


Free Energies of Quantum Particles: The Coupled-Perturbed Quantum Umbrella Sampling Method

William J. Glover, Jennifer R. Casey, and Benjamin J. Schwartz*

Department of Chemistry and Biochemistry, University of California, Los Angeles, Los Angeles, California 90095-1569, United States

ABSTRACT: We introduce a new simulation method called Coupled-Perturbed Quantum Umbrella Sampling that extends the classical umbrella sampling approach to reaction coordinates involving quantum mechanical degrees of freedom. The central idea in our method is to solve coupled-perturbed equations to find the response of the quantum system's wave function along a reaction coordinate of interest. This allows for propagation of the system's dynamics under the influence of a quantum biasing umbrella potential and provides a method to rigorously undo the effects of the bias to compute equilibrium ensemble averages. In this way, one can drag electrons into regions of high free energy where they would otherwise not go, thus enabling chemistry by fiat. We demonstrate the applicability of our method for two

Quantum PMF for interfacial hydrated electron



condensed-phase systems of interest. First, we consider the interaction of a hydrated electron with an aqueous sodium cation, and we calculate a potential of mean force that shows that an $e^-:\text{Na}^+$ contact pair is the thermodynamically favored product starting from either a neutral sodium atom or the separate cation and electron species. Second, we present the first determination of a hydrated electron's free-energy profile relative to an air/water interface. For the particular model parameters used, we find that the hydrated electron is more thermodynamically stable in the bulk rather than at the interface. Our analysis suggests that the primary driving force keeping the electron away from the interface is the long-range electron–solvent polarization interaction rather than the short-range details of the chosen pseudopotential.

1. INTRODUCTION

Much of chemistry is governed by activated processes that involve surmounting the free energy barriers that separate reactants from products. To make dynamical or mechanistic predictions of such processes, one first needs to identify the barriers and determine the probability of finding a system in such activated regions. This can be achieved by computing the free energy F of the system along a proposed reaction coordinate q , which is known as the potential of mean force (PMF):

$$F(q) = F_0 - k_B T \ln P(q) \quad (1)$$

where $P(q)$ is the probability of finding the system at coordinate q , k_B is Boltzmann's constant, T is the temperature, and F_0 is an additive constant that can be taken to be the free energy minimum along the reaction coordinate.

There is a vast literature on sampling techniques designed to efficiently compute PMFs under the assumption that classical mechanics holds;¹ however, few methods² have been developed to handle reaction coordinates involving quantum mechanical degrees of freedom (such as the location of an electron during a charge-transfer reaction). A notable exception is when the quantum particles are treated via classical isomorphism, such as in path integral molecular dynamics.³ However, due to issues associated with the small mass of the electron and the requirement of Fermi statistics, path integral descriptions of

many-electron systems are not computationally practical for routine use.⁴ Instead, wave function- or density functional-based methods are what have received wide application. When quantum systems are described by wave functions, the ideal situation would be to alter the wave function *a priori*, dragging the quantum mechanical electrons to locations of interest and calculating the associated free energy cost. In other words, what we desire is a way to understand quantum systems where the electrons are far from equilibrium by constructing a PMF based on moving electrons to regions where they otherwise would rarely go.

In this article, we work toward this goal by presenting a new method for extending classical umbrella sampling to quantum systems represented by wave functions. Our method makes use of coupled-perturbed equations to solve for the wave-function response to motion along a (classical) reaction coordinate, so we name it Coupled-Perturbed Quantum Umbrella Sampling (CP-QUMB). Our method allows one to restrain quantum particles, such as electrons, to regions of high free energy that would otherwise rarely be sampled under equilibrium conditions, opening up the way to do chemistry by fiat, and to understand how pushing electrons into different locations changes the free energy of a quantum system. We begin our presentation in Section 2 by discussing the theoretical

Received: August 12, 2014

Published: September 10, 2014

background and development of the coupled-perturbed quantum umbrella sampling method. We then compare the accuracy and efficiency of our CP-QUMB method to a related QUMB method developed previously² in Section 3. We next demonstrate how the CP-QUMB method can be applied to two systems of interest: the interaction of a hydrated electron and aqueous sodium cation in Section 4.1, and the free energy profile of a hydrated electron relative to an air/water interface in Section 4.2. Finally, after summarizing our results in Section 5, we present details of the gradient potential matrix elements needed by our method in the Appendix.

2. THEORETICAL DEVELOPMENT

2.1. Umbrella Sampling along a Quantum Mechanical Coordinate. To generate a potential of mean force along a reaction coordinate, the most naïve approach one could take is to apply eq 1 directly and simply take the logarithm of the number of times the system visits each point of the reaction coordinate during equilibrium molecular dynamics. This approach is problematic, however, because regions of high free energy are rarely visited, giving rise to large uncertainties or even gaps in PMFs calculated this way unless exhaustive sampling is used.⁵ As a result, many ways to improve the sampling of regions of high free energy have been developed,¹ and we focus here on the method of Umbrella Sampling.^{1b,6} We note that our coupled-perturbed method could easily be extended to other enhanced sampling schemes that make use of gradients in the reaction coordinate.

In Umbrella Sampling, one adds a biasing potential, U^{bias} , to the system's Hamiltonian that is designed to enhance sampling in certain desired regions of configuration space. The biasing potential typically takes the form of a harmonic confining potential:

$$U^{\text{bias}} = \frac{1}{2}k(q(\mathbf{R}^N) - \zeta)^2 \quad (2)$$

where q is a reaction coordinate of interest that may depend on some or all of the degrees of freedom, \mathbf{R}^N , and ζ is the biasing target value of the reaction coordinate. For a classical system, one can rigorously “undo” the effects of the biasing potential to calculate equilibrium properties in the original unbiased ensemble. In particular, for a classical system described by Hamiltonian H , the canonical ensemble average of any observable, A , is given by

$$\begin{aligned} \langle A \rangle &= \frac{\text{Tr}(Ae^{-\beta H})}{\text{Tr}(e^{-\beta H})} \\ &= \frac{\text{Tr}(Ae^{\beta U^{\text{bias}}} e^{-\beta(H+U^{\text{bias}})})}{\text{Tr}(e^{\beta U^{\text{bias}}} e^{-\beta(H+U^{\text{bias}})})} \\ &= \frac{\langle Ae^{\beta U^{\text{bias}}} \rangle_{U^{\text{bias}}}}{\langle e^{\beta U^{\text{bias}}} \rangle_{U^{\text{bias}}}} \end{aligned} \quad (3)$$

where $\text{Tr}()$ denotes the trace over all degrees of freedom, $\beta = 1/k_B T$, and $\langle \rangle_{U^{\text{bias}}}$ indicates an ensemble average taken with Hamiltonian $H + U^{\text{bias}}$. The second equality in eq 3 results from multiplying and dividing by $\exp[\beta U^{\text{bias}}]$, while the last equality shows how to relate ensemble averages calculated with the biasing potential to the (desired) unbiased values.

Unfortunately, when the reaction coordinate of interest involves a quantum degree of freedom, the approach outlined above can no longer be applied directly. The reason is that the

second equality of eq 3 uses $\exp[-\beta H] \times \exp[-\beta U] = \exp[-\beta(H + U)]$, which is true only for systems where H and U commute. H and U always commute for classical systems, but they almost never commute in cases where they are quantum mechanical operators. This is because to confine a quantum system to a certain region of space, the confining potential U needs to be a function of the position operator. Because of the noncommutation of position and momentum operators, however, there exists no function, $U(\hat{r})$, that commutes with the kinetic energy operator of a quantum Hamiltonian.

