Oak Ridge National Laboratory

From muons to supercomputers

Stefano Tognini

Research Associate

Oak Ridge National Laboratory

ORNL is managed by UT–Battelle, LLC for the US Department of Energy

U.S. DEPARTMENT OF ENERGY

Summary

- **Part 1 (past)**
  - MINOS/MINOS+
  - NOvA

- **Part 2 (present)**
  - Computing in HEP
  - Celeritas
  - Takeaways

Part 1
Past work

Background

- Undergrad + grad school at Federal University of Goias (UFG)
  - Somehow, iNSPIRE calls it Goias University...

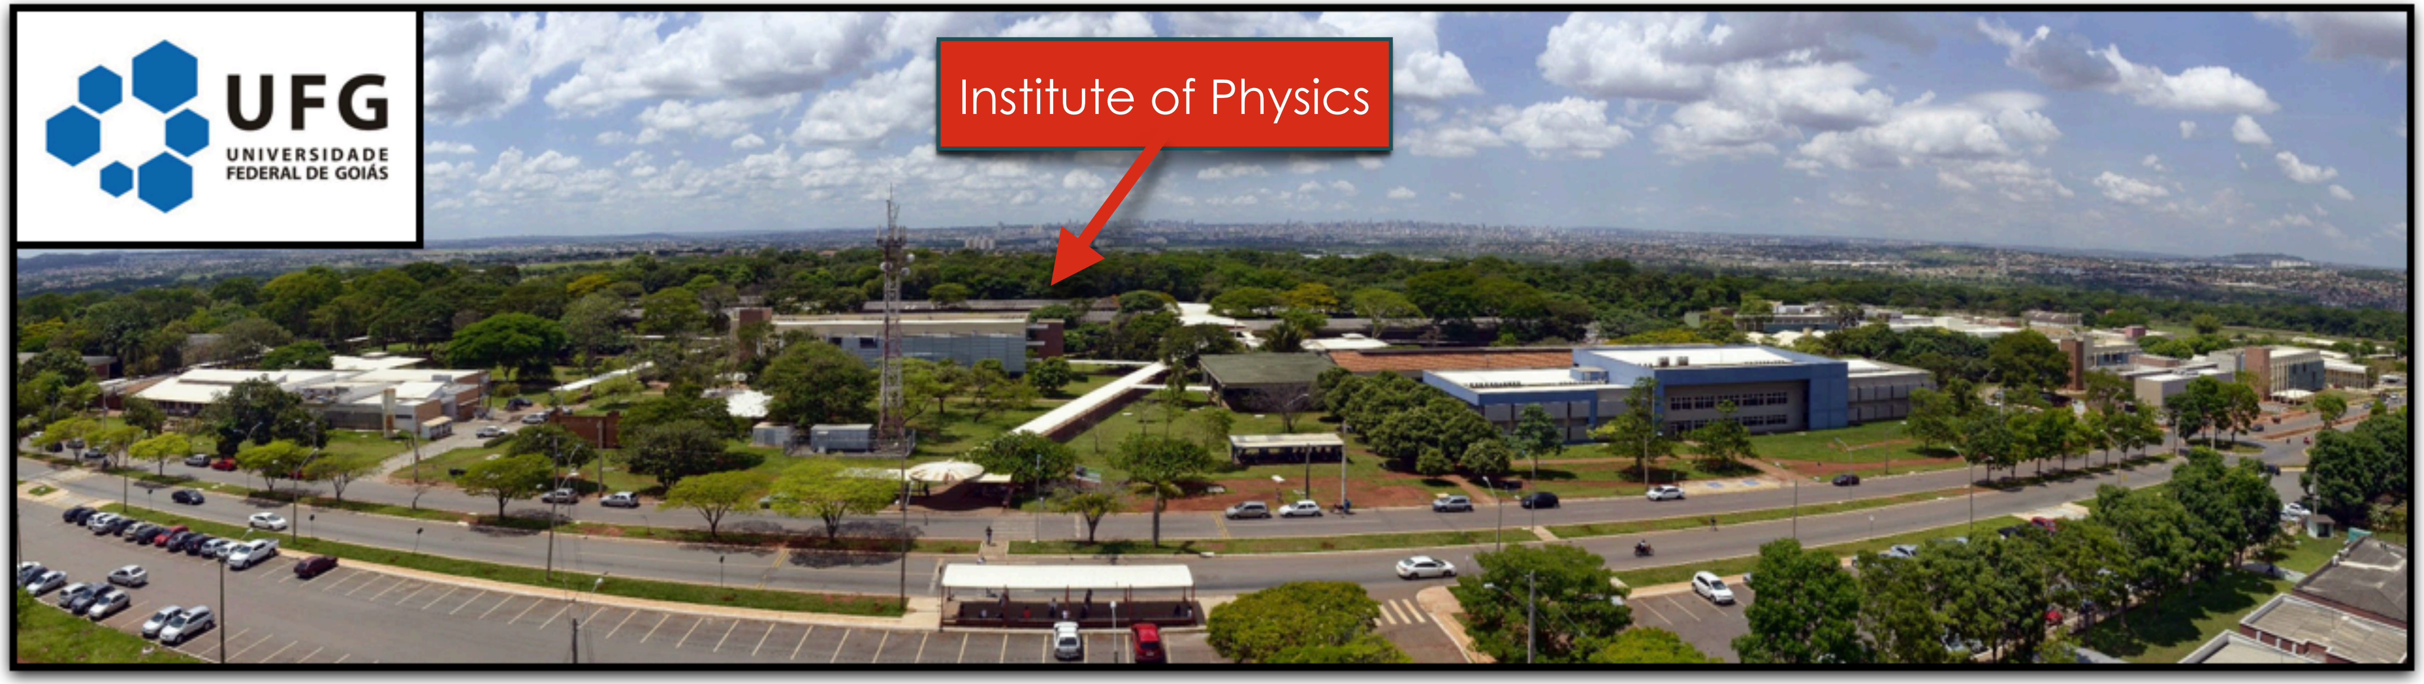

Background

- **Undergrad**
  - Senior thesis on neutral hyperon semileptonic decays

- **Master's**
  - Cosmic ray MC using CORSIKA
    - Wrote a CORSIKA binary to ROOT converter
    - Deployed CORSIKA in our small CPU cluster at UFG
    - Thesis:
      - Simulation of atmospheric temperature effects on cosmic ray muons

MINOS/MINOS+

- Started on MINOS in 2011 (I guess I'm old...)

- Main service: **Remote Operation Centers**
  - Get MINOS ROC scripts up and running
  - Write documentation
  - Create certification procedures and hand-on/hand-off rules for shifters
    - Allegedly, all neutrino experiment certifications are based on MINOS

MINOS/MINOS+

- Our ROC was the first certified MINOS ROC in the World

Certified ROC's

CERTIFIED ON	STATION	INSTITUTION'S NAME	PHONE	CURRENT RESPONSIBLE
2013-05-27	Goias	Federal University of Goias	+55 62 3521 1122 ext. 217	Stefano Tognini (stognini@fnal.gov)
2013-07-22	Warsaw	University of Warsaw	+48 22 5532817	Katarzyna Grzelak (grzelakk@fnal.gov)
2013-09-??	Duluth	Univ of Minnesota Duluth	+1 218 726 7214	Alec Habig (ahabig@umn.edu)
2013-09-25	Tufts	Tufts University	+1 617 627 4373	Tony Mann (mann@fnal.gov)
2013-11-27	UMN	Univ of Minnesota	+1 612 624 4546	Marvin Marshak (marshak@umn.edu)
2013-12-04	WM	College of William and Mary	+1 757 221 5485	Alena Devan (avgavrilenko@email.wm.edu)
2014-03-17	UCL	University College London	+44 20 7679 3425	Leigh Whitehead (l.whitehead@ucl.ac.uk)
2015-09-30	UT	University of Texas at Austin	+1 512 475 7285	Karol Lang (lang@physics.utexas.edu)
2015-10-06	Cincinnati	University of Cincinnati	+1 513 556 0480	Alex Sousa (absousa@gmail.com)
uncert	BNL	Brookhaven National Lab	+1 631 344 8954	?

MINOS/MINOS+

- I set up the MINOS ROCW... we were the first to go live

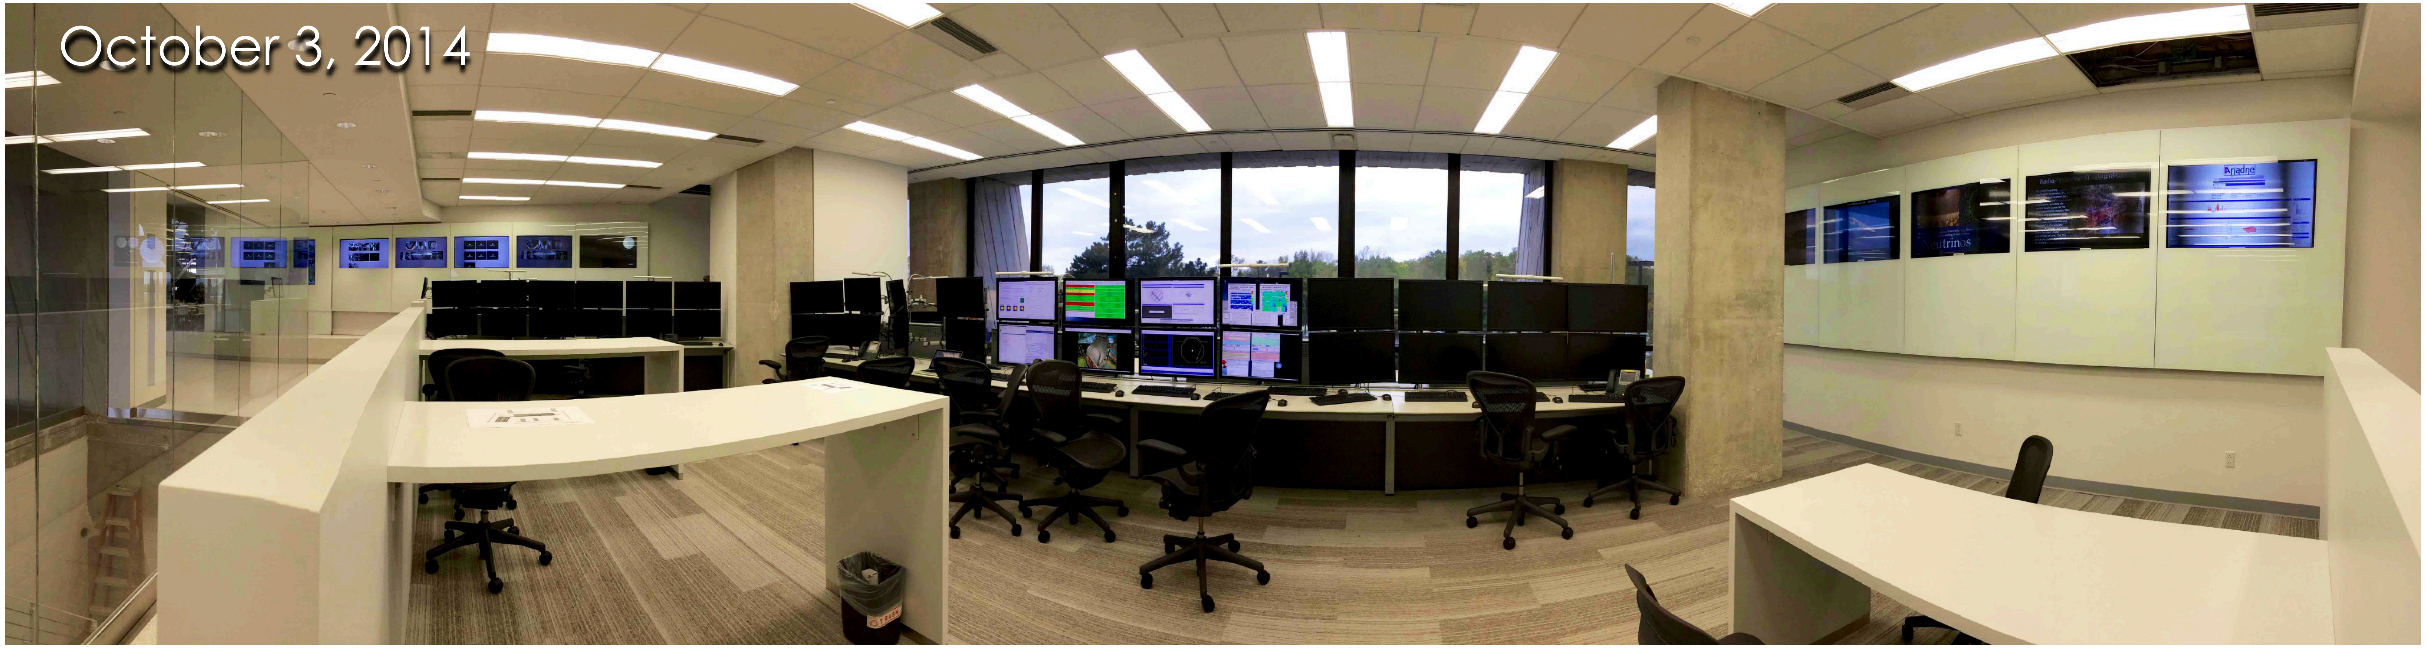

- Around that time, I became liaison for all ROCs, including ROCW

MINOS/MINOS+

- Highly involved on 2 analyses

PHYSICAL REVIEW D **91**, 112006 (2015)

Observation of seasonal variation of atmospheric multiple-muon events in the MINOS Near and Far Detectors

PHYSICAL REVIEW D **93**, 052017 (2016)

Measurement of the multiple-muon charge ratio in the MINOS Far Detector

MINOS/MINOS+

PHYSICAL REVIEW D **93**, 052017 (2016)

Measurement of the multiple-muon charge ratio in the MINOS Far Detector

- Cosmic ray primaries are mostly protons (+1), therefore

$$R \equiv \frac{N_{\mu^+}}{N_{\mu^-}} > 1$$

- Knowing **R** helps
  - Predicting the atmospheric $\nu/\bar{\nu}$ rate
  - Tuning models

MINOS/MINOS+

PHYSICAL REVIEW D **91**, 112006 (2015)

Observation of seasonal variation of atmospheric multiple-muon events in the MINOS Near and Far Detectors

