

Towards Unification of Computation and Communication

Tightly Coupled Accelerators & Accelerators in Switch

Large-scale Deep Learning using Supercomputer

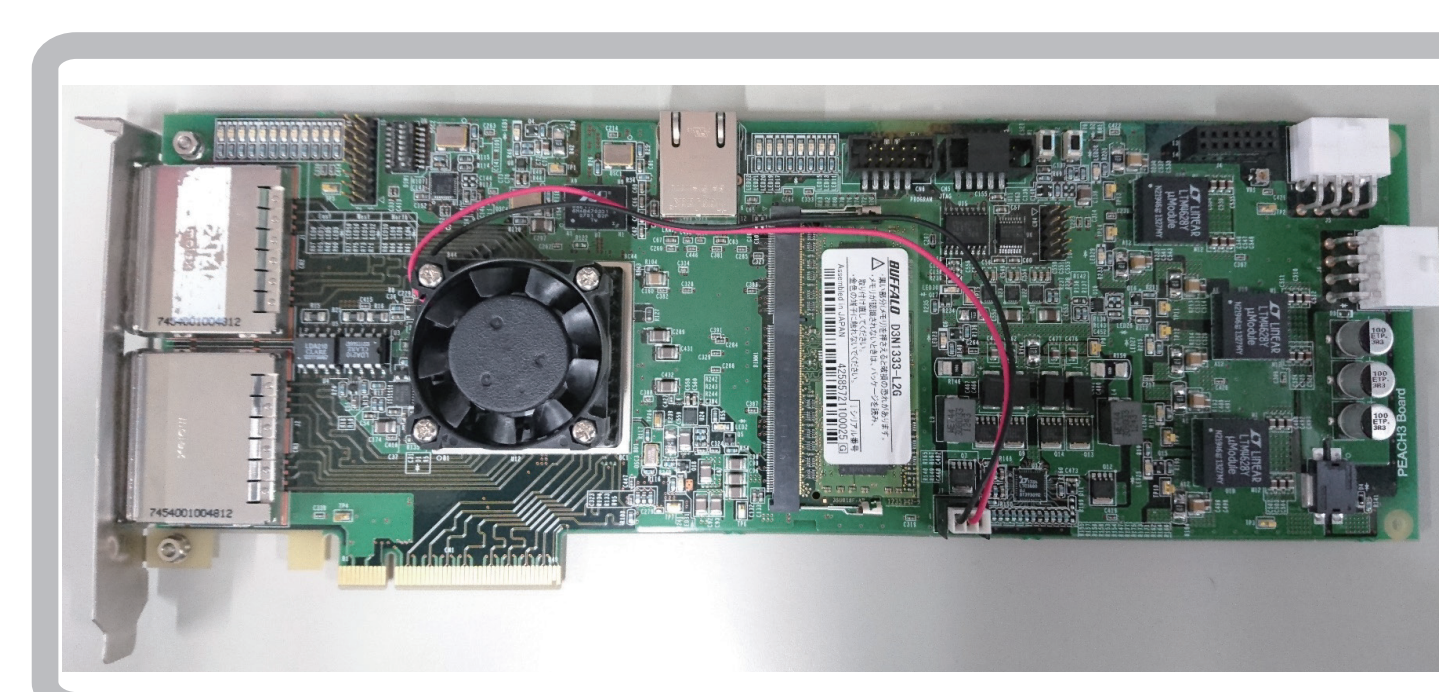
Tightly Coupled Accelerators (TCA) Architecture

GPGPU is now widely used for accelerating scientific and engineering computing to improve performance significantly with less power consumption.

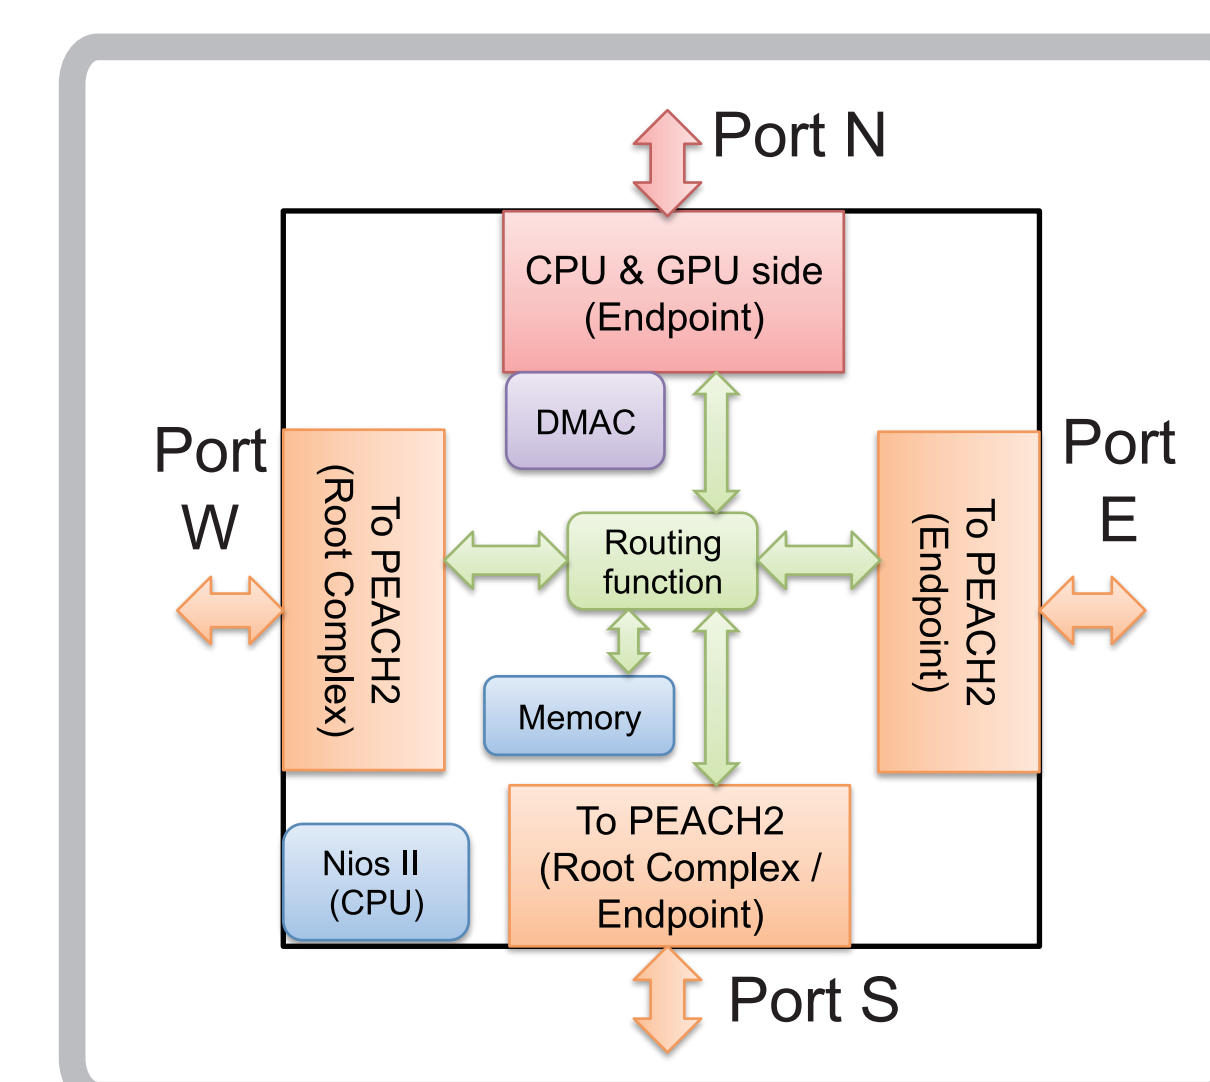
However, I/O bandwidth bottleneck causes serious performance degradation on GPGPU computing. Especially, latency on inter-node GPU communication significantly increases by several memory copies. To solve this problem, **TCA (Tightly Coupled Accelerators)** enables direct communication among multiple GPUs over computation nodes using PCI Express.

