Divide-Conquer-Recombine: An Algorithmic Pathway toward Metascalability

Aiichiro Nakano,[a] Shinnosuke Hattori,[a] Rajiv K. Kalia,[a] Weiwei Mou,[a] Ken-ichi Nomura,[a] Pankaj Rajak,[a] Priya Vashishta,[a] Kohei Shimamura,[b] Fuyuki Shimojo,[b] Manaschai Kunaseth,[c] Satoshi Ohmura,[d] Paul C. Messina,[e] and Nichols A. Romero[e]

[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] National Nanotechnology Center (NANOTEC), Thailand Science Park, Pathumthani 12120, Thailand

[d] Faculty of Engineering, Hiroshima Institute of Technology, Hiroshima 731-5193, Japan

[e] Argonne Leadership Computing Facility, Argonne National Laboratory, Argonne, IL 60439, USA

(anakano, sh_423, rkalia, wmou, knomura, rajak, priyav)@usc.edu, (143d9003, shimojo)@kumamoto-u.ac.jp, manaschai@nanotec.or.th, s.ohmura.m4@cc.it-hiroshima.ac.jp, (messina, naromero)@alcf.anl.gov

ABSTRACT

We propose an extension of the divide-and-conquer (DC) algorithmic paradigm called divide-conquer-recombine (DCR) in order to develop $O(N)$ applications that will continue to scale on future parallel supercomputing, *i.e.*, making them metascalable (or "design once, scale on new architectures"). In DCR, the DC phase constructs globally informed local solutions, which in the recombine phase are synthesized into a global solution. Innovative recombination algorithms allow the synthesis of a variety of global properties in broad applications. To enable large spatial-scale molecular dynamics (MD) simulations, DCR-in-space is empowered by globally scalable and locally fast (GSLF) hybrid solvers based on spatial locality. In addition, DCR-in-time is used to predict long-time dynamics based on temporal locality, while utilizing space-time-ensemble parallelism (STEP). We have used DCR to perform quantum molecular dynamics (QMD) and reactive molecular dynamics (RMD) simulations that encompass unprecedented spatiotemporal scales. Our 50.3 million-atom QMD benchmark achieved a parallel efficiency of 0.984 and 50.5% of the peak floating-point performance on 786,432 IBM Blue Gene/Q cores. Production QMD simulation involving 16,661 atoms for 21,140 time steps (or 129,208 self-consistent-field iterations) revealed a novel nanostructural design for on-demand hydrogen production from water, advancing renewable energy technologies. Nonadiabatic QMD simulation of photoexcitation dynamics involving 6,400 atoms reached the experimental time scales, elucidating molecular mechanisms of a novel singlet-fission phenomenon to realize low-cost, high-efficiency solar cells. Our billion-atom RMD simulation revealed the role of focused nanojet for the damage of solid surface caused by shock-induced collapse of nanobubbles in water, and suggested how to mitigate the damage by filling the bubble with inert gas.

Beowulf '14, October 13 - 14, 2014, Annapolis, MD, USA
Copyright 2014 ACM 978-1-4503-3031-2/14/10…$15.00
http://dx.doi.org/10.1145/2737909.2737911

1. INTRODUCTION

There is a growing need for scalable applications to keep pace with the ever-increasing degree of parallelism in high-end parallel supercomputers [1]. A fundamental question is: Is there a class of applications that will continue to scale on future supercomputers, *i.e.*, are there "metascalable" (or "design once, scale on new architectures") algorithms [2, 3]?

Divide-and-conquer (DC) is a highly scalable and versatile algorithmic paradigm, which has been applied successfully to design linear-scaling algorithms for broad computational problems ranging from the formally $O(N^2)$ N-body problem [4-7], to the $O(N^3)$ eigenvalue problem [8, 9] and linear systems [10], to the exponentially complex quantum N-body problem [11-17]. Our own research has focused on molecular dynamics (MD) simulations, which follow the trajectories of all atoms to study material properties and processes [18]. These include reactive molecular dynamics (RMD) simulations to study chemical reactions [19], and quantum molecular dynamics (QMD) simulations in which interatomic interaction is described quantum mechanically from first principles [3].

To extend the applicability of DC to even broader problems, we have recently proposed its extension named divide-conquer-recombine (DCR) [20]. In DCR, the DC phase constructs globally informed local solutions, which in the recombine phase are synthesized into a global solution encompassing large spatiotemporal scales. To efficiently implement DCR simulation algorithms on massively parallel supercomputers, we have also developed globally scalable and locally fast (GSLF) solvers that hybridize, *e.g.*, a global real-space multigrid with local plane-wave bases [3]. In the context of QMD, for example, DC electronic wave functions have been used to synthesize: (1) high-order inter-molecular-fragment interactions [21, 22]; (2) global frontier (*i.e.* highest occupied and lowest unoccupied) molecular orbitals [23, 24]; and (3) global charge-migration dynamics [25, 26].

The above DCR-in-space algorithms encompass large length scales based on data locality principles such as the quantum nearsightedness principle [27]. A harder problem is to predict long-time dynamics of these systems, because the sequential

bottleneck of time precludes efficient parallelization. Recently, various statistical approaches have been developed [28, 29] to utilize an ensemble of mutually uncorrelated MD trajectories to synthesize longer-time trajectories based on temporal locality of physical processes [30]. This DCR-in-time, combined with space-time-ensemble parallelism (STEP) [31], maximally exposes concurrency and data locality, and allows the prediction of material processes at the length and time scales that are otherwise impossible to reach. We have used DCR-in-time to study the long-time global kinetics of photoexcited electrons and holes in an exciton-flow network [32].

The global-local separation through DCR, enhanced with GSLF and STEP techniques, can be generalized into broad applications and will likely make them metascalable, with a minimal architectural assumption that a tree network topology (involving progressively reduced communication volume at upper tree levels) will be supported.

This paper is organized as follows. The next section introduces the DCR algorithmic paradigm. Section 3 demonstrates the scalability and performance of DCR algorithms. Applications of DCR-based QMD and RMD simulations are presented in section 4, and finally section 5 contains conclusion.

2. DIVIDE-CONQUER-RECOMBINE (DCR)

2.1 DCR-in-Space

In a divide-conquer-recombine (DCR) algorithm, the divide-and-conquer (DC) phase constructs globally informed local solutions in $O(N)$ time based on tree-based algorithms; see Fig. 1(a) [20, 32]. Subsequently in the recombine phase, these local solutions are synthesized into a global solution that encompasses large spatiotemporal scales; see Fig. 1(b) [20, 32]. Figures 1(a) schematically shows a tree data structure used in DCR-based molecular dynamics (MD) simulations, in which the root of the tree represents the entire simulation volume Ω. The entire volume is recursively subdivided into subsystems (or cells) of equal volume. Recursive subdivision is repeated until each cell at the leaf level defines a DC domain Ω_α.

