



# MEMORIE ASSOCIATIVE PER SELEZIONE DI EVENTI ONLINE

**Alberto Stabile** on beyond of  
AMchip design team

# CONTEXT AND MAIN GOAL

At hadron colliders:

- **Common problem:** identification of particle tracks in vertex detector

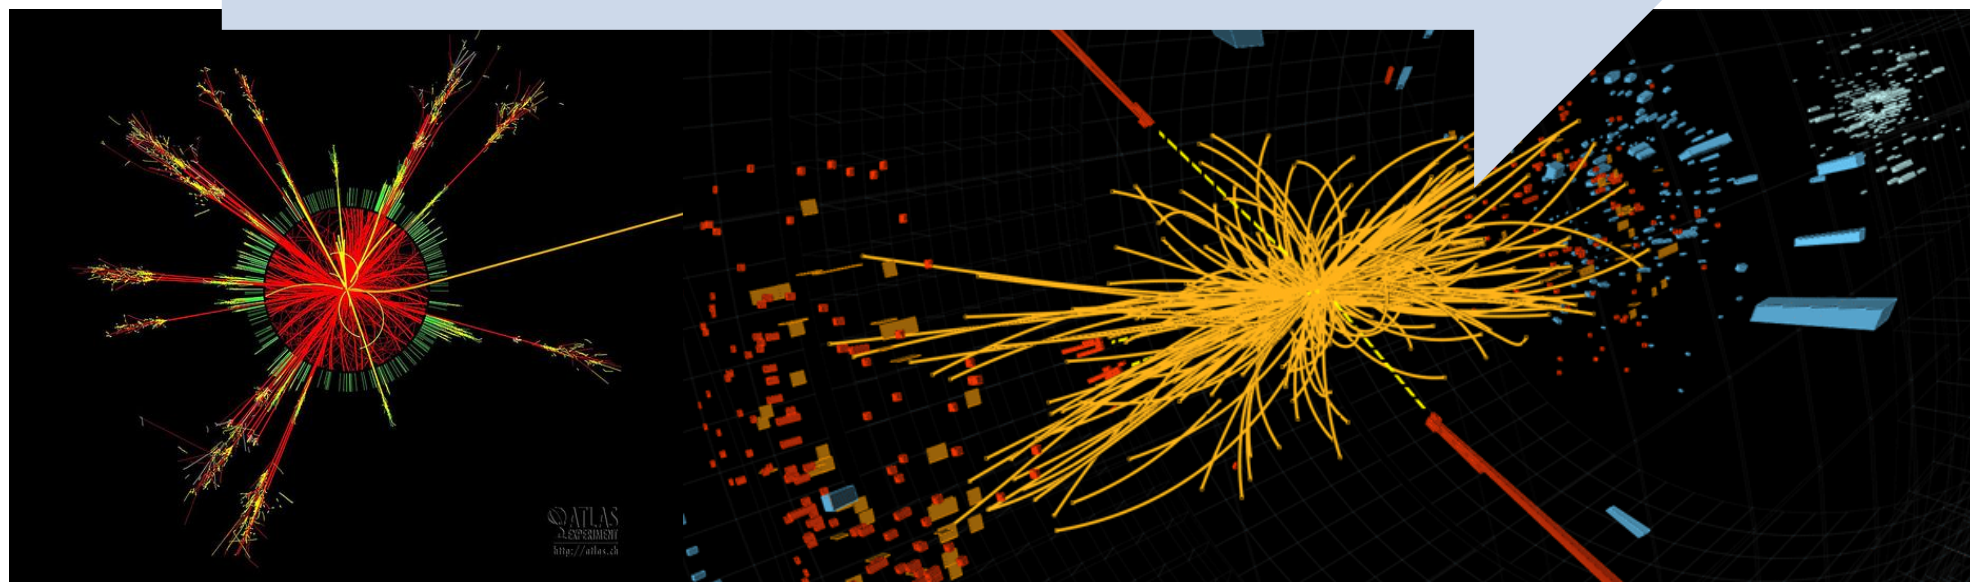
Huge amount of produced data

- **Limited** amount of events can be transferred

Data reduction must be performed

Trigger system

- Particle track recognition in **real time**



# DEDICATED HARDWARE

Pattern recognition problem can be solved by an **Associative Memory (AM) chip**

Integrated circuit

Content Addressable memory

Maximum level of parallelism

Comparison of input data with a set of precomputed patterns stored in a memory

Select few tracks among several tracks



**1990: AM chip v.1** – ASIC CMOS 700 nm with 128 patterns

**1998: AM chip v.2** – FPGA same of AM chip v.1

**2004: AM chip v.3** – ASIC CMOS 180 nm with 5000 patterns

# AMCHIP APPROACH

Two common memory devices are RAMs and CAMs. The Associative Memory is an evolution over the concept of CAM

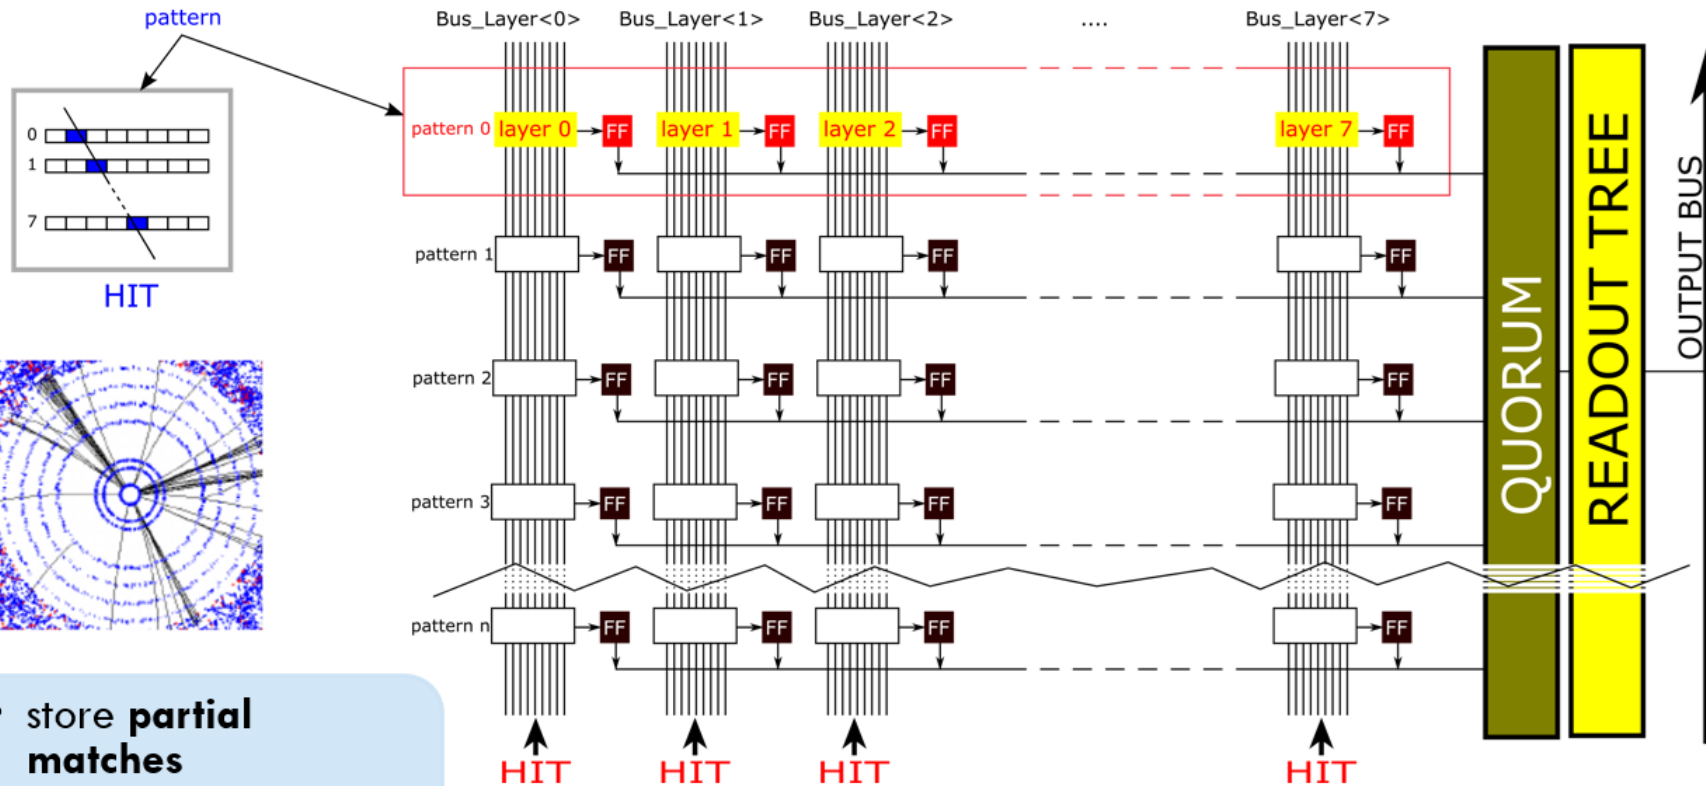
Type	Function	Application
RAM	write data at address read data from address	common memory device used for data storage in information technology
CAM	write data at address find addresses that match data	sparse database search, cache, routing tables
AM	write <b>segmented data</b> at address find addresses that match a combination of segments within a <b>data sample</b>	combinatorial pattern matching, CDF SVT, ATLAS FTK, in future CMS, ATLAS FTK++, and interdisciplinary application with IMPART and IAPP

It is more than a memory device, it is an engine to solve a class of combinatorial problem


# THE AM CHIP ARCHITECTURE

For each bus and for each pattern there is a small **CAM cell array (layer x)**

- It compares its own content with all data received. If it matches a **memory is set (FF)**
- The partial matches are analyzed by **Quorum logic** and compared to the desired threshold
- A **readout encoder (Fischer Tree)** reads the matched patterns in order




The AM and commercially available CAMs differs substantially



The AM provides the unique capability:

- store **partial matches**
- find correlations at **different times**

# AMCHIP TREND



Since 2010:

- **AM chip v. 4** designed and characterized 2012 – area: 14 mm<sup>2</sup> – cost: 50 k€ – 8 kpattens



65 nm

65 nm


2015-2017  
IMPART+RDPHASE\_2

- **AM07 chips** will be designed and characterized within the project for several disciplines:
  - **Image analysis**
  - **DNA sequencing**
  - **Trigger DAQ (CMS\* + ATLAS)**