PEACH2/3 (PCI Express Adaptive Communication Hub ver. 2/3) chip was developed and it had been evaluated using HA-PACS/TCA cluster, and so on. PEACH2/3 achieved high-performance thanks to PCI Express (PCIe) Gen2/3 x8. However, PEACH2/3 scalability was naturally restricted by PCIe spec.

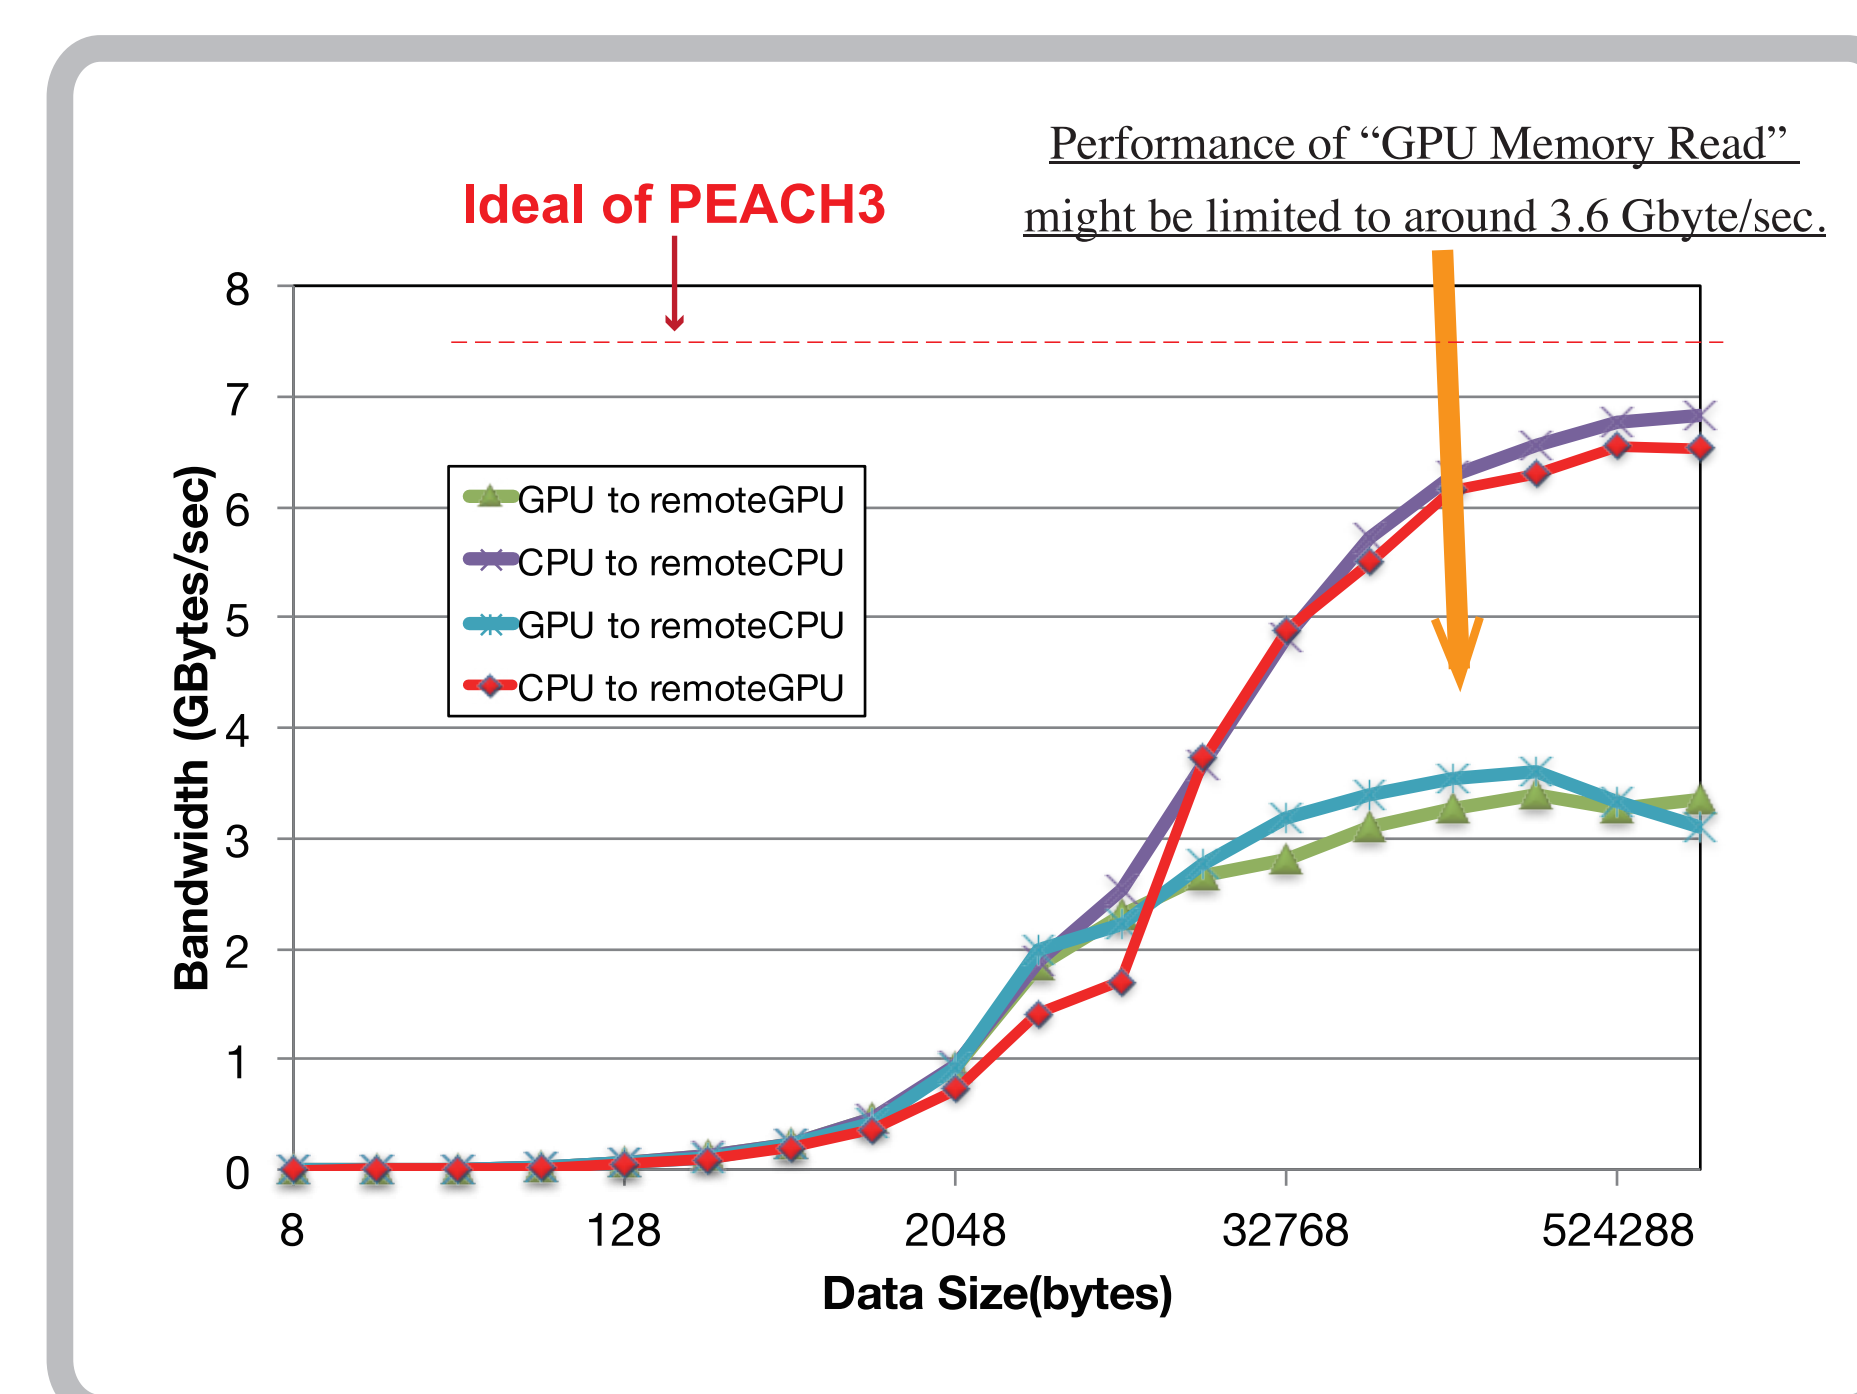
Moreover, FPGA was used only for the communication device in the TCA implementation. Thus, many unused logics in FPGA are helpful as small accelerator with reconfiguration for on-the-fly computation thru CPU and accelerators. Accelerators in Switch (AiS) is the concept to build such accelerators for TCA-like architecture.



PEACH3 Communication Board
(PCIe CEM Spec., single height)



Block diagram of PEACH2/3 Chip



Ping-pong Bandwidth using DMA
(CPU: Intel Xeon E5-2680v2,
GPU: NVIDIA K40 with PCIe Gen3 x16)

Reference

T. Hanawa, et al., "Improving Strong-Scaling on GPU Cluster Based on Tightly Coupled Accelerators Architecture," IEEE Cluster 2015, pp.88-91, 2015.

