Welcome!

Virtual tutorial starts at 15:00 GMT

Please leave feedback afterwards at:
www.archer.ac.uk/training/feedback/online-course-feedback.php

archer

epcc

THE UNIVERSITY OF EDINBURGH

Multi-resolution modelling of biological systems in LAMMPS

ARCHER Virtual Tutorial, 19th Oct 2016

Iain Bethune ibethune@epcc.ed.ac.uk
Oliver Henrich ohenrich@epcc.ed.ac.uk

Reusing this material

archer

epcc

THE UNIVERSITY OF EDINBURGH

Outline

- ARCHER eCSE programme

- Implementation of Dual Resolution Simulation Methodology in LAMMPS
  - ELBA force-field
  - Implementation in LAMMPS
  - Performance testing
  - Summary

ARCHER eCSE programme

- Funding for the ARCHER user community to develop software
  - Implementation of algorithmic improvements within an existing code
  - Improving the scalability of software on higher core counts
  - Improvements to code which allows new science to be carried out
  - Porting and optimising a code to run efficiently on ARCHER
  - Adding new functionalities to existing codes
  - Code development to take a code from a Tier-2 (Regional) or local university cluster to Tier-1 (National) level bringing New Communities onto ARCHER
- Projects typically 3 months – 1 year
- Next call closes 31st Jan 2017

ARCHER eCSE programme

- More information on the ARCHER website:
  - https://www.archer.ac.uk/community/eCSE/
  - Project Reports
  - How To Apply
  - List of funded projects

- Webinar from last month:
  - https://youtu.be/WRGsNKWrNIc

Implementation of Dual Resolution Simulation Methodology in LAMMPS

- eCSE04-7 (January 2015)
- PI: Prof. Jonathan Essex, Southampton
- 6 person-months funded: August 2015 – August 2016

- Objective: enable fast and reliable calculations with the ELBA force-field in LAMMPS
  - New integrators
  - Parallel load balancing

ELBA Force-field

- ELBA = ELectrostatics-BAsed coarse grained forcefield
  - Orsi & Essex, *PLoS ONE* 6(12) 2011
- Originally for studying lipids
  - Also applied to other biomolecules
- Explicit solvent
  - One dipolar bead per water molecule
- Allows for atomistic detail e.g.
  - Using CHARMM parameters

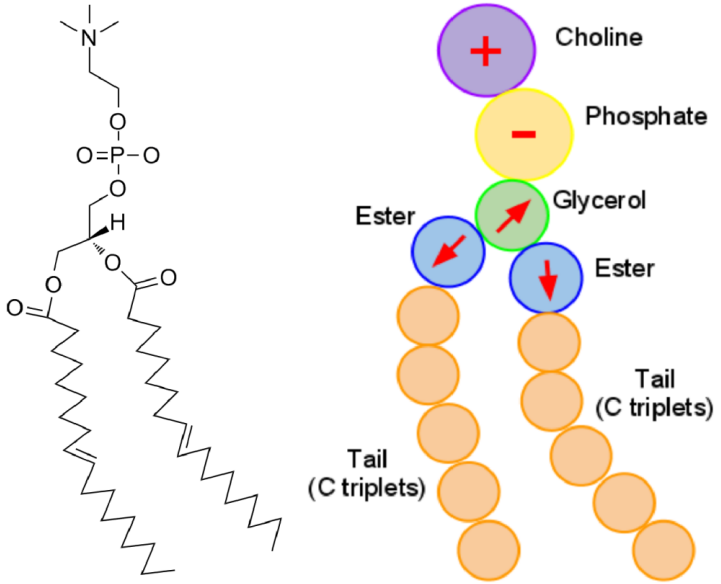

138 atoms -> 15 CG beads

archer

epcc

THE UNIVERSITY OF EDINBURGH

ELBA Force-field

- Implemented in BRAHMS-MD (Biomembrane Reduced-ApproacH Multiresolution Simulator for Molecular Dynamics):
  - https://code.google.com/archive/p/brahms-md/
  - Limited user base, single developer -> not sustainable
  - No parallelisation -> small systems

- Why LAMMPS?
  - Main interaction types already implemented
  - Support for spherical particles
  - r-RESPA multiple timestepping
  - Flexible, scalable, large user base

http://lammps.sandia.gov

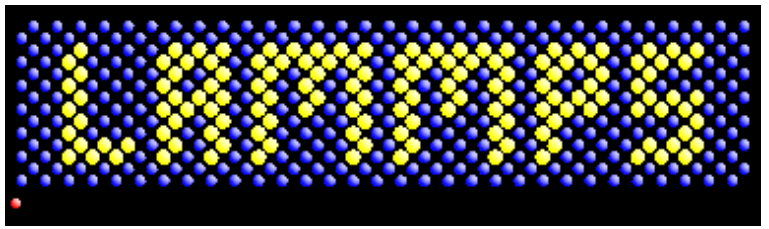

archer

epcc

THE UNIVERSITY OF EDINBURGH

Implementation in LAMMPS

- LAMMPS `fix nve/sphere` integrator does not conserve energy well
- Better scheme to integrate rotational d.o.f. - DLM
  Dullweber, Leimkuhler and McLachlan, *JCP* 107(15) 1997

1. Construct rotation matrix Q from dipole (taken as the body-fixed z-axis)
2. In body-space, apply rotations around each local axis:

$$\omega_b = Q\omega_s$$

$$R_1 = R_x(\frac{\delta t}{2}\omega_1), \quad \omega = R_1\omega, \quad Q = R_1^T Q$$

$$R_2 = R_y(\frac{\delta t}{2}\omega_2), \quad \omega = R_2\omega, \quad Q = R_2^T Q$$

$$R_3 = R_z(\delta t\omega_3), \quad \omega = R_3\omega, \quad Q = R_3^T Q$$

$$R_4 = R_y(\frac{\delta t}{2}\omega_2), \quad \omega = R_4\omega, \quad Q = R_4^T Q$$

$$R_5 = R_x(\frac{\delta t}{2}\omega_1), \quad \omega = R_5\omega, \quad Q = R_5^T Q$$

3. Finally, compute the new dipole: $\mu_s = Q^T[001] \cdot \|\mu\|$

Implementation in LAMMPS

- 4000 ELBA water beads, 10fs timestep, 20ps NVT, 20ps NVE

```
fix thermostat all langevin 303 303 200 48279 omega yes
```

