

Metascalable molecular dynamics simulation of nano-mechano-chemistry

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2008 J. Phys.: Condens. Matter 20 294204

(<http://iopscience.iop.org/0953-8984/20/29/294204>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 129.252.86.83

The article was downloaded on 29/05/2010 at 13:33

Please note that [terms and conditions apply](#).

Metascalable molecular dynamics simulation of nano-mechano-chemistry

F Shimojo^{1,2}, R K Kalia¹, A Nakano¹, K Nomura¹ and P Vashishta¹

¹ Collaboratory for Advanced Computing and Simulations, Department of Computer Science, Department of Physics and Astronomy, Department of Chemical Engineering and Materials Science, University of Southern California, Los Angeles, CA 90089-0242, USA

² Department of Physics, Kumamoto University, Kumamoto 860-8555, Japan

E-mail: shimojo@kumamoto-u.ac.jp, rkalia@usc.edu, anakano@usc.edu, knomura@usc.edu and priyav@usc.edu

Received 6 January 2008, in final form 12 February 2008

Published 24 June 2008

Online at stacks.iop.org/JPhysCM/20/294204

Abstract

We have developed a metascalable (or ‘design once, scale on new architectures’) parallel application-development framework for first-principles based simulations of nano-mechano-chemical processes on emerging petaflops architectures based on spatiotemporal data locality principles. The framework consists of (1) an embedded divide-and-conquer (EDC) algorithmic framework based on spatial locality to design linear-scaling algorithms, (2) a space–time-ensemble parallel (STEP) approach based on temporal locality to predict long-time dynamics, and (3) a tunable hierarchical cellular decomposition (HCD) parallelization framework to map these scalable algorithms onto hardware. The EDC-STEP-HCD framework exposes and expresses maximal concurrency and data locality, thereby achieving parallel efficiency as high as 0.99 for 1.59-billion-atom reactive force field molecular dynamics (MD) and 17.7-million-atom (1.56 trillion electronic degrees of freedom) quantum mechanical (QM) MD in the framework of the density functional theory (DFT) on adaptive multigrids, in addition to 201-billion-atom nonreactive MD, on 196 608 IBM BlueGene/L processors. We have also used the framework for automated execution of adaptive hybrid DFT/MD simulation on a grid of six supercomputers in the US and Japan, in which the number of processors changed dynamically on demand and tasks were migrated according to unexpected faults. The paper presents the application of the framework to the study of nanoenergetic materials: (1) combustion of an Al/Fe₂O₃ thermite and (2) shock initiation and reactive nanojets at a void in an energetic crystal.

1. Introduction

The ever-increasing capability of high-end computing platforms is enabling unprecedented scales of first-principles based simulations to predict system-level behavior of complex systems [1]. An example is large-scale molecular dynamics (MD) simulation involving a million to a billion atoms, in which interatomic forces are computed quantum mechanically to accurately describe chemical reactions [2]. Such simulations can couple chemical reactions at the atomistic scale and mechanical processes at the mesoscopic scale to solve broad mechano-chemistry problems of great societal impact such as stress corrosion cracking, where chemical reactions at the crack tip are inseparable from long-range stress fields, and the sensitivity

of high-energy-density nanomaterials, in which chemical reactions sustain shock waves.

Specifically, several computing platforms offer great promise. One is petaflops (10¹⁵ floating-point operations per second) computers to be built in the next few years [3, 4]. Already, the 212 992-processor IBM BlueGene/L supercomputer at the Lawrence Livermore National Laboratory has achieved 0.478 petaflops for solving a linear system of equations [5]. The second platform is a Grid of geographically distributed supercomputers [6]. For example, the US TeraGrid connects a number of high-end supercomputers via a high-speed optical-fiber network [7, 8]. Also promising are many-core microprocessors. Computer industry is facing a historical shift, in which Moore’s law due to

ever-increasing clock speeds has been subsumed by increasing numbers of cores in microchips [4, 9]. The number of cores per microprocessor is expected to double at each generation, reaching a thousand in ten years. We can thus envision a Grid of 10^2 parallel computing clusters, each consisting of 10^5 compute nodes, where each node contains 10^3 cores.

The many-core revolution will mark the end of the free-ride era (i.e., legacy software will run faster on newer chips), resulting in a dichotomy—subsiding speed-up of conventional software and exponential speed-up of scalable parallel applications [9]. Recent progress in high-performance technical computing has led to key technologies for parallel computing with portable scalability. An example is an embedded divide-and-conquer (EDC) algorithmic framework to design linear-scaling algorithms for broad scientific and engineering applications (e.g. equation solvers, constrained optimization, search, visualization, and graphs involving massive data) based on spatiotemporal locality principles [2, 10]. The EDC framework maximally exposes concurrency and data locality, thereby achieving reusable ‘design once, scale on new architectures’ (or metascalable) applications. It is expected that such metascalable algorithms will continue to scale on future many-core architectures. ‘Seven dwarfs’ (a dwarf is an algorithmic method that captures a pattern of computation and communication), which were first identified by Phillip Colella of the Lawrence Berkeley National Laboratory, have been used widely to develop scalable parallel programming models and architectures [4]. We expect that the EDC algorithmic framework will serve as a ‘metascalable dwarf’ to represent broad large-scale scientific and engineering applications.

