東京大学情報基盤センター
INFORMATION TECHNOLOGY CENTER, THE UNIVERSITY OF TOKYO
東京大学
THE UNIVERSITY OF TOKYO

Innovative Scientific Computing by Integration of (Simulation+Data+Learning)

Wisteria BDEC-01

Hierarchical, Hybrid, Heterogeneous
h3-Open-BDEC
Big Data & Extreme Computing

SC22
Dallas, TX | hpc accelerates.

Kengo Nakajima
Information Technology Center
The University of Tokyo

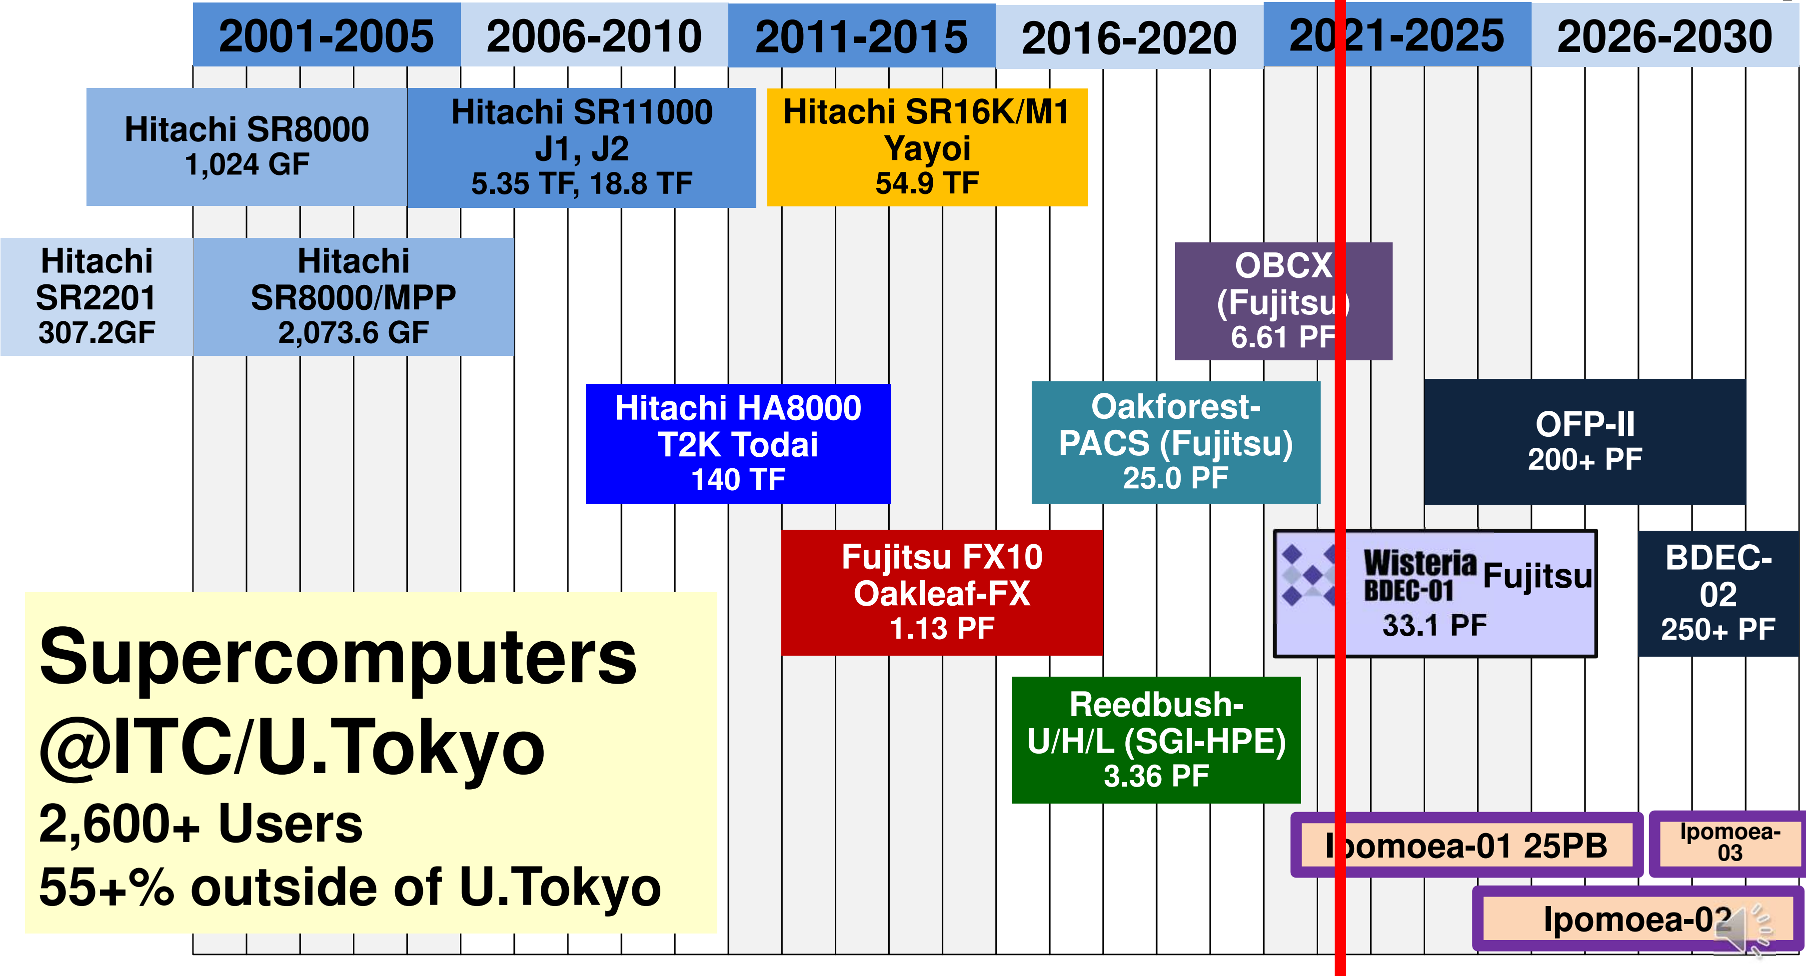
2001-2005
2006-2010
2011-2015
2016-2020
2021-2025
2026-2030
Hitachi SR8000
1,024 GF
Hitachi SR11000
J1, J2
5.35 TF, 18.8 TF
Hitachi SR16K/M1
Yayoi
54.9 TF
Hitachi
SR2201
307.2GF
Hitachi
SR8000/MPP
2,073.6 GF
OBCX
(Fujitsu)
6.61 PF
Hitachi HA8000
T2K Todai
140 TF
Oakforest-
PACS (Fujitsu)
25.0 PF
OFP-II
200+ PF
Fujitsu FX10
Oakleaf-FX
1.13 PF
Wisteria
BDEC-01
Fujitsu
33.1 PF
BDEC-
02
250+ PF
Reedbush-
U/H/L (SGI-HPE)
3.36 PF
Ipomoea-01 25PB
Ipomoea-
03
Ipomoea-02
Supercomputers
@ITC/U.Tokyo
2,600+ Users
55+% outside of U.Tokyo

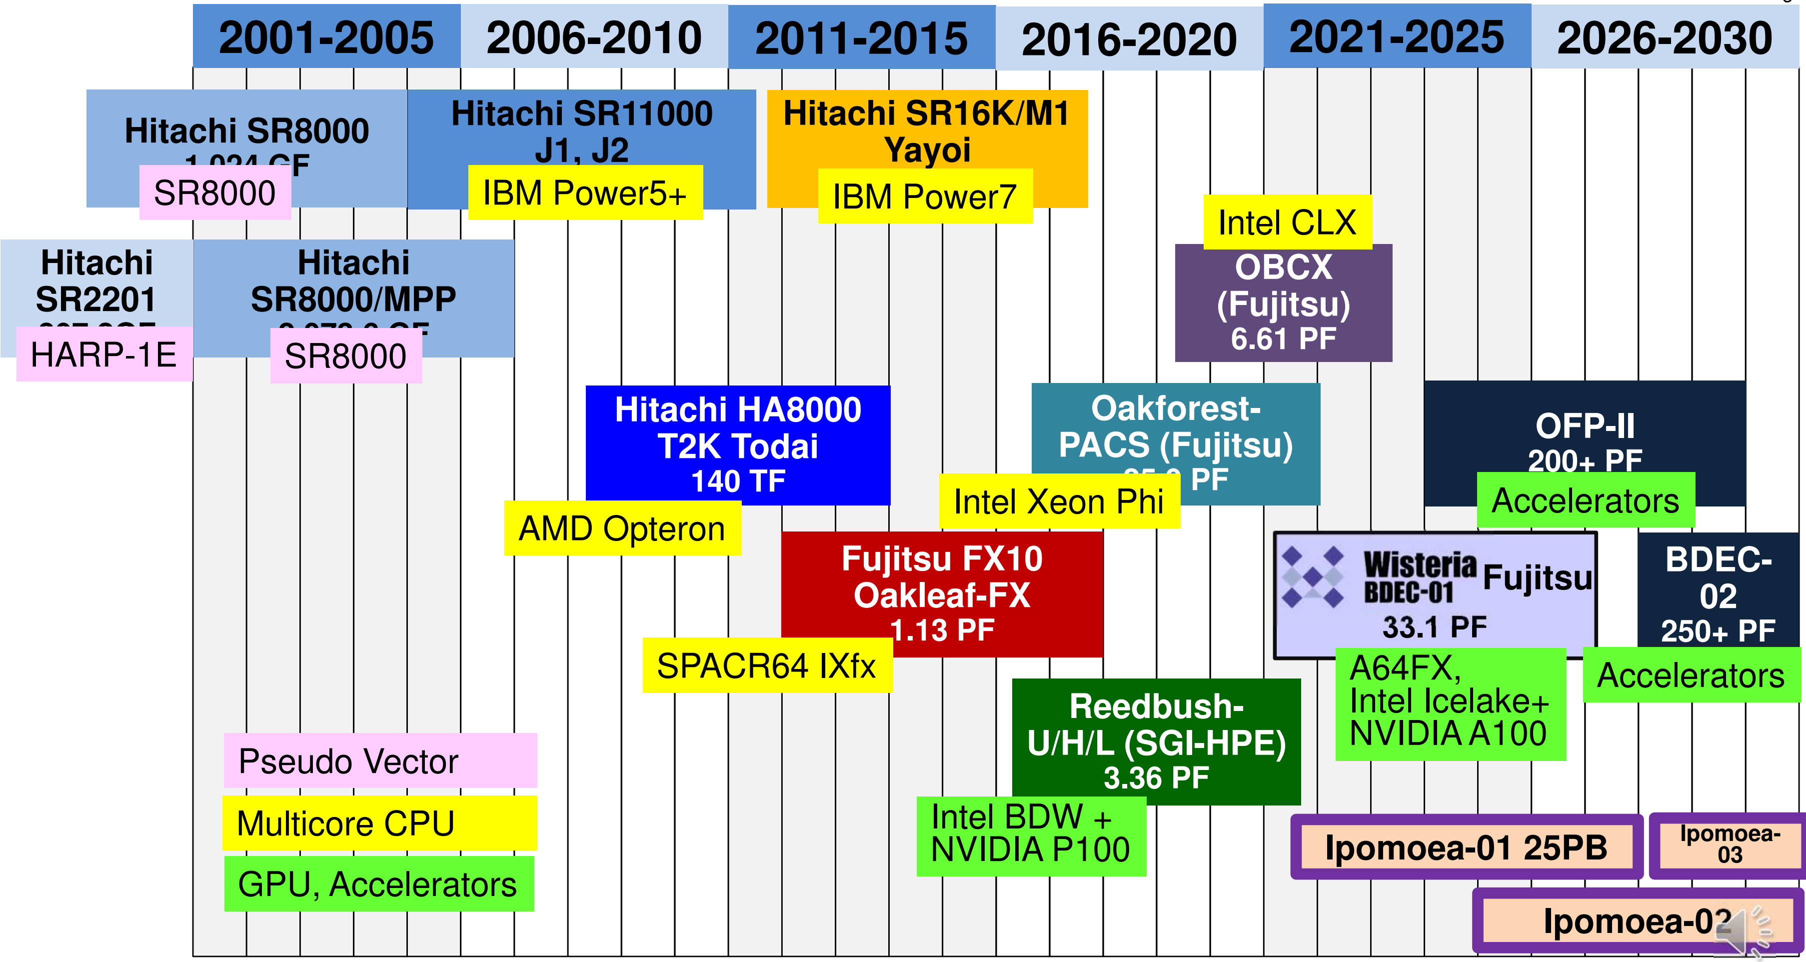
2001-2005
2006-2010
2011-2015
2016-2020
2021-2025
2026-2030
Hitachi SR8000
SR8000
Hitachi SR11000
J1, J2
IBM Power5+
Hitachi SR16K/M1
Yayoi
IBM Power7
Intel CLX
OBCX
(Fujitsu)
6.61 PF
Hitachi
SR2201
HARP-1E
Hitachi
SR8000/MPP
SR8000
Hitachi HA8000
T2K Todai
140 TF
AMD Opteron
Oakforest-
PACS (Fujitsu)
PF
Intel Xeon Phi
OFP-II
200+ PF
Accelerators
Fujitsu FX10
Oakleaf-FX
1.13 PF
SPACR64 IXfx
Wisteria
BDEC-01
Fujitsu
33.1 PF
A64FX,
Intel Icelake+
NVIDIA A100
BDEC-
02
250+ PF
Accelerators
Reedbush-
U/H/L (SGI-HPE)
3.36 PF
Intel BDW +
NVIDIA P100
Pseudo Vector
Multicore CPU
GPU, Accelerators
Ipomoea-01 25PB
Ipomoea-
03
Ipomoea-02

Future of Supercomputing

- Various Types of Workloads
  - Computational Science & Engineering: Simulations
  - Big Data Analytics
  - AI, Machine Learning …

