

The Associative Memory Serial Link Processor for the Fast TrackKer (FTK) at ATLAS

This content has been downloaded from IOPscience. Please scroll down to see the full text.

2014 JINST 9 C11006

(<http://iopscience.iop.org/1748-0221/9/11/C11006>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 188.184.3.56

This content was downloaded on 10/06/2015 at 07:36

Please note that [terms and conditions apply](#).

WORKSHOP ON INTELLIGENT TRACKERS,
14–16 MAY 2014,
UNIVERSITY OF PENNSYLVANIA, U.S.A.

The Associative Memory Serial Link Processor for the Fast TracKer (FTK) at ATLAS

A. Andreani,^{a,b} A. Annovi,^{c,2} R. Beccherle,^{d,e} M. Beretta,^c N. Biesuz,^{f,3} W. Billereau,^g R. Cipriani,^{e,f} S. Citraro,^{e,f} M. Citterio,^a A. Colombo,^b J.M. Combe,^g F. Crescioli,^d D. Dimas,^h S. Donati,^{e,f} C. Gentsos,ⁱ P. Giannetti,^e K. Kordas,ⁱ A. Lanza,^j V. Liberali,^{a,b,1} P. Luciano,^{e,f} D. Magalotti,^{k,l} P. Neroutsos,ⁱ S. Nikolaidis,ⁱ M. Piendibene,^{e,f} E. Rossi,^{f,4} A. Sakellariou,^h S. Shojai,^{a,b} C.-L. Sotiropoulou,ⁱ A. Stabile^a and P. Vulliez^g

^aINFN — Sezione di Milano, Via G. Celoria 16, 20133 Milano, Italy

^bDipartimento di Fisica, Università degli Studi di Milano, Via G. Celoria 16, 20133 Milano, Italy

^cINFN — LNF, Via E. Fermi 40, 00044 Frascati, Italy

^dLaboratoire de Physique Nucléaire et de Hautes Energies (LPNHE), 4 place Jussieu, 75252 Paris, France

^eINFN — Sezione di Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy

^fDipartimento di Fisica, Università degli Studi di Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy

^gCERN, 1211 Geneva 23, Switzerland

^hPrisma Electronics SA, El. Venizelou 128, Nea Smyrni, 17123 Athens, Greece

ⁱDepartment of Physics, Aristotle University of Thessaloniki, 54124 Thessaloniki, Greece

^jINFN — Sezione di Pavia, Via A. Bassi 6, 27100 Pavia, Italy

^kUniversità di Modena e Reggio Emilia, Via Università 4, 41121 Modena, Italy

^lINFN — Sezione di Perugia, Via A. Pascoli, 06100 Perugia, Italy

E-mail: valentino.liberali@mi.infn.it

ABSTRACT: The Fast TracKer (FTK) is an extremely powerful and very compact processing unit, essential for efficient Level 2 trigger selection in future high-energy physics experiments at the LHC. FTK employs Associative Memories (AM) to perform pattern recognition; input and output data are transmitted over serial links at 2 Gbit/s to reduce routing congestion at the board level. Prototypes of the AM chip and of the AM board have been manufactured and tested, in preparation of the imminent design of the final version.

KEYWORDS: VLSI circuits; Trigger concepts and systems (hardware and software); Digital electronic circuits

¹Corresponding author.

²Now at: INFN — Sezione di Pisa, 56127 Pisa, Italy.

³Now at: INFN — Sezione di Pisa and Dipartimento di Fisica, Università degli Studi di Pisa, 56127 Pisa, Italy.

⁴Now at: INFN — LNF, 00044 Frascati, Italy.

Contents

1	Introduction	1
2	The FTK system	1
2.1	The AM system	2
2.2	The FTK AM integrated circuit	3
2.3	The AM05	4
2.3.1	XORAM Memory Layer	5
2.3.2	LV NAND-NOR Memory Layer	6
3	Serial link characterization and current consumption measurements	6
3.1	Serial link measurements	6
3.2	Current measurements	8
4	The FTK AM board	8
4.1	FTK mini-LAMB prototype and final LAMB design	8
4.2	Mini-LAMB prototype test	9
5	Conclusion and future work	10

1 Introduction

Experiments at the LHC, such as ATLAS [1], produce a huge amount of data. Since a limited amount of data can be transferred to a storage system for subsequent off-line processing, an enormous data reduction must be performed. To this end, a trigger system is used to recognize interesting events in real time [2]. Tracking devices play an essential role in this trigger selection task and in particular the silicon devices that are becoming the preponderant tracking technology.

