

Molecular Dynamics with Multiple Time Scales: How to Avoid Pitfalls

Joseph A. Morrone,[†] Ruhong Zhou,[‡] and B. J. Berne^{*,†,‡}

Department of Chemistry, Columbia University, 3000 Broadway, MC 3103, New York, New York 10027, and IBM Thomas J. Watson Research Center, 1101 Kitchawan Road, Yorktown Heights, New York 10598

Received January 27, 2010

Abstract: Multiple time scale methodologies have gained widespread use in molecular dynamics simulations and are implemented in a variety of ways across numerous packages. However, performance of the algorithms depends upon the details of the implementation. This is particularly important in the way in which the nonbonded interactions are partitioned. In this work, we show why some previous implementations give rise to energy drifts, and how this can be corrected. We also provide a recipe for using multiple time step methods to generate stable trajectories in large scale biomolecular simulations, where long trajectories are needed.

1. Introduction

Molecular dynamics is a ubiquitous tool for simulating a wide variety of large scale systems, ranging from the materials to the biological sciences. Schemes that increase the efficiency of such simulations are of great interest. In standard techniques, the time step of the generated trajectory is limited by the fastest motions present in the system. However, realistic systems contain a broad spectrum of frequencies. Multiple time scale (MTS) methods partition the computation into “slow” and “fast” portions, assigning appropriate time steps to each segment. This methodology may be exploited in systems with disparate masses,¹ high frequency oscillators in slowly evolving baths,² and distance based schemes that partition the nonbonded interactions into short- and long-range components.^{3–5}

The reversible reference system propagator algorithm (r-RESPA)⁶ is one of the most powerful implementations of the multiple time scale concept. r-RESPA integrators are readily derived from factorization of the Liouville propagator.^{6,7} It therefore provides an integration scheme that is reversible in time and evolves in a symplectic and area preserving fashion, thereby preserving these attributes of an exact solution to Hamilton’s classical equations of motion. Furthermore, a variety of different multiple time scale

partitionings may be readily recovered from this framework, including related integrators.⁵ This algorithm has been widely implemented in simulation packages such as IMPACT,⁸ NAMD,⁹ AMBER,¹⁰ and DESMOND.¹¹

The widespread availability of fast multicore computer clusters, massively parallel supercomputers, and the improvements of parallel algorithms have facilitated the simulation of longer trajectories on the order of tens of nanoseconds to microseconds for large biomolecular systems. Since the majority of papers reporting tests of the stability of multiple times scale methods were published before such advances, it is important to evaluate the validity of MTS algorithms using much longer times scales. A more recent work by Han et al.¹² studies disparate time scales in a simulation of a biomolecular system in a Langevin bath over several nanoseconds. Here, we focus on the stability of the integrator as measured by its energy conservation in the microcanonical simulation. This is due to the fact that coupling the system to the thermostats and barostats necessary to generate other ensembles may obscure or complicate the evaluation of the integrator’s stability.

Symplectic integrators such as r-RESPA can successfully generate long stable trajectories. In addition, for such integrators, there exists a modified or shadow Hamiltonian^{13,14} which is exactly conserved as the system is propagated, although it is only known approximately for realistic systems.^{15,16} However, despite these desirable properties, r-RESPA and related integrators are known to suffer from

* Corresponding author e-mail: bb8@columbia.edu.

[†] Columbia University.

[‡] IBM Thomas J. Watson Research Center.




Figure 1. (Top) Deviation of the total energy from the initial energy ($E_0 = -240\,068$ kcal/mol) plotted against trajectory length for a solvated lysozyme using the NAMD2 package. The lengths of all covalent bonds to hydrogen are constrained in all runs. The default RESPA2/C1 (red line) and BONDED/NON-BONDED (labeled B/NB, green line) schemes are plotted against a standard velocity Verlet run (black line). The bottom panel shows a NAMD2 simulation of the same system. The red and black lines correspond to the same plots as in the top panel. The blue line utilizes the RESPA2 type splitting with our fix, which has been implemented as the “long-Splitting=c2” option in NAMD2, version 2.7b2 (see section 3).

resonance instabilities,^{5,17–19} which bound the size of the time step of the slowest motions relative to the size of the faster modes. Resonance phenomena engender the building up of energy in the system, thereby giving rise to drifts in average properties and inaccurate sampling. Schemes have been developed to alleviate this problem,^{18,20–27} although some are only suitable for sampling and not investigations of system dynamics.

