

Quantum Dynamics at Scale: Ultrafast Control of Emergent Functional Materials

Subodh C Tiwari
Collaboratory for Advanced
Computing and Simulations
University of Southern California
Los Angeles, CA, U.S.
sctiwari@usc.edu

Putt Sakdhnagool
NSTDA Supercomputer Center
National Science and Technology
Development Agency
Pathum Thani, Thailand
putt.sakdhnagool@nectec.or.th

Shogo Fukushima
Department of Physics
Kumamoto University
Kumamoto, Japan
189d8065@st.kumamoto-u.jp

Rajiv K. Kalia
Collaboratory for Advanced
computing and Simulations
University of Southern California
Los Angeles, CA, U.S.
rkalia@usc.edu

Aravind Krishnamoorthy
Collaboratory for Advanced
Computing and Simulations
University of Southern California
Los Angeles, CA, U.S.
kris658@usc.edu

Manaschai Kunaseth
NSTDA Supercomputer Center
National Science and Technology
Development Agency
Pathum Thani, Thailand
manaschai.kun@nectec.or.th

Aiichiro Nakano
Collaboratory for Advanced
Computing and Simulations
University of Southern California
Los Angeles, CA, U.S.
anakano@usc.edu

Ken-ichi Nomura
Collaboratory for Advanced
Computing and Simulations
University of Southern California
Los Angeles, CA, U.S.
knomura@usc.edu

Pankaj Rajak
Argonne Leadership Computing
Facility
Argonne National Laboratory
Lemont, IL, U.S.
prajak@anl.gov

Fuyuki Shimojo
Department of Physics
Kumamoto University
Kumamoto, Japan
shimojo@kumamoto-u.ac.jp

Ye Luo
Argonne Leadership Computing
Facility
Argonne National Laboratory
Lemont, IL, U.S.
yeluo@anl.gov

Priya Vashishta
Collaboratory for Advanced
Computing and Simulations
University of Southern California
Los Angeles, CA, U.S.
priyav@usc.edu

ABSTRACT

Confluence of extreme-scale quantum dynamics simulations (i.e. quantum@scale) and cutting-edge x-ray free-electron laser experiments are revolutionizing materials science. An archetypal example is the exciting concept of using picosecond light pulses to control emergent material properties on demand in atomically-thin

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2020 Association for Computing Machinery. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

HPCAsia2020, January 15–17, 2020, Fukuoka, Japan
© 2020 Association for Computing Machinery.
ACM ISBN 978-1-4503-7236-7/20/01...\$15.00
<https://doi.org/10.1145/3368474.3368489>

layered materials. This paper describes efforts to scale our quantum molecular dynamics engine toward the United States’ first exaflop/s computer, under an Aurora Early Science Program project named “Metascalable layered material genome”. Key algorithmic and computing techniques incorporated are: (1) globally-scalable and locally-fast solvers within a linear-scaling divide-conquer-recombine algorithmic framework; (2) algebraic ‘BLASification’ of computational kernels; and (3) data alignment and loop restructuring, along with register and cache blocking, for enhanced vectorization and efficient memory access. The resulting weak-scaling parallel efficiency was 0.93 on 131,072 Intel Xeon Phi cores for a 56.6 million atom (or 169 million valence-electron) system, whereas the various code transformations achieved 5-fold speedup. The optimized simulation engine allowed us for the first time to establish a significant effect of substrate on the dynamics of layered material upon electronic excitation.

CCS CONCEPTS

- Theory of computation → Massively parallel algorithms
- Applied computing → Physics

KEYWORDS

Quantum molecular dynamics simulation, Divide-conquer-recombine algorithm, Globally-scalable and locally-fast solver, Algebraic BLASification, Register and cache blocking, Vectorization, Photo-excitation, Layered material, Substrate effect

1 Introduction

Over the past 30 years, quantum molecular dynamics (QMD) has been developed into a powerful simulation methodology, with which we can now study complex material dynamics incorporating their intrinsically quantum nature [1-3]. For example, recent advancements in ultrafast materials science have raised the exciting possibility of using picosecond light pulses to switch emergent material properties on demand in various functional materials [4]. An archetype is optically-induced semiconductor-to-metal phase transition in atomically-thin layered materials (LMs) [5], where nanoscale metal wires are lithographed using a laser beam to fabricate electronic circuits [6]. Here, state-of-the-art x-ray free-electron laser (XFEL) [7, 8] and ultrafast electron diffraction (UED) [9, 10] experiments provide unprecedented capability to observe such light-induced material dynamics with sub-nanometer spatial and sub-picosecond time resolutions. Synergy between leadership-scale QMD simulations [11] and ultrafast XFEL/UED experiments provides fundamental understanding of quantum dynamics, thereby revolutionizing ultrafast control of quantum-material properties. For example, recent QMD simulations not only reproduced surprising near-100% quantum efficiency for picosecond (ps) conversion of photo-induced electronic-excitation energy to crystalline lattice motions observed in XFEL and UED experiments, but also determined their structural transition pathways and electronic origins, e.g., electronic excitation-induced Fermi-surface nesting and phonon softening [8, 9, 12].

In order to perform extreme-scale QMD simulations (*i.e.* quantum@scale) on the United States’ first exaflop/s computer, Aurora A21, in tandem with the next-generation XFEL experiments at the forthcoming LCLS-II (Linac Coherent Light Source) facility at SLAC National Accelerator Laboratory, the Aurora Early Science Program (ESP) project, “Metascalable layered material genome” [13], is porting our QMD simulation engine to A21. The key computational challenge here is to simultaneously achieve high scalability on a large number of computing nodes with high floating-point (FP) performance on massive cores within each node. We address this challenge through a concept of global-local separation. Based on a divide-conquer-recombine (DCR) algorithmic framework [14], which is metascalable (or “design once, scale on new architectures”) with minimal architectural assumptions [15], our globally-scalable and locally-fast (GSLF) electronic-structure solvers hybridize highly-scalable global real-space multigrids with local plane-wave bases for fast Fourier transforms (FFT). This global-local separation

allows us to focus on local performance optimization within each node, while leaving internode message passing intact, when porting the QMD engine to a new computer architecture like A21.

This paper describes computational techniques and preliminary performance results in our Aurora ESP project. In preparation for the Intel A21 architecture, we have transformed computations using (1) algebraic ‘BLASification’ of computational kernels and (2) data alignment and loop restructuring, along with register and cache blocking, for enhanced vectorization and efficient memory access. Preliminary benchmark tests were performed on an Intel Xeon Phi Knights Landing (KNL) platform, with special focus on single-instruction multiple-data (SIMD) data parallelism.

2 Simulation Method

QMD follows the trajectories of all atoms while computing interatomic forces quantum mechanically from first principles in the framework of density functional theory (DFT) [16, 17]. By solving N one-electron problems self-consistently instead of directly solving one N -electron problem, DFT reduces the exponential complexity of the quantum N -body problem to $O(N^3)$. QMD simulations are typically limited to small systems involving a few hundred atoms due to its $O(N^3)$ complexity.

