GPU-accelerated calculation of electrostatic interactions with zero-dipole summation method

Tadaaki Mashimo[1,2], Yoshifumi Fukunishi [2], Ikuo Fukuda[3,4], Narutoshi Kamiya[4], Yu Takano[4], Haruki Nakamura[4]

1 Research & Development Dept., JBIC, Aomi, Koto-ku, Tokyo 135-0064, Japan
2 BIRC, AIST, 2-3-26, Aomi, Koto-ku, Tokyo 135-0064, Japan
3 Computational Science Research Program, RIKEN, 2-1 Hirosawa, Wako, Saitama 351-0198, Japan
4 Institute for Protein Research, Osaka University, 3-2 Yamadaoka, Suita, Osaka 565-0871, Japan

A major bottleneck in molecular dynamics (MD) simulation is the calculation of electrostatic interactions. Our recently developed zero-dipole summation method[1], in which electrostatic interactions are calculated only in a neutralized subset (total charge = zero, total dipole = zero) in a given cutoff sphere, along with a truncation of the reciprocal space term of the Ewald method, enables the calculation of electrostatic interactions with high accuracy and low computational cost. The algorithm of this method is represented by a simple pair-wise function sum and has high portability between CPU and GPU (general purpose Graphics Processing Unit), which is expected to provide the high computational performance for CPU/GPU applications.

We implemented the zero-dipole summation method into our spatial subdivision MD code, myPresto[2]/psygene, and developed its GPU version (psygene-G) with a GPU port of the pair-wise non-bonded interaction calculations which consist of the van der Waals interaction and the zero-dipole summation method calculation. The GPU port subdivides assigned subspaces, executes local search, and calculates the pair-wise non-bonded interactions in the cutoff as is used in the GPU-based smoothed particle hydrodynamics fluid simulation[3].

We carried out tests of simulation quality and performance of myPresto/psygene-G with MD simulations for three protein systems: EGFR (38453atoms), AQP4 (104415atoms) and Dynein (1004847atoms). These tests were performed on HA-PACS[4] and OLABB GPU cluster which consists of 7 cluster nodes each with 4 Tesla M2050/M2070, 2 intel Xeon5600 (6cores), 24GB memory, CentOS5.6 (64bit), CUDA4.0, intel compiler11.1, and intel MPI4.0.2.

As a result of these tests, relative RMS force error and relative energy error fluctuation are within acceptable numerical accuracy, and sufficient acceleration performance is confirmed by these scale MD simulations on HA-PACS.

This work was supported by NEDO & METI of Japan.

Relative RMS Force errors (8, 27, 64 GPUs)

Protein	Atoms	Relative RMS Force Error
EGFR	38453	2.7455×10^{-6}
AQP4	104415	3.2896×10^{-6}
Dynein	1004847	1.2384×10^{-5}

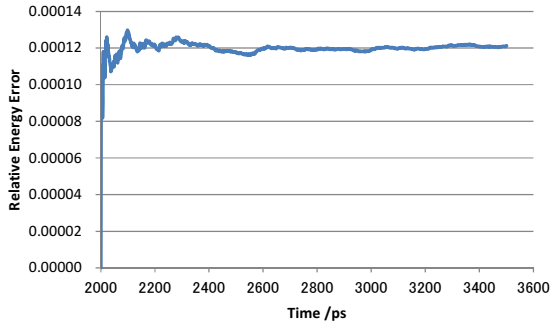

Relative Total Energy error (EGFR, $E_0 = E_{2ns}$)

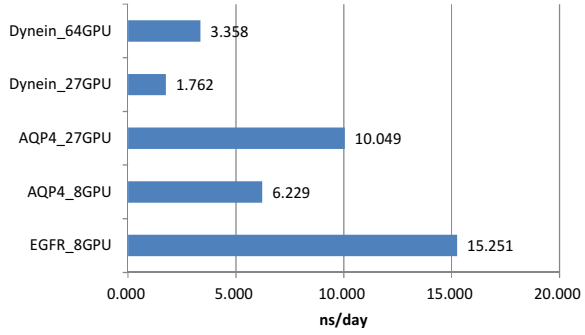

Psygene-G Performance on HA-PACS

References

1) I. Fukuda, Y. Yonezawa, and H. Nakamura, J. Chem. Phys. 134, 164107 (2011).
2) myPresto: [Osaka-Univ.] http://presto.protein.osaka-u.ac.jp/myPresto4, [MEDALS] http://medals.jp/myPresto/
3) http://sa08.idav.ucdavis.edu/CUDA_physx_fluids.Harris.pdf
4) http://www.ccs.tsukuba.ac.jp/CCS/research/project/ha-pacs/cluster