

Benchmarks with Tensorflow

Updated Jan 25, 2019

ykchoi@harsper.co.kr

Overview

This benchmark is a measure of the GPU's computing power using TensorFlow, a machine learning (deep learning) library provided by google.

Tools

TensorFlow provides benchmark tools on the Github.

```
https://github.com/tensorflow/benchmarks/tree/master/scripts/tf\_cnn\_benchmarks
```

This benchmark tool is based on TensorFlow, CUDA and cuDDN, and performs several kinds of image classification models (Inception V3, ResNet-50, etc.).

Environment

- GPU : NVIDIA GTX 1070 (8GB) / 1080 (8GB)
- OS : Windows Server 2016 x64
- CUDA/cuDDN : 9.0 / 7.3.1
- tf-nightly-gpu : 1.13.0.dev20181104
- Python : 3.6.7 x64
- DataSet : Synthetic


- Configurations for scripts:

```
--batch_size=64 --model=resnet50 --variable_update=parameter_server --local_parameter_device=cpu
```

Results

With ResNet-50 model

GPUs	NVIDIA DGX-1 (Tesla P100)	Amazon EC2 (Tesla K80)	MG140-G1 (GTX 1070)	MG140-G1 (GTX 1080)
1	219	52	111	130
2	422	99	213	259
4	852	195	423	496
8	1734	387	-	-



For information on Tesla P100 and Tesla K80, see following link.

<https://www.tensorflow.org/guide/performance/benchmarks>