Time-Reversible Multiple Time Scale ab Initio Molecular Dynamics

Douglas A. Gibson and Emily A. Carter*

Department of Chemistry and Biochemistry, University of California, Los Angeles, 405 Hilgard Avenue, Los Angeles, California 90024-1569

Received: September 15, 1993; In Final Form: October 28, 1993®

We have developed previously ab initio molecular dynamics (AIMD) algorithms utilizing local molecular wave functions with atom-centered bases. The expense of such methods led us to explore means to employ multiple time step schemes to reduce costs. Herein we report our second implementation of such a scheme. In particular, our previous multiple time step AIMD method has been improved substantially by the implementation of the time-reversible reference system propagator algorithm (RESPA) of Berne and co-workers. The time-reversible version of RESPA treats multiple time steps rigorously without introducing the limitations inherent in our previous non-time-reversible AIMD RESPA implementation. We are thus simultaneously able to obtain substantial additional savings in computer time and to maintain high accuracy, as shown in test calculations on a sodium cluster.

Introduction

Ab initio molecular dynamics (AIMD) involves classical molecular dynamics of atomic nuclei to which forces derived from a quantum mechanical electronic wave function are applied. This electronic wave function then follows the motion of the nuclei. In our case, a Car–Parrinello type algorithm[1] is used to propagate the electronic degrees of freedom; a fictitious mass is assigned to the coefficients in the wave function, and they are then treated as classical degrees of freedom. Alternative approaches in the same vein include time-dependent Hartree–Fock-based methods.[2] While AIMD is quite computationally intensive, it has the advantage that a priori knowledge of the entire potential surface is not required; the energy gradient is calculated only at points on the actual trajectory of interest.

Since we have developed the AIMD method utilizing Hartree–Fock (HF)[3,4] and generalized valence bond (GVB)[5,6] wave functions expanded in an atom-centered Gaussian basis set, the integrals and gradient must be completely recalculated every time the nuclei move. Thus, the most time-consuming portions of our molecular quantum mechanical-based version of AIMD are the calculation of the two-electron integrals over the basis functions, which account for approximately 20% of the total time required for a single time step, and the analytic energy gradient, which uses almost all of the remaining 80% of that time. The full analytic energy gradient is necessary because approximations such as using the Hellmann–Feynman forces are not sufficiently accurate; for example, in some cases, even the sign of the Hellmann–Feynman forces is incorrect. In the traditional Car–Parrinello (CP) method,[1] density functional theory (DFT) wave functions expanded in a plane wave basis set are used, which do not depend explicitly on the positions of the nuclei. As a result, recalculations of the integrals evaluated over plane waves are not necessary. In addition, cheaper Hellmann–Feynman forces are sufficiently accurate for plane wave bases. Thus, the traditional CP DFT method is significantly faster than our GVB AIMD method. However, DFT cannot treat excited states rigorously correctly, and it remains unclear that the DFT method is appropriate for nonequilibrium configurations of finite systems, despite recent impressive improvements utilizing gradient-corrected functionals.[7] We therefore believe pursuing molecular wave function-based AIMD methods is still fruitful. Given this objective, it is especially critical to find ways to reduce the amount of computer time used in any AIMD calculation with an atom-centered basis set in order to begin to make it competitive with traditional DFT-based methods. This is the aim of our current work.

One possible way to speed up AIMD is to take advantage of the multiple time scales inherent in the system being studied. That has been done, for example, by Micha and co-workers in the time-dependent Hartree–Fock formalism.[8] The electronic degrees of freedom are known to change much more rapidly than the nuclear degrees of freedom; this is the basis of the Born–Oppenheimer approximation. If a single dynamical time step is used, it must be small enough that the electronic wave function (the fast degrees of freedom) changes by only a small amount each time step in order to retain accuracy while integrating the equations of motion. However, if the nuclei are propagated with the same time step, much time is wasted in recalculating forces on the nuclei, because those forces are changing much more slowly. Since most of the time used in an AIMD simulation is required for the calculation of the forces on the nuclei, we can save a large amount of computer time by calculating these forces less often.

