# ABCI 2.0: opportunities and challenges of an open research platform for AI/ML workload

National Institute of Advanced Industrial Science and Technology
Ryousei Takano

2nd ML Hardware Workshop, 3 July 2021

# Introduction

- Modern Artificial Intelligence (i.e., deep learning) and big data are enablers of digital transformation in multiple industries.
  - Classical AI used logical rules to model human intelligence.
  - Modern AI used a set of data and machine learning techniques to solve many different classes of problems, including smart cities, etc.

- How can the computer systems support modern AI?
- How can we make the computer systems more environmentally friendly and inclusive?

# AI is now a supercomputing problem

The amount of computation required to train DL models has been increasing exponentially with a 3.4 month doubling time [OpenAI]

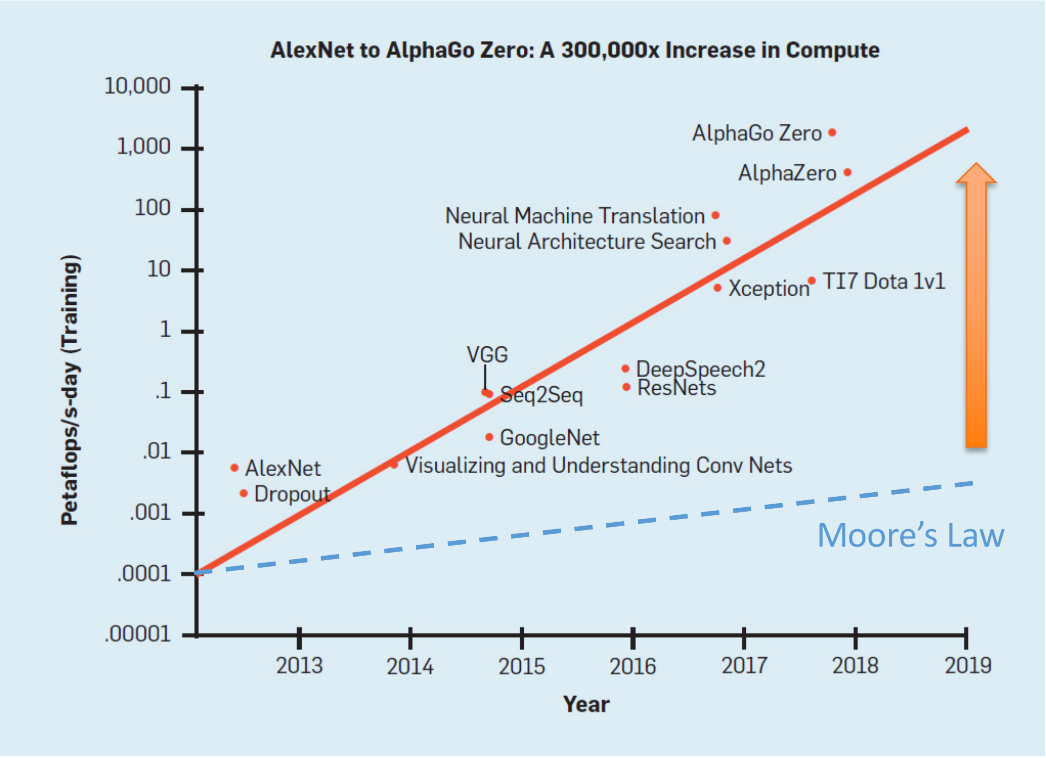


R. Schwartz, et al., "Green AI," CACM, Dec. 2020

The growth in effective computation available largest experiments (estimate)

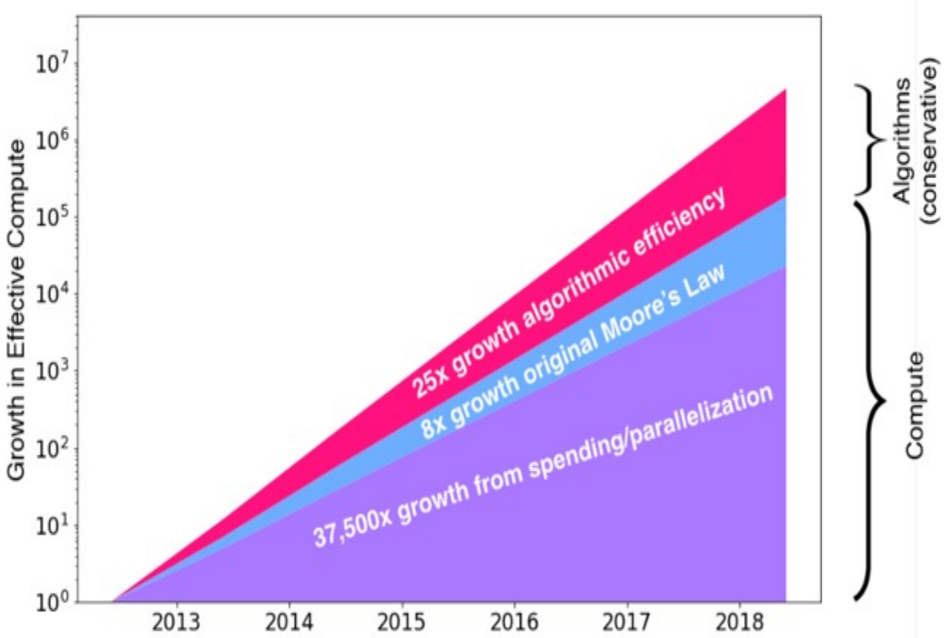


D. Hernandez and T. B. Brown, "Measuring the Algorithmic Effi... ...cy of Neural Networks," https://arxiv.org/pdf/2005.04305.pdf

# ABCI: The World’s First Large-Scale Open AI Infrastructure

Univ. Tokyo / Kashiwa II Campus

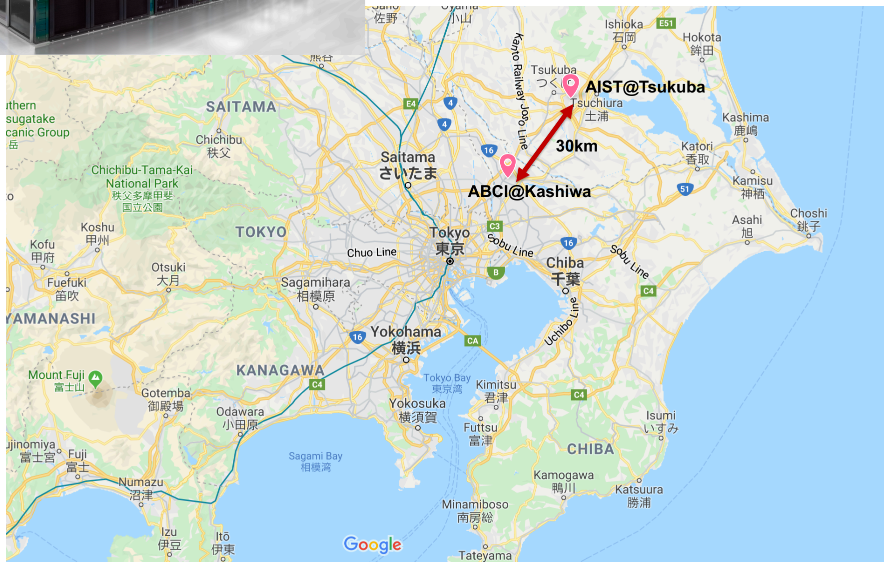


ABCI AI Bridging Cloud Infrastructure

- World Top-Level compute and data process capability
- **Open, Public, and Dedicated** infrastructure for AI & Big Data Algorithms, Software, and Applications
- **Open Innovation Platform** to accelerate joint academic-industry R&D for AI

**Free access to ABCI to support research related to COVID-19.**

**Operation: October 2018~**
**Upgrade to 2.0: May 2021~**

# AI Infrastructure for Everyone

## Expert

- ABCI Grand Challenge: Demonstration of highly challenging academic and/or industrial themes using the whole ABCI resources for 24 hours
- ABCI Data Challenge: Competition of accelerating open science using open data

## Advanced & Intermediate

- Up to 512-node computing resource is available for everyone
- Software, datasets, and pre-trained models are ready for use
- High computing capability enables to accelerate AI R&D and promote social implementation

## Beginner

- User friendly WebUI based IDE for supporting beginners of deep learning
- PoC platform of B2B2C

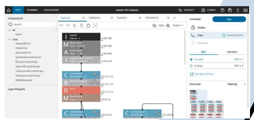

Building ecosystem of AI R&D from beginners to state-of-the-art researches

**Promoting the use by over 100 institutions and over 1000 researchers/engineers**

# ABCI: The numbers of users and projects

- Since the start of operation, the number of users and projects has steadily increased.
- The total number of users is about 2500 (external user of 86%) and the total number of projects is about 360 (external use of 60%).
- The utilization and revenue have increased more than four times since FY2018.

