

Celeritas: HEP detector simulation on GPUs

arXiv
white paper

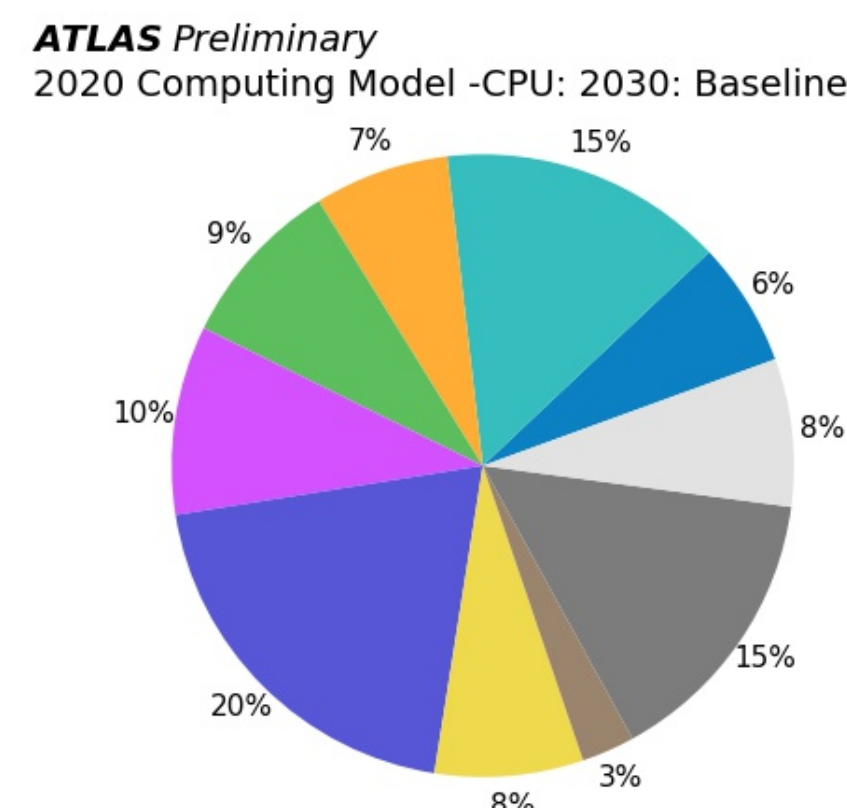
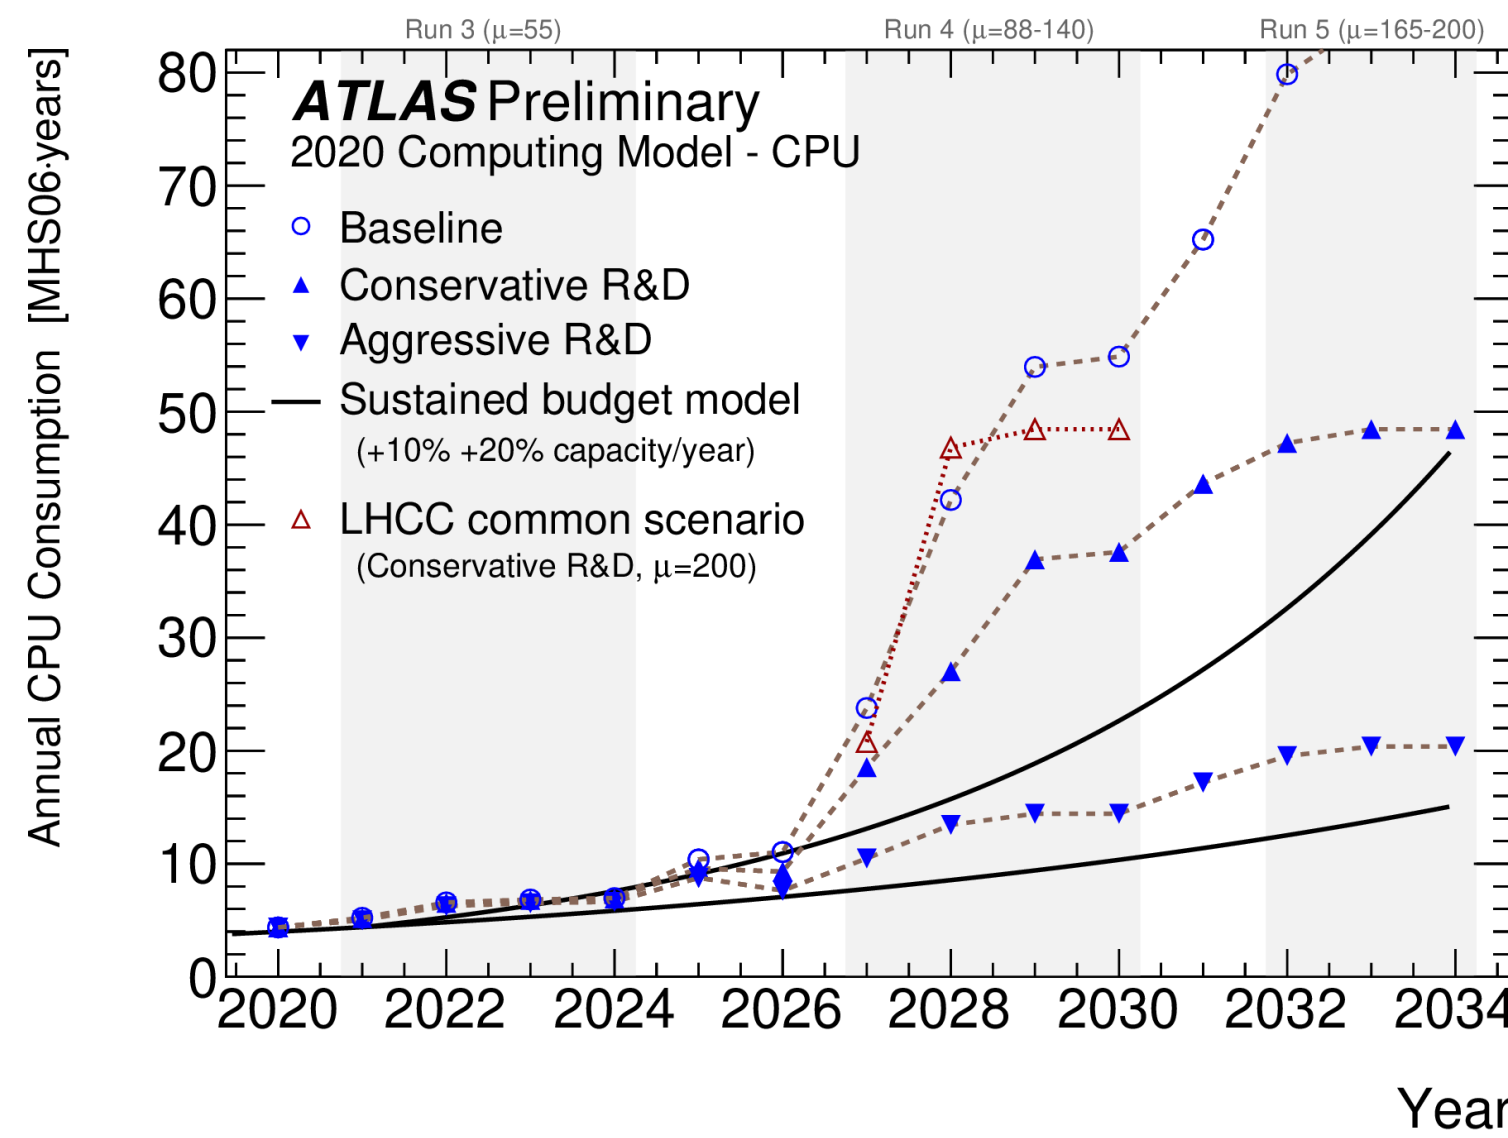
S. C. Tognini, S. R. Johnson and T. M. Evans
P. K. Romano and A. L. Lund
P. Canal, G. Lima and S. Y. Jun
V. R. Pascuzzi