In this work, however, we do not focus on resonance phenomena in particular, but rather on optimizing the splitting of the long-range nonbonded electrostatics in order to ensure long time numerical stability. In periodic systems, nonbonded interactions may be split in two ways. We will refer to these schemes as RESPA1 (split by real-space and k-space in Ewald summation) and RESPA2 (split by intrinsic time scale; see section 2). Another consideration is how the interactions are smoothed at the boundary between the long- and short-range interactions.^{4,6} This is facilitated by use of a switching function, the details of which may be crucial to generating a stable trajectory.

As with any numerical scheme, the performance of r-RESPA depends on the details of its implementation. In order to illustrate this point, we present in Figure 1 the deviation of the total energy from its initial energy over the course of a microcanonical simulation of lysozyme in an 8 M urea solution performed utilizing several integration schemes, as implemented in the NAMD2 simulation package. The simulations were run for up to 200 ns, but only the first 15 ns of data are shown. It can be seen that the standard velocity Verlet integrator is stable with a time step of 1 fs. In the r-RESPA schemes, the bonded interaction is evaluated with a time step of 1 fs. However, the nonbonded interaction

is handled in two different ways, each utilizing an additional time scale of 2 fs. In one version, the entire nonbonded interaction is evaluated at the larger time step. We refer to this as “RESPA:B/NB”, and the results are reported in the top panel of Figure 1. Also plotted in the top panel is the case when the nonbonded interactions are split across the two time scales according to the default implementation in NAMD2. This is denoted as “RESPA2:C1” (as we will discuss in section 2, this implementation in NAMD2 is similar to RESPA2 as presented in the literature^{28,29} but is not exactly the same). It can be seen that, whereas RESPA:B/NB is relatively stable, a significant energy drift is present in the RESPA2:C1 result. Counterintuitively, therefore, the larger drift is seen in the case where a greater portion of the interaction is integrated at the shorter time step.

It has been shown that seemingly reasonable implementations can give rise to unexpected energy drifts. We are therefore motivated to test the stability of the various schemes of applying the r-RESPA algorithm in microcanonical simulations over long trajectories. In particular, we are interested in investigating the details of the partitioning of the nonbonded interactions. Stability depends upon the details of the interaction split, such that the incorrect choice of parameters can lead to unstable trajectories. We find that the large energy drift of RESPA2:C1 as shown in the top panel of Figure 1 is engendered by the choice of switching function that facilitates the partitioning of the nonbonded interactions (electrostatic interactions to be specific). The bottom panel of this figure shows the total energy for the trajectory for the same RESPA2 scheme, except that the switching function has been “fixed” (denoted RESPA2:C2). This modification has been recently ported into NAMD2, and its details will be explained below.

This article is organized as follows: Section 2 reviews the different ways to decompose the nonbonded interactions. In section 3, the sensitivity of the nonbonded splittings is tested for a simple water system, and what is learned here is applied to a biomolecular system in section 4. Conclusions are given in section 5.

2. Choosing the Force Splitting

Of crucial importance to the nature of the algorithm is the way in which the multiple time scales are defined. In a typical empirical potential, this may be done by splitting the force into a set of terms that are evaluated at different time steps. It is typical for the fastest motions to be chosen as the stretching and bending terms of the force field. The torsional terms of the potential may be included here, or treated at another level “outside” the stretching and bending interaction. The nonbonded interaction may be split into two or more parts according to their relative intrinsic time scales (fast or slow, on the basis of pair distances). In this work, we will only consider nonbonded potentials of the following form which act in a periodic simulation cell with vector of periodicity \mathbf{n} :

$$V = \sum_{\mathbf{n}} \sum_i \sum_{j \geq i} (1 - \delta_{ij}^0) \left[\phi(r_{ij}^{\mathbf{n}}) + \frac{q_i q_j}{r_{ij}^{\mathbf{n}}} \right] \quad (1)$$

where the sum is over all periodic images and all pairs that do not correspond to the same atomic site. The function, ϕ , is a short-ranged potential effectively accounting for repulsive and dispersion interactions (typically taken to be of the Lennard-Jones form), and the second term is the electrostatic interactions of fixed point charges. The electrostatic potential is long-range and may be treated via the Ewald summation technique^{3,30,31}

$$\sum_{\mathbf{n}} \sum_i \sum_{j \geq i} (1 - \delta_{ij}^0) \frac{q_i q_j}{r_{ij}^n} = V_{\text{scr}} + V_{\text{KS}} \quad (2)$$

where V_{scr} is a short-ranged, screened potential and V_{KS} is a smooth, slowly varying potential that is most efficiently computed in reciprocal space:

$$V_{\text{scr}} = \sum_{\mathbf{n}} \sum_i \sum_{j \geq i} (1 - \delta_{ij}^0) q_i q_j \frac{\text{erfc}(\alpha r_{ij}^n)}{r_{ij}^n} \quad (3)$$

$$V_{\text{KS}} = \sum_{\mathbf{n}} \sum_i \sum_{j \geq i} (1 - \delta_{ij}^0) q_i q_j \frac{\text{erf}(\alpha r_{ij}^n)}{r_{ij}^n} \quad (4)$$

$$= \frac{1}{L^3} \sum_{\mathbf{k} \neq 0} \frac{2\pi}{k^2} e^{-k^2/4\alpha^2} |S(\mathbf{k})|^2 - V_{\text{self}} \quad (5)$$

where $S(\mathbf{k}) = \sum_i^N q_i e^{i\mathbf{k} \cdot \mathbf{r}_i}$. The term V_{self} subtracts out the interaction between the same sites that is implicit in the reciprocal space sum. As it is position-independent, it will not contribute to the forces and will be neglected for the rest of this discussion. The parameter α determines the degree of screening and is chosen in accordance with the real space interaction cutoff r_{cut} . The reciprocal space part of the Ewald summation may be computed directly or by means of methods that utilize fast Fourier transforms (FFT) such as particle mesh Ewald (PME),³² smooth particle mesh Ewald (SPME),³³ particle-particle particle-mesh Ewald³⁴ (P3ME), and the fast multipole method³⁵ (FMM). Such techniques have been implemented alongside r-RESPA^{29,36–38}

Within the multiple time scale framework, the nonbonded forces may be split into two or more partitions.^{29,39,40} In this study, we will restrict ourselves to splitting the non-bonded force into two parts. When choosing a splitting for the forces into near and far contributions, a natural choice would be to utilize the explicitly short-ranged potentials as the near force and the reciprocal space sum as the long force. This choice also has the utility that the more computationally expensive reciprocal space part is computed less frequently. Following the nomenclature of ref 29, this choice is referred to as “RESPA1.”

$$V_{\text{near}}^{(1)} = \sum_{\mathbf{n}} \sum_i \sum_{j \geq i} (1 - \delta_{ij}^0) \left[\phi(r_{ij}^n) + q_i q_j \frac{\text{erfc}(\alpha r_{ij}^n)}{r_{ij}^n} \right] \times (1 - \Theta(r_{ij}^n - r_{\text{cut}}))$$

$$V_{\text{far}}^{(1)} = V_{\text{KS}} \quad (6)$$

The factor $(1 - \Theta(r_{ij}^n - r_{\text{cut}}))$ cuts off the interaction in real space, where $\Theta(r)$ is the Heavyside function. Smoother

functions may be utilized to facilitate improved energy conservation.

As noted in previous work,^{28,29} this is not the optimal split, as some fast components are screened out of the potential in eq 3 and are therefore present in the reciprocal space term (eq 5). One may subtract this portion from the “far” potential and add it into the “near” potential, yielding a split that we will refer to (as in ref 29) as “RESPA2.”

$$V_{\text{near}}^{(2)} = V_{\text{near}}^{(1)} + \sum_{\mathbf{n}} \sum_i \sum_{j \geq i} (1 - \delta_{ij}^0) (1 - \Theta(r_{ij}^n - r_{\text{cut}})) \times \frac{q_i q_j \frac{\text{erf}(\alpha r_{ij}^n)}{r_{ij}^n}}{r_{ij}^n} \quad (7)$$

$$= \sum_{\mathbf{n}} \sum_i \sum_{j \geq i} (1 - \delta_{ij}^0) (1 - \Theta(r_{ij}^n - r_{\text{cut}})) \times \left[\phi(r_{ij}^n) + \frac{q_i q_j}{r_{ij}^n} \right]$$

$$V_{\text{far}}^{(2)} = V_{\text{KS}} - \sum_{\mathbf{n}} \sum_i \sum_{j \geq i} (1 - \delta_{ij}^0) (1 - \Theta(r_{ij}^n - r_{\text{cut}})) \times \frac{q_i q_j \frac{\text{erf}(\alpha r_{ij}^n)}{r_{ij}^n}}{r_{ij}^n} \quad (8)$$

