AI-for-Scienceを加速しその先へと導く初のエクサスケールスーパーコンピューター

(The First Exascale Supercomputer Accelerating AI-for-Science and Beyond)

モハマド ワヒブ (Mohamed WAHIB)
理研 計算科学研究センター (RIKEN-CCS)
高性能人工知能システム研究チームリーダー

27-10-2022

HPCIコンソーシアム

Prelude

- **AI/ML is not magic**
  - **Important to distinguish hype from real prospect**
- **There is prospect for AI/ML in HPC**

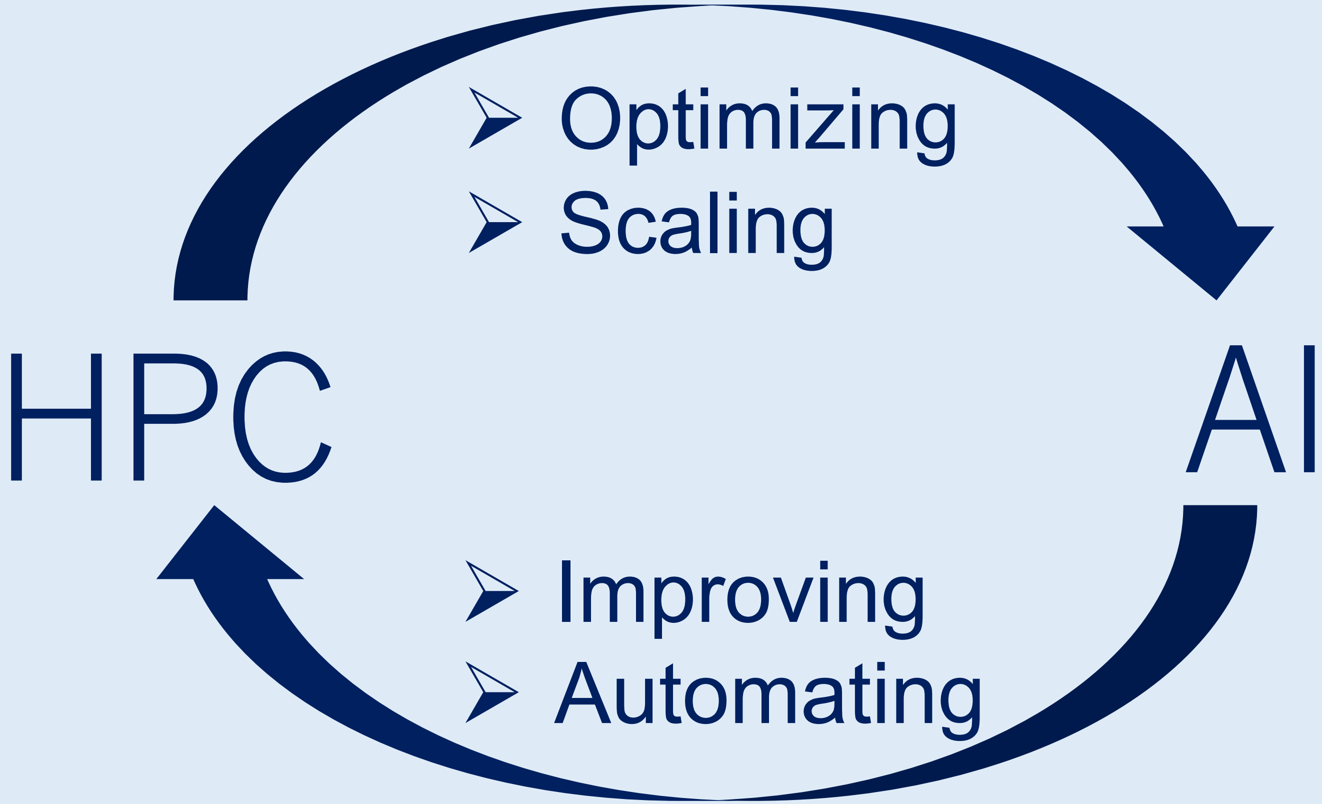

Fugaku Supercomputer

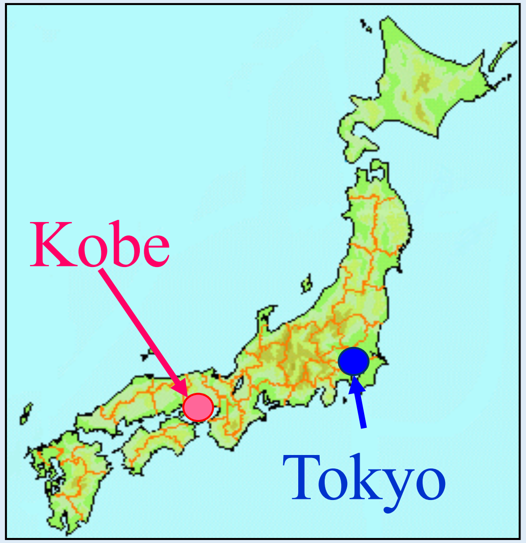

423 km (263 miles)
west of Tokyo

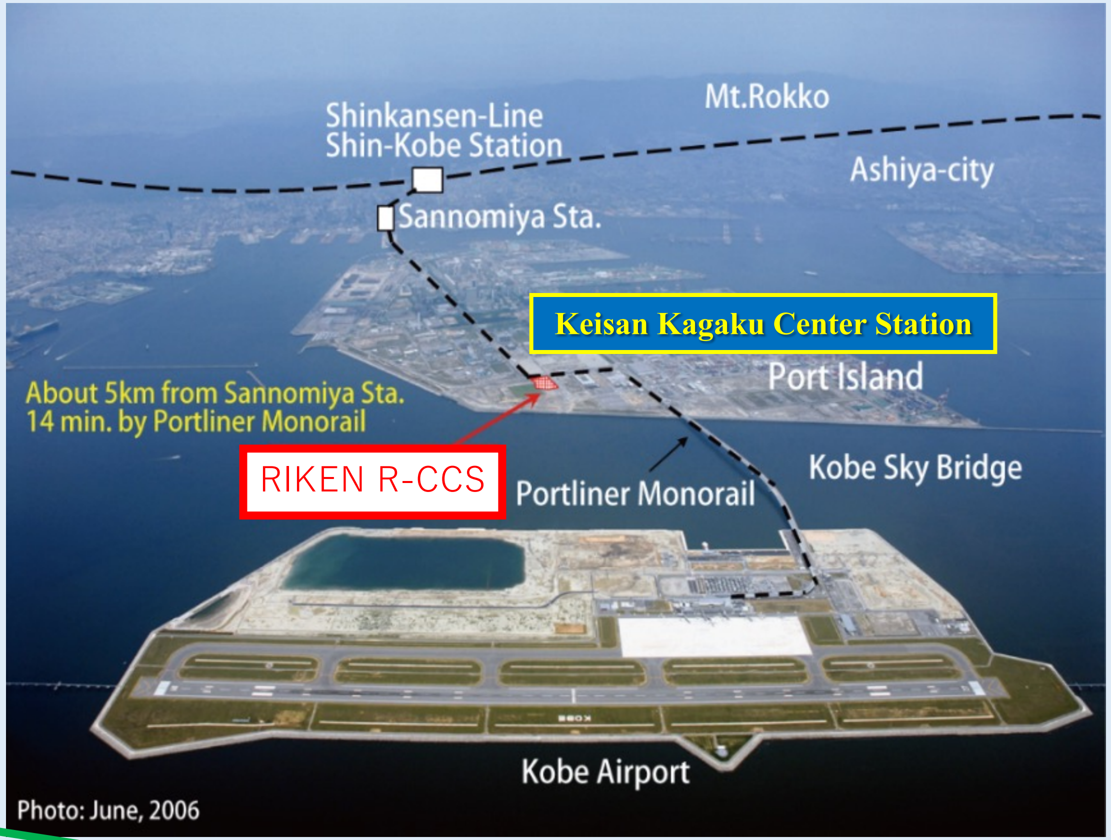

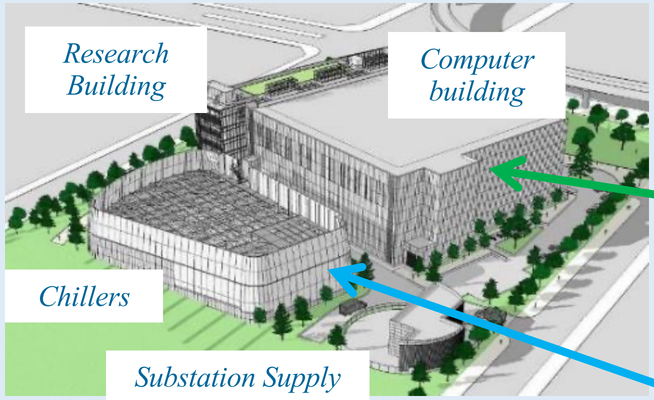

Computer room 50 m x 60 m = 3,000 m^2
Electric power up to 37 MW
Water cooling system

