A divide-conquer-recombine algorithmic paradigm for large spatiotemporal quantum molecular dynamics simulations

Fuyuki Shimojo, Shinnosuke Hattori, Rajiv K. Kalia, Manaschai Kunaseth, Weiwei Mou, Aiichiro Nakano, Ken-ichi Nomura, Satoshi Ohmura, Pankaj Rajak, Kohei Shimamura, and Priya Vashishta

Citation: The Journal of Chemical Physics **140**, 18A529 (2014); doi: 10.1063/1.4869342
View online: http://dx.doi.org/10.1063/1.4869342
View Table of Contents: http://scitation.aip.org/content/aip/journal/jcp/140/18?ver=pdfcov
Published by the AIP Publishing

A divide-conquer-recombine algorithmic paradigm for large spatiotemporal quantum molecular dynamics simulations

Fuyuki Shimojo,[1,2] Shinnosuke Hattori,[1,2] Rajiv K. Kalia,[1] Manaschai Kunaseth,[1,3] Weiwei Mou,[1] Aiichiro Nakano,[1] Ken-ichi Nomura,[1] Satoshi Ohmura,[1,2,4] Pankaj Rajak,[1] Kohei Shimamura,[1,2,5] and Priya Vashishta[1]

[1]*Collaboratory for Advanced Computing and Simulations, Department of Physics and Astronomy, Department of Computer Science, and Department of Chemical Engineering and Materials Science, University of Southern California, Los Angeles, California 90089-0242, USA*
[2]*Department of Physics, Kumamoto University, Kumamoto 860-8555, Japan*
[3]*National Nanotechnology Center, Pathumthani 12120, Thailand*
[4]*Department of Physics, Kyoto University, Kyoto 606-8502, Japan*
[5]*Department of Applied Quantum Physics and Nuclear Engineering, Kyushu University, Fukuoka 819-0395, Japan*

(Received 25 November 2013; accepted 12 March 2014; published online 28 March 2014)

We introduce an extension of the divide-and-conquer (DC) algorithmic paradigm called divide-conquer-recombine (DCR) to perform large quantum molecular dynamics (QMD) simulations on massively parallel supercomputers, in which interatomic forces are computed quantum mechanically in the framework of density functional theory (DFT). In DCR, the DC phase constructs globally informed, overlapping local-domain solutions, which in the recombine phase are synthesized into a global solution encompassing large spatiotemporal scales. For the DC phase, we design a lean divide-and-conquer (LDC) DFT algorithm, which significantly reduces the prefactor of the $O(N)$ computational cost for N electrons by applying a density-adaptive boundary condition at the peripheries of the DC domains. Our globally scalable and locally efficient solver is based on a hybrid real-reciprocal space approach that combines: (1) a highly scalable real-space multigrid to represent the global charge density; and (2) a numerically efficient plane-wave basis for local electronic wave functions and charge density within each domain. Hybrid space-band decomposition is used to implement the LDC-DFT algorithm on parallel computers. A benchmark test on an IBM Blue Gene/Q computer exhibits an isogranular parallel efficiency of 0.984 on 786 432 cores for a 50.3×10^6-atom SiC system. As a test of production runs, LDC-DFT-based QMD simulation involving 16 661 atoms is performed on the Blue Gene/Q to study on-demand production of hydrogen gas from water using LiAl alloy particles. As an example of the recombine phase, LDC-DFT electronic structures are used as a basis set to describe global photoexcitation dynamics with nonadiabatic QMD (NAQMD) and kinetic Monte Carlo (KMC) methods. The NAQMD simulations are based on the linear response time-dependent density functional theory to describe electronic excited states and a surface-hopping approach to describe transitions between the excited states. A series of techniques are employed for efficiently calculating the long-range exact exchange correction and excited-state forces. The NAQMD trajectories are analyzed to extract the rates of various excitonic processes, which are then used in KMC simulation to study the dynamics of the global exciton flow network. This has allowed the study of large-scale photoexcitation dynamics in 6400-atom amorphous molecular solid, reaching the experimental time scales. © *2014 AIP Publishing LLC*. [http://dx.doi.org/10.1063/1.4869342]

I. INTRODUCTION

Divide-and-conquer (DC) is a highly scalable 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,[1–4] to the $O(N^3)$ eigenvalue problem[5,6] and linear systems,[7] to the exponentially complex quantum N-body problem.[8–14] With the advent of multicore revolution in computer architecture, DC software on emerging exaflop/s computers[15] will provide an unprecedented capability to solve complex problems, only if the software continues to scale on the millions of cores in an exaflop/s 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. Such a formidable challenge can only be addressed based on a solid theoretical foundation to guarantee provable scalability, adapting to evolving architectures, i.e., transforming the DC algorithmic framework to be "metascalable" (or "design once, scale on new architectures").[16]

For the study of material properties and processes involving electrons, the density functional theory (DFT)[17] has become a common computational method. There is growing interest in large quantum molecular dynamics (QMD) simulations involving thousands of atoms,[18] in which interatomic forces are computed quantum mechanically[19,20] in the framework of DFT. Such large QMD simulations

on high-end parallel supercomputers can describe the coupling of chemical reactions, atomistic processes, and long-range stress phenomena for broad applications. The major bottleneck of such large DFT calculations is the asymptotic $O(N^3)$ complexity for an N-electron system. To overcome this bottleneck, various $O(N)$ DFT algorithms[21] have been designed on the basis of the data locality principle called quantum nearsightedness.[22–24] Among them, the DC-DFT algorithm[8–14] pioneered by Weitao Yang[8] is highly scalable[16,25] on massively parallel computers.[26] Publicly available DC-DFT codes include the OpenMX software,[11] and implementation of DC-DFT codes on massively parallel computers is discussed in Refs. 16 and 27. The DC-DFT algorithm represents the three-dimensional space as a union of spatially localized domains, and global physical properties are computed as linear combinations of local domain properties. However, it is only in the past several years that the DC-DFT algorithm, especially with large basis sets ($>10^4$ unknowns per electron, which is necessary for the transferability of accuracy), has attained controlled error bounds, robust convergence properties, and adequate energy conservation[28] for its use in QMD simulations, to make large DFT-based QMD simulations practical.[10,13,16,25,29]

A major remaining problem associated with DC-DFT is the large prefactor of its $O(N)$ computational cost, which makes it a challenge to perform large QMD simulations involving over 10^4 atoms for more than 10^5 time steps. This large prefactor arises from a thick buffer layer that surrounds each computational domain in order to minimize the effect of artificial boundary conditions imposed at domain peripheries.[23] In this paper, we design a light overhead $O(N)$ DFT algorithm called lean divide-and-conquer (LDC) DFT. On the basis of complexity and error analyses of DC-DFT, LDC minimizes the $O(N)$ prefactor through: (1) optimization of DC parameters; and (2) a density-adaptive boundary condition. We use a hybrid real-reciprocal (HR^2) space approach that combines: (1) a plane-wave basis for electronic wave functions and charge density within each DC domain; and (2) a real-space multigrid to represent the global charge density. Hybrid space-band (HSB) decomposition is used to implement the algorithm on massively parallel computers. The parallel LDC-DFT code has achieved a parallel efficiency of 0.984 on 786 432 IBM Blue Gene/Q cores for a 50.3×10^6-atom SiC system. LDC-DFT has also been use for 16 661-atom QMD simulation to study on-demand production of hydrogen gas from water using LiAl alloy particles.

One advantage of DC-DFT is the possibility to use local electronic structures from DC domains as a basis set to construct various physical properties at the global level. For example, DC electronic wave functions have been used to calculate: (1) high-order inter-molecular-fragment interactions;[30,31] (2) global frontier (i.e., highest occupied and lowest unoccupied) molecular orbitals;[32,33] (3) global charge-migration dynamics;[34,35] and (4) global kinetics of photoexcited electrons and holes in an exciton flow network;[36] at the length and time scales that are otherwise impossible to reach. In this paper, we present an extension of DC named divide-conquer-recombine (DCR) as a metascalable algorithmic paradigm that will continue to scale on future computing architectures. In DCR, the DC phase constructs globally informed, overlapping local-domain solutions,[8] which in the recombine phase are synthesized into a global solution conforming to the correct global geometry and boundary conditions.[36] As a specific example, DCR is applied to the study of photoexcitation dynamics in amorphous molecular solid, specifically the fission of a spin-singlet exciton into spin-triplet excitons in amorphous diphenyl tetracene (DPT).[37] In contrast to singlet-fission (SF) processes in crystals and molecular dimers,[38] SF in amorphous molecular solid has not been studied theoretically. This is largely due to the required large quantum-mechanical calculations that capture nanostructural features. To address this challenge, we adopt a DCR approach, in which we perform nonadiabatic quantum molecular dynamics (NAQMD) simulations[39–43] embedded in amorphous DPT involving 6400 atoms in a DC manner. These local wave functions are used as a basis set to describe electronic excitations in the framework of linear-response time-dependent density functional theory (LR-TDDFT).[44] The NAQMD simulations describe coupled electron-ion dynamics involving nonadiabatic transitions between excited electronic states based on a surface-hopping approach.[45–48] NAQMD results on exciton dynamics are then augmented with time-dependent perturbation calculation of SF rates to provide inputs to kinetic Monte Carlo (KMC) simulation[49–52] of a global exciton-flow network to reach the experimental time scales.

This paper is organized as follows. Section II describes the DCR approach using NAQMD-KMC simulation as a specific example. Section III presents the LDC-DFT algorithm to accelerate the computation in the DC phase, along with scalable parallel implementation of the DCR algorithm. Numerical results are presented in Sec. IV, and finally Sec. V contains conclusion.