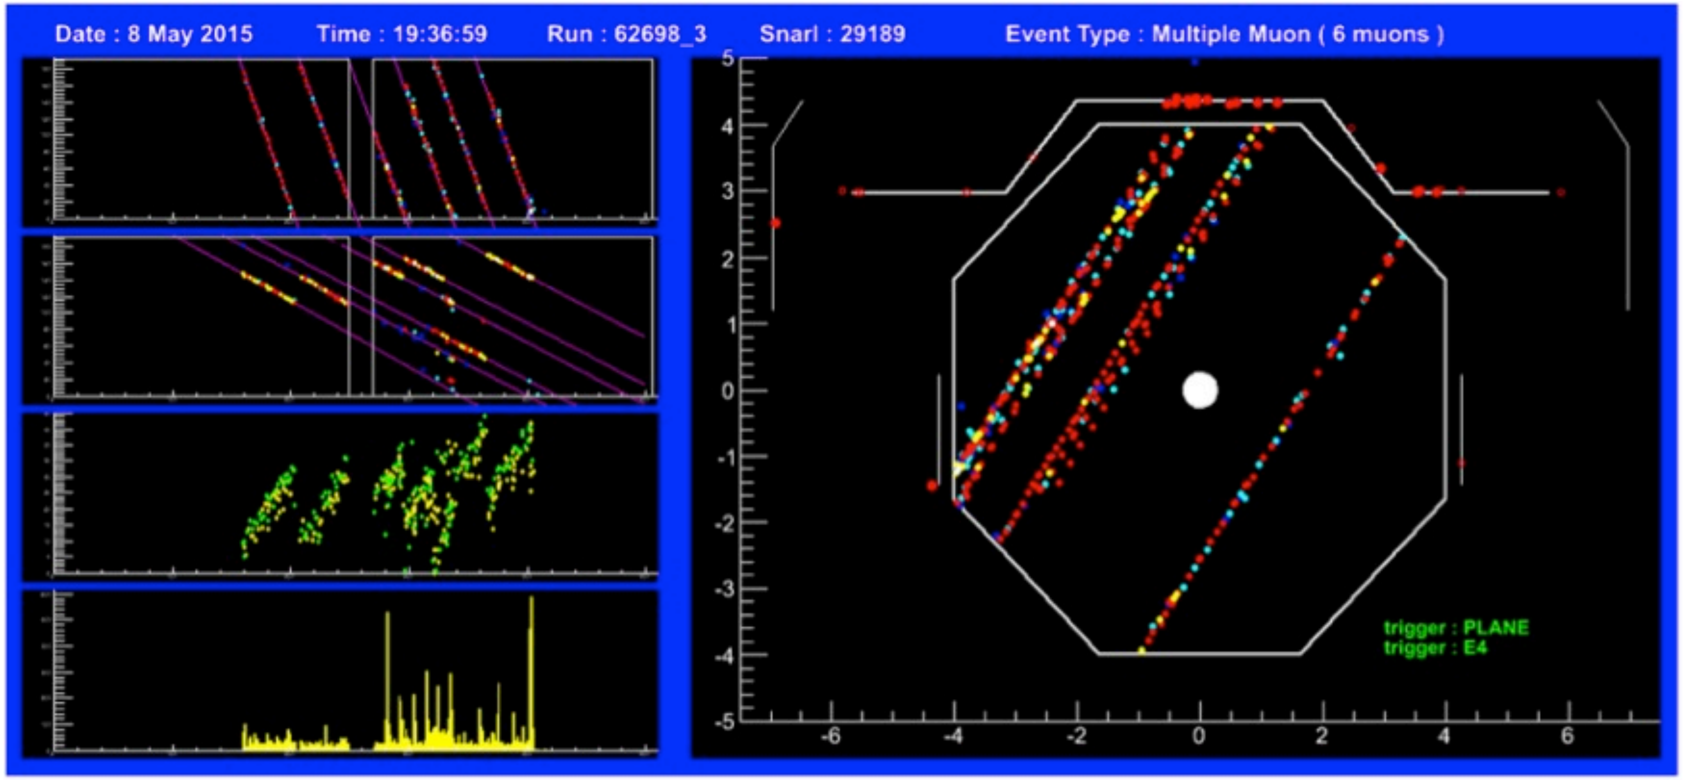

Single-muons
Low E

Multiple-muons
High E

https://www.fnal.gov/pub/today/archive/archive_2015/today15-05-15.html

MINOS/MINOS+

PHYSICAL REVIEW D **91**, 112006 (2015)

Observation of seasonal variation of atmospheric multiple-muon events in the MINOS Near and Far Detectors

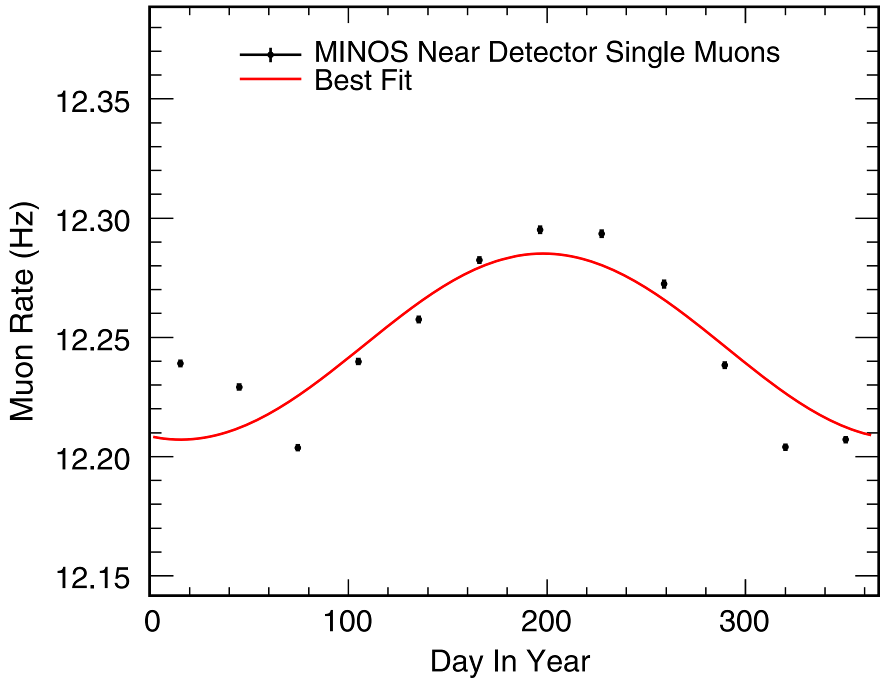

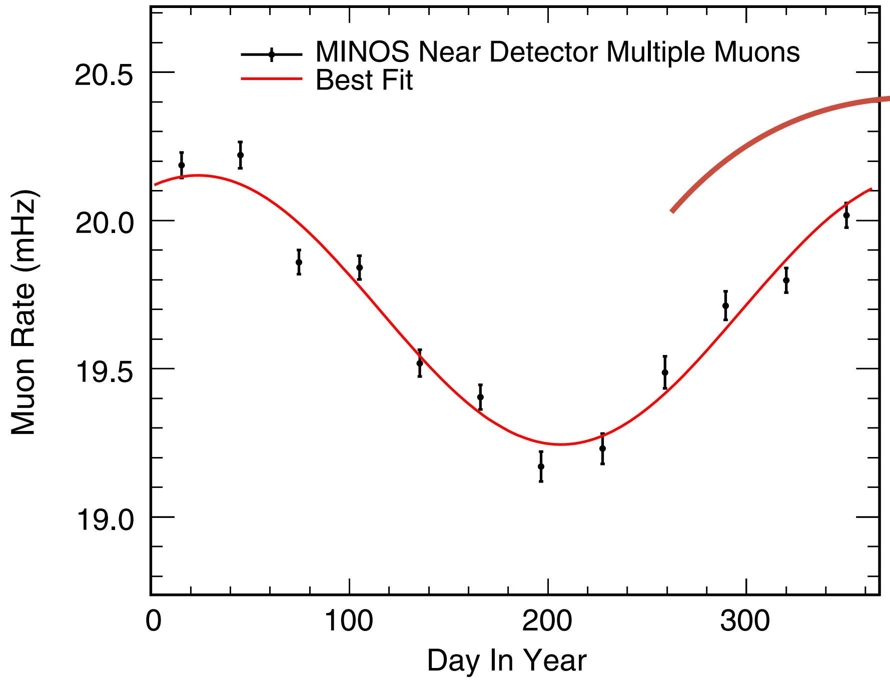

Don't know why

4 hypotheses

One is very likely to be the explanation

NOvA (Ph.D.)

PHYSICAL REVIEW D **99**, 122004 (2019)

Observation of seasonal variation of atmospheric multiple-muon events in the NOvA Near Detector

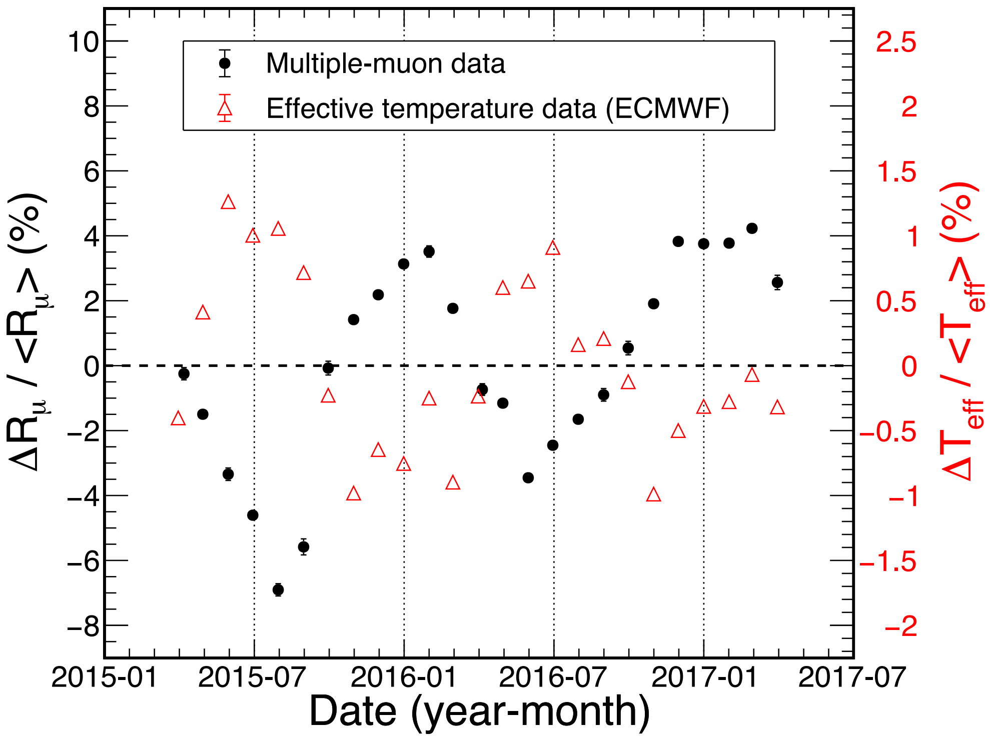

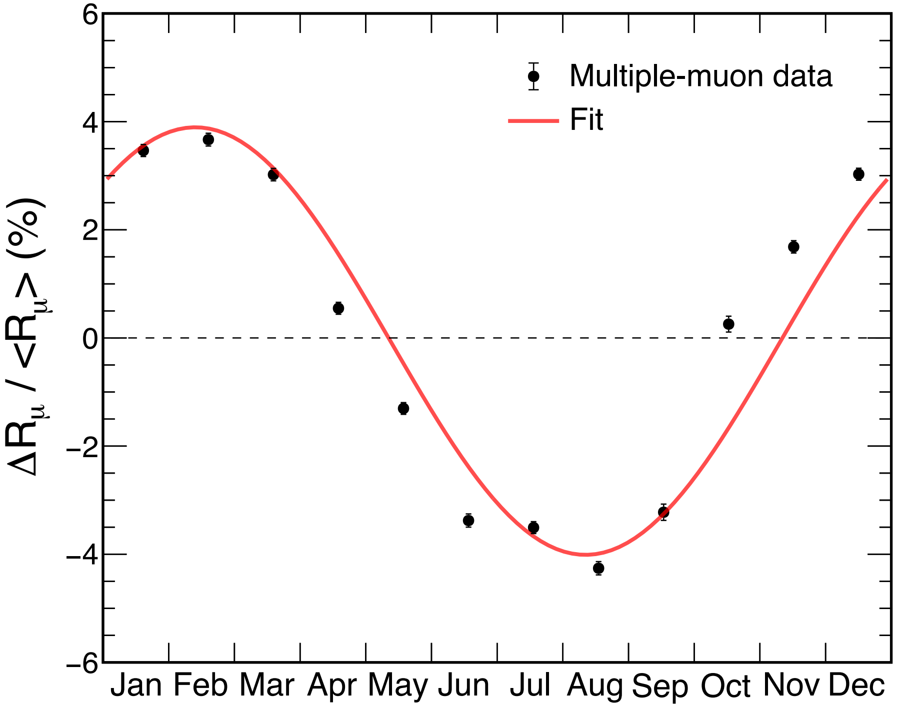

OAK RIDGE National Laboratory

NOvA (Ph.D.)

PHYSICAL REVIEW D **99**, 122004 (2019)

Observation of seasonal variation of atmospheric multiple-muon events in the NOvA Near Detector

- Lots of computing work, almost all single-handed:
  - Reconstruction algorithm written almost from scratch
  - Convert ECMWF (atmospheric temperature) GRIB data to ROOT
  - Tweak CRY MC to produce multiple-muons
  - MC validation, data production
  - Analysis files, scripts, and macros

NOvA

- Computing services:
  - More ROC work (now liaison of MINOS+ ROCs & NOvA ROC @ UFG)

- Yes, we've got a bigger room!
- 2016: MINOS+ shifts were over
- Active ROCS:
  - NOvA & LArIAT

NOvA

- Computing services:
  - Add CORSIKA to NOvASoft (later on to LArSoft)

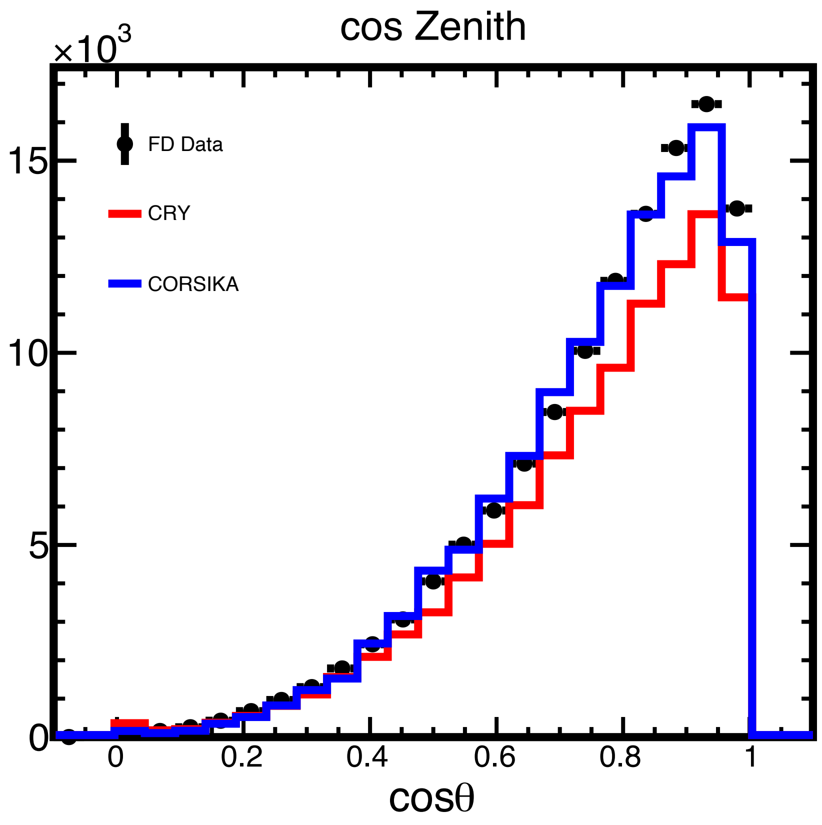

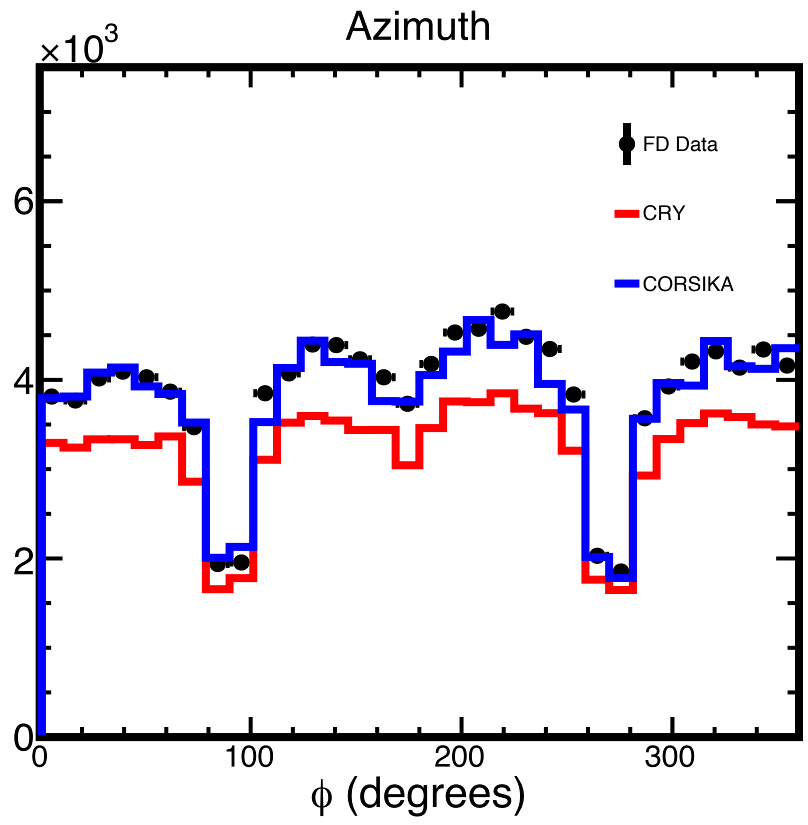

CRY matched "well" if ×1.2

Except neutrons (low flux)

