



Perspectives & Outlook on Hardware Tracking


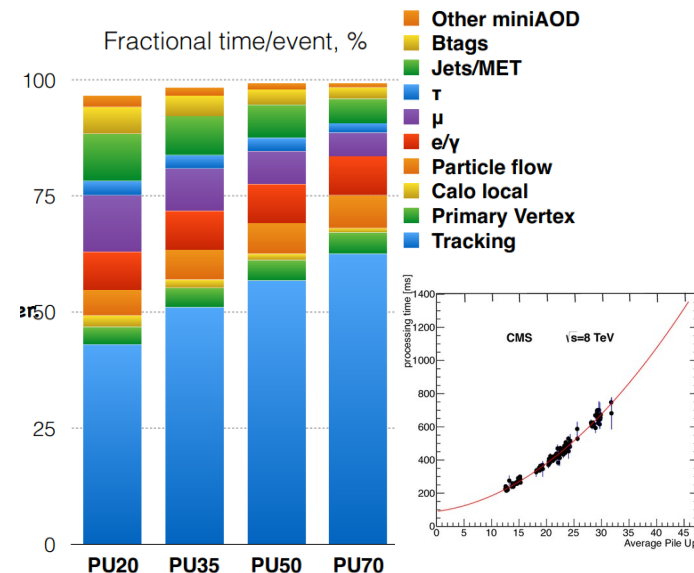
Kristian Hahn
Northwestern University

Snowmass 2022, Seattle
Instrumentation Frontier – TDAQ
7.19.22

Tracking

The most computationally intensive aspect of reconstruction at collider experiments


- Entails :
 - **Track Finding** – pattern recognition, the association of detector hits to a single trajectory
 - **Track Fitting** – obtaining kinematic variables from a trajectory fit
- Typically iterative, dominates overall reconstruction time
- Complicated by overlapping collisions (pileup)
 - Complexity / processing time for pattern recognition scales as a power of #vertices
 - Pileup grows as instantaneous luminosity increases




Online Tracking

Situation is even more challenging for *online* reconstruction

- Tracks are needed by trigger selection algorithms
- The online environment imposes latency constraints on reconstruction processes, eg:
 - 4 us total latency in the Run 2 CMS L1 hardware trigger
 - And an average of ~260 ms / event in the Run 2 software-based HLT
- Online algorithms generally operate with reduced inputs and/or precision
 - Either because of latency concerns (eg: truncation)
 - Or due to data availability (eg: lack of pixel readout)



HLT tracking




Hardware Tracking

Hardware tracking as a way to improve the physics performance of online reconstruction and selection

- In hardware (L1) trigger systems
- As custom co-processors for the software trigger
- On heterogeneous commodity accelerators in software trigger systems

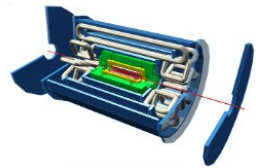
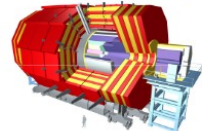
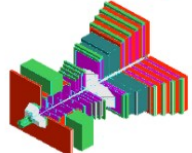
Will review recent ATLAS/CMS/LHCb history

- Much development (and disruption!) in online tracking / trigger over the past several years
- From this, will try to draw some conclusions for hardware tracking for future experiments & upgrades



Run-1 Trigger Systems

ATLAS/CMS/LHCb started Run-1 with similar, “conventional” triggers

	No.Levels Trigger	Level-0,1,2 Rate (Hz)	Event Size (Byte)	Readout Bandw.(GB/s)	HLT Out MB/s (Event/s)
	3	LV-1 10^5 2.5 μ s LV-2 3×10^3	1.5×10^6	4.5	300 (2×10^2)
	2	LV-1 10^5 4 μ s	10^6	100	100 (10^2)
	2	LV-0 10^6	3×10^4	30	40 (2×10^2)

- First level h/w trigger : calorimeter & muon only, tracking (incl pixel) only in the HLT
- Some differences in the details
 - L1 ROIs in ATLAS, tracking within ROI in L2 (CPU), full tracking at HLT
 - LHCb : larger rate out of L1, smaller event size

L2/HLT physically merged in Run-2




LHC / HL-LHC Plan



LHCb Run-3

Hardware trigger (L0) removed!