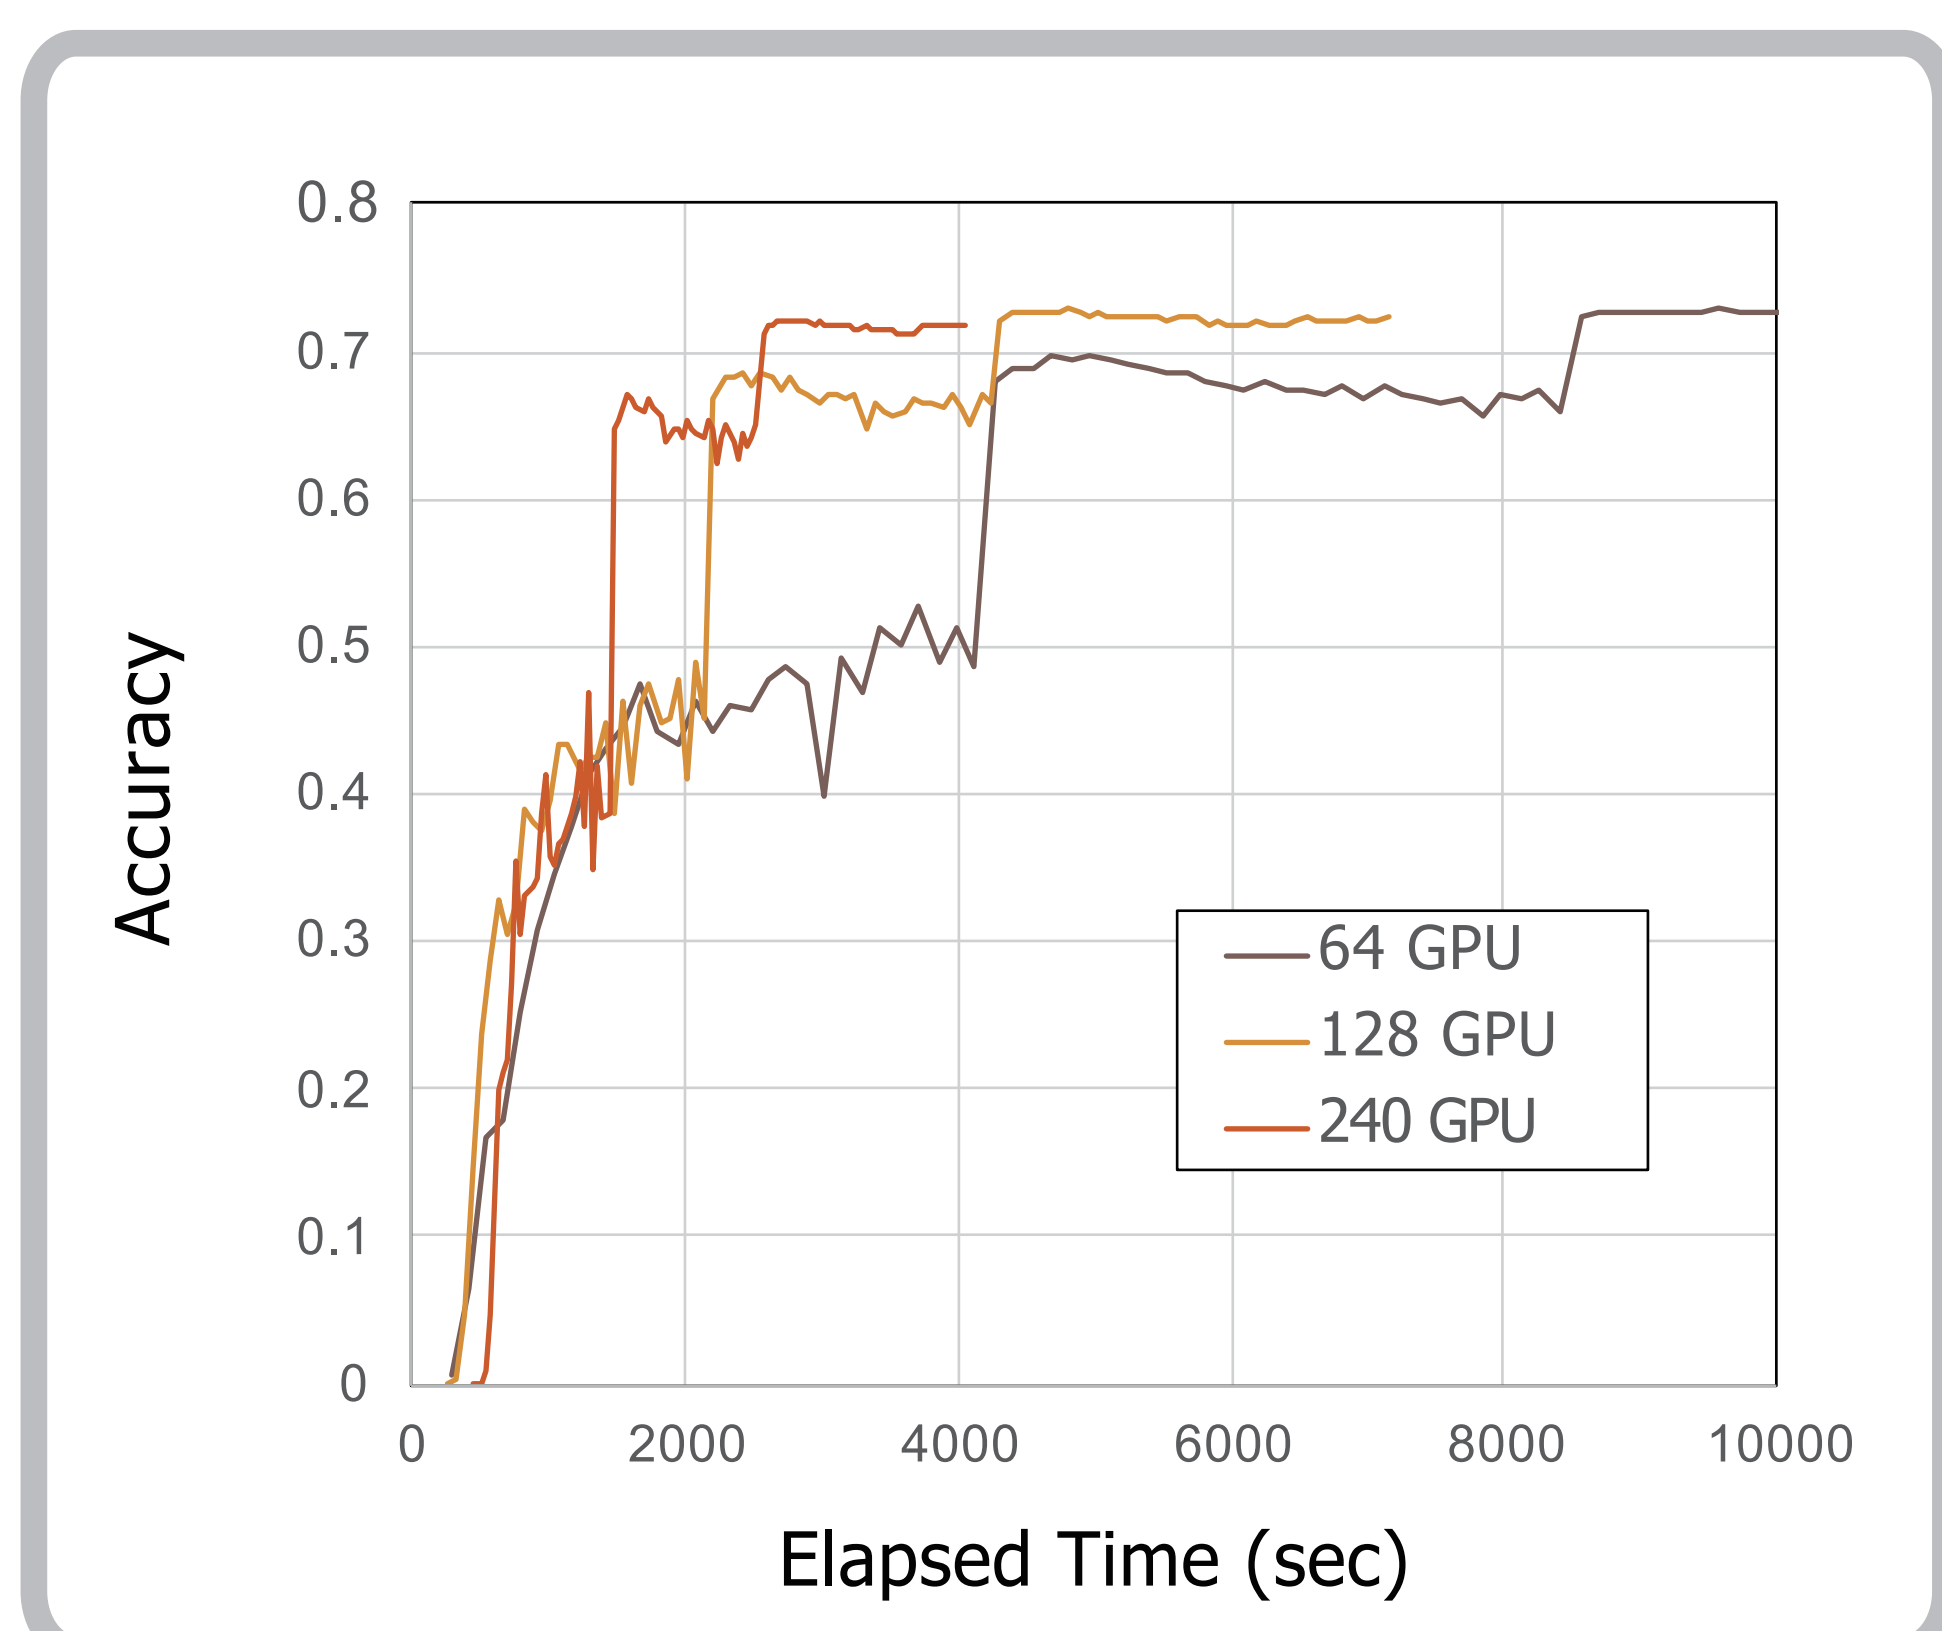
Ongoing & Future Work

- **Accelerators in Switch & New Technologies**
 - OpenCL & other programming method for AiS
 - Intel Xeon+FPGA platform under HARP2 (Hardware Accelerator Research Project)
 - New Interface Technologies like OpenCAPI

