

Scalable atomistic simulation algorithms for materials research

Aiichiro Nakano^a, Rajiv K. Kalia^a, Priya Vashishta^a, Timothy J. Campbell^b, Shuji Ogata^c, Fuyuki Shimojo^d and Subhash Saini^e

^a*Department of Computer Science, Department of Materials Science & Engineering, Department of Physics & Astronomy, Department of Biomedical Engineering, University of Southern California, Los Angeles, CA 90089, USA, and Concurrent Computing Laboratory for Materials Simulations, Biological Computation and Visualization Center, Department of Computer Science, Department of Physics & Astronomy, Louisiana State University, Baton Rouge, LA 70803, USA*

Tel.: +1 225 578 1342; Fax: +1 225 578 5855; E-mail: {nakano, kalia, priyav}@bit.csc.lsu.edu

^b*Naval Oceanographic Office, Stennis Space Center, MS 39529, USA*

E-mail: tjcamp@navo.hpc.mil

^c*Department of Applied Sciences, Yamaguchi University, Ube 755-8611, Japan E-mail:*

ogata@po.cc.yamaguchi-u.ac.jp

^d*Faculty of Integrated Arts and Sciences, Hiroshima University, Higashi-Hiroshima 739-0046, Japan*

E-mail: shimojo@minerva.ias.hiroshima-u.ac.jp

^e*IT Modeling and Simulation, NASA Ames Research Center, Moffett Field, CA 94035-1000, USA*

E-mail: saini@nas.nasa.gov

Abstract. A suite of scalable atomistic simulation programs has been developed for materials research based on space-time multiresolution algorithms. Design and analysis of parallel algorithms are presented for molecular dynamics (MD) simulations and quantum-mechanical (QM) calculations based on the density functional theory. Performance tests have been carried out on 1,088-processor Cray T3E and 1,280-processor IBM SP3 computers. The linear-scaling algorithms have enabled 6.44-billion-atom MD and 111,000-atom QM calculations on 1,024 SP3 processors with parallel efficiency well over 90%. The production-quality programs also feature wavelet-based computational-space decomposition for adaptive load balancing, spacefilling-curve-based adaptive data compression with user-defined error bound for scalable I/O, and octree-based fast visibility culling for immersive and interactive visualization of massive simulation data.

Keywords: Parallel computing, molecular dynamics, variable-charge molecular dynamics, quantum mechanics, density functional theory, load balancing, data compression

1. Introduction

Modern design of high-performance materials and devices focuses on controlling structures at diverse length scales from atomic to macroscopic [28]. Rich variety of atomistic simulation methods -ranging from empirical molecular-dynamics (MD) simulations to ab initio quantum-mechanical (QM) calculations- are expected to play an important role in scaling down engineering concepts to nanometer scales. Recent advances

in computational methodologies and massively parallel computers have made it possible to carry out 10–100 million atom simulations of real materials and devices typically on 10–100 processors [1,10,38].

This paper describes our efforts to enable very large-scale atomistic simulations involving multibillion atoms by designing scalable and portable simulation algorithms [16]. In the next section, we describe linear-scaling parallel algorithms for MD and QM calculations. Section 3 discusses software tools to support

billion-atom simulations. Results of benchmark tests are given in Section 4, and Section 5 contains conclusions.

2. Parallel atomistic simulation algorithms

We have developed a suite of scalable MD and QM algorithms for materials simulations. The linear-scaling algorithms encompass a wide spectrum of physical reality: i) classical MD based on a many-body interatomic potential model; ii) environment-dependent, variable-charge MD; and iii) self-consistent QM calculation based on the density functional theory (DFT). These algorithms deal with the following three problems that are common in many scientific and engineering applications: i) all-pairs function evaluation in the N -body problem; ii) dense linear system of equations in the variable N -charge problem; and iii) exhaustive combinatorial enumeration in the quantum N -body problem. This section describes general algorithmic techniques to obtain approximate solutions to these problems in $O(N)$ time, including i) clustering, ii) hierarchical abstraction, and iii) the analysis of asymptotic solution properties.

2.1. Multiresolution molecular dynamics algorithm

In the MD approach, one obtains the phase-space trajectories of the system (positions and velocities of all atoms at all time) [31]. Atomic force laws for describing how atoms interact with each other is mathematically encoded in the interatomic potential energy, $E_{MD}(\mathbf{r}^N)$, which is a function of the positions of all N atoms, $\mathbf{r}^N = \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}^N\}$, in the system. In our many-body interatomic potential scheme, $E_{MD}(\mathbf{r}^N)$ is expressed as an analytic function that depends on relative positions of atomic pairs and triples [38]. Time evolution of \mathbf{r}^N is governed by a set of coupled ordinary differential equations. For interatomic potentials with finite ranges, the computational cost can be made $O(N)$ using a linked-cell-list approach [31].

