HEPscore Workshop: Study of MLPF-based GPU benchmarking

Eduard Cuba, David Southwick
Eric Wulff, Maria Girone

September 20, 2022

The role of AI workloads in HEP

Increasing interest in ML for event reconstruction due to

- Increased data volume and processing requirements
- Possibly better scalability
- Data-driven
- Use-cases at ATLAS, CMS, LHCb, ALICE

→ Changing paradigm and new requirements for HPC

→ Resource intensive training, efficient inference

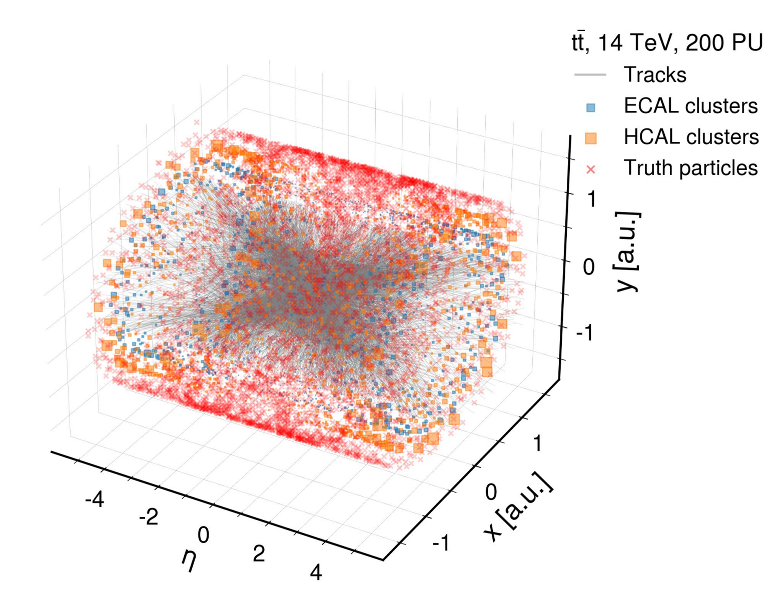

Example event

Figure source: [1] Pata et. al. (2021): https://arxiv.org/abs/2101.08578

Matrix multiplication on GPU

Bottleneck: Large matrix multiplications

→ Accelerators required

→ GPUs: Optimized for parallel computing

Massive parallelism (6912 cores @ Nvidia A100)

- divide and conquer mechanism using high number of computing cores (more than provided by CPUs)
- possibly multiple GPUs per node

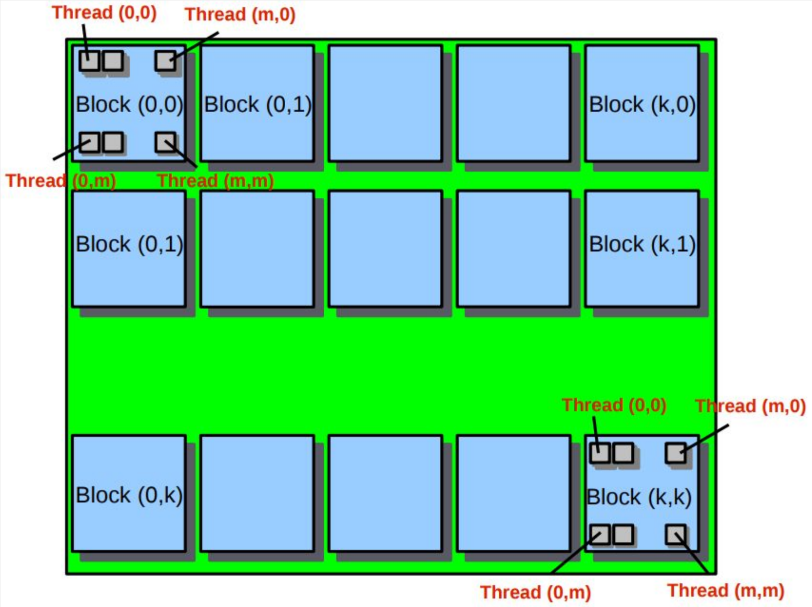

Figure source: GPU Programming with CUDA @ JSC, 24. – 26. April 2017

Benchmark definition

ML workflow: First training, then inference.

Training goal: learn a function from training data, such that it generalizes well to unseen data.

