

Event Generator Update and Future Plans

Taylor Childers (ANL) and Stefan Hoeche (FNAL)

Two Ongoing Projects

Pepper (was Blockgen) for Sherpa

ANL: Taylor Childers, Rui Wang

FNAL: Walter Giele, Stefan Höche, Joshua Isaacson

Göttingen: Enrico Bothmann, Max Knobbe

Designing a leading-order matrix element integrator and event generator for GPUs from the ground up. HEP-CCE contributing by porting to Kokkos.

MadGraph for GPU

ANL: Taylor Childers, Walter Hopkins, Nathan Nichols,

CERN: Laurence Field, Stefan Roiser, David Smith, Andrea Velassi

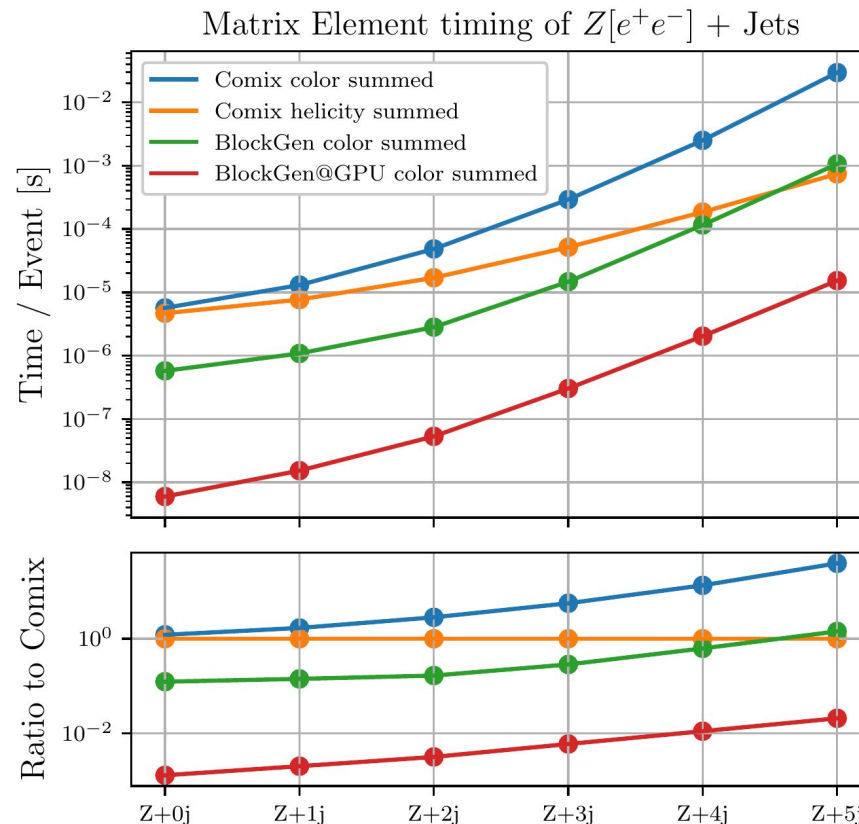
UC-Louvain: Olivier Mattelaer


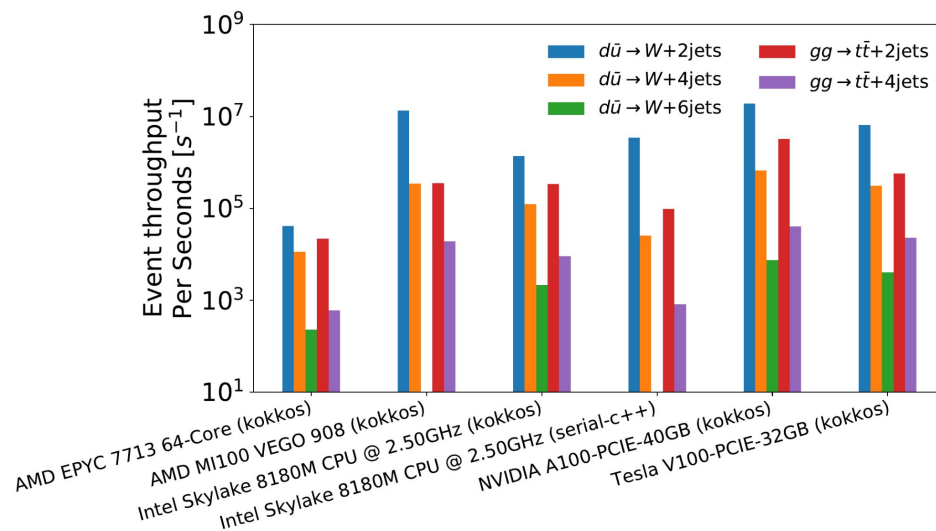
Retrofit of leading-order MadGraph Fortran algorithm to run on GPUs. CERN team developing CUDA version. HEP-CCE team + ALCF developing Kokkos and Sycl version.