To circumvent this problem of the noncommutation of the biasing potential with the quantum Hamiltonian, Borgis and Staib (B&S) developed an approach that uses a biasing potential that is a function of the quantum position expectation value, $r_e = \langle \hat{r} \rangle$.² Since the quantum degrees of freedom have been integrated out in the expectation value, the biasing potential is a function only of classical degrees of freedom, and eq 3 can be used as written. Since most electronic quantities of interest are usually taken to be expectation values, such as the distance of an electron from a particular molecule or interface, applying biasing potentials directly on expectation-value coordinates is a natural choice.

For systems involving both quantum and classical particles, the canonical ensemble sampling indicated by eq 3 is carried out most efficiently with molecular dynamics (MD) since the quantum wave function (whose determination is often the most expensive part of the calculation) can be updated with the simultaneous propagation of all classical degrees of freedom, as opposed to a Monte Carlo approach, which moves particles one by one. To perform MD, we require gradients of the biasing potential of eq 2 with respect to the classical particles' positions, \mathbf{R}^N . Calculating this force, however, is not trivial since the quantum expectation value is determined parametrically from the classical positions (via their coupling in the Schrodinger equation). Thus, when taking the gradient of eq 2, there are two terms that contribute to the derivative of q .

$$\frac{dq}{dR_\alpha} = \frac{\partial q}{\partial R_\alpha} + \nabla_r q \cdot \frac{dr_e}{dR_\alpha} \quad (4)$$

where α denotes a particular classical degree of freedom. The first term of eq 4 comes from the explicit dependence of q on the positions of each of the classical particles and is usually trivial to calculate. The second term encompasses the dependence of the quantum position expectation value on the positions of all of the classical particles, and this term is what has prevented quantum umbrella sampling from being as straightforward as classical umbrella sampling.

Thus, to make progress, we need to focus on this expectation value gradient:

$$\frac{dr_e}{dR_\alpha} = 2\langle \psi^I | \hat{r}_e | d\psi^I / dR_\alpha \rangle = 2 \sum_j (c_j^I)^* r_j \frac{dc_j^I}{dR_\alpha} \quad (5)$$

where I is an index denoting the occupied quantum electronic state, and the second equality results from an expansion of the wave function in an orthonormal grid basis (so that the wave function for state I has value c_j^I at grid point j). Since the position operator is diagonal in a grid basis and acts only on the quantum degrees of freedom, the gradient with respect to the classical positions arises solely from the response of the wave function. B&S used a perturbation expansion to find a solution to the wave function response:²

$$\frac{dr_e}{dR_\alpha} \approx 2 \sum_{J \neq I}^{N^{\text{St}}} (E_I - E_J)^{-1} \langle \psi^I | \hat{r} | \psi^J \rangle \langle \psi^J | \frac{\partial U}{\partial R_\alpha} | \psi^I \rangle \quad (6)$$

where E_J is the electronic energy of state J , and U is the potential energy operator. We refer to this approach as sum-over-states quantum umbrella sampling (SOS-QUMB). Equation 6 becomes exact if the sum is carried out over a complete set of states. This type of sum was possible in the first application of the method due to the use of a small basis size, of order $N^{\text{basis}} = 100$.² More recent applications of SOS-QUMB have used basis set sizes of order $N^{\text{basis}} = 1000$, which necessitates a truncated sum-over-states ($N^{\text{St}} \ll N^{\text{basis}}$) to avoid a complete diagonalization of the Hamiltonian, whose cost scales as $(N^{\text{basis}})^3$.⁷

The use of a truncated expansion in eq 6 makes most practical applications of SOS-QUMB approximate in nature. Although the energy denominator and transition dipole matrix elements should give the greatest weight to contributions from states close to the occupied state, I , justifying a truncated sum, no study of the convergence properties of this equation has previously been reported. Furthermore, for systems exhibiting substantial asymmetry, such as for ions or electrons near an interface, the spatial and energetic distribution of the electronic states may require a very large or fully complete set of states to converge the sum in eq 6. For this reason, it is clearly desirable to find an alternative formulation to solve eq 5 without requiring the calculation of any excited states. As we show in the next section, an analytical solution to this problem is possible, with the final result being:

$$\frac{dr_e}{dR_\alpha} = -2 \langle Z | \frac{\partial U}{\partial R_\alpha} - \frac{dE}{dR_\alpha} | \psi^I \rangle \quad (7)$$

where Z is a solution to the coupled-perturbed equations (eq 18 below). The key advantages of eq 7, which forms the heart of CP-QUMB, are that the required matrix elements are no more complicated than those in SOS-QUMB, and the need for calculating electronic excited states is avoided entirely.

2.2. Development of Coupled-Perturbed Quantum Umbrella Sampling. Our formulation of Coupled-Perturbed Quantum Umbrella Sampling is generally applicable to any electronic structure method, but for the purposes of presentation, we choose to focus on the specific case of an adiabatic mixed quantum/classical simulation (MQC) involving a single ground-state quantum particle embedded in a classical subsystem. We note that extension to adiabatic excited-state free energy surfaces is trivial providing the reaction coordinate of interest does not bring the system to conical intersections with other electronic states. In MQC MD, the quantum and classical degrees of freedom are coupled by pseudopotentials,⁸ and we assume without loss of generality that the quantum system's wave function is solved on a grid by diagonalizing the Hamiltonian in this basis:

$$\sum_j H_{ij} c_j = E c_i \quad (8)$$

The wave function response can then be found by taking derivatives of eq 8 with respect to the classical degrees of freedom:

$$\frac{d}{dR_\alpha} \left[\sum_j (H_{ij} - \delta_{ij} E) c_j \right] = 0 \quad (9)$$

Since H , E , and c are known (following a solution of the eigenvalue problem), finding the wave function response amounts to inverting eq 9. This equation, however, is singular, so to proceed one must modify its eigenvalue away from zero without changing the response coefficients. Following Osamura et al.,⁹ this is achieved by adding the normalization condition of the wave function:

$$\frac{d}{dR_\alpha} \left\{ \sum_j [H_{ij} - \delta_{ij} E - \delta_{ij} (1 - \sum_k |c_k|^2)] c_j \right\} = 0 \quad (10)$$

Collecting terms in eq 10 produces the Coupled-Perturbed Response equations:

$$\sum_j [H_{ij} - \delta_{ij} E + 2c_i c_j^*] \frac{\partial c_j}{\partial R_\alpha} = - \sum_j \left[\delta_{ij} \left(\frac{\partial U_j}{\partial R_\alpha} - \frac{dE}{dR_\alpha} \right) \right] c_j \quad (11)$$

Equation 11 makes clear that the response of the wave function coefficients on the left-hand side comes from the coupling between quantum and classical particles through the gradients of the potential on the right-hand side.

Although eq 11 could be solved for every R_α , there are $3N$ (where N is the number of atoms) such perturbations, making this approach computationally too expensive. Instead, we use the Z -vector method of Handy and Schaefer and solve an analog of eq 11 for a single perturbation.¹⁰ We start by writing eq 11 in matrix form:

$$\underline{D} \frac{\partial \underline{c}}{\partial R_\alpha} = -\underline{M}^\alpha \quad (12)$$

where

$$D_{ij} = H_{ij} - \delta_{ij} E + 2c_i c_j^* \quad (13)$$

and

$$M_j^\alpha = \left(\frac{\partial U_j}{\partial R_\alpha} - \frac{dE}{dR_\alpha} \right) c_j \quad (14)$$

Introducing a Lagrangian for the position operator:

$$L_i = c_i r_i \quad (15)$$

it is clear from eq 5 that the gradient of the position operator is a simple inner product:

$$\frac{dr_e}{dR_\alpha} = 2\underline{L} \cdot \frac{\partial \underline{c}}{\partial R_\alpha} \quad (16)$$

So inverting eq 12 and plugging into eq 16 gives

$$\frac{dr_e}{dR_\alpha} = -2\underline{L} \cdot \underline{D}^{-1} \underline{M}^\alpha \quad (17)$$

The Z -vector approach then corresponds to solving a single Coupled-Perturbed equation

$$\underline{D} \underline{Z} = \underline{L} \quad (18)$$

for each of the three Cartesian directions of eq 17. With these solutions, the gradients of the position expectation value can be recovered by inserting eq 18 into eq 17:

$$\frac{dr_e}{dR_\alpha} = -2\underline{Z} \cdot \underline{M}^\alpha \quad (19)$$

which is the matrix form of eq 7.

