



HAL
open science

A Short Description of DL_POLY

William Smith, Ilian Todorov Todorov

► **To cite this version:**

William Smith, Ilian Todorov Todorov. A Short Description of DL_POLY. *Molecular Simulation*, 2007, 32 (12-13), pp.935-943. 10.1080/08927020600939830 . hal-00515003

HAL Id: hal-00515003

<https://hal.science/hal-00515003>

Submitted on 4 Sep 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Short Description of DL_POLY

Journal:	<i>Molecular Simulation/Journal of Experimental Nanoscience</i>
Manuscript ID:	GMOS-2006-0162
Journal:	Molecular Simulation
Date Submitted by the Author:	02-Aug-2006
Complete List of Authors:	Smith, William; Daresbury Laboratory, CSED Todorov, Ilian; Daresbury Laboratory, CSED
Keywords:	DL_POLY, molecular dynamics, scientific software

SCHOLARONE™
Manuscripts

A Short Description of DL_POLY

W. Smith and I.T. Todorov

Computational Science and Engineering Department, Daresbury Laboratory,
CCLRC, Daresbury, Warrington WA4 4AD, United Kingdom.

Abstract

DL_POLY is a general purpose molecular dynamics simulation package with in-built parallel algorithms. It may be run on a wide selection of distributed memory parallel computers, from national supercomputers with thousands of processors, to single processor workstations and can simulate small systems with order 100 atoms, to systems with millions of atoms. This introduction provides an outline of the features of the package and the underlying methodology.

Keywords: DL_POLY, molecular dynamics, scientific software.

Word count: 5743

1. Introduction

The DL_POLY molecular dynamics (MD) package has been under continual development at Daresbury Laboratory since 1994. It was first released to the academic community in 1996, so this special issue of Molecular Simulation marks its 10th anniversary as a public code. Its prime purpose on first release was to provide a simulation package for the UK CCP5 community [1] that was capable of exploiting the emerging parallel computers, of which the Intel IPSC 860 at Daresbury was a prime example. Since then it has gained popularity all over the world and is in considerable demand. It currently exists in two forms: DL_POLY_2, which is based on Replicated Data parallelism and DL_POLY_3, which is based on Domain Decomposition parallelism. Both versions are run on major parallel platforms all over the world.

DL_POLY was arguably the first public general purpose MD packages to be written specifically for parallel computers. Practical parallel platforms began to appear in the late 1980s and Daresbury was active in developing parallel algorithms at this early stage [2]. When the demand appeared for a new MD program for CCP5 we had already examined a number of possible strategies including Replicated Data (RD) [3], Systolic Loops [4] and Domain Decomposition [5,6]. The RD strategy was initially chosen because it offered the simplest approach to complex force fields, with reasonable scaling properties on platforms with up to 100 processors. The first version incorporating this strategy was developed in-house as DL_POLY_1 in 1994 and was circulated among close collaborators for testing and early

exploitation. The first generally available package was eventually released in 1996 as DL_POLY_2 and was written by T. Forester and W. Smith. These codes were intended for *distributed memory* machines and this assumption underpins all DL_POLY codes to this day.

In the following sections we outline the features available in the DL_POLY codes, including the force field specification, the parallel strategies that the codes are based on, techniques for modelling electrostatic interactions, the implementation of rigid bonds for parallel processing and the integration algorithms.

2. The DL_POLY Force Field

The DL_POLY package does not provide any particular set of force field parameters to describe the interatomic interactions, as with more specific packages such as AMBER [7], GROMOS [8] and CHARMM [9]. This is impractical given that the simulation code caters for widely disparate kinds of molecular system. It does, however, implement an enormous selection of functional forms for the interaction potentials arising in many of the force fields commonly used in molecular simulation. It is also easy, due to the structure of the software, for the user to add new potential functions and it should be noted that the user may use many different kinds of potential in the same simulation, which is not therefore confined to purely Lennard-Jones or purely Buckingham descriptions for example.

The total potential energy for DL_POLY is expressed by the formula:

$$\begin{aligned}
 V(\vec{r}_1, \vec{r}_2, \dots, \vec{r}_N) = & \sum_{i,j}^{N'} U_{pair}(r_{ij}) + \frac{1}{4\pi\epsilon} \sum_{i,j}^{N'} \frac{q_i q_j}{r_{ij}} + \sum_{i,j,k}^{N'} U_{3-body}(\theta_{ijk}) + \\
 & \sum_{i,j,k,n}^{N'} U_{4-body}(\phi_{ijkn}) + \left\{ \frac{1}{2} \sum_{i,j}^{N'} U_{FS}(r_{ij}) + \sum_i^{N'} F_{FS}(\rho_i) \right\} + \frac{1}{2} \sum_{i,j}^{N'} \{U_{Ter}(r_{ij}) + \gamma_{ij} V_{Ter}(r_{ij})\} + \\
 & \sum_{i_{bond}}^{N_{bond}} U_{bond}(i_{bond}, r_{ab}) + \sum_{i_{angle}}^{N_{angle}} U_{angle}(i_{angle}, \theta_{abc}) + \sum_{i_{dihed}}^{N_{dihed}} U_{dihed}(i_{dihed}, \phi_{abcd}) + \\
 & \sum_{i_{invers}}^{N_{invers}} U_{invers}(i_{invers}, \psi_{abcd}) + \sum_{i=1}^N \Phi_{external}(\vec{r}_i)
 \end{aligned} \tag{1}$$