Gas-turbine co-generation 5 MW x 2

Fugaku Supercomputer

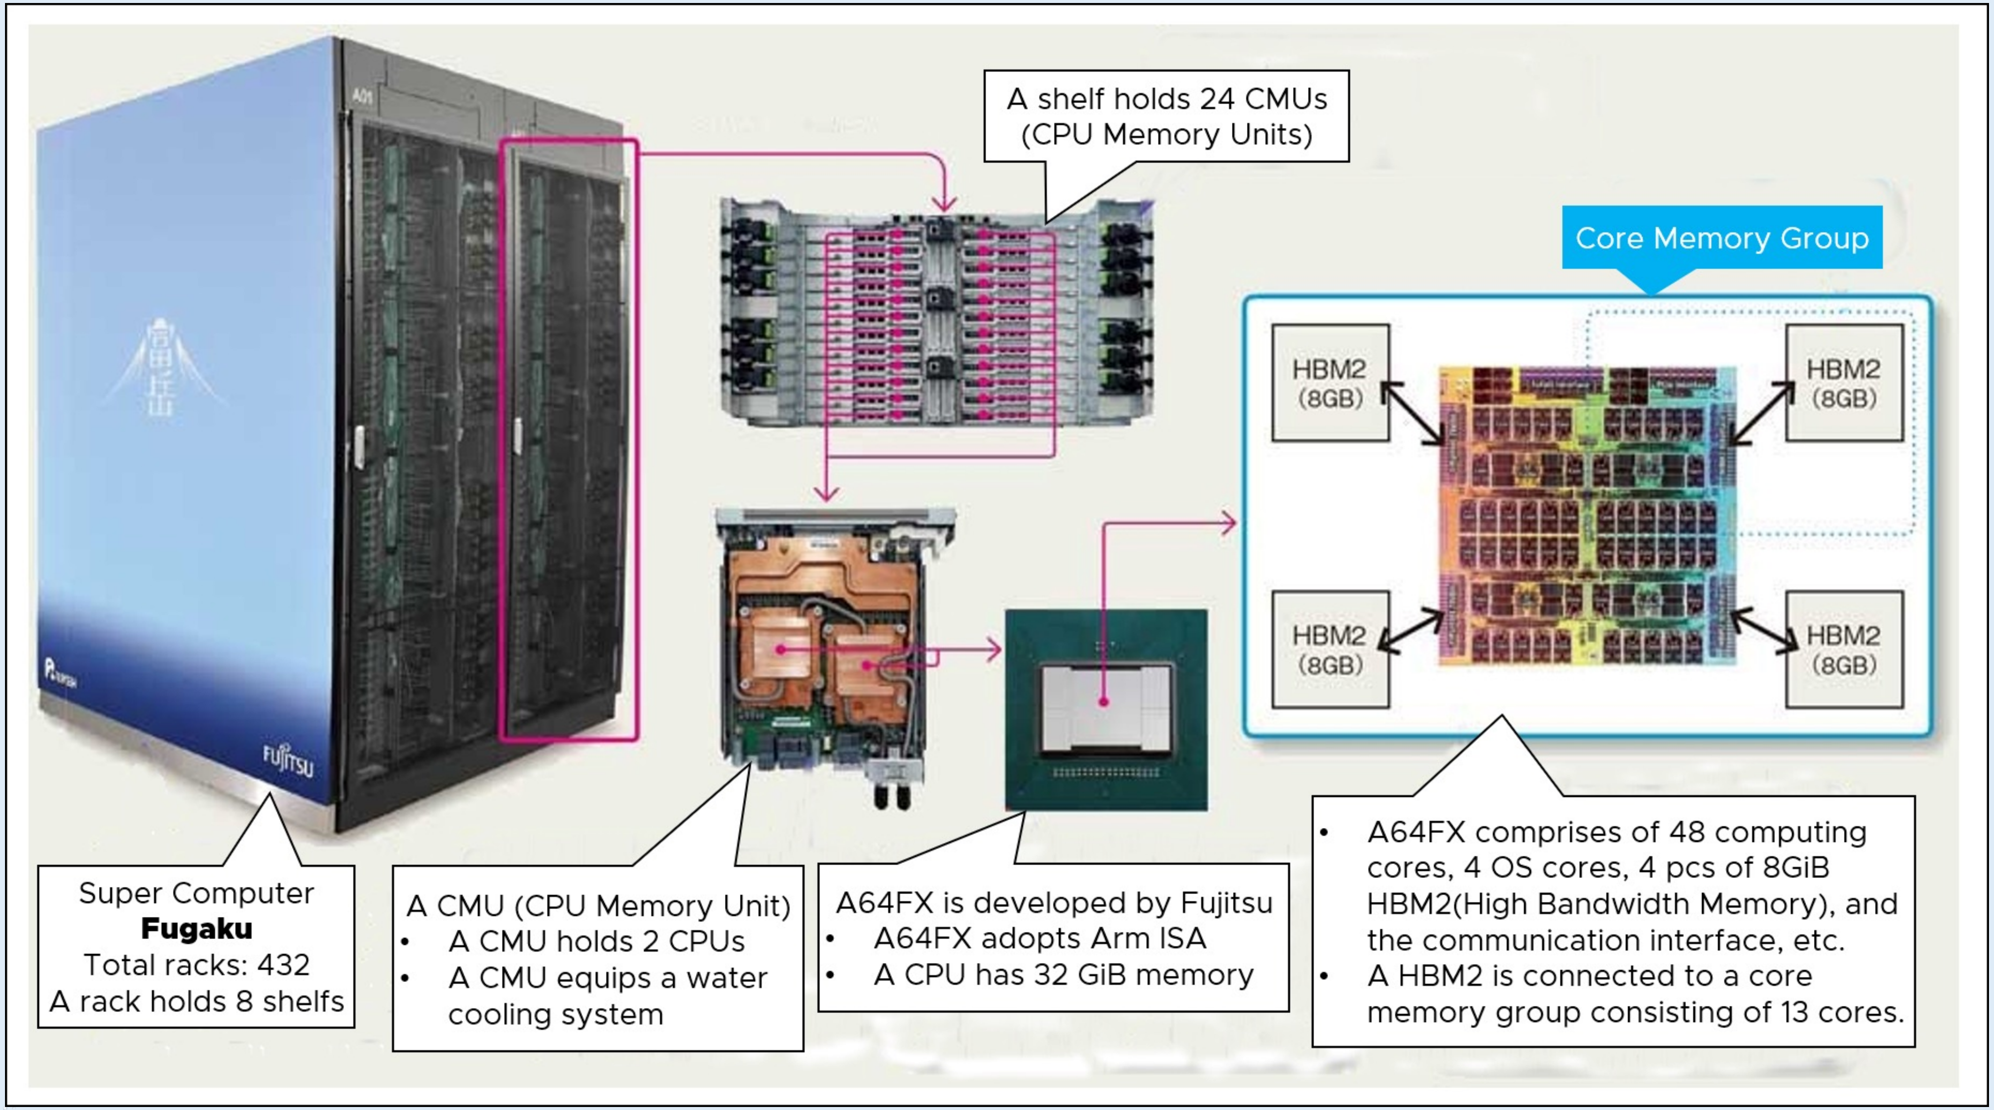

Fugaku Research Highlights (FY21)

Achieved world's first five titles

(4 consecutive terms + new ML Perf HPC 1st place)

In four HPC performance rankings (Top500, HPCG, HPL-AI, Graph500), Fugaku won four titles consecutively from June 2020. In November 2021, also awarded first place in ML Perf HPC, a new overall performance evaluation of AI processing.

Fugaku's high overall performance in a wide range of fields, as well as its ability to make a significant contribution to the realization of Society 5.0/SDGs.

Gordon Bell Special Prize

Fight against COVID-19

Successfully developed a detailed and quantitative COVID-19 droplet and aerosol dispersion model using "Fugaku" for the first time in the digital transformation of infectious disease epidemiology. **Visualizing arised awareness of the importance of understanding droplet and aerosol infection changing behaviour not only in Japan, but also around the world.**

ITU-AJ Special Achievement Award

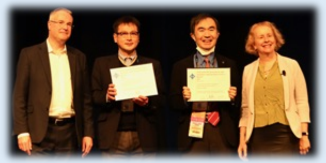

Data Assimilation Research

(Prediction of Sudden downpours, COVID-19 infection)

Using big data from weather radar, **a real-time of ultra-fast precipitation forecasting, was conducted in the Tokyo area using Fugaku during the Tokyo Olympic and Paralympic Games.**

Data assimilation methods developed in numerical weather forecasting were **applied to the forecasting of COVID-19 infections**

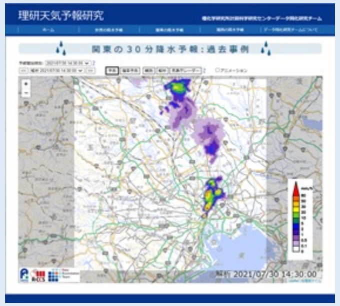

“GENESIS” new version released

Molecular dynamics (MD) simulations

A new version of GENESIS, optimized for 'Fugaku' by co-design, **more than 125 times faster** and with many new features, has been released as free software in 2020. **Work on the dynamic**

structure of spike proteins on the surface of COVID-19 has analized successfully. RIKEN EIHO Award (RIKEN Significant Achievement Award)

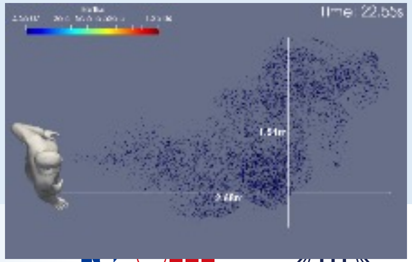

AI on Fugaku

- **World 2nd fastest Supercomputer (Top500: June 2022)**
  - 1st on HPCG and Graph500
- **Fujitsu A64FX Arm v8.2-A + SVE CPUs (32GB HBM memory)**
- **Creating DL ecosystem for Fugaku: https://github.com/dl4fugaku**

Scaling AI/ML on Fugaku

- **Past and on-going effort on scaling ML on Fugaku goes into benchmarks**
- **MLPerf (https://mlcommons.org/en/)**
  - **Subcategory: MLPerf Training HPC**

Area	Benchmark	Dataset	Quality Target	Reference Implementation Model
Scientific	Climate segmentation	CAM5+TECA simulation	IOU 0.82	DeepCAM
Scientific	Cosmology parameter prediction	CosmoFlow N-body simulation	Mean average error 0.124	CosmoFlow
Scientific	Quantum molecular modeling	Open Catalyst 2020 (OC20)	Forces mean absolute error 0.036	DimeNet++

Scaling AI/ML on Fugaku

PERFORMANCE METRICS (TIME TO SOLUTION IN MINUTES) FROM SUBMISSIONS IN CLOSED AND OPEN DIVISIONS

Division	System	Submission	Software	#Processors	#Accelerators	Parallelism†	CosmoFlow	DeepCAM
Closed	Piz Daint	Piz-Daint-128	TensorFlow 2.2.0	128	128	2 s/1 GPU	461.01	-
	Piz Daint	Piz-Daint-256	TensorFlow 2.2.0	256	256	2 s/1 GPU	327.01	-
	ABCI	ABCI-1024	PyTorch 1.6.0	512	1,024	2 s/1 GPU	-	11.71
	ABCI	ABCI-512	TensorFlow 2.2.0	256	512	1 s/1 GPU	34.42	-
	Fugaku	Fugaku-512	TensorFlow 2.2.0 + Mesh TensorFlow	512	-	1 s/1 CPU	268.77	-
	Fugaku	Fugaku-8192	TensorFlow 2.2.0 + Mesh TensorFlow	8,192	-	1 s/16 CPUs	101.49	-
	Cori-GPU	Cori-GPU-64	PyTorch 1.6.0	16	64	2 s/1 GPU	-	139.29
	Cori-GPU	Cori-GPU-64	TensorFlow 1.15.0	16	64	1 s/1 GPU	364.73	-
	Cori-KNL	Cori-KNL-512	TensorFlow 1.15.2	512	-	1 s/1 CPU	536.06	-
	HAL	HAL-64	TensorFlow 1.15.0	32	64	1 s/1 GPU	265.59	-
	Frontera-RTX	Frontera-RTX-64	TensorFlow 1.15.2	32	64	1 s/1 GPU	602.23	-
Open	ABCI	⋆ABCI-1024	PyTorch 1.6.0	512	1,024	2 s/1 GPU	-	10.49
	ABCI	⋆ABCI-2048	TensorFlow 2.2.0	1,024	2,048	1 s/1 GPU	13.21	-
	Fugaku	⋆Fugaku-16384	TensorFlow 2.2.0 + Mesh TensorFlow	16,384	-	1 s/4 CPUs	30.07	-
	Cori-KNL	⋆Cori-KNL-1024	TensorFlow 1.15.2	1,024	-	1 s/1 CPU	419.69	-

Farrell et al, "MLPerf HPC: A Holistic Benchmark Suite for Scientific Machine Learning on HPC Systems", WS col-located with SC'21

AI-for-Science on Fugaku*

by R-CCS Research Teams & Priority application programs (collected by Kento Sato)

* To Appear: Book Chapter "The First Exascale Supercomputer Accelerating AI-for- Science and Beyond "
Satoshi Matsuoka, Mohamed Wahib, Aleksandr Drozd, Kento Sato

Nonlinear Dimensionality Reduction for Three-dimensional Flow Field Around Circular Cylinder with Distributed Parallel Machine Learning on Fugaku (Kazuto Ando et al. [1])

- Mode decomposition is an approach used to decompose flow fields into physically important flow structures known as modes, POD (Proper Orthogonal Decomposition) is a powerful method
- MD-CNN-AE [1] extended to support **3-D flow**
  - MD-CNN-AE [2]: A neural network for a mode decomposition method for **2-D flow** field around circular cylinder using a convolutional neural network
- A hybrid parallelism method combining the distribution of network structure (model parallelism) and the training data (data parallelism) using up to **10,500 nodes on Fugaku was employed for learning**.

