Ab initio molecular dynamics simulations using a Chebyshev-filtered subspace iteration technique

K. H. Khoo,[1] M. Kim,[3] G. Schofield,[1] and James R. Chelikowsky[1,2,3,4]
[1]*Center for Computational Materials, Institute for Computational Engineering and Sciences, University of Texas, Austin, Texas 78712, USA*
[2]*Department of Chemistry and Biochemistry, University of Texas, Austin, Texas 78712, USA*
[3]*Department of Chemical Engineering, University of Texas, Austin, Texas 78712, USA*
[4]*Department of Physics, University of Texas, Austin, Texas 78712, USA*
(Received 8 June 2010; revised manuscript received 19 July 2010; published 6 August 2010)

A method for performing *ab initio* molecular dynamics simulations using a real-space pseudopotential density-functional approach is presented in this work. The Kohn-Sham equation is represented on a real-space grid and solved using higher-order finite differencing. A crucial ingredient of our method is the replacement of "standard" diagonalizations at each self-consistent-field iteration by a Chebyshev subspace filtering step, resulting in a speed up of an order of magnitude or more. This increase in computational efficiency allows molecular dynamics simulations to be carried out for large systems. As an illustration, we apply this method to the study of liquid Al.

DOI: 10.1103/PhysRevB.82.064201 PACS number(s): 31.15.E−, 61.20.Ja

I. INTRODUCTION

With the advent of improved algorithms and increased computational power, molecular dynamics (MD) simulations have become an indispensable tool for the study of a wide variety of systems such as liquids, disordered solids, and complex molecules.[1–6] Owing to the large system sizes and long simulation times typically encountered in MD simulations, one is often restricted to describing systems using "classical" rather than "quantum" forces.[7,8] However, the mapping of inherently quantum phenomena such as hybridization changes and charge transfer in chemical bonds onto classical potentials is not transparent or straightforward. Here, we adopt an approach known as "Born-Oppenheimer MD" (BOMD), where atomic forces are extracted from a fully quantum-mechanical solution at every time step.[9] The advantage of this method is that uncontrolled errors arising from the use of *ad hoc* classical forces are eliminated.

A highly efficient and widely used technique for evaluating quantum-mechanical forces is given by pseudopotential density-functional theory (DFT).[10,11] In this formalism, the many-electron problem is transformed into a single-electron self-consistent eigenvalue problem given by the Kohn-Sham equation, and the pseudopotential approximation allows strong all-electron ionic potentials to be replaced by significantly weaker pseudopotentials. Basis sets such as plane waves[12,13] and atomic orbitals[14] have been used to represent the Kohn-Sham wave functions. In our work, however, we have chosen to forgo the use of an explicit basis and represent quantities on a uniform real-space grid, owing to the numerous advantages of working in such a representation.[15] First, the grid points are independent of ionic positions and convergence can be straightforwardly achieved by reducing the grid spacing. This avoids problems encountered in atom-centered basis sets where artifacts such as fictitious Pulay forces have to be accounted for.[16] The Hamiltonian matrix on a real-space grid is sparse, which means that quadratic scaling of matrix-vector multiplications can be readily attained. In addition, the matrix is inherently local on the grid representation, which makes implementation on parallel computing platforms simpler and more efficient. This can be contrasted with plane-wave methods where matrix-vector multiplications are performed using fast Fourier transforms, a procedure that involves significantly more nonlocal operations and increased global communications among processors. Another advantage of the real-space method is that localized systems can be treated without the use of supercells.[17] This eliminates the problem of spurious interactions occurring between the system and replicated images in adjacent cells. We would like to point out that there has been previous work on molecular dynamics based on real-space DFT; however, our simulations are based on BOMD, while previous work employed Car-Parinello molecular dynamics.[18–20]

Recently, we made a significant algorithmic advance to our real-space density-functional method using the Chebyshev-filtered subspace iteration technique.[21,22] In this new algorithm, the "standard" diagonalization at each self-consistent field iteration (SCF) is replaced by a Chebyshev subspace filtering step, resulting in a speed up of an order of magnitude or more. This method may be thought of as an approach to solve the nonlinear Kohn-Sham equation using a nonlinear subspace iteration technique, shifting emphasis away from the intermediate linearized Kohn-Sham eigenvalue problem. In general, MD simulations employing DFT are limited to relatively small systems owing to the prohibitive cost of computing atomic forces for a large number of MD steps. However, due to the significant speed up provided by our Chebyshev subspace filtering algorithm, we are now able to perform MD simulations on relatively large systems at a reasonable computational cost. As an illustration of our method, we have applied it to the study of liquid Al, and our findings are compared to existing experimental and theoretical literature to verify the accuracy of our method.

