HPC at CCS

University of Tsukuba

Masayuki Umemura
(Center for Computational Sciences, University of Tsukuba)

Organization of CCS

- Director
  - Steering Council
  - Strategic Planning Headquarters
  - Administrative Council
  - Researchers Meeting
  - Deputy Director
  - Director for Organization Strategy
  - Research Project Offices
    - Project Office for Exascale Computational Sciences
    - Project Office for Exascale Computing System Development
    - Project Office for HPCI Strategic Program
    - Liaison Office for Multidisciplinary Computational Sciences
    - Promotion Office for Computational Sciences
  - Bureau of Public Relations
  - Director for Cooperative Research
    - Board of Cooperative Research
    - Administrative Committee for Cooperative Research
  - Administrative Committee for Computer Systems
  - Joint Center of Advanced HPC (with U. Tokyo)
  - Research Divisions
    - Division of Particle Physics
    - Division of Astrophysics & Nuclear Physics
    - Division of Quantum Condensed Matter Physics
    - Division of Life Science
    - Division of Global Environmental Science
    - Division of High Performance Computer Systems
    - Division of Computational Informatics

Faculty (full-time)	30
Professor	11
Associate Prof.	12
Lecturer	5
Assistant Prof.	2
Faculty (cooperation)	14
Professor	10
Lecturer	1
Assistant Prof.	3

T2K Tsukuba System

(T2K: Tsukuba, Tokyo and Kyoto Open Supercomputer Alliance)

648 nodes
Peak performance : 95.4TF
(95.4 trillion 10^{12} operations/sec)
Memory 20TB
Disk 800TB
Linpack Result:
Rpeak = 92.0TF (625 nodes)
Rmax = 76.5TF (83.15 %)
20th position in 2008 Jun Top500

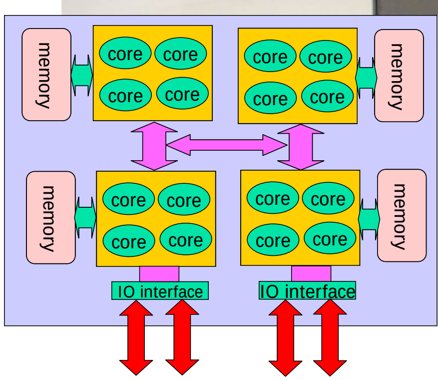

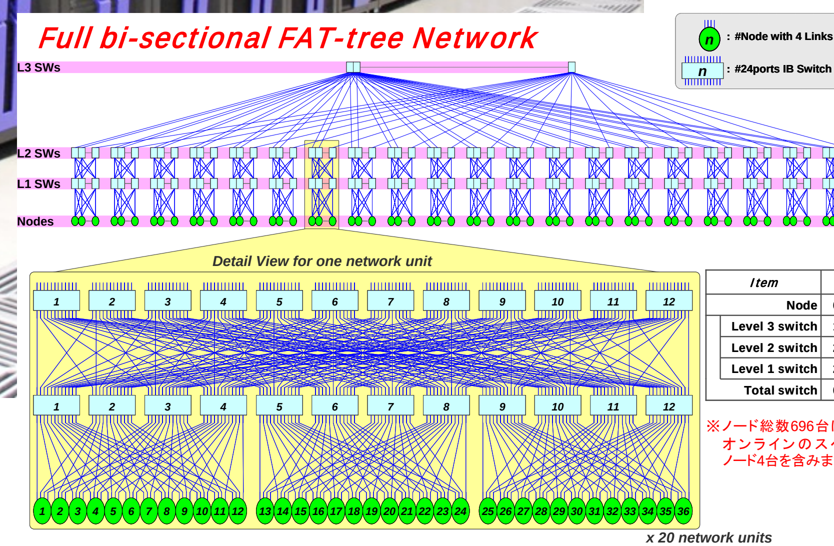

Item	#
Node	696
Level 3 switch	144
Level 2 switch	240
Level 1 switch	232
Total switch	616

※ノード総数696台にはオンラインのスペアノード4台を含みます。

HA-PACS system (base-cluster)

- System spec.
  - 268 nodes
  - CPU 89TFLOPS + GPU 713TFLOPS = total 802TFLOPS
  - Memory 34TByte、memory bandwidth 26TByte/sec
  - Bi-section bandwidth 2.1TByte/秒
  - Storage 504TByte
  - Power 408kW
  - 26 ranks, Installed on Jan, 2012
  - Operation started from Feb, 2012

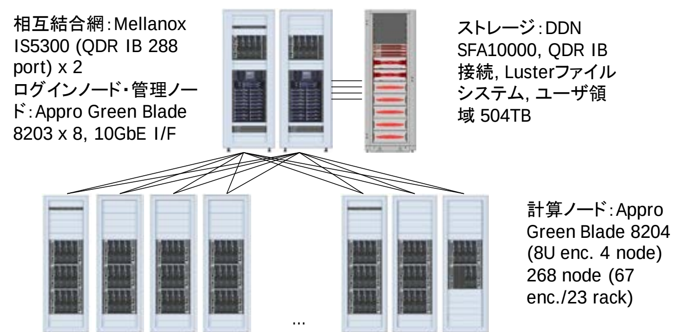

- Spec of compute nodes
  - CPU x2 + GPU x4/node

  （4 GPU: 2660 GFLOPS + 2 CPU: 332 GFLOPS~ 3 TFLOPS / node）

  - Advanced CPU：Intel SandyBridge： high-peak performance enhanded by 256bit AVX instruction, and high memory bandwidth by 1600MHz DDR3 (2.6GHz SandyBridge-EP (8 core) = 166.4 GFLOPS, 51.2GB/s memory bandwitdh, 128GB)
  - x40 lane for CPU direct I/O of PCIe Gen3
  - Advanced GPU：NVIDIA M2090: M2070 512core enhance version：peak performance 665GFLOPS
  - Interconnect network
    - Infiniband QDR x 2 rail (trunk)
    - Connected by PCIe Gen3 x8 lane

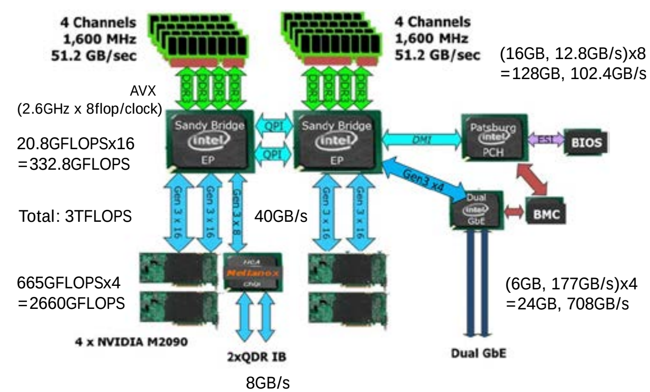

Organization of HA-PACS Project

Project Office for Exascale Computational Sciences

Project Office for Exascale Computing System Development

Masayuki Umemura	CCS, U. Tsukuba	Astrophysics , *Chief*
Kazuhiro Yabana	CCS, U. Tsukuba	Nuclear Physics
Hiroshi Tanaka	CCS, U. Tsukuba	Geoenvironment
Hiroyuki Kitagawa	CCS, U. Tsukuba	Database
Yoshinobu Kuramashi	CCS, U. Tsukuba	Particle Physics
Tong Xiao-Min	CCS, U. Tsukuba	Materials Science
Hiroyuki Kusaka	CCS, U. Tsukuba	Geoenvironment
Toshiyuki Amagasa	CCS, U. Tsukuba	Database
Daisuke Takahashi	CCS, U. Tsukuba	Computer Science
Toshihiro Hanawa	CCS, U. Tsukuba	Computer Science
Kohji Yoshikawa	CCS, U. Tsukuba	Astrophysics
Mitsuo Shoji	CCS, U. Tsukuba	Bioscience
Koji Terasaki	CCS, U. Tsukuba	Geoenvironment
Ken'ichi Ishikawa	Hiroshima U.	Particle Physics
Yasuteru Shigeta	Osaka U.	Bioscience