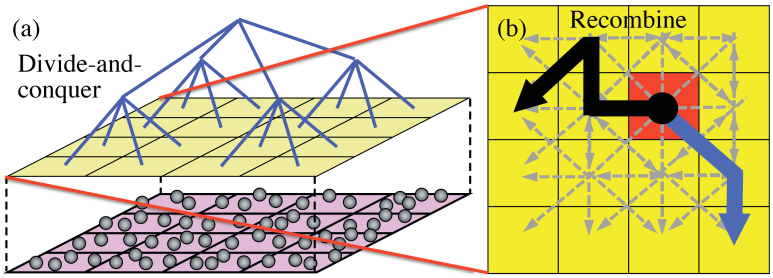

Figure 1. DCR algorithm. (a) Two-dimensional schematic of the DC phase, in which atoms (spheres in the bottom plane) within a DC domain (each small magenta parallelogram in the bottom plane) are abstracted by collective variables such as a finite-difference representation of the charge distribution. The DC domains are combined recursively to form a tree data structure consisting of progressively coarser cells, until reaching the entire simulation volume at the root of the tree. A DC algorithm performs $O(N)$ computation by traversing the tree both upward and downward. (b) Typical triplet and quadruplet computations involving one of the DC domains (colored red) in the recombine phase are shown by blue and black lines, respectively.

2.1.1 Divide-and-Conquer

We have designed a number of $O(N)$ DC algorithms (N is the number of atoms) for MD simulations.

Space-time multiresolution molecular dynamics (MRMD) reduces the $O(N^2)$ complexity of the N-body problem to $O(N)$ [18]. 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) [4, 7] to reduce the computational complexity to $O(N)$ based on spatial locality. We also use multiresolution in time, where temporal locality is utilized by computing forces from further atoms with less frequency [6].

Fast ReaxFF (F-ReaxFF) algorithm solves the $O(N^3)$ variable N-charge problem in chemically reactive molecular dynamics (RMD) simulations in $O(N)$ time [19]. 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. To describe charge transfer, ReaxFF uses a charge-equilibration scheme, in which atomic charges are determined at every MD step to minimize the electrostatic energy with the charge-neutrality constraint. This variable N-charge problem amounts to solving a dense linear system of equations, which requires $O(N^3)$ operations. F-ReaxFF uses 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 [10]. 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.

Lean divide-and-conquer density functional theory (LDC-DFT) is a new $O(N)$ algorithm for quantum molecular dynamics (QMD) simulations, which follow the trajectories of all atoms while computing interatomic forces quantum mechanically from first principles [3, 20]. Density functional theory (DFT) approximately reduces the exponential complexity of the quantum N-body problem to $O(N^3)$ by solving N one-electron problems self-

consistently instead of directly solving one N-electron problem. We solve the Kohn-Sham (KS) equations within each DC domain Ω_α to obtain local electronic wave functions $\{\psi_n^\alpha(\mathbf{r})\}$ iteratively using the MPCG method (n indexes multiple wave functions per domain). The KS equations depend on the global electrostatic potential, which is computed from the global electron density $\rho(\mathbf{r})$ in $O(N)$ time using a tree-based multigrid method. The global density in turn is a sum of domain densities calculated from $\{\psi_n^\alpha(\mathbf{r})\}$. The global $\rho(\mathbf{r})$ and local $\{\psi_n^\alpha(\mathbf{r})\}$ are determined iteratively in a global-local self-consistent-field (SCF) iteration. On the basis of complexity and error analyses, LDC-DFT minimizes the $O(N)$ prefactor through: (1) optimization of DC computational parameters; and (2) a density-adaptive boundary condition [3, 20].

2.1.2 *Recombine*

The key idea of DCR is to utilize DC solutions as compactly supported basis functions, with which global properties are synthesized using various recombination algorithms. The recombine phase typically performs range-limited n-tuple computations [33] among DC domains to account for higher inter-domain correlations that are not captured by the tree topology used in the DC phase. Examples of triplet (n = 3) and quadruplet (n = 4) computations are illustrated by arrows in Fig. 1(b). An example of n-tuple computations in the recombine phase (Fig. 1(b)) is the computation of the effective inter-molecular-fragment interaction energy, for which up to 4-tuple corrections have been incorporated in a perturbative manner [21]. Furthermore, even higher-order screening effects have been included through a posteriori recipe based on statistical mechanics, in which the self-consistent Ornstein-Zernike equation was solved within the Percus-Yevick [21] and hypernetted-chain [22] approximations. Linear combination of DC electronic wave functions can also be used in the recombine phase to construct the highest occupied molecular orbitals (HOMO) and lowest unoccupied molecular orbitals (LUMO) of the entire system [23]. The computational cost for obtaining these global frontier orbitals is drastically reduced by including only a small subset of DC orbitals near the Fermi energy [23]. DC wave functions were also used to describe global charge-migration dynamics by constructing coarse-grained electronic Hamiltonians with the use of quantum-dynamical [25] or bridge Green function [26] methods.

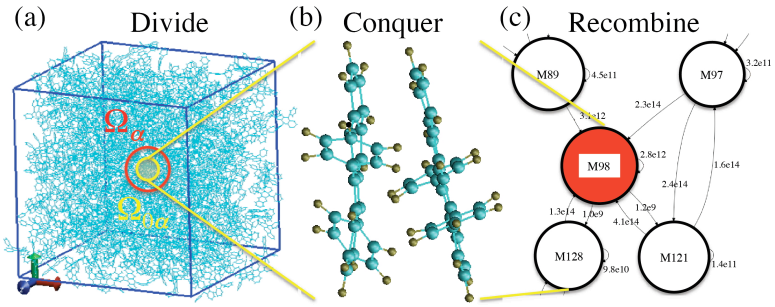

Figure 2. DCR to reach the experimental length and time scales for the exciton dynamics in amorphous molecular solid. (a) An entire configuration is subdivided into non-overlapping domain cores $\Omega_{0\alpha}$, each of which is augmented with neighbor molecules to form a domain Ω_α. (b) NAQMD simulations are performed in each domain to extract the rate constants of various excitonic processes. (c) In the recombine phase, KMC simulations are performed to describe the exciton dynamics that reflects the global geometry and topology of the exciton-flow network. Here, each circle represents a molecule, while a directed edge between molecules and a loop pointing to itself are labeled by exciton hopping and recombination rates, respectively.

