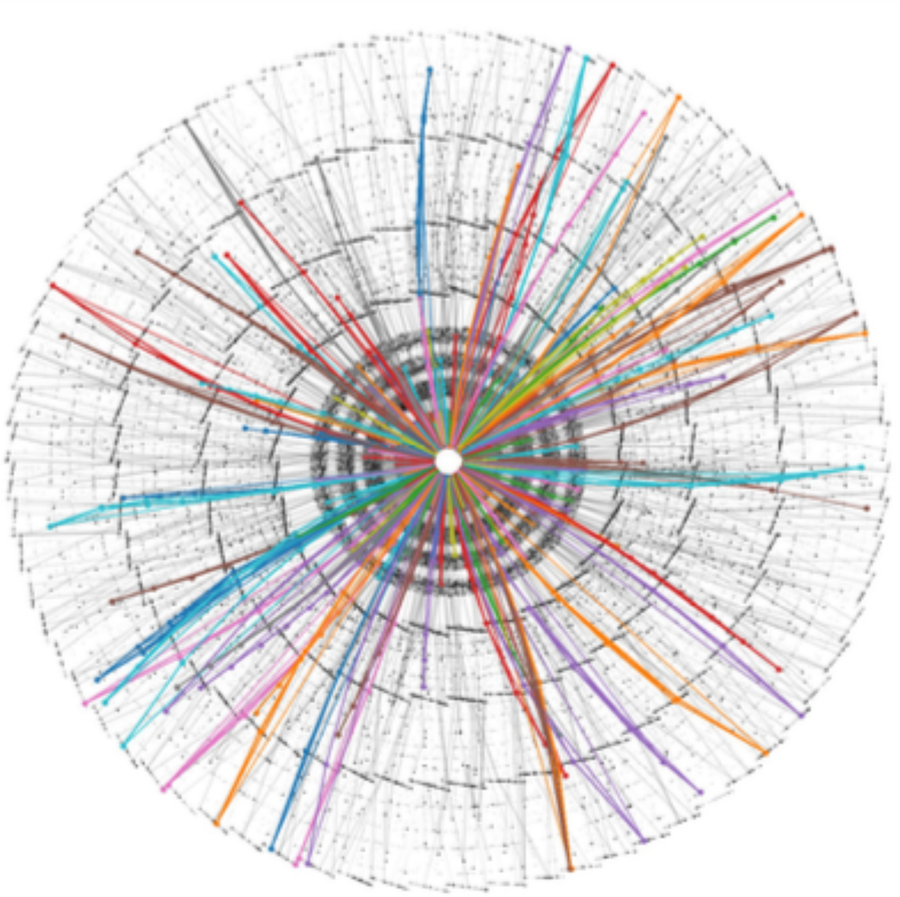

ATLAS HLT/offline strategy

A. Salzburger (CERN)

HL-LHC timeline

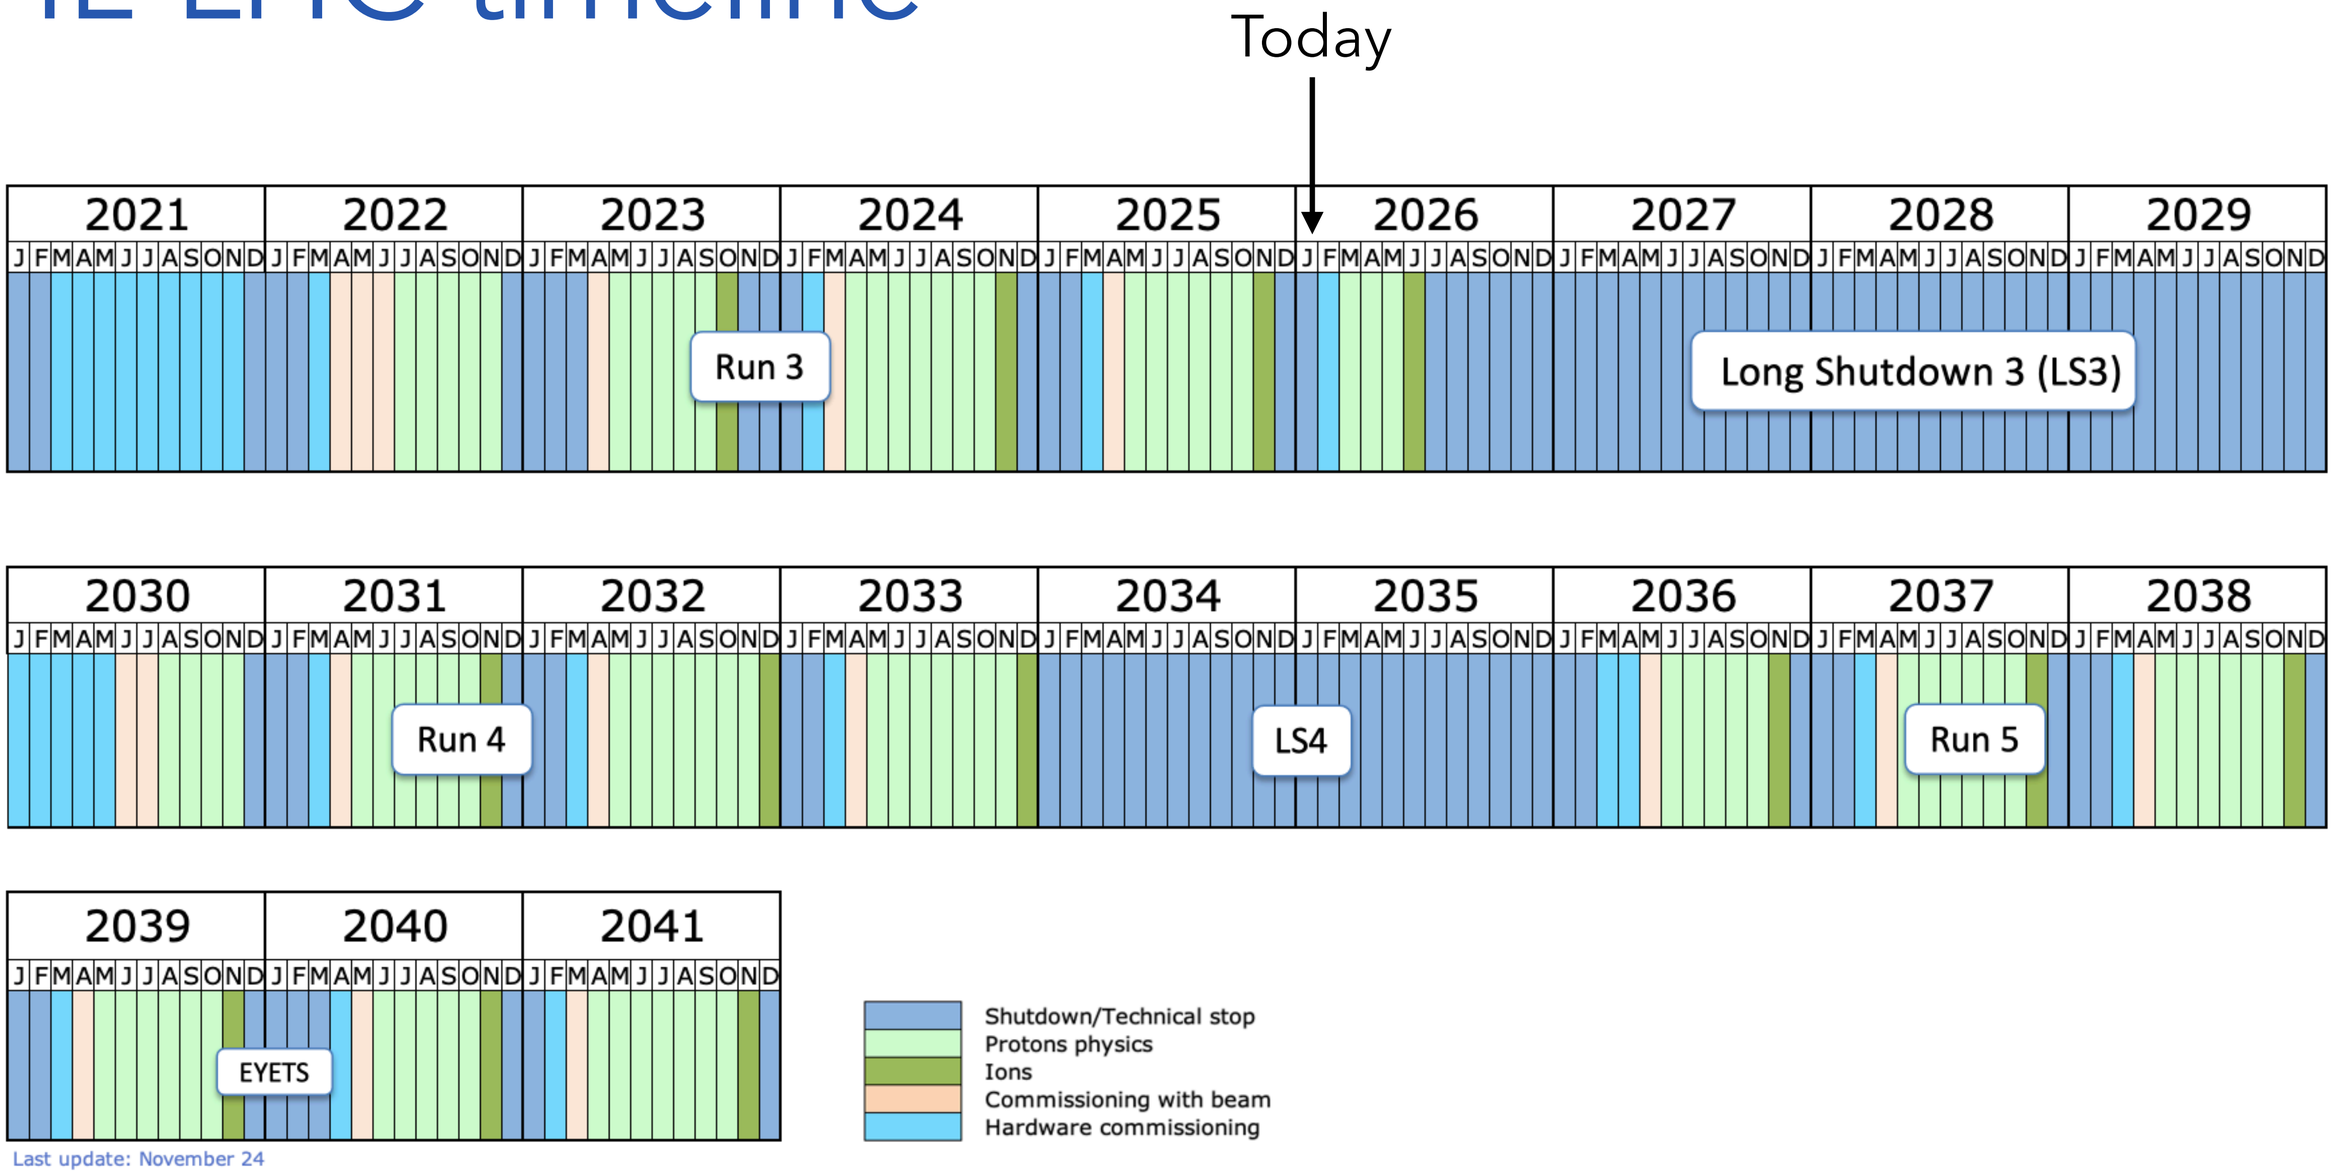

https://lhc-commissioning.web.cern.ch/schedule/LHC-long-term.htm

Inner Tracker (ITk) System

All new full silicon Inner Tracking system

- 5 layer (barrel) Pixel system with inclined and ring end-cap sections
- 4 layer (barrel) double sided stereo Strip system, with 6 end-cap disks