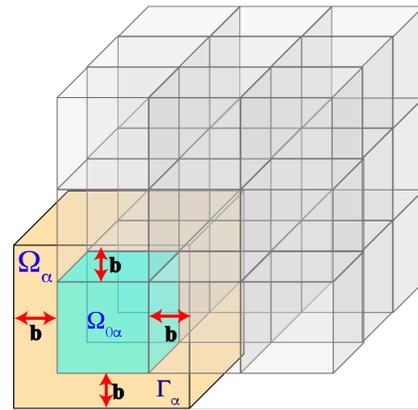


Figure 1: 3D schematic of divide-and-conquer (DC) spatial decomposition. The simulation cell Ω is subdivided into non-overlapping cores Ω_{α} (colored in cyan). Buffer region (colored yellow), Γ_{α} , of thickness b is added on all six sides of each core. A spatially-localized domain, $\Omega_{\alpha} = \Omega_{\alpha} \cup \Gamma_{\alpha}$, is union of core and buffer region.

To overcome the $O(N^3)$ bottleneck, various $O(N)$ DFT algorithms have been designed [18, 19] on the basis of the data-locality principle called quantum nearsightedness [20, 21]. Among them, the divide-and-conquer (DC) DFT algorithm [22-24] pioneered by Weitao Yang [22] is highly scalable [25] on massively parallel computers. DC-DFT represents the three-dimensional space Ω as a union of spatially-localized domains, $\Omega = \cup_{\alpha} \Omega_{\alpha}$ (Ω_{α} is the α -th domain), and global physical properties are computed as linear combinations of local domain properties. While the electronic Kohn-Sham (KS) wave functions $\{\psi_n(\vec{r})\}$ and charge

density $\rho(\vec{r}) = \sum_n |\psi_n(\vec{r})|^2$ are obtained in a self-consistent-field (SCF) iteration in conventional $O(N^3)$ DFT [26], in $O(N)$ DC-DFT, local KS wave functions $\{\psi_n^\alpha(\vec{r})\}$ and global charge density

$$\rho(\vec{r}) = \sum_\alpha \sum_n p_\alpha(\vec{r}) |\psi_n^\alpha(\vec{r})|^2 \quad (1)$$

are determined instead in a “global-local SCF” iteration [15]. Here, $p_\alpha(\vec{r})$ is a partition of unity, $\sum_\alpha p_\alpha(\vec{r}) = 1$ at every spatial position \vec{r} , which is compactly supported (i.e., $p_\alpha(\vec{r}) = 0$ if $\vec{r} \notin \Omega_\alpha$). In Eq. (1), $\psi_n^\alpha(\vec{r})$ is the n -th KS wave function in the α -th domain (Fig. 1).

Numerically, local KS wave functions are obtained by solving an eigenvalue problem with the preconditioned conjugate-gradient (PCG) method concurrently among all domains. Here, the Hamiltonian operator acting on each $\psi_n^\alpha(\vec{r})$ contains the electrostatic Hartree field $V_H(\vec{r})$. This global field in turn is obtained from the global electron density by solving Poisson’s equation:

$$\nabla^2 V_H(\vec{r}) = -4\pi\rho(\vec{r}) \quad (2)$$

The local KS wave functions $\psi_n^\alpha(\vec{r})$, global density $\rho(\vec{r})$ and Hartree field $V_H(\vec{r})$, along with the global chemical potential μ , are thus obtained in a fixed-point iteration in the global-local SCF iteration. Workflow of DC-DFT is shown in Fig. 2.

3 Computational Techniques

In this section, we describe a series of algorithmic and computational techniques that make our QMD engine scalable and achieve high FP performance.

3.1 Divide Conquer Recombine

We have proposed an extension of divide-and-conquer (DC) algorithm called divide-conquer-recombine (DCR) [11, 14, 15]. The DCR algorithm consists of two phases: DC and recombine. In the DC phase, the physical space is subdivided into spatially-localized domains as in Fig. 1. These domains are embedded in a global field, which is computed efficiently using tree-based algorithms such as real-space multigrids. The DC phase constructs globally-informed local solutions, which, in the second recombine phase, are used as compactly-supported bases to synthesize a global solution encompassing large spatiotemporal scales. The recombine phase typically involves range-limited n -tuple computations to account for higher interdomain correlations that are not captured by the tree network topology used in the DC phase [27, 28].

Lean divide-and-conquer (LDC): For QMD, we have developed a LDC-DFT algorithm, which significantly reduces the prefactor of the $O(N)$ computational cost based on complexity and error analyses. Specifically, LDC-DFT applies an adaptive boundary condition, which is proportional to the difference between local ($\rho_\alpha(\vec{r}) = \sum_n |\psi_n^\alpha(\vec{r})|^2$) and global ($\rho(\vec{r})$) electron densities, so as to minimize the contamination of local solutions at domain peripheries $\partial\Omega_\alpha$ [14].

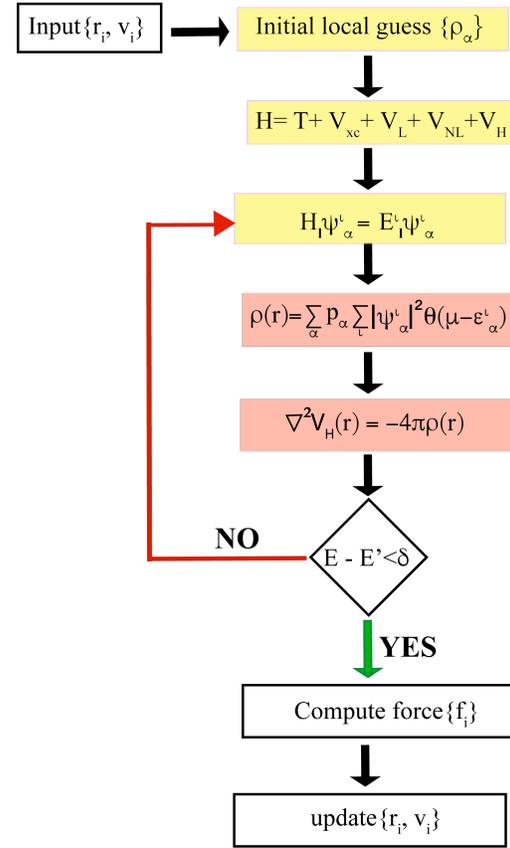


Figure 2: Workflow of DC-DFT in $O(N)$ QMD simulation to obtain Kohn-Sham eigenstates. Upon input of initial atomic positions and atomic velocities, first steps include making an initial guess for local charge density $\rho_\alpha(r)$, and computation of the local Hamiltonian H_i for each domain, including kinetic energy T , exchange-correlation potential V_{xc} , local potential V_L and nonlocal pseudopotential V_{NL} . Next, global charge density $\rho(r)$ is computed and Hartree potential V_H is obtained using multigrid method. SCF iteration is based on multiple convergence criteria. If criteria are met, atomic forces are computed using Hellmann-Feynman theorem and atomic positions and velocities are updated based on velocity Verlet algorithm. If criteria are not met, new SCF iteration is performed using a new charge density $\rho(r)$ created by mixing old charge density and new guess.

Globally scalable and locally fast (GSLF) solver: Our GSLF solver combines: (1) fast intradomain computation—FFT-based spectral method is used to obtain local KS wave functions $\psi_n^\alpha(\vec{r})$ with high numerical efficiency [29]; and (2) scalable interdomain computation—octree-based multigrids [30] are employed to solve Eq. (2) and obtain the global Hartree field $V_H(\vec{r})$ from the global electron density $\rho(\vec{r})$ in a highly scalable manner on massively parallel computers. Use of plane-wave bases in spatially-localized domains is akin to the discontinuous Galerkin DFT method [31-

33]. Here, it is further combined hierarchically with the multigrid method.