II. DIVIDE-CONQUER-RECOMBINE ALGORITHM

In a DCR-DFT algorithm, the DC phase constructs globally informed local solutions (Fig. 1(a)), which in the recombine phase are synthesized into a global solution (Fig. 1(b)).[36] Figure 1(a) shows a tree data structure, in which the root of the tree represents the entire simulation volume. The entire volume is subdivided into $2 \times 2 \times 2$ (or 2×2 in the 2D example in Fig. 1) subsystems (or cells) of equal volume. Recursive subdivision is repeated until each cell at the leaf level defines a DC domain. The long-range electrostatic interaction is computed in the DC phase using octree-based algorithms such as the fast multipole method (FMM)[1–4] and the multigrid method (MGM).[13,53,54] In FMM, each tree node is abstracted as a collection of multipole expansion coefficients of the charge distribution $\rho(\mathbf{r})$ and local expansion coefficients of the electrostatic potential $\phi(\mathbf{r})$. After computing the multipoles of all DC domains at the leaf level, the tree is traversed upward to compute the multipoles of all cells at all upper levels by combining those of lower levels. Subsequently, Taylor expansions of $\phi(\mathbf{r})$ around the centers of all cells at all levels are computed from the multipoles, starting from the root of the tree and traversing the tree downward. For each subsystem at each tree level, the multipoles of only a constant number of neighbor cells are used for the computation of $\phi(\mathbf{r})$ in order

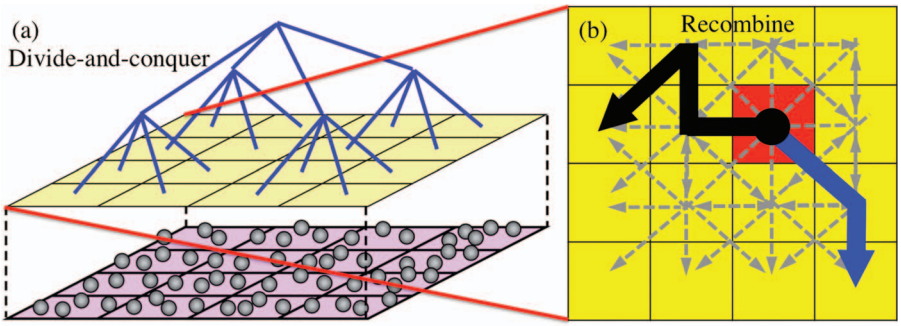

FIG. 1. Divide-conquer-recombine algorithm. (a) Two-dimensional schematic of the divide-and-conquer (DC) phase, in which atoms (spheres in the bottom plane) within a DC domain (each small magenta parallelogram in the bottom plane) are abstracted by collective variables such as a finite-difference representation of the charge distribution. The DC domains are recursively combined to form a tree data structure consisting of progressively coarser cells, until reaching the entire simulation volume at the root of the tree. Blue lines in the upper panel represent the merger of 2×2 daughter cells at a tree level to form a mother cell at the next upper level. A DC algorithm performs computation in $O(N)$ time 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.

to satisfy a prescribed error bound. Since the number of tree nodes is $O(N)$ and each node carries constant amount of computation, the computational complexity of FMM is $O(N)$. In MGM, each tree node is abstracted as a finite-difference representation of $\rho(\mathbf{r})$ and $\phi(\mathbf{r})$. The tree is traversed upward to obtain coarser representations of the Poisson equation to be iteratively solved, and downward to interpolate approximate coarser solutions of the $\phi(\mathbf{r})$ in $O(N)$ time. Electronic structure calculations solve local problems in different DC domains independently from each other. These local electronic structures are globally informed through the global electrostatic potential calculated on the tree, as well as the globally determined chemical potential for the electrons,[8] allowing charge transfer across DC domains.

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[55] among DC domains to account for higher inter-domain correlations that are not included in the tree data structure 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.[30] 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[30] and hypernetted-chain[31] 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.[32] 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.[32] DC wave functions were also used to describe global charge-migration dynamics by constructing coarse-grained electronic Hamiltonians with the use of quantum-dynamical[34] or bridge Green function[35] methods. 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.[36]

In this section, we first present a standard $O(N)$ DC-DFT algorithm. We then describe the recombination of local DC-DFT electronic structures to synthesize the global exciton dynamics with the use of NAQMD and KMC simulation methods.

A. Divide-and-conquer density functional theory

1. Domain decomposition

The DC-DFT algorithm[8–14] represents the three-dimensional space Ω as a union of overlapping spatial domains, $\Omega = \cup_\alpha \Omega_\alpha$, and physical properties are computed as linear combinations of domain properties (Fig. 2). Each domain Ω_α is further decomposed into its sub-volumes, $\Omega_\alpha = \Omega_{0\alpha} \cup \Gamma_\alpha$. Here, $\Omega_{0\alpha}$ is a non-overlapping core covering Ω (i.e., $\Omega = \cup_\alpha \Omega_{0\alpha}$ and $\Omega_{0\alpha} \cap \Omega_{0\beta} = 0$ $(\alpha \neq \beta)$), whereas Γ_α is a buffer layer that surrounds $\Omega_{0\alpha}$.

2. Partition of unity

For each domain α, we define a domain support function $p_\alpha(\mathbf{r})$ from $\mathbf{r} \in \Re^3$ ($\Re$ is the set of real numbers) to the

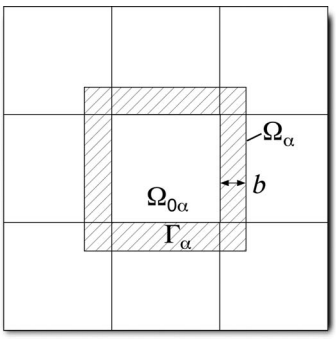

FIG. 2. Two-dimensional schematic of the DC-DFT algorithm. The physical space Ω is a union of overlapping domains, $\Omega = \cup_\alpha \Omega_\alpha$. Each domain Ω_α is further decomposed into a non-overlapping core $\Omega_{0\alpha}$ and a buffer layer Γ_α (see the shaded area). The thickness of the buffer layer is b.

unit interval [0, 1], which is compactly supported within the domain, $p_\alpha(\mathbf{r}) = 0$ ($\mathbf{r} \notin \Omega_\alpha$). The domain support functions constitute a partition of unity, i.e., they satisfy the sum rule, $\Sigma_\alpha p_\alpha(\mathbf{r}) = 1$, at every spatial position $\mathbf{r}$. The partition of unity allows the valence electron number density $\rho(\mathbf{r})$ to be exactly decomposed into

$$\rho(\mathbf{r}) = \sum_\alpha \rho_\alpha(\mathbf{r}), \tag{1}$$

where $\rho_\alpha(\mathbf{r}) = p_\alpha(\mathbf{r})\rho(\mathbf{r})$ is the partial contribution to the electron density from domain α. For the support function, we use a cubic interpolation function such that both the function value and its derivative is continuous at $p_\alpha(\mathbf{r}) = 0$ and 1.[10,13]

3. Subspace Hamiltonian approximation

The key approximation in DC-DFT is the replacement of the self-consistent Kohn-Sham Hamiltonian $\hat{H}$ by its subspace approximation $\hat{H}_\alpha$,[8]

$$\begin{aligned}\rho_\alpha(\mathbf{r}) &= p_\alpha(\mathbf{r})\langle\mathbf{r}|\frac{2}{\exp[(\hat{H}-\mu)/k_\mathrm{B}T]+1}|\mathbf{r}\rangle \\ &\cong p_\alpha(\mathbf{r})\langle\mathbf{r}|\frac{2}{\exp[(\hat{H}_\alpha-\mu)/k_\mathrm{B}T]+1}|\mathbf{r}\rangle,\end{aligned} \tag{2}$$

where k_B is the Boltzmann constant, T is the temperature, and the chemical potential μ is determined from the number of valence electrons N through the relation, $N = \int d\mathbf{r}\rho(\mathbf{r})$. The subspace Hamiltonian is defined through projection,

$$\hat{H}_\alpha = \int_{\Omega_\alpha} d\mathbf{r} \int_{\Omega_\alpha} d\mathbf{r}' |\mathbf{r}\rangle\langle\mathbf{r}|\hat{H}|\mathbf{r}'\rangle\langle\mathbf{r}'|, \tag{3}$$

where $|\mathbf{r}\rangle$ is the coordinate eigenstate.

We solve the Kohn-Sham (KS) equation within each domain,

$$\hat{H}_\alpha\psi_s^\alpha(\mathbf{r}) \equiv \left(-\frac{1}{2}\nabla^2 + \hat{V}_\mathrm{ion} + V_\mathrm{H} + \hat{V}_\mathrm{xc}\right)\psi_s^\alpha(\mathbf{r}) = \varepsilon_s^\alpha\psi_s^\alpha(\mathbf{r}), \tag{4}$$

with the orthonormality constraints, $\int d\mathbf{r}\psi_s^{\alpha*}(\mathbf{r})\psi_t^\alpha(\mathbf{r}) = \delta_{s,t}$ $= 1$ ($s = t$); 0 ($s \neq t$). In Eq. (4), $\psi_s^\alpha(\mathbf{r})$ is the sth KS orbital with the energy eigenvalue ε_s^α, ∇^2 is the Laplacian operator, $\hat{V}_\mathrm{ion}$ and $\hat{V}_\mathrm{xc}$ are the electron-ion and exchange-correlation (xc) potential operators, respectively, and the Hartree potential V_H is given by

$$V_\mathrm{H}(\mathbf{r}) = \int d\mathbf{r}' \frac{\rho(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|}. \tag{5}$$

The electronic ground state is determined self-consistently, i.e., the electron density is obtained iteratively until the input density $\rho_\mathrm{in}(\mathbf{r})$ becomes equal to the output density $\rho_\mathrm{out}(\mathbf{r})$ within a prescribed tolerance. Here, $\rho_\mathrm{in}(\mathbf{r})$ is used to calculate the KS potential, $\hat{V}^\mathrm{in} = \hat{V}_\mathrm{ion} + V_\mathrm{H}^\mathrm{in} + \hat{V}_\mathrm{xc}^\mathrm{in}$, whereas $\rho_\mathrm{out}(\mathbf{r})$ is calculated from Eqs. (1) and (2) using the KS orbitals, $\{\psi_s^\alpha(\mathbf{r})\}$, obtained by solving the KS equations, Eq. (4). The KS equations are solved iteratively using the conjugate-gradient method.[10,13] It should be noted that the local domain KS orbitals are globally informed through the global KS potential and chemical potential.

In order to apply the DC-DFT algorithm to problems involving coupled electronic excitations and nuclei motions, we incorporate several modifications: (1) long-range exact exchange correction (LC) for the xc potential for asymptotically correct description of electron-hole interactions;[56] (2) description of electronic excitations within Casida's LR-TDDFT,[40,41,44,57,58] using the ground-state LC-KS orbitals as a basis set; (3) efficient calculation of excited-state forces using a non-self-consistent approach;[59] and (4) surface-hopping approaches to describe nonadiabatic transitions between excited electronic states.[45–48,60] Below, these modifications are applied in a domain-by-domain manner, including the Coulombic interaction with frozen charge densities of the other domains. This approach is justified quantitatively for some systems as explained in Sec. IV C. Nevertheless, the recombine phase is indispensable for describing global properties (e.g., nonlocal correlations,[30,31] global electronic wave functions,[32,33] and charge[34,35] and exciton[36] transport) using these local electronic structures as a basis set.