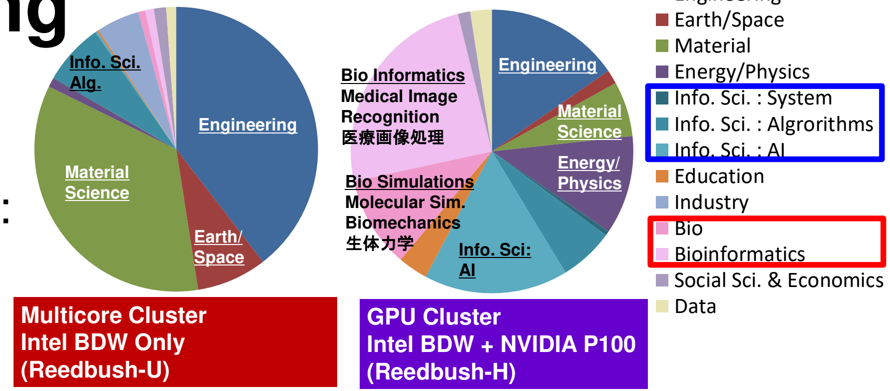

Future of Supercomputing

- Various Types of Workloads
  - Computational Science & Engineering: Simulations
  - Big Data Analytics
  - AI, Machine Learning …
- Integration/Convergence of (Simulation + Data + Learning) (S+D+L) is important towards Society 5.0
  - Super Smart & Human-centered Society by Digital Innovation (IoT, Big Data, AI etc.) and by Integration of Cyber Space & Physical Space

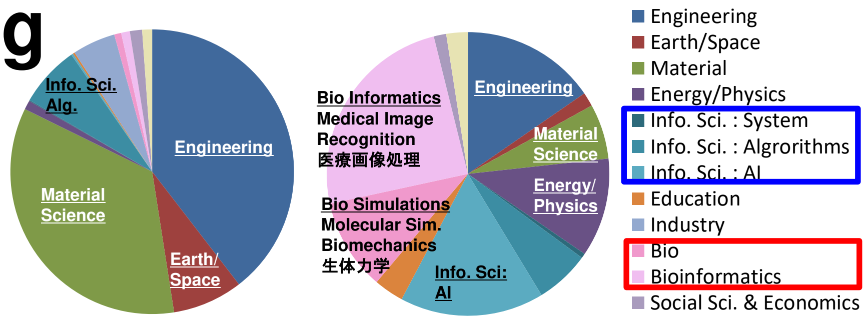

Multicore Cluster
Intel BDW Only
(Reedbush-U)

GPU Cluster
Intel BDW + NVIDIA P100
(Reedbush-H)

Future of Supercomputing

- Various Types of Workloads
  - Computational Science & Engineering: Simulations
  - Big Data Analytics
  - AI, Machine Learning …
- **Integration/Convergence of (Simulation + Data + Learning) (S+D+L) is important towards Society 5.0**

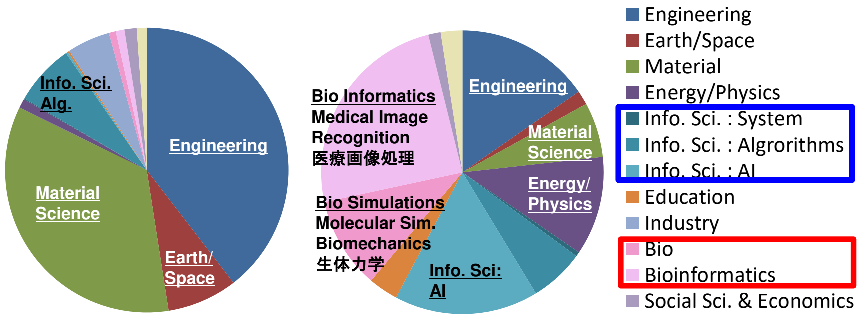

Multicore Cluster
Intel BDW Only
(Reedbush-U)

GPU Cluster
Intel BDW + NVIDIA P100
(Reedbush-H)

- **BDEC (Big Data & Extreme Computing)**
  - **Platform for Integration of (S+D+L)**
  - **Focusing on S (Simulation)**
    - **AI for HPC, AI for Science, Digital Twins**
  - **Planning started in 2015**

BDEC (Big Data & Extreme Computing)

S + D + L

Wisteria/BDEC-01

The 1st BDEC System (Big Data & Extreme Computing) Platform for Integration of (S+D+L)

- Operation starts on May 14, 2021
- 33.1 PF, 8.38 PB/sec by **Fujitsu**
  - ~4.5 MVA with Cooling, ~360m²
- 2 Types of Node Groups
  - Hierarchical, Hybrid, Heterogeneous (h3)
  - Simulation Nodes: Odyssey
    - Fujitsu PRIMEHPC FX1000 (A64FX), 25.9 PF
      - 7,680 nodes (368,640 cores), Tofu-D
      - General Purpose CPU + HBM
      - Commercial Version of "Fugaku"
  - Data/Learning Nodes: Aquarius
    - Data Analytics & AI/Machine Learning
    - Intel Xeon Ice Lake + NVIDIA A100, 7.2PF
      - 45 nodes (90x Ice Lake, 360x A100), IB-HDR
    - Some of the DL nodes are connected to external resources directly
- File Systems: SFS (Shared/Large) + FFS (Fast/Small)

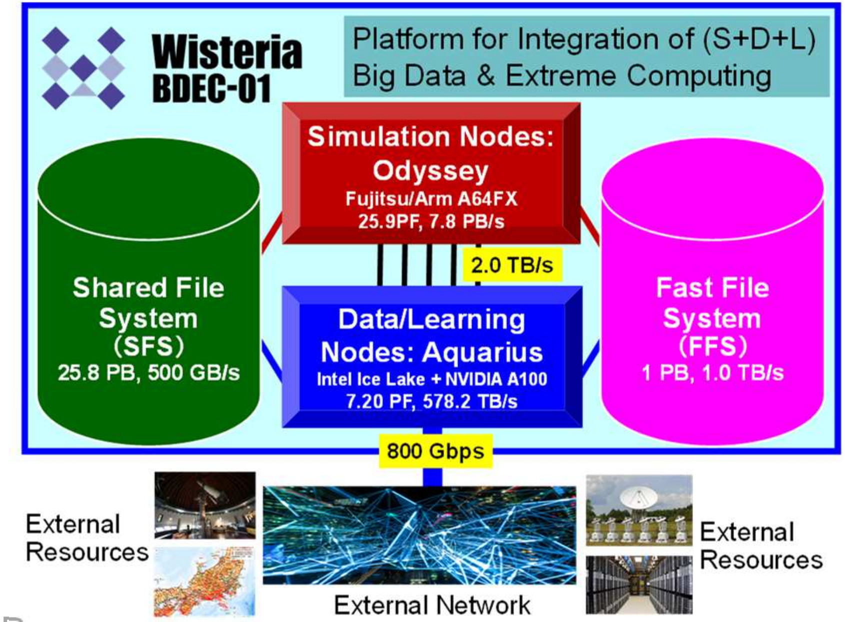

Wisteria/BDEC-01

The 1st BDEC System (Big Data & Extreme Computing) Platform for Integration of (S+D+L)

- Operation starts on May 14, 2021
- 33.1 PF, 8.38 PB/sec by **Fujitsu**
  - ~4.5 MVA with Cooling, ~360m^2
- **2 Types of Node Groups**
  - Hierarchical, Hybrid, Heterogeneous (h3)
  - **Simulation Nodes: Odyssey**
    - **Fujitsu PRIMEHPC FX1000 (A64FX), 25.9 PF**
      - **7,680 nodes (368,640 cores), Tofu-D**
      - **General Purpose CPU + HBM**
      - **Commercial Version of "Fugaku"**
  - Data/Learning Nodes: Aquarius
    - Data Analytics & AI/Machine Learning
    - Intel Xeon Ice Lake + NVIDIA A100, 7.2PF
      - 45 nodes (90x Ice Lake, 360x A100), IB-HDR
    - Some of the DL nodes are connected to external resources directly
- File Systems: SFS (Shared/Large) + FFS (Fast/Small)

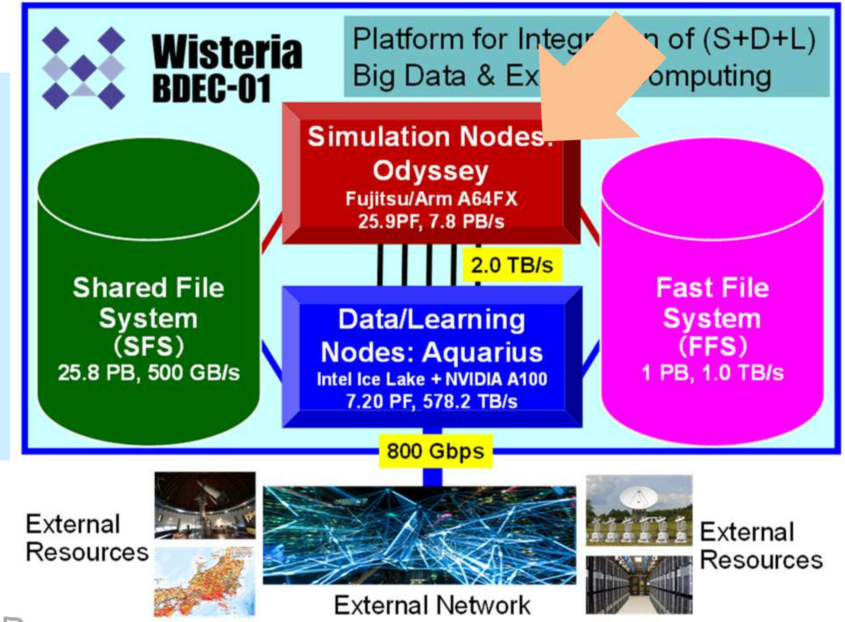

Wisteria/BDEC-01

The 1st BDEC System (Big Data & Extreme Computing)
Platform for Integration of (S+D+L)

- Operation starts on May 14, 2021
- 33.1 PF, 8.38 PB/sec by **Fujitsu**
  - ~4.5 MVA with Cooling, ~360m^2
- **2 Types of Node Groups**
  - Hierarchical, Hybrid, Heterogeneous (h3)
  - **Simulation Nodes: Odyssey**
    - **Fujitsu PRIMEHPC FX1000 (A64FX), 25.9 PF**
      - **7,680 nodes (368,640 cores), Tofu-D**
      - **General Purpose CPU + HBM**
      - **Commercial Version of "Fugaku"**
  - **Data/Learning Nodes: Aquarius**
    - **Data Analytics & AI/Machine Learning**
    - **Intel Xeon Ice Lake + NVIDIA A100, 7.2PF**
      - **45 nodes (90x Ice Lake, 360x A100), IB-HDR**
    - **Some of the DL nodes are connected to external resources directly**
- File Systems: SFS (Shared/Large) + FFS (Fast/Small)

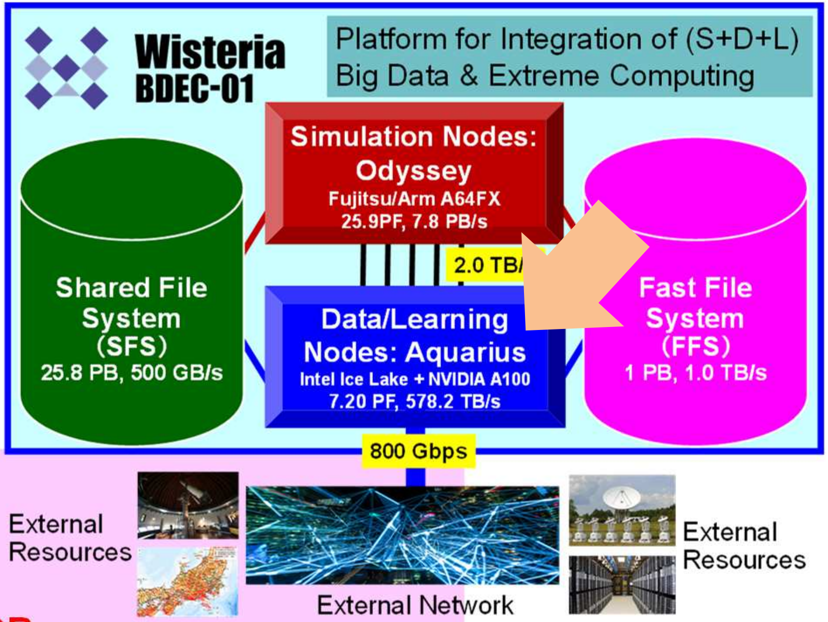

Rankings@ISC 2022
June 2022

ISC High Performance
The HPC Event.

ISC HIGH PERFORMANCE 2021 DIGITAL
JUNE 24 - JULY 2, 2021
ISC-HPC.COM

	Odyssey	Aquarius
TOP 500	20	115
Green 500	34	21
HPCG	10	62
Graph 500 BFS	3	-
HPL-AI	10	-

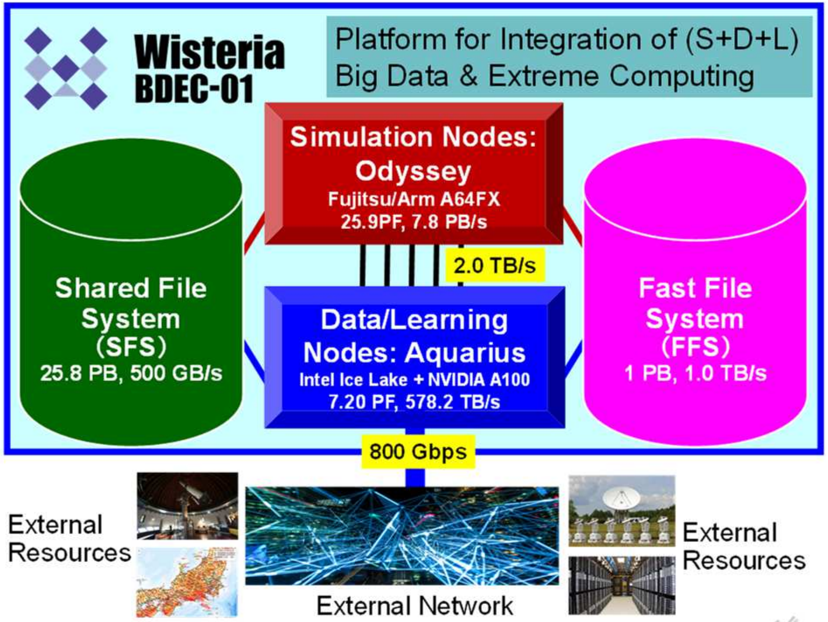

Wisteria/BDEC-01

Platform for Integration of (Simulation+Data+Learning) (S+D+L)