Hybrid band-space-domain (BSD) decomposition: On parallel computers, we utilize hybrid parallelization. We first decompose the computation into those of different domains. For higher parallelism, we further employ band decomposition (i.e., assigning the calculations of different KS wave functions to different processors) and spatial decomposition (i.e., distributing real-space or reciprocal-space grid points among processors) within each domain.

3.2 Algebraic Computational Transformation

The numerical kernel of QMD is the solution of local KS equations (Fig. 2), i.e., iterative minimization of the energy as a functional of $\psi_n^\alpha(\vec{r})$ using the PCG method [2]. Our original QMD code used band-by-band PCG minimization, in which KS wave functions (or bands) are optimized one at a time sequentially in ascending order of KS energy. To take maximal advantage of the highly-optimized basic linear algebra subprograms (BLAS) library, we instead adopt an all-band approach, in which PCG minimization is performed simultaneously for all KS wave functions at once. The original band-by-band computation is expressed in terms of matrix-vector operations using the DGEMV subroutine in level 2 BLAS (BLAS2). By combining all bands together into a matrix and operating the KS Hamiltonian at once, we transform the entire PCG minimization to BLAS3 matrix-matrix computations using the DGEMM subroutine. This exposes more parallelism and drastically increases the FP performance. We have formulated these computational transformations algebraically at the abstract level of governing equations. Consequently, key computational kernels such as nonlocal pseudopotential calculation are now performed solely as matrix-matrix operations using the DGEMM subroutine in BLAS3. Here, KS wave functions are orthonormalized using parallel Cholesky decomposition of the overlap matrix. While all-band PCG computation may have numerical-stability issues in very large systems, our approach applies it only in domain-by-domain fashion. Consequently, we have not encountered any stability issue.

3.3 Performance Profiling and Code-Level Optimization

To ensure optimal performance on target architecture platform, we performed performance profiling and code-level optimization of the QMD engine, called QXMD, for Intel-based architecture. To identify the most promising candidate subroutine for performance optimization, we first performed a hardware-mode hotspot profiling using Intel VTune Amplifier 2019. The profiling was performed on a compute node with dual 20-core Intel Xeon Scalable Gold 6,148 processors at NSTD Supercomputer Center (ThaiSC). For the test configuration, we performed a SCF iterations on a 12-atom molybdenum ditelluride (MoTe₂) system using 40 MPI tasks. The profiled QXMD code was compiled with “-lm -mkl=sequential -O3 -axCORE-AVX512 -qopt-report=5 -qopt-

report-phase=vec” flags, whereas the non-uniform memory access (NUMA) allocation and CPU binding were also enforced during runtime. Note that the -qopt-report flag provides additional automatic vectorization information from compiler for every source files, which were subsequently used to analyze the degree and method of vectorization carried out by the compiler.

Table 1: Hotspot profiling results.

Subroutine	CPU time (%)	% of FP Ops	% Vectorized FP Ops
<i>mbyv_p</i>	18.8%	34.8%	0.9%
<i>precmp</i>	9.7%	7.9%	7.2%
<i>cgrad</i>	5.7%	54.8%	100.0%
<i>hcvand_g</i>	5.0%	16.0%	99.3%
<i>rfft3</i>	2.8%	1.8%	100.0%

Table 1 shows the FP profile for the most time-consuming subroutines obtained from hotspot profiling result. Here, *mbyv_p* subroutine was the most time consuming (18.8% CPU time), whereas 99.1% of its FP operations were not vectorized. The other subroutines in Table 1 either contained small fraction of FP operations, or were already vectorized. Therefore, we focused our optimization on *mbyv_p* subroutine.

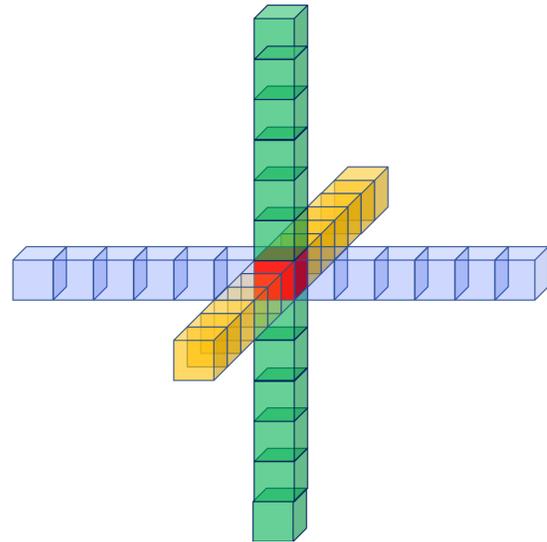


Figure 3: 6-th order, 37-point stencil, whose footprint is 13^3 on a three-dimensional mesh. Blue, yellow and green colors show the x, y and z locations of memory in 3D space, respectively. Red color shows the target cell.

The main function of subroutine *mbyv_p* is to calculate the Laplacian (i.e. second derivative) of a scalar variable on a three-dimensional mesh in the finite-difference (FD) approximation, which is mainly used to solve Poisson equation to obtain Hartree field in the entire space. The operations in *mbyv_p* can be expressed as the multiplication between a sparse matrix A and a vector x . Here, the QMD code utilizes high-order stencil computation (Fig. 3) to solve the FD approximation. To better understand the computing

characteristics of this subroutine, we further performed an in-depth performance profiling using HPC performance characterization mode. The result showed that 48.6% of pipeline executions within the main stencil computation kernel in *mbyv_p* was stalled, of which 19.9% specifically originated from L1 miss/L2 hit stalls.

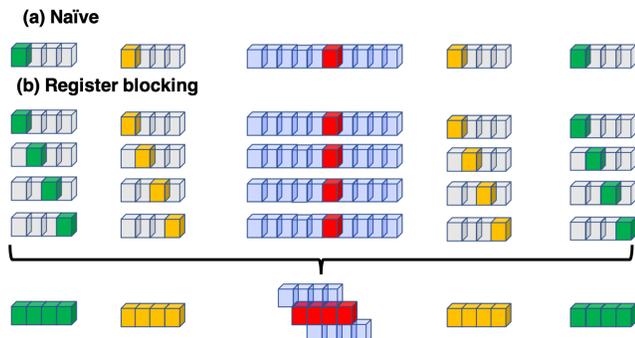


Figure 4: Memory access pattern of register blocking compared to the naïve access pattern of stencil computation. In the naïve stencil kernel (a), red block signifies target grid point, while blue, yellow and green blocks show locations in memory space of the stencil in x , y and z directions, respectively. On the other hand, CRB algorithm (b) targets on a chunk of grid points in x -stride rather than a single grid point as in naïve algorithm, thereby optimally utilizing register and data cache.

To improve the performance of the stencil-computation kernel in *mbyv_p*, we employ a cache and register-blocking (CRB) algorithm [34], which improves the high-order stencil performance by reducing the computation and memory access of the high-order strides (e.g., y -stride and z -stride). The CRB algorithm packs y - and z -strides data into aligned $nBlock$ chunk and processes them at once (see Fig. 4), thereby improving data reuse within the cacheline in y - and z -strides, increasing the vectorization and reducing redundant micro-operations.