4. Long-range exact exchange correction

In QMD simulations, where the electrons remain in their ground state, the generalized gradient approximation (GGA)[61] is typically used for the xc potential. For the description of excited electronic states, we include the LC through a range-separated hybrid exact exchange functional,[62] starting from the self-consistent GGA KS orbitals in the DC domains as explained above.[59] In a range-separated functional, the Coulomb-repulsion operator $1/r_{12}$ is divided into short-range and long-range parts using the error function:

$$\frac{1}{r_{12}} = \frac{1-\mathrm{erf}(\eta r_{12})}{r_{12}} + \frac{\mathrm{erf}(\eta r_{12})}{r_{12}}, \tag{6}$$

where $r_{12} = |\mathbf{r}_1 - \mathbf{r}_2|$ is the distance between two electrons at $\mathbf{r}_1$ and $\mathbf{r}_2$, and η is a range-separation parameter.[62] The range-separated xc energy functional is then given by

$$E_\mathrm{xc}'^\alpha = E_\mathrm{c,GGA}^\alpha + E_\mathrm{x,GGA}^{\mathrm{SR},\alpha} + E_\mathrm{x,HF}^{\mathrm{LR},\alpha}, \tag{7}$$

where $E_\mathrm{c,GGA}^\alpha$ is the GGA correlation energy functional, $E_\mathrm{x,GGA}^{\mathrm{SR},\alpha}$ is the short-range part of the GGA exchange energy functional, and $E_\mathrm{x,HF}^{\mathrm{LR},\alpha}$ is the long-range part of the Hartree-Fock (HF) exchange integral (i.e., using the second term in Eq. (6) as the Coulomb-repulsion operator).

In the non-self-consistent (NSC) approximation by Zhang *et al.*,[63] long-range corrected Hamiltonian matrix elements are constructed from the self-consistent GGA KS orbitals as

$$\begin{aligned}H_{st}'^\alpha = \delta_{st}\varepsilon_t^\alpha &- \left\langle\psi_s^\alpha\middle|V_\mathrm{x,GGA}^{\mathrm{LR},\alpha}[\rho(\mathbf{r})]\middle|\psi_t^\alpha\right\rangle \\ &- \sum_{i\in\{\mathrm{occupied}\}}\left[\psi_s^{\alpha*}\psi_i^\alpha|\mathrm{erf}(\mu r)/r|\psi_i^{\alpha*}\psi_t^\alpha\right],\end{aligned} \tag{8}$$

where $V_\mathrm{x,GGA}^{\mathrm{LR},\alpha}(\mathbf{r}) = \delta E_\mathrm{x,GGA}^{\mathrm{LR},\alpha}/\delta\rho(\mathbf{r})$, with $E_\mathrm{x,GGA}^{\mathrm{LR},\alpha}$ being the long-range part of the GGA exchange energy functional, and the Coulomb-like integral is defined as $[f.|h(r)|.g] \equiv \iint d\mathbf{r}d\mathbf{r}'f(\mathbf{r})h(|\mathbf{r}-\mathbf{r}'|)g(\mathbf{r}')$. We diagonalize the Hamiltonian

matrix, Eq. (8), as

$$\sum_t H'^{\alpha}_{st} U^{\alpha}_{tu} = U^{\alpha}_{su} \varepsilon'^{\alpha}_{u}, \tag{9}$$

to obtain long-range corrected KS (LC-KS) energies $\varepsilon'^{\alpha}_{s}$ and orbitals, $\psi'^{\alpha}_{s}(\mathbf{r}) = \sum_t \psi^{\alpha}_{t}(\mathbf{r}) U^{\alpha}_{ts}$.

In the NSC approximation, the compute-intensive exchange integrals in the last term of Eq. (8) are computed only once, thereby significantly reducing the excessive computational cost of the self-consistent range-separated hybrid exact exchange scheme. Here, the long-range interaction is computed using the reciprocal-space formalism of Martyna and Tuckerman[64] to avoid the interaction with electrons in the periodically repeated image cells.

5. Electronic excitations

We describe electronic excited states as a linear combination of electron-hole pairs within Casida's LR-TDDFT,[40,41,44,57,58] using the ground-state LC-KS orbitals as a basis set. It should be noted that, when considering the electronic excitation within each DC domain, the electronic charge density arising from all the other domains is kept frozen. In LR-TDDFT, electronic excitation energies are calculated from the poles of an electron-hole pair response function. This amounts to solving a non-Hermitian eigenvalue problem,[44]

$$\begin{pmatrix} \mathbf{A}^{\alpha} & \mathbf{B}^{\alpha} \\ \mathbf{B}^{\alpha *} & \mathbf{A}^{\alpha *} \end{pmatrix} \begin{pmatrix} \mathbf{X}^{\alpha}_{I} \\ \mathbf{Y}^{\alpha}_{I} \end{pmatrix} = \omega^{\alpha}_{I} \begin{pmatrix} \mathbf{1} & \mathbf{0} \\ \mathbf{0} & -\mathbf{1} \end{pmatrix} \begin{pmatrix} \mathbf{X}^{\alpha}_{I} \\ \mathbf{Y}^{\alpha}_{I} \end{pmatrix}, \tag{10}$$

where the eigenvalue ω^{α}_{I} is the Ith excitation energy in domain α, with the corresponding eigenvectors $\mathbf{X}^{\alpha}_{I}$ and $\mathbf{Y}^{\alpha}_{I}$. In Eq. (10), the elements of the $\mathbf{A}^{\alpha}$ and $\mathbf{B}^{\alpha}$ matrices are given by

$$A^{\alpha}_{ai\sigma,bj\tau} = \delta_{a,b}\delta_{i,j}\delta_{\sigma,\tau}\left(\varepsilon'^{\alpha}_{a\sigma} - \varepsilon'^{\alpha}_{i\sigma}\right) + K^{\alpha}_{ai\sigma,bj\tau}, \tag{11}$$

$$B^{\alpha}_{ai\sigma,bj\tau} = K^{\alpha}_{ai\sigma,jb\tau}, \tag{12}$$

where the indices i, j and a, b are used for occupied and virtual orbitals, respectively, σ, τ are spin variables, and $\varepsilon'^{\alpha}_{i\sigma}$ is the ith LC-KS orbital energy with spin σ. For a range-separated xc functional,[62,65] the coupling matrix elements in Eqs. (11) and (12) are given by

$$\begin{aligned} K^{\alpha}_{ai\sigma,bj\tau} = &\left[\psi'^{\alpha *}_{a\sigma}\psi'^{\alpha}_{i\sigma}|1/r|\psi'^{\alpha *}_{j\tau}\psi'^{\alpha}_{b\tau}\right] \\ &- \delta_{\sigma,\tau}\left[\psi'^{\alpha *}_{a\sigma}\psi'^{\alpha}_{b\tau}|\mathrm{erf}(\mu r)/r|\psi'^{\alpha *}_{j\tau}\psi'^{\alpha}_{i\sigma}\right] \\ &+ \int d\mathbf{r} \int d\mathbf{r}' \psi'^{\alpha *}_{a\sigma}(\mathbf{r})\psi'^{\alpha}_{i\sigma}(\mathbf{r}) \\ &\times \frac{\delta^2\left(E^{\alpha}_{\mathrm{xc,GGA}} - E^{\mathrm{LR},\alpha}_{\mathrm{x,GGA}}\right)}{\delta\rho_{\sigma}(\mathbf{r})\delta\rho_{\tau}(\mathbf{r}')} \psi'^{\alpha *}_{j\tau}(\mathbf{r}')\psi'^{\alpha}_{b\tau}(\mathbf{r}'), \end{aligned} \tag{13}$$

where $E^{\alpha}_{\mathrm{xc,GGA}}$ is the xc functional within GGA and $\rho_{\sigma}(\mathbf{r})$ is the electron density with spin σ.

According to the assignment ansatz by Casida,[44] the many-body wave function of the Ith excited state is given by[63,66]

$$\left|\Phi^{\alpha}_{I}\right\rangle = \sum_{i\in\{\mathrm{occupied}\}} \sum_{a\in\{\mathrm{unoccupied}\}} \sum_{\sigma} \frac{X^{\alpha}_{I,ai\sigma} + Y^{\alpha}_{I,ai\sigma}}{\sqrt{\omega^{\alpha}_{I}}} \hat{c}^{\alpha +}_{a\sigma}\hat{c}^{\alpha}_{i\sigma}\left|\Phi^{\alpha}_{0}\right\rangle, \tag{14}$$

where $|\Phi^{\alpha}_{0}\rangle$ is the Slater determinant of the occupied LC-KS orbitals, and $\hat{c}^{\alpha +}_{s\sigma}$ and $\hat{c}^{\alpha}_{s\sigma}$ are the creation and annihilation operators acting on the sth LC-KS orbital of spin σ in domain α. It should be noted that the binding of an electron-hole pair (or exciton) is described by summing ladder diagrams in many-body perturbative expansion, while the ring diagrams in the random phase approximation (RPA) describe electronic screening and collective excitations such as plasmons.[59] The two types of diagrams are interrelated through the exchange operation, i.e., by swapping the labels of electronic orbitals as is done between the first two terms in the right-hand side of Eq. (13). Both effects are thus included in LR-TDDFT using a range-separated hybrid xc functional that incorporates long-range exact exchange correction. We typically include 8 to 160 single-particle wave functions to represent each many-body wave function in Eq. (14).

6. Molecular dynamics

In QMD simulations, the equations of motion for atoms are integrated numerically, where interatomic forces are computed quantum mechanically based on the Hellmann-Feynman theorem.[67,68] For an excited electronic state, we use a NSC method to evaluate accurate forces at a moderate computational cost.[59] This method is an extension of the Harris-Foulkes approach adopted by Ref. 69 in a different context.

