

Efficient Implementation of Constant pH Molecular Dynamics on Modern Graphics Processors

Evan J. Arthur¹ and Charles L. Brooks III²

Correspondence to: Charles L. Brooks III (E-mail: brookscl@umich.edu)

ABSTRACT

The treatment of pH sensitive ionization states for titratable residues in proteins is often omitted from molecular dynamics (MD) simulations. While static charge models can answer many questions regarding protein conformational equilibrium and protein-ligand interactions, pH-sensitive phenomena such as acid-activated chaperones and amyloidogenic protein aggregation are inaccessible to such models. Constant pH molecular dynamics (CPHMD) coupled with the Generalized Born with a Simple sWitching function (GBSW) implicit solvent model provide an accurate framework for simulating pH sensitive processes in biological systems. Although this combination has demonstrated success in predicting pK_a values of protein structures, and in exploring dynamics of ionizable side-chains, its speed has been an impediment to routine application. The recent availability of low-cost graphics processing unit (GPU) chipsets with thousands of processing cores, together with the implementation of the accurate GBSW implicit solvent model on those chipsets [E.J. Arthur and C.L. Brooks III, *J. Comp. Chem.* **37**:927, **2016**], provide an opportunity to improve the speed of CPHMD and ionization modeling greatly. Here we present a first implementation of GPU-enabled CPHMD within the CHARMM-OpenMM simulation package interface. Depending on the system size and non-bonded force cutoff parameters, we find speed increases of between one and three orders of magnitude. Additionally, the algorithm scales better with system size than the CPU-based algorithm, thus allowing for larger systems to be modeled in a cost effective manner. We anticipate that the improved performance of this methodology will open the door for broad-spread application of CPHMD in its modeling pH-mediated biological processes.

¹ Department of Chemistry, University of Michigan, 930 N. University Ave, Ann Arbor, MI 48109

² Department of Chemistry and Biophysics Program, University of Michigan, 930 N. University Ave, Ann Arbor, MI 48109

Introduction

Proteins typically maintain their native structure and optimal functionality under a narrow range of pH.¹⁻³ Consequently, many biological systems tightly control local solvent pH to tune the effectiveness of enzymes, or to promote a useful protein conformation.^{1,4,5} Mitochondrial ATP synthase utilizes a trans-membrane proton gradient to power its rotary catalysis mechanism,⁶⁻⁸ and the departure from a normal pH range is known to be a driving force in forming the amyloid fibrils associated with Alzheimer's disease.^{9,10} Additional examples of pH driven processes include the proton-activated gate mechanism of the KcsA potassium channel,¹¹ and the catalytic pathway of dihydrofolate reductase.¹² Finally, a notable survey by Aguilar et al. showed that about 60% of the protein-ligand complexes indicated that at least one titratable residue of the protein assumed a different protonation state between bound and unbound states.¹³ Although important to many biological processes, pH-dependence in bio-macromolecule simulations remains a nonstandard tool that awaits both wider acceptance, and finer tuning of its models.

Typical molecular dynamics (MD) simulations fix all amino acid protonation states to those of isolated residues in a neutral pH environment. While this pH-insensitive approach is sufficient to fold some proteins and observe their conformational equilibria,¹⁴ it arguably fails to capture phenomena dependent on local ionization effects of side-chains or perturbations to a residue's pK_a .^{15,16} This failure is particularly problematic for histidine residues, in that they have two hydrogen atoms that titrate with near-neutral pH. This ionizability indicates that in biologically-relevant pH environments histidine's protonation state and tautomeric configuration are often unclear.¹⁷

In recent decades a series of models of varying complexity and accuracy promise to bring accurate pH responsiveness to MD simulations.

Protonation-state modeling of amino acids in MD simulations is based on setting up a pH-sensitive extended Hamiltonian that modifies the force field parameters and structure of a given molecule. This began by discretely titrating protons, and progressing a simulation using instantaneous switches between protonated and unprotonated states. Mertz and Pettitt used an open system Hamiltonian to model the titration of acetic acid,¹⁸ and Sham et al. applied a linear response approximation through the protein-dipoles Langevin-dipoles model to calculate lysozyme residue pK_a values.¹⁹ Additional work has been done where Monte Carlo (MC) sampling guides the protonation state of an otherwise classical MD simulation. Baptista et al. used explicitly represented solvent molecules with an implicit solvent Poisson-Boltzmann (PB) function to determine protonation states.^{20,21} Meanwhile, Mongan et al. utilized generalized Born (GB) implicit solvation to add a solvation free energy component to the protonation function.²² While all these discrete models can predict pK_a values for individual amino acids to within one pK unit, they are computationally expensive. Whether the expense stems from the need to relax numerical instabilities caused by instantaneous protonation / deprotonation events, or from the MC algorithms' ability to titrate only one hydrogen at a time, such methods may require an unreasonable amount of time to study large systems with many titratable groups.

One possible solution to these issues with discrete titration methods is to use continuous titration of H^+ atoms. Brooks and co-workers developed one such method called constant pH molecular dynamics (CPHMD), which uses λ -dynamics coupled to transitions between protonation states.^{23,24} This method uses the Generalized Born implicit solvent model with a Simple sWitching function (GBSW) model,²⁵ or the related Generalized Born with Molecular Volume (GBMV) model,²⁴ to efficiently couple the protonation state to the solvation free energy of the molecule. Khandogin and Brooks then introduced proton tautomerism

capabilities to this method, which allows multi-site titrating residues, such as histidines, to be modeled more accurately.²⁶ Since the method is continuous, there are no instantaneous protonation/deprotonation events, and multiple residues can titrate simultaneously. Additionally, such continuous titration methods allow for the efficient coupling of protonation states among neighboring residues. The result is a pH simulation method that can calculate pK_a values of protein structures to within 0.75 pK units,¹⁶ and can resolve the dominant folding pathway of the pH-sensitive HdeA homodimer.¹⁵

CPHMD's efficiency, however, is bound by the rate-limited component of the calculation: the GBSW solvent model. As such, when running on a single-core central processing unit (CPU), CPHMD achieves on the order of 1 nanosecond (ns) of simulation time per day when simulating a solute system of about 1,000 atoms. Since typical uses of CPHMD, such as predicting pK_a shifts of protein residues, may require many nanoseconds of simulation time,¹⁶ even smaller proteins, such as lysozymes, may require about a week to converge on useful results. Larger systems, such as asymmetric viral capsid subunits with tens of thousands of atoms, may require unreasonably long simulation times if captured in full atomic detail.^{27,28} Fortunately, the GBSW solvent model has recently been refactored to function on new, parallel graphics processing unit (GPU) hardware, and is now between 1 and 2 orders of magnitude faster than its CPU counterpart.^{16,29} By incorporating the CPHMD model into the GPU-GBSW algorithm, there holds the promise of speeding up pH simulations substantially.

This study represents an increment in the ongoing adaptation of efficient and useful algorithms onto parallel-processing GPUs. Such chipsets can contain thousands of processing cores, and are able to process C-like languages

such as Open Computing Language (OpenCL) and Compute Unified Device Architecture (CUDA). This combination of features has opened up a new frontier of parallel processing where expensive computer clusters can be replaced with single, affordable graphics cards. Simulation packages such as CHARMM,³⁰ AMBER,³¹ OpenMM,³² GROMACS,³³ and NAMD³⁴ all offer GPU-accelerated options for many types of studies, and most of those options receive speed increases of greater than an order of magnitude over their CPU counterparts.

Due to OpenMM's effectiveness in harnessing the capabilities of GPUs with a wide variety of hardware, a CHARMM-OpenMM interface was developed to combine the strengths of both simulation packages.^{30,32} CHARMM's robust algorithms can be used to design and parameterize a simulation, and OpenMM's efficient GPU-based algorithms can be used to propagate dynamics.^{30,32} Now with the recent incorporation of the GBSW solvent model into the CHARMM-OpenMM interface, many of CHARMM's algorithms parameterized for use with GBSW, such as CPHMD, can be adapted for parallel processing on GPUs as well. In this study we take advantage of the recent incorporation of GBSW onto GPUs, and discuss the adaptation of CPHMD onto this new parallel architecture. First we explain the underlying theory behind λ -dynamics: how a λ coordinate is used to represent the titration state of a residue, and how that coordinate is propagated. Then we delve into how it was originally implemented for CHARMM, and examine fitting CPHMD into the GBSW algorithm. Here we discuss the algorithmic improvements, and show how many force contributions on λ are calculated alongside the free energy of solvation. Finally we present benchmarks achieved by the new algorithm, and comment on future directions for pH simulations.

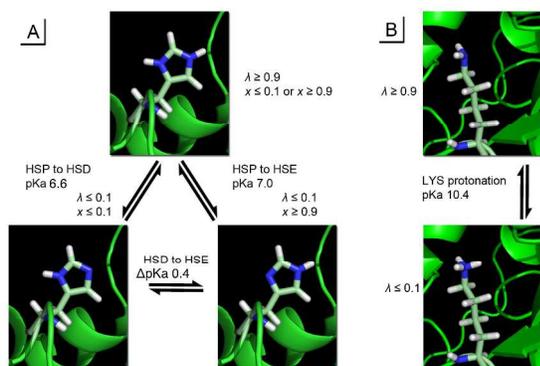


Figure 1. Shown are Cartoons of the protonated and unprotonated states of A) histidine and B) lysine. Also noted are the reference pK_a values of each transition when occurring in an isolated residue, as well as the λ values at each state.

Methods

The λ coordinates and their underlying energy function for single-site titration

For clarity in following discussions, we present the underlying theory of CPHMD. We start by setting up the framework for a single residue with one titrating hydrogen. The rudimentary picture of titration events is an equilibrium association/dissociation reaction of a model compound $A_{(aq)}$ in aqueous solution from a titrating proton.



Here, the protonation free energy is defined by

$$\Delta G^{\text{exp}}(\text{model}) = -k_B T \ln 10(pK_a^{\text{exp}} - pH) \quad (2)$$

where k_B is Boltzmann's constant, and T is the temperature. We can approximate the above equations through classical simulations by interpreting the protonation interaction as a change in free energies:

$$\begin{aligned} & \Delta G^{\text{exp}}(\text{protein}) - \Delta G^{\text{exp}}(\text{model}) \\ &= \Delta G^{\text{classical}}(\text{protein}) - \Delta G^{\text{classical}}(\text{model}) \end{aligned} \quad (3)$$

This relationship then leads to an estimate of experimental free energy of protonation for a single titrating site:

$$\begin{aligned} \Delta G^{\text{exp}}(\text{protein}) &= \Delta G^{\text{classical}}(\text{protein}) \\ &- \Delta G^{\text{classical}}(\text{model}) + \Delta G^{\text{exp}}(\text{model}) \end{aligned} \quad (4)$$

From this perspective, we infer that titratable groups have an intrinsic free energy of protonation that is perturbed by the protein environment mainly through non-bonded interactions. We model this perturbation by extending the system's Hamiltonian with a non-geometric dimension of λ . As mentioned in the introduction, the CPHMD model uses a series of λ coordinates where each λ value tracks the progress of protonation-deprotonation events at a single titration site. For a particular residue i , these coordinates are generated from

$$\lambda_i = \sin^2(\theta_i) \quad (5)$$

where i is the residue being titrated. In this form the θ variable is bound to all real numbers, and λ is bound to the continuous range $0 \leq \lambda_i \leq 1$. The sine-squared function then favors λ values near the boundary protonated (1) and unprotonated (0) states. Because λ is only physically relevant as it nears these boundary states, we impose cutoffs on interpreting λ . In CPHMD an unprotonated state is $\lambda_i \leq 0.1$, a protonated state is $\lambda_i \geq 0.9$, and a mixed state is $0.1 < \lambda_i < 0.9$. Figure 1 illustrates the protonation states and their corresponding λ values. Potentials and their derivative forces on λ are then interpreted as potentials and forces on θ .

The potential energy that governs protonation states contains five λ -dependent components. We start with the pH dependence of the deprotonation free energy as follows from ΔG^{exp} in Equation 4. This potential connects λ to the pK_a of a residue in its isolated, reference state:

$$U^{pH}(\lambda_i) = \lambda_i (pK_a(i) - pH) (k_B T \ln 10) \quad (6)$$

Here, $pK_a(i)$ is the pK_a of titrating group i . Next we have the potential of mean force (PMF) along the λ coordinate from ΔG^{model} . This term corresponds to the negative of free energy needed to deprotonate a model residue:

$$U^{model}(\lambda_i) = A_i (\lambda_i - B_i)^2 \quad (7)$$

Equation 7 is a quadratic fit to the thermodynamic work potential of deprotonating a model compound, and it splits the protonation state into two low-energy wells that represent the protonated and unprotonated states. Then a barrier potential is added that disfavors mixed states of λ :

$$U^{barrier}(\lambda_i) = 4\beta_i (\lambda_i - 1/2)^2 \quad (8)$$

The barrier scaling parameter β_i is an empirical coefficient designed to tune the propensity for a λ value to remain in either protonated or unprotonated states, while facilitating transitions between them. In the current iteration of CPHMD, β_i assumes a value of either 2.5 or 1.75 kcal/mol, depending on the residue. Finally, we arrive at the two charge-dependent potentials: the Coulombic and generalized Born. The classical Coulombic potential is

$$U^{elec}(\lambda_i) = \sum_{a,i} \sum_b K^{elec} \frac{q_{a,i}(\lambda_i) q_b}{r_{ab}} \quad (9)$$

Here K^{elec} is Coulomb's constant, q_a and q_b are the partial charges of atoms a and b respectively, and r_{ab} is the distance between those atoms. Note that this potential for residue i includes the interactions between all atoms a in residue i to all other atoms in the system. Meanwhile, $q_a(\lambda_i)$ is a λ -dependent charge of atom a , which follows the form

$$q_{a,i}(\lambda_i) = \lambda_i q_{a,i}^{unprot} + (1 - \lambda_i) q_{a,i}^{prot} \quad (10)$$

where charges on titrating atom a can be in protonated ($q_{a,i}^{prot}$) and unprotonated ($q_{a,i}^{unprot}$) states. We note that in an effective charge model of pH, titrating residues are allowed to interact. As such, any atom b from a titrating residue j interacting with residue i has its own $q_{b,j}^{prot}$ and $q_{b,j}^{unprot}$. Thus the partial charge q_b follows one of two possibilities:

$$q_b = \begin{cases} q_b & \text{non-titrating} \\ \lambda_j q_{b,j}^{unprot} + (1 - \lambda_j) q_{b,j}^{prot} & \text{titrating} \end{cases} \quad (11)$$

That is if atom b lies in a non-titrating residue, that atom's partial charge is simply the standard partial charge from that residue's force field. If atom b lies in a titrating residue j and its charge is affected by the protonation state of j , then its partial charge is derived from the same λ -dependent relationship from Equation 10. Since atoms near a titrating site can have their partial charges affected by titration states, many more than the titrating hydrogen atoms can possess a λ -dependent charge state. We also note that at times $j = i$ if we observe the Coulombic interaction between two atoms on the same titrating residue. The final λ -dependent potential is that from the GB solvent model as expressed in the Still equation:³⁵

$$U^{GB}(\lambda_i) = \sum_{a,i} \sum_b \tau \frac{q_{a,i}(\lambda_i) q_b}{f_{ab}^{GB}} \quad (12)$$

where

$$f_{ab}^{GB} = \left[r_{ab}^2 + R_a^{Born} R_b^{Born} \exp\left(-r_{ab}^2 / \left(4R_a^{Born} R_b^{Born}\right)\right) \right]^{1/2} \quad (13)$$

Here, $q_a(\lambda_i)$ and q_b follow the same form as in Eqn 10 and 11 respectively; r_{ab} is the distance between atoms a and b ; τ is the factor that

scales the Born energy by the difference in dielectric values at the dielectric boundary and by any contributing salt effects,³⁶ and the values R_a^{Born} and R_b^{Born} represent the Born radii of atoms a and b respectively. The Born radii are the effective distance between an atom and the solute-solvent dielectric boundary, and they are calculated through volumetric integration following the GBSW implicit solvent model.²⁵

If we pull together the complete potential for a titrating residue i from Eqns 6 through 13, then we arrive at the form

$$U_i^{total}(\lambda_i) = U_i^{pH}(\lambda_i) + U_i^{model}(\lambda_i) + U_i^{barrier}(\lambda_i) + U_i^{elec}(\lambda_i) + U_i^{GB}(\lambda_i) + U_i^{VDW} + U_i^{internal} \quad (14)$$

The so-called “internal energy” term ($U^{internal}$) corresponds to the bond, angle, and torsional energy terms of a classical energy force field. In this model, the titration state is dynamically independent of this potential. Although several models of CPHMD include a λ -dependent van der Waals term (U^{VDW}),^{26,37,38} during this study it was found that at most it contributes a minimal amount to a given residue’s force on λ , while it nearly doubles the calculation time of CPHMD. This term is negligible compared to the force on λ from other effects, and omitting it from the calculation showed no effect on the accuracy of CPHMD. Thus, in the interest of speeding up the original algorithm, the λ -dependent potential U^{VDW} was ignored in this implementation of CPHMD.

Although we now have the proper setup for addressing residues with a single titration site, such as in lysine, we need to address how CPHMD handles tautomerization in residues such as in aspartic acid and histidine.

Proton tautomerism

Similarly to how one λ variable is used to track the progress of titration states of a residue,

Khandogin and Brooks incorporated tautomeric behavior into CPHMD by providing residues with a second λ variable, called x , to track the progress of tautomeric states.²⁶ This arrangement is illustrated in Figure 1a with histidine. Just as in λ dynamics for titration states, transitions between tautomeric states are linearly interpolated using the x variable. What results are potentials that become bivariate in λ and x , and each tautomeric residue has four charge states: tautomer A in protonated and unprotonated states, and tautomer B in protonated and unprotonated states. What we shall see later is that residues can have equivalent states in this setup. Histidine’s protonated state, for example, is a residue saturated with protons. As such tautomers A and B of the protonated state are equivalent. We now review the influence of including two λ parameters for a tautomeric titrating residue.

