

Quantum Dynamics Simulation of Electrons in Materials on High-Performance Computers

André Schleife, Erik W. Draeger, Victor Anisimov, Alfredo A. Correa, and Yosuke Kanai

André Schleife:

André Schleife is a *Blue Waters* Assistant Professor in the Department of Materials Science and Engineering at the University of Illinois at Urbana-Champaign. He obtained his Diploma (2006) and Ph.D. (2010) in physics at Friedrich-Schiller-University in Jena, Germany for his theoretical work on transparent conducting oxides. Before he started at UIUC he was a Directorate Postdoctoral Researcher at Lawrence Livermore National Laboratory (2011-2013) working on a description of non-adiabatic electron-ion dynamics. His current research focus is on excited electronic states and their real-time dynamics in various material systems using accurate, first-principles computational methods and making use of modern super computers. He is a member of APS, MRS, DFG, and SPIE. Contact him at schleife@illinois.edu.

Complete contact information:

Department of Materials Science and Engineering
University of Illinois at Urbana-Champaign
1304 W. Green St.
Urbana, IL 61801
Tel : (217)-244-0339
Fax : (217)-333-2736

Erik W. Draeger:

Erik Draeger is a staff scientist at the Center for Applied Scientific Computing, in the Computation Directorate at Lawrence Livermore National Laboratory. He received his Bachelor's degree in physics from the University of California, Berkeley in 1995 and his Ph.D. in theoretical physics from the University of Illinois, Urbana-Champaign in 2001. His primary research interest is the development of massively-parallel scientific applications, with a focus on materials science and biological applications. He currently leads the development of `qb@ll`, the LLNL branch of the Qbox first-principles molecular dynamics code. Contact him at draeger1@llnl.gov

Complete contact information:

Center for Applied Scientific Computing
7000 East Ave., L-561
Livermore, CA 94550

Tel : (925)423-3650

Victor Anisimov:

Victor Anisimov is a Senior Research Programmer in the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign. He holds Ph.D. degree in Physical Chemistry from Institute of Chemical Physics Russian Academy of Sciences (1997), followed by 5 years of computational chemistry software development with Fujitsu. He performed postdoctoral work in University of Maryland at Baltimore (2003-2008) and in University of Texas at Houston (2008-2011). Dr. Anisimov specializes in high-performance computing of soft matter and in performance optimization of quantum chemistry codes. He is an author of the linear scaling semiempirical method LocalSCF. Contact him at anisimov@illinois.edu.

Complete contact information:

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign
1205 W. Clark St.
Urbana, IL 61801
Tel : (217)-333-2799
Fax : (217)-244-9757

Alfredo A. Correa:

Alfredo A. Correa is Staff Scientist at Lawrence Livermore National Laboratory, Condensed Matter and Materials Division where he was previously a Lawrence Fellow during 2009-2012. He obtained a B.A. from Instituto Balseiro in 2001 and a Ph.D. in Physics from University of California at Berkeley in 2008, and he was subsequently a post-doctoral scholar at Stanford University, Mechanical Engineering department for 2009. His current research focus is on atomistic simulation of electron-ion dynamics, material at extreme conditions and high energy density physics modeling. Contact him at correaa@llnl.gov.

Complete contact information:

Lawrence Livermore National Laboratory
7000 East Avenue
Livermore, CA 94550
Tel : (925)-422-0102
Fax: (925)-422-6594

Yosuke Kanai:

Yosuke Kanai is Assistant Professor at the University of North Carolina at Chapel Hill in Department of Chemistry and Visiting Scientist in Condensed Matter and Materials Division at Lawrence Livermore National Laboratory, where he was previously a Lawrence Fellow and also served as director of Computational Chemistry and Materials Science Summer

Institute for 2009-2011. He obtained a Ph.D. in theoretical chemistry from Princeton University in 2006, and he was subsequently a Berkeley Nanoscience and Nanoengineering Institute post-doctoral scholar at University of California at Berkeley for 2006-2009. His research interests are on development and application of electronic structure methods based on first-principles theory for investigating novel phenomena in materials and condensed matter systems. He is a member of APS, MRS, and ACS. Contact him at ykanai@unc.edu.