Track detection can be performed by comparing data from detectors with a set of pre-computed patterns stored in a memory. The pattern recognition problem can be solved by a dedicated system based on Associative Memories (AM) [3], which exploits parallelism to the highest possible level.

This paper presents the ongoing design of the Associative Memory Serial Link Processor (AMSLP), based on AM chips and boards, with high-speed serial links for internal data exchange.

2 The FTK system

The LHC accelerator will deliver increased instantaneous luminosity in the coming years [4]. The main purpose is to increase the physics output of the LHC experiments. As a side effect it will be more challenging to perform online data reduction. In order to maintain current High Level Trigger (HLT) performance under the harsher conditions, the ATLAS Trigger and Data Acquisition

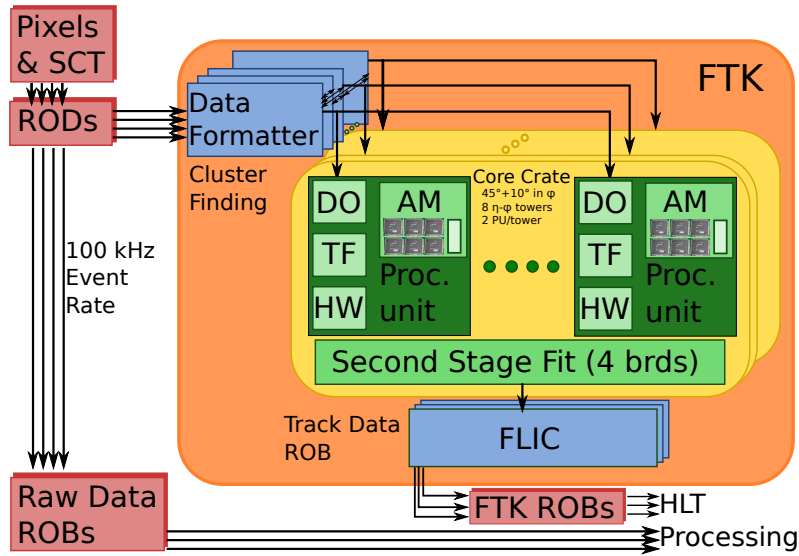


Figure 1. FTK architecture.

(TDAQ) is undergoing several upgrades [5] including the addition of the ATLAS Fast Tracker (FTK) Processor [6]. The FTK processor is designed to provide massive computing power and to minimize the on-line execution time of complex tracking algorithms. The FTK will provide the ATLAS High Level Trigger (HLT) with a complete list of tracks for particles with transverse momentum above 1 GeV/c. It will process all events that are accepted by the level-1 trigger, with an event rate up to 100 kHz and a latency of the order of 100 μ s.

The FTK system is illustrated in figure 1. It is made of 48 Data Formatters (DF), 128 Processing Units, 32 Second Stage Fit Boards, and an interface to the Level 2 Trigger. Each processing unit is composed of an AM Board with 8 millions patterns and a rear transition module with Data Organizer (DO), Track Fitter (TF), and Hit Warrior (HW).

2.1 The AM system

The AM system [7] is the core of the FTK. The whole AM system stores 1 billion (10^9) AM patterns for pattern recognition, it performs pattern matching using the hit information from the ATLAS silicon tracker, and it finds track candidates at low resolution that are seeds for full resolution track fitting.

The AM compares stored patterns to input data received as 16 bit words at a 100 MHz rate in parallel over 8 input channels. At the maximum speed, the overall system will be able to perform $8 \cdot 10^{17}$ comparisons per second in parallel between 16-bit words. The total size of the AM is $8 \cdot 18 \cdot 10^9$ bits.

The number of patterns to be stored in a single AM chip can be calculated as follows. The whole FTK system requires 1 billion patterns. Hence, 8 million patterns need to be stored on a single board (10^9 patterns / 128 boards). Finally, 128 thousand patterns per chip are required (8 million patterns per board / 64 AM chips per board).