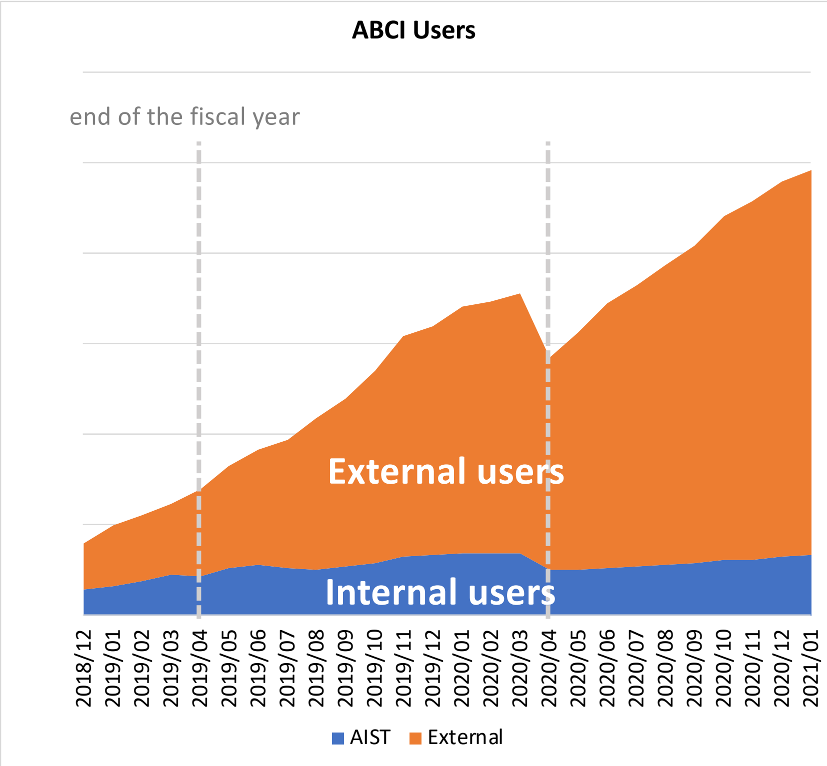


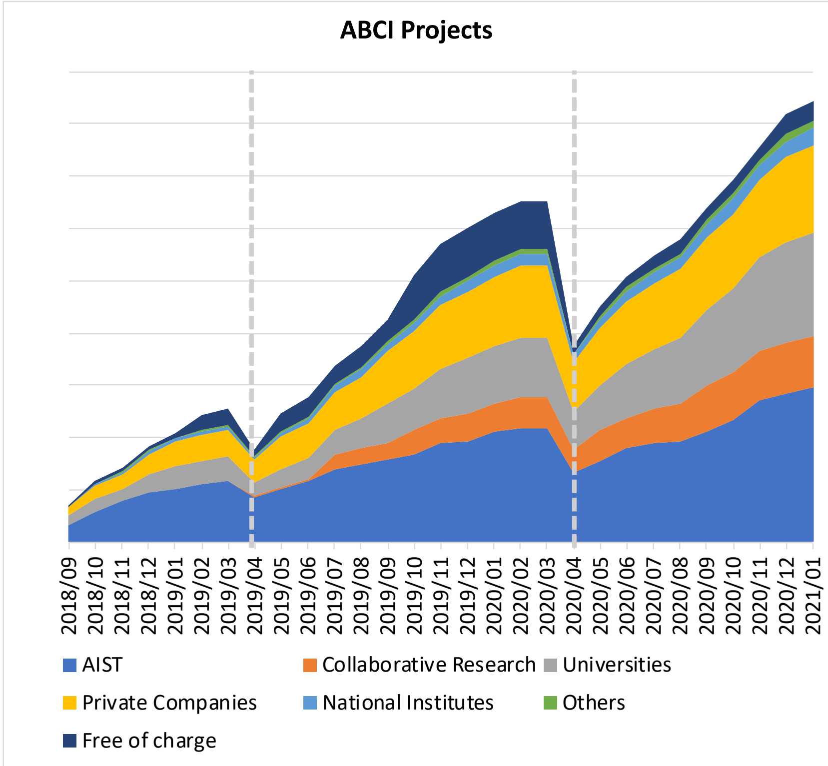

# ABCI: Long-tailed distribution of application fields

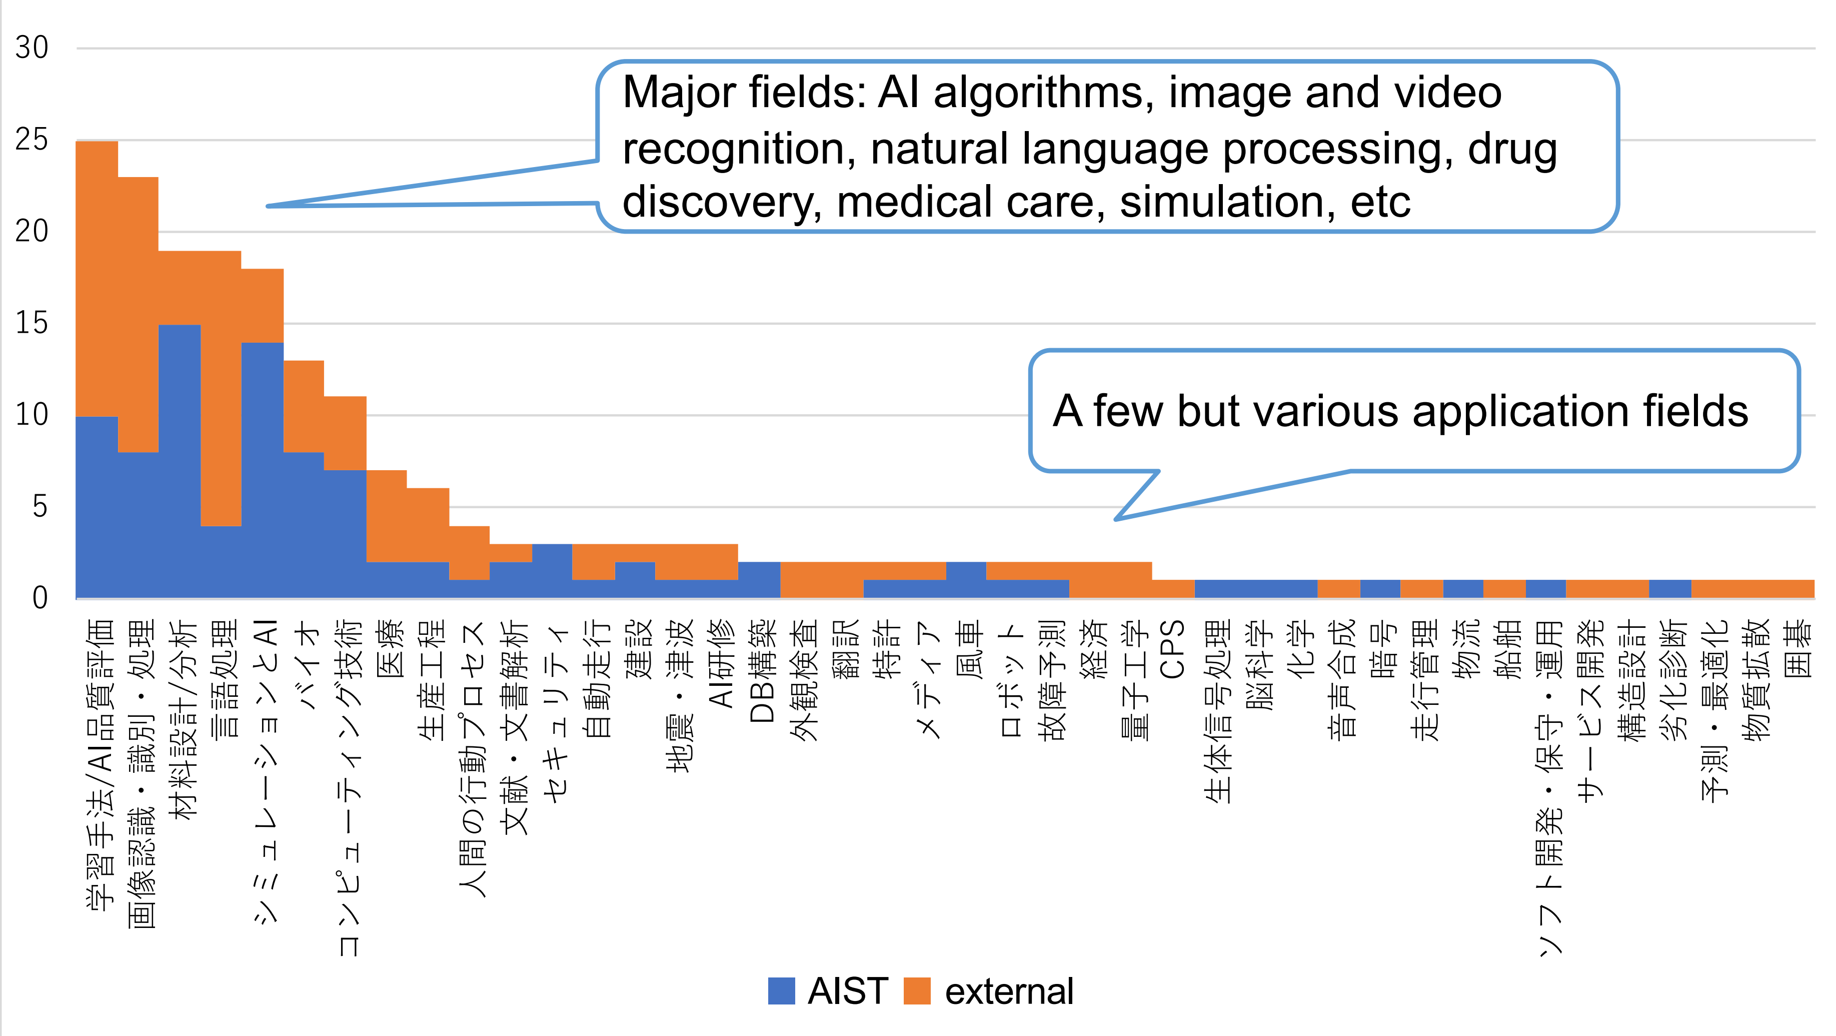

# MLPerf Training v0.6 – Image Classification

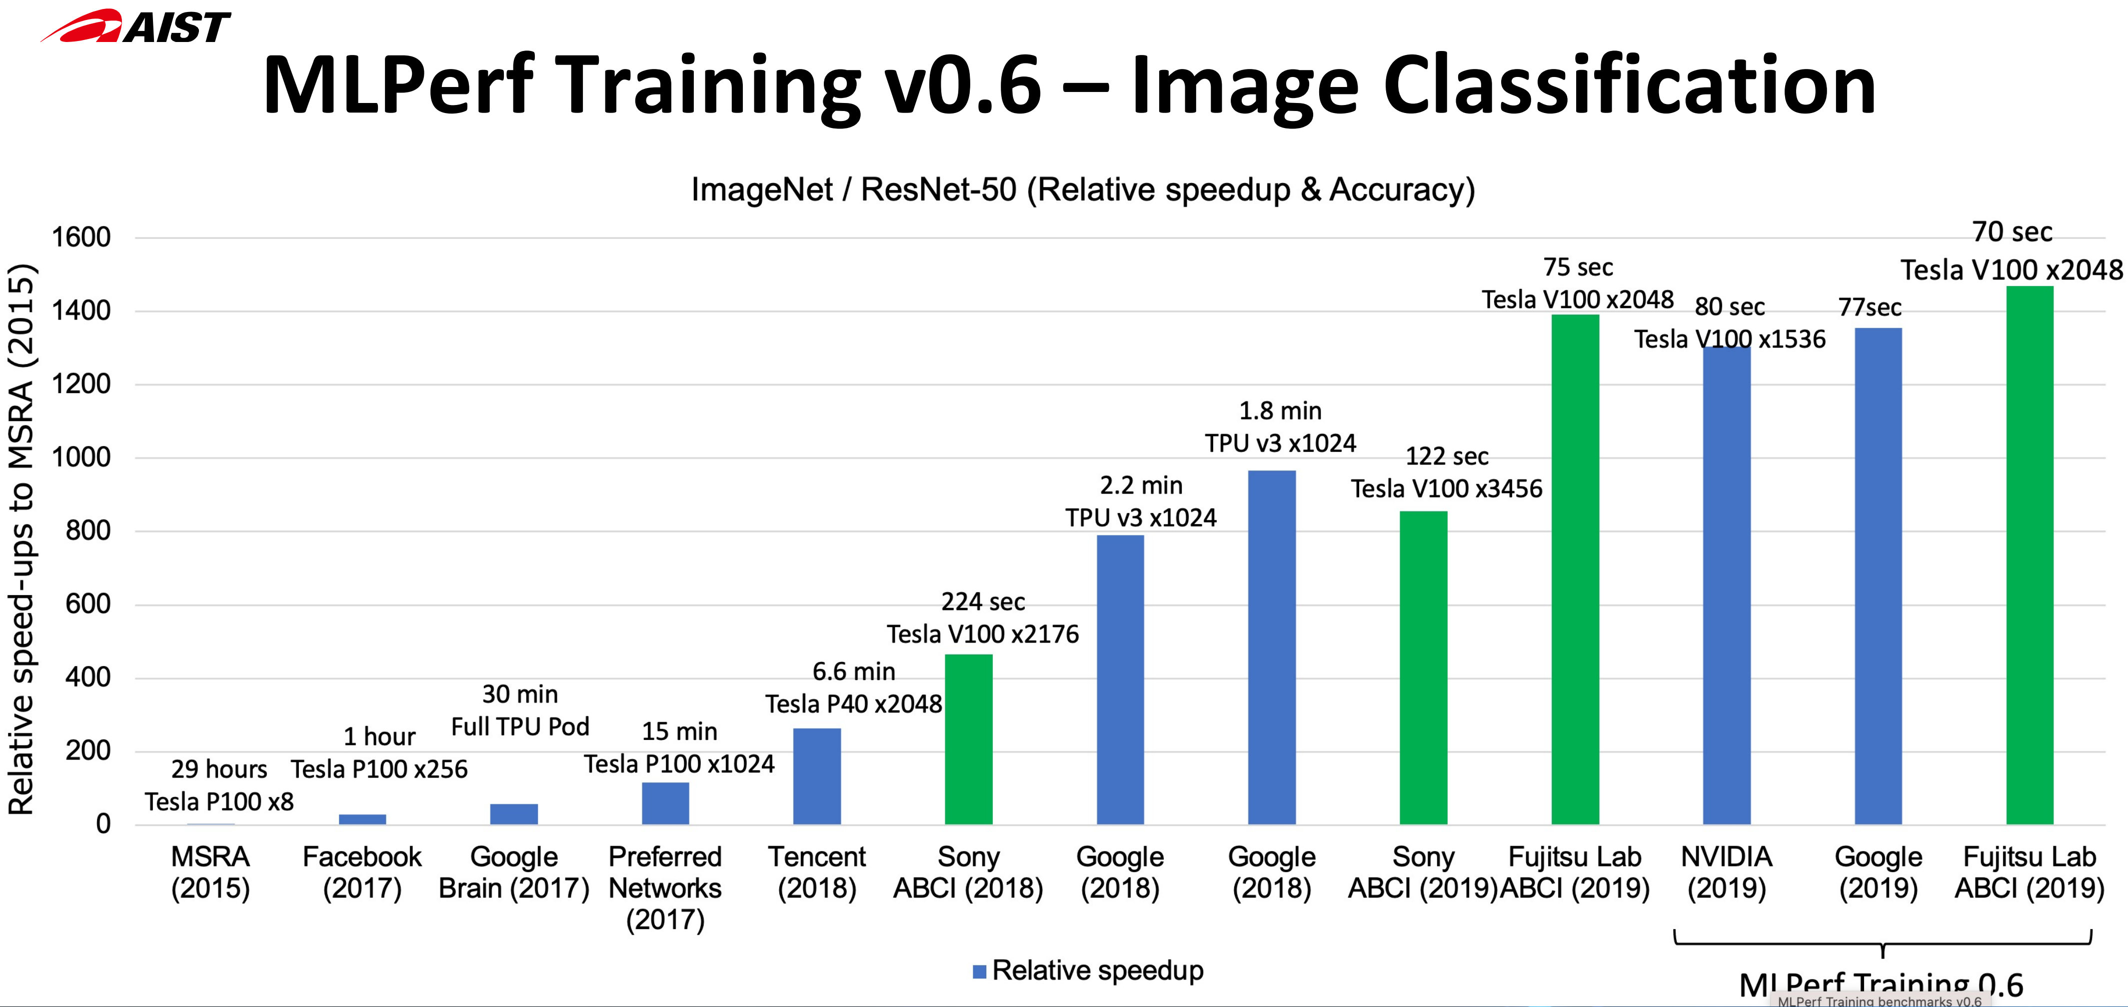


**Fujitsu achieved 70 seconds on ABCI, which was the world record at the time. ABCI is a powerful infrastructure to compete with hyper giants for AI R&D**

AIST

# MLPerf HPC Training v0.7 (Dec. 2020)

ABCI achieved 20x speedup for CosmoFlow and more than 12x speedup for DeepCam compared to other supercomputer systems.

CosmoFlow: Cosmology parameter prediction

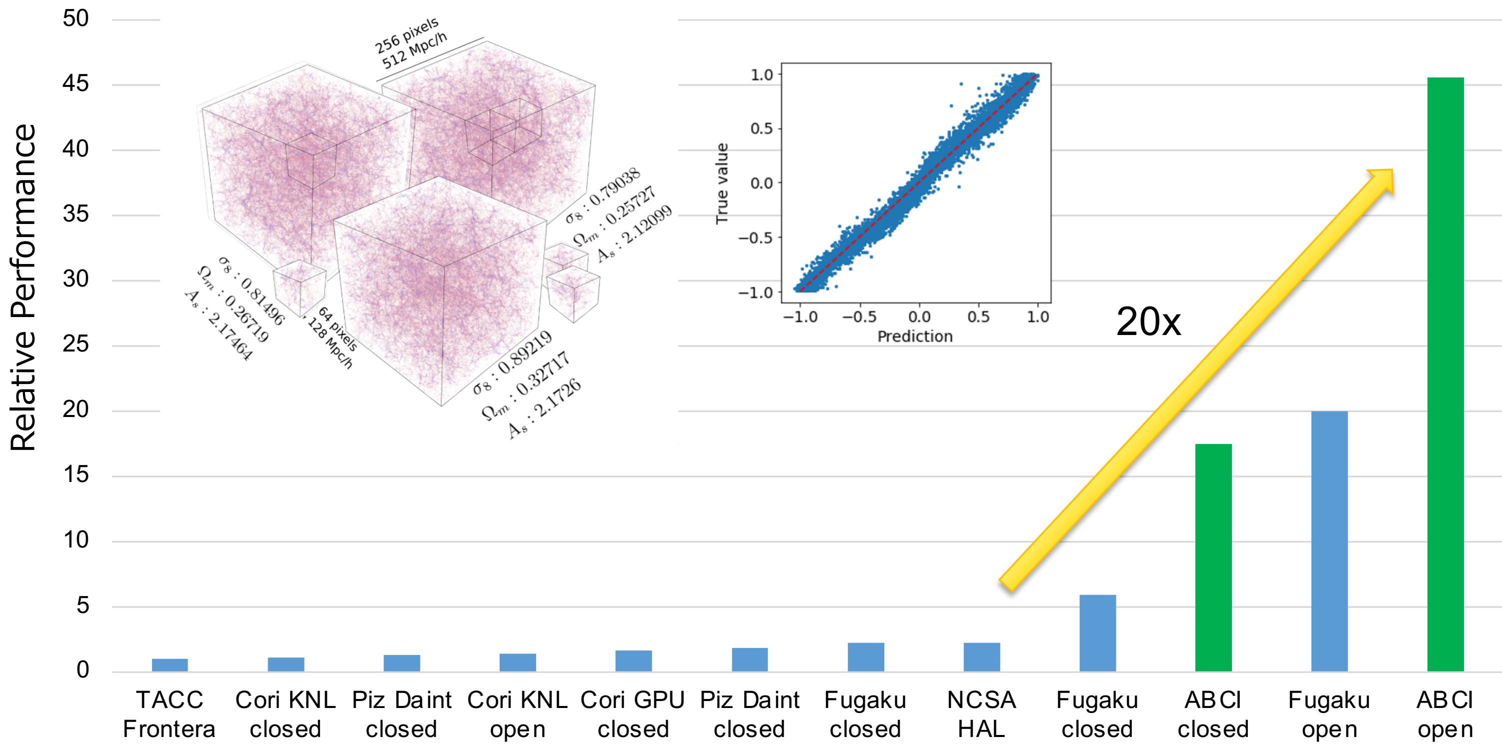


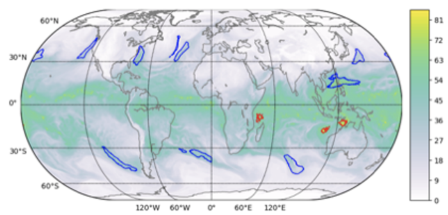

DeepCam: Climate segmentation

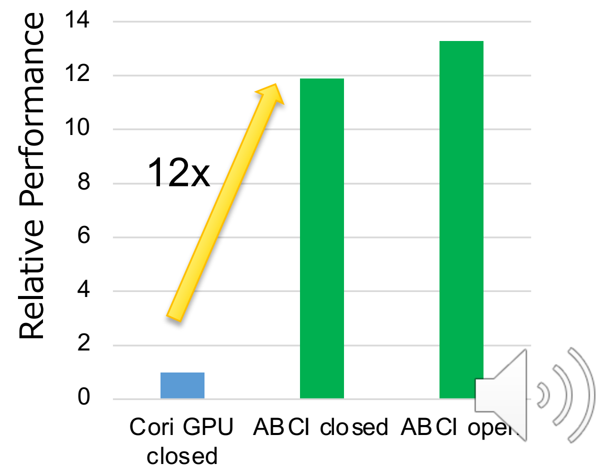

# High-resolution Image Reconstruction on ABCI

- Background
  - High-resolution Compute Tomography (CT) is widely used
    - Medical diagnosis, non-invasive inspection, reverse engineering
