Published in final edited form as:

J Chem Theory Comput. 2012 December 11; 8(12): 4863–4875. doi:10.1021/ct3006437.

Computationally Efficient Multiconfigurational Reactive Molecular Dynamics

Takefumi Yamashita^{1,‡}, Yuxing Peng^{2,‡}, Chris Knight^{3,‡}, and Gregory A. Voth^{2,3,*}

¹Laboratory for Systems Biology and Medicine, Research Center for Advanced Science and Technology, the University of Tokyo, 4-6-1 Komaba, Tokyo, Japan

²Department of Chemistry, James Franck Institute, and Computation Institute, University of Chicago, Chicago, IL 6063

³Computing, Environment, and Life Sciences, Argonne National Laboratory, Argonne, IL 60439

Abstract

It is a computationally demanding task to explicitly simulate the electronic degrees of freedom in a system to observe the chemical transformations of interest, while at the same time sampling the time and length scales required to converge statistical properties and thus reduce artifacts due to initial conditions, finite-size effects, and limited sampling. One solution that significantly reduces the computational expense consists of molecular models in which effective interactions between particles govern the dynamics of the system. If the interaction potentials in these models are developed to reproduce calculated properties from electronic structure calculations and/or ab initio molecular dynamics simulations, then one can calculate accurate properties at a fraction of the computational cost. Multiconfigurational algorithms model the system as a linear combination of several chemical bonding topologies to simulate chemical reactions, also sometimes referred to as “multistate”. These algorithms typically utilize energy and force calculations already found in popular molecular dynamics software packages, thus facilitating their implementation without significant changes to the structure of the code. However, the evaluation of energies and forces for several bonding topologies per simulation step can lead to poor computational efficiency if redundancy is not efficiently removed, particularly with respect to the calculation of long-ranged Coulombic interactions. This paper presents accurate approximations (effective long-range interaction and resulting hybrid methods) and multiple-program parallelization strategies for the efficient calculation of electrostatic interactions in reactive molecular simulations.

Keywords

reactive molecular dynamics; coarse-graining; multiple-program parallelization; effective long-range interaction

*Corresponding author: gavoth@uchicago.edu.

‡These authors contributed equally.

Supporting Information Available: Further discussion on accuracy of methods with comparison of energies and total energy conservation. This information is available free of charge via the Internet at <http://pubs.acs.org>.

1. Introduction

Chemical reaction dynamics take place in many important systems ranging from bulk solutions to biological environments and nanomaterial systems. In order to obtain atomistic level insight into chemical reaction dynamics, the molecular dynamics (MD) simulation technique has proven a uniquely powerful tool; however, standard molecular mechanics force fields cannot describe the chemical bond transformations associated with chemical reactions. Although recent increases in computational power have enabled the utilization of electronic structure methods to calculate forces “on-the-fly” during a simulation (e.g., Car-Parrinello and Born-Oppenheimer MD),¹ it remains computationally difficult to sufficiently sample important reactive degrees of freedom when coupled to slow motions of the system. For example, in the process of transporting excess protons in biological systems, the fast motion of the proton dynamics (femtosecond to picosecond timescales) can be strongly coupled to domain-domain fluctuations in a protein and even influenced by the surrounding membrane (e.g., in the case of cytochrome *c* oxidase).² The requirement to sufficiently sample these multiple length-scale motions makes it computationally intractable to simulate the entire system with electronic structure methods without resorting to significant approximations. In recent years, researchers have successfully applied hybrid quantum mechanics methods, such as QM/MM, which treat only the portion of the system one is interested in (usually a small fraction of the entire system) with accurate electronic structure methods, while the rest of the system is modeled by using molecular mechanics force fields.^{3, 4} QM/MM methods present unique challenges in themselves, however, as there can be difficulties in choosing the boundary and interactions between the QM and MM regions to reduce artifacts, although significant progress has been made in recent years.^{5–7} However, in order to converge statistical properties across multiple time and length scales using a reactive simulation MD methodology, the method must possess two properties: (1) the capability to accurately model chemical bond formation and breaking for general chemical reactions and (2) computational efficiency and, ideally, also the incorporation of highly scalable algorithms.

Among the many approaches for reactive MD simulations^{8, 9} with a reduced computational cost compared with that of full quantum mechanics simulations, the multistate empirical valence bond (MS-EVB) method,^{10–12} which belongs to the general class of “multiconfigurational” reactive MD algorithms, has successfully modeled chemical reactions in condensed phase environments. (We note here that the term “multistate” does not imply multiple real electronic states and with it nonadiabatic dynamics; rather, the “states” model the various bonding configurations of the adiabatic ground electronic state of the system.) Examples include proton and hydroxide transport in bulk water^{13–23} and at the water/vapor interface,^{24–27} proton transport in water-alcohol²⁸ and water-acetone mixtures,²⁹ proton transport in proteins^{11, 30–34} and lipid membrane interfaces,^{35, 36} proton solvation and transport in fuel cell polymer electrolyte membranes,^{37, 38} intramolecular isomerization/rearrangement,³⁹ and nucleophilic substitution.⁴⁰ The computational efficiency of the multistate reactive method, combined with its ability to accurately reproduce quantum mechanical reference data, enables the method to reliably converge

statistical properties and gain atomistic insight into complex, condensed phase reactive dynamics.

An important generalization of the original N-state MS-EVB method^{41, 42} is the ability to dynamically identify reactants of chemical reactions over the course of a simulation, thus enabling one to simulate chemical processes involving many molecules and spanning variable length separations. Moreover, the framework of MS-EVB (and other multistate, i.e., multiconfigurational methods) is sufficiently general to apply to general chemical reactions, as discussed earlier. A particular advantage of these multistate methods is that previously established molecular mechanics force fields can be utilized to describe all portions of the system not directly involved in the chemical reactions, thereby eliminating the need to reparameterize models containing many parameters. Those portions of the model describing the chemical reactions are parameterized by using *ab initio* calculations on clusters¹⁶ and condensed phase *ab initio* molecular dynamics (AIMD) simulations.¹⁸ Recently, the empiricism of using predefined interaction potentials for the reactive parts of the model was eliminated with the development of a force-matching algorithm to determine the effective interactions between nuclei as flexible tabulated expressions determined from condensed phase AIMD simulations.²² This new multiscale reactive molecular dynamics (MS-RMD) methodology was successfully applied to the case of the hydrated excess proton and hydroxide ion.²² The MS-RMD method provides a general framework for multistate reactive simulations for which MS-EVB is a specific application in which empirical functions are used to define the Hamiltonian matrix elements.

In multistate methods, a quantum-like Hamiltonian matrix is evaluated and diagonalized and the ground state eigenvector is used in combination with the Hellman-Feynman theorem to calculate the forces on the atomic nuclei. The diagonal elements of the Hamiltonian matrix correspond to different chemical bond topologies (diabatic states) included in the calculation at a given simulation step, and the off-diagonal couplings enable the transition between two states forming the reactants and products of a chemical reaction. The molecules in the system as described with a variable bond topology form the reactive complex, while all other molecules with fixed bonding topologies form the environment. Importantly, this definition of reactive complex and environment is dynamic and adjusted at every timestep in the simulation.

The diagonal matrix elements in the MS-RMD (or MS-EVB) method are typically defined by, but not limited to, a molecular mechanics force field with a diabatic correction. The off-diagonal couplings are typically defined to be local functions of those atoms directly involved in the chemical reaction. In multistate models for the excess proton and hydroxide ion, the off-diagonal has been defined to include an electrostatic term, in order to include effects from the environment on the reactive complex.^{16, 22} Because all nonzero elements of the Hamiltonian matrix need to be calculated at each step in the simulation, these reactive simulations inherently require more computational effort than do the corresponding nonreactive simulations. Significant computational savings, however, can be gained by evaluating all interactions involving the environment particles only once per simulation step. The remaining interactions involving atoms in the reactive complex are then evaluated for each nonzero element of the Hamiltonian matrix. The removal of redundancy in the

calculation of energies and forces is particularly important for reducing the computational cost of calculating long-range electrostatic interactions,⁴³ which are typically treated with Ewald summation or Ewald-type mesh algorithms that rely on fast fourier transforms (FFTs). Despite these optimizations, however, the calculation of long-range electrostatics for each nonzero Hamiltonian matrix element remains the main cause of performance degradation at high processor counts. Thus, it is highly advantageous to develop accurate effective interaction approximations and parallelization strategies in order to efficiently calculate electrostatic interactions in multistate reactive MD simulations, such as those discussed in this paper.