DL_POLY contains all the commonly used pair potentials ($U_{pair}(r_{ij})$), including Lennard-Jones, Buckingham, 12-6, N-M and Morse potentials. The electrostatic interactions, indicated by the $q_i q_j / r_{ij}$ terms, are available as point charge and polarisable shell models, for which a variety of summation techniques may be selected (see section 5). In DL_POLY polarisation is treated with the shell model of Dick and Overhauser [10] by the adiabatic method of Fincham [11] and the relaxation model of Lindan [12]. The three-body ($U_{3-body}(\bullet_{ijk})$) and four-body ($U_{4-body}(\bullet_{ijkn})$) interactions are non-specific angular potentials (suitable for glass simulations) and again offer a variety of

1
2
3 functional forms. Many-body interactions, an increasing common
4 requirement for modelling complex systems, are available in the Finnis-
5 Sinclair form [13] for metals (terms $U_{FS}(r_{ij})$ and $F_{FS}(\bullet_{ij})$) and Tersoff form [14] for
6 covalent systems (terms $U_{Ter}(r_{ij})$, \bullet_{ij} and $V_{Ter}(r_{ij})$). All these mentioned forms are
7 for non-bonded *inter*-molecular interactions.
8
9

10 For *intra*-molecular interactions DL_POLY has a wide selection of bond
11 potentials (U_{bond}), angle potentials (U_{angle}), dihedral angle potentials (U_{dihed}), and
12 inversion angle potentials (U_{invers}). The forms of these are taken from many
13 published force fields including AMBER [7], GROMOS [8] and CHARMM [9]
14 and Dreiding [15].
15
16

17
18 Lastly, DL_POLY permits the user to implement external force fields. This
19 capability is useful for modelling transport (e.g. conduction), or containment
20 (e.g. pores) or mechanical intervention (e.g. shearing).
21
22

23 24 3. Integration Algorithms 25

26
27 The integration algorithms in DL_POLY handle the dynamics of the system
28 being simulated. From the current positions of the atoms, the forces may be
29 calculated from the first derivatives of the potential functions outlined above
30 and used to update the atomic velocities and positions. The integration
31 progresses in a sequence of finite steps in time, each time step being of the
32 order 1~10 fs. The algorithms for this purpose in DL_POLY are based on the
33 Verlet leapfrog (LF) and velocity Verlet (VV) schemes [16].
34
35

36
37 In addition to providing a numerical solution to the equations of motion, the
38 integration algorithm also defines the thermodynamic ensemble. At the base
39 level, both LF and VV provide the NVE (constant energy ensemble), but we
40 have also implemented in DL_POLY the NVT (canonical) ensembles of Evans
41 [17], Hoover [18] (after Melchionna *et al.* [19]) and the pseudo canonical
42 ensemble of Berendsen [20]. For constant pressure work the isotropic
43 isothermal-isobaric (NPT) ensemble has available in both Hoover and
44 Berendsen forms and complemented by the anisotropic forms (NST) for
45 simulation of phase transitions in solids.
46
47

48
49 As well as the integration/ensemble algorithms DL_POLY also accepts
50 molecular structures defined by rigid bonds and, in the case of DL_POLY_2
51 only, rigid bodies. The types of molecular structures that may be
52 accommodated in a DL_POLY simulation are shown in Figure 1. It is
53 important to note that all such structures may be present in one simulation!
54
55

56
57 Rigid bonds adapt easily within the framework of LF and VV, though the well
58 known algorithms SHAKE (LF) [21] and RATTLE (VV) [22] and we have
59 devised versions of these for DL_POLY that are both parallel and appropriate
60 for the above ensembles. Considerations pertaining to the parallel SHAKE
algorithm are described in section 6.

1
2
3
4
5 Rigid bodies may be used to represent structures like aromatic hydrocarbons
6 and their derivatives, which arise in all branches of chemistry. In DL_POLY_2
7 the dynamical treatment of such entities is based on Euler's prescription [23]
8 augmented by a quaternion treatment of the orientation [24]. For the LF
9 integration scheme DL_POLY_2 employs the Fincham implicit quaternion
10 algorithm [25] and for the VV scheme the NOSQUISH algorithm of Miller *et*
11 *al.* [26] is used. The latter algorithm has the advantage of being symplectic
12 and therefore stable for long time integrations [26].
13
14

15
16 The presence of both rigid bonds and rigid bodies in the same systems raises
17 the possibility of rigid bodies linked by rigid bonds. A suite of integration
18 routines are available for this situation in DL_POLY_2. These routines are
19 derived from the QSHAKE algorithm [27] which was devised by us and is
20 able also to generate any of the ensembles described above.
21
22

