Parallel 3B: Computing Performance

Sept. 26, 2023 (2:00-3:30 p.m.)

co-chaired by G. Folger and S.Y. Jun

ALICE	*Ivana Hrivnacova et al.*
Hokkaido University, Room B	14:00 - 14:10
ATLAS	*Marilena Bandieramonte*
Hokkaido University, Room B	14:10 - 14:20
CMS	*Prof. Vladimir Ivantchenko*
Hokkaido University, Room B	14:20 - 14:30
LHCb	*Witold Pokorski*
Hokkaido University, Room B	14:30 - 14:40
IF/Neutrino	*Hans-Joachim Wenzel*
Hokkaido University, Room B	14:40 - 14:50
IBS/Underground Physics	*Eunju Jeon*
Hokkaido University, Room B	14:50 - 15:00
Medical Physics	*Susanna Guatelli*
Hokkaido University, Room B	15:00 - 15:10
Discussion	
Hokkaido University, Room B	15:10 - 15:30

Keywords

- Optimization
- FastSim/ML
- GPU: (Offloading EM particles, optical photons)

ALICE: Sandro Wenzel (CERN) and Ivana Hrivnacova (IJClab IN2P3/CNRS)

- Geant4 11.0.4 within VMC since LHC Run3
  - FTFP_BERT_EMV+optical + NystromRK4
  - First benchmarks from the Run3 production system: TGeo (22%), field (12%), stepper(14%)
- Optimization
  - Automatically tuning “production cuts” but keeping the hit (30% CPU reduction)
- ML
  - Avoid ZDC transport (2-3x slower) with replacing predicted output by 2D image (GAN) based on particle properties

