

THE ATLAS FAST TRACKER PROCESSOR DESIGN

VERTEX 2015 - Santa Fe

G. Volpi, University and INFN of Pisa

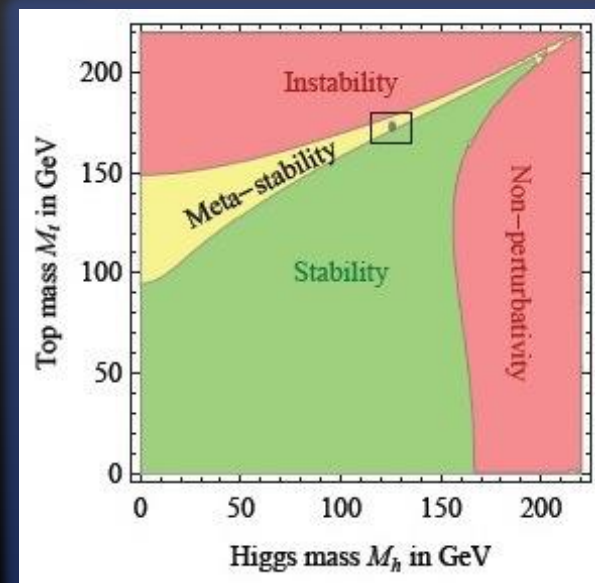

on behalf of ATLAS collaboration and FTK team



Introduction

- ATLAS and the LHC achieved important results during the Run I
 - Higgs boson discovery
 - Strong constraints on new physics phenomena
- Many questions still remain unsolved
 - Search for dark matter candidates
 - Naturalness of the Higgs boson and more
- The data taking conditions for the next running period will offer opportunity for new discoveries:
 - The beams' energy will grow from 4 to 6.5 TeV, increasing the production cross-section of heavier particles
 - The luminosity will increase, reaching 2-3 times the LHC design

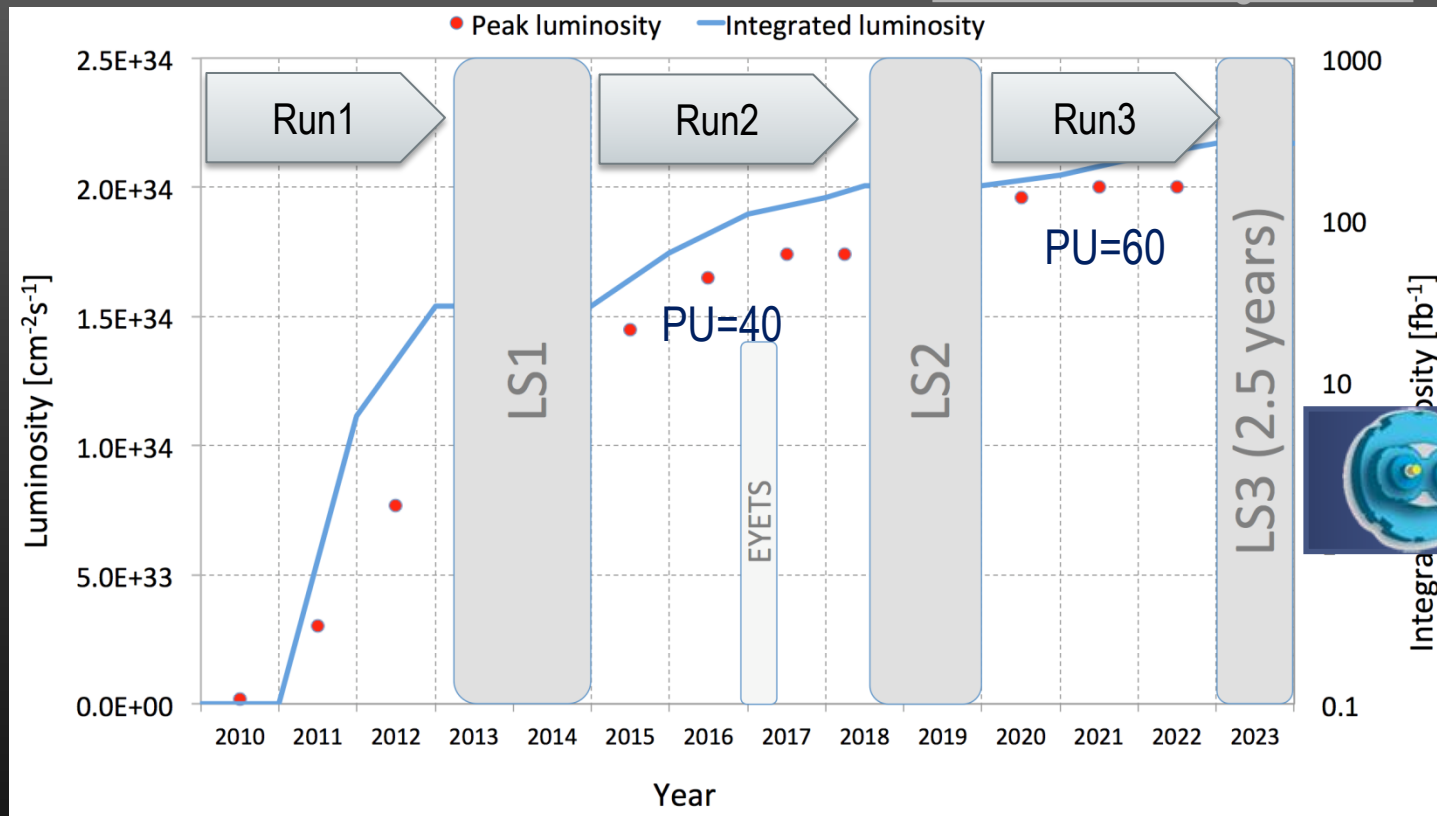
[Phys. Rev. D 91, 012006 \(2015\)](#)



