CERN openlab

Evaluating IBM POWER Architecture for Deep Learning in High-Energy Physics

CERN openlab Workshop 2019

Ahmad Hesam

23/01/2019

About Me

- PhD Student from Delft University of Technology
- Quantum and Computer Engineering department
- Research on DL in HEP at CERN openlab in collaboration with IBM

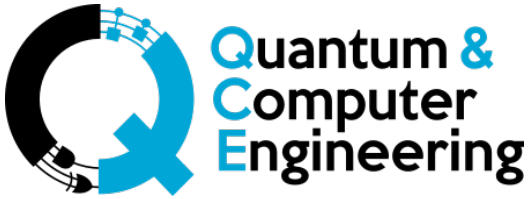

3D GAN Architecture

~ 25h training on single P100!

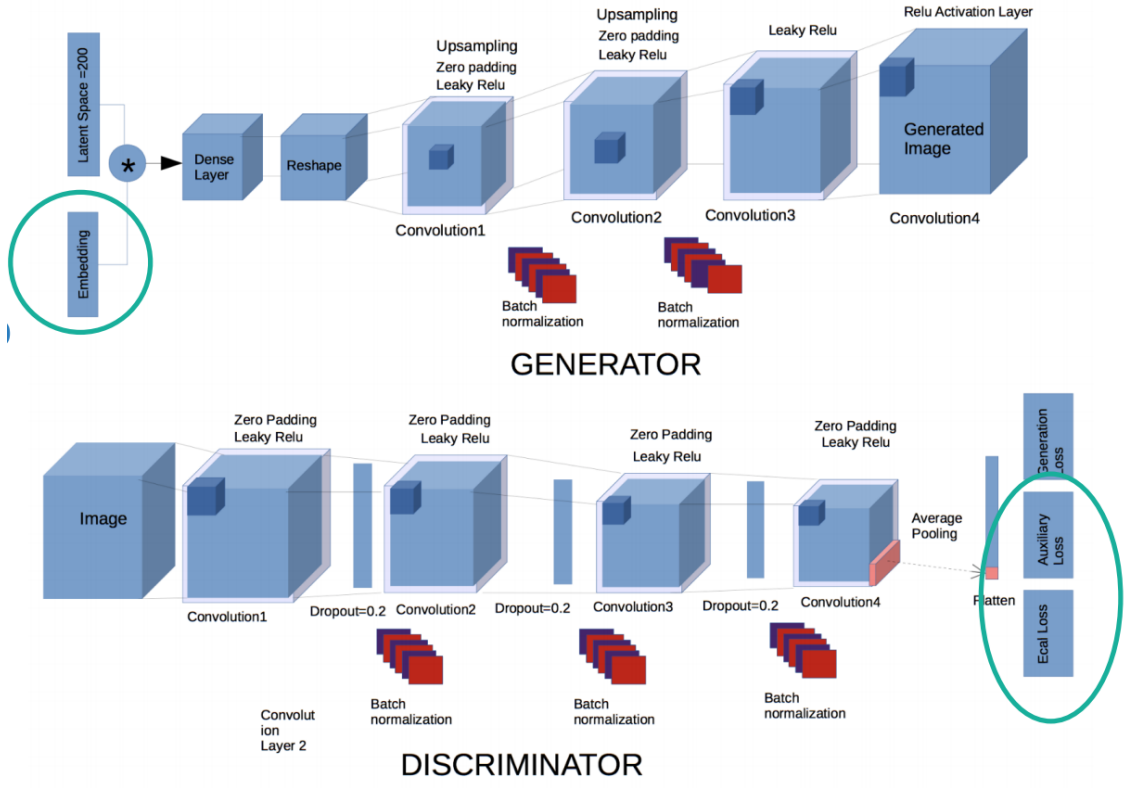

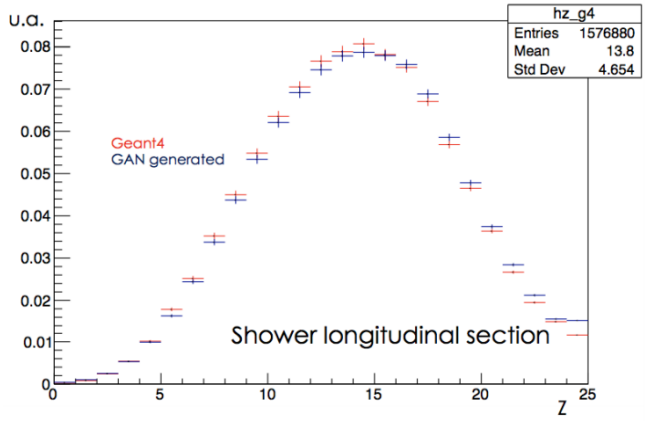

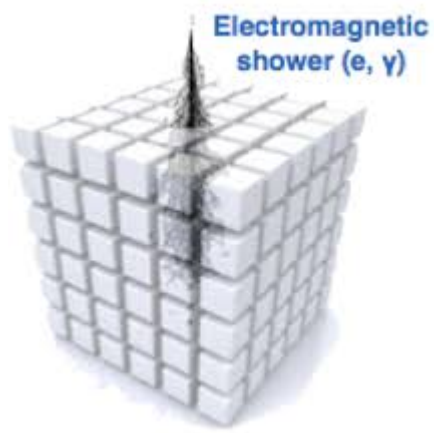

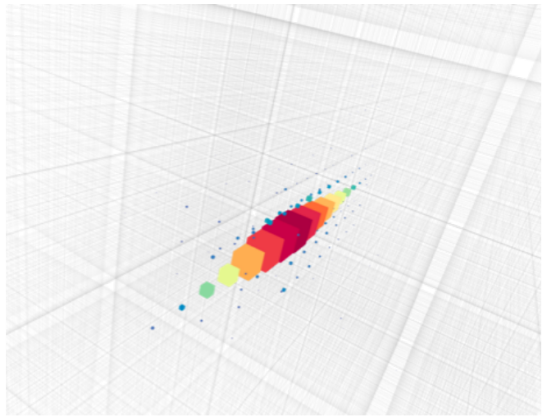

IBM Minsky

1x IBM SC821LC (login node)

- 1x POWER8 socket (=8 cores)
- 64 GB DDR4

3x IBM SC822LC (worker nodes)