- **Collaboration between Computational Science and Computer Science**
- **Collaboration among Computational Sciences**
  - Algorisms
  - GPU computing
  - Science goals

Computational Sciences by HA-PACS

Particle Physics

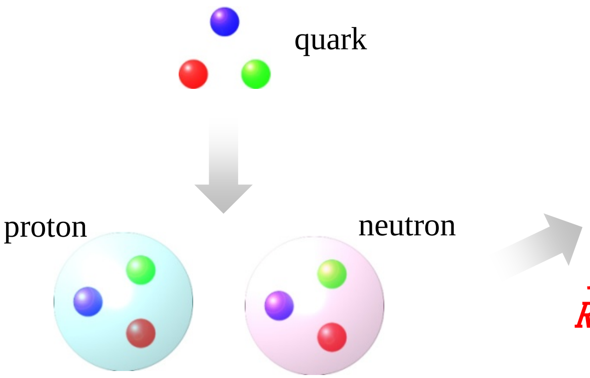

Nuclear Physics

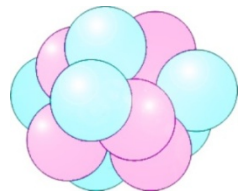

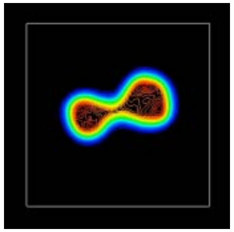

Quantum Many-body System

Materials Science

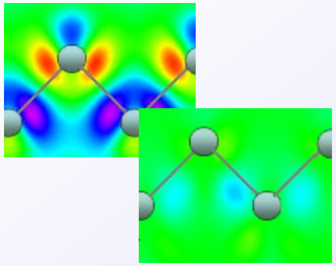

Nuclear Reactions

Particle Reactions

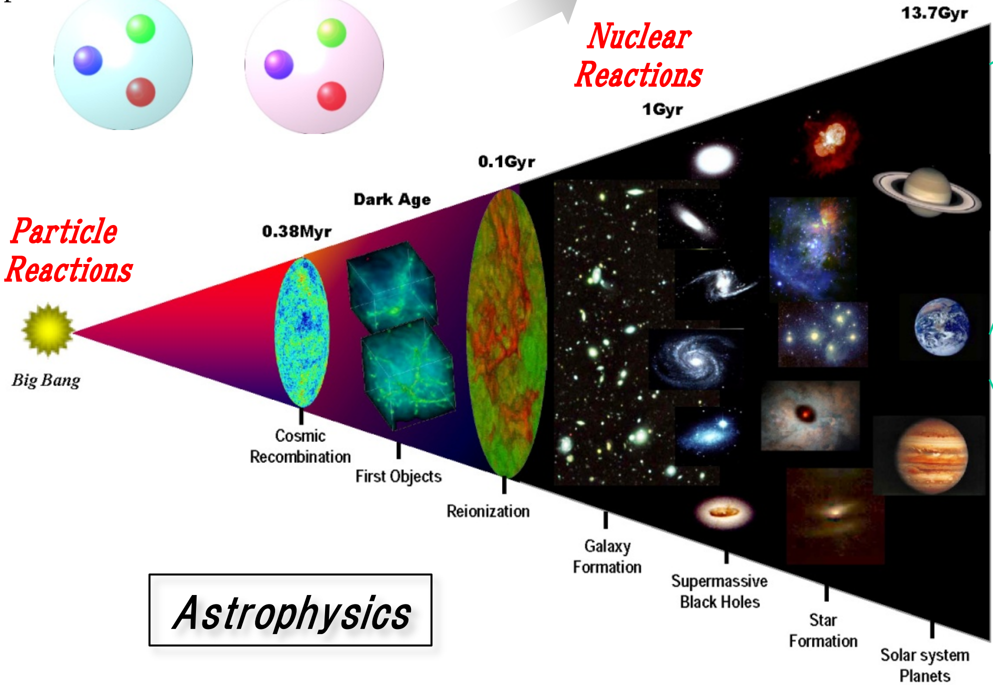

Geoenvironment

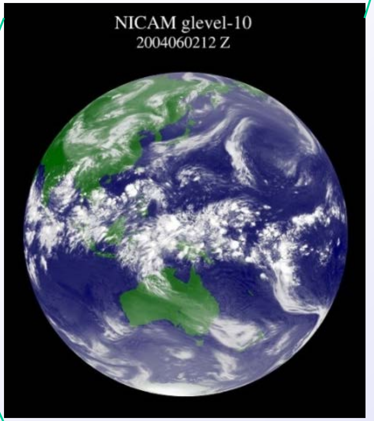

Bioscience

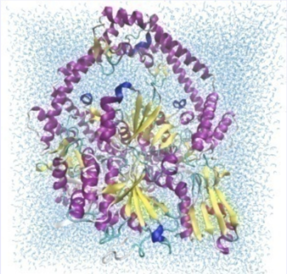

Astrophysics

Database

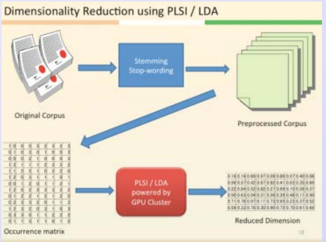

Hydrodynamics, Radiation, Chemical Reactions

Particle Physics

Multi-scale physics

Investigate hierarchical properties via direct construction of nuclei in lattice QCD

GPU to solve large sparse linear systems of equations

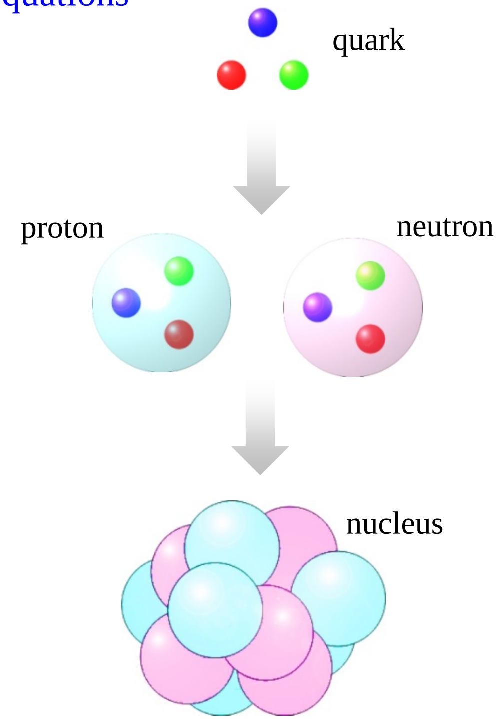

Finite temperature and density

Phase analysis of QCD at finite temperature and density

GPU to perform matrix-matrix product of dense matrices

Expected QCD phase diagram

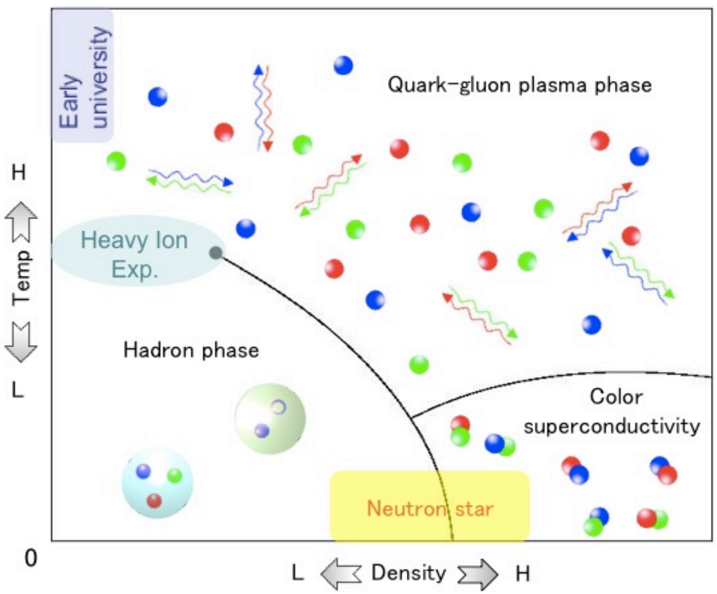

Astrophysics

(A) 6-Dimensional Radiation Hydrodynamics

3-Dimenasional Hydrodynamics + 6-Dimensional Radiation Transfer

<u>Goals</u> Elucidation of Galaxy Formation
Elucidation of Cosmic Reionization

➢ Acceleration by GPU/TCA of ray tracing and chemical reactions, which are of strong scaling.
➢Realization of radiation hydrodynamics

