Cygnus-BD: the Big Memory Supercomputer for HPC, Big Data and AI

Taisuke Boku

Director, Center for Computational Sciences

University of Tsukuba, *http://www.ccs.tsukuba.ac.jp*

(with Osamu Tatebe, CCS)

taisuke@ccs.tsukuba.ac.jp, tatebe@cs.tsukuba.ac.jp

CCS at University of Tsukuba

- Center for Computational Sciences
- Established in 1992
  - 12 years as Center for Computational Physics
  - Reorganized as Center for Computational Sciences in 2004
- Daily collaborative researches with two kinds of faculty members (45 in total)
  - Computational Scientists
    who have NEEDS (applications)
  - Computer Scientists
    who have SEEDS (system & solution)
- One of national supercomputer centers under MEXT, but we are Research Center (others are service centers)

History of PACS (PAX) series development at CCS

- 1977: research started by T. Hoshino and T. Kawai
- 1978: PACS-9 (with 9 nodes) completed
- 1996: CP-PACS, the first vendor-made supercomputer at CCS, ranked as #1 in TOP500

1978
1st gen: PACS-9

1980
2nd gen. PACS-32

1989
5th gen, QCDPAX

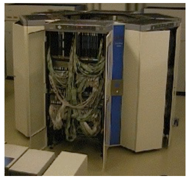

1996
6th gen: CP-PACS
Ranked #1 in TOP500

2006
7th gen: PACS-CS

2012~2013
8th gen: GPU cluster HA-PACS

2014
9th gen: COMA

2019
10th gen: Cygnus

Year	Name	Performance
1978	PACS-9	7 KFLOPS
1980	PACS-32	500 KFLOPS
1983	PAX-128	4 MFLOPS
1984	PAX-32J	3 MFLOPS
1989	QCDPAX	14 GFLOPS
1996	CP-PACS	614 GFLOPS
2006	PACS-CS	14.3 TFLOPS
2012~13	HA-PACS	1.166 PFLOPS
2014	COMA (PACS-IX)	1.001 PFLOPS
2019	Cygnus (PACS-X)	2.5 PFLOPS

- *co-design* by computer scientists and computational scientists toward “practically high speed computer”
- Application-driven development
- Sustainable development experience

HPC technology contributing to AI: Computation

- Neural Network Processing
  - early ’80s, Neural Network started to be studied for machine learning
  - supported by only poor processing power (CPU), no special hardware
  - just one middle layer, and results are not sufficient (low computation power)
    ⇨ not a Deep Learning (deep layered CNN-base machine learning)
- Accelerators, especially GPU
  - after GPU became attractive for numerical computing, GPU started to be introduced for NN
  - very regular and large capacity of computing ⇨ good for SIMD implementation
  - GPU is now the main player of AI (DL) and GPU vendors continue to build AI-oriented GPU as well as HPC use: NVIDIA and AMD (followed by Intel too)
- CPU instruction set
  - to support ML, FP16 (16-bit half precision) and BFloat16 (long mantissa) are introduced
    ⇨ recently with FP8
  - SIMD vector instruction is good to support them in up to 512bit

AI contribution to HPC: efficient data analysis and reduction

- efficient parameter space search
  - reducing the parameter search space by machine learning
    - climate simulation
    - astrophysics
    - life science
  - data matching on large data space
    - text base docking for creation of medicine
- efficient data analysis
  - finding the characteristics of phenomena
  - machine learning base data sorting and collection
- surrogate modeling

HPC for AI ⇔ AI for HPC

Storage technologies

- traditionally
  - large capacity but low performance (bandwidth, latency)
- TAPE ⇒ HDD ⇒ SSD ⇒ NVRAM (Non-Volatile RAM)
  - medium capacity and high performance (storage view)
    ⇨ large capacity and a bit low performance (main memory view)
  - NVRAM usability reaches to the practical use (Persistent Memory)
- data science driven by large capacity memory and high speed shared storage
  - filling the gap between memory capacity and computation performance
  - in-situ processing
  - high performance ad hoc storage

Why we need Big Memory

- CPU performance 50x, but memory size 3.8x in 8 years
- It matters for Data-driven and AI-driven Science
  - Memory size and Storage performance are really important
- Introduce Persistent Memory (NVRAM)
  - Memory mode for memory size and direct mode for storage performance

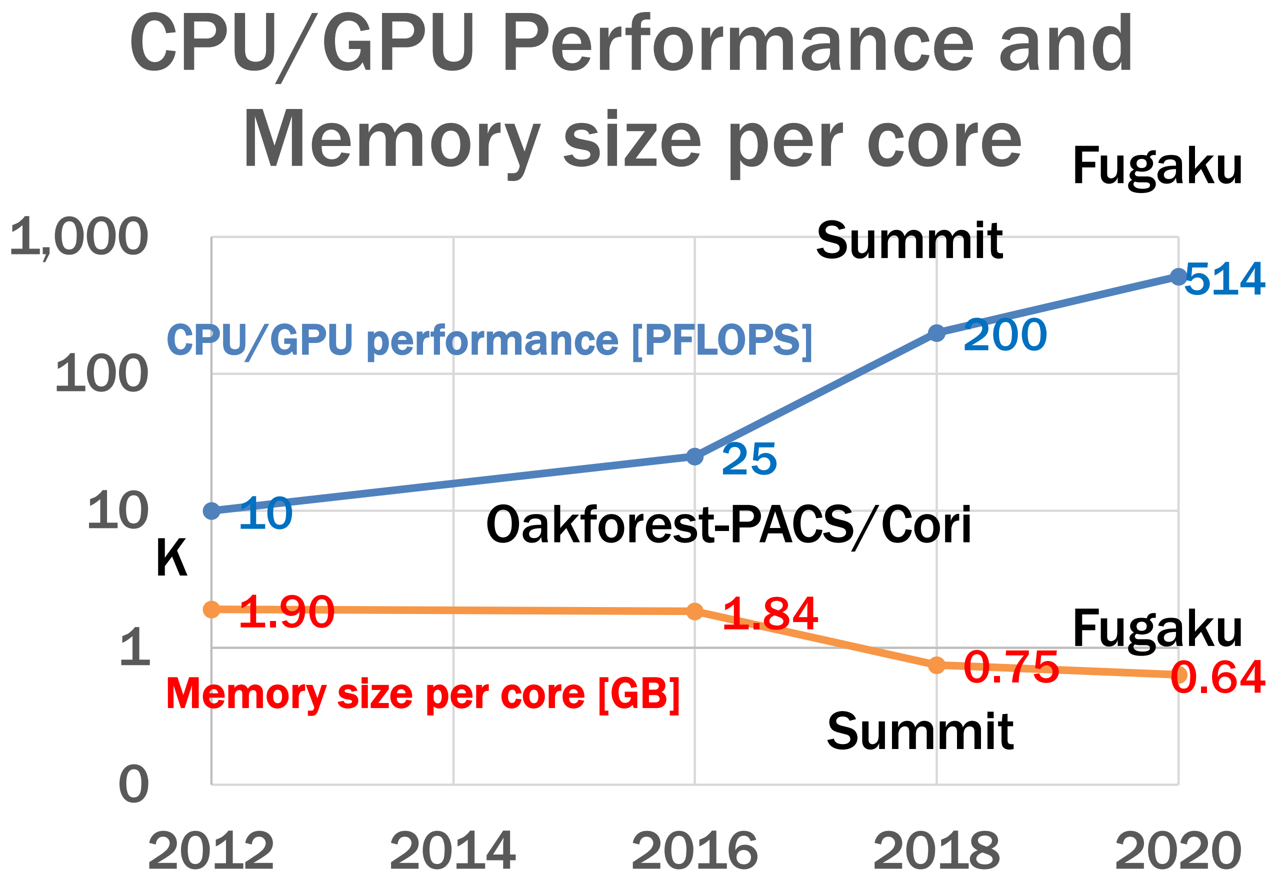

Persistent Memory

- One order better price/capacity
- Minimum latency is ~60 ns (similar to DRAM)
- ~Half of bandwidth
- Memory mode
  - Larger memory space without much performance penalty
  - Possible to use DRAM as last level cache
- App direct mode
  - Direct access to byte-addressable persistent memory and high-performance storage

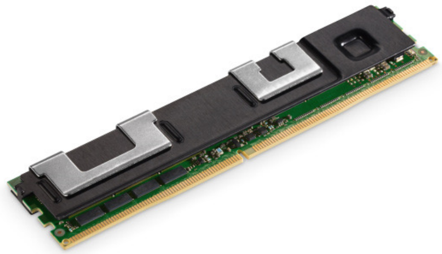

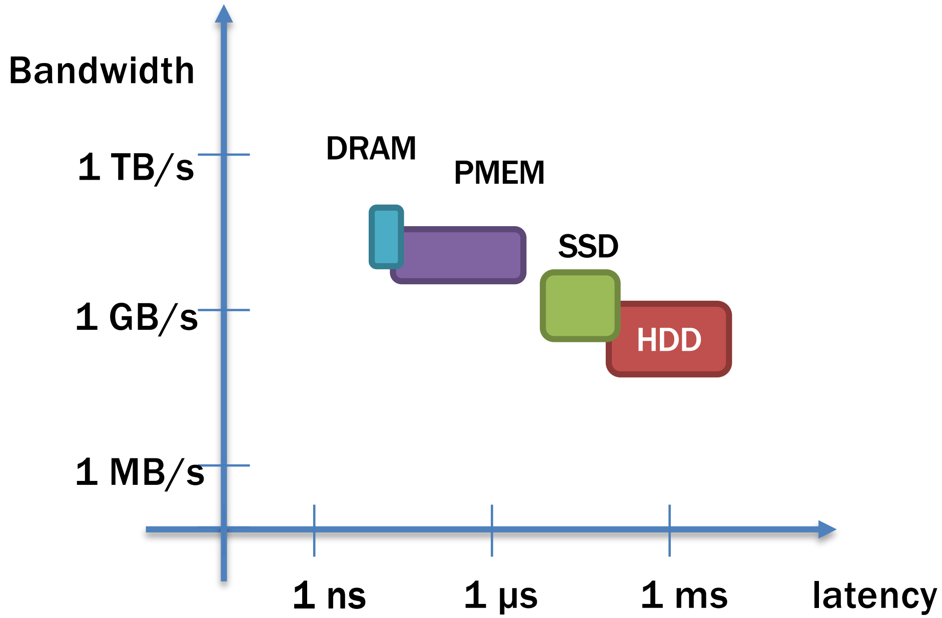

A Consideration of Future of Persistent Memory

- For HPC and Big Data AI/ML, large memory dramatically improves the app performance
  - Large buffer cache that critically improve I/O performance
  - No need to spill out to NVMe SSD
- Pmem is 10x more cost effective and 6.4x less power consumption than DRAM
  - DRAM (and flash devices) cannot replace Pmem.
- This device technology evolves with CXL (Compute Express Link)
- Research and industry activities are increasing
  - SNIA Persistent Memory + Computational Storage Summit on May 24-25, 2022
  - CXL: Getting Ready to Take-off on Aug 2, 2022
  - USENIX FAST Technical Papers