23 24 25 **4. Parallelisation Strategies**

26 27 **4.1 Replicated Data Parallelism**

28
29 In the RD strategy [3] each processor of the parallel machine maintains a
30 replica of the configuration of the simulated system i.e. the coordinates $\{\underline{r}_i\}$,
31 velocities $\{\underline{v}_i\}$ and forces $\{\underline{f}_i\}$ for all atoms $\{i=1,\dots,N\}$ in the system. Each
32 processor may be thought of as running the same simulation, but replication
33 of the computational effort is avoided by assigning to each processor a subset
34 of the tasks involved. The simulation as a whole is established by
35 communicating the results of these concurrent tasks to all processors, so that
36 every processor can continue to maintain a full replica of the simulation. The
37 communication of these data is inevitably a global operation, since all
38 processors need all the data. Key points at which this global operation is
39 necessary are in the computation of forces and in the integration of the
40 equations of motion.
41
42
43
44

45
46 The principal expense in MD lies in the computation of atomic forces and this
47 is where most of the effort lay in developing DL_POLY. Calculation of *intra*-
48 molecular forces is handled through bookkeeping arrays that store the
49 identities of interacting atoms and the relevant force forms (bonds, angles,
50 dihedrals *etc.*). Parallelism is achieved by simply allocating each processor a
51 subset of bookkeeping arrays (Figure 2). The overall efficiency of this
52 approach is extremely high. The basic approach is suitable for both RD and
53 DD implementations, though there are additional complications with DD (see
54 section 4.2).
55
56

57
58 A more difficult task is the parallel distribution of *inter*-molecular forces, of a
59 type similar to van der Waals (VDW) interactions. The approach adopted is to
60 construct a distributed Verlet Neighbour List [16] based on the Brode-
Ahlricks decomposition of the pair force matrix [29], which is built

1
2
3
4 independently on each processor so that no pair interaction is replicated
5 (Figure 3). This is almost ideal parallelism and therefore highly efficient, since
6 each processor can build a list without communicating with others, and the
7 resulting list is a fraction of that for the whole system. In keeping with the
8 Verlet method, the list must be updated at intervals during the simulation,
9 which is accomplished by monitoring the distances atoms move and updating
10 the list when a tolerance is exceeded. The neighbour list enables efficient
11 identification of interacting atom pairs. Intrinsic to our implementation is a
12 strategy for omitting VDW interactions from the neighbour list if the atom
13 pair concerned is already part of an *intra*-molecular interaction.
14
15

16
17 For more complicated intermolecular interactions, such as the Finnis-Sinclair
18 [13] or Tersoff [14], which are density dependent, DL_POLY_2 employs a
19 Link Cell [29] approach. More problematical however, is the treatment of long
20 ranged forces via the Ewald method [30]. This merits special consideration
21 and is described in section 5.
22
23

24
25 It is apparent that distribution of the force calculations under RD results in a
26 partial description of the full force field on any one processor. Thus a global
27 communication step is required to establish full replication of the force data
28 everywhere. In DL_POLY_2 this is accomplished by a global summation of
29 the atomic force arrays. This is a relatively expensive step in communication
30 terms, and is a principal reason why the RD strategy is not recommended for
31 computers with high processor counts. However, experience shows that the
32 impact of this step is very dependent on the nature and size of the simulation
33 being undertaken and simulations that are large (approaching say 30,000
34 atoms) are known to scale quite well, sometimes up to 250 processors with
35 good communication hardware.
36
37

38
39 The second opportunity for exploiting parallelism under RD occurs in the
40 numerical integration of the equations of motion. Thus in DL_POLY_2 each
41 processor integrates the motion for a subset of atoms only. Replication of the
42 coordinate and velocity arrays is established once again by a global
43 communication using systolic loops. Integration of the equations of motion
44 frequently involves employing the SHAKE algorithm for systems with rigid
45 bonds [21]. This is an issue which is dealt with in section 6.
46
47
48

49
50 Though DL_POLY_2 proved itself beyond its original design in terms of
51 system sizes and processor counts, it became apparent in the early 2000's that
52 the underlying RD strategy was not appropriate for the emerging platforms
53 with thousands of processors, so work began on a DD version that would
54 scale more efficiently on such platforms. This eventually became the
55 DL_POLY_3 code that appeared in 2004 and was written by I.T. Todorov and
56 W. Smith [31].
57
58

59 4.2 Domain Decomposition Parallelism 60

1
2
3 The DD strategy is radically different from RD. Under the DD approach the
4 simulation cell is divided spatially into quasi-independent domains which are
5 allocated to individual processors.
6

7 If follows immediately that the simulated system must be reasonably isotropic
8 if a reasonable degree of load balancing amongst the processors is to be
9 achieved. The spatial division naturally does not recognise molecular entities,
10 which are therefore usually divided between processors, creating special
11 communication difficulties. The implementation of DD in DL_POLY_3 is
12 based on Hockney and Eastwood's Link Cell algorithm [29], which was
13 adapted for parallel use by Pinches *et al.* [5] and Rapaport [6]. A Link Cell
14 approach is not entirely essential for DD, but it provides useful constructs to
15 aid its implementation and yields order N scaling for large numbers of atoms,
16 N. The structural aspects of DD are shown in Figure 4.
17
18
19