- 2x POWER8 sockets (=16 cores)
- 4x NVIDIA P100
- CPU ↔ GPU NVLink
- 256GB DDR4

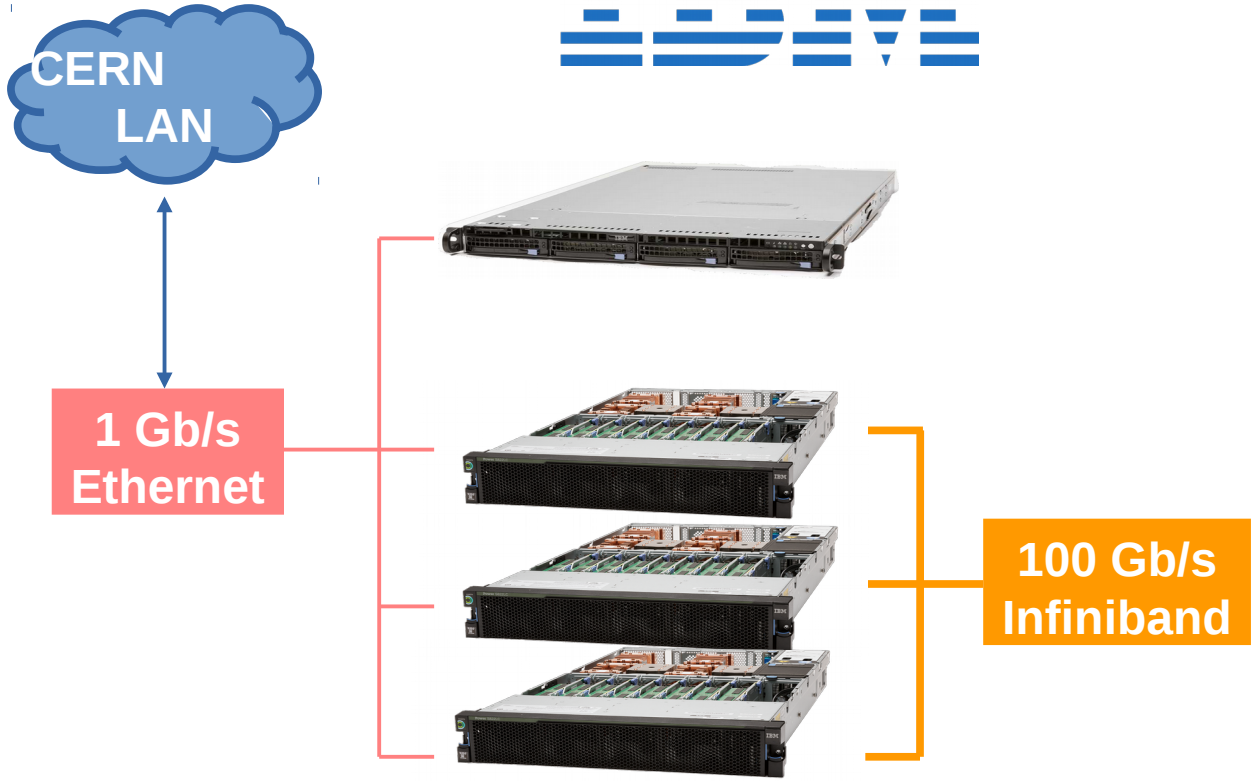

IBM Minsky

1x IBM SC821LC (login node)

- 1x POWER8 socket (=8 cores)
- 64 GB DDR4

3x IBM SC822LC (worker nodes)

- 2x POWER8 sockets (=16 cores)
- 4x NVIDIA P100
- CPU ↔ GPU NVLink
- 256GB DDR4

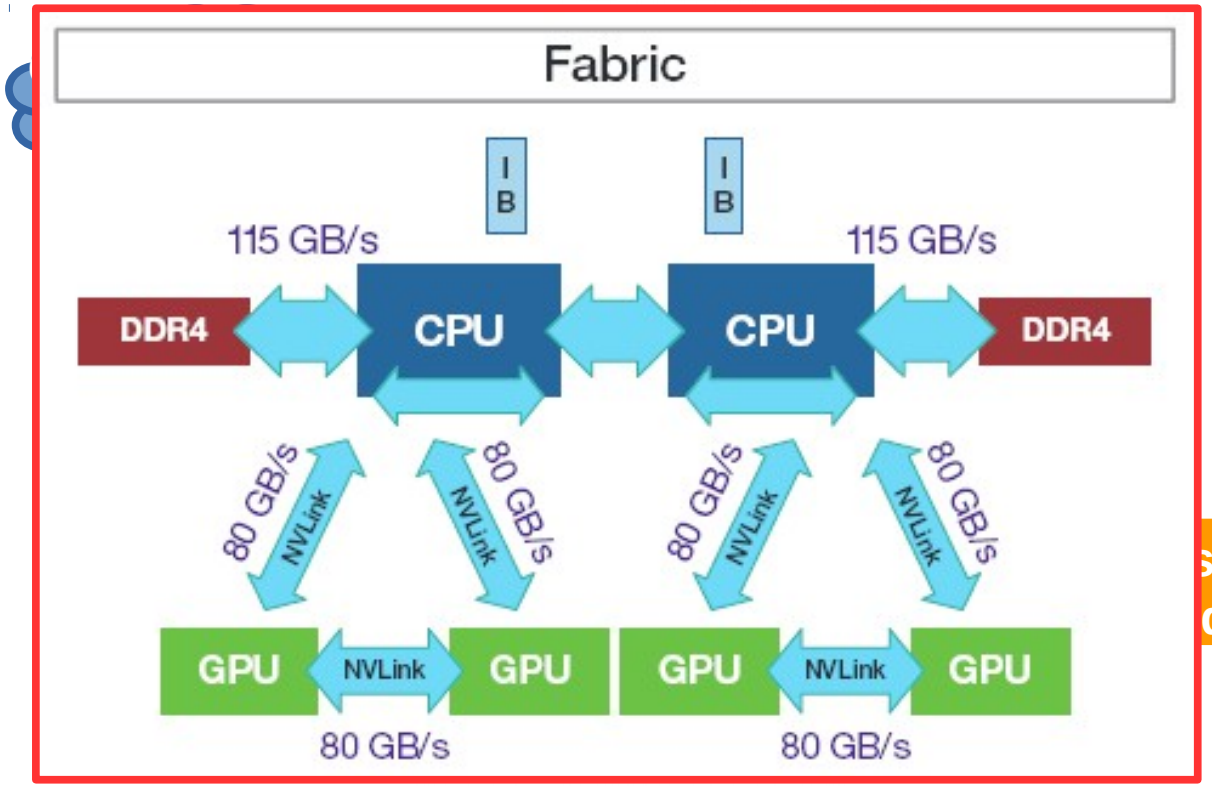

http://developercongress2017.openpowerfoundation.org/wp-content/uploads/2017/05/Porting-Applications-to-OpenMP4.5.pdf