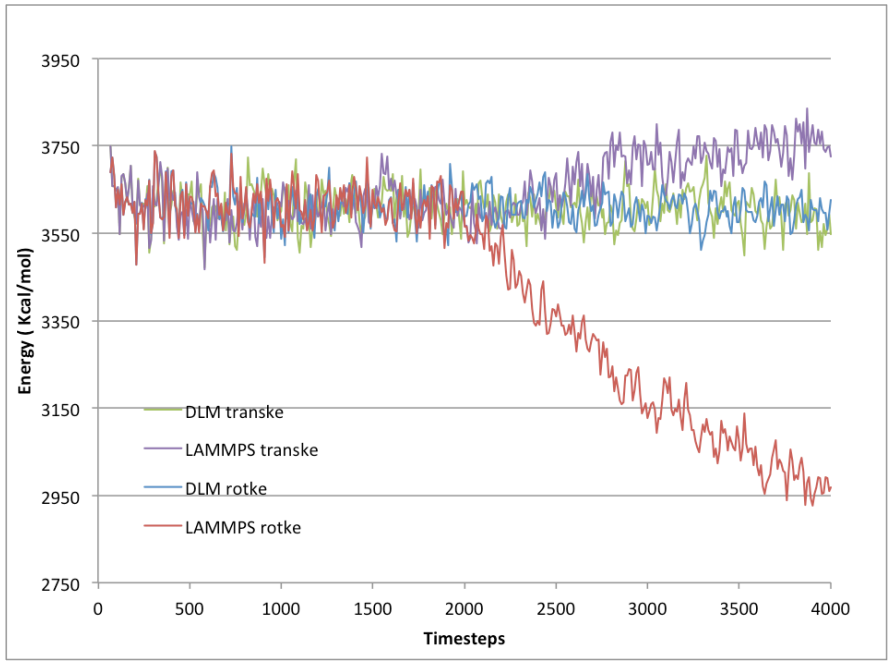

archer

epcc

THE UNIVERSITY OF EDINBURGH

Implementation in LAMMPS

- 128 DPMC molecules in water, 75ps NVT, 100ps NVE

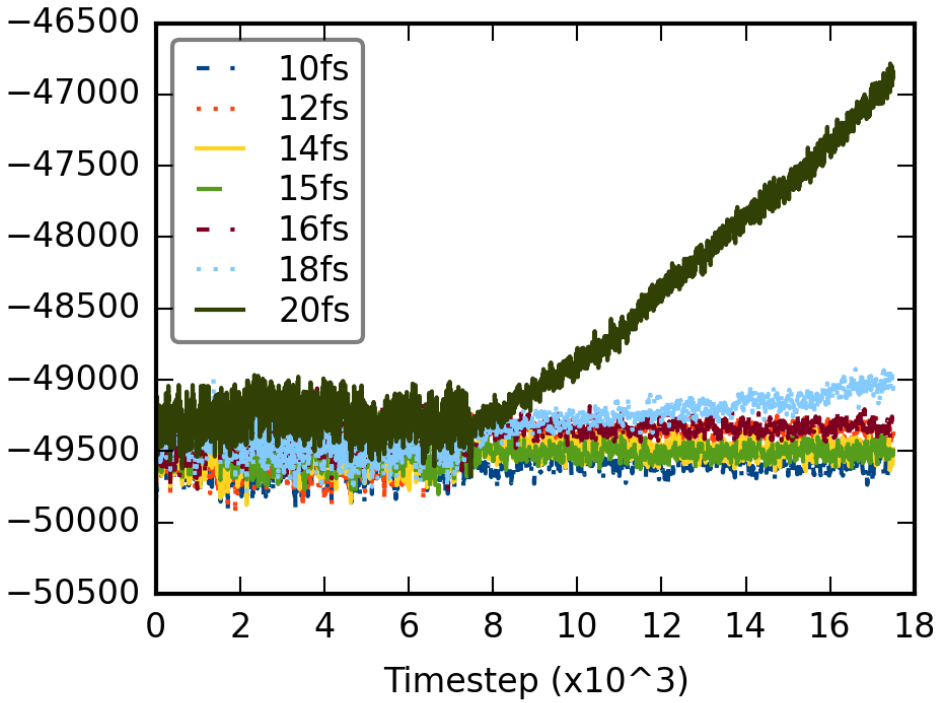

Image from Sam Genheden

archer

epcc

THE UNIVERSITY OF EDINBURGH

Implementation in LAMMPS

- DLM integrator enabled by an optional argument:
  - `fix nve/sphere … update dipole/dlm`
- Also for other ensembles:
  - Constant temperature / NVT (Nosé-Hoover)
  - `fix nvt/sphere … update dipole/dlm`
  - Isothermal-isobaric (Nosé-Hoover / Parrinello-Rahman)
  - `fix npt/sphere … update dipole/dlm`
  - Isenthalpic (Parrinello-Rahman)
  - `fix nph/sphere … update dipole/dlm`

Implementation in LAMMPS

- Load balancing schemes:

http://lammps.sandia.gov/doc/balance.html

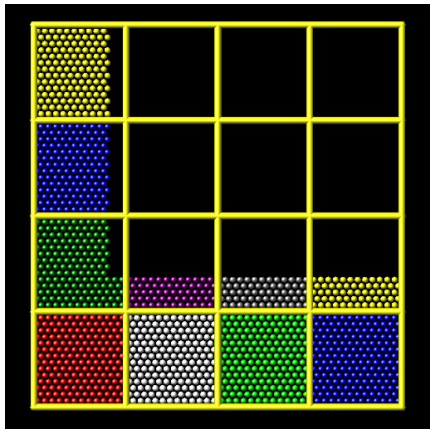

None

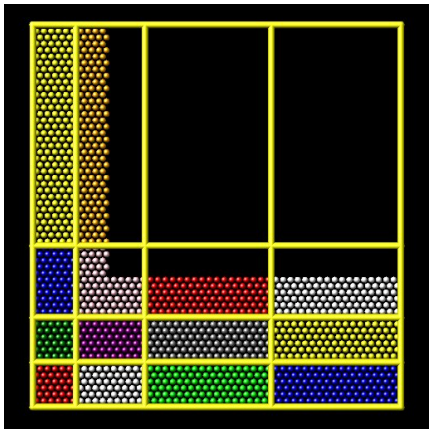

Shift

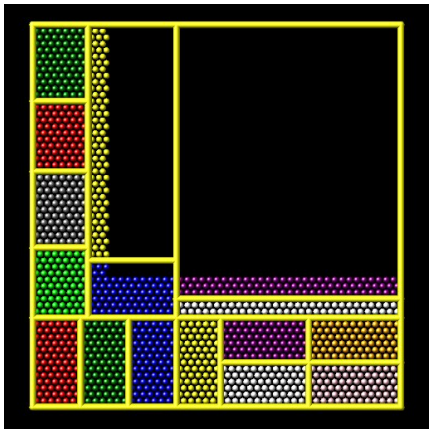

RCB

- Problem for dual-resolution simulations!
  - 90% of computational cost is force evaluation
  - Not all particles are the same
  - r-RESPA – some forces are computed more frequently than others

Implementation in LAMMPS

- New load balancing metrics:
  - Weighting by particle groups
    - Uses LAMMPS existing `group` command e.g.

```
group solute type > 1
group water type 1
```

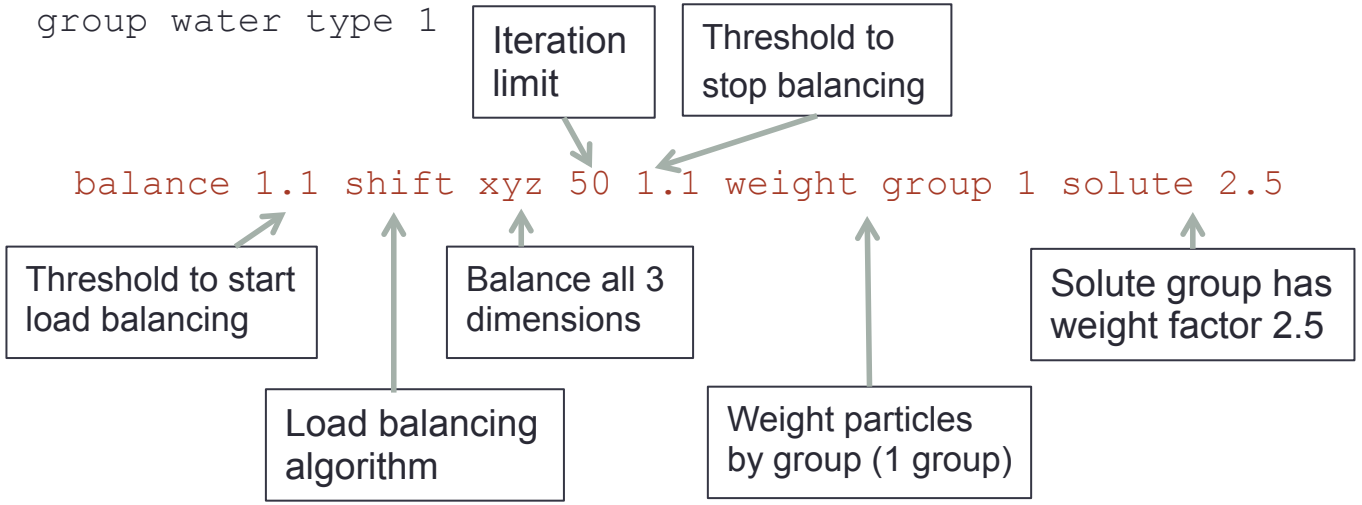

Implementation in LAMMPS

- New load balancing metrics (subsequently added by LAMMPS developers):
  - Weighting by number of neighbors – `weight neigh`
  - Weighting by compute time – `weight time`
    - Doesn't account for the different particles types contributing to different parts of the computation (pair, bond, kspace, neigh)
  - Weighting by arbitrary user-defined variables – `weight var`

Performance testing

- Bovine Pancreatic Trypsin Inhibitor (BPTI) dual-resolution model:
  - 882 atoms, CHARMM force-field
  - 6136 water molecules, ELBA beads

