TDAQ EF Tracking Decision

B. Gorini presentation at the May P2UG in-depth review

Where we stand

Task-force documents and Review Committee Report — *end May/ beg. June*

↓

extended **TDAQ SG decision** — *June 18th*

↓

TDAQ IB (TDIB) endorsement — *July 9th*

↓

TDR Amendment preparation — *Jul through Sep*

↓

Internal Circulation and TDIB endorsement — *TDAQ week Sep 20-24*

ATLAS Circulation — ***Sep 24- Oct 7***

↓

EB approval submission — ***Oct 8***

↓ ***Oct 15***

report to CB

↓ ***Oct - Mar***

LHCC/UCG review process (in parallel)

↓ ***Nov 8-10***

P2UG in-depth review (for the remaining part of TDAQ + info)

↓ ***Feb 21***

CB approval final TDR amendment

↓ ***Mar 16***

CERN RB approval

↓

MoU update to DRC — ***~mid to end Mar***

↓

RRB report — ***Apr 25***

↓

P2UG summary review — ***early May***

- Kickoff meeting — ***~Oct 22***
- Report to LHCC week — ***Nov 15-17***
- Iteration meeting — ***~Dec***
- (2nd Iteration meeting - if needed) — ***~Jan***
- Review — ***~mid Feb***
- Recommendations LHCC week — ***Mar 7-10***

TDR Assumptions: HTT in a nutshell

- TDAQ Phase-II TDR baseline for EF Tracking was Hardware Tracking for the Trigger (HTT)
  - ✦ Regional Tracking (rHTT) based on pattern recognition on AM ASICs
  - ✦ Global Tracking (gHTT) w/ FPGA linear fitting
- Organized in HTT units (6 AMTP + 1 SSTP) each serving a fixed (ηxϕ) region of the detector
- 4 HW boards:
  - ✦ Common TP blade
  - ✦ PRM mezzanine (w/ AM) for pattern recognition (→ AMTP)
  - ✦ TFM mezzanine (FPGA based) for reconstruction/fitting (→ SSTP)
  - ✦ Rear module (RTM) for I/O
- HTTIF based on FELIX hardware (sort of reverse FELIX) to route input raw data to HTT

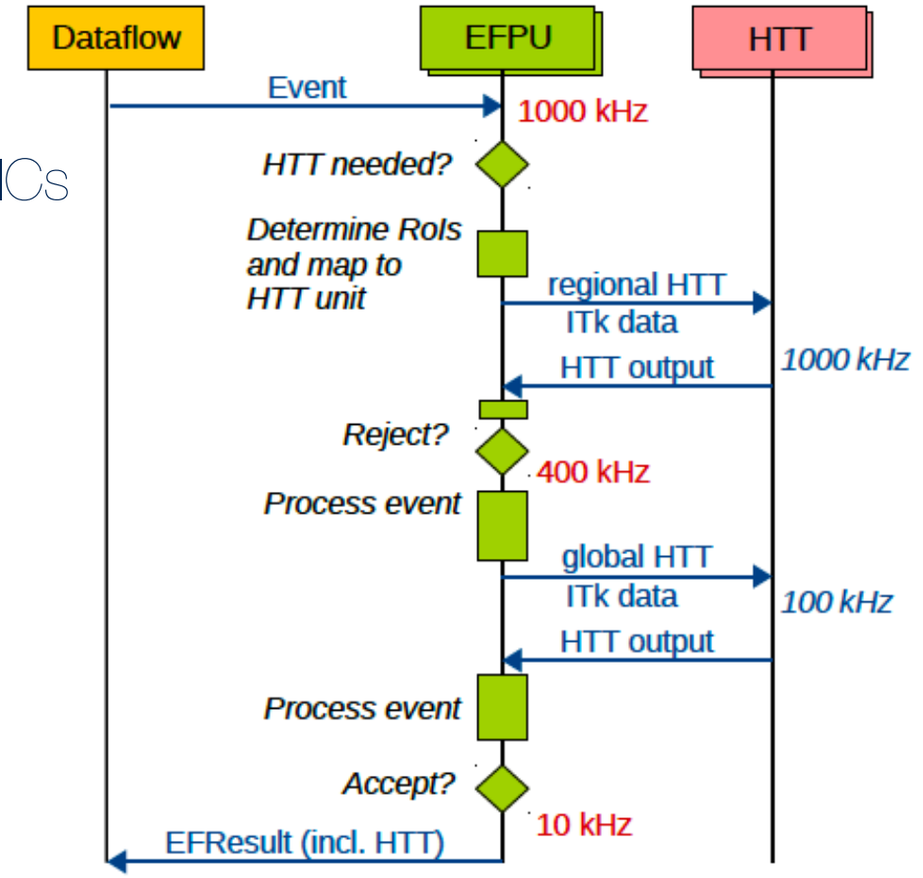

	HTT
Number of HTTIF	24
Number of ATCA shelves for AMTPs	48
Number of AM/SSTP boards per shelf	12
Total number of AMTPs	576
Number of PRMs per AMTP	2
Total number of PRMs	1152
Number of AM ASICs per PRM	12
Total number of AM ASICs	13824
Number of ATCA shelves for SSTPs	8
Total number of SSTPs	96
Number of TFMs per SSTP	2
Total number of TFMs	192

