

Highly Parallelized Pattern Matching Execution for the ATLAS Experiment

A. Annovi, F. Bertolucci, N. Biesuz, D. Calabrò, S. Citraro, F. Crescioli, D. Dimas, M. Dell'Orso, S. Donati, C. Gentsos, P. Giannetti, S. Gkaitatzis, V. Greco, P. Kalaitzidis, K. Kordas, N. Kimura, A. Lanza, P. Luciano, B. Magnin, I. Maznas, K. Mermikli, H. Nasimi, S. Nikolaidis *Senior Member IEEE*, M. Piendibene, A. Sakellariou, D. Sampsonidis, C.-L. Sotiropoulou *Member IEEE*, G. Volpi, G. Xiotidis.

Abstract— The Associative Memory (AM) system of the Fast TracKer (FTK) processor has been designed to perform pattern matching using as input the data from the silicon tracker in the ATLAS experiment. The AM is the primary component of the FTK system and is designed using ASIC technology (the AM chip) to execute pattern matching with a high degree of parallelism. The FTK system finds track candidates at low resolution that are seeds for a full resolution track fitting. The AM system implementation is named “Serial Link Processor” and is based on an extremely powerful network of 2 Gb/s serial links to sustain a huge traffic of data.

This paper reports on the design of the Serial Link Processor consisting 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 which hosts four LAMB daughterboards. We also report on the performance of the prototypes (both hardware and firmware) produced and tested in the global FTK integration, an important milestone to be satisfied before the FTK production.

I. INTRODUCTION

THE implementation presented in this paper is developed for the Fast TracKer Processor (FTK) [1], which is an approved ATLAS trigger upgrade. The FTK processor [2] executes a very fast tracking algorithm organized in a 2-level pipelined architecture. The AM system implements the pattern matching algorithm in the first stage of the pipeline. It uses a large bank of stored patterns of trajectory points, the AM bank. It compares the detector data of resolution reduced by FTK preprocessing stage, 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 high resolution input data to perform high resolution track fitting at the AM output rate.

A key role in the FTK architecture is played by high-performance field programmable gate arrays (FPGAs), while most of the computing power is provided by full-custom ASICs named AM chip [3]. Powerful highly parallel dedicated hardware provides excellent performance, reaching resolutions, efficiencies typical of high quality algorithms executed on CPU farms. The FTK system is optimized for short latencies (few tens of microseconds), low power usage (the AM chip, a device able to execute 1 Million comparisons of 16 bit words each 10 ns, has a power consumption below 3 W). In addition, the system is small (4 racks of electronics are able to perform a task that would need a farm of thousands of commercial CPUs).

The AM system for the ATLAS experiment is organized into 128 Processing Units (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 bits each. The PU consists of a 9U VME card, the AM board and a Rear Transition Module, named AUX card. The AM board holds 64 AM chips. The AUX card is placed in the same slot of the VME core crate (see Figure 1) as the AM board and communicates with the AM board through a high density high speed connector providing the input data and collecting the fired patterns.

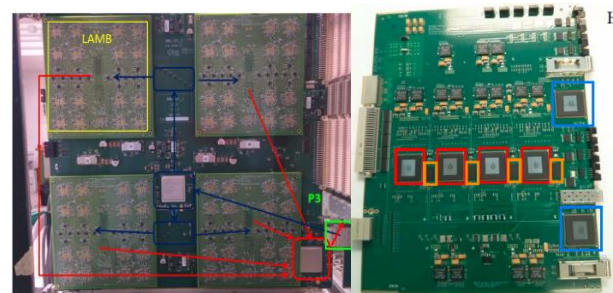


Figure 1: AMBSLP and AUX card

The design of the AM system is a challenging task, due to the following factors: (1) the high pattern density, 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

Manuscript received November 15, 2015. This work receives support from Istituto Nazionale di Fisica Nucleare. The project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement n.324318.

N. Biesuz, S. Citraro, M. Dell'Orso, S. Donati, M. Piendibene, H. Nasimi, C.-L. Sotiropoulou are with the University of Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy.

A. Annovi, F. Bertolucci, P. Giannetti, P. Luciano, G. Volpi are with the Sezione di Pisa INFN, Largo Bruno Pontecorvo 3, 56127 Pisa, Italy.

F. Crescioli is with Laboratoire de Physique Nucléaire et de Hautes Energies, Couloir 12-22 étage 4 Place Jussieu, 75005 Paris, France.

D. Dimas, P. Kalaitzidis, K. Mermikli, A. Sakellariou are with Prisma Electronics SA, El Venizelou 128 Nea Smyrni, 17123 Athens, Greece.

V. Greco, B. Magnin are with CERN, CH-1211 Geneva 23, Switzerland.

A. Lanza, D. Calabrò are with Sezione di Pavia INFN, Via Agostino Bassi 6, 27100 Pavia, Italy.

C. Gentsos, S. Gkaitatzis, K. Kordas, N. Kimura, I. Maznas, S. Nikolaidis, D. Sampsonidis, G. Xiotidis are with Aristotle University of Thessaloniki, 54124 Thessaloniki, Greece.

system. The 8000 AM chips are fit in 8 VME crates and 4 racks and the power should not exceed 250 W per AM board.

II. THE BOARDS: PUTTING CHIPS TOGETHER

A 9U-VME board populated with 64 AM chips can hold 8 million patterns, each one made of 128 bits. To simplify input/output operations, the AM chips are grouped into AM units composed of 16 chips each, called Little Associative Memory Boards (LAMB, Figure 2).

