

A divide-and-conquer/cellular-decomposition framework for million-to-billion atom simulations of chemical reactions

Aiichiro Nakano ^{a,*}, Rajiv K. Kalia ^a, Ken-ichi Nomura ^a, Ashish Sharma ^a,
Priya Vashishta ^a, Fuyuki Shimojo ^{a,b}, Adri C.T. van Duin ^c,
William A. Goddard ^c, Rupak Biswas ^d, Deepak Srivastava ^d

^a Collaboratory for Advanced Computing and Simulations, Department of Computer Science, Department of Physics & Astronomy, Department of Chemical Engineering & Materials Science, University of Southern California, Los Angeles, CA 90089-0242, USA

^b Department of Physics, Kumamoto University, Kumamoto 860-8555, Japan

^c Materials and Process Simulation Center, Division of Chemistry and Chemical Engineering, California Institute of Technology, Pasadena, CA 91125, USA

^d NASA Advanced Supercomputing (NAS) Division, NASA Ames Research Center, Moffett Field, CA 94035, USA

Received 22 February 2006; accepted 21 April 2006

Abstract

To enable large-scale atomistic simulations of material processes involving chemical reactions, we have designed linear-scaling molecular dynamics (MD) algorithms based on an embedded divide-and-conquer (EDC) framework: first principles-based fast reactive force-field (F-ReaxFF) MD; and quantum-mechanical MD in the framework of the density functional theory (DFT) on adaptive multigrids. To map these $O(N)$ algorithms onto parallel computers with deep memory hierarchies, we have developed a tunable hierarchical cellular-decomposition (THCD) framework, which achieves performance tunability through a hierarchy of parameterized cell data/computation structures and adaptive load balancing through wavelet-based computational-space decomposition. Benchmark tests on 1920 Itanium2 processors of the NASA Columbia supercomputer have achieved unprecedented scales of quantum-mechanically accurate and well validated, chemically reactive atomistic simulations—0.56 billion-atom F-ReaxFF MD and 1.4 million-atom (0.12 trillion grid points) EDC–DFT MD—in addition to 18.9 billion-atom non reactive space–time multiresolution MD. The EDC and THCD frameworks expose maximal data localities, and consequently the isogranular parallel efficiency on 1920 processors is as high as 0.953. Chemically reactive MD simulations have been applied to shock-initiated detonation of energetic materials and stress-induced bond breaking in ceramics in corrosive environments.

© 2006 Elsevier B.V. All rights reserved.

PACS: 02.70.–c; 02.70.Ns; 71.15.–m

Keywords: Molecular dynamics; Reactive force field; Quantum mechanics; Density functional theory; Parallel computing

1. Introduction

There is growing interest in large-scale molecular dynamics (MD) simulations [1–4] involving million-to-billion atoms, in which interatomic forces are computed quantum mechanically [5] to accurately describe chemical

reactions. Such large reactive MD simulations would for the first time provide requisite coupling of chemical reactions, atomistic processes, and macroscopic materials phenomena, to solve a wide spectrum of problems of great societal impact. Examples of technological significance include: stress corrosion cracking (corrosion-related direct costs make up 3% of the gross domestic product in the US), where chemical reactions at the crack tip are inseparable from long-range stress fields [6]; energetic nanomaterials to boost the impulse of rocket fuels, in which

* Corresponding author. Tel.: +1 213 821 2657; fax: +1 213 821 2664.
E-mail address: anakano@usc.edu (A. Nakano).

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 APR 2006	2. REPORT TYPE	3. DATES COVERED 00-00-2006 to 00-00-2006	
4. TITLE AND SUBTITLE A divide-and-conquer/cellular-decomposition framework for million-to-billion atom simulations of chemical reactions		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) California Institute of Technology, Materials and Process Simulation Center, MC 139-74, Pasadena, CA, 91125		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			
14. ABSTRACT To enable large-scale atomistic simulations of material processes involving chemical reactions, we have designed linear-scaling molecular dynamics (MD) algorithms based on an embedded divide-and-conquer (EDC) framework: first principles-based fast reactive force-field (F-ReaxFF) MD; and quantum-mechanical MD in the framework of the density functional theory (DFT) on adaptive multigrids. To map these O(N) algorithms onto parallel computers with deep memory hierarchies, we have developed a tunable hierarchical cellular-decomposition (THCD) framework, which achieves performance tunability through a hierarchy of parameterized cell data/computation structures and adaptive load balancing through wavelet-based computational-space decomposition. Benchmark tests on 1920 Itanium2 processors of the NASA Columbia supercomputer have achieved unprecedented scales of quantum-mechanically accurate and well validated, chemically reactive atomistic simulations?0.56 billion-atom F-ReaxFF MD and 1.4 million-atom (0.12 trillion grid points) EDC?DFT MD?in addition to 18.9 billion-atom non reactive space?time multiresolution MD. The EDC and THCD frameworks expose maximal data localities, and consequently the isogranular parallel efficiency on 1920 processors is as high as 0.953. Chemically reactive MD simulations have been applied to shock-initiated detonation of energetic materials and stress-induced bond breaking in ceramics in corrosive environments.			
15. SUBJECT TERMS			
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	
			18. NUMBER OF PAGES 11
			19a. NAME OF RESPONSIBLE PERSON

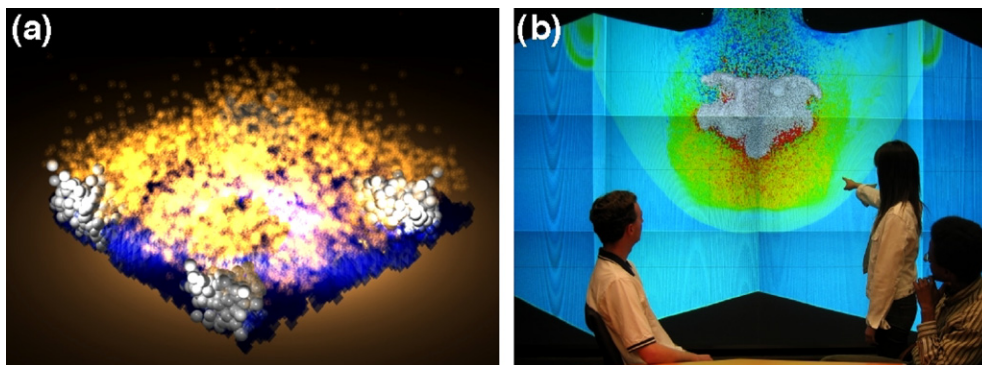


Fig. 1. (a) Reactive force-field molecular dynamics simulation of shock-initiated combustion of an energetic nanocomposite material (nitramine matrix [7] embedded with aluminum nanoparticles [9]). (b) Rendering of a molecular dynamics simulation on a tiled display at USC, showing hypervelocity impact damage of a ceramic plate with impact velocity 15 km/s, where one quarter of the system is cut to show the internal pressure distribution (the projectile is shown in white). Such simulations help design thermal and radiation protection systems of aerospace vehicles, which are tolerant to micrometeorite impacts (where impact speeds are as high as 40 km/s).

chemical reactions sustain shock waves (see Fig. 1(a)) [7]; and micrometeorite impact damages to the thermal and radiation protection layers of aerospace vehicles, understanding of which is essential for safer space flights (Fig. 1(b)). Emerging petaflops computers could potentially extend the realm of quantum-mechanical (QM) simulation [8] to the macroscopic scales, but only if scalable parallel simulation technologies were developed.

