FUJITSU

shaping tomorrow with you

Mitigating the Impact of Tail Latency of Storage Systems on Scalable Deep Learning Applications

Hiroki Ohtsuji[1], Erika Hayashi[1], Naoto Fukumoto[1], Eiji Yoshida[1]
Takuya Okamoto[2], Takeru Kuramoto[3], Osamu Tatebe[3]
[1]Fujitsu Laboratories Ltd.
[2]Fujitsu Ltd.
[3]University of Tsukuba

Massive-scale Fast Distributed Deep Learning

FUJITSU

- ResNet-50/ImageNet (1000x1000 Images, 120 GB) training in **a minute**
  - All training processes read the data sets during the training
  - Can shared storage systems provide enough access performance to those workloads?

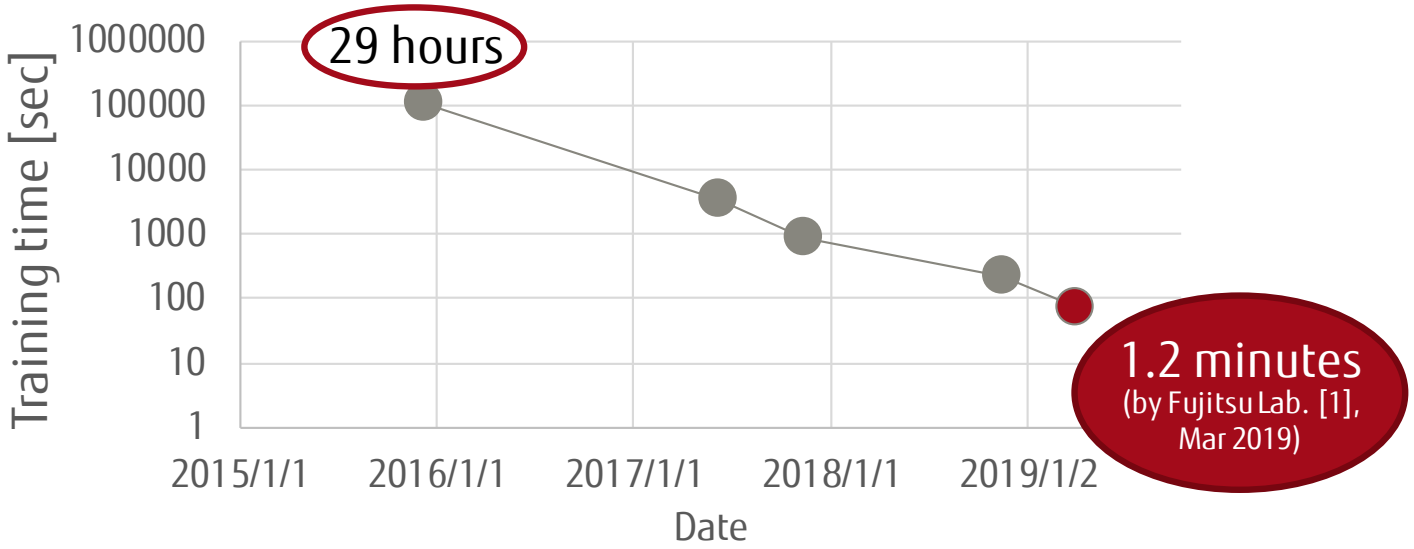

[1] Masafumi Yamazaki et al. Yet Another Accelerated SGD: ResNet-50 Training on ImageNet in 74.7 seconds. *CoRR*, abs/1903.12650, 2019.

