

The Associative Memory Boards for the FTK Processor at ATLAS

D. Calabro¹, R. Cipriani, S. Citraro, S. Donati, P. Giannetti, A. Lanza, P. Luciano, D. Magalotti and M. Piendibene

Abstract— The Associative Memory (AM) system, a major component of the FastTracker (FTK) processor, is designed to perform pattern matching using the information from the silicon tracking detectors of the ATLAS experiment. It finds track candidates at low resolution that are sent to the track fitting stage. The system has to support challenging data traffic, handled by a group of modern low-cost FPGAs, the Xilinx Artix 7 chips, which have Low-Power Gigabit Transceivers (GTPs). Each GTP is a combined transmitter and receiver capable of operating at data rates up to 7 Gb/s.

The paper reports on the design and initial tests of the most recent version of the AM system, based on the new AM chip design which uses serialized I/O. An estimation of the power consumption of the final system is also provided and the cooling system design is described. The first cooling test results are reported.

I. INTRODUCTION

THE FTK processor, organized in a pipelined architecture, executes a very fast algorithm based on the use of a large bank of pre-stored track patterns, the AM bank. The AM system carries out pattern matching using the detector hits at low resolution. The track candidates found by the AM are fit using the full resolution hits.

The AM system consists of the AM chip, an ASIC designed and optimized for this particular application, and two types of boards: the Local Associative Memory Board (LAMB), a mezzanine where the AM chips are mounted, and the 9U VME Associative Memory Board (AMB), which holds the LAMBS. Both the AM chip [1] and the two boards have a long development history [2]-[5]. We report here about the latest versions, specifically developed for FTK and denoted LAMBSLP (LAMB with Serial Link Processor) and AMBSLP, designed for the final AM chip which will entirely use serialized input/outputs.

We also report on preliminary tests of power consumption and cooling, both executed and planned.

Manuscript received November 15, 2013. This work was supported by the Istituto Nazionale di Fisica Nucleare (I.N.F.N.), Italy, and by the FTK 324318 FP7-PEOPLE-2012-IAPP Grant, European Union.

D. Calabro¹ (domenico.calabro@pv.infn.it) and A. Lanza (agostino.lanza@pv.infn.it) are with I.N.F.N. Pavia, via A. Bassi 6, 27100 Pavia, Italy.

R. Cipriani (ciprianircicardo@gmail.com), S. Citraro (saveriocitraro86@gmail.com), S. Donati (simone.donati@gmail.com), P. Luciano (pierluigiluciano@gmail.com) and M. Piendibene (marco.piendibene@pi.infn.it) are with University of Pisa and I.N.F.N. Pisa, largo Pontecorvo 3, 56127 Pisa, Italy.

P. Giannetti (paola.giannetti@pi.infn.it) is with I.N.F.N. Pisa, largo Pontecorvo 3, 56127 Pisa, Italy.

D. Magalotti (daniel.magalotti@pg.infn.it) is with Modena and Reggio Emilia University, via Universita' 4, 41121 Modena, Italy

II. THE AMBSLP AND LAMBSLP

The LAMBFTK, designed in the past for the AM chip with parallel I/O, had a challenging layout with a high degree of complexity, shown in Fig. 1, because of the presence of the parallel buses. Moreover, it required many FPGAs (the light boxes in Fig. 1) for handling the input fan-out and the output merging. In order to minimize the impact of these weaknesses on the system functionality, a new approach was developed: the LAMBSLP.

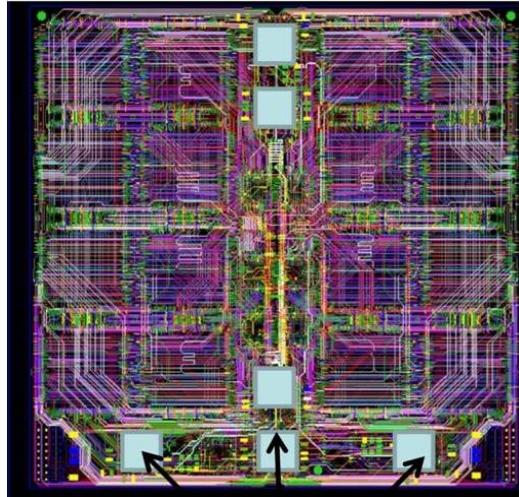


Fig. 1. Layout of the old LAMBFTK.

