
Eco-Comp: Towards Responsible Computing in Materials Science

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Bridging the time and length scales and the use of large molecular dynamics (MD)
2 simulations in material science is expected to surge in the next few years, parti-
3 tially due to the development of highly accurate machine learning inter-atomic
4 potentials that enable the simulation of multi-million atomic systems. We also
5 expect a high demand for material science simulations using multiple nodes within
6 high-performance computing facilities (HPCs) due to their computational intensity.
7 Through the analysis of catalysis simulation setups consisting of bulk metallic
8 systems with adsorbed molecular species on the surface, we identified various
9 factors that affect parallel computing efficiency. To foster sustainable and ethical
10 computing practices, this study employs the Large-scale Atomic/Molecular Mas-
11 sively Parallel Simulator (LAMMPS) to find the optimal allocation of computing
12 resources based on the simulation input. We thus propose guidelines to promote
13 responsible computing within HPC architecture: Eco-Comp is a user-friendly
14 automated Python tool that allows material scientists to optimize the power con-
15 sumption of their simulations using one command. This tutorial gives a broad
16 overview of the Eco-Comp software and its potential use for the material science
17 community through an interactive guide.

1 Introduction

19 High-Performance Computing, or HPC, has revolutionized the field of science, enabling researchers
20 and scientists to perform simulations that were once impossible. Molecular dynamics (MD) simulation
21 has become a key aspect of focus in material science, as HPC are utilized to simulate, for example,
22 catalytic reactions, enzyme active site exploration, and mechanical properties on the atomic scale. The
23 emergence of easy-to-use machine learning potentials (MLP), such as the Machine-Learned Spectral
24 Neighbor Analysis Potential (ML-SNAP) is indicative of the growing accessibility and versatility of
25 this tool in the scientific community. It allows scientists and engineers to understand key chemical
26 reactions at the atomic level relevant to the development of green energy and environment solutions
27 such as catalysts, batteries, and solar cell. Running MD simulations with millions of atoms and
28 time-steps is extremely time-consuming and resource-demanding. Thus, implementing responsible
29 high-performance computing requires users to pay more attention to allocating the optimal computing
30 resources for parallel atomic simulations. The Large-scale Atomic/Molecular Massively Parallel
31 Simulator (LAMMPS) is renowned for its extensible and substantial documentation, and has seen a
32 surge in adoption owing to the incorporation of new packages and functionalities tailored to ML-driven
33 simulations. While this expansion offers users exciting opportunities to explore complex biological
34 and physical systems, it also underscores the importance of benchmarking and resource optimization.
35 This work focuses on using LAMMPS with the aim of finding the optimal computing resource usage.
36 We start by analyzing and profiling bulk metallic systems to understand the relationship between the
37 number of nodes used in the calculation and the LAMMPS parallelization efficiency. The Kokkos

38 package and other kernel libraries are utilized to assess their impact on computational speed and
39 efficiency.

40 We find that surface reactions involving a metallic slab and a vacuum with reacting molecules (in this
41 case, carbon monoxide) suffer significant degradation of the parallel efficiency upon increasing the
42 number of nodes. Parallel efficiency drops by 30% if ten nodes are used on atomic systems of 20-60
43 thousand atoms. Profiling analysis indicates wasteful looping within specific functions and an increase
44 in delays of their execution time as a primary cause behind performance degradation. After a series
45 of calculations related to parallel efficiency, we build a set of guidelines for responsible computing to
46 be used in HPC environments. Performance analysis of these calculations have been used to develop
47 an automated tool that can recognize, evaluate, and recommend choices on resource usage to a user
48 based on the comparison, enabling them to save precious compute time and energy. This work aims
49 to lower the carbon footprint of computational loads used in developing some technologies of our
50 future and ensures that it will be built upon a foundation that prioritizes sustainability and the ethics
51 of responsible computing.

52 2 Software

53 Eco-Comp is a software that finds the optimal allocation of computing resources through bench-
54 marking statistics run in an automated manner. With great emphasis on ease of use, especially
55 for new users who do not possess the technical prowess of the architecture of the software, the
56 present-day best practice programming language of Python was chosen as our foundation. After a
57 simple installation, the software takes inputs from the user regarding details of the supercomputer and
58 the resources available. Using this information, a unique Python script is created for the user, which
59 can be used to run their simulation data. It automatically makes submissions via a job scheduler hence
60 benchmarking on the user’s specific production setup. This process, that lasts few seconds, provides
61 an accurate metric related to parallel efficiency. In addition, based on the data extracted from the job
62 submissions and the optimal configuration of computing resources, a ready-to-use customized job
63 submission file (Slurm scheduler bash file) is created for the user tailored for their simulation needs.

