CSE Computational Science & Engineering Department

The DL_POLY_3 Package: Scalability and Performance

I.T. Todorov[1,2] and W. Smith[1]

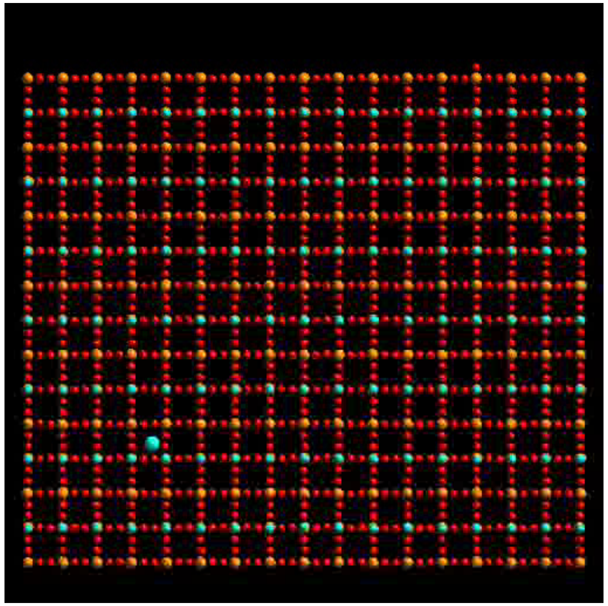

1. CSE Department
CCLRC Daresbury Laboratory

2. Department of Earth Sciences
University of Cambridge

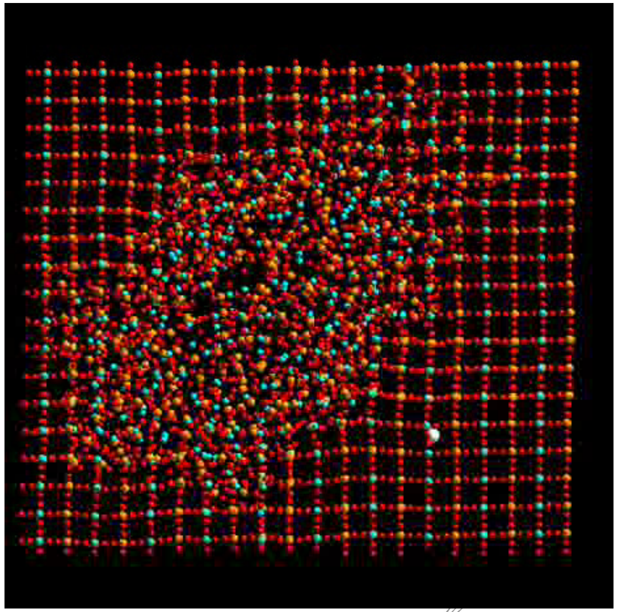

CCLRC

What is DL_POLY

- General purpose Molecular Dynamics simulation software
- Provided in open source code to licence holders
- Licence is free of charge to academia researchers
- Provided with a comprehensive, interactive manual
- Provided with test cases
- It is a PROJECT
  - Support and Development
  - wiki forum

DL_POLY Flavours

- DL_POLY_2 (version 16)
  - **Replicated Data** parallelisation, limits up to ≈30,000 atoms with good parallelisation to 16-64 (count $2\times k$) processors (system dependent)
  - Full force field and molecular description
  - Written in modularised F77/F90 (+MPI), no external dependencies
  - Fixed format reading with somewhat rigid semantics
- DL_POLY_3 (version 6)
  - **Domain Decomposition** parallelisation, limits up to $\approx 2.1\times10^9$ atoms with inherent parallelisation (any high processor count 2^k)
  - Full force field and molecular description but no rigid body description
  - Written in modularised free formatted F90 (+MPI) with fairly rigorous code syntax (NAGWare verified), no external library dependencies
  - Free format semantically approached reading with some fail-safe features (but fully fool-proofed)
  - CML enabled version exists but depends on external software

Parallel Computing Guide-lines

- Load Balancing
  - Equal division of work among processors
  - Equal division of memory requirement
  - Concurrent processor use

- Communications
  - Minimise number of messages
  - Maximise information per message
  - Local communication rather than global
  - Asynchronous communication