In this paper, we focus on the recombination of DC wave functions to describe the kinetics of photoexcited electrons and holes, for which the global topology of a large-scale exciton-flow network was found to play an essential role [32]. We first perform nonadiabatic quantum molecular dynamics (NAQMD) simulations to describe coupled electron-ion dynamics involving nonadiabatic transitions between excited electronic states. Let us consider a specific example of amorphous diphenyl tetracene (DPT) consisting of 128 DPT molecules (each DPT molecule in turn consists of 50 atoms). To enable larger NAQMD simulations than have been performed previously (*i.e.* less than 1,000 atoms) [34, 35], the entire simulation box Ω is subdivided into M non-overlapping spatial domains $\Omega_{0\alpha}$; see Fig. 2(a). Here, each DPT molecule constitutes a non-overlapping domain $\Omega_{0\alpha}$, thus M = 128. We augment $\Omega_{0\alpha}$ by surrounding it with a buffer layer consisting of the k nearest-neighbor molecules in terms of the intermolecular C-C distance averaged over the backbone π-orbital planes (we use k = 2), so that the augmented domains Ω_α are mutually overlapping. For NAQMD simulation within each augmented domain Ω_α (Fig. 2(b)), the rest of the system is represented by a fixed charge density. In each domain Ω_α consisting of k+1 molecules, the charge density from the other M–k–1 molecules is used to form a global KS potential in Ω, including a non-additive contribution to the kinetic energy within an embedded cluster scheme [13]. Each NAQMD simulation starts from an electronic excited state corresponding to the excitation of an electron from the HOMO to the LUMO. Next, NAQMD trajectories are analyzed to obtain exciton-hopping rates between DPT molecules. The overlapping domains in the DC approach allow the construction of a graph data structure that spans the entire amorphous DPT solid; see Fig. 2(c). In the graph (or exciton-flow network), each DPT molecule constitutes a node, and the nodes are interconnected by directed edges labeled by the corresponding exciton-hopping rates obtained by the NAQMD simulations. The nonadiabatic coupling is also used to compute the exciton-annihilation rate, at which each exciton recombines to the electronic ground state. In addition to the phonon-assisted contribution to electronic transitions computed by NAQMD, we include the spontaneous emission contribution calculated within the transition dipole approximation. In addition to the exciton hopping and annihilation rates, we estimate the singlet-fission (SF) rate of each singlet excitonic state using a time-dependent

perturbation theory. Next, we perform first-principles kinetic Monte Carlo (KMC) simulations of exciton dynamics using the calculated hopping rates between DPT molecules as well as the SF and annihilation rates. Each KMC simulation starts by placing an exciton on a randomly selected DPT molecule. At each KMC step, the exciton either: (1) hops to one of the *k*-neighbor DPT molecules; (2) splits into two triplet excitons via SF; or (3) annihilates. The event to occur is chosen stochastically with the probability proportional to the corresponding rate. Each simulation continues until either SF or annihilation event occurs.

2.1.3 Globally Scalable and Locally Fast (GSLF) Solver

To efficiently solve the global-local SCF problem in LDC-DFT, our globally scalable and locally fast (GSLF) electronic-structure solver combines a local plane-wave basis within each DC domain for high numerical efficiency and a global real-space multigrid for scalability on massively parallel computers (Fig. 3) [3, 20]. Here,

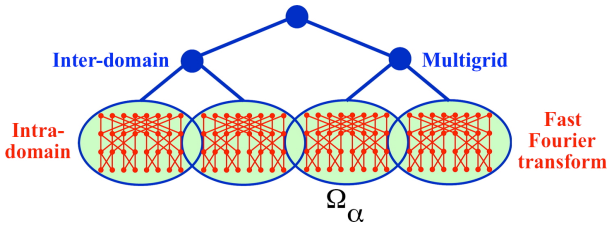

Figure 3. Globally scalable and locally fast, hybrid electronic-structure solver. Blue lines represent a tree network topology for multigrid-based inter-domain computation, whereas red lines show the butterfly network for fast Fourier transform within each domain.

fast intra-domain computation utilizes fast Fourier transform (FFT) to operate the kinetic-energy and potential-energy operators on the electronic wave functions in the sparsest form, respectively, in the real and reciprocal spaces. On the other hand, a real-space multigrid method is used to solve the Poisson equation to obtain the global electrostatic potential based on the locality preserving octree data structure; see Fig. 1(a).

2.2 DCR-in-Time

2.2.1 Molecular Kinetics

A challenging problem is to predict long-time dynamics because of the sequential bottleneck of time [36, 37]. 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 [30, 38]. We also introduce a discrete abstraction based on graph data structures, so that combinatorial techniques can be used for the search [38].

We have developed a directionally heated nudged elastic band (DH-NEB) method to search for thermally activated events starting from a given initial state without the knowledge of final states [31]. In DH-NEB, a nudged elastic band (NEB) consisting of a sequence of S states [39], $\mathbf{R}_s \in \Re^{3N}$ (s = 0,...,S–1) ($\Re$ is the set of real numbers), at different temperatures searches for a transition event (Fig. 4(a)). We use an ensemble consisting of B bands to perform long-time simulation — molecular kinetics (MK) simulation — in the framework of KMC [31]. At each MK step, the B bands concurrently find escape events from a common initial state, thereby constructing a list of multiple events. The activation barrier of each event is used to calculate the rate of the event, and the event to occur is chosen stochastically according to the rate. The time is incremented stochastically according to the Poisson distribution corresponding to the sum of the rates of all the events.

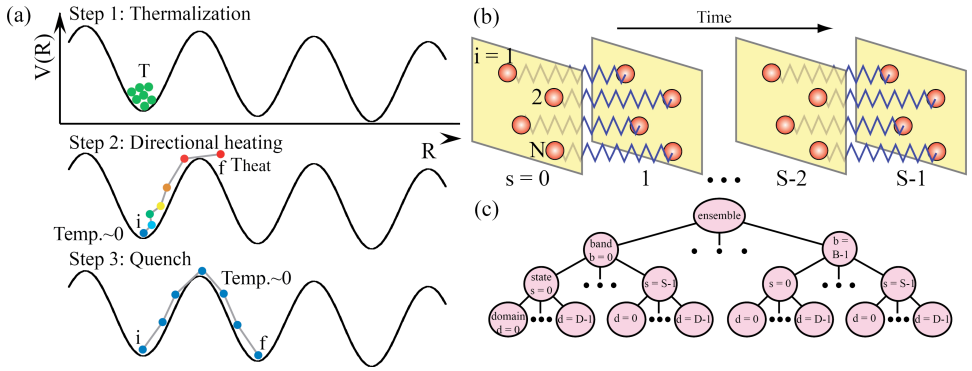

Figure 4. Schematic of the DH-NEB and STEP approaches. (a) Algorithmic steps of the DH-NEB method consist of thermalization (step 1), directional heating (step 2), and quench (step 3) of a band. Black solid curves represent the potential energy surface $V(\mathbf{R})$, whereas circles (with color-coded 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. (b) A nudged elastic band consists of a sequence of S states (gray parallelograms), $\mathbf{R}_s$ (s = 0,...,S–1), where each state consists of N atoms (white spheres), i = 1,...,N. Corresponding atoms in consecutive states interact via harmonic forces represented by wavy lines. (c) Tree-structured processor organization in the STEP approach. An ensemble consists of B bands, each consisting of S states; each state in turn contains D spatial domains.