It was realized soon after the initial implementation of the MS-EVB algorithm that the k-space contribution to electrostatic interactions would be a computationally intensive calculation for all but the smallest systems.⁴³ The reason is largely that the (reciprocal) lattice summations in Ewald and mesh-type algorithms are global operations typically requiring communication amongst all processors. In Figure 1, the time cost of the real and k-space portions of an MS-RMD calculation are plotted versus the number of states in the simulation of an excess proton and 30,000 water molecules. For a nonreactive simulation (1 state), the k-space component takes up ~17% of the total run-time, but it quickly becomes dominant (~81% of run-time) by the time 16 EVB states are included. By comparison, the cost of evaluating the real-space interactions is nearly independent of the number of states relative to the dominant cost of the k-space calculations, with only a slight increase in cost out to 16 states. This low cost of the real-space interactions with respect to the number of states is largely because the interactions between the environment atoms are calculated only once per MD step and because the reactive complex represents a small fraction of the total system. These results indicate that compared with nonreactive simulations, the major component of the higher computational cost in MS-RMD (or MS-EVB) simulations is mostly due to the k-space calculations. Thus, an efficient evaluation of electrostatic interactions and a reduction in the number of required FFTs will lead to a significant reduction in the cost of the reactive MD simulations, approaching that of a nonreactive MD simulation.

In an effort to remove the necessity of the k-space calculation while retaining its accuracy, an effective-interaction MS-EVB (EI-MS-EVB) method⁴⁴ was recently developed where all nonbonded interactions in the MS-EVB3 model¹⁶ for the excess proton were mapped into short-range effective potentials using the coarse-graining in interaction space (CGIS) methodology.^{45, 46} With all nonbonded interactions described by effective short-range tabulated potentials, the (reciprocal) lattice summation is no longer necessary during simulations, thus leading to efficient, highly scalable calculations of energies and atomic forces. The analytic expressions derived by using CGIS⁴⁶ were also recently used to define reactive models for the excess proton and hydroxide ion in order to approximate the electrostatic energy and forces in the off-diagonal couplings.²²

In spite of the efficient calculation of forces in the EI-MS-EVB method, the requirement of first generating a sufficiently long reactive simulation using the exact method to generate the necessary data to force-match the effective interactions may not always be convenient; therefore, the analytic CGIS approximation to the Coulomb interaction, which was

previously derived for bulk water, was examined in simulations with an excess proton. Although it was found that a number of properties calculated by using the analytic CGIS model agreed with those from the original simulation for bulk water, the observed differences in the dipole autocorrelation function for water may have implications on the hydrated proton dynamics.⁴⁶ In the present work, by introducing the analytic CGIS expression for the electrostatic term only in the off-diagonal elements, the bulk of the system, where the forces are dominated by intraenvironment interactions, should be negligibly affected.

In addition to algorithms such as CGIS, recent successful efforts have been directed at developing efficient algorithms for evaluating nonbonded interactions in path integral quantum simulations.^{47, 48} In these methods, an efficient ring polymer contraction scheme is used where slow-varying interactions are evaluated by using ring polymers containing fewer path integral quasiparticles. For long-ranged k-space interactions, each ring polymer was replaced with the corresponding centroid, leading to a significant reduction in the computational cost of the electrostatics calculation while retaining an accurate description of the energy and atomic forces. In this work, an idea similar in spirit is introduced for multistate reactive MD methods where the exact k-space forces are approximated by a single k-space force calculation using a set of effective charges determined from the ground state eigenvector of the multistate Hamiltonian. This approximate method is referred to as the effective long-ranged interaction (ELRI) method, whereby the exact Hellman-Feynman k-space forces are approximated. This idea of using effective charges is also similar to the algorithm used for the self-consistent interactive MS-EVB (SCI-EVB) simulations,⁴⁹ where the exact electrostatic interaction between multiple reactive complexes is approximated by using effective charges.

Although the productivity of a multistate reactive MD simulation is increased by using the methods discussed here, simply reducing the number of k-space force calculations per MD step does not address the performance issues at high processor counts. Thus, the design and implementation of new parallelization strategies are among the most essential issues for performing reactive MD simulations at the petascale and beyond. One approach recently implemented in the GROMACS⁵⁰ MD code is a multiple-program (MP) strategy in which a small subset of processors is assigned the sole task of calculating k-space interactions. The remaining processors in this parallelization model calculate the intramolecular and short-range interactions, which tend to scale well with the number of processors. This MP parallelization strategy enables one to control the load balance between real and k-space calculations with more flexibility and (as discussed below) can straightforwardly be applied to accelerate reactive multistate simulations. Through a collaborative effort with XSEDE Extended Collaborative Support (ECS),⁵¹ the MP approach was implemented in LAMMPS⁵² and was recently found to increase productivity of simulations that made use of GPU acceleration.⁵³

In the next section, the analytic CGIS⁴⁶ expression for electrostatic energies and forces is reviewed and its implementation in reactive multistate MD simulations discussed. This discussion is followed by discussion of the ELRI method as well as the resulting hybrid methods formed by coupling the CGIS and ELRI methods. Physical properties from

simulations of three different systems with each of the approximate methods are discussed, including an analysis of the forces calculated with each of the methods. An MP parallelization strategy and its implementation are then presented, extending the parallel scaling of reactive multistate MD simulations. The paper concludes with a summary of results and a discussion of how the MP strategy can be used to further improve the parallel scaling of larger-scale reactive simulations.

2. Methods

In MD simulations, an accurate calculation for the long-range electrostatic interactions in periodic systems is assumed to require use of the Ewald method or one of its variants. The Ewald method splits the slow-converging Coulomb sum into a sum of two rapidly convergent terms: a short-range, pairwise component, which can be evaluated similar to conventional real-space pairwise interactions, and a long-range component, which is efficiently evaluated as a lattice (k-space) sum. The k-space component requires the Fourier summation over all atomic charges in the MD system with an optimal computational complexity of $O(N^{3/2})$,⁵⁴ where N is the number of charged particles. This scaling with respect to charged particles can lead to degraded parallel performance in large-scale simulations. In order to remedy this problem, mesh-type algorithms were developed; two popular methods are particle-mesh-ewald (PME)⁵⁵ and particle-particle particle-mesh (PPPM),⁵⁶ which discretize the charge distribution onto a grid and take advantage of FFTs, which scale as $O(N \log N)$. While improving the computational efficiency of the original Ewald method, however, these FFT-based algorithms also suffer from degraded parallel efficiency at higher processor counts, because of the communication intensive 3D FFT operations involving all-to-all communication, and largely limit the parallel efficiency of a simulation code. These performance issues are compounded in multistate simulations where multiple k-space calculations (3D FFTs) are required at each step in the simulation. The next sections discuss accurate approximations for efficiently calculating electrostatic interactions in multistate reactive MD simulations, followed by parallelization strategies that further extend the parallel scaling of MD codes and are equally applicable to both nonreactive and reactive simulations.

2.1 Reactive MD Simulations

In order to demonstrate the accuracy of the approximate electrostatic methods discussed below, simulations were calculated in the constant NVT ensemble for the hydrated excess proton using a Nose-Hoover thermostat with a relaxation time of 0.5 ps, and the equations of motion were integrated by using a 0.5 fs timestep. The system consists of the 216 water molecules and a single excess proton, with all interactions defined by using the MS-EVB3 model.¹⁶ The simulation cell was cubic with side lengths 18.621 Å; a pairwise cutoff of 9.0 Å was used, and, where appropriate, the Ewald was used with a precision of 10^{-6} . Additional simulations were calculated in the constant NVE ensemble in order to determine dynamic properties and examine any observed drifts in total energy (see Supporting Information).

