

Next generation Associative Memory devices for the FTK tracking processor of the ATLAS experiment

Alessandro Andreani, Alberto Annovi, Roberto Beccherle, Matteo Beretta, Mauro Citterio, Francesco Crescioli, Alessandro Colombo, Paola Giannetti, Valentino Liberali, Jafar Shojaii, Alberto Stabile

Abstract—The AMchip is a VLSI device that implements the associative memory function, a special content addressable memory specifically designed for high energy physics applications and first used in the CDF experiment at Tevatron. The 4th generation of AMchip has been developed for the core pattern recognition stage of the Fast TrackKer (FTK) processor: a hardware processor for online reconstruction of particle trajectories at the ATLAS experiment at LHC. We present the architecture, design considerations, power consumption and performance measurements of the 4th generation of AMchip. We present also the design innovations toward the 5th generation and the first prototype results

Index Terms—Associative Memory, FTK, Trigger, ATLAS

I. INTRODUCTION

FTK is a dedicated processor that rapidly finds and fits tracks in the ATLAS inner detector silicon layers for every event that passes the Level-1 Trigger. It uses all silicon layers (4 pixel, including IBL, and 4 SCT stereo and axial for a total of 12 layers) over the full rapidity range covered by the barrel and the disks. It receives a parallel copy of the silicon detector data at the full speed of 100 kHz of the readout from the detector front end to the read-out subsystem following a Level-1 Trigger.

The FTK processor operates with a two step algorithm: in the first step coarse tracks are found then in the second step for each coarse track fits is performed to find high resolution tracks. The first step is very computing intensive: many hits combinations on the silicon layers must be analyzed in order to find all coarse tracks present in the event. The Associative Memory (AM) chips are hardware devices capable of solving this very demanding computing task in real time comparing all possible hits combinations with a pre-computed set of possible tracks and finding the matches.

A. The Associative Memory working principle

The FTK AM is an evolution of the concept of Content Addressable Memory (CAM). A CAM is a device that implements the inverse function of a Random Access Memory (RAM). A RAM stores data at given addresses and retrieves the corresponding data if queried with the address,

F. Crescioli is with LPNHE IN2P3 CNRS, 4 place Jussieu, 75252 Paris Cedex 05, France. e-mail: francesco.crescioli@lpnhe.in2p3.fr

A. Andreani, M. Citterio, A.Colombo, V. Liberali, J. Shojaii and A. Stabile are with INFN Milano, Italy

R. Beccherle and P. Giannetti are with INFN Pisa, Italy

A. Annovi and M. Beretta are with INFN Laboratori Nazionali di Frascati, Italy

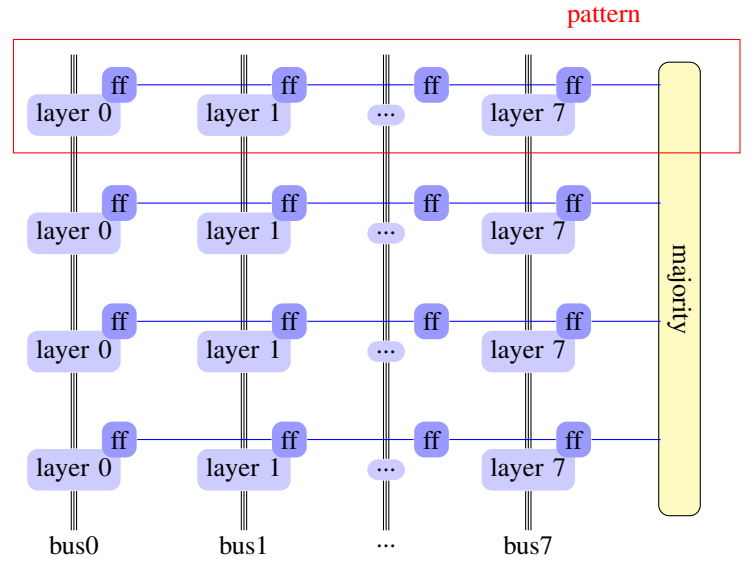


Fig. 1. Schematic internal structure of an Associative Memory

a CAM stores data at given addresses and retrieves a list of addresses matching the queried data. CAMs are commonly used in networking devices to implement routing tables or in CPU to implement translation lookaside buffers.

