EPS Conference, July 12th 2019, Ghent (Belgium)

THE ATLAS HARDWARE TRACK TRIGGER DESIGN TOWARDS FIRST PROTOTYPES

F.Pastore (RHUL)

on behalf of the ATLAS TDAQ Collaboration

ATLAS & HTT COLLABORATION

Status: April 2019

HTT Collaboration

INFN Pisa
INFN Bologna
INFN Pavia
INFN Milano
U. Melbourne
U. Thessaloniki
U. Uppsala
U. Bergamo
LPNHE Paris
UC London
RHUL, London
UC Irvine
U. Illinois
U. Geneva
U. Chicago
U. Heidelberg
Niels Bohr Institute
Argonne Nat. Lab.
U. Arizona
U. Ohio State
LIP - Portugal
U. Sussex (UK)
U. Pennsylvania

ATLAS Collaboration

183 institutions (232 individual institutes) from 38 countries

ATLAS EXPERIMENT

LHC BECOMING IMPRESSIVELY LUMINOUS

European Council (2014): "CERN is the strong European focal point for particle physics in next 20 years"

LHC / HL-LHC Plan

HiLumi HL-LHC PROJECT

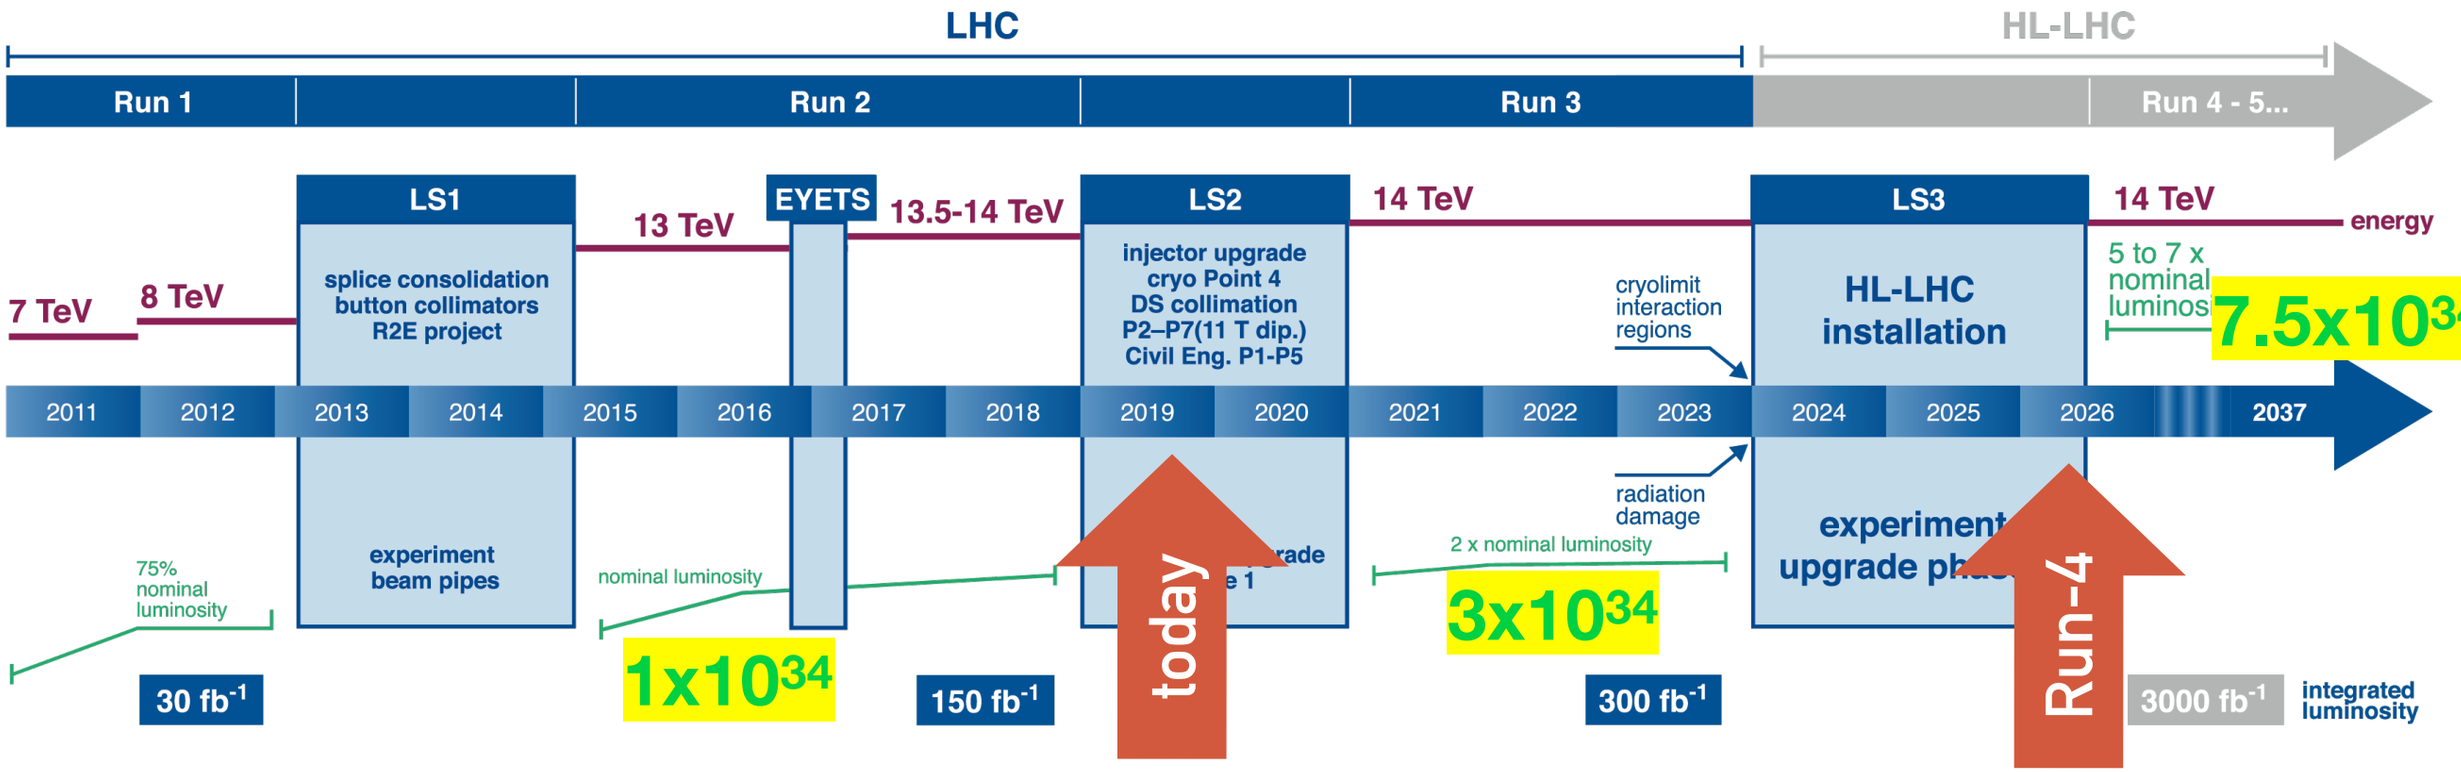

- **High-Luminosity LHC project will increase LHC peak Luminosity ~ x 7.5**
  - 3000 fb^{-1} of data at a centre of mass energy of 14 TeV
  - **ATLAS** detector requirements will go beyond original design specifications
- **Exploit full potential of upgraded detectors and large data sample for EW and QCD precision measurements, Higgs boson properties, flavour physics, BSM searches**

ONE EVENT AT HIGH-LUMINOSITY (L=7.5X10^{34} /CM^2/S)

ATLAS EXPERIMENT

HL-LHC $t\bar{t}$ event in ATLAS ITK
at $\langle\mu\rangle$=200

- ➡ 200 pile-up collisions per bunch crossing (every 25 ns), compare to 40 in Run2
- ➡ ~ 10 000 particles per event
- ➡ Mostly low p_T particles due to low transfer energy interactions

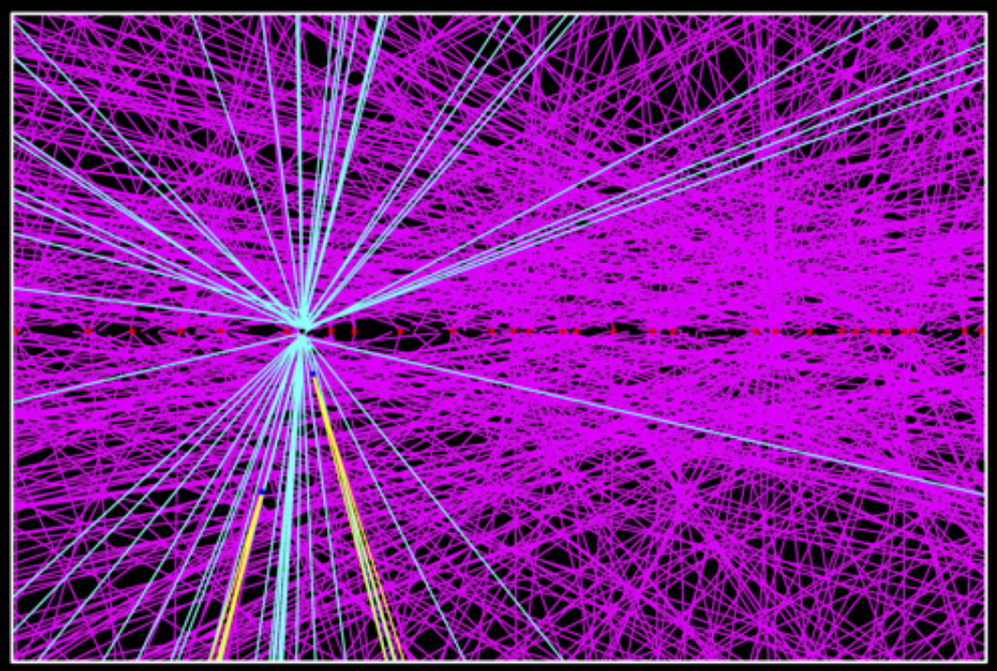

Design Luminosity x 7.5

ATLAS TRIGGER STRATEGY AT THE HL-LHC

- **Maintain physics acceptance at EW scale**
  - focus on high-p_T leptons
- **Accepting higher rates…**
  - Level-1 trigger: 0.1 ➠ 1 MHz
  - DAQ throughput: 1 ➠ 50 Tbps
- **…and reduced rejection due to high pile-up**
- **Trigger rejection (after first-level trigger) can be increased by using tracks reconstructed in the (new) full-silicon Inner Tracker (ITk)**
  - Better p_T resolution rejecting low-p_T leptons
  - Identification of primary vertex will help selections of hadronics (jets/MET) and multi-objects from the same interaction

