Software Integration Plans for 2023: some Highlights
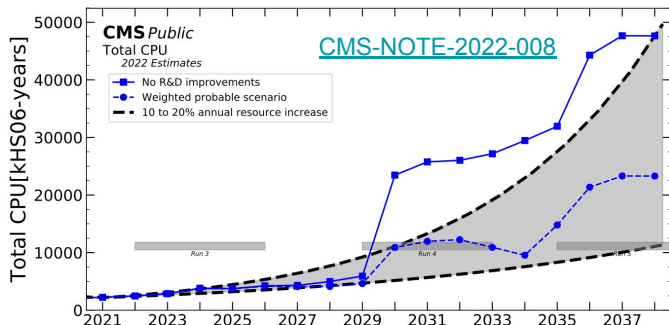
D. Piparo (SFT & CMS), J. Letts (UCSD)

AF 1-9-`22

CMS
O&C

Offline Software and Computing

CERN

UCSD

Run 3 is not "More of the same"

- Very successful start of the Run, for CMS as a whole and for offline (and online) data processing
  - Successfully processed streams meant for parking promptly: reconstructed at the Tier-0 @ >5 kHz during long fills
- Bleeding edge sw integrated in production releases is working flawlessly
  - E.g. new tracking (mkfit, 2x faster), ROOT 6.24, G4 10.7, GCC 10, CUDA code at HLT (+74% throughput thanks to GPUs! 2 NVidia T4 for each 2x64 cores AMD Milan node)
  - LZMA compression for RAW data: a flag for ROOT, -10/15% RAW event size for CMS
- Prompt NanoAOD on new data acquired: analyses already using those!
- Opportunity to anticipate solutions originally imagined for Phase-2
  - E.g. Software for heterogeneous architectures, substantial detector-specific improvements of high-fidelity simulation, aggressive technical improvements.

CMS *Public*
Total CPU
2022 Estimates
CMS-NOTE-2022-008
- No R&D improvements
- Weighted probable scenario
- 10 to 20% annual resource increase

This ambitious approach also requires continuous early integration of bleeding edge software components: CMS aims to do a lot for CMSSW 13_X series, 2023 data taking release

CMSSW 12_X = year 2022, CMSSW 13_X = year 2023

Core Software in the 2023 release

S. Muzaffar, M. Kortelainen

- Update GCC from 10 to 11

- Update ROOT from 6.24 to 6.26

- Add Link Time Optimization flags to all CMSSW libraries: 2-5% more throughput for ALL applications (sim-reco-HLT)
  - LTO to be introduced after GCC 11

- All the above require careful validation of physics results

- Alpaka already available in CMSSW as Run 3 performance portability layer
  - See later slides about offline reco code for GPU

- Profile Guided Optimization studied in detail
  - PGO+LTO → 10% improvement for sim/reco/hlt
  - Detailed study: veritable matrices of different processes built, in general one process is enough to optimize all others (e.g. TTBar for ALL sim processes, JetHT dataset for all real data)
  - Still discussing the best way to put PGO in production

Would it be interesting to hear about this in more detail?

Simulation

- For CMSSW 12_X series (2022), Geant 4 10.7.2: candidate for mass production in 2023

  - Possible rare bug fixes in G4

  - Gamma general process **2% speedup** - ready

  - CMS improved DormandPrince stepper **3% speedup** - ready

  - Tracking cut from 1 keV to 25 keV **1-10% speedup** - validation required, no difference expected

  - LTO **3-5% speedup** (see Core slide) - validation required, no difference expected

  - G4FElowE in ECAL **~15% speedup** - validation not yet converged

- For 13_X series (2023), Geant 4 11.1.1: candidate for mass production in 2024

  V. Ivantchenko

  - Updated muon system geometry (CMS upgrades that during the run)

  - A naturally faster and more robust G4 potentially enhanced with:

    - Neutron general processes, combined multiple scattering/tracking for electrons, G4HepEM, GPU aided HGCal simulation (Adept/Celeritas R&Ds)

Run 3 and Phase-2 Reconstruction

- Alpaka is the performance portability layer of CMS for Run 3
  - Already in CMSSW_12_X
- Port current offloading framework from CUDA to Alpaka - by the end 2022
- Detectors having GPU code move to Alpaka - during H1 2023
  - Mid 2023: all CUDA modules and framework is replaced by Alpaka
- GPU code for offline/online (besides what is available now)
  - H1 2023 (directly in Alpaka): HGCal
  - Early 2023: Particle Flow in CUDA
  - Tracking
    - H1 2023: partial offload of tracking (Cellular Automaton) and primary vertex for offline (~10% of the Run 3 reco sequence, substantial for Phase-2 as well)
    - Will need massive validation during 2023
- Side note: code interpreted with ROOT might get in contact with Alpaka headers. This is new, especially in the context of a big experiment setup. A lot of potential and, perhaps, some hurdles.

A. Di Florio. P. Srimanobhas