

Multimillion Atom Reactive Simulations of Nanostructured Energetic Materials

Priya Vashishta,* Rajiv K. Kalia,* and Aiichiro Nakano*
University of Southern California, Los Angeles, California 90089-0242

and

Barrie E. Homan and Kevin L. McNesby
U.S. Army Research Laboratory, Aberdeen Proving Ground, Maryland 21005

DOI: 10.2514/1.25651

For large-scale atomistic simulations involving chemical reactions to study nanostructured energetic materials, we have designed linear-scaling molecular dynamics algorithms: 1) first-principles-based fast reactive force field molecular dynamics, and 2) embedded divide-and-conquer density functional theory on adaptive multigrids for quantum-mechanical molecular dynamics. These algorithms have achieved unprecedented scales of quantum-mechanically accurate and well validated, chemically reactive atomistic simulations [0.56 billion-atom first principles-based fast reactive force field molecular dynamics and 1.4 million-atom (0.12 trillion grid points) embedded divide-and-conquer density functional theory molecular dynamics] in addition to 18.9 billion-atom nonreactive space-time multiresolution molecular dynamics, with parallel efficiency as high as 0.953 on 1920 Itanium2 processors. These algorithms have enabled us to perform reactive molecular dynamics simulations to reveal various atomistic processes during 1) the oxidation of an aluminum nanoparticle, 2) the decomposition and chemisorption of an RDX (1, 3, 5-trinitro-1, 3, 5-triazine) molecule on an aluminum surface, and 3) shock-initiated detonation of energetic nanocomposite material (RDX crystalline matrix embedded with aluminum nanoparticles.

I. Introduction

PETAFLIPS computers [1] to be built in the near future and associated massive data analysis [2] will offer tremendous opportunities for high-end computer simulations of nanostructured energetic materials (NSEMs). The computing power will enable unprecedented scales of first-principles-based predictive simulations to understand microscopic mechanisms that govern macroscopic materials behavior, thereby enabling rational design of material compositions and microstructures to achieve high specific impulse and insensitivity.

Specifically, we focus on thermomechanical properties and microscopic mechanisms of detonation processes of RDX (1, 3, 5-trinitro-1, 3, 5-triazine, $C_3N_6O_6H_6$) crystalline matrix embedded with aluminum nanoparticles, including shock-induced initiation of detonation, chemically sustained shock wave, and shock-induced flow initiation at the interfaces. Aluminum powders are widely used as propellants, because their combustion products, such as Al_2O_3 , are accompanied by a large amount of heat release. Burn rates of propellants can be accelerated by reducing the size of Al particles, thereby increasing the surface to volume ratio and the rate of chemical reactions. A major technical difficulty for such small reactant particles is the dead weight of oxide layers. The thickness of the oxidized layer in an Al particle is known to be a few nanometers regardless of the particles' size. Therefore, the ratio of the oxidized layer, which is not effective as a propellant, to the reactive portion increases for the smaller Al particles. This dead-weight problem in nanoscale reactant particles may be overcome by encapsulating the particles within complementary reactive materials such as RDX.

Prerequisite to petaflops-scale simulations of such NSEMs is a hierarchy of algorithms that are scalable to 10^5 processors. The

multitude of length and time scales and the associated wide solution space have thus far precluded such first-principles large-scale approaches. Toward achieving first-principles-based atomistic simulations of NSEMs at the petaflops-scale, we have developed an embedded divide-and-conquer (EDC) algorithmic framework to design linear-scaling algorithms for approximately solving hard simulation problems with tight error control.

This paper describes our scalable parallel EDC simulation algorithms to accurately describe the coupling of chemical reactions, atomistic processes, and macroscopic materials phenomena and presents their applications to several NSEMs. In the next section, we describe the EDC algorithmic framework along with results of benchmark tests. Section III describes applications of large-scale chemically reactive molecular dynamics (MD) simulations to three test cases: 1) the oxidation of an aluminum nanoparticle (nAl), 2) the decomposition and chemisorption of an RDX (1, 3, 5-trinitro-1, 3, 5-triazine) molecule on an Al surface, and 3) shock-initiated detonation of energetic nanocomposite material (RDX crystalline matrix embedded with nAl). Finally, Sec. IV contains conclusions.

