JOURNAL OF CHEMICAL PHYSICS VOLUME 114, NUMBER 21 1 JUNE 2001

Rapid estimation of electronic degrees of freedom in Monte Carlo calculations for polarizable models of liquid water

Michael W. Mahoney
Department of Physics, Yale University, New Haven, Connecticut 06520

William L. Jorgensen[a]
Department of Chemistry, Yale University, New Haven, Connecticut 06520

(Received 9 August 2000; accepted 14 March 2001)

The inclusion of electronic polarization within Monte Carlo calculations of simple models of molecular liquids is hampered, relative to its inclusion within molecular dynamics calculations, by the need to fully determine the variables that specify the electronic configuration every time each molecule is moved, i.e., N times per cycle, rather than once per cycle. Classical statistical mechanical Monte Carlo calculations on two models of liquid water have been performed. For each of the models, electronic degrees of freedom are modeled by polarizable sites; thus it is the components of the induced dipole vector that must be determined at each step. Commonly used approximation methods have been characterized and found to be inadequate. Efficient procedures have been devised to estimate the dipole vector and have been tested on reproducing electronic, thermodynamic, and structural properties of the two polarizable water models. The most promising procedure, considering both computational time saved and accuracy at reproducing pure liquid properties, involves approximating the induced dipoles at each step by an initial perturbative modification of the dipoles from the previous step, followed by an iteration of the induced dipoles on only the moved molecule. With this procedure, the CPU time is dramatically reduced, and the thermodynamic and structural properties are estimated correctly to within a few percent. They are reproduced more rapidly and with greater accuracy than in calculations in which the dipole vector is estimated by a single iterative cycle starting with the dipoles from the previous Monte Carlo step.
[DOI: 10.1063/1.1370083]

I. INTRODUCTION

Most calculations on polarizable models of molecular liquids, i.e., those models that include an explicit multibody term to describe rapidly responsive electronic polarization, have been performed with molecular dynamics algorithms.[1–50] The inclusion of polarizability in molecular dynamics leads to calculations that are slower, although not prohibitively so.[21,50] Monte Carlo (MC) techniques have also proven successful as a method for evaluating the high dimensional integrals required to characterize the statistical mechanical description of models of molecular liquids.[51–59] These techniques typically use the Metropolis algorithm to construct a Markov chain of accessible configurations[51,60] and are particularly efficient for systems that involve a potential energy function that is pairwise decomposable,[60] i.e., a potential energy function that can be written as the sum of one- and two-body terms. In this case, since Monte Carlo trial moves typically involve moving a single molecule, the calculation of the change in energy between the trial state and the previous state involves only the recalculation of the energy of the moved molecule interacting with every other molecule. This order N, $O(N)$, calculation is rapid. Although some Monte Carlo calculations have been performed on polarizable liquids,[61–83] a disadvantage of Monte Carlo, as compared to molecular dynamics, for such systems is the relatively larger increase in required CPU time due to the lack of pairwise decomposability in the energy terms describing the polarizable degrees of freedom.[75,76]

An exact description of the polarizable degrees of freedom is typically formulated as a matrix equation that must be solved at each step.[3,21,78] The exact solution of this equation corresponds to the adiabatic separation of nuclear and electronic degrees of freedom represented by the Born–Oppenheimer approximation.[24,25,75,76,84,85] Although matrix inversion is an $O(N^3)$ process,[86] the matrices generated for physically realistic molecular configurations can be iterated, an $O(N^2)$ procedure,[87,88] to solution relatively rapidly, i.e., in a number of steps $M \ll N$, where N is the size of the matrix.[3] While different properties exhibit a differential sensitivity to incomplete iteration, in general the iteration must be carried out to high accuracy to ensure that detailed balance is satisfied.[25,32,75,76,89] Since there are N Monte Carlo steps, as opposed to one molecular dynamics step, per cycle, a molecular dynamics algorithm will require $O(N^2)$ steps per cycle for a polarizable model, albeit with a larger prefactor than for a nonpolarizable model, while a Monte Carlo algorithm will require $O(N^3)$ steps per cycle for a polarizable model, as opposed to $O(N^2)$ for a nonpolarizable model.[90]

The computational expense associated with an exact description of polarizable systems has led to the introduction of

---

[a] Electronic mail: william.jorgensen@yale.edu

numerous methods to approximate the variables describing the polarizable degrees of freedom. One class of approximations involves including those variables as dynamical variables within an extended Lagrangian frame work.[19,21,24,30,31,50,75,76] Another class of approximations involves a partial iteration starting from a reasonable initial guess, or the modification of the polarizable degrees of freedom on just the moved molecule at each Monte Carlo step, or the modification of all the polarizable degrees of freedom at a subset of the Monte Carlo steps.[64,65,69,74,78]

Most calculations on polarizable models of molecular liquids have been on models of liquid water. Water is both highly polar and highly polarizable, with both a relatively large gas phase dipole moment[91] and condensed phase induced dipole.[92] In addition, although fixed charge models have been successful at reproducing a wide range of experimental data to high accuracy,[93–96] several of the areas of greatest weakness for fixed charge models, i.e., ion–water interactions,[3–5,36] benzene–water, and ion–benzene interactions,[42,97–99] and the interactions involving dipolar solutes in low dielectric media,[69] involve water. Given this, it is desirable to develop and characterize the realm of applicability of methods for the rapid estimation of the variables that determine the electronic degrees of freedom in Monte Carlo calculations of polarizable models of liquid water. Naturally, there is a tension between the speed and the accuracy of an approximation. Approximations can lead to significant errors in computed thermodynamic quantities.[65] Parameters could be developed around a rapid approximation method,[42] and those approximations may be useful in determining whether polarization is a contributing factor to certain phenomena.[69] Nevertheless, we are interested in the more restrictive class of approximations that approximate the exact solution well, i.e., those approximations that model the adiabatic separation of electronic and nuclear degrees of freedom sufficiently well that errors in electronic, thermodynamic, and structural quantities are minimized.

In the present paper, several approximation methods are described and characterized. Monte Carlo calculations in which the electronic degrees of freedom are determined by one of the approximation methods are compared with calculations in which the electronic degrees of freedom are determined exactly. Since the models used in these calculations represent the variable electronic degrees of freedom with point polarizable sites on a subset of the atomic sites,[1,2] the variables to be estimated are the components of the induced dipole vector. It should be noted, however, that the techniques presented are not limited to this class of models, but only require that each Monte Carlo step involve the trial move of at most a small number of molecules. Thus they are applicable with minor modification to models that represent the variable electronic degrees of freedom in other ways, e.g., by fluctuating charges. In Sec. II the two models used to evaluate the quality of the approximations are described. Section III presents some computational details, and Sec. IV presents the results of the exact calculations. In Sec. V several methods to approximate the electronic variables are presented. Section VI presents preliminary considerations and initial results for several methods, and Sec. VII presents in more detail the results of the most promising methods. Finally, Sec. VIII presents a discussion and conclusion.

TABLE I. Parameters for the Dang93 and Dang97 models.

	Dange93	Dang97
$q_H(e^-)$	0.365	0.519
$q_0(e^-)$	-0.730	
$q_M(e^-)$		-1.038
σ_0(Å)	3.205	3.234
ε_0 (kcal/mol)	0.160	0.1825
α_H (Å^3)	0.170	
α_0 (Å^3)	0.528	
α_M (Å^3)		1.444
r_{OH} (Å)	1.00	0.9572
θ_{HOH} (deg)	109.47	104.5
r_{OM} (Å)		0.215
$\mu^{\mathrm{gas}}(D)$	2.02	1.84
$\mu^{\mathrm{liq}}(D)$[a]	2.60 ± 0.01	2.75 ± 0.05

[a] As calculated by Dang in Refs. 1 and 2.

II. REVIEW OF THE MODELS AND THE THEORY

In order to test the approximation methods, calculations were performed on the two models summarized in Table I. The Dang93[1] model is a three-site model with three polarizable sites; its gas-phase dipole moment is 2.02 D, and the total dipole moment in pure liquid calculations at 303 K is ca. 2.6 D.[1] The Dang97[2] model is a four-site model with permanent charges constructed to reproduce the gas-phase dipole moment of the monomer. It has a single polarizable site on the massless charge site, a polarizability equal to the average molecular polarizability, and a total dipole moment in pure liquid calculations at 25 °C of ca. 2.75 D.[2] Given the computational expense associated with including explicit polarization in liquid phase calculations, it would be desirable if a single polarizable site could be used to model liquid water accurately. Nevertheless, in the absence of the demonstration that such models are generally superior to three-polarizable-site models, the response of both models to the approximations was examined.

The potential part of the model Hamiltonian is given in Eq. (1),[1,2,100]

$$
\begin{aligned}
U&=U_{\mathrm{POL}}+\sum_{a\neq b}U_{ab},\\
U_{ab}&=\sum_{ij}\frac{q_iq_je^2}{r_{aibj}}+4\varepsilon_0\left[\left(\frac{\sigma_0}{r_{a0b0}}\right)^{12}-\left(\frac{\sigma_0}{r_{a0b0}}\right)^{6}\right],\\
U_{\mathrm{POL}}&=-\frac{1}{2}\sum_a \mathbf{p}_a\cdot\mathbf{E}_a^q,
\end{aligned}
\tag{1}
$$

where a and b are dummy indices representing the molecules, i and j are dummy indices representing the atoms, $\mathbf{E}^q$ is the electric field due to the charges, $\mathbf{p}$ is the induced dipole vector, and the remaining variables represent their standard quantities.

