Detector Simulation in SWIFT-HEP

Ben Morgan

WARWICK
THE UNIVERSITY OF WARWICK

Detector Simulation R&D @ November 2023

- **AdePT Project (CERN-SFT)**
  - *https://github.com/apt-sim*
- **Celeritas Project (ECP: ORNL, FNAL, Argonne, LBL)**
  - *https://github.com/celeritas-project*
- **Vecgeom/ORANGE Surface Based Geometry (CERN, Celeritas/ORNL)**
  - *https://gitlab.cern.ch/VecGeom/VecGeom (See surface_model branch)*
  - *https://github.com/celeritas-project/celeritas/tree/develop/src/orange*
  - ***ExaTEPP grant from UKRI ExCALIBUR enabling contribution here***
- Short overview of status today - follow links to recent presentations at Geant4 Collaboration Workshop (1, 2) for in depth details

Reminder on Objectives of AdePT and Celeritas

- ***Understand usability of GPUs for general particle transport simulation***
  - *Prototype e+/e−/γ EM shower simulation on GPU, evolve to realistic use-cases*
  - *Focus on **EM physics** given computational cost in HEP workflows*
- ***Implement GPU-targeted components for physics, geometry, field, with data models and workflow***
  - *Integrate components in a **hybrid CPU-GPU Geant4 workflow ("Fast Sim" approach)***
  - *Offload tracks to GPU/CPU when preconditions like particle type or geometric region met*
  - *Most realistic short-term objective to allow testing/use in existing experiment code*
- **Ensure correctness and reproducibility**
  - *Validate GPU-only, CPU+GPU off/onload against pure CPU Geant4*
- **Understand bottlenecks and blockers limiting performance**
  - *Feasibility and future effort required for efficient simulation workflows on GPU*
- **Celeritas also have a longer term objective to include full hadronic, optical photon physics**
  - ***See Seth's talk***

Track-parallel Stepping Workflow on GPUs

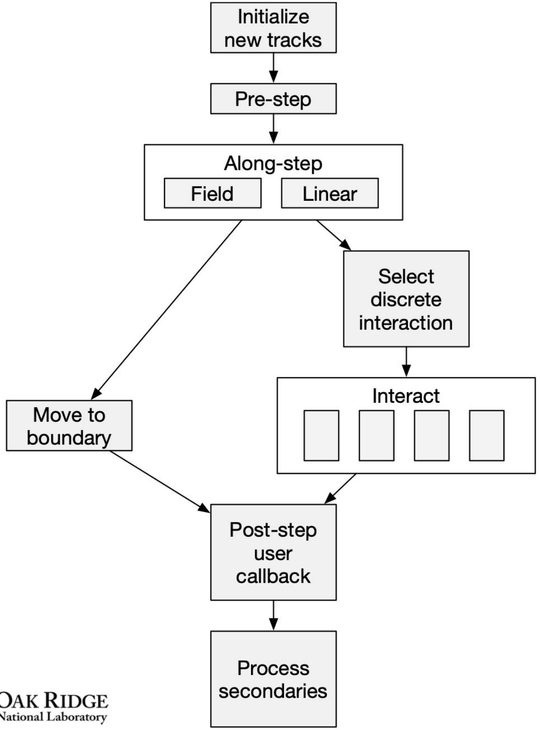


```
extend_from_primaries          ▷ Copy primaries to device, create track initializers
while Tracks are alive do
    initialize_tracks                          ▷ Create new tracks in empty slots
    pre_step                     ▷ Sample mean free path, calculate step limits
    along_step                                 ▷ Propagation, slowing down
    boundary                                   ▷ Cross a geometry boundary
    discrete_select                            ▷ Discrete model selection
    launch_models              ▷ Launch interaction kernels for applicable models
    extend_from_secondaries         ▷ Create track initializers from secondaries
end while
```

- CPU: parallel Events, *sequential* Tracks
- CPU+GPU: parallel Events, *parallel* Tracks
  - ***Action*** *based control flow*
  - *Kernels determine next* ***Action****, or perform an* ***Interaction***
  - *Example from Celeritas, AdePT's is similar though with larger, per-particle, kernels*