II. ELECTRONIC-STRUCTURE PROBLEM

In DFT, the total energy of a system containing electrons and ions (in positions $\{\boldsymbol{R}_a\}$) is a unique functional of the electron density ρ given by

$$E_{tot}[\rho] = T[\rho] + E_{ion}(\{R_a\},[\rho]) + E_H[\rho] + E_{xc}[\rho] + E_{ion\text{-}ion}(\{R_a\}), \quad (1)$$

where $T[\rho]$ is the kinetic energy, $E_{ion}(\{\boldsymbol{R_a}\},[\rho])$ the electron-ion energy, $E_H[\rho]$ the Hartree potential energy, $E_{xc}[\rho]$ the exchange-correlation energy, and $E_{ion\text{-}ion}$ $(\{\boldsymbol{R_a}\})$ the classical electrostatic energy between ions. The ground-state electron density that minimizes the total-energy functional satisfies the Kohn-Sham equation

$$H\psi_n(\boldsymbol{r}) = \left\{ -\frac{1}{2}\nabla^2 + V_{ion}(\boldsymbol{r}) + V_H[\rho(\boldsymbol{r})] + V_{xc}[\rho(\boldsymbol{r})] \right\} \psi_n(\boldsymbol{r}) = \varepsilon_n \psi_n(\boldsymbol{r}), \quad (2)$$

$$\rho(\boldsymbol{r}) = \sum_n f(\varepsilon_n, T)|\psi_n(\boldsymbol{r})|^2, \quad (3)$$

where ε_n are the energy eigenvalues, $f(\varepsilon_n, T)$ the Fermi-Dirac distribution, $\psi_n(\boldsymbol{r})$ the Kohn-Sham wave functions, $V_{ion}(\boldsymbol{r})$ the ionic potential, $V_H[\rho(\boldsymbol{r})]$ the Hartree potential, and $V_{xc}[\rho(\boldsymbol{r})]$ the exchange-correlation potential given by $\delta E_{xc}[\rho]/\delta\rho$.[10] (In this work, we employ atomic units where $e=m=\hbar=1$.) To solve the self-consistent Kohn-Sham eigenvalue problem, the first step is to construct an approximate $V_H[\rho(\boldsymbol{r})]$ and $V_{xc}[\rho(\boldsymbol{r})]$ from the superposition of atomic electron densities. Equation (2) is then solved using these approximate potentials and the solution wave functions and charge densities are used to update the potentials. This process is repeated until the "input" and "output" potentials agree to within a specified tolerance and a self-consistent solution is realized. The total energy is then calculated using

$$E_{tot}[\rho] = \sum_{n=1}^{n_{occ}} f(\varepsilon_n, T)\varepsilon_n - \frac{1}{2}\int \rho(\boldsymbol{r}) V_H[\rho(\boldsymbol{r})] d^3r + E_{xc}[\rho] - \int \rho(\boldsymbol{r}) V_{xc}[\rho(\boldsymbol{r})] d^3r \quad (4)$$

which can be obtained by combining Eqs. (1) and (2), and the atomic forces are calculated using the Hellman-Feynman theorem[23]

$$F^a = -\sum_{n=1}^{n_{occ}} f(\varepsilon_n, T)\langle\psi_n|\nabla_{R_a} H|\psi_n\rangle - \nabla_{R_a} E_{ion-ion}. \quad (5)$$

Only the electron-ion interaction contributes to the first term and the second term originates from ion-ion interactions.

In our density-functional method, potentials and wave functions are expressed on a uniform, orthogonal grid within a fixed domain.[24] The wave functions outside this domain are required to obey periodic boundary conditions owing to translational symmetry. The Laplacian term in Eq. (2) is calculated using a finite-difference approach

$$\frac{\partial^2\psi}{\partial x^2} = \sum_{n=-N}^{N} C_n \psi(x_i + nh, y_j, z_k) + O(h^{2N+2}), \quad (6)$$

where h is the grid spacing and N is an integer representing the order of finite differencing. For calculating V_{ion}, we employ norm-conserving ionic pseudopotentials cast in the Kleinman-Bylander form.[25,26] The contribution to the ionic potential from a single atom can be expressed as the sum of a local term and an angular-momentum-dependent nonlocal term

$$V^a_{ion}(r)\psi_n(r) = V_{loc}(r_a)\psi_n(r) + \sum_{lm} G^a_{n,lm} u_{lm}(r_a)\Delta V_l(r_a), \quad (7)$$

where u_{lm} is the atomic pseudopotential wave function, $\boldsymbol{r_a} = \boldsymbol{r} - \boldsymbol{R_a}$ is the displacement from the atomic center, $\Delta V_l = V_l - V_{loc}$ is the difference between V_l (the lth component of the ionic pseudopotential) and the local potential V_{loc}, and $G^a_{n,lm}$ is the projection coefficient given by

$$G^a_{n,lm} = \frac{1}{\langle \Delta V^a_{lm} \rangle} \int u_{lm}(r_a)\Delta V_l(r_a)\psi_n(r) d^3r,$$

$$\Delta V^a_{lm} = \int u_{lm}(r_a)\Delta V_l(r_a) u_{lm}(r_a) d^3r. \quad (8)$$