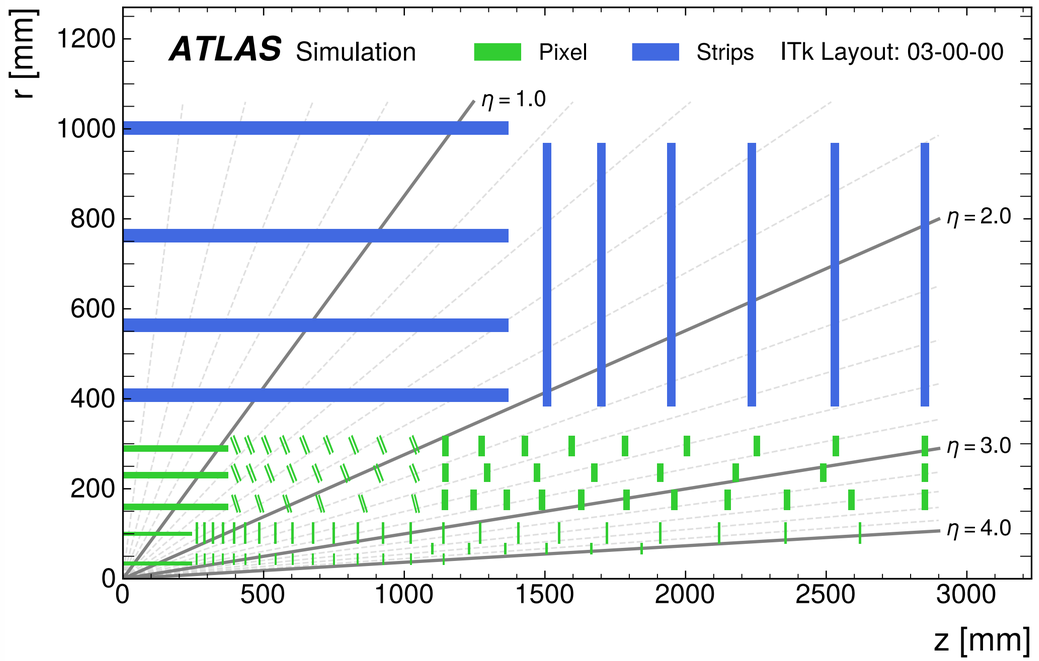

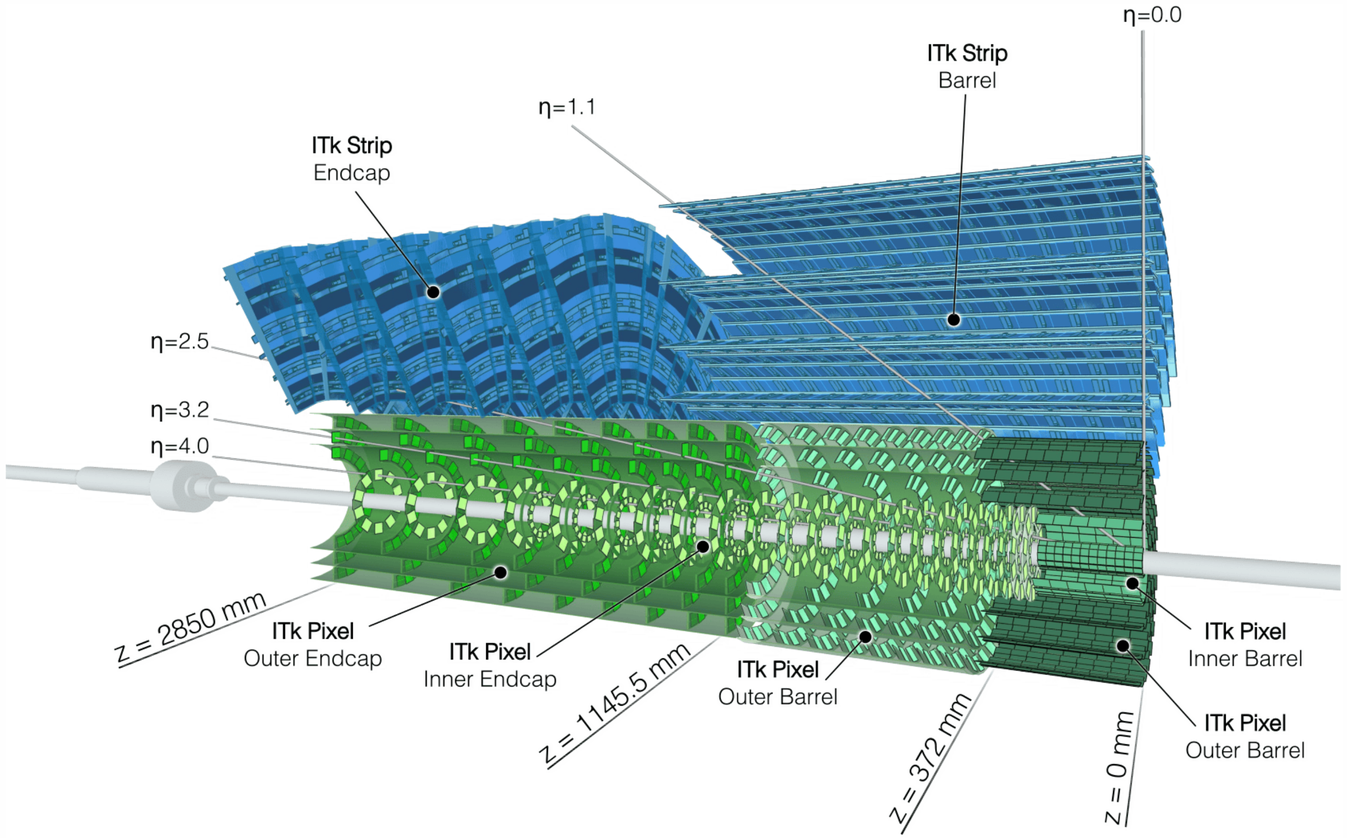

https://doi.org/10.1088/1748-0221/20/02/P02018

Run-4 Trigger/DAQ

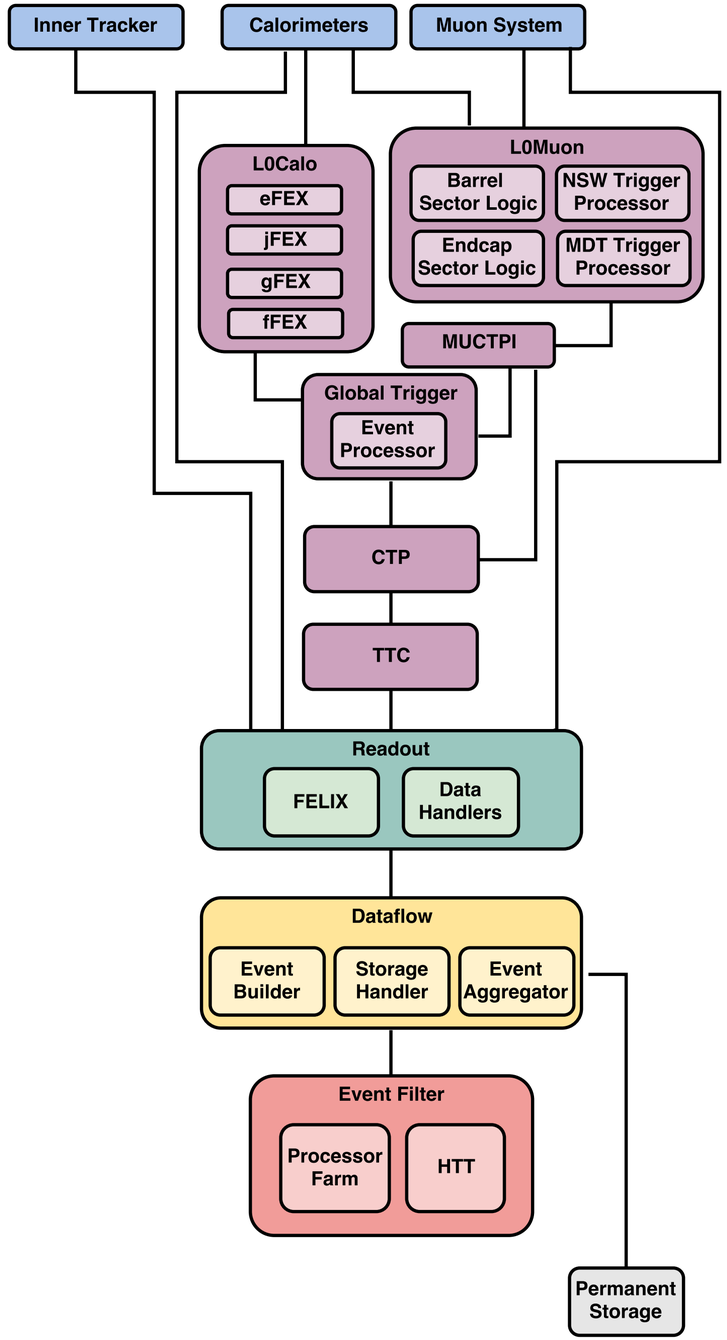

TDAQ Phase-II Upgrade Project

Level 0 Trigger

L0Calo & L0Muon feed into Central Trigger Processor into Readout (FELIX Cards)

Latency: 10 μs

→ reduces rate to1 MHz

Readout & DAQ

Event building at 1 MHz &compression, storage, and transfer to offline at 10 kHz

Event Filter

→ reduces rate further to 10 kHz

https://cds.cern.ch/record/2285584

Event Filter - Requirements

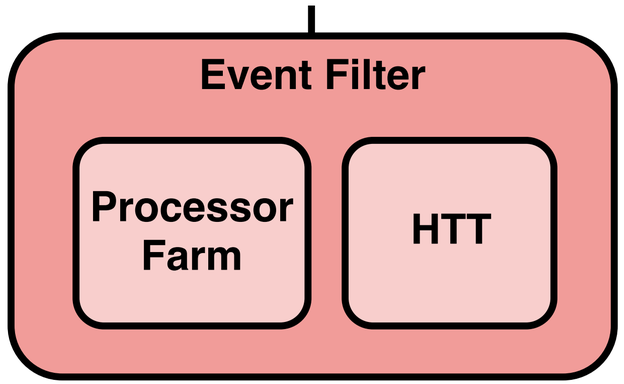

Event Filter

→ reduces rate further to 10 kHz

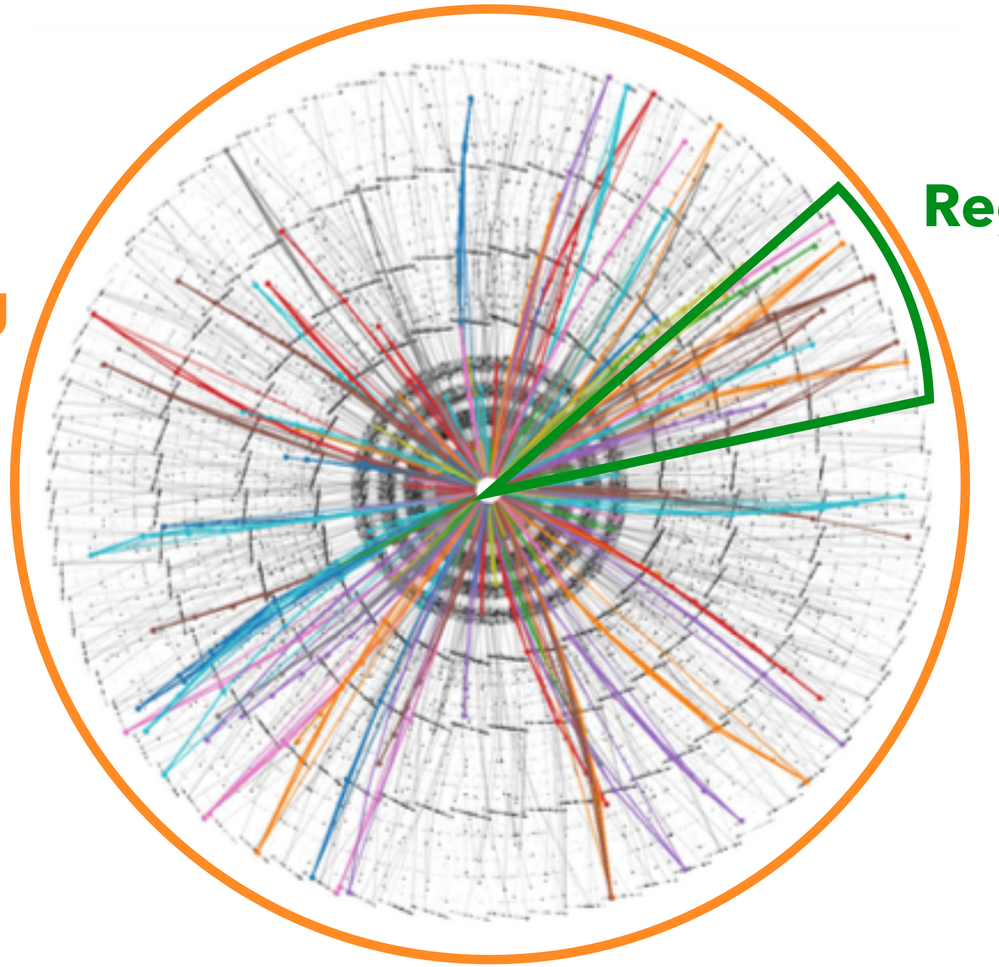

Global tracking

Using full ITk information at **150 kHz**

Regional tracking

In defined regions of interest (ROIs) at **1 MHz**

https://cds.cern.ch/record/2285584

Event Filter - v0 (TDR 2017)