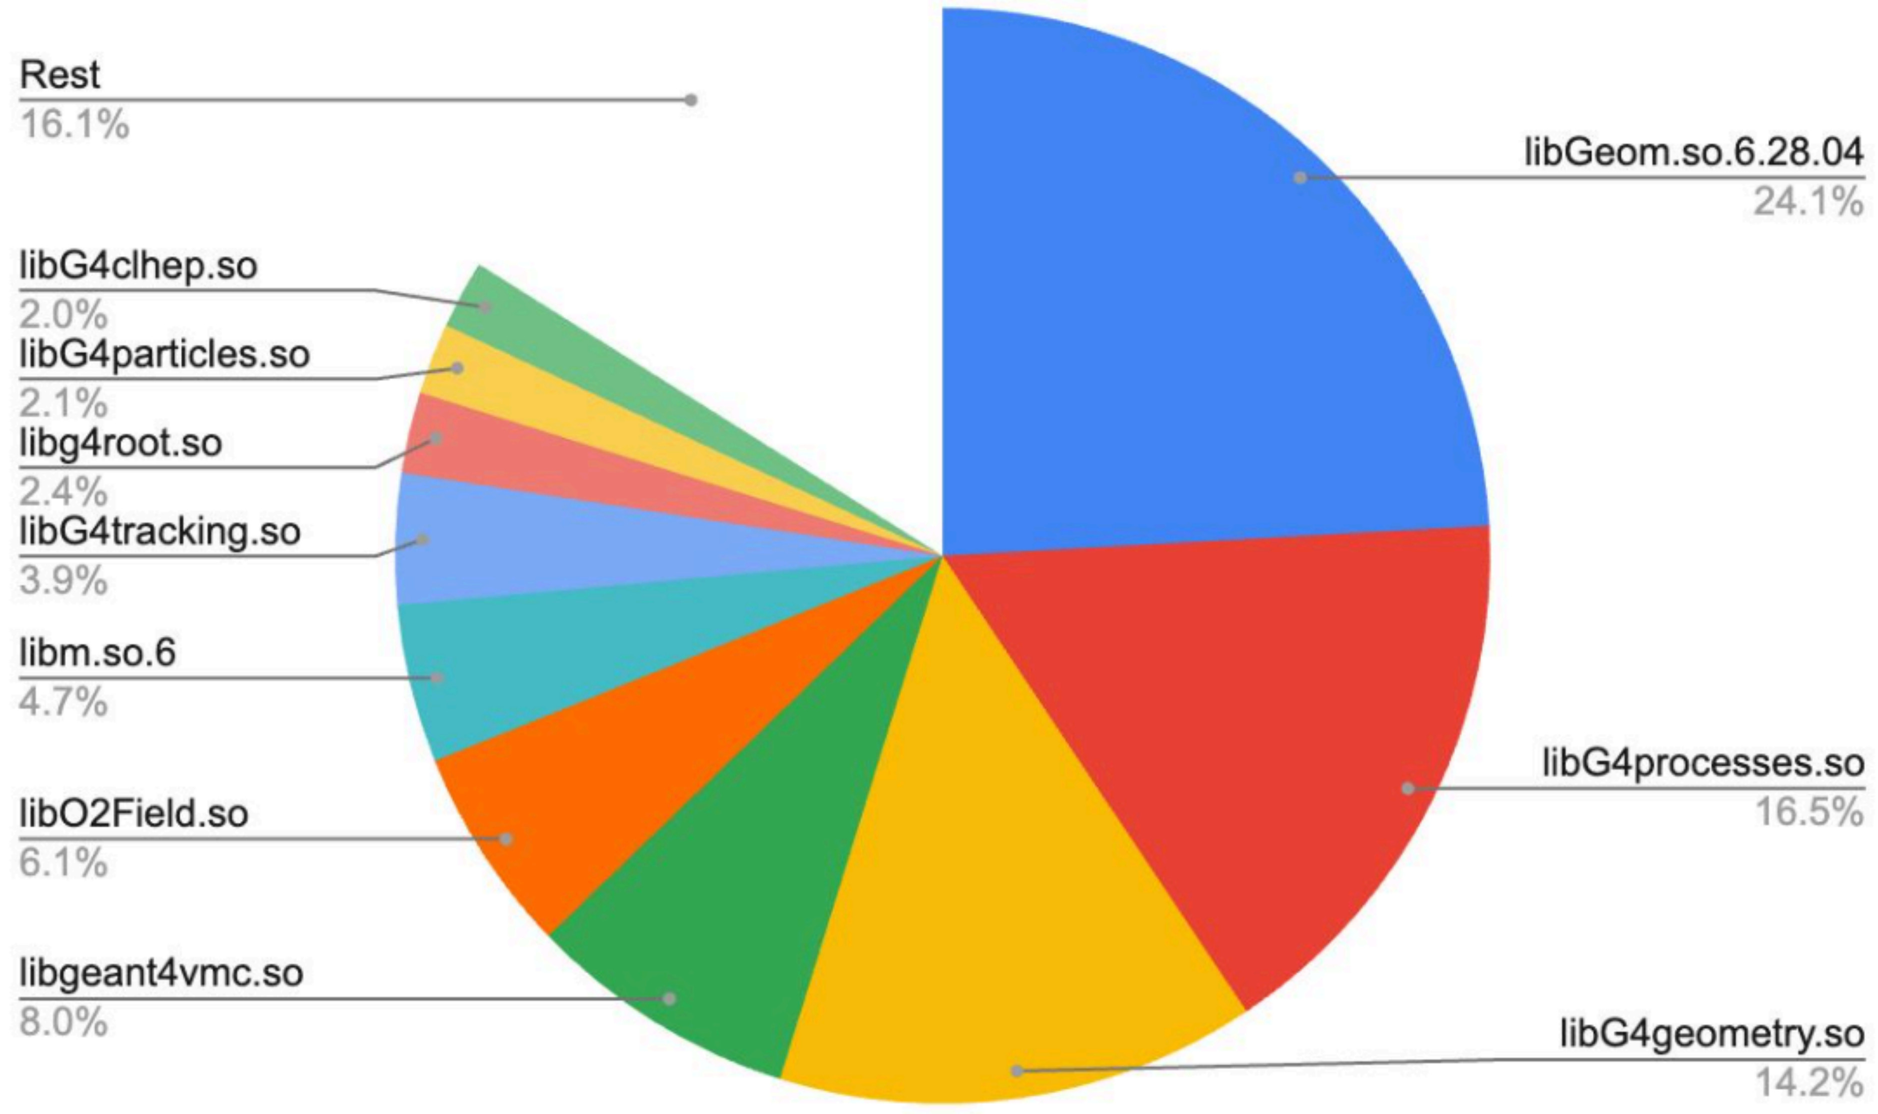

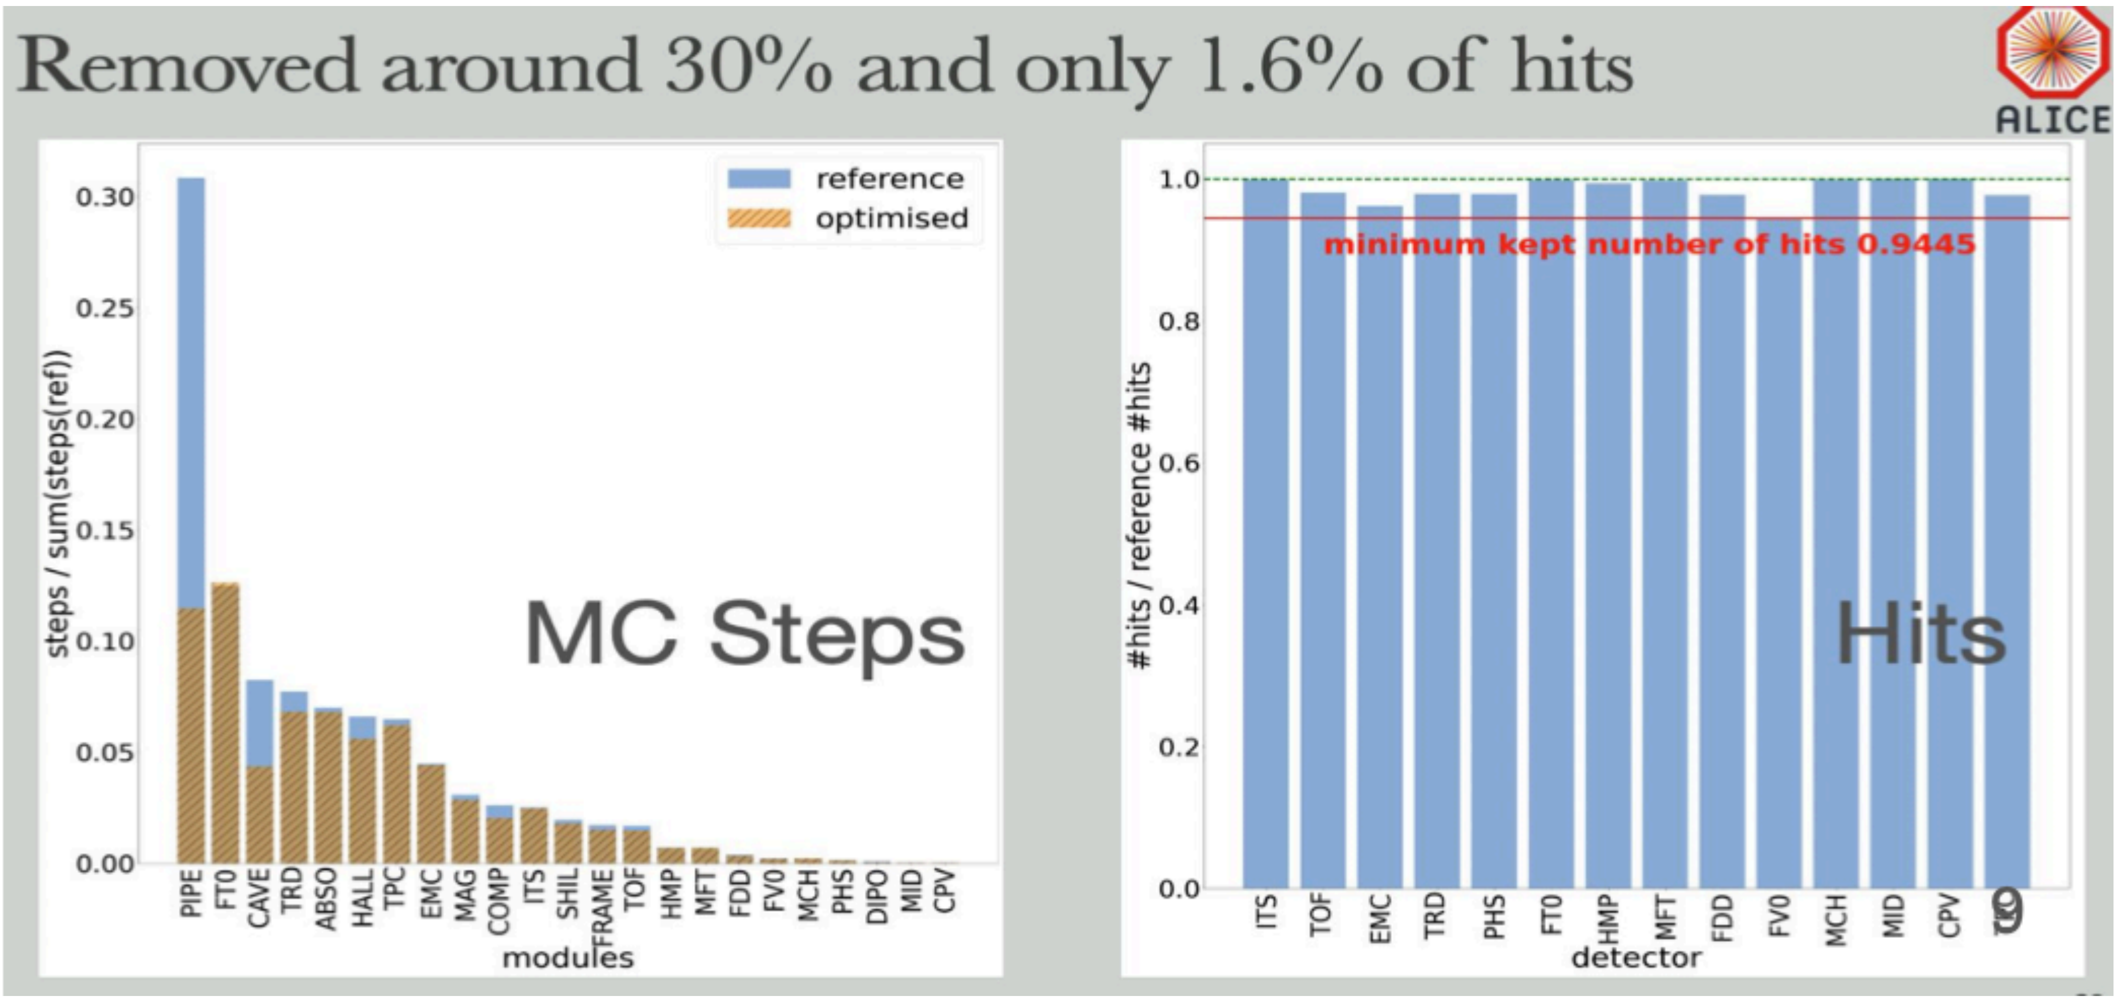

ATLAS: Marilena Bandieramonte (University of Pittsburgh)

- Geant4 10.6.p03 + ATLAS patches for Run3
- Run3 optimization
  - VeGeom, Woodcock, EMEC, G4HepEM, etc.
  - Overall optimization: mc21 (33%), mc23 (48%)
  - (LTO/PGO, FDO, Auto-FDO: ~5%)
- GPUs
  - GPU friendly EMEC implementation
  - Initial demonstration with TileCal
  - Celeritas or AdePT as a FullSimLight plugin

