JCTC Journal of Chemical Theory and Computation

Article

pubs.acs.org/JCTC

ff14ipq: A Self-Consistent Force Field for Condensed-Phase Simulations of Proteins

David S. Cerutti,* William C. Swope, Julia E. Rice, and David A. Case*

Department of Chemistry and Chemical Biology and BioMaPS Institute, Rutgers University, 610 Taylor Road, Piscataway, New Jersey 08854-8066, United States

Supporting Information

ABSTRACT: We present the ff14ipq force field, implementing the previously published IPolQ charge set for simulations of complete proteins. Minor modifications to the charge derivation scheme and van der Waals interactions between polar atoms are introduced. Torsion parameters are developed through a generational learning approach, based on gas-phase MP2/cc-pVTZ single-point energies computed of structures optimized by the force field itself rather than the quantum benchmark. In this manner, we sacrifice information about the true quantum minima in order to ensure that the force field maintains optimal agreement with the MP2/cc-pVTZ benchmark for the ensembles it will actually produce in simulations. A means of making the gas-phase torsion parameters compatible with solution-phase IPolQ charges is presented. The ff14ipq model is an alternative to ff99SB and other Amber force fields for protein simulations in programs that accommodate pair-specific Lennard−Jones combining rules. The force field gives strong performance on α-helical and β-sheet oligopeptides as well as globular proteins over microsecond time scale simulations, although it has not yet been tested in conjunction with lipid and nucleic acid models. We show how our choices in parameter development influence the resulting force field and how other choices that may have appeared reasonable would actually have led to poorer results. The tools we developed may also aid in the development of future fixed-charge and even polarizable biomolecular force fields.

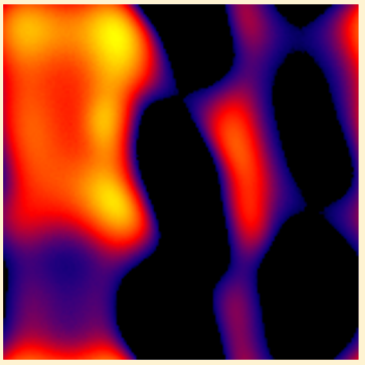

1. INTRODUCTION

Molecular simulations of biomolecules are expressions of their underlying force fields, sampling the interactions of chemically bonded atoms through Newtonian approximations of the quantum systems. The typical model, given in eq 1, is well-known. It approximates the energy of a molecular system as a thoroughly decomposable, easily differentiable sum of terms involving harmonic bonds, harmonic angles, Lennard−Jones repulsion/dispersion terms, electrostatic interactions, and additional parameters guiding the dihedral preferences of the model.[1] The "additional parameters" are both the least physically grounded and the most frequently edited part of most force fields; numerous forms of the expressions modifying the dihedral potential energy include screening factors applied between the electrostatic and Lennard−Jones interactions of atoms connected by chains of three bonds, Fourier series in the dihedral angles made by such atoms, and coupling terms between Fourier series of consecutive dihedral atom chains.[2,3] Changes in these parameters that have made considerable improvements in numerous force fields[4−8] should be viewed in light of the fact that their adjusments to the system energy are small. The terms describe a regime between the high-frequency motions of bonded atoms and the low-frequency rearrangements of nonbonded chemical groups and appear to be a sort of keystone in numerous molecular models.

$$U = \sum_b \frac{1}{2}k_b \| \mathbf{r}_{b,0} - \mathbf{r}_b \|^2 + \sum_a \frac{1}{2}k_a \| \eta_{a,0} - \eta_a \|^2 + \sum_h k_h[1 + \cos(n_h\theta_h + \phi_h)] + \sum_{ij}\left[\left(\frac{\sigma_{ij}}{\| \mathbf{r}_{ij} \|}\right)^{12} - \left(\frac{\sigma_{ij}}{\| \mathbf{r}_{ij} \|}\right)^{6}\right] + \sum_{ij} \frac{q_i q_j}{4\pi\varepsilon_0 \| \mathbf{r}_{ij} \|} \quad (1)$$

Many molecular force fields can be classified into several lineages.[9] While some force fields have been developed to reproduce bulk liquid properties, new models today tend to emerge as new levels of quantum mechanical theory become accessible and also as updates based on more extensive fitting data or inconsistencies with known biochemical data. Within each lineage, the models tend to evolve while retaining a significant portion of their parameters from previous work. Even comparatively minor changes can take years to gain acceptance, however, and any jump to parameter development based on a new level of quantum theory also requires significant effort to validate. The lineages of force fields are a natural consequence of the economics of force field development, particularly in the academic community where novelty and publications are essential. This environment has also driven the evolution of

Received: July 20, 2014
Published: September 18, 2014

   dx.doi.org/10.1021/ct500643c | J. Chem. Theory Comput. 2014, 10, 4515−4534

composite models where the various energy terms are shaped by different levels of quantum theory.

The charge model developed in 1995[10] by Cornell et al., incorporated first into the Amber ff94 force field, has remained in service for nearly 20 years. These charges were developed according to the Kollmann REsP method,[11] using HF/6-31G* quantum calculations to provide the target electrostatic potentials. Numerous other force fields have adopted the charge model, most of them distinguished by new dihedral Fourier series,[7,12] typically based on quantum calculations at the MP2 level with either a cc-pvTZ or 6-31++G basis set.[13−17] While the mismatch in quantum targets implies that the torsion terms are effectively modifying the short-ranged behavior of the charge set to fit a different potential energy surface, the force fields derived in this way have led to impressively stable protein behavior. In 2003, Duan et al. derived an alternative charge model, based on B3LYP/cc-pvTZ quantum calculations in a polarizable continuum (PCM) solvent intended to mimic the interior of a protein.[18] Backbone torsion Fourier series were derived specifically for this new charge set at the MP2/cc-pvTZ level of theory, also in the context of PCM solvent, to complete the Amber ff03 force field. While the charges and single-point energies are derived with similar styles of QM theory, the standard torsion fitting procedure does not incorporate its own PCM solvent, and it is not trivial to extract the energy of the PCM-polarized wave functions in vacuum. Molecular mechanics gas-phase energies computed with charges derived in the context of PCM solvent have accordingly been shown to double-count polarization effects, as the statically polarized charges are likely to be further affected by direct interaction with the solvent.[19] It is unclear how this imbalance may have affected the quality of simulations performed with ff03, but the force field has not become as widely used or further refined as ff94 and its derivatives.