CORSIKA showed good results

Part 2
Current work

Stefano C. Tognini
Research Associate in Computational Particle Physics
HPC Methods for Nuclear Applications
Nuclear Energy and Fuel Cycle Division
Fusion and Fission Energy and Science Directorate

togninis@ornl.gov
www.ornl.gov

1 Bethel Valley Road
P.O. Box 2008, MS 6170
Oak Ridge, TN, 37831

DOE Exascale Computing Project

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.
indico.cern.ch/event/834880/

Oak Ridge Leadership Computing Facility

Petascale

Exascale

2012

2018

2022

Titan

AMD Opteron
NVIDIA Kepler
27 PF **35% from GPU**

Summit

IBM Power9
NVIDIA Volta
200 PF **95% from GPU**

Frontier

AMD EPYC
AMD Instinct 250X
> 1.5 EF

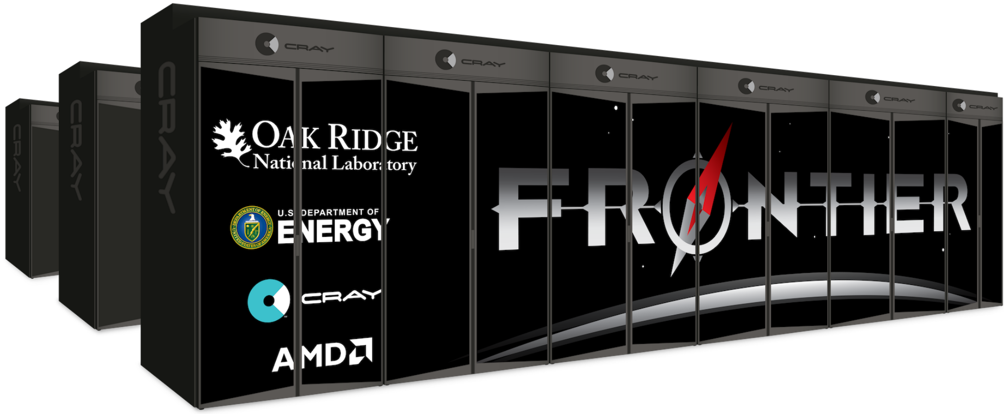

For more, visit www.olcf.ornl.gov

ExaSMR / Shift

- ExaSMR: *Coupled MC Neutronics and Fluid Flow Simulation of SMRs*
  - Shift – GPU MC radiation transport code

Nearly perfect parallel scaling efficiency on Summit

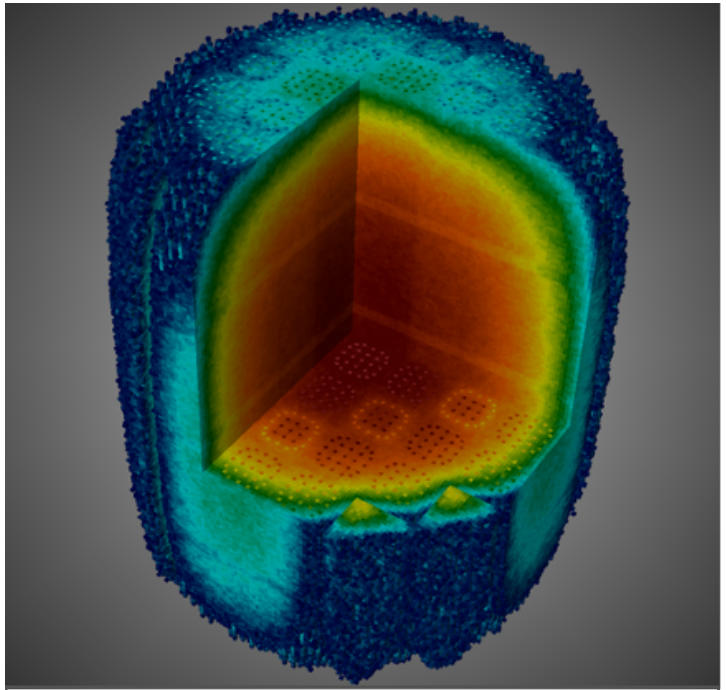

Total reaction rate in a
Small Modular Reactor (SMR) core

~160× speedup

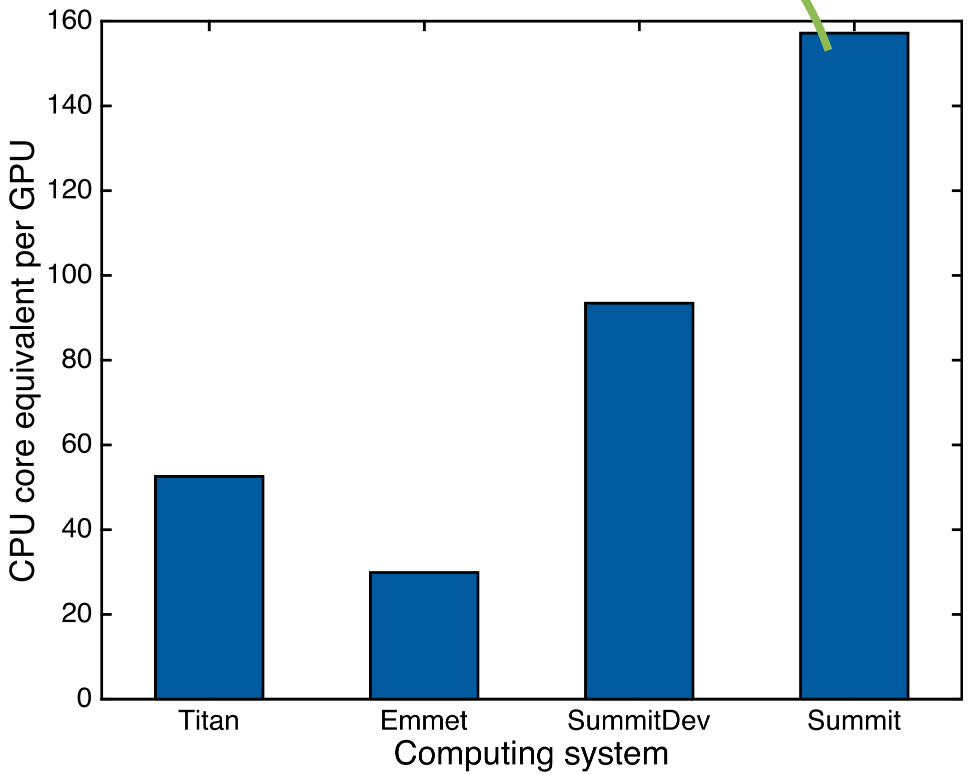

Annals of Nuclear Energy, vol. **128**, pp. 236-247 (2019)