The second system used for validating the approximate electrostatic algorithms involves examining the potential of mean force for deprotonating a histidine amino acid in bulk water. The simulation cell was orthorhombic, with side lengths $24.257 \times 23.887 \times 27.538$ Å, and contained a single Ace/Nme-terminated histidine residue, 500 water molecules, and a single excess proton, similar to the setup previously used when the reactive MS-EVB amino acid model was originally developed.⁵⁷ A pairwise cutoff of 9.5 Å was used for all short-range interactions, and a precision of 10^{-5} was used for k-space interactions calculated using PPPM, where appropriate. The equations of motion were integrated by using a 1.0 fs timestep in the constant NVT ensemble using the Nose-Hoover thermostat with a relaxation time of 0.5 ps. The histidine deprotonation potential of mean force (PMF) was calculated by using the umbrella sampling method, with a total of 32 windows evenly spaced between 0.8 and 7.0 Å for the reaction coordinate taken to be the distance between excess proton charge defect the center of excess charge (CEC) and the nitrogen atom. The sampling for each window was 2 ns, and the umbrella force constant was 40 kcal/mol/Å². The WHAM method was used to construct the potential of mean force.^{58, 59}

In order to investigate the possible introduction of artifacts from the present approximations when the systems being studied incorporate an interface, the affinity for the excess proton to the water/vapor interface was also examined. For this system, constant NVT MD simulations ($T = 298.15$ K) were calculated by using a simulation setup similar to that used in previous work,^{24, 27} where it was found that the excess proton resided preferentially at the interface and displayed “amphiphilic” characteristics. In the tests discussed here, a system consisting of 216 water molecules and one excess proton was prepared. The Nose-Hoover thermostat was used with relaxation time of 0.2 ps. The lattice constants for the simulation cell were $18.6121 \times 18.621 \times 93.621$ Å, and the water slab had a length of ~ 18.5 Å along the z-axis. For each of the methods examined, six 3.5 ns trajectories (21 ns in total) were used to calculate the probability distribution for the CEC coordinate, which has a function of distance normal to the water/vapor interface.

2.2 Coarse-graining in Interaction Space Effective Interactions

In the CGIS method, the long-range electrostatic interactions are mapped into a set of effective short-range interaction potentials, and therefore the large computational cost associated with the k-space summation is eliminated in molecular simulations.^{45, 46} Similar to recently proposed real-space approximations to electrostatics, the CGIS potential energy and force smoothly tend to zero at the cutoff and are defined to be

$$f_{eff}(r_{ab}) = \begin{cases} q_a q_b \left[\frac{1}{r_{ab}^2} - B r_{ab} \right], & 0 < r_{ab} < r_c \\ q_a q_b \left[\frac{1}{r_{ab}^2} + A(r_{ab} - R_{cut})^2 + B(r_{ab} - R_{cut}) + C \right], & r_c \leq r_{ab} < R_{cut} \end{cases} \quad (1)$$

and

$$V_{eff}(r_{ab}) = \begin{cases} q_a q_b \left[\frac{1}{r_{ab}} + \frac{B}{2} r_{ab}^2 - E \right], & 0 < r < r_c \\ q_a q_b \left[\frac{1}{r_{ab}} - \frac{A}{3} (r_{ab} - R_{cut})^3 - \frac{B}{2} (r_{ab} - R_{cut})^2 - C(r_{ab} - R_{cut}) - D \right], & r_c \leq r_{ab} < R_{cut} \end{cases} \quad (2)$$

where the coefficients are

$$A = \frac{1}{R_{\text{cut}}^2 r_c (R_{\text{cut}} - r_c)}, B = \frac{1}{R_{\text{cut}}^2 r_c}, C = -\frac{2}{R_{\text{cut}}^1}, D = \frac{1}{R_{\text{cut}}}, \quad (3)$$

$$E = \frac{A}{3} (r_c - R_{\text{cut}})^3 + \frac{B}{2} [(r_c - R_{\text{cut}})^2 + r_c^2] + C (r_c - R_{\text{cut}}) + D, \quad (4)$$

respectively.⁴⁶ Here, R_{cut} is the cutoff length where the potential and force are both zero, q_a is the charge of atom a , r_{ab} is the distance between atoms a and b , and the parameter r_c is set to be $0.81R_{\text{cut}}$ as in the original paper.⁴⁶ Importantly, these analytic functions were derived by using a variational force-matching (FM) method and were uniquely determined from bulk water simulations using an Ewald method for the electrostatic interactions.

In the first case studied here, the analytic CGIS expression is used to evaluate the electrostatic contribution for each nonzero Hamiltonian matrix element; this method will be referred to as CGISa in the remainder of this paper. It is similar to the effective interactions developed in the EI-MS-EVB method,⁴⁴ which map all nonbonded (Coulombic and vdW) interactions into short-range, pairwise tabulated functions; however, the method as used here alters only the electrostatic interactions, leaving the other nonbonded vdW terms intact. Similar to the EI-MS-EVB method, this application of CGIS and defining all interaction potentials as short-range, real-space potentials leads to the fastest reactive MD simulation algorithm. The choice of using an analytic expression with the original partial charges from the model improves transferability over the use of tabulated expressions developed for a specific system, but the use of real-space approximations still needs to be validated when interfaces or low-dielectric media are present, in order to ensure that artifacts are not introduced or at least are kept minimal.⁶⁰ According to the numerical tests for the hydrated excess proton (discussed below in more detail), these CGIS expressions prove to be sufficiently accurate, not only for the case of bulk water, but also for reactive simulations of the hydrated proton and amino acids.

In the second scheme studied in this paper, the CGIS expressions are applied to only the electrostatic term in the off-diagonal coupling. In this case, the electrostatic calculation for the diagonal elements remains being evaluated using a k-space method, such as Ewald or PPPM. The computational cost for this scheme, called CGISo, will of course be higher than that of CGISa but should be a more accurate approximation of the original energies and forces. In fact, by approximating the electrostatic term only in the off-diagonal coupling, the intraenvironment interactions, which constitute the vast majority of the system, remain unchanged from calculation with the exact method. A comparison of these CGIS variants is discussed below, but it is important to note that the use of CGIS for the off-diagonal electrostatics has been successful in the recent development of reactive MD models for the hydrated excess proton and hydroxide ion, which generated physical properties in agreement with the original AIMD data used in the parameterization.²²

2.3 Effective Long-Range Interaction Method

It is well known that a simple, truncated real-space approximation to Coulomb's law can potentially lead to unphysical results in molecular simulations of bulk and interfacial systems,⁶¹ although there has been recent success with the use of damped shifted force approximations and charge neutralization within the cutoff sphere.^{62, 63} The long-range contribution to electrostatics is a slowly varying function of distance compared with the short-range contribution, and so it is natural to consider whether the exact long-range k-space interaction in an Ewald-type method could be approximated by an effective interaction, which could then be calculated with less computational effort without introducing significant error. The success of CGIS in mapping long-range electrostatics onto a set of effective short-range potentials provides some support for this idea.^{45, 46} The slow variation with distance was also the motivation behind the recently introduced contraction schemes for path integral simulations, where instead of performing an Ewald calculation for each bead in the ring polymer, a single Ewald calculation is performed using only the centroid coordinate of each ring polymer.⁴⁸ In SCI-MS-EVB simulations of multiple reactive complexes, the electrostatic potential energies and forces for each diabatic state for a given reactive complex are also evaluated in the effective field setup by the other complexes.⁴⁹ The effective field from each reactive complex is defined by a set of effective charges determined from the ground state eigenvector of the corresponding Hamiltonian matrix:

$$q_{\alpha}^{eff} = \sum_{ij} c_i c_j q_{\alpha,ij}, \quad (5)$$

where the effective charge on the α th atom is calculated as a sum of charges of the α th atom in each ij element of the Hamiltonian matrix. One might then ask whether a similar set of effective charges would prove useful in approximating energies and forces in a multistate MD simulation using a single k-space calculation as opposed evaluating k-space forces for each element of the Hamiltonian matrix. This is exactly idea behind the ELRI method introduced here; and, as discussed below, the errors introduced have proved negligible for all systems examined thus far.