The FTK AM implements the CAM function, but the data is segmented into independent busses and the matching of an address is decided by a majority logic unit over the partial matches of each segment that are stored in FlipFlops. With this feature the FTK AM is able to find matches between its own content and any combination of the input data solving a complex combinatorial problem in real time.

In figure 1 is shown the internal structure of the Associative Memory. Each row represents a pre-calculated trajectory stored in the memory (pattern) in the form of one hit per silicon detector layer. The hits coming from different layers of the tracking detector (up to 8 busses in the current ASIC implementations) are processed in parallel by the AM for pattern recognition. Each layer is compared independently by a CAM cell for each pattern and for each layer. A local memory (“ff” in the figure) latches the result of the comparison. The pattern is matched if there are layer matches above a programmable threshold (majority).

II. THE ASSOCIATIVE MEMORY CHIP

The first device implementing the FTK AM function described in the previous section was developed in the 90s: a VLSI ASIC for the SVT tracking processor at the CDF experiment at Fermilab [1]. It had a capacity of 128 patterns and was designed in full custom $1.5 \mu\text{m}$ technology.

In 2006 an improved ASIC called AMchip03 was developed for the SVT upgrade [2]. AMchip03 was designed using UMC $0.18 \mu\text{m}$ Standard Cells library, it had a capacity of 5k patterns, 6 input busses and a power consumption of 1.8 W at an operating frequency of 40 MHz.

The requirements for the ATLAS FTK application are more demanding than those for SVT: a bigger silicon detector with higher granularity requires more patterns and more input busses. Higher trigger frequency requires higher operating frequency while the total power consumption must be contained. The AMchip04, the 4th generation of AM devices and the first prototype aimed at FTK application, introduced a mixed architecture: full custom blocks for the CAM cells, standard cell logic for everything else (JTAG, input / output logic, majority logic, priority encoder). The technology chosen for AMchip04 is TSMC 65 nm. The use of full custom CAM cells enabled a higher pattern density with respect to AMchip03 and also the use of advanced techniques to reduce power consumption, more than what expected from simple node scale from 180 nm to 65 nm.

Another very important feature was introduced with the AMchip04 and continued in the newer generation: ternary logic bits. Some bits in the CAM cell can store ternary values (1, 0, don't care) and they can be used to achieve a variable resolution pattern. The idea of variable resolution pattern is essential in FTK to have a high efficiency pattern bank without increasing the capacity of the AM system over the foreseen one billion patterns [3].

The AMchip04 produced in 2012 had the following key characteristics:

- 8192 patterns storage capacity
- 8 input busses, 15 bits wide (parallel)
- 3 bits with ternary logic (configurable up to 6 bits with 12 bit wide inputs)
- 100 MHz operating frequency

The AMchip05 presented in this paper will utilize the same TSMC 65 nm technology, but with further improvements in the direction of low power consumption.

A. AMchip04 CAM cell architecture

The AMchip04 core is based on CAM memory architecture whose working principle is described in [4]. The matchline, see figure 2, power dissipation, due to its continuous charge and discharge, is one of the major sources of power consumption in CAM.