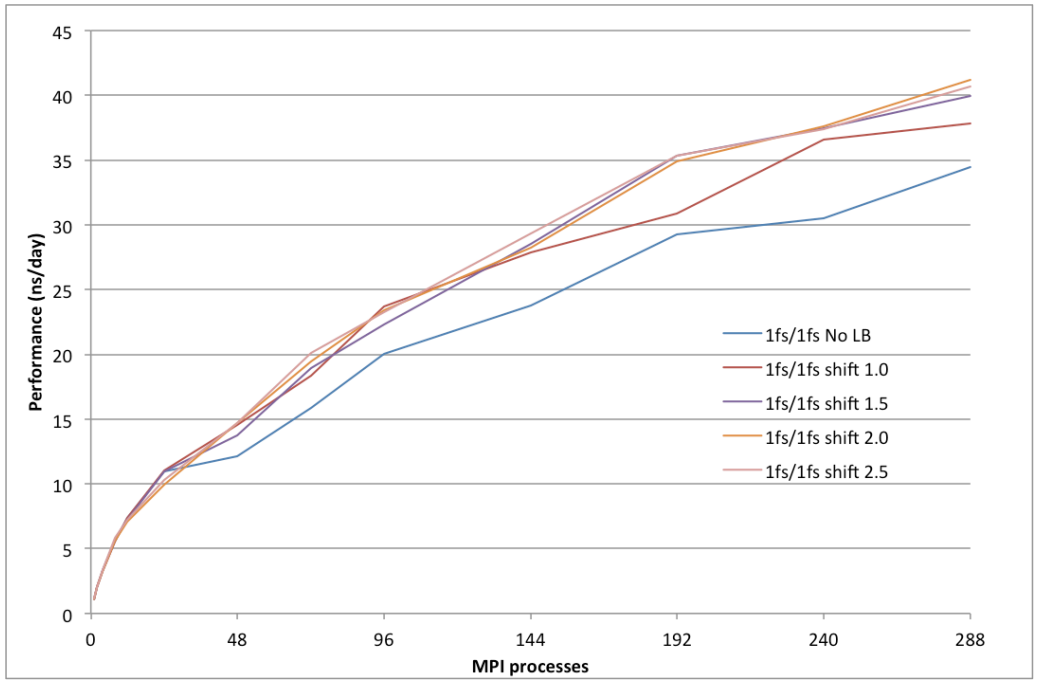

- No r-RESPA
- 1fs timestep
- Up to 10% speedup over non-weighted balance

archer

epcc

THE UNIVERSITY OF EDINBURGH

Performance testing

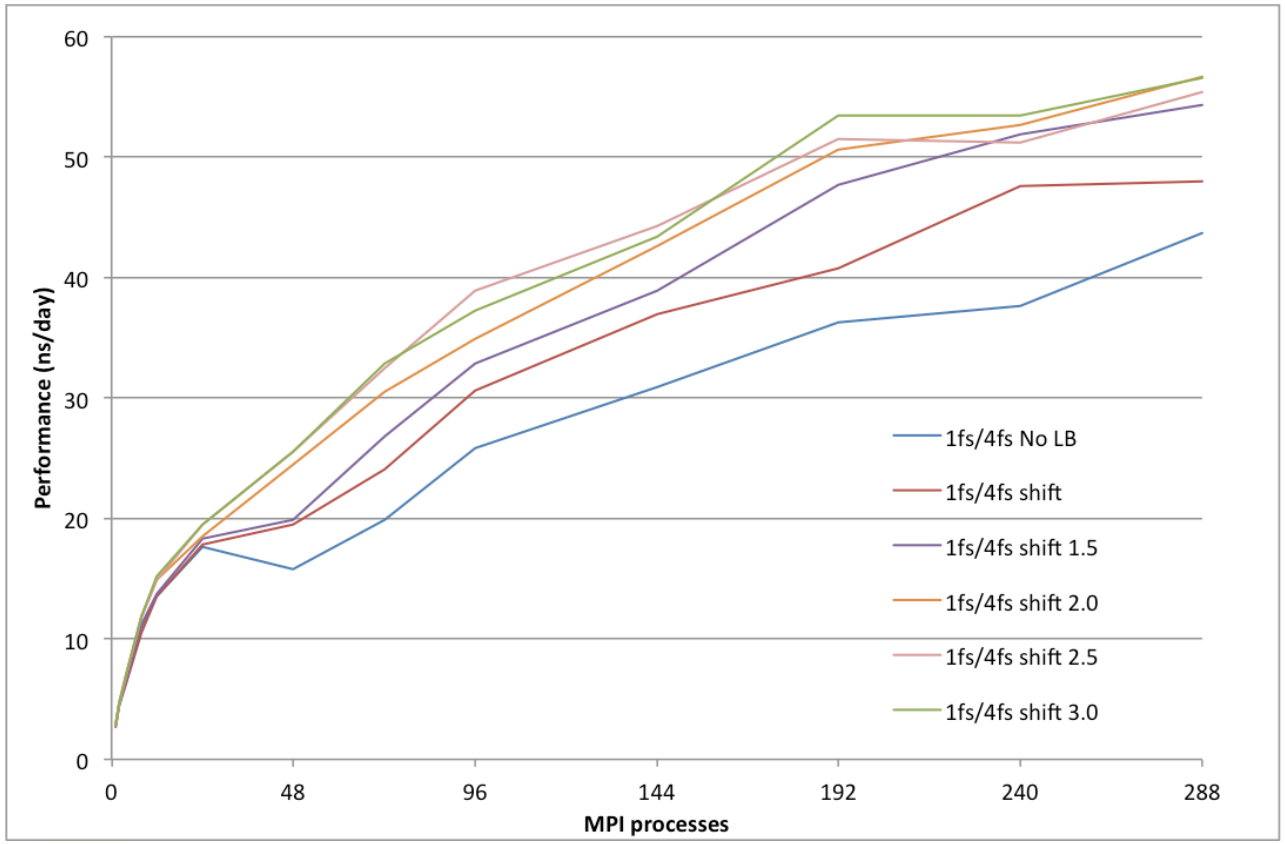

- 1:4 r-RESPA ratio
- Water pair forces + dihedral forces computed every fourth step
- Larger weightings better (2.5-3.0)
- Up to 36% speedup over non-weighted balance

archer

epcc

THE UNIVERSITY OF EDINBURGH

Performance testing

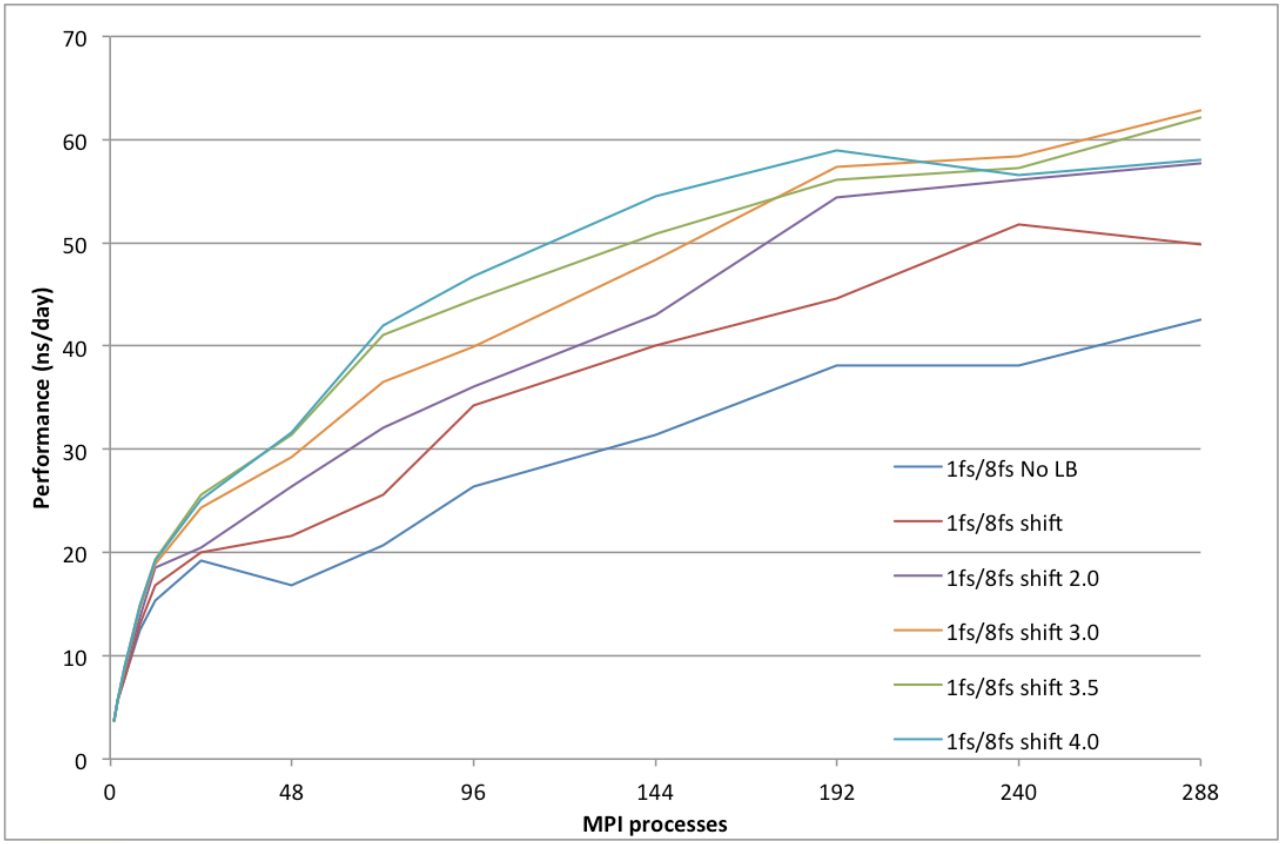

- 1:8 r-RESPA ratio
- Larger weightings better (3.0-4.0)
- Up to 65% speedup over non-weighted balance

archer

epcc

THE UNIVERSITY OF EDINBURGH

Summary

- DLM integrator for NVE/NVT/NPT/NPH dynamics
  - Stable for water up to 16fs timestep
  - Included in LAMMPS stable release 30 Jul 2016

Installed on ARCHER

```
module load
   lammps/elba
```

- New load balancing metrics
  - Better performance for r-REPSA and hybrid pair force
  - Include in LAMMPS patch release 27 Sep 2016

