

The Associative Memory Serial Link Processor of the ATLAS Fast Tracker Processing System

Calliope-Louisa Sotiropoulou, *Member, IEEE*, on behalf of the ATLAS Collaboration

Abstract— The upgrade to the Trigger and Data Acquisition (TDAQ) system of the ATLAS experiment at the LHC will improve the capability of the detector to select the events with the greatest scientific potential. The Fast Tracker (FTK) is one of the ATLAS TDAQ upgrades that is presently under commissioning. FTK is a custom hardware system that provides the High Level Trigger (HLT) with charged particle tracks reconstructed from hits in silicon detectors at the rate of 10^5 events per second. The main processing element of FTK is the Associative Memory (AM) system that is used to perform pattern matching with a high degree of parallelism. Its implementation is called the AM Board Serial Link Processor (AMBSLP) and it is a very efficient pattern matching machine that handles in parallel massive data samples.

The AMBSLP consists of two types of boards: the Little Associative Memory Board (LAMB), a mezzanine where the AM chips are mounted, and the Associative Memory Board (AMB), a 9U VME motherboard that hosts four LAMB daughter-boards. We report on the complex FPGA firmware design that has been implemented to operate the high degree of parallelism of the board. We also report on the commissioning status of the AMBSLP and on the performance of the Processor during early data taking.

I. INTRODUCTION

THE Associative Memory Serial Link Processor (AMBSLP) implementation developed for the Fast Tracker System (FTK) [1], an ATLAS trigger upgrade [2], is presented in this paper. The FTK system is built to execute a very fast tracking algorithm organized in a 2-level pipelined architecture. The input stage of the FTK system is implemented on the Data Formatter board and the Input Mezzanines that cluster the hits arriving from the silicon detectors. The clustered hits are then forwarded to the Fast Tracker main Processing Unit (PU) that includes the AMBSLP. The AMBSLP implements the pattern matching in the first stage of the pipeline. It uses a large bank of pre-stored patterns of trajectory points, the Associative Memory (AM) bank. It compares low resolution data, coming from the detector in parallel data links, with the AM bank and finds track candidates in real time during the detector readout phase. The second stage receives the track candidates and the full resolution input data to perform full resolution track fitting at the AM output rate.

The AMBSLP is part of the Associative Memory system [3] for the ATLAS experiment. This system is organized into 128 PUs that process the tracker data in parallel, working on data input coming from different sections of the detector. The whole AM system stores 1 billion AM patterns, made of 128 bit each. The PU is made of a 9U VME card, the AM board and a Rear Transition Module, named AUX card (Figure 1). The AMBSLP is the AM board with its four daughter Little Associative Memory Boards (LAMBs). Each VME crate holds 16 PUs and 4 Second Stage Boards (SSBs), one SSB per 4 PUs. The SSB executes the full resolution track fitting.

The ability of the AMBSLP to execute very fast and efficient pattern matching with an enormous degree of parallelism derives from the AM chip [4], an ASIC that operates like a Content-Addressable-Memory with advanced functionality, specifically designed for pattern matching applications. In addition, high performance Field Programmable Gate Arrays (FPGAs) are used for data flow and control. This powerful highly parallel dedicated hardware provides excellent performance, reaching resolutions, efficiencies and fake track rejection approaching full resolution algorithms executed on CPU farms.

The design of the AM system is a challenging task, due to the following factors: (1) the high pattern density of approximately 8 million patterns per board, which requires a large silicon area; (2) the I/O signal congestion at the board level, which requires the use of serial links; and (3) the power limitation due to the cooling system. The system is based on very high density logic: 8000 AM chips are placed in 8 VME crates, organized in 4 racks. The system requires well-designed power distribution and cooling at the level of 250 W per AM board. The system specifications are challenging and lead to a complex integration process that requires rigorous verification techniques.

II. THE AMBSLP




Figure 1: AMBSLP (AMBoard with a LAMB) and AUX card

Manuscript received November 17, 2017.

C.-L. Sotiropoulou is with the University of Pisa and INFN Sezione di Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy. (e-mail: c.sotiropoulou@cern.ch).

The AMBSLP board [5] contains 64 AM chips. To simplify input/output operations, the AM chips are grouped into AM units composed of 16 chips each, the LAMBs (Figure 2).




Figure 2: LAMB: Little Associative Memory Board

The AMBSLP motherboard has been implemented to hold 4 such units (Figure 1). The LAMB and the motherboard communicate through a high frequency and high pin-count connector placed in the center of the LAMB. A network of high speed serial links handles the data distribution to the 64 AM chips and collects the output.

The data traffic is handled by 2 Artix-7 Xilinx FPGAs [6] with 16 Gigabit Transceivers (GTP), each providing fast data transmission. Two separate Xilinx Spartan-6 FPGAs implement the data control logic. The 12 input serial links are merged into the 8 buses received by each AM chip, one bus for each detector layer used for pattern matching. The data distribution and communication with the 64 chips is very challenging. A huge quantity of data must be distributed at high rate (2 Gb/s on each serial link, for a total of 16 Gb/s maximum rate), with extremely large fan-out to the 64 AM chips, the 16 Gb/s becomes 1024 Gb/s.