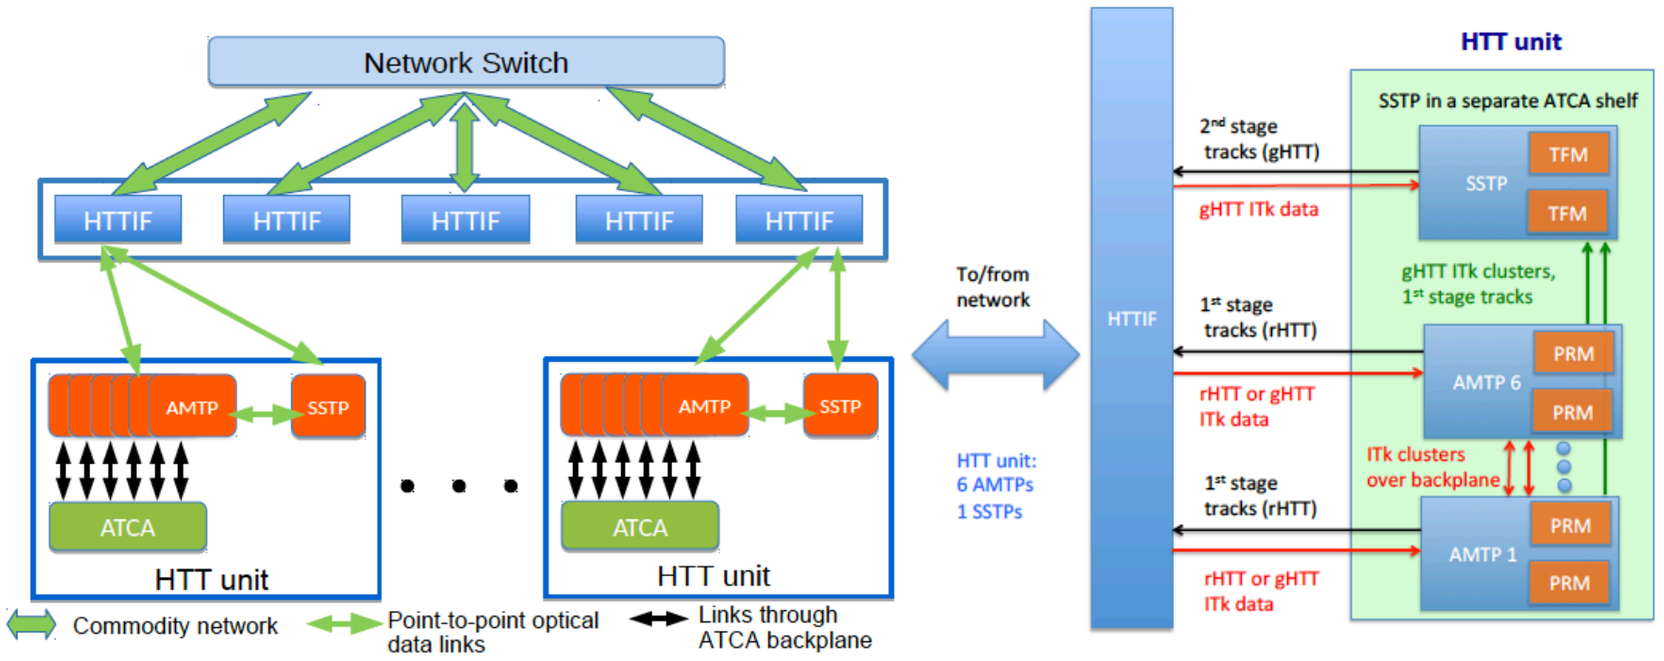

TDR Assumptions

- CORE cost ~17MCHF

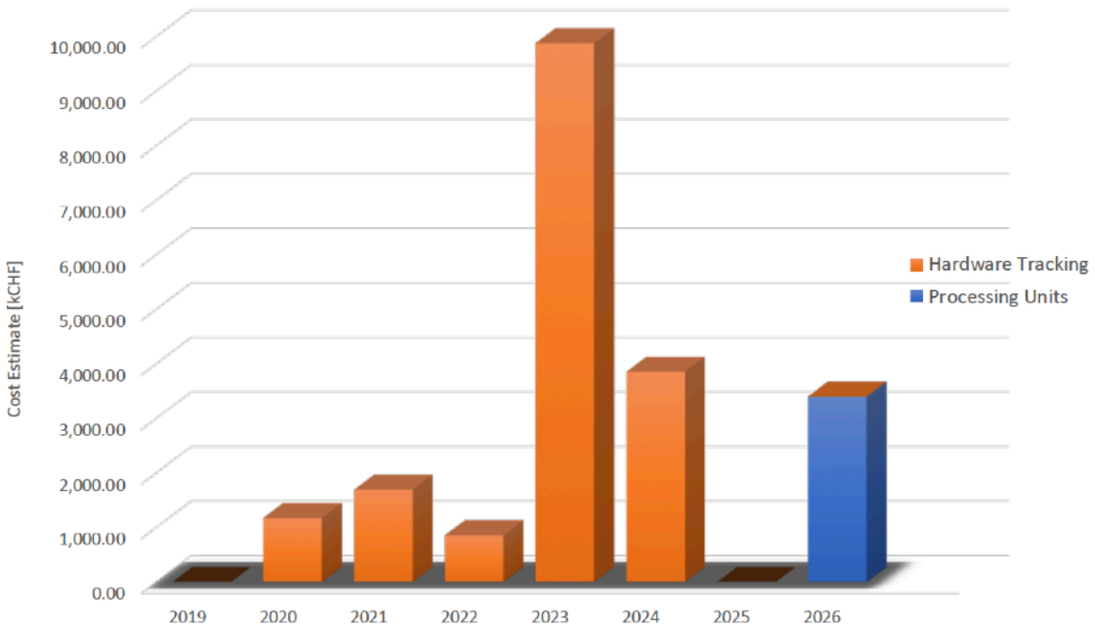

- Required effort: ~190 FTE-years

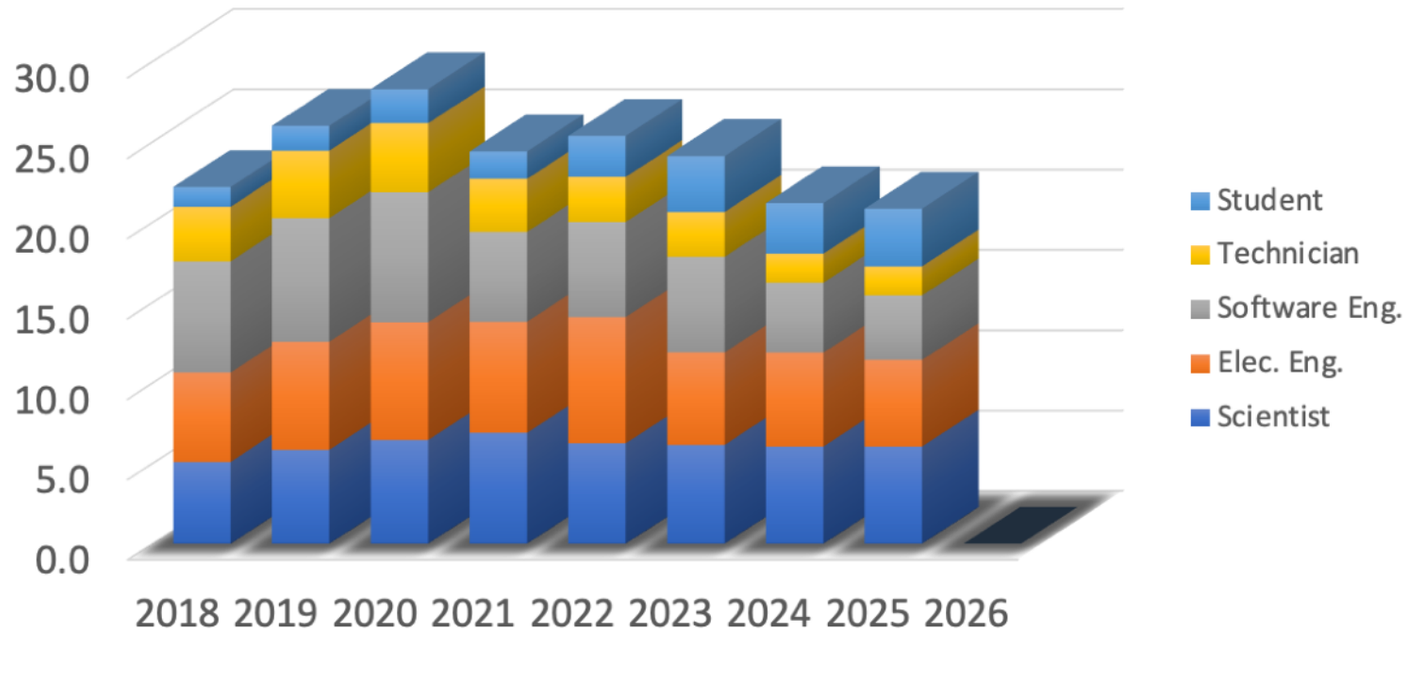

- Main reasons for the TDR baseline choice:
  - Cost effectiveness and the independence of its cost from the commodity computing market
  - Lower power budget and less demanding space requirements,
  - Short latency
  - Capability to evolve the HTT system for use in the hardware based Level-1 trigger should ATLAS need to change to a dual L0/Level-1 Trigger (L1)
  - Extensive experience with the technology within ATLAS

CPU requirements w/o HTT in theTDR

Offline tracking studies with the ITk layout estimated 270 HS06 x seconds per event required to perform full-event tracking at μ = 200 and with a p_T cut of 900 MeV.

- Without global HTT:
  - ✦ 100 kHz of events → 27 MHS06 required
- Without regional HTT:
  - ✦ Full 1 MHz input rate on 5-10% of the entire tracking volume (average RoI size)
  - ✦ Based on the L0 trigger menu with a 2 GeV p_T cut
  - ✦ 13.4 MHS06 required

➡ ~40 MHS06 total

CERN-IT extrapolations of CPU cost (per HS06): average between optimistic and pessimistic scaling scenarios (-10%, -20% per year):

- ✦ 2 CHF/HS06 in 2026

➡ ~80 MCHF CORE

Studies on other commodity accelerator technologies (i.e. GPUs) didn't demonstrate significant advantages cf. CPU only - limited mainly by the overhead spent in preparing/packaging data when offloading the algorithms to the accelerators.

- ➡ The cost-effectiveness of adding GPGPU to the EF depends on the evolution of CPU and GPGPU in terms of price, performance and packaging.

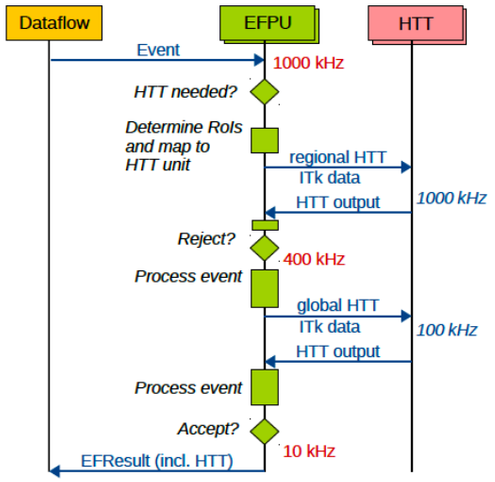

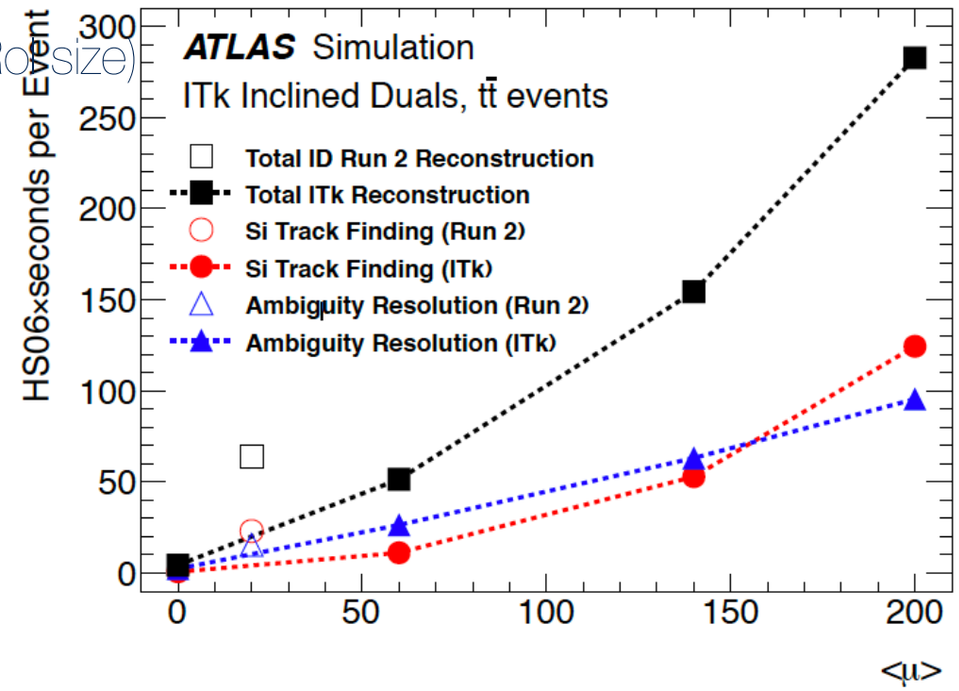

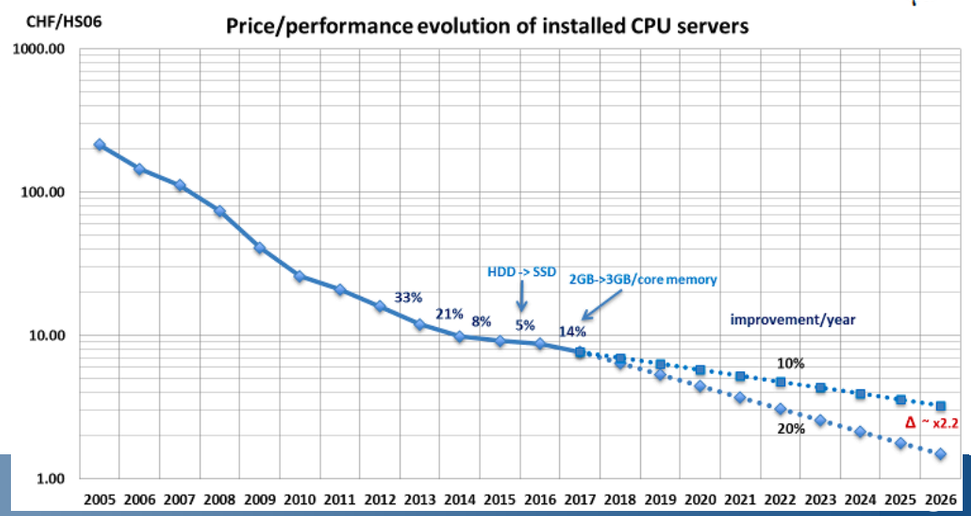

Rebaseline motivations

- Several paradigms changed since TDR:
  - Decision in October to remove the option to evolve to a L0/L1 trigger architecture
    - Removing the low latency constraints and common HW to handle L1 and EF Track
  - Software Tracking improvements
    - 7x faster than assumed in TDR
  - Significant cost reduction of CPU servers
  - New studies and optimizations for both SW and custom HW tracking solution
    - Showing potential for significant cost reductions

- New baseline was required
  - Process established late 2020/early 2021 by TDAQ Phase-II management to drive such a decision
  - To plan effort and CORE spending accordingly

Track Reco Improvements

- Functional **fast ITk track reconstruction prototype** developed ~end of 2019.
  - Released as PUB note (ATL-PHYS-PUB-2019-041)
- Prototype implemented now in the latest 21.9 release w/ the latest layout (e.g. final radius of innermost layer, pixel size) and geometry description (services)
  - Near-offline quality tracking performance with much reduced CPU
- Fully exploits the essential features of the ITk layout:
  - Preserving optimized tracking performance under the harshest pileup conditions
  - High purity tracking for b-tagging and pile-up jet rejection as required for ATLAS HL-LHC physics programme
  - Keep CPU for tracking under control - otherwise, exponential growth would be major issue for offline computing model and for online reconstruction for trigger selection
  - Design choices of Pixel system as documented in the TDR also minimise CPU for reconstruction
  - Sufficient redundancy to allow for some detector defects and irradiation causing loss of sensor efficiency (robust tracking)

Fast ITk Reconstruction

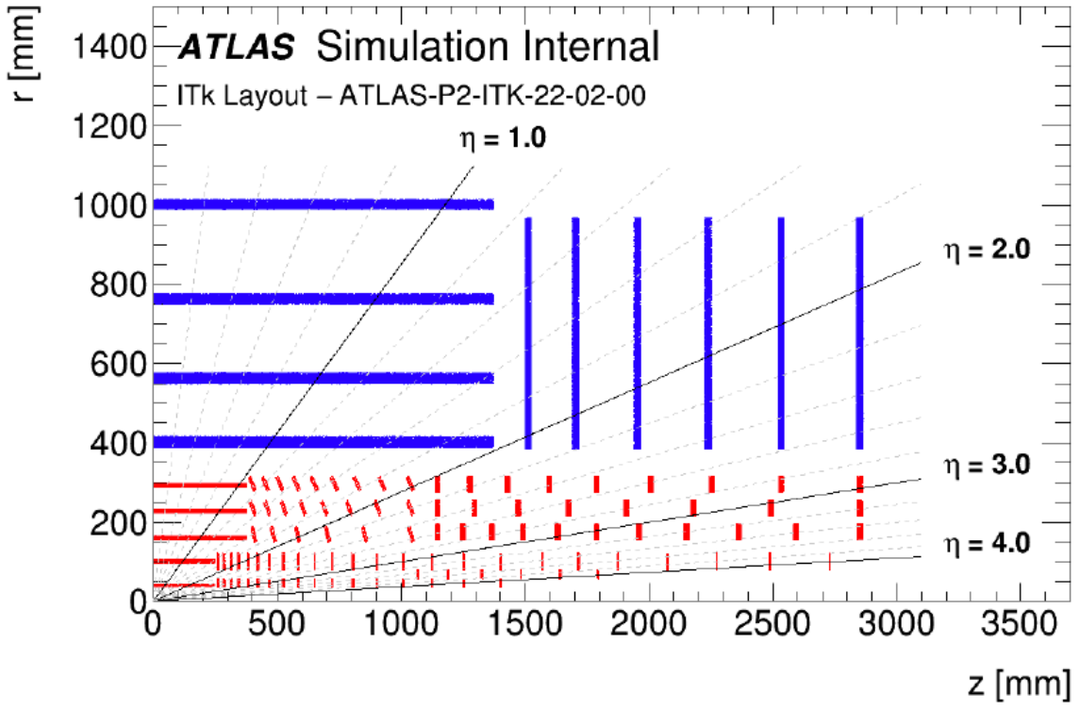

Key to the Fast ITk Reconstruction is the 5 layer Pixel system

- Pixel system covers full η range
- Seed finding in 5 Pixel layers (**redundancy**) allows to drop iteration of seed finding in Strips
- 5 layers allow to **confirm 3 layer seeds in 4th layer** with good efficiency, resulting in high efficiency (**performance**) and high purity (**CPU**) seeding
- High purity Pixel seeded track finding (and improvements in track fitting during track finding) allows to also **drop Ambiguity Resolution** step without major performance impact, leading to an overall CPU **gain of a factor ~4-5 w.r.t. default tracking (in release 21.9)**