Our multiresolution molecular dynamics (MRMD) algorithm [25] also uses an approach called the multiple time-scale (MTS) method [21,23,37]. The MTS method uses different force-update schedules for different force components, i.e., forces from the nearest-neighbor atoms are computed at every MD step, and forces from farther atoms are computed with less frequency. This not only reduces the computational cost but also enhances the data locality, and accordingly the

parallel efficiency is increased. These different force components are combined using a reversible symplectic integrator [37], and the resulting algorithm consists of nested loops to use forces from different spatial regions. It has been proven that the phase-space volume occupied by atoms is a simulation-loop invariant in this algorithm [37], and this loop invariant results in excellent long-time stability of the solutions.

For parallelization of MD simulations, we use spatial decomposition [25,31]. The total volume of the system is divided into P subsystems of equal volume, and each subsystem is assigned to a processor in an array of P processors. To calculate the force on an atom in a subsystem, the coordinates of the atoms in the boundaries of neighbor subsystems are “cached” from the corresponding processors. After updating the atomic positions due to a time-stepping procedure, some atoms may have moved out of its subsystem. These atoms are “migrated” to the proper neighbor processors. With the spatial decomposition, the computation scales as N/P while communication scales in proportion to $(N/P)^{2/3}$ for an N -atom system.

2.2. Variable-charge molecular dynamics

Physical realism of MD simulations is greatly enhanced by incorporating variable atomic charges that dynamically adapt to the local environment [4,35]. However, the increased realism of this variable-charge molecular dynamics (VCMD) is accompanied by increased computational complexity, $O(N^3)$, for solving a dense linear system to determine atomic charges at every MD step. We have reduced this complexity to $O(N)$ [20] by combining: i) the fast multipole method (FMM) [12,25] to perform matrix-vector multiplications with $O(N)$ operations; and ii) an iterative minimization approach to initialize the solution using the previous MD step’s charges. These algorithms reduce the amortized computational cost averaged over simulation steps to $O(N)$.

In the FMM, the first essential idea is clustering, i.e., instead of computing interactions between all atomic pairs, atoms are clustered and cluster-cluster interactions are computed. At the source of interaction, cluster information is encapsulated in terms of the multipoles of the charge distribution, and this multipole expansion provides a well-defined error bound. At the destination, on the other hand, the electrostatic potential is expanded in terms of local terms, which is similar to the Taylor expansion. The second essential idea is to use larger clusters for longer distances, in order to reduce

the computation and keep the error constant. This is achieved by recursively subdividing the simulation box into smaller cells to form an octree data structure. The $O(N)$ algorithm traverses this tree twice. In the upward pass, multipoles are computed for all cells at all levels. First the multipoles of the leaf cells are computed using atomic charges and coordinates, and the multipoles of these children cells are shifted and combined to obtain the multipoles of the parent cells. This procedure is repeated until the root of the octree is reached. In the downward pass, these multipoles are translated to local terms for all cells at all levels, starting from the root. For a given cell at each level, only the multipoles of a constant number of interactive cells contribute to the local terms. Contributions from farther cells have already been computed at the previous coarse level, and they are inherited from the parent cell. On the other hand, the contributions from the nearest-neighbor cells will be computed at the next fine level. This procedure is repeated until we reach the leaf level. Finally, the nearest-neighbor-cell contributions at the leaf level are computed by direct summation over atoms. Since constant computation is performed at each of the $O(N)$ octree nodes, the complexity of the FMM algorithm is $O(N)$.

On parallel computers, upper-level octree cells are handled globally on all P processors if the number of cells is less than P , and this introduces an $O(\log P)$ overhead. At the lower octree levels, on the other hand, spatial decomposition is employed, so that the computation on each processor is $O(N/P)$. The spatial decomposition necessitates the caching of surface-cell information from the nearest-neighbor processors, which involves $O((N/P)^{2/3})$ communication. For typical coarse-grained applications ($N/P \sim 10^5 - 10^6$ and $P < 10^3$), both the global ($\log P$) and the nearest-neighbor ($(N/P)^{2/3}$) communications are negligible compared with the $O(N/P)$ computation.

To further accelerate the convergence of the iterative solution mentioned above, we have developed a multilevel preconditioned conjugate-gradient (MPCG) method [20] by splitting the Coulomb-interaction matrix into short- and long-range components and using the sparse short-range matrix as a preconditioner. The extensive use of the sparse preconditioner enhances the data locality, and accordingly the parallel efficiency is increased.