In multistate MD methods, such as MS-RMD or MS-EVB, the total atomic forces are calculated by using the Hellmann-Feynman theorem, such that

$$\mathbf{F}_{\alpha} = - \left\langle \Psi_0 \left| \frac{\partial \mathbf{H}}{\partial \mathbf{r}_{\alpha}} \right| \Psi_0 \right\rangle = \sum_{ij} c_i c_j \mathbf{F}_{\alpha}^{ij}. \quad (6)$$

The long-ranged k-space contribution, acting on atom α from atom β , is given by

$$\mathbf{F}_{\alpha\beta}^{long} = \sum_{ij} c_i c_j \mathbf{F}_{\alpha\beta}^{ij, long},$$

where

$$\mathbf{F}_{\alpha\beta}^{ij, long} = q_{\alpha}^{ij} q_{\beta}^{ij} \mathbf{f}_{\alpha\beta}, \quad (7)$$

the charges are defined for each atom and element of the Hamiltonian and $\mathbf{f}_{\alpha\beta}$ is a function only of distance between the two atoms, typically involving a summation over wave vectors. In the ELRI method, the explicit charges on atoms are replaced with a set of effective charges, thus approximating the exact Hellman-Feynman k-space interactions with a single k-space force calculation, which for a pair of atoms would give given by

$$\mathbf{F}_{\alpha\beta}^{long} \approx \mathbf{F}_{\alpha\beta}^{ELRI} \equiv q_{\alpha}^{eff} q_{\beta}^{eff} \mathbf{f}_{\alpha\beta}. \quad (8)$$

The computational savings for this method can be significant given the reduction in the number of wavevector summations (FFTs) required for the k-space forces at each step in the simulation. As discussed below, this approach has been found so far to introduce only small differences in the atomic forces compared with the exact forces.

If atom α is located not in the reactive complex but instead within the environment region, then the value for the effective charge simply takes the value defined in the molecular mechanics force field as is done in the case of the exact method. If both atoms α and β are found in the environment region, then no approximation is introduced, and the exact k-space force acting between these two particles is retained. This partitioning of atoms as belonging to either the reactive complex or the environment region was the motivation for the original k-space splitting method implemented for MS-EVB simulations to significantly reduce the computational cost of the Ewald calculations.⁴³ With the ELRI method, however, there is no need to distinguish between atoms residing in one or the other, except for the calculation of effective charges for those atoms within the reactive complex, which is commonly only a small fraction of the total number of atoms in the simulation cell.

In the ELRI method, the accuracy of the k-space forces, apart from Ewald-related parameters, is determined solely by the accuracy of the ground state eigenvector obtained after diagonalization of the Hamiltonian matrix. If an accurate approximation to the Hamiltonian matrix is found, then additional computational savings can be gained, particularly by reducing the computational cost of the k-space energy calculations. In fact, this is similar to the case of the EI-MS-EVB Hamiltonian in which all nonbonded interactions were mapped into a set of effective short-range potentials.⁴⁴ Once the ground state eigenvector for this new effective Hamiltonian, $\mathbf{c}^{eff}$, is obtained, then a set of effective charges,

$$q_{\alpha}^{eff} = \sum_{ij} c_i^{eff} c_j^{eff} q_{\alpha}^{ij}, \quad (9)$$

can be determined in analogous fashion, and a single k-space calculation can be used to approximate the forces. For this effective short-range Hamiltonian, the corresponding energies and forces are given by

$$E = \sum_{ij} c_i^{eff} c_j^{eff} H^{ij,short} + E^{ELRI}(\mathbf{c}^{eff}) \quad (10)$$

$$\mathbf{F}_\alpha = \sum_{ij} c_i^{eff} c_j^{eff} \mathbf{F}_\alpha^{ij,short} + \mathbf{F}_\alpha^{ELRI}(\mathbf{c}^{eff}). \quad (11)$$

Here, $H^{ij,short}$ is an element of the short-range Hamiltonian matrix, and $\mathbf{F}_\alpha^{ij,short}$ is the force on atom α . The quantities $E^{ELRI}(\mathbf{c}^{eff})$ and $\mathbf{F}_\alpha^{ELRI}(\mathbf{c}^{eff})$ are the same long-range interaction energy and force calculated above, but this time with charges derived from the eigenvector of the effective short-range Hamiltonian.

Although eliminating the need for k-space energy calculations in the construction of the Hamiltonian matrix will lead to considerable computational savings, as seen for the EI-MS-EVB models,⁴⁴ it was also found through the testing of several simple short-range approximations that the resulting reactive MD simulations yielded unacceptable drifts in total energy (see Supporting Information). However, a compromise in efficiency between a full k-space calculation and that of a short-range real-space approximation was found by using a coarse k-space grid for the long-range contribution to the electrostatic energies in the diagonals. Although the use of a coarse k-space Hamiltonian is not as computationally efficient as the real-space effective Hamiltonians, there is an expected larger degree of transferability. By reducing the number of wave vectors in each dimension only by a factor of 2, considerable computational savings can be obtained in the calculation of k-space energies for each diagonal element of the Hamiltonian matrix. As discussed below, the physical properties calculated from simulations using these approximate k-space energies are in excellent agreement with those calculated using the exact method, implying that the use of a coarse k-space grid does not significantly alter the relative energy differences of the diabatic states. For Ewald-mesh type algorithms, such as PPPM or SPME, the computational cost of the simulations can be significantly lowered because of reduced volume of communication involved in the 3D FFTs. For example, consider the case of N k-space calculations (one for each of the N nonzero Hamiltonian matrix elements) with M k-space grid points along each axis. The cost of a 3D FFT operation is proportional to $NM^3 \log M$. With the use of a coarse k-space grid for the diagonal energy calculations, the cost is decreased to $1/8NM^3 \log(M/2) + M^3 \log M$, corresponding to the cost of the N coarse k-space energy calculations and a single full-grid force calculation. For the discussion that follows, this method of using a coarse k-space grid will be abbreviated as ELRICK.

2.4 Hybrid Electrostatic Methods

As discussed above, the CGISO method is computationally slower than the EI-MS-EVB method because the energies and forces for the diagonal elements of the Hamiltonian matrix are still evaluated by using a k-space method. In order to improve further the computational efficiency, the ELRI method, which focuses on the diagonal elements, can be combined with the CGISO method, which focuses on the off-diagonal couplings. The combination of these two methods (CGISO+ELRI) will be referred to as CE (computationally efficient) in the

discussion that follows. In the CGISO Hamiltonian, only the diagonal matrix elements involve a k-space calculation; thus, only the diagonal contribution to the effective charges should be included when calculating the k-space forces in an ELRI-type approximation using

$$q_{\alpha}^{eff} = \sum_i |c_i|^2 q_{\alpha}^{ij}. \quad (12)$$

In other words, since the long-range Coulombic interaction has already been completely accounted for in the off-diagonal couplings using a short-range approximation, the remaining contribution to the long-range electrostatic forces is derived only from the diagonal elements; thus they are the only ones that contribute to the effective charges. As discussed below, the combination of these two methods leads to an increased computational efficiency by significantly reducing the cost of the electrostatic calculations, but at the same time this combination remains sufficiently accurate to yield thermodynamic and dynamic results in agreement with the original “exact” reactive MD simulations calculated using full k-space methods for all nonzero Hamiltonian matrix elements. The use of a coarse k-space grid for the diagonal energies in the ELRIck method can similarly be combined with CGISO for the off-diagonals, which leads to a further reduction in computational cost. This combination of methods, referred to as CEck in the discussion that follows, provides the fastest of the approximate electrostatic algorithms introduced thus far that still accurately reproduce a number of physical properties compared with full multistate reactive MD simulations with exact energies and forces. Again, as with all of the approximations introduced, the interactions involving only atoms within the environment region are unaffected.

