Infrastructure for Big Data and AI
System Software and Application Research

Exploring Next-Gen Memory Hierarchy

To achieve higher performance and larger capacity on recent and future architectures, we need to explore next-gen memory hierarchy, including heterogeneous devices. Also placement of many-cores and memory devices can be reconsidered; 3D stacking of cores and memory chips may go mainstream in HPC/Big-data area.

Under such assumptions, understanding memory performance will be more challenging for data and address traffic routed among many-cores, which may often conform 2D mesh.

To analyze its effects, we have conducted preliminary measurements of memory latency from every core of Xeon Phi. We see 9% difference in the current architecture; the effect will be expanded in future architecture where chip cores are also stacked.

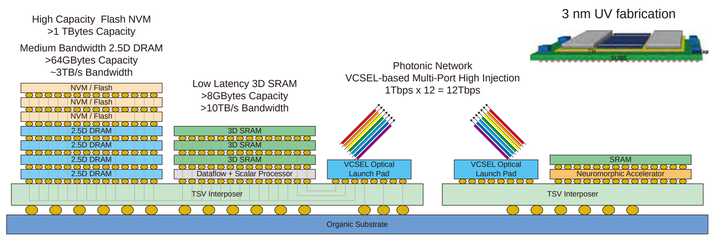

An Example of Future Architecture with 3D Stacking Technology and Heterogeneous Memory

								5 161.2	4 161.2	1 163.1	0 163.2
33 170.9	32 171.0	25 168.1	24 168.1					7 161.5	6 161.5	3 163.8	2 163.7
35 171.3	34 171.4	27 168.2	26 168.3	19 166.5	18 166.5	13 162.9	12 162.9	9 162.0	8 162.0		
		29 168.9	28 168.9	21 166.8	20 166.8	15 163.6		11 162.9	10 162.9		
		31 169.9	30 169.9	23 167.8	22 167.8	17 164.7	16 164.7	41 165.1	40 165.1		
63 174.1	62 174.1	59 171.8	58 171.8	53 168.8	52 168.9	47 165.9	46 165.8	43 166.2	42 166.2	37 167.7	36 167.8
65 175.7	64 175.6	61 172.8	60 172.8	55 170.5	54 170.6	49 167.3	48 167.3	45 167.5	44 167.5	39 168.8	38 168.8
67 176.7	66 176.6			57 171.6	56 171.6	51 168.4	50 168.7				

max: 176.7 (ns)
min: 161.2 (ns)
delta: 15.5 (ns)

Memory access latency from every core of Xeon Phi 7285. As the target of access, one of MCDRAM modules, placed close to core 4, 5 is used.

Acknowledgments. This research is supported by the New Energy and Industrial Technology Development Organization (NEDO), Japan.

Inter-Cloud Infrastructure for Big Data Analysis

Building a Testbed Infrastructure on Overlay Cloud

- Using SINET5 network infrastructure (100Gbps Network)
- Cooperation with Cloud and Supercomputer
- Providing the Science Data Repository
  - Testbed: Petabytes class object storage
  - Real System: Exabytes class object storage

Cloud and Supercomputer Federation

- Development of collaboration technology between cloud and supercomputer using TSUBAME 3.0 as a test case

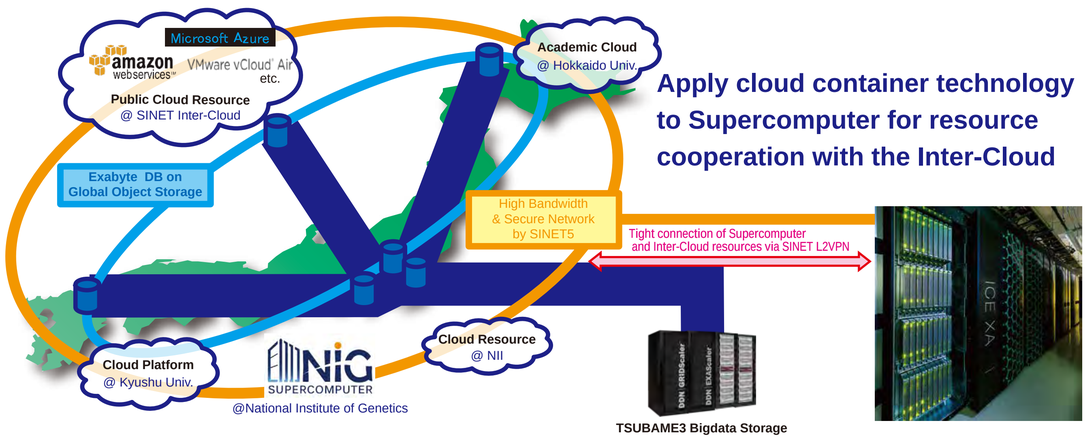

Apply cloud container technology to Supercomputer for resource cooperation with the Inter-Cloud

Acknowledgments. These researches are supported by CREST, JST *CREST Grant Numbers JPMJCR1303 and JPMJCR1501.*
(Research Area: Advanced Core Technologies for Big Data Integration).

Efficient Deep Learning for Video Processing

Fine-to-Coarse Convolutional Neural Network for 3D Human Action Recognition

Current convolutional action recognition models struggle to fully capture the temporal dimension due to the required network depth. To avoid this issue, we investigate a coarse-to-fine network architecture to learn long-distance temporal dependencies. The information of each body part is aggregated while respecting the structure of the body. We achieve state-of-the-art performance and improve it for two person interactions.

