



An extended-Lagrangian scheme for charge equilibration in reactive molecular dynamics simulations



Ken-ichi Nomura, Patrick E. Small, Rajiv K. Kalia, Aiichiro Nakano*, Priya Vashishta

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

ARTICLE INFO

Article history:

Received 16 February 2015
 Accepted 21 February 2015
 Available online 2 March 2015

Keywords:

Reactive molecular dynamics
 Extended Lagrangian scheme
 Parallel computing

ABSTRACT

Reactive molecular dynamics (RMD) simulations describe chemical reactions at orders-of-magnitude faster computing speed compared with quantum molecular dynamics (QMD) simulations. A major computational bottleneck of RMD is charge-equilibration (QEq) calculation to describe charge transfer between atoms. Here, we eliminate the speed-limiting iterative minimization of the Coulombic energy in QEq calculation by adapting an extended-Lagrangian scheme that was recently proposed in the context of QMD simulations, Souvatzis and Niklasson (2014). The resulting XRMD simulation code drastically improves energy conservation compared with our previous RMD code, Nomura et al. (2008), while substantially reducing the time-to-solution. The XRMD code has been implemented on parallel computers based on spatial decomposition, achieving a weak-scaling parallel efficiency of 0.977 on 786,432 IBM Blue Gene/Q cores for a 67.6 billion-atom system.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Molecular dynamics (MD) simulations follow time evolution of the positions, $\mathbf{r}^N = \{\mathbf{r}_i \mid i = 1, \dots, N\}$, of N atoms by numerically integrating Newton's equations of motion, where the atomic force law is mathematically encoded in the interatomic potential energy $E(\mathbf{r}^N)$ [1]. Reliable interatomic potentials are key to accurately describing thermomechanical properties of materials. To describe material processes involving chemical reactions, in particular, quantum molecular dynamics (QMD) simulations [2–5] determine the force law by minimizing the potential, $E(\mathbf{r}^N, \psi^{N\text{el}})$, as a functional of electronic wave functions $\psi^{N\text{el}}(\mathbf{r}) = \{\psi_n(\mathbf{r}) \mid n = 1, \dots, N_{\text{el}}\}$ (N_{el} is the number of wave functions) usually in the framework of density functional theory (DFT) [6]. Despite remarkable progresses in $O(N)$ DFT algorithms [7–10], the largest QMD simulations to date have been limited to $N \sim 10^4$ for the duration of 10^{-11} s [11,12]. To extend the limited spatiotemporal scales covered by QMD simulations, reactive molecular dynamics

(RMD) simulation methods [13] have been developed based on a first principles-based reactive force-field (ReaxFF) approach [14, 15]. The ReaxFF approach significantly reduces the computational cost of simulating chemical reactions, while reproducing the energy surfaces and barriers as well as charge distributions of DFT calculations. RMD simulations describe formation and breakage of chemical bonds using reactive bond orders [14,16,17]. The most intensive computation in RMD simulations arises from a charge-equilibration (QEq) scheme [18–20] to describe charge transfer between atoms. QEq treats variable atomic charges as dynamic variables, $q^N = \{q_i \mid i = 1, \dots, N\}$. The charges and the resulting force law are determined by minimizing the potential, $E(\mathbf{r}^N, q^N)$, with respect to q^N at every MD step. This variable N -charge problem is commonly solved iteratively with the conjugate gradient (CG) method [21,22]. Though recent advancements in parallel ReaxFF algorithms [23–25] have enabled large RMD simulations [26–28] involving multimillion atoms, QEq computation remains to be the major bottleneck toward achieving billion-atom RMD simulations based on ReaxFF. The problem is that an excessively large number of CG iterations are required to reach sufficient convergence of charges q^N to guarantee the conservation of the total energy as a function of time. Insufficiently converged charges act as a heat sink of energy, and the resulting broken time reversibility causes the total energy to drift over time. A similar

* Corresponding author at: Collaboratory for Advanced Computing and Simulations, University of Southern California, Los Angeles, CA 90089-0242, USA. Tel.: +1 213-821-2657; fax: +1 213-821-2664.