The nonlocal term in Eq. (7) is short ranged and can be directly evaluated on the real-space grid. The local potential V_{loc} has a long-ranged Coulomb tail that leads to convergence issues along periodic directions. This has been dealt with by evaluating the potential in reciprocal space.[24]

The Hartree potential V_H is obtained from the electronic charge density by solving the Poisson equation $\nabla^2 V_H = -4\pi\rho$. This is accomplished by first setting the total charge in the supercell to zero by adding a uniform neutralizing charge density. The problem is then expressed as a matrix equation using finite-differencing techniques and solved using periodic boundary conditions for V_H. The exchange-correlation potential V_{xc} is evaluated directly on the real-space grid using the local-density approximation.[27]

III. ALGORITHMS FOR SOLVING THE KOHN-SHAM EQUATION

In typical implementations of DFT, the most time-consuming part of the calculation is solving the self-consistent Kohn-Sham equation, owing to the high computational cost of matrix diagonalizations. Numerous minimization-based[28,29] and iterative eigensolver-based[30,31] methods have been used to tackle this problem. Here, we employ a method that avoids solving the eigenvalue problem explicitly by using damped Chebyshev polynomial filtered subspace iterations.[21,22] In this method, only the first iteration necessitates solving an eigenvalue problem, and this is carried out using the Chebyshev-Davidson algorithm.[31] The purpose of this step is to provide a good initial subspace (or initial approximation to the wave functions) for subsequent filtering. Compared to highly efficient restarted eigensolvers such as ARPACK (Ref. 32) and TRLAN (Thick-Restart Lanczos),[33] our filtering method utilizes less memory and consumes significantly less computational time performing operations such as orthogonalization of the subspace.

The main idea of our proposed method is to start with a good initial basis $\{\psi_n\}$ corresponding to states of the initial

Hamiltonian that have non-negligible occupation numbers based on the Fermi-Dirac distribution, and then improve adaptively the subspace by polynomial filtering. That is, at a given self-consistent step, a polynomial filter, $P_m(t)$ of order m is constructed for the current Hamiltonian H. As the basis gets updated, the polynomial will be different at each SCF step since H will change. The goal of the filter is to make the subspace spanned by $\{\hat{\psi}_n\}=P_m(H)\{\psi_n\}$ approximate the eigensubspace corresponding to the occupied states of H. There is no need to make the new subspace, $\{\hat{\psi}_n\}$ approximate the wanted subspace of H to high accuracy at intermediate steps. Instead, the filtering is designed so that the new subspace obtained at each self-consistent iteration step will progressively approximate the wanted eigenspace of the final Hamiltonian when self-consistency is reached.

This can be efficiently achieved by exploiting the Chebyshev polynomials, C_m, for the polynomials P_m. Specifically, we wish to exploit the dramatic increase in magnitude of the polynomial outside the $[-1,1]$ interval. All that is required to obtain a good filter at a given SCF step is to provide a lower bound and an upper bound for an interval in the spectrum of the current Hamiltonian H. The lower bound can be readily obtained by setting it to be slightly above the largest Ritz value computed in the previous SCF step and the upper bound can be inexpensively obtained by a very small number of Lanczos steps. Hence the main cost of this approach at each iteration is in obtaining products of the Hamiltonian to generate the filtering polynomial and multiplying the polynomial to the basis vectors.

To construct a Chebyshev polynomial "damped" on the interval $[a,b]$, we start from one that is damped between $[-1,1]$ and use an affine mapping such that

$$l(t) = \frac{2t-(a+b)}{b-a}. \tag{9}$$

The energy interval is chosen to contain the eigenspace to be dampened, i.e., energies above the eigenvalues of occupied states. The filtering operation can then be expressed as

$$\{\hat{\psi}_n\} = C_m[l(H)]\{\psi_n\}. \tag{10}$$

This computation is accomplished by exploiting the convenient three-term recurrence property of Chebyshev polynomials

$$C_0(t)=1, \quad C_1(t)=t, \quad \text{and} \quad C_{m+1}(t)=2tC_m(t)-C_{m-1}(t). \tag{11}$$

An example of a damped Chebyshev polynomial as defined by Eqs. (10) and (11) is given in Fig. 1, where we have taken the lower bound as $a=0.2$ and the upper bound as $b=2.0$. In this example, the filtering would enhance the eigenvalue components in the shaded region.