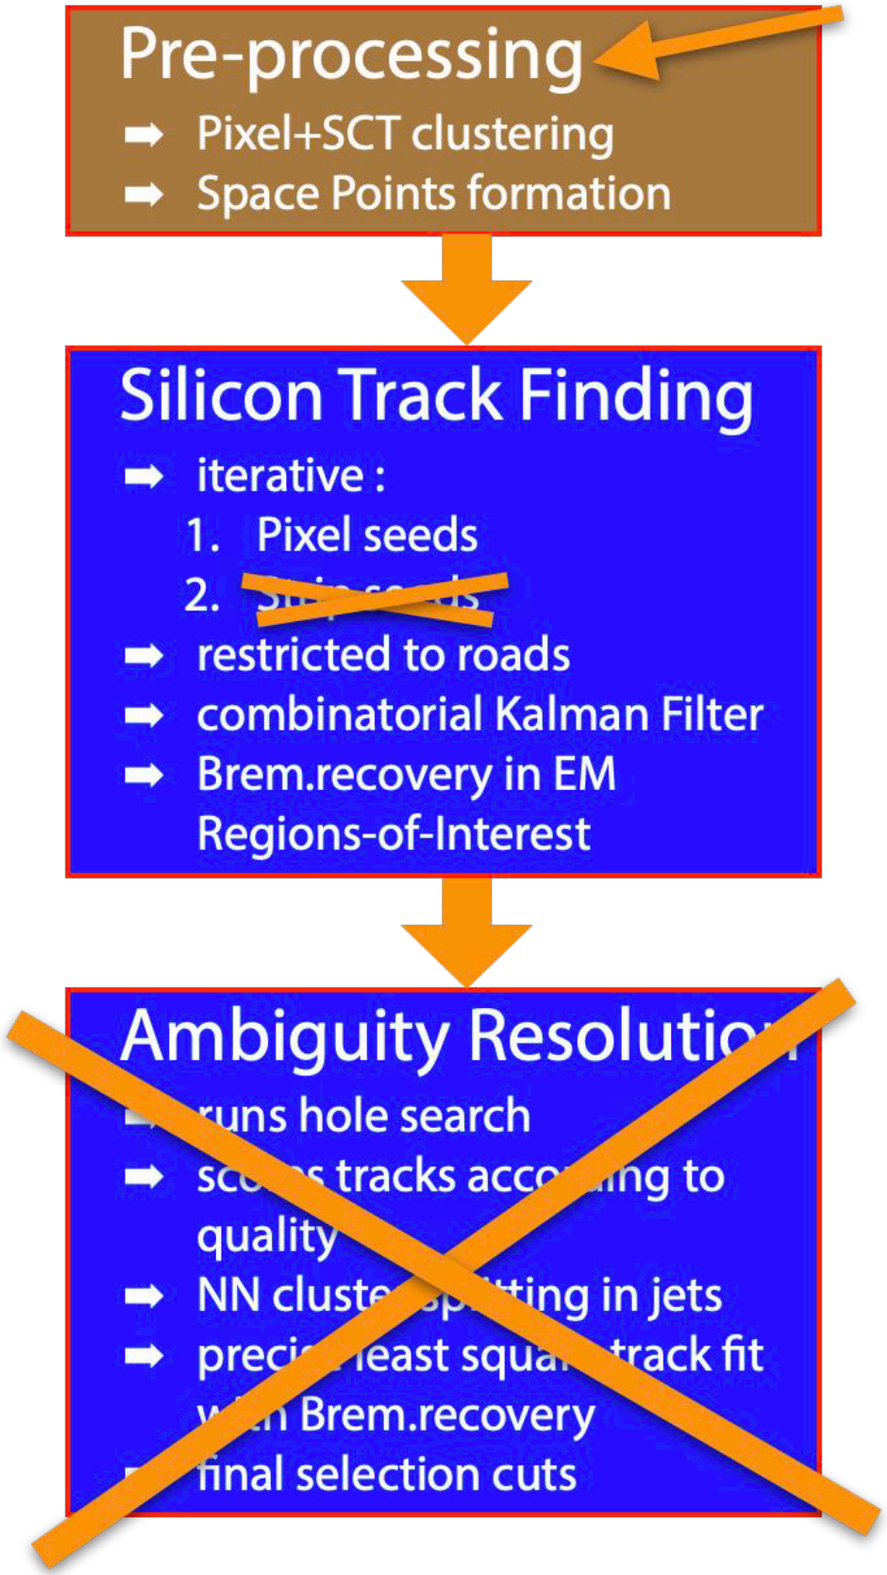

Fast ITk Reconstruction

- CPU time required in HS06 x seconds to reconstruct a $t\bar{t}$ event for the ITk

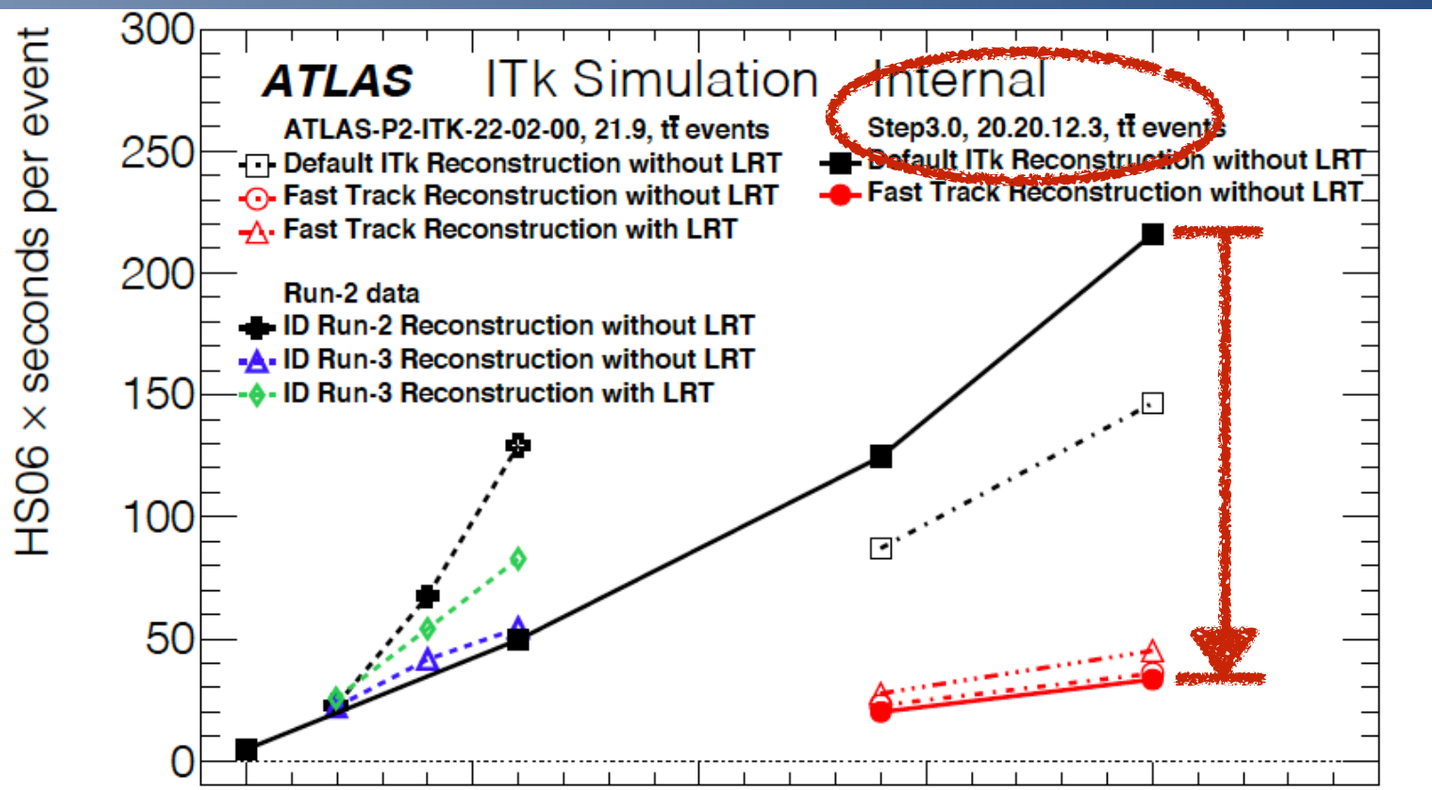

<μ>	Tracking	Release	Byte Stream Decoding	Cluster Finding	Space Points	Si Track Finding	Ambiguity Resolution	Total ITk
140	default	20.20	2.2	17.1	6.0	41.1	58.2	124.6
	fast			4.5	0.9	12.4	-	20.0
200	default		3.2	26.3	8.6	85.8	92.0	**215.9**
	fast			6.3	1.2	22.6	-	**33.3**
140	default	21.9	2.2	6.4	3.5	31.6	43.4	87.1
	fast			6.1	1.0	13.4	-	22.7
200	default		3.2	8.3	4.9	66.1	64.1	**146.6**
	fast			8.1	1.2	23.2	-	**35.8**

Fast ITk Track Reco: performance comparison

- Irradiation leads to various effects, in particular to random loss of hits, causing loss of clusters and hence creating "holes" in the tracks
- Introduce both sensor inefficiencies due to irradiation and sensor failures randomly

- **Defects randomly introduced:**
  - ✦ inefficiencies: 3% Pixel + 1% Strip
  - ✦ 15% sensor failures in both Pixels and Strips
- **Default tracking:** reduction of ~10% for in the barrel section ($|\eta|<1.6$)
- **Fast tracking:** stronger effect, reduction up to ~20%