The polarization energy, U_{POL}, is calculated in the following manner,[100] Upon expanding the classical electrostatic potential in a multipole series

$$\Phi = \sum_i e_i \Phi(r_i) = \sum_i e_i \Phi(0) + \mathbf{r}_i \cdot \nabla \Phi(0) + \cdots$$
$$= \frac{e_{\mathrm{TOT}}}{r} + \frac{\mathbf{m}\cdot\mathbf{r}}{r^3} + \cdots, \tag{2}$$

where e_{TOT} and $\mathbf{m}$ are the total charge and dipole moment of the charge distribution, and retaining the first two terms, the electric field becomes

$$\begin{aligned}\mathbf{E} &= -\nabla\Phi = -\nabla \frac{e_{\mathrm{TOT}}}{r} - \nabla \frac{\mathbf{m}\cdot\mathbf{r}}{r^3} \\ &= \frac{\mathbf{r}}{r^3} - 3\left(\frac{\mathbf{m}\cdot\mathbf{r}}{r^5}\right)\mathbf{r} - \frac{\mathbf{m}}{r^3} \\ &= \frac{\mathbf{r}}{r^3} - \overleftrightarrow{T}\cdot\mathbf{m},\end{aligned} \tag{3}$$

where the dipole–dipole interaction tensor between two polarizable sites is written as

$$\overleftrightarrow{T} = \frac{3\mathbf{r}\cdot\cdot\mathbf{r}}{r^5} - \frac{1}{r^3} = \frac{1}{r^5}\begin{pmatrix} 3x^2-r^2 & 3xy & 3xz \\ 3xy & 3y^2-r^2 & 3yz \\ 3xz & 3yz & 3z^2-r^2 \end{pmatrix}. \tag{4}$$

This electric field will induce a dipole in a molecule; thus for calculations such as those performed on liquid systems, where there are a set of charges and polarizabilities, $\{q_i, \alpha_i\}$, at positions $\{\mathbf{r}_i\}$, the following dipoles will be induced at those sites:

$$\mathbf{p}_i = \alpha_i \cdot \mathbf{E}_i = \alpha_i \left(\mathbf{E}_i^q + \sum_{j \neq i} \overleftrightarrow{T}_{ij} \cdot \mathbf{p}_j \right). \tag{5}$$

Here, $\mathbf{E}_i^q = \Sigma_{j\neq i}(q_j \mathbf{r}_{ij}/r_{ij}^3)$ is the electric field at the ith site due to permanent charges at other sites, and $\overleftrightarrow{T}_{ij} = (3\mathbf{r}\cdot\cdot\mathbf{r}/r^5) - (1/r^3)$ is the dipole–dipole interaction tensor between the ith site and the jth site. In the calculation of $\mathbf{E}^q$ and $\overleftrightarrow{T}$, intramolecular contributions are set to zero,[1,2] and interactions beyond a cutoff value are truncated, as was performed by Dang.[1,2] Equation (5) can be rearranged as the following matrix equation

$$\begin{pmatrix} \alpha_1^{-1} & -T_{12} & \cdots & -T_{1N} \\ -T_{21} & \alpha_2^{-1} & & \vdots \\ \vdots & & \ddots & \vdots \\ -T_{N1} & \cdots & \cdots & \alpha_N^{-1} \end{pmatrix} \begin{pmatrix} p_1 \\ p_2 \\ \vdots \\ p_N \end{pmatrix} = \begin{pmatrix} E_1^q \\ E_2^q \\ \vdots \\ E_N^q \end{pmatrix}. \tag{6}$$

For realistic configurations of molecular liquids, the dominant elements of this matrix are the diagonal elements and, thus, Eq. 6 may be solved for $\mathbf{p}$ by iteration.[87,88] A high tolerance is achieved in a number of iterative steps $M \ll N_{\mathrm{mat}}$, where N_{mat} is the size of the matrix and is given by $N_{\mathrm{mat}} = 3N_{\mathrm{psite}}N_{\mathrm{mol}}$, where N_{psite} is the number of polarizable sites on a given molecule and N_{mol} is the number of molecules in the system. It is the rapid estimation of the vector $\mathbf{p}$ at each step in the Monte Carlo calculations that is our focus.

III. COMPUTATIONAL DETAILS

All calculations were run under conditions as close to those in the original papers as possible. For the Dang97 model,[2] the system consisted of 550 molecules in the isothermal-isobaric (*N-P-T*) ensemble at 25 °C and 1 atm, and for the Dang93 model,[1] the system consisted of 216 molecules in the *N-P-T* ensemble at 303 K and 1 atm. For both models, a molecule–molecule based truncation procedure was employed at an intermolecular separation of $r_{\mathrm{OO}} = 9.0$ Å.[1,2] No long range correction for the electrostatics was included. The merits of various methods for dealing with boundary conditions are important issues, but the purpose of the present work is to analyze methods for rapidly approximating polarizable degrees of freedom and not to deal with boundary condition issues. Thus the procedure of Dang is followed.[1,2,101]

The calculations were performed with a modified version of the Monte Carlo program BOSS, version 3.8.[102] Code was added to calculate the polarization matrix and electric field at each step. On solvent moves, this could be performed with an $O(N)$ calculation of the change in the polarization matrix and the electric field between successive steps. Upon examination of several iterative packages, it was found that for the most expensive calculations an iterator[103] using either a conjugate gradient method or one of several methods derived therefrom yielded the fastest calculations. In addition, Jacobi preconditioning[87,103] was performed since it improved the convergence rate, and a matrix vector multiplication routine was written to take advantage of the symmetry and relative sparsity of the polarization matrix. The iteration began with an initial guess for the solution, generally taken to be the dipole vector from the previous step, and continued until $\delta X \leqslant \mathrm{TOL}$, where $\delta X = |b - A\cdot x_i|/|b|$, x_i is the estimate of the induced dipole vector at the ith step of the iteration, and TOL is a user specified tolerance.[86–88,103] This residual based criterion is more robust[87,103] than the typical criterion of requiring the rms difference between the dipole vector at successive steps in the iteration to be less than some tolerance,[3,27,78] but it was found that for typical liquid state configurations, both criteria yielded identical results with an appropriate choice for the respective tolerances.

The approximation methods described in the next section are inapplicable to volume moves, i.e., when the dimensions of the periodic cell are changed and the positions of the molecular centers are scaled accordingly in *N-P-T* calculations. Thus volume moves have been performed in the following manner: first, iterate the matrix equation for the induced dipoles from the previous step to an exact solution; then, perform the volume move, including a recalculation of the polarization matrix and the electric field vector; and finally, iterate the polarization matrix for the induced dipoles at that move to an exact solution. Handling volume moves asymmetrically by, e.g., performing a full iteration only after the move would lead to a disproportionate number of volume moves being accepted. The failure to fully minimize the dipoles at every solvent step leads to a systematic drag effect, which the approximations attempt to minimize. An analogous effect caused by volume moves would be much less tractable due to the sensitive dependence of the energy on the volume and to the relatively long runs required for volume equilibration.[57] One may claim that the calculations should be implemented in the *N-V-T* ensemble since the volume moves ''reset'' the electronic degrees of freedom. *N-P-T* calculations were performed since the *N-P-T* ensemble is the ensemble of primary experimental interest and since the original calculations of Dang were performed in that ensemble. Nonconverged *N-V-T* calculations[104] indicate that the results presented here for the quality of the approximations are not sensitive to the frequency of volume moves; thus it is expected that similar conclusions hold for *N-V-T* calculations.

For both models, the exact calculation began with a box of nonpolarizable molecules with the same geometry and Lennard-Jones parameters, the charges of which were adjusted to give approximately the correct energy and density. Initial conditions for the other calculations came from a partially equilibrated box in which the dipoles were determined either exactly or with a single full iterative cycle. In all calculations, moves involved either the rotation and translation of a single molecule or the uniform scaling of the volume. For solvent moves, the magnitude of the translations and rotations were adjusted to yield acceptance rates of ca. 40%. It was difficult to achieve such high acceptance rates for the volume moves, in particular for the one-polarizable-site model. Smaller volume moves were required than for similarly sized nonpolarizable systems, and acceptance rates for volume moves were ca. 20% to 30%. Error estimates were calculated with the means batch procedure.[60] Note that since this procedure assumes that the underlying time series is stationary and Gaussian, for some of the less well converged calculations discussed below the quoted errors are likely underestimated.

IV. RESULTS OF EXACT MONTE CARLO CALCULATIONS

The energy and density for both the Dang93 and Dang97 models as a function of length along the Markov chain are presented in Figs. 1 and 2. Averages are presented in Table II along with the original results of Dang.[1,2] Two features of the exact calculations deserve particular mention.

First, both models require an equilibration period significantly longer than for calculations on similar nonpolarizable models. To examine this, Figs. 1 and 2 also present the energy and density for 550 molecules of TIP4P water[56] run at 25 °C and 1 atm using a protocol that was similar to that for the polarizable calculations. The initial condition was an arbitrary configuration that was worse, i.e., farther from equilibrium, than the initial conditions for the polarizable calculation, and the results presented are typical of nonpolarizable models. The number of Monte Carlo steps required for equilibration is clearly much larger for models with polarizable degrees of freedom. Figures 3 and 4 present the polarizable and nonpolarizable components of the energy for the Dang93 and Dang97 models, respectively, and show that the dominant contribution to the slow rate of convergence is the slow relaxation of the polarizable degrees of freedom. There is also a coupling between the polarizable and nonpolarizable components of the energy and so, e.g., the nonpolarizable component of the energy relaxes more slowly for the smaller box of 216 molecules with the three-polarizable-site Dang93 model than for 550 molecules with the one-polarizable-site Dang97 model. The slow relaxation is no doubt related to the method of generating the Markov chain by performing single particle moves. This method appears to be relatively disadvantageous for systems with multiparticle potentials, as compared to systems with pairwise decomposable potentials.[105] Thus in addition to the extra time at each step required to determine the values of the variables specifying the electronic degrees of freedom, polarizable models require more steps to equilibrate.