The filtering procedure for the self-consistent cycle is illustrated in Fig. 2. Unlike traditional methods, the cycle only requires one explicit diagonalization step. Instead of repeating this step again within the self-consistent loop, a filtering operation is used on the previous set of states to create a new basis in which the desired states are enhanced. After the new

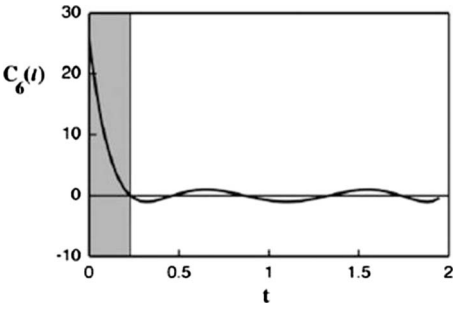

FIG. 1. Schematic example of a damped Chebyshev polynomial, C_6. The shaded area corresponds to a hypothetical eigenvalue spectrum regime that will be enhanced by the filtering operation (see the text).

basis is formed, it is orthonormalized by a Gram-Schmidt algorithm followed by a Rayleigh-Ritz procedure. The Ritz pairs are used just as a set of converged eigenpairs would be, but the filtering iteration need not check for convergence of the Ritz pairs, as the procedure has been found to be self-correcting. The quality of the initial set of vectors used in filtering only determines the number of steps needed for convergence, but the desired subspace can always be found by successive iteration. The orthonormalization and Rayleigh-Ritz steps scale approximately as the cube of the number of occupied states and as such, this method is not an "order-N" method. However, the prefactor is sufficiently small that the method is much faster than previous implementations of real-space methods. The cycle is repeated until the input and

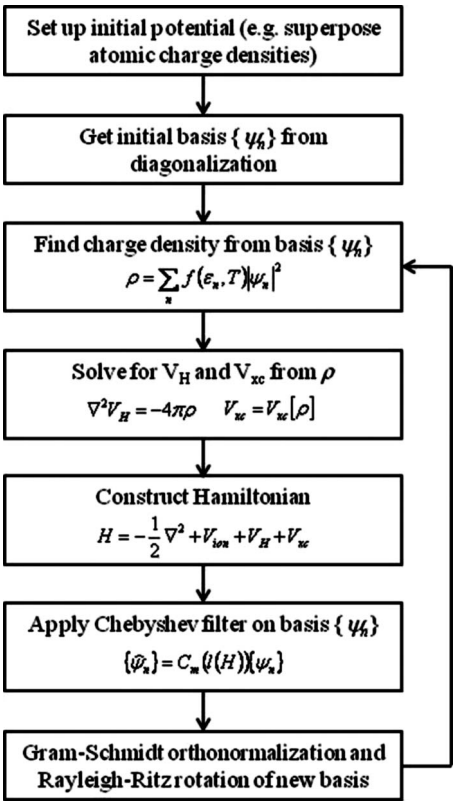

FIG. 2. Schematic of the self-consistent cycle using our Chebyshev filtering algorithm.

TABLE I. Performance of various algorithms on a 200-atom liquid-Al system. Listed are the number of iterations required to reach self-consistency and time needed for one SCF iteration.

Solver	Number of iterations	Single iteration time (s)
Chebyshev filtering	18	203.4
ARPACK (Ref. 29)	18	4879.5
TRLan (Ref. 30)	18	2116.1

output density is unchanged to within a specified tolerance.

In an MD simulation, the self-consistent cycle outlined above has to be repeated for every MD step. To speed up these calculations, we take advantage of the fact that geometry changes are small between time steps, and the converged wave function from a previous step may be used as the starting basis for filtering in the current step. Using this scheme, *diagonalization only has to be performed for the first iteration of the first* MD *step*, and this gives rise to an appreciable reduction in the computational cost of our calculations. In addition, it has been found that the upper and lower bounds of the Chebyshev filter remain relatively constant throughout the MD simulation and their determination needs to be performed only once (at the start of the simulation), leading to further savings in computational time.