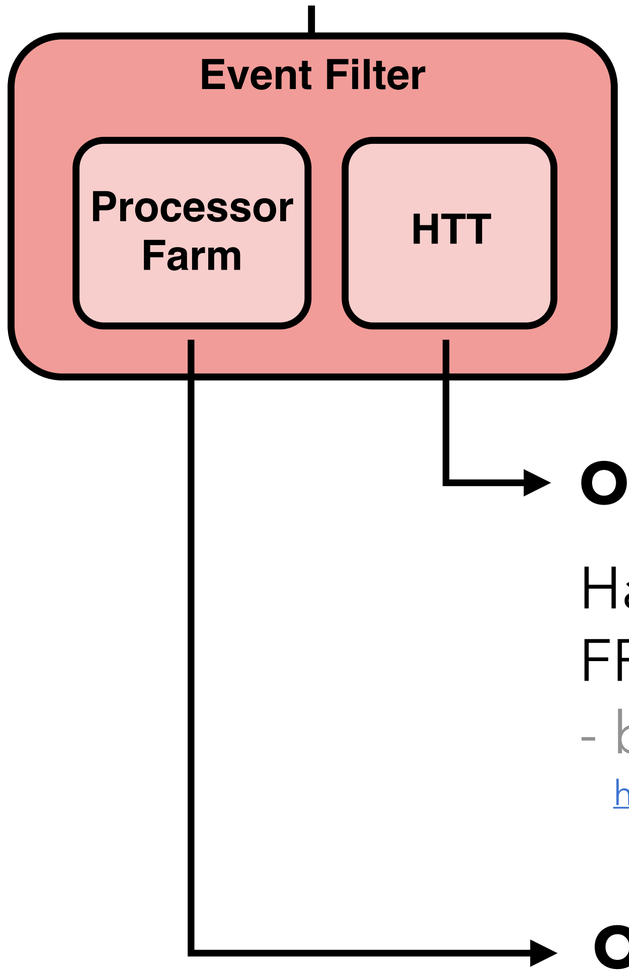

Option 1: Hardware based Track Tigger (HTT)

Hardware system with ASICS for pattern recognition/ FPGAs for track fitting

- built on ATLAS' FTK project experience (disc. 2019)
  https://iopscience.iop.org/article/10.1088/1748-0221/16/07/P07006

Option 2: Commodity Farm with CPU

- possible adding accelerators (GPU, FPGA)

https://cds.cern.ch/record/2285584

HTT - evolution of FTK concept

Option 1: Hardware based Track Tigger (HTT)

- initial baseline: hard based tracking pre-procssing step
  - evolution of the ATLAS FTK demonstrator
- Associative Memory ASICS for **pattern recognition**
  - detector granularity modeled as super-strips
  - Pattern recognition done by matching to pre-computed super-strip patterns
- FPGAs for first level **track fitting**

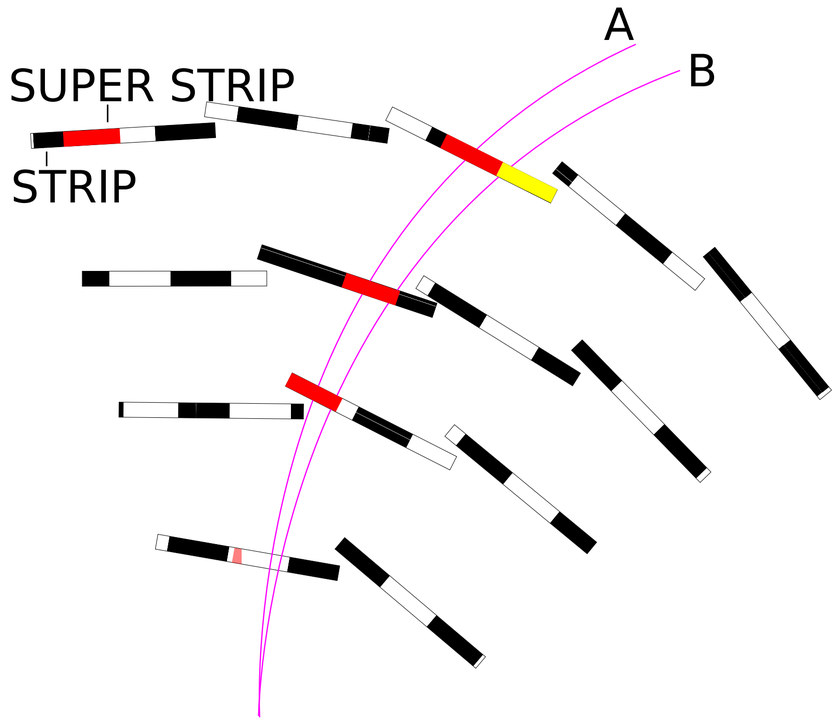

https://cds.cern.ch/record/2285584

HTT - evolution of FTK concept

Option 1: Hardware based Track Tigger (HTT)

- initial baseline: hard based tracking pre-procssing step
  - evolution of the ATLAS FTK demonstrator
- Associative Memory ASICS for **pattern recognition**
- FPGAs for first level **track fitting**
  - linearized x2 fitter in two stages
  - executed on FGPAs
  - output interfaced to EF farm via network switches

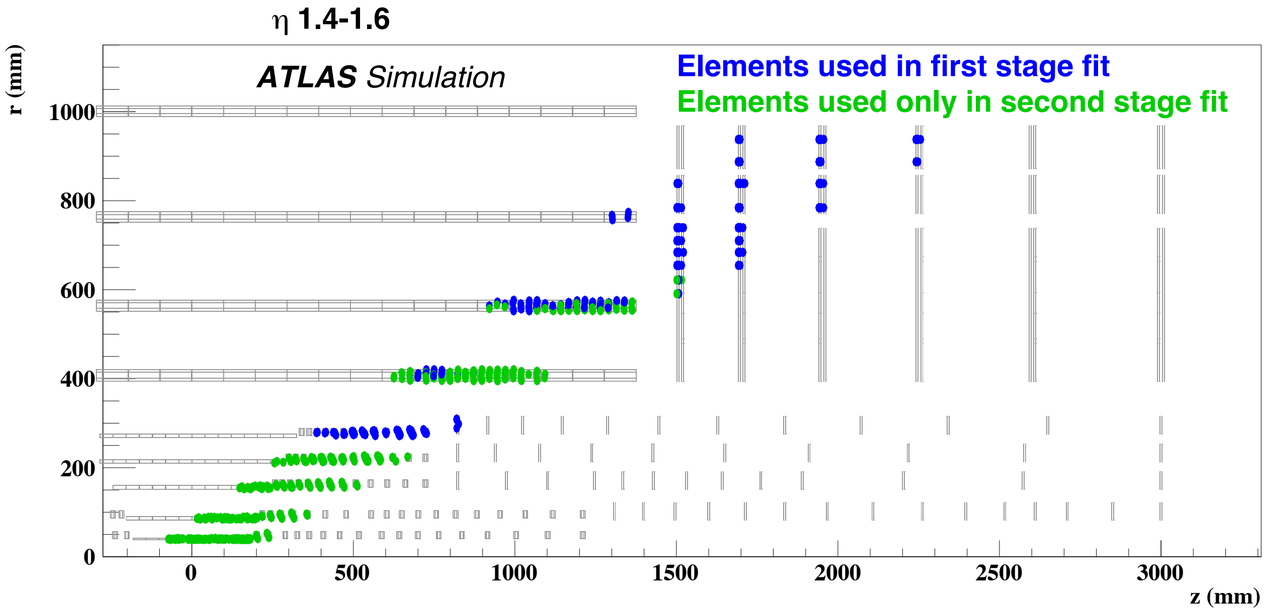

https://cds.cern.ch/record/2285584

HTT - proof of concept

Option 1: Hardware based Track Tigger (HTT)

- could efficiently find tracks (match patterns)
- track fit ok for trigger decision (though worse than offline)
- Proof of concept with FTK on ALTAS data

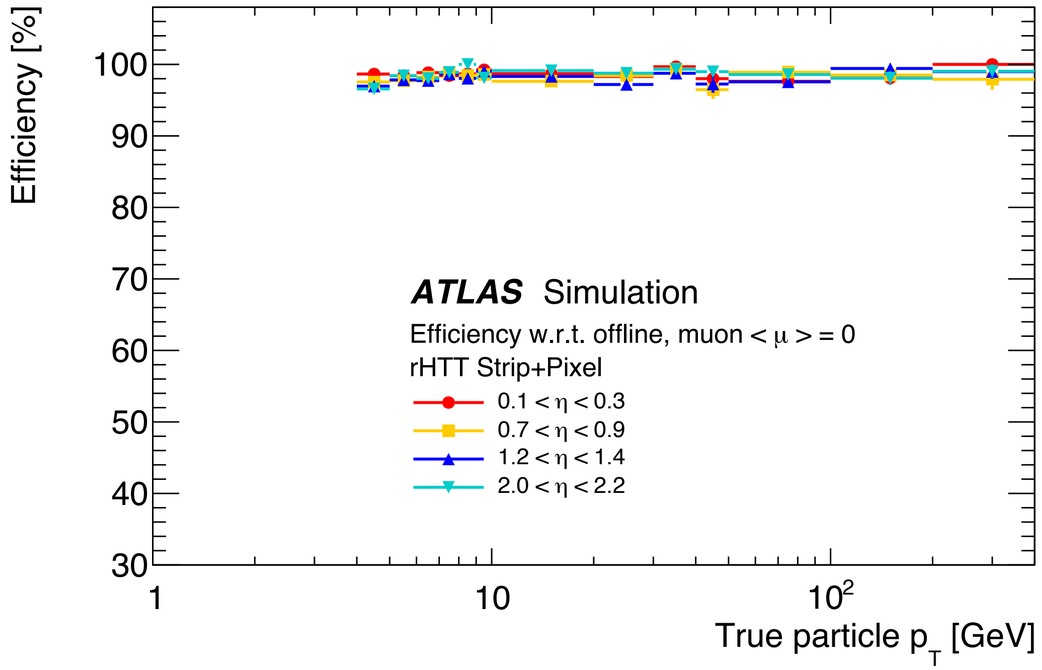

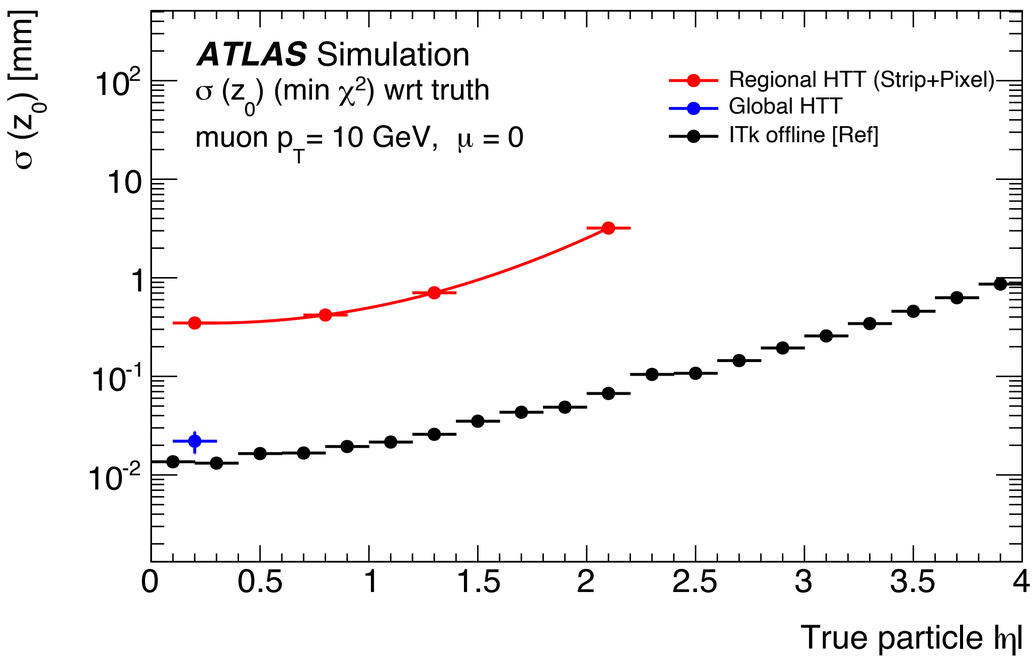

https://cds.cern.ch/record/2285584

HTT - shortcomings

Option 1: Hardware based Track Tigger (HTT)

- Single tracks usually created many match candidates
  - Resolving by offline/EF CPU code did not lead to any compute improvements
  - Mitigation strategy on FPGA: candidate reduction (HitWarrior)

particle	min p_T	Eff. (%)	# roads	# fits	# tracks $\chi^2 < 40$	# tracks HitWarrior
muon	1 GeV	99.5	144	1115	55	4.6
muon	2 GeV	99.1	79	586	23	1.9
muon	4 GeV	99.2	48	313	16	1.2

- HTT simulation for MC studies was unsolved
  - Most progressed approach was a parametric simulation (not really satisfying)
- HTT ASICS and FPGAs were practically unused resources during down time
  - Commodity farm stood in strong contrast to it

https://cds.cern.ch/record/2285584

EF Tracking Addendum (2022)

Option 2: Commodity Farm with CPU (+ possible GPGPUs)

