

Use of massively parallel molecular dynamics simulations for radiation damage in pyrochlores

Ilian T. Todorov · Neil L. Allan · J. A. Purton ·
Martin T. Dove · William Smith

Received: 4 January 2006 / Accepted: 1 March 2006 / Published online: 25 January 2007
© Springer Science+Business Media, LLC 2007

Abstract DL_POLY_3 is a general purpose molecular dynamics (MD) simulation package designed to simulate systems of the order of tens of millions of particles and beyond by efficiently harnessing the power of modern computer clusters. Here we discuss the package design, functionality and report on performance and capability limits. We then report the application of DL_POLY_3 to study radiation cascades in $\text{Gd}_2\text{Ti}_2\text{O}_7$ and $\text{Gd}_2\text{Zr}_2\text{O}_7$, potential materials for high-level radioactive waste storage and discuss problems associated with the analysis of the cascades. We see little direct amorphisation but rather the start of a transition to the fluorite structure which is more pronounced for the Zr than the Ti compound.

Introduction

Nowadays researchers use molecular dynamics (MD) as a standard theoretical tool to model the detailed atomistic behaviour of various types of systems over

time scales from picoseconds to microseconds as well as to examine equilibrium thermodynamic properties and free energies along reaction paths. In an MD simulation, the classical equations of motion governing the microscopic time evolution of a many body system are solved numerically. This information can be used to monitor the microscopic mechanisms of energy and mass transfer in chemical and physical processes, and dynamical properties such as absorption spectra, rate constants and transport properties can be calculated.

Another area of application, where MD is perceived as the chief theoretical tool available, is the simulation of strongly non-equilibrium processes. Of particular importance in this context is the simulation of radiation damage where the process is initiated by the creation of a 'primary knock-on' atom (PKA). The PKA dissipates its energy through interactions with 'target' atoms and when the PKA has sufficient kinetic energy (and mass) there is an atomic collision cascade [1]. The final damaged state results from the initial cascade (timescale of femtoseconds), nearly instantaneous relaxation of the structure (picoseconds) and the subsequent evolution of the system. Simulating such processes introduces particular problems associated with the size of the simulation (depending on the PKA energy), efficiency of parallelisation and simulation speed steering. In this paper we describe the new program DL_POLY_3 with particular emphasis on recently developed software features that make it particularly suitable for radiation damage studies and present a brief example of its use.

Our example is a set of two gadolinium pyrochlores of interest for radiation waste storage. $\text{Gd}_2\text{Ti}_2\text{O}_7$ is a primary candidate material, but studies employing heavy ion bombardment [2], simulating the results of

I. T. Todorov (✉) · J. A. Purton · W. Smith
Department of Computational Science and Engineering,
CCLRC, Daresbury Laboratory, Warrington WA4 4AD,
UK
e-mail: I.T.Todorov@dl.ac.uk

N. L. Allan
School of Chemistry, University of Bristol, Cantock's Close,
Bristol BS8 1TS, UK

I. T. Todorov · M. T. Dove
Department of Earth Sciences, University of Cambridge,
Downing Street, Cambridge CB2 3EQ, UK

α -decay, have shown this undergoes amorphisation which in turn leads to an increase in the leach rate of Pu by an order of magnitude. Also using heavy ion bombardment, Wang et al. [3] have demonstrated a striking variation in radiation tolerance in $\text{Gd}_2(\text{Zr}_x\text{Ti}_{1-x})_2\text{O}_7$ ($x = 0-1$) such that radiation tolerance increases with increasing Zr content and $\text{Gd}_2\text{Zr}_2\text{O}_7$ is predicted to withstand radiation damage for millions of years (cf. 100s of years for $\text{Gd}_2\text{Ti}_2\text{O}_7$). $\text{Gd}_2(\text{Zr}_x\text{Ti}_{1-x})_2\text{O}_7$ and other pyrochlores have now received considerable experimental attention [4–10]. The Gd zirconate end member ($x = 1$) is indeed largely insensitive to irradiation damage remaining highly crystalline to high doses even at very low temperatures undergoing a radiation-induced transition to a defect fluorite structure in which the Gd^{3+} and Zr^{4+} cations are disordered and which is itself highly radiation resistant. We use very large simulation cells in this paper and DL_POLY_3 to compare displacement cascades in $\text{Gd}_2\text{Ti}_2\text{O}_7$ and $\text{Gd}_2\text{Zr}_2\text{O}_7$. For a comprehensive recent review of pyrochlores in the context of nuclear waste disposal see Ewing et al. [11].

DL_POLY_3 Foundations

DL_POLY_3 [12] is a modern MD code conceived to harness the power of HPC clusters by embedding the domain decomposition strategy with linked cells methodology. It has recently been re-engineered in Fortran 90 taking full advantage of the modularisation concept to separate and distribute logically common (science-, maths- and semantics-wise) sets of variable declarations, methods and initialisations in modules. Adopting modularisation allows a lego-like build of further enhancements and new implementations such as force fields, scientific methodologies and numerical algorithms. The inter-CPU communication is implemented using MPI as most of the communication in the code is implicit, based on dedicated functions and subroutines developed as methods in a communication dedicated module.

