

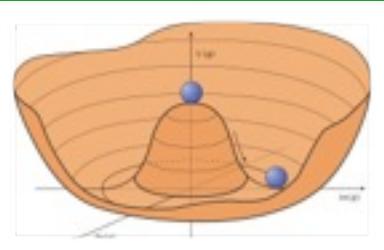
ATLAS challenging trigger in the High Luminosity LHC



ROYAL
HOLLOWAY
UNIVERSITY
OF LONDON

past, present & future

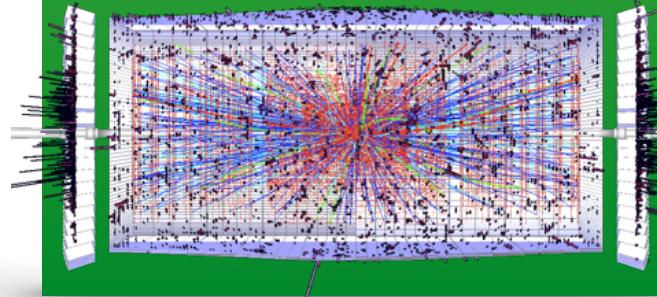
1964



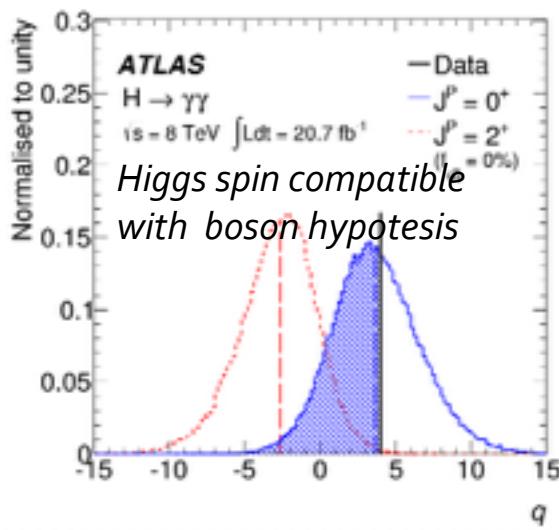
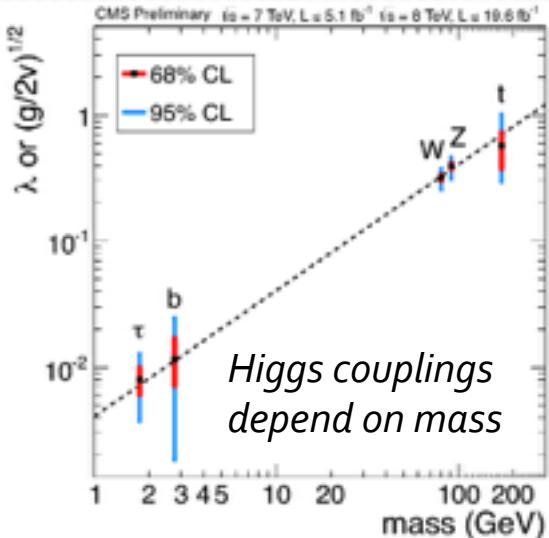
2013



~2030



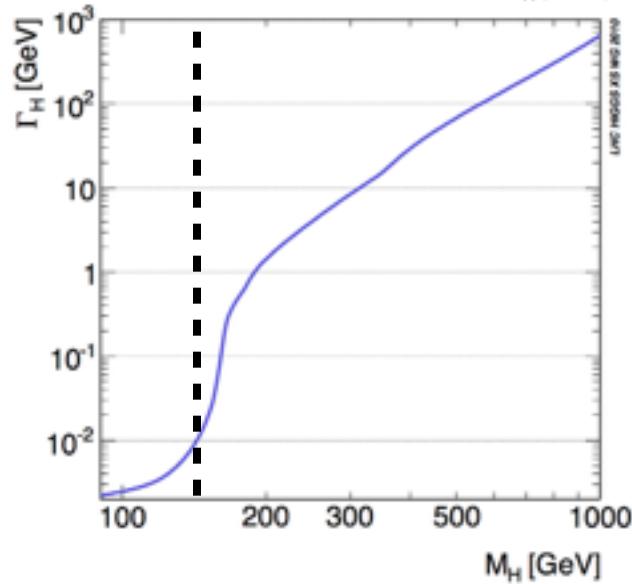
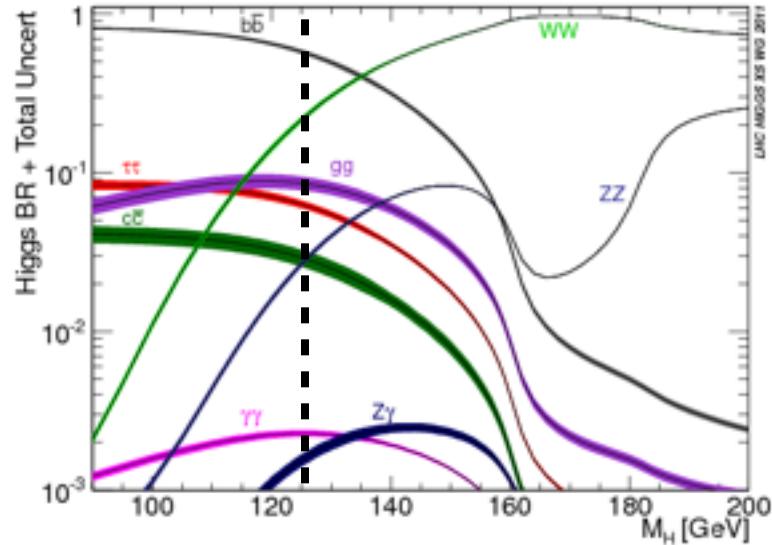
The Higgs discovery at 125 GeV



*ATLAS and CMS
verified two
fingerprints of the
Higgs Boson*

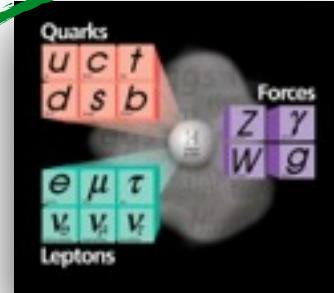
■ Higgs mass at 125 GeV

- Opens access to many Higgs decay channels, and then to the measurements of its coupling constants
- Pushes performance of our detectors since Higgs width is only few MeV



LHC for future frontiers of particle physics

- The existence of a Boson at low mass scale sets a limit to the validity of the SM (~**TeV scale?**) and its extension must answer still open questions
 - mass hierarchy, cosmological questions (DM, DE, inflation), gravity and CP violation



- The key answers are hidden in the **properties of the Higgs boson....**
 - spin, self-interaction, multiplicity,... to be studied through the interactions with other particles

Standard Model is completed!
We have no evidence of New Physics!

- Precise measurements and rare process discovery need at least **3000 fb⁻¹** of data
 - ATLAS and CMS collected 30 fb⁻¹ so far

Expected &
unexpected

- ...and in the deep exploration of the **TeV scale and rare processes**
 - investigating B-physics, top decays, gauge bosons scattering

- European Council: "CERN is the strong European focal point for particle physics in next 20 years"
 - https://cds.cern.ch/record/1551933/files/Strategy_Report_L_R.pdf?version=1

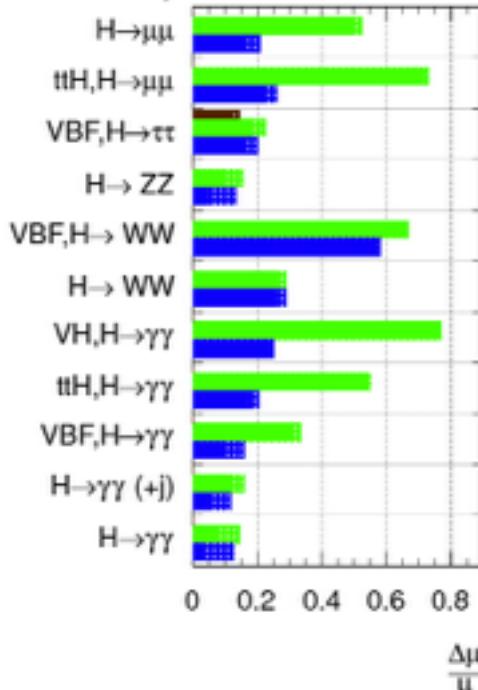
What can we do with the Higgs Factory?

Physics at a High-Luminosity LHC with ATLAS (<http://arxiv.org/abs/arXiv:1307.7292>)

Expected precision on signal strength
at 300 and 3000 fb^{-1}

ATLAS Simulation

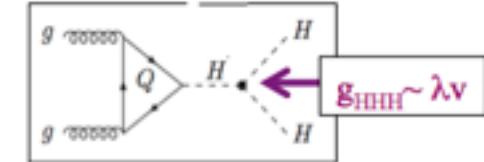
$\sqrt{s} = 14 \text{ TeV}$: $\int \text{Ldt} = 300 \text{ fb}^{-1}$; $\int \text{Ldt} = 3000 \text{ fb}^{-1}$
 $\int \text{Ldt} = 300 \text{ fb}^{-1}$ extrapolated from 7+8 TeV



- Measurement of Higgs couplings by
 - increased precision on already observed
 - access to rare ($H \rightarrow \mu\mu$, $ttH \rightarrow \tau\tau\gamma\gamma$)

3000 fb^{-1} !

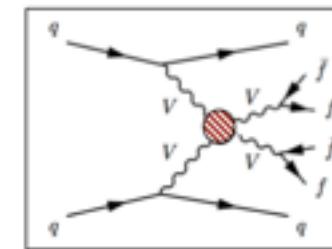
More than 3M Higgs events
for precise measurements
(\geq ILC/CLIC/TLEP)



- measure Higgs self-coupling (giving access to lambda)

Require:

- SM energy scale
- Lepton signatures
- Increased importance of tau and b-quark



- Forward jets give clear signature (possible extension of trackers to $|\eta| < 4$)

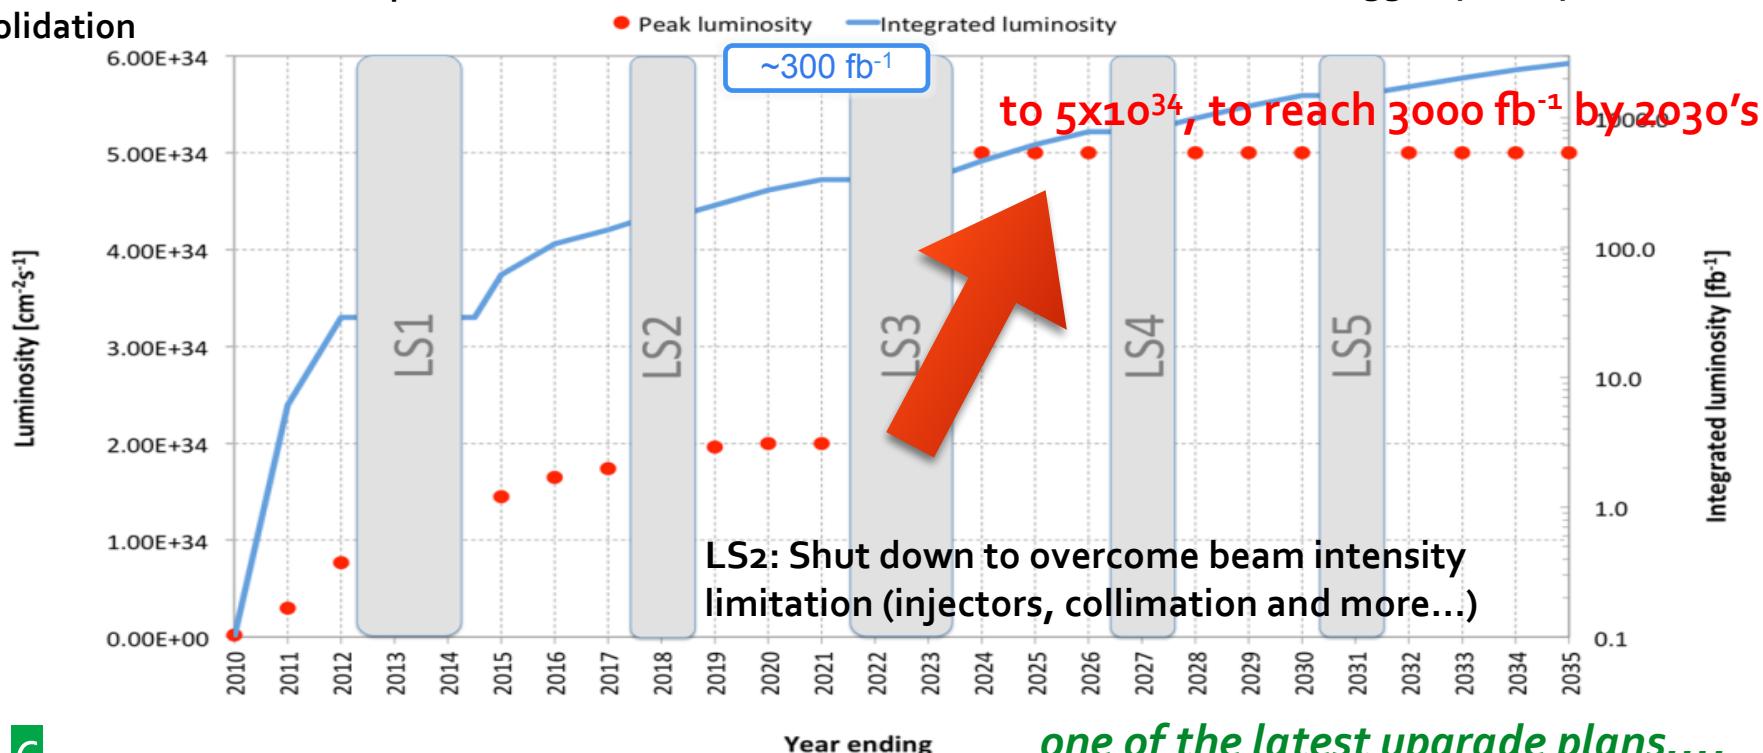
- verify that the Higgs boson fixes the SM problems with W/Z scattering at high energy

LHC becoming impressively luminous: HL-LHC

- LHC plans for next 10 years are approved (LS1 and LS2). Next: **HL-LHC starts in 2024**
- New project to upgrade large part of the accelerator complex
 - Linac4, Booster, SPS, Interaction regions
- Collect $300 \text{ fb}^{-1}/\text{year}$, peak luminosity increases by factor 5 w.r.t. the design value (+ Luminosity leveling)

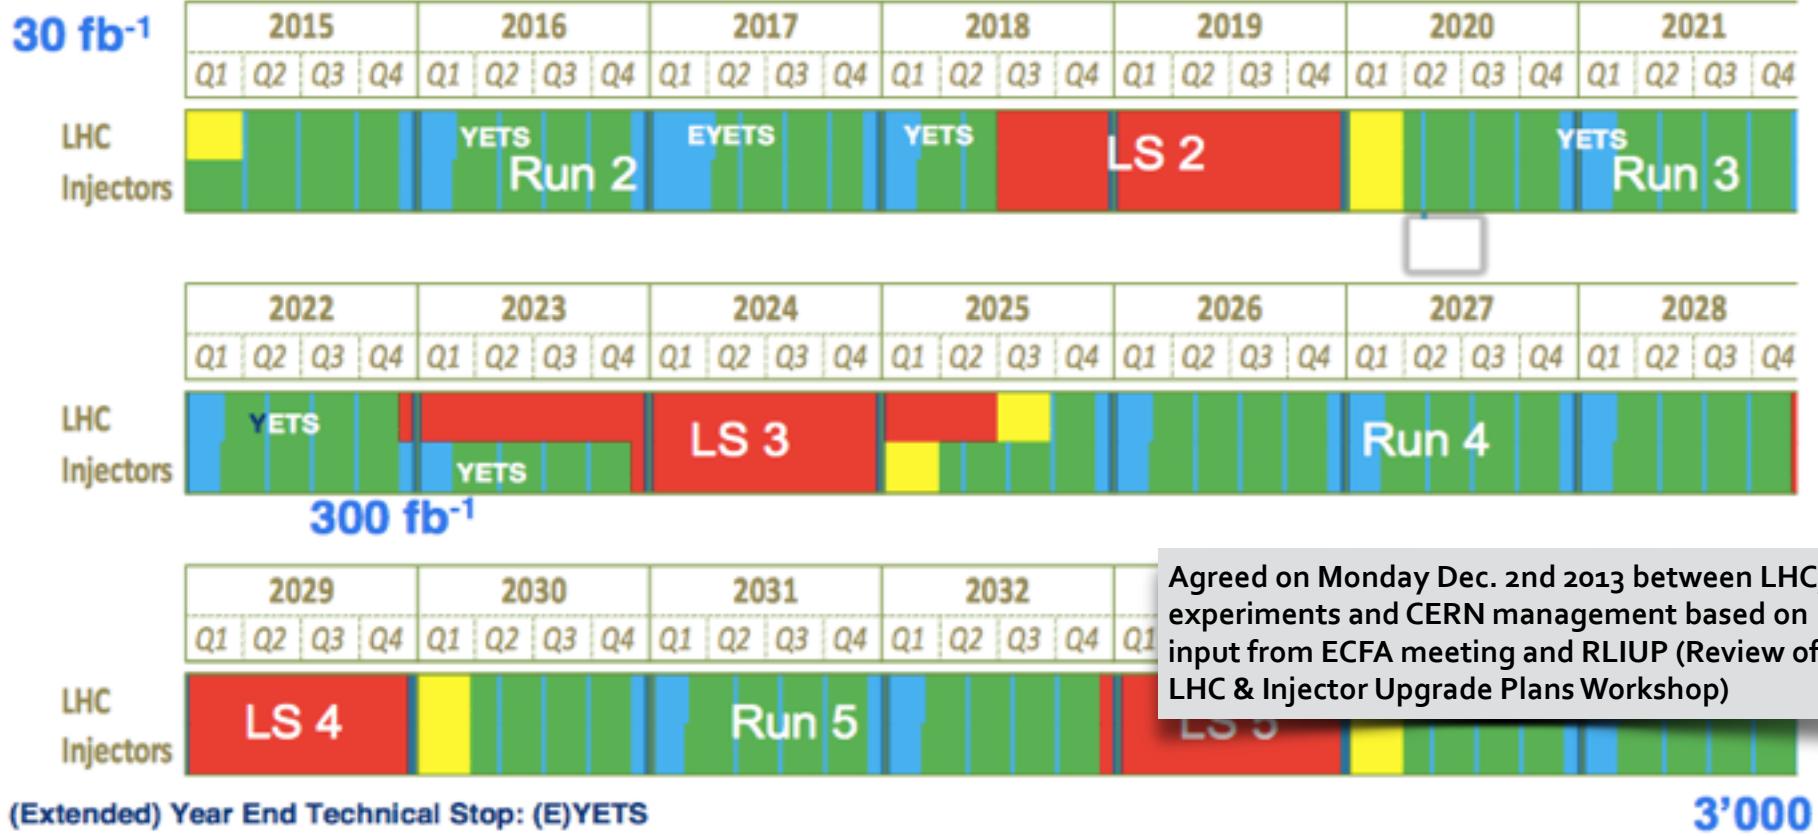
LS1: repair interconnects to overcome energy limitation (LHC incident of Sept. 2008) and consolidation

LS3: full upgrade (new magnet technology for the IR, new bigger quadrupoles)



