# ParaGAN: A Cloud Training Framework for Generative Adversarial Networks



Ziji Shi<sup>1,2</sup>, Fuzhao Xue<sup>1</sup>, Jialin Li<sup>1</sup>, Yang You<sup>1</sup> <sup>1</sup>National University of Singapore, <sup>2</sup>Alibaba Group

# (-) Alibaba Cloud

# Challenge: Large-Scale GAN Training

#### a. Numerical instability

- Generative Adversaria Network has two subnetworks:
  - Generator (**G**): synthesize images
  - Discriminator (D): distinguish fake images from the real ones
- Convergence is often not guaranteed
  - Two subnetworks are optimized for contradicting goals, and the convergence is defined as a Nash Equilibrium
  - **G** may be overly optimized for **D**, producing poor image results



#### Fig. 1: GAN training process.

Hardware

8 V100 GPUs

8 V100 GPUs

8 V100 GPUs

8 V100 GPUs

trained on ImageNet 2012 dataset.

GANs

SNGAN [40]

SAGAN [54]

BigGAN [4]

ContraGAN [24]

 Table 1: Training Time and Parameters Number for GANs

Training Time Parameters

81.44M

81.47M

158.42M

160.78M

3d 13.6h

10d 18.7h

30d 6.6h

5d 3.5h

## Asynchronous Training

#### a. Motivation

- Previous works optimize both generator and discriminator in the same manner
- Can we decouple it by optimizing **G** and **D** separately? •

### b. Asynchronous training

- Let **G** and **D** run on different nodes within the same TPU Chip (each chip has two nodes)
- Sync gradient globally every other step, and cache the generated fake image
- Why it works?
  - Add perturbation to the gradient (no long update each other immediately)
  - Force **G** not to mimic a limited set of modes

Also save memory and communication => enables the training of larger models



- Successful GAN models take days to weeks to train
- Issues with datacenter-level distributed training:
  - Network congestion prolongs unpredictable latency
  - Each HW accelerator has different characteristics
- Existing works emphasize little on training speed

#### Approach: Numerical and System Optimization Co-Design

- a. Numerical optimizations for training stability
- Asynchronous Training
- Asymmetric Optimization

#### a. System optimizations for scalability

- Model size is rather small => Data parallelism suffices
- Congestion aware data pipeline
  - Dynamically adjust the prefetch buffer size •
  - Moderate the impact of network jittering

Training loss of BigGAN with different optimizers





-Sync G-128 D-128 Async G-128 D-128 Async G-128 D-64 100-50(a) CIFAR-10 Sync G-1024 D-1024 140Sync G-512 D-512 Async G-512 D-512 120Async G-512 D-256 100€ 80 60 ----- $\cdot 10^{4}$ (b) Tiny-ImageNet

-Sync G-256 D-256

Fig. 6: Async update scheme.

Fig. 7: FID of both update schemes. Lower is better.

## Evaluation

- Hardware-aware layout transformation
  - Improve accelerator utilization
    - > minimize zero padding
    - > convert tensor to HW-friendly data layouts
  - Support CPU/GPU/TPU



Fig. 3: Profiling results of individual operators on TPU. Model shown is BigGAN.



Fig. 4: Overview of ParaGAN.

## Example and Implementation

import paragan as pg

- 1. BigGAN end-to-end training: reduced from 15 days to 14 hours.
- 2. Scales to 1024 TPU nodes at 91% efficiency.
- 3. Enables direct generation of 1024x1024 images.





Fig. 9: Image generated at 1024x1024 resolution.



Fig. 8: Scaling on 1024 TPU nodes.



64

class Generator: def model\_fn(latent\_var, y): # generator model return output

class Discriminator: def model\_fn(x, y): # discriminator model return output, out\_logit

scale\_mgr = pg.ScalingManager(config=cfg, bs=2048, num\_workers=128)

g = Generator()d = Discriminator() gan = pg.Estimator(g,d)

#### # train

for step in cfg.max\_steps: scale\_mgr.train(gan)

# evaluate scale\_mgr.eval(metric='fid')

Fig. 5: ParaGAN example.

#### a. Scaling manager

- In charge of hyper-parameter tuning
- Scaling strategy is set based on the profile collected from the single-worker

#### b. Network backbone

- Implements several popular GAN backbones
- Easily extensible to other architectures

c. Evaluation metrics

- Standardized evaluation metrics
  - Frechet Inceptions Distance(FID),
  - Inception Score (IS)
  - • •

b) Images per second.

Number of accelerators

128 256 512 1024

Fig. 10: Weak scaling on BigGAN (128x128)

Limitation & Future Work

. More GAN architectures shall be evaluated. 2. Scale on larger clusters made of GPU 3. Adapt to stable diffusion models

#### ISCA'23 ML for Computer Architecture and Systems (MLArchSys) Workshop | Orlando, FL, USA