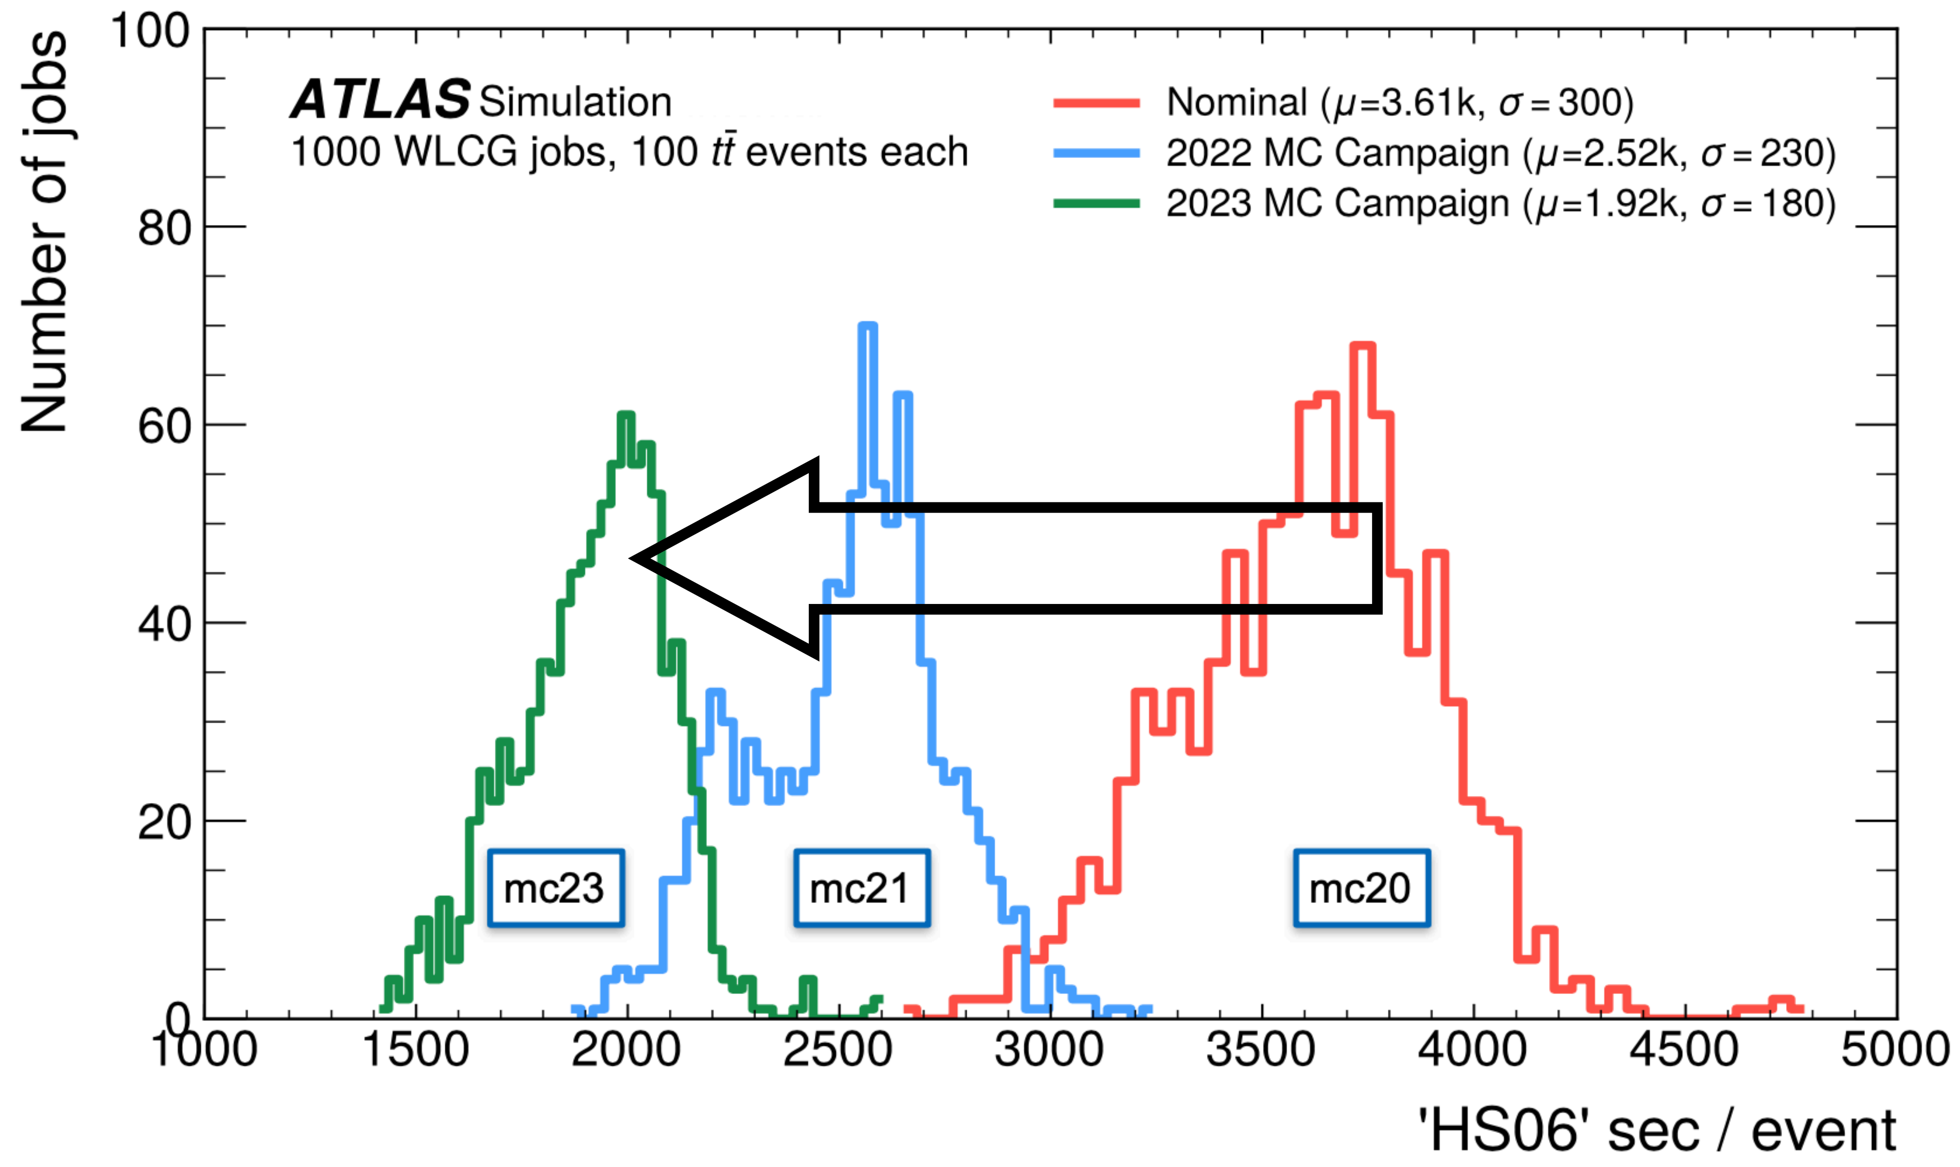

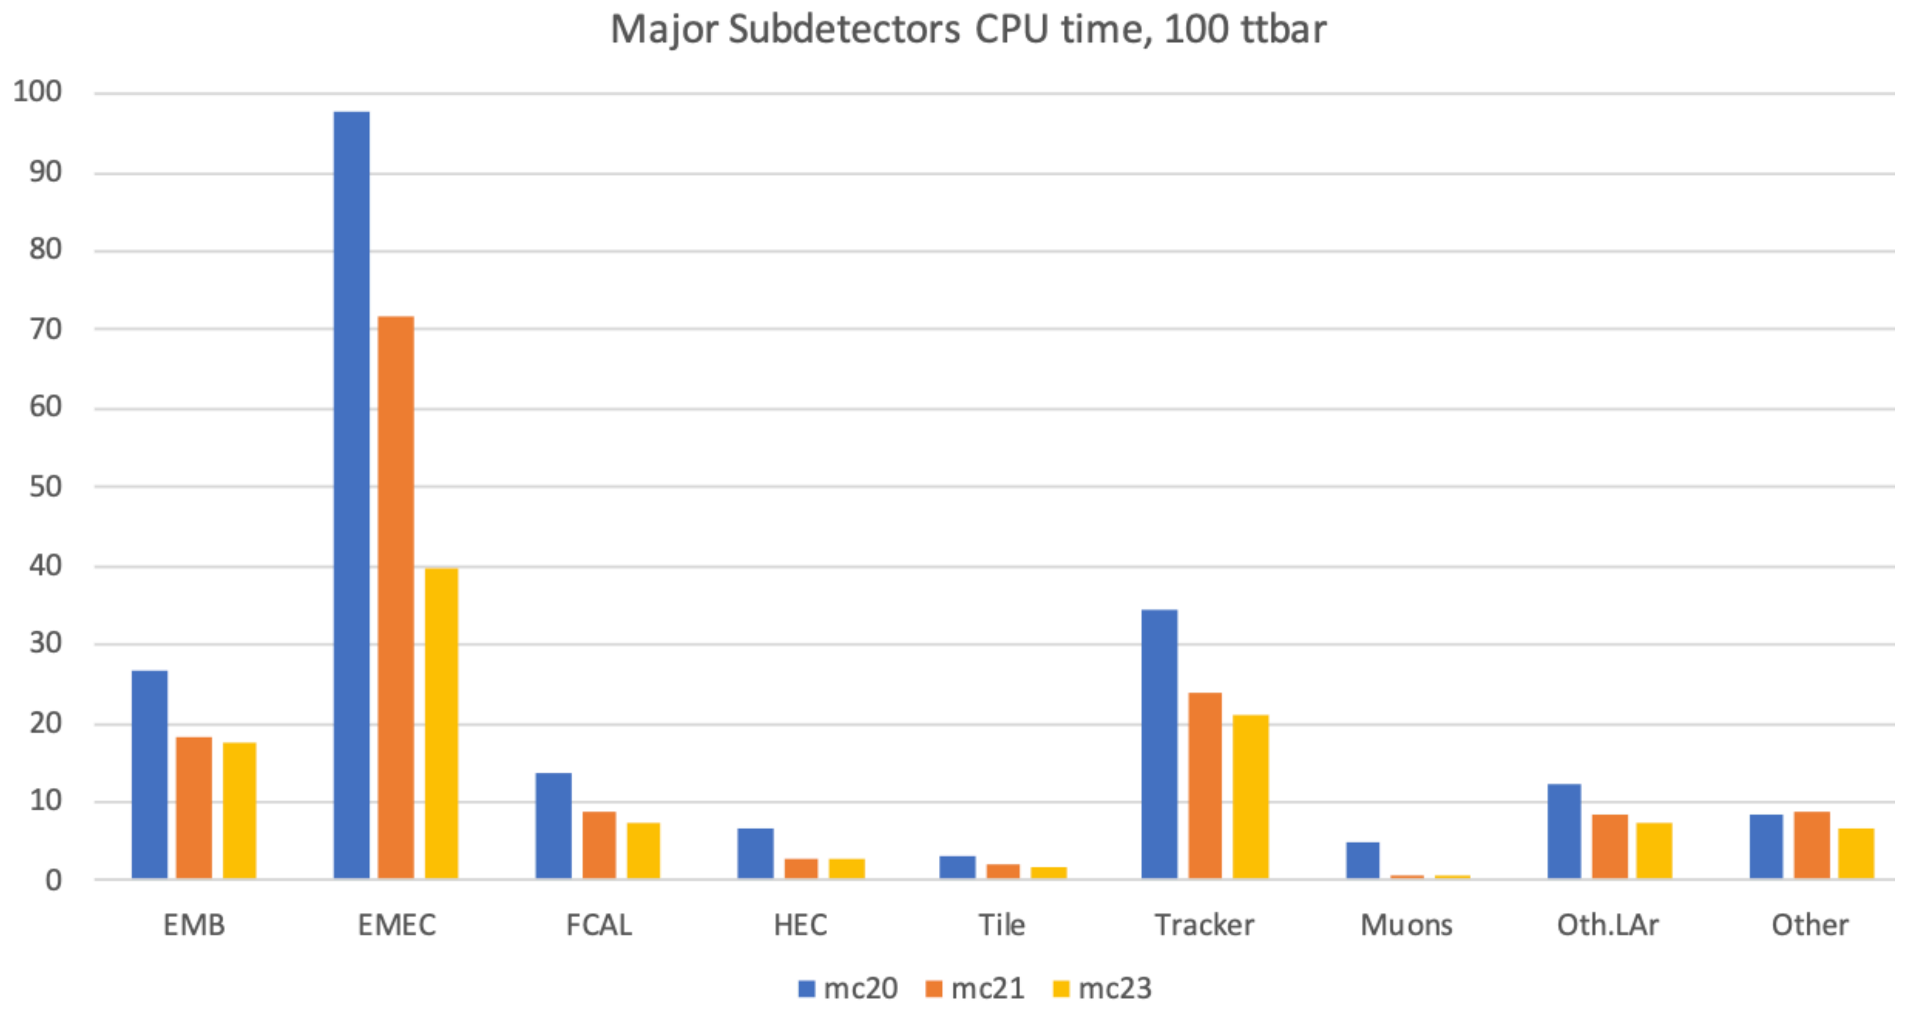

CMS: Vladimir Ivantchenko (CERN/Princeton) and *et. al.*

- CPU improvements of Run3 MC Production w.r.t. Run2
  - 2022-2023: G4 10.7.p02 + CMS patches (25%)
  - 2023: G4 11.1.p01 + LTO + Gamma General
  - 2024: G4 11.1.p02 + G4TransportationWithMsc (30%)
  - Physics samples: Minimum bias (36%), ttbar (32%), BSM T1tttt (27%) and Z -> ee (32%).
- Results of test-beam analysis are good in general
  - Pion and proton are within uncertainty (~3 %)
  - Anti-proton above 5 GeV is overestimated for ~5%
  - Kaon below 10 GeV is underestimated for ~10 %