Tree-based algorithms such as the FMM have been used extensively to perform massively parallel computer simulations of gravitational systems in astrophysics [39]. The FMM has also been used in conjunc-

tion with the MTS in parallel MD simulations of materials [25] and biomolecular systems [15,32]. Although these space-time multiresolution algorithms deal with $O(N^2)$ problems, their basic algorithmic structures can be reused for the $O(N^3)$ VCMD. For example, the computational cost of our VCMD code is further amortized by taking advantage of the algorithmic similarity between the MTS and the multilevel-preconditioning algorithms, i.e., by reusing a doubly nested loop with associated neighbor-list construction for both the MTS method for time-stepping and multilevel preconditioning for determining charges [20].

2.3. Linear-scaling quantum-mechanical calculation based on the density functional theory

Empirical interatomic potentials used in MD simulations fail to describe chemical processes. Instead, interatomic interaction in reactive regions needs to be calculated by a QM method that can describe breaking and formation of bonds. An atom consists of a nucleus and surrounding electrons, and quantum mechanics explicitly treats the electronic degrees-of-freedom. Since each electron's wave function is a linear combination of many states, the combinatorial solution space for the many-electron problem is exponentially large. The density functional theory (DFT) avoids the exhaustive enumeration of many-electron correlations by solving M single-electron problems in a common average environment (M is the number of independent wave functions and is on the order of N).¹ As a result, the problem is reduced to a self-consistent matrix eigenvalue problem, which can be solved with $O(M^3)$ operations [13,18,27]. The DFT problem can also be formulated as a minimization of the energy, $E_{QM}(\mathbf{r}^N, \psi^M)$, with respect to electron wave functions, $\psi^M(\mathbf{r}) = \{\psi_1(\mathbf{r}), \psi_2(\mathbf{r}), \dots, \psi_M(\mathbf{r})\}$, subject to orthonormalization constraints between the wave functions,

$$\int d\mathbf{r} \psi_m^*(\mathbf{r}) \psi_n(\mathbf{r}) = \begin{cases} 1 & (m = n) \\ 0 & (m \neq n) \end{cases}. \quad (1)$$

Efficient parallel implementation of DFT is possible with real-space approaches based on higher-order finite differencing [6] and multigrid acceleration [3,9]. We include electron-ion interaction using norm-conserving pseudopotentials [36] and the exchange-correlation energy associated with electron-electron interaction in a

¹Walter Kohn received a 1998 Nobel chemistry prize for the development of the DFT.

generalized gradient approximation [29]. For larger systems ($M > 1,000$), however, the $O(M^3)$ orthonormalization becomes the bottleneck.

For scalable DFT calculations, linear-scaling algorithms are essential [11]. We have implemented [34] an $O(M)$ algorithm [19] based on unconstrained minimization of a modified energy functional and a localized-basis approximation. This algorithm is based on the observation that, for most materials at most temperatures, the off-diagonal elements of the density matrix,

$$\rho(\mathbf{r}, \mathbf{r}') = \sum_{n=1}^M \psi_n^*(\mathbf{r})\psi_n(\mathbf{r}'), \quad (2)$$

decays exponentially [11], i.e., $\rho(\mathbf{r}, \mathbf{r}') \propto \exp(-C|\mathbf{r} - \mathbf{r}'|)$ for $|\mathbf{r} - \mathbf{r}'| \rightarrow \infty$ (C is a constant). Such a diagonally-dominant matrix can be represented by maximally localizing each wave function, $\psi_n(\mathbf{r})$, by a unitary transformation and then truncating it with a finite cut-off radius. A Lagrange-multiplier-like technique is also used to perform unconstrained minimization, avoiding the $O(M^3)$ orthonormalization procedure. In the parallel linear-scaling density functional theory (LSDFT) algorithm, the computation time scales as $O(M/P)$ on P processors, whereas the communication scales as $O((M/P)^{2/3})$. This is in contrast to the $O(M(M/P)^{2/3})$ communication in the conventional parallel real-space DFT algorithm. Global communication for calculating overlap integrals of the wave functions (which scales as $M^2 \log P$ in the conventional DFT algorithm) is unnecessary in the linear-scaling algorithm.

3. Software tools

Practical simulations involving multibillion atoms are associated with a number of computational challenges, which have been addressed by a number of software tools.

3.1. Wavelet-based adaptive computational-space decomposition for load balancing

Many MD simulations are characterized by irregular atomic distribution and associated load imbalance. For irregular data structures, the number of atoms assigned to each processor varies significantly, and this load imbalance degrades the parallel efficiency. The load-balancing problem can be stated as an optimization problem, i.e., one minimizes the load-imbalance

cost as well as the size and the number of messages:

$$\begin{aligned} E = & t_{\text{comp}}(\max_p |\{i|\mathbf{r}_i \in p\}|) \\ & + t_{\text{comm}}(\max_p |\{i|\|\mathbf{r}_i - \partial p\| < r_c\}|), \quad (3) \\ & + t_{\text{latency}}(\max_p [N_{\text{message}}(p)]) \end{aligned}$$

where the three terms are the load-imbalance cost, the size of messages, and the number of messages, respectively. In Eq. (3), ∂p and $N_{\text{message}}(p)$ denote the boundary surface of the physical volume assigned to processor p , and the number of messages per MD step for p , respectively. The expression, $\max_p f(p)$, denotes the maximum value of function $f(p)$ over all the processors, and r_c is the range of the interatomic potential. The prefactors, t_{comp} , t_{comm} and t_{latency} , are constants related to the processor speed, communication bandwidth and latency, respectively, and they are determined experimentally by test runs on the parallel computer under consideration.

