

OPEN ACCESS

Next generation associative memory devices for the FTK tracking processor of the ATLAS experiment

To cite this article: M Beretta *et al* 2014 *JINST* **9** C03053

View the [article online](#) for updates and enhancements.

Related content

- [FTK: a Fast Track Trigger for ATLAS](#)
J Anderson, A Andreani, A Andreatza et al.
- [ATLAS FTK a - very complex - custom super computer](#)
N Kimura and ATLAS Collaboration
- [Performance and development plans for the Inner Detector trigger algorithms at ATLAS](#)
Stewart Martin-Haugh and the Atlas Collaboration

Recent citations

- [M. Ali Mirzaei et al](#)
- [A Pattern Recognition Mezzanine based on Associative Memory and FPGA technology for Level 1 Track Triggers for the HL-LHC upgrade](#)
D. Magalotti *et al*
- [Dawei Li et al](#)



IOP | ebooks™

Bringing you innovative digital publishing with leading voices to create your essential collection of books in STEM research.

Start exploring the collection - download the first chapter of every title for free.

TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS 2013,
23–27 SEPTEMBER 2013,
PERUGIA, ITALY

Next generation associative memory devices for the FTK tracking processor of the ATLAS experiment

**M. Beretta,^{a,1} A. Annovi,^a A. Andreani,^b M. Citterio,^b A. Colombo,^b V. Liberali,^b
S. Shojaii,^b A. Stabile,^b R. Beccherle,^c P. Giannetti^c and F. Crescioli^d**

^a*INFN-Laboratori Nazionali di Frascati,
Via E. fermi, 40, 00044 Frascati (Roma), Italy*

^b*INFN-Milano,
Via Celoria, 16, 20133 Milano, Italy*

^c*INFN-Pisa,
Edificio C - Polo Fibonacci Largo B. Pontecorvo, 3 - 56127 Pisa, Italy*

^d*LPNHE IN2P3 CNRS,
4 place Jussieu, 75252 Paris Cedex 05, France*

E-mail: matteo.beretta@lnf.infn.it

ABSTRACT: Higher LHC energy and luminosity increase the challenge of track reconstruction for the ATLAS trigger. To effectively handle the very high data rate, a dedicated hardware-based system has been designed. The Fast Track Trigger (FTK) will provide high quality track reconstruction over the entire detector volume to be run after the first level trigger has accepted an event. It will help to improve the efficiency and background rejection for triggers on tau leptons and b-hadrons by the second level trigger and help reduce the luminosity dependence of isolation requirements for electrons and muons. In this paper we present the status of associative memory design and its future development.

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

¹Corresponding author

Contents

1	ATLAS FTK Architecture	1
2	STD cells vs full custom design	2
3	AMchip04 core architecture	2
4	Power reduction strategies	4
4.1	Low voltage CAM cells	4
4.2	XOR + RAM architecture	5
4.3	Architectures comparison	6
5	Serialized and deserialized input and output	6
6	AMchip05	7
7	Future evolution of the associative memory chip	7
8	Conclusions	8

1 ATLAS FTK Architecture

Fast Track Trigger (FTK) is an electronics system that rapidly finds and fits tracks in the ATLAS [8] inner detector silicon layers (pixel and SCT) for every event that passes the Level-1 Trigger (figure 1). It uses all 12 silicon layers over the full rapidity range covered by the barrel and the disks. It receives a parallel copy of the pixel and silicon strip (SCT) data at the full data transfer speed from the detector front end to the read-out subsystem following a Level-1 Trigger rate, about 100KHz. The FTK algorithm consists of two sequential steps. In step 1, pattern recognition is carried out by a dedicated device called the Associative Memory (AM) [5], which finds track candidates in coarse-resolution roads using 8 of the silicon layers. When a road has hits in at least 7 silicon layers, step 2 is carried out in which the full resolution hits within the road are fit to determine the track helix parameters and a goodness of the fit. Tracks that pass all these steps are kept. The first step uses massive parallelism to carry out what is usually the