[arXiv:1205.6497](#)

LHC Upgrade Schedule




lhccommissioning.web.cern.ch



- Current shutdown phase just completed
- The next to LHC runs will bring more luminosity but increase the number of contemporary collisions (pileup, PU) up to 60 or more
- Preparation for a planned Phase II of the project called High Luminosity LHC

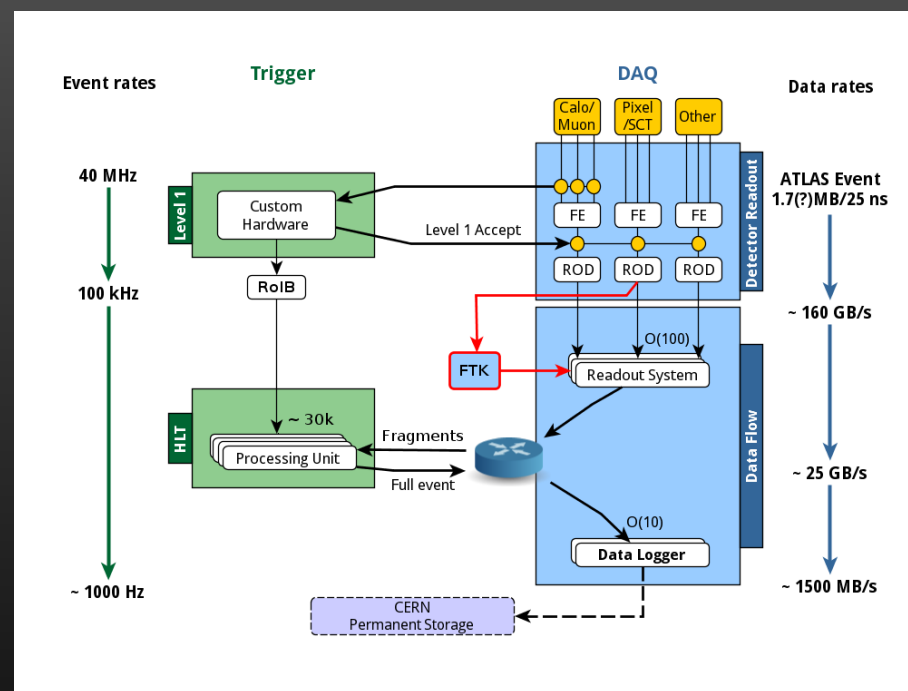
Use of tracking in the trigger

- Data taking conditions will be more difficult
 - More stable particles produced per collision
 - More contemporary collisions (pileup) in each bunch
- The use of tracks many object selections important in Run II
 - Helps to resolve the topology of b- and τ -jets
 - Determines the number and position of the primary vertex
 - Improves robustness in jets and missing energy selections in high pileup events
- Important ingredient for Run II and beyond
 - The increase in event complexity requires more advanced selections





ATLAS Trigger upgrade and the Fast Tracker

- The ATLAS High Level Trigger has undergone several updates for Run II
 - The internal subdivision between Level-2 and HLT has been removed
 - The network infrastructure has been improved
- Software based full tracking will still be limited to a fraction of the Level-1 triggers
 - Full tracking will be provided for every Level-1 trigger (up to 100 KHz)
 - Any trigger selection will be able to exploit the track information
 - FTK tracks can be used to bootstrap other tracking algorithms
- The inclusion of the Fast Tracker (FTK) processor will fill this gap
 - Full tracking will be provided for every Level-1 trigger (up to 100 KHz)
 - Any trigger selection will be able to exploit the track information
 - FTK tracks can be used to bootstrap other tracking algorithms



FTK will receive data in parallel to the HLT farm. Its output will be available to the HLT before the algorithms start.