28 nm + FPGA

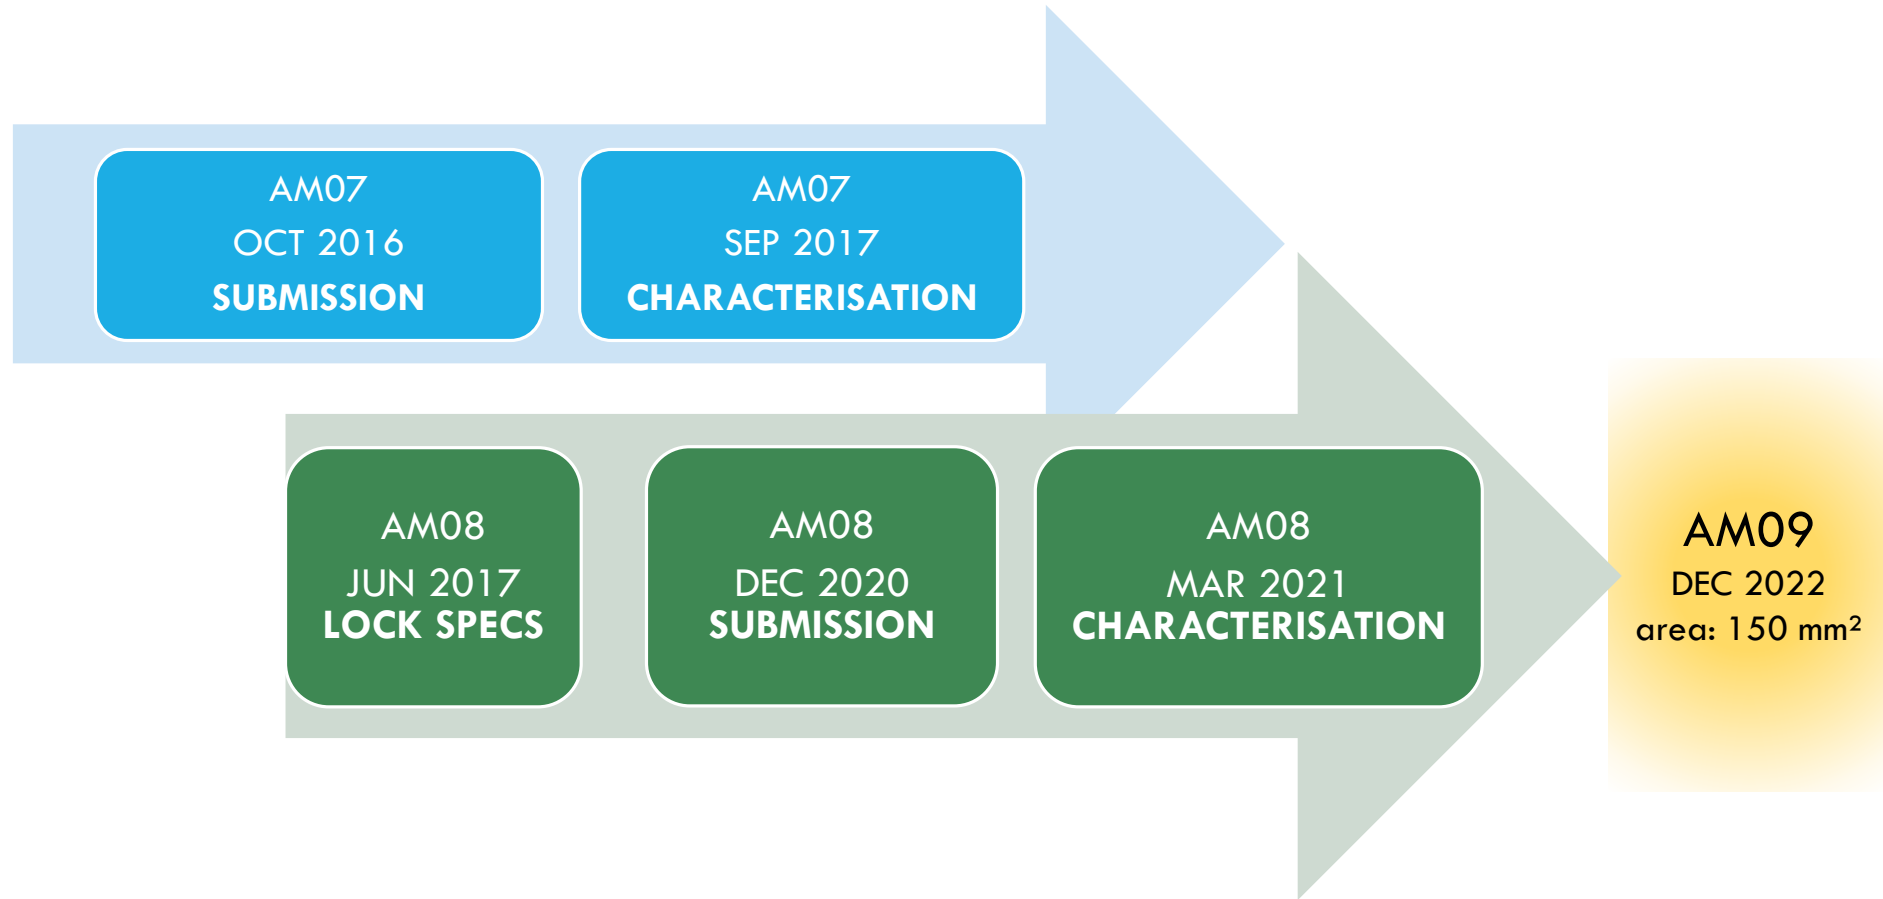
Since 2011:

- **AM chip v. 5-mini@sic** designed and characterized 2012 – area: 4 mm<sup>2</sup> – cost: 20 k€ – 256 pattens for cell test
- **AM chip v. 5-mpw** designed and characterized 2013 – area: 12 mm<sup>2</sup> – cost: 75 k€ – 5 kpattens
- **AM chip v. 6** under design; to be submitted in Dec. 2014 – area: 160 mm<sup>2</sup> – cost: 550 k€ – 128 kpattens - **421 millions transistors**



\*[https://cds.cern.ch/record/2263760/files/CR2017\\_117.pdf](https://cds.cern.ch/record/2263760/files/CR2017_117.pdf)

# SCHEDULE FOR THE NEXT CHIPS




# AMCHIP DESIGN COMPLEXITY VS CPUS


Chip name	Transistor count	Year	Brand	Technology	Area
<a href="#">Core 2 Duo</a> Conroe	291,000,000	2006	Intel	65 nm	143 mm <sup>2</sup>
<a href="#">Itanium 2</a> Madison 6M	410,000,000	2003	Intel	130 nm	374 mm <sup>2</sup>
<a href="#">Core 2 Duo</a> Wolfdale	411,000,000	2007	Intel	45 nm	107 mm <sup>2</sup>
<b>AM06</b>	<b>421,000,000</b>	<b>2014</b>	<b>AMteam</b>	<b>65 nm</b>	<b>168 mm<sup>2</sup></b>
<a href="#">Itanium 2</a> with 9 MB cache	592,000,000	2004	Intel	130 nm	432 mm <sup>2</sup>
<a href="#">Core i7</a> (Quad)	731,000,000	2008	Intel	45 nm	263 mm <sup>2</sup>
Quad-core <a href="#">z196</a> <sup>[20]</sup>	1,400,000,000	2010	IBM	45 nm	512 mm <sup>2</sup>
Quad-core + GPU <a href="#">Core i7 Ivy Bridge</a>	1,400,000,000	2012	Intel	22 nm	160 mm <sup>2</sup>
Quad-core + GPU <a href="#">Core i7 Haswell</a>	1,400,000,000	2014	Intel	22 nm	177 mm <sup>2</sup>
<b>AM09</b>	<b>1,684,000,000</b>	<b>2019</b>	<b>AMteam</b>	<b>28 nm</b>	<b>150 mm<sup>2</sup></b>
Dual-core <a href="#">Itanium 2</a>	1,700,000,000	2006	Intel	90 nm	596 mm <sup>2</sup>