Domain decomposition (DD) is widely used to parallelise various sorts of problems. The MD cell is divided equi-spatially into geometrically identical regions (domains) and each region is allocated to a node (a CPU). The mapping of the domains on the array of nodes is a non-trivial problem in general, although specific solutions for certain numbers of nodes (e.g., hypercubes) are much easier to obtain. The domains are divided into sub-cells (link-cells) with width slightly greater than the potential cut off for the system force-field. The coordinates of the atoms in

link-cells adjacent to the boundaries of each domain are passed onto neighbouring nodes sharing the same boundaries (*exchange of boundary data*) to create the so called *domain halo*. After this, each node may proceed to calculate all pair interactions in its region *independently* using the linked-cell method. No further communication between the nodes is necessary until after the equations of motion have been integrated.

If bond constraints are present in the system, the equations of motions are modified to include constraint solvers, RATTLE [13] for Velocity Verlet and SHAKE [14] for Leapfrog Verlet integrators [15]. These constraint solvers are iterative algorithms that add incremental corrections to the positions, velocities and forces of constrained particles until the bond lengths for all constraints in the system equal the corresponding pre-defined constraint bond lengths to within a given tolerance. Constraint algorithms involve extra communication at each iteration when constraint bonds cross domains. A domain crossing constraint bond exists on two domains and one of the constrained particles is on the domain halo of the other. Since constraint particle positions, velocities and forces change at each iteration it is necessary to refresh the constraint particles that lie in the halo by updating their positions, velocities and forces. Thus systems with constraints are bound to have lower parallelisation efficiency than systems without constraints. After the equations of motion have been integrated the particles which have moved out of their original domain must be reallocated to a new domain.

In the exchange of boundary data to build domain halos, it is crucial that data are only passed in successively complementary directions (in 3D: north-south and back, east-west and back and up-down and back) at any given instant and in between exchanges, the exchanged data must be re-sorted before the next exchange. This is necessary to ensure the corner and edge link-cell data are correctly exchanged between domains sharing edges and corners, rather than faces.

The linked-cell method (LC) [16, 17] is a simple serial algorithm linearly dependent on the number of particles comprising the domain and its surrounding halo. It assigns each atom to its appropriate sub-cell and a *linked list* is used to construct a logical chain identifying common cell members. A subsidiary *header* list identifies the first member of the chain. This allows all the atoms in a cell to be located quickly. The calculation of the forces is treated as a sum of interactions between sub-cells, in the course of which all pair forces are calculated. It is straightforward to allow for periodic boundary conditions. The algorithm greatly reduces the time spent in locating interacting

particles when the potential cut off is very short in relation to the size of the domain. This makes the LC algorithm particularly powerful as it enables extremely large systems to be simulated very cost-effectively.

Although the DD algorithm is originally designed for systems with short-range forces, it can also be used for systems with Coulombic forces. DL_POLY_3 relies on a new DD adaptation of the Smoothed Particle Mesh Ewald method (SPME) [18] for calculating long range forces in molecular simulations. In this adaptation [19] two strategies are employed to optimise the traditional Ewald sum. The first is to calculate the (short-ranged) *real space* contributions to the sum using the DD method as outlined above. The second is to use a fine-grained mesh in *reciprocal space* and replace the Gaussian charges by finite charges on mesh points. The mesh permits the use of 3D Fast Fourier Transforms (3D FFTs) [20]. DL_POLY_3 uses a novel, fully memory distributed, parallel implementation of the 3D FFT—the Daresbury Advanced Fourier Transform (DAFT) [21]—that exploits the DD concept.

DL_POLY_3 also offers a great number of control options [22]. However, it suffices here to outline only the few that are of particular use to radiation damage simulations. The *variable timestep* option requires the user to specify an initial guess for a reasonable timestep for the system (in picoseconds). The simulation is unlikely to retain this as the operational timestep as the latter may change in response to the dynamics of the system. The option is used in conjunction with two variables *mindis* (default 0.03 Å) and *maxdis* (default 0.1 Å; note *maxdis* ≥ 2.5 *mindis*). These distances serve as control values in the variable timestep algorithm, which calculates the greatest distance a particle has travelled in any timestep during the simulation. If the maximum distance is exceeded, the timestep variable is halved and the step repeated. If the greatest move is less than the minimum allowed, the timestep variable is doubled and the step repeated. In this way the integration timestep self-adjusts in response to the dynamics of the system; the simulation slows down to account accurately for the dynamics during the initial stages in the radiation damage cascade simulations far from equilibrium and then speeds up to utilise CPU time effectively when the system cools down.

The *defect detection* tool uses an algorithm that compares the simulated MD cell to a reference MD cell. The former defines the actual positions of the particles and their atom types and the latter is taken here to be the structure of the undamaged lattice. If a particle, p , is located in the vicinity of a site, s , defined by a sphere with centre this site and a user defined

radius, $0.3 \text{ \AA} \leq R_{\text{def}} \leq 1.3 \text{ \AA}$ (default value 0.75 Å), then the particle is a *first hand* claimer of s , and the site is not vacant. Otherwise the site is presumed vacant and the particle is presumed a general interstitial. If a site, s , is claimed while another particle, p' , is located within the sphere around it, then p' becomes an interstitial associated with s . After all particles and all sites are considered, it is clear which sites are vacancies. Finally, for every claimed site, distances between the site and its *first hand* claimer and interstitials are compared and the particle with the shortest one becomes the *real* claimer. If a *first hand* claimer of s is not the *real* claimer it becomes an interstitial associated with s . At this stage it is clear which particles are interstitials. The sum of interstitials and vacancies gives the total number of defects in the simulated MD cell. Note that the algorithm *cannot* be applied safely if R_{def} is larger than half the shortest interatomic distance within the reference MD cell since a particle may claim or/and be an interstitial associated with more than one site. Low values of R_{def} are likely to lead to slight overestimation of the number of defects. If the simulation and reference MD cell have the same number of atoms then the total number of interstitials is always equal to the total number of vacancies.