Data parallel Distributed Deep Learning

FUJITSU

- More than 1K processes read training data sets from a shared storage

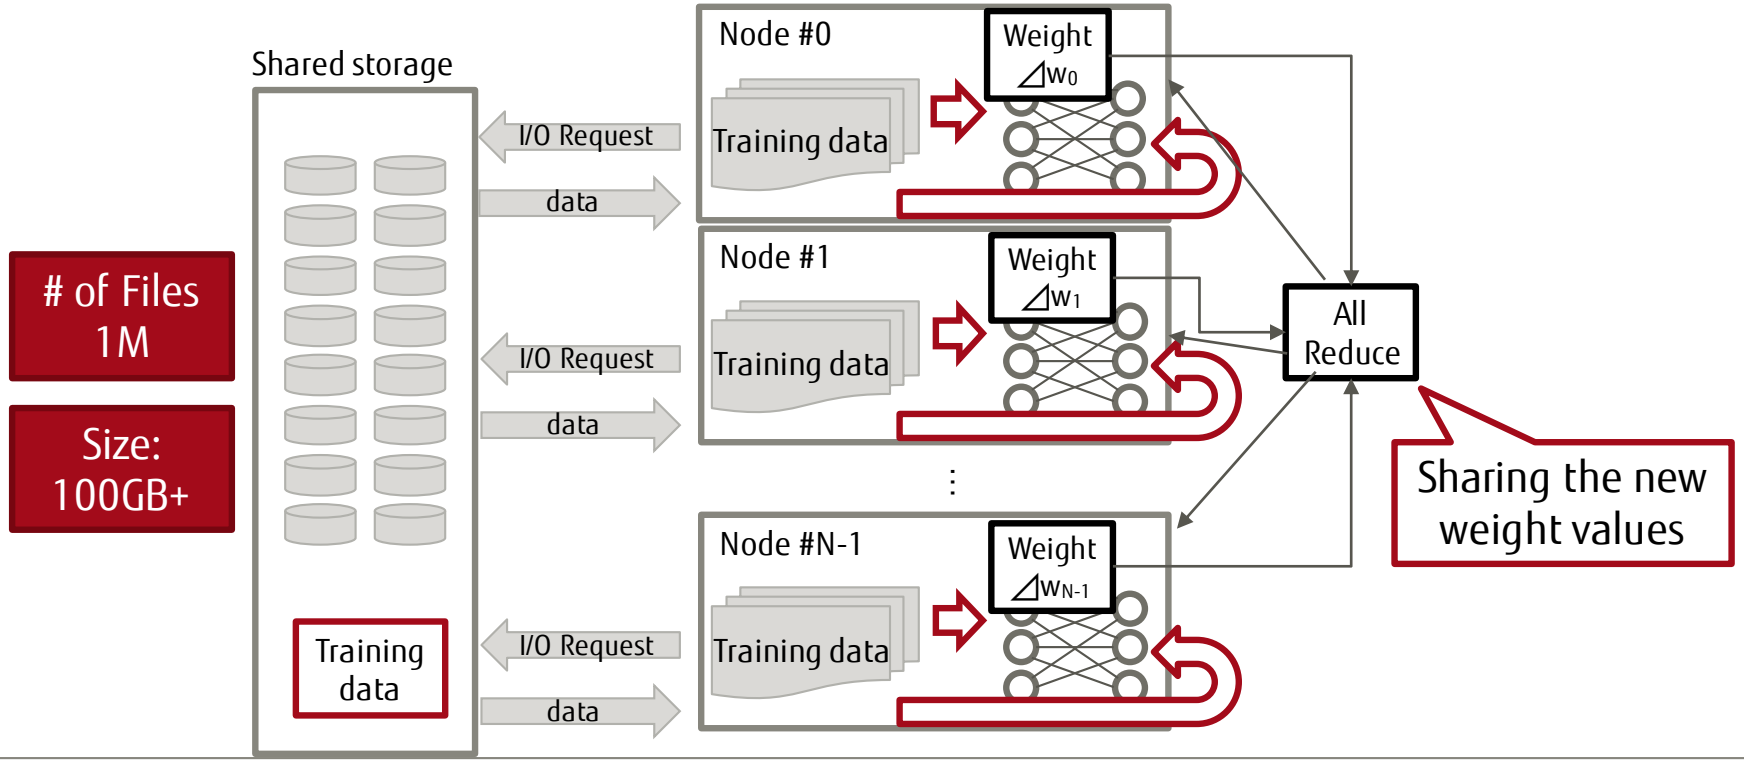

The Impact of Delayed I/O Requests

FUJITSU

- A single delayed I/O can block the entire training process
  - A few seconds of tail latency is harmful due to the shortened training time
    - If the training takes **26 hours**, a 2 seconds delay does not matter.
    - However, a 2 seconds delay in a **1-minute** training process is **significant**.