- More information
  - Technical Report: http://www.archer.ac.uk/documentation/white-papers/lammps-elba/lammps-ecse.pdf
  - Tutorials, references, discussion: https://sgenheden.github.io/Elba/


archer


epcc
THE UNIVERSITY OF EDINBURGH

ARCHER eCSE05-10 Project

- **Adding Multiscale Models of DNA to LAMMPS** (09/2015 – 08/2016)
- Dr Oliver Henrich (PI, UoE), Prof Davide Marenduzzo (Co-I, UoE), Dr Thomas Ouldridge (Co-I, Imperial College London)
- Overview:
  - Implemented **oxDNA coarse-grained DNA model** for single- and double- stranded DNA into LAMMPS code
  - Implemented new **Langevin-type rigid-body integrators**
  - Software available
    from **CCPForge** (https://ccpforge.cse.rl.ac.uk/gf/project/cgdna)
    soon as **LAMMPS USER-package**
  - Currently adapting utility software of oxDNA standalone version

From DNA to Chromosomes

- Haploid human genome contains 3.2 billion base pairs organised in 23 chromosomes
  - Diameter of DNA strand = 2 x 10^{-9}m
  - Total length of DNA in human cell = 2m
  - Diameter of spherical ‘blob’ of DNA in human cell = 2 x 10^{-6}m

- Hierarchical organisation
  - Histone octamer
  - Nucleosome core particles 200 bps
  - 10 nm beads-on-a-string
  - 30 nm chromatin fibre
  - smallest loop in chromatin fibre 50,000 bps

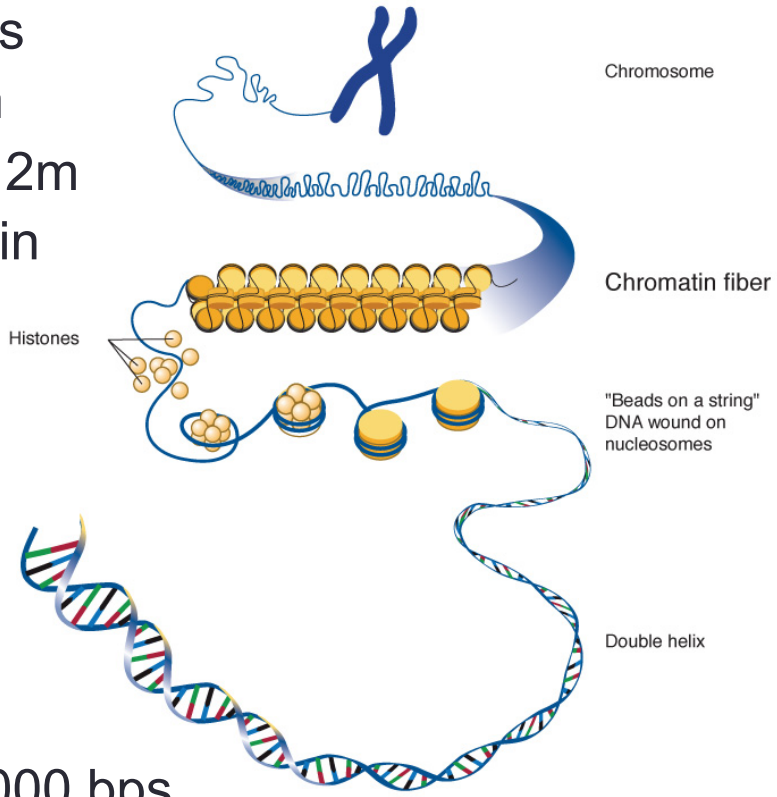

archer

epcc

THE UNIVERSITY OF EDINBURGH

Atomistic Simulation of DNA

- Good for capturing fast conformational fluctuations and protein-DNA binding
- Usually limited to a few 1000 base pairs
- Phenomena on large time and length scales, e.g. DNA supercoiling or nucleosome positioning are permanently out of reach

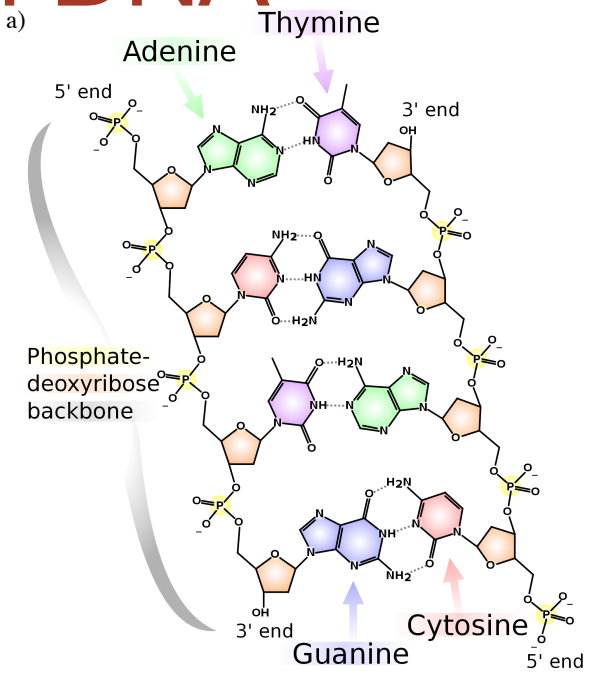

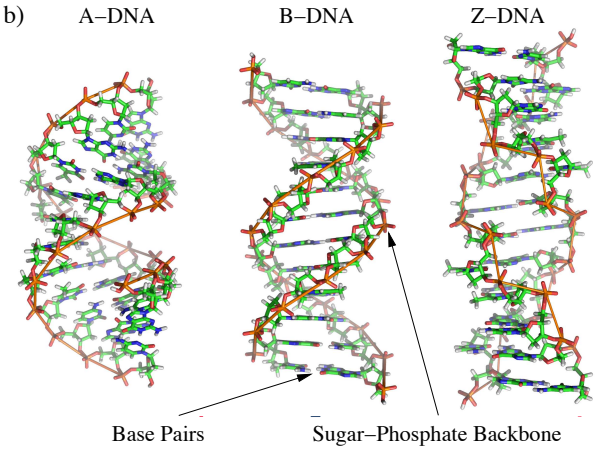

archer

THE UNIVERSITY OF EDINBURGH

Coarse-grained simulation with oxDNA

- Must deliver correct longitudinal and torsional persistence length, electrostatics, if sequence-specific correct melting temperature, elastic constants …
- oxDNA: top-down approach of a CG model, nucleotides modelled as rigid bodies (DOF are COM & quaternion)
- Parameterize interaction between nucleotides with 6 independent interactions (7 for implicit electrostatics)

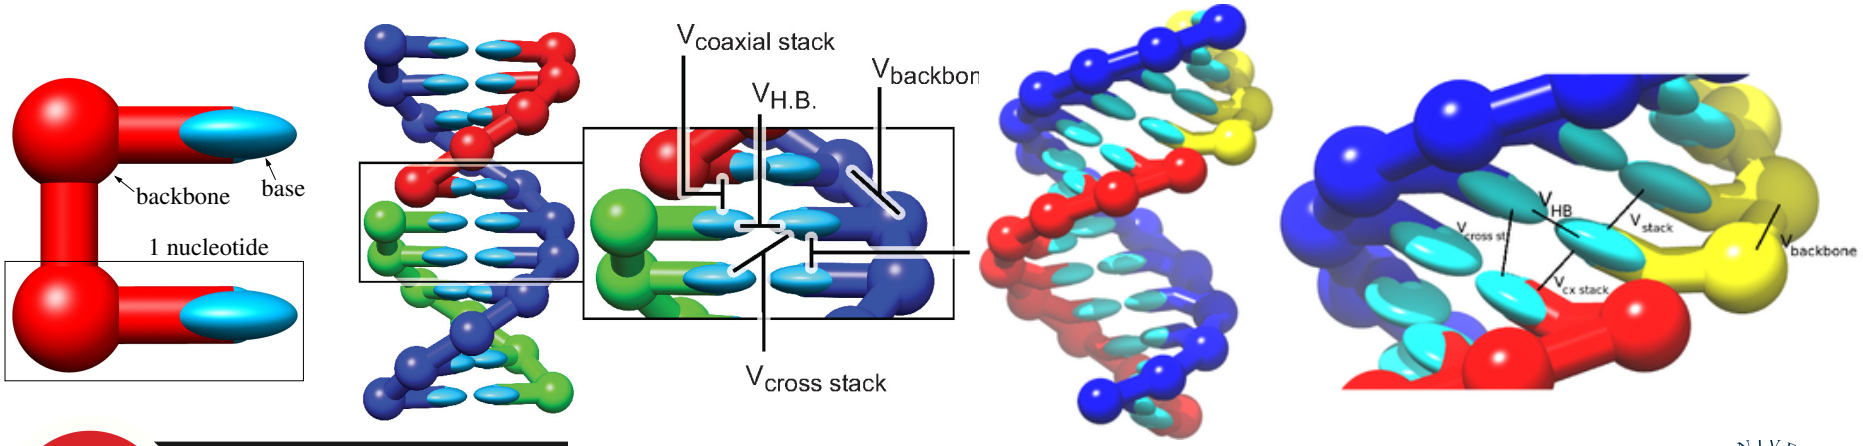

archer

epcc

THE UNIVERSITY OF EDINBURGH

oxDNA Force Field

- **Backbone**: FENE (finite extensible nonlinear elastic) springs
- **Excluded volume**: Lennard-Jones potential
- **Stacking**: harmonic angle x Morse potential
- **Cross-stacking**: harmonic angle x harmonic distance potential
- **Coaxial stacking**: harmonic angle x harmonic distance potential
- **Hydrogen bonding**: harmonic angle x Morse potential