This new mezzanine is designed to match the new AM chip I/O, which is totally serialized. The FPGAs can be replaced by simple low-cost serial links and high-fan-out chips. Its motherboard, the AMBSLP, is a 9U VME board on which 4 LAMBSLPs are installed. Fig. 2 shows the AMBSLP, highlighting the LAMBSLP mezzanine positions in yellow. A network of high speed serial links characterizes the bus distribution on the AMBSLP: 12 input serial links (in blue) that carry the silicon detector hits from the high-frequency ERNI P3 connector (in orange) to the LAMBSLPs, and 16 output serial links (each red arrow represents 4 links) that carry the matched road numbers from the LAMBSLPs to P3.

These buses are connected through P3 to a board denoted AUX that sits in the same slot on the back side of the 9U VME crate. The board on the back side performs the full resolution track fitting, refining the AMBSLP tracks.

The data rate is up to 2 Gb/s on each serial link. Thus the AMBSLP has to handle a challenging data traffic rate: a huge number of silicon detector hits must be distributed at high rate

with very large fan-out to all patterns (8 million patterns will be located on 64 AM chips on a single AMBSLP) and a similarly large number of roads must be collected and sent back to the FTK full resolution track fitting function.

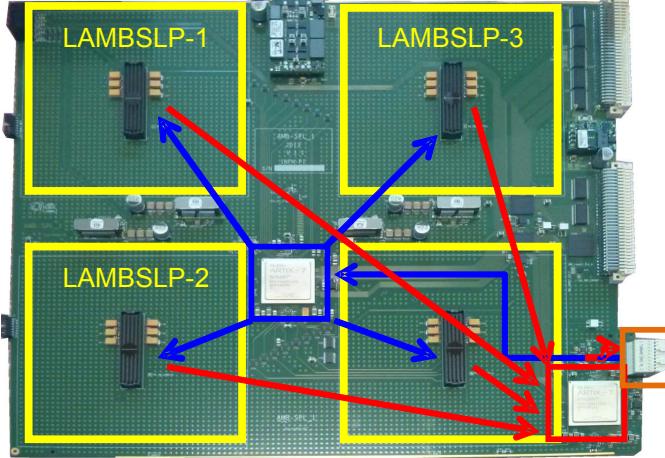


Fig. 2. The AMBSLP with the positions of the four LAMBSLP highlighted in yellow. The data traffic is represented with blue arrows for the input, and with red arrows for the output.

The AMBSLP has flexible control logic placed inside a group of FPGA chips visible in Fig. 2. They are 2 large Xilinx Artix 7 [6] for data transfer and Xilinx Spartan 6 [6] for control and VME logic. The Artix 7 have 16 GTPs operating at data rates up to 7 Gb/s. This ultra-fast data transmission requires specialized, dedicated on-chip circuitry and differential I/O capable of coping with the signal integrity issues.

The incoming hits are received by the GTPs in the input FPGA (blue box) and saved in large de-randomizing FIFOs. The Control FPGA handles the event processing and error conditions. When the Control chip starts processing an event, hits are popped in parallel from all the hit input FIFOs and are simultaneously sent to the four LAMBs. An End Event (EE) word, which includes the event tag, separates hits belonging to different events. Data in different streams have to be synchronized to guarantee that hits belonging to the same event are being processed by the AM patterns. If occasionally a FIFO becomes Almost Full (AF), a HOLD signal is sent to the upstream board, suspending the data flow until more FIFO locations become available. The AF threshold and the size of the FIFOs are set to give enough time to the upstream board to react and avoid frequent backpressure.

The LAMBSLP and the AMBSLP communicate through a high-pin-count, high-performance connector (HPC) type ASP-134488-01, placed in the center of the LAMBSLP (the connector is inside the 4 yellow boxes in Fig. 2).

Each LAMBSLP contains 16 AM chips but no FPGAs, as shown in Fig. 3, simplifying the board layout and the functionality tests.

For each LAMBSLP, large currents (up to ~ 32 A at 1.0 V, ~ 8 A at 2.5 V and ~ 1.8 A at 1.2 V) are provided by three types of DC-DC converters through the HPC connector. The total power dissipation for each LAMBSLP is about 50 W.