LHC schedule beyond LS1

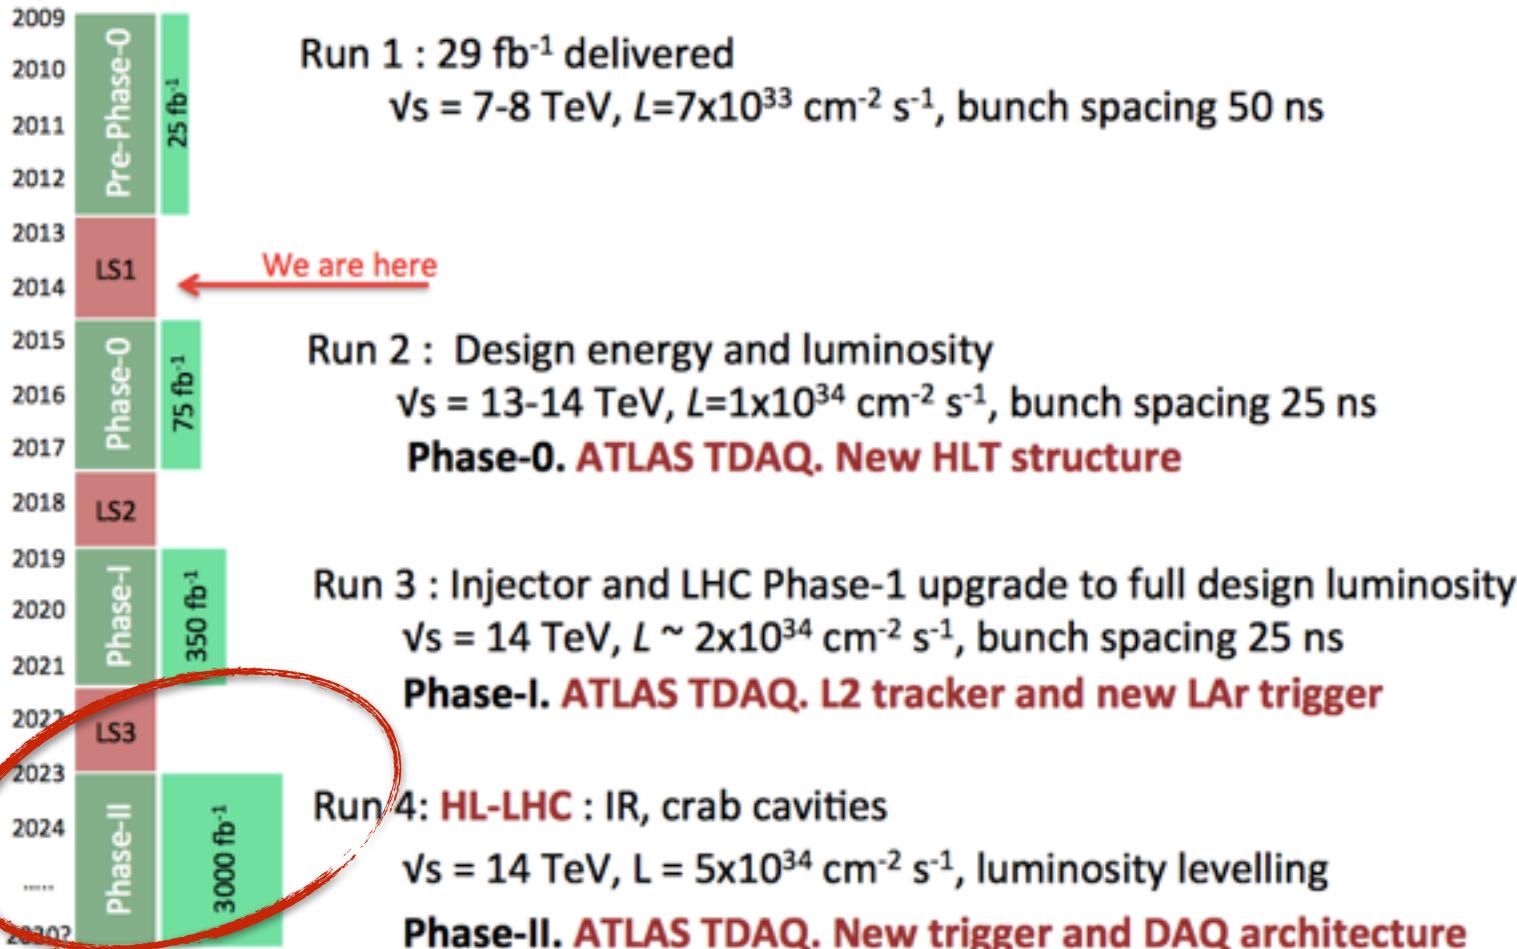
LS2 starting in **2018 (July)** => **18 months + 3 months BC**
 LS3 LHC: starting in **2023** => **30 months + 3 months BC**
 Injectors: in **2024** => **13 months + 3 months BC**



The CERN Roadmap
 Frédéric Bordry
 Future Circular Collider Kick-off Meeting – Geneva . 12th February 2014

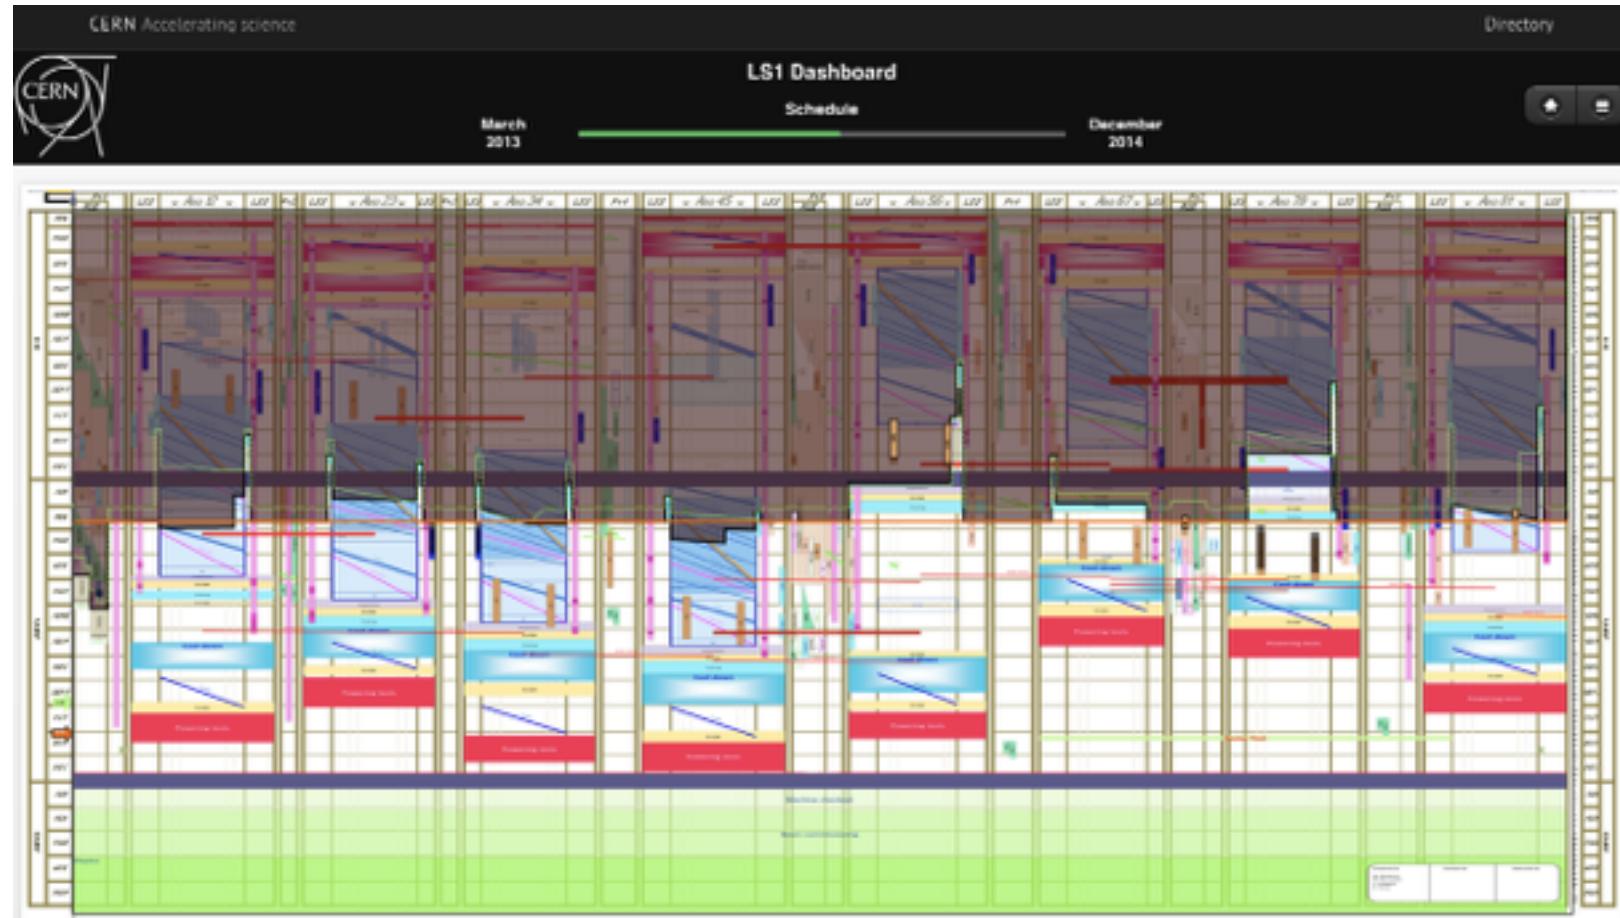
LHC schedule approved by CERN management and LHC experiments spokespersons and technical coordinators (December 2013)

ATLAS upgrade steps (focused on Trigger-DAQ)



Ls1 is on schedule!

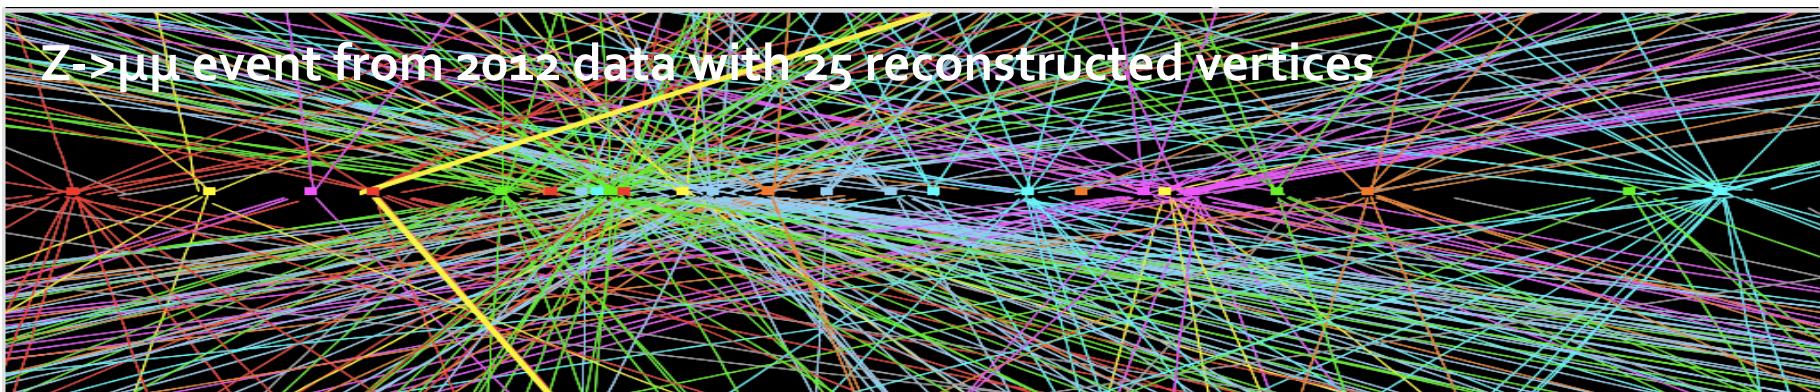
look at <http://cern.ch/ls1dashboard>



Most recent overall summary was Aix les Bain meeting (<http://Indico.cern.ch/event/252045>)

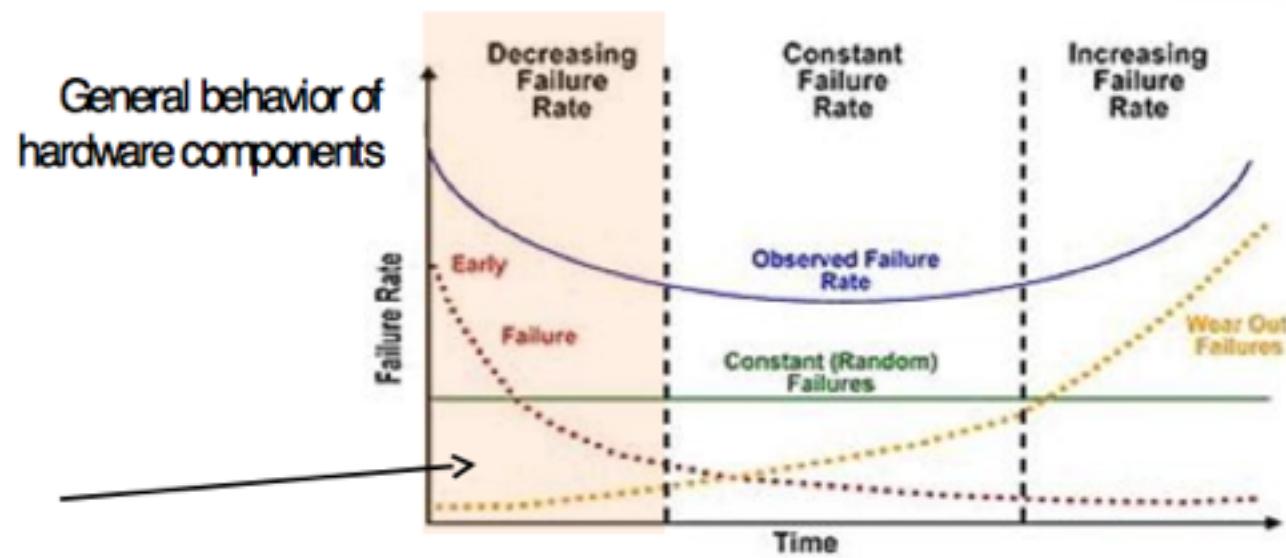
Well, a dirty Higgs factory!

- **HL-LHC: 25 ns bunch crossing, $L=5 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$**
- Higher luminosity is reached by increasing the number of interactions/collision, but new future techniques (leveling, crab cavities, crab kissing...?) can modify the interaction region and help in maintaining the number of overlapping interactions low
 - Pessimistic view: experiments must deal with 140 <interactions per collision>, maximum 200
- Detectors requirements will go **beyond the current design** specifications:
 - Higher **peak luminosity** means increased density of interactions in space and time and higher detector occupancy: need higher resolutions
 - Higher **integrated luminosity** pose limits of irradiation damage and activation of materials



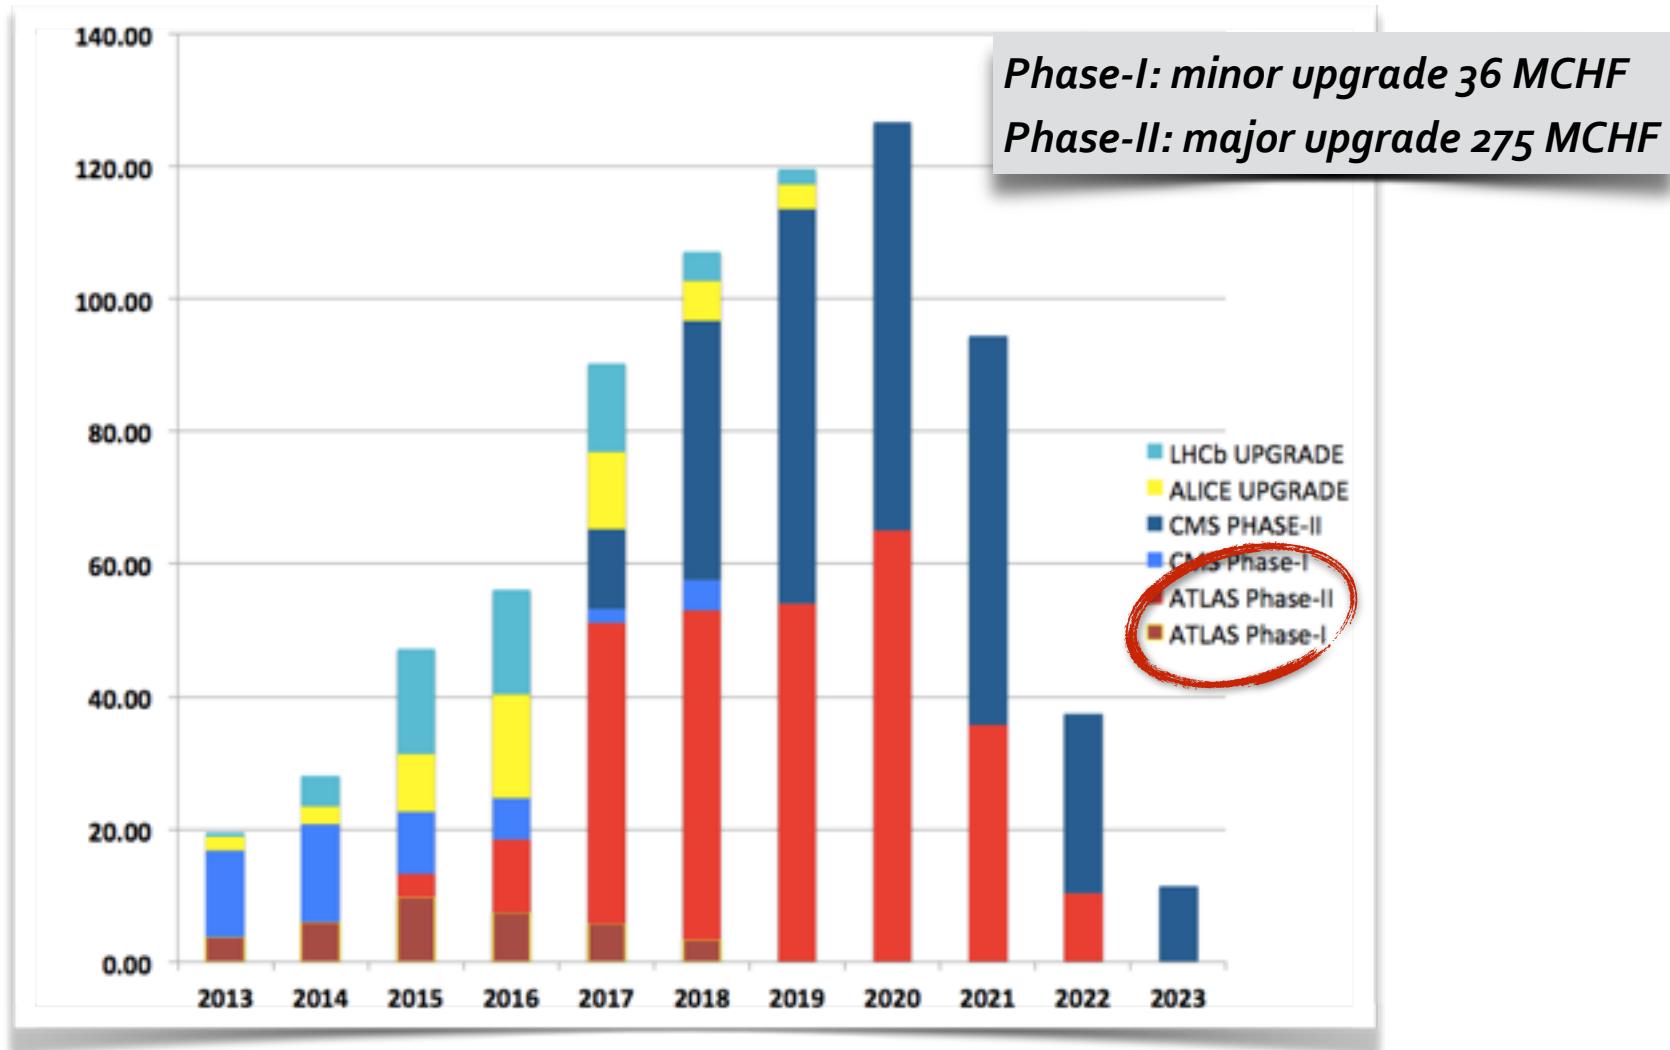
Real life.... another reason to “upgrade”

- A lot of hardware components become old
- System reliability decreases
 - It makes sense to replace PCs and network equipments every 5 years
 - Custom hardware is usually kept longer... by of course it also starts breaking