OAK RIDGE National Laboratory

Credit: *[Seth Johnson (ORNL)]*

Implementing Geant4 CPU+GPU hybrid applications

- AdePT and Celeritas only target e-/e+/g physics at present, so cannot be used standalone for simulating a **full** EM+hadronic experiment
- Instead, use them as a “service” to offload tracks to the GPU according to preconditions such as particle type, or geometric region.
  - ***Basically the same as “Fast Simulation” methods***
- Use Geant4 Fast Simulation or other hooks allowing Track modification, but all have same challenges:
  - *Minimizing number/size of on/offload actions*
  - *Allow user-defined actions on GPU, such as scoring/hits*
  - *Synchronization between CPU/GPU (event boundaries)*
  - *Handing back particles (e.g. exiting particles, hadrons from photonuclear processes) from GPU to CPU*

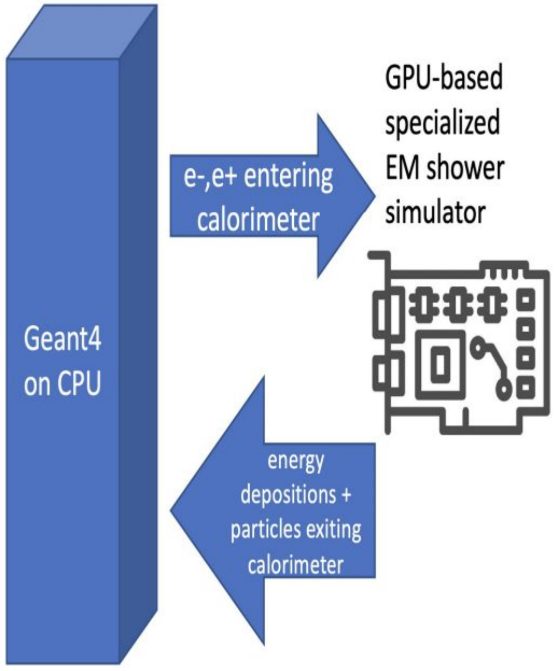

Physics Validation

- G4HepEM in AdePT, CPU/GPU implementation of Geant4 models/data in Celeritas.
- ***Excellent agreement with Geant4, but studies ongoing across problem space***