Oak Ridge National Laboratory
Argonne National Laboratory
Fermi National Accelerator Laboratory
Brookhaven National Laboratory



MOTIVATION

Unprecedented computing demand on both HL-LHC and DUNE



ATLAS HL-LHC Computing Conceptual Design Report. CERN-LHCC-2020-015/LHCC-G-178 (2020)

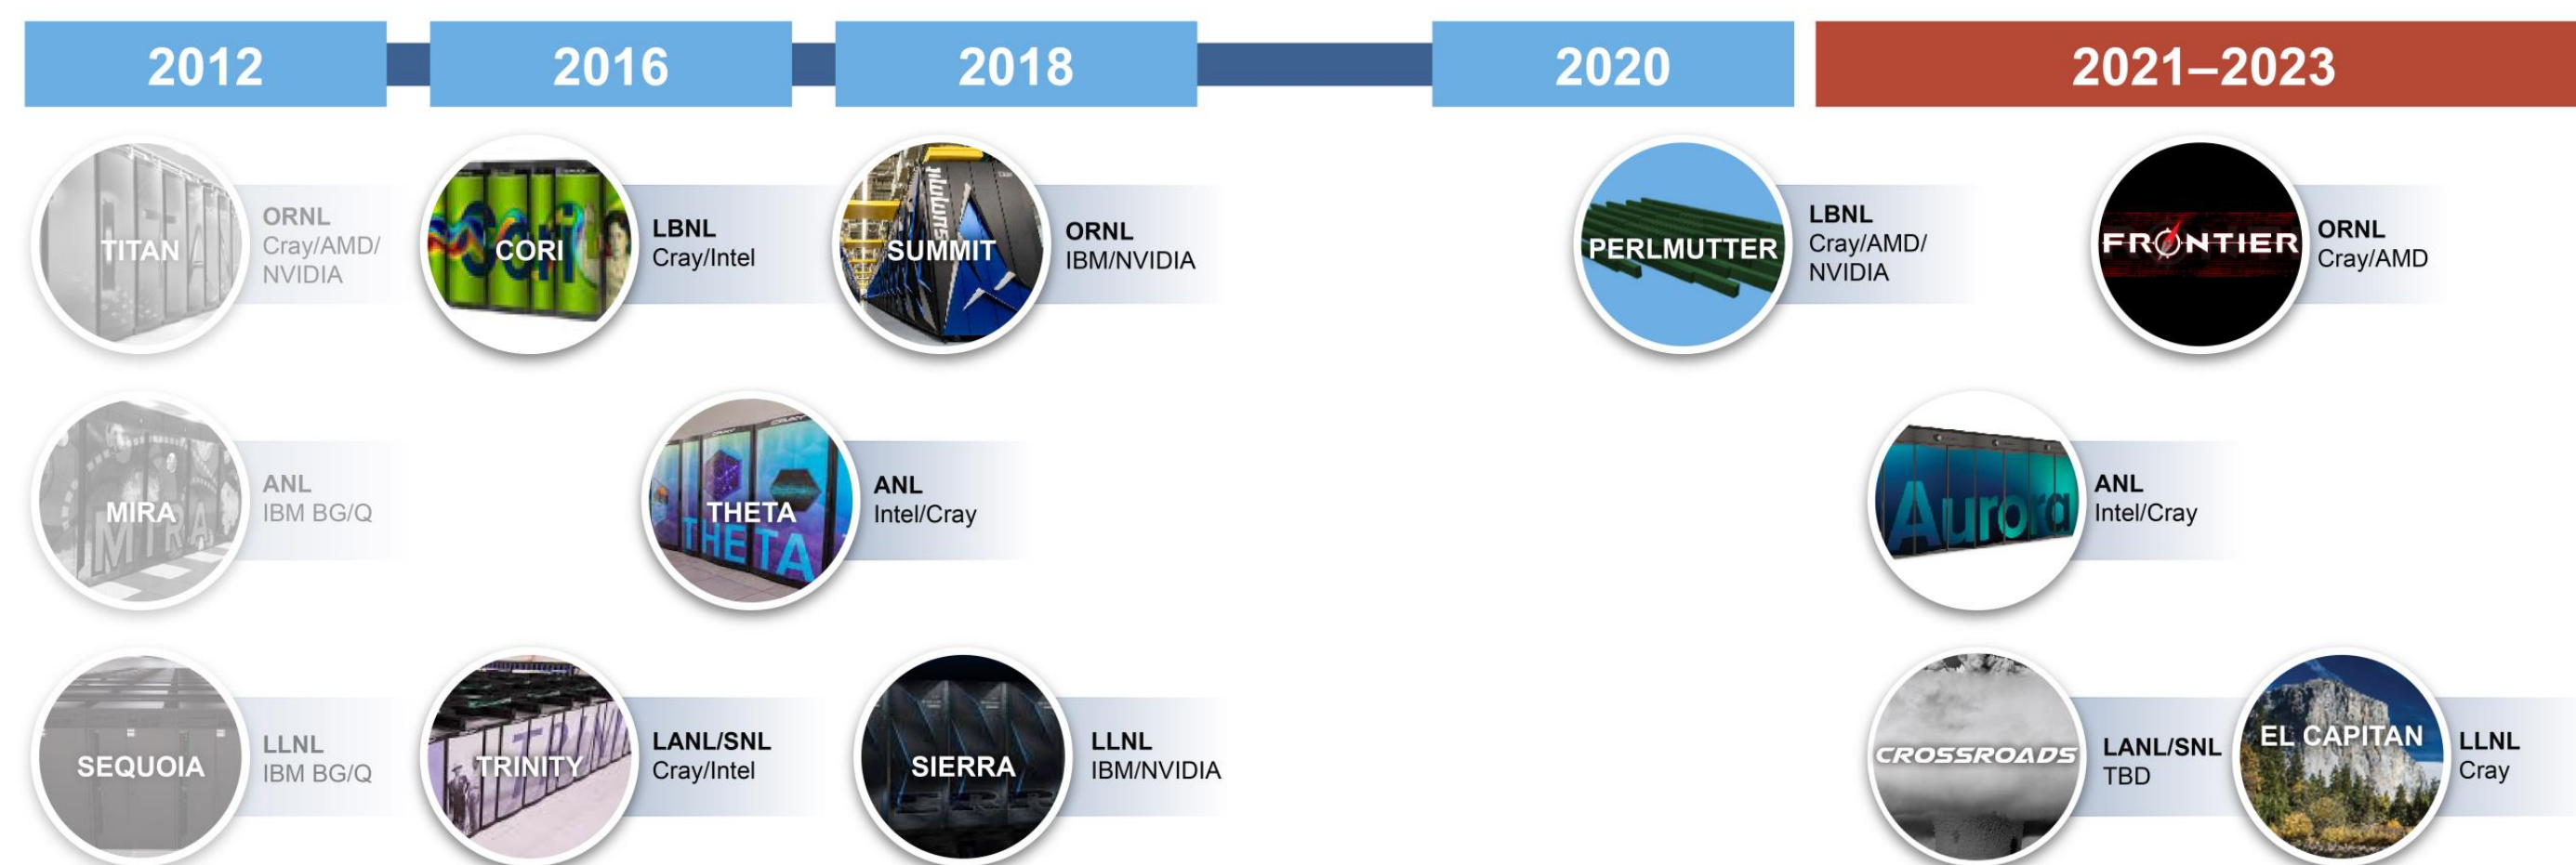
Detector simulation can take up a large fraction of the total CPU time