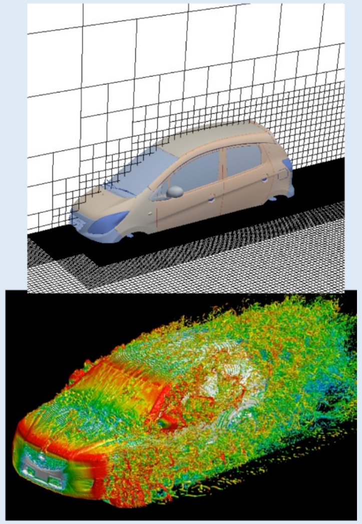

Schematics of 3D-extended version of MD-CNN-AE

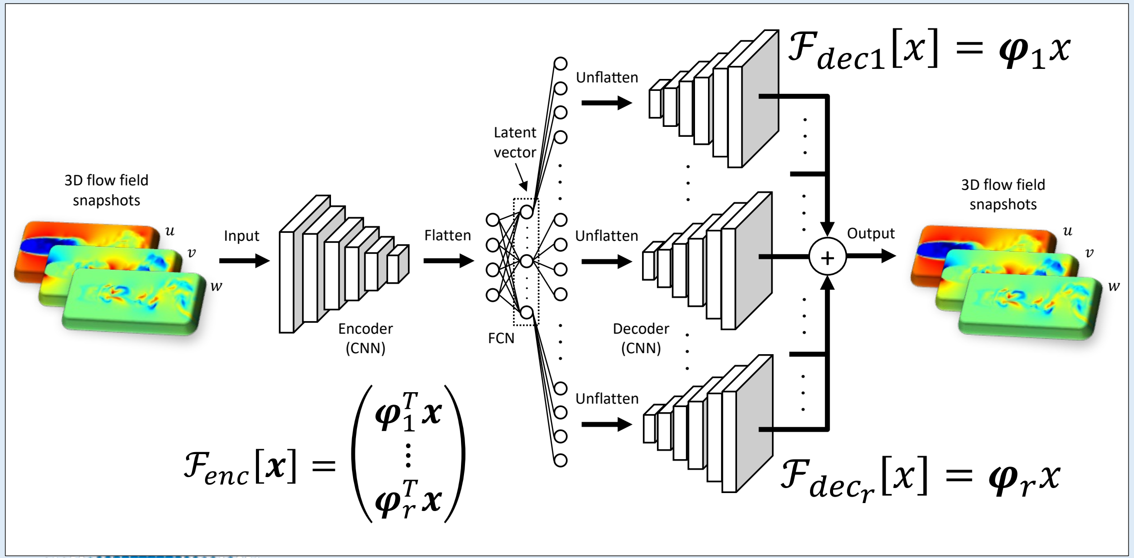

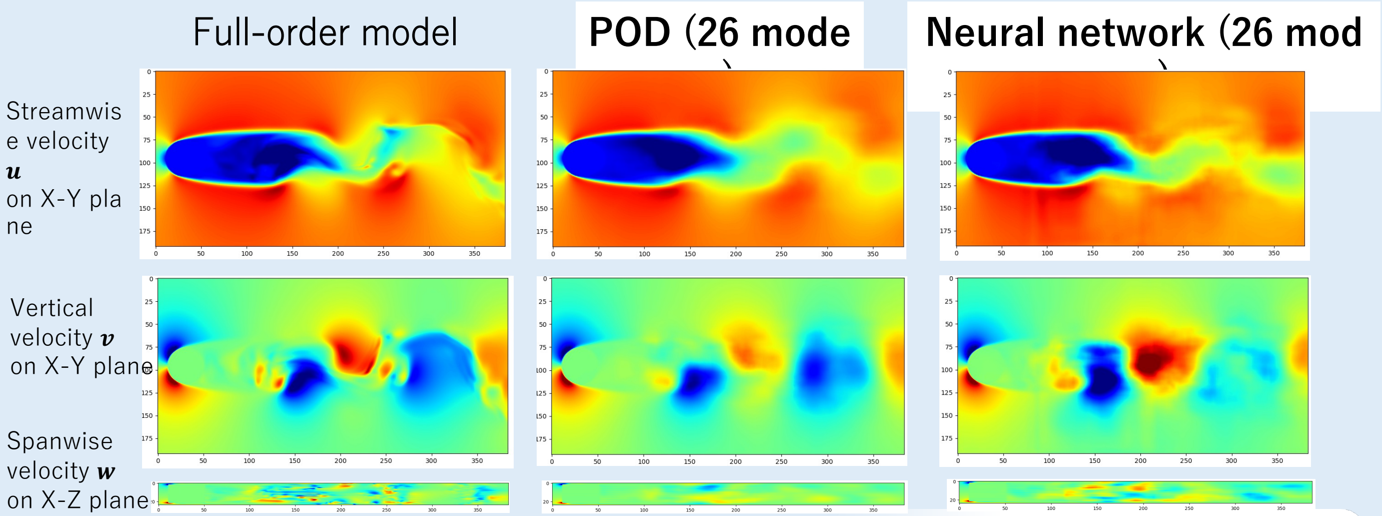

[1] Ando, K., Onishi, K., Bale, R., Tsubokura, M., Kuroda, A., Minami, K. (2021). Nonlinear Mode Decomposition and Reduced-Order Modeling for Three-Dimensional Cylinder Flow by Distributed Learning on Fugaku. In: Jagode, H., Anzt, H., Ltaief, H., Luszczek, P. (eds) High Performance Computing. ISC High Performance 2021. Lecture Notes in Computer Science(), vol 12761. Springer, Cham. https://doi.org/10.1007/978-3-030-90539-2_8
[2] T. Murata, K. Fukami, and K. Fukagata, “Nonlinear mode decomposition with convolutional neural networks for fluid dynamics,” J. Fluid Mech., vol. 882, A13., 2020, doi:10.1017/jfm.2019.822.

Development of NN for High-resolution, Real-Time Tsunami Flood Prediction (Fumihiko Imamura group [1])

- Tsunami simulations to generate training data
  - Training Input data: Tsunami waveform in offshore areas
  - Training Output data: Flooding conditions in coastal areas
- Training an AI model to predict flooding condition in coastal areas from Tsunami wave format in offshore

→ This approach makes it possible to accurately and rapidly obtain detailed flooding forecast before landfall of Tsunami

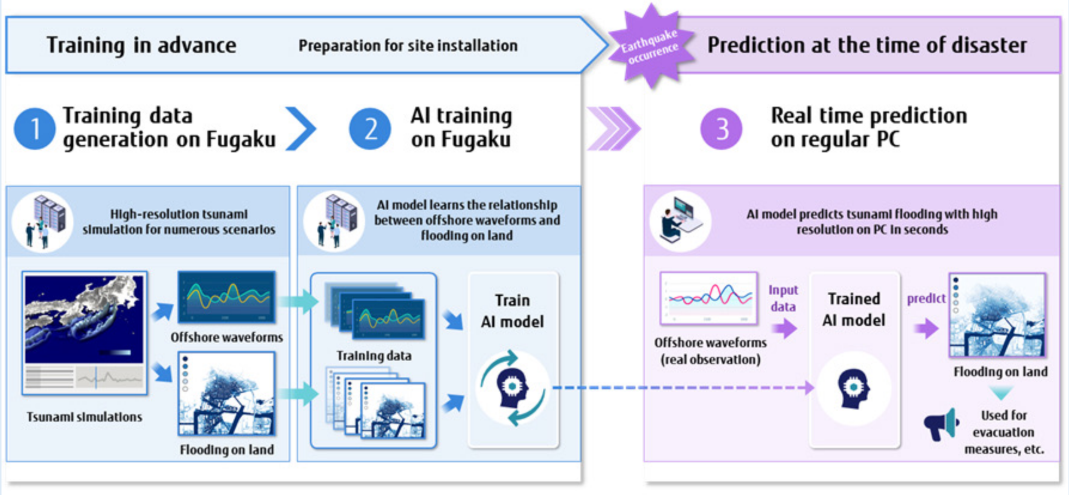

Fig. 1 Overview of tsunami prediction with AI

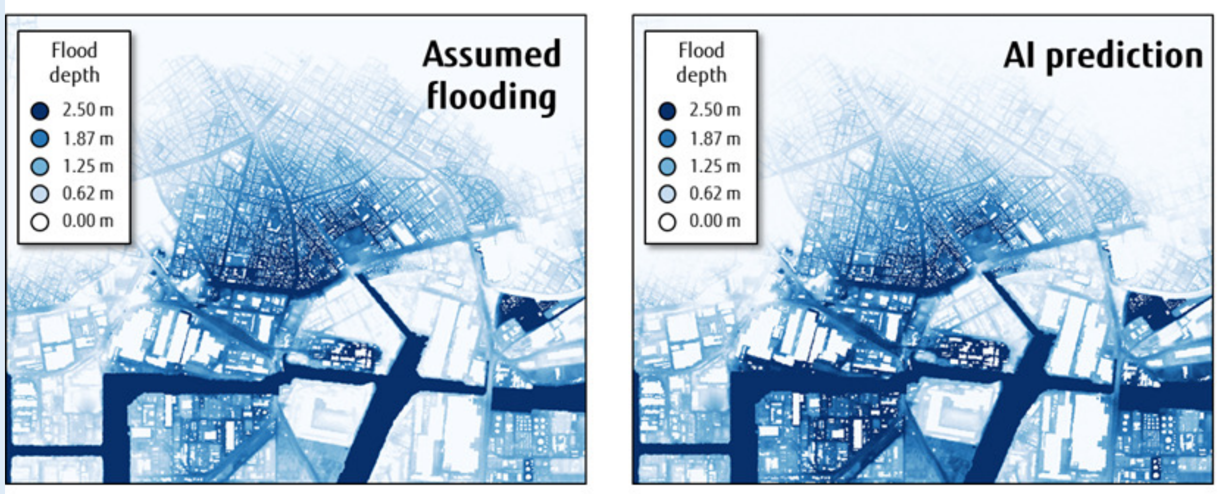

Fig 2. Comparison between anticipated flooding (tsunami source model created by Cabinet Office of Japan with tripled wave heights) of Nankai Trough Megathrust Earthquake and prediction results of newly developed AI

[1] (Press release) International Research Institute of Disaster Science, Tohoku University, Earthquake Research Institute, The University of Tokyo, Fujitsu Laboratories Ltd.Fujitsu Leverages World's Fastest Supercomputer 'Fugaku' and AI to Deliver Real-Time Tsunami Prediction in Joint Project

Elucidation of the Cause and Diversity of Cancer using Large-scale Data Analysis and AI Tech. [1]

- Even if a patient receives a targeted cancer drug therapy, the appearance of drug-resistant cancer cells represents an ongoing threat to full remission
  - The mechanism for how certain cancers become drug resistant remains unclear
- Fujitsu implemented parallel conditional and causal algorithms
  - Utilizing Fujitsu's "Wide Learning" AI technology to extract combinations of potential genes relating to the emergence of drug resistance based on statistical information
  - Also, maximizing computational performance with the supercomputer Fugaku
- Fugaku, Fujitsu and TMDU were able to search the entire human genome for conditions and causality within a single day and determine the genes that cause resistance to drugs used to treat lung cancer

Causality of Gefitinib resistance in cancer cells suggested by the new technology

Gene names and expression levels	Indications regarding genes in existing medical literature	Possible findings from existing medical literature
Low expression level of *ZNF 516*	Contributes to the suppression of *EGFR* expression. Suppresses cell proliferation by inhibiting the EGFR pathway*	A low expression level of *ZNF 516* may indicate activation of the EGFR pathway targeted by Gefitinib.
Low expression level of *E2F6*	*E2F1* and *E2F2* stimulate cell proliferation by cycling the cell cycle pathway. *E2F6* inhibits this function.	A low expression of *E2F6* may promote cell proliferation (not covered by the EGFR pathway).
Low expression level of *EMX1*	Known to suppress the Wnt pathway (which promotes cell cycle pathways).	A low expression level of *EMX1* may not suppress the Wnt pathway and may promote cell proliferation (not covered by the EGFR pathway).