In practice, we solve eq 18 iteratively using the Conjugate Gradient algorithm,¹¹ as implemented in PETSc 3.4.4,¹² which requires only matrix-vector products and thus eliminates the need to store the complete Hamiltonian matrix. This procedure thus gives us the forces we need to restrain the expectation value of a quantum object with respect to the classical coordinates, and allows us to use eq 3 to directly determine the free energetic cost associated with this restraint. In this way, we can do chemistry by fiat, and construct PMFs for quantum mechanical particles by forcing the system to have a desired expectation value of its quantum wave function. A step-by-step approach to propagating CP-QUMB dynamics is as follows:

- 1) From a set of classical particle coordinates, update the Hamiltonian matrix elements and solve the Time-Independent Schrödinger Equation (eq 8) for the electron's energy and wave-function coefficients.
- 2) Find the electronic forces on the classical particles from the Hellmann–Feynman theorem (eq 21).
- 3) Construct the left- and right-hand sides of the Coupled-Perturbed equations (eq 18) from the position operator matrix elements, the Hamiltonian matrix elements, the energy eigenvalue, and the wave-function coefficients. Solve this iteratively for a Z vector for each component of the quantum position operator.
- 4) Find the nuclear gradients of the quantum position expectation value according to eq 19 and eq 14 using the wave-function coefficients from step 1, energy gradients from step 2, and Z vectors from step 3. From these expectation-value gradients, find the gradients of the umbrella potential (eq 2) via the chain rule.
- 5) Propagate the classical positions and velocities one time step using the sum of their pair-potential forces, the electronic forces from step 2, and the umbrella forces from step 4.
- 6) Return to step 1 for the next MD time step.

3. COMPARING QUANTUM UMBRELLA SAMPLING METHODS

3.1. Computational Efficiency. Now that we have a formalism in place that allows for quantum umbrella sampling with exact forces on the classical particles to restrain the quantum expectation value, the next thing we investigate is how practical this formalism is to implement. The computational cost of our CP-QUMB method over regular MQC dynamics arises from two aspects. First, one must solve the coupled-perturbed equations for the wave-function response. As detailed above, this involves solving three linear equations (one for each Cartesian component of the quantum position operator),¹³ which given the large size and sparsity of the Hamiltonian is best performed with an iterative solver such as the method of Conjugate Gradients.¹¹ Second, according to eq 7, one must evaluate up to $9N$ additional matrix elements of the gradient potential energy operator between the Z -vector solutions to the coupled-perturbed equations and the occupied state of interest. This is on top of the $3N$ gradient potential matrix elements needed for Hellman–Feynman forces between the quantum and classical particles (which need to be evaluated with or without QUMB). As we show below, however, evaluating these gradient potential matrix elements does not consume a significant portion of the computational cost per time step. In the Appendix we give details on how these gradient potential matrix elements readily can be evaluated in a Fourier grid basis.

Although there is a lot to calculate to successfully perform CP-QUMB, it is still quite a bit less computationally expensive than SOS-QUMB, which is an approximate method and for which the cost grows with the number of states used in the sum-over-states. This is because SOS-QUMB requires the calculation of excited-state wave functions, which makes the iterative diagonalization of the Hamiltonian more costly than solving the coupled-perturbed equations in CP-QUMB. Moreover, according to eq 6, the number of additional gradient potential matrix elements that must be calculated when using SOS-QUMB is $3N \times N^{\text{st}}$, which can quickly become the slowest part of the calculation for the typical number of states needed to reasonably converge this method.

To explore the computational cost of the QUMB methods in more detail, we evaluated CPU timings for one of the umbrella sampling simulations of Section 4.2, which involves a hydrated electron restrained so that its center of mass lies ~ 4 Å above an air/water interface, a height that corresponds to an electron “floating” on the water surface. In Figure 1a we show a typical molecular dynamics snapshot of such a floating electron, where it is clear the electron is localized to within a solvation shell of four water molecules, despite being partially exposed to the vacuum. We performed timings on one core of a dual quad-core Intel Xeon X5560 2.8-GHz computer, averaging over 50 water configurations. To remove any bias, we used the same water configurations for each method and refactored our code to keep implementation details between the methods as similar as possible.

The timing results from this exercise are collected in Table 1, where we see that CP-QUMB is faster than SOS-QUMB with only five states and introduces only a small computational overhead relative to MQC dynamics without QUMB (first row). The cost of solving for the coupled-perturbed wave-function responses of eq 18 is comparable to the iterative diagonalization of the Hamiltonian and introduces roughly three times as many matrix-vector products as the diagonalization. Although the time spent evaluating gradient potential energy matrix elements is roughly 4 times larger in CP-QUMB compared to trajectories without QUMB, this part of the calculation is a small fraction of the total CPU time, most of which for this particular system is taken up evaluating the electron–water potential energy matrix elements, a cost common to all methods. Finally, as expected, the cost of SOS-QUMB scales roughly linearly with the number of states in the sum-on-states, which for this particular umbrella window necessitated $N^{\text{st}} = 25$ to converge to the point of accurate determination of the PMF. Thus, for this problem, the CP-QUMB method is over a factor of 2 faster than the approximate SOS-QUMB method.

3.2. Accuracy of CP-QUMB and SOS-QUMB. Now that we know that CP-QUMB is generally faster than SOS-QUMB, we next explore the accuracy of the two methods. In particular, we examine the convergence properties of the sum-over-states in SOS-QUMB since recent applications of this method have been forced to use a truncated sum,⁷ but no convergence studies have been published. To this aim, we ran CP-QUMB and SOS-QUMB dynamics for the same system considered in the previous section: a hydrated electron restrained to float on a water surface. One way to detect errors in QUMB methods is to monitor the total energy conservation during molecular dynamics: errors in the gradient of the quantum position operator expectation value will propagate into errors in the umbrella sampling forces, resulting in energy nonconservation.




Figure 1. Molecular dynamics snapshots of the hydrated electron near an air/water interface obtained from the Coupled-Perturbed Quantum Umbrella Sampling method. In all panels, the translucent blue surface represents the location of the air/water interface (Gibbs dividing surface), while the dark and light gray contour surfaces enclose 50% and 95% of the excess electron's charge density, respectively. (a, b) An electron with its expectation value position restrained 4 Å above and 1 Å below the air/water interface, respectively. (c) An electron located 9 Å below the interface, where essentially bulk solvation is observed. For clarity, in (c) we represent the water molecules by licorice bonds.

