

MULTI-MILLION ATOM MOLECULAR DYNAMICS SIMULATIONS OF SHOCKED MATERIALS

William D. Mattson and Betsy M. Rice
U. S. Army Research Laboratory, AMSRD-ARL-WM-BD
Aberdeen Proving Ground, MD 21005-5069

ABSTRACT

We describe a new extensible software system that performs multi-million atom molecular dynamics simulations. This new software system is based on a program environment and a set of components stored in shared object libraries that are executed by the program environment as specified by an XML description of the simulation. Fixed and scaled speedups for this software are demonstrated in molecular dynamics simulations of α -quartz, using up to 256 processors and up to 4.4 million atoms. Scaled speed up is over 80% of optimal on 256 processors and fixed speed up is also over 80% on 64 processors for system sizes as small as 1,000 atoms per processor. The outstanding scaling for both small and large systems is somewhat unusual, since other multi-million atom simulation codes tend to perform well only for very large systems.

1. INTRODUCTION

Materials of interest to Future Force weapons (e.g. energetic materials, advanced armor) are destined to be subjected to extreme impacts leading to substantial physical or chemical changes in the atomic structure of the material. Molecular Dynamics (MD), an atomistic simulation method that provides a temporal description of the behavior of a system subjected to initiating events, is a particularly effective means to examine complex changes occurring in condensed phase materials after shock, including reaction and failure waves, phase transitions, and defect formation and propagation. Such a detailed examination cannot be performed experimentally; however, the detailed information obtained from MD simulations can easily be used to either design experiments that will transform the material to a desired state upon shock, or to design new materials with tailored shock properties.

With this in mind, we have established a computational framework that will allow easy integration of evolving software required to produce molecular dynamics simulations of various materials of interest to the Army under various conditions, with rapid turnaround. We have implemented into this framework a variety of methods to simulate shocked materials under conditions of constant pressure or constant volume, and for a variety of different materials. Commercial molecular simulation codes for reactive systems are

limited in material prediction and are designed to treat systems no larger than 10,000 atoms. Such a small system is not sufficient to study shock, reaction, or failure wave phenomena, which requires a sufficiently large simulation cell through which a wave can initiate and reach a steady state. Additionally, when studying effects of defects or grain boundaries, a significantly larger number of atoms is required in the simulation. Finally, realistic interaction potentials used in the MD simulations are often extremely complex and computationally expensive, thus requiring sophisticated dynamics algorithms to minimize computational cost associated with large-scale simulations and costly interaction potentials. Toward this end, we have developed and introduced several new methods for efficiently accommodating the extremely large systems and dramatically speed up the calculations of the interaction potential.

As in most large-scale MD codes, we utilize a domain decomposition scheme, in which the physical space of the simulation (and atoms contained therein) is distributed among processors. Thus, each processor handles the dynamics corresponding to a subset of atoms in the system. As the atoms move during the simulation they migrate from the physical space contained on one processor to that on another, frequently oscillating about this boundary. For the case of shocked materials, this introduces substantial communication overhead. Thus, we have introduced a just-in-time redistribution of atoms across processors to assist in minimizing the communication required for large systems. This just-in-time method makes these boundaries fuzzy to allow atoms to stay on the same processor much longer, minimizing the overhead involved in the communication and tracking of atomic motion. Additionally we have adopted the mid-point method of Bowers from particle communication (Bowers et. al., 2006), which can substantially decrease the required communication at each time step, and we have implemented the high resolution cell algorithm from our previous work (Mattson and Rice, 1999). To increase the efficiency of the complex interaction potentials we have implemented a new multi-resolution sorted pair/triplet list giving significantly better performance for these complex potentials than previously obtained (Yao et. al. 2004).