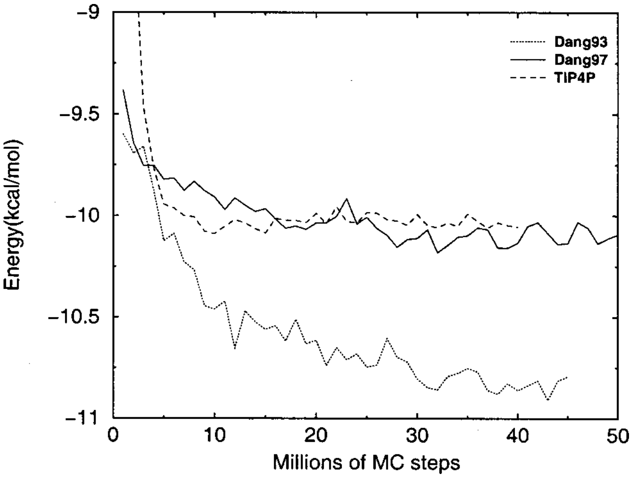

FIG. 1. Intermolecular energy for the Dang93 and Dang97 models as a function of length along the Markov chain. Averages are taken over 1 mil MC steps. Also presented are the results for the nonpolarizable model TIP4P.

Second, although Table II indicates that the results for the Dang97 model agree with the original results[2] to within ca. 1% for the density and ca. 2.5% for the energy, there is substantial disagreement between the values calculated for the Dang93 model and the values originally reported.[1,101] This disagreement appears to be due to the increased sensitivity of small polarizable systems to changes in such param-

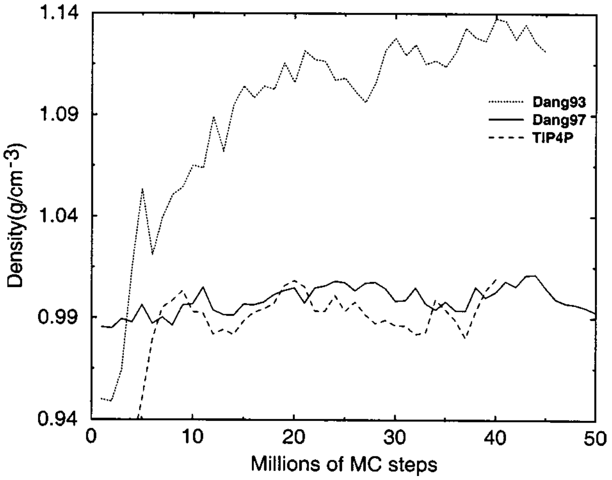

FIG. 2. Density for the Dang93 and Dang97 models as a function of length along the Markov chain. Averages are taken over 1 mil MC steps. Also presented are the results for the nonpolarizable model TIP4P.

TABLE II. Thermodynamic properties calculated for the Dang93[a] and Dang97[b] models at 303 K and 25 °C, respectively, and 1 atm.

	Length of average[c]	Energy (kcal/mol)	Density (g/cm^3)	Dipole (D)
Dang93[a] MD		-9.84 ± 0.08	1.003 ± 0.015	2.62
Dang93, exact MC	15	-10.82 ± 0.02	1.125 ± 0.003	2.68
Dang97[b] MD		-9.84 ± 0.07	0.995 ± 0.007	2.75
Dang97, exact MC	20	-10.10 ± 0.01	1.001 ± 0.001	2.78

[a] Reference 1.
[b] Reference 2.
[c] Millions of MC steps.

eters as the number of molecules and the intermolecular cutoff distance; this is discussed in more detail in Appendix A.

V. METHODS TO RAPIDLY ESTIMATE THE INDUCED DIPOLE VECTOR

Previous work[76,78] and our initial results indicate that any reasonable $O(N^2)$ method for estimating the induced dipole vector, e.g., a single full iteration, would still be prohibitively expensive for large scale computation on polarizable models of molecular liquids. Attention was thus focused on developing $O(N)$ methods, three examples of which are described.

A. Iteration over many iterations

When a given molecule moves, it is expected that the set of molecules that experience the largest change in induced dipoles will be strongly correlated with the set of molecules near the moved molecule. This intuition has been confirmed by an analysis, not presented in detail, of the rms difference between induced dipoles at successive steps.[104] In addition, the analysis indicates that molecules that cross the 9 Å cutoff at which nonbonded interactions are truncated also have a tendency to experience a change in induced dipole similar in magnitude to the change experienced by molecules within ca. 4 Å of the moved molecule.[104] The simplest approximation using this idea is to modify the dipole on only the moved molecule.[77] A more expensive suggestion to modify the dipoles on the moved molecule and those molecules that will likely experience a large change in induced dipole, i.e., those molecules near the moved molecule, where ‘‘near’’ may be defined to be those molecules with an r_{OO} less than some value, e.g., 4 Å, and perhaps those molecules that cross the truncation boundary. A still more expensive suggestion, which is not $O(N)$, is to modify the dipoles on the moved molecule and those molecules near the moved molecule and then to modify the dipoles on each of the molecules near each of those molecules, etc., iteratively. In the limit of a large number of these cycles with the iterations performed to a high tolerance, this leads to an exact solution but does so more slowly than using an efficient iterator.[104]

B. Perturbation theoretic estimation

Another class of approximation methods, which to our knowledge has never been presented, seeks to approximate the solution of the matrix equation at a given step as a perturbation of the matrix equation from the previous step. If the solution of the matrix equation $A \cdot x = b$ in Eq. (6) is known

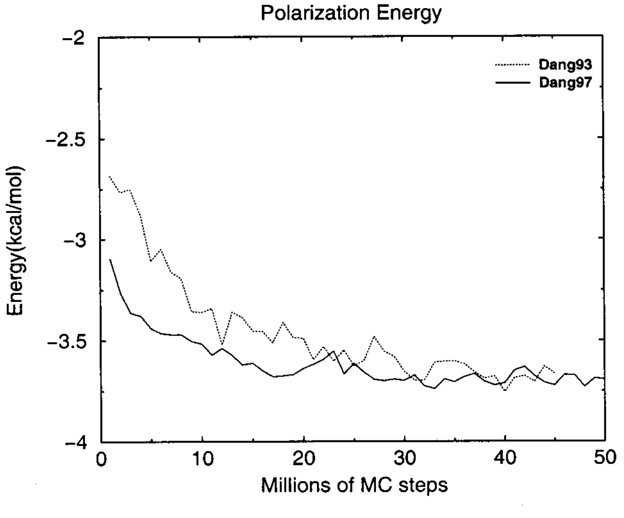

FIG. 3. Polarizable component of the intermolecular energy for the Dang93 and Dang97 models as a function of length along the Markov chain. Averages are taken over 1 mil MC steps.

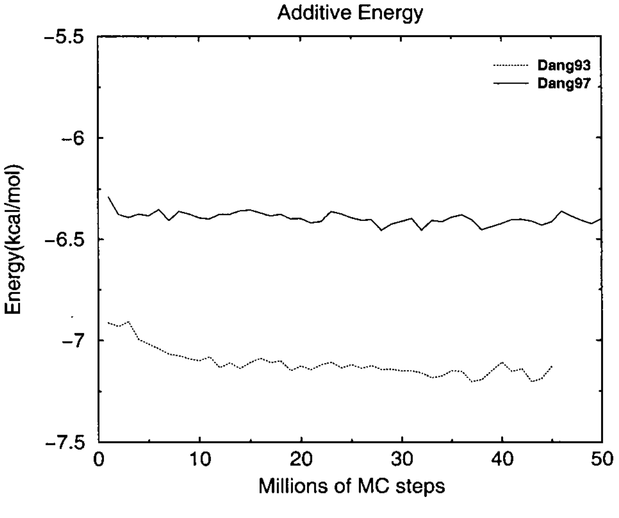

FIG. 4. Nonpolarizable component of the intermolecular energy for the Dang93 and Dang97 models as a function of length along the Markov chain. Averages are taken over 1 mil MC steps.

exactly or approximately at the ith step, then the matrix equation at the $(i+1)$th step may be written as $A'\cdot x'=b'$, where

$$A'=A+\delta A=A+\delta A_{\mathrm{rows}}+\delta A_{\mathrm{cols}},$$
$$b'=b+\delta b,\ \text{and} \tag{7}$$
$$x'=x+\delta x.$$

For Monte Carlo moves that consist of attempting to move a single molecule, i.e., for solvent as opposed to volume moves, the matrix A' differs from the matrix A along a set of $3N_{\mathrm{psite}}$ rows and $3N_{\mathrm{psite}}$ columns. Defining for an arbitrary square matrix A,

$$A\equiv D+N, \tag{8}$$

where D is a diagonal matrix consisting of the diagonal elements of A and N contains the off-diagonal elements of A, with zeros along the diagonal, then $A'\cdot x'=b'$ can be rewritten, after subtracting $A\cdot x=b$ from both sides, as

$$D\cdot\delta x+N\cdot\delta x+\delta A_{\mathrm{rows}}\cdot\delta x+\delta A_{\mathrm{cols}}\cdot\delta x$$
$$=\delta b-\delta A_{\mathrm{rows}}\cdot x-\delta A_{\mathrm{cols}}\cdot x. \tag{9}$$

The operations on the right-hand side of Eq. (9) are $O(N)$, as is multiplying by the diagonal matrix D. Ignoring the $\delta A_{\xi}\cdot\delta x$ terms and the $N\cdot\delta x$ term, which is small relative to the $D\cdot\delta x$ term, and performing the $O(N)$ inversion of the diagonal matrix D yields an estimate for the change in induced dipoles from one step to the next:

$$\delta x=D^{-1}(\delta b-\delta A_{\mathrm{rows}}\cdot x-\delta A_{\mathrm{col}x}\cdot x). \tag{10}$$

This approximation is the lowest order term in an expansion of the inverted matrix. Within this approximation, the change in the induced dipoles at each step is caused by each dipolar site interacting with the change in local electric field due to the permanent charges and with the change in dipoles elsewhere, as they existed at the previous step. The changes do not back interact to self-consistency. If they did, an exact solution to the matrix equation would be achieved. Keeping higher terms in the expansion would correspond to more back interactions, but would not be $O(N)$.