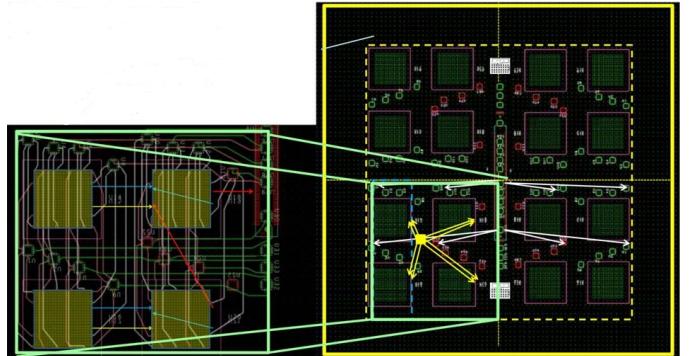


Fig. 3. Layout of the LAMBSLP showing four AM chips and no FPGAs.

Final prototypes of this new mezzanine will be available in the next months. In the meanwhile, we have realized a mini-ASIC AM chip using the final serialized I/O, which can be installed in a new mezzanine denoted miniLAMB, specifically designed to test the network of serial connections and the compatibility between the FPGAs, the new fan-out chips and the new AM chip I/O.

III. THE AMBSLP AND MINILAMB TESTS

The new miniLAMB mezzanine is shown in the Fig. 4. Four small mini-ASIC AM chips placed in QFN64 packages, since they are too small to use the final BGA package, are visible at the bottom right. In the center of this group of mini-ASICs, there are the low-jitter oscillator and the fan-out chip that distributes the 100 MHz clock to them. Above the group there are the fan-out chips for distribution of the hit serial buses. Two FPGAs are on board, one to program the mini-ASICs, the other for test purposes.

The serial links of the mini-ASICs were tested in the lab. Their signals were correctly transmitted to the miniLAMB, correctly received by the FPGA on the board and correctly transmitted to the mini-ASICs through the fan-out. The next step will be the test of the small AM bank (128 patterns) inside the mini-ASICs.

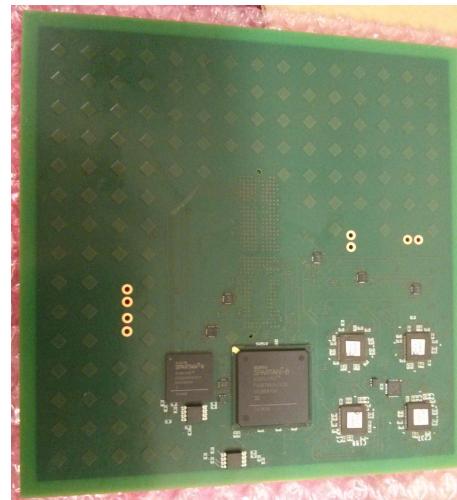


Fig. 4. The miniLAMB prototype developed to test the GTP network on four mini-ASIC AM chips.

IV. THE COOLING TESTS

The final FTK AM system will be composed of 512 LAMBSLPs installed on 128 AMBSLPs. They will be contained in eight 9U VME core crates, each one with 16 AMBLSPs.

Since the back side of the core crates is occupied by the AUX boards, the power supply (PS) of each core crate must be positioned in the rack on top of the crate, increasing the rack space necessary to host a core crate to 19U at least, including the fan tray and the heat exchangers on top of the crate and the PS. As a result, no more than two core crates can be contained in one rack, meaning four racks for the full AM system.

Fig. 5 represents the proposed core crate layout (a) and the rack layout (b) for the final AM system.