E-mail address: anakano@usc.edu (A. Nakano).

trade-off between the computational speed and energy conservation is encountered in QMD simulations, where insufficient convergence of the iterative refinement of wave functions ψ^{Nel} causes serious energy drift. Niklasson proposed an extended Lagrangian scheme [29–31] that achieves excellent long-time energy conservation with drastically reduced number of iterations. In fact, an extended Lagrangian scheme with no iteration (i.e. requiring only one evaluation of energy gradient) has recently been demonstrated [32]. The key idea is to introduce auxiliary wave functions as dynamic variables that are numerically integrated by reversible, symplectic integration schemes to address the broken reversibility problem, while the auxiliary wave functions are constrained to iteratively determined wave functions by a harmonic potential.

In this paper, the extended Lagrangian scheme [29–32] is adapted to RMD simulations, thereby eliminating speed-limiting iterations in QEq calculation. The resulting XRMD simulation code drastically improves the energy conservation compared with our previous RMD code [23], while substantially reducing the time-to-solution. This paper presents key features and implementation details of XRMD. The rest of the paper is organized as follows. The next section describes the computational method. Benchmark results are presented in Section 3, and Section 4 contains conclusions.

2. Methods

2.1. Charge-equilibration (QEq) method

The interatomic potential energy in ReaxFF is composed of a number of bonded and nonbonded terms [14,15]. The nonbonded terms are Coulombic and van der Waals energies. The Coulombic energy is expressed as

$$E_{\text{Coulomb}}(\mathbf{r}^N, q^N) = \sum_i \chi_i q_i + \frac{1}{2} \sum_i \sum_j q_i H(r_{ij}) q_j, \quad (1)$$

where \mathbf{r}_i , χ_i , and q_i are the position, electronegativity, and charge of the i th atom, respectively, and r_{ij} is the distance between a pair of atoms i and j . In Eq. (1), the Coulombic interaction is defined as

$$H(r_{ij}) = J_i \delta_{ij} + \frac{T(r_{ij})}{[r_{ij}^3 + \gamma_{ij}^{-3}]^{1/3}} (1 - \delta_{ij}), \quad (2)$$

where J_i is the self-Coulomb repulsion coefficient, γ_{ij} is a parameter for the smeared Coulombic function, and the Kronecker delta is $\delta_{ij} = 1$ ($i = j$) or 0 ($i \neq j$). The Coulombic interaction is screened using a taper function, $T(r)$, which has a finite range with a cutoff length of r_c .

In ReaxFF, atomic charges q_i are variables that change dynamically in time. When atomic positions are updated during RMD simulation, the QEq subroutine updates charge distribution q^N by minimizing E_{Coulomb} subject to charge-neutrality constraint, $\sum_i q_i = 0$. With the Lagrange-multiplier method, the constrained energy minimization is equivalent to solving the electronegativity equalization problem [18,22,33,19],

$$g_i \equiv -\frac{\partial E_{\text{Coulomb}}}{\partial q_i} = -\mu, \quad (3)$$

where μ is the electrochemical potential. We solve this problem iteratively using the CG method [21,23,34].

2.2. Extended Lagrangian QEq scheme

Our extended Lagrangian QEq scheme introduces auxiliary variables, $\theta^N = \{\theta_i \mid i = 1, \dots, N\}$. System dynamics for extended Lagrangian reactive molecular dynamics (XRMD) simulations is derived from the extended Lagrangian,

$$L_{\text{XRMD}} = L_{\text{RMD}} + \frac{\mu}{2} \sum_i \dot{\theta}_i^2 - \frac{\mu\omega^2}{2} \sum_i (\theta_i - q_i)^2, \quad (4)$$

Table 1

Energy deviation during MD simulation for 1 ps.

