




ASSOCIATIVE MEMORY COMPUTING POWER AND ITS SIMULATION

G. Volpi on behalf of the authors
Univ. and INFN of Pisa




The role of tracking for the LHC Run II and beyond




- After the success if the Run I LHC is undergoing major updates
 - Collisions energy will increase from 8 to 13-14 TeV
 - Beam intensity will increase to a luminosity $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and beyond
- The reconstruction of tracks generated by charged particles will be crucial for background rejection
 - Tracking is a computational intensive task
- ATLAS is planning to use a specific tracking processor: Fast Tracker (FTK)
 - Use the power of Associative Memories and FPGAs

Tracking detector allows to separate different collisions and the originated objects.
Specific topologies (i.e. tau- or b-jet) can be efficiently identified with high quality track reconstruction

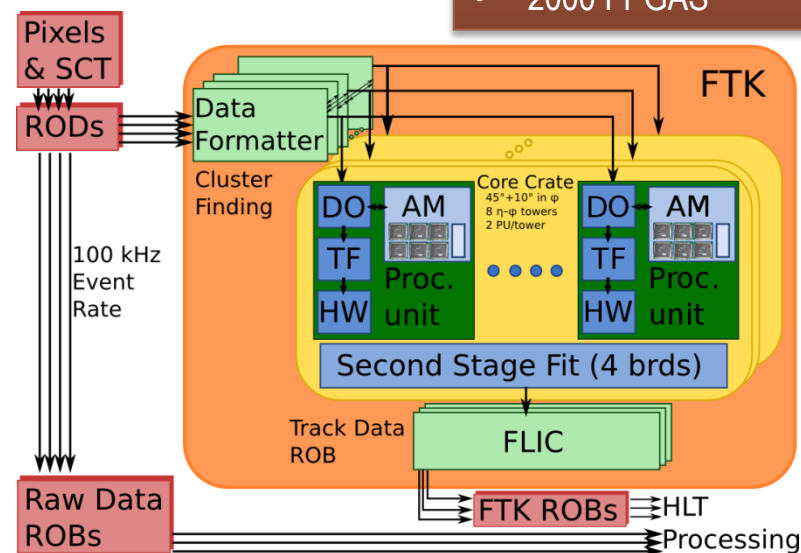


Fast Tracker Processor Reminder

- The FTK processor will provide full tracking for each Lvl1 accept (100 kHz)
 - Expected latency $< 100 \mu\text{s}$
 - Reconstruct all tracks with $p_T > 1 \text{ GeV}$, $|\eta| < 2.5$, $|d_0| < 2 \text{ mm}$ from the BL
- Data are obtained receiving raw data send from the Pixel and strip detectors read-out drivers
 - FTK works in parallel with the existing ATLAS DAQ infrastructure
- Highly parallel system organized in 64 η - ϕ towers
 - Internal data transmission uses high speed serial links



- 2 9U VME boards/tower
- 4 tracking engines each
- 8192 AM Chips
- ~2000 FPGAs





See P. Luciano [talk](#) on Tuesday@III.c

FTK Main Algorithms

FTK has a custom clustering algorithm, running on FPGAs

See. L.C. Sotiropoulou [talk](#) on Tue. III.c

The data are geometrically distributed to the processing units and compared to existing track patterns.