Replicated Data - I

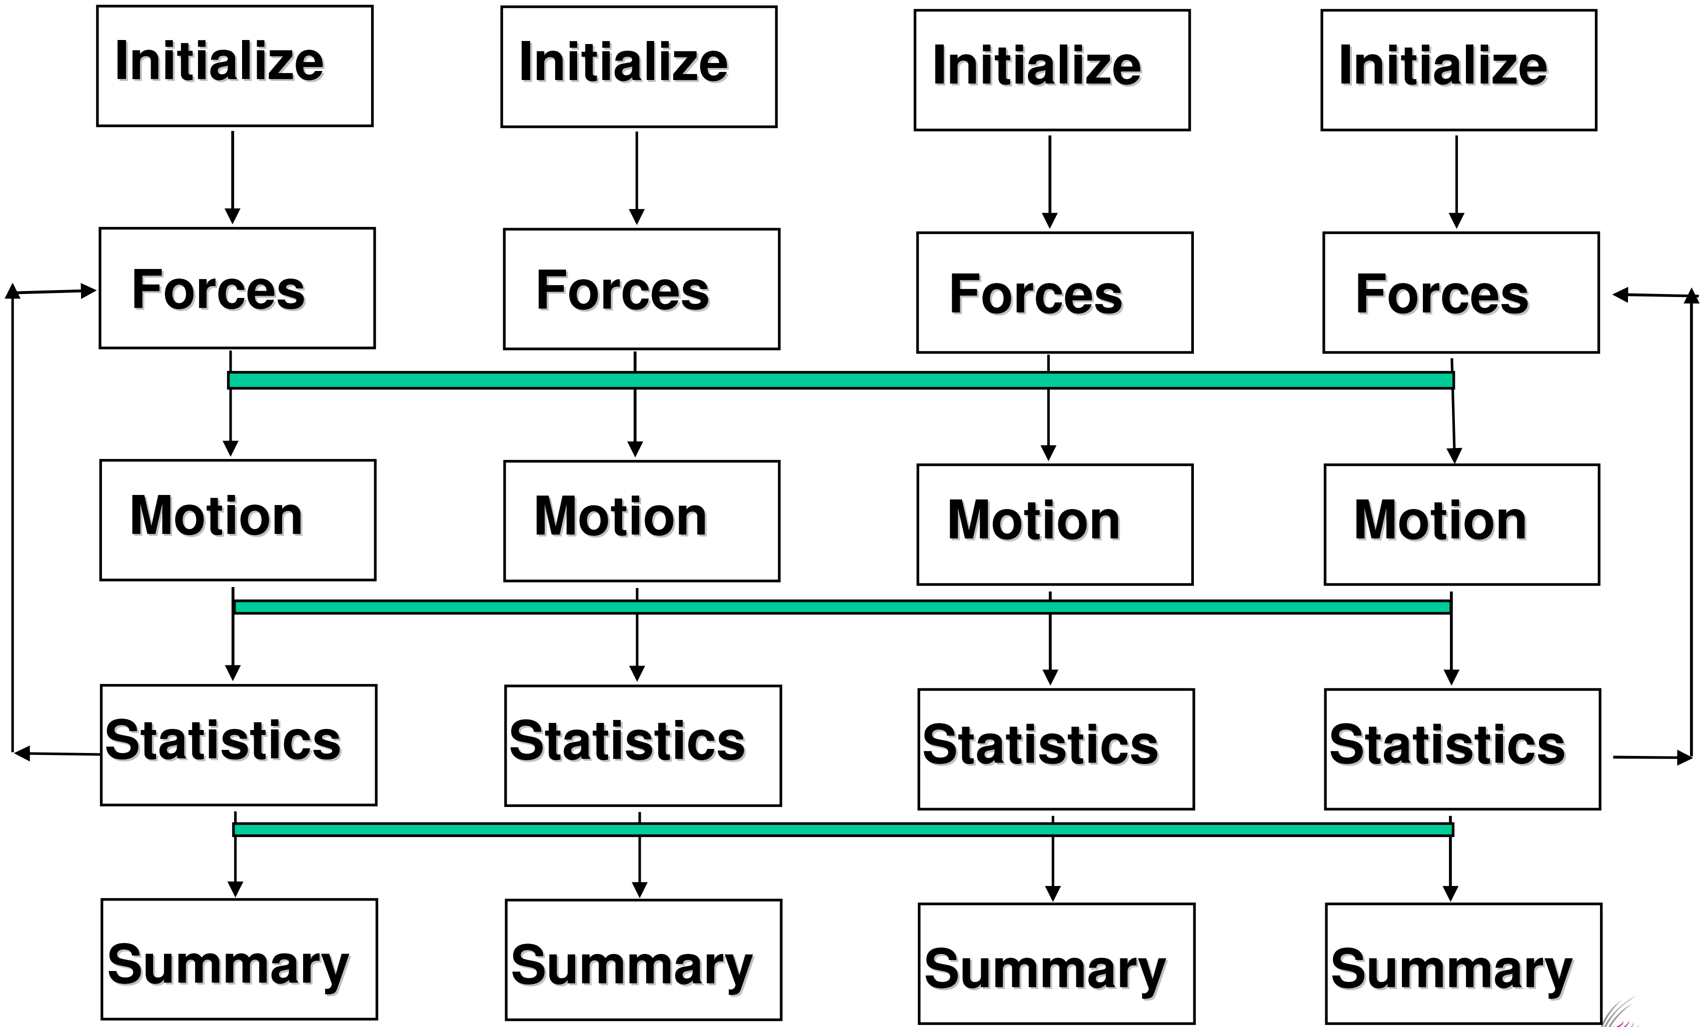

Replicated Data - II

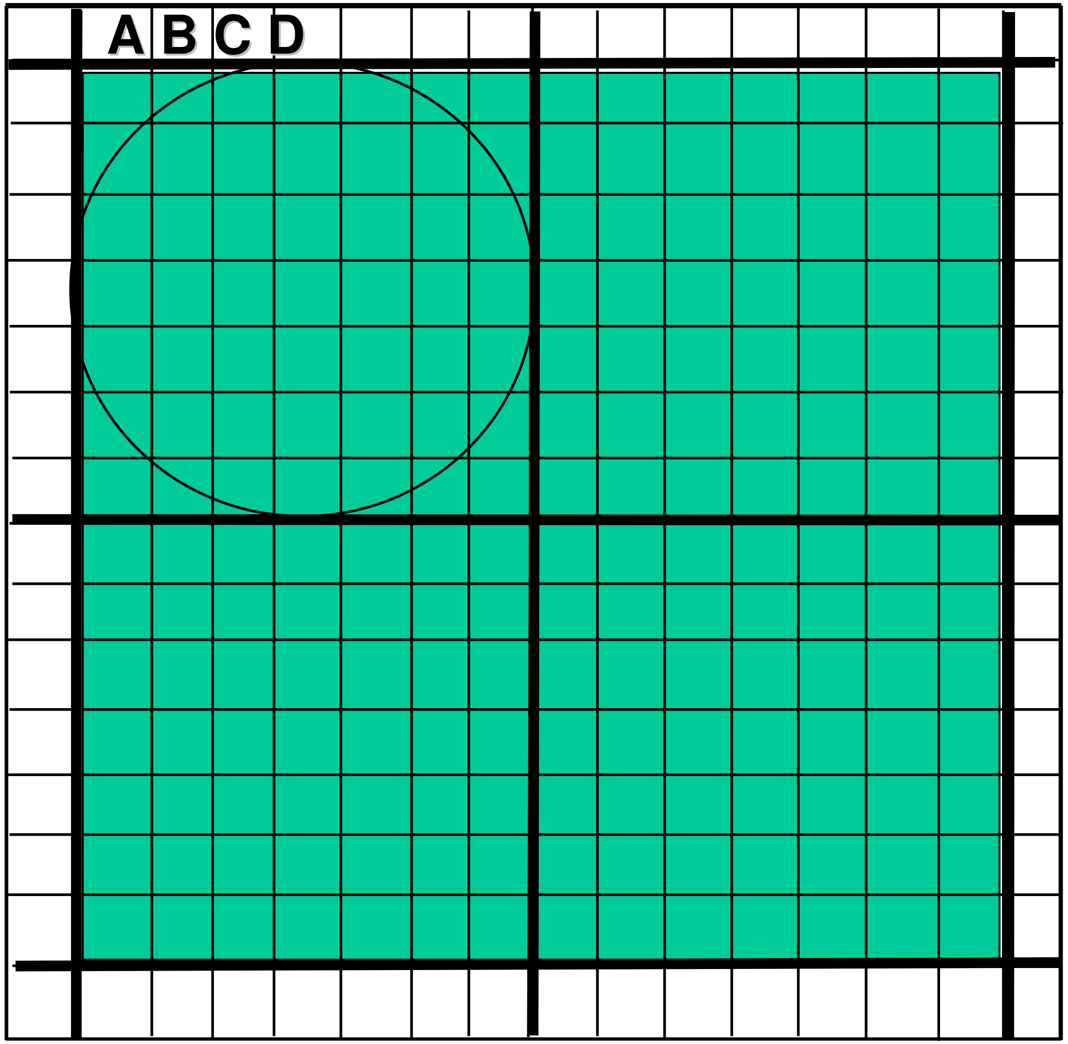

- No memory distribution (performance overheads and limitations increase with increasing system size)
- Functional decomposition of the workload
- Cutoff ≤ 0.5 min system width
- Extensive global communication (extensive overheads increase with increasing system size)