In recent work, we devised the implicitly polarized charge model (IPolQ),[20] based on a new quantum mechanical method that integrates the condensed-phase environment due to explicit water molecules into quantum calculations at the MP2 level of theory. In this method, the appropriate atomic partial charges of a nonpolarizable model are estimated to be halfway between the charges of the system in vacuum and those of a system fully polarized in the presence of a condensed-phase reaction field potential.[21] The averaging accounts for the polarization energy and instantaneous rearrangements of dipoles in otherwise nonpolarizable models. The choice of TIP4P-Ew[22] to represent the solution-phase environment while developing the protein force field supports the use of this water model for later simulations. Furthermore, the IPolQ method provides a basis for understanding why essentially all fixed-charge water models carry a dipole moment of roughly 2.3 D: the dipole is halfway between the true condensed phase dipole of water, estimated to be 2.6 to 2.9 D,[23] and the 1.85 D dipole of water in the gas phase. Charges for solutes can therefore be derived under physical approximations that also obtain very similar results to existing water models. While the IPolQ approach offers numerous benefits for creating a self-consistent force field, the process of deriving charges is laborious, and our previous work did not assign torsion Fourier terms to complete the model. In this article, we will extend the IPolQ method to facilitate derivation of complete models, and we will also describe automation of the procedure, which greatly reduces the human effort needed to create new force fields for proteins and drug-like molecules.

2. THEORY

The Amber ff14ipq force field is intended to be a direct alternative to other nonpolarizable Amber force fields and contains no new functional forms. However, ff14ipq is also designed to consistently adhere to a specific level of quantum theory, MP2/cc-pvTZ. Over the course of this study, we found two major obstacles to fitting a molecular mechanics (MM) expression to a quantum mechanical (QM) potential energy surface (PES) and devised our own solutions to each of them. The first concerns the means for adapting the charge set and torsion Fourier series terms to work in concert. The second focuses on whether to make corrections for the inability of common MM models to reproduce certain high-energy features of the QM PES.

2.1. The Implicitly Polarized Charge Model and an Extension That Facilitates Derivation of a Complete Protein Force Field. The implicitly polarized charge method produces charges that, given certain assumptions detailed earlier,[20] reflect the appropriate partial atomic charges of a nonpolarizable model intended for simulations in aqueous solution. The atomic partial charges of our IPolQ amino acid charge set approximate electrostatic potentials around dipeptides; these potentials are halfway between the potential calculated for an unpolarized dipeptide (in numerous conformations) in vacuo and the potential calculated for the dipeptide (in the same conformations) after its electron density has been polarized by a solvent reaction field potential due to the time-averaged TIP4P-Ew water density sampled around each conformation. The electrostatic potentials needed for fitting charges in the condensed phase are laborious but straightforward to compute. It is less obvious how to extract the internal potential energy of the dipeptides in the condensed phase or make an equivalent averaging to derive an appropriate PES for fitting torsion parameters. Zgarbová and colleagues[19] solved this problem by computing quantum mechanical energies of solutes at the Hartree−Fock level of theory in the presence of the COSMO continuum solvent model, removing the solute−solvent interaction energy and then comparing to molecular mechanics potential energies computed with a PB solvent model. In effect, they compared the internal energies of the QM and MM systems in the presence of equivalent continuum solvents. In our case, it is not as straightforward to decompose single-point energies at the MP2 level of theory, nor is it as certain whether a quantum-mechanical implicit solvent model could be substituted for TIP4P-Ew. We instead chose to extend the IPolQ methodology and offer an alternative method for deriving torsion Fourier series from gas-phase PESs when the accompanying charge sets approximate polarization in a condensed-phase environment.

It would be simple to derive torsion parameters for a set of charges derived strictly from the electrostatic potentials for solutes in vacuum: compute the single-point energies of many additional solute conformations at the same level of quantum theory, again in vacuum, and fit the MM energies of each conformation to match the QM results. Furthermore, the role of the torsion Fourier series is to artificially correct errors in the nonbonded electrostatic and van der Waals interactions between atoms on either end of a dihedral; the PES of a rotatable bond is mostly captured by the nonbonded parameters. We therefore sought a way to express the set of IPolQ charges for a solute of interest, Q^{IPol}, as a perturbation ΔQ of charges Q^{vac} derived for the solutes using only the vacuum quantum calculations:

$$Q^{\mathrm{IPol}} = Q^{\mathrm{vac}} + \Delta Q \tag{2}$$

It is simple to compute torsion parameters to match an MM PES Q^{vac} to a QM PES computed in vacuum. Changes in the PES of a molecule's rotatable bond due to immersing it in water could then be assumed to be adequately represented by changes in local dipoles as expressed by ΔQ. This appears to be a safe assumption, given that $\Delta Q \ll Q^{\mathrm{vac}}$ and that the van der Waals parameters, which also strongly influence the potential energy changes due to rotation about a bond, are constant across both phases. However, the linear least-squares fit from which our IPolQ charges Q^{IPol} and our proposed vacuum charges Q^{vac} are derived is able to obtain many different solutions with similar levels of accuracy, particularly for buried atoms whose charges are not well determined. In our earlier work, we used large numbers of conformations of each solute in order to guard against this behavior. However, because residual indeterminacy could exist in either Q^{IPol} or Q^{vac}, the difference ΔQ could still be amplified. We therefore extended the IPolQ least-squares fitting procedure to fit both Q^{IPol} and Q^{vac} simultaneously, with additional restraint equations to keep ΔQ small. The results are solutions for each of Q^{IPol} and Q^{vac} related by a minimal, smooth perturbation ΔQ.

The original matrix equation constructed in the IPolQ fitting procedure may be expressed in shorthand:

$$\begin{bmatrix} A \\ V \end{bmatrix}[Q] \approx \begin{bmatrix} u \\ v_{\mathrm{T}} \end{bmatrix} \tag{3}$$

Here, A is the fitting matrix whose elements are derived from the molecular coordinates and the kernel of Coulomb's law. Each row of A describes one linear equation through which the partial charges of a molecule in some conformation create an electrostatic potential at a nearby point in space. The electrostatic potential found by the target quantum method is stored in the vector u. Each column of A solves for an independent charge variable; if multiple atoms of a molecule are constrained to have the same charges, then they contribute to the same column. Appended to A, the matrix V contains additional restraint equations that penalize the least-squares fit when particular charges stray from target values encoded in the vector v_{T} appended to u. A more detailed statement of the linear least-squares problem is as follows:

$$\begin{bmatrix} A_{(1,1)1} & \cdots & A_{(1,1)n} \\ \vdots & \ddots & \vdots \\ A_{(1,m)1} & \cdots & A_{(1,m)n} \\ \vdots & \ddots & \vdots \\ A_{(k,1)1} & \cdots & A_{(k,1)n} \\ \vdots & \ddots & \vdots \\ A_{(k,m)1} & \cdots & A_{(k,m)n} \\ V_{\mathrm{NC}}\delta_{(1,1)} & \cdots & V_{\mathrm{NC}}\delta_{(1,n)} \\ \vdots & \ddots & \vdots \\ V_{\mathrm{NC}}\delta_{(k,1)} & \cdots & V_{\mathrm{NC}}\delta_{(k,n)} \\ V_{\mathrm{SA},1} & 0 & \cdots \\ \vdots & \ddots & \vdots \\ 0 & \cdots & V_{\mathrm{SA},n} \end{bmatrix} \begin{bmatrix} Q_1^{\mathrm{IPol}} \\ \vdots \\ Q_n^{\mathrm{IPol}} \end{bmatrix} \approx \begin{bmatrix} u_{(1,1)} \\ \vdots \\ u_{(1,m)} \\ \vdots \\ u_{(k,1)} \\ \vdots \\ u_{(k,m)} \\ V_{\mathrm{NC}}q_1 \\ \vdots \\ V_{\mathrm{NC}}q_k \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{4}$$