C. Use of a smaller cutoff criteria or a weaker tolerance

A third class of approximation methods involves the use of a smaller cutoff criterion for the dipole–dipole interaction than for the other interactions.[66] This is similar in spirit to multiple time step methods[98,99] in molecular dynamics in which the interactions are divided between slowly changing interactions and the more rapidly changing interactions that generally occur between molecules that are spatially proximate. The dipole–dipole interaction energy between spatially proximate molecules is recalculated at every step, and the more distant interactions are recalculated less frequently. The more slowly changing interactions are included as an additive cutoff correction to the energy. They are added to each step and recomputed by doing a full iteration to a high tolerance every $M\approx N_{\mathrm{mol}}$ steps, taking the difference between the energy estimates before and after the full iteration at that step as the correction. Thus at the steps at which the cutoff correction is recomputed, both of the following calculations are performed:

$$A_S\cdot x_S=b_S, \tag{11}$$

and

$$A_B\cdot x_B=b_B, \tag{12}$$

where A_ξ, x_ξ, and b_ξ are the polarization matrix, the induced dipole vector, and the electric field due to the charges, respectively, corresponding to the bigger and smaller cutoff for $\xi=B,S$, respectively. These correspond to two versions of Eq. (6) in which two different cutoffs are used. Equation (12) can be rewritten as

$$(A_S+\delta A)\cdot(x_S+\delta x)=b_S+\delta b,\ \text{or}$$
$$A_S\cdot x_S+A_S\cdot\delta x+\delta A\cdot x_S+\delta A\cdot\delta x=b_S+\delta b. \tag{13}$$

Note that these difference matrices are not the same as those in Sec. V B. The difference in the polarization component of the intermolecular energy using the two cutoffs is

$$\Delta U=-\tfrac{1}{2}x_B\cdot b_B--\tfrac{1}{2}x_S\cdot b_S. \tag{14}$$

Adding the cutoff correction directly to the energy leads to estimates for the induced dipole vector that are incorrect. An alternative formulation of this method would be to store not the energy difference but the vectors of differences of dipoles and electric fields, i.e., δb and δx. Then the following energy would be added at each step:

$$\Delta U=-\tfrac{1}{2}x_B\cdot b_B--\tfrac{1}{2}x_S\cdot b_S$$
$$=-\tfrac{1}{2}(x_S+\delta x)\cdot(b_S+\delta b)+\tfrac{1}{2}x_S\cdot b_S,$$

or, upon performing the algebra,

$$\Delta U=-\tfrac{1}{2}(x_S\cdot\delta b+\delta x\cdot b_S+\delta x\cdot\delta b). \tag{15}$$

In this case, the x_S and b_S are stored from the previous update step, although it may be possible to achieve a better estimate of the energy difference by using those values calculated at each step.

With box sizes and cutoffs of typical sizes, approximately half of the components of the polarization matrix are identically zero. Thus an additional advantage to imposing a smaller cutoff criterion would be that the polarization matrix would be sparser, leading to a larger relative advantage for matrix-vector multiplication routines written to take advantage of the sparsity of the matrix. For such a method to be effectively $O(N)$ the cutoff could have to be quite small. For example, decreasing the cutoff from 9 to 6.5 or 4.5Å decreases the number of nonzero elements by approximately 62% and 88%, respectively. Thus only the most drastic cutoffs lead to significant computational savings.

VI. PRELIMINARY CONSIDERATIONS AND INITIAL RESULTS

In order to determine the appropriate tolerance for use in the exact calculations, a series of calculations was performed in which iteration continued until $\delta X\leqslant\mathrm{TOL}$, where $\mathrm{TOL}=10^{-n}$, with $n=1,2,\ldots$. These calculations indicate that for a typically sized system of either Dang93 or Dang97 molecules run under typical conditions for tens of thousands of steps, the energies and densities varied in a systematic way as n was increased. The Markov chains with $n=5$ and $n=6$ were identical to those with larger n for several thousand steps and several tens of thousands of steps, respectively. At the end of the latter calculations, the components of the induced dipole vector differed in the eighth significant digit relative to the dipole vectors in similar calculations with larger n. Although the chain with $n=6$ will ultimately diverge from chains with larger n, the identity of the chains and the similarity of their dipole vectors provide confidence that differences between the fully iterated solution and the true solution are minute.

Since the approximation methods are designed to be rapid, calculations were performed to estimate relative times and to eliminate those methods that did not yield significant time savings. The relative times for methods of estimating the induced dipole vector are dependent on the number of molecules and the frequency of volume moves. Timings relative to nonpolarizable code depend on the efficiency of the nonpolarizable code as well as the efficiency of the polarizable code not involved in the determination of the induced dipole vector. Due to these considerations and since these factors differ widely between different codes, our actual timing results for each of the approximation methods are not presented. However, the timing studies, a summary of which is presented in Appendix B, indicate that the only approximation methods that are not too computationally expensive are those that involve an iteration over the moved molecule and its neighbors and/or an initial perturbative modification. (Methods that involve multiple cutoffs with very short inner cutoffs also have acceptable timing but were otherwise unacceptable as described below.) The timing studies also indicate that the most promising approximation method, i.e., ''perturb_nbrit_0A,'' has a time speedup, relative to the exact calculation, of a factor of ca. 70 for systems consisting of 216 molecules of the Dang93 model or 550 molecules of the Dang97 model. In this case, the CPU time increase for a polarizable calculation, relative to a nonpolarizable calculation, is about an order of magnitude for the three-polarizable-site model and less than an order of magnitude for the one-polarizable-site model.

The approximation methods were designed to approximate the exact solution. A necessary condition for this is that the dipoles in an approximation method stay close to the exact dipoles that represent the adiabatically separated Born–Oppenheimer surface. To examine this, a series of short calculations was performed using a variety of approximation methods in which at each step the exact induced dipole vector was also calculated. The sequence of states was identical to the sequence obtained when the exact calculation was not performed since the exactly calculated dipoles were stored separately. Figure 5 presents, as a function of step number along the calculation, the rms difference between the two vectors for representative calculations on Dang93. Similar results were obtained from calculations on the Dang97 model and are not presented. The results for the case where the dipoles were determined from a single full iteration are presented for reference. It is notable that the use of a perturbative modification of the dipoles at each step leads to dipoles that fail to stay near the Born–Oppenheimer surface reliably. In addition, the dipoles fail to stay near the Born–Oppenheimer surface reliably when they are estimated by updating the dipole on only the moved molecule and also when the dipoles are updated with the use of multiple cutoffs with the cutoff correction recomputed every ca. N_{mol} steps, although for clarity these results are not presented in the figure. Methods for which the dipoles fail to stay near the Born–Oppenheimer surface reliably were eliminated as possible approximation methods, although some were used as control calculations.

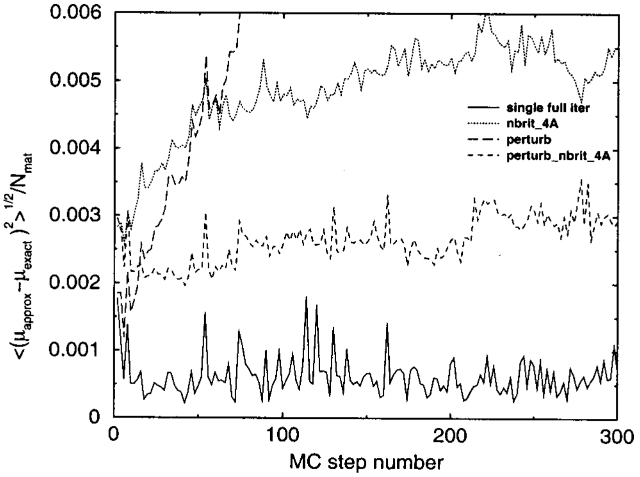

FIG. 5. The induced dipole drift as estimated by the rms difference between approximated dipole vector and the corresponding fully iterated dipole vector during short Monte Carlo calculations on 216 molecules of Dang93 water. The two vectors were stored separately; thus the Markov chain generated is identical to that generated with the approximation by itself, i.e., without the associated exact calculation. N_{mat} is the length of the vector of induced dipoles. See the beginning of Sec. VII for a description of the approximation methods and a definition of the corresponding labels.