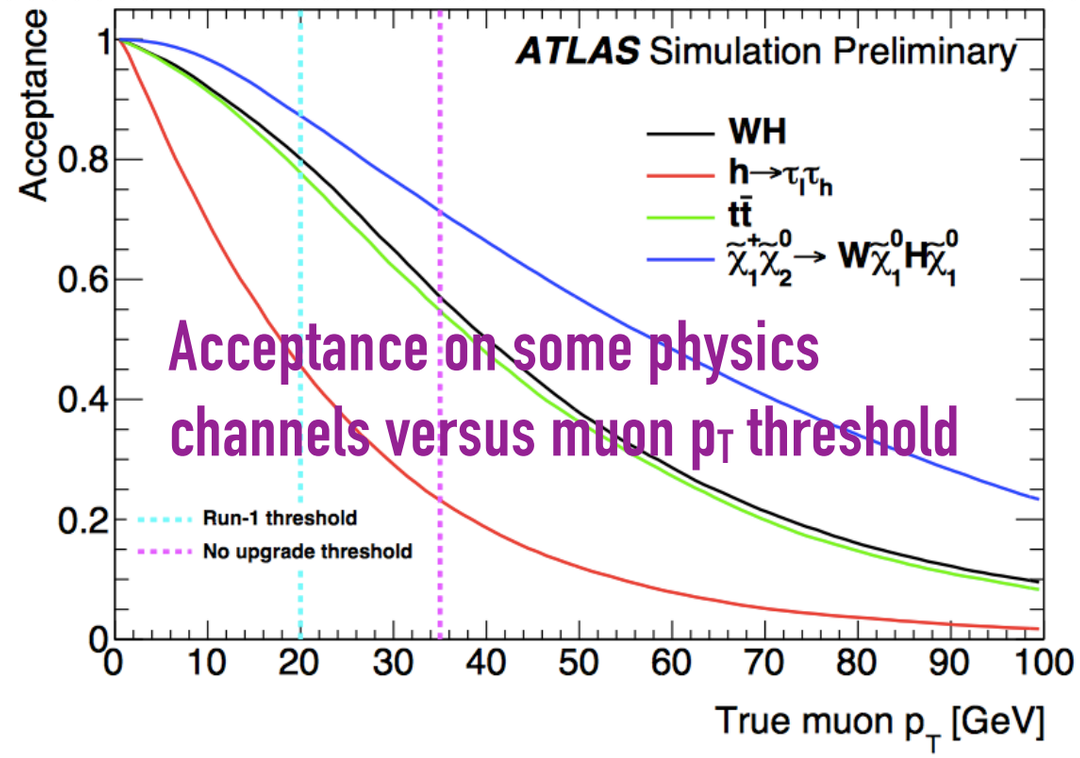

dataflow and processing time do not scale linearly with Luminosity

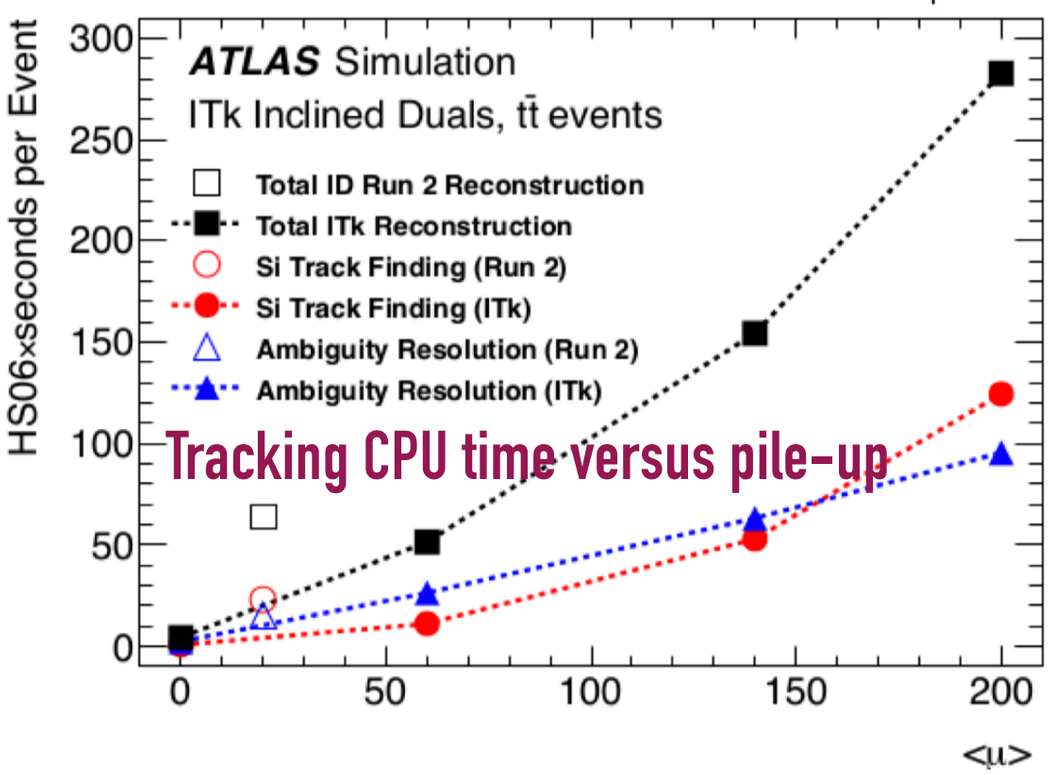

Luminosity x10, complexity x100: a hardware tracking can help!

THE HARDWARE TRACK-TRIGGER PROJECT (HTT)

