


AM06: the Associative Memory chip for the Fast Tracker in the upgraded ATLAS detector

Valentino Liberali on behalf of the AM06 Design & Test Team



V. Liberali — INFN Milano and Department of Physics, Università degli Studi di Milano
Via Celoria, 16 — 20133 Milano — Italy
valentino.liberali@mi.infn.it


TWEPP 2016 — Karlsruhe, Germany — Sept. 2016

- 1 Introduction: the FTK System
- 2 The AM06 Chip
- 3 The AM06 Test Setup
- 4 Conclusion

Motivation

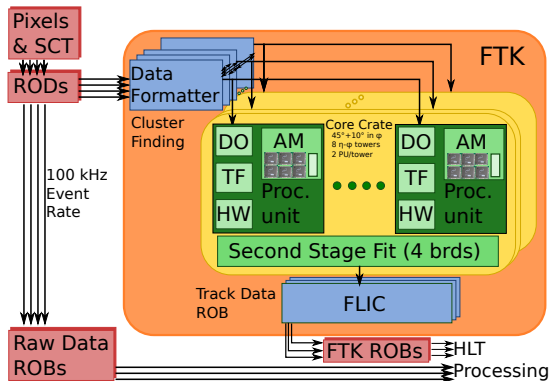
The High-Luminosity LHC will reach a luminosity up to $3 \cdot 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ and experiments will produce a huge amount of data

→ A **tight selection of events** will be required to choose data to be transferred to the mass storage system



See also the previous talk from Nikolina Ilic: *Installation, Commissioning, and Running of the ATLAS Fast Tracker Hardware System*

FTK Architecture




The FTK system is made of:

- 32 Data Formatters (DF)
- 128 Processing Units
- 32 Second Stage Fitting Boards
- Interface towards Level 2 Trigger

Each processing unit is composed of:

- an **AM Board** for pattern matching
- a rear card (AUX Board) with
 - Data Organizer (DO)
 - Track Fitter (TF)
 - Hit Warrior (HW)

The AM System



The **Associative Memory (AM)** system is the core of the FTK:

- it stores 1 billion (10^9) patterns for pattern recognition
- it performs pattern matching using the hit information of the ATLAS silicon tracker, and finds track candidates at low resolution

Number of stored patterns:

- 8 M patterns per board, 128 k patterns per chip ($128 \text{ boards} \times 64 \text{ chips}$)
- 1 pattern = 8 layers, each of them with 18-bit coordinates

Major concerns:

- Large silicon area required to store patterns
- 2 Gbit/s serial links needed to reduce I/O signal congestion at board level
- The rack cooling system limits the maximum power to **250 W per AM board**

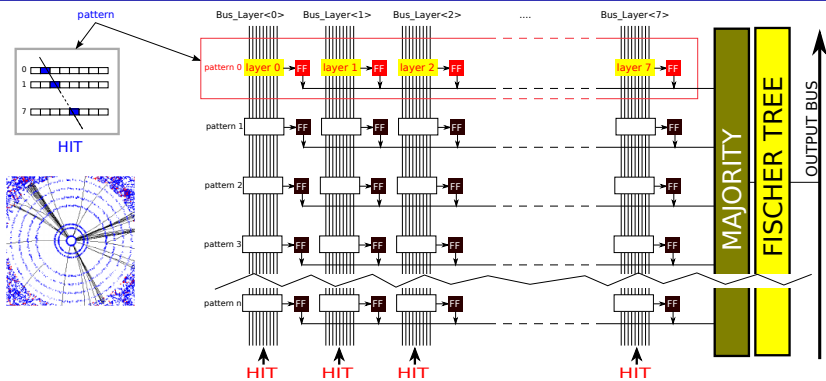
A. Andreani et al., "The AMchip04 and the processing unit prototype for the FastTracker," *IOP J. Instr.* **7** (2012) C08007

The Associative Memory integrated circuit (AMChip)

Vers.	Design	Tech.	Area	Patterns	Package
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 mm ²	0,5 k	QFP
5	+ IP blocks	65 nm	12 mm ²	3 k	BGA
6	Std cells + Full custom + IP blocks	65 nm	168 mm²	128 k	BGA
7	Std cells + Full custom	28 nm	10 mm ²	16 k	BGA, SiP

blue: The AM07 is under design.