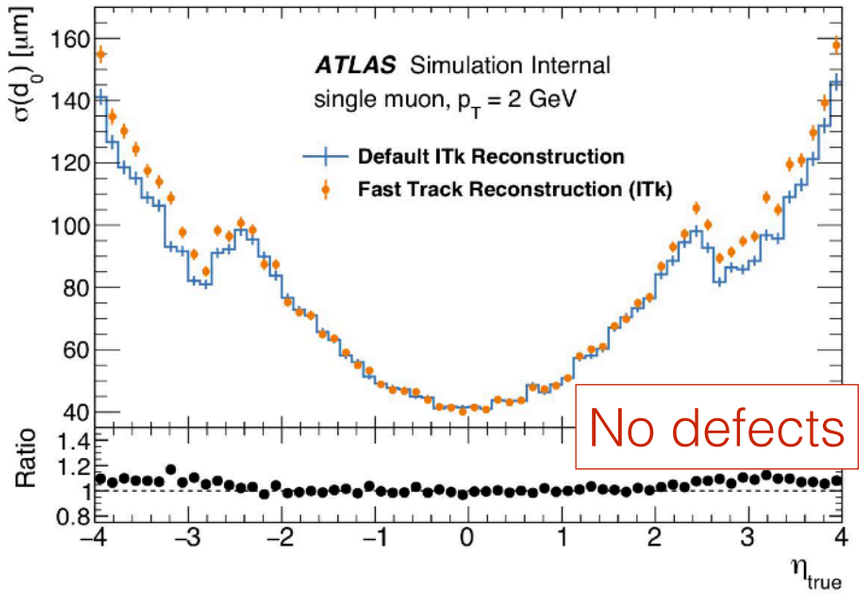

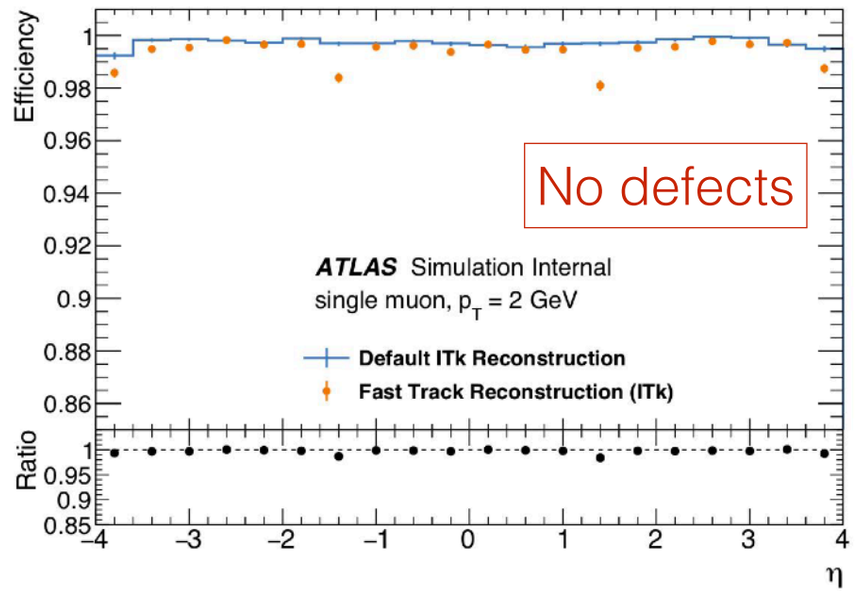

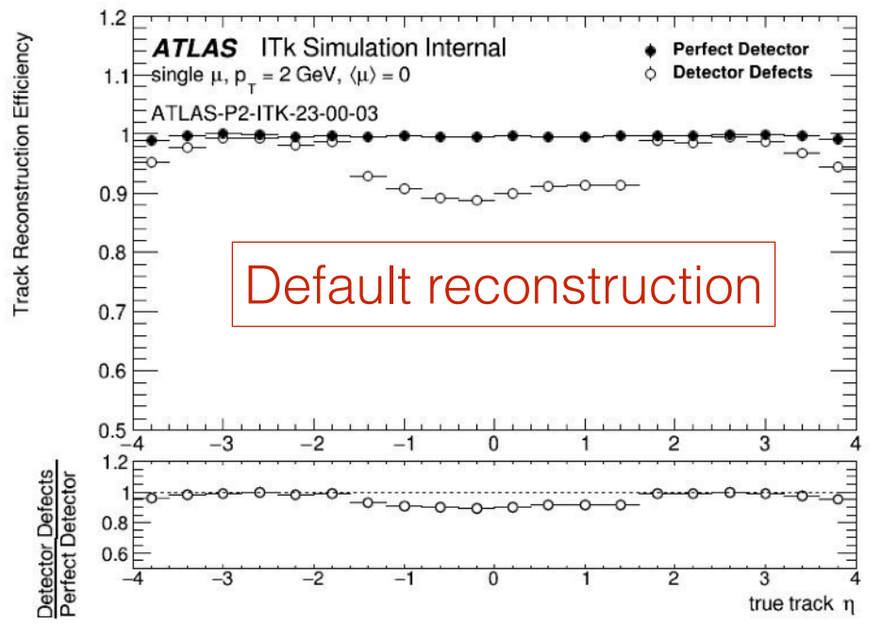

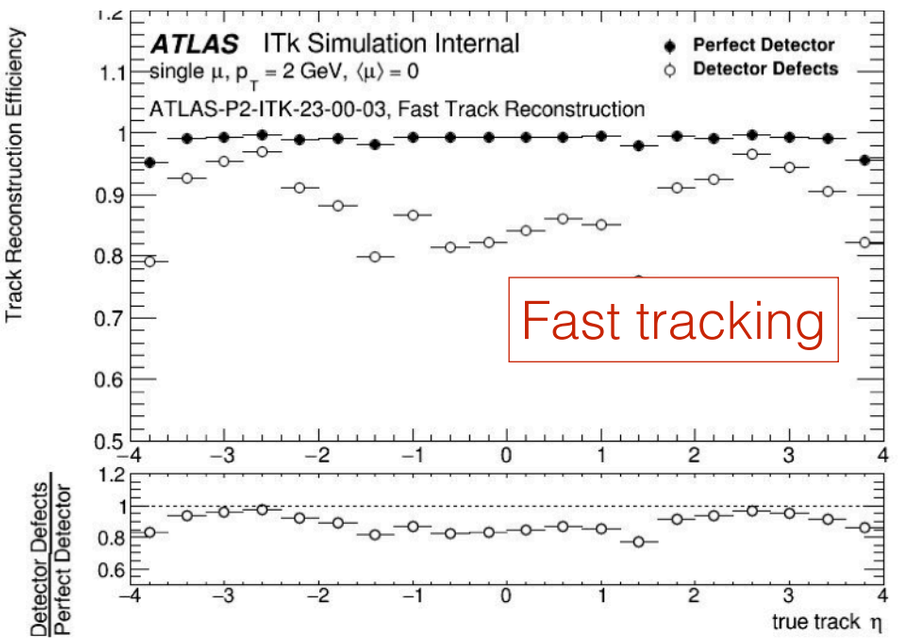

CPU server reduction cost

- Price volatility of commodity components introduce large uncertainties on predictions for the CPU server costs
  - ✦ Slowdown in feature size reduction in processors (14nm → 10nm → 7nm)
  - ✦ Intel issues with 10nm foundries in 2018/2019 still persists
- In particular RAMs and SSD devices

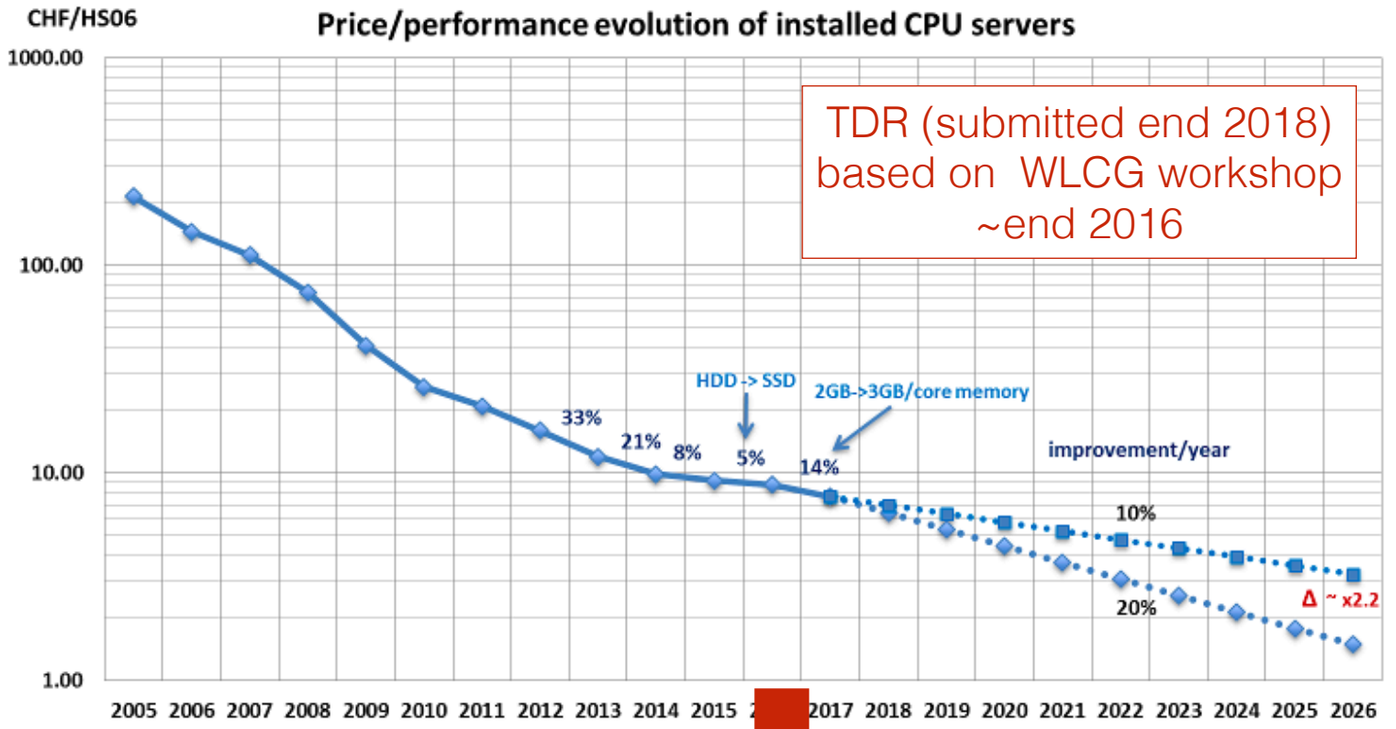

Samsung DIMM 16GB, DDR4-2133, CL15-15-15, reg ECC

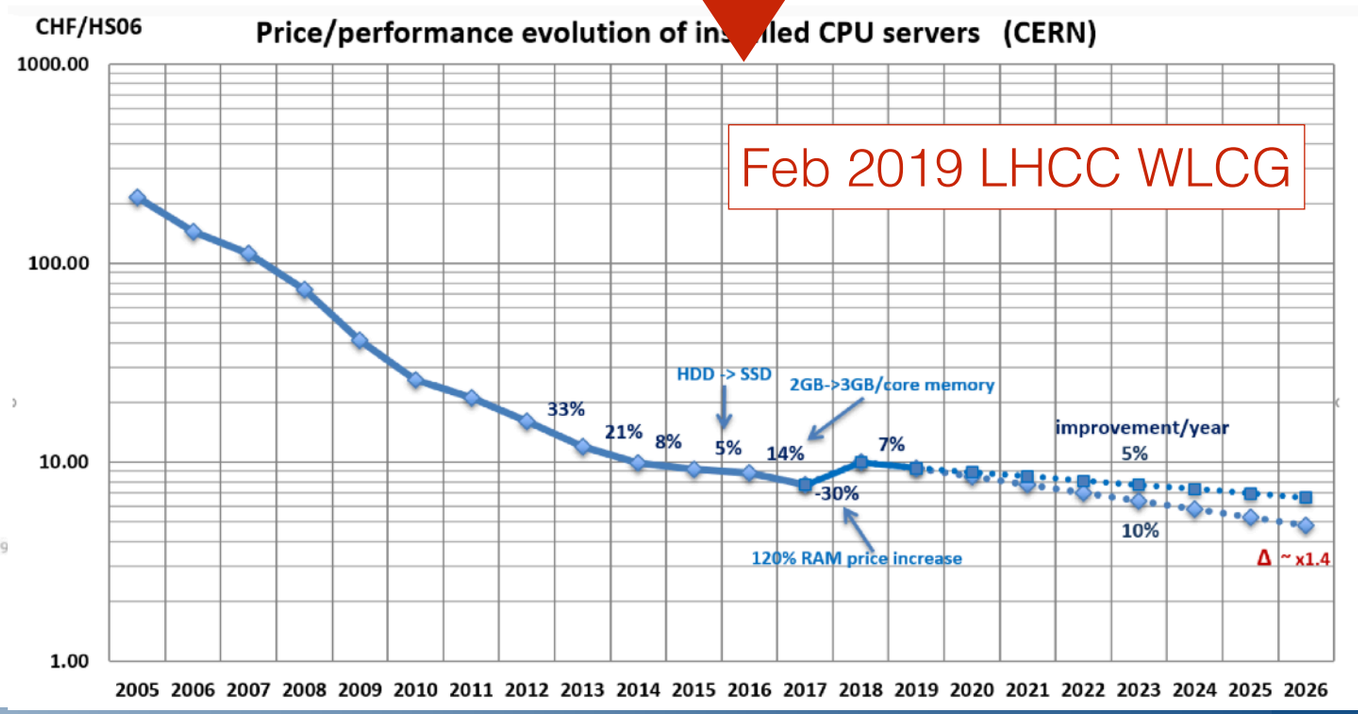

CPU cost reductions

- CERN-IT new contract w/ AMD has lowered significantly the server costs and future predictions
  - ✦ AMD already using 7nm technology in server applications (Intel still 14 nm, now migrating to TSMC)
  - ✦ In 2022 AMD release Zen-4 architecture servers in 5 nm technology - same process as Apple M1 in laptops/desktops
  - ✦ TSMC confirmed volume production w/ 3 nm lines already in 2nd half of 2022