On the advanced computing platforms mentioned above, we will use a hierarchy of atomistic simulation methods. In MD simulation, the system is represented by a set of N point atoms whose trajectories are followed to study material properties. Quantum mechanical (QM) simulation further treats electronic wavefunctions explicitly to describe chemical reactions. To seamlessly couple MD and QM simulations, we have found it beneficial to introduce an intermediate layer, the first-principles based reactive force field (ReaxFF) approach, in which interatomic interaction adapts dynamically to the local environment to describe chemical reactions [2, 10, 11]. The ReaxFF is trained by performing thousands of small QM calculations.

This paper is organized as follows. Section 2 describes metascalable computing technologies for billion-atom simulations of chemical reactions based on data locality principles. First, we develop an EDC framework to design linear-scaling algorithms for broad applications based on spatial locality. Second, we develop space–time-ensemble parallelism (STEP) to predict long-time dynamics based on temporal locality. We also develop a tunable hierarchical cellular decomposition (HCD) framework to map these $O(N)$ algorithms onto parallel computers with deep memory hierarchy. In section 3, we discuss the application of these computational techniques to mechano-chemical processes at the nanoscale: combustion and shock-induced initiation of energetic materials.

2. A metascalable dwarf

2.1. Embedded divide-and-conquer (EDC) algorithmic framework

We have designed linear-scaling algorithms for large-scale atomistic simulations based on a unified algorithmic framework. In the embedded divide-and-conquer (EDC) algorithms, the physical system is divided into spatially localized computational cells [2, 10]. These cells are embedded in a global field that is computed efficiently with tree-based algorithms (figure 1).

Within the EDC framework, we have designed a number of $O(N)$ algorithms (N is the number of atoms). For example, we have designed a space–time multiresolution MD (MRMD) algorithm to reduce the $O(N^2)$ complexity of the N -body problem to $O(N)$ [12]. MD simulation follows the trajectories of N point atoms by numerically integrating coupled ordinary differential equations. The hardest computation in MD simulation is the evaluation of the long-range electrostatic potential at N atomic positions. Since, each evaluation involves contributions from $N - 1$ sources, direct summation requires $O(N^2)$ operations. The MRMD algorithm uses the octree-based fast multipole method (FMM) [13–15] to reduce the computational complexity to $O(N)$ based on spatial locality. We also use multiresolution in time [16], where temporal locality is utilized by computing forces from further atoms with less frequency [14, 17].

We have also designed a fast ReaxFF (F-ReaxFF) algorithm to solve the $O(N^3)$ variable N -charge problem in chemically reactive MD in $O(N)$ time [18]. To describe chemical bond breakage and formation, the ReaxFF potential energy is a function of the positions of atomic pairs, triplets and quadruplets as well as the chemical bond orders of all constituent atomic pairs [11]. To describe charge transfer, ReaxFF uses a charge-equilibration (QEq) scheme, in which atomic charges are determined at every MD step to minimize the electrostatic energy with the charge-neutrality constraint [19–21]. This variable N -charge problem amounts to solving a dense linear system of equations, which requires $O(N^3)$ operations. The F-ReaxFF algorithm uses the FMM to perform the matrix–vector multiplications with $O(N)$ operations. It further utilizes the temporal locality of the solutions to reduce the amortized computational cost averaged over simulation steps to $O(N)$. To further speed up the solution, we use a multilevel preconditioned conjugate gradient (MPCG) method [22]. This method splits the Coulomb interaction matrix into far-field and near-field matrices and uses the sparse near-field matrix as a preconditioner. The extensive use of the sparse preconditioner enhances the data locality, thereby increasing the parallel efficiency.

To approximately solve the exponentially complex quantum N -body problem, we use a divide-and-conquer (DC) density functional theory (DFT) algorithm [23–25]. The DFT reduces the exponential complexity to $O(N^3)$, by solving N_{el} one-electron problems self-consistently instead of one N_{el} -electron problem (the number of electrons, N_{el} , is of the order of N) [26, 27]. The DFT problem can be formulated as a minimization of an energy functional with respect to electronic

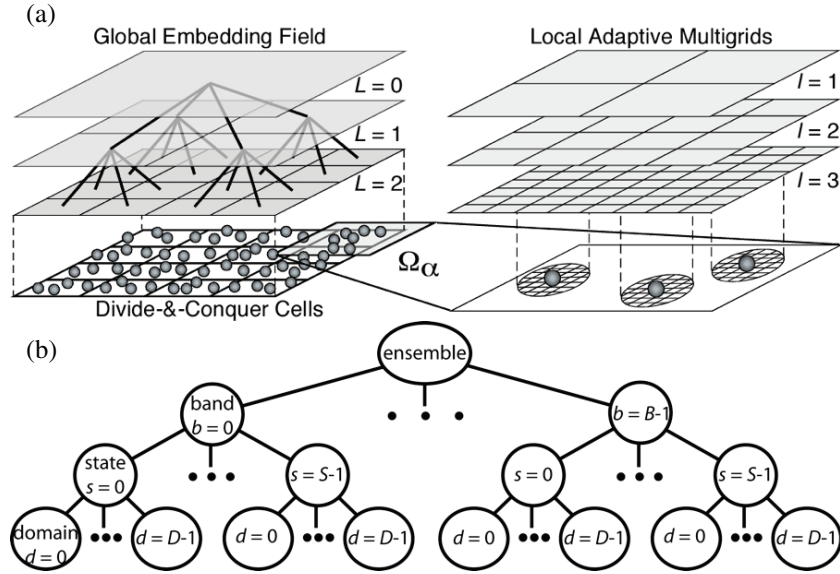


Figure 1. (a) Schematic diagram of embedded divide-and-conquer (EDC) algorithms. The physical space is subdivided into spatially localized cells, with local atoms constituting subproblems, which are embedded in a global field solved with tree-based algorithms. To solve the subproblem in domain Ω_α in the divide-and-conquer density functional theory (DC-DFT) algorithm, coarse multigrids are used to accelerate iterative solutions on the original real-space grid (corresponding to the grid refinement level $l = 3$). Fine grids are adaptively generated near the atoms to accurately operate the ionic pseudopotentials on the electronic wavefunctions. (b) Tree-structured processor organization in the hierarchical space–time-ensemble parallelization (STEP). An ensemble consists of B bands, each consisting of S states. Each state in turn contains D spatial domains.