To minimize the number of messages, we preserve the 3D mesh topology, so that message passing is performed in a structured way in only 6 steps [25]. To minimize the load imbalance cost as well as the message size, we have developed a computational-space decomposition scheme [22,24]. The main idea of this scheme is that the computational space shrinks where the workload density is high and expands where the density is low, so that the workload is uniformly distributed in the computational space. To implement the curved computational space, we introduce a curvilinear coordinate transformation,

$$\xi = \mathbf{x} + \mathbf{u}(x). \quad (4)$$

where \mathbf{x} is a position in the physical Euclidean space and $\mathbf{u}(\mathbf{x})$ is a deformation field. We then use regular 3D mesh topology in the computational space, ξ , to map atom i to processor p in an array of $P_x \times P_y \times P_z$ processors:

$$\begin{cases} p(\xi_i) = p_x(\xi_{ix})P_yP_z + p_y(\xi_{iy})P_z + p_z(\xi_{iz}) \\ p_\alpha(\xi_{i\alpha}) = \lfloor \xi_{i\alpha}P_\alpha/L_\alpha \rfloor \quad (\alpha = x, y, z) \end{cases}, \quad (5)$$

where $\xi_i = (\xi_{ix}, \xi_{iy}, \xi_{iz})$ is the coordinate of atom i and L_α is the simulation box size in the α direction in the computational space. This regular 3D mesh partition in the computational space results in curved partition boundaries in the physical space, \mathbf{x} . The load-imbalance and communication costs are minimized as a functional of the coordinate transformation, $\xi(\mathbf{x})$, using simulated annealing. We have found that wavelet representation leads to compact representation of curved partition boundaries, and accordingly to fast convergence of the minimization procedure [22].

3.2. Spacefilling-curve-based adaptive data compression for scalable I/O

A 1.5-billion-atom MD simulation we are currently performing produces 150 GB of data per frame (or per minute), including atomic species, positions, velocities, and stresses. For scalable input/output (I/O) of such large datasets, we have designed a data compression algorithm [26]. It uses octree indexing and sorts atoms accordingly on the resulting spacefilling curve. By storing differences between successive atomic coordinates, the I/O requirement for the same error tolerance level reduces from $O(N \log N)$ to $O(N)$. An adaptive, variable-length encoding scheme is used to make the scheme tolerant to outliers and optimized dynamically. An order-of-magnitude improvement in the I/O performance was achieved [26] for actual MD data with user-controlled error bound [40].

Another important issue in data management is the analysis of simulation results.² For atomistic simulations of materials, a challenge is to extract topological defects such as dislocations from massive data with large thermal noises. Graph data structures have played an important role in analyzing atomistic data [2, 7, 14], where vertices and edges represent atoms and bonds, respectively. Recently, we have used a shortest-path ring analysis to study intermediate-range orders in amorphous materials [8] and an edge-based indexing to detect grain boundaries in semiconductors nanocrystals [17].

3.3. Octree-based fast visibility culling for immersive and interactive visualization

Interactive exploration of large-scale atomistic simulations is important for identifying and tracking atomic features that are responsible for macroscopic phenomena, and an immersive and interactive virtual environment is an ideal platform for such explorative visualization, see Fig. 1.

We have developed a scalable visualization system to allow the viewer to walk through multibillion atoms [33]. The system uses fast visibility culling based on the octree data structure to reduce the number of atoms sent to the graphics pipeline. Multiresolution rendering is used to further speed up the rendering process. The visibility-culling task is offloaded to a

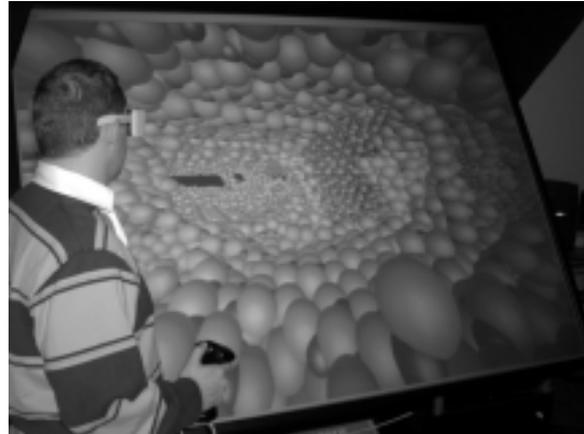


Fig. 1. A researcher investigating a fracture in a 1.5-billion-atom model of a ceramic fiber composite material rendered in an ImmersaDesk virtual environment at our Concurrent Computing Laboratory for Materials Simulations.