To demonstrate the capability of this code, we present multi-million atom simulations of a shock wave

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 01 NOV 2006		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Multi-Million Atom Molecular Dynamics Simulations Of Shocked Materials				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U. S. Army Research Laboratory, AMSRD-ARL-WM-BD Aberdeen Proving Ground, MD 21005-5069				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADM002075., The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

passing through silica glasses. We chose this system for two reasons: First, accurate and widely tested atomistic models are available (van Beest et. al., 1990; Yuan and Cormack, 2001; Cormack and Du, 2001). Secondly, these silica glasses have demonstrated failure waves after shock, i.e. waves propagate into these materials after compressive shock, behind which the material experiences a complete loss of tensile strength and substantial loss in shear strength. The result is the comminution (pulverization) of the material (Bless et. al., 1992). The detailed atomic level mechanism is not yet determined; however, one hypothesis is that the phenomenon is due to phase transformation of the amorphous glassy material to a crystalline state (Raiser and Clifton, 1994), something that can be readily observed using MD.

It is possible that the mechanism of failure of the material will be influenced by the dynamic introduction of a shear into the compressed metastable solid or in the presence of defects. The contributions of dynamic shear or various point defects (i.e. impurities, vacancy or interstitial defects) to failure of silicate glasses can be readily explored using MD. Our goal is to identify the key factors initiating and controlling the failure waves, and explore, through performing computer experiments, ways in which the behavior can be manipulated or defeated.

The multi-million atom simulations reported here are the largest classical MD simulations ever performed at the US Army Research Laboratory, and are used to exemplify a dramatic new modeling and simulation capability to the Army. Most importantly, though, the software used for the dynamics presented here can easily be used to model other energetic and non-energetic solids that are being developed to meet armor, armament, and materiel requirements that assure supremacy in future warfare. The results of such modeling will produce a molecular understanding of DOD structural materials and fatigue and failure mechanisms (micro-crack formation and propagation, twinning, cleavage), composites, light-weight structural materials, metals, ceramics, surface coatings (diamond coatings, nitride coatings) and nano-materials.

1.1 Potential Energy Surface (PES)

We have identified a model of silica (van Beest et. al., 1990) that has been extensively tested and that predicts several polymorphic states (Saika-Voivod et. al., 2004). This model, hereafter denoted as the BKS model after its authors (van Beest et. al., 1990), assumes that the potential energy for a system of N atoms can be described as the sum of interatomic interaction terms consisting of the superposition of Buckingham (6-exp)

(repulsion and dispersion) and Coulombic potentials of the form:

$$V^{Total} = \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \{A_{ij} \exp(-B_{ij}r) - C_{ij} / r^6 + \frac{q_i q_j}{4\pi\epsilon_0 r}\} \quad (1)$$

where r is the interatomic distance between atoms i and j , q_i and q_j are the electrostatic charges on the atoms, and ϵ_0 is the dielectric permittivity constant of free space. This same functional form is also used to describe sodium silicate (Cormack and Du, 2001).

2. THE SOFTWARE PARADIGM AND ENVIRONMENT

An important design goal of this project is to develop a suite of efficient, easy-to-use, extensible, portable and scalable molecular simulation tools for simulations of materials of critical DoD interest. To create an easy to use and extensible environment, this suite of software, entitled CoreXMD, relies on a standards based execution environment to require only a minimal knowledge of the details of the underlying software. At the same time the execution environment is rich enough to allow the expert to create unique advanced simulations. The code is executed by running the code executive which reads in an XML file, which hierarchically describe the simulation to be run, it then loads shared object libraries containing the needed components and executes them in the order specified by the XML file. The shared object libraries provide the extensibility as many may be used and each can come from different developers and they are independent of each other and the code executive.

CoreXMD libraries of subroutines, that provide all of the basic functionality for efficient highly scalable molecular dynamics simulations, include a variety of potentials from the simple to the chemically realistic.

The atomic potentials of interest are either short ranged, in that there is a distance (the cutoff radius) at which the interaction of the atoms falls to zero, or they have a short range component. The efficient calculation of these short-range interactions forms the basis of the solid performance and scalability of CoreXMD. We shall now examine several of these methods.