# DESIGN METHODOLOGY

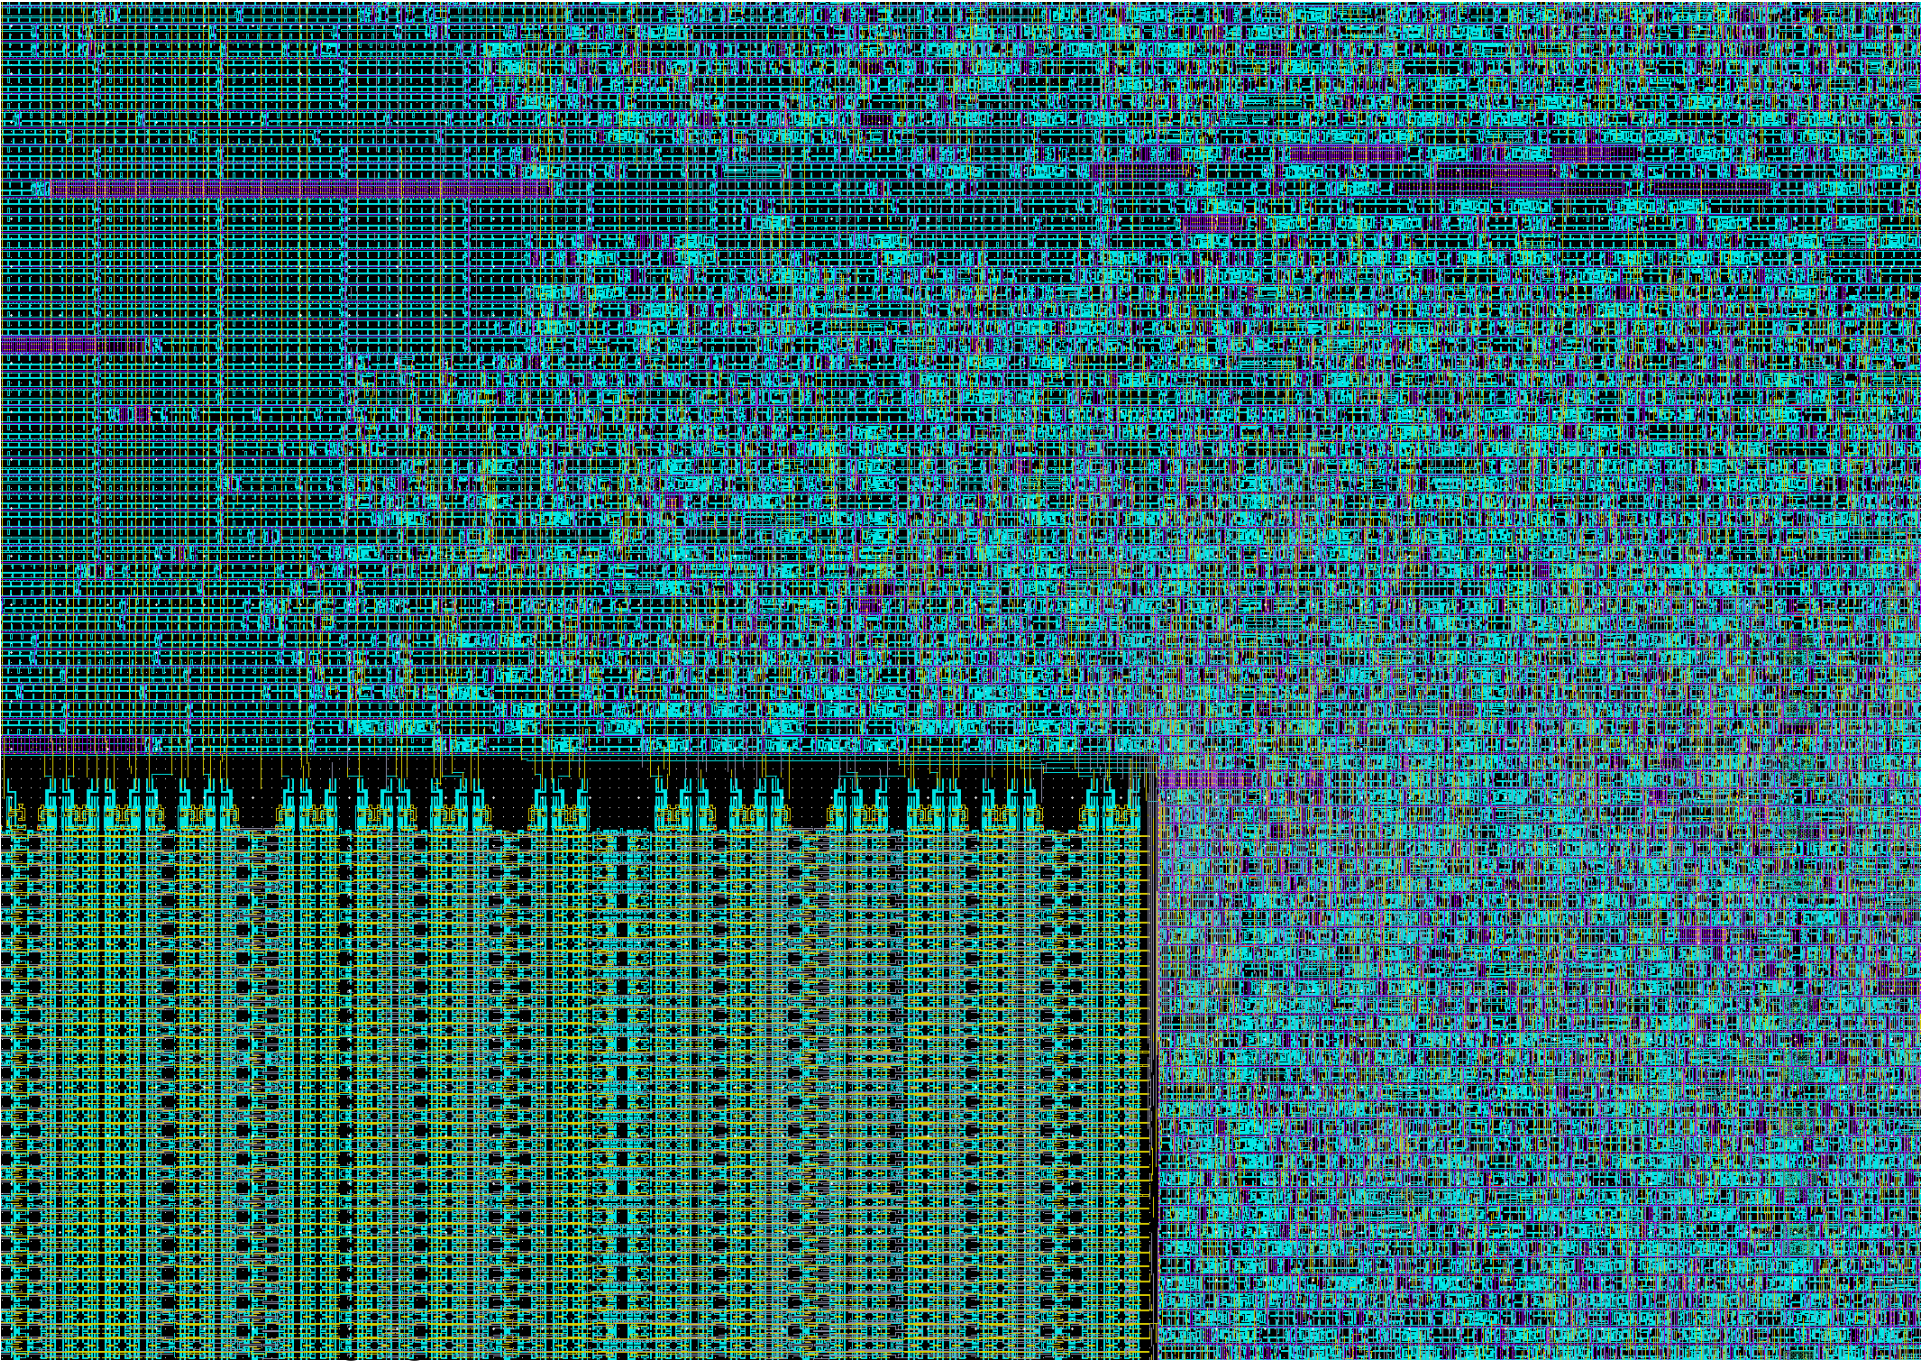


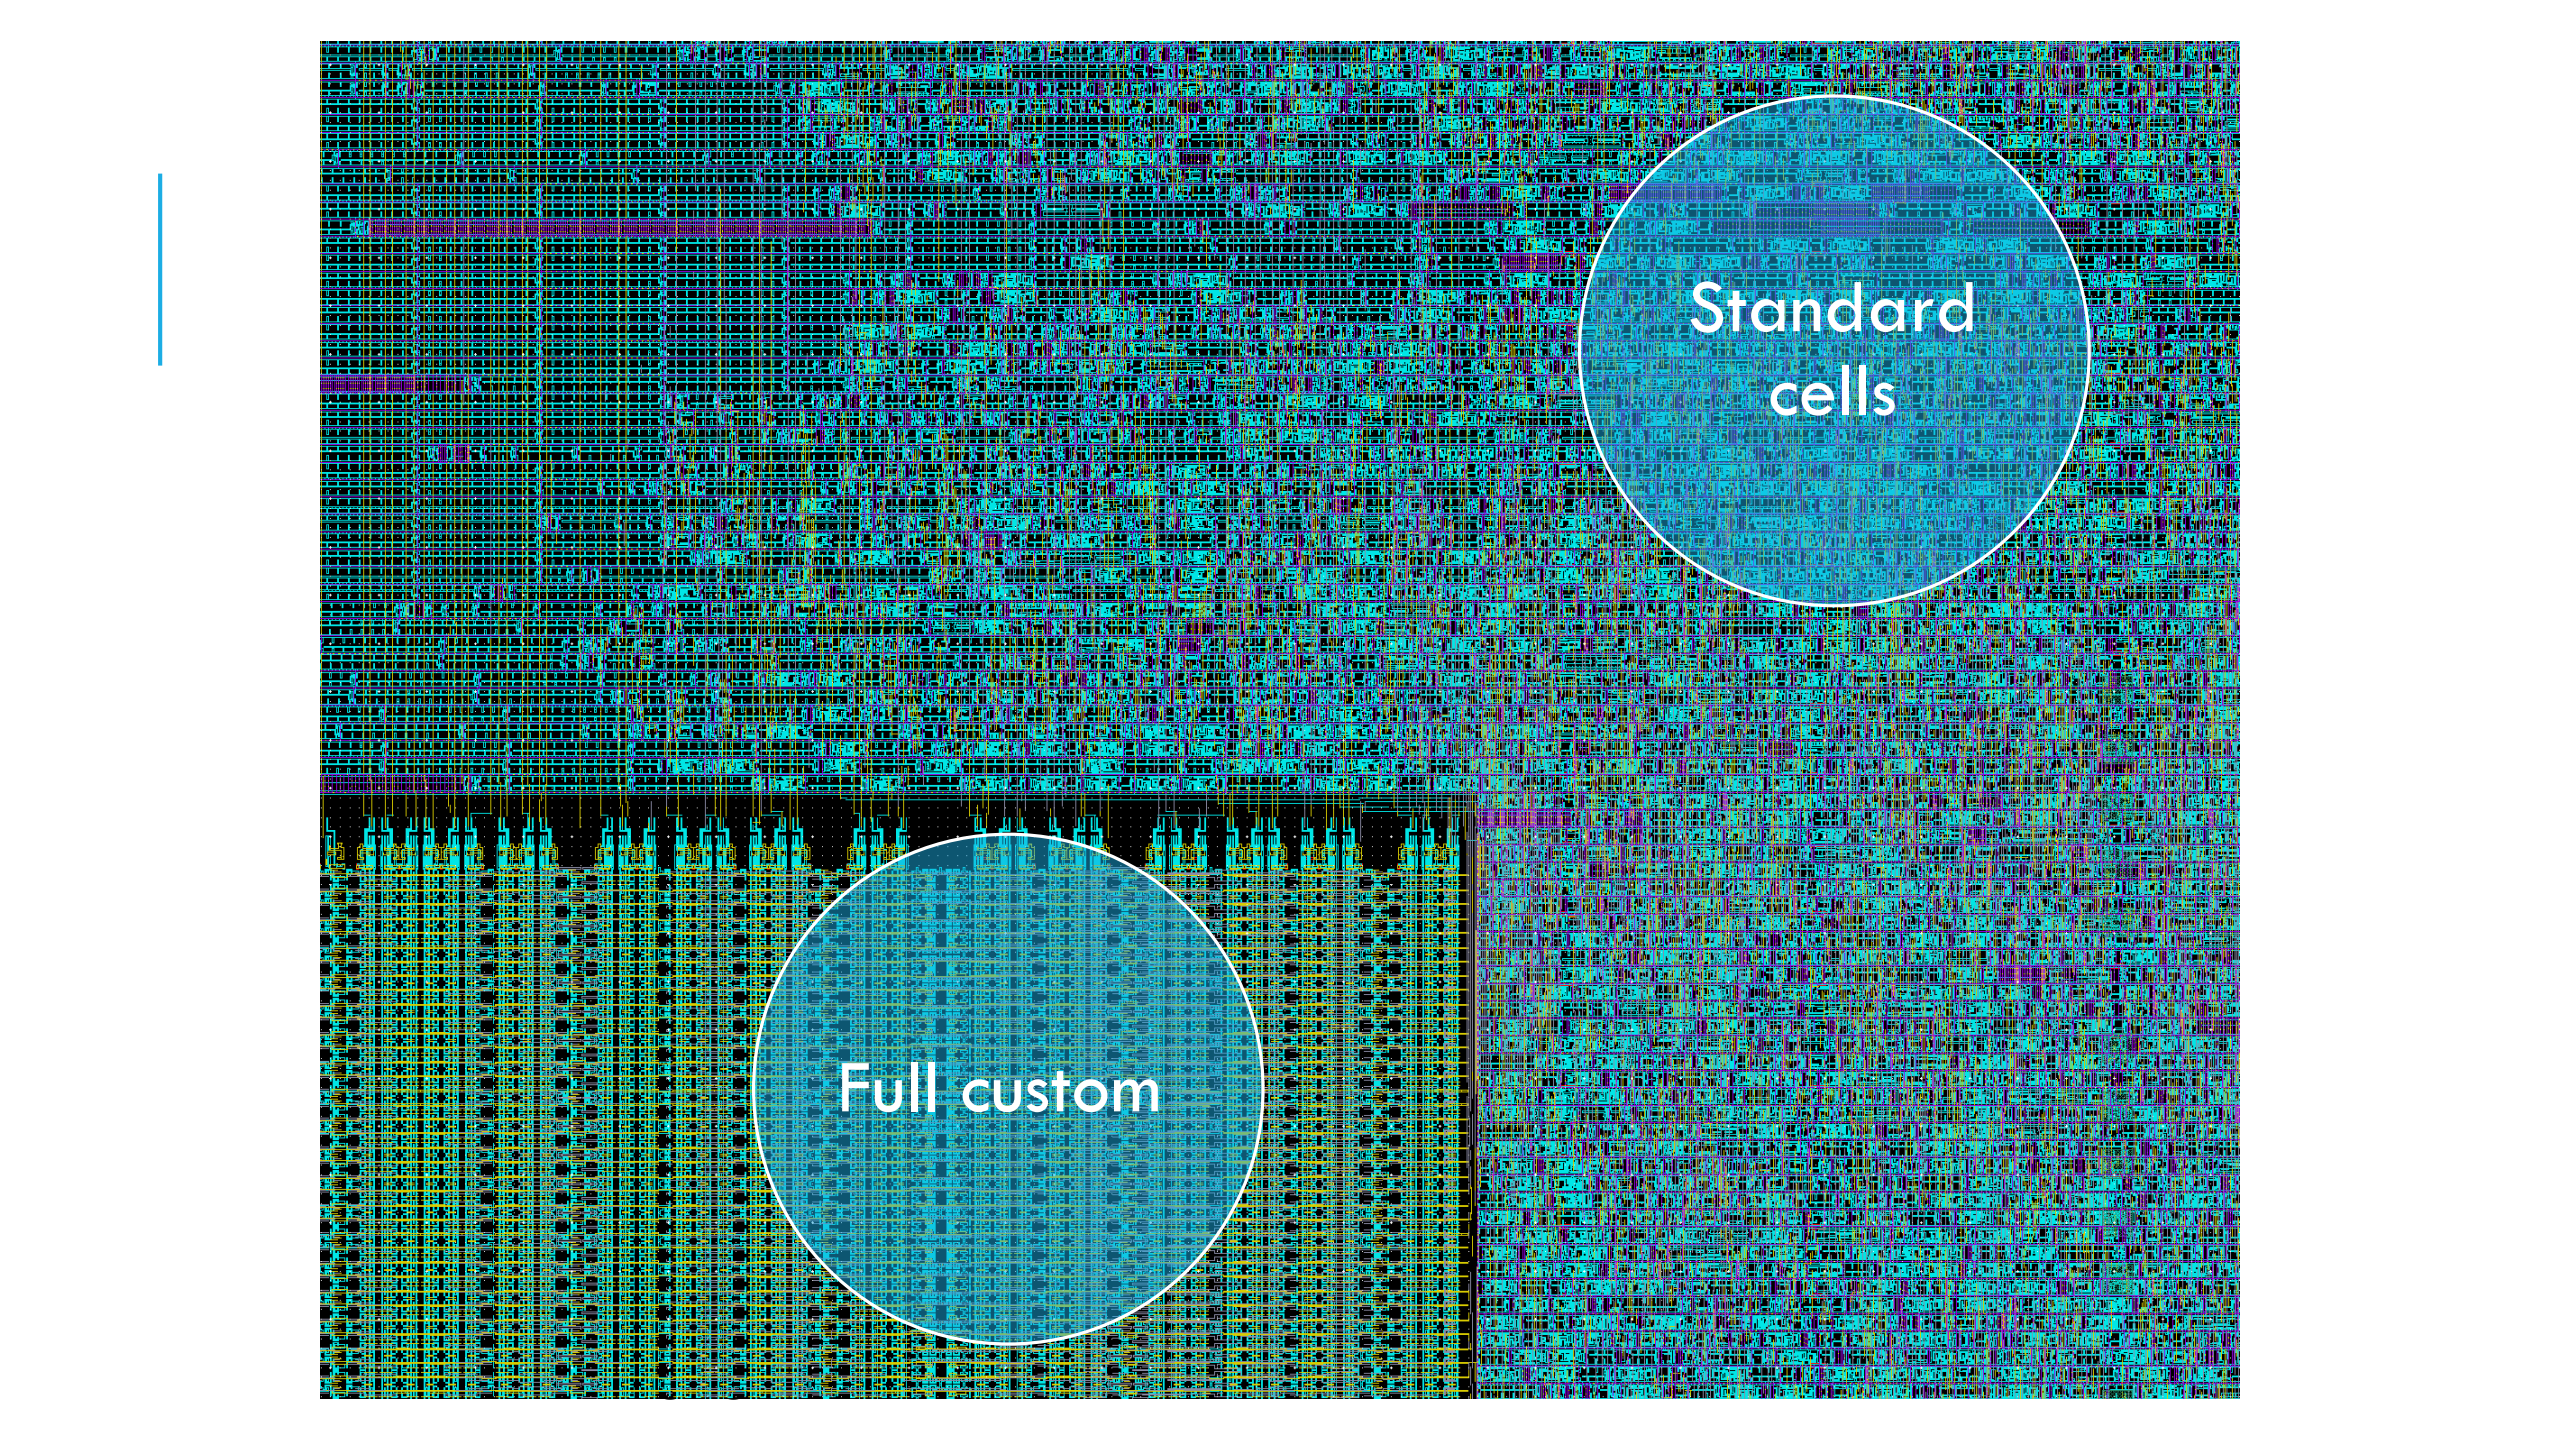
More repetitive parts have been design “by hand” with a full custom approach