  - Computations are very intensive
    - Filtering computation, $O(Log(N)N^2)$
    - Back-projection, $O(N^4)$
  - High-resolution image is often required but not attainable
    - $(4K)^3$, $(6K)^3$, $(8K)^3$, etc.
- Proposed algorithm
  - Take advantage of the heterogeneity of GPU-accelerated supercomputer
    - GPUs are used for back-projection
    - CPUs are used for filtering computation
  - Design a parallel computing pipeline
  - Employ distributed system to tackle the out-of-core problem
  - Use advanced MPI for inter-node communication
- Impact to the real-world applications
  - Regardless of resolution of 3D images
  - Efficiently generate 3D images for many purposes
    - Training sample for Deep Learning
    - Improve image quality with iterative reconstruction

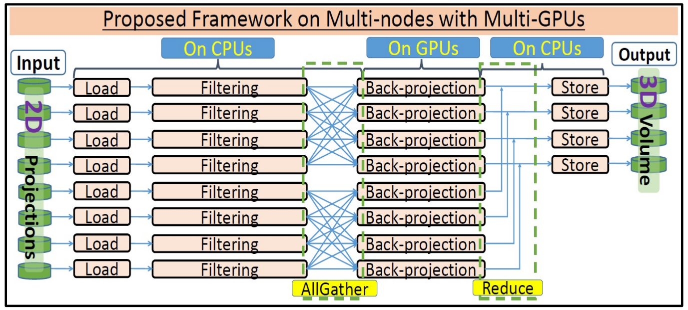


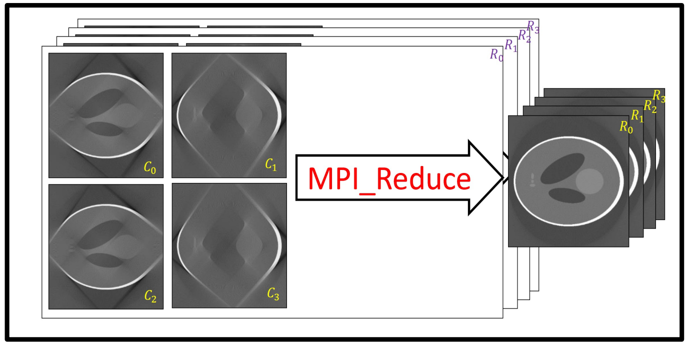


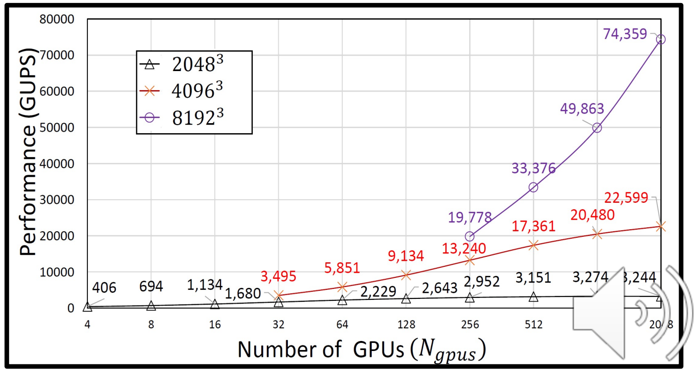


[1] P. Chen, et al., "iFDK: A Scalable Framework for Instant High-resolution Image Reconstruction", SC'19.