2.1 Cellular Decomposition

The first method used to enhance the performance of the code for both medium and large scale simulations is to decompose the simulation into smaller cells. Each atom is then assigned to a specific cell according to its position in the simulation space. Each process is assigned a set of cells by the spatial decomposition

algorithm and the atoms for the set of cells on a process reside on that process. To find all atoms within the cutoff radius for an atom of a given cell, all cells containing any point within the cutoff radius of any point within the given cell must be searched. These neighboring cells will often be assigned to other processes. Therefore, neighboring cells on other processes (and the atoms contained therein) will need to be communicated. The cellular decomposition provides a ready mapping of which atoms need to be communicated. To minimize the communications a heuristic in the spatial decomposition chooses one of three methods to divide up space among the processes. Slicing (dividing the simulation like a loaf of bread with one slice per process) of the simulation space has been shown to be the most efficient method for spatial decomposition (as long as each slice is at least as wide as the cutoff radius). In this method each process then only has to communicate with the two processes on either side of it in the spatial decomposition. For smaller systems wherein the slices would be thin enough that all of the atoms in the slice would need to be communicated with more than two other processes, it is more efficient to “julienne” (like French fries) or “dice” the simulation space. For the simulations reported here slicing is always used. The cellular decomposition simplifies the book keeping required for communication since only the atoms in a certain set of cells need to be communicated to another processor.

2.2 Sorting the Atomic Data by Cells

Substantial performance enhancements have been reported by Yao and coworkers (Yao et. al., 2006) by sorting the atomic data so that interacting atoms are stored more closely in physical memory. However they report only sorting into slices or layers, instead of sorting them into cells. When using a cellular decomposition method, sorting the atomic data by the cell to which the atom is assigned is, on its face, more productive. The main purpose of the sorting is to keep data that is needed for all interacting atoms close together in memory. When comparing the atoms in one cell to those in another sorting the data by cell allows vector operations, since the data for the atoms in each of the cells is a contiguous vector. This both minimizes cache misses and allows the use of highly optimized machine specific libraries for the vector operations, substantially enhancing the performance. The type of data sorting is an approximate sort, since the atoms are only sorted by the cell to which they belong. No order is set for atoms within a given cell, and as such the sort has a complexity of $O(N)$. Additionally for a cell based sort the sort order is already known from that atom assignment to the cells reducing the overall cost of the sort.

2.3 High Resolution Cellular Decomposition

Traditionally cellular decomposition has been performed with cells of dimensions equal to or larger than the cutoff radius (since the simulation space must be divided into a whole number of cells and may not be an exact multiple of the cutoff radius in size). However a simple analysis shows this to be quite inefficient. If the cell is cubic with each dimension being the cutoff radius, the actual search volume (the volume that must be searched to find all interaction atoms) is twenty-seven times the cube of the cutoff radius. The actual volume of interaction is $4/3\pi$ times the cube of the cutoff radius, or less than sixteen percent of the actual search volume. For cells with large dimensions the efficiency is even worse. By simply using a cell with dimensions that are an integral division of the cell used above, the efficiency can be improved. For example by cutting the cell dimension in half, the actual search space becomes 15.625 times the cube of the cutoff radius, for an efficiency of twenty-seven percent. However, increasing the number of cells increases the amount of memory required and the computational cost of using the cells. The optimal cell size has been found to range between one-half to one-fifth of the cutoff radius, depending on the density (Sutmann and Stegailov, 2006) and whether multiple cutoff radii are used in a simulation (Mattson and Rice, 1999). For the simulations reported here, the cell size is about one-half of the cutoff radius.

2.4 Just In Time Cellular Decomposition

A common method used to improve the efficiency of short range interaction calculations is the method of Verlet pair lists (Verlet, 1967). In this method a list of pairs of atoms that are within the cutoff radius plus some extra distance (usually referred to as the skin thickness) is constructed and only updated when the atoms have moved farther than one-half of the skin thickness. The skin acts like a buffer for atoms that may soon move within the cutoff radius. Verlet pair lists provide substantial performance increases for small systems, but are problematic for large systems because they are usually constructed in a brute force method, i.e., by calculating the distance from each atom to every other atom in the simulation. A cellular decomposition method can be used to construct the Verlet pair list which substantially improves performance for large systems sizes (Yao et. al., 2006). We have implemented such a cellular decomposition created Verlet pair list.