7. Nonadiabatic electron-ion dynamics

In NAQMD simulations with electronic transitions using Tully's fewest-switches surface-hopping (FSSH) method[45,46] along with the LC-KS representation of TDDFT, we calculate the time evolution of the many-electron wave function between consecutive nuclei-position updates. These equations are derived by expanding the electronic state $\Psi(t)$ at time t in terms of the electronic excited states $\Phi^{\alpha}_{J}(\mathbf{R}(t))$ in LR-TDDFT corresponding to the atomic configuration $\mathbf{R}(t)$ at time t:

$$|\Psi^{\alpha}(t)\rangle = \sum_J C^{(I),\alpha}_{J}(t)\left|\Phi^{\alpha}_{J}(\mathbf{R}(t))\right\rangle, \quad C^{(I),\alpha}_{J}(0) = \delta_{I,J}\,. \tag{15}$$

The time evolution of the expansion coefficients $C^{(I),\alpha}_{J}(t)$ is governed by

$$\frac{d}{dt}C^{(I),\alpha}_{J}(t) = -\sum_k \left(i\omega^{\alpha}_{K}\delta_{JK} + D^{\alpha}_{JK}\right)C^{(I),\alpha}_{k}(t), \tag{16}$$

where the nonadiabatic coupling (NAC) elements are defined as[40,41]

$$D^{\alpha}_{JK} = \left\langle \Phi^{\alpha}_{J}\left|\frac{\partial}{\partial t}\right|\Phi^{\alpha}_{K}\right\rangle. \tag{17}$$

The NAC elements are calculated using finite differencing from a pair of excited-state wave functions at consecutive time steps in an adiabatic QMD simulation. The switching probability from the Ith to the Jth adiabatic states is given by $P^{(I),\alpha}_{J}(t) = |C^{(I),\alpha}_{J}(t)|^2$, which is used to stochastically incur

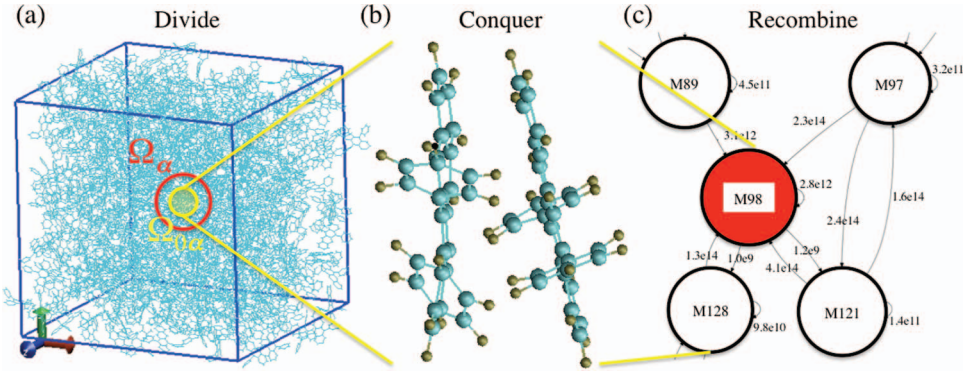

FIG. 3. Divide-conquer-recombine simulation to reach the experimental length and time scales for the exciton dynamics in amorphous molecular solid. (a) An entire configuration is subdivided into non-overlapping domain cores $\Omega_{0\alpha}$, each of which is augmented with neighbor molecules to form a domain Ω_α. (b) Nonadiabatic quantum molecular dynamics simulations are performed in each domain to extract the rate constants of various excitonic processes. (c) In the recombine phase, kinetic Monte Carlo 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.

interstate transitions. The use of the many-body wave function, Eq. (14), in Eq. (17) has been shown to be exact between ground and singly excited states, as well as between any pair of excited states when the Tamm-Dancoff approximation is used.[66] Various approaches have been proposed to improve FSSH for describing coupled electron-ion dynamics involving nonadiabatic electronic processes.[47,48] These approaches incorporate, e.g., quantum uncertainty for better accuracy. Specifically, we have implemented the decoherence-induced surface hopping (DISH) approach.[60]

B. Recombination

In order to explain the recombination phase of the DCR-NAQMD-KMC simulation, we use a specific example of amorphous solid 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 1000 atoms),[42,43] the entire simulation box Ω is subdivided into M non-overlapping spatial domains $\Omega_{0\alpha}$ (Fig. 3). 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, the rest of the system is represented by a fixed charge density. (A similar charge patching method[70] was used in large electronic-structure calculations in combination with a fragment method.[71]) 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.[10,14,72–74] Each NAQMD simulation starts from an electronic excited state corresponding to the excitation of an electron from the HOMO to the LUMO. Each NAQMD simulation is run typically for 200 fs.

1. Exciton-flow network

NAQMD trajectories are analyzed to obtain exciton-hopping rates between DPT molecules, where nonadiabatic electronic transitions between excited states are described by a surface hopping approach as described above. Specifically, when a nonadiabatic transition of an exciton occurs from molecule i to molecule j (the position of an exciton is determined as the center-of-mass position of the quasielectron and quasihole), the inverse of the preceding residence time on molecule i is used as the exciton-hopping rates $r^{\mathrm{hop}}_{i\to j}$. The overlapping domains in the DC approach allow the construction of a graph data structure that spans the entire amorphous DPT solid. In the graph, 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 NAC is also used to compute the exciton annihilation rate r_i^{AN}, at which an exciton residing on molecule i recombines to the electronic ground state. Note that the surface-hopping approach monitors accumulated transition probabilities from the current electronic excited state to the electronic ground state, from which the exciton annihilation rate is obtained. In addition to the phonon-assisted contribution to electronic transitions computed by NAQMD, we include the spontaneous emission contribution to r_i^{AN} calculated within a dipole approximation.[42] We also use time-dependent perturbation to compute the SF rate r_i^{SF}.[36]

2. Kinetic Monte Carlo

We perform first-principles KMC simulation[49–52] of exciton dynamics[42] using the hopping rates between DPT molecules as well as SF and annihilation rates, which are obtained from the NAQMD simulations and time-dependent perturbation calculations.[36] Each KMC simulation starts by placing an exciton on a randomly selected DPT molecule and resetting the time to be 0. At each KMC step, the exciton either: (i) hops to one of the nearest neighbor DPT molecules; (ii) annihilates to the ground state; or (iii) splits into two triplet excitons via SF. Let i be the molecule on which the

spin-singlet exciton resides at a given KMC step, and r_i^{AN} and r_i^{SF} be the annihilation and SF rates, respectively, for the exciton on the molecule. In addition, we have exciton-hopping rates $r_{i\to j}^{\mathrm{hop}}$ for all molecules j in the set of neighbor molecules of molecule i, *neighbor*(i). Let us denote the total rate of events as

$$r_i = r_i^{\mathrm{AN}} + r_i^{\mathrm{SF}} + \sum_{j\in neighbor(i)} r_{i\to j}^{\mathrm{hop}}. \tag{18}$$

The event to occur is chosen stochastically with the probability proportional to the corresponding rate:

$$i = \min_j \left\{ \sum_{k=1}^{j} \frac{r_k}{r} > u_1 \right\}, \tag{19}$$

where the cumulative rate is given by

$$r = \sum_k r_k. \tag{20}$$

After one of the hopping events is chosen, the time is incremented by

$$\Delta t = -\ln(\xi)/r, \tag{21}$$

where ξ is a uniform random number in the range [0, 1]. The simulation then continues by replacing molecule i by the destination molecule j of the chosen hopping. Else, if the exciton annihilates, the simulation terminates, and the number of singlet excitons is decreased by 1. Else (i.e., SF is chosen), the simulation also terminates. For the SF event, the number of singlet excitons is decreased by 1, while that of triplet excitons is increased by 2.

III. LEAN DIVIDE-AND-CONQUER ALGORITHM AND SCALABLE PARALLEL IMPLEMENTATION

In this section, we first present our LDC-DFT algorithm to significantly reduce the $O(N)$ prefactor of the DC-DFT algorithm. We then discuss scalable implementation of the LDC-DFT algorithm as well as its extension to DCR-DFT using a NAQMD-KMC approach on massively parallel computers.

A. Optimization of DC domain size through computational complexity analysis

In order to reduce the prefactor of the $O(N)$ computational cost of DC-DFT, we first optimize the size of the DC domains, based on an analysis of its computational cost.[23] The result is summarized in the following theorem.

Theorem. Consider a cubic system of side length L, and let the core length and the buffer thickness of a cubic domain in the DC-DFT algorithm be l and b, respectively (see Fig. 2). The computational complexity of the DFT computation within each domain is assumed to be the νth power of the system size. Then, the optimal domain size l_*, which incurs the minimal computational cost, is given by

$$l_* = \frac{2b}{\nu - 1}. \tag{22}$$

Proof. The number of domains is $N_{\mathrm{domain}} = (L/l)^3$, and the computational cost per domain can be written as

$$T_{\mathrm{domain}} = c(l + 2b)^{3\nu}, \tag{23}$$

with a prefactor c. The total computational cost as a function of l is thus

$$T_{\mathrm{comp}}(l) = T_{\mathrm{domain}} \cdot N_{\mathrm{domain}} = c\left(\frac{L}{l}\right)^3 (l + 2b)^{3\nu}. \tag{24}$$

The optimal domain size to minimize the computational cost is given by minimizing Eq. (24) with respect to l:

$$l_* = \arg\min[T_{\mathrm{comp}}(l)] = \frac{2b}{\nu - 1}, \tag{25}$$

which proves the theorem. □

The computational complexity of the DFT problem is $O(n^2)$ for typical domain sizes, where the number of atoms per domain is $n < 100$,[18] and thus $l_* = 2b$. Namely, the buffer size is usually chosen as half the core length of each domain. The asymptotic complexity, which has rarely been encountered in practical DFT calculations,[18] arises from the orthonormalization of KS orbitals and is $O(n^3)$. In this limit, $l_* = b$.

B. Optimization of DC buffer thickness through error analysis and a density adaptive boundary condition: LDC-DFT

The choice of the buffer thickness b is dictated by accuracy requirement. The quantum nearsightedness principle[22] indicates that the error involved in the DC-DFT algorithm (which is due to the artificial boundary condition imposed at the domain boundary $\partial\Omega_\alpha$) decays exponentially as a function of b.[23] Due to the artificial boundary condition at $\partial\Omega_\alpha$, the domain density $\rho_\alpha(\mathbf{r})$ deviates from the total density $\rho(\mathbf{r})$. Let λ be the exponential decay constant of the density perturbation, $\Delta\rho_\alpha(\mathbf{r}) = \rho_\alpha(\mathbf{r})/p_\alpha(\mathbf{r}) - \rho(\mathbf{r})$, away from $\partial\Omega_\alpha$. In the case of an insulating material, λ is related to that of the Wannier function of the highest occupied band.[23] Suppose that an error tolerance of $\varepsilon\langle\rho_\alpha(\mathbf{r})\rangle$ is imposed on $|\Delta\rho_\alpha(\mathbf{r})|$ at the periphery of $\Omega_{0\alpha}$, where $\langle\rho_\alpha(\mathbf{r})\rangle$ is the average density in Ω_α. To satisfy the error tolerance, the buffer depth needs to be as large as

$$b = \lambda \ln\left(\frac{\max\{|\Delta\rho_\alpha(\mathbf{r})||\mathbf{r}\in\partial\Omega_\alpha\}}{\varepsilon\langle\rho_\alpha(\mathbf{r})\rangle}\right). \tag{26}$$

Since the computational complexity of the DC-DFT algorithm scales with the buffer depth asymptotically as $b^{3\nu} = b^6 \sim b^9$ (see Eq. (23)), the large b value required for obtaining a sufficient accuracy in energy (e.g., 10^{-3} a.u. per atom) represents a major computational bottleneck. According to Eq. (26), key to reducing the prefactor of the $O(N)$ computational cost of DC-DFT is to minimize $\Delta\rho_\alpha(\mathbf{r})$ at $\partial\Omega_\alpha$. We address this problem through an improved treatment of domain boundaries as explained below.

