INTERNATIONAL JOURNAL OF QUANTUM CHEMISTRY, VOL. 45, 59–70 (1993)

Multiple Time Scale Hartree–Fock Molecular Dynamics

BERND HARTKE, 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

Abstract

This is the first application of a rigorous, established multiple time-step method to ab initio molecular dynamics. The resulting algorithm is conceptually simple and easy to implement, but very effective. It translates the large mass differences present in ab initio molecular dynamics into substantial savings in computer time while retaining high accuracy. This transforms ab initio molecular dynamics from a desirable but prohibitively expensive possibility into a viable method, at least for short-time phenomena in small systems or for otherwise inaccessibly complicated potential energy surfaces. © 1993 John Wiley & Sons, Inc.

1. Introduction

For the purpose of this paper, we define ab initio molecular dynamics (AIMD) as classical molecular dynamics of atomic nuclei governed by first-principles forces derived from a molecular quantum mechanical wave function for the electrons, which, in turn, is made to follow the movements of the nuclei. This method is computationally very demanding, but also extremely attractive since the potential energy surface for the nuclear motion is calculated only at the points actually visited by the trajectory. In recent years, several approaches have been developed, differing mostly in the way the electronic wave function is treated. For example, Landman and co-workers [1] propagate the wave function quantum mechanically using density functional theory (DFT); Micha and co-workers [2] employ time-dependent Hartree–Fock (TDHF) theory; and there are even attempts to join nuclear and electronic motion without invoking the Born–Oppenheimer approximation [3]. In this paper, however, our focus is on the group of algorithms often termed Car–Parrinello propagations [4–6]. Their main idea is to treat the time evolution of the wave function within the framework of the Born–Oppenheimer separation by a classical mechanical propagation of wave-function parameters on the same footing as the nuclear propagation [7–9]. These algorithms have seen many applications within DFT [10], semiempirical methods [11], tight-binding models [12], and Hartree–Fock (HF) theory, both without [13] and with nuclear dynamics [14, 15] for geometry optimization. We have recently introduced implementations of AIMD for free dynamics and geometry optimization, both on the HF [16] and generalized valence bond (GVB) [17] levels of theory.

 CCC 0020-7608/93/010059-12

In spite of these promising developments, AIMD computations are still very time-consuming. But there is one feature of the method that allows for straightforward improvement of performance: To propagate the wave-function parameters, fictitious masses have to be assigned to them. These masses are typically orders of magnitude smaller than the masses of the nuclei. The fast movement of the ultralight wave-function parameters [9] makes a very small time step necessary, while the movement of the nuclei could well be described using a much larger time step. We present here a multiple time-step algorithm that exploits the fact that nuclear motion is so much slower than electronic motion, leading to considerable CPU time savings. Some workers have utilized different time steps for these different systems in related theoretical approaches (e.g., Landman and co-workers [1] and Micha and co-workers [2]), but to our knowledge, this idea has not yet been systematically used in AIMD methods of the Car–Parrinello type.

Previous tests of our AIMD implementation have shown [16, 17] that while retaining a very small time step dictated by the fictitious masses one can get away with recomputing basis function integrals and forces on the nuclei only every 5–10th step. This gains a factor of 7 in overall performance, but it is a compromise between making the computations faster and decreasing their accuracy; furthermore, it is a purely empirical approach without further theoretical foundation.

Yet, the problem of disparate time scales has already been solved in the domain of traditional molecular dynamics by several multiple time-step methods: Some authors [18–20] divide many-particle systems into "primary" and "secondary" neighbor regions and calculate forces arising from the secondary region less frequently than those from the primary region. These methods are able to gain a factor of 3–10 in speedup depending on the system and technical details about the neighbor list. Obviously, such a method is not advantageous for a system of a few atoms (most of them will be "primary" anyway) and a lot of wave-function parameters (which have no clear-cut definition of "neighborhood"). Another group of methods [21–23] focuses directly on the issue of disparate masses. These are directly applicable to the present problem, even for small systems. Of these latter methods, we chose the RESPA (reference system propagator algorithm) procedure of Berne et al. [23], since it includes an inherent self-correction: Not only are the systems of heavy masses (here, the nuclei) and light masses (here, the wave-function coefficients) propagated quasi-separately, but also is the difference between them. This difference is then used to improve upon the quasi-separate propagation, without incurring significant additional computational cost. The method of Swindoll and Haile [22] is also more accurate than that of Teleman and Jönsson [21], but it needs higher-order spatial derivatives of the potential. This has to be avoided here, since most of the time is already spent calculating first derivatives.