More complex logics have been design with automatic tools based on standard cells (synthesis, place & route)

MIXED APPROACH

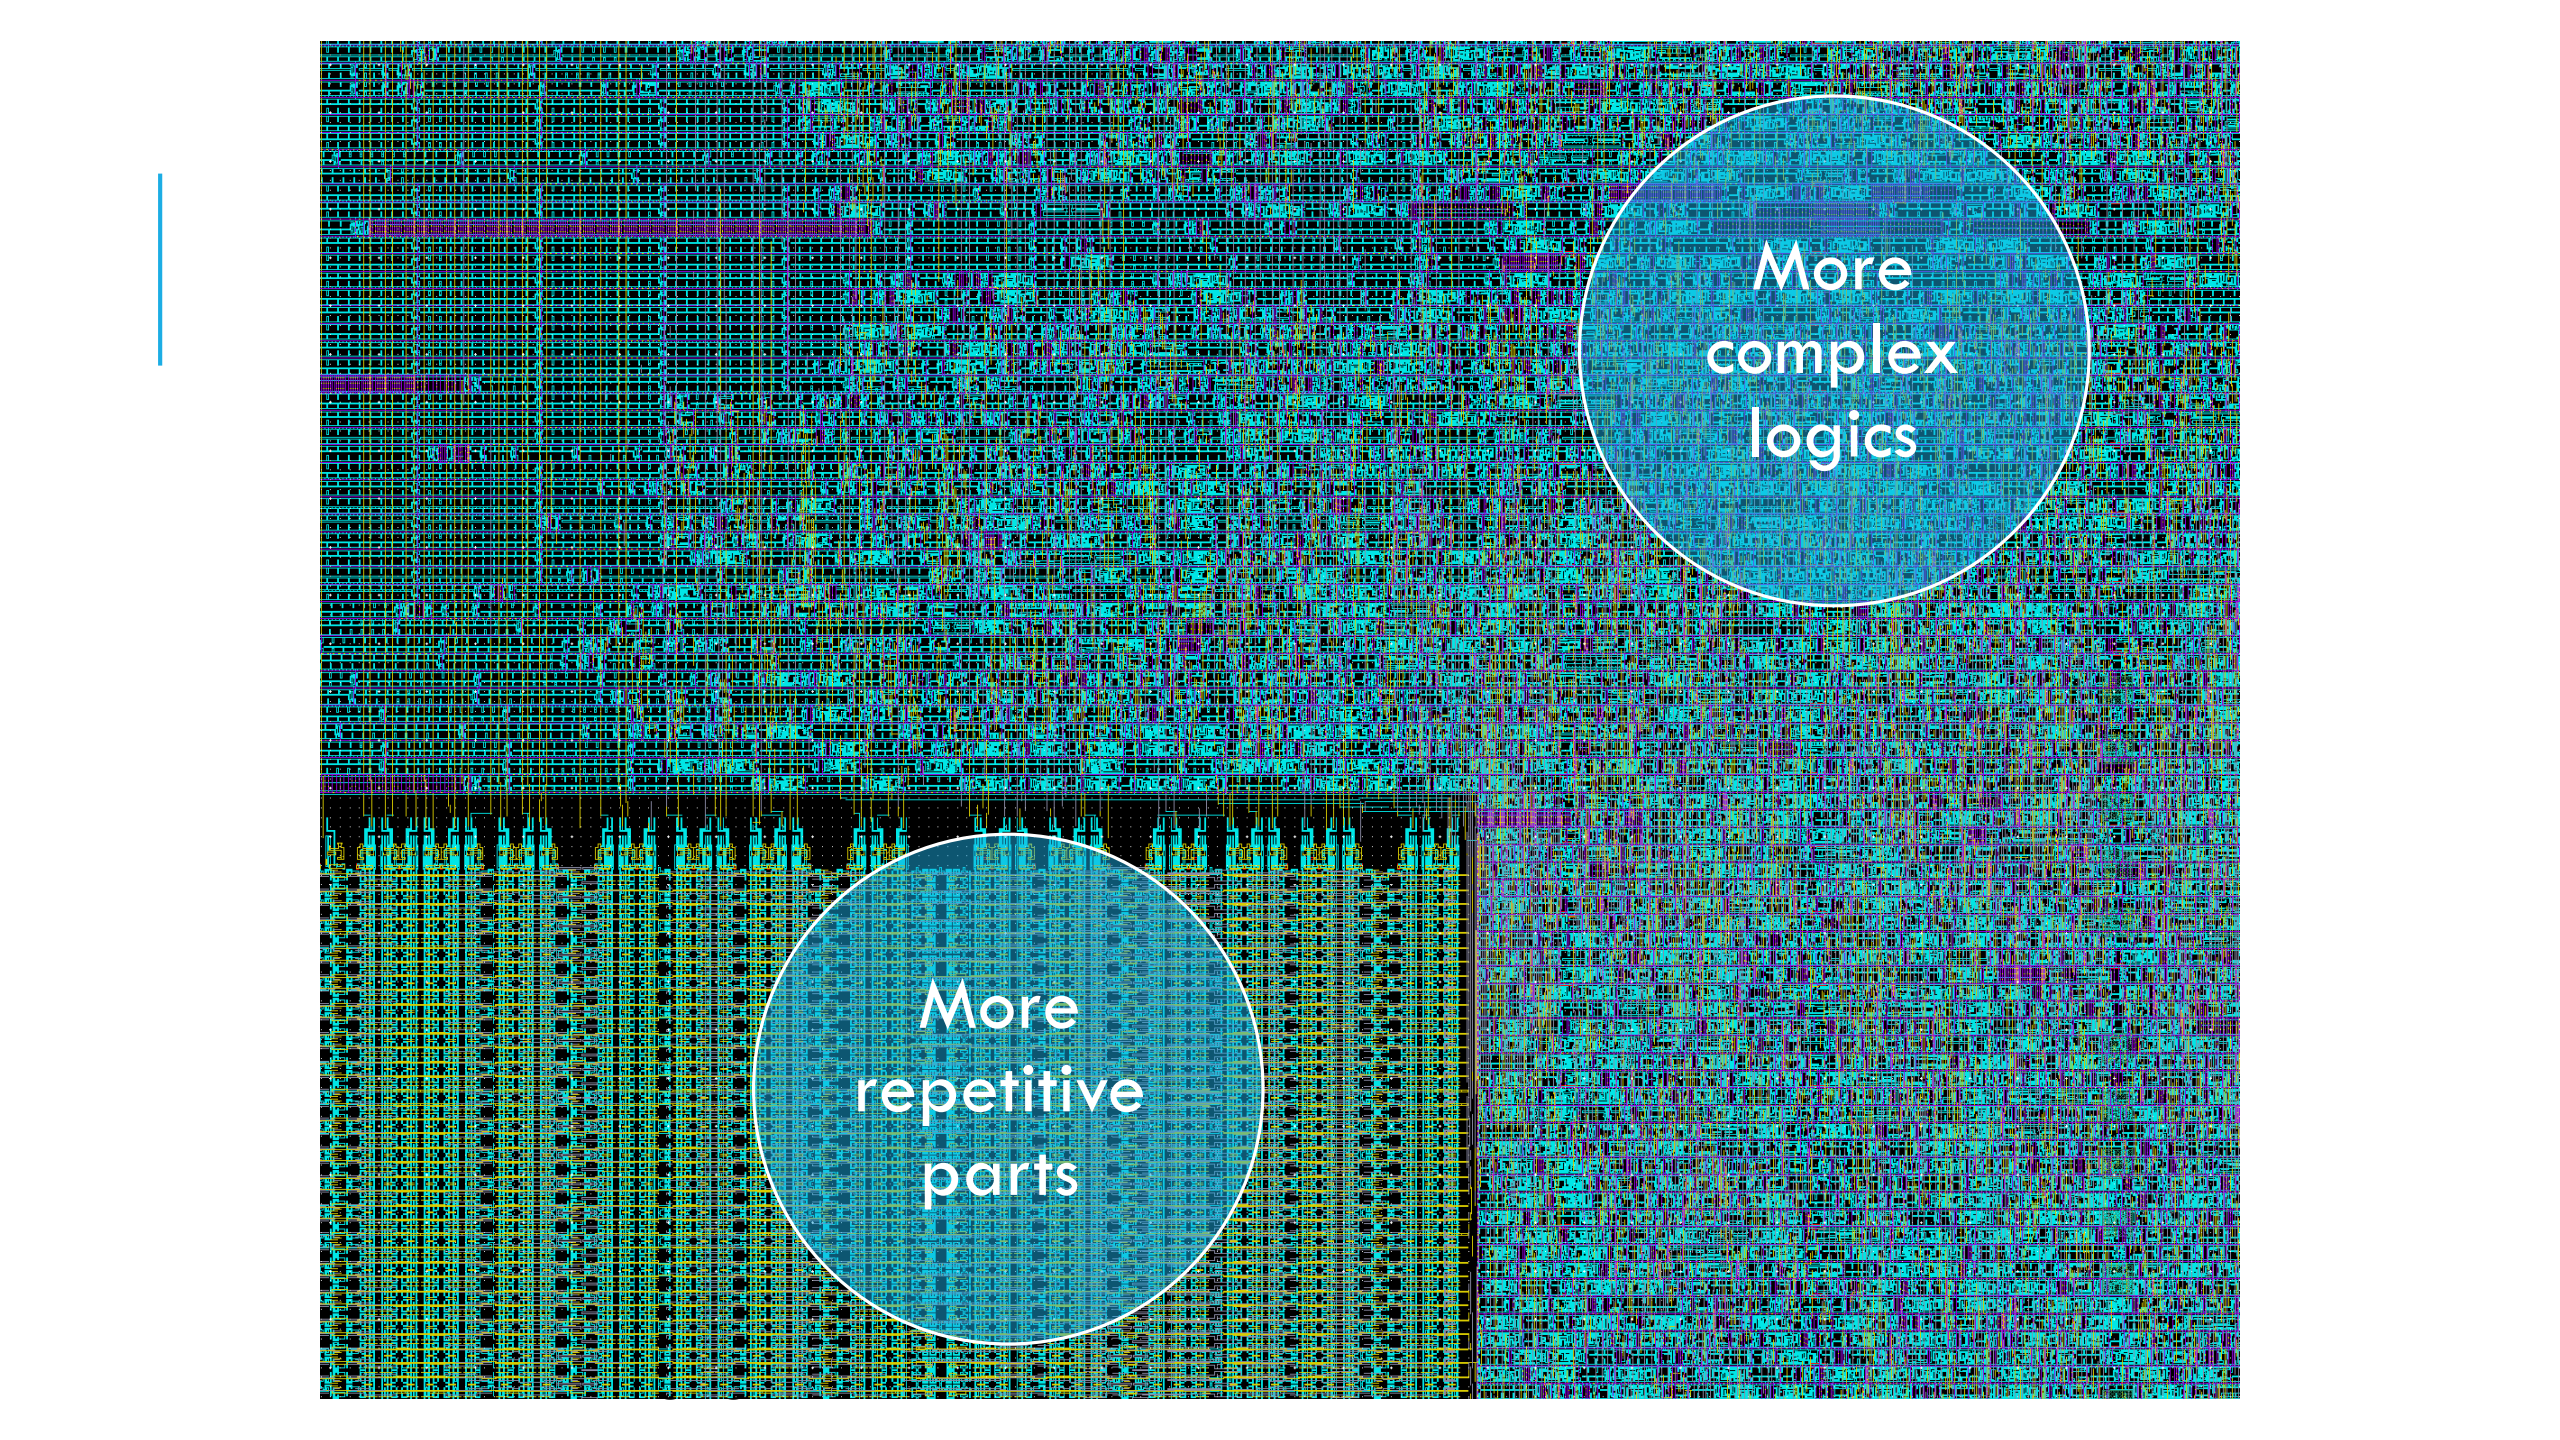




Standard  
cells



Full custom



More  
complex  
logics

More  
repetitive  
parts

# AM09 COMPLEXITY

**AM06 and AM09 are two of the most complex chips designed within CERN collaboration**

**Comparison rate:**

AM09 about 0.2 zeta comparisons per second per chip


AM06 about 15 exa comparisons per second per chip

# AM04: THE FIRST PROTOTYPE FOR FTK

The design of the first prototype of the Processing Unit for the FTK processor had to face the most challenging aspects of this technology: a huge number of detector clusters (“hits”) distributed at high rate with large fan-out to all patterns (10 million patterns will be located on 128 chips placed on a single board) and the large number of roads collected and sent back to the FTK post-pattern-recognition functions.

The network of high speed serial links used to solve the data distribution problem has been experimentally verified.

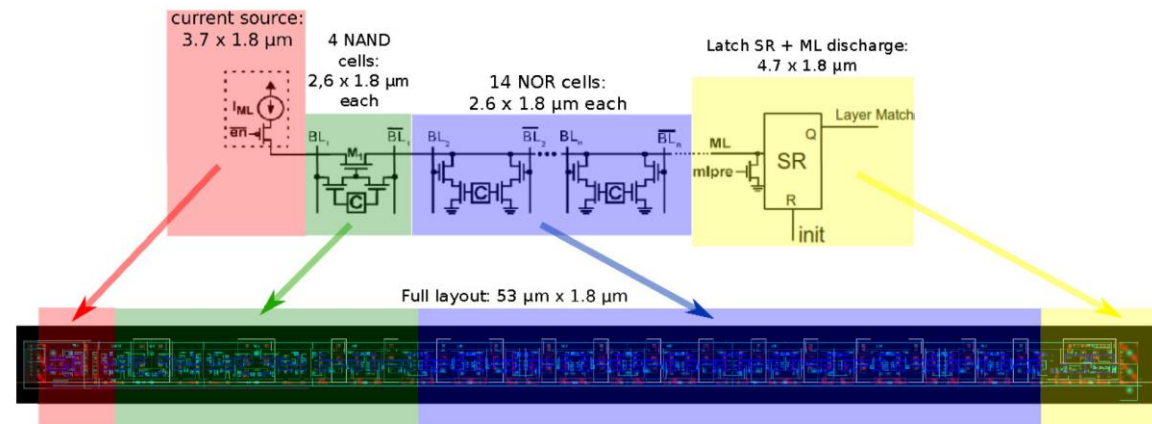
The AMchip04 prototypes were successfully tested, and their performance and current consumption is fully compatible with the AMBFTK. However, further reduction in current consumption will be mandatory in the design of the final AM chip.




(a)

address	data					
	meas.	sim.				
C8600F98	1D00	<= 1D00	0 <= 0	0 <= 0	1 <= X	
C8602530	1D40	<= 1D40	0 <= 0	0 <= 0	1 <= X	
C8603AC8	1D80	<= 1D80	0 <= 0	0 <= 0	1 <= X	
C8601060	1DC0	<= 1DC0	0 <= 0	0 <= 0	1 <= X	
C86025F8	1E00	<= 1E00	0 <= 0	0 <= 0	1 <= X	
C8603B90	1E40	<= 1E40	0 <= 0	0 <= 0	1 <= X	
C8601128	1E80	<= 1E80	0 <= 0	0 <= 0	1 <= X	
C86026C0	1EC0	<= 1EC0	0 <= 0	0 <= 0	1 <= X	
C8603C58	1F00	<= 1F00	0 <= 0	0 <= 0	1 <= X	
C86011F0	1F40	<= 1F40	0 <= 0	0 <= 0	1 <= X	
C8602788	1F80	<= 1F80	0 <= 0	0 <= 0	1 <= X	
C8603D20	1FC0	<= 1FC0	0 <= 0	0 <= 0	1 <= X	