- 30 MHz rate into the HLT
- Full tracking performed on GPU
 - Co-located with event building
 - Simple track selections at this stage (HLT1)
 - Full offline-quality reconstruction in CPU (HLT2)

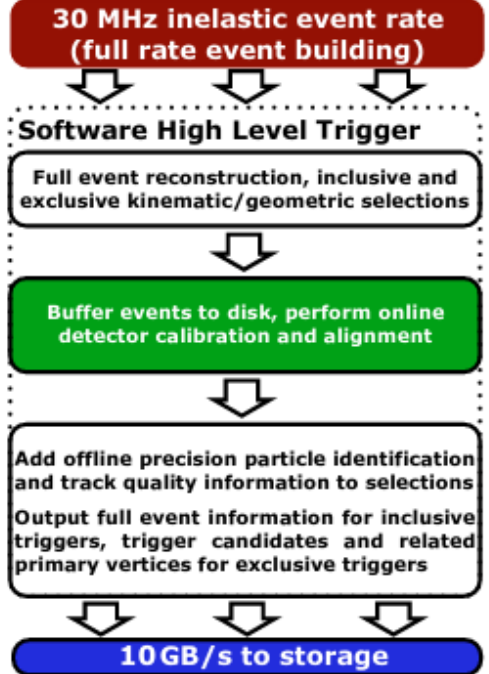


Object filtering applied for ~2/3 of the output data stream

- Significantly reducing output rate to tape
- Re-reconstruction not possible for these events

We describe a fully GPU-based implementation of the first level trigger for the upgrade of the LHCb detector, due to start data taking in 2021. We demonstrate that our implementation, named Allen, can process the 40 Tbit/s data rate of the upgraded LHCb detector and perform a wide variety of pattern recognition tasks. These include finding the trajectories of charged particles, finding proton–proton collision points, identifying particles as hadrons or muons, and finding the displaced decay vertices of long-lived particles. We further demonstrate that Allen can be implemented in around 500 scientific or consumer GPU cards, that it is not I/O bound, and can be operated at the full LHC collision rate of 30 MHz. Allen is the first complete high-throughput GPU trigger proposed for a HEP experiment.

LHCb Upgrade Trigger Diagram



“Allen: A High-Level Trigger on GPUs for LHCb”,
[Computing and Software for Big Science \(2020\) 4:7](#)

<https://lhcb.github.io/starterkit-lessons/first-analysis-steps/dataflow-run3.html>

ATLAS FTK

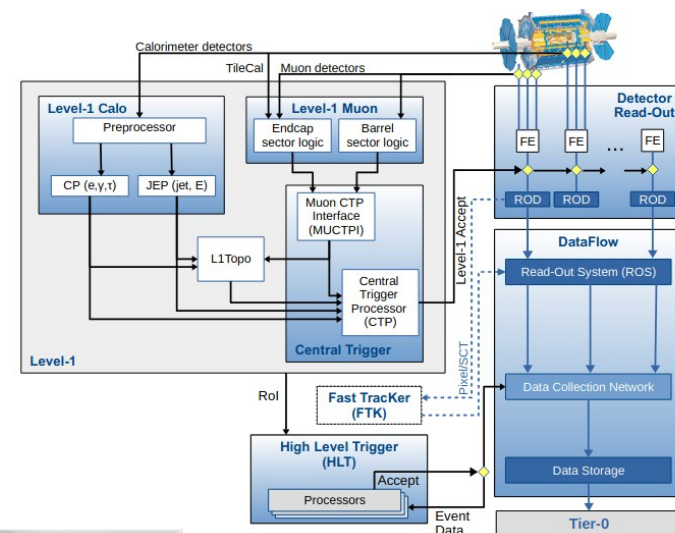
Phase-1 upgrade : custom hardware for bringing tracks into the HLT

Designed to :

- Operate with full tracker coverage, 8 SCT + 4 Pixel layers
- Perform global tracking to on every L1 accept (100 kHz)
- Provide tracks > 1 GeV to HLT within a 100 us latency

Pattern recognition based on associative memory

- Custom ASIC : AM06, 65nm
- Pattern matching of SCT hits with 128k pattern banks
 - 8k AM06 chips in the system, 1B patterns
- On successful match, linearized track fit using full resolution hits from 12 layers



ATLAS FTK TDR - CERN-LHCC-2013-007

"AM06: the Associative Memory chip for the Fast Tracker in the upgraded ATLAS detector", A. Annovi et al 2017 JINST 12 C04013

"The ATLAS Fast Tracker System", JINST 16 (2021) P07006


ATLAS FTK/HTT

A sizeable backend system was foreseen

- Successful slice testing performed in 2018
- **Project canceled in 2019**

this decision were the lower than expected pile-up due to cryogenic limits of the LHC [13], significant gains from optimization of the HLT software-tracking algorithm, and potential resource shortages. Instead, it is envisioned to perform more tracking in the HLT for signatures that benefit the most, e.g. jets and E_T^{miss} .

