Optimizing the ATLAS Geant4 Detector Simulation Software

Dr. Evangelos Kourlitis

on behalf of the ATLAS Geant4 Optimization Task Force

October 24th 2022

LHC Analysis Workflow

real data

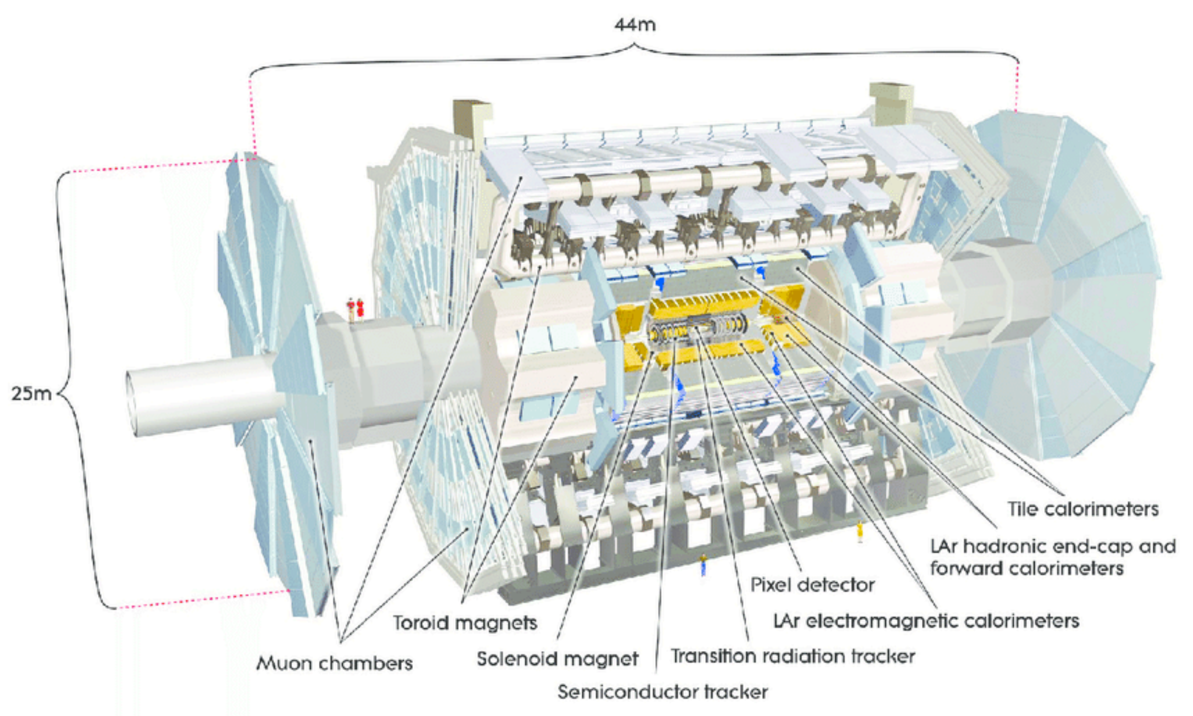

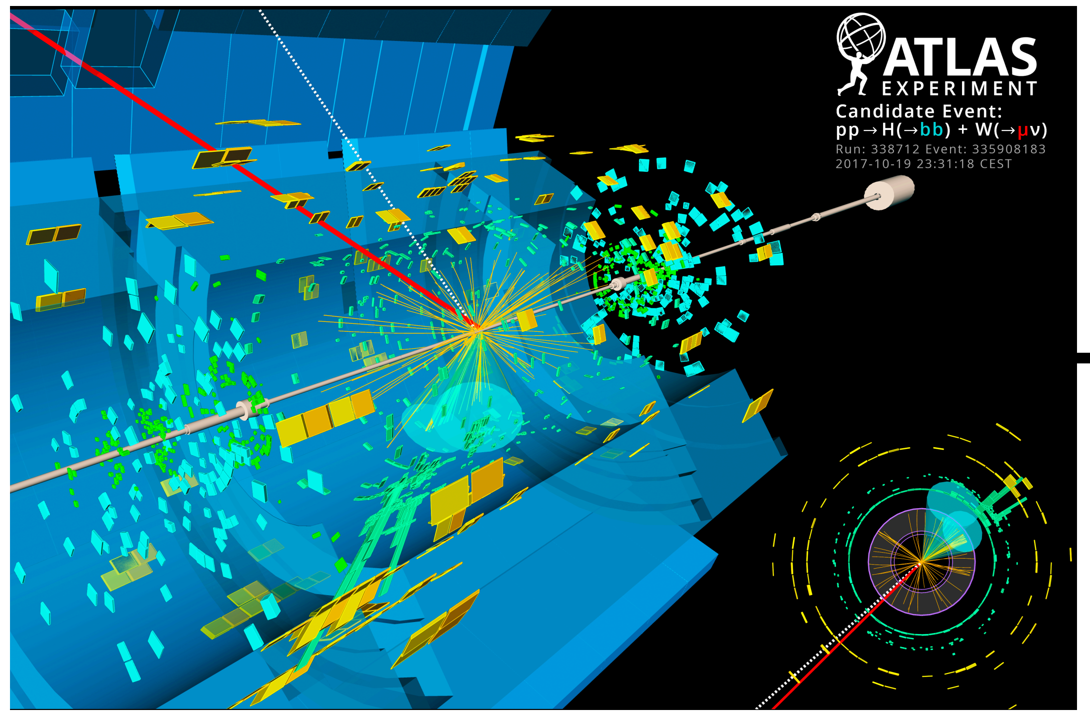

RAW Data

Reconstruction & Analysis

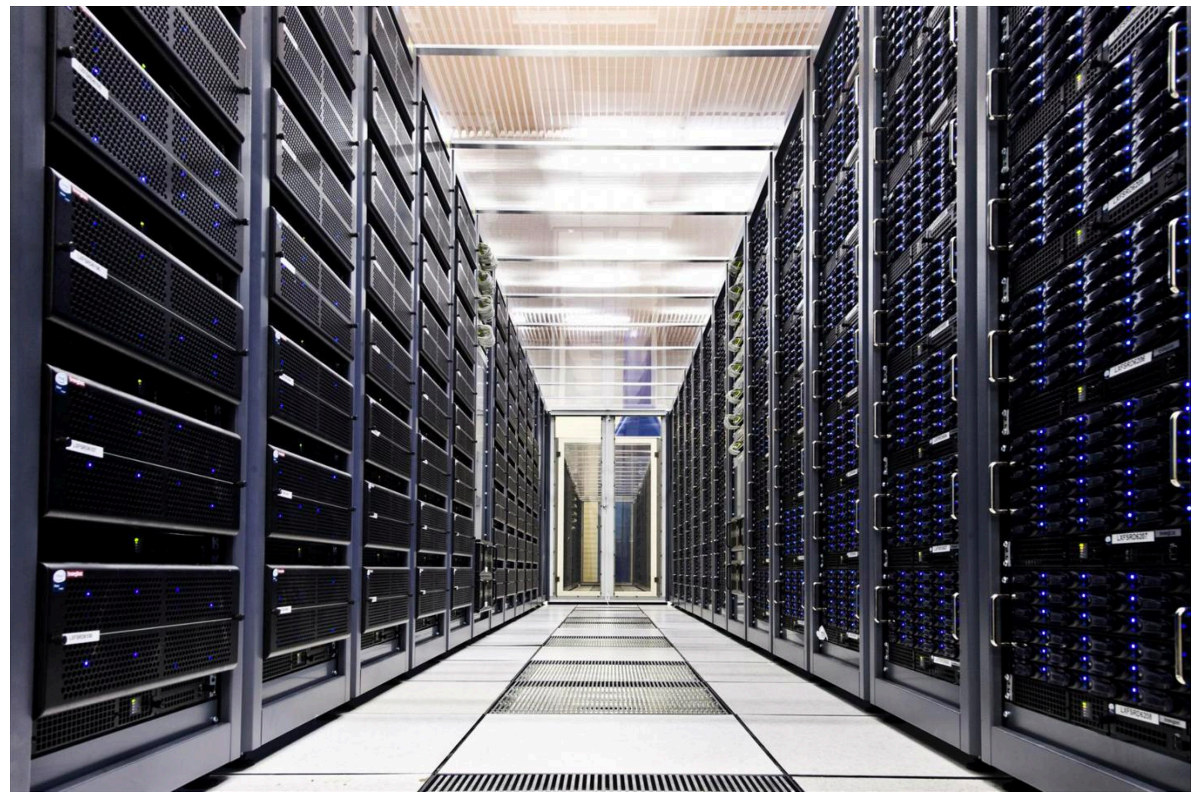

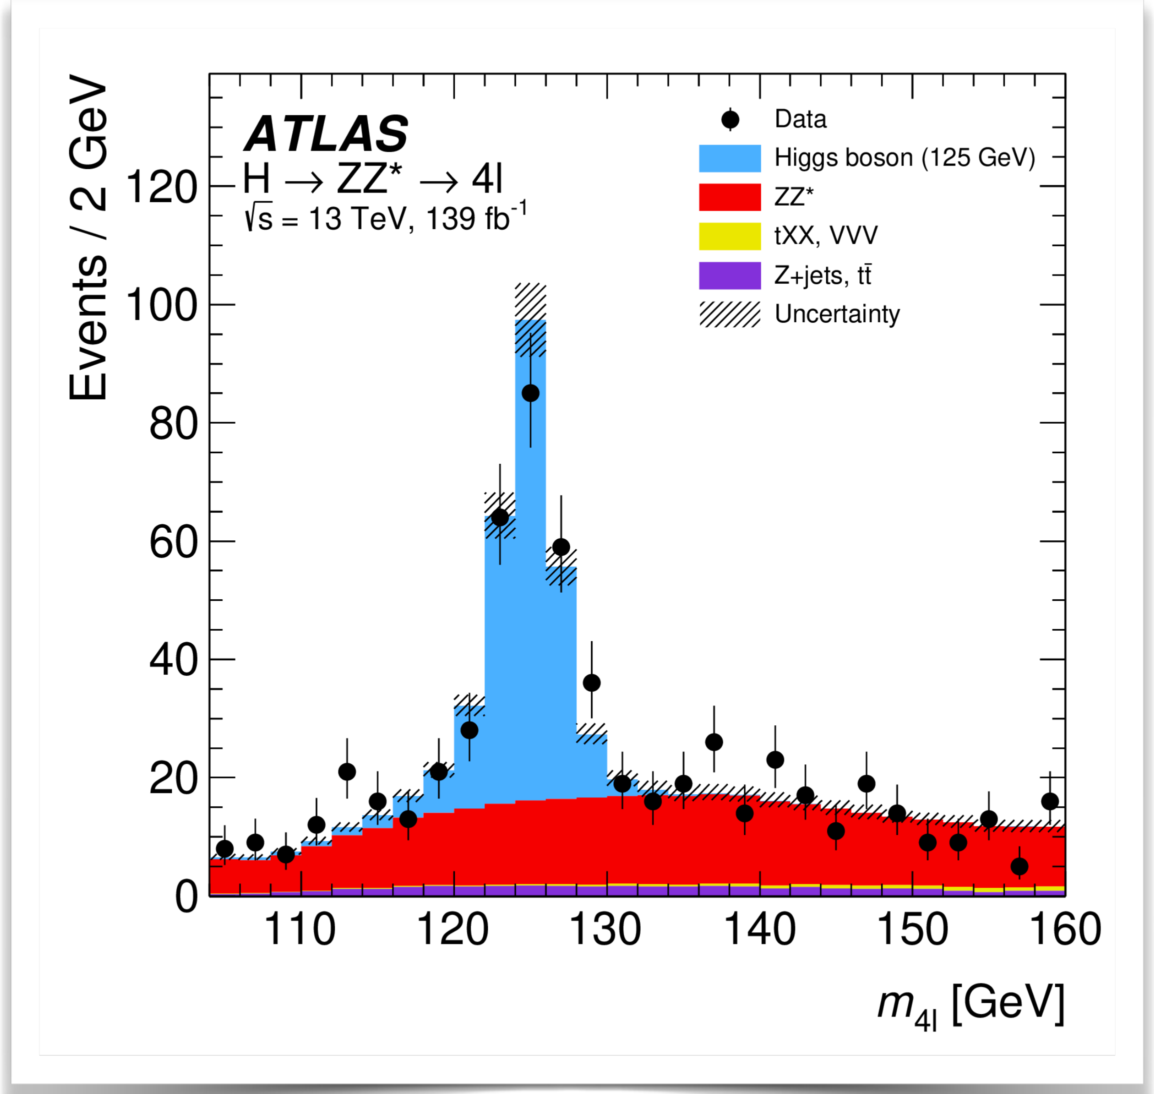

arXiv:2207.00320

sim data

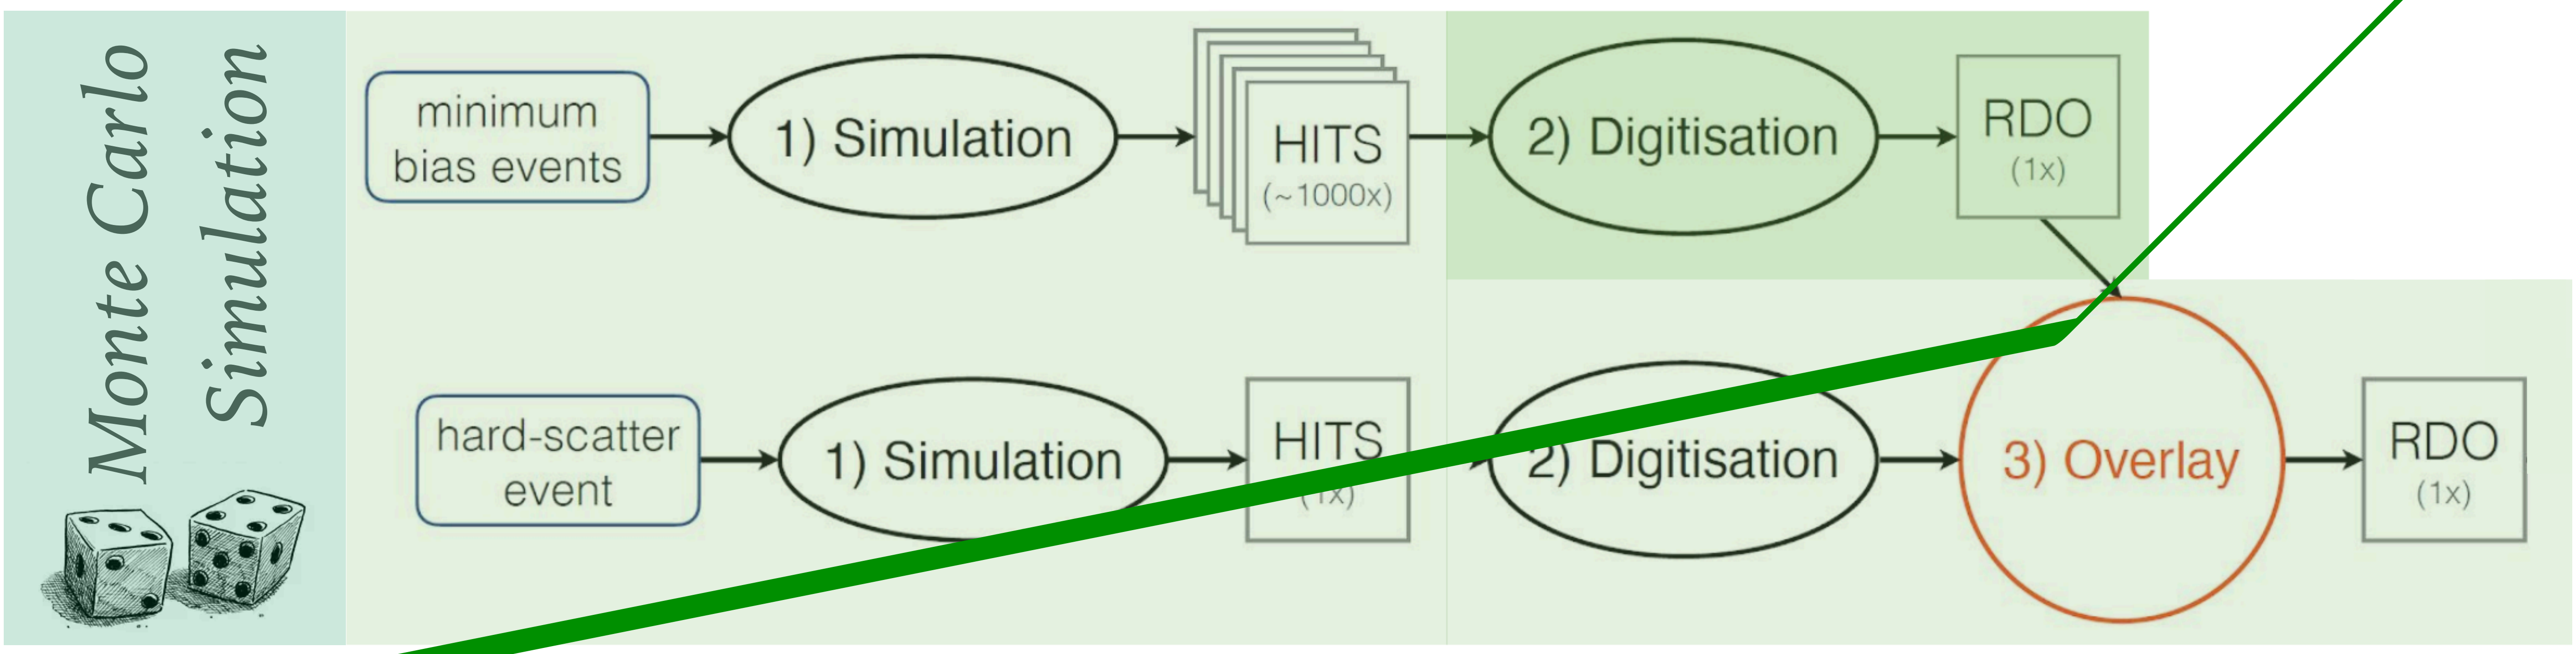

LHC Analysis Workflow

real data

RAW Data

sim data

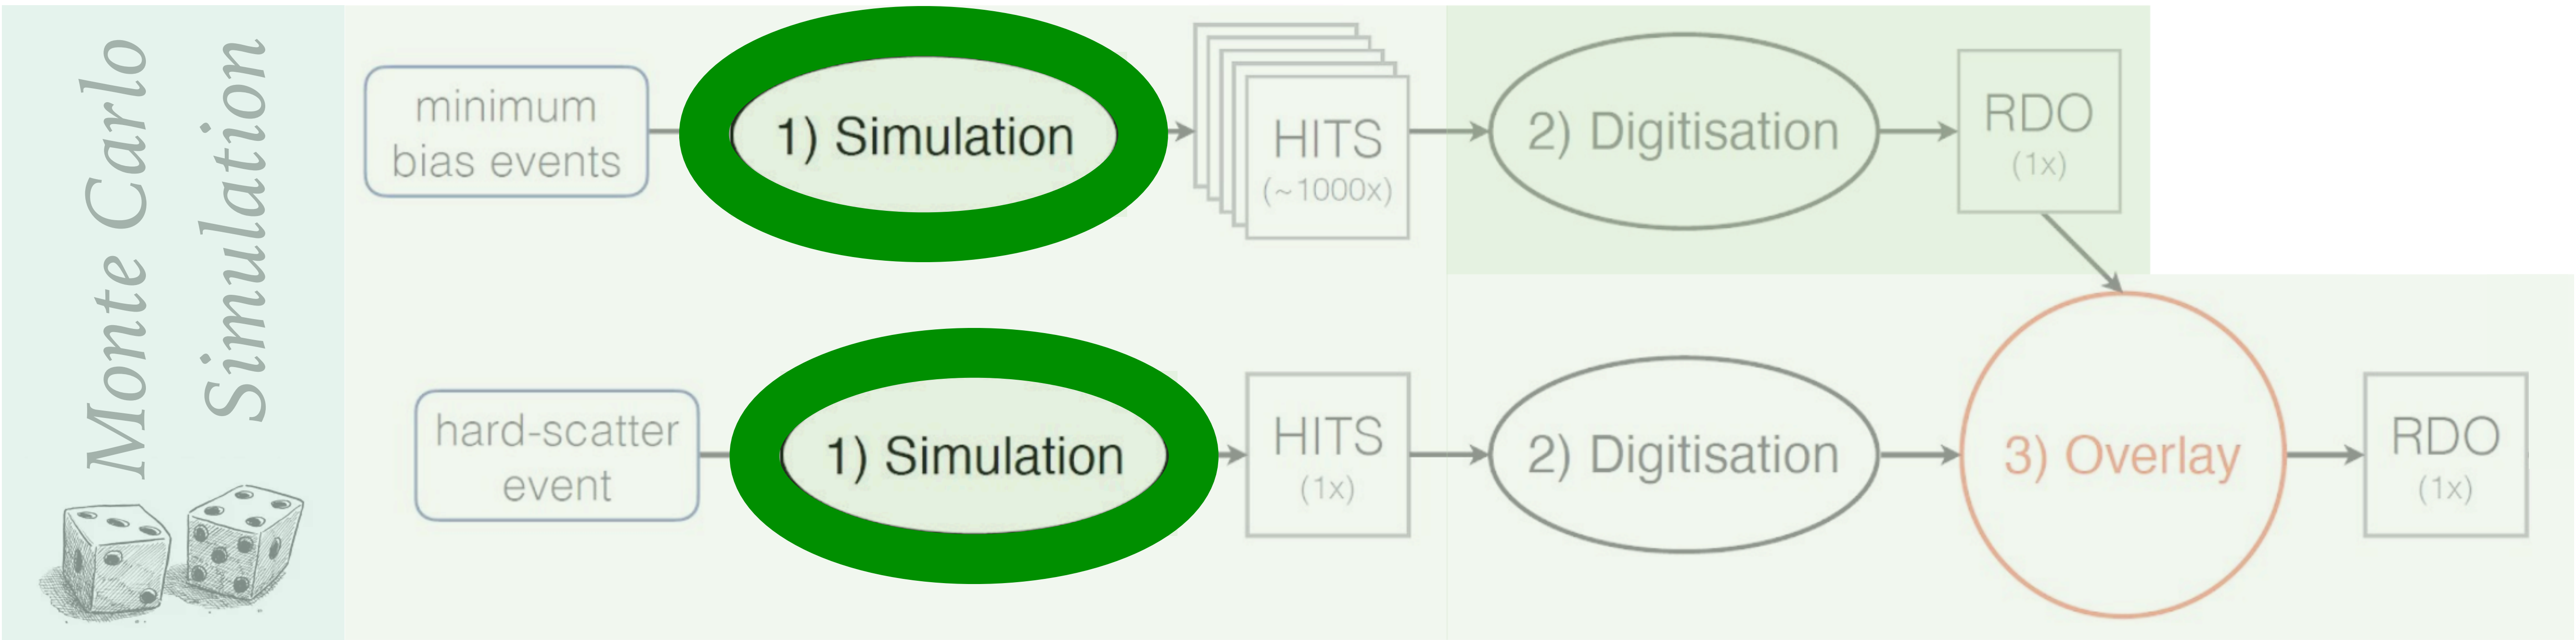

Reconstruction & Analysis

arXiv:2207.00320

ATLAS Detector Simulation

Facts

1. Full detector simulations (aka FullSim - full Geant4 tracking) are **accurate but the largest CPU consumer**
