Mission of CCS

The CCS promotes "**multidisciplinary computational science**" on the basis of the fusion between computational science and computer science. For the purpose, the CCS develops high-performance computing systems by the "co-design". The scientific research areas cover particle physics, astrophysics, nuclear physics, nano-science, life science, environmental science, and information science.

The CCS was reorganized in April, 2004, from the preceding center, Center for Computational Physics that was established in 1992. The CCS is the institute for the above-mentioned research fields and also the joint-use facility for outside researchers. Since 2010, the CCS has been approved as a national core-center, Advanced Interdisciplinary Computational Science Collaboration Initiative (AISCI), by the Ministry of Education, Culture, Sports, Science and Technology (MEXT). The CCS aims at playing a significant role for the development of the Multidisciplinary Computational Science.

Oakforest-PACS II at JCAHPC

JCAHPC (Joint Center for Advanced High Performance Computing) has been established in 2013 under agreement between **CCS, Univ. of Tsukuba** and **ITC, Univ. of Tokyo**, to design and manage a next-generation supercomputer system.
JCAHPC has operated Oakforest-PACS supercomputer from 2016 to 2022, which had the highest performance in Japan, to promote advanced computational sciences, and contributes to the field of academia, science, and technology. We are currently in the process of procuring a new **Oakforest-PACS II** supercomputer, which will be operated from April 2024.

Target Minimum Performance

- Peak Performance 200 PFlops
- GPU Node: FP64 240 TFlops, 256 GiB HBM, 9.6 TB/s
- CPU Node: 8 TFlops, 64 GiB HBM, 2.4 TB/s
- Storage
  - Node local NVMe SSD: 3 TB
  - Parallel File System: 25 PB, 1.2 TB/s

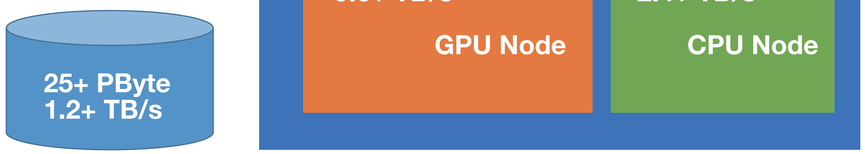

Supercomputer at CCS: Cygnus
Multi-Hybrid Accelerated Computing Platform

Combining goodness of different type of accelerators: GPU + FPGA

- GPU is still an essential accelerator for simple and large degree of parallelism to provide ~10 TFLOPS peak performance
- FPGA is a new type of accelerator for application-specific hardware with programmability and speeded up based on pipelining of calculation
- FPGA is good for external communication between them with advanced high speed interconnection up to 100Gbps x4 chan.

GPU: NVIDIA Tesla V100

FPGA design plan

- Router
  - For the dedicated network, this impl. is mandatory.
  - Forwarding packets to destinations
- User Logic
  - OpenCL kernel runs here.
  - Inter-FPGA comm. can be controlled from OpenCL kernel.
- SL3
  - SerialLite III : Intel FPGA IP
  - Including transceiver modules for Inter-FPGA data transfer.
  - Users don’ t need to care

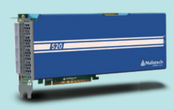

FPGA: Nallatech 520N

Specification of Cygnus

Item	Specification
Peak performance	2.4 PFLOPS DP (GPU: 2.2 PFLOPS. CPU: 0.2 PFLOPS, FPGA: 0.6 PFLOPS SP) ⇨ enhanced by mixed precision and variable precision on FPGA
#of nodes	81 (32 Albireo (GPU+FPGA) nodes, 49 Deneb (GPU-onlu) nodes)
Memory	192 GiB DDR4-2666/node = 256GB/s, 32GiB x 4 for GPU/node = 3.6TB/s
CPU / node	Intel Xeon Gold (SKL) x2 sockets
GPU / node	NVIDIA V100 x4 (PCIe)
FPGA / node	Intel Stratix10 x2 (each with 100Gbps x4 links/FPGA and x8 links/node)
Global File System	Luster, RAID6, 2.5 PB
Interconnection Network	Mellanox InfiniBand HDR100 x4 (two cables of HDR200 / node) 4 TB/s aggregated bandwidth
Programming Language	CPU: C, C++, Fortran, OpneMP, GPU: OpenACC, CUDA, FPGA: OpenCL, Verilog HDL
System Vendor	NEC