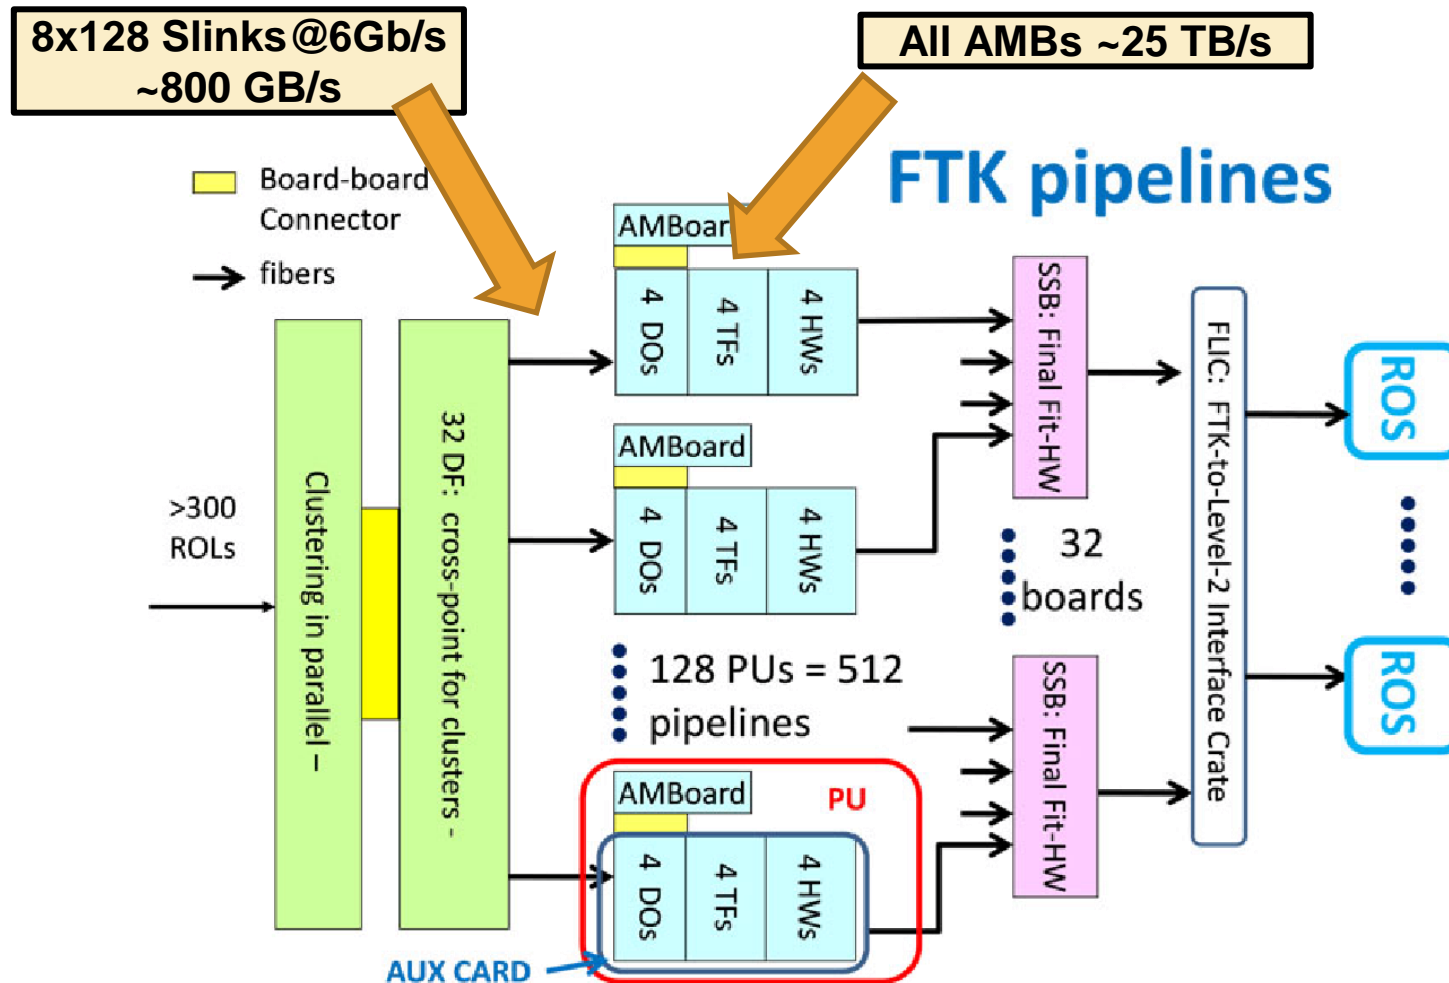
Good 8-layer tracks are extrapolated to additional layers, improving the fit

$$p_i = \sum_j C_{ij} \cdot x_j + q_i$$

Pattern matching limited to 8 layers: 3 pixels + 5 SCTs. Hits compared at reduced resolution.