The FTK basic ideas

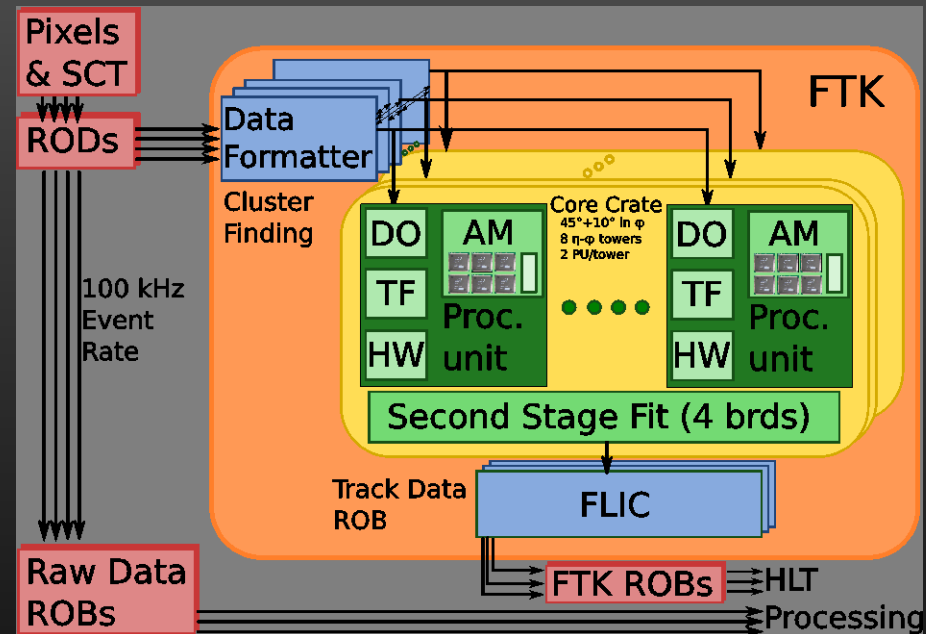


- FTK has to access full silicon detector data and prepare the full list of tracks before the HLT algorithms start
- Tracking is divided into two main sequential steps
- Pattern matching performed using a large set of predetermined tracks' trajectories
 - FTK tests 1 billion of patterns
 - Coarse resolution tracks candidates
 - FTK filters and organizes inner detector clusters
- Full resolution hits associated to patterns are used for track fitting
 - The track's parameters and χ^2 are evaluated through the full resolution hits using a linear Principal Component Analysis algorithm (j.nima.2003.11.078)




Fast Tracker Implementation

- Computing power organized in many independent processors
 - 64 η - ϕ independent towers
 - 8 independent core processors each
 - Increases throughput and reduces latency
- Track finding exploits custom VLSI chips and high-end FPGAs
 - 8192 AM chips and more than 2000 FPGAs will be installed in a combination of VME and ATCA boards





FTK Algorithms pipeline

FTK has a custom clustering algorithm, running on FPGAs



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


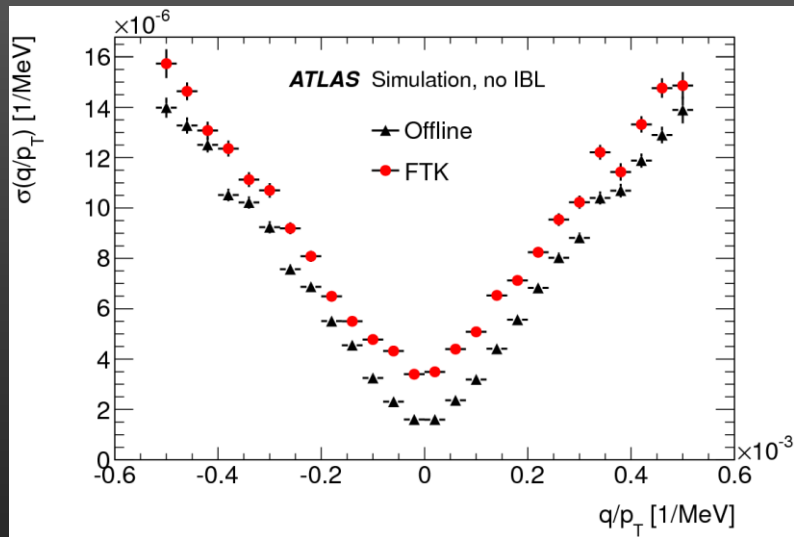
$$p_i = \sum_j C_{ij} \cdot x_j + q_i$$

$$\chi^2 = \sum_i \left(\sum_j A_{ij} \cdot x_j + k_i \right)^2$$

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 system expected tracking performance