- Advances in reconstruction made SW solution on Commodity HW farm feasible:
  - new seeding strategy
  - ITk layout (optimized also for speed)
- Possibilities with CPU only,
  but also with accelerated components
- Start of the EF Tracking initiative of
  ATLAS to address the SW and HW
  landscape

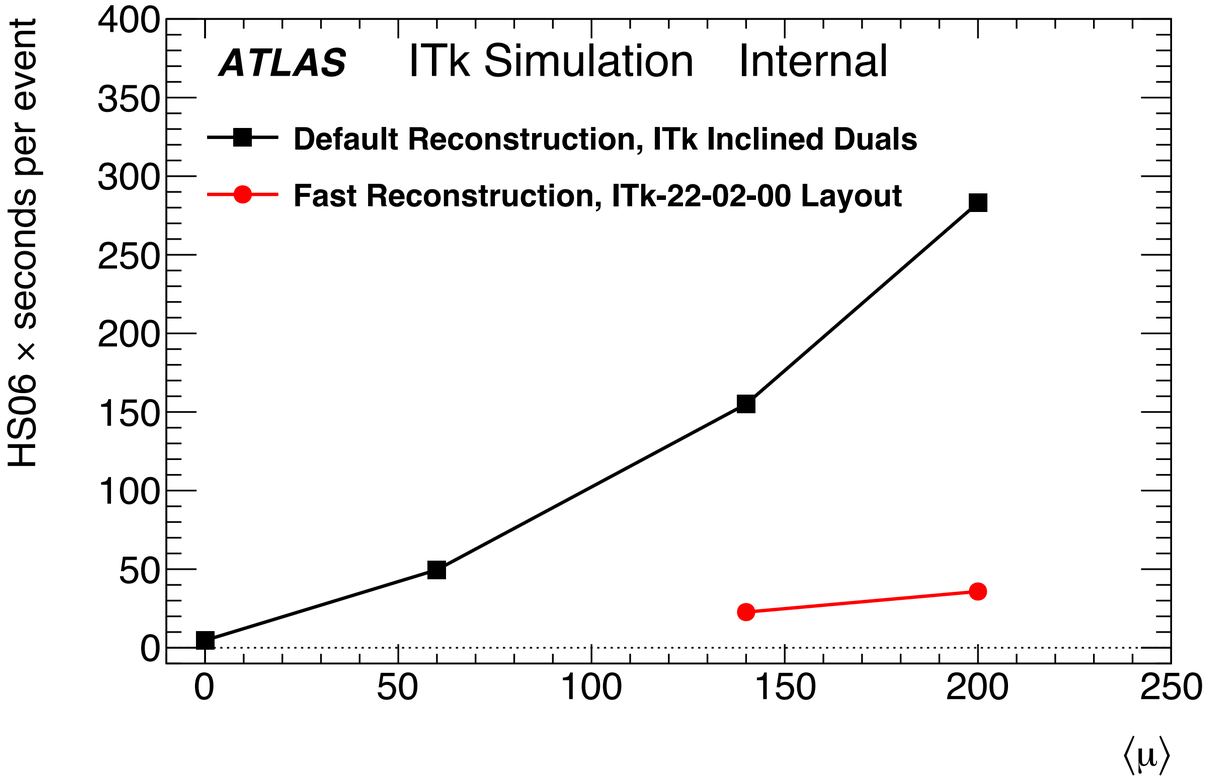

https://cds.cern.ch/record/2802799

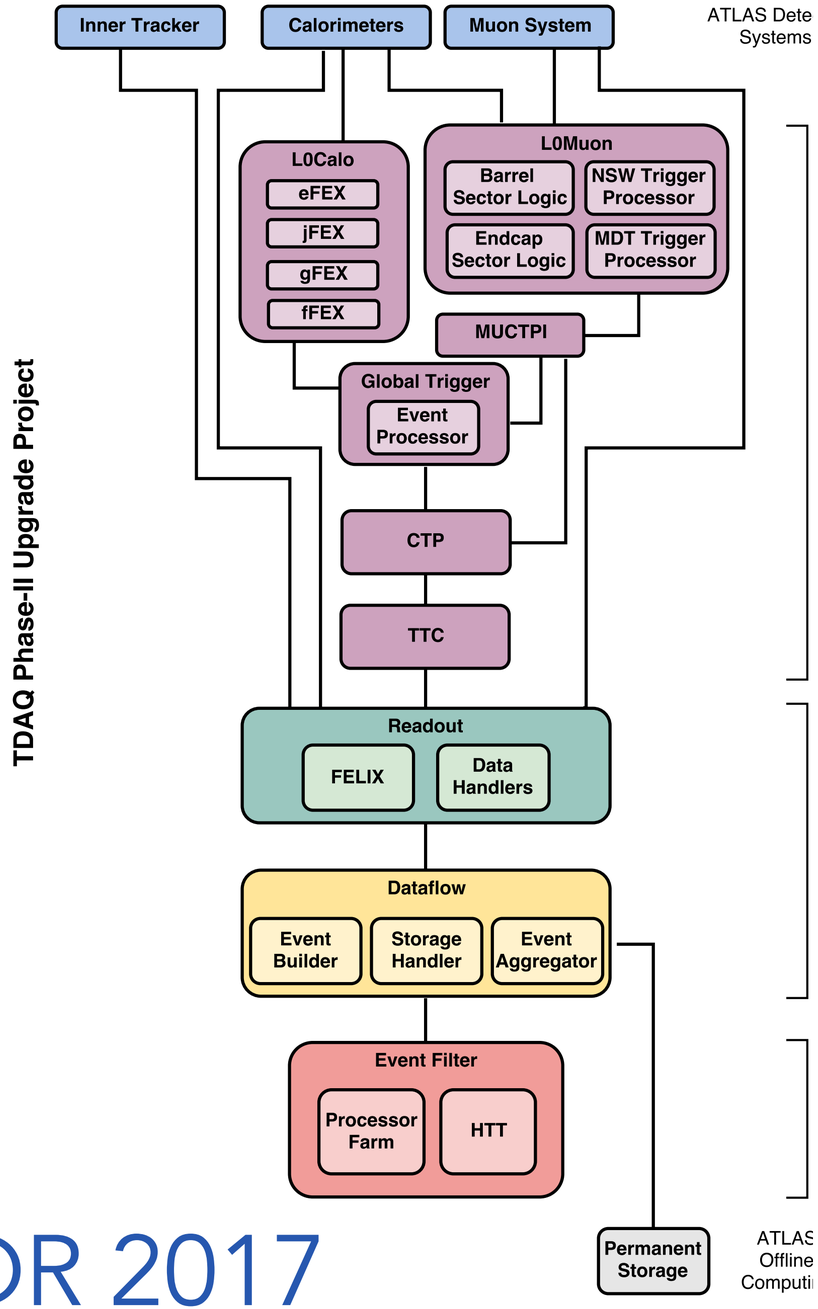

TDR 2017

https://cds.cern.ch/record/2285584

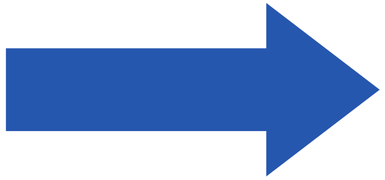

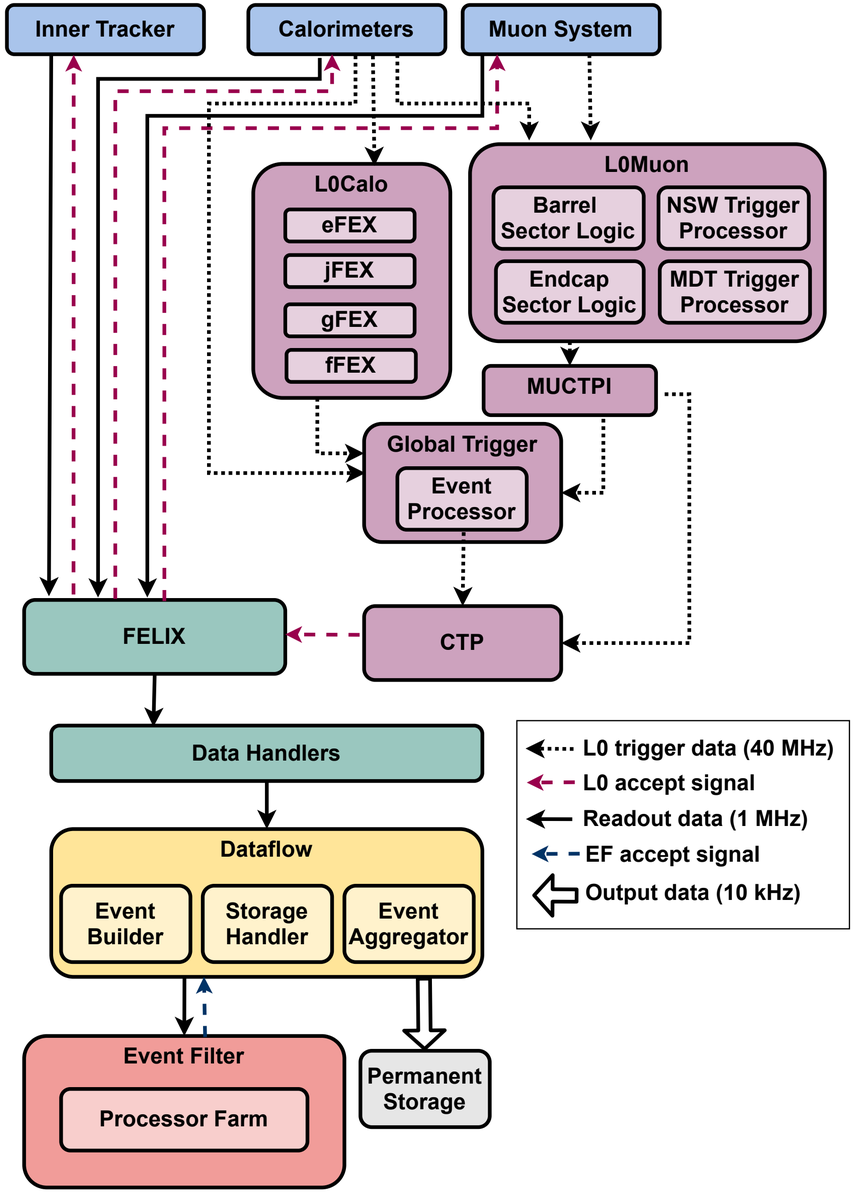

TDR Addendum 2022

https://cds.cern.ch/record/2802799

Event Filter - v1 (TDR Addendum 2022)

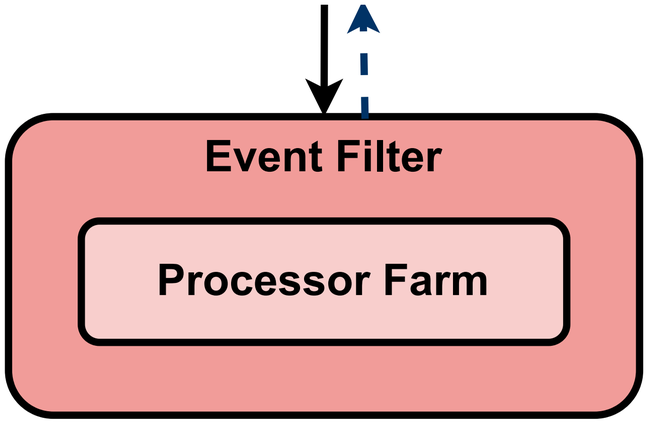

New baseline: CPU farm within cost-/power bracket

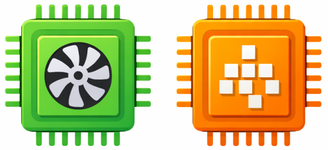

How much is to **gain** by offloading?