HEP computing challenges | LHC

HL-LHC

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

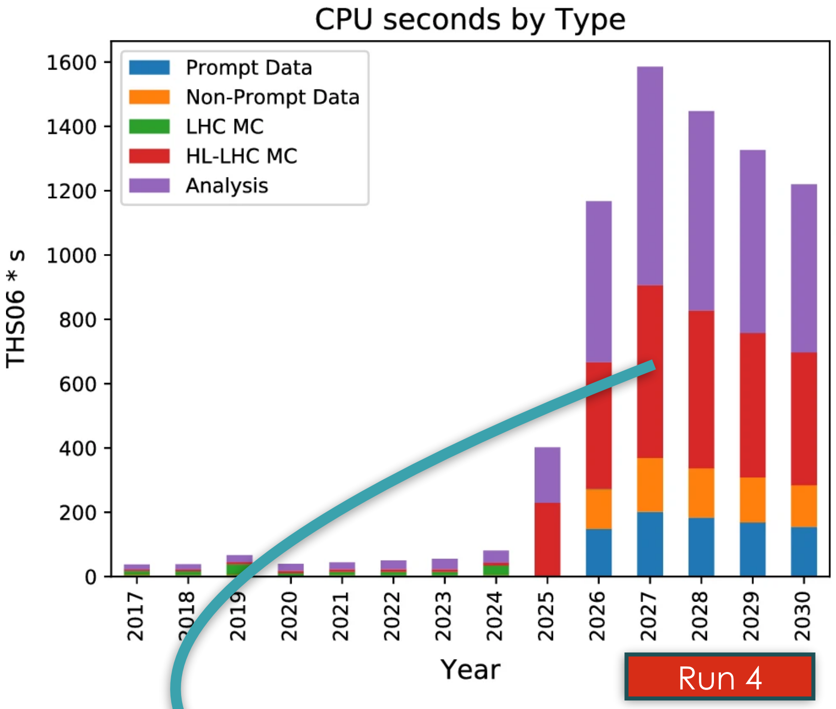

MC could take ~30% of the HL-LHC CPU time

ATLAS

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

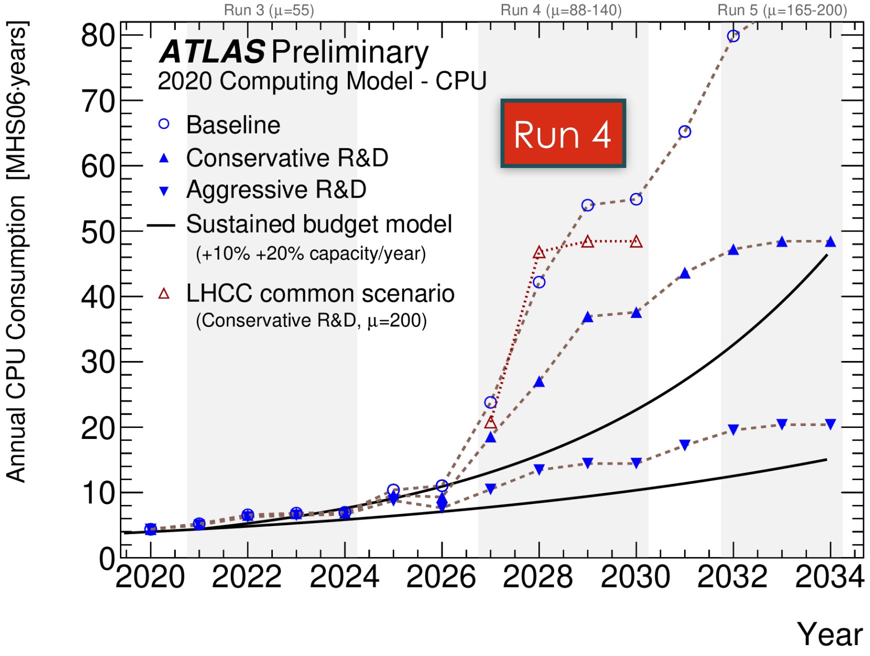

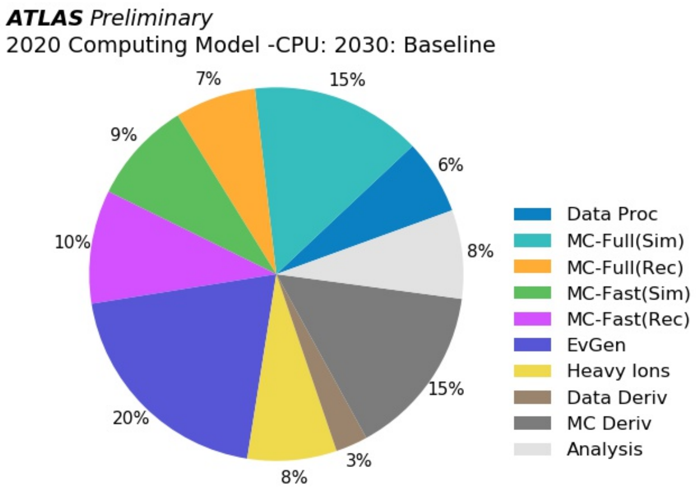

- Evt Gen: 20%
- MC: 39%
- Reco: TBD

Opportunities

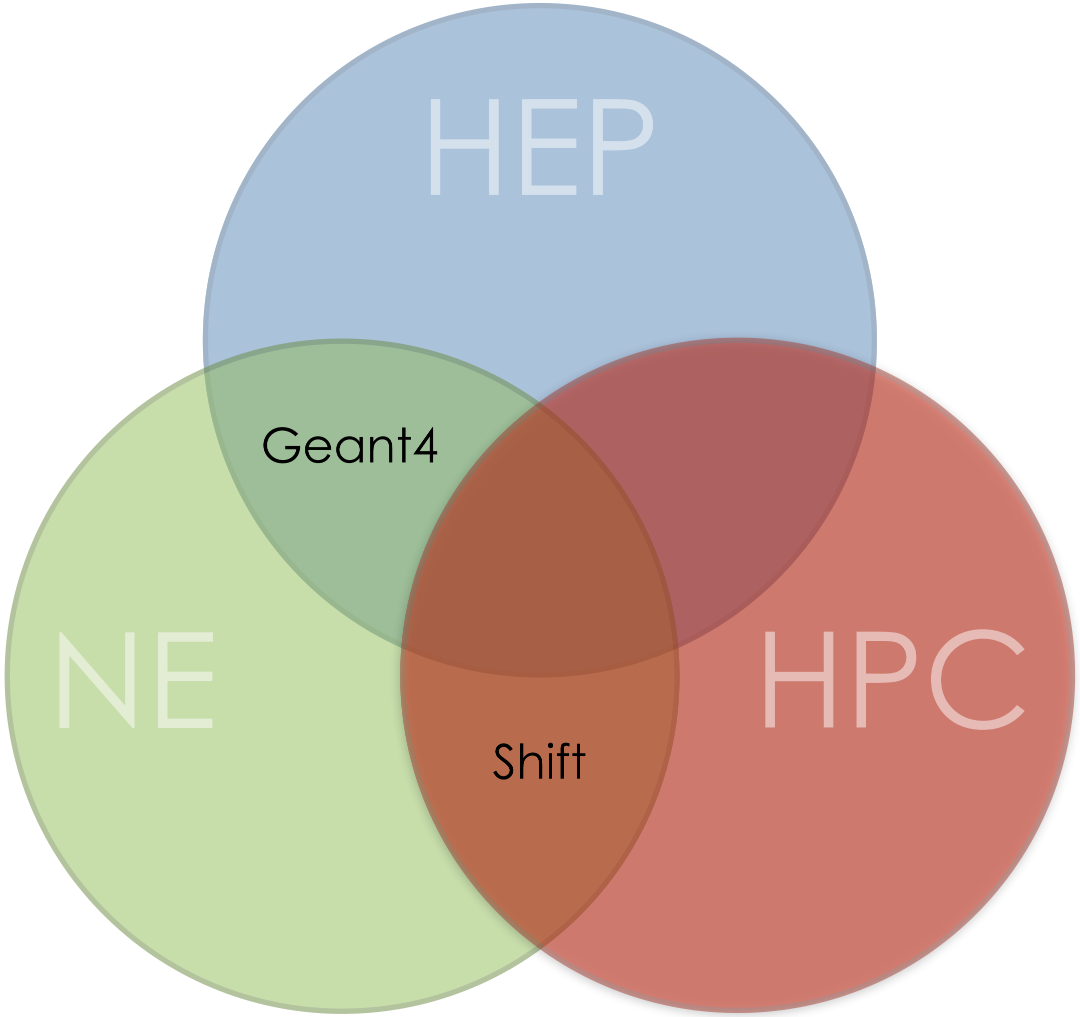

Oak Ridge National Laboratory

Opportunities

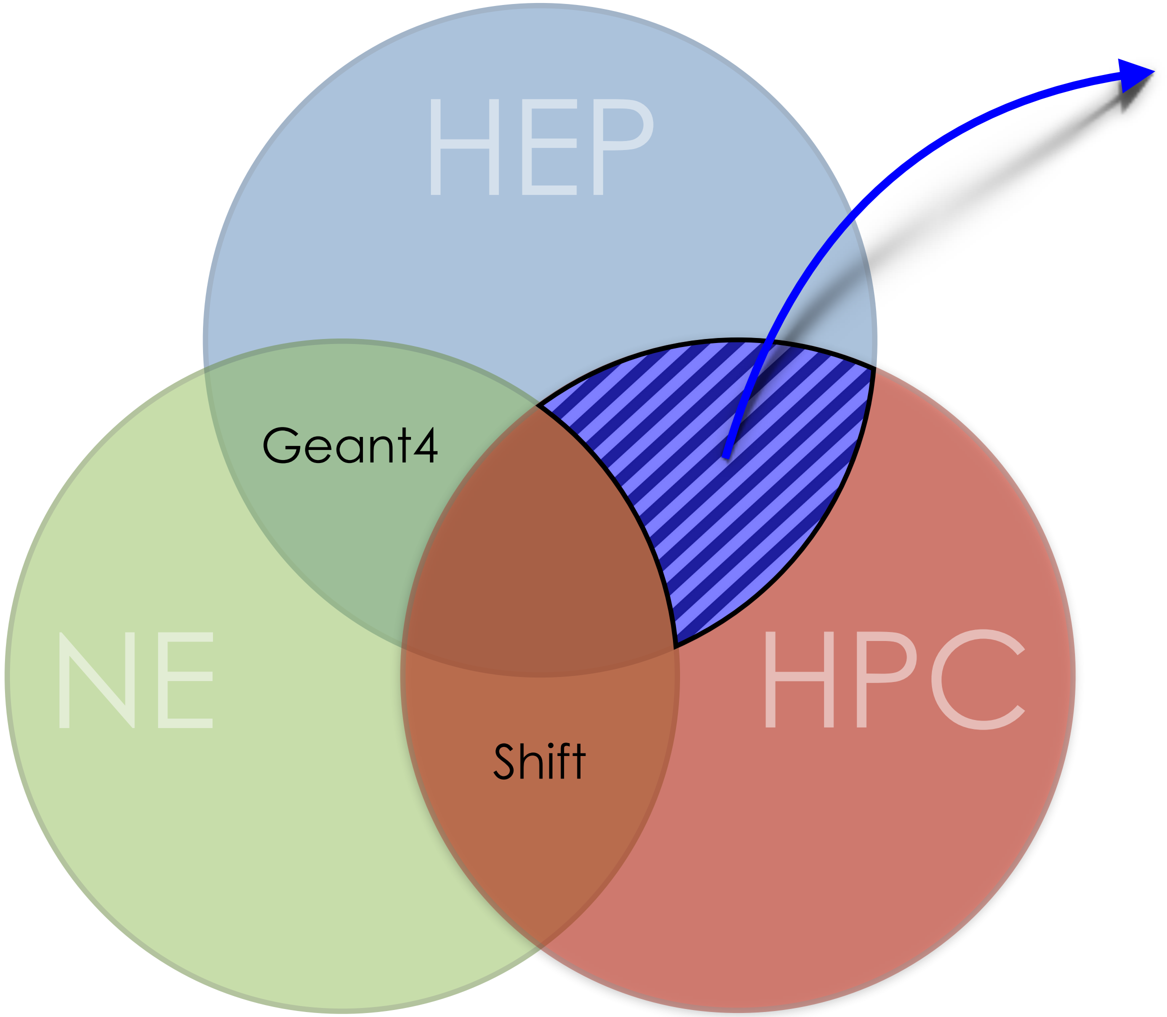

HEP ∩ HPC

- Previous efforts
  - GeantV
  - GeantX
- Current efforts
  - AdePT [github.com/apt-sim/AdePT]
  - Opticks [EPJ **214**, 02027 (2019)]

Opportunities

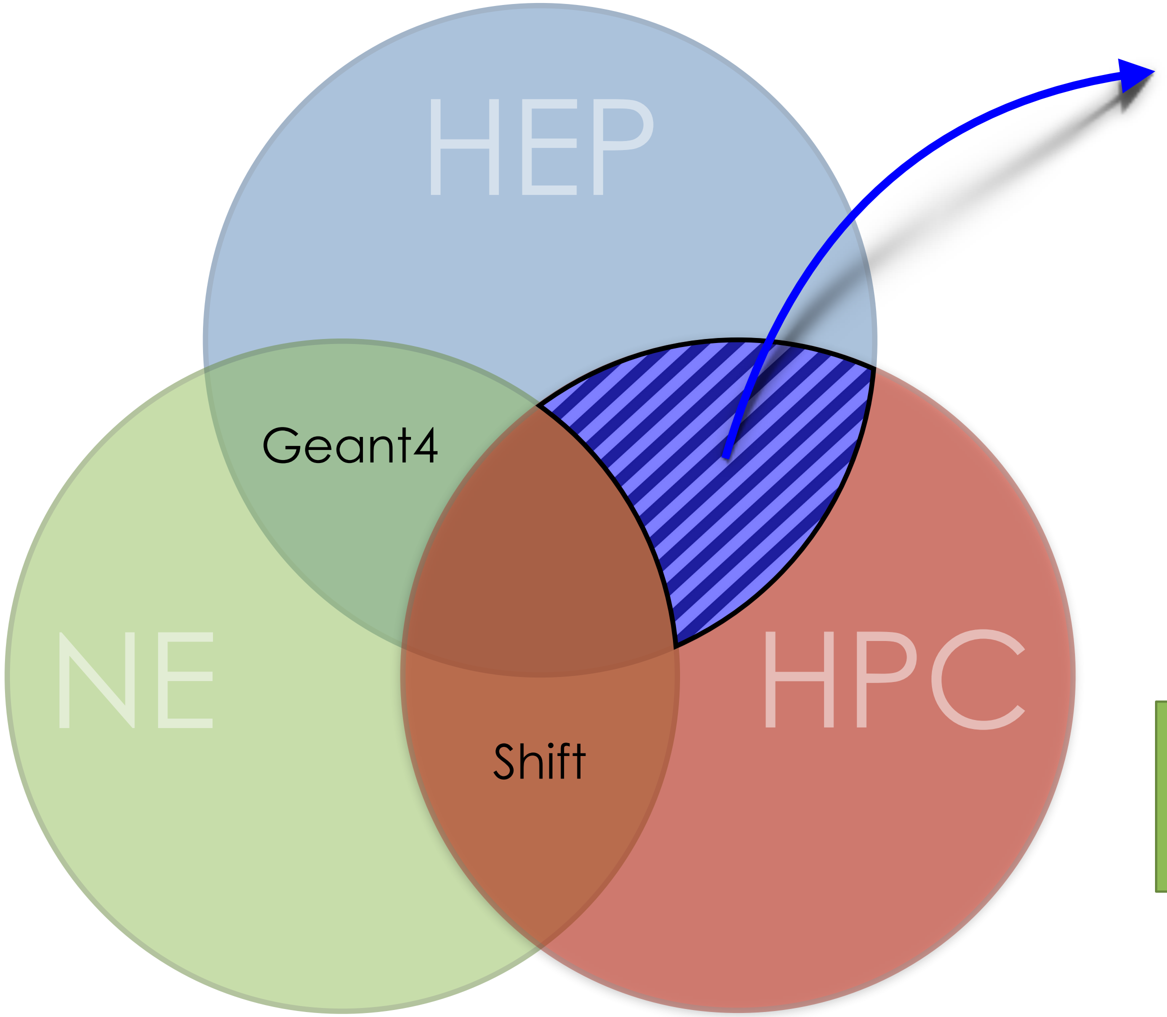

Impact

- **Summit (ORNL)**
  - 27,648 GPUs $\xrightarrow{\times 160}$ 4,423,680 CPUs
  - WLCG (2017): 500,000 CPUs*

Adding LCFs to HEP will vastly expand its current computing capacity

* EPJ Web of Conferences **245**, 07035 (2020)

Celeritas

github.com/celeritas-project

A GPU Monte Carlo detector simulation code for HEP

Will **NOT** replace Geant4, but could massively speed up HEP production runs

- Core team
  - **ORNL** Tom Evans, Seth Johnson, Stefano Tognini
  - **ANL** Paul Romano, Amanda Lund
  - **FNAL** Philippe Canal, Guilherme Lima, Soon Yung Jun
  - **BNL** Vincent Pascuzzi