Method	Energy deviation (10^{-3} kcal/mol/atom)
XRMD	6.8
RMD (CG tolerance: 10^{-8})	4.1
RMD (one CG iteration per MD step)	499

where L_{RMD} is the RMD Lagrangian,

$$L_{\text{MD}} = \sum_i \frac{m_i}{2} \dot{\mathbf{r}}_i^2 - E(\mathbf{r}^N), \quad (5)$$

with m_i being the mass of the i th atom and the dot denoting time derivative. In Eq. (4), μ and ω are fictitious mass and frequency parameters for the auxiliary charge degrees of freedom.

The time evolution of the dynamical system described by L_{XRMD} is determined by Euler–Lagrange equations of motion. In the limit of $\mu \rightarrow 0$, the equations become

$$m_i \ddot{\mathbf{r}}_i = -\frac{\partial}{\partial \mathbf{r}_i} E(\mathbf{r}^N), \quad (6)$$

$$\ddot{\theta}_i = \omega^2 (q_i - \theta_i). \quad (7)$$

We numerically integrate Eqs. (6) and (7) using the velocity Verlet algorithm with a unit time step of δt . At each time step, we perform just one step of CG iteration toward minimizing the Coulombic energy, Eq. (1), as a function of q^N using the gradient in Eq. (3). We use the auxiliary variable θ^N as an initial guess for the CG optimization of q^N . For extended Lagrangian-based QMD simulations, the recommended value for the dimensionless parameter, $K = \omega^2 \delta t^2$, is 2 [31]. We have found that the same value works for XRMD as well, and $K = 2$ will be used in the numerical tests discussed below.

We have implemented the extended Lagrangian QEq scheme in our scalable parallel RMD simulation code [23]. All XRMD computations are parallelized using spatial decomposition, where the simulated system is decomposed into spatially localized subsystems and each processor is assigned computations associated with one subsystem. Message passing is used to exchange necessary data for the computations utilizing the message passing interface (MPI) library. The XRMD program is written in Fortran 90.

3. Results

We test our extended Lagrangian QEq scheme and the parallel XRMD code using oxidation of a silicon carbide nanoparticle (n-SiC) as an example. We adopt ReaxFF parameterization by Newsome et al. [35,36]. A n-SiC composed of 25 silicon (Si) and 25 carbon (C) atoms is placed in oxygen environment. We place 50 oxygen (O_2) molecules randomly around the n-SiC using a Monte Carlo procedure. The total number of atoms is 150 in a cubic box of side 16.289 Å. Fig. 1(a) shows the initial configuration of the n-SiC + O_2 system. The equations of motion are integrated using a unit time step of 0.25 fs. We first thermalize the system at temperature 300 K by velocity scaling. Subsequently, we switch to the microcanonical ensemble to test the energy conservation. Single MPI process is used for this test.

Fig. 1(b) shows the total energy as a function of time during MD simulation. We compare three cases: (1) XRMD code; (2) RMD code, where CG iterations at each MD step are continued until the change in the Coulombic energy E_{Coulomb} per iteration falls below 10^{-8} of $|E_{\text{Coulomb}}|$; and (3) RMD code performing only one CG iteration per MD step. Total energy conservation of XRMD is comparable to converged RMD, while RMD with single CG step exhibits significant energy drift. Table 1 summarizes total energy drift during 1 ps MD simulation with the three methods.