Full hits precision restored in good roads.
Fits reduced to scalar products.

FTK Electronic Pipeline I/O

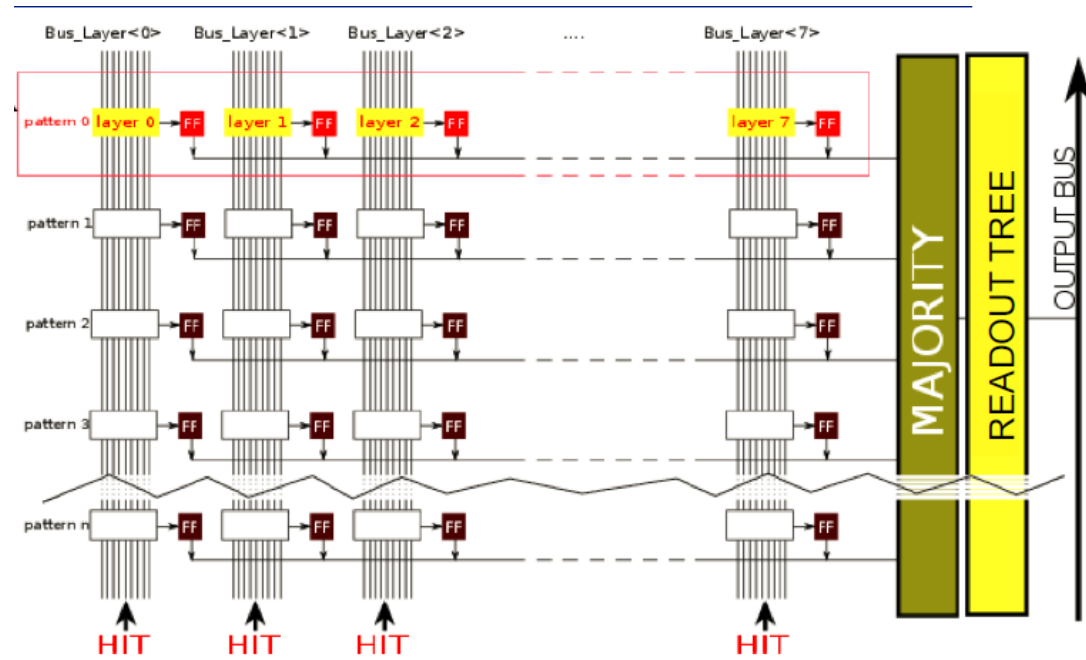


Pattern Matching with AM chip

AM CONSUMPTION:

~ 2.5 W for
128 kpatterns

- The AM chip is a special CAM chip
- The AM identifies the presence of stored patterns in the incoming data
 - Input data arrives through independent busses
 - Patterns with enough matching data are selected
 - Threshold can be reprogrammed
 - DC feature (next slide) allow different match precision
- The chips are installed in boards able to send data to all the chips in parallel
 - At every clock incoming data can be compared with all the patterns




AM COMPUTING POWER

Each pattern can be seen as 4 32 bits comparators, operating at 100 MHz.
 $50 \cdot 10^6 \text{ MIP/chip} \rightarrow 4 \cdot 10^{11} \text{ MIP in the whole AM system}$


Don't care features

- The LSB bit in AM match lines can use up to 6 ternary bits
 - Each bit allow 0, 1 and X (don't care)
 - *K. Pagiamtzis and A. Sheikholeslami, Solid-State Circuits, IEEE Journal of, vol. 41, no. 3, 2006*
 - The DC bits allow to reduce the match precision where required
- The use of DC solves the problem in balancing the match precision
 - Low resolution patterns allow smaller pattern bank size (less chips, less cost), but the probability of random coincidences grows
 - High resolution fakes increase the filtering power at the price of larger banks
 - DC allows to merge similar pattern in favored configurations (less patterns) maintaining high-resolution and rejection power where convenient