Baseline scenario

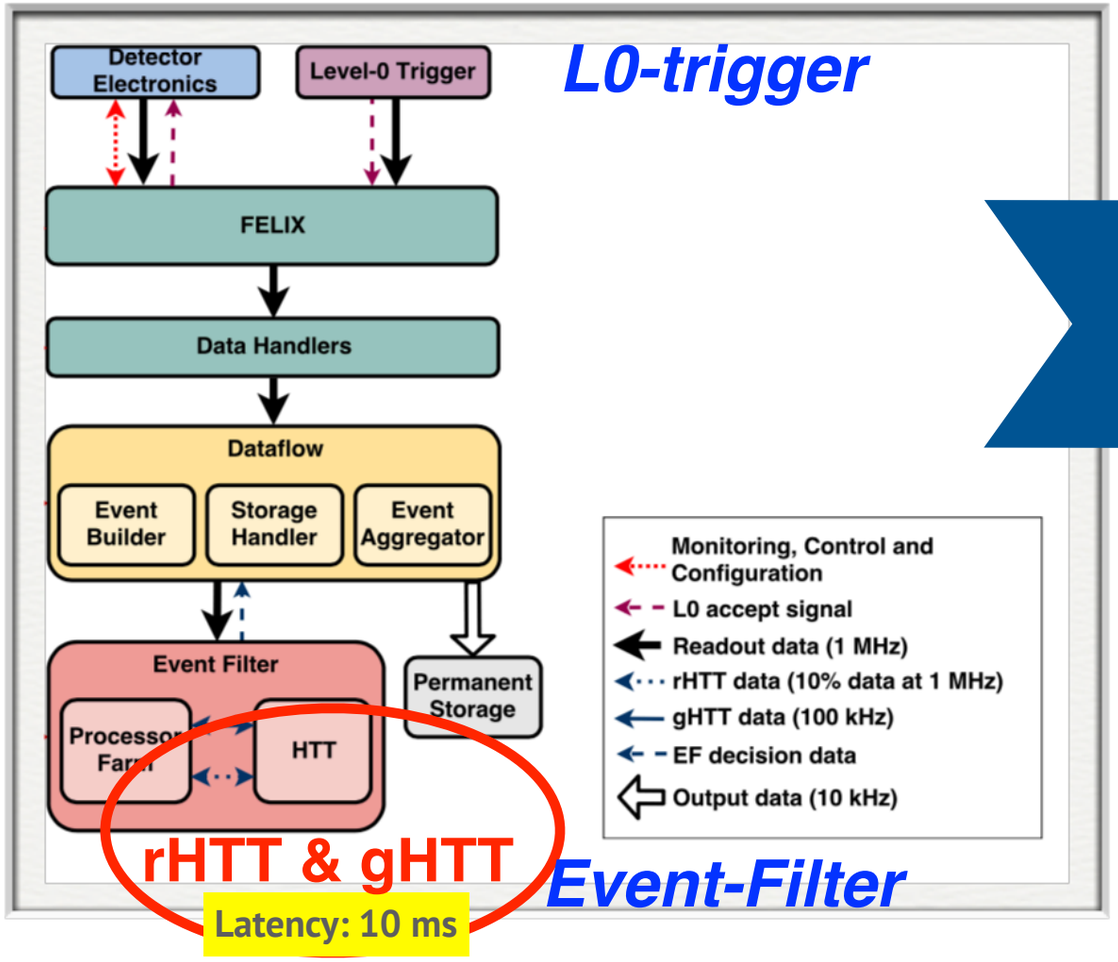

Evolved scenario

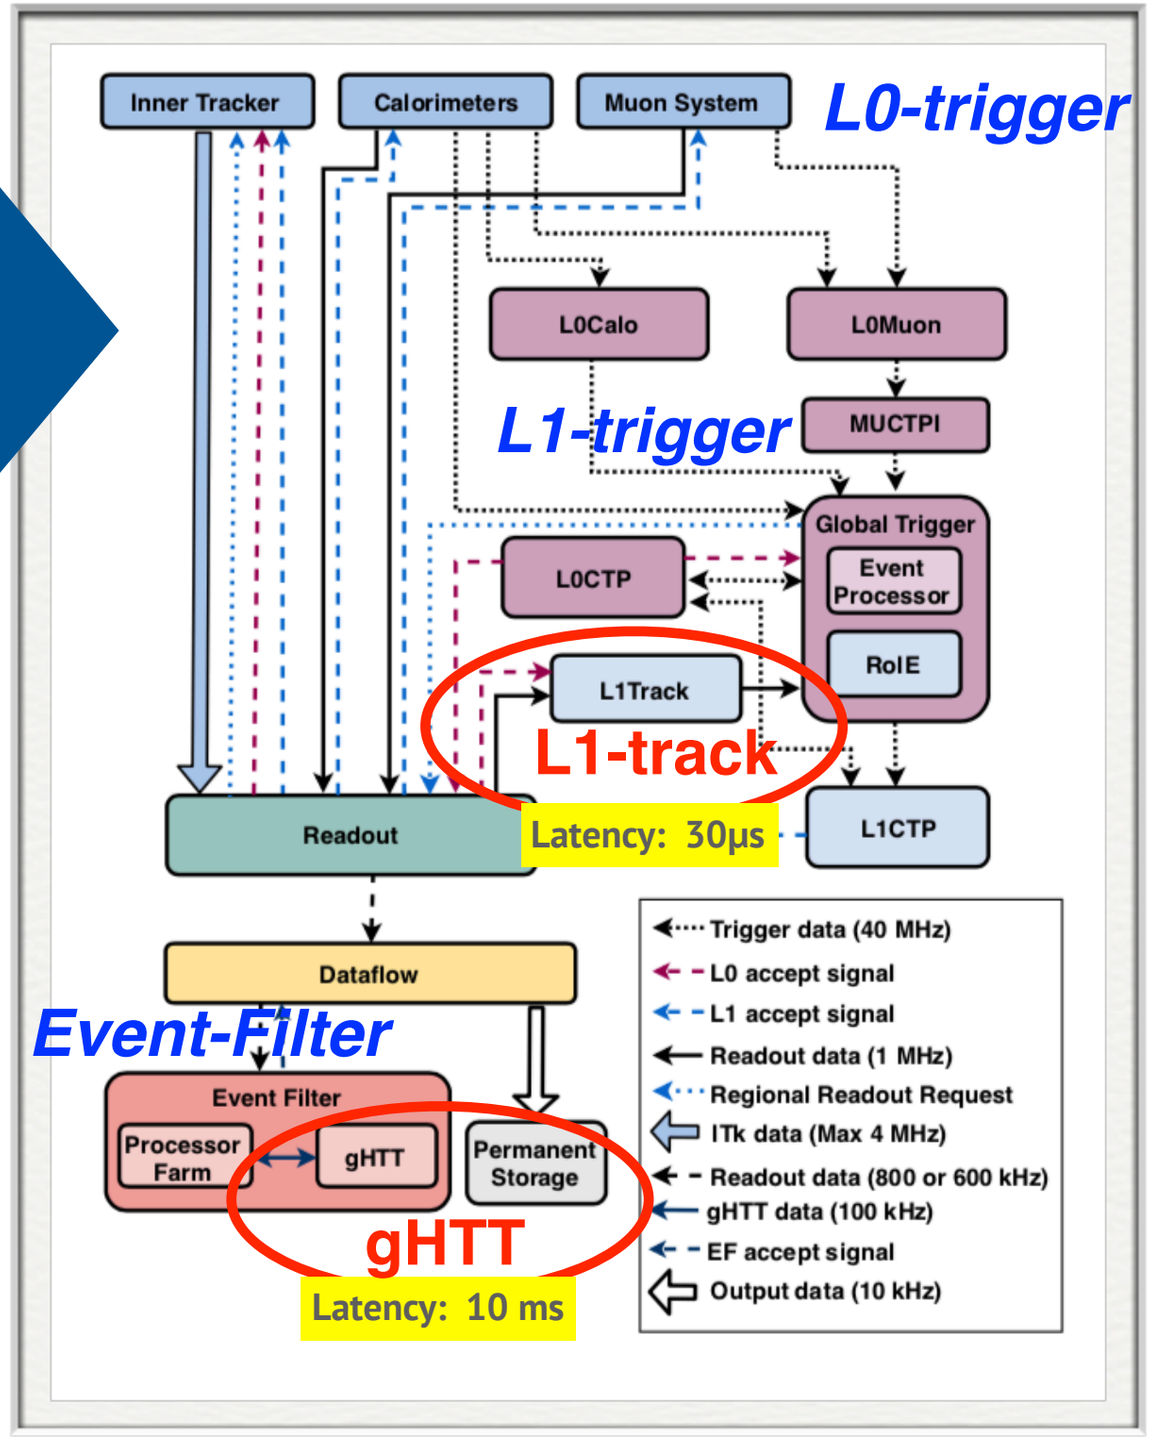

- Massively parallel, with O(500) boards @1-4 MHz Level-0 input rate
- Two basic T/DAQ architectures, exploiting both regional (rHTT) and global (gHTT) tracking
  - Designed on Baseline, including needs of Evolved
  - rHTT latency <30 μs in the evolved scenario

With HTT as co-processor, the Event Filter farm can be reduced by x10 in size

ATLAS AS PIONEER OF HW TRACKING PROCESSORS

CDF- SVX II

ATLAS FTK Run2-Run 3

ATLAS HTT Run 4

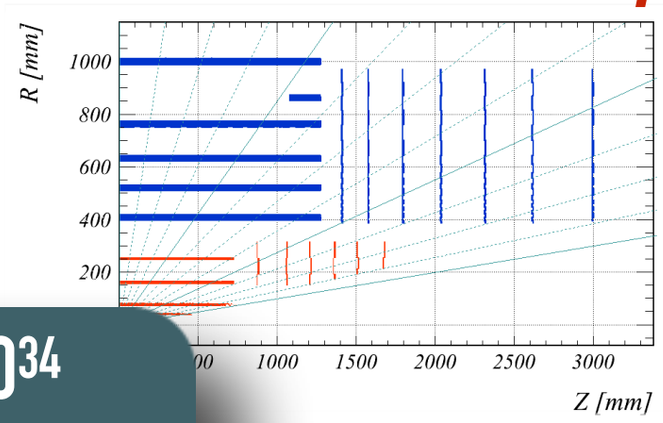

Luminosity:	$3x10^{32}$	→	$2/3x10^{34}$	→ $7x10^{34}$
L1 rate:	30 kHz	→	100 kHz	→ 1 MHz
Tracker channels:	0.2M	→	100M	→ 800M
Crates:	10 VME	→	13 ATCA	→ 50 ATCA

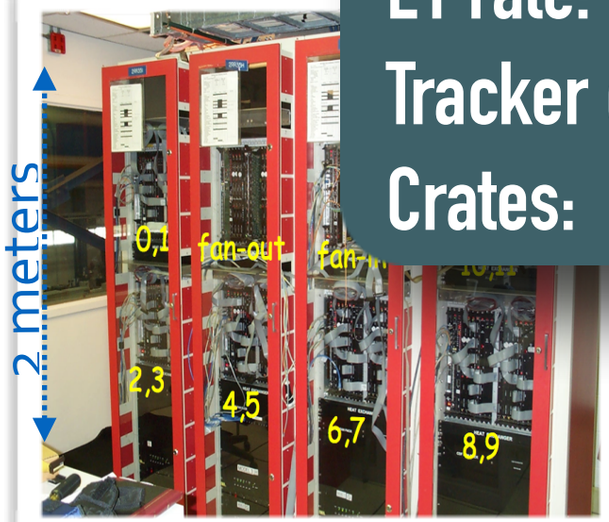

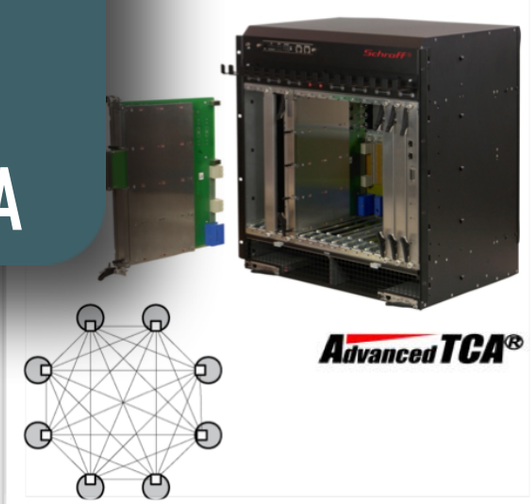

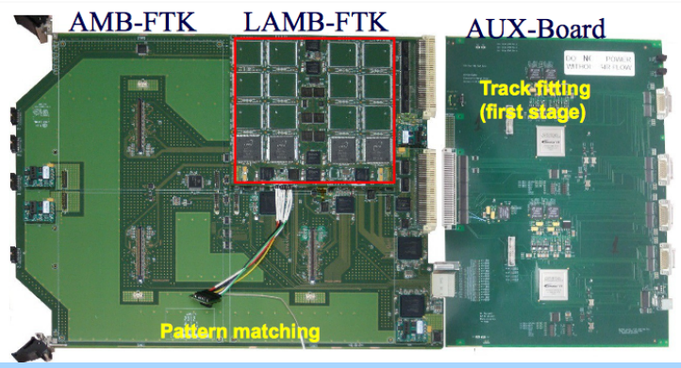

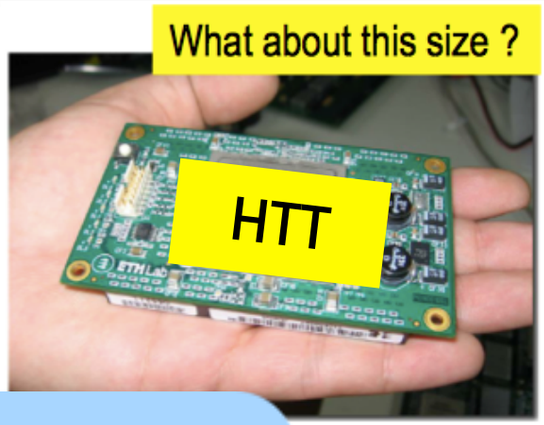

HTT is based on current **ATLAS FTK** system

HTT PROCESSING AND EXPECTED PERFORMANCE

- **Pattern-recognition in AM ASICs (rHTT)**
  - pre-calculated patterns of 8-layer hits
- **1st stage track-fitting in FPGAs (rHTT)**
  - 8 outer layers, performances limited by short lever arm and distance to impact point
    - Based on Principal component analysis (j.nima. 2003.11.078, and H. Wind, CERN-EP-INT-81-12-REV, 1982)
- **2nd stage track-fitting in FPGAs (gHTT)**
  - all 13 layers, quasi-offline resolution