We originally did this in a relatively crude fashion, by simply calculating the forces on the nuclei every 5–10 time steps and either keeping the forces constant in between calculations or using a linear extrapolation scheme.[4,6] While this is effective in reducing the amount of computer time expended, it is an ad hoc method and has no rigorous basis in theory. A rigorous separation of time scales (such as that provided by the RESPA integrator[9]) is preferable. The RESPA (reference system propagator algorithm) integrator integrates the equations of motion for a reference system that holds some degrees of freedom fixed. Those degrees of freedom are then propagated afterward, along with propagation of a correction term that accounts for the artificial separation of the degrees of freedom. In our case, the natural separation is to adopt a reference system where we hold the nuclei fixed while propagating the electronic degrees of freedom for a number of short time steps. Then the nuclei and the analytically-derived correction to the wave function are propagated in one long time step. This correction thus accounts for the fact that the nuclei should have been moving while the electronic wave function was being propagated. We have recently reported the implementation of the RESPA integrator in AIMD simulations.[10] Because of the great difference between the time scales of the nuclear and electronic degrees of freedom present in our AIMD simulations, we observed a speedup factor of about 7 in test simulations of a cluster of four Na atoms, with better accuracy than was obtainable

* To whom correspondence should be addressed.
® Abstract published in *Advance ACS Abstracts*, December 1, 1993.

0022-3654/93/2097-13429$04.00/0

by simply skipping force recalculations. However, no systematic means of obtaining even better performance was available, because it turns out that the form of the analytic correction to the wave function is akin to an exponential growth expression, and it did indeed grow exponentially if the long time step was increased above a certain limit (approximately a 10:1 ratio between the long and short time steps).

In this paper we present an alternative method that eliminates the previous limitations. We have chosen to implement the time-reversible form of the RESPA integrator[11] that has been recently developed by Berne and co-workers for Lennard-Jones particles. This method has proven to be superior for AIMD to the first RESPA integrator, because the correction to the wave function, which diverged in the previous nonreversible RESPA algorithm when the long time step became too large, is present only implicitly in time-reversible RESPA and does not have an equation of motion of its own. Therefore, exponential divergence of the correction should not be a problem in time-reversible RESPA AIMD. The next section describes the simple Verlet adaptation of time-reversible RESPA to GVB AIMD. In sections 3 and 4, we present our initial test results and conclusions.

Method

The details of our AIMD implementation have been published previously.[4,6,10] The system is described by the nuclear coordinates $\{R_I\}$, which describe the geometry of the molecule, the SCF coefficients $\{c_{\mu i}\}$, and the GVB-CI coefficients, $\{\alpha_k\}$, which, in conjunction with the basis set, describe the electronic wave function. The evolution of the system is determined by the equations of motion:[6]

$$M_I\ddot{R}_I = -\frac{\partial E}{\partial R_I} \tag{1}$$

$$m_{\mathrm{SCF}}\ddot{c}_{\mu i} = -\frac{\partial E}{\partial c_{\mu i}} - \sum_{\nu j}\lambda_{ij}c_{\nu j}S_{\mu\nu} \tag{2}$$

$$m_{\mathrm{GVB}}\ddot{\alpha}_k = -\frac{\partial E}{\partial \alpha_k} \tag{3}$$

where M_I are the atomic masses, E is the total potential energy, $S_{\mu\nu}$ is the overlap matrix element between basis functions μ and ν, λ_{ij} are the associated Lagrange multipliers needed to maintain orthonormality between the orbitals, m_{SCF} is the fictitious mass of the SCF coefficients, and m_{GVB} is the fictitious mass of the GVB-CI coefficients. More detailed expressions for each of these terms may be found in ref 6.

We can now integrate these equations of motion numerically with a long time step Δt for the nuclei and a short time step δt for the SCF coefficients, using a simple Verlet version of the time-reversible RESPA integrator. The idea behind the time-reversible RESPA scheme is fairly simple.[11] One integrates the fast degrees of freedom for half of the long time interval appropriate for the slow degrees of freedom. Then the slow degrees of freedom are propagated for a full long time step, where the forces on the slow degrees of freedom are evaluated subject to the half-interval values of the fast degrees of freedom. Then the fast degrees of freedom are integrated to the end of the long time interval subject to the updated coordinates of the slow degrees of freedom. In this way, both the slow and fast degrees of freedom see average values for the complementary set of degrees of freedom. This averaging implicitly accounts for the correction term that was explicitly included in the previous RESPA algorithm. We expect the implicit form of the correction should lead to greater numerical stability. Furthermore, it is obvious from the structure of the algorithm that it is time-reversible. We now outline its implementation for AIMD.