When the dipoles are determined by iterating over the molecules whose dipoles are likely to change significantly (''nbrit_4A''), i.e., those on the moved molecule, on molecules within 4 Å of the moved molecule, and those molecules that cross the truncation cutoff, the rms difference is larger than with a single iteration, and it is unclear whether the approximated dipoles stay near the Born–Oppenheimer surface for runs much longer than those presented. When this iteration is coupled with an initial perturbative modification (''perturb_nbrit_4A''), the rms difference is decreased and the dipoles do stay near the Born–Oppenheimer surface. In addition, when coupled with an initial perturbative modification, the iteration over only the moved molecule, i.e., ''perturb_nbrit_0A,'' leads to dipoles that stay near the Born–Oppenheimer surface, although for clarity this is not presented. These latter two methods have turned out to be the most promising. When the iteration over the moved molecule and its neighbors within 4 Å is followed by an iteration over the neighbors of each of those neighbors, the rms difference is reduced relative to iterations over only the moved molecule and its neighbors. Nevertheless, pure liquid calculations indicate that the accuracy of the calculated thermody-

TABLE III. Thermodynamic properties calculated for the Dang97[a] model at 25 °C and 1 atm.

	Length of average[b]	Energy (kcal/mol)	Density (g/cm^3)	Dipole (D)
exact calculation	20	-10.10 ± 0.010	1.001 ± 0.001	2.78
one full iterative cycle	20	-9.632 ± 0.002	1.014 ± 0.001	2.70
perturb_nbrit_4A	50	-9.771 ± 0.004	0.992 ± 0.001	2.73
perturb_nbrit_0A	50	-9.735 ± 0.004	1.010 ± 0.001	2.72
nbrit_4A	25	-9.450 ± 0.005	0.962 ± 0.001	2.69
nbrit_0A	15	-9.068 ± 0.006	1.067 ± 0.001	2.63
perturbative	15	-9.823 ± 0.006	1.076 ± 0.001	2.74

[a] Reference 2.
[b] Millions of MC steps.

namic properties is not improved enough to justify the increased computational expense.[104] In addition, several values of the neighbor cutoff distance other than 0 and 4 Å were examined. Above 4 Å the calculations become too expensive because too many molecules are included in the iteration and below 4 Å a slight but systematic decrease in the quality of the results is seen.[104]

VII. ANALYSIS OF THE THERMODYNAMIC AND STRUCTURAL RESULTS FOR THE APPROXIMATIONS

In addition to Monte Carlo calculations for both the Dang93 and Dang97 models in which the dipoles were determined exactly, calculations on both models were performed in which the induced dipole vector was approximated in one of several ways.

The approximation methods were to start with the dipole vector from the previous step and to:

(1) perform a single full iterative cycle (''single_full_iter'');
(2) modify the dipoles on those molecules whose dipoles were most likely to change significantly, i.e., the moved molecule, molecules within 4 Å of the moved molecule, and molecules that crossed the cutoff after an initial perturbative modification (''perturb_nbrit_4A'');
(3) modify the dipoles on those molecules whose dipoles were most likely to change significantly, i.e., the moved molecule, molecules within 4 Å of the moved molecule, and molecules that crossed the cutoff (''nbrit_4A'');
(4) modify the dipoles on the moved molecule after an initial perturbative modification (''perturb_nbrit_0A'');
(5) modify the dipoles on the moved molecule (''nbrit_0A''); or
(6) perform a perturbative modification (''perturb'').

A. Calculations on the Dang97 model

The results calculated for the energy, density, and total dipole moment for 550 molecules of the Dang97 model are presented in Table III. Other thermodynamic quantities such as the isobaric heat capacity, the isothermal compressibility, and the thermal expansivity were calculated but are not presented since they converge more slowly.[57] Figures 6 and 7 present the energy and density as a function of Markov chain length, in which the averages are taken in batches of 1 mil MC steps. The exact calculations appear to be sufficiently well converged for reliable conclusions to be drawn. Performing a single full iteration leads to an energy that is ca.

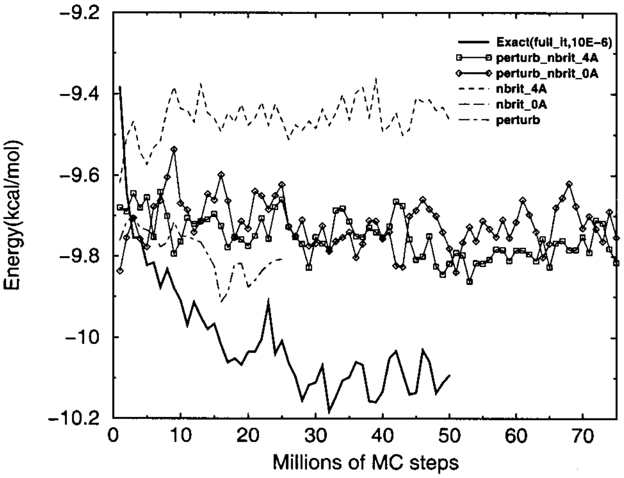

FIG. 6. Intermolecular energy for the Dang97 model at 25 °C and 1 atm for several methods of approximating the induced dipoles. Note from Table III that the results for ''nbrit_0A'' are too high to appear on the figure. See the beginning of Sec. VII for a description of the approximation methods and a definition of the corresponding labels.

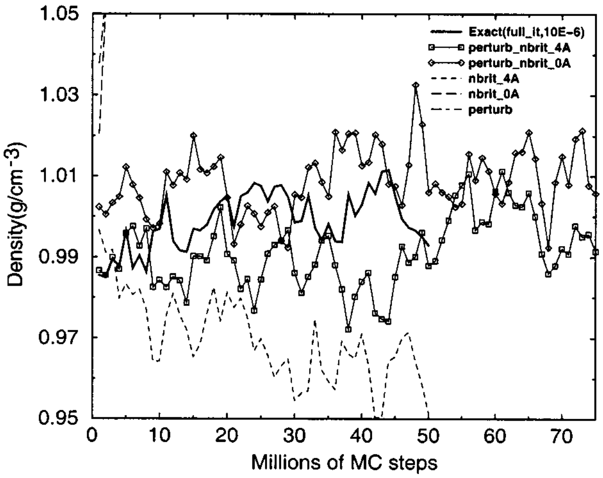

FIG. 7. Density for the Dang97 model at 25 °C and 1 atm for several methods of approximating the induced dipoles. See the beginning of Sec. VII for a description of the approximation methods and a definition of the corresponding labels.

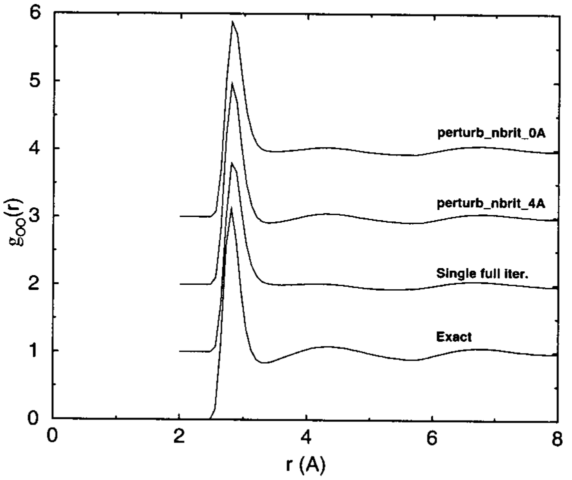

FIG. 8. OO radial distribution function for the Dang97 model at 25 °C and 1 atm for several methods of approximating the induced dipoles.

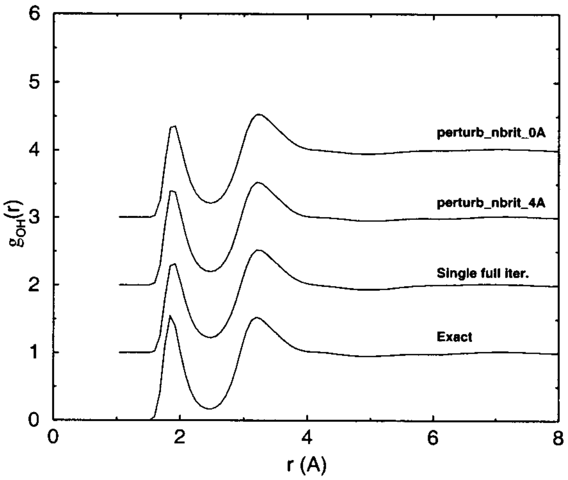

FIG. 9. OH radial distribution function for the Dang97 model at 25 °C and 1 atm for several methods of approximating the induced dipoles.

5% too high and a density that is ca. 1% too high. The most promising rapid approximation methods are ''perturb_nbrit_4A'' and ''perturb_nbrit_0A.'' For the former, the density is fairly well reproduced, being ca. 1% too low, and the energy is almost as well reproduced, being ca. 3% too high. For the latter, results are almost identical; the density is ca. 1% too high and the energy is 3% to 4% too high. These approximation methods are very rapid and lead to estimates for the energy and density that are more accurate than those achieved with a single full iteration in much less time.

As a control, the results are also presented for calculations in which the approximation involves only a perturbative modification or only a modification over the moved molecule. The density for these two methods is ca. 7% greater than the density for the exact solution. Interestingly, although the density for these two methods is quite similar, the energies are very different. Modifying only the dipoles on the moved molecule leads to energies which are ca. 10% too high, while the perturbative method leads to energies which are only about 3% too high. Clearly, these two control methods are inferior to methods for which the dipoles stay near the Born–Oppenheimer surface.