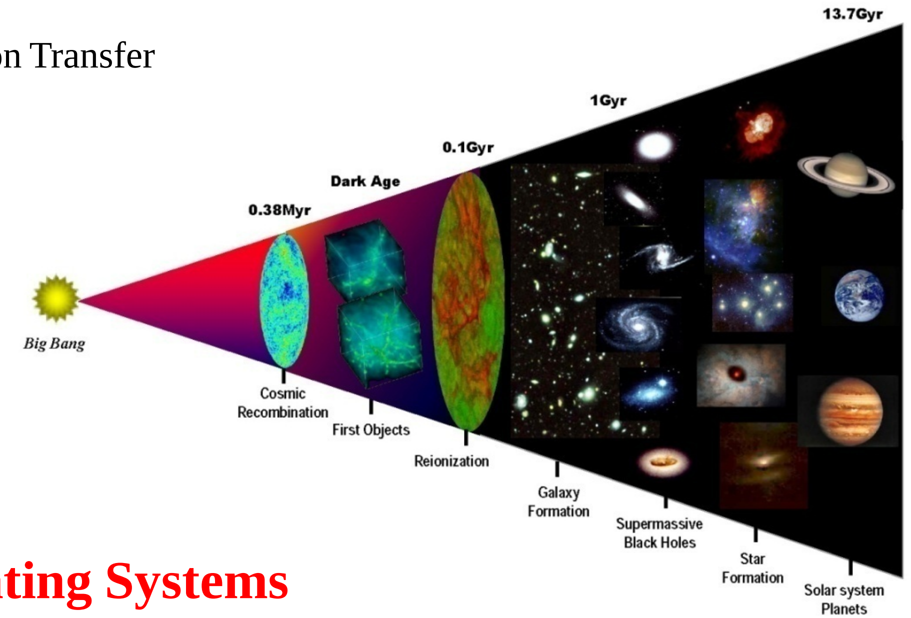

(B) 6D Vlasov Simulation of Self-Gravitating Systems

<u>Goals</u> Dark Matter Dynamics
Collisionless plasma

➢ A direct integration of collisionless Boltzmann equation
➢ Not suffer from two-body relaxation which is inevitable in N-body simulation

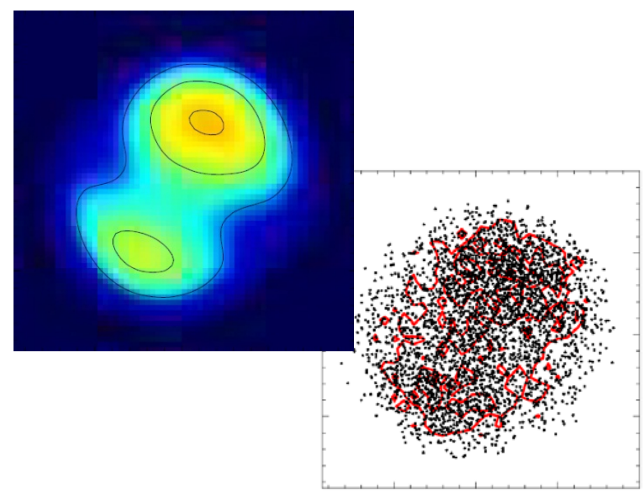

Nuclear Physics

Simulation with real-time and real-space method for many-fermion systems

Nuclear response and reaction dynamics relevant to nucleosynthesis

- Nuclear transfer reaction to produce r-process nuclei experimentally
- Fusion reaction of light nuclei
- Systematic investigation of nuclear response function

Application of nuclear methods to other fields

- First-principles calculation for light-matter interaction
- Propagation of ultra-intense laser pulse
- Simulation for atto-second electron dynamics

Methodology : Time-dependent mean-field theory (TDDFT, TDHF, TDHFB) with real-time and 3D real-space method

Merit of GPU calculation : High performance calculation for the operation of Hamiltonian on orbital wave functions

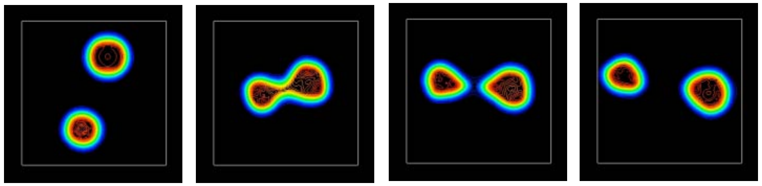

TDHF simulation to produce neutron-rich nuclei by multi-nucleon transfer reaction

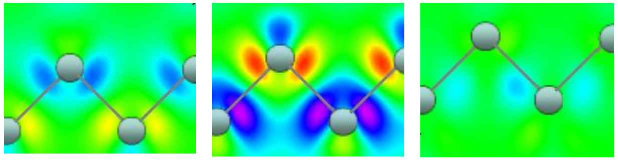

Atto-second electron dynamics in solid induced by ultrashort laser pulse

Materials Science

Develop a general numerical method to solve the time-dependent Schrödinger equation for many-electron quantum systems and use it to

$$i\frac{\partial}{\partial t}\Psi = H\Psi$$

- ***understand*** atomic, molecular and materials structures and their dynamics
- ***search*** a way to ***control*** the structures and dynamics in femtosecond (10^{-15} s) or even attosecond (10^{-18} s) time scales.

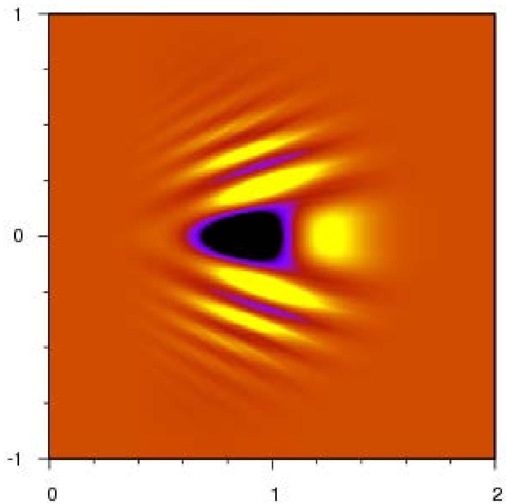

Holographic image of an electron wavepacket colliding with ionic core.