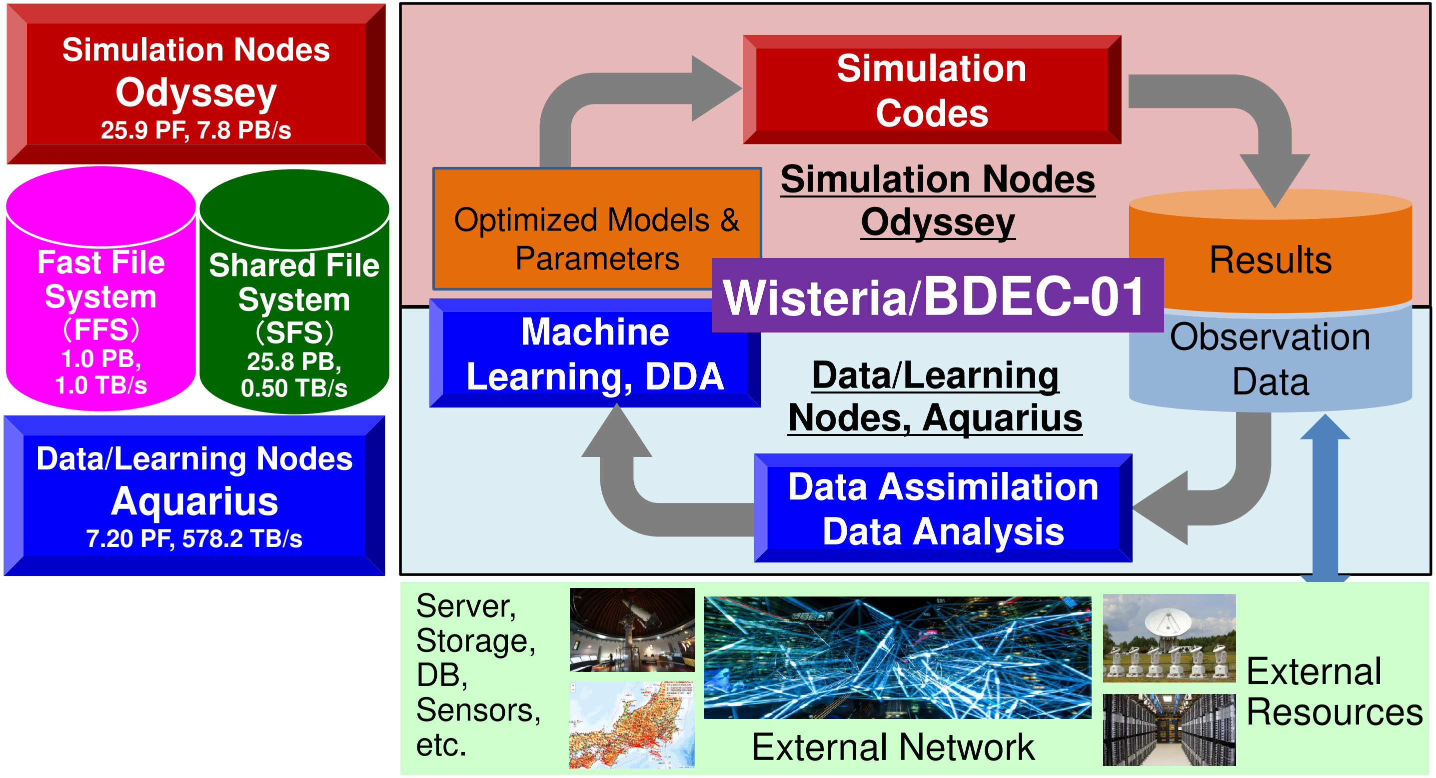

Wisteria/BDEC-01

Platform for Integration of (Simulation+Data+Learning) (S+D+L)

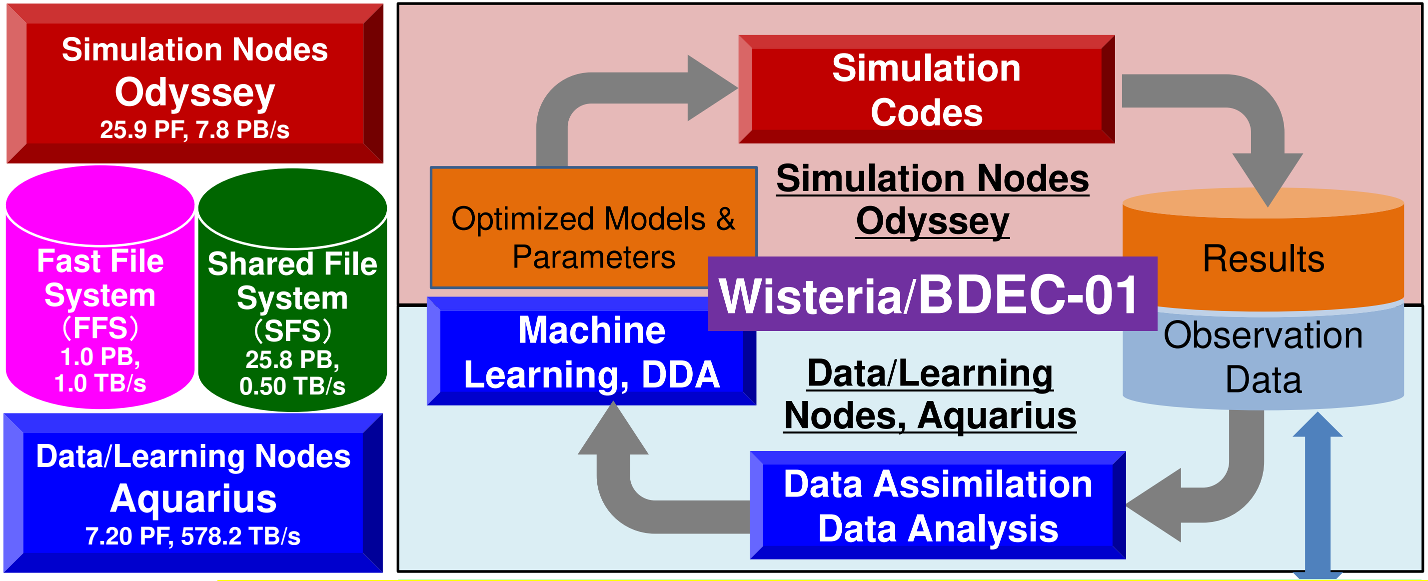

Optimization of Models/Parameters for Simulations by Data Analytics & Machine Learning (S+D+L)

h3-Open-BDEC: Innovative Software Platform for Integration of (S+D+L) on the BDEC System, such as Wisteria/BDEC-01

- 5-year project supported by Japanese Government (JSPS) since 2019
- Leading-PI: Kengo Nakajima (The University of Tokyo)
- Total Budget: 1.41M USD

Hierarchical, Hybrid, Heterogeneous
h3-Open-BDEC
Big Data & Extreme Computing

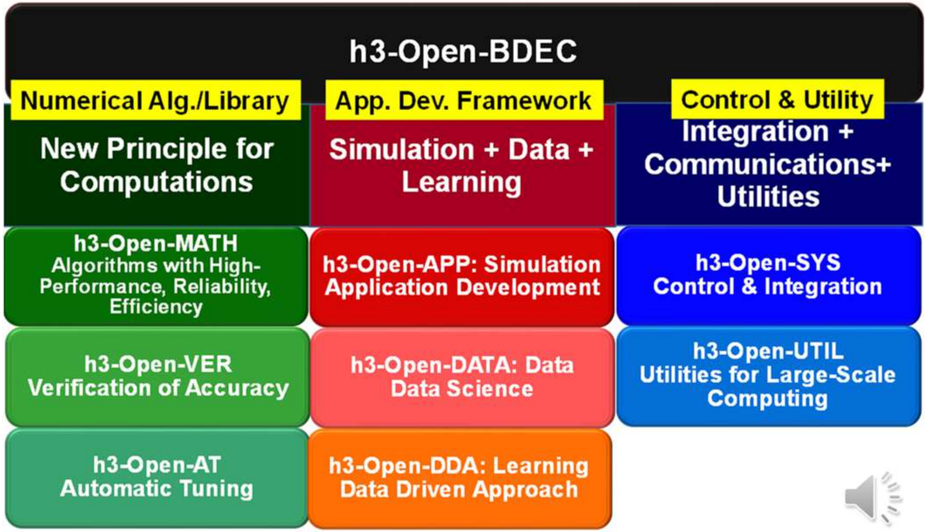

東京大学 THE UNIVERSITY OF TOKYO

RIKEN R-CCS

国立研究開発法人 国立環境研究所 National Institute for Environmental Studies

北海道大学 HOKKAIDO UNIVERSITY

Tokyo Woman's Christian University 東京女子大学

名古屋大学 NAGOYA UNIVERSITY

h3-Open-BDEC: Innovative Software Platform for Integration of (S+D+L) on the BDEC System, such as Wisteria/BDEC-01

- "Three" Innovations
  - New Principles for Numerical Analysis by Adaptive Precision, Automatic Tuning & Accuracy Verification
  - Hierarchical Data Driven Approach (*h*DDA) based on Machine Learning
  - Software & Utilities for Heterogenous Environment, such as Wisteria/BDEC-01

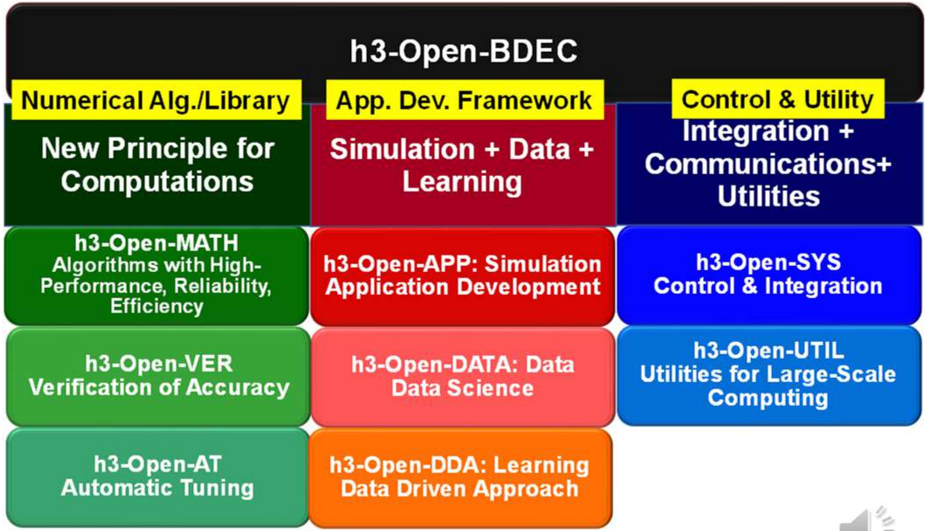

Adaptive Precision Computing with FP42/FP21

Masatoshi Kawai (kawai@cc.u-tokyo.ac.jp)

Heat Conduction with Heterogeneous Material Property

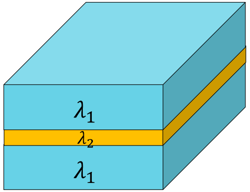

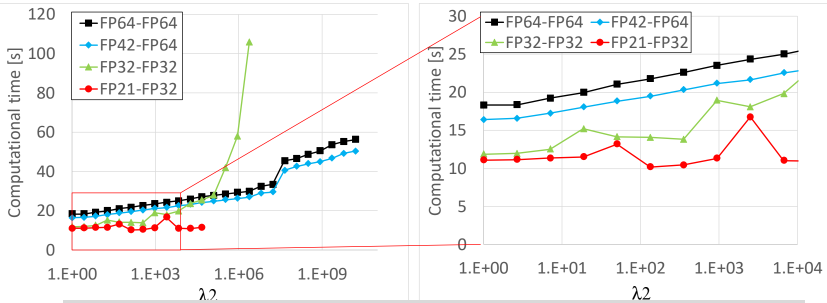

Computation Time for ICCG Solver
Various Types of Precisions on Intel Xeon Cascadelake

In recent years, the usefulness of low-precision floating-point representation has been studied in various fields such as machine learning. Low accuracy can be expected to have effects such as shortening calculation time and reducing power consumption. For example, in an application with a memory bandwidth bottleneck, the effect of reducing the calculation time by reducing the amount of memory transfer is significant. However, in fields such as iterative methods, it is common to use FP64 because the calculation accuracy strongly affects the convergence, and there are few application examples of low-precision arithmetic. This study investigates the applicability of low-precision representation to the Krylov subspace and stationary iterative methods. In this research, we focus on the FP32, FP16, and FP42, FP21, which are not standardized by IEEE754. Developed method has been evaluated for ICCG solver, which solves linear equations derived from 3D FVM code for steady-state head conduction with heterogeneous material property ($\lambda_1=10^0$, $\lambda_2=10^0\sim10^9$). Generally, computation with lower precision (e.g. FP32-FP32, FP21-FP32) becomes unstable, if condition number of the coefficient matrix is larger (λ_2 is larger), FP21-FP32 provides the best performance if λ_2 is up to 10^4. ("FP21-FP32" means "matrices are in FP21, and vectors are in FP32)

Prediction of CFD Simulation by Deep Learning

Takashi Shimokawabe (shimokawabe@cc.u-tokyo.ac.jp)

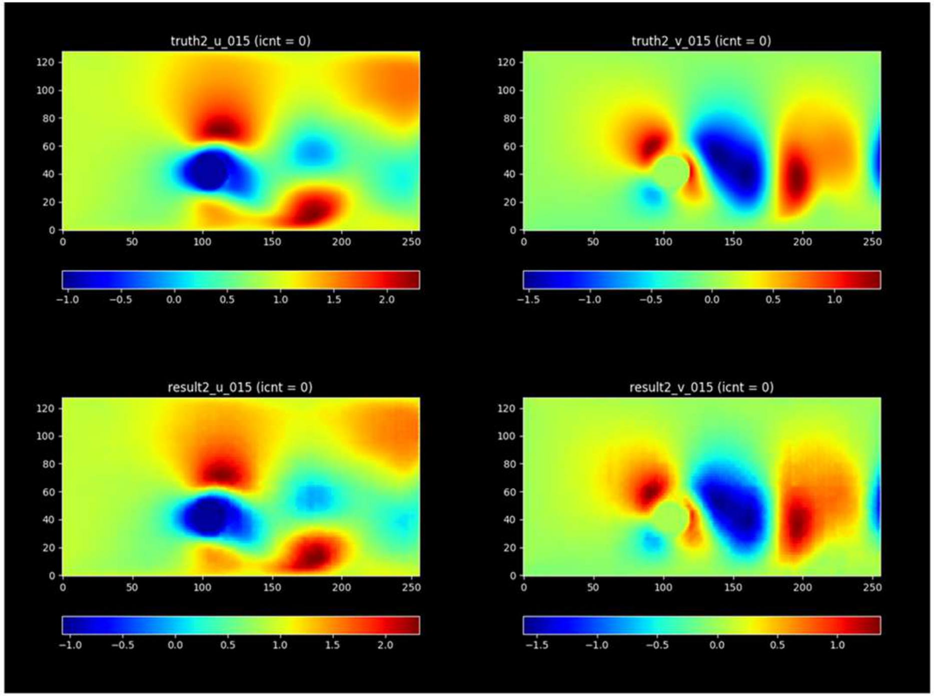

Comparison of the flow velocity results obtained by the conventional simulation (upper) and the prediction of these results by deep learning (lower)