Pattern in AM chip w/o the DC




Pattern in AM chip w/ the DC



Fast Tracker Emulation Goal

- The FTK emulation represents the logic of all the different parts to support:
 - Development and design of the architecture
 - Evaluate the physics performance at the trigger level
 - System monitoring during the system installation and operations
- FTK emulation results need to be combined with HW development data and ATLAS High Level Trigger simulation
 - Combining the FTK results with data with firmware and boards development allows precise prediction of the system workflow
 - Feeding the FTK emulation output in the HLTframework allows to characterize the use of FTK tracks and design trigger algorithm
- Need to combine different goals
 - HW design require extremely accurate reproduction of the workflow in the different boards, limited data samples are adequate for the task
 - Physics performance requires realistic representation and million of simulated events to evaluate rejection of background events

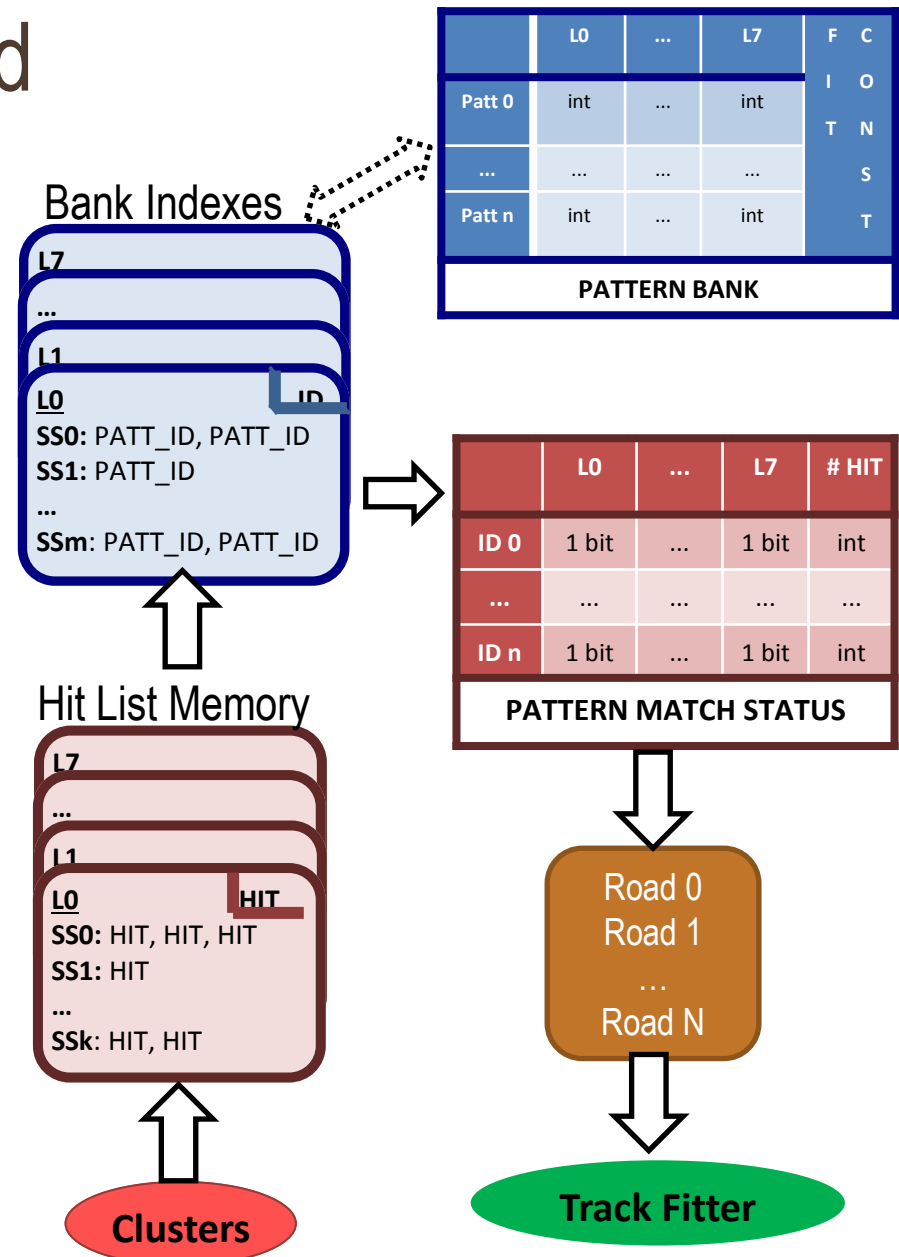
FTK Emulation structure



- The emulation reproduces the entire FTK pipeline
 - Raw hits can be read from MC or real data
 - All the internal steps are reproduced with the possibility to change all internal parameters
 - The impact of firmware or boards' design decisions can be carefully evaluated
- The use of algorithms tailored for a large electronic system is resource demanding
- Memory size and access are the main challenges
 - FTK can have **1 billion patterns** requires about **35 GB**
 - **2 million fit constants** require about **2 GB**