Domain Decomposition - I

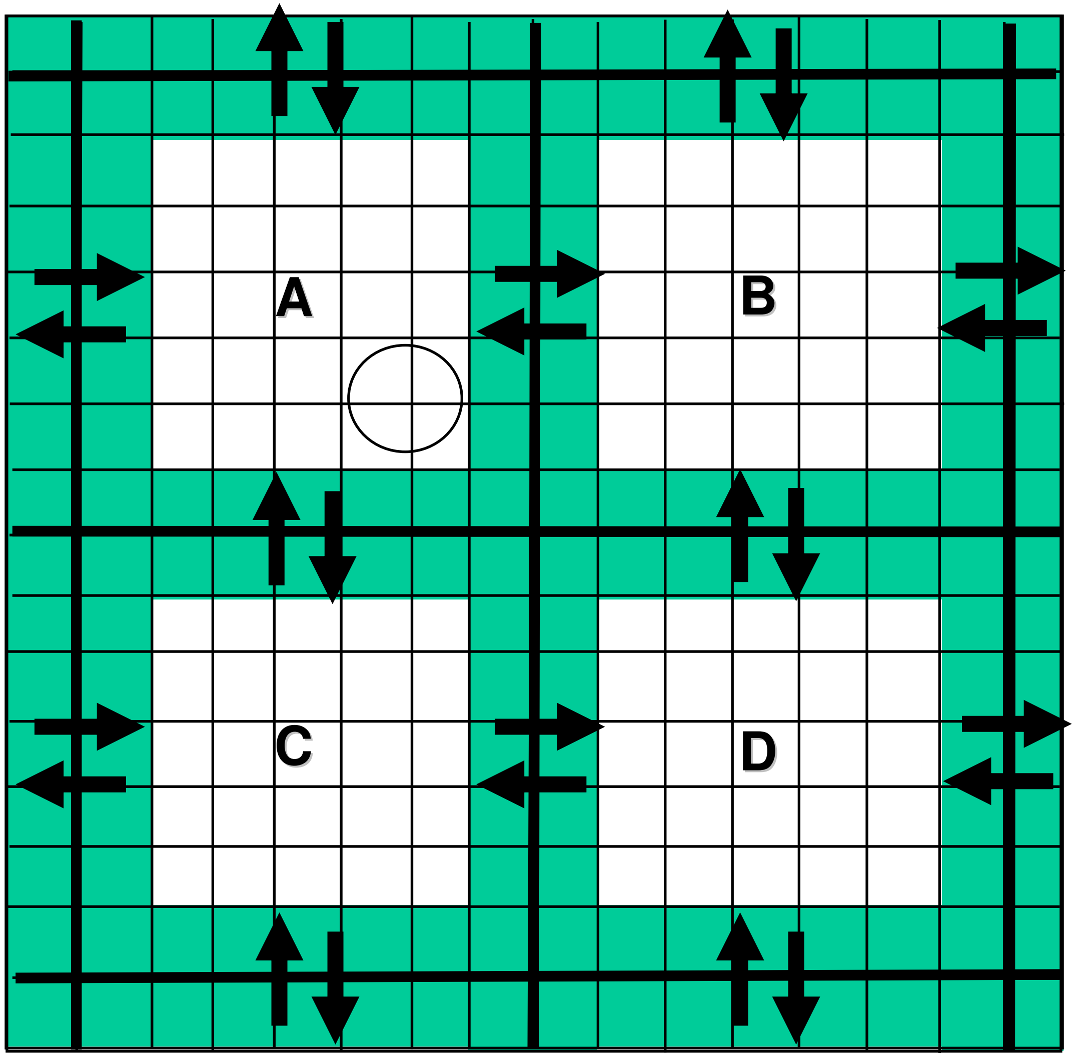

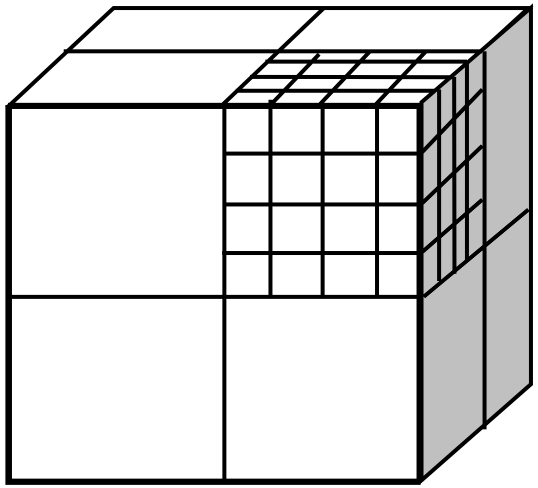

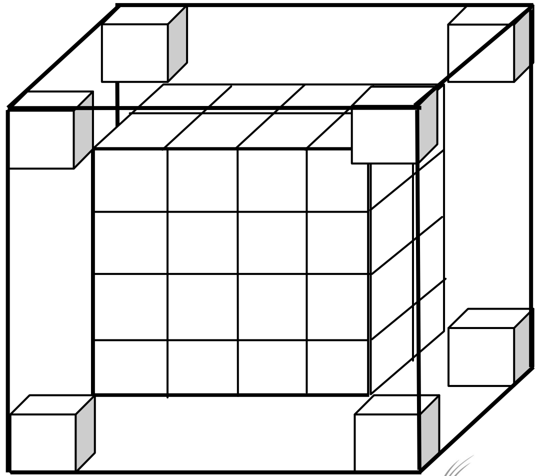

Domain Decomposition – II

- Linked lists provide an elegant way to scale short-ranged two body interactions from $O(N^2/2)$ to $\approx O(N)$. The efficacy increases with increasing link cell partitioning – as a rule of thumb best efficacy is achieved for cubic-like partitioning with number of link-cells per domain ≥ 4 for any dimension.
- Linked lists can be used with the same efficiency for 3-body (bond-angles) and 4-body (dihedral & improper dihedral & inversion angles) interactions. For these the linked cell halo is double-layered and since $cutoff^{3/4\text{-}body} \leq 0.5*cutoff^{2\text{-}body}$ which makes the partitioning more effective than that for the 2-body interactions.
- Our SPM Ewald summation uses 3D FFT procedures (DAFT by Ian Bush) implemented to fit in the DD and LC spliting of the MD box so that overall scaling is $\approx O(N \log N)$.

Functionality Differences

DL_POLY_2

- Rigid bodies
- Standard Ewald, HK_Ewald, Neutral groups
- Multiple timestep (not precise for NED)
- Truncated octahedral and rhombic dodecahedral periodic boundaries (geometries cannot be DD partitioned)
- Can run on any $2 \times k$ number of processors

DL_POLY_3

- Variable timestep, Defect Detection, Pseudo Thermostat, Extended Coulombic eXclusion, densvar, rlxtol, mxshak, PMF to all ensembles (CONTROL)