We first calculate the forces $F_c\{c(t),R(t)\}$ on the SCF and GVB-CI coefficients c at time t_0. We then propagate the SCF and GVB-CI coefficients for $n/2$ short time steps δt, where n is equal to $\Delta t/\delta t$, recomputing the force F_c before each step, and utilizing the usual Verlet propagation algorithm:

$$c(t+\delta t) = 2c(t) - c(t-\delta t) + \frac{(\delta t)^2}{m_{\mathrm{SCF/GVB}}}F_c\{c(t),R(t_0)\} \tag{4}$$

where t_0 is the time at the beginning of the long time step. This brings the SCF and GVB-CI coefficients c to time $t_0 + \Delta t/2$ while leaving the nuclear positions R at time t_0. We then compute the force $F_R\{c(t+\Delta t/2),R(t_0)\}$ on the nuclei and move the nuclei in one large time step Δt:

$$R(t+\Delta t) = 2R(t) - R(t-\Delta t) + \frac{(\Delta t)^2}{M}F_R\{c(t+\Delta t/2),R(t_0)\} \tag{5}$$

The nuclei have then been advanced to time $t_0 + \Delta t$ while the SCF and GVB-CI coefficients are left at time $t_0 + \Delta t/2$. We then compute new forces F_c on the SCF and GVB-CI coefficients and integrate them for $n/2$ more short time steps δt using the equation

$$c(t+\delta t) = 2c(t) - c(t-\delta t) + \frac{(\delta t)^2}{m_{\mathrm{SCF/GVB}}}F_c\{c(t),R(t_0+\Delta t)\} \tag{6}$$

at the end of which the entire system will be at time $t_0 + \Delta t$.

We emphasize again the fact that half of the short time steps are performed with the nuclei at time t_0 and half are performed with the nuclei at time $t_0 + \Delta t$. The end result of this is that the explicit correction found in the nonreversible version of RESPA disappears, because the nuclear motion is now treated directly in the propagation of the SCF and GVB-CI coefficients instead of being assumed negligible during that portion of the propagation. Note also that when the nuclei are moved, the force is derived from the wave function at time $t_0 + \Delta t/2$, and not at time t_0, giving a better (averaged) representation of the wave function over the entire long time step.

Since we only need to compute the two-electron integrals over basis functions and the forces on the nuclei once every long time step Δt, and these operations are by far the most expensive operations in an AIMD simulation, we expect significant savings in computer time. Furthermore, since the reversible RESPA integrator was derived by Berne and co-workers from the Trotter factorization of the Liouville propagator, which has error terms of order t^3, the same order as the error terms in the Verlet integrator itself, there should be no appreciable loss of accuracy due to the use of reversible RESPA instead of a simple, single time step Verlet method. This factorization also results in a rigorously time-reversible integrator, unlike the previous version of RESPA.

Test Results

We have applied this simple Verlet reversible RESPA algorithm to the propagation of a cluster of four sodium atoms in the ground (singlet) electronic state with full forces at the GVB-PP (2/4) level of theory. The basis set and effective core potential are the same as in our previous studies.[4,6,10] The initial geometry for all cases is a planar rectangle with sides of length 4.4 and 3.6 Å, with slight distortion to remove symmetry. The nuclei are initially at rest. All trajectories are approximately 350 fs in length (2880 small time steps δt).

We assigned the SCF and GVB-CI coefficients fictitious masses of 3000 au, approximately 14 times smaller than the mass of a sodium nucleus. For comparison purposes, a non-RESPA control trajectory with a single time step of 5 au (~0.12 fs) was run, where the full forces and all two-electron integrals were computed each time step. The electronic wave function was reconverged approximately every 15 fs to ensure that the system remained