Celeritas | Code architecture & challenges

- CPU vs. GPU programming

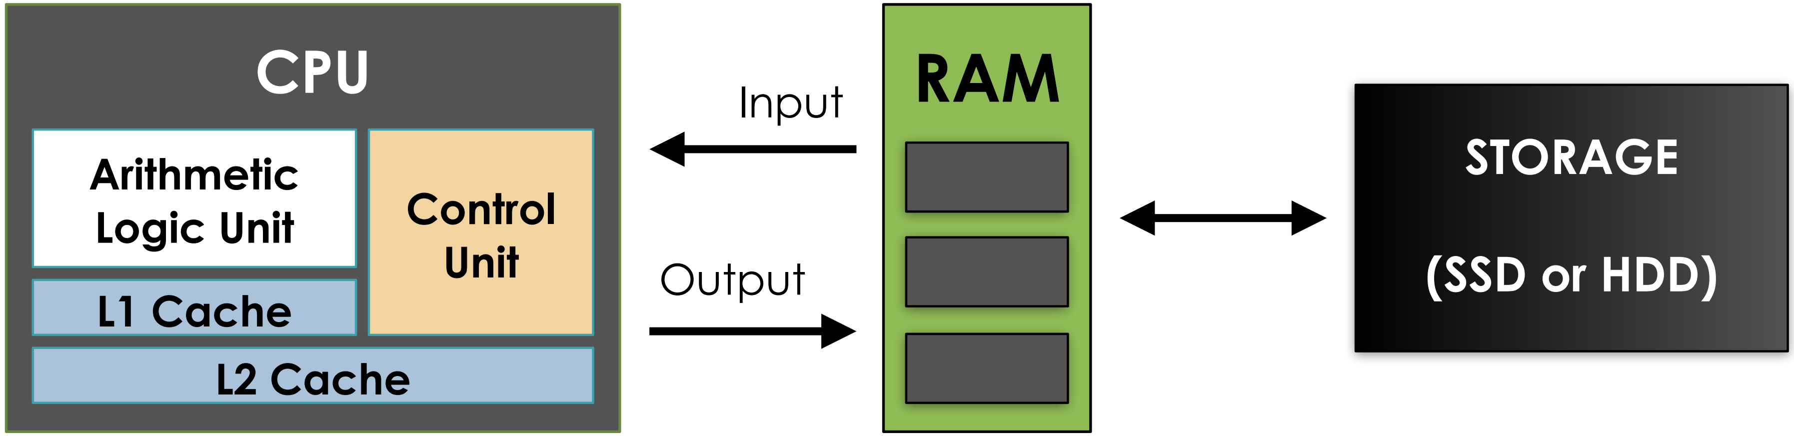

Celeritas | Code architecture & challenges

- CPU vs. GPU programming

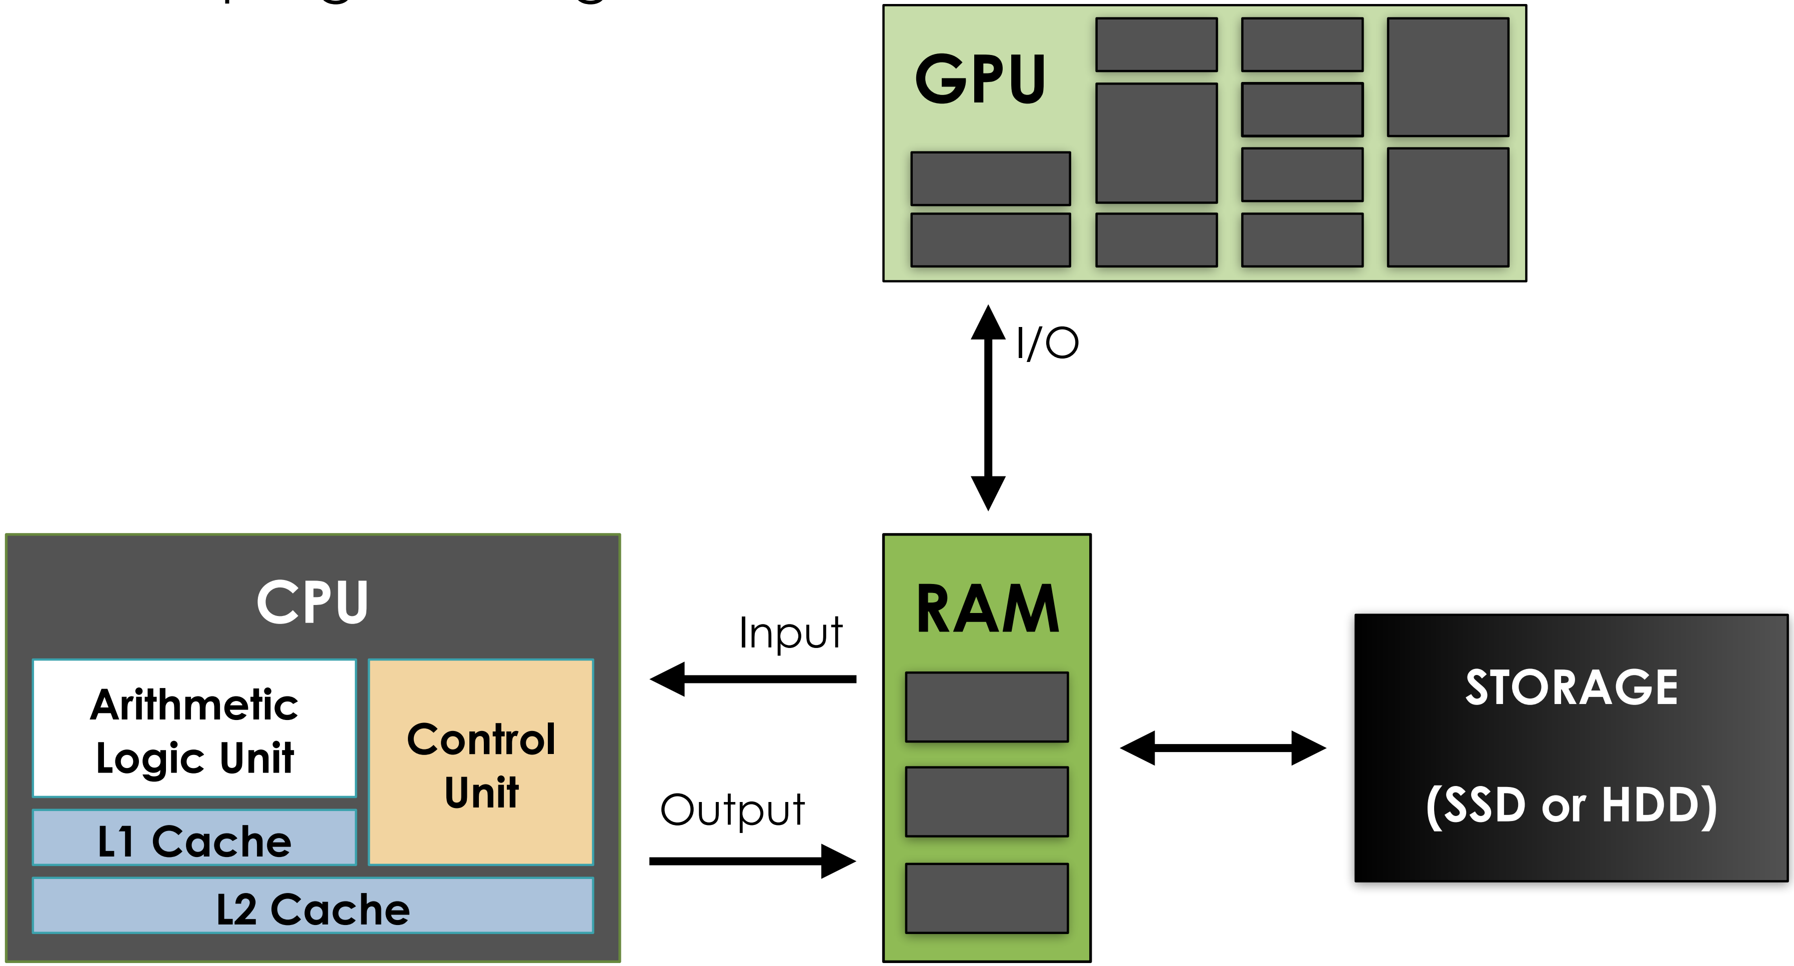

- CPU vs. GPU programming

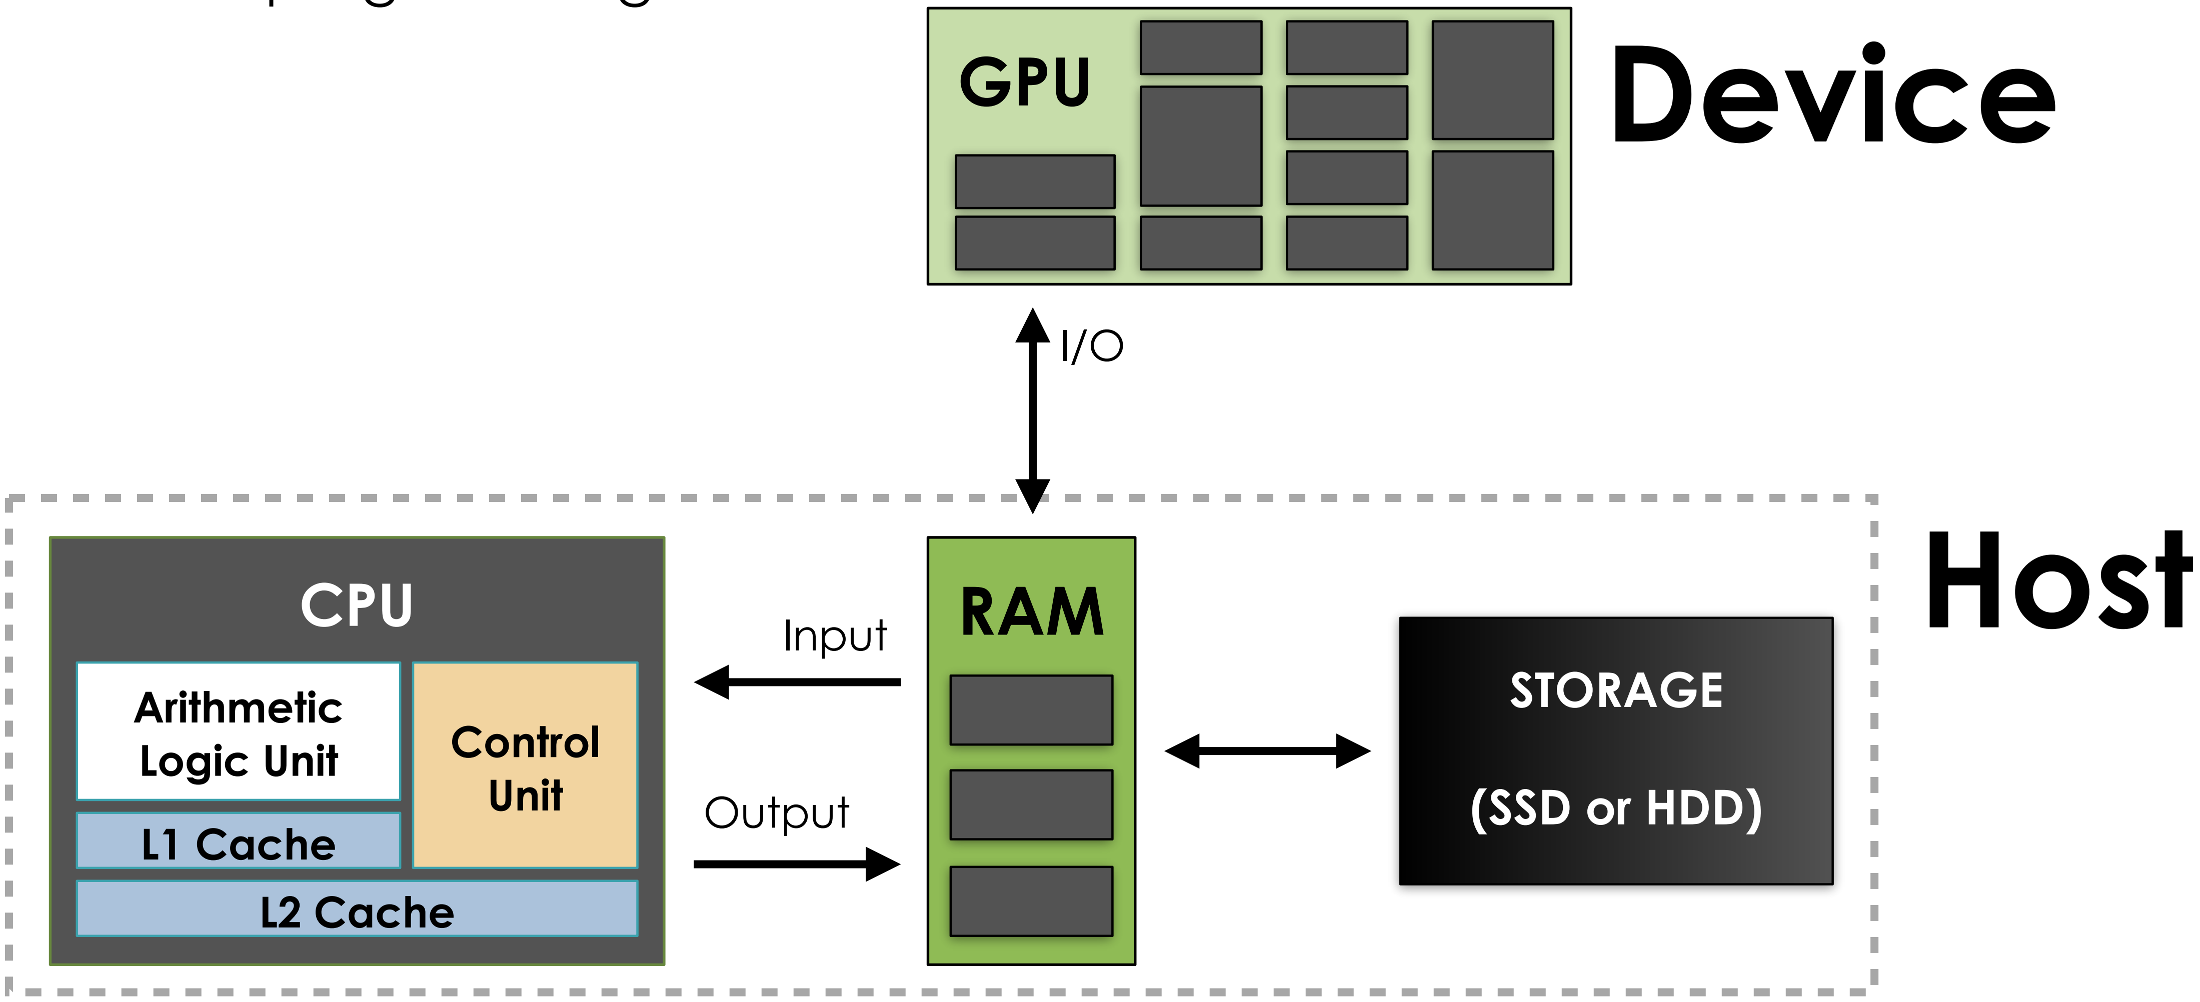

Celeritas | Code architecture & challenges

- CPU vs. GPU programming

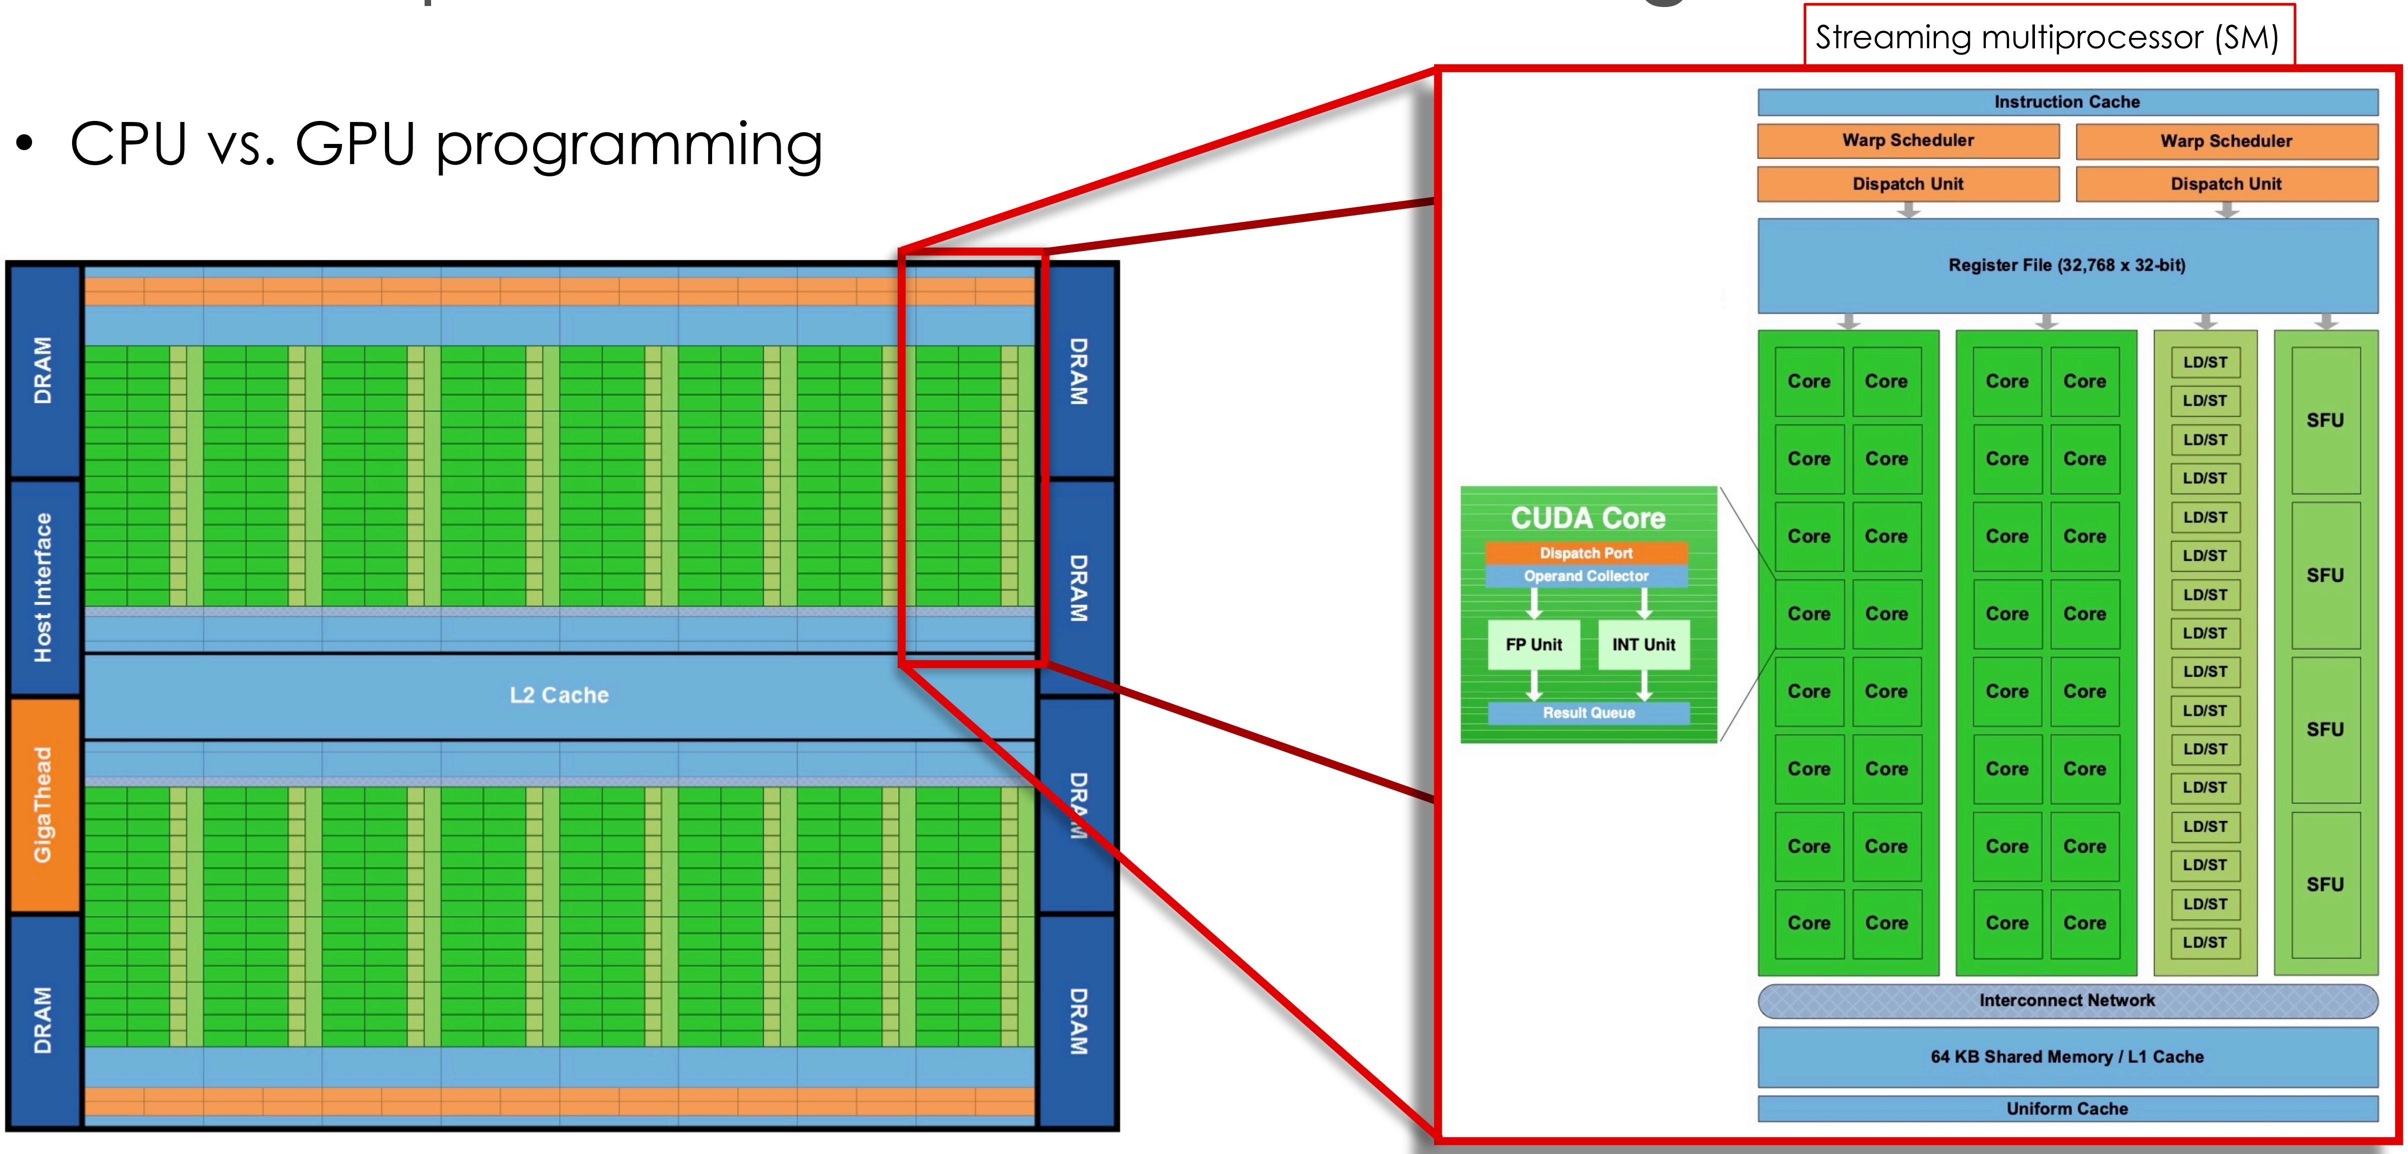

https://www.nvidia.com/content/PDF/fermi_white_papers/NVIDIA_Fermi_Compute_Architecture_Whitepaper.pdf