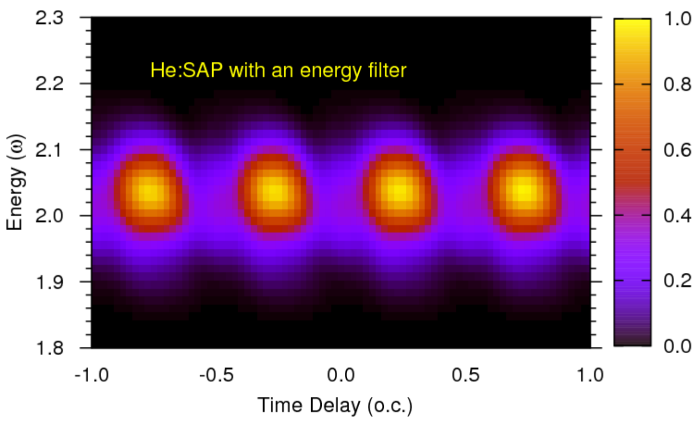

Controlling the XUV transparency by IR laser in attosecond time scale

Bioscience

New developments in computational biosciences

GPU acceleration
- **Molecular Dynamics (MD)**
- **Quantum Mechanics (QM)**

①

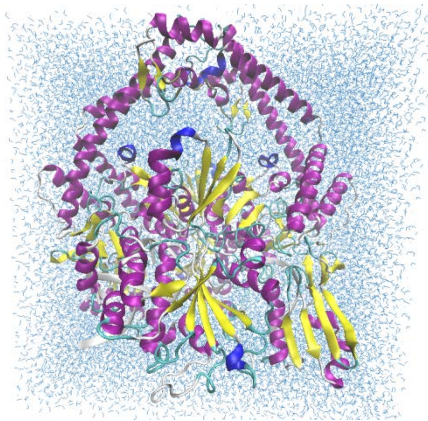

Macroscale MD
DNA-protein complexes)

②

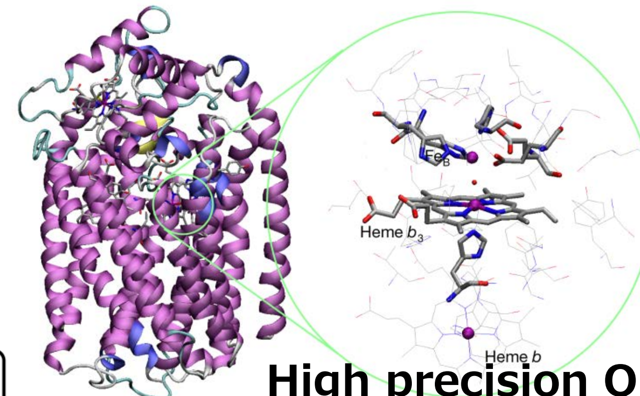

Large QM region
> 100 atoms

High precision QM calculations for reaction mechanisms
（QM/MM-MD, RSDFT-CPMD, PW-CPMD）

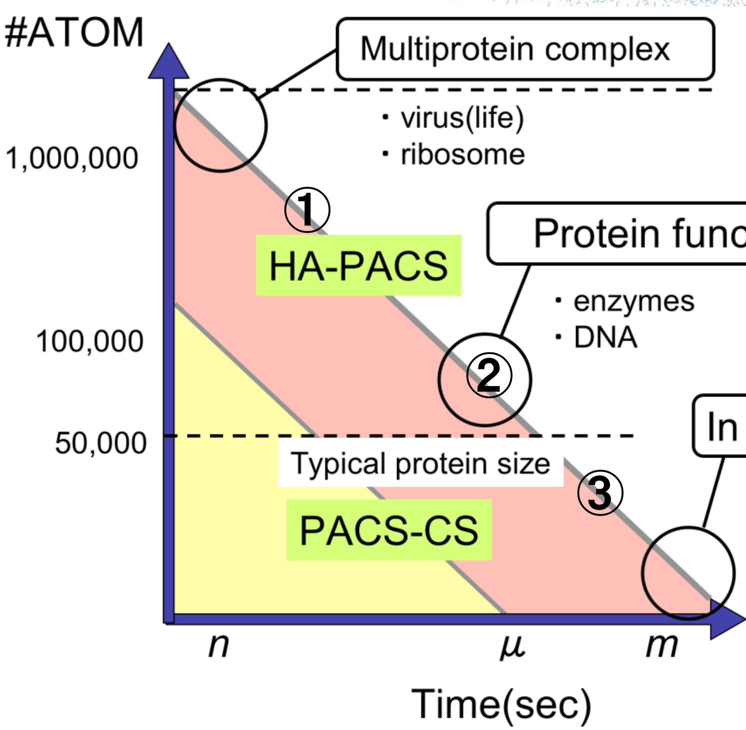

③

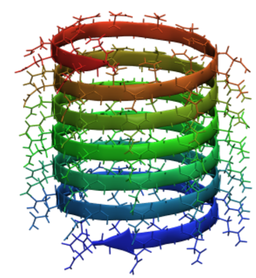

Long time MD
(Prion Proteins)

Geoenvironment

Objectives

✓GPU application to the Next-Generation Atmospheric
General Circulation Modell NICAM
✓GPU application to the Large Eddy Simulation (LES)
✓GPU application to the 3D Normal Mode Expansion
of the atmospheric state variables

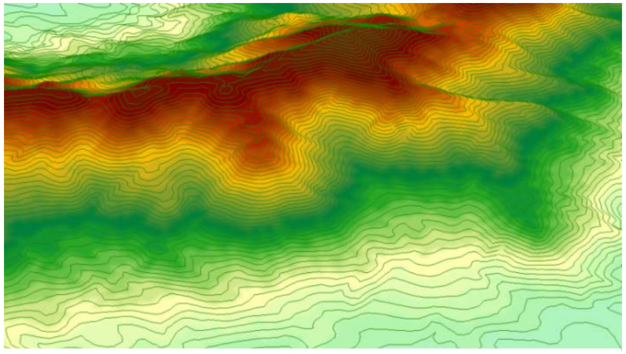

Topography with a 50 m grids

Expected Products

- LES model with 10 m spatial resolution is developed by the GPU application
- NICAM physical processes is efficiently calculated by the GPU application
- Energetics analysis of the high-resolution atmospheric GCM is possible by the GPU application

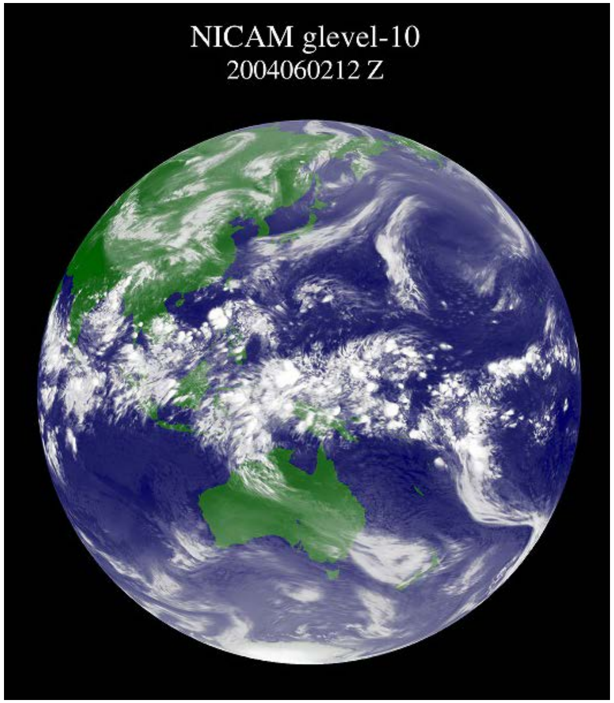

Merit of the GPU/TCA application

Weather forecasting model by a grid discretization is a type of stencil computation. The memory access is therefore simple, and the computational acceleration up to 10 times speed is possible by the GPU/TCA application.

Data Mining of Big Data based on GPGPU

- Research objective and plan
  - **Accelerating data mining** from big data using **GPU**
  - Target mining algorithms
    - Document clustering
      - PLSI (Probabilistic Latent Semantic Indexing)
      - LDA (Latent Dirichret Allocation)
    - Probabilistic association-rule mining
  - Developed algorithms for single-GPU.
  - Develop multi-GPU versions for GPU-cluster environment based on the current algorithms.

- Expected results and breakthrough
  - Application of GPU-cluster to problems other than numerical analysis or simulation.
  - Few existing works have applied GPU-cluster to data mining problems so far.
  - Promote the use of GPU-cluster as a platform for big data analysis.

- Applicability of GPU
  - Some data mining algorithms are suitable for GPU, but others may not.
  - A technical challenge is to combine CPU- and GPU-based computation taking account of the algorithmic characteristics.
- Scale of computation
  - Under consideration
  - Aiming at processing big datasets such that GPU-cluster is necessary.