20
21 Spatial partitioning for DL_POLY_3 demands that the number of processors P
22 must have the form $P=2^n$ where n is an integer. This is a requirement arising
23 from the Ewald calculations described in section 5.2. The MD cell is most
24 often divided into near-cubic domains, though exception is made for systems
25 with slab geometries to help achieve load balance. Each domain is then sub-
26 divided into link cells according to the normal prescription, in which the
27 width of a link cell must be greater than the cut-off distance applied to all
28 *interatomic* interactions. (In the context of DL_POLY_3, this criterion must
29 also include the 1-4 distance in the *intramolecular* dihedral potentials.)
30 Ideally, these requirements should lead to better than a 3x3x3 link cell
31 partitioning of the domain in the three principal directions. DL_POLY_3 can
32 handle fewer link cells per domain than this, but it raises major efficiency
33 issues arising from the construction of the 'halo data'.
34
35
36
37

38 The 'halo data' represents the construction around each domain of a partial
39 image of all neighbouring domains so that calculation of all the forces
40 relevant to a domain can take place (Figure 5). In DL_POLY_3 this amounts to
41 the transfer of the atomic coordinates of all atoms located in link cells at the
42 boundaries of a domain to the processors managing the neighbouring
43 domains. This is a six-fold transfer operation that moves data in directions
44 North, South, East, West, Up and Down of each domain. These six transfers
45 do not happen concurrently, since some data sorting is necessary to populate
46 the 'corners' of the halo data. It is apparent from the nature of the link cell
47 method, that these transfers are sufficient for a complete calculation of the
48 forces on all atoms in any domain. It is also apparent that if the domains have
49 relatively few link cells (or their shape is far from cubic), then the transfer of
50 the halo data represents the transfer of a major proportion of the contents of a
51 domain, which implies a large, possibly prohibitive, communication cost.
52 This can be avoided by running the program on fewer processors.
53
54
55
56
57

58 The transfer of the halo data is the main communication cost of the basic DD
59 strategy. After the transfer, the atomic forces may be calculated and the
60 equations of motion integrated independently on each processor. Atoms that
move sufficiently far may then be reallocated to a new domain.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Computation of *intra*-molecular forces can be accomplished, in principle, using the partitioning scheme described in Figure 2. However, there are particular complications arising from the DD scheme. There are two aspects to this: firstly the description of the molecular structures (commonly called the topology) is 'broken' by the decomposition into domains; and secondly the evolution of the system demands that the topology be partially reconstructed every time atoms move from one domain to another. In order to accomplish this, the package of data transported with each atom that leaves a domain contains not only its configurational data (position, velocity and force), but also a topological description of the bonding terms associated with the atom.

5. The Treatment of Long Ranged Electrostatic Forces

5.1 The Standard Ewald Sum

The treatment of long ranged electrostatic forces represents a particular challenge in molecular simulation. Direct summation of the Coulomb pair interactions is rarely adequate, except for the treatment of atomic clusters, so more sophisticated treatments have evolved. The main methods used in DL_POLY are based on the Ewald sum [30].

The Ewald sum casts the sum of Coulomb pair interactions into two separate sums (plus a correction term, which is computationally trivial). The first sum is a screened Coulomb sum, which resembles the Coulomb formula but each term is weighted by a screening function (the complementary error function *erfc*) which compels the sum to converge in a finite range. The second sum is a sum of structure factors, which are calculated from reciprocal space vectors, and which are again weighted by a screening function (this time a Gaussian) which guarantees a finite sum. The first sum is therefore set in *real* space, while the second is set in *reciprocal* space. The convergence of both sums is governed by a single parameter α , which defines the range of both convergence functions and is known as the Ewald convergence parameter.

The original implementation of the Ewald sum in DL_POLY_2 was a RD adaptation [32] devised by us. Later this was augmented by a similar adaptation of the Hautman-Klein-Ewald (HKE) method [33] for systems with 2D periodicity, and a partially distributed adaptation of the Smoothed Particle Mesh Ewald (SPME) [34]. A fully distributed SPME version was implemented as the primary method in DL_POLY_3.

The RD adaptation of Ewald's method [32] requires no special modifications for calculating the *real*-space components. These are treated in the same way as the van der Waals terms described above. The *reciprocal*-space terms, which are derived from a Fourier transform of the system charge density, are parallelised through an atomic decomposition: each processor is made

1
2
3 responsible for a fixed set of atoms. The method involves the global
4 summation of the structure factors associated with each reciprocal space
5 vector. The cost of this was minimised in later versions by summing the
6 structure factors for all vectors simultaneously. The same approach is taken
7 with the HKE method.
8
9

10 5.2 The Smoothed Particle Mesh Ewald Sum

11
12
13 The RD adaptation of the SPME method employs the same treatment of the
14 *real*-space terms as for the standard RD Ewald approach. However the key
15 difference is in the treatment of *reciprocal*-space, which is an interpolation of
16 the charge distribution, based on Cardinal B-splines [33], on a regular 3D grid.
17 This permits the use of a 3D Fast Fourier Transform (FFT) to calculate the
18 structure factors, which accelerates the process enormously. An important
19 consideration is how to parallelise the method. Distributing the central FFT
20 operation risks impairing the supreme efficiency of the FFT algorithm. In
21 DL_POLY_2 the decision was made to replicate the full FFT operation on all
22 processors, though the construction of the 3D charge array needed for the
23 calculation is constructed in a distributed fashion and completed by a single
24 global sum operation. This strategy is inevitably expensive in memory terms
25 but is simple to implement and does not disrupt the FFT algorithm, which
26 therefore retains its efficiency. The method has proved to work with
27 acceptable efficiency, scaling reasonably well for large simulations and
28 processor counts of order 256.
29
30
31
32
33