As stated in Sec. II, the key approximation of DC-DFT is the replacement of the global KS Hamiltonian $\hat{H}$ by a local KS Hamiltonian $\hat{H}_\alpha$ for each spatially localized domain α. The local Hamiltonian approximation in Eq. (3) can be interpreted as either projection or a boundary potential. In the

projection formalism, the domain Hamiltonian

$$\hat{H}_\alpha = \sum_{s,t\in\Omega_\alpha} |s\rangle\langle s|\hat{H}|t\rangle\langle t| \tag{27}$$

is expanded by basis functions $|s\rangle$ that are localized in the domain Ω_α. In projection approaches using molecular orbitals, purification of dangling bonds is often performed by filtering out eigenstates with outlier eigenvalues.[32] However, this requires case-by-case treatment with prior knowledge of the expected energy spectrum.

To have general recipes independent of material, geometry, and basis, we take another view of the local Hamiltonian, i.e., a boundary condition on the electronic wave functions $\psi_s^\alpha(\mathbf{r})$ imposed at the domain boundary $\partial\Omega_\alpha$. The boundary condition can alternatively be considered as a boundary potential included in the domain KS Hamiltonian. An example of the boundary potential is a hard-wall (or Dirichlet) boundary condition:

$$v_\alpha^{\rm bc}(\mathbf{r}) = \begin{cases} 0 & (\mathbf{r}\in\Omega_\alpha) \\ \infty & (\mathbf{r}\notin\Omega_\alpha) \end{cases}. \tag{28}$$

An improved boundary potential may be obtained as follows. According to the Hohenberg-Kohn theorem,[17] the local density $\rho_\alpha(\mathbf{r})/p_\alpha(\mathbf{r})$ corresponds to a unique external potential $v(\mathbf{r})$, which is distinct from that corresponding to the global density $\rho(\mathbf{r})$. To reduce the discrepancy $\Delta\rho_\alpha(\mathbf{r}) = \rho_\alpha(\mathbf{r})/p_\alpha(\mathbf{r}) - \rho(\mathbf{r})$, we use a linear-response formula for the boundary potential,

$$v_\alpha^{\rm bc}(\mathbf{r}) = \int d\mathbf{r}' \frac{\partial v(\mathbf{r})}{\partial\rho(\mathbf{r}')}\Delta\rho_\alpha(\mathbf{r}'). \tag{29}$$

We adopt a local approximation,[75]

$$\frac{\partial v(\mathbf{r})}{\partial\rho(\mathbf{r}')} = \frac{\delta(\mathbf{r}-\mathbf{r}')}{\xi}, \tag{30}$$

so that

$$v_\alpha^{\rm bc}(\mathbf{r}) = \frac{\Delta\rho^\alpha(\mathbf{r})}{\xi}. \tag{31}$$

The same density-template potential as Eq. (31) was used by Ohba *et al.*[14] to augment the hard-wall boundary potential, Eq. (28). Here, we instead use the periodic boundary condition on the KS wave functions, incorporating the boundary potential, Eq. (31), in the KS Hamiltonian

$$\hat{H}_\alpha = -\frac{1}{2}\nabla^2 + \hat{V}_{\rm ion} + V_{\rm H} + \hat{V}_{\rm xc} + v_\alpha^{\rm bc}(\mathbf{r}) \tag{32}$$

to replace Eq. (4). The local approximation may be justified by the quantum nearsightedness principle[22] as formulated by Prodan and Kohn.[23] Namely, the response kernel $\partial v(\mathbf{r})/\partial\rho(\mathbf{r}')$ is short-ranged with respect to $|\mathbf{r}-\mathbf{r}'|$. In the numerical tests in this paper, we take the adjustable parameter ξ in Eq. (31) to be 0.333 in the atomic unit.[14]

C. Efficient and scalable numerical implementation and parallelization

1. Hybrid real-reciprocal (HR^2) space approach

Our implementation of the LDC-DFT algorithm 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:

1. A plane-wave basis[20] is used to represent local electronic wave functions and charge density within each domain (containing ∼100 atoms), which takes advantage of a highly efficient numerical implementation based on fast Fourier transform.[76]
2. A real-space multigrid is used to represent the global charge density of the total system, which is highly scalable on massively parallel computers due to the locality preserving octree data structure.[10,13]

Within each domain, the electronic pseudo-wave functions and the pseudo-charge density are expanded by plane waves. As in our previous DC-DFT algorithm,[10,13] the Hartree potential in the KS potential, Eq. (32), is obtained globally from the total density $\rho(\mathbf{r})$ using a real-space multigrid method. To prevent charge sloshing in large systems, we combine the Pulay charge mixing method with a real-space formulation of the Kerker preconditioning.[77]

2. Hybrid space-band (HSB) decomposition on parallel computers

The LDC-DFT QMD simulation code is implemented on massively parallel computers by employing two levels of parallelism. At the coarser level, we use spatial decomposition among overlapping domains Ω_α. The program is implemented using the message passing interface (MPI) library for interprocessor communications, and each domain is assigned a dedicated MPI communicator (i.e., programming construct that combines a process group and a system-defined context identifier) using a MPI_COMM_SPLIT call. At the finer level, the plane-wave-basis calculations within each domain are further parallelized by a hybrid approach combining spatial decomposition (i.e., distributing real-space or reciprocal-space grid points among processors) and band decomposition (i.e., assigning the calculations of different KS orbitals to different processors)[67,68] within the communicator assigned to the domain. The iterative band-by-band minimization is carried out by band decomposition, where the eigenenergy of each band is minimized in each processor independently from the others under an approximate orthonormal condition. The electron density is also calculated by band decomposition. On the other hand, the KS orbitals are orthonormalized by first constructing an overlap matrix between them using reciprocal-space decomposition, where the Fourier components of the wave functions are distributed among multiple processors. This is followed by parallel Cholesky decomposition of the overlap matrix, which introduces an additional parallelization axis. To switch between the spatial and band decompositions, all-to-all communications are required only within the communicator. In addition, global communication

within the communicator is necessary to calculate the scalar products between the wave functions.

3. Parallelization of KMC

The naive KMC algorithm presented in Sec. II B is not scalable to larger system sizes. The cumulative event rate in Eq. (20) grows as $O(N)$, and accordingly the time scale of the simulation determined by its inverse becomes progressively smaller in larger systems. To overcome this scaling problem, KMC simulation is parallelized approximately in a DC fashion as well, using a synchronous formulation and graph coloring to avoid conflicting events.[78] Here, the cumulative rate in Eq. (20) is computed for each non-overlapping domain $\Omega_{0\alpha}$ as r_α, and then the global maximum rate is defined as $r_{\max} = \max_\alpha\{r_\alpha\}$. KMC simulation is performed concurrently among all the domains similarly to the sequential algorithm in Sec. II B, except that a null event (where nothing occurs) is generated with the rate of $r_{0\alpha} = r_{\max} - r_\alpha$. This allows global synchronous time evolution, where the time is incremented by $\Delta t = -\ln(\xi)/r_{\max}$ at each KMC step. To avoid conflicting events to occur between neighboring domains, the domains are colored so that no domains of the same color are adjacent to each other. Each KMC simulation step then consists of a loop over colors, where at each color step events occur only in the domains of the chosen color.

IV. NUMERICAL RESULTS

A. Numerical convergence and computational cost

We first test the convergence of calculation with respect to the buffer length b, which controls the data locality of the algorithm (see Fig. 2). Figure 4 shows the calculated potential energy as a function of b for an amorphous cadmium selenide (CdSe) system containing 512 atoms in a cubic simulation box of length 45.664 a.u. The side length l of each cubic domain is fixed as 11.416 a.u. To study the effect of boundary conditions on the energy convergence with respect to the localization control parameter b, we compare results of two calculations: (1) our original DC-DFT algorithm with the Dirichlet boundary condition, where the KS wave functions $\psi_s^\alpha(\mathbf{r})$ in domain α are made zero at the domain boundary $\partial\Omega_\alpha$; and (2) the LDC-DFT algorithm that uses the periodic boundary condition within Ω_α along with the boundary potential, Eq. (31). We see that the LDC-DFT calculation converges much more rapidly than the DC-DFT calculation. The LDC-DFT potential energy converges within 10^{-3} a.u. per atom above $b = 4$ a.u.

In order to assess the impact of the improved convergence of LDC-DFT on the computational cost, let us consider an example of the error tolerance of 5×10^{-3} a.u. for the energy. According to Fig. 4, the buffer length b to achieve this convergence criterion is decreased from 4.73 a.u. for DC-DFT to 3.57 a.u. for LDC-DFT. According to Eq. (24), this amounts to the computational speedup by a factor of $[(11.416 + 2 \times 4.72)/(11.416 + 2 \times 3.57)]^{3\nu} = 2.03$ (for $\nu = 2$) $- 2.89$ (for $\nu = 3$).

The convergence property of the LDC-DFT algorithm has been tested for other material systems as well. Figure 5 shows the calculated potential energy as a function of the buffer length b for crystalline and amorphous silicon carbide (SiC) systems. For crystalline SiC, we consider both the zero-temperature configuration and a QMD configuration at a temperature of 300 K. Here, the total system contains 128 atoms in a rectangular cell of side lengths $16.61 \times 16.61 \times 33.22$ a.u., and the side length l of each cubic domain is fixed as 8.305 a.u. The potential energy converges within 10^{-3} a.u. per atom above $b = 4$ a.u. for all systems, where more rapid convergence is observed for the amorphous system. In all cases, the LDC-DFT calculation values converge to conventional plane-wave DFT calculation values for large b.

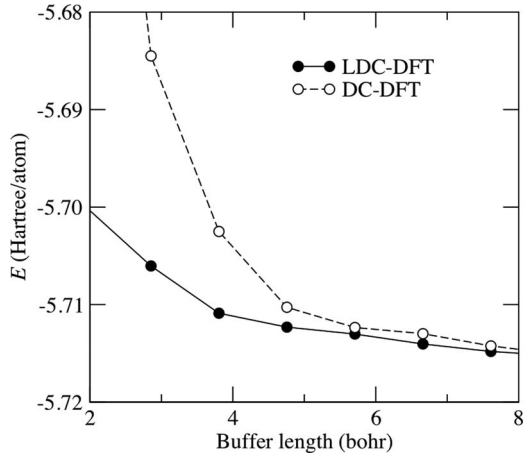

FIG. 4. Effects of boundary conditions on the energy convergence with respect to the localization control parameter. Potential energy is plotted as a function of the buffer length b for an amorphous CdSe system (512 atoms in a cubic simulation box of side length 45.664 a.u.). The solid and open circles represent results for the LDC-DFT and original DC-DFT algorithms, respectively. The domain size is fixed as 11.416 a.u. The atomic units are used for both energy and length.