	Model	Process	TDP/CPU	Real cores/node	HS06	Watt/HS06
Run-2 EF	2x Intel 2660v5 (2 GHz, 64GB)	14 nm	105 W	2x14	638	0.38
Run-3 EF	2x AMD EPYC 7302 (3 GHz, 256 GB)	7 nm	155 W	2x16	1158	0.33
	2x AMD EPYC 7502 (2.5 GHz, 512 GB)	7 nm	180 W	2x32	2010	0.26
	2x AMD EPYC 7702 (2 GHz, 1024 GB)	7 nm	200 W	2x64	3050	0.19

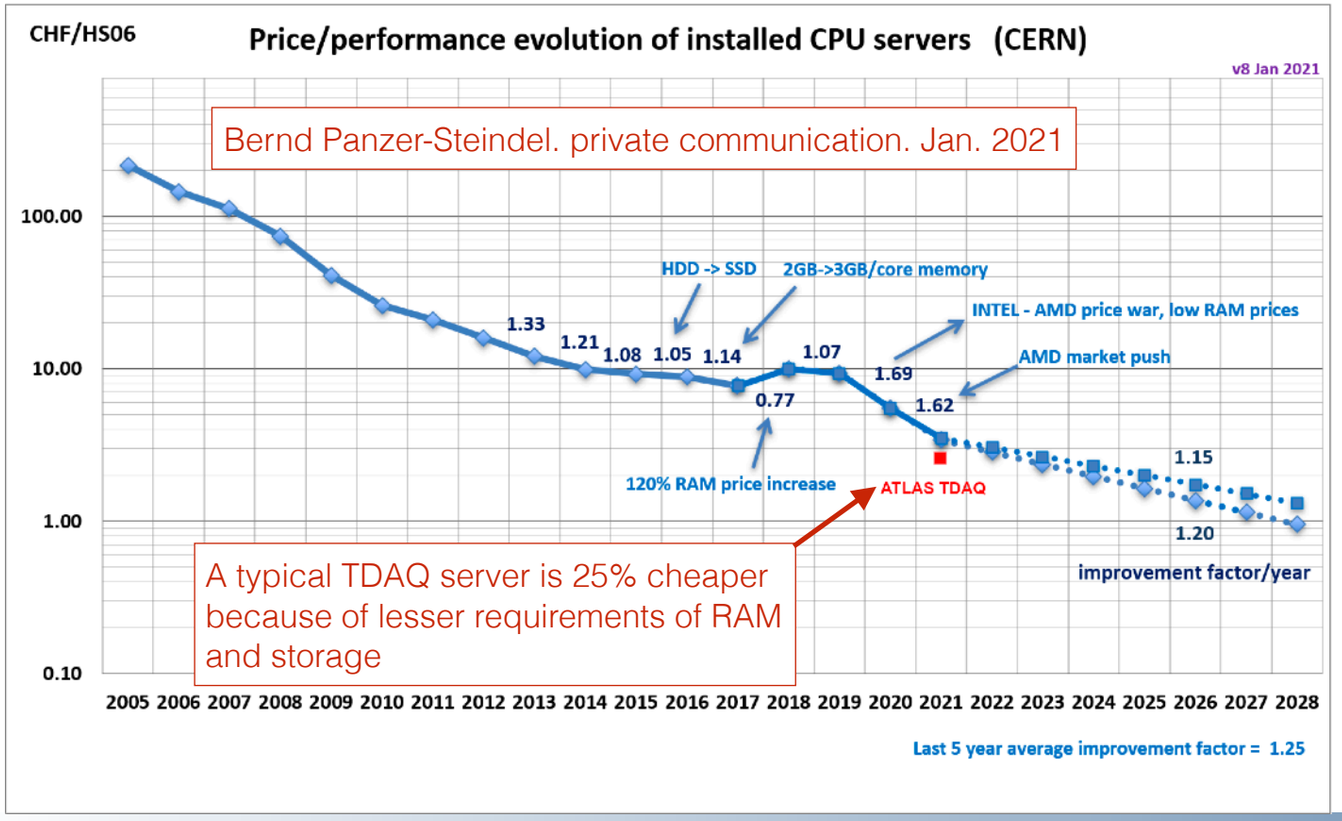

	improvement (%/year)	CHF/HS06 2028	CHF/HS06 2032
pessimistic	-10	1.3	0.8
realistic	-15	1.0	0.5
optimistic	-20	0.7	0.2

CPU costs extrapolations for EF Farms

	Expected Compute		
	2027	**2032**	**Total Run-4/5**
Tracking Requirements [MHS06]	3.1	4.8	
Other Reconstruction [MHS06]	1.9	2.3	
Total [MHS06]	**5.0**	**7.1**	
Compute Cost [CHF/HS06]	1.0	0.5	
Cost Range [CHF/HS06]	[0.7-1.3]	[0.2-0.8]	
Cost Tracking [MCHF]	**3.1**	**0.9**	**4.0 [2.5-5.4]**
Cost Other Reco. [MCHF]	1.9	0.2	2.1 [1.4-2.8]
Cost Total [MCHF]	**5.0**	**1.1**	**6.1 [3.9-8.2]**
Existing Farm through M&O Rolling Replacement [MHS06]	2.9	5.9	
Upgrade extension Tracking [MHS06]	1.4	0.9	
Upgrade extension Other Reconstruction [MHS06]	0.8	0.4	
Total Upgrade extension [MHS06]	**2.2**	**1.3**	
Cost Tracking [MCHF]	**1.4**	**0.5**	**1.9 [1.1-2.2]**
Cost Other Reco. [MCHF]	0.8	0.2	2.0 [0.6-1.2]
Cost Total [MCHF]	**2.2**	**0.7**	**2.9 [1.7-3.4]**
TDR Total Cost (HTT+EF) [MCHF]			**20.8 (17.4+3.4)**

My estimates from SW-only report
Rounding errors
See slide #28 for 'official' assessment
and CORE estimates by TFs and committee

Options considered

Process launched in December 2020 w/ two task-forces where engineering studies were needed (options #1, #2), while development inside existing project's organisation for option #3

1. Optimised custom HW architecture
   - Exploit unconstraint latency requirement
   - Re-optimize AM pattern banks
   - FPGA as alternative (to AM) for pattern recognition
   - Simplify HW design maintaining ATCA as base platform
   - Led by A. Annovi (former HTT co-coord.)

2. Heterogeneous architecture
   - Commodity platform of CPUs and accelerators
   - Through initial assessment based on Hough Transform implemented in FPGAs
   - Led by S. Majewski (PPES co-coord.)

3. Software Tracking
   - Optimized SW fast tracking (i.e. near-offline quality tracking performance with much reduced CPU)
   - Revision of CORE cost extrapolation based on IT prediction released early 2021
   - Led by F. Winklmeier (Event Filter Farm coord.)

Mandate

- The three coordinators synchronised weekly with the UPR mgmt. and reported at each eTDSG
- Open meetings were organised ~monthly to to inform and collect feedback from the whole TDAQ community
- Produce (for each option) an engineered solution to be used to prove the feasibility of the specific approach and for comparison between alternatives
- Each TF and the EF team (SW-only) were charged to study all the following points:
  - Technical feasibility
  - Estimated tracking performance
  - Operational procedures: calibration, alignment, monitoring, etc.
  - Opportunities for improvement
  - Risks
  - Resource requirements
- Performance requirements were reviewed by the PPES in the early spring and original mandate amended accordingly

Optimization of custom HW solution

- Custom hardware system based on an optimized HTT design
  - ✦ An array of HTT units (ATCA shelves) each covering a detector eta-phi region
  - ✦ A single ATCA card vs 4 cards in the original design
  - ✦ Each blade performs all tracking algorithms
  - ✦ Clusters are found and shared within a shelf

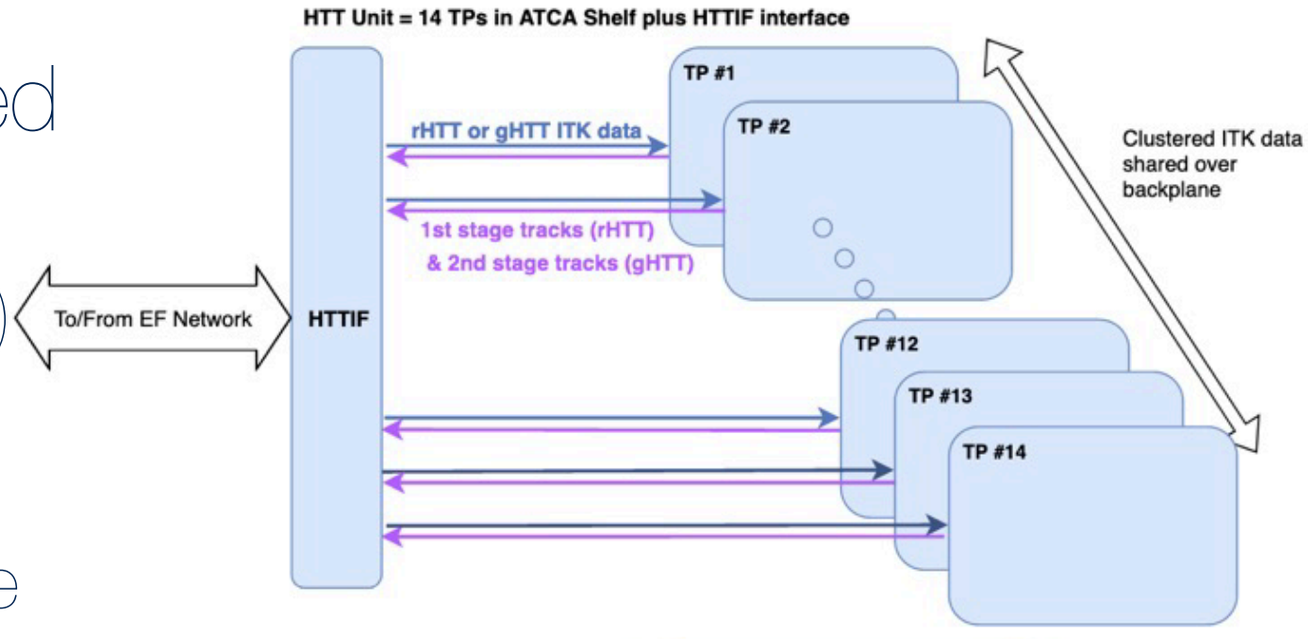

- Two approaches considered for pattern recognition
  - ✦ AM ASIC vs Hough Transform on FPGA