EXASCALE COMPUTING SYSTEMS & PROJECT

Pre-Exascale Systems



Future Exascale Systems



D. B. Kothe, et al. **The US DOE Exascale Computing Project (ECP) Perspective for the HEP Community**
A Coordinated Ecosystem for HL-LHC Computing R&D. Washington D.C. (2019)

- 24 open science applications have been ported and deployed on LCFs—most resulted in complete rewrites
- ExaSMR: Coupled Monte Carlo Neutronics and Fluid Flow Simulation of Small Modular Reactors*
- Summit: 1 GPU \equiv 160 CPU cores[†]**

[†]S. P. Hamilton, et al. **Continuous-energy Monte Carlo neutron transport on GPUs in the Shift code**
Annals of Nuclear Energy, vol. **128**, pp. 236–247 (2019)

IMPACT

Celeritas is designed to take full advantage of DOE LCFs

- ORNL alone is home of both Summit and Frontier
- If *Celeritas* reaches a CPU to GPU factor of 160 (such as ExaSMR)
 - Summit: 27,648 GPUs \rightarrow 4,423,680 CPU cores**
 - The Worldwide LHC Computing Grid (WLCG) had \sim 500,000 CPU cores in 2017^{††}
 - Summit can reach an equivalent computing power of 8 WLCG**

