power needed for full scan is not compatible with the available CPU resources.

The Fast TrackKer (FTK) [3] is a newly installed electronic hardware system that reconstructs tracks rapidly for an entire event at the trigger level. The FTK consists of several electronics circuit boards controlled by Field Programmable Gate Arrays (FPGAs). A functional overview of the system is shown in Fig. 3. The FTK is implemented between the L1 trigger and the HLT. It receives hit information and calculates track information for entire detector region for all the events which pass the L1 trigger. With the FTK system, the HLT does not need to calculate track information, and then can utilize time and track information provided by the FTK for more sophisticated algorithms which achieves reduction of data rate without raising threshold on transverse momentum or energy.

Fig. 3. Functional overview of the FTK system. [3]

II. ASSOCIATIVE MEMORY BOARD

One of the key algorithms implemented on the FTK for the fast track reconstruction is “pattern matching”, where track candidates are found among pre-calculated track patterns. The ATLAS tracking detector has a layer structure, and each pattern is a set of the detector unit (strip/chip) from each silicon detector layer. The algorithm works as following.

- 1) Prepare track patterns (more than 1 billion patterns) in advance of the operation. Each pattern consists of coarse resolution hit information which is called Super-Strip (SS).
- 2) At the operation, compare the input SS and the SSs in the prepared patterns, and mark all of the matched SSs in the latter.
- 3) Detect patterns which have marked a pre-defined number of layers of the SSs.
- 4) Iterate processes 2) to 4) until all the hit information in an event is readout.

A schematic view of the pattern matching algorithm is shown in Fig. 4. In the left figure, the blue dots stand for detector units, the red dots for units on which charged particles pass in the event (fired), the blue rectangle for detector modules, and the yellow rectangle for SSs. The right figure shows a schematic view of prepared patterns. Here 10 patterns are shown as an example. If a charged particle passes as a red line shown in Fig. 4 left, the SSs of number 6, 8, 11, and 13 are fired. In this case, pattern #4 is detected as a track candidate since all the SSs in the pattern are marked.

Fig. 4. A schematic view of the pattern matching algorithm.

When the input SS matches with the SS in the stored pattern, Set-Reset Flip-Flop (SR-FF) becomes high. Since it is just taking logical “AND”, the SS comparison with all patterns finishes at a single clock cycle after the hit is loaded. Thus pattern matching is completed at the next clock after the final hit in the event is received.

The pattern matching algorithm is performed by the Associative Memory (AM) system which is organized with Associative Memory Board (AMB) [5], [6], Local Associative Memory Board (LAMB), and AM chips. AM chips are ASIC chips designed and optimized specifically for the FTK. The

AMB is a 9U VME board and 4 LAMBs can be mounted on a single AMB. Figure 5 shows a photo of an AMB with 4 LAMBs mounted (left) and a LAMB (right). 4 FPGAs are implemented on an AMB which control dataflow. A network of high speed serial links determines bus distribution on an AMB. These buses are connected to the backplane card through high frequency ERNI P3 connector. The data rate for each serial link is 7 Gb/s at the maximum. 8192 chips are used in total on 128 AMBs.

Fig. 5. A photo of an AMB with 4 LAMBs mounted (left) and a LAMB (right).

The patterns are stored in AM chips [4]. The AM chip has been developed and improved significantly for recent years. The development of the AM chip is summarized in Tab. I. The FTK utilizes AM06 which works at 100 MHz clock cycle and has an ability to store 128k patterns per chip, thus more than 1 billion patterns can be stored to cover entire detector coverage.

TABLE I
A BRIEF SUMMARY OF AM CHIP SPECIFICATION.

Chip	Technology	#Patterns	Clock (MHz)	Purpose
AM03	180 nm	5k	50	CDF SVT (2003)
AM04	65 nm	8k	100	FTK R&D
AM06	65 nm	128k	100	FTK (2018)

III. PERFORMANCE STUDIES OF THE AMB

There are 64 AMBs, 256 LAMBs, and a few spares already produced. This achieves the production goal in 2018. The rest of the boards will be produced during long shutdown of the LHC starting from 2019, and will be ready before the start of FTK operation in 2021. The produced boards are installed in ATLAS counting room underground, which enables performance study in realistic environment.