The pH dependent potential becomes

$$U^{pH}(\lambda_i, x_i) = \lambda_i [x_i (pK_a^A(i) - pH) + (1 - x_i) (pK_a^B(i) - pH)] (k_B T \ln 10) \quad (15)$$

where the pK_a values of tautomers A and B are pK_a^A and pK_a^B respectively. While these pK_a values for aspartic acid and glutamic acid are equivalent and only serve as a sampling expedient,²⁶ in residues with asymmetric titrating sites such as histidine they are not. The PMF for protonation becomes a bivariate polynomial from Equation 7, which then expands into the general form

$$U^{model}(\lambda_i, x_i) = a_0 \lambda_i^2 x_i^2 + a_1 \lambda_i^2 x_i + a_2 \lambda_i x_i^2 + a_3 \lambda_i x_i + a_4 \lambda_i^2 + a_5 x_i^2 + a_6 \lambda_i + a_7 x_i + a_8 \quad (16)$$

The barrier potential is simply a summation of terms that disfavor the mixed states of both λ and x , and follows the form

$$U^{\text{barrier}}(\lambda_i, x_i) = 4\beta_i^\lambda (\lambda_i - 1/2)^2 + 4\beta_i^x (x_i - 1/2)^2 \quad (17)$$

Note that there are two barrier scaling parameters β_i^λ and β_i^x for λ and x . Although different biases for tautomeric and protonation transitions are possible in this equation, in the discussed CPHMD model they are identical for all titrating residues.

The charge-dependent potentials in Eqns 9 and 12 are only modified in that charges for atoms can now be dependent on the new x coordinate. The Coulombic and generalized Born potentials then follow the forms

$$U^{\text{elec}}(\lambda_i, x_i) = \sum_{a,i} \sum_b K^{\text{elec}} \frac{q_{a,i}(\lambda_i, x_i) q_b}{r_{ab}} \quad (18)$$

and

$$U^{\text{GB}}(\lambda_i, x_i) = \sum_{a,i} \sum_b \tau \frac{q_{a,i}(\lambda_i, x_i) q_b}{f_{ab}^{\text{GB}}} \quad (19)$$

respectively. The bivariate charge $q_{a,i}(\lambda_i, x_i)$ then follows the form

$$q_{a,i}(\lambda_i, x_i) = \lambda_i \left[x_i q_{a,i}^{\text{A,unprot}} + (1-x_i) q_{a,i}^{\text{B,unprot}} \right] + (1-\lambda_i) \left[x_i q_{a,i}^{\text{A,prot}} + (1-x_i) q_{a,i}^{\text{B,prot}} \right] \quad (20)$$

Where charges on titrating atom a are derived from the protonated and unprotonated variants of both A and B tautomers, $q_{a,i}^{\text{A,prot}}$, $q_{a,i}^{\text{A,unprot}}$, $q_{a,i}^{\text{B,prot}}$, and $q_{a,i}^{\text{B,unprot}}$. Similarly, the charge on atom b emerges as

$$q_b = \begin{cases} q_b & \text{non-titrating} \\ \lambda_j \left[x_j q_{b,j}^{\text{A,unprot}} + (1-x_j) q_{b,j}^{\text{B,unprot}} \right] + (1-\lambda_j) \left[x_j q_{b,j}^{\text{A,prot}} + (1-x_j) q_{b,j}^{\text{B,prot}} \right] & \text{titrating} \end{cases} \quad (21)$$

We now arrive at a general-purpose setup for evaluating the underlying potential for continuous transitions among various charge states of a particular residue. Deriving the forces with respect to λ and x , while important, serves little purpose for illuminating the topics explored in the remainder of this study. With the framework above, we now can discuss the construction of the original algorithm, and the changes made to refactor it for efficient parallel processing on GPUs.

Refactoring CPHMD

The original CPHMD model was built with mathematical precision and function portability in mind. It is a stand-alone module that can be applied to both implicit and explicit solvent systems, and except for atom coordinate and Born radii updates, it receives no input from other functions during a simulation. During the course of a timestep, each titrating coordinate λ_i is scanned to identify the residue type (such as whether the residue has one or two titrating hydrogens), and then an appropriate functional is applied to calculate its pH (Eqn 6 and 15), model (Eqn 7 and 16), and barrier (Eqn 8 and 17) potentials. Next, neighboring atom-atom interactions are scanned for whether one or both atoms reside in titrating groups. If a titrating atom-atom pair is found, then contributions to the electrostatic (Eqn 9 and 18) and GB (Eqn 12 and 19) potentials are integrated. Neighboring atom-atom pairs are then scanned again to calculate the VDW potential (ignored in this new iteration of CPHMD). Finally, the force on θ is calculated, and λ via θ is advanced a timestep using Langevin dynamics.³⁹ In this setup there are several opportunities presented to us for improving the algorithm both in the efficiency of its execution in parallel, and by weaving portions of the calculation into existent functions elsewhere in the simulation.

We first note that the majority of clock cycles used for calculating λ dynamics are spent on neighboring atom-atom interactions when

accumulating the electrostatic and GB potentials. While the calculations required for each atom pair are computationally cheap, the large number of interatomic interactions in a protein containing thousands of atoms can make this multitude of cheap calculations altogether expensive. As shown in Figure 2a, about 12% of a 2000-atom simulation is spent only on this calculation.

Both CPHMD and the GBSW solvent model require calculating the Still equation (Eqn 12 and 13) to address part of the neighboring atom potential, so a significant speed improvement can be made by placing all of CPHMD's atom-atom processes inside the neighboring atom process of the GBSW solvent model. This way, as GBSW produces the solute molecule's electrostatic solvation free energy and its derivative force on atoms, CPHMD processes neighboring atom potentials on λ simultaneously. Thus the large number of redundant atom-atom distance calculations can be reduced significantly during a simulation. This setup gains additional speedup through GBSW by using OpenMM's efficient parallel processing of neighboring-atom interactions. As shown in Figure 2, by combining the CPHMD and GBSW algorithms we see that pH modeling with CPHMD accounts for a much smaller fraction of the overall simulation time.

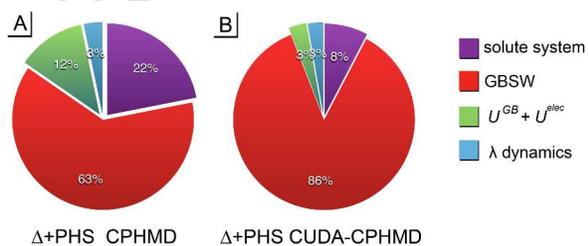


Figure 2. The approximate distributions of CPU time spent on running simulation components of Δ +PHS staphylococcal nuclease molecule. This protein contains 2132 atoms and 37 titrating residues. A) run using the original algorithm on a single processing core in CHARMM. B) run using the newly refactored CUDA-CPHMD algorithm.

Due to the nature of parallel processing, bottlenecks are often created from the longest

portions of non-parallel code. While a single-core process can be sped up dramatically by creating a case-by-case set of calculations, navigating through the additional overhead to make the situation-specific decision can slow parallel processes down. Regarding the equations described earlier, a titrating residue with one tautomer requires fewer calculations than a titrating residue with two. As we place each residue's force calculations in parallel processes, however, the speed of the code is improved by regarding all titrating residues as possessing two tautomeric states. In this new implementation of CPHMD, single-titration residues, such as lysine, are given extraneous x coordinates. Lysine then uses the barrier potential from Eqn 16, where the x -coupled coefficients a_0 , a_1 , a_2 , a_3 , and a_5 are set to a value of 0.0. Without the overhead for residue identification, the longest calculation required, that is calculating the force on θ for a residue with two tautomeric states, is shortened. What results is a speed improvement when calculating all components of the total potential on λ coordinates. As shown in Figure 2b, using the parallel CUDA-CPHMD algorithm for a small system impacts the processing time by approximately 6%, as opposed to 15% for the original algorithm.

Benchmarking CUDA-CPHMD

We finally reach an efficient setup where using the CPHMD model results in little slowdown of the overall simulation time. We chose several systems to benchmark the new algorithm, and explore the speed benefits it offers. We chose the naja atra snake cardiotoxin (PDB: 1CVO),⁴⁰ the Δ +PHS hyperstable variant of staphylococcal nuclease (PDB: 3BDC),⁴¹ and the asymmetric subunit of the bacteriophage HK97 head capsid (PDB: 2FT1).⁴² This trio provided a range of system sizes and residue configurations. To add additional statistics, the 7 proteins of the HK97 capsid were assembled into 6 additional subsystems, all of which appear in Figure 3 to show for a range of system sizes the speed dependence on system size. All simulations

were using the CHARMM22 force field^{43,44} using the Langevin integrator⁴⁵ with a timestep of 2 femtoseconds. These were NT (constant particle number and temperature) simulations at 298K in unbounded volumes using the CUDA-GBSW solvent model, and CUDA-CPHMD to model titration states and advance λ coordinates. Atomic radii for the GBSW solvent model were provided through work by Chen et al.⁴⁶ The hardware specifications of the computer used appear in Table S11 of the Supporting Information. We found speed improvements of between 1 and 3 orders of magnitude in the CUDA-CPHMD algorithm over its CPU counterpart.

As we combine the improved efficiency and parallel execution of both GBSW and CPHMD (shown in Figure 3a to 3d), substantial speed gains are found in this new version of pH modeling over its predecessor. For smaller

1,000-atom systems, we see a speed improvement of over 20-fold when comparing a 12-threaded CPHMD simulation to the new CUDA-CPHMD, and an improvement of over 150-fold when compared to the single-core algorithm (shown in Figure 3a). For larger 29,000 atom systems, we see a speed improvement of over 1,000-fold (shown in Figure 3c). Since the neighboring-atom component does not scale linearly with system size, larger systems experience a greater calculation time penalty than smaller ones. Fortunately, simple changes, such as using non-bonded cutoffs, can mitigate such problems. For instance, a non-bonded cutoff of 14 Å sped up the large viral capsid simulation to 6.7 ns/day (a 270% speed increase versus the no cutoff case).

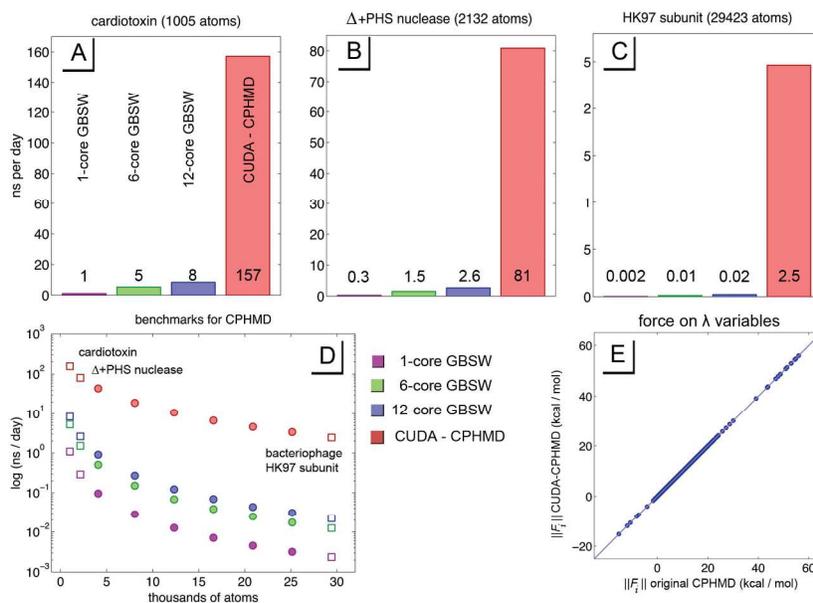


Figure 3. The benchmarks for the new CUDA-CPHMD algorithm. The individual systems tested were A) the naja atra snake cardiotoxin (PDB: 1CVO); B) the Δ +PHS hyperstable variant of staphylococcal nuclease (PDB: 3BDC); and C) the asymmetric subunit of the bacteriophage HK97 head capsule (PDB: 2FT1). As shown, the new algorithm is substantially faster than the original CPU algorithm by up to 3 orders of magnitude. In D) the same benchmarks from earlier are shown (squares) alongside subsystems from the 7 proteins of the bacteriophage subunit (circles). Notice that the CUDA algorithm scales more linearly with system size than its CPU-based counterpart. E) compares the force on λ as calculated on all 595 λ coordinates from both CPHMD algorithms. There is less than a 0.00017 (kcal/mol) AUE (average unsigned error) between the two algorithms.

Accuracy of the new CUDA-CPHMD algorithm

Single-residue Systems

Speed gains in implementing CPHMD are an important goal both for increasing the algorithm's applicability to a wider range of system sizes, and for its ability to converge on useful results more rapidly. Its accuracy, however, must not be compromised as we reconfigure the execution of the algorithm. In Figure 3e we show that there is little difference between the original CPHMD and CUDA-CPHMD algorithms when calculating the force on λ . We maintain an average unsigned error (AUE) of less than 0.00017 kcal/mol in this force. We also note that 99.9% of the AUE between the two CPHMD methods is from the slight differences in Born radii calculated from the original and CUDA implementations of GBSW. Thus, we conclude that CUDA-CPHMD accurately reproduces the original algorithm's force on λ .

While CUDA-CPHMD may be able to produce the force on λ coordinates, we ran additional tests to see whether or not residue protonation states are also reproduced. Due to each residue's pH-dependent biasing potential, a single residue alone in solution presumably should find an optimal protonation state depending on the environmental pH. At pH environments below a residue's pK_a the residue should favor a protonated state ($\lambda_i \leq 0.1$), and conversely a residue exposed to a pH above its pK_a should favor an unprotonated state ($\lambda_i \geq 0.9$). By calculating the fraction of protonated to unprotonated states of residues at various pH values and fitting the results to the Henderson-Hasselbalch equation of states, we expect the point of inflection to reproduce the pK_a of that residue.

We ran simulations of aspartic acid, glutamic acid, histidine, and lysine to calculate their protonation states, as shown in Figure 4. These residues were simulated using the same setup from the benchmarking section as NT simulations in an unbound volume, and CUDA-CPHMD was used both to model titration states and advance λ coordinates. The backbone atom ends were capped with the ACE and CT2 patches in CHARMM. Each dot in Figure 4 represents the average residue titration state from 200 ps of simulation time, and the residues ran at an average speed of 690 ns/day.

We find that without optimizing the simulations for speed, accuracy, or convergence of protonation states, that the pK_a values could be captured to within 0.5 pK units. Interestingly, all states reported a small, systematic overestimation of the pK_a , and the exact source of this discrepancy remains unclear. The CUDA-GBSW solvent model overestimates solvation energy by an average of approximately 0.16 kcal/mol. However, this overestimation of energy should bias deprotonation events to occur slightly more often, and thus lower the calculated pK_a . What is clear from these data, though, is that like its predecessor, the CUDA-CPHMD algorithm models the pH dependence of titration well for single residue systems. Next we explore multi-residue titration and the influence of protein conformation on pK_a values.

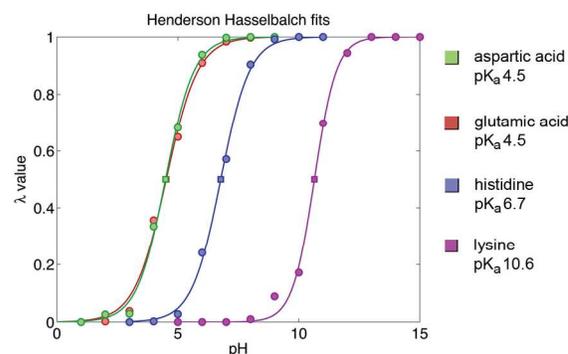


Figure 4. The pK_a calculations for 4 single residues: aspartic acid, glutamic acid, histidine, and lysine. The protonation state (dots) was calculated from the fraction of λ values in pure unprotonated and protonated states. The point of

inflection (boxes) of Henderson Hasselbalch equation fits (lines) indicates the calculated pK_a values. Even without optimizing for efficiency, convergence of data, or simulation parameters, we find the calculated pK_a values match those from the force field to within 0.5 pK units.

Multiple-residue Systems

The end purpose for CPHMD is to enable the study of complex pH-coupled phenomena of biological systems, such as pH-dependent protein conformation and cooperative titration effects among neighboring residues. As such, we test the accuracy of the CUDA-CPHMD algorithm by its ability to recapitulate residue pK_a values from both experiments and previous replica exchange studies, as shown in Figure 5. We study 9 model protein systems here: barnase^{2,47,48} (PDB code 1A2P); the serine protease inhibitor CI-2 from barley seeds^{47,49} (PDB code 2CI2); the hyperstable variant of staphylococcal nuclease, Δ +PHS^{3,16} (PDB code 3BDC); hen egg white lysozyme^{47,50,51} (PDB code 1LSA); the N-terminal domain of ribosomal protein L9^{47,52} (PDB code 1CQU); turkey ovomucoid^{47,53,54} (PDB code 1OMU); ribonuclease A^{47,55} (PDB code 7RSA); ribonuclease H from *Escherichia coli*^{47,56,57} (PDB code 2RN2); and *Bacillus circulans*

xylanase^{47,58,59} (PDB code 1BCX). Each protein was simulated in 11 pH windows from the pH -1 to the pH 9. Within each window, the proteins were simulated for 80 ps in 10 independent trajectories, which resulted in a total of 4.4 ns of simulation time per structure. All titrating residues were allowed to change protonation state using the new CUDA-CPHMD algorithm (Figure 5a, 5d, and 5e) and the original CPHMD algorithm (Figure 5b and 5d); and salt concentrations were added using concentrations that corresponded to the experiments.^{16,47} The simulations were run using the CHARMM22 force field^{43,44} with the Langevin integrator⁴⁵ with an integration timestep of 2 fs. These were NT (constant particle number and temperature) simulations each in an unbounded volume at a temperature of 298K using a Langevin heat bath. Atomic radii were optimized through work by Chen et al.⁴⁶ Similarly to the single-residue simulations, pK_a values were calculated by fitting the Henderson-Hasselbalch equation of states to the average protonation state λ of each titrating residue. Again, the point of inflection of the fit corresponds to the pK_a value of that residue. We report these values in Figure 5.