II. Scalable Parallel Molecular Dynamics Simulation Algorithms

In the past several years, we have developed a unified embedded divide-and-conquer algorithmic framework based on data locality principles to design linear-scaling algorithms for broad scientific applications with tight error control [3,4]. In the EDC algorithms, spatially localized subproblems are solved in a global embedding field, which is efficiently computed with tree-based algorithms. Examples of the embedding field are 1) the electrostatic field in MD simulations [5], 2) the self-consistent Kohn–Sham potential [6] in the density functional theory (DFT) [7] in quantum-mechanical (QM) simulations [4], and 3) coarser but less computer-intensive models in hierarchical simulations that embed high-accuracy but computer-intensive simulations within a coarse simulation only when and where high-fidelity is required [8,9].

We have used the EDC framework to develop a suite of linear-scaling MD algorithms, in which interatomic forces are computed with varying accuracy and complexity. The Appendix describes the three EDC simulation algorithms that are used in our simulations:

Received 5 June 2006; revision received 31 January 2007; accepted for publication 30 March 2007. Copyright © 2007 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0748-4658/07 \$10.00 in correspondence with the CCC.

*Collaboration by Advanced Computing and Simulations, Department of Chemical Engineering and Materials Science, Department of Physics and Astronomy, and Department of Computer Science.

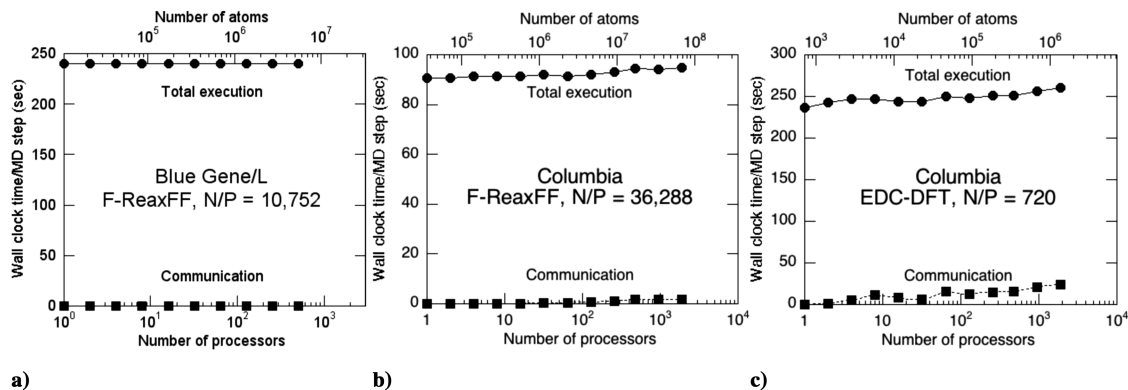


Fig. 1 Total execution and communication times (per MD time step) of EDC MD simulation algorithms as a function of the number of processors P : a) F-ReaxFF MD for scaled workloads: 10, 752 P atom RDX systems on Blue Gene/L; b) F-ReaxFF MD for 36, 288 P atom RDX systems on Columbia; and c) EDC-DFT MD for 720 P atom alumina systems on Columbia.




Fig. 2 Controlled convergence of the potential energy for amorphous CdSe by localization parameters: a) the domain size; b) the buffer length. Numerals are the number of self-consistent iterations for the convergence; and c) total energy conservation in EDC-DFT-based MD simulation of liquid Rb (solid curve) compared with the time variation of the potential energy (dashed curve).

1) Algorithm MRMD: space-time multiresolution MD to solve the formally $O(N^2)$ N -body problem in classical MD in $O(N)$ computing time [10].

2) Algorithm F-ReaxFF: first-principles-based fast reactive force field to solve the $O(N^3)$ variable N -charge problem in semiclassical reactive force field (ReaxFF) MD in $O(N)$ time [3].

3) Algorithm EDC-DFT: embedded divide-and-conquer density functional theory to approximately solve the exponentially complex quantum N -body problem in quantum-mechanical MD in $O(N)$ time [3,4].

The EDC algorithms are portable and have been run on various high-end parallel supercomputers such as the Intel Itanium2-based NASA Columbia and IBM Blue Gene/L (Fig. 1). The EDC algorithms expose maximal data locality and thus achieve high parallel efficiency. For example, the F-ReaxFF algorithm on Columbia has achieved the parallel efficiency 0.953[†] on 1920 Itanium2 processors and over 0.999 on 512 Blue Gene/L processors.[‡]

The largest benchmark tests of the EDC simulation algorithms on Columbia include 18,925,056,000-atom MRMD, 557,383,680-atom F-ReaxFF, and 1,382,400-atom (121,385,779,200 electronic degrees of freedom) EDC-DFT calculations [3]. Careful analysis of the benchmark test results shows perfect linear scaling for all three algorithms.