64 3 Methods

65 Eco-Comp can extract the required information from simulation data, and run bench-marking based
66 on the system complexity. This tool should then extract data from job submissions, suggest the
67 optimal configuration of use to the user, and, more importantly, provide them with a ready-to-use
68 customized job scheduler submission file based on their needs.

69 We used LAMMPS to evaluate the impact of characteristics such as vacuum space in different atom
70 systems, type & number of atoms present, the impact loaded packages such as Kokkos, ReaxFF, and
71 the Machine Learning SNAP (ML-SNAP) potential. Parallel efficiency using strong-scaling speedup
72 is calculated to understand performance degradation across various types of simulations, after which
73 Google Performance Tools are used to analyze potential bottlenecks from the profiling data.

74 To calculate the strong scaling speedup on N nodes:

$$\text{Strong Scaling Speedup on } N \text{ nodes} = \frac{\text{Time on 1 node}}{\text{Time on } N \text{ nodes}}$$

75 To calculate the parallel efficiency with N nodes as a percentage:

$$\text{Parallel Efficiency with } N \text{ nodes (\%)} = \frac{(\text{Strong Scaling Speedup on } N \text{ nodes}) \times 100}{\text{Number of nodes}}$$

76 This provided various metrics that were used to understand how these subtle changes affected runs
77 and their parallel efficiency. Tests showed us that utilizing the Kokkos package improved efficiency
78 and speed by almost two-fold. When a vacuum was present in the atomic system, there was an
79 increase of 30% computational speed when using the ReaxFF potential, whereas there was a greater
80 45% increase when using the Machine Learning Snap Potential. Through Google Performance
81 Tools, we established 66% as a threshold after profiling analysis; any efficiency below this value is
82 unsustainable and wastes resources and power.

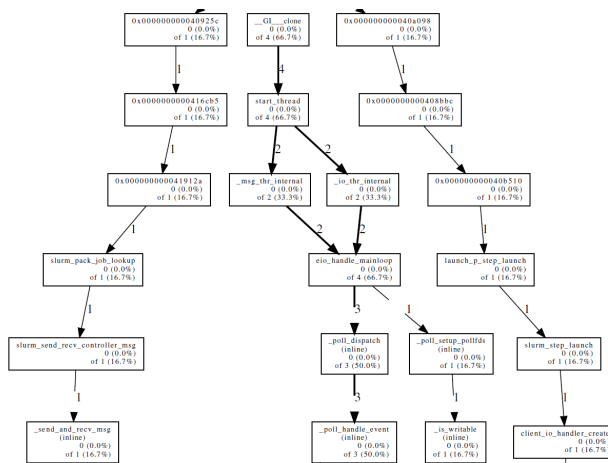


Figure 1: Bottleneck analysis on Google Performance Tools - visual call-graph output of a run at sub-par parallel efficiency




Figure 2: Impact of parallel efficiency on simulation runs using ReaxFF and ML-SNAP

83 4 Workflow of Software

```

JOBID USER ACCOUNT NAME ST REASON START_TIME TIME TIME_LEFT NODES CPUS
479014 ebentria p20004 20 R None 2023-02-22T14:54:08 0:08 6-23:59:52 2 80
JOBID USER ACCOUNT NAME ST REASON START_TIME TIME TIME_LEFT NODES CPUS
479014 ebentria p20004 20 R None 2023-02-22T14:54:08 0:11 6-23:59:49 2 80
JOBID USER ACCOUNT NAME ST REASON START_TIME TIME TIME_LEFT NODES CPUS
Total wallclock time: 6.58428
Number of cores: 80
Submitted batch job 479015
Slurm file submitted successfully.
479016
JOBID USER ACCOUNT NAME ST REASON START_TIME TIME TIME_LEFT NODES CPUS
479016 ebentria p20004 20 PD None N/A 0:00 7-00:00:00 3 120
JOBID USER ACCOUNT NAME ST REASON START_TIME TIME TIME_LEFT NODES CPUS
479016 ebentria p20004 20 R None 2023-02-22T14:54:23 0:02 6-23:59:58 3 120
JOBID USER ACCOUNT NAME ST REASON START_TIME TIME TIME_LEFT NODES CPUS
479016 ebentria p20004 20 R None 2023-02-22T14:54:23 0:05 6-23:59:55 3 120
JOBID USER ACCOUNT NAME ST REASON START_TIME TIME TIME_LEFT NODES CPUS
479016 ebentria p20004 20 R None 2023-02-22T14:54:23 0:08 6-23:59:52 3 120
JOBID USER ACCOUNT NAME ST REASON START_TIME TIME TIME_LEFT NODES CPUS
479016 ebentria p20004 20 R None 2023-02-22T14:54:23 0:11 6-23:59:49 3 120
JOBID USER ACCOUNT NAME ST REASON START_TIME TIME TIME_LEFT NODES CPUS
Total wallclock time: 5.54258
Number of cores: 120
[10.0, 40.0, 80.0, 120.0]
[24.9739, 11.5758, 6.58428, 5.54258]
File plot.png has been saved in your directory successfully.

The best option to use is running 80 cores because the parallel efficiency drops below 66% parallel efficiency threshold
after this:

```

Figure 3: Implementation of Eco-Comp on HPCs

84 Typically, the functionality of Eco-Comp is split into two simple and intuitive steps: (1) installation
85 of the Eco-Comp software, (2) running the Eco-Comp software and the analysis of the results. The
86 following paragraphs will delve into each of the three steps in detail.