Performance and capability

To evaluate DL_POLY_3 performance and scalability a set of test MD simulations were run on the HPCx (IBM SP4 cluster—<http://www.hpcx.ac.uk>) super-cluster at Daresbury Laboratory (the UK's 1st and world's 46th fastest¹). The tests were based on three model systems; (i) Solid Ar, (ii) NaCl and (iii) SPC Water, with increasing complexity of their force-fields as outlined in Table 1 and carried out at conditions as outlined in Table 2. All test cases were set up using values for the simulation cell parameters obtained from previous equilibration runs using the same force-fields as listed in Table 1(a–c) and at the same simulation conditions as in Table 2. The size-per-CPU and cut off values for each system were chosen to ensure that all systems have the same domain halo volume on average, so that relatively the same volume of MPI messaging for domain boundary data exchange is required between neighbouring domains for each system at any timestep. Thus the difference in parallelisation performance between the three systems is

¹ Ranking in <http://www.top500.org> at the time of writing in January 2006.

Table 1 Force-fields for the tested systems as discussed in the text

(a) Solid Ar with mass of 39.95 Da and zero charge						
Interaction	Type	ϵ [eV]	σ [Å]			
Ar–Ar	Lennard–Jones	4.2	0.001			
(b) NaCl with $m(\text{Na}) = 22.9898$ Da, $q(\text{Na}) = 1e$, $m(\text{Cl}) = 35.453$ Da and $q(\text{Cl}) = -e$						
Interaction	Type	A [eV]	B [Å ⁻¹]	σ [Å]	C [eV Å ⁶]	D [eV Å ⁸]
Na–Na	Born–Huggins–Meyer	2544.35	3.1545	2.340	10117.0	4817.7
Na–Cl		2035.48	3.1545	2.755	67448.0	83708.0
Cl–Cl		1526.61	3.1545	3.170	698570.0	1403200.0
(c) SPC Water with $m(\text{H}) = 1.00797$ Da, $q(\text{H}) = 0.365e$, $m(\text{O}) = 16.9994$ Da and $q(\text{O}) = -0.73 e$, and constraints $\text{H}_{1,2}\text{--O} = 1.03$ Å and $\text{H}_1\text{--H}_2 = 1.681982$ Å						
Interaction	Type	ϵ [eV]	σ [Å]			
O–O	Lennard–Jones	0.16	3.196			

The Lennard–Jones and Born–Huggins–Meyer potentials have forms defined by

$$U(r_{ij}) = \left[\left(\frac{\sigma}{r_{ij}} \right)^{12} - \left(\frac{\sigma}{r_{ij}} \right)^6 \right] \text{ and } U(r_{ij}) = A \exp[B(\sigma - r_{ij})] - \frac{C}{r_{ij}^6} - \frac{D}{r_{ij}^8} \text{ respectively}$$

Table 2 Simulation parameters of the model systems tested using DL_POLY_3

System	Size per CPU [particles]	Ensemble [type]	Short-range potential cut off [Å]	Equilibrium temperature [K]	Equilibrium pressure [k atm]
Solid Ar	32,000	NVE	9	4.2	0.001
NaCl	27,000	NVE	12	500	0.001
SPC water	20,736	NPT Berendsen 0.5 0.75	8	300	0.001

NPT ensembles are characterised by thermostat and barostat relaxation times in picoseconds

based on the complexity of the different force fields and on the additional communication these involve.

To detect and compare the parallelisation efficiency of different systems and of different processor counts the following construction was employed. Whenever the number of CPUs was doubled the simulated systems were also doubled in size,² ensuring that the link-cell algorithms and the domain halo volume for each system remained the same for any processor count. Thus, if parallelism were ideal the simulation time-per-timestep for each system would be the same for any processor count.

Table 3(a–c) present simulation performance data for systems (i)–(iii). The tables list the time-per-timestep and system size as a function of processor count. Also listed for comparison are the dimensionless time-per-timestep, τ , and the speed gain, γ , defined by

$$\tau(n) = \frac{\Delta t_n}{\Delta t_1}, \quad \gamma(n) = \frac{\Delta t_1}{\Delta t_n} \tag{1}$$

where Δt_n is the time-per-time step at processor count n . τ and γ for all three systems are plotted as a function of the processor count in Fig. 1a, b respectively. Perfect parallelisation corresponds to $\tau(n) = 1$, good parallelisation corresponds to 25% increase at each doubling of the processor count. Figure 1a shows that for all three systems τ increases logarithmically until 32 CPUs and then almost linearly. The logarithmic increase in τ at low processor counts is associated with the population of the first logical partition (LPAR) of HPCx. HPCx comprises 38 LPARs, each running its own copy of the AIX operating system, and connected via high performance ‘federation’ switches. The linear increase of τ at high processor counts is related to the incremental increase of (i) time for MPI organisation and (ii) global (collective) communication operations. The results show that parallelisation is excellent for all three systems.

The SPME summation accuracy (10^{-6}) was kept constant³ for all NaCl and SPC water simulations

² System sizes were doubled cyclically in the a , then b and then c directions.

