GPUs in Triggers and Reconstruction

Tom Boettcher

on behalf of the ALICE, ATLAS, CMS, and LHCb Collaborations

LHCP
June 9, 2021

Computing Challenges at the LHC

Run 3 (2022): "Triggerless" data collection at ALICE and LHCb

- LHCb: *pp* at 40 MHz, 5 TB/s
- ALICE: PbPb at 50 kHz, 3.5 TB/s

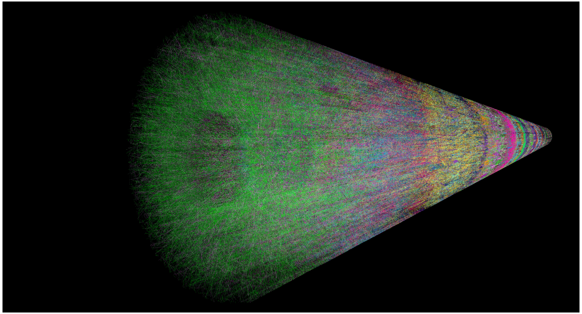

Run 4 (2027): The HL-LHC

- ATLAS and CMS will collect *pp* data at a pileup of 200
- Will increase hardware trigger rate from 100 kHz to 750 – 1000 kHz

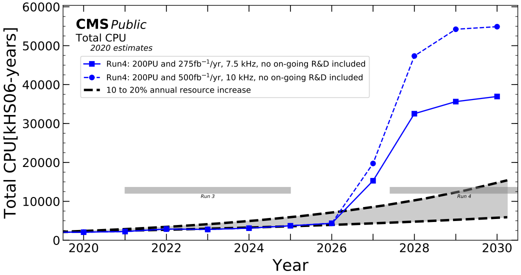

Why use GPUs

Affordable computing power in a compact package

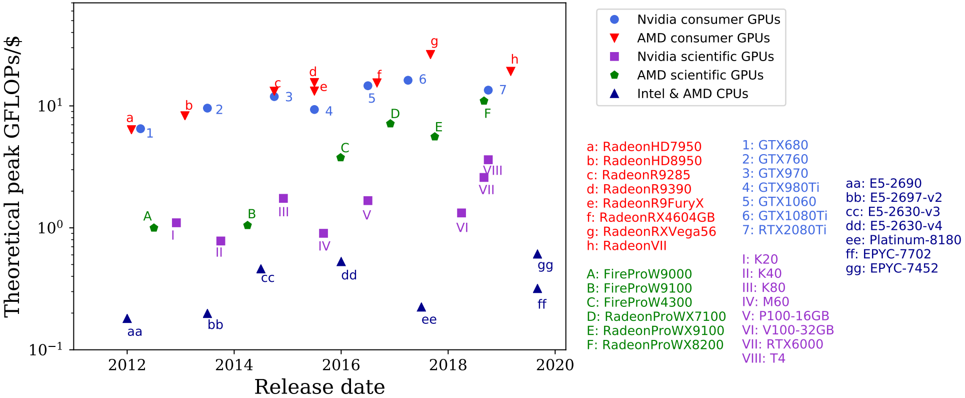

Courtesy of Dorothea Vom Bruch, arXiv:2003.11491

Also see talk by Ole Schmidt on June 8 for more on hardware acceleration in general

What is a GPU?

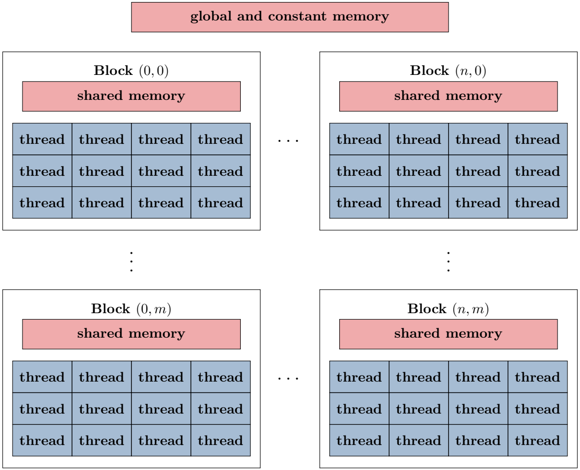

- Highly parallel processors with thousands of cores (threads)
- Many trigger and reconstruction tasks parallelize nicely (with some work)