Figure 2. Boards.

Table 1. AM chip versions.

<i>Vers.</i>	<i>Design approach</i>	<i>Technology</i>	<i>Area</i>	<i>Patterns</i>	<i>Package</i>
1	Full custom	700 nm		128	QFP
2	FPGA	350 nm		128	QFP
3	Std cells	180 nm	100 mm ²	5 k	QFP
4	Std cells + Full custom	65 nm	14 mm ²	8 k	QFP
mini-5	Std cells + Full custom +	65 nm	4 mm²	0,5 k	QFP
5	+ SERDES + IP blocks		12 mm²	3 k	BGA
6¹	Std cells + Full custom +	65 nm	150 mm²	128 k	BGA
	+ SERDES + IP blocks				

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 system [8]: as we are fitting 8 000 AM chips in 8 VME crates and 4 racks, the power should not exceed 250 W per AM board.

The Associative Memory board is a 9U VME card, connected to a Rear Transition Module (RTM), which is placed in the same slot of the VME core crate, as shown in figure 2. An ERNI 973028 ERmet ZD high-speed connector (labelled P3 in figure 2) enables the data transfer between cards through serial links at 2 Gbit/s.

2.2 The FTK AM integrated circuit

The Associative Memory integrated circuit (AM chip) is a dedicated device specifically designed to achieve maximum parallelism during operation. Each pattern has a dedicated comparator, and track searching is performed during detector readout.

The AM chip has been previously designed in several versions. Table 1 lists the main features of the various versions of the AM chip.

¹Under design (area and no. of patterns are estimated).

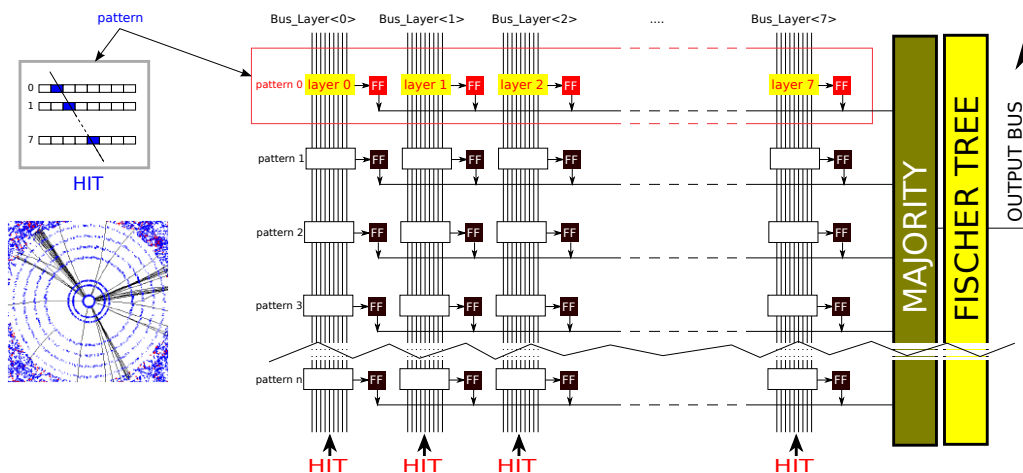


Figure 3. Scheme of an associative memory.