2. FullSim usage is **critical** (CP calibrations, FastSim training, etc.)
3. **LAr EM calorimeters** dominate the simulation load:
   a. low-energy photons from electron scattering
   b. highly-segmented geometry

Geant4 Optimisation Task Force responsible for optimising the performance of the ATLAS Geant4 simulation software

Formed in September 2020 with the target to achieve for Run 3 >30% CPU time speed up w.r.t the comparable Run 2 simulation

ATLAS CPU hours used by various activities in 2018

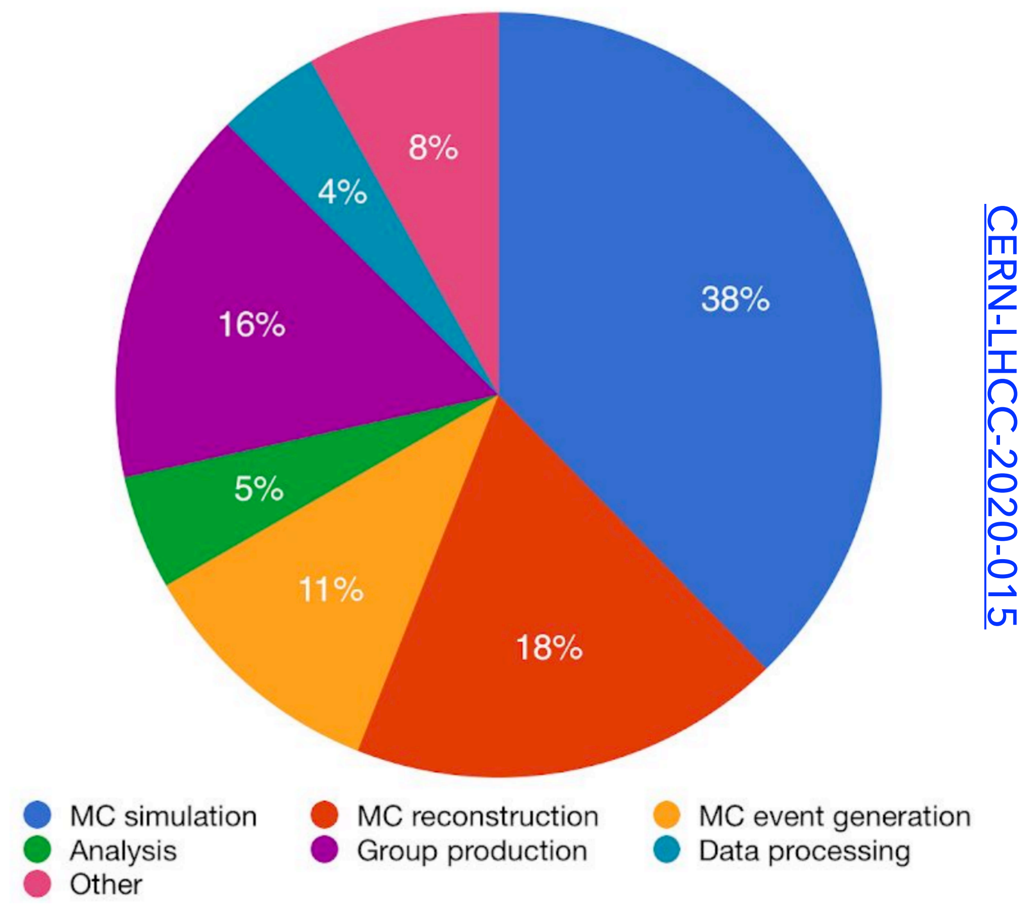

SIM-2019-001

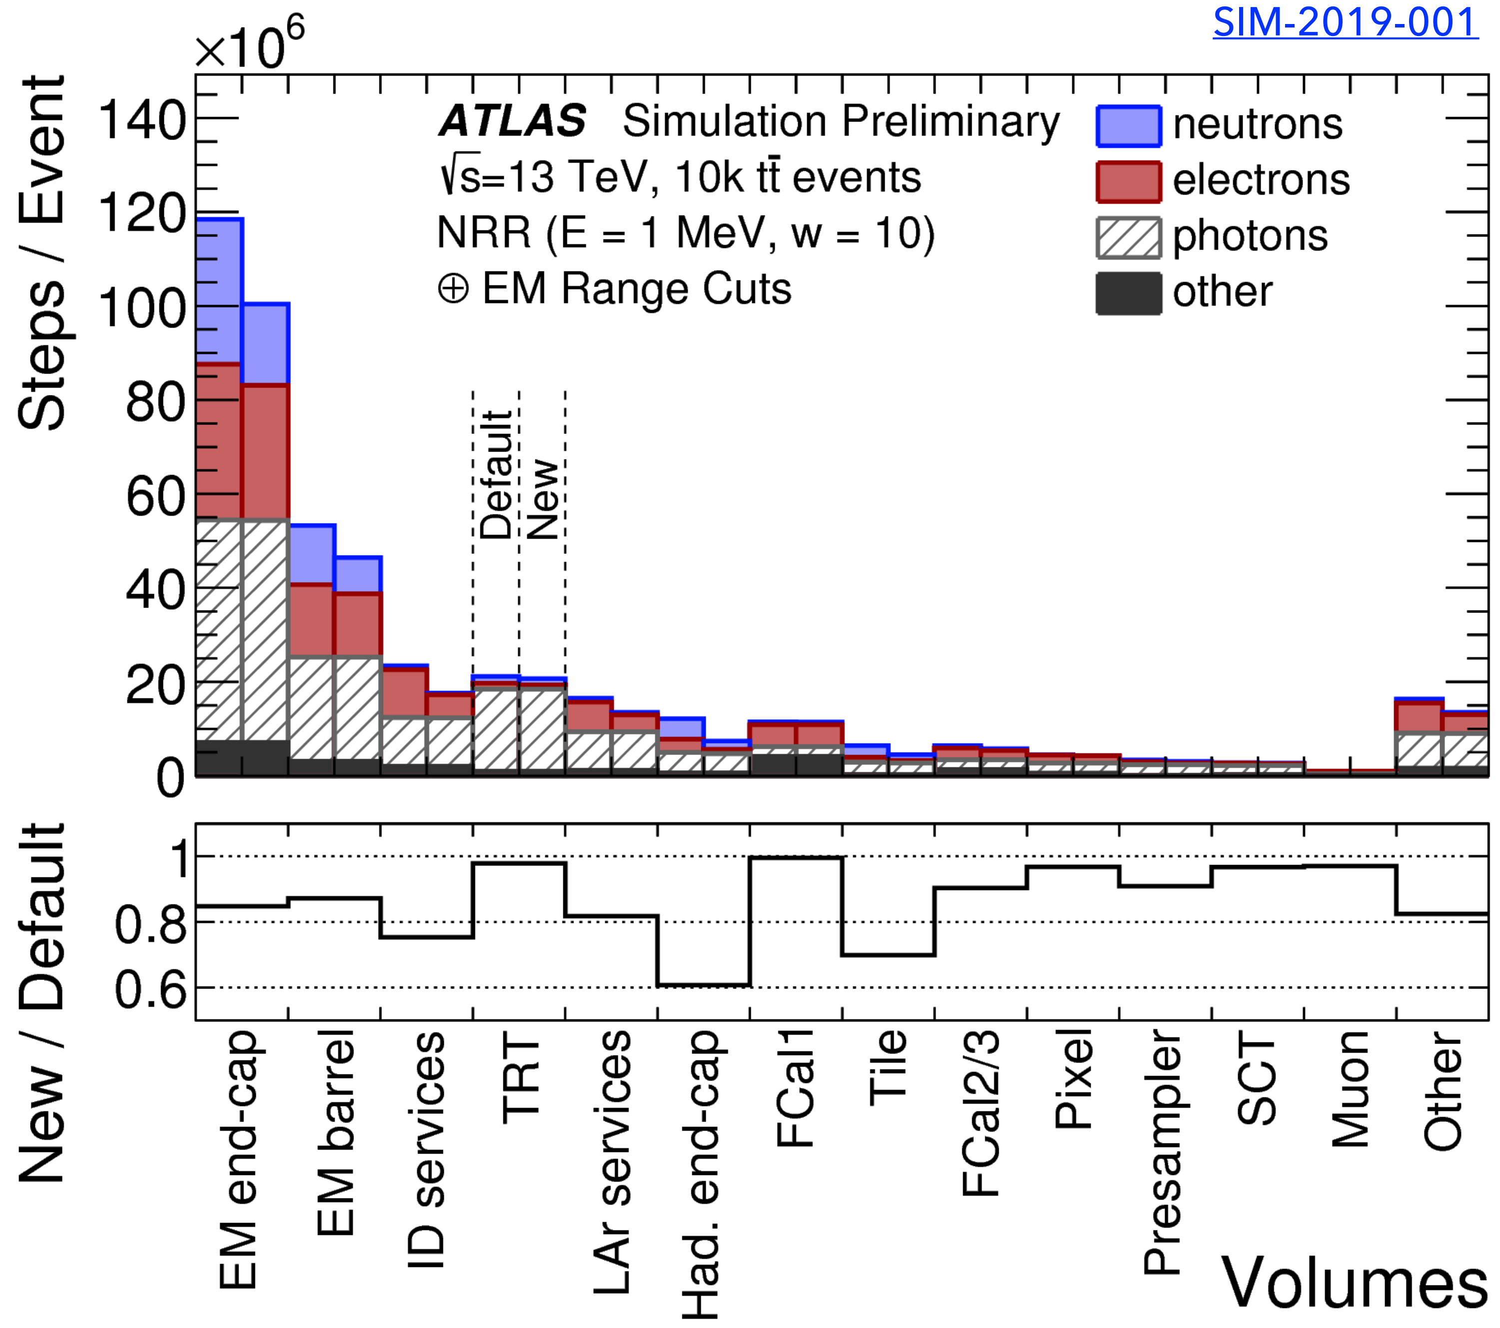

Synopsis

Geant4 Optimizations for Run 3

1. Production-ready
2. Ongoing R&D selection

LHC schedule [link]