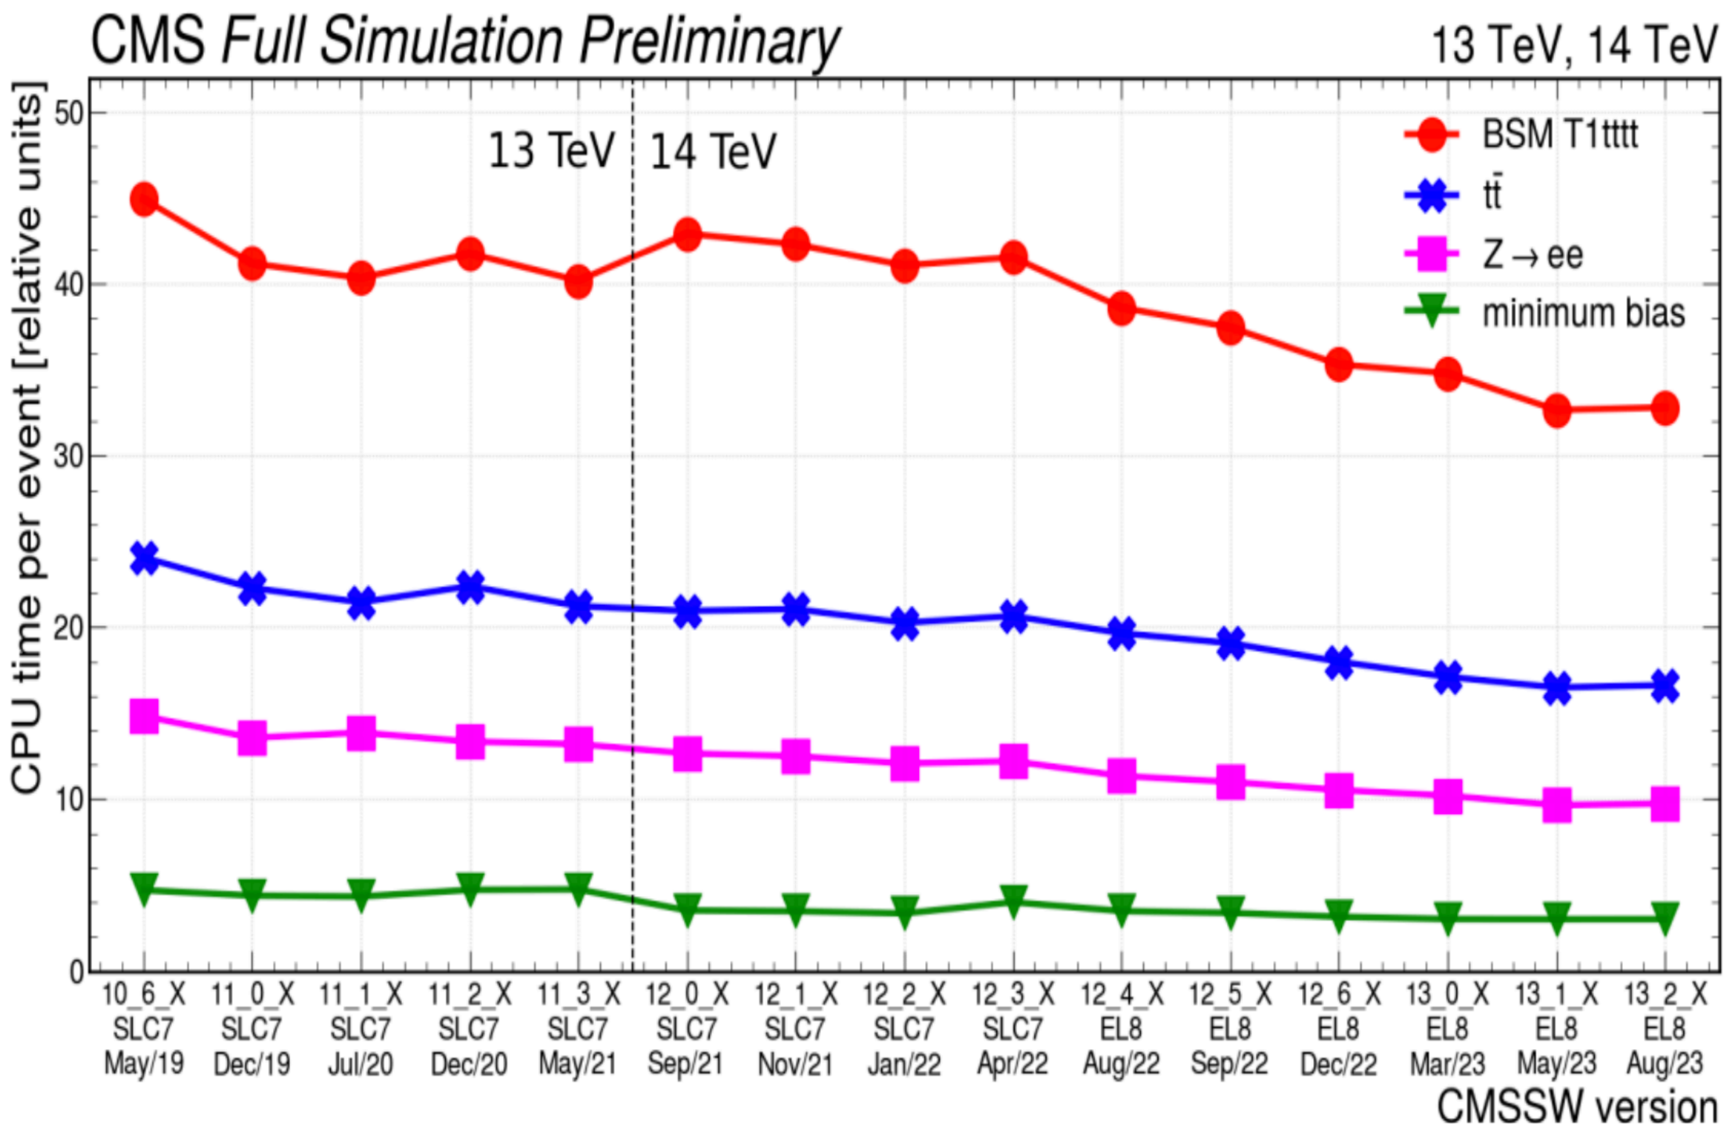

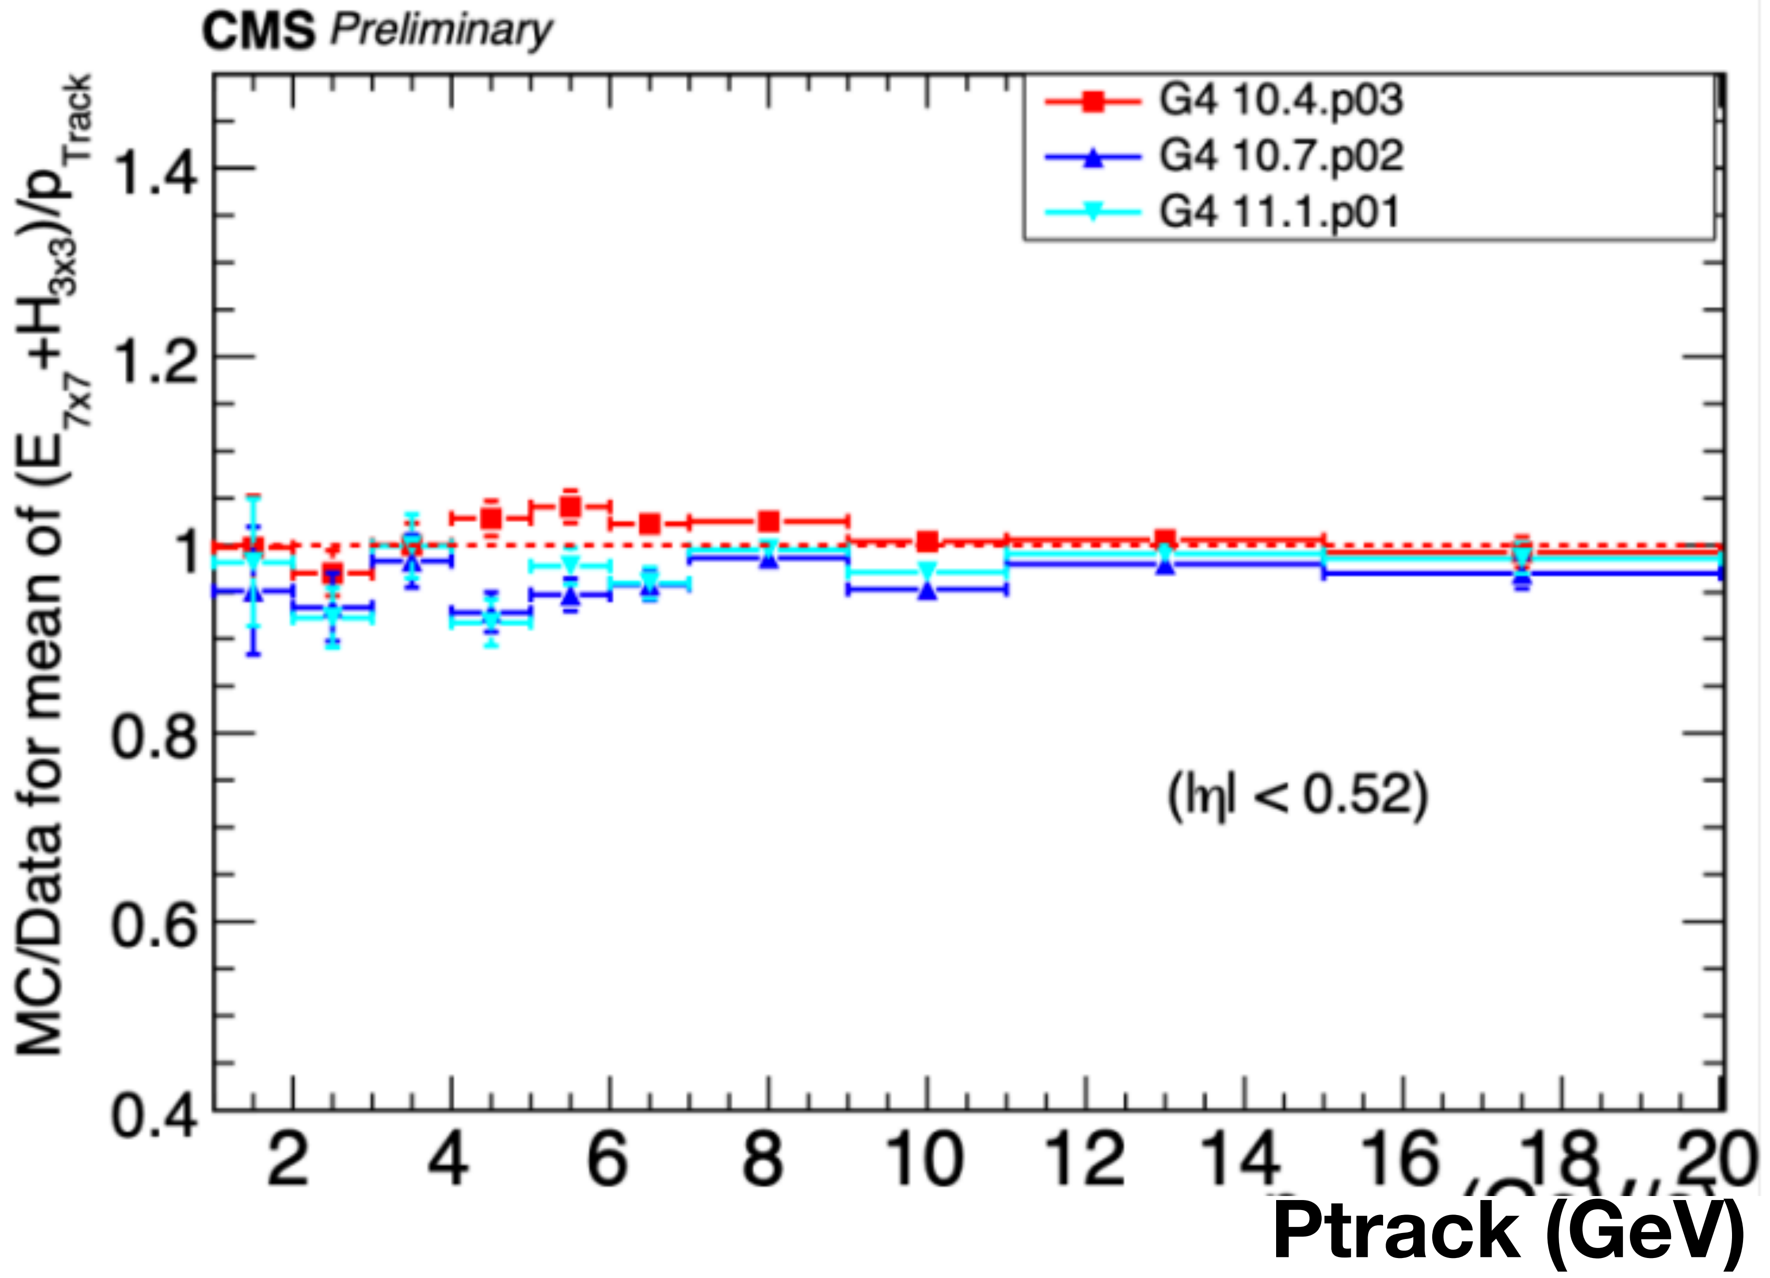

LHCb: G. Corti, M. Mazurek, W. Pokorski, D. Popov, M. Veltri

- Modernization of the whole LHCb software in Run3
  - MT, code vectorization, algorithmic optimization, ..
- Geant4 10.6.4 in production, 10.7.3 in nightly, plan to add 11 in near future
  - Performance regression tests (CPU/memory, …)
  - Profiling with flamegraphs to pinpoints CPU hotspots (WIP)
- ML with Gaussino (Geant4+ML4Sim): CaloChallenge
- Investigating hybrid workflows with GPUs
  - Electromagnetic showers - AdePT Integration
  - Optical photons in RICH (Opticks, Mitsuba3)