PC cluster so that the graphics server is dedicated to rendering. The resulting visualization system renders a billion-atom system at nearly interactive frame rates on a dual processor SGI Onyx2 with an InfiniteReality2 graphics pipeline connected to four PCs running Linux 6.2 each with an 800 MHz Pentium III processor and 512 MB RAM.

4. Performance tests

Benchmark tests of the three parallel algorithms – MRMD, VCMD, and LSDFT – have been performed on the Cray T3E and the IBM SP3 computers at the U.S. Naval Oceanographic Office (NAVO) Major Shared Resource Center. All the three programs are written using MPI (Message Passing Interface) for message passing. The T3E at NAVO at the time of the benchmark tests consisted of 1,088 Digital Alpha processors with clock speed 450 MHz and 256 GB memory. The IBM SP3 at NAVO is configured with 375 MHz Power3 CPUs and has 334 nodes with 4 CPUs and 4 GB of memory per node.

Figure 2 shows the execution time of the MRMD algorithm for silica material as a function of the number of processors, P . In this algorithm, the interatomic potential energy is split into the long-range and short-range contributions, where the long-range contribution is computed after 10 steps. We scale the system size linearly with the number of processors, so that the number of atoms, $N = 648,000 P$. On the T3E, execution time increases only slightly as a function of P , and

²For an extensive list of analysis tools for MD simulations, see <http://www.ks.uiuc.edu/Research/MMTools>.

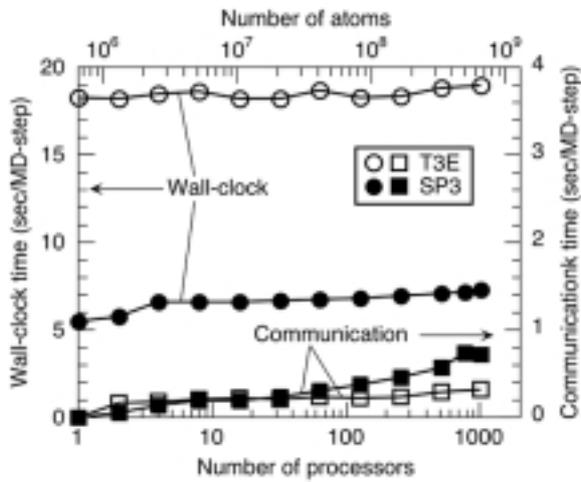


Fig. 2. Wall-clock (circles) and communication (squares) times per time step of the MD algorithm with scaled workloads – 648,000 P atom silica systems on P processors ($P = 1, \dots, 1,024$) of Cray T3E (open symbols) and IBM SP3 (solid symbols).

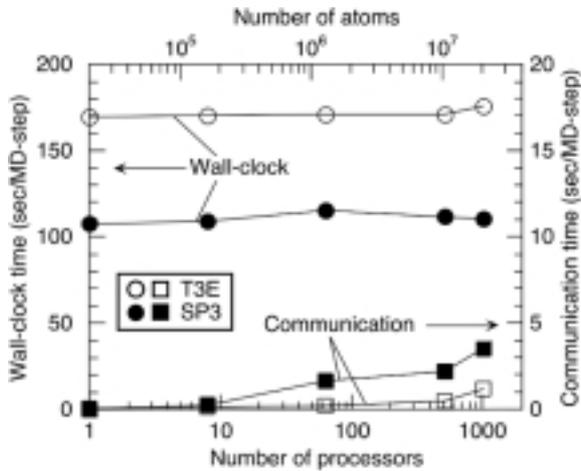


Fig. 3. Wall-clock (circles) and communication (squares) times for the VCMD algorithm with scaled workloads – 20,160 P atom alumina systems on P processors ($P = 1, \dots, 1,024$) of Cray T3E (open symbols) and IBM SP3 (solid symbols).

this signifies an excellent parallel efficiency. On 1,024 processors, the parallel efficiency is as high as 97%. The computational time on the SP3 is significantly less than that on the T3E, but with increased communication time.