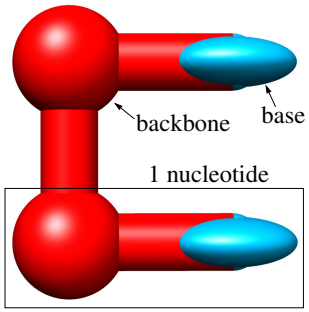

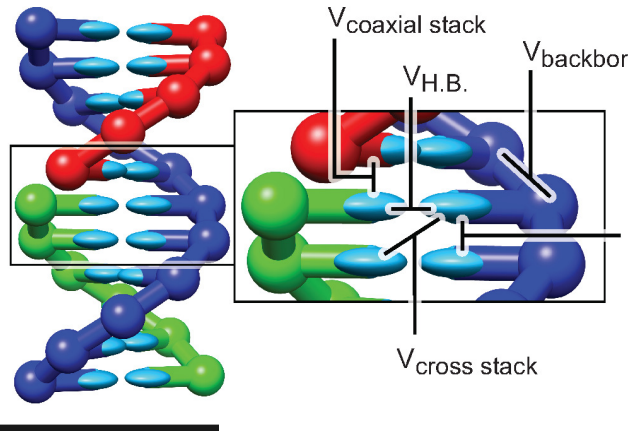

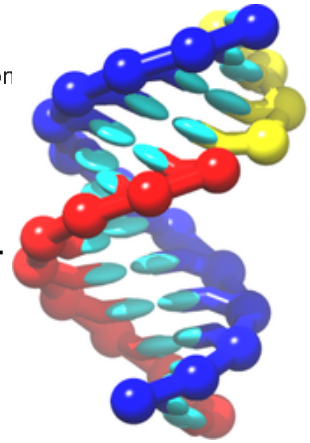

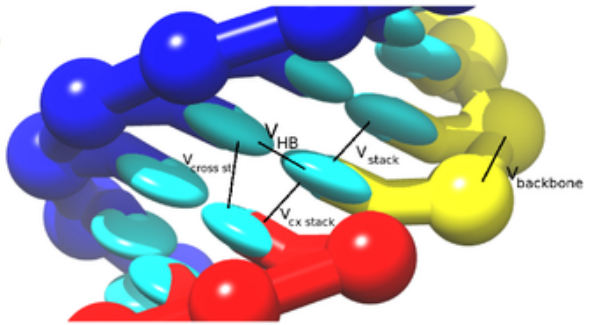

archer

epcc

THE UNIVERSITY OF EDINBURGH

oxDNA Force Field

- Smoothed, truncated and modulated forms of the above
- 1 bonded interaction (backbone), 5 pair interactions (excluded volume, stacking, hydrogen boding, cross-stacking, coaxial stacking)

- FENE spring (used to connect backbones):

$$V_{\mathrm{FENE}}(r,\epsilon,r^0,\Delta) = -\frac{\epsilon}{2}\ln\left(1-\frac{(r-r^0)^2}{\Delta^2}\right)$$

- Morse potential (used for stacking and H-bonding):

$$V_{\mathrm{Morse}}(r,\epsilon,r^0,a) = \epsilon\left(1-\exp\left(-(r-r^0)a\right)\right)^2$$

- Harmonic potential (used for cross-stacking and coaxial stackin

$$V_{\mathrm{harm}}(r,k,r^0) = \frac{k}{2}\left(r-r^0\right)^2.$$

- Lennard - Jones potential (used for soft repulsion):

$$V_{\mathrm{LJ}}(r,\epsilon,\sigma) = 4\epsilon\left(\left(\frac{\sigma}{r}\right)^{12}-\left(\frac{\sigma}{r}\right)^{6}\right).$$

- Quadratic terms (used for modulation):

$$V_{\mathrm{mod}}(\theta,a,\theta^0) = 1-a(\theta-\theta^0)^2.$$

- Quadratic smoothing terms for truncation:

$$V_{\mathrm{smooth}}(x,b,x^c) = b(x^c-x)^2.$$

- The radial part of the stacking and hydrogen-bonding potentials:

$$f_1(r) = \begin{cases} V_{\mathrm{Morse}}(r,\epsilon,r^0,a) - V_{\mathrm{Morse}}(r^c,\epsilon,r^0,a)) & \text{if } r^{low} < r < r^{high}, \\ \epsilon V_{\mathrm{smooth}}(r, b^{low}, r^{c,low}) & \text{if } r^{c,low} < r < r^{low}, \\ \epsilon V_{\mathrm{smooth}}(r, b^{high}, r^{c,high}) & \text{if } r^{high} < r < r^{c,high}, \\ 0 & \text{otherwise.} \end{cases} \quad (2.7)$$

- The radial part of the cross-stacking and coaxial stacking potentials:

$$f_2(r) = \begin{cases} V_{\mathrm{harm}}(r,k,r^0) - V_{\mathrm{harm}}(r^c,k,r^0) & \text{if } r^{low} < r < r^{high}, \\ kV_{\mathrm{smooth}}(r, b^{low}, r^{c,low}) & \text{if } r^{c,low} < r < r^{low}, \\ kV_{\mathrm{smooth}}(r, b^{high}, r^{c,high}) & \text{if } r^{high} < r < r^{c,high}, \\ 0 & \text{otherwise.} \end{cases} \quad (2.8)$$

- The radial part of the excluded volume potential:

$$f_3(r) = \begin{cases} V_{\mathrm{LJ}}(r,\epsilon,\sigma) & \text{if } r < r^*, \\ \epsilon V_{\mathrm{smooth}}(r,b,r^c) & \text{if } r^* < r < r^c, \\ 0 & \text{otherwise.} \end{cases} \quad (2.9)$$

- The angular modulation factor used in stacking, hydrogen-bonding, cross-stacking and coaxial stacking:

$$f_4(\theta) = \begin{cases} V_{\mathrm{mod}}(\theta,a,\theta^0) & \text{if } \theta^0 - \Delta\theta^* < \theta < \theta^0 + \Delta\theta^*, \\ V_{\mathrm{smooth}}(\theta,b,\theta^0 - \Delta\theta^c) & \text{if } \theta^0 - \Delta\theta^c < \theta < \theta^0 - \Delta\theta^*, \\ V_{\mathrm{smooth}}(\theta,b,\theta^0 + \Delta\theta^c) & \text{if } \theta^0 + \Delta\theta^* < \theta < \theta^0 + \Delta\theta^c, \\ 0 & \text{otherwise.} \end{cases} \quad (2.10)$$

- Another modulating term which is used to impose right-handedness (effectively a one-sided modulation):

$$f_5(\phi) = \begin{cases} 1 & \text{if } x > 0, \\ V_{\mathrm{mod}}(x,a,0) & \text{if } x^* < x < 0, \\ V_{\mathrm{smooth}}(x,b,x^c) & \text{if } x^c < x < x^*, \\ 0 & \text{otherwise.} \end{cases} \quad (2.11)$$

$$V_{\mathrm{backbone}} = V_{\mathrm{FENE}}(\delta r_{\mathrm{backbone}}, \epsilon_{\mathrm{backbone}}, \delta r^0_{\mathrm{backbone}}, \Delta_{\mathrm{backbone}}).$$

$$\begin{aligned} V_{\mathrm{exc}} &= f_3(\delta r_{\mathrm{backbone}}, \epsilon_{\mathrm{exc}}, \sigma_{\mathrm{backbone}}, \delta r^*_{\mathrm{backbone}}) + f_3(\delta r_{\mathrm{base}}, \epsilon_{\mathrm{exc}}, \sigma_{\mathrm{base}}, \delta r^*_{\mathrm{base}})) \\ &\quad + f_3(\delta r_{\mathrm{back-base}}, \epsilon_{\mathrm{exc}}, \sigma_{\mathrm{back-base}}, \delta r^*_{\mathrm{back-base}}) \\ &\quad + f_3(\delta r_{\mathrm{base-back}}, \epsilon_{\mathrm{exc}}, \sigma_{\mathrm{back-base}}, \delta r^*_{\mathrm{back-base}}). \end{aligned}$$