- Triggerless PbPb data taking at 50 kHz
- TPC must be reconstructed in real time for compression and calibration
- TPC reco on GPUs since Run 1
- Aim for full barrel track reco on GPUs

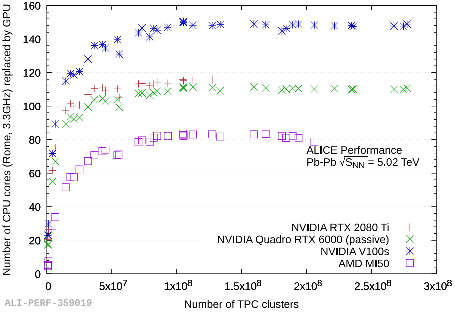

PoS (LHCP2020) 053

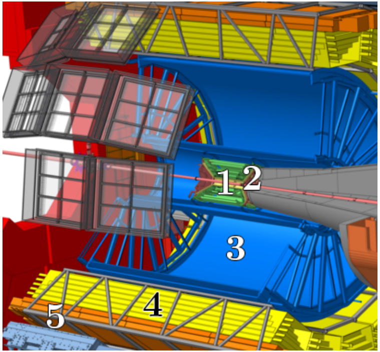

1. **ITS**
2. V0 and T0
3. **TPC**
4. **TRD**
5. **TOF**

GPU reco ready

GPU reco in development

ALICE Run 3 GPU Scenarios

Baseline: TPC reconstruction + whatever else is ready
Optimistic: Full barrel track reconstruction on GPUs

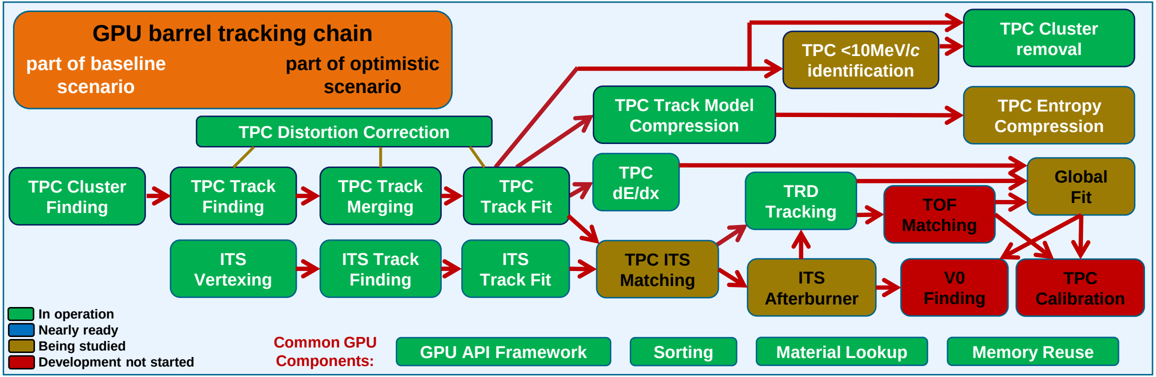

See David Rohr's vCHEP talk

GPUs at ALICE vs. ATLAS and CMS

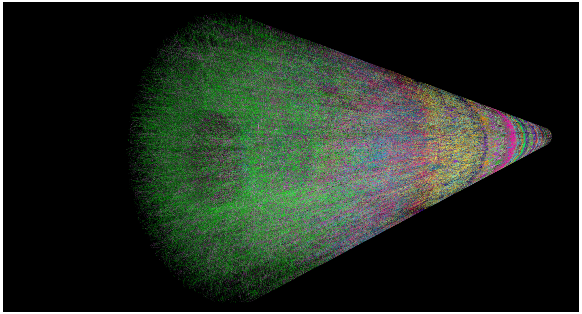

ALICE

- TPC reconstructions dominates
- Process 10 ms “timeframes”, $\mathcal{O}$(10 GB)
- $\sim 80\times$ speedup on GPUs
- Clear use case for GPUs