In the past few years, several promising approaches have emerged toward achieving million-to-billion atom simulations of chemical reactions. One computational approach toward QM-based million-atom MD simulations is to perform a number of small density functional theory (DFT) [10,11] calculations “on the fly” to compute interatomic forces quantum mechanically. We have recently designed a linear-scaling divide-and-conquer DFT algorithm on adaptive multigrids, which achieves robust convergence, controlled errors, and energy conservation during MD simulations [12]. Here we present the first million-atom DFT-based MD simulation, where electronic wave functions are represented on 10^{11} grid points. Alternative to this concurrent DFT-based MD approach is a sequential DFT-informed strategy, which employs environment-dependent interatomic potentials based on: (1) variable atomic charges to describe charge transfers; and (2) reactive bond orders to describe chemical bond formation and breakage. In our first principles-based reactive force-field (ReaxFF) approach [7,13], the parameters in the interatomic potential are “trained” to best fit thousands of DFT calculations on small (the number of atoms, $N \sim 10$) clusters of various atomic-species combinations. This paper presents a new $O(N)$ parallel ReaxFF algorithm, which for the first time enables billion-atom MD simulations of chemical reactions.

The first contribution of this paper is a unified embedded divide-and-conquer (EDC) algorithmic framework for designing linear-scaling parallel algorithms for broad scientific and engineering problems, with specific applications to two reactive atomistic simulation methods: ReaxFF MD and DFT-based MD. Mapping these $O(N)$

algorithms onto multi-teraflops to petaflops parallel computers, however, poses a number of challenges, e.g., achieving high scalability for irregularly distributed billion atoms, and enabling performance portability [14] for a wide range of parallel architectures. To overcome these challenges, the second contribution of this paper is a tunable hierarchical cellular-decomposition (THCD) framework, which is aware of deep memory hierarchies, by maximally exposing data locality and exploiting parallelism at each decomposition level. The framework features topology-preserving computational-space decomposition and wavelet-based adaptive load balancing. To ensure performance portability, a hierarchy of cell data/computational structures are parameterized and tuned on each platform. The major accomplishment of this paper is the unprecedented scales of quantum-mechanically accurate and well validated, chemically reactive atomistic simulations on the NASA Columbia supercomputer. Benchmark tests on 1920 Itanium2 processors have achieved 0.56 billion-atom ReaxFF and 1.4 million-atom (0.12 trillion degrees of freedom) DFT simulations, in addition to 18.9 billion-atom MD simulation, with isogranular parallel efficiency as high as 0.953.

This paper is organized as follows. In the next section, we describe the EDC algorithmic framework. Section 3 discusses the THCD parallel computing framework. Results of benchmark tests are given in Section 4, and Section 5 describes applications of large-scale chemically reactive MD simulations to shock-initiated detonation of energetic materials and stress-induced bond breaking in ceramics in corrosive environments. Finally, Section 6 contains conclusions.

2. Linear-scaling embedded divide-and-conquer simulation algorithms

We have developed a unified algorithmic framework to design linear-scaling algorithms for broad scientific and engineering applications, based on data locality principles. In the *embedded divide-and-conquer (EDC) algorithms*,

spatially localized subproblems are solved in a global embedding field, which is efficiently computed with tree-based algorithms (see Fig. 2). Examples of the embedding field are the electrostatic field in molecular dynamics (MD) simulations and the self-consistent Kohn–Sham potential in the density functional theory (DFT).

Specifically, we have used the EDC framework to develop a suite of linear-scaling MD simulation algorithms for materials research, in which interatomic forces are computed with varying accuracy and complexity. The linear-scaling algorithms encompass a wide spectrum of physical reality: (1) classical MD based on a many-body interatomic potential model, which involves the formally $O(N^2)$ N -body problem; (2) environment-dependent, reactive force-field (ReaxFF) MD, which involves the $O(N^3)$ variable N -charge problem; and (3) quantum-mechanical (QM) calculation based on the DFT, to provide approximate solutions to the exponentially complex quantum N -body problem.

2.1. Space–time multiresolution molecular dynamics algorithm

We have designed chemically reactive $O(N)$ MD simulation algorithms on the basis of their nonreactive predecessor, *space–time multiresolution molecular dynamics (MRMD)* algorithm [3]. In the MD approach, one obtains the phase-space trajectories of the system (positions and velocities of all atoms at all time). Atomic force law for describing how atoms interact with each other is mathematically encoded in the interatomic potential energy, $E_{\text{MD}}(\mathbf{r}^N)$, which is a function of the positions of all N atoms, $\mathbf{r}^N = \{\mathbf{r}_i | i = 1, \dots, N\}$, in the system. In our many-body interatomic potential scheme, the energy, $E_{\text{MD}}(\{\mathbf{r}_{ij}\}, \{\mathbf{r}_{ijk}\})$, is an analytic function that depends on relative positions of atomic pairs, \mathbf{r}_{ij} , and triplets, \mathbf{r}_{ijk} . 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 is made $O(N)$ using a linked-list cell approach [3]. For the long-range electrostatic interaction, we use the fast multipole method (FMM) [15–17] to reduce the $O(N^2)$ computational com-

plexity of the N -body problem to $O(N)$. 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 the 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 [17]. The MRMD algorithm also utilizes temporal locality through multiple time stepping (MTS), which uses different force-update schedules for different force components [16,18–20]. Specifically, forces from the nearest-neighbor atoms are computed at every MD step, whereas forces from farther atoms are updated less frequently.

For parallelization of MD simulations, we use spatial decomposition [3]. The total volume of the system 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 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 nodes. 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. Tree-based algorithms such as the FMM incur an $O(\log P)$ overhead, which is negligible for coarse-grained ($N/P \gg P$) applications.