The EDC algorithms have a well-defined set of localization parameters, which control the computational cost and the accuracy.

[†]Columbia is a cluster of 20 Altix boxes, each with 512 processors, and the interbox parallel efficiency with the NUMalink4 interconnect, from 480 processors in one box to 1920 processors in four boxes, is 0.995.

[‡]More recent benchmark results include 134 billion-atom MRMD, 1.06 billion-atom F-ReaxFF MD, and 11.8 million-atom (1.04 trillion grid points) EDC-DFT MD simulations, with the parallel efficiency as high as 0.998 on 131,072 Blue Gene/L processors.

Figures 2a and 2b show the rapid convergence of the EDC-DFT energy as a function of its localization parameters (the size of a domain and the length of a buffer layer that augments each domain to avoid artificial boundary effects). The EDC-DFT MD algorithm has also overcome the energy drift problem [11], which plagues most $O(N)$ DFT-based MD algorithms, especially with large basis sets ($>10^4$ unknowns per electron, necessary for the transferability of accuracy) (Fig. 2c) [4]. It also has robust convergence properties in contrast to many $O(N)$ DFT-based MD algorithms that have runaway solutions and fail to converge for some problems [11].

III. Reactive MD Simulations of High-Energy Materials

A. Oxidation of an Aluminum Nanoparticle

Oxidation plays a critical role in the burning of energetic materials. We have performed the first successful MD simulation of oxidation of an Al nanoparticle (diameter 200 Å) [12,13]. The reactive MD simulation is based on the interaction scheme developed by Streitz and Mintmire [14], which can successfully describe a wide range of physical properties of both metallic and ceramic systems.[§] This scheme is capable of treating bond formation and bond breakage and changes in charge transfer as the atoms move and their local environments are altered.

In our microcanonical MD simulation, energy released from Al-O bond formation is rapidly transported into the nanocluster resulting in disordering of the Al nanocrystal and outward expansion of the oxide region (see Fig. 3). The thickness of the oxide region increases linearly with time and does not saturate. By 50 ps the thickness and temperature of the oxide region are 35 Å and 2500 K, respectively.

[§]The form of the Streitz–Mintmire interatomic potential is a subset of the ReaxFF described in the Appendix.

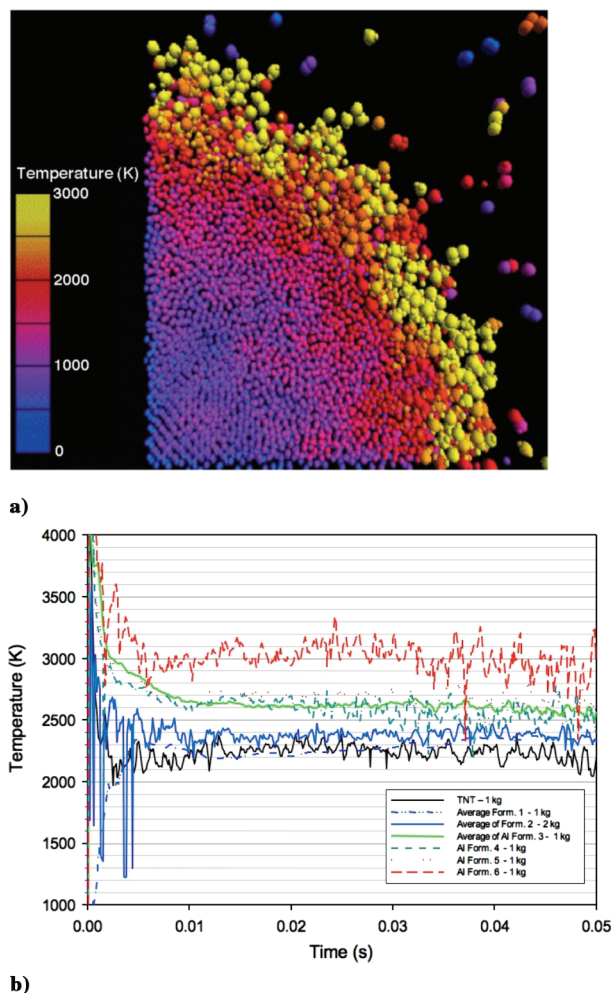


Fig. 3 Temperature during combustion of energetic materials: a) Temperature of atoms during the microcanonical MD simulation of Al nanoparticle in an oxygen environment. The larger spheres correspond to oxygen and smaller spheres to aluminum (color represents the temperature); b) Experimental pyrometry data, showing the temperature vs time during the combustion of energetic materials of various formulations and weights.