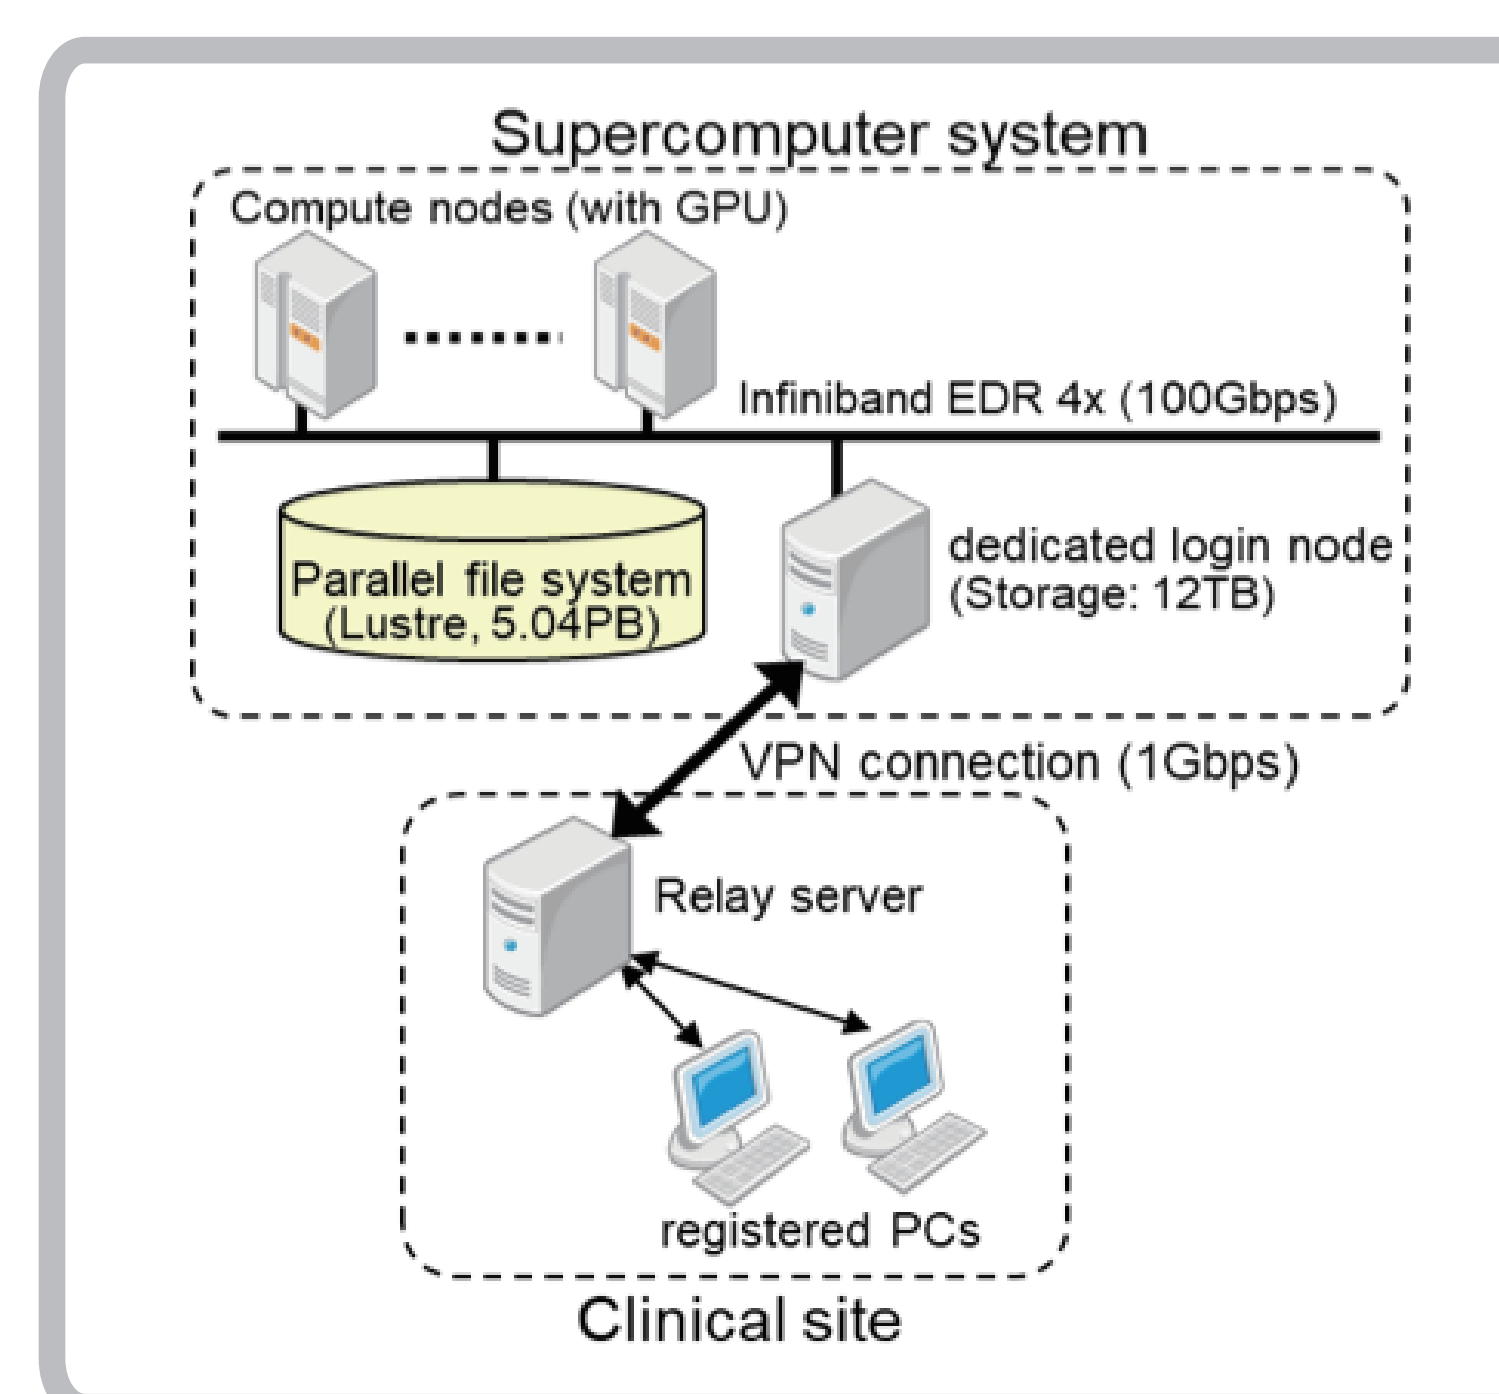
Large-scale Deep Learning using Supercomputer

We collaborate with Univ. Tokyo Hospital to realize high efficient medical image analysis using ITC's supercomputers. Deep learning requires large amounts of computational power, and numerous hyper-parameter optimization has great influence on the performance of deep learning. Thus, we have been developing a framework for training deep learning with hyper-parameter optimization on the supercomputer system.