wavefunctions. In the DC-DFT algorithm, the physical space is a union of overlapping domains (figure 1),

$$\Omega = \sum_\alpha \Omega_\alpha, \quad (1)$$

and physical properties are computed as linear combinations of domain properties that in turn are computed from local electronic wavefunctions. For example, the electronic density $\rho(\mathbf{r})$ is calculated as

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

where the support function $p^\alpha(\mathbf{r})$ vanishes outside domain Ω_α and satisfies the sum rule, $\sum_\alpha p^\alpha(\mathbf{r}) = 1$, and $f(\varepsilon_n^\alpha)$ is the Fermi distribution function corresponding to the energy ε_n^α of the n th electronic wavefunction (or Kohn–Sham orbital) $\psi_n^\alpha(\mathbf{r})$ in Ω_α [24]. For DFT calculation within each domain, we use a real-space approach based on finite differencing [28], where iterative solutions are accelerated using the multigrid preconditioning [29]. The multigrid is augmented with high-resolution grids that are adaptively generated near the atoms to accurately operate atomic pseudopotentials (figure 1) [30].

The DC-DFT algorithm has an algorithmic parameter that controls the data locality, i.e. the depth of a buffer layer that augments each domain to reduce the effect of artificial domain boundaries [23–25]. Figure 2(a) shows that the energy converges rapidly as a function of the localization parameter. The algorithm also conserves the total energy during MD simulation (figure 2(b)), in contrast to many $O(N)$ DFT algorithms that suffer energy-drift problems. We have also validated our DC-DFT based MD simulation against experimental data. For example, the calculated pair

distribution function of liquid rubidium in figure 2(c) agrees well with recent x-ray diffraction data [31].

The EDC framework has also been used to develop multiscale QM/MD simulation, in which DC-DFT calculation is embedded within MD simulation [32]. We use an additive hybridization scheme [33], in which the energy is a linear combination of MD and QM energies,

$$E_{\text{QM/MD}} = E_{\text{MD}}^{\text{system}} + \sum_\alpha (E_{\text{QM}}^\alpha - E_{\text{MD}}^\alpha). \quad (3)$$

Here, $E_{\text{MD}}^{\text{system}}$ is the MD energy of the entire system, whereas E_{QM}^α and E_{MD}^α are the QM and MD energies of the α th domain, respectively. The buffer layer scheme in DC-DFT makes the accuracy of QM/MD simulation insensitive to the location of embedded QM domains. This is essential to an adaptive QM/MD simulation, in which a DFT calculation is dynamically embedded only when and where high accuracy is required [8]. We have estimated the errors in our approach by performing QM/MD simulations with various sizes of the QM region. For example, the maximum errors in interatomic forces and atomic displacements are 0.01 and 0.05 au, respectively, for Si crystal involving 70-atom DFT calculation [34].

We have applied QM/MD simulations extensively to covalent systems [8, 35]. For metallic systems, a similar multiscale simulation approach has been developed, in which DFT calculation is embedded within an orbital-free DFT (OFDFT) calculation that is in turn embedded within MD simulation based on an embedded atom method (EAM) [36–38].

Starting with a pioneering work by Abraham *et al* [39, 40] multiscale simulations combining QM/MD and continuum calculations have widely been applied to materials

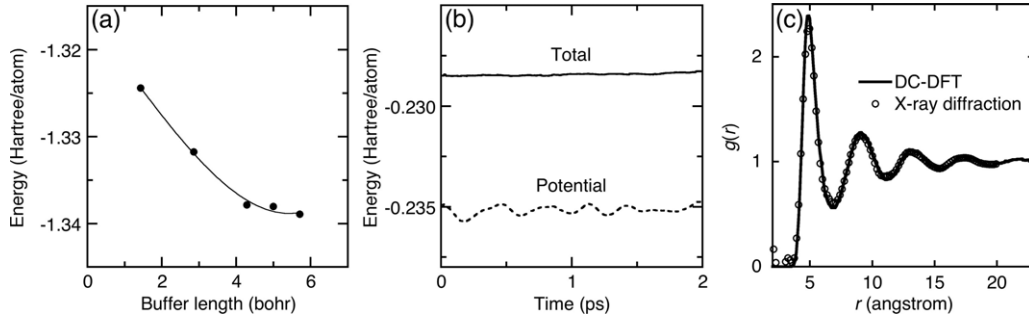


Figure 2. Validation of the DC-DFT algorithm. (a) Controlled convergence of the potential energy of amorphous CdSe by buffer depth (with the domain size fixed as 11.416 bohr). (b) Energy conservation in DC-DFT based MD simulation of liquid Rb. The domain size and buffer depth are 16.424 and 8.212 bohr, respectively. (c) Radial distribution function of liquid Rb at 350 K. The solid curve shows a 432-atom DC-DFT based MD result, which agrees well with recent x-ray diffraction data (open circles).

research [32]. Recent progress in this field is found in a review article [41]. A notable example is the ‘learning on the fly’ approach, in which the interatomic potential in MD simulation is adaptively refined during simulation to better fit QM calculation [42].

