

Benchmarks on ARIS

Tsapatsaris Panagiotis(Ntua)

Supervisors:

K. Iliakis(Ntua,Cern)

H. Timko(Cern)

S. Xydis(Ntua)

D. Soudris(Ntua)

August 2020

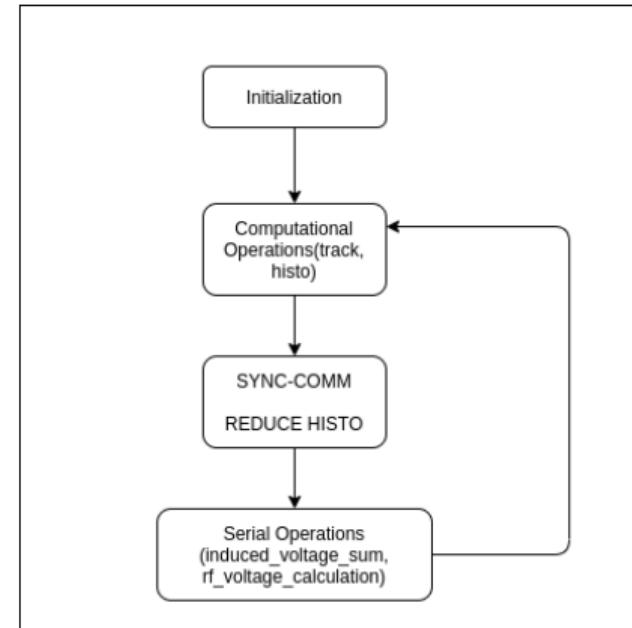
Classification of Operations

- Computation or COMP : dependent on the number of particles (kick, drift, linear_interpolation_kick, histogram)
- Communication or COMM : operations needed for distributed implementation (all reduce for histogram)
- Serial : independent of the number of particles (beam_phase, induced_voltage_sum)

Type of Benchmarks

Data Parallelism Version

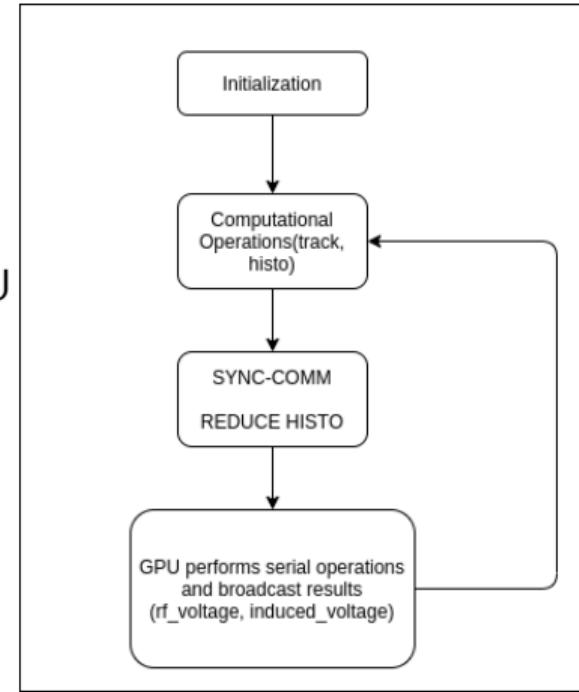
- *Serial + Comp* equal for all workers
- Can control only Computation Time