2.2. Fast reactive force-field molecular dynamics algorithm

The density functional theory (DFT) [10,11] has reduced the exponentially complex quantum-mechanical problem to $O(M^3)$, by solving M one-electron problems self-consistently instead of one M -electron problem. Unfortunately, DFT-based MD simulations [5] are rarely performed over $N \sim 10^2$ atoms because of its $O(N^3)$ computational complexity, which severely limits their scalability. The first principles-based ReaxFF approach [7,13] significantly reduces the computational cost of simulating chemical




Fig. 2. Schematic of an embedded divide-and-conquer (EDC) algorithm. (Left) The physical space is subdivided into spatially localized cells, with local atoms constituting subproblems (bottom), which are embedded in a global field (shaded) solved with a tree-based algorithm. (Right) To solve the subproblem in domain Ω_α in the EDC–DFT algorithm, coarse multigrids (gray) are used to accelerate iterative solutions on the original real-space grid (corresponding to the grid refinement level, $l = 3$). The bottom panel shows fine grids adaptively generated near the atoms (spheres) to accurately operate the ionic pseudopotentials on the electronic wave functions.

reactions. However, because of its $O(N^3)$ complexity associated with the variable N -charge problem and the multitude of atomic n -tuple information ($n = 2$ – 6) required to compute interatomic forces, parallelization of the ReaxFF has only seen limited success, and the largest ReaxFF MD simulations to date have involved $N < 10^4$ atoms.

The $O(N^3)$ complexity of ReaxFF arises from the dense linear system of equations to determine atomic charges, $\{q_i | i = 1, \dots, N\}$, at every MD step, i.e., the variable N -charge problem. We have developed a *fast reactive force-field (F-ReaxFF) MD* algorithm, which reduces the complexity to $O(N)$ by combining the FMM [15–17] based on spatial locality and an iterative minimization approach to utilize the temporal locality of the solutions [21]. To further accelerate the convergence, we use a multilevel preconditioned conjugate-gradient (MPCG) method [9,21], 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 thereby improves the parallel efficiency [21].

The chemical bond order, B_{ij} , is an attribute of an atomic pair, (i, j) , and changes dynamically adapting to the local environment. In the ReaxFF, the interatomic 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 [13]. Force calculations in ReaxFF MD thus involve up to atomic 4-tuples explicitly, and require information on 6-tuples implicitly due to chain-rule differentiations through the bond orders. To efficiently handle the resulting multiple interaction ranges, our parallel F-ReaxFF algorithm employs a multilayer cellular-decomposition (MCD) scheme for caching atomic n -tuple ($n = 2$ – 6) information (see Section 3).

The F-ReaxFF calculation of RDX (1,3,5-trinitro-1,3,5-triazine) in this paper has an extensive validation database against DFT calculations, which include not only 1600 equilibrated molecular fragments but also 40 key reaction energies [7].

2.3. Divide-and-conquer density functional theory algorithm on adaptive multigrids

The concurrent DFT-based MD approach is best implemented with a divide-and-conquer algorithm [22], which is based on a data locality principle called quantum near-sightedness [23], and naturally leads to $O(N)$ DFT calculations [12,22,24–27]. 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 [12,25–27]), have attained controlled error bounds, robust convergence properties, and energy conservation during MD simulations, to make large DFT-based MD simulations practical [12,26]. For example, we have recently 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 [12].

The DFT can be formulated as a minimization of the energy functional, $E_{\text{QM}}(\mathbf{r}^N, \psi^M)$, with respect to electronic wave functions (or Kohn–Sham orbitals), $\psi^M(\mathbf{r}) = \{\psi_n(\mathbf{r}) | n = 1, \dots, M\}$, subject to orthonormality constraints (M is the number of independent electronic wave functions and is on the order of N). The EDC-DFT algorithm represents the physical system as a union of overlapping spatial domains, $\Omega = \cup_{\alpha} \Omega_{\alpha}$ (see Fig. 2), 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 [12,28].

The DFT calculation in each domain is performed using a real-space approach [29], in which electronic wave functions are numerically represented on grid points (see Fig. 2). The real-space grid is augmented with coarser multigrids to accelerate the convergence of iterative solutions [12,28,30]. Furthermore, a finer grid is adaptively generated near every atom, in order to accurately operate ionic pseudopotentials for calculating electron–ion interactions [12]. We include electron–ion interactions using norm-conserving pseudopotentials [31] and the exchange–correlation energy in a generalized gradient approximation [32].

The EDC-DFT algorithm on the hierarchical real-space grids is implemented on parallel computers based on spatial decomposition [12]. Each compute node contains one or more domains of the EDC algorithm. For each domain, its electronic structure is computed independently, with little information needed from other compute nodes (only the global density but not individual wave functions is communicated). The resulting large computation/communication ratio makes this approach highly scalable on parallel computers.

The convergence of the new algorithm has been verified for nontrivial problems such as amorphous CdSe and liquid Rb [12]. The EDC-DFT calculation for alumina in this paper (with the domain size $9.0 \times 7.8 \times 8.2$ a.u. and the buffer lengths 4.5, 3.9, and 4.1 a.u.) reproduces an $O(N^3)$ DFT energy within 0.001 a.u. per atom. The EDC-DFT MD algorithm has also overcome the energy drift problem, which plagues most $O(N)$ DFT-based MD algorithms [12].

3. Tunable hierarchical cellular-decomposition parallelization framework

Data locality principles are key to developing a scalable parallel computing framework as well. We have developed a *tunable hierarchical cellular-decomposition (THCD)*

framework to map the above $O(N)$ algorithms onto massively parallel computers with deep memory hierarchies. The THCD maximally exposes data locality and exploits parallelism at multiple decomposition levels, while providing performance tunability through a hierarchy of parameterized cell data/computation structures (see Fig. 3).

At the finest level, EDC algorithms consist of computational cells—linked-list cells (which are identical to the octree leaf cells in the FMM) [15–17] in MRMD and F-ReaxFF, or domains in EDC–DFT [12] (see Fig. 3). In the THCD framework, each compute node (often comprising multiple processors with shared memory) of a parallel computer is identified as a subsystem (P_π^j in Fig. 3) in spatial decomposition, which contains a large number of computational cells. Our EDC algorithms are implemented as hybrid message-passing and shared-memory (MPI + OpenMP) programs, in which inter-node communication for caching and migrating atoms between P_π^j 's is handled with messages, whereas loops over cells within each P_π^j (or MPI process) are parallelized with threads (denoted as dots with arrows in Fig. 3). To avoid performance-degrading critical sections, the threads are ordered by blocking cells, so that the atomic n -tuples being processed by the threads share no common atom.

On top of the computational cells, cell blocks, and spatial-decomposition subsystems, the THCD framework introduces a coarser level of decomposition by defining process groups as MPI Communicators ($PG^j = \cup_\pi P_\pi^j$ in Fig. 3). This provides a mechanism to optimize EDC applications distributed over a loosely coupled collection of parallel computers, e.g., a Grid of globally distributed parallel computers [33,34]. Our programs are designed to minimize global operations across PG^j 's and to overlap computations with inter-group communications [34].