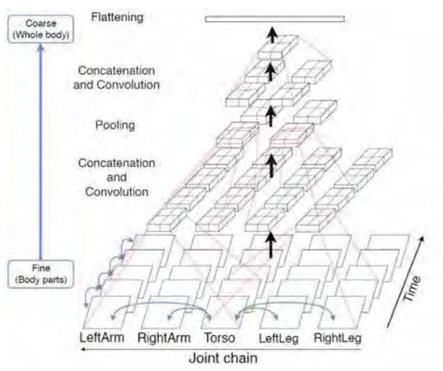

High-resolution Image Reconstruction on Supercomputers

CT is a widely used technology that requires compute-intense algorithms for image reconstruction. We exploit the heterogeneity of GPU-accelerated systems by overlapping the filtering and back-projection stages on CPUs and GPUs, respectively. We also propose a distributed framework for high-resolution image reconstruction on state-of-the-art GPU-accelerated supercomputers. We demonstrate the scalability and instantaneous CT capability of the distributed framework by using up to 2,048 V100 GPUs to solve 4K and 8K problems within 30 seconds and 2 minutes, respectively.

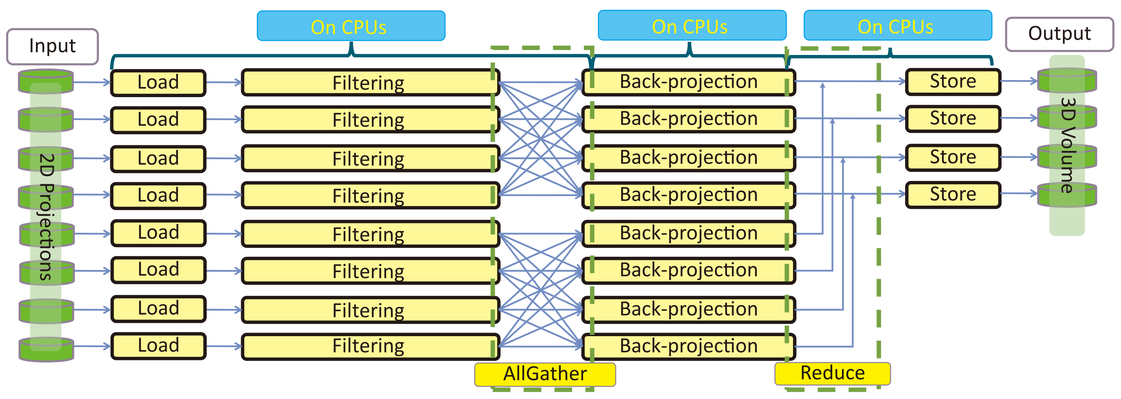

Presentation in SC19 Technical Program
iFDK: A Scalable Framework for Instant High-Resolution Image Reconstruction
Thursday 3:30pm-4pm @ Room 301-302-303

Second Order Optimization for Deep

Distributed deep learning on thousands of GPUs forces the batch stochastic descent methods to operate in a regime where the increasing batch size starts to have a detrimental effect on the convergence and generalization. We investigate the possibility of using second-order optimization methods with proper regularization as an alternative to conventional stochastic gradient descent methods.

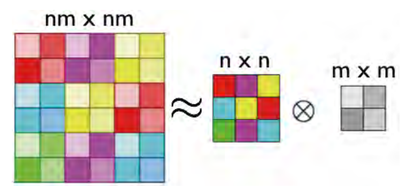

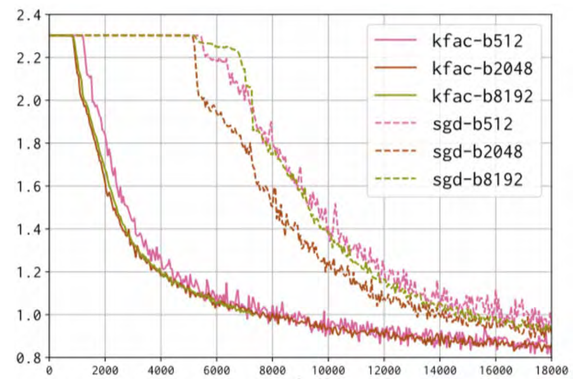

Scene Graph Prediction

Understanding higher-level concepts such as recognizing a scene in a video stream requires algorithms that can leverage information beyond the pixel level. Graphs have been used to represent relationships between objects in a given image. A scene graph is a compact representation of an image that can represent both objects present and relationships between them. We investigate the application of graph neural networks based methods for designing efficient scene graph prediction models.

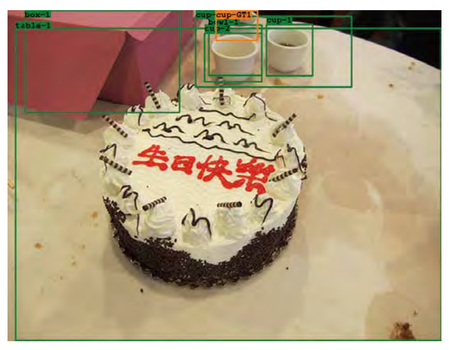

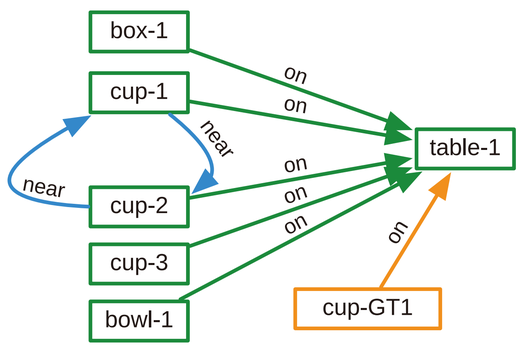

Acknowledgments.
This research is supported by CREST, JST *CREST Grant Number JPMJCR19F5.*

https://www.gsic.titech.ac.jp/sc19