Complete contact information:

Department of Chemistry
Campus Box 3290
The University of North Carolina at Chapel Hill
Chapel Hill, NC 27599-3290
Tel : (919)-962-3891
Fax : (919)-962-2388

Abstract:

Advancement in high-performance computing allows us to calculate properties of increasingly complex materials with unprecedented accuracy. At the same time, in order to take full advantage of modern leadership-class supercomputers, the calculations need to scale well on hundreds of thousands of processing cores. We demonstrate such high scalability of our recently developed implementation of Ehrenfest non-adiabatic electron-ion dynamics up to 1 million floating-point processing units on two different leadership-class computing architectures. As a representative example of material properties that derive from quantum dynamics of electrons, we demonstrate the accurate calculation of electronic stopping power, which characterizes the rate of energy transfer from a high-energy particle to electrons in materials. We discuss the specific case of crystalline gold with a hydrogen atom as the high-energy particle, and we illustrate detailed scientific insights that can be obtained from the quantum dynamics simulation at electronic structure level.

Keywords: Quantum electron dynamics, High-performance computing, Computational Materials Science

Introduction

Over the past half-century, computation has become an indispensable tool in science and engineering. On the road to exascale, supercomputers with multi-peta-FLOP per second (PFlop/s) performance have become increasingly available. The continued advancement of high-performance computational capabilities allows us to solve mathematical problems of increasing complexity with greater accuracy, providing detailed scientific understanding and insights for technological innovation. At the same time, the increasing scale and complexity of modern supercomputers have exposed the need for a greater emphasis on development of algorithms and methodologies that are tailored for massively-parallel multi-core machines in order to fully take advantage of these capabilities. Computational materials research has traditionally benefited greatly from high-performance computing, and increased effort on addressing this issue is becoming crucial for the continued advancement of the field.

In order to systematically design new materials or even to improve existing ones, obtaining structure–function–property relationships with predictive accuracy is essential. Oftentimes this requires a deep understanding of material properties that emerge from many-body interactions of electrons and ions (i.e. atoms with electrons missing) that constitute the material. At the microscopic scale of electrons, quantum mechanical effects play a central role, and the mathematical description of many-body interactions becomes highly non-trivial. Furthermore, many challenging problems in materials research call for predictive calculations without the need for adjustable empirical parameters from experiments (i.e. *first-principles* approaches). High-performance computing is indispensable for tackling this challenging goal.

The behavior of quantum-mechanical particles such as electrons is governed by the many-body Schrödinger equation, a multi-dimensional differential equation. Because a direct solution of this equation has exponential complexity [1], it is computationally intractable for all but the smallest and simplest molecules even on today’s supercomputers. Scientists thus rely on physical and numerical approximations for computing electronic structure to predict material properties. Naturally, the utility of quantum-mechanical calculations depends on the accuracy of these underlying approximations that are involved in first-principles computational approaches. Recent advances of quantum-mechanical calculations are indeed closely tied to the emergence of high-performance computers in the last few decades, and they allow us to employ better approximations which have become highly complicated for numerical calculations.

Among various quantum mechanical computational methods for obtaining electronic structure, density functional theory (DFT) has and continues to play a central role in computational materials research [1,2]. Enabled by efficient implementations of DFT methods, “high-throughput” computational searches for optimal materials are becoming popular for various technological applications such as batteries and photo-electrochemical cells, for instance, within the Materials Project [3]. Computational prediction of optical and electronic properties of materials also advanced greatly in recent years with many-body perturbation theory and quantum Monte Carlo methodologies, especially with increasing research effort for solar energy conversion [4]. As for calculating time-dependent properties, dynamical behavior of complex materials and condensed phase systems can be simulated today with first-principles molecular dynamics approaches [5].

In first-principles molecular dynamics approaches, the quantum dynamics of electrons is coarse-grained out to reduce the computational cost. The electronic evolution is designed to simply follow the movement of ions instantaneously (known as the adiabatic Born-Oppenheimer approximation [6]). Unfortunately, in this approximation the physics associated with the quantum dynamics of electrons is neglected and various interesting properties are therefore not accessible. Electron dynamics is indeed responsible for many novel phenomena that are essential for a wide range of technological applications—ranging from photocatalytic cells to radiation shielding in space. Consequently, an accurate description of electron dynamics through time-dependent quantum-mechanical theory is an important challenge in computational materials physics and chemistry today. While modern high-performance supercomputers help us tackle this great challenge, their massively parallel, hybrid-paradigm (traditional central-processing units, graphic chips, many integrated cores) architectures presents new challenges for numerical simulations at the same time. Theoretical

and algorithmic methods need to be developed accordingly in order for electronic-structure theory to take full advantage of current and future supercomputers.