Computational fluid dynamics (CFD) is widely used in science and engineering. However, since CFD simulations requires a large number of grid points and particles for these calculations, these kinds of simulations demand a large amount of computational resources such as supercomputers. Recently, deep learning has attracted attention as a surrogate method for obtaining calculation results by CFD simulation approximately at high speed. We are working on a project to develop a parallelization method to make it possible to apply the surrogate method based on the deep learning to large scale geometry. Unlike the model parallel computing, the method we are currently developing predicts large-scale steady flow simulation results by dividing the input geometry into multiple parts and applying a single small neural network to each part in parallel. This method is developed based on considering the characteristics of CFD simulation and the consistency of the boundary condition of each divided subdomain. By using the physical values on the adjacent subdomains as boundary conditions, applying deep learning to each subdomain can predict simulation results consistently in the entire computational domain. It is possible to predict the simulation results in about 36.9 seconds by the developed method, compared to about 286.4 seconds by the conventional numerical method. In addition to this, we are also attempting to develop a method for fast prediction of time evolution calculations using deep learning.

h3-Open-BDEC: Innovative Software Platform for Integration of (S+D+L) on the BDEC System, such as Wisteria/BDEC-01

Hierarchical, Hybrid, Heterogeneous
h3-Open-BDEC
Big Data & Extreme Computing

JAPAN SOCIETY FOR THE PROMOTION OF SCIENCE · FOUNDED 1932 ·

- "Three" Innovations
  - New Principles for Numerical Analysis by Adaptive Precision, Automatic Tuning & Accuracy Verification
  - Hierarchical Data Driven Approach (*h*DDA) based on Machine Learning
  - Software & Utilities for Heterogenous Environment, such as Wisteria/BDEC-01

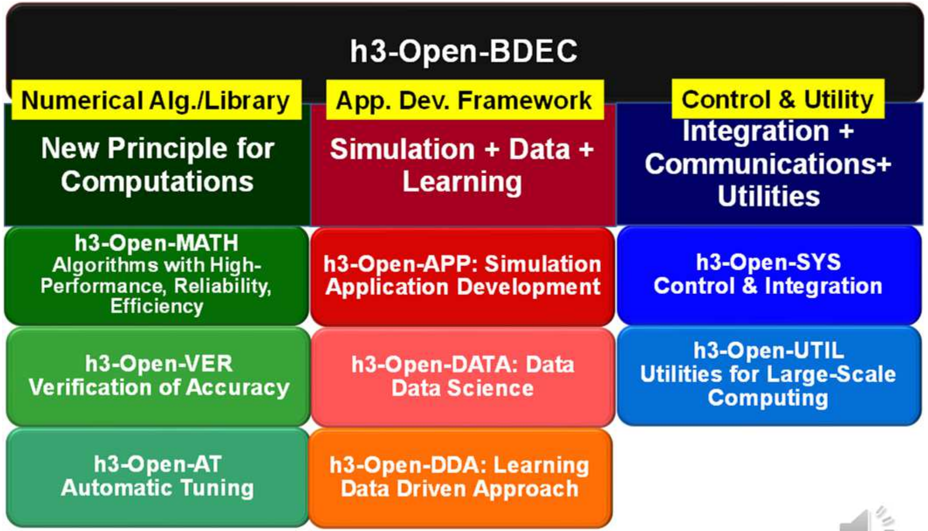

Wisteria/BDEC-01: The First "Really Heterogenous" System in the World

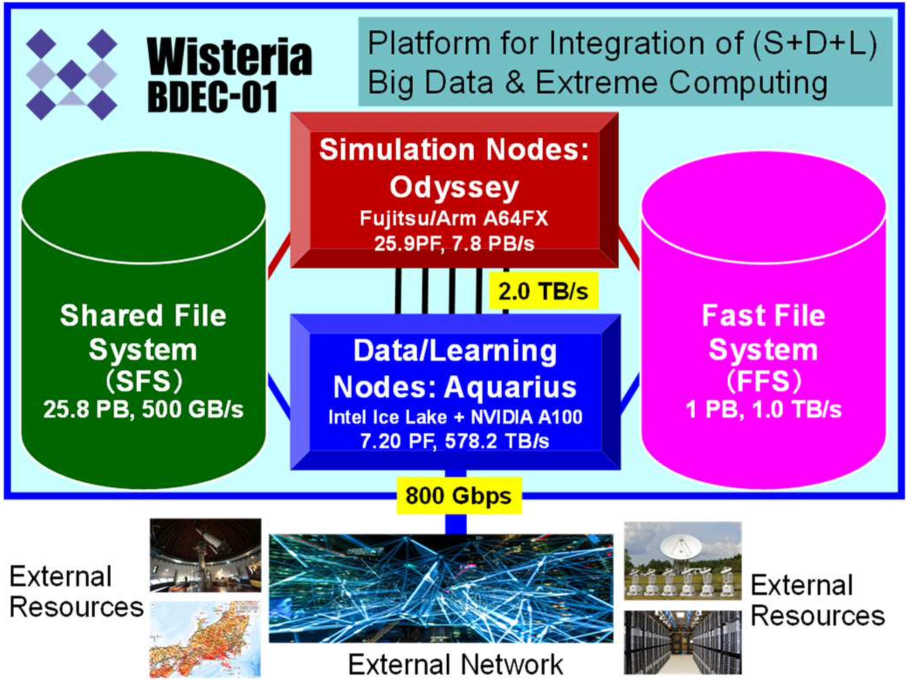

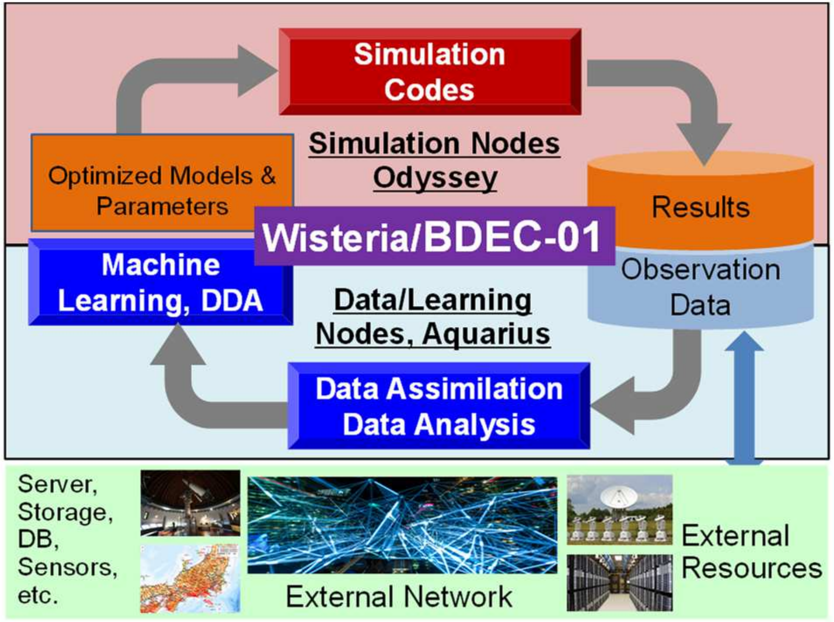

h3-Open-UTIL/MP

Multilevel Coupler/Data Assimilation

- Current Coupler: ppOpen-MATH/MP
  - Weak-Coupling of Multiple (usually two) Applications
    - Each application does a single computation

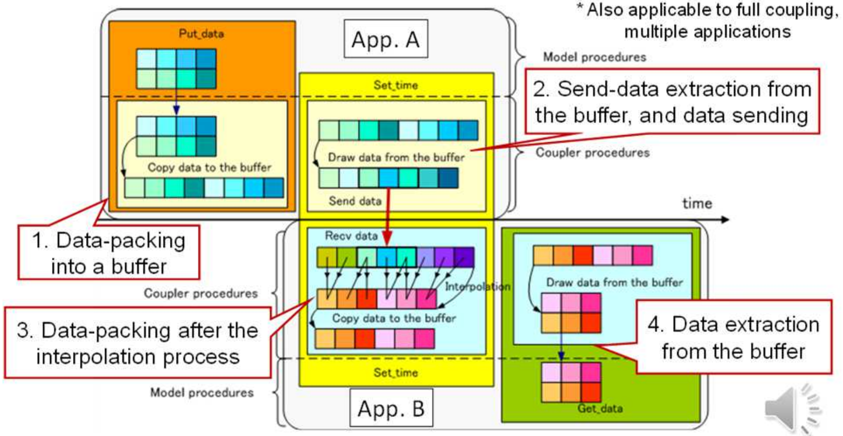

h3-Open-UTIL/MP

Multilevel Coupler/Data Assimilation

- Current Coupler: ppOpen-MATH/MP
  - Weak-Coupling of Multiple (usually two) Applications
    - Each application does a single computation
- h3-Open-UTIL/MP
  - Data Assimilation (Multiple Computations: Ensemble)
  - Assimilation of Computations with Different Resolutions
    - h3-Open-DATA, h3-Open-APP
  - Data Assimilation by Coupled Codes
    - e.g. Atmosphere-Ocean
- Data Assimilation: h3-Open-DATA
  - Karman Filter, Particle Karman Filter
  - LETKF
  - Adjoint Method
- Generation of Simplified Models in hDDA

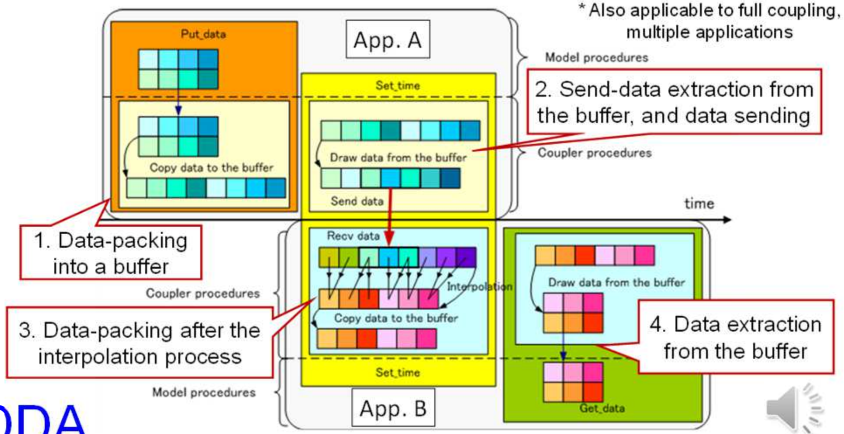

h3-Open-UTIL/MP (h3o-U/MP)

(HPC+AI) Coupling

[Dr. H. Yashiro, NIES]

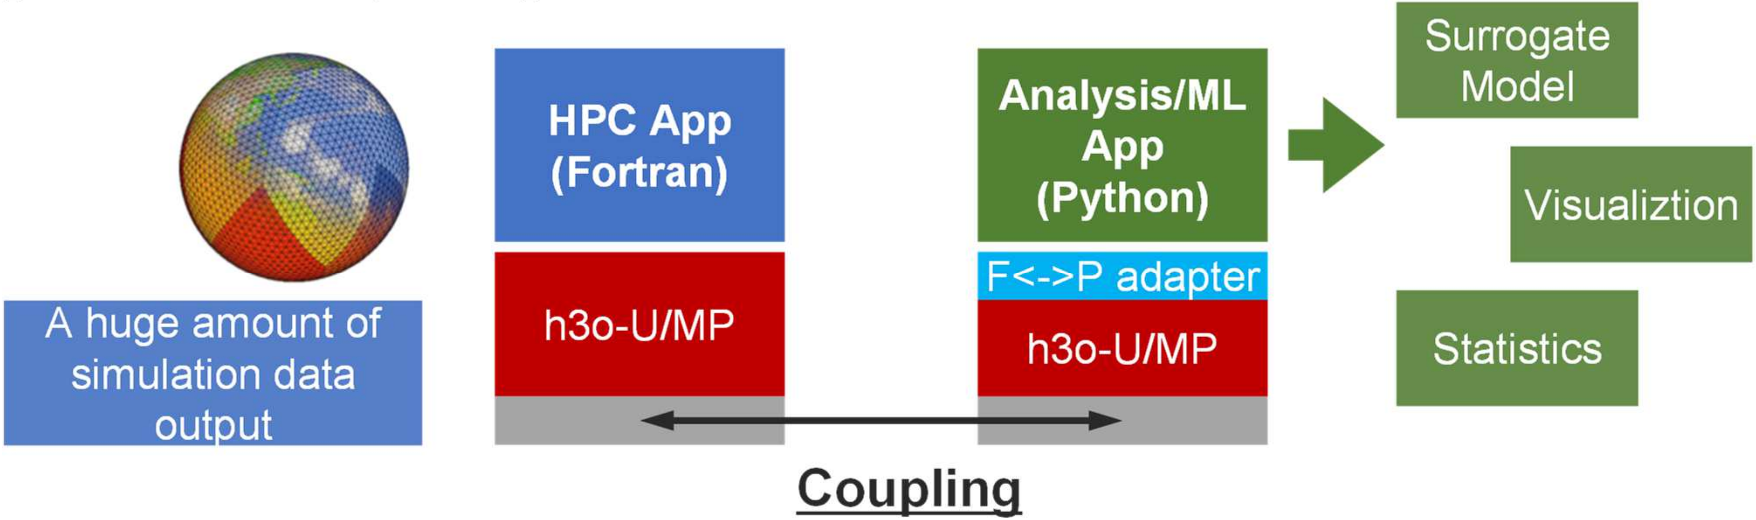

- Providing on-the-fly input/output/training data to the Analysis/ML tools
  - Easy to apply to existing HPC applications
  - Easy access to existing Python-based tools for AI/ML

