

Benchmark Mar_26_2018

References

Tensorflow Official Benchmarks (May 2017, GitHub source): <https://www.tensorflow.org/performance/benchmarks>

IBM Power9 benchmark results (Nov 2017, 1.4.0): <https://developer.ibm.com/linuxonpower/perfcol/perfcol-mld/>

Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour, Facebook (Jun 2017): <https://research.fb.com/wp-content/uploads/2017/06/imagenet1kin1h5.pdf>

Benchmark Source Code

<https://github-dev.cs.illinois.edu/kindrtnk/DL>

Official TF Benchmark System Characteristics

- **Instance type:** NVIDIA® DGX-1™
- **GPU:** 8x NVIDIA® Tesla® P100
- **OS:** Ubuntu 16.04 LTS with tests run via Docker
- **CUDA / cuDNN:** 8.0 / 5.1
- **TensorFlow GitHub hash:** b1e174e
- **Benchmark GitHub hash:** 9165a70
- **Build Command:** `bazel build -c opt --copt=-march="haswell" --config=cuda //tensorflow/tools/pip_package:build_pip_package`
- **Disk:** Local SSD
- **DataSet:** ImageNet
- **Test Date:** May 2017

Our System Characteristics (more details in GitHub Repo)

- **Instance type:** IBM Power9, 8335-GTG AC922 server
- **CPU:** 2x 20-core IBM POWER9 CPU @ 2.00GHz
- **SDRAM:** 512G DDR4
- **GPU:** 4x NVIDIA® Tesla® V100, 5120 cores, 16 GB HBM 2
- **Disk:** Local SSD
- **OS:** Red Hat Enterprise Linux Server release 7.4
- **Python Distribution:** Anaconda python 3.6.2
- **CUDA / cuDNN:** 9.1/7.0.5
- **TensorFlow Version:** 1.5.0
- **DataSet:** ImageNet (synthetic)
- **Precision:** floating point 32 and 16
- **Test Date:** Mar 25 2018

The following table is the result of running with the same configurations as the official Tensorflow benchmark mentioned in "Reference" section above:

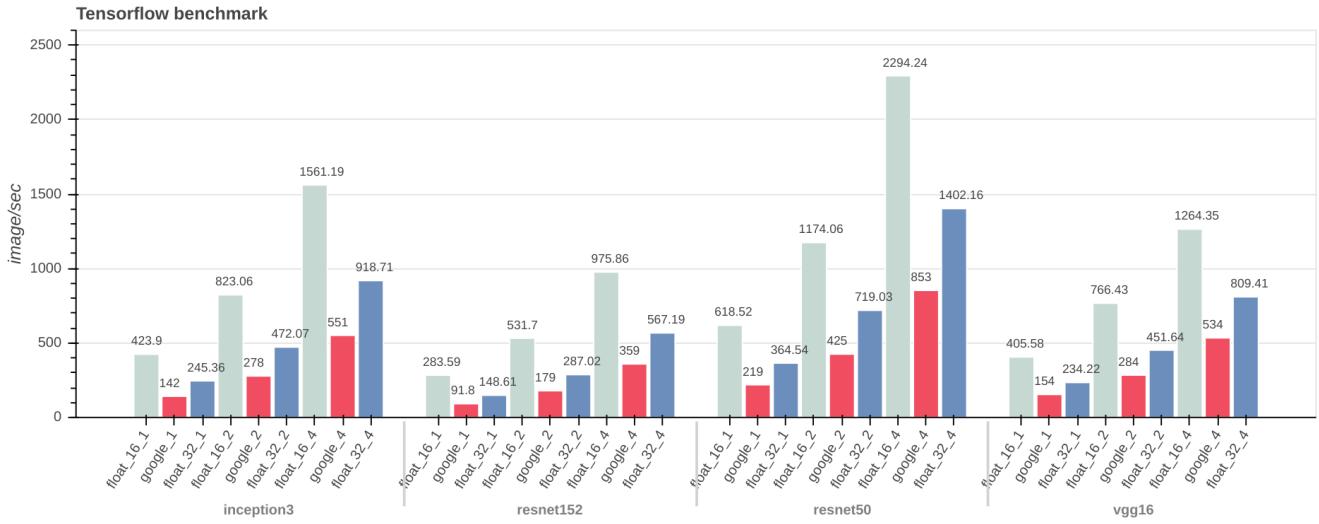
 Unknown macro: 'csv'

This figure compares the result we get with Tensorflow official ones.

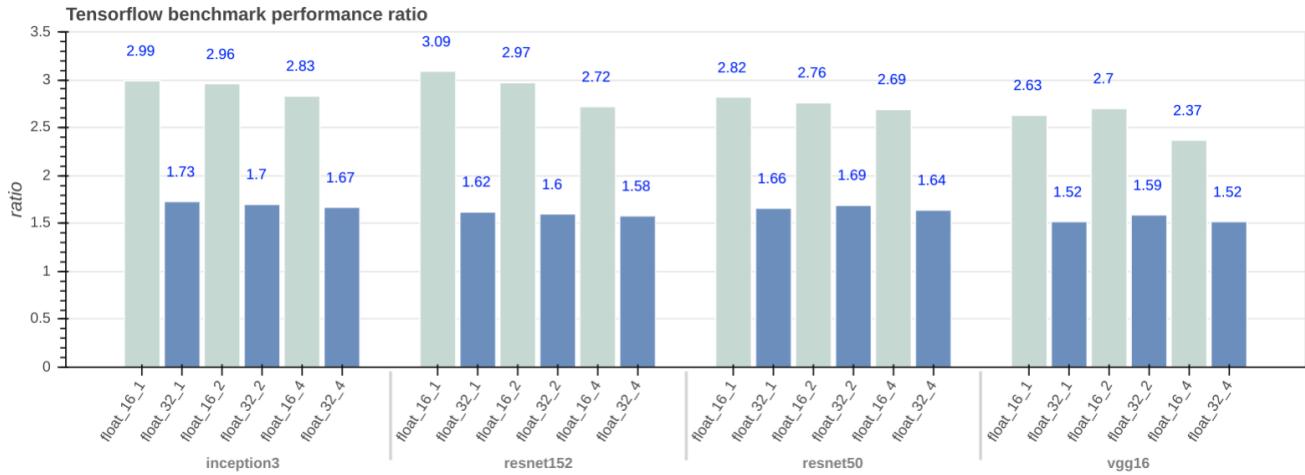
Green bars stand for our benchmark results using floating point 16.

Red bars are the official Tensorflow result.

Blue bars stand for our benchmark results using floating point 32.



This figure shows the performance ratio of our floating point 16 and 32 benchmarks with respect to Tensorflow official results:



The following table provides a more comprehensive benchmark results on our system:

Unknown macro: 'csv'

POWER8 (p8)