In this paper, we discuss our recent effort on developing first-principles computational methodology to simulate the quantum dynamics of electrons in materials on massively parallel supercomputers. We illustrate the approach with a prototypical example of combined electron-ion dynamics for calculating *electronic stopping*, i.e. the non-equilibrium energy transfer from high-energy particles (such as atoms and ions) to the electrons in materials as the projectile particle traverses through it. We demonstrate the computational capability to make quantitative predictions of electronic stopping that are needed for various technological developments such as nuclear reactor materials design and radiation hardening of microelectronics in space.

Theoretical and Numerical Approach

The Runge-Gross theorem [7] allows us to circumvent a prohibitively expensive task of solving the $3N$ -dimensional, time-dependent (TD) Schrödinger equation for N electrons. After mapping the problem onto a system of non-interacting Kohn-Sham (KS) particles, a set of coupled non-linear partial differential equations, known as TD-KS equations, has to be solved for the time evolution of the quantum system. TD-KS equations read

$$i\hbar \frac{\partial}{\partial t} \phi_i(\mathbf{r}, t) = \left\{ -\frac{1}{2} \nabla^2 + V_{\text{ext}}(\mathbf{r}, t) + V_s[n](\mathbf{r}, t) \right\} \phi_i(\mathbf{r}, t) \quad (1)$$

where \mathbf{r} is the 3-dimensional spatial coordinate, t is time, and $V_{\text{ext}}(\mathbf{r}, t)$ is the external potential that describes the Coulomb interaction of electrons and ions. The term $V_s[n](\mathbf{r}, t)$ represents the sum of electrostatic (Hartree) electron-electron interaction, $\int \frac{n(\mathbf{r}', t)}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}'$, and quantum-mechanical exchange-correlation (XC) potential, $\frac{\delta E_{\text{XC}}[n]}{\delta n(\mathbf{r}, t)}$. The terms between the brackets in Eq. (1) are collectively called the “Kohn-Sham Hamiltonian”, and it drives the time evolution of the electronic system. Importantly, this Hamiltonian depends on the electronic density $n(\mathbf{r}, t)$, a function of three spatial and one temporal variables, that is given as the sum of TD-KS wave functions, $n(\mathbf{r}, t) = \sum_i f_i |\phi_i(\mathbf{r}, t)|^2$, weighed by their respective occupation f_i . The *exact* XC potential, albeit unknown, is defined through a functional derivative and an approximation is used in practice for the XC functional. By numerically integrating the TD-KS equations (1) in time, we obtain the quantum dynamics of electrons from first principles.

In order to combine the quantum dynamics of electrons with the movement of classical ions, we rely on Ehrenfest dynamics approach [8]. In this framework, forces on classical ions are obtained from the total energy functional of the time-dependent electronic density. This results in the following equation of motion for the ions:

$$M_I \frac{d^2 \mathbf{R}_I}{dt^2} = -\nabla_{\mathbf{R}_I} E[\{\mathbf{R}_I\}, n(\mathbf{r}, t)] \quad (2)$$

where M_I and \mathbf{R}_I are the mass and position of the I -th ion.

The most time-consuming part of employing Ehrenfest dynamics as described above is the numerical integration of TD-KS equations. The numerical approach has to be accurate,

stable, and efficient, and importantly it must perform well on highly parallelized computers with hundreds of thousands of processing cores. Consequently, a suitable numerical integrator must be balanced in terms of integration step size, numerical stability, accuracy, and scalability for massive parallelization. We have recently shown that a generalized 4th order Runge-Kutta explicit integration scheme is highly suitable when combined with a plane-wave expansion of TD-KS wave functions [9], and we also find it convenient for computations with large numbers of processors, as demonstrated in the next section. The numerical integration is conditionally stable, and an integration step size on the order of one atto-second can be used for the plane-wave basis set sizes that are typically needed for investigating materials. The plane-wave expansion allows for a systematic convergence of the basis set and also an efficient computation of key quantities (e.g. kinetic energy) in reciprocal space using highly optimized fast Fourier transform routines. Here, we demonstrate that the new algorithm [9] can be efficiently implemented.