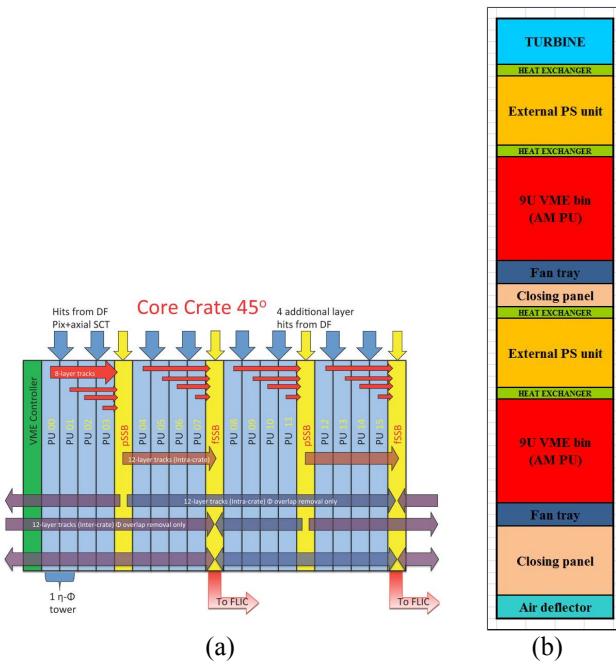


Fig. 5. Core crate layout (a) and rack layout (b)

The present baseline for the rack layout is to use one PS for each crate, but we are studying the feasibility of using a single PS for the two crates, increasing the transparency of the stack to the turbine forced ventilation [7].

The estimated core crate power consumption for the final system is about 5k W, which makes the design of the rack cooling system challenging.

In order to measure the cooling performance of the present fan trays, and to study possible improvements, we have built a dedicated setup in the lab. The overall performance in the ATLAS electronics room will be investigated in the near future.

The presently used PS is a 6k W Wiener UEP6021 unit which provides 5VDC@345A, 3.3VDC@115A and 48V@81A. The cooling tests were performed with 15 resistive 9U VME load boards. On eight of them we placed six T sensors (National Semiconductor LM35) in different positions, manually read by means of a voltmeter. Fig. 6 shows one of the load boards equipped with the six T sensors, denoted UF

(Upper Front), UR (Upper Rear), UC (Upper Center), LC (Lower Center), LF (Lower Front) and LR (Lower Rear).

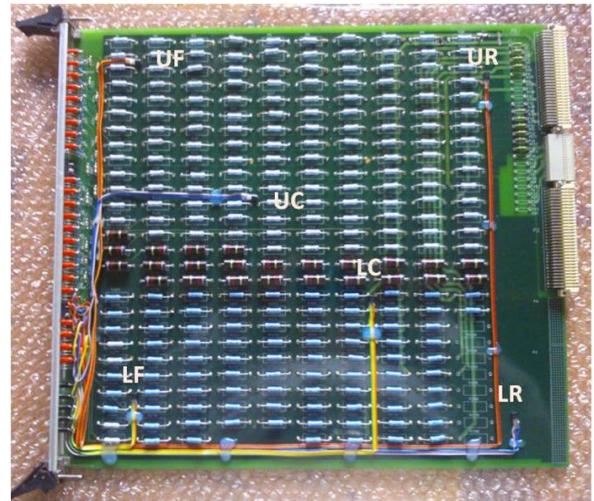


Fig. 6. 9U VME load board used for tests with the six installed T sensors.

PS currents were detected by three power shunts in series with the cables connecting the PS with the backplane, and read with voltmeters. The crate was cooled by one heat exchanger supplied by a chiller. The test setup is visible in Fig. 7.



Fig. 7. The setup used for the test. To the left there are the voltmeters reading the PS currents and the T sensors, in the center the rack with, from the bottom, the fan tray, the bin filled in with the load boards, the heat exchanger and the PS, and to the right the chiller.

We measured the performance of two different fan trays, the first an old one borrowed from the CDF experiment and equipped with high-flux fans, the second a Wiener UEL6020A equipped with six standard fans.

We used the CDF fan tray as the reference, after having optimized its performance adding one fan in the center front region where there was insufficient ventilation.

The Wiener fan tray performance was compared to the modified CDF fan tray. After the first tests we were able to replace one of its six standard fans with a special Hyper-Blower fan, provided by Wiener, which delivers three times the air flux with respect to the standard fans. We are waiting for a new Wiener fan tray fully equipped with Hyper-Blower fans.

Fig. 8 shows the modified CDF fan tray (a) and the modified Wiener fan tray (b).



(a)



(b)

Fig. 8. The reference CDF fan tray (a) with the addition of the small fan to the left, and the Wiener fan tray (b) with an Hyper-Blower fan to the bottom left.