Here, the elements of the fitting matrix $A_{(p,i)j}$ pertain to the influence of all instances of the jth adjustable charge on the ith fitting point around the pth molecular conformation. The solution vector u on the right-hand side is filled with values of the electrostatic potential measured at many points $\mathbf{r}_{p,i}$. As was mentioned, there may be multiple instances of the same adjustable charge in each molecule. A general definition of the elements of A is then

$$A_{(p,i)j} = \sum_c \left[\frac{1}{4\pi\varepsilon_0 \| \mathbf{r}_{(p,i)} - \mathbf{r}_c \|} \right] \tag{5}$$

where the summation runs over all atoms c bearing the fitted charge described by variable j. The various molecular conformations may be different poses of the same molecule or even a collection of poses of many different molecules. In the latter case (which best describes the original IPolQ fit we performed), the matrix A becomes a sparse matrix due to the fact that not all systems contain all charge variables, but it is not often sparse enough to warrant a special storage format. This can require a large amount of memory, but simultaneous solution of all charge variables permits multiple systems to share the same charges. In our IPolQ fit, for instance, we chose to follow the Cornell charge set convention of giving similar charges to backbone N, H, C, and O atoms involved in the peptide bonds for neutral, positively charged, and negatively charged amino acids, compressing what could have been more than 80 charge variables into just 12, in the hope that this would expedite development of a common set of backbone torsion potentials. The restraint matrix V is appended to the fitting matrix A in order to keep the charges of certain atoms small and also to restrain the overall charges of certain groups of atoms: the fit is heavily penalized by $V_{\mathrm{NC}} = 1.0 \times 10^5$ kcal/mol·e^2 if the sum of atomic partial charges of neutral or ionic residues differs from the appropriate net charge (NC), q_k for the kth system. In eq 5, $\delta_{p,j}$ is 1 if the pth system contains the jth charge variable. To put all restraint constants $V_{\mathrm{SA},j}$ for specific atomic charges on the same scale, these constants were set proportional to the number of times N_j the jth charge variable appeared in A

$$V_{\mathrm{SA},j} = V_{\mathrm{SA}}^0 N_j \tag{6}$$

In our IPolQ fit, we set V_{SA}^0 to 1.0×10^{-2} kcal/mol·e^2.

The extended least-squares fit for deriving IPolQ charges as a perturbation of vacuum charges solves for twice as many variables based on the same molecular conformations. The extended problem to solve is simply

$$\begin{bmatrix} A & 0 \\ A & A \\ V & 0 \\ 0 & V_{\mathrm{GP}}I \end{bmatrix} \begin{bmatrix} Q^{\mathrm{vac}} \\ \Delta Q \end{bmatrix} \approx \begin{bmatrix} u^{\mathrm{vac}} \\ u^{\mathrm{IPol}} \\ v_{\mathrm{T}} \\ 0 \end{bmatrix} \tag{7}$$

In this system of equations, the original fitting matrix is replicated into several blocks of the extended system, and the solution vector contains electrostatic potentials computed from each molecular conformation in vacuum as well as the vacuum potential averaged with the electrostatic potential computed in the condensed phase u^{IPol}. The original restraint equations are still present, but they apply only to the vacuum charges Q^{vac}; additional restraint equations are added to strongly force the sum of all perturbation charges in each system to zero. An additional block matrix loosely restrains perturbation charges individually

dx.doi.org/10.1021/ct500643c | J. Chem. Theory Comput. 2014, 10, 4515−4534

Table 1. Differences in the Molecular Mechanics (MM) Energy Components for 628 Conformations of Various Dipeptides after Optimization by MM or MP2/6-31++G[a]

dipeptide	bond	angle	1−4 LJ	1−4 elec	other LJ	other elec	bonded sum[b]	nonbonded sum[c]
Ash	2.99 ± 0.34	0.66 ± 0.98	0.96 ± 0.17	2.20 ± 1.25	0.28 ± 0.41	−1.88 ± 1.00	3.65 ± 1.15	1.57 ± 1.06
Asn	2.08 ± 0.45	0.56 ± 0.90	0.93 ± 0.15	4.16 ± 1.28	0.24 ± 0.40	−2.88 ± 1.25	2.64 ± 2.06	2.45 ± 1.01
Asp	1.92 ± 0.38	1.08 ± 0.55	0.88 ± 0.16	2.60 ± 1.01	0.58 ± 0.34	−1.98 ± 1.16	3.00 ± 0.65	2.08 ± 0.48
Cys	1.63 ± 0.24	1.46 ± 0.77	0.46 ± 0.26	2.71 ± 0.89	0.43 ± 0.14	−2.45 ± 0.80	3.09 ± 0.83	1.15 ± 0.44
Hid	3.21 ± 0.33	3.02 ± 1.03	0.83 ± 0.17	2.27 ± 0.81	0.24 ± 0.31	−2.06 ± 0.75	6.23 ± 1.05	1.28 ± 0.53
Hie	2.93 ± 0.26	2.62 ± 1.04	0.92 ± 0.17	2.71 ± 0.79	0.34 ± 0.31	−2.32 ± 0.75	5.55 ± 1.05	1.65 ± 0.47
Hip	3.17 ± 0.45	2.99 ± 1.17	1.10 ± 0.16	3.43 ± 0.95	0.70 ± 0.46	−2.61 ± 0.90	6.16 ± 1.42	2.61 ± 1.06
Ile	1.36 ± 0.29	0.89 ± 0.61	0.89 ± 0.17	2.20 ± 0.62	0.62 ± 0.58	−2.04 ± 0.49	2.25 ± 0.70	1.67 ± 0.72
Leu	1.47 ± 0.30	1.13 ± 0.59	0.64 ± 0.14	3.23 ± 0.89	0.31 ± 0.45	−2.78 ± 0.73	2.60 ± 0.70	1.40 ± 0.56
Phe	1.77 ± 0.26	0.92 ± 0.69	1.73 ± 0.18	2.34 ± 0.76	0.45 ± 0.52	−2.21 ± 0.65	2.69 ± 0.76	2.32 ± 0.62
Ser	1.59 ± 0.26	0.48 ± 0.33	0.95 ± 0.36	2.04 ± 0.73	0.49 ± 0.09	−1.83 ± 0.67	2.07 ± 0.42	1.64 ± 0.49
Thr	1.59 ± 0.27	0.46 ± 0.33	0.99 ± 0.42	2.84 ± 0.99	0.47 ± 0.12	−2.30 ± 0.86	2.05 ± 0.43	2.00 ± 0.63
Trp	2.39 ± 0.28	2.85 ± 1.02	1.64 ± 0.19	2.64 ± 0.78	0.48 ± 0.80	−2.31 ± 0.79	5.24 ± 1.07	2.44 ± 0.81
Tyr	2.16 ± 0.26	1.36 ± 0.70	1.88 ± 0.18	2.37 ± 0.76	0.43 ± 0.59	−2.39 ± 0.66	3.52 ± 0.79	2.29 ± 0.70
Val	1.42 ± 0.25	0.75 ± 0.25	0.70 ± 0.13	2.71 ± 0.54	0.40 ± 0.12	−2.32 ± 0.46	2.17 ± 0.30	1.48 ± 0.28