RESPA was originally formulated based on the velocity Verlet algorithm [23], but later modified and extended to a wider class of propagation methods [24]. Since our AIMD implementation is based on the simple Verlet procedure, we decided to adapt the original RESPA formulation to our needs. The next section describes this modification for HF–AIMD; the extension to GVB–AIMD is obvious and straightforward. In Sections 3 and 4, we present our first test calculations and conclusions.

2. Computational Method

Technical details of our AIMD implementation have already been published elsewhere [16, 17]. The method is based on a description of the dynamical system by two sets of parameters, the nuclear coordinates $\{R_I\}$ that determine the geometry of the molecule or cluster and the SCF coefficients $\{c_{\mu i}\}$, which, together with the traditional atom-fixed Gaussian basis set, determine the shape of the electronic wave function. The dynamics are then given by the equations of motion:

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

$$m_{\text{SCF}} \ddot{c}_{\mu i} = -\frac{\partial E}{\partial c_{\mu i}}, \tag{2}$$

where M_I are the atomic masses; E, the total potential energy, and m_{SCF}, the fictitious mass assigned to the SCF coefficients.

These equations of motion can then be propagated numerically with a long time step Δt for the nuclei and a short time step δt for the SCF coefficients, using the following RESPA variant based upon the simple Verlet algorithm:

First, we calculate the forces $F_c\{c(t), R(t)\}$ and $F_c\{c(t - \delta t), R(t)\}$ on the SCF coefficients $\{c_{\mu i}\}$ at times t and $t - \delta t$, and the force $F_R\{c(t), R(t)\}$ on the nuclei at geometry $R(t)$. Then, we propagate the nuclei for one large time step Δt:

$$R(t + \Delta t) = 2R(t) - R(t - \Delta t) + \frac{(\Delta t)^2}{M} F_R\{c(t), R(t)\}, \tag{3}$$

followed by propagation of the reference system $\{c_0\}$ for n small time steps δt (where $n \times \delta t = \Delta t$), subject to the initial conditions $c_0(t) = c(t)$ and $\dot{c}_0(t) = \dot{c}(t)$ and with the nuclei fixed at $R(t)$:

$$c_0(t + \delta t) = 2c_0(t) - c_0(t - \delta t) + \frac{(\delta t)^2}{m_{\text{SCF}}} \bar{F}_c\{c_0(t), R(t)\}. \tag{4}$$

Finally, we correct the movement of the reference system by the difference between the reference system and the full system, to obtain the true movement of the $c_{\mu i}$ with moving nuclei:

$$c(t + \Delta t) = c_0(t + \Delta t) + D(t + \Delta t) \tag{5}$$

$$c(t + \Delta t - \delta t) = c_0(t + \Delta t - \delta t) + D(t + \Delta t - \delta t). \tag{6}$$

The corrections D are given by

$$D(t + \Delta t) = \frac{(\Delta t)^2}{m_{\text{SCF}}} [F_c\{c(t), R(t)\} - \bar{F}_c\{c_0(t), R(t - \Delta t)\}] \tag{7}$$

$$D(t + \Delta t - \delta t) = \frac{(\Delta t)^2}{m_{\text{SCF}}} [F_c\{c(t - \delta t), R(t)\} - \bar{F}_c\{c_0(t - \delta t), R(t - \Delta t)\}], \tag{8}$$

where $\bar{F}_c\{c_0(t), R(t - \Delta t)\}$ is the force on the reference system with the nuclei held fixed at $R(t - \Delta t)$ and with $c_0(t)$ referring to the coefficients of the reference sys-

tem propagated from $t - \Delta t$ to t and where the approximation $R(t - \delta t) = R(t)$ is assumed to be valid. This approximation enables us to calculate the integrals over basis functions, which enter the forces F_c on the SCF coefficients via the Fock matrix, only for those nuclear geometries obtained through propagation of Eq. (3), i.e., only once per large time step Δt. Also, obviously, the expensive calculation of the forces F_R on the nuclei has to be done only once per large time step Δt. Therefore, we can expect considerable savings in computer time. Because of the correction D, there is no loss in accuracy other than from the finite time differences and the above-mentioned approximation.

We have tested this version of simple Verlet RESPA in a Lennard–Jones system with mass disparity 100:1. These tests were done on a system of eight heavy particles and two light particles that were otherwise identical, and time step ratios of up to 500:1 were tried. Typically, we found that time-step ratios $\Delta t/\delta t$ of 50:1 can be used without appreciable loss of accuracy in the trajectories or in the total energy conservation.