The load boards were installed in the VME crate together with two old AMBFTK boards in the positions visible in Fig. 9 by the black panels. The total power consumption measured by the shunts, and confirmed by the Wiener fan tray monitor, is 4.5k W.

The plots in Fig. 10 show the temperature measured at the six T sensor positions for the eight load boards, identified with different colors, respectively for the CDF modified fan tray (a) and for the standard Wiener fan tray (b). Notice the different temperature scales. The results of the Wiener fan tray with one Hyper-Blower fan are shown in Fig. 11 together with the comparison plot with respect to the standard fan tray.

The very hot spot at position UF of the Load 8 – Slot 11 was not present in the previous tests and was not understood, but that T sensor was shown to work fine after the test.

The Hyper-Blower showed a good cooling performance for the load boards under which it was installed.

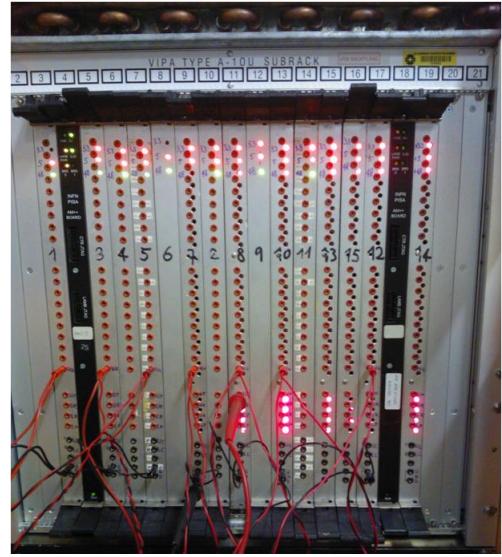


Fig. 9. Positions of the load boards inside the VME crate. Numbers on the front panels correspond to the Load #. The total power consumption of this crate was measured to be 4.5k W.

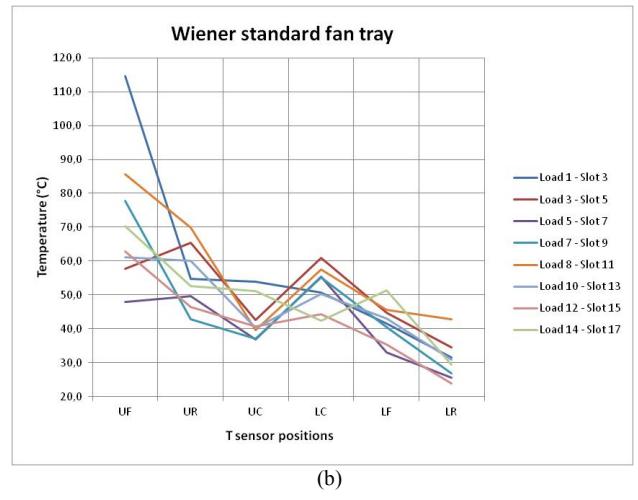
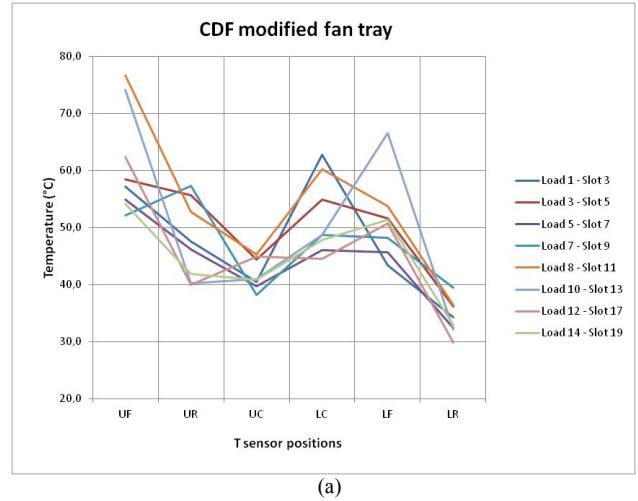
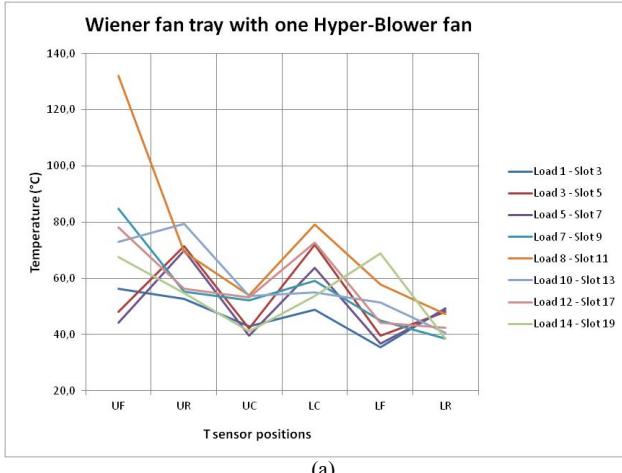
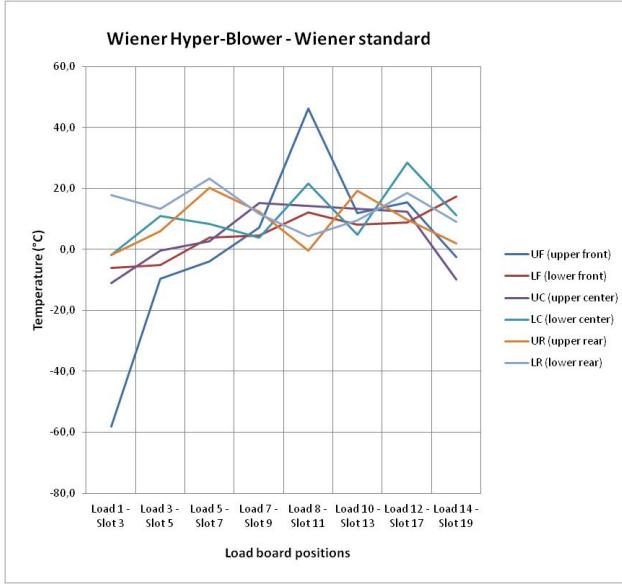