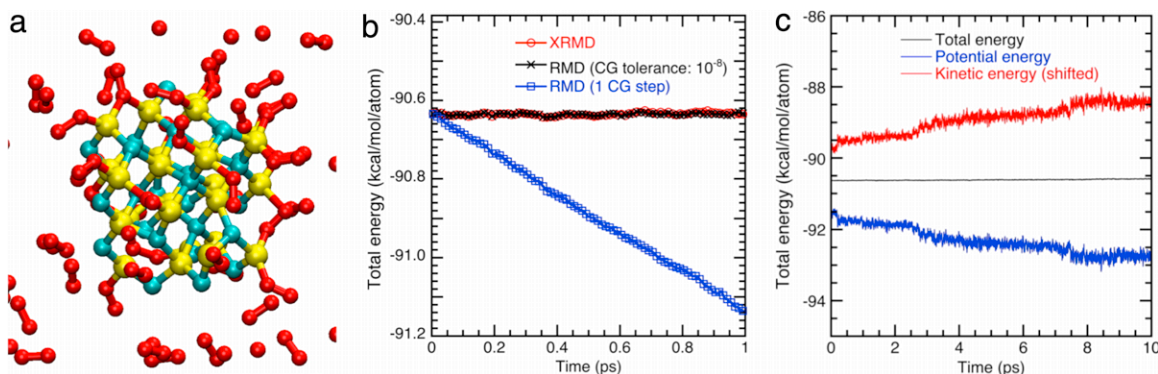


Fig. 1. (a) A snapshot of the simulated n-SiC surrounded by O₂ molecules. Silicon, carbon, and oxygen atoms are colored cyan, yellow, and red, respectively. (b) Total energy as a function of time for XRMD code compared with two RMD results: with (i) converged CG iterations with a tolerance of 10⁻⁸; and (ii) one CG iteration per MD step. (c) Total, potential, and kinetic energies obtained with XRMD code. The kinetic energy is shifted by a constant value (i.e., the total energy at time 0). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 2. Total energy of converged RMD and that of XRMD during 1 ps.

Fig. 1(c) plots time evolutions of the total, potential and kinetic energies obtained with XRMD code. The kinetic energy is shifted by the total energy at time $t = 0$ to better compare the changes of the three energy terms. It should be noted that the n-SiC + O₂ system is highly reactive, where the surface of n-SiC is continuously oxidized during simulation. These exothermic reactions release a large amount of heat, resulting in decreased potential energy and increased kinetic energy. Compared with the changes in the potential and kinetic energies, the total energy is well conserved with the XRMD method.

To magnify small deviations in the total energy for both XRMD and converged RMD simulations, Fig. 2 shows time evolution of the total energy over 1 ps. The two energy values are nearly identical up to 500 fs and subsequently exhibit very small difference even at this fine scale.

To validate atomic trajectories obtained by XRMD, Fig. 3 plots the difference of atomic coordinates after 1 ps of microcanonical MD simulation between XRMD and converged RMD with CG tolerance of 10⁻⁸. In the figure, the horizontal axis is atom ID $\in \{1, 2, \dots, 150\}$. The difference in atomic coordinates between the two methods mostly falls within ± 0.1 (Å) except for a few atoms.

Fig. 4 plots the difference of atomic charges after 1 ps of microcanonical MD simulation between XRMD and converged RMD methods. The largest difference in atomic charge is 0.081e for one of the carbon atoms on the n-SiC surface. The simulated system undergoes extensive reduction and oxidation reactions with associated large changes in atomic charges. The figure shows that charge dynamics is accurately described by XRMD despite much less computing time compared with fully converged RMD. To quantify the overall accuracy of XRMD, Table 2 summarizes the average and

Table 2

Average and standard deviation of the absolute difference between XRMD and converged RMD methods.

	X (Å)	Y (Å)	Z (Å)	Charge (e)
Average	0.084	0.080	0.074	0.0092
Standard deviation	0.112	0.116	0.104	0.0134

standard deviation of the absolute difference in atomic coordinates and charges obtained with XRMD compared with those in converged RMD with CG tolerance of 10⁻⁸.