Figure 3 shows the performance of the VCMD algorithm with scaled workloads – 20,160 P -atom alumina systems on P processors ($P = 1, \dots, 1,024$). Multipoles up to $l_p = 6$ are taken and the largest number of leaf octree cells is 8^6 ($P = 1,024$) in the FMM. The wall-clock time increases only slightly as a function of

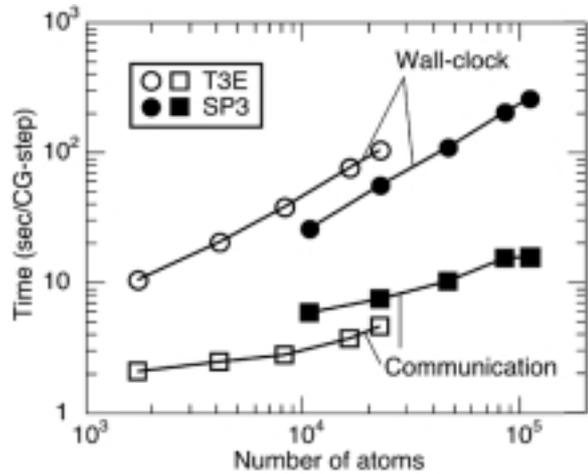


Fig. 4. Wall-clock (circles) and communication (squares) times per CG step as a function of the number of atoms for the parallel LSDFT algorithm on Cray T3E and IBM SP3 computers. The system is gallium arsenide crystal in the zinc-blende structure. The number of processors is 1,024.

P , and the memory-bound parallel efficiency is 0.96 on 1,024 processors on the T3E.

In the LSDFT calculations for gallium arsenide material, the localization region for the wave functions is defined as a spherical space with radius 4.4 Å. Figure 4 shows the wall-clock and communication times per CG iteration on 1,024 T3E and SP3 processors. The wall-clock time scales linearly with N above $N \sim 10,000$ (the number of wave functions, $M = 2N$). For the largest system ($N = 110,592$), the parallel efficiency is estimated to be 93%. The interprocessor communication scales as $O(N^{0.6})$ for $N > 10,000$.

Major design parameters for MD simulations of materials include the number of atoms in the simulated system and the methodologies to compute interatomic forces (classically in MRMD, semiempirically in VCMD, or quantum-mechanically [5] in LSDFT). Figure 5 shows a design-space diagram for classical and quantum-mechanical MD simulations on 1,024 T3E and SP3 processors. (For the LSDFT, one MD step involves 3 self-consistent DFT iterations each consisting of 20 CG steps.) The figure demonstrates linear scaling for all the three algorithms, with prefactors spanning seven-orders-of-magnitude. The largest benchmark tests in this study include 6.44-billion-atom MRMD and 111,000-atom LSDFT calculations on 1,024 SP3 processors.

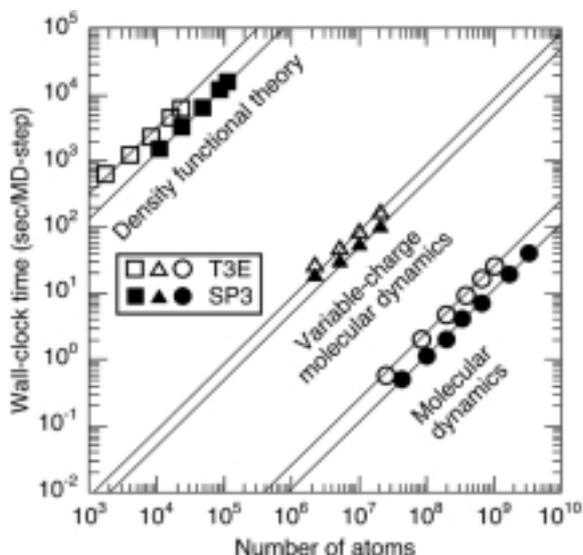


Fig. 5. Design-space diagram for MD and QM simulations on 1,024 Cray T3E processors (open symbols) and on 1,024 IBM SP3 processors (solid symbols). The figure shows wall-clock time per MD step as a function of the number of atoms for three linear-scaling algorithms: Classical MD (MRMD, circles); environment-dependant variable-charge MD (VCMD, triangles); and, quantum-mechanical MD based on the DFT (LSDFT, squares). Lines show $O(N)$ scaling.

5. Conclusions

Modern MD simulations of materials started in 1964 when Aneesur Rahman simulated 864 argon atoms on a CDC 3600 computer [30]. Assuming a simple exponential growth, the number of atoms that can be simulated in classical MD simulations has doubled every 19 months to reach 6.44 billion atoms in this study. Similarly, the number of atoms in DFT-based ab initio MD simulations (started by Roberto Car and Michelle Parrinello in 1985 for 8 Si atoms [5]) has doubled every 12 months to 111,000 atoms in this study. Petaflop computers anticipated to be built in the next ten years will maintain the growth rates in these “MD Moore’s Laws”, and we will be able to perform 10^{12} -atom classical and 10^7 -atom quantum MD simulations on such computers. Multiresolution approaches used in our algorithms, combined with cache-conscious techniques, will be essential to achieve scalability on petaflop architectures.

Atomistic simulations have now reached a scale such that they must be performed in a metacomputing environment of geographically-distributed multiple supercomputers. Such efforts are underway on NASA’s Information Power Grid (www.ipg.nasa.gov).