When deciding to use a cell list or a Verlet pair list, certain considerations must be made. Verlet pair lists do not need to be updated at every time step and are efficient for small systems. Cell lists, on the other hand, are not efficient for very small systems and traditionally require updating at every time step. Our just-in-time

method only updates the cell lists as required by the motion of the atoms, in the spirit of the Verlet lists through inclusion of “skin thickness” into the definition of the cell size. The just-in-time cell method can be used directly or to generate Verlet pair lists, depending on simulation requirements (such as limited memory or complex many-body potentials requiring multiple uses of pair information per time step). This just in time cellular decomposition also makes communication more efficient as atomic migration from one cell on one process to another needs to be performed when the atoms are resorted. This reduces the communication frequency for migration from every time step to that of the cell updates. Because of the large system sizes and memory requirements and use of a two body potential for the simulations reported herein, the just in time cellular decomposition has been used directly and the pair list has not been used.

3. VALIDATION OF THE METHOD AND MODEL

Before proceeding with large-scale simulations, we decided it necessary to compare our calculations against those generated using a well-established MD code, DL_POLY, Version 2. While this is a very sophisticated MD package and has been enhanced to treat large systems ($\sim 10,000$ atoms), it has not been demonstrated for multimillion atom simulations, nor does it have the capability to perform the specialized computer simulations associated with shock waves. The calculations run using DL_POLY and CoreMD used a supercell consisting of 864 atoms of α -quartz arranged in the experimental condition in the ambient state (Fig. 1).




Figure 1. Graphical depiction of a representative simulation supercell of α -quartz consisting of 540 atoms arranged in the experimental configuration at ambient conditions.

Periodic boundary conditions were imposed in all dimensions, and the interaction potential cutoff was set at 9 Å, due to lack of convergence of the Coulombic interactions at smaller distances. Thermostatting and barostatting rates were set to 2 and 1 ps-1, respectively. Ewald parameters were selected to produce an error in

the real and reciprocal space sums of no greater than 5×10^{-6} eV. Conditions of the simulation were $T=300$ K, $P = 1$ bar. At the beginning of the simulation, the atoms were arranged in the experimental configuration at room conditions; atomic velocities are selected from a Maxwell-Boltzmann distribution. An equilibration trajectory consisting of a large number of steps is performed in order to move the system away from the initial configuration. At the end of the equilibration period, a new trajectory is initiated using the final coordinates and velocities from the equilibration trajectory; time-averages of various properties are taken over the duration of this trajectory and given in Table 1.

Table 1. MD Simulation Results of α -quartz at $T=300$ K, $P= 0.001$ kbar.

	DLPOLY	CoreXMD
Time (ps)	2.250	2.545137
T (K)	300.57	301.00
P (kbar)	0.25562	0.046442
U (eV) (Total)	-16788	-16788.12
U/atom (eV)	-19.43	-19.43
Volume/atom ($\text{\AA}^3$)	12.74	12.75
a ($\text{\AA}$)	4.9392 (4.9141)	4.9352 (4.9141)
b ($\text{\AA}$)	4.9313 (4.9141)	4.9352 (4.9141)
c ($\text{\AA}$)	5.4402 (5.4060)	5.4421 (5.4060)
α ($\text{\AA}$)	90.0 (90)	90.0 (90)
β ($\text{\AA}$)	90.0 (90)	90.0 (90)
γ ($\text{\AA}$)	120.0 (120)	120.0 (120)
Density (atom/ $\text{\AA}^3$)	0.0784	0.0784
Energy Conservation	3.5×10^{-2}	2.8×10^{-5}

Expt. values from (Lager et. al., 1982) are in parentheses.

We next needed to confirm that we understood the model interaction potential. Very detailed molecular dynamics studies were performed by Tse and Klug (Tse and Klug, 1991) using the BKS model to explore the pressure-induced amorphization in α -quartz. In that study, isothermal-isobaric (NPT) MD simulations were performed for pressures ranging from 1 atm to 80 GPa; these correctly predicted static properties of several polymorphs and the crystalline-to-amorphous phase transition. The next series of calculations presented herein are attempts to reproduce the Tse and Klug work (Tse and Klug, 1991). While Tse and Klug provide little details of their calculations, such as initial conditions for the sequence of simulations, thermostatting and barostatting parameters, Ewald summation parameters and interaction potential cutoff, they report using a simulation cell consisting of only 576 atoms. In order to perform the same size of calculations using NPT-MD, the interaction potential could be no greater than 7.5 Å at 1 atm. The simulation cell would have to be considerably smaller for the simulations at higher pressures. It is well established that the Coulombic