- MTK NPT and NσT ensembles, Langevin thermostats
- RDF look-up (FIELD)
- Atom displacements from original positions in HISTORY
- Can run "only" on 2^k number of processors

DL_POLY_3.CML

- Benefits
  - Metadata information – data about the data
  - Built in semantics and native hierarchy of data structures allows for faster search of data
  - Avoids fixed data/document formatting
    - Provides common ground for scientific simulation codes and supplementary software – INTEROPERABILITY
    - Avoids the use of parsers
  - Facilitates of database incorporation
  - Can be transformed to XHTML and browsed

- Drawbacks
  - Not (intended to be) human readable
  - Increases the volume of data - filesize

Benchmarking on HPC*x*

- **Solid Ar** : L-J interactions between Ar-Ar
- **NaCl** : vdW interactions between Na^+-Na^+, Na^+-Cl^- and Cl^--Cl^-, and Coulomb forces between the ions.
- **SPC Water** : L-J interactions between O^--O^-, Coulomb forces between the ions and three constraints per water molecule: O-H1, O-H2 and H1-H2.
- To keep the numeric algorithms the same with different processor counts all test cases have been constructed recursively from previously equilibrated initial structures so that the work load and memory distribution per CPU is kept constant.
- Constant Ewald precision and constraint algorithms do not scale linearly with system size.