# Massively Parallel Causal Inference of Whole Brain Dynamics at Single Neuron Resolution

- Understand how the brain works by finding the causal relationships between neurons
- **Develop and design the EDM** framework to be **optimized for the GPU supercomputer**
- Using 512 ABCI nodes, mpEDM achieves 1,530x speedup compared to the original CPU-based implementation
- **Analyze the largest causal of whole brain dynamics at single neuron resolution to this date**

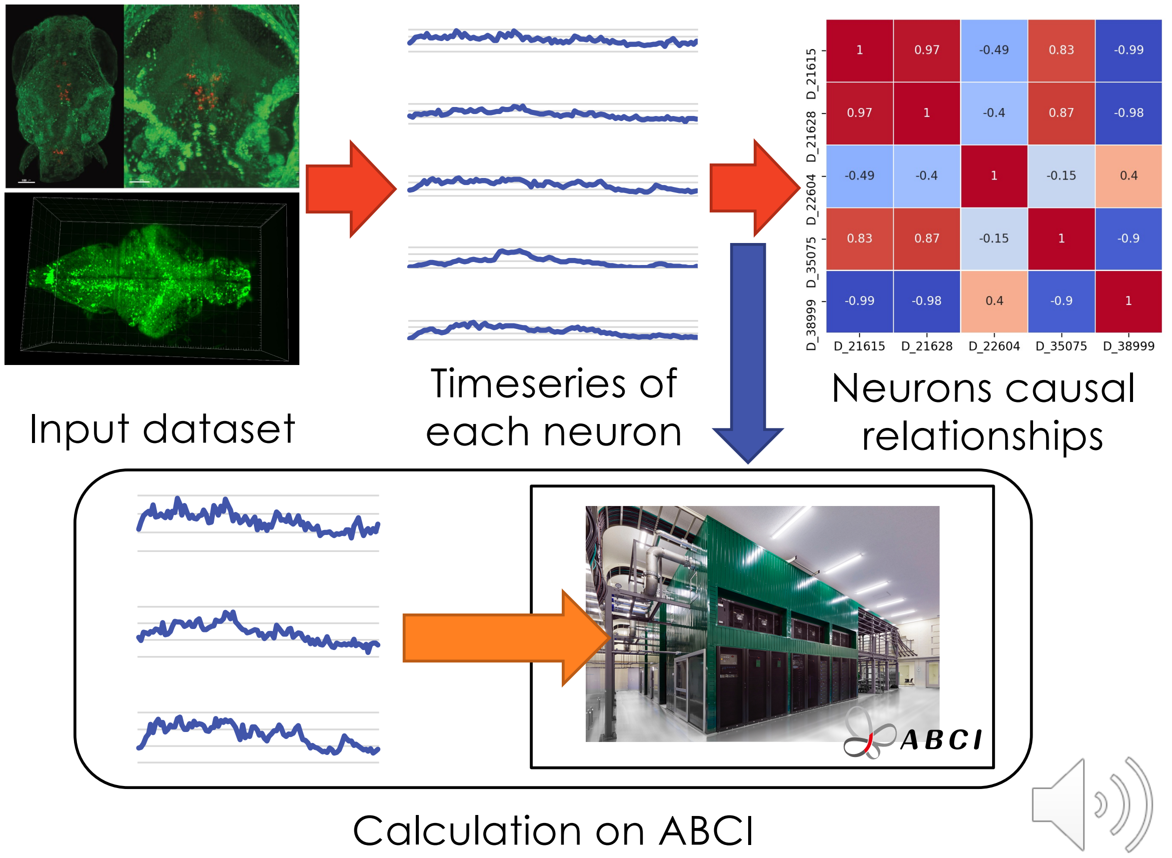



[1] **W. Watanakeesuntorn, et. al.,** Massively Parallel Causal Inference of Whole Brain Dynamics at Single Neuron Resolution, IEEE ICPADS 2020

AIST

# ABCI: Optimized for large-scale DL with commodity hardware

## High-Performance Computing System
**550 PFlops (FP16), 37.2 PFlops (FP64)**
**476 TiB memory, 1.74 PB NVMe SSD**

### Compute Node (w/ GPU) x 1,088

| | |
|---|---|
| GPU | NVIDIA Tesla V100 SXM2 x 4 (Total 4,352) |
| CPU | Intel Xeon Gold 6148 x 2 (Total 2,176) |
| Memory | 384GiB |
| Local Storage | Intel SSD DC P4600 (NVMe) 1.6TB x 1 |
| Interconnect | InfiniBand EDR x 2 |

**Multi Platform Node (w/o GPU) x 10**
- Intel Xeon Gold 6132 (2.6GHz/14cores) x 2
- 768GiB Memory, 3.8TB NVMe SSD

**Interactive Node x 4**

**Management and Gateway Nodes x 15**

## Large Volume Storage System
**22 PB GPFS + 1 PB All-Flash Storage +**
**17 PB Cloud Storage**

**DDN SFA14K (w/ SS8462 Enclosure x 10) x 3**
- 12TB 7.2Krpm NL-SAS HDD x 2400
- 3.84TB SAS SSD x 216
- NSD S

**DDN S**
- 7.68TI
- MDS S

**HPE A**
- 12TB SATA HDD x 1440
- 3.2TB SSD x 24

## Interconnect Network (Infiniband EDR)
- Mellanox CS7500 x 2
- Mellanox SB7890 x 229

## Service Network (10GbE)

## Ga
- Nexsus 3232C x2
- FortiGate 1500D x2
- FortiAnalyzer 400E x1

SINET5

### ■ Point 1: Large-scale and State-of-the-art GPU farm
Total 4352 GPUs. Users can run a large-scale AI/deep learning jobs with friction-less.

### ■ Point 2: Supercomputer-based network technology
All GPUs and CPUs can communicate with each other through ultra low latency (less than 2.5 us) and wide bandwidth (200Gbps) network.

# ABCI 2.0 Upgrade: Computing resources

## ABCI 2.0 (2021Q2-)

### ABCI 1.0 (2018Q2-)

550 PF (FP16), 37.2 PF (FP64)
476 TiB Memory, 1.74 PB NVMe SSD

**Compute Nodes (V) x 1088**

- GPU: NVIDIA Tesla V100 SXM2 x 4
- CPU: Intel Xeon Gold 6148 (2.4GHz/20cores) x 2
- Memory: 384 GiB
- Local Storage: Intel SSD DC P4600 (NVMe) 1.6TB x 1
- Interconnect: InfiniBand EDR x 2 (25 GB/sec)

+

### ABCI Expansion (2021Q2-)

300 PF (FP16), 19.3 PF (FP64)
97.5 TiB Memory, 384 TB NVMe SSD

**Compute Nodes (A) x 120**

- GPU: NVIDIA A100 x 8
- CPU: Intel Xeon SP (Ice Lake) x 2
- Memory: 512 GiB
- Local Storage: Intel SSD DC P4610 (NVMe) 1.6TB x 2
- Interconnect: InfiniBand HDR x 4 (100 GB/sec)

| Precision | ABCI 1.0 | Expansion | ABCI 2.0 | Scale-up |
|---|---|---|---|---|
| FP32/TF32 | 75 PF | 150 PF | 225 PF | x 3 |
| TF32 w/ Sparsity | ↑(*1) | 300 PF | 375 PF | x 5 |
| FP16/BF16 | 550 PF | 300 PF | 850 PF | x 1.55 |
| FP16/BF16 w/ Sparsity | ↑(*1) | 600 PF | 1.15 EF | x 2.09 |
| FP64 | 37.2 PF | 19.3 PF | 56.5 PF | x 1.52 |

(*1) V100 tensor core does not support Sparsity.

