Implementation of Dual Resolution Simulation Methodology in LAMMPS

Iain Bethune [1]*, Sophia Wheeler [2], Samuel Genheden [2,3] and Jonathan W. Essex [2]
* ibethune@epcc.ed.ac.uk

[1] EPCC, The University of Edinburgh [2] Department of Chemistry, The University of Southampton [3] University of Gothenburg

Overview

The ELBA force-field is an ELectrostatics-BAsed coarse-grained (CG) model developed by the Essex group for Molecular Dynamics simulations of lipid membrane systems. In the model, groups of atoms such as water molecules or choline, phosphate and glycerol regions of a phospholipid are represented as spherical beads, which may be charged and/or dipolar (see Figure 1). As well as the electrostatic interactions of the the beads, bonds may be defined, represented by harmonic bond and bond-angle potentials, and dispersion interactions are captured by a shifted-force Lennard-Jones potential. By aggregating several (typically 3-10) atoms into a single bead computational speedups of up to 200x may be obtained over all-atom simulations, while correctly predicting structural and dynamic properties [1].

The *dual-resolution (hybrid)* aspect of ELBA is that parts of the system may be represented atomistically, using the standard Amber force-field. Applications include atomistic bio-molecules solvated in CG water [2] and atomistic proteins embedded in a CG membrane environment [3]. ELBA has also been recently applied to compute solvation free energies of a range of small molecules in the SAMPL5 challenge [4].

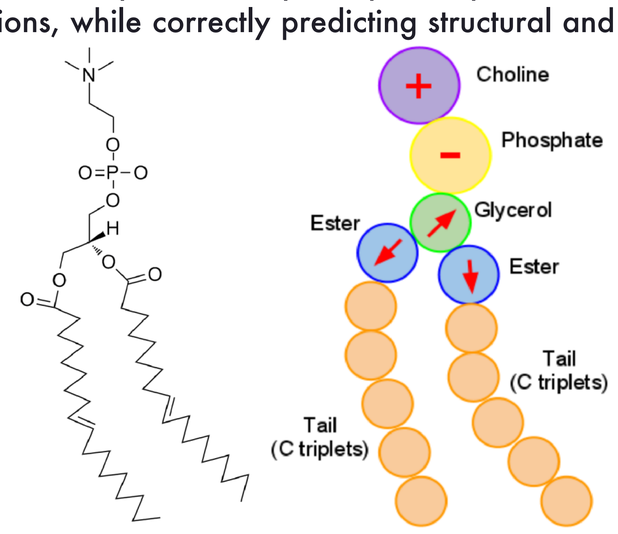

Figure 1: CG representation of a dioleoylphosphatidylcholine (DOPC) lipid in ELBA

ELBA is implemented in BRAHMS-MD – which is a serial code with a small user base. To enable ELBA for larger systems, making use of High Performance Computing and to make it available to a much larger user community, we have now implemented extensions to the LAMMPS package to support dual-resolution simulations. Full details of the implementation are available as an ARCHER white paper [5], and the key points are discussed here.

Rotational Integration

Stable long-timescale molecular dynamics requires using an integrator which is *symplectic* and *time-reversible*. For point particles, Verlet integration is sufficient, but to treat the rotational motion of the dipolar CG beads correctly, we can use the Dullweber-Leimkuhler-McLachlan (DLM) scheme [6]. An orientation matrix is constructed from the space-frame dipole vector, transformed into a local frame of reference and rotated according to the local-frame torque around each local axis, before an inverse transform recovers the updated space-frame dipole.

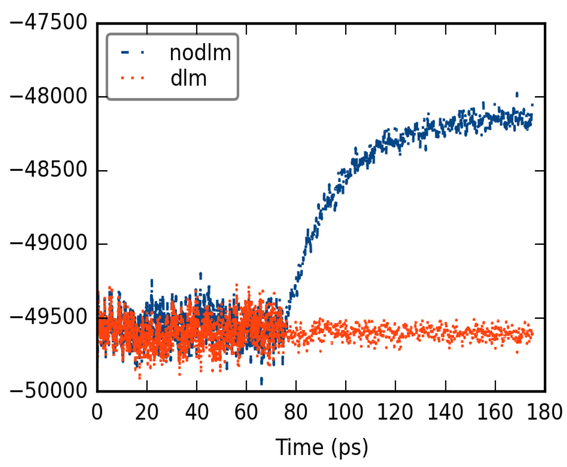

Figure 2: Improved energy conservation using the DLM integrator compared with the existing integrator in LAMMPS

Figure 2 shows the total energy of a system of 128 DMPC lipids during 175 ps of MD. During the first 75ps a Langevin thermostat is used, but once this is turned off the total energy quickly drifts unless the DLM integrator is used.