Computing on Wisteria/BDEC-01

- Wisteria/BDEC-01
  - Aquarius (GPU: NVIDIA A100)
  - Odyssey (CPU: A64FX)
- Combining Odyssey-Aquarius
  - Single MPI Job over O-A is impossible
  - Actually, O-A are connected through IB-EDR with 2TB/sec.
  - h3-Open-SYS/WaitIO-Socket
    - Library for Inter-Process Communication through IB-EDR with MPI-like interface
  - h3-Open-UTIL/MP
    - Multiphysics Coupler

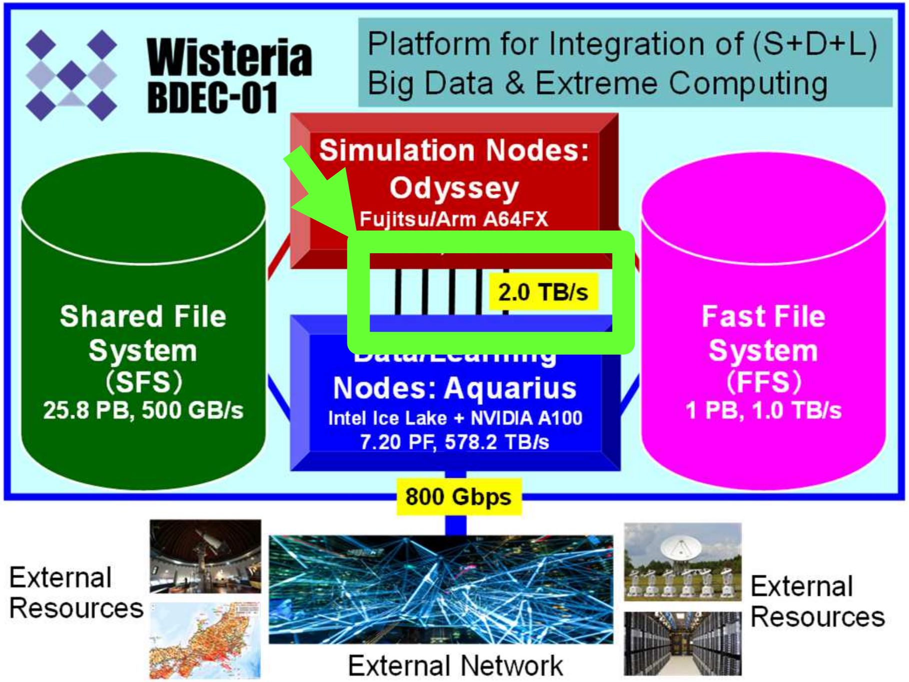

Hierarchical, Hybrid, Heterogeneous
h3-Open-BDEC
Big Data & Extreme Computing

h3-Open-UTIL/MP + h3-Open-SYS/WaitIO-Socket

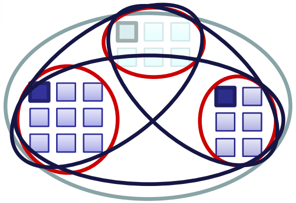

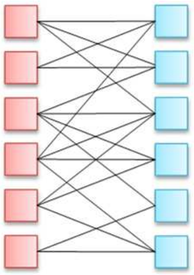

- Single MPI Job (May 2021)
- Direct Communication between Odyssey-Aquarius through IB-EDR by h3-Open-SYS/WaitIO, which provides MPI-like Interface

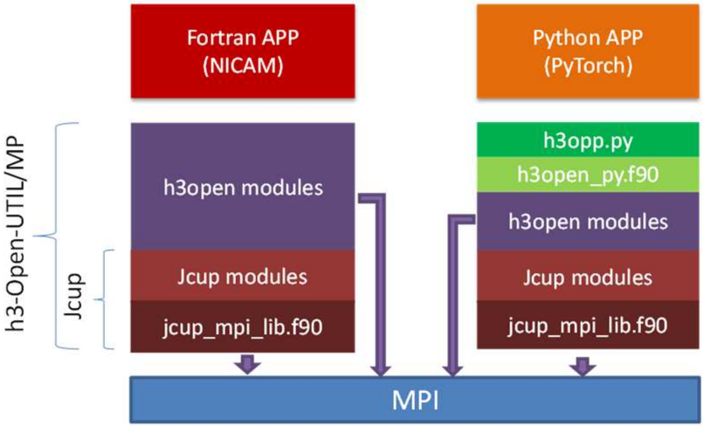

May 2021 (Initial): Single MPI Job

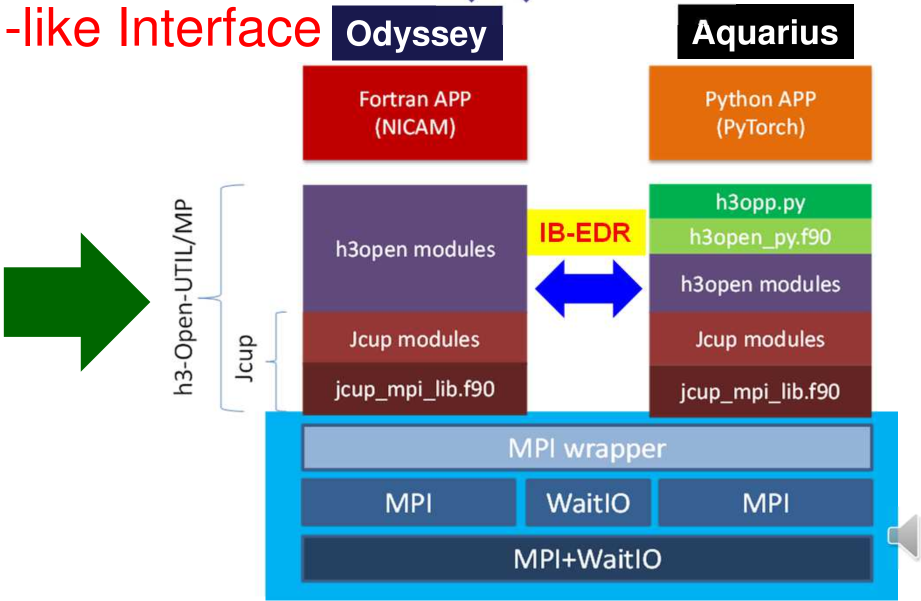

API of h3-Open-SYS/WaitIO-Socket PB (Parallel Block): Each Application

WaitIO API	Description
waitio_isend	Non-Blocking Send
waitio_irecv	Non-Blocking Receive
waitio_wait	Termination of waitio_isend/irecv
waitio_init	Initialization of WaitIO
waitio_get_nprocs	Process # for each PB (Parallel Block)
waitio_create_group waitio_create_group_wranks	Creating communication groups among PB's
waitio_group_rank	Rank ID in the Group
waitio_group_size	Size of Each Group
waitio_pb_size	Size of the Entire PB
waitio_pb_rank	Rank ID of the Entire PB

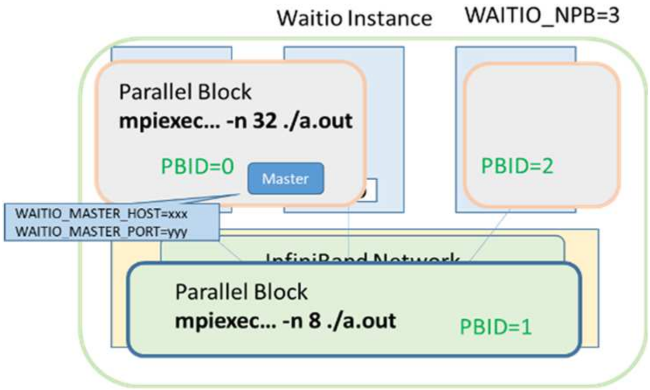

[Sumimoto et al. 2021]

h3-Open-UTIL/MP + h3-Open-SYS/WaitIO-Socket

Available in June 2022

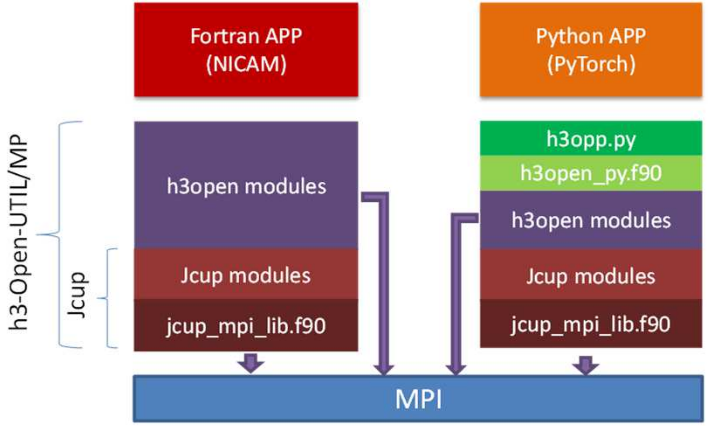

May 2021: MPI Only

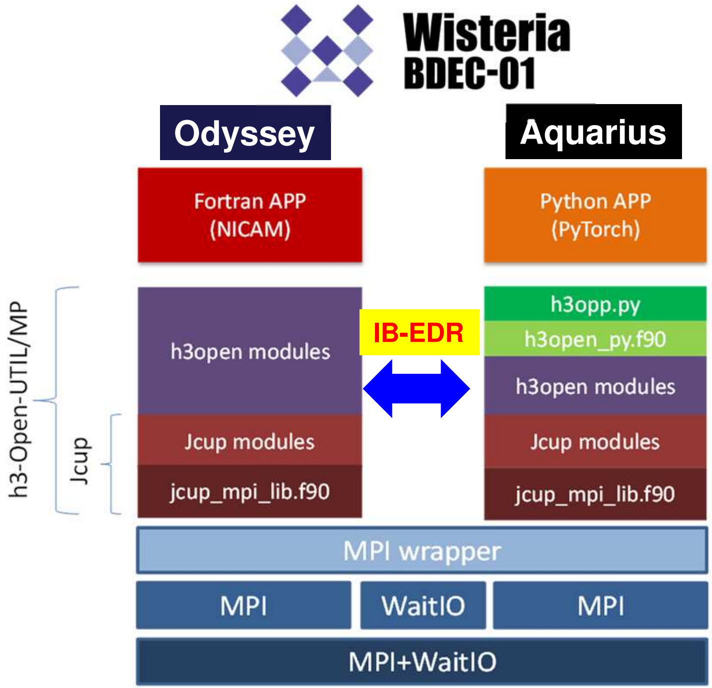

June 2022: Coupler＋WaitIO

3D Earthquake Simulation with Real-Time Data Observation/Assimilation

Simulation of Strong Motion (Wave Propagation) by 3D FDM

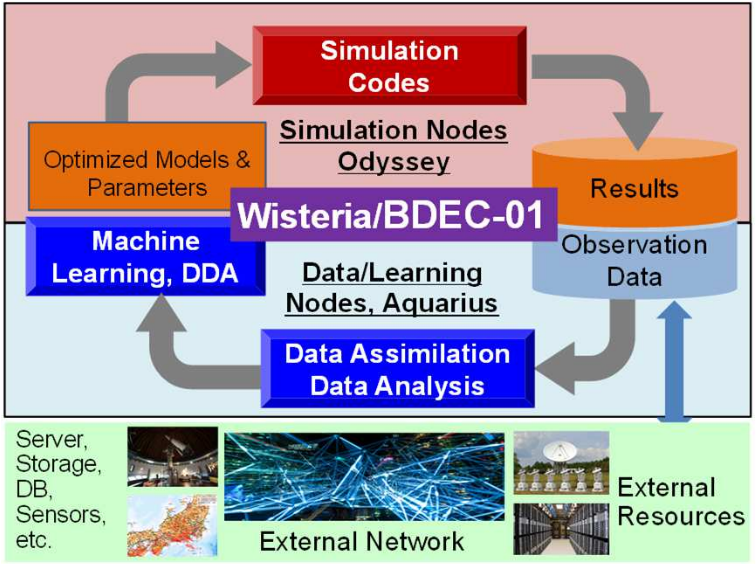

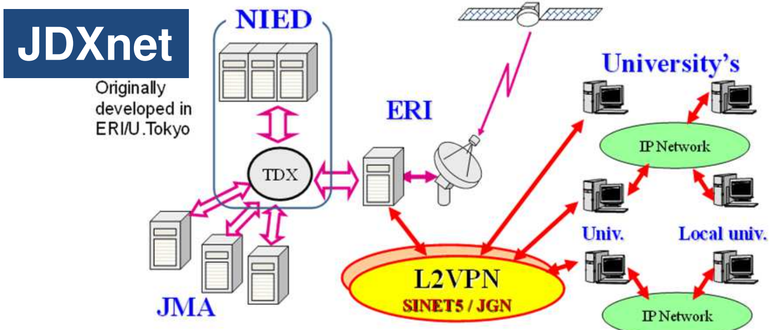

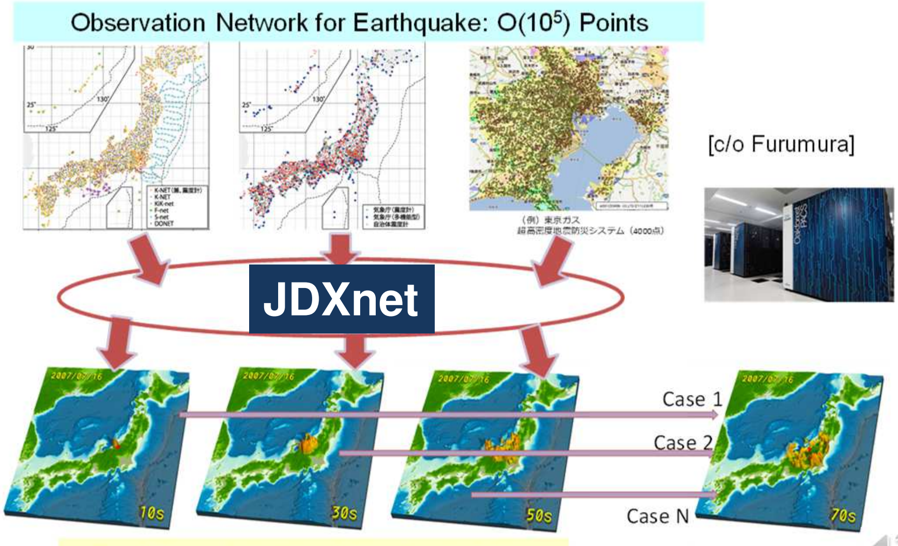

[c/o Prof. T.Furumura (ERI/U.Tokyo)]