$$\begin{aligned} V'_{\mathrm{exc}} &= f_3(\delta r_{\mathrm{base}}, \epsilon_{\mathrm{exc}}, \sigma_{\mathrm{base}}, \delta r^*_{\mathrm{base}}) \\ &\quad + f_3(\delta r_{\mathrm{back-base}}, \epsilon_{\mathrm{exc}}, \sigma_{\mathrm{back-base}}, \delta r^*_{\mathrm{back-base}}) \\ &\quad + f_3(\delta r_{\mathrm{base-back}}, \epsilon_{\mathrm{exc}}, \sigma_{\mathrm{back-base}}, \delta r^*_{\mathrm{back-base}}). \end{aligned}$$

$$\begin{aligned} V_{\mathrm{stack}} &= f_1(\delta r_{\mathrm{stack}}, \epsilon_{\mathrm{stack}}, a_{\mathrm{stack}}, \delta r^0_{\mathrm{stack}}, \delta r^{c,low}_{\mathrm{stack}}, \delta r^{c,high}_{\mathrm{stack}}, \delta r^{low}_{\mathrm{stack}}, \delta r^{high}_{\mathrm{stack}}) \\ &\times f_4(\theta_4, a_{\mathrm{stack},4}, \theta^0_{\mathrm{stack},4}, \Delta\theta^*_{\mathrm{stack},4}) \\ &\times f_4(\theta_{5'}, a_{\mathrm{stack},5}, \theta^0_{\mathrm{stack},5}, \Delta\theta^*_{\mathrm{stack},5})\ f_4(\theta_{6'}, a_{\mathrm{stack},6}, \theta^0_{\mathrm{stack},6}, \Delta\theta^*_{\mathrm{stack},6}) \\ &\times f_5(-\cos(\phi_1), a_{\mathrm{stack},1}, -\cos(\phi_1)^*_{\mathrm{stack}})\ f_5(-\cos(\phi_2), a_{\mathrm{stack},2}, -\cos(\phi_2)^*_{\mathrm{stack}}). \end{aligned}$$

$$\begin{aligned} V_{\mathrm{HB}} &= f_1(\delta r_{\mathrm{HB}}, \epsilon_{\mathrm{NB}}, a_{\mathrm{HB}}, \delta r^0_{\mathrm{HB}}, \delta r^{c,low}_{\mathrm{HB}}, \delta r^{c,high}_{\mathrm{HB}}, \delta r^{low}_{\mathrm{HB}}, \delta r^{high}_{\mathrm{HB}}) \\ &\times f_4(\theta_1, a_{\mathrm{HB},1}, \theta^0_{\mathrm{HB},1}, \Delta\theta^*_{\mathrm{HB},1})\ f_4(\theta_2, a_{\mathrm{HB},2}, \theta^0_{\mathrm{HB},2}, \Delta\theta^*_{\mathrm{HB},2}) \\ &\times f_4(\theta_3, a_{\mathrm{HB},3}, \theta^0_{\mathrm{HB},3}, \Delta\theta^*_{\mathrm{HB},3})\ f_4(\theta_4, a_{\mathrm{HB},4}, \theta^0_{\mathrm{HB},4}, \Delta\theta^*_{\mathrm{HB},4}) \\ &\times f_4(\theta_7, a_{\mathrm{HB},7}, \theta^0_{\mathrm{HB},7}, \Delta\theta^*_{\mathrm{HB},7})\ f_4(\theta_8, a_{\mathrm{HB},8}, \theta^0_{\mathrm{HB},8}, \Delta\theta^*_{\mathrm{HB},8}). \end{aligned}$$

$$\begin{aligned} V_{\mathrm{cross\_stack}} &= f_2(\delta r_{\mathrm{HB}}, k_{\mathrm{cross}}, \delta r^0_{\mathrm{cross}}, \delta r^{c,low}_{\mathrm{cross}}, \delta r^{c,high}_{\mathrm{cross}}, \delta r^{low}_{\mathrm{cross}}, \delta r^{high}_{\mathrm{cross}}) f_4(\theta_1, a_{\mathrm{cross},1}, \theta^0_{\mathrm{cross},1}, \Delta\theta^*_{\mathrm{cross},1}) \\ &\times f_4(\theta_2, a_{\mathrm{cross},2}, \theta^0_{\mathrm{cross},2}, \Delta\theta^*_{\mathrm{cross},2})\ f_4(\theta_3, a_{\mathrm{cross},3}, \theta^0_{\mathrm{cross},3}, \Delta\theta^*_{\mathrm{cross},3}) \\ &\times \left(f_4(\theta_4, a_{\mathrm{cross},4}, \theta^0_{\mathrm{cross},4}, \Delta\theta^*_{\mathrm{cross},4}) + f_4(\pi-\theta_4, a_{\mathrm{cross},4}, \theta^0_{\mathrm{cross},4}, \Delta\theta^*_{\mathrm{cross},4})\right) \\ &\times \left(f_4(\theta_7, a_{\mathrm{cross},7}, \theta^0_{\mathrm{cross},7}, \Delta\theta^*_{\mathrm{cross},7}) + f_4(\pi-\theta_7, a_{\mathrm{cross},7}, \theta^0_{\mathrm{cross},7}, \Delta\theta^*_{\mathrm{cross},7})\right) \\ &\times \left(f_4(\theta_8, a_{\mathrm{cross},8}, \theta^0_{\mathrm{cross},8}, \Delta\theta^*_{\mathrm{cross},8}) + f_4(\pi-\theta_8, a_{\mathrm{cross},8}, \theta^0_{\mathrm{cross},8}, \Delta\theta^*_{\mathrm{cross},8})\right). \end{aligned}$$

$$\begin{aligned} V_{\mathrm{coax\_stack}} &= f_2(\delta r_{\mathrm{stack}}, k, \delta r^0_{\mathrm{coax}}, \delta r^{c,low}_{\mathrm{coax}}, \delta r^{c,high}_{\mathrm{coax}}, \delta r^{low}_{\mathrm{coax}}, \delta r^{high}_{\mathrm{coax}})\ f_4(\theta_4, a_{\mathrm{coax},4}, \theta^0_{\mathrm{coax},4}, \Delta\theta^*_{\mathrm{coax},4}) \\ &\times \left(f_4(\theta_1, a_{\mathrm{coax},1}, \theta^0_{\mathrm{coax},1}, \Delta\theta^*_{\mathrm{coax},1}) + f_4(2\pi-\theta_1, a_{\mathrm{coax},1}, \theta^0_{\mathrm{coax},1}, \Delta\theta^*_{\mathrm{coax},1})\right) \\ &\times \left(f_4(\theta_5, a_{\mathrm{coax},5}, \theta^0_{\mathrm{coax},5}, \Delta\theta^*_{\mathrm{coax},5}) + f_4(\pi-\theta_5, a_{\mathrm{coax},5}, \theta^0_{\mathrm{coax},5}, \Delta\theta^*_{\mathrm{coax},5})\right) \\ &\times \left(f_4(\theta_6, a_{\mathrm{coax},6}, \theta^0_{\mathrm{coax},6}, \Delta\theta^*_{\mathrm{coax},6}) + f_4(\pi-\theta_6, a_{\mathrm{coax},6}, \theta^0_{\mathrm{coax},6}, \Delta\theta^*_{\mathrm{coax},6})\right) \\ &\times f_5(\cos(\phi_3), a_{\mathrm{coax},3'}, \cos(\phi_3)^*_{\mathrm{coax}})\ f_5(\cos(\phi_4), a_{\mathrm{coax},4'}, \cos(\phi_4)^*_{\mathrm{coax}}). \end{aligned}$$

$$\begin{aligned} V &= \sum_{\mathrm{nn}} \left(V_{\mathrm{backbone}} + V_{\mathrm{stack}} + V'_{\mathrm{exc}}\right) \\ &+ \sum_{\mathrm{other\ pairs}} \left(V_{\mathrm{HB}} + V_{\mathrm{cross\_stack}} + V_{\mathrm{coaxial\_stack}} + V_{\mathrm{exc}}\right). \end{aligned}$$