3. Test Results

We have applied this simple Verlet RESPA algorithm to the propagation of a cluster of four sodium atoms in the singlet electronic state with Hellmann–Feynman forces at the HF level of theory. The basis set and effective potential used here are the same as in [16, 17]. The initial geometry for all cases is a planar rectangle with side lengths of 4.82 and 2.89 Å, with slight random distortion to avoid a symmetric situation. The initial nuclear kinetic energy is 0.025 Hartrees, with all velocities pointed inward with equal magnitudes. All runs attempted have been set up to run for the same length of time, approximately 38 fs.

We set the fictitious mass of the SCF coefficients to 3000 atomic units (au). A non-RESPA control run with a time step of 5 au (0.12 fs) and RESPA trajectories with the short time step held fixed at 5 au and long time steps of 25 au (henceforth referred to as 5:1) and 50 au (10:1) were then computed. Plots of the results of these three runs appear in the figures.

From Figure 1, we can see clearly the effect of the different times at which the first time step is completed, as the potential energy does not change until that first long time step is finished. In this plot, as in most of the rest, the deviation from the control trajectory is two to three times as large for the 10:1 curve as for the 5:1 curve, with overall deviations of less than 0.0005 Hartrees.

Figure 2 displays the component in the x direction of the acceleration of a selected atom in the system. Whereas the 5:1 curve is quite close to the non-RESPA control curve, we see that the curve for the 10:1 case deviates considerably more in places. The small peaks in the control curve and the 5:1 curve at $\approx$12 and $\approx$24 fs are due to the reconvergence of the wave function. The 10:1 case was also reconverged at these times, but the peaks due to this reconvergence are not visible against the background of the larger deviations. Note, however, that the RESPA force appears to oscillate nicely about the control force values.

In Figure 3, we have plotted the internuclear distance for a pair of nearest neighbors throughout the simulations. Although the trajectories are close, they

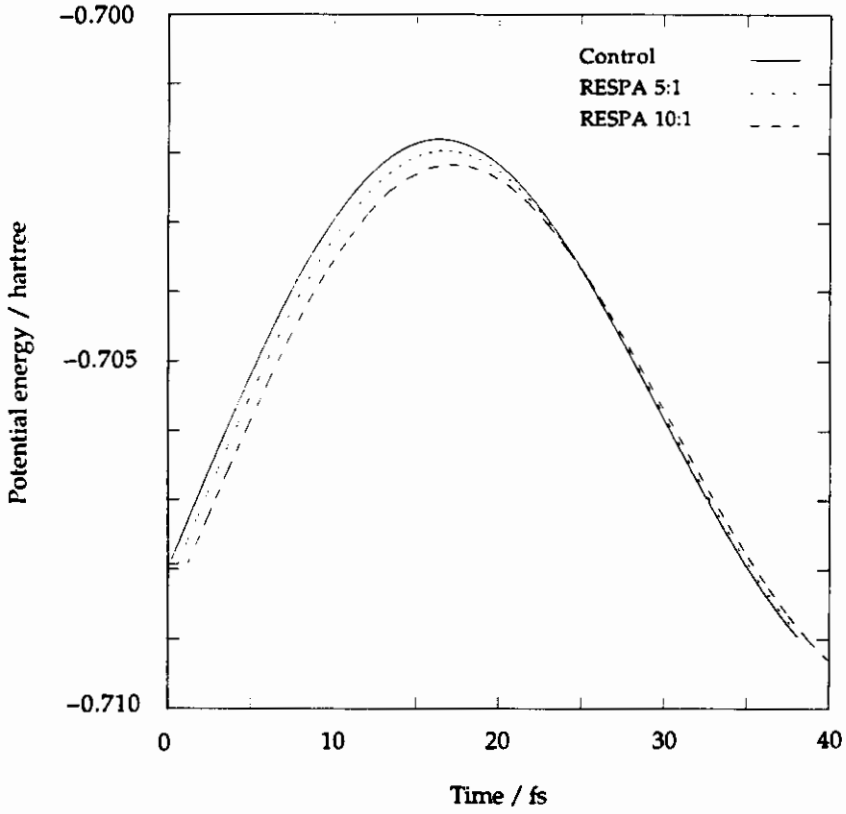

Figure 1. Potential energy in singlet Na_4 propagation: This plot shows the temporal behavior of the potential energy (electronic energy plus internuclear repulsion) as seen by the nuclei at each instant of the propagation for each of three runs. The strange behavior at the beginning of the run is an artifact of the start-up process; see text. In each figure, the curve labeled "Control" is the non-RESPA control run, and the curves labeled "RESPA 5:1" and "RESPA 10:1" are the RESPA test runs with the same short time step as the control run and a long time step that is, respectively, 5 or 10 times larger than in the control run.

are not exactly identical, with maximum deviations of ≈0.01 Å. Note, however, that while the trajectories are diverging slightly near the beginning of the run the difference between the trajectories remains approximately constant for the second half of the run.