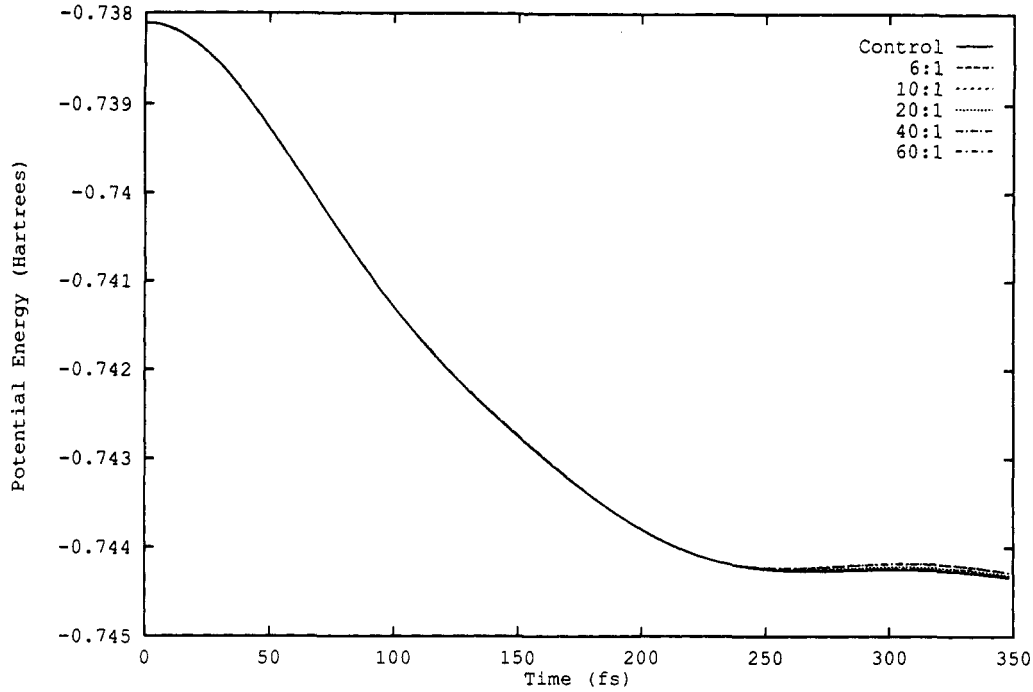

Figure 1. Potential energy in singlet Na_4 propagation. This plot shows the behavior of the potential energy (electronic energy plus nuclear repulsion) as felt by the nuclei at each long time step in each of six trajectories. In each figure, the curve labeled "Control" is the non-RESPA control ("exact") trajectory, and the curves labeled, for instance, 6:1, are the time-reversible RESPA trajectories with the ratio between long and short time steps indicated.

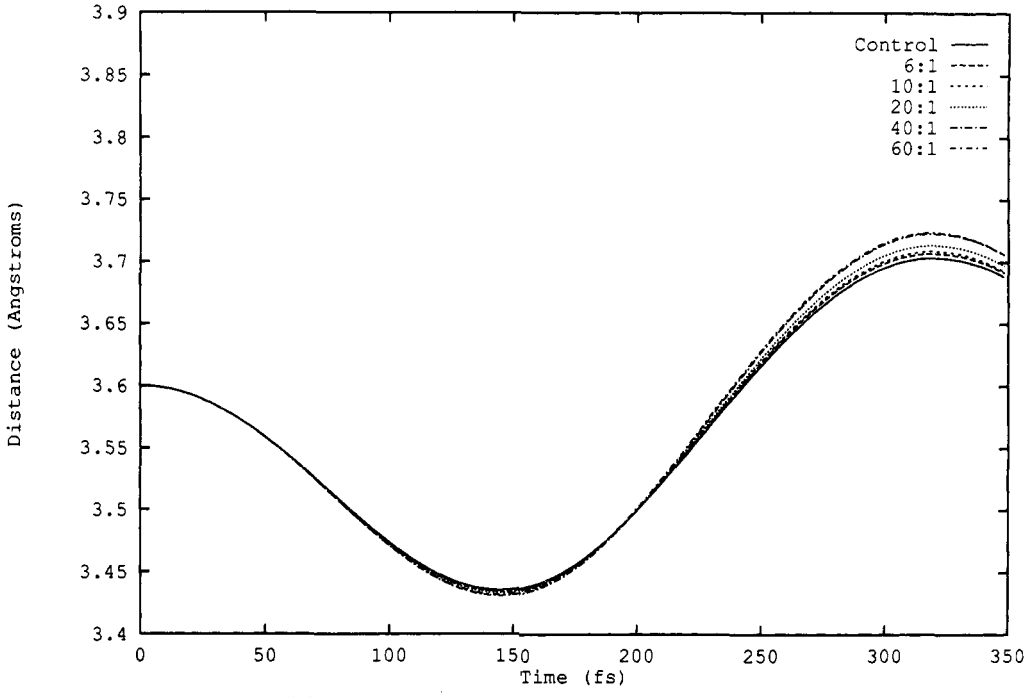

Figure 2. Distance between a pair of nearest neighbors in the singlet Na_4 propagation. This plot shows the distance between a selected pair of Na atoms as a function of time in each trajectory, illustrating the similarity of the trajectories among the test cases.

near the Born–Oppenheimer potential surface. All trajectories kept the short time step constant at 5 au, while increasing the long time step by the appropriate factor; e.g., a time step ratio of 6:1 means that the long time step is 30 au. The time-reversible RESPA test trajectories utilized time step ratios of 6:1, 10:1, 20:1, 40:1, and 60:1. The figures discussed below illustrate the behavior of various properties of the cluster as a function of the time step ratio.