C86006BC	3FC0	<= 3FC0	0 <= 0	1 <= 1	1 <= X	

(b)




# THE XORAM CELL



Based on the previous 65 nm XORAM cell

Based on the XOR boolean function, instead of the NAND and NOR functions

Is made of a 6T SRAM cell connected to a 6T-XOR gate



# AM05: THE LAST PROTOTYPE FOR FTK

and tested. In particular, tests performed on the mini-AM05 demonstrated the correct operation of the new XORAM cell and excellent performance of serial links at 2 Gbit/s. The current consumption was measured in different modes. As a significant fraction of the power dissipation is due to the input data distribution inside the chip, board level and crate level consumption are still a concern. For this reason, the AM05 was completely redesigned at layout level, to improve power performance. AM05 prototypes are now under test.

High speed serial links at 2 Gbit/s have also been successfully tested on the mini-LAMB.


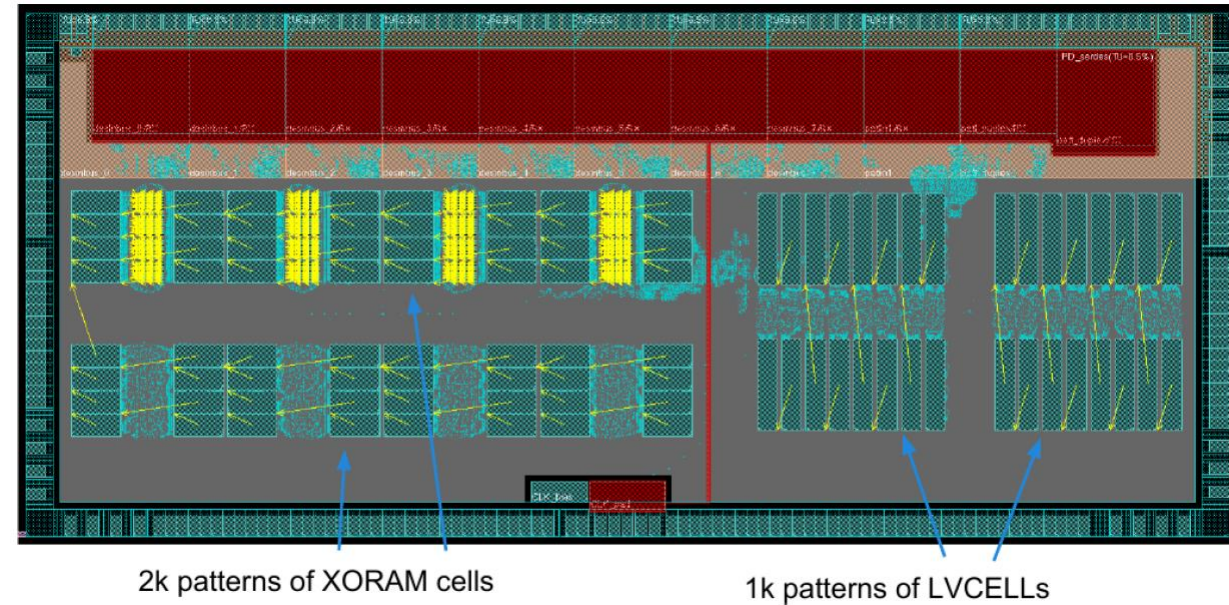


Figure 6. Layout of the XORAM block in a 65 nm CMOS technology.





# AM06: THE INDUSTRIALIZED CHIP FOR FTK

## 5 Conclusion

The AM06 has been successfully designed and fabricated. The prototypes are working and no redesign is needed. Tests on the first production batch show a high yield.

The AM06 current consumption exhibits peaks when the chip performs parallel comparisons, and special care is needed in package and board design to reduce the supply voltage ripple. To cope with this issue, the LAMB design is being improved. Test results indicate that the optimal voltage supply for the AM06 chip is 1.1 V, and this value is being used in FTK.

## Acknowledgments


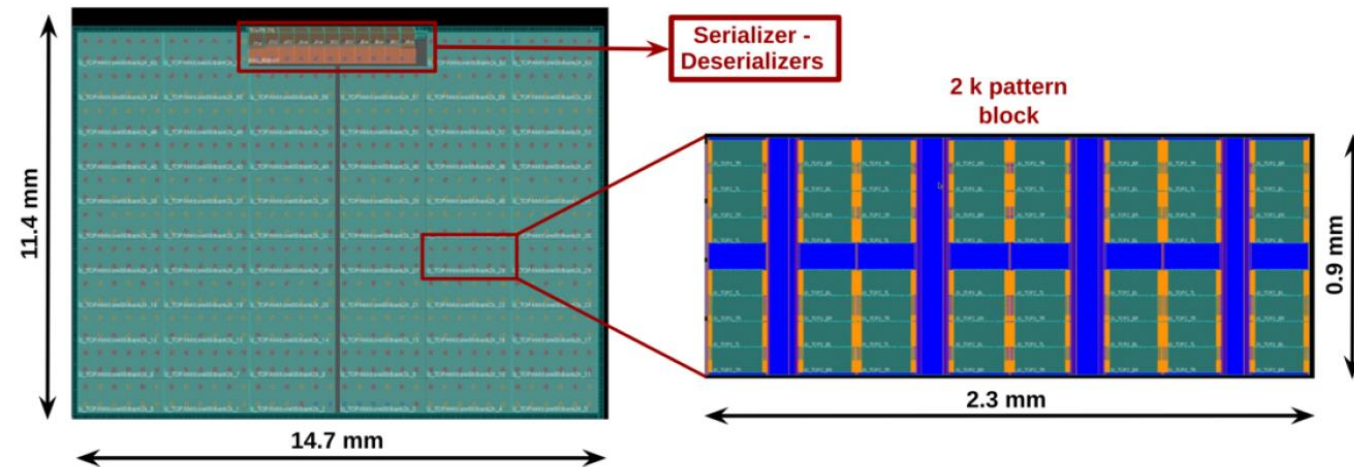


Figure 6. Measured eye diagram on a 2 Gbit/s serial link.

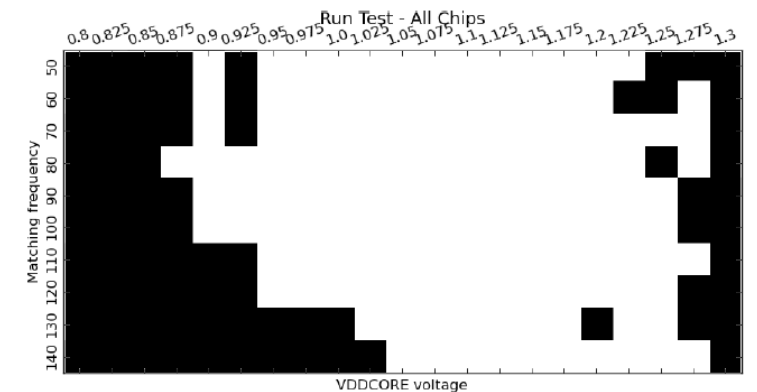


Figure 8. Shmoo plot showing the chip functionality (white region) as a function of the  $V_{DD,CORE}$  voltage and of the frequency.