2.2.2 *Space-Time-Ensemble Parallelism (STEP)*

The ensemble-based MK simulations are implemented efficiently on massively parallel computers based on a space-time-ensemble parallel (STEP) approach [31]. STEP combines a hierarchy of concurrency, *i.e.*, the number of processors is $P = BSD$: (1) spatial decomposition within each state (D is the number of spatial subsystems); (2) temporal parallelism across S states within each band; and (3) ensemble parallelism over B bands (Fig. 4(c)). The multiple parallelization axes of STEP greatly enhance the scalability of MK simulations on massively parallel computers. The program is implemented using the message passing interface (MPI) library for interprocessor communications, and each domain is assigned a dedicated MPI communicator using an MPI_COMM_SPLIT call.

2.2.3 *Super-state Parallel Replica Dynamics (SPRD)*

Another widely-used DCR-in-time method is parallel replica dynamics (PRD) that samples rare events [40]. PRD reduces sequential long time dynamics (Fig. 5(a)) to statistically independent parallel trajectory runs (or replicas); see Fig. 5(b). To handle widely disparate activation barriers in complex reaction pathways, an extension of PRD called super-state parallel replica dynamics (SPRD) groups a number of microscopic states into a super-state, and only rare transitions between those super-states (or super-events) are handled with PRD [28].

Super-states are usually pre-defined based on some a priori knowledge [28, 41]. For complex reaction dynamics, for which such a priori definition is unavailable, we instead employ machine-learning approaches to automatically define super-states on the fly. Rare super-events are detected using cross-correlation between time series of replica trajectories based on time-series anomaly detection and Bayesian analysis [29, 42]. Conventional SPRD simulations run for a block of MD time steps (*e.g.* 1,000 steps), followed by cross-time-series statistical analysis to detect super-events performed in a separate post-processing phase, where simulation trajectory data are stored in files and later analyzed (Fig. 5(c)). We are currently developing *in situ* SPRD method, where correlation analysis is done as data are produced so as to reduce the need to store data in files for post-processing (Fig. 5(d)).

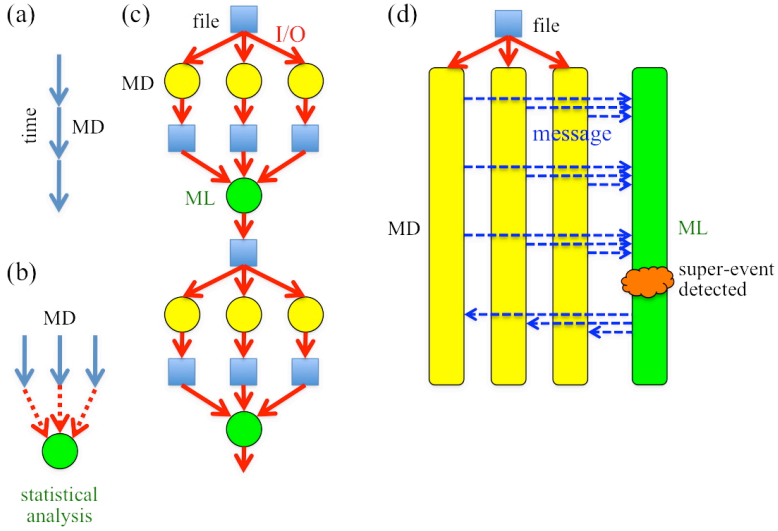

Figure 5. DCR-in-time. (a) Conventional MD simulation has sequential time dependence. (b) PRD predicts the long-time behavior through statistical analysis of multiple parallel MD trajectories. Conventional file-based (c) and new *in situ* (d) PRD simulations. ML represents machine-learning tasks.

3. SCALABILITY AND PERFORMANCE

DCR algorithms are characterized by excellent scalability, time-to-solution, and portable floating-point performance. This section presents benchmark results on our QMD and RMD simulation codes. The programs are written in Fortran 90 with MPI for message passing combined with OpenMP for multithreading.

Weak scaling: We have benchmarked our LDC-DFT code for QMD simulations on the 786,432-code Blue Gene/Q at the Argonne National Laboratory. We first performed a weak-scaling benchmark, in which the number of atoms per core N/P is kept constant. Figure 6(a) shows the wall-clock time per QMD simulation step with scaled workloads — $64P$-atom SiC system on P cores. The execution time includes 3 SCF iterations to determine the electronic wave functions and the global density, with 3 MPCG iterations per SCF cycle to refine each wave function. By increasing the number of atoms linearly with the number of cores, the wall-clock time remains nearly constant, indicating excellent scalability. To quantify the parallel efficiency, we first define the speed of the LDC-DFT program as a product of the total number of atoms and the number of QMD simulation steps executed per second. The isogranular speedup is given by the ratio between the speed on P cores and that on 16 cores (*i.e.* one computing node) as a reference system. The weak-scaling parallel efficiency is the isogranular speedup divided by $P/16$. With the granularity of 64 atoms per core, the parallel efficiency was 0.984 on $P = 786,432$ for a 50,331,648-atom (or 201,326,592-electron) SiC system. This computation involved 39,815,773,421,568 electronic wave-function and charge-density values sampled on mesh points. The result demonstrates the very

high scalability of the LDC-DFT algorithm.

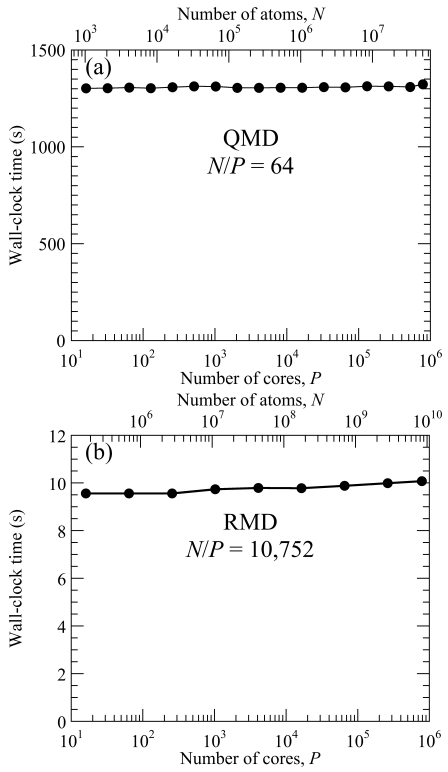

Figure 6. Weak scaling of DCR algorithms on Blue Gene/Q. (a) Wall-clock time per QMD simulation step of the parallel LDC-DFT algorithm, with scaled workloads — 64*P*-atom SiC system on *P* cores (*P* = 16, ..., 786,432) of Blue Gene/Q. (b) Wall-clock time per RMD simulation step of the parallel F-ReaxFF algorithm, with scaled workloads — 10,752*P*-atom RDX system on *P* cores (*P* = 16, ..., 786,432).