³ This corresponded in a $64 \times 64 \times 64$ grid for the FFT per domain (CPU).

Table 3 DL_POLY_3.04 scaling performance on HPCx

CPUs	Speed gain	Dimensionless time-per-timestep	Time-per-timestep [s]	System size [atoms]
(a) Solid Ar				
1	1	1.00	0.59	32,000
2	1.97	1.02	0.60	64,000
4	3.87	1.03	0.61	128,000
8	7.96	1.00	0.59	256,000
16	15.69	1.02	0.60	512,000
32	29.96	1.07	0.63	1,024,000
64	59.14	1.08	0.64	2,048,000
128	116.50	1.10	0.65	4,096,000
256	231.86	1.10	0.65	8,192,000
512	448.44	1.14	0.67	16,384,000
1024	838.36	1.22	0.72	32,768,000
(b) Solid NaCl				
1	1	1.00	2.82	27,000
2	1.95	1.02	2.89	54,000
4	3.80	1.05	2.97	108,000
8	7.69	1.04	2.93	216,000
16	14.96	1.07	3.01	432,000
32	27.79	1.15	3.24	864,000
64	54.07	1.18	3.34	1,728,000
128	105.12	1.22	3.43	3,456,000
256	206.24	1.24	3.50	6,912,000
512	388.41	1.32	3.71	13,824,000
1024	705.81	1.45	4.09	27,648,000
(c) SPC water				
1		1.00	1.90	20,736
2	1.81	1.11	2.10	41,472
4	3.36	1.19	2.26	82,944
8	7.11	1.13	2.14	165,888
16	13.54	1.18	2.24	331,776
32	23.53	1.36	2.58	663,552
64	43.80	1.46	2.77	1,327,104
128	84.24	1.52	2.89	2,654,208
256	158.97	1.61	3.06	5,308,416
512	285.41	1.79	3.41	10,616,832
1,024	495.35	2.07	3.93	21,233,664

Speed gain, dimensionless time-per-timestep, real time-per-timestep (in seconds, averaged over 10 timesteps) and system size are listed as a function of number of CPUs used in parallel. Details of individual simulations (a)–(c) are given in the text and Table 1

together with the constraint tolerance (relative, 10^{-5}) for all SPC water simulations. This imposes some small extra memory overheads as well as more expensive SPME electrostatics due to 3D FFT calculations which scale as $N \log N$ with the number of particles N . This is shown well in Fig. 1b, which plots the speed gain, $\gamma(n)$, as a function of processor count, n , where the performance for NaCl and SPC water is poorer than that for solid Ar. As expected, the SPC water system exhibits the worst parallelisation performance. This is due to the existence of constraint bonds involving (i) iterative integration algorithms [13, 14] and (ii) extra MPI communications for constraint bonds crossing domain boundaries, which, as we have already commented, need to be refreshed during each iteration cycle of the constraint algorithms.

Maximum load tests involving the three test systems were also run on HPCx. Each system size load on one CPU was increased incrementally until execution

failed. The maximum load for the solid Ar system was $\approx 700,000$ particles per 1 GB. This on 1,024 CPUs on HPCx corresponds to ≈ 610 million particles, which is well below the limit (2,147,483,647) DL_POLY_3 can handle on conventional 32-bit machines. Even this limit, of course, can be exceeded with care, but this requires the package is compiled in a 64-bit mode and run on a 64-bit (enabled) platform (with theoretical limit of system sizes of 9.2×10^{18} particles) and substantial hard disk space is needed. For a one million particle system the size of a configuration file (lattice parameters, positions, velocity and forces) in text format is ≈ 0.25 GB.

The maximum loads per CPU for the NaCl and SPC water systems were $\approx 220,000$ and $\approx 210,000$ ions per 1 GB memory respectively. These cannot be scaled directly to 1,024 CPUs since some extra memory per CPU will be needed for the 3D FFTs in the SPME summations driven by the constant SPME precision.




Fig. 1 DL_POLY_3 (a) dimensionless time-per-timestep and (b) speed gain from the simulation data in Table 1(a–c) plotted as a function of processor count. The green dotted line indicates the parallelisation limit also called perfect or embarrassing parallelisation. The red line indicates the standard for a good parallelisation (time-per-timestep increase of 25% for each doubling of the processor count)

However, our estimate is a 1,000 fold load on 1,024 CPUs. The reduced maximum load per CPU for the latter two systems relative to that in the solid Ar simulation is readily explained. The maximum load per CPU is dependent on the available memory per CPU after force field arrays are allocated. The larger the complexity and/or the higher the accuracy of the force field description, the lower the limit of the maximum load per CPU.

Overall, DL_POLY_3 exhibits excellent parallelisation performance over large processor counts and ability to utilise available memory extremely efficiently allowing loads as large as $\approx 220,000$ ions per 1 GB memory for systems with force-fields as complex as that for SPC water.

Cascade simulations

In this section we report molecular dynamics simulations of collision cascades for PKAs with energies $E \leq 3$ keV, $3 \text{ keV} < E < 10$ keV and $E = 10$ keV using cubic simulation cells of 88,000, 152,064 and 193,336 atoms respectively. These large cell sizes ensure that a cascade should not normally interact with its periodic images for all practical purposes in our simulations. Initial simulation cell lengths were set to the appropriate multiple of the experimental lattice parameter, corresponding to 10, 12 and 13 unit cells in each direction respectively. The initial temperature of the simulation was set to 300 K and the simulation allowed to equilibrate for 10 ps within the NPT-ensemble, using a timestep of 2 fs.