IBM Deep Learning Stack

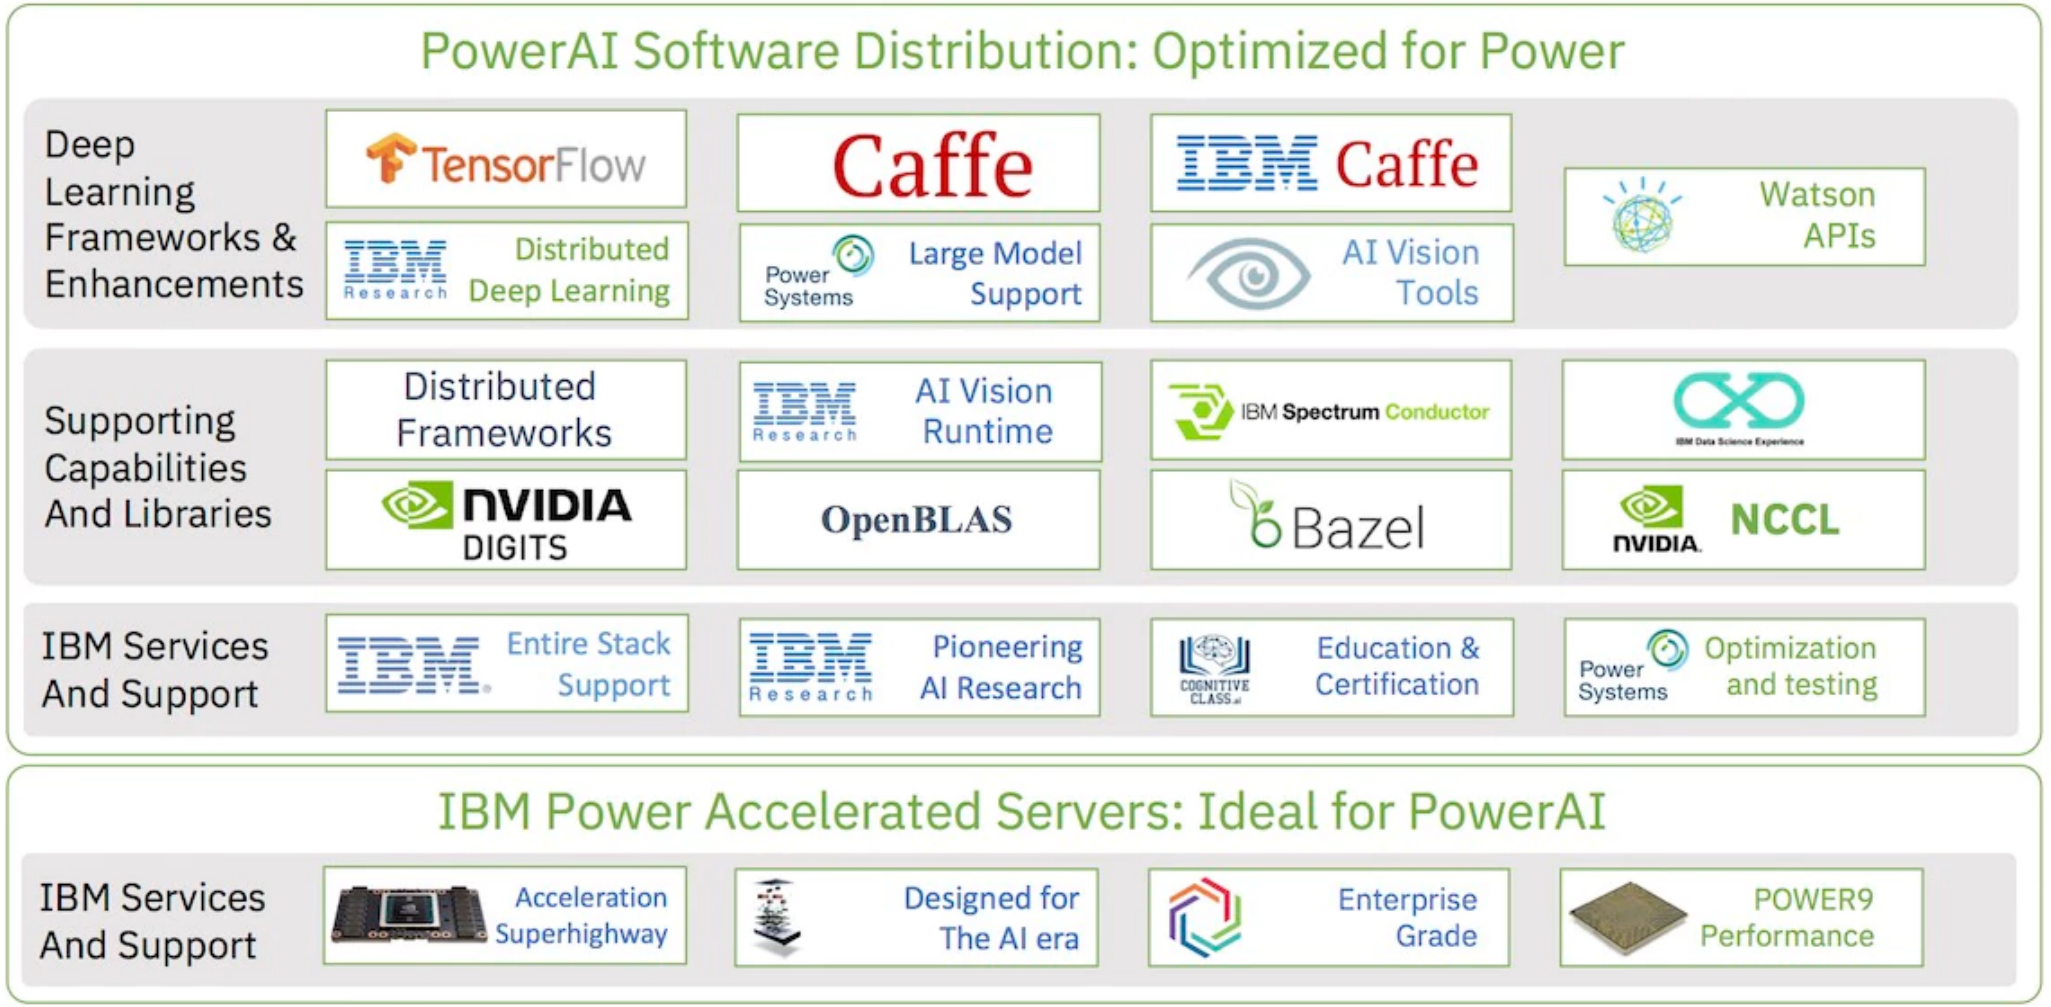

Model Parallelism

vs.

