GPU benchmarking

From Benchmarking WG

Aritz Brosa

ML workloads

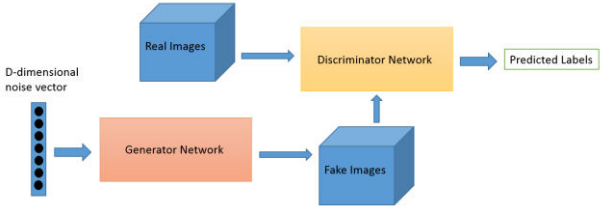

CNN classification

- ✓ Supervised
- ✓ Single convolutional neural network (with many layers)
- ✓ Convergence happens easily
- ✓ Measure: Convergence time

GAN

- ✓ Unsupervised
- ✓ Two neural networks: Generator (reverse CNN)& Discriminator (CNN)
- ✓ Not easy for getting to convergence (long time)
- ✓ Measure: Generated image/sec

Benchmark candidates

CNN

- Atlas_dl

Particle physics event classifier

https://github.com/eracah/atlas_dl/tree/micky

GAN

- cosmoGAN (ATLAS)

Cosmological model convergence maps.

https://github.com/MustafaMustafa/cosmoGAN

- 2DGAN (SFT)

2D calorimeter fast simulation

https://github.com/svalleco/3Dgan/blob/svalleco/2Dtest/keras/Ecal2DBenchmark.py

Not ML: Alice tpc tracker

Platforms' specs

- K20

-Nvidia driver: 390.46
-CUDNN: 7.1.3
-CUDA: 9.1
-tensorflow: 1.8.0
-keras: 2.0.8

- K40

-Nvidia driver: 387.26
-CUDNN: 5.1.5
-CUDA: 8.0
-tensorflow: 1.4.0
-keras: 2.1.1

- **Nvidia driver:** Nvidia proprietary driver for GPUs.
- **CUDA:** Nvidia's framework for parallel computing
- **CUDNN:** CUDA Deep Neural Network library.
- **Tensorflow:** Google's open source Machine Learning framework.
- **Keras:** High-level neural network API, able to run over different frameworks.

- GTX1080

-Nvidia driver: 367.48
-CUDNN: 5.1.5
-CUDA: 8.0
-tensorflow: 1.2.1
-keras: 2.0.4

- GTX1080Ti

-Nvidia driver: 384.81
-CUDNN: 3.0.8
-CUDA:9.0
-tensorflow:1.2.1
-keras: 2.0.4

Benchmark results for GANs

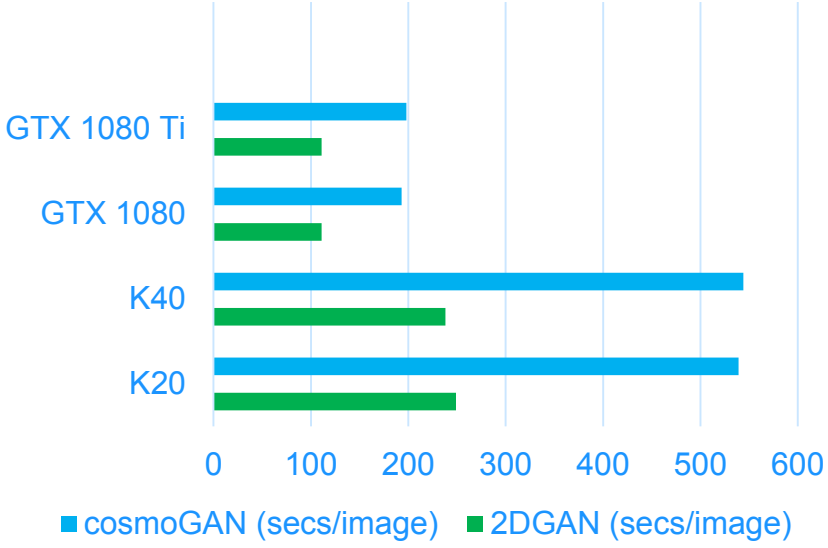

Results atlas_dl

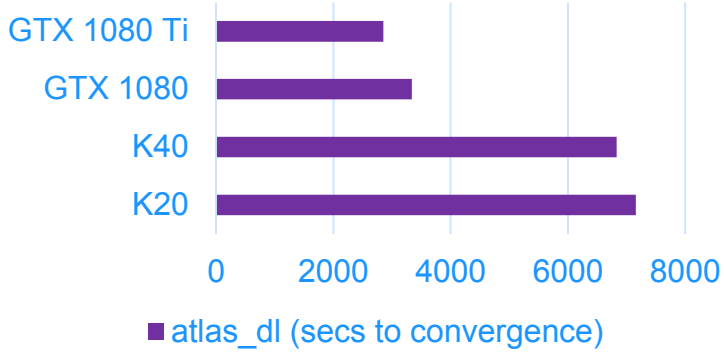

Results Alice tpc

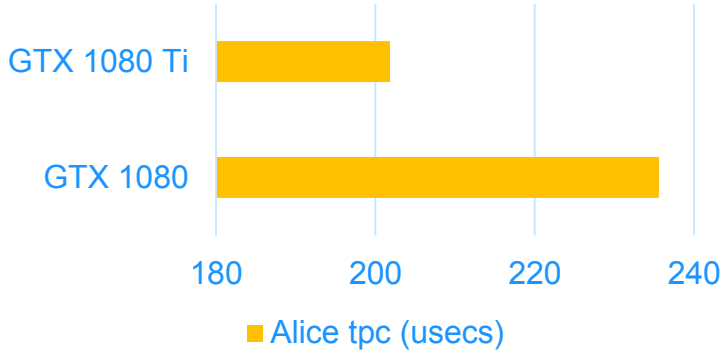


CERN