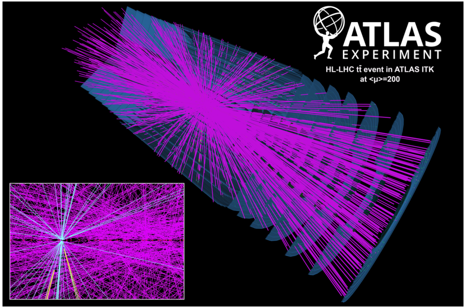

ATLAS and CMS

- No single subdetector dominates
- Much higher rates of smaller events
- Advantages of GPUs compete with challenges: portability, overhead, etc.

ATLAS: GPU R&D in ACTS

A Common Tracking Software (ACTS)

- Targeting end-to-end tracking on GPUs: traccc
- Includes R&D studies of data structures and memory management: VecMem

Track Seeding Algorithm

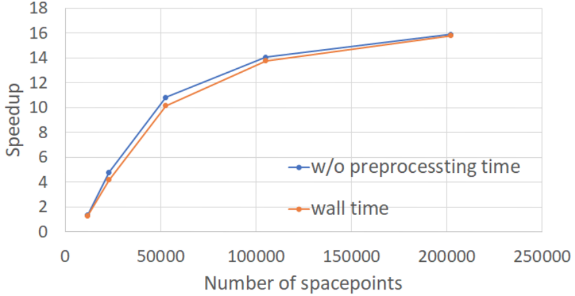

arXiv:2007.01239

Kalman Filter

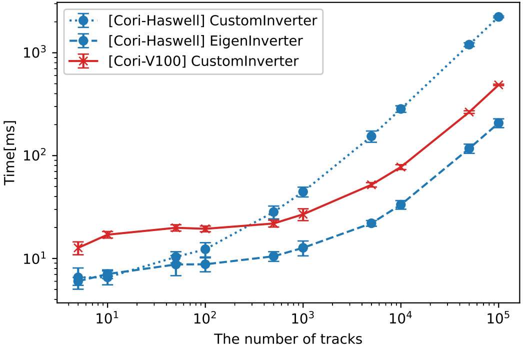

arXiv:2105.01796

CMS: Pixel Track and Calorimeter Reconstruction

CMS reconstruction on GPUs in Run 3

- Pixel track reconstruction: Front. Big Data 3 (2020) 601728
- Primary vertices from pixel tracks
- Calorimeter local reconstruction
- Out-of-time pileup subtraction
- Fully integrated into CMSSW. See arXiv:2004.04334

Other GPU R&D activities

- Particle Flow for Run 3
- HGCAL reco for Phase 2. See Front. in Big Data 3 (2020) 591315 and Bruno Alves' vCHEP talk