System on Wisteria/BDEC-01 using WaitIO

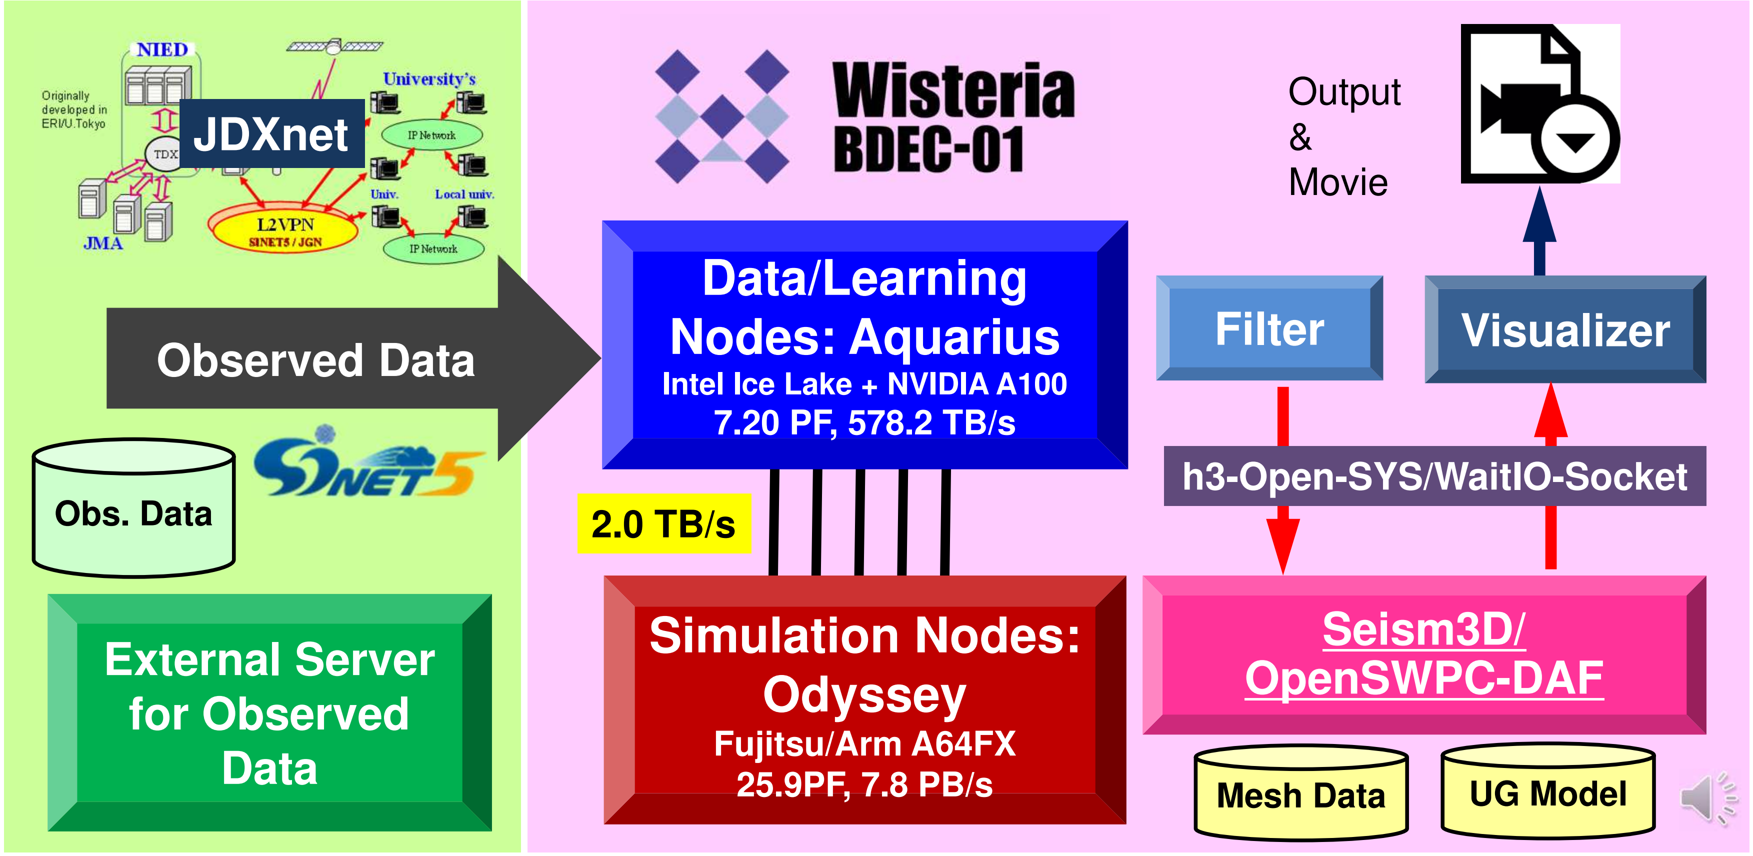

Communications by WaitIO-Socket

[Kasai et al. 2021]

Aquarius: SEND

```
program dmy_filter
<省略: 型宣言等>
call mpi_init (ierr)
call mpi_comm_size (MPI_COMM_WORLD, nprocs, ierr)
call mpi_comm_rank (MPI_COMM_WORLD, myrank,  ierr)
call WAITIO_CREATE_UNIVERSE (WAITIO_COMM_UNIVERSE, ierr)

if (myrank==0) then
  open(100,file='./obsfile_list.txt', form='formatted', status='old', iostat=ierr)
  do i=1,300
   <省略: obsデータ読み込み処理>
    print *,"Send obs data ...... "
    call WAITIO_MPI_ISEND (NTMAX1_o, 1,           WAITIO_MPI_INTEGER,  2,1, WAITIO_COMM_UNIVERSE,req(1,1), ierr)
    call WAITIO_MPI_ISEND (DT_o,     1,           WAITIO_MPI_FLOAT,    2,2, WAITIO_COMM_UNIVERSE,req(1,2), ierr)
    call WAITIO_MPI_ISEND (NST_o,    1,           WAITIO_MPI_INTEGER,  2,3, WAITIO_COMM_UNIVERSE,req(1,3), ierr)
    call WAITIO_MPI_ISEND (AT_o,     1,           WAITIO_MPI_FLOAT,    2,4, WAITIO_COMM_UNIVERSE,req(1,4), ierr)
    call WAITIO_MPI_ISEND (T0_o,     1,           WAITIO_MPI_FLOAT,    2,5, WAITIO_COMM_UNIVERSE,req(1,5), ierr)
    call WAITIO_MPI_ISEND (ISO_X_o,  NSMAX,       WAITIO_MPI_INTEGER,  2,6, WAITIO_COMM_UNIVERSE,req(1,6), ierr)
    call WAITIO_MPI_ISEND (ISO_Y_o,  NSMAX,       WAITIO_MPI_INTEGER,  2,7, WAITIO_COMM_UNIVERSE,req(1,7), ierr)
    call WAITIO_MPI_ISEND (ISO_Z_o,  NSMAX,       WAITIO_MPI_INTEGER,  2,8, WAITIO_COMM_UNIVERSE,req(1,8), ierr)
    call WAITIO_MPI_ISEND (ISTX_o,   NST,         WAITIO_MPI_INTEGER,  2,9, WAITIO_COMM_UNIVERSE,req(1,9), ierr)
    call WAITIO_MPI_ISEND (ISTY_o,   NST,         WAITIO_MPI_INTEGER,  2,10,WAITIO_COMM_UNIVERSE,req(1,10),ierr)
    call WAITIO_MPI_ISEND (ISTZ_o,   NST,         WAITIO_MPI_INTEGER,  2,11,WAITIO_COMM_UNIVERSE,req(1,11),ierr)
    call WAITIO_MPI_ISEND (STC_o,    6*NST,       WAITIO_MPI_CHAR,     2,12,WAITIO_COMM_UNIVERSE,req(1,12),ierr)
    call WAITIO_MPI_ISEND (VxAll_obs,NST*NOBS_LEN,WAITIO_MPI_FLOAT,    2,13,WAITIO_COMM_UNIVERSE,req(1,13),ierr)
    call WAITIO_MPI_ISEND (VyAll_obs,NST*NOBS_LEN,WAITIO_MPI_FLOAT,    2,14,WAITIO_COMM_UNIVERSE,req(1,14),ierr)
    call WAITIO_MPI_ISEND (VzAll_obs,NST*NOBS_LEN,WAITIO_MPI_FLOAT,    2,15,WAITIO_COMM_UNIVERSE,req(1,15),ierr)
    call WAITIO_MPI_WAITALL (15,req, status, ierr)
    call sleep(1)
  enddo
  close (100)
endif
call WAITIO_FINALIZE (ierr)
call mpi_finalize (ierr)
end
```

Odyssey: RECV

```
call WAITIO_MPI_IRECV (NTMAX1_o, 1,           WAITIO_MPI_INTEGER, 0,1, WAITIO_COMM_UNIVERSE,...)
call WAITIO_MPI_IRECV (DT_o,     1,           WAITIO_MPI_FLOAT,   0,2, WAITIO_COMM_UNIVERSE,...)
call WAITIO_MPI_IRECV (NST_o,    1,           WAITIO_MPI_INTEGER, 0,3, WAITIO_COMM_UNIVERSE,...)
call WAITIO_MPI_IRECV (AT_o,     1,           WAITIO_MPI_FLOAT,   0,4, WAITIO_COMM_UNIVERSE,...)
call WAITIO_MPI_IRECV (T0_o,     1,           WAITIO_MPI_FLOAT,   0,5, WAITIO_COMM_UNIVERSE,...)
call WAITIO_MPI_IRECV (ISO_X_o,  NSMAX,       WAITIO_MPI_INTEGER, 0,6, WAITIO_COMM_UNIVERSE,...)
call WAITIO_MPI_IRECV (ISO_Y_o,  NSMAX,       WAITIO_MPI_INTEGER, 0,7, WAITIO_COMM_UNIVERSE,...)
call WAITIO_MPI_IRECV (ISO_Z_o,  NSMAX,       WAITIO_MPI_INTEGER, 0,8, WAITIO_COMM_UNIVERSE,...)
call WAITIO_MPI_IRECV (ISTX_o,   NST,         WAITIO_MPI_INTEGER, 0,9, WAITIO_COMM_UNIVERSE,...)
call WAITIO_MPI_IRECV (ISTY_o,   NST,         WAITIO_MPI_INTEGER, 0,10,WAITIO_COMM_UNIVERSE,...)
call WAITIO_MPI_IRECV (ISTZ_o,   NST,         WAITIO_MPI_INTEGER, 0,11,WAITIO_COMM_UNIVERSE,...)
call WAITIO_MPI_IRECV (STC_o,    6*NST,       WAITIO_MPI_CHAR,    0,12,WAITIO_COMM_UNIVERSE,...)
call WAITIO_MPI_IRECV (VxAll_obs,NST*NOBS_LEN,WAITIO_MPI_FLOAT,   0,13,WAITIO_COMM_UNIVERSE,...)
call WAITIO_MPI_IRECV (VyAll_obs,NST*NOBS_LEN,WAITIO_MPI_FLOAT,   0,14,WAITIO_COMM_UNIVERSE,...)
call WAITIO_MPI_IRECV (VzAll_obs,NST*NOBS_LEN,WAITIO_MPI_FLOAT,   0,15,WAITIO_COMM_UNIVERSE,...)
```