2.2. Space–time-ensemble parallelism (STEP) for predicting long-time dynamics

A challenging problem is to predict long-time dynamics because of the sequential bottleneck of time. Due to temporal locality, however, the system stays near local minimum-energy configurations most of the time, except for rare transitions between them. In such cases, the transition state theory (TST) allows the reformulation of the *sequential long-time dynamics* as computationally more efficient *parallel search* for low-activation-barrier transition events [43–45]. We also introduce a discrete abstraction based on graph data structures, so that combinatorial techniques can be used for the search [45]. We have developed a directionally heated nudged elastic band (DH-NEB) method [46], in which an NEB consisting of a sequence of S states [47, 48], $\mathbf{R}_s \in \mathfrak{R}^{3N}$ ($s = 0, \dots, S-1$, and \mathfrak{R} is the set of real numbers), at different temperatures searches for transition events (figure 3):

$$\mathbf{M} \frac{d^2}{dt^2} \mathbf{R}_s = \mathbf{F}_s - \mathbf{M} \gamma_s \frac{d}{dt} \mathbf{R}_s, \quad (4)$$

where $\mathbf{M} \in \mathfrak{R}^{3N \times 3N}$ is the diagonal mass matrix and γ_s is a friction coefficient. Here, the forces are defined as

$$\mathbf{F}_s = \begin{cases} \frac{\partial V}{\partial \mathbf{R}_s} \Big|_{\perp} + \mathbf{F}_s^{\text{spr}} \Big|_{\parallel} & (1 \leq s \leq S-2) \\ -\frac{\partial V}{\partial \mathbf{R}_s} & (s = 0, S-1), \end{cases} \quad (5)$$

where $V(\mathbf{R})$ is the interatomic potential energy, $\mathbf{F}_s^{\text{spr}}$ are spring forces that keep the states equidistant, and \perp and \parallel denote respectively the projections of a $3N$ -element vector perpendicular and parallel to the tangential vector connecting the consecutive states.

We use an ensemble consisting of B bands to perform long-time simulation in the framework of kinetic Monte Carlo




Figure 3. Schematic diagram of the directionally heated nudged elastic band (DH-NEB) method. (a) An NEB consists of a sequence of S states (gray parallelograms) R_s ($s = 0, \dots, S-1$), where each state consists of N atoms (spheres), $i = 1, \dots, N$. Corresponding atoms in consecutive states interact via harmonic forces represented by wavy lines. (b) Abstraction of an NEB consisting of S states (circles) connected by harmonic forces (lines). (c) Algorithmic steps of the DH-NEB method: thermalization, directional heating, and quench of a band. Black solid curves represent the potential energy surface $V(\mathbf{R})$, whereas circles (with gray-scaled temperature) are the states interconnected by harmonic forces (gray lines) to form the band. The letters i and f mark the initial and final ends of the band.

simulation [46, 49]. Here, our space–time-ensemble parallel (STEP) approach combines a hierarchy of concurrency, i.e., the number of processors is $P = \text{BSD}$: (1) spatial decomposition within each state (D is the number of spatial subsystems); (2) temporal parallelism across the states within each band; and (3) band-ensemble parallelism (figure 1).

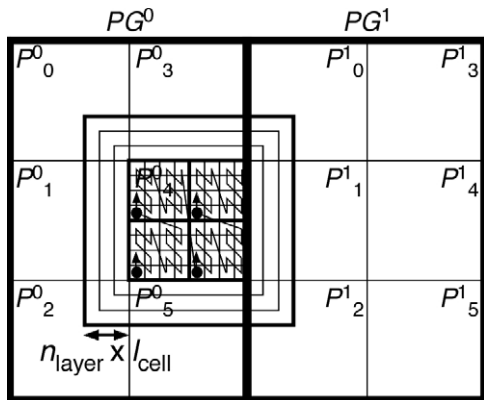


Figure 4. In tunable hierarchical cellular decomposition (HCD), the physical volume is subdivided into spatial subsystems that are assigned to process groups PG^y , each of which is spatially decomposed into processes P^y_π . Each process consists of a number of computational cells (e.g. linked-list cells in MD or domains in DC-DFT) of size l_{cell} , which are traversed concurrently by threads (denoted by dots with arrows) to compute blocks of cells. P^y_π is augmented with n_{layer} layers of cached cells from neighbor processes.

2.3. Tunable hierarchical cellular decomposition (HCD) for mapping metascalable algorithms onto parallel computers

To map the above $O(N)$ algorithms onto parallel computers, we have developed a tunable hierarchical cellular decomposition (HCD) framework (figure 4) [10]. Our parallelization is based on spatial decomposition, in which each spatial subsystem is assigned to a compute node in a parallel computer. At the finest level, EDC algorithms consist of computational cells, such as linked-list cells in MD and domains in DC-DFT. On a multicore–multiprocessor compute node, a block of cells is assigned to a thread. Furthermore, a group of spatial subsystems is assigned to each parallel computer on a Grid of parallel computers.

The hierarchy of computational cells provides an efficient mechanism for performance optimization—we make both the layout and size of the cells as tunable parameters that are optimized on each computing platform. Our EDC algorithms are implemented as hybrid MPI + OpenMP programs, in which the numbers of MPI processes and OpenMP threads are also tunable parameters. The HCD framework thus exposes data locality and concurrency. We are currently collaborating with compiler and artificial intelligence (AI) research groups to use (1) knowledge-representation techniques for expressing the exposed concurrency and (2) machine-learning techniques for optimally mapping the expressed concurrency to hardware [50].

