Associative Memory Design for the FastTrack Processor (FTK) at ATLAS

A. Annovi, M. Beretta, E. Bossini, F. Crescioli, M. Dell'Orso, P. Giannetti, M. Piendibene, I. Sacco, L. Sartori, R. Tripiccione

***Abstract*– We propose a new generation of VLSI processor for pattern recognition based on Associative Memory architecture, optimized for on-line track finding in high-energy physics experiments. We describe the architecture, the technology studies and the prototype design of a new R&D Associative Memory project: it maximizes the pattern density on ASICs, minimizes the power consumption and improves the functionality for the Fast Tracker (FTK) proposed to upgrade the ATLAS trigger at LHC. Finally we will focus on possible future applications inside and outside High Physics Energy (HEP).**

I. Introduction

THE Track reconstruction in high-energy physics experiments requires large online computing power. The Fast Tracker for ATLAS triggers [1] is an evolution of the Silicon Vertex Tracker (SVT) in CDF [2], [3] the only state of the art online processor that tackles and solves the full track reconstruction problem at a hadron collider.

The SVT track fitting system approaches the offline tracking precision with a processing time of the order of tens of microseconds, compatible with 30 kHz input event rates. This task can be performed with negligible time delay by a content addressable memory (Associative Memory, AM), i.e. a device that compares in parallel the event hits with all the stored pre-calculated low resolution track candidates (patterns) and return the addresses of the matching locations. A second processor receives the matching patterns and their related full-resolution hits to perform the final track fitting (Track Fitter, TF).

A critical figure of merit for the AM-based track reconstruction system is the number of patterns that can be stored in the bank. For the SVT upgrade [2], [4] we developed a version of the AM chip (AMchip03 processor) [5] using a 180 nm CMOS technology and strictly standard-cell based VLSI design approach. The AM chip upgrade increased the number of patterns stored in chip from 128 to 5×10^3 and it could work at 50 MHz frequency.

A. Annovi, M. Beretta, and E. Bossini are with Istituto Nazionale di Fisica Nucleare, Laboratori Nazionali di Frascati, via E. Fermi 40, 00044 Frascati Italy (telephone: +39 06 94031).

E. Bossini, F. Crescioli, M. Dell'Orso, P. Giannetti, and M. Piendibene are with Istituto Nazionale di Fisica Nucleare, Sezione di Pisa and University of Pisa, Largo B. Pontecorvo, 3 - 56127 Pisa (telephone: +39 - 050 2214 000)

I. Sacco is with Scuola Superiore Sant'Anna, Piazza Martiri della Libertà, 33 - 56127 Pisa Italia (telephone: +39 050 88.31.11)

L. Sartori, and R. Tripiccione are with Istituto Nazionale di Fisica Nucleare, Sezione di Ferrara and University of Ferrara, Via Saragat, 1. 44100 Ferrara. Italy. (telephone: +39 0532 974280)

ATL-DAQ-PROC-2010-013
21 June 2010

The FTK processor proposed for the ATLAS experiment is much more ambitious than SVT. In fact a very high efficiency and high quality track reconstruction, already shown possible by SVT, must be achieved in a much more complex detector. Moreover, the higher luminosity (10^{34} cm^{-2} s^{-1}) will increase the complexity of events. As a consequence a very large bank is necessary: candidate tracks have to cover with more than 95% efficiency the whole tracking detector ($|\eta|<2.5$), with high efficiency on transverse momentum down to 1 GeV, and the pattern recognition has to be extended to 11 silicon detector layers (3 pixel layers and 8 SCT layers) with a reasonable resolution to reduce drastically the number of found fakes and the track fitting processing time.

II. New Associative Memory

A full-custom cell is the most important goal of the R&D devoted to a new ASIC associative memory device. The new chip, with extremely high pattern density, will need also to be enriched of new functional elements and will need to be faster (at least a factor 2) than the previous version. A 65 nm technology is the best choice today to design a chip characterized by the speed, density and flexibility required by the results of the FTK performance studies. We start with 90 nm technology, available as mini ASIC already in 2009, to test early the full custom cell advantages, but probably we will move the project to 65 nm before producing the final chip.