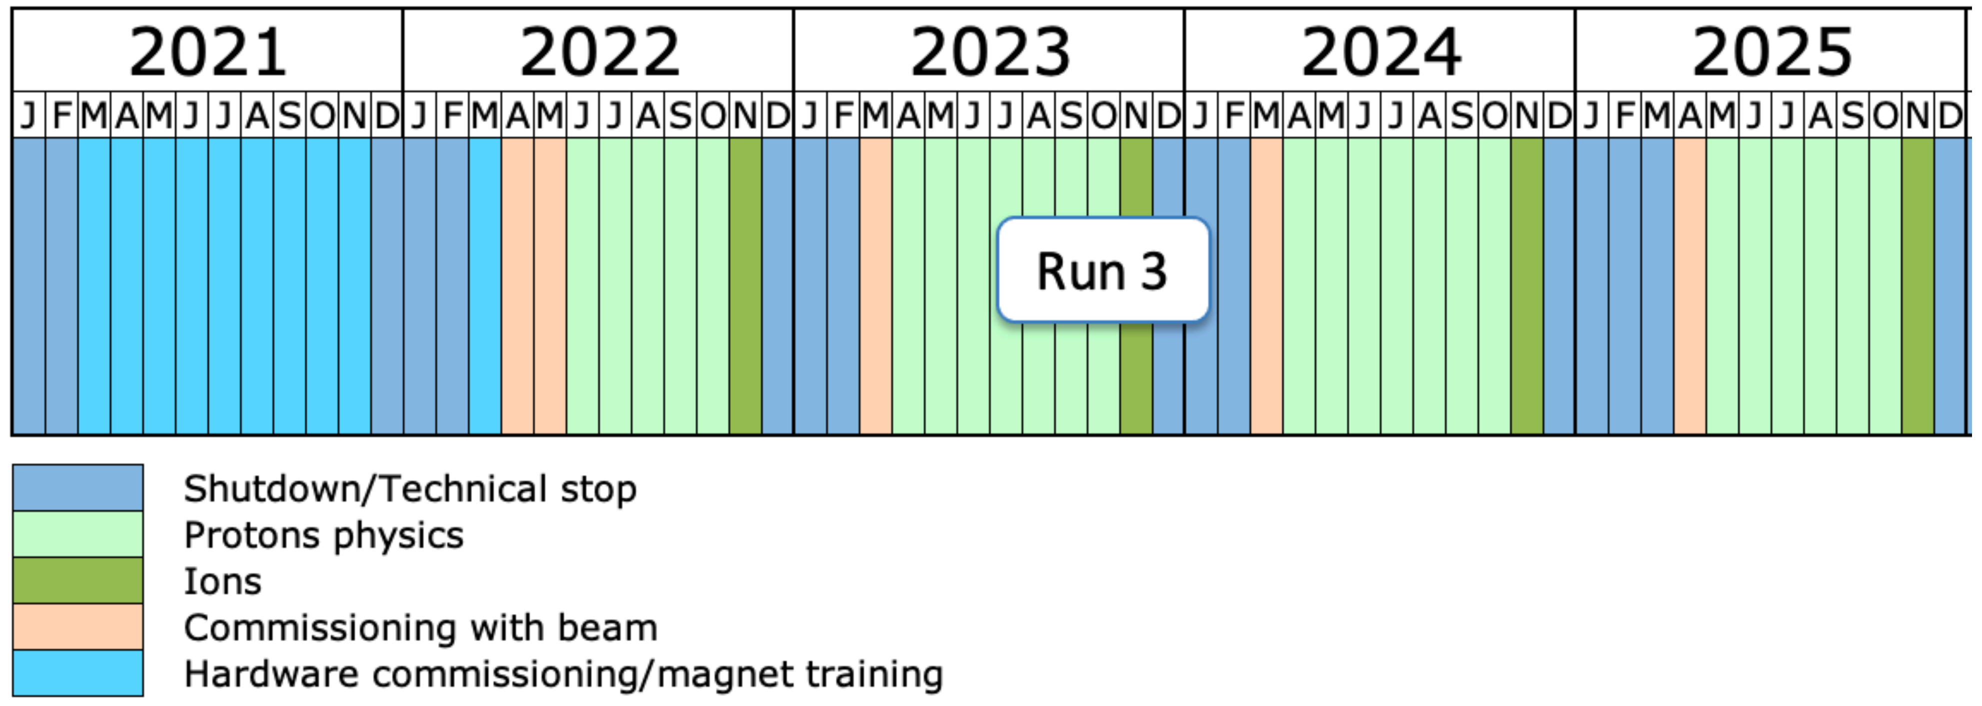

ATLAS Simulation Production Schedule

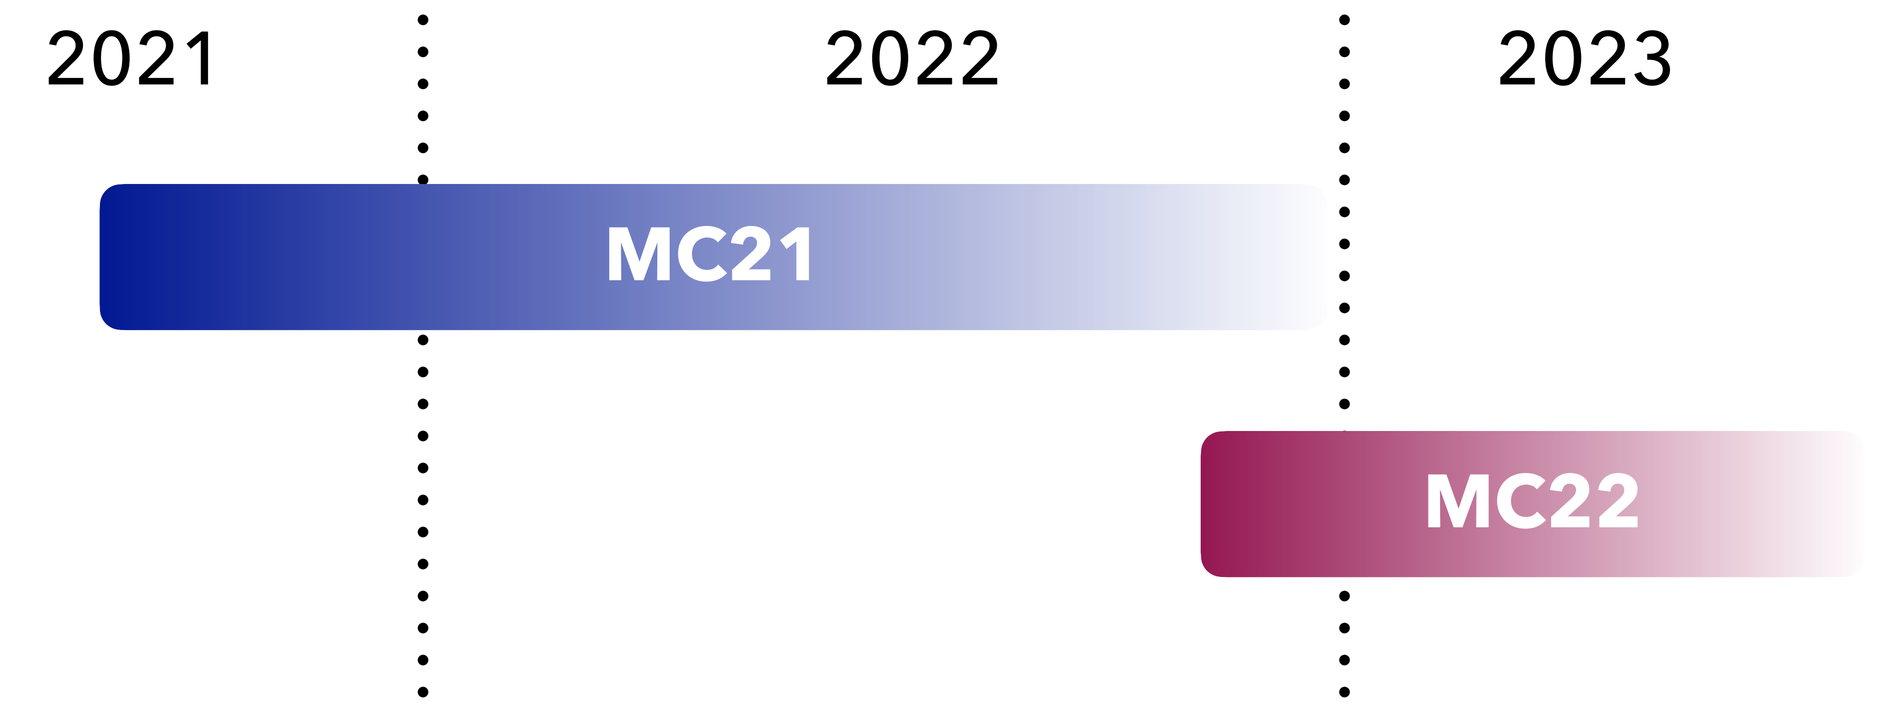

Production-Ready Optimizations

Big Library / Static Linking

Use Geant4 as static library to avoid "trampolines" (Lookup Table delay)

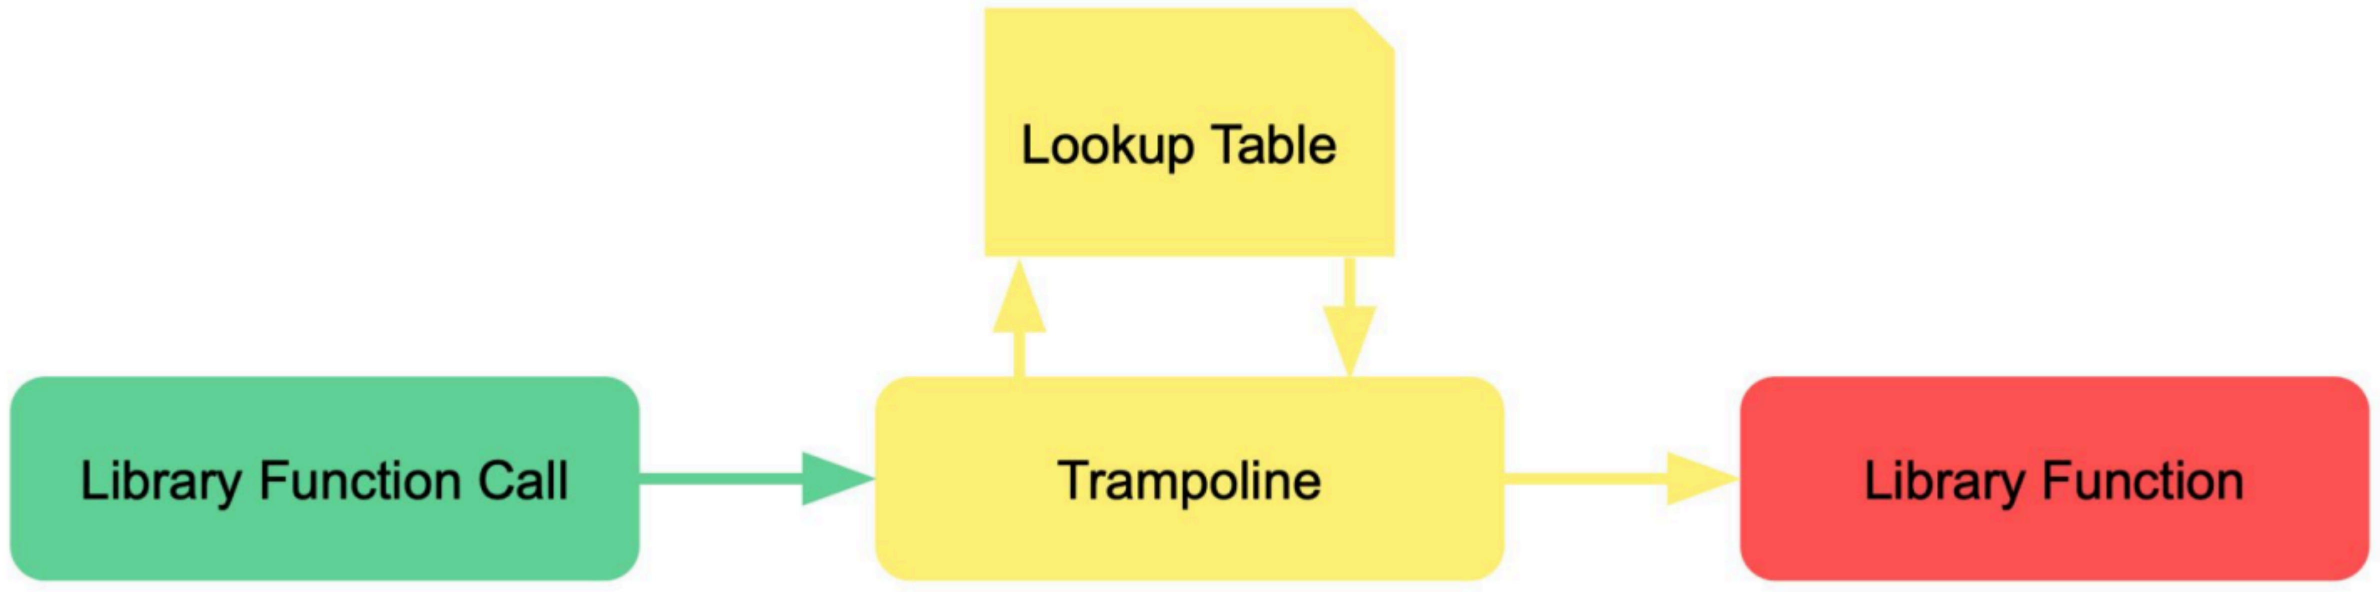

Define a `BigSimulation` shared library, by grouping of all libraries from Athena packages that use Geant4