The cellular data structure offers an effective abstraction mechanism for performance optimization. We optimize




Fig. 3. In tunable hierarchical cellular decomposition (THCD), the physical volume is subdivided into process groups, PG^j , each of which is spatially decomposed into processes, P_π^j . Each process consists of a number of computational cells (e.g., linked-list cells in MD or domains in EDC–DFT) of size l_{cell} , which are traversed concurrently by threads (denoted by dots with arrows) to compute interatomic forces in blocks of cells. P_π^j is dynamically augmented with n_{layer} layers of cached cells from neighbor processes.

both data layouts (atoms are sorted according to their cell indices and the linked lists) and computation layouts (force computations are re-ordered by traversing the cells according to a spacefilling curve, a mapping from the 3D space to a 1D list) [35]. Cells are traversed along either a Morton curve (Fig. 3) or a Hilbert curve, instead of the traditional raster-scan order. In a multi-threading case, the Morton curve ensures maximal separation between the threads and thus eliminates critical sections. Furthermore, the cell size is made tunable [14] to optimize the performance. The computational cells are also used in the multilayer cellular-decomposition (MCD) scheme for inter-node caching of atomic n -tuple ($n = 2$ –6) information (see Fig. 3), where n changes dynamically in the MTS or MPCG algorithm (see Section 2). The Morton curve also facilitates a data compression algorithm based on data locality to reduce the I/O. The algorithm uses octree indexing and sorts atoms accordingly on the resulting Morton curve [36]. By storing differences between successive atomic coordinates, the I/O requirement for a given 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 for MD data with user-controlled error bound [36].

The THCD framework includes a topology-preserving computational spatial-decomposition scheme to minimize latency through structured message passing [37] and load-imbalance/communication costs through a novel wavelet-based load-balancing scheme [38]. 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 [39]. To minimize the number of messages, we preserve the 3D mesh topology, so that message passing is performed in a structured way in only six steps. To minimize the load-imbalance cost as well as the message size, we have developed a computational-space decomposition scheme [37]. The main idea of this scheme is that the computational space shrinks where the workload density is high, so that the workload is uniformly distributed in the computational space. To implement the curved computational space, we introduce a curvilinear coordinate transformation. The sum of load imbalance and communication costs is then minimized as a functional of the coordinate transformation, using simulated annealing. We have found that wavelet representation leads to compact representation of curved partition boundaries, and accordingly speeds up the convergence of the minimization procedure [38].

4. Performance tests

The two new parallel reactive MD algorithms, F-ReaxFF and EDC–DFT, as well as their nonreactive predecessor, MRMD, are portable and have been run on various platforms, including Intel Itanium2, Intel Xeon, AMD

Opteron and IBM Power4-based parallel computers. This section presents performance tests of the three algorithms on some of the architectures.

4.1. Performance tunability

We have tested the performance tunability of the THCD parallel implementations of the MRMD, F-ReaxFF, and EDC–DFT algorithms. For example, we typically observe ~10% performance improvement by the data re-ordering and computation re-ordering based on the Morton curve for MRMD.

Fig. 4c shows the effect of the linked-list cell size (in unit of the average volume per atom, Ω/N) on the CPU time of the MRMD algorithm for a 331,776-atom silica material on a 1.4 GHz Pentium III processor. The CPU time takes the minimal value at a cell size of 2.1, which can be understood as a trade-off between the increasing number of floating-point operations (Fig. 4a) and the decreasing number of L2 cache misses (Fig. 4b) as a function of the cell size. Here, we have used the Performance Application Programming Interface (PAPI) to measure the number of cache misses (see <http://icl.cs.utk.edu/papi>).




Fig. 4. (a) The number of floating-point instructions, (b) the number of L2 cache misses, and (c) the CPU time per MD step, as a function of the linked-list cell size, for the MRMD algorithm, to study a 331,776-atom silica material on a 1.4 GHz Pentium III processor.

There is also a trade-off between spatial-decomposition/message-passing (MPI) and thread (OpenMP) parallelisms [40–42] in our hybrid MPI + OpenMP programs. While spatial decomposition involves extra computation on cached cells from neighbor subsystems, its disjoint memory subspaces are free from shared-memory protocol overhead. A marked contrast in this trade-off is observed between the MRMD and F-ReaxFF algorithms. The MRMD is characterized by a low inter-node caching overhead in spatial decomposition and a small number of floating-point operations per memory access, since it mostly consists of look-ups for pre-computed interatomic force tables. In contrast, F-ReaxFF has a large inter-node caching overhead for 6-tuple information and a large computation/memory-access ratio.

Table 1 shows the execution time of the MRMD algorithm for an 8,232,000-atom silica material and that of the F-ReaxFF algorithm for a 290,304-atom RDX crystal on $P = 8$ processors in an 8-way 1.5 GHz Power4 node. (The test was performed on the Iceberg Power4 system at the Arctic Region Supercomputing Center.) We compare different combinations of the number of OpenMP threads per MPI process, n_{td} , and that of MPI processes, n_p , while keeping $P = n_{td} \times n_p$ constant. The optimal combination of (n_{td}, n_p) with the minimum execution time is (1, 8) for the MRMD and is (4, 2) for the F-ReaxFF.

4.2. Scalability

Scalability tests of the two new parallel MD algorithms, F-ReaxFF and EDC–DFT, as well as MRMD, on which they are based, have been performed on the 10,240-processor Columbia supercomputer at the NASA Ames Research Center. The SGI Altix 3000 system uses the NUMAflex global shared-memory architecture, which packages processors, memory, I/O, interconnect, graphics, and storage into modular components called bricks. The computational building block of Altix is the C-Brick, which consists of four Intel Itanium2 processors (in two nodes), local memory, and a two-controller application-specific integrated circuit called the Scalable Hub (SHUB). Each SHUB

Table 1
Performance of hybrid MPI + OpenMP programs

Number of OpenMP threads, n_{td}	Number of MPI processes, n_p	Execution time/MD time step (s)	
		MRMD	F-ReaxFF
1	8	4.19	62.5
2	4	5.75	58.9
4	2	8.60	54.9
8	1	12.5	120

Execution time per MD time step on $P = n_{td} \times n_p = 8$ processors in an 8-way 1.5 GHz Power4 node, with different combinations of the number of OpenMP threads per MPI process, n_{td} , and that of MPI processes, n_p , for: (1) the MRMD algorithm for an 8,232,000-atom silica system; and (2) the F-ReaxFF algorithm for a 290,304-atom RDX system. The minimum execution time for each algorithm is typed in boldface.

interfaces to the two CPUs within one node, along with memory, I/O devices, and other SHUBs. The Altix cache-coherency protocol implemented in the SHUB integrates the snooping operations of the Itanium2 and the directory-based scheme used across the NUMAflex interconnection fabric. A load/store cache miss causes the data to be communicated via the SHUB at the cache-line granularity and automatically replicated in the local cache.