Take an ML model (e.g. neural network) and select:

- architecture (layers, activation functions)
- dataset
- optimizer, loss function
- number of epochs (passes through the dataset)
- batch size (samples per weight update step)

→ We get an execution graph

→ Perform a fixed number of executions

→ Compute benchmark metrics

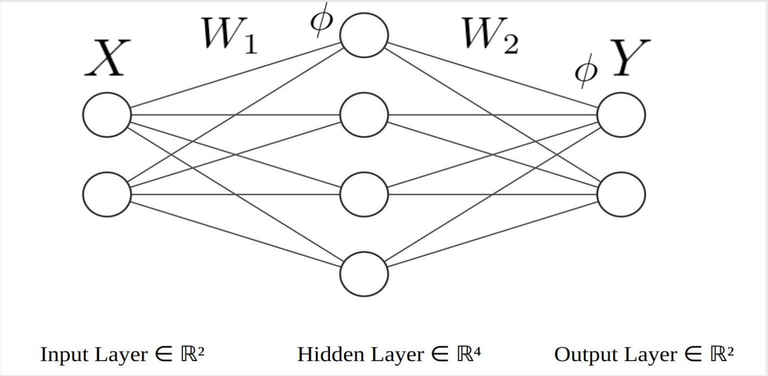

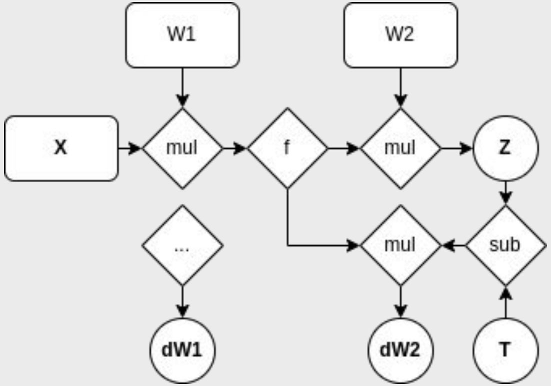

Benchmarking goals

Measure how fast and efficiently we can train on the given hardware

- does it fit? can this HW train the model?
- how fast does it train?
- how much does it cost?

Metrics of interest:

- training time (on a fixed dataset and num. of epochs)
- events per second (throughput)

→ Find the best hardware for training the model

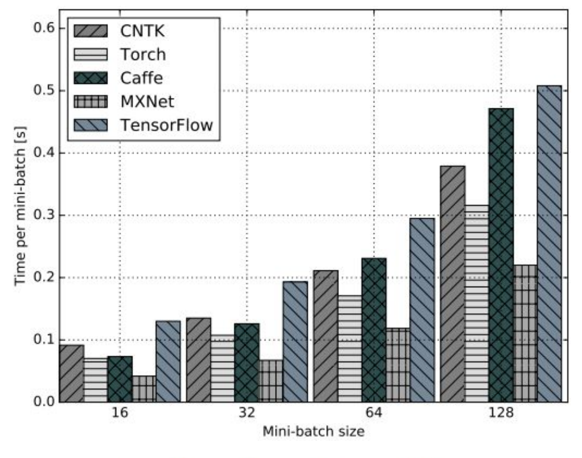

(a) Results on Tesla K80.

Benchmarking of AlexNet-S on TeslaK80 GPU using different ML frameworks.

Figure source: https://arxiv.org/abs/1608.07249

MLPF: Machine-Learned Particle-Flow

RAISE Center of Excellence

Graph Neural Network (GNN) for particle-flow reconstruction.

- ➢ Representative AI workflow in HEP
- ➢ Experience with training and hyperparameter tuning within the RAISE project [2]
  - ○ 4 NVIDIA A100 series GPUs per node took ~75k core-hours, 83 hours on 12 nodes

MLPF-based benchmark shows feasibility of training HEP AI applications on HPC hardware.