Performance on HPC*x* - *1*

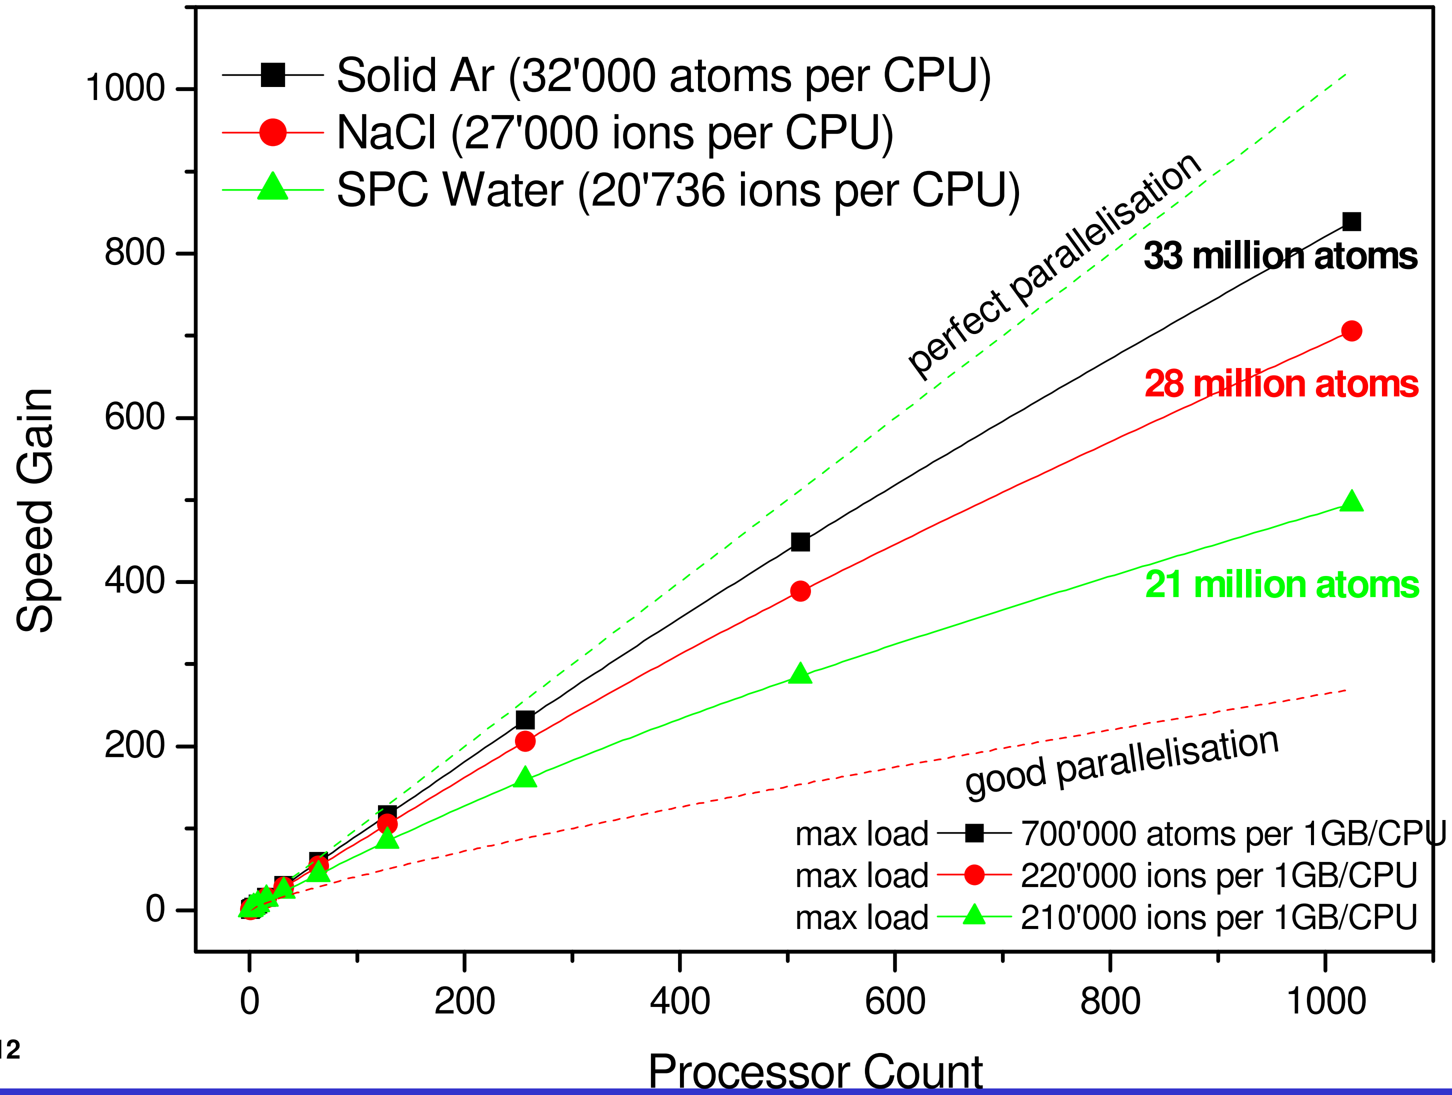

Performance on HPC*x* - 2

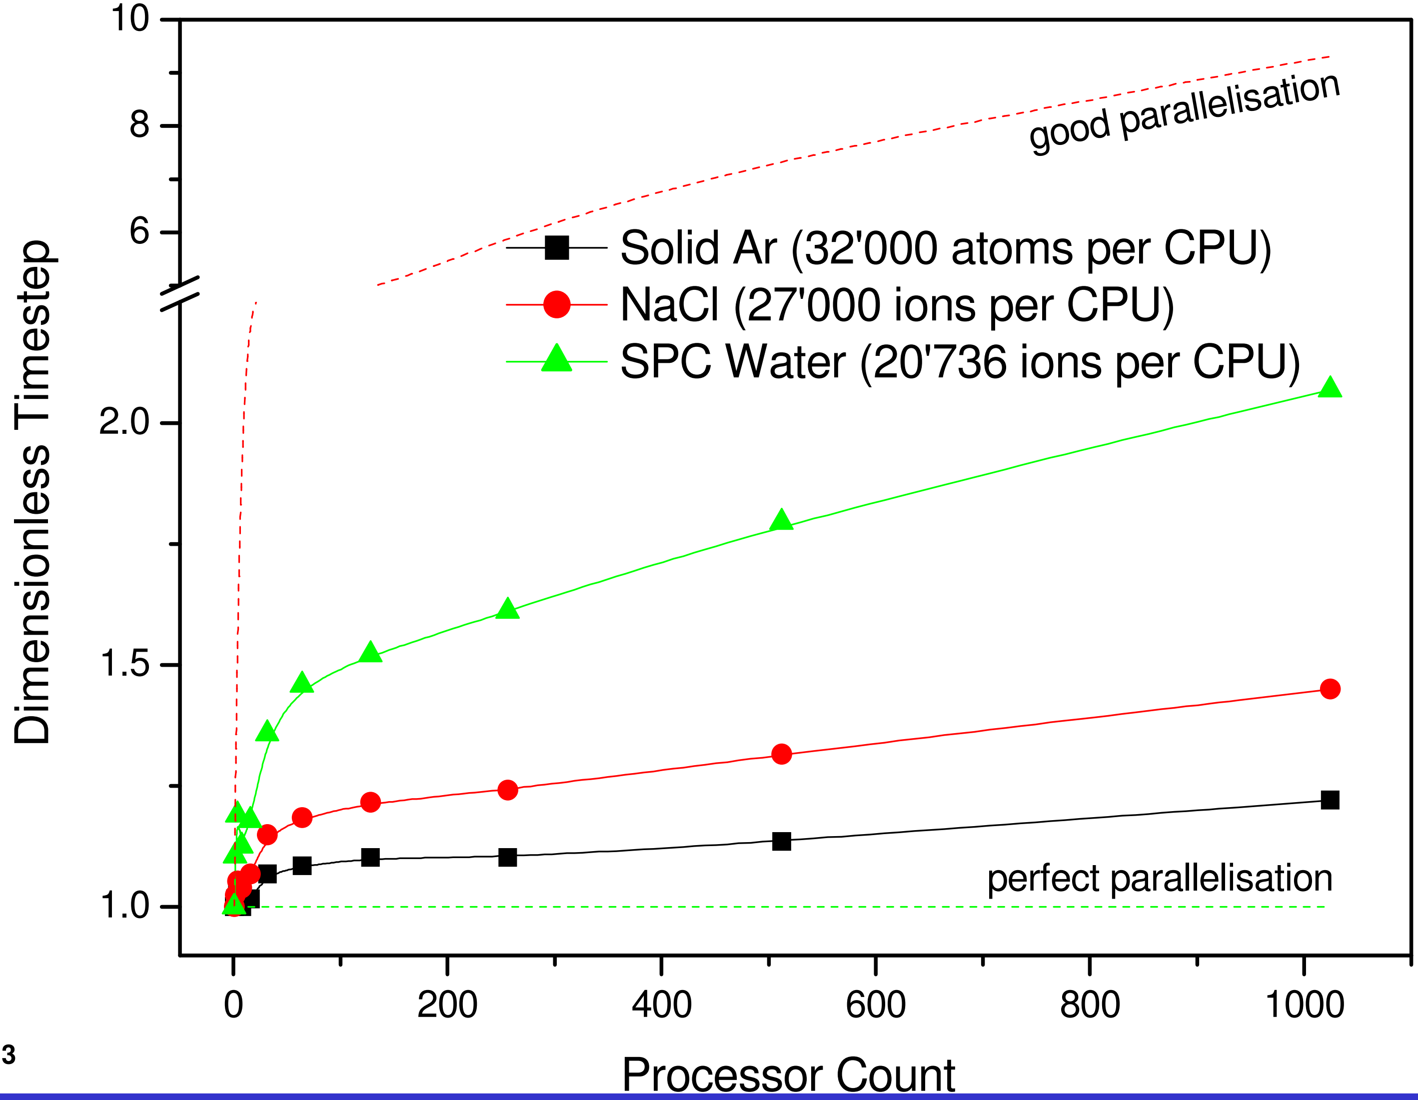

Linear loss regime from 64 to 1024 CPUs.
3.6×10^{-8} timesteps per short range interacting particle per CPU.
4.9×10^{-8} timesteps per long range interacting particle per CPU.
19.1×10^{-8} timesteps per constraint per CPU.