Celeritas | Code architecture & challenges

- CPU vs. GPU programming

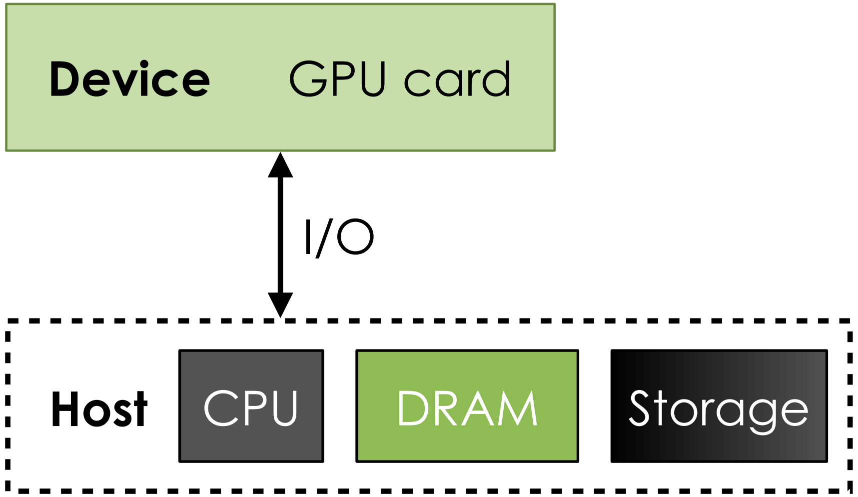

- Host and device memory are independent
- Host can read, but NOT edit data on device
- Host/device I/O is slow and non-trivial
- Device dynamic memory allocation is non-trivial
- Poor runtime polymorphism support
- Many libraries do not have a device-equivalent counterpart (e.g. std::string)

Celeritas | Code architecture & challenges

- CPU vs. GPU programming

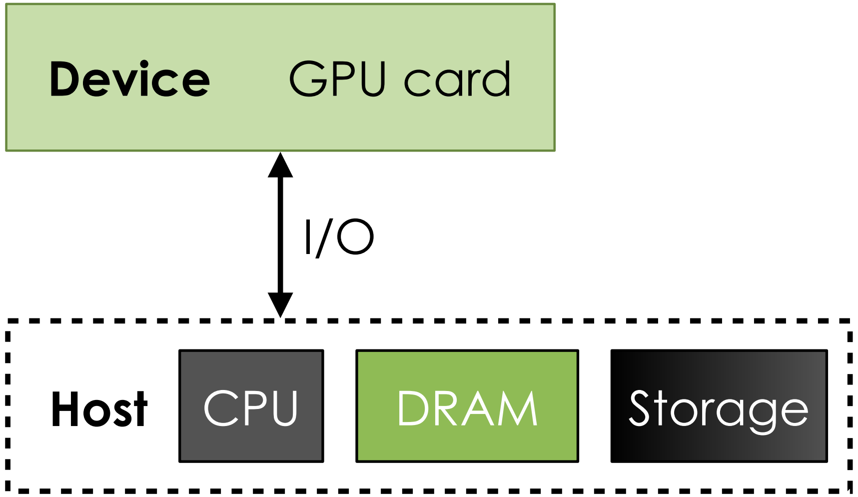

- Host and device memory are independent
- Host can read, but NOT edit data on device
- Host/device I/O is slow and non-trivial
- Device dynamic memory allocation is non-trivial
- Poor runtime polymorphism support
- Many libraries do not have a device-equivalent counterpart (e.g. std::string)

Geant4 is fundamentally built on these

Celeritas | Code architecture & challenges

- **Started almost from scratch**

- A LOT of development to do
  - Geometry import & navigation; physics models; XS data; EM fields; I/O
  - CPU and GPU compatible
  - Reproducible
  - Multiplatform (AMD, NVIDIA, Intel)
  - Scalable

Celeritas | Physics

Implemented

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
$\mu^{\pm}$	muon bremsstrahlung	Muon Bremsstrahlung

Complete validations are still ongoing

Planned

Physics	Process	Particle(s)
EM	photon conversion	γ
	pair annihilation	$e^{\pm}$
	photoelectric effect	γ
	ionization	charged leptons, hadrons, and ions
	bremsstrahlung	charged leptons and hadrons
	Rayleigh scattering	γ
	Compton scattering	γ
	Coulomb scattering	charged leptons, hadrons
	multiple scattering	charged leptons, hadrons
	continuous energy loss	charged leptons, hadrons, and ions
Decay	two body decay	$\mu^{\pm}$, $\tau^{\pm}$, hadrons
	three body decay	$\mu^{\pm}$, $\tau^{\pm}$, hadrons
	n-body decay	$\mu^{\pm}$, $\tau^{\pm}$, hadrons
Hadronic	photon-nucleus	γ
	lepton-nucleus	leptons
	nucleon-nucleon	p, n
	hadron-nucleon	hadrons
	hadron-nucleus	hadrons
	nucleus-nucleus	hadrons

Celeritas | Preliminary results

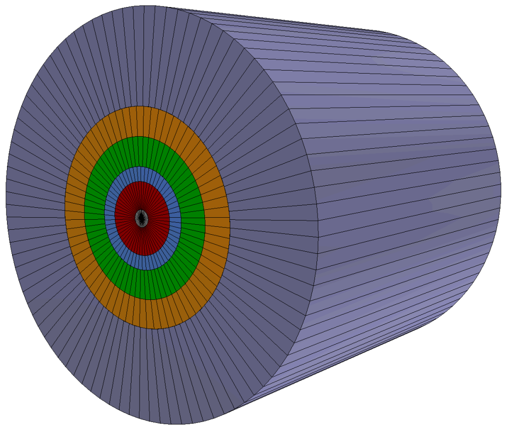

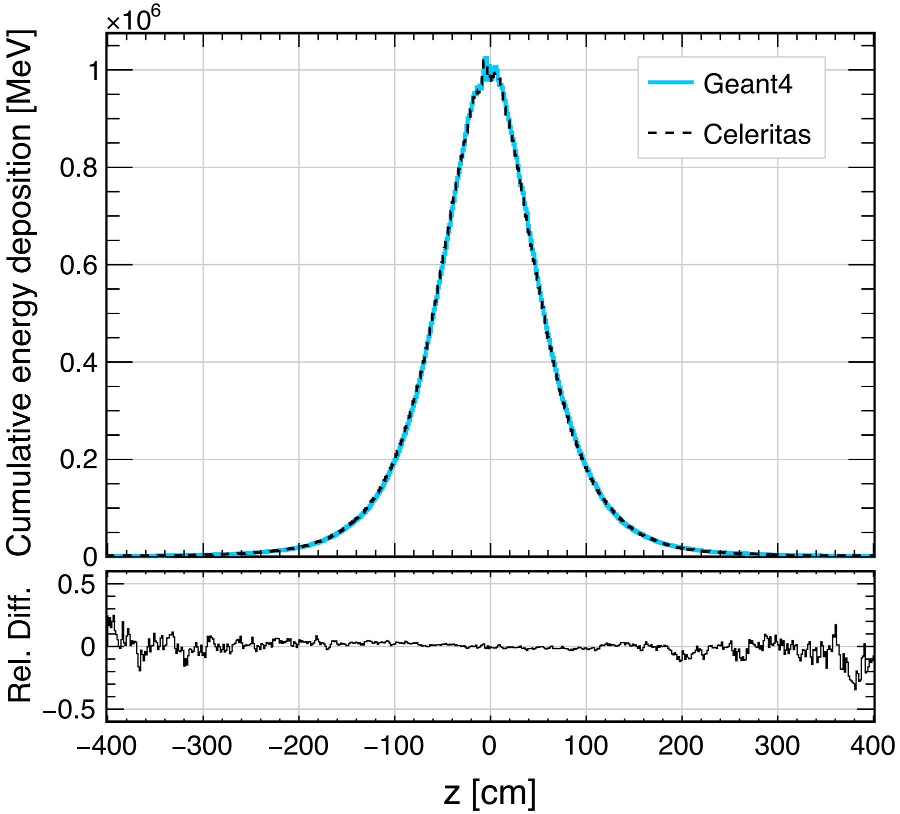

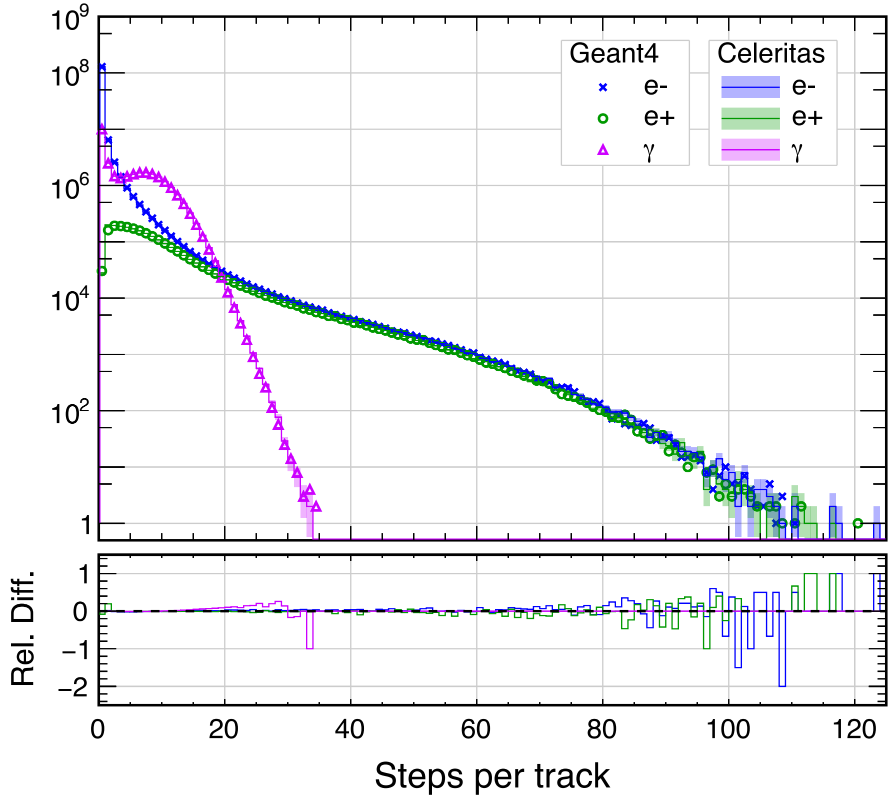

- ~~Spherical~~ Cylindrical cow in a vacuum version of CMS

- Single-element concentric cylinders of **Si**, **Pb**, **C**, **Ti**, and **Fe**
- Uses all implemented physics for e± and γ

Isotropic source
100k photon primaries
1 GeV each
Vertex at the origin

Celeritas | Preliminary results

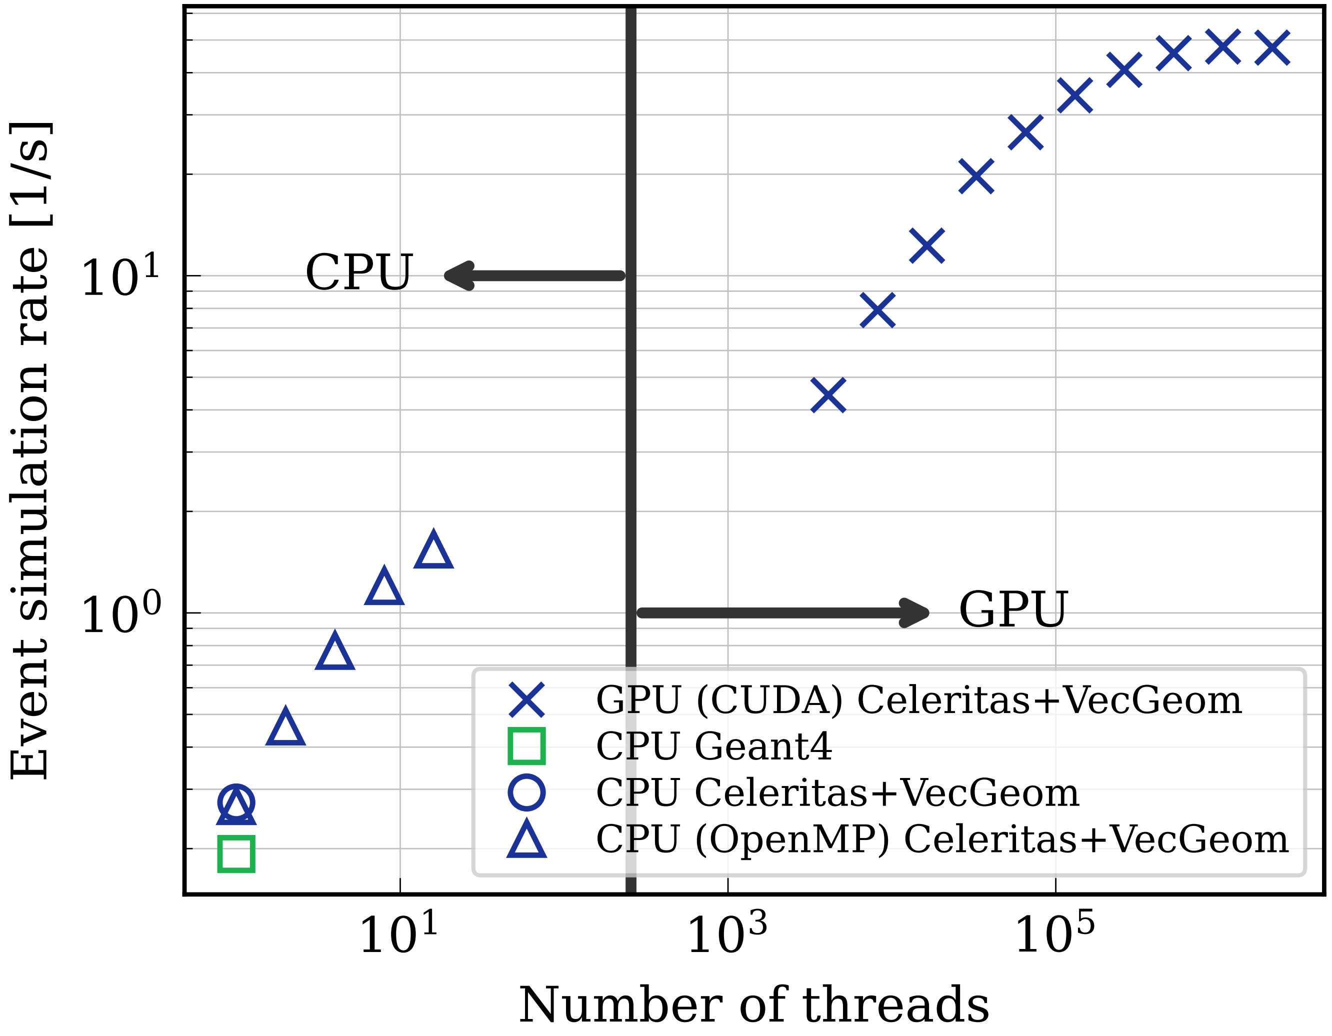


```
CUDA 11.5   -O3 --use_fast_math
GCC 8.5     -O3 -march=skylake-avx512 -mtune=skylake-avx512
```