34 For the DD approach in DL_POLY_3, a fully distributed implementation of
35 the SPME method was essential. This was accomplished through a distributed
36 3D FFT algorithm devised by Bush [35]. Known as the Daresbury Advanced
37 Fourier Transform (DAFT), this FFT employs a domain decomposition of the
38 3D FFT arrays which maps neatly on to the DD structure of DL_POLY_3. This
39 means that all computations necessary to build the (partial) arrays can take
40 place without inter-processor communication. Furthermore all
41 communication required by the FFT algorithm is handled internally. While
42 the insertion of communication processes into the heart of the FFT algorithm
43 inevitably affects the efficiency of the FFT calculation, DAFT nevertheless
44 possesses excellent scaling characteristics and the associated economies in
45 data management resulting from its use makes the DL_POLY_3 SPME
46 implementation a highly efficient algorithm [36]. Radiation damage
47 simulations of order 2 million atoms (and larger) are regularly performed
48 with this program [37].
49
50
51
52
53
54
55
56

57 6. The Parallel SHAKE Algorithm

58
59 An essential requirement for all molecular dynamics codes intended for
60 modelling complex systems is a means of handling rigid interatomic bonds
(also called constraint bonds). Not only is this necessary to permit practical

1
2
3 (i.e. not too short) time step intervals, but also to remove the problem
4 *nonergodicity* (poor coupling of the system degrees of freedom) that delays, or
5 even prevents, the onset of equilibrium. Both versions of DL_POLY
6 implement the SHAKE algorithm for rigid bonds [21].
7
8

9
10 The principles of the SHAKE algorithm are well known. In the first stage the
11 motions of atoms are integrated without consideration of the rigid bonds. In
12 the second stage, a correction to the displacement of the atoms is applied to
13 restore the required bond length. The correction is applied to each bond in
14 turn and is applied iteratively, so that perturbations to each bond due to
15 corrections applied to neighbouring bonds may be allowed for. The iteration
16 ceases when all the bonds have converged, to within a tolerance of order $\sim 10^{-5}$
17 of the required length.
18
19

20
21 A parallel implementation of the SHAKE algorithm introduces new
22 considerations. In order to distribute the work load over a number of
23 processors it is sensible to allocate independent sets of constraint bonds to
24 different processors. Since atoms are generally linked into more than one
25 bond by virtue of their valency, this inevitable means that some bonded
26 atoms will be handled by more than one processor. It follows (see Figure 6)
27 that during the iterative stage of SHAKE, it is necessary to communicate
28 between processors to ensure that the corrections applied to individual atoms
29 take account of the full connectivity of the molecular structure. This problem
30 was first solved by the RD_SHAKE algorithm [38] which was intended for RD
31 implementation, but the techniques can be carried over to DD. In fact, the
32 communication overhead in DD is much less than in RD, since data
33 replication at intermediate stages is not required.
34
35
36
37

38 The RD implementation of SHAKE therefore requires firstly that the rigid
39 bonds for which the processor is responsible be identified and a list compiled
40 of the number of such bonds each atom participates in. This list is circulated
41 to all other processors, from which each processor may establish a 'shared
42 atom' list which records which atoms are shared with which processor. Then
43 during the SHAKE iteration, changes in the positions of shared atoms are
44 communicated to the appropriate processor, thus avoiding the problem
45 indicated in Figure 6. In the RD case the list of shared atoms needs to be
46 compiled only once, at the start of the simulation. For DD implementation
47 however, the shared atom list must be continually updated during the
48 simulation, as atoms move between processors. Fortunately, this is purely a
49 data transfer issue and does not involve a repeat of the reconstruction of the
50 complete shared atom arrays.
51
52
53
54

55 Similar to SHAKE, the RATTLE algorithm was also developed to treat
56 constraints, but in the velocity Verlet scheme. It has two parts: the first is an
57 iterative correction to the constrained atom positions as in SHAKE; and the
58 second is an additional iteration procedure to constrain atom velocities so that
59 the component of their relative velocity along the constraint bond is zero,
60

1
2
3 within a given tolerance. Both parts are similar in implementation to the RD
4 SHAKE scheme.
5
6

7. General Comments

7.1 Extension of the Code

13 As described above, DL_POLY_2 and DL_POLY_3 contain a common set of
14 functionality despite of their differences in strategies of parallelisation.
15 (Though we emphasise again that certain aspects of molecular topology e.g.
16 rigid bodies, present considerable difficulties in the DD implementation and
17 are therefore not currently available.) However, each of the packages also
18 contains unique features which were developed in response to demands from
19 users. Thus DL_POLY_3 supports defect detection tools and a variable
20 timestep algorithm, suitable for highly non-equilibrium simulations such as
21 radiation damage studies (with correspondingly large systems), which
22 require such features. More information about these can be found in the
23 respective manuals [39,40]. However, our vision for the future is to keep
24 these two packages as mutually compatible as possible.
25
26
27
28