First Benchmarks in Gassino

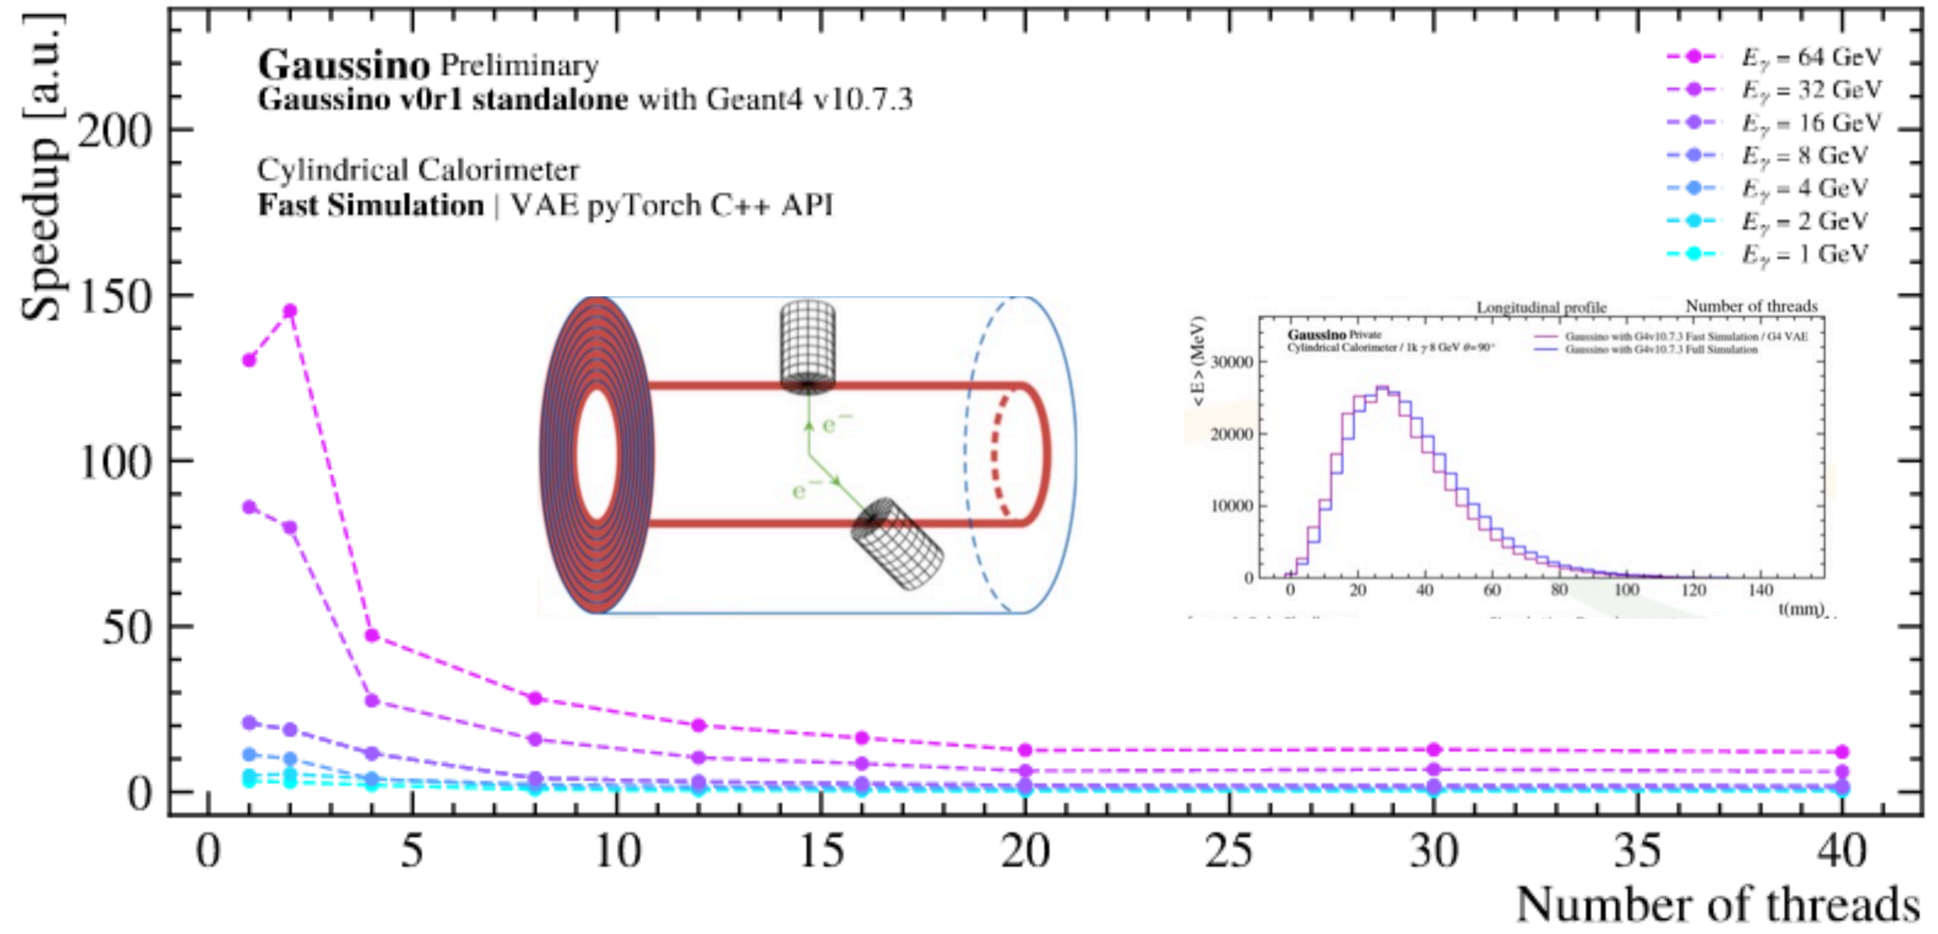

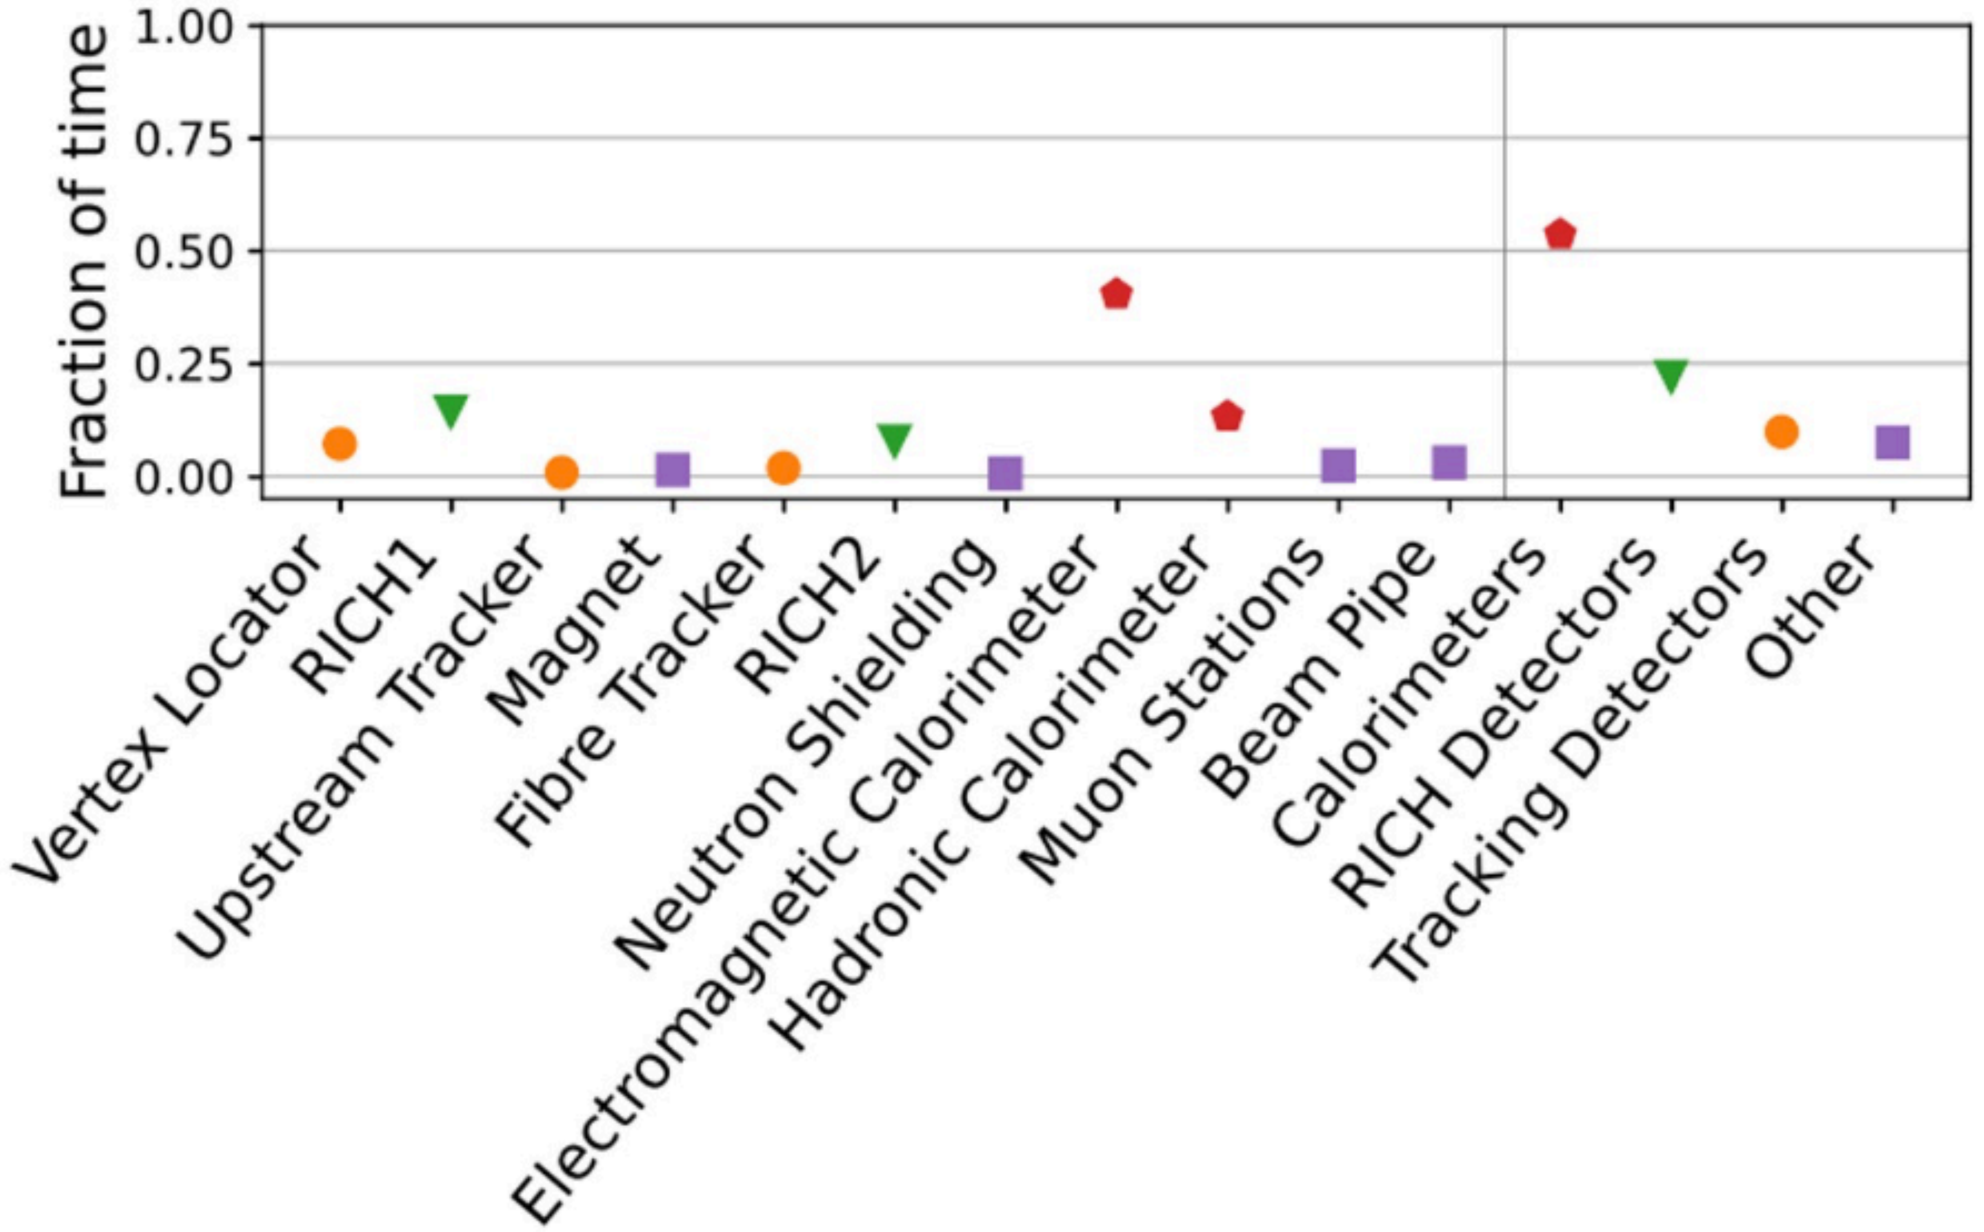

Intensity Frontier and Neutrino Experiments @Fermilab: Hans Wenzel

- Geant4 versions (10.3 and 11.1) with the art framework
- Intensity Frontier
  - Mu2e: MT (2 threads, ~3.4Gb memory)
  - g-2: no major issue (tested the G4 symplectic stepper)
- LArSoft/LArTPCs (DUNE, ICARUS, SBND, MicroBooNE, …)
  - Refactored Geant4 simulation into LArG4
  - CPU fraction of Geant4 is relatively small (%-level)
  - Optimization for memory (geometry, step limit, I/O)
- GPU: ~50K optical photons/MeV energy loss in LAr
  - Lookup table —> Full Geant4 simulation
  - Integrating Opticks into artg4k (100x-1000x speedup)