FAST	2018	2019	2020	2021	2022
# pmem papers / # total papers	1 / 23	6 / 26	1 / 23	5 / 28	6 / 28

Large dataset science (1): Astrophysics

- Data-pipelining on analysis with Subaru Telescope

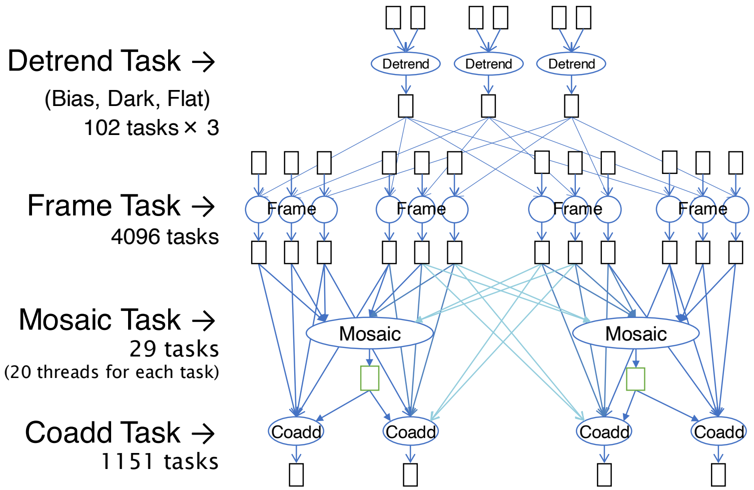

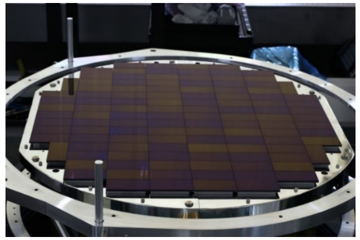

© NAOJ

- Currently done by staging of computation (through file)
  ⇒ greatly enhanced by large capacity memory

Large dataset science (2): Genomic data accumulation and analysis

- environmental genome analysis ⇒ specific genome analysis
- large dataset ⇒ distributed processing with query
- large data capacity on computation node enhances performance

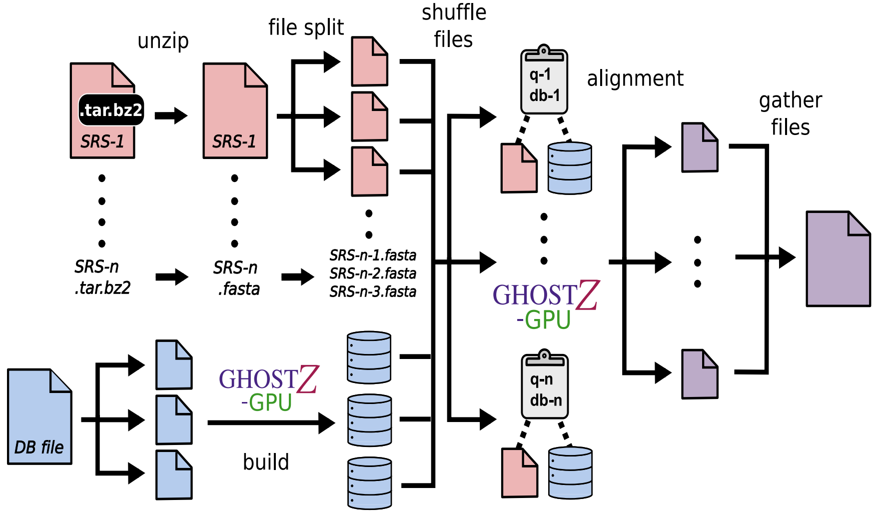

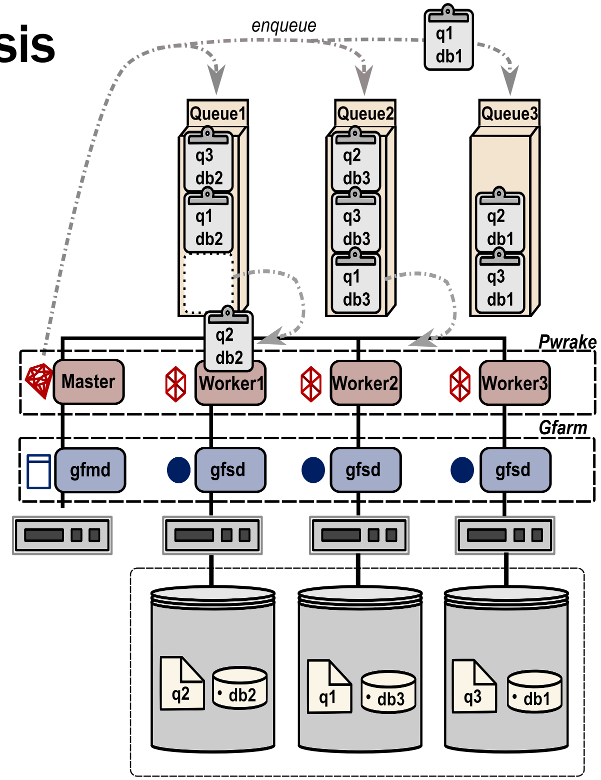

Large dataset science (3): Local climate simulation

Ultra high resolution of local climate simulation by multi-physics

- City-LES: urban climate simulation by LES, solar effect, building structure, surface material

TOKYO2020 model around Tokyo Station, 1m grid

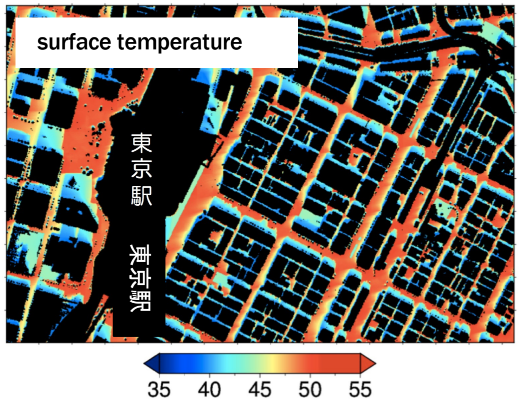

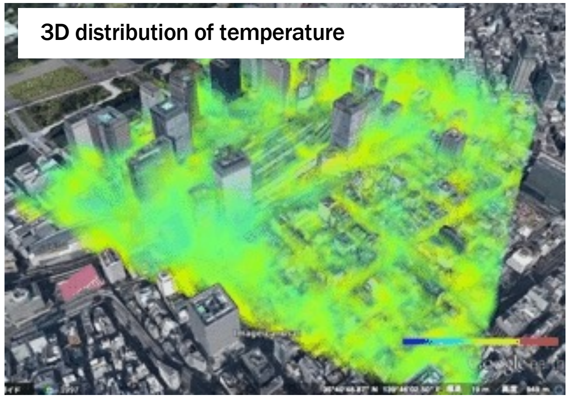

- large scale in-transit analysis on large capacity date
- completely GPUized (up to 15x performance of CPU)

Data Science for human life
Automated Sleep Stage Scoring using In-home EEG Device

In-home EEG Device

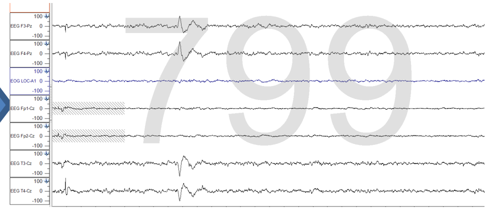

Obtained EEG

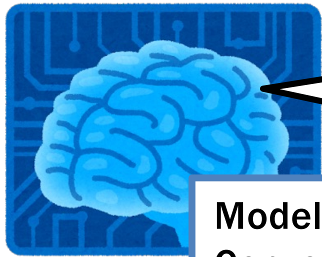

She is in REM sleep

Model based on Convolutional NN

- Training : 700 whole-night records
- Testing : 48 whole-night records & scoring results by sleep experts