Farm Size/Cost (CPU only)

	Run-4 (2027)	Run-5 (2032)
CPU for tracking [MHS06]	5.90	9.17
Other reconstruction [MHS06]	1.86	2.27
Required CPU [MHS06]	7.76	11.44
Available CPU [MHS06]	2.85	7.76
Compute cost [CHF/HS06]	1.3	0.6
Farm extension cost [MCHF]	6.4	2.2
Total CORE cost [MCHF]	8.8	

2022 estimates with 1.3 CHF/HS06 for Run-4 (2027)
and 0.6 CHF/HS06 for Run-5 (2032)

Power estimates (CPU only)

	Pile-up 140		Pile-up 200	
	full-scan	regional	full-scan	regional
Rate [MHz]	0.15	1.0	0.15	1.0
CPU Resource requirement [MHS06]	3.41	2.49	5.36	3.81
Tracking resource requirement [MHS06]	5.90		9.17	
Tracking power requirement [MW]	1.47		1.83	

https://cds.cern.ch/record/2802799

Event Filter - v1 (TDR Addendum 2022)

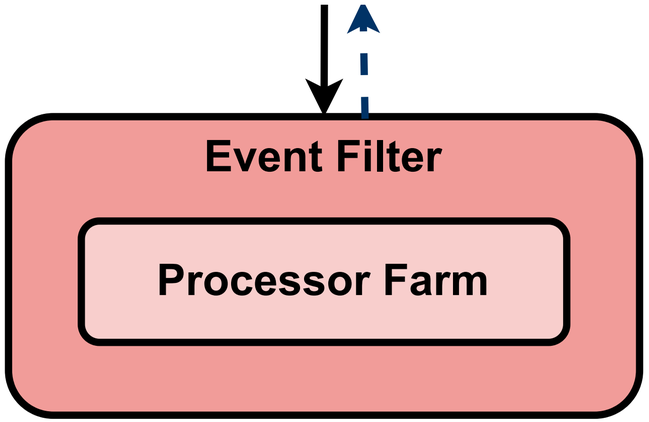

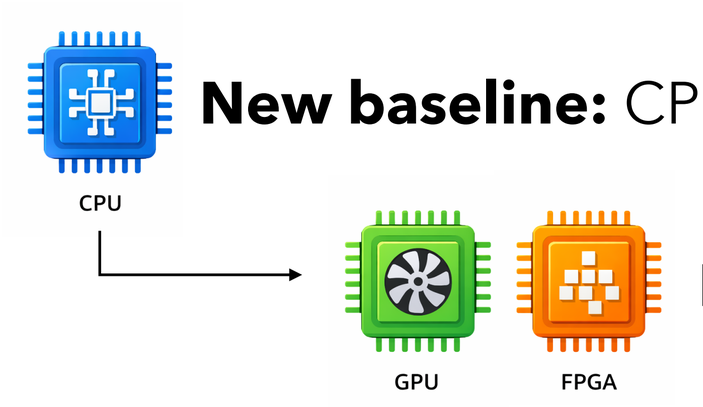

New baseline: CPU farm within cost-/power bracket

How much is to **gain** by offloading?

- throughput
- power, cooling

How much is to **lose** by offloading?

- tracking performance (CPU is close to offline quality)
- system complexity, network, cooling

Additional considerations

How can the farm be used in down time?

Geant4 on GPUs?
https://indico.cern.ch/event/1526077 (WLCG workshop)

https://cds.cern.ch/record/2802799

EF Tracking Initiative

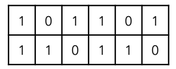

Data Preparation

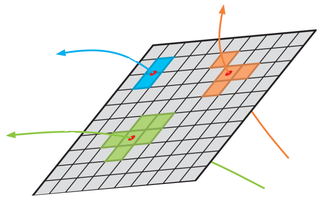

Clustering

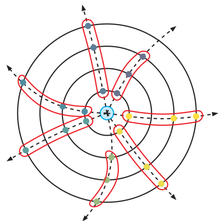

Seeding

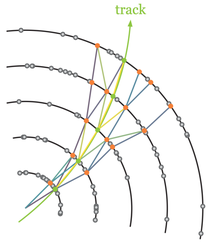

Track Finding

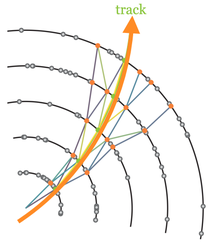

Track Fitting

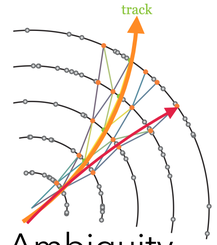

Ambiguity Solving

EDM Preparation

CPU

Athena	ACTS	ACTS	ACTS	ACTS	ACTS	Athena

GPU

AthenaGPU	traccc	ACTS	ACTS	ACTS	ACTS	Athena
AthenaGPU	traccc	traccc	ACTS	ACTS	ACTS	Athena
AthenaGPU	traccc	traccc	traccc	traccc	traccc	Athena

FPGA

AthenaFPGA	AthenaFPGA	ACTS	ACTS	ACTS	ACTS	Athena
AthenaFPGA	AthenaFPGA	AthenaFPGA	ACTS	ACTS	ACTS	Athena

EF Tracking Initiative - pipelines

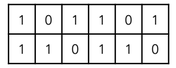

Data Preparation

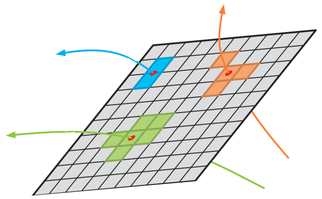

Clustering

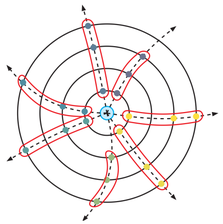

Seeding

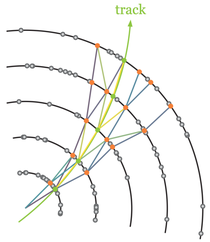

Track Finding

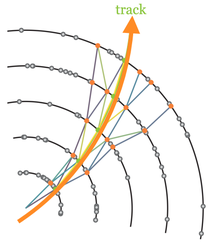

Track Fitting

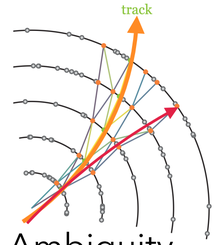

Ambiguity Solving

EDM Preparation

CPU

Athena	ACTS	ACTS	ACTS	ACTS	ACTS	Athena

C-100

GPU

AthenaGPU	traccc	ACTS	ACTS	ACTS	ACTS	Athena
AthenaGPU	traccc	traccc	ACTS	ACTS	ACTS	Athena
AthenaGPU	traccc	traccc	traccc	traccc	traccc	Athena

FPGA

AthenaFPGA	AthenaFPGA	ACTS	ACTS	ACTS	ACTS	Athena
AthenaFPGA	AthenaFPGA	AthenaFPGA	ACTS	ACTS	ACTS	Athena

EF Tracking Initiative - pipelines

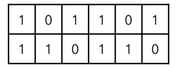

Data Preparation

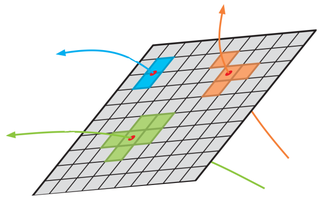

Clustering

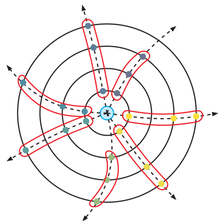

Seeding

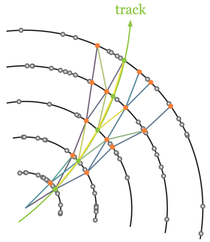

Track Finding

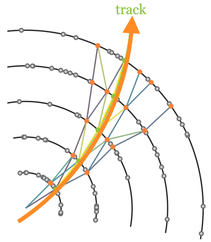

Track Fitting

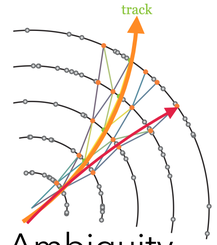

Ambiguity Solving

EDM Preparation

CPU	Athena	ACTS	ACTS	ACTS	ACTS	ACTS	Athena	C-100
GPU	AthenaGPU	traccc	ACTS	ACTS	ACTS	ACTS	Athena	G-050
	AthenaGPU	traccc	traccc	ACTS	ACTS	ACTS	Athena	G-100
	AthenaGPU	traccc	traccc	traccc	traccc	traccc	Athena	G-200
FPGA	AthenaFPGA	AthenaFPGA	ACTS	ACTS	ACTS	ACTS	Athena	F-100
	AthenaFPGA	AthenaFPGA	AthenaFPGA	ACTS	ACTS	ACTS	Athena	F-150

Offloading potential - at scale

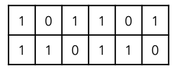

Data Preparation

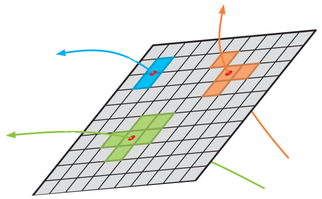

Clustering

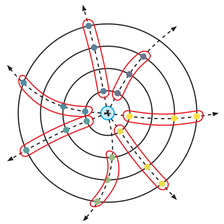

Seeding

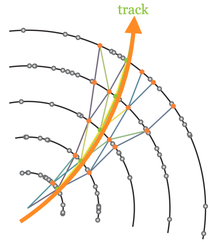

Track Finding, Fitting & Ambiguity Solving

EDM Preparation

CPU

GPU

FPGA

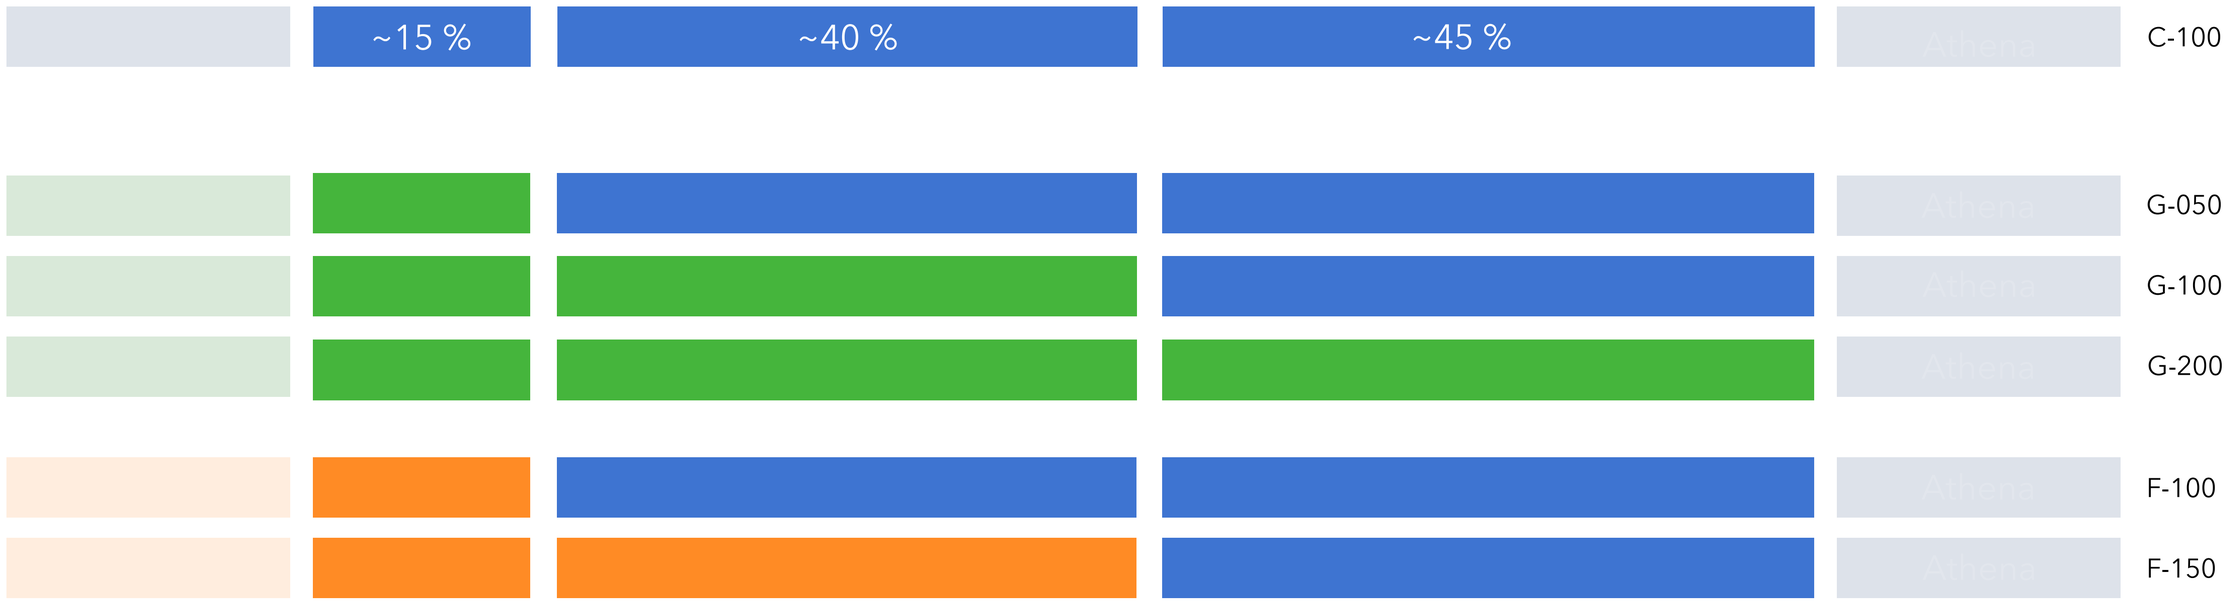

CPU Baseline solution

ACTS Fast Tracking: based on ATLAS legacy tracking philosophy

- implemented in ACTS, wrapped into Athena (see talk by Corentin)