Type of Benchmarks

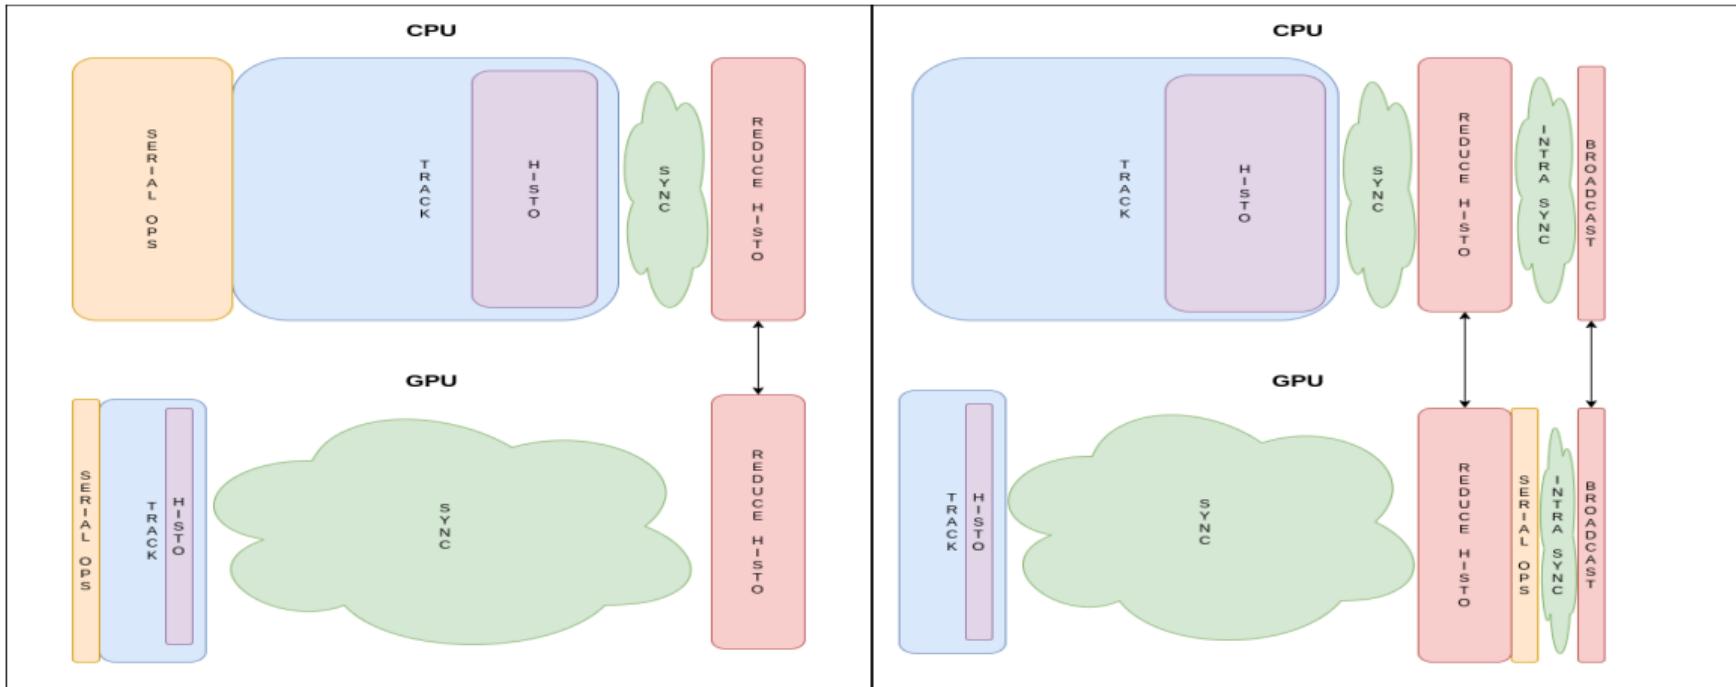
Task Parallelism Version

- Goal : Minimize Sync Time
- Some Serial Operations can be done in parallel
- GPU serial operations are faster
- We can share the results with the slower CPU
- *Comp* needs to be the same for all workers