As demonstrated above, XRMD provides sufficient accuracy that is comparable to fully converged RMD but at much less computational cost. To test the difference in computing time in realistic simulations, we perform benchmark tests using one of our production simulations. The system is a n-SiC of diameter 460 Å in oxygen environment with the total number of atoms, 10,007,652 (Fig. 5). We use 1280 MPI processes that are assigned to $20 \times 8 \times 8$ spatial sub-systems in the x, y and z directions. Two sets of benchmark tests are performed on different platforms: (1) dual quad-core AMD 2.3 GHz Optron central processing units (CPUs) with 10 Gbit/s Myrinet interconnect; and (2) dual octa-core Intel Xeon E5-2665 2.4 GHz CPUs with 56 Gbit/s Infiniband interconnect. We use Intel compiler version 14.0 linked with OpenMPI version 1.8.1. The identical executable is used for both measurements. 1280 MPI processes are run on 107 computing nodes on Myrinet and 80 nodes on Infiniband. We measure timing of the QEq subroutine over 5 MD steps. Table 3 summarizes the measured wall-clock time of the four benchmarks (i.e., XRMD and RMD with CG tolerance of 10⁻⁸ on both platforms). Computation time of XRMD is 8.6 times or 4.4 times smaller than that of RMD, respectively, on the Myrinet and Infiniband platforms. QEq computation involves heavy communication in order to exchange atomic charges within a cut off radius of $r_c \sim 10$ Å from subsystem boundaries at every CG iteration. The latter platform has much faster network interconnect, and accordingly, it penalizes extensive QEq computations in RMD less severely.

Next, we test the scalability of the parallel XRMD code on massively parallel settings using RDX crystal as an example. Numerical tests are performed on the IBM Blue Gene/Q computer, Mira, at the Argonne National Laboratory. Mira consists of 48 racks each with 1024 nodes. Each node has a 16-core processor operating at 1.6 GHz for computation, 32 MB of L2 cache with hardware transactional memory and speculative operation functionality [12,37,38], and 16 GB DDR3 main RAM. The processor employs PowerPC A2 architecture that supports quad floating operation units, 16 kB of L1 instruction and data cache, and 4-way multithreading per core enabling 64 concurrent threads on one node. Though it is highly energy efficient (55 Watts per node) thanks to the relatively low




Fig. 3. Difference in atomic coordinates after 1 ps of MD simulation between XRMD and converged RMD methods in (a) x, (b) y, and (c) z directions.

Table 3
Wall-clock time for 5 MD steps of 10 million-atom simulation using 1280 MPI processes. XRMD and fully converged RMD simulations are run on both Myrinet and Infiniband platforms.

Network	Number of nodes	Wall-clock time (s)	
		Converged RMD (CG tolerance 10^{-8})	XRMD
10 Gbit/s Myrinet	107	69.9	8.09
56 Gbit/s Infiniband	80	13.4	3.03




Fig. 4. Atomic charges after 1 ps of MD simulation with converged RMD and XRMD methods.

clock speed, the Blue Gene/Q chip delivers a peak performance of 204.8 Gflop/s. Each node has 11 links—10 links to connect computing nodes and one link to I/O node. Each link can simultaneously transmit and receive data at 2 GB/s, amounting to a total bandwidth of 44 GB/s. A 5-dimensional torus network is used for peer-to-peer communications.

We first perform an isogranular-scaling test of XRMD code on the Blue Gene/Q, in which the number of atoms per core N/P is kept constant. Fig. 6 shows the wall-clock time per XRMD simulation step with scaled workloads—86,016 P -atom RDX system on P cores of Blue Gene/Q. By increasing the number of atoms linearly with the number of cores, the wall-clock time remains almost constant, indicating excellent scalability. To quantify the parallel efficiency, we first define the speed of the XRMD code as a product of the total number of atoms and the number of MD steps executed per second. The isogranular speedup is given by the ratio between the speed of P cores and that of 8192 cores as a reference system. The weak-scaling parallel efficiency is the isogranular speedup divided by $P/8192$. With the granularity of 86,016 atoms per core, the parallel efficiency is 0.977 on $P = 786,432$ for a 67,645,734,912-atom system. This demonstrates a very high scalability of the XRMD code.