AM working principle



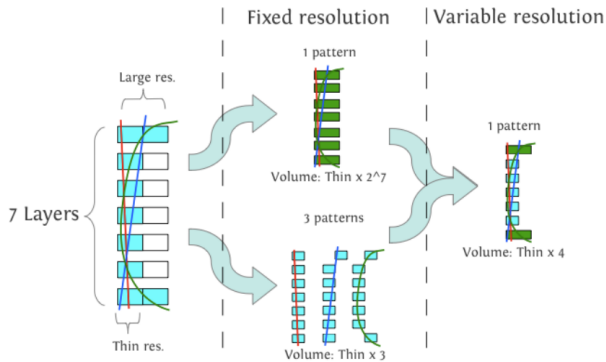
- 18-bit CAM cell for each bus and for each pattern
- The CAM cell compares its own content with the hits received
- Matching result (1 or 0) is stored into a Flip-flop (FF)
- The “majority” logic compares the number of matched layers with the desired threshold (ALL layers, (ALL - 1) layers, (ALL - 2) layers)
- A “priority encoder” reads the matched addresses, orders them, and sends them to the output

Variable resolution AM

Ternary logic (1, 0 and “don’t care”) can improve performance

A “variable resolution pattern” has:

- wider areas in extreme layers
- smaller areas in central layers



Variable resolution allows to save memory while reducing the number of fake trajectories

Pixels:

0	1	2	3	4	5	6	7
8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23
24	25	26	27	28	29	30	31

Using binary format

“01010” selects bin 10

“0001x” selects bins 2 or 3



“1x000” selects bins 16 or 24

“0x11x” selects bins 6,7,14, or 15

“111xx” selects bins 28 to 31

The XORAM cell

The new CAM cell, called **XORAM** (= XOR + RAM), is based on the XOR function and made of a 6T SRAM cell merged with an XOR gate.



Schematic and layout of the XORAM cell in a 65 nm CMOS technology


The XORAM cell has a low power consumption: ≈ 1 fJ/bit per comparison

The NOR18 cell

The memory array is organised in “words”.

Each “word” has **18 bit** and identifies a group of neighbouring pixels (“superstrip”).

18 bit = 15 bit of data + 3 optional ‘don’t-care’ bits for variable resolution




An 18-input NOR cell provides a ‘1’ at the output (O) when all the input bits (I<17>, ..., I<0>) match.