Associative Memories (AM)

combinatorics (10^4 hits/BC) reduced with pattern-recognition

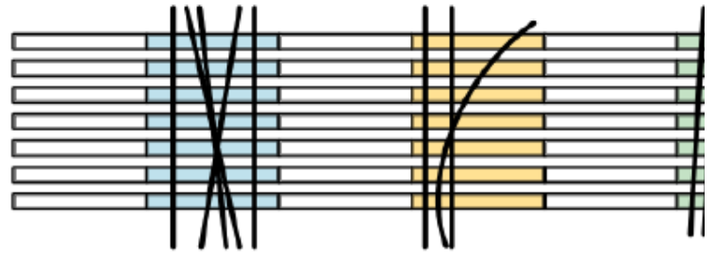

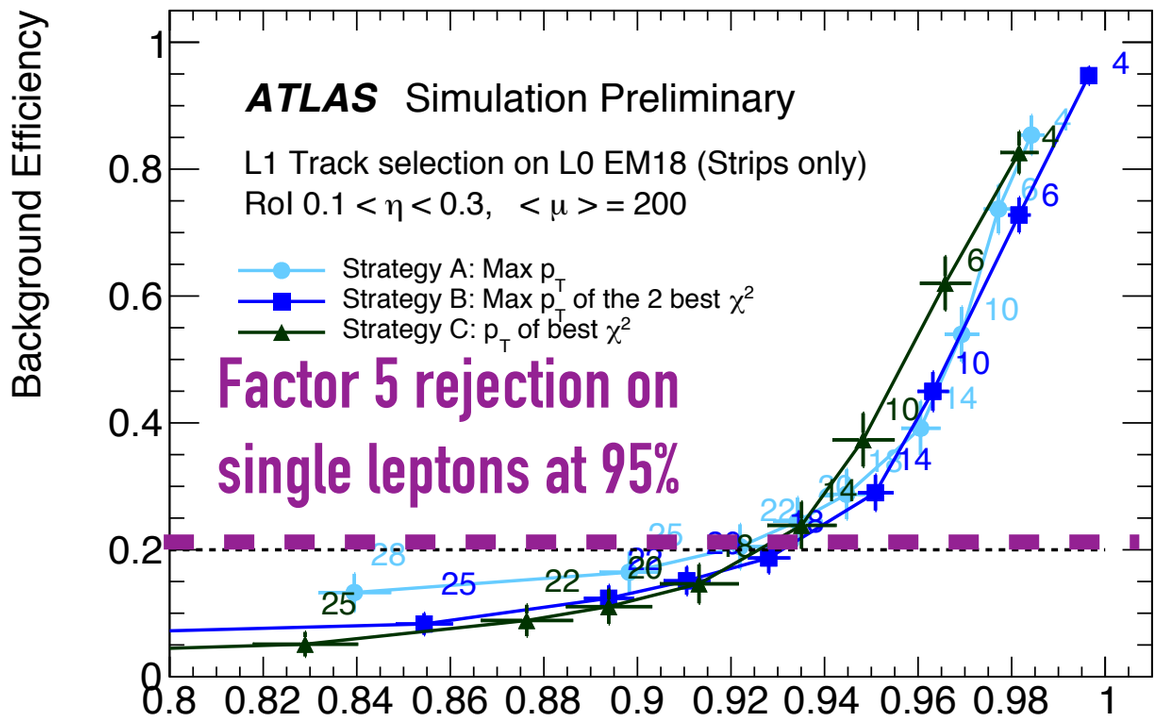

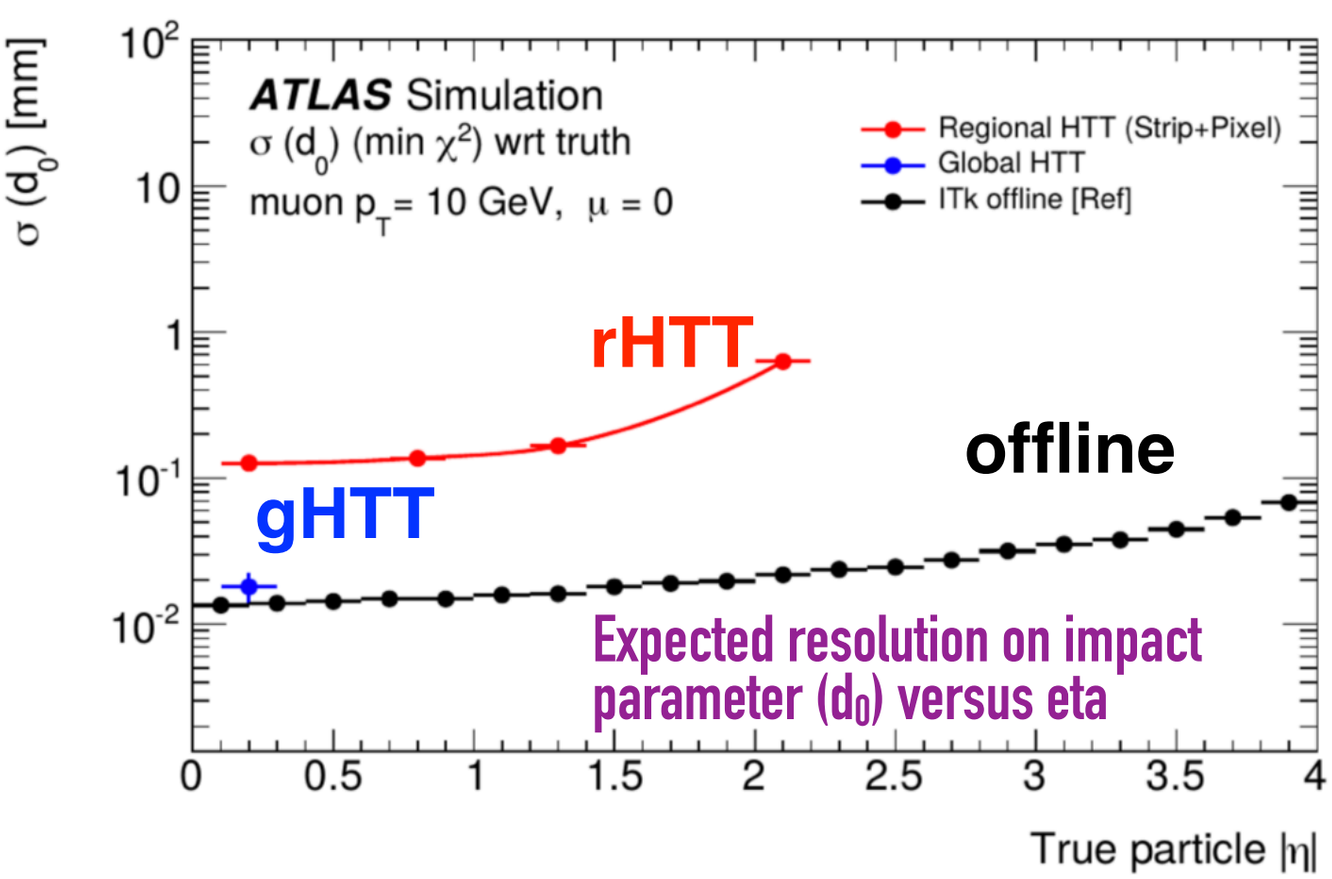

- **Tracking efficiency**
  - ~100% flat over the full p_T range for muons/pions while drops at $p_T < 10$ GeV for electrons because of Bremsstrahlung
  - At large eta, performance on electrons worsens due to increase in material

HTT OVERVIEW

- **Basic design principles**
  - Do not interfere with the design of the new silicon tracker detector
    - Only ensure fast readout on the front-end (for 30 μs latency)
  - Transparently respond to both regional and global readout requests
  - High modularity: different configurations reuse the same hardware

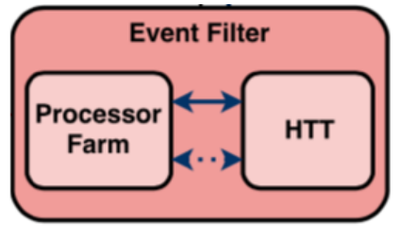

- **Boards grouped in HTT-units**
  - Interface to the Event Filter (CPU farm) via dedicated servers (**HTTIF**)
- **All boards based on the same Trigger Processor (TP) motherboard**
  - which only differ by mounting different mezzanines
- **Two types:**
  - Associative Memories TP (**AMTP**)
  - Second Stage TP (**SSTP**)