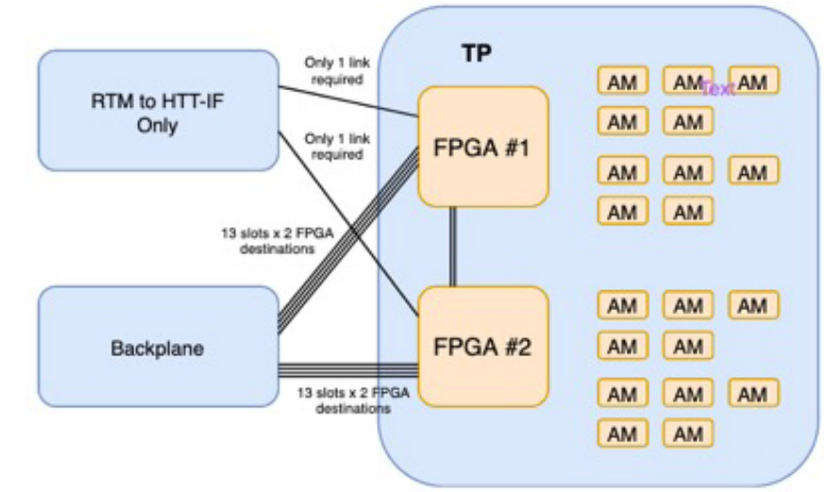

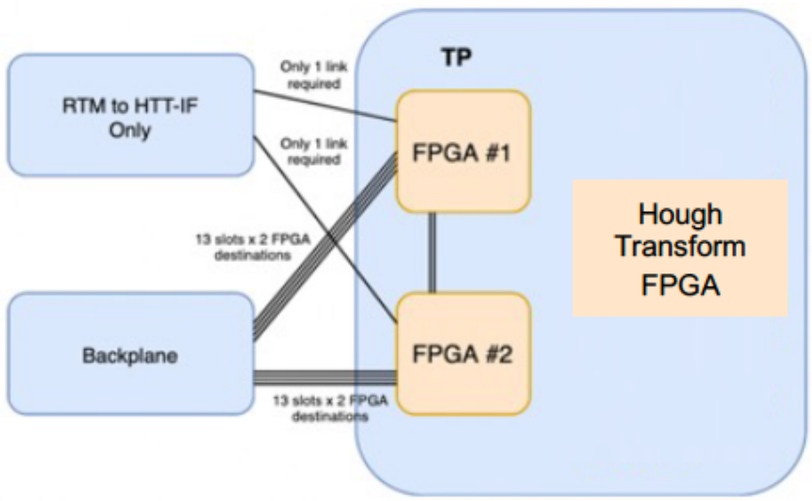

Optimization of custom HW solution

- 224 TP modules max. >596 AMTPs + 96 SSTPs modules in TDR
- 8 racks vs. 24 in TDR
- Power estimate for the AM and FPGA-based options (135-190 kW vs. 310-400 kW in the TDR)

Item	Numbers				
Number of shelves	**8**	**10**	**12**	**15**	**16**
Number TP per shelf	14	14	14	14	14
Total number of TP	**112**	**140**	**168**	**210**	**224**
Shelves per rack	2	2	2	2	2
HTT-IF PC size (U)	2	2	2	2	2
Number TP per HTT-IF PC	7	7	7	7	7
HTT-IF PCs per rack	4	4	4	4	4
Total HTT-IF PCs	**16**	**20**	**24**	**30**	**32**
ConMon PC size (U)	1	1	1	1	1
Number TP per ConMon PC	14	14	14	14	14
ConMon PCs per rack	2	2	2	2	2
Total ConMon PCs	**8**	**10**	**12**	**15**	**16**
Total racks (rounded up)	**4**	**5**	**6**	**8**	**8**
Spare racks (MPV only)	1	1	1	1	1

Item (W)	CBE	MPV
TP blade		
- FPGA (x2)	121	157
- 20 AM-ASICs	70	98
- Other (incl. SoC)	13	18
- Low V DC/DCs Inefficiency (15%)	37	49
- 48/12V DC/DC Inefficiency (5%)	12	16
TP blade total	**254**	**339**
RTM	5	7
Total	**259**	**346**
Total for 14 slots	**3626**	**4844**

Item (kW)	CBE	MPV
HTT ATCA (8 shelves)	29	38.8
Fans (estimated)	8	11.2
ATCA Rack 48V AC/DC - 95.5% eff. (3 sh/rack)	0.4	0.6
HTTIF PCs (0.3kW CBE ea, x16)	4.8	6.7
ConMon PCs (0.15kW ea, x8)	1.2	1.7
EFPUs (0.37MHS06 @0.25W/HS06 CBE)	92.5	129.5
Total	**135.9**	**188.4**

Item (W)	CBE	MPV
TP blade		
- FPGA (x2)	150	162
- Other (incl. SoC)	13	18
- Low V DC/DCs Inefficiency (15%)	30	33
- 48/12V DC/DC Inefficiency (5%)	10	11
TP blade total	**203**	**224**
RTM	5	7
Total	**208**	**231**
Total for 14 slots	**2912**	**3234**

Item (kW)	CBE	MPV
HTT ATCA (12 shelves)	34.9	38.8
Fans (estimated)	12	16.8
ATCA Rack 48V AC/DC - 95.5% eff. (2 sh/rack)	0.4	0.4
HTTIF PCs (0.3kW CBE ea, x24)	7.2	10.1
ConMon PCs (0.15kW ea, x12)	1.8	2.5
EFPUs (0.37MHS06 @0.25W/HS06 CBE)	92.5	99.9
Total	**148.8**	**168.5**

Heterogeneous Arch.

- Goal was to establish the viability of a heterogeneous commodity approach using FPGAs as coprocessors

- Technical choice for (only!) proof-of-concept study dictated by the limited time available
  - ✦ Full ITk event loaded into single commercially available accelerator card, then processed using Hough Transform in firmware
  - ✦ Precision track fit using the fast tracking Kalman filter developed for Phase-II

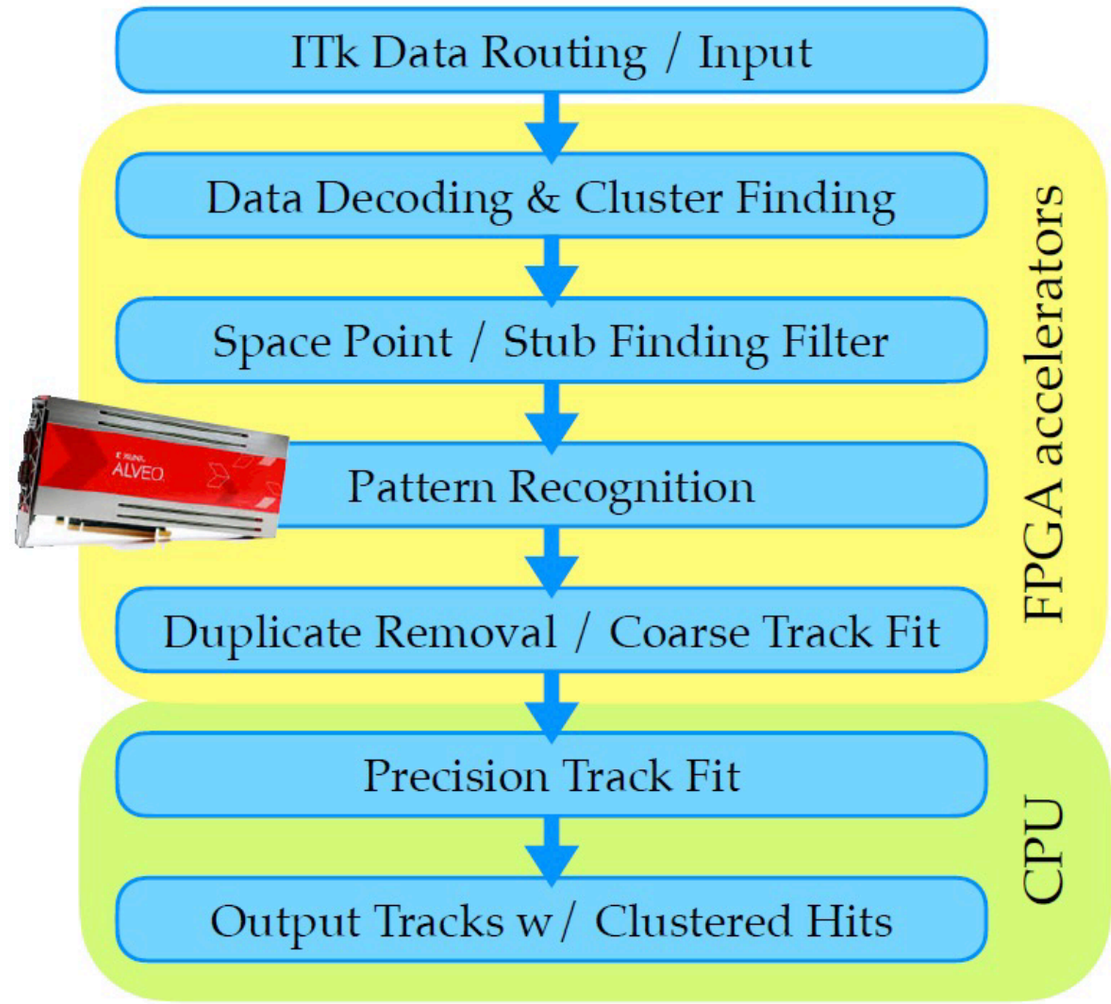

Heterogeneous Architecture

- Two independent implementations of Hough Transforms (Penn/UCI and Bologna) and 4 scenarios (HT configurations - binning)
  - ✦ Realistic FPGA resource usages

Firmware Block			LUT (%)	FF (%)	BRAM (%)	URAM (%)	DSP (%)
			Logic Functions		Memory Functions		
PCIe			0.6	0.6	0.3	–	–
Clustering			1–4	0.14–0.51	1.3–5.4	–	–
Stub-Finding			0.2	0.05	0.1	–	–
Slicing Engine			0.1	0.07	–	13	–
Pattern Recognition:	ANL (0.2×0.2 region)	s1	11	3	1.5	–	1.8
		s2	39	10	5	–	1.8
		s3	39	10	5	–	1.8
	Bologna/Uppsala (0.2×0.2 region)	s1	15	9	1	–	8
		s2	59	30	1	–	21
		s3	52	28	1	–	8
	Penn/UCI (0.2×0.8 region)	s4	12	7	27	–	1
Fake Rejection (NN)			8	1	0.02	–	29
Duplicate Removal			1	1	–	–	–
Track Fitting			~ 10		–	–	~ 10
Monitoring (IPBus)			~ 1		–	–	–
Case 1: Track Extension and Fitting in CPU							
Totals (Case 1)		s1	33–41	18–24	3–7	13	41–47
		s2	61–84	25–45	3–11		41–60
		s3	61–77	25–43	3–11		41–47
		s4	34–37	22	29–33		40
Case 2: Track Extension and Fitting in FPGA, only precision fit in CPU							
2nd-Stage Fitting			~ 10		~ 30	–	~ 15
Totals (Case 2)		s1	43–51	28–34	33–37	13	56–62
		s2	71–94	35–55	33–41		56–75
		s3	71–87	35–53	33–41		56–62
		s4	44–47	32	59–63		55

- System size

Table 2.9: The size of accelerator part of the system.