 - Other configuration data are negligible while important
 - Initial FTK configuration will be smaller (256 MP) but with still important memory requirements


AM Emulation Structure and Workflow

- The AM emulation completely reproduce the matching criteria of the real chip
- Workflow differences are required to improve the algorithm performance and allow tests
 - Pattern bank scan uses a large indexing structure to optimize the access time
 - The DC is implemented as filter for the roads
- The pattern bank indexing structure has a large memory footprint
 - Indexing structure is large, almost as the pattern bank
 - Gain w.r.t. the linear scan makes it unavoidable



Emulating a massively parallel tracking processor

- The use of common ATLAS resources has strict resource request
 - Single threaded/process jobs
 - 2 GB/core of real memory
 - Preferred maximum execution time of several hours for few 100s of events
 - Single job uses few GB configuration data
- Natural in FTK to distribute the computation in independent jobs
 - FTK configuration is naturally divided in towers, with further internal segmentation possible
 - Each event is processed many time, using a different partition of the configuration
 - Final data require merging stages, typically organized in 2 steps.
- Single jobs can be sequentially executed on a single node or multiple nodes
 - Usually distributed in parallel to different Grid nodes



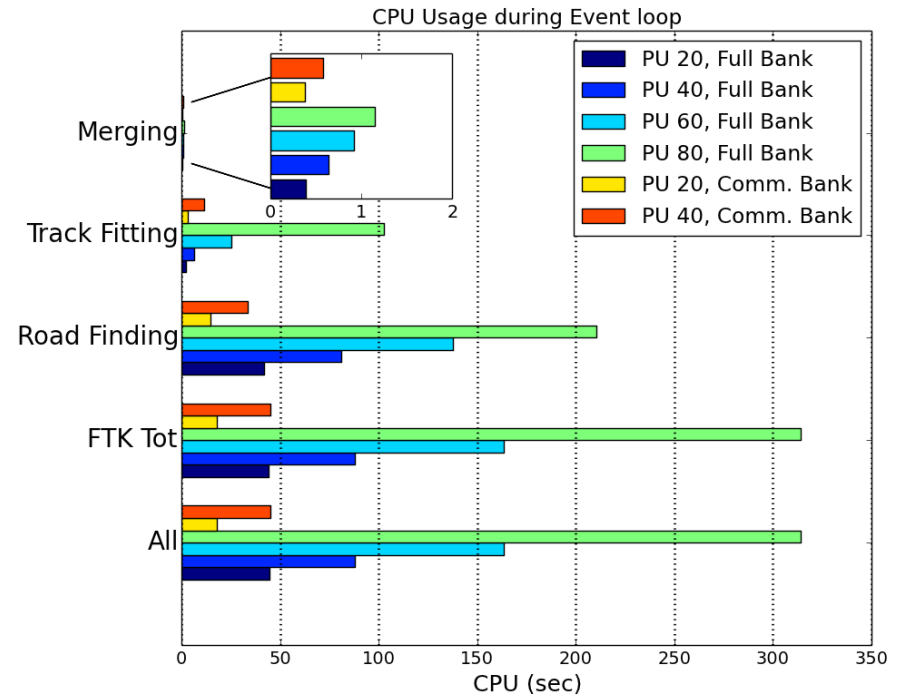
Each tower is segmented to reduce the resource footprint and allow a more parallel use of the resources.