[a]Each difference is given as an average ± standard deviation, in kcal/mol. [b]Includes bond and angle contributions. [c]Includes Lennard−Jones (LJ) and electrostatic (elec) contributions.

to zero using stiffness constants V_{GP} scaled by the number of instances of each (perturbation) charge variable N_j; numerical values of V_{GP} are discussed below. In eq 7, I is the identity matrix; every perturbation charge is individually restrained. The extended system, which requires slightly more than four times the memory of the original problem and eight times the computation cost, yields Q^{IPol} by applying eq 2.

2.2. The Traditional Torsion Fitting Approach. The traditional least-squares approach for fitting torsion parameters is composed as shown in eq 8. The fitting matrix is composed of the kernels $T_{(p,i)j}$ of the jth torsion term to be fitted for the ith conformation of the pth system, while other columns store energy adjustment constants C_p applied to the pth chemical system present in the fitting data. The solution vector u in this equation contains the QM single-point energies of each conformation, less the average single-point energy of all systems with the same chemical composition, less the MM energy arising from other terms in the force field. The torsion terms are fitted so that the sum of their contributions compensates for inaccuracies in the molecular mechanics model to bring the relative energies estimated for different conformations of the same molecule into agreement with the QM PES.

$$\begin{bmatrix} T_{(1,1)1} & \cdots & T_{(1,1)n} & 1 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ T_{(1,m(1))1} & \cdots & T_{(1,m(1))n} & 1 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ T_{(k,1)1} & \cdots & T_{(k,1)n} & 0 & \cdots & 1 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ T_{(k,m(k))1} & \cdots & T_{(k,m(k))n} & 0 & \cdots & 1 \\ V_{NC,1} & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & V_{NC,n} & 0 & \cdots & 0 \end{bmatrix} \begin{bmatrix} k_1 \\ \vdots \\ k_n \\ C_1 \\ \vdots \\ C_k \end{bmatrix} \approx \begin{bmatrix} u_{(1,1)} \\ \vdots \\ u_{(1,m)} \\ \vdots \\ u_{(k,1)} \\ \vdots \\ u_{(k,m)} \\ 0 \\ \vdots \\ 0 \end{bmatrix} \quad (8)$$

In principle, it would be feasible to leave the average single-point energy of each chemical system in u, but this would lead to a poorly conditioned matrix, as the energy adjustment constants C_{adj}, here $C_1 \ldots C_k$, would need to take on very large values. It is also possible to compose u based on system conformational energies relative to the structure with the lowest quantum mechanical energy and to omit the constants C_{adj} for each system. However, no quantum method is perfect, and omission of C_{adj} would amplify the influence of a particular conformation over the fitted parameters.

2.3. On the Inconsistency of Quantum Mechanical and Molecular Mechanical Potential Energy Surfaces in High-Frequency Degrees of Freedom. It is well-known to force field developers that the bond and angle terms that approximate high-frequency motions in molecular simulations are not consistent with the QM target models; aside from the breakdown of the harmonic approximation with increasing strain, the equilibrium bond length or angle is dependent on the chemical context. The inconsistency elevates the MM energies for structures optimized by QM methods and is therefore a sort of contaminant when fitting lower-frequency degrees of freedom, in particular torsion Fourier series terms. Concern about introducing error by this method has led numerous groups to fit torsion Fourier series by positing that a "conformation" of a molecule is defined strictly by its torsional degrees of freedom and making two slightly different variants of that conformation: relaxing all bond, angle, and nonbonded degrees of freedom by MM and QM approximations, respectively. The objective then becomes to fit torsion parameters such that the coordinates optimized by MM produce an energy most like the single-point energy found for the nuclear coordinates optimized according to QM. This method, hereafter the tandem optimization approach, relaxes much of the inconsistency arising from high-frequency degrees of freedom in the two PESs, but it also introduces a new source of error in the nonbonded interactions between the different coordinate sets. While many investigators have accepted the trade-off, we performed an independent analysis, which led us to reject the approach in favor of a direct, one-to-one mapping between coordinates and energies.

We analyzed the trade-off between error removed by relaxing high-frequency degrees of freedom and error introduced by changes in nonbonded interactions by computing MM energies for MM and QM optimized variants of 648 conformations in each of 15 amino acid dipeptides (over 9000 pairs of energies).

dx.doi.org/10.1021/ct500643c | J. Chem. Theory Comput. 2014, 10, 4515−4534

The conformations sampled χ_1 and χ_2 at 20° intervals while the backbone was held in an α-helical or β-sheet conformation; χ_1,χ_2, ϕ, and ψ dihedrals were held fixed during both optimizations. MM optimizations and all MM energy evaluations were performed with a new variant of the Amber ff99SB force field provided to us courtesy of James Maier and Professor Carlos Simmerling's research group. QM optimizations, also performed for us by James Maier, were performed at the MP2 level with the 6-311++G basis set.[24] The MM and QM approximations are similar to those we have chosen for development of Amber ff14ipq.

First, we focused on the difference in the molecular mechanics energy computed for each variant of a given conformation; by construction, the MM optimized variant was always lower in energy, as scored by Amber ff99SB. Table 1 gives the differences computed for 628 conformations of each dipeptide. By decomposing the MM energy, it is apparent that the MM and QM models do disagree about the optimal bond lengths and angles, and the energy differences in there terms are greater, sometimes much greater, than the energy differences arising from nonbonded terms. At face value, this result would appear to support the tandem optimization approach. However, a closer inspection weakens that conclusion.