Our parallelization framework includes load balancing capability. For irregular data structures, the number of atoms assigned to each processor varies significantly, and this load imbalance degrades the parallel efficiency. Load balancing can be stated as an optimization problem. We minimize the load imbalance cost as well the size and the number of messages. Our topology-preserving spatial decomposition allows message passing to be performed in a structured way in only six steps, so that the number of messages is minimized.

To minimize the load imbalance cost and the size of messages, we have developed a computational space decomposition scheme [51]. The main idea is that the computational space shrinks in a region with high workload density, so that the workload is uniformly distributed. The sum of load imbalance and communication costs is minimized as a functional of the computational space using simulated annealing. We have found that wavelets allow compact representation of curved partition boundaries and thus speed up the optimization procedure [52].

2.4. Locality in massive data visualization and analysis

Data locality is also utilized for massive dataset visualization. We use octree data structures to efficiently extract atoms within the field of view. We then use a probabilistic approach to remove far atoms that are occluded. We off-load these culling tasks to a parallel computer using a hierarchical distributed depth buffer [53], so that the graphics server is dedicated to the rendering of a reduced subset. As a result, we have succeeded in interactively visualizing a billion-atom dataset in an immersive virtual environment [54].

Our visualization system is combined with various data analysis tools. For example, we use a graph algorithm called shortest-path circuit analysis to identify and track topological defects in multimillion-vertex chemical bond networks, where atoms are vertices of a graph and bonds are its edges (figure 5) [55]. We have used this method to identify a dislocation network created by hypervelocity impact on ceramics (figure 5) [56–58].

2.5. Scalability test results

We have tested the scalability of our MD and QM algorithms on a number of parallel supercomputers such as IBM BlueGene/L, SGI Altix 3000, and AMD Opteron cluster [2]. Figure 6(a) shows that the execution times for the MRMD, F-ReaxFF, and DC-DFT algorithms scale linearly with the number of atoms on all three platforms. The largest benchmark tests include 201-billion-atom MRMD, 1.59 billion-atom F-ReaxFF MD, and 17.7-million-atom (or 1.56 trillion electronic degrees of freedom) DC-DFT based MD simulations. The EDC algorithms expose maximal data locality and consequently the parallel efficiency is as high as 0.99 on 196 608 processors.

We have also performed multiscale QM/MD simulation on a Grid of six supercomputer centers in the US and Japan [8]. To allow the number of processors to increase or decrease at runtime, we use a hybrid Grid remote procedure call + message passing interface Grid computing scheme. The simulation involved 150 000 processor hours, where the number of processors changed dynamically on demand; see figure 6(b). Furthermore, computations were automatically migrated between remote supercomputers due to schedules and unexpected faults. The Grid QM/MD simulation involved high-energy beam oxidation of silicon in the fabrication of low-power microprocessors.




Figure 5. Left, circuits for vertex x in a chemical bond network. Only the paths labeled a , b , and c constitute shortest-path circuits. Right, dislocations and brittle cracks propagating to the surface in a 300-million-atom MD simulation of SiC ceramics under hypervelocity projectile impact. Only atoms that participate in non-six-membered circuits are visualized. Here the color represents the pressure value. (A perfect crystal has only six-membered circuits.)

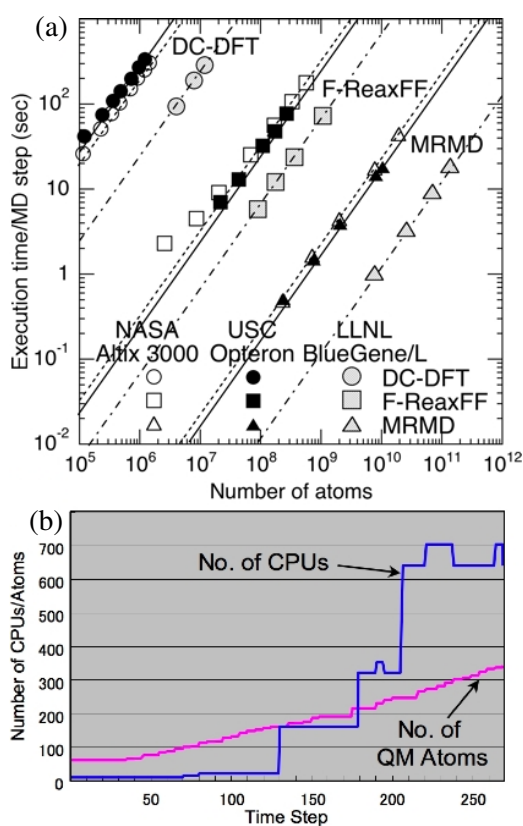


Figure 6. (a) Benchmark tests of MD and QM simulations on 1920 Itanium2 processors of Altix 3000 (open symbols), 2000 Opteron processors (solid symbols), and 196 608 BlueGene/L processors (shaded symbols). The execution time per MD step is shown as a function of the number of atoms for quantum mechanical MD based on the divide-and-conquer density functional theory (DC-DFT, circles); fast reactive force field MD (F-ReaxFF, squares); and space-time multiresolution MD (MRMD, triangles). Lines show ideal $O(N)$ scaling. (b) Time evolution of the number of QM atoms and the number of processors used for QM calculations in QM/MD simulation on a US-Japan Grid.