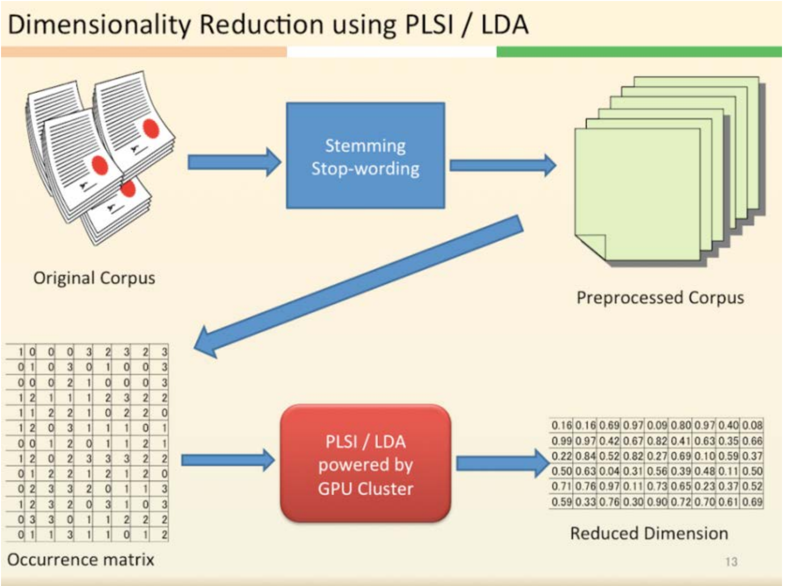

K computer

- SPARC64™ VIIIfx 2.0GHz octcore (128Gflops / core)
- 16 GB memory / core
- 6D torus network
- Total 82944 nodes (663552 CPU core)
- 1.3PB memory
- 10.6 Pflops peak speed

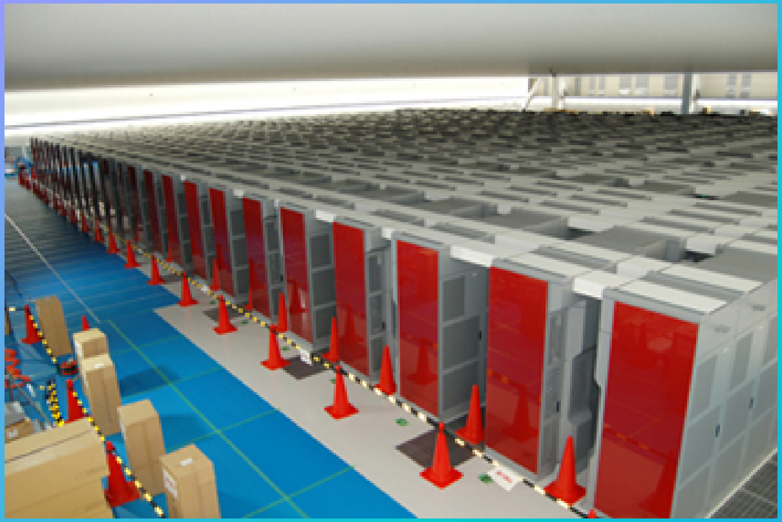

Prediction of Electron States of Si Nanowires with **100,000 atoms on K Computer**

Gordon Bell Prize 2011

Si Nanowire, a booster
in the next-generation semiconductor technology
More Moore → More than Moore

planar transistor

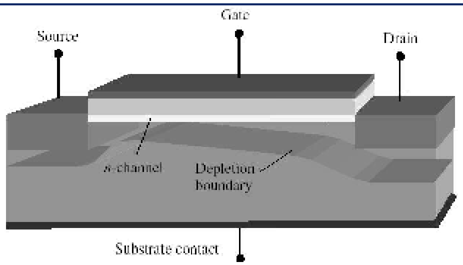

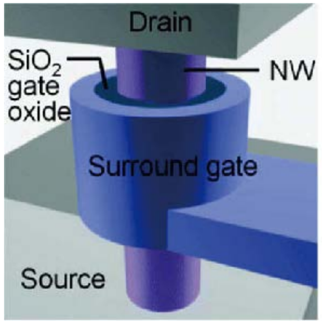

Surrounding gate transistor

Gate Controllability
→ Suppress short-channel effects
Suppress leaks at off state
→ save energy

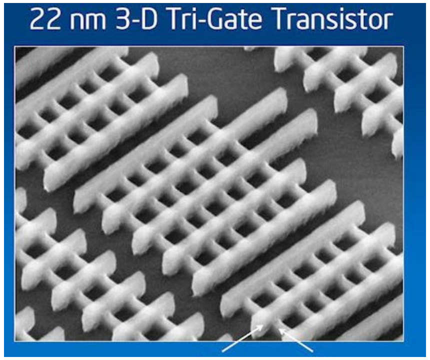

Actually tri-gate by Intel in 2011

Number of atoms in SiNW channels
→ 10,000 - 100,000 atoms !

Collaborators

- Yukihiro Hasegawa (RIKEN)
- Jun-Ichi Iwata (The University of Tokyo)
- Miwako Tsuji (University of Tsukuba)
- Daisuke Takahashi (University of Tsukuba)
- Atsushi Oshiyama (The University of Tokyo)
- Kazuo Minami (RIKEN)
- Taisuke Boku (University of Tsukuba)
- Fumiyoshi Shoji (RIKEN)
- Atsuya Uno (RIKEN)
- Motoyoshi Kurokawa (RIKEN)
- Hikaru Inoue (Fujitsu Limited)
- Ikuo Miyoshi (Fujitsu Limited)
- Mitsuo Yokokawa (RIKEN)

Solving Kohn-Sham equation -*RSDFT*-

J.-I. Iwata *et al.*, J. Comp. Phys. (2010)

Real Space

Blue : Si atom
Yellow: electron density

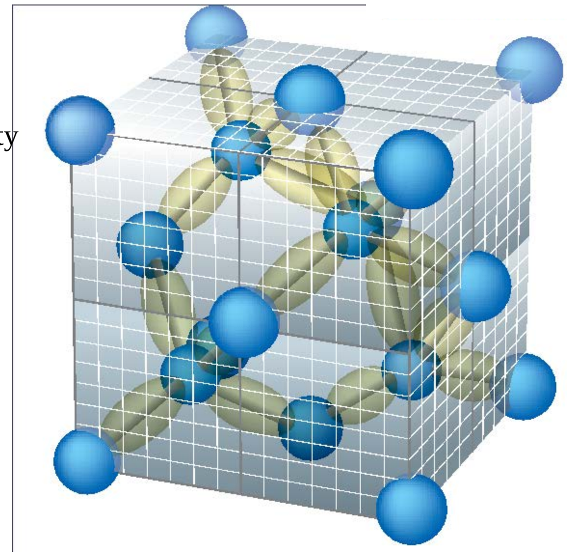

CPU space

MPI for inter-CPUs
Open-MP for intra-CPU

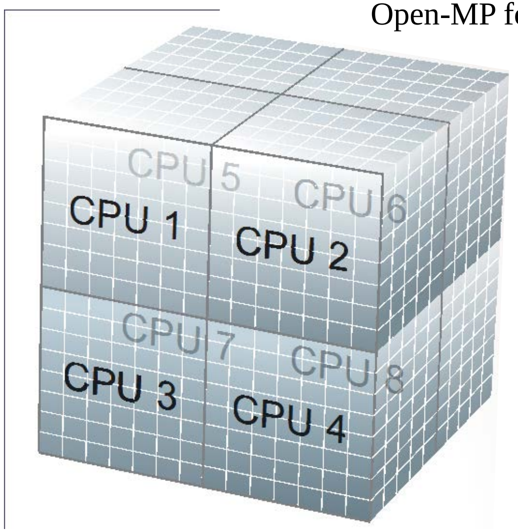

Kohn-Sham differential equation is converted to M-th order finite-difference equation

$$\frac{\partial^2}{\partial x^2}\varphi_j(x,y,z) = \sum_{m=-M}^{M} c_m \ \varphi_j(x+mH,y,z)$$

<u>Advantages</u>

- Almost <u>free from FFT</u>, reducing communication burden ⇒ high efficiency
- <u>No periodic boundary condition</u> to wave-functions
  ⇒ molecules, clusters, surfaces, etc.
- ……