interaction energy treated by Ewald summations is an approximation and the quality of the approximation is strongly dependent on choice of parameters and cutoff range of the real and reciprocal space sums in the technique. We found upon varying the cutoff range that a choice of a cutoff of less than 9 Å produced different energies that were dependent on the cutoff value; only at distances greater than 9 Å did the Coulombic energy converge to the proper value. Thus, it is possible that the Tse and Klug treatment of the Coulombic interactions using this model might be inadequate.

Figure 2 shows results from our NPT-MD calculations for pressure ranging from 1 atm up to over 40 GPa.




Figure 2. The calculated volume change (upper frame) and total energy versus pressure (lower frame) for α -quartz. The solid lines are values digitized from Fig. 1 of Tse and Klug (Tse and Klug, 1991).

One of our goals was to determine if we could simulate the pressure-induced amorphization of the crystal and have results in reasonable agreement with Tse and Klug. In these figures, we have superimposed digitized curves of the Tse and Klug calculations (Tse and Klug, 1991). The upper and lower frames of Fig. 2

show the volume change and total energy per atom as a function of pressure, respectively. The calculate volume changes with increasing pressure are in reasonable agreement with Tse and Klug with the exception of the location of the transition pressure (discussed hereafter). Table 2 shows the comparison of values of lattice parameters calculated at ambient conditions; these were only one set of two numerical values given by Tse and Klug in their work (Tse and Klug, 1991). Experimental values for the cell edge lengths are given in parentheses below the predicted values.

Table 2: NPT-MD predictions of lattice parameters of α -quartz at T=300 K, P=1 bar using BKS model.

	P (GPa)	a (Å)	b (Å)	C (Å)
Tse and Klug	0.0001	4.98 (4.9141)	4.98 (4.9141)	5.48 (5.4060)
CoreXMD	0.0001	4.9352 (4.9141)	4.9352 (4.9141)	5.4421 (5.4060)

Our predictions are in better agreement with experiment than Tse and Klug; discrepancies could be due to choices of simulation parameters such as choice of Ewald parameters and cutoff range.

In our calculations, the transition pressure from the crystalline to the amorphous state occurs around 19 GPa; Tse and Klug report that value to be 22 GPa. The energy per atom curves, while showing similar behavior, differ in magnitude by ~ 4 kJ/mol at the ambient state, with the Tse and Klug calculations being higher in energy than ours (note that their energies plotted along the ordinate of their Fig. 2 runs from high to low, rather than low to high). In other words, a first glance at this figure would suggest that the ambient state is higher in energy than the amorphous state; a careful examination of the ordinate indicates that such is not the case. The energies are in better agreement at pressures between 10 and 20 GPa, particularly at the transition pressure. However, we were concerned that our transition pressure was lower than that reported by Tse and Klug (Tse and Klug, 1991). Therefore, using the information they provided us for cell parameters at 22 GPa, we calculated thermodynamic averages for a fixed simulation cell shape and size using constant-volume-constant temperature (NVT) MD. Our calculations indicate that the pressure for these cell parameters is 20.5 GPa rather than the 22 GPa reported by Tse and Klug (Tse and Klug, 1991). Further our energy is -1856 kJ/mol, which appears to be ~ 6 kJ/mol lower in energy than the Tse and Klug results. Again, we suspect the discrepancies can be traced to choices of Ewald parameters, interaction potential cutoff range, and other simulation details that were not provided by Tse and Klug.

A final calculation was performed to determine whether the pressure-induced phase transition of α -quartz to the amorphous state at pressures > 20 GPa was reversible. To determine this, the RDFs for the ambient crystal and the amorphous state are compared with the RDF generated after releasing the pressure of the high-pressure system. Figure 3 shows such a comparison. In this, RDFs for the ambient crystal and amorphous states are given. The lack of long-range structure in the RDF for the amorphous structure denotes that the order associated with a crystal has disappeared. Also given is the RDF for the system after release of the pressure. In this simulation, the initial state is the simulation cell corresponding to the amorphous structure at 25 GPa; during equilibration, the pressure is relaxed from 25 GPa to 1 bar; a subsequent trajectory is then run to evaluate thermal averages and the RDF. Clearly, the RDF of this latter system more closely resembles the amorphous state at 25 GPa, indicating that the transformation is irreversible.