	ANL (Scenario 2)	Bologna/Uppsala (Scenario 2)	Penn/UCI (Scenario 4)
# of Accelerator Cards	510	281	202
# of Accelerator Cards (LRT 5 kHz/20 kHz)	34/136	19/75	13/54
# of PC servers	64	36	26
# of PC servers (LRT 5 kHz/20 kHz)	4/17	2/10	2/7
# of Racks	7	4	3
# of Racks (LRT 5 kHz/20 kHz)	0/2	0/1	0/1
# of Accelerator Card Total	544/646	300/356	215/256
# of PC server Total	68/81	38/46	28/33
# of Rack Total	7/9	4/5	3/4
Power Total (kW)	137/178	78/98	59/78

Table 2.10: The size estimation of the CPU

	Run 4	Run 5
CPU needed [MHS06]	0.7-0.9	1-1.4
CPU needed (LRT 5kHz/20 kHz) [MH06]	0.025/0.10	0.048/0.19
# of dual-socked servers	484/534-617/667	583/662-805/884
# of Racks	11/12-14/15	13/15-18/20
Power [MW]	0.18/0.20-0.23/0.25	0.21/0.24-0.29/0.32

Review Process

- A review committee has been charged to review and make a comparative analysis of the three options
  - Chair: G. Broojimans
  - 4 members from TDAQ:
    - N. Ellis, K. Nagano, I. Riu, D. Sankey
  - 5 full members from outside TDAQ to cover various areas of expertise:
    - P. Calafiura, H. Chen, S. McMahon, N. Pettersson, A. Polini
  - Ex-officio:
    - TDAQ UPL, PL, EF manager, TDAQ UPR Resource Coordinator, TDAQ UPR TC, TDIB chair, PMO Leader and Coordinators of the different studies
- The committee has been an integral part of the process
- Operated since December alongside the study groups to ensure coherence
- Presented the conclusions by early June to the extended TDAQ Steering Group (eTDSG)

Committee Findings and Comments

Performance were evaluated cf. efficiencies, resolutions and fake rates

Efficiency

$\|\eta\|$	Muons, p_T = 10 GeV			Pions, p_T = 10 GeV			Electrons, p_T = 10 GeV		
	0.1–0.3	0.7–0.9	2.0–2.2	0.1–0.3	0.7–0.9	2.0–2.2	0.1–0.3	0.7–0.9	2.0–2.2
Offline	> 99.9%	> 99.9%	> 99.9%	> 96.1%	> 94.2%	> 90.5%	> 98.2%	> 95.5%	> 94.1%
Software	99.6%	99.7%	99.7%	95.0%	92.8%	90.0%	96.2%	95.8%	94.0%
Custom (1st stage)	98.8%	97.9%	98.7%	92%	88%	80%	90%	85%	80%
Custom (2nd stage)	98.1%	-	-	89%	-	-	89%	-	-
Heterogeneous (best)	99.6%	99.6%	99.0%	96%	94%	80%	92%	85%	80%
Heterogeneous (worst)	98.5%	98.5%	97.6%	94%	92%	-	88%	86%	-

Impact parameter resolution: Heterogeneous option argues it should be close to fast software as final fit done using same SW - true *if* the same hits are used

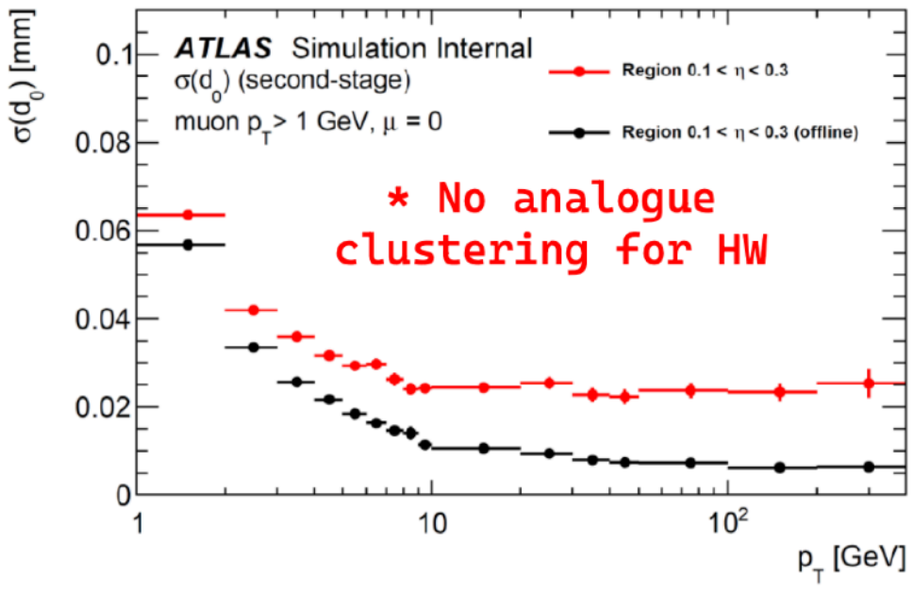

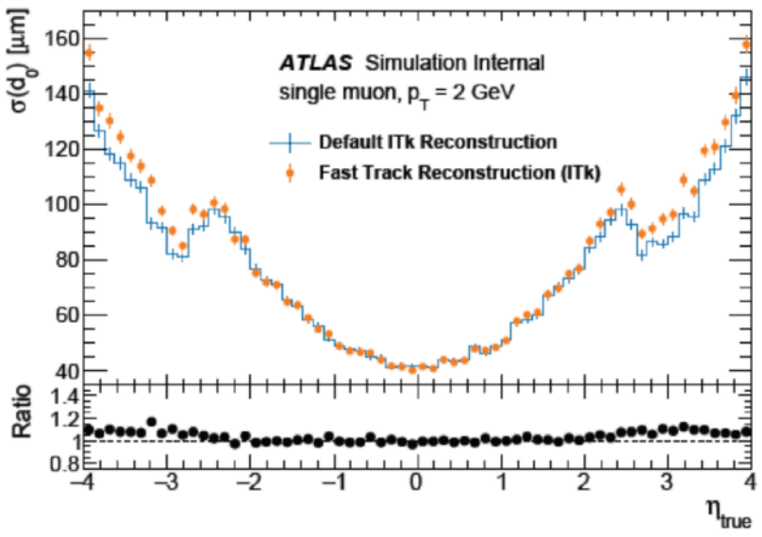

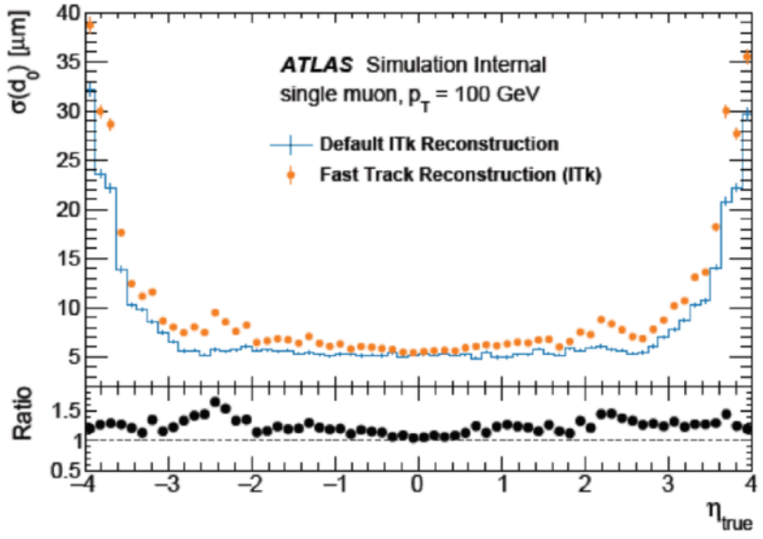

Committee Findings and Comments

Fake rates:

- Number of tracks (fakes or duplicates, with different consequences downstream)
- Custom TF report shows ~3% loss in efficiency for muons aber second stage tracking
  - Recoverable by optimization?
- Heterogeneous commodity TF focused on implementing a well-known algorithm (Hough transform) given short time available
  - Hough transform is known to be particularly susceptible to fakes
  - Fake rejection NN for heterogeneous leads to 2-4% loss in efficiency !
- Software track rate very similar to offline, which has ~10^{-5} fake rate