Layout of the 18-bit CAM cell, made of 18 XORAM cells and a 18-bit NOR gate

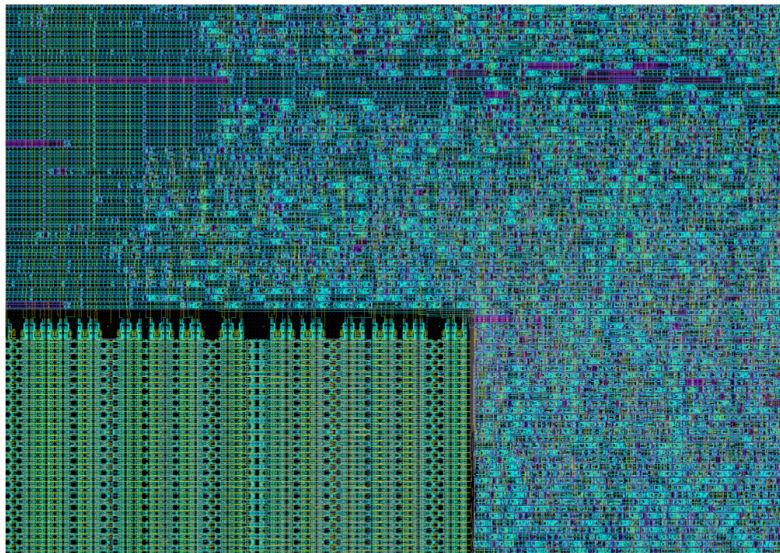
AM06 floorplan



Silicon area: **168 mm²**
Complexity: **421 million transistors**
Comparison rate: **1.88 Pbit/s**

I/O: **2 Gbit/s serial links (LVDS)**
Internal clock: **100 MHz**
Supply voltage: $V_{DD} = (1.0 \text{ to } 1.2) \text{ V}$

AM06 layout detail



memory array (full custom design) ← || → std cell logic (VHDL)

Package and current consumption



The AM06 has been packaged into an **HFC-BGA**: High-performance Flip Chip Ball Grid Array (High-performance = thermally enhanced)


- 529 balls on the external side of the substrate (23 × 23 array)
- 1178 bumps from the substrate to the chip

Most of the balls and bumps are used for power and ground interconnections
→ **reduction of parasitic inductances and resistances** in series to ground and power supply

The **AM06 current consumption** is a critical issue: **up to 2.5 A per chip** in normal operation mode (comparison)


→ a special effort was spent to carefully design the boards

AM06 test board



Test setup:

- an FPGA evaluation board from HiTechGlobal –the **red board**–
- a custom board equipped with a ZIF socket (BGA Ironwood Electronic Clamshell socket) –the **green board**–
- a box with only the ZIF socket accessible to the operator, to prevent accidental changing the test board configuration switches




A computer-based test procedure performs the test of one chip in about 2 min


The industrial test is performed by a company (Microtest srl, Altopascio, Italy)

Test board details: capacitors

The test board contains many capacitors, in order to limit the voltage drop when the AM06 switches to the 'compare' operation mode.




Electrolytic capacitors on the top side of the board



Ceramic capacitors with different capacitance values and with low ESR-ESL on the bottom side of the board (in the picture, the ZIF socket was removed)

Measurements: Eye diagram

Serial data link operation:




Measured eye diagram with LVDS serial data at 2 Gbit/s

- Jitter is ≈ 25 ps

Measurements: VDD ripple

Transition between 'write' mode and 'compare' mode:

- the AM core average current rises from 0.1 A to 2.2 A in about 0.1 ns
- current peaks are synchronous with the 100 MHz clock rising edge



The ripple of the V_{DD} core voltage is about ± 50 mV (measured on the test board, close to the socket).

To reduce problems, we operate the AM06 chip with $V_{DD} = 1.1$ V, instead of 1 V.

Conclusion

- AM06 successfully designed and fabricated
- The 'pilot run' consisted of 9 wafers in 'split lots':
 - 3 wafers with 'typical' transistors
 - 3 wafers with 'slow' transistors
 - 3 wafers with 'fast' transistors
 - Each wafer contains ≈ 300 AM06 chips
- The prototypes are working; **no redesign is needed**
- Tests on the samples from pilot run show a **high yield** ($> 80\%$ of chips with zero defects, one or two catastrophic faults)
- The industrial test setup is ready for production
- **The current consumption is a critical issue** and requires special care in package and board design

- AM06 production starting in Oct. 2016

THANK YOU !

The AM06 Design & Test Team:

Alberto Annovi¹, Matteo M. Beretta², Giovanni Calderini³, Francesco Crescioli³,
Luca Frontini^{4,5}, Valentino Liberali^{4,5}, Seyed Ruhollah Shojaii⁴, Alberto Stabile⁴

¹ INFN Pisa

² INFN Frascati

³ LPNHE Paris

⁴ INFN Milano

⁵ Università degli Studi di Milano