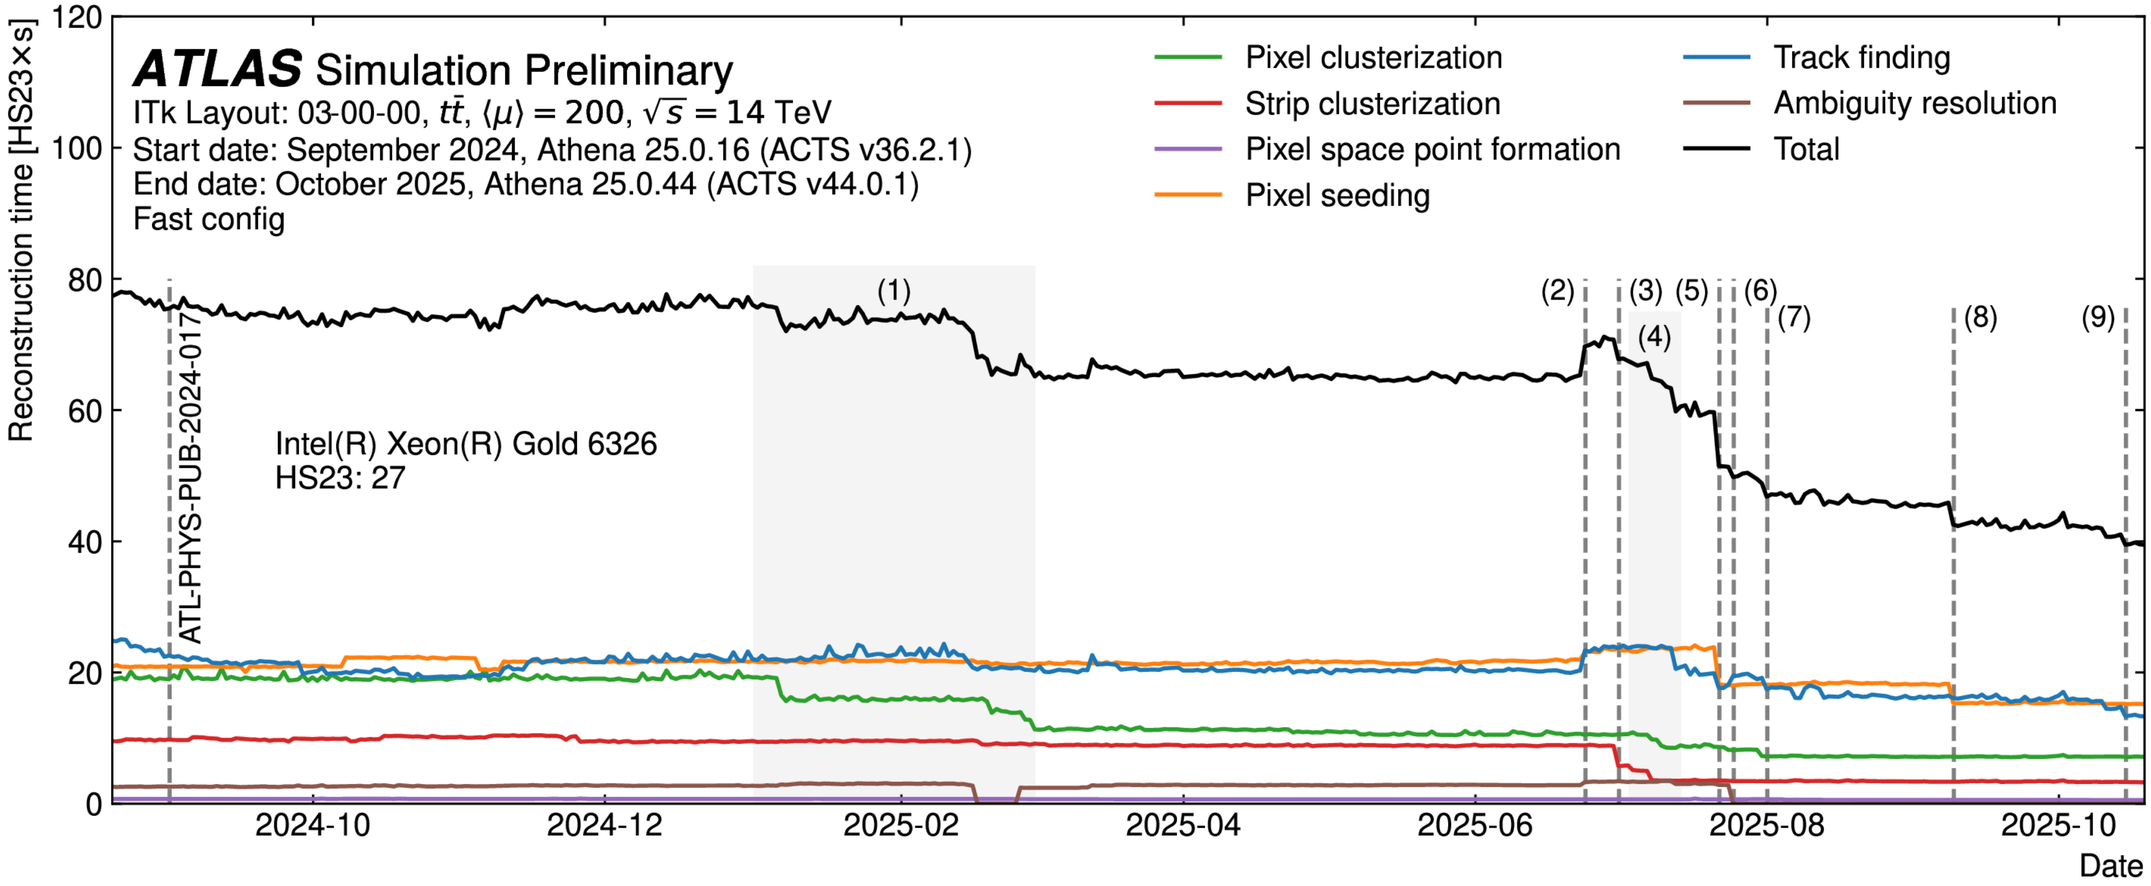

https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/IDTR-2025-06/

CPU pipelines - flash results

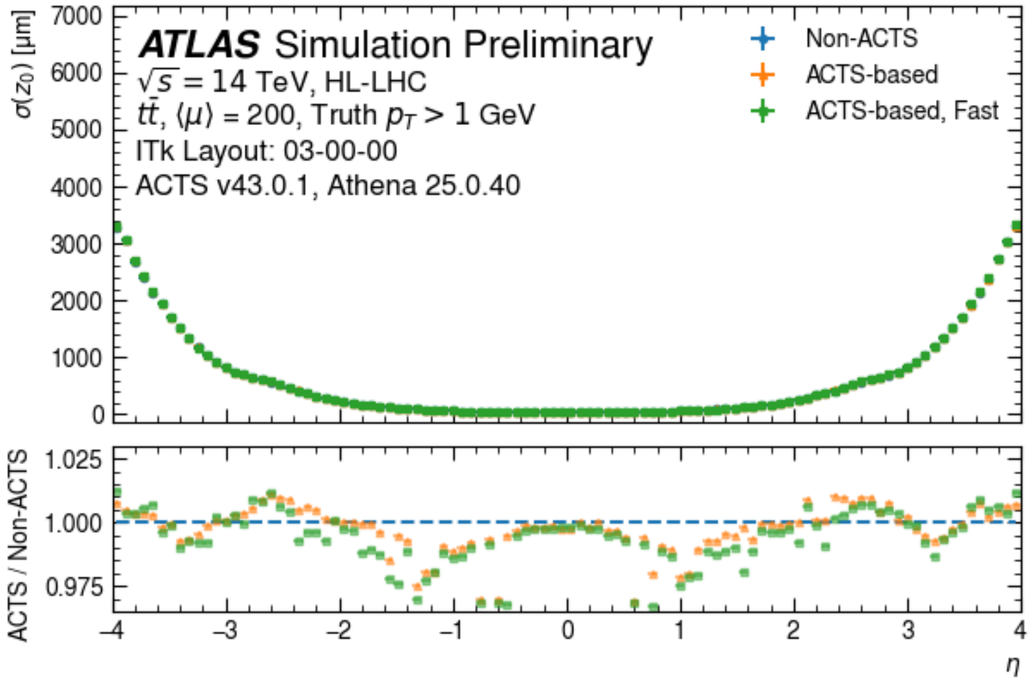

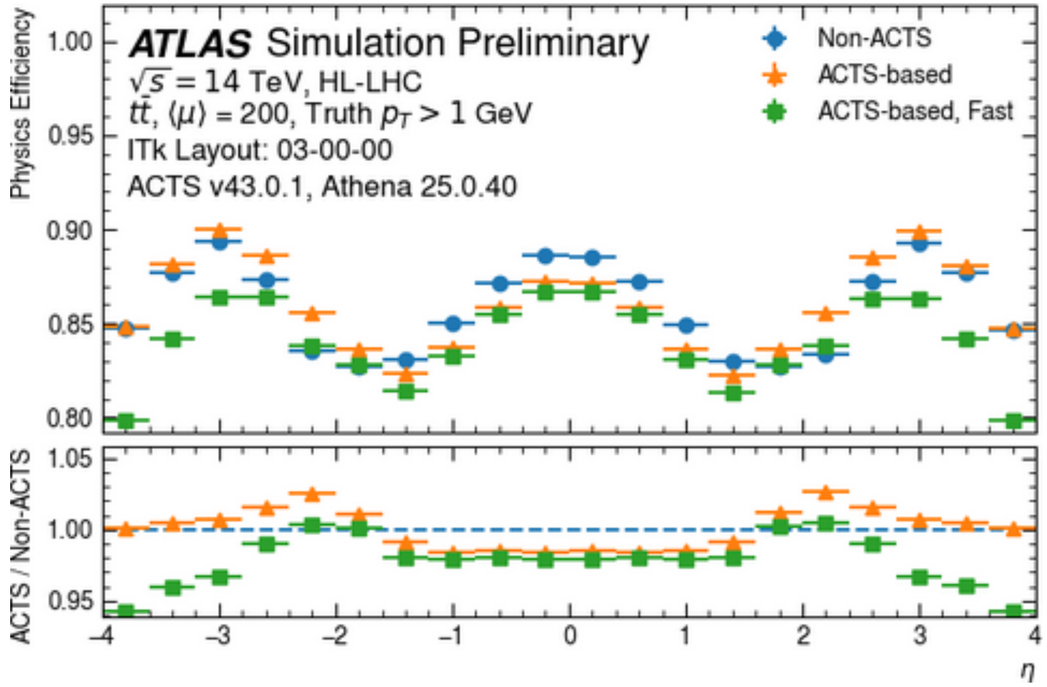

ACTS based Tracking is reaching legacy performance (limited by detector material)

- plan in 2026 to switch to ACTS Tracking as ATLAS default for ITk

https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/IDTR-2025-04/

CPU/GPU hybrid setup

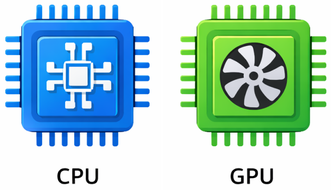

ACTS/traccc coupled project strategy:

- traccc was established as the GPU R&D line of ACTS (see talk by Stephen)
- basic strategy was to re-implement the ACTS tracking concepts for GPUs
- philosophy: no compromise on physics performance
  - automated geometry transcript ACTS/detray
  - same detail of material description
  - Same detail of magnetic field description (covfie)
- Aim to bring CPU/GPU code as close as possible
- 2026: start of re-integration of traccc code into ACTS (and code sharing where possible)