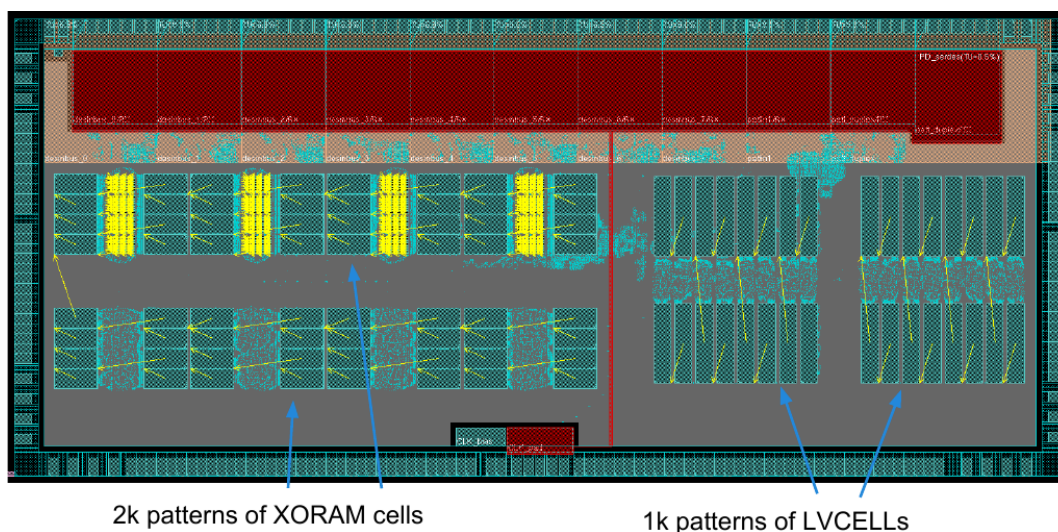


Figure 4. Floorplan of the AM05.

Figure 3 illustrates the scheme of an associative memory array. Detector layers produce “hit” signals due to colliding particles (figure 3, left). The set of hits is sent by the Data Formatter to the AM, which compares its own content with the data received. Matching results (1 or 0) are stored in Flip-Flops (FF), and partial matches are analyzed by the majority logic and compared to the desired threshold. Finally, a priority encoder reads the matched patterns in order (figure 3, right), using a modified Fischer Tree algorithm [9].

2.3 The AM05

The version 5 of the Associative Memory chip (AM05) has been designed in 65 nm CMOS technology. Figure 4 shows the floorplan of the whole chip, which occupies a total area of 12 mm². The AM chip includes ternary cells that allow the width of a pattern to vary layer by layer. This increases the effectiveness of the AM chip by the equivalent of a factor ≈ 5 in the number of patterns [10].




Figure 5. XORAM schematic.




Figure 6. Layout of the XORAM block in a 65 nm CMOS technology.

The main purpose of the AM05 is to evaluate three different options for the associative memory: (1) a new type of AM cell based on XOR logic gate (XORAM); (2) the same XORAM cell, with full-custom majority; and (3) a low-voltage (LV) NAND-NOR memory logic, based on a modified version of the scheme presented in [11].

The following subsections summarize the characteristics of the two memory cell options. Details on the design of the AM05 are presented in [12].

2.3.1 XORAM Memory Layer

The XORAM cell has been described in [13]. It is based on the XOR logic function, and it is made of a conventional 6T SRAM cell merged with a pass-transistor XOR gate. Figure 5 shows the CMOS schematic diagram, and figure 6 illustrates the layout of a 1-bit cell.

The single bit cell output (OUT) is equal to zero when the stored bit (A) matches the bit-line (BL), and is equal to one when they are different. The comparison on the 18-bit words is made simply by taking the logic NOR of the 18 AM cell output bits.

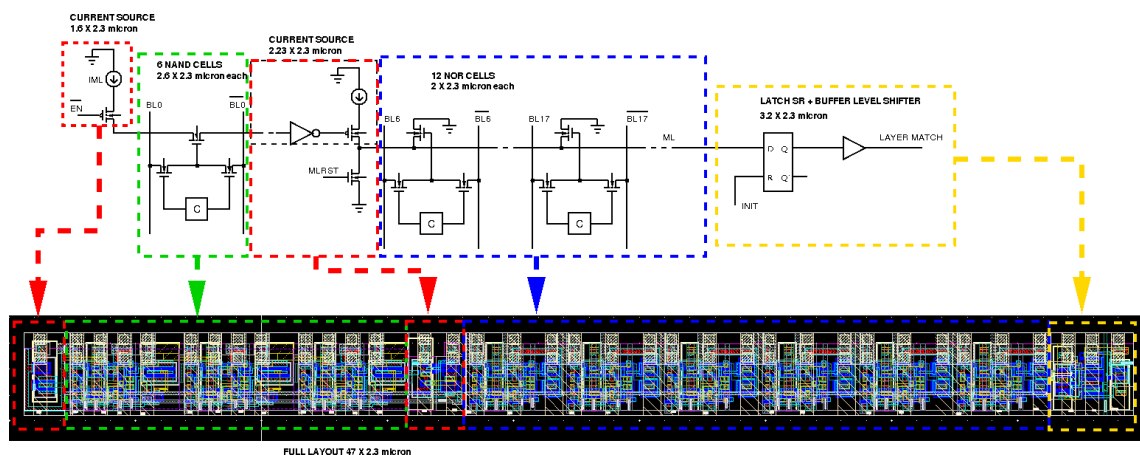


Figure 7. Schematic diagram and layout of one layer of LV cells in the AM05.