Data Parallelism

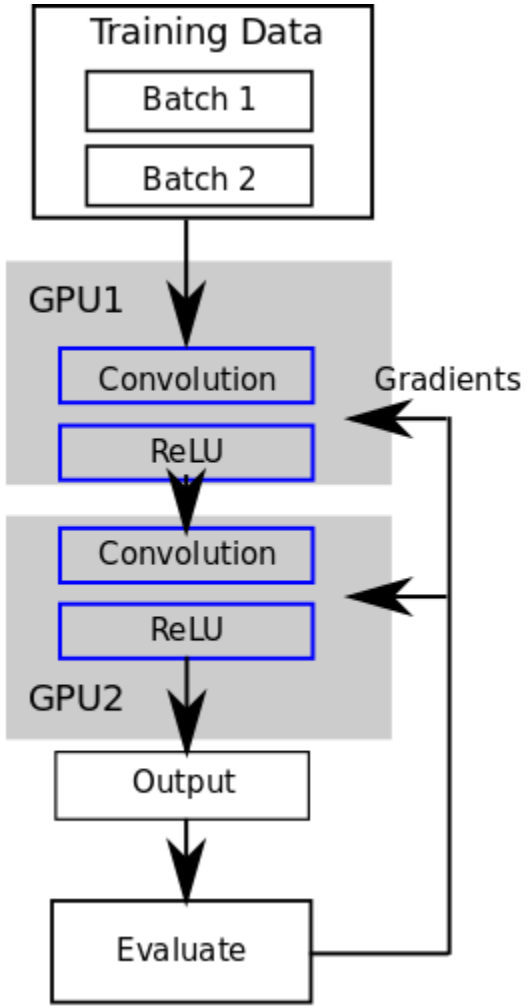

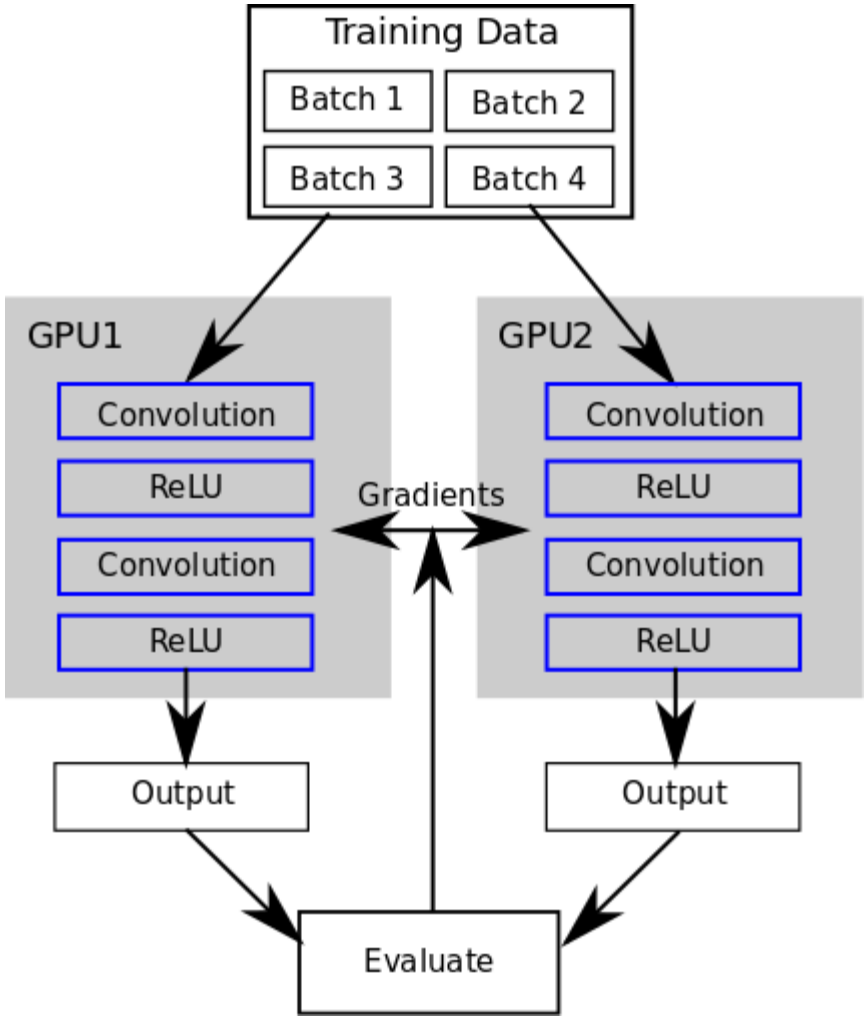

https://blog.rescale.com/deep-learning-with-multiple-gpus-on-rescale-torch/

Model Parallelism

vs.