4 Benchmark Tests

Numerical tests are performed on Intel Xeon Phi Knights Landing (KNL) based supercomputer [35] named Theta at Argonne Leadership Computing Facility (ALCF). Theta consists of 24 racks totaling 4,392 Intel Xeon phi nodes. Each node has 64 physical cores. Each core operates with 32 KB L1 instruction cache, 32 KB L1 data cache with two independent floating points operation units capable of executing 512-bit wide SIMD instruction set. Each tile consists of two cores and a shared 1MB L2 cache. Tiles are connected by 2D mesh networks-on-chips. Each node also supports 16 GB high bandwidth in-package multichannel DRAM memory (MCDRAM) and 192 GB DDR4 memory. There are three memory mode supported on Intel Xeon Phi. One of the cache mode, flat mode and hybrid mode is selected at booting. In cache mode, MCDRAM is used as a cache to DRAM; while in flat mode, both MCDRAM and DRAM operates as a single linear memory space. Hybrid mode allows us to choose part of MCDRAM as cache and

the remaining memory space as linear memory space. Following benchmarks are performed in cache mode [36].

4.1 Weak and Strong Scaling on Theta

Weak scaling on Theta: We performed a weak-scaling benchmark of the QMD code on Theta, in which the number of atoms per core N/P is kept constant. We measured the wall-clock time per QMD simulation step with scaled workload—216 P -atom MoTe₂ system on P cores of Theta. In this test, we set the number of domains to be P . The execution time includes 3 SCF iterations to determine the electronic wave functions and the global density, with 3 PCG 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 QMD code 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 of the speed on P cores to that on 512 cores as a reference system. The weak-scaling parallel efficiency is the isogranular speedup divided by $P/512$ (Fig. 5). With the granularity of 216 atoms per core, the parallel efficiency remains nearly unity, as shown in Fig. 5, within measurement fluctuation up to $P = 261,144$ for a 56,623,104-atom MoTe₂ system. This result demonstrates a high scalability of the QXMD code.

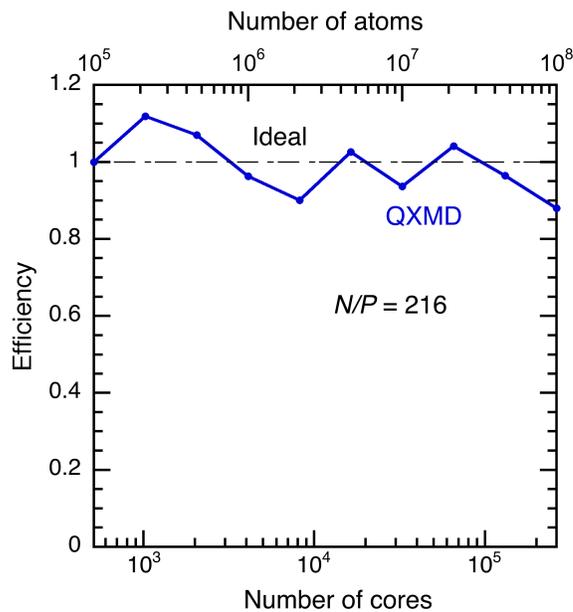


Figure 5: Weak-scaling parallel efficiency of the parallel QXMD code based on LDC-DFT algorithm, with scaled workloads — 216 P -atom MoTe₂ system on P Intel KNL cores ($P = 512, \dots, 262,144$) of Theta.

Figure 5 exhibits an network-interference pattern observed in earlier study [37]. Theta employs dragonfly network based on the Cray Aries technology [38]. Aries is a system-on-chip device that consists of four network interface cards (NICs) via 16x PCI-Express Gen3 interface connecting the KNL nodes and a 48-port

router tile with 4.7-5.25 GB/s optical and electrical bidirectional links. A dragonfly network topology can be seen as a two-level all-to-all mesh network between computing node groups. The cost-effective electrical link is used to connect the NICs to other local routers in a cabinet and the optical links serve as global interconnect across cabinets. Theta system consists of 20 cabinets in total, 192 KNL nodes via 48 Aries chips per cabinet realizing a high-bandwidth and low-diameter network for a wide range of applications [39]. A number of benchmarks on dragonfly network have shown that the communication performance is highly dependent on job-placement strategy, data routing policy and the intensity of data transfer of each application, resulting in variations in their performance [37]. The degree of shared links among concurrent jobs is an important factor to avoid performance degradation caused by interference between multiple jobs.

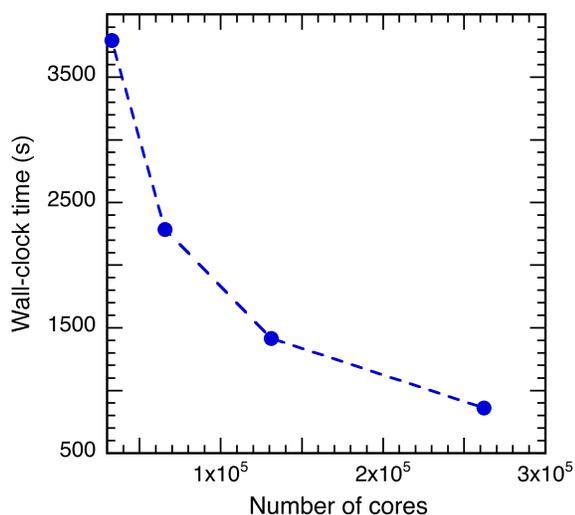


Figure 6: Wall-clock time of QMD with a fixed workload — 7,077,888-atom MoTe₂ system on P cores ($P = 512, \dots, 262,144$) of Intel KNL cores.

Strong scaling on Theta: Next, we performed a strong-scaling test by simulating a MoTe₂ system containing a total of 7,077,888 atoms. In this test, the number of cores ranges between $P = 32,768$ and 262,144, while keeping the total problem size constant. Figure 6 shows the wall-clock time per QMD simulation step as a function of P . The time-to-solution is reduced by a factor of 4.4 on 262,144 cores compared with the 32,768-core run (i.e., using 8-times larger number of cores). The corresponding strong-scaling parallel efficiency is 55%. It is more difficult to achieve high strong-scaling parallel efficiency compared with weak-scaling parallel efficiency. This is due to the surge of communication/computation ratio as the workload per rank shrinks proportionally. This is partly understood by analyzing the parallel efficiency η as a function of the number of cores P and that of atoms N . For the weak-scaling parallel efficiency with constant granularity ($n = N/P$), $\eta = 1/[1 + \alpha n^{-\frac{1}{3}} + \beta n^{-1} \log P]$, exhibiting a very weak logarithmic dependence on P . For the strong-scaling parallel efficiency with

constant N , in contrast, $\eta = 1/[1 + \alpha (\frac{P}{N})^{-\frac{1}{3}} + \beta N^{-1} P \log P]$, which exhibits much stronger power dependence on P . With 8-times smaller system size of the weak-scaling test, the observed strong-scaling efficiency is considered decent.