Application	Execution	Speedup
Geant4 (v10.7)	CPU (serial)	1
Celeritas	CPU (serial)	1.4
Celeritas	CPU (OpenMP)	40
Celeritas	GPU	280

- **Geant4 scales linearly**
- 30 cores ≈ 30× serial execution
- Single NVIDIA V100 ≈ 280 cores

CPU Intel Xeon Gold 5218 @ 2.3 GHz (Cascade Lake)
GPU NVIDIA V100 @ 1.53 GHz (80 symmetric multiproc. 64 cores each 16 GB of memory)

OAK RIDGE National Laboratory

Celeritas | Integration paths and challenges

Acceleritas

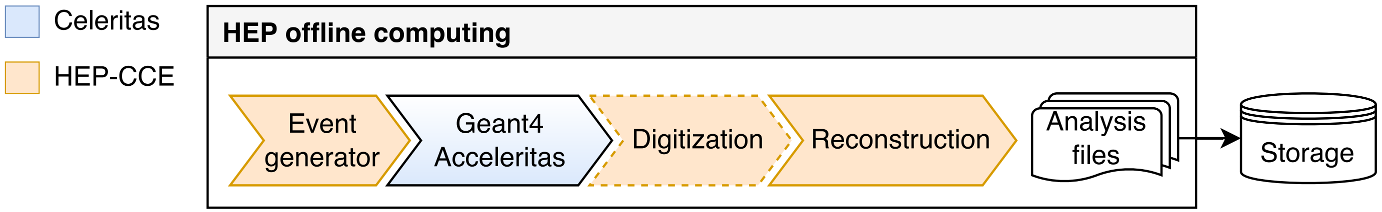

- Acceleritas library provides a streamline integration with relatively small changes
  - Cons:
    - Considerably smaller performance impact
    - Many (most?) offline HEP working nodes do not have dedicated GPU hardware

Celeritas | Integration paths and challenges

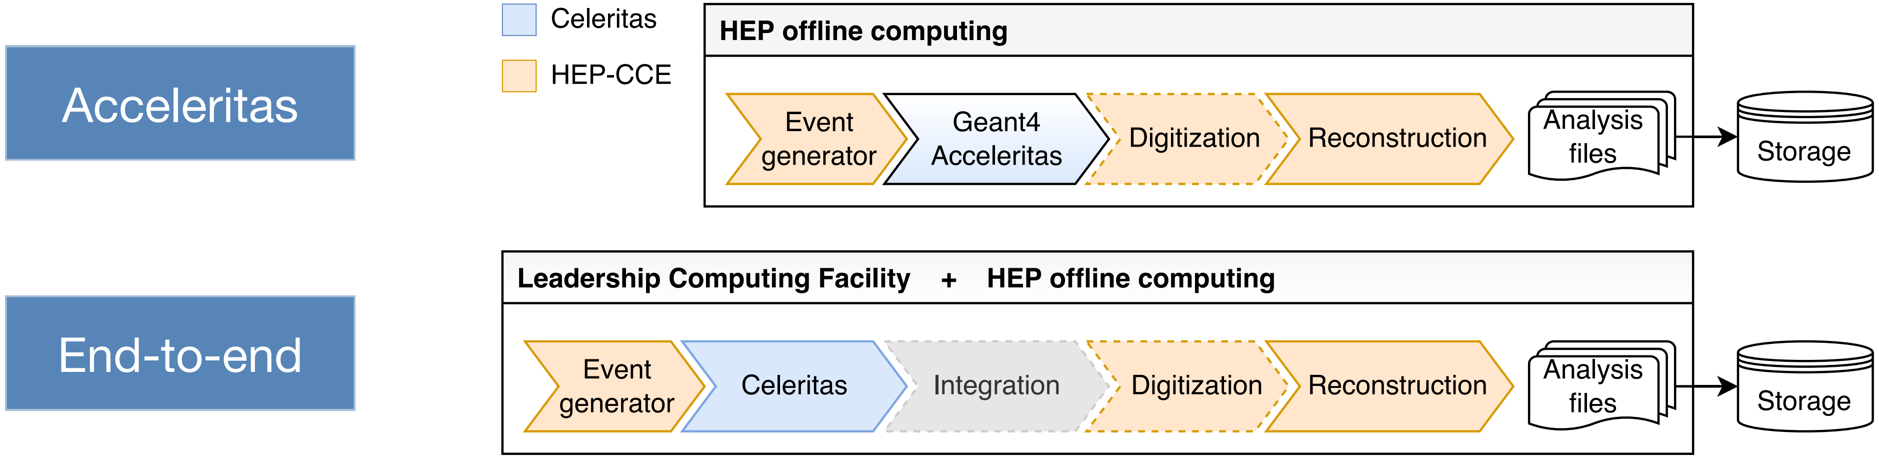

- End-to-end is envisioned after decay + hadronic physics are available
- Integration between LCFs and HEP-EX is a gray area still
  - I/O bottlenecks; ROOT integration; what is processed where
  - LCFs are mostly GPU (Summit is >95%); is local CPU post-processing worth it?
  - Network transfers can become another bottleneck; ...

Lessons learned

- HPC vs. HEP-Ex workflows are drastically different
- Mindset change: physics results vs. code performance & quality
- Lots of learning: CMake, C++, CUDA, HIP, GIT, CI, QA unit-tests...
  - Lots of new jargons
- Networking goes both ways (you're an outsider; end up meeting new people)

Backup

HEP computing challenges

- Detector triggering and reconstruction
- Event generators
- **Detector simulation** our focus
- Frameworks
- Data analysis
- Software dev. tools and packaging

- HEP Software Foundation [hepsoftwarefoundation.org]
- Snowmass [snowmass21.org]

HSF HEP Software Foundation

Snowmass 2021

Oak Ridge National Laboratory

HEP computing challenges | DUNE

- **Signal processing**
  - Noise filtering
  - Triggering and data processing/compression (raw events: 6 GB – 115 PB)
- **MC**
  - EM showers are a big bottleneck in LArTPCs
- **Reconstruction** [FERMILAB-CONF-20-074-SCD]
  - Hit based (e.g. Pandora); Image based (CNNs); hybrid
  - Well suited for AI/ML and GPUs

HEP computing challenges | LHC

CMS

https://twiki.cern.ch/twiki/bin/view/CMSPublic/CMSOfflineComputingResults

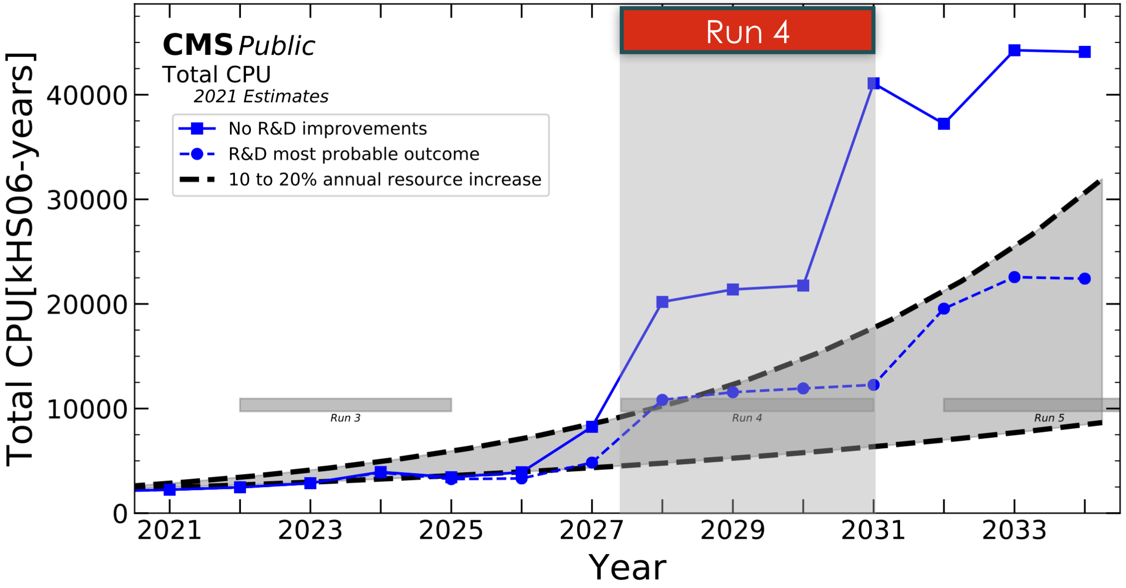

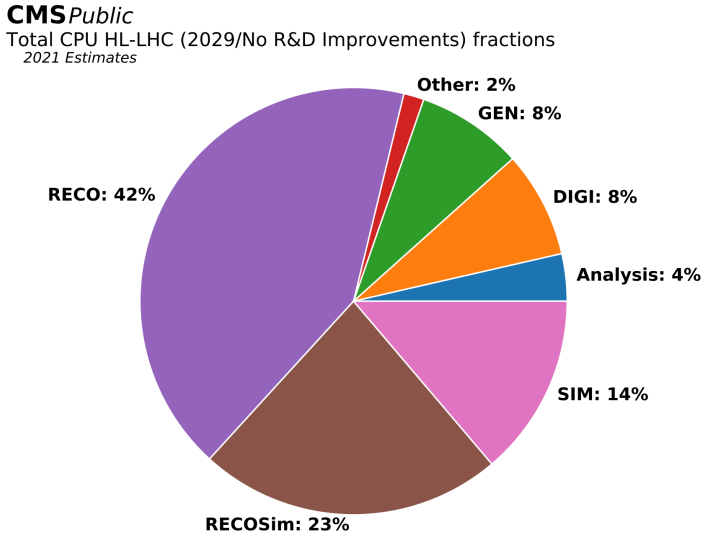

- MC: 14%
- Reco:
  - 42% (data)
  - 65% (data + MC)

- Reconstruction takes 65% of CPU time
- Sim takes 14% — or 42% of what's left after reconstruction

