

## The DL\_POLY Project

W. Smith and T. Forester  
TCS Division, Daresbury Laboratory,  
Daresbury, Warrington WA4 4AD

### Background

The DL\_POLY project, one of the last to be funded by the former Science and Materials Computing Committee of the EPSRC, is now beginning to make a real impact on molecular dynamics simulation in the UK. The project was conceived as a flagship for CCP5, with the aim of developing a general purpose parallel simulation package to exploit the new generation of distributed memory parallel computers that were appearing with great promise in the early 1990s.

Work began on DL\_POLY in 1992 at Daresbury Laboratory, where the Intel iPSC/860 computer was readily available as a development platform. The SMCC stepped in later that year and sponsored a PDRA to complete its development. This permitted two years of intensive effort, resulting in the first public release of DL\_POLY (Version 1.1) which is described below. At the time of the SMCC award little was known about the procurement which later became the T3D supercomputer at Edinburgh. The appearance of that machine, which is ideally suited to the programming strategy of DL\_POLY, encouraged the SMCC to support the project for a further two years. We are already well on the way to producing an enhanced version of the original code, as we describe below.

### DL\_POLY Version 1.1

As mentioned above DL\_POLY is a parallel Molecular Dynamics simulation package capable of simulating a wide range of molecular and atomistic systems. Version 1.1 was made generally available in October 1994 and currently there are over 40 registered users world wide. The code was based on the Intel iPSC/860 architecture, which is a distributed memory machine with an underlying hypercube communications network. This communication network permits highly efficient global summation algorithms and this encouraged the adoption of a *Replicated Data* strategy for the parallelisation of the molecular dynamics. This strategy, we anticipated, would be reasonably efficient up to perhaps 100 nodes (we had 64) provided the simulated system was significantly large (~2000 atoms or more). However hardware developments have overtaken us here and we have since found the method to be highly efficient on 256 nodes of the T3D (see performance graph). Added to this, the RD strategy is very easy to scale down to a single processor, and in fact most of our users run the code on single processor machines and have found the adaptation straightforward. We hope that the temptation to upgrade to a multiprocessor will be overwhelming, since it can be

accomplished trivially easily - the same data files can be used!

The range of systems DL\_POLY can simulate is large. We are adding new features all the time. However, Version 1.1 is static and has the following capabilities

- all common periodic boundary conditions (including non-periodic)
- all common short-ranged pair potentials (Lennard-Jones, Born-Huggins-Meyer etc).
- atomic partial charges handled by the Ewald sum, direct summation or a distant dependent dielectric.
- rigid bond constraints (the SHAKE algorithm)
- extensible bonds
- three body and four body potentials (i.e. valence angle and dihedral angle potentials)
- external fields
- microcanonical and canonical ensembles.

Typical systems that can be simulated using version 1.1 are

- Simple atomic systems and mixtures e.g. Ne, Ar, Kr, etc.
- Simple unpolarisable point ions e.g. NaCl, KCl, MgO, etc.
- Rigid molecules with partial charges e.g. H<sub>2</sub>O, CCl<sub>4</sub>, SF<sub>6</sub> etc.
- Polymers with partial charges and a mixture of extensible and rigid bonds e.g. C<sub>n</sub>H<sub>(2n+2)</sub>, proteins etc.
- Macromolecules, Zeolites, and biological systems.
- Mixtures of all the above

Readers interested in finding out more about DL\_POLY are invited to access the World Wide Webb page for the package as described in the following section.

In Version 1.1 there are no facilities for rigid body dynamics, beyond using simple bond constraints. This deficiency will be dealt with in later releases. The other point worth mentioning is that DL\_POLY does support a multiple timestep algorithm, an essential for large scale simulations.

## Obtaining the Source Code

To obtain a copy of DL\_POLY it is first necessary to obtain a licence from Daresbury Laboratory. A copy of the licence form may be obtained in two ways: either by selecting the licence button on the World Wide Web page

[http://www.dl.ac.uk/TCSC/Software/DL\\_POLY/main.html](http://www.dl.ac.uk/TCSC/Software/DL_POLY/main.html)

and downloading and printing the file; or by using ftp to copy the postscript file from the CCP5 Program Library at Daresbury Laboratory in the following manner:

```
-> move to the desired directory on YOUR machine,  
-> type: ftp 148.79.80.10 or ftp ftp.dl.ac.uk  
-> enter userid: anonymous  
-> enter passwd: (use your name and site)  
-> change to the CCP5 directory: cd ccp5  
-> change to the DL_POLY directory: cd DL_POLY  
-> type: binary (for postscript files)  
-> type: get LICENCE.ps.Z  
-> type: quit
```

The licence file will need to be uncompressed (using the unix uncompress command) before printing.