When fitting torsion Fourier series to match a molecular mechanics PES to a QM PES, the objective is correct relative energies; the MM PES is otherwise very far removed from the QM PES, which includes factors such as electron−nuclei interactions and nuclear repulsion. The mean energy of the QM PES is therefore subtracted, and an additional constant C_{adj} is included in the fit to arbitrarily adjust the energy of a particular molecule, regardless of conformation, up or down, to bring the two PESs into agreement. Because of this, the mean energy differences between the MM and QM optimized variants of each conformation will "fall through" the fit, absorbed into C_{adj}. If the molecular mechanics and quantum models optimize bonds to different lengths, but the disagreement is consistent across all conformations of the molecule, then the mismatch in these high-frequency terms will have no effect on the fitted torsion parameters. Standard deviations of the difference, the variability of the disagreement between the QM and MM PES, is a much better indicator of possible contamination in a parameter fit. Table 1 shows that bond energy differences have a small deviation but angle energy differences carry a much higher deviation and therefore might contaminate the torsion parameter fit. Comparing the standard deviations in the energy differences arising from the sum of bond and angle terms to those arising from all nonbonded terms suggests that the tandem optimization approach offers a marginal benefit in most cases but is detrimental in cases such as Ser and Thr.

The total nonbonded energy differences appear to be smaller than those of bonds, but the nonbonded energy is a sum of many interactions. The major Amber molecular dynamics engines conveniently break nonbonded interactions into "1−4" contributions between atoms at either end of a dihedral group and "all other" contributions. These two parts of the Lennard−Jones interactions are uncorrelated, but the sum of electrostatic 1−4 interactions tends to be strongly anticorrelated with the sum of all other electrostatic contributions, as shown in Table 2. As a consequence, the standard deviations of the differences in the electrostatic 1−4" nonbonded terms are as large or larger than the deviations in total nonbonded energy differences. (The deviations only get larger by folding in the Lennard−Jones short-ranged interactions.) These 1−4 interactions form the base of the torsional PES around each rotatable bond and, of all nonbonded interactions, are the most strongly connected to the torsion potentials. This suggests that the contaminant introduced by the tandem optimization method, a mismatch in the nonbonded interactions, is actually worse than the contaminant being removed, the mismatch in energies due to high-frequency modes.

Table 2. Correlations between 1−4 Nonbonded Terms and All Other Nonbonded Terms in 628 Conformations of Many Dipeptide Systems[a]

dipeptide	Lennard−Jones	electrostatic
Ash	0.34	−0.67
Asn	0.20	−0.74
Asp	0.30	−0.96
Cys	−0.12	−0.91
Hid	−0.37	−0.77
Hie	−0.06	−0.91
Hip	0.12	−0.37
Ile	0.36	−0.99
Leu	−0.19	−0.99
Phe	0.08	−0.97
Ser	0.40	−0.94
Thr	0.46	−0.95
Trp	0.12	−0.88
Tyr	0.10	−0.95
Val	0.53	−0.98

[a]Like Table 1, this table compares MM energies computed for each of two optimized variants of a dipeptide conformation.

We conclude that it is at least as sane to accept errors that might be introduced by a mismatch between molecular mechanics bond and angle geometries and those of the quantum target as to accept errors arising from a mismatch in electrostatic or steric interactions. For this reason, we chose to submit the MM optimized coordinates, which reflect the states that the molecular mechanics model will actually explore in simulations, to single-point quantum calculations and demand that our fitted molecular mechanics model reproduce the QM single-point energies calculated for precisely the same set of nuclear positions; details of our procedure for generating the actual fitting data for torsion Fourier series can be found in the Methods. Due to the one-to-one mapping of coordinates and energies, the fitting data suffers some contamination from angle terms in the MM approximation being inadequate to describe the QM PES. These terms are roughly 1 order of magnitude higher in energy than the torsion terms or nonbonded interactions in the MM model; as such, they could make erroneous contributions to the energy with the risk of torsion parameters becoming fitted against noise. However, the majority of angle strain is orthogonal to the torsional subspace of molecular motions. With adequate sampling of the torsional degrees of freedom in the fitting set, and with a consistent optimization of each set of coordinates relative to one model or the other, the risk is limited.

3. METHODS

We designed the Amber ff14ipq force field to be similar to previous Amber nonpolarizable force fields. Bond and angle parameters were taken from the existing ff99SB force field,[7] along with most Lennard−Jones parameters. Atomic partial charges were refitted according to the updated IPolQ fitting procedure described in the Theory section, using the same data produced in our previous study,[20] but for future force field development, we automated the IPolQ fitting cycle in the Amber mdgx program.[1]

dx.doi.org/10.1021/ct500643c | J. Chem. Theory Comput. 2014, 10, 4515−4534

In our derivation of IPolQ charges for amino acid side-chain analogues, we adjusted the Lennard−Jones σ parameters of polar atoms in order to bring the computed hydration free energies into agreement with experiment. In this respect, the development philosophy of ff14ipq resembles that of the new Gromos 54A8 force field.[25] However, the interactions of these atoms with water were essentially the only changes that were of consequence during our hydration free energy calculations, and we found that larger σ radii fitted for most of the atom types intensified 1−4 nonbonded repulsion, making torsion parameters much more difficult to fit. To bring the Lennard−Jones changes into ff14ipq, we made them applicable only between the polar atom types and the TIP4P-Ew oxygen type. The Lennard−Jones Lorentz−Berthelot mixing rule is broken for the interactions of these atom types with water, in the manner that CHARMM36 makes use of NBFIX terms.[5] With the bonded and nonbonded parameters established, the majority of the development in ff14 focused on torsion Fourier series terms.