29 The user community is encouraged to extend the current functionality of the
30 packages for their own benefit. For this purpose the software is supplied in
31 source form. Ideally such modifications would find their way back into the
32 standard versions of the programs, but they must comply with the high
33 standards of coding and documentation of the overall package. Unfortunately
34 we can make no commitment to verify, extend or develop extensions
35 contributed by users, but at our discretion we may develop certain
36 contributions if we judge they would benefit the wider DL_POLY
37 community.
38
39
40
41

7.2 Porting Issues

42 The DL_POLY packages are written in highly modularised FORTRAN 90 and
43 do not make use of any external libraries. Thus the packages are fully self
44 contained. Users can, of course, substitute the DL_POLY FFT routines with
45 local or vendor specific versions at compile time, though we do not
46 recommend this for DL_POLY_3 on account of the unique matching of the 3D
47 FFT to the DD force calculations.
48
49
50
51
52

53 The packages are supplied with template makefiles to handle assembly and
54 compilation in serial or parallel modes in either UNIX compatible or UNIX
55 emulated environments (e.g. Windows with CygWin [41]). Ideally users
56 would find these makefiles entirely sufficient, but a working knowledge of
57 compiler flags and optimisation issues is an advantage. While compilation in
58 serial mode is straight forward, parallel mode compilation raises additional
59 considerations. Inter-processor communications in DL_POLY are
60 implemented through FORTRAN MPI calls. For successful compilation and

1
2
3 flawless execution users must ensure that their communication hardware has
4 stable builds of the MPI software compiled with respect to the corresponding
5 FORTRAN 90 compiler.
6
7

8 7.3 The DL_POLY Java GUI 9

10 The DL_POLY suite also features a basic GUI for managing some aspects of
11 code use. Written in Java, and therefore highly portable, the GUI can help
12 with construction of input data, job submission and analysis of the simulation
13 output. Given the high degree of versatility of DL_POLY it has not been
14 possible to develop a GUI that satisfies every user need. For those willing to
15 experiment however, it can prove a valuable aid in exploiting DL_POLY. The
16 GUI is supplied with the DL_POLY source and is fully documented [42].
17
18
19

20 7.4 Obtaining DL_POLY 21

22 The DL_POLY programs are available free of charge (to academic researchers)
23 from the DL_POLY website [43]. The software is supplied under a licence
24 protecting the commercial rights of Daresbury Laboratory and the potential
25 user must agree to these terms before downloading the source.
26
27
28
29

30 8. Conclusion 31

32 The DL_POLY package provides a powerful and versatile set of programs for
33 molecular dynamics simulation of complex molecular systems. The potential
34 range of applications is vast, as this issue of Molecular Simulation
35 demonstrates. Systems as small as 100 atoms and as large as 30 million atoms
36 can be simulated with it. The code is free (to academics) and the source is
37 open to inspection, verification and extension. Scientists with an intention to
38 simulate large or complex systems should seriously consider what it has to
39 offer.
40
41
42
43
44

45 References 46

47 [1] For information on CCP5 see the project website at:
48 <http://www.ccp5.ac.uk>.

49
50

51 [2] W. Smith, "Molecular Dynamics on Distributed (MIMD) Parallel
52 Computers", *Theoretica. Chim. Acta.* **84** (1993) 385.
53
54

55 [3] W. Smith, "Molecular Dynamics on Hypercube Parallel Computers",
56 *Comput. Physics Comm.* **62** (1991) 229.
57
58

59 [4] A.R.C. Raine, D. Fincham and W. Smith, "Systolic Loop Methods for
60 Molecular Dynamics using Multiple Transputers", *Comp. Phys. Commun.* **55**
(1989) 13.