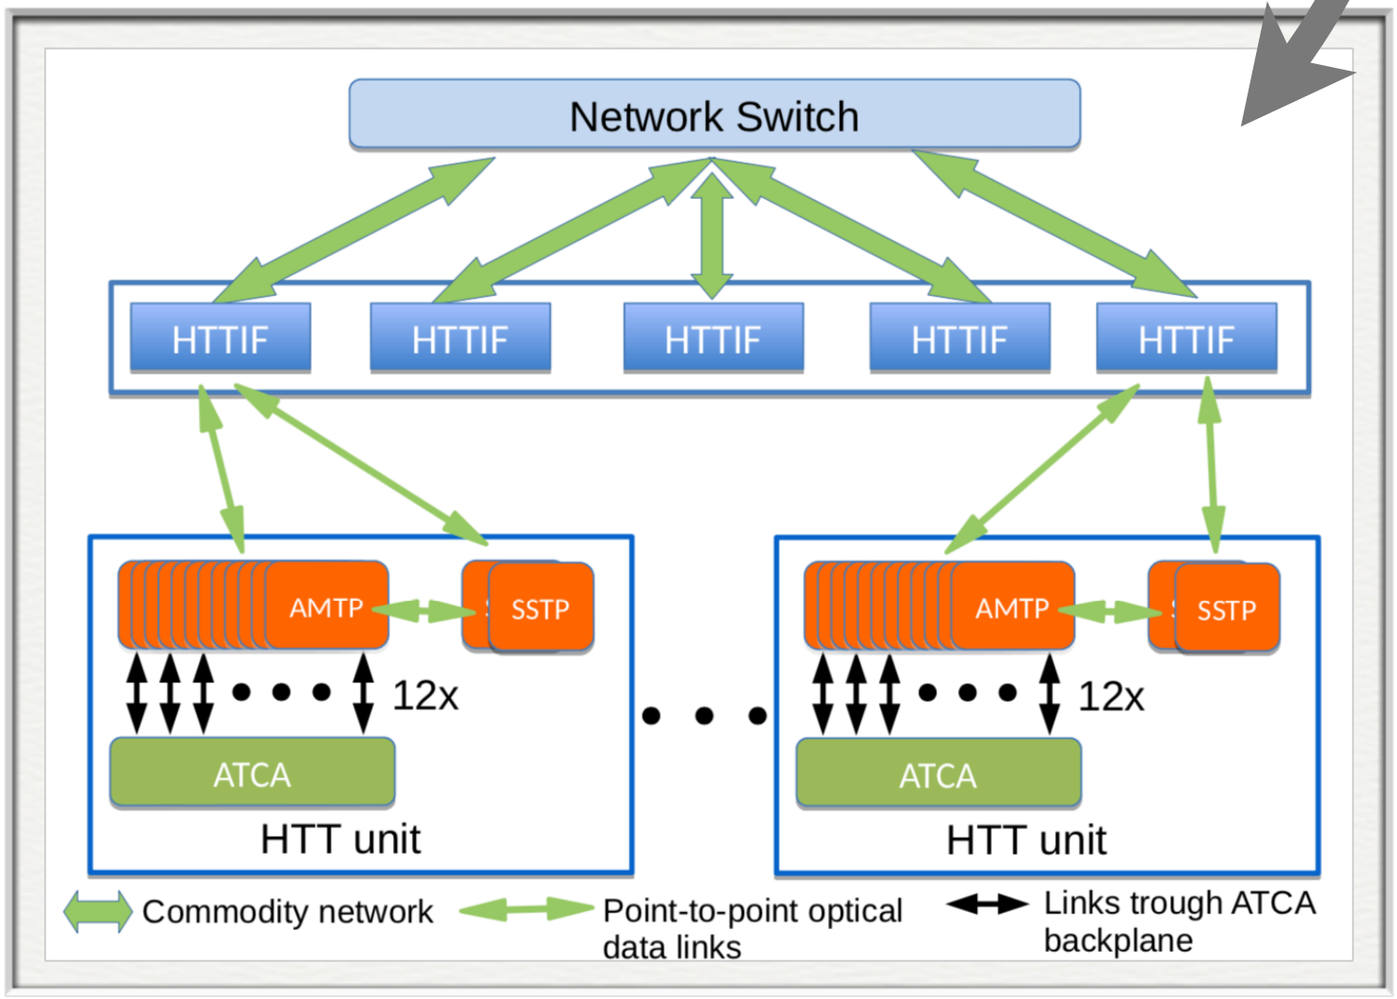

Project status:

Specifications are under finalisation, going towards first demonstrators

TRACKING PROCESSOR BOARD (TP)

- **High-Bandwidth motherboard with:**
  - 2 small mezzanines or 1 large mezzanine
  - 10 Gbps links
  - Large FPGA
  - SoC to support DCS and monitoring
  - ATCA support
- **Functions:**
  - data communication to/from **HTTIF** and **mezzanine** cards
  - data sharing
  - specific algorithms (Pixel clustering, duplicate track removal)

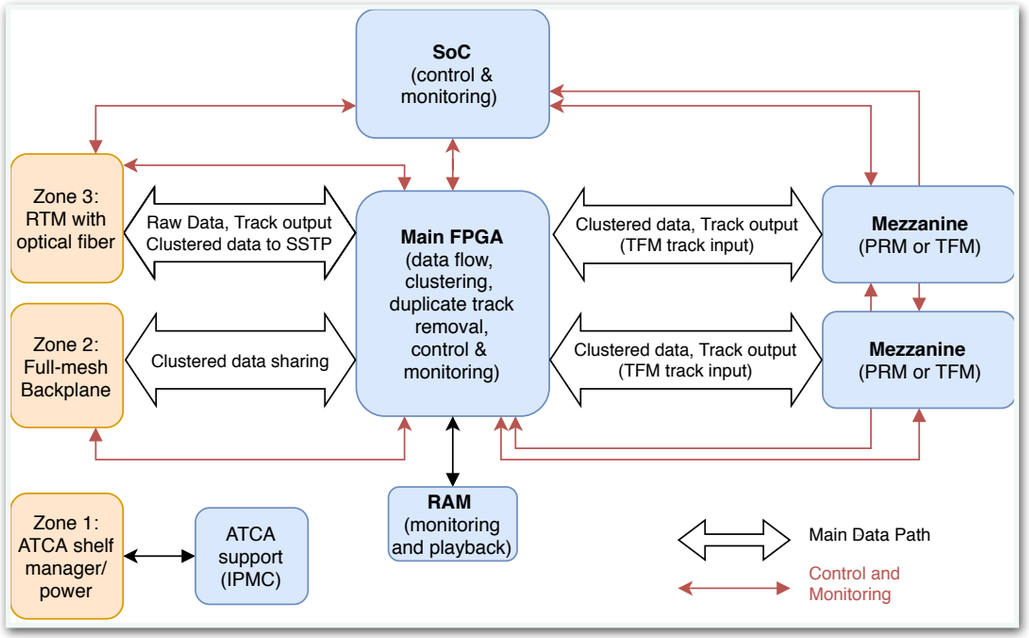

Firmware blocks are designed

Main Layout of motherboard

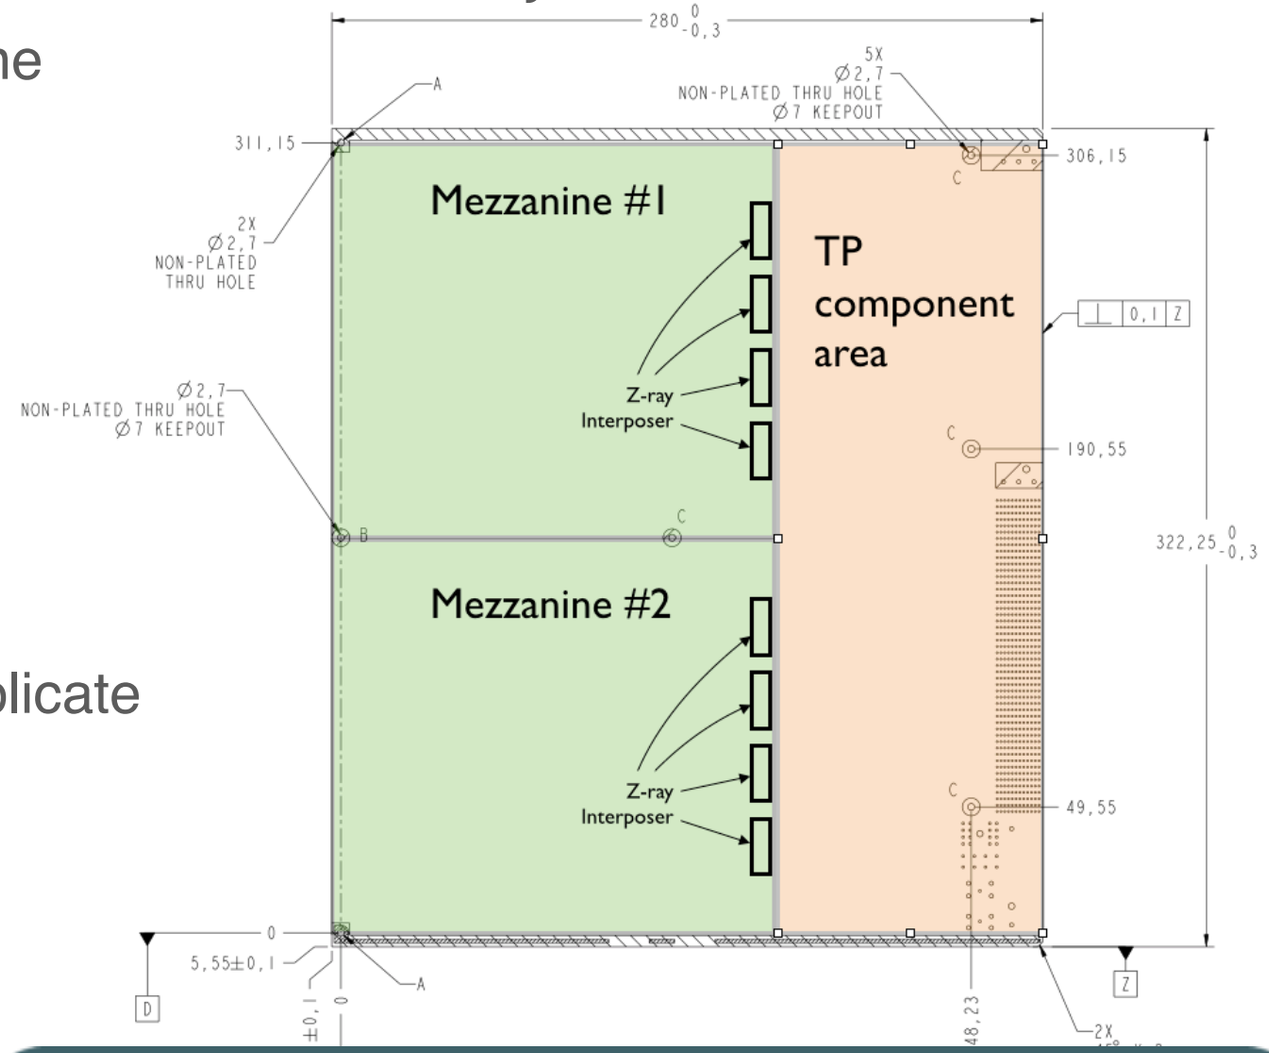

- **Demonstrator ready by October 2019, to validate**
  - Z-ray connector for mezzanines
  - Thermal and mechanical model
  - High-speed links
  - Readiness for integration tests and cooling tests at CERN (August 2020-)