Flow of the SCF calculation

Kohn-Sham equation

$$\left[-\frac{1}{2}\nabla^2 + v_{\mathrm{nucl}}(\mathbf{r}) + \int \frac{n(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|}d\mathbf{r}' + \frac{\delta E_{XC}[n]}{\delta n(\mathbf{r})}\right]\varphi_i(\mathbf{r}) = \varepsilon_i\,\varphi_i(\mathbf{r})$$

Electron density

$$n(\mathbf{r}) = \sum_{i:\ \mathrm{occupied}} |\varphi_i(\mathbf{r})|^2$$

Self-Consistent Field procedure

		Computational complexity (Number of atoms : N)
1	(CG) Conjugate-Gradient Method	$O(N^2)$
2	(GS) Gram-Schmidt ortho-normalization	$O(N^3)$
3	Density, Potential update	$O(N)$
4	(SD) Subspace Diagonalization (an eigenvalue problem)	$O(N^3)$

GS and SD parts are hot spot

Scalability of parallelization in grid spaces and orbitals

SiNW, 19848 atoms, Grid spaces:320x320x120, Orbitals:40992

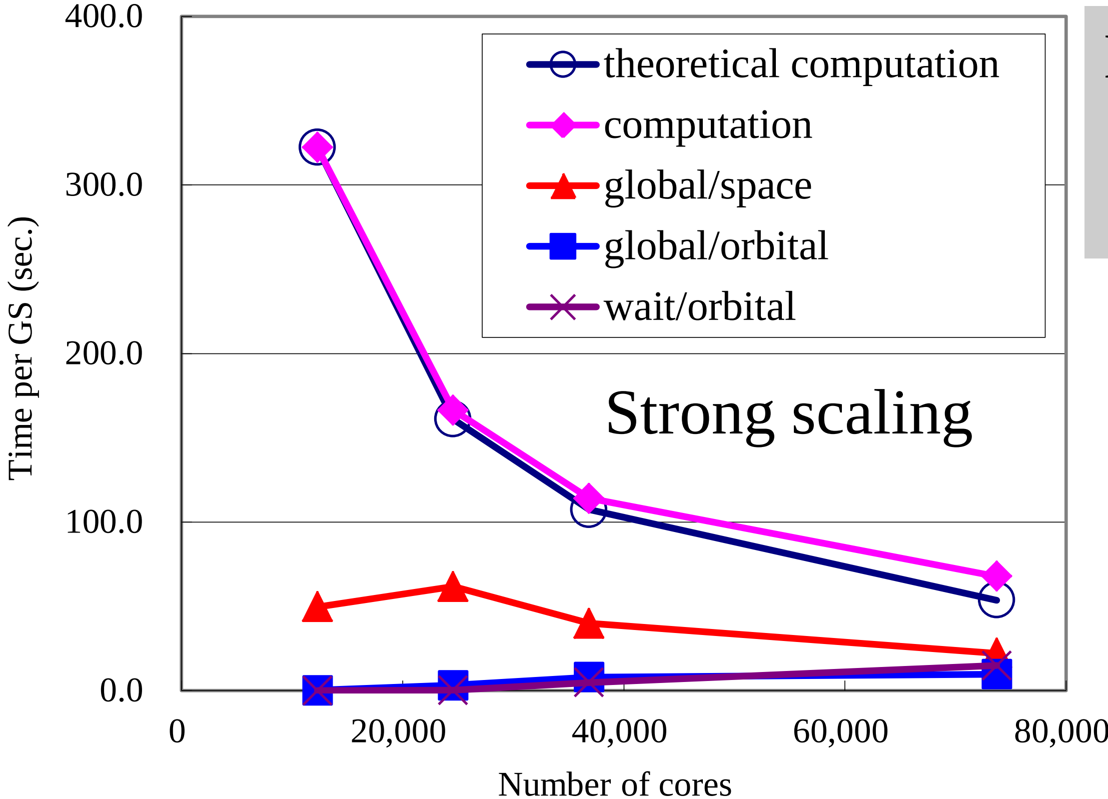

Number of parallel tasks
- Spaces: 1536(fix)
- Orbitals:1, 2, 3 and 6

Scalability is better than our expectation

- Similarly, other procedure blokes also has a good scalability.

Overall Performance: Gordon Bell Prize 2011

- Sustained performance is **3.08 PFLOPS** /SCF.
- **43.6 %** of the peak performance.
- Communication cost is 19% of all execution times.
- One iteration time of SCF is 5500 sec.

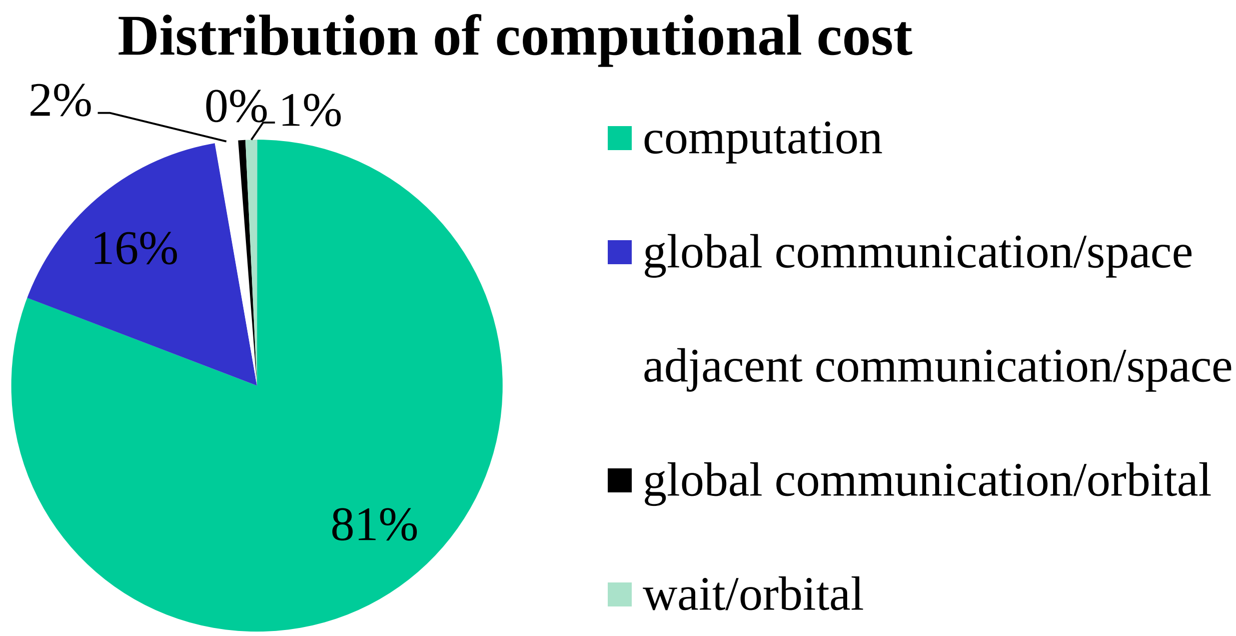

Trillion-body Simulations of Dark Matter Universe

on K-Computer

Ishiyama (Tsukuba), Makino (TiTech), Nitadori (AICS, Riken)

Gordon Bell Prize 2012

Visualization by Takeda (CfCA, NAO)