Once you have obtained the licence form you must sign it and return it to: Dr W. Smith or Dr. T.R. Forester, by post, at the following address.

DL\_POLY Program Library  
Theory and Computational Science  
EPSRC Daresbury Laboratory  
Daresbury  
Warrington WA4 4AD  
England

When the signed licence has been received DL\_POLY source code will be sent by ftp. To enable this we will require a *guest id and password on your machine*. Please note we will not create accounts on any of our machines for this purpose.

These restrictions do not apply to the DL\_POLY documentation. The DL\_POLY Manual will be freely available via World Wide Web or ftp, in the same manner as the licence form. The DL\_POLY Reference Manual will be available by the above ftp procedure only.

Please note that Daresbury Laboratory is the sole centre for the distribution of DL\_POLY and copies obtained from elsewhere will be regarded as illegal and will not be supported.

## DL\_POLY and the Cray T3D

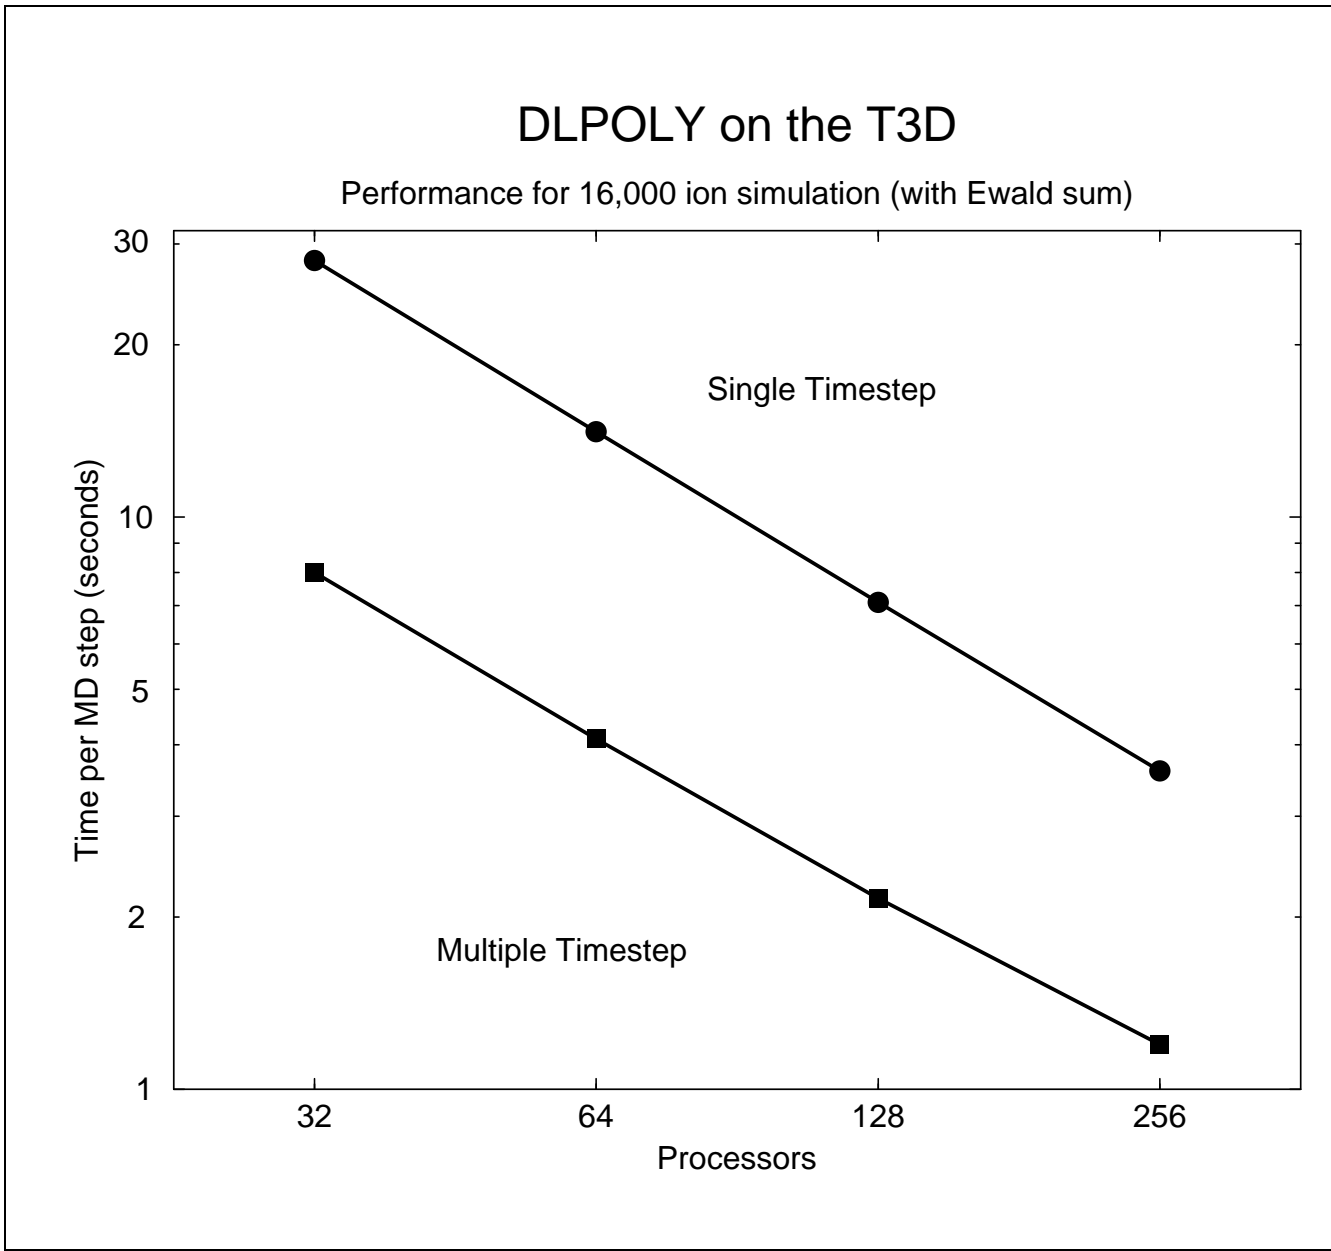
Over the past few months most of our work on the DL\_POLY project has focussed on a version for the Edinburgh Cray-T3D. The adaptation of the code for the T3D was greatly accelerated by the assistance of the Edinburgh Parallel Computing Centre, in the person of Alan Simpson. In the event, the practice we had adopted in DL\_POLY, of localising the communications inside a handful of short and easily identified routines, greatly aided the process. A full SHMEM version has been implemented and for suitably large systems (i.e. those with approx 15,000 atoms or more) the scaling performance up to the full capacity of the machine (256 processors) is very encouraging. With 256 processors a parallel efficiency of approximately 85% was obtained for a 16,000 atom case that included an Ewald sum converged to 6 significant figures. Higher efficiencies were obtained with smaller number of processors (*viz.* 128 or 64). The program runs at approximately 10-15 MFlops per processor, giving a performance of between 2-3 GFlops on 256 processors allowing a turn around of approximately 1.2 seconds per MD time-step for the 16000 ion Ewald sum simulation.