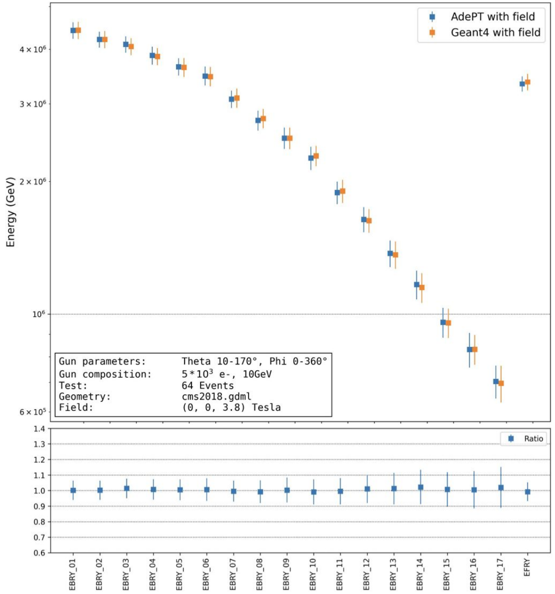

Example: CMS ECAL
← AdePT: 10GeV e⁻

Celeritas: 14TeV $t\bar{t}$ →

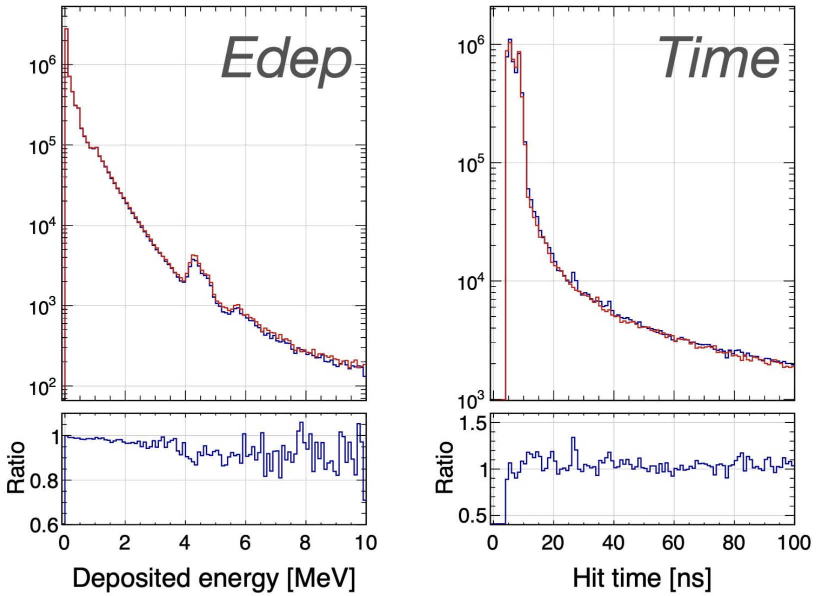

Credits: Jonas Hahnfeld (CERN), Seth Johnson (ORNL), Amanda Lund (Argonne)

Some Benchmarking Results: AdePT w/CMS2018

- Speedup of the ECAL simulation and overall event speedup (10GeV e-)
  - *AdePT does not affect the rest of the simulation, 1:1 ratio in the time spent outside ECAL*
- Vary number of Geant4 worker threads
  - *Decreasing AdePT speedup as the GPU becomes more saturated*

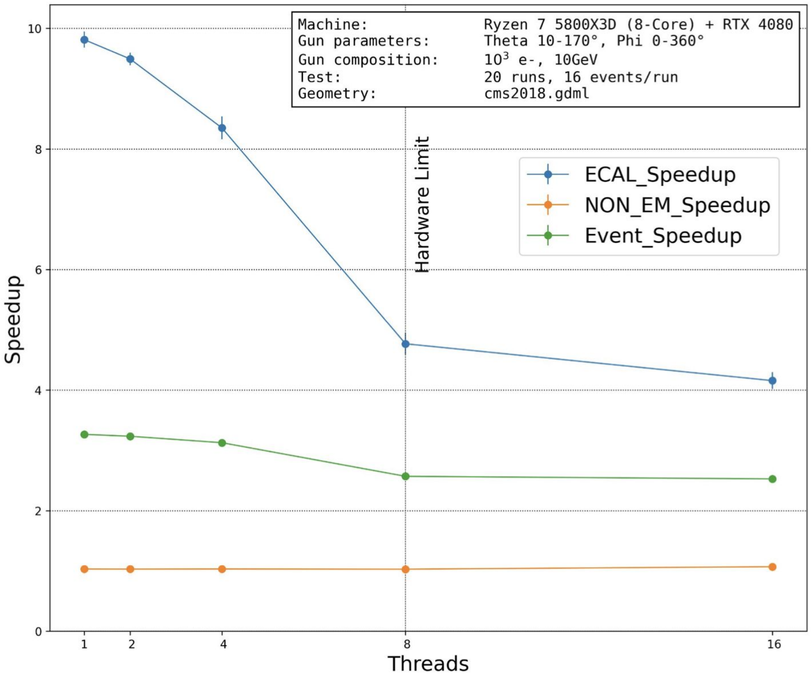

Speedup of the ECAL simulation and overall per-event speedup

Credit: [Jonas Hahnfeld (CERN)]

Some Benchmarking Results: Celeritas w/CMS Run3

- Initial performance comparison in standalone Geant4+Celeritas application
  - *CMS GDML geometry/SDs*
- 8CPU+1GPU standalone simulation with 14TeV tt **17-87% faster**
  - *Theoretical maximum speedup (all e^-/e^+/g tracks take zero time) in full CMSSW ~230%*