HTT CORE COMPONENT: THE ASICS AM09

- **Low-power CMOS device for bit-wise comparison of incoming data (clustered hits) with pre-stored patterns**
  - About 30 peta comparisons per second per chip
  - Low power and small size allow high density of chips on board

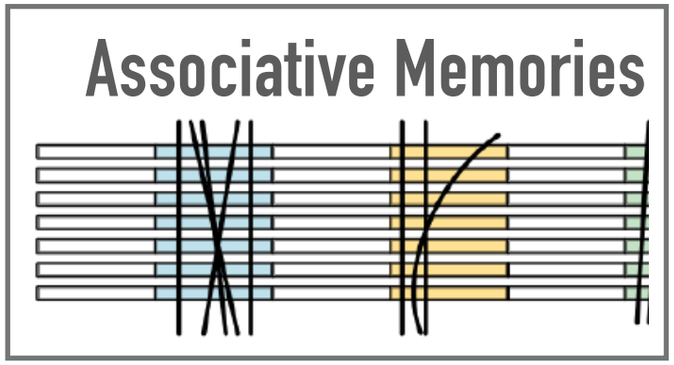

- **AM09 is chosen as production chip for HTT, evolution of FTK AM06, improving # of patterns, I/O bandwidths, power**
  - Power driven by bit-comparison: *P = 1W+ < inputRate >*0.05W/MHz*
    - 16 bits x 50% bit-flip for data @ 50 MHz, on 8 buses ==> **2.5 W**
  - 8 inputs (one per **ITk** layer) at 1 Gbps (LVDS electrical protocol)

ASIC	year	technology	patterns	clock	power	cell type
		[nm]	[k]	MHz	[fJ/comp/bit]	
AM06	2015	65	128	100	1,11	XORAM
AM07	2017	28	16	200	0.75	DOXORAM, KOXORAM
AM08	2020	28	16	250	0.42	KOXORAM+
AM09	2021	28	3x128	250	0.42	KOXORAM+

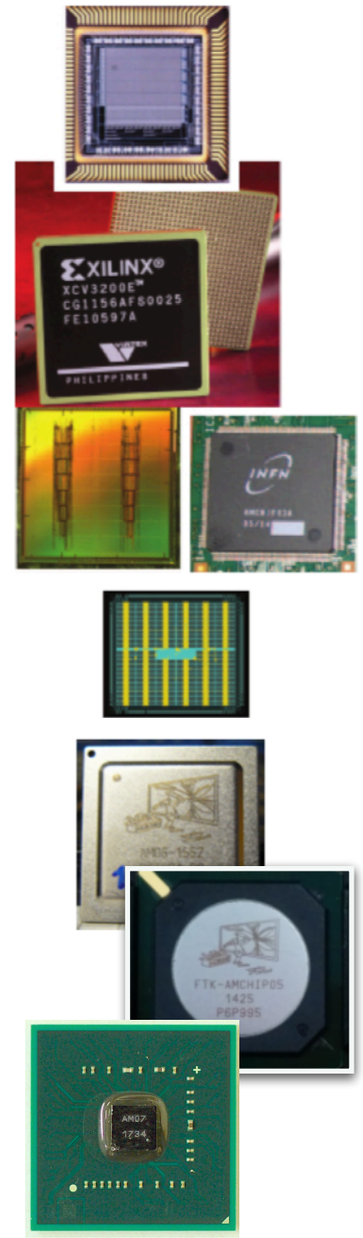

Current understanding of **AM07** tests and **AM08** simulations shows that **AM09** will be within the allocated budget both in terms of area and power consumption

PATTERN RECOGNITION MEZZANINE (PRM)

- **Single board mezzanine on AMTP board**
  - 4 Blocks of 6 **AM-ASICs** (9.2M patterns/**PRM**)
    - peak cluster rate/layer @250MHz
  - One large **FPGA** for data sharing and 1st stage track fitting
    - ~1GHz fits/board
  - High Speed Lines @10 Gbps

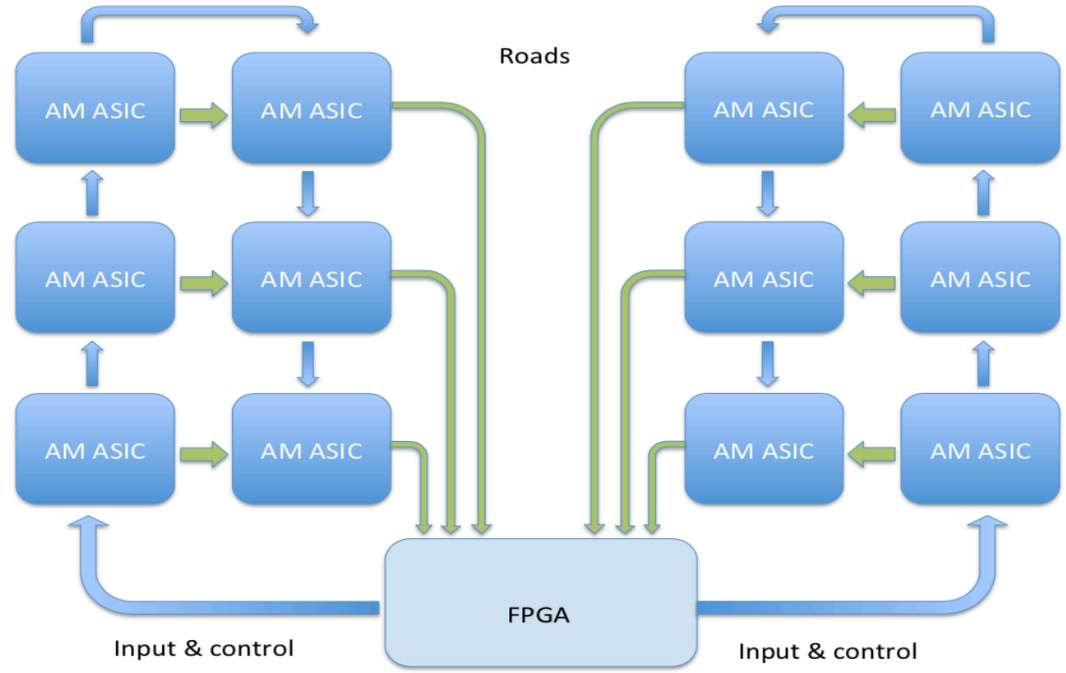

- **Demonstrator ready by May 2020**
  - ongoing firmware development, power scheme and selection of components
- **Number of ASICs/board limited by power and space density on board**
  - evaluating to reduce from 24 to 20 AM/PRM
- **Choice of the FPGA**
  - candidate (Intel® Stratix® 10) with High Bandwidth Memory (HBM) for accessing fit constants
    - investigating other track-fit techniques (here)

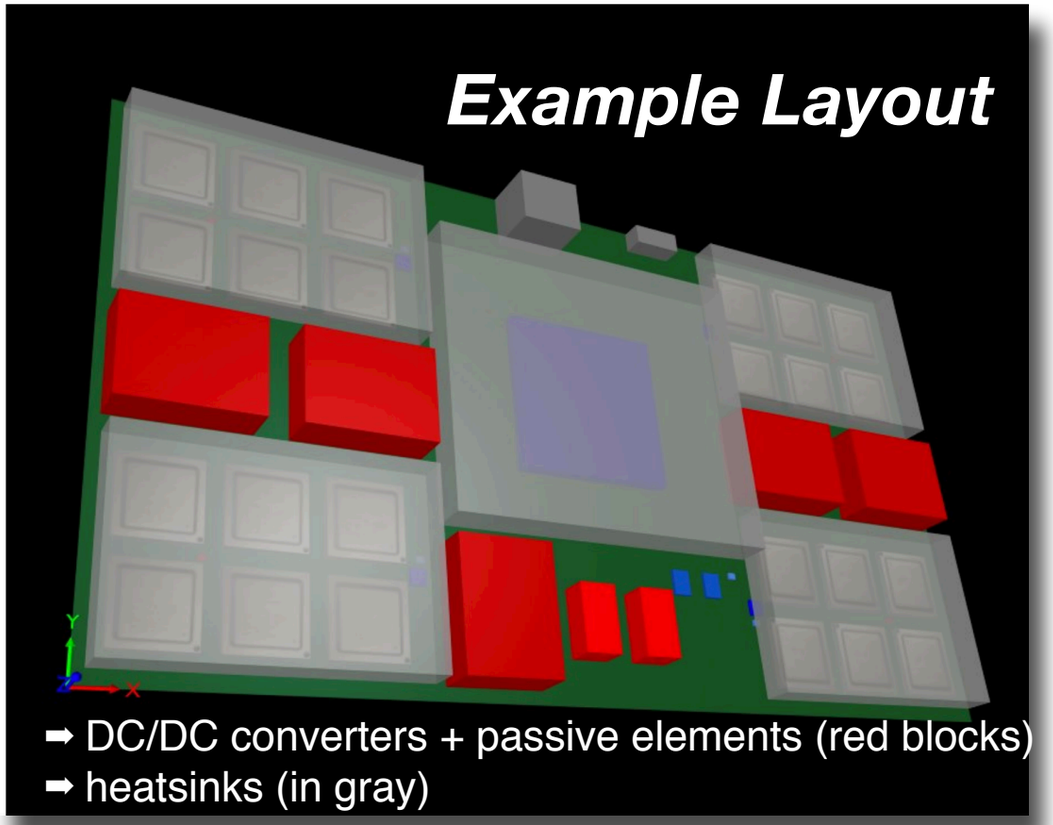

- DC/DC converters + passive elements (red blocks)
- heatsinks (in gray)

TRACK FITTING MEZZANINE (TFM)