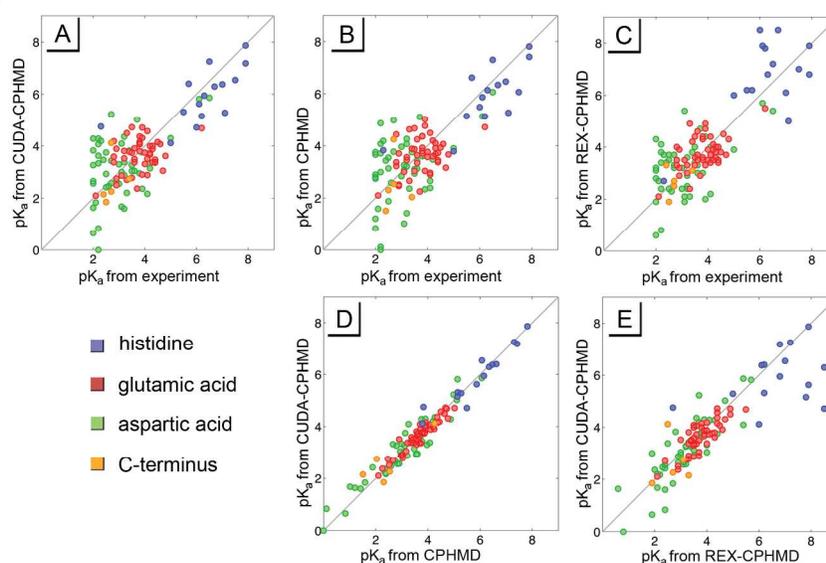


Figure 5. The pK_a calculations for all histidine (blue), glutamic acid (red), aspartic acid (green), and titrating C-terminus (orange) residues in all 9 of the test proteins. Each dot corresponds to a pK_a value resulting from fitting the Henderson Hasselbalch

equation to the fraction of λ values in pure unprotonated states. We present comparisons between pK_a values from A) CUDA-CPHMD and experiment; B) CPHMD and experiment; C) CPHMD with replica exchange (REX-CPHMD) and experiment; D) CUDA-CPHMD and CPHMD; and E) CUDA-CPHMD and REX-CPHMD. Even without optimizing the simulations to accommodate various titration equilibria for each protein, the CUDA-CPHMD algorithm successfully recapitulates experimental pK_a values to within 0.79 pK units of AUE. The experimental and REX-CPHMD results are from reference⁴⁷ and papers cited therein.

The AUE for all residues using CUDA-CPHMD was 0.79 pK units, which compares favorably to the AUE of 0.97 pK units using the null approximation (all pK_a values correspond to their reference values). This was the same 0.79 pK units of AUE that the original algorithm achieved, which further supports CUDA-CPHMD's accurately representing its CPU counterpart. Interestingly, while the average accuracy of CUDA-CPHMD and CPHMD were less than the 0.75 pK units of AUE achieved using the replica exchange methods from earlier studies, the non-replica-exchange pK_a calculations had a smaller standard deviation of error and fewer outlying predictions.^{16,47} Additional accuracy should be possible by coupling CUDA-CPHMD with the enhanced sampling of replica exchange in temperature or pH.^{16,47} This result holds great promise in establishing dynamic titration as a common feature of protein simulations.

Conclusions

In this study we present a significantly faster version of the CPHMD algorithm adapted for parallel processing in the CHARMM-OpenMM interface. While algorithmically the new CUDA-CPHMD algorithm represents little change over its predecessor, the speed improvements are so great that previously-unreasonable simulations are now straightforward to perform. For instance, what may have been a year-long simulation of the HK97 head capsule can now be performed in about 160 minutes. With this newfound speed is an opportunity to fine-tune the CPHMD titration model for a variety of protein systems, and to explore the impact of pH environments on side-chain dynamics both at the microsecond timescale and with all-atom detail.

Similarly to GBSW, the CPHMD model carries with it over a decade of research and parameterization.^{26,47,52} One model of particular interest is pH replica exchange (REX),⁶⁰ which has been shown to predict pK_a values of protein structures within single nanoseconds of simulation time.^{60,61} Coupled with the improved speed of CPHMD, adapting REX will enable a useful and rapid method for characterizing the chemical environment of protein interiors.

Acknowledgements

The authors gratefully acknowledge support from the NIH through grants GM103695 and GM107233 to CLB.

Keywords: implicit solvation, solvation, solvent model, CUDA, parallelization

References

1. Nielsen, J. E.; McCammon, J. A. *Protein Science : A Publication of the Protein Society* 2003, 12(9), 1894-1901.
2. Sali, D.; Bycroft, M.; Fersht, A. R. *Nature* 1988, 335(6192), 740-743.
3. Cannon, B.; Isom, D.; Robinson, A.; Seedorff, J.; Garcia-Moreno, B. *Biophysical Journal* 2007, 403A-403A.
4. Rabbani, G.; Ahmad, E.; Zaidi, N.; Fatima, S.; Khan, R. *Cell Biochemistry and Biophysics* 2012, 62(3), 487-499.
5. Wagner, G. R.; Payne, R. M. *Journal of Biological Chemistry* 2013, 288(40), 29036-29045.
6. Baker, L. A.; Watt, I. N.; Runswick, M. J.; Walker, J. E.; Rubinstein, J. L. *Proceedings of the National Academy of Sciences* 2012, 109(29), 11675-11680.
7. Cain, B. D.; Simoni, R. D. *Journal of Biological Chemistry* 1986, 261(22), 10043-10050.
8. Rastogi, V. K.; Girvin, M. E. *Nature* 1999, 402(6759), 263-268.
9. Clippingdale, A. B.; Wade, J. D.; Barrow, C. J. *Journal of Peptide Science* 2001, 7(5), 227-249.

10. Dobson, C. M. *Nature* 2003, 426(6968), 884-890.
11. Cuello, L. G.; Cortes, D. M.; Jogini, V.; Somporsut, A.; Perozo, E. *FEBS letters* 2010, 584(6), 1126-1132.
12. Howell, E. E.; Villafranca, J. E.; Warren, M. S.; Oatley, S. J.; Kraut, J. *Science* 1986, 231(4742), 1123-1128.
13. Aguilar, B.; Anandakrishnan, R.; Ruscio, J. Z.; Onufriev, A. V. *Biophysical journal* 2010, 98(5), 872-880.
14. Lindorff-Larsen, K.; Piana, S.; Dror, R. O.; Shaw, D. E. *Science* 2011, 334(6055), 517-520.
15. Ahlstrom, L. S.; Law, S. M.; Dickson, A.; Brooks III, C. L. *Journal of Molecular Biology* 2015, 427(8), 1670-1680.
16. Arthur, E. J.; Yesselman, J. D.; Brooks III, C. L. *Proteins: Structure, Function, and Bioinformatics* 2011, 79(12), 3276-3286.
17. Bashford, D.; Case, D. A.; Dalvit, C.; Tennant, L.; Wright, P. E. *Biochemistry* 1993, 32(31), 8045-8056.
18. Mertz, J. E.; Pettitt, B. M. *International Journal of Supercomputer Applications and High Performance Computing* 1994, 8(1), 47-53.
19. Sham, Y. Y.; Chu, Z. T.; Warshel, A. *Journal of Physical Chemistry B* 1997, 101(22), 4458-4472.
20. Baptista, A. M.; Martel, P. J.; Petersen, S. B. *Proteins-Structure Function and Genetics* 1997, 27(4), 523-544.
21. Baptista, A. M.; Teixeira, V. H.; Soares, C. M. *Journal of Chemical Physics* 2002, 117(9), 4184-4200.
22. Mongan, J. T.; Case, D. A.; McCammon, J. A. *Abstracts of Papers American Chemical Society* 2005, 229(Part 1), U768.
23. Kong, X. J.; Brooks III, C. L. *Journal of Chemical Physics* 1996, 105(6), 2414-2423.
24. Lee, M. S.; Salsbury, F. R.; Brooks III, C. L. *The Journal of Chemical Physics* 2002, 116(24), 10606-10614.
25. Im, W.; Lee, M. S.; Brooks III, C. L. *Journal of computational chemistry* 2003, 24(14), 1691-1702.
26. Khandogin, J.; Brooks, C. L., III. *Biophysical Journal* 2005, 89(1), 141-157.
27. Zeng, X.; Mukhopadhyay, S.; Brooks, C. L., III. *Proceedings of the National Academy of Sciences* 2015, 112(7), 2034-2039.
28. May, E. R.; Arora, K.; Brooks, C. L. I. *Journal of the American Chemical Society* 2014, 136(8), 3097-3107.
29. Arthur, E. J.; Brooks, C. L., III. *Journal of computational chemistry* 2016, 37(10), 927-939.
30. Brooks, B. R.; Brooks III, C. L.; Mackerell, A. D.; Nilsson, L.; Petrella, R. J.; Roux, B.; Won, Y.; Archontis, G.; Bartels, C.; Boresch, S.; Caffisch, A.; Caves, L.; Cui, Q.; Dinner, A. R.; Feig, M.; Fischer, S.; Gao, J.; Hodoscek, M.; Im, W.; Kuczera, K.; Lazaridis, T.; Ma, J.; Ovchinnikov, V.; Paci, E.; Pastor, R. W.; Post, C. B.; Pu, J. Z.; Schaefer, M.; Tidor, B.; Venable, R. M.; Woodcock, H. L.; Wu, X.; Yang, W.; York, D. M.; Karplus, M. *Journal of computational chemistry* 2009, 30(10), 1545-1614.
31. Case, D. A.; Cheatham, T. E.; Darden, T.; Gohlke, H.; Luo, R.; Merz, K. M.; Onufriev, A.; Simmerling, C.; Wang, B.; Woods, R. J. *Journal of computational chemistry* 2005, 26(16), 1668-1688.
32. Eastman, P.; Friedrichs, M. S.; Chodera, J. D.; Radmer, R. J.; Bruns, C. M.; Ku, J. P.; Beauchamp, K. A.; Lane, T. J.; Wang, L.-P.; Shukla, D.; Tye, T.; Houston, M.; Stich, T.; Klein, C.; Shirts, M. R.; Pande, V. S. *Journal of Chemical Theory and Computation* 2013, 9(1), 461-469.
33. Hess, B.; Kutzner, C.; van der Spoel, D.; Lindahl, E. *Journal of Chemical Theory and Computation* 2008, 4(3), 435-447.
34. Phillips, J. C.; Braun, R.; Wang, W.; Gumbart, J.; Tajkhorshid, E.; Villa, E.; Chipot, C.; Skeel, R. D.; Kalé, L.; Schulten, K. *Journal of computational chemistry* 2005, 26(16), 1781-1802.
35. Still, W. C.; Tempczyk, A.; Hawley, R. C.; Hendrickson, T. *Journal of the American Chemical Society* 1990, 112(16), 6127-6129.
36. Srinivasan, J.; Trevathan, M. W.; Beroza, P.; Case, D. A. *Theoretical Chemistry Accounts* 1999, 101(6), 426-434.
37. Donnini, S.; Tegeler, F.; Groenhof, G.; Grubmüller, H. *Journal of Chemical Theory and Computation* 2011, 7(6), 1962-1978.
38. Lee, M. S.; Salsbury, F. R.; Brooks III, C. L. *Proteins-Structure Function and Bioinformatics* 2004, 56(4), 738-752.
39. Uhlenbeck, G. E.; Ornstein, L. S. *Physical Review* 1930, 36(5), 0823-0841.
40. Singhal, A. K.; Chien, K. Y.; Wu, W. G.; Rule, G. S. *Biochemistry* 1993, 32(31), 8036-8044.
41. Castañeda, C. A.; Fitch, C. A.; Majumdar, A.; Khangulov, V.; Schlessman, J. L.; García-Moreno, B. E. *Proteins: Structure, Function, and Bioinformatics* 2009, 77(3), 570-588.

42. Gan, L.; Speir, J. A.; Conway, J. F.; Lander, G.; Cheng, N.; Firek, B. A.; Hendrix, R. W.; Duda, R. L.; Liljas, L.; Johnson, J. E. *Structure* 2006, 14(11), 1655-1665.
43. MacKerell, A. D.; Bashford, D.; Bellott, M.; Dunbrack, R. L.; Evanseck, J. D.; Field, M. J.; Fischer, S.; Gao, J.; Guo, H.; Ha, S.; Joseph-McCarthy, D.; Kuchnir, L.; Kuczera, K.; Lau, F. T. K.; Mattos, C.; Michnick, S.; Ngo, T.; Nguyen, D. T.; Prodhom, B.; Reiher, W. E.; Roux, B.; Schlenkrich, M.; Smith, J. C.; Stote, R.; Straub, J.; Watanabe, M.; Wiórkiewicz-Kuczera, J.; Yin, D.; Karplus, M. *The Journal of Physical Chemistry B* 1998, 102(18), 3586-3616.
44. MacKerell, A. D.; Feig, M.; Brooks III, C. L. *Journal of computational chemistry* 2004, 25(11), 1400-1415.
45. Verlet, L. *Physical Review* 1967, 159(1), 98-103.
46. Chen, J. H.; Im, W. P.; Brooks III, C. L. *Journal of the American Chemical Society* 2006, 128(11), 3728-3736.
47. Khandogin, J.; Brooks, C. L., III. *Biochemistry* 2006, 45(31), 9363-9373.
48. Oliveberg, M.; Arcus, V. L.; Fersht, A. R. *Biochemistry* 1995, 34(29), 9424-9433.
49. Tan, Y.-J.; Oliveberg, M.; Davis, B.; Fersht, A. R. *Journal of Molecular Biology* 1995, 254(5), 980-992.
50. Takahashi, T.; Nakamura, H.; Wada, A. *Biopolymers* 1992, 32(8), 897-909.
51. Bartik, K.; Redfield, C.; Dobson, C. M. *Biophysical Journal* 1994, 66(4), 1180-1184.
52. Kuhlman, B.; Luisi, D. L.; Young, P.; Raleigh, D. P. *Biochemistry* 1999, 38(15), 4896-4903.
53. Schaller, W.; Robertson, A. D. *Biochemistry* 1995, 34(14), 4714-4723.
54. Forsyth, W. R.; Gilson, M. K.; Antosiewicz, J.; Jaren, O. R.; Robertson, A. D. *Biochemistry* 1998, 37(24), 8643-8652.
55. Baker, W. R.; Kintanar, A. *Archives of Biochemistry and Biophysics* 1996, 327(1), 189-199.
56. Oda, Y.; Yoshida, M.; Kanaya, S. *Journal of Biological Chemistry* 1993, 268(1), 88-92.
57. Oda, Y.; Yamazaki, T.; Nagayama, K.; Kanaya, S.; Kuroda, Y.; Nakamura, H. *Biochemistry* 1994, 33(17), 5275-5284.
58. Plesniak, L. A.; Connelly, G. P.; McIntosh, L. P.; Wakarchuk, W. W. *Protein Science* 1996, 5(11), 2319-2328.
59. Joshi, M. D.; Hedberg, A.; McIntosh, L. P. *Protein Science* 1997, 6(12), 2667-2670.
60. Sabri Dashti, D.; Meng, Y.; Roitberg, A. E. *The Journal of Physical Chemistry B* 2012, 116(30), 8805-8811.
61. Nymeyer, H.; Gnanakaran, S.; García, A. E. In *Methods in Enzymology*; Ludwig, B.; Michael, L. J., Eds.; Academic Press, 2004, p 119-149.

Efficient Implementation of Constant pH Molecular Dynamics on Modern Graphics Processors