In Figure 2a we plot a portion of the total energy conservation for CP-QUMB (black curve), where we see excellent energy conservation, with no noticeable drift, and root mean square (RMS) fluctuations of 0.015 eV, comparable to what is obtained for mixed quantum/classical dynamics for this system in the absence of QUMB. This level of energy

Table 1. Computational Cost of Quantum Umbrella Sampling Methods^a

method	tot. time	grad. time ^b	diag. time ^c	CG time ^d	av # mat-vec ^e
no QUMB ^f	2.10	0.12	0.11		70
CP-QUMB	2.69	0.53	0.11	0.17	256
SOS $N^{\text{St}} = 5$	2.73	0.58	0.30		278
SOS $N^{\text{St}} = 10$	3.58	1.29	0.55		486
SOS $N^{\text{St}} = 25$	5.81	3.23	1.01		731

^aCPU times in seconds per time step for QUMB simulations of a hydrated electron restrained so that its average position lies 4 Å above an air/water interface; the times reported are averaged over 50 water configurations. ^bCPU time to evaluate all matrix elements involving the gradient potential energy operator, i.e., eq 21. ^cCPU time to iteratively diagonalize the electronic Hamiltonian. ^dCPU time to iteratively solve the Z-vector of eq 18 with the Conjugate Gradient method. ^eTotal number of matrix-vector multiplies per time step, from both the iterative diagonalization and conjugate gradient routines. ^fCPU times for pure MQC MD, with all QUMB routines disabled.

conservation is maintained for the entire >100-ps simulation, giving us confidence in the accuracy of CP-QUMB. Also plotted in Figure 2a are total energy traces for SOS-QUMB with three different levels of truncation in the sum-over-states. Including only five states (red curve) gives very poor energy conservation, with an unacceptably large drift of >0.6 eV in the course of 10 ps. Increasing the number of states to 10 (blue curve) improves the drift somewhat, but only when $N^{\text{St}} = 25$ (green curve) is energy conservation comparable to CP-QUMB, and even then the RMS fluctuations are still somewhat larger at 0.018 eV. We note $N^{\text{St}} = 25$ is what was used by Boutin and co-workers in their studies of electron:cation contact pairs,⁷ so we expect the results of ref 7 to be reasonably accurate, even though the computational cost was ~ 3 times larger than what it would have been with CP-QUMB.

The observation of total energy drifts in the SOS-QUMB method implies that the truncation of the sum-over-states introduces a systematic error in the umbrella force. This translates to an incorrect sampling of the umbrella coordinate and thus leads to errors in the calculated potential of mean force. To explore these sampling errors, we calculated a (partial) PMF from a single umbrella sampling window that restrains a hydrated electron to lie ~ 4 Å above the air/water interface for each of the methods (the full PMF calculated via CP-QUMB is presented in Section 4.2, below); the results are displayed in Figure 2b. As expected, SOS-QUMB with only five states (red curve, circles) exhibits a substantial error, overestimating the free-energy difference between 3.0 and 4.0 Å above the interface by $\sim 3 k_{\text{B}}T$ relative to the exact CP-QUMB method. This is because for this number of states, the approximate umbrella sampling forces underestimate their correct values and are therefore unable to properly restrain the hydrated electron's expectation value. When the number of states used in SOS-QUMB is increased to 10 (blue curve, crosses), the error goes the other way, and now the free energy difference is underestimated by $\sim 2 k_{\text{B}}T$ relative to CP-QUMB. This nonmonotonic convergence of SOS-QUMB therefore makes it difficult to know *a posteriori* what should be the correct number of states to use. For this particular system, we find it takes $N^{\text{St}} = 25$ in SOS-QUMB to converge the PMF into agreement with CP-QUMB (not shown in Figure 2b for clarity), but, as discussed in the previous section, increasing N^{St} in SOS-QUMB costs nearly 3 times as much as CP-QUMB.




Figure 2. Convergence properties of SOS-QUMB with the number of electronic states and comparison to CP-QUMB for a hydrated electron restrained to have its expectation value lie 4 Å above the air/water interface, i.e., a “floating” electron. (a) Total energy conservation, relative to the first time step, during quantum umbrella sampling simulations. (b) Partial PMFs of the hydrated electron as a function of its quantum expectation value position relative to the air/water interface, calculated from the simulations of panel (a). SOS-QUMB with $N^{\text{St}} = 25$ (not shown) agrees with CP-QUMB to within $\pm 1 k_B T$. Error bars in this panel and in all subsequent figures represent 95% confidence intervals after subsampling the molecular dynamics data to remove correlations.¹⁴

Since the convergence of the sum-over-states will depend on the particular system and the nature of its excited states, taken alone, there is no simple way to guarantee that even $N^{\text{St}} = 25$ will suffice to converge a QUMB calculation with the SOS-QUMB method. For these reasons, we strongly prefer the use of CP-QUMB over SOS-QUMB.

4. APPLICATIONS OF CP-QUMB

In this section, we apply CP-QUMB to two realistic MQC systems. To this end, we calculate potentials of mean force for moving electrons off sodium atoms in liquid water and for moving hydrated electrons close to the air/water interface.

4.1. Hydrated Electron–Aqueous Cation Contact Pair.

As a first application of the CP-QUMB method, we consider the interaction of a hydrated electron with a cationic species in

water, specifically Na^+ . We chose this as our initial test system for CP-QUMB because Boutin and co-workers have previously studied electron– Na^+ interactions quite extensively with SOS-QUMB with $N^{\text{St}} = 25$, as discussed above.⁷ These researchers found that the hydrated electron forms a contact pair with aqueous Na^+ : there is a minimum in the potential of mean force at around 2 Å separation, and the contact pair is stable to dissociation by $\sim 3 k_B T$.^{7a} Boutin and co-workers did not explore the region of the PMF where the electron completely overlaps the cation (i.e., becomes a neutral atom), however, so the question as to whether the neutral alkali atoms are (quasi)-stable in liquid water, or if they spontaneously form electron:cation contact pairs is still open.

In addition to serving as a test of our method, this system is chemically interesting for at least two reasons. First, the structure of the hydrated electron has received considerable attention recently.^{8e,f,15} Most theoretical studies treat the excess electron as an isolated species, but experimentally, counter-cations are usually present, and the influence of these cations on the properties of hydrated electrons has not been fully elucidated. Second, in charge-transfer-to-solvent (CTTS) reactions, photoexcitation of an atomic anion generates a neutral atom and a solvated electron.¹⁶ The neutral atom left behind following CTTS was originally considered an inert partner,¹⁷ but more recent work has revealed that the neutral atom can undergo a subsequent solvent-induced ionization, producing a solvated e^- :cation contact pair.¹⁸ Thus, exploring the aqueous e^- – Na^+ PMF can further our understanding of CTTS reactions, the nature of solvated atoms, and the nature of the hydrated electron, all of which are topics of current interest.

Our simulation protocol for calculating an e^- – Na^+ PMF in water was chosen to follow closely the previous SOS-QUMB simulations of Boutin and co-workers.^{7a} In particular, we performed umbrella sampling along the electron center of mass–sodium cation distance with 24 windows centered at distances ranging from 0 to 6 Å, separated by 0.25 Å. The force constant of the harmonic restraining potential was set to 1.4 $\text{eV}/\text{Å}^2$ to give sufficient overlap in the electron–sodium distance fluctuations between neighboring windows. After a period of equilibration, adiabatic MQC MD was propagated on the electronic ground-state surface for at least 30 ps in each window. We used the SPC/Flex potential¹⁹ to describe 499 classical water molecules and a single sodium cation with Lennard-Jones parameters taken from Aqvist.²⁰ The Turi–Borgis (TB) pseudopotential was used to represent the electron–water interactions,^{8b} while the electron–sodium pseudopotential was taken from our previous work.^{18a} The quantum hydrated electron’s wave function was represented on a Fourier grid with 20^3 points spanning 2/3 of the box length, and the grid was periodically recentered on the electron in a manner described elsewhere.^{18a} The wave functions were found by diagonalizing the quantum Hamiltonian using the Davidson algorithm²¹ as implemented in SLEPc 3.4.4.²² Velocities were rescaled according to a global Langevin thermostat²³ to sample the canonical ensemble at a temperature of 298 K, and the box size of 24.64 Å was chosen to ensure a density of 1.0 g/cm^3 . Finally the PMFs from each umbrella sampling window were unbiased and stitched together using the Multistate Bennett Acceptance Ratio method.²⁴