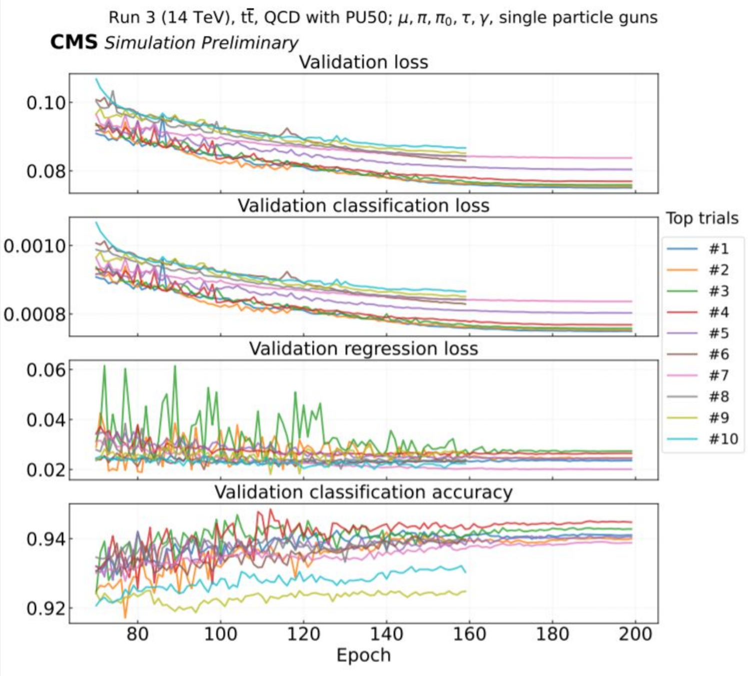

Previous hypertuning results from ACAT 2021, thanks to the CMS collaboration

Figure source: [2] Wulff et. al. (2022): https://arxiv.org/abs/2203.01112

Benchmark integration

Approach ML/AI workloads as repeatable benchmark

- Containerized in similar manner to traditional CPU benchmarks (Docker, Singularity)
- Support (multi) GPU accelerators
- Allow configuration (number of GPUs, batch size, etc.)
- Report metrics in a HEPscore compatible format
- Small dataset portion (~1GB) within the container
- Using a subset of a public Delphes dataset (Pata et. al., 2021 [3], https://zenodo.org/record/4559324)

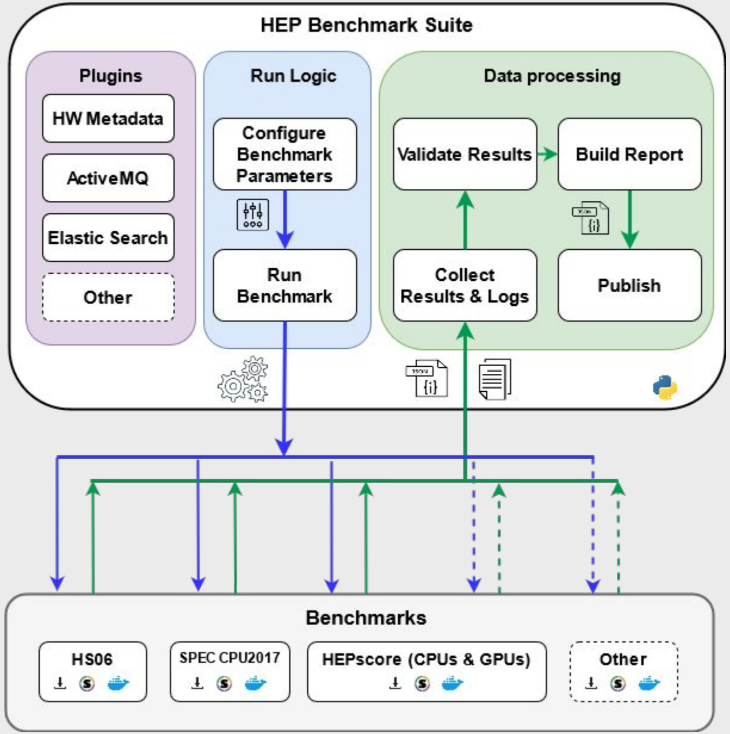

Figure source: https://gitlab.cern.ch/hep-benchmarks/hep-benchmark-suite

Benchmark metrics

RAISE Center of Excellence

Inspired by the events/second processed used in HEPiX CPU jobs

→ events as training samples in the ML context

→ throughout (training samples/second)

→ = batch_size * batches_per_second

Additionally, report epoch time

→ epoch: one pass through the dataset

Hardware	Mean Epoch time	Mean Throughput
Intel Xeon Gold 6148 40-Core	365.3 s	3.01 eps
2 x AMD EPYC 7742 64-Core	780.7 s	1.41 eps
Intel Xeon Platinum 8358 + 4 x NVIDIA A100 40GB	13.48 s	**128.5** eps

Reproducibility

Same computational graph on each platform

- results are platform independent (in model accuracy)
  - as long as the hyper-parameters are the same (batch size, lr, optimizer...)
- network training uses randomness
  - fix the random seeds for weight initialization and optimizer
  - same workflow executed -> same results
- minor differences in FP representations
  - potentially different order of operations (math kernel, massive parallelism)

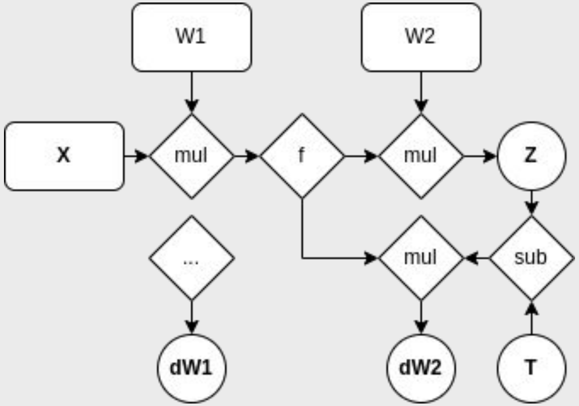

Reproducibility results

Validation accuracy (1100 events, batch size 16, 5 epochs), 2 runs:

Metric	Nvidia V100	Nvidia A100	Intel Skylake 40-core	AMD Rome 128-core
Validation Accuracy	0.9576 0.9576	0.957**9** 0.957**8**	0.9576 0.9576	0.9576 0.9576
Mean throughput [e/s]	16.70 16.70	45.68 45.60	3.02 3.01	2.00 1.90
Mean epoch time [s]	65.87 65.88	24.08 24.12	364.47 365.65	549.32 579.43

~ 0.01% difference in validation accuracy on GPUs

Tests performed on native (AMD), Intel and Nvidia tensorflow backends.

