

Horovod Cheatsheet

Installation and Setup

pip install horovod[tensorflow,pytorch,mxnet]
Install Horovod with frameworks support

horovodrun --check-build
Verify Horovod installation and supported frameworks

Basic Usage

horovodrun -np 4 -H localhost:4 python script.py
Run script with 4 processes on local machine

**horovodrun -np 16 -H
server1:4,server2:4,server3:4,server4:4 python
script.py**
Run script on 4 servers with 4 processes each

Framework Integration

import horovod.tensorflow as hvd
Import Horovod for TensorFlow

import horovod.torch as hvd
Import Horovod for PyTorch

import horovod.mxnet as hvd
Import Horovod for MXNet

hvd.init()
Initialize Horovod

hvd.size()
Get number of processes

hvd.rank()
Get rank of current process

hvd.local_rank()
Get local rank within node

Distributed Operations

hvd.allreduce(tensor, name='allreduce')
Average tensor across all processes

hvd.allgather(tensor, name='allgather')
Gather tensors from all processes

**hvd.broadcast(tensor, root_rank=0,
name='broadcast')**
Broadcast tensor from root rank to all processes

**hvd.broadcast_parameters(model.state_dict(),
root_rank=0)**
Broadcast model parameters (PyTorch)

hvd.broadcast_variables(tf_variables, root_rank=0)
Broadcast variables (TensorFlow)

Optimizer Wrapping

hvd.DistributedOptimizer(optimizer)
Wrap optimizer for distributed training

**opt = hvd.DistributedOptimizer(opt,
backward_passes_per_step=1)**
Set backward passes per step

Advanced Options

horovodrun --timeline-filename timeline.json
Generate timeline for performance analysis

horovodrun --verbose
Enable verbose logging

horovodrun --gloo
Force using Gloo as communication backend

horovodrun --mpi
Force using MPI as communication backend

horovodrun --nccl
Force using NCCL as communication backend

Environment Variables

export HOROVOD_GPU_OPERATIONS=NCCL
Set GPU operations backend

export HOROVOD_CPU_OPERATIONS=MPI
Set CPU operations backend

export HOROVOD_TIMELINE=timeline.json
Enable timeline recording

export HOROVOD_FUSION_THRESHOLD=67108864
Set tensor fusion threshold (bytes)