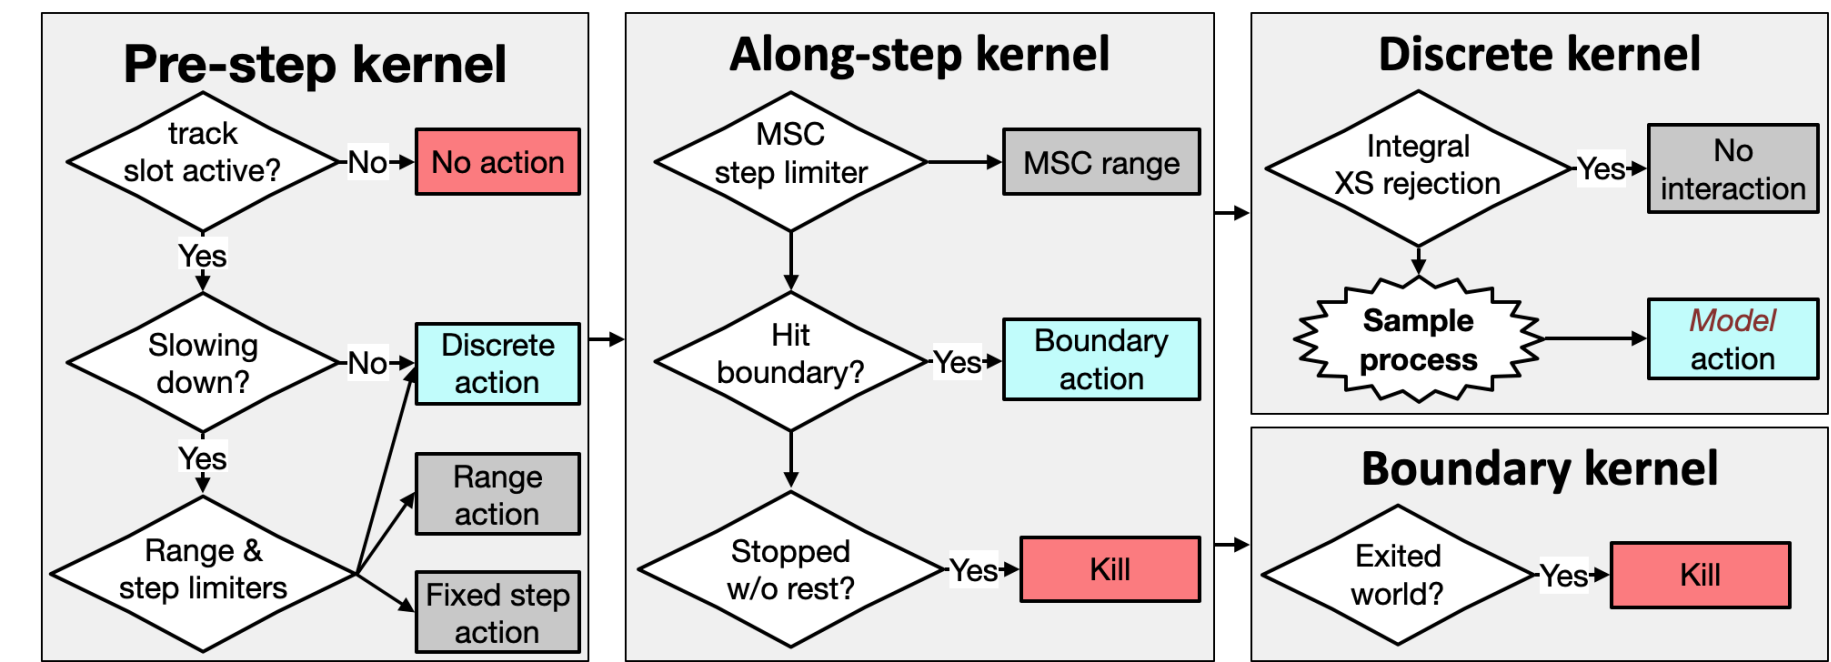
^{††} A Roadmap for HEP Software and Computing R&D for the 2020s. Comput. Softw. Big Sci. **3**, 7 (2019)

CODE ARCHITECTURE & PHYSICS

- Data-oriented programming
- Composition-based classes
- Extensive unit-testing

Geometries:

- ORANGE** (Oak Ridge Advanced Nested Geometry Engine)
- VecGeom**
- Multi-architecture:
- NVIDIA & AMD** currently



Particle	Process	Model(s)
γ	photon conversion	Bethe-Heitler
	Compton scattering	Klein-Nishina
	photoelectric effect	Livermore
	Rayleigh scattering	Livermore
e^\pm	ionization	Møller-Bhabha
	bremsstrahlung	Seltzer-Berger, relativistic
	pair annihilation	EPlusGG
	multiple scattering	Urban, WentzelVI