Evan J. Arthur¹ and Charles L. Brooks III²

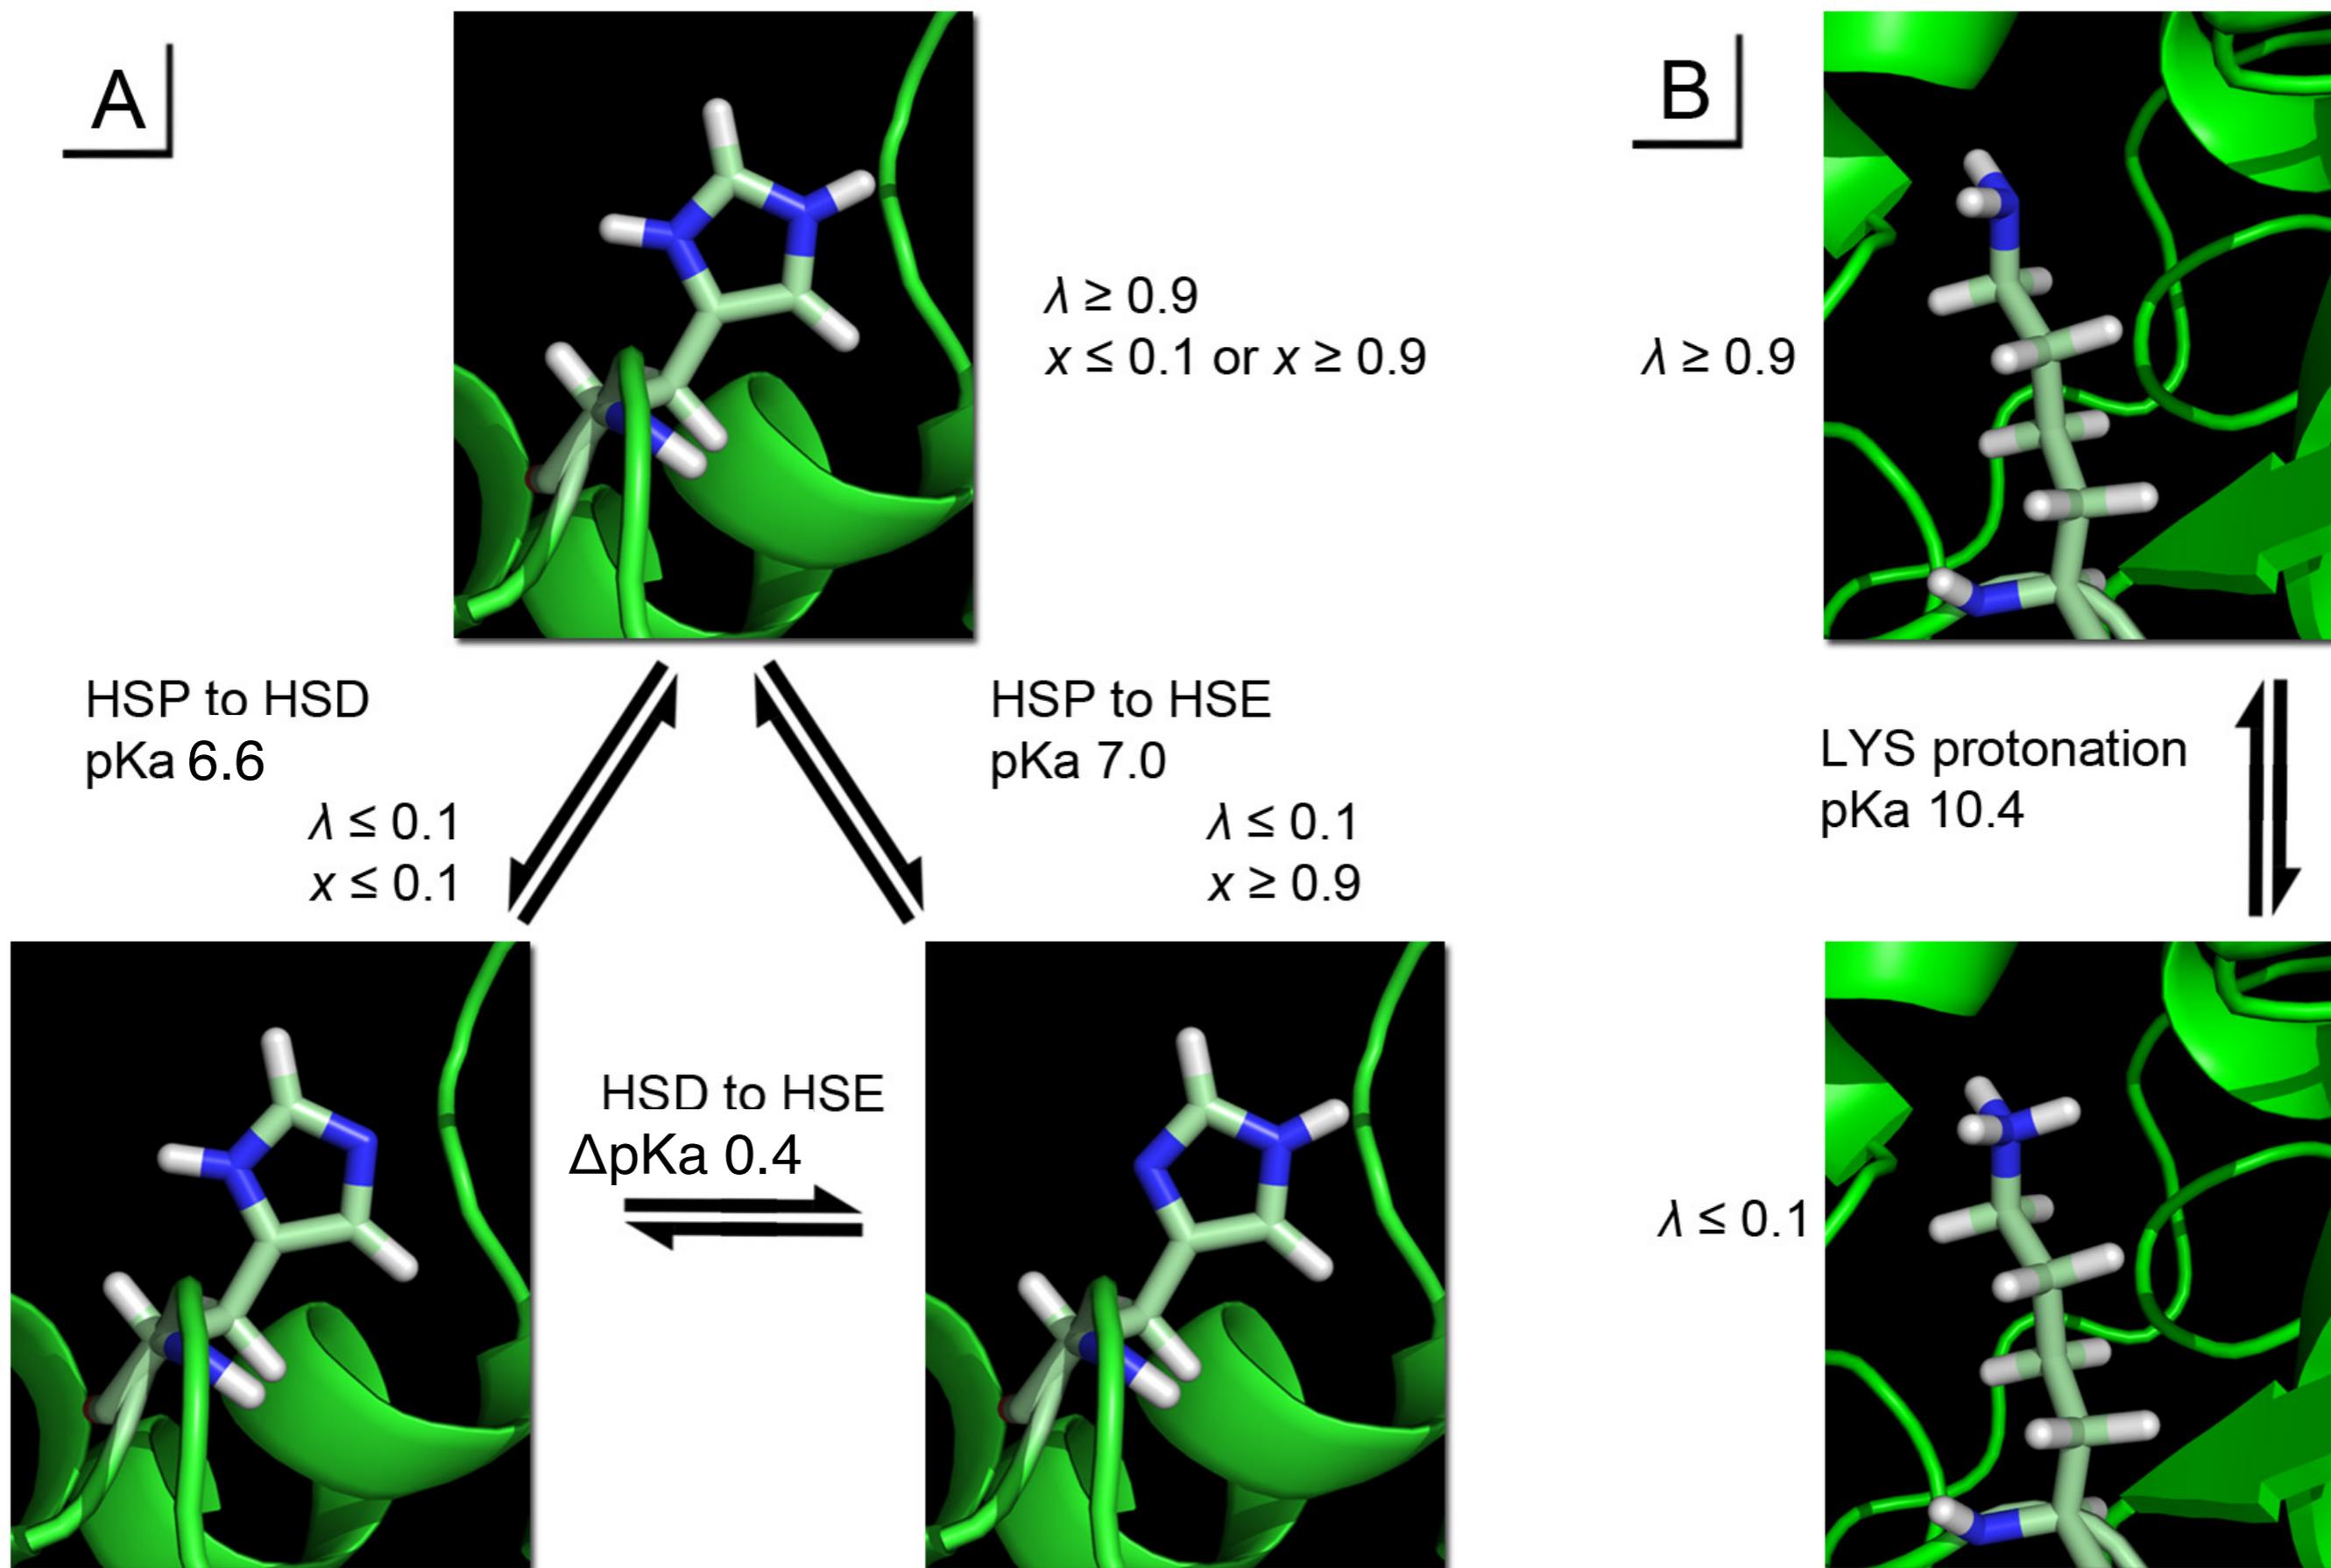
Content

Table SI-1. Computer specifications used to benchmark the new GPU-enabled CPHMD algorithm.

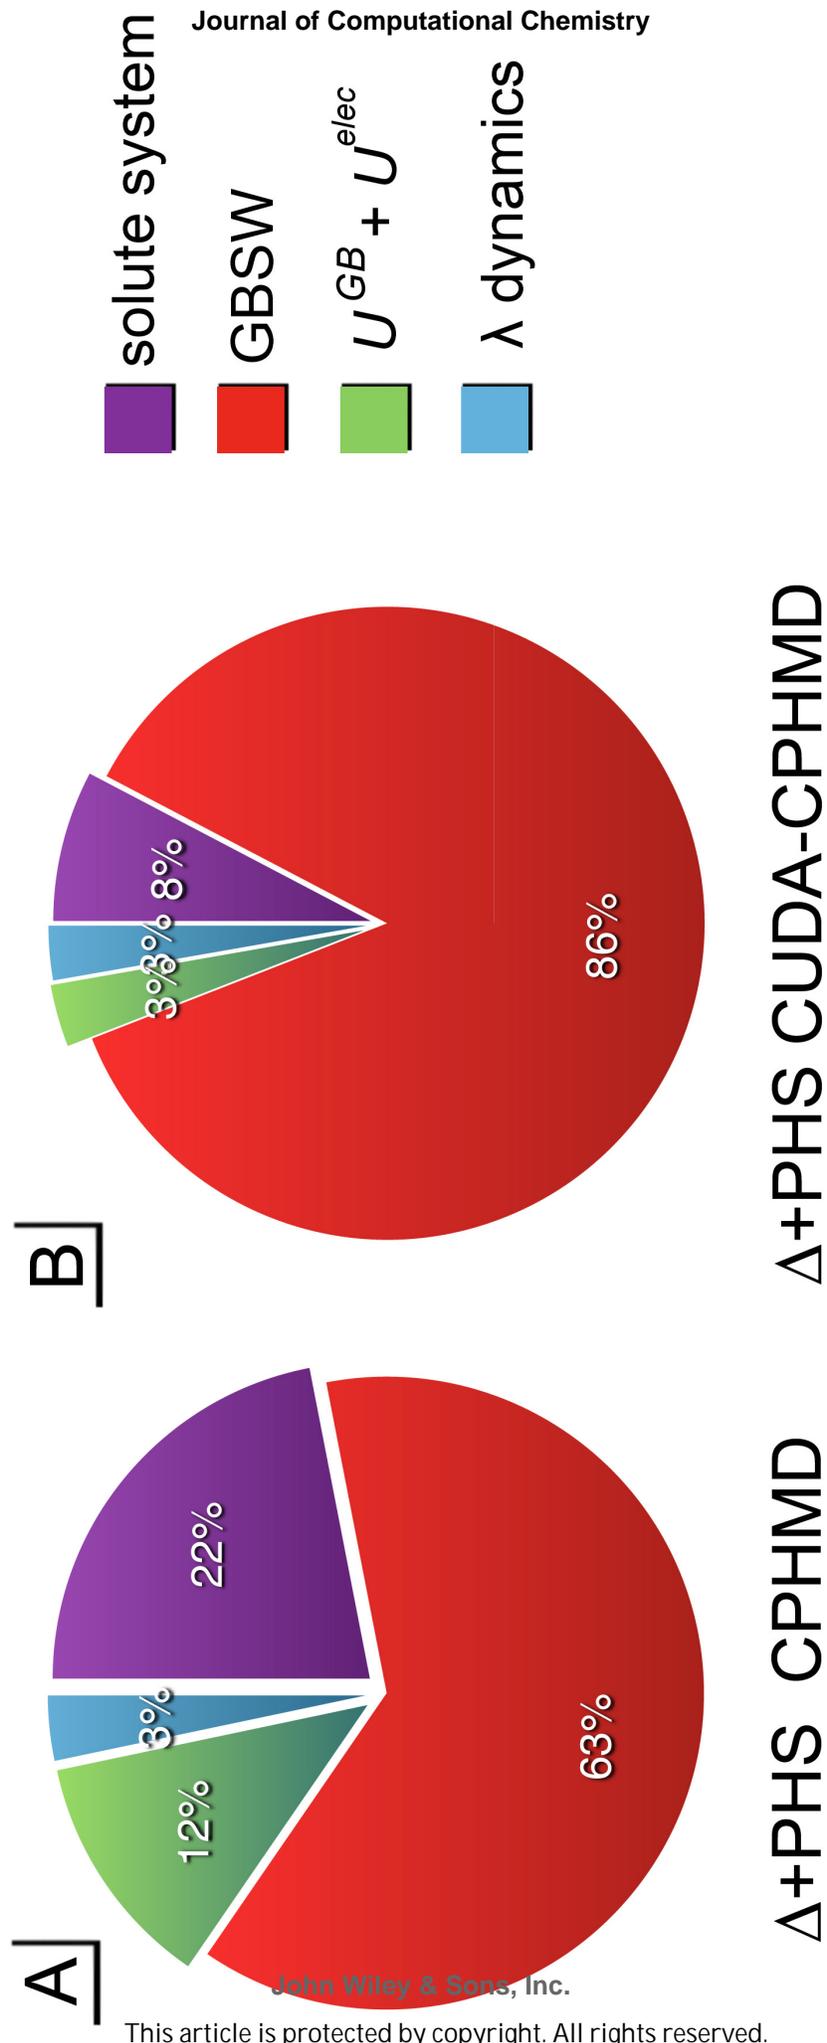
part	description	notes
main processor	2 x Intel Xeon E5-2620	2.0GHz, 6 cores each, 12 cores total
GPU	GeForce GTX 780 Ti	875 MHz, 2880 cores, 3.1 GB memory
motherboard	ATX Server Motherboard Dual LGA 2011	DDR3 1600 MHz
power supply	CORSAIR AX1200i	1200W
RAM	8 x 4GB Ripjaws Z Series DDR3	240-Pin, 1600 MHz, 32GB total
hard drive	Crucial M500 480GB SATA	480GB of storage
heatsink	2 x Supermicro SNK-P0050AP4	

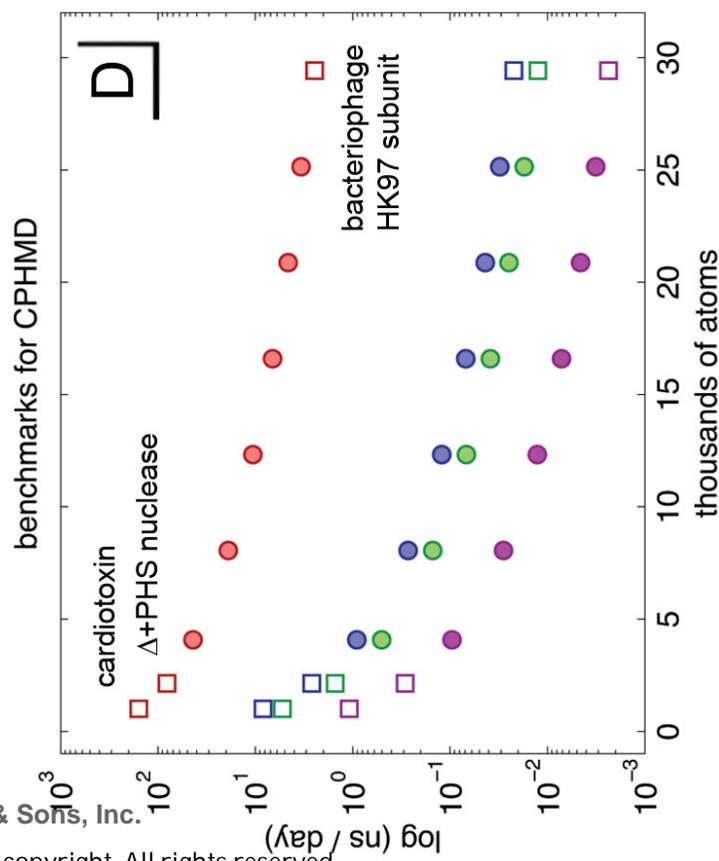
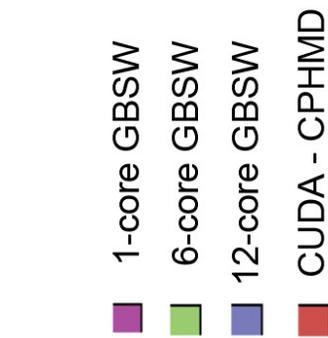
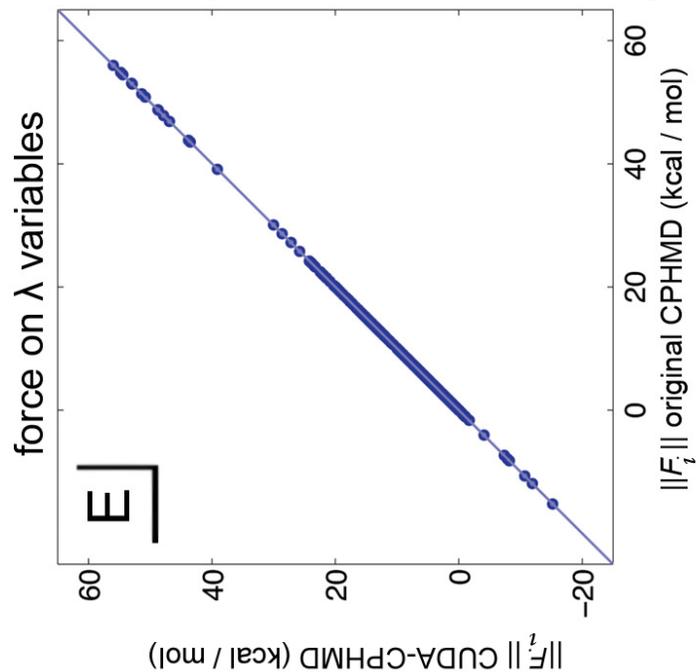
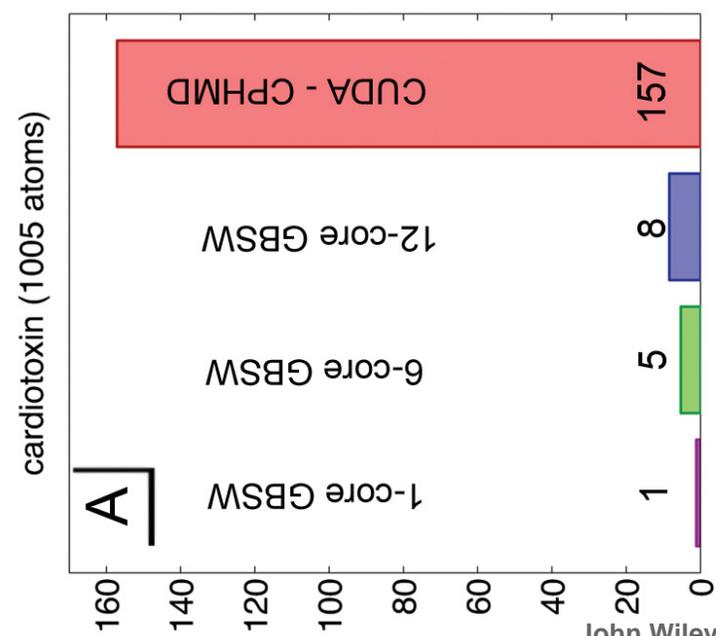
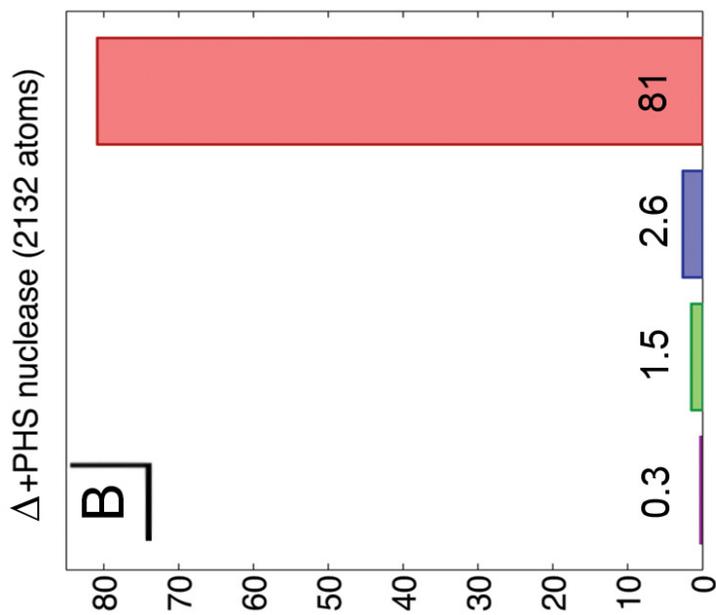
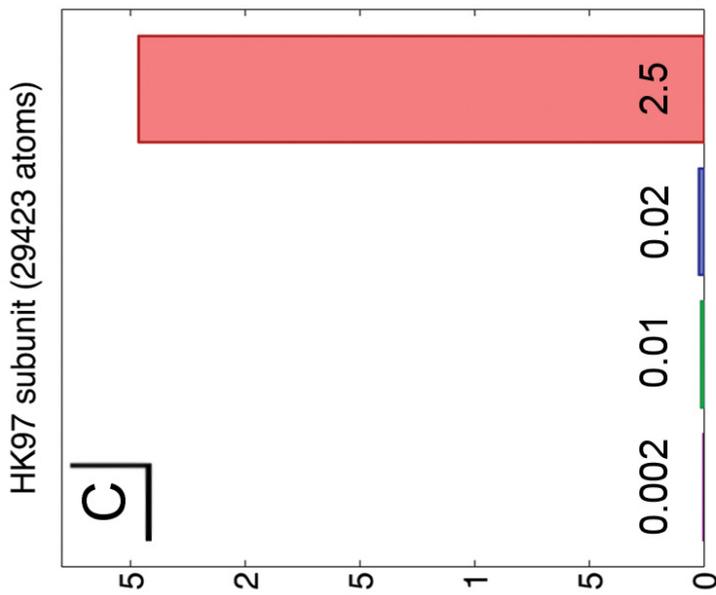
¹ Department of Chemistry, University of Michigan, 930 N. University Ave, Ann Arbor, MI 48109