3. Nano-mechano-chemistry applications

We have applied the EDC-STEP-TCD simulation framework to study how atomistic processes determine material proper-

ties. Examples include nanoindentation on nanocomposite materials [59], oxidation of nanoenergetic materials [21], hypervelocity impact damage [56], fracture [60], and the interaction of voids and nanoductility [61]. We have also studied colloidal [62, 63] and epitaxial [64–66] quantum dots and rods for optoelectronic applications. In this section, we illustrate the use of chemically reactive MD simulations for the study of nano-mechano-chemical processes in energetic materials.

3.1. Combustion of nanoenergetic materials

Chemical reactions in energetic materials with nanometer-scale microstructures (or nanoenergetic materials) are very different from those in conventional energetic materials. For example, in conventional thermite materials made of aluminum and iron oxide, the combustion front propagates at a speed of $\sim\text{cm s}^{-1}$. In nanothermites of Al nanoparticles embedded in iron oxide, the combustion speed is accelerated to $\sim\text{km s}^{-1}$ [67]. Such rapid reactions cannot be explained by conventional diffusion based mechanisms.

We have performed DFT based *ab initio* MD simulation to study electronic processes during a thermite reaction. Here, the reactants are Al and Fe_2O_3 , and the products are Al_2O_3 and Fe (figure 7). The QM simulation allows quantitative study of reaction rates. In figure 7, bond-overlap population analysis shows that an oxygen atom changes its character from iron oxide type to aluminum oxide type within 0.3 ps, as it crosses the Al/ Fe_2O_3 interface. Here, the partial sum of the bond-overlap population (SBOP) $O_i^\alpha(t)$ ($\alpha = \text{Fe}$ or Al) for the i th oxygen atom is defined as

$$O_i^\alpha(t) = \sum_{j \in \alpha} O_{ij}(t), \quad (6)$$

where $O_{ij}(t)$ is the bond-overlap population between atoms i and j [68]. The total SBOP is calculated as $O_i(t) = O_i^{\text{Fe}}(t) + O_i^{\text{Al}}(t)$. We are currently performing reactive MD simulation of the flash heating of an oxidized Al nanoparticle in oxygen atmosphere to study how mechanical failure of the oxide shell causes rapid reactions.




Figure 7. DFT-based *ab initio* MD simulation of a thermite reaction. Initial (left) and final (center) configurations of atoms, where green, red, and gray spheres show the positions of Fe, O, and Al atoms, respectively. (Right) Time evolution of the total and partial bond-overlap populations, $O_i(t)$ and $O_i^\alpha(t)$, associated with an oxygen atom. The black, red, and blue curves show $O_i(t)$, $O_i^{\text{Fe}}(t)$, and $O_i^{\text{Al}}(t)$, respectively. Shown below the plot are atomic configurations near the oxygen atom of interest (pointed by yellow arrows) at different times.

3.2. Shock-induced initiation of energetic crystals

Initiation of an energetic crystal is highly sensitive to microstructures such as voids and grain boundaries. We have simulated shock initiation of an RDX crystal, which is a molecular crystal comprising $\text{C}_3\text{H}_6\text{O}_6\text{N}_6$ molecules [69]. To study the effect of microstructures, the crystal contains a void of diameter 8 nm [70]. In the shock simulation, an RDX crystal impacts a flat hard wall with a particle velocity V_p ($=1\text{--}5\text{ km s}^{-1}$), which generates a shock wave that propagates back with a velocity V_s . In figure 8(a), the color-coded velocities of RDX molecules show the formation of a hot spot at the void. We have also observed the formation of a nanojet and the acceleration of the nanojet velocity to $3V_p$. This may be due to the focusing of the nanojet illustrated in figure 8(a), where the arrows represent the molecular velocities.

We have also found that the nanojet-focusing catalyzes chemical reactions that do not occur otherwise [70]. Specifically, we observe two distinct reaction regimes as a function of time. Figure 8(b) shows the number of molecular fragments generated at the void surface as a function of time. Initially, as the void collapses, NO_2 fragments are produced. When the nanojet hits the downstream wall of the void at 2.6 ps, we instead observe the production of other molecules such as N_2 and H_2O .

4. Concluding remarks

In summary, we have developed high-end reactive atomistic simulation programs within common algorithmic and computational frameworks based on spatiotemporal data locality principles. They have enabled billion-atom simulations of mechano-chemical processes, which have applications in broad areas such as energy and environment.

An important issue is the timescale 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 span T . On petaflops computers,