It has been shown in previous tests that our code scales well up to several thousand processors on parallel computing platforms.[22] In Table I, we compare the timings for performing a single SCF iteration on a periodic supercell of 200 liquid-Al atoms (equilibrated at 1000 K) using our filtering method along with explicit diagonalization solvers TRLAN (Ref. 33) and ARPACK.[32] The numerical runs were performed on a linux-PC with a 1.8 GHz AMD Athlon chip. Although the number of iterations needed for self-consistency is the same for all three methods, the time required for one SCF iteration using the filtering algorithm is at least an order of magnitude smaller than that required using the two other solvers. Similar improvements in timings have also been observed in applications involving other systems such as Si nanocrystals, as detailed in previous publications.[21,22]

IV. MOLECULAR DYNAMICS SIMULATIONS

By combining the significant speed up enabled by our filtering algorithm and the quality of forces derived from DFT calculations, our real-space method is well suited for performing large-scale BOMD simulations. To demonstrate the accuracy of our approach, we performed an MD simulation on liquid Al at a temperature of 1000 K. The number density was set to $\rho_n = 0.0526$ Å^{-3}, a value derived from x-ray diffraction experiments.[34] Liquid Al may be considered to be a simple metal in which the core electrons are clearly distinct from the valence electrons. However, the study of this system is far from trivial, as it takes fairly sophisticated model potentials to reproduce experimental data on the static structure factor.[35] The demands on accuracy are even more stringent for dynamical properties such as the self-diffusion coefficient, where even *ab initio* MD simulations yield a fairly wide range of values.[36–39] These observations point to the need of a realistic description of liquid Al that requires quantum-mechanical simulations. Here, we compare the results of our simulation with available experimental data and previous theoretical work.

We considered a system of 300 atoms in a cubic supercell with a cell size of $L=33.76$ a.u., and the real-space grid was constructed with a spacing of $h=0.70$ a.u. Norm-conserving pseudopotentials were generated using the reference configuration $[\mathrm{Ne}]3s^2 3p^1$, and a cutoff radius of 2.4 a.u. was selected for both s and p potentials with the p potential chosen to be the local component. The local-density approximation as parametrized by Ceperley and Alder[27] was used for exchange correlation and the Γ point was sampled in the Brillouin zone. The order of the Chebyshev polynomial used in this calculation is 8.

In our simulation, ion dynamics was generated using Hellman-Feynman forces together with the Beeman algorithm[40] using a time step of 248 a.u. (6 fs). Also, an efficient extrapolation of the charge density and wave functions was implemented to provide a better initial basis for filtering at every MD step.[41] To set up the calculation, atoms were placed at random inside the cubic supercell while taking care to prevent any two atoms from getting too close to each other. The system was then coupled to a virtual heat bath via the Langevin equation of motion,[42] at a temperature far above the target value, and propagated in time to eradicate any memory of its initial configuration. The system was then cooled down gradually to the target temperature and decoupled from the heat bath where a microcanonical simulation was carried out for 670 time steps ($\sim$3.6 ps). The data collected from the microcanonical part of our simulation are then used for computing the properties of liquid Al.

A microcanonical molecular dynamics simulation constitutes a stringent test of the accuracy of calculated ionic forces because the trajectory of the system through configuration space is deterministic, and any systematic error in the force calculations will prevent conservation of the total energy. For our simulations, it can be seen in Fig. 3 that the total energy drift is negligible (less than 1 meV/ps). The kinetic and potential energies perform bounded oscillations around stable mean values and the oscillations of one cancel those of the other, resulting in very good conservation of the total energy. Also, the average temperature of our microcanonical simulation is 974 K, differing only slightly from our target temperature of 1000 K. It should be mentioned that we were able to use relatively long time steps for integrating the equations of motion because our simulation is performed on the Born-Oppenheimer surface. This is in contrast to the Car-Parrinello MD method,[43] where full self-consistency is avoided at each step by using fictitious dynamics for the electrons, but accurate integrations of the equations of motion requires time steps an order of magnitude smaller than those employed in BOMD.

The static properties of liquid Al obtained from our simulation are plotted in Figs. 4 and 5. The radial distribution function $g(r)$, shown in Fig. 4(a), agrees fairly well with results obtained from previous *ab initio* MD simulations utilizing plane-wave basis sets.[38] The slight discrepancy be-

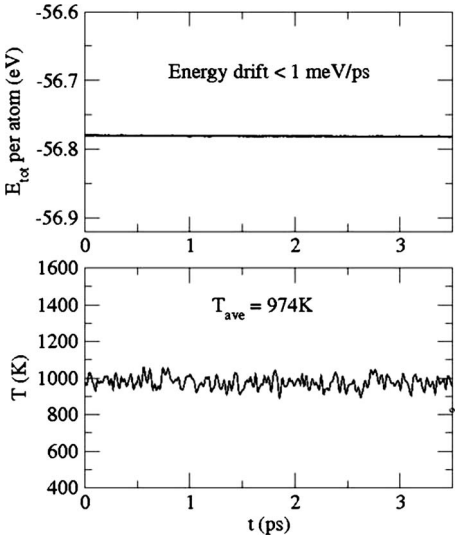

FIG. 3. Total energy per atom and simulation temperature for microcanonical MD simulation of liquid Al with target temperature of 1000 K.