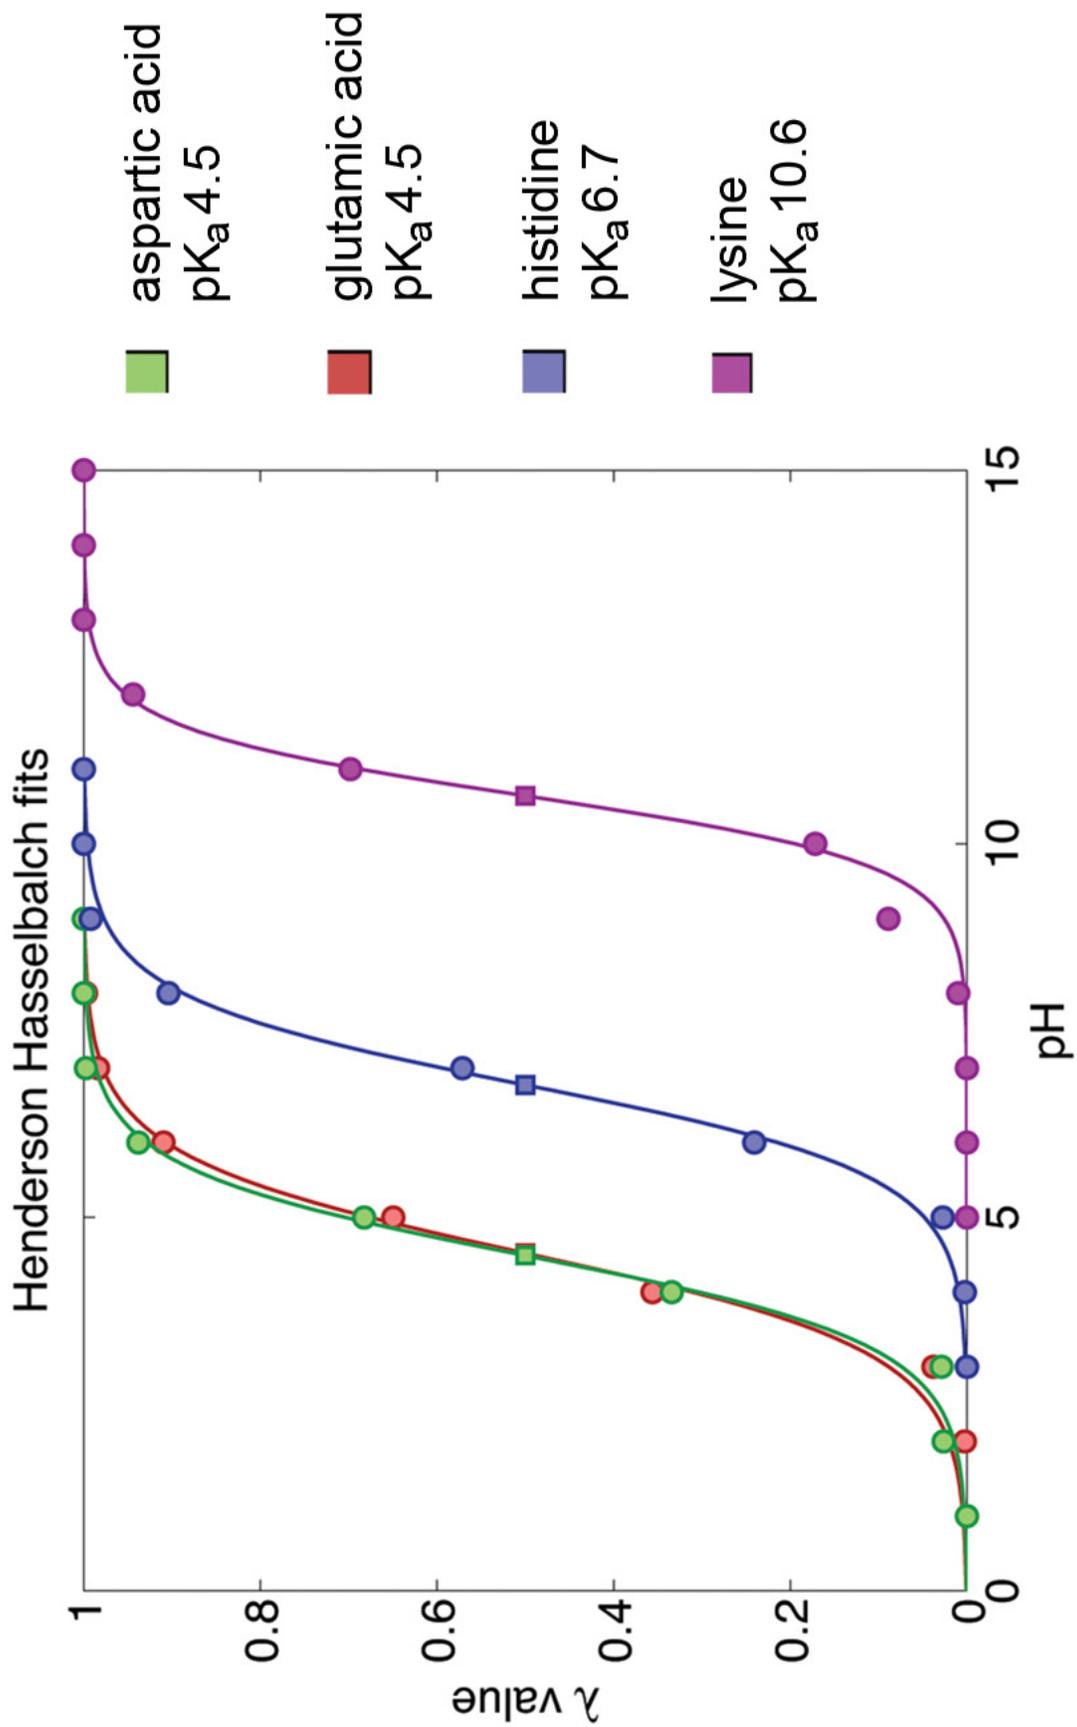
² Department of Chemistry and Biophysics Program, University of Michigan, 930 N. University Ave, Ann Arbor, MI 48109

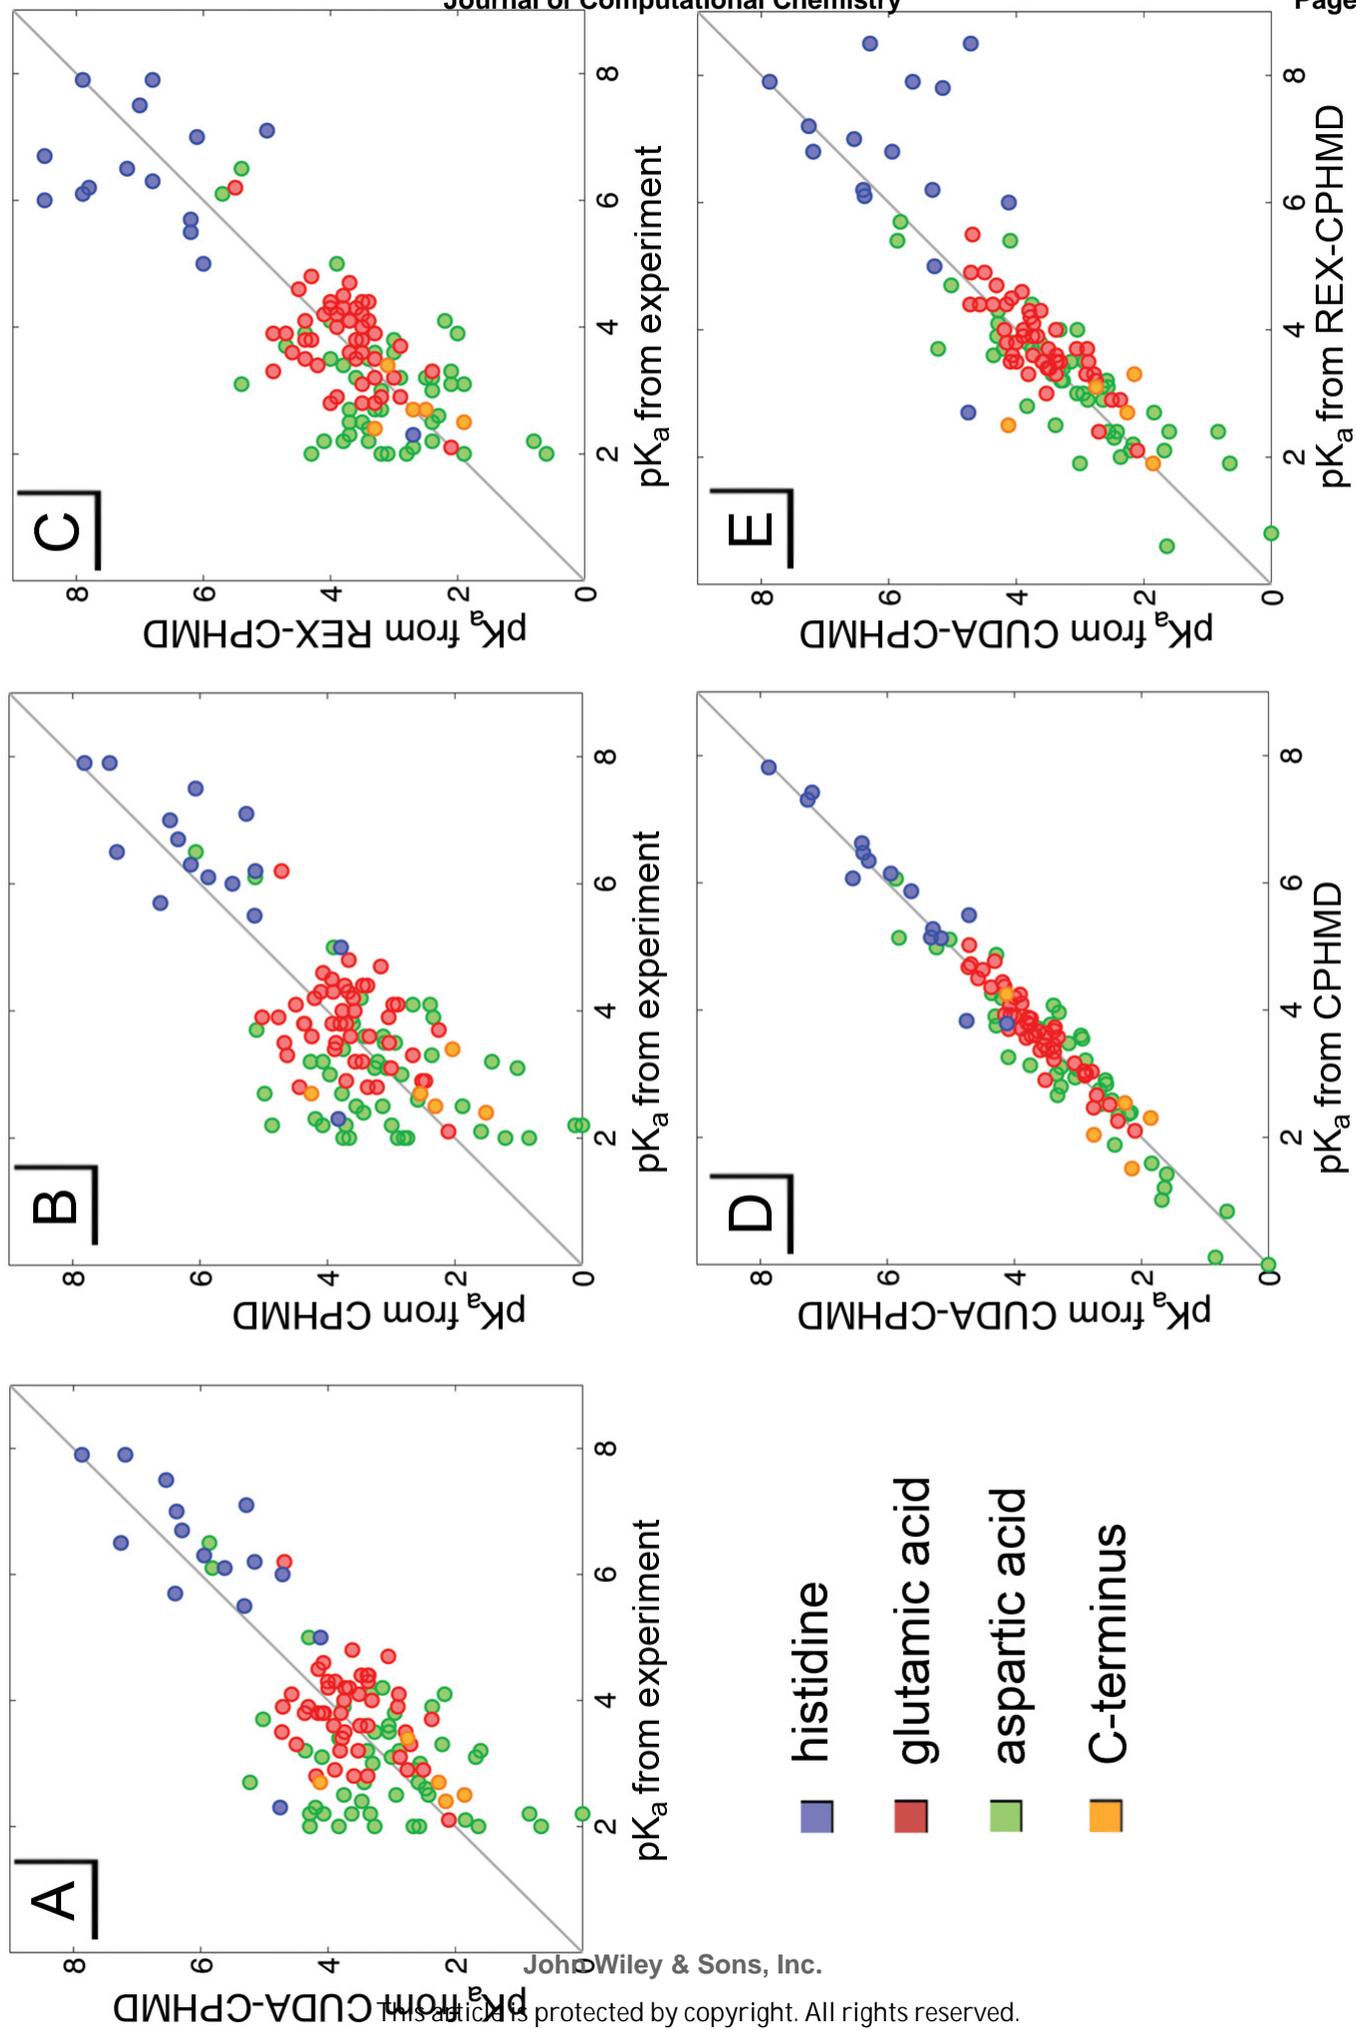


Accepted Article

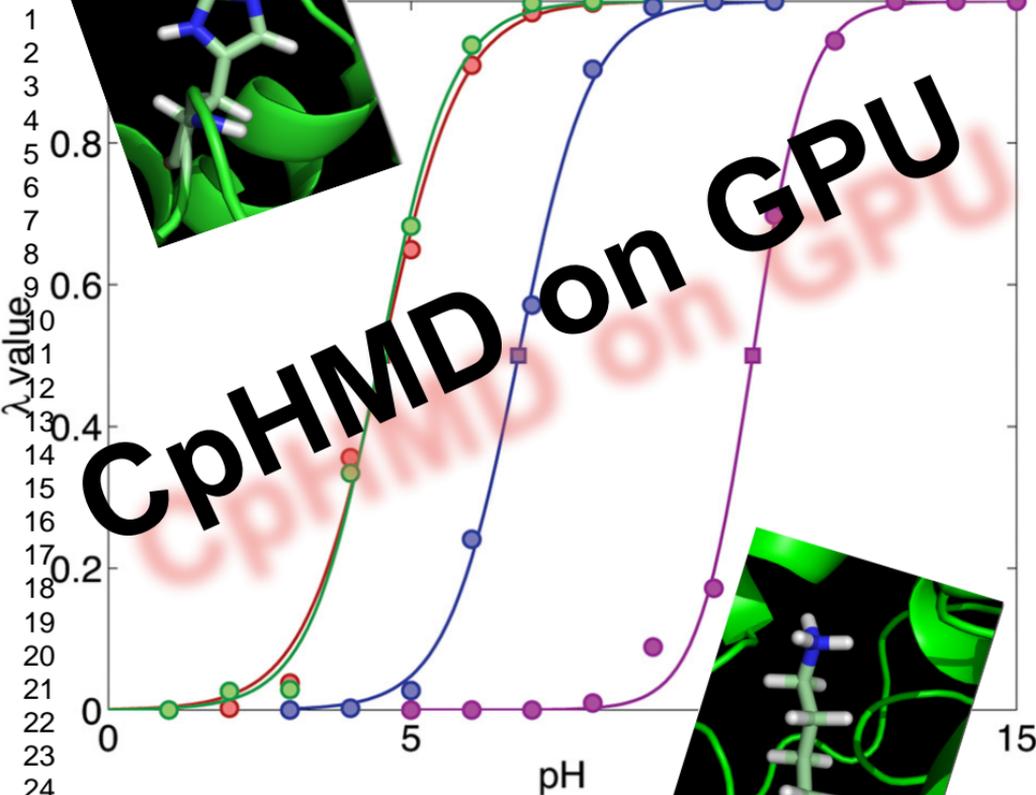








Henderson Hasselbalch fits



TOC Text: Novel implementation of constant pH molecular dynamics (CpHMD) on GPUs utilizing GPU-based GBSW implicit solvent enables fast constant pH simulations of proteins.