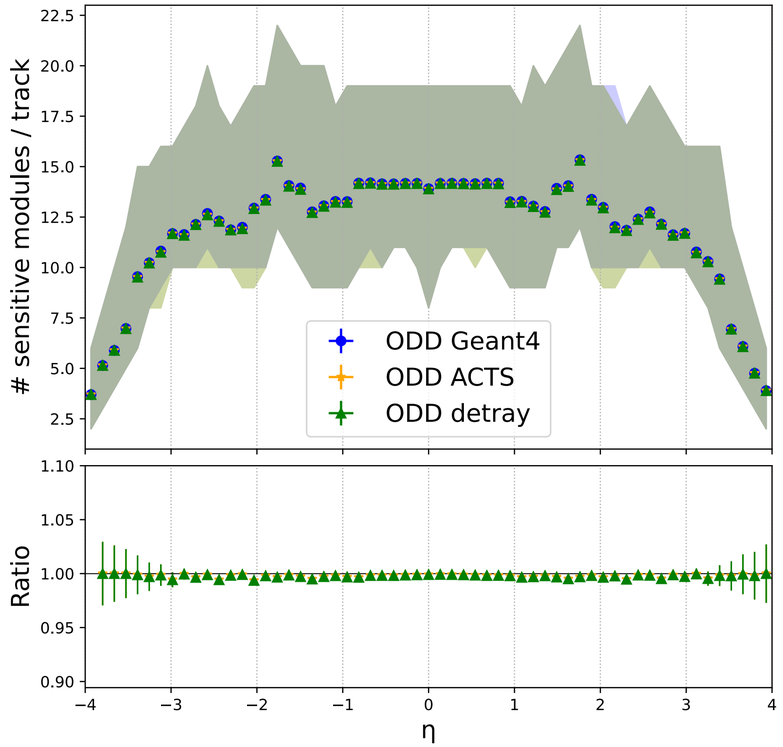

https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-DAQ-PUB-2025-003/

GPU pipelines - flash results

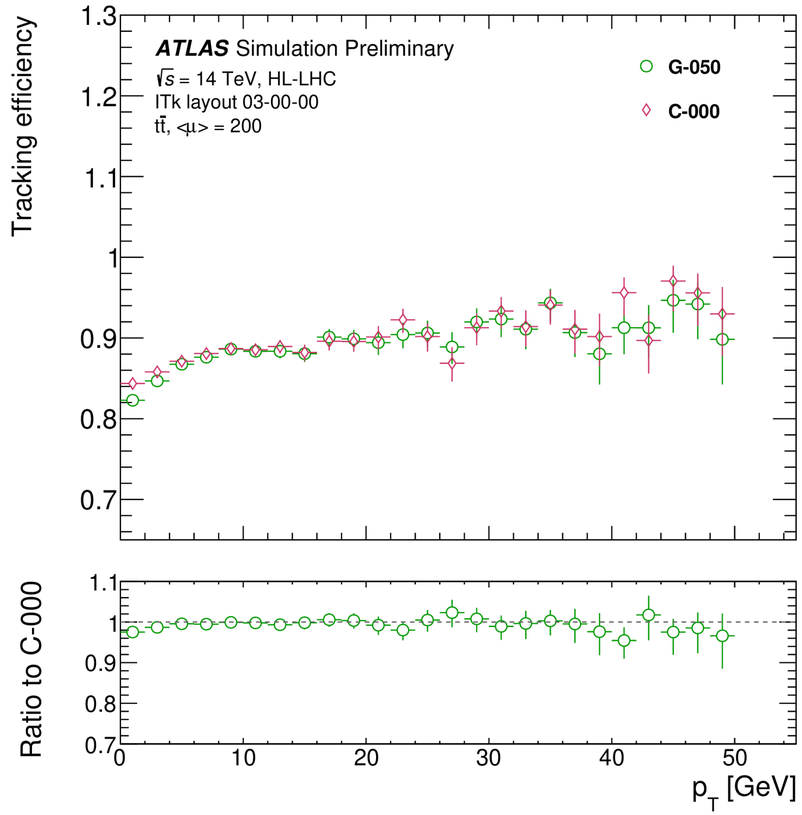

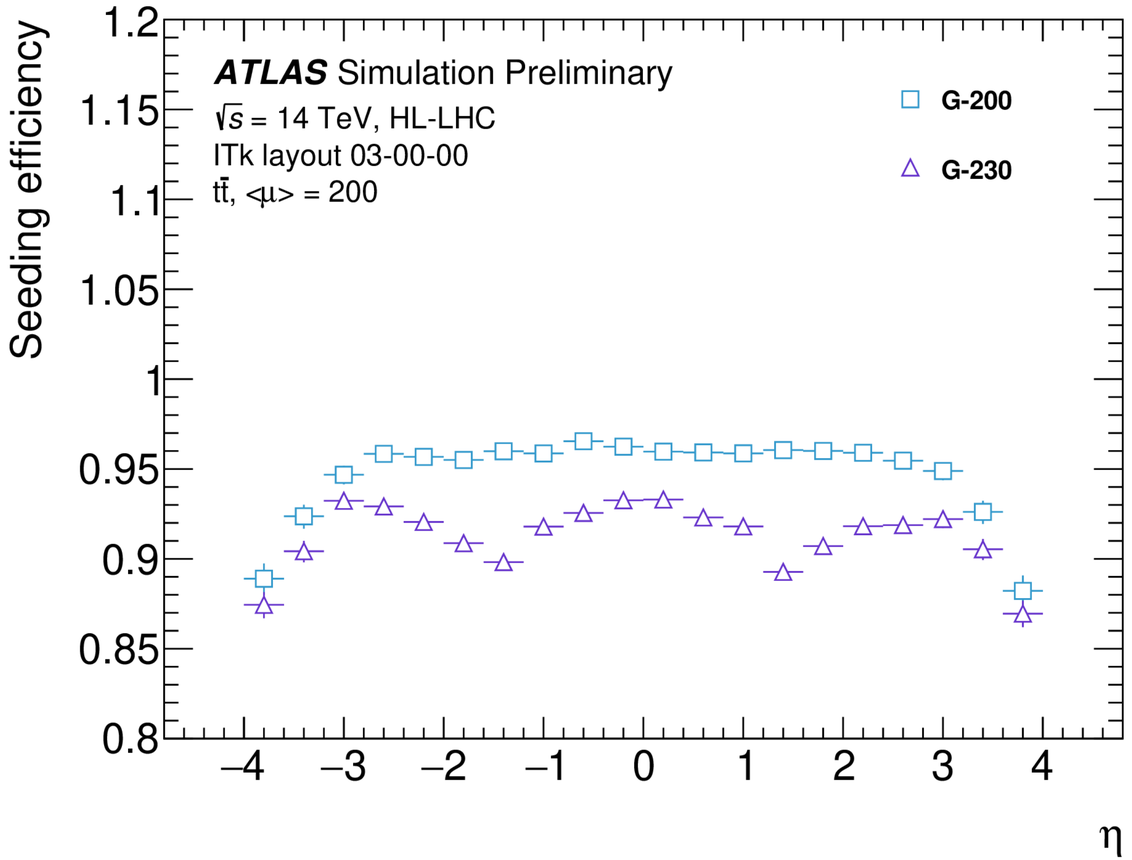

traccc based Tracking is becoming an attractive alternative

- goal to be as close as possible to CPU based results

https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-DAQ-PUB-2025-003/

Offloading to FPGAs

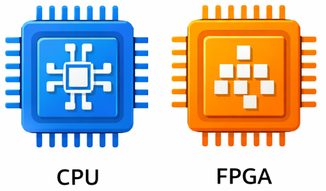

First of its kind R&D program:

- Full integration into Athena established
  - Including CI, performance evaluation, latency
- FPGAs as possible, energy efficient alternatives for HLT Tracking demonstrated
- AMD/Xilinx Alveo PCIe accelerator cards
  - U250 used for testing and development
  - challenge: FPGAs are resource constrained
  - 1.7M LUTs, 3.4M Regs, 97 Mb BRAM, 360 Mb URAM, 12.2k DSPs for U250
  - Installed on CERN hosted testbed, interfaced with Xilinx Runtime