* mechanism that activates cell growth when stimulated by a biomolecule called epidermal growth factor (EGF)

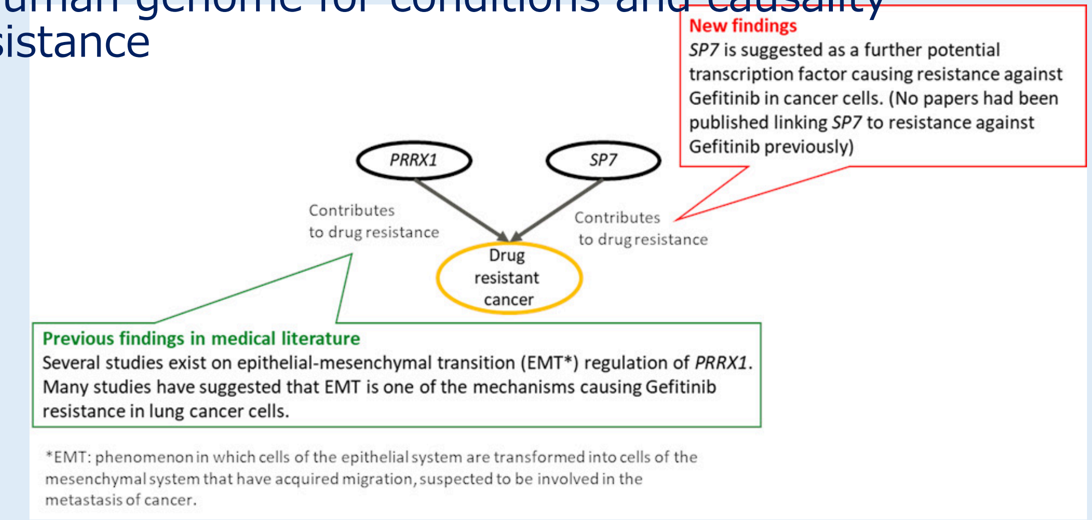

[1] (Press release) Fujitsu Limited, Tokyo Medical and Dental University," Fujitsu and Tokyo Medical and Dental University leverage world's fastest supercomputer and AI technology for scientific discovery to shed light on drug resistance in cancer treatment, March 7, 2022 (https://www.fujitsu.com/global/about/resources/news/press-releases/2022/0307-01.html)

Forecasting the Expansion of Supernova Shells Using Deep Learning toward High-Resolution Galaxy Simulations (Kaiya Hirashima et al. [1])

- Goal: achieve the highest-resolution galaxy simulation which can resolve individual stars on the supercomputer Fugaku.
- One of the bottlenecks is the calculation including small scall phenomena (e.g., supernovae)
- Just using massively parallel computers cannot solve it.
- And algorithm that uses deep learning to resolve the bottlenecks.

→ The deep learning model to forecast the expansion of supernova shells

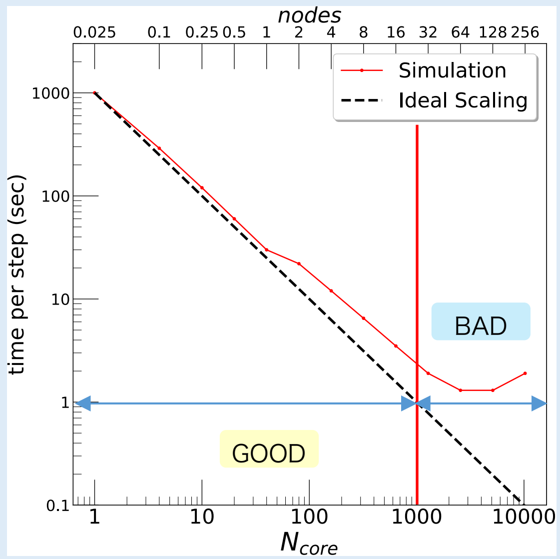

Strong Scaling of GADGET-4
(Based on [5] Figure 63)

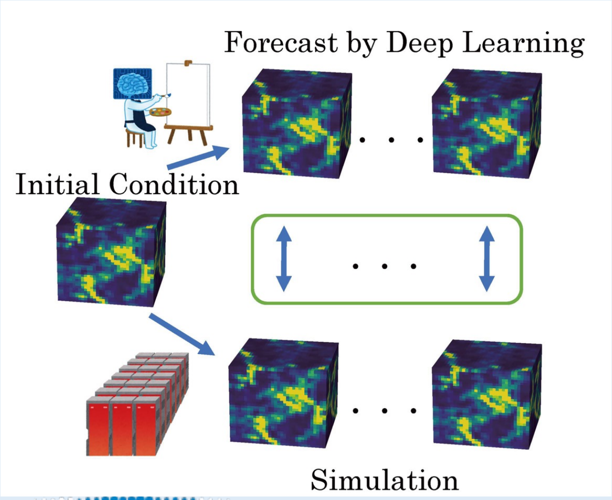

❖ **They applied the framework for video predictions to forecasting the results of simulations**
- Prediction of spatial features: CNN
- Prediction of temporal changes: LSTM

❖ **Training data**
- Input: 1 image, Density distribution just before the explosion
- Output: 19 images, Spatiotemporal changes in density after the explosion

❖ **Training & Prediction**
- The training takes several days using an NVIDIA GeForce RTX 3090 GPU.
- The estimation itself takes only a few second.

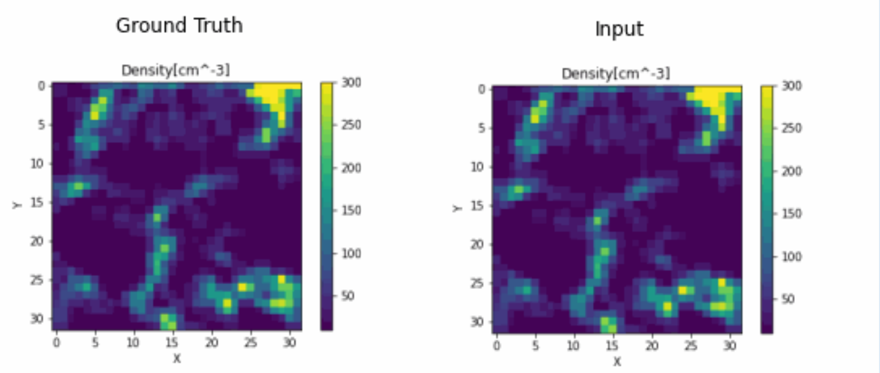

[1] K. Hirashima and K. Moriwaki and M. S. Fujii and Y. Hirai and T. Saitoh and J. Makin, “Predicting the Expansion of Supernova Shells for High-Resolution Galaxy Simulations Using Deep Learning”, Journal of Physics: Conference Series, March, 2022
[2] Springel et al. (2021)

Other AI-based Scientific Achievements (Fugaku is not used)

by R-CCS research team & Priority application programs

Remote Sensing by Deep Learning From Simulated Data for Flood and Debris-Flow Mapping for flood disaster damage estimation (Naoto Yokoya et al., 2020 [1])

- Yokoyama et al. proposed a framework that estimates the inundation depth (maximum water level) and debris-flow-induced topographic deformation
  - From remote sensing imagery by integrating deep learning and numerical simulation.
  - A water and debris-flow simulator generates training data for various artificial disaster scenarios.
- They showed that regression models based on Attention U-Net and LinkNet architectures trained on such synthetic data can predict the maximum water level and topographic deformation from a remote sensing-derived change detection map and a digital elevation model
- The proposed framework has an inpainting capability, thus mitigating the false negatives that are inevitable in remote sensing image analysis qualitatively

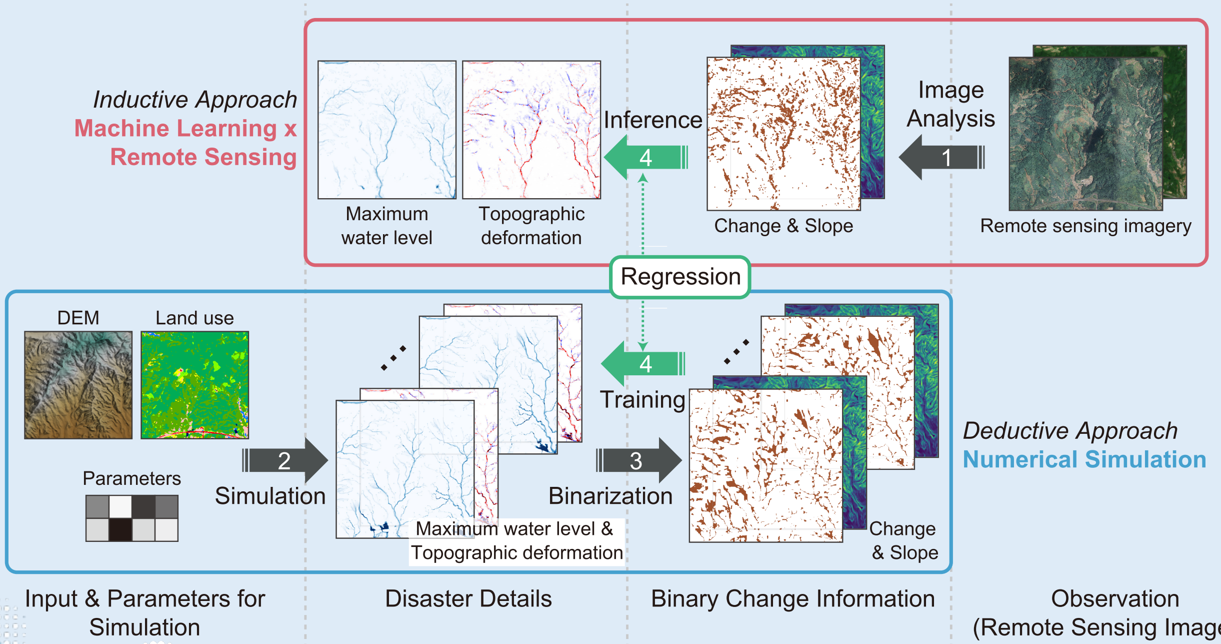

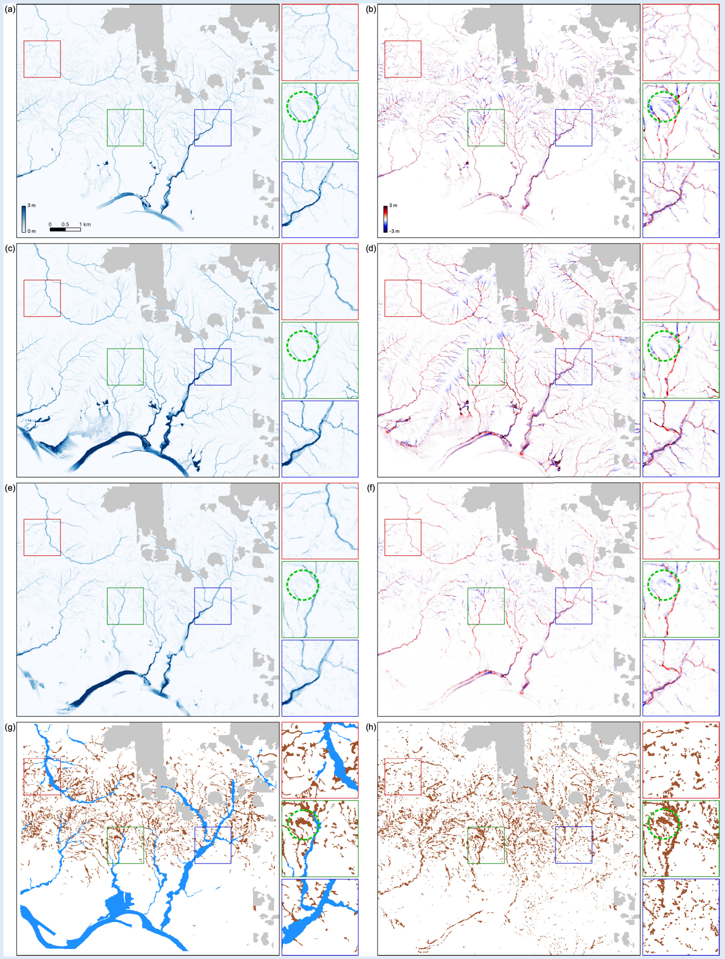

Fig. 8. (a), (c), and (e) Maximum water level and (b), (d), and (f) topographic deformation of (a) and (b) reference simulation, (c) and (d) average simulation, and (e) and (f) our method with (g) reference map of (blue) flood and (brown) debris-flow extent and (h) binary change detection map. Mask for clouds and no data is shown in gray.