Each ion in the pyrochlores is assigned its formal charge, i.e. 3+, 4+ and 2– for Gd, Ti/Zr and O respectively. Ions can approach each other closely particularly at the start of the simulations and it is vital to use potentials that are likely to be accurate over a wide range of internuclear separations and especially at short distances. The potentials here have been calculated using the modified Gordon–Kim electron-gas model [23]. They have been used previously in simulation studies of a wide range of binary and ternary oxides (e.g., Ref. [24]) including problems involving defects where the interatomic distances close to the defect after relaxation may be very different from those in the perfect lattice at equilibrium [25]. The potential parameters for O–O, Ti–O, Gd–O and Zr–O, collected together in Table 4, were generated by fitting an exponential of the form $A \exp(-r/\rho)$ (where A and ρ are constants) to the electron-gas interaction energies. A comparison of experimental and calculated lattice parameters and bond lengths for the two pyrochlores is given in Ref. [26]. Devanathan and Weber have recently used potentials for the Ti and Zr compounds very similar to our own [27].

The production of a cascade was started from the energetic recoil of a single atom, the PKA. All simulations used for the PKA a U atom substituted

Table 4 Short range potential parameters used for the pyrochlores and the primary knock-on' atom (PKA) with cut off of 8 Å

Interaction	A/eV	$\rho/\text{Å}$
O–O	249.3764	0.3621
Ti–O	3878.4225	0.2717
Gd–O	4651.9633	0.2923
Zr–O	8769.5930	0.2619
U–O	4792.4400	0.3009

The form of the Buckingham potential is $V_{ij} = A \exp(-r_{ij}/\rho)$

for a native Gd ion, with the U–O potential also given in Table 4. The subsequent collisions, displacements and recombination of atoms with vacant sites were followed, so that the complete structure and evolution of the cascade could be monitored. A 5 ps evolution of a 5 keV PKA cascade in an MD cell of 152,064 particles takes 3.5 h of CPU time on 16 processors of the Dirac cluster at the University of Bristol (48 dual processor Pentium 4 Xeon on 2 GHz nodes with Myranet 2000 interconnect). A 10 ps evolution of a 10 keV PKA cascade for the largest MD cell used in this study takes 9 h of CPU time on 16 processors on the same cluster. Here in all simulations the initial velocity of the PKA was along a body-diagonal of the MD cell (eight directions). The creation of a radiation cascade results in an initial sharp rise in the global temperature of the simulation cell, no larger than 1,150 K for the simulation cells in this paper, and the temperature then decreases during the simulation. Timestep lengths vary during the development of a cascade. Typically for the 10 keV cascades reported in this paper, timesteps are very short (approximately a few 10^{-3} fs) at the beginning of the cascade simulation (up to 0.25 ps of the simulation time) and remain relatively unchanged from after a few picoseconds ($1 \div 5$ ps of simulation time) to the end of the simulation (total simulation time 8–12 ps).

Analysis of the cascades can be carried out in several useful ways. We start as in Ref. [28] by defining a ‘defect’ in terms of displacements, i.e., as an ion that has moved more than half the average of the first-neighbour interatomic distances. This is determined with respect to the position of an ion after the equilibration phase of the simulation. Alternative definitions [29] are discussed below. Calculated limiting values of the mean-square displacements for the parent pyrochlores (no cascades) indicate that self-diffusion processes are negligible after the final formation of the cascade and so do not influence the number of defects on this relatively short timescale.

For a given PKA energy the number of such ‘defects’ is larger for $\text{Gd}_2\text{Zr}_2\text{O}_7$ than $\text{Gd}_2\text{Ti}_2\text{O}_7$, as is clear in Fig. 2, which plots the running totals of ‘defects’ as a function of time during a 10 keV cascade. In Table 5 we break down this total for the 10 keV cascades. The numbers of each type of atom which have been displaced by a given distance are listed, together with the total distance travelled by the PKA. More atoms are displaced in the Zr than the Ti compound. The PKA has travelled further in the Zr case, and fewer heavier Zr ions are displaced than lighter Ti ions. Nevertheless, the number of cations




Fig. 2 Instantaneous number of ‘defects’ plotted as a function of time during a 10 keV cascade (U PKA) in the Ti and Zr pyrochlores (Gd–Ti–O triangles, Gd–Zr–O squares). There are separate plots for the damage on the cation (empty symbols) and anion (filled symbols) sublattices. As discussed in the text we here define a ‘defect’ as an ion that has moved more than half the average of the first-neighbour interatomic distances

Table 5 Structural damage data from 10 keV cascades in Ti and Zr pyrochlores at 300 K and zero pressure

PKA directions $\langle 111 \rangle$	$\text{Gd}_2\text{Ti}_2\text{O}_7$			$\text{Gd}_2\text{Zr}_2\text{O}_7$		
	Species			Species		
Displaced:	Gd	Ti	O	Gd	Zr	O
Band (\AA)						
<1	0	0	0	0	0	0
1–2	0	1	6	0.3	0	153
2–3	1	2	212	3	2	232
3–4	8	8	49	6	6	132
4–5	0.5	1	30	3	2	51
5–6	0	0.5	13	0.8	1	18
6–7	0.3	0.8	4	0.8	1	8
7–8	0	0.8	4	0	0	2
8–9	0.5	0	3	0.5	0.3	2
9–10	0	0	1	0	0	1
...
Total	11	15	329	15	12	610
Final distance of PKA from starting position (\AA)	84			95		