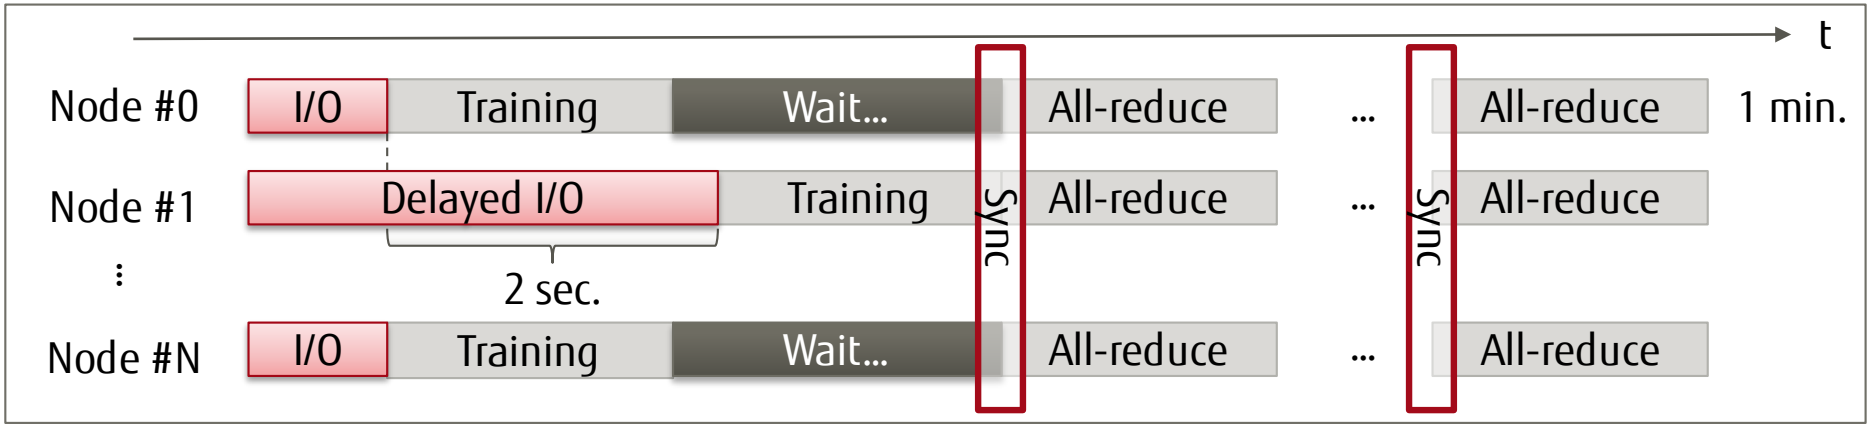

Synchronized processing makes the workload more sensitive to delayed I/Os.

Observation: Tail Latency of Storage Systems

FUJITSU

- Latency distribution when 16 processes read the ImageNet data set from the shared storage system
  - Measured on Cygnus Supercomputer at Univ. of Tsukuba

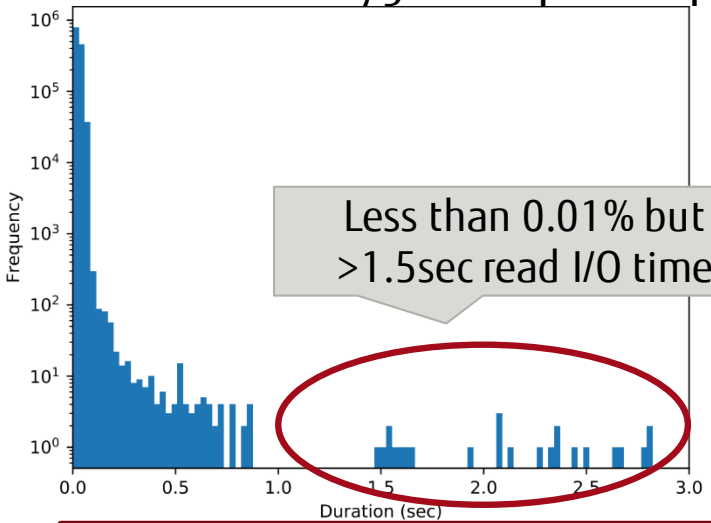

>99.99% of requests complete in **0.5 sec**
<0.01% of requests take more than **1.5 sec**

There are very rare (<0.01%) long-delayed I/O requests on a shared storage system

I/O Latency Aware Dynamic I/O Replacement

FUJITSU

- Monitoring all read I/O operations and...
  - Discarding/replacing the delayed I/O with alternative data sets
  - Managing the history of data replacement to minimize the impact on the training.