Contributors to



HL-LHC evolution : HTT

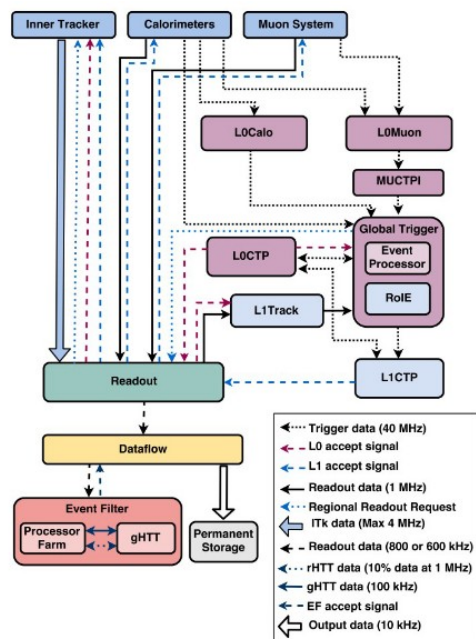
- AM09 at 28 nm, 3-4 x 128k patterns, streamlined system design
- Regional (1 MHz), global (100 kHz) tracking co-processor for EF
 - Optional *regional L1 (ie: L2) tracking*, 30 us latency, 2-4 MHz input rate
- **Project canceled in 2021**

The decision not to pursue the “evolution” scenario was based primarily on the risks in the development of the ITk Pixel FE ASIC, the reduced physics motivations due to the stronger than anticipated limitations of throughput for the ITk Pixel Detector readout, risks related to L1 latency and the technical challenges in commissioning an L0/L1 TDAQ system alongside the L0 TDAQ system.

An independent ATLAS committee reviewed the reports from these three task forces and recommended that “ATLAS commit to a commercial solution for EF tracking at HL-LHC,” including further optimizations targeting the use of accelerators to potentially mitigate risks related to power and cost: “TDQA should continue investigating using hardware accelerators to optimize the EF farm.” The Heterogeneous commodity task force has largely

Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System, CERN-LHCC-2017-020

Technical Design Report for the Phase-II Upgrade of the ATLAS Trigger and Data Acquisition System - EF Tracking Amendment, CERN-LHCC-2022-004



CMS Phase-2

Run-2,3 : tracking acceleration in HLT w/ GPUs

HL-LHC : hardware tracking in the L1 trigger

- Double-sided p_T modules for Phase-2 Outer Tracker

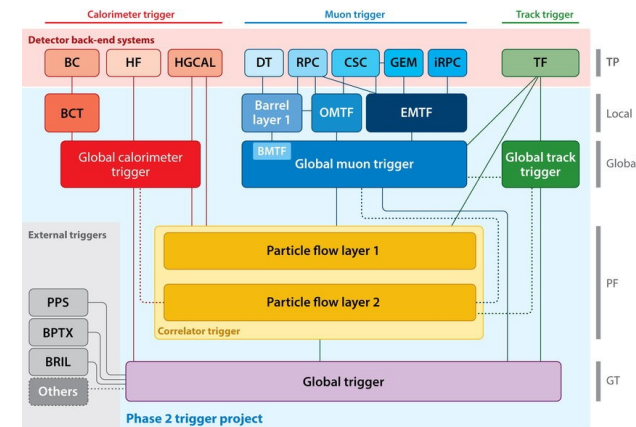


- Readout ASICs (CBC,SSA) identify correlated hits (“stubs”) consistent with >2 GeV tracks
- Stubs (40 MHz) multiplexed with DAQ hits (L1A rate, 750 kHz) and streamed to backend electronics

- 80% of OT's data stream consists of trigger info!