Computational Details

The Qbox code [10] is an efficient, scalable implementation of DFT using a plane-wave expansion of the KS wave functions, $\phi_i(\mathbf{r}, t) = \frac{1}{\sqrt{\Omega}} \sum_{\mathbf{G}} C_i(\mathbf{G}, t) e^{-i\mathbf{G}\cdot\mathbf{r}}$. Written in C++, Qbox uses MPI to carry out communication between compute nodes and a mix of OpenMP and threaded kernels to maximize usage of on-node resources. For this work, we used the Qb@ll branch of the code [10] currently developed and maintained at Lawrence Livermore National Laboratory (LLNL).

A typical quantum dynamics simulation consists of many individual time steps and the same operations are performed for each step. Hence, we investigate the scaling behavior for a representative 1600-atom gold system (27,200 electrons) at Γ point in Brillouin zone integration by showing the wall time for one individual step in Fig. 1. We note that a 256 atom super cell is sufficient to make predictions agreeing with the experimental data for this specific system (see below). The benchmark, however, shows that the code and the existing high-performance computing machines are capable of treating much more complex problems. Two different computational architectures are compared: Sequoia, the 20 PFlop/s IBM Blue Gene/Q supercomputer at LLNL having theoretical peak performance of 204.8 GF/node, and Blue Waters, the 13.3 PFlop/s Cray Blue Waters supercomputer at the National Centre for Supercomputing Applications (NCSA) having theoretical peak performance of 313.6 GF/node for XE node. On both Sequoia and Blue Waters each node has 16 floating-point processing units and this is what we use as reference for the scaling in Fig. 1. The timings for successive steps (maximum over all processes) using a given number of floating-point units are similar to within less than 1%.

The calculations on the Blue Gene/Q system used parallel linear algebra as implemented in the SCALAPACK library and one-dimensional Fourier transforms are computed with either the FFTW or vendor-supplied libraries such as IBM's ESSL library. For the calculations on the Cray system, parallel linear algebra is handled by SCALAPACK as included in the vendor-optimized Cray LIBSCI library. On the Cray platform, one-dimensional Fourier transforms are computed with the FFTW library. Computations were performed on Cray XE6 nodes (i.e. no GPUs were used for any of the calculations reported here).

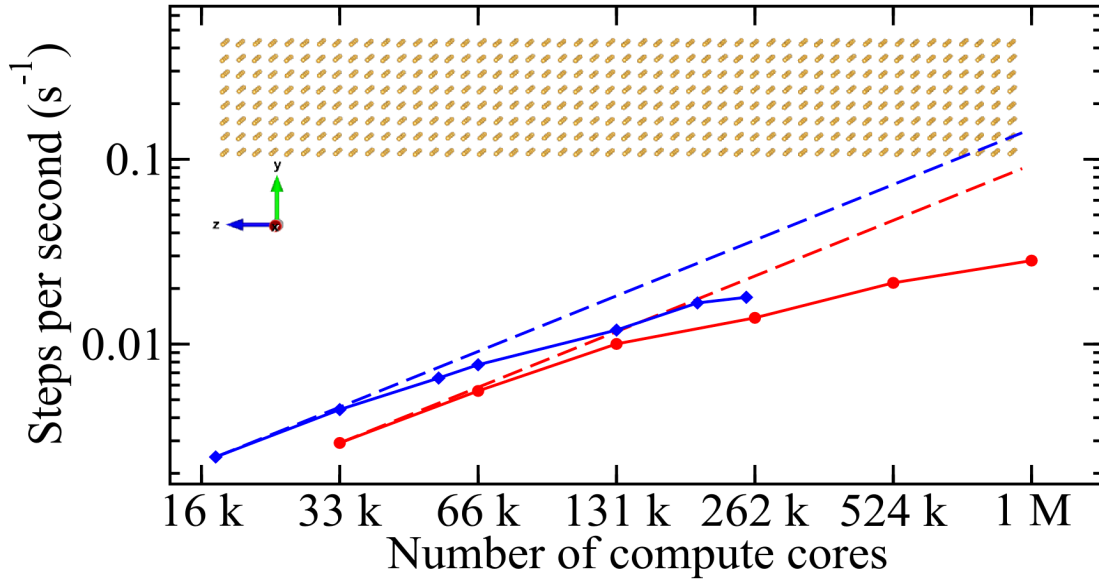


Figure 1: Scaling of the TD-KS calculation for 1600 atom gold (27,200 electrons) on the IBM Blue Gene/Q Sequoia machine (red circles) and Cray Blue Waters machine (blue diamonds). Maximum performance is 1.5 and 0.8 PFlop/s respectively. The inset shows the super cell that this simulation is based on.