The 64-bit Itanium2 architecture operates at 1.5 GHz and is capable of issuing two multiply–add operations per cycle for a peak performance of 6Gflops. The memory hierarchy consists of 128 floating-point registers and three on-chip data caches (32KB L1, 256KB L2, and 6MB L3). The Itanium2 cannot store floating-point data in L1, making register loads and spills a potential source of bottlenecks; however, a relatively large register set helps mitigate this issue. The superscalar processor implements the Explicitly Parallel Instruction set Computing (EPIC) technology, where instructions are organized into 128-bit VLIW bundles. The Altix platform uses the NUMalink3 interconnect, a high-performance custom network in a fat-tree topology, in which the bisection bandwidth scales linearly with the number of processors. Columbia runs 64-bit Linux version 2.4.21. Our experiments use a 6.4TB parallel XFS file system with a 35-fiber optical channel connection to the CPUs.

Columbia is configured as a cluster of 20 Altix boxes, each with 512 processors and 1TB of global shared-access memory. Of these 20 boxes, 12 are model 3700 and the remaining eight are BX2—a double-density version of the 3700. Four of the BX2 boxes are linked with NUMalink4 technology to allow the global shared-memory constructs to significantly reduce inter-processor communication latency. This 2048-processor subsystem within Columbia provides a 13Tflops peak capability platform, and was the basis of the computations reported here.

Fig. 5 shows the execution time of the F-ReaxFF MD algorithm for RDX material as a function of the number of processors, P . In this and following figures, we set n_{td} to 1. We scale the problem size linearly with the number of processors, so that the number of atoms, $N = 36,288P$ ($P = 1, \dots, 1920$). The computation time includes three conjugate-gradient (CG) iterations to solve the electronegativity equalization problem for determining atomic charges at each MD time step. The execution time increases only slightly as a function of P , and this signifies an excellent parallel efficiency. We define the speed of an MD program as a product of the total number of atoms and time steps executed per second. The speedup is the ratio between the speed of P processors and that of one processor. The parallel efficiency is the speedup divided by P . On 1920 processors, the isogranular parallel efficiency of the F-ReaxFF algorithm is 0.953. A better measure of the inter-box scaling efficiency based on NUMalink4 is the speedup from 480 processors in one box to 1920 processors in four boxes, divided by the number of boxes. On 1920 processors, the measured inter-box scaling efficiency is

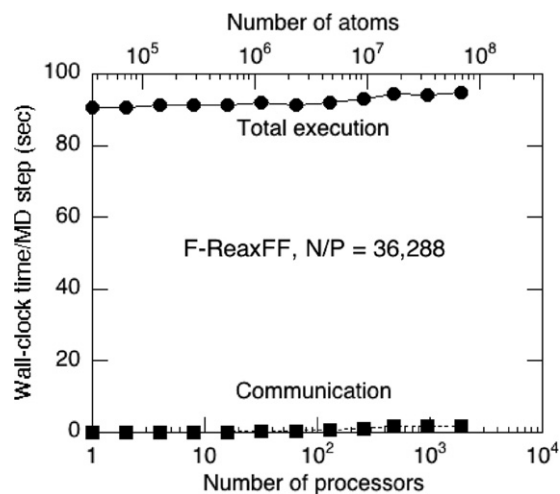


Fig. 5. Total execution (circles) and communication (squares) times per MD time step as a function of the number of processors for the F-ReaxFF MD algorithm with scaled workloads—36,288 P atom RDX systems on P processors ($P = 1, \dots, 1920$) of Columbia.

0.995. Also the algorithm involves very small communication time (see Fig. 5).

Fig. 6 shows the performance of the EDC–DFT-based MD algorithm with scaled workloads—720 P atom alumina systems on P processors ($P = 1, \dots, 1920$). In the EDC–DFT calculations, each domain of size $6.66 \times 5.76 \times 6.06 \text{ \AA}^3$ contains 40 electronic wave functions, where each wave function is represented on $28^3 = 21,952$ grid points. The execution time includes three self-consistent (SC) iterations to determine the electronic wave functions and the Kohn–Sham potential, with three CG iterations per SC cycle to refine each wave function iteratively. The largest calculation on 1920 processors involves 1,382,400 atoms and 5,529,600 electronic wave functions on

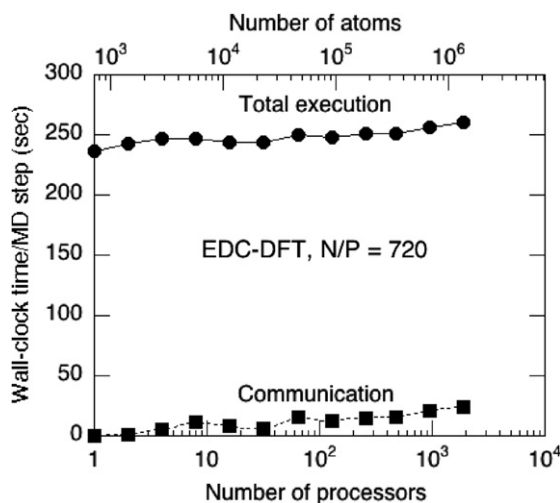


Fig. 6. Total execution (circles) and communication (squares) times per MD time step as a function of the number of processors for the EDC–DFT MD algorithm with scaled workloads—720 P atom alumina systems on P processors ($P = 1, \dots, 1920$) of Columbia.

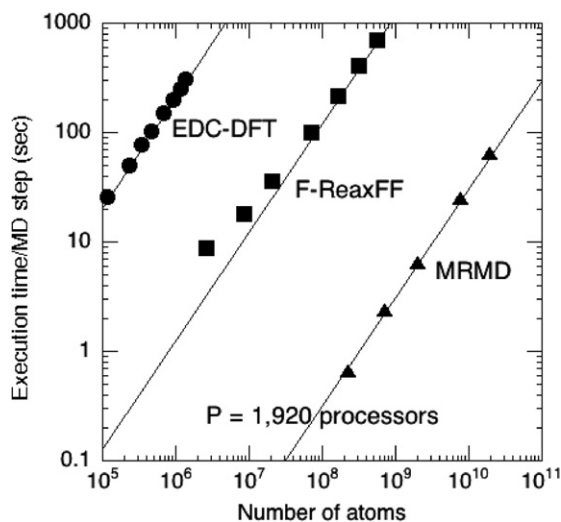


Fig. 7. Design-space diagram for reactive and nonreactive MD simulations on 1920 Itanium2 processors of Columbia. The figure shows the total execution time per MD step as a function of the number of atoms for three linear-scaling algorithms: quantum-mechanical MD based on the embedded divide-and-conquer density functional theory (EDC–DFT, circles); fast reactive force-field MD (F-ReaxFF, squares); space–time multi-resolution MD (MRMD, triangles). Lines show ideal $O(N)$ scaling.

121,385,779,200 grid points, for which the isogranular parallel efficiency is 0.907. The inter-box scaling efficiency between 480 and 1920 processors is 0.966.

Major design parameters for reactive and nonreactive MD simulations of materials include the number of atoms in the simulated material and the method to compute interatomic forces (classically in MRMD, semi-empirically in F-ReaxFF MD, or quantum mechanically in EDC–DFT MD). Fig. 7 shows a design-space diagram for classical and quantum-mechanical MD simulations on 1920 Itanium2 processors of Columbia. The largest benchmark tests in this study 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. The figure demonstrates perfect linear scaling for all the three algorithms, with prefactors spanning five orders-of-magnitude. Only exception is the F-ReaxFF algorithm below 100 million atoms, where the execution time scales even sub-linearly.