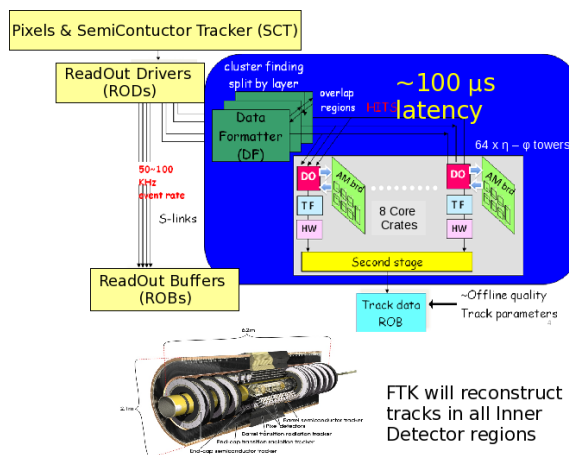


Figure 1. FTK system architecture.

most CPU-intensive aspect of tracking by processing hundreds of millions of roads nearly simultaneously as the silicon data pass through FTK. This step is performed by the associative memory chips, which contain roads consistent with particle trajectories. The AM chips compare these roads with the data coming from the ATLAS inner tracker.

2 STD cells vs full custom design

In this section we describe the project constraints and architectural solutions that have brought AMchip asic design from a purely standard cell layout to a full custom layout of the memory core. The starting point for the design of the new generation of AMchips was the AMchip03 [2]. This chip was developed in UMC 0.18 μm technology. The goal was to reach a density of 5000 patterns per chip consuming 1W at a 40MHz clocking frequency. This chip was designed using a full standard cell architecture both for the memory core and control logic. Amchip03 was successfully used in CDF experiment at Fermilab. For the FTK upgrade of the ATLAS experiment at CERN we had to change the approach of the memory core design. In this application the AMchip has to process events coming from the inner tracker (pixel and SCT detectors) at the Level-1 trigger rate, about 100KHz. The new design constraints were

- Number of stored patterns: 128000
- Number of inputs: 8 parallel buses 18 bits each + control lines
- Ternary memory cells (allows storing 0, 1, don't care)
- Frequency: 100MHz
- Power consumption: $< 2\text{W}$

To meet all these constraints it was not possible to design the new AM chip using only standard cells, because of the area and power limitations. Instead we decided to use a mixed approach in which the control logic is designed using standard cells, while the memory core is a full custom design. This has the advantage that a full custom design usually occupies less area than a standard cell design [8], this is mainly due to the fact that std cells have fixed size and cannot be scaled down to reduce their dimensions. Moreover, we could implement power saving techniques that are impossible to realize with standard cells. AMchip04 was designed with this philosophy in order to contain the power consumption and increase the number of patterns that could be stored in the memory.

3 AMchip04 core architecture

The AMchip04 core, designed in TSMC 65nm LP process, is based on CAM memory architecture whose working principle is described in [3]. The power dissipation of the matchline (shown in figure 2) is a major source of power consumption, since it is charged and discharged during every clock cycle. To reduce matchline activity, we chose to implement two power reduction techniques: selective precharging, and current race scheme.

The former performs a match operation on the first few bits of a word before activating the search of the remaining bits [4]. For example, in our 18-bit word memory, selective precharge

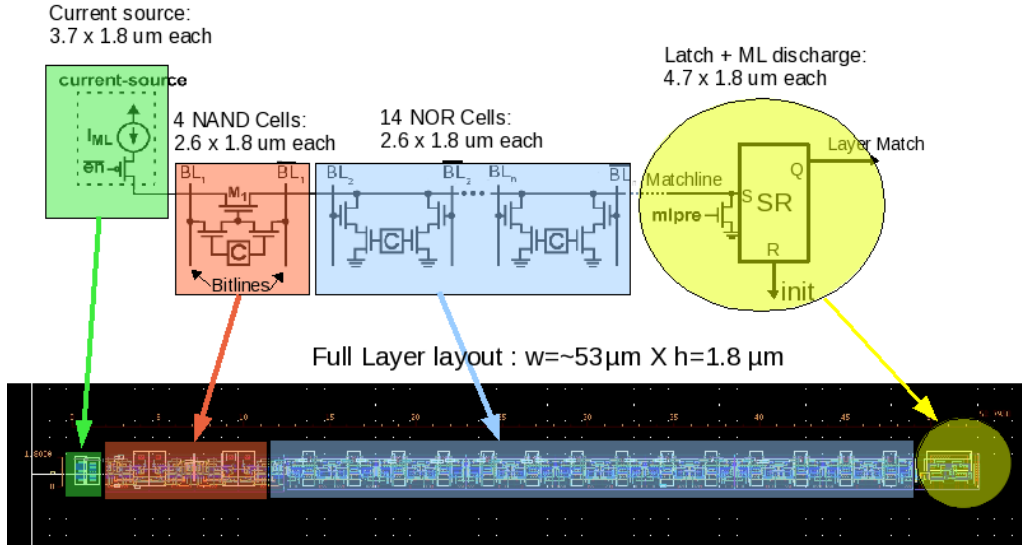


Figure 2. AMchip04 layer architecture.