Agreement rate with experts : **83.2%**

Kappa Statistics : **0.77**

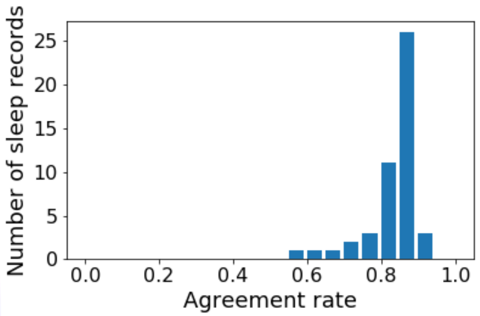

Agreement rate among Experts ➧ **82.0%**

Applicable for Medical purpose

Hypnogram Example（Median Data, Agreement rate: 86.4%）

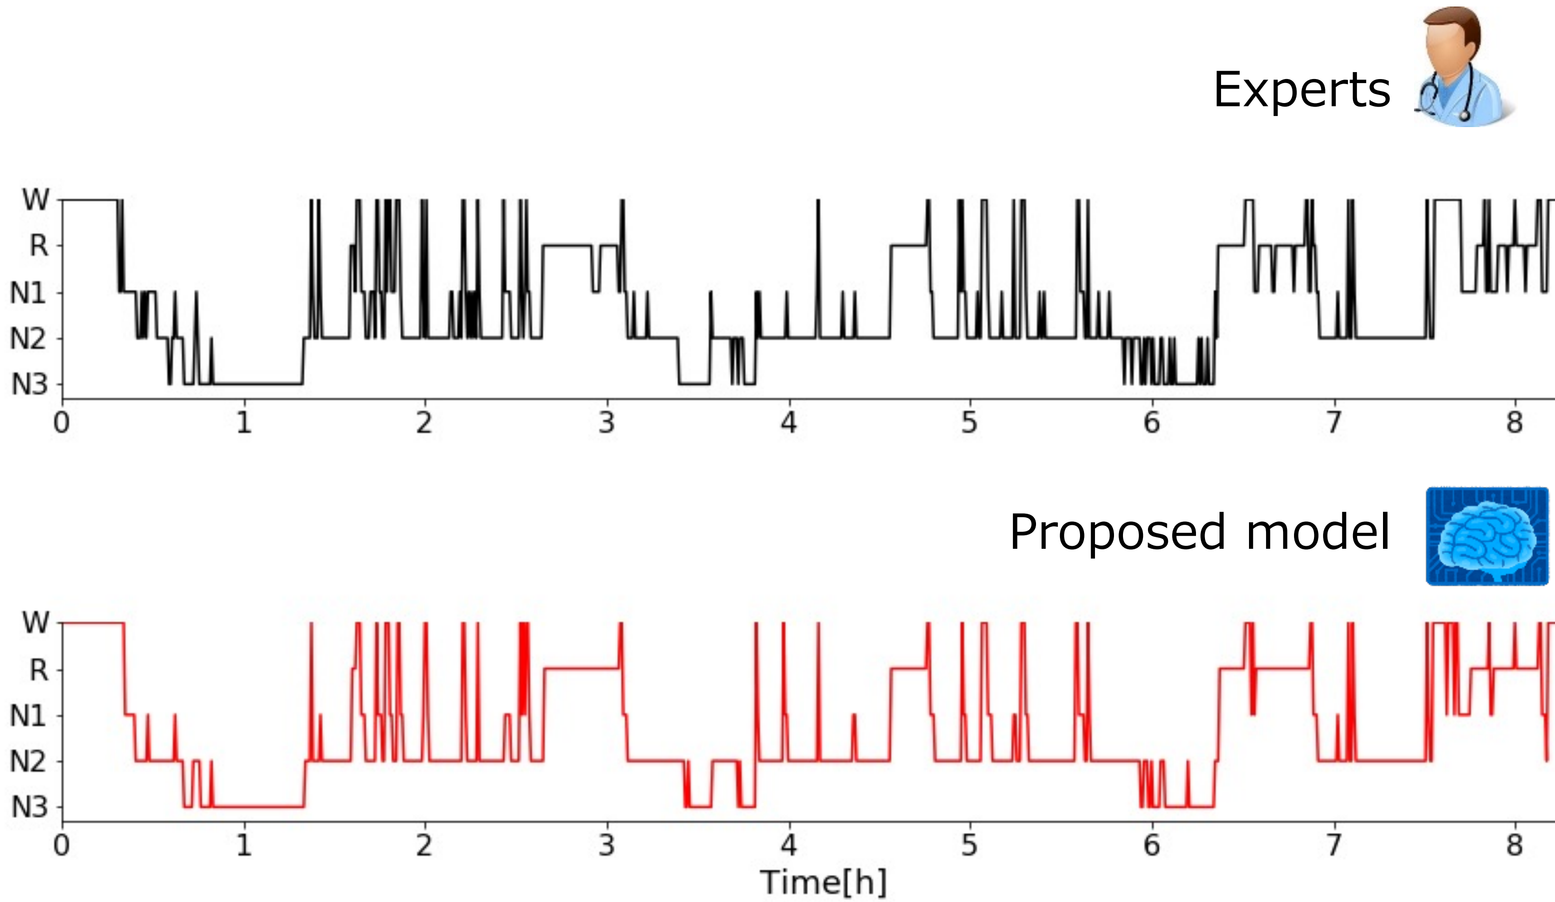

Big Data & AI for HPC in Bio-Science & Drug Discovery

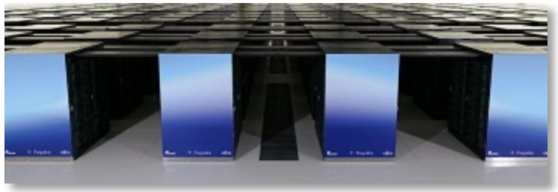

Data Science & Machine Learning

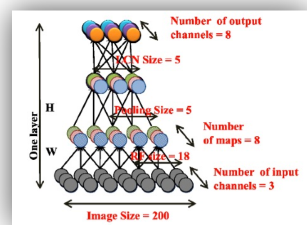

high speed
high accuracy

data
generation

Simulation (MD, docking)

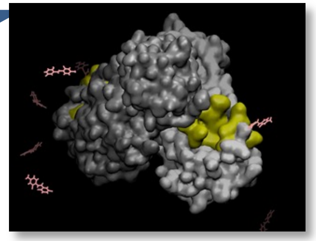

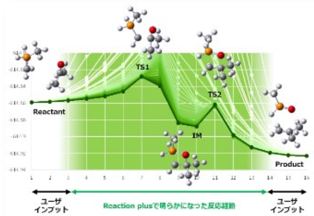

minimum real data by medical research

Big Data (medical experiment and database)

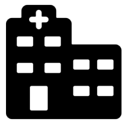

(slide by Yasushi Okuno, Kyoto U.)

Alpha-Fold: Protein structure prediction by AI with

Protein structure prediction by ultra high-performance MD → AI replaces

deep learning

molecular simulation

amino acid array

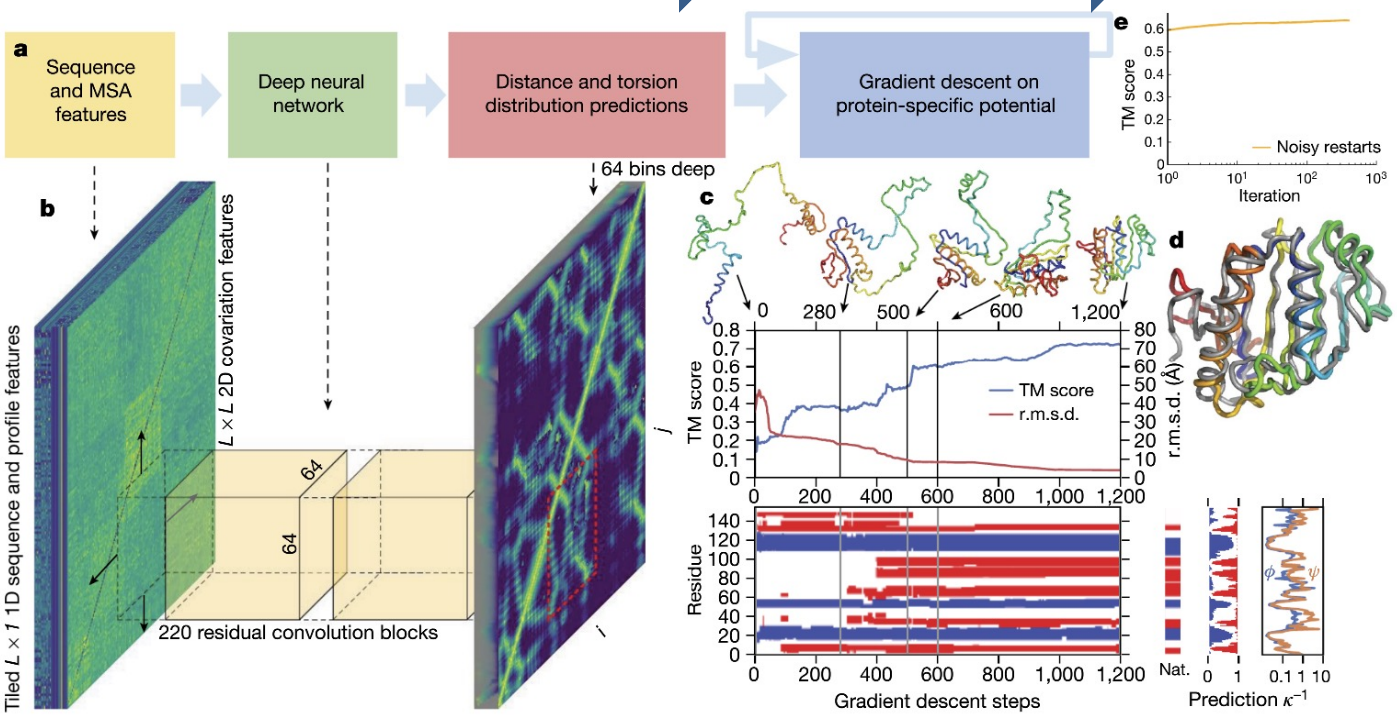

protein structure

(slide by Yasushi Okuno, Kyoto U.)