Memory requirement of the single reduced to 1 GB, not dominated by the pattern bank content.

Input data can already be organized in towers reducing the time to retrieve the data.

Performance of the emulation

- The FTK emulation performance is under constants scrutiny
 - Performance evaluated using virtual machines based on Scientific Linux 6
 - VM score ~1800 SpecInt
 - Results obtained summing the results of all the independent jobs and steps for few hundred of events
- CPU change according the FTK conditions
 - Full pattern bank, 1024 MP
 - Commissioning bank, 256 MP
 - Foreseen for the period after the installation
- Execution time change according the expected LHC condition
 - Increasing LHC luminosity will bring more overlapping collisions per bunch crossing (pileup collisions, PU)
- Total execution time <50 sec/evt for the small FTK configuration
- Full FTK configuration requires about 300 sec/evt





Road finding time dominated by the AM emulation, ~66% of the time.

The AM emulation time largely depends from the bank size.



Dependence with the pileup is about linear, showing the efficiency of the match scheme.

Accelerating the AM emulation

- The emulation of the AM logic with standard RAM chips has showed to be difficult
 - Using parallel coprocessor should allow better results
- The AM test board can open the use of AM-based cards as accelerator
 - The test board is based on a Kintex-7 test board
 - Mezzanines with AM chip can be plugged using standard connectors
 - Could use the same mezzanine plugged in the FTK processing units
- Early development to test the different parts
 - Communication with a host PC possible through Ethernet
 - Other protocols can be used, i.e. PCI-express
- Future test will evaluate how to integrate more closely the board with the emulation





FTK system expect performance (latency)




The expected latency of the FTK pipeline has been carefully emulated highlighting how full tracking reconstruction can be achieved within 100 μ s.

Results obtained combining the emulation results with the parameters from the boards' design.

FTK system expect performance: tracking performance



The emulation allowed to evaluate the quality and efficiency of the tracks obtained from FTK. Track quality and efficiency is close to state-of-the-art algorithms, based on CPU and much larger latencies. The overall quality is verified in the use of tracks found by FTK in complex identification algorithms, as b-tagging.



Conclusions


- The FTK will be an important technological update for the ATLAS trigger and data acquisition system
- The use of custom hardware based on AM chips and FPGAs allows to use a tremendous amount of computing power
 - FTK will allow full track reconstruction for the full tracking geometry within 100 μ s
 - The implementation of the algorithms uses the intrinsic parallelism in those devices
- To help the design of the system a detailed simulation of the system has been developed
 - The emulation of such large system has been extremely challenging
 - Reproducing algorithms specifically thought for parallel devices on commercial CPU has revealed several many issues
 - The emulation of the AM remains extremely demanding despite the many efforts to reduce the bottlenecks

Acknowledgements


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.