- Stubs sent to L1 tracking system on 25 Gbps links

- Aggregate stub bandwidth ~ 50 Tbps
- Full tracking with the outer detector (>2 GeV) in $4+1$ μ s
- Fit tracks sent to L1 trigger, used in particle flow algorithms and forwarded to HLT

“Performance of Phase-2 HLT Reconstruction and GPU offloading benchmarks”,
CMS-DP-2021-013



Pyd A. Skinnari L. 2020
Annu. Rev. Nucl. Part. Sci. 70:171-95

CMS L1 Tracking

Gains from L1 Tracking in CMS are well established

- Improving resolution, sharpening efficiency turn-ons, lowering thresholds. LLP potential

“The Phase-2 Upgrade of the CMS Level-1 Trigger”,
[CERN-LHCC-2020-004](#)

Implementation





“CMS Hardware Track Trigger: New Opportunities for Long-Lived Particle Searches at the HL-LHC”,
Y. Gershtein, [arxiv:1705.04321](#)

- 3 demo'ed proposals : 1 AM-based, 2 fully-FPGA
 - AM options : 28 nm planar, 3D stacked with TSV
 - 2017 decision to not pursue the AM option**

While the AM-based method requires a challenging chip development in a novel (for HEP) technology, carrying many unknowns, a method based only on commercial FPGAs can be considered as low-risk, provided that the required developments are properly planned and managed.

- Current baseline: a “hybrid” of the two FPGA algorithms
 - “Tracklet” road search track finding + Kalman Filter fit
 - Run on ~200 ATCA blades, each 2x Xilinx VU9/13P

The Apollo ATCA Design for the CMS Track Finder and the Pixel Readout at the HL-LHC, A. Albert et al,
[arXiv:2112.01566](#)



Thoughts on the Future



FPGAs vs ASICs

Modern FPGAs have become incredibly powerful ...


High-speed SERDES links are now commonplace

- Significant processing power can be brought to bear on a lot of data ... very quickly


Diminishing the need for custom ASICs in *backend* TDAQ

- A stark difference with respect to not too long ago ...
- No backend ASICs foreseen in HL-LHC ATLAS/CMS

	VU19P
System Logic Cells (K)	8,938
GTY/GTM Transceivers (Number of 32.75Gb/s and 58.0Gb/s)	80 / 0
PCIe Interfaces	8 Gen4 x 8/CCIX
In-package Memory	224Mb
DSP Slices	3,840
I/O	2,072



RCT Receiver card

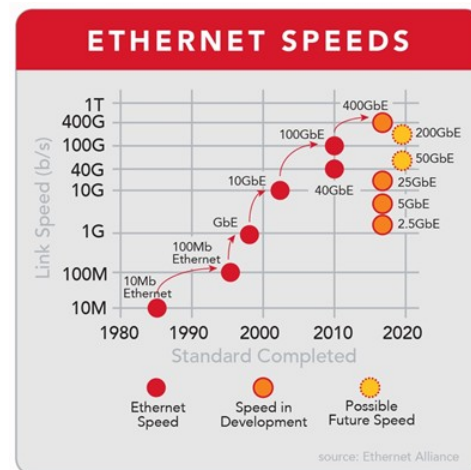


RCT Jet/Summary card



RCT Electron isolation card

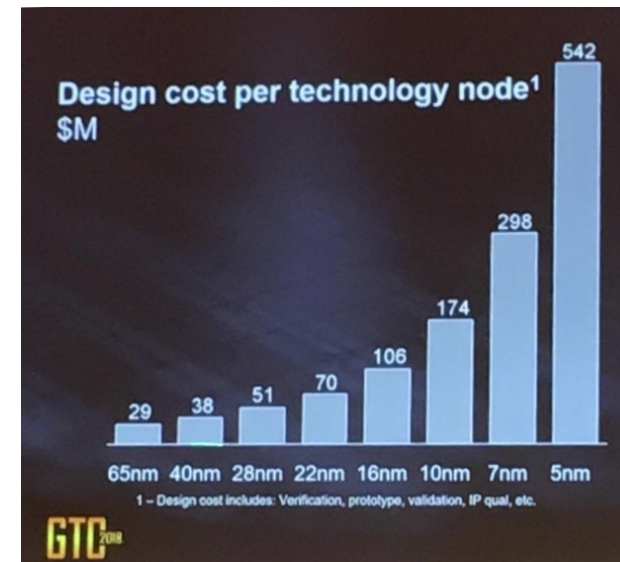
Some of the Run-1 CMS L1 electronics



More on ASICs

HEP remains on lagging technology nodes

- Eg: 130 nm, 65 nm for HL-LHC
- Designs become more complex & costly at higher nodes
 - In particular, **significant effort needed for design verification**
- Many design teams already under strain
 - As reflected by establishment of CERN **CHIPS** service
 - And by collaborative design efforts (eg: RD53)



We're a smaller player in a rough market

- **Considerable schedule/budget risks associated with ASICs in the current climate**
- Can't completely avoid (eg: for frontends), but there's a strong desire to minimize these risks ...