Figure 6(b) shows the execution time of the F-ReaxFF based RMD algorithm for cyclotrimethylenetrinitramine (RDX) material as a function of P, where the number of atoms is $N = 10{,}752P$. The computation time includes 3 MPCG iterations to solve the electronegativity equalization problem for determining atomic charges at each MD time step. On 786,432 Blue Gene/Q cores, the isogranular parallel efficiency of the F-ReaxFF algorithm is 0.957.

Strong scaling: Next, we performed a strong-scaling test by simulating a LiAl nanoparticle immersed in water containing a total of 77,889 atoms ($Li_{2136}Al_{2136}$ immersed in 24,539 H_2O molecules). In this test, the number of cores ranged from P = 49,152 to 786,432, while keeping the total problem size constant. Figure 7 shows the wall-clock time per QMD simulation step as a function of P. The time-to-solution was reduced by a factor of 12.85 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 12.85, with the corresponding strong-scaling parallel efficiency of 0.803.

Time-to-solution: The LDC-DFT algorithm significantly reduces the $O(N)$ prefactor of its computational cost with a given energy tolerance, and consequently the time-to-solution. To quantify the time-to-solution, let us consider the largest computation shown in Fig. 6(a), *i.e.*, a 50,331,648-atom SiC system on the entire 786,432 cores of the Blue Gene/Q. One SCF iteration using LDC-DFT in this case took 441 seconds, with a speed of 114,000 atom•iteration/s. This represents a several orders-of-magnitude improvement over the previous state-of-the-art in terms of the time-to-solution [3].

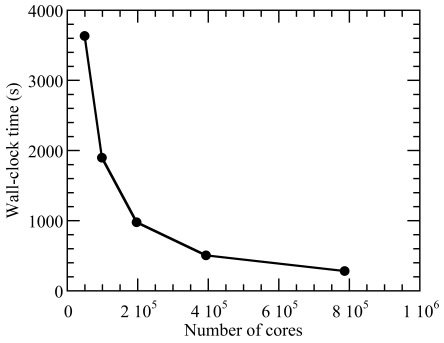

Figure 7. Wall-clock time per QMD simulation step of the parallel LDC-DFT algorithm with strong scaling — 77,889-atom LiAl-water system on *P* cores (*P* = 49,152, ..., 786,432) of Blue Gene/Q.

Floating-point performance: We have measured the floating-point (FP) performance of LDC-DFT on the Blue Gene/Q. The Blue Gene Performance Monitoring (BGPM) API was used to measure the FP performance of the entire program by linking the threaded HPM library (libmpihpm_smp). Here, a 6,291,456-atom SiC system was simulated on the entire 48 racks (or 786,432 cores). We used 16 MPI ranks per node, where each MPI rank spawns 4 threads. The measured performance was 50.5% of the peak (*i.e.* 5.08 PFLOP/s).

4. SCIENCE APPLICATIONS

DCR algorithms have enabled some of the largest QMD and RMD simulations. This section presents several examples.

4.1 Hydrogen-on-Demand for Renewable Energy

Hydrogen production from water using aluminum (Al) particles could provide a renewable energy cycle [43]. However, its practical application is hampered by the low reaction rate and poor yield. Using the largest ever QMD simulations on 786,432 Blue Gene/Q cores based on the LDC-DFT algorithm (Fig. 8(a)), we have shown that orders-of-magnitude faster reactions with higher yields can be achieved by alloying Al particles with lithium (Li); see Fig. 8(b) [44]. Here, the orders-of-magnitude improvement of the time-to-solution over the previous state-of-the-art shown in section 3 has enabled QMD simulations encompassing unprecedented spatio-temporal scales, involving 16,661 atoms (or 43,708 electrons) for 21,140 time steps (or 129,208 SCF iterations) with a unit time step of 0.242 fs.

Based on the simulation data, a key nanostructural design has been identified as the abundance of neighboring Lewis acid-base pairs, where water-dissociation and hydrogen-production require very small activation energies. These reactions are facilitated by wide charge pathways across Al atoms that collectively act as a "superanion". Furthermore, dissolution of Li atoms into water produces a corrosive basic solution that inhibits the formation of a reaction-stopping oxide layer on the particle surface, thereby increasing the yield. We have also found a surprising catalytic behavior of bridging oxygens (similar to Ref. [45]) that connect

Al and Li. Namely, Li-O-Al is not merely an inert reaction product but instead plays an unexpectedly active role in the oxidation process by assisting the breakage of O-H and formation of Al-O bonds.

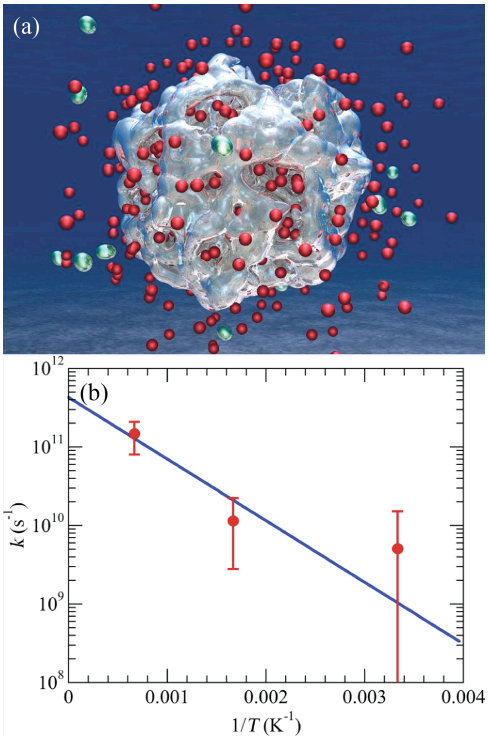

Figure 8. QMD simulation of H_2 production from water using a $Li_{441}Al_{441}$ particle. (a) The valance electron density (silver isosurface) is centered around Al atoms, whereas some of the Li atoms (red spheres) are dissolved into water. Produced H_2 molecules are represented by green ellipsoids. Water molecules are not shown for clarity. (b) H_2 production rate per LiAl pair as a function of the inverse temperature.

Not only this microscopic understanding explains recent experimental findings in similar alloy systems (*e.g.* alloy composition-dependent reactivity and a remarkable pH change associated with H_2 production) [46], but it also predicts a specific nanostructural design for rapid high-yield production of hydrogen on demand, which is expected to scale up to industrially relevant particle sizes. This work thus lays a foundation for future studies on fundamental science toward rational nanostructural design of renewable energy technologies.