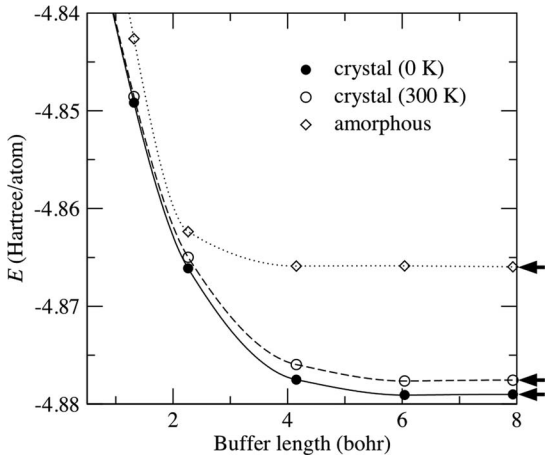

FIG. 5. Potential energy as a function of the buffer length b for crystalline and amorphous SiC systems, where each system contains 128 atoms in a rectangular cell of side lengths, $16.61 \times 16.61 \times 33.22$ a.u. The cubic domain size is fixed as 8.305 a.u. Arrows indicate conventional plane-wave calculation values with cutoff energies of 30 and 250 Ry, respectively, for the electronic pseudo-wave functions and the pseudo-charge density.

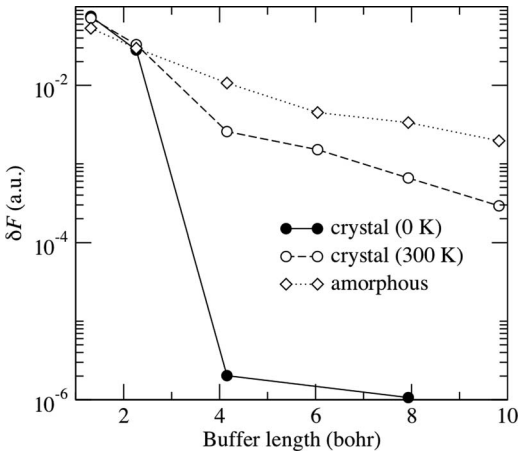

FIG. 6. Effects of boundary conditions on the force convergence with respect to the localization control parameter. The average of the deviations of the atomic forces $\{\delta\mathbf{F}_i = |\mathbf{F}_i - \mathbf{F}_i^{\mathrm{ref}}|\}$ obtained by the LDC-DFT method for crystal and amorphous SiC is shown as a function of the buffer length b.

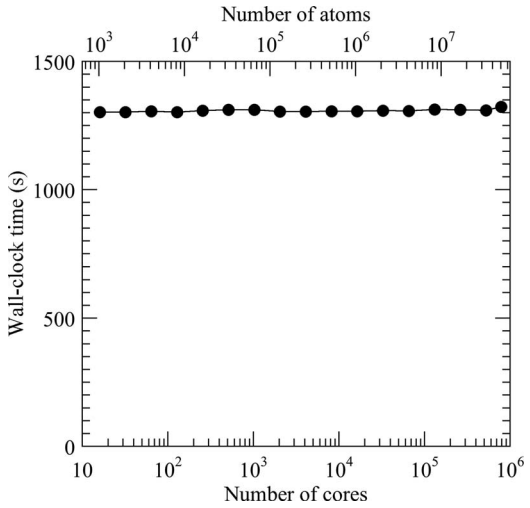

FIG. 7. Wall-clock time per QMD simulation step of the parallel LDC-DFT algorithm, with scaled workloads—64P-atom SiC system on P cores ($P = 16, \ldots, 786\,432$) of Blue Gene/Q.

We also analyze the accuracy of atomic forces in the LDC-DFT method. We define the deviations of the forces on atoms $\delta\mathbf{F}_i = |\mathbf{F}_i - \mathbf{F}_i^{\mathrm{ref}}|$ from the reference values $\mathbf{F}_i^{\mathrm{ref}}$ calculated by a conventional plane-wave DFT method with the same plane-wave cutoff. Figure 6 shows the average of $\delta\mathbf{F}_i$ over all atoms i as a function of the buffer length b for crystalline and amorphous SiC systems. Both maximum and average of $\delta\mathbf{F}_i$ decrease significantly as b increases. The maximum of $\delta\mathbf{F}_i$ is less than 0.01 a.u. for $b = 8.7$ a.u. for all systems.

To test the applicability of the LDC-DFT algorithm to QMD simulations, we have calculated the total energy as a function of time for a 432-atom liquid rubidium system, where the equations of motion for the nuclei are integrated using the velocity Verlet algorithm with a time step of $\Delta t = 4.8$ fs. The result is comparable (but with much less computational time) to that with our previous DC-DFT algorithm.[13] The total energy is conserved within 2×10^{-4} a.u./ps per atom, which is two orders-of-magnitude smaller than the variation in the potential energy.

B. Scalability on massively parallel computers

Next, we test the scalability of the parallel LDC-DFT algorithm for SiC crystal. The program is written in Fortran 90 with MPI for massage passing. Numerical tests are performed on the IBM Blue Gene/Q computer at the Argonne National Laboratory. The Blue Gene/Q 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,[79] 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 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 perform an isogranular scaling benchmark of LDC-DFT code on the Blue Gene/Q, in which the number of atoms per core N/P is kept constant. Figure 7 shows the wall-clock time per QMD simulation step with scaled workloads—64P-atom SiC system on P cores of Blue Gene/Q. The execution time includes 3 self-consistent (SC) iterations to determine the electronic wave functions and the KS potential, with 3 CG iterations per SC cycle to refine each wave function. 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 LDC-DFT algorithm as a product of the total number of atoms and the number of self-consistent iterations executed per second. The isogranular speedup is given by the ratio between the speed of P cores and that of 16 cores as a reference system. The weak-scaling parallel efficiency is the isogranular speedup divided by P. With the granularity of 64 atom per core for SiC systems, the parallel efficiency is 0.984 on $P = 786\,432$ for a 50 331 648-atom SiC system. This demonstrates the high scalability of the LDC-DFT algorithm.

Compared with the above benchmark tests, the number of atoms involved in production QMD simulations is orders-of-magnitude smaller. An example is our 16 661-atom QMD simulation of a $Li_{441}Al_{441}$ particle in liquid water (Fig. 8(a)). Hydrogen production using Al particles in water could provide a renewable energy cycle, with a potential application to on-board hydrogen production for hydrogen-powered vehicles.[67] However, the major problems are its low reaction rate and poor yield. We are investigating the possibility of accelerated high-yield production of H_2 from water using LiAl alloy particles instead. Figure 8(b) shows the number of H_2 molecules produced during the QMD simulation at a temperature of 1500 K. In total of 42 H_2 molecules are

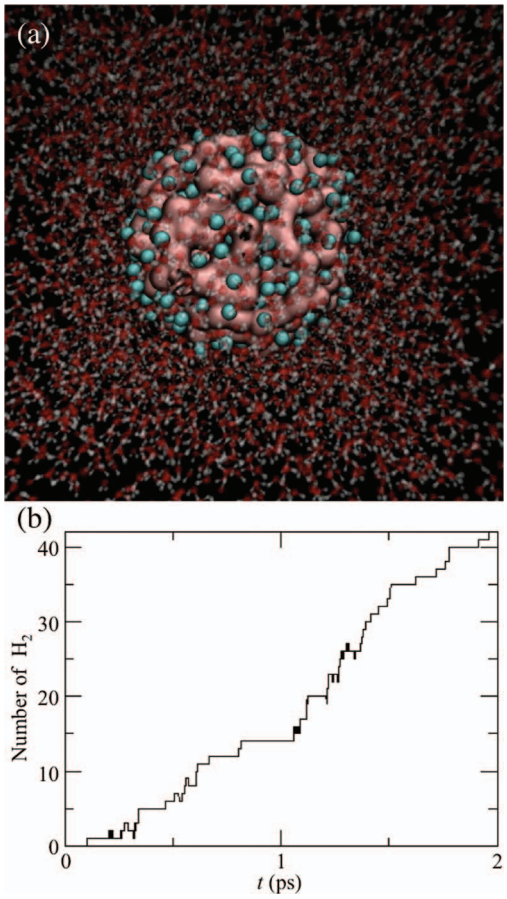

FIG. 8. 16 661-atom QMD simulation on Blue Gene/Q based on LDC-DFT. (a) A snapshot at 1 ps, where white, red and cyan spheres are H, O and Li atoms, respectively, whereas the valence charge density colored in magenta is centered at Al atoms. (b) The number of produced H_2 molecules as a function of time.

produced within 2 ps. From this plot, the H_2 production rate per LiAl pair is estimated to be $k = (5.94 \pm 0.02) \times 10^{10}$ (s^{-1}). We have also studied the dependence of this rate on the temperature T using a smaller particle, $Li_{30}Al_{30}$. The calculated rate is fitted well with the Arrhenius relation, $k(T) = k_0 \exp(-\Delta/k_B T)$, where Δ is the activation barrier and k_0 is a constant. The best fit yields $\Delta = 0.0676$ eV. This is an order-of-magnitude lower than the calculated activation barrier of 0.3 eV for the H_2 production from water using Al nanoparticles, where the water-splitting reaction was the rate limiting process.[67] The estimated H_2 production rate from this fit is 1.04×10^9 (s^{-1}) at 300 K, which is much higher than 10^7 s^{-1} calculated for Al.[67] The simulation results thus show a dramatically accelerated reaction of Li-alloyed Al particles compared with pure Al particles for hydrogen production from water.

C. Application to singlet fission in amorphous organic solid

We have applied the DCR paradigm to the study of photoexcitation dynamics in amorphous DPT solid. The system consisting of 6400 atoms (or 128 DPT molecules) is prepared in a cubic simulation box with the side length of 43.3 Å by the melt-quench procedure in molecular dynamics (MD) simulation. Periodic boundary conditions are applied in all Cartesian directions. To provide a basis set for NAQMD simulations, we first obtain electronic ground states using the projector-augmented-wave method.[80,81] Projector functions are generated for the 2s and 2p states of C, and the 1s state of H. The GGA[61] is used for the exchange-correlation energy with nonlinear core corrections.[82] The electronic pseudo-wave functions and the pseudo-charge density are expanded by plane waves with cutoff energies of 30 and 250 Ry, respectively. We then include self-interaction correction based on a range-separated hybrid functional that includes a long-range exact exchange potential.[62] Starting from the amorphous DPT configuration prepared by the MD simulation, NAQMD simulations are carried out in the canonical ensemble. The equations of motion for nuclei are integrated numerically with a time step of 20 a.u. (~0.48 fs), whereas the time-dependent Schrödinger equation, Eq. (16), for electrons is integrated with a time step of 0.04 a.u. NAQMD trajectories are analyzed to obtain exciton-hopping rates between DPT molecules. The NAC 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.[42] In addition to the exciton hopping and annihilation rates, we estimate the SF rate of each singlet excitonic state using a time-dependent perturbation theory.[83] We then perform first-principles KMC simulations of exciton dynamics using the calculated hopping rates between DPT molecules as well as the SF and annihilation rates.[19] In total of 5000 KMC simulations are performed to take statistics.