The data arrive at the input stage of FTK grouped in “events”. These events are fed to the AMBSLP at a maximum rate of 100 kHz. On average, every 10 μ s 8 thousand 16 bit words must reach the patterns through 8 buses and a similarly large number of output words must be collected and sent in output (32 Gb/s maximum output rate). Simulations have shown a typical occupancy of the available bandwidths around 50%. Each input word has to reach the 8 million patterns on the board. The large input fan-out is obtained through 3 levels of serial fan-out chips to reach each of the 64 AM chips and a very powerful data distribution tree inside each AM chip itself. The AM chip compares 8 input 16 bit words with 128 k patterns every 10 ns. Each LAMB has 40 1:4 fan-outs. The placement of chips on the LAMB has been studied and optimized with the goal of minimizing the crossing of the serial links.

Output words are collected from the 16 AM chips in 4 daisy chains. Each AM device has the capability to receive outputs from two other AM chips and to merge them internally with the output found in the chip itself. Each daisy chain has a single output that goes directly to the connector. Each quartet also shares a 100 MHz low jitter clock necessary for the 11 serial links handled by each AM chip.

III. AMBSLP FIRMWARE FUNCTIONS, INTEGRATION AND RESULTS

A. Firmware functions

Due to the high complexity of the FTK system a number of advanced monitoring functions have been implemented in the AMBSLP firmware. These are:

- **Parallel input stream synchronization control:** checks that all input streams propagate the same event and in the case this is not true it tags as valid the event located at the majority of the input streams.
- **Event number monitoring:** checks if an event is skipped or missing.
- **Back-pressure control for board-to-board communication:** if the board processing is slowed down by a following board a signal is sent to the AMBSLP and pauses the data flow.
- **Real time temperature monitoring:** temperature is checked in real time through hardware and firmware. There are two temperature sensors on each LAMP. Temperature monitoring can stop AMBSLP operation in case a preset threshold is passed.
- **Real time error monitoring of the system:** control of the input data format.
- **Timing measurements:** latency in data distribution and processing is measured from the input to the AM chips and from the AM chips to the AMBSLP exit.

All the above features are supported by software that allows remote control and monitoring of the AMBSLP.

B. Integration

A full FTK slice processing up to 12 layer tracks is under commissioning. The AMBSLP is currently being integrated at CERN. A rigid testing procedure is implemented before each new component is included in the system. Testbeds are set up in different locations, in Pisa and at CERN, to test hardware, firmware and software combinations. At CERN a full processing chain FTK test system that includes one board per type has been assembled. The firmware is first tested in Pisa, then tested in the full slice in CERN. If all tests are successful, it is then deployed in the primary system in the ATLAS service cavern.

C. Results

A partial FTK system slice including the AMBSLP (from the FTK input interface card up to the AMBSLP) can process real data for hours. This system includes the input system that clusters the particle hits, the pattern matching function of the AMBSLP and the first level track fitting process that is executed in the AUX. Further commissioning is ongoing.

Currently there are ~200 LAMBs at CERN and ~70 AMBs. The target for the initial FTK system is 64 AMBs available with 256 LAMBs and spare boards for back up. This will cover the full ATLAS inner tracking detector for Run 2 with half of the full FTK computing potential.

CONCLUSIONS AND FUTURE GOALS

The AMBSLP of the ATLAS FTK system is a complex pattern matching processor that is currently being

commissioned for the ATLAS TDAQ system upgrade. It consists of a mother board, the AMB, and 4 daughter boards the LAMBs. A rigid testing procedure has been set up to validate hardware, firmware and supporting software for the system. A partial FTK system that includes the AMBSLP can currently process data sent from the ATLAS tracking detectors for hours. The current goals of the FTK system are to establish stable full slice processing and validate 8-layer pattern recognition output with data and simulation. For 2018 the goal for the system is to operate during ATLAS data taking with 50% of the final PU computing potential.

REFERENCES

- [1] The FTK collaboration, "Technical Design Report Fast TracKer (FTK)" CERN-LHCC-2013-007 ATLAS-TDR-021 available online: <https://cds.cern.ch/record/1552953/files/ATLAS-TDR-021.pdf>
- [2] The ATLAS Collaboration, "The ATLAS Experiment at the CERN Large Hadron Collider," *Journal of Instrumentation* 3 S08003, 2008.
- [3] C. L. Sotiropoulou *et al.*, "The Associative Memory System Infrastructures for the ATLAS Fast Tracker," in *IEEE Trans. on Nuclear Science*, vol. 64, no. 6, pp. 1248-1254, June 2017.
- [4] A. Annovi, et al., "AM06: the Associative Memory chip for the Fast TracKer in the upgraded ATLAS detector," 2017 *JINST* **12** C04013.
- [5] S. Citraro *et al.*, "Highly Parallelized Pattern Matching Hardware for Fast Tracking at Hadron Colliders," in *IEEE Trans. on Nuclear Science*, vol. 63, no. 2, pp. 1147-1154, April 2016.
- [6] Xilinx Inc, "7 Series FPGAs Data Sheet: Overview", available online: https://www.xilinx.com/support/documentation/data_sheets/ds180_7Series_Overview.pdf