The primary knock-on’ atom (PKA) is a U ion. The table lists how many atoms of each species have moved a particular distance given in the first column from their original lattice positions. The second column and the third column list the mean values from all eight cascades in $\langle 111 \rangle$ directions for Ti and Zr pyrochlores respectively

displaced overall is comparable in both compounds, while more anions are displaced in the Zr pyrochlore. This appears to be in conflict with results from high-energy ion bombardment experiments [3], where the resistance to amorphization in $\text{Gd}_2\text{Ti}_{2-x}\text{Zr}_x\text{O}_7$ increases with increasing Zr content. We have

therefore examined the damage distribution in more detail. We note first that the increase in damage on the oxygen sublattice is largely due to a substantial increase in the Zr compound of the numbers of O ions moving between 1 Å and 2 Å (by over 150) and between 3 Å and 4 Å, relative to those in the Ti compound. The numbers of oxygen atoms moving between 2 Å and 3 Å or more than 4 Å from their initial positions are broadly similar. We obtain further insight into the oxygen disorder after each cascade by plotting the number of ions displaced vs. distance rather than tabulating this information in broad bands as in Table 5. This plot is shown in Fig. 3 for the two materials, and it is clear that this plot is highly structured with well defined peaks up to ≈ 6 Å for each compound. The differences in the number of atoms displaced at short distances (<2.7 Å) between the Zr and Ti compounds are particularly striking in this figure, with pronounced peaks at just below and just above 2 Å for the former which are almost completely absent in the curve for the latter. Many oxygens are thus displaced a small amount in the Zr compound. The differences in the two curves for larger displacement distances (>3.5 Å) are smaller and the sign of the difference between the two curves changes with distance.

It is particularly instructive to consider not just the displacements of individual ions but to consider the structure of the damaged compound itself. We examine the number of atoms not located at lattice sites, using the defect detection algorithm described in the previous section, thus changing our operating definition of a defect and avoiding the previous somewhat arbitrary definition of a defect in terms of displacements. Table 6 lists the number of defects, both interstitials

Table 6 Averaged final numbers of interstitials and vacancies by atom type and antisite defects in the simulation cell associated with cascades in each pyrochlore $Gd_2X_2O_7$ ($X = Ti, Zr$) for a primary knock-on' atom (PKA) of 10 keV with initial temperature 300 K

Compound	Defects							Gd-X
	Gd		X		O			
	<i>int</i>	<i>vac</i>	<i>int</i>	<i>vac</i>	<i>int</i>	<i>vac</i>	<i>int'</i>	
$X = Ti$	2.3	2.8	4.0	3.0	16.0	16.5	15.8	8.3
$X = Zr$	4.3	5.0	3.3	2.8	244.0	243.8	74.8	12.0

The final column for oxygen interstitials (headed *int'*) is the total number recalculated for the three compounds classifying the 8a site as a lattice rather than interstitial site, as described in the text

and vacancies, defined in this new way, in the simulation cell (193,336 atoms) for a PKA of 10 keV for each compound. The variation in the number of interstitial defects from one atom type to another and from the Ti to the Zr pyrochlore are the same as for the totals of each atom type displaced in each compound (e.g., Table 5).

A measure based on a *point* defect description can be misleading in certain situations. For example in the presence of dislocation-type *extended* defects every atom associated with the extended defect will count towards the defect total even in a region where there is significant *local* order. This consideration has led us to calculate the oxygen–oxygen radial distribution functions (RDFs) for the undamaged and damaged materials. These and the differences between the two are plotted in Fig. 4a–b. Although some oxygen atoms move by a small amount (Fig. 3) this shows that, Fig. 4 does not mean that there are any such short oxygen–oxygen separations in the damaged materials. Overall, as indicated by the sharp negative peaks in Fig. 4, there is a loss of order and there are clear differences between the Ti and Zr compounds. One way of investigating this a little more quantitatively is to sum the absolute values of the areas under the difference curves in Fig. 4a–b for each compound. In Fig. 5 the integral of the absolute difference between the oxygen–oxygen RDFs of the undamaged and damaged material is plotted as a function of distance. There is a larger loss of order in the Zr pyrochlore at shorter separations (3–5 Å), whereas at larger distances there appears to be a slightly larger loss of order for the Ti compound. The total integral is approximately 2% larger for the Ti than for the Zr compound. This is of course a small difference but nevertheless this measure, unlike those based on oxygen *displacements* considered earlier, at least suggests that overall the loss of oxygen order is at least comparable in the Ti to




Fig. 3 Mean final number of oxide ions displaced as a function of distance moved for all 10 keV cascades (U PKA) in the Ti and Zr pyrochlores (Gd–Ti–O triangles, Gd–Zr–O squares)




Fig. 4 Calculated oxygen–oxygen radial distribution functions (RDF) for the two undamaged (thick solid line, black) and damaged (dashed line, red) (a) Ti and (b) Zr pyrochlores. In each plot the difference between the undamaged and damaged RDFs is also plotted (thin solid line, green)

that in the Zr compound, due to the nature of the RDFs at larger rather than shorter separations.