- **Two mezzanines per SSTP board, for 2nd stage tracking (gHTT)**
  - Extrapolate 1st-stage (**PRM**) tracks into the remaining **ITk** layers and find nearby hits
  - Fit track candidates: calculate and cut on chi^2, calculate helix parameters.
- **Physical description**
  - One large **FPGA** with high-speed links to **TP**
  - Same connectors and similar heat sinks as **PRM**
- **Power within limits (nominal power = 76W; maximum power = 92W)**

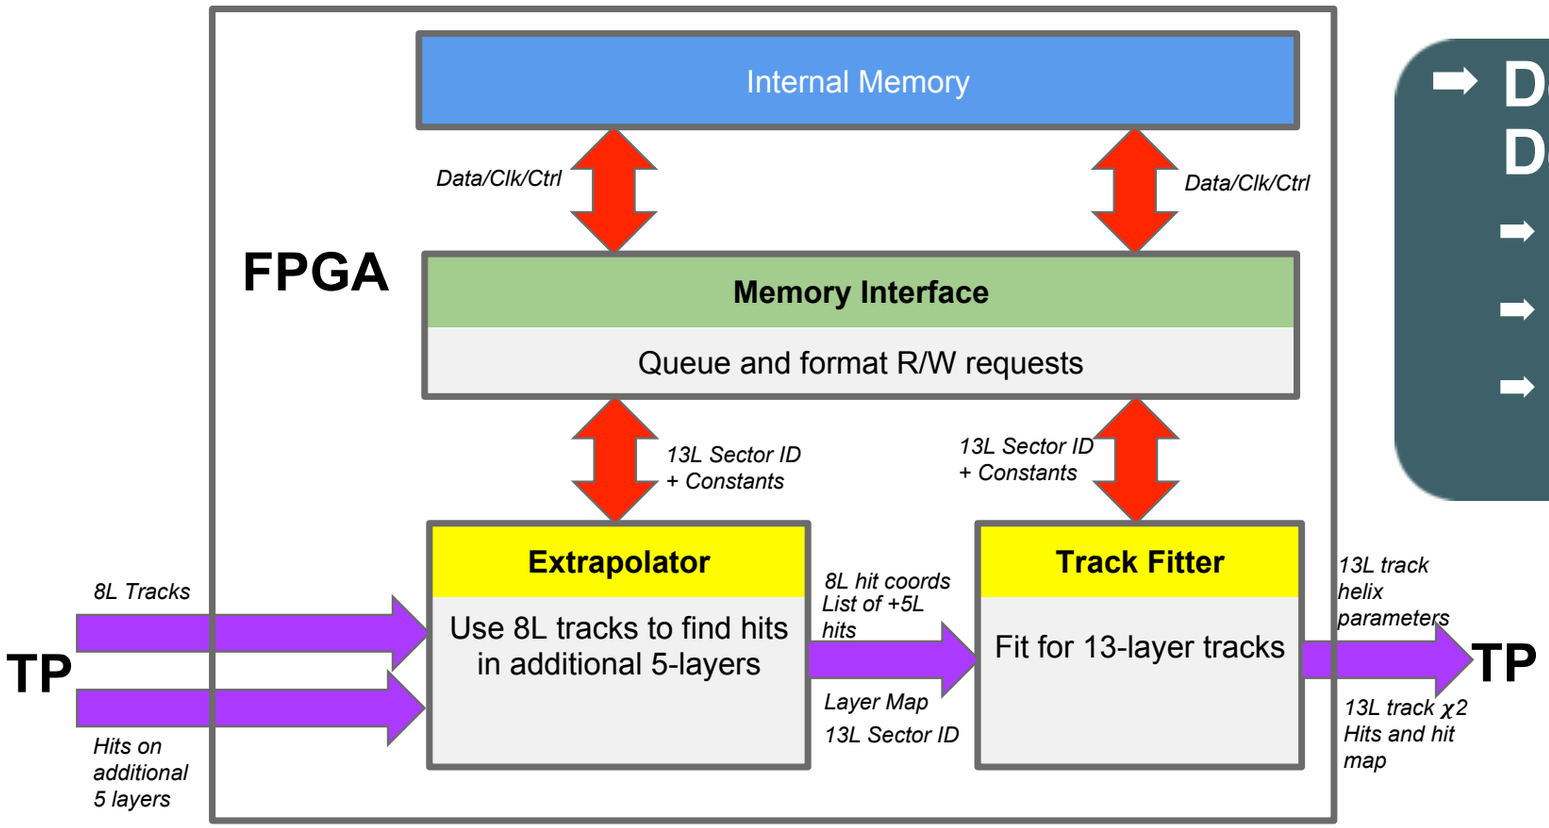

- **Demonstrator ready by December 2019**
  - Detailing schematic design
  - Firmware under development
  - Track fitting firmware first iteration based on the **FTK** AUX firmware

SYSTEM CONSIDERATIONS

- **Dataflow and power estimates, derived from detailed simulations, drive design choices to match limits and resource availability**
- **The size of the system is driven by the processing power density per board**
  - number of patterns —> number of chips
  - dataflow —> number of links
  - processing and input rates —> power allocation

- **Current estimated size: 48 AMTP + 8 SSTP ATCA shelves, with 12 blades/shelf**
  - total 13824 **ASICs** and 1440 **FPGAs**
- **Total > 1400 cards and > 5.3 G patterns**

Power estimates

Item	Number	Each	Total
ATCA shelves (12*426W blades)	62	5.1kW	317kW
Fans (per shelf (15*140W)	62	2.1kW	131kW
Rack 48V (5% of 3 shelves)	23	1.1kW	23kW
Networking (in rack only)	23	1.5kW	35kW
HTTIF PCs	26	0.3kW	8kW
ConMon PCs	58	0.3kW	18kW
Total			**532kW**

** Specs under review*

PRM dataflow estimates

	rHTT/event	gHTT/event	HTT rate	available
# Cluster /PRM (layer average)	200	260	46 MHz	60 MHz
# Roads/PRM	170	270	45 MHz	400 MHz
# Constants read/PRM	90	140	23 MHz	30 MHz
Fits/PRM	1500	2250	400 MHz	1 GHz
Tracks after χ^2/PRM	80	280	36 MHz	
Tracks after HitWarrior/AMTP	10	35	4.5 MHz	
rHTT output bandwidth /AMTP	640 Mb/s			
<Tracks after HitWarrior>/AMTP	7	20	2.6 MHz	
Total output bandwidth	250 Gb/s	750 Gb/s	1 Tb/s	
Average event size	30kB	900kB		

** TDR*

SUMMARY

- **The Hardware Track Trigger system is under design, being a crucial component of the ATLAS trigger upgrades in HL-LHC**
  - System described in 2018 TDR and User Requirement Document
- **Good flexibility and modularity, to run as both regional and global tracking co-processor in the High Level Trigger, at 1 MHz and 100 kHz respectively (baseline)**
  - If required, the system can evolve to run as L1-track-trigger, with regional requests up to 4MHz L0A, and increased p_T thresholds
- **The required performance, in terms of efficiency and resolution, has driven the design choices, and the experience in the chosen technologies allows for detailed estimates of resources and project planning**
  - Power budget and dataflow remain challenging and are key drivers of the system design

REFERENCES

- **The CFD Silicon Vertex Tracker**
  - http://inspirehep.net/record/535348/files/fermilab-conf-00-238.pdf
- **The ATLAS Fast Tracker Technical Design Report**
  - https://cds.cern.ch/record/1552953
- **The ATLAS TDAQ for Phase II Technical Design Report**
  - https://cds.cern.ch/record/2285584
- **Figures:**
  - https://twiki.cern.ch/twiki/bin/view/AtlasPublic/UpgradeEventDisplays
  - https://twiki.cern.ch/twiki/bin/view/AtlasPublic/PhysicsAndPerformancePhaseIIUpgradePublicResults
  - https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/UPGRADE/CERN-LHCC-2017-021/ch02_fig_015.png
  - https://twiki.cern.ch/twiki/bin/view/AtlasPublic/L1TrackPublicResults
  - https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/UPGRADE/CERN-LHCC-2017-020/fig_179a.png