In Figure 4, we see that the deviations in the kinetic energy of the nuclei are most apparent in the latter part of the propagation, reaching a maximum deviation of ≈0.0006 Hartrees. Close inspection of the local minima of these three curves reveals that the 10:1 curve reaches its minimum slightly later than the 5:1 curve, which, in turn, reaches its minimum later than the control curve. This is because the potential energy starts changing 45 au later in the 10:1 case than in the control case, as observed in Figure 1.

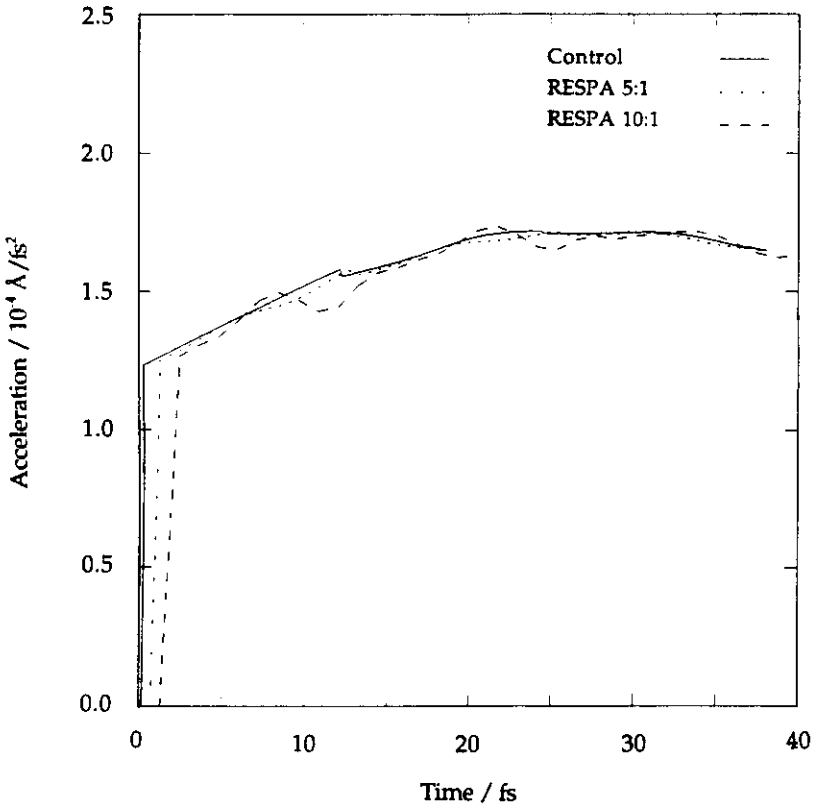

Figure 2. Component in the x direction of the acceleration of a selected atom in the singlet Na_4 propagation: This plot shows the temporal behavior of the x component of the acceleration of a selected atom in each simulation, illustrating how the forces felt by the atoms differ among the test cases. The strange behavior at the beginning of the run is because this value is found by finite difference and thus cannot be computed until two long time steps have been completed.

Figure 5 illustrates the poor performance of the Hellmann–Feynman forces. Even in the non-RESPA control run, energy is conserved quite poorly. Note that (especially in the case of the 5:1 run) the difference between the total energy of the RESPA run and that of the non-RESPA run is nearly constant. It is thus apparent that most of this deviation occurs in the first step. The deviation in the first step is an artifact of the start-up process, since the potential energy does not change during the first long time step.

The most important feature of Figure 6, the kinetic energy of the SCF coefficients, is that this energy is orders of magnitude smaller than the real energies with which we are concerned. Note also that the gross shapes of these three curves are similar. The sudden jumps at ≈12, 24, and 36 fs, which are most obvi-