2.3.2 LV NAND-NOR Memory Layer

A new low voltage (LV) current race AM cell has been designed, suitable for 0.8 V supply (lower than 1.2 V used for standard cells). It is based on a current race and selective precharge scheme, and it contains 6 NAND type cells (with 9 transistors each) and 12 NOR type cells (with 9 transistors each). The schematic diagram and the layout of the LV NAND-NOR AM are illustrated in figure 7.

3 Serial link characterization and current consumption measurements

The mini-AM05 (“mini-5” in table 1) is a small integrated circuit prototype containing 512 patterns. It has been designed to verify the functionality of the XORAM cell, to test serial links at 2 Gbit/s, and to measure current consumption for different operations.

In the FTK system, chips will communicate through serial links at 2 Gbit/s, to avoid routing congestion at the board level and to reduce crosstalk. Serializers and deserializers (SERDES) have been included in the test chip, using IP blocks provided by Silicon Creation [14].

Figure 8 shows the test setup. The AMChip is inserted into a zero insertion force (ZIF) socket, supplied by Yamaichi and designed for high-frequency applications. The ZIF socket is mounted onto a passive printed circuit board (PCB), called “mezzanine card”, where signals coming from the AMChip through ZIF pins have been routed to a high pin count VITA 57.1 connector.

Power supply lines are routed on the mezzanine card and are connected to 4-pin connectors, for 4-wire measurements of current consumption in different parts of the chip core, I/O, and SERDES IP blocks in different configurations. Data are sent and collected by a Xilinx Virtex-6 FPGA mounted on an evaluation board supplied by HiTechGlobal. The firmware to program the FPGA hardware has been written in VHDL, and the software in Python.

3.1 Serial link measurements

Serial links are based on the LVDS electrical protocol, with a voltage swing of 400 mV and an average value $V_{\text{avg}} = 1.8$ V. Each link has a coupling resistance with a value of 100 Ω .



Figure 8. Test setup for the mini-AM05.

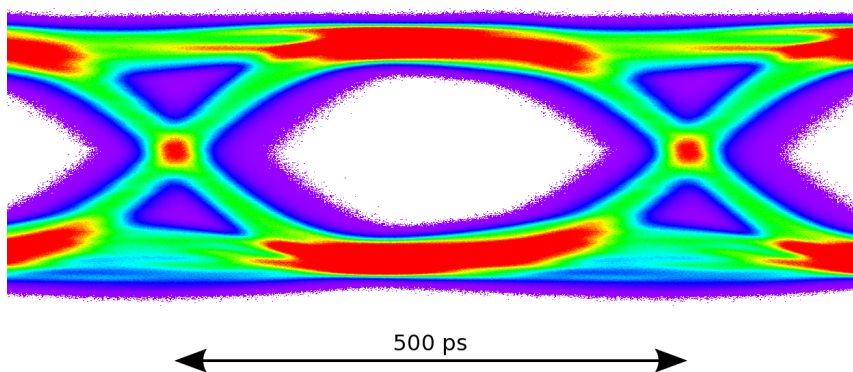


Figure 9. Eye diagram.

A pseudo-random bit sequence (PRBS) generator inside the FPGA transmits data on the serial links towards the AM chip, which has been configured in a parallel loopback mode to send data back to the FPGA. The analog waveforms of the LVDS link have been acquired by a LeCroy digitizing oscilloscope sampling at 40 Gsample/s, through a differential analog probe, and jitter analysis has been performed to characterize the quality of serial data links.