Acknowledgements

This work was partially supported by AFOSR, ARL, DOE, NASA, NSF, and Louisiana Board of Regents. Benchmark tests were performed using the 1,088-node Cray T3E and the 1,280-processor IBM SP3 computers at Department of Defense’s Naval Oceanographic Office (NAVO) Major Shared Resource Center under a DoD Challenge Project. Programs have been developed using parallel computers at the NASA Ames Research Center and the 166-processor PC cluster at the Concurrent Computing Laboratory for Materials Simulations (CCLMS) at Louisiana State University. Visualization was performed using the ImmersaDesk virtual environment at the CCLMS in collaboration with Mr. Andy Haas, Dr. Xinian Liu, Mr. Paul Miller, Mr. Ashish Sharma, and Mr. Wei Zhao.

References

- [1] F.F. Abraham, Portrait of a crack: Rapid fracture mechanics using parallel molecular dynamics, *IEEE Computational Science & Engineering* **4**(2) (1997), 66–77.
- [2] M.J. Alava, P.M. Duxbury, C.F. Moukarzel and H. Rieger, Exact combinatorial algorithms: Ground states of disordered systems, in: *Phase Transitions and Critical Phenomena*, (Vol. 18), J. Lebowitz and C. Domb, eds, Academic Press, London, 2000.
- [3] T.L. Beck, Real-space mesh techniques in density-functional theory, *Reviews of Modern Physics* **72** (2000), 1041–1080.
- [4] T.J. Campbell, R.K. Kalia, A. Nakano, P. Vashishta, S. Ogata and S. Rodgers, Dynamics of oxidation of aluminum nanoclusters using variable charge molecular-dynamics simulations on parallel computers, *Physical Review Letters* **82** (1999), 4866–4869.
- [5] R. Car and M. Parrinello, Unified approach for molecular dynamics and density functional theory, *Physical Review Letters* **55** (1985), 2471–2474.
- [6] J.R. Chelikowsky, Y. Saad, S. Ögüt, I. Vasiliev and A. Stathopoulos, Electronic structure methods for predicting the properties of materials: Grids in space, *Physica Status Solidi* **B217** (2000), 173–195.
- [7] D.J. Cook, L.B. Holder, G. Galal and R. Maglothlin, Approaches to parallel graph-based knowledge discovery, *Journal of Parallel and Distributed Computing* **61** (2001), 427–446.
- [8] I. Ebbsjö, R.K. Kalia, A. Nakano, J.P. Rino and P. Vashishta, Topology of amorphous gallium arsenide on intermediate length scales: A molecular dynamics study, *Journal of Applied Physics* **87** (2000), 7708–7711.
- [9] J.-L. Fattebert and J. Bernholc, Towards grid-based $O(N)$ density-functional theory methods: Optimized nonorthogonal orbitals and multigrid acceleration, *Physical Review* **B62** (2000), 1713–1722.
- [10] T.C. Germann and P.S. Lomdahl, Recent advances in large-scale atomistic materials simulations, *Computing in Science and Engineering* **1**(2) (1999), 10–11.
- [11] S. Goedecker, Linear scaling electronic structure methods, *Reviews of Modern Physics* **71** (1999), 1085–1123.

