OAK RIDGE National Laboratory

Experiment integration

Stefano Tognini

HPC Methods for Nuclear Applications

Oak Ridge National Laboratory

HSF Detector Simulation on GPU Community Meeting
May 3-6, 2022

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

U.S. DEPARTMENT OF ENERGY

***Celeritas* core team:**

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

Caveats

- Full assessment of the code maturity entails more than FOMs
  - Easiness to integrate on current workflows
  - LCF + HEP computing network integration
    - Storage and network transfer of simulated data on LCFs to HEP computing centers

Code adaptation + new production workflows + performance gain must be seen as worthwhile by experiments

Integration paths and challenges

Acceleritas

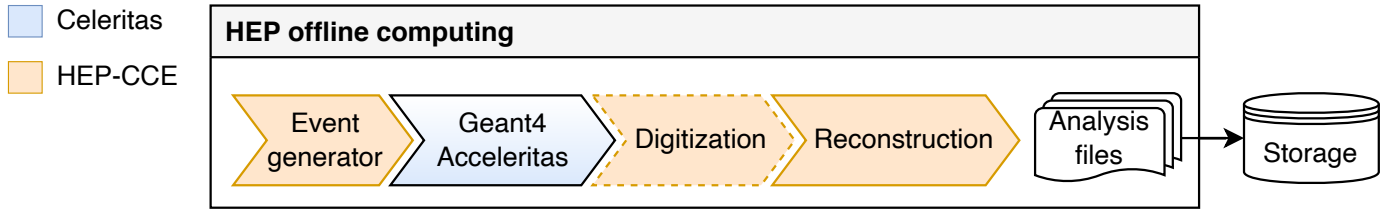

- Acceleritas library provides a streamline integration with relatively small changes
  - Cons:
    - Considerably smaller performance impact
    - Many (most?) offline HEP working nodes do not have dedicated GPU hardware