[1] Yokoya, Naoto & Yamanoi, Kazuki & He, Wei & Baier, Gerald & Adriano, Bruno & Miura, Hiroyuki & Oishi, Satoru. (2020). Breaking Limits of Remote Sensing by Deep Learning From Simulated Data for Flood and Debris-Flow Mapping. IEEE Transactions on Geoscience and Remote Sensing. PP. 1-15. 10.1109/TGRS.2020.3035469.

Compression of Time Evolutionary Image Data through Predictive Deep Neural Networks (Rupak Roy et al., 2021 [1])

- Roy et al. proposed new AI-drive data compressor (TEZIP) for time evolutionary data
  - We train PredNet to learn how pixels move and how fast by inputting a number of time evolutional frames
  - When compressing frames, we predict future frames from original base data, compute delta values and apply series of encoding

→ We only store (1) base frame data and (2) compressed data

- They achieved higher compression ration compared to existing video encoder (Zstd, HFYU, FFV1, x.265)

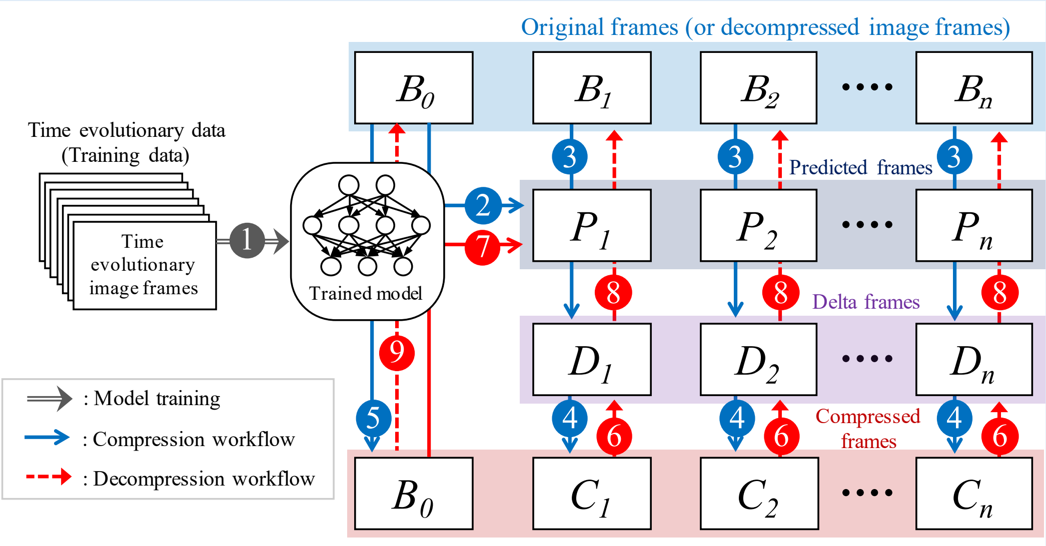

Fig. 1. Workflows of TEZIP (de)compression

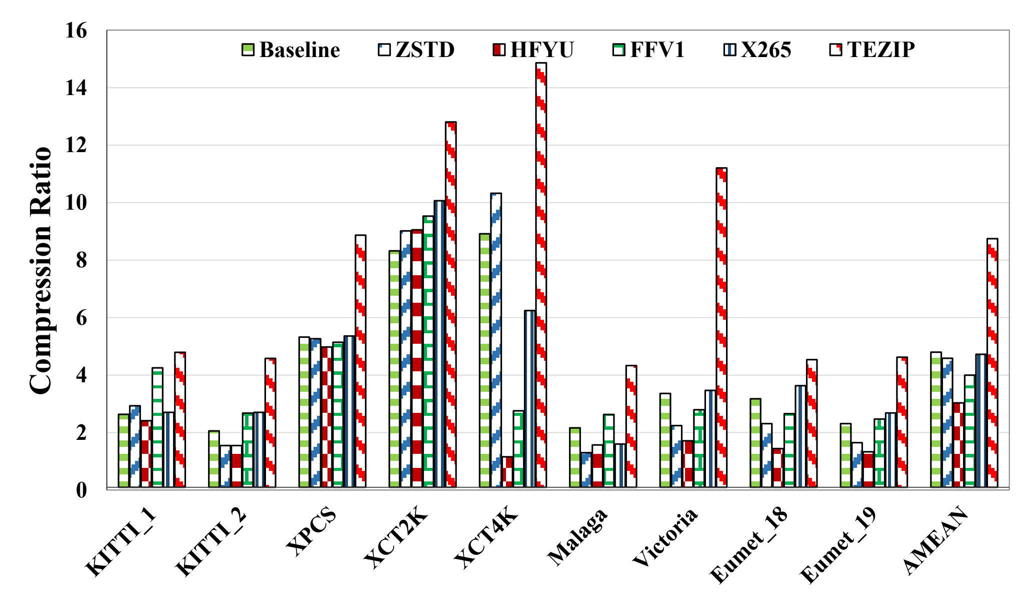

Fig. 8. Compression ratio with lossless compressors.

[1] Rupak Roy, Kento Sato, Subhadeep Bhattacharya, Xingang Fang, Yasumasa Joti, Takaki Hatsui, Toshiyuki Hiraki, Jian Guo and Weikuan Yu, "Compression of Time Evolutionary Image Data through Predictive Deep Neural Networks", In the proceedings of the 21 IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGrid 2021), May, 2021

Automatic Quantum Circuit Encoding of a given Arbitrary Quantum State (Tomonori Shirakawa et al. [1])

- Quantum computer can in principle treat an exponentially large amount of data, however, a major challenge in practice to represent classical data such as a classical image into a quantum circuit
- Inspired by a tensor network method, they have proposed a quantum-classical hybrid algorithm to construct an optimal quantum circuit for classical data that is represented as a quantum state by the amplitude encoding
- The proposed algorithm employs as an objective function the absolute value of fidelity $F = \langle 0|\hat{C}^\dagger|\Psi\rangle$, which is maximized iteratively to construct an optimal quantum circuit $\hat{C}$ with controlled accuracy.

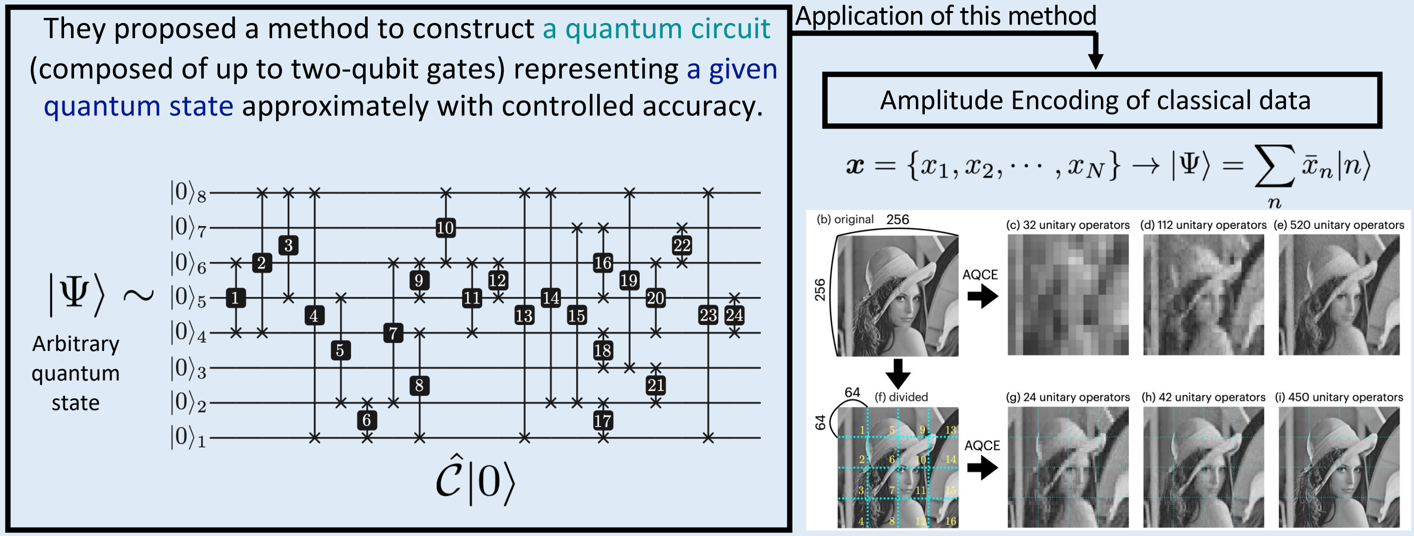

Fig 1: Quantum circuit encoding of a gray scale. (a) Labeling of two-dimensional classical data (with 8 × 8 pixels, as an example). (b) Original picture with 256 × 256 pixels. (c)-(e) Pictures reconstructed by decoding the quantum circuit states. (f) Original picture divided into 16 pieces. (g)-(i) Pictures reconstructed by decoding each quantum c ircuit state

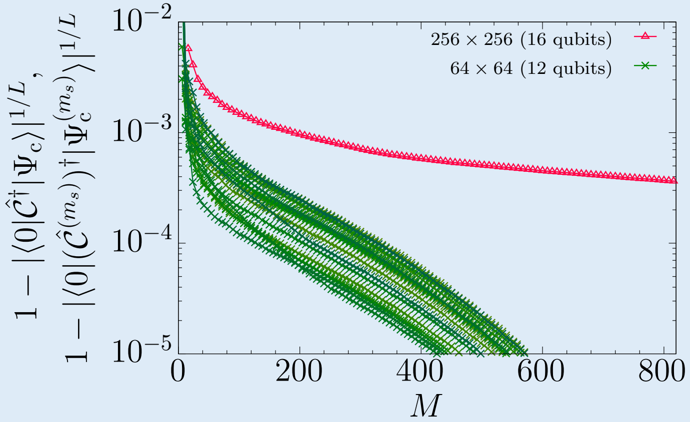

Fig 2: Fidelity per site between the quantum state qubits representing the original picture in Fig. 1(b) [the msth s egment of the original picture in Fig.1(f)] and the quantu m circuit state indicated by red triangles (green crosses) .