The approximation methods tend to smooth out peaks and valleys in the radial distribution functions (RDFs), although the positions of the peaks and minima are not altered. This is most prominent in the OO RDF, presented in Fig. 8. The methods ''perturb_nbrit_4A'' and ''perturb_nbrit_0A'' clearly yield a second peak in their OO RDF. For the latter, this is reduced slightly more, relative to the exact calculations, than for the former. Both are superior to a single full iterative cycle, for which structure beyond the first peak is largely lost. Although less prominent, similar effects occur in the OH and HH RDFs presented in Figs. 9 and 10. The energy distributions and energy pair distributions are not presented but indicate similar relative destructuring effects. For example, there is a slight tendency towards a less pronounced hydrogen bonding peak in the pair distributions among approximation methods. Although this is more severe for a single full iterative cycle than for either ''perturb_nbrit_4A'' or ''perturb_nbrit_0A,'' the characteristic hydrogen bonding peak[56] is present for all the methods.[104]

B. Calculations on the Dang93 model

The results calculated for energy, density, and total dipole moment for 216 molecules of the Dang93 model using various approximation procedures are presented in Table IV. Figures 11 and 12 present the energy and density as a function of Markov chain length, in which the averages are taken in batches of 1 mil MC steps.

Comparison of these figures with the corresponding figures for the Dang97 model indicates that the trends noted above regarding the quality of various approximation methods are reproduced for the Dang93 model. The trends are less dramatic since the three-polarizable-site model has more flexibility in its electronic degrees of freedom with which to absorb environmental changes and since the individual elements of the dipole vector are significantly smaller than with the Dang97 model. The Dang97 model, with its larger dipoles and larger changes in dipoles from step to step, seems

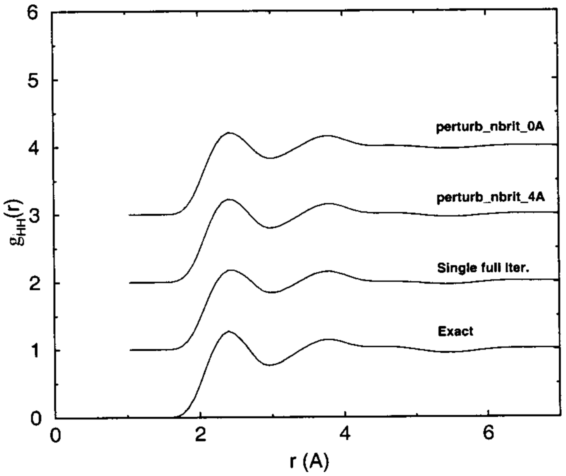

FIG. 10. HH radial distribution function for the Dang97 model at 25 °C and 1 atm for several methods of approximating the induced dipoles.

TABLE IV. Thermodynamic properties calculated for the Dang93[a] model at 303 K and 1 atm.

	Length of average[b]	Energy (kcal/mol)	Density (g/cm^3)	Dipole (D)
exact calculation	15	-10.82 ± 0.02	1.125 ± 0.003	2.68
one full iterative cycle	12	-10.44 ± 0.01	1.105 ± 0.001	2.65
perturb_nbrit_4A	50	-10.72 ± 0.01	1.124 ± 0.001	2.68
perturb_nbrit_0A	50	-10.54 ± 0.01	1.131 ± 0.001	2.67
nbrit_4A	25	-9.95 ± 0.01	1.062 ± 0.002	2.61
nbrit_0A	15	-9.91 ± 0.01	1.156 ± 0.001	2.59
perturbative	15	-10.75 ± 0.01	1.168 ± 0.001	2.68

[a] Reference 1.

[b] Millions of MC steps.

more amenable to the types of approximation methods examined here. Definitive statements on precisely how far the approximation methods are from the exact results are hampered by the poor convergence of the exact results for the Dang93 model.

As with the Dang97 model, the most promising rapid approximation methods are ‘‘perturb_nbrit_4A’’ and ‘‘perturb_nbrit_0A.’’ The density for the former is well reproduced and the energy is ca. 1% too high. Again, the results for ‘‘perturb_nbrit_0A’’ are almost identical to those for ‘‘perturb_nbrit_4A.’’ The density is ca. 1% too high and the energy is ca. 3% too high. In addition, ‘‘perturb_nbrit_0A’’ yields results that are similar to those with a single full iteration at a time savings relatively more significant than for the Dang97 model. Results presented for the control calculations indicate that they are inferior, although not as relatively inferior as for the Dang97 model. The densities for the two methods ‘‘perturb’’ and ‘‘nbrit_0A’’ are ca. 4% too high. As with the Dang97 model, although the densities for these two methods are quite similar, the energies are very different. Modifying only the dipoles on the moved molecule leads to energies that the approximately 9% too high, while the perturbative method leads to energies that are ca. 1% too high. If the iteration is performed over the moved molecule and those molecules near it, the error in the computed energy is 9% and the error in the density is 6%, although, as before, the sign of the density error is opposite that of an iteration over only the moved molecule.

VIII. DISCUSSION AND CONCLUSION

The importance of developing theoretically rigorous algorithms for rapidly determining the variables that specify the electronic degrees of freedom and thus for rapidly calculating properties of interest should be emphasized. *Ad hoc* measures, while applicable to particular cases, will not be generally applicable, and if there exists no method to systematically improve an *ad hoc* set of approximations, they will be limited to the particular cases for which they have been tested. Results presented, as well as other results not presented, indicate that naively making an approximation less drastic does not necessarily lead to improvement in the reproduction of thermodynamic and structural quantities.

Several methods have been presented and characterized to rapidly estimate the vector of induced dipoles in classical statistical mechanical Monte Carlo calculations of polarizable models of liquid water. Although these models represent electronic degrees of freedom with variable dipoles, the techniques are not limited to systems that model the variable

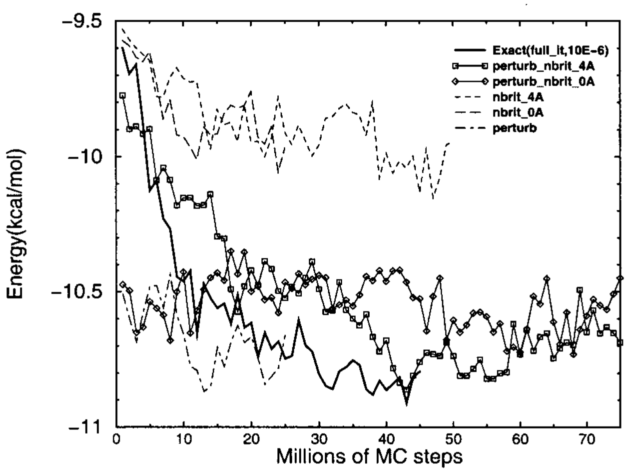

FIG. 11. Intermolecular energy for the Dang93 model at 303 K and 1 atm for several methods of approximating the induced dipoles. See the beginning of Sec. VII for a description of the approximation methods and a definition of the corresponding labels.

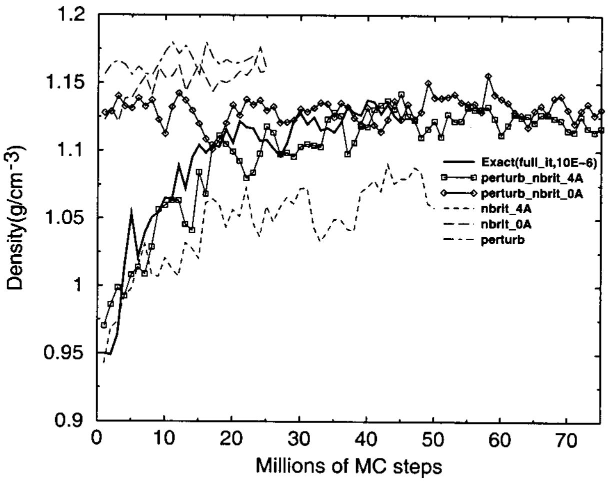

FIG. 12. Density for the Dang93 model at 303 K and 1 atm for several methods of approximating the induced dipoles. See the beginning of Sec. VII for a description of the approximation methods and a definition of the corresponding labels.

electronic degrees of freedom in this way. They are applicable, for example, to systems modeling polarization with a set of fluctuating charges, although the quality of the particular approximations would have to be tested in these other cases. These methods require that successive steps in the Monte Carlo calculation be generated in such a way that only $O(N)$ elements of the polarization matrix change from step to step, e.g., by moving only a single molecule in a trial step.

The most promising approximation method, considering both computational efficiency and the quality of reproduction of thermodynamic and structural properties of the liquid relative to the exact results, was labeled ''perturb_nbrit_0A'' in the text and figures. This method involves first performing an initial perturbative modification of the dipoles from the previous step and then modifying just the dipoles on the moved molecule, both in accordance with the procedures describes in Sec. V. Note that the method ''perturb_nbrit_4A'' performed slightly better but was also slightly more computationally expensive.

For these methods, the calculated energy and density in the *N*-*P*-*T* ensemble differed from the exact results by only a few percent and the radial distribution functions showed minor deterioration. The results were in better agreement with the results of the exact calculations than were the results of calculations in which the induced dipoles were estimated with a single iterative cycle, using the dipole vector from the previous step as an initial guess. The savings in CPU time are nearly two orders of magnitude for the systems considered; the exact factor depends on the number of molecules, the number of polarizable sites per molecule, and on the frequency of volume moves. The importance of minimizing the CPU time required for estimation of the electronic degrees of freedom in these polarizable calculations is particularly important in light of the longer equilibration period for polarizable models relative to similar nonpolarizable models.

It is interesting to note that the calculated energies, densities, radial, and energy distribution functions for ''perturb_nbrit_4A'' and ''perturb_nbrit_0A'' are superior to those from a single full iterative cycle. This is to be contrasted with results that indicate that the dipoles for the latter method are closer to the exact dipoles in an rms sense. Thus while being close to the exact dipoles is a necessary condition for an acceptable approximation method, it does not follow that methods that have dipoles that are closer to the exact dipoles will yield superior results for bulk computed quantities. In particular, when a molecule moves, the dipoles near it will likely change significantly, and it is the values of these changes that are the most important determinants of computed properties. The ''drag'' introduced by a failure to fully iterate the dipole to solution will be most severe for these dipoles.