4.2 Singlet Fission for Efficient Solar Cells

Singlet fission (SF) is a process, in which a spin-singlet exciton in an organic semiconductor is split into two spin-triplet excitons [32]. If high SF yield is realized in mass-produced disordered organic solid, it could revolutionize low-cost fabrication of efficient solar cells. Recently, an experimental breakthrough was made by observing SF in amorphous diphenyl tetracene (DPT). However, atomistic mechanisms that enable efficient SF in amorphous molecular solid remain elusive, largely due to the required large quantum-mechanical calculations that capture nanostructural features.

To address this multiscale challenge, we have applied DCR to a NAQMD study of photoexcitation dynamics in amorphous DPT composed of 6,400 atoms (Fig. 9(a)) [32]. We identified the key molecular geometry (*i.e.*, molecular dimers with close twisted stacking of π-orbital planes) and exciton-flow-network topology (*i.e.*, a large number of reverse nearest neighbor molecules) for SF "hot spots", where SF occurs preferentially. The simulation revealed the molecular origin of experimentally observed two time scales in exciton population dynamics (Fig. 9(b)), and may pave a way to nanostructural design of efficient low-cost solar cells from first principles.

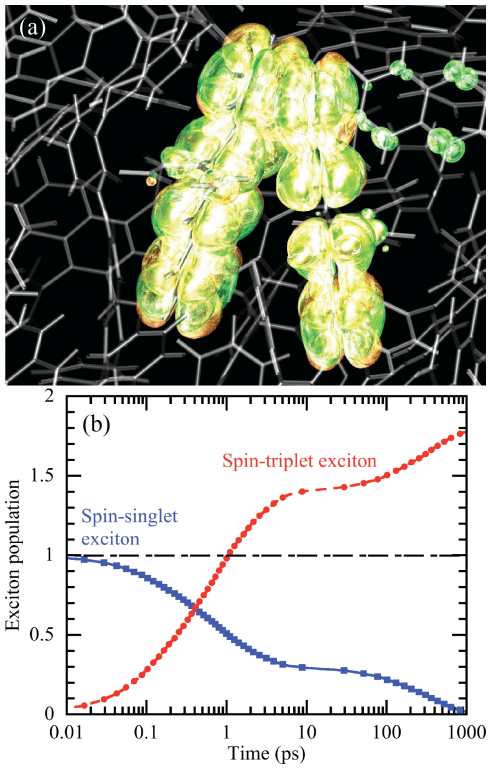

Figure 9. NAQMD simulation of photo-induced excitons in amorphous DPT. (a) Orange and green isosurfaces represent quasi-electron and quasi-hole wave functions, whereas gray rods represent DPT molecules. (b) The calculated population dynamics of spin-singlet and spin-triplet excitons.

4.3 Cavitation Bubble Collapse

Cavitation bubbles occur in fluids subjected to rapid changes in pressure. We used billion-atom RMD simulations on a 163,840-processor Blue Gene/P supercomputer to investigate damage caused by shock-induced collapse of nanobubbles in water near an amorphous silica surface [47]; see Fig. 10. Collapse of an empty bubble generated a high-speed nanojet, which caused pitting on silica surface. We found that pit radii are close to bubble radii, and experiments also indicate linear scaling between them. The gas-filled bubbles underwent partial collapse and consequently the damage on the silica surface was mitigated.

5. CONCLUSION

With the advent of multicore revolution in computer architectures, DC software on emerging exascale computers could provide an unprecedented capability to solve complex problems. But this is

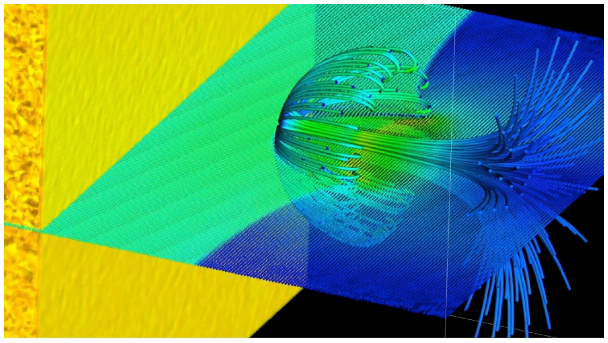

Figure 10. Billion-atom RMD simulation of cavitation bubble collapse in water. A nanojet (stream lines) hits and damages silica surface (yellow). The color represents pressure.

true only if the software continues to scale on the many millions of cores expected in a future exascale computer. This is an enormous challenge, since we do not even know the architecture of such platforms. The primary challenge is to sustain DC's scalability on rapidly evolving parallel computing architectures. Light-overhead DC algorithms such as LDC-DFT are expected to be "metascalable", assuming only that a tree network topology (involving progressively small communication volume at upper tree levels) will be supported. The communication requirement is further reduced algorithmically via dimensionality reduction — abstracting the global information with much reduced dimensionality, *e.g.*, using only one global function (density) and one scalar value (chemical potential) instead of the entire $O(N)$ wave functions in LDC-DFT. As an extension of DC, we have proposed a highly scalable computational approach named DCR to describe various large spatiotemporal-scale material processes using local atomistic and electronic solutions as a compact basis set. The highly scalable parallel DCR approach has broad applicability for multiscale material problems, where electronic structures and chemical reactions are inseparably coupled to microstructures and long-range stress fields. Innovative algorithms in the recombine phase typically perform range-limited n-tuple computations [33] among DC domains to account for higher inter-domain correlations that are not captured by the tree topology used in the DC phase, so that a wider variety of hard problems can be addressed by DCR. For example, innovative recombination of DC solutions can be incorporated into high-throughput screening to explore a large combinatorial search space for discovering new materials [48, 49]. Implemented with metascalable DC algorithms like LDC-DFT, in conjunction with scalable GSLF and STEP computational techniques, the DCR algorithmic framework promises to scale on future systems to address many challenging scientific and engineering problems. The DCR algorithmic framework is thus expected to provide a metascalable-computing paradigm on future systems, which is generalizable to very broad applications.

6. ACKNOWLEDGMENTS

This research was supported by the U.S. Department of Energy (DOE), Office of Science, Basic Energy Sciences, Materials Science and Engineering Division, Theoretical Condensed Matter Physics Program, Grant # DE-FG02-04ER-46130. We used resources of the Argonne Leadership Computing Facility at Argonne National Laboratory, which is supported by the Office of Science of the U.S. DOE under contract DE-AC02-06CH11357. The initial code development and the performance portability test were conducted at the Center for High-Performance Computing and Communication at the University of Southern California. We thank Dr. Bob Walkup of IBM for his help on performance monitoring and optimization, Dr. Franz Franchetti of Carnegie Mellon University and Dr. Brian Duff of SpiralGen for their help on the use of the Spiral FFT software, and Dr. Vitali A. Morozov of ALCF for his help on debugging our code on Blue Gene/Q.