In this work we adopt the memory management and data distribution schemes of the Qb@ll code: The expansion coefficients, $C_i(\mathbf{G},t)$, of the KS electronic wave functions are distributed on a logical 2D process grid, with electronic states distributed across process rows and plane-wave basis functions distributed across process columns (see Fig. 2). The TD-KS equations are well-suited for such a distribution, as nearly all communication can be restricted to sub-communicator collectives within a process row or column. Four additional sets of the wave function are stored in memory for each integration step in the fourth-order Runge-Kutta integration scheme we have implemented [9]. Note that once the ground state of the quantum system is computed (e.g. initial condition for the time-propagation), TD-KS equation does not require re-orthogonalization of the wave functions or subspace diagonalization every iteration like traditional Born-Oppenheimer dynamics. Our TD-KS implementation is therefore free of the global communication bottlenecks that could limit the strong scalability of standard DFT calculations.

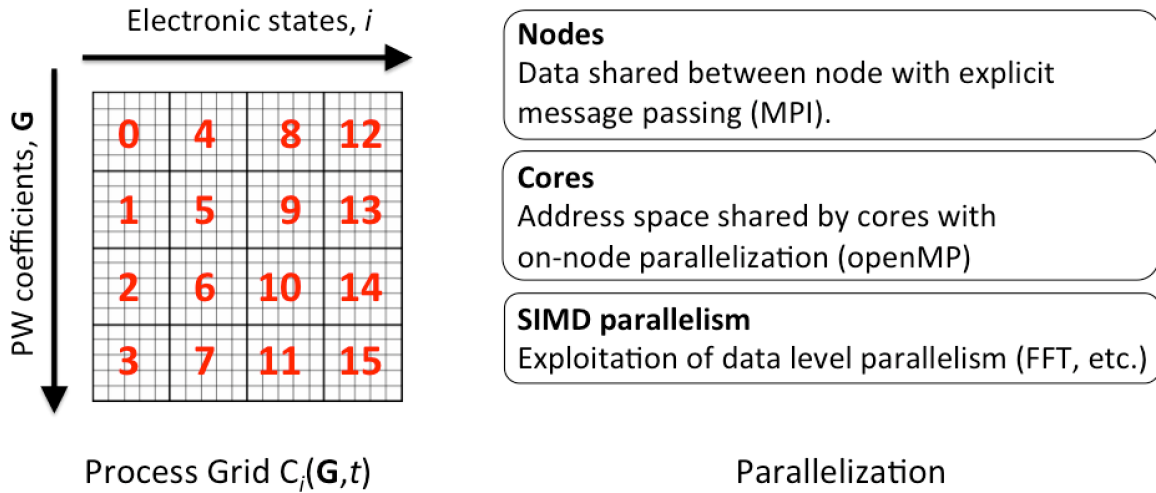


Figure 2: A schematic illustrating representation of the Kohn-Sham electronic wave functions on two-dimensional logical process grid: Expansion coefficients $C_i(\mathbf{G},t)$ for plane waves are ordered in terms of the state index, i and the expansion index \mathbf{G} . Red numbers indicate MPI process numbers.

Modern high-performance computing often requires keen awareness of the underlying hardware: For example, on Blue Gene/Q, the theoretical peak performance assumes full usage of the four-way single-instruction-multiple-data (SIMD) vector units and hardware threads. If one does not take advantage of these features, the maximum performance of the machine drops from 20 PFlop/s to 2.5 PFlop/s. For Qbox, optimized single-node kernels were found to be essential in minimizing the time to solution, most notably matrix multiplication and one-dimensional fast Fourier transforms. On Sequoia we were able to achieve up to 35% of peak (on 32,768 cores) and about 11% on 1 million cores. On Blue Waters’ XE nodes this run achieved 33.1% of peak (not counting GPUs) on 32,768 cores and 20.8% on 196,608 cores.

Test calculations on Sequoia provided insight into the timings and scaling of the three main communication operations (“MPI_Reduce”, “MPI_Allreduce”, and “MPI_Alltoallv”). The total time spent on reduction operations depends only weakly on the number of cores since these are typically carried out across rows on the logical process grid, and the actual amount of data being moved does not change significantly. The timing slightly decreases, depending on how spread out the data is. The “alltoallv” operations, on the other hand, are carried out within process columns as part of the Fourier transforms to compute the charge density. For those operations, the number of electronic states on each process column decreases linearly with the increasing number of nodes, leading to excellent scaling.