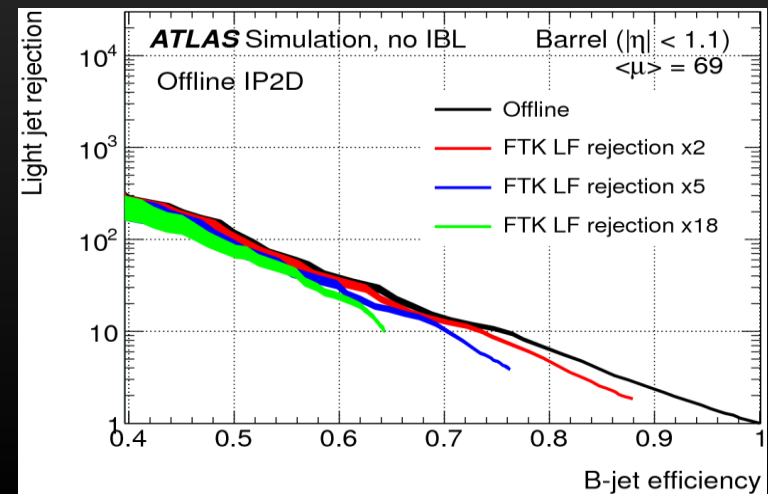


FTK will provide reconstruction for all tracks with $p_T > 1$ GeV in about 100 μ s.

Tracks are expected to be close to those of the offline.

Further quality improvements if the tracks are refitted using offline algorithms or used as seed. Allows the HLT to increase the use of selections based on tracks, mitigating the effect of the large pileup expected in Run II.


Large benefits in complex objects, as b- or τ -jets.




→→Details in the following slides→→

FTK boards pipeline



Data Formatter
(32)





Auxiliary Card
(128)



AM Board
(128)



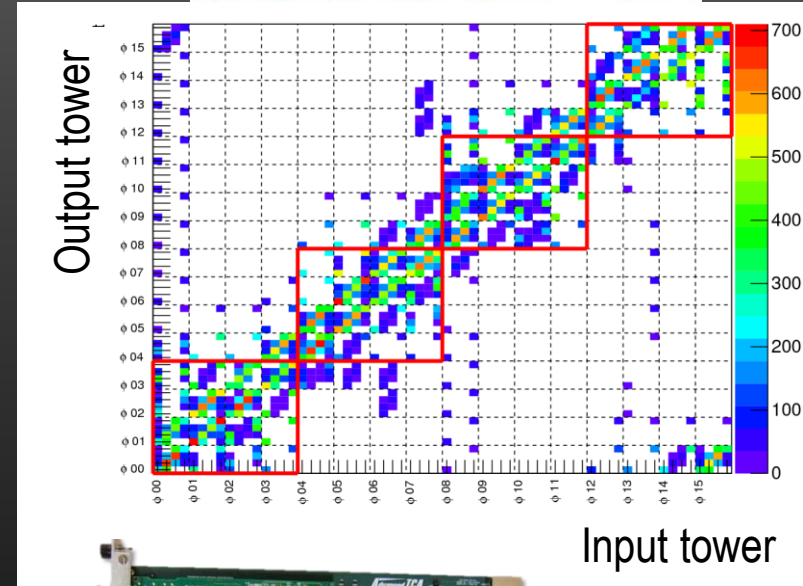

Second Stage Board
(32)




FTK to Level-2 Interface Card
(2)

Data Formatting

- 32 ATCA boards receive data from the whole ATLAS inner detector
 - Input from ATLAS ID read-out drivers
 - Up to 16 links per board through 4 input mezzanines
 - Data need to be organized for the parallel processors
- Data are organized in 64 overlapping towers
 - Each tower covers a $\Delta\phi \times \Delta\eta \approx 32^\circ \times 1.2$
- Board based on a Xilinx Virtex 7 FPGA
 - Large data amount exchanged through a full-mesh backplane
 - Extensive use of high-speed links



Input mezzanine and clustering algorithms



- Up to 128 (32x4) Input Mezzanines would receive approximately 700 Gbps in total


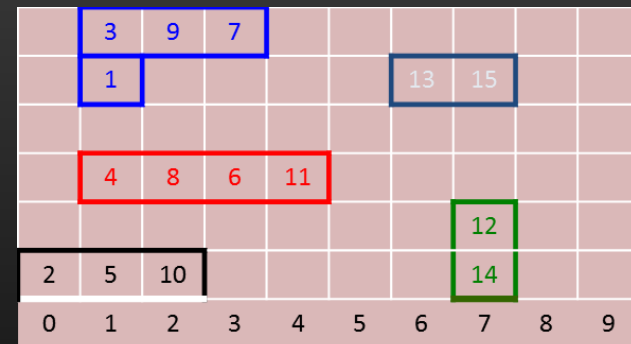

- Each card can receive 4 input links
 - 2 SCT + 2 pixel links

- IM goal is to build the list of clusters at 40 MHz input rate

- 2D Pixel clustering uses a sliding window approach
- Firmware designed and tested



- Computation done by 2 FPGAs

- Each FPGA receives two links: 1 SCT + 1 Pixel
- Two Xilinx FPGA models will be used: Spartan-6 and Artix-7



Auxiliary card

- The AUX (Auxiliary card) is the entry point of the processing units
- It receives clusters from the DF compatible with the tower
 - A pair of AUXs works on a specific eta-phi tower
- The incoming hits are stored according to a coarse resolution identifier (the SS)
- AUX sends SS to the AM board and receives back the full list of roads
 - Road = coarse resolution track candidate
- Calculates the χ^2 for any possible combination of one hit per layer



Main processing units 4 Altera Arria V FPGAs.
Ability to perform 1 fit in 1 ns in each FPGA.