Dynamic domain decomposition

△ Space filling curve

○ multi section

Allows each node to know easily
where to perform short communications

x equal #particles

△ equal #interactions

◎ equal <#interactions + correction>

= equal calculation time

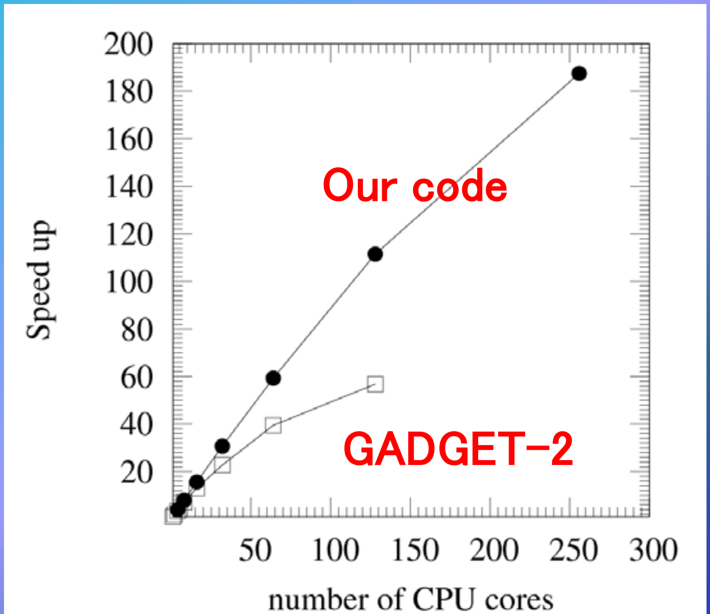

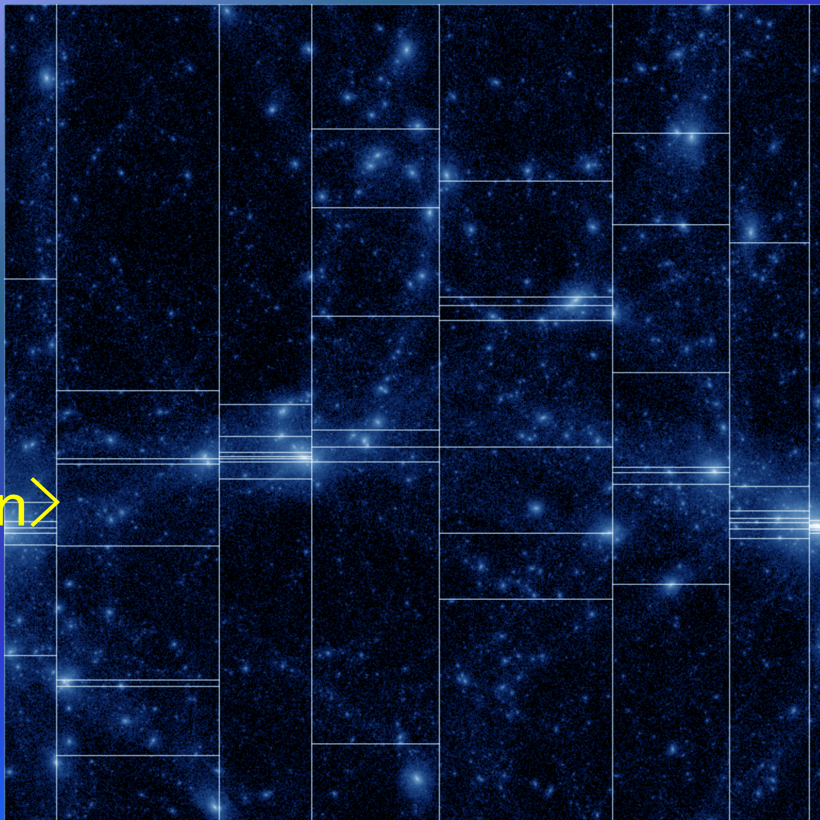

Performance results

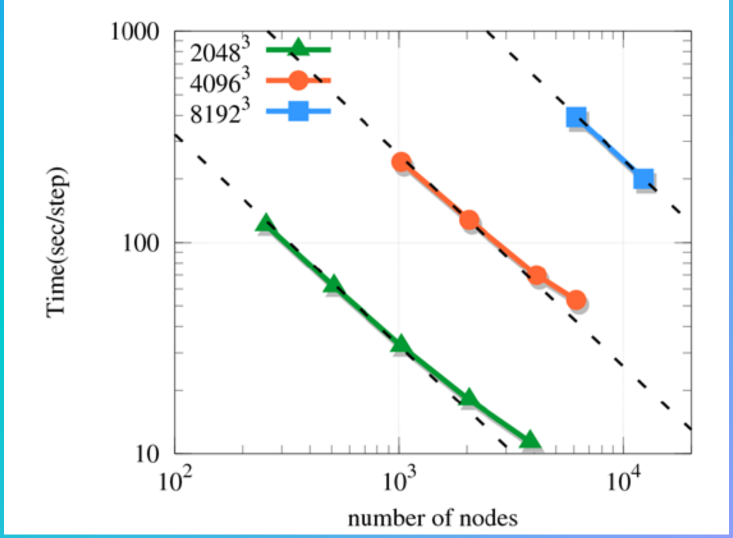

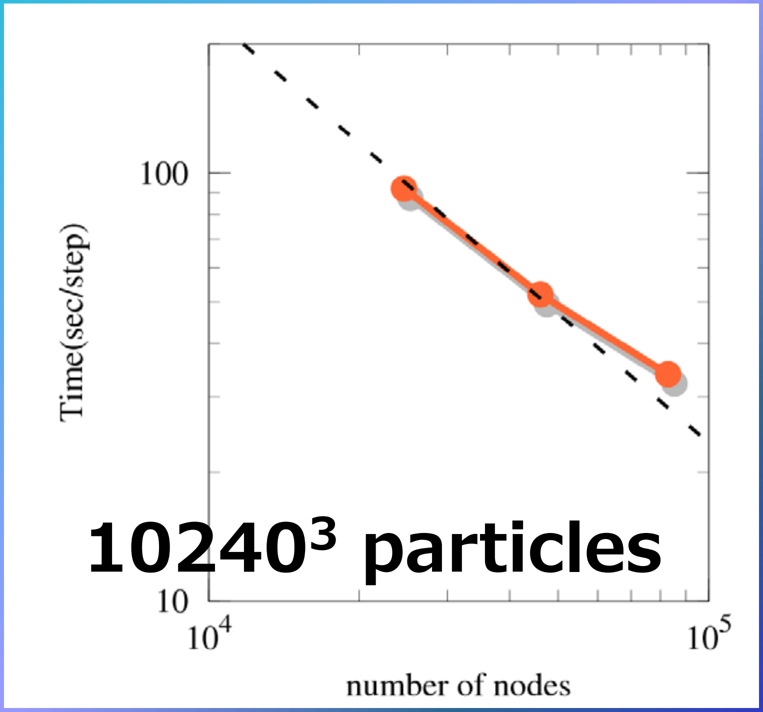

- **Scalability (2048^3 - 10240^3)**
  - Excellent strong scaling
  - 10240^3 simulation is well scaled from 24576 to 82944 (full) nodes of K computer
- **Performance (12600^3)**
  - The average performance on fullsystem is ~**5.67Pflops**, which correspond to ~**55**% of the peak speed