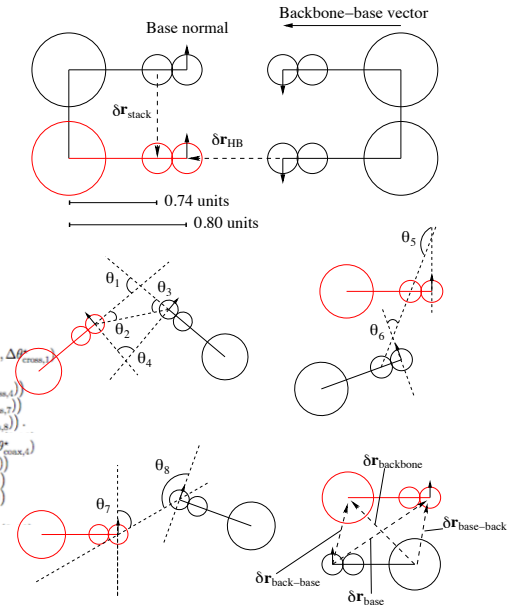

For full details see **Thomas Ouldridge, Coarse-grained modelling of DNA and DNA self-assembly, DPhil, University of Oxford, 2011.**

archer

epcc

THE UNIVERSITY OF EDINBURGH

Parallel Performance

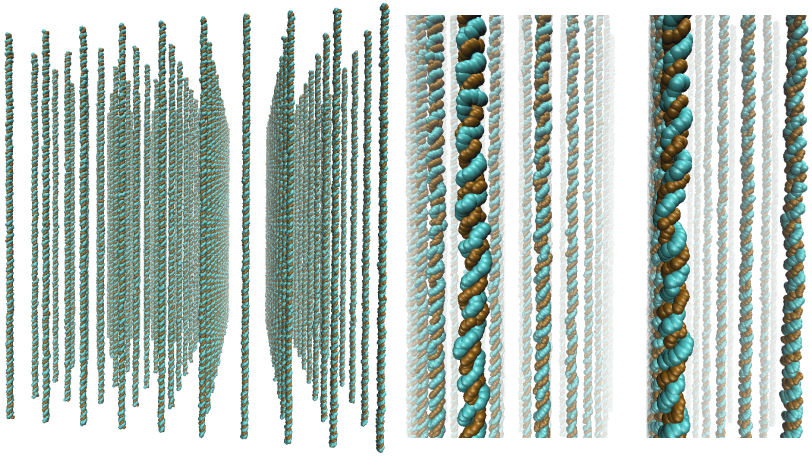

Benchmarks:

- Low density: array of 100 DNA duplexes with 600 base pairs long (see above) = 60 kbp
- High density: array of 1600 duplexes = 960 kbp

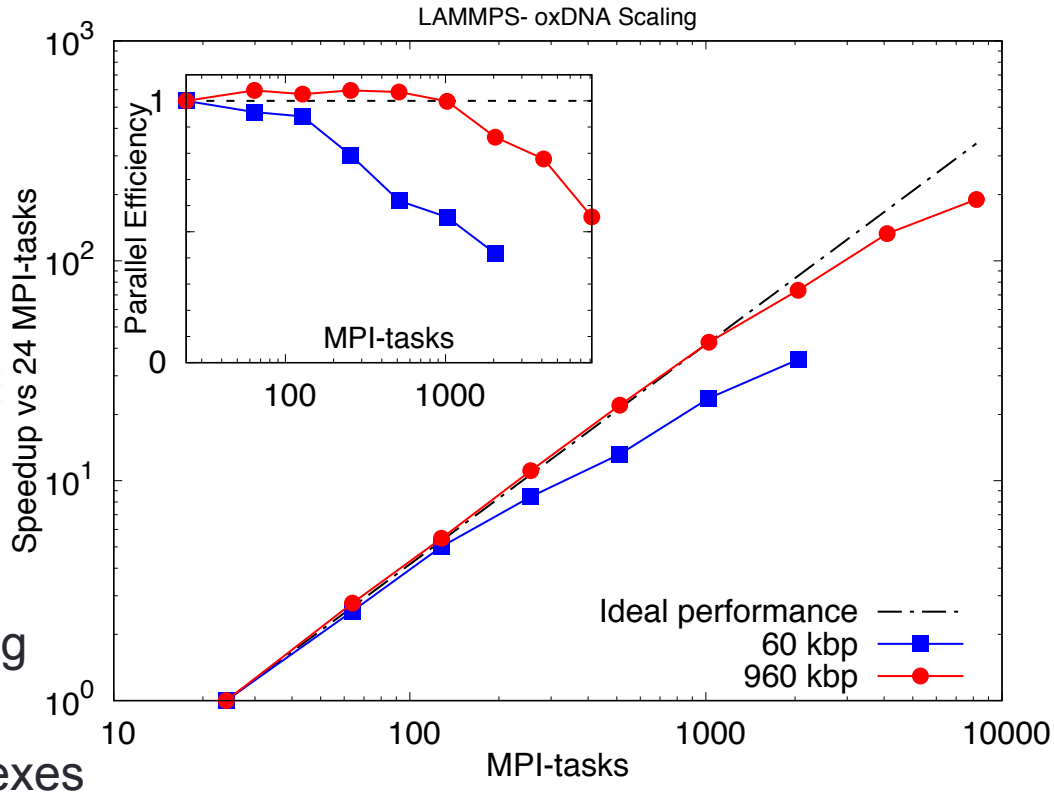

Strong scaling for 60kbp and 960 kbp up to 8192 MPI-tasks

archer

epcc

THE UNIVERSITY OF EDINBURGH

Craypat Performance Analysis

60 kbp benchmark

- Single node
  - Nlocal 5000
  - Nghost 1300
  - MPI < 5% (LMP)
  - compute 86% (LMP)
  - acos 12%
  - q_to_exyz 11%
- 2048 MPI-tasks
  - Nlocal 60
  - Nghost 225
  - MPI > 50%
  - compute 43% (LMP)

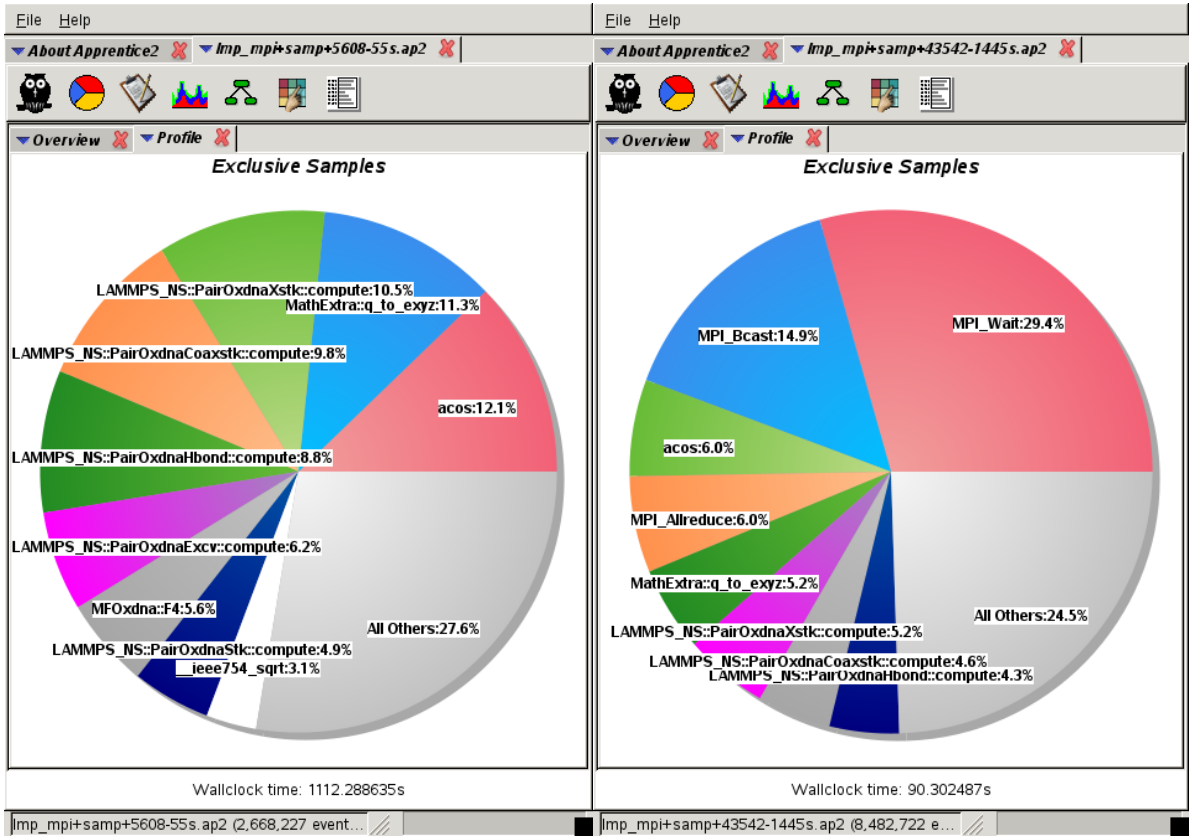

Single node (24 MPI-tasks)

2048 MPI-tasks

archer

epcc

THE UNIVERSITY OF EDINBURGH

Craypat Performance Analysis

960 kbp benchmark

- Single node
  - Nlocal 80000
  - Nghost 8300
  - MPI < 3%
  - compute 88% (LMP)
  - acos 12%
  - q_to_exyz 12%
- 2048 MPI-tasks
  - Nlocal 940
  - Nghost 480
  - MPI < 13%
  - compute 82% (LMP)