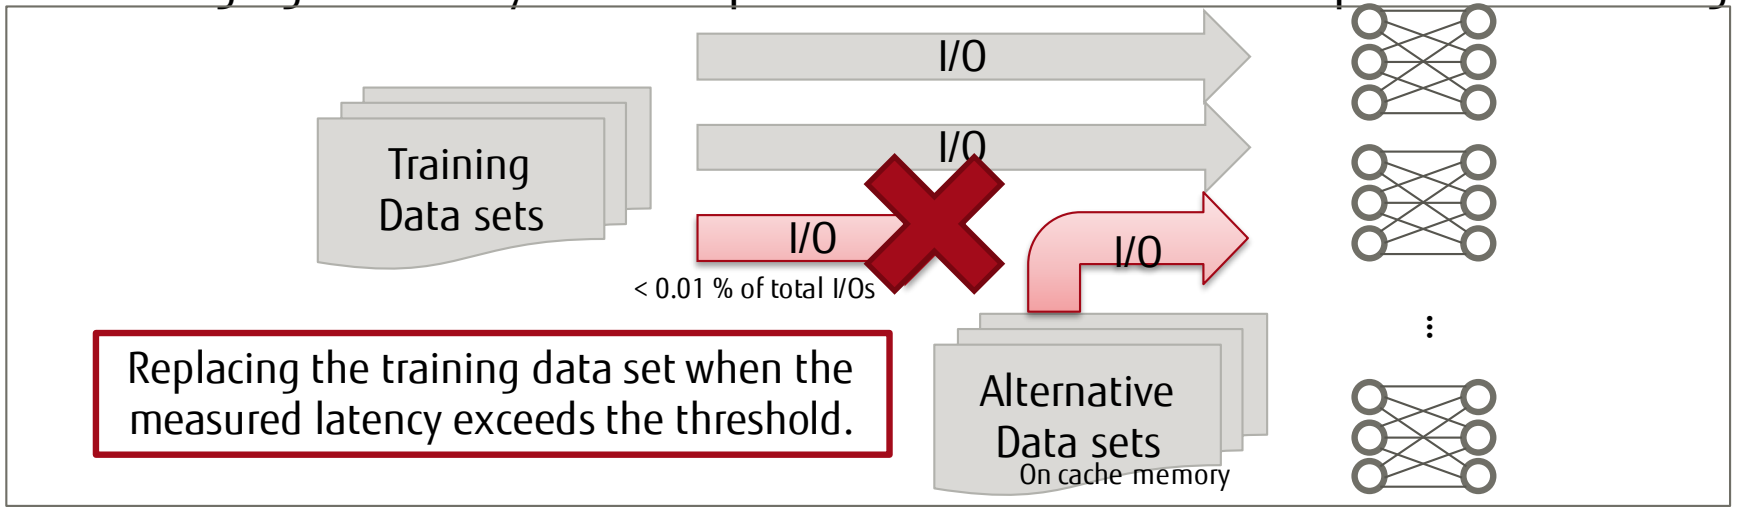

Avoiding the impact of (rare) long-delayed I/O
with the negligible drawback in precision of the trained model

Conclusion and Future work

FUJITSU

- I/O tail latency has a significant impact on the performance of massive scale fast deep learning applications.
- Shared storage systems have very rare long tail latency, that affects to the performance of distributed deep learning workloads.
- We are developing a method to avoid the impact of delayed I/O requests on training processes.
  - Dynamic I/O Replacement
    - Replacing less than 0.01 % of total I/O operations.
    - Adaptive data replacement to minimize the drawbacks in the precision of the trained model.

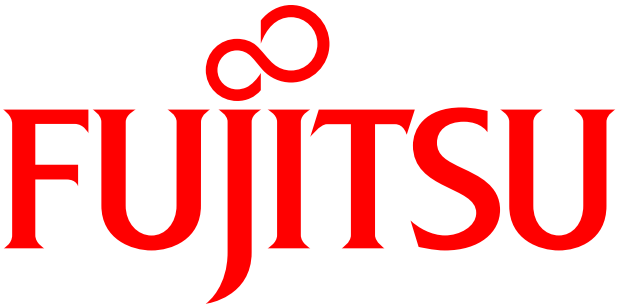

shaping tomorrow with you