tween our $g(r)$ and the earlier result may be due to the fact that the previous simulation used only 64 atoms. Integrating $g(r)$ up to its first minima R_{min} gives the value of the average coordination number as $N_C \sim 11.4$, identical to the previously simulated value of 11.4.[39] For a comparison of simulations to experiments, however, it is preferable to look at the static structure factor $S(q)$

$$S(q) = \frac{1}{N}\left\langle \sum_i \sum_j e^{-iq(R_i - R_j)} \right\rangle, \tag{12}$$

where sums over i,j are taken over atoms in the unit cell. The angular brackets denote averaging over time steps as well as $\boldsymbol{q}$ vectors of constant magnitude q. In a diffraction experiment, $S(q)$ is the quantity obtained directly from measurements whereas $g(r)$ has to be obtained from the Fourier transform of $S(q)$ and is more prone to numerical errors. In Fig. 4(b), we have plotted our calculated $S(q)$ together with data measured using x-ray diffraction at $T=1023$ K.[34] As can be seen, the agreement between our results and experiment is nearly perfect, and we correctly predict all the successive maxima and minima of the function.

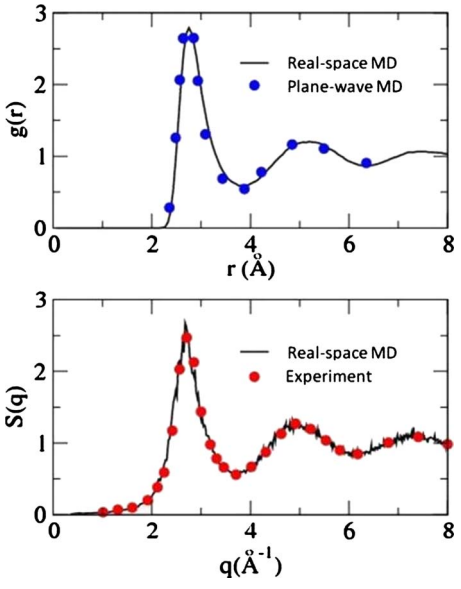

FIG. 4. (Color online) (a) Radial distribution function and (b) static structure factor obtained from MD simulations and x-ray diffraction experiments for liquid Al at ~1000 K.

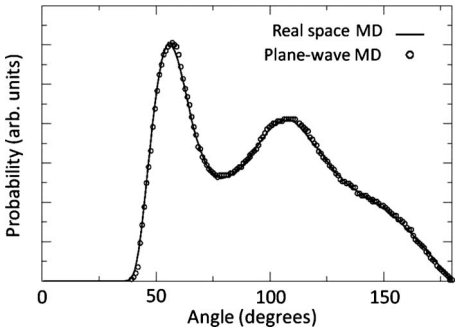

FIG. 5. Bond-angle distribution function $P(\theta, R_{min})$ calculated using our method (solid line) and *ab initio* MD simulations using a plane-wave basis (open circles) for liquid Al at 1000 K.

Another quantity we have looked at is the bond-angle distribution function $P(\theta, R_{min})$, which represents the distribution of bond angles between atoms separated by distances smaller than R_{min}. Plotted in Fig. 5 are results obtained from our work and those obtained from a prior *ab initio* MD simulation utilizing a plane-wave basis.[39] The two sets of results are in nearly perfect agreement, and the peaks at 56° and 110° are due to the presence of quasi-icosahedral short-range ordering, a feature that is known to occur in liquid Al close to the melting point.[44]

For dynamical properties of liquid Al, we examine the mean-square displacement $\langle \Delta R^2(t) \rangle$ defined by

$$\langle \Delta R^2(t) \rangle = \frac{1}{N}\left\langle \sum_i [R_i(t) - R_i(0)]^2 \right\rangle, \tag{13}$$

where angular brackets represent an average over time steps. It can be shown that the mean-square displacement obeys the relation

$$\langle \Delta R^2(t) \rangle = 6Dt + c \quad \text{for large} \quad t, \tag{14}$$

where D is the self-diffusion coefficient and c is a constant.[45] Our calculated mean-square displacement is plotted in Fig. 6(a), and it is almost perfectly linear, in agreement with Eq. (14). The self-diffusion coefficient obtained from the slope of $\langle \Delta R^2(t) \rangle$ is 0.54 Å^2/ps and it is within the range of values obtained from previous MD simulations 0.52–0.70 Å^2/ps.[36–39] Another quantity that we have calculated is the normalized velocity autocorrelation function defined by