Pepper Update

ACAT 2022 Results

- Extension from pure QCD to V+jets using novel color decomposition
- Realistic setup for ME computation
Include all sub-processes, no PDF
- Compare dedicated C++ Version with dedicated Cuda version
- Phase Space: Rambo
CPU: i3-8300 CPU @ 3.70GHz
GPU: Tesla V100S
- Excellent performance compared to current Sherpa standalone (Comix)







ACAT2022 Results


- Kokkos portability framework to access all modern (co-)processors
- Very good performance in comparison to C++/CUDA prototypes
- Currently being integrated into event processing framework

MadGraph Update


- CERN team leads CUDA development
- ANL team leads Kokkos and Sycl development
- Historically MadGraph uses a python user interface, which generates, compiles, and runs FORTRAN code to generate events.
- Process began with a CUDA implementation of the Matrix Element (ME) calculations which are used to calculate the interaction cross-section for a process and are the biggest computational piece.
- Kokkos port followed, then Sycl, then Alpaka, and the CUDA port was extended to also offer OpenMP



Results at ACAT2022



Variable GPU-grid size (throughput scan)

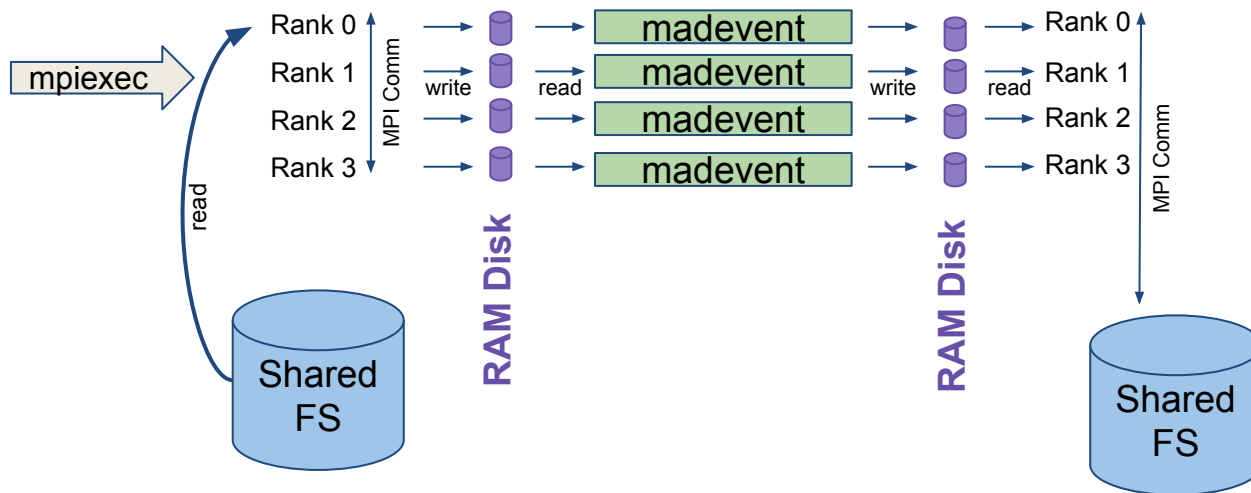


Presented results at ACAT2022

- *tt*+2jet process has reasonable computational needs on an accelerator.
- Shows the power of using portability frameworks to access all modern silicon.
- XE-HPC is an early Aurora-type GPU.

Proceedings draft on [arXiv:2303.18244](https://arxiv.org/abs/2303.18244)



Setup for multi-node running



- MadEvent is a serial Fortran code calling our parallelized-C++ Sycl version of the leading-order matrix element calculations.
- We've developed a set of wrapping applications that can run MadEvent using MPI across multiple nodes of an HPC

Scaling Tests on Polaris with Sycl Version Only


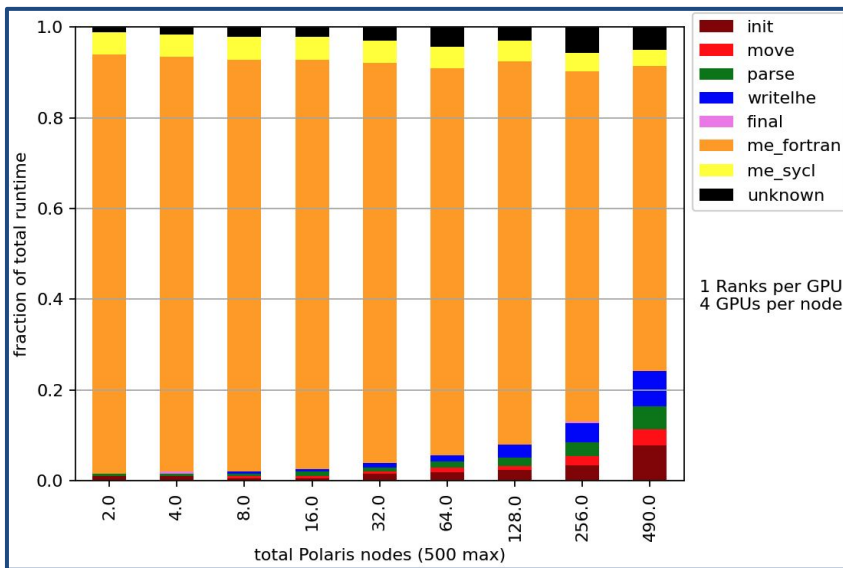
HEP-CCE



- Matrix element calculation rate is stable, but overall run time is increasing
- At 490 nodes of polaris, we are producing 12,126,203 tt+2jet events.
- Polaris is a 500 node system at ALCF (4x A100 per node).

Scaling Tests on Polaris with Sycl Version Only

HEP-CCE



- Matrix element calculation rate is stable, but overall run time is increasing
- At 490 nodes of polaris, we are producing 12,126,203 tt+2jet events.
- The distribution of time spent at each scale shows that the IO parts are increasing in overall run proportion, driving up run times.

Future Plans

Pepper / Sherpa

- Combination with parton-level integrator developed during SciDAC-4.
- Integration into event generation framework developed during SciDAC-4.
- Integration into ATLAS/CMS workflow.
- Tests & performance tuning on available hardware.

MadGraph for GPU

- Aiming for alpha release by CHEP in May.
- Focusing on Sycl port and dropping support for Kokkos
- Continuing to test different processes.
- Walter working with Nathan to produce SUSY signal samples for an analysis.
- Aurora scaling testing will be done when machine is accessible.

Longer Term:

- Begin NLO algorithmic R&D