5. Applications

The scalable parallel MD algorithms, F-ReaxFF and EDC–DFT, have been applied to the study of a number of chemically reactive material processes, and this section briefly describes some of the applications.

5.1. Shock-induced detonation of energetic materials

We have performed F-ReaxFF MD simulations to study shock-initiated detonation of RDX (1,3,5-trinitro-1,3,5-triazine, $C_3N_6O_6H_6$) matrix embedded with aluminum nanoparticles (Fig. 1a). 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.

The 1.01 million-atom F-ReaxFF MD simulation has been performed on 256 Xeon processors to study shock-initiated detonation of RDX crystal/oxidized Al nanoparticle (n -Al) composite. In the simulation, two slabs of the RDX/ n -Al 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 n -Al 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 [43].

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. Fig. 8 compares MD and experimental results on the shock velocity as a function of the particle velocity that drives the shock [7]. The MD and experimental data agree very well. Furthermore, the simulation shows

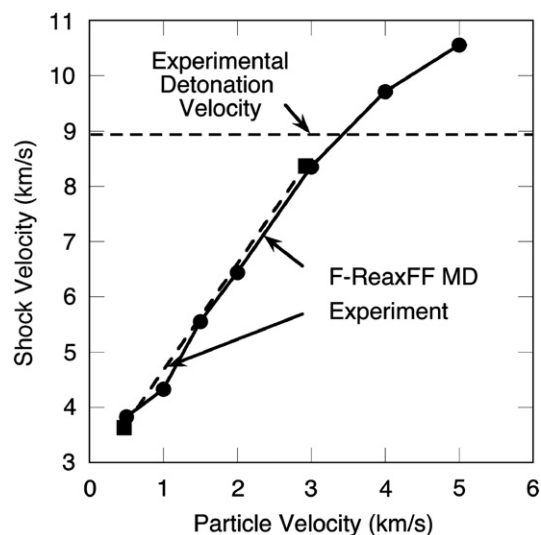


Fig. 8. F-ReaxFF MD and experimental data on the planar shock velocity in RDX as a function of the particle velocity. An experimental detonation velocity in Ref. [44] is also shown.




Fig. 9. (a) Close-up of EDC–DFT MD simulation of an α - Al_2O_3 (0001) surface immersed in water. White, red, and green spheres represent hydrogen, oxygen, and aluminum atoms, respectively. (b) Indentation of amorphous silica covered with adsorbed water layer. White and yellow spheres are hydrogen and silicon atoms, whereas blue and red spheres are oxygen atoms originated from water and silica substrate, respectively. The indenter is shown in green.

a sudden increase of the number of molecular products such as HONO, N_2 , and OH above a shock velocity ~ 9 km/s, which is consistent with an experimental detonation velocity [44].

5.2. Stress-induced bond breaking in ceramics in corrosive environments

We have also started EDC–DFT MD simulations involving 10,368 atoms on 240 dual-processor/dual-core AMD Opteron nodes (in total of 960 cores), to study the effect of applied stress on the bond breaking at an α - Al_2O_3 (0001) surface immersed in water. The alumina substrate is $57.0 \times 67.2 \times 12.2 \text{ \AA}^3$ in the $[2\bar{1}\bar{1}0]$, $[10\bar{1}0]$ and $[0001]$ directions, covered with 13 \AA layer of adsorbed water. The performance and lifetime of materials widely used in industrial applications is often severely limited by corrosion of these systems in an environment containing oxygen and water. Most critical here is premature and

catastrophic failure of materials resulting from chemically influenced corrosion. The basic requirements for the operation of structural systems exposed to corroding conditions under stress loads are safety and reliability. Such safe and reliable operation is endangered by the uncertainties in stress corrosion cracking (SCC) [45,46]. To prevent SCC and to predict the lifetime beyond which SCC may cause failure requires that we understand the atomistic mechanisms underlying SCC; that is the conditions influencing initiation of SCC and the dynamics and growth rates. Fig. 9a shows a close-up of the simulation at zero stress, which shows the dissociation of water molecules consistent with the mechanism found by Hass et al. [47]. Under a uniaxial strain of 0.15 in the $[10\bar{1}0]$ direction, the simulation exhibits significant relaxation of the alumina surface.

We are also studying the effect of adsorbed layers of water and hydrazine (N_2H_4) on the indentation behavior of amorphous silica using F-ReaxFF MD simulations (Fig. 9b). Indentation is a unique local probe to measure




Fig. 10. A hierarchy of MD simulation methods. (Top) Reactive MD methods and their largest implementations include DFT-based MD of metals involving 1000 atoms (or $N_{el} = 10,000$ electrons) [50], fragment molecular orbital (FMO) method calculation of proteins involving 20,000 atoms [51], $O(N)$ DFT (ONDFT, the present paper), reactive force field (ReaxFF, the present paper), whereas the largest nonreactive MD simulations based on effective force fields (EFF) include a 160 billion-atom MD of metals [52]. (Bottom) The spatio-temporal scale, NT , of MD simulations. Our large MD simulations of ceramics [48,53,54] had $NT = 0.03$ – 0.04 atom·s, whereas one of the longest biomolecular simulations on the HP-36 peptide had $NT = 0.0096$ [55].

mechanical properties of materials [48]. Experimental microhardness measurements by indentation exhibit loading-time dependence of the crack initiation stress in water, while no time dependence is observed in nonaqueous liquids such as hydrazine [49]. The silica substrate in our simulation consists of 39,201 atoms and is $106 \times 113 \times 50 \text{ \AA}^3$, covered with a monolayer of adsorbed water or hydrazine. The simulations reveal significant diffusion of water into the silica substrate under the indenter.

6. Conclusions

Since we demonstrated a hundred thousand-atom reactive molecular dynamics simulation based on the density functional theory in 2001 [3], significant progresses have been made in simulation methods (e.g., first principles-based reactive force-field molecular dynamics), linear-scaling algorithms (e.g., embedded divide-and-conquer density functional theory algorithm on adaptive multigrids), and scalable parallel computing technologies (e.g., tunable hierarchical cellular decomposition with wavelet-based adaptive load balancing). See Fig. 10 for a hierarchy of MD simulation methods. Based on these innovations, this paper has demonstrated million-to-billion atom reactive molecular dynamics simulations, which show considerable promise for atomistic simulations of chemical reactions with unprecedented scales and accuracy on emerging petaflops-scale computer architectures.