Figure 9 shows the “eye diagram” with serial data at 2 Gbit/s, and figure 10 shows the “bathtub” diagram. After 18 h of data acquisition, the deterministic jitter and the periodic jitter are 55 ps and 83 ps, respectively.

The resulting bit error ratio has been estimated as $\text{BER} \approx 10^{-21}$.




Figure 10. Bathtub diagram.

Table 2. Current consumption of the mini-AM05.

test mode	current consumption	
	1.0 V	1.2 V
baseline (all cells are disabled)	3.3 mA	4.0 mA
clock propagation inside XORAM array	0.9 mA	1.0 mA
clock propagation inside NAND-NOR array	0.9 mA	1.0 mA
matching of 64 patterns for XORAM array	2.7 mA	3.2 mA
matching of 64 patterns for NAND-NOR array	1.9 mA	2.4 mA

3.2 Current measurements

Table 2 shows the current drawn by the mini-AM05 in different operation modes.

When the input data is changing, the AM chip is active and performs the comparison between input and stored data in parallel. Dynamic power consumption due to the input data buses has been identified as the major contribution to the overall power. In particular, the XORAM array exhibits a larger current, due to the aspect ratio of the XORAM cell, which leads to a higher parasitic capacitance of the parallel input data wires.

For this reason, the AM05 (in figure 4) has been completely redesigned with different shape and arrangement of cells, aiming at lowering input wire capacitance.

4 The FTK AM board

Figure 11 illustrates the new AM board prototype, showing: (1) the input 2 Gbit/s serial links, i.e., the hit paths from the ERNI 973028 ERmet ZD connector to the “Little AM Boards” (LAMBs); (2) the output 2 Gbit/s serial links, i.e., the road paths from the LAMBs to the ERNI 973028 ERmet ZD connector; (3) the two ARTIX-7 FPGAs (one for the input and one for the output).

Each AM board will hold 4 LAMBs, each with 16 AM chips.

4.1 FTK mini-LAMB prototype and final LAMB design

A mini-LAMB prototype, shown in figure 12, was designed and manufactured to verify serial link performance. The mini-LAMB has the same size and the same connector as the LAMB, the main difference being the number of AM chips: 4 mini-AM05 are mounted onto the mini-LAMB, instead of 16.




Figure 11. New AM board prototype.




Figure 12. Mini-LAMB prototype.

The final LAMB has also been designed taking into account the features of the new AM chip BGA package.

4.2 Mini-LAMB prototype test

Figure 13 shows the test setup for the mini-LAMB: one mini-LAMB is mounted onto a dedicated mezzanine card, which is connected to the Xilinx FPGA evaluation board.

Test performed on the mini-LAMB prototypes demonstrated that serial links at 2 Gbit/s are working: JTAG commands were successfully transmitted, and PRBS tests were performed. Figure 14 shows the results after ≈ 10 h of PRBS testing; the estimated bit error ratio is $\text{BER} < 10^{-15}$.

The mini-LAMB has also been mounted on the motherboard shown in figure 11 and successfully tested also in this condition.




Figure 13. Setup for the mini-LAMB prototype tests.

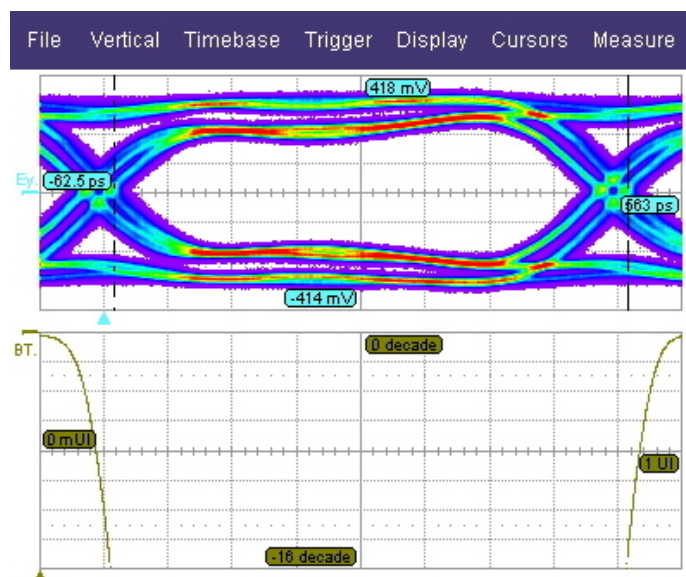


Figure 14. Mini-LAMB prototype test results: “eye” diagram (top), and “bathtub” diagram (bottom).