It is a very fortunate accident history that the Cray T3D became available just as the first version of DL\_POLY was reaching fruition. As a result, the code became the natural choice for simulators seeking T3D time. The T3D code is an integral part of the “Long time-scale Molecular Dynamics” group of the Materials Consortium. (We should point out however, that for operational reasons this version of the code is currently available to *Consortium members only*. With this code good progress is being made on a number of key scientific projects including:

1. the capture and transport of potassium by the naturally occurring antibiotic valinomycin at a model membrane / water interface (Tim Forester, Daresbury);
2. diffusion of hydrocarbons in Zeolites (Paul Barram, Bath);
3. diffusion of alkali ions in silicate glasses (the “mixed alkali” effect) (Bill Smith, Daresbury);  
and
4. amorphous polyelectrolyte systems (Rob Bell, Royal Institution).

These projects are at different stages, though a great deal has been done on the valinomycin and glass projects, about which we hope to report in later newsletters.

In addition to these projects, others are waiting to commence which will also require DL\_POLY as a simulation platform, these include;



1. the phase diagram for clathrate hydrates (Mark Rodger, Reading);
2. the study of defects in ceramic materials (Donnell MacDonaill, Dublin);
3. the structure and thermodynamics of oxide crystallites (Bill Mackrodt, St-Andrews); and
4. adsorption on surfaces (D.J. Tildesley, Southampton)

With the excellent projects currently running and those waiting in the wings, it is apparent that DL\_POLY is set to become a major code in UK simulations. An outcome we hope fully justifies the resources made available to it. This is clearly an exciting era for large scale molecular dynamics simulation.