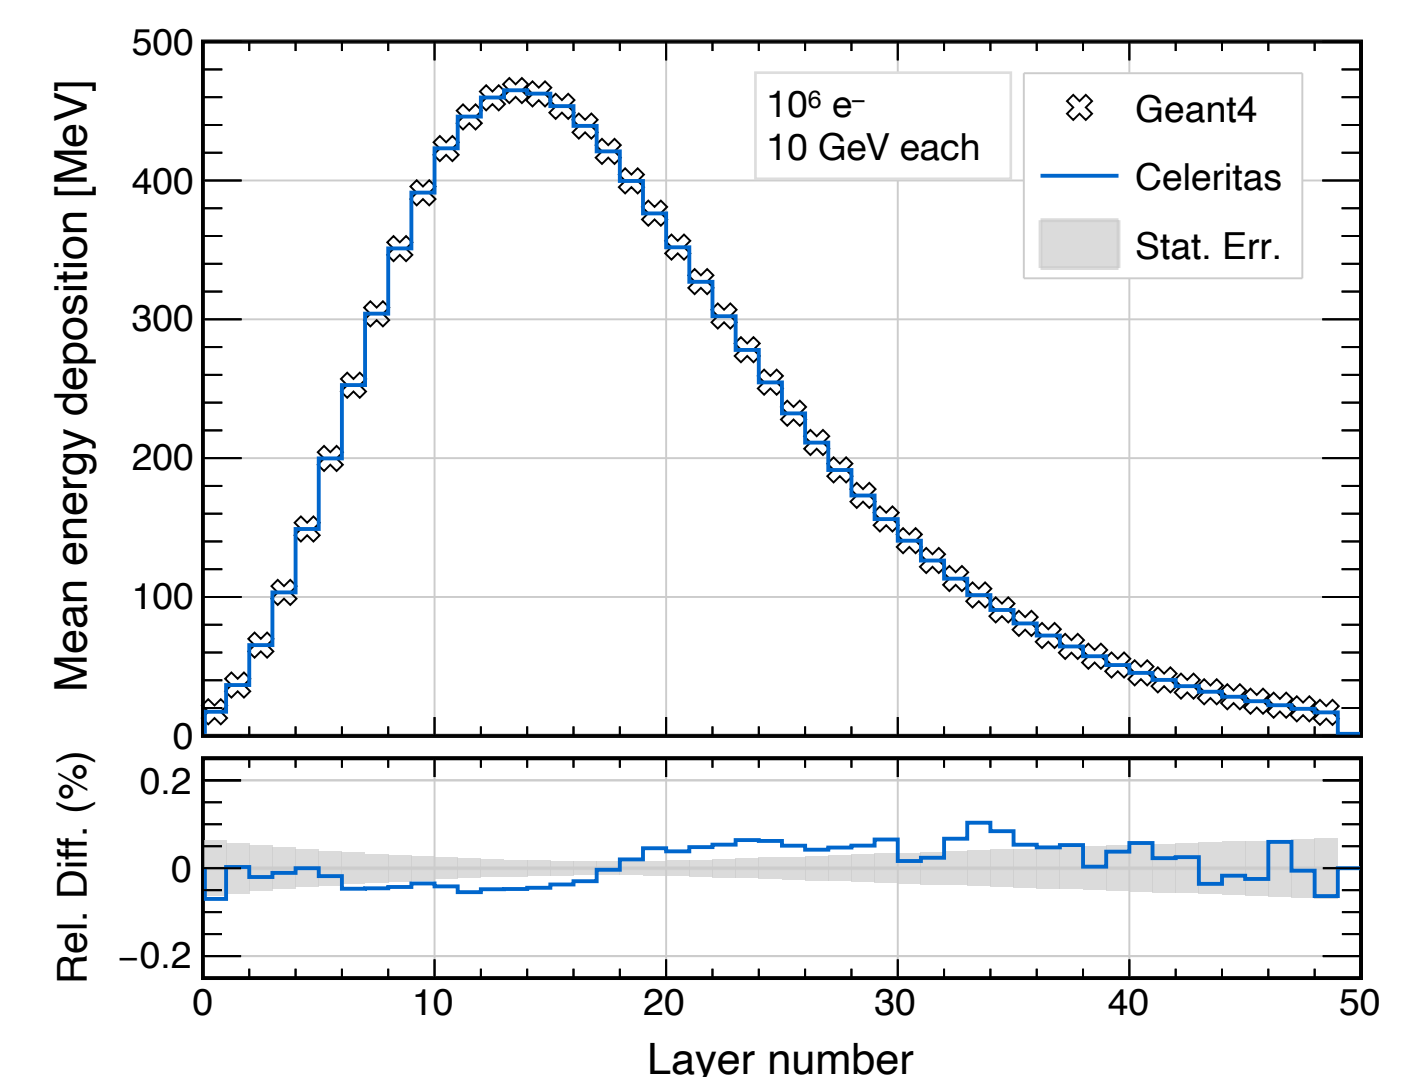
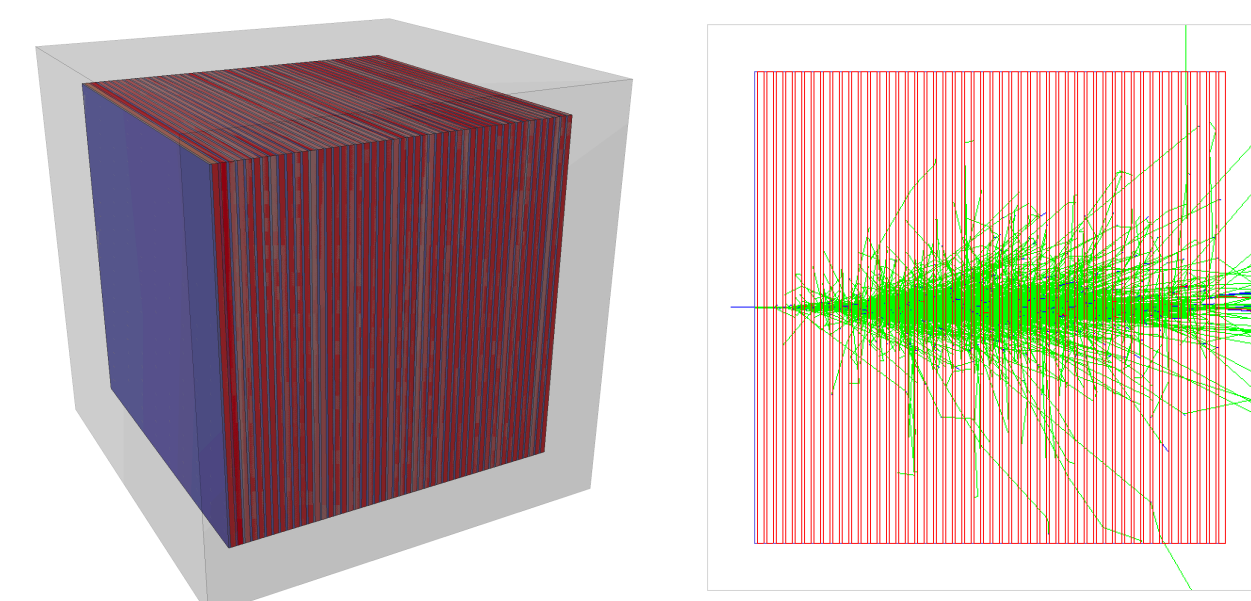
- Comprehensive EM physics
- Tag release v0.1 imminent**