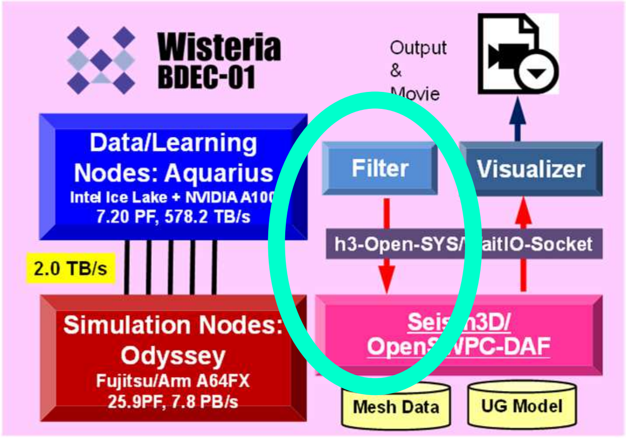

h3-Open-UTIL/MP (h3o-U/MP) + h3-Open-SYS/WaitIO-Socket

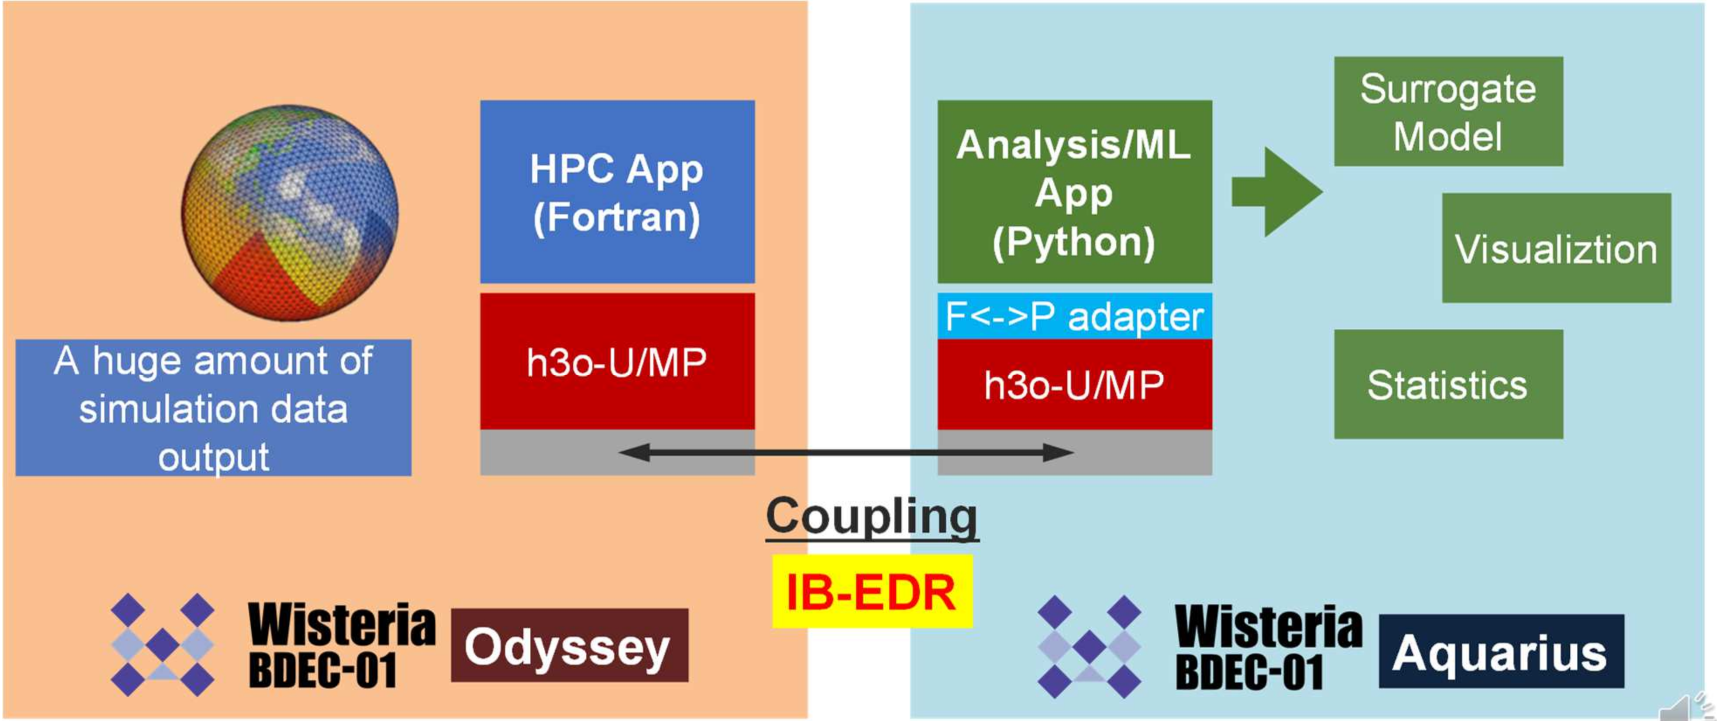

Atmosphere-ML Coupling [Yashiro (NIES), Arakawa (ClimTech/U.Tokyo)]

- Motivation of this experiment
  - Tow types of Atmospheric models: Cloud resolving VS Cloud parameterizing
  - Could resolving model is difficult to use for climate simulation
  - Parameterized model has many assumptions
  - Replacing low-resolution cloud processes calculation with ML!

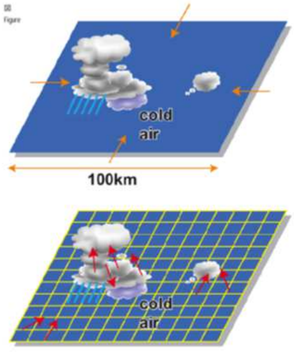

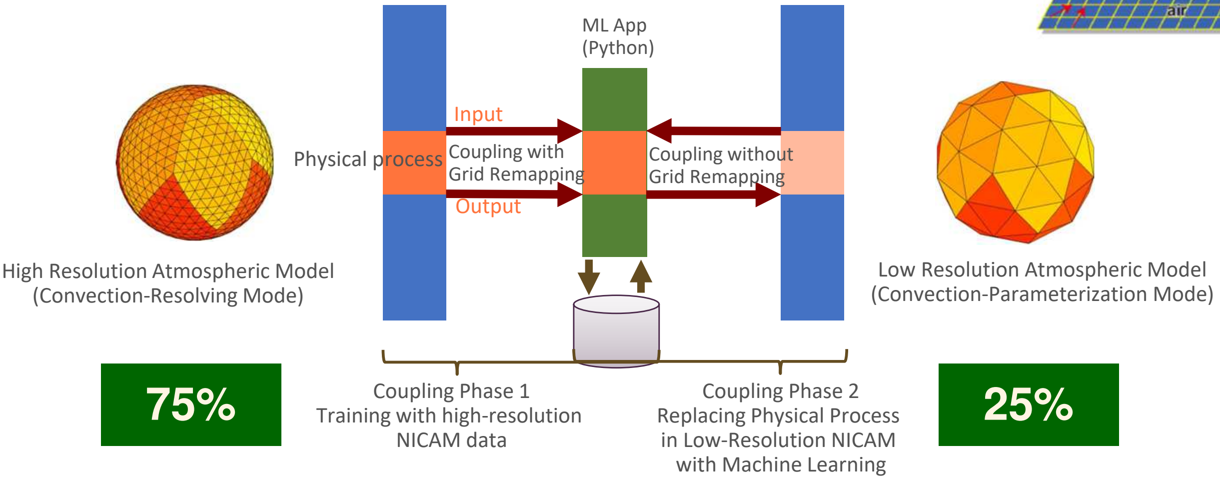

Diagram of applying ML to an atmospheric model

Atmosphere-ML Coupling

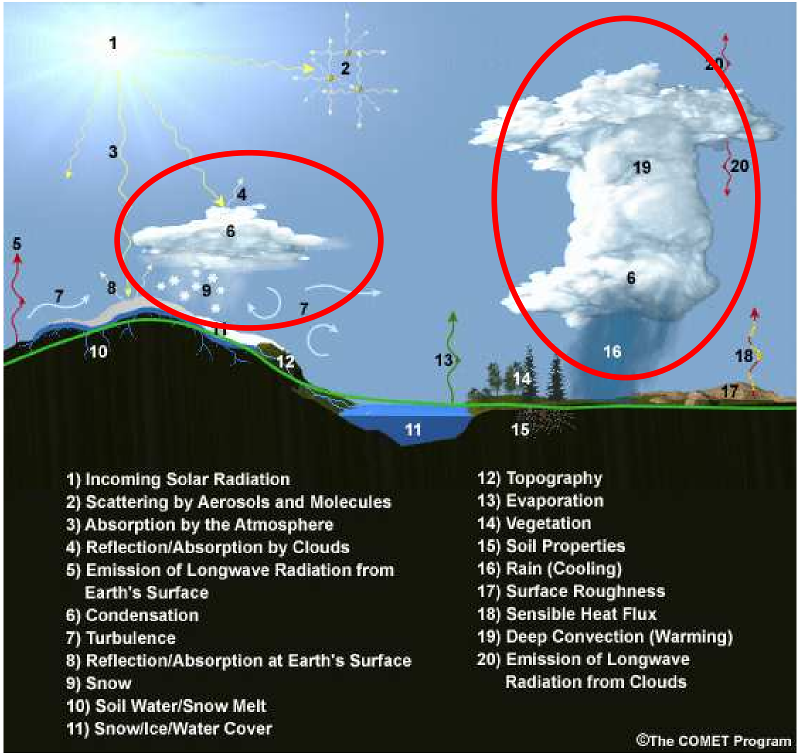

- Model component emulation (surrogation)
  - The emulation target in this study is cloud microphysical processes (phase changes, collision, coagulation, and precipitation)
  - Atmospheric pressure, temperature, and vertical distribution of water will change between before and after computing the cloud microphysical processes
  - The data-driven cloud model predicts atmospheric state changes per unit of time

Experimental Design

- Atmospheric model on Odyssey
  - NICAM : global non-hydrostatic model with an icosahedral grid
  - Resolution : horizontal : 10240, vertical : 78
- ML on Aquarius
  - Framework : PyTorch
  - Method : Three-Layer MLP
  - Resolution : horizontal : 10240, vertical : 78
- Experimental design
  - Phase1: PyTorch is trained to reproduce output variables from input variables of cloud physics subroutine.
  - Phase2:Reproduce the output variables from Input variables and training results
- Training data
  - Input : total air density (rho), internal energy (ein), density of water vapor (rho_q)
  - Output : tendencies of input variables computed within the cloud physics subroutine $\frac{\Delta rho}{\Delta T}$ $\frac{\Delta ein}{\Delta T}$ $\frac{\Delta rho\_q}{\Delta T}$

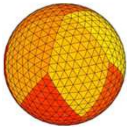

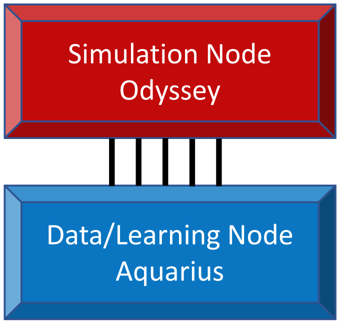

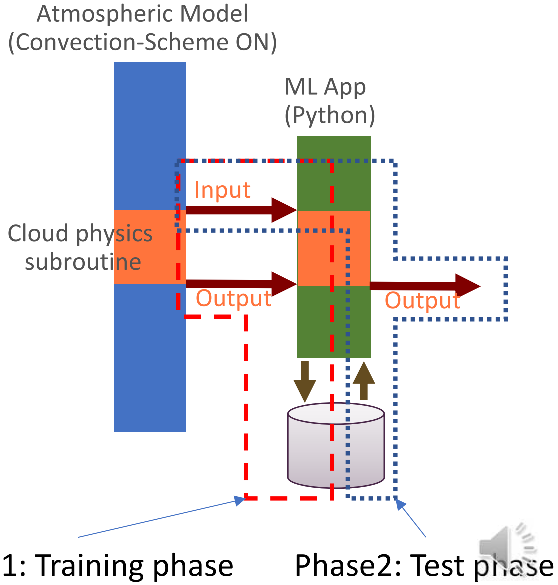

Test calculation

- Compute output variables from input variables and PyTorch
  - The rough distribution of all variables is well reproduced
  - The reproduction of extreme values is no good