We also perform a strong-scaling test by simulating a RDX crystal containing a total of 4,227,858,432 atoms. In this test,




Fig. 5. Snapshot of 10 million-atom simulation. The color code is the same as that for Fig. 1(a). O₂ molecules are not shown for the clarity of presentation. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the number of cores ranges from $P = 49,152$ to 786,432, while keeping the total problem size constant. Fig. 7 shows the wall-clock time per MD simulation step as a function of P . The time-to-solution is reduced by a factor of 14.17 on 786,432 cores compared with the 49,152-core run (i.e., using 16-times larger number of cores). This signifies a strong-scaling speedup of 14.17, with the corresponding strong-scaling parallel efficiency of 0.886. It is more difficult to achieve high strong-scaling parallel efficiency compared with weak-scaling parallel efficiency. This is due to decreasing granularity, and accordingly increasing communication/computation ratio for larger number of processors, in the former. With 16 times smaller system size of the weak-scaling test, the observed strong-scaling parallel efficiency is considered excellent.

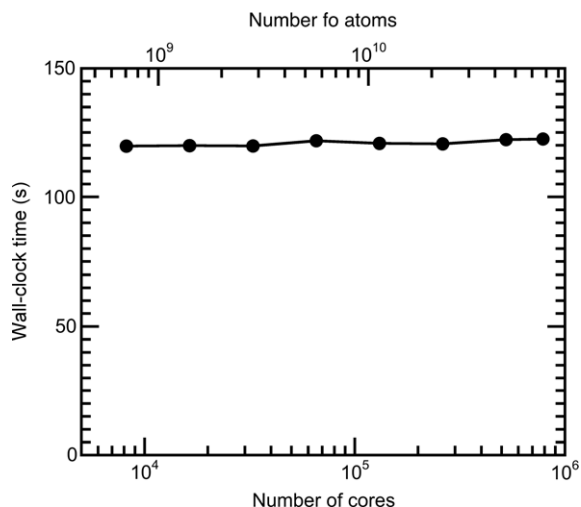


Fig. 6. Wall-clock time per MD simulation step of parallel XRMD code, with scaled workloads—86,016 P -atom RDX system on P cores ($P = 8192, \dots, 786,432$) of Blue Gene/Q.




Fig. 7. Wall-clock time per MD simulation step of parallel XRMD code with strong scaling—4,227,858,432-atom RDX system on P cores ($P = 49,152, \dots, 786,432$) of Blue Gene/Q.

4. Conclusion

We have eliminated speed-limiting CG iterations in QEq calculation for RMD simulations by adapting the extended Lagrangian scheme that has recently been developed for QMD simulations. The resulting XRMD code has drastically improved the energy conservation compared with our precious RMD code, while substantially reducing the time-to-solution. The XRMD code has also demonstrated high parallel efficiency, 0.977 on 786,432 Blue Gene/Q processors. These results demonstrate a wide applicability of the extended Lagrangian scheme beyond the original QMD domain.

Acknowledgments

Development of the simulation method was supported by the Department of Energy, Office of Science, Basic Energy Sciences, Materials Science and Engineering Division, Grant No. DE-FG02-04ER-46130. Scalable implementation of the XRMD method on massively parallel computers was supported by the National Science Foundation, Cyber-enabled Discovery and Innovation Program, Grant No. 1125317. Numerical tests were performed at the Center for High Performance Computing of the University of Southern California. Scalability tests were performed on the 786,432-core Blue Gene/Q computer at the Argonne Leadership Computing

Facility under a DOE INCITE project. We thank Paul Messina and Nichols Romero for their help on the use of the Blue Gene/Q.