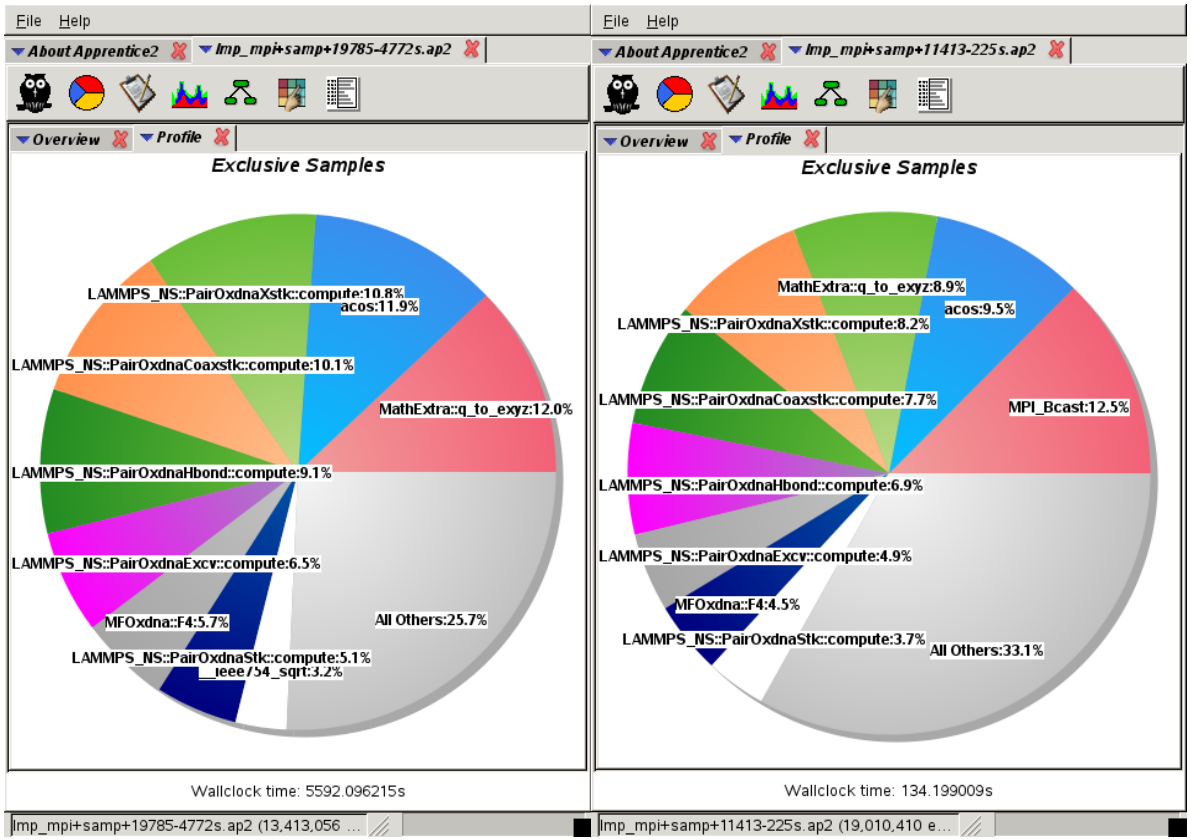

Single node
(24 MPI-tasks)

2048 MPI-tasks

archer

epcc

THE UNIVERSITY OF EDINBURGH

Applications

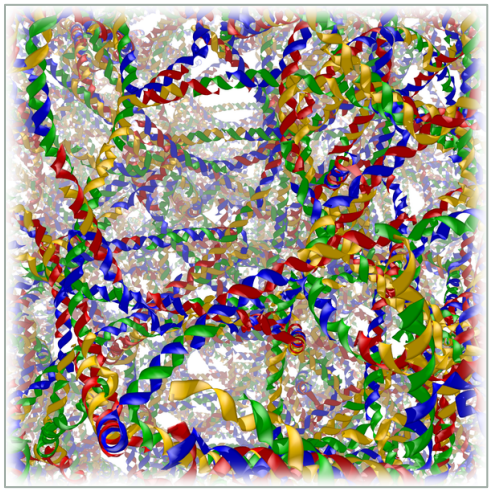

DNA gel

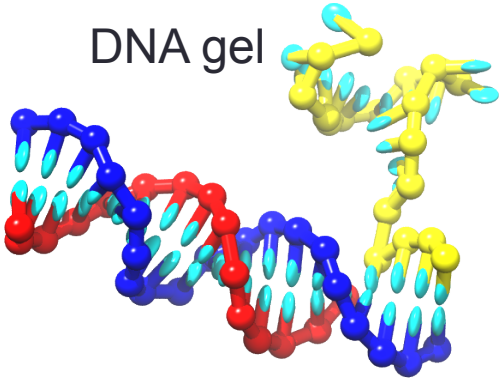

Stand displacement

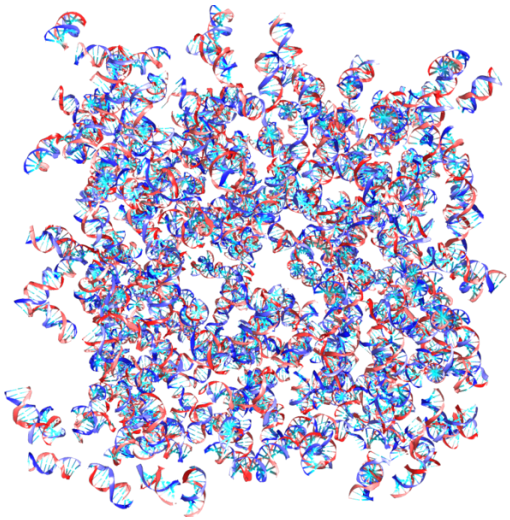

Liquid-crystalline states of DNA

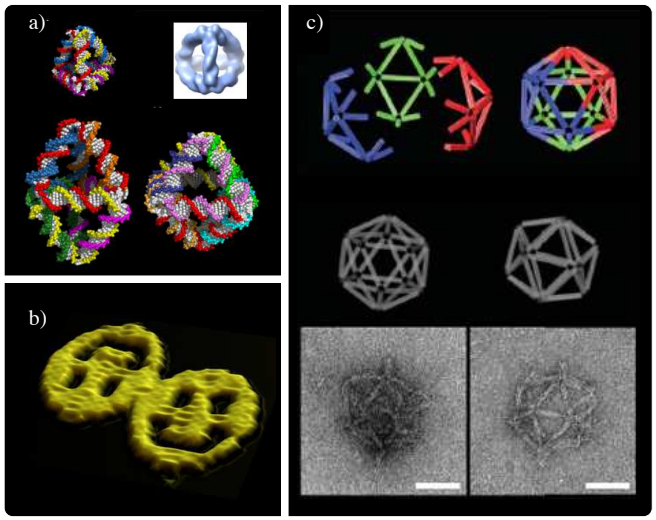

DNA nanostructures

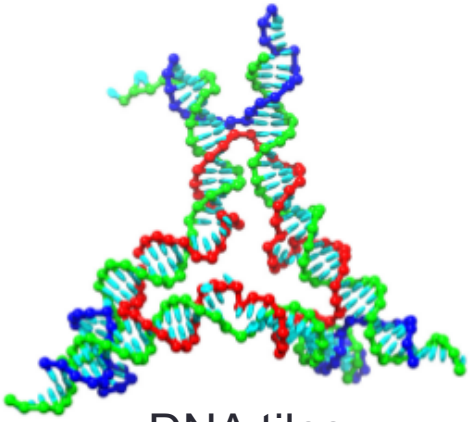

DNA tiles

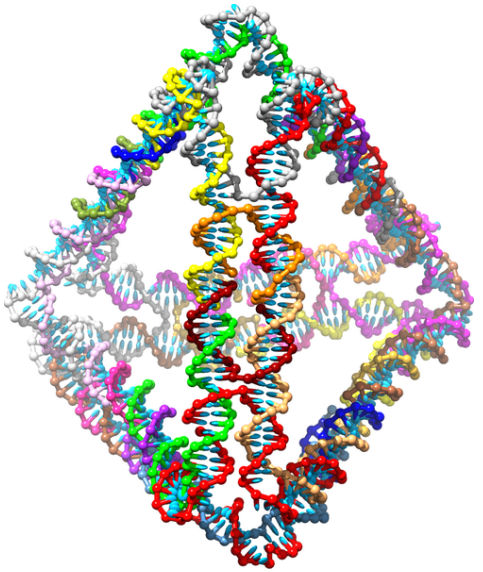

DNA tetrahedra

archer

epcc

THE UNIVERSITY OF EDINBURGH

Code Distribution

- LAMMPS version via **CCPForge**
  - https://ccpforge.cse.rl.ac.uk/gf
  - Project: **Coarse-Grained DNA Simulation** (cgdna)
  - Anonymous subversion access

    svn checkout https://ccpforge.cse.rl.ac.uk/svn/cgdna
- In the near future also as **LAMMPS USER-package** with extended documentation
- Standalone version from https://dna.physics.ox.ac.uk

http://www.archer.ac.uk/training/

- Face-to-face courses
  - timetable, information and registration
  - material from all past courses

- Virtual tutorials & webinars
  - https://www.archer.ac.uk/training/virtual/
  - timetable plus slides and recordings
  - please leave feedback on previous tutorials after viewing material

Goodbye!

Thanks for attending

Please leave feedback at:
www.archer.ac.uk/training/feedback/online-course-feedback.php

archer

epcc

THE UNIVERSITY OF EDINBURGH