- [12] L. Greengard and V. Rokhlin, A fast algorithm for particle simulations, *Journal of Computational Physics* **73** (1987), 325–348.
- [13] P. Hohenberg and W. Kohn, Inhomogeneous electron gas, *Physical Review* **136** (1964), B864–B871.
- [14] D.J. Jacobs and B. Henderson, An algorithm for two-dimensional rigidity percolation: The pebble game, *Journal of Computational Physics* **137** (1997), 346–365.
- [15] L. Kalé, R. Skeel, M. Bhandarkar, R. Brunner, A. Gursoy, N. Krawetz, J. Phillips, A. Shinozaki, K. Varadarajan and K. Schulten, NAMD2: Greater scalability for parallel molecular dynamics, *Journal of Computational Physics* **151** (1999), 283–312.
- [16] R.K. Kalia, T.J. Campbell, A. Chatterjee, A. Nakano, P. Vashishta and S. Ogata, Multiresolution algorithms for massively parallel molecular dynamics simulations of nanostructured materials, *Computer Physics Communications* **128** (2000), 245–259.
- [17] S. Kodiyalam, R.K. Kalia, H. Kikuchi, A. Nakano, F. Shimojo and P. Vashishta, Grain boundaries in gallium arsenide nanocrystals under pressure: A parallel molecular-dynamics study, *Physical Review Letters* **86** (2001), 55–58.
- [18] W. Kohn and P. Vashishta, General density functional theory, in: *Inhomogeneous Electron Gas*, N.H. March and S. Lundqvist, Plenum, New York, 1983, pp. 79–184.
- [19] F. Mauri and G. Galli, Electronic-structure calculations and molecular-dynamics simulations with linear system-size scaling, *Physical Review* **B50** (1994), 4316–4326.
- [20] A. Nakano, Parallel multilevel preconditioned conjugate-gradient approach to variable-charge molecular dynamics, *Computer Physics Communications* **104** (1997), 59–69.
- [21] A. Nakano, Fuzzy clustering approach to hierarchical molecular dynamics simulation of multiscale materials phenomena, *Computer Physics Communications* **105** (1997), 139–150.
- [22] A. Nakano, Multiresolution load balancing in curved space: The wavelet representation, *Concurrency: Practice and Experience* **11** (1999), 343–353.
- [23] A. Nakano, A rigid-body based multiple time-scale molecular dynamics simulation of nanophase materials, *The International Journal of High Performance Computing Applications* **13** (1999), 154–162.
- [24] A. Nakano and T.J. Campbell, An adaptive curvilinear-coordinate approach to dynamic load balancing of parallel multi-resolution molecular dynamics, *Parallel Computing* **23** (1997), 1461–1478.
- [25] A. Nakano, R.K. Kalia and P. Vashishta, Multiresolution molecular dynamics algorithm for realistic materials modeling on parallel computers, *Computer Physics Communications* **83** (1994), 197–214.
- [26] A. Omeltchenko, T.J. Campbell, R.K. Kalia, X. Liu, A. Nakano and P. Vashishta, Scalable I/O of large-scale molecular-dynamics simulations: A data-compression algorithm, *Computer Physics Communications* **131** (2000), 78–85.
- [27] M.C. Payne, M.P. Teter, D.C. Allan, T.A. Arias and J.D. Joannopoulos, Iterative minimization techniques for ab initio total energy calculations: molecular dynamics and conjugate gradients, *Reviews of Modern Physics* **64** (1992), 1045–1097.
- [28] A. Pechenik, R.K. Kalia and P. Vashishta, *Computer-Aided Design of High-Temperature Materials*, Oxford Univ. Press, Oxford, 1999.
- [29] J.P. Perdew, K. Burke and M. Ernzerhof, Generalized gradient approximation made simple, *Physical Review Letters* **77** (1996), 3865–3868.
- [30] A. Rahman, Correlations in the motion of atoms in liquid argon, *Physical Review* **136**(2A) (1964), A405–A411.
- [31] D.C. Rapaport, *The Art of Molecular Dynamics Simulation*, Cambridge Univ. Press, Cambridge, 1995.
- [32] T. Schlick, R.D. Skeel, A.T. Brunger, L.V. Kalé, J.A. Board, J. Hermans and K. Schulten, Algorithmic challenges in computational molecular biophysics, *Journal of Computational Physics* **151** (1999), 9–48.
- [33] A. Sharma, P. Miller, X. Liu, A. Nakano, R.K. Kalia, P. Vashishta, W. Zhao, T.J. Campbell and A. Haas, Immersive and interactive exploration of billion-atom systems, in: *Proceedings of 2002 IEEE Virtual Reality Conference*, IEEE Computer Society, Los Alamitos, CA, 2002, pp. 217–223. [A shorter version has appeared in: A. Nakano, R.K. Kalia and P. Vashishta, Scalable molecular-dynamics, visualization, and data-management algorithms for materials simulations, *Computing in Science & Engineering* **1**(5) (1999), 39–47.]
- [34] F. Shimojo, T.J. Campbell, R.K. Kalia, A. Nakano, P. Vashishta, S. Ogata and K. Tsuruta, A scalable molecular-dynamics-algorithm suite for materials simulations: Design-space diagram on 1,024 Cray T3E processors, *Future Generation Computer Systems* **17** (2000), 279–291.
- [35] F.H. Streitz and J.W. Mintmire, Electrostatic potentials for metal-oxide surfaces and interfaces, *Physical Review* **B50** (1994), 11996–12003.
- [36] N. Troullier and J.L. Martins, Efficient pseudopotentials for plane-wave calculations, *Physical Review* **B43** (1991), 1993–2006.
- [37] M.E. Tuckerman, D.A. Yarne, S.O. Samuelson, A.L. Hughes and G.J. Martyna, Exploiting multiple levels of parallelism in molecular dynamics based calculations via modern techniques and software paradigms on distributed memory computers, *Computer Physics Communications* **128** (2000), 333–376.
- [38] P. Vashishta, R.K. Kalia and A. Nakano, Large-scale atomistic simulation of dynamic fracture, *Computing in Science & Engineering* **1**(5) (1999), 56–65.
- [39] M.S. Warren and J.K. Salmon, A portable parallel particle program, *Computer Physics Communications* **87** (1995), 266–290.
- [40] D.-Y. Yang, A. Grama and V. Sarin, Bounded-error compression of particle data from hierarchical approximation methods, in: *Proceedings of SC99*, ACM/IEEE, 1999.