Accepted Article

Response to Reviewers

1. *Figure 1: Shown are cartoons -> Cartoons*

Modification was made.

2. *In Figure 1, are these pKa values correct for Histidine? Should HSD to HSE be pKa = 7?*
Thank you for pointing out this oversight. We corrected the figure.

3. *Page 4: from Equation 4 -> in Equation 4.*

Modification was made.

4. *Page 6: the potential UVDW was ignored in this implementation of CPHMD. -> the λ – dependent potential UVDW was ignored in this implementation of CPHMD*

Modification was made.

5. *Page 7: “a complete force derivation is given in the Supporting Information Eqns S11-S17.”
There is only Table SI-1 in the SI file.*

This was an oversight, since the derivatives were so simply we decided not to include them.

This text was removed from the manuscript.

6. *Figure 2: Shown are the -> The*

Modification was made.

7. *Page 9: These were NT simulations at 298K -> These were NT (constant particle number and temperature) simulations at 298K*

Modification was made.

8. *Page 9: component doesn't scale -> component does not scale*

Modification was made.

9. *page 9: “to 6.7 ns/day (a 270% speed increase).” Compared to what cutoff value?*

Manuscript text was modified here to make clear this is compared to the no cutoff case.

10. *Figure 3: Shown are the -> The*

Modification was made.

11. *In Figure 3: AUE -> AUE (average unsigned error)*

Added text.

12. *Page 10: “Interestingly, all states reported a small, systematic overestimation of the pKa “
Do you see the same result in corresponding CPU calculations?”*

I believe the text clearly states this is an artifact of the GPU implementation, most likely arising from small differences in GBSW.

13. *Figure 4: Above are the -> The*

Modification was made.

14. *Figure 5: Above are the -> The*

Modification was made.

15. *In Figure 5, and CPHMD with replica exchange -> and REX-CPHMD*

Modification was made.

16. *Page 12: “to the 0.97 pK units of AUE found using the null approximation.” What is the null approximation in this context?*

Manuscript was changed to clarify that the null model was no pKa shift from the reference values.

17. *Page 12: “What is clear, though, is additional accuracy should be possible by coupling*

CUDA-CPHMD with the enhanced sampling of replica exchange in temperature or pH.” If REX-CPHMD produces more scattered larger standard deviation of errors, why this sentence is correct?

We believe this is a result of the different sampling protocols and the lesser amount of sampling from the earlier studies. It is not our anticipation that REX-CPHMD will produce more scatter. Hence we left this sentence as is.

18. Page 12: *“Most importantly, though, is that like its predecessor, CUDA-CPHMD locates most residues with highly-perturbed pKa values.” Check the sentence. And, "though" is also in the previous sentence.*

This sentence and the one that preceded it were confusing. We have rewritten these two sentences and removed the last.

Accepted Article

Efficient Implementation of Constant pH Molecular Dynamics on Modern Graphics Processors

Evan J. Arthur¹ and Charles L. Brooks III²

Correspondence to: Charles L. Brooks III (E-mail: brookscl@umich.edu)

ABSTRACT

The treatment of pH sensitive ionization states for titratable residues in proteins is often omitted from molecular dynamics (MD) simulations. While static charge models can answer many questions regarding protein conformational equilibrium and protein-ligand interactions, pH-sensitive phenomena such as acid-activated chaperones and amyloidogenic protein aggregation are inaccessible to such models. Constant pH molecular dynamics (CPHMD) coupled with the Generalized Born with a Simple sWitching function (GBSW) implicit solvent model provide an accurate framework for simulating pH sensitive processes in biological systems. Although this combination has demonstrated success in predicting pK_a values of protein structures, and in exploring dynamics of ionizable side-chains, its speed has been an impediment to routine application. The recent availability of low-cost graphics processing unit (GPU) chipsets with thousands of processing cores, together with the implementation of the accurate GBSW implicit solvent model on those chipsets [E.J. Arthur and C.L. Brooks III, *J. Comp. Chem.* **37**:927, **2016**], provide an opportunity to improve the speed of CPHMD and ionization modeling greatly. Here we present a first implementation of GPU-enabled CPHMD within the CHARMM-OpenMM simulation package interface. Depending on the system size and non-bonded force cutoff parameters, we find speed increases of between one and three orders of magnitude. Additionally, the algorithm scales better with system size than the CPU-based algorithm, thus allowing for larger systems to be modeled in a cost effective manner. We anticipate that the improved performance of this methodology will open the door for broad-spread application of CPHMD in its modeling pH-mediated biological processes.

¹ Department of Chemistry, University of Michigan, 930 N. University Ave, Ann Arbor, MI 48109

² Department of Chemistry and Biophysics Program, University of Michigan, 930 N. University Ave, Ann Arbor, MI 48109

Introduction

Proteins typically maintain their native structure and optimal functionality under a narrow range of pH.¹⁻³ Consequently, many biological systems tightly control local solvent pH to tune the effectiveness of enzymes, or to promote a useful protein conformation.^{1,4,5} Mitochondrial ATP synthase utilizes a trans-membrane proton gradient to power its rotary catalysis mechanism,⁶⁻⁸ and the departure from a normal pH range is known to be a driving force in forming the amyloid fibrils associated with Alzheimer's disease.^{9,10} Additional examples of pH driven processes include the proton-activated gate mechanism of the KcsA potassium channel,¹¹ and the catalytic pathway of dihydrofolate reductase.¹² Finally, a notable survey by Aguilar et al. showed that about 60% of the protein-ligand complexes indicated that at least one titratable residue of the protein assumed a different protonation state between bound and unbound states.¹³ Although important to many biological processes, pH-dependence in bio-macromolecule simulations remains a nonstandard tool that awaits both wider acceptance, and finer tuning of its models.

Typical molecular dynamics (MD) simulations fix all amino acid protonation states to those of isolated residues in a neutral pH environment. While this pH-insensitive approach is sufficient to fold some proteins and observe their conformational equilibria,¹⁴ it arguably fails to capture phenomena dependent on local ionization effects of side-chains or perturbations to a residue's pK_a .^{15,16} This failure is particularly problematic for histidine residues, in that they have two hydrogen atoms that titrate with near-neutral pH. This ionizability indicates that in biologically-relevant pH environments histidine's protonation state and tautomeric configuration are often unclear.¹⁷

In recent decades a series of models of varying complexity and accuracy promise to bring accurate pH responsiveness to MD simulations.

Protonation-state modeling of amino acids in MD simulations is based on setting up a pH-sensitive extended Hamiltonian that modifies the force field parameters and structure of a given molecule. This began by discretely titrating protons, and progressing a simulation using instantaneous switches between protonated and unprotonated states. Mertz and Pettitt used an open system Hamiltonian to model the titration of acetic acid,¹⁸ and Sham et al. applied a linear response approximation through the protein-dipoles Langevin-dipoles model to calculate lysozyme residue pK_a values.¹⁹ Additional work has been done where Monte Carlo (MC) sampling guides the protonation state of an otherwise classical MD simulation. Baptista et al. used explicitly represented solvent molecules with an implicit solvent Poisson-Boltzmann (PB) function to determine protonation states.^{20,21} Meanwhile, Mongan et al. utilized generalized Born (GB) implicit solvation to add a solvation free energy component to the protonation function.²² While all these discrete models can predict pK_a values for individual amino acids to within one pK unit, they are computationally expensive. Whether the expense stems from the need to relax numerical instabilities caused by instantaneous protonation / deprotonation events, or from the MC algorithms' ability to titrate only one hydrogen at a time, such methods may require an unreasonable amount of time to study large systems with many titratable groups.

One possible solution to these issues with discrete titration methods is to use continuous titration of H^+ atoms. Brooks and co-workers developed one such method called constant pH molecular dynamics (CPHMD), which uses λ -dynamics coupled to transitions between protonation states.^{23,24} This method uses the Generalized Born implicit solvent model with a Simple sWitching function (GBSW) model,²⁵ or the related Generalized Born with Molecular Volume (GBMV) model,²⁴ to efficiently couple the protonation state to the solvation free energy of the molecule. Khandogin and Brooks then introduced proton tautomerism

capabilities to this method, which allows multi-site titrating residues, such as histidines, to be modeled more accurately.²⁶ Since the method is continuous, there are no instantaneous protonation/deprotonation events, and multiple residues can titrate simultaneously. Additionally, such continuous titration methods allow for the efficient coupling of protonation states among neighboring residues. The result is a pH simulation method that can calculate pK_a values of protein structures to within 0.75 pK units,¹⁶ and can resolve the dominant folding pathway of the pH-sensitive HdeA homodimer.¹⁵

CPHMD's efficiency, however, is bound by the rate-limited component of the calculation: the GBSW solvent model. As such, when running on a single-core central processing unit (CPU), CPHMD achieves on the order of 1 nanosecond (ns) of simulation time per day when simulating a solute system of about 1,000 atoms. Since typical uses of CPHMD, such as predicting pK_a shifts of protein residues, may require many nanoseconds of simulation time,¹⁶ even smaller proteins, such as lysozymes, may require about a week to converge on useful results. Larger systems, such as asymmetric viral capsid subunits with tens of thousands of atoms, may require unreasonably long simulation times if captured in full atomic detail.^{27,28} Fortunately, the GBSW solvent model has recently been refactored to function on new, parallel graphics processing unit (GPU) hardware, and is now between 1 and 2 orders of magnitude faster than its CPU counterpart.^{16,29} By incorporating the CPHMD model into the GPU-GBSW algorithm, there holds the promise of speeding up pH simulations substantially.

This study represents an increment in the ongoing adaptation of efficient and useful algorithms onto parallel-processing GPUs. Such chipsets can contain thousands of processing cores, and are able to process C-like languages

such as Open Computing Language (OpenCL) and Compute Unified Device Architecture (CUDA). This combination of features has opened up a new frontier of parallel processing where expensive computer clusters can be replaced with single, affordable graphics cards. Simulation packages such as CHARMM,³⁰ AMBER,³¹ OpenMM,³² GROMACS,³³ and NAMD³⁴ all offer GPU-accelerated options for many types of studies, and most of those options receive speed increases of greater than an order of magnitude over their CPU counterparts.

Due to OpenMM's effectiveness in harnessing the capabilities of GPUs with a wide variety of hardware, a CHARMM-OpenMM interface was developed to combine the strengths of both simulation packages.^{30,32} CHARMM's robust algorithms can be used to design and parameterize a simulation, and OpenMM's efficient GPU-based algorithms can be used to propagate dynamics.^{30,32} Now with the recent incorporation of the GBSW solvent model into the CHARMM-OpenMM interface, many of CHARMM's algorithms parameterized for use with GBSW, such as CPHMD, can be adapted for parallel processing on GPUs as well. In this study we take advantage of the recent incorporation of GBSW onto GPUs, and discuss the adaptation of CPHMD onto this new parallel architecture. First we explain the underlying theory behind λ -dynamics: how a λ coordinate is used to represent the titration state of a residue, and how that coordinate is propagated. Then we delve into how it was originally implemented for CHARMM, and examine fitting CPHMD into the GBSW algorithm. Here we discuss the algorithmic improvements, and show how many force contributions on λ are calculated alongside the free energy of solvation. Finally we present benchmarks achieved by the new algorithm, and comment on future directions for pH simulations.

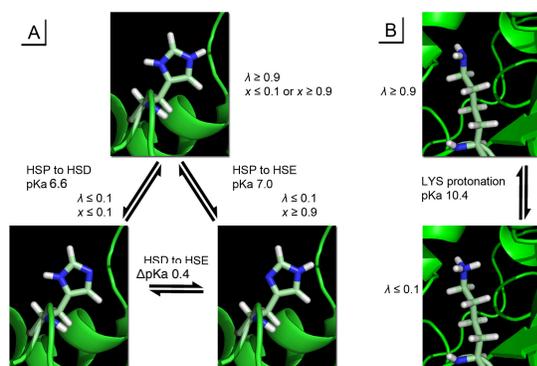


Figure 1. Shown are Cartoons of the protonated and unprotonated states of A) histidine and B) lysine. Also noted are the reference pK_a values of each transition when occurring in an isolated residue, as well as the λ values at each state.

Methods

The λ coordinates and their underlying energy function for single-site titration

For clarity in following discussions, we present the underlying theory of CPHMD. We start by setting up the framework for a single residue with one titrating hydrogen. The rudimentary picture of titration events is an equilibrium association/dissociation reaction of a model compound $A_{(aq)}$ in aqueous solution from a titrating proton.



Here, the protonation free energy is defined by

$$\Delta G^{\text{exp}}(\text{model}) = -k_B T \ln 10(pK_a^{\text{exp}} - pH) \quad (2)$$

where k_B is Boltzmann's constant, and T is the temperature. We can approximate the above equations through classical simulations by interpreting the protonation interaction as a change in free energies:

$$\begin{aligned} & \Delta G^{\text{exp}}(\text{protein}) - \Delta G^{\text{exp}}(\text{model}) \\ &= \Delta G^{\text{classical}}(\text{protein}) - \Delta G^{\text{classical}}(\text{model}) \end{aligned} \quad (3)$$

This relationship then leads to an estimate of experimental free energy of protonation for a single titrating site:

$$\begin{aligned} \Delta G^{\text{exp}}(\text{protein}) &= \Delta G^{\text{classical}}(\text{protein}) \\ &- \Delta G^{\text{classical}}(\text{model}) + \Delta G^{\text{exp}}(\text{model}) \end{aligned} \quad (4)$$

From this perspective, we infer that titratable groups have an intrinsic free energy of protonation that is perturbed by the protein environment mainly through non-bonded interactions. We model this perturbation by extending the system's Hamiltonian with a non-geometric dimension of λ . As mentioned in the introduction, the CPHMD model uses a series of λ coordinates where each λ value tracks the progress of protonation-deprotonation events at a single titration site. For a particular residue i , these coordinates are generated from

$$\lambda_i = \sin^2(\theta_i) \quad (5)$$

where i is the residue being titrated. In this form the θ variable is bound to all real numbers, and λ is bound to the continuous range $0 \leq \lambda_i \leq 1$. The sine-squared function then favors λ values near the boundary protonated (1) and unprotonated (0) states. Because λ is only physically relevant as it nears these boundary states, we impose cutoffs on interpreting λ . In CPHMD an unprotonated state is $\lambda_i \leq 0.1$, a protonated state is $\lambda_i \geq 0.9$, and a mixed state is $0.1 < \lambda_i < 0.9$. Figure 1 illustrates the protonation states and their corresponding λ values. Potentials and their derivative forces on λ are then interpreted as potentials and forces on θ .