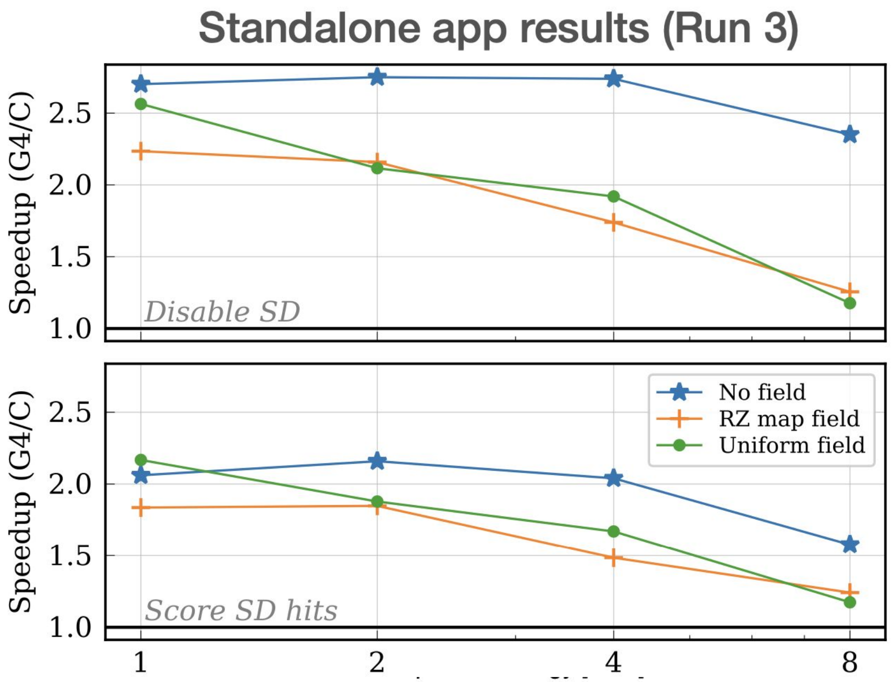

Hardware: Intel Xeon Gold 6152 CPU 22c 2.10GHz + NVIDIA Tesla V100 SXM2
Geometry: CMS detector (Run 3 configuration)
Input: 8 $t\bar{t}$ events @ 14 TeV from LHC pp collision

Credits: *Seth Johnson (ORNL), Amanda Lund (Argonne)*

Surface-based Geometry: VecGeom, ORANGE

- Current CSG model of VecGeom known bottleneck for GPU kernels in AdePT and Celeritas
  - *Divergence from different algorithmic complexity in different solids, etc*
- Effort to develop and use surface- based geometry models, navigation
  - *Reduce divergence from smaller number of surfaces, simpler algorithms*
  - ***VecGeom**: bounded surfaces (explore potential for work reduction in LHC-complexity geometries by reducing checks on "virtual" crossings)*
  - ***ORANGE**: unbounded surfaces (approach from nuclear engineering codes for reactor geometries)*
- Defer to the following presentations at the Geant4 Collaboration meeting for details of developments and results:
  - *Surface-based GPU model in VecGeom, Andrei Gheata et al*
  - *ORANGE surface geometry progress, Seth Johnson et al*

Integration/Testing in Experiments: ATLAS

- TileCal test beam standalone Geant4 application as testbed
    - *Code: https://github.com/lopezzot/ATLTileCalTB (see presentation)*
- **AdePT**: Initial integration by Davide:
    - *See https://indico.cern.ch/event/1215829/contributions/5306569/*
- **Celeritas**: integration both standalone and as FullSimLight plugin
    - *Ongoing work to integrate AdePT as FullSimLight plugin, using new functionality to score on host side to allow direct comparison with pure Geant4 and Celeritas integration*
- New EMEC geometry compatible with VecGeom on GPU now available
    - *Second stage testbed in FullSimLight, allowing full ATLAS geometry to be tested before bigger integration into Athena*
- Topics being worked on in coordination with ATLAS Full Simulation WG.