Figure 9(a) shows the calculated population dynamics of singlet and triplet excitons. Our first-principles KMC

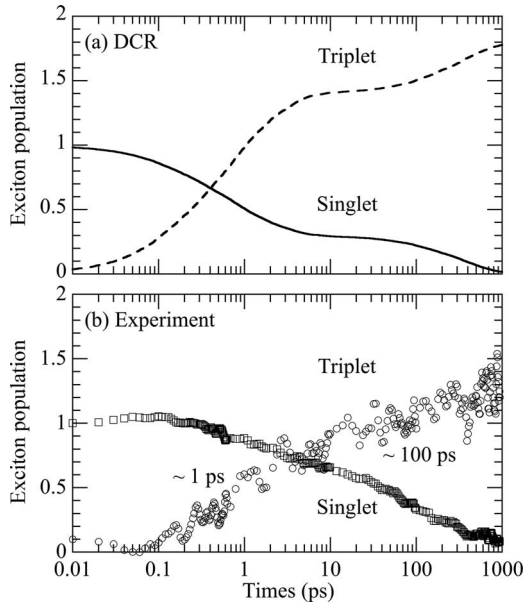

FIG. 9. (a) Time evolution of the population of singlet (solid line) and triplet (dashed line) excitons obtained by the DCR approach using NAQMD-informed KMC simulation. (b) The same quantity observed experimentally for singlet (squares) and triplet (circles) excitons in Ref. 37. Adapted with permission from W. Mou *et al.*, Appl. Phys. Lett. **102**(17), 177301 (2013). Copyright 2013 American Institute of Physics.

result captures key features in the experimental data in Ref. 37; see Fig. 9(b). In particular, the NAQMD-informed KMC data reproduces the two time constants observed experimentally. Namely, rapid singlet-exciton decay and triplet-exciton generation occur within ∼1ps, followed by slower processes on the time scale of ∼100 ps. In both simulation and experimental data, the number of generated triplet excitons is larger than that of the initial singlet excitons, signifying efficient SF in amorphous DPT.

The observed double exponential behavior in Fig. 9(b) was previously interpreted by postulating the presence of a subset of molecules only where SF can occur.[37] Singlet excitons photoexcited near these sites rapidly undergo fission in ∼1 ps, while those generated elsewhere must diffuse to these sites via a slow diffusion process. Our first-principles KMC result also exhibits a similar two-stage population dynamics without any fitting parameter. Analysis of the simulation data confirms the existence of the postulated SF hot spots. Namely, 91% of all the SF events are accounted for by 3.9% of the DPT molecules. Detailed analyses of the data revealed geometric characteristics of SF hot spots: Twist-stack conformation of DPT molecular dimers. In addition to the geometric indicators for high SF rates, we found that the topology of the exciton-flow network influences the hot spots. Namely, a molecule with a high reverse k-nearest neighbor degree (i.e., to how many other molecules, a given molecule is within the k-th nearest neighbors) acts as a hub of the exciton-flow network, to which a large number of excitons flow into. When a network hub coincides with a high SF-rate site, the site acts as a SF hot spot. This result highlights the effectiveness of the DCR approach. While the exciton-flow network has been constructed locally in a DC manner, the calculated DCR exciton dynamics on the combined graph data structure reflects a system-wide network-topology such as the distribution of network hubs.[36]

In this simulation, quantum coherence was incorporated by solving the time-dependent Schrödinger equation, Eq. (16), using the adiabatic basis, while decoherence due to coupling to nuclei was taken into account using the DISH approach.[60] The calculated decoherence time was on the order of 10 fs, indicating the lack of extensive quantum coherence. This is consistent with experimental observations, where both optical absorption and emission spectra for vapor deposited DPT films are nearly identical to those of single DPT molecules.[37]

The domain-by-domain description of electronic excitations described in Sec. II A is well justified in the amorphous DPT system. To do so, we quantify exciton localization using the participation number, $n_{\mathrm{p}} = 1/\Sigma_i\ p_i^2$ (p_i is the existing probability of the exciton in the ith molecule). This quantity reflects the number of molecules, over which the exciton is spread.[36] It mostly takes a value near unity, indicating that the exciton is localized on one DPT molecule, except for short transient times when the exciton extends over 2 molecules (i.e., $n_{\mathrm{p}} \sim 2$). The above mentioned optical absorption and emission spectra also confirm the highly localized nature of excitons in amorphous DPT.[37]

V. CONCLUSION

We have developed a highly scalable computational approach named DCR to describe large spatiotemporal-scale material processes using local electronic structures from the DC-DFT algorithm as a basis set. This paper presented a LDC-DFT algorithm that reduced the prefactor of the $O(N)$ computational cost by a factor of 2−3 over that of our previous DC-DFT algorithm. The LDC-DFT algorithm achieved an isogranular parallel efficiency of 0.984 on 786 432 IBM Blue Gene/Q cores for a 50.3×10^6-atom system. The algorithm was also used for 16 661-atom QMD simulation of hydrogen production from water using LiAl alloy particles. The DCR approach was demonstrated for SF of excitons in a 6400-atom amorphous DPT with the use of NAQMD and KMC simulation methods. The calculated exciton population dynamics agreed well with ultrafast transient absorption measurements, and revealed the molecular origin of the experimentally observed two time scales.

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.[84–86] These problems can be addressed by multiscale modeling such as the QM/MM method,[87,88] in which more accurate models are embedded in coarser models only where high accuracy is required. We have used the DC algorithmic framework to develop an adaptive multiscale dynamics approach that combines QMD, MD, and finite-element dynamics.[89] In this approach,[90] DC-DFT based QMD simulations are invoked within coarser simulations on demand on the basis of an additive hybridization scheme.[91] Furthermore, the number of atoms in the QMD regions changes dynamically in response to the progress of chemical reactions.[90] In these multiscale dynamics applications, reactive MD approaches based on environment-dependent reactive force fields have become an essential intermediate layer between QMD and MD methods.[92–96] In addition to the multiscale modeling, innovative recombination of DC solutions can be incorporated into high-throughput screening to explore a large combinatorial search space for discovering new materials.[97,98]

ACKNOWLEDGMENTS

This research was supported by the Department of Energy (DOE), Office of Science, Basic Energy Sciences, Materials Science and Engineering Division, Grant No. DE-FG02-04ER-46130. The algorithms were developed at the Center for High Performance Computing and Communications of the University of Southern California. Benchmark tests and production runs 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, Nichols Romero, and Robert Walkup for their help on the use of the Blue Gene/Q.

[1]L. Greengard and V. Rokhlin, J. Comput. Phys. **73**(2), 325–348 (1987).
[2]C. A. White and M. Headgordon, J. Chem. Phys. **101**(8), 6593–6605 (1994).