The forces may be obtained by taking the negative gradient of the associated piece of the potential and then splitting it according to the r-RESPA algorithm.⁶ In general, a cutoff different than the overall real space interaction cutoff, $r_{\text{cut}}^{\text{res}}$, may be employed for the division between the “near” and “far” forces. The forces may be decomposed as follows:^{4,6}

$$f_{\text{inner}} = f_{\text{near}} S(r; r_{\text{cut}}^{\text{res}}, \lambda)$$

$$f_{\text{outer}} = f_{\text{far}} + f_{\text{near}} (1 - S(r; r_{\text{cut}}^{\text{res}}, \lambda)) \quad (9)$$

where S is a switching function that softens the transition from fast to slow forces that occurs at a given cutoff. The switching function is made to act over a healing length, λ . Too harsh a transition can lead to errors in the region of the cutoff, thereby introducing instabilities in the propagation that accumulate over the integration time. In some implementations,^{9,22} the entirety of short ranged potential ϕ is placed in the inner loop and only the electrostatic interactions are split. It is also possible to instead apply a switching function directly to the potential.^{22,41} If S is applied to the potential, then the corresponding switching function \tilde{S} will act upon the force, thereby replacing S in eq 9 with the following expression:

$$\tilde{S}(r) = S(r) + \frac{V(r)}{V'(r)} S'(r) \quad (10)$$

where the prime indicates a derivative with respect to distance r . Note that, since the forces and not the potentials are utilized to generate the trajectories, it is the smoothness of the switching function that acts on the force which will impact the integration stability. In principle, equivalent switching functions, either on the potential or on the force, will generate the same smooth trajectories.

The switching function, S , used here is given in the following piecewise form:

$$S(r; r_{\text{cut}}^{\text{res}}, \lambda) = \begin{cases} 1 & r \leq r_{\text{cut}}^{\text{res}} - \lambda \\ g(r; r_{\text{cut}}^{\text{res}}, \lambda) & r_{\text{cut}}^{\text{res}} - \lambda < r < r_{\text{cut}}^{\text{res}} \\ 0 & r \geq r_{\text{cut}}^{\text{res}} \end{cases} \quad (11)$$

where the function, g , is chosen so as to smoothly transition from 1 to 0. In this work, we will consider three forms of this function: a cubic spline, g_3 , that was utilized on the force in the original formulation of r-RESPA,^{4,6} a quintic spline, g_5 , that was utilized to cut off the electrostatic potential in a different context,⁴² and a different form of cubic spline (denoted as the C1 spline), g_{C1} , which is the default choice in NAMD.⁹

$$g_3(r; r_{\text{cut}}^{\text{res}}, \lambda) = 1 + u^2(2u - 3) \quad (12)$$

$$g_5(r; r_{\text{cut}}^{\text{res}}, \lambda) = 1 + u^3(15u - 6u^2 - 10) \quad (13)$$

$$g_{\text{C1}}(r; r_{\text{cut}}^{\text{res}}, \lambda) = 1 + \frac{u}{2}(u^2 - 3) \quad (14)$$

where:

$$u = \frac{1}{\lambda}(r - r_{\text{cut}}^{\text{res}} + \lambda) \quad (15)$$

The quintic spline has the benefit of being a smoother function than the other choices, although it is also marginally more computationally expensive. Of course, we make no pretense of making the optimal choice, and other functions may be appropriate. The sensitivity of the algorithm performance to the nonbonded interaction splitting will be considered in the next section.

3. Testing the Splitting

In order to test the performance of the schemes delineated above, we perform a series of simulations and monitor the total energy conservation via the deviation of the energy from its initial energy, $E - E_0$, as a function of time, where E is the total energy and E_0 is the initial energy. This plot monitors the drift and fluctuations of the conserved quantity.

We carried out simulations utilizing the PINY_MD package.⁴³ The PINY_MD package contains a multifaceted r-RESPA implementation to which we have added the RESPA2 splitting and the quintic switching function. The particle mesh Ewald method is utilized for computing the long-range interactions.³³ An overall real space cutoff of 10 Å and an $r_{\text{cut}}^{\text{res}}$ of 8 Å were utilized in all of our PINY_MD studies in this and the succeeding section. Using a relatively small cutoff for $r_{\text{cut}}^{\text{res}}$ is presently computationally advantageous due to the fact that it shifts a greater burden onto the less frequently evaluated “far” interactions. It has been shown by Han et al.¹² that the cutoffs may be increased to yield larger differences between the inner and outer time steps such that the algorithm efficiency may be optimized according to the features of the simulation package and available hardware. The switching function is applied directly to the force as in eq 9.




Figure 2. The deviation, $E - E_0$ (where $E_0 = -7259$ kcal/mol), of the conserved energy for TIP3P water when the RESPA1 (black lines) and RESPA2 (red lines) schemes are utilized. The top panel depicts the results if a cubic switching function with a healing length of 2 Å is employed, whereas in the bottom panel a smoother choice for the switch (a quintic function with a healing length of 4 Å) is made.