[1] Tomonori Shirakawa, Hiroshi Ueda, Seiji Yunoki, "Automatic Quantum Circuit Encoding of a given Arbitrary Quantum State ", arXiv:2112.14524, Dec, 2021

Other AI researches by Priority Application Projects (Fugaku is not used)

- **Comprehensive gene network analysis by XAI for uncovering molecular mechanism of disease (Heewon Park et al., 2020 [1])**
  - Although various approaches were developed to gene network analysis, comprehensive interpretation of the massive networks remains a challenge
  - They conducted comprehensive analysis of the massive gene networks with the use of explainable artificial intelligence (XAI) approaches
    - DeepTensor and Tensor Reconstruction-based Interpretable Prediction (TRIP), proposed by Maruhashi et al. 2018; 2020
  - The use of the XAI enables us to overcome limitation of existing gene network analysis, i.e., narrow angle in the large-scale gene network analysis, and this leads to a better understanding of molecular interplay involved to disease
- **Machine learning by neural network for quantum many-body solver and experimental data analysis (Yusuke Nomura et al., 2021 [2])**
  - Quantum many-body problems are known to show NP hard difficulty
  - We have developed a solver by using Boltzmann machines to approximate solutions accurately and efficiently
  - It was applied to a challenging problem of a quantum spin model and established the existence of quantum spin liquid phase, which bears a long-ranged quantum entanglement. On another application of machine learning, we have extracted electron self-energy, which is hidden in the photoemission experimental data. It has established the existence of prominent resonance peaks which are responsible for the high-temperature superconductivity.

[7] Park, Heewon Maruhashi, Koji, Yamaguchi, Rui, Imoto, Seiya, Miyano, Satoru, "Global gene network exploration based on explainable artifi cial intelligence approach", Nov., 2020
[8] Nomura, Yusuke and Imada, Masatoshi, "Dirac-Type Nodal Spin Liquid Revealed by Refined Quantum Many-Body Solver Using Neural-Net work Wave Function, Correlation Ratio, and Level Spectroscopy", American Physical Society, 10.1103/PhysRevX.11.031034, Nov., 2021

Other AI researches by Priority Application Projects (Fugaku is not used)

- **Facilitating ab initio configurational sampling of multicomponent solids using an on-lattice neural network model and active learning (Kasamatsu et al., 2020 [1])**
  - Kasamatsu et al. proposed a scheme for ab initio configurational sampling in multicomponent crystalline solids using Behler-Parinello type neural network potentials (NNPs)
  - The NNPs are trained to predict the energies of relaxed structures from the perfect lattice with configurational disorder instead of the usual way of training to predict energies as functions of continuous atom coordinates.
    - Training set bias is avoided through an active learning scheme
    - This enables bypassing of the structural relaxation procedure which is necessary when applying conventional NNP approaches to the lattice configuration problem
  - This idea is demonstrated on the calculation of the temperature dependence of the degree of A/B site inversion in $MgAl_2O_4$, which is a multivalent system requiring careful handling of long-range interactions
  - The present scheme may serve as an alternative to cluster expansion for 'difficult' systems, e.g., complex bulk or interface systems with many components and sublattices that are relevant to many technological applications today
- **Density functional theory from supervised learning (Ryo Nagai et al. 2021 [2])**
  - Density functional theory (DFT) is the standard electronic structure theory and is widely used as a basis for materials design
  - DFT is based on the Hohenberg-Kohn theorem that there is a one-to-one correspondence between particle density and energy, a relationship that should be machine learnable, but the lack of this relationship makes the accuracy limited
  - They recently proposed a method to establish this relationship using neural network methods
  - They found that the DFT developed in this study can improve the accuracy not only for molecular systems, for which the training set can be obtained by a very accurate quantum chemical method, but also for solids, for which the training set is unavailable but theoretically derived physical conditions work alternatively
  - This paves a systematic way for further accuracy and will serve as a tool for developing the ultimate material database

[1] Kasamatsu, Shusuke and Motoyama, Yuichi and Yoshimi, Kazuyoshi and Matsumoto, Ushio and Kuwabara, Akihide and Ogawa, Takafumi, "Facilitating {¥it ab initio} configurational sampling of multicomponent solids using an on-lattice neural network model and active learning, 10.48550/ARXIV.2008.02572"
[2] Nagai, Ryo and Akashi, Ryosuke and Sugino, Osamu, "Machine-Learning-Based Exchange-Correlation Functional with Physical Asymptotic Const

Highlight of Efforts in Bridging AI <-> HPC in R-CCS (Fugaku)

by R-CCS High Performance Artificial Intelligence Systems Research Team

R-CCS HPAIS Team

- **PI:** Mohamed Wahib
- **Senior Research Scientists:** Jun Igarashi
- **Research Scientists:** Aleksandr Drozd, Emil Vatai
- **Visiting Scientists:** Rio Yokota (TokyoTech), Balazs Gerofi (Intel)
- **Interns:** Soyturk (Koc U., Turkey), Zhang (Hokkaido U), Pucceti (SNS)

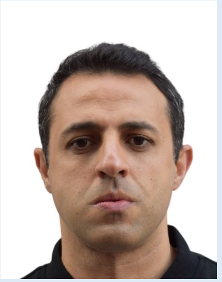

Mohamed Wahib

Jun Igarashi

Aleksandr Drozd

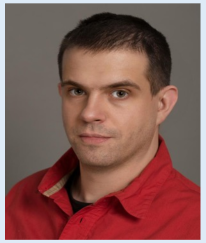

Emil Vatai

Rio Yokota

Balasz Gerofi

Muhammet Soyturk

Enzhi Zhang

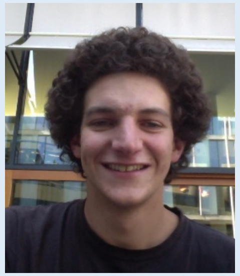

Giovanni Pucceti

Bottlenecks in Scaling AI

Unlike traditional HPC applications, AI workloads have multiple bottlenecks

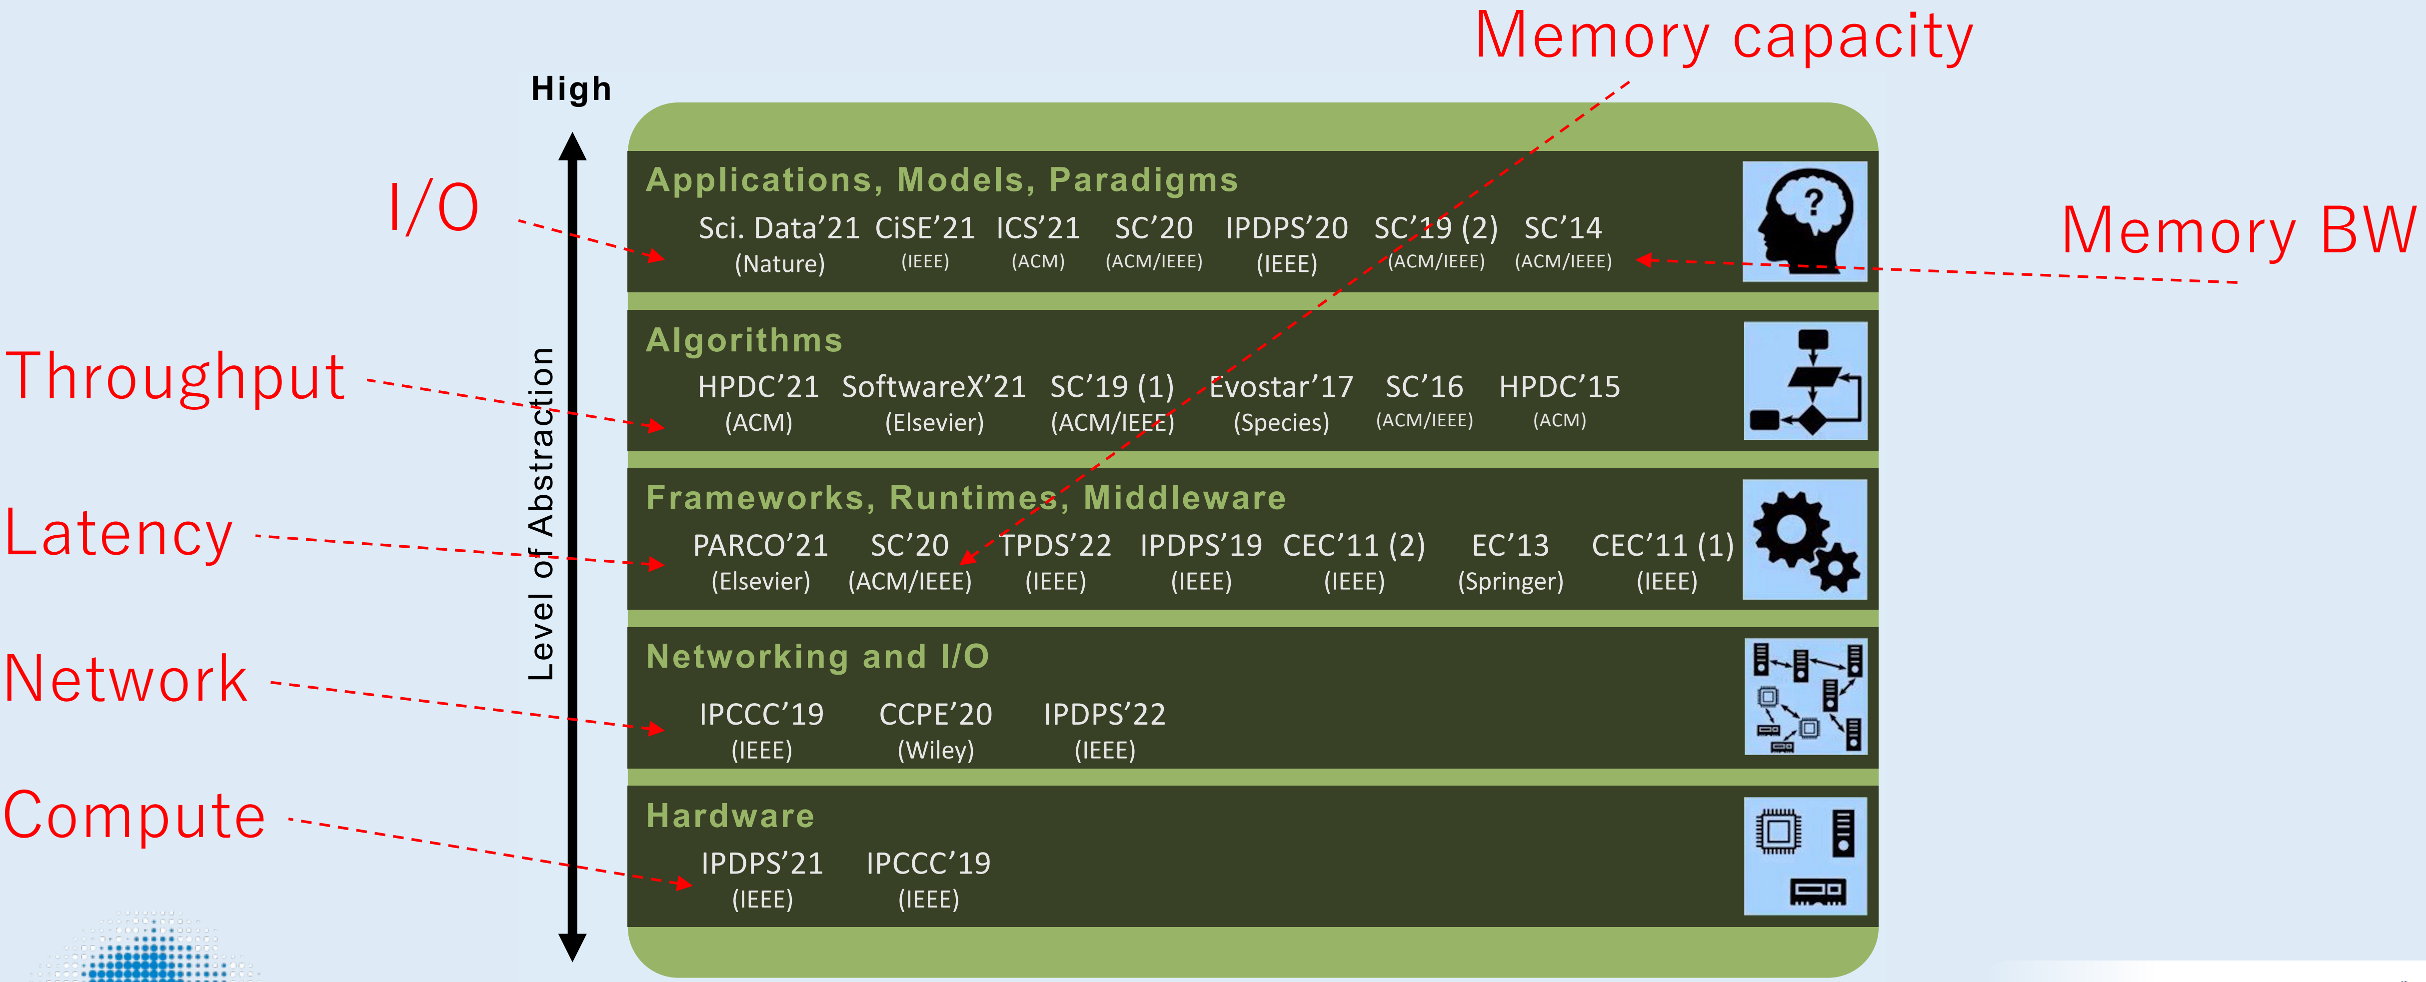

Overview of previous research: perspective of AI+HPC computing stack

ParDNN: An Oracle for Characterizing and Guiding Large-Scale Training of Deep Neural Networks