The new full custom cell includes all the hardware necessary for the elementary functions of a single pattern layer: latches, comparators, JTAG logic for storage element read, write and debugging. In the previous AM chip these functions were implemented putting together a set of standard cells, unavoidably more expensive in terms of silicon area.

The gain we can get for the layer size is assessed to be roughly 2. Since the bank patterns occupy 80% of the full AM chip, it is reasonable to think that the new layout will produce at least a factor 2 more patterns per chip. This factor combined with the gain due to the technology scaling from 180 nm to 90 nm produces an estimate global increment of a factor 8 for the number of patterns. We can also expand the used silicon area since the package can house a 1,5×1,5 cm^2 chip, while the CDF chip is just 1cm^2. Considering that for ATLAS we will need 8 layer patterns, while in CDF we used 6 layer patterns, the final gain factor will be 12. We expect to produce a chip with a final number of 60×10^3 pattern/chip. Each pattern is provided of the logic necessary to perform pattern recognition using up to 8-layers.

The full custom cell offers also the possibility to implement important new strategies to reduce the power consumption of the chip. This is very important issue because the pattern density growth is limited by the consumption. The clock cycle is limited to 10 ns by the board complexity: each hit found in the detector has to be distributed to 128 AM chips, with a very high fan-out. However inside the 90 or 65 nm chip the 10 ns clock cycle is very conservative and this time can be exploited to do the pattern comparison with the event into 2 steps: first we compare the 4 least significant bits of each layer word, then only if they match we can continue and compare the rest of the pattern bits. This "pre-match" technique can save up to 80% power consumption.

III. The Mini-Asic Project

The features of the new AMchip04 and in particular the performances of the full custom cell must be tested in order to validate our expectations and to gain experience for further improvement of the Associative Memory technology. For these purpose we are design a mini ASIC prototype of the new AM chip with 90nm technology. The main goal is to verify that the full custom associative memory cell works properly and to verify the expected gain in terms of pattern density and power saving.

The mini ASIC area is 3.6*1.8 mm^2, which is enough to fit a pattern block of 1024 patterns (green area) with its majority logic (blue). Fig. 1 compares the AMchip04 mini-ASIC with the full AMchip03 where the pattern blocks are shown in black.

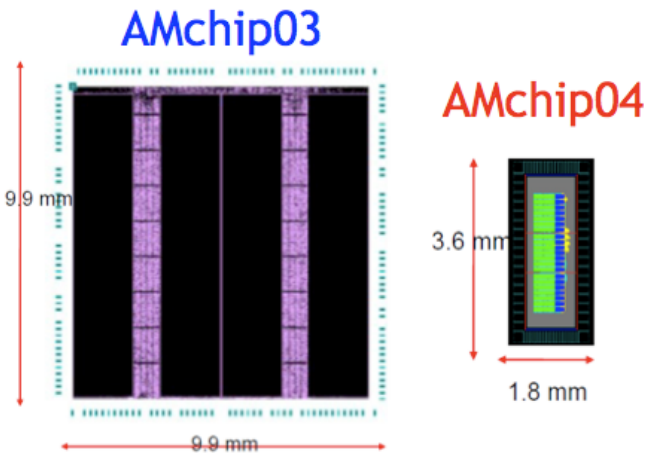

Fig. 1. Comparison of full AMchip03 with AMchip04 mini-ASIC prototype.

The full custom design is used only for the green pattern block that occupies a large fraction of the device area. Each pattern is made of 8 layers cells. Each layer cell compares 15 input bits with a stored data word. Fig. 2 shows the single bit NAND comparison cell.

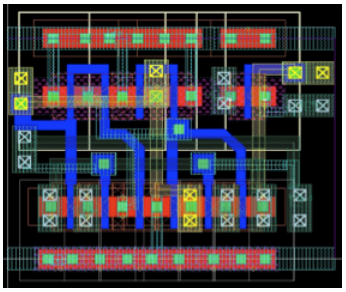

Fig. 2. Layout of a NAND CAM cell.

