TSUBAME Grand Challenge program

Adopted projects in the spring 2017 program

TSUBAME Grand Challenge Summary

This program is only chance to use all nodes of TSUBAME2.5/3.0 exclusively, because TSUBAME2.5/3.0 is shared by thousands of users. There are two categories:

- Category A: The large scale application aims high peak-performance. All of TSUBAME2.5/3.0 nodes are available.
- Category B: The large scale application aims scientifically meaningful results. A large portion (1/3) of TSUBAME2.5 is available.

Table Number of Adopted Projects in the TSUBAME Grand-Challenge Program

Category	FY2017		FY2016		FY2015		FY2014		FY2013		FY2012		FY2011		Total
	Fall	Spr.	Fall	Spr.	Fall	Spr.	Fall	Spr.	Fall	Spr.	Fall	Spr.	Fall	Spr.	
A	2	0	1	1	1	2	1	2	0	1	2	2	3	4	22
B	0	1	0	1	1	3	2	2	1	1	0	0	2	-	14
Total	2	1	1	2	2	5	3	4	1	2	2	2	5	4	36

We started this program since FY2011, and keep on carrying out twice in each year.
Under this program, we have adopted total 36 fruitful projects, some of which were awarded prizes as below.

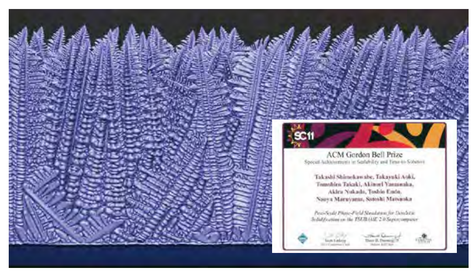

2011 ACM Gordon Bell Prize: Special Achievements in Scalability and Time-to-Solution
"Peta-scale Phase-Field Simulation for Dendritic Solidification on the TSUBAME 2.0 Supercomputer"

2011 ACM Gordon Bell Prize: Honorable Mention
Large scale biofluidics simulations on TSUBAME2

2011 Graph500 Challenge on TSUBAME 2.0

The First Adopted Projects on TSUBAME3.0 in Fall 2017

Deep Learning for Fusion Simulation Using the Fusion Recurrent Neural Network

FRNN (Fusion Recurrent Neural Net) is the deep learning framework developed for fusion energy applications. It is implementing a distributed data-parallel approach to train deep neural networks (in particular, stacked LSTMs). With this approach, a replica of the model is kept on each worker, processing different mini-batches of the training dataset in parallel. The model parameters from each worker are collected using MPI, and synchronized via parameter averaging with learning rate adjusted after each epoch to improve convergence. This produces a global set of parameters, which are then broadcasted to each model replica. The stochastic gradient descent (SGD) method is used for large-scale optimization with parallelization via mini-batch training to reduce communication costs. In FRNN, we integrate Keras (TensorFlow and Theano backends) and MPI to enable training across multiple GPU nodes using high-speed interconnects.

The starting phase of this proposed 2017 TSUBAME 3.0 Grand Challenge project will be to initiate scaling studies on the powerful TSUBAME 3.0 supercomputer featuring thousands of Pascal P-100 GPU's will begin with the goal of replicating the excellent FRNN scaling observed on "Titan" as well as Pascal P-100 systems at NVIDIA and Princeton University. We look forward to exploring challenges associated with engaging the full capability of TSUBAME 3.0. In particular, since the FRNN software uses the steepest gradient descent (SGD) with mini-batch training to reduce communication costs, it will be very interesting to observe the behavior of this DL software with scaling studies on TSUBAME 3.0 to examine if the convergence rate saturates/decreases with increasing mini-batch size to thousands of P-100 GPU's.

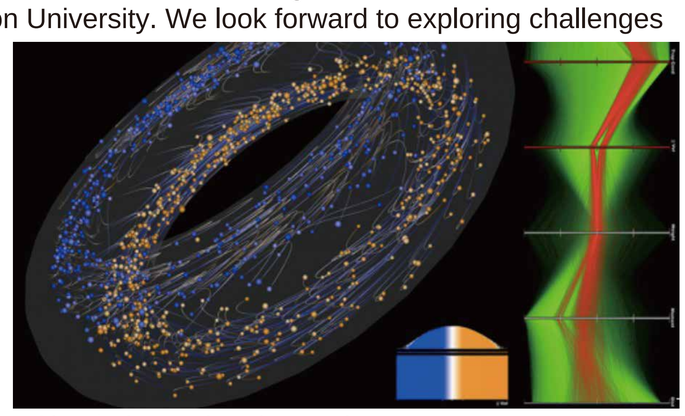

Fusion Simulation Using the Fusion Recurrent Neural Network

Realizing Extreme Large-Scale Matrix Computations with a Memory Management Library Utilizing Fast Flash SSDs

The objective of our project is to achieve extremely high-speed and large-scale matrix computation on TSUBAME3.0. Here we focus on large scale Cholesky decomposition, which is an important kernel in solving semi-definite programming (SDP) problem. In order to support even larger matrices than aggregated host memory capacity of available nodes, we harness large capacity of NVMe SSDs, which is 8 times larger than DRAM capacity in TSUBAME3.0. This should be achieved with less development effort; thus we have developed a global address space runtime library, named vGASNet. On vGASNet, the aggregated capacity of SSDs distributed among compute nodes are virtually visible as a single address space. For high-performance computing, caching mechanism using DRAM is necessary. Moreover, for scalable data movement, it supports "cooperative caching mechanism" to avoid bottleneck in access congestion. On top of vGASNet, we also have implemented the parallel Cholesky decomposition algorithm. While it is based on data-driven execution and tile-based dynamic scheduling to suppress global synchronization, the implementation is relatively simple and clean owing to the global address space view. Inside each task, we use CUBLAS and MAGMA library in order to harness high performance of NVIDIA P100 GPUs. Taking this co-design approach, we aim to achieve peta-scale matrix computation.

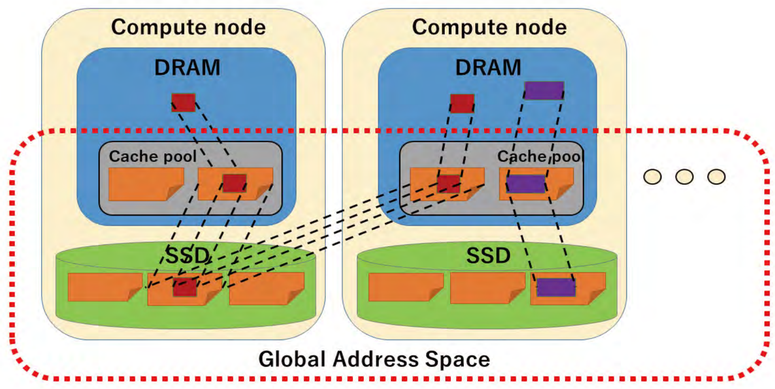

Global Address Space Model of vGASNet

DC-DFTB-metadynamics for primary proton transfer of bacteriorhodopsin

Minori Imai[1], Junichi Ono[1], Yoshifumi Nishimura[2], Hiromi Nakai[1-4]
([1]Grad. Sch. of Adv. Sci. & Eng., Waseda Univ., [2]RISE, Waseda Univ., [3]JST-CREST, [4]ESICB, Kyoto Univ.)

Bacteriorhodopsin: light-driven proton pump

Bacteriorhodopsin (BR) is one of the most famous photoreceptive proteins, which actively translocates proton across the membrane with the aid of light (Fig. 1a) [1]. The resulting proton gradient can produce ATP, i.e., energy resources in biological systems. Thus, elucidation of microscopic mechanism of the proton transfers in BR is important for understanding the biological function of the light-energy conversion.

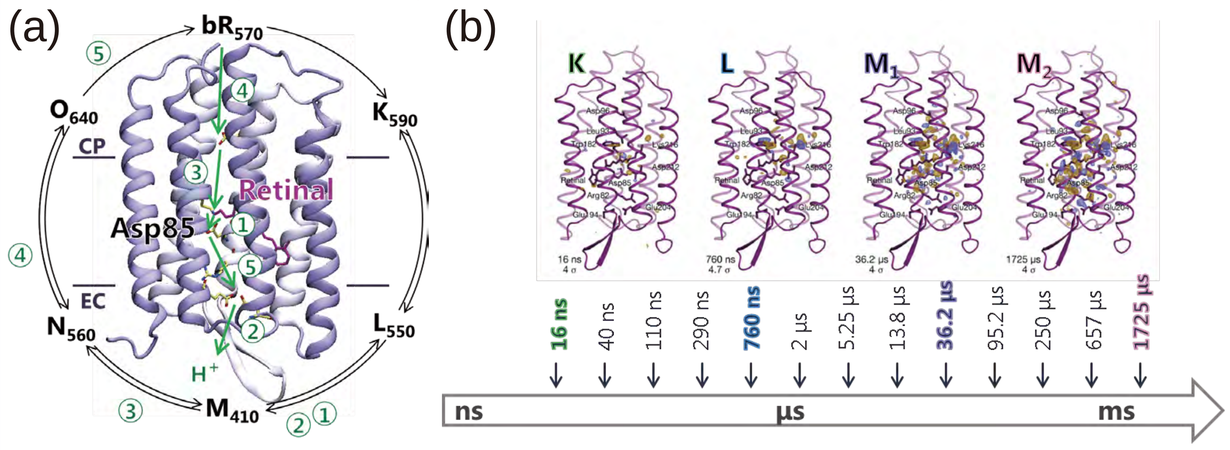

Figure 1. (a) Proton transfers of BR and (b) heavy-atom snapshots of BR captured with X-ray free electron laser along the photocycle [2].

[1] O. P. Ernst, *et al*., *Chem. Rev*., **114**, 126 (2014).
[2] E. Nango *et al*., *Science*, **345**, 1552 (2016).

Multiple pathways of primary proton transfer

In 2016, successive snapshots of BR along the photocycle captured with X-ray free electron laser are reported (Fig. 1b) [2]. Due to the limitation of the experiment, however, only heavy-atom positions are determined. Therefore, no proton transfer has been observed, and the pathway and the timing of the proton transfers have still been unclear.

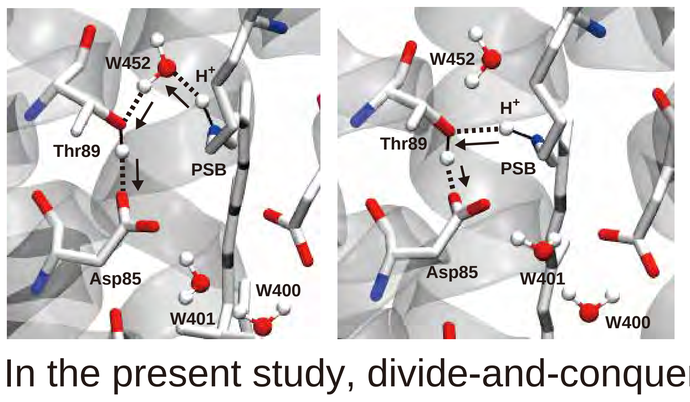

Figure 2. Multiple pathways of the primary proton transfer from the protonated Schiff base (PSB) to Asp85 on the photocycle of BR: (a) path A via water (W452) and Thr89, and (b) path B only via Thr89.

In the present study, divide-and-conquer density-functional tight-binding (**DC-DFTB**) metadynamics simulations were performed for the primary proton transfer, starting from the snapshots. It is found that **the proton transer can occur via two pathways**: path A and B (Fig. 2). Moreover, the former two L-state structures (**760 ns** & **2 µs** in Fig. 1b) are found to be relevant for the primary proton transfer.