BACKUP SLIDES

The AM chip history




CDF



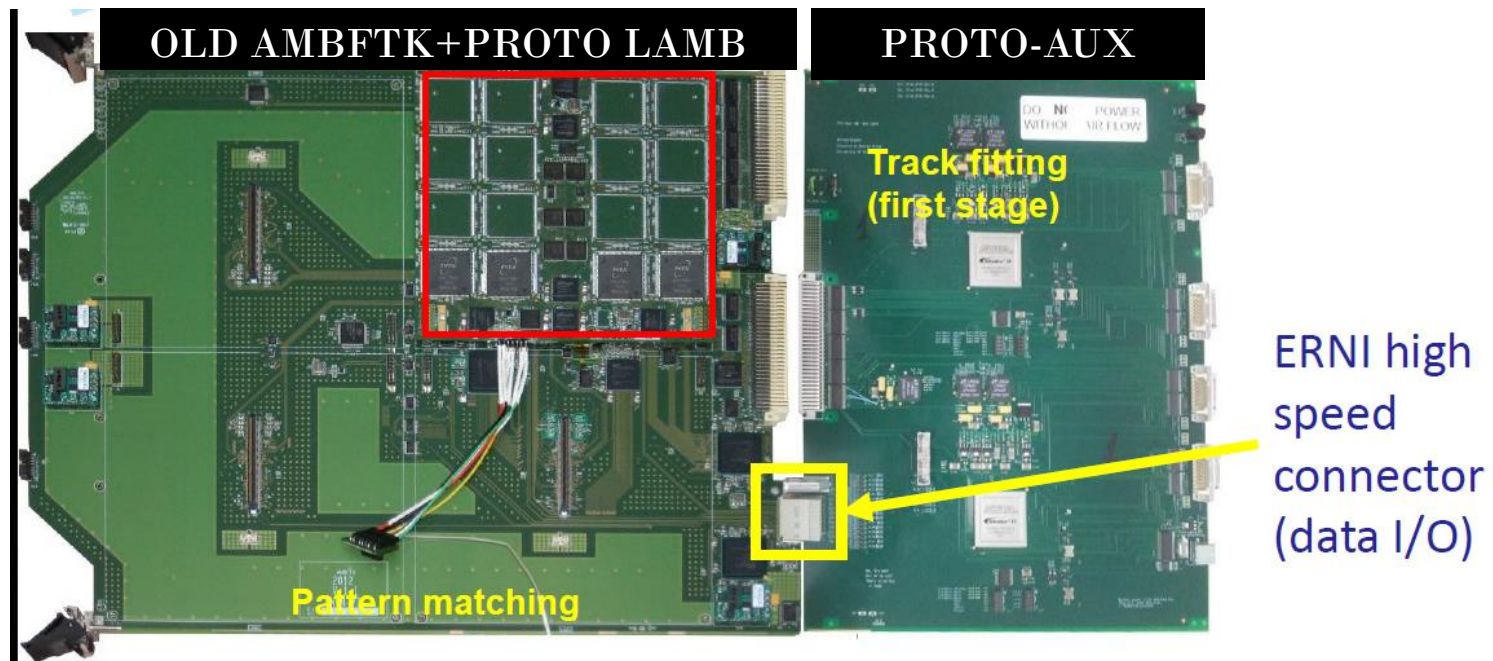
- ▣ **90's Full custom VLSI chip – 0,7 mm AMS (INFN-Pisa) 128 patterns**, 6x12 bit words each (F. Morsani et al., The AM chip: a Full-custom MOS VLSI Associative memory for Pattern Recognition, IEEE Trans. on Nucl. Sci., vol. 39, pp. 795-797 (1992).) **25 MHz clock**
- ▣ **1998 FPGA (Xilinx 5000)** for the same AMchip (P. Giannetti et al., A Programmable Associative Memory for Track Finding, Nucl. Instr. and Meth., vol. A 413/2-3, pp.367-373, (1998)).
- ▣ **1999 first standard cell** project presented at LHCC
- ▣ **2006 AMChip 03 Standard Cell UMC 0,18 mm, 5k patterns in 100 mm²** for CDF SVT upgrade total: AM patterns (L. Sartori, A. Annovi et al., A VLSI Processor for Fast Track Finding Based on Content Addressable Memories, IEEE TNS, Vol 53, Issue 4, Part 2, Aug. 2006). **50 MHz clock**

ATLAS



- ▣ **2012 AMchip04 (Full custom/Std cell) TSMC 65 nm LP technology, 8k patterns in 14mm²** Pattern density x12. First variable resolution implementation. (F. Alberti et al, 2013 JINST & C01040, doi:10.1088/1748-0221/8/01/C01040) **100 MHz**
- ▣ **2013 AMchip05, 4k patterns in 12 mm²** a further step towards final AMchip version. **Serialized I/O buses** at 2 Gbs, further power reduction approach. BGA 23x23 package.
- ▣ **2014 AMchip06: 128k patterns in 180 mm²**. Final version of the AMchip for the ATLAS experiment.

Associative memory board (AMBFTK) and Large Area Mezzanine board (LAMB)



- Associative memory board receives coarse resolution hits from the AUX using high speed connector
- Each board is composed of 4 LAMBs with AM chips
 - Each LAMB-FTK will contain 16 chips, $\sim 10^6$ patterns/LAMB