Data Parallelism

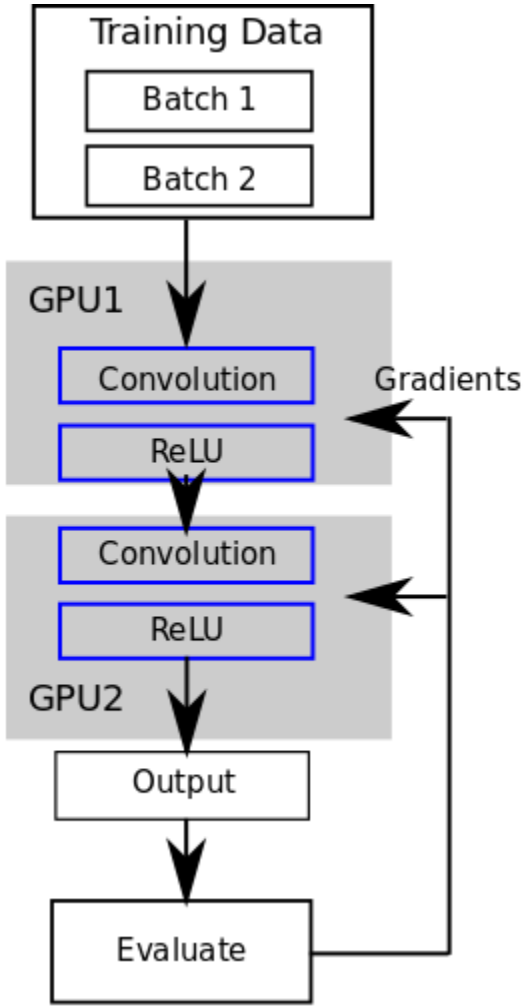

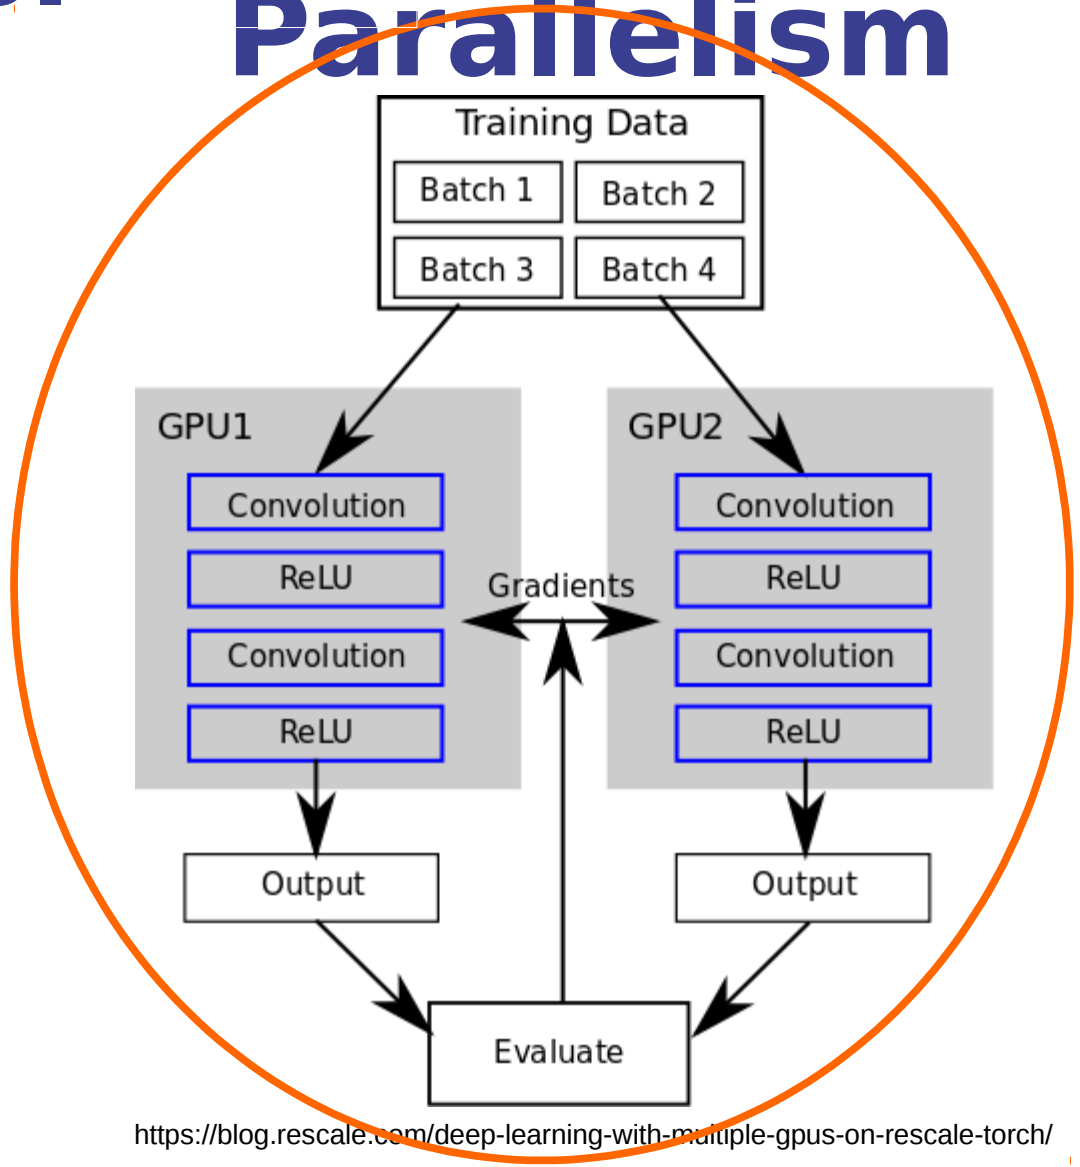

https://blog.rescale.com/deep-learning-with-multiple-gpus-on-rescale-torch/

Distributed Deep Learning