In the next section, the discussion turns to the design of efficient parallelization strategies for multistate reactive MD simulations to efficiently calculate k-space interactions. For these simulations, all approximate methods developed thus far can be utilized.

2.5 Multiple-Program Parallelization

The main bottleneck at high processor counts in the calculation of electrostatic interactions with Ewald mesh-type algorithms (i.e., PME and PPPM) is the calculation of 3D FFTs, which involves communication among all processors in typical parallelization strategies. For MD codes that utilize a hybrid OpenMP/MPI strategy, the performance of the 3D FFTs can be improved with the reduced number of MPI ranks communicating the required FFT data. In order to improve the performance even further, a multiple-program (MP) scheme can be adopted to further increase the flexibility of the number of processors involved in the k-space electrostatic calculation.

Conventionally, MD codes (as most others) are parallelized with a single-program-multiple-data scheme, in which all processors execute the same program instructions but with different blocks of data assigned to each processor. In a multiple-program strategy, instructions that a processor executes may be different from the instruction set another processor executes. In one possible scheme for a multiple-program strategy for calculating electrostatics, the processors assigned to the calculation are first divided into two separate

partitions. The first partition, referred to as the k-space partition, is assigned the sole task of computing the k-space component of electrostatic interactions, principally involving 3D FFTs. The other partition, referred to as the real-space partition, is assigned all remaining effort for the simulation, such as calculating the remaining forces (e.g., pairwise and bonded interactions), which tend to scale well with the number of processors, integrating Newton's equations of motion, and writing output. A typical ratio for the number of processors in each partition is 3:1, in which only a fourth of all processors are involved in the calculation of the communication intensive 3D FFTs. The optimal ratio will vary between computing architectures and details of the systems being studied, but it is a parameter that can be optimized for each particular case.^{50, 51} Some performance results for the implementation of a 2-partition MP strategy in the LAMMPS code are discussed below for both nonreactive and reactive multistate MD simulations.

With the MP strategy alone, the communication-intensive 3D FFT operations would eventually reclaim their status as the major bottleneck in simulations at larger processor counts. One can simply increase the ratio of real-space to k-space processors; but depending on details of the implementation, it can start to become tedious to determine the optimal processor mappings for each partition, so that the processor grids are roughly uniform and themselves remain an efficient choice for a domain-decomposition strategy, such as that employed in the LAMMPS MD code.⁵² However, by coupling the approximate methods developed here, which reduce the cost associated with the k-space calculations, one can further extend the parallelization of a simulation while remaining at a 3:1 ratio. All approximate electrostatic methods discussed in this paper are fully compatible with this MP parallelization strategy, and their coupling will further improve the scaling performance of reactive simulations. As discussed below, however, the ability to go beyond just two partitions will be important in the design of a highly scalable multistate reactive MD simulation code.

3. Results

3.1 Excess Proton in Water

The first system used to validate and compare the different electrostatic approximations is that of the hydrated proton in bulk water using the MS-EVB3 model.¹⁶ A comparison of properties calculated using the CGIS, ELRI, and hybrid methods is shown in Figures 2–4, respectively. As can be seen, all approximate methods yield results that are in excellent agreement with those results calculated with an exact MS-EVB method. Radial distribution functions for each of the four pairs of atoms between the hydronium cation and water molecules are shown in plots a–d for each of Figures 2–4, where the atom labels are defined in the figure caption. All RDFs calculated by using approximate methods were found to be similar to the curves calculated by using the exact method, indicating that equilibrium configurational properties are not significantly affected. This result is further confirmed in the comparison of pK_a values for the hydronium cation as calculated using the RDF between oxygen atoms and the CEC. As shown in Table 1, all calculated values for the pK_a are identical to the value calculated for the exact method, which is slightly underestimated compared with the experimental value of -1.74 .

Other important properties, which are specific to the multistate methodology, are the distributions of the largest ($|c_{\max}|^2$) and the second largest amplitudes ($|c_{\max2}|^2$) of the ground state eigenvector. In the distributions of these quantities, the region of $|c_{\max}|^2 \approx 0.6$ and $|c_{\max}|^2 - |c_{\max2}|^2 \approx 0.5$ corresponds to the dynamically distorted Eigen-type (H_9O_4^+) complex^{13, 14} in which the excess proton is dominantly associated with a single water molecule, with three other water molecules strongly hydrogen-bonded to it. The region of $|c_{\max}|^2 \approx 0.45$ and $|c_{\max}|^2 - |c_{\max2}|^2 \approx 0$ corresponds to the Zundel-type complex (H_5O_2^+), which is more closely associated with transitional complexes for the proton transport process.¹⁴ Again, properties calculated using the approximate methods are in excellent agreement with those calculated from the exact method, with deviations in the ELRI-related curves in Figure 3 most noticeable and tending to a more localized charged defect.

In addition to configurational properties and equilibrium distributions of coefficients for the ground state eigenvector, dynamical properties were found to be largely unaffected by the approximations made in each of the different methods discussed. As shown in Table 2, the diffusion coefficient of the excess proton CEC for all methods is in excellent agreement with the value obtained from the exact MS-EVB3 reactive MD calculation. The root-mean-squared-deviations between the total atomic forces for the CGISo and CEck methods with the forces from the exact method are 0.007 and 0.012 (kcal/mol)/Å, respectively, as determined from 500 configurations equally sampled over the course of a 500 ps trajectory. The maximum absolute deviation in a component of the total atomic force is 0.282 and 0.557 (kcal/mol)/Å for the CGISo and CEck methods, respectively, which are small fractions of the average magnitude of forces observed in the simulation, 13.476 (kcal/mol)/Å (a detailed comparison of the energies is discussed in the Supplemental Information). In addition, the total energy in constant NVE simulations was not significantly altered from the drift previously reported for the MS-EVB3 model of -3.4 kcal/mol per nanosecond.¹⁶

3.2 Histidine Deprotonation

The next system examined to validate the electrostatic approximations developed here was the deprotonation of an Ace/Nme histidine residue in water. For this system, the PMF was calculated for three of the methods (Figure 5), and all were found to agree well with the PMF calculated from simulations with the exact MS-EVB reactive MD method. The minimum in the PMF for the undissociated amino acid and the barrier for deprotonation are in good agreement for all methods. Noticeable differences appear for the CGISo method near the position of the second shallow minimum (near 2.72 Å), and the second barrier is reduced by ~ 0.33 kcal/mol. When CGISo is combined with either of the ELRI methods, however, the resulting PMFs match well with that calculated from the exact method. For configurations corresponding to the protonated histidine, the average number of diabatic states was five, and the CGISo, CE, and CEck methods were found to be 1.34, 1.56, and 1.64 times faster than the case of using a k-space method for all electrostatic interactions. The corresponding speedups were increased to 1.5, 1.88, and 1.99 times faster for the case of the excess proton in the bulk water environment, far from the residue, where the average number of diabatic states was 25. This timing result is in agreement with the expected trend that as the number of states increases, the benefits of reducing the cost of the 3D FFTs will become more substantial (Table 3). The same is true that the resulting speedups will increase

with increasing system size as the reactive complex becomes a smaller and smaller fraction of the total system.