7. REFERENCES

[1] Dongarra, J., Beckman, P., Moore, T., Aerts, P., Aloisio, G., Andre, J. C., Barkai, D., Berthou, J. Y., Boku, T., Braunschweig, B., Cappello, F., Chapman, B., Chi, X. B., Choudhary, A., Dosanjh, S., Dunning, T., Fiore, S., Geist, A., Gropp, B., Harrison, R., Hereld, M., Heroux, M., Hoisie, A., Hotta, K., Jin, Z., Ishikawa, Y., Johnson, F., Kale, S., Kenway, R., Keyes, D., Kramer, B., Labarta, J., Lichnewsky, A., Lippert, T., Lucas, B., Maccabe, B., Matsuoka, S., Messina, P., Michielse, P., Mohr, B., Mueller, M. S., Nagel, W. E., Nakashima, H., Papka, M. E., Reed, D., Sato, M., Seidel, E., Shalf, J., Skinner, D., Snir, M., Sterling, T., Stevens, R., Streitz, F., Sugar, B., Sumimoto, S., Tang, W., Taylor, J., Thakur, R., Trefethen, A., Valero, M., Van Der Steen, A., Vetter, J., Williams, P., Wisniewski, R., and Yelick, K., 2011. The international exascale software project roadmap. *International Journal of High Performance Computing Applications 25*, 1 (Feb), 3-60.

[2] Nomura, K., Dursun, H., Seymour, R., Wang, W., Kalia, R. K., Nakano, A., Vashishta, P., Shimojo, F., and Yang, L. H., IEEE, 2009. A metascalable computing framework for large spatiotemporal-scale atomistic simulations. *Proceedings of the International Parallel and Distributed Processing Symposium, IPDPS 2009*.

[3] Nomura, K., Kalia, R. K., Nakano, A., Vashishta, P., Shimamura, K., Shimojo, F., Kunaseth, M., Messina, P. C., and Romero, N. A., IEEE/ACM, 2014. Metascalable quantum molecular dynamics simulations of hydrogen-on-demand. *Proceedings of Supercomputing, SC14*.

[4] Greengard, L. and Rokhlin, V., 1987. A fast algorithm for particle simulations. *Journal of Computational Physics 73*, 2 (Dec), 325-348.

[5] White, C. A. and Headgordon, M., 1994. Derivation and efficient implementation of the fast multipole method. *Journal of Chemical Physics 101*, 8 (Oct 15), 6593-6605.

[6] Nakano, A., Kalia, R. K., and Vashishta, P., 1994. Multiresolution molecular-dynamics algorithm for realistic materials modeling on parallel computers. *Computer Physics Communications 83*, 2-3 (Dec), 197-214.

[7] Ogata, S., Campbell, T. J., Kalia, R. K., Nakano, A., Vashishta, P., and Vemparala, S., 2003. Scalable and portable implementation of the fast multipole method on parallel computers. *Computer Physics Communications 153*, 3 (Jul 1), 445-461.

[8] Cuppen, J. J. M., 1981. A divide and conquer method for the symmetric tridiagonal eigenproblem. *Numerische Mathematik 36*, 2, 177-195.

[9] Gansterer, W. N., Ward, R. C., and Muller, R. P., 2002. An extension of the divide-and-conquer method for a class of symmetric block-tridiagonal eigenproblems. *ACM Transactions on Mathematical Software 28*, 1 (Mar), 45-58.

[10] Nakano, A., 1997. Parallel multilevel preconditioned conjugate-gradient approach to variable-charge molecular dynamics. *Computer Physics Communications 104*, 1-3 (Aug), 59-69.

[11] Yang, W. T., 1991. Direct calculation of electron-density in density-functional theory. *Physical Review Letters 66*, 11 (Mar 18), 1438-1441.

[12] Dixon, S. L. and Merz, K. M., 1997. Fast, accurate semiempirical molecular orbital calculations for macromolecules. *Journal of Chemical Physics 107*, 879.

[13] Shimojo, F., Kalia, R. K., Nakano, A., and Vashishta, P., 2005. Embedded divide-and-conquer algorithm on hierarchical real-space grids: parallel molecular dynamics simulation based on linear-scaling density functional theory. *Computer Physics Communications 167*, 3 (May 1), 151-164.

[14] Ozaki, T., 2006. *O*(*N*) Krylov-subspace method for large-scale ab initio electronic structure calculations. *Physical Review B 74*, 24 (Dec 1), 245101.

[15] Kobayashi, M. and Nakai, H., 2008. Extension of linear-scaling divide-and-conquer-based correlation method to coupled cluster theory with singles and doubles excitations. *Journal of Chemical Physics 129*, 4 (Jul 28), 044103.

[16] Shimojo, F., Kalia, R. K., Nakano, A., and Vashishta, P., 2008. Divide-and-conquer density functional theory on hierarchical real-space grids: parallel implementation and applications. *Physical Review B 77*, 8 (Feb 15), 085103.

[17] Ohba, N., Ogata, S., Kouno, T., Tanmura, T., and Kobayashi, R., 2012. Linear scaling algorithm of real-space density functional theory of electrons with correlated overlapping domains. *Computer Physics Communications 183*, 8 (Aug), 1664-1673.

[18] Nakano, A., Kalia, R. K., Vashishta, P., Campbell, T. J., Ogata, S., Shimojo, F., and Saini, S., ACM/IEEE, 2001. Scalable atomistic simulation algorithms for materials research. *Proceedings of Supercomputing, SC01*.

[19] Nomura, K., Kalia, R. K., Nakano, A., and Vashishta, P., 2008. A scalable parallel algorithm for large-scale reactive force-field molecular dynamics simulations. *Computer Physics Communications 178*, 2 (Jan 15), 73-87.

[20] Shimojo, F., Kalia, R. K., Kunaseth, M., Nakano, A., Nomura, K., Ohmura, S., Shimamura, K., and Vashishta, P., 2014. A divide-conquer-recombine algorithmic paradigm for multiscale materials modeling. *Journal of Chemical Physics 140*, 18 (May 14), 18A529.

[21] Okiyama, Y., Tsukamoto, T., Watanabe, C., Fukuzawa, K., Tanaka, S., and Mochizuki, Y., 2013. Modeling of peptide-silica interaction based on four-body corrected fragment molecular orbital (FMO4) calculations. *Chemical Physics Letters 566*(Apr 12), 25-31.

[22] Tanaka, S., Watanabe, C., and Okiyama, Y., 2013. Statistical correction to effective interactions in the fragment molecular orbital method. *Chemical Physics Letters 556*(Jan 29), 272-277.

[23] Tsuneyuki, S., Kobori, T., Akagi, K., Sodeyama, K., Terakura, K., and Fukuyama, H., 2009. Molecular orbital calculation of biomolecules with fragment molecular orbitals. *Chemical Physics Letters 476*, 1-3 (Jul 7), 104-108.