Due to the large difference between fast OH stretches and the slower librations and translational motions, liquid water is a natural system on which to test the multiple time scale approach. However, the high frequency of the OH stretch induces resonance instabilities (see section 1) at rather small outer time steps.²² The resonance barrier may be simply postponed by constraining the lengths of all the covalent bonds to hydrogens, and this is the approach that we follow here. Therefore, we utilize a rigid model of water. Since all covalent bond lengths are being constrained, the forces are only split between near and far nonbonded contributions.

We simulate a system of 905 water molecules in a periodic cubic cell with a side 30 Å in length, and the TIP3P model⁴⁴ is utilized to describe the interactions. The internal geometry of each molecule is constrained.⁴⁵ The near forces are updated every 1 fs, whereas the far forces are updated every 5 fs. The Ewald screening parameter utilized is $\alpha = 0.4$ Å. In this regime, previous studies have shown that the multiple time scale splitting is stable for rigid water models.^{22,29} Simulations of 4 ns in length were carried out in order to test the sensitivity of RESPA1 and RESPA2 to the choice of switching function and healing length.

This model will serve as our test of the sensitivity to choice of switching function parameters and functional forms. Healing lengths of 2 and 4 Å, as well as the cubic and quintic forms, were studied. In Figure 2, $E - E_0$ is plotted versus time for selected choices of the switching function using both the RESPA1 and RESPA2 schemes. Additional stability data for all runs are given in the Appendix. It can be seen that the energy conservation dramatically improves as the switching function is made “smoother” for RESPA2, although RESPA1 is relatively insensitive to this change. The switching function can be made smoother by increasing its order and by increasing the distance over which the function acts (λ). As is evident from the data in the Appendix, the careful choice of both of these aspects is necessary to optimize performance. Sensitivity of the RESPA2 scheme to the

choice of healing length has been previously noted in ref 24. The RESPA1 scheme is relatively insensitive to this choice due to the fact that the interactions are already damped in the cutoff region because the screened and not the bare Coulomb potential is utilized for the “near interaction” (see eq 6). Furthermore, one may note that the drift observed in the RESPA2 scheme for a poor choice of switching function can go unnoticed over the first several hundred picoseconds of the simulation. This observation underlines the importance of monitoring longer trajectories in order to assess the performance of any integration scheme.

We now return to the original question posed by Figure 1. As noted in the Introduction, the nonbonded portion is split according to RESPA2 within NAMD2.⁹ By default, the C1 switching function (eq 14), with a hard wired healing length of $\lambda = r_{\text{cut}}$, is applied directly to the *potential*. Numerical stability therefore depends upon the smoothness of $\tilde{S}(r)$ as defined in eq 10 (see section 2). In the bottom panel of Figure 1, we plot the results shown in the top panel of this figure against what occurs if the default C1 switching function is simply replaced with our smoother quintic switching function for the RESPA2 scheme.⁴⁶ One can see that this single modification largely alleviates the drift in the conserved energy and provides a fix, albeit not an optimal one.

4. Appropriate Settings for Biomolecular Simulations

In current research, the primary utility of multiple time scale methods is to increase the computational efficiency of biomolecular simulations. To this end, we make a careful study of the performance of r-RESPA for such systems. As a test case, we simulate a lysozyme surrounded by 6328 TIP3P water molecules in a periodic cubic box with a side of length 61.5 Å using the CHARMM22⁴⁷ force field. The Ewald screening parameter is set to 0.37 Å. We chose a time step of 1 fs for the near nonbonded interactions and torsional terms, and we integrate the stretching and bending terms of the protein with a time step of 0.5 fs. The outer time step that splits the nonbonded interactions is varied. We utilize a quintic switching function with a healing length of 4 Å in all runs. All water molecules are taken to be rigid, and all bonds to hydrogen (except hydrogen bonds) within the protein are also constrained, so as to delay resonance instabilities (see sections 1 and 3). We equilibrate the system using the RESPA2-4FS protocol for over 4 ns.

Several runs are presented in Figure 3. It is shown that the splitting is strictly stable up to 4 fs, whereas larger steps exhibit some degree of drifting due to the onset of resonance effects discussed in section 1. When the switching function is properly set, it can be seen that the RESPA2 scheme outperforms RESPA1. This is indicated by the smaller energy fluctuations in the RESPA2 runs. For an outer time step of 6 fs, the RESPA2 result also possesses a significantly smaller drift than the corresponding RESPA1 run and, in fact, appears to be fairly stable, as shown in Figure 3. The small drifts become more apparent, however, as the simulation progresses beyond 8 ns (results not shown).