PRELIMINARY PERFORMANCE RESULTS

Summit

CPU IBM Power9 22C @ 3.07 GHz
GPU NVIDIA Tesla V100 @ 1.53 GHz
(80 SM, 64 cores each, 16 GB mem.)

- Geometry: 50 layers of LAr + Pb



- CPU: 7 cores multithread**
- GPU: 1 core per GPU**

Celeritas per-node performance

- \sim 40 \times faster on GPU**
- 1 GPU \equiv \sim 280 CPU cores**

New preliminary results; **NOT** published yet

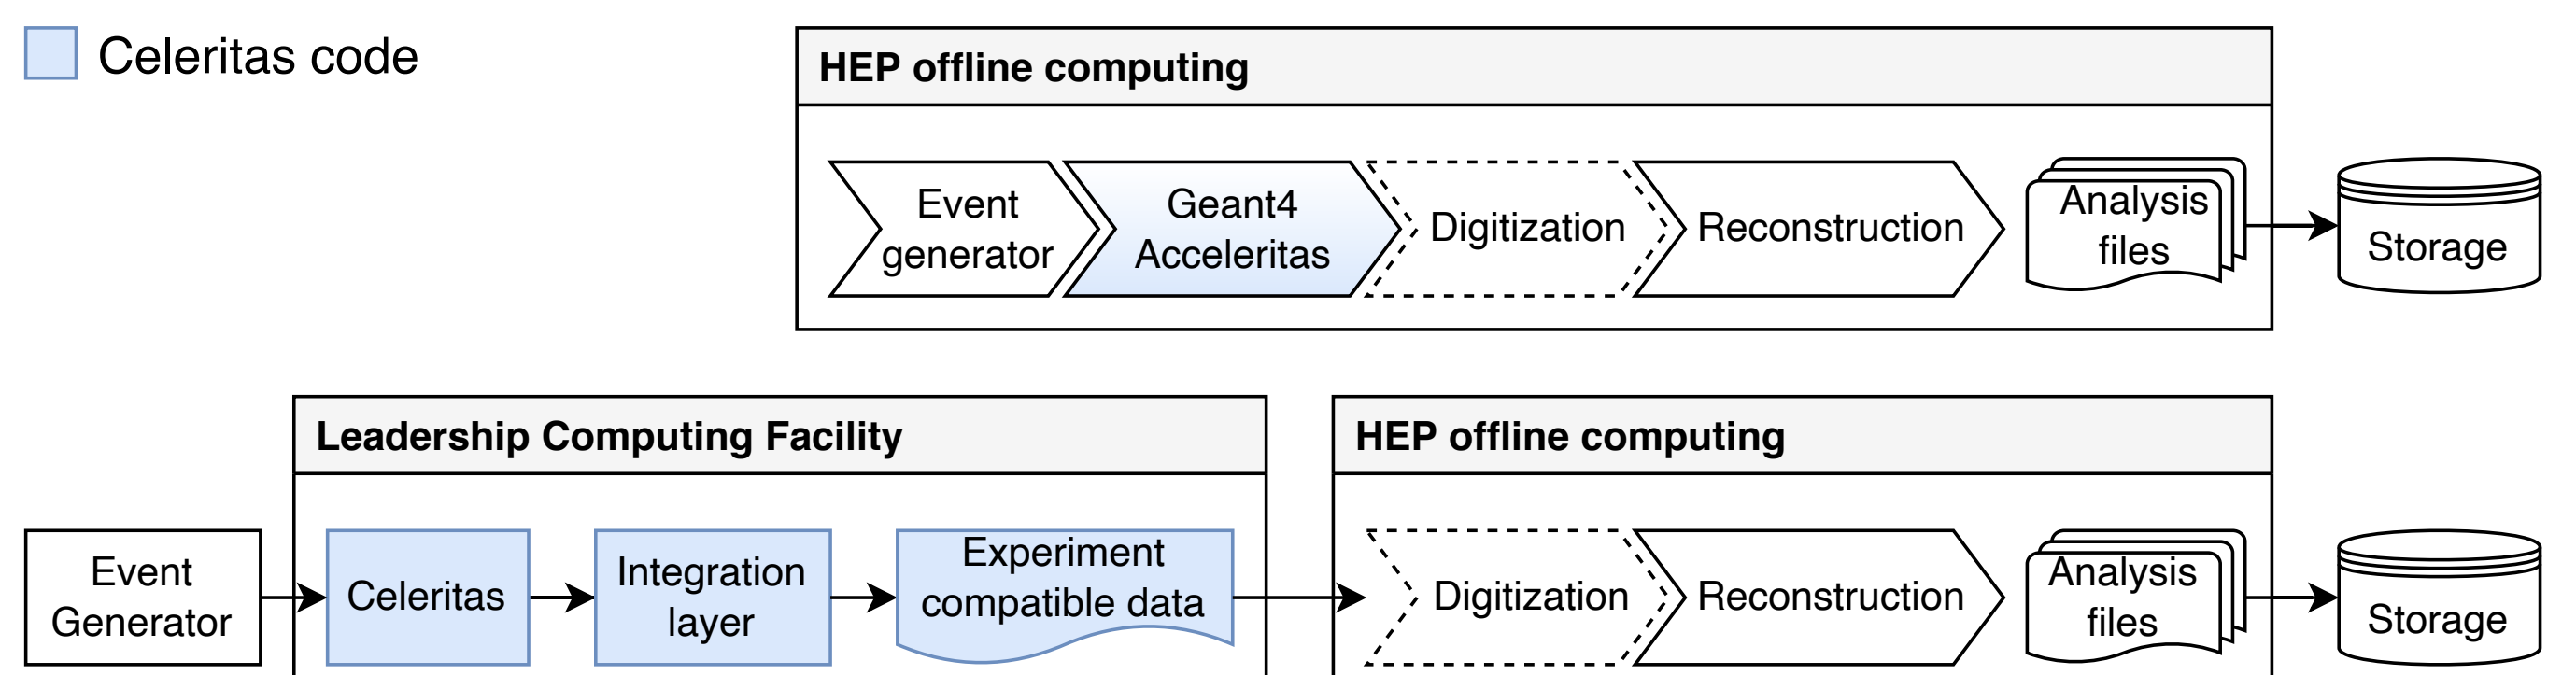
Input		
Application	Arch.	Events
Geant4 (v10.7.1)	CPU	10^4
Celeritas	CPU	10^3
	GPU	10^5

Wall time per primary			
Application	Geo.	Arch.	Mean (ms)
Geant4 (10.7.1)	Geant4	CPU	2.9 ± 0.1
		GPU	0.046 ± 0.001
Celeritas	ORANGE*	CPU	1.95 ± 0.04
	VecGeom	GPU	0.0627 ± 0.0004

*Oak Ridge Advanced Nested Geometry Engine

EXPERIMENT INTEGRATION

- Working with experiments to provide a proof-of-principle workflow
- Acceleritas*: Geant4-Celeritas hybrid that offloads EM physics to GPU



ACKNOWLEDGEMENTS

Work for this paper was supported by Oak Ridge National Laboratory (ORNL), which is managed and operated by UT-Battelle, LLC, for the U.S. Department of Energy (DOE) under Contract No. DEAC05-00OR22725 and by Fermi National Accelerator Laboratory, managed and operated by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy (Fermilab publication number for this paper is FERMILAB-FN-1159-SCD). This research was supported by the Exascale Computing Project (ECP), project number 17-SC-20-SC. The ECP is a collaborative effort of two DOE organizations, the Office of Science and the National Nuclear Security Administration, that are responsible for the planning and preparation of a capable exascale ecosystem— including software, applications, hardware, advanced system engineering, and early testbed platforms—to support the nation's exascale computing imperative. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.