3.3 Excess Proton at Water/Vapor Interface

In order to examine how these approximations handle the case of charged excess proton defects next to an interface, the water/vapor interface is the next system examined. As shown in Figure 6, the distributions from all approximate methods are consistent with the probability distribution calculated by using the exact method. Qualitatively, all the present approximations can reproduce the surface propensity very well. The CGIS-related methods reproduce the position of the peak in the probability distribution, but the CGISa method leads to a slight overestimation of the surface propensity of the proton, while the effect is reduced somewhat in the CGISo distribution improving the agreement with the exact result. The improved agreement of CGISo over CGISa is expected because only the off-diagonal couplings are altered in the CGISo method, leaving all water-water interactions in the diagonal Hamiltonian matrix elements unchanged. In Figure 6b, the peak in the probability distributions for the ELRI-related methods is found to be the slightly shifted toward the slab center. A similar shift toward the center of the slab is also observed in the hybrid methods in Figure 6c, but the shift is reduced in comparison. In addition, the probability for the proton to be found near the center of the slab is increased relative to the results from the exact method. A similar increase in probability near the center of the slab was found for the CGIS-related methods in Figure 6a, but to a smaller extent. The corresponding free energy differences to bring the proton from the interface toward the bulk region for all methods span the range of 1.5–2.0 kcal/mol, with the hybrid methods having the lowest values. The free energy stabilization for the nonhybrid methods is closer to the previously calculated value of 1.8 kcal/mol at 300 K.²⁷

3.4 Multiple-Program Parallelization

As discussed earlier, an efficient parallelization strategy for large-scale molecular simulations is crucial in order to properly sample the slow degrees of freedom in a system. In order to highlight the improvements from the MP parallelization strategy in LAMMPS, the first system considered is the Hv1 biomolecular⁶⁴ system containing ~40,000 atoms, running on the Ranger (Sun Constellation Linux Cluster) at the Texas Advanced Computing Center. For a single, 8-core compute node, the real-space calculations in LAMMPS using a standard molecular mechanics force field and short-range cutoff of 10 Å takes 83.4% of the total run-time, while the k-space calculation with PPPM and a precision of 10^{-4} only takes 8.9%. The remaining portion of the run-time corresponds to the simulations setup, I/O, integrating the equations of motion, and so forth. When the same simulation is scaled up to 256 cores, however, the computational cost of the k-space calculations increases to 56.9%, while the relative cost of the real-space calculation decreases to 38.1%. This trend of the relative cost of k-space portion of the calculation increasing at higher core counts is general among MD codes; it results primarily from the degraded performance of 3D FFTs and associated all-to-all communication. In domain-decomposition strategies, such as that employed in the LAMMPS MD code,⁵² the real-space portion of the force calculation tends to scale well with the number of processors until the number of atoms per MPI rank has decreased below a certain threshold value depending on details of the system under study. In

order to better understand the scaling properties of the k-space calculation, the parallel efficiency from benchmarks of a larger system, Cytochrome *c* Oxidase (CcO)³⁰ solvated with water, containing ~159,000 atoms was examined. The parallel efficiency was calculated by using

$$E_p = \frac{T_1}{p \times T_p}, \quad (13)$$

where T_1 and T_p are the times associated with simulations on a single-core and p cores. In the nonreactive simulation on 32 nodes (512 cores), the real-space portion of the force calculation is 98.9% efficient, displaying the nearly ideal parallel scaling mentioned earlier. In contrast, the k-space portion of the calculation is only 3.0% efficient on 32 nodes. For this setup, primarily because of the performance of the k-space calculation, adequate performance is observed only up to 4 nodes. As expected, this performance degradation is largely due to the communication of 3D FFTs, which themselves have poorer efficiency (0.4%) at 32 nodes. By reducing the number of MPI ranks involved in the communication of the 3D FFTs, through either a hybrid MPI/OpenMP or MP strategy (or combination of the two), a larger number of compute nodes can be used in the simulation before the k-space performance bottlenecks noticeably degrade overall scaling. An additional advantage of implementing an MP strategy is the performance boost associated with the simultaneous evaluation of forces on the separation partitions.

The simulation productivities of the CcO benchmark for both single-program (SP) and MP algorithms in the LAMMPS and RAPTOR (MS-EVB/RMD code coupled to LAMMPS) codes are shown in Figure 7. As discussed above for this particular setup, the degradation in performance starts near 4 compute-nodes. However, by assigning fewer cores to the task of calculating the k-space interactions (FFTs) using a 3:1 partition ratio, the productivities of both nonreactive and reactive simulations continue to increase up to at least 512 cores. The original parallelization method, SP, reaches maximum performance at 128 cores, with ~4.73 million-step/day with the k-space calculations dominating the run-time. For the MP method, the productivity at 512 cores is ~11.28 million-step/day, which is ~2.38 times larger with the possibility of increased productivities at still higher processor counts.

In the multistate reactive MD simulations, the computational burden due to k-space calculations is much heavier, since multiple k-space energy and force calculations are required at each step in the simulation. A comparison of the productivities for reactive simulations is shown in Figure 7 for both the SP and MP strategies. Similar trends are observed for the RAPTOR code in which the default SP algorithm efficiently scales only to 64 cores, with a productivity of ~0.33 million-step/day, but the performance of the MP algorithm continues to efficiently scale up to 256 cores with a productivity of ~1.07 million-steps/day. At 256 cores, the reactive simulations are only a factor of 3.3 times slower than the corresponding nonreactive simulations when the MP algorithm is used.

In Table 3, a general comparison of some systems for each of the approximate methods is given, comparing the costs of the multistate reactive MD relative to that of nonreactive MD simulations using identical simulation setups with respect to k-space precision, short-range

cutoffs, and frequency of updating the neighborlists. The exact ratios depend on several details, but the trend is clear: as the number of reactive MD states increases, the benefits from the approximate methods developed here become increasingly more significant. The same is true as the volume of the simulation cell increases and the reactive complex becomes a smaller and smaller fraction of the total system. In the best-case scenarios for the systems with +150K atoms, the reactive multistate simulations are only about five times more expensive than the corresponding nonreactive MD simulations. While a factor of 5 is not negligible, these reactive MD simulations are still several orders of magnitude faster than purely AIMD or QM/MM MD simulations. (The latter are still invaluable, e.g., in bridging quantum mechanical information to the MS-EVB and MS-RMD simulations.^{18, 22})

The productivity of reactive simulations can also be further improved by combining the approximate methods for electrostatic interactions with the MP parallelization strategy. In preliminary benchmarks on Intrepid (Blue Gene/P) at the Argonne Leadership Computing Facility (ALCF), the ~159K atom CcO benchmark system was successfully scaled up to 8,192 cores using a 3:1 ratio of MP partitions (Figure 8). With the exception of the CEck method (and other ELRI-related methods), the MP approach greatly extends the parallel efficiency of the approximate methods discussed here. The lack of a significant speedup of the ELRI-related methods compared with just using CGISo is that there are only six diabatic states on average in this CcO benchmark. With the CGISo approximation, the k-space partition of processors waits for the real-space partition to finish its workload at each simulation step; thus, there are no additional speedup gains by further reducing the cost of the 3D FFTs. As the number of states increases (along with the cost of the k-space calculation), the performance of the ELRI-related methods becomes more noticeable, as observed in preliminary benchmark data for an excess proton in a simulation cell of 50K water molecules (data not shown). Further improvements in performance are expected by increasing the ratio of real-space to k-space processors and addressing load imbalance issues between the two partitions.

4. Conclusions

In order to accelerate multistate reactive MD simulations, particularly the evaluation of electrostatic energies and forces, without sacrificing the accuracy of the original multistate reactive models, several methods and techniques have been developed and validated in the present paper. In this work, the CEck method, which is the combination of the CGISo and the ELRIck methods, leads to the fastest algorithm without significantly altering calculated thermodynamic and dynamic properties, even in the presence of abrupt interfaces such as the water/vapor interface. These algorithms in combination with improved parallelization strategies greatly reduce the computational time of the multistate reactive MD simulations, which in some cases are comparable to that of nonreactive simulations. In order to check the accuracy and efficiency of the methods developed here, several physical properties were compared for excess proton solvation and transport in bulk water and at the water/vapor interface. The results from these tests showed that the newly developed methods can reproduce all calculated properties, in most cases quantitatively, as compared with those calculated with the exact method, while at the same time leading to a significant reduction in computational cost. The parallel performance of all simulation methods discussed here was

further improved by using a multiple program parallelization strategy that evaluated the k-space portion of the calculation on a separate, smaller partition of processors.