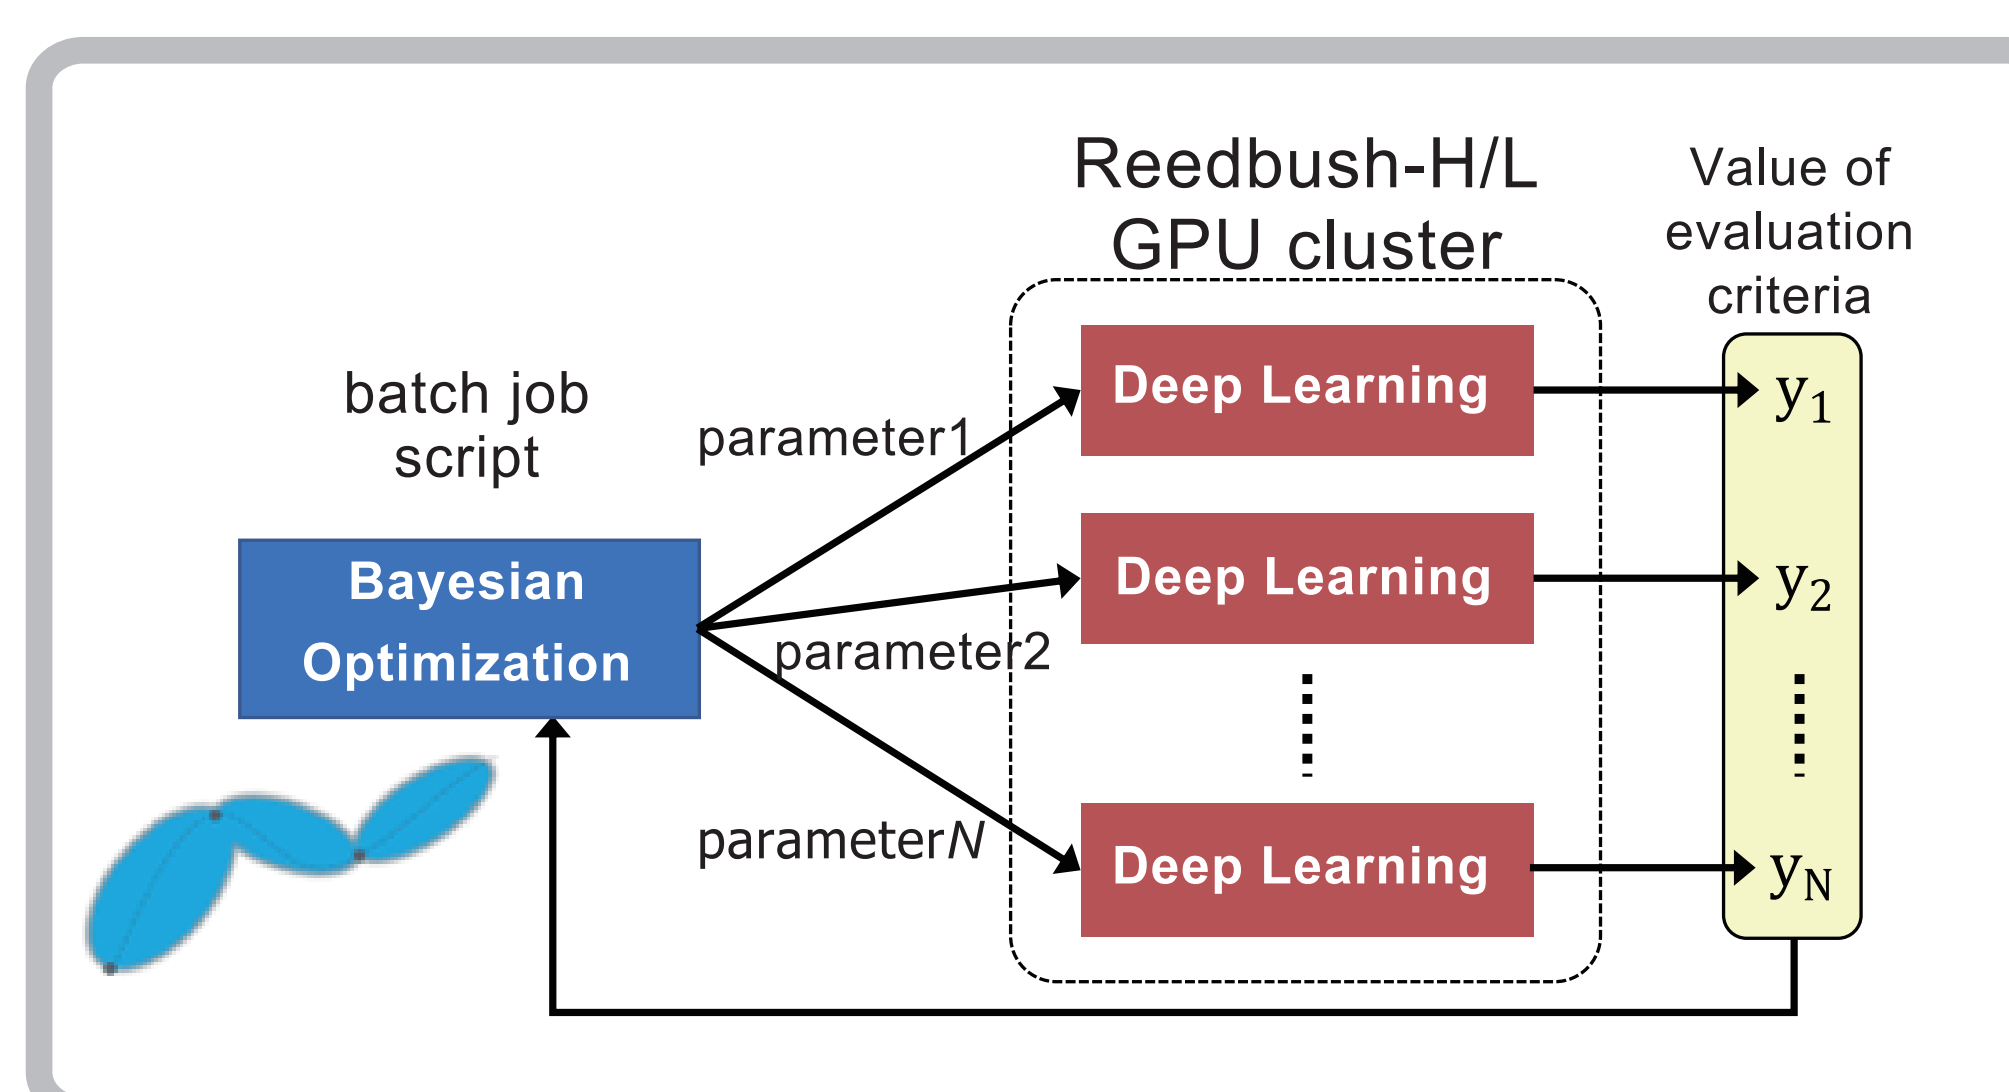
In addition, we are developing large-scale deep learning environment using Oakforest-PACS, which consists of 8208 Intel Xeon Phi compute nodes.