The first technique, called selective precharge, performs a match operation on the first few bits of a word using a slower but more energy efficient technique (NAND bit cell) before activating the search of the remaining bits with faster but more power demanding technique (NOR bit cell) [5]. For example, in our 18 bit word memory, selective precharge

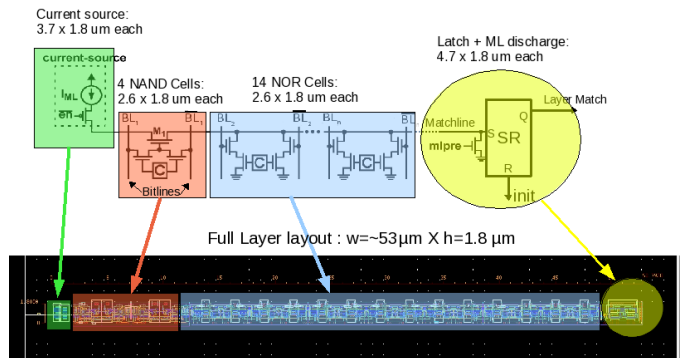


Fig. 2. AMchip04 CAM cell scheme

TABLE I
AMCHIP04 POWER CONSUMPTION

	Measured @ 100MHz	extrapolated to 128K
Baseline, leakage (mA)	7	112
clock distribution (mA)	30	480
std, not bitline propagation (mA)	6	96
bitline propagation (mA)	82	1312
AM cells (mA)	70	1120
Total Core (mA)	195	3120
Voltage (V)	1.2	1.2
Total Core (W)	0.234	3.744

initially searches only the first 4 bits and the search the remaining 14 bits only for words that matched the first 4 bits. The current race scheme precharge the matchline low, instead of high as in conventional schemes, and evaluates the matchline state by charging the matchline with a current supplied by a current source. The benefit of this schemes, apart from the power consumption reduction, is the simplicity of their circuitry.

The ternary logic feature can be implemented using two NOR cell bits: the internal 18 bit word with 3 bits configured to store ternary logic values corresponds to an 15 bit wide input bus.

B. AMchip04 power consumption

In table I are reported the power consumption measurements done on the AMchip04 prototype chip with 8K pattern at 100 MHz and the extrapolation for the final 128K pattern chip needed for FTK.

Comparing the AMchip03 power consumption of $P_{AMchip03} = 1 \mu\text{W} / \text{pattern/layer/MHz}$ at 40MHz with the AMchip04 power consumption of $P_{AMchip04} = 0.036 \mu\text{W} / \text{pattern/layer/MHz}$ at 100 MHz we can see that

$$\frac{P_{AMchip03}}{P_{AMchip04}} = 28 \quad (1)$$

The power saving techniques implemented in AMchip04 has allowed to gain about a factor 28 in power consumption with respect to the previous version. With this memory core architecture a full 128K pattern AMchip should consume 3.7 W. This value that seems to be very promising still remain




Fig. 3. AMchip05 CAM NOR cell layout

too high for our purposes, because, including FPGA and other components on the board, bring the power consumption greater than 5 KW per crate. It is needed a further power reduction in the AMchip, our goal is to reach 2 W per chip.

III. THE 5TH GENERATION AM CHIP

To further increase the power reduction we can apply several different strategies, that is:

- 1) Reduce the full custom core power supply voltage from 1.2 V to 1.0 V or 0.8 V
- 2) Reduce the CAM layer matchline capacity (length)
- 3) Reduce the CAM layer matchline voltage swing from 1.2 V to 1.0 V or 0.8 V
- 4) Reduce the bitline capacity (length)
- 5) Reduce the bitline voltage swing from 1.2 V to 0.8 V
- 6) Reduce the standard logic supply voltage from 1.2 V to 1.0 V.

What we would like to avoid is reducing the operating frequency of 100 MHz. As can be seen from this list, there are two different ways to take under control the power consumption, reduce the net capacity and their voltage swing. This is because the dynamic power consumption is dominated by the charging and discharging of line capacitance, and can be expressed by the following equation:

$$P_{Cap} = C_{line} V_{DD}^2 f \quad (2)$$

where C_{line} is the capacitance of the net, V_{DD} is the power supply voltage and f is the switching frequency of the line. Two different CAM cell architectures have been developed: LV AM and XOR+RAM.