These methods are able to achieve estimates for the energy and density that are comparable to or better than those achieved with a single full iterative cycle since they focus their efforts on those elements of the dipole vector that change the most, while at the same time accounting for the minor modification of the dipoles at every other site. While simply modifying the dipoles on the moved molecules and its neighbors within 4 Å leads to acceptable results, doing this after ''tweaking'' the values of the dipoles at all of the other sites with the perturbative modification leads to much better results with little additional computational cost. Simply modifying the dipoles on the moved molecule leads to unacceptable results, which is not surprising considering that no account is taken of the changes that are induced in the dipoles at the other sites. When this is coupled with the method that perturbatively takes into account the changes in the dipoles at the other sites, a rapid and accurate method is achieved. This focusing of computational effort where major changes are expected should be contrasted with a Jacobi iterative procedure, which naively improves every dipole slightly and systematically.

ACKNOWLEDGMENTS

Gratitude is expressed to National Science Foundation for support of this work. The authors are also grateful to Dr. Julian Tirado-Rives and Dr. John C. Tully for helpful discussions and to Dr. J. Ilja Siepmann and Dr. Liem X. Dang for preprints and unpublished results.

APPENDIX A

As mentioned in the text, there is substantial disagreement between the values calculated for the Dang93 model and the values originally reported.[1,101] This disagreement appears to be due to the increased sensitivity of small polarizable systems to changes in such parameters as the number of molecules and the intermolecular cutoff distance for the following reason. The values originally reported for Dang97 water are reproduced reasonably well in Monte Carlo calculations with 550 molecules and have been reproduced with systems of up to 2500 molecules in molecular dynamics calculations.[106] They are also reasonably well reproduced in Monte Carlo calculations with 216 molecules if the cutoff distance is 8 Å, but if the cutoff in Monte Carlo calculations with 216 molecules is set to 9 Å, the density and magnitude of the energy increase so dramatically that the box fails to equilibrate before it shrinks to less than 18 Å on a side, exactly as is seen with the Dang93 results presented in Sec. IV.[104] It should be emphasized that this density increase is not due to an equilibrium fluctuation. A box of 216 molecules at standard density is ca. 18.62 Å on a side. To shrink to less than 18.0 Å would require a density fluctuation and thus an enthalpic cost much larger than the ca. 3% that is the maximum observed in extremely long equilibrium runs.[104] Thus for polarizable models it appears to be important to choose a cutoff distance significantly less than one-half the length of a side of the periodic cube. This sensitivity is to be contrasted with nonpolarizable models, where the energy and density exhibit a gradual dependence on the number of molecules and cutoff distance down to systems with fewer than 100 molecules and for cutoffs near half the box length.[58,107] For example, analogous calculations on 216 molecules of the TIP5P model of water[96,104] indicate that, while there are small changes in the energy and density upon changing the cutoff from 8 to 9 Å, they are similar to the changes induced upon changing the cutoff from 7 to 8 Å. Of course, the correct solution for dealing with this situation is to have a suf-

TABLE V. Relative timing results for several methods.[a]

Model	Dang93	Dang93	Dang97	Dang97	Dang97	Dang97
number of molecules	216	216	216	216	550	550
thousands of steps	1	10	1	10	1	10
no polarization	1.0	1.0	1.0	1.0	1.0	1.0
nbrit_4A	30	40	8.3	9.5	6.6	7.4
perturb_nbrit_0A	16	20	6.7	7.1	5.7	6.3
perturb_nbrit_4A	28	35	7.8	8.7	6.4	7.2
exact calculation	1100	1400	110	130	400	540

[a]See Appendix B for discussion.

ficiently large box. To our knowledge, a systematic examination of this has not been performed for polarizable models, whereas it has for nonpolarizable models.[58,107] It is unclear from the present results whether a different treatment of the boundary conditions would help the situation for a box of 216 polarizable molecules.

APPENDIX B

As indicated in Sec. VI, the only approximation methods that are not too computationally expensive are those that involve an iteration over the moved molecule and its neighbors and/or an initial perturbative modification. Representative timing results for the exact calculation are presented in Table V and compared with the results for several of the most promising approximate methods as well as results for systems that were equivalent except that there were no polarizable degrees of freedom. The calculations were in the N-V-T ensemble and involved 216 molecules for the Dang93 model and 216 and 550 molecules for the Dang97 model. The results are reported to two significant digits and are an average over five identical runs of length either 1000 or 10 000 steps. Due to the frequency of volume moves in the N-P-T calculations reported in this paper and the ''start-up'' costs in these calculations, the most reliable comparison is to the calculations here with 1000 steps; longer calculations are included to show that a relatively greater improvement is expected for N-V-T calculations. The code was not specifically optimized for speed, but instead represents a compromise between speed and memory requirements. Thus the results should not be taken as definitive, but rather representative.

Clearly, doing Monte Carlo on polarizable systems is significantly more difficult than doing molecular dynamics, as indicated by the factor of 2 to 3 orders of magnitude decrease relative to similar nonpolarizable systems. Note that doing a perturbative modification before iterating over neighbors within 4 Å (''perturb_nbrit_4A'') leads to calculations that are consistently faster than calculations without the initial perturbative modification (''nbrit_4A''); this is because performing the initial perturbative modification leads to a better initial guess for the iteration over the neighbors. Note also the dramatic dependence on the number of polarizable sites, as seen both by comparing the results for the Dang97 model at two different box sizes as well as the results for the two models in calculations with the same number of molecules. Table V indicates that the most promising approximation method, i.e., ''perturb_nbrit_0A,'' has a time speedup, relative to the exact calculation, of a factor of ca. 70 for systems consisting of 216 molecules for the Dang93 model or 550 molecules for the Dang97 model. The method ''perturb_nbrit_4A'' leads to slightly slower calculations, but is also much faster than the exact calculations. In the fastest approximate case, the CPU time increase for a polarizable calculation, relative to a nonpolarizable calculation, is a little more than an order of magnitude for the three-polarizable-site model and less than an order of magnitude for the one-polarizable-site model.

[1] L. X. Dang and B. C. Garrett, J. Chem. Phys. **99**, 2972 (1993).
[2] L. X. Dang and T. M. Chang, J. Chem. Phys. **106**, 8149 (1997).
[3] J. Caldwell, L. X. Dang, and P. A. Kollman, J. Am. Chem. Soc. **112**, 9144 (1990).
[4] L. X. Dang, J. E. Rice, J. Caldwell, and P. A. Kollman, J. Am. Chem. Soc. **113**, 2481 (1991).
[5] L. X. Dang, J. Chem. Phys. **96**, 6970 (1992).
[6] L. X. Dang and D. E. Smith, J. Chem. Phys. **99**, 6950 (1993).
[7] L. X. Dang, J. Chem. Phys. **97**, 2659 (1992).
[8] D. E. Smith and L. X. Dang, J. Chem. Phys. **100**, 3757 (1994).
[9] T. M. Chang and L. X. Dang, Chem. Phys. Lett. **263**, 39 (1996).
[10] L. X. Dang, J. Phys. Chem. B **102**, 620 (1998).
[11] T. P. Straatsma and J. A. McCammon, Mol. Simul. **5**, 181 (1990).
[12] C. J. Burnham, J. C. Li, S. S. Xantheas, and M. Leslie, J. Chem. Phys. **110**, 4566 (1999).
[13] E. L. Pollock and B. J. Alder, Phys. Rev. Lett. **39**, 299 (1977).
[14] E. L. Pollock and B. J. Alder, Phys. Rev. Lett. **41**, 903 (1978).
[15] A. Wallqvist and B. J. Berne, J. Phys. Chem. **97**, 13841 (1993).
[16] S. B. Zhu, S. Yao, J. B. Zhu, S. Singh, and G. W. Robinson, J. Phys. Chem. **95**, 6211 (1991).
[17] S. B. Zhu, S. Singh, and G. W. Robinson, J. Chem. Phys. **95**, 2791 (1991).
[18] A. Wallqvist, Chem. Phys. **148**, 439 (1990).
[19] M. Sprik, J. Phys. Chem. **95**, 2283 (1991).
[20] P. Ahlstrom, A. Wallqvist, S. Engstrom, and B. Jonsson, Mol. Phys. **68**, 563 (1989).
[21] S. W. Rick, S. J. Stuart, and B. J. Berne, J. Chem. Phys. **101**, 6141 (1994).
[22] A. A. Chialvo and P. T. Cummings, J. Chem. Phys. **105**, 8274 (1996).
[23] J. Brodholt, M. Sampoli, and R. Vallauri, Mol. Phys. **86**, 149 (1995).
[24] R. Car and M. Parrinello, Phys. Rev. Lett. **55**, 2471 (1985).
[25] D. K. Remler and P. A. Madden, Mol. Phys. **70**, 921 (1990).
[26] R. D. Mountain, J. Chem. Phys. **103**, 3084 (1995).
[27] D. N. Bernardo, Y. B. Ding, K. Krogh-Jespersen, and R. M. Levy, J. Phys. Chem. **98**, 4180 (1994).
[28] R. A. Bryce, M. A. Vincent, N. O. J. Malcolm, I. H. Hillier, and N. A. Burton, J. Chem. Phys. **109**, 3077 (1998).
[29] S. B. Zhu and C. F. Wong, J. Phys. Chem. **98**, 4695 (1994).
[30] M. L. Saboungi, A. Rahman, J. W. Halley, and M. Blander, J. Chem. Phys. **88**, 5818 (1988).
[31] M. Wilson and P. A. Madden, J. Phys.: Condens. Matter **5**, 2687 (1993).
[32] Y. B. Ding, D. N. Bernardo, K. Krogh-Jespersen, and R. M. Levy, J. Phys. Chem. **99**, 11575 (1995).
[33] D. van Belle and S. J. Wodak, J. Am. Chem. Soc. **115**, 647 (1993).