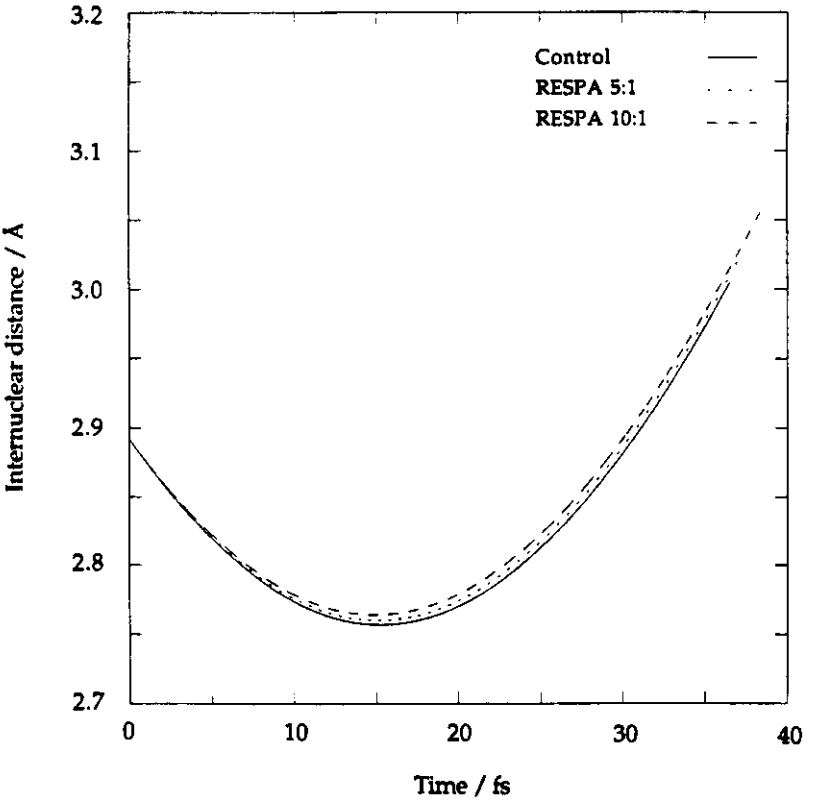

Figure 3. Distance between a pair of nearest neighbors in the singlet Na_4 propagation: This plot shows the distance between a selected pair of nearest neighbors in each simulation, illustrating the similarity of the trajectories of the nuclei among the test cases.

ous in the control curve, are again due to wave-function reconvergence, as in Figure 2.

Although we have observed some deviations from the non-RESPA trajectory, they appear to have little effect on the overall dynamics. For example, while the forces on the nuclei (as observed in Fig. 2) are not quite the same in the RESPA trajectories as in the non-RESPA trajectory, they oscillate around the forces in the non-RESPA trajectory and are, on average, quite close. Second, the positions of the nuclei in the RESPA trajectories diverge from the non-RESPA positions at first, but after a short induction period, this divergence ceases and the nuclear positions in the RESPA cases maintain a fairly constant distance from those positions in the non-RESPA case. Finally, although there is an observable difference in the total energies of the different cases, this manifests itself almost entirely in the first step, after which the difference in the total energies remains nearly constant.

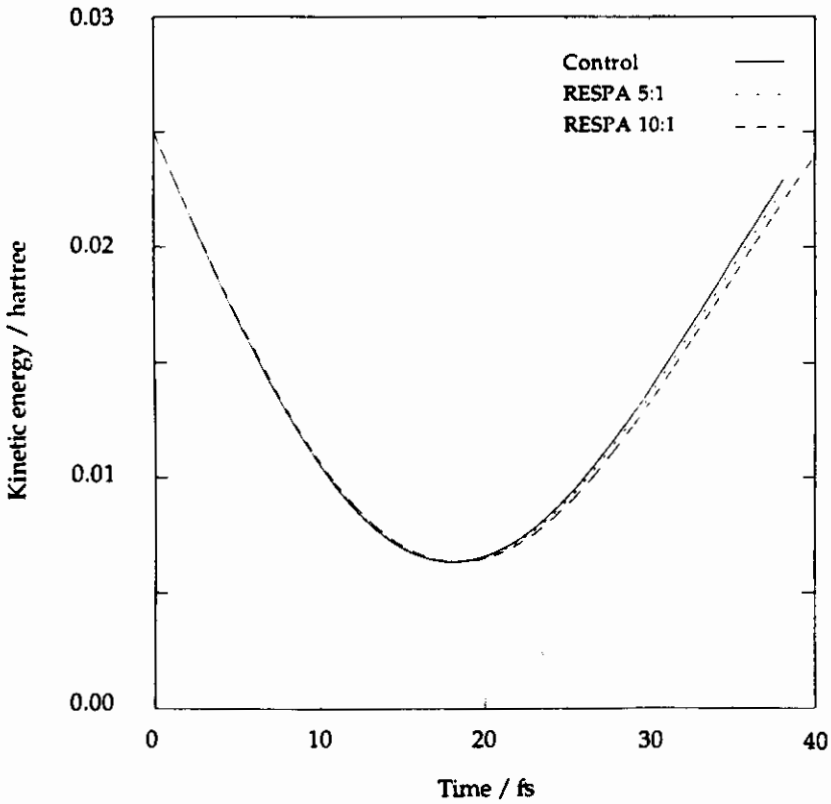

Figure 4. Total kinetic energy of the nuclei in the singlet Na_4 propagation: This plot shows how the kinetic energy of the nuclei changes with time in each of the three simulations.