4.2 Performance Tests

To test the effects of various performance-optimization techniques described in Sec. 3, we measured the cumulative performance improvement of our optimized code (named QXMD-G-CRB) with respect to the original code. The notation of our codes is as follows. Nonlocal-pseudopotential calculation performed in momentum space is labeled QXMD-G, whereas calculation in real space is labeled QXMD-R. Algebraic transformation with inclusion of BLAS, and calculation of non-local potential in momentum space is labeled QXMD-G-BL, and the code implementing calculation of non-local potential in momentum space, algebraic transformation and register/cache blocking is labeled QXMD-G-CRB. The test was performed on Intel Xeon Phi cores of Theta with MoTe₂ system containing 192 atoms. Figure 7 shows that nonlocal pseudopotential calculation performed in real space (QXMD-R) is faster than that in reciprocal space (QXMD-G). However, calculation in real space resulted in small errors using equivalent computational parameters. From our optimization, we have achieved a cumulative 5x speedup with respect to QXMD-G and 2.5x speed with respect to QXMD-R, while avoiding the error associated with the real-space formalism. Furthermore, with increased number of cores, speedup obtained remained almost constant, suggesting a scalable speedup.

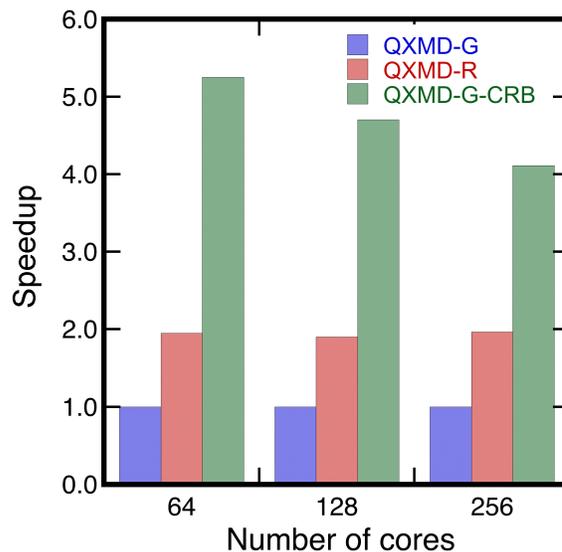


Figure 7: Performance comparison of the QMD code before optimization (QXMD-G, QXMD-R) and after optimization with nonlocal pseudopotential calculation in momentum space (QXMD-G-CRB) for the number of cores ranging from 64 to 256.

To further delineate the effect of algebraic transformation from that of machine-level optimization, we compared the performance of QXMD-G-BL and QXMD-G-CRB codes on Intel Xeon Phi processors and dual 20-core Intel Xeon Scalable Gold 6,148 processors with 12-atom MoTe₂ system.

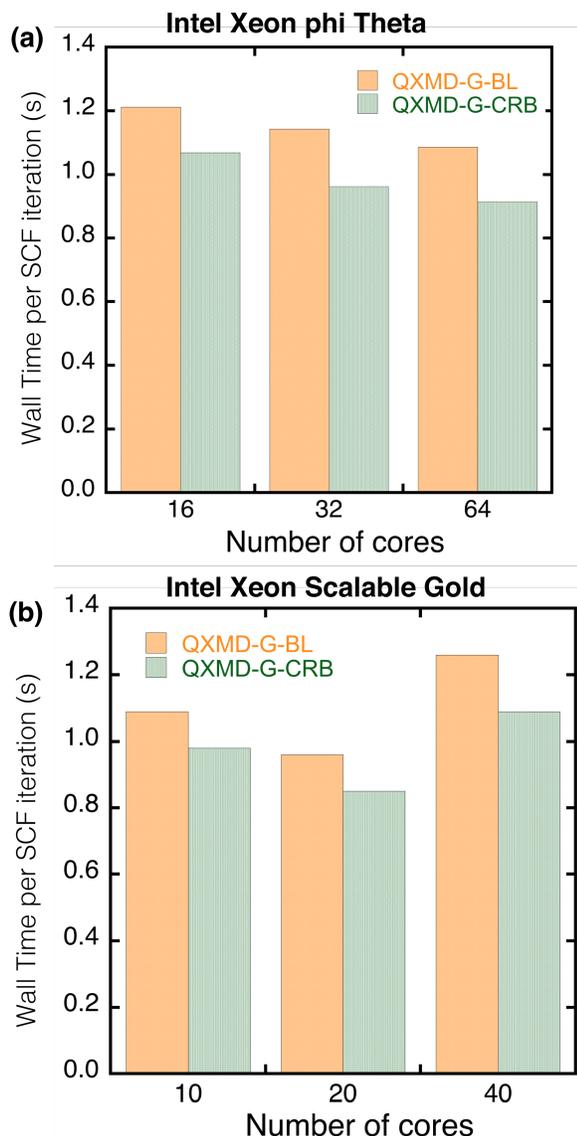


Figure 8: Performance of QMD codes before (orange) and after (green) machine-level performance optimization on Intel Xeon Phi (a) and Intel Xeon Scalable Gold 6,148 processors (b).

Figure 8 compares runtime of the QMD code before (QXMD-G-BL) and after CRB optimization (QXMD-G-CRB) with varying numbers of cores. In addition to Theta, we have also performed the same test on another platform to assess the portability of our

performance-optimization techniques. Figure 8a shows that we have gained 13-18 % improvement with CRB optimization on Intel Xeon Phi KNL cores, consistently for $P=16-64$ cores. For example, runtime per SCF iteration decreased from 1.14 s to 0.96 s on 32 cores. Similarly, Fig. 8b shows runtime of the code before and after CRB optimization on Intel Xeon Scalable gold 6,148 processors. Runtime before CRB optimization [15] is 1.09, 0.96 and 1.26 s, respectively, on 10, 20 and 40 cores. On the other hand, runtime after CRB optimization decreased to 0.98, 0.85 and 1.09 s on 10, 20 and 40 cores, respectively, which amount to 11-16% speedup. In terms of the number of cores used, the shortest runtime was observed on 20 cores. This is likely due to the MPI synchronization on the dual 20-core processors.

To identify the source of performance improvement, we compared the profiling results for memory and computing characteristics of the QMD code before and after optimization. First, we found that the character of FP operations in *mbyv_p* was improved from no vectorization to fully AVX(256)/FMA(256) vectorized after optimization. As a result, the computing time spent in *mbyv_p* subroutine was reduced by 39.0%, and the number of memory-load instructions was significantly reduced by 65.0%. These improvements are directly associated with the improvement of computing pattern and memory alignment of CRB scheme within the stencil kernel. In terms of memory access, we found that cache hit/miss profiles are similar before and after optimization. We suspect that this is due to the relatively small mesh size such that it could not benefit adequately from improved memory layout.

5 APPLICATION

Our scalable, performance optimized QMD code provides new capabilities to study nontrivial excitation dynamics in material systems of experimentally relevant length scales. To demonstrate this new functionality, we model the excited-state dynamics of monolayer WSe₂ crystals supported on alumina (Al₂O₃) substrates. WSe₂ is an atomically thin layered material belonging to the class of transition-metal dichalcogenides (TMDC), which are being actively investigated for their high carrier mobility and valley- and spin-dependent optical and electronic properties not observed in their 3D counterparts. This makes them suitable for tunable semiconducting to metallic structural phase change materials and synthesis of new materials with remarkable functionality. Photoresponse of these material systems is an ongoing topic of interest owing to its potential in enabling optically-induced material functionalization and structural transformations [40]. Non-trivial photo-induced dynamics of a WSe₂ monolayer on alumina was recently demonstrated in XFEL experiments [8]. However, the lack of computational resources has constrained QMD simulations to the study of only suspended crystals (*i.e.* no substrate) and therefore the effect of substrate on photodynamics remains to be understood.