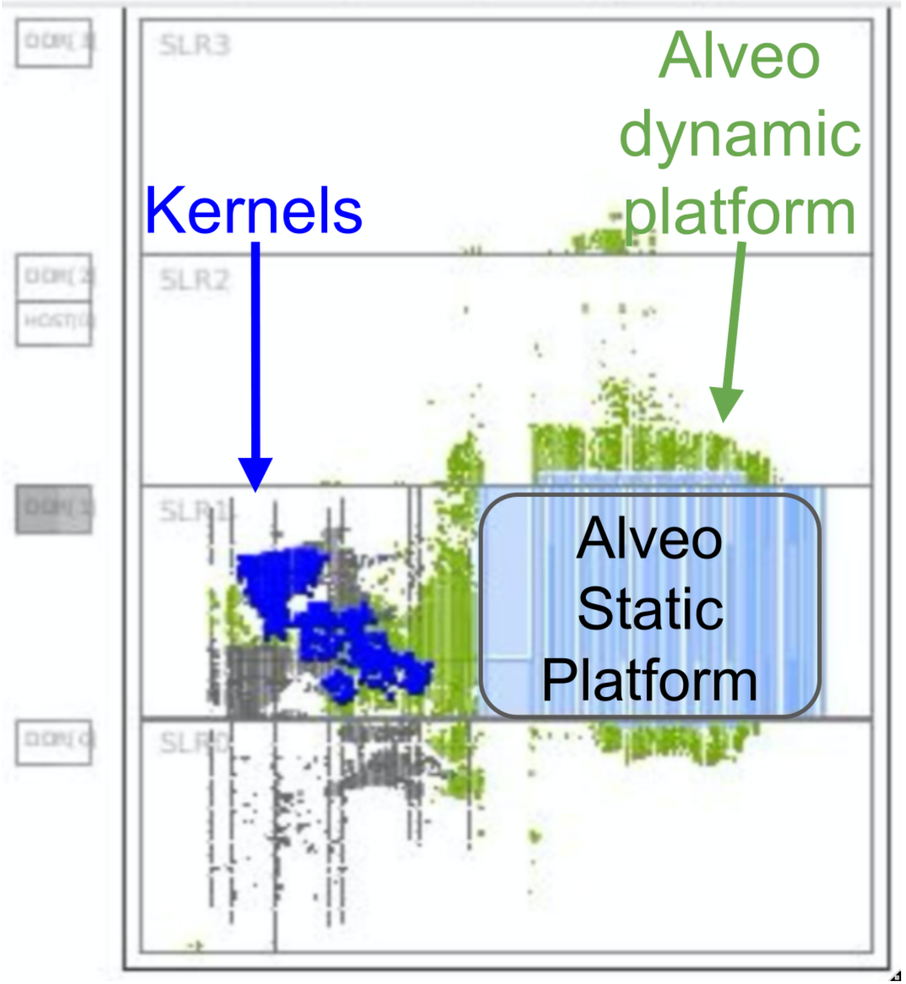

https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-DAQ-PUB-2025-002/

FPGA pipelines - flash results

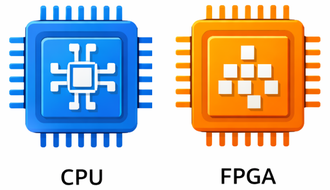

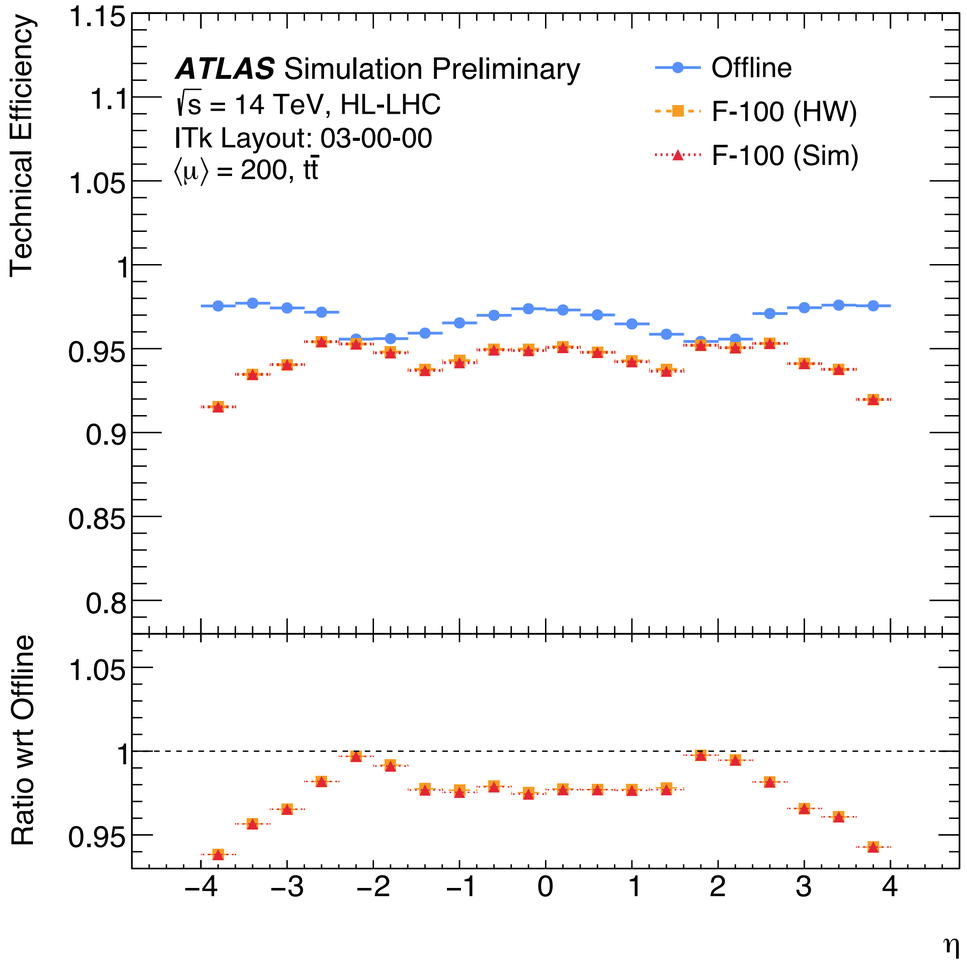

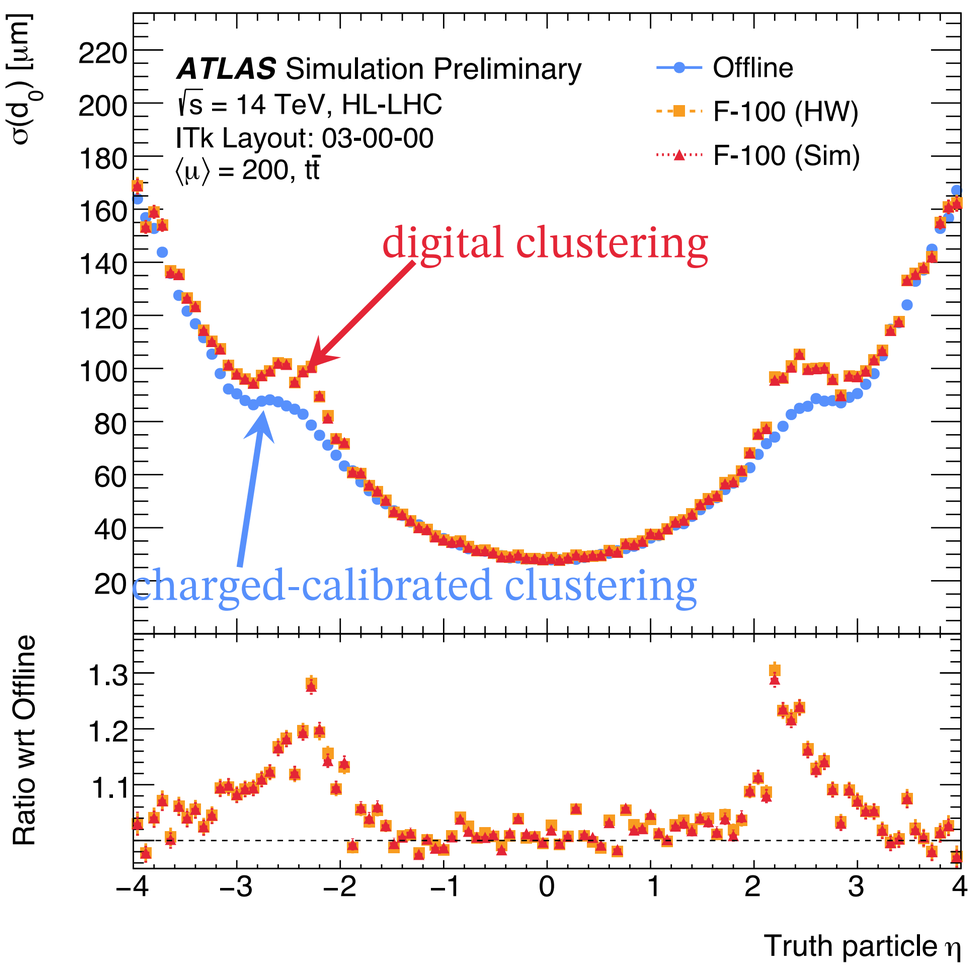

https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-DAQ-PUB-2025-002/

EF Tracking Decision Process

Technology Pipeline reports (internal) handed in to ATLAS in Dec 2025

- Include Tracking performance, cost & power estimates, maintenance & robustness studies

Evaluation of these reports by Technology Choice Committee (TCC)

- Input from Tracking performance group, Computing Coordination, ATLAS at Large

Q1 2026: TCC will give recommendation to TDAQ project

- Upgrade TDAQ project will formulate a Technology choice decision

ATLAS will review the TDAQ decision

HL-LHC timeline

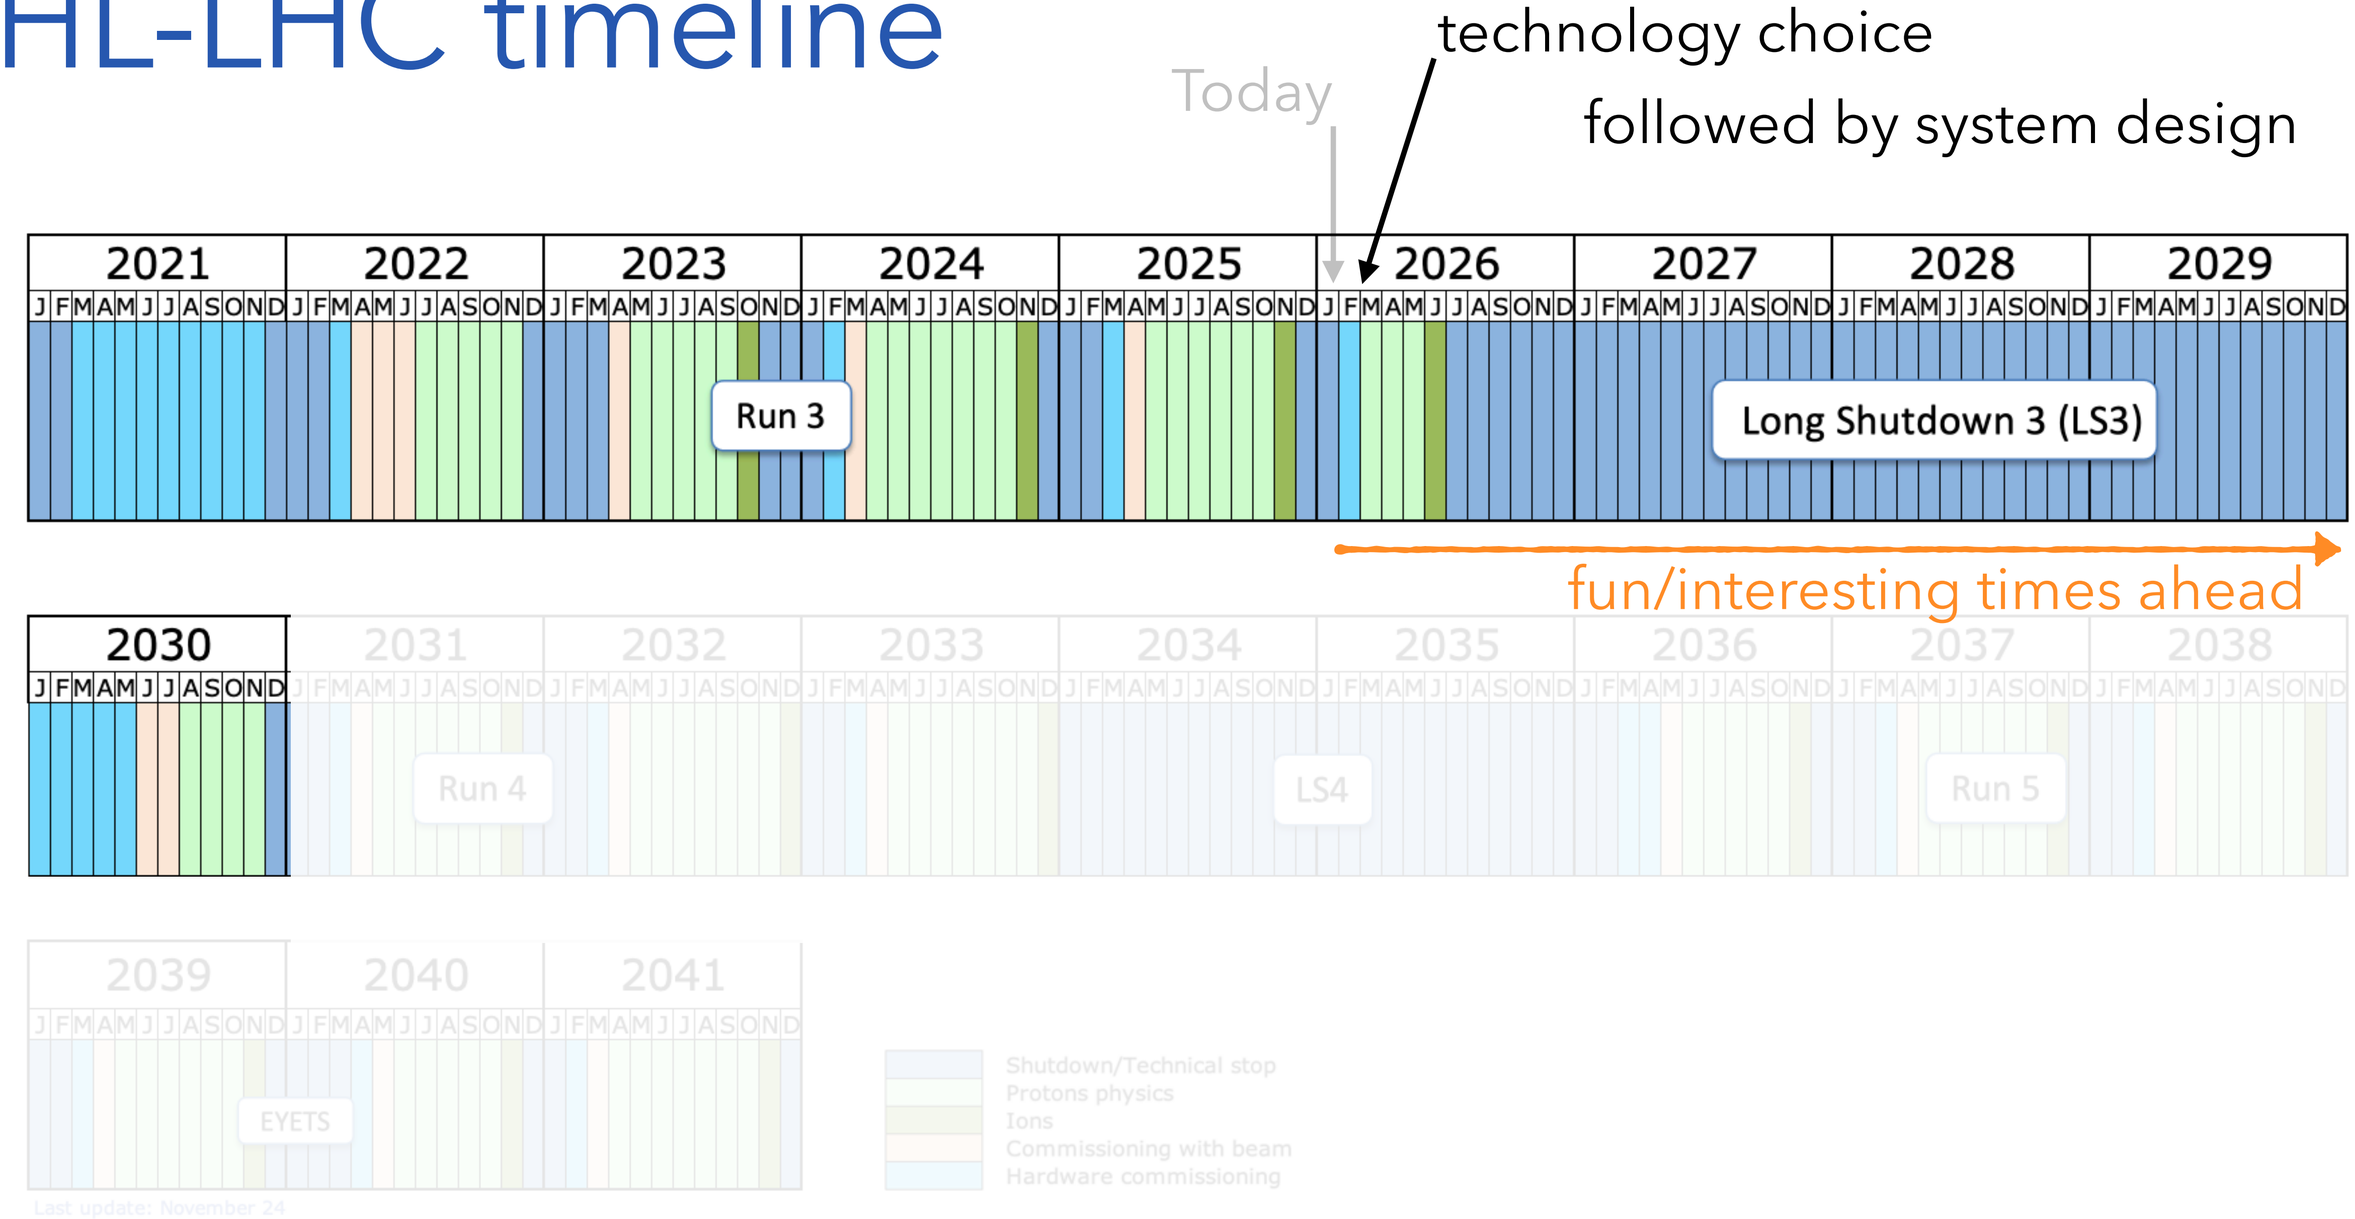

https://lhc-commissioning.web.cern.ch/schedule/LHC-long-term.htm