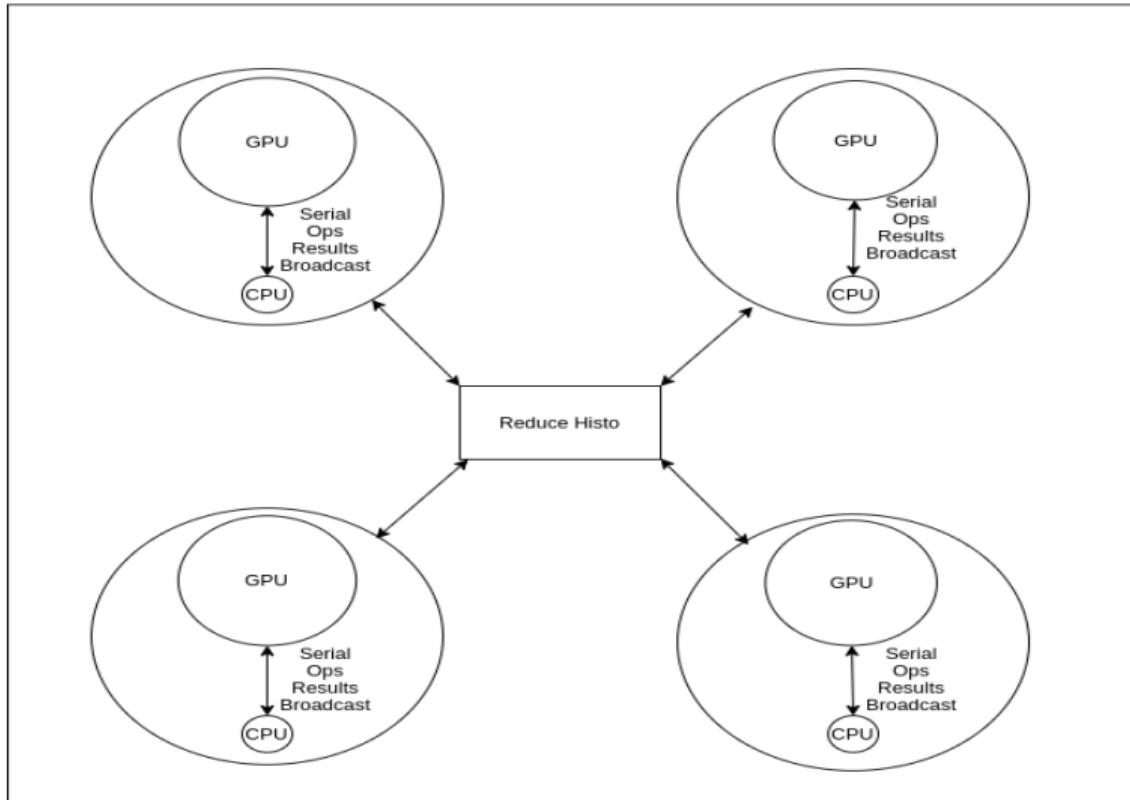


One Node CPU-GPU

Simple Version - Task Parallelism Version



Multi-Node Task Parallelism CPU-GPU

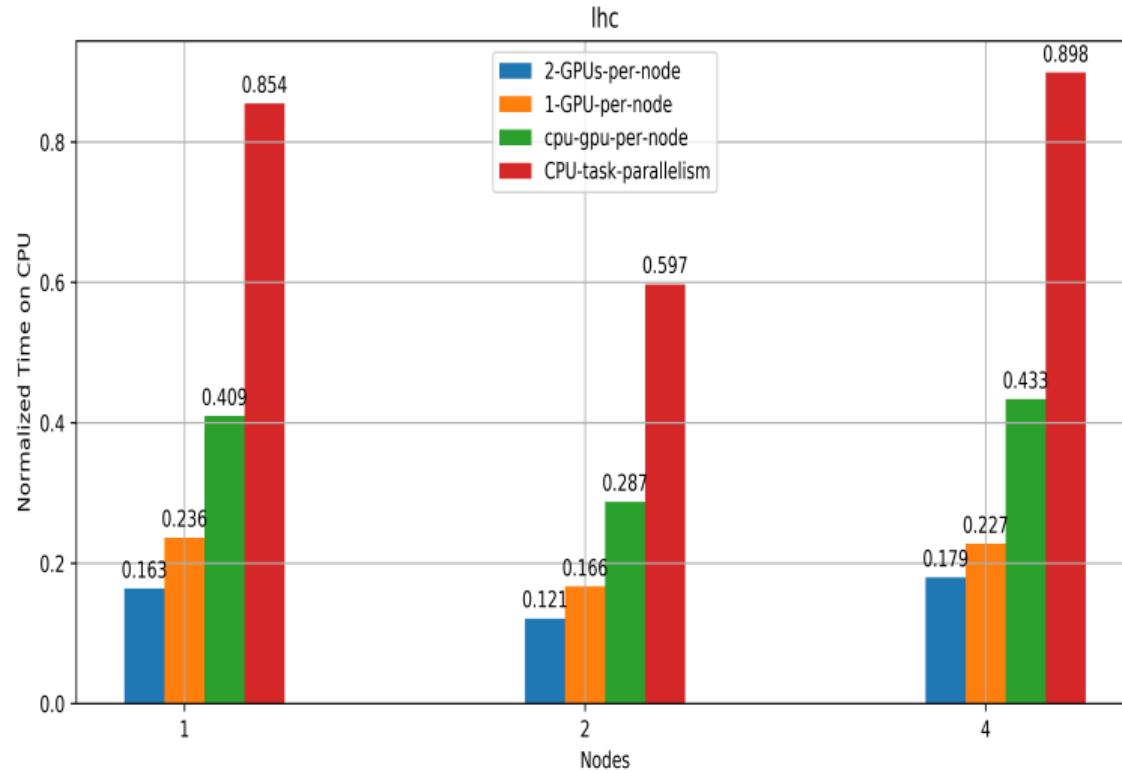


CPU vs GPU

Our benchmarks:

- nodes : [1,2,4]
- particles per node : 48m for LHC, 84m for PS
- slices : 1000
- ARIS CPUs per Node : 2 Intel Xeon(R) E5-2660v3
- ARIS GPUs per Node : 2 Nvidia K40
- 2-GPUs-per-node, 1-GPU-per-node, 2-CPUs-per-node, 2-CPUs-per-node-tp,
CPU-GPU-per-node-tp

CPU vs GPU



CPU vs GPU