The hierarchy of molecular dynamics simulation algorithms developed in this paper can be integrated seamlessly into a hierarchical simulation framework, which embeds accurate but compute-intensive simulations in coarse simulations only when and where high fidelity is required [6,56–58]. Such a hybrid approach is complementary to the approach in this paper. The hybrid approach applies to a class of problems, e.g., certain aspects of stress corrosion cracking, in which localized chemical reactions are studied on moderate computational resources [6], whereas other classes of problems, e.g., combustion of nanoenergetic materials [7], require full quantum-mechanical simulations developed in this paper to study extended chemical reactions on the highest end computers. Our hierarchical simulation framework consists of: (1) hierarchical division of the physical system into subsystems of decreasing sizes and increasing quality-of-solution (QoS) requirements, $S_0 \supset S_1 \supset \dots \supset S_n$; and (2) a suite of simulation services, M_α ($\alpha = 0, 1, \dots, n$), of ascending order of accuracy (e.g., MRMD \prec F-ReaxFF \prec EDC-DFT). In our additive hybridization scheme [6,57], an accurate estimate of the energy of the entire system is obtained from the recurrence relation, $E_\alpha(S_i) = E_{\alpha-1}(S_i) + E_\alpha(S_{i+1}) - E_{\alpha-1}(S_{i+1})$. The scalable parallel simulation techniques presented in this paper, with such a dynamically extensible hierarchical simulation framework, should open up enormous opportunities for scientific computing on high-end computers.

Data locality also plays a critical role in designing scalable data visualization and mining techniques within the

THCD framework. We have developed a scalable visualization system, Atomviewer, to allow the viewer to walk through a billion atoms [59]. The system uses: the octree data structure as an efficient abstraction mechanism to extract atoms within the field-of-view (view frustum culling); a novel probabilistic approach to remove far atoms that are hidden by other atoms (occlusion culling); parallel/distributed processing of these culling tasks on a Linux cluster connected to a graphics server; a machine-learning approach to predict the user's next movement and prefetch data from the Linux cluster to the graphics server; and multiresolution rendering. The resulting system renders a billion-atom dataset at nearly interactive frame rates on a dual-processor SGI Onyx2 with an InfiniteReality2 graphics pipeline, connected to a 4-processor Linux cluster. We have also developed a data mining approach based on a graph algorithm (i.e., the shortest-path circuit analysis) to detect and track topological anomalies in multimillion-node chemical bond networks in materials [48]. At the Collaboratory for Advanced Computing and Simulations (CACs) at USC, the Atomviewer is used on an 8 ft \times 14 ft tiled display wall (Fig. 1b) driven by a 26-processor Linux cluster and an immersive and interactive virtual environment called ImmersaDesk. Ultrascale simulations presented in this paper, combined with these massive data visualization and mining approaches, promise to bring in fundamental advances in science.

Finally 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 [48,53,54] simulate sub-billion atoms for sub-nanosecond, resulting in $NT = 0.03\text{--}0.04$ atoms s. Biomolecular simulations, on the other hand, involve much smaller number of atoms ($N \sim 10^4$) but for longer times ($T \sim 10^{-6}$ s). One of the largest NT values, 0.0096, in biomolecular simulations was achieved by Duan and Kollman in their MD simulation of the HP-36 peptide [55]. Petaflops computers are expected to push the spatiotemporal envelope to $NT \sim 1$ and beyond, thereby bringing in further new scientific knowledge.

Acknowledgements

This work was partially supported by AFOSR-DUR-INT, ARO-MURI, DARPA-PROM, DOE, and NSF. Benchmark tests were performed using the NASA Columbia supercomputer at the NASA Ames Research Center, and at Department of Defense's Major Shared Resource Centers under a DoD Challenge Project. Programs have been developed using the 2000-processor (6Tflops) Xeon/Opteron/Apple G5 cluster at the Collaboratory for Advanced Computing and Simulations and the 4000-processor (13.8Tflops) Linux cluster at the High Performance Computing Center at the USC. The authors thank Walter Kohn and Emil Prodan for discussions on $O(N)$ density functional theory, Zhen Lu for tunability tests, and Davin

Chan, Johnny Chang, Bob Ciotti, Edward Hook, Art Lazanoff, Bron Nelson, Charles Niggley, and William Thigpen for technical discussions on Columbia.