Gordon Bell Prize 2012

Organization for Collaborative Research on Computational Astrobiology

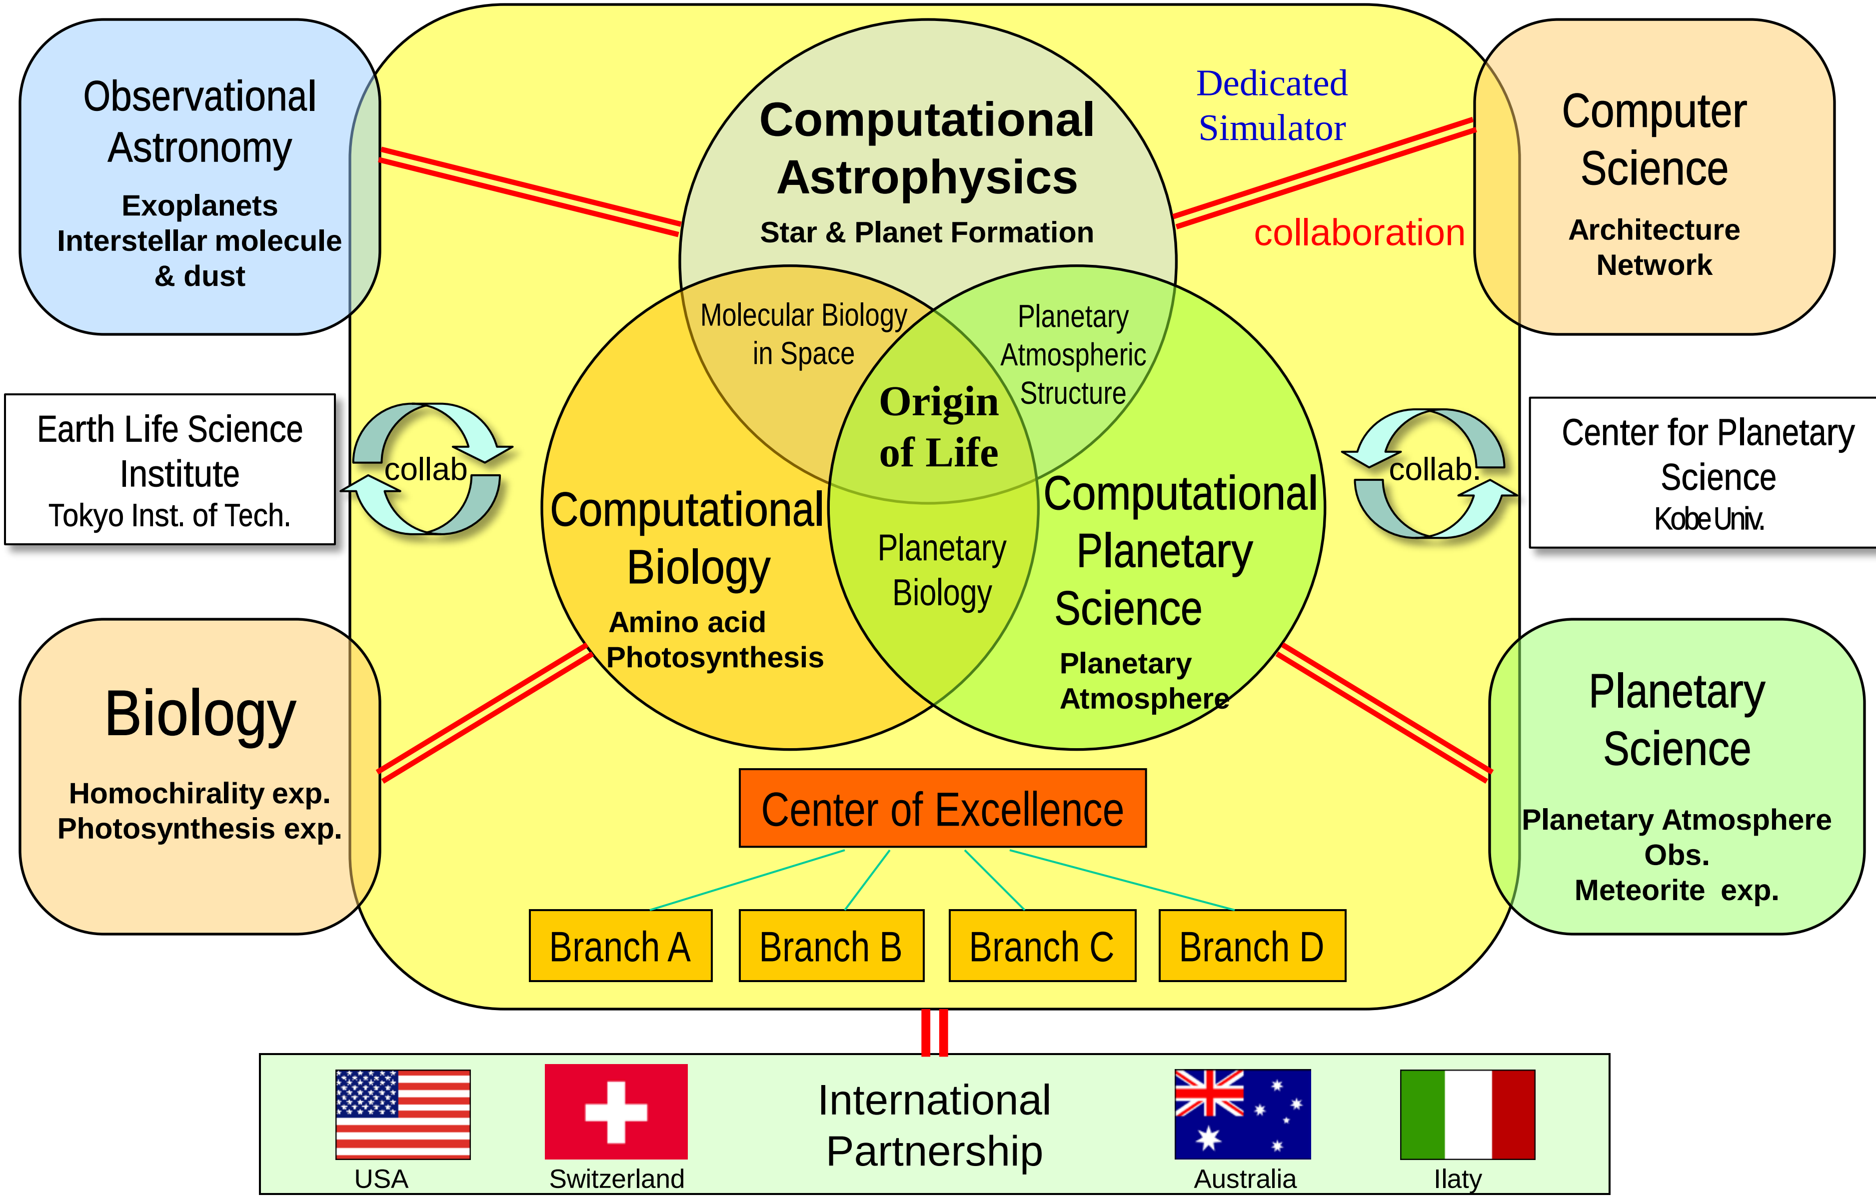

Computational Astrobiology

Collaboration of Astrophysics, Biophysics, and Planetary Science

Cosmic origin of L-amino acid

Enantiomeric excess of Amino Acid by interstellar circular-polarized light

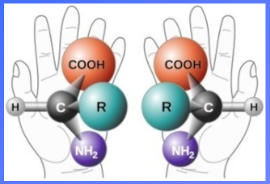

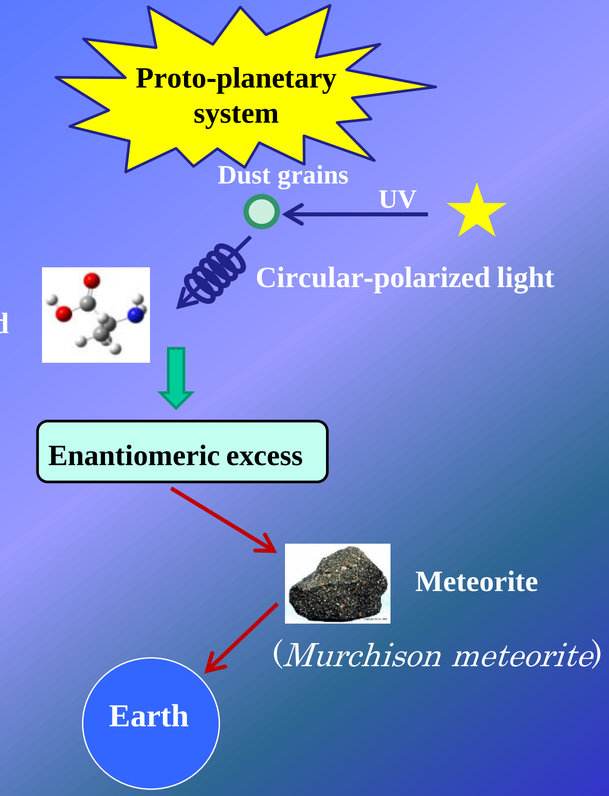

Photosynthesis on extra-solar planets

Red edge as a biomerker

TDDFT calculations of chlorophylls

Chlorophyll dynamics QM/MM