In Figure 1 we can see that the deviation in the potential energy as a function of time is quite small as the time step ratio increases; only in the last 100 fs do the curves begin to separate appreciably, and even then they remain quite close to each other (within 10^{-4} hartrees). The nuclear trajectories are clearly very close to the control (which is in turn very near the true trajectory) in all cases. The deviation is larger for larger time step ratios.

Figure 2 shows the distance between two atoms throughout the simulations. The deviation from the control trajectory is so small for all test trajectories that the curves are nearly indistinguishable for the first 200 fs. The largest deviations (0.025 Å) occur near the end, although the trajectories appear to be coalescing at later times.

The time evolution of the nuclear kinetic energy is plotted in Figure 3. Most of the deviation from the control trajectory occurs in the latter half of the trajectory, where the kinetic energy is larger. This is no surprise; any integrator of this sort will become less accurate as the nuclei begin to move more rapidly. We also

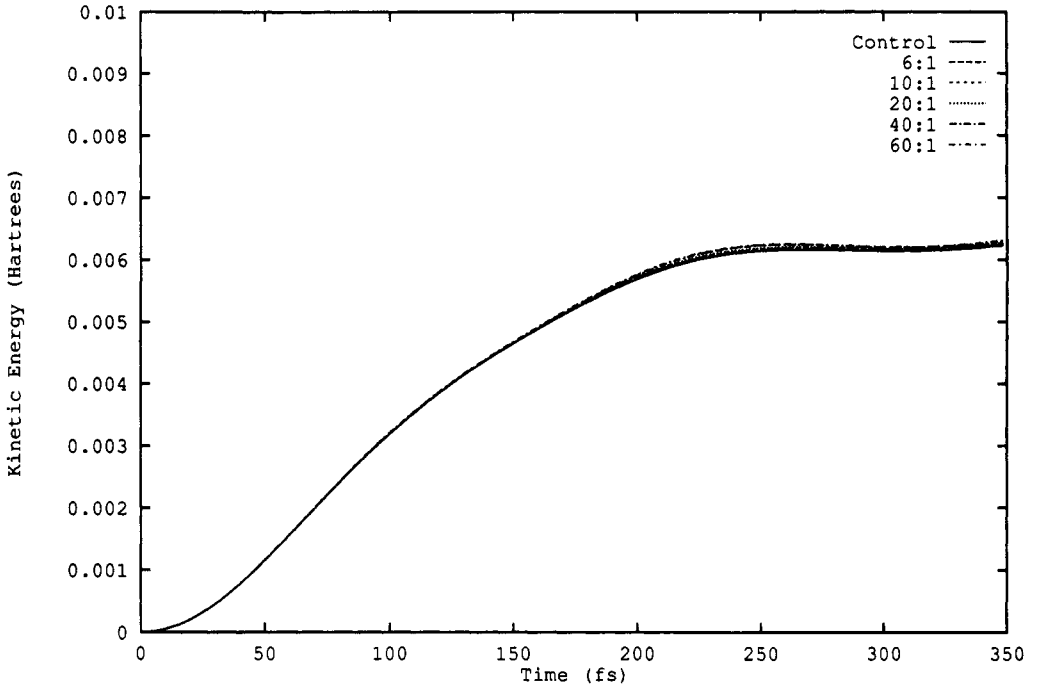

Figure 3. Total kinetic energy of the nuclei in the singlet Na_4 propagation. This plot shows how the kinetic energy of the nuclei changes with time in each trajectory.