The potential energy that governs protonation states contains five λ -dependent components. We start with the pH dependence of the deprotonation free energy as follows from ΔG^{exp} in Equation 4. This potential connects λ to the pK_a of a residue in its isolated, reference state:

$$U^{pH}(\lambda_i) = \lambda_i (pK_a(i) - pH) (k_B T \ln 10) \quad (6)$$

Here, $pK_a(i)$ is the pK_a of titrating group i . Next we have the potential of mean force (PMF) along the λ coordinate from ΔG^{model} . This term corresponds to the negative of free energy needed to deprotonate a model residue:

$$U^{model}(\lambda_i) = A_i (\lambda_i - B_i)^2 \quad (7)$$

Equation 7 is a quadratic fit to the thermodynamic work potential of deprotonating a model compound, and it splits the protonation state into two low-energy wells that represent the protonated and unprotonated states. Then a barrier potential is added that disfavors mixed states of λ :

$$U^{barrier}(\lambda_i) = 4\beta_i (\lambda_i - 1/2)^2 \quad (8)$$

The barrier scaling parameter β_i is an empirical coefficient designed to tune the propensity for a λ value to remain in either protonated or unprotonated states, while facilitating transitions between them. In the current iteration of CPHMD, β_i assumes a value of either 2.5 or 1.75 kcal/mol, depending on the residue. Finally, we arrive at the two charge-dependent potentials: the Coulombic and generalized Born. The classical Coulombic potential is

$$U^{elec}(\lambda_i) = \sum_{a,i} \sum_b K^{elec} \frac{q_{a,i}(\lambda_i) q_b}{r_{ab}} \quad (9)$$

Here K^{elec} is Coulomb's constant, q_a and q_b are the partial charges of atoms a and b respectively, and r_{ab} is the distance between those atoms. Note that this potential for residue i includes the interactions between all atoms a in residue i to all other atoms in the system. Meanwhile, $q_a(\lambda_i)$ is a λ -dependent charge of atom a , which follows the form

$$q_{a,i}(\lambda_i) = \lambda_i q_{a,i}^{unprot} + (1 - \lambda_i) q_{a,i}^{prot} \quad (10)$$

where charges on titrating atom a can be in protonated ($q_{a,i}^{prot}$) and unprotonated ($q_{a,i}^{unprot}$) states. We note that in an effective charge model of pH, titrating residues are allowed to interact. As such, any atom b from a titrating residue j interacting with residue i has its own $q_{b,j}^{prot}$ and $q_{b,j}^{unprot}$. Thus the partial charge q_b follows one of two possibilities:

$$q_b = \begin{cases} q_b & \text{non-titrating} \\ \lambda_j q_{b,j}^{unprot} + (1 - \lambda_j) q_{b,j}^{prot} & \text{titrating} \end{cases} \quad (11)$$

That is if atom b lies in a non-titrating residue, that atom's partial charge is simply the standard partial charge from that residue's force field. If atom b lies in a titrating residue j and its charge is affected by the protonation state of j , then its partial charge is derived from the same λ -dependent relationship from Equation 10. Since atoms near a titrating site can have their partial charges affected by titration states, many more than the titrating hydrogen atoms can possess a λ -dependent charge state. We also note that at times $j=i$ if we observe the Coulombic interaction between two atoms on the same titrating residue. The final λ -dependent potential is that from the GB solvent model as expressed in the Still equation:³⁵

$$U^{GB}(\lambda_i) = \sum_{a,i} \sum_b \tau \frac{q_{a,i}(\lambda_i) q_b}{f_{ab}^{GB}} \quad (12)$$

where

$$f_{ab}^{GB} = \left[r_{ab}^2 + R_a^{Born} R_b^{Born} \exp\left(-r_{ab}^2 / (4R_a^{Born} R_b^{Born})\right) \right]^{1/2} \quad (13)$$

Here, $q_a(\lambda_i)$ and q_b follow the same form as in Eqn 10 and 11 respectively; r_{ab} is the distance between atoms a and b ; τ is the factor that

scales the Born energy by the difference in dielectric values at the dielectric boundary and by any contributing salt effects,³⁶ and the values R_a^{Born} and R_b^{Born} represent the Born radii of atoms a and b respectively. The Born radii are the effective distance between an atom and the solute-solvent dielectric boundary, and they are calculated through volumetric integration following the GBSW implicit solvent model.²⁵

If we pull together the complete potential for a titrating residue i from Eqns 6 through 13, then we arrive at the form

$$U_i^{total}(\lambda_i) = U_i^{pH}(\lambda_i) + U_i^{model}(\lambda_i) + U_i^{barrier}(\lambda_i) + U_i^{elec}(\lambda_i) + U_i^{GB}(\lambda_i) + U_i^{VDW} + U_i^{internal} \quad (14)$$

The so-called “internal energy” term ($U^{internal}$) corresponds to the bond, angle, and torsional energy terms of a classical energy force field. In this model, the titration state is dynamically independent of this potential. Although several models of CPHMD include a λ -dependent van der Waals term (U^{VDW}),^{26,37,38} during this study it was found that at most it contributes a minimal amount to a given residue’s force on λ , while it nearly doubles the calculation time of CPHMD. This term is negligible compared to the force on λ from other effects, and omitting it from the calculation showed no effect on the accuracy of CPHMD. Thus, in the interest of speeding up the original algorithm, the λ -dependent potential U^{VDW} was ignored in this implementation of CPHMD.

Although we now have the proper setup for addressing residues with a single titration site, such as in lysine, we need to address how CPHMD handles tautomerization in residues such as in aspartic acid and histidine.

Proton tautomerism

Similarly to how one λ variable is used to track the progress of titration states of a residue,

Khandogin and Brooks incorporated tautomeric behavior into CPHMD by providing residues with a second λ variable, called x , to track the progress of tautomeric states.²⁶ This arrangement is illustrated in Figure 1a with histidine. Just as in λ dynamics for titration states, transitions between tautomeric states are linearly interpolated using the x variable. What results are potentials that become bivariate in λ and x , and each tautomeric residue has four charge states: tautomer A in protonated and unprotonated states, and tautomer B in protonated and unprotonated states. What we shall see later is that residues can have equivalent states in this setup. Histidine’s protonated state, for example, is a residue saturated with protons. As such tautomers A and B of the protonated state are equivalent. We now review the influence of including two λ parameters for a tautomeric titrating residue.

The pH dependent potential becomes

$$U^{pH}(\lambda_i, x_i) = \lambda_i [x_i (pK_a^A(i) - pH) + (1 - x_i)(pK_a^B(i) - pH)] (k_B T \ln 10) \quad (15)$$

where the pK_a values of tautomers A and B are pK_a^A and pK_a^B respectively. While these pK_a values for aspartic acid and glutamic acid are equivalent and only serve as a sampling expedient,²⁶ in residues with asymmetric titrating sites such as histidine they are not. The PMF for protonation becomes a bivariate polynomial from Equation 7, which then expands into the general form

$$U^{model}(\lambda_i, x_i) = a_0 \lambda_i^2 x_i^2 + a_1 \lambda_i^2 x_i + a_2 \lambda_i x_i^2 + a_3 \lambda_i x_i + a_4 \lambda_i^2 + a_5 x_i^2 + a_6 \lambda_i + a_7 x_i + a_8 \quad (16)$$

The barrier potential is simply a summation of terms that disfavor the mixed states of both λ and x , and follows the form

$$U^{\text{barrier}}(\lambda_i, x_i) = 4\beta_i^\lambda (\lambda_i - 1/2)^2 + 4\beta_i^x (x_i - 1/2)^2 \quad (17)$$

Note that there are two barrier scaling parameters β_i^λ and β_i^x for λ and x . Although different biases for tautomeric and protonation transitions are possible in this equation, in the discussed CPHMD model they are identical for all titrating residues.

The charge-dependent potentials in Eqns 9 and 12 are only modified in that charges for atoms can now be dependent on the new x coordinate. The Coulombic and generalized Born potentials then follow the forms

$$U^{\text{elec}}(\lambda_i, x_i) = \sum_{a,i} \sum_b K^{\text{elec}} \frac{q_{a,i}(\lambda_i, x_i) q_b}{r_{ab}} \quad (18)$$

and

$$U^{\text{GB}}(\lambda_i, x_i) = \sum_{a,i} \sum_b \tau \frac{q_{a,i}(\lambda_i, x_i) q_b}{f_{ab}^{\text{GB}}} \quad (19)$$

respectively. The bivariate charge $q_{a,i}(\lambda_i, x_i)$ then follows the form

$$q_{a,i}(\lambda_i, x_i) = \lambda_i \left[x_i q_{a,i}^{\text{A,unprot}} + (1-x_i) q_{a,i}^{\text{B,unprot}} \right] + (1-\lambda_i) \left[x_i q_{a,i}^{\text{A,prot}} + (1-x_i) q_{a,i}^{\text{B,prot}} \right] \quad (20)$$

Where charges on titrating atom a are derived from the protonated and unprotonated variants of both A and B tautomers, $q_{a,i}^{\text{A,prot}}$, $q_{a,i}^{\text{A,unprot}}$, $q_{a,i}^{\text{B,prot}}$, and $q_{a,i}^{\text{B,unprot}}$. Similarly, the charge on atom b emerges as

$$q_b = \begin{cases} q_b & \text{non-titrating} \\ \lambda_j \left[x_j q_{b,j}^{\text{A,unprot}} + (1-x_j) q_{b,j}^{\text{B,unprot}} \right] + (1-\lambda_j) \left[x_j q_{b,j}^{\text{A,prot}} + (1-x_j) q_{b,j}^{\text{B,prot}} \right] & \text{titrating} \end{cases} \quad (21)$$

We now arrive at a general-purpose setup for evaluating the underlying potential for continuous transitions among various charge states of a particular residue. Deriving the forces with respect to λ and x , while important, serves little purpose for illuminating the topics explored in the remainder of this study. With the framework above, we now can discuss the construction of the original algorithm, and the changes made to refactor it for efficient parallel processing on GPUs.

Refactoring CPHMD

The original CPHMD model was built with mathematical precision and function portability in mind. It is a stand-alone module that can be applied to both implicit and explicit solvent systems, and except for atom coordinate and Born radii updates, it receives no input from other functions during a simulation. During the course of a timestep, each titrating coordinate λ_i is scanned to identify the residue type (such as whether the residue has one or two titrating hydrogens), and then an appropriate functional is applied to calculate its pH (Eqn 6 and 15), model (Eqn 7 and 16), and barrier (Eqn 8 and 17) potentials. Next, neighboring atom-atom interactions are scanned for whether one or both atoms reside in titrating groups. If a titrating atom-atom pair is found, then contributions to the electrostatic (Eqn 9 and 18) and GB (Eqn 12 and 19) potentials are integrated. Neighboring atom-atom pairs are then scanned again to calculate the VDW potential (ignored in this new iteration of CPHMD). Finally, the force on θ is calculated, and λ via θ is advanced a timestep using Langevin dynamics.³⁹ In this setup there are several opportunities presented to us for improving the algorithm both in the efficiency of its execution in parallel, and by weaving portions of the calculation into existent functions elsewhere in the simulation.

We first note that the majority of clock cycles used for calculating λ dynamics are spent on neighboring atom-atom interactions when

accumulating the electrostatic and GB potentials. While the calculations required for each atom pair are computationally cheap, the large number of interatomic interactions in a protein containing thousands of atoms can make this multitude of cheap calculations altogether expensive. As shown in Figure 2a, about 12% of a 2000-atom simulation is spent only on this calculation.

Both CPHMD and the GBSW solvent model require calculating the Still equation (Eqn 12 and 13) to address part of the neighboring atom potential, so a significant speed improvement can be made by placing all of CPHMD's atom-atom processes inside the neighboring atom process of the GBSW solvent model. This way, as GBSW produces the solute molecule's electrostatic solvation free energy and its derivative force on atoms, CPHMD processes neighboring atom potentials on λ simultaneously. Thus the large number of redundant atom-atom distance calculations can be reduced significantly during a simulation. This setup gains additional speedup through GBSW by using OpenMM's efficient parallel processing of neighboring-atom interactions. As shown in Figure 2, by combining the CPHMD and GBSW algorithms we see that pH modeling with CPHMD accounts for a much smaller fraction of the overall simulation time.

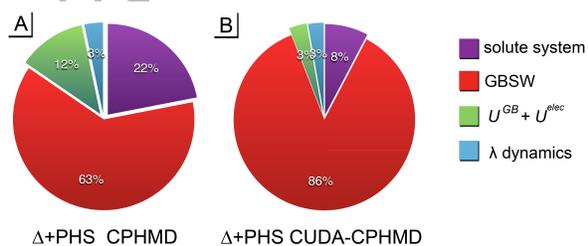


Figure 2. The approximate distributions of CPU time spent on running simulation components of Δ +PHS staphylococcal nuclease molecule. This protein contains 2132 atoms and 37 titrating residues. A) run using the original algorithm on a single processing core in CHARMM. B) run using the newly refactored CUDA-CPHMD algorithm.

Due to the nature of parallel processing, bottlenecks are often created from the longest

portions of non-parallel code. While a single-core process can be sped up dramatically by creating a case-by-case set of calculations, navigating through the additional overhead to make the situation-specific decision can slow parallel processes down. Regarding the equations described earlier, a titrating residue with one tautomer requires fewer calculations than a titrating residue with two. As we place each residue's force calculations in parallel processes, however, the speed of the code is improved by regarding all titrating residues as possessing two tautomeric states. In this new implementation of CPHMD, single-titration residues, such as lysine, are given extraneous x coordinates. Lysine then uses the barrier potential from Eqn 16, where the x -coupled coefficients a_0 , a_1 , a_2 , a_3 , and a_5 are set to a value of 0.0. Without the overhead for residue identification, the longest calculation required, that is calculating the force on θ for a residue with two tautomeric states, is shortened. What results is a speed improvement when calculating all components of the total potential on λ coordinates. As shown in Figure 2b, using the parallel CUDA-CPHMD algorithm for a small system impacts the processing time by approximately 6%, as opposed to 15% for the original algorithm.

Benchmarking CUDA-CPHMD

We finally reach an efficient setup where using the CPHMD model results in little slowdown of the overall simulation time. We chose several systems to benchmark the new algorithm, and explore the speed benefits it offers. We chose the naja atra snake cardiotoxin (PDB: 1CVO),⁴⁰ the Δ +PHS hyperstable variant of staphylococcal nuclease (PDB: 3BDC),⁴¹ and the asymmetric subunit of the bacteriophage HK97 head capsid (PDB: 2FT1).⁴² This trio provided a range of system sizes and residue configurations. To add additional statistics, the 7 proteins of the HK97 capsid were assembled into 6 additional subsystems, all of which appear in Figure 3 to show for a range of system sizes the speed dependence on system size. All simulations

were using the CHARMM22 force field^{43,44} using the Langevin integrator⁴⁵ with a timestep of 2 femtoseconds. These were NT (constant particle number and temperature) simulations at 298K in unbounded volumes using the CUDA-GBSW solvent model, and CUDA-CPHMD to model titration states and advance λ coordinates. Atomic radii for the GBSW solvent model were provided through work by Chen et al.⁴⁶ The hardware specifications of the computer used appear in Table SI1 of the Supporting Information. We found speed improvements of between 1 and 3 orders of magnitude in the CUDA-CPHMD algorithm over its CPU counterpart.

As we combine the improved efficiency and parallel execution of both GBSW and CPHMD (shown in Figure 3a to 3d), substantial speed gains are found in this new version of pH modeling over its predecessor. For smaller

1,000-atom systems, we see a speed improvement of over 20-fold when comparing a 12-threaded CPHMD simulation to the new CUDA-CPHMD, and an improvement of over 150-fold when compared to the single-core algorithm (shown in Figure 3a). For larger 29,000 atom systems, we see a speed improvement of over 1,000-fold (shown in Figure 3c). Since the neighboring-atom component does not scale linearly with system size, larger systems experience a greater calculation time penalty than smaller ones. Fortunately, simple changes, such as using non-bonded cutoffs, can mitigate such problems. For instance, a non-bonded cutoff of 14 Å sped up the large viral capsid simulation to 6.7 ns/day (a 270% speed increase versus the no cutoff case).

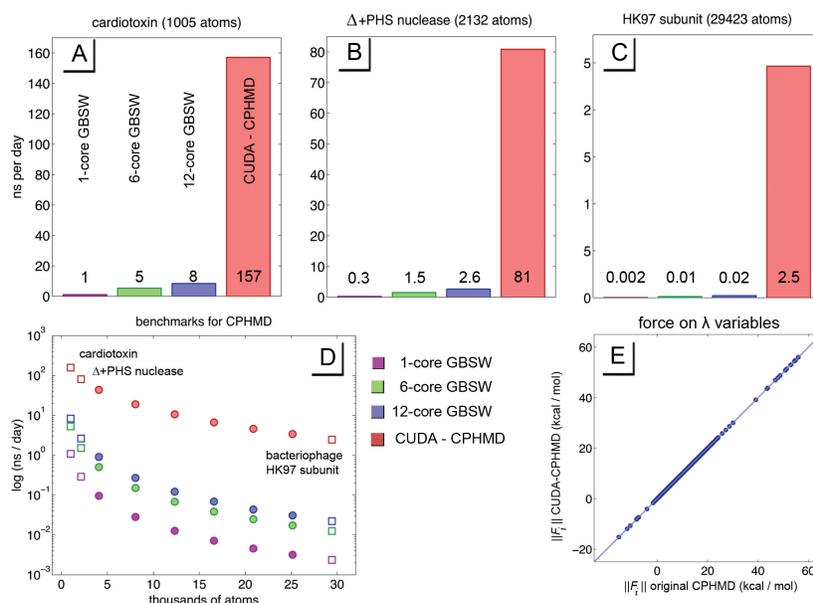


Figure 3. The benchmarks for the new CUDA-CPHMD algorithm. The individual systems tested were A) the naja atra snake cardiotoxin (PDB: 1CVO); B) the Δ +PHS hyperstable variant of staphylococcal nuclease (PDB: 3BDC); and C) the asymmetric subunit of the bacteriophage HK97 head capsule (PDB: 2FT1). As shown, the new algorithm is substantially faster than the original CPU algorithm by up to 3 orders of magnitude. In D) the same benchmarks from earlier are shown (squares) alongside subsystems from the 7 proteins of the bacteriophage subunit (circles). Notice that the CUDA algorithm scales more linearly with system size than its CPU-based counterpart. E) compares the force on λ as calculated on all 595 λ coordinates from both CPHMD algorithms. There is less than a 0.00017 (kcal/mol) AUE (average unsigned error) between the two algorithms.

Accuracy of the new CUDA-CPHMD algorithm

Single-residue Systems

Speed gains in implementing CPHMD are an important goal both for increasing the algorithm's applicability to a wider range of system sizes, and for its ability to converge on useful results more rapidly. Its accuracy, however, must not be compromised as we reconfigure the execution of the algorithm. In Figure 3e we show that there is little difference between the original CPHMD and CUDA-CPHMD algorithms when calculating the force on λ . We maintain an average unsigned error (AUE) of less than 0.00017 kcal/mol in this force. We also note that 99.9% of the AUE between the two CPHMD methods is from the slight differences in Born radii calculated from the original and CUDA implementations of GBSW. Thus, we conclude that CUDA-CPHMD accurately reproduces the original algorithm's force on λ .

While CUDA-CPHMD may be able to produce the force on λ coordinates, we ran additional tests to see whether or not residue protonation states are also reproduced. Due to each residue's pH-dependent biasing potential, a single residue alone in solution presumably should find an optimal protonation state depending on the environmental pH. At pH environments below a residue's pK_a the residue should favor a protonated state ($\lambda_i \leq 0.1$), and conversely a residue exposed to a pH above its pK_a should favor an unprotonated state ($\lambda_i \geq 0.9$). By calculating the fraction of protonated to unprotonated states of residues at various pH values and fitting the results to the Henderson-Hasselbalch equation of states, we expect the point of inflection to reproduce the pK_a of that residue.

We ran simulations of aspartic acid, glutamic acid, histidine, and lysine to calculate their protonation states, as shown in Figure 4. These residues were simulated using the same setup from the benchmarking section as NT simulations in an unbound volume, and CUDA-CPHMD was used both to model titration states and advance λ coordinates. The backbone atom ends were capped with the ACE and CT2 patches in CHARMM. Each dot in Figure 4 represents the average residue titration state from 200 ps of simulation time, and the residues ran at an average speed of 690 ns/day.

We find that without optimizing the simulations for speed, accuracy, or convergence of protonation states, that the pK_a values could be captured to within 0.5 pK units. Interestingly, all states reported a small, systematic overestimation of the pK_a , and the exact source of this discrepancy remains unclear. The CUDA-GBSW solvent model overestimates solvation energy by an average of approximately 0.16 kcal/mol. However, this overestimation of energy should bias deprotonation events to occur slightly more often, and thus lower the calculated pK_a . What is clear from these data, though, is that like its predecessor, the CUDA-CPHMD algorithm models the pH dependence of titration well for single residue systems. Next we explore multi-residue titration and the influence of protein conformation on pK_a values.