Figure 3. The deviation, $E - E_0$, of the energy from the initial energy is plotted for selected runs of lysozyme utilizing RESPA1 with an outer time step of 4 fs (black line) and 6 fs (purple line) and RESPA1 with an outer time step of 4 fs (red line) and 6 fs (blue line). For this system, $E_0 = -53\,377$ kcal/mol.

Table 1. Comparison of the Average Total Energy, Potential Energy, and Temperature for Selected Simulations of the Solvated Lysozyme System^a

system	E (kcal/mol)	T (K)	V (kcal/mol)
NO-SPLIT-1FS	-53374 (0.17)	299.45 (1.7)	-66138 (73)
RESPA1-4FS	-53373 (0.73)	299.24 (1.7)	-66127 (72)
RESPA2-4FS	-53374 (0.30)	299.64 (1.7)	-66147 (73)

^a The standard deviation of each quantity is given in parentheses.

Furthermore, it is shown in Table 1 that the r-RESPA runs yield equivalent averages to runs where the nonbonded interactions are not split and evaluated every femtosecond. It may be possible to utilize larger outer time steps by either increasing the real space cutoff¹² or by splitting the nonbonded interaction into more than two portions,²⁹ where different time scales may be used to characterize near, intermediate, and long-range nonbonded interactions.

5. Conclusion

Multiple time scale molecular dynamics techniques can be an important tool for the creation of optimized molecular dynamics integrators. Although splitting the nonbonded interactions according to their intrinsic space or time scales may be readily accomplished, the proper division of the intermediate and long-range interactions so as to ensure stability and energy conservation can be full of pitfalls. In this work, we have performed a detailed study of the accuracy of nonbonded splitting schemes, in particular, schemes where the electrostatic interactions are split into real and reciprocal space (RESPA1) parts or distance based (RESPA2) contributions. It is found that, while RESPA2 outperforms RESPA1, it has a greater dependence upon the details of the function that switches between the two contributions. This dependence can lead to rather significant drifts in the total energy over the course of long simulations, as is found in some previous implementations. To this end, we have provided some guidance for nonbonded splitting

Table 2. Simulation Details as well as Energy Conservation Measures E_{conv} and R for the Simulations of Water (section 3) and the Solvated Lysozyme (section 4)^a

system	split	Δt_{outer}		length		$\log(E_{\text{conv}})$	$R (\times 10^{-2})$
		(fs)	sw_{ord}	$\lambda(\text{\AA})$	(ns)		
water	RESPA1	5	3	2.0	4.0	-4.42	1.37
water	RESPA1	5	5	4.0	4.0	-4.42	1.37
water	RESPA2	5	3	2.0	4.0	-2.96	21.8
water	RESPA2	5	5	2.0	4.0	-4.38	1.48
water	RESPA2	5	3	4.0	4.0	-3.75	5.63
water	RESPA2	5	5	4.0	4.0	-4.65	0.628
lysozyme	NO-SPLIT	1			2.0	-5.58	0.231
lysozyme	RESPA1	4	5	4.0	8.0	-4.95	1.02
lysozyme	RESPA1	6	5	4.0	8.0	-3.64	10.4
lysozyme	RESPA2	4	5	4.0	8.0	-5.23	0.408
lysozyme	RESPA2	6	5	4.0	8.0	-4.59	1.35

^a Entries are categorized according to the size of the outer loop time step, the type of non-bonded splitting employed (RESPA1, RESPA2, or NO-SPLIT), the details of the switching function (order (sw_{ord}) and healing length (λ)), and trajectory length.

schemes and implemented these in selected simulation packages (NAMD2 and PINY_MD). Even though these implementations are applied to particular simulation packages, we believe that these findings are of broader applicability to multiple time scale methods and will be particularly useful for modern biomolecular simulations where long trajectories on the order of tens of nanoseconds to microseconds are needed.

Acknowledgment. This research was supported from a grant to B.J.B from the National Science Foundation via grant (NSF-CHE-0910943). R.Z. acknowledge support from IBM Blue Gene Science Program. We acknowledge Mark Tuckerman for his assistance with regards to our modifications of PINY_MD. We also acknowledge Jingyuan Li and Thomas Markland for useful discussions.