Training of ImageNet with ResNet-50 using Chainer-MN on Reedbush-H system
(Chainer-MN 1.0.0)

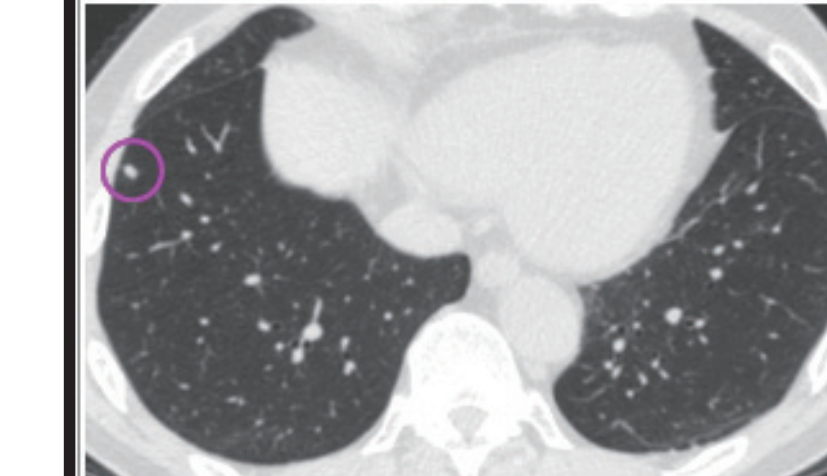


Configuration of Medical Image Analysis System with Reedbush system



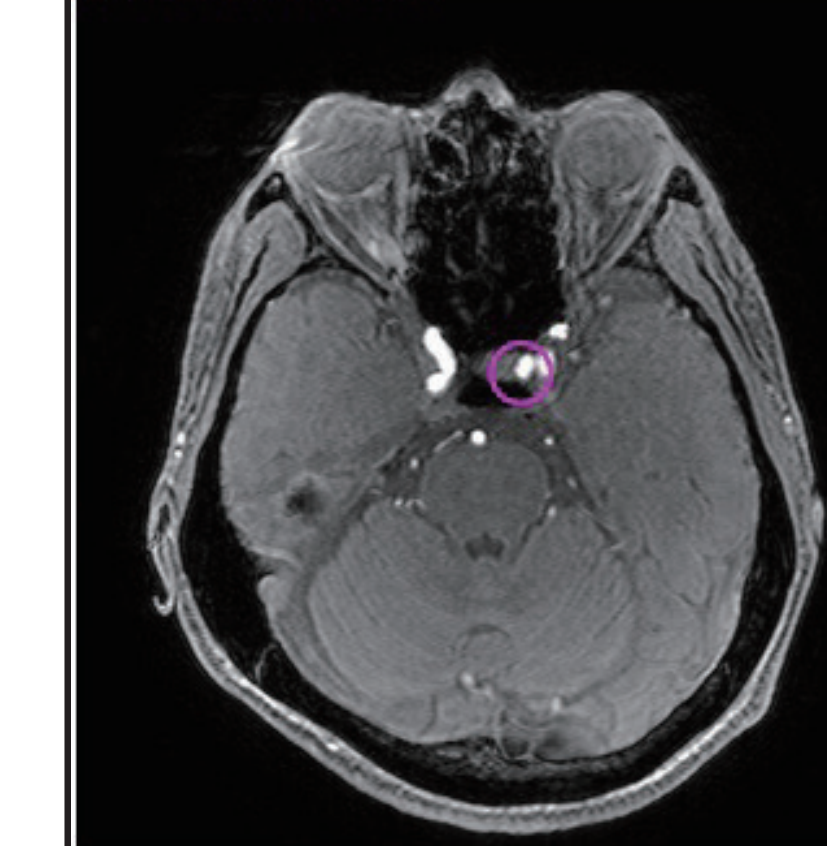
Automated Hyper-parameter Tuning by Bayesian Optimization on Reedbush-H

Image No.: 265
Slice Location: 282.5 [mm]
Volume: 89.91 [mm³]
Confidence: 0.967783



Lung Nodule by Chest CT
(○:detected lesion)

Image No.: 66
Slice Location: -18.1 [mm]
Volume: 56.67 [mm³]
Confidence: 0.529881



Lung Nodule by Head CT
(○:detected lesion)

Examples of Lung Nodule detection
[Courtesy: Dr. Y. Nomura (UTokyo Hospital)]

Reference

Y. Nomura, et al, Preliminary development of training environment for deep learning on supercomputer system, 32nd International Congress and Exhibition on Computer Assisted Radiology (CARS 2018), 2018