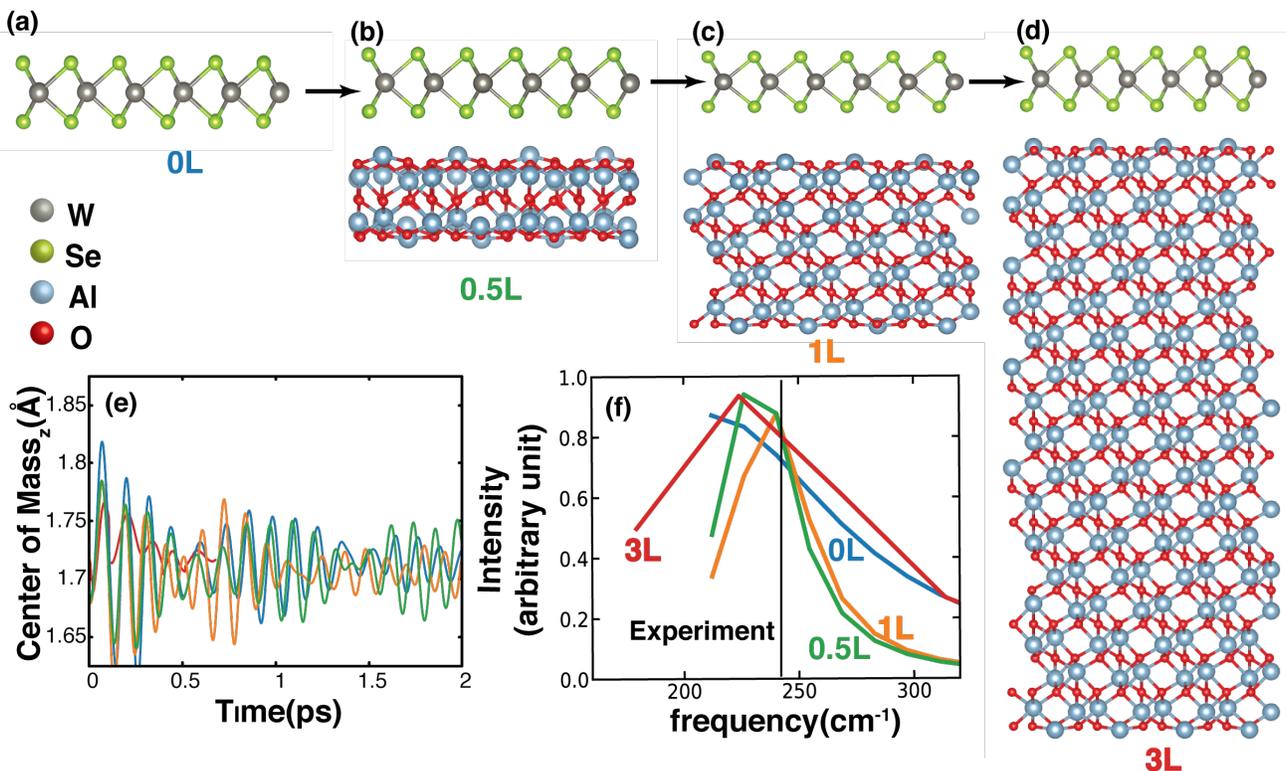


Figure 9: (a-d) Schematics of four atomic systems (0L, 0.5L, 1L, 3L) used to study the influence of dielectric substrate thicknesses on the photodynamics of the supported WSe₂ monolayer. (e) The measured center of mass of the lower plane of Se atoms is independent of the thickness of the substrate. (f) Discrete Fourier transform of the atomic displacements of the Se atoms in the in-plane (i.e. x - y) directions shows that the substrate does not modify the frequency of the dominant mode of lattice vibrations.

We use our DCR-QMD program to understand excited-state dynamics of WSe₂ on alumina (001) surface. In order to systematically study the substrate effect, we consider 4 distinct systems (Fig 9, a-d) – an unsupported WSe₂ monolayer with no substrate (a), WSe₂ supported on an Al₂O₃ (001) surface that is 0.5 unit cells thick (b), WSe₂ supported on 1-unitcell-thick Al₂O₃ (c) and WSe₂ supported on a 3-unitcell-thick Al₂O₃ substrate (d) that approximates the behavior of an infinitely thick substrate. Systems are labeled based on the thickness of substrate, such as system with no substrate is labeled 0L and systems with 0.5, 1.5 and 3 unit cells of the Al₂O₃ substrate are respectively labeled 0.5L, 1L and 3L. The largest system considered in this study (3L) contains 1,014 atoms. All the systems were first relaxed to the ground state. After relaxation, we perform excited-state dynamics by promoting electrons to the unoccupied orbitals by manually modifying the occupancy of the Kohn-Sham eigenlevels according to a Fermi distribution at high temperatures. We studied excitation dynamics up to 2 ps.

The inclusion of the Al₂O₃ substrate breaks the symmetry between the two Se atomic planes resulting in two different type of Se atoms. The lower plane of Se atoms closer to the substrate are denoted as Se_L, whereas the upper plane of Se atoms further away

from the surface are denoted as Se_U. There is no anisotropy in the atomic dynamics of the Se_U and Se_L planes, indicating a weak influence of the substrate on the atomic dynamics. While the magnitude of atomic displacements decreases between 0L to 3L systems, there is no measurable variation in the average center of mass of the Se_L plane as a function of time in the presence of the substrate (Fig. 9e). Further, the dominant vibrational frequency for the Se_L atoms is unaffected by the presence of the substrate. We also performed discrete Fourier transformation of mean square displacement in x - y plane as shown in Fig. 9f. Once again, the weak influence of the substrate is manifested in the lack of modulation in the peak at 240 cm⁻¹ (± 13 cm⁻¹) corresponding to active E_g^2 Raman mode.

These results, along with the minimal charge transfer between alumina substrate and WSe₂ materials, suggest the relatively weak influence of the substrate on photodynamics in TMDC materials. While dielectric substrates have been found to drastically modify electronic properties of TMDC LMs like band gap [41], effect on excitation dynamics involving both electrons and atoms are minor on substrate [8].

6 CONCLUSION

We have developed a series of algorithmic and computational techniques to speed up our linear-scaling quantum molecular dynamics simulation code based on divide-and-conquer density functional theory. The improved code on an Intel Xeon Phi KNL-based supercomputer allowed us for the first time to reveal essential effects of dielectric substrate on photo-induced dynamics involving electrons and atoms. Quantitative understanding enabled by such quantum@scale will be catapulted to the next level on the exaflop/s A21, likely filling fundamental knowledge gap in more complex functional materials such as magic-angle twisted LMs [42] and emergent phases such as polar Skyrmions [43].

ACKNOWLEDGEMENT

This research was supported as part of the Computational Materials Sciences Program funded by the U.S. Department of Energy (DOE), Office of Science, Basic Energy Sciences, under Award Number DE-SC0014607. An award of computer time was provided by the Aurora Early Science Program. This research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. Computational resources were partially provided by NSTDA Supercomputer Center (ThaiSC).