References

- [1] F.F. Abraham, R. Walkup, H.J. Gao, M. Duchaineau, T.D. de la Rubia, M. Seager, *Proc. Natl. Acad. Sci.* 99 (2002) 5777–5782.
- [2] K. Kadau, T.C. Germann, P.S. Lomdahl, B. Lee Holian, *Science* 296 (2002) 1681–1684.
- [3] A. Nakano, R.K. Kalia, P. Vashishta, T.J. Campbell, S. Ogata, F. Shimojo, S. Saini, *Sci. Prog.* 10 (2002) 263–270.
- [4] L. Kalé, R. Skeel, M. Bhandarkar, R. Brunner, A. Gursoy, N. Krawetz, J. Phillips, A. Shinozaki, K. Varadarajan, K. Schulten, *J. Comput. Phys.* 151 (1999) 283–312.
- [5] R. Car, M. Parrinello, *Phys. Rev. Lett.* 55 (1985) 2471–2474.
- [6] S. Ogata, F. Shimojo, R.K. Kalia, A. Nakano, P. Vashishta, *J. Appl. Phys.* 95 (2004) 5316–5323.
- [7] A. Strachan, A.C.T. van Duin, D. Chakraborty, S. Dasgupta, W.A. Goddard III, *Phys. Rev. Lett.* 91 (2003) 098301:1–098301:4.
- [8] R.A. Kendall, E. Apra, D.E. Bernholdt, E.J. Bylaska, M. Dupuis, G.I. Fann, R.J. Harrison, J. Ju, J.A. Nichols, J. Nieplocha, T.P. Straatsma, T.L. Windus, A.T. Wong, *Comput. Phys. Commun.* 128 (2000) 260–283.
- [9] T.J. Campbell, R.K. Kalia, A. Nakano, P. Vashishta, S. Ogata, S. Rodgers, *Phys. Rev. Lett.* 82 (1999) 4866–4869; T.J. Campbell, G. Aral, S. Ogata, R.K. Kalia, A. Nakano, P. Vashishta, *Phys. Rev. B* (2005) 205413:1–205413:14.
- [10] P. Hohenberg, W. Kohn, *Phys. Rev.* 136 (1964) B864–B871.
- [11] W. Kohn, P. Vashishta, in: N.H. March, S. Lundqvist (Eds.), *Inhomogeneous Electron Gas*, Plenum, New York, 1983, pp. 79–184.
- [12] F. Shimojo, R.K. Kalia, A. Nakano, P. Vashishta, *Comput. Phys. Commun.* 167 (2005) 151–164.
- [13] A.C.T. van Duin, S. Dasgupta, F. Lorant, W.A. Goddard III, *J. Phys. Chem. A* 105 (2001) 9396–9409.
- [14] R.C. Whaley, A. Petitet, J.J. Dongarra, *Par. Comput.* 27 (2001) 3–35.
- [15] L. Greengard, V. Rokhlin, *J. Comput. Phys.* 73 (1987) 325–348.
- [16] A. Nakano, R.K. Kalia, P. Vashishta, *Comput. Phys. Commun.* 83 (1994) 197–214.
- [17] S. Ogata, T.J. Campbell, R.K. Kalia, A. Nakano, P. Vashishta, S. Vemparala, *Comput. Phys. Commun.* 153 (2003) 445–461.
- [18] M.E. Tuckerman, D.A. Yarne, S.O. Samuelson, A.L. Hughes, G.J. Martyna, *Comput. Phys. Commun.* 128 (2000) 333–376.
- [19] A. Nakano, *Comput. Phys. Commun.* 105 (1997) 139–150.
- [20] A. Nakano, *Int. J. High Perf. Comput. Appl.* 13 (1999) 154–162.
- [21] A. Nakano, *Comput. Phys. Commun.* 104 (1997) 59–69.
- [22] W. Yang, *Phys. Rev. Lett.* 66 (1991) 1438–1441.
- [23] W. Kohn, *Phys. Rev. Lett.* 76 (1996) 3168–3171, E. Prodan, W. Kohn, private communication.
- [24] S. Goedecker, *Rev. Mod. Phys.* 71 (1999) 1085–1123.
- [25] F. Shimojo, R.K. Kalia, A. Nakano, P. Vashishta, *Comput. Phys. Commun.* 140 (2001) 303–314.
- [26] J.-L. Fattebert, F. Gygi, *Comput. Phys. Commun.* 162 (2004) 24–36.
- [27] C.-K. Skylaris, P.D. Haynes, A.A. Mostofi, M.C. Payne, *J. Chem. Phys.* 122 (2005) 084119:1–084119:10.
- [28] T.L. Beck, *Rev. Mod. Phys.* 72 (2000) 1041–1080.
- [29] J.R. Chelikowsky, Y. Saad, S. Ögüt, I. Vasiliev, A. Stathopoulos, *Phys. Status Solidi (b)* 217 (2000) 173–195.
- [30] J.-L. Fattebert, J. Bernholc, *Phys. Rev. B* 62 (2000) 1713–1722.
- [31] N. Troullier, J.L. Martins, *Phys. Rev. B* 43 (1991) 1993–2006.
- [32] J.P. Perdew, K. Burke, M. Ernzerhof, *Phys. Rev. Lett.* 77 (1996) 3865–3868.
- [33] G. Allen, T. Damlitsch, I. Foster, N.T. Karonis, M. Ripeanu, E. Seidel, B. Toonen, in: *Proceedings of SC2001*, ACM, New York, 2001.
- [34] H. Kikuchi, R.K. Kalia, A. Nakano, P. Vashishta, F. Shimojo, S. Saini, in: *Proceedings of SC2002*, IEEE, Los Alamitos, 2002.
- [35] J. Mellor-Crummey, D. Whalley, K. Kennedy, *Int. J. Par. Prog.* 29 (2001) 217–247.
- [36] A. Omeltchenko, T.J. Campbell, R.K. Kalia, X. Liu, A. Nakano, P. Vashishta, *Comput. Phys. Commun.* 131 (2000) 78–85.
- [37] A. Nakano, T.J. Campbell, *Par. Comput.* 23 (1997) 1461–1478.
- [38] A. Nakano, *Concurrency: Practice and Experience* 11 (1999) 343–353.
- [39] A. Nakano, R.K. Kalia, A. Sharma, P. Vashishta, in: K.Y. Lam, H.P. Lee (Eds.), *Computational Methods in Large Scale Simulation*, World Scientific, Singapore, 2005, pp. 229–243.
- [40] D.S. Henty, in: *Proceedings of SC2000*, IEEE, Los Alamitos, 2000.
- [41] H. Shan, J.P. Singh, L. Oliker, R. Biswas, *J. Par. Distrib. Comput.* 62 (2002) 241–266.
- [42] H. Shan, J.P. Singh, L. Oliker, R. Biswas, *Par. Comput.* 29 (2003) 167–186.
- [43] N. Umezawa, R.K. Kalia, A. Nakano, P. Vashishta, F. Shimojo, *J. Chem. Phys.*, in press.
- [44] R.D. Cowan, W. Fickett, *J. Chem. Phys.* 24 (1956) 932–939.
- [45] T.A. Michalske, B.C. Bunker, S.W. Freiman, *J. Am. Ceram. Soc.* 69 (1986) 721–724.
- [46] T. Zhu, J. Li, X. Lin, S. Yip, *J. Mech. Phys. Solids* 53 (2005) 1597–1623.
- [47] K.C. Hass, W.F. Schneider, A. Curioni, W. Andreoni, *Science* 282 (1998) 265–268.
- [48] I. Szlufarska, A. Nakano, P. Vashishta, *Science* 309 (2005) 911–914.
- [49] K. Hirao, M. Tomozawa, *J. Am. Ceram. Soc.* 70 (1987) 497–502.
- [50] F. Gygi, E.W. Draeger, B.R. de Supinski, R.K. Yates, F. Franchetti, S. Kral, J. Lorenz, C.W. Ueberhuber, J.A. Gunnels, J.C. Sexton, in: *Proceedings of SC05*, ACM, New York, 2005.
- [51] T. Ikegami, T. Ishida, D.G. Fedorov, K. Kitaura, Y. Inadomi, H. Umeda, M. Yokokawa, S. Sekiguchi, in: *Proceedings of SC05*, ACM, New York, 2005.
- [52] T.C. Germann, K. Kadau, P.S. Lomdahl, Technical Report, Los Alamos National Laboratory, 2005.
- [53] Z. Lu, K. Nomura, A. Sharma, W. Wang, C. Zhang, A. Nakano, R.K. Kalia, P. Vashishta, E. Bouchaud, C.L. Rountree, *Phys. Rev. Lett.* 95 (2005) 135501:1–135501:4.
- [54] P.S. Branicio, R.K. Kalia, A. Nakano, P. Vashishta, *Phys. Rev. Lett.* 96 (2006) 065502:1–065502:4.
- [55] Y. Duan, P.A. Kollman, *Science* 282 (1998) 740–744.
- [56] J.Q. Broughton, F.F. Abraham, N. Bernstein, E. Kaxiras, *Phys. Rev. B* 60 (1999) 2391–2403.
- [57] S. Ogata, E. Lidorikis, F. Shimojo, A. Nakano, P. Vashishta, R.K. Kalia, *Comput. Phys. Commun.* 138 (2001) 143–154.
- [58] G. Lu, E.B. Tadmor, E. Kaxiras, *Phys. Rev. B* 73 (2006) 024108:1–024108:4.
- [59] A. Sharma, A. Nakano, R.K. Kalia, P. Vashishta, S. Kodiyalam, P. Miller, W. Zhao, X. Liu, T.J. Campbell, A. Haas, *Presence: Teleoperators Virtual Environ.* 12 (2003) 85–95.