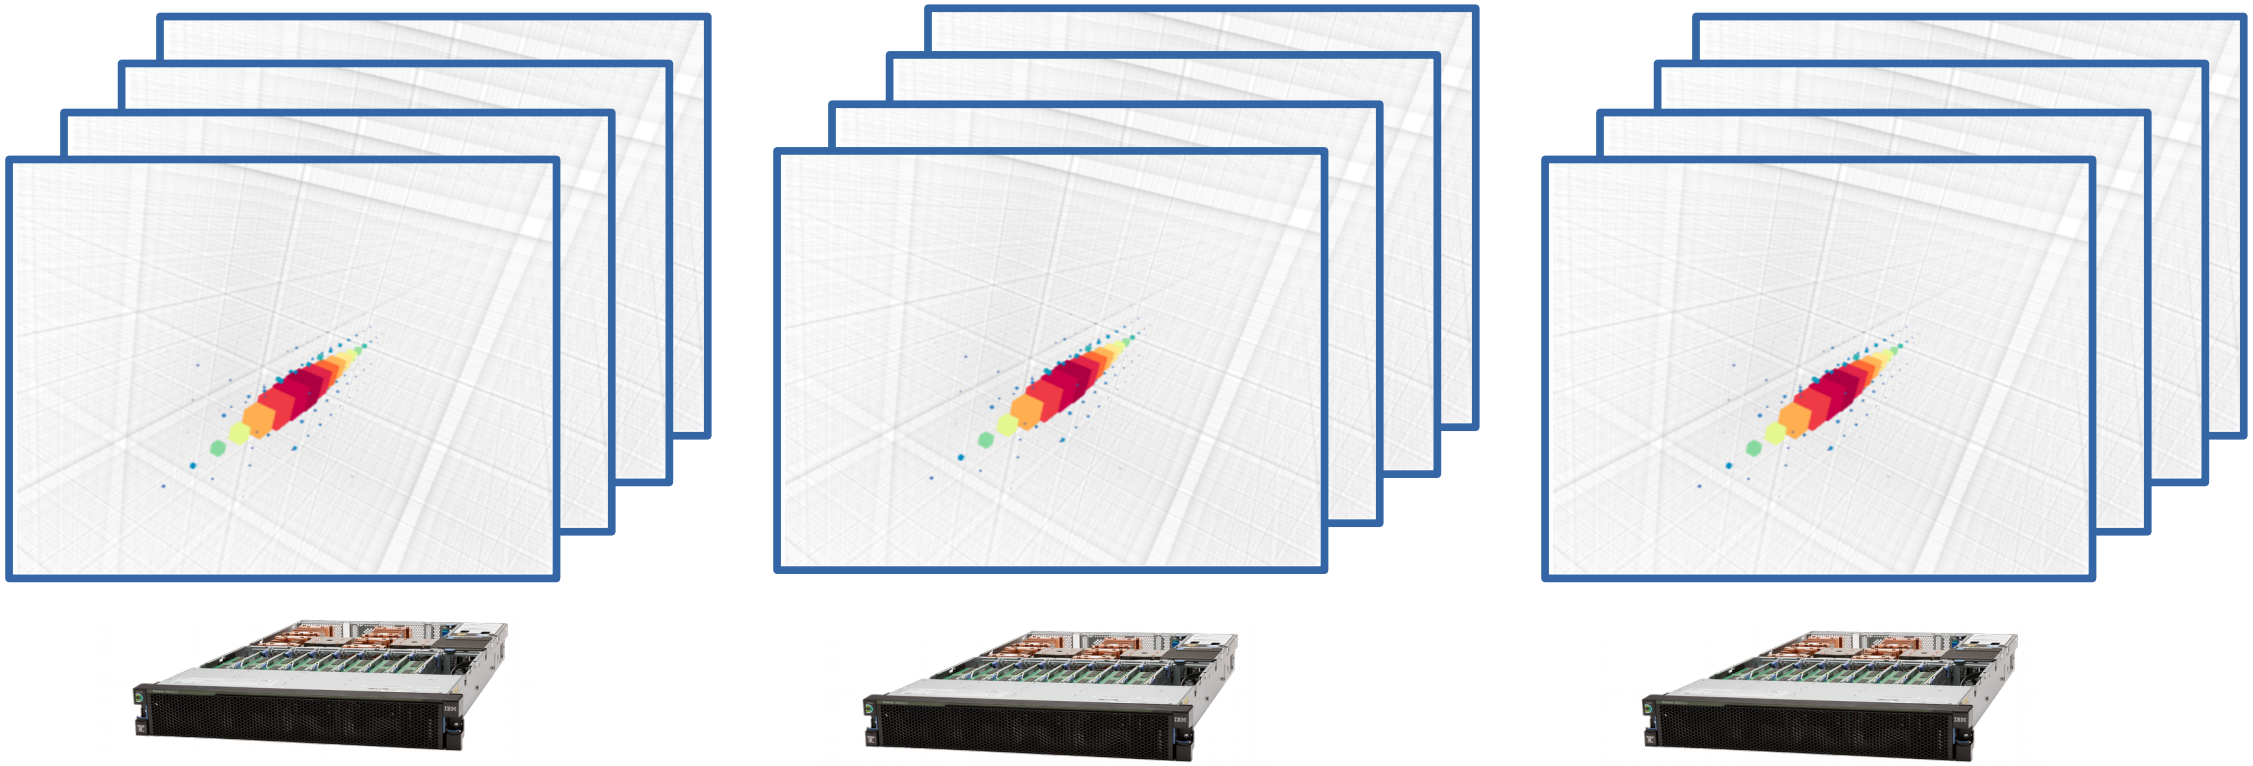

v2.2.2 v1.12.0 v0.15.1

Performance Results: Training

Scaling behavior over multiple nodes / GPUs on IBM Minsky

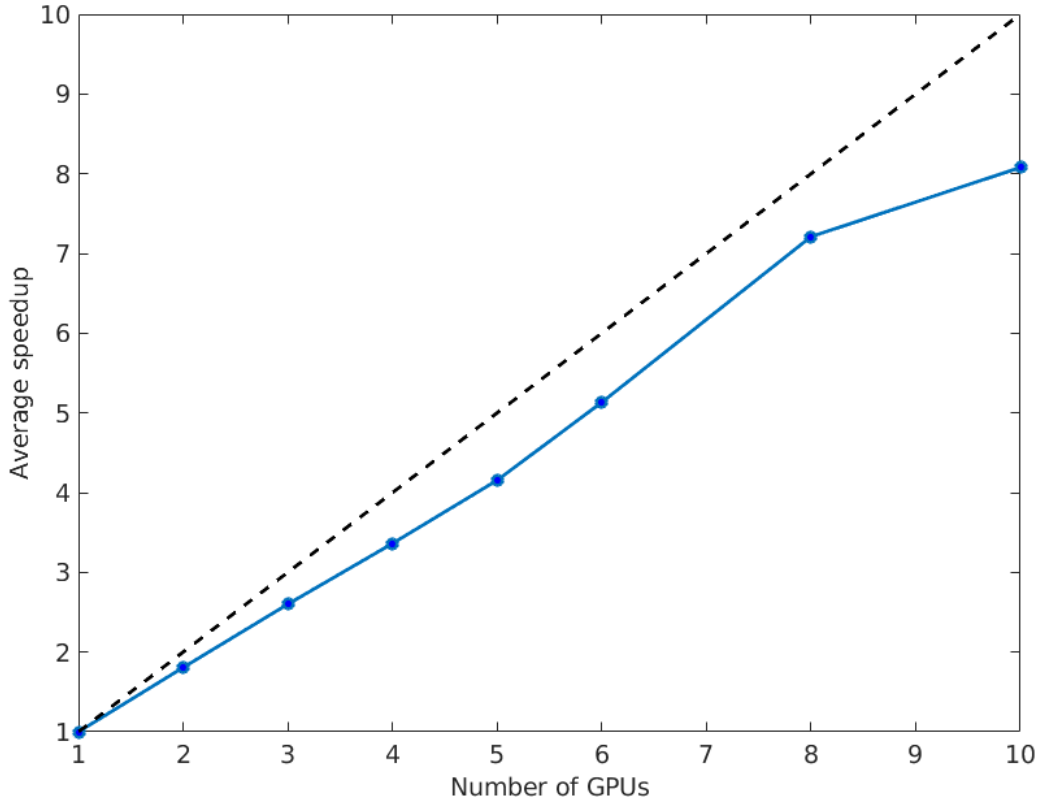