# NEW OPTIMIZED CELLS

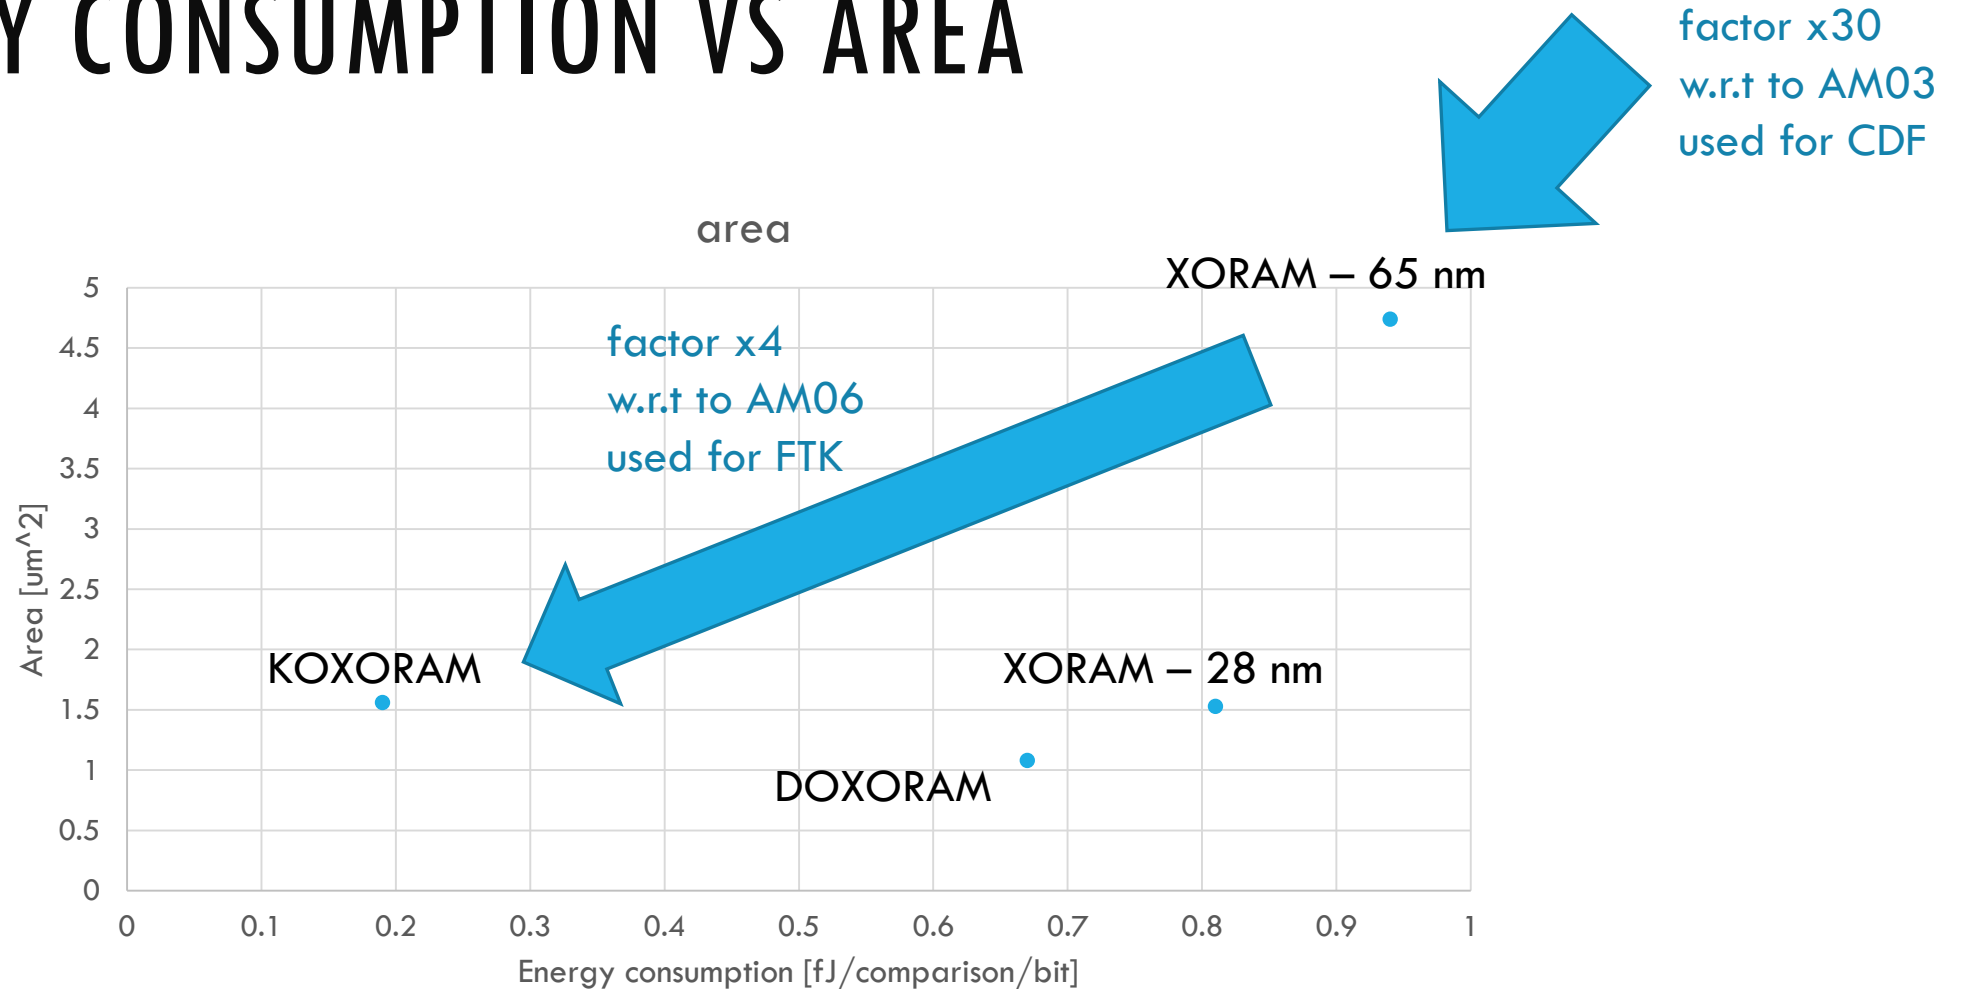
With similar power save methods we designed two new cell tech:

DOXORAM


KOXORAM

**Italian Patent:** A. Annovi, L. Frontini, V. Liberali, A. Stabile, "MEMORIA CAM", UA2016A005430

# FULL CUSTOM CELL: ENERGY CONSUMPTION VS AREA



# AM07: THE PROTOTYPE FOR INTERDISCIPLINARY APPLICATION



Compared to the AM06 chip, the AM07 exhibits a power consumption reduced by a factor of 1.7 and a density increased by a factor of 2.9. In the AM07 design, we have noticed that the automatically designed logic circuitry does not scale as predicted by the Moore's law. For this reason, in the future we plan to re-design the Quorum circuit with a full-custom approach, to reduce the overall area. With this approach, we aim to improve also the power consumption and the maximum




Fig. 8. Photograph of the characterization setup at CERN.

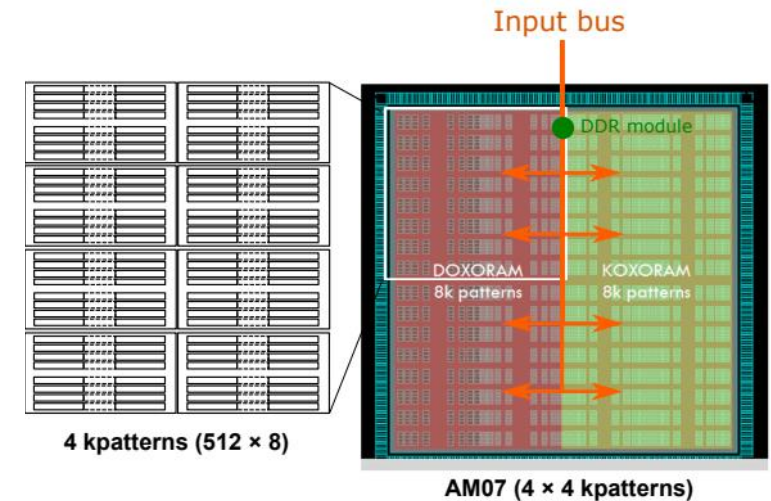




Fig. 2. Arrangement of the 16 kpatterns cell array.

# COMPUTER VISION FOR SMART CAMERAS AND MEDICAL IMAGING APPLICATIONS

**Smart cameras** capture high-level description of a scene and perform **real-time extraction of meaningful information**

- Current compression algorithms: **few seconds** are required
- **For safety-critical applications (e.g., transports, or personnel tracking in a dangerous environment), latency could lead to serious problems.**

Del Viva et al algorithm<sup>1</sup> studied how to reproduce initial stage of the brain visual processing: find contours



<sup>1</sup>M. Del Viva, G. Punzi, and D. Benedetti. *Information and Perception of Meaningful Patterns*. *PloS one* 8.7 (2013): e69154.

# FUTURE DEVELOPMENTS

## Medical application

### Automated medical diagnosis:

- Huge amount of image data
  - time-varying images
  - very accurate resolution

### Real-time applications for MRI fingerprint in collaboration with the INFN-Pisa research group

- Guido Buonincontri's CSN5 funded project in 2015 and PUMA project

### IMPART-based system performance:

- **Human exome:** 1.5 % subset of the human genome (25 million nucleotide pairs)
- **Nucleotide encoding:** FASTA format (at least 4 bits are needed)
- Whole exome alignment with **this device:** ~ 4 s

### Commercial machines performance:

- Bowtie based machines: **1 CPU hour**

**Speed improvement factor is about 900x**

**DNA application**

# AM08 AND AM09

- **AM08 prototype:** small area MPW prototype to test all the full custom features, the VHDL logic and the I/O. This chip must be fully functional with smaller memory area than the final ASIC;
- **AM09pre pre-production:** full area ASIC to be fabricated with a full-mask set pilot run. Production corner wafers will be created;
- **AM09 production:** full area ASIC with refinements for the mass production.

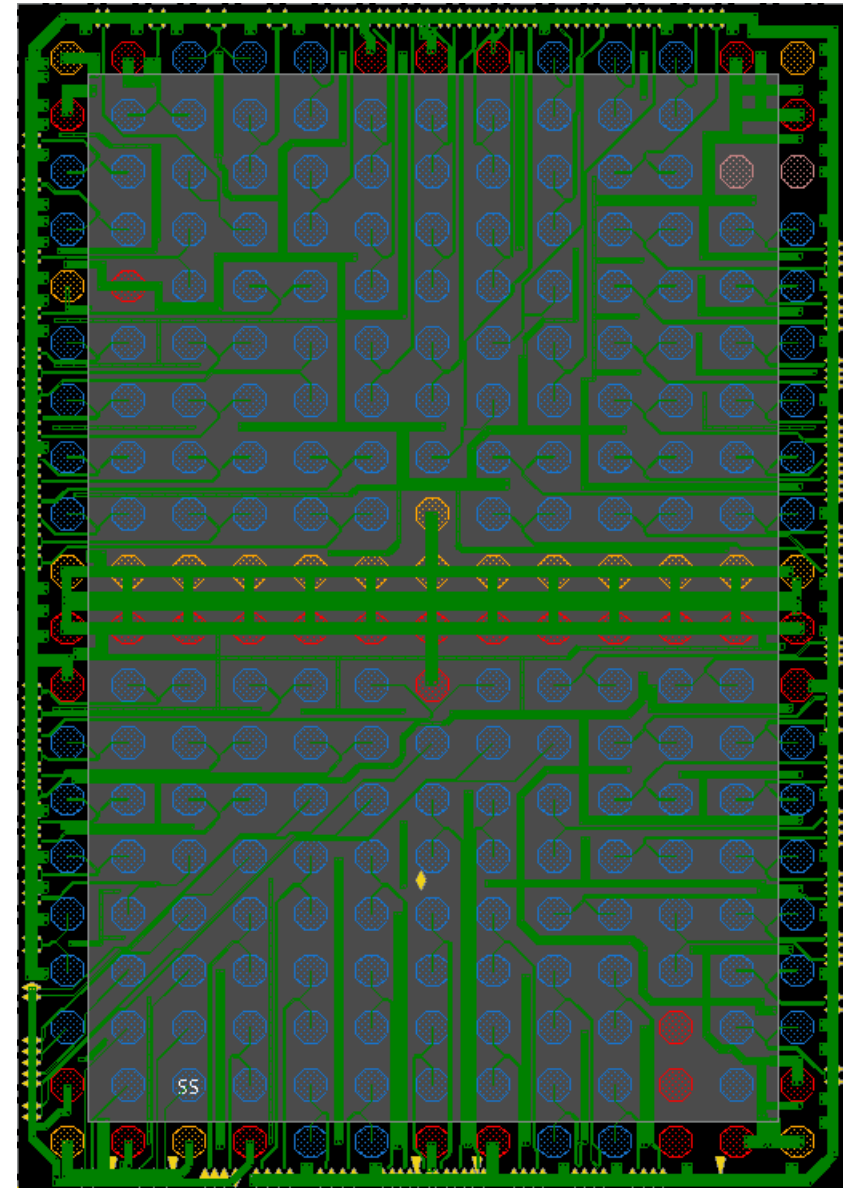
AM08 is a **12 kpatterns**<sup>1</sup> and AM09 is a **3×128 kpatterns** low power CMOS associative memories organized as 3 (AM08) or 96 (AM09) **4 kpatterns** cores respectively.