Associative Memory chip design

- The AM chip is a special CAM chip
 - VLSI design using both full custom and standard cells at 65 nm
 - Effort in having low voltage/power device
- The AM identifies the presence of stored patterns in the incoming data
 - Input data arrive through 8 independent busses
 - Minimum number of matched layers programmable
- “Don’t care” (DC) feature allows to change the match precision
 - The match precision is set independently for each layer in each pattern
 - Ability to decrease random coincidences while keeping a limited number of patterns


AM CONSUMPTION:

~ 2.5 W for 128 k patterns

Performing 10^{14} parallel comparisons at 16 bits per second




Pattern in AM chip w/ the DC





Associative Memory Board

- The AM chips are installed in the AMBSLP board
 - They control the pattern matching phase
- SS are received from the AUX and distributed to the AM chip
 - Chips are distributed in 4 Local AMB (LAMB) mezzanines
 - In each mezzanine the incoming SSs are then distributed to 16 chips
 - The same input is distributed to all the 64 chips in parallel through high speed serial links
- Found roads and match information are distributed back to the AUX
- I/O through high speed serial links
- Large computing power
 - Each pattern can be seen as 4 comparators at 32 bits, operating at 100 MHz
 - $50 \cdot 10^6$ MIP/chip \rightarrow $4 \cdot 10^{11}$ MIP in the whole AM system



Second Stage Board (SSB)

- Receives track candidates from the AUX cards and attempts to improve the precision with additional hits
- Each incoming 8-layer track is extrapolated to the 4 layers which are not used in the pattern matching
 - The extrapolation uses linear approximation
 - The additional hits results in better fake rejection and improved track resolution
- Performs track duplicates removal
- Each board works with 2 towers: 4 AUX \rightarrow 1 SSB



FTK to Level 2 Interface Card (FLIC)

- FLIC receives the final tracks from the SSB
- Interface board with the HLT farm
 - FTK tracks will be represented using the same format used by the CPU-based algorithms
- SSB data format describes track parameters and associated silicon clusters
 - Format designed to save bandwidth
 - Global detector element identifiers removed in favor of local ID
 - Track parameters represented as fixed point precisions
 - Global detector element identifiers for the clusters restored using lookup tables
 - Track parameters converted to floating points in standard units



Expected FTK Installation Schedule

Step	IM	DF	AUX	AMB	Chip v.	SSB	FLIC	Milestones	Expected
1 st	4-16	1-4	1	1	x05	1	2	Full Slice Test	09/15
2 nd	128	32	16	1	x06	8	2	Full Slice Test	11/15
3 rd	128	32	16	16	x06	8	2	Full Barrel, $\mu=40$	02/2016
4 th	128	32	32	32	x06	16	2	Full coverage, $\mu=40$	08/2016
Final	128	32	128	128	x06	32	2	TDR Specs	2018/ Lumi-driven

- FTK is a Phase I upgrade project with installation during Run 2
- The installation of the complete slice is expected during the summer
 - Allows testing of complete FTK pipeline with ATLAS
 - The installation of the different parts will be completed according to availability
- All the boards have final prototypes or have been declared ready for production
- First barrel coverage completion in early 2016 with complete detector volume coverage 6 months later
 - Complete installation will be driven by the luminosity evolution of the machine

Conclusions

- Many trigger selections will be able to exploit the additional capabilities provided by FTK
 - Full use of tracks will be beneficial for b- or τ -jet identification
- The ATLAS Fast Tracker processor will start taking data in 2016
 - Designed on a combination of ATCA and VME boards, custom and commercial chips
 - Allows full track reconstruction for the beginning of HLT processing
- All boards are in final development stage or at the beginning of production
 - Full slice tests expected during 2015
 - First full integration in data taking expected during summer 2016
- The core technologies used by FTK (AM chips and FPGAs) will play a key role in HL-LHC
 - ATLAS Level 1 Track Trigger upgrade for the HL-LHC is expected to use AM chips for pattern recognition
 - Under consideration the possibility to use an upgraded version of FTK during this period