Appendix: Quantifying the Integrator Stability

We utilize two standard measures in order to assess the stability of the run:^{48–50}

$$E_{\text{conv}} = \frac{|E - E_0|}{|E_0|} \quad (16)$$

$$R = \frac{\langle (E - \langle E \rangle)^2 \rangle^{1/2}}{\langle (T - \langle T \rangle)^2 \rangle^{1/2}} \quad (17)$$

where E is the total energy and T is the kinetic energy. In general, E_{conv} is more sensitive to the drift in energy, whereas R is more directly related to its fluctuation. The details of all the simulations performed with PINY_MD are given alongside these measures of stability in Table 2. It can be seen that RESPA2 does not outperform RESPA1 until both a suitable smooth switching function form and appropriate healing length are chosen.

References

- (1) Tuckerman, M.; Berne, B.; Rossi, A. *J. Chem. Phys.* **1991**, *94*, 1465.
- (2) Tuckerman, M.; Martyna, G.; Berne, B. *J. Chem. Phys.* **1990**, *93*, 1287.
- (3) Allen, M.; Tildesley, D. *Computer Simulation of Liquids*; Oxford University Press: Oxford, 1987.
- (4) Tuckerman, M.; Berne, B.; Martyna, G. *J. Chem. Phys.* **1991**, *94*, 6811.
- (5) Grubmuller, H.; Heller, H.; Windemuth, A.; Schulten, K. *Mol. Simul.* **1991**, *6*, 121.
- (6) Tuckerman, M.; Berne, B.; Martyna, G. *J. Chem. Phys.* **1992**, *97*, 1990.
- (7) Sexton, J.; Weingarten, D. *Nucl. Phys. B* **1992**, *380*, 665.
- (8) Banks, J.; Beard, H.; Cao, Y.; Cho, A.; Damm, W.; Farid, R.; Felts, A.; Halgren, T.; Mainz, D.; Maple, J.; Murphy, R.; Philipp, D.; Repasky, M.; Zhang, L.; Berne, B.; Friesner, R.; Gallicchio, E.; Levy, R. *J. Comput. Chem.* **2005**, *26*, 1752.
- (9) Phillips, J.; Braun, R.; Wang, W.; Gumbart, J.; Tajkhorshid, E.; Villa, E.; Chipot, C.; Skeel, R.; Kale, L.; Schulten, K. *J. Comput. Chem.* **2005**, *2*, 1781.
- (10) Case, D.; Darden, T.; Cheatham, I.; Simmerling, C.; Wang, J.; Duke, R.; Luo, R.; Crowley, M.; Walker, R.; Zhang, W.; Merz, K.; Wang, B.; Hayik, S.; Roitberg, A.; Seabra, G.; Kolossvary, I.; Wong, K. F.; Paesani, F.; Vanicek, J.; Wu, X.; Brozell, S.; Steinbrecher, T.; Gohlke, H.; Yang, L.; Tan, C.; Mongan, J.; Hornak, V.; Cui, G.; Mathews, D.; Seetin, M.; Sagui, C.; Babin, V.; Kollman, P. *AMBER 10*; University of California: San Francisco, 2008.
- (11) Bowers, K.; Chow, E.; Xu, H.; Dror, R.; Eastwood, M.; Kolossvary, B.; Moraes, M.; Sacerdoti, F.; Salmon, J.; Shan, Y.; Shaw, D. Scalable Algorithms for Molecular Dynamics Simulations on Commodity Clusters. Proceedings of the ACM/IEEE Conference on Supercomputing (SC06), Tampa, FL, 2006.
- (12) Han, G.; Deng, Y.; Glimm, J.; Martyna, G. *Comput. Phys. Commun.* **2007**, *176*, 271.
- (13) Skeel, R.; Zhang, G.; Schlick, T. *Siam J. Sci. Comput.* **1997**, *18*, 203.
- (14) Gans, J.; Shalloway, D. *Phys. Rev. E* **2000**, *61*, 4587.
- (15) Skeel, R.; Hardy, D. *Siam J. Sci. Comput.* **2001**, *23*, 1172.
- (16) Engle, R.; Skeel, R.; Drees, M. *J. Comput. Phys.* **2005**, *206*, 432.
- (17) Biesadeki, J.; Skeel, R. *J. Comput. Phys.* **1993**, *109*, 318.
- (18) Schlick, T.; Mandziuk, M.; Skeel, R.; Srinivas, K. *J. Comput. Phys.* **1998**, *140*, 1.
- (19) Ma, Q.; Izaguirre, J.; Skeel, R. *Siam J. Sci. Comput.* **2003**, *24*, 1951.
- (20) Barth, E.; Schlick, T. *J. Chem. Phys.* **1998**, *109*, 1617.
- (21) Sandu, A.; Schlick, T. *J. Comput. Phys.* **1999**, *151*, 74.
- (22) Izaguirre, J.; Reich, S.; Skeel, R. *J. Chem. Phys.* **1999**, *110*, 9853.
- (23) Izaguirre, J.; Catarello, D.; Wozniak, J.; Skeel, R. *J. Chem. Phys.* **2001**, *114*, 2090.
- (24) Qian, X.; Schlick, T. *J. Chem. Phys.* **2002**, *116*, 5971.
- (25) Ma, Q.; Izaguirre, J. *Multiscale Model. Simul.* **2003**, *2*, 1.
- (26) Minary, P.; Tuckerman, M.; Martyna, G. *Phys. Rev. Lett.* **2004**, *93*, 150201.
- (27) Sweet, C.; Petrone, P.; Pande, V.; Izaguirre, J. *J. Chem. Phys.* **2008**, *128*, 145101.