The pyrochlore $\text{Gd}_2\text{Zr}_2\text{O}_7$ shows enhanced conductivity relative to that of the defect fluorite phases, which bound it on both sides in the phase diagram [30]. The pyrochlore structure not only contains A and B ions on separate sublattices but the seven oxygen ions per formula unit in the unit cell also fully occupy two sublattices ($48f$ and $8b$). Unlike the defective fluorite structure there is an ordered arrangement of oxygen vacancies, with the so-called $8a$ site vacant. The ionic conductivity is due to intrinsic disorder of oxygen ions between the occupied $48f$ and empty $8a$ sites [30]. The loss of order at short distances (3–5 Å) is associated with the partial occupation of this $8a$ site. Classifying this position as an additional available lattice rather interstitial site, reduces very substantially the number




Fig. 5 The integrated absolute difference, I , between the oxygen–oxygen radial distribution functions (RDFs) of the undamaged and damaged compounds as a function of distance. $\text{Gd}_2\text{Ti}_2\text{O}_7$ red (dashed) line, $\text{Gd}_2\text{Zr}_2\text{O}_7$ black (thick solid) line

of oxygen defects in $\text{Gd}_2\text{Zr}_2\text{O}_7$ by a factor of 3 (Table 6), while the total in the Ti compound is virtually unchanged. For $\text{Gd}_2\text{Ti}_2\text{O}_7$ intrinsic oxygen disorder and hence ionic conductivity is much less [31].

Thus in both $\text{Gd}_2\text{Ti}_2\text{O}_7$ and $\text{Gd}_2\text{Zr}_2\text{O}_7$, the total number of oxygen atoms displaced is much larger than that of the cations (as also noted [32] in simulations of cascades in $\text{La}_2\text{Zr}_2\text{O}_7$) consistent with a much lower threshold displacement energy for O. With the PKAs considered here, the final total of displaced cations is low. Most of the displaced cations occupy an equivalent crystallographic site. Only very few cations (at most one in each cell) move to occupy an oxygen site. Cation displacement threshold energies in $\text{Gd}_2\text{Ti}_2\text{O}_7$ appear to be significantly larger than in $\text{Gd}_2\text{Zr}_2\text{O}_7$, in agreement with Ref. [27]. As the RDFs show, for the cascades we have considered here both long and short range order are largely preserved, as also observed [32] for $\text{La}_2\text{Zr}_2\text{O}_7$. Although many oxygen atoms are displaced from their equilibrium positions most end up occupying equivalent crystallographic sites. Anion disordering involving the greater occupancy of the $8a$ sites is greater in the Zr compound as shown by the larger occupancy of this $8a$ site; the oxygen vacancies are disordered over the oxygen sites in addition to the cation disorder. Figure 5 hints at a possible larger loss of longer range oxygen order in the Ti compound, although the number of oxygens displaced is much smaller. The characteristics of the cation and anion disorder differ from compound to compound, with different timescales, and largely need to be considered independently. Overall, the primary damage by our

low energy events shows the beginnings of the formation of the fluorite-like phase, rather than amorphisation per se, though this is less for the Ti compound.

Sickafus et al. [33] have proposed that oxygen-deficient fluorites, with the same basic $A_2B_2O_7$ structure as pyrochlores but a random arrangement of the cations, have a greater propensity for resisting radiation damage (see also Ref. [34]). The lower the cation antisite defect energy (e.g., $Gd_2Zr_2O_7$) the more readily will a transition under irradiation to the disordered fluorite structure take place and the more radiation resistant the compound. Alternatively, if the antisite defect is large (e.g., $Gd_2Ti_2O_7$) then there is an irradiation-induced crystalline to amorphous transition. The cation antisite defect energies are smaller the closer these ions are in size [35, 36].

How do such trends relate to the simulations reported here? Overall our calculations suggest the connection between ease of damage and the formation of defect cascades over the first few ps is more complex than anticipated hitherto. Different measures of damage (displacements, defect totals, RDFs) need careful consideration. The simulations see no direct amorphisation but rather a transition to the fluorite structure which is more pronounced for the Zr compound (greater occupancy of the 8a site) than the Ti system, which also contains a smaller number of antisite defects. The total number of displaced atoms is larger for the Zr pyrochlore, but the loss of long range order as indicated by the change in the oxygen–oxygen radial distribution indicates this is comparable or less for the Zr than the Ti pyrochlore.

There are a number of caveats. Our simulations have been limited to small initial PKA energies because of the problem of cascade overlap. The creation of a defect cascade is only part of the process of the formation of the damaged state. Our simulations are not able to probe the healing process and determine the ultimate fate on longer timescales of the interstitial ions and vacancies (e.g., formation of defect clusters, dislocations). Nevertheless, we might expect thermal annealing via order-disorder in $Gd_2Zr_2O_7$ (which is an ionic conductor) to be faster than for structural recovery in $Gd_2Ti_2O_7$. It will be interesting in future work to try and build up correlations between ease of damage formation and ease of amorphisation with measures calculated in molecular dynamics simulations of the type presented here, and the ability to study ever larger cells will permit the study of more energetic PKAs.

Conclusions

DL_POLY_3 is a new generation molecular dynamics software with inherent parallelisation design based on domain decomposition and linked cells methodologies. It exhibits excellent parallel performance and ability to handle simulations of systems of order of tens of millions of particles and beyond on any high processor counts. DL_POLY_3 offers a wide variety of tools and controls, including additional functionality to help highly non-equilibrium simulations. Re-engineered in a modular, free-format FORTRAN 90 manner, it guarantees full portability (completely self-contained) and allows for easier support and user development. Application of the code to study radiation cascades in $Gd_2Ti_2O_7$ and $Gd_2Zr_2O_7$, has revealed marked differences between the two compounds, and we have discussed different methods for the analysis of the extensive disorder produced by these cascades.