Subsequently, numerous small Al_xO_y fragments are ejected from the nanocluster surface, indicating that the nanocluster is exploding. This behavior under closed conditions has also been observed experimentally.

Such time-resolved temperature data in reactive MD simulations provide indispensable insight into combustion mechanisms underlying, e.g., experimental pyrometry data by real-time optical measurements of the combustion of high-energy density materials (Fig. 3). The pyrometry data were obtained using a three-color method based on Planck's law of blackbody radiation [15], employing a fiber-coupled pyrometer built inhouse at the U.S. Army Research Laboratory, with temporal resolution $\sim 11 \mu\text{s}$. The temperatures in Fig. 3 are of the surface of the expanding fire ball.

In our simulations, local stresses in the oxide scale cause rapid diffusion of aluminum and oxygen atoms. The local stresses are examined by separating the contributions from the electrostatic and nonelectrostatic forces. We find that the attractive Coulomb forces between aluminum and oxygen contribute a large negative pressure localized in the oxide. The electrostatic pressure contribution increases in magnitude toward the middle of the oxide where charge transfer is the highest. The large attractive forces are partially offset by steric repulsion, which gives rise to a positive nonelectrostatic contribution to the local pressure in the oxide. Analyses of the local stresses reveal large stress gradients throughout the nanocluster with the oxide largely under negative pressure and the metal core under positive pressure. Local pressures range between -1 and 1 GPa.




Fig. 4 Snapshots of DFT-based MD simulation on the decomposition of an RDX molecule on an Al (111) surface.

B. Decomposition of an RDX Molecule on an Aluminum Surface

We have also performed an MD simulation to study the decomposition of an RDX molecule on an Al (111) surface, in which interatomic forces are computed quantum-mechanically according to the Hellmann–Feynman theorem in the framework of DFT (Fig. 4). The DFT calculation is based on the norm-conserving pseudopotentials by Troullier and Martins [16] and the generalized gradient correction by Perdew et al. [17], and it is implemented on a real-space grid and accelerated with the multigrid method [4]. The simulation shows the decomposition of NO_2 fragments in RDX and subsequent formation of Al–O bonds on the Al surface. This suggests that, to stop the premature reaction of RDX with a pure Al surface, a nanoscale buffer layer may be necessary on the Al surface.

Our simulation reveals drastically different RDX decomposition pathways due to the presence of the Al surface. Strong attractive forces between oxygen and aluminum atoms break both N–O and N–N bonds in the RDX and, subsequently, the dissociated oxygen atoms and NO molecules oxidize the Al surface. In addition to these Al surface-assisted decompositions, ring cleavage of the RDX molecule is also observed. These reactions occur spontaneously without potential barriers, and result in the attachment of the rest of the RDX molecule to the surface.

C. Shock-Induced Detonation of RDX/Aluminum Nanoparticles Composite

On the basis of the preceding simulations of Al nanoparticles (nAl) and RDX–Al systems, we have recently performed 1.01 million-atom F-ReaxFF MD simulation to study shock-initiated detonation of RDX crystal/oxidized nAl composite. In the simulation, two slabs of the RDX/nAl composite, each of size $482 \times 353 \times 65 \text{ \AA}^3$ in the x , y , and z directions, are impacted with the impact velocity of 5 km/s in the z direction. Each oxidized nAl consists of 707 atoms. The simulation reveals atomistic processes of shock compression and subsequent explosive reaction. Strong attractive forces between oxygen and aluminum atoms break N–O and N–N bonds in the RDX and, subsequently, the dissociated oxygen atoms and NO molecules oxidize Al, which has also been observed in our DFT-based MD simulation.

The F-ReaxFF MD method has been validated by comparing calculated shock wave velocities in RDX with experimental data, where a shock wave is generated by a planar impactor. Figure 5 compares MD and experimental results on the shock velocity as a function of the particle velocity that drives the shock [18]. The MD and experimental data agree very well. Furthermore, the simulation shows a sudden increase of the number of molecular products such as HONO, N_2 , and OH above a shock velocity $\sim 9 \text{ km/s}$, which is consistent with an experimental detonation velocity [19].

We have found that the large simulation system size is essential for quantitative study of shock velocity and shock structure, because it takes time for a steady shock-front structure to be established. Our simulation reveals a dynamic transition of the shock structure (from a diffuse shock front with well-ordered molecular dipoles behind it to a disordered dipole distribution behind a sharp front) as the shock velocity increases.