Table 1. AMchip04 power consumption measurement and estimation for the 128K pattern chip.

	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 then searches the remaining 14 bits only for words that matched the first 4 bits.

The current race scheme precharges the matchline low, instead of high as in conventional schemes, and evaluates the matchline state by charging the matchline with a current I_{ML} supplied by a current source. The benefit of this scheme, which is partly responsible for the power reduction, is the simplicity of the circuitry that is composed only by two types of memory cells (NAND and NOR type) a current generator and an output SR latch. Table 1 shows the power consumption measurements done on the AMchip04 prototype chip with 8K patterns at 100MHz and the extrapolation to the final 128K patterns chip.

Comparing the AMchip03 power consumption of $P_{AMchip03} = 1 \mu\text{W}/\text{pattern}/\text{layer}/\text{MHz}$ [2] with the AMchip04 power consumption of $P_{AMchip04} = 0.036 \mu\text{W}/\text{pattern}/\text{layer}/\text{MHz}$ we can see

that

$$\frac{P_{AMchip03}}{P_{AMchip04}} = 27.7 \quad (3.1)$$

The power saving techniques implemented in AMchip04 provided about a factor of 28 reduction in power consumption with respect to the previous version. With this memory core architecture a full 128K pattern AMchip should consume 3.7W. This value however is still too high for our purposes, because, including FPGAs and the other components on the board, the power consumption would be greater than 5KW per crate. Further power reduction is needed in the AMchip and our goal is to reach less than 2W per chip.

4 Power reduction strategies

To further reduce the power we could apply several different strategies:

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

What we wanted to avoid is reducing the clock frequency below 100MHz. As can be seen from this list, there are two different ways to 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, which can be expressed by the following equation:

$$P_{Cap} = C_{line} V_{DD}^2 f \quad (4.1)$$

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.

4.1 Low voltage CAM cells

Reducing the capacitance of the nets means reducing their length and their coupling with neighboring nets, in particular the power supply (VDD) and ground (GND). This can be achieved by properly designing the CAM cell layout. figure 3 shows a layout example of the NOR CAM cell, in which the two memory cells are designed one over the other instead side by side as in the AMchip04 memory core. The match line path is shown in white.

The other way to reduce power consumption is to decrease the power supply voltage. However this raises speed problems, since reducing VDD increase the MOS channel resistance and reduces the MOS speed. In order to maintain the




Figure 3. Nor type cell layout for AMchip05.

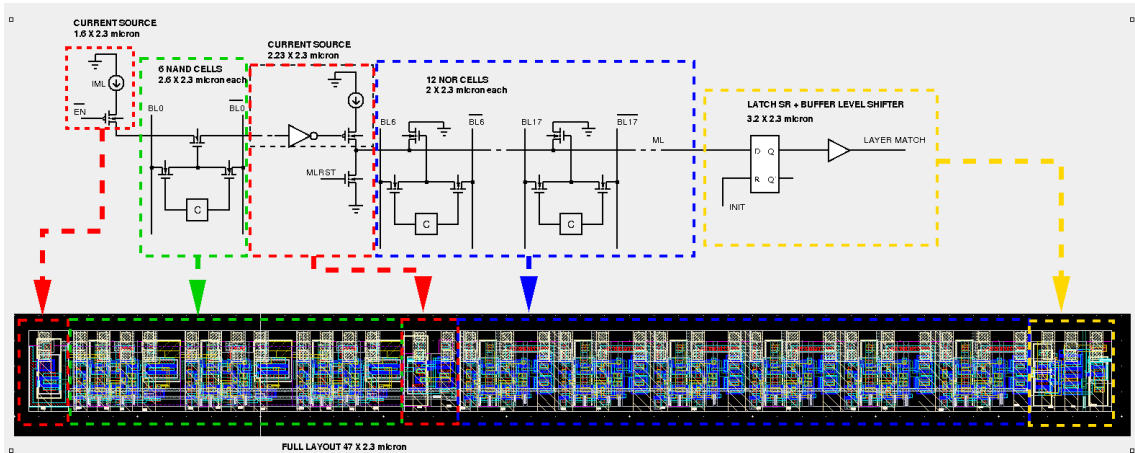


Figure 4. AMchip05 CAM layer schematic and layout.