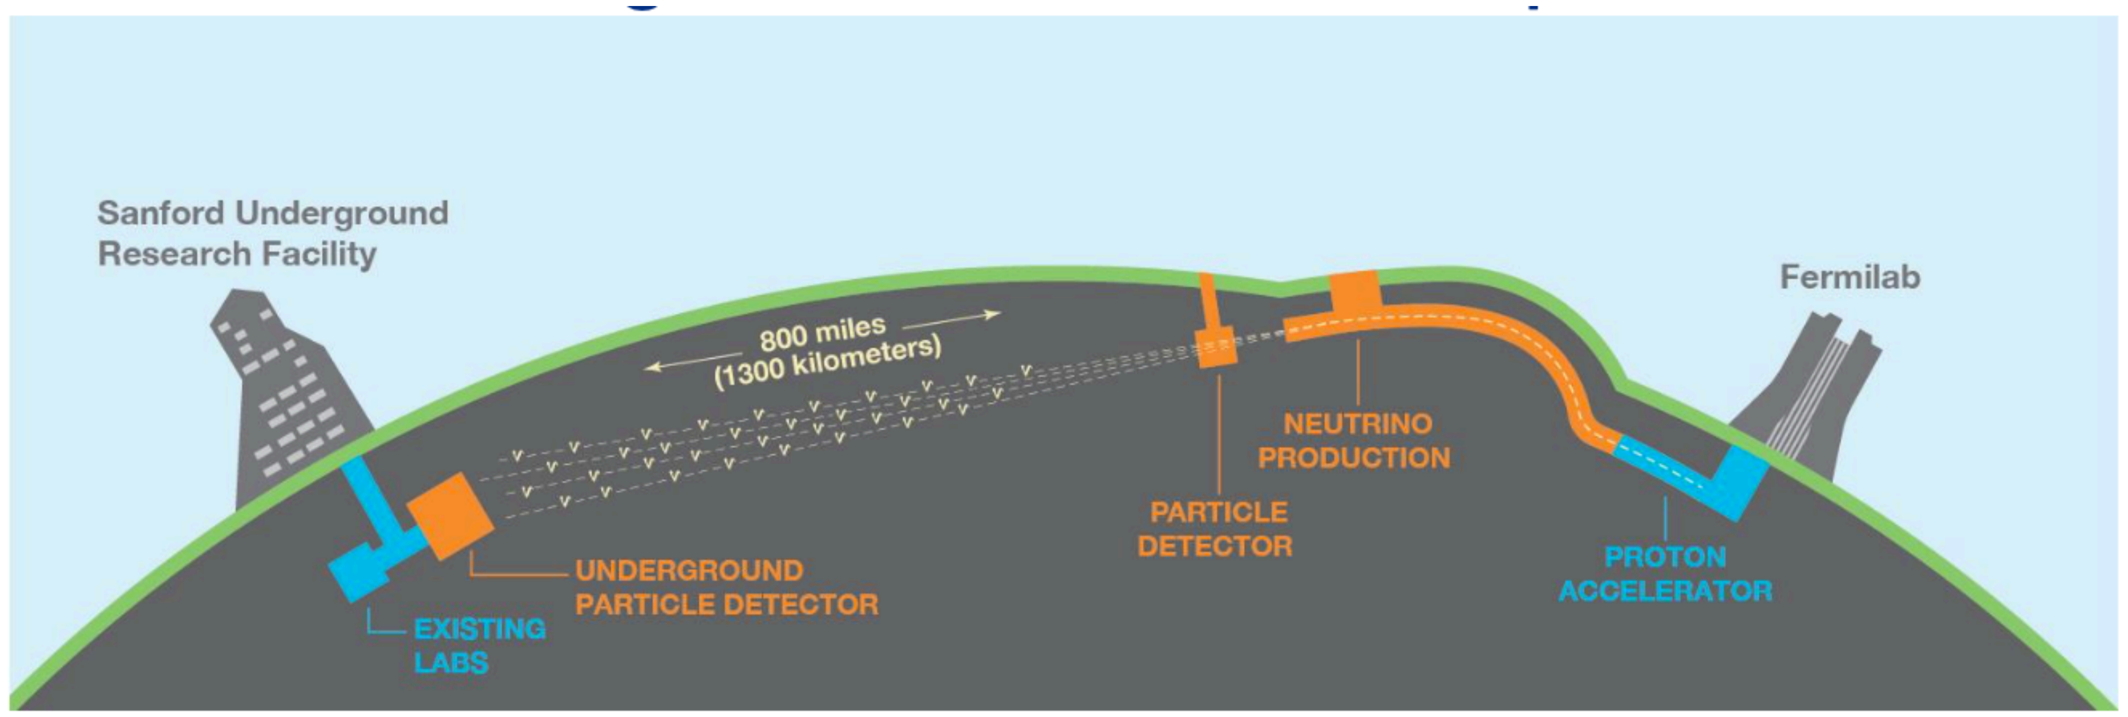

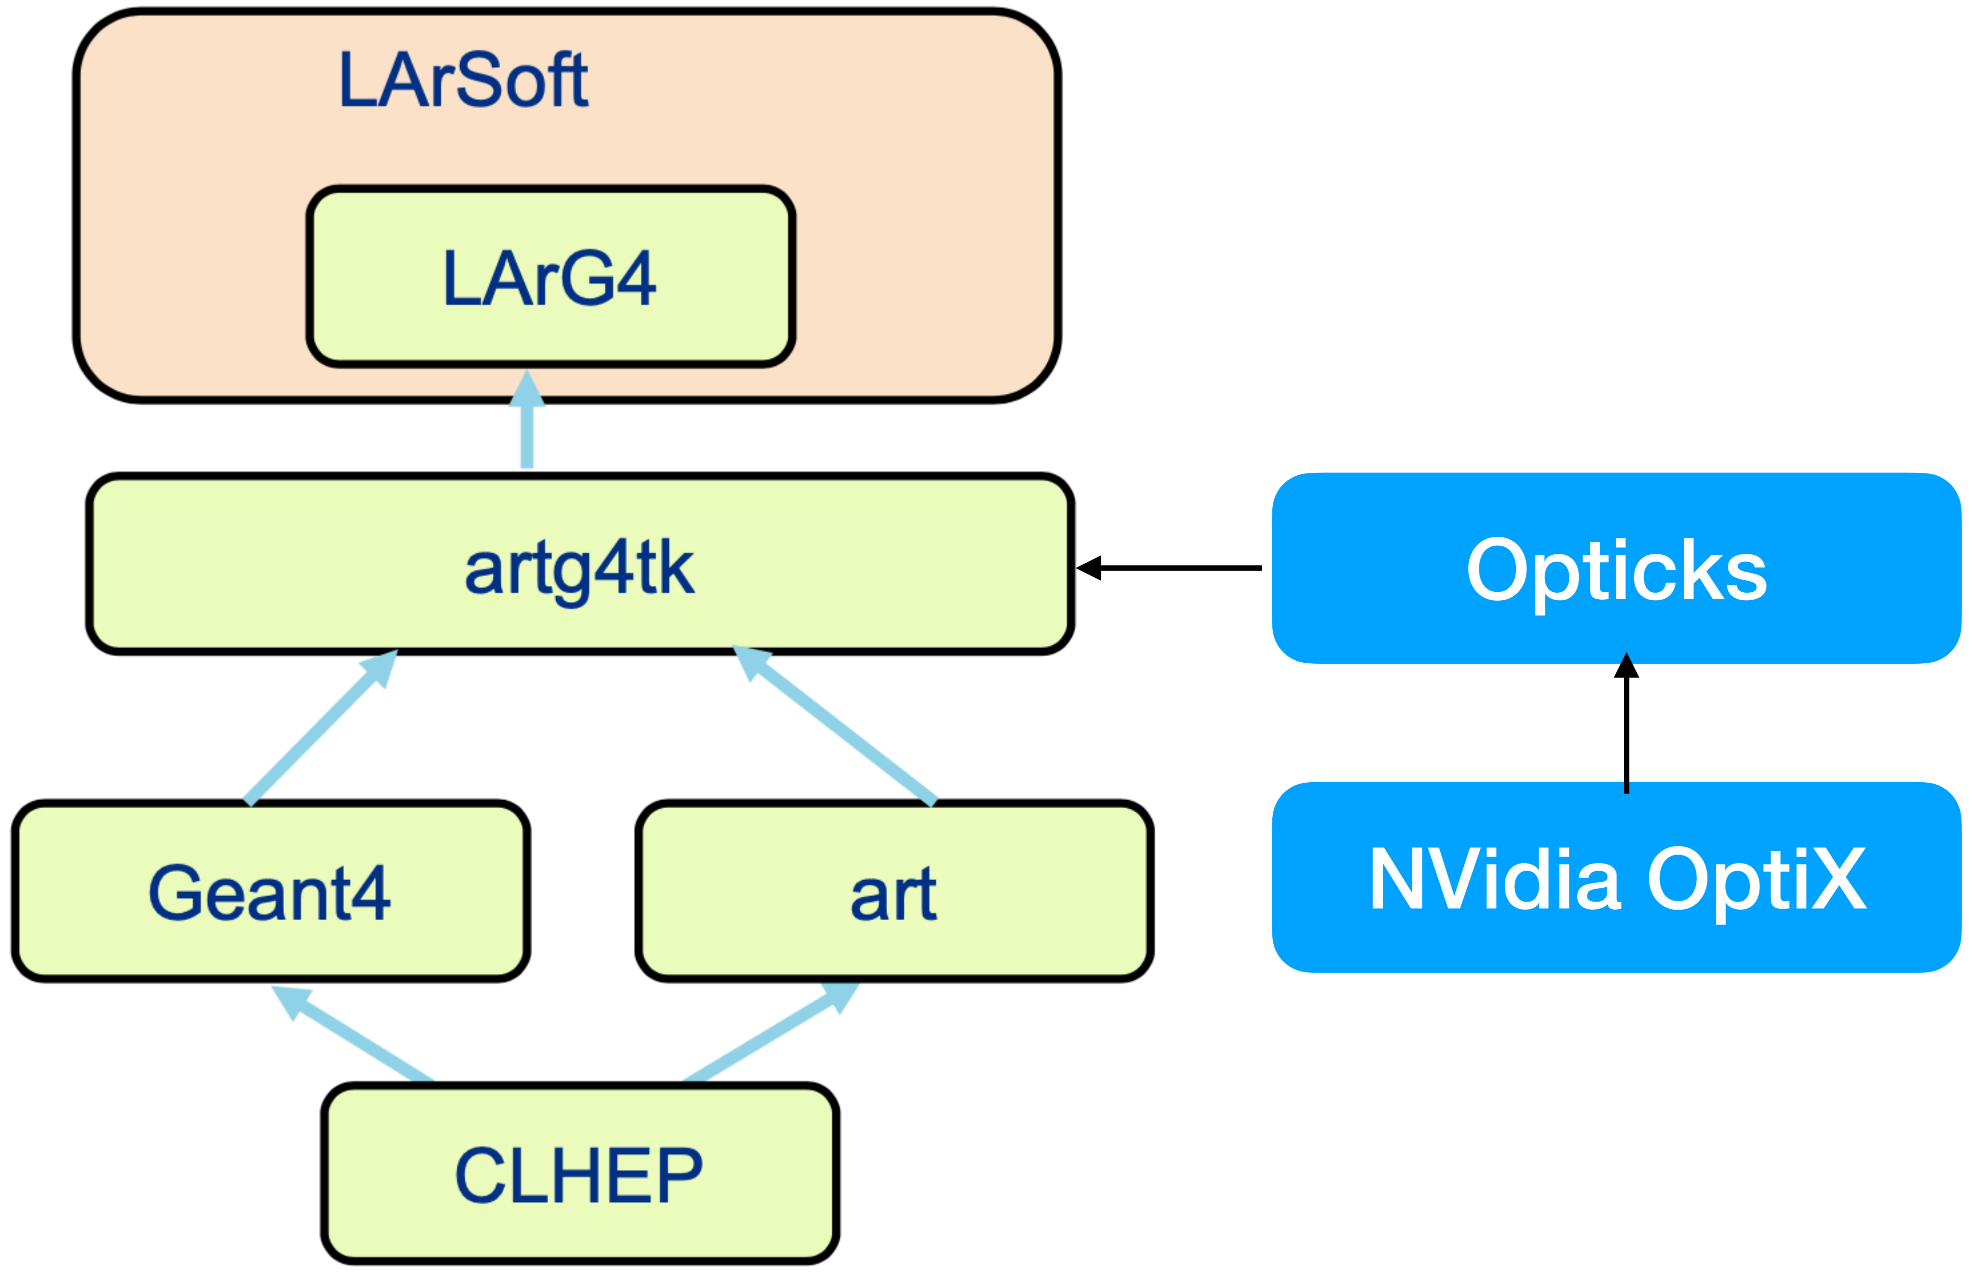

Underground Physics at CUP: Eunju Jeon (CPU/IBS)

- Underground physics
  - DM search, 0νββ decay, etc.
- Major concerns in simulating underground physics
  - How to estimate ultra-low backgrounds with shielding configuration without consuming a lot of CPU time
  - Gamma/neutron , cosmic muons backgrounds (comparisons to MUSIC, FLUKA)