- **<u>Different forms of model parallelism are emerging</u>**
- **We provide a model-driven analysis and a utility to help in detecting the limitations and bottlenecks of different parallelism approaches at scale**

Overview of ParDNN

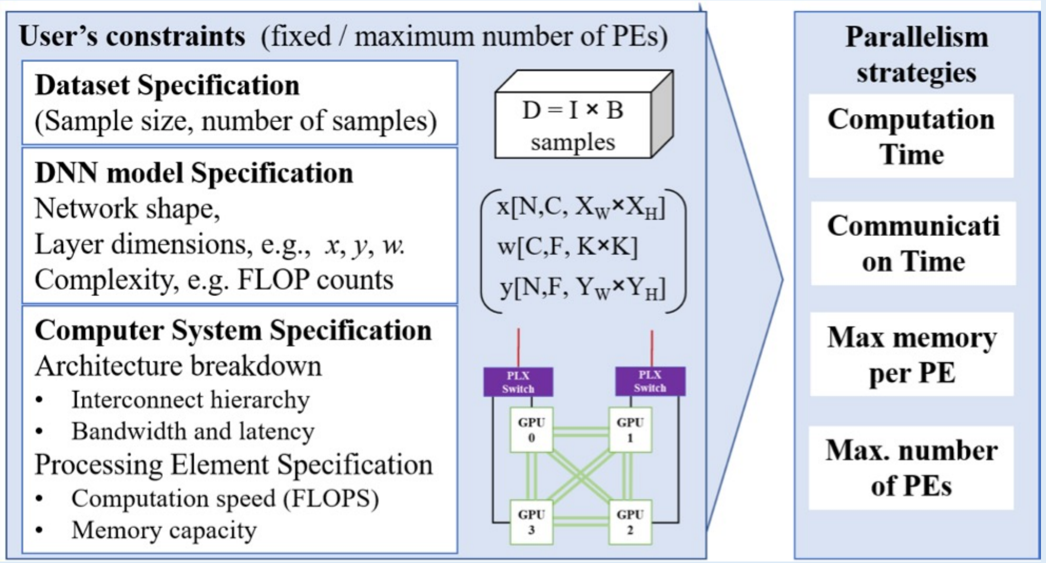

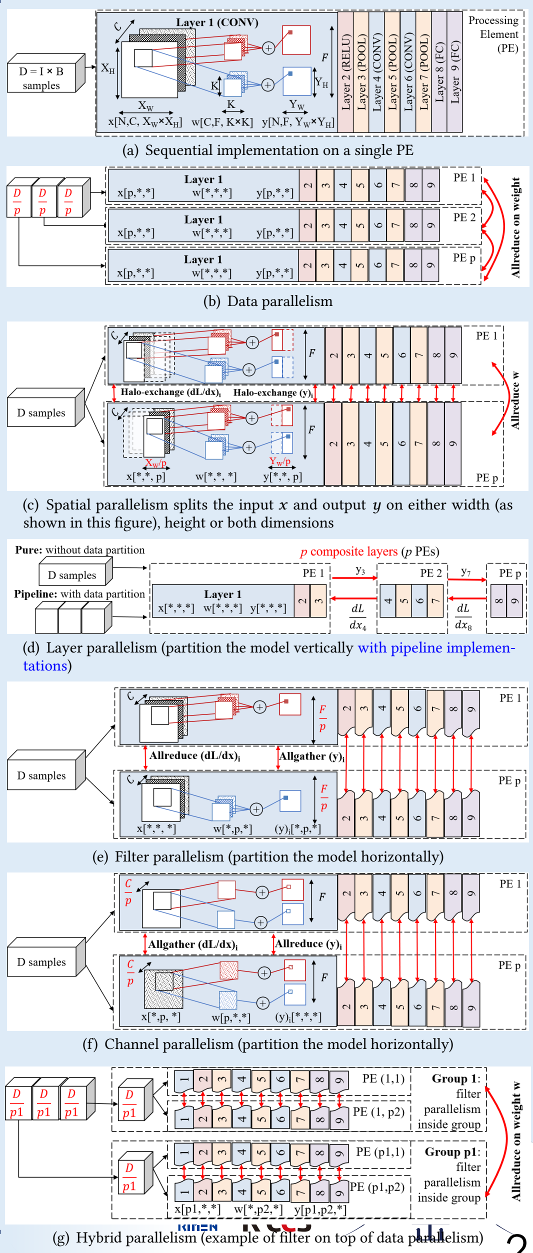

ACM HPDC'21 A. Kahira, T. Nguyen,, L. Gomez, Ryousei Takano, R. Badia, M. Wahib: ParDNN: An Oracle for Characterizing and Guiding Large-Scale Training of Deep Neural Networks, ACM Proceedings of the International Symposium on High-Performance Parallel and Distributed Computing 2021 (ACM HPDC'21)

GTOPX Space Mission Benchmarks

- **Landscape Analysis for ML**
  - **Important to understand importance of parameters**
  - **Space mission benchmarks**

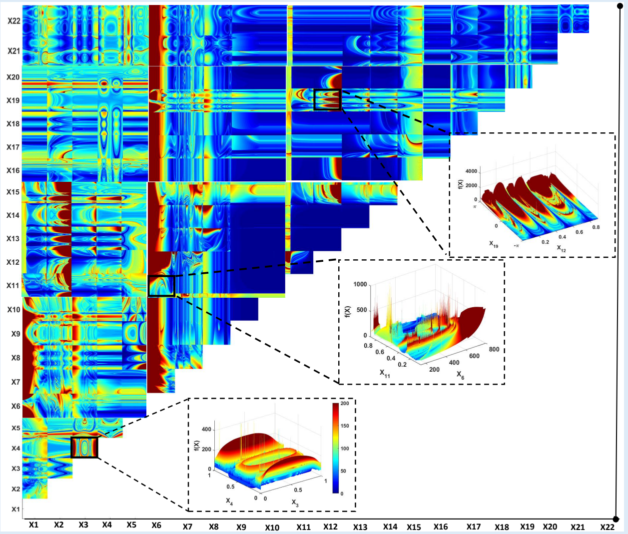

(b)

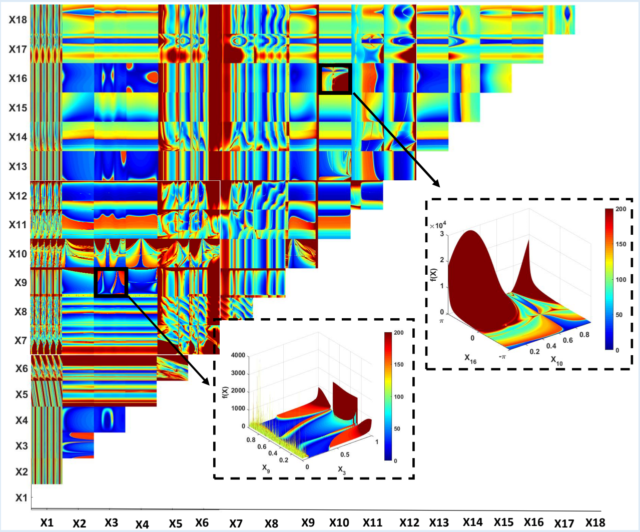

(c)

Elsevier SoftwareX M. Schlueter, M. Neshat, M. Wahib, M. Munetomo, M. Wagner: GTOPX Space Mission Benchmarks Elsevier SoftwareX Volume 14, June 2021

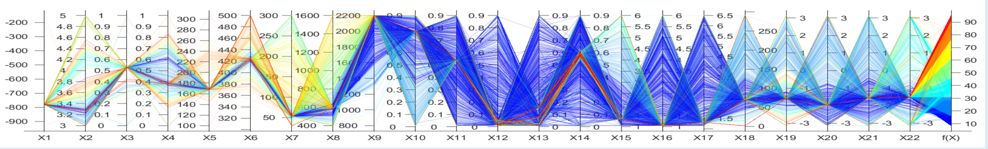

(a)

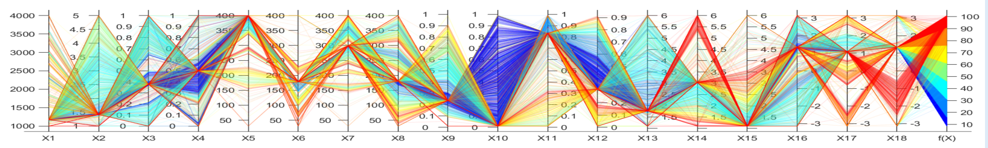

(b)

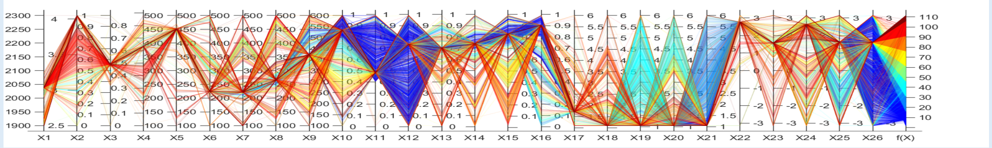

(c)

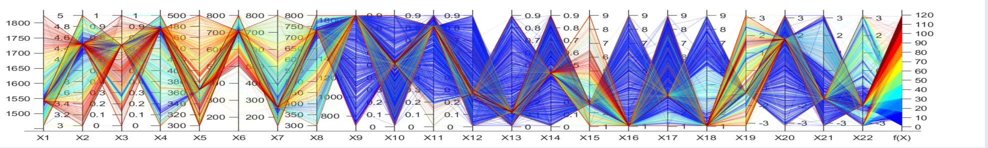

(d)

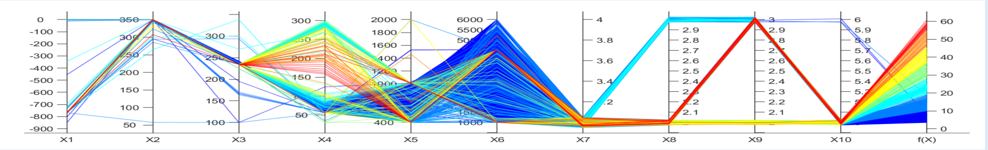

(e)