Convergence of losses for distributed learning on IBM Minsky

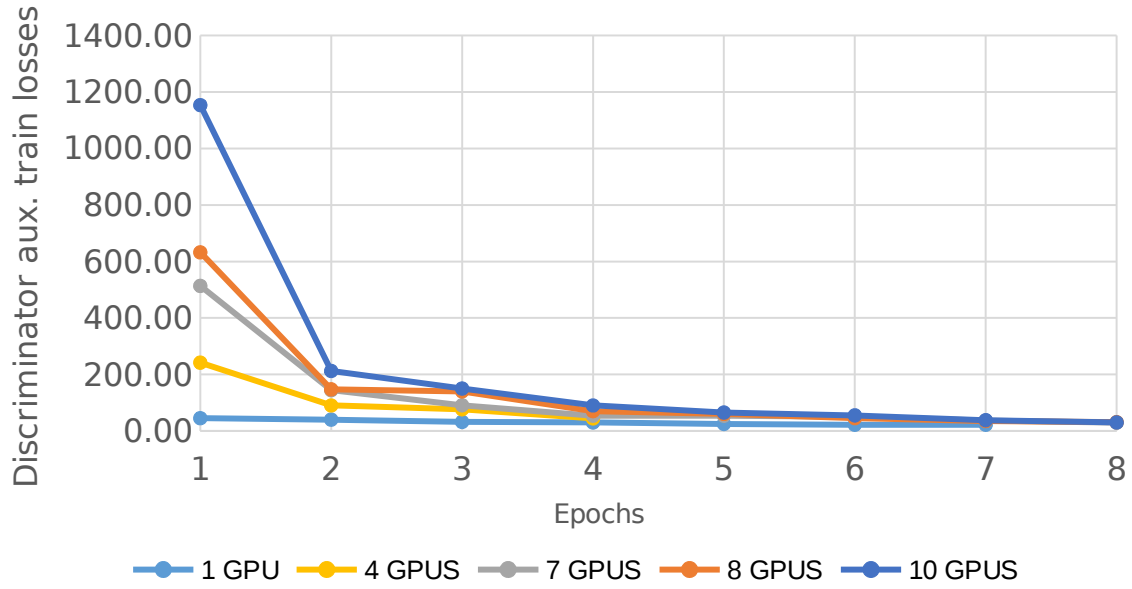

Performance Results: Inference

- Inference on single P100 GPU
- 17 sec → 0.8 ms
- 4 orders of magnitude improvement