no physics validation was needed

Benchmark and *Athena**

showed ~7% speed up

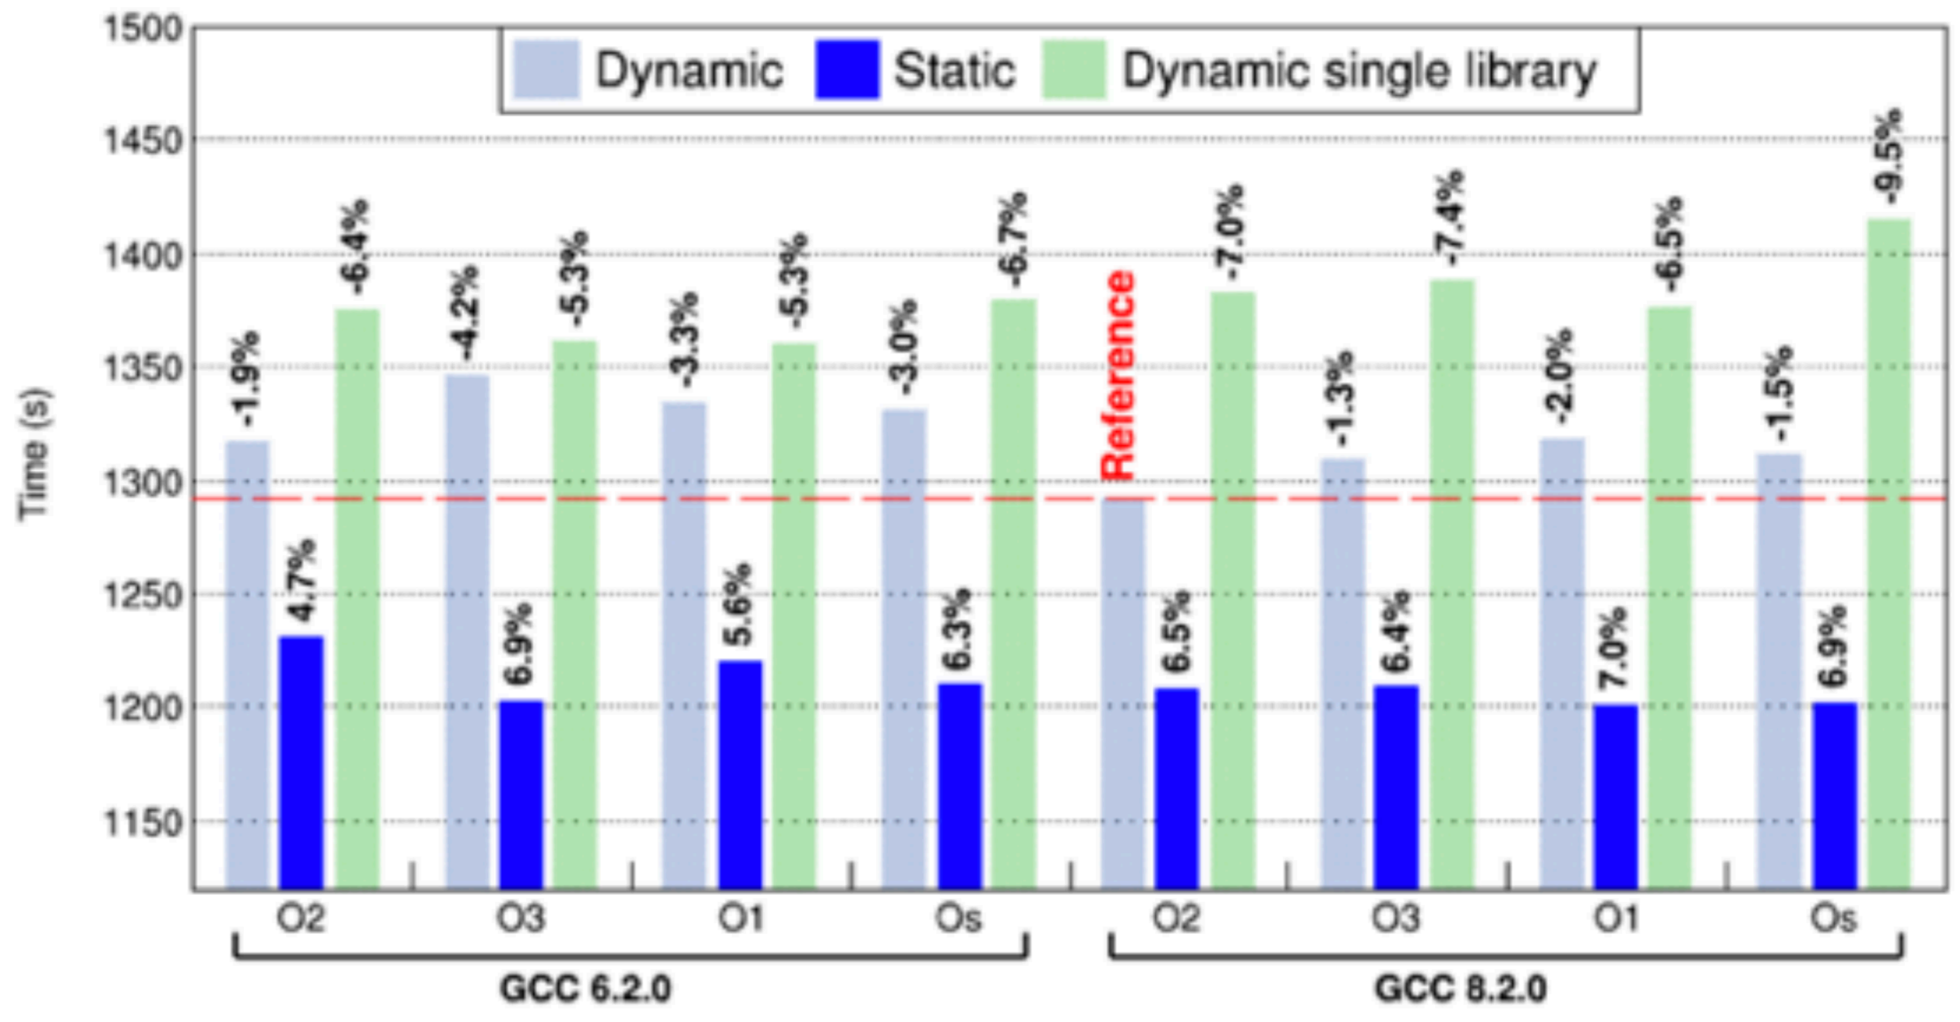

[HepExpMT](#) *stand-alone Geant4 benchmark*

Athena: *The ATLAS Experiment's main offline software*

Reducing Operations

Gamma General Process

Use only one, collective physics process for photons → reduced number of instructions/ calculations on geometry boundary crossings

~4.5% speed up

on 100 ttbar events in Athena

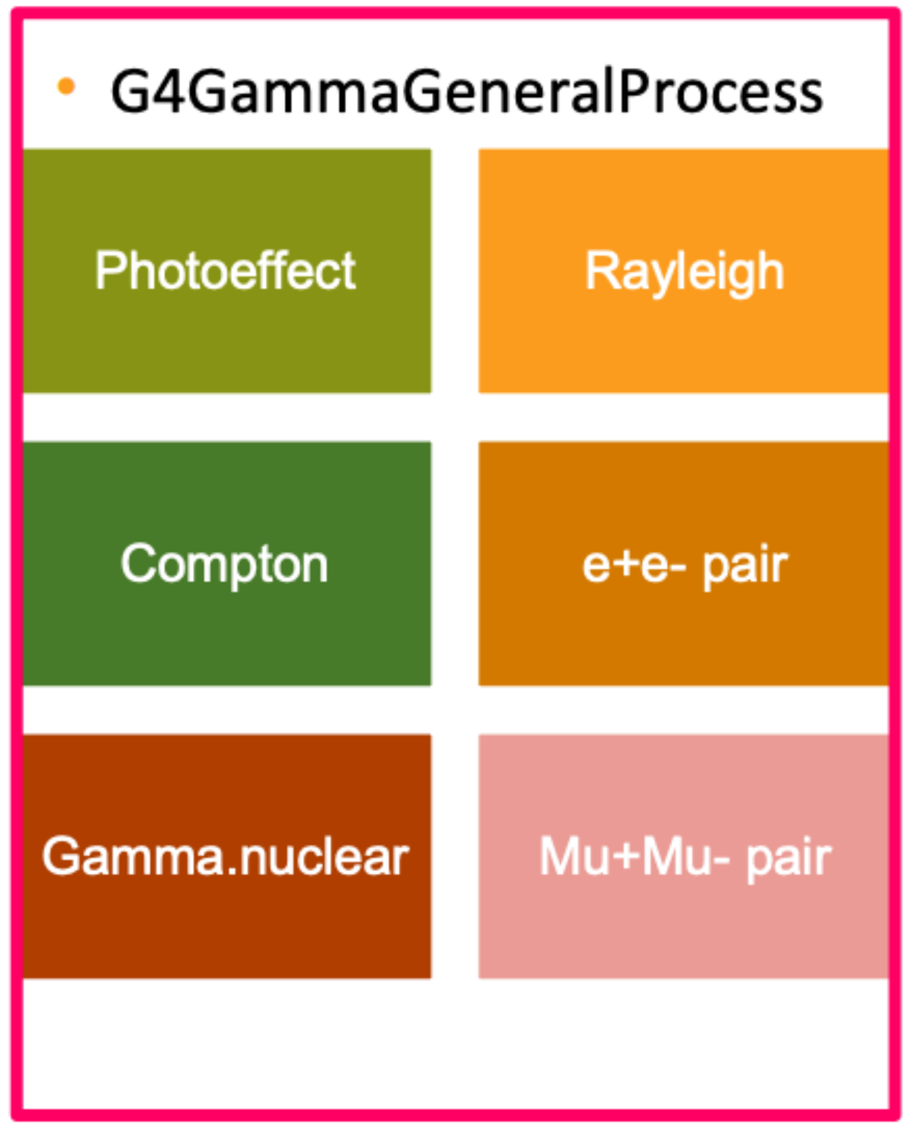

Magnetic Field Tailored Switch-OFF

Speed up observed when switching-off magnetic field in LAr calorimeter (except for muons) without affecting shower shapes

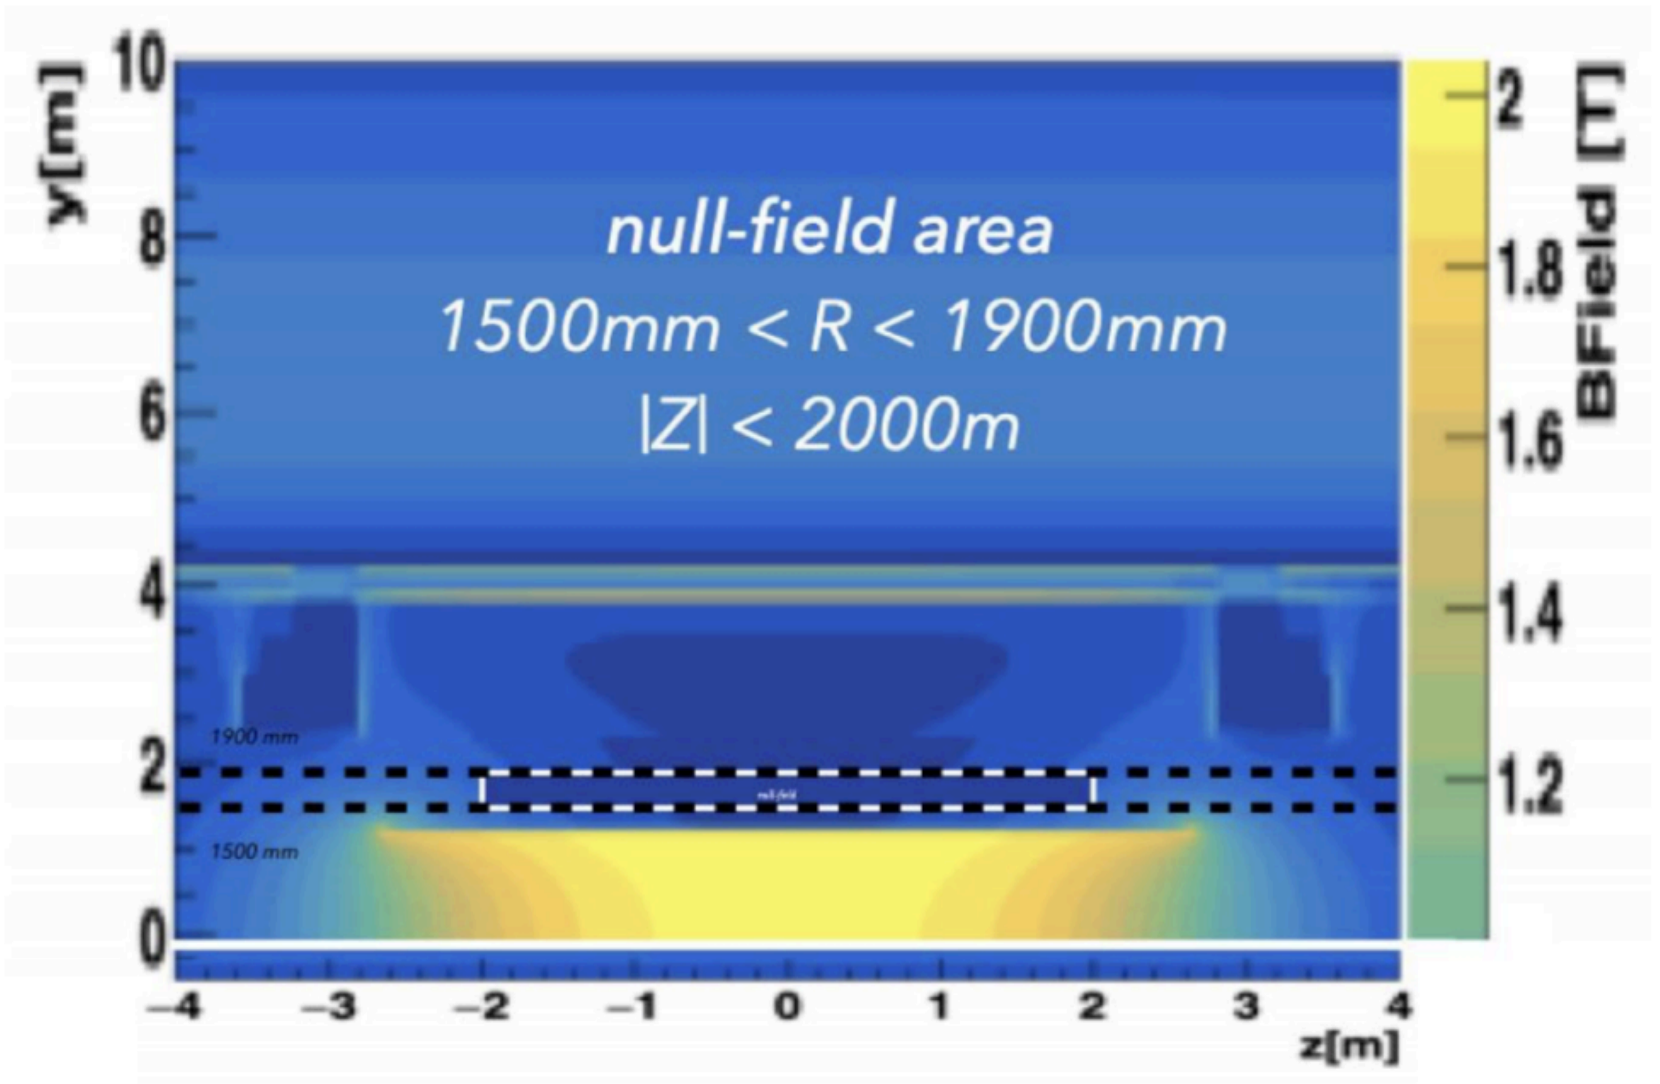

~3% speed up on 200 ttbar events in Athena

Possibility to extend solution to other detector regions too

Accelerating Geometries

EMEC Custom Solid

Described by a custom Geant4 solid using `G4Polycone` for internal calculations.
Re-Implemented custom solid variants:

- Wheel: the default with `G4Polycone`
- Cone: improved shape using `G4ShiftedCone` – outer wheel divided into two conical-shaped sections
- Slices: new `LArWheelSliceSolid` – each wheel is divided into many thick slices along Z axis

Slices variant showed **5-6% speed up**

Cone ***Slices***

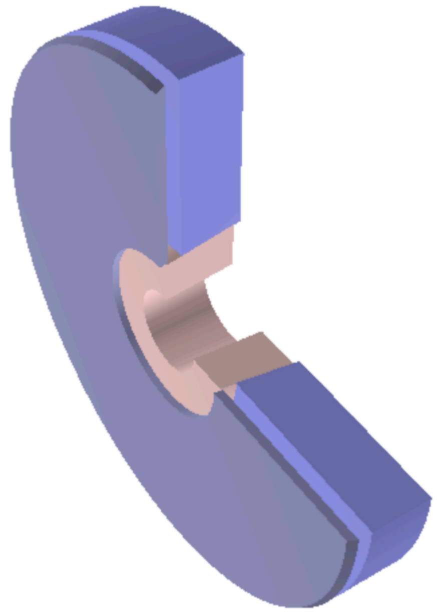

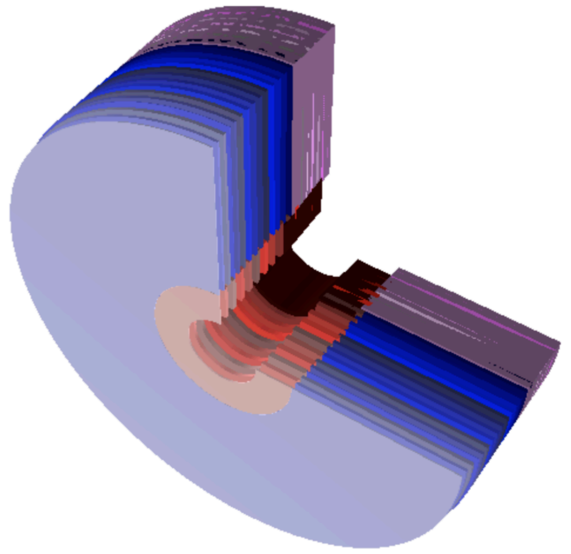

VecGeom Integration

A geometry modeller/navigation library providing fast geometry primitives as well as vectorised (SIMD) navigation.