# Comparison of ABCI Compute Nodes

## Compute Node (V)

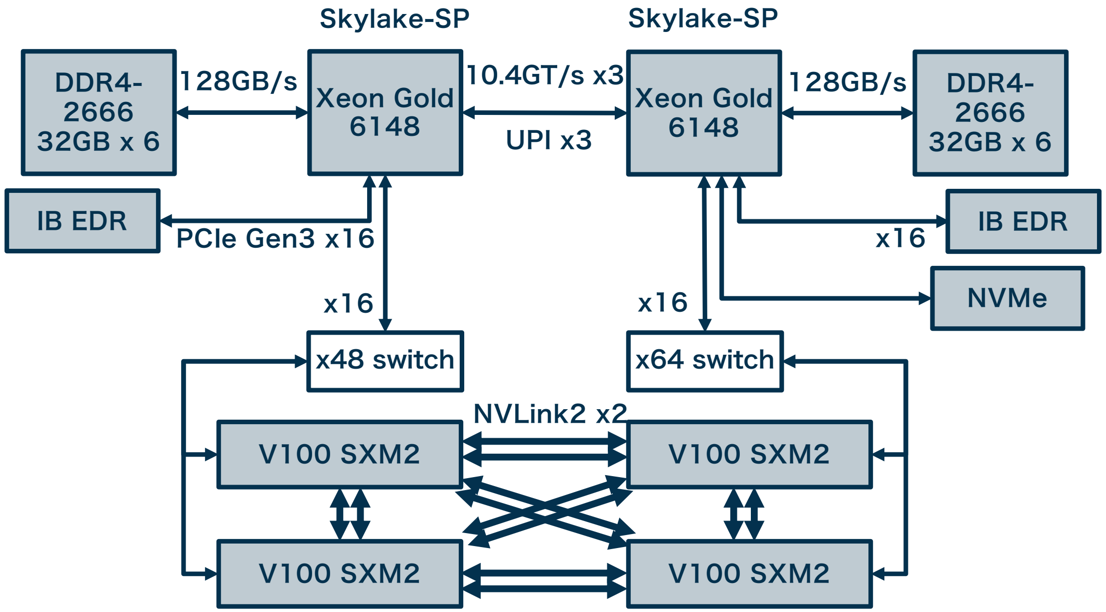


**CPU blade**

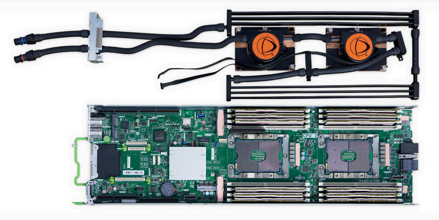

**GPU blade**

## Compute Node (A)

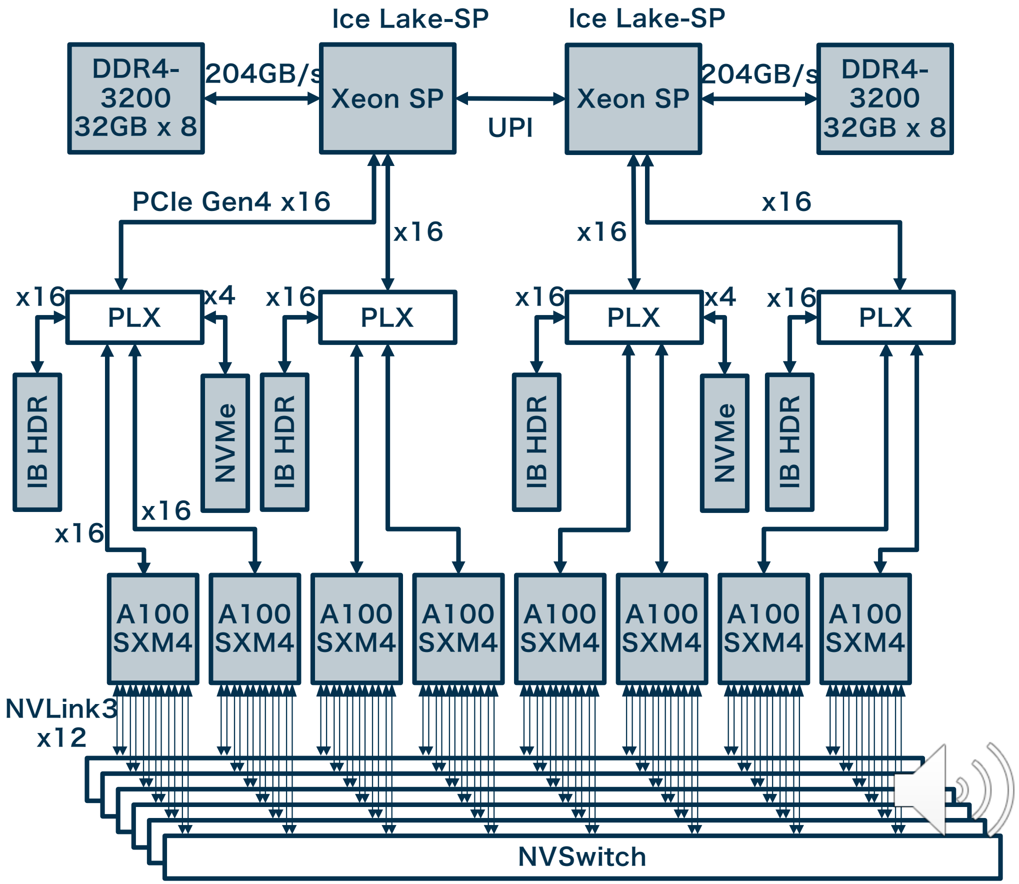

# ABCI Compute Network (A)

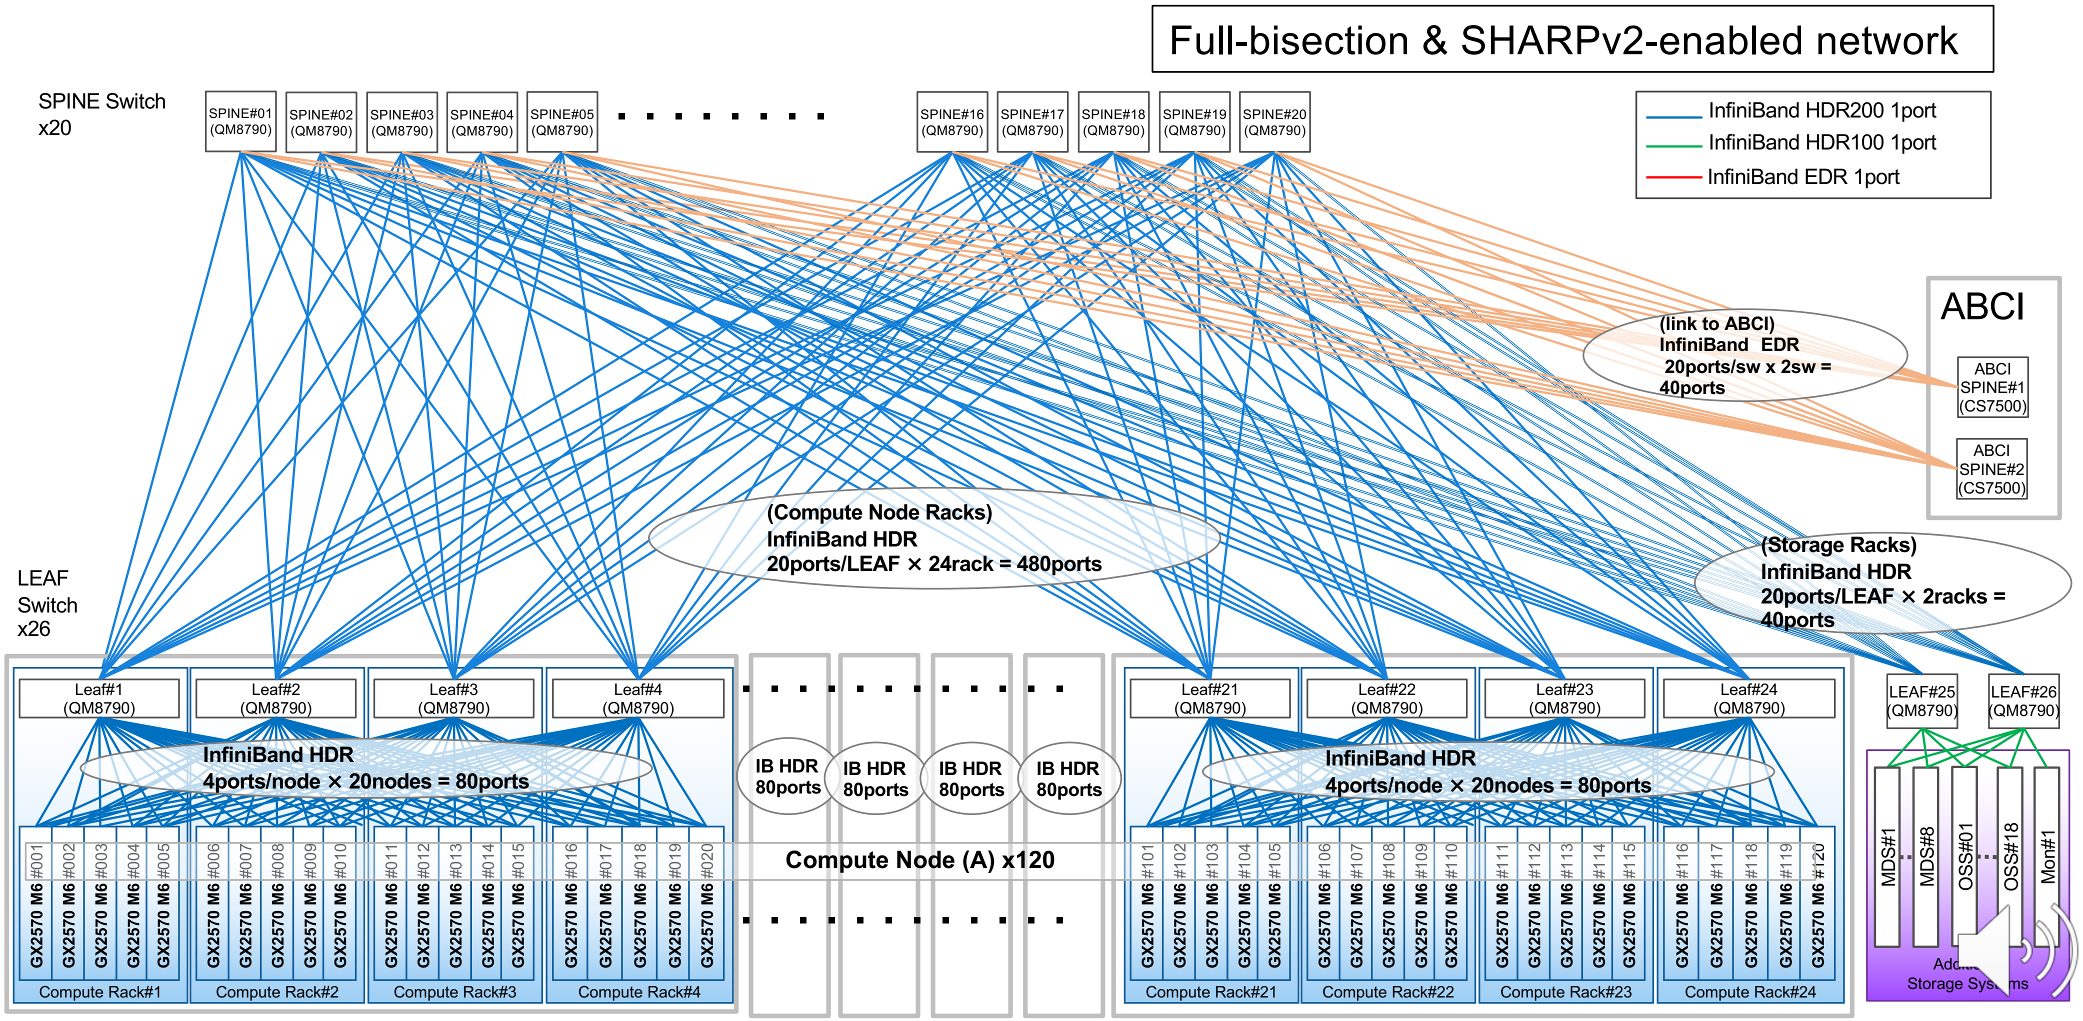

# ABCI 2.0 Upgrade: Storage IO

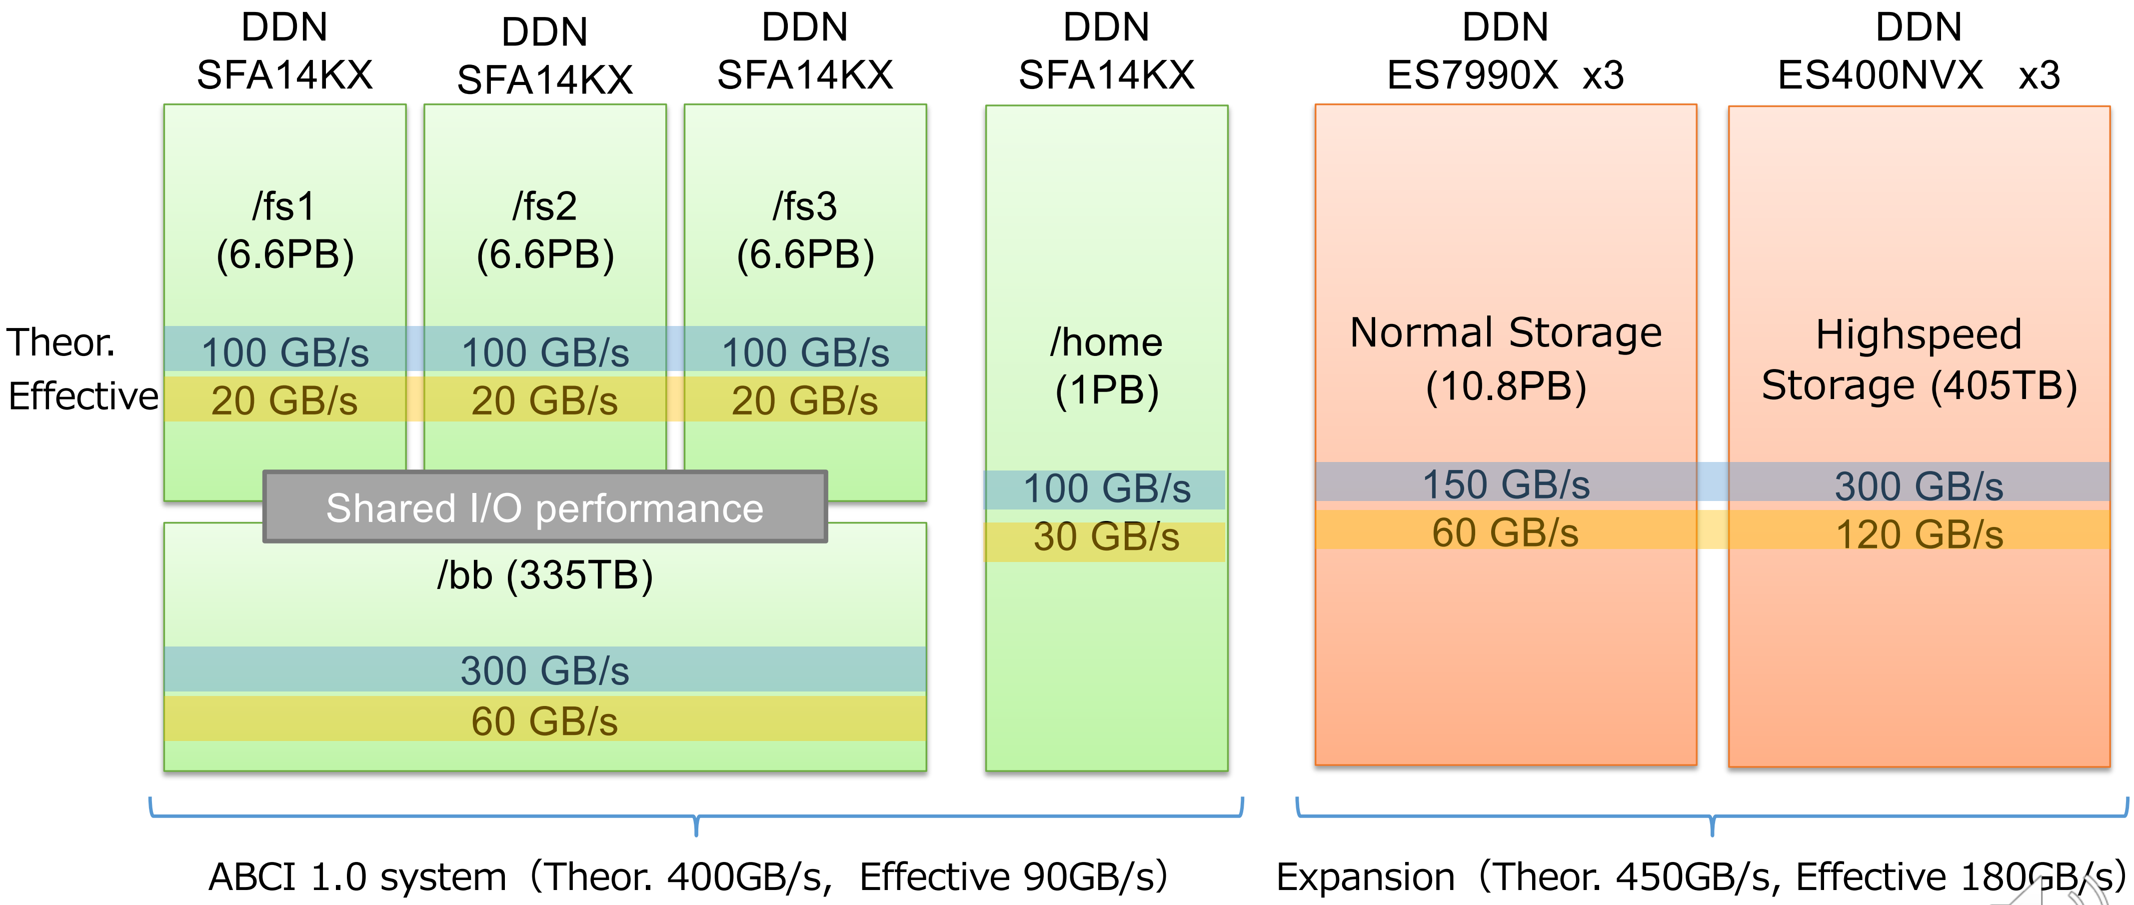

# ABCI Software Stack

- User-defined software stack
  - Container-based software deployment
  - Users can build their own environment with venv, pip, anaconda, pyenv, etc.
  - ISV applications can also be freely installed in the user environment.

- Basic software already installed and updated with short cycles:
  - CUDA SDK, cuDNN, NCCL, NVIDIA HPC SDK, Intel Parallel Studio XE Cluster Edition, PGI Professional Edition
  - Open MPI, Intel MPI, MVAPICH2