  - Example of Geant4 regression tests: monitoring (n, gamma) capture cross section
- GPU
  - Optical simulation to increase its speed by using a package of Opticks
- Phonon simulation: use/test of G4CMP

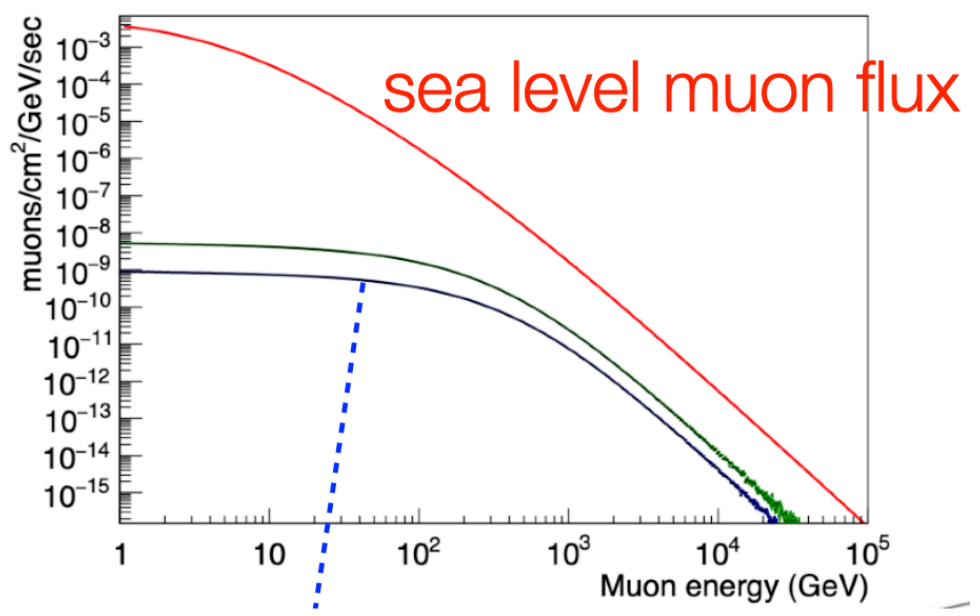

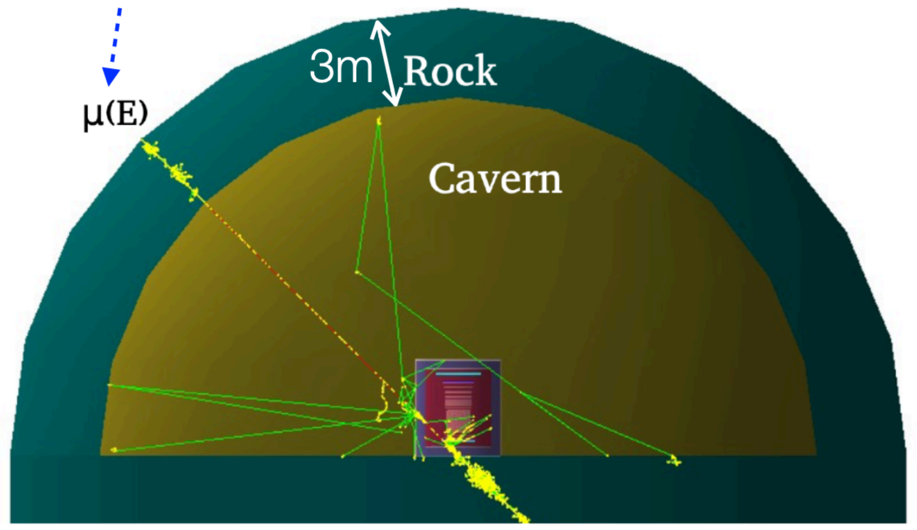

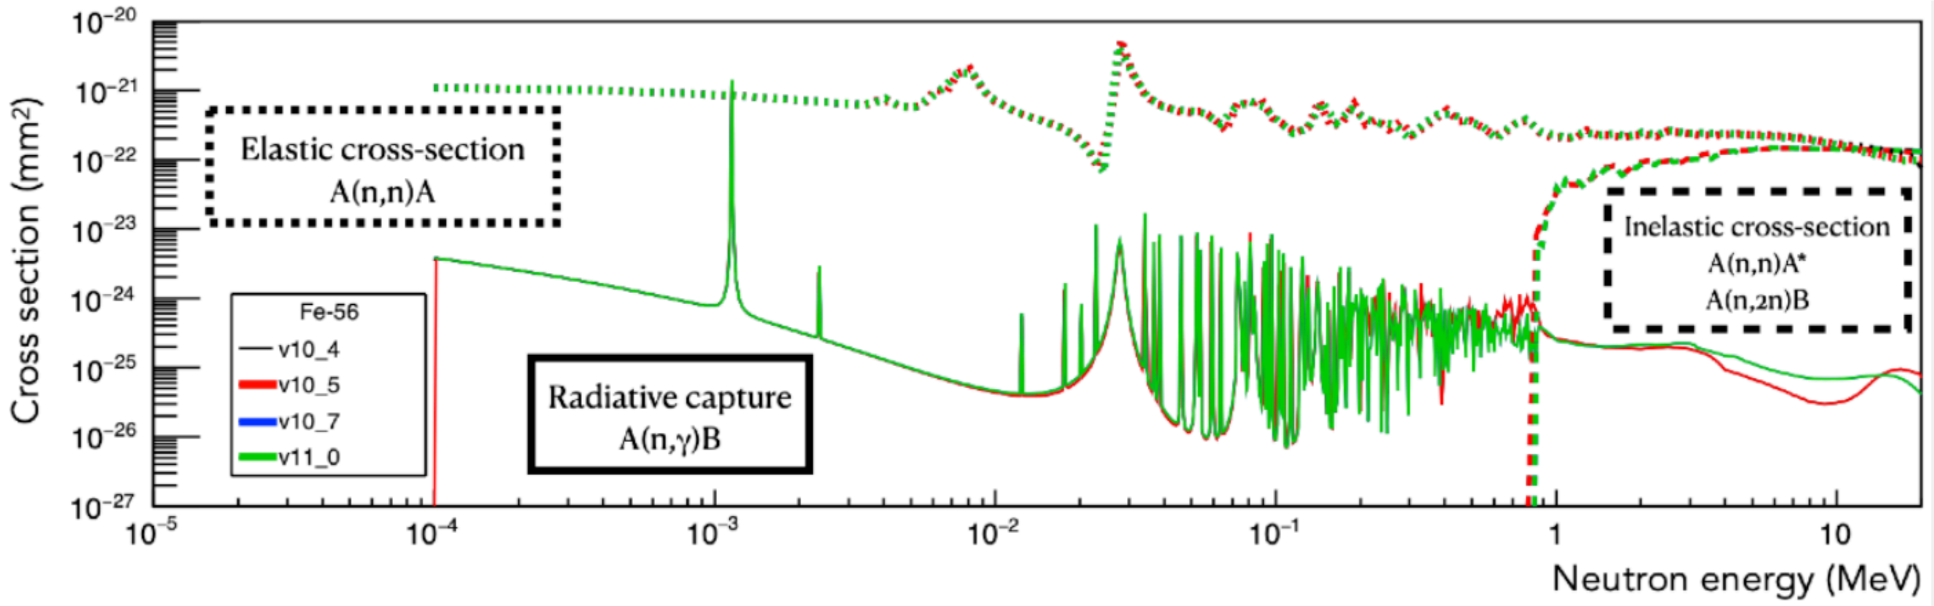

Medical Physics: Susanna Guatelli (University of Wollongong)