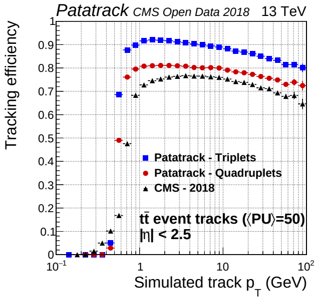

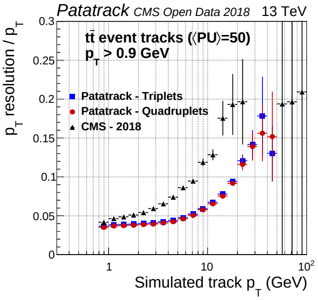

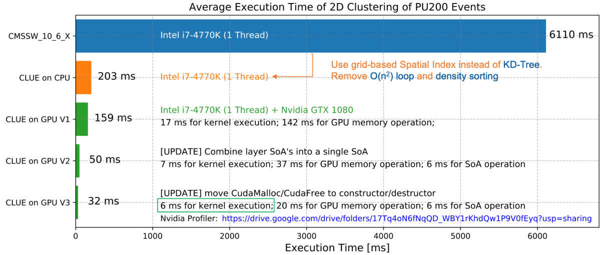

CMS: Run 3 Performance

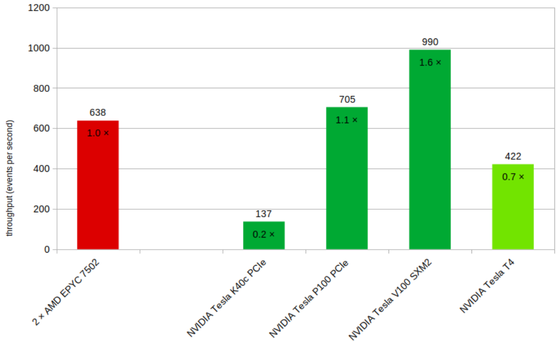

CPU only

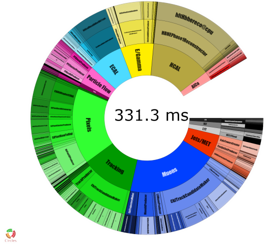

CPU w/ GPU offload

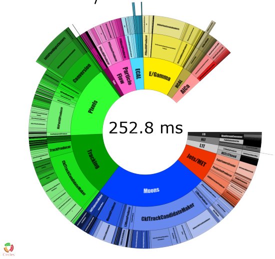

- CPU w/ GPU offload yields $\sim 25\%$ increase in HLT throughput
- Tesla T4 does about 70% of the work of $2\times$AMD Rome CPUs
- Costs $< 25\%$ as much
- Uses $< 20\%$ as much power

GPUs at ATLAS and CMS: Performance Portability

Major obstacle of GPU computing in HEP

- Different languages for and algorithm design for GPUs and CPUs
- Different languages between GPU vendors (e.g. CUDA, HIP)
- New accelerators may appear

Want native performance on many architectures with one codebase → implement GPU algorithms with "performance portability" APIs

- Kokkos
- alpaka
- oneAPI/SYCL

NB: ALICE and LHCb use their own lightweight wrappers for performance portability

GPUs at ALICE, ATLAS, and CMS vs. LHCb

ATLAS and CMS

- GPUs as coprocessors in software triggers
- Perform reconstruction for a few individual detectors

ALICE

- No software trigger
- Use GPUs for reconstruction and data compression

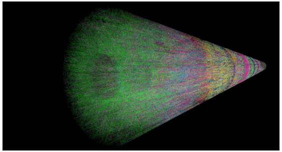

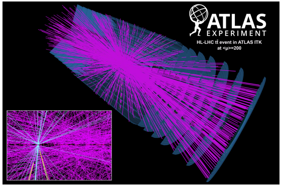

LHCb

- Full first level software trigger (HLT1) on GPUs
- 30 MHz of small events → unique I/O challenge

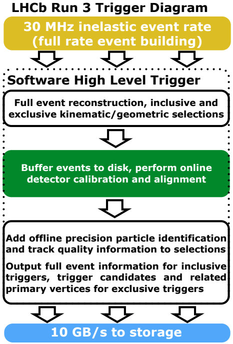

GPUs at LHCb: Allen (CSBS 4, 7 (2020), LHCB-TDR-021)

Allen is LHCb's GPU-based first level software trigger (HLT1)

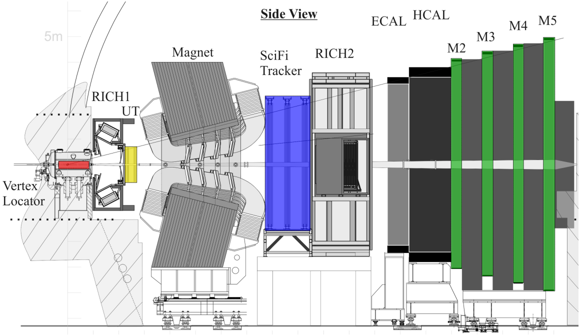

- Decode data from the **VELO**, **UT**, **SciFi**, and **Muon** systems
- Cluster detector data into "hits"
- Build tracks (**VELO**, **UT**, and **SciFi**)
- Find primary vertices (PVs) (**VELO**)
- Match tracks to **Muon** hits