It is fabricated using very high performance, high reliability **CMOS technology at 28 nm** (HPC – 10 metal layers + RDL).


The AM08 and AM09 devices are well **designed for high energy applications**, and particularly well suited for ATLAS trigger applications.

The AM08 and AM09 operate with a nominal power supply of 1.0V and all **data inputs and outputs are fully LVDS18** compatible.

The LVDS18 I/O have been designed to works at **1 Gbps**.



# CONCLUSION AND SOCIAL IMPACT




This innovative systems ameliorate the efficiency of many **HEP trigger system**



**Smart cameras** with this system could be installed in remote environments (forests or mountains)



**DNA sequencing** could benefit from the project



The system could be also used to **better filtering the fingerprint magnetic resonance images (MRI)**

Several applications could benefit from the project outcome.





# BACKUP

A. Stabile

# FAST TRACKER (FTK)

Il trigger sul momento trasverso mancante è la chiave di numerose misure di Fisica: gli articoli riportano rispettivamente la prima evidenza del decadimento  $H \rightarrow b\bar{b}$  e la ricerca di nuovi fenomeni fisici

- FasTrack (FTK) è il processore usato per il run2 di LHC e permetterà di migliorare i trigger calorimetrici usati finora aggiungendovi l'identificazione dei vertici primari per sopprimere l'effetto del pile-up.

La selezione di muoni isolati è critica per la ricerca di nuova fisica, come ad esempio in SUSY dove i muoni ad alto momento trasverso o i decadimenti ad alta massa ( $Z' \rightarrow \mu\mu$ ). O meglio ancora per lo studio di processi del modello standard come ad esempio  $W \rightarrow \mu\nu$  o  $Z \rightarrow \mu\mu$ .

Con le informazioni di tracciamento disponibili in anticipo da FTK, si può calcolare l'isolamento usando solamente le tracce che puntano allo  $z_0$  delle tracce dei muoni.

- Questo isolamento basato su tracce rimuove ogni necessità di utilizzo di dati energetici mantenendo un'alta efficienza per i muoni singoli/isolati in un ambiente ad alto pileup.

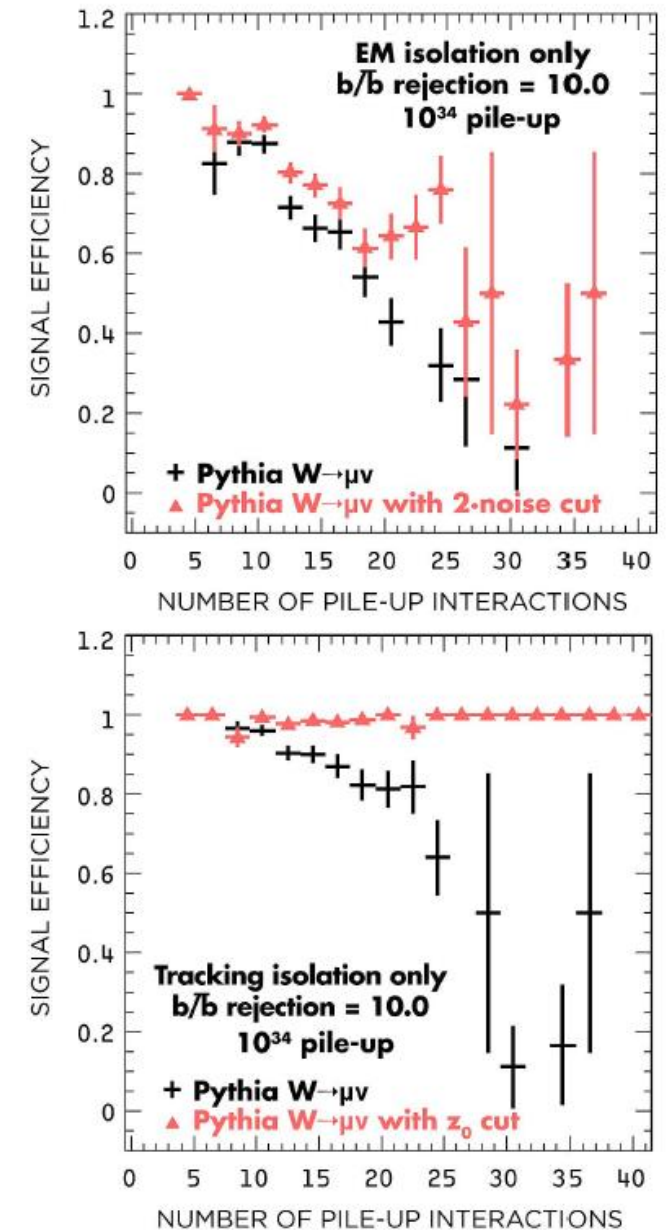


Fig. 13. Isolated muon efficiency using the EM calorimeter with two cell energy thresholds (top) or with (light) a  $\Delta z_0$  cut as a function of the number of pile-up interactions in the event (bottom). The isolation cut is selected to provide a  $b\bar{b}$  rejection factor of 10.

# THE FTK SYSTEM

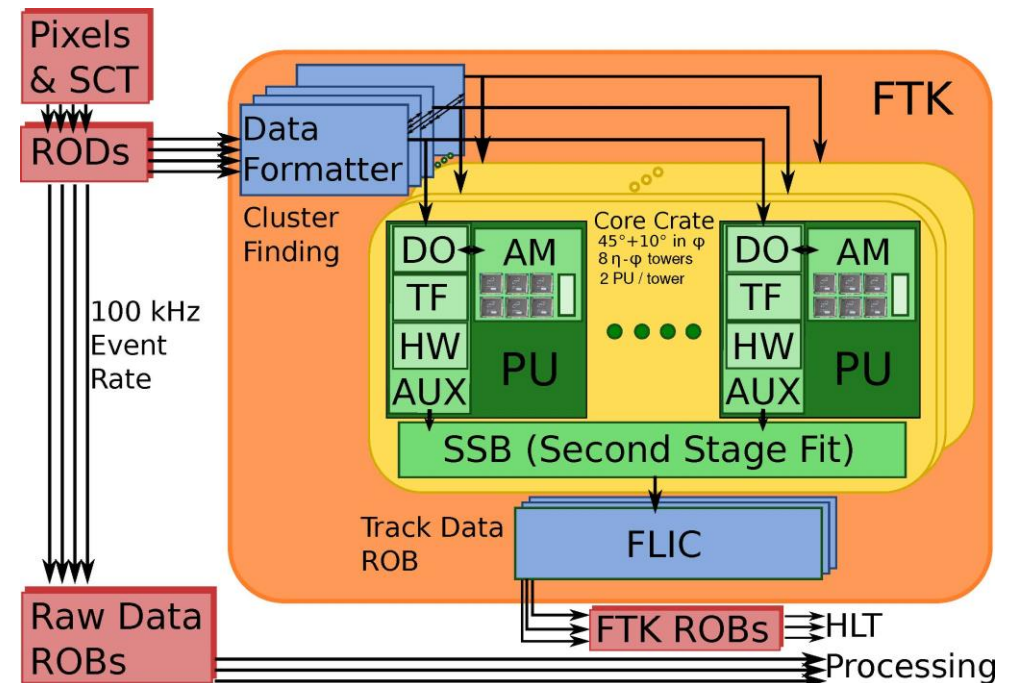
The whole FastTraKer (FTK) system stores **one billion ( $10^9$ ) patterns (31 Ebit/s)**

- **8 Mpatterns** per board (128 boards)
- **128 kpatterns** per chip (64 AM chips / board)
- A pattern is composed by  $18 \text{ bits} \times 8 \text{ words}$

## Major concerns:

- high pattern density
  - ▶ **large silicon area**
- I/O signal congestion at board level (**solution: 2 Gbit/s serial links**)
- Maximum power limited by cooling (because we are fitting 8192 AMchips in 8 VME crates): **250 W per AM board**

<sup>1</sup>A. Andreani et al., "The AMchip04 and the processing unit prototype for the FastTracker," *IOP J. Instr.* 7 (2012) C08007

