



Excalibur-HEP: “EMCuda” Test bench of secondary generation and handling on GPUs

Ben Morgan

Overview

- Excalibur-HEP: 15 month project since May 1st in UK funded under [UKRI Excalibur High Priority Use Cases Phase 1](#)
- Work package on simulation, focussing on use of GPUs for EM Physics (typically highest CPU consumer)
- Two lines of investigation:
 - *Secondary generation/management, understanding where/how to avoid divergences, Host/GPU transfer*
 - *Contribute to investigations on applicability of RT methods to charged particle transport*
- Strong overlap with Portability Work Package (oneAPI, Alpaka, etc)

Global Picture

- Established initial contacts with others working on these areas:
- **Geant4 R&D Task Force:** *projects on EM/RT*
- **HSF Simulation Working Group:** *participating in meetings, inc. brainstorming sessions*
- **US Exascale Project:** initial contacts through Excalibur to Celeritas project at ORNL/FNAL
- Links with other Excalibur Working Groups expected soon

Work So Far

- Ray Tracing work planned for second half of project
- Started with “simpler” problem of EM physics/showers
- Sandbox on CERN GitLab for:
 - *Information/Literature collation*
 - *Development of “Fisher-Price” test benches to explore algorithms/techniques/limitations on/of GPUs (or rather SIMT) for EM physics/showers*
 - *Enumerate issues found to identify “go/no-go” areas for further development or retirement*
 - *Shared codebase for portability studies*
- Interested parties are welcome to join/contribute!

“Fisher-Price” EM Physics in CUDA

- *No geometry/transport - but provide “stubs” for where the calls would occur*
- *Toy particles/physics: “e+/e-/g”, continuous energy loss, scatter, compton/pair (gradually add realistic models, utilizing work from CERN/FNAL)*
- *Start one particle on GPU*
 - *“Step” until it is killed by some threshold, “Score” energy losses*
 - *Key item: when secondaries are generated, how to handle/schedule them on the device?*
 - *Key item: how best to hand back info to host (“hits/scores”, maybe particles we can’t handle like hadrons)?*
- *Extraordinarily simple, but that’s the idea: to focus on the aspects, challenges, scaling for this type of MC on a GPU*

Summary

- Excalibur-HEP started, and linked up with worldwide efforts
- Starting simple
- Aim is not a full implementation, but a series of test benches to clearly identify traps and pitfalls
- This is also to allow exploration/comparison of using “portable APIs”
 - *Initially use CUDA to implement*
 - *Then try porting to oneAPI/Alpaka etc and compare performance*