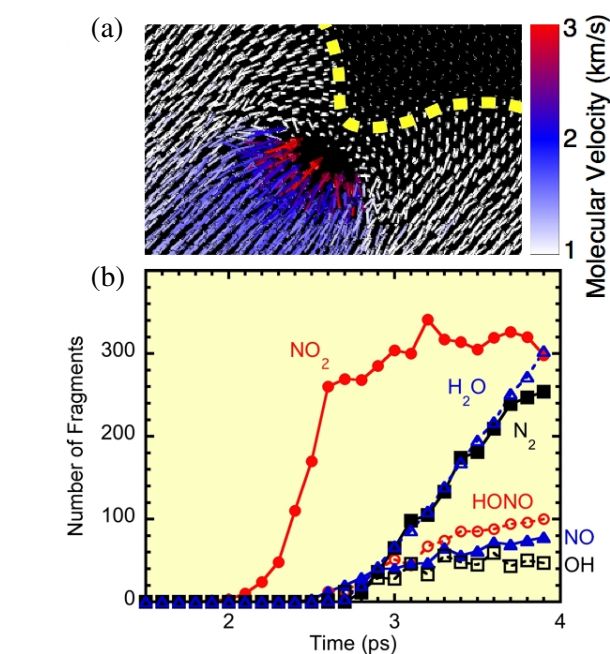


Figure 8. Shock initiation of an RDX crystal with a nanovoid. (a) A snapshot of the velocity distribution of RDX molecules around the void, where the arrows representing the velocities of RDX molecules are color coded by the magnitude of the velocity. The yellow dotted curve indicates the position of the shock front. (b) Number of molecular fragments near the void surface as a function of time. From the arrival of the shock wave until the void closure (~ 2.6 ps), a rapid production of NO_2 is observed. Shortly after this, when molecules strike the downstream wall of the void (2.6–3.9 ps), various chemical products such as N_2 , H_2O , and HONO are produced.

direct EDC MD simulations will be performed for $NT = 1\text{--}10$ atoms (i.e. multibillion-atom simulation for several nanoseconds or multimillion-atom simulation for several microseconds). STEP molecular-kinetics simulations will push the spatiotemporal envelope to $NT \sim 10$ and beyond, but they

need to be fully validated against MD simulations at $NT = 1-10$. Such large spatiotemporal scale atomistic simulations are expected to further advance scientific knowledge.

Acknowledgments

This work was partially supported by ARO MURIs, DTRA, DOE-BES/SciDAC, and NSF-ITR/PetaApps/CSR. FS acknowledges the support of a Grant-in-Aid for Scientific Research on Priority Area 'Nanoionics (439)' from MEXT, Japan. Simulations were performed at the University of Southern California using the 5384-processor Linux cluster at the Research Computing Facility and the 2048-processor Linux cluster at the Collaboratory for Advanced Computing and Simulations. We thank R Biswas, D Srivastava, and L H Yang for their help with scalability tests on the IBM BlueGene/L computer at the Lawrence Livermore National Laboratory and the SGI Altix 3000 computer at the NASA Ames Research Center.