Celeritas | Partnerships

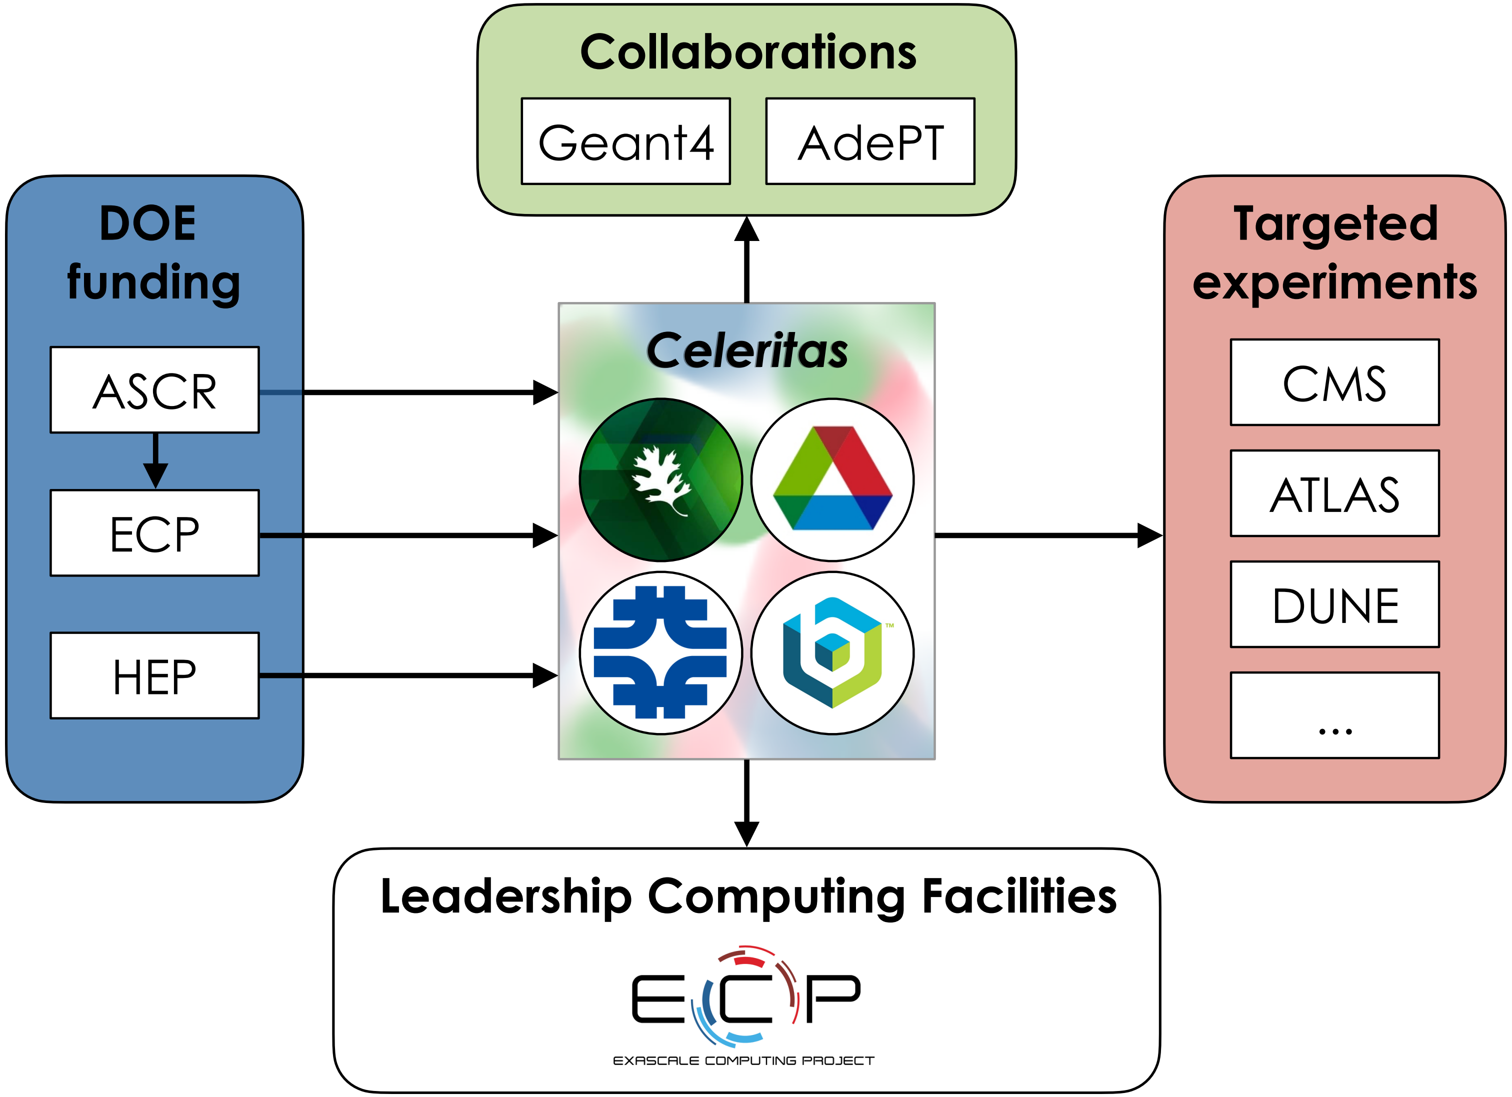

Celeritas | Dependencies

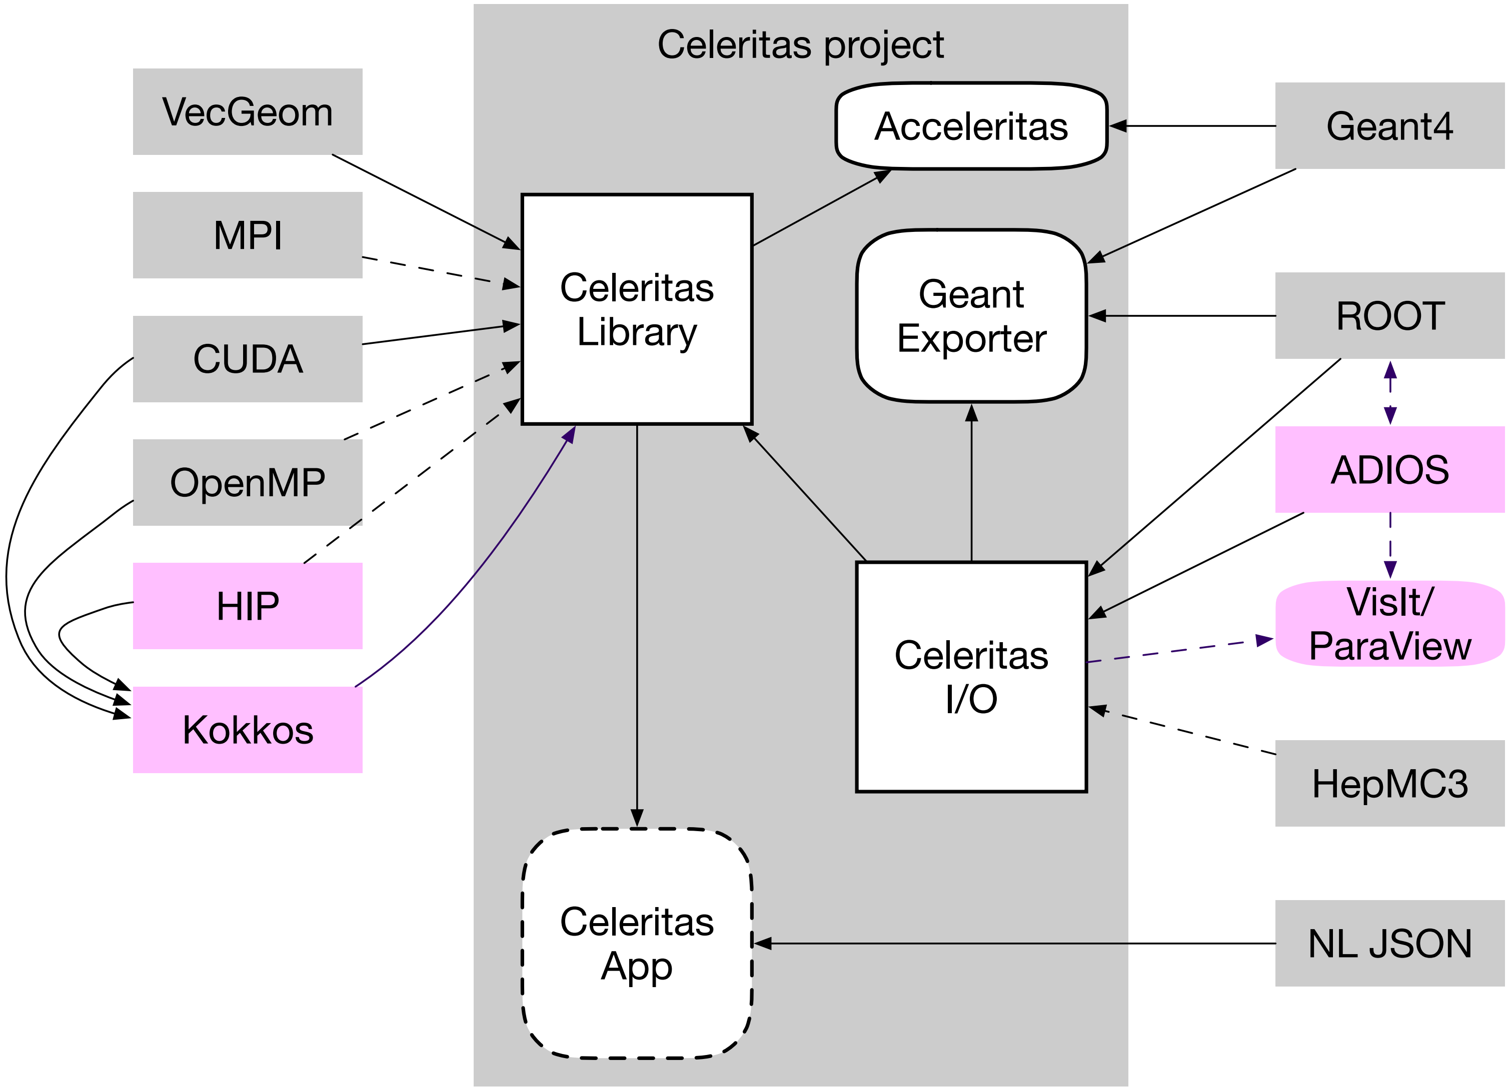

Celeritas | Code architecture & challenges

- GPU vs. CPU programming

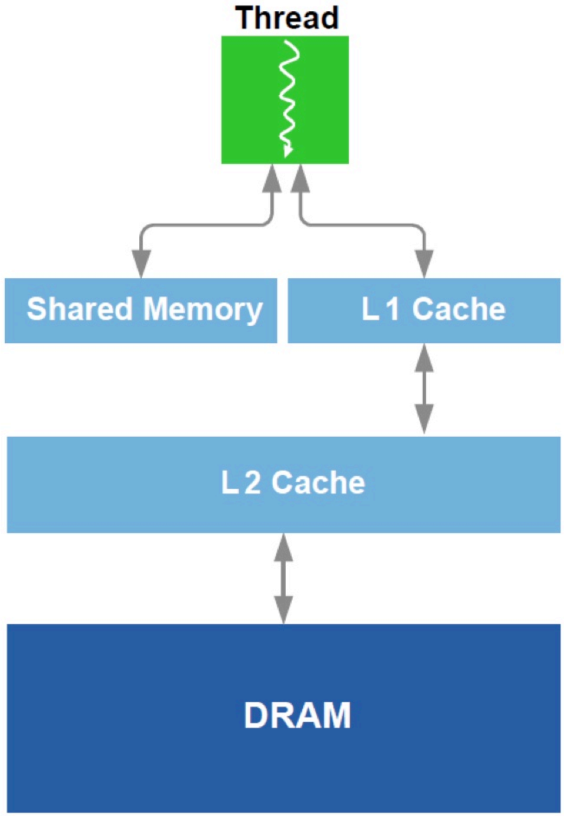

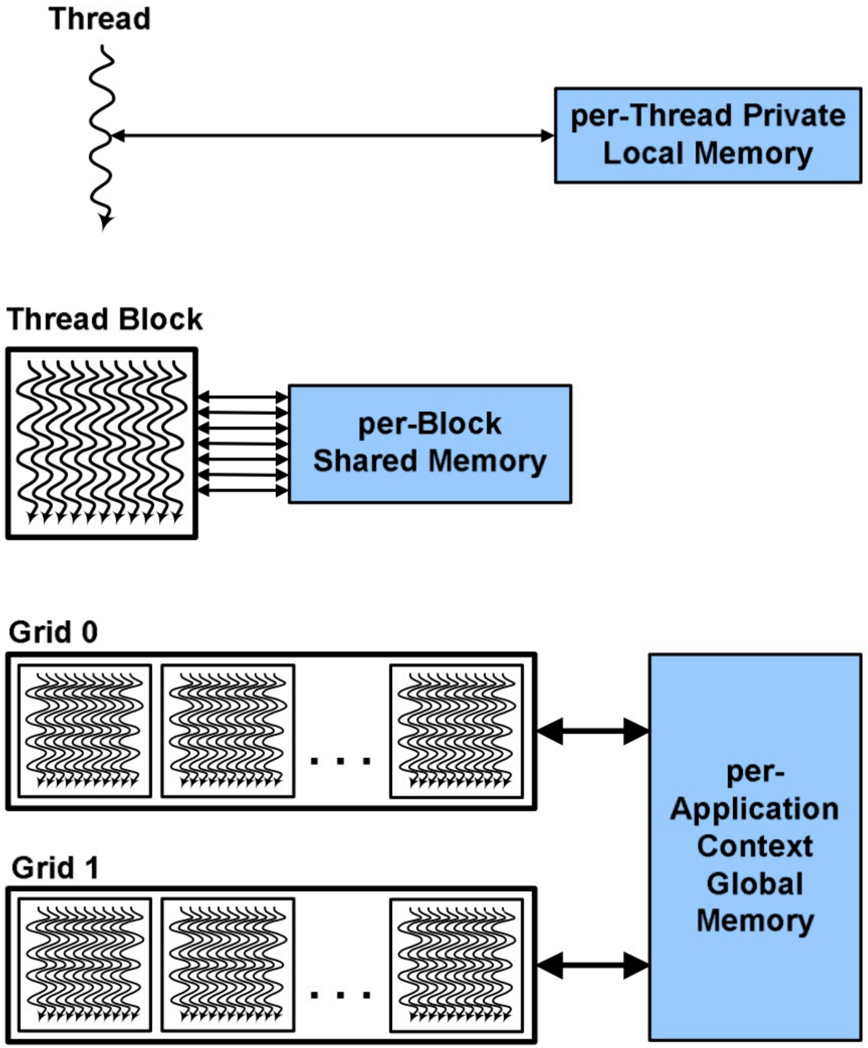

https://www.nvidia.com/content/PDF/fermi_white_papers/NVIDIA_Fermi_Compute_Architecture_Whitepaper.pdf

Celeritas | Geometry navigation

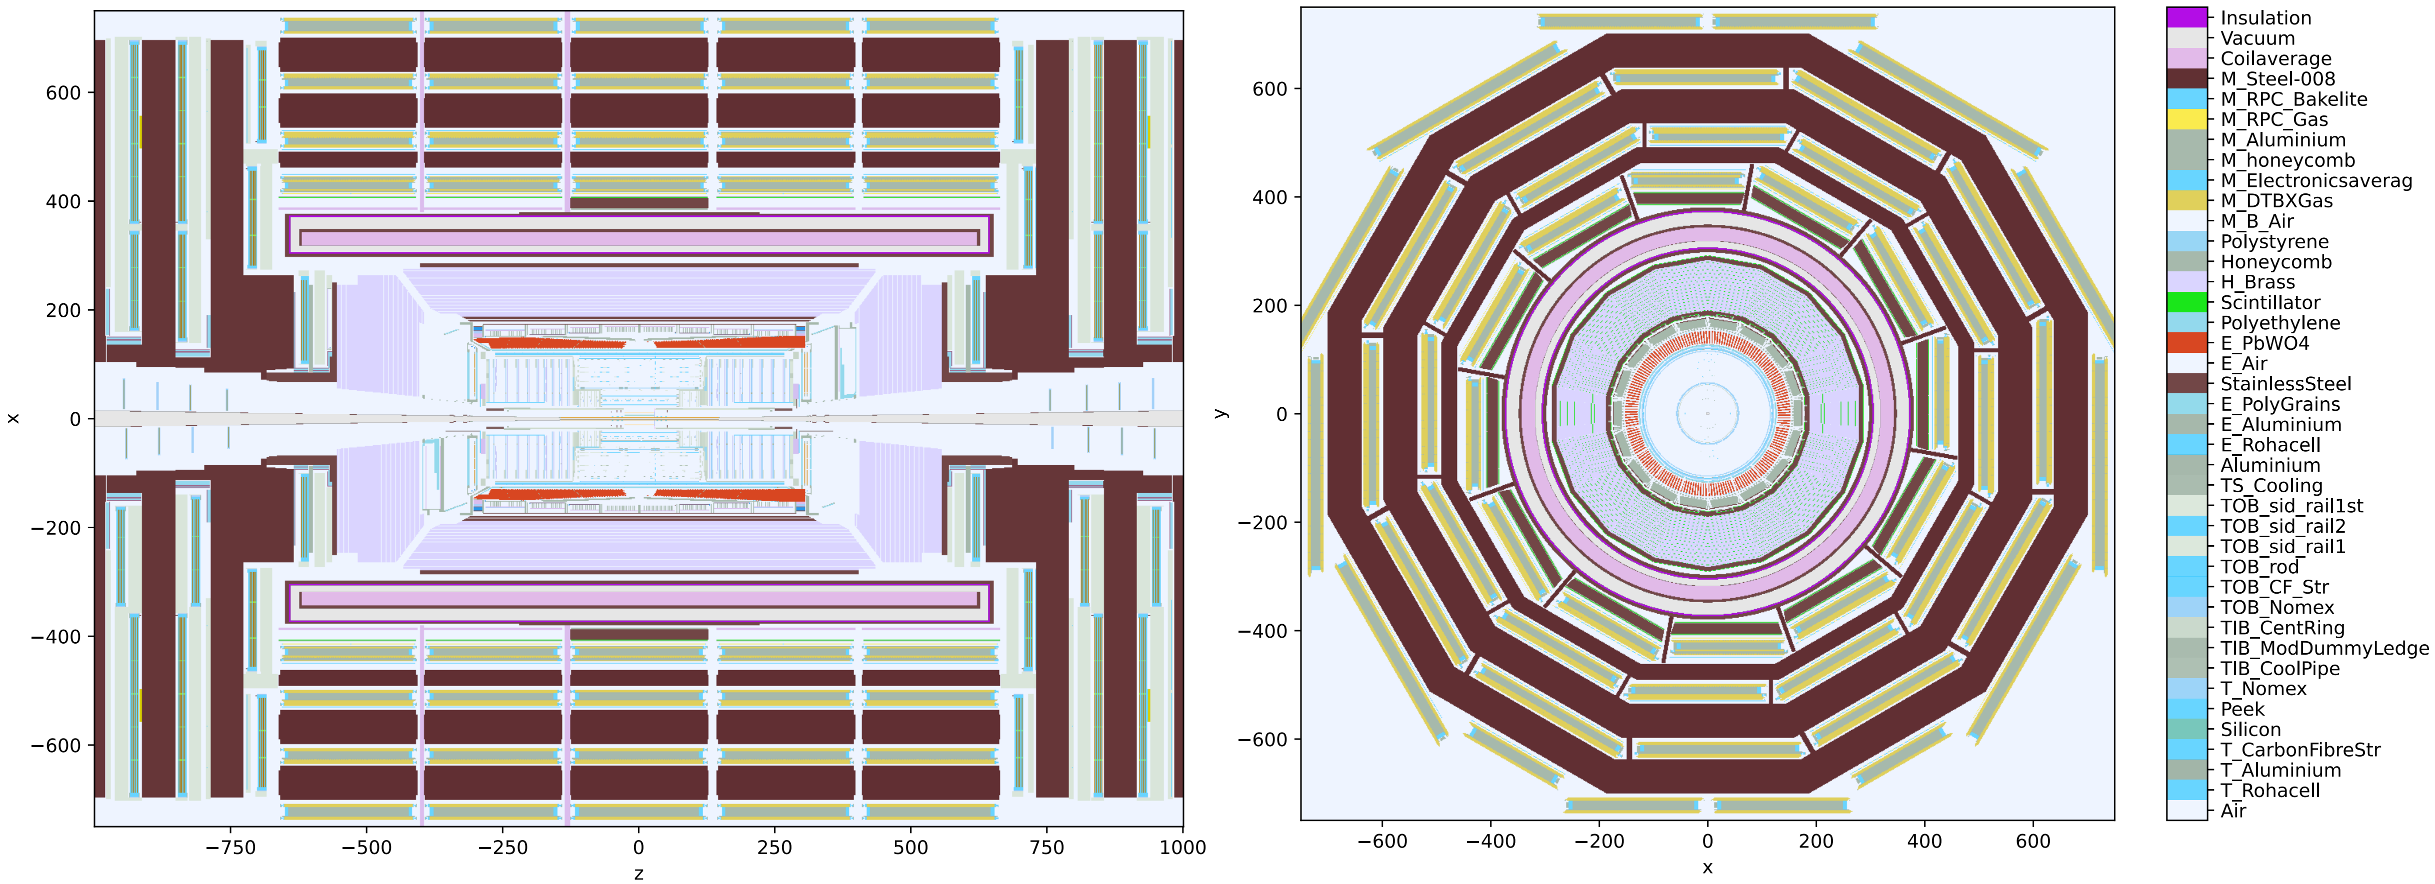

- VecGeom
- ORANGE (Oak Ridge Adaptable Nested Geometry Engine)

Celeritas | Preliminary results