Figure 3 displays the PMF for the e^- – Na^+ distance in water computed using CP-QUMB (solid black curve). The results show a well-defined minimum at $r \approx 2$ Å, corresponding to an e^- : Na^+ contact pair stabilized by $\sim 3 k_B T$ relative to dissociation,




Figure 3. Potential of mean force between a hydrated electron and sodium cation in liquid water calculated with the coupled-perturbed quantum umbrella sampling method (solid black curve) and, for the section of the PMF near $r = 0$, with the sum-on-states quantum umbrella sampling method (dashed red curve) with $N^{\text{st}} = 5$. The SOS-QUMB PMF was displaced vertically by $-1 k_{\text{B}}T$ for clarity. The minimum in the PMF at 2 Å corresponds to an $e^-:\text{Na}^+$ contact pair, and formation of an aqueous neutral sodium atom is clearly unstable.

in excellent agreement with the previous work of Boutin and co-workers using SOS-QUMB.^{7a,25} To explore the convergence properties of SOS-QUMB for this particular system, we calculated a partial PMF using the SOS-QUMB from the window centered at $r = 0$ Å using five electronic states in the sum-over-states. This PMF is plotted as the dashed red curve in Figure 3. Surprisingly, we see excellent agreement between this highly truncated SOS-QUMB calculation and the exact CP-QUMB to within $\pm 0.5 k_{\text{B}}T$. It just so happens that for this particular system, the convergence of SOS-QUMB with the number of states is rapid. This favorable convergence arises because of the electronic structure of the sodium atom, which has transition dipole matrix elements in the SOS that are dominated by just a few $3s-3p$ transitions.

Going beyond the previous work on this system,⁷ we also extended the aqueous $e^-:\text{Na}^+$ PMF all the way to the origin ($r = 0$ Å). Figure 3 shows that at the origin, the PMF is at a local maximum. This means that a neutral sodium atom in water has a strong thermodynamic driving force to form an $e^-:\text{Na}^+$ contact pair, which explains our previous observations of such contact pairs in simulations of the CTTS reaction of $\text{Na}^+_{(\text{aq})}$.^{18a} For the case of sodium atoms in liquid tetrahydrofuran (THF), we previously found that the driving force to form an $e^-:\text{Na}^+$ contact pair from a newly created Na atom was a favorable complexation of the partially exposed Na^+ ion by the oxygen atoms of a few first-solvation-shell THF molecules.^{18c} We believe that a similar mechanism is operative in aqueous solutions, and indeed, Coudert et al. have shown that the sodium cation of aqueous $e^-:\text{Na}^+$ contact pairs is solvated on average by the O atoms of two nearby water molecules.^{7b}

4.2. Hydrated Electrons at the Air/Water Interface.

Having demonstrated the application of our method on the test case of a hydrated electron:sodium cation contact pair, we next apply CP-QUMB to a system that has not previously been explored with umbrella sampling: a hydrated electron near the air/water interface. This is a system of importance both in

terms of a wider understanding of surface ionic adsorption,^{5,26} which has consequences for atmospheric aerosol photochemistry and in providing insight into hydrated electron interactions with aqueous/biological interfaces. Recently, there has been considerable debate as to whether or not hydrated electrons reside near the water surface as the heavy halide ions do.²⁷ Part of the reason the heavy halides prefer the interface is that polarization of such anions by dielectric interfaces contributes to their surface enhancement relative to the bulk.^{26a} In this sense, the hydrated electron, which is one of the most polarizable solutes,²⁸ might be expected to have a particularly large surface enhancement. Experimental evidence for surface-bound hydrated electrons, however, is inconclusive.^{27b-e,g} Although one photoelectron spectroscopy study identified spectral features assigned to a quasi-stable surface-bound electron with a distinct ionization potential of 1.6 eV (compared to the bulk value of 3.3 eV),^{27b} subsequent studies under similar experimental conditions were unable to reproduce the features associated with surface-bound electrons.^{27e,g}

Theoretically, the interfacial stability of hydrated electrons has received only limited attention. Rodriguez and Laria were first to simulate the properties of hydrated electrons at the air/water interface, but they did not explore its stability relative to the bulk.²⁹ Madarász, Rossky, and Turi (MRT) studied the localization and stabilization of electrons introduced at the air/water interface. These authors found that under ambient conditions the electron diffuses into the bulk on a ~ 10 -ps time scale while the reverse process was absent, suggesting that the bulk is preferred.^{27a} More recently, Uhlig, Marsalek, and Jungwirth (UMJ) studied the localization dynamics and properties of a hydrated electron at the air/water interface with *ab initio* QM/MM simulations.^{27f} They found that an initially created interfacial electron persisted for at least 10 ps, but the high computational expense of their method precluded longer simulation runs, so quasi-stability of the interfacial electron could not be ruled out. To date, there has been no attempt to calculate a potential of mean force for the electron relative to the air/water interface that could address directly the question of the relative stability of bulk versus interfacial solvation and provide insight into the energetic contributions favoring one solvation motif over the other.^{3,26b}

To use CP-QUMB to generate a PMF for a hydrated electron relative to the air/water interface, we chose to use identical potential interactions as MRT to facilitate a direct comparison to their previous dynamical studies,^{27a} including use of the Turi–Borgis (TB) electron-water pseudopotential.^{8b} Thus, we constructed a simulation slab geometry of 499 waters in a vacuum with an $L = 24.64$ Å simulation cell length in the x and y directions and a length 5 times that in the z direction. We then ran 32 umbrella sampling window simulations with a spring constant of $0.4 \text{ eV}/\text{Å}^2$ to restrain the z -component of the electron's position to be centered between 0.0 to 16.0 Å relative to the center of mass of all the water molecules in the z -direction. The Gibbs dividing surface (GDS, corresponding to the height at which the average water density falls to half of the bulk value) was located at $z = 12.3$ Å, and the 90% and 10% water density range fell between $11.0 < z < 14.1$ Å. For the highest umbrella window (~ 4 Å above the GDS), the radius of the electron in the z -direction was found to be 1.9 Å; this window therefore corresponds to an electron “floating” on the water surface, as seen in Figure 1a. To accommodate the increasing size of the hydrated electron as it approached and

traversed the interface, we had to adapt the spatial extent of the Fourier grid used to represent the electron. For umbrella windows centered between $13 \leq z \leq 16$ Å, we found a Fourier grid of 22^3 points spanning the length L was necessary (as previously observed,^{27a,f} the electron's ground-state wave function did not extend significantly into the vacuum). For $9 < z < 13$ Å, we used an 18^3 point grid spanning $9/11 L$, and for $z \leq 9$ Å, a 14^3 grid spanning $7/11 L$ was sufficient to describe the effectively bulk hydrated electron. At least 100 ps of data were accumulated in each umbrella sampling window after an initial period of equilibration.

Figure 4a shows the CP-QUMB potential of mean force of the hydrated electron relative to the air/water interface (as




Figure 4. Energetic properties of a hydrated electron as a function of distance from the air/water interface. Negative distances indicate aqueous solvation of the electron below the average height of the interface (Gibbs dividing surface), positive distances are when the electron center of mass is above the interface, outside the bulk. (a) Helmholtz free energy, i.e., the potential of mean force, calculated via CP-QUMB. Arrows indicate where snapshots were taken for Figure 1. (b) (left axis, solid blue curve) Electron's average quantum energy eigenvalue (negative vertical binding energy). (right axis, dashed red curve) Electron-solvent polarization energy. Note that unlike panel (a), the energy units of this panel are in eV.

defined by the GDS). For the chosen TB potential,^{8b} the hydrated electron is solvated preferentially in bulk water away from the interface. This is consistent with the nonequilibrium simulations by MRT, who using this same model found that interfacial solvated electrons moved into the bulk on a ~ 10 -ps time scale.^{27a} Furthermore, the PMF for dragging an electron above the water surface is strongly repulsive, such that "floating" surface states of the hydrated electron are highly unstable by $\sim 8 k_B T$ relative to the bulk. The MD snapshot shown in Figure 1a hints at the origin of this behavior: rather than simply floating on the surface of water, the hydrated electron in this environment prefers to drag a number of water molecules with it to above the GDS, which carries a large free energy cost due to the high surface tension of water.