This indicates that once the RESPA algorithm has completed a full step, it conserves energy about as well as the single time-step Verlet integrator does. We expect that further testing with the exact forces on the nuclei (which conserve energy and enable us to look at the independent criterium of total energy conservation, unlike Hellmann–Feynman forces) will further verify the accuracy of our implementation of the RESPA algorithm.

One other interesting result has been observed in testing the RESPA algorithm for AIMD. It turns out that when RESPA is applied to AIMD, $\Delta t/\delta t$ is not the only factor that affects the performance of the RESPA algorithm relative to the simple Verlet algorithm. Comparing trajectories with a short time step of 5 au, we find that, for this particular system, time-step ratios of 5:1 and 10:1 work for a fictitious mass of 3000 au, but if the fictitious mass is reduced to 300 au, the RESPA algorithm fails completely for time-step ratios of even 4:1 and 5:1. With a fictitious mass of 3000 au, a time-step ratio of 20:1 (long time step = 100 au) also

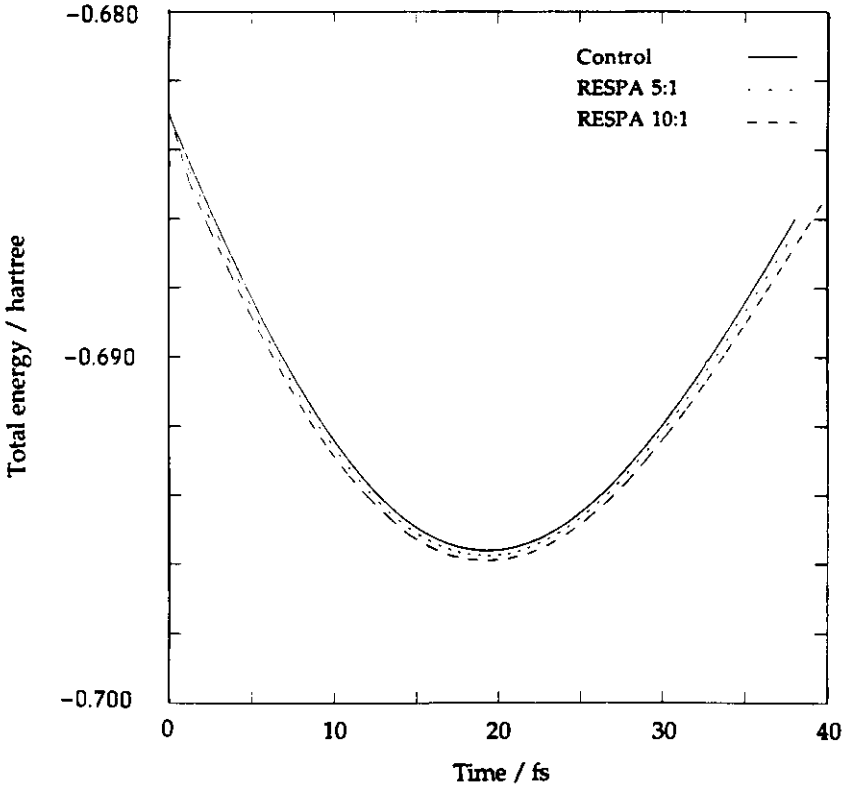

Figure 5. Total energy of the system in the singlet Na_4 propagation: This plot compares the energy conservation of the three test runs. The poor overall conservation is due to the use of Hellmann–Feynman forces. Note also that most of the differences among the three curves come in at the beginning and are an artifact of the start-up process; see text.

fails completely. This can be traced back to the observation that when the factor $(\Delta t)^2/m_{\text{SCF}}$ is larger than approximately unity (our test runs have been sufficient only to bound this number by 0.83 below and 1.33 above for this particular system), the corrections D grow until they are of the same order of magnitude as the SCF coefficients c, at which point the algorithm used to enforce the orthonormality of the orbitals breaks down and the simulation fails. Note that this factor is the prefactor in Eqs. (7) and (8), which bear some similarity to the differential equations that describe exponential growth, when we consider that the forces are obtained by multiplying a Fock matrix by the SCF coefficient vector and that the correction vector D is the difference between the two coefficient vectors that appear in those equations. Thus, the choice of time-step ratio is constrained by the choice of fictitious mass(es), which, in turn, is bounded from above by the desire to minimize the fictitious kinetic energy contribution to the total energy. How-