References

- [1] A. Rahman, Correlations in the motion of atoms in liquid argon, *Phys. Rev.* 136 (1964) A405–A411.
- [2] R. Car, M. Parrinello, Unified approach for molecular-dynamics and density-functional theory, *Phys. Rev. Lett.* 55 (1985) 2471–2474.
- [3] M.C. Payne, M.P. Teter, D.C. Allan, T.A. Arias, J.D. Joannopoulos, Iterative minimization techniques for ab initio total-energy calculations—molecular-dynamics and conjugate gradients, *Rev. Modern Phys.* 64 (1992) 1045–1097.
- [4] F. Shimojo, S. Ohmura, R.K. Kalia, A. Nakano, P. Vashishta, Molecular dynamics simulations of rapid hydrogen production from water using aluminum clusters as catalyzers, *Phys. Rev. Lett.* 104 (2010) 126102.
- [5] K. Shimamura, F. Shimojo, R.K. Kalia, A. Nakano, P. Vashishta, Bonding and structure of ceramic–ceramic interfaces, *Phys. Rev. Lett.* 111 (2013) 066103.
- [6] P. Hohenberg, W. Kohn, Inhomogeneous electron gas, *Phys. Rev.* 136 (1964) B864–B871.
- [7] W.T. Yang, Direct calculation of electron-density in density-functional theory, *Phys. Rev. Lett.* 66 (1991) 1438–1441.
- [8] S. Goedecker, Linear scaling electronic structure methods, *Rev. Modern Phys.* 71 (1999) 1085–1123.
- [9] D.R. Bowler, T. Miyazaki, $O(N)$ methods in electronic structure calculations, *Rep. Progr. Phys.* 75 (2012) 036503.
- [10] F. Shimojo, R.K. Kalia, M. Kunaseth, A. Nakano, K. Nomura, S. Ohmura, K. Shimamura, P. Vashishta, A divide-conquer-recombine algorithmic paradigm for multiscale materials modeling, *J. Chem. Phys.* 140 (2014) 18A529.
- [11] K. Shimamura, F. Shimojo, R.K. Kalia, A. Nakano, K. Nomura, P. Vashishta, Hydrogen-on-demand using metallic alloy nanoparticles in water, *Nano Lett.* 14 (2014) 4090–4096.
- [12] K. Nomura, R.K. Kalia, A. Nakano, P. Vashishta, K. Shimamura, F. Shimojo, M. Kunaseth, P.C. Messina, N.A. Romero, Metascale quantum molecular dynamics simulations of hydrogen-on-demand, in: *Proc SC14, IEEE/ACM*, 2014, pp. 661–673.
- [13] S.B. Sinnott, D.W. Brenner, Three decades of many-body potentials in materials research, *MRS Bull.* 37 (2012) 469–473.
- [14] A.C.T. van Duin, S. Dasgupta, F. Lorant, W.A. Goddard, ReaxFF: a reactive force field for hydrocarbons, *J. Phys. Chem. A* 105 (2001) 9396–9409.
- [15] A. Nakano, R.K. Kalia, K. Nomura, A. Sharma, P. Vashishta, F. Shimojo, A.C.T. van Duin, W.A. Goddard, R. Biswas, D. Srivastava, L.H. Yang, De novo ultrascale atomistic simulations on high-end parallel supercomputers, *Int. J. High Perform. Comput. Appl.* 22 (2008) 113–128.
- [16] J. Tersoff, Modeling solid-state chemistry: interatomic potentials for multi-component systems, *Phys. Rev. B* 39 (1989) 5566–5568.
- [17] D.W. Brenner, Empirical potential for hydrocarbons for use in simulating the chemical vapor-deposition of diamond films, *Phys. Rev. B* 42 (1990) 9458–9471.
- [18] A.K. Rappe, W.A. Goddard, Charge equilibration for molecular-dynamics simulations, *J. Phys. Chem.* 95 (1991) 3358–3363.
- [19] F.H. Streitz, J.W. Mintmire, Electrostatic potentials for metal–oxide surfaces and interfaces, *Phys. Rev. B* 50 (1994) 11996–12003.
- [20] S.W. Rick, S.J. Stuart, B.J. Berne, Dynamical fluctuating charge force-fields—application to liquid water, *J. Chem. Phys.* 101 (1994) 6141–6156.
- [21] A. Nakano, Parallel multilevel preconditioned conjugate-gradient approach to variable-charge molecular dynamics, *Comput. Phys. Comm.* 104 (1997) 59–69.
- [22] T.J. Campbell, R.K. Kalia, A. Nakano, P. Vashishta, S. Ogata, S. Rodgers, Dynamics of oxidation of aluminum nanoclusters using variable charge molecular-dynamics simulations on parallel computers, *Phys. Rev. Lett.* 82 (1999) 4866–4869.
- [23] K. Nomura, R.K. Kalia, A. Nakano, P. Vashishta, A scalable parallel algorithm for large-scale reactive force-field molecular dynamics simulations, *Comput. Phys. Comm.* 178 (2008) 73–87.
- [24] H.M. Aktulga, J.C. Fogarty, S.A. Pandit, A.Y. Grama, Parallel reactive molecular dynamics: numerical methods and algorithmic techniques, *Parallel Comput.* 38 (2012) 245–259.
- [25] H.M. Aktulga, S.A. Pandit, A.C.T. van Duin, A.Y. Grama, Reactive molecular dynamics: numerical methods and algorithmic techniques, *SIAM J. Sci. Comput.* 34 (2012) C1–C23.
- [26] K. Nomura, R.K. Kalia, A. Nakano, P. Vashishta, A.C.T. van Duin, W.A. Goddard, Dynamic transition in the structure of an energetic crystal during chemical reactions at shock front prior to detonation, *Phys. Rev. Lett.* 99 (2007) 148303.
- [27] H. Chen, R.K. Kalia, E. Kaxiras, G. Lu, A. Nakano, K. Nomura, A.C.T. van Duin, P. Vashishta, Z. Yuan, Embrittlement of metal by solute segregation-induced amorphization, *Phys. Rev. Lett.* 104 (2010) 155502.
- [28] M. Vedadi, A. Choubey, K. Nomura, R.K. Kalia, A. Nakano, P. Vashishta, A.C.T. van Duin, Structure and dynamics of shock-induced nanobubble collapse in water, *Phys. Rev. Lett.* 105 (2010) 014503.
- [29] A.M.N. Niklasson, Extended Born–Oppenheimer molecular dynamics, *Phys. Rev. Lett.* 100 (2008) 123004.
- [30] M.J. Cawkwell, A.M.N. Niklasson, Energy conserving, linear scaling Born–Oppenheimer molecular dynamics, *J. Chem. Phys.* 137 (2012) 134105.
- [31] A.M.N. Niklasson, M.J. Cawkwell, Generalized extended Lagrangian Born–Oppenheimer molecular dynamics, *J. Chem. Phys.* 141 (2014) 164123.