- 1
2
3
4
5 [5] M.R.S. Pinches, D. Tildesley and W. Smith, "Large Scale Molecular
6 Dynamics on Parallel Computers using the Link Cell Algorithm", *Molecular*
7 *Simulation*, **6** (1991) 51.
8
- 9
10 [6] D.C. Rapaport, "Multi-million particle molecular dynamics. II. Design
11 considerations for distributed processing", *Comput. Phys. Commun.* **62** (1991)
12 217.
13
- 14 [7] S.J. Weiner, P.A. Kollman, D.T. Nguyen and D.A. Case, "An All Atom
15 Force Field for Simulations of Proteins and Nucleic Acids", *J. Comp. Chem.* **7**
16 (1986) 230.
17
- 18 [8] W.F. van Gunsteren and H.J.C. Berendsen, "Groningen Molecular
19 Simulation (GROMOS) Library Manual", published by BIOMOS, Nijenborgh,
20 9747 Ag Groningen, The Netherlands,
21 (1987).
22
- 23 [9] A. D. MacKerell, Jr., B. Brooks, C. L. Brooks, III, L. Nilsson, B. Roux, Y.
24 Won, and M. Karplus, "CHARMM: The Energy Function and its
25 Parameterization with an Overview of the Program", in: *The Encyclopedia of*
26 *Computational Chemistry*, Vol. **1**, P. v. R. Schleyer, N. L. Allinger, T. Clark, J.
27 Gasteiger, P. A. Kollman, H. F. Schaefer
28
- 29 [10] B.G. Dick and A.W. Overhauser, "Theory of dielectric constants of alkali
30 halide crystals", *Phys. Rev. B* **112** (1958) 90.
31
- 32 [11] D. Fincham P.J. Mitchell, "Shell model simulations by adiabatic
33 dynamics", *J. Phys. Condens. Matter*, **5** (1993) 1031.
34
- 35 [12] P.J.D. Lindan and M.J. Gillan, "Shell-model molecular-dynamics
36 simulation of superionic conduction in CAF_2 ", *J. Phys. Condens. Matter*, **5**
37 (1993) 1019
38
- 39 [13] M.W. Finnis and J.E. Sinclair, "A simple empirical N body potential for
40 transition metals", *Philos. Mag. A*, **50** (1984) 45.
41
- 42 [14] J. Tersoff, "Modelling solid state chemistry: Interaction potentials for
43 multicomponent systems", *Phys. Rev. B*, **39** (1989) 5566.
44
- 45 [15] S.L. Mayo, B.D. Olafson and W.A.Goddard, "DREIDING: A generic force
46 field for molecular simulations", *J. Phys. Chem.*, **94** (1990) 8897.
47
- 48 [16] M.P. Allen and D.J. Tildesley, "Computer Simulation of Liquids",
49 Clarendon Press, Oxford (1989)
50
- 51 [17] D.J. Evans and G.P. Morriss, "Non-Newtonian molecular dynamics",
52 *Computer Physics Reports*, **1** (1984) 297.
53
54
55
56
57
58
59
60

- 1
2
3
4
5 [18] W.G. Hoover, "Canonical dynamics: Equilibrium phase-space
6 distributions", *Phys. Rev. A*, **31** (1985) 1695.
7
- 8 [19] S. Melchionna, G. Ciccotti and B. Holian, "Hoover NPT dynamics for
9 systems varying in shape and size", *Molecular Physics*,
10 **78** (1993) 533.
11
- 12 [20] H.J.C Berendsen, J.P.M. Postma, W.F. van Gunsteren, A. DiNola and J.R.
13 Haak, "Molecular dynamics with coupling to an external bath", *J. Chem.*
14 *Phys.*, **81** (1984) 3684).
15
- 16 [21] J.P. Ryckaert, G. Ciccotti. and H.J.C. Berendsen, "Numerical integration
17 of the Cartesian equations of motion of
18 a system with constraints: molecular dynamics of n-alkanes", *J. Comput.*
19 *Phys.*, **23** (1977) 327.
20
21
22
23
- 24 [22] H.C. Andersen, "Rattle: a velocity version of the SHAKE algorithm for
25 molecular dynamics calculations", *J. Comput. Phys.*, **52** (1983) 24.
26
27
- 28 [23] H. Goldstein, "Classical Mechanics", Addison Wesley (1980).
29
- 30 [24] D.J. Evans, "On the representation of orientation space",
31 *Molecular Physics*, **34** (1977) 317.
32
33
- 34 [25] D. Fincham, "Leapfrog rotational algorithms", *Molecular Simulation*, **8**
35 (1992) 165.
36
37
- 38 [26] T.F. Miller, M. Eleftheriou, P. Pattnaik, A. Ndirango, D. Newns, and G.J.
39 Martyna, "Symplectic quaternion scheme for biophysical molecular
40 dynamics", *J. Chem. Phys.*, **116** (2002) 8649.
41
42
- 43 [27] T.R. Forester and W. Smith, "Shake, Rattle and Roll: Efficient constraint
44 algorithms for linked rigid bodies", *J Computational Chemistry*, **19** (1998) 102.
45
46
- 47 [28] S. Brode and R. Ahlrichs, "An optimised MD program for a vector
48 computer Cyber 205", *Comput. Phys. Commun.*, **42** (1986) 41.
49
50
- 51 [29] R.W. Hockney and J.W. Eastwood, "Computer Simulation Using
52 Particles", McGraw-Hill International (1981).
53
54
- 55 [30] C. Kittel, *Solid State Physics*, John Wiley and Sons (1986).
56
- 57 [31] I.T. Todorov, W. Smith, K. Trachenko and M.T. Dove, "DL_POLY_3: new
58 dimensions in molecular dynamics simulations via massive parallelism", *J.*
59 *Mater. Chem.* **16** (2006) 1911.
60

1
2
3 [32] W. Smith, "A replicated data molecular dynamics strategy for the parallel
4 Ewald sum", *Comput. Phys. Commun.*, **67** (1992) 392.

5
6
7 [33] J. Hautman, J. and M.L. Klein, "An Ewald summation method for planar
8 surfaces and interfaces", *Molecular Physics*, **75** (1992) 379.

9
10
11 [34] U. Essmann, L. Perera, M.L. Berkowitz, T. Darden,
12 H. Lee, and L.G. Pedersen, "A smooth particle mesh Ewald method", *J.*
13 *Chem. Phys.*, **103** (1995) 8577.

14
15
16 [35] I.J. Bush, "The Daresbury Advanced Fourier Transform", Daresbury
17 Laboratory (1999).