Custom HW

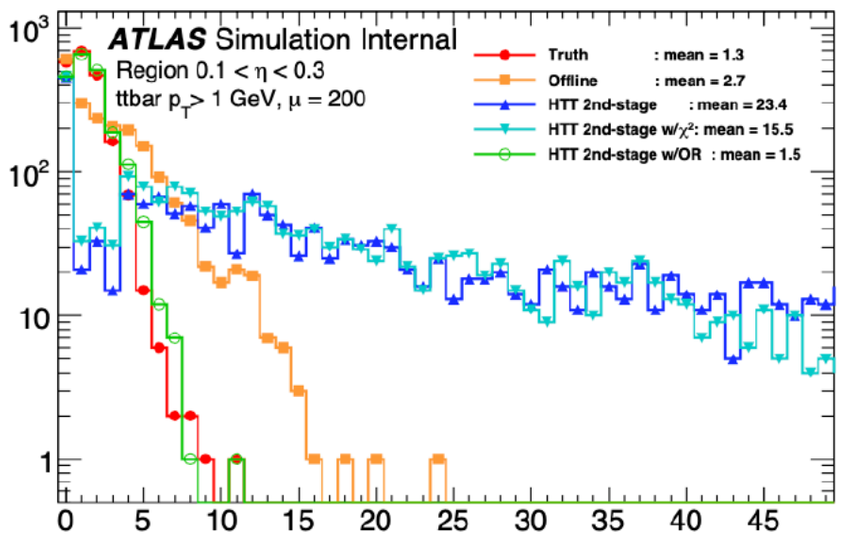

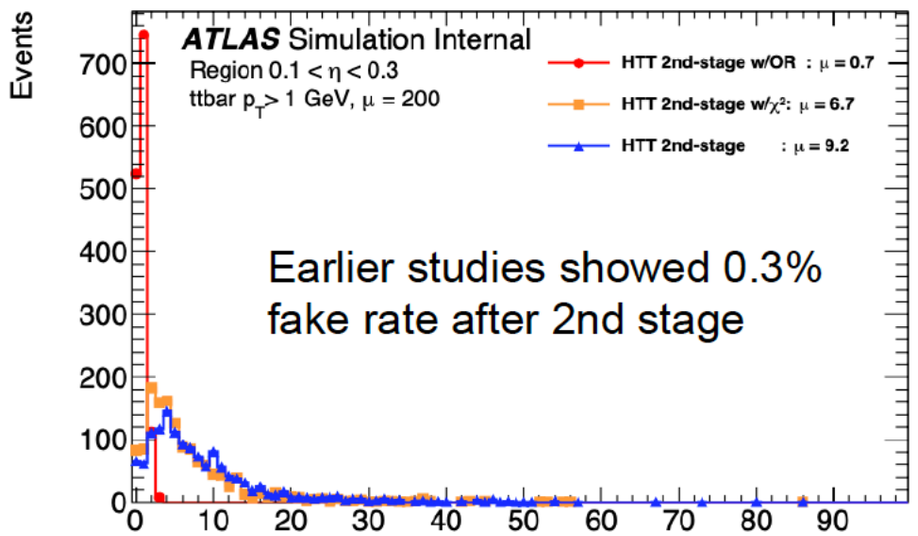

Heterogeneous

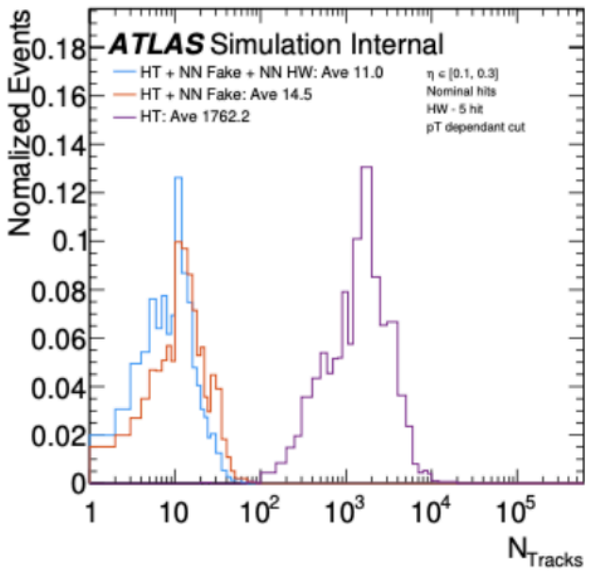

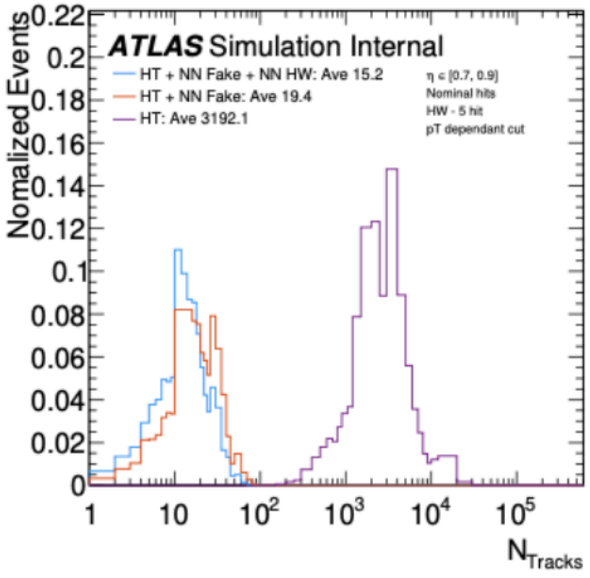

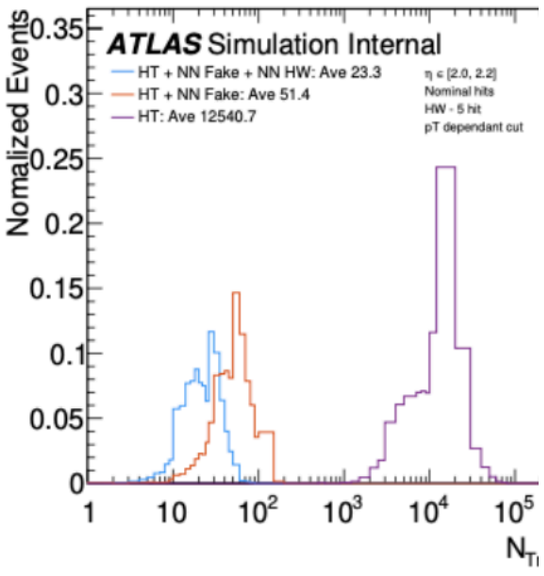

Committee Findings and Comments

On Performance

- **SW-only** approach clearly yields best performance, meets all requirements
- **Hardware-based** options may get to similar efficiencies
  - But fake/duplicate rates are high, and reduction comes at a cost in efficiency
  - Impact of worse resolution downstream non-trivial
    - Eg for vertexing, b-tagging, …
  - d0 resolution certainly worse for custom HW - crucial for b-tagging (large fraction of regional tracking trigger menu)
  - In the current reports, hardware-based options do not meet all requirements
- **Heterogeneous Commodity** TF argues resolutions will be similar to SW-only fast tracking, but this will only be true if correct hits are fed to software tracking
  - They do budget for fast tracking Kalman filter in CPU
- Event Filter is not Level-1: need to provide precision tracking with best possible efficiency and resolution, in all areas
  - Direct impact on all physics objects!

Committee Findings and Comments

On Feasibility

- Custom
  - Challenging data flow, fake/duplicate rates are high and fake rejection via χ2 drives system size and throughput
  - Regional resource assignments limit flexibility or very hard constraint for Associative Memory-based pattern recognition, but FPGA-based still needs it for χ^2 fit
- Heterogeneous commodity
  - Simpler data flow: each FPGA sees a full event
  - High fake/duplicate rate from Hough transform, but not enough time to study other (eg machine learning-based) algorithms
    - Recent papers show good promise for eg Graph Neural Networks
- Software
  - Speed-optimized offline reconstruction, proven

Committee Findings and Comments

On Maturity

- Custom
  - Design has concentrated on pattern recognition
  - Dataflow synchronization across the shelf is likely just as hard
  - Fake/duplicate reduction while maintaining efficiency key
  - Unfortunately not much design time left
- Heterogeneous commodity
  - Impressed by all the work done, maturity after 5 months
    - Proof of concept largely achieved, but lots of work to do to demonstrate solution will meet specifications
- Software
  - Fully functional version meeting the specs exists
  - Further improvements likely, including potential use of accelerators (incl. GPUs) following offline R&D work

Committee Findings and Comments

On Risks

- Custom
  - ✦ Custom solutions are intrinsically riskier than commercial ones
  - ✦ AM ASIC is complex, AM09 is a significant step up from AM08 in size
    - ❖ Sizable risk that more cycles than planned necessary, with cost & schedule implications
  - ✦ Demanding dataflow issues
  - ✦ Any system sizing issues require remapping the system, complex

- Heterogeneous commodity
  - ✦ Biggest risk is if resources to cover full η-ϕ range don't fit in single FPGA
    - ❖ Would lead to huge increase in complexity
  - ✦ Fake rate reduction to be figured out Proof of concept largely achieved, but lots of work to do to demonstrate solution will meet specifications

- Software
  - ✦ Cost risk if CPU more expensive, or need lower power CPU
    - ❖ But assumptions made are conservative - CPUs at needed power level already exist
  - ✦ We often struggle to find enough people to work on low-level software

Committee Findings and Comments

On operations

- Compared for different phases
  - ✦ e.g.: maintenance requirement for HW failures or changes to detector/ operations conditions

	Custom-AM	Custom-FPGA	Heterogeneous	Software
Installation				
Commissioning				
Operations				
Maintenance & Upgrades				
Power & Cooling				
Simulation				

Key	Hardest			Easiest

Committee Findings and Comments

On costs

- Taking credit for existing Run-3 farm (value ~2.8 MCHF), and using same rates for regional tracking as in the TDR

HW 5%: 5% of 8*2π, includes margins
SW 5%: 5% of detector elements, as in TDR

	Custom-AM	Custom-FPGA	Heterogeneous	Software
CORE Cost - HW 5% def	4.55 MCHF	3.25 MCHF	2.8 MCHF	5.44 MCHF
CORE Cost - SW 5% def	4.55 MCHF	2.71 MCHF	2.0 MCHF	3.28 MCHF

- Uncertainties: ~30%
  - ✦ For reference in the TDR: HTT (17.4 MCHF / EF farm: 3.4 MCHF)
- Including additional scope, i.e. +20 kHz for Large Radius Tracking (LRT) and +50 kHz of full scan

	Custom-AM	Custom-FPGA	Heterogeneous	Software
LRT: 20 kHz	0.54 MCHF	0.59? MCHF	0.8 MCHF	0.14 MCHF
Full scan + 50 kHz	0.17 MCHF	0.45 MCHF	1.17 MCHF	1.5 MCHF

On required effort

- Effort estimates, on top of existing EF tracking effort, are substantially different

	Effort [FTE-years]	Comments
SW-only	14	Based on past experience with software projects, this number seems to be underestimated. Certainly does not leave room for additional
Heterogeneous	80	~evenly split between engineers and scientific personnel. Comparison with other projects suggests this is somewhat underestimated
Custom HW	130-140	~evenly split between engineers and scientkfic personnel

Committee Findings and Comments

On Power and Cooling

	Custom-AM	Custom-FPGA	Heterogeneous	Software
Tracking only	0.14 MW	0.15 MW	0.28 MW	0.97 MW
Total EF, incl LRT & 150 kHz full scan	0.59 MW	0.6 MW	0.91 MW	1.82 MW

- Assumes 0.2 W/HS06 which is somewhat conservative
- More complex simulation for physics for AM-based, and possibly FPGA based options could be CPU-intensive; AM ASIC pattern regeneration when conditions change may require non-negligible CPU resources if frequency higher than planned

Committee Assessment

- On all fronts but power and cooling the SW option is the strongest
  - And there is ample power and cooling margin for a SW system that meets the specs
  - *If we had to make a final choice today, there is little doubt we'd (have to) go with the SW option: lowest risk, best physics performance*
    - *EF is not L1, need the best possible efficiency and resolution: a 10% efficiency loss in EF is like running at 900 kHz L0 accept instead of 1 MHz, or waste of 10% LHC running time*
- The downside of the SW option is power consumption
  - A big reason FPGA-on-PCIe cards are commercialized is reducing data center power usage
- Heterogeneous commodity approach offers significant reduction in power
  - More time is needed to study more complex algorithms
  - R&D should be continued
- No real advantage to the custom solution
  - But significant risk and a large investment would need to be made soon
- Recommends that
  - ATLAS commits to a commercial solution for EF tracking at HL-LHC
  - TDAQ should continue investigating using hardware accelerators to optimize the EF farm
    - Heterogeneous commodity TF has largely demonstrated proof-of-concept
    - A heterogeneous solution (incl. FPGAs and/or GPUs) could lead to substantial power and cost savings

eTDSG Decision

- Considering the reports [about] alternative scenarios, the report of the Review Committee, and folding in additional project considerations, the eTDSG decided [...] that the new baseline [for EF Tracking] will be based on commercial hardware.

- This change of baseline technology has been precipitated by revisions to the underlying assumptions since the time of the TDR [...]. The custom-based solution now has no clear competitive advantage compared to a commercial solution. Conversely, it carries a significantly higher risk than commercial options, which is inherent to all custom developments and systems.

- An ambitious program will be undertaken to further develop and optimise efficient algorithms for commodity platforms (CPUs and accelerators) and to follow and evaluate commodity computing technologies.
  - A variety of high performance accelerator technologies, system architectures, and implementation languages shall be investigated. The results of these studies will contribute to the final EF tracking technology choice

Current Status

- TDIB almost unanimously approved the eTDSG decision (3 abstentions) on July 9th
- TDAQ UPL and EF Level-2/deputy appointed two L3 Coordinators for the new EF Tracking activity: V. Boisvert, S. Majewski
- Proceeding towards TDR amendment to be endorsed by the TDAQ IB on 24.09.2021
- An editorial team has been active since the endorsement
  - Liza Brost (chair), Jahred Adelman, Markus Elsing, Elliot Lipeles, and Frank Winklmeier
  - Ex-officio: UPL, EF L2, EF Tracking L3, TDAQ Resource and Technical Coordinators

Contents

Next Steps

Task-force documents and Review Committee Report — *end May/ beg. June*

↓

extended **TDAQ SG decision** — *June 18th*

↓

TDAQ IB (TDIB) endorsement — *July 9th*

↓

TDR Amendment preparation — *Jul through Sep*

↓

Internal Circulation and TDIB endorsement — *TDAQ week Sep 20-24*

ATLAS Circulation — ***Sep 24- Oct 7***

↓

EB approval submission — ***Oct 8***

↓ ***Oct 15***

report to CB

↓ ***Oct - Mar***

LHCC/UCG review process (in parallel)

↓ ***Nov 8-10***

P2UG in-depth review (for the remaining part of TDAQ + info)

↓ ***Feb 21***

CB approval final TDR amendment

↓ ***Mar 16***

CERN RB approval

↓

MoU update to DRC — ***~mid to end Mar***

↓

RRB report — ***Apr 25***

↓

P2UG summary review — ***early May***

- Kickoff meeting — ***~Oct 22***
- Report to LHCC week — ***Nov 15-17***
- Iteration meeting — ***~Dec***
- (2nd Iteration meeting - if needed) — ***~Jan***
- Review — ***~mid Feb***
- Recommendations LHCC week — ***Mar 7-10***