5 Conclusion and future work

The FTK system is currently under design. AM board and chip prototypes have been manufactured and tested. In particular, tests performed on the mini-AM05 demonstrated the correct operation of the new XORAM cell and excellent performance of serial links at 2 Gbit/s. The current consumption was measured in different modes. As a significant fraction of the power dissipation is due to the input data distribution inside the chip, board level and crate level consumption are still a concern. For this reason, the AM05 was completely redesigned at layout level, to improve power performance. AM05 prototypes are now under test.

High speed serial links at 2 Gbit/s have also been successfully tested on the mini-LAMB.

Future work will include the test of AM05, which is expected to provide information for the optimal design of the final AM06. AM05, which is pin compatible with AM06, will also be used to test the new LAMB and to integrate the AM system in the FTK demonstrator in which system tests will be carried out prior to production.

Acknowledgments

The Fast Tracker project receives support from Istituto Nazionale di Fisica Nucleare; the US National Science Foundation and Department of Energy; Grant-in-Aid for Scientific Research from the Japan Society for the Promotion of Science and MEXT, Japan; the Bundesministerium für Bildung und Forschung, FRG; the Swiss National Science Foundation; and the European community FP7 People grant FTK 324318 FP7-PEOPLE-2012-IAPP.

References

- [1] ATLAS collaboration, *The ATLAS experiment at the CERN Large Hadron Collider*, 2008 *JINST* **3** S08003.
- [2] W.H. Smith, *Triggering at LHC experiments*, *Nucl. Instrum. Meth. A* **478** (2002) 62.
- [3] M. Dell’Orso and L. Ristori, *VLSI structures for track finding*, *Nucl. Instrum. Meth. A* **278** (1989) 436.
- [4] L. Rossi, *LHC upgrade plans: Options and strategy*, in *Proceedings of the International Particle Accelerator Conference (IPAC)*, San Sebastián, Spain (2011), pg. 908.
- [5] R. Bartoldus et al., *Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System*, CERN-LHCC-2013-018, ATLAS-TDR-023 (2013).
- [6] A. Andreani et al., *The FastTracker real time processor and its impact on muon isolation, tau and b-jet online selections at ATLAS*, *IEEE Trans. Nucl. Sci.* **59** (2012) 348.
- [7] A. Andreani et al., *The AMchip04 and the processing unit prototype for the FastTracker*, 2012 *JINST* **7** C08007.
- [8] D. Calabrò et al., *The associative memory boards for the FTK processor at ATLAS*, in *Proceedings of IEEE Nuclear Science Symposium and Medical Imaging Conference*, Seoul Korea (2013).
- [9] P. Fischer, *First implementation of the MEPHISTO binary readout architecture for strip detectors*, *Nucl. Instrum. Meth. A* **461** (2001) 499.
- [10] A. Annovi et al., *A new variable-resolution associative memory for high energy physics*, in *Proceedings of IEEE International Conference on Advancements in Nuclear Instrumentation Measurement Methods and their Applications (ANIMMA)*, Ghent Belgium (2011).
- [11] K. Pagiamtzis and A. Sheikholeslami, *Content-addressable memory (CAM) circuits and architectures: A tutorial and survey*, *IEEE J. Solid-State Circ.* **41** (2006) 712.
- [12] A. Andreani et al., *Next generation associative memory devices for the FTK tracking processor of the ATLAS experiment*, in *Proceedings of IEEE Nuclear Science Symposium and Medical Imaging Conference*, Seoul Korea (2013).
- [13] L. Frontini, S. Shojaii, A. Stabile and V. Liberali, *A new XOR-based Content Addressable Memory architecture*, in *Proceedings of the International Conference on Electronics, Circuits and Systems (ICECS)*, Seville Spain (2012), pg. 701.
- [14] Silicon Creation, *Serializer - Deserializer IP*, online: <http://www.siliconcr.com/products/serdes-interfaces>.