FPGAs

FPGAs are not without risk ...

- Unclear if L1 triggering/tracking will continue to be as well supported as FPGAs shift to co-processors
- **Clear movement toward the server : heterogeneous computing for AI / ML / "analytics"**
 - Will the market constrain our flexibility to design custom systems?
 - Will it still be advantageous/desirable to do so?
 - Will we eventually need to work fully in popular AI paradigms?
 - How well will single-minded algorithms perform on h/w optimized for AI?
 - Vendor, platform, and toolkit lock-in?



ALTERA
now part of Intel



At the heart of the Xilinx Alveo U200 and U250 accelerator cards are specially screened FPGAs that run optimally (and exclusively) on Alveo. The Alveo U200 card features the XCU200 FPGA and the Alveo U250 card features the XCU250 FPGA. Both XCU200 and XCU250 FPGAs use Xilinx stacked silicon interconnect (SSI)



Back to ASICs

Integration of trigger capabilities in CMS frontend OT ASICs crucial for L1 Tracking

- Factor of 10 rate reduction to L1Tk from FE stub finding
- Tight integration facilitates detector optimization simultaneously for trigger & DAQ performance

Further integration of real time tracking functionalities?

- So far tracking ASICs have dealt with Outer Trackers ... extend to pixels
- Include timing → 4D tracking
 - Will be required at FCC to cope with extreme pileup
- Inter-module correlations? Hardware triplets/Tracklets?
 - Via wireless? 6G targeting O(10 us) latencies
- In general : push more logic into the frontend ASICS (ie: “Intelligent Trackers”)

Must consider on-detector power & mass needed to achieve these




Measurement	Technical requirement
Tracking for e^+e^-	Granularity: $25 \times 50 \mu\text{m}^2$ pixels
	$5 \mu\text{m}$ single hit resolution
	Per track resolution of 10 ps
Tracking for 100 TeV pp	Generally the same as e^+e^-
	Radiation toleran up to $8 \times 10^{17} \text{ n/cm}^2$
	Per track resolution of 5 ps

FPGAs (again)


Toolkits helping to promote physicist engagement

- CMS Hybrid algorithm utilizes Xilinx High Level Synthesis (HLS)
 - Write C++, tool converts to RTL
 - HLS4ML used throughout the CMS L1 trigger
- SYCL / oneAPI allows a single codebase to run on heterogeneous accelerators (GPU/FPGA/AI)
- Helps to broaden involvement beyond EE/CS experts



System on Chip?


- FPGA logic + CPU embedded system (eg: Zynq, Versal)
- Planned use in HL-LHC for blade control & interfacing
- Affordable, multi-core (including RT), low power
- New devices also with on-die “AI engines”
- Can online tracking algorithms benefit? OS latencies ...
 - See also : ATLAS [gFEX](#), uses Zynq PL



Finally : Future Experiments

FCC(-hh)

- Some specs :
 - 30 GHz pp collisions, 4 THz track rate
 - 1000 <pileup>, 125 μm <vertex separation>
 - Compare : 200 pileup & 1mm for HL-LHC
- Poses similar design choices as the (HL)-LHC
 - L1 + HLT \rightarrow high-performance on-detector processing
 - Multi-level trigger \rightarrow significant on-detector buffering
 - Triggerless \rightarrow huge bandwidth




Will HL-LHC experience / tech road maps simplify the choice?

Must consider on-detector power & mass needed to achieve these

A world exists beyond colliders ...

- Hardware tracking needed in smaller HEP experiments
 - Eg: The MuonE proposal

Status of the MuonE Project,
G. Abbiendi



A photograph of a server rack with a semi-transparent dark blue overlay in the center containing the text 'Thank You!'. The server rack is filled with various components, including network switches and cables. The cables are primarily blue, with some red and green ones visible. The server rack is mounted in a blue metal frame. The text 'Thank You!' is written in a white, sans-serif font and is centered within the dark blue overlay. The background shows the internal components of the server rack, including network ports and cables. The overall image has a professional and technical feel.

Thank You!