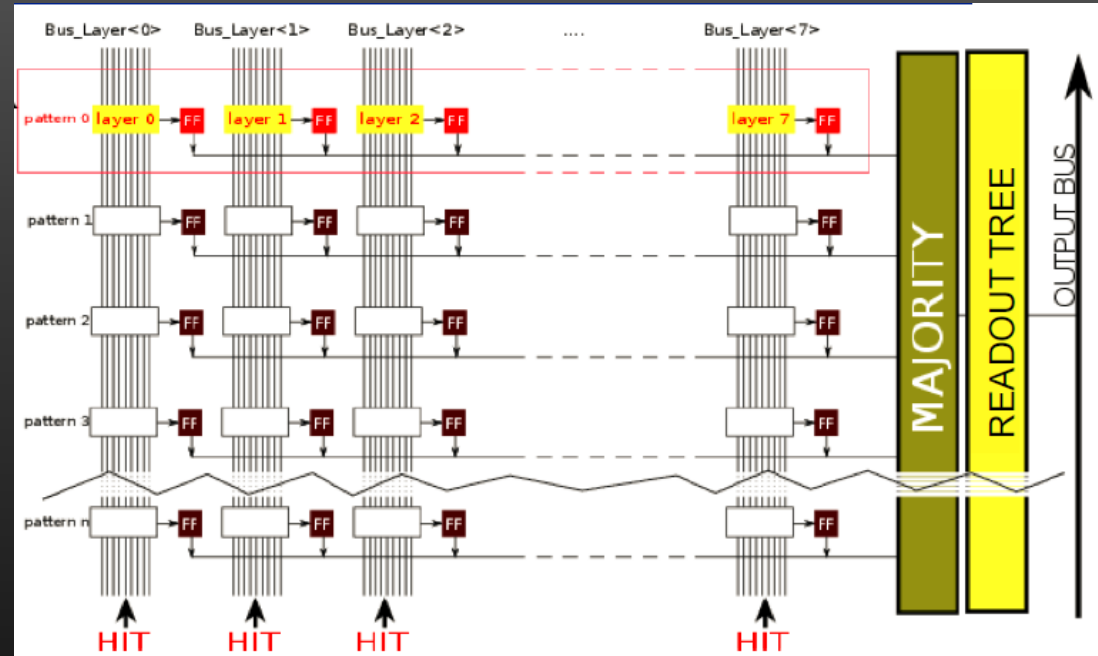
BACKUP SLIDES

Pattern Matching with AM chip

- The AM chip is a special CAM chip
- The AM identify 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 CONSUMPTION:

~ 2.5 W for
128 kpatterns



AM COMPUTING POWER


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

The AM chip history



- **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**

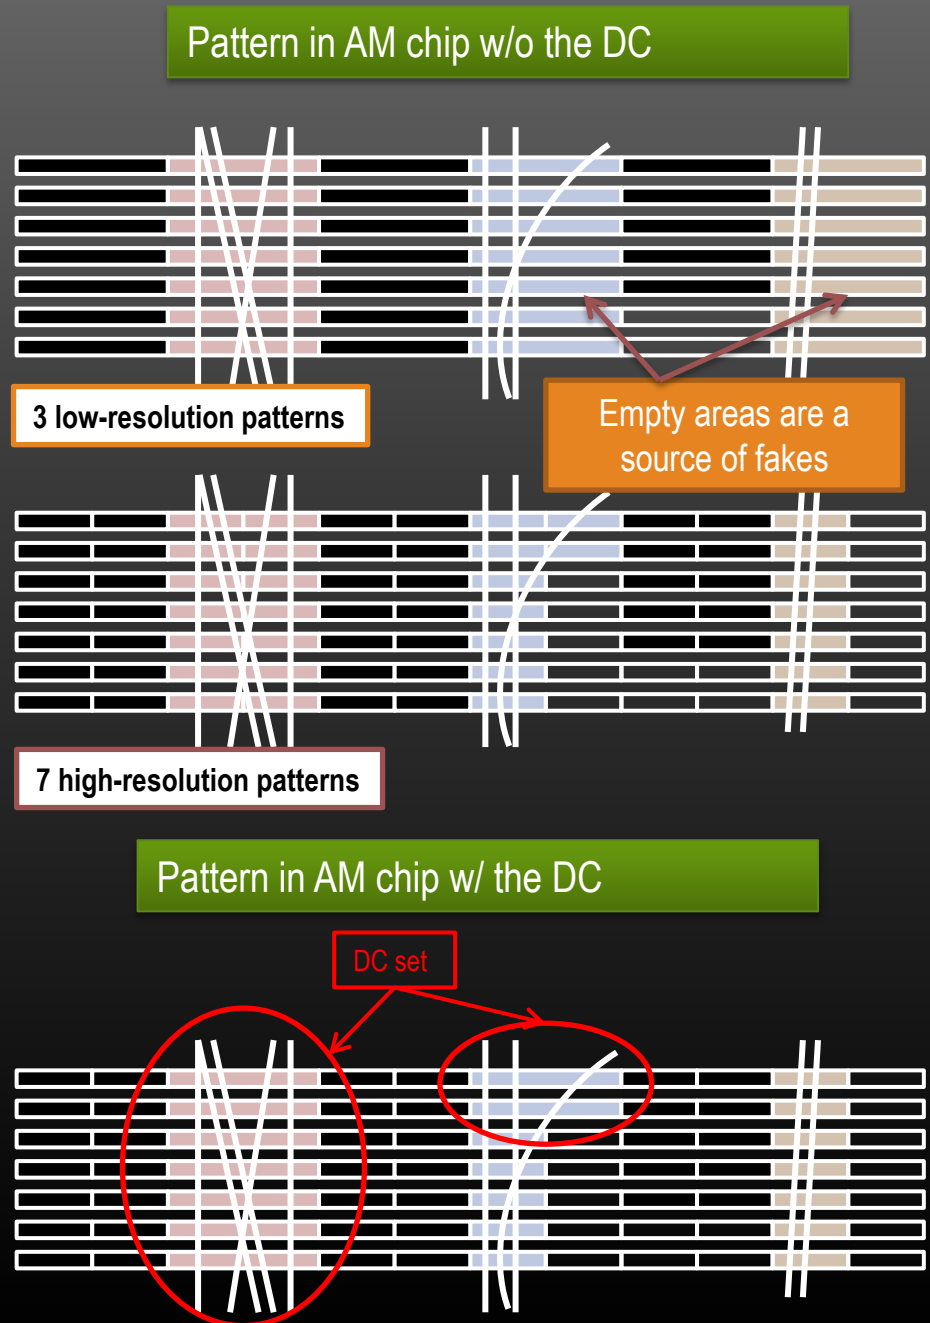
ATLAS


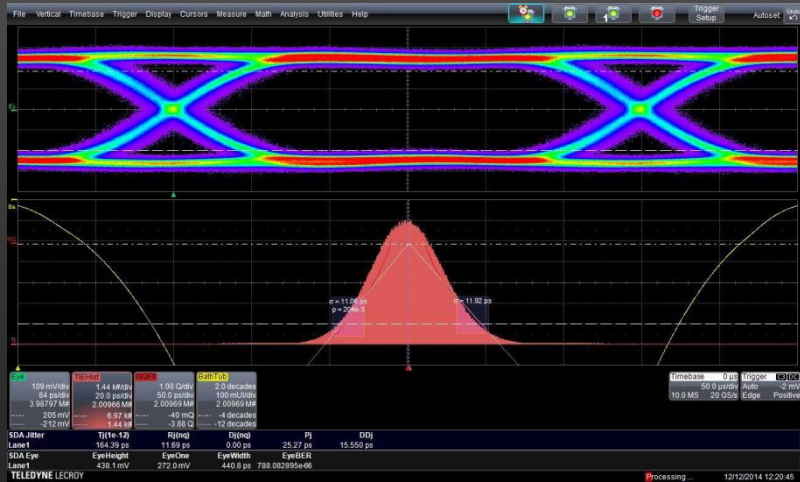



- **2012 AMchip04 (Full custom/Std cell) TSMC 65 nm LP technology, 8k patterns in 14mm²** Pattern density x12. First variable resolution implementation. (F. Alberti et al, 2013 JINST & C01040, doi:10.1088/1748-0221/8/01/C01040) **100 MHz**
- **2013 AMchip05, 4k patterns in 12 mm²** a further step towards final AMchip version. **Serialized I/O buses** at 2 Gbs, further power reduction approach. BGA 23x23 package.
- **End 2015 AMchip06: 128k patterns in 180 mm²**. Final version of the AMchip for the ATLAS experiment.