- [32] P. Souvatzis, A.M.N. Niklasson, First principles molecular dynamics without self-consistent field optimization, *J. Chem. Phys.* 140 (2014) 044117.
- [33] W.J. Mortier, S.K. Ghosh, S. Shankar, Electronegativity equalization method for the calculation of atomic charges in molecules, *J. Am. Chem. Soc.* 108 (1986) 4315–4320.
- [34] B. Bayraktar, T. Bernas, J.P. Robinson, B. Rajwa, A numerical recipe for accurate image reconstruction from discrete orthogonal moments, *Pattern Recognit.* 40 (2007) 659–669.
- [35] D.A. Newsome, D. Sengupta, H. Foroutan, M. Francis Russo, A.C.T. van Duin, Oxidation of silicon carbide by O₂ and H₂O: a ReaxFF reactive molecular dynamics study: Part I, *J. Phys. Chem. C* 116 (2012) 16111–16121.
- [36] D.A. Newsome, D. Sengupta, A.C.T. van Duin, High-temperature oxidation of SiC-based composite: rate constant calculation from ReaxFF MD simulations, part II, *J. Phys. Chem. C* 117 (2013) 5014–5027.
- [37] M. Kunaseth, R.K. Kalia, A. Nakano, P. Vashishta, D.F. Richards, J.N. Glosli, Performance characteristics of hardware transactional memory for molecular dynamics application on Blue Gene/Q: toward efficient multithreading strategies for large-scale scientific applications, in: *Proc PDSEC-13, IEEE*, 2013.
- [38] M. Kunaseth, R.K. Kalia, A. Nakano, K. Nomura, P. Vashishta, A scalable parallel algorithm for dynamic range-limited n-tuple computation in many-body molecular dynamics simulation, in: *Proc Supercomputing, SC13, ACM/IEEE*, 2013.