Integration/Testing in Experiments: CMS

- **AdePT**: G4HepEM on CPU integrated as optional physics list, permitting direct comparison against AdePT physics
  - *Validation and additional options in progress, discussions on scoring requirements for HGCal*
- **Celeritas**: integration in CMSSW working with:
  - *Complete offload of EM particles to GPU*
  - *R-Z field map preprocessed*
  - *Reconstruction of tracks hitting SDs on host side for scoring*
- Work and validation ongoing together with CMS Simulation, some caveats
  - *No support for MC truth*
  - *Celeritas reproduces "standard" Geant4 models/data, but CMSSW has many fine-grained tunes to physics/tracking compared to this*

Integration/Testing in Experiments: LHCb

- **AdePT**: Combine standalone application example with Gauss-on-Gaussino machinery
  - *Fill AdePT pipeline with particles entering LHCb Calo region*
  - *Generate Gauss hits from AdePT (to give equivalence with plain Geant4)*
  - *Working with Juan Bernardo Benavides (LHCb Doctoral Student)*
- **Celeritas**: **Request to the UK community for help/time or contacts within LHCb to try out integration!**

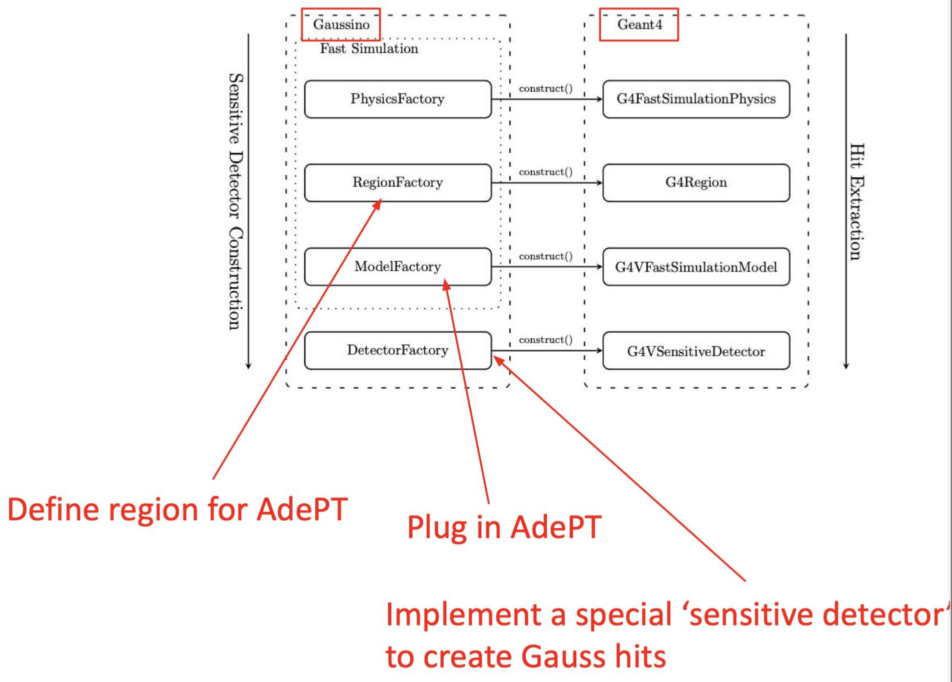

Geant4 Review of GPU R&D Projects

- With a huge amount of progress by both projects, and growing test program with experiments, it was felt a good opportunity for a technical review with a broader panel of Geant4/Simulation experts
  - *@CERN December 13-14th*
  - *Forum for presentation of status and results, discussion of technical topics between projects and experts*
- **Another key item will be finding common code/patterns between Celeritas/AdePT with the aim of converging towards a single project**
  - *Meetings and hackathons planned during review week on specific areas such as physics, workflows, and geometry*
  - *Initial discussions/ideas already in progress through monthly stand ups and other channels*

Summary

- AdePT and Celeritas continuing to demonstrate feasibility of detector simulation on GPU
  - *Near full EM physics validated*
  - ***Working with ATLAS, CMS, LHCb(*) on integration, benchmarking, and scoring in their frameworks.***
  - ***(*) Any UK interest in testing Celeritas in LHCb?***
- GPU friendly geometry modeling/navigation using surface models is an **ongoing task that will be key for performance**
- **Technical review in December, together with meetings/hackathons to discuss convergence towards a single project**

WARWICK
THE UNIVERSITY OF WARWICK