enabled in Athena only for `G4Cons` & `G4Polycones`

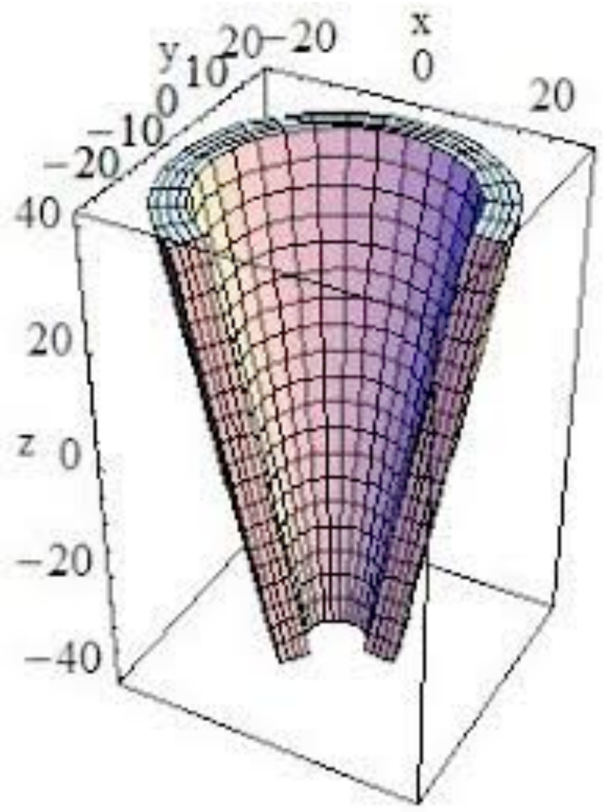

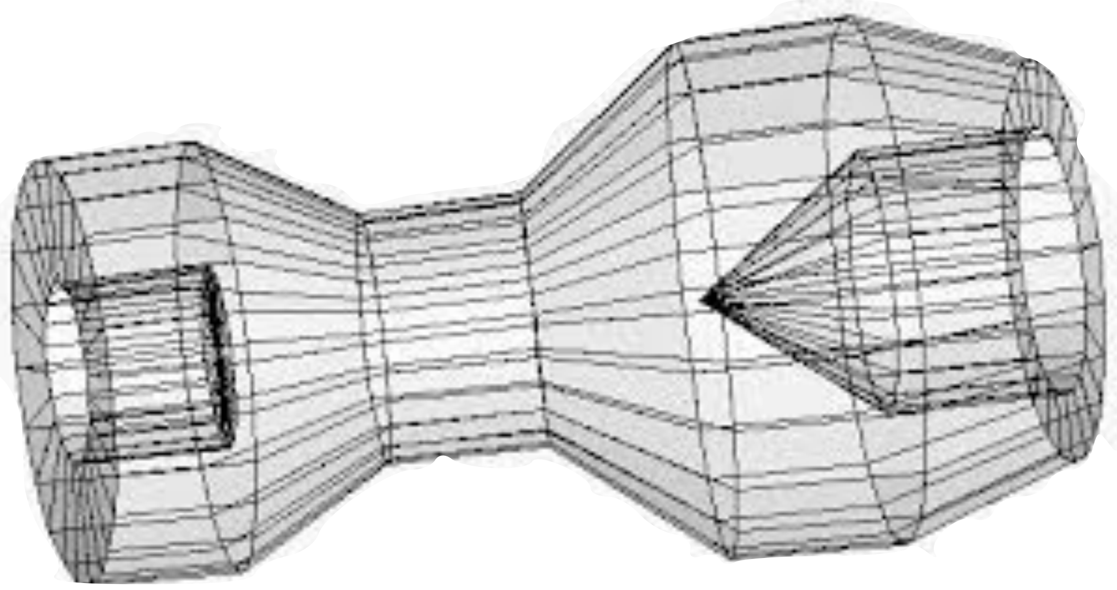

3 - 4.5% speed up on top of the rest of Run3Opt

Russian Roulettes & EM Range Cuts

Russian Roulette

- Neutrons and photons take majority of CPU time
  EM calorimeters most resource intensive
- Photon/Neutron Russian Roulette (PRR/NRR): **randomly discard particles below energy threshold and weight the energy deposits of remaining particles accordingly**
- NRR performance: 10% speed up with 2 MeV threshold for neutrons

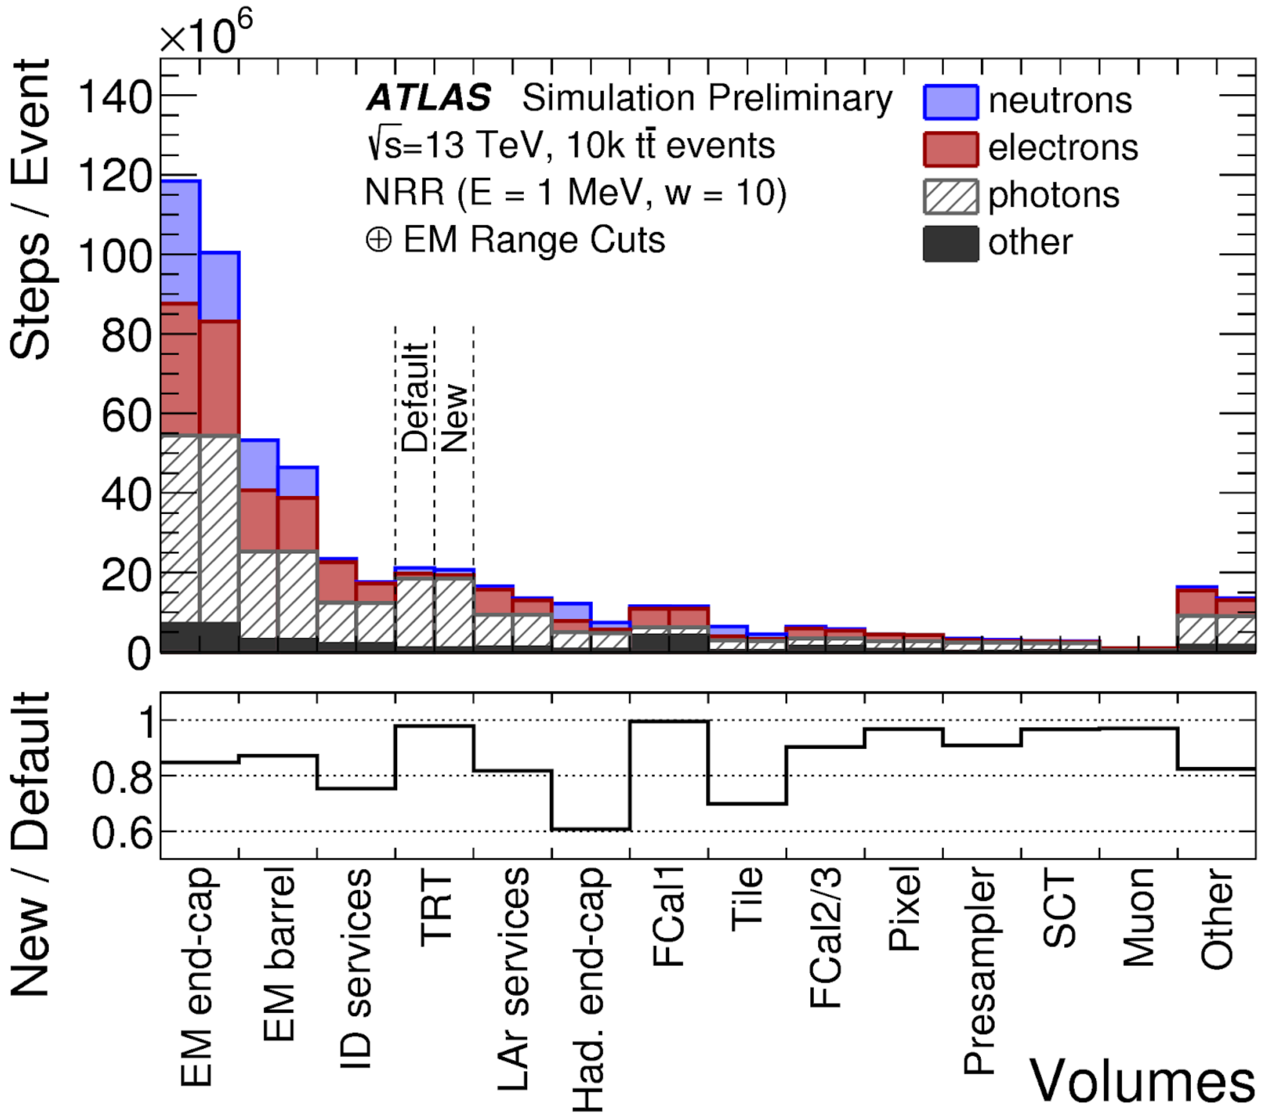

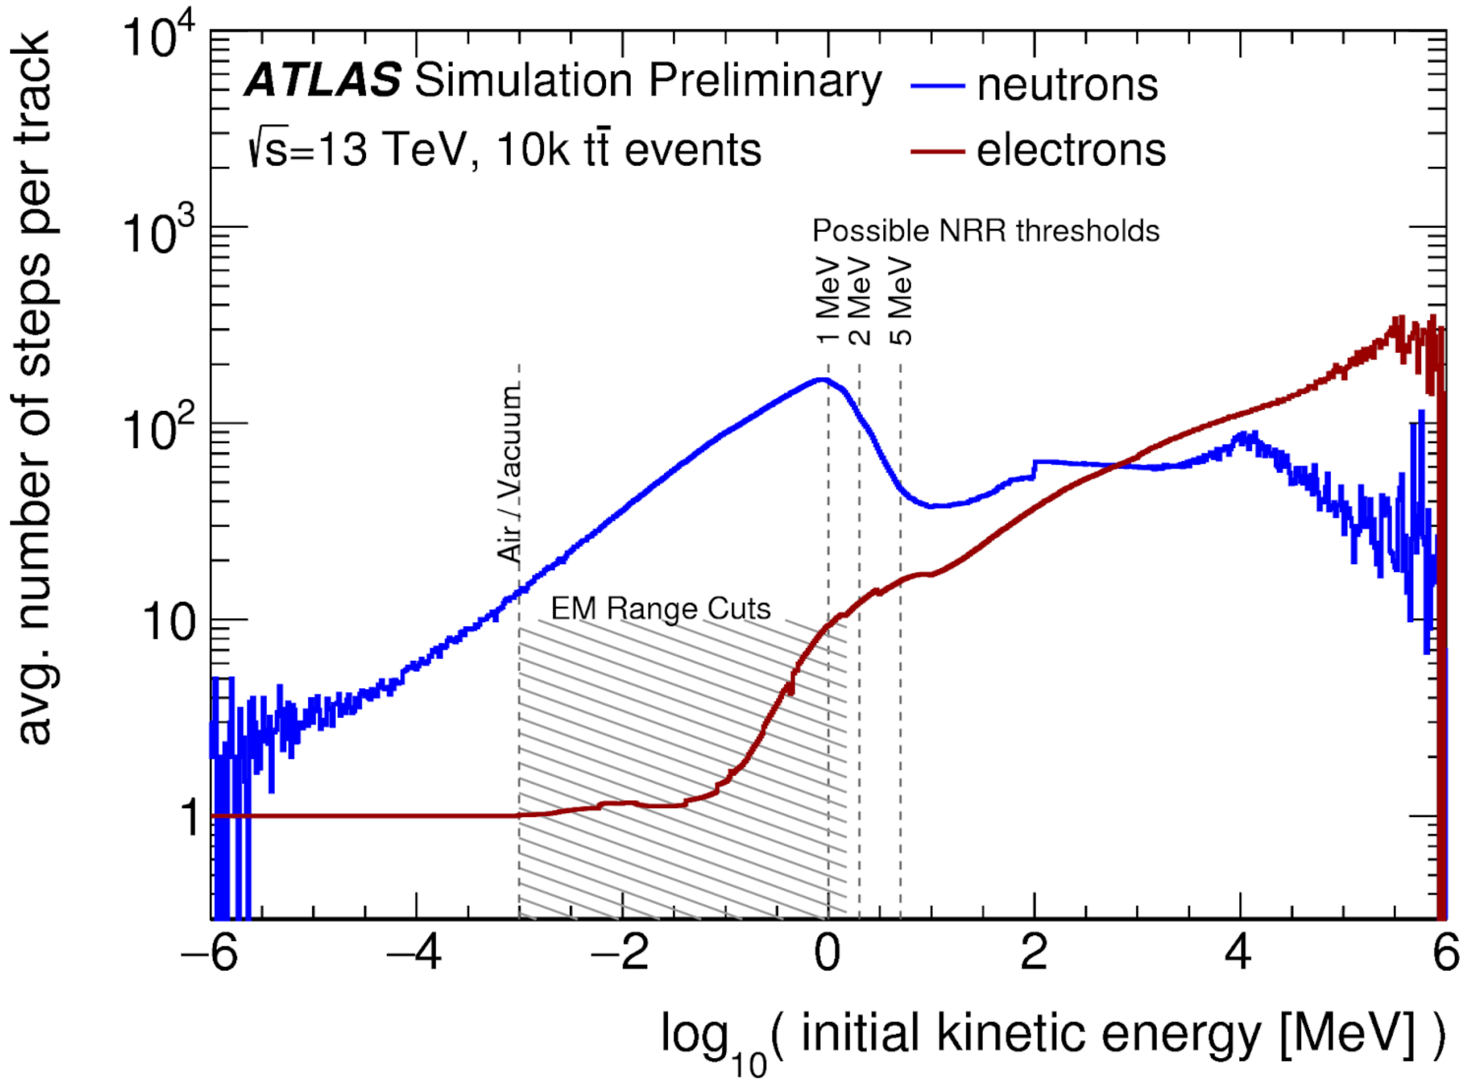

EM Range Cuts

- **Avoid creating secondaries if their energy is below a certain value**
- OFF by default for three processes: **Compton, conversion, photo-electric effect**
- Turning them on provides ~6-7% speed up with negligible impact on physics