The 15 bits layer implements selective pre-charge [6]. The cell is made of 4 NAND cells and 11 NOR cells as shown in Fig. 3. The match line that runs through the 15 bits is pre-charged only if the 4 least significant bits match. This means a large power consumption reduction due to charge/discharge of the match line. The full custom layer cell occupies half the area of a standard cell equivalent with the same 90nm technology and save u to 80% power consumption.

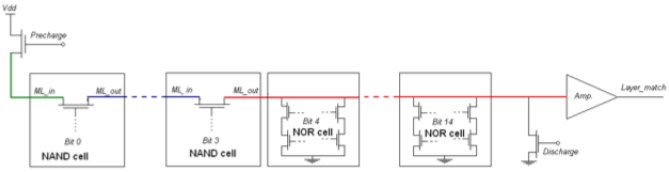

Fig. 3. Layout of a layer cell with 15 bits. The 4 least significant bits (left most) are NAND cell, while the most significant bits are NOR cells.

The remaining logic of the AMchip04 mini-ASIC is implemented with standard cells and is very similar to the AMchip03 logic. The main change is that the new chip will perform pattern recognition with 8 layers instead of 6 or 12.

Table I compares the performance of AMchip03 with the expected performance of a full AMchip04 extrapolated from the current mini-ASIC design.

TABLE I. AMCHIP PARAMETERS

	AMchip03	AMchip04	Effect
Technology	180nm	90nm	x4 pattern density
Clock freq.	50MHz	100MHz	faster, higher power cons.
Die size	10x10mm^2	15x15mm^2	x2 patterns
Core voltage	1.8V	1V	lower power consumption
Selec. Prech.	No	Yes	~80% power saving
Full custom	No	Yes	x2 pattern density
Layers	6 (or 12)	8	
Patters/chip	5k	60k	

IV. Future Improvements

The Associative Memory has a long history with the first design started in the '80s. There is room for improvement and we have ideas to push the technology beyond the performances reported in table I. Among these ideas, we are considering increasing the available area using 3D technology to stack multiple dies in the same package. We plan to switch to the 65nm technology as soon as it becomes affordable. The full custom cell layout can be further improved to optimize area and power consumption. Finally we can add to a few bits (e.g. 1 or 2 per layer) the "don't care" option. Using the "don't

care" as in ternary CAMs, we can allow for variable size patterns. When a bit is "don't care" the effective pattern size for that bit is doubled because it will match two numbers. This will allow a smarter and more efficient use of the AM patterns.

V. CONCLUSIONS

We are designing a new Associative Memory device (AMchip04) that exploits full custom CAM cells along with 90nm technology to increase the patterns available per chip by a factor 12 with respect to the AMchip03. At the same time we use low-power techniques to achieve these goals while keeping the power consumption below 2W per chip. The AMchip04 design is well underway and a mini-ASIC prototype will be produced soon. This is the first prototype of the AMchip being developed for the Fast TracKer processor.

REFERENCES

[1] A. Annovi et al., "Hadron Collider Triggers with High-Quality Tracking at Very High Event Rates", *IEEE Trans. Nucl. Sci.*, vol. 51, pp 391, 2004.

[2] J. Adelman et al., "The Silicon Vertex Trigger upgrade at CDF", *Nucl. Intsr. and Meth. in Physics Research A*, vol. 572, Issue 1, pp 361-364, March 2007.

[3] J. Adelman et al., "Real time secondary vertexing at CDF", *Nucl. Intsr. and Meth. in Physics Research A*, vol. 569, pp 111–114, 2006.

[4] J. Adelman et al., "On-line tracking processors at hadron colliders: the SVT experience at CDF II and beyond", *Nucl. Instr. and Meth. in Physics Research A,* vol. 581, pp 473–475, 2007.

[5] A. Annovi at al., "A VLSI Processor for Fast Track Finding Based on Content Addressable Memories", *IEEE Trans. Nucl. Sci.,* vol. 53, pp *24*28, 2006.

[6] K Pagiamtzis and A Sheikholeslami, "Content-Addressable Memory (CAM) Circuits and Architectures: A Tutorial and Survey", *IEEE journal of SSC*, vol. 41, n. 3, pp 2006.