A. Low Voltage AM cell

The Low Voltage (LV) AM cell utilizes the same architecture of AMchip04 with some improvements. Reducing nets capacitance means reducing their length and their coupling with neighboring nets in particular power supply (VDD) and ground (GND). This can be achieved by properly design the layout of the CAM cells. In figure 3 the is a layout example of the NOR CAM cell, in which the two memory cells are designed one over the other instead of side by side as in the AMchip04 memory core. In green is shown the match line path.

This cell is designed to run at low voltage, nominally 0.8 V. Lowering the power supply voltage brings speed problems, in

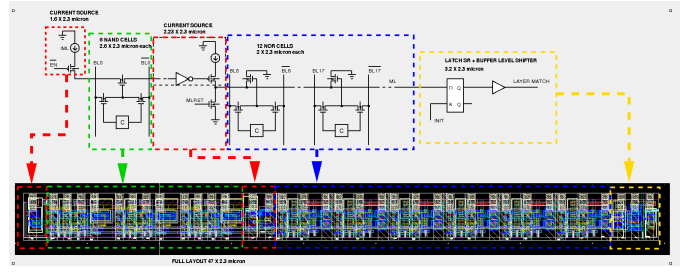


Fig. 4. AMchip05 CAM cell scheme

fact reducing the VDD increase the MOS channel resistance and the MOS speed. In order to maintain the circuit speed it is needed to use low threshold transistors instead of standard ones. Moreover, to reduce the voltage drop through the MOS channel resistance it is necessary to increase the MOS channel width and then increase the transistor area. The scheme and the layout of the new CAM memory layer, shown in figure 4, has been designed in order to satisfy all these design constraints.

B. XOR+RAM AM cell

The XOR+RAM architecture is a completely different approach to CAM cell design: the matchline is a combinatorial logic network that perform the comparison of the memory content with the data preset on the bit lines. This can be achieved combining together an SRAM memory cell and an XOR combinatorial network that perform the comparison. The result is shown in figure ??.

As can be seen from the scheme, this is a completely digital approach to the CAM memory design. With this kind of architecture, every read cycle the circuit doesn't have to charge and discharge the matchline to perform the comparison between data present on the bit lines with data stored in the memory, the XOR network is devoted to this job. Only when there is a match all the XOR gate are activated, otherwise only a small fraction of these gate is involved.

In order to reduce power consumption also the XOR+RAM has been designed to minimize the bitline capacity. This architecture is more compact in area with respect to AMchip04 and LV AM cell architecture and all bits can be paired as ternary logic bits.

C. Power estimation for LV AM and XOR+RAM cells

From simulation results it is possible to compare the architecture presented in the previous subsections.

Both architectures seems to be very promising, as reported in table II, and can achieve a 35-45% power saving in the typical case with respect to the AMchip04 reaching the desired about 2 W / chip core power consumption. Power estimation for other conditions (best speed, worst speed, worst power, ...) are in progress.

These two architectures will be tested in an upcoming prototype chip MPW submission (AMchip05 MPW) with 2K patterns with LV AM technology and 2K patterns with XOR+RAM technology. The technology for the final FTK AM

TABLE II
LV AM AND XOR+RAM POWER CONSUMPTION PREDICTION FROM
SIMULATION RESULTS FOR 128K PATTERNS CHIP

	LV AM	XOR+RAM
Baseline, leakage (mA)	112	112
clock distribution (mA)	720	576
std, not bitline propagation (mA)	144	115.2
bitline propagation (mA)	1022	1112
AM cells (mA)	524	455
Total Core (mA)	2523	2370
Voltage (V)	0.8	1.0
Total Core (W)	2.02	2.4

chip (AMchip06) will be based on the measurements results on this prototype.

