AHM Closeout

Meeting Goals → achieved!

- Carry out a self-assessment of where we are at this point as a project (roughly two and a half years after the start of the planned activities) for PPS, IOS and EG
- Plan in detail area activities until the end of the current funding cycle (FY23), keeping in mind possible activities beyond FY23.
  Our FY23 Planning document will be the starting point for the detailed planning

IOS Progress

Impressive!

HEP-CCE

Total throughput (Events/sec)

- Total throughput calculated using event processing time and number of events
- Performance scaling with the number of nodes
- Using 4 ranks per node (16 threads/rank) gave the best performance, ~51% improvement as compared with 1 rank per node.
- ~97K events/sec is the maximum throughput observed for 256 ranks on 64 nodes.
- This translates to 8GB/s bandwidth we were able to consume for the largest set.

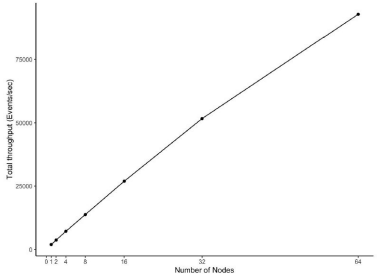

HEP-CCE

Darshan handling of `fork()`

To address this Darshan was modified to properly handle apps that call `fork()`

- Use `pthread_atfork()` to register a callback that is executed before passing control to child process
- Darshan atfork callback re-initializes the library to clear all parent process state
- Child processes maintain mapping to corresponding parent process ID, allowing Darshan logs to capture process relationships

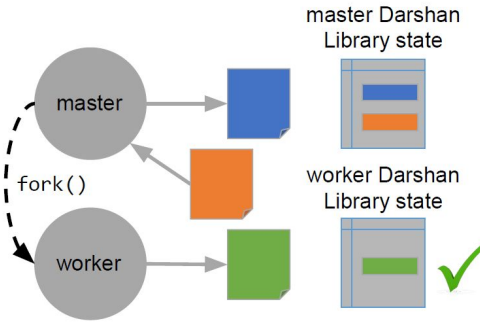

U.S. DEPARTMENT OF ENERGY | Office of Argonne BROOKHAVEN Fermilab BERKELEY LAB

POSIX Operations

- In AthenaMP each worker does the writes while SharedWirter took over these

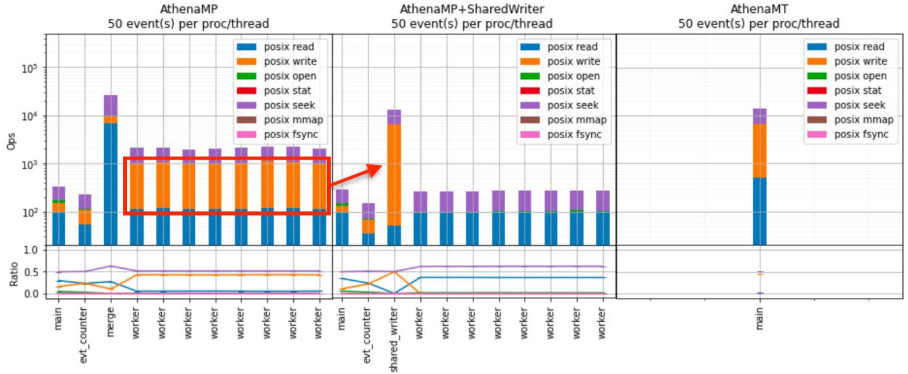

Data Structure of ProtoDUNE Raw MC

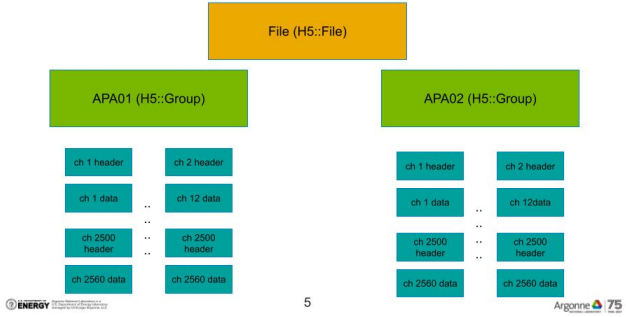

PPS +EG Progress

Relentless!