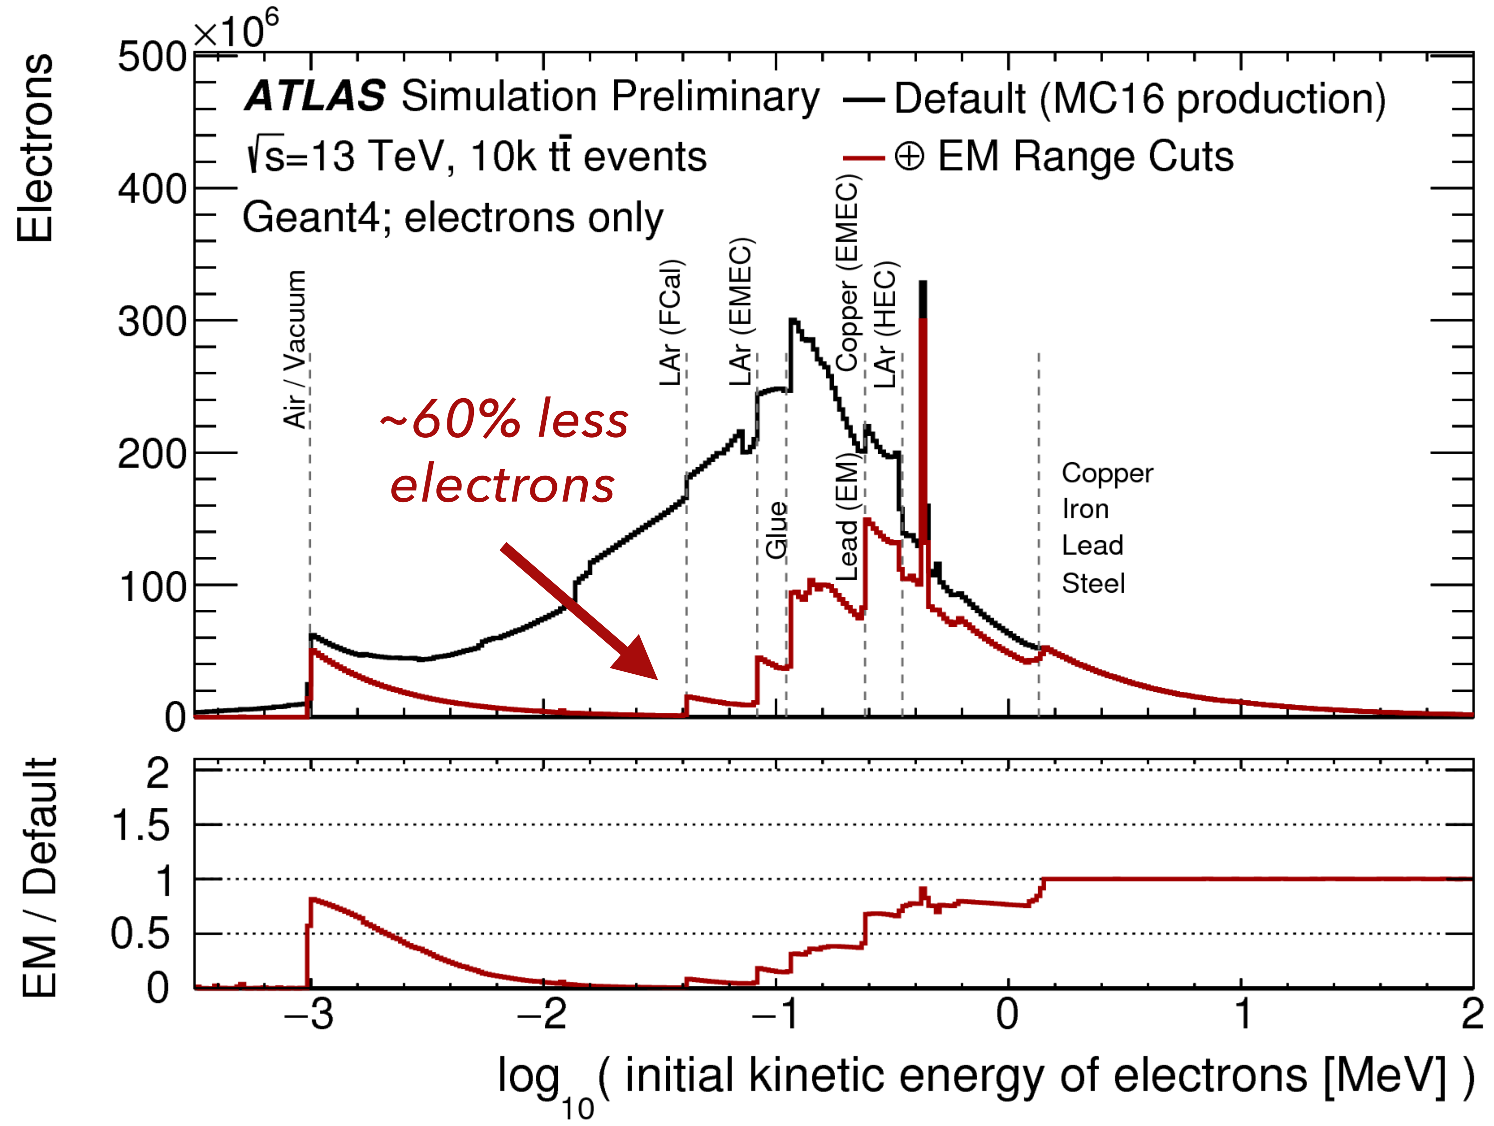

SIM-2019-001

Geant4 Optimizations Speedup Timeline

Current optimisations into production surpass the target performance improvement

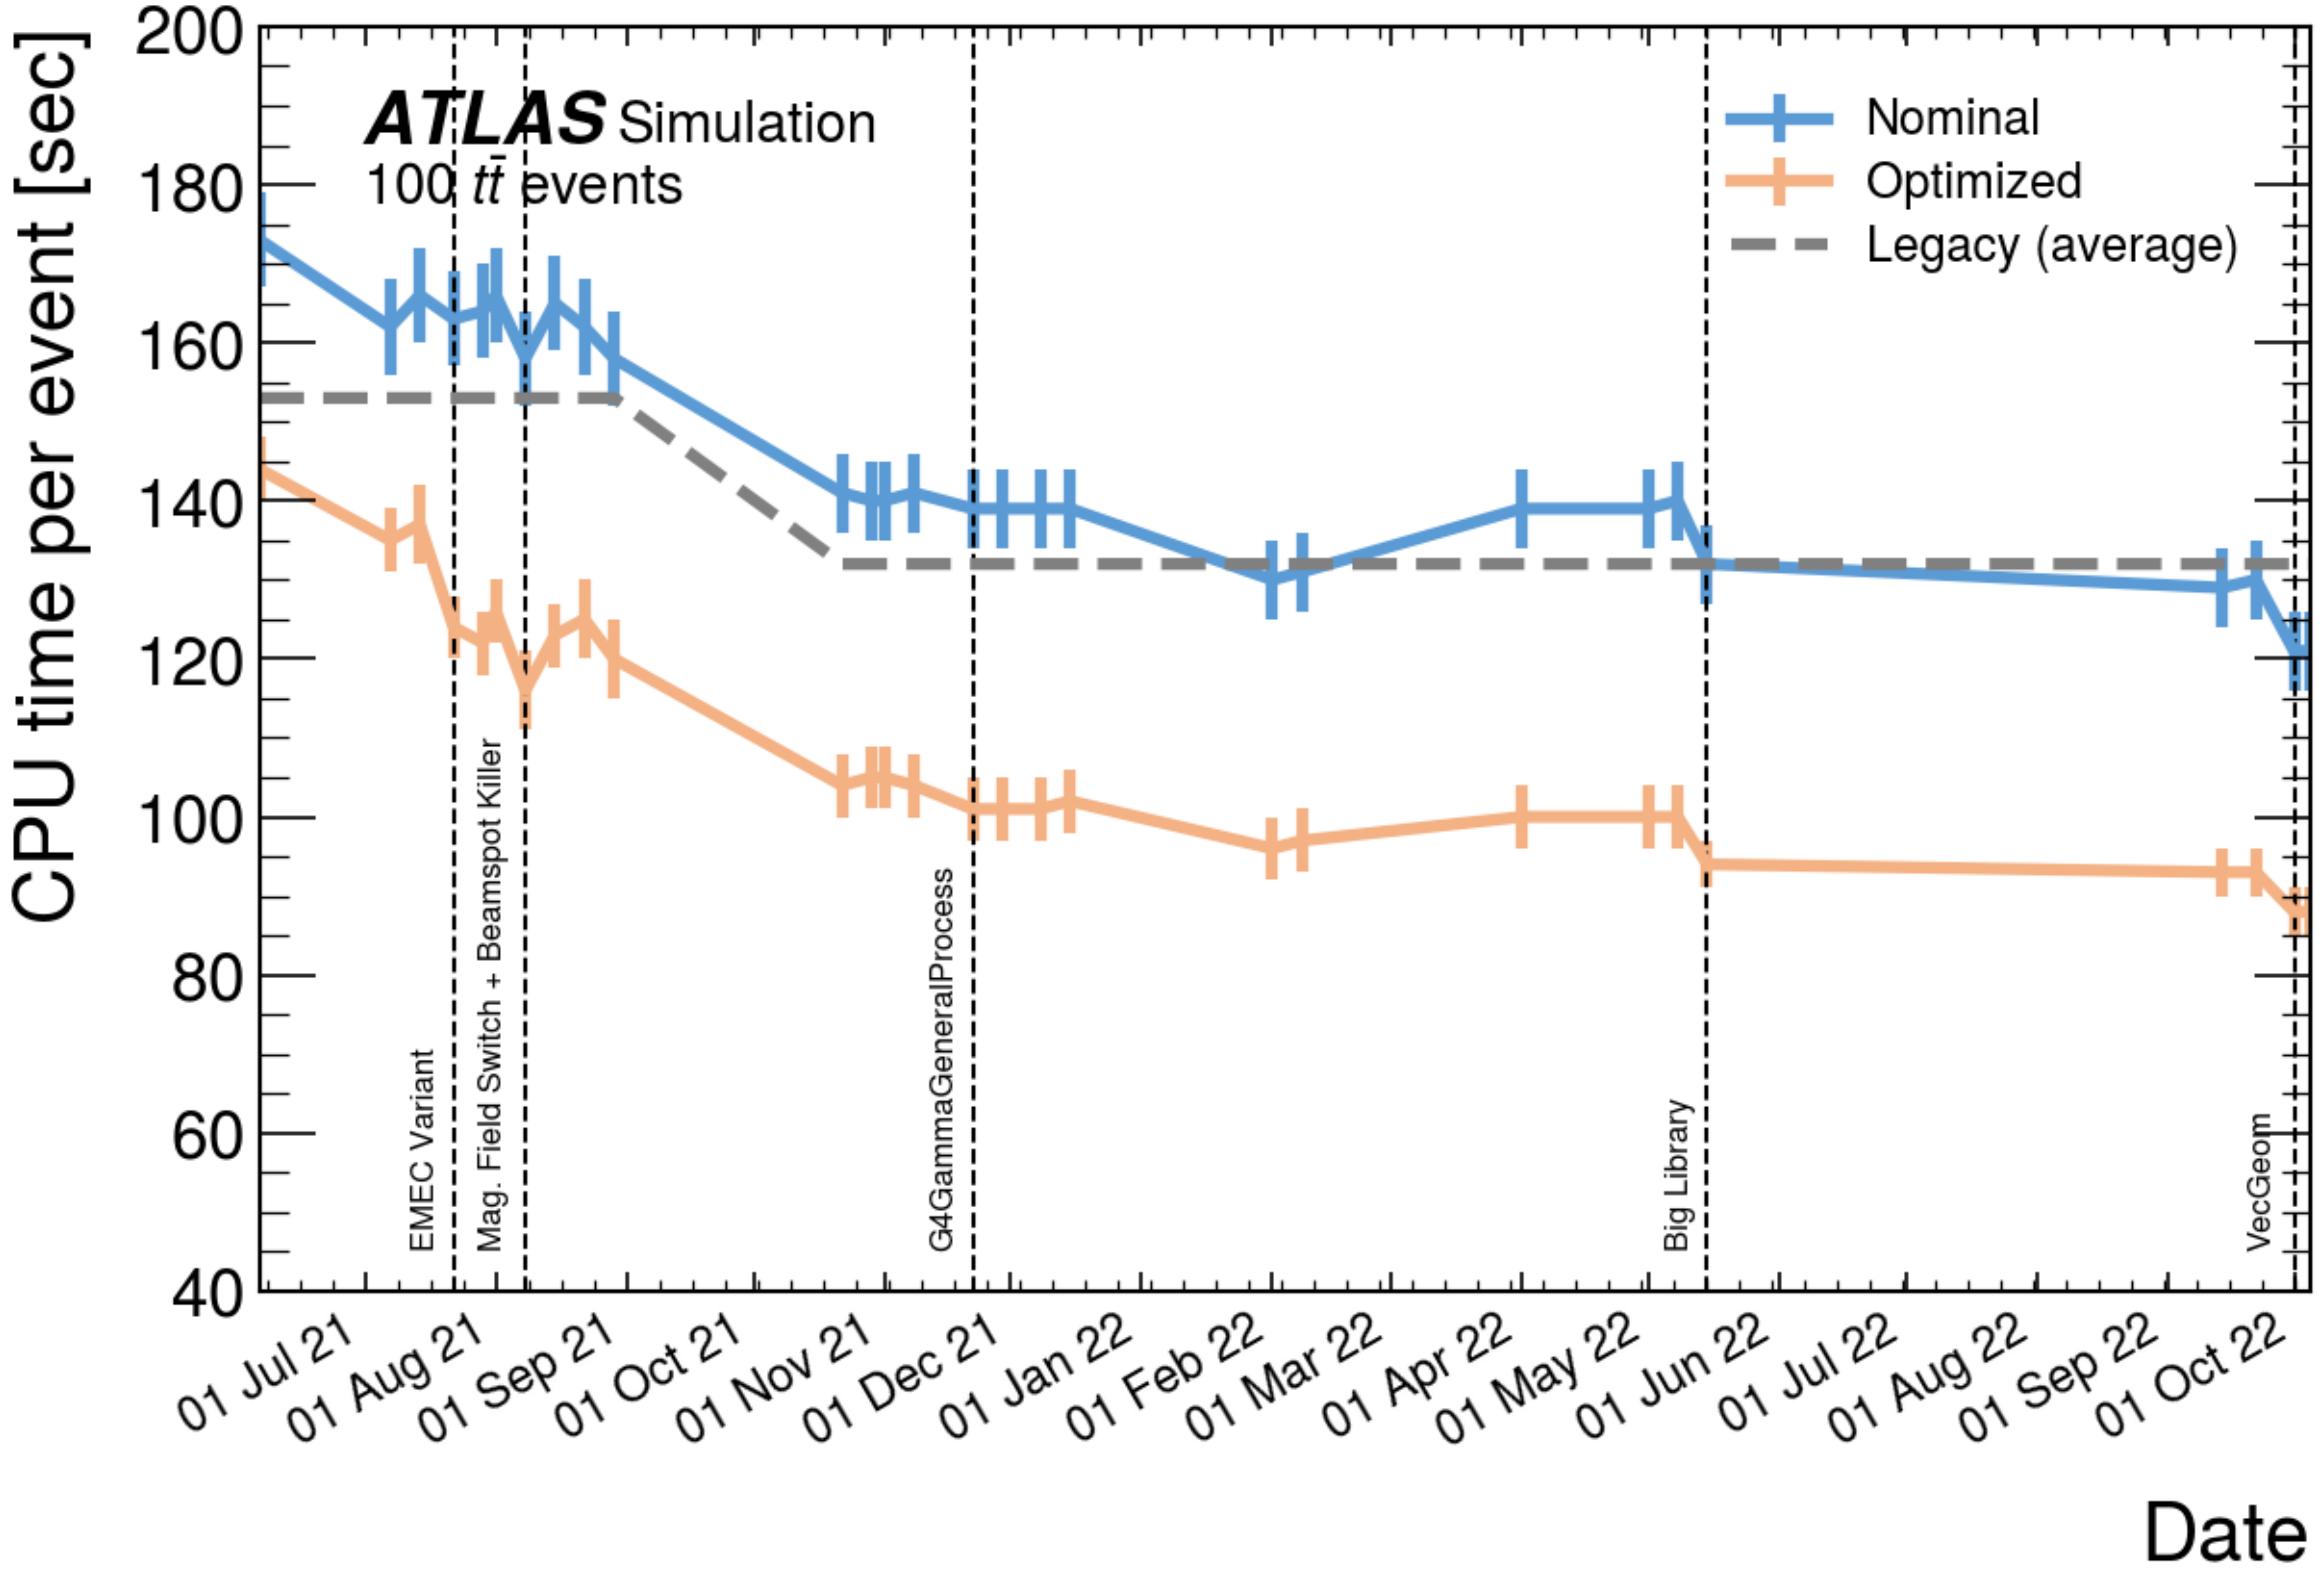

Disclaimer: optimizations are correlated - improvement does not add up linearly

wrt Legacy

33% CPU Speedup

50% Higher Throughput

37% Higher Throughput

27% CPU Speedup

wrt Nominal

Realistic Benchmarks

Improvement is also observed in "realistic" production conditions – LWCG sites

Brookhaven Natl. Lab (WLCG Tier 1)