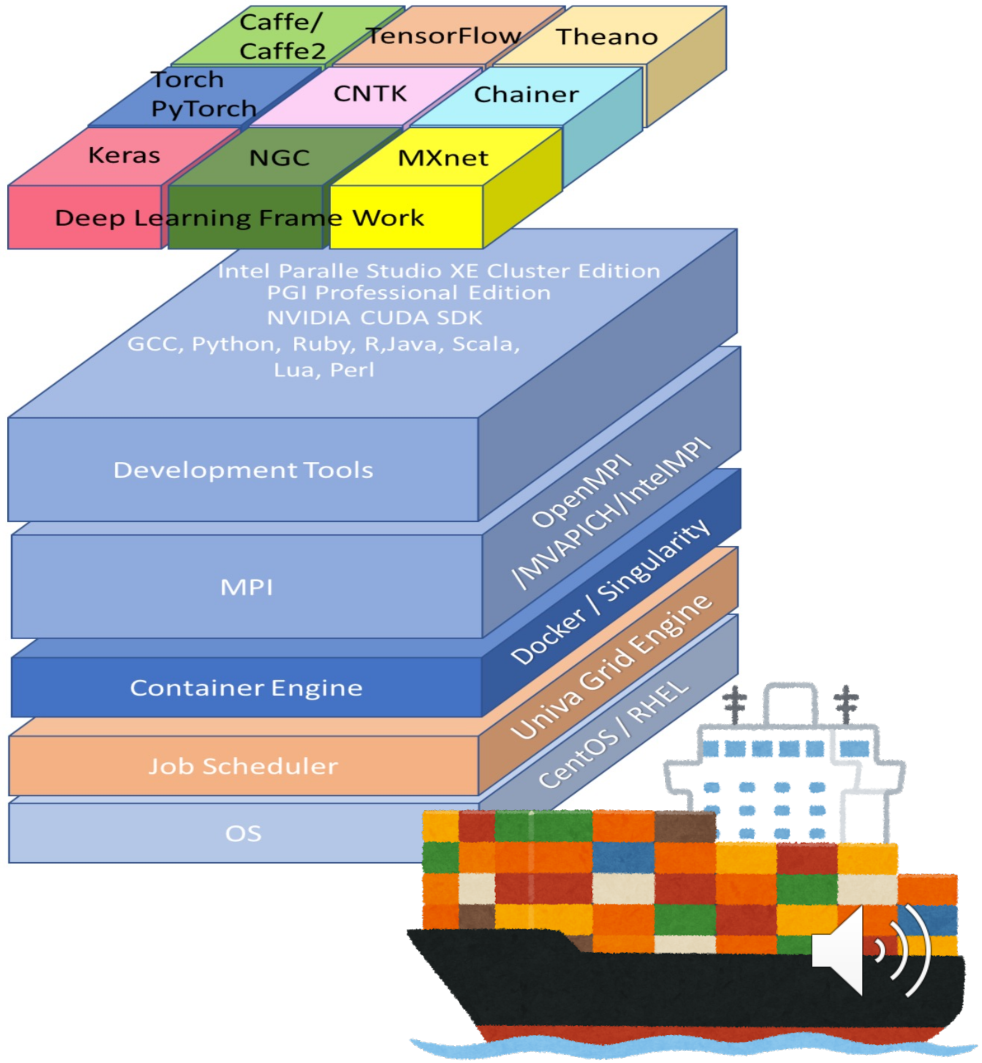

# Preliminary performance evaluation of ABCI 2.0

- Image classification training using ResNet-50/ILSVRC2012 to compare compute node (V) with node (A)
- A100 has almost double the performance of V100
- In terms of the scalability, node (A) has room to improve the performance

| Software | Version |
|---|---|
| Singularity Pro | 3.7 |
| CUDA | 11.2.2 |
| NCCL | 2.8.4-1 |
| OpenMPI | 4.0.5 |
| NVIDIA PyTorch Image | 21.04-py3 |
| Horovod | 0.22.0 |

Single node training

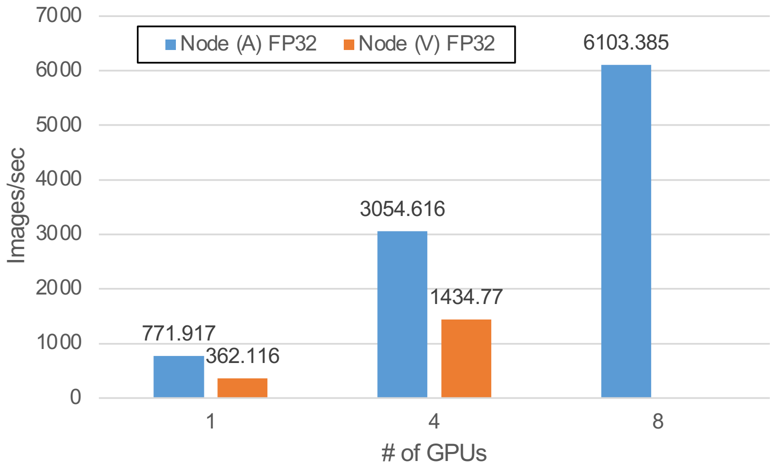


Multi node training

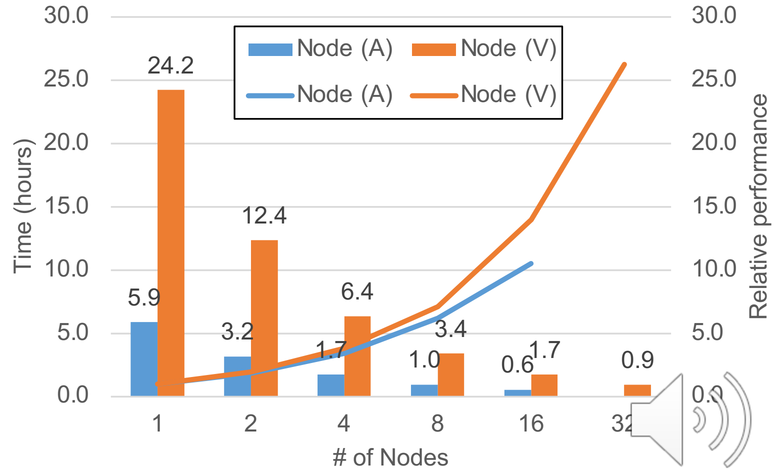

# ParaDL: Performance/Memory Estimating Model

- ✓ Formally defines possible pure parallelism strategies for Distributed Deep Learning:
  - Data, Spatial, Layer, Channel, Filter, Hybrid (combine of data with others)