Performance on HPC*x - 4*

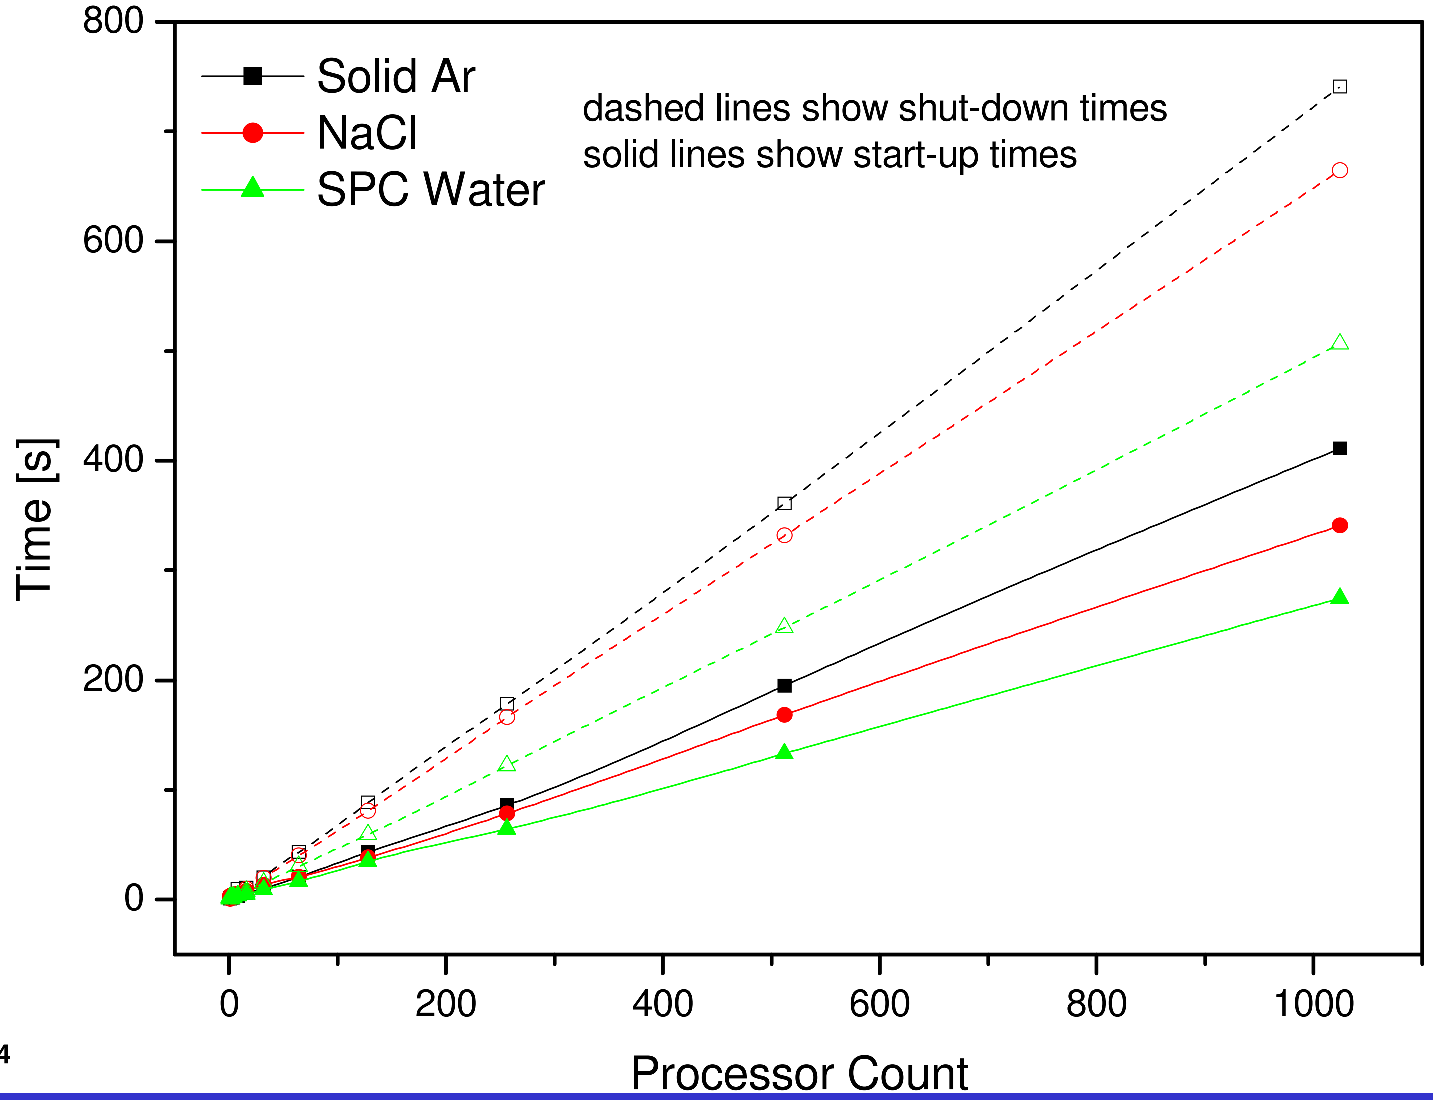

Phase2A Proof of Concept

- 300 763 000 NaCl with full SPME electrostatics evaluation
- 1024 processor units
- Start up time 1 hour
- Timestep time 67.3 seconds
  - FFT evaluation 55.3 seconds
  - vdw evaluation 2 seconds
  - integration 4 seconds
  - DD comms 6 seconds
- In theory the system can be seen by the eye. Although you would need a very good microscope – the MD cell size for this system is 2μm along the side and as the wavelength of the visible light is 0.5μm it should be theoretically possible.

FURTHER INFORMATION

- Information

  http://www.cse.clrc.ac.uk/msi/software/DL_POLY/index.shtml
- Forum

  http://www.cse.clrc.ac.uk/disco/forums.shtml