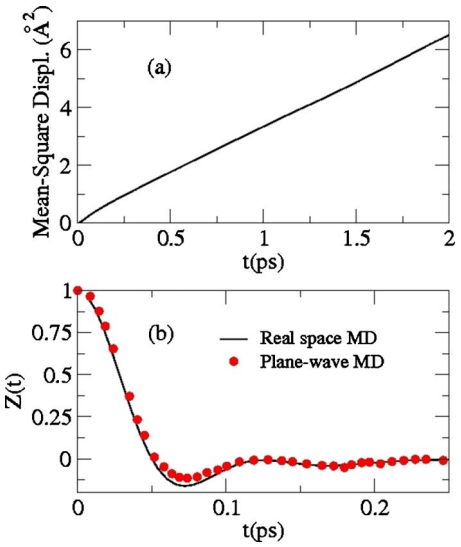

FIG. 6. (Color online) (a) Mean-square displacement and (b) velocity autocorrelation function obtained from real-space (black line) and plane-wave (red dots) *ab initio* MD simulations for liquid Al at $T=1000$ K.

$$Z(t) = \frac{\left\langle \sum_i v_i(t) \cdot v_i(0) \right\rangle}{\left\langle \sum_i v_i(0) \cdot v_i(0) \right\rangle} \tag{15}$$

and our results are plotted in Fig. 6(b) along with those obtained in previous *ab initio* plane-wave MD simulations.[39] As can be seen, the curves exhibit the "cage-effect," a feature characteristic of simple liquid metals near their melting point.[44] This is when the atoms surrounding any given atom allow the latter to move over short distances but then reflect it, causing $Z(t)$ to change sign. Also, the $Z(t)$ functions from our real space and previous plane-wave simulations show good agreement, apart from a slight discrepancy at the minima.

Similar to the mean-square displacement, the velocity autocorrelation function is also related to the self-diffusion coefficient D.[45]

$$D = \frac{k_B T}{M} \int_0^\infty Z(t) dt. \tag{16}$$

By comparing the values of D obtained from two different methods [slope of $\langle \Delta R^2(t) \rangle$ and integral of $Z(t)$], we get an indication of the quality of our simulation. As it turns out, both Eqs. (14) and (16) yield a value of 0.54 Å^2/ps for D in our simulations, and this agreement gives us confidence in the accuracy of our results.

V. CONCLUSIONS

In this paper, we have presented a method for performing MD simulations employing the real-space pseudopotential density-functional approach. Central to this approach is an algorithm that avoids explicit eigenvalue calculations in solving the Kohn-Sham equation through the use of Chebyshev polynomial filters.[21,22] Only the initial self-consistent field iteration of the first MD step requires solving an eigenvalue problem, and this is to provide a good initial subspace for filtering. In the remaining iterations, no eigensolvers are involved. Instead, Chebyshev polynomial filtering is used to refine the subspace and each filtering step is at least an order of magnitude faster than solving a corresponding eigenproblem even with the most efficient eigensolvers. Moreover, the subspace iteration reaches self-consistency within roughly the same number of steps as an eigensolver-based approach.

By taking advantage of the speed up enabled by our filtering algorithm and the quality of forces derived from DFT calculations, our real-space method is well suited for performing large-scale BOMD simulations. As an illustration of our method, we have applied it to the study of liquid Al. Our findings are in good agreement with existing experimental and *ab initio* simulation results for a variety of static and dynamic properties, verifying the accuracy of our method.

ACKNOWLEDGMENTS

This work is supported by the National Science Foundation under Grant No. DMR 09-41645 and The Welch Foundation under Grant No. F-1708. Calculations were performed at the Texas Advanced Computing Center (TACC). We thank M. M. G. Alemany for many helpful discussions.

[1] M. M. G. Alemany, R. C. Longo, L. J. Gallego, D. J. Gonzalez, L. E. Gonzalez, M. L. Tiago, and J. R. Chelikowsky, Phys. Rev. B **76**, 214203 (2007).
[2] E. Ko, M. Jain, and J. R. Chelikowsky, J. Chem. Phys. **117**, 3476 (2002).
[3] V. V. Godlevsky, J. J. Derby, and J. R. Chelikowsky, Phys. Rev. Lett. **81**, 4959 (1998).
[4] A. Pasquarello, M. S. Hybertsen, and R. Car, Nature (London) **396**, 58 (1998).
[5] Y. Duan, C. Wu, S. Chowdhury, M. C. Lee, G. Xiong, W. Zhang, R. Yang, P. Cieplak, R. Luo, T. Lee, J. Caldwell, J. Wang, and P. Kollman, J. Comput. Chem. **24**, 1999 (2003).
[6] J. Sarnthein, A. Pasquarello, and R. Car, Phys. Rev. Lett. **74**, 4682 (1995).
[7] K. A. Jackson, J. Cryst. Growth **198-199**, 1 (1999).
[8] T. Sinno, J. Cryst. Growth **303**, 5 (2007).
[9] J. R. Chelikowsky, J. J. Derby, V. V. Godlevsky, M. Jain, and J. Y. Raty, J. Phys.: Condens. Matter **13**, R817 (2001).
[10] W. Kohn and L. J. Sham, Phys. Rev. **140**, A1133 (1965).
[11] D. R. Hamann, M. Schluter, and C. Chiang, Phys. Rev. Lett. **43**, 1494 (1979).
[12] A. Canning and D. Raczkowski, Comput. Phys. Commun. **169**, 449 (2005).
[13] G. Kresse and J. Furthmüller, Phys. Rev. B **54**, 11169 (1996).