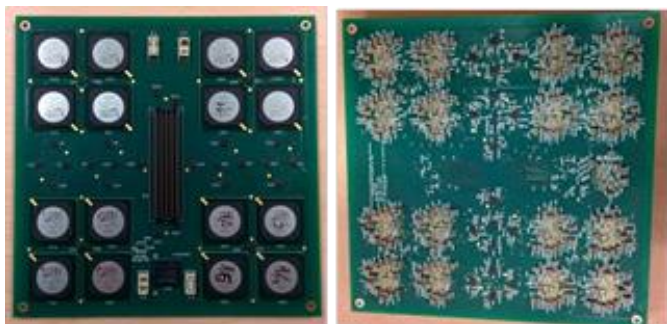


Figure 2: LAMB: Little Associative Memory Board

A 9U-VME 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 made of ~750 point-to-point connections handles the data distribution to the 64 AM chips and collects the output. Twelve input serial links (shown in blue) carry the detector data from the P3 (high density connector in the green box Fig. 1), and 16 output serial links (4 links from each LAMB, represented by a red arrow in the figure) carry the identified patterns from the LAMBS to P3.

The data traffic is handled by 2 Artix-7 Xilinx FPGAs with 16 Gigabit Transceivers (GTP) each, providing ultra-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 to each AM chip, one bus for each detector layer used for pattern matching. The data distribution is very challenging. A huge quantity of data must be distributed at high rate with extremely large fan-out to the 64 AM chips. The global data rate to the chips is 1024 Gb/s.

The data arrive at the input stage of FTK packed as events. These events are fed to the board at a maximum rate of 100 kHz. On average, every 10 μ s eight thousand words (128 bits) must reach the patterns through 8 buses and a similarly large number of output words must be collected and sent back to the P3 (32 Gb/s maximum output rate). Each input 128 bit 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 locations 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 other two AM chips and 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.

Particular care has been given to the PCB routing, in particular for the many serial links (~750 links) as keeping the relative impedance fixed at 100 Ω and to minimize the cross talk. The PCB has 12 layers, where signal planes and power-GND planes are alternated. The serial links are all routed into internal layers, so that they are isolated between two metal planes. In addition they are shielded from other lines in the same plane by metal ground fill.

III. RESULTS

All the serial links internal to the AM board were tested and also the ones connecting the AM with the AUX inside the PU before producing the final prototype. The dependence of the signal quality on the link length and design method was observed and the production PCB was optimized in this respect. The eye diagram of the typical link after the optimization process can be seen in Figure 3. The bit error rate (BER) was tested directly using PRBS-7 generator and was found to be less than 10^{-14} (estimation by the oscilloscope is $BER \sim 10^{-22}$).

To test the global functionality of the system a Random Test was used, that generates events containing random input data in order to test rare conditions that could escape standard tests. This test is important because it performs a realistic simulation of the AM system dataflow and provides a tool that allows comparing the observed fired patterns with the expected ones. The AM Board has been successfully tested with the random events data received from the AUX card. After successful tests the board has been integrated in the FTK Global Integration test at CERN. Production system will be installed in the experiment to take data for the first time at the end of 2015.

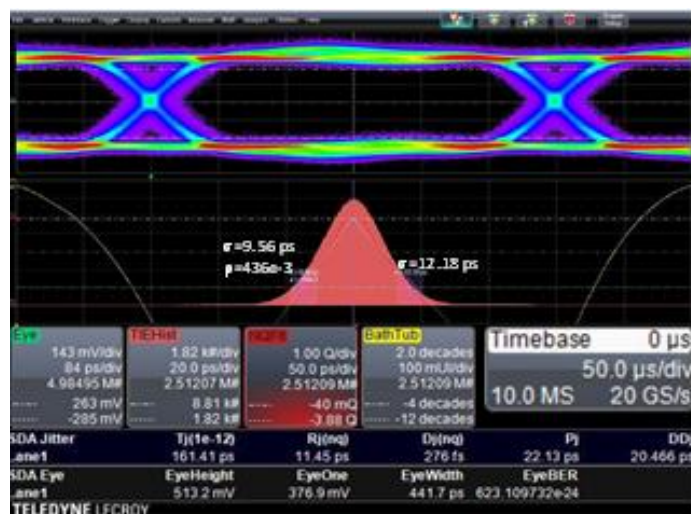


Figure 3: Serial data link analysis

IV. CONCLUSIONS

A powerful, highly parallelized pattern matching system is presented. The system exploits dedicated hardware to provide excellent performance, reaching resolutions, efficiencies typical of high quality algorithms executed on CPU farms. The system achieves very short latencies (few tens of microseconds) and is able to execute 1 billion comparisons of 128 bit words each 10 ns. The system is much more compact than its CPU equivalent (4 racks of electronics are able to perform a task that would need a farm of thousands of commercial CPUs).

REFERENCES

- [1] The ATLAS Collaboration: Fast TracKer (FTK) Technical Design Report; CERN-LHCC-2013-007; ATLAS-TDR-021 <https://cds.cern.ch/record/1552953>
- [2] Andreani et al., The FastTracker Real Time Processor and Its Impact on Muon Isolation, Tau and b-Jet Online Selections at ATLAS, 2012 TNS Vol.: 59 , Issue:2, pp, 348 – 357
Annovi, A. et al., A VLSI Processor for Fast Track Finding Based on Content Addressable Memories, *IEEE Trans. Nucl. Sci.*, vol. 53, pp 2428, 2006.