REFERENCES

[1] R. Car and M. Parrinello. 1985. Unified approach for molecular-dynamics and density-functional theory. *Physical Review Letters*, 55 (22), 2471-2474 <https://doi.org/10.1103/PhysRevLett.55.2471>

[2] M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias and J. D. Joannopoulos. 1992. Iterative minimization techniques for ab initio total-energy calculations - molecular-dynamics and conjugate gradients. *Reviews of Modern Physics*, 64 (4), 1045-1097 <https://doi.org/10.1103/RevModPhys.64.1045>

[3] F. Shimojo, S. Fukushima, H. Kumazoe, M. Misawa, S. Ohmura, P. Rajak, K. Shimamura, L. Bassman, S. C. Tiwari, R. K. Kalia, A. Nakano and P. Vashishta. 2019. QXMD: an open-source program for nonadiabatic quantum molecular dynamics. *SoftwareX*, 10, 100307 <https://doi.org/10.1016/j.softx.2019.100307>

[4] D. N. Basov, R. D. Averitt and D. Hsieh. 2017. Towards properties on demand in quantum materials. *Nature Materials*, 16 (11), 1077-1088 <https://doi.org/10.1038/NMAT5017>

[5] A. K. Geim and I. V. Grigorieva. 2013. Van der Waals heterostructures. *Nature*, 499 (7459), 419-425 <https://doi.org/10.1038/nature12385>

[6] A. Krishnamoorthy, L. Bassman, R. K. Kalia, A. Nakano, F. Shimojo and P. Vashishta. 2018. Semiconductor-metal structural phase transformation in MoTe2 monolayers by electronic excitation. *Nanoscale*, 10 (6), 2742-2747 <https://doi.org/10.1039/c7nr07890k>

[7] C. Bostedt, S. Boutet, D. M. Fritz, Z. R. Huang, H. J. Lee, H. T. Lemke, A. Robert, W. F. Schlotter, J. J. Turner and G. J. Williams. 2016. Linac Coherent Light Source: the first five years. *Reviews of Modern Physics*, 88 (1), 015007 <https://doi.org/10.1103/RevModPhys.88.015007>

[8] I. Tung, A. Krishnamoorthy, S. Sadasivam, H. Zhou, Q. Zhang, K. L. Seyler, G. Clark, E. M. Mannebach, C. Nyby, F. Ernst, D. Zhu, J. M. Glownia, M. E. Kozina, S. Song, S. Nelson, H. Kumazoe, F. Shimojo, R. K. Kalia, P. Vashishta, P. Darancet, T. F. Heinz, A. Nakano, X. Xu, A. M. Lindenberg and H. Wen. 2019. Anisotropic structural dynamics of monolayer crystals revealed by femtosecond surface x-ray scattering. *Nature Photonics*, 13, 425-430 <https://doi.org/10.1038/s41566-019-0387-5>

[9] M. F. Lin, V. Kochat, A. Krishnamoorthy, L. Bassman, C. Wening, Q. Zheng, X. Zhang, A. Apte, C. S. Tiwary, X. Z. Shen, R. K. Li, R. Kalia, P. Ajayan, A. Nakano, P. Vashishta, F. Shimojo, X. J. Wang, D. M. Fritz and U. Bergmann. 2017. Ultrafast non-radiative dynamics of atomically thin MoSe2. *Nature Communications*, 8, 1745 <https://doi.org/10.1038/s41467-017-01844-2>

[10] M. Z. Mo, Z. Chen, R. K. Li, M. Dunning, B. B. L. Witte, J. K. Baldwin, L. B. Fletcher, J. B. Kim, A. Ng, R. Redmer, A. H. Reid, P. Shekhar, X. Z. Shen, M. Shen, K. Sokolowski-Tinten, Y. Y. Tsui, Y. Q. Wang, Q. Zheng, X. J. Wang and S. H. Glenzer. 2018. Heterogeneous to homogeneous melting transition visualized with

ultrafast electron diffraction. *Science*, 360 (6396), 1451-1454 <https://doi.org/10.1126/science.aar2058>

[11] N. A. Romero, A. Nakano, K. Riley, F. Shimojo, R. K. Kalia, P. Vashishta and P. C. Messina. 2015. Quantum molecular dynamics in the post-petaflops era. *IEEE Computer*, 48 (11), 33-41 <https://doi.org/10.1109/MC.2015.337>

[12] L. Bassman, A. Krishnamoorthy, H. Kumazoe, M. Misawa, F. Shimojo, R. K. Kalia, A. Nakano and P. Vashishta. 2018. Electronic origin of optically-induced sub-picosecond lattice dynamics in MoSe2 monolayer. *Nano Letters*, 18 (8), 4653-4658 <https://doi.org/10.1021/acs.nanolett.8b00474>

[13] R. F. Service. 2018. Design for U.S. exascale computer takes shape. *Science*, 359 (6376), 617-618 <https://doi.org/10.1126/science.359.6376.617>

[14] F. Shimojo, R. K. Kalia, M. Kunaseth, A. Nakano, K. Nomura, S. Ohmura, K. Shimamura and P. Vashishta. 2014. A divide-and-conquer-recombine algorithmic paradigm for multiscale materials modeling. *Journal of Chemical Physics*, 140 (18), 18A529 <https://doi.org/10.1063/1.4869342>

[15] K. Nomura, R. K. Kalia, A. Nakano, P. Vashishta, K. Shimamura, F. Shimojo, M. Kunaseth, P. C. Messina and N. A. Romero. 2014. Metascalable quantum molecular dynamics simulations of hydrogen-on-demand. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'14)*, IEEE, New Jersey, NJ, USA, 661-673 <https://doi.org/10.1109/SC.2014.59>

[16] P. Hohenberg and W. Kohn. 1964. Inhomogeneous electron gas. *Physical Review*, 136 (3B), B864-B871 <https://doi.org/10.1103/PhysRev.136.B864>

[17] W. Kohn and L. J. Sham. 1965. Self-consistent equations including exchange and correlation effects. *Physical Review*, 140 (4A), A1133-A1138 <https://doi.org/10.1103/PhysRev.140.A1133>

[18] S. Goedecker. 1999. Linear scaling electronic structure methods. *Reviews of Modern Physics*, 71 (4), 1085-1123 <https://doi.org/10.1103/RevModPhys.71.1085>

[19] D. R. Bowler and T. Miyazaki. 2012. O(N) methods in electronic structure calculations. *Reports on Progress in Physics*, 75 (3), 036503 <https://doi.org/10.1088/0034-4885/75/3/036503>

[20] W. Kohn. 1996. Density functional and density matrix method scaling linearly with the number of atoms. *Physical Review Letters*, 76 (17), 3168-3171 <https://doi.org/10.1103/PhysRevLett.76.3168>

[21] E. Prodan and W. Kohn. 2005. Nearsightedness of electronic matter. *Proceedings of the National Academy of Sciences*, 102 (33), 11635-11638 <https://doi.org/10.1073/pnas.0505436102>

[22] W. T. Yang. 1991. Direct calculation of electron-density in density-functional theory. *Physical Review Letters*, 66 (11), 1438-1441 <https://doi.org/10.1103/PhysRevLett.66.1438>

[23] F. Shimojo, R. K. Kalia, A. Nakano and P. Vashishta. 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), 151-164 <https://doi.org/10.1016/j.cpc.2005.01.005>