[14] J. M. Soler, E. Artacho, J. D. Gale, A. Garcia, J. Junquera, P. Ordejon, and D. Sanchez-Portal, J. Phys.: Condens. Matter **14**, 2745 (2002).
[15] J. R. Chelikowsky, N. Troullier, and Y. Saad, Phys. Rev. Lett. **72**, 1240 (1994).
[16] B. Delley, J. Chem. Phys. **94**, 7245 (1991).
[17] M. L. Cohen, M. Schluter, J. R. Chelikowsky, and S. G. Louie, Phys. Rev. B **12**, 5575 (1975).
[18] R. Schmid, J. Comput. Chem. **25**, 799 (2004).
[19] R. Schmid, M. Tafipolski, P. H. König, and H. Köstler, Phys. Status Solidi B **243**, 1001 (2006).
[20] T. D. Kühne, M. Krack, F. R. Mohamed, and M. Parrinello, Phys. Rev. Lett. **98**, 066401 (2007).
[21] Y. Zhou, Y. Saad, M. L. Tiago, and J. R. Chelikowsky, Phys. Rev. E **74**, 066704 (2006).
[22] Y. Zhou, Y. Saad, M. L. Tiago, and J. R. Chelikowsky, J. Comput. Phys. **219**, 172 (2006).
[23] R. P. Feynman, Phys. Rev. **56**, 340 (1939).
[24] M. M. G. Alemany, M. Jain, L. Kronik, and J. R. Chelikowsky, Phys. Rev. B **69**, 075101 (2004).
[25] N. Troullier and J. L. Martins, Phys. Rev. B **43**, 1993 (1991).
[26] L. Kleinman and D. M. Bylander, Phys. Rev. Lett. **48**, 1425 (1982).
[27] D. M. Ceperley and B. J. Alder, Phys. Rev. Lett. **45**, 566 (1980).
[28] M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos, Rev. Mod. Phys. **64**, 1045 (1992).
[29] P. Pulay, Chem. Phys. Lett. **73**, 393 (1980).
[30] J. L. Martins and M. L. Cohen, Phys. Rev. B **37**, 6134 (1988).
[31] Y. Saad, A. Stathopoulos, J. Chelikowsky, K. Wu, and S. Ogut, BIT **36**, 563 (1996).
[32] D. C. Sorensen, SIAM J. Matrix Anal. Appl. **13**, 357 (1992).
[33] K. Wu, A. Canning, H. D. Simon, and L. W. Wang, J. Comput. Phys. **154**, 156 (1999).
[34] Y. Waseda, *The Structure of Non-Crystalline Materials* (McGraw-Hill, New York, 1980).
[35] G. Jacucci, R. Taylor, A. Tenenbaum, and N. van Doan, J. Phys. F: Met. Phys. **11**, 793 (1981).
[36] D. Alfè and M. J. Gillan, Phys. Rev. Lett. **81**, 5161 (1998).
[37] D. J. González, L. E. González, J. M. López, and M. J. Stott, Phys. Rev. B **65**, 184201 (2002).
[38] P. E. Blöchl and M. Parrinello, Phys. Rev. B **45**, 9413 (1992).
[39] M. M. G. Alemany, L. J. Gallego, and D. J. González, Phys. Rev. B **70**, 134206 (2004).
[40] D. Beeman, J. Comput. Phys. **20**, 130 (1976).
[41] D. Alfè, Comput. Phys. Commun. **118**, 31 (1999).
[42] N. Binggeli, J. L. Martins, and J. R. Chelikowsky, Phys. Rev. Lett. **68**, 2956 (1992).
[43] R. Car and M. Parrinello, Phys. Rev. Lett. **55**, 2471 (1985).
[44] U. Balucani and M. Zoppi, *Dynamics of the Liquid State* (Clarendon, Oxford, 1994).
[45] J. P. Hansen and I. R. McDonald, *Theory of Simple Liquids* (Academic, London, 1986).