Backpropagation

How to use Big Memory (PMEM) for HPC/BD/AI ?

- Ultra large capacity of medium~high speed memory
  - astrophysics, climate, life science: requiring large capacity of working set with relatively low FLOPS requirement
- In-Situ processing
  - bypassing the slow in/out of large capacity of file data
    ⇨ definitely speeding-up entire processes
  - so far, memory capacity was a big problem
    ⇨ PMEM technology support
- High-speed distributed storage
  - even as a distributed storage, much faster solution is possible (shown later)

Cygnus (PACS-X): CCS latest supercomputer, with GPU+FPGA

The world first practical supercomputer with Multi-Hybrid (GPU + FPGA) Accelerating Architecture: 320 GPUs (V100) + 64 FPGAs (Stratix10) in 80 nodes

Single node configuration (Albireo)

- All nides in Cygnus are equipped with both IB EDR and FPGA-direct network
- Some nodes (Albireo) are equipped with both FPGAs and GPUs, and other nodes (Deneb) are with GPUs only
- GPU: NVIDIA V100 x4
  FPGA: Intel Stratix10 x2
- InfiniBand - MVAPICH2-ready and officially supported (we have been traditionally using MVAPIC2 multi-rail solution supported by OSU)

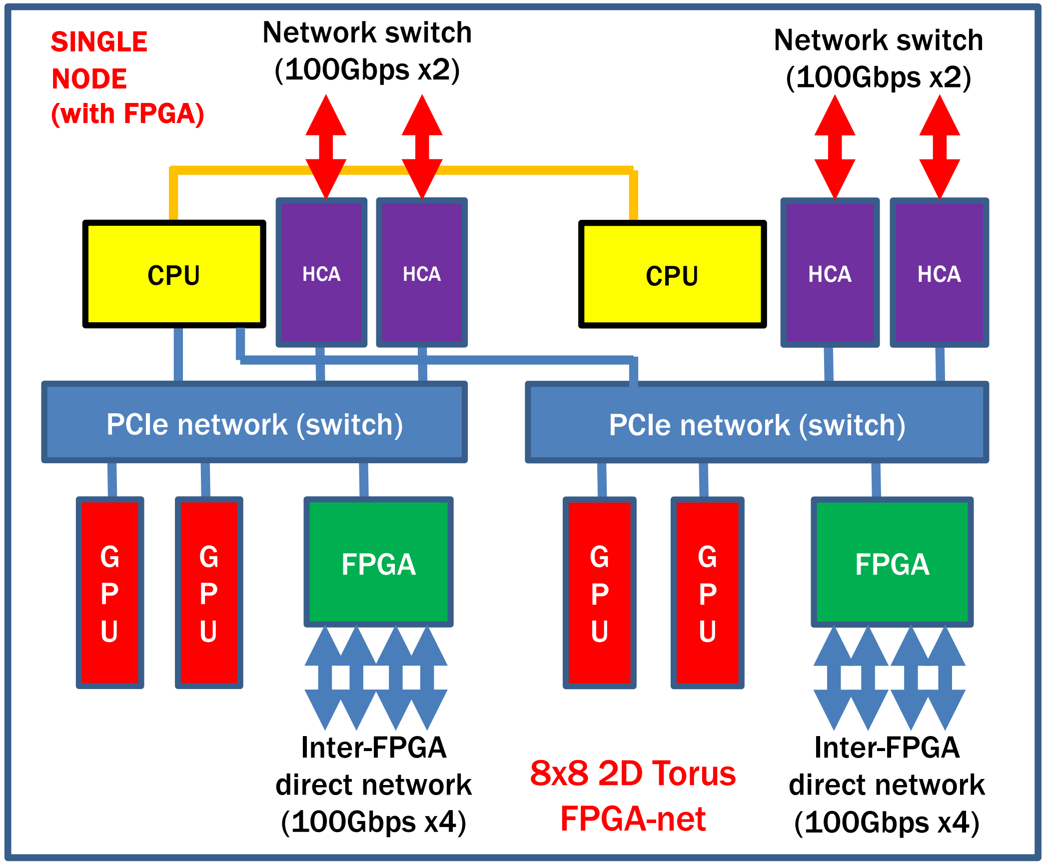

CHARM: Cooperative Heterogeneous Acceleration with Reconfigurable Multi-devices

multi-physics/multi-scale complicated problem

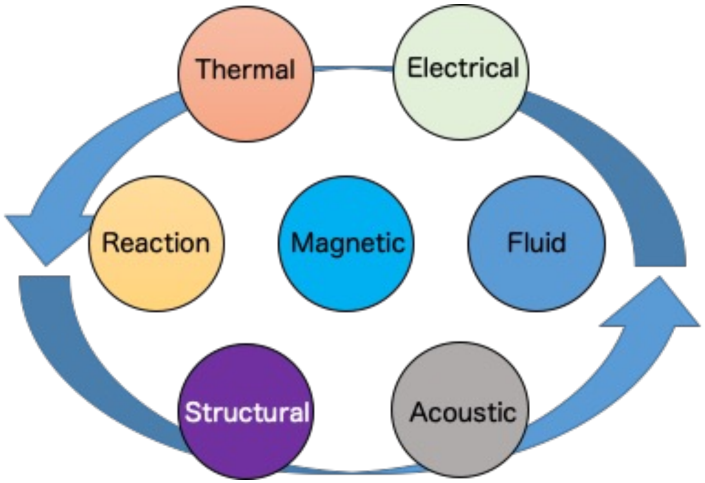

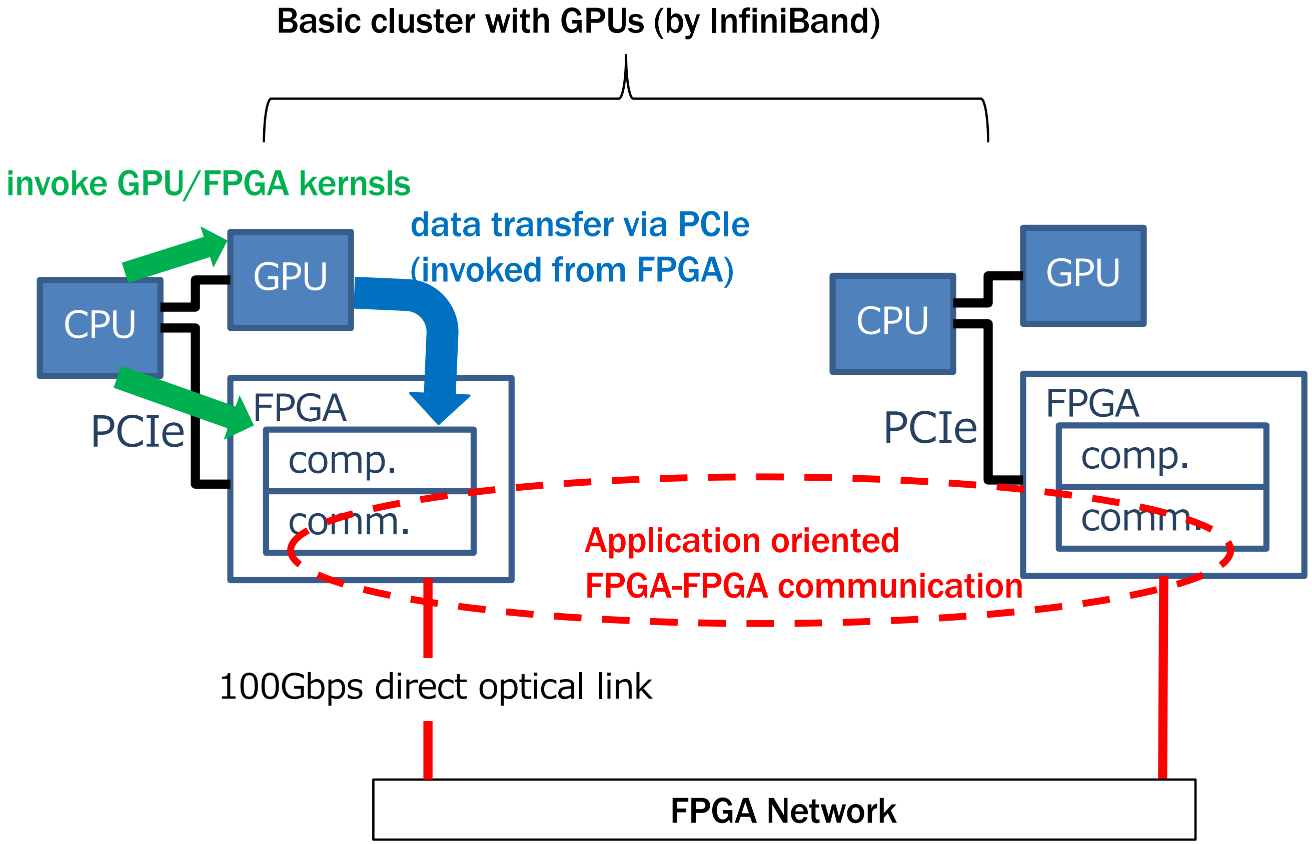

Our next challenge with Big Memory Supercomputer

- Strategy of current Cygnus
  - accelerating traditional HPC, especially for multi-physical simulation with multiple phenomena, by coupling of GPU + FPGA
  - GPU: SIMD-type of spatial parallelism
    FPGA: pipelining x spatial parallelism
  - GPU and FPGA compensate with each other to fill the gap for various parallelism in various applications or even in an application
  - AI (deep learning) is mainly done on GPU, and FPGA partially support (ex. sorting on database search)
- New concept of Big Memory Supercomputer
  - much larger simulation on HPC applications: astrophysics, climate, bioscience
  - much faster distributed file system with large scale cluster computing