Fig. 5 F-ReaxFF MD and experimental data on the planar shock velocity in RDX as a function of the particle velocity.

IV. Conclusions

In the past several years, significant progress has been made in simulation methods, linear-scaling algorithms, and scalable parallel computing technologies to enable first-principles-based predictive simulations to study macroscopic properties of nanostructured energetic materials with atomistic descriptions of chemical reactions. In addition to the number of atoms in such simulations, an important issue is the time scale studied by MD simulations. We define the spatiotemporal scale NT of an MD simulation as the product of the number of atoms N and the simulated time T . Our large MD simulations of ceramics simulate subbillions of atoms for subnanoseconds, resulting in $NT = 0.03\text{--}0.04$ atoms \cdot s [20,21]. Petaflops computers are expected to push the spatiotemporal envelope to $NT \sim 1$ and beyond, thereby bringing in further new scientific knowledge on energetic materials.

Appendix: Embedded Divide-and-Conquer Simulation Algorithms

I. Algorithm MRMD: Space-Time Multiresolution Molecular Dynamics

The MRMD algorithm is used as a template for developing broad particle and continuum simulation algorithms. The MD approach follows the time evolution of the positions, $\mathbf{r}^N = \{\mathbf{r}_i | i = 1, \dots, N\}$, of N atoms by solving coupled ordinary differential equations [10]. Atomic force law is mathematically encoded in the interatomic potential energy $E_{\text{MD}}(\mathbf{r}^N)$, which is often an analytic function $E_{\text{MD}}(\{\mathbf{r}_{ij}\}, \{\mathbf{r}_{ijk}\})$ of atomic pair \mathbf{r}_{ij} and triplet \mathbf{r}_{ijk} positions. For the long-range electrostatic interaction, we use the fast multipole method (FMM) to reduce the $O(N^2)$ computational complexity of the N -body problem to $O(N)$ [22]. In the FMM, the physical system is recursively divided into subsystems to form an octree data structure, and the electrostatic field is computed recursively on the octree with $O(N)$ operations, while maintaining spatial locality at each recursion level. Our scalable parallel implementation of the FMM has a unique feature to compute atomistic stress tensor components based on a complex charge method [23]. MRMD also uses temporal locality through multiple time stepping, which uses different force-update schedules for different force components [10,24,25]. Specifically, forces from neighbor atoms are computed at every MD step, whereas forces from farther atoms are updated less frequently. For parallelization, we use spatial decomposition. The total volume is divided into P subsystems of equal volume, and each subsystem is assigned to a node in an array of P compute nodes. 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 nodes. After updating the atom positions due to time stepping, some atoms may have moved out of its subsystem.

These atoms are “migrated” to the proper neighbor nodes. With spatial decomposition, the computation scales as N/P , whereas communication scales as $(N/P)^{2/3}$. The FMM incurs an $O(\log P)$ overhead, which is negligible for coarse-grained ($N/P \gg P$) applications.

II. Algorithm F-ReaxFF: Fast Reactive Force Field Molecular Dynamics

In the past five years, chemists have developed a first-principles-based reactive force field (ReaxFF) approach to significantly reduce the computational cost of simulating chemical reactions [18,26]. However, its parallelization has seen only limited success, with the previously largest ReaxFF MD involving $N < 10^4$ atoms. We have developed F-ReaxFF to enable ReaxFF MD involving 10^9 atoms [3,27]. The variable N -charge problem in ReaxFF amounts to solving a dense linear system of equations to determine atomic charges $\{q_i | i = 1, \dots, N\}$ at every MD step [12,28]. F-ReaxFF reduces its $O(N^3)$ complexity to $O(N)$ by combining the FMM based on spatial locality and iterative minimization to use the temporal locality of the solution. To accelerate the convergence, we use a multilevel preconditioned conjugate-gradient (MPCG) method that splits the Coulomb-interaction matrix into short- and long-range parts and uses the sparse short-range matrix as a preconditioner [29]. The extensive use of the sparse preconditioner enhances the data locality and thereby improves the parallel efficiency. The chemical bond order B_{ij} is an attribute of atomic pair (i, j) and changes dynamically, adapting to the local environment. In ReaxFF, the potential energy $E_{\text{ReaxFF}}(\{\mathbf{r}_{ij}\}, \{\mathbf{r}_{ijk}\}, \{\mathbf{r}_{ijkl}\}, \{q_i\}, \{B_{ij}\})$ between atomic pairs \mathbf{r}_{ij} , triplets \mathbf{r}_{ijk} , and quadruplets \mathbf{r}_{ijkl} depends on the bond orders of all constituent atomic pairs. Force calculations in ReaxFF thus involve up to atomic 6-tuples due to chain-rule differentiations through B_{ij} . To efficiently handle the multiple interaction ranges, the parallel F-ReaxFF algorithm employs a multilayer cellular decomposition scheme for caching atomic n -tuples ($n = 2\text{--}6$) [3].