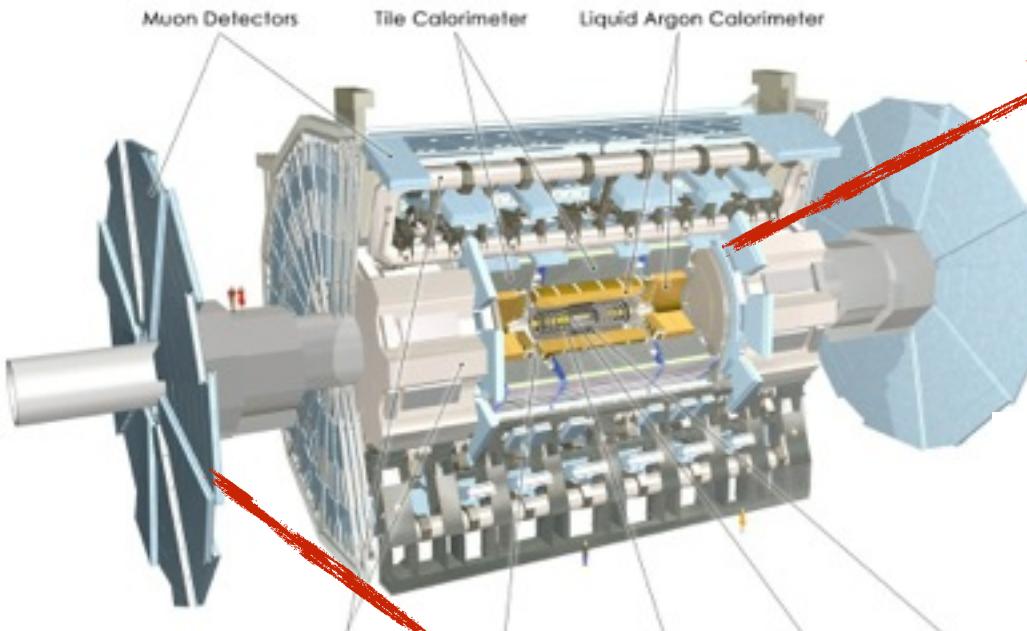
: Failure rate versus t

Cost of all LHC upgrades

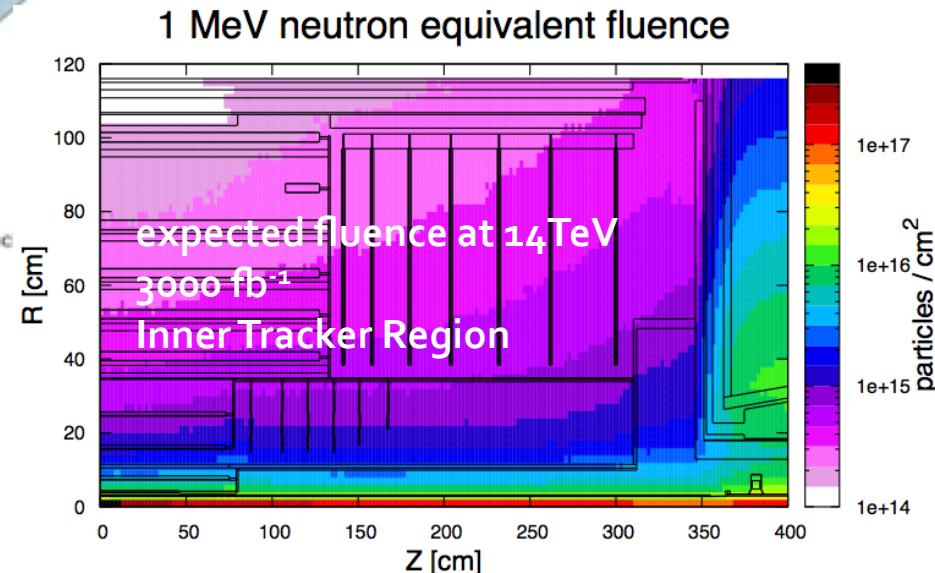


What ATLAS will change for HL-LHC

Phase-II Lol: <https://cds.cern.ch/record/1502664?ln=en>

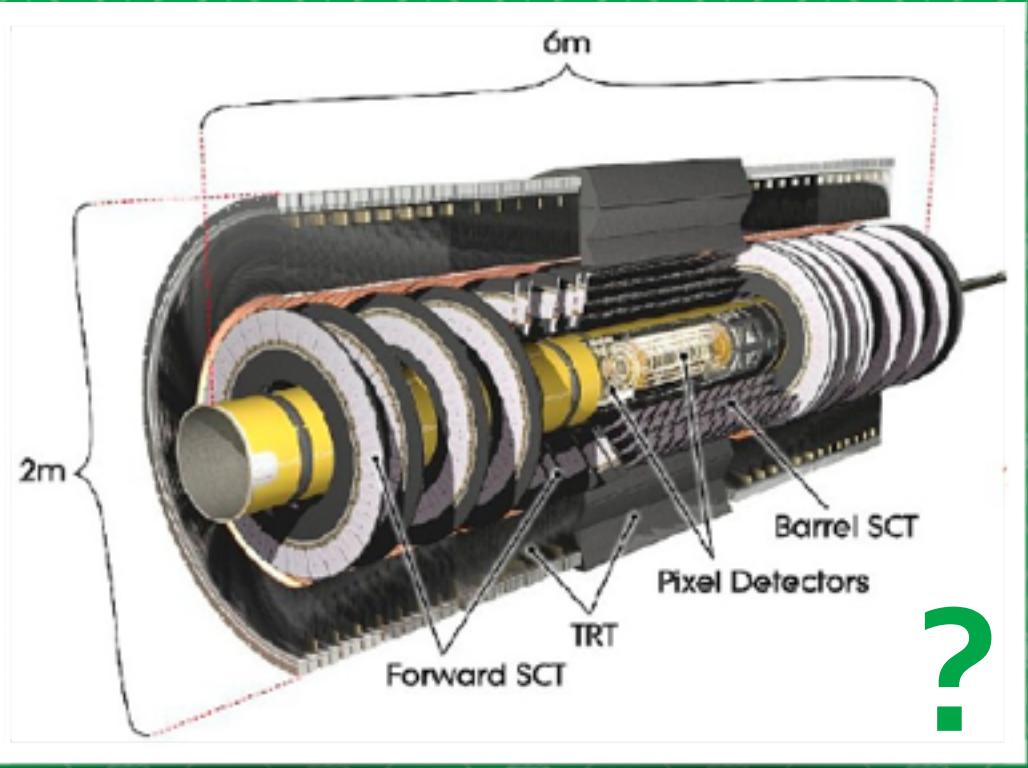


- Inner components ($R < 1$ m) will suffer from radiation damage and high occupancy
 - New silicon tracker, current one would not survive
 - New calorimeter FE electronics



- Outer components ($R > 1$ m) will suffer from pile-up and high occupancy: reduce sensor size, increase redundancy and fast time response
 - Some new muon chambers (inner endcaps) will be installed already in LS2

Higher trigger rates will impose a new design of the trigger and DAQ system (TDAQ)



The silicon trackers evolution

Inner Tracker: key issues for the Upgrades

Increment of leakage currents with int. Luminosity



- **Radiation damage** with more integrated Luminosity, observed in ATLAS, CMS and LHCb (**RD50 R&D project**)

- Increase in leakage current and S/N degradation
- Projections demonstrate that the tracker will survive 500 fb^{-1} if operated at -20 C after LS1
- Must replace the full tracker after LS3

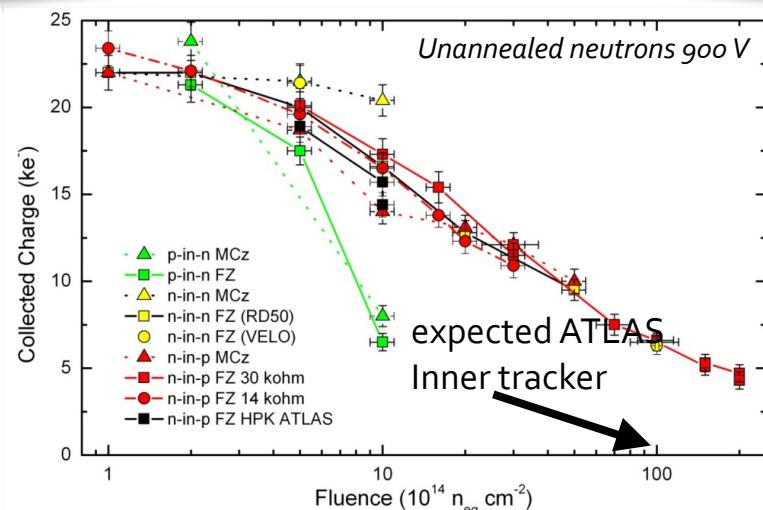
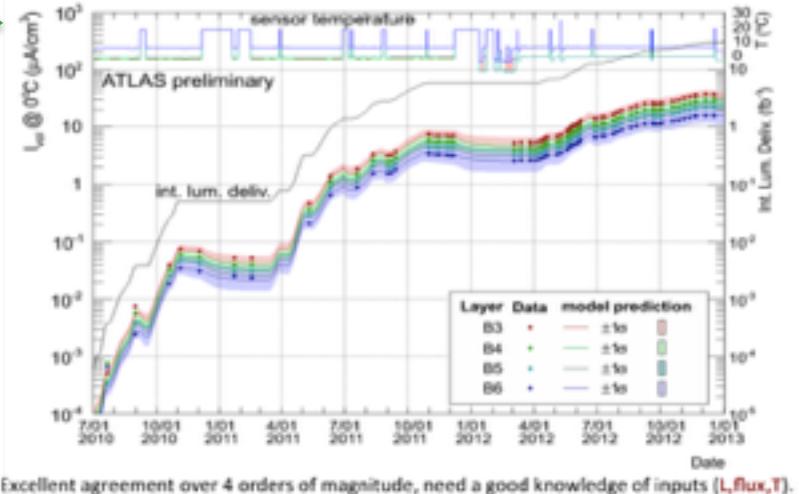
- **Increased performance**

- Higher granularity
- Lower material budget

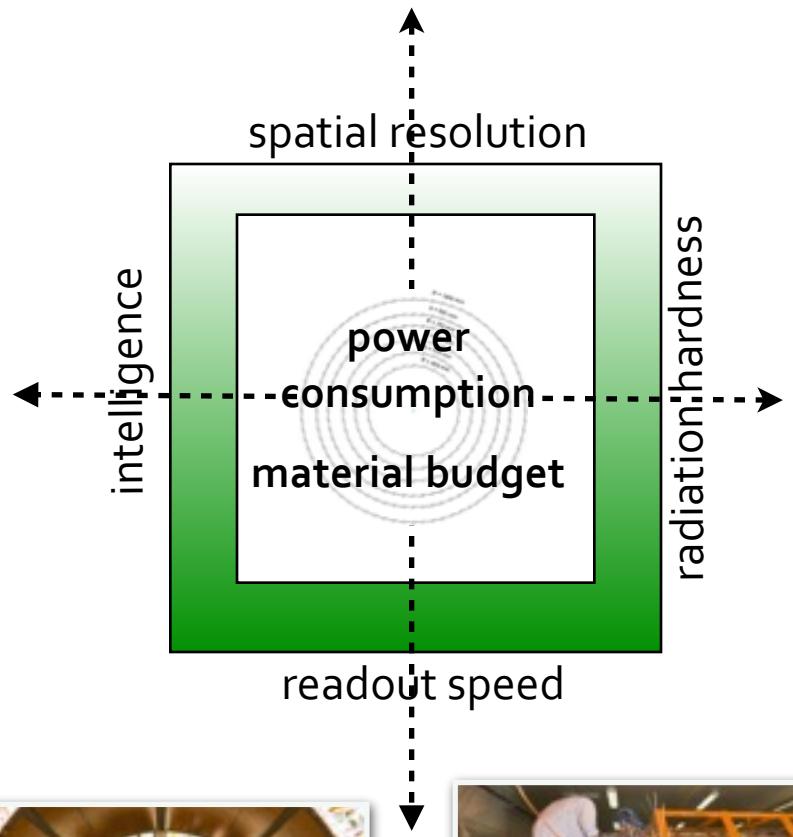
- **Control and minimize cost**

- Large areas & stable/timely production

Signal amplitude decreases with irradiation



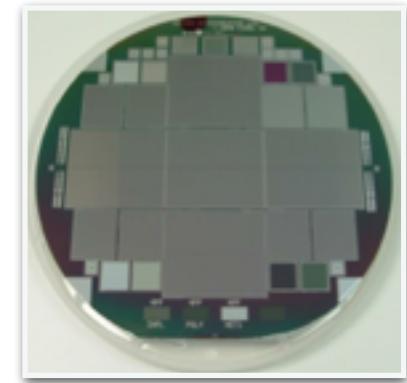
The quadrature of vertex detectors



Microstrip Stave Prototype



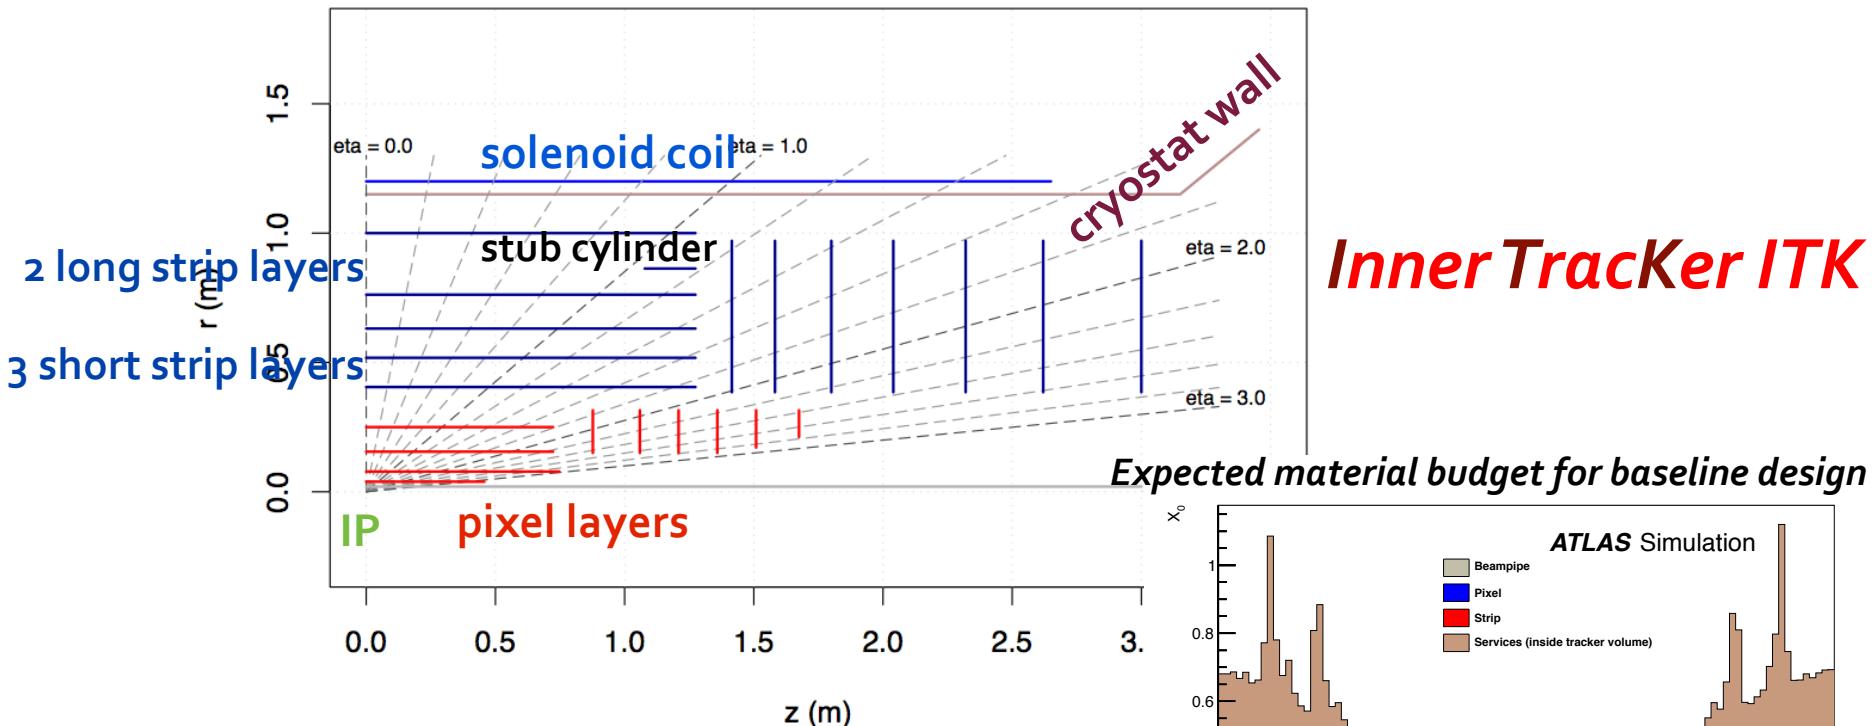
Quad Pixel Sensor Wafer



- In p-p environments, the high level of radiation and hit occupancy imposes **struggling requirements**
- Adding features and performance means increasing the number of chips, then power consumption and additional material

The future all-silicon Inner Tracker at HL-LHC

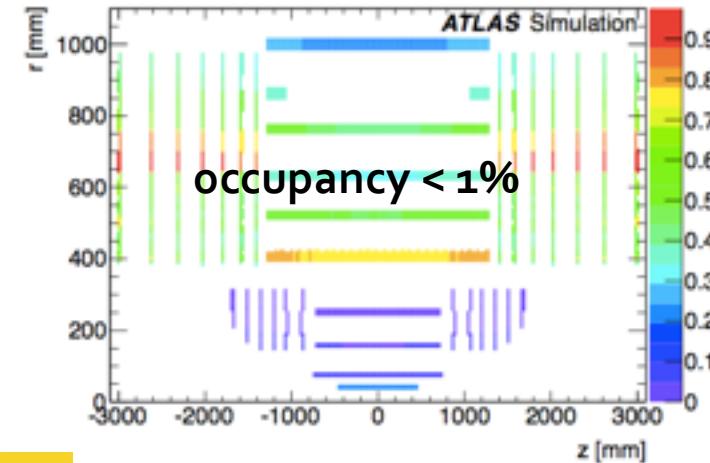
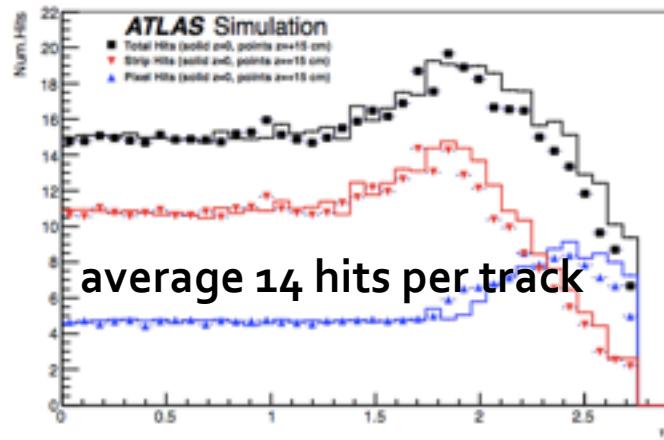
- Full silicon tracker: barrel cylinders and endcap disks, with different granularity
- Baseline layout to maintain optimal tracking performance (and cost)



Careful study of the material budget, with consequences on:

- tracking performance
- more flounce in the full detector

Expected performance of the baseline layout



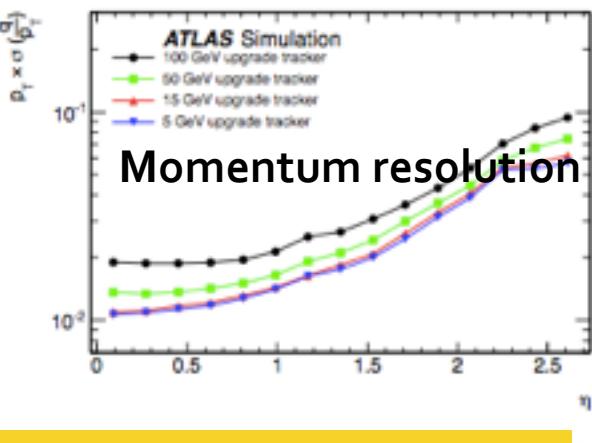
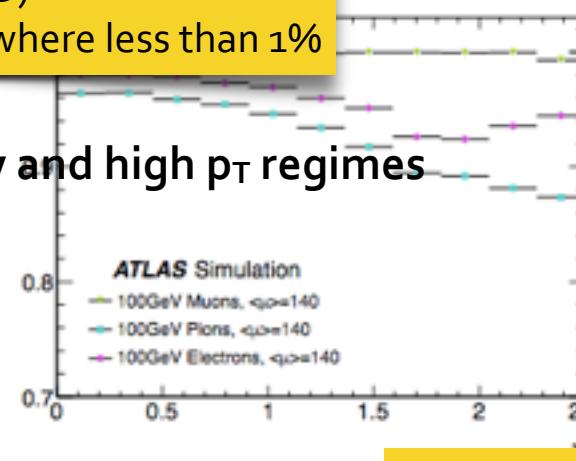
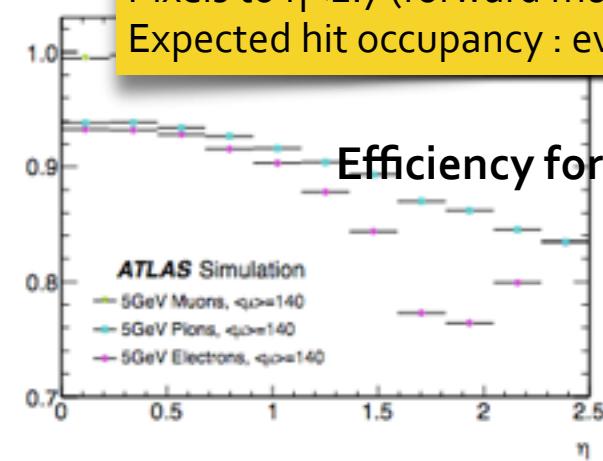
Robust tracking:

total of 14 hits with full coverage to $\eta=2.5$

Pixels to $\eta < 2.7$ (forward muon ID)

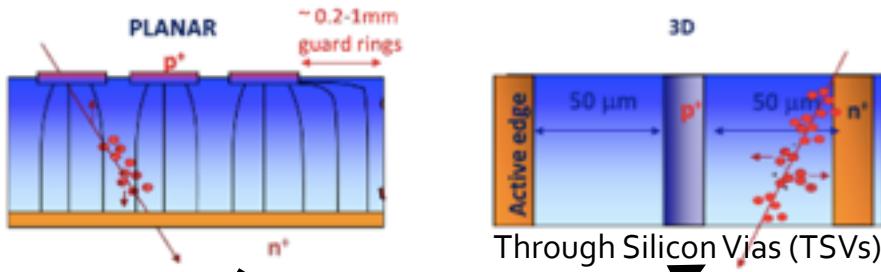
Expected hit occupancy : everywhere less than 1%

Efficiency for low and high p_T regimes

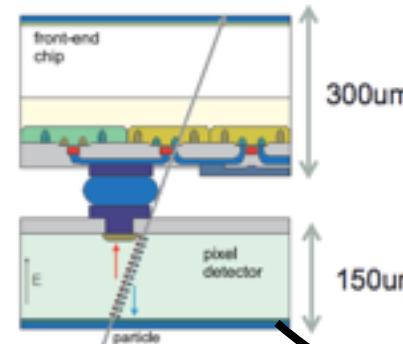


But other layouts are under study

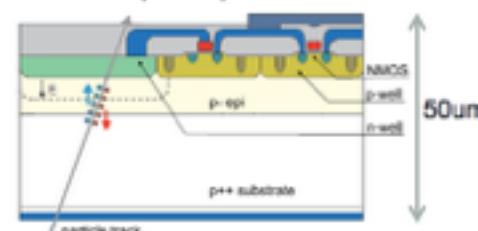
All-silicon sensors evolution (few words)



Hybrid Pixel Detector



CMOS (Pixel) Detector



Planar silicon-sensors

- n-in-p : Single-sided process (less expensive)
- n+-in-n : Double-sided (more expensive)
- Both can work at HL-LHC radiation levels
 - If carefully designed...
 - And if they are kept cold $\sim 20\text{ C}$