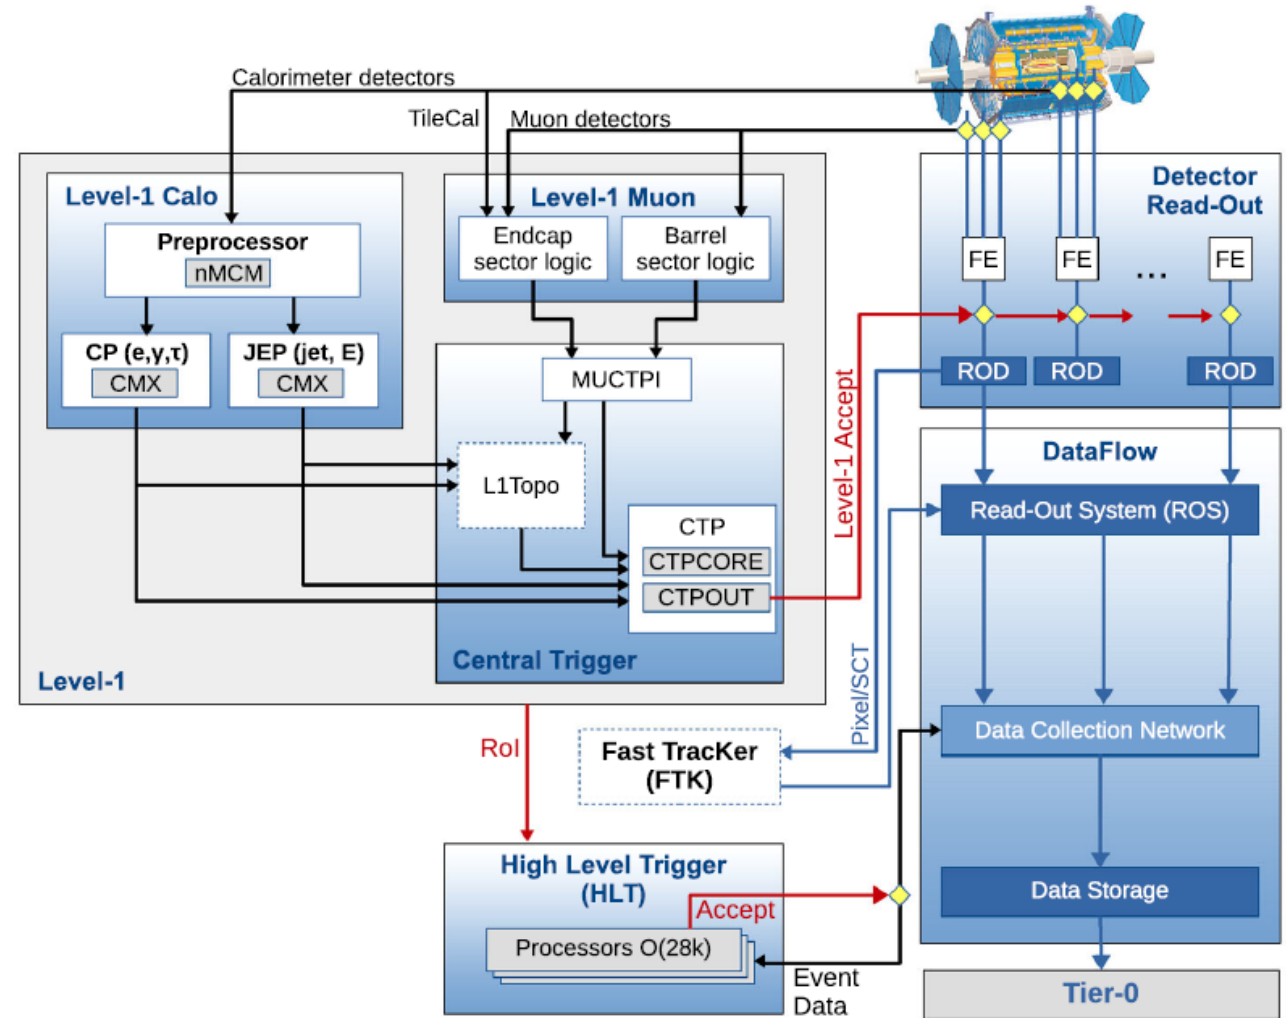


# Performance of the ATLAS trigger system in 2015

ATLAS Collaboration\*  
CERN, 1211 Geneva 23, Switzerland

Questo paper racconta di come FTK migliora le performance di trigger

A new Fast TracKer (FTK) system [18] will provide global ID track reconstruction at the L1 trigger rate using lookup tables stored in custom associative memory chips for the pattern recognition. Instead of a computationally intensive helix fit, the FPGA-based track fitter performs a fast linear fit and the tracks are made available to the HLT. This system will allow the use of tracks at much higher event rates in the HLT than is currently affordable using CPU systems. This system is currently being installed and expected to be fully commissioned during 2017.



# HARDWARE FOR THE TRACK TRIGGER (HTT)

Trigger Selection	Object Multiplicity	Region Size (% detector)	Regional Tracking (rHTT) Need Per Event			
			% $\eta - \phi$ coverage	Detector Data Fraction		
				Pixel Layer 5	Strip Layer 1	Strip Layer 4
isolated single $e$	1	$0.2 \times 0.2$	0.13%	1.0%	1.3%	0.4%
isolated single $\mu$	Not Used					
single $\gamma$	Not Used					
forward $e$	1	$0.2 \times 0.2$	0.13%	1.0%	1.3%	0.4%
di- $\gamma$	Not Used					
di- $e$	2	$0.2 \times 0.2$	0.25%	2.0%	2.6%	0.8%
di- $\mu$	2	$0.2 \times 0.2$	0.25%	2.0%	2.6%	0.8%
$e - \mu$	2	$0.2 \times 0.2$	0.25%	2.0%	2.6%	0.8%
single $\tau$	Not Used					
di- $\tau$	2	$0.2 \times 0.2$	0.25%	2.0%	2.6%	0.8%
single jet	Not Used					
large- $R$ jet	Not Used					
four-jet	5	$0.8 \times 0.8$	10.2%	23.0%	25.5%	14.5%
$H_T$	5	$0.8 \times 0.8$	10.2%	23.0%	25.5%	14.5%
$E_T^{\text{miss}}$	3	$0.8 \times 0.8$	5.7%	12.9%	14.3%	8.1%
VBF inclusive	2	$0.8 \times 0.8$	4.1%	9.2%	10.2%	5.8%
Supporting Trigs	10% of total rate		0.2%	0.5%	0.6%	0.3%
Averages per event weighted by rates			2.3%	6.0%	6.8%	3.5%

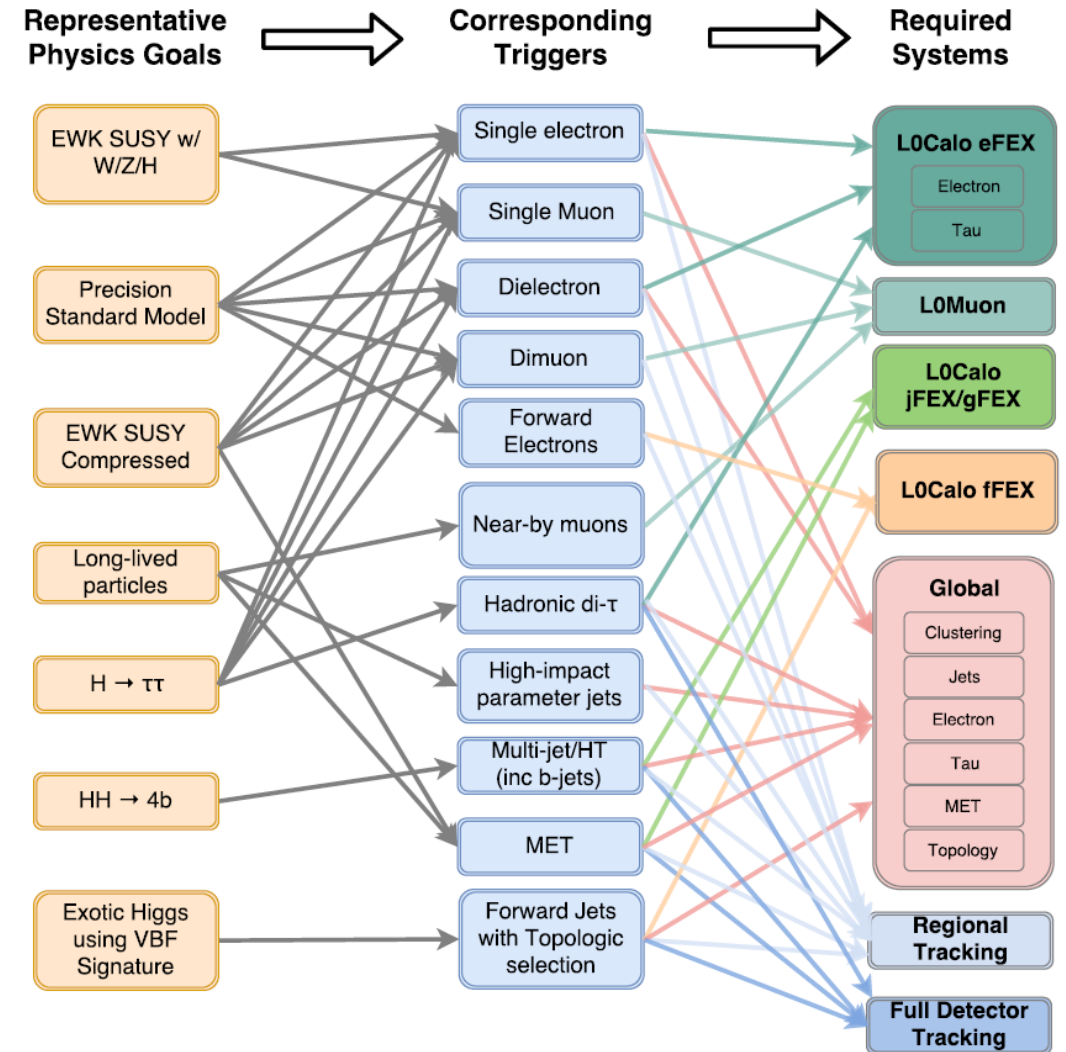


Figure 2.1: Schematic summary of the flow from the representative set of physics goals described in this section (left column) to the hardware systems (right column) needed to achieve these goals. The middle column lists the corresponding triggers required.

# Evidence for the $H \rightarrow b\bar{b}$ decay with the ATLAS detector

*from the TDAQ TDR of 2019*

Without the [HTT](#), the nominal tracking requirements would require  $\approx 10$  times more CPUs than the baseline design with [HTT](#) (see Section 12.4). If less tracking is available, then the thresholds for the objects requiring tracking would need to be raised to reduce the rates. For example to reduce the tracking needs for the  $b$ -tagged four-jet the Level-0 threshold would have to be raised to an effective offline threshold of 85 GeV, which would raise the limit on the search for  $HH \rightarrow 4b$  by approximately 50%.

All cases show substantial degradation with increased minimum jet  $p_T$  requirements. Table 6.7 shows the impact on this analysis for three scenarios with reduced upgrades: a) no [Global Trigger](#), b) no [HTT](#), and c) no-upgrade as previously defined to mean the Phase-I system with an output rate for 100 kHz. Without the [Global Trigger](#) system, the threshold for the lowest  $p_T$  jet in a four-jet trigger would have to be raised to approximately 75 GeV instead of 65 GeV to maintain the same Level-0 output rate. In addition, the loss of acceptance for near-by jets would cause a further efficiency loss of 10-15%. The impact of that reduction on the  $HH \rightarrow 4b$  analysis would be to reduce the cross-section sensitivity by  $\approx 25\%$ . [Without the HTT, the tracking would be CPU-limited, and the Level-0 trigger rate would need to be reduced by a factor of  \$\approx 10\times\$  to allow CPUs in the Event Filter to do the required tracking. Such a rate reduction corresponds to a 85 GeV threshold, and an  \$\approx 45\%\$  loss of sensitivity.](#) With no upgrade at all the loss is greater than  $\approx 65\%$ . A scenario with

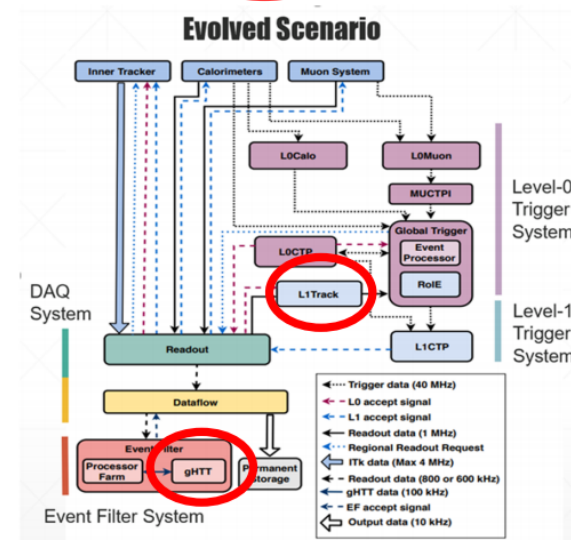
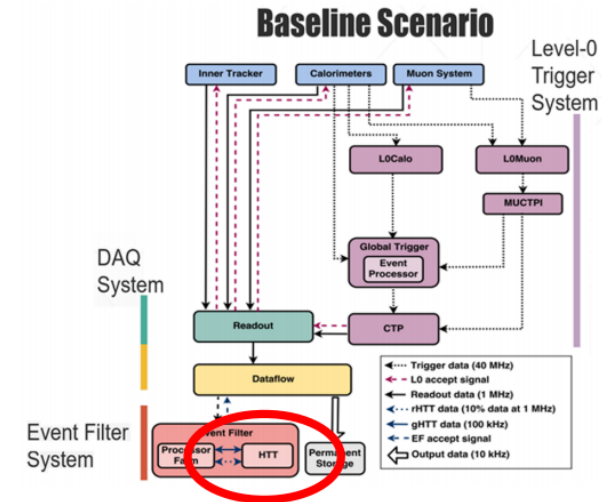
# HTT IN TDAQ

## Baseline scenario

- HTT performs **global tracking (gHTT)** at 100 kHz and **regional tracking (rHTT)** at 1 MHz on up to 10% of the ITk data
- Event-data is provided by the EF processing unit and tracks are returned to it (within 10 ms )

## Evolved scenario,

- HTT performs **global tracking (gHTT)** at 100 kHz on L1-accepted events, and **regional tracking (L1Track)** at up to 4 MHz in up to 10% of the ITk data
- The L1Track system processes event-data directly from ITk-FELIXs and tracks are sent to L1 Global (within 6  $\mu$ s)



3

# HTT ARCHITECTURE

Several reasons motivate this decision, which include considerable experience in the AM technology, the potential for short latency, a lower power budget and less demanding space requirements compared to other technologies, its cost effectiveness and the independence of its cost from the commodity computing market, availability of in-house expertise, and the capability to evolve the HTT system for use in the hardware-based Level-1 trigger, should ATLAS need to change to a dual L0/L1 trigger system