- (28) Stuart, S.; Zhou, R.; Berne, B. *J. Chem. Phys.* **1996**, *105*, 1426.
- (29) Zhou, R.; Harder, E.; Xu, H.; Berne, B. *J. Chem. Phys.* **2001**, *115*, 2348.
- (30) Ewald, P. *Ann. Phys.* **1921**, *64*, 253.
- (31) Frenkel, D.; Smit, B. *Understanding Molecular Simulation*, 2nd ed.; Academic Press: London, 2002.
- (32) Darden, T.; York, D.; Pedersen, L. *J. Chem. Phys.* **1993**, *98*, 10089.
- (33) Essmann, U.; Perera, L.; Berkowitz, M.; Darden, T.; Lee, H.; Pedersen, L. *J. Chem. Phys.* **1995**, *103*, 8577.
- (34) Luty, B.; Tironi, I.; van Gunsteren, W. *J. Chem. Phys.* **1995**, *103*, 3014.
- (35) Greengard, L.; Rokhlin, V. *J. Comput. Phys.* **1985**, *60*, 187.
- (36) Zhou, R.; Berne, B. *J. Chem. Phys.* **1995**, *103*, 9444.
- (37) Figueirido, F.; Levy, R.; Zhou, R.; Berne, B. *J. Chem. Phys.* **1997**, *106*, 9835.
- (38) Procacci, P.; Darden, T.; Marchi, M. *J. Phys. Chem.* **1996**, *100*, 10464.
- (39) Procacci, P.; Marchi, M. *J. Chem. Phys.* **1996**, *104*, 3003.
- (40) Procacci, P.; Marchi, M.; Martyna, G. *J. Chem. Phys.* **1998**, *108*, 8799.
- (41) Procacci, P.; Berne, B. *J. Chem. Phys.* **1994**, *101*, 2421.
- (42) Lau, K.; Alper, H.; Thacher, T.; Stouch, T. *J. Phys. Chem.* **1994**, *98*, 8785.
- (43) Tuckerman, M.; Yarne, D.; Samuelson, S.; Hughes, A.; Martyna, G. *Comput. Phys. Commun.* **2000**, *128*, 333.
- (44) Jorgensen, W.; Chandrasekhar, J.; Madura, J.; Impey, R.; Klein, M. *J. Chem. Phys.* **1983**, *79*, 926.
- (45) Rychaert, J.; Ciccotti, G.; Berendsen, H. *J. Comput. Phys.* **1977**, *23*, 327.
- (46) Our implementation of the quintic spline function has been ported into the official NAMD2 package, version 2.7b (<http://www.ks.uiuc.edu/Research/namd/>, accessed May 2010) and may be activated by setting "longSplitting=c2".
- (47) MacKerell, A.; Bashford, D.; Bellott, M.; Dunbrack, R.; Evanseck, J.; Field, M.; Fischer, S.; Gao, J.; Guo, H.; Ha, S.; Joseph-McCarthy, D.; Kuchnir, L.; Kuczera, K.; Lau, F.; Mattos, C.; Michnick, S.; Ngo, T.; Nguyen, D.; Prodhom, B.; Reiher, W.; Roux, B.; Schlenkrich, M.; Smith, J.; Stote, R.; Straub, J.; Watanabe, M.; Wiorkiewicz-Kuczera, J.; Yin, D.; Karplus, M. *J. Chem. Phys. B* **1998**, *102*, 3586.
- (48) van Gunsteren, W.; Berendsen, H. *Mol. Phys.* **1977**, *34*, 1311.
- (49) Watanabe, M.; Karplus, M. *J. Chem. Phys.* **1993**, *99*, 8063.
- (50) Humphreys, D.; Friesner, R.; Berne, B. *J. Phys. Chem.* **1994**, *98*, 6885.

CT100054K