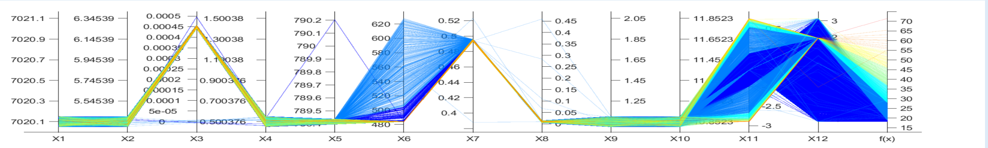

Scaling Distributed Deep Learning Workloads beyond the Memory Capacity with KARMA

Performance on V100. For all figures, only the first reported mini-batch size (x-axis) fits in memory

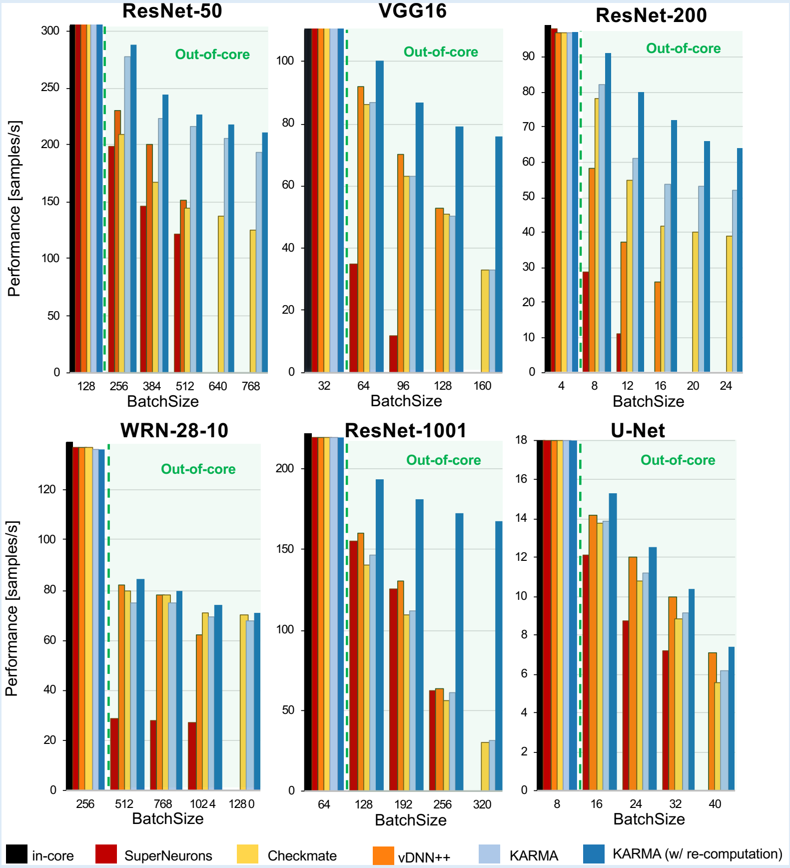

- ➢Concurrency-driven out-of-Core
  - ➢Capacity-based interleaved with recompute
- ➢1.52x over state-of-the-art (single GPU)
- ➢First out-of-core to support multi-GPU
  - ➢Heterogeneity and careful orchestration
- ➢Outperforming DP+MD with out-of-core
  - ➢Experiments with up to 2,048 GPUs

ACM/IEEE SC'20 Mohamed Wahib, Haoyu Zhang, Truong Thao Nguyen, Aleksandr Drozd, Jens Domke, Lingqi Zhang, Ryousei Takano, Satoshi Matsuoka, Scaling Distributed Deep Learning Workloads beyond the Memory Capacity with KARMA, ACM/IEEE Proceedings of Supercomputing 2020 (ACM/IEEE SC'20)

Optimizing Network Architecture for Deep Learning

Problem: Number of involved GPUs (*P*) becomes too big, 1000s GPUs.

- ➢ Ring-based algorithm O(*P*)α : latency factor increase.
- ➢ Halving-Doubleling algorithms O(log*P*)α : network congestion of communication.

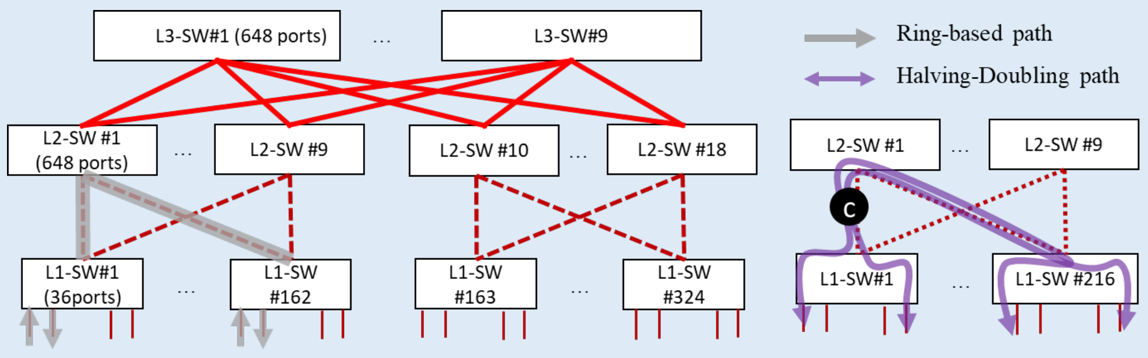

(a) Non-blocking Three-level Fat-tree (up to 5832 NICs)

(b) Non-blocking Two-level Fat-tree (up to 3888 NICs)

Proposal: Distributed Loop Network topology

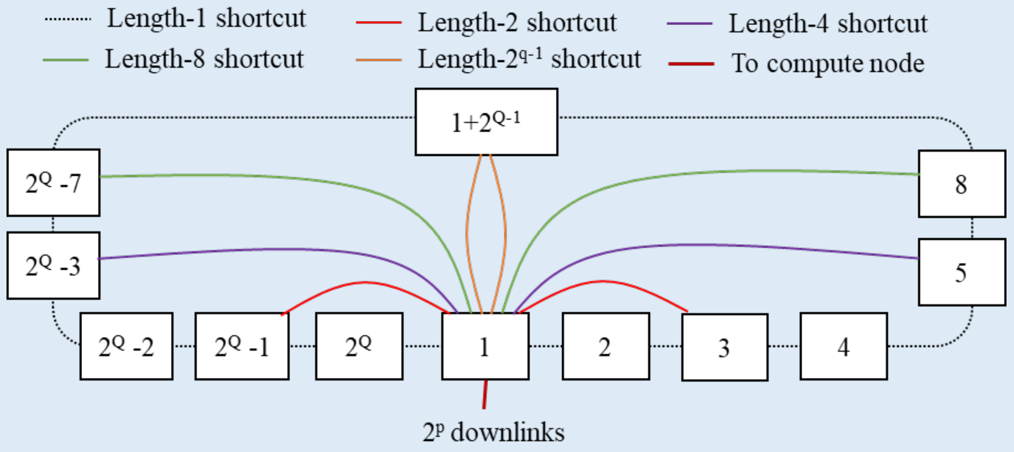

Latency factor

$$2\log(k)\alpha_{intra} + 2\log\left(\frac{P}{k}\right)\alpha_{inter}$$

Bandwidth factor

$$2\frac{(k-1)}{k}N\beta_{intra} + \frac{2(P-k)}{P}\frac{N}{k}\beta_{inter}$$

Pipelined Halving – Doubling Algorithm

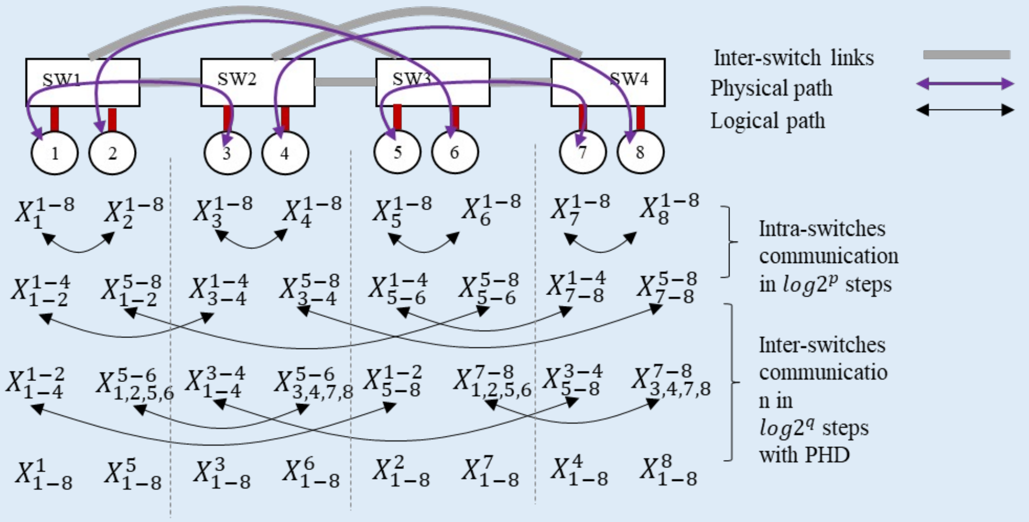

CCGRID'21 Truong Thao Nguyen, Mohamed Wahib: An Allreduce Algorithm and Network Co-design for Large-Scale Training of Distributed Deep Learning, In proceedings of the IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing 2021

AI for Scientific Codes

Use AI for code auto generation to aid scientific programmers in producing HPC programs.

Compiler Technologies

- ✓ Code auto-generation
- ✗ High-level context lost
- ✗ Constrained time budget

High-level Code-Transformation

- ✓ High-level context
- ✗ Not scalable
- ✗ Problem-specific

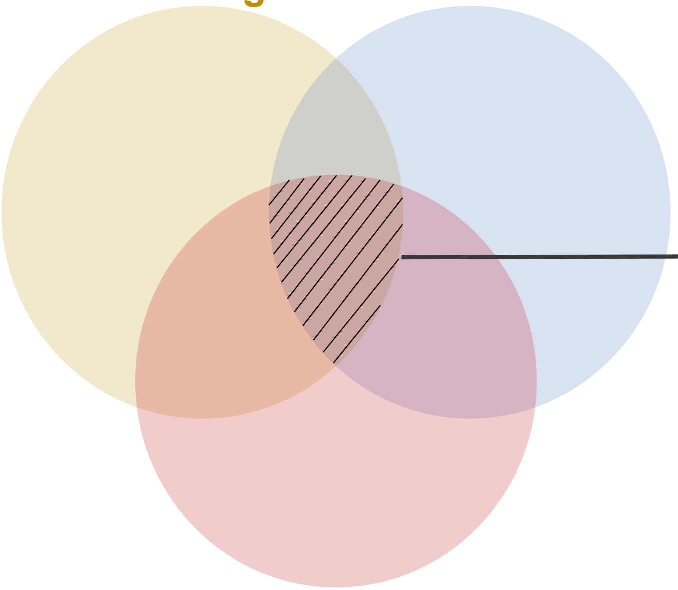

Auto-tuning

- ✓ Scalable exploration; general
- ✗ No code auto-generation

This Proposal

- Compiler tech.
- High-level trans.
- Auto-tuning

+ Machine Learning

- ✓ Scalable exploration; general
- ✓ High-level context
- ✓ Code auto-generation
- ✓ Elastic time budget
- ✓ General and expandable approach

RIKEN R-CCS 富岳

Fugaku is driving innovation in bridging HPC <-> AI, and will continue to do!

Thank you.

We are hiring!