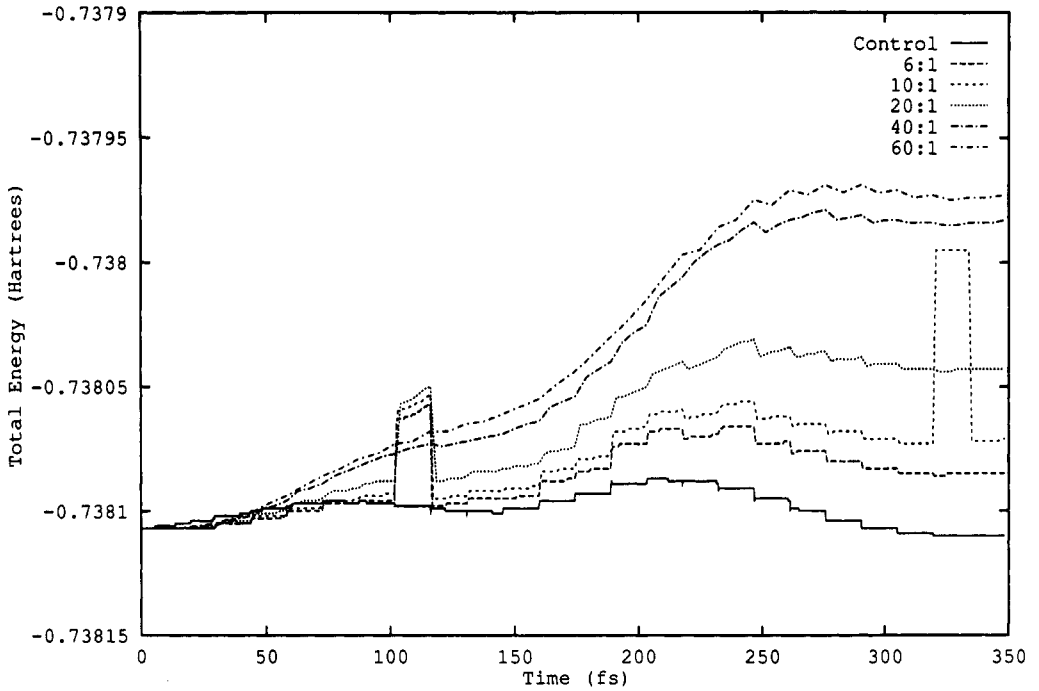

Figure 4. Total energy of the system in the singlet Na_4 propagation. This plot compares the energy conservation of the test trajectories. The abrupt changes observed are due to nonconservation of the kinetic energy of the SCF and GVB coefficients when the wave function is reconverged (see Figures 5 and 6).

observe, as expected, that the deviation is larger for larger time step ratios. However, the deviation from the control is still less than 10^{-4} hartrees.

Figure 4 displays the time evolution of the total energy of the system. One measure of the accuracy of a molecular dynamics simulation is conservation of energy. In the worst case, 60:1, we see that the energy deviation grows as large as 1.4×10^{-4} hartrees. The magnitude of this deviation is approximately proportional to the time step ratio except for an anomaly in the 60:1 trajectory, which for reasons unknown has an energy deviation which is only slightly larger than that of the 40:1 trajectory.

Figure 5 shows the kinetic energy of the SCF coefficients versus time. Since this is a purely fictitious energy, we wish it to be as small as possible so that energy is not drained from the real portion of the system. Here it is only the magnitude of the kinetic energy of the SCF coefficients and not the deviation from the control that is important, because this is a purely fictitious quantity. The most important piece of information contained in Figure 5 is that, in all cases, this fictitious energy is 2 orders of magnitude smaller than the real kinetic energies with which we are concerned. If we reduce the fictitious mass to 300 au as in our original studies (which could not be done with normal RESPA because of the behavior of the correction equation but can be done with reversible RESPA), the observed fictitious kinetic energies will become even smaller.

Figure 6 shows the kinetic energy of the GVB-CI coefficients. As for the kinetic energy of the SCF coefficients, we wish this to be small compared to the real kinetic energy in the system. We