References

- [1] Emmott S and Rison S 2006 *Towards 2020 Science* (Cambridge: Microsoft Research)
- [2] Nakano A *et al* 2008 *Int. J. High Perform. Comput. Appl.* **22** 113
- [3] Dongarra J J and Walker D W 2001 *Comput. Sci. Eng.* **3** 32
- [4] Asanovic K *et al* 2006 *The Landscape of Parallel Computing Research: A View From Berkeley* (Berkeley, CA: University of California)
- [5] <http://www.top500.org>
- [6] Foster I and Kesselman C 2003 *The Grid 2: Blueprint for a New Computing Infrastructure* (San Mateo, CA: Morgan Kaufmann University)
- [7] <http://www.teragrid.org>
- [8] Takemiya H, Tanaka Y, Sekiguchi S, Ogata S, Kalia R K, Nakano A and Vashishta P 2006 *Proc. Supercomput. 2006* (Tampa, FL: IEEE/ACM)
- [9] Dongarra J, Gannon D, Fox G and Kennedy K 2007 *CTWatch Q.* **3** 11
- [10] Nakano A *et al* 2007 *Comput. Mater. Sci.* **38** 642
- [11] van Duin A C T, Dasgupta S, Lorant F and Goddard W A 2001 *J. Phys. Chem. A* **105** 9396
- [12] Nakano A, Kalia R K, Vashishta P, Campbell T J, Ogata S, Shimojo F and Saini S 2002 *Sci. Prog.* **10** 263
- [13] Greengard L and Rokhlin V 1987 *J. Comput. Phys.* **73** 325
- [14] Nakano A, Kalia R K and Vashishta P 1994 *Comput. Phys. Commun.* **83** 197
- [15] Ogata S, Campbell T J, Kalia R K, Nakano A, Vashishta P and Vemparala S 2003 *Comput. Phys. Commun.* **153** 445
- [16] Martyna G J, Tuckerman M E, Tobias D J and Klein M L 1996 *Mol. Phys.* **87** 1117
- [17] Nakano A 1997 *Comput. Phys. Commun.* **105** 139
- [18] Nomura K, Kalia R K, Nakano A and Vashishta P 2008 *Comput. Phys. Commun.* **178** 73
- [19] Rappe A K and Goddard W A 1991 *J. Phys. Chem.* **95** 3358
- [20] Streitz F H and Mintmire J W 1994 *Phys. Rev. B* **50** 11996
- [21] Campbell T J, Kalia R K, Nakano A, Vashishta P, Ogata S and Rodgers S 1999 *Phys. Rev. Lett.* **82** 4866
- [22] Nakano A 1997 *Comput. Phys. Commun.* **104** 59
- [23] Yang W T 1991 *Phys. Rev. Lett.* **66** 1438
- [24] Shimojo F, Kalia R K, Nakano A and Vashishta P 2005 *Comput. Phys. Commun.* **167** 151
- [25] Ozaki T 2006 *Phys. Rev. B* **74** 245101
- [26] Kohn W and Sham L J 1965 *Phys. Rev.* **140** A1133
- [27] Kohn W and Vashishta P 1983 *Inhomogeneous Electron Gas* ed N H March and S Lundqvist (New York: Plenum) p 79
- [28] Chelikowsky J R, Troullier N, Wu K and Saad Y 1994 *Phys. Rev. B* **50** 11355
- [29] Briggs E L, Sullivan D J and Bernholc J 1996 *Phys. Rev. B* **54** 14362
- [30] Ono T and Hirose K 1999 *Phys. Rev. Lett.* **82** 5016
- [31] Matsuda K, Tamura K and Inui M 2007 *Phys. Rev. Lett.* **98** 096401
- [32] Ogata S, Lidorikis E, Shimojo F, Nakano A, Vashishta P and Kalia R K 2001 *Comput. Phys. Commun.* **138** 143
- [33] Dapprich S, Komáromi I, Byun K S, Morokuma K and Frisch M J 1999 *J. Mol. Struct. (Theochem)* **461/462** 1
- [34] Ogata S 2005 *Phys. Rev. B* **72** 045348
- [35] Ogata S, Shimojo F, Kalia R K, Nakano A and Vashishta P 2004 *J. Appl. Phys.* **95** 5316
- [36] Choly N, Lu G, E W and Kaxiras E 2005 *Phys. Rev. B* **71** 094101
- [37] Lu G, Tadmor E B and Kaxiras E 2006 *Phys. Rev. B* **73** 024108
- [38] Zhang X and Lu G 2007 *Phys. Rev. B* **76** 245111
- [39] Abraham F F, Broughton J Q, Bernstein N and Kaxiras E 1998 *Comput. Phys.* **12** 538
- [40] Broughton J Q, Abraham F F, Bernstein N and Kaxiras E 1999 *Phys. Rev. B* **60** 2391
- [41] Csanyi G, Albaret T, Moras G, Payne M C and De Vita A 2005 *J. Phys.: Condens. Matter* **17** R691
- [42] Csanyi G, Albaret T, Payne M C and De Vita A 2004 *Phys. Rev. Lett.* **93** 175503
- [43] Truhlar D G, Garrett B C and Klippenstein S J 1996 *J. Phys. Chem.* **100** 12771
- [44] Voter A F, Montalenti F and Germann T C 2002 *Annu. Rev. Mater. Res.* **32** 321
- [45] Nakano A 2007 *Comput. Phys. Commun.* **176** 292
- [46] Nakano A 2008 *Comput. Phys. Commun.* **178** 280
- [47] Henkelman G and Jonsson H 2000 *J. Chem. Phys.* **113** 9978
- [48] Jonsson H, Mills G and Jacobsen K 1998 *Classical and Quantum Mechanics in Condensed Phase Simulations* (Singapore: World Scientific)
- [49] Voter A F 2006 *Radiation Effects in Solids* ed K E Sickafus, E A Kotomin and B P Uberuaga (The Netherlands: Springer)
- [50] Bansal B *et al* 2007 *Proc. Next Generation Software Workshop, Int. Parallel and Distributed Processing Symp.* (Piscataway, NJ: IEEE)
- [51] Nakano A and Campbell T J 1997 *Parallel Comput.* **23** 1461
- [52] Nakano A 1999 *Concurrency: Pract. Exp.* **11** 343
- [53] Zhang C, Callaghan S, Jordan T, Kalia R K, Nakano A and Vashishta P 2007 *Int. J. Comput. Sci.* **1** 407
- [54] Sharma A *et al* 2003 *Presence-Teleoperators and Virtual Environ.* **12** 85
- [55] Zhang C, Bansal B, Branicio P S, Kalia R K, Nakano A, Sharma A and Vashishta P 2006 *Comput. Phys. Commun.* **175** 339
- [56] Branicio P S, Kalia R K, Nakano A and Vashishta P 2006 *Phys. Rev. Lett.* **96** 065502
- [57] Zhang C, Kalia R K, Nakano A and Vashishta P 2007 *Appl. Phys. Lett.* **91** 071906
- [58] Zhang C, Kalia R K, Nakano A and Vashishta P 2007 *Appl. Phys. Lett.* **91** 121911
- [59] Szlufarska I, Nakano A and Vashishta P 2005 *Science* **309** 911
- [60] Lu Z *et al* 2005 *Phys. Rev. Lett.* **95** 135501
- [61] Chen Y C, Lu Z, Nomura K I, Wang W, Kalia R K, Nakano A and Vashishta P 2007 *Phys. Rev. Lett.* **99** 155506
- [62] Kodiyalam S, Kalia R K, Nakano A and Vashishta P 2004 *Phys. Rev. Lett.* **93** 203401
- [63] Lee N J, Kalia R K, Nakano A and Vashishta P 2006 *Appl. Phys. Lett.* **89** 093101
- [64] Su X T, Kalia R K, Nakano A, Vashishta P and Madhukar A 2001 *Appl. Phys. Lett.* **79** 4577

- [65] Makeev M A, Kalia R K, Nakano A, Vashishta P and Madhukar A 2005 *J. Appl. Phys.* **98** 114313
- [66] Lidorikis E, Bachlechner M E, Kalia R K, Nakano A, Vashishta P and Voyiadjis G Z 2001 *Phys. Rev. Lett.* **87** 086104
- [67] Levitas V I, Asay B W, Son S F and Pantoya M 2006 *Appl. Phys. Lett.* **89** 071909
- [68] Umezawa N, Kalia R K, Nakano A, Vashista P and Shimojo F 2007 *J. Chem. Phys.* **126** 234702
- [69] Nomura K, Kalia R K, Nakano A, Vashishta P, van Duin A C T and Goddard W A 2007 *Phys. Rev. Lett.* **99** 148303
- [70] Nomura K, Kalia R K, Nakano A and Vashishta P 2007 *Appl. Phys. Lett.* **91** 183109