[34] M. H. New and B. J. Berne, J. Am. Chem. Soc. **117**, 7172 (1995).
[35] S. W. Rick and B. J. Berne, J. Am. Chem. Soc. **118**, 672 (1996).
[36] S. J. Stuart and B. J. Berne, J. Phys. Chem. **100**, 11934 (1996).
[37] T. P. Straatsma and J. A. McCammon, Chem. Phys. Lett. **177**, 433 (1991).
[38] J. W. Caldwell and P. A. Kollman, J. Phys. Chem. **99**, 6208 (1995).
[39] E. C. Meng and P. A. Kollman, J. Phys. Chem. **100**, 11460 (1996).
[40] E. C. Meng, J. W. Caldwell, and P. A. Kollman, J. Phys. Chem. **100**, 2367 (1996).
[41] J. W. Halley, J. R. Rustad, and A. Rahman, J. Chem. Phys. **98**, 4110 (1993).
[42] B. Roux, Chem. Phys. Lett. **212**, 231 (1993).
[43] M. A. Carignano, G. Karlstrom, and P. Linse, J. Phys. Chem. B **101**, 1142 (1997).
[44] K. Ramnarayan, B. G. Rao, and U. C. Singh, J. Chem. Phys. **92**, 7057 (1990).
[45] S. B. Zhu and G. W. Robinson, J. Chem. Phys. **94**, 1403 (1991).
[46] R. Veldhuizen and S. W. de Leeuw, J. Chem. Phys. **105**, 2828 (1996).
[47] S. Lin and P. C. Jordan, J. Chem. Phys. **89**, 7492 (1988).
[48] I. M. Svishchev, P. G. Kusalik, J. Wang, and R. J. Boyd, J. Chem. Phys. **105**, 4742 (1996).
[49] F. J. Vesely, J. Comput. Phys. **24**, 361 (1977).
[50] M. Sprik and M. L. Klein, J. Chem. Phys. **89**, 7556 (1988).
[51] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, J. Chem. Phys. **21**, 1087 (1953).
[52] I. R. McDonald, Chem. Phys. Lett. **3**, 241 (1969).
[53] J. A. Barker and R. O. Watts, Chem. Phys. Lett. **3**, 144 (1969).
[54] J. A. Barker and R. O. Watts, Mol. Phys. **26**, 789 (1973).
[55] R. O. Watts, Mol. Phys. **28**, 1069 (1974).
[56] W. L. Jorgensen, J. Chandrasekhar, J. D. Madura, R. W. Impey, and M. L. Klein, J. Chem. Phys. **79**, 926 (1983).
[57] W. L. Jorgensen, Chem. Phys. Lett. **92**, 405 (1982).
[58] W. L. Jorgensen and J. D. Madura, Mol. Phys. **56**, 1381 (1985).
[59] W. L. Jorgensen and J. Tirado-Rives, J. Phys. Chem. **100**, 14508 (1996).
[60] M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids* (Oxford University Press, Oxford, 1987).
[61] I. R. McDonald, J. Phys. C **7**, 1225 (1974).
[62] G. N. Patey, G. M. Torrie, and J. P. Valleau, J. Chem. Phys. **71**, 96 (1979).
[63] D. J. Adams and E. M. Adams, Mol. Phys. **42**, 907 (1981).
[64] J. M. Goodfellow, Proc. Natl. Acad. Sci. U.S.A. **79**, 4977 (1982).
[65] J. A. C. Rullmann and P. T. Van Duijnen, Mol. Phys. **63**, 451 (1988).
[66] P. Cieplak, P. Kollman, and T. Lybrand, J. Chem. Phys. **92**, 6755 (1990).
[67] J. S. Cao and B. J. Berne, J. Chem. Phys. **99**, 6998 (1993).
[68] J. S. Cao and B. J. Berne, J. Chem. Phys. **100**, 4720 (1994).
[69] W. L. Jorgensen, N. A. McDonald, M. Selmi, and P. R. Rablen, J. Am. Chem. Soc. **117**, 11809 (1995).
[70] J. Gao, D. Habibollazadeh, and L. Shao, J. Phys. Chem. **99**, 16460 (1995).
[71] J. Gao, J. J. Pavelites, and D. Habibollazadeh, J. Phys. Chem. **100**, 2689 (1996).
[72] J. Gao, J. Comput. Chem. **18**, 1061 (1997).
[73] J. Gao, J. Chem. Phys. **109**, 2346 (1998).
[74] M. Hloucha and U. K. Deiters, Mol. Phys. **90**, 593 (1997).
[75] M. G. Martin, B. Chen, and J. I. Siepmann, J. Chem. Phys. **108**, 3383 (1998).
[76] B. Chen and J. I. Siepmann, Theor. Chem. Acc. **103**, 87 (1999).
[77] M. Medeiros and M. E. Costas, J. Chem. Phys. **107**, 2012 (1997).
[78] P. Jedlovszky and J. Richardi, J. Chem. Phys. **110**, 8019 (1999).
[79] K. Kiyohara, K. E. Gubbins, and A. Z. Panagiotopoulos, Mol. Phys. **94**, 803 (1998).
[80] G. J. Martyna and B. J. Berne, J. Chem. Phys. **90**, 3744 (1989).
[81] R. A. Kuharski and P. J. Rossky, Chem. Phys. Lett. **103**, 357 (1984).
[82] A. Wallqvist and B. J. Berne, Chem. Phys. Lett. **117**, 214 (1985).
[83] R. O. Weht, J. Kohanoff, D. A. Estrin, and C. Chakravarty, J. Chem. Phys. **108**, 8848 (1998).
[84] P. Barnes, J. L. Finney, J. D. Nicholas, and J. E. Quinn, Nature (London) **282**, 459 (1979).
[85] G. Pastore, E. Smargiassi, and F. Buda, Phys. Rev. A **44**, 6334 (1991).
[86] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes: The Art of Scientific Computing* (Cambridge University Press, New York, 1986).
[87] C. W. Ueberhuber, *Numerical Computation 2* (Springer-Verlag, Berlin, 1997).
[88] G. H. Golub and C. F. Van Loan, *Matrix Computations* (Johns Hopkins University Press, Baltimore, 1983).
[89] D. Frenkel and B. Smit, *Understanding Molecular Simulation* (Academic, San Diego, 1996).
[90] Care must be taken when applying ''big-O'' arguments to systems for which the relevant variables do not increase to infinity. Not only may the magnitude of the prefactors and coefficients of lower order terms dominate at any given finite value of N, but also increasing the number of molecules without, for example, increasing the solvent–solvent cutoff distance can lead to algorithms which require effectively $O(N)$ steps per cycle, albeit generally with large prefactors.
[91] S. A. Clough, Y. Beers, G. P. Klein, and L. S. Rothman, J. Chem. Phys. **59**, 2254 (1973).
[92] P. L. Silvestrelli and M. Parrinello, Phys. Rev. Lett. **82**, 3308 (1999).
[93] W. L. Jorgensen, D. S. Maxwell, and J. Tirado-Rives, J. Am. Chem. Soc. **118**, 11225 (1996).
[94] W. D. Cornell, P. Cieplak, C. I. Bayly, I. R. Gould, K. M. Merz, D. M. Ferguson, D. C. Spellmeyer, T. Fox, J. W. Caldwell, and P. A. Kollman, J. Am. Chem. Soc. **117**, 5179 (1995).
[95] G. Kaminski and W. L. Jorgensen, J. Phys. Chem. **100**, 18010 (1996).
[96] M. W. Mahoney and W. L. Jorgensen, J. Chem. Phys. **112**, 8910 (2000).
[97] B. Roux and M. Karplus, J. Comput. Chem. **16**, 690 (1995).
[98] M. E. Tuckerman, G. J. Martyna, and B. J. Berne, J. Chem. Phys. **93**, 1287 (1990).
[99] M. E. Tuckerman, B. J. Berne, and G. J. Martyna, J. Chem. Phys. **94**, 6811 (1991).
[100] C. J. F. Bottcher, *Theory of Electric Polarization*, 2nd ed. (Elsevier, Amsterdam, 1973, Vol. 1); H. Frohlich, *Theory of Dielectrics* (Oxford University Press, Oxford, 1986).
[101] Since the calculations for the Dang93 model led to higher densities than those reported by Dang, for these calculations the intermolecular cutoff distance was decreased from 9 to 8.9 Å and then to 8.75 Å. This was not done to try to keep the box at the correct density. It was done in order to implement the intermolecular cutoff in a manner consistent with the periodic boundary conditions. See Appendix A for a more detailed discussion.
[102] W. L. Jorgensen, BOSS, Version 3.8, Yale University, New Haven, CT, 1997.
[103] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. M. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. Van der Vorst, *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods* (SIAM, 1993).
[104] M. W. Mahoney and W. L. Jorgensen (unpublished results).
[105] M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics* (Clarendon, Oxford, 1999).
[106] L. X. Dang (personal communication).
[107] C. Pangali, M. Rao, and B. J. Berne, Mol. Phys. **40**, 661 (1980).