3.1. Automation of the IPolQ Charge Derivation. A new module was added to the mdgx program to aid users in computing IPolQ charges for arbitrary molecules. For a standard REsP procedure, researchers must have a set of conformations of their molecule of interest; the conformations serve as inputs to tools such as the RED server, which manages the necessary quantum calculations and performs the restrained charge fit.[26,27,11] With the new IPolQ module, researchers must have a set of conformations of their molecule immersed in the solvent of interest. The mdgx program will read the solvated conformations as restart files along with an appropriate topology and begin dynamics with the solute molecule held in a fixed position. The mdgx program automates the process of collecting the solvent charge density, computing the solvent reaction field potential (SRFP), preparing inputs to a quantum program, and launching the calculations to ultimately produce grids of electrostatic potential computed for the molecule in vacuum and in the influence of the SRFP. The mdgx IPolQ module improves on the original protocol, always applying a shell of point charges around the average solvent charge density taken from the simulation. The charges in the shell are fitted to reproduce the SRFP due to infinite electrostatics present in the simulation in the context of a quantum calculation on an isolated system. Furthermore, users can specify up to three concentric charge shells to increase the accuracy of the SRFP, evaluate the SRFP at additional sites throughout the solute volume, and even specify an interior shell of charges to reproduce the SRFP in and around the solute without including point charges nearer than an arbitrary distance from solute atoms. The final feature could be useful for researchers who are concerned about QM basis functions adversely interacting with solvent charges, although we did not notice any effects on the IPolQ results for some systems we tested with MP2/cc-pvTZ calculations used in the original IPolQ protocol (data not shown). In summary, mdgx manages a cycle of the IPolQ procedure for a conformation of the molecule of interest; the output of independent runs on multiple solute conformations can be pooled and sent to the mdgx charge fitting module to derive a new charge set, which can then be used to update the solvated system's topology and start another round of calculations with the IPolQ module. In this manner, mdgx manages nearly all aspects of the IPolQ procedure until the solute charge model converges. The module supports both Gaussian[28] and Orca[29] quantum chemistry packages, performs dynamics with modest parallelism and efficient CPU execution, and launches quantum chemistry programs for parallel execution in accord with the molecular dynamics run to avoid wasting CPU cycles.

3.2. Rederivation of the IPolQ Charges. IPolQ charges were rederived using the expanded matrix method described in Theory. Data for the fitting matrix A was derived from the same quantum data as the original IPolQ charge set. The same protocol was followed for selecting fitting points from the electrostatic potentials evaluated around each conformation, but, due to memory constraints in the much larger matrix problem, only 3750 points per conformation were selected. (Our earlier work showed that anywhere from 3000 to 5000 points per conformation yielded convergent results in the fitted charges.) This method implies two charge sets, one valid in vacuo and the other in solution. While the complete release version of ff14ipq takes the charges valid in solution, $Q^{vac} + \Delta Q$ in eq 7, we also considered a variant based on Q^{vac} in some studies of systems in vacuo, hereafter named V-ff14 and not distributed for general use.

3.3. Generating the Fitting Data for Torsion Fourier Series Terms. The fitting set for Amber ff14ipq was created by molecular simulations and energy optimizations performed with the Amber ff99SB force field and new variants under development in the Simmerling group. All of these MM models are based on the Cornell charge set and have been edited over more than a decade. The most significant feature of the force fields, to us, was the similarity in form and some parameters to the proposed Amber ff14ipq force field. While it is impossible to completely sample the available conformational space in the fitting data, we chose the most recent Amber force fields to simulate each molecule in the hope that the these models would sample the unrestrained degrees of freedom in a manner that reveals their biases so that those biases could be eliminated in ff14ipq. Nevertheless, critical dihedral angles near the center of each molecule were restrained to various values as we populated the fitting data set; the exact details of the simulations, such as cutoff, time step, and solvent model, are therefore of secondary importance to the choice of protein-like molecules, the array of restraints, and the number of snapshots collected for single-point energy calculations.

The most common structures included in our fitting data set were blocked dipeptides, Ace-XX-Nme, where XX is an amino acid, although we also included Ace-Ala-Ala-Ala-Nme and Ace-Gly-Gly-Gly-Nme tetrapeptides to sample backbone ϕ and ψ angles and Ace-XX-XX-Nme tripeptides to further reduce the influence of the blocking groups on these backbone dihedrals. In all, the torsion fitting data set consisted of nearly 28 000 structures whose single-point energies were computed in vacuum by MP2/cc-pvTZ calculations.

To sample side chain conformations, we obtained a set of some 17 000 dipeptide conformations for neutral Asp, Asn, deprotonated Asp, Cys, δ-protonated His, ε-protonated His, ionic His, Ile, Leu, Phe, Ser, Thr, Trp, Tyr, and Val from James Maier and the Simmerling group. Each amino acid was sampled in 628 conformations, exploring χ_1 and χ_2 angles at 20° intervals while holding the backbone in either an α-helical or β-sheet conformation. These amino acid conformations were among those used in our analysis of the effects of unrelaxed high-frequency degrees of freedom described in the Theory section. To obtain a better sampling of backbone conformations and to reduce the risk of side chain dihedrals becoming coupled to particular backbone conformations, we generated another 180 conformations of each dipeptide by sampling ϕ and ψ individually at 20° intervals with no restraints on any other

degrees of freedom. To decorrelate the unrestrained degrees of freedom, we ran molecular dynamics simulations of each dipeptide while progressively advancing the ϕ and ψ restraints. Dynamics were run in baths of TIP4P-Ew water at 300 K, with 100 ps intervals between each snapshot. ϕ or ψ were incremented with every snapshot, but each angle was completely rotated five times; as a result, separate snapshots at the same ϕ or ψ coordinate were spaced by 1.8 ns of dynamics.

To sample backbone conformations, reduce the dependence of the fitted parameters on Ace and Nme blocking groups, and increase side-chain sampling, we included a set of over 7000 tripeptides. We sampled all 20 amino acids with Ala, Gly, or Ser adjacent in either direction in the peptide chain, restraining ϕ and ψ individually at 15° intervals with no other restraints and running dynamics on each system at 450 K to generate multiple conformations at each value of ϕ or ψ. Because there were so many of these systems, a generalized Born solvent[30] was used to generate the conformations. Dynamics were run for 300 ps between snapshots selected for subsequent MP2 calculations.

To further sample backbone conformations of two key amino acids, alanine and glycine tetrapeptides were sampled in 1296 conformations each, sampling the central residue's ϕ and ψ space at 10° intervals on a two-dimensional grid. Aside from simultaneous restraints on ϕ and ψ, no other restraints were used as dynamics were performed for 100 ps between snapshots to decorrelate other degrees of freedom.

Snapshots of molecular dynamics trajectories are not suitable for direct incorporation into quantum calculations because of the potential for bond and angle high-frequency energy terms to contaminate the fitting data, as explained in the Theory section. However, if the contributions from these terms are relaxed out with respect to either a MM or QM approximation, then the single-point QM energies calculated for each conformation then constitute an acceptable fitting set. We optimized each conformation with restraints on dihedrals involving four heavy atoms (but not hydrogens) in order to preserve the diversity of conformations created in our simulations while relaxing as much angle and bond strain as possible. We chose to optimize conformations with respect to the MM approximation because the calculations were cheap, allowing us to devote more time to evaluating single-point energies. Also, we felt that it was more instructive to evaluate the local minima of an MM model with respect to the MP2/cc-pvTZ target, to examine states that an MM model might have a propensity to populate and to verify their energies with QM in the fitting data.