The focus throughout this paper has been mostly on multistate reactive MD simulations of a single reactive complex, such the hydrated excess proton; however, all the algorithms and parallelization strategies discussed are immediately applicable to simulations within the general MS-RMD framework. For example, the methods presented here can be straightforwardly applied to quantum MS-RMD simulations based on imaginary-time path integral methods. In addition, these algorithms are currently being used to accelerate simulations with multiple reactive complexes in SCI-MS-EVB⁴⁹ simulations. Importantly, these methods are independent of the reactive species being investigated, thus enabling one to take full advantage of the generality of multistate reactive MD algorithms to simulate chemical reactions in complex condensed phase systems. Future efforts will be directed to applying the methods discussed in this paper to a broad spectrum of chemical reactions, as well as further optimizing the parallel efficiency of these reactive simulations by extending the MP strategy to more than two partitions. A highly efficient parallel strategy for multistate reactive MD simulations can in principle be achieved by distributing subsets of Hamiltonian matrix elements for each reactive complex to independent sets of partitions for their simultaneous evaluation, in addition to the reduced costs of the k-space calculations using the methods discussed here. By coupling our general multistate framework for reactive MD simulations with computationally efficient algorithms to take advantage of current and future computing technologies, transformative large-scale reactive MD simulations will become possible enabling researchers to address key challenges in chemistry, biochemistry, biophysics, and materials science.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This research was supported by the National Science Foundation (NSF grant CHE-1214087) and the National Institutes of Health (NIH grant R01-GM053148). T.Y. acknowledges support from the Japan Society for the Promotion of Science (JSPS) through the “Funding Program for World-Leading Innovative R&D on Science and Technology (FIRST Program),” initiated by the Council for Science and Technology Policy (CSTP). We also acknowledge support from the U.S. Department of Energy under Contract DE-AC02-06CH11357 and an Argonne National Laboratory Computational Science Postdoctoral Fellowship to C.K. The computations in this work were supported in part by a grant of computer time from the DOD High Performance Computing Modernization Program at the Engineer Research and Development Center (ERDC) DoD Supercomputing Resource Center. The computations also utilized the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number OCI-1053575. An award of computer time was further provided by the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. This research used resources of the Argonne Leadership Computing Facility at Argonne National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under contract DE-AC02-06CH11357. We thank Phil Blood, Lonnie Crosby, and XSEDE Extended Collaborative Support (ECS) for assistance in optimizing the multiple-program parallelization in the LAMMPS and RAPTOR codes.

REFERENCES

1. Kirchner B, Dio P. J. d. Hutter J. Top. Curr. Chem. 2012; 307:109–154. [PubMed: 21842358]
2. Öjemyr LN, Ballmoos C. v. Faxén K, Svahn E, Brzezinski P. Biochemistry. 2012; 51:1092–1100. [PubMed: 22257086]

3. Senn HM, Thiel W. *Angew. Chem. Int. Ed.* 2009; 48:1198–1229.
4. Lin H, Truhlar DG. *Theor. Chem. Acc.* 2007; 117:185–199.
5. Kamerlin SCL, Haranczyk M, Warshel A. *J. Phys. Chem. B.* 2009; 113:1253–1272. [PubMed: 19055405]
6. Higashi M, Truhlar DG. *J. Chem. Theory Comput.* 2008; 4:790–803.
7. Bernstein N, Várnai C, Solt I, Winfield SA, Payne MC, Simon I, Fuxreiter M, Csányi G. *Phys. Chem. Chem. Phys.* 2012; 14:646–656. [PubMed: 22089416]
8. Duin, A. C. T. v.; Dasgupta, S.; Lorant, F.; Goddard, WA, III. *J. Phys. Chem. A.* 2001; 105:9396–9409.
9. Frauenheim T, Seifert G, Elstner M, Niehaus T, Köhler C, Amkreutz M, Sternberg M, Hajnal Z, Carlo AD, Suhai S. *J. Phys.: Condens. Matter.* 2002; 14:3015–3047.
10. Voth GA. *Acc. Chem. Res.* 2006; 39:143–150. [PubMed: 16489734]
11. Swanson JMJ, Maupin CM, Chen H, Petersen MK, Xu J, Wu Y, Voth GA. *J. Phys. Chem. B.* 2007; 111:4300–4314. [PubMed: 17429993]
12. Knight C, Voth GA. *Acc. Chem. Res.* 2012; 45:101–109. [PubMed: 21859071]
13. Schmitt UW, Voth GA. *J. Chem. Phys.* 1999; 111:9361–9381.
14. Markovitch O, Chen H, Izvekov S, Paesani F, Voth GA, Agmon N. *J. Phys. Chem. B.* 2008; 112:9456–9466. [PubMed: 18630857]
15. Xu J, Zhang Y, Voth GA. *J. Phys. Chem. Lett.* 2011; 2:81–86.
16. Wu Y, Chen H, Wang F, Paesani F, Voth GA. *J. Phys. Chem. B.* 2008; 112:7146.
17. Wang F, Izvekov S, Voth GA. *J. Am. Chem. Soc.* 2008; 130:3120–3126. [PubMed: 18275191]
18. Knight C, Maupin CM, Izvekov S, Voth GA. *J. Chem. Theory Comput.* 2010; 6:3223–3232.
19. Xu J, Zhang Y, Voth GA. *J. Phys. Chem. Lett.* 2011; 2:81–86.
20. Park K, Lin W, Paesani F. *J. Phys. Chem. B.* 2012; 116:343–352. [PubMed: 22107267]
21. Ufimtsev IS, Kalinichev AG, Martinez TJ, Kirkpatrick RJ. *Phys. Chem. Chem. Phys.* 2009; 11:9420–9430. [PubMed: 19830325]
22. Knight C, Lindberg G, Voth GA. *J. Chem. Phys.* 2012; 137:22A525.
23. Chen H, Voth GA, Agmon N. *J. Phys. Chem. B.* 2010; 114:333–339. [PubMed: 19961199]
24. Petersen MK, Iyengar SS, Day TJF, Voth GA. *J. Phys. Chem. B.* 2004; 108:14804–14806.
25. Wick CD, Dang LX. *J. Phys. Chem. A.* 2009; 113:6356–6364. [PubMed: 19391589]
26. Wick CD. *J. Phys. Chem. C.* 2012; 116:4026–4038.
27. Iuchi S, Chen H, Paesani F, Voth GA. *J. Phys. Chem. B.* 2009; 113:4017–4030. [PubMed: 18821788]
28. Xu J, Izvekov S, Voth GA. *J. Phys. Chem. B.* 2010; 114:9555–9562. [PubMed: 20593833]
29. Semino R, Laria D. *J. Chem. Phys.* 2012; 136:194503. [PubMed: 22612099]
30. Yamashita T, Voth GA. *J. Am. Chem. Soc.* 2012; 134:1147–1152. [PubMed: 22191804]
31. Li H, Chen H, Steinbronn C, Wu B, Beitz E, Zeuthen T, Voth GA. *J. Mol. Biol.* 2011; 407:607–620. [PubMed: 21277313]
32. Chen H, Ilan B, Wu Y, Zhu F, Schulten K, Voth GA. *Biophys. J.* 2007; 92:46–60. [PubMed: 17056733]
33. Chen H, Wu Y, Voth GA. *Biophys. J.* 2006; 90:L73–L75. [PubMed: 16581846]
34. Chen H, Wu Y, Voth GA. *Biophysical Journal.* 2007; 93:3470–3479. [PubMed: 17693473]
35. Tepper HL, Voth GA. *Biophys. J.* 2005; 88:3095–3108. [PubMed: 15695636]
36. Yamashita T, Voth GA. *J. Phys. Chem. B.* 2010; 114:592–603. [PubMed: 19924872]
37. Petersen MK, Voth GA. *J. Phys. Chem. B.* 2006; 110:18594–18600. [PubMed: 16970488]
38. Feng S, Voth GA. *J. Phys. Chem. B.* 2011; 115:5903–5912. [PubMed: 21510678]
39. Guo Y, Thompson DL. *J. Chem. Phys.* 2003; 118:1673–1678.
40. Nelson KV, Benjamin I. *J. Phys. Chem. C.* 2010; 115:1154–1163.
41. Warshel A, Weiss RM. *J. Am. Chem. Soc.* 1980; 102:6218–6226.