Pegasus (called Cygnus-BD so far)

- additional supercomputer for more tight binding of HPC and Data Science
  - large scale memory for ultra large scale HPC simulations to balance GPU&CPU performance and data capacity
  - large capacity of data science (data access & AI) by high speed distributed file system
- Pegasus a.k.a. Cygnus-BD (Cygnus for Big Data)
  - brother system of current Cygnus system
  - introducing persistent memory for both large capacity memory and high speed distributed file system (ad hoc)
  - system development to support these features
- Cygnus & Pegasus (Cygnus-BD)
  - Cygnus ⇨ Extreme Computing
  - Pegasus ⇨ Big Data


Pegasus

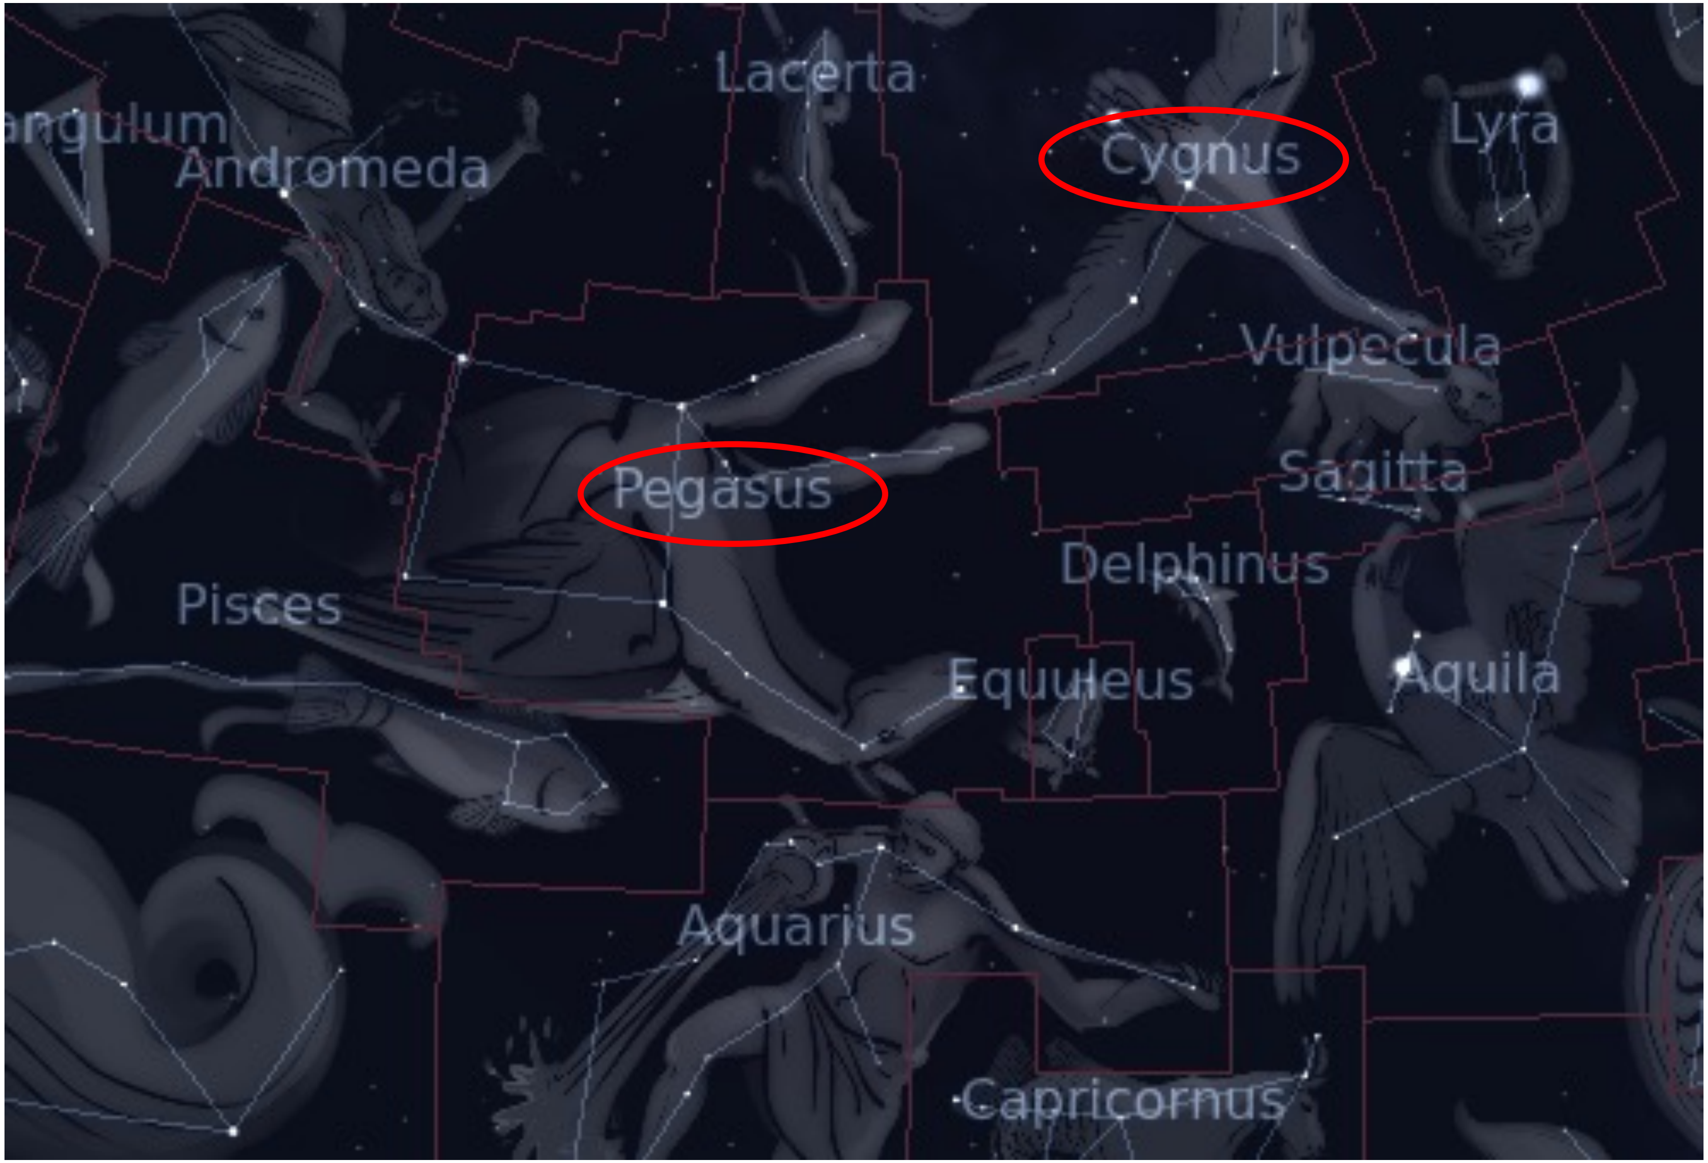
Lacerta
angulum
Andromeda
Cygnus
Lyra
Vulpecula
Pegasus
Sagitta
Delphinus
Pisces
Equuleus
Aquila
Aquarius
Capricornus

Pegasus: world first combination of PMEM + SRP + H100

- will be delivered on Oct. 2022
- total Performance
  - 120 nodes, 6 PFlops, 240 TiB
- node components
  - NVIDIA Tesla H100 PCIe
  - Intel next gen Xeon (code Sapphire Rapids)
  - Intel next gen Optane
- aggregated performance
  - ~0.24 PFlops (CPU), ~5.8 PFLOPS (GPU)
  - 128 GB DRAM, 2 TiB PMEM
  - 6.4 TB NVMe SSD
- Interconnection Network
  - 200 Gbps full bisecion
    (NDR200: NVIDIA Quantum-2 IB)
- parallel file system (DDN)
  - 7.1 PByte, 40 GB/s

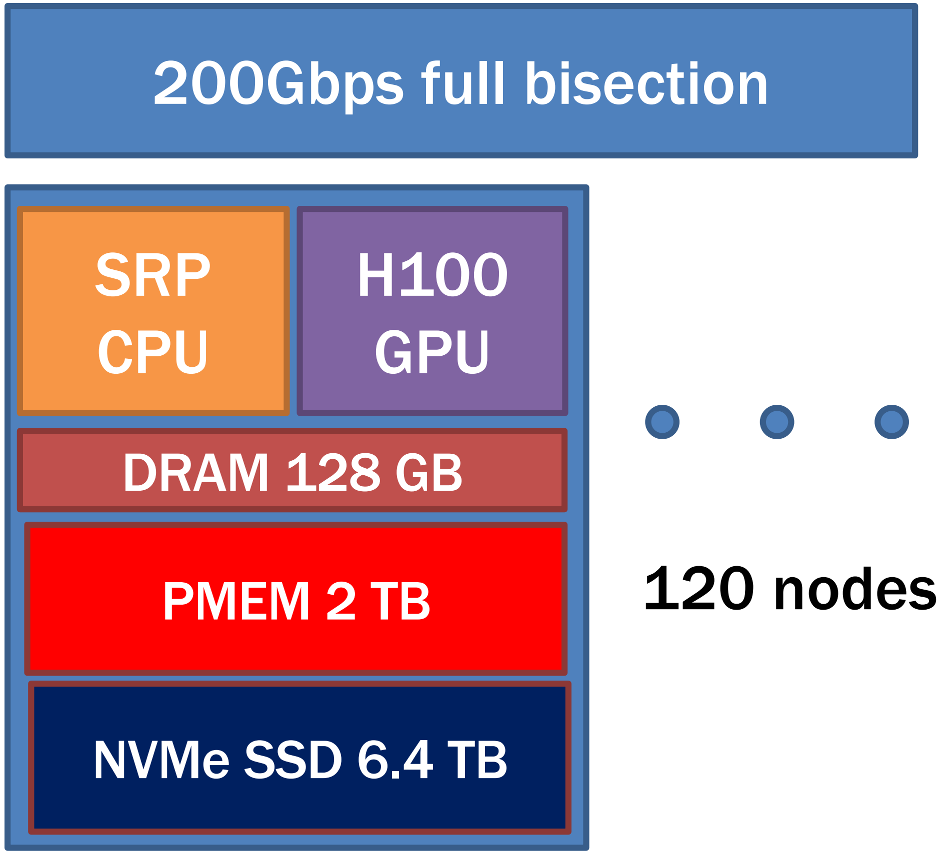

Comparison of Cygnus and Pegasus

	Cygnus (2019)	Pegasus (2022)
PFLOPS (DP)	2.3	6.0 (2.6x)
CPU	0.16	2.4 (1x)
GPU	2.18	5.76 (2.6x)
FPGA (SP)	0.64	0
DRAM (TiB)	10.2	15.4 (1.5x)
PMEM (TiB)	0	240
Storage (PB)	2.4	7.1 (3x)