Scalability studies

Benchmarking is focused on single nodes

- possibly with multiple GPUs

Benchmarking of multi-node jobs is not supported

We explored scalability effects in training to large numbers of nodes in actual training

- scaling is near linear for smaller number of nodes

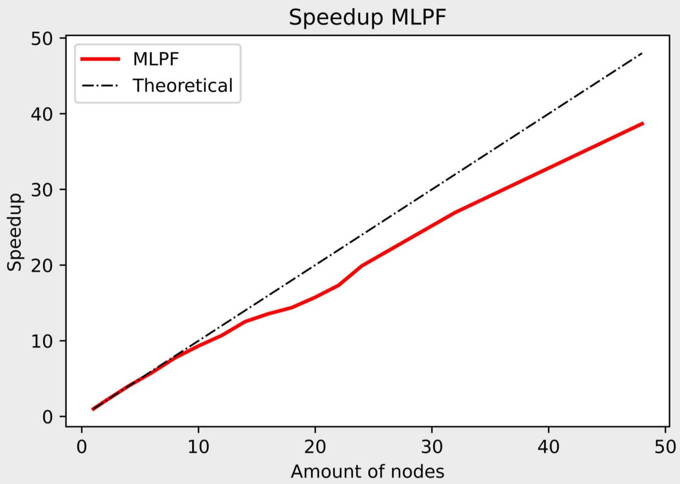

Data-parallel training of MLPF using Horovod on JUWELS Booster, 4xGPU per node

Discussion

Now: Applicable on single nodes (with multiple GPUs)

Open questions:

- Adjusting configuration for data-parallelism
    - batch size for multiple GPUs
    - batch size for large or small memory GPUs
    - leads to different model results?
    - dataset needs to be large enough

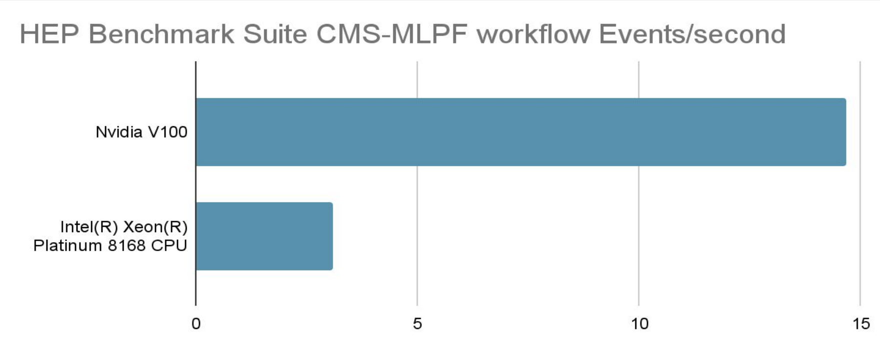

Conclusions

First ML/AI workloads for HEPiX benchmark working group introduced

- train the same AI model (e.g. MLPF) in a fixed setting on different platforms
- test feasibility of training / tuning HEP-driven AI applications on HPC hardware
- growing support for heterogeneous workloads in the benchmarking suite

We can investigate questions like:

- how fast and efficiently can we train this AI models on this platform
- what is the required compute time
- what is the impact of HW settings on the performance (e.g. clock speed)

References

[1] Pata, J., Duarte, J., Vlimant, J.-R., Pierini, M., and Spiropulu, M., “MLPF: efficient machine-learned particle-flow reconstruction using graph neural networks”, European Physical Journal C, vol. 81, no. 5, 2021. doi:10.1140/epjc/s10052-021-09158-w.

[2] Wulff, E., Girone, M., and Pata, J., “Hyperparameter optimization of data-driven AI models on HPC systems”, arXiv e-prints, 2022. https://arxiv.org/abs/2101.08578

[3] Pata, Joosep, Duarte, Javier Mauricio, Vlimant, Jean-Roch, Pierini, Maurizio, & Spiropulu, Maria. (2021). Simulated particle-level events of ttbar and QCD with PU200 using Pythia8+Delphes3 for machine learned particle flow (MLPF) (v1.1) [Data set]. Zenodo. https://doi.org/10.5281/zenodo.4559324

drive. enable. innovate.

RAISE
Center of Excellence

The CoE RAISE project has received funding from the European Union's Horizon 2020 – Research and Innovation Framework Programme H2020-INFRAEDI-2019-1 under grant agreement no. 951733

Follow us: