Supporting Information:

Scalable Parallel Algorithm for Graph Neural Network Interatomic Potentials in Molecular Dynamics Simulations

Yutack Park,[†] Jaesun Kim,[†] Seungwoo Hwang,[†] and Seungwu Han[*,†,‡]

[†] *Department of Materials Science and Engineering and Research Institute of Advanced Materials, Seoul National University, Seoul 08826, Korea*

[‡] *Korea Institute for Advanced Study, Seoul 02455, Korea*

E-mail: hansw@snu.ac.kr

GPU utilization with respect to the number of atoms

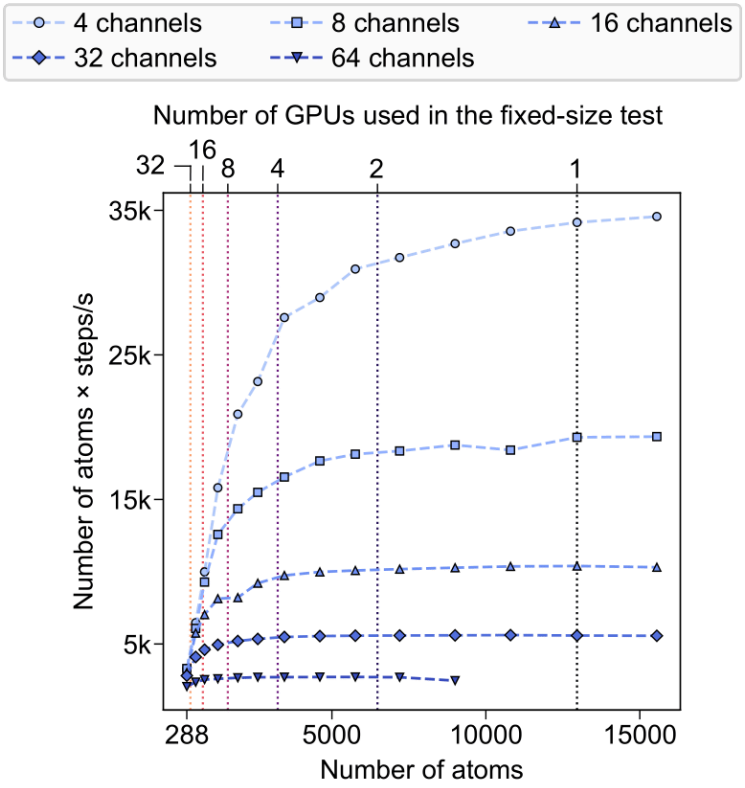

Figure S1. Single GPU utilization curve of models with four message-passing layers and 4, 8, 16, 32, and 64 channels. The dotted lines indicate the number of atoms used in the fixed-size test.

Figure S1 shows how the different number of channels and the number of atoms affects GPU utilization. The y-axis, the number of atoms multiplied by timesteps per second, indicates the amount of work (number of atoms multiplied by MD step) a GPU can yield in a unit of time (second). A plateau in this figure signifies full GPU utilization, achieved when a sufficient number of atoms is allocated. Conversely, a ballistic region observed at lower atom counts represents the GPU is in a suboptimal state. Notably, as the number of channels decreases, a greater number of atoms is needed to transition from this suboptimal state to the fully utilized state.

Single GPU performance of SevenNet and NequIP

Considering the original NequIP is the baseline of our implementation, we compare a single GPU performance of SevenNet and NequIP. We use the same model hyperparameters of 32 channels, 4 message-passing layers, a maximum degree of representation ($l_{\max}$) of 3, and a cutoff radius of 4.0 Å. The same benchmark script used in the scaled-size tests section is utilized. The benchmark runs span 210 MD steps, and we measure wall-clock time over the final 100 steps. The benchmark system involves alpha-quartz SiO_2 containing 4,608 atoms. SevenNet, and NequIP achieved 1.99, and 1.12 timesteps per second, respectively.

For NequIP, we used the original implementation for training[1] and inference[2]. The SevenNet pair style is compiled with LAMMPS (23 Jun 2022 - Update 4), while the NequIP pair style uses a different version of LAMMPS (29 Sep 2021 - Update 2). Both pair styles utilize LibTorch from PyTorch/1.12.1. The compilation was done with nvcc compiler from CUDA/11.4, and gcc/8.3.0. The benchmarks are conducted on NVIDIA A100 80GB GPU in the KISTI Neuron cluster.

Scaling of SevenNet and MACE

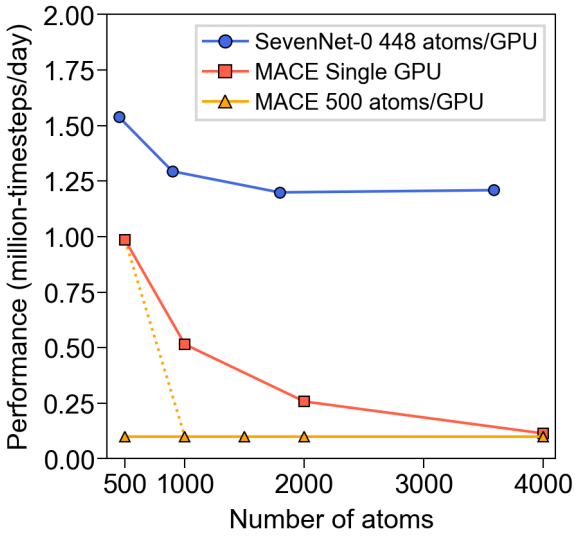

Figure S2. Weak-scaling performance of SevenNet-0 and MACE-MP-0[3]. The blue circles represent the performance of SevenNet in 448-atom amorphous Si_3N_4, from single to 8 GPUs. For MACE-MP-0, the red rectangle and yellow triangles indicate single-GPU and multi-GPU performance, respectively, measured with the high-entropy alloy system (500 atoms per GPU). The dashed line signifies the transition to multi-GPU in MACE. It is seen that the multi-GPU implementation of MACE becomes favorable only for very large-scale simulations due to significant additional costs. In contrast, SevenNet-0 well maintains the parallel efficiency regardless of the GPU number.

References

(1) Batzner, S.; Musaelian, A.; Sun, L.; Geiger, M.; Mailoa, J. P.; Kornbluth, M.; Molinari, N.; Smidt, T. E.; Kozinsky, B. *NequIP*, v0.5.6; Github:https://github.com/mir-group/nequip (accessed 2024-04-08).

(2) Batzner, S.; Musaelian, A.; Sun, L.; Geiger, M.; Mailoa, J. P.; Kornbluth, M.; Molinari, N.; Smidt, T. E.; Kozinsky, B. *Pair-NequIP,* v0.5.2; Github:https://github.com/mir-group/pair_nequip (accessed 2024-04-08).

(3) Batatia, I.; Benner, P.; Chiang, Y.; Elena, A. M.; Kovács, D. P.; Riebesell, J.; Advincula, X. R.; Asta, M.; Baldwin, W. J.; Bernstein, N.; Bhowmik, A.; Blau, S. M.; Cărare, V.; Darby, J. P.; De, S.; Pia, F. D.; Deringer, V. L.; Elijošius, R.; ElMachachi, Z.; Fako, E.; Ferrari, A. C.; Genreith-Schriever, A.; George, J.; Goodall, R.E. A.; Grey, C. P.; Han, S.; Handley, W.; Heenen, H. H.; Hermansson, K.; Holm, C.; Jaafar, J.; Hofmann, S.; Jakob, K. S.; Jung, H.; Kapil, V.; Kaplan, A. D.; Karimitari, N.; Kroupa, N.; Kullgren, J.; Kuner, M. C.; Kuryla, D.; Liepuoniute, G.; Margraf, J. T.; Magdău, I.-B.; Michaelides, A.; Moore, J. H.; Naik, A. A.; Niblett, S. P.; Norwood, S. W.; O'Neill, N.; Ortner, C.; Persson, K. A.; Reuter, K.; Rosen, A. S.; Schaaf, L. L.; Schran, C.; Sivonxay, E.; Stenczel, T. K.; Svahn, V.; Sutton, C.; Oord, C.v. d.; Varga-Umbrich, E.; Vegge, T.; Vondrák, M.; Wang, Y.; Witt, W. C.; Zills, F.; Csányi, G. A foundation model for atomistic materials chemistry. *arXiv* **2023**, arXiv:2401.00096. https://arxiv.org/abs/2401.00096 arXiv.org ePrint archive