BACK-UP SLIDES

KOXORAM/KOXORAM+

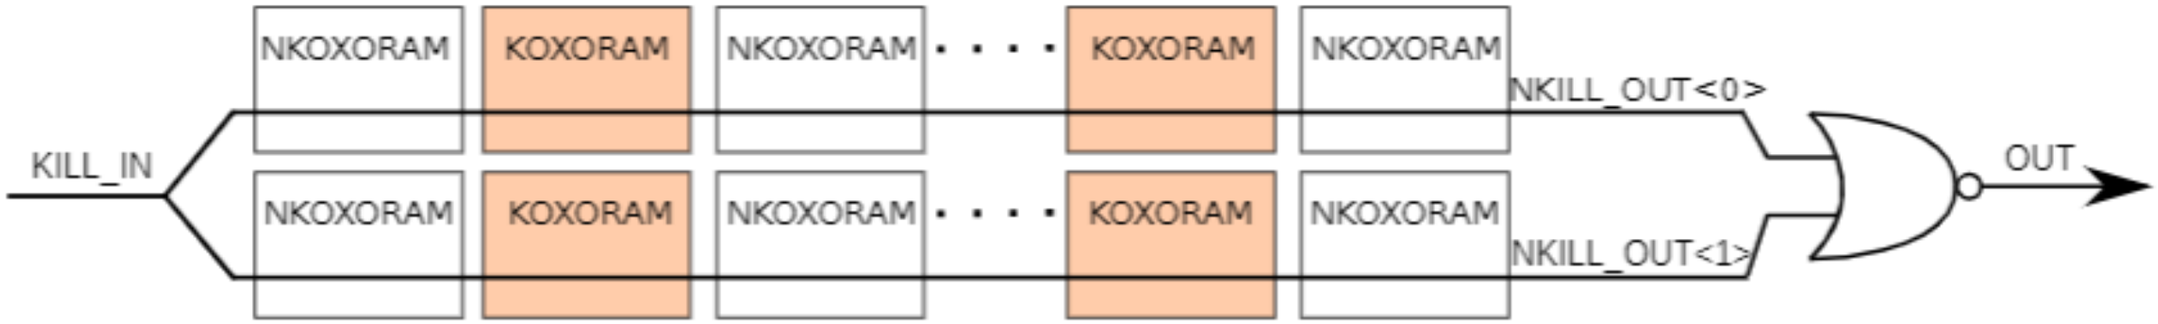

Figure 3.6: Connection of KOXORAM cells operating on the bits of a word

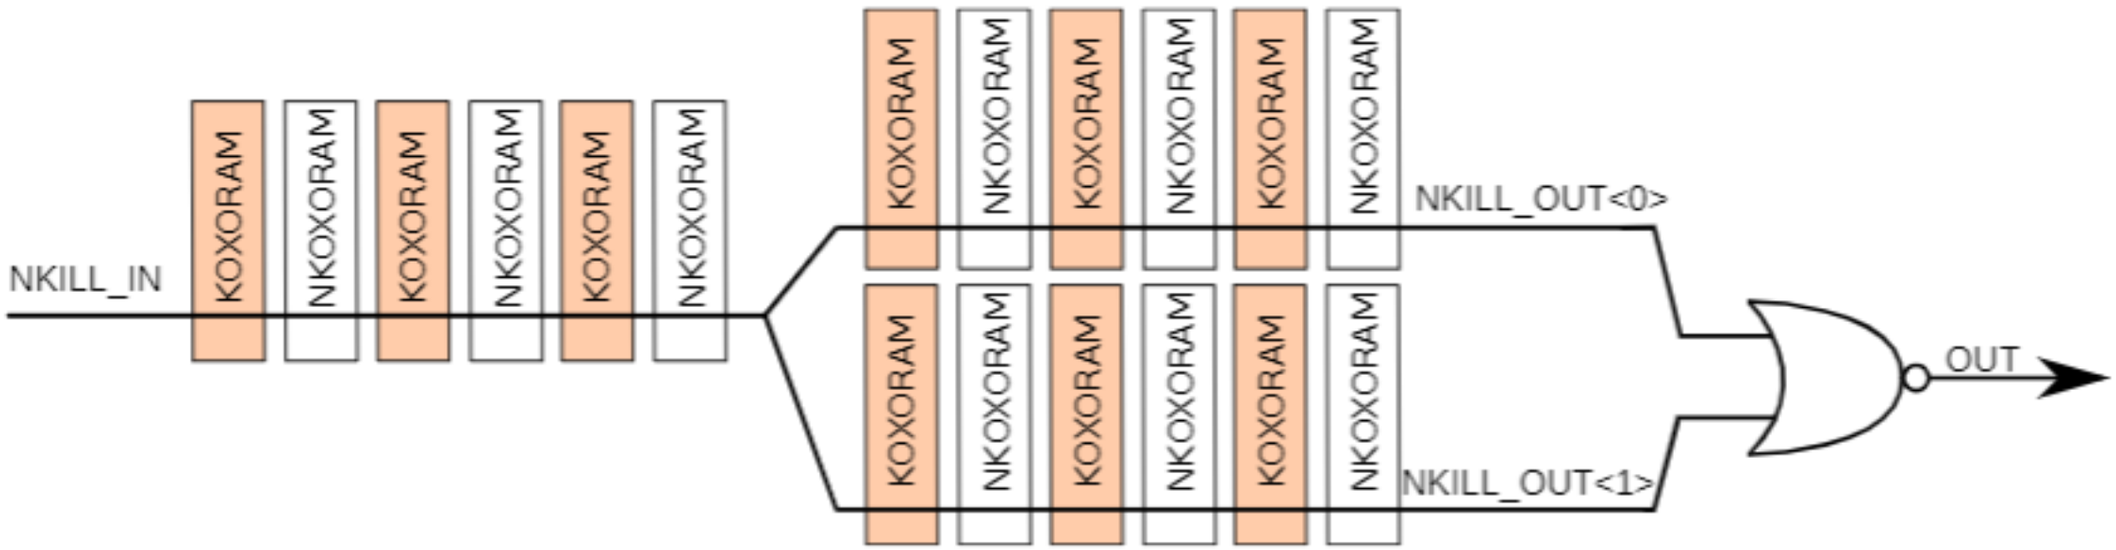

Figure 3.13: Architecture of a single row of KOXORAM+ cell

- **KOXORAM for 28 nm HPL**
- **KOXORAM+ for 28 nm HPC**
- **different cell organization for 18-bit words: maximise probability of switching cells off, even if not all serially**

SUMMARY OF HTT EXPECTED SYSTEM SIZE

Item	Number
Number of HTTIF PCs	24
Number of ATCA shelves for AMTP	48
Number of AMTP blades per shelf	12
Number of AMTP blades per HTTIF	24
Total number of AMTP	576
Number of PRM per AMTP	1
Total number of PRM	576
Number of AM ASIC per PRM	24
Total number of AM ASIC	13824
Number of ATCA shelves for SSTP	8
Number of SSTP blades per shelf	12
Number of SSTP blades per HTTIF	4
Total number of SSTP	96
Number of TFM per SSTP	2
Total number of TFM	192
Number of ConMon PCs per ATCA shelf	1
Total number of ConMon PCs	56

COST ESTIMATE (AS ESTIMATED IN TDR)

Table 18.7: *CORE Cost Estimate of the EF system in the TDAQ UPR. CORE values of the system's components represent current best estimates.*

PBS Code	Item	Cost [kCHF]	Cost Quality
1.3	**Event Filter**	**20,847**	
1.3.1	**Tracking Hardware**	**17,448**	
1.3.1.1	Tracking Processor (TP)	6,451	QF3
1.3.1.2	TP Rear transition Module (RTM)	695	QF2
1.3.1.3	Pattern Recognition Mezzanine (PRM)	4,678	QF3
1.3.1.4	AM ASIC	3,293	QF4
1.3.1.5	Track fitting Mezzanine (TFM)	1,261	QF3
1.3.1.6	Infrastructure	658	QF2
1.3.1.7	Hardware Tracking Interface	411	QF2
1.3.2	**Processing Units**	**3,399.0**	
1.3.2.1	Servers	3,399.0	QF3

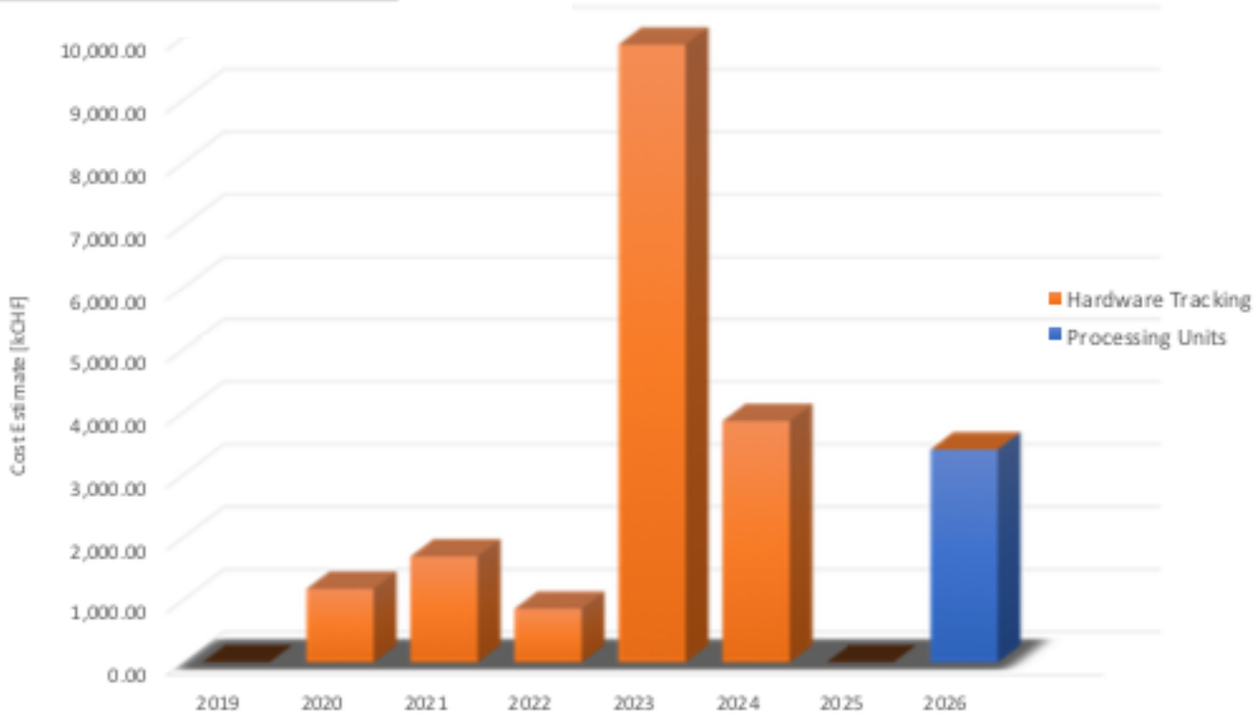

AMCHIP DESIGN COMPLEXITY VS CPUS

Chip name	*Transistor count*	*Year*	*Brand*	*Technology*	*Area*
Core 2 Duo Conroe	291,000,000	2006	Intel	65 nm	143 mm²
Itanium 2 Madison 6M	410,000,000	2003	Intel	130 nm	374 mm²
Core 2 Duo Wolfdale	411,000,000	2007	Intel	45 nm	107 mm²
AM06	**421,000,000**	**2014**	**AMteam**	**65 nm**	**168 mm²**
Itanium 2 with 9 MB cache	592,000,000	2004	Intel	130 nm	432 mm²
Core i7 (Quad)	731,000,000	2008	Intel	45 nm	263 mm²
Quad-core z196[20]	1,400,000,000	2010	IBM	45 nm	512 mm²
Quad-core + GPU Core i7 Ivy Bridge	1,400,000,000	2012	Intel	22 nm	160 mm²
Quad-core + GPU Core i7 Haswell	1,400,000,000	2014	Intel	22 nm	177 mm²
AM09	**1,684,000,000**	**2019**	**AMteam**	**28 nm**	**150 mm²**
Dual-core Itanium 2	1,700,000,000	2006	Intel	90 nm	596 mm²

***Ref**: http://repodip.fisica.unimi.it/cdip17/talks/8_2_AStabile-compressed.pdf*