One of the important factors to be studied is power consumption since it affects board temperature. If it gets too large value, the board temperature also gets high and the board would be damaged. The power consumption is defined by the number of AMBs running in a single identical VME crate as well as the number of bit-flips inside the AMB logic caused by the input data to be processed. The larger the number of AMBs and bit-flips, the larger the power consumption value becomes.

Power consumption and board temperature were measured with full 16 AMBs populating a single VME crate, and using input data which has the expected number of bit-flips as LHC running from 2021. The result of power consumption measurement of a LAMB is shown in Fig. 6. Power consumption of an AMB is 100 W, and the mean power consumption of a LAMB is 37 W so that the total power consumption on a single AM system reaches $100 + 37 \times 4 = 248$ W. With this setup, temperature of LAMBs is measured and shown in Fig. 7. This figure expresses the LAMBs of the highest temperature on each AMB as a function of time. A board would be damaged if the temperature exceeds 100 °C, while the measured values are well below that temperature. Thus board temperature is found to be kept in safe range even in a realistic situation.

LAMB Power consumption

Fig. 6. Power consumption of a LAMB. [7]

LAMB temperatures

Fig. 7. Temperature of LAMBs as a function of time. [7]

The dataflow and the quality of the output of the AMB are tested using a full FTK chain in LHC collisions. As a current result, the dataflow inside AMB is found to be stable and pattern matching performance has rather small error rate. This is achieved by several improvements on all of the hardware,

software and firmware. For example: firmware updates to avoid clock crossing the clock domain, updates of monitoring system in the software to facilitate commissioning, and implementation of an IPC semaphore to avoid conflict between monitoring system and VME control. The measured error rate of pattern matching is below 10^{-4} which has negligible impact on the entire FTK track reconstruction performance. Similar results are obtained with standalone processing on separate 32 AMBs. Further investigations and improvements are being performed.

IV. SUMMARY

FTK is a new electronic hardware system that reconstructs tracks rapidly in the trigger system. One of the key algorithms for the fast track reconstruction is pattern matching which extracts track candidates with using pre-calculated track patterns. The pattern matching is implemented on the AM system which consists of AMB, LAMB, and AM chips. Production and installation of the AM system are going as scheduled.

The performance of AM system has been studied. The power consumption and board temperature are measured and found to be kept controlled with the realistic setup and data occupancy. The dataflow and the quality of the output have been tested and improved. The error rate of pattern matching is less than 10^{-4} which is already small and has negligible impact for the outcomes provided by the FTK system.

REFERENCES

- [1] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, JINST, 3, 08, S8003, 2008.
- [2] <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2>.
- [3] ATLAS Collaboration, *Fast TracKer (FTK) Technical Design Report*, CERN-LHCC-2013-007, ATLAS-TDR-21, 2013.
- [4] A. Annovi and M. Beretta and G. Calderini and F. Crescioli and L. Frontini and V. Liberali and S.R. Shojaii and A. Stabile, *AM06: the Associative Memory chip for the Fast TracKer in the upgraded ATLAS detector*, JINST, 12, 04, C04013, 2017.
- [5] C.-L. Sotiropoulou, I. Maznas, S. Citraro, A. Annovi, L. S. Ancu, R. Beccherle, F. Bertolucci, N. Biesuz, D. Calabr, F. Crescioli, D. Dimas, M. DellOrso, S. Donati, C. Gentsos, P. Giannetti, S. Gkaitatzis, J. Gramling, V. Greco, P. Kalaitzidis, K. Kordas, N. Kimura, T. Kubota, A. Iovene, A. Lanza, P. Luciano, B. Magnin, K. Mermikli, H. Nasimi, A. Negri, S. Nikolaidis, M. Piendibene, A. Sakellariou, D. Sampsonidis, and G. Volpi, *The Associative Memory System Infrastructures for the ATLAS Fast Tracker*, IEEE Transactions on Nuclear Science, vol. 64, no. 6, pp. 1248-1254, June 2017.
- [6] S. Citraro, N. Biesuz, P. Giannetti, P. Luciano, H. Nasimi, M. Piendibene, C.-L. Sotiropoulou and G. Volpi, *Highly Parallelized Pattern Matching Execution for Event Real-Time Reconstruction*, IEEE Transactions on Nuclear Science, vol. 63, no. 2, pp. 1147-1154, April 2016.
- [7] <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/FTKPublicResults>.