System integrator: NEC

(subject to change)

Research of ad hoc parallel file system

- Temporal parallel file system using node-local storage
- Fill the performance gap between CPU/GPU and storage

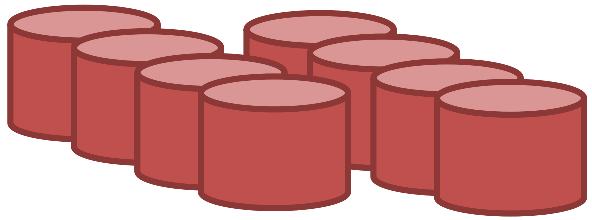

- We are developing CHFS (Consistent Hash File System) ad hoc file system to utilize persistent memory
  - No metadata server, no sequential processing for performance and scalability

* O. Tatebe, et. al., **"CHFS: Parallel Consistent Hashing File System for Node-local Persistent Memory", HPC Asia 2022**

Design goal of CHFS

- Utilizing persistent memory performance
  - In-memory persistent key-value store (not block-based file system)
- Reduce metadata overhead and achieve scalable performance improvement
  - No dedicated metadata server
  - No sequential execution
  - Based on highly parallel distributed key-value store without any central data structure
- Improve single-shared-file performance
  - File is divided into fixed-size chunks to distribute a single file among servers

Design of File System

- All data is stored in a highly parallel distributed KV store
- A single Key-Value format in CHFS for a file chunk and a directory

Key	Value	
Full path and chunk number	Metadata	File data

Metadata (64 bytes)
mode, uid, gid, size chunk size mtime, ctime

- For a directory, there is no chunk number and no file data

  - No file-level information in KV store, such as total chunk numbers and total file size to avoid sequential processing

Design Features

- A file is divided into fixed-size chunks
  - Users define chunk size for each file, because optimal size is different for each file
  - File chunking is critical to access a single file in parallel without lock conflicts
- No dedicated metadata server
  - File chunk and directory can be accessed directly
  - Metadata is included in the key and the value
    - Some metadata is redundant, but the percentage is low when the chunk size is large
  - Design does not include any central data structure
- No file-level information
  - Each chunk is updated perfectly in parallel
  - File size, number of chunks, time stamps are not managed to avoid sequential processing that kills scalability
  - Exact file size requires to count the number of chunks

System Architecture of CHFS

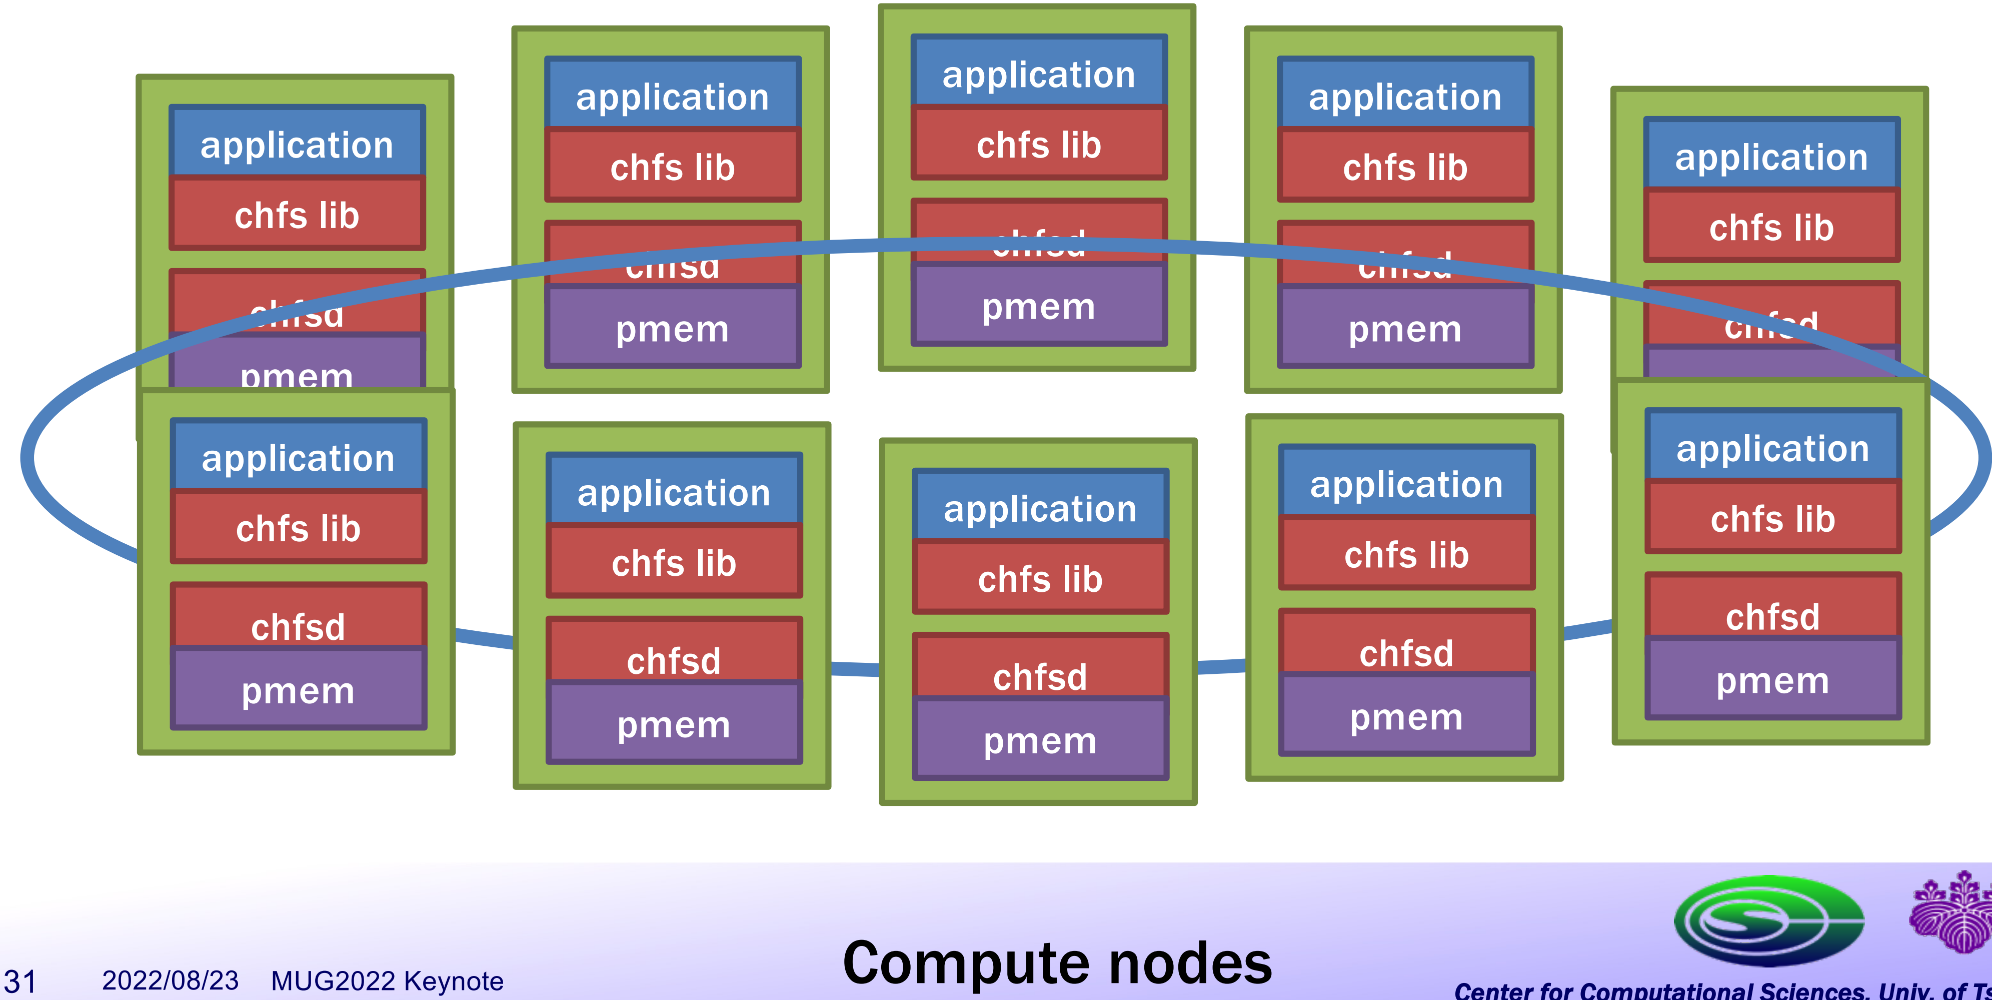

Implementation of CHFS

- Mochi-Margo [JCST 2020]
  - https://mochi.readthedocs.io/en/latest/
  - Communication library using Mercury and Argobots
- Mercury [Cluster 2013]
  - Async RPC, RDMA communication library
  - libfabric, CCI, shared memory plug-ins
- Argobots [IEEE TPDS 2018]
  - Light-weight thread library
- pmemkv
  - cmap – concurrent hash map