For hydrated electron solvation below the air/water interface, interestingly the PMF shows the presence of a plateau extending from the GDS to ~ 4 Å below the surface. This suggests that rather than floating on the water surface, the TB model of an interfacial hydrated electron is solvated preferentially below the surface. Examining MD snapshots, such as the one shown in Figure 1b, reveals that in this plateau region the hydrated electron can acquire essentially a complete first solvation shell, similar to the findings of UMJ for their QM/MM model of the interfacial hydrated electron.^{27f} Furthermore, the presence of the plateau in the PMF, immediately below the water surface suggests that the interfacial electron is quasi-stable, so that it might be possible to trap the electron close to the interface if diffusion were sufficiently slowed. Indeed, MRT found that at a simulation temperature of $T = 200$ K, a hydrated electron introduced above a supercooled water surface persisted in the interfacial region for at least 180 ps.^{27a}

One of the large advantages of CP-QUMB is that it not only provides the PMF in Figure 4a but also allows for analysis of the driving forces to bulk solvation for this (TB) model of the hydrated electron. These driving forces can be revealed by calculating ensemble-averaged quantities as a function of distance from the interface, according to

$$A(z) = \frac{\langle \delta(q^z(\mathbf{R}^N) - z)A \rangle}{\langle \delta(q^z(\mathbf{R}^N) - z) \rangle} \quad (20)$$

where A is the observable of interest, and q^z is the reaction coordinate corresponding to the distance of the electron from the interface. We evaluated eq 20 by histogram-bin averaging the quantity A (e.g., the electron's eigenvalue) for each MD snapshot according to the electron's distance, z , from the interface. Ensemble-averaged quantities were then unbiased for each umbrella window according to eq 3. In Figure 4b we plot the hydrated electron's average quantum energy eigenvalue calculated this way as a function of distance from the air/water interface (solid blue curve). The quantum energy shows a clear energetic stabilization of bulk solvation relative to the interface ($z = 0$) of 0.77 eV = $30 k_B T$. The electron's energy eigenvalue can be related to the negative of its vertical binding energy (VBE). Our simulations therefore predict that the interfacial hydrated electron would have a binding energy of ~ 3 eV, in stark contrast to the controversial interpretations of experimental photoelectron spectra, which assigned a VBE of 1.6 eV to the interfacial electron.^{27b} The somewhat small shift in VBE from interface to bulk is consistent with the findings of UMJ, who found that their QM/MM model of the interfacial electron is essentially fully hydrated with bulklike properties. All of this means that if this model of the interfacial electron is correct, photoelectron experiments would be unable to distinguish surface-bound electrons from their bulk counterpart.^{27e-g}

We can further decompose the energetic contributions to bulk stabilization by specifically examining the electron-solvent polarization interaction (dashed red curve in Figure 4b), that is, the energy associated with polarization of the water molecules by the hydrated electron charge density. The data show clearly that solvent polarization is the dominant driving force for bulk solvation, contributing 0.43 eV = $17 k_B T$ of energy. Perhaps of greatest interest, from the bulk to just below the interfacial region at $z = -2$ Å, the electron's eigenvalue closely tracks the electron-solvent polarization energy (dashed red curve), indicating that in this region, polarization is the *only* part of

the electron–water interaction that drives the electron away from the interface.

Note that the electron–water polarization potential we used (which was chosen to match that in MRT's study) is rather crude in nature, comprising a simple pairwise additive potential of the form $U^{\text{pol}} = -f(r) \alpha / (2r^4)$, where f is a short-range damping function that removes the polarization catastrophe at the origin, and α is the isotropic polarizability of a water molecule, centered on the oxygen atom.^{8b} It is well-known that such pairwise additive polarization potentials tend to overestimate the overall polarization energy, particularly in water, due to the neglect of many-body dipole–induced dipole interactions that screen the polarization.³⁰ Within the pairwise additive approximation, it is perhaps not surprising that bulk solvation is favored over the interface since the polarization potential is purely attractive, favoring electron solvation in regions of high water density. Given the magnitude of the electron–solvent polarization energy, it is clear that a more careful treatment of this contribution is necessary before quantitative predictions of interface versus bulk stability can be made: for all the argument in the literature over the proper choice of electron–water pseudopotential,^{8e,31} for this problem of interfacial electrons, the oft-neglected many-body polarization term is clearly the most important.

Overall, our CP-QUMB simulations of the hydrated electron's PMF relative to the air/water interface provides valuable insight into the properties of interfacial hydrated electrons within the constraints of commonly used electron–water interactions. The development of CP-QUMB allows us to make qualitative comparisons between the interfacial properties of different models of the hydrated electron,^{8b,e,27a} a topic that we will address in the near future.

5. CONCLUSIONS

In this article, we introduced a new simulation method called Coupled-Perturbed Quantum Umbrella Sampling that extends the method of classical umbrella sampling to reaction coordinates involving the expectation values of quantum mechanical degrees of freedom. Using our method, one can restrain quantum degrees of freedom, for example, moving electrons into desired configurations to understand free energy barriers, differences, and nonequilibrium behavior. In this way, CP-QUMB not only allows for the construction of potentials of mean force for quantum coordinates, but also provides for chemistry by fiat, since it allows quantum systems to be placed into configurations where they might otherwise never go under equilibrium conditions. Finally, we were able to show that our CP-QUMB method is both more accurate and computationally cheaper than the previous Sum-Over-States formulation of Quantum Umbrella Sampling,² and we were able to demonstrate sampling of quantum reaction coordinates at tens of $k_B T$ in free energy.

We applied our new method to explore the potential of mean force for two condensed phase mixed quantum/classical systems of interest. First, we found that a neutral sodium atom in water is highly unstable and has a barrierless free energy profile to dissociation into an $e^-:\text{Na}^+$ contact pair and that the equilibrium structure of the pair was in good agreement with that previously determined by Spezia et al.^{7a} The features of the PMF we calculated also explain the experimental observation of spontaneous electron:cation contact pair formation following the charge-transfer-to-solvent reaction of atomic anions.^{18a,b,d}

In the second case, we used CP-QUMB to study hydrated electrons near the air/water interface and found that bulk solvation of the electron is thermodynamically favored over interfacial solvation, at least for the TB model of the hydrated electron.^{8b} This is consistent with previous calculations that TB hydrated electrons initiated at the interface diffuse into the bulk on a ~ 10 -ps time scale. Furthermore, we found that the main component of the electron–water interaction potential responsible for driving electrons into the bulk is the electron–solvent polarization interaction. Our results clearly show that any quantitative assessment of the bulk versus surface stability of the hydrated electron will require an improved description of electron–solvent polarization.

For both the $e^-:\text{Na}^+$ and interfacial electron systems, we found that the quantum position expectation value served as a useful reaction coordinate. This is because quantum fluctuations of the hydrated electron were largely unchanged upon dragging it away from a sodium cation or away from an air/water interface: the electron remained a localized ball of charge. In contrast, electron transfer (ET) reactions represent a system where quantum fluctuations are expected to be important.^{3b,32} In particular, at the ET transition state, the electron is delocalized over both donor and acceptor sites, giving rise to a substantial increase in quantum fluctuations relative to either the reactant or product species. Our CP-QUMB formalism could be directly applied to other quantum expectation values, such as the second moment, $\langle r^2 \rangle$, which is a measure of the quantum fluctuations. Thus, extending quantum umbrella sampling to multiple quantum expectation value coordinates could address the role of quantum fluctuations in a reaction coordinate, such as that during ET.