Don't care features

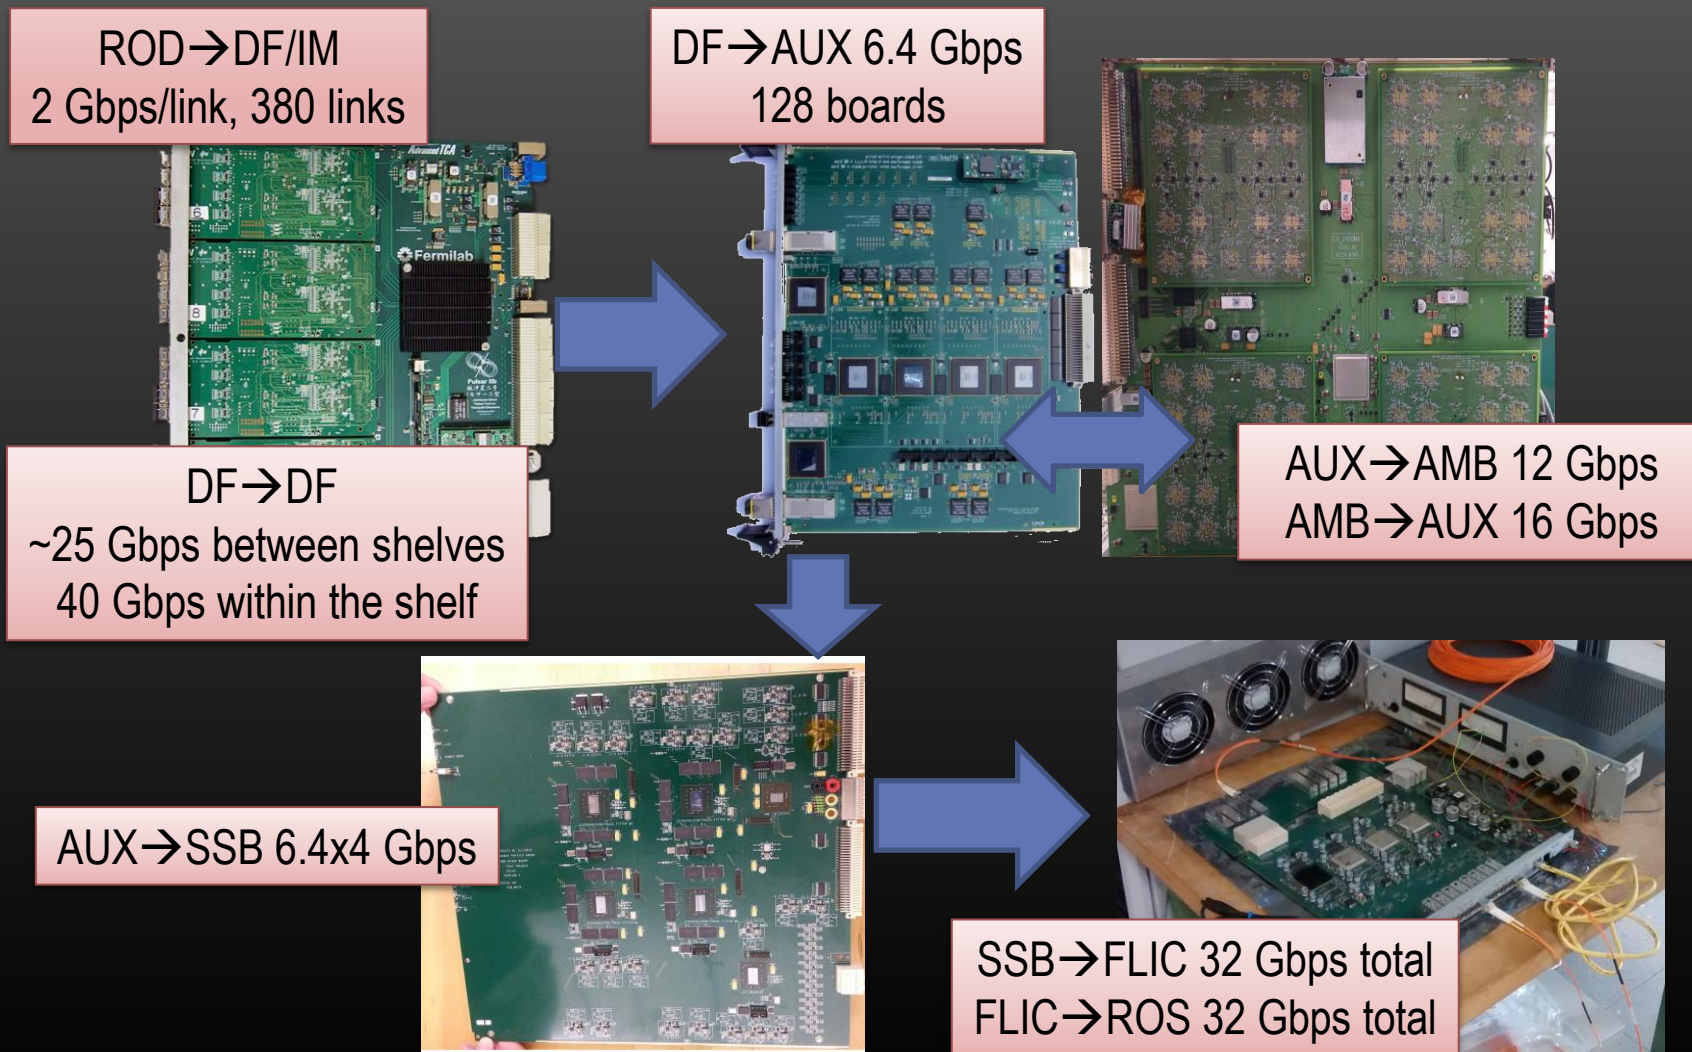
- The LSB bit in AM match lines can use up to 6 ternary bits
 - Each bit allow 0, 1 and X (don't care)
 - K. Pagiamtzis and A. Sheikholeslami, Solid-State Circuits, IEEE Journal of, vol. 41, no. 3, 2006*
 - The DC bits allow to reduce the match precision where required
- The use of DC solves the problem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









FTK Pipeline Bandwidth Summary



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.