circuit speed it is necessary 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. Taking into account all of these design considerations, we arrived at the design of figure 4.

As can be seen in the figure, in order to maintain the circuit speed an additional current generator is inserted between the NAND and NOR memory cells. Previous versions of the memory layer (figure 2) had only one current generator at the input of the NAND type memory cells. Due to all these improvements the area is increased by about 12 %.

4.2 XOR + RAM architecture

Another approach to reduce power consumption is to substitute the matchline with a combinatorial logic network that performs the comparison of the memory content with the data present on the bit lines. This can be achieved by combining an SRAM memory cell and an XOR combinatorial network that performs the comparison [7]. The result is shown in figure 5. As can be seen from the schematic, this is a completely digital approach to the CAM memory design. With this architecture, the circuit does not have to charge and discharge the matchline every read cycle to perform the comparison between the data on the bitlines and the data stored in memory. The XOR network is devoted to this job. Only when there is a match is the XOR gate activated, otherwise only a small fraction of these gates are on.

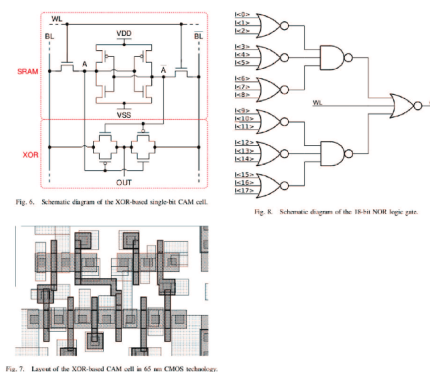


Figure 5. XOR+RAM cell scheme and layout.

Table 2. Power consumption comparison for the XOR+RAM and LV CAM cell architectures. The values reported in the table are extrapolated to a 128K patterns chip.

	XOR + RAM	LV CAM cells
Baseline, leakage (mA)	112	112
clock distribution (mA)	720	720
std, not bitline propagation (mA)	144	144
bitline propagation (mA)	1278	1022
AM cells (mA)	377	524
Total Core (mA)	2631	2523
Voltage (V)	0.8	0.8
Total Core (W)	2.1	2.02

4.3 Architectures comparison

Simulation results for the architectures presented in the previous subsections are shown in table 2. Both architectures are very promising, showing a 40% power saving with respect to the AMchip04. The XOR+RAM architecture and the AMchip04 CAM layer architecture have been implemented in TSMC 65nm technology and are currently under test.

5 Serialized and deserialized input and output

In this section we explain why we chose to put serialized and deserialized (SERDES) input and output instead of parallel buses as in AMchip04, and we present the SERDES characteristics.

The main reason for the change is that the reduction in the core VDD from the standard 1.2V to 0.8V meant several different power domains within the chip. This implies many more power pins and a package change from the TQFP208 in the AMchip04 to FBGA23x23 for AMchip05. Most of these pins are devoted to the different power supply voltages and ground, as shown in figure 6. In the few remaining pins we have to accommodate all the input and output buses, the clock and the control pins. There are not enough pins to accommodate parallel buses, so we have to use high speed serialized input and output. The main requirements for the SERDES are:



Figure 6. Sketch of the AMchip05 BGA package. The green rectangles are signal pins while all others are grounds and power supplies.

- data rate of at least 2Gbps
- separate serializer and deserializer macros
- 32bit input/output buses

- driver and receiver circuits compatible with standard LVDS
- 8b/10b encode/decode capability
- comma detection and word alignment
- BIST capability for fast debugging
- low power

We bought a SERDES core IP from Silicon Creations. To test it we designed a miniasic with 5 DES, 1 SER, their control logic and our AMchip04 memory core plus XOR+RAM with only a few banks. The miniasic is fully working and its characterization is in progress.