- G4-MED tests with Geant4 11.1
    - Computing performance along with regression tests (EMStandard opt3, G4-DNA opt2, etc.)
    - A set of tests (e.g., Total attenuation coeff. test in water, Brachytherapy, Fano Cavity, CCCSTest, etc.)
- ML: Fast dose predictions with adapted 3D U-Net
    - 3D density matrix -> 3D energy deposition matrix (Geant4 to train, validate and test)
    - Speedup of the ML-based dose engine vs. Geant4: 10^6 considering 1 computing unit A
- DL: Graph Neural Network (e.g., emulate BLOB)

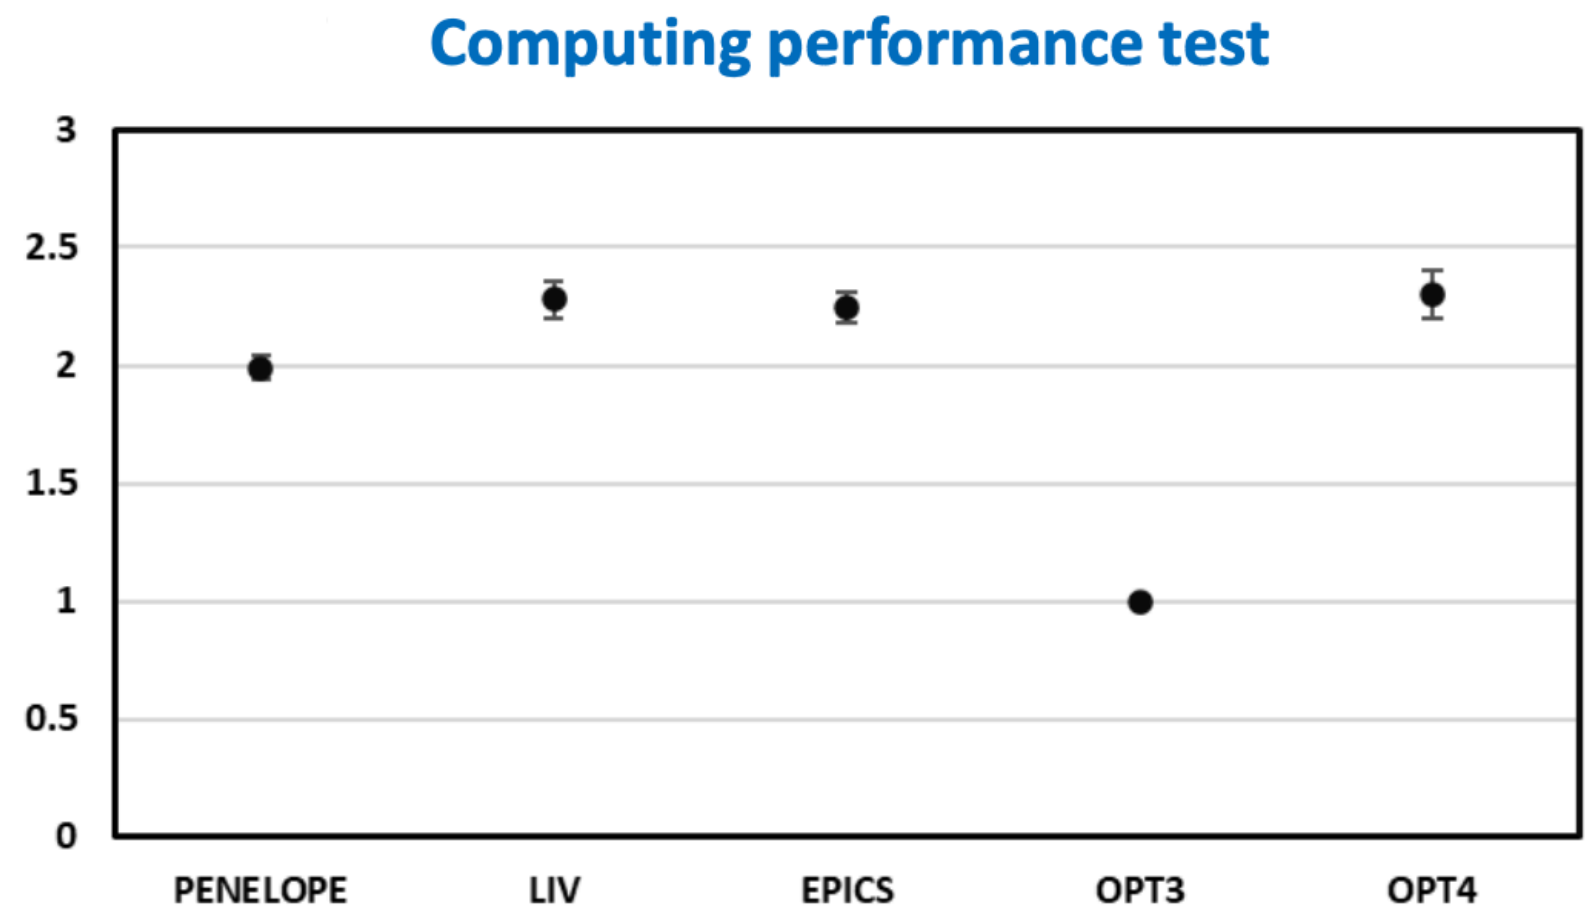

{G4 data} Adapted 3D U-Net ~0.1sec

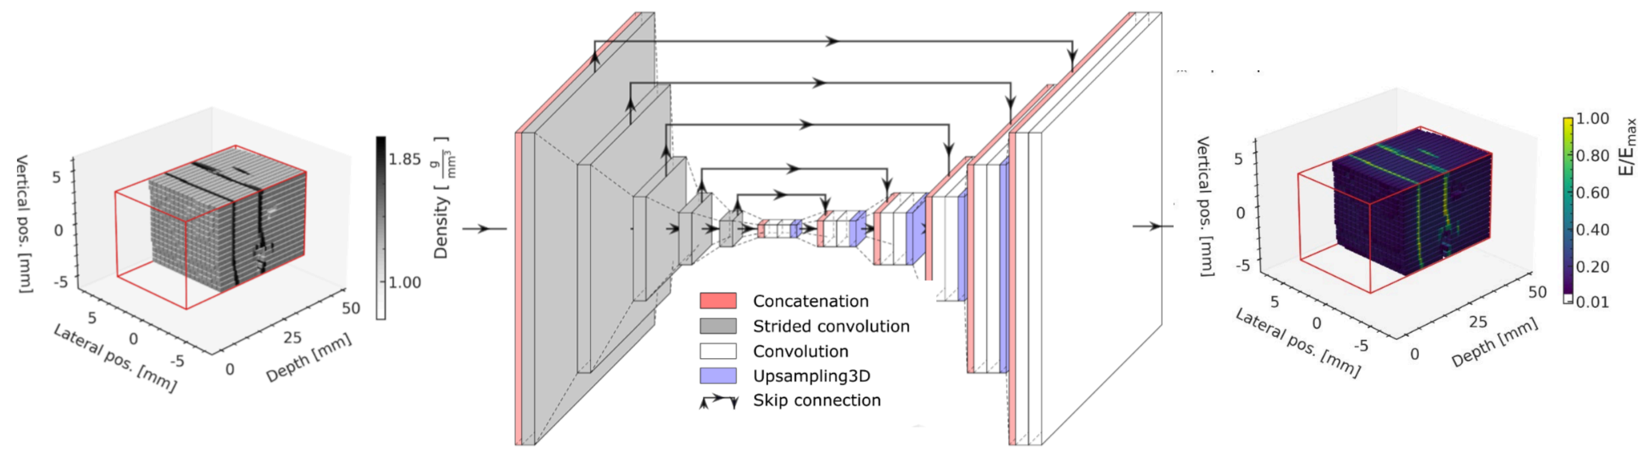

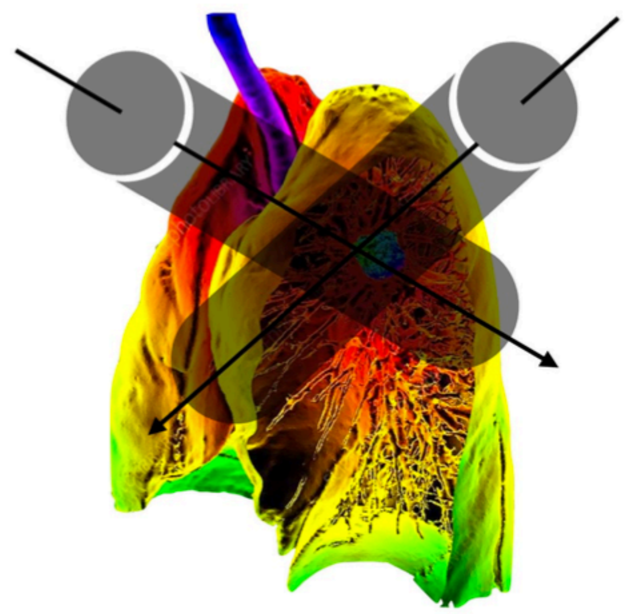

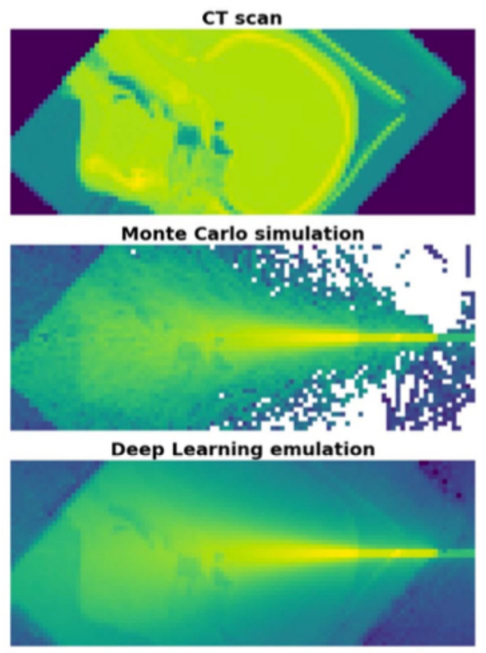

Geant4

GNN

Summary

- Continuing efforts of high-level and/or fine-grain optimizations for Geant4 use cases
- Increasing demands for offloading EM particles and optical photons to GPUs
- Exploring more sophisticated FastSim, ML/DL approaches for Geant4 applications

Many thanks to all speakers and contributors!