III. Algorithm EDC-DFT: Embedded Divide-and-Conquer Density Functional Theory on Adaptive Multigrids for Quantum-Mechanical Molecular Dynamics

The EDC-DFT algorithm describes chemical reactions with a higher quantum-mechanical accuracy than ReaxFF. The DFT problem is formulated as a minimization of the energy functional $E_{\text{QM}}(\mathbf{r}^N, \psi^{N_{\text{el}}})$ with respect to electronic wave functions (or Kohn–Sham orbitals) $\psi^{N_{\text{el}}}(\mathbf{r}) = \{\psi_n(\mathbf{r}) | n = 1, \dots, N_{\text{el}}\}$, subject to orthonormality constraints (N_{el} is the number of wave functions on the order of N) [7]. The data locality principle called quantum nearsightedness [30] in DFT is best implemented with a divide-and-conquer algorithm [31,32], which naturally leads to $O(N)$ DFT calculations [11]. However, it is only in the past several years that $O(N)$ DFT algorithms, especially with large basis sets ($> 10^4$ unknowns per electron, necessary for the transferability of accuracy), have attained controlled error bounds, robust convergence properties, and energy conservation during MD simulations, to make large DFT-based MD simulations practical [4,33]. We have designed an embedded divide-and-conquer density functional theory (EDC-DFT) algorithm, in which a hierarchical grid technique combines multigrid preconditioning and adaptive fine mesh generation [4]. The EDC-DFT algorithm represents the physical system as a union of overlapping spatial domains, $\Omega = \cup_{\alpha} \Omega_{\alpha}$, and physical properties are computed as linear combinations of domain properties. For example, the electronic density is expressed as

$$\rho(\mathbf{r}) = \sum_{\alpha} p^{\alpha}(\mathbf{r}) \sum_n f_n^{\alpha} |\psi_n^{\alpha}(\mathbf{r})|^2$$

where $p^{\alpha}(\mathbf{r})$ is a support function that vanishes outside the α th domain Ω_{α} , and f_n^{α} and $\psi_n^{\alpha}(\mathbf{r})$ are the occupation number and the wave function of the n th Kohn–Sham orbital in Ω_{α} . The domains are embedded in a global Kohn–Sham potential, which is a functional of $\rho(\mathbf{r})$ and is determined self-consistently with $\{f_n^{\alpha}, \psi_n^{\alpha}(\mathbf{r})\}$. We use the

multigrid method to compute the global potential in $O(N)$ time. The DFT calculation in each domain is performed using a real-space approach [34], in which electronic wave functions are represented on grid points. The real-space grid is augmented with coarser multigrids to accelerate the iterative solution. Furthermore, a finer grid is adaptively generated near every atom to accurately operate ionic pseudopotentials for calculating electron-ion interactions. The EDC-DFT algorithm on the hierarchical real-space grids is implemented on parallel computers based on spatial decomposition. Each compute node contains one or more domains of the EDC algorithm. Then only the global density, but not individual wave functions, needs to be communicated. The resulting large computation/communication ratio makes this approach highly scalable.

Acknowledgments

This work was partially supported by the Army Research Office–Multidisciplinary University Research Initiative, U.S. Air Force Office of Scientific Research–Defense University Research Initiative on Nanotechnology, U.S. Department of Energy, Defense Threat Reduction Agency, and National Science Foundation. Benchmark tests were performed using the NASA Columbia supercomputers at the NASA Ames Research Center. The simulations were carried out in collaboration with Gurcan Aral, Rupak Biswas, Tomothy J. Campbell, Adri van Duin, William A. Goddard III, Ken-ichi Nomura, Fuyuki Shimojo, and Deepak Srivastava. We thank the following collaborators for valuable discussions: Richard Yetter, David Allala, Kenneth Kuo, Vigor Yang, Dana Dlott, Ralph Nuzzo, Brad Forch, Shashi Karna, Betsy Rice, Margaret Hurley, and William Wilson.