[24] Kobori, T., Sodeyama, K., Otsuka, T., Tateyama, Y., and Tsuneyuki, S., 2013. Trimer effects in fragment molecular orbital-linear combination of molecular orbitals calculation of one-electron orbitals for biomolecules. *Journal of Chemical Physics 139*, 9 (Sep 7), 094113.

[25] Gollub, C., Avdoshenko, S., Gutierrez, R., Berlin, Y., and Cuniberti, G., 2012. Charge migration in organic materials: can propagating charges affect the key physical quantities controlling their motion? *Israel Journal of Chemistry 52*, 5 (May), 452-460.

[26] Kitoh-Nishioka, H. and Ando, K., 2012. Fragment molecular orbital study on electron tunneling mechanisms in bacterial photosynthetic reaction center. *Journal of Physical Chemistry B 116*, 43 (Nov 1), 12933-12945.

[27] Prodan, E. and Kohn, W., 2005. Nearsightedness of electronic matter. *Proceedings of the National Academy of Sciences 102*, 33 (Aug 16), 11635-11638.

[28] Kum, O., Dickson, B. M., Stuart, S. J., Uberuaga, B. P., and Voter, A. F., 2004. Parallel replica dynamics with a heterogeneous distribution of barriers: application to n-hexadecane pyrolysis. *Journal of Chemical Physics 121*, 20 (Nov 22), 9808-9819.

[29] Pande, V. S., Beauchamp, K., and Bowman, G. R., 2010. Everything you wanted to know about Markov state models but were afraid to ask. *Methods 52*, 1 (Sep), 99-105.

[30] Voter, A. F., Montalenti, F., and Germann, T. C., 2002. Extending the time scale in atomistic simulation of materials. *Annual Review of Materials Research 32*, 321-346.

[31] Nakano, A., 2008. A space-time-ensemble parallel nudged elastic band algorithm for molecular kinetics simulation. *Computer Physics Communications 178*, 280-289.

[32] Mou, W., Hattori, S., Rajak, P., Shimojo, F., and Nakano, A., 2013. Nanoscopic mechanisms of singlet fission in amorphous molecular solid. *Applied Physics Letters 102*, 17 (Apr 29), 173301.

[33] Kunaseth, M., Kalia, R. K., Nakano, A., Nomura, K., and Vashishta, P., ACM/IEEE, 2013. A scalable parallel algorithm for dynamic range-limited n-tuple computation in many-body molecular dynamics simulation. *Proceedings of Supercomputing, SC13*.

[34] Zhang, X., Li, Z., and Lu, G., 2011. First-principles simulations of exciton diffusion in organic semiconductors. *Physical Review B 84*, 23 (Dec 22), 235208.

[35] Mou, W., Ohmura, S., Shimojo, F., and Nakano, A., 2012. Molecular control of photoexcited charge transfer and recombination at a quaterthiophene/zinc oxide interface. *Applied Physics Letters 100*, 20 (May 18), 203306.

[36] Phillips, J. C., Zheng, G., Kumar, S., and Kale, L. V., 2002. NAMD: biomolecular simulation on thousands of processors. In *Proceedings of the Proceedings of Supercomputing (SC02)* (2002), ACM/IEEE.

[37] Shaw, D. E., Deneroff, M. M., Dror, R. O., Kuskin, J. S., Larson, R. H., Salmon, J. K., Young, C., Batson, B., Bowers, K. J., Chao, J. C., Eastwood, M. P., Gagliardo, J., Grossman, J. P., C.R. Ho, Ierardi, D. J., KolossváRy, I., Klepeis, J. L., Layman, T., Mcleavey, C., Moraes, M. A., Mueller, R., Priest, E. C., Shan, Y., Spengler, J., Theobald, M., Towles, B., and Wang, S. C., 2007. Anton, a special-purpose machine for molecular dynamics simulation. *ACM SIGARCH Computer Architecture News 35*, 2, 1-12.

[38] Nakano, A., 2007. Pathfinder: a parallel search algorithm for concerted atomistic events. *Computer Physics Communications 176*, 4 (Feb 15), 292-299.

[39] Henkelman, G. and Jonsson, H., 2000. Improved tangent estimate in the nudged elastic band method for finding minimum energy paths and saddle points. *Journal of Chemical Physics 113*, 22 (DEC 8), 9978-9985.

[40] Voter, A. F., 1998. Parallel replica method for dynamics of infrequent events. *Physical Review B 57*, 22 (Jun 1), R13985-R13988.

[41] Joshi, K. L., Raman, S., and Van Duin, A. C. T., 2013. Connectivity-based parallel replica dynamics for chemically reactive systems: from femtoseconds to microseconds.

Journal of Physical Chemistry Letters 4, 21 (Nov 7), 3792-3797.

[42] Hinrichs, N. S. and Pande, V. S., 2007. Calculation of the distribution of eigenvalues and eigenvectors in Markovian state models for molecular dynamics. *Journal of Chemical Physics 126*, 24 (Jun 28), 244101.

[43] Shimojo, F., Ohmura, S., Kalia, R. K., Nakano, A., and Vashishta, P., 2010. Molecular dynamics simulations of rapid hydrogen production from water using aluminum clusters as catalyzers. *Physical Review Letters 104*, 12 (Mar 26), 126102.

[44] Shimamura, K., Shimojo, F., Kalia, R. K., Nakano, A., Nomura, K., and Vashishta, P., 2014. Hydrogen-on-demand using metallic alloy nanoparticles in water. *Nano Letters*, accepted. DOI= http://dx.doi.org/10.1021/nl501612v.

[45] Wu, C. J., Fried, L. E., Yang, L. H., Goldman, N., and Bastea, S., 2009. Catalytic behaviour of dense hot water. *Nature Chemistry 1*, 1 (Apr 1), 57-62.

[46] Chen, X., Zhao, Z., Hao, M., and Wang, D., 2013. Hydrogen generation by splitting water with Al-Li alloys. *International Journal of Energy Research 37*, 13 (Oct 25), 1624-1634. DOI= http://dx.doi.org/10.1002/er.2978.

[47] Shekhar, A., Nomura, K., Kalia, R. K., Nakano, A., and Vashishta, P., 2013. Nanobubble collapse on a silica surface in water: Billion-atom reactive molecular dynamics simulations. *Physical Review Letters 111*, 18 (Oct 31), 184503.

[48] Wilmer, C. E., Leaf, M., Lee, C. Y., Farha, O. K., Hauser, B. G., Hupp, J. T., and Snurr, R. Q., 2012. Large-scale screening of hypothetical metal-organic frameworks. *Nature Chemistry 4*, 2 (Feb), 83-89.

[49] Virshup, A. M., Contreras-Garcia, J., Wipf, P., Yang, W. T., and Beratan, D. N., 2013. Stochastic voyages into uncharted chemical space produce a representative library of all possible drug-like compounds. *Journal of the American Chemical Society 135*, 19 (May 15), 7296-7303.