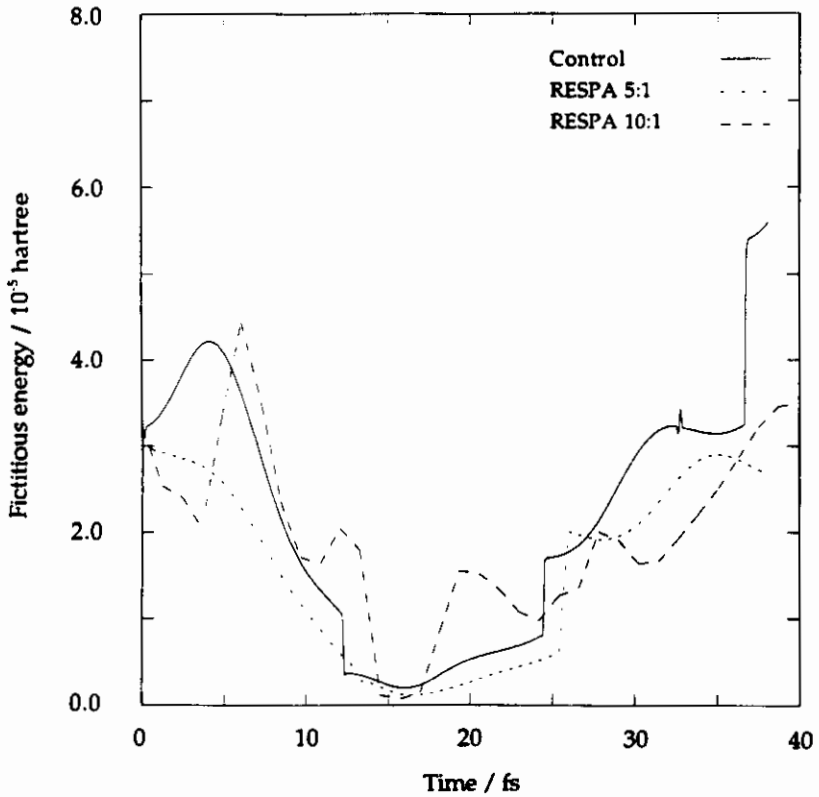

Figure 6. Fictitious kinetic energy of the SCF coefficients in the singlet Na_4 propagation: This plot shows the temporal behavior of the kinetic energy of the SCF coefficients for each of the test cases. This energy arises from the treatment of the SCF coefficients as particles in the propagation. Note that it is orders of magnitude smaller than the real energies observed in Figures 4 and 5.

ever, this does not seriously impair the computer time savings that can be expected for other systems. For example, a mixed system containing lighter atoms forces a choice of a smaller fictitious mass, but the lighter atoms will also be moving faster; this necessitates the choice of shorter time steps δt (to accommodate the faster changing wave function) and Δt, and a shorter Δt may, in turn, reduce the critical factor $(\Delta t)^2/m_{SCF}$ back to an acceptable value.

4. Conclusions

We have presented a simple Verlet AIMD RESPA algorithm. Our test calculations have suggested that for this system time-step ratios of up to 10:1 can be safely used and lead to a savings in computer time of up to a factor of 7. Other systems with different parameters may necessitate a reassessment of this limit, particu-

larly when the concept of rescaling nuclear masses [17] is used to accelerate the dynamics during geometry optimization.

Obviously, this is a static scheme, but the optimum time-step ratio clearly depends not only on the system but also on its temperature and its location on the potential energy surface, especially on the ratio of kinetic energy in the dynamical subsystems of the nuclei and of the wave function parameters (i.e., the heavy and light "particles"). A dynamic adaption of the time-step ratio to the momentary characteristics of the system could therefore lead to even better accuracy and larger savings in computer time.

Acknowledgments

This work was supported by the Office of Naval Research (Grant No. N00014-89-J-1492). E. A. C. also acknowledges support from the National Science Foundation and the Camille and Henry Dreyfus Foundation through their Presidential Young Investigator and Teacher-Scholar Award programs. D. A. G. thanks the National Science Foundation for a predoctoral fellowship. B. H. thanks the Fonds der Chemischen Industrie for a Liebig scholarship.