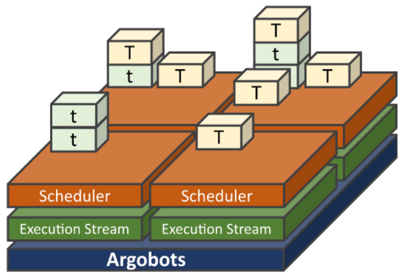

Environment

- 4-node persistent memory cluster (extended part of Cygnus)

InfiniBand HDR 100 (100 Gbps, verbs 11.5 GB/s)

IOR/MDtest	IOR/MDtest	IOR/MDtest	IOR/MDtest
2.3GHz 16core Xeon 96GB DDR4-3200 **768GB Optane PM**	2.3GHz 16core Xeon 96GB DDR4-3200 **768GB Optane PM**	2.3GHz 16core Xeon 96GB DDR4-3200 **768GB Optane PM**	2.3GHz 16core Xeon 96GB DDR4-3200 **768GB Optane PM**

IO500 Bandwidth (CHFS/DAOS/GekkoFS)

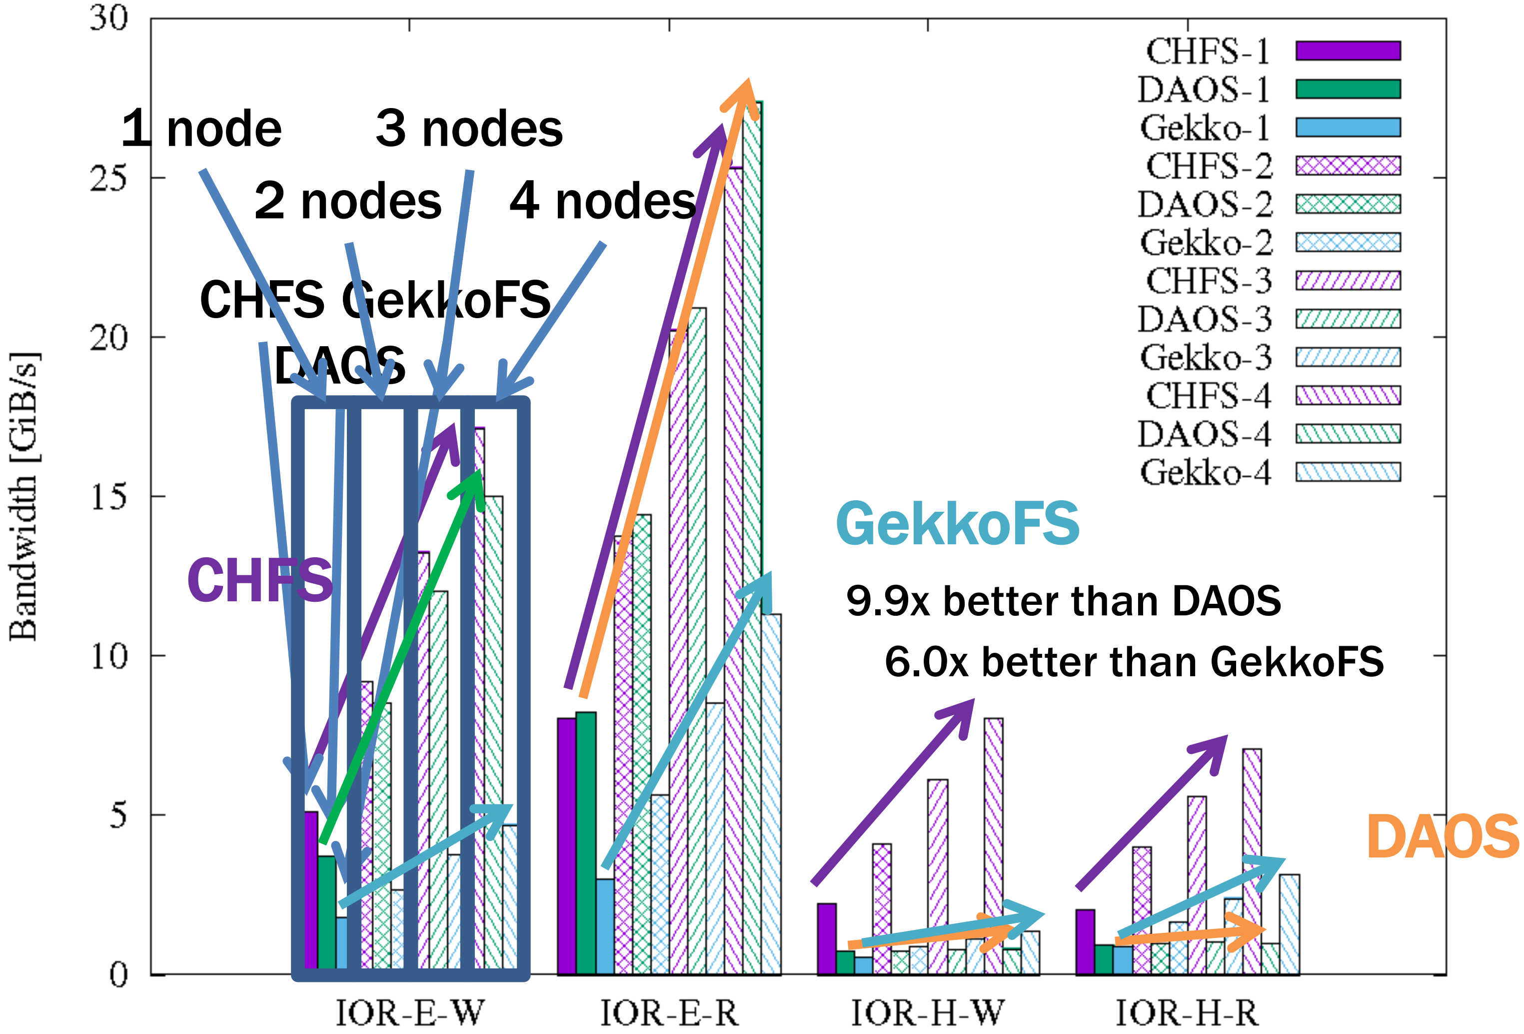

IO500 Metadata Easy (CHFS/DAOS/GekkoFS)