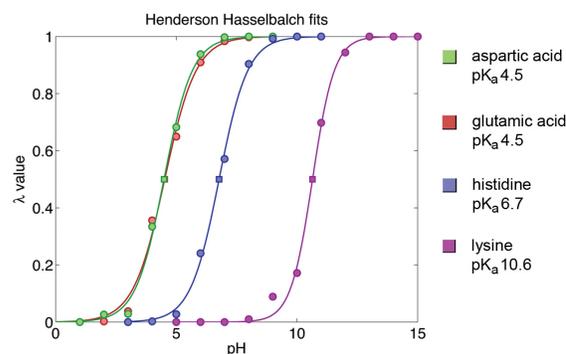


Figure 4. The pK_a calculations for 4 single residues: aspartic acid, glutamic acid, histidine, and lysine. The protonation state (dots) was calculated from the fraction of λ values in pure unprotonated and protonated states. The point of

inflection (boxes) of Henderson Hasselbalch equation fits (lines) indicates the calculated pK_a values. Even without optimizing for efficiency, convergence of data, or simulation parameters, we find the calculated pK_a values match those from the force field to within 0.5 pK units.

Multiple-residue Systems

The end purpose for CPHMD is to enable the study of complex pH-coupled phenomena of biological systems, such as pH-dependent protein conformation and cooperative titration effects among neighboring residues. As such, we test the accuracy of the CUDA-CPHMD algorithm by its ability to recapitulate residue pK_a values from both experiments and previous replica exchange studies, as shown in Figure 5. We study 9 model protein systems here: barnase^{2,47,48} (PDB code 1A2P); the serine protease inhibitor CI-2 from barley seeds^{47,49} (PDB code 2CI2); the hyperstable variant of staphylococcal nuclease, Δ +PHS^{3,16} (PDB code 3BDC); hen egg white lysozyme^{47,50,51} (PDB code 1LSA); the N-terminal domain of ribosomal protein L9^{47,52} (PDB code 1CQU); turkey ovomucoid^{47,53,54} (PDB code 1OMU); ribonuclease A^{47,55} (PDB code 7RSA); ribonuclease H from *Escherichia coli*^{47,56,57} (PDB code 2RN2); and *Bacillus circulans*

xylanase^{47,58,59} (PDB code 1BCX). Each protein was simulated in 11 pH windows from the pH -1 to the pH 9. Within each window, the proteins were simulated for 80 ps in 10 independent trajectories, which resulted in a total of 4.4 ns of simulation time per structure. All titrating residues were allowed to change protonation state using the new CUDA-CPHMD algorithm (Figure 5a, 5d, and 5e) and the original CPHMD algorithm (Figure 5b and 5d); and salt concentrations were added using concentrations that corresponded to the experiments.^{16,47} The simulations were run using the CHARMM22 force field^{43,44} with the Langevin integrator⁴⁵ with an integration timestep of 2 fs. These were NT (constant particle number and temperature) simulations each in an unbounded volume at a temperature of 298K using a Langevin heat bath. Atomic radii were optimized through work by Chen et al.⁴⁶ Similarly to the single-residue simulations, pK_a values were calculated by fitting the Henderson-Hasselbalch equation of states to the average protonation state λ of each titrating residue. Again, the point of inflection of the fit corresponds to the pK_a value of that residue. We report these values in Figure 5.

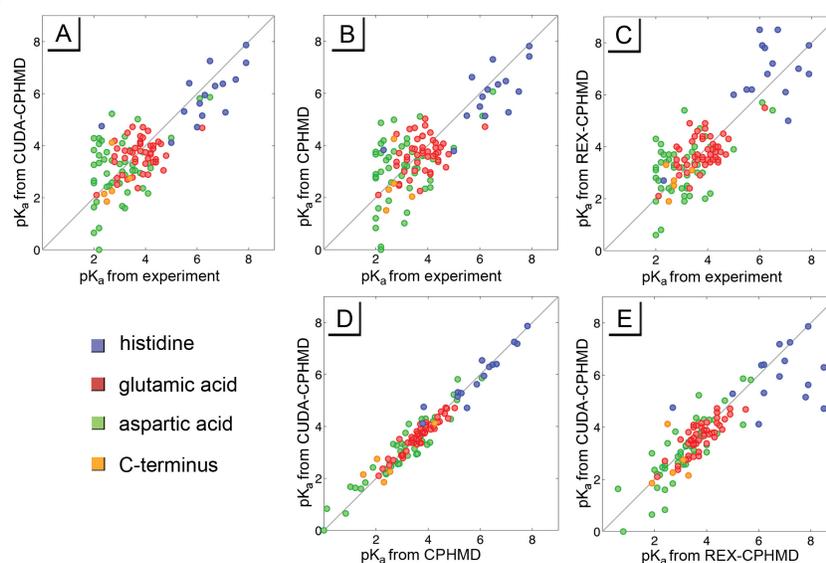


Figure 5. The pK_a calculations for all histidine (blue), glutamic acid (red), aspartic acid (green), and titrating C-terminus (orange) residues in all 9 of the test proteins. Each dot corresponds to a pK_a value resulting from fitting the Henderson Hasselbalch

equation to the fraction of λ values in pure unprotonated states. We present comparisons between pK_a values from A) CUDA-CPHMD and experiment; B) CPHMD and experiment; C) CPHMD with replica exchange (REX-CPHMD) and experiment; D) CUDA-CPHMD and CPHMD; and E) CUDA-CPHMD and REX-CPHMD. Even without optimizing the simulations to accommodate various titration equilibria for each protein, the CUDA-CPHMD algorithm successfully recapitulates experimental pK_a values to within 0.79 pK units of AUE. The experimental and REX-CPHMD results are from reference⁴⁷ and papers cited therein.

The AUE for all residues using CUDA-CPHMD was 0.79 pK units, which compares favorably to the AUE of 0.97 pK units using the null approximation (all pK_a values correspond to their reference values). This was the same 0.79 pK units of AUE that the original algorithm achieved, which further supports CUDA-CPHMD's accurately representing its CPU counterpart. Interestingly, while the average accuracy of CUDA-CPHMD and CPHMD were less than the 0.75 pK units of AUE achieved using the replica exchange methods from earlier studies, the non-replica-exchange pK_a calculations had a smaller standard deviation of error and fewer outlying predictions.^{16,47} Additional accuracy should be possible by coupling CUDA-CPHMD with the enhanced sampling of replica exchange in temperature or pH.^{16,47} This result holds great promise in establishing dynamic titration as a common feature of protein simulations.

Conclusions

In this study we present a significantly faster version of the CPHMD algorithm adapted for parallel processing in the CHARMM-OpenMM interface. While algorithmically the new CUDA-CPHMD algorithm represents little change over its predecessor, the speed improvements are so great that previously-unreasonable simulations are now straightforward to perform. For instance, what may have been a year-long simulation of the HK97 head capsule can now be performed in about 160 minutes. With this newfound speed is an opportunity to fine-tune the CPHMD titration model for a variety of protein systems, and to explore the impact of pH environments on side-chain dynamics both at the microsecond timescale and with all-atom detail.

Similarly to GBSW, the CPHMD model carries with it over a decade of research and parameterization.^{26,47,52} One model of particular interest is pH replica exchange (REX),⁶⁰ which has been shown to predict pK_a values of protein structures within single nanoseconds of simulation time.^{60,61} Coupled with the improved speed of CPHMD, adapting REX will enable a useful and rapid method for characterizing the chemical environment of protein interiors.

Acknowledgements

The authors gratefully acknowledge support from the NIH through grants GM103695 and GM107233 to CLB.

Keywords: implicit solvation, solvation, solvent model, CUDA, parallelization

References

1. Nielsen, J. E.; McCammon, J. A. *Protein Science : A Publication of the Protein Society* 2003, 12(9), 1894-1901.
2. Sali, D.; Bycroft, M.; Fersht, A. R. *Nature* 1988, 335(6192), 740-743.
3. Cannon, B.; Isom, D.; Robinson, A.; Seedorff, J.; Garcia-Moreno, B. *Biophysical Journal* 2007, 403A-403A.
4. Rabbani, G.; Ahmad, E.; Zaidi, N.; Fatima, S.; Khan, R. *Cell Biochemistry and Biophysics* 2012, 62(3), 487-499.
5. Wagner, G. R.; Payne, R. M. *Journal of Biological Chemistry* 2013, 288(40), 29036-29045.
6. Baker, L. A.; Watt, I. N.; Runswick, M. J.; Walker, J. E.; Rubinstein, J. L. *Proceedings of the National Academy of Sciences* 2012, 109(29), 11675-11680.
7. Cain, B. D.; Simoni, R. D. *Journal of Biological Chemistry* 1986, 261(22), 10043-10050.
8. Rastogi, V. K.; Girvin, M. E. *Nature* 1999, 402(6759), 263-268.
9. Clippingdale, A. B.; Wade, J. D.; Barrow, C. J. *Journal of Peptide Science* 2001, 7(5), 227-249.

10. Dobson, C. M. *Nature* 2003, 426(6968), 884-890.
11. Cuello, L. G.; Cortes, D. M.; Jogini, V.; Somporsit, A.; Perozo, E. *FEBS letters* 2010, 584(6), 1126-1132.
12. Howell, E. E.; Villafranca, J. E.; Warren, M. S.; Oatley, S. J.; Kraut, J. *Science* 1986, 231(4742), 1123-1128.
13. Aguilar, B.; Anandakrishnan, R.; Ruscio, J. Z.; Onufriev, A. V. *Biophysical journal* 2010, 98(5), 872-880.
14. Lindorff-Larsen, K.; Piana, S.; Dror, R. O.; Shaw, D. E. *Science* 2011, 334(6055), 517-520.
15. Ahlstrom, L. S.; Law, S. M.; Dickson, A.; Brooks III, C. L. *Journal of Molecular Biology* 2015, 427(8), 1670-1680.
16. Arthur, E. J.; Yesselman, J. D.; Brooks III, C. L. *Proteins: Structure, Function, and Bioinformatics* 2011, 79(12), 3276-3286.
17. Bashford, D.; Case, D. A.; Dalvit, C.; Tennant, L.; Wright, P. E. *Biochemistry* 1993, 32(31), 8045-8056.
18. Mertz, J. E.; Pettitt, B. M. *International Journal of Supercomputer Applications and High Performance Computing* 1994, 8(1), 47-53.
19. Sham, Y. Y.; Chu, Z. T.; Warshel, A. *Journal of Physical Chemistry B* 1997, 101(22), 4458-4472.
20. Baptista, A. M.; Martel, P. J.; Petersen, S. B. *Proteins-Structure Function and Genetics* 1997, 27(4), 523-544.
21. Baptista, A. M.; Teixeira, V. H.; Soares, C. M. *Journal of Chemical Physics* 2002, 117(9), 4184-4200.
22. Mongan, J. T.; Case, D. A.; McCammon, J. A. *Abstracts of Papers American Chemical Society* 2005, 229(Part 1), U768.
23. Kong, X. J.; Brooks III, C. L. *Journal of Chemical Physics* 1996, 105(6), 2414-2423.
24. Lee, M. S.; Salsbury, F. R.; Brooks III, C. L. *The Journal of Chemical Physics* 2002, 116(24), 10606-10614.
25. Im, W.; Lee, M. S.; Brooks III, C. L. *Journal of computational chemistry* 2003, 24(14), 1691-1702.
26. Khandogin, J.; Brooks, C. L., III. *Biophysical Journal* 2005, 89(1), 141-157.
27. Zeng, X.; Mukhopadhyay, S.; Brooks, C. L., III. *Proceedings of the National Academy of Sciences* 2015, 112(7), 2034-2039.
28. May, E. R.; Arora, K.; Brooks, C. L. I. *Journal of the American Chemical Society* 2014, 136(8), 3097-3107.
29. Arthur, E. J.; Brooks, C. L., III. *Journal of computational chemistry* 2016, 37(10), 927-939.
30. Brooks, B. R.; Brooks III, C. L.; Mackerell, A. D.; Nilsson, L.; Petrella, R. J.; Roux, B.; Won, Y.; Archontis, G.; Bartels, C.; Boresch, S.; Caffisch, A.; Caves, L.; Cui, Q.; Dinner, A. R.; Feig, M.; Fischer, S.; Gao, J.; Hodoscek, M.; Im, W.; Kuczera, K.; Lazaridis, T.; Ma, J.; Ovchinnikov, V.; Paci, E.; Pastor, R. W.; Post, C. B.; Pu, J. Z.; Schaefer, M.; Tidor, B.; Venable, R. M.; Woodcock, H. L.; Wu, X.; Yang, W.; York, D. M.; Karplus, M. *Journal of computational chemistry* 2009, 30(10), 1545-1614.
31. Case, D. A.; Cheatham, T. E.; Darden, T.; Gohlke, H.; Luo, R.; Merz, K. M.; Onufriev, A.; Simmerling, C.; Wang, B.; Woods, R. J. *Journal of computational chemistry* 2005, 26(16), 1668-1688.
32. Eastman, P.; Friedrichs, M. S.; Chodera, J. D.; Radmer, R. J.; Bruns, C. M.; Ku, J. P.; Beauchamp, K. A.; Lane, T. J.; Wang, L.-P.; Shukla, D.; Tye, T.; Houston, M.; Stich, T.; Klein, C.; Shirts, M. R.; Pande, V. S. *Journal of Chemical Theory and Computation* 2013, 9(1), 461-469.
33. Hess, B.; Kutzner, C.; van der Spoel, D.; Lindahl, E. *Journal of Chemical Theory and Computation* 2008, 4(3), 435-447.
34. Phillips, J. C.; Braun, R.; Wang, W.; Gumbart, J.; Tajkhorshid, E.; Villa, E.; Chipot, C.; Skeel, R. D.; Kalé, L.; Schulten, K. *Journal of computational chemistry* 2005, 26(16), 1781-1802.
35. Still, W. C.; Tempczyk, A.; Hawley, R. C.; Hendrickson, T. *Journal of the American Chemical Society* 1990, 112(16), 6127-6129.
36. Srinivasan, J.; Trevathan, M. W.; Beroza, P.; Case, D. A. *Theoretical Chemistry Accounts* 1999, 101(6), 426-434.
37. Donnini, S.; Tegeler, F.; Groenhof, G.; Grubmüller, H. *Journal of Chemical Theory and Computation* 2011, 7(6), 1962-1978.
38. Lee, M. S.; Salsbury, F. R.; Brooks III, C. L. *Proteins-Structure Function and Bioinformatics* 2004, 56(4), 738-752.
39. Uhlenbeck, G. E.; Ornstein, L. S. *Physical Review* 1930, 36(5), 0823-0841.
40. Singhal, A. K.; Chien, K. Y.; Wu, W. G.; Rule, G. S. *Biochemistry* 1993, 32(31), 8036-8044.
41. Castañeda, C. A.; Fitch, C. A.; Majumdar, A.; Khangulov, V.; Schlessman, J. L.; García-Moreno, B. E. *Proteins: Structure, Function, and Bioinformatics* 2009, 77(3), 570-588.

42. Gan, L.; Speir, J. A.; Conway, J. F.; Lander, G.; Cheng, N.; Firek, B. A.; Hendrix, R. W.; Duda, R. L.; Liljas, L.; Johnson, J. E. *Structure* 2006, 14(11), 1655-1665.
43. MacKerell, A. D.; Bashford, D.; Bellott, M.; Dunbrack, R. L.; Evanseck, J. D.; Field, M. J.; Fischer, S.; Gao, J.; Guo, H.; Ha, S.; Joseph-McCarthy, D.; Kuchnir, L.; Kuczera, K.; Lau, F. T. K.; Mattos, C.; Michnick, S.; Ngo, T.; Nguyen, D. T.; Prodhom, B.; Reiher, W. E.; Roux, B.; Schlenkrich, M.; Smith, J. C.; Stote, R.; Straub, J.; Watanabe, M.; Wiórkiewicz-Kuczera, J.; Yin, D.; Karplus, M. *The Journal of Physical Chemistry B* 1998, 102(18), 3586-3616.
44. MacKerell, A. D.; Feig, M.; Brooks III, C. L. *Journal of computational chemistry* 2004, 25(11), 1400-1415.
45. Verlet, L. *Physical Review* 1967, 159(1), 98-103.
46. Chen, J. H.; Im, W. P.; Brooks III, C. L. *Journal of the American Chemical Society* 2006, 128(11), 3728-3736.
47. Khandogin, J.; Brooks, C. L., III. *Biochemistry* 2006, 45(31), 9363-9373.
48. Oliveberg, M.; Arcus, V. L.; Fersht, A. R. *Biochemistry* 1995, 34(29), 9424-9433.
49. Tan, Y.-J.; Oliveberg, M.; Davis, B.; Fersht, A. R. *Journal of Molecular Biology* 1995, 254(5), 980-992.
50. Takahashi, T.; Nakamura, H.; Wada, A. *Biopolymers* 1992, 32(8), 897-909.
51. Bartik, K.; Redfield, C.; Dobson, C. M. *Biophysical Journal* 1994, 66(4), 1180-1184.
52. Kuhlman, B.; Luisi, D. L.; Young, P.; Raleigh, D. P. *Biochemistry* 1999, 38(15), 4896-4903.
53. Schaller, W.; Robertson, A. D. *Biochemistry* 1995, 34(14), 4714-4723.
54. Forsyth, W. R.; Gilson, M. K.; Antosiewicz, J.; Jaren, O. R.; Robertson, A. D. *Biochemistry* 1998, 37(24), 8643-8652.
55. Baker, W. R.; Kintanar, A. *Archives of Biochemistry and Biophysics* 1996, 327(1), 189-199.
56. Oda, Y.; Yoshida, M.; Kanaya, S. *Journal of Biological Chemistry* 1993, 268(1), 88-92.
57. Oda, Y.; Yamazaki, T.; Nagayama, K.; Kanaya, S.; Kuroda, Y.; Nakamura, H. *Biochemistry* 1994, 33(17), 5275-5284.
58. Plesniak, L. A.; Connelly, G. P.; McIntosh, L. P.; Wakarchuk, W. W. *Protein Science* 1996, 5(11), 2319-2328.
59. Joshi, M. D.; Hedberg, A.; McIntosh, L. P. *Protein Science* 1997, 6(12), 2667-2670.
60. Sabri Dashti, D.; Meng, Y.; Roitberg, A. E. *The Journal of Physical Chemistry B* 2012, 116(30), 8805-8811.
61. Nymeyer, H.; Gnanakaran, S.; García, A. E. In *Methods in Enzymology*; Ludwig, B.; Michael, L. J., Eds.; Academic Press, 2004, p 119-149.