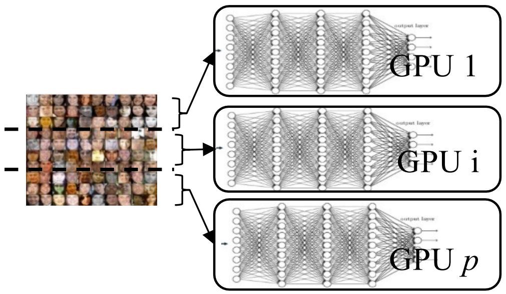


**Data parallelism**
split the dataset (input)

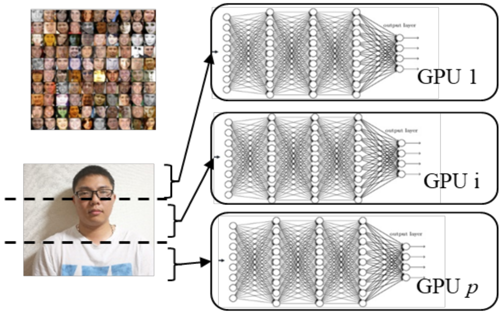


**Spatial parallelism**
split a sample (input)

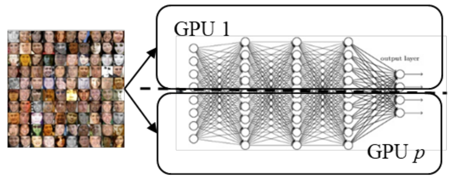


**Channel/Filter parallelism:**
vertically split the DNN model

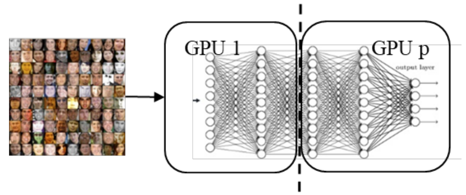


**Layer (Pipe) parallelism:**
horizontally split the DNN model

- ✓ Proposed an initial analysis/estimation model
  - Support many CNN, transformer (GEMM)
  - Based on an **Ideal Parameterization**

| | Computation Time $T_{comp}$ | Communication Time $T_{comm}$ |
|---|---|---|
| **Serial** | $D\sum_{l=1}^{G}\left(FW_l + BW_l\right) + \frac{D}{B}\sum_{l=1}^{G}(WU_l)$ | 0 |
| **Data** | $\frac{D}{p}\sum_{l=1}^{G}(FW_l + BW_l) + \frac{D}{B}\sum_{l=1}^{G}(WU_l)$ | $2\frac{D}{B}(p-1)\left(\alpha + \frac{\sum_{l=1}^{G}\lvert w_l\rvert}{p}\delta\beta\right)$ |
| **Spatial** | $\frac{D}{p}\sum_{l=1}^{G}\left(FW_l + BW_l\right) + \frac{D}{B}\sum_{l=1}^{G}(WU_l)$ | $2\frac{D}{B}\Big((p-1)(\alpha + \frac{\sum_{l=1}^{G}\lvert w_l\rvert}{p}\delta\beta) + \sum_{l=1}^{G}\left(2\alpha + B(\text{halo}(\lvert x_l\rvert) + \text{halo}(\lvert \frac{dL}{dy_l}\rvert))\delta\beta\right)\Big)$ |

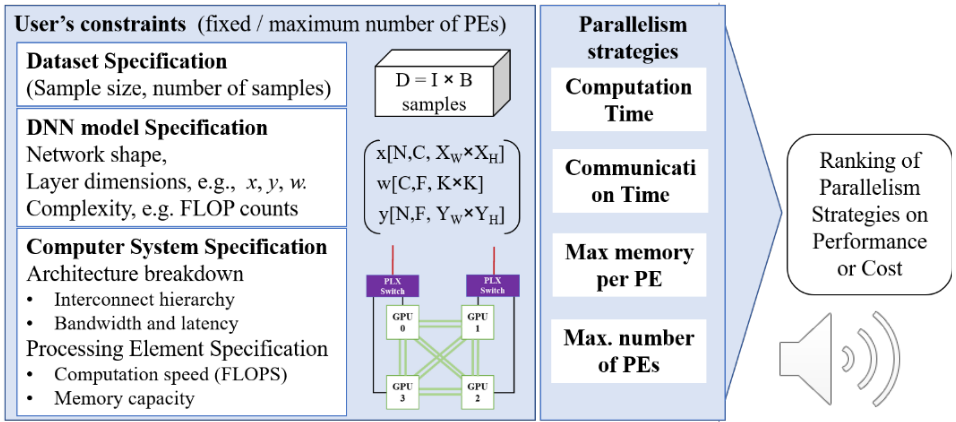

# ParaDL: Improving the Estimation Accuracy?

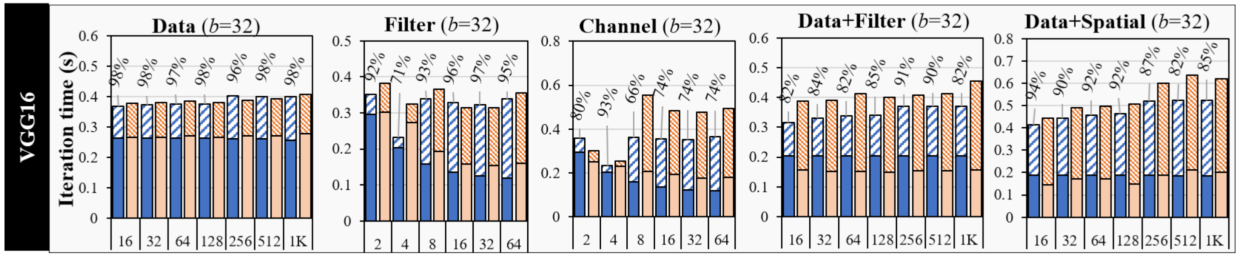


Time breakdown of our analytical model (ParaDL) in comparison with measured runs for different parallelism strategies of DL training. The x-axis is the number of GPUs. The label above each column shows the projection accuracy **(up to 97% correctness, and 86.7% on average)**

Different latency $\alpha$ and bandwidth factors $\beta$ → **Empirical Parameterization**

- Computation parameters ($FW_l$, $BW_l$, and $WU_l$)
- Communication parameter ($\alpha$ and $\beta$): use NCCL-test, OSU benchmarks

Network contention → **Self-contention modeling**

- Introduce contention penalty coefficient $\phi$,
- Using dynamic contention graphs [2] to detect contention

Network congestion→ **Detach it in the empirical result**

- Report minimum communication time
- Congestion impact factor using benchmark [3]

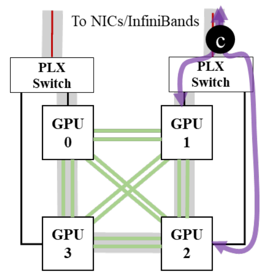


Self-contention in hybrid parallelism with $\phi = 2$


[2] Maxim Martinasso et al. 2011. A Contention-Aware Performance Model for HPC-Based Networks: A Case Study of the InfiniBand Network. In Euro-Par 2011 Parallel Processing.
[3] Sudheer Chunduri et al. 2019. GPCNeT: Designing a Benchmark Suite for Inducing and Measuring Contention in HPC Networks (SC '19)
[**HPDC**] **Truong Thao Nguyen**, et al. "An Oracle for Characterizing and Guiding Large-Scale Training of Deep Neural Networks," ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC2021) (**to be appeared**)

# An Efficient 2D Convolution Algorithm on GPUs

- **Background**
  - GPU is one of the most important accelerators
  - 2D convolution is a fundamental operation
    - Signal/Image processing
    - Deep learning
- **Proposed algorithms**
  - Use register files as a user-managed cache
  - Employ shuffle intrinsic for efficient intra-warp communication
- **Impact to real-world applications**
  - Accelerate Image processing applications
  - Speed up Deep Learning workload

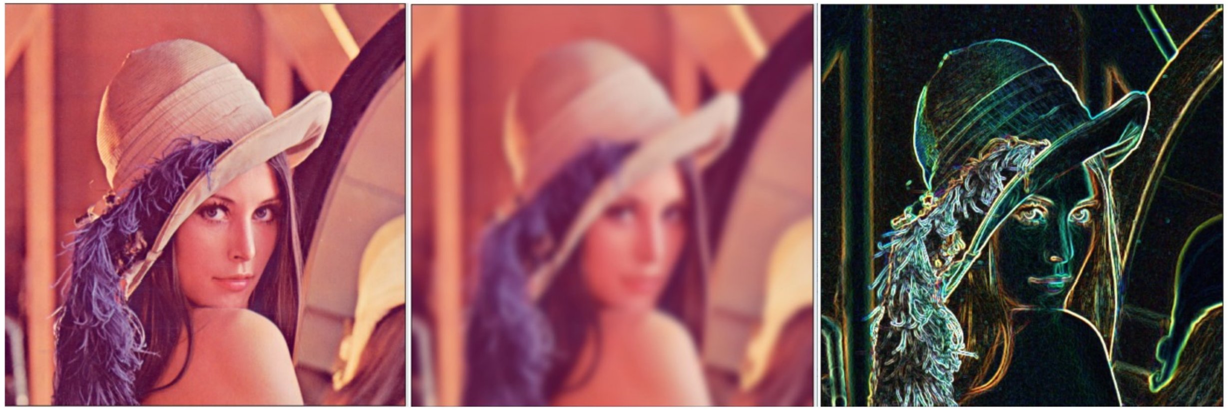

(a) Original (b) Gaussian (c) Edge detection

[4] P. Chen, et al., "A Versatile Software Systolic Execution Model for GPU Memory-Bound Kernels“, SC’19. (Nominated as Best paper and Best Student paper finalist)