References

- [1] Dongarra, J., and Walker, D., "Quest for Petascale Computing," *Computing in Science and Engineering*, Vol. 3, No. 3, 2001, pp. 32–39.
- [2] Sharma, A., Nakano, A., Kalia, R. K., Vashishta, P., Kodiyalam, S., Miller, P., Zhao, P., Liu, X., Campbell, T. J., and Haas, A., "Immersive and Interactive Exploration of Billion-Atom Systems," *Presence: Teleoperators and Virtual Environments*, Vol. 12, No. 1, 2003, pp. 85–95.
- [3] Nakano, A., Kalia, R. K., Nomura, K., Sharma, A., Vashishta, P., Shimojo, F., van Duin, A. C. T., Goddard, W. A., Biswas, R., and Srivastava, D., "Divide-and-Conquer/Cellular-Decomposition Framework for Million-to-Billion Atom Simulations of Chemical Reactions," *Computational Materials Science*, Vol. 38, No. 4, 2007, pp. 642–652.
- [4] Shimojo, F., Kalia, R. K., Nakano, A., and Vashishta, P., "Embedded Divide-and-Conquer Algorithm on Hierarchical Real-Space Grids: Parallel Molecular Dynamics Simulation Based on Linear-Scaling Density Functional Theory," *Computer Physics Communications*, Vol. 167, No. 3, 2005, pp. 151–164.
- [5] Nakano, A., Kalia, R. K., Vashishta, P., Campbell, T. J., Ogata, S., Shimojo, F., and Saini, S., "Scalable Atomistic Simulation Algorithms for Materials Research," *Proceedings of Supercomputing 2001*, Association for Computing Machinery, New York, 2001.
- [6] Kohn, W., and Sham, L. J., "Self-Consistent Equations Including Exchange and Correlation Effects," *Physical Review*, Vol. 140 A, 1965, pp. 1133–1138.
- [7] Hohenberg, P., and Kohn, W., "Inhomogeneous Electron Gas," *Physical Review*, Vol. 136 B, 1964, pp. 864–871.
- [8] Ogata, S., Lidorikis, E., Shimojo, F., Nakano, A., Vashishta, P., and Kalia, R. K., "Hybrid Finite-Element/Molecular-Dynamics/Electronic-Density-Functional Approach to Materials Simulations on Parallel Computers," *Computer Physics Communications*, Vol. 138, No. 2, 2001, pp. 143–154.
- [9] Ogata, S., Shimojo, F., Kalia, R. K., Nakano, A., and Vashishta, P., "Environmental Effects of H₂O on Fracture Initiation in Silicon: A Hybrid Electronic-Density-Functional/Molecular-Dynamics Study," *Journal of Applied Physics*, Vol. 95, No. 10, 2004, pp. 5316–5323.
- [10] Nakano, A., Kalia, R. K., Vashishta, P., Campbell, T. J., Ogata, S., Shimojo, F., and Saini, S., "Scalable Atomistic Simulation Algorithms for Materials Research," *Scientific Programming*, Vol. 10, No. 4, 2002, pp. 263–270.
- [11] Goedecker, S., "Linear Scaling Electronic Structure Methods," *Reviews of Modern Physics*, Vol. 71, No. 4, 1999, pp. 1085–1123.
- [12] Campbell, T., Kalia, R. K., Nakano, A., Vashishta, P., Ogata, S., and Rodgers, S., "Dynamics of Oxidation of Aluminum Nanoclusters Using Variable Charge Molecular-Dynamics Simulations on Parallel Computers," *Physical Review Letters*, Vol. 82, No. 24, 1999, pp. 4866–4869.
- [13] Campbell, T. J., Aral, G., Ogata, S., Kalia, R. K., Nakano, A., and Vashishta, P., "Oxidation of Aluminum Nanoclusters," *Physical Review B*, Vol. 71, No. 20, 2005, pp. 205413: 1–14.
- [14] Streitz, F. H., and Mintmire, J. W., "Electrostatic Potentials for Metal-Oxide Surfaces and Interfaces," *Physical Review B*, Vol. 50, No. 16, 1994, pp. 11,996–12,003.
- [15] Felice, R. A., "Spectropymeter: a Practical Multi-Wavelength Pyrometer," *AIP Conference Proceedings*, Vol. 684, 2003, pp. 711–716.
- [16] Troullier, N., and Martins, J. L., "Efficient Pseudopotentials for Plane-Wave Calculations 2: Operators for Fast Iterative Diagonalization," *Physical Review B*, Vol. 43, No. 11, 1991, pp. 8861–8869.