18
19
20 [36] I.J. Bush, I.T. Todorov and W. Smith, "A DAFT DL_POLY Distributed
21 Memory Adaptation of the Smoothed Particle Mesh Ewald Method",
22 Accepted for publication *Comput. Phys. Commun.* (2006).

23
24
25 [37] K. Trachenko, M.T. Dove, E. Artacho, I.T. Todorov and W. Smith,
26 "Atomistic simulations of resistance to amorphization by radiation damage",
27 *Phys. Rev. B*, **73** (2006) 174207.

28
29
30 [38]] W. Smith and T.R. Forester, "Parallel Macromolecular Simulations and
31 the Replicated Data Strategy II: The RD-SHAKE Algorithm", *Comput. Phys.*
32 *Comm.* **79** (1993) 63.

33
34 [39] W. Smith, I.T. Todorov, T.R. Forester and M. Leslie, "The DL_POLY_2
35 User Manual", Daresbury Laboratory (2006) available from the DL_POLY
36 website[43].

37
38
39 [40] I.T. Todorov and W. Smith, "The DL_POLY_3 User Manual", Daresbury
40 Laboratory (2006) available from the DL_POLY website[43].

41
42
43 [41] The CygWin Linux API emulator is available from:
44 <http://sources.redhat.com/cygwin/>

45
46 [42] W. Smith, "The DL_POLY Java Graphical User Interface II", Daresbury
47 Laboratory (2003) available from the DL_POLY website[44].

48
49
50 [43] The DL_POLY website is located at: http://www.ccp5.ac.uk/DL_POLY/

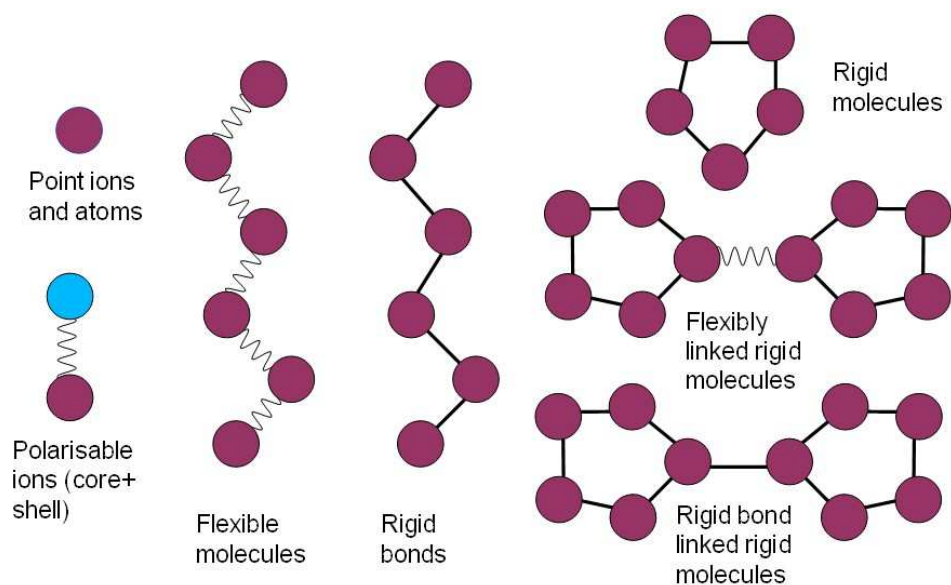


Figure 1

Molecular structures supported by DL_POLY. Any or all of such structures may be present in a given model at the same time. Rigid bodies however (right of diagram) are not available in DL_POLY_3.




Figure 2

Division of intramolecular force terms over processors. Here, bond angle terms A_1 to A_{16} are evenly allocated to specific processors for computation.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60




Figure 3

The parallel distribution of the Verlet neighbour list in DL_POLY_2 on a 4 processor machine. The pair force matrix is restructured according to the scheme of Brode and Ahlrichs and rows are assigned to processors P_0 to P_3 according to the colour scheme, to achieve a reasonable load-balance across processors.




Figure 4

Domain Decomposition. The MD cell (large cube) is divided into equal domains (middle cubes) each of which is allocated to a specific processor. Each domain is divided into link-cells (small cubes) the width of which must be greater than the radius of cut-off applied to the inter-atomic force terms. The sphere above represents the cut-off sphere defining the interaction range.




Figure 5

Halo data construction in Domain Decomposition. The central cube represents a spatial domain that is allocated to a single processor, where it is divided into link-cells (small cubes). Surrounding the domain it is necessary to add the halo data, which is one link-cell in width (indicated by the isolated small cube), as it is composed of the coordinates of atoms found in the link-cells at the boundaries of the neighbouring domains. It is apparent from this construction that the smaller the link cells, the more efficient will be the overall algorithm, since less data will need to be transferred.




Figure 6

The SHAKE algorithm on multiple processors. Atoms 1,2 and 3 represent part of a molecular structure. The bond 1-2 is handled by processor A and bond 2-3 is handled by processor B. Stage 1 of the shake algorithm will carry atom 1 to 1', 2 to 2' and 3 to 3'. Stage 2 iterates these positions until the final positions 1'',2'' and 3'' have conserved bond lengths. It is apparent however, that unless processors A and B communicate during iteration, atom 2 will relax to different positions on the two processors.