6 AMchip05

Starting from the good results obtained in the miniasic test, we began to design the AMchip05 in TSMC 65nm technology. In this chip there will be 8 hit buses, 2 pattern-in and 1 pattern-out buses, one input 100MHz LVDS clock plus single-ended control signals: JTAG Init, Dtest, Holds.

As stated in the previous section all the input and output buses are serialized and de-serialized at 2Gbs. Moreover due to the high number of power supply regions we have to pay particular attention to the floor plan of the chip. Figure 7 shows the floor plan for the AMchip05 in which the various blocks in the chip are seen. In particular the high frequency input and output buses are aligned on the top, while the bottom is devoted to the various memory core architectures that we want to test.

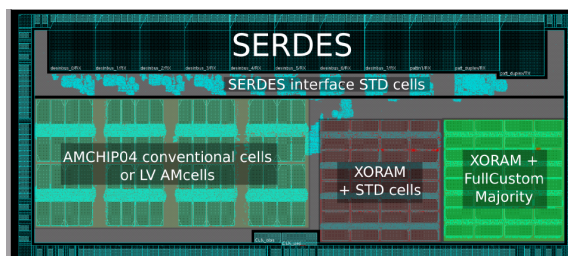


Figure 7. Floor plan of the AMchip05.

7 Future evolution of the associative memory chip

The future evolution of the associative memory chip will be to increase the pattern density while maintaining an acceptable power consumption. The first step is a technique we call 2.5D. AMchip04 has been designed to be horizontally symmetric. In/out buses for the pattern output pipeline can change direction. Moreover buses are swapped internally to maintain consistency. In this way the symmetry of the chip helps in designing and routing mezzanines for 2D chips, but also enables vertical stacking, that is the 2.5D chip architecture.

The internal design of associative memory makes it a good candidate for full 3D implementation to increase further the pattern density and decrease the footprint. Stacking of dies allows the matchline to be shortened, thus increasing speed and decreasing capacitance and power consumption [5]. The bit line will also be shorter, contributing to a further power reduction.

8 Conclusions

This paper has shown the recent development of the associative memory chip, describing in particular the power reduction techniques and architectural solutions implemented in the final chip. Two different solutions are proposed for achieving the power consumption and speed requirements. Moreover, because several different power supply voltages are needed, a change in the associative memory package was necessary. We also changed from parallel input and output buses to high speed serial differential lines. In the future, we are going towards 2.5D, i.e. stacking of several dies. Then we can start to design a real 3D chip which has several advantages in terms of power consumption, memory capacity and speed.

Acknowledgments

We would like to acknowledge Silicon Creation for their SERDES design. Moreover we would like to acknowledge the technicians who worked with us in designing, producing and testing the mezzanine boards and the chip prototypes.

References

- [8] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, 2006 *JINST* **3** S08003.
- [2] A. Annovi et al., *A VLSI processor for fast track finding based on content addressable memories*, *IEEE Trans. Nucl. Sci.* **1** (2005) 259.
- [3] K. Pagiamtzis et al., *Content-Addressable Memory (CAM) Circuits and Architectures: A Tutorial and Survey* *IEEE J. Solid-State Circ.* **41** (1234) 712.
- [4] C.A. Zukowski and S.Y. Wang, *Use of selective precharge for low power content-addressable memories*, *Proc. IEEE Int. Symp. Circuits Syst.* **3** (1234) 1788.
- [5] E.C. Oh and P.D. Franzon, *Design Considerations and Benefits of Three-Dimensional Ternary Content Addressable Memory*, *IEEE Custom Integr. Circ. Conf.* (2007) 591.
- [6] M. Dell'Orso and L. Ristori, *VLSI structures for track finding*, *Nucl. Instrum. Meth. A* **278** (1989) 436.
- [7] L. Frontini, S. Shojaii, A. Stabile and V. Liberali, *A new XOR-based Content Addressable Memory architecture*, *Proc. Int. Conf. Electron. Circ. Sys.*, Seville, Spain (2012) 701.
- [8] H. Eriksson, P. Larsson-Edefors, T. Henriksson and C. Svensson, *Full-custom vs. standard-cell design flow - an adder case study*, *Proc. Design Automation Conference Asia and South Pacific*, (2003) 507.