# Implementation of the proposed algorithms

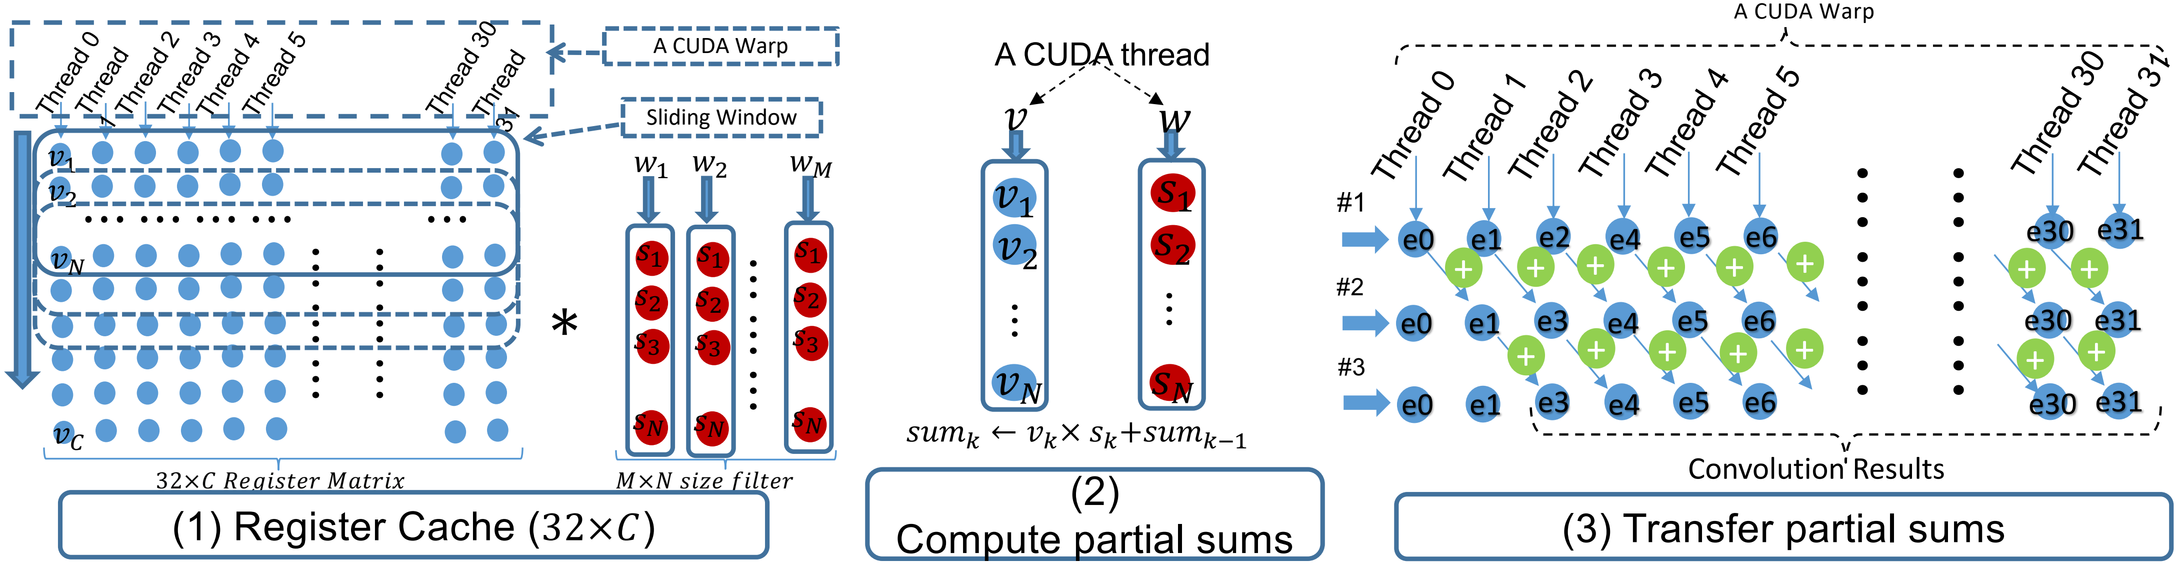


- **Evaluation on Tesla P100/V100 GPUs**

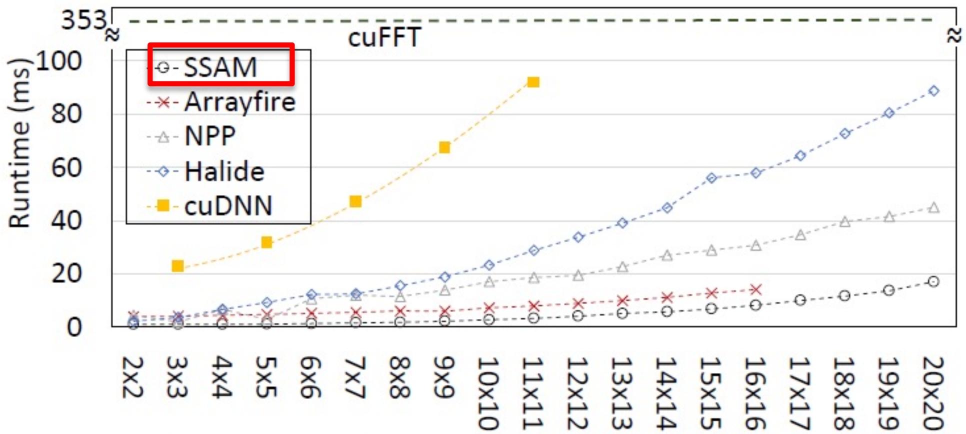


(a) Single precision on P100. cuFFT method is constant as 353 ms.

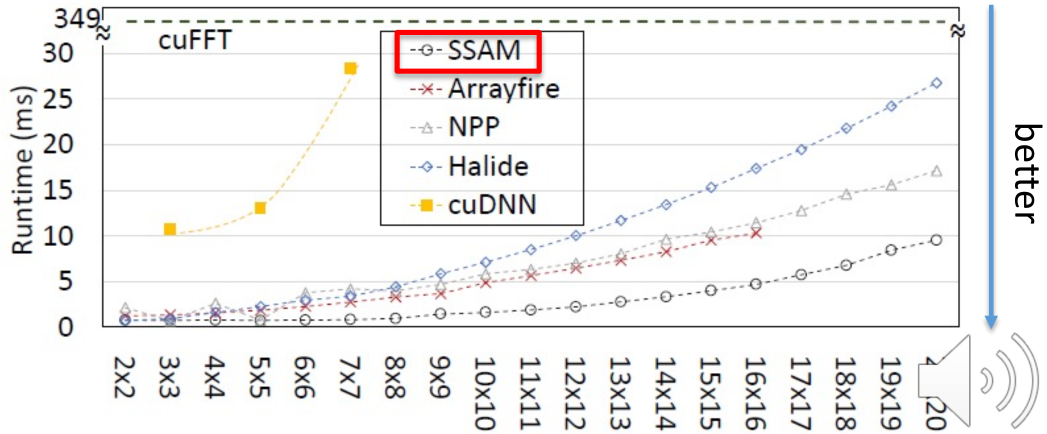


(b) Single precision on V100. cuFFT method is constant as 349 ms.

# AI can have a substantial energy consumption

- Opportunities to improve energy efficiency and $CO_2$ emission
  - Algorithm/program improvement
    - Green AI algorithms
  - Processor improvement
    - Domain-specific accelerators
  - Datacenter improvement
    - Better PUE
  - Energy mix improvement
    - Clean energy purchase

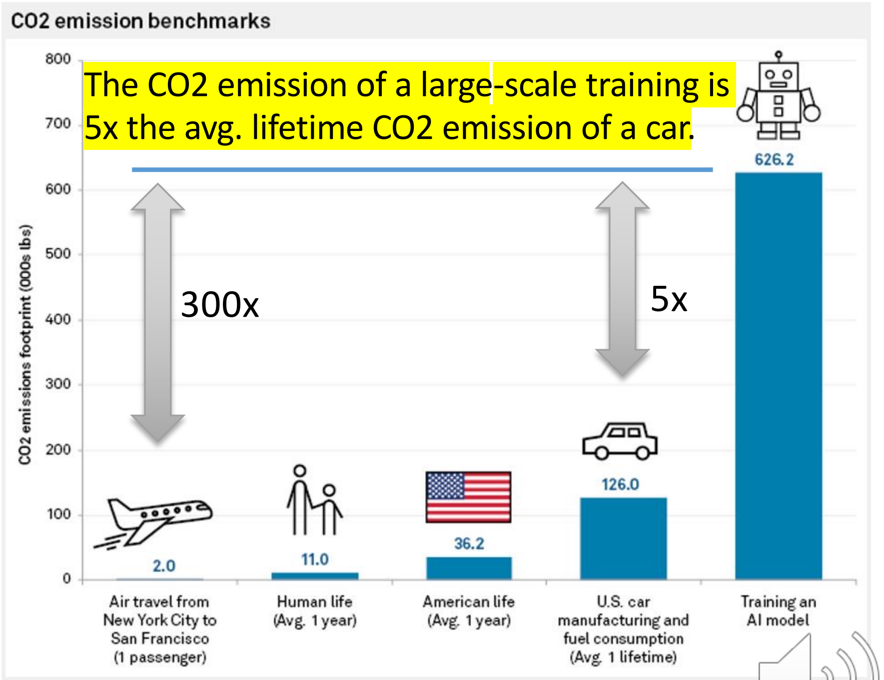


D. Patterson, et al., "Carbon Emissions and Large Neural Network Training," https://arxiv.org/ftp/arxiv/papers/2104/2104.10350.pdf

G. Gow, "Environmental Sustainability and AI," Forbes, Aug. 2020. https://www.forbes.com/sites/glenngow/2020/08/21/environmental-sustainability-and-ai/

# ABCI Ultra Green Datacenter

**"Commoditizing supercomputer cooling technologies to Cloud"**

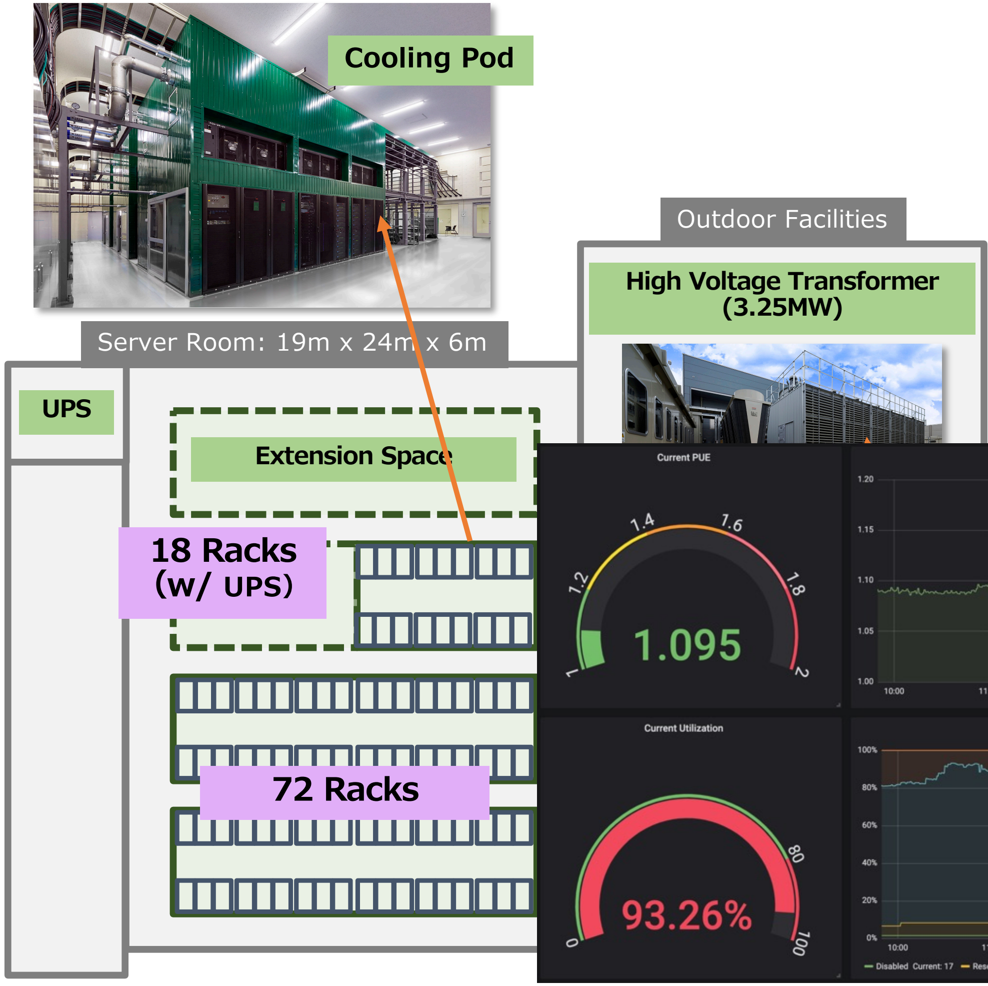


- **Single floor, cost effective building**
- **Hard concrete floor 2t/m² weight tolerance for racks and cooling pods**
- **Number of Racks**
  - Initial: 90 (ABCI uses 41 racks)
  - Max: 144
- **Power capacity: 3.25 MW**
  - MW max
- **y: 3.2MW**
  - kW water＋10kW air
- **fectiveness: < 1.1**

PUE) = $\frac{\text{Total Facility Energy}}{\text{IT Equipment Energy}}$

dard, 1.4 Good, 1.0 Best)

## Summary

- **AI is now a Supercomputing problem.**
- **ABCI is an open and large-scale infrastructure** to push the envelope of AI research and development.
- There are many opportunities and challenges. E.g.,
  - Pursuing an optimized parallelism strategy (data/model/hybrid parallelism)
  - Improving energy efficiency and $CO_2$ emission from algorithms to datacenter infrastructure

AIST

ABCI

AI Bridging Cloud Infrastructure

https://abci.ai/