HEP-CCE

ICHEP2022 SLIDES

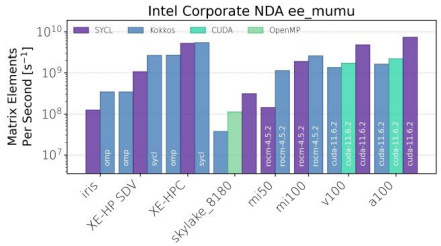

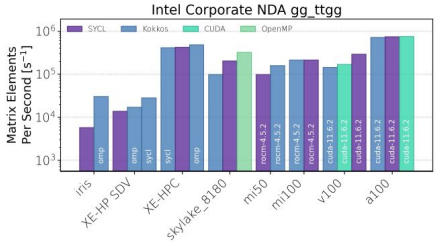

Performance Comparison

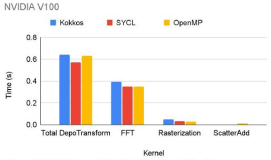

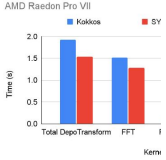

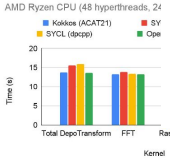

- **FFT :** SYCL/OpenMP better than Kokkos:
  - Due to optimized normalization kernel (fft) and S*F etc kernels from 2D to 1D Compared to Kokkos code
- **Rasterization:** SYCL/OpenMP better than Kokkos
  - Mainly due to RNG . Other kernels are similar in timing
- **ScatterAdd:**
  - SYCL 5x slower than Kokkos for CUDA backend
  - OpenMP 15x slower than Kokkos for CUDA backend (should be much better in clang-15)
  - SYCL same as Kokkos for HIP(AMD)

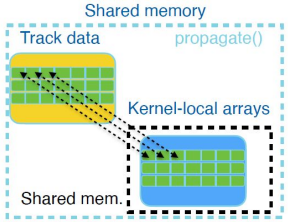

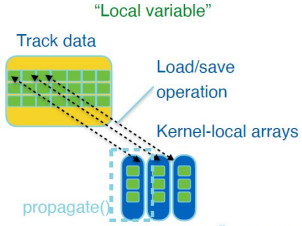

- Patatrack
  - A frozen, standalone version of CMS heterogeneous pixel track and vertex reconstruction
    - "End-to-end", with mock framework and build system
  - Current status

	Implementations							
	CPU Serial	CUDA (original)	HIP	Kokkos	Alpaka (by CERN team)	std::par	SYCL (by CERN team)	OpenMP
NVIDIA			Eigen does not support					
AMD				Crashes randomly				
Intel				Does not compile (Eigen)				
CPU				Serial, POSIX threads	Serial, TBB			

Completed
In progress
Not started

Performance Studies: std::par Group

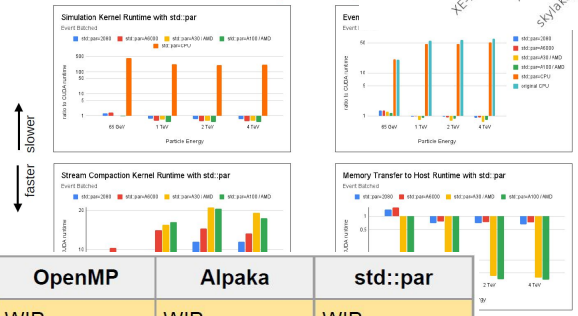

Particle Energy	CUDA Kernel Runtime
65 GeV	385 µs
1 TeV	987 µs
2 TeV	1093 µs
4 TeV	1210 µs

	Kokkos	SYCL	OpenMP	Alpaka	std::par
Patatrack	Done	Done	WIP	WIP	WIP
Wirecell	Done	WIP	WIP	stretch goal	TBD
FastCaloSim	regular: done group: WIP	regular: done group: no	regular: WIP group: TBD	regular: WIP group: TBD	regular: done group: done
P2R	done	Done	Need local var	Done	Done

FY23 Planning (some initial, personal observations)

Nobody complained loudly about the ~~draft~~ plan

PPS: lots has been accomplished, focus on reporting

IOS: Finish (and publish) HDF-5 event store studies

EG: Lots on their plate with madgraph4GPU and BlockGen prototype

CW: ATLAS, DUNE interest in funcX and Parsl. We will follow up in new series of workflow meetings soon

Bonus: ESnet discussion on data caches started informal follow-ups during ESnet-6 workshop

Group Dinner Tonight

Great China at 7PM

- 2190 Bancroft Way - SW corner of Bancroft and Oxford
- 20 minute (downhill) walk, or take the LBL shuttle to BART and walk 3 blocks

Room Change for Thursday

Tomorrow morning we will be in **Bld 59**
(Shy Wang hall, where Cori/Perlmutter live)
in **4102** (one floor up)

same zoom meeting coordinates

Thanks

To **Jason, Charles** and the rest of the LOC
(too bad I could not taste those cookies)

To all who joined and contributed in the room or via zoom.

Special thanks to our guest speakers from ESnet and NVIDIA

Enjoy Great China!