Bibliography

[1] R. N. Barnett, U. Landman, A. Nitzan, and G. Rajagopal, J. Chem. Phys. **94**, 608 (1991).
[2] K. Runge, D. A. Micha, and E. Q. Feng, Int. J. Quantum Chem., Quantum Chem. Symp. **24**, 781 (1990).
[3] E. Deumens, A. Diz, H. Taylor, and Y. Öhrn, J. Chem. Phys. **96**, 6820 (1992).
[4] P. Bendt and A. Zunger, Phys. Rev. Lett. **50**, 1684 (1983).
[5] R. Car and M. Parrinello, Phys. Rev. Lett. **55**, 2471 (1985).
[6] M. Head-Gordon and J. A. Pople, J. Phys. Chem. **92**, 3063 (1988).
[7] R. Car and M. Parrinello, in *Simple Molecular Systems at Very High Density,* NATO ASI Series, Series B: Physics, Vol. 186, A. Polian, P. Loubeyre, and N. Boccara, Eds. (Plenum Press, New York, 1989); R. O. Jones, Angew. Chem. Int. Ed. Engl. **30**, 630 (1991).
[8] D. Remler and P. A. Madden, Mol. Phys. **70**, 921 (1990).
[9] G. Pastore, E. Smargiassi and F. Buda, Phys. Rev. A **44**, 6334 (1991).
[10] D. Hohl, R. O. Jones, R. Car, and M. Parrinello, J. Chem. Phys. **89**, 6823 (1988); M. R. Pederson, B. M. Klein, and J. Q. Broughton, Phys. Rev. B **38**, 3825 (1988); M. Sprik and M. L. Klein, J. Chem. Phys. **89**, 1592 (1988); *Ibid.* **90**, 7614 (1989); *Ibid.* **91**, 5665 (1989); M. C. Payne, M. P. Teter, and D. C. Allan, J. Chem. Soc., Faraday Trans. **86**, 1221 (1990); J. Hafner and M. C. Payne, J. Phys. Condensed Matter **2**, 221 (1990); F. Ancilotto, W. Andreoni, A. Selloni, R. Car, and M. Parrinello, Phys. Rev. Lett. **65**, 3148 (1990); R. Kawai and J. H. Weare, Phys. Rev. Lett. **65**, 80 (1990); A. St-Amant and D. R. Salahub, Chem. Phys. Lett. **169**, 387 (1990); C. Tsoo, D. A. Estrin, and S. J. Singer, J. Chem. Phys. **93**, 7187, 7201 (1990); D. Hohl and R. O. Jones, Phys. Rev. B Condensed Matter, **43**, 3856 (1991); U. Röthlisberger and W. Andreoni, J. Chem. Phys. **94**, 8129 (1991); *Ibid.* **96**, 1248 (1992); W. Andreoni, F. Gygi, and M. Parrinello, Phys. Rev. Lett. **68**, 823 (1992); *Ibid.* Chem. Phys. Lett. **189**, 241 (1992); *Ibid.* **190**, 159 (1992).
[11] J. Broughton and F. Khan, Phys. Rev. B **40**, 12098 (1989); A. Caro, S. Ramos de Debiaggi, and M. Victoria, Phys. Rev. B **41**, 913 (1990); M. J. Field, Chem. Phys. Lett. **172**, 83 (1990).
[12] O. F. Sankey and D. J. Niklewski, Phys. Rev. B **40**, 3979 (1989); O. F. Sankey, D. J. Niklewski, D. A. Drabold, and J. D. Dow, Phys. Rev. B **41**, 12750 (1990).
[13] H. Chacham and J. R. Mohallem, Mol. Phys. **70**, 391 (1990).
[14] M. J. Field, J. Phys. Chem. **95**, 5104 (1991).

[15] J. Mohallem, R. O. Vianna, and H. Chacham, Poster presented at the 32nd Sanibel Symposia in Florida, March 14–21, 1992; submitted.
[16] B. Hartke and E. A. Carter, Chem. Phys. Lett. **189**, 358 (1992).
[17] B. Hartke and E. A. Carter, J. Chem. Phys., in press (1992).
[18] M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids* (Oxford University Press, Oxford, 1987).
[19] W. B. Streett, D. J. Tildesley, and G. Saville, Mol. Phys. **35**, 639 (1978).
[20] D. E. Smith, Y.V. Kalyuzhnyi, and A. D. J. Haymet, J. Chem. Phys. **95**, 9165 (1991).
[21] O. Teleman and B. Jönsson, J. Comp. Chem. **7**, 58 (1986).
[22] R. D. Swindoll and J. M. Haile, J. Comp. Phys. **53**, 289 (1984).
[23] M. E. Tuckerman, B. J. Berne, and A. Rossi, J. Chem. Phys. **94**, 1465 (1991).
[24] M. E. Tuckerman, B. J. Berne, and G. J. Martyna, preprint (1992).

Received July 14, 1992
Revised manuscript received August 11, 1992
Accepted for publication August 12, 1992