Figure 3. RDFs for α -quartz at 1 bar and 25 GPa at $T=300$ K. The curve denoted “amorphous” denotes the RDF calculated for a system in which the initial condition of the simulation corresponded to 25 GPa, and under which pressure was gradually released to 1 bar.

For multimillion atom simulations, the standard Ewald treatment of the Coulombic interaction is computationally infeasible since the reciprocal space requires an inordinately large number of k points. Various parallel scheme exist to treat Ewald summations for large systems, which we intend to incorporate into our codes at this time, however, we choose to use an approximation proposed by Linse and Andersen (Linse and Andersen, 1986) in which the reciprocal-space term is not used in the Ewald summation. Although this approximation, denoted as the spherical Ewald truncation method, has been known to fail for other systems, we tested this approximation using NPT-MD simulations on a small quartz system; the results show that differences in total energies are less than 0.02%.

4. SPEED-UP: SCALED AND FIXED

As proof of scalability and performance we have analyzed the performance of the software for this system under a variety of different scaling conditions. All calculations were performed on the ASC MSRC SGI Altix 3700, eagle. All calculations were based on an equilibrated 1080 atom unit cell of α -quartz with full periodic boundary conditions that was replicated in a super cell to meet the needs of the specific test. First, the scaled speed up is a measure of how the computation scales as the number of atoms per process is constant. For this we created a super cell from the initial simulation cell of 1080 atoms with 16 of these simulation cells in a slice, with one slice per process. This corresponds to the approximately 4.4 million atom system on 256 processes that we use for the shock wave calculation. The calculations were run for 10,000 time steps, with full output of atomic data (coordinates, velocity, forces, etc.) at every 1000 time steps. In Table 3 we show the number of processors being used, the total number of CPU seconds (the sum of the time used on each process in seconds), the wall clock time (the maximum time of all processes), and the percentage of optimal scaling. We can see that even at 256 processors a scaling of over eighty percent of optimum is achieved. The largest factor in the deviation from perfect scaling is the output of atomic data which is essentially a serial portion of the process that does not scale.

Table 3. Scaled speed-up of CoreXMD with 17280 atoms per process.

Processes	CPU Time(s)	Wall Clock(s)	% of Optimal
1	42,361	42361.00	100.00%
2	90,792	45396.00	93.31%
4	181,300	45325.00	93.46%
8	364,400	45550.00	93.00%
16	747,700	46731.25	90.65%
32	1,574,500	49203.13	86.09%
64	3,254,000	50843.75	83.32%
128	6,566,700	51302.34	82.57%
256	13,299,000	51949.22	81.54%

The fixed speed up denotes how the computation scales for a fixed system size or in other-words the same number of atoms for each simulation regardless of the number of processes being used. In our calculations the 1080 atom initial simulation cell was replicated sixty-four times, giving 1080 atoms per process for 64 processes and 69120 atoms per process on the single process run. This represents medium to small system sizes, depending on the number of processes being used. We only performed scaling calculations for up to 64 processors for fixed speedup since for larger number of

processors we would either have to have a ridiculously small number of atoms per processor, or increase the system size for all calculations giving an absurdly large system for a single processor. A one thousand atom simulation is quite small by today's standards for classical molecular dynamics. These calculations were performed for only 100 time steps with one atomic data output. In Table 4 we again show the number of processors, CPU Time, Wall Clock Time, and the percentage of optimal speed up. The fixed speed up for 64 processes is above eighty percent which is remarkable given the system size.

Table 4. Fixed speed-up of CoreXMD with 69120 atoms.