87 **1. Installation** All the user is required to fork from GitHub then install the program and execute it.
88 The user will be prompted for various information required for future calculation purposes, such as

the type of supercomputer in use, nodes and cores available, maximum wall-time, and the directory of the LAMMPS executable. Based on this information, a run.py script is automatically generated, which can be added to the bash script of the user.

2. Execution The user can then run the Python script within the directory within which their simulation data exists. This will prompt the user for the input file and data file names, then it calculates the system's complexity and categorizes it based on the aforementioned parameters. Depending on the complexity, the script automatically submits 4 runs to the job scheduler system using different number of cores, and collects CPU time once simulations are completed. The simulation is run for 100 MD steps such that the entire process finishes in a matter of seconds. Once simulated, the program reads the data from the output file, and calculates and prints the optimal configuration for the user to use. A plot is generated by Eco-Comp, visualizing the benchmarking times, and this data is also saved in a .JSON file for the user's future reference.

5 Conclusions

In conclusion, we have successfully built Eco-Comp, a user-friendly Python tool that optimizes the computing power of simulations for material scientists. By using just one command, one could considerably reduce the HPC carbon footprint within the material science community. This approach in the automatic bench-marking of software and their respective simulations is not limited to LAMMPS, and further efforts can be focused on other popular packages within the materials science community, such as the Vienna Ab initio Simulation Package (VASP). In the future, an expansion to the such as Graphical Processing Units (GPU) or hybrid CPU/GPU systems would be of interest.

6 Acknowledgment

This work is supported by the Qatar National Research Fund (QNRF) through the National Priorities Research Program (NPRP) under project number NPRP12S-0209-190063. The advanced computing facilities of Hamad Bin Khalifa University ("Hazeem Supercomputer") and "Raad2" in Texas A&M University at Qatar are used for all calculations.

References

- [1] Armstrong, Brian & Kim, Seon & Eigenmann, Kim. (2000). *Quantifying Differences between OpenMP and MPI Using a Large-Scale Application Suite*. "Bondchk Failed with Reax/c." Materials Science Community Discourse, 7 Sept. 2016.
- [2] Eichstädt, J. & Green, M. & Turner, M. & Peiró, J., & Moxey, D. (2018, April 5). *Accelerating high-order mesh optimisation with an architecture-independent programming model*. Computer Physics Communications.
- [3] Google. "Google/Pprof." GitHub, <https://github.com/google/pprof>.
- [4] He, Helen. "Performance Engineering of Reactive Molecular Dynamics Simulations." MIT Libraries, Massachusetts Institute of Technology, <https://dspace.mit.edu/bitstream/handle/1721.1/139047/He-hmhe-meng-eecs-2021-thesis.pdf?sequence=1>
- [5] K. Cha, "Performance Evaluation of LAMMPS on Multi-core Systems," 2013 IEEE 10th International Conference on High Performance Computing and Communications & 2013 IEEE International Conference on Embedded and Ubiquitous Computing, 2013, pp. 812-819, doi: 10.1109/HPCC.and.EUC.2013.117.
- [6] LAMMPS Documentation. *LAMMPS Documentation* (23 Jun 2022 version) - LAMMPS documentation. (n.d.). <https://docs.lammps.org/Manual.html>
- [7] LAMMPS molecular dynamics simulator. *LAMMPS Molecular Dynamics Simulator*. (n.d.).
- [8] "Main GPerfTools Repository." GitHub, <https://github.com/gperftools/gperftools>.
- [9] "Parallel Performance." Parallel Performance - Parallel Computing for Beginners, <https://www.learnpdc.org/PDCBeginners/introduction/3.performance.html> . [10] "Performance Scaling." Archer - Writing Scalable Parallel Applications with MPI.
- [11] Yan, Beichuan & Regueiro, Richard. (2018). *Comparison between pure MPI and hybrid MPI-OpenMP parallelism for Discrete Element Method (DEM) of ellipsoidal and poly-ellipsoidal particles*. Computational Particle Mechanics. 6. 10.1007/s40571-018-0213-8.