At larger task counts, the strong scalability is limited not by communication costs, but by the lack of parallelization over terms that depend on data that is duplicated across tasks e.g. ionic positions. At one million cores, over half of the total time is spent in these heretofore insignificant calculations. The dominance of these terms is a consequence of the unprecedented scalability of the rest of the code. Additional development is currently underway to further distribute the calculations of these terms and achieve greater scalability. Nevertheless, Fig. 1 shows that we obtain excellent scaling up to very large numbers of cores on both of these leadership-class machines and, in addition, that we can apply this approach to complex materials with thousands of atoms and tens of thousands of electrons.

Scientific Application: Electronic Stopping




Figure 3: Time evolution of the difference of the time-dependent electronic charge density and the density at $t=0$ for a hydrogen atom moving through gold at a velocity of $v=2$ a.u. The snapshots are taken at $t=0.02$ a.u. (a and d), $t=0.42$ a.u. (b and e), and $t=0.82$ a.u. (c and f). The top three images show a channeling projectile and the bottom three ones an off-channeling projectile. Red color indicates $n(t)<n(0)$ and blue indicates $n(t)>n(0)$.

We now demonstrate an application of the Ehrenfest dynamics implementation for computing an important material property that is governed by coupled electron-ion dynamics and, hence, explicitly depends on the quantum dynamics of electrons; Electronic stopping characterizes the rate of non-equilibrium energy transfer from a projectile (e.g. atoms, ions, etc.) to the electronic system of a target material as the projectile moves through the material. The energy transfer results in excitations of the electrons in the material and this process strongly depends on charge and velocity of the projectile, the composition of the target material, and also on the trajectory of the projectile through the material. Here we demonstrate the computational capability for quantitatively calculating electronic stopping for the prototypical case of a fast hydrogen (effectively a proton in bulk gold) penetrating crystalline gold.

Detailed knowledge of electronic stopping such as its velocity and geometrical dependence is highly important because it describes the interaction of high-energy projectiles with materials under various extreme conditions. Quantitative understanding is essential, for example, in predicting structural damages in fusion/fission reactor materials as they face significant bombardment by high-energy protons and other particle radiation. Radiation shielding of microelectronics from high-energy protons in space and the interaction of high-energy protons with biological molecules as in radiation therapy are other important areas of materials research where detailed characterization of electronic stopping is crucial.




Figure 4: Spatially-averaged electronic stopping S as a function of the projectile velocity. The blue (solid) line represents averaged experimental data from the SRIM database [11]. The black (square symbols) and the red (diamond symbols) lines are for channeling and for off-channeling projectile, respectively.

The projectile deposits energy into the electronic system as it moves through the material at a given (constant) velocity. Our simulations provide us with the energy E (defined here as the *total* energy minus the kinetic energy of all ions) and the electronic charge density at each step of the time integration of the TD-KS equations. Electronic stopping is defined as the increase of E as a function of projectile displacement x ,

$$S(x) = dE(x)/dx. \quad (3)$$

This stopping power $S(x)$ has the dimension of a force and can also be understood as a “drag force” acting on the projectile [12].

In order to compute the electronic stopping of gold with a hydrogen atom as the projectile, we use a simulation cell of 256 gold atoms with periodic boundary conditions, and only the Γ point is used for Brillouin zone integrations. More than 840,000 plane waves are used in the expansion of each KS wave function (130 Rydberg for the plane wave energy cutoff), and norm-conserving pseudopotentials (with $5p$, $5d$, and $6s$ electrons of gold as valence electrons) are used to describe the electron-ion interaction. We use the adiabatic local-density approximation to describe exchange and correlation. The integration time step of $\Delta t=0.24$ atto-seconds is used for accurate integration of the time-dependent KS equations for all the projectile velocities.

The influence of the stopping geometry is quantified by comparing results for a projectile along a *channeling* trajectory (an ideal [001] channel is used in this work) to the results computed for a projectile along a randomly-oriented off-channeling path as in most experiments. The latter is approximated by placing the direction of the projectile’s velocity vector to be incommensurate with the lattice vectors of the bulk crystalline gold. The two stopping geometries are shown in Fig. 3. In addition, Fig. 3 also shows the time evolution of

the difference of the time-dependent electronic charge density (i.e. as the projectile passes through gold) and the equilibrium electronic density at $t=0$. The density changes induced in this process are clearly visible. We observe that the density oscillation lags behind the projectile and its magnitude depends strongly on the velocity. The oscillation remains rather localized in the orthogonal direction of the projectile trajectory. These insights into the electronic stopping process at microscopic level, obtained without any fitting parameter to experiment, help us understand the initial mechanism of material damages induced by high-energy projectiles.