Works as a standalone application or as part of LHCb's software stack

Can be compiled for CPU or GPU with CUDA or HIP

GPUs at LHCb: Performance

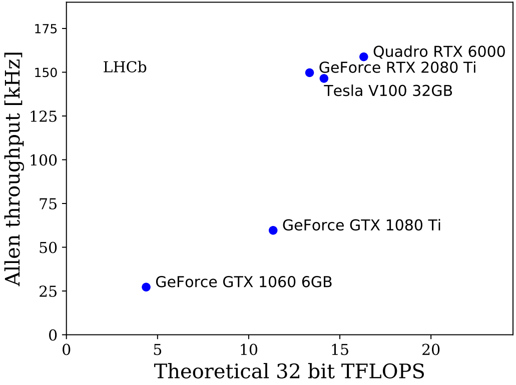

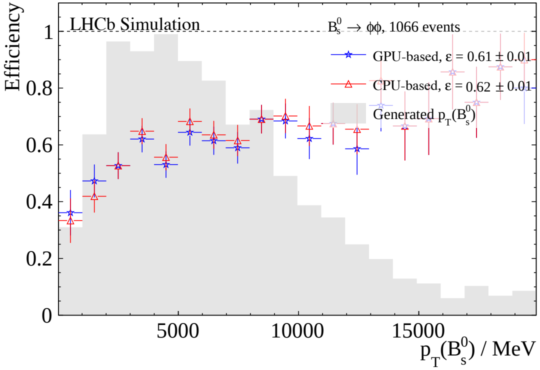

- Can process the full LHC event rate at LHCb with $\sim$ 200 RTX 2080 Tis
- Achieves the baseline HLT1 physics goals for Run 3

GPUs at LHCb: The GPU vs. CPU Decision (arXiv:2105.04031)

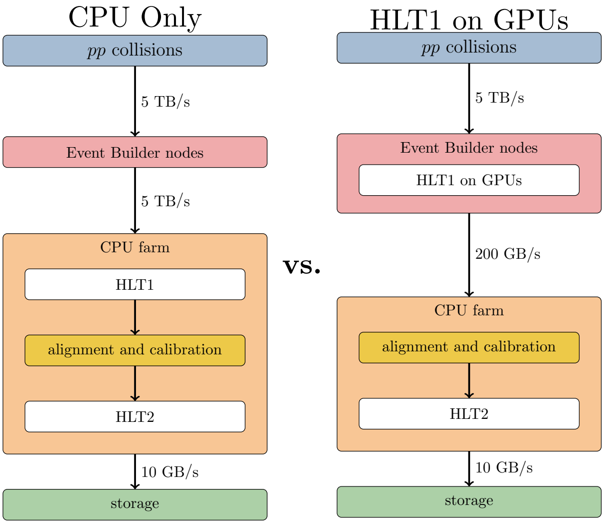

- Both CPU and GPU software triggers meet the basic requirements for HLT1
- CPU-GPU decision requires a detailed cost benefit analysis
- GPU solution leads to cost savings on processors and networking
- Enough throughput headroom for additional features
- **The final verdict:** A GPU-based software trigger will allow LHCb to expand its physics reach in Run 3 and beyond.

Conclusions

- All four LHC experiments are using or studying GPUs in triggers and reconstruction
- The experiments utilize GPUs in very different ways
- GPUs have the potential to expand the reach of the entire LHC physics program

Thank you!

Special thanks to Xiaocong Ai, Andrea Bocci, Attila Krasznahorkay, Felice Pantaleo, David Rohr, Ole Schmidt, and Dorothea vom Bruch for their help preparing this talk!

Backup

ATLAS: GPU R&D

GPU studies for Run 3

- ATL-DAQ-PROC-2012-006
- ATL-DAQ-PROC-2016-035
- Not cost-effective for Run 3 trigger
- Re-evaluating for Run 4 and beyond

Heterogeneous Computing Forum

- Accelerator R&D for online and offline
- ML projects
- FastCaloSim: arXiv:2103.14737
- **A** Common **T**racking **S**oftware (ACTS): See e.g. arXiv:1910.03128