Integration paths and challenges

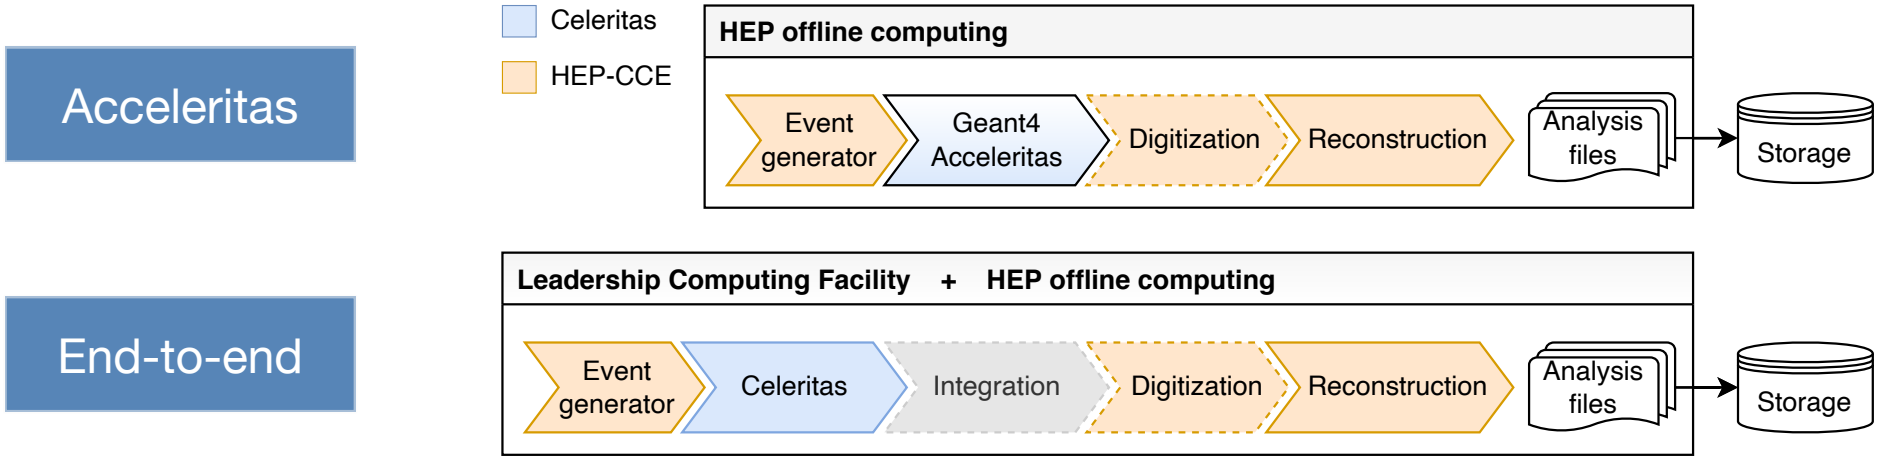

- End-to-end is envisioned after decay + hadronic physics are available
- Integration between LCFs and HEP-EX is a gray area still
  - Celeritas must provide I/O and in memory integration
  - Event-by-event workflows vs. Celeritas' track-based approach
  - No clear delimitation on what is processed where (LCFs vs. HEP computing systems)

Integration paths and challenges

Acceleritas

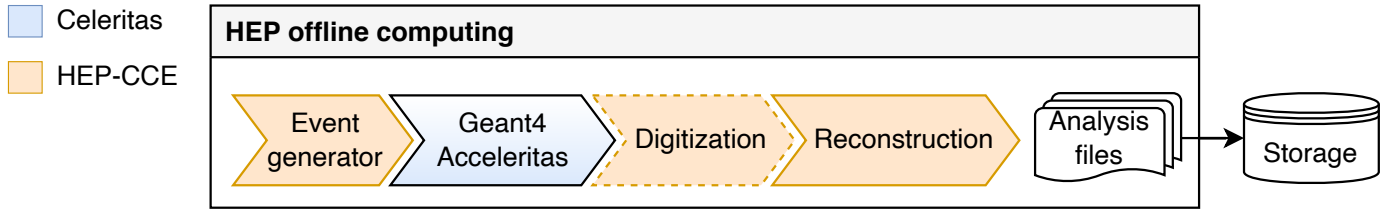

End-to-end

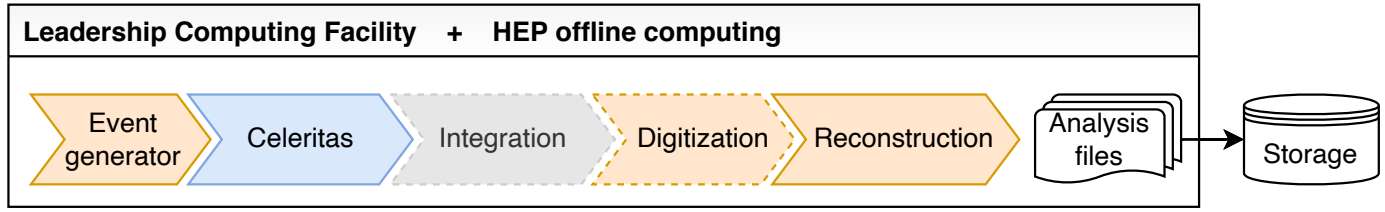

- End-to-end is envisioned after decay + hadronic physics are available
- Integration between LCFs and HEP-EX is a gray area still (cont.)
  - LCFs are mostly GPU (e.g. Summit is >95%); is CPU post-processing worth it?
  - A non-negligible amount of simulation happens outside production campaigns
  - Experiments might store more data (i.e. digi.root + reco1.root + reco2.root + ...)

Integration paths and challenges

- Best way to start answering these questions is to start interacting with experiments
  - Access to offline working nodes or
  - Being able to deploy their framework and run their workflow
    - Start with Acceleritas
    - Try to run an end-to-end hypothetical workflow (even with limited physics) to know what we need