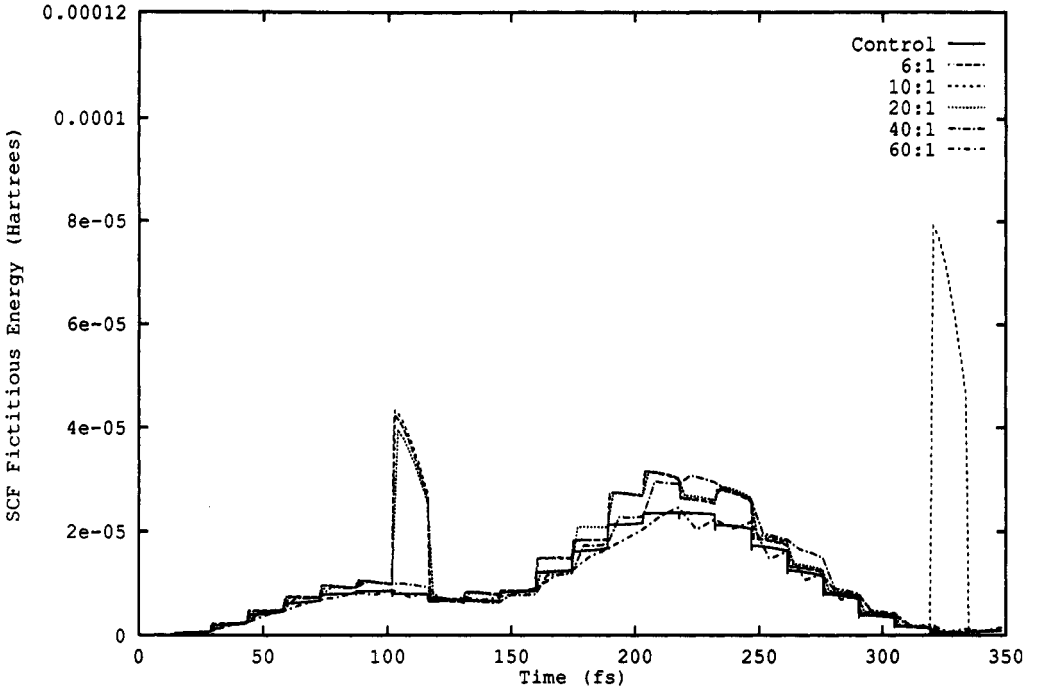

Figure 5. Fictitious kinetic energy of the SCF coefficients in the singlet Na_4 propagation. This plot shows the behavior of the kinetic energy of the SCF coefficients in each trajectory. This energy is a consequence of the treatment of the electronic wave function as classical degrees of freedom during propagation.

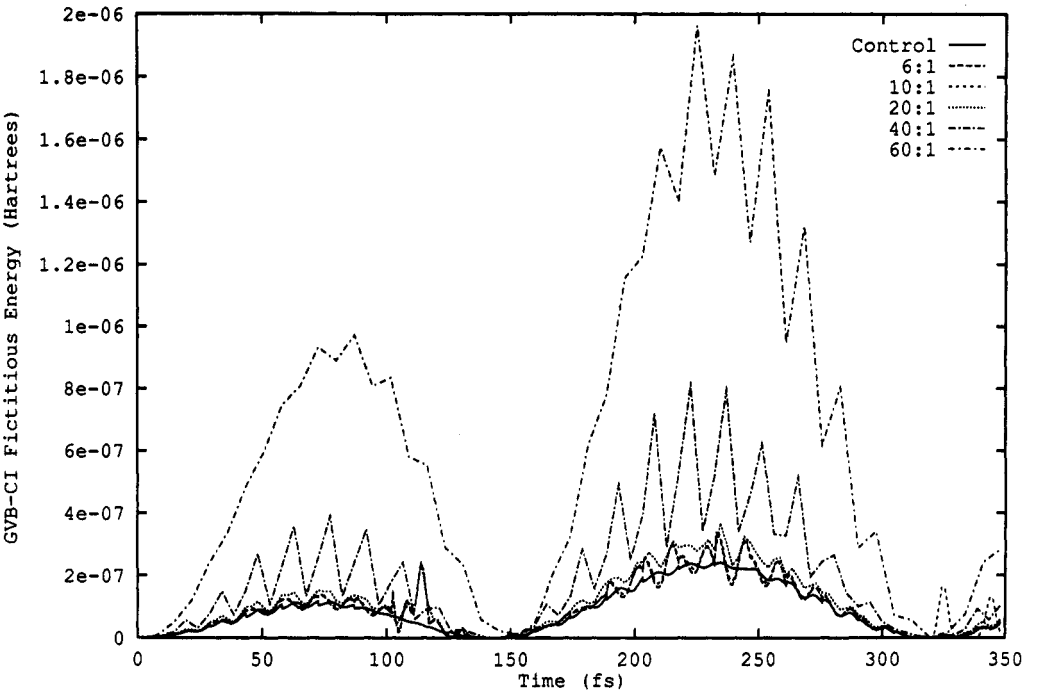

Figure 6. Fictitious kinetic energy of the GVB-CI coefficients in the singlet Na_4 propagation. This plot shows the behavior of the kinetic energy of the GVB-CI coefficients in each trajectory. Again, this energy is a consequence of the treatment of the electronic wave function as classical degrees of freedom.

see from this figure that the kinetic energy of the GVB-CI coefficients is 3–4 (depending on the trajectory in question) orders of magnitude smaller than the nuclear kinetic energy, exactly as desired.

We have observed only small deviations from the non-RESPA control trajectory in various properties as a function of time. For example, we note that the total energy deviation observed in the worst (60:1) case, which peaks at about 1.4×10^{-4} hartrees, is ~5000 times smaller than the total energy (0.02% error). This minor error in energy conservation is a small price to pay for the time savings in computation. In particular, the speed increases we observed for the 6:1, 10:1, 20:1, 40:1, and 60:1 time-reversible RESPA runs relative to the control ("exact" dynamics) run were factors of 5.0, 7.6, 12, 18, and 21, respectively. It appears that at sufficiently large time step ratios the computation of the forces on the nuclei ceases to be the primary use of computer time, since we see only a small increase in speed when the time step ratio is increased from 40:1 to 60:1.