42. Warshel, A. Computer Modeling of Chemical Reactions in Enzymes and Solutions. John Wiley and Sons; New York, NY, USA: 1991. p. 1-236.
43. Smondyrev AM, Voth GA. Biophys. J. 2002; 82:1460–1468. [PubMed: 11867461]
44. Chen H, Liu P, Voth GA. J. Chem. Theory Comput. 2010; 6:3039–3047.
45. Izvekov S, Swanson MJ, Voth GA. J. Phys. Chem. B. 2008; 112:4711–4724. [PubMed: 18366209]
46. Shi Q, Liu P, Voth GA. J. Phys. Chem. B. 2008; 112:16230–16237. [PubMed: 19367797]
47. Markland TE, Manolopoulos DE. J. Chem. Phys. 2008; 129:024105. [PubMed: 18624514]
48. Markland TE, Manolopoulos DE. Chem. Phys. Lett. 2008; 464:256–261.
49. Wang F, Voth GA. J. Chem. Phys. 2005; 122:144105. [PubMed: 15847514]
50. Hess B, Kutzner C, Spoel D. v. d. Lindahl E. J. Chem. Theory Comput. 2008; 4:435–447.
51. Peng, Y.; Knight, C.; Blood, P.; Crosby, L.; Voth, GA. XSEDE12. ACM; Chicago, IL: 2012.
52. Plimpton SJ. J. Comp. Phys. 1995; 117:1–19.
53. Brown WM, Nguyen TD, Fuentes-Cabrera M, Fowlkes JD, Rack PD, Berger M, Bland AS. Procedia Comput. Sci. 2012; 9:186–195.
54. Perram JW, Petersen HG, Leeuw SWD. Mol. Phys. 1988; 65:875–893.
55. Darden T, York D, Pedersen L. J. Chem. Phys. 1993; 98:10089–10092.
56. Hockney, RW.; Eastwood, JW. Computer Simulation Using Particles. Taylor & Francis Group; New York: 1988. p. 1-540.
57. Maupin CM, Wong KF, Soudackov AV, Kim S, Voth GA. J. Phys. Chem. A. 2006; 110:631–639. [PubMed: 16405335]
58. Kumar S, Rosenberg JM, Bouzida D, Swendsen RH, Kollman PA. J. Comput. Chem. 1994; 16:1339–1350.
59. Roux B. Comput. Phys. Commun. 1995; 91:275–282.
60. Mendoza FN, López-Lemus J, Chapela GA, Alejandre J. J. Chem. Phys. 2008; 129:024706. [PubMed: 18624551]
61. Feller SE, Pastor RW, Rojnuckarin A, Bogusz S, Brooks BR. J. Phys. Chem. 1996; 100:17011–17020.
62. Fennell CJ, Gezelter JD. J. Chem. Phys. 2006; 124:234104. [PubMed: 16821904]
63. Viveros-Méndez PX, Gil-Villegas A. J. Chem. Phys. 2012; 136:154507. [PubMed: 22519336]
64. Ramsey S, Mokrab Y, Carvacho I, Sands ZA, Sansom MSP, Clapham DE. Nat. Struct. Mol. Biol. 2010; 17:869–875. [PubMed: 20543828]

Figure 1.

A time cost analysis of multistate simulations for a system containing 30,000 water molecules and a hydrated excess proton as a function of the number of diabatic MS-EVB states showing the relative costs of the real and k-space calculations. This analysis was performed on a single processor for 10 MD steps. The left vertical axis indicates the time cost for the real-space (blue squares) and k-space (red circles) portions of the calculation as a function of the number of diabatic states. The right vertical axis indicates the fraction of time spent evaluate the k-space interaction (dashed green line) as a function of the number of diabatic states.

Figure 2.

Properties of the hydrated excess proton in bulk water calculated by the CGISa/CGISo: The radial distribution functions (RDFs) between (a) OT and OW, (b) OT and HW, (c) HT and OW, and (d) HT and HW are shown as functions of distance. These atom labels are defined as follows: hydronium oxygen (OT), hydronium hydrogen (HT), water oxygen (OW), and water hydrogen (HW). The distributions of MS-EVB coefficients (e) $|c_{\max}|^2$ and (f) $|c_{\max}|^2 - |c_{\max2}|^2$ are also shown. See text for more details.

Figure 3.

Properties calculated by the ELRI/ELRICK: The RDFs between (a) OT and OW, (b) OT and HW, (c) HT and OW, and (d) HT and HW are shown as functions of distance. These atom labels are defined as follows: hydronium oxygen (OT), hydronium hydrogen (HT), water oxygen (OW), and water hydrogen (HW). The distributions of MS-EVB coefficients (e) $|c_{\max}|^2$ and (f) $|c_{\max}|^2 - |c_{\max2}|^2$ are also shown. See text for more details.

Figure 4.

Properties calculated by the hybrid CE/CEck: The radial distribution functions between (a) OT and OW, (b) OT and HW, (c) HT and OW, and (d) HT and HW are shown as functions of distance. These atom labels are defined as follows: hydronium oxygen (OT), hydronium hydrogen (HT), water oxygen (OW), and water hydrogen (HW). The distributions of (e) $|c_{\max}|^2$ and (f) $|c_{\max}|^2 - |c_{\max2}|^2$ are also shown. See text for more details.

Figure 5. Potentials of mean force (PMFs) for histidine deprotonation as a function of separation between the excess proton center of excess charge (CEC) and nitrogen atom are shown for the exact k-space method (solid black), CGISo (short-dashed red), CE (long-dashed green), and CEck (dashed-dotted blue) methods.

Figure 6. Probability distributions for excess hydrated proton as a function of distance near the water/vapor interface for (a) CGIS, (b) ELRI, and (c) hybrid methods.

Figure 7. Productivity of (a) nonreactive and (b) multistate reactive MD simulations of a 159K atom CcO system calculated on Diamond (SGI ICE) at ERDC for SP (blue diamonds) and MP (red circles) parallelization strategies.

Figure 8. Scaling results of exact (circles), CGISo (squares), and CEck (triangles) methods for multistate reactive simulations using both SP (solid lines) and MP (dashed lines) parallelization strategies calculated on Intrepid (Blue Gene/P) at ALCF.

Table 1

Comparison of calculated pK_a values for hydronium in bulk water.

Method	pK_a
Original MS-EVB3	-1.68
ELRI	-1.68
ELRIck	-1.68
CGISa	-1.68
CGISo	-1.68
CE (CGISo+ELRI)	-1.68
CEck (CGISo+ELRIck)	-1.68
Experiment	-1.74

All the standard errors are less than 0.001.

The experimental value is estimated by definition with the water density (water concentration).

Table 2

Comparison of calculated diffusion coefficients for excess proton in bulk water estimated from the mean squared displacement in the time interval from 1 to 10 ps.

Method	Diffusion coefficient ($\text{\AA}^2/\text{ps}$)
Exact MS-EVB3 Model	0.29
ELRI	0.29
ELRIck	0.30
CGISa	0.30
CGISo	0.29
CE (CGISo+ELRI)	0.29
CEck (CGISo+ELRIck)	0.30

All the standard errors are less than $0.005 \text{ \AA}^2/\text{ps}$.

Table 3

Performance of approximate methods for each hydrated excess proton system. Given are the number of atoms, average number of diabatic MS-EVB states, and computational cost (ratio) relative to nonreactive MD simulations in LAMMPS. The benchmarks were calculated on Chugach (Cray XE6) located at the Engineer Research and Development Center (ERDC) DoD Supercomputing Resource Center.

System	# of Atoms	# of States	Exact	CGISO	CE	CEck
Liquid/vapor interface	3001	23	31.36	13.37	7.31	6.36
CcO	158982	6	2.07	2.03	1.85	1.78
50,000 H ₂ O	150001	23	23.34	8.84	5.20	4.75