As an example of quantitative characterization, Fig. 4 compares our computed results for the velocity dependence of averaged electronic stopping to experimental data from the SRIM database that summarizes a large number of experimentally measured results [11]. We are able to accurately calculate the electronic stopping over a wide range of projectile velocities in very good agreement with experimental measurements. In addition, we can obtain detail atomistic insights into the influence of the projectile trajectories, and we find that the channeling trajectory results in less energy transfer per unit distance for the projectile velocities higher than 1.0 atomic units. We note that each data point in Fig. 4 requires between 400 and 24000 time steps, depending on the projectile velocity. Accurate calculations of the stopping power for materials such as gold is computationally quite expensive because a large number of (semi-) core electrons need to be taken into account. In addition to the large computational cost associated with obtaining accurate averages for the stopping power, the increased memory requirement as discussed in the previous section makes it necessary to perform these calculations on highly-parallelized computers.

Conclusions and Perspective

Computational materials research has greatly benefited from recent advancements in high-performance computing. Coupled with new developments in electronic structure theory, various technologically important material properties can now be computed from first principles without empirical parameters from experiments. At the same time, computing non-adiabatic properties of materials that derive from the quantum dynamics of electrons remains a significant challenge, and leadership-class supercomputers allow us to meet this challenge for many complex materials of great technological importance.

In this article, we described our recent effort in developing a new highly scalable first-principles approach that can explicitly model electron dynamics in complex materials by efficiently utilizing modern supercomputers. We presented performance data of our implementation on two leadership-class computing architectures and we obtained scientific results for the electronic stopping power of crystalline gold with hydrogen atom. We envision that collaborative interactions between physical science and computer science communities will be increasingly important for computational materials research in the future.

Acknowledgments

This research is part of the Blue Waters sustained-petascale computing project, which is supported by the National Science Foundation (award number OCI 07-25070) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications. Computing support for this work came from the Lawrence Livermore National Laboratory (LLNL) Institutional Computing Grand Challenge program. Part of this work was performed under the auspices of the U.S. Department of Energy at Lawrence Livermore National Laboratory under Contract DE-AC52-07A27344.

References

- [1] W. Kohn, Rev. Mod. Phys. 71, 1253 (1999).
- [2] R. M. Martin, Electronic Structure: Basic Theory and Practical Methods (Cambridge University Press, 2008).
- [3] A. Jain, G. Hautier, C. J. Moore, S. P. Ong, C. C. Fischer, T. Mueller, K. A. Persson, and G. Ceder, Comput. Mater. Sci. 50, 2295 (2011).
- [4] Y. Kanai, J. B. Neaton, and J. C. Grossman, Computing in Science Engineering 12, 18 (2010).
- [5] D. Marx and H. Jurg, Ab Initio Molecular Dynamics: Basic Theory and Advanced Methods (Cambridge University Press, 2009).
- [6] M. Born and J. R. Oppenheimer, Ann. Phys. 389, 457 (1927).
- [7] E. Runge and E. K. U. Gross, Phys. Rev. Lett. 52, 997 (1984).
- [8] P. Ehrenfest, Z. Phys. A Hadron. Nucl. 45, 455 (1927).
- [9] A. Schleife, E. W. Draeger, Y. Kanai, and A. A. Correa, J. Chem. Phys. 137, 22A546 (2012).
- [10] F. Gygi, Qbox open source code project, Tech. Rep. (University of California, Davis) <http://eslab.ucdavis.edu/>; E. W. Draeger and F. Gygi, Qbox code, Qb@ll version, Tech. Rep. (Lawrence Livermore National Laboratory).
- [11] M. D. Z. James F. Ziegler, Jochen P. Biersack, The Stopping and Range of Ions in Matter (2008); <http://www.srim.org/SRIM/SRIMPICS/STOP01/STOP0179.gif>
- [12] A. A. Correa, J. Kohanoff, E. Artacho, D. Sánchez-Portal, and A. Caro, Phys. Rev. Lett. 108, 213201 (2012).