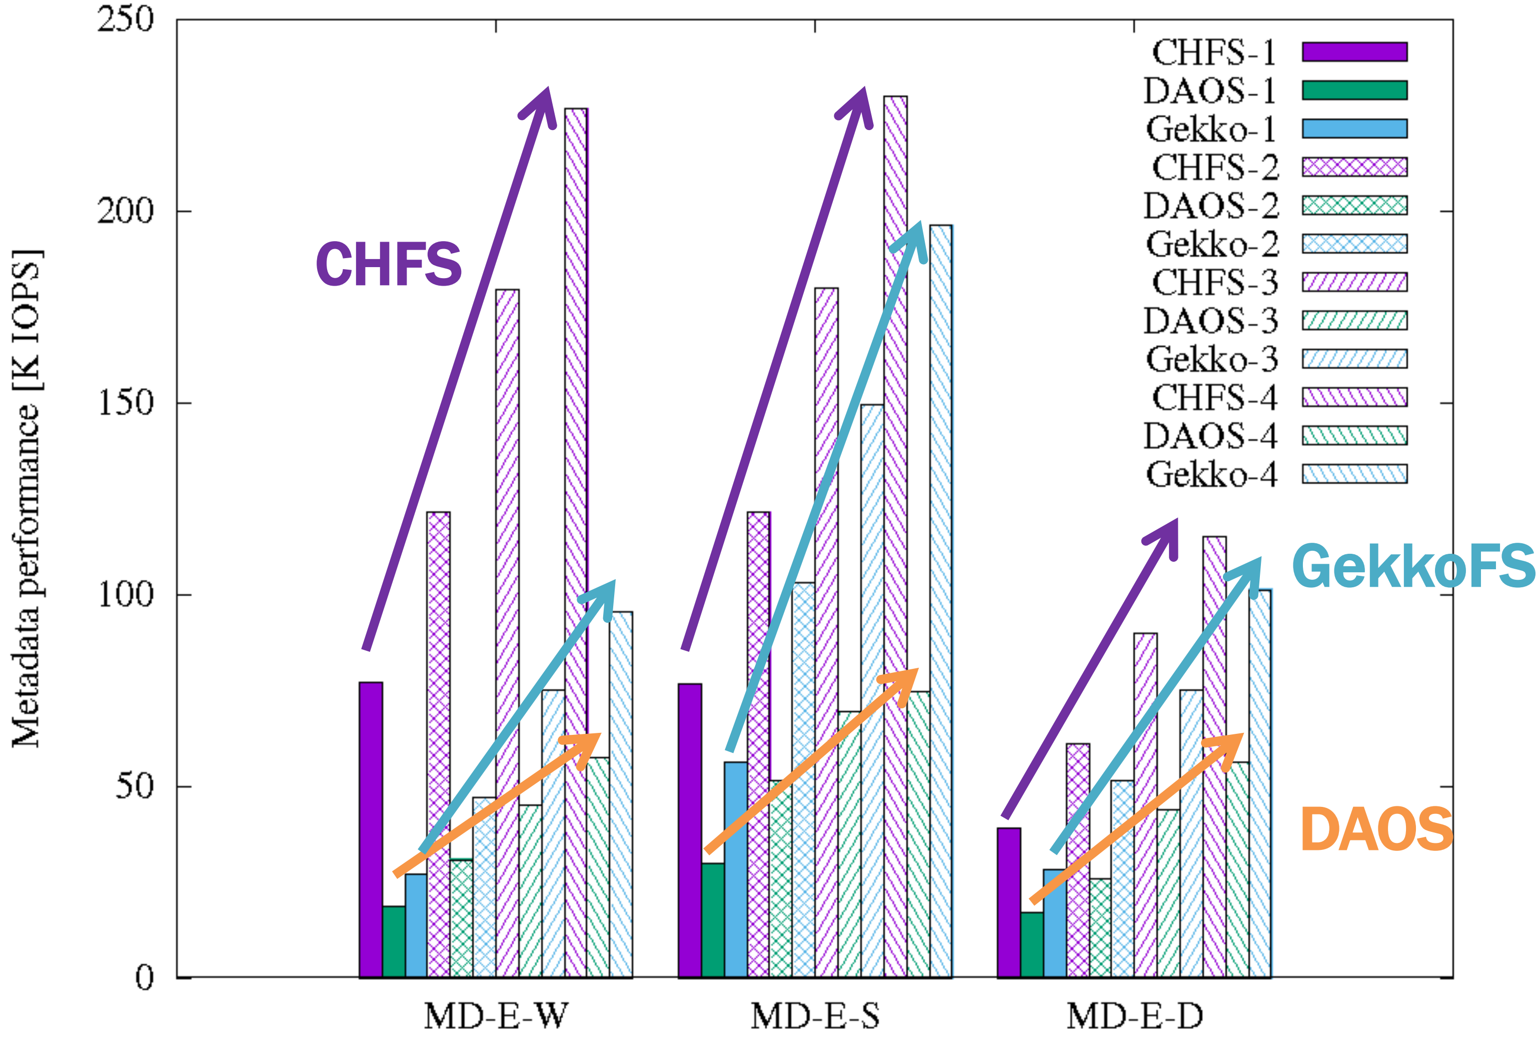

IO500 Metadata Hard (CHFS/DAOS/GekkoFS)

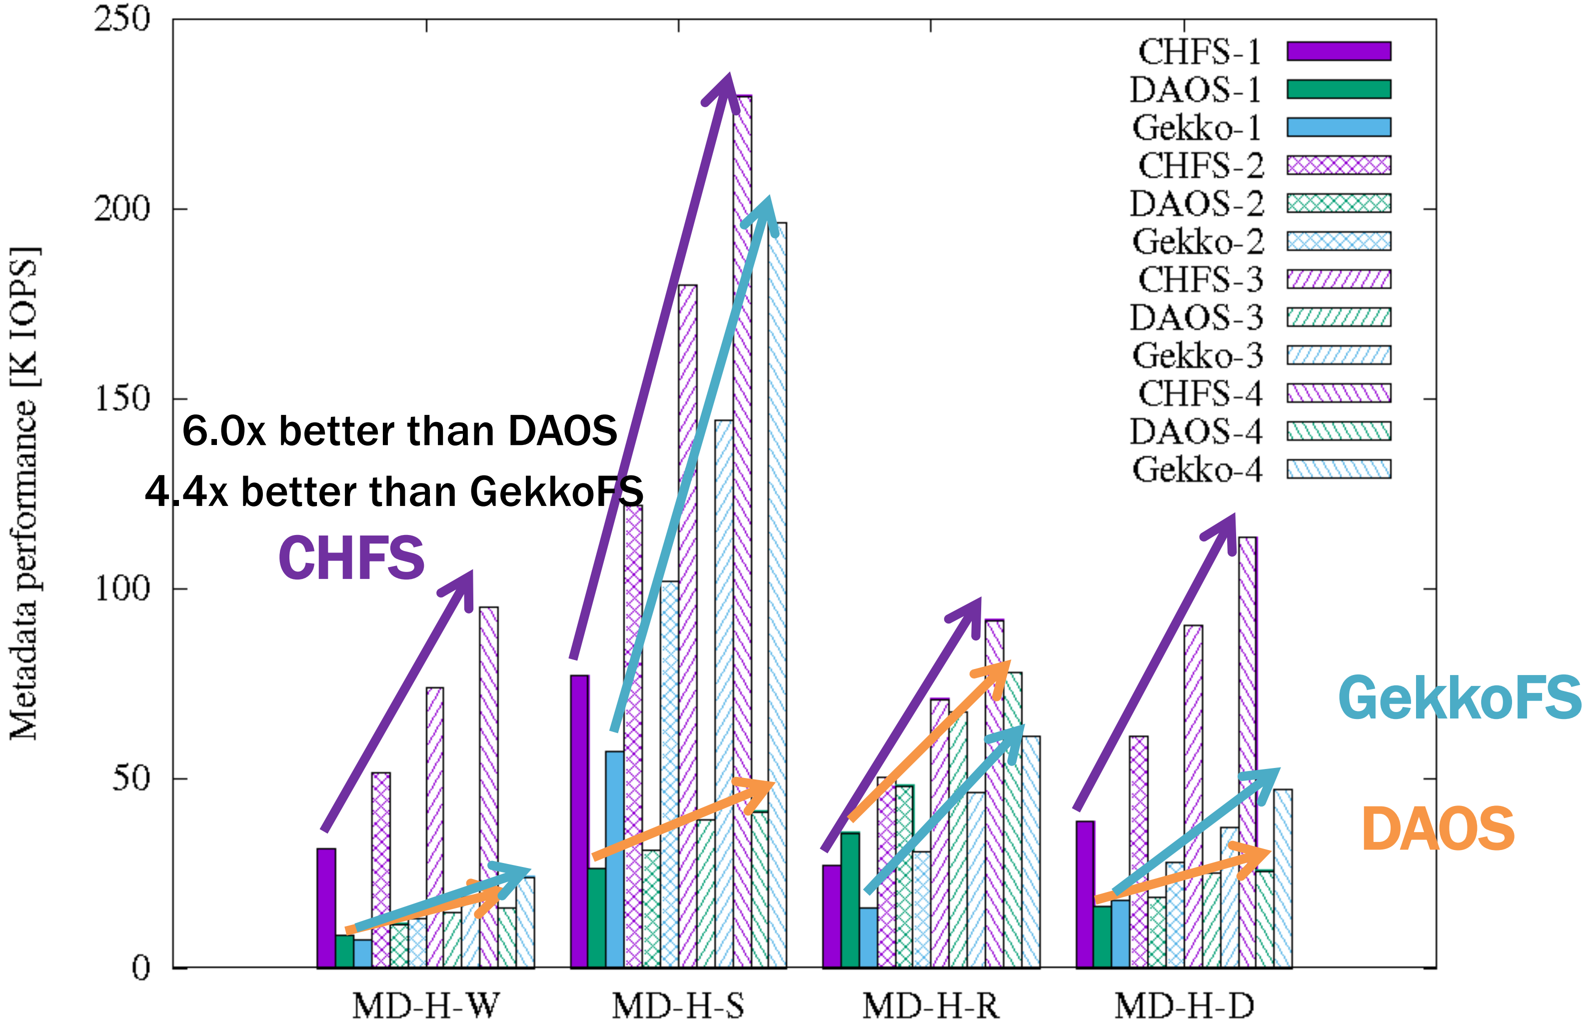

Current status

- Pegasus deployment: Oct. 2022
  - H100 + next gen Xeon (SRP) + next gen Optane (ver3)
  - high bandwidth interconnection: 200Gbps IB
- CHFS development on PMEM
  - NVMe version represents a good performance
    ⇒ enhanced on PMEM by App Direct Mode
- Covering large capacity memory and high performance shared storage
  - coupling Memory Mode and App Direct Mode on PMEM
  - ex: 512MB with Memory Mode + 1.5TB with App Direct Mode for 2TB total
    ⟶ but we expect most of users prefer full-memory or full-AppDirect
    (controllable by dynamic node reconfiguration on batch job script)

Recent issue on "Optane"

- Intel recently announced about "Production shut-down of Optane series"
  - It is true, Intel will stop the production of Optane series modules around 2025
  - No future continuation based on 3DXpoint technology
    ⇨ perhaps switching to CXL based memory devices
- U. Tsukuba and Intel tentatively agreed to the following conditions
  - U. Tsukuba will introduce Pegasus with Optane as on scheduled
  - Intel's announcement does not mean immediate shutdown of Optane product
  - Intel support the current Optane as well as operation of Pegasus for 5 years with full warranty
  - Intel, NEC and Supermicro will support Pegasus project and its operation as usual

MVAPICH2 test on Pegasus

- CCS, U. Tsukuba will provide a special early access to Pegasus for OSU especially for testing MVAPICH2
  - new combination of products: H100 (GDR), SRP, Optane-mapped memory
  - we will do shake-down test just after the deployment, and OSU access will be included
  - it may be the world first large scale MVAPICH2 test with H100+SRP+Optane
- Deployment
  - it may be most likely to the real end of October
  - we will schedule the shake-down test in early November

Summary

- AI is the latest important application and gets growing rapidly as a practical application in human life
- HPC technologies have been contributing to AI so far, and now it's time to use AI technologies for efficient HPC solutions
- Gap between computation performance and memory capacity is so serious today, and one key solution is Persistent Memory
- Utilizing PMEM both for large capacity memory and high performance shared file system (ad hoc) simultaneously, including efficient in-situ processing
- Variation of coverage is so wide to support large scale data science and simulation, based on HPC/BD/AI and GPU/Storage combination