- [17] Perdew, J. P., Burke, K., and Ernzerhof, M., "Generalized Gradient Approximation Made Simple," *Physical Review Letters*, Vol. 77, No. 18, 1996, pp. 3865–3868.
- [18] Strachan, A., van Duin, A. C. T., Chakraborty, D., Dasgupta, S., and Goddard, W. A., "Shock Waves in High-Energy Materials: the Initial Chemical Events in Nitramine RDX," *Physical Review Letters*, Vol. 91, No. 9, 2003, pp. 098301: 1–4.
- [19] Cowan, R. D., and Fickett, W., "Calculation of the Detonation Properties of Solid Explosives with the Kistiakowsky–Wilson Equation of State," *Journal of Chemical Physics*, Vol. 24, No. 5, 1956, pp. 932–939.
- [20] Branicio, P. S., Kalia, R. K., Nakano, A., and Vashishta, P., "Shock-Induced Structural Phase Transition, Plasticity, and Brittle Cracks in Aluminum Nitride Ceramic," *Physical Review Letters*, Vol. 96, No. 6, 2006, pp. 065502: 1–4.
- [21] Szlufarska, I., Nakano, A., and Vashishta, P., "Crossover in the Mechanical Response of Nanocrystalline Ceramics," *Science*, Vol. 309, No. 5736, 2005, pp. 911–914.
- [22] Greengard, L., and Rokhlin, V., "Fast Algorithm for Particle Simulations," *Journal of Computational Physics*, Vol. 73, No. 2, 1987, pp. 325–348.
- [23] Ogata, S., Campbell, T. J., Kalia, R. K., Nakano, A., Vashishta, P., and Vemparala, S., "Scalable and Portable Implementation of the Fast Multipole Method on Parallel Computers," *Computer Physics Communications*, Vol. 153, No. 3, 2003, pp. 445–461.
- [24] Martyna, G. J., Tuckerman, M. E., Tobias, D. J., and Klein, M. L., "Explicit Reversible Integrators for Extended Systems Dynamics," *Molecular Physics*, Vol. 87, No. 5, 1996, pp. 1117–1157.
- [25] Schlick, T., Skeel, R. D., Brunger, A. T., Kale, L. V., Board, J. A., Hermans, J., and Schulten, K., "Algorithmic Challenges in Computational Molecular Biophysics," *Journal of Computational Physics*, Vol. 151, No. 1, 1999, pp. 9–48.
- [26] van Duin, A. C. T., Dasgupta, S., Lorant, F., and Goddard, W. A., "ReaxFF: a Reactive Force Field for Hydrocarbons," *Journal of Physical Chemistry A*, Vol. 105, No. 41, 2001, pp. 9396–9409.
- [27] Vashishta, P., Kalia, R. K., and Nakano, A., "Multimillion Atom Simulations of Dynamics of Oxidation of an Aluminum Nanoparticle and Nanoindentation on Ceramics," *Journal of Physical Chemistry B*, Vol. 110, No. 8, 2006, pp. 3727–3733.
- [28] Rappe, A. K., and Goddard, W. A., "Charge Equilibration for Molecular-Dynamics Simulations," *Journal of Physical Chemistry*, Vol. 95, No. 8, 1991, pp. 3358–3363.
- [29] Nakano, A., "Parallel Multilevel Preconditioned Conjugate-Gradient Approach to Variable-Charge Molecular Dynamics," *Computer Physics Communications*, Vol. 104, Nos. 1–3, 1997, pp. 59–69.
- [30] Kohn, W., "Density Functional and Density Matrix Method Scaling Linearly with the Number of Atoms," *Physical Review Letters*, Vol. 76, No. 17, 1996, pp. 3168–3171.
- [31] Yang, W., "Direct Calculation of Electron Density in Density-Functional Theory," *Physical Review Letters*, Vol. 66, No. 11, 1991, pp. 1438–1441.
- [32] Yang, W., and Lee, T.-S., "Density-Matrix Divide-and-Conquer Approach for Electronic Structure Calculations of Large Molecules," *Journal of Chemical Physics*, Vol. 103, No. 13, 1995, pp. 5674–5678.
- [33] Fattebert, J.-L., and Gygi, F., "Linear Scaling First-Principles Molecular Dynamics with Controlled Accuracy," *Computer Physics Communications*, Vol. 162, No. 1, 2004, pp. 24–36.
- [34] Chelikowsky, J. R., Saad, Y., Ögüt, S., Vasilev, I., and Stathopoulos, A., "Electronic Structure Methods for Predicting the Properties of Materials: Grids in Space," *Physica Status Solidi B: Basic Research*, Vol. 217, No. 1, 2000, pp. 173–195.