Finally, we note that while we have presented CP-QUMB in the framework of a grid-based one-electron electronic structure, there is no fundamental reason that it cannot be extended to many-electron *ab initio* electronic structure calculations. Indeed, analytical derivatives of the necessary matrix elements in CP-QUMB recently have been published at the configuration interaction with singles (CIS) level of theory for atom-centered Gaussian basis functions.³³ Thus, combining *ab initio* electronic structure with quantum umbrella sampling opens up the exciting possibility of performing chemistry by fiat.

■ APPENDIX

A.1. Matrix Elements Involving the Coulomb Potential

Both the SOS-QUMB (eq 6) and CP-QUMB (eq 7) methods require matrix elements of the derivative of the potential operator. In mixed/quantum classical schemes, the potential operator is typically represented by pseudopotentials that couple the quantum and classical particles,^{8b–e} which in a grid basis have derivatives that are readily evaluated:

$$\langle \psi^I | \nabla_{\mathbf{R}_\alpha} U | \psi^I \rangle = \sum_j (c_j^I)^* c_j^I \nabla_{\mathbf{R}_\alpha} U(\vec{r}_j) \quad (21)$$

In the SOS-QUMB method, the bra corresponds to an electronic eigenstate, while in CP-QUMB it corresponds to a Z -vector solution of eq 18; that is, $c_j^I = Z_j$. To handle electrostatic interactions, we employ the standard Ewald method and split the Coulomb operator part of the pseudopotential into a short-range (SR) term that decays rapidly in real space and a long-range (LR) term that converges rapidly in reciprocal space.^{8a,34} The SR term is best evaluated in real space according to eq 21, with

$$U^{\text{SR}}(\vec{r}_j) = -\sum_{\alpha} \frac{q_{\alpha} \operatorname{erfc}(\kappa|\vec{R}_{\alpha} - \vec{r}_j|)}{|\vec{R}_{\alpha} - \vec{r}_j|} \quad (22)$$

where q_{α} is the partial charge of classical site α , and κ is the Ewald splitting parameter that controls how fast the SR term decays in real space. The LR term, on the other hand, is best evaluated in reciprocal space:

$$\langle \psi^I | \nabla_{\mathbf{R}_{\alpha}} U^{\text{LR}} | \psi^I \rangle = -\frac{1}{V} \sum_{\vec{k} \neq 0} \frac{\vec{k}}{k^2} e^{-k^2/4\kappa^2} \operatorname{Im}[\tilde{\rho}_{\alpha}^{\text{Cl}}(\vec{k})(\tilde{\rho}^{\text{QJ}}(\vec{k}))^*] \quad (23)$$

where V is the cell volume, and k is a reciprocal lattice vector. The classical and quantum reciprocal-space charge densities are given by

$$\tilde{\rho}_{\alpha}^{\text{Cl}}(\vec{k}) = q_{\alpha} e^{i\vec{k} \cdot \vec{R}_{\alpha}} \quad (24)$$

$$\tilde{\rho}^{\text{QJ}}(\vec{k}) = -\sum_j (c_j^I)^* c_j^I e^{i\vec{k} \cdot \vec{r}_j} \quad (25)$$

Since the wavefunction coefficients, c , are stored on a regular real-space grid, the quantum reciprocal-space charge densities from eq 25 could be found by fast Fourier transform (FFT). However, in the systems of interest here, the quantum subsystem is localized to a region much smaller than the total system size. Furthermore, for reasonable values of κ , we found that the largest reciprocal lattice vector, k , required to converge eq 23 was smaller than the largest wavevector spanned by the quantum grid. Use of an FFT in eq 25 would therefore require a costly extrapolation of the wavefunctions to a grid commensurate with the full system size, even though many of the reciprocal wavefunction coefficients are subsequently discarded in eq 23. We therefore chose to evaluate eq 25 as a discrete Fourier transform, including only the reciprocal lattice vectors required to converge the LR force, the cost of which scales as $N_c N_k^{1/3}$, where N_c is the number of wave function coefficients, and N_k is the number of reciprocal lattice vectors. FFTs scale as $N_c \log(N_c)$, so their use may become favorable for large N_c and $N_k \approx N_c$.

AUTHOR INFORMATION

Corresponding Author

*E-mail: schwartz@chem.ucla.edu.

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

This work was supported by U.S. National Science Foundation Grant No. CHE-1212951. Computational resources were supplied by the UCLA Institute for Digital Research and Education and the California NanoSystems Institute.

REFERENCES

- (1) (a) Frenkel, D.; Smit, B. *Understanding Molecular Simulation*, 2nd ed.; Academic Press: San Diego, CA, 2002. (b) Kästner, J. *WIREs Comput. Mol. Sci.* **2011**, *1*, 932–942. (c) Abrams, C.; Bussi, G. *Entropy* **2013**, *16*, 163–199.
- (2) Borgis, D.; Staib, A. *J. Chem. Phys.* **1996**, *104*, 4776–4783.
- (3) (a) Sprik, M.; Impey, R. W.; Klein, M. L. *Phys. Rev. Lett.* **1986**, *56*, 2326–2329. (b) Bader, J. S.; Kuharski, R. A.; Chandler, D. *J. Chem. Phys.* **1990**, *93*, 230–236. (c) Major, D. T.; Gao, J. *J. Chem. Theory Comput.* **2007**, *3*, 949–960. (d) Collepardo-Guevara, R.; Craig, I. R.; Manolopoulos, D. E. *J. Chem. Phys.* **2008**, *128*, 144502. (e) Wong, K.

F.; Sonnenberg, J. L.; Paesani, F.; Yamamoto, T.; Vaníček, J.; Zhang, W.; Schlegel, H. B.; Case, D. A.; Cheatham, T. E.; Miller, W. H.; Voth, G. A. *J. Chem. Theory Comput.* **2010**, *6*, 2566–2580. (f) Menzeleev, A. R.; Ananth, N.; Miller, T. F. *J. Chem. Phys.* **2011**, *135*, 074106. (g) Pérez, A.; von Lilienfeld, O. A. *J. Chem. Theory Comput.* **2011**, *7*, 2358–2369.

(4) (a) Feynman, R. P.; Hibbs, A. R. *Quantum mechanics and path integrals*; McGraw-Hill: New York, 1965. (b) Ceperley, D. M. *Path Integral Monte Carlo Methods for Fermions*. In *Monte Carlo and molecular dynamics of condensed matter systems*, Proceedings of the Euroconference on Computer Simulation in Condensed Matter Physics and Chemistry, Como, Italy, 3–28 July, 1995; Binder, K., Ciccotti, G., Eds.; Italian Physical Society: Bologna, Italy, 1996.

(5) Stern, A. C.; Baer, M. D.; Mundy, C. J.; Tobias, D. J. *J. Chem. Phys.* **2013**, *138*, 114709.

(6) Torrie, G. M.; Valleau, J. P. *J. Comput. Phys.* **1977**, *23*, 187–199.

(7) (a) Spezia, R.; Nicolas, C.; Archirel, P.; Boutin, A. *J. Chem. Phys.* **2004**, *120*, 5261–5268. (b) Coudert, F.-X.; Archirel, P.; Boutin, A. *J. Phys. Chem. B* **2006**, *110*, 607–615.

(8) (a) Ewald, P. P. *Ann. Phys. (Berlin, Ger.)* **1921**, *369*, 253–287.

(b) Turi, L.; Borgis, D. *J. Chem. Phys.* **2002**, *117*, 6186–6195.