We have implemented the DLM integrator as an option to the existing LAMMPS rotational integration fixes, so it may be used to generate an NVE ensemble, or in combination with Nosé-Hoover thermostats and barostats to produce the NVT, NPT or NPH ensembles.

The DLM integrator is available in the LAMMPS 30 Jul 2016 stable release. Use `fix nve/sphere` (or `nvt/npt/nph`) with the keywords `update dipole/dlm`.

Example (integrating all particles using the DLM scheme):

```
fix integrate all nve/sphere update dipole/dlm
```

Load Balancing for Parallel Performance

Dual-resolution simulations obtain best performance when using the r-RESPA multiple time-stepping scheme [7]. Rapidly varying forces including the atomistic pair forces and electrostatics are computed with a short *inner* time-step of a few fs and other forces such as the CG pair forces are only computed every *outer* step, where the outer time-step is typically 2-8 times the inner time-step. Thus large numbers of CG beads can be included and integrated very cheaply without sacrificing accuracy for the atomistic integration.

In practice, this causes load-imbalance which can limit parallel scaling. For the purposes of dividing up the simulation domain over MPI processes, LAMMPS assumes every particle has the same computational cost! Not only do CG beads use different force calculations to atomistic particles, but they are computed much less frequently. Since atomistic particles are usually clustered in a particular part of the domain (where the molecules of interest are located), the default LAMMPS load balancer produces sub-optimal partitions in most cases.

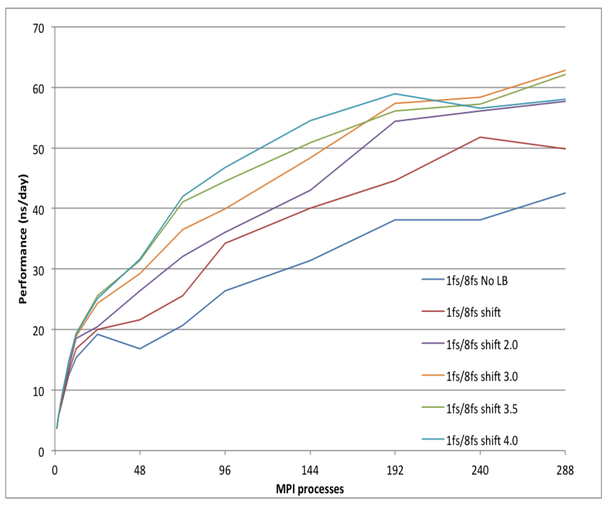

Figure 3: Parallel scaling of hybrid BPTI system with 8:1 timestep ratio for different values of the solute load weighting on ARCHER.

We have extended LAMMPS so that groups of particles can be assigned a weight for the purposes of determining the optimal load balance. For example, an atomistic solute particle might be assigned a load of 4 compared with a CG water bead with a load of 1. Figure 3 shows that for an atomistic BPTI molecule (882 atoms) solvated in 6136 ELBA beads with a time-step ratio of 8:1 speedups of up to 28% on a single ARCHER node (24 cores) or 65% on 72 cores can be achieved using a solute load factor of 4. This option is included in the ICMS branch of LAMMPS and will be made available in the next stable release (likely late 2016).

Example (load balancing using the `shift` algorithm to within 10% using a weighting of 4.0 for particles in the `solute` group):

```
group solute type > 1
balance 1.1 shift xyz 50 1.1 weight group 1 solute 4.0
```

ELBA Tutorials & Instructions

A comprehensive set of tutorials, force-field parameters, example LAMMPS input files and discussion of how to set up simulations can be found on the web:

https://sgenheden.github.io/Elba

References

[1] Orsi M, Essex JW (2011). *PLoS ONE* 6(12): e28637
[2] Orsi M, Ding W and Palaiokostas M (2014) *JCTC* 10(10)
[3] Genheden S, Essex JW (2015) *JCTC* 11(10)
[4] Genheden S, Essex JW (2016) *J Comput Aided Mol Des* (Online)
[5] Bethune I *et al* (2016) ARCHER White Paper *http://tinyurl.com/LAMMPS-ELBA*
[6] Dullweber A, Leimkuhler B and McLachlan R (1997) *JCP* 107(15)
[7] Tuckerman M, Berne B J and Martyna G J (1992) *JCP* 97(3)

Acknowledgements

This work was funded under the embedded CSE programme of the ARCHER UK National Supercomputing Service. This work used the ARCHER UK National Supercomputing Service. (http://www.archer.ac.uk)

epcc | EPSRC Engineering and Physical Sciences Research Council | archer