IV. AMCHIP05 HIGH SPEED SERIAL I/O

The AMchip03 and AMchip04 used parallel busses for I/O. This led to extreme complexity in the design of the mezzanine boards to host the AM chips each board hosting 16 or 32 AM chips. Furthermore for AMchip05 is it foreseen to use different power domains (0.8 V or 1.0 V for the AM core, 1.0 V or 1.2 V for the standard cells, 1.2 V and 2.5 V for I/O) increasing again the routing complexity of the board.

In order to solve this board routing issue and be able to produce a reliable and relatively simple mezzanine boards we decided to switch from parallel busses to high speed serial busses. The package of the AM chip also changed from TQFP208 to BGA 23x23 in order to have many pins for the many power domains and a small number of pins for the serial I/O.

The main features required for the AM chip serial links (SERDES) are:

- data rate at least 2 Gbps to match 16 bit @ 100 MHz
- separate serializer and deserializer macro (the AM chip has many input busses but one output bus for patterns)
- 32bit input/output bus
- driver and receiver circuits compatible with LVDS standard
- 8b/10b encode/decode capabilities
- comma detection and word alignment
- BIST capabilities for fast debugging
- Low power

We have bought SERDES IP by Silicon Creations meeting all our requests.

A. AMchip05 mini@sic

To test the SERDES IP we have designed a mini@sic with a scaled down version of the AMchip05 serial architecture: 4 bus inputs + one pattern input (5 DES) and pattern output (1 SER). This mini@sic implemented a fully functional AM chip with 128 patterns made with AMchip04 technology and 128 patterns made with XOR+RAM logic. The XOR+RAM layout in this mini@sic ram is not optimized for low power consumption as it will be in the AMchip05 MPW run, but the cell logic is the same.




Fig. 5. AMchip mini@sic test setup

The mini@sic chips were received in early summer 2013 and promptly tested with a FMC socket board hosted on an FPGA evaluation board (Xilinx ML-605) as shown in figure 5.

The AMchip mini@sic was connected to the high speed GTX transceivers of the Xilinx FPGA and we were able to fully test the serial links at the desired 2 Gbps rate. We directly tested with a PRBS-7 generator the bit error rate to be less than 10^{-14} (estimation from bathtub plot is $BERR=10^{-21}$). The eye diagram of the mini@sic SER can be seen in figure 6.

The FMC board with the socket to test mini@sic was designed in order to provide and monitor the various needed power supply (AM core, SERDES core, SERDES I/O) from separated external power supplies.

We developed an FPGA firmware with the capability to load hits in the FPGA BlockRAMs and then run simple tests (finite or looped) to emulate real use cases. The structure of the tests is described in figure 7. We were able to fully verify the functionality of the new XOR+RAM logic.

We measured power consumption of 128 pattern bank with XOR+RAM logic and compared it with the simulation prediction. At 1 V and 100 MHz of operating frequency the XOR+RAM logic cells power consumption is 2.69 mA and