Fig. 10. Temperature performance vs T sensor positions of the modified CDF fan tray (a) and the standard Wiener fan tray (b).



(a)



(b)

Fig. 11. Results of the modified Wiener fan tray (a) and temperature difference between the modified and the standard fan trays (b).

The test setup was recently moved to the ATLAS electronics room and installed in a rack fully equipped with the ATLAS cooling system, as shown in Fig. 12.



Fig. 12. The cooling test setup installed in the ATLAS electronics room USA15.

V. CONCLUSIONS

We are testing the new AM system, integrating the new boards (AMBSLP and LAMBSLP) and the new AM chip.

A new LAMB board, the miniLAMB, was built in order to test the performance of the GTP links of the new AM chip, and passed the first lab tests. The next step will be the test of one complete AMBSLP + LAMBSLP board.

The final crate and rack layout was defined. In order to prove the effectiveness of the cooling system, a test setup was built and several tests on standard and modified fan trays were conducted. The partial results obtained with a high-flux Wiener fan are good, but the tests will be repeated when a full fan tray equipped with these fans become available. The test stand is now installed in the ATLAS electronics room to perform these tests.

REFERENCES

- [1] A. Annovi et al., "A VLSI Processor for Fast Track Finding Based on Content Addressable Memories", *IEEE Trans. Nucl. Sci.*, vol. 53, pp 2428, 2006.
- [2] A. Annovi et al., "The AM++ board for the silicon vertex tracker upgrade at CDF", *Nuclear Science Symposium Conference Record, 2005, NSS '05, IEEE Volume 1*, Page(s): 598 – 602.
- [3] G. Batignani et al., "The associative memory for the self-triggered SLIM5 silicon telescope", *Nuclear Science Symposium Conference Record, 2008. NSS '08, IEEE*, Page(s): 2765 – 2769.
- [4] A. Andreani et al., "The AMchip04 and the processing unit prototype for the FastTracker", *IOP J. Instr.* 7 (2012) C08007.
- [5] A. Annovi et al., "Next generation Associative Memory devices for the FTK tracking processor of the ATLAS experiment", submitted to this conference.
- [6] <http://www.xilinx.com/>.
- [7] G. Thomas, P. Vive Roux Fontaine, V. Pittin, "LHC cooling measurements", CERN Electronics Pool report, 2002.