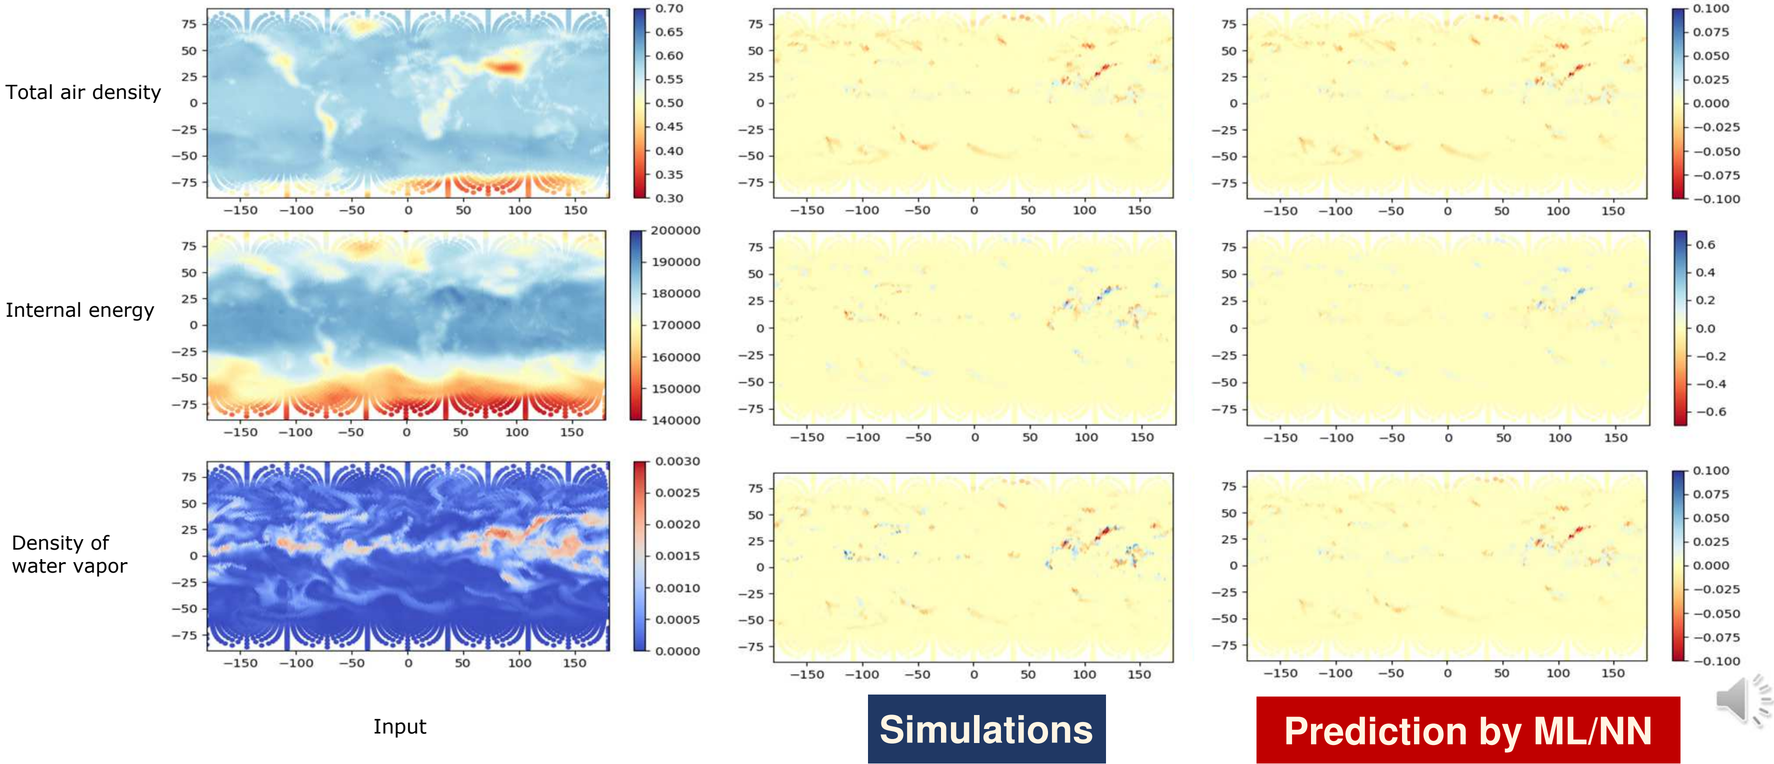

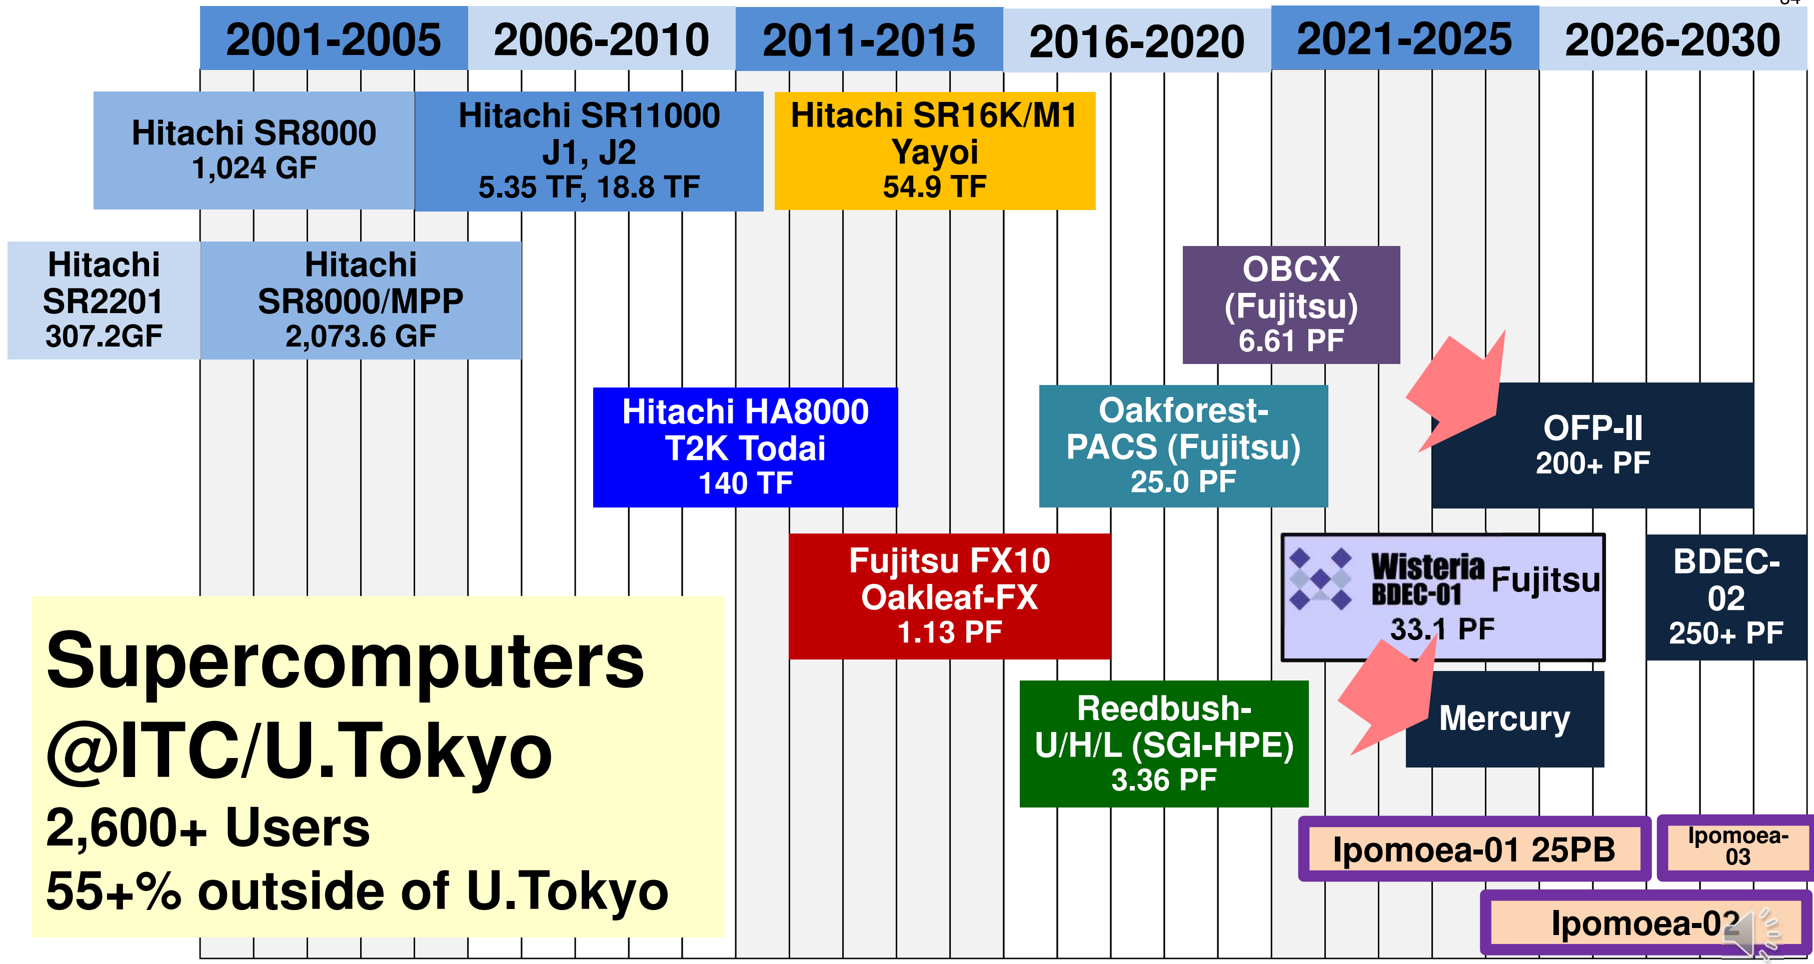
2001-2005
2006-2010
2011-2015
2016-2020
2021-2025
2026-2030
Hitachi SR8000
1,024 GF
Hitachi SR11000
J1, J2
5.35 TF, 18.8 TF
Hitachi SR16K/M1
Yayoi
54.9 TF
Hitachi
SR2201
307.2GF
Hitachi
SR8000/MPP
2,073.6 GF
OBCX
(Fujitsu)
6.61 PF
Hitachi HA8000
T2K Todai
140 TF
Oakforest-
PACS (Fujitsu)
25.0 PF
OFP-II
200+ PF
Fujitsu FX10
Oakleaf-FX
1.13 PF
Wisteria
BDEC-01
Fujitsu
33.1 PF
BDEC-
02
250+ PF
Mercury
Reedbush-
U/H/L (SGI-HPE)
3.36 PF
Ipomoea-01 25PB
Ipomoea-
03
Ipomoea-02
Supercomputers
@ITC/U.Tokyo
2,600+ Users
55+% outside of U.Tokyo

Near Future Plan: Shifting to GPUs

- U.Tokyo is shifting to GPUs/Accelerators in next 10 years
  - Maximum performance under constraint of power consumption
- Wisteria-Mercury (October 2023)
  - GPU Cluster, for supporting “Aquarius”
  - Prototype of OFP-II (128+ GPU’s, 32+ nodes)
- OFP-II (April 2024)
  - Successor of OFP (JCAHPC, U.Tsukuba & U.Tokyo), 200+PF
  - Group-A (CPU+GPU), Group-B (Only CPU)
    - Same GPUs as those of Mercury
    - CPUs in Group-A and Group-B could be different

Wisteria BDEC-01

JCAHPC

筑波大学 University of Tsukuba

東京大学 THE UNIVERSITY OF TOKYO

- Porting codes of 3,000+ users of OFP to GPU is the most critical issue
  - Starting this Fall

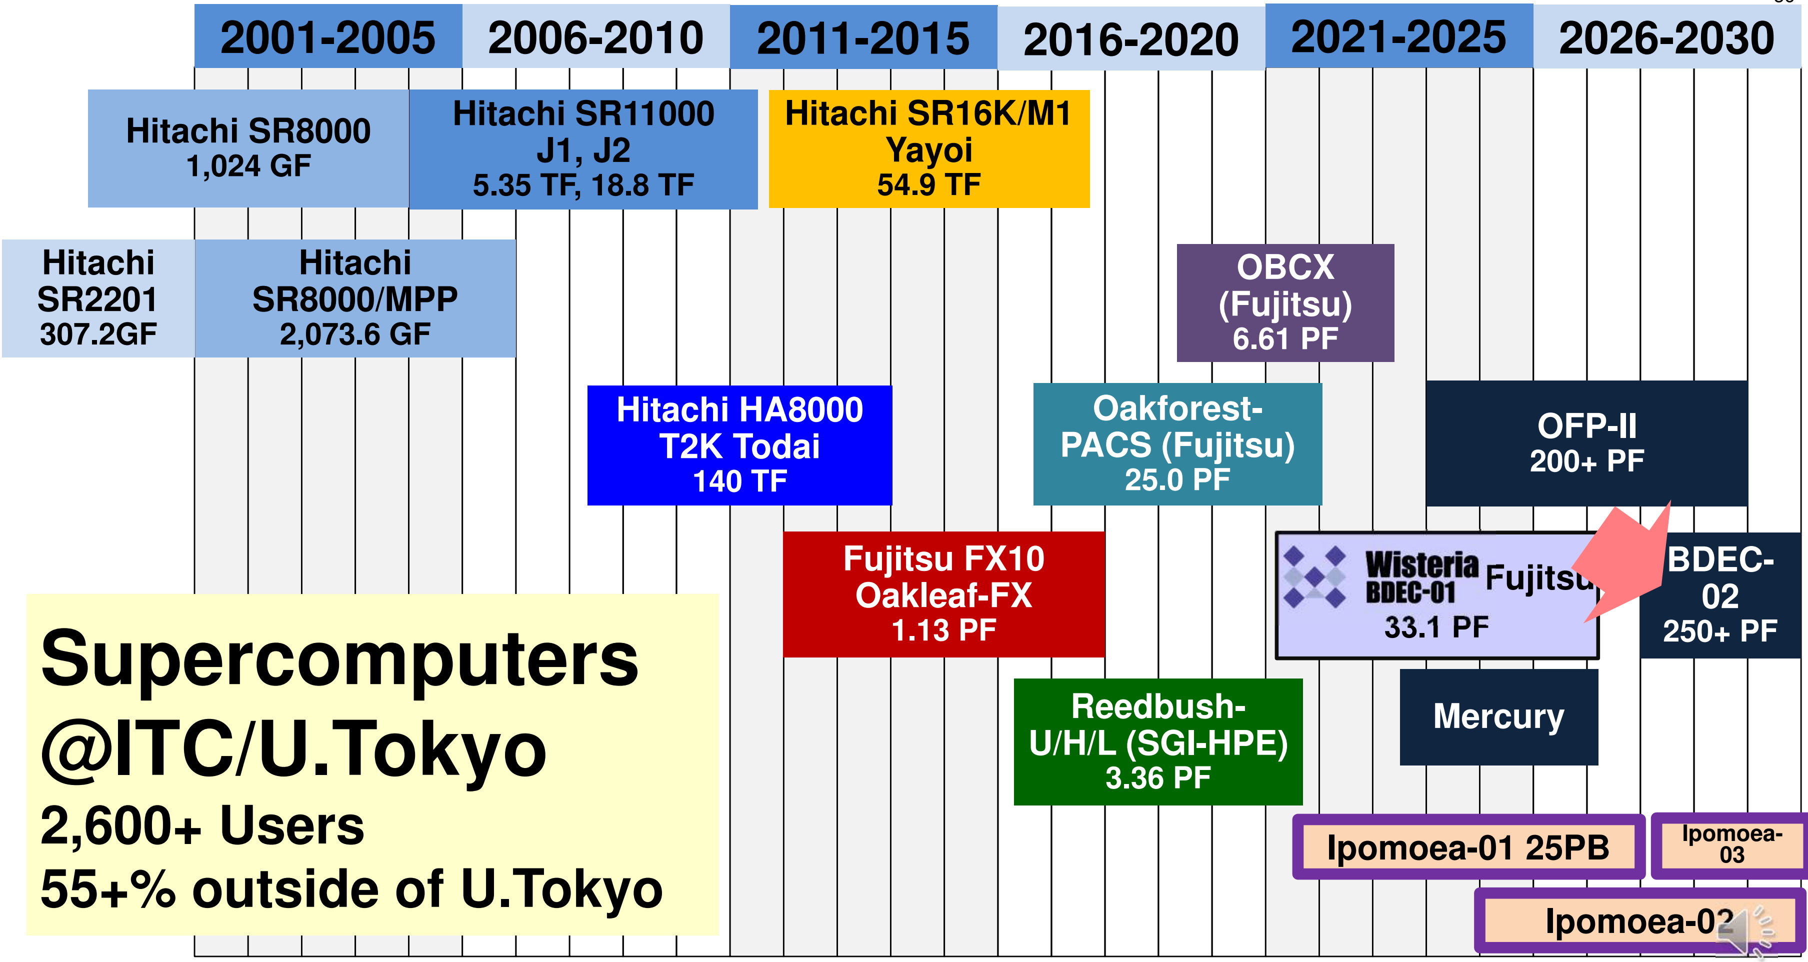
2001-2005
2006-2010
2011-2015
2016-2020
2021-2025
2026-2030
Hitachi SR8000 1,024 GF
Hitachi SR11000 J1, J2 5.35 TF, 18.8 TF
Hitachi SR16K/M1 Yayoi 54.9 TF
Hitachi SR2201 307.2GF
Hitachi SR8000/MPP 2,073.6 GF
OBCX (Fujitsu) 6.61 PF
Hitachi HA8000 T2K Todai 140 TF
Oakforest-PACS (Fujitsu) 25.0 PF
OFP-II 200+ PF
Fujitsu FX10 Oakleaf-FX 1.13 PF
Wisteria BDEC-01 Fujitsu 33.1 PF
BDEC-02 250+ PF
Reedbush-U/H/L (SGI-HPE) 3.36 PF
Mercury
Ipomoea-01 25PB
Ipomoea-03
Ipomoea-02
Supercomputers @ITC/U.Tokyo
2,600+ Users
55+% outside of U.Tokyo

BDEC-02 (Fall 2027-Spring 2028)

- Platform for "Digital Twin", "S+D+L"
- GPU Cluster + CPU Cluster: GPU-Focused
- We are thinking about introducing DPU, IPU, Quantum-Inspired Devices etc. for supporting workloads for (D+L)
  - ✓ We are also considering the introduction of multiple types of GPUs
- We have been using Fujitsu's Digital Annealer since 2019: Combinatorial Optimization
- Programming Environment & Communication Library for Integration of HPC and Such Devices are needed.
  - ✓ We can extend the idea of h3-Open/SYS-WaitIO

Hierarchical, Hybrid, Heterogeneous
h3-Open-BDEC
Big Data & Extreme Computing