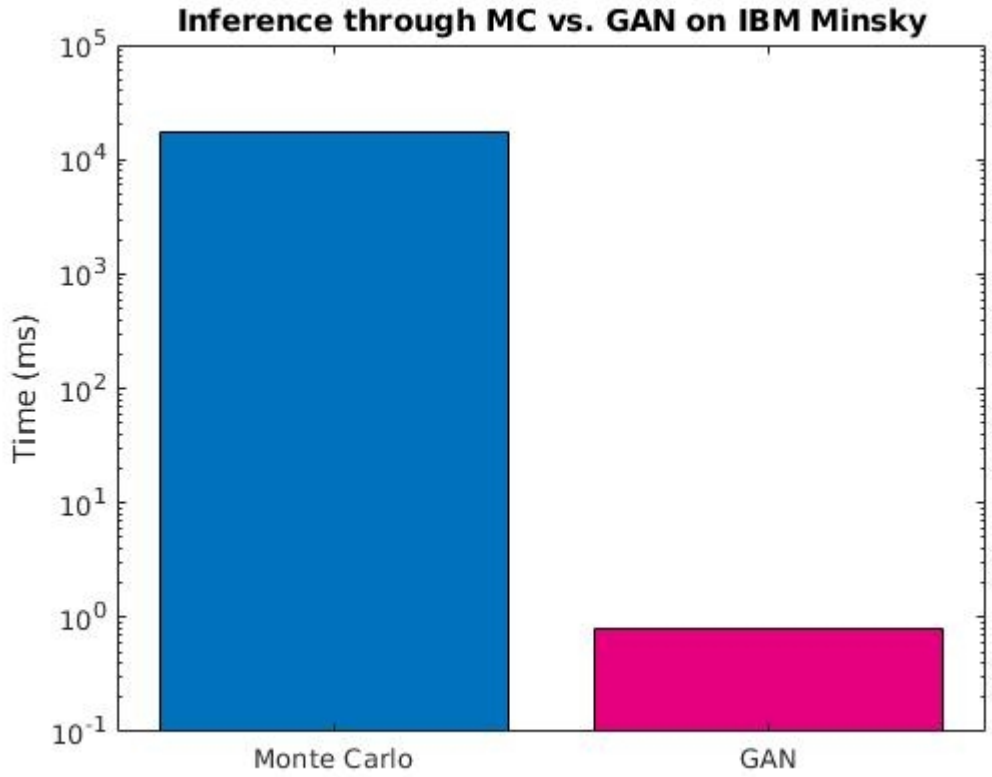

Performance Results: NVLink

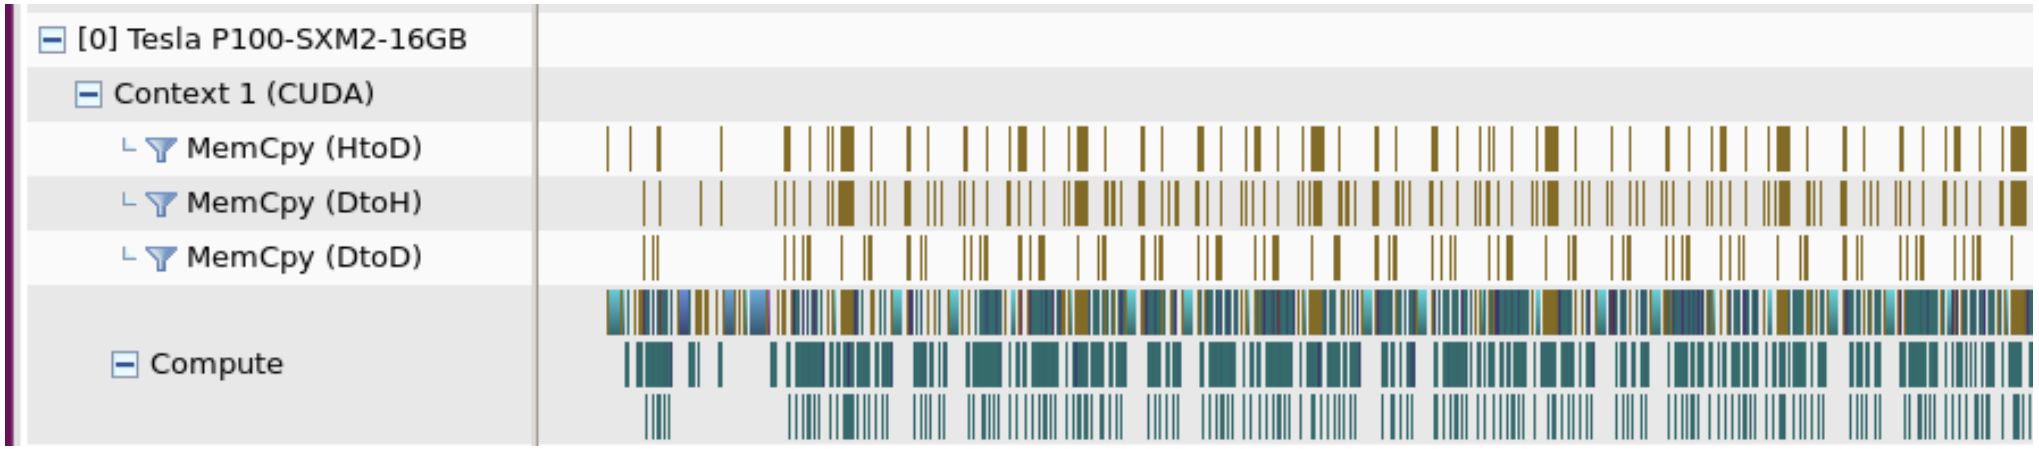

High bandwidth interconnects alleviates stalls for memcpy operations → higher throughput

Conclusion: POWER for DL

- 25 hours of training → 3 hours of training on 10 GPUs over 3 nodes
- Convergence of losses for varying number of GPUs
- NVLink between CPU ↔ GPU enables high throughput

Upcoming PhD Research

- Deep Learning approach for offline event reconstruction in DUNE neutrino experiments
- Find novel methods of harnessing the cutting-edge in HPC
- Large data sets → strain on memory bandwidth → IBM is adapting its architecture towards this trend

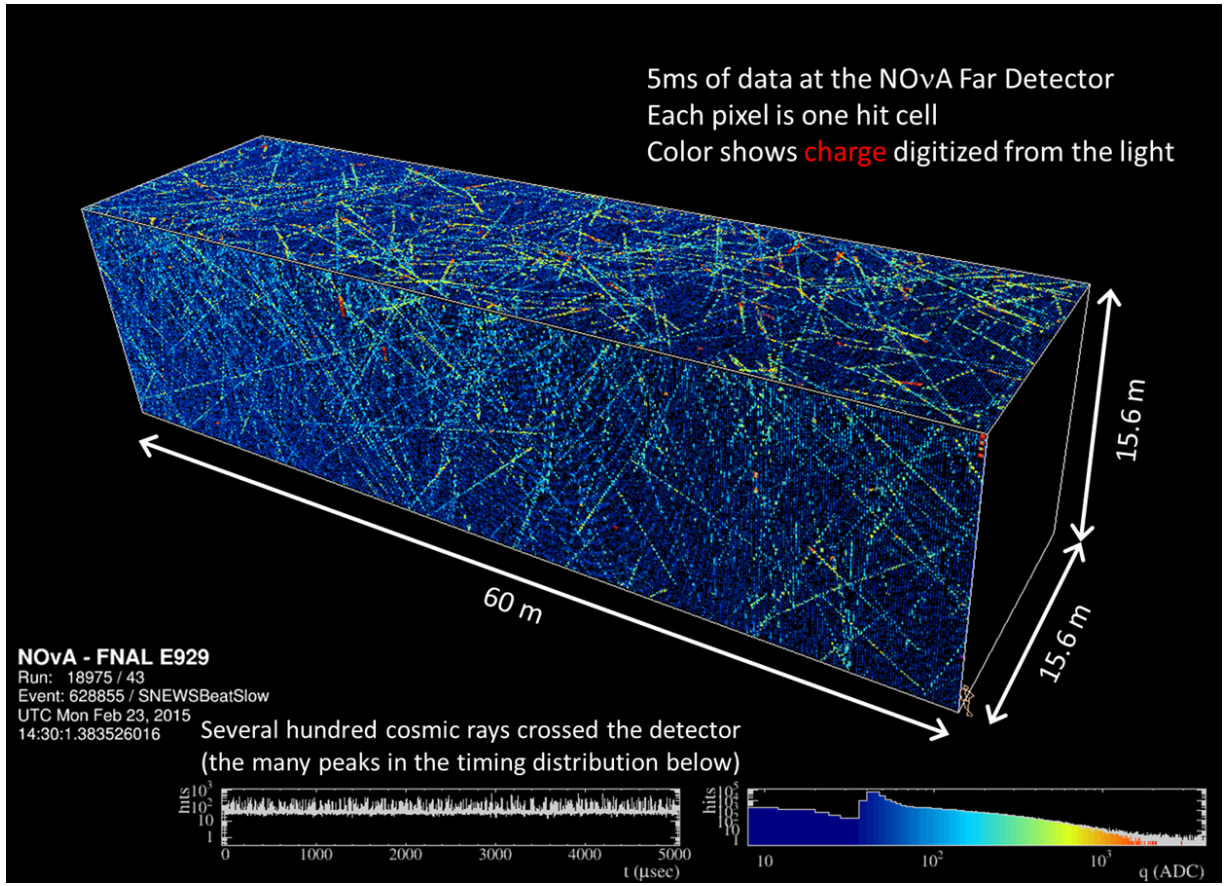

Example of a neutrino event data display (NovA Experiment)

QUESTIONS?

Ahmad Hesam

ahmad.hesam@cern.ch