3.4. Fitting Torsion Fourier Series Terms. All torsion parameters for the Amber ff14ipq force field were fitted simultaneously. While the length of each Fourier series and also the phase angles were taken from ff99SB, new atom types added by the Simmerling group as well as new atom types required by the IPolQ charge set for amino acids were included, and the glycine C_α and proline backbone nitrogen atoms were given their own unique atom types distinct from other amino acids. Including a new atom type implied replicating all bond, angle, and Fourier series terms pertaining to the original type, thereby creating additional parameters to fit. All Fourier series amplitudes were reoptimized by the standard linear least-squares approach[7,31] in a single matrix equation to obtain the best overall fit for terms that appear in different contexts across multiple residues, as described in the Theory section. Another new module of the mdgx program was created to perform the atom type branching, set up the matrix equation, and perform the linear least-squares fit.

While it was an advantage to have the extensive fitting set described in the preceding section, all torsion Fourier terms were still restrained loosely toward zero by a penalty of 2.0×10^{-4} kcal/mol (multiplied by the number of times each parameter appeared in the fitting matrix) to keep the amplitudes small and avoid overfitting. Torsion amplitudes were optimized to make molecular mechanics energies of the di-, tri-, and tetrapeptide systems computed in vacuo with charges from Q^{vac} agree with MP2/cc-pvTZ single-point energies of the peptides also computed in vacuo. As explained in the Theory section, we assume these torsion parameters to be transferable to describe the behavior of solvated peptides when paired with Q^{IPol}.

3.5. Energy Minimization of Small Peptides in Vacuo for Iterative Refinement of the Force Field. As a preliminary test of force field performance, we evaluated fidelity to the underlying MP2/cc-pvTZ benchmark over the course of molecular mechanics energy minimization. The molecular conformations in the training set were already optimized with respect to a very similar MM model, but in the presence of one or more restraints on dihedral angles. Energy minimization of each conformation of all systems found in the fitting set was performed in vacuo with V-ff14, the variant of ff14ipq substituting Q^{vac} for the implicitly polarized charges found in the release version. New QM energy calculations were then performed on the resulting structures optimized by V-ff14. These new conformations and energies were added to the training set, and new energy optimizations were performed in an iterative fashion until V-ff14 could score structures created by its own optimization consistently with those in its training set.

3.6. Simulations of Peptides, Oligopeptides, and Proteins. Building on the fitting data of small peptides, we computed potentials of mean force (PMFs) for pairs of dihedral angles in blocked dipeptides at 298 K with a standard two-dimensional umbrella sampling technique. Octahedral boxes enclosed each dipeptide in roughly 1400 water molecules. We collected data in 1296 windows spaced by 10° in ϕ and ψ or in χ_1 and χ_2, depending on the amino acid and PMF. Each window was seeded from a continuous, incrementally restrained simulation similar to that used to generate conformations for the corresponding tetrapeptides prior to MP2/cc-pvTZ calculations. After seeding, each window was sampled for 4 ns following 0.5 ns equilibration. Dihedral angles were initially restrained by an 8 kcal/mol·rad^2 harmonic penalty function, but if some of the 10° bins were left undersampled, then more windows were added on a 5° grid with 16.0 kcal/mol·rad^2 dihedral restraints to completely fill out the PMF.

We performed long equilibrium simulations of a variety of peptides and proteins. We began with simulations of penta-alanine (hereafter, Ala(5)), the α-helix K19, and β-hairpins chignolin (starting structure PDB entry 1UAO[32]) and the "GB1" hairpin from the C-terminal fragment of Protein G. Furthermore, we performed microsecond length simulations of the globular proteins GB3 (starting structure PDB entry 1P7E,[33] similar to protein G and containing a motif homologous to the GB1 hairpin) and lysozyme (4LZT[34]) in solution. Each system was equilibrated with protein backbone atoms held under progressively decreasing restraints for up to 7 ns, depending on the size of the system. All simulations were performed with TIP4P-Ew water (the water model used to develop our solution phase charge set) in sufficient quantity to enclose the peptide and solvated protein systems by at least 10.0 Å within octahedral boxes after equilibration under constant pressure dynamics. Nonbonded interactions were calculated with a 10.0 Å cutoff on

Lennard−Jones interactions, a homogeneity approximation for long-ranged van der Waals interactions, and smooth particle-mesh Ewald electrostatics with direct and reciprocal space accuracies near 1 part in 200 000 (this corresponds to a direct sum tolerance of 5.0×10^{-6} and a direct space cutoff of 9.0 Å with the default mesh grid spacing of up to 1 Å). A 2 fs time step was used in all simulations, along with the SHAKE[35] and SETTLE[36] algorithms to constrain the lengths of bonds to hydrogen. Most of the small systems were simulated at 277 K, by a Langevin thermostat[37] with collision frequency 3/ps, to replicated NMR conditions. The GB1 hairpin system was simulated at 298 K to investigate β-sheet stability at room temperature. Globular proteins were simulated at 300 K, again with the Langevin thermostat.

4. RESULTS

4.1. Rederivation of IPolQ Charges as a Minimal Perturbation from Charges Appropriate to the Vacuum Phase. The updated IPolQ procedure described in the Theory section uses the same fitting data as the original method to derived Q^{IPol}, but it also derives Q^{vac}, a set of charges appropriate for modeling electrostatics in vacuo. These two charge sets would be independent, and the updated Q^{IPol} would be the same as the original, but for the fact that the restraint equations used to temper the fit of the original Q^{IPol} instead temper values in Q^{vac}, and the perturbation ΔQ that relates Q^{IPol} to Q^{vac} by eq 2 is subject to its own set of restraints. In our protocol, the size of ΔQ is controlled by setting the parameter V_{GP} in eq 7. Larger values of ΔQ could make it difficult to transfer torsion parameters developed for systems in vacuum to simulations in solution, but too small a ΔQ would imply $Q^{IPol} \simeq Q^{vac}$, and neither set of charges would describe its intended environment well. We chose to set V_{GP} to 0.005 kcal/(mol·e^2-instance), where e is the charge of a proton. This is half the stiffness of restraints holding charges of underdetermined sites to zero in this and earlier IPolQ derivations.

The fitting matrix is not well conditioned in either case: for main-chain amino acids, we found the (2-norm) condition number of the matrix A in eq 3 to be 1.8×10^6 and that of the extended matrix in eq 7 to be 2.2×10^6. The possibility of a worst-case loss of precision compels us to use double precision arithmetic for what is already a very large matrix, but the results are nonetheless consistent. The earlier IPolQ protocol, for instance, judged convergence of the charge set by the point at which successive generations changed the fitted charges by less than the inclusion of one of the fitting conformations. (These changes were on the order of 0.01e.) As before, we are more concerned with the effect that a different set of constraints might affect the outcome.