3D sensors

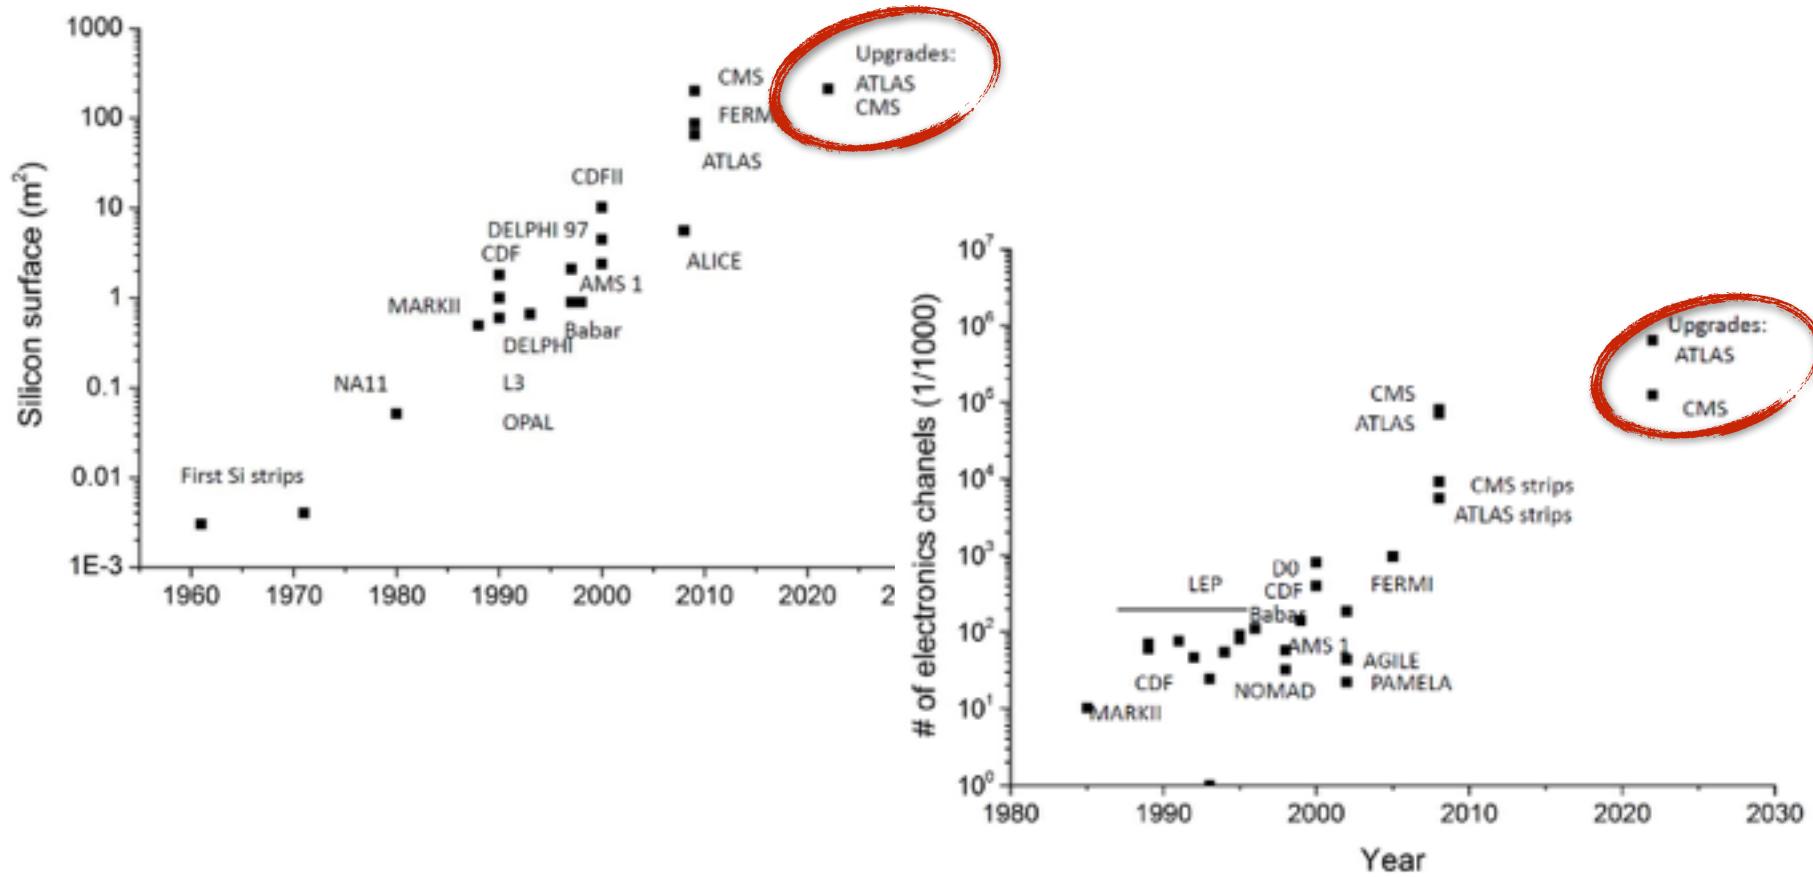
- Very good performance at high fluences
- Production time and complexity to be investigated for larger scale production
- Used in ATLAS IBL (LS1 upgrade)

CMOS sensors

- Contain sensor and electronics combined in one chip
- Standard CMOS processing (many foundries, lower cost/area)
- Prominent advantages: high granularity, low material, high data throughput

Upgrades	Area	Baseline sensor type
ALICE ITS	10.3 m^2	CMOS
ATLAS Pixel	8.2 m^2	tbd
ATLAS Strips	193 m^2	n-in-p
CMS Pixel	4.6 m^2	tbd
CMS Strips	218 m^2	n-in-p
LHCb VELO	0.15 m^2	tbd
LHCb UT	5 m^2	n-in-p

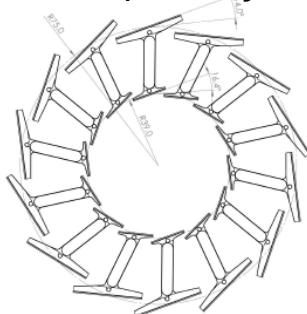
Time evolution of highly segmented silicon detectors



The tracker elements

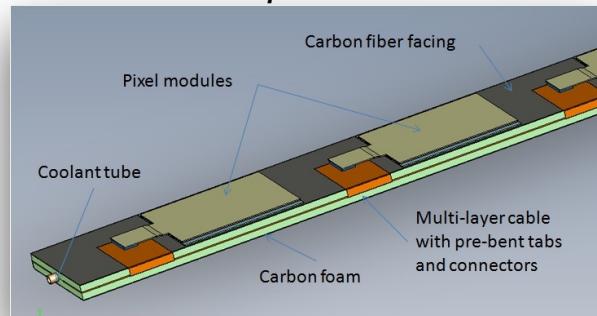
- **Robustness:** detector modules are integrated and fully functional packages, called **staves**, that can be produced in parallel and fully tested before assembly
- **Reduce material:** services are included in the module (cooling, monitoring, control....)
- **Data links** challenge the high radiation level and the high data throughput (~1Gbps)
 - **pixels** use twisted micro-cables to send LVDS data to a dedicated optical board
 - **strips** use dedicated optical links (CERN Versatile) up to 5Gb/s

Inner pixel layers

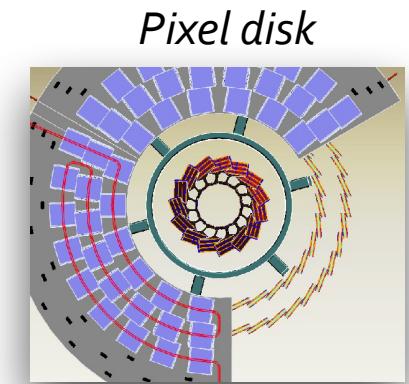


I-beam shape and clamshell design (n+-in-n sensors)

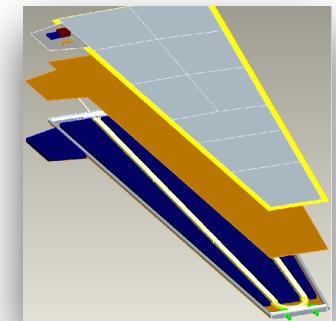
Outer pixel stave



5mm thick staves, with modules on both sides (n-in-p)



Pixel disk

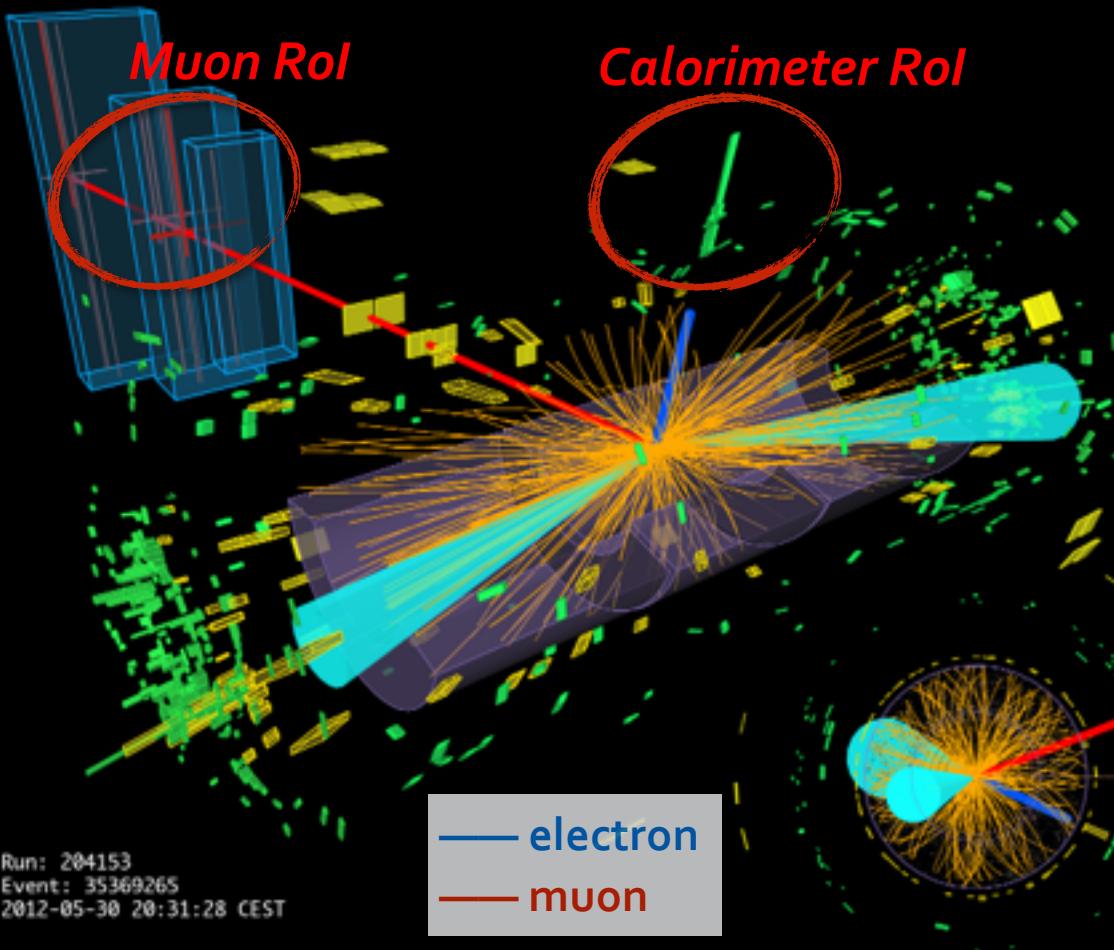


Endcap strip petal

Commercial copper cables can transmit several Gb/s over tens of meters. However, the diameters of these cables are too large for the pixel detector.

The trigger upgrade strategy

Legacy from today's trigger system

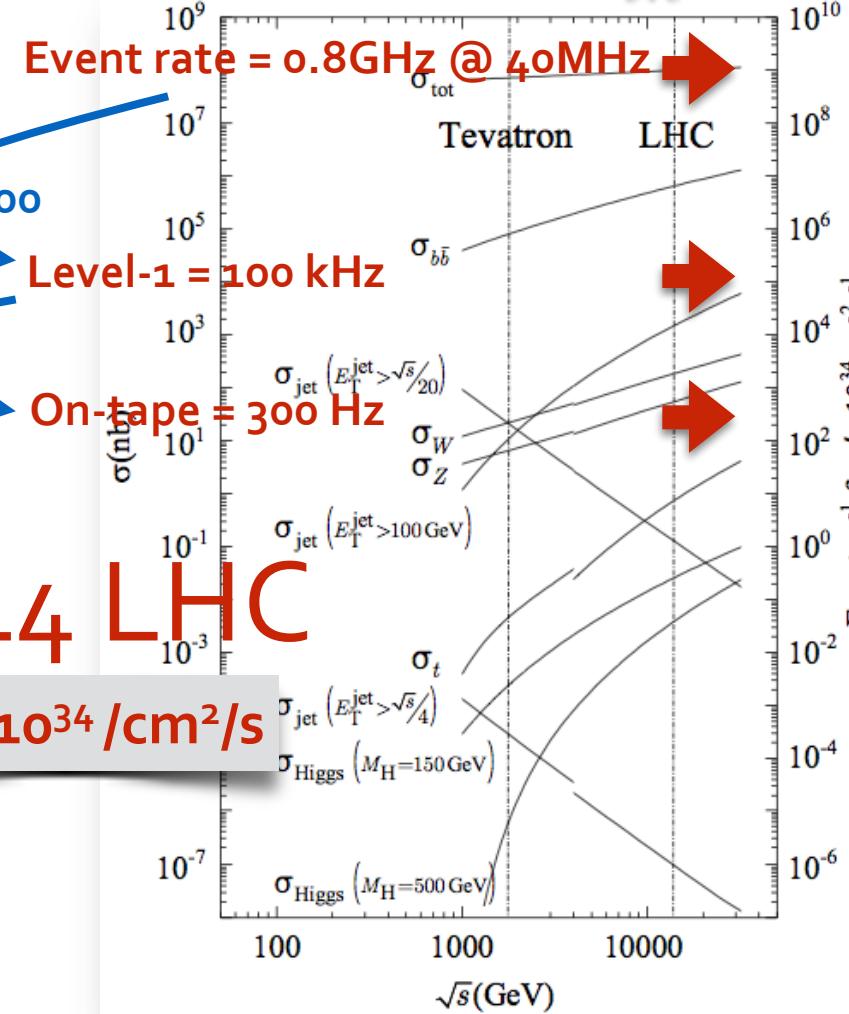


- **First-level trigger (L1)**
 - Synchronous at 40MHz, with fixed latency: **2.5 μ s**
 - Identifies **Region-of-Interest (Rol)** in the muon spectrometer and/or in the calorimeter, with coarse resolution
 - No tracking information can be used due to limited latency
- **High-level trigger (HLT)**
 - Handles complexity with custom fast software on commercial CPUs
 - Accessing the full resolution of all the detectors (both Rols and full event)

Event display of a 2-tau event in the ATLAS detector. Run number: 204153, Event number: 35369265. The taus decay into an electron (blue line) and a muon (red line).

Expected trigger rates at HL-LHC

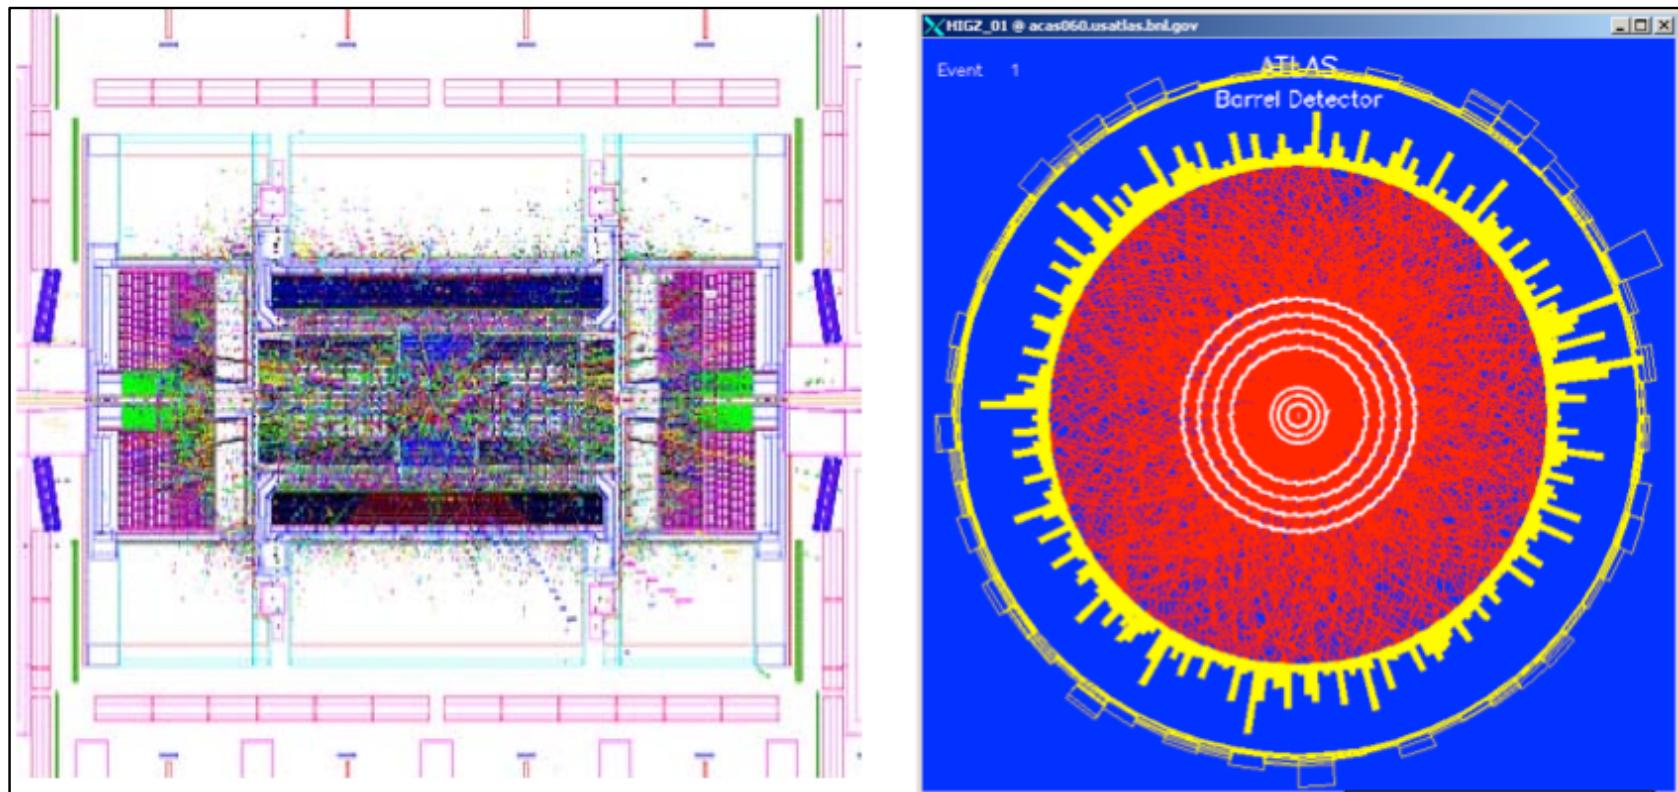
$$R = \sigma_{in} \times L$$



- Change of FE buffer size
- Maintain ~same rejection factor on HLT
- Event size will increase : 1.5MB to 2MB
- Challenging storage: 10-20 GB/s
- Moore's law can handle this!

But in one ATLAS event at High-Luminosity ($L=5\times10^{34} \text{ cm}^2/\text{s}$)

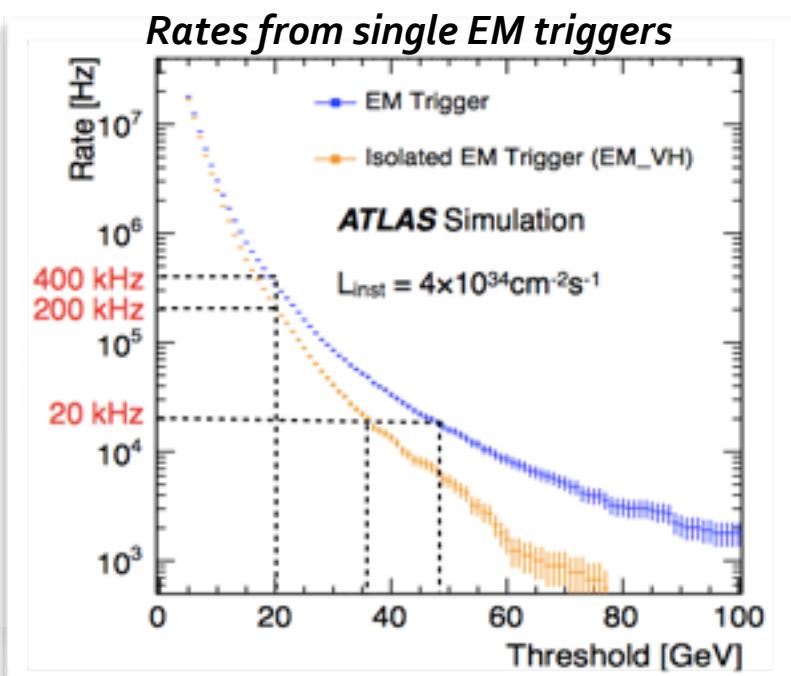
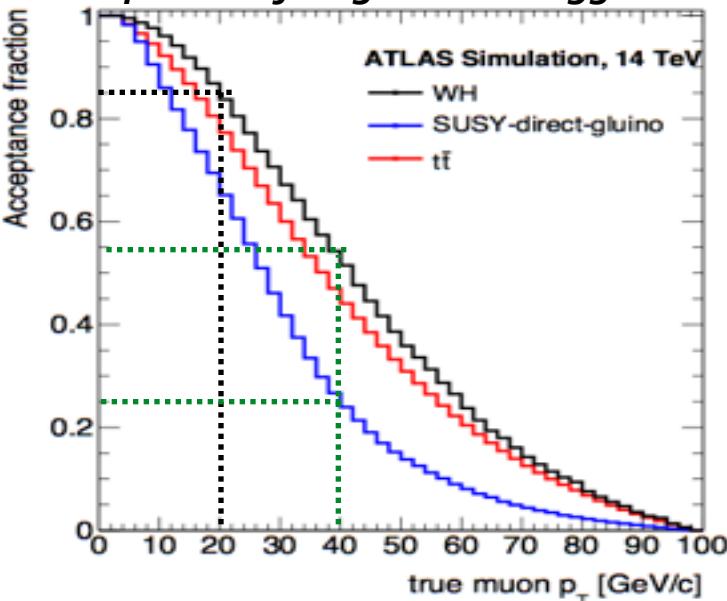
- 200 collisions per bunch crossing (any 25 ns)
- $\sim 10\,000$ particles per event
- Mostly low momentum (p_T) particles due to low transfer energy interactions



The trigger selection will become harder and harder

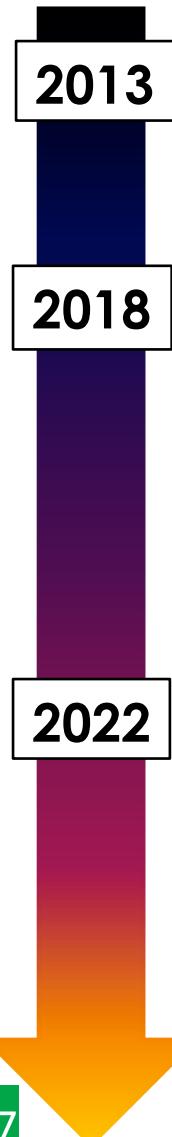
- Trigger strategy: maintain adequately wide trigger selections at the Electroweak scale:
 - Inclusive single leptons with thresholds ~LHC
 - Exclusive / multi-object triggers
 - Increased importance of tau and adronic (MET) triggers

acceptance of single muon triggers



- Higher occupancies in the detectors bring:
 - Increased fake rate
 - Jets mimicking electrons
 - High radiation in the forward regions
 - Reduced rejection power of the algorithms
 - Worse resolution in calorimeters
 - Less effective isolation and pattern recognition

Phases of the L1 trigger evolution: become more intelligent!



- Phase-0: be prepared for $L = 10^{34}/\text{cm}^2/\text{s}$ (PU~25)
 - Complete detector & consolidate operations
 - Allow L1 topological criteria / more exclusive selections
- Phase-1: be prepared for $L = 3 \times 10^{34}/\text{cm}^2/\text{s}$ (PU~40)
 - Add more flexibility, without major architectural changes:
 - Additional coincidence layers in the forward muon spectrometer
 - Increased granularity in the calorimeter
- Phase-2: be prepared for $L = 5 \times 10^{34}/\text{cm}^2/\text{s}$ (PU~140)
 - Major upgrade for HL-LHC era: ensure appropriate rejection
 - Expected L1 rates over the limit allowed by detector FE
 - A new tracker will be available...

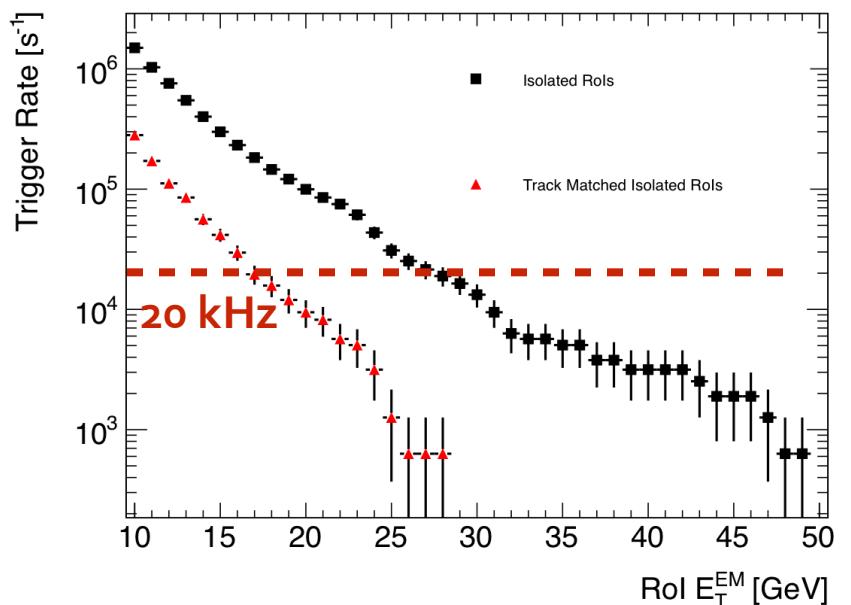
Any component installed in Phase-I must be fully operational also through Phase-II

First-level trigger in Phase-2: deriving ideas from HLT...

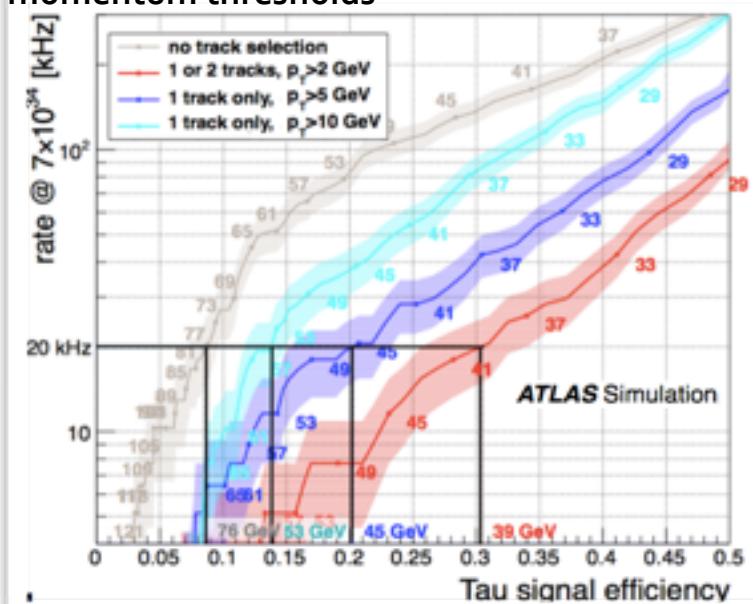
■ Tracking information at L1: adding flexibility

- Combines calorimeter/muon with tracks, to remove mis-reconstructed or fake objects
- Provides track isolation and multiplicity for τ , impact parameter for b-tagging
- Vertex information for multi-object triggers (multi-jet)
-

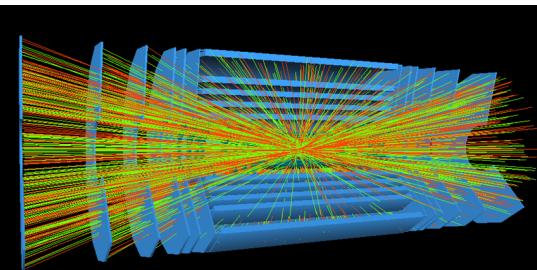
EM rate reduction when applying the track match @14TeV, L=3e34, $\langle \mu \rangle = 70$



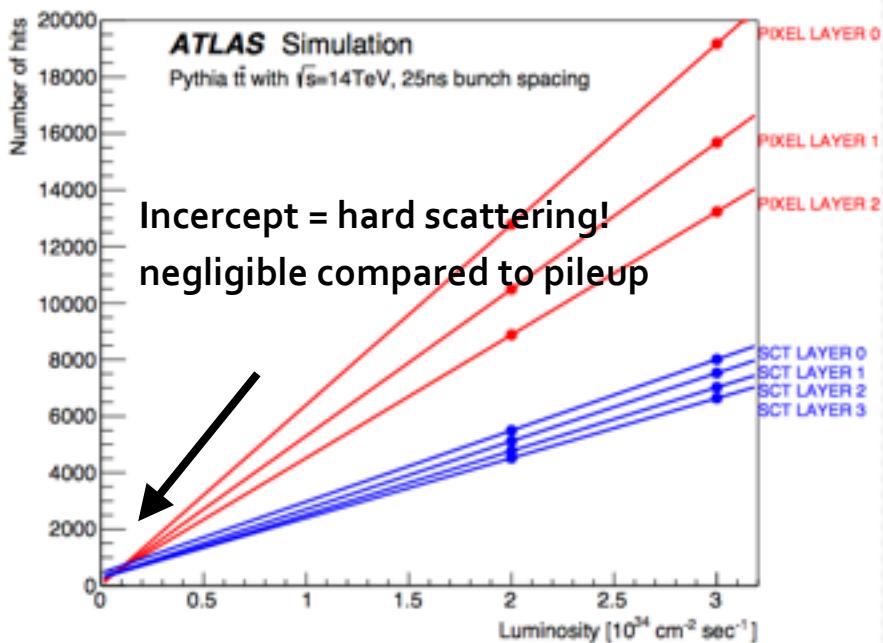
Tau rate reduction and efficiency with different selections based on track multiplicity and momentum thresholds



Triggering on tracks at L1 may be difficult!

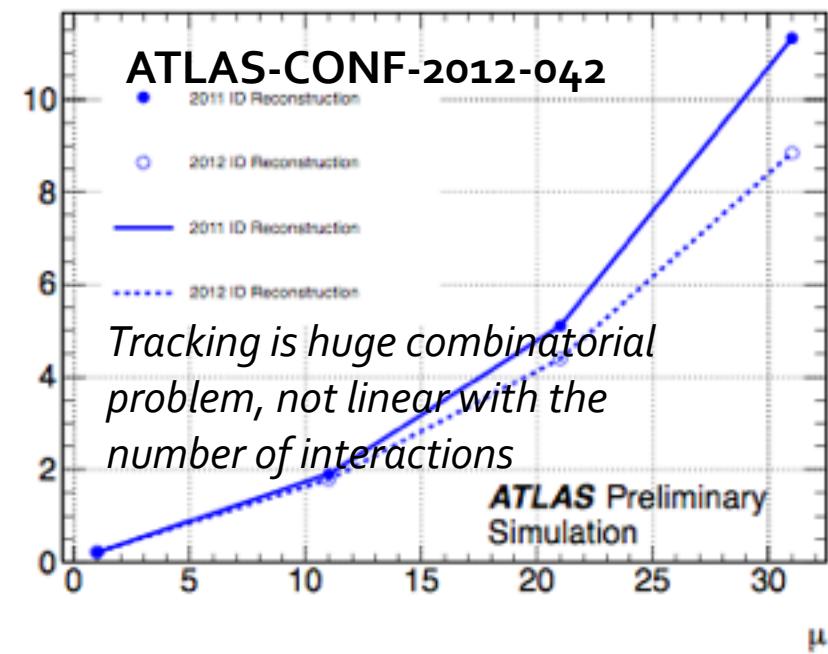


- Cannot readout the full tracker at 40 MHz
- Already $\sim 10\text{-}20 \text{ Gbs/cm}^2$ per layer at L1



- Data reduction/reformatting

- Reconstruction complexity/timing naively scale with the number of tracks....



- Longer latencies/larger Front End buffers

- Faster data transmission and processing (Increase parallelism, network and trigger CPU needs)

The ATLAS L1Track project

So far....

- Simulation studies to define upgrade requirements and evaluate detector and physics performance at high Luminosity
 - Even modest (p_T , η , ϕ) resolution on tracking information can provide sufficient rejection
 - Rejection **x3** for muons and **x10** for electrons, with only small efficiency losses
 - Double-lepton signatures are under control
 - Minimum track p_T can be ~ 17 GeV for single leptons, few GeV for double signatures and taus

Next steps.....

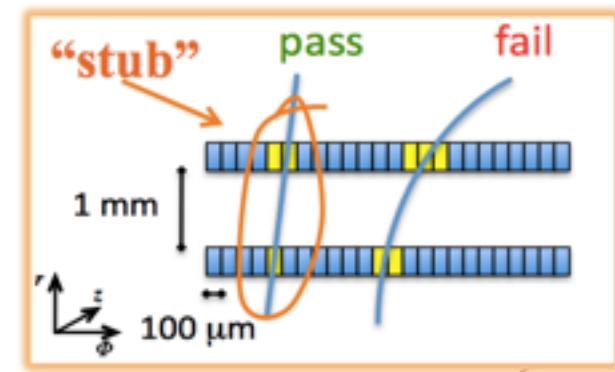
- Development of conceptual design and technical solutions during next coming years, in connections with the Tracker upgrade (tracker construction will start in 2016)
 - Good view of the L1Track system design for the Initial Design Review in **2015**
 - Document the overall scope in a Technical Design Proposal around **2016** (same time ITK TDR)

L1Track is effective in reducing the rates in two momentum regimes: high- p_T single leptons, low- p_T double leptons and taus.

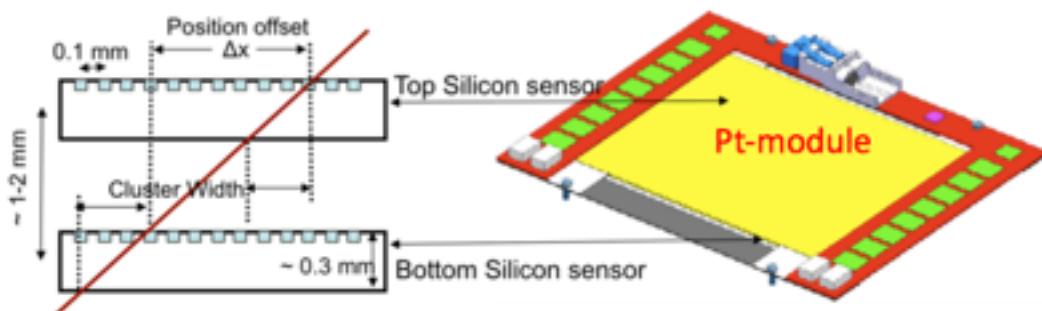
Object(s)	Trigger	Estimated Rate	
		no L1Track	with L1Track
e	EM20	200 kHz	40 kHz
γ	EM40	20 kHz	10 kHz*
μ	MU20	> 40 kHz	10 kHz
τ	TAU50	50 kHz	20 kHz
ee	2EM10	40 kHz	< 1 kHz
$\gamma\gamma$	2EM10	as above	~ 5 kHz*
$e\mu$	EM10_MU6	30 kHz	< 1 kHz
$\mu\mu$	2MU10	4 kHz	< 1 kHz
$\tau\tau$	2TAU15I	40 kHz	2 kHz
Other	JET + MET	~ 100 kHz	~ 100 kHz
Total		~ 500 kHz	~ 200 kHz

CMS approach: low- p_T track filtering

- CMS is designing tracker and FE modules with p_T discrimination capability
 - Reject low- p_T tracks, reducing data volume by one order of magnitude (40 MHz to \sim MHz)
- Correlate signals in two closely-spaced sensors, exploiting the strong magnetic field of CMS, with two steps:
 - Cluster width approach:** preselection of hits according to their cluster width
 - Stacked tracker:** correlation between preselected hits in nearby sensors



6

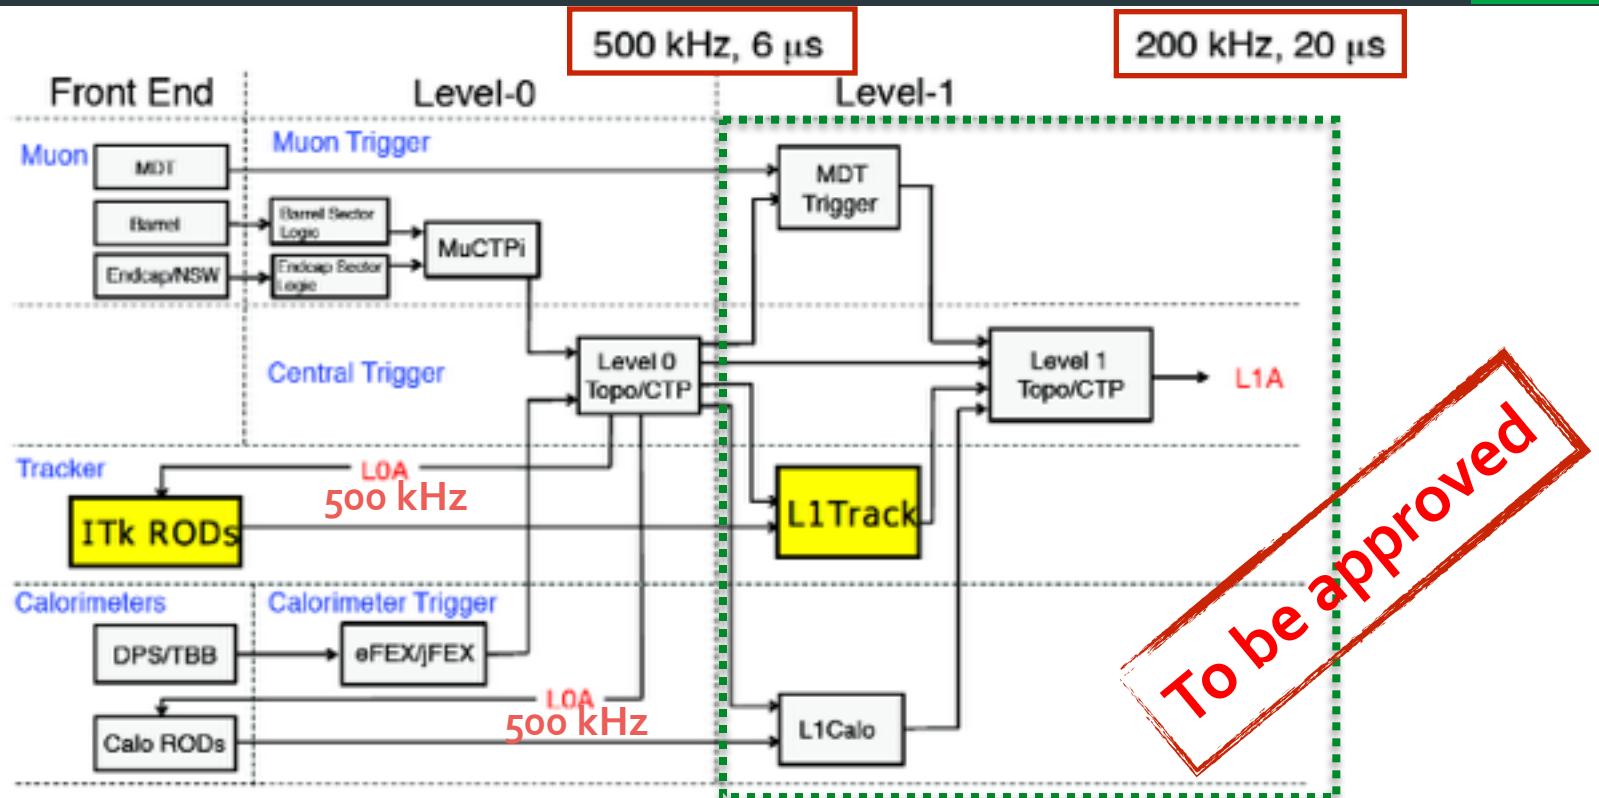


Main challenges:

- L1 latency $< 10 \mu\text{s}$**
- L1 requirements affect the design of the tracker**

- Different geometries are under study, to have coherent p_T threshold over the entire volume
- Material may affect resolution at low- p_T due to MS

Increase latency: a new trigger scheme for ATLAS Phase-II

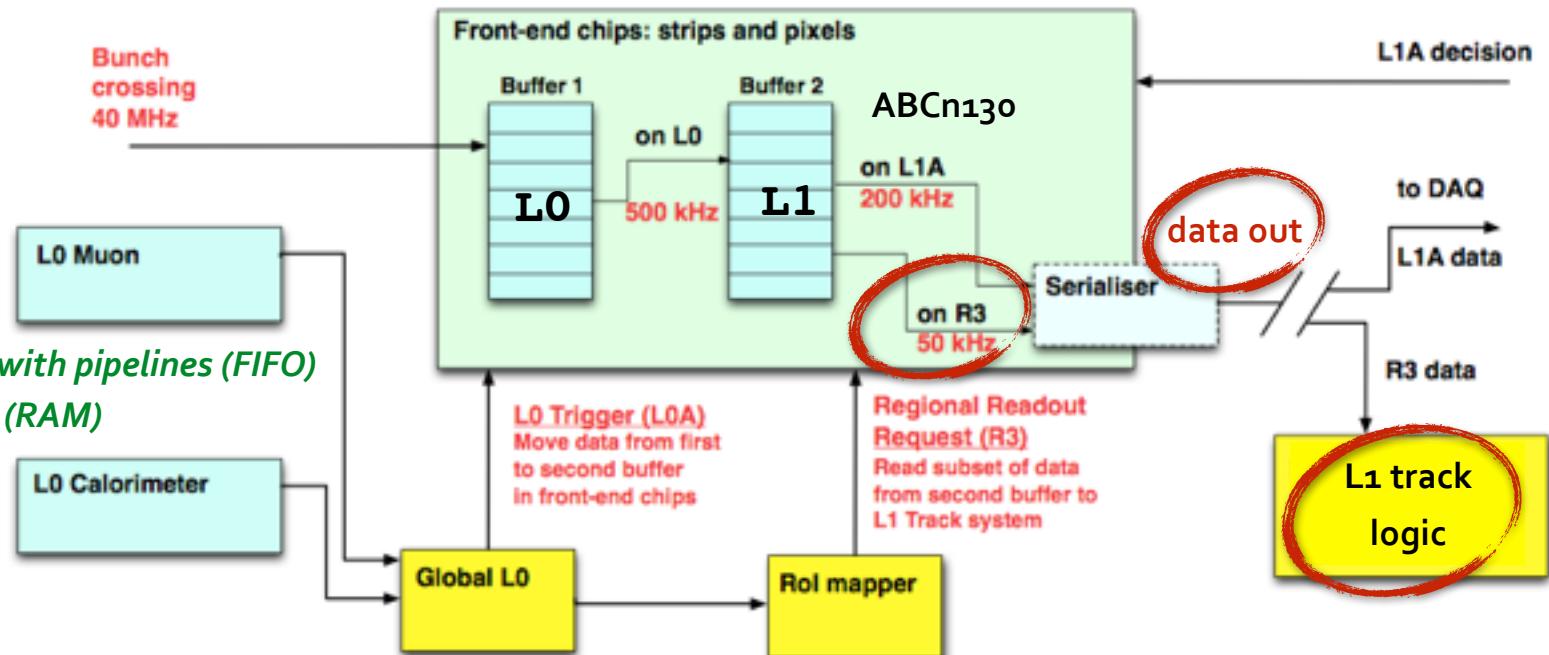


- Exploit the Region-of-interest mechanism!
- Add one trigger level, with extended latency (20 μ s) to include the tracker information at L1
- Scale down current L1 to become L0, with extended latency (from 2.5 to 6 μ s) and increased accept rate (0.5MHz, maximum 1MHz)

Double-buffer readout strategy

ABCn130 FE chip (on the tracker staves):

Analog Binary Chip, 130nm CMOS ASICs, with 256 readout channels and double-buffer architecture



- Need to reduce data throughput from Lo buffer (LoA rate ~0.5/1 MHz)

- Exploit Region-of-interest mechanism: only 10% of the chips are readout after a LoA, via a **Regional Readout Request (R3)**: any chip has a reduced data request rate at ~50 kHz (10% of 500 kHz)
- Data can be reformatted for trigger purpose (reduced information, like filter clusters)

- Need anyhow to increase the latency: everything must be completed within <20 μ s

L1Track latency budget

- To stay within 20 μ s latency, crucial limits on the readout (6 μ s) and on the L1TT algorithm (6 μ s) timings are imposed
- **Can we handle this?**
- Readout data size contributes hugely to latency, but tracking doesn't need complete data
 - If necessary, data size can be reduced
 - Different formatting strategies under study

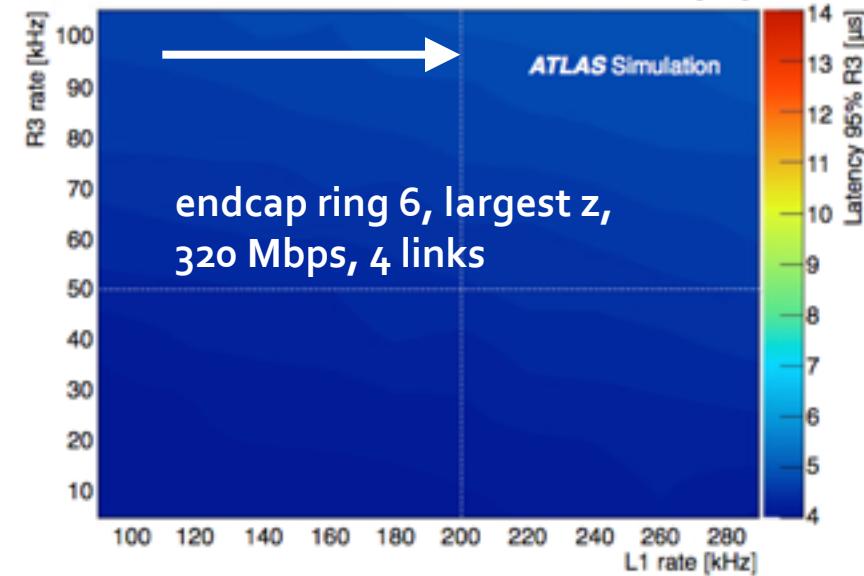
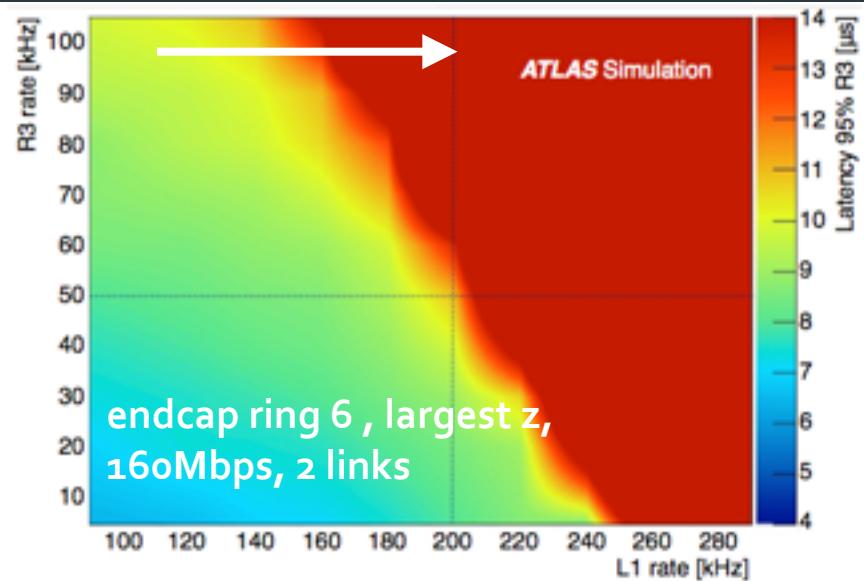
	Latency [μ s]	cum. Latency [μ s]
formation of L0A	3.0	
map R0L-ITK and send R0Ls to ITK	1.25	4.25
ITK readout in R0L regions	6.00	10.25
transmit to L1TT	2.00	12.25
L1TT algorithm	6.00	18.25
L1A formation from track+L1MU+L1Calo	1.00	19.25

- Different deadlines for decisions
 - All aspects connected to the actual detector (front-end, communication etc.) fixed before construction of ITK starts
 - **Data formatting** studies to be completed in parallel
 - **Trigger processor technology** can be decided later

What about robustness?

- What if different unexpected conditions (increase of Lo/L1 rates, occupancy) occur?
 - Final cost increases as LoA rate increases (cost of links and processors)
 - If LoA rate increases and the allowed L1A rate cannot increase, more processing is needed and cost/complexity of Level-1 increases

- For the tracker design:
 - **ABC130 buffers**: increase is not costly
 - **Latency**: increase is negligible if extra links are added
- See the latency maps for Endcap ring 6 at largest z
 - on some critical endcap rings, latency can change rapidly if the bandwidth is not appropriately increased



Readout challenges for the ITK detector

Pixel Readout

- Readout the full event at 1MHz is possible with IpGBTx, new protocols are under study to include contingency

- Requirements on material and financial cost

Strips Readout

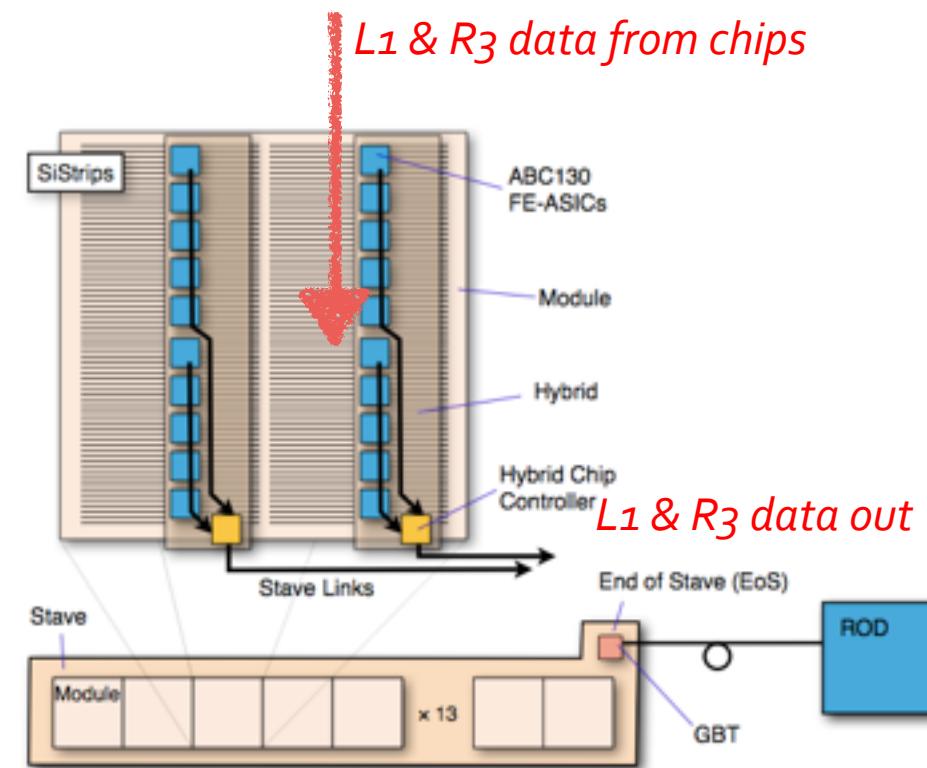
- Long staves require large data bandwidth, 1MHz readout possible using reduced regional requests (R₃)

- More links? More buffers? Different strategies to control traffic are under study, without undue additional material (and power)

- redundancy vs reliability: call for engineering!

- ABC130 chip prototype already prepared

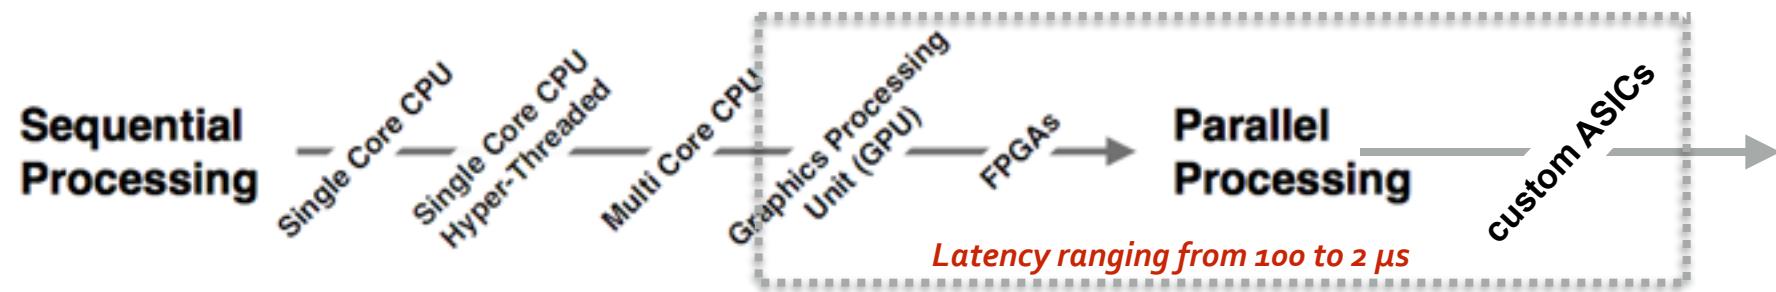
- Hybrid Chip Controller (HCC) still under design



There is a delicate rate/latency balance : queuing buffers absorb peaks of rate but cause deadtime

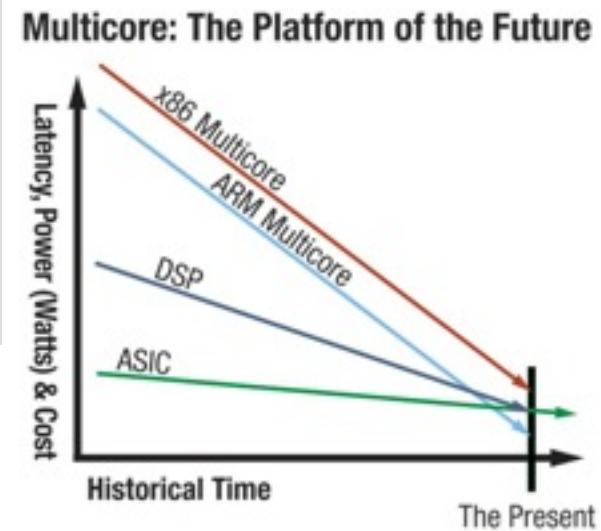
L1Track trigger logic challenge

Level-1 trigger processors (a small excursion on technology trends)

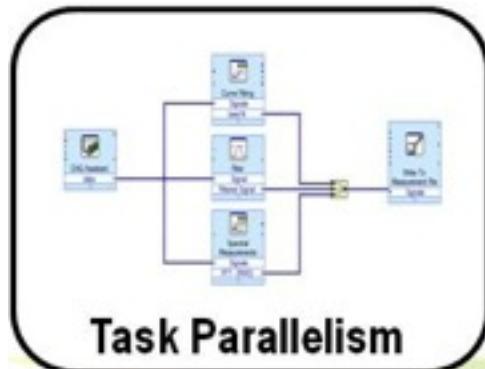


- Follow the evolution of digital integrated circuits on a single chip (SoC)
 - Request of higher complexity \Rightarrow higher chip density \Rightarrow smaller structure size (for transistors and memory size): **32 nm \Rightarrow 10 nm**
- Custom ASICs or Off-the-Shelf component (COTS):
 - Specific microprocessors (CPUs, DSPs=Digital Signal Processors,...)
 - Programmable logic devices (FPGAs)
 - high throughput, flexible, parallel
 - but development requires more effort

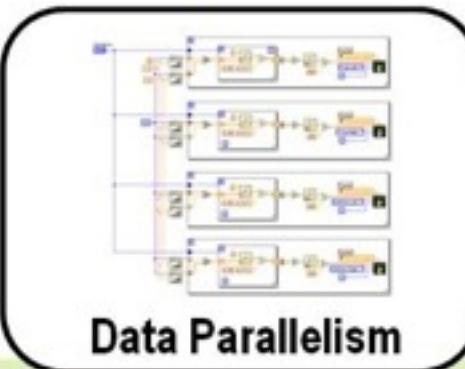
Over recent years, the latency and performance gap between multicore processors has been closing to the point that many of the functions that required the specialized hardware properties of DSPs and FPGAs can now be done in software in General Purpose Processors (GPP).



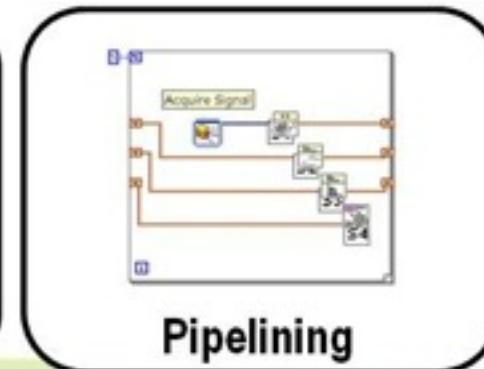
Trends: combined technology



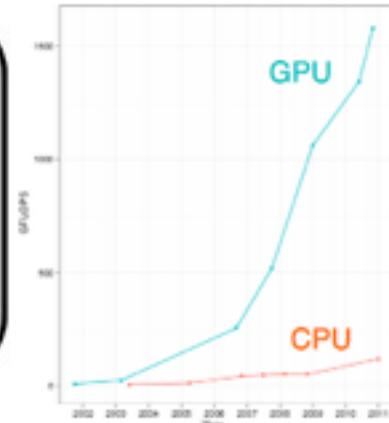
Task Parallelism



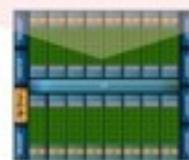
Data Parallelism



Pipelining



**Multicore
Processors**



GPUs*



FPGAs

**Nvidia GPUs:
3.5 B transistors**

**Virtex-7 FPGAs:
6.8 B transistors**

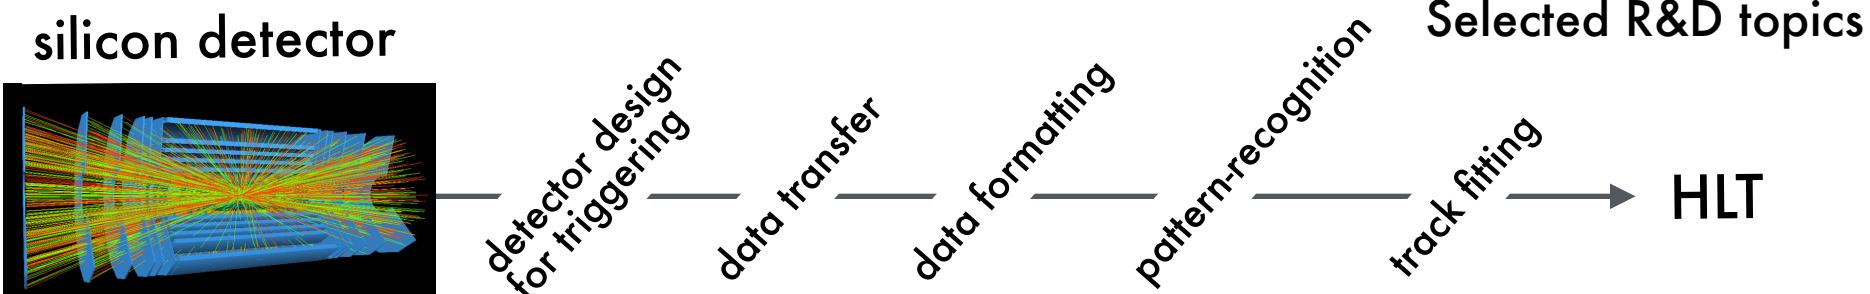
*can implement multiple
DSP algorithms*

(*) Access to the nVIDIA® GPUs through the CUDA and CUBLAS toolkit/library using the NI LabVIEW GPU Computing framework.

The right choice can be to combine the best of both worlds by analyzing which strengths of FPGA, GPU and CPU best fit the different demands of the application.

ATLAS Upgrade: take full advantages of modern real-time technology

- The current technology using fiber data transfer, FPGAs, custom chips and modern PCs could not be scaled in a simple manner to accommodate all the tracking trigger demands
 - Significant improvements, or breakthroughs, will be probably needed. In other words: **aggressive R&D**



The golden time for "easy" digital electronics is over

- Very high clock frequency (20 MHz to 20 GHz and beyond)
- Analog interference on digital electronics becomes important (noise, cross-talk, signal reflection)
 - Major challenges for system design, from power distribution, PCB layout,
- Cannot just buy some FPGAs, write some VHDL code and claim to have an electronics board!
 - see [High-Speed Digital Design: A Handbook of Black Magic](#)

Past, present and future of hardware track-trigger systems

CDF- SVX II



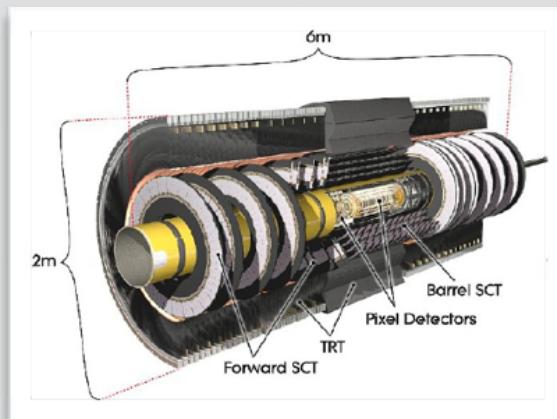
CDF SVX II

- peak $L = 3 \times 10^{32}$, 10 PU
- BC = 396 ns
- $L_1 = 30$ kHz
- $L_2 = 750$ Hz

fast tracking for L_2

ATLAS FTK Run2-Run3

Fast TrackEr over current detector

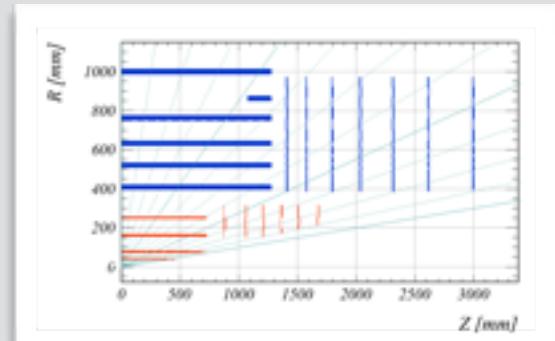


- peak $L = 3 \times 10^{34}$, 69 PU
- BC = 25 ns
- $L_1 = 100$ kHz
- $L_2 = 10$ kHz

fast tracking for L_2

time

ATLAS L1Track Run4



- peak $L = 5 \times 10^{34}$, 200 PU
- BC = 25 ns
- $L_0 = 0.5/1$ MHz
- $L_1 = 200$ kHz

fast tracking for L_1

An evolution of methods and technologies for fast tracking

CDF- SVX II



CDF SVX II

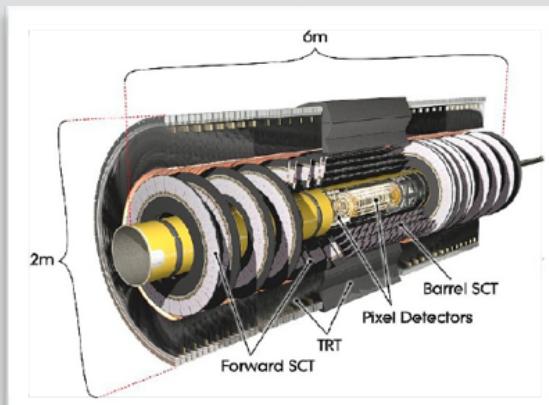
- ❖ ~ 0.2 millions channels
- ❖ L2 decision: ~ 20 μ s, 30 kHz
- ❖ tracks with offline-like resolution: i.e. 35 μ m on the impact parameter

fast tracking for L2

- ❖ Other relevant aspects:
 - ❖ Symmetrical design or not
 - ❖ Materials
 - ❖ Cabling map

ATLAS FTK Run2-Run3

Fast Tracker built over current detector

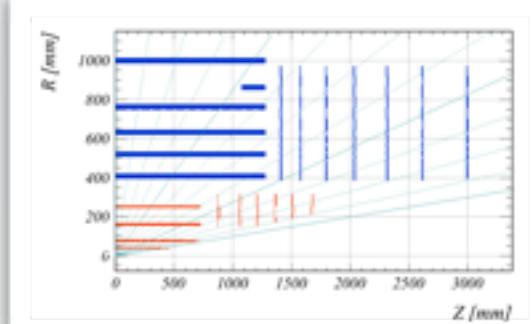


- ❖ 80 M (Pixel) + 6 M (SCT) channels
- ❖ L2 decision :~ 10 μ s, 100 kHz
- ❖ ~offline quality tracks with $p_T > 1$ GeV

fast tracking for L2

time

ATLAS L1Track Run4



- ❖ 638 M (pixel)+ 74 M (strip) channels
- ❖ L1 decision: ~20 μ s, 1 MHz
- ❖ ~offline quality tracks with $p_T >$ few GeV

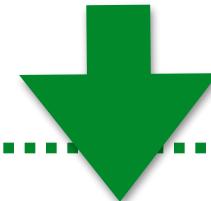
fast tracking for L1

Tracking trigger approach from the past

Same requirements

- Tracks reconstruction close to the offline
- Highly parallelism
- Reduce combinatorics by use of multiple step processing

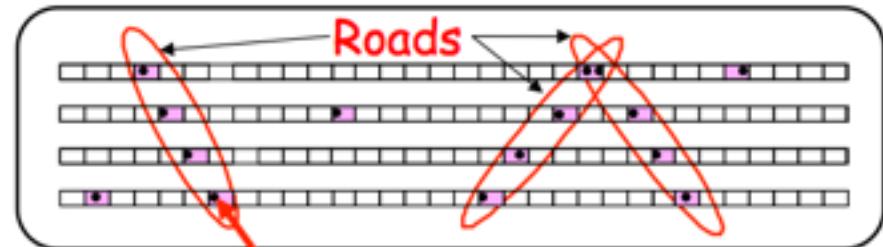
1. Find low resolution track candidates called “roads”
 - Solve most of the combinatorial problems



2. Then fit tracks inside roads
 - Thanks to 1st step, this is much easier
 - A linear approximation gives near ideal precision

A very successful approach at CDF for RunII: SVT (Silicon Vertex Trigger) based on Associative Memory, in turn made of CAM

- APS Panofsky Prize to Aldo Menzione and Luciano Ristori



Pattern recognition w/ Associative Memory
Originally: M. Dell'Orso, L. Ristori, NIM A 278, 436 (1989)



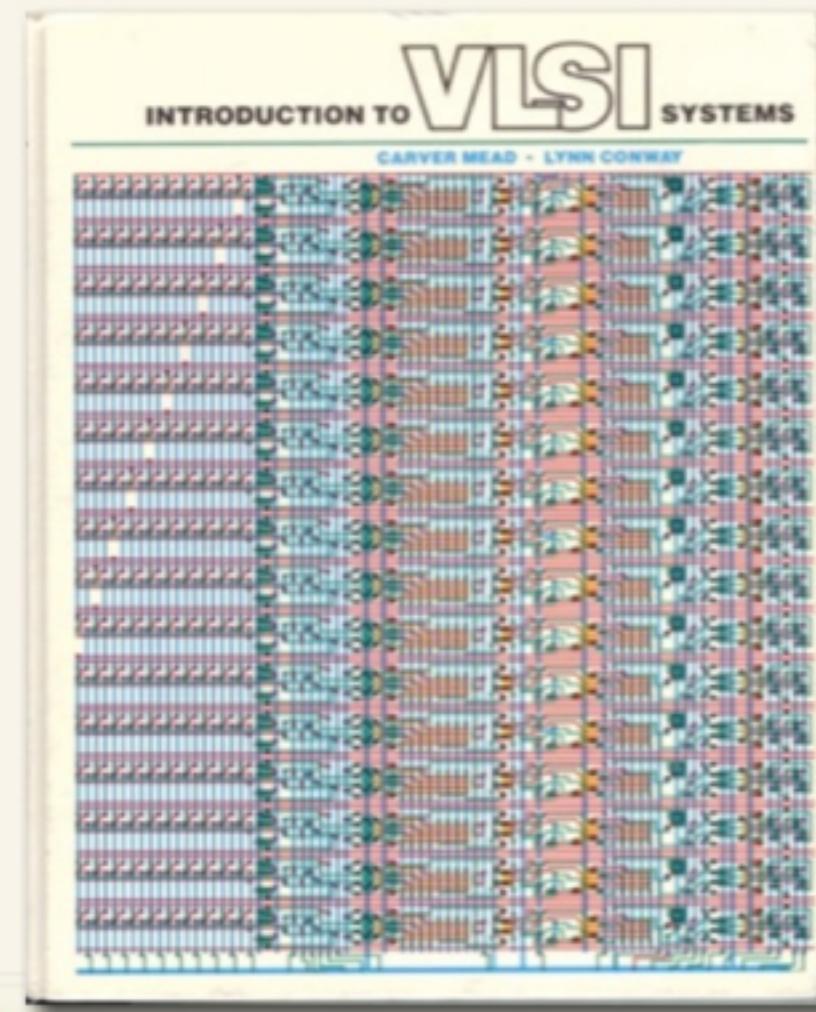
http://www.pi.infn.it/~orso/ftk/IEEECNF2007_2115.pdf

Very Large Scale Integration the revolution

*Stories on some technological
innovations at CDF in the 1980s-1990s*

in the '80s the technology of
VLSI design becomes
available to the universities
and to small research
projects

Carver Mead & Lynn Conway



A slide from Luciano Ristori
at TIPP 2011 conference

October 24, 1988

VLSI STRUCTURES FOR TRACK FINDING

Mauro DELL'ORSO

Dipartimento di Fisica, Università di Pisa, Piazza Torricelli 2, 56100 Pisa, Italy

Luciano RISTORI

INFN Sezione di Pisa, Via Vecchia Lavorosa 582a, 56010 S. Piero a Grado (PI), Italy

Received 24 October 1988

We discuss the architecture typical of high energy physics "machine" is implemented as "patterns". All the patterns are read out.

Stories on some technological innovations at CDF in the 1980s-1990s

1. Introduction

The quality of results from present and future high energy physics experiments depends to some extent on the implementation of fast and efficient track finding algorithms. The detection of *heavy flavor* production, for example, depends on the reconstruction of secondary vertices generated by the decay of long lived particles, which in turn requires the reconstruction of the majority of the tracks in every event.

Particularly appealing is the possibility of having detailed tracking information available at trigger level even for high multiplicity events. This information could be used to select events based on impact parameter or secondary vertices. If we could do this in a sufficiently short time we would significantly enrich the sample of events containing heavy flavors.

Typical events feature up to several tens of tracks, each of them traversing a few position sensitive detector layers. Each layer detects many hits and we must correctly correlate hits belonging to the same track on different layers before we can compute the parameters

2. The detector

In this discussion we will assume that our detector consists of a number of layers, each layer being segmented into a number of *hits*. When charged particles cross the detector they hit one bin per layer. No particular assumption is made on the shape of trajectories: they could be straight or curved. Also the detector layers need not be parallel nor flat. This abstraction is meant to represent a whole class of real detectors (drift chambers, silicon microstrip detectors etc.). In the real world the coordinate of each hit will actually be the result of some computation performed on "raw" data: it could be the center of gravity of a cluster or a charge division interpolation or a drift-time to space conversion depending on the particular class of detector we are considering. We assume that all these operations are performed upstream and that the resulting coordinates are "binned" in some way before being transmitted to our device.

We discuss the architecture of a device based on the concept of *associative memory* designed to solve the track finding problem, typical of high energy physics experiments, in a time span of a few microseconds even for very high multiplicity events. This "machine" is implemented as a large array of custom VLSI chips. All the chips are equal and each of them stores a number of "patterns". All the patterns in all the chips are compared in parallel to the data coming from the detector while the detector is being read out.

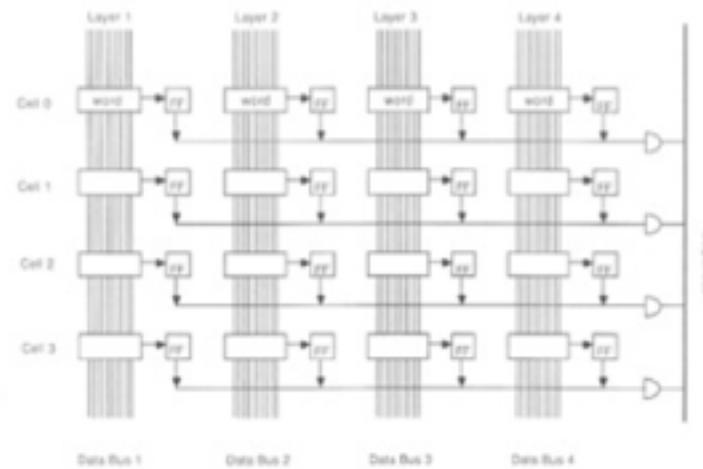


Fig. 3. Associative memory architecture.

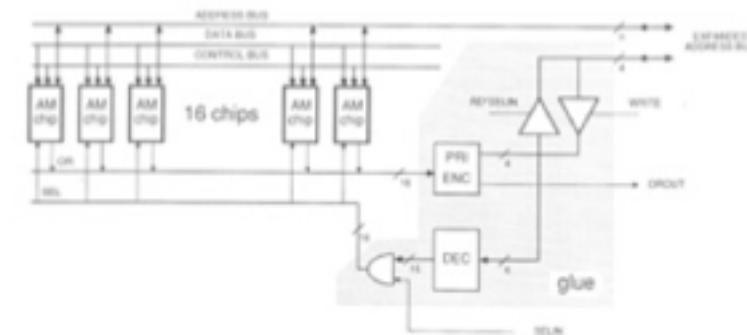
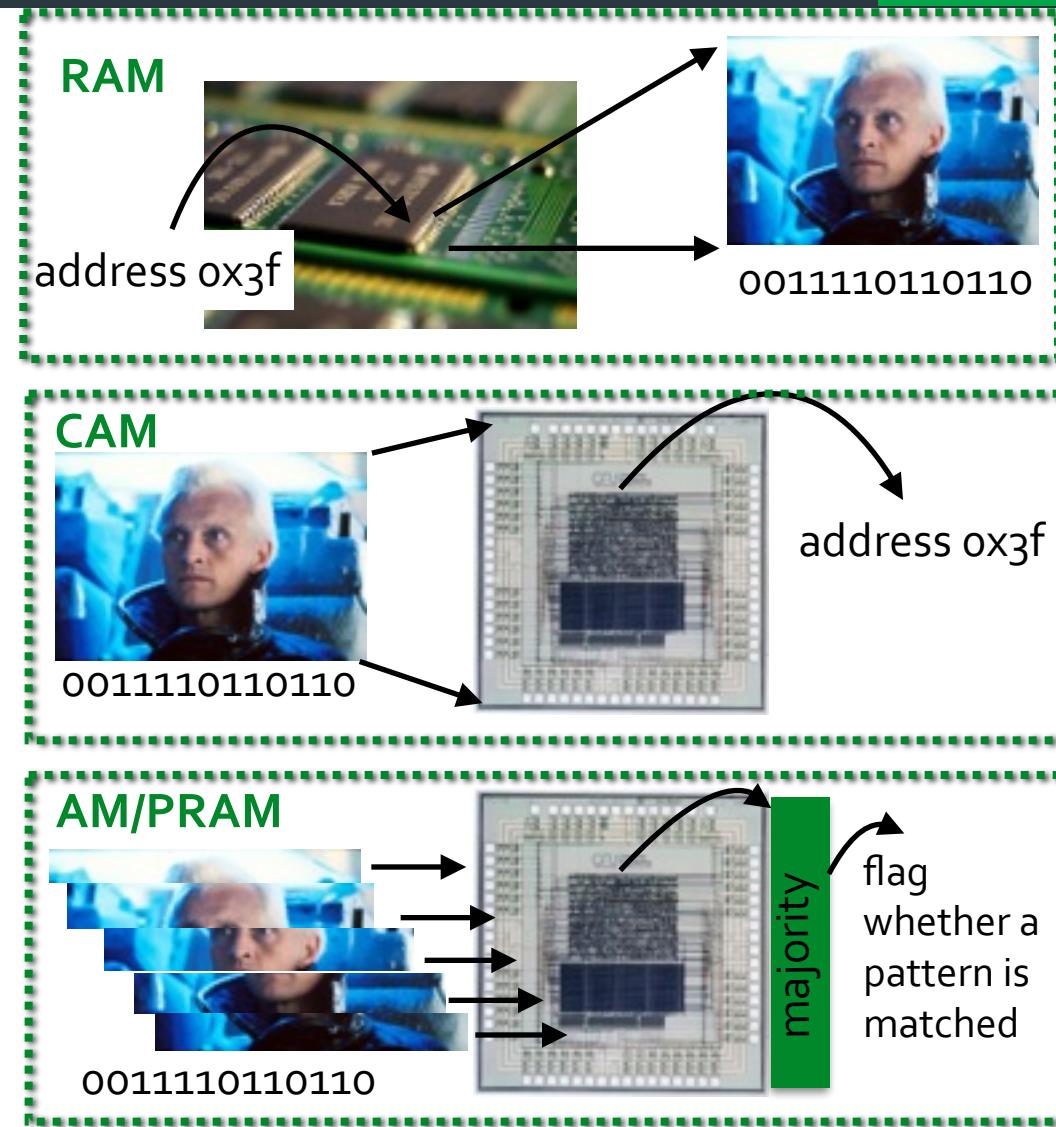


Fig. 5. 16 AM chips tied by the "glue".

Main ingredient: the Associative memory (AM)

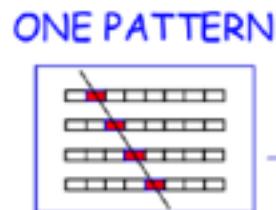
- RAM (Random-Access-Memory): memory address \rightarrow the data word stored at that address
- CAM (Content-Addressable-Memory): data word \rightarrow **searches its entire memory in one single operation** and return the address
 - much faster than RAM
 - commonly used for networking and computing: transform IP address, data compression, cache tag (parallel RAM access)
- AM or PRAM (Pattern Recognition Associative Memory) CAM based
 - Pattern recognition stops when all hits arrive
 - Use majority logic
 - Can be Ternary CAM: 3 states (1/o/x) with the addition of a "don't care" bit



AM based tracking system

- Dedicated device: maximum parallelism
- Each pattern with private comparator
- Track search during detector readout

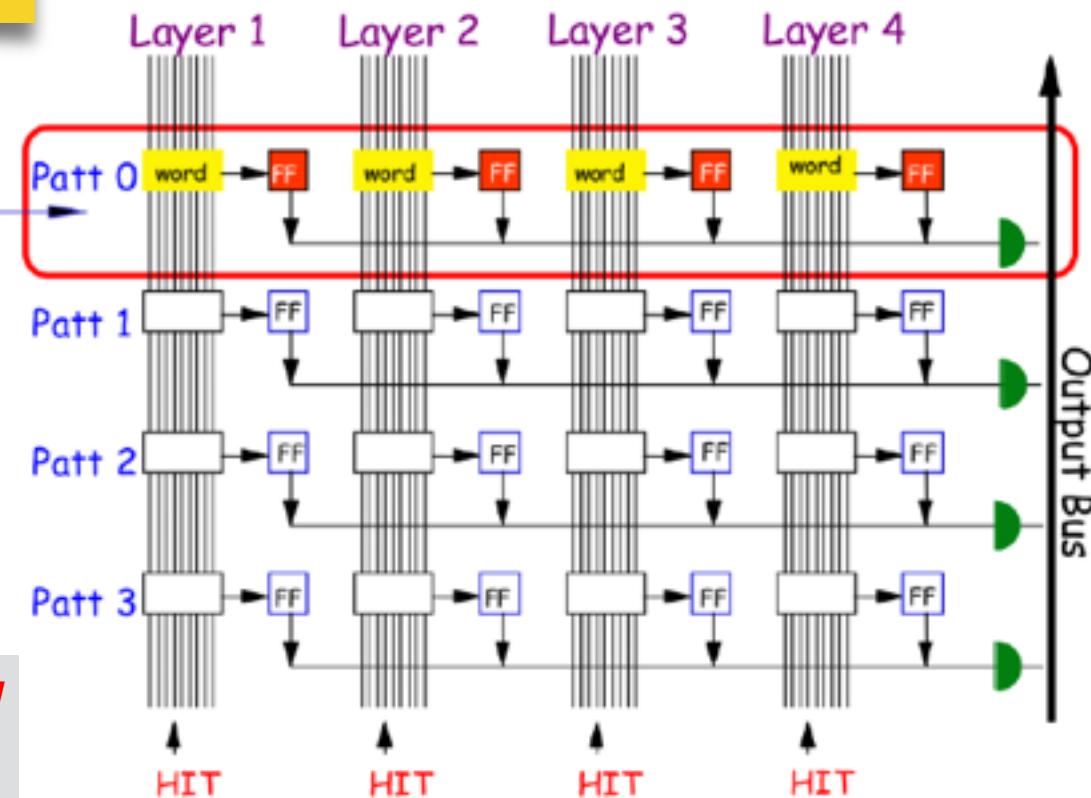
When a pattern is matched, the corresponding hits are selected for the following step



Inputs are compared to pre-calculated patterns of valid tracks originating from the interaction vertex

Pattern-matching done in few 10 ns!

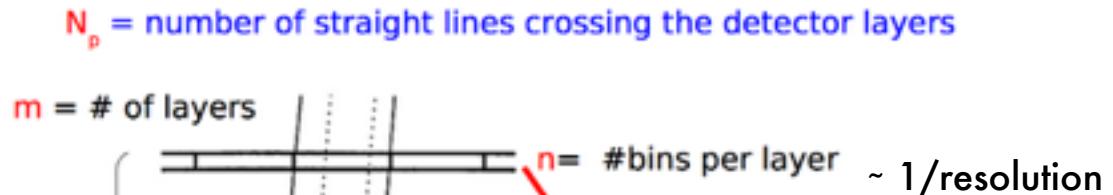
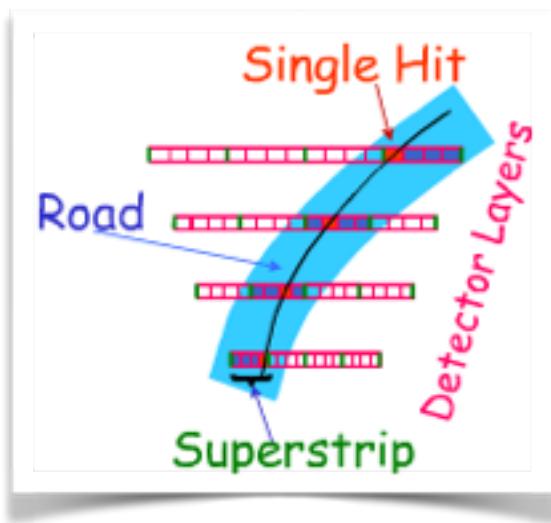
data readout, data distribution and data formatting takes longer.....



AM inputs are hits from different layers

Pattern bank & resolution

Finite number of patterns (pattern-bank): given finite resolution, different tracks generate the same pattern



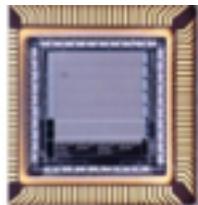
$$N_{\text{Pattern}} \propto N_{\text{Pileup}} \cdot \frac{1}{P_t} \cdot N_{\text{Layer}} \cdot \frac{N_{\text{Strip}}}{d_{\text{Strip}}^2}$$

FTK->L1TT x2-3 1->2 GeV 8->? X>6

- Higher resolution and rejection, if more patterns can be stored (and if more CAM cells/chips are available)
 - Preferred approach: 90% efficiency in a low fake scenario (to control the workload to the fitting step)
- To add flexibility, the resolution can be variable - with the use of "don't care" bits (Ternary CAMs)

AM evolution, to increase pattern density

SVT
AM chip



- 100 s) **Full custom VLSI chip** - 0.7 μ m (INFN-Pisa)
- 128 patterns, 6x12bit words each, 30MHz

Successfull! High pattern density, high speed and low power consumption



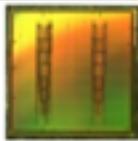
Alternative **FPGA** implementation of SVT AM chip

P. Giannetti et al., Nucl. Instr. and Meth., vol. A413/2-3, (1998)

G Magazzù, 1st std cell project presented @ LHCC (1999)

design with state-of-the-art technology

SVT upgrade

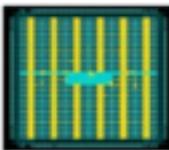


Standard Cell 0.18 μ m \rightarrow 5000 pattern/AM chip

SVT upgrade total: 6M pattern, 40MHz

A. Annovi et al., **IEEE TNS**, Vol 53, Issue 4, Part 2, 2006

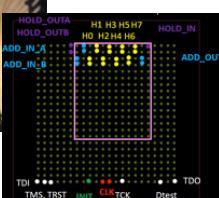
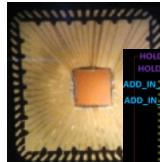
FTK R&D



AMchip04 –65nm technology, std cell & full custom, 100MHz
Power/pattern/MHz ~30 times less. Pattern density x12.
First variable resolution implementation!

F. Alberti et al 2013 **JINST 8 C01040**, doi:10.1088/1748-0221/8/01/C01040

SVT upgrade ready for LHC performance



FTK R&D in progress:

AMchip05: switched to serialized IO (11*2Gbs)

AMchip06 prototype: the FTK AM chip with 128k patterns/chip

New technologies for L1 Track?

Limits of the AM approach

- Performance fundamentally limited by Moore's Law
- AMChip near limit of conventional associative memory densities
 - Earlier studies demonstrated that ternary CAMs can be used with 10 billion patterns or more, doing a pattern lookup in < 200 ns

A challenge for HL-LHC

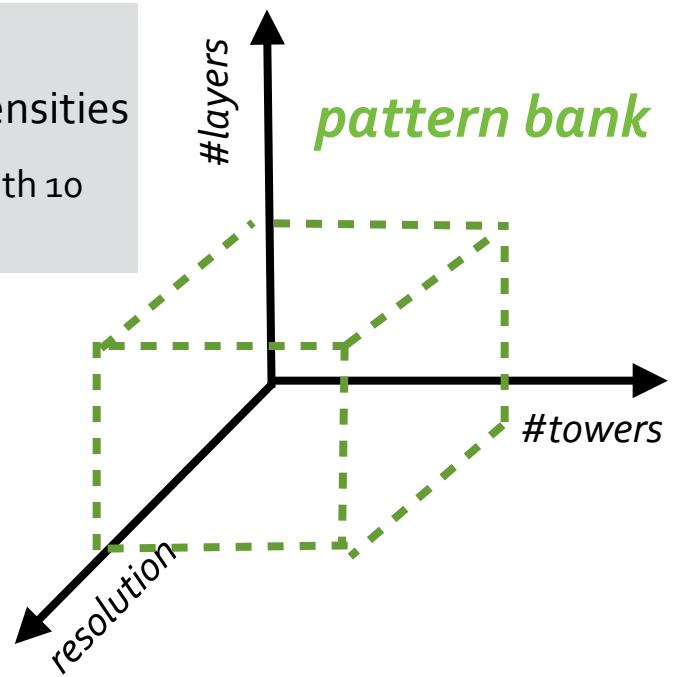
Increase the patterns density by **2** orders of magnitude

Increase the speed by a factor of $>\sim 3$

while keeping similar power consumption

or

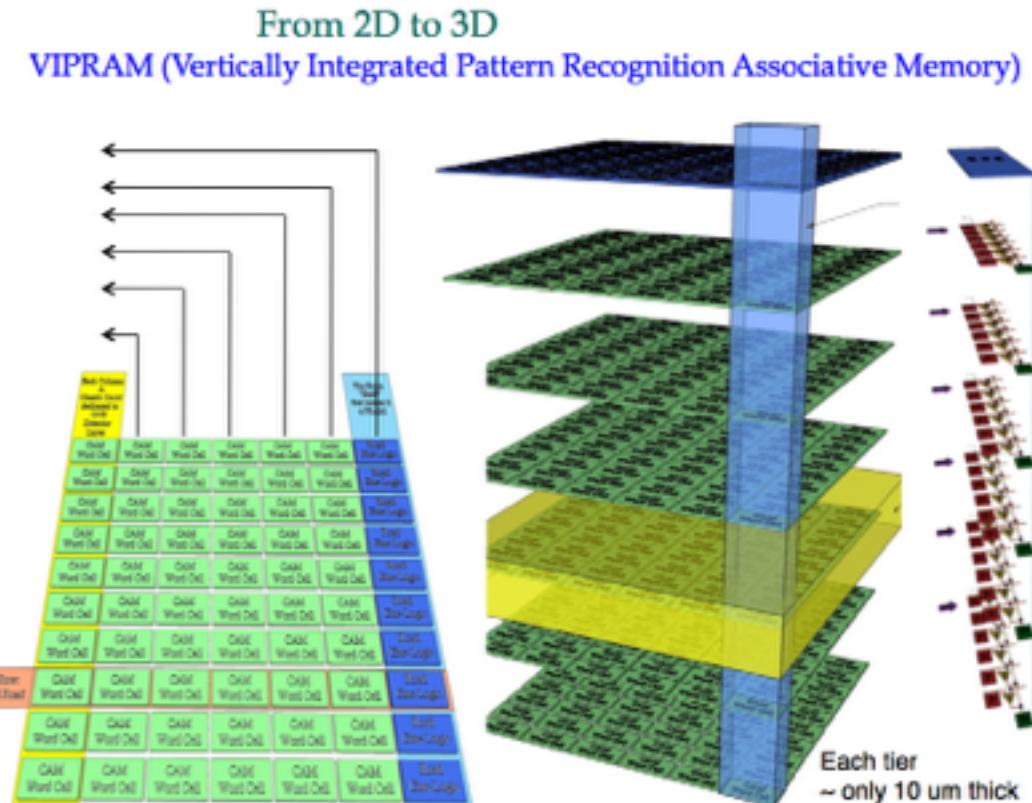
go to higher dimensions



AM evolution: 3D approach

VIPRAM: Fermilab project using 3D vertical integration technology (*TIPP 2011 pre-print*)

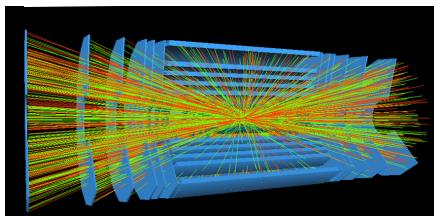
- One cell can process N layers in about one CAM cell size \Rightarrow density increased by N
 - 2D with 65 nm: $\sim 50K$ patterns/cm² (AMchip04)
 - 3D with 130 nm: $\sim 200K$ patterns/cm²
- Reduced connections \Rightarrow higher speed and less power density
- More flexible design



Physical detector layers \leftrightarrow silicon layers

Track fitting techniques

silicon detector



detector design
for triggering

data transfer

data formatting

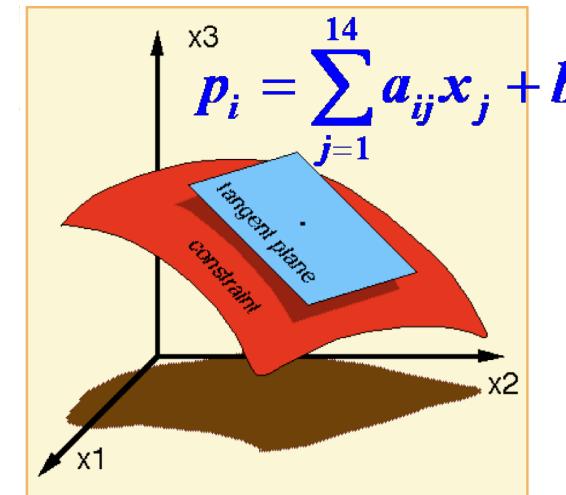
pattern-recognition

Selected R&D topics

track fitting

HLT

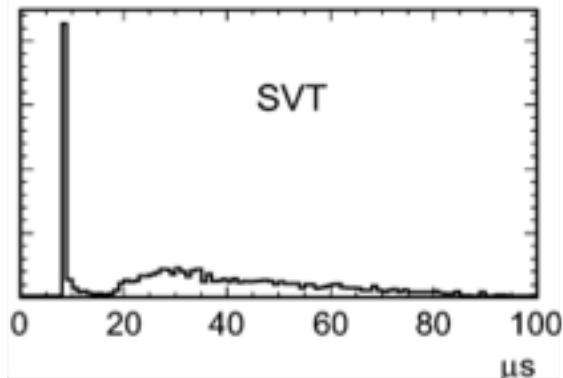
- Simple algorithm performed on any good combination of hits \Rightarrow can be massively parallelized
- Linear approximation on a limited region: get a set of linear equations (instead of solving helix) \rightarrow fast multiplications with pre-computed constants
 - Use of Look-up-tables (LUTs) with precalculated values (5 track parameters and the χ^2) stored in a table and interpolated



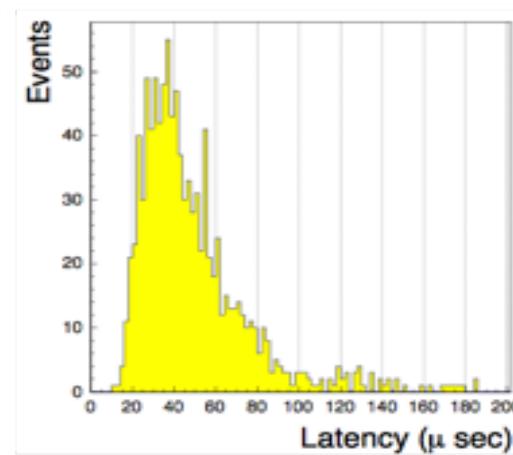
Due to short latencies and huge number of inputs, use of more complex algorithms, like Kalman Filters and Hough Transforms, (used for image processing) is limited

Track fitting technology evolution

CDF- SVX II



ATLAS FTK Run2-Run3



- Dedicated AM hardware combined with a dual-processor PC running an optimized Linux quasi-realtime kernel

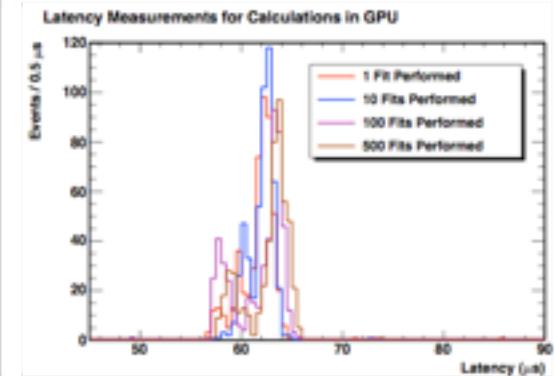
■ IEEETrans.Nucl.Sci. 53(2006)653–658.doi: 10.1109/TNS.2006.871782

- FPGAs w/ many Digital Signal Processors (DSPs):
→ ~ 1 fit/ns
- Constraints due to limited bandwidth and processing power
 - #AM patterns $< 16.8 \times 10^6$
 - #fits/event $< 80 \times 10^3$

time



ATLAS L1Track Run4

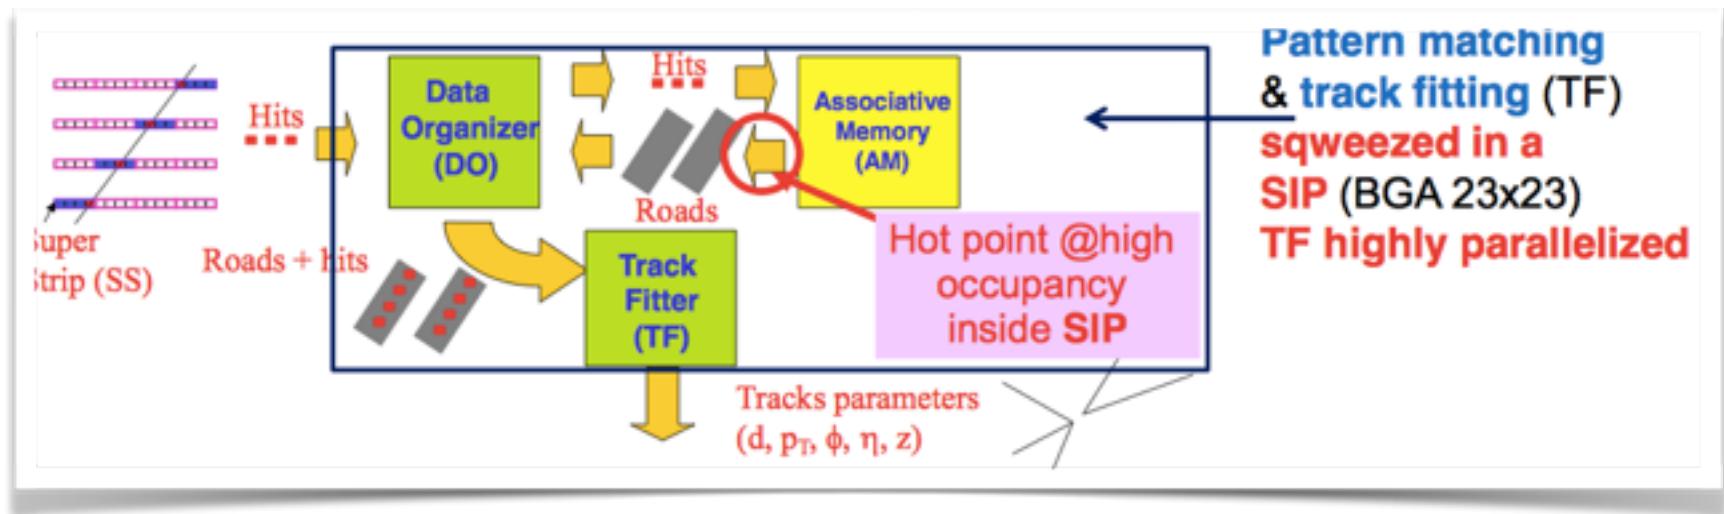


- GPUs is promising candidate: constant performance with increasing # of fits

■ Little is known about GPU performance, both in terms of speed and latency overheads, in low-latency environments

■ FERMILAB-CONF-11-710-PPD (2012)

Future: miniaturize, larger AM chips, integrate!



- If the AM stage and the Track Fitting can be integrated
 - latency is reduced
 - bandwidth is under control
- 3D Technology could help here (in the future)
- New generation of FPGAs with **stacked silicon interconnect (SSI) technology**: break through the limitations of Moore's law
 - Xilinx SSI technology

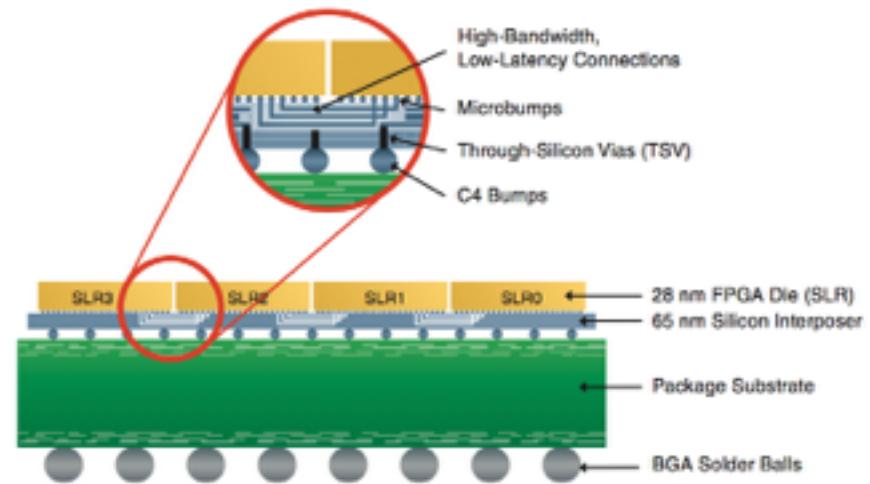
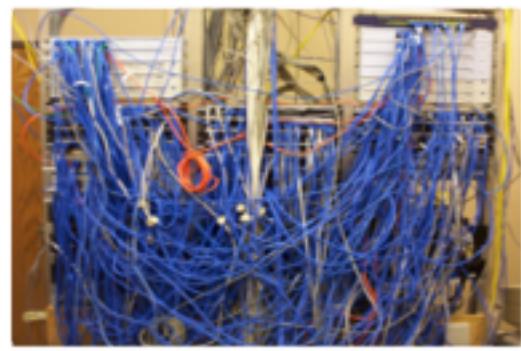


Figure 1: Virtex®-7 2000T FPGA Enabled by SSI Technology

Data sharing technique

Millions of channels to be readout.....

Past...

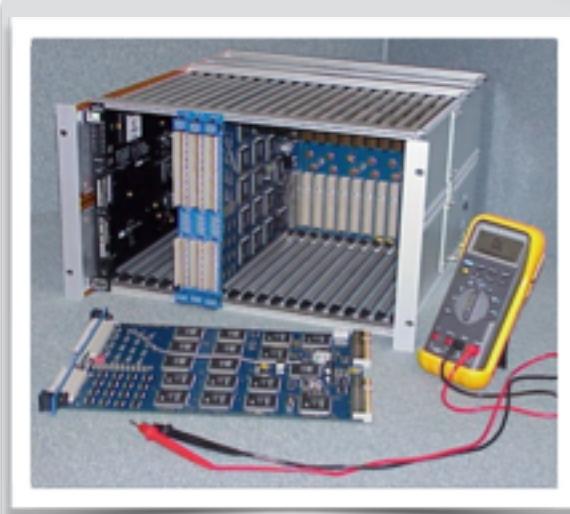


- ❖ Jumper cables

- ❖ Flexible, but ugly and difficult to maintain
- ❖ Still requires custom backplane

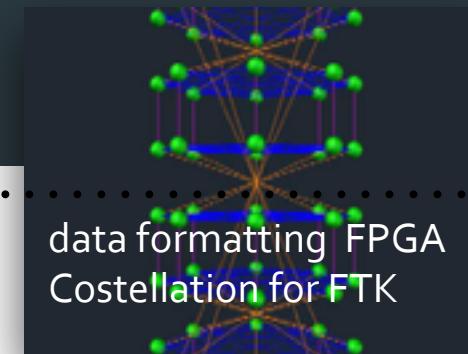
Data transmission technology advanced quickly

CDF- SVX II



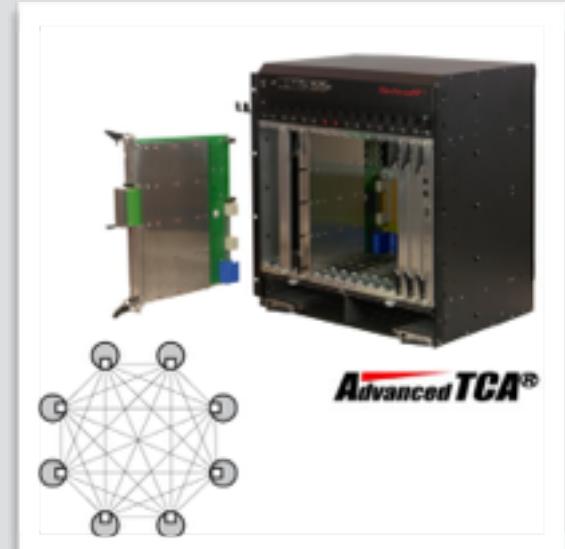
- ❖ Dedicated traces on the backplane

- ❖ Custom backplane
- ❖ Each crate may be different
- ❖ Inflexible design



ATLAS FTK Run2-Run3

- ❖ Modern ATCA with full-mesh



ATLAS L1Track Run4

- ❖ Patterns ~ Billion / crate/shelf

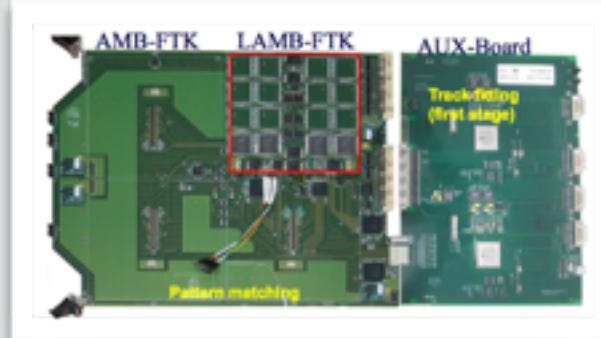
How could the L1Track may appear?

CDF- SVX II



- CDF original SVT system had ~400K patterns total: 128 patterns per AMchip
- Test state-of-the-art CAD tools
- Commissioned around ~2001

ATLAS FTK Run2-Run3



- 16400 AM chips + 2000 FPGAs @ 100 MHz for 16-bit words (2 Gbs)
- #AM patterns < 16.8 millions, with variable resolution
- Schedule:
 - Integration with limited coverage in Run2 (2015)
 - 2016: full coverage

ATLAS L1Track Run4



- Aim to reach ~500K patterns/cm² for VIPRAM chip
- Or other technology? GPU?
- Schedule: ready for 2022

Outline

- ▣ I'm tempted to say: There are no conclusions, future is open
- ▣ L1Track project will deal with the possibility of triggering in HL-LHC
 - ▣ Need to weigh potential physics gains against added material and possible cost for the trackers
 - ▣ A lot of demanding (electronics) developments
 - ▣ The R&D programs started some years ago
- ▣ Old stories from the past can help us in seeing how it could be, if much effort is concentrated in understanding the requirements
- ▣ We must maintain wide open sight to what we can steal from the technology market, that has somehow similar demands on large data-processing, in short time, on large systems.... Steal from your cell-phone!

References

- ▣ ECFA workshop in Aix-les-Bains
- ▣ Review of LHC & Injector Upgrade Plans Workshop (RLIUP)
- ▣ References for tracker evolution
 - ▣ Strip CMOS task force
 - ▣ High-performance Signal and Data Processing Workshop 2014
 - ▣ WIT workshop on intelligent trackers
 - ▣ 2012: <https://indico.cern.ch/conferenceTimeTable.py?confId=154525#all.detailed>