[3]A. Nakano, R. K. Kalia, and P. Vashishta, Comput. Phys. Commun. **83**(2–3), 197–214 (1994).
[4]S. Ogata, T. J. Campbell, R. K. Kalia, A. Nakano, P. Vashishta, and S. Vemparala, Comput. Phys. Commun. **153**(3), 445–461 (2003).
[5]J. J. M. Cuppen, Numer. Math. **36**(2), 177–195 (1981).
[6]W. N. Gansterer, R. C. Ward, and R. P. Muller, ACM Trans. Math. Software **28**(1), 45–58 (2002).
[7]A. Nakano, Comput. Phys. Commun. **104**(1–3), 59–69 (1997).
[8]W. T. Yang, Phys. Rev. Lett. **66**(11), 1438–1441 (1991).
[9]S. L. Dixon and K. M. Merz, J. Chem. Phys. **107**, 879 (1997).
[10]F. Shimojo, R. K. Kalia, A. Nakano, and P. Vashishta, Comput. Phys. Commun. **167**(3), 151–164 (2005).
[11]T. Ozaki, Phys. Rev. B **74**(24), 245101 (2006).
[12]M. Kobayashi and H. Nakai, J. Chem. Phys. **129**(4), 044103 (2008).
[13]F. Shimojo, R. K. Kalia, A. Nakano, and P. Vashishta, Phys. Rev. B **77**(8), 085103 (2008).
[14]N. Ohba, S. Ogata, T. Kouno, T. Tanmura, and R. Kobayashi, Comput. Phys. Commun. **183**(8), 1664–1673 (2012).
[15]An exaflop/s computer can operate 10^{18} arithmetic operations per second.
[16]K. Nomura, H. Dursun, R. Seymour, W. Wang, R. K. Kalia, A. Nakano, P. Vashishta, F. Shimojo, and L. H. Yang, *Proceedings of the International Parallel and Distributed Processing Symposium*, IPDPS 2009 (IEEE, 2009).
[17]P. Hohenberg and W. Kohn, Phys. Rev. **136**(3), B864–B871 (1964).
[18]F. Gygi, E. Draeger, B. R. de Supinski, R. K. Yates, F. Franchetti, S. Kral, J. Lorenz, C. W. Ueberhuber, J. A. Gunneis, and J. C. Sexton, *Proceedings of Supercomputing*, SC05 (ACM/IEEE, 2005).
[19]R. Car and M. Parrinello, Phys. Rev. Lett. **55**, 2471–2474 (1985).
[20]M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos, Rev. Mod. Phys. **64**(4), 1045–1097 (1992).
[21]D. R. Bowler and T. Miyazaki, Rep. Prog. Phys. **75**(3), 036503 (2012).
[22]W. Kohn, Phys. Rev. Lett. **76**(17), 3168–3171 (1996).
[23]E. Prodan and W. Kohn, Proc. Natl. Acad. Sci. U.S.A. **102**(33), 11635–11638 (2005).
[24]M. Benzi, P. Boito, and N. Razouk, SIAM Rev. **55**(1), 3–64 (2013).
[25]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, and L. H. Yang, Intl. J. High Performance Comput. Appl. **22**(1), 113–128 (2008).
[26]A related $O(N)$ algorithm called density fragment interaction has been developed for large molecular systems by Weitao Yan and collaborators: X. Q. Hu, Y. D. Jin, X. C. Zeng, H. Hu, and W. T. Yang, Phys. Chem. Chem. Phys. **14**(21), 7700–7709 (2012).
[27]T. V. T. Duy and T. Ozaki, Comput. Phys. Commun. **185**(3), 777–789 (2014).
[28]M. J. Cawkwell and A. M. N. Niklasson, J. Chem. Phys. **137**(13), 134105 (2012).
[29]F. Shimojo, A. Nakano, R. K. Kalia, and P. Vashishta, Appl. Phys. Lett. **95**(4), 043114 (2009).
[30]Y. Okiyama, T. Tsukamoto, C. Watanabe, K. Fukuzawa, S. Tanaka, and Y. Mochizuki, Chem. Phys. Lett. **566**, 25–31 (2013).
[31]S. Tanaka, C. Watanabe, and Y. Okiyama, Chem. Phys. Lett. **556**, 272–277 (2013).
[32]S. Tsuneyuki, T. Kobori, K. Akagi, K. Sodeyama, K. Terakura, and H. Fukuyama, Chem. Phys. Lett. **476**(1–3), 104–108 (2009).
[33]T. Kobori, K. Sodeyama, T. Otsuka, Y. Tateyama, and S. Tsuneyuki, J. Chem. Phys. **139**(9), 094113 (2013).
[34]C. Gollub, S. Avdoshenko, R. Gutierrez, Y. Berlin, and G. Cuniberti, Isr. J. Chem. **52**(5), 452–460 (2012).
[35]H. Kitoh-Nishioka and K. Ando, J. Phys. Chem. B **116**(43), 12933–12945 (2012).
[36]W. Mou, S. Hattori, P. Rajak, F. Shimojo, and A. Nakano, Appl. Phys. Lett. **102**(17), 173301 (2013).
[37]S. T. Roberts, R. E. McAnally, J. N. Mastron, D. H. Webber, M. T. Whited, R. L. Brutchey, M. E. Thompson, and S. E. Bradforth, J. Am. Chem. Soc. **134**(14), 6388–6400 (2012).
[38]M. B. Smith and J. Michl, Ann. Rev. Phys. Chem. **64**, 361–386 (2013).
[39]C. F. Craig, W. R. Duncan, and O. V. Prezhdo, Phys. Rev. Lett. **95**(16), 163001 (2005).
[40]C. P. Hu, H. Hirai, and O. Sugino, J. Chem. Phys. **127**(6), 064103 (2007).
[41]E. Tapavicza, I. Tavernelli, and U. Rothlisberger, Phys. Rev. Lett. **98**(2), 023001 (2007).
[42]X. Zhang, Z. Li, and G. Lu, Phys. Rev. B **84**(23), 235208 (2011).
[43]W. Mou, S. Ohmura, F. Shimojo, and A. Nakano, Appl. Phys. Lett. **100**(20), 203306 (2012).
[44]M. E. Casida, in *Recent Advances in Density Functional Methods (Part I)*, edited by D. P. Chong (World Scientific, Singapore, 1995), pp. 155–192.
[45]J. C. Tully, J. Chem. Phys. **93** (2), 1061–1071 (1990).
[46]J. R. Schmidt, P. V. Parandekar, and J. C. Tully, J. Chem. Phys. **129**(4), 044104 (2008).
[47]O. V. Prezhdo, J. Chem. Phys. **111**(18), 8366–8377 (1999).
[48]A. W. Jasper, S. N. Stechmann, and D. G. Truhlar, J. Chem. Phys. **116**(13), 5424–5431 (2002).
[49]A. B. Bortz, M. H. Kalos, and J. L. Lebowitz, J. Comput. Phys. **17**(1), 10–18 (1975).
[50]D. T. Gillespie, J. Comput. Phys. **22**(4), 403–434 (1976).
[51]K. A. Fichthorn and W. H. Weinberg, J. Chem. Phys. **95**(2), 1090–1096 (1991).
[52]A. F. Voter, in *Radiation Effects in Solids*, edited by K. E. Sickafus, E. A. Kotomin, and B. P. Uberuaga (Springer, Dordrecht, The Netherlands, 2006).
[53]A. Brandt, Math. Comput. **31**(138), 333–390 (1977).
[54]A. Nakano, P. Vashishta, and R. K. Kalia, Comput. Phys. Commun. **83**(2–3), 181–196 (1994).
[55]M. Kunaseth, R. K. Kalia, A. Nakano, K. Nomura, and P. Vashishta, *Proceedings of Supercomputing*, SC13 (ACM/IEEE, 2013).
[56]A. Dreuw, J. L. Weisman, and M. Head-Gordon, J. Chem. Phys. **119**(6), 2943–2946 (2003).
[57]M. Walter, H. Hakkinen, L. Lehtovaara, M. Puska, J. Enkovaara, C. Rostgaard, and J. J. Mortensen, J. Chem. Phys. **128**(24), 244101 (2008).
[58]M. E. Casida and M. Huix-Rotlant, Annu. Rev. Phys. Chem. **63**, 287–323 (2012).
[59]F. Shimojo, S. Ohmura, W. Mou, R. K. Kalia, A. Nakano, and P. Vashishta, Comput. Phys. Commun. **184**(1), 1–8 (2013).
[60]H. M. Jaeger, S. Fischer, and O. V. Prezhdo, J. Chem. Phys. **137**(22), 22A545 (2012).
[61]J. P. Perdew, K. Burke, and M. Ernzerhof, Phys. Rev. Lett. **77**(18), 3865–3868 (1996).
[62]Y. Tawada, T. Tsuneda, S. Yanagisawa, T. Yanai, and K. Hirao, J. Chem. Phys. **120**(18), 8425–8433 (2004).
[63]X. Zhang, Z. Li, and G. Lu, J. Phys.: Condens. Matter **24**(20), 205801 (2012).
[64]G. J. Martyna and M. E. Tuckerman, J. Chem. Phys. **110**(6), 2810–2821 (1999).
[65]J. Heyd, G. E. Scuseria, and M. Ernzerhof, J. Chem. Phys. **118**(18), 8207–8215 (2003).
[66]I. Tavernelli, B. F. E. Curchod, A. Laktionov, and U. Rothlisberger, J. Chem. Phys. **133**(19), 194101 (2010).
[67]F. Shimojo, S. Ohmura, R. K. Kalia, A. Nakano, and P. Vashishta, Phys. Rev. Lett. **104**(12), 126102 (2010).
[68]K. Shimamura, F. Shimojo, R. K. Kalia, A. Nakano, and P. Vashishta, Phys. Rev. Lett. **111**(6), 066103 (2013).
[69]A. S. Torralba, D. R. Bowler, T. Miyazaki, and M. J. Gillan, J. Chem. Theory. Comput. **5**(6), 1499–1505 (2009).
[70]N. Vukmirovic and L. W. Wang, J. Chem. Phys. **134**(9), 094119 (2011).
[71]G. J. O. Beran and S. Hirata, Phys. Chem. Chem. Phys. **14**(21), 7559–7561 (2012).
[72]T. A. Wesolowski and A. Warshel, J. Phys. Chem. **97**(30), 8050–8053 (1993).
[73]N. Govind, Y. A. Wang, and E. A. Carter, J. Chem. Phys. **110**(16), 7677–7688 (1999).
[74]P. Elliott, K. Burke, M. H. Cohen, and A. Wasserman, Phys. Rev. A **82**(2), 024501 (2010).
[75]A. Nakano and S. Ichimaru, Phys. Rev. B **39**(8), 4930–4937 (1989).
[76]M. Frigo and S. G. Johnson, Proc. IEEE **93**(2), 216–231 (2005).
[77]Y. Shiihara, O. Kuwazuru, and N. Yoshikawa, Model Simul. Mater. Sci. Eng. **16**(3), 035004 (2008).
[78]E. Martinez, J. Marian, M. H. Kalos, and J. M. Perlado, J. Comput. Phys. **227**(8), 3804–3823 (2008).
[79]M. Kunaseth, R. K. Kalia, A. Nakano, P. Vashishta, D. F. Richards, and J. N. Glosli, *Proceedings of the International Workshop on Parallel and Distributed Scientific and Engineering Computing*, PDSEC-13 (IEEE, 2013).
[80]P. E. Blochl, Phys. Rev. B **50**(24), 17953–17979 (1994).
[81]G. Kresse and D. Joubert, Phys. Rev. B **59**(3), 1758–1775 (1999).
[82]S. G. Louie, S. Froyen, and M. L. Cohen, Phys. Rev. B **26**(4), 1738–1742 (1982).
[83]Y. Fu, Y. H. Zhou, H. B. Su, F. Y. C. Boey, and H. Agren, J. Phys. Chem. C **114**(9), 3743–3747 (2010).

[84]I. Szlufarska, A. Nakano, and P. Vashishta, Science **309**(5736), 911–914 (2005).
[85]H. Chen, R. K. Kalia, E. Kaxiras, G. Lu, A. Nakano, K. Nomura, A. C. T. van Duin, P. Vashishta, and Z. Yuan, Phys. Rev. Lett. **104**(15), 155502 (2010).
[86]Z. Yuan and A. Nakano, Nano Lett. **13**(10), 4925–4930 (2013).
[87]A. Warshel and M. Karplus, J. Am. Chem. Soc. **94**(16), 5612–5625 (1972).
[88]A. Warshel and M. Levitt, J. Mol. Biol. **103**(2), 227–249 (1976).
[89]S. Ogata, E. Lidorikis, F. Shimojo, A. Nakano, P. Vashishta, and R. K. Kalia, Comput. Phys. Commun. **138**(2), 143–154 (2001).
[90]H. Takemiya, Y. Tanaka, S. Sekiguchi, S. Ogata, R. K. Kalia, A. Nakano, and P. Vashishta, *Proceedings of Supercomputing*, SC06 (IEEE/ACM, 2006).
[91]S. Dapprich, I. Komáromi, K. S. Byun, K. Morokuma, and M. J. Frisch, J. Mol. Struct. **461–462**, 1–21 (1999).
[92]A. C. T. van Duin, S. Dasgupta, F. Lorant, and W. A. Goddard, J. Phys. Chem. A **105**(41), 9396–9409 (2001).
[93]K. Nomura, R. K. Kalia, A. Nakano, P. Vashishta, A. C. T. van Duin, and W. A. Goddard, Phys. Rev. Lett. **99**(14), 148303 (2007).
[94]M. Vedadi, A. Choubey, K. Nomura, R. K. Kalia, A. Nakano, P. Vashishta, and A. C. T. van Duin, Phys. Rev. Lett. **105**(1), 014503 (2010).
[95]S. B. Sinnott and D. W. Brenner, MRS Bull. **37**(5), 469–473 (2012).
[96]A. Shekhar, K. Nomura, R. K. Kalia, A. Nakano, and P. Vashishta, Phys. Rev. Lett. **111**(18), 184503 (2013).
[97]C. E. Wilmer, M. Leaf, C. Y. Lee, O. K. Farha, B. G. Hauser, J. T. Hupp, and R. Q. Snurr, Nat. Chem. **4**(2), 83–89 (2012).
[98]A. M. Virshup, J. Contreras-Garcia, P. Wipf, W. T. Yang, and D. N. Beratan, J. Am. Chem. Soc. **135**(19), 7296–7303 (2013).