Processes	CPU Time(s)	Wall Clock(s)	% Optimal
1	1788	1788.00	100.00%
2	1790	895.00	99.89%
4	1820	455.00	98.24%
8	1870	233.75	95.61%
16	1971	123.19	90.72%
32	1993	62.27	89.73%
64	2187	34.17	81.76%

CONCLUSIONS

In this paper, we described and demonstrated CoreXMD, a multi-million atom molecular dynamics code designed for ease of portability, extensibility and usability. Fixed and scaled speedups for this software are demonstrated in molecular dynamics simulations of α -quartz, using up 256 processors and up to 4.4 million atoms. Scaled speed up is over 80% of optimal on 256 processors and fixed speed up is also over 80% on 64 processors for system sizes as small as 1,000 atoms per processor. The outstanding scaling for both small and large systems is somewhat unusual, since other multi-million atom simulation codes tend to perform well only for very large systems. This code significantly enhances multi-scale modeling capabilities of complex dynamic processes in materials of interest to the Army.

ACKNOWLEDGMENTS

This work was supported by the DOD High Performance Computing Modernization Program, under CHSSI projects CCM-5 and MBD-4 and the Office of Naval Research. We would also like to acknowledge the assistance of Dr. John Brennan in providing the initial atomic coordinates for the simulation cell.

REFERENCES

- Bless, S. J. and Brar, N. S. and Kanel, G. and Rosenberg, Z., 1992: Failure Waves in Glass, *J. Am. Ceramic Soc.*, **75**, 1002-1005.
- Bowers, K.J. Dror, R.O. and Shaw, D.E., 2006: The Midpoint Method for Parallelization of Particle Simulations, *J. Chem. Phys.*, **124**, 184109, 11 pages.
- Cormack, A.N. and Du, J., 2001: Molecular Dynamics Simulations of Soda-lime-silicate glasses, *J. Non-Crystalline Solids*, **293-295**, 283-289.
- Lager, G.A. and Jorgensen, J.D. And Rotella, F.J., 1982: Crystal Structure and Thermal Expansion of α -quartz SiO_2 at Low Temperatures, *J. Appl. Phys.*, **53**, 6751-6756.
- Linse, P. and Andersen, H.C., 1986: Truncation of Coulombic Interactions in Computer Simulations of liquids, *J. Chem. Phys.*, **85**, 3027-3041.
- Mattson, W. and Rice, B. M., 1999: Near-neighbor Calculations Using a Modified Cell-linked List Method, *Comput. Phys. Commun.*, **119**, 135-148.
- Raiser, G. and Clifton, R. J., 1994: Failure Waves in Uniaxial Compression of an Aluminosilicate Glass, *High Pressure Science and Technology*. Edited by S. C. Schmidt. Joint AIRAPT-APS Conference, Colorado Springs, CO, 1039-1042.
- Saika-Voivod, and I., Sciortino, and F. Grande, T. and Poole, P. H., 2004: Phase Diagram of Silica from Computer Simulation, *Phys. Rev. E*, **70**, 061507, 8 pages
- Sutmann, G. and Stegailov, V., 2006: Optimization of Neighbor List Techniques in Liquid Matter Simulations, *J. Mol. Liq.*, **125**, 197-203
- Tse, J. S. and Klug, D. D., 1991: The Structure and Dynamics of Silica Polymorphs Using a Two-body Effective Potential Model, *J. Chem. Phys.*, **95**, 9176-9185; Tse, J. S. and Klug, D. D., 1991: Mechanical Instability of α -quartz: A Molecular Dynamics Study, *Phys. Rev. Lett.*, **67**, 3559-3562.
- van Beest, B. W. H. and Kramer, G. J. and van Santen, R. A., 1990: Force Fields for Silicas and Aluminophosphates Based on Ab Initio Calculations, *Phys. Rev. Lett.*, **64**, 1955-1958.

Verlet, L., 1967: Computer "Experiments" on Classical Fluids. I. Thermodynamical Properties of Lennard-Jones Molecules, *Phys. Rev.*, **159**, 98-103.

Yao, Z. and Wang, J. and Liu, G. and Cheng, M., 2004: Improved Neighbor List Algorithm in Molecular Simulations Using Cell Decomposition and Data

Sorting Method, *Comput. Phys. Commun.*, **161**, 27-35.

Yuan, X. and Cormack, A. N., 2001: Local structures of MD-modeled vitreous silica and sodium silicate glasses, *J. Non-Crystalline Solids*, **283**, 69-87.