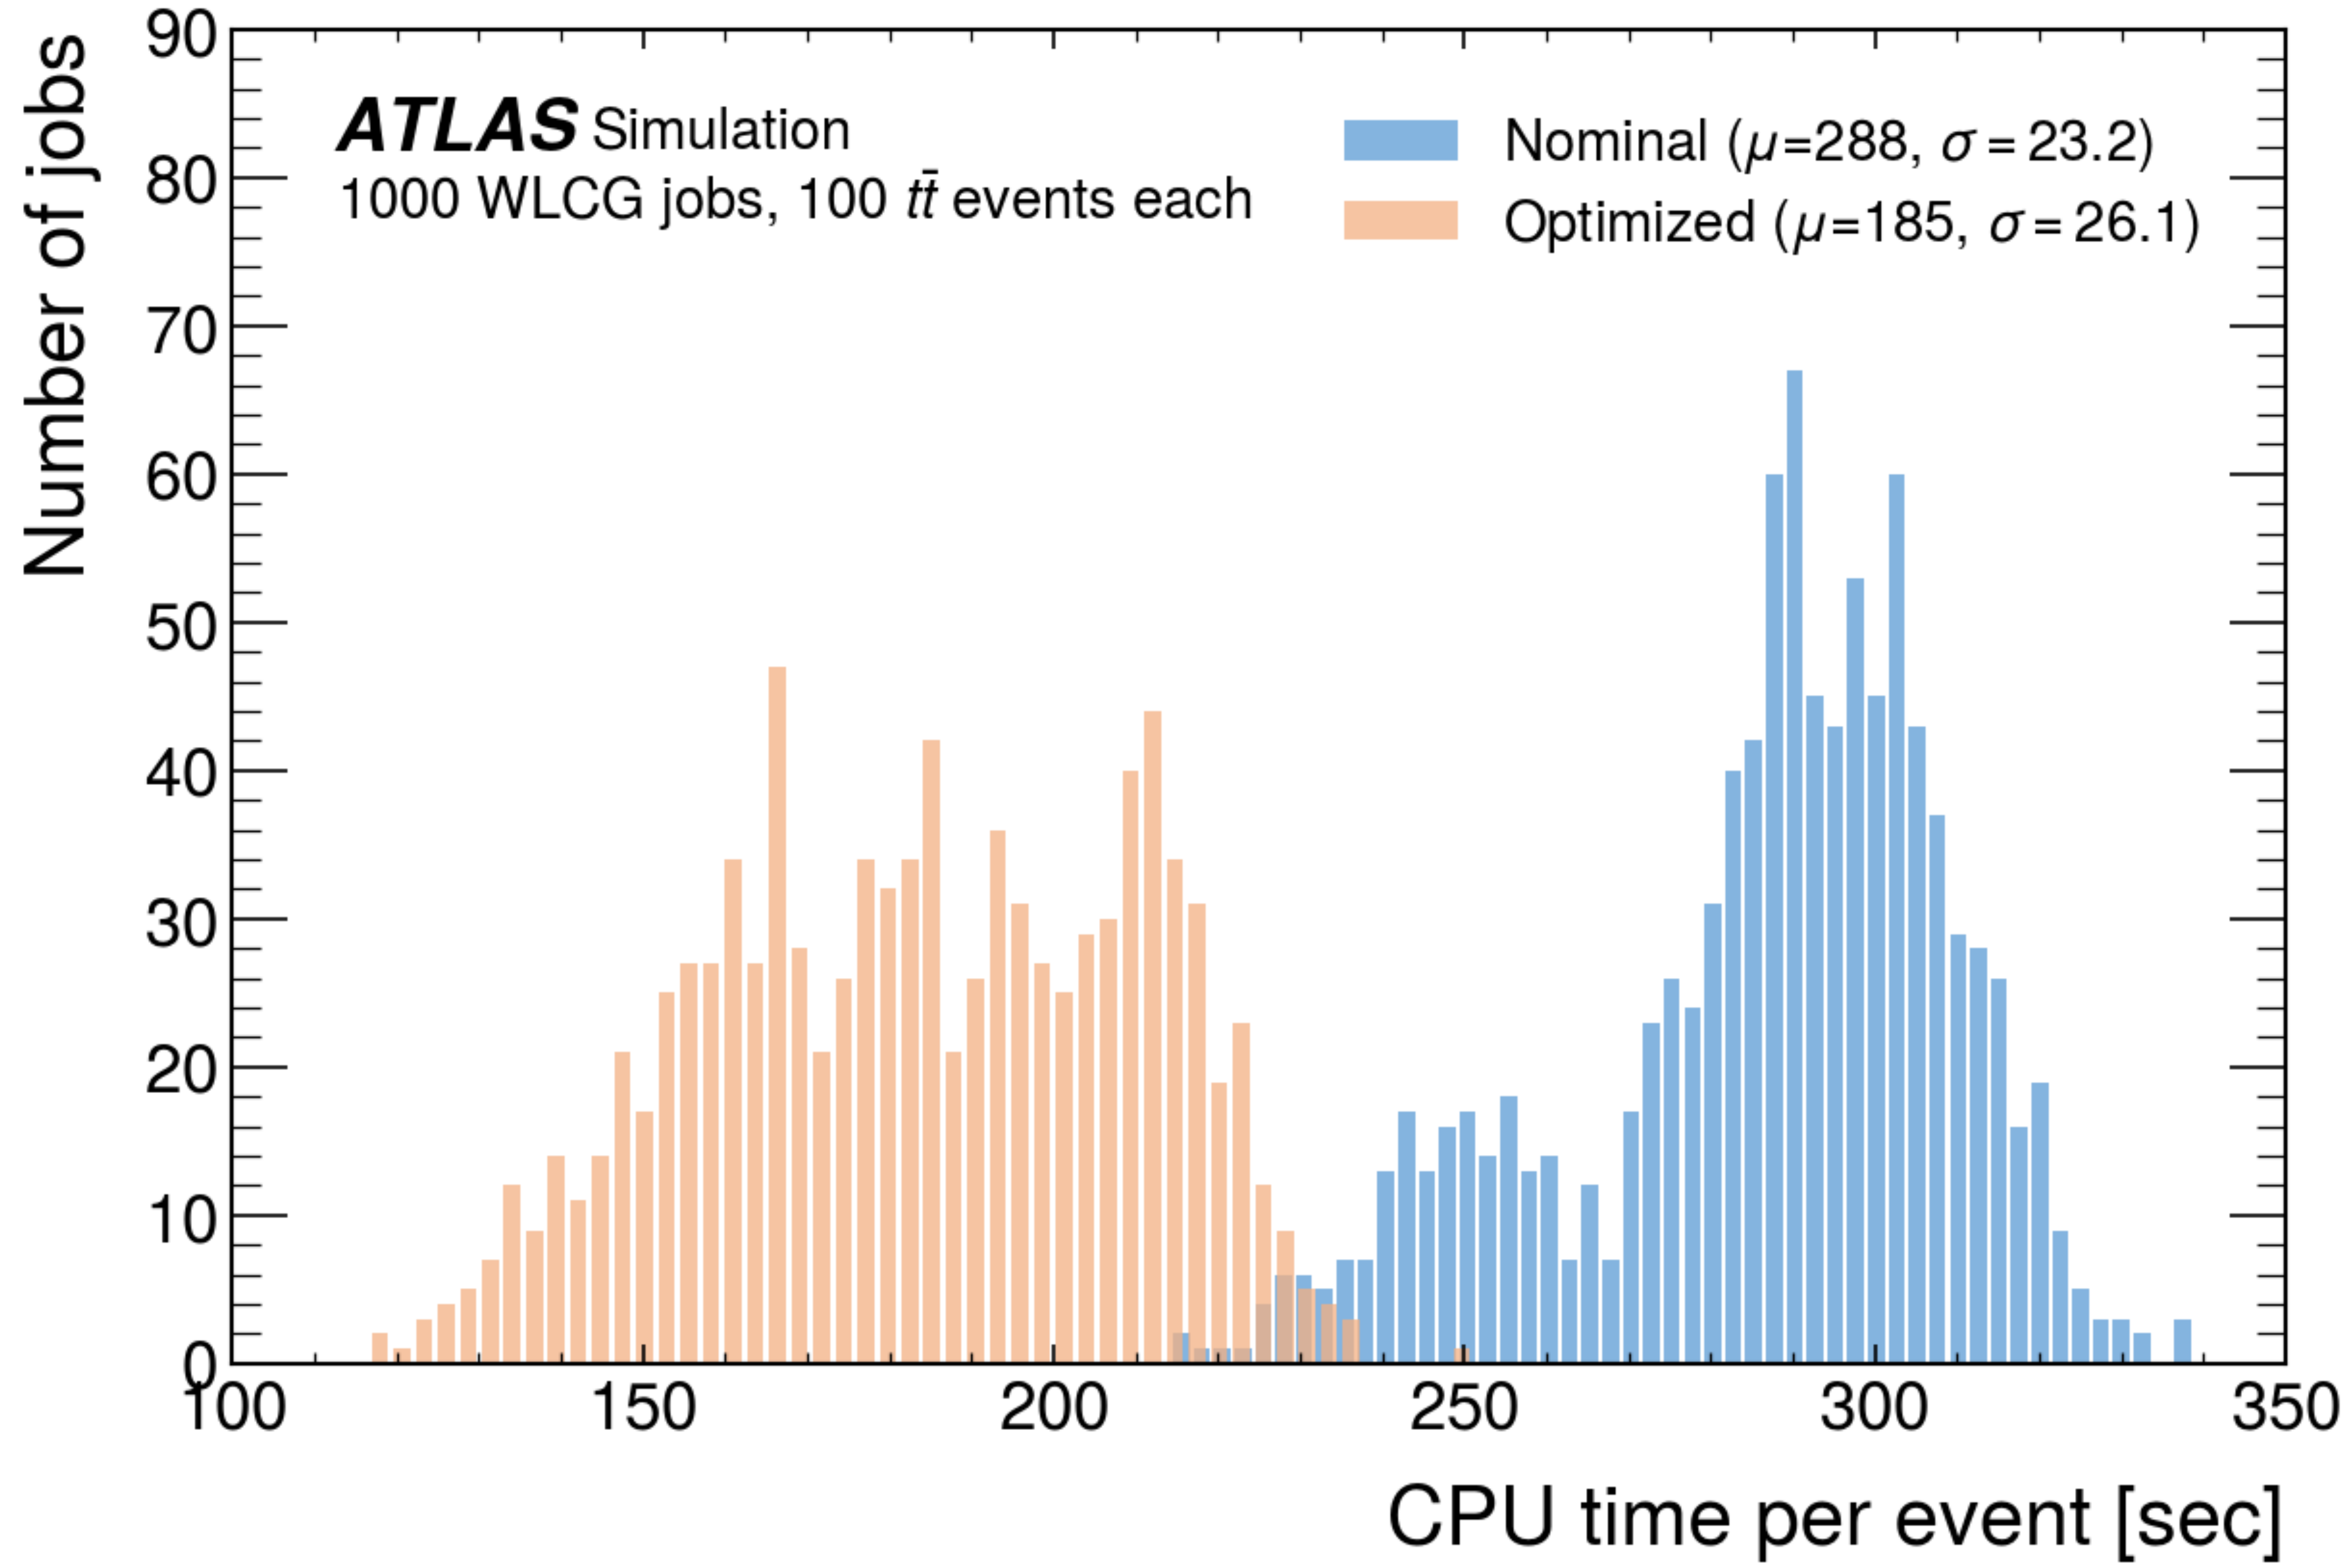

	CPU Time [sec]	σ
Nominal	288	23.2
Optimized	185	26.1

36% CPU Speedup

56% Higher Throughput

*Big Library optimization is missing

Ongoing R&D

Voxel Density Tuning

Geant4 tracking is optimized by voxelization

The density of the voxels can be tuned

- **Goal:** Optimize the voxel density value for a balance between memory usage for the detector description and CPU time for tracking

- No significant effect on simulation accuracy is expected – to be validated

Optimal values show a preliminary:

- ~40% memory usage reduction
- for the ~same CPU time

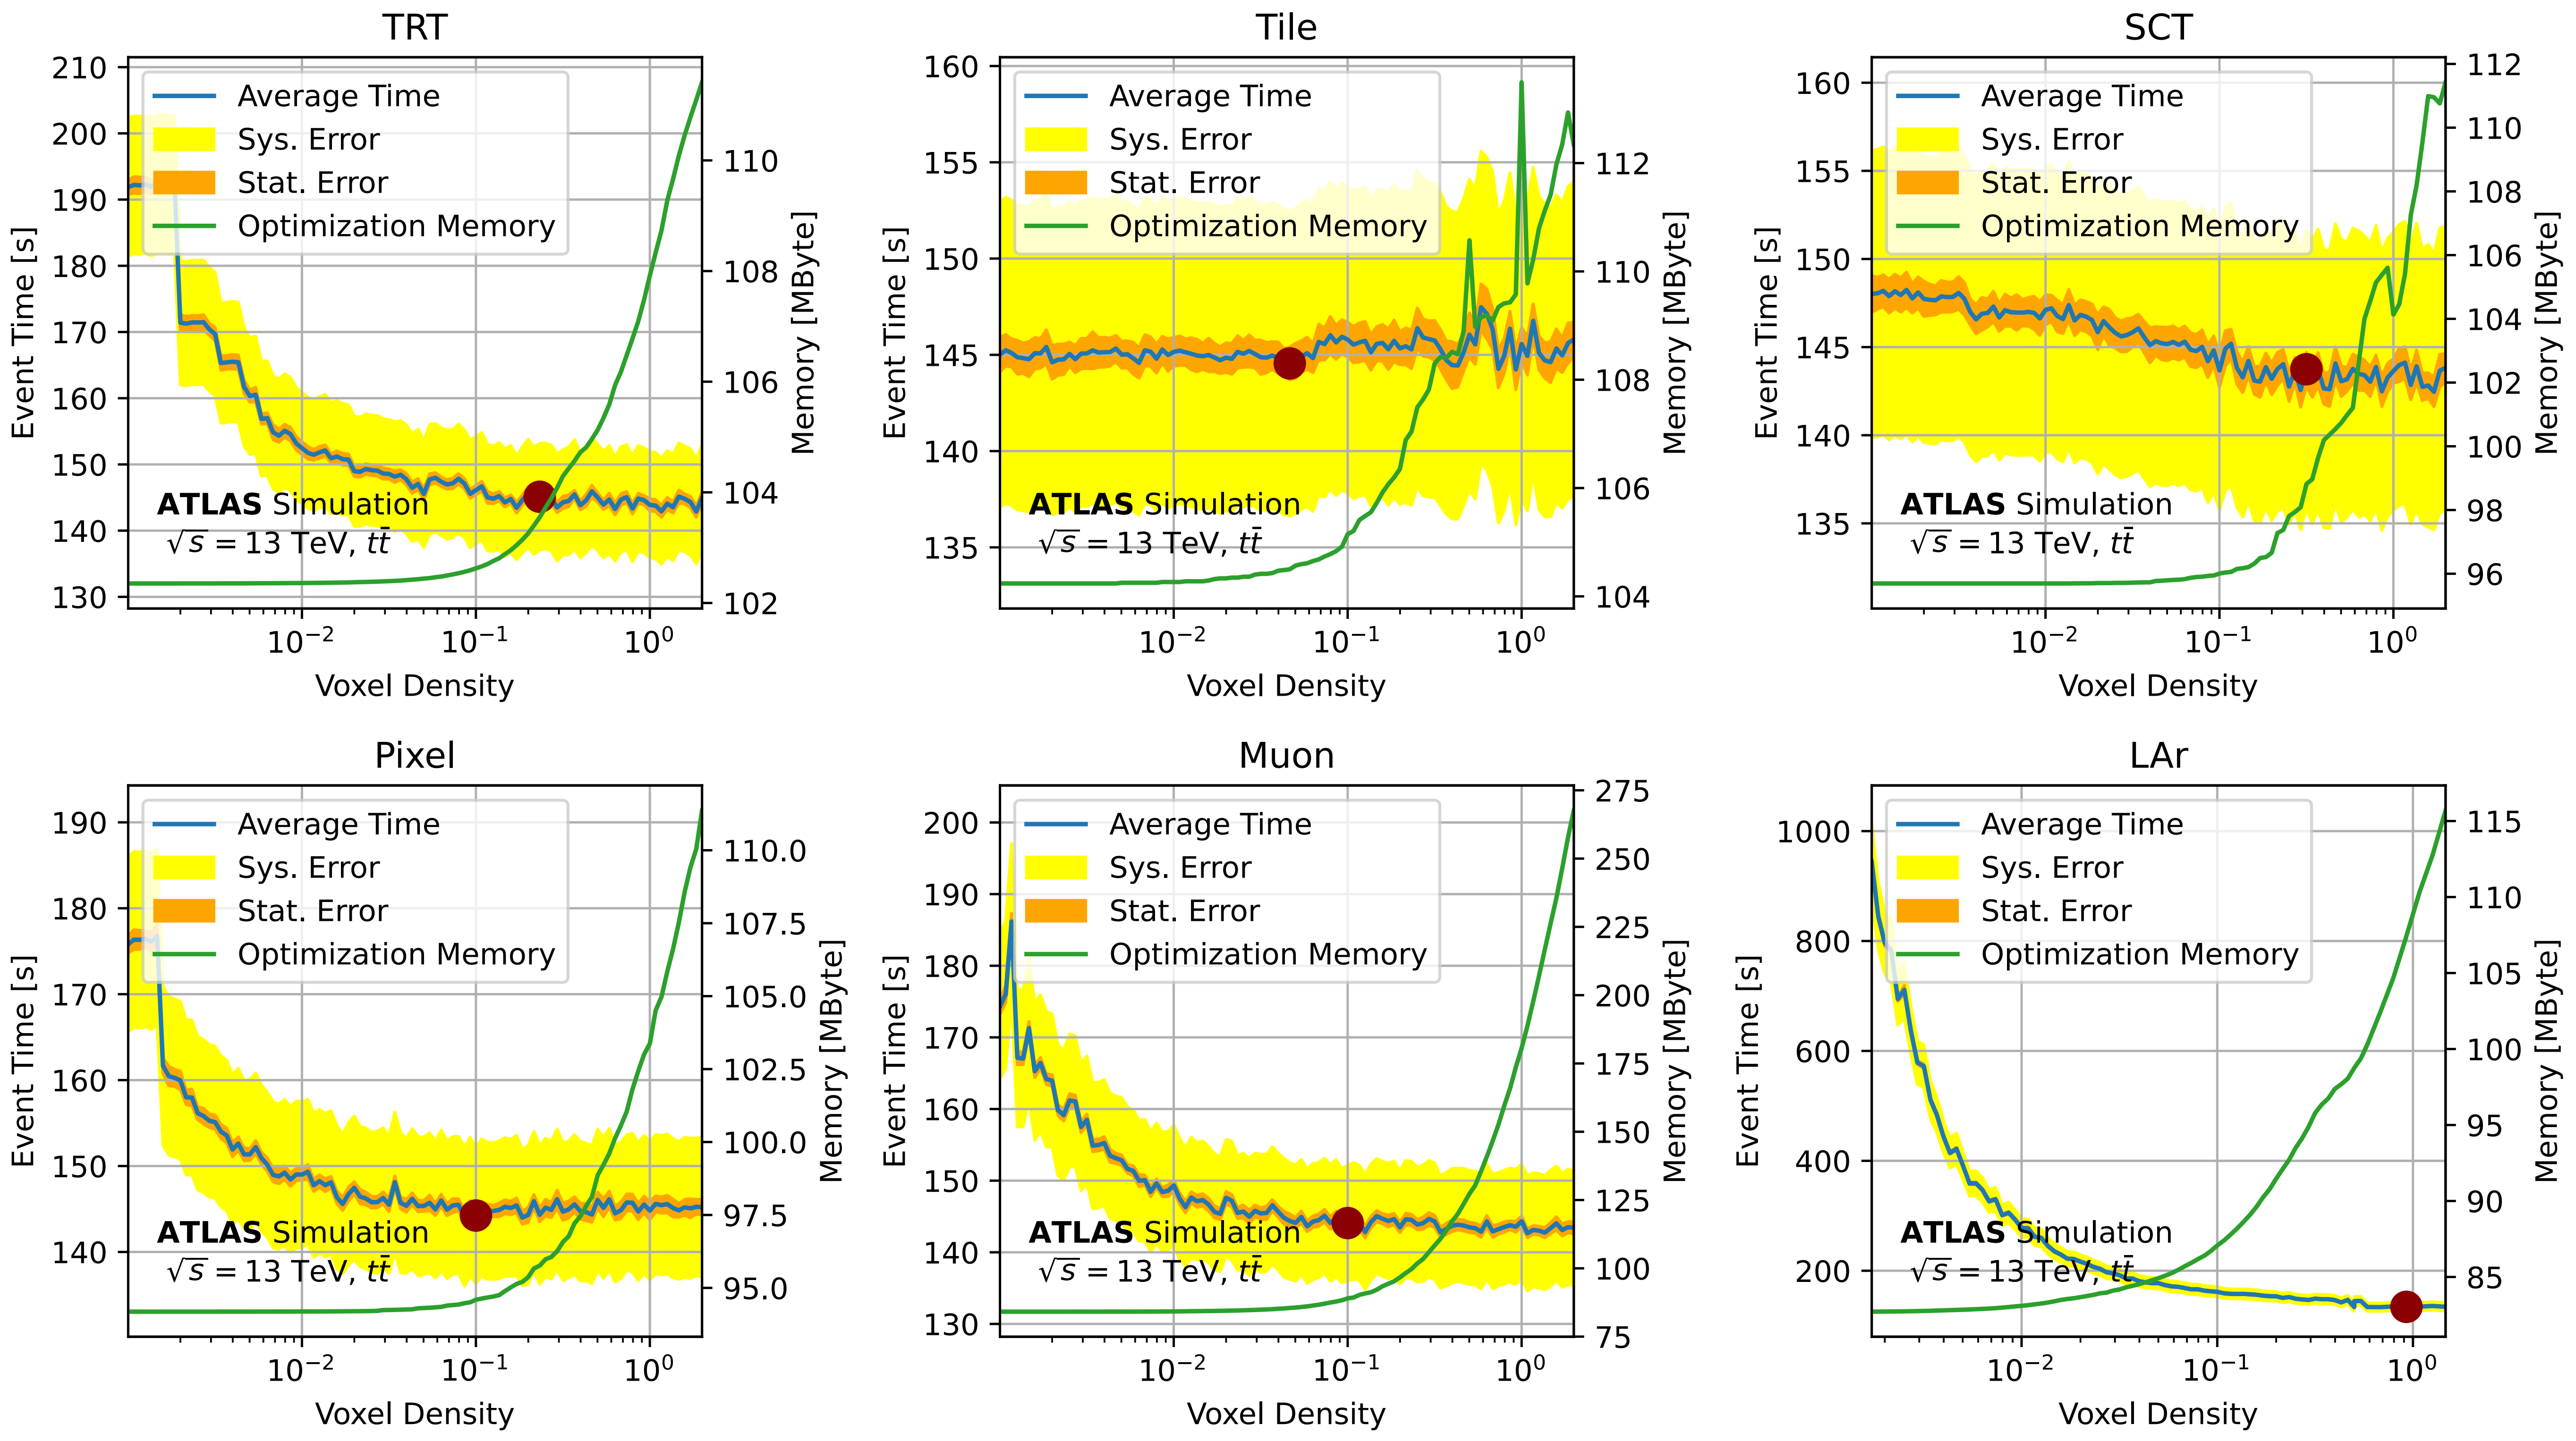

larger/coarser voxels ⟷ *smaller/finer voxels*

Woodcock Tracking

Reducing simulation steps (and CPU time) without approximations

Woodcock, E., et al. "Techniques used in the GEM code for Monte Carlo neutronics calculations in reactors and other systems of complex geometry." Proc. Conf. Applications of Computing Methods to Reactor Problems. Vol. 557. No. 2. 1965.

- Especially powerful in highly granular detectors (e.g, the EMEC) where geometric boundaries limit steps, rather than interactions
- Performs **tracking in one material without boundaries**: the densest (Pb)
  - Avoids many steps caused by geometric boundaries (*Transportation*)
- **Interaction probability is proportional to the cross section ratio between the real material and Pb**
- Implemented as a modification of the `G4GammaGeneralProcess`

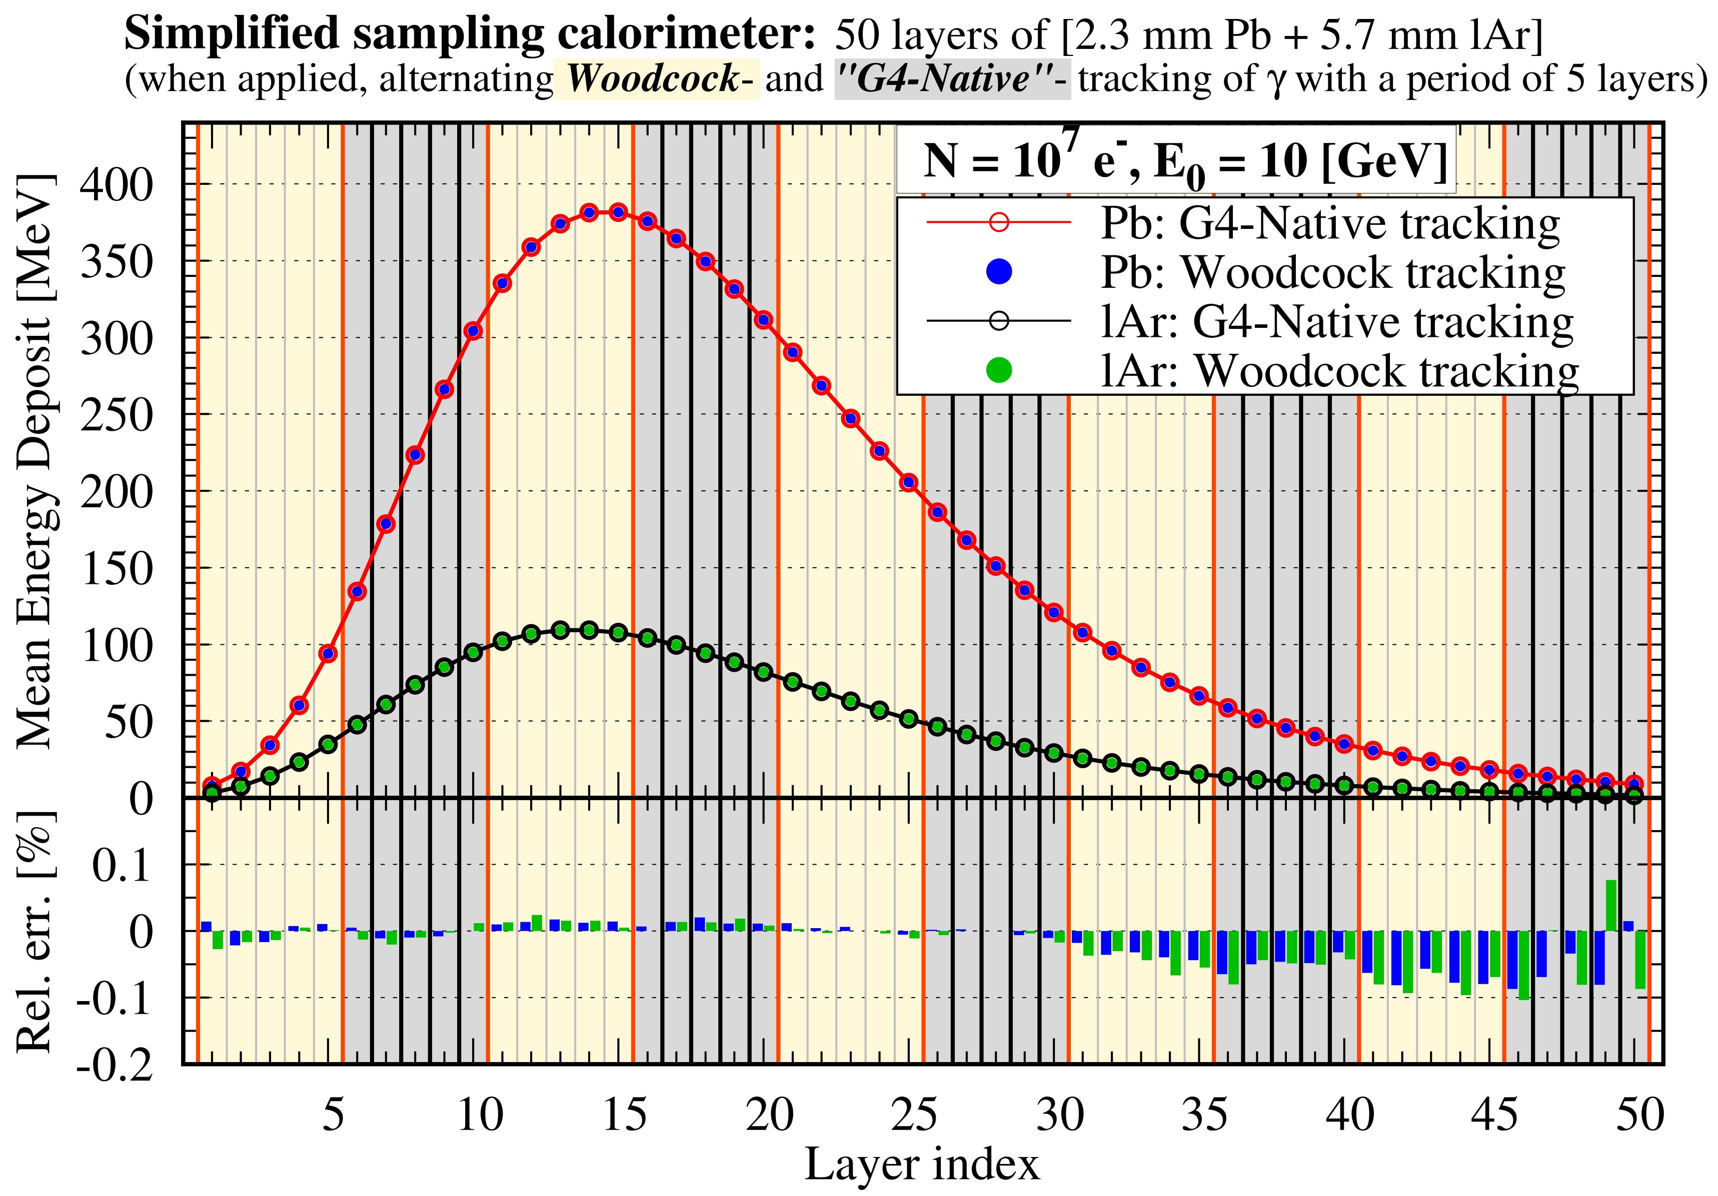

Results from stand-alone simulation of the ATLAS detector with Woodcock Tracking enabled in EMEC

		FTFP_BERT_ATL	_WDCK *Woodcock*(EMEC)
#steps	**charged**	3.548e+06	3.550e+06
	neutral	**8.501e+06**	**4.215e+06**

- ~50% EMEC step reduction
- 8-9% overall CPU speedup

Longer R&D Projects

1. **G4HepEM library**
   New compact and GPU-friendly Geant4 HEP EM library provides significant speed up with specialized tracking
2. **GPU-friendly EMEC description**
   Description of the EMEC with standard Geant4/VecGeom shapes provides speedup due to internal vectorization and portability to accelerator hardware
3. **TRT Geometry Optimization**
   Reduce the usage of Boolean shapes/operations using alternative descriptions
4. **New beam-pipe particle filter**
   Kill primary particles generating secondaries close to the beam-pipe at 5-6 m that will never cause any energy deposition to the detector
5. **Quantized State System Stepper**
   Efficient handling of discontinuities (geometric boundaries) in the simulation
6. **Thread Local Storage optimization**
   Reduce TLS usage/bottlenecks in magnetic field handling
7. **Machine Learning Correction for Aggressive Range Cuts**
   Fast production and correction of low-fidelity EM showers

Conclusions

Optimizations delivered **>50% higher throughput** so far

1.5 times more events can be simulated using the same computational resources compared to Run 2

More optimizations in R&D phase

ongoing and new ideas target both the upcoming Run 3 and future Run 4 MC campaigns

Excellent collaboration with the Geant4 team allowed R&D projects to turn into production