(c) Smallwood, C. J.; Larsen, R. E.; Glover, W. J.; Schwartz, B. J. *J. Chem. Phys.* **2006**, *125*, 074102. (d) Smallwood, C. J.; Mejia, C. N.; Glover, W. J.; Larsen, R. E.; Schwartz, B. J. *J. Chem. Phys.* **2006**, *125*, 074103. (e) Larsen, R. E.; Glover, W. J.; Schwartz, B. J. *Science* **2010**, *329*, 65–69. (f) Herbert, J. M.; Jacobson, L. D. *J. Phys. Chem. A* **2011**, *115*, 14470–14483.

(9) Osamura, Y.; Yamaguchi, Y.; Schaefer, H., III *Theoret. Chim. Acta* **1987**, *72*, 71–91.

(10) Handy, N. C.; Schaefer, H. F. *J. Chem. Phys.* **1984**, *81*, 5031–5033.

(11) Hestenes, M. R.; Stiefel, E. *J. Res. Natl. Bur. Stand.* **1952**, *49*, 409.

(12) Balay, S.; Adams, M. F.; Brown, J.; Brune, P.; Buschelman, K.; Eijkhout, V.; Gropp, W. D.; Kaushik, D.; Knepley, M. G.; McInnes, L. C.; Rupp, K.; Smith, B. F.; Zhang, H. *PETSc Users Manual*; Technical Report ANL-95/11, Revision 3.4, for the Argonne National Laboratory: Argonne, IL, 2013.

(13) For the case of a reaction coordinate involving one dimension of the position operator, such as a single Cartesian component or a radial coordinate, one need only solve a single coupled-perturbed equation for the coordinate of interest. However, for the purposes of benchmarking our method in Section 3.1, we evaluated the computational cost of solving the wave function response to all three components of the position operator.

(14) Chodera, J. D.; Swope, W. C.; Pitera, J. W.; Seok, C.; Dill, K. A. *J. Chem. Theory Comput.* **2006**, *3*, 26–41.

(15) (a) Uhlig, F.; Marsalek, O.; Jungwirth, P. *J. Phys. Chem. Lett.* **2012**, *3*, 3071–3075. (b) Casey, J. R.; Kahros, A.; Schwartz, B. J. *J. Phys. Chem. B* **2013**, *117*, 14173–14182.

(16) Chen, X.; Bradforth, S. E. *Annu. Rev. Phys. Chem.* **2008**, *59*, 203–231.

(17) (a) Barthel, E. R.; Martini, I. B.; Schwartz, B. J. *J. Chem. Phys.* **2000**, *112*, 9433–9444. (b) Smallwood, C. J.; Bosma, W. B.; Larsen, R. E.; Schwartz, B. J. *J. Chem. Phys.* **2003**, *119*, 11263–11277.

(18) (a) Glover, W. J.; Larsen, R. E.; Schwartz, B. J. *J. Chem. Phys.* **2008**, *129*, 164505. (b) Bragg, A. E.; Cavanagh, M. C.; Schwartz, B. J. *Science* **2008**, *321*, 1817–1822. (c) Glover, W. J.; Larsen, R. E.; Schwartz, B. J. *J. Phys. Chem. B* **2010**, *114*, 11535–11543. (d) Bragg, A. E.; Glover, W. J.; Schwartz, B. J. *Phys. Rev. Lett.* **2010**, *104*, 233005. (e) Glover, W. J.; Larsen, R. E.; Schwartz, B. J. *J. Phys. Chem. A* **2011**, *115*, 5887–5894.

(19) Toukan, K.; Rahman, A. *Phys. Rev. B* **1985**, *31*, 2643–2648.

(20) Aqvist, J. *J. Phys. Chem.* **1990**, *94*, 8021–8024.

(21) Davidson, E. R. *J. Comput. Phys.* **1975**, *17*, 87–94.

(22) Hernandez, V.; Roman, J. E.; Vidal, V. *ACM Trans. Math. Software* **2005**, *31*, 351–362.

(23) Bussi, G.; Donadio, D.; Parrinello, M. *J. Chem. Phys.* **2007**, *126*, 014101.

(24) Shirts, M. R.; Chodera, J. D. *J. Chem. Phys.* **2008**, *129*, 124105.

(25) Note that, while we followed closely the simulation protocols of ref 7a, some modifications to accommodate a Fourier grid had to be made. In particular, the sodium pseudopotential used in ref 7a contains a Coulomb singularity, which cannot be represented on a finite Fourier grid, so instead we used our previously developed sodium pseudopotential from ref 18a. Coudert et al. have demonstrated that the e^- - Na^+ PMF is somewhat sensitive to the details of this interaction potential;^{7b} in particular, after optimizing the sodium pseudopotential, they found a smaller equilibrium e^- - Na^+ separation of $r = 1.7$ Å and a higher barrier to dissociation.

(26) (a) Jungwirth, P.; Tobias, D. J. *J. Phys. Chem. B* **2002**, *106*, 6361–6373. (b) Otten, D. E.; Shaffer, P. R.; Geissler, P. L.; Saykally, R. J. *Proc. Natl. Acad. Sci. U.S.A.* **2012**, *109*, 701–705.

(27) (a) Madarász, Á.; Rosicky, P. J.; Turi, L. *J. Chem. Phys.* **2007**, *126*, 234707. (b) Siefermann, K. R.; Liu, Y.; Lugovoy, E.; Link, O.; Faubel, M.; Buck, U.; Winter, B.; Abel, B. *Nat. Chem.* **2010**, *2*, 274–279. (c) Tang, Y.; Shen, H.; Sekiguchi, K.; Kurahashi, N.; Mizuno, T.; Suzuki, Y.-I.; Suzuki, T. *Phys. Chem. Chem. Phys.* **2010**, *12*, 3653–3655. (d) Sagar, D. M.; Bain, C. D.; Verlet, J. R. R. *J. Am. Chem. Soc.* **2010**, *132*, 6917–6919. (e) Buchner, F.; Schultz, T.; Lubcke, A. *Phys. Chem. Chem. Phys.* **2012**, *14*, 5837–5842. (f) Uhlig, F.; Marsalek, O.; Jungwirth, P. *J. Phys. Chem. Lett.* **2013**, *4*, 338–343. (g) Elkins, M. H.; Williams, H. L.; Shreve, A. T.; Neumark, D. M. *Science* **2013**, *342*, 1496–1499.

(28) Fonseca, T. L.; Castro, M. A.; Cabral, B. J. C.; Coutinho, K.; Canuto, S. *Chem. Phys. Lett.* **2009**, *481*, 73–77.

(29) Rodriguez, J.; Laria, D. *J. Phys. Chem. B* **2004**, *109*, 6473–6478.

(30) (a) Stampfli, P. *Phys. Rep.* **1995**, *255*, 1–77. (b) Voora, V. K.; Ding, J.; Sommerfeld, T.; Jordan, K. D. *J. Phys. Chem. B* **2012**, *117*, 4365–4370.

(31) (a) Larsen, R. E.; Glover, W. J.; Schwartz, B. J. *Science* **2011**, *331*, 1387. (b) Jacobson, L. D.; Herbert, J. M. *Science* **2011**, *331*, 1387. (c) Turi, L.; Madarász, Á. *Science* **2011**, *331*, 1387.

(32) Kuharski, R. A.; Bader, J. S.; Chandler, D.; Sprik, M.; Klein, M. L.; Impey, R. W. *J. Chem. Phys.* **1988**, *89*, 3248–3257.

(33) Fatehi, S.; Alguire, E.; Subotnik, J. E. *J. Chem. Phys.* **2013**, *139*, 124112–124118.

(34) Allen, M. P.; Tildesley, D. J. *Computer Simulation of Liquids*; Oxford University Press: New York, 1989.