[24] F. Shimojo, R. K. Kalia, A. Nakano and P. Vashishta. 2008. Divide-and-conquer density functional theory on hierarchical real-space grids: parallel implementation and applications. *Physical Review B*, 77 (8), 085103 <https://doi.org/10.1103/PhysRevB.77.085103>

[25] K. Nomura, H. Dursun, R. Seymour, W. Wang, R. K. Kalia, A. Nakano, P. Vashishta, F. Shimojo and L. H. Yang. IEEE, 2009. A metascalable computing framework for large spatiotemporal-scale atomistic simulations. In *Proceedings of the International Parallel and Distributed Processing Symposium, IPDPS 2009, IEEE, New Jersey, NJ, USA*, 1-10 <https://doi.org/10.1109/IPDPS.2009.5160992>

[26] N. D. Woods, M. C. Payne and P. J. Hasnip. 2019. Computing the self-consistent field in Kohn-Sham density functional theory. *Journal of Physics: Condensed Matter*, 31, 453001 <https://doi.org/10.1088/1361-648X/ab31c0>

[27] M. Kunaseth, R. K. Kalia, A. Nakano, K. Nomura and P. Vashishta. 2013. A scalable parallel algorithm for dynamic range-limited n-tuple computation in many-body molecular dynamics simulation. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'13)*, ACM, New York, NY, USA, <https://doi.org/10.1145/2503210.2503235>

[28] M. Kunaseth, S. Hannongbua and A. Nakano. 2019. Shift/collapse on neighbor list (SC-NBL): fast evaluation of dynamic many-body potentials in molecular dynamics simulations. *Computer Physics Communications*, 235, 88-94 <https://doi.org/10.1016/j.cpc.2018.09.021>

[29] L. Lin, J. Lu, L. Ying and W. N. E. 2019. Adaptive local basis set for Kohn-Sham density functional theory in a discontinuous Galerkin framework I: Total energy calculation. *J. Comput. Phys.*, 231 (4), 2140-2154 <https://doi.org/10.1016/j.jcp.2011.11.032>

[30] A. Brandt. 1977. Multilevel adaptive solutions to boundary-value problems. *Mathematics of Computation*, 31 (138), 333-390 <https://doi.org/10.1090/S0025-5718-1977-0431719-X>

[31] W. Hu, L. Lin and C. Yang. 2015. DGDFT: A massively parallel method for large scale density functional theory calculations. *The Journal of Chemical Physics*, 143 (12), 124110 <https://doi.org/10.1063/1.4931732>

[32] A. S. Banerjee, L. Lin, P. Suryanarayana, C. Yang and J. E. Pask. 2018. Two-level Chebyshev Filter Based Complementary Subspace Method: Pushing the Envelope of Large-Scale Electronic Structure Calculations. *Journal of Chemical Theory and Computation*, 14 (6), 2930-2946 <https://doi.org/10.1021/acs.jctc.7b01243>

- [33] L. Lin, J. F. Lu, L. X. Ying, and W. N. E. 2012. Adaptive local basis set for Kohn-Sham density functional theory in a discontinuous Galerkin framework I: total energy calculation. *Journal of Computational Physics*, 231, 4 (Feb 20), 2140-2154. [10.1016/j.jcp.2011.11.03](https://doi.org/10.1016/j.jcp.2011.11.03)
- [34] H. Dursun, M. Kunaseth, K. Nomura, J. Chame, R. F. Lucas, C. Chen, M. Hall, R. K. Kalia, A. Nakano and P. Vashishta. 2012. Hierarchical parallelization and optimization of high-order stencil computations on multicore clusters. *Journal of Supercomputing*, 62 (2), 946-966 <https://doi.org/10.1007/s11227-012-0764-z>
- [35] A. Sodani, R. Gramunt, J. Corbal, H. S. Kim, K. Vinod, S. Chinthamani, S. Hutsell, R. Agarwal and Y. C. Liu. 2016. Knights Landing: second-generation Intel Xeon Phi product. *IEEE Micro*, 36 (2), 34-46 <https://doi.org/10.1109/MM.2016.25>
- [36] K. Liu, S. Hong, R. K. Kalia, A. Nakano, K. Nomura, P. Rajak, S. Tiwari, P. Vashishta, Y. Luo, N. A. Romero, S. S. Naserifar, W. A. Goddard and M. Kunaseth. IEEE, 2018. Shift-collapse acceleration of generalized polarizable reactive molecular dynamics for machine learning-assisted computational synthesis of layered materials. *Proceedings of Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems, ScalA18*, 41-48 <https://doi.org/10.1109/ScalA.2018.00009>
- [37] X. Yang, J. Jenkins, M. Mubarak, R. B. Ross and Z. Lan. 2016. Watch out for the bully! job interference study on dragonfly network. *Proceedings of Supercomputing, SC16* <https://doi.org/10.1109/SC.2016.63>
- [38] B. Alverson, E. Froese, L. Kaplan and D. Roweth. 2012. *Cray XC series network. Cray Research White Paper.*
- [39] S. Parker, V. Morozov, S. Chunduri, K. Harms, C. Knight and K. Kumaran. 2017. Early evaluation of the Cray XC40 Xeon Phi system 'Theta' at Argonne. *Sci Tech Connect*
- [40] A. Krishnamoorthy, M.-F. Lin, C. Weninger, R. Ma, A. Britz, C. S. Tiwari, V. Kochat, A. Apte, J. Yang, S. Park, R. Li, X. Shen, X. Wang, R. K. Kalia, A. Nakano, F. Shimojo, D. Fritz, U. Bergmann, P. Ajayan and P. Vashishta. 2019. Semiconductor-metal structural phase transformation in MoTe2 monolayers by electronic excitation. *Nano Letters*, 19 (8), 4981-4989 <https://doi.org/10.1021/acs.nanolett.9b01179>
- [41] M. M. Ugeda, A. J. Bradley, S. F. Shi, F. H. da Jornada, Y. Zhang, D. Y. Qiu, W. Ruan, S. K. Mo, Z. Hussain, Z. X. Shen, F. Wang, S. G. Louie and M. F. Crommie. 2014. Giant bandgap renormalization and excitonic effects in a monolayer transition metal dichalcogenide semiconductor. *Nature Materials*, 13 (12), 1091-1095 <https://doi.org/10.1038/NMAT4061>
- [42] Y. Cao, V. Fatemi, S. Fang, K. Watanabe, T. Taniguchi, E. Kaxiras and P. Jarillo-Herrero. 2018. Unconventional superconductivity in magic-angle graphene superlattices. *Nature*, 556 (7699), 43-50 <https://doi.org/10.1038/nature26160>
- [43] S. Das, Y. L. Tang, Z. Hong, M. A. P. Goncalves, M. R. McCarter, C. Klewe, K. X. Nguyen, F. Gomez-Ortiz, P. Shafer, E. Arenholz, V. A. Stoica, S. L. Hsu, B. Wang, C. Ophus, J. F. Liu, C. T. Nelson, S. Saremi, B. Prasad, A. B. Mei, D. G. Schlom, J. Iniguez, P. Garcia-Fernandez, D. A. Muller, L. Q. Chen, J. Junquera, L. W. Martin and R. Ramesh. 2019. Observation of room-temperature polar skyrmions. *Nature*, 568 (7752), 368-372 <https://doi.org/10.1038/s41586-019-1092-8>