Summary and Conclusions

We have presented a simple Verlet time-reversible RESPA algorithm for ab initio molecular dynamics (AIMD). Exploitation of the natural separation of time scales between nuclear and electronic motion and recognition that the major bottleneck in AIMD utilizing atom-centered bases is the calculation of the

forces on the nuclei has allowed us to minimize these calculations by using the time-reversible RESPA multiple time step integration scheme. Our test calculations show that a considerable increase in speed (up to a factor of $\sim$20) can be gained at very little cost in accuracy, with the total energy easily conserved to within 10^{-4} hartrees and the nuclear trajectory itself showing little deviation from "exact" Verlet dynamics. Furthermore, upon extension of the most poorly behaved (60:1) trajectory to 700 fs, we find that the total energy deviation observed is increasing in an approximately linear fashion (disregarding relatively minor oscillations) as a function of time. In fact, the total energy is still conserved to within 0.26 mhartree over this time period, well within chemical accuracy.

By the time one reaches a time step ratio of 60:1, we see less than a 20% improvement in speed over a time step ratio of 40:1, with similar errors incurred, suggesting that we have reached nearly the limit of the maximum efficiency of this algorithm. The time step ratio best for any particular system will depend on the accuracy required and the nature of the system's potential energy surface and the temperature imposed. Future work will be to develop a dynamically changing multiple time step procedure that imposes the same level of accuracy in the total energy conservation throughout the dynamics, while optimizing the time step ratio on the fly as the system evolves on the potential surface and as the system is heated or cooled.

Acknowledgment. This work was supported by the Office of Naval Research. E.A.C. also acknowledges support from the National Science Foundation, the Camille and Henry Dreyfus Foundation, and the Alfred P. Sloan Foundation, through their Presidential Young Investigator, Teacher–Scholar, and Research Fellow Award programs. D.A.G. thanks the National Science Foundation for a predoctoral fellowship.

References and Notes

(1) Car, R.; Parrinello, M. *Phys. Rev. Lett.* **1985**, *55*, 2471.

(2) (a) Runge, K.; Micha, D. A.; Feng, E. Q. *Int. J. Quantum Chem. Symp.* **1990**, *24*, 781. (b) Feng, E. Q.; Micha, D. A.; Runge, K. *Int. J. Quantum Chem.* **1991**, *40*, 545.

(3) Roothaan, C. C. J. *Rev. Mod. Phys.* **1960**, *32*, 179.

(4) Hartke, B.; Carter, E. A. *Chem. Phys. Lett.* **1992**, *189*, 358.

(5) Bobrowicz, F. W.; Goddard, W. A. III In *Methods of Electronic Structure Theory*; Schaefer, H. F., Ed.; Plenum: New York, 1977; p 79.

(6) Hartke, B.; Carter, E. A. *J. Chem. Phys.* **1992**, *97*, 6569.

(7) (a) Perdew, J. P.; Wang, Y. *Phys. Rev. B* **1986**, *33*, 8800. (b) Lee, C.; Yang, W.; Parr, R. G. *Phys. Rev. B* **1988**, *37*, 785. (c) Becke, A. D. *Phys. Rev. A* **1988**, *38*, 3098. (d) Perdew, J. P. *Proceedings of the 21st Annual International Symposium on "Electronic Structure of Solids"*, Dresden, 11–15 March 1991. (e) Becke, A. D. *J. Chem. Phys.* **1993**, *98*, 1372. (f) Becke, A. D. *J. Chem. Phys.* **1993**, *98*, 5648.

(8) (a) Micha, D. A.; Runge, K. In *Time-Dependent Quantum Molecular Dynamics*; Broeckhove, J., Lathouwers, L., Eds.; Plenum: New York, 1992; p 247. (b) Micha, D. A.; Runge, K. *Abstracts of Papers*, ACS National Meeting in Washington, DC., 1992; No. 204.

(9) Tuckerman, M. E.; Berne, B. J.; Rossi, A. *J. Chem. Phys.* **1991**, *94*, 1465.

(10) Hartke, B.; Gibson, D. A.; Carter, E. A. *Int. J. Quantum Chem.* **1993**, *45*, 59.

(11) Tuckerman, M. E.; Berne, B. J.; Martyna, G. J. *J. Chem. Phys.* **1992**, *97*, 1990.