Acknowledgements ITT acknowledges support from NERC. The work was also supported by two JREI awards (the Mott (RAL) and Dirac (Bristol) supercomputer clusters).

References

1. Weber WJ, Ewing RC, Catlow CRA, Diaz de la Rubia T, Hobbs LW, Kinoshita C, Matzke HJ, Motta AT, Nastasi M, Salje EKH, Vance ER, Zinkle SJ (1998) *J Mater Res* 13:1434
2. Wang SX, Wang LM, Ewing RC, Govindan Kutty KV (2000) *Nucl Instrum Methods Phys Res B* 169:135
3. Wang SX, Begg BD, Wang LM, Ewing RC, Weber WJ, Govindan Kutty KV (1999) *J Mater Res* 14:4470
4. Gupta HC, Brown S, Rani N, Gohel VB (2001) *J Raman Spect* 32:41
5. Hess NJ, Begg BD, Conradson SD, McCready DE, Gassman PL, Weber WJ (2002) *J Phys B* 106:4663
6. Lian J, Wang L, Chen J, Sun K, Ewing RC, Matt Farmer J, Boatner LA (2003) *Acta Materialia* 51:1493
7. Lian J, Chen J, Wang LM, Ewing RC, Matt Farmer J, Boatner LA, Helean KB (2003) *Phys Rev B* 68:134107
8. Nachimuthu P, Thevuthasan S, Engelhard MH, Weber WJ, Shuh DK, Hamdan NM, Mun BS, Adams EM, McCready DE, Shutthanandan V, Lindle DW, Balakrishnan G, Paul DM, Gullikson EM, Perera RCC, Lian J, Wang LM, Ewing RC (2004) *Phys Rev B* 70:100101
9. Nachimuthu P, Thevuthasan S, Shutthanandan S, Adams EM, Weber WJ, Begg BD, Shuh DK, Lindle DW, Gullikson EM, Perera RCC (2005) *J Appl Phys* 97:033518
10. Nachimuthu P, Thevuthasan S, Adams EM, Weber WJ, Begg BD, Mun BS, Shuh DK, Lindle DW, Gullikson EM, Perera RCC (2005) *J Phys Chem B* 109:1337
11. Ewing RC, Weber WJ, Lian J (2004) *J Appl Phys* 95:5949
12. Todorov IT, Smith W (2004) *Phil Trans R Soc Lond A* 362:1835

13. Andersen HC (1983) *J Comput Phys* 52:24
14. Ryckaert JP, Ciccotti G, Berendsen HJC (1977) *J Comp Phys* 23:327
15. Allen MP, Tildesley DJ (2002) In: *Computer simulation of liquids*. Clarendon Press, Oxford
16. Pinches MRS, Tildesley D, Smith W (1991) *Mol Simulation* 6:51
17. Hockney RW, Eastwood JW (1981) In: *Comp simulation using particles*. McGraw-Hill, New York
18. Essmann U, Perera L, Berkowitz ML, Darden T, Lee H, Pedersen LG (1995) *J Chem Phys* 103:8577
19. Bush IJ, Todorov IT, Smith W (2006) *Comput Phys Commun* 175:323
20. Brigham EO (1988) In: *The fast Fourier transform and its applications*. Prentice Hall, Singapore
21. Bush IJ (1999) In: *The Daresbury advanced Fourier transform*. Daresbury Laboratory
22. Todorov IT, Smith W (2005) In: *The DL_POLY_3 user manual, version 3.05*, http://www.cse.clrc.ac.uk/msi/software/DL_POLY/
23. Allan NL, Mackrodt WC (1994) *Phil Mag B* 69:871
24. Allan NL, Mackrodt WC (1993) *Adv Solid State Chem* 3:221
25. Allan NL, Mackrodt WC (1994) *Mol Simul* 12:89
26. Purton JA, Allan NL (2002) *J Mat Chem* 12:2923
27. Devanathan R, Weber WJ (2005) *J Appl Phys* 98:0861110
28. Deng HF, Bacon DJ (1996) *Phys Rev B* 53:11376
29. Todorov IT, Allan NL, Purton JA, Dove MT (2006) *J Phys Condens Matter* 18:2217
30. van Dijk MP, de Vries KJ, Burggraaf AJ (1983) *Solid State Ionics* 9/10:913
31. Tuller HL (1994) *J Phys Chem Solids* 55:1393
32. Chartier A, Meis C, Crocombette J-P, René Corrales L, Weber WJ (2003) *Phys Rev B* 67:174102
33. Sickafus KE, Minervini L, Grimes RW, Valdez JA, Ishimaru M, Li F, McClellan KJ, Hartmann T (2000) *Science* 289:748
34. Sickafus KE, Minervini L, Grimes RW, Valdez JA, Hartmann T (2001) *Rad Effects Def Sol* 155:133
35. Minervini L, Grimes RW, Sickafus KE (2000) *J Am Ceram Soc* 83:1873
36. Minervini L, Grimes RW, Tabira Y, Withers RL, Sickafus KE (2002) *Phil Mag A* 82:123