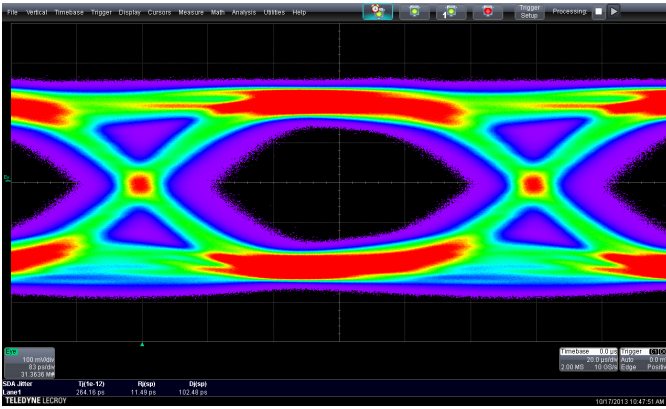


Fig. 6. AMchip mini@sic SER eye diagram at 2 Gbps

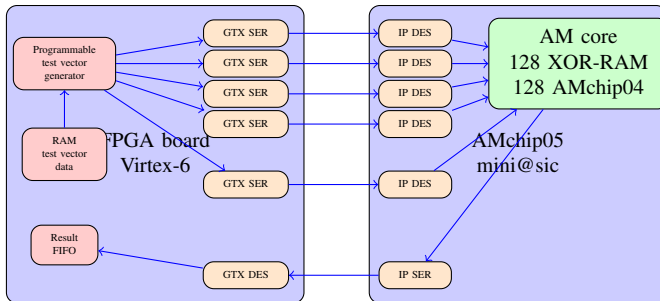


Fig. 7. AMchip mini@sic test setup scheme

the prediction from simulation is 2.79 mA.

The correspondence between the simulation and the actual measurement is very good and validates the strategy for AMchip05.

V. CONCLUSION

In this paper we described the concept of Associative Memory as it is used by the Fast Tracker processor at the ATLAS experiment. We described the 4th generation of Associative Memory. It is the first in 65 nm and includes innovative features with respect to the previous generations. We also described the path toward the final FTK chip and the planned step of development with the 5th generation of AM chips. Particular attention is devoted to explain the power reduction techniques applied and the architectural design solution to implement them into the final chip. Two different solutions are proposed both matching the power consumption and speed requirements. Moreover, because several different power supply voltages are needed a change in the associative memory package has been necessary. In order to adapt to the few available data signal pins in the new package and ease the design of the mezzanine board the 5th generation of AM chips started to use high speed serial I/O. We described the results obtained with the first test chip of the 5th generation, a mini@sic run, in particular regarding the serial link and the power consumption estimation reliability.

We foreseen a submission of a bigger chip (AMchip05) with the two low-power cells developed. We will choose the

technology for the final FTK chip in 2014 based on the AMchip05 tests results.

ACKNOWLEDGMENT

We would like to acknowledge Silicon Creation for their availability in supporting us for the design of the chip with serialized and deserialize input and output busses.

REFERENCES

- [1] S. R. Amendolia, S. Galeotti, F. Morsani, D. Passuello, L. Ristori, G. Sciacca, and N. Turini, "The AMchip: a full-custom CMOS VLSI associative memory for pattern recognition," *Nuclear Science, IEEE Transactions on*, vol. 39, no. 4, pp. 795–797, 1992.
- [2] A. Annovi, A. Bardi, M. Bitossi, S. Chiozzi, C. Damiani, M. Dell'Orso, P. Giannetti, P. Giovacchini, G. Marchiori, I. Pedron, M. Piendibene, L. Sartori, F. Schifano, F. Spinella, S. Torre, and R. Tripiccione, "A vlsi processor for fast track finding based on content addressable memories," *Nuclear Science, IEEE Transactions on*, vol. 53, no. 4, pp. 2428–2433, 2006.
- [3] A. Annovi, S. Amerio, M. Beretta, E. Bossini, F. Crescioli, M. Dell'Orso, P. Giannetti, J. Hoff, T. Liu, D. Magalotti, M. Piendibene, I. Sacco, A. Schoening, H. Soltveit, A. Stabile, R. Tripiccione, V. Liberali, R. Vitillo, and G. Volpi, "A new variable-resolution associative memory for high energy physics," in *Advancements in Nuclear Instrumentation Measurement Methods and their Applications (ANIMMA), 2011 2nd International Conference on*, 2011, pp. 1–6.
- [4] K. Pagiamtzis and A. Sheikholeslami, "Content-addressable memory (cam) circuits and architectures: a tutorial and survey," *Solid-State Circuits, IEEE Journal of*, vol. 41, no. 3, pp. 712–727, 2006.
- [5] C. Zukowski and S.-Y. Wang, "Use of selective precharge for low-power on the match lines of content-addressable memories," in *Memory Technology, Design and Testing, 1997. Proceedings., International Workshop on*, 1997, pp. 64–68.