The value of V_{GP} we chose seems to have minor effects on the overall accuracy of Q^{IPol} and Q^{vac}. Averaged over all systems and conformations, the original Q^{IPol} reproduces the target electrostatic potential with 1.82 kcal/mol·e root-mean-squared error (rmse), whereas $Q^{vac} + \Delta Q$ reproduces the same potentials with 1.89 kcal/mol·e rmse. The effects on selected systems are shown in the lower panel of Figure 1. Over the range we tested V_{GP}, the effects on electrostatic potential rmse appear to begin to approach an asymptotic limit (half the difference between the electrostatic potentials computed in vacuo and in the condensed phase). If V_{GP} is tuned even lower, then the rmse of $Q^{vac} + \Delta Q$ actually improves over the original Q^{IPol} due to the fact that the perturbation charges are then much less restrained than the original charges. However, tuning V_{GP} too low leads to undesirable effects such as Ala C_β taking a charge of −0.21 at V_{GP} = 0.0025, as opposed to −0.04 in the original Q^{IPol}. Accuracies and charge perturbations for all amino acids at V_{GP} = 0.005 are shown in Table 3.

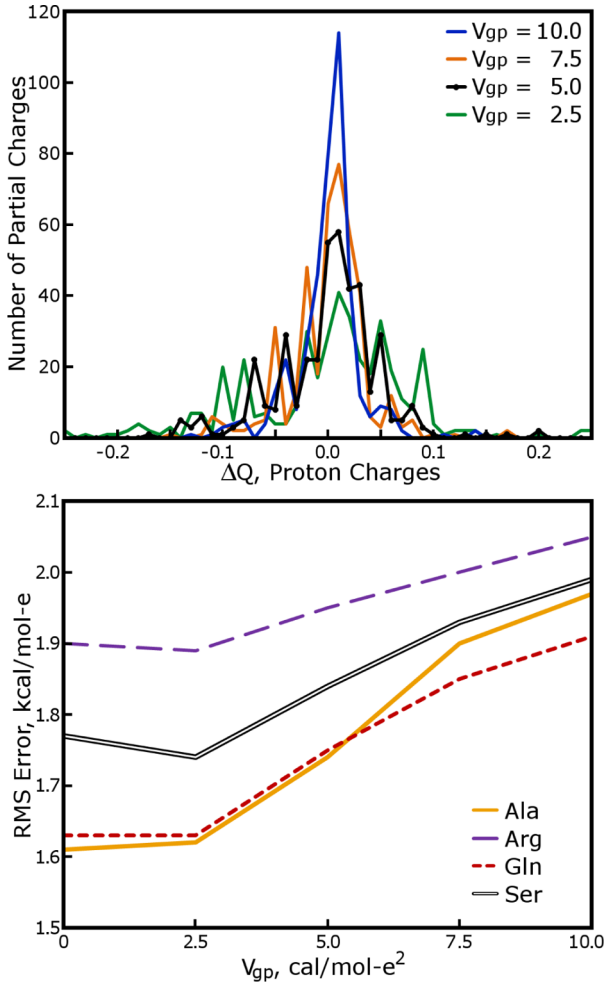

Figure 1. Effects of the coupling constant V_{GP} on ΔQ and the accuracy of electrostatic potentials for condensed-phase systems. V_{GP} determines the strength of the harmonic penalty restraining all charges ΔQ toward zero. Values of the root-mean-squared error (rmse) describe electrostatic potentials projected by each dipeptide's molecular mechanics charge set relative to the QM target. V_{GP} = 0 corresponds to partial charges by the original IPolQ method[20] before the extension, which allows us to express them as a perturbation to charges appropriate for simulations in vacuo.

At V_{GP} = 0.005, the majority of values in ΔQ are smaller than 0.05e, as shown by the top panel in Figure 1; the larger values tend to be in buried methyl carbons, whose charges are small to begin with, and the atoms of polar head groups. Neither of these changes are likely to make torsion parameters developed with Q^{vac} less transferable because the electrostatic potential energy surface will change only if atoms with large values of ΔQ can rotate around nearby atoms possessing large of charges of their own. More concerning is the effect of ΔQ on the backbone. The perturbation does not, in fact, bring $Q^{vac} + \Delta Q$ to the polarity of the original Q^{IPol}: the carbonyl carbon atom becomes less positively charged by as much as 0.05e, the oxygen becomes less negatively charged by 0.03e, and the polarity of the N−H backbone group is also decreased. However, the perturbation does make the backbone significantly more polar than Q^{vac} alone

dx.doi.org/10.1021/ct500643c | J. Chem. Theory Comput. 2014, 10, 4515−4534

Table 3. Properties of the Original and Extended IPolQ Charge Sets[a]

dipeptide	accuracy[b] original $Q^{IPol,c}$	$Q^{vac,d}$	$Q^{vac} + \Delta Q^{c}$	max ΔQ, atom[e]
Ala	1.57	1.81	1.74	0.06, CB
Arg	1.90	2.42	1.95	0.05, CA
Asn	1.65	2.01	1.72	0.10, CA
Asp	1.91	2.58	1.97	0.08, CB
Cys	2.39	2.76	2.46	0.07, CA
Gln	1.63	1.99	1.75	0.09, OE1
Glu	1.79	2.44	1.87	0.08, CA
Gly	1.57	1.93	1.73	0.01, CA
Hie	1.85	2.20	1.95	0.12, ND1
Ile	1.75	2.05	1.77	0.11, CG2
Leu	1.77	2.00	1.81	0.15, CG
Lys	1.92	2.43	1.94	0.05, CE
Met	2.12	2.35	2.20	0.12, CE
Phe	1.61	1.96	1.67	0.08, CA
Pro	1.75	1.79	1.86	0.13, C
Ser	1.77	2.16	1.84	0.09, OG
Thr	1.77	2.11	1.81	0.12, OG1
Trp	1.69	2.06	1.77	0.06, CA
Tyr	1.64	1.97	1.76	0.08, OH
Val	1.56	1.85	1.65	0.17, CB

[a]Errors and charge variations are expressed for a restraint of V_{GP} = 0.005 kcal/(mol·e^2-instance) applied to all perturbation charges. [b]Root-mean-squared error (rmse) of fitted MM charges in replicating the QM target electrostatic potential. [c]The target is the average electrostatic potential of the solute's MP2/cc-pvTZ wave function in vacuum and in the solvent reaction field potential due to TIP4P-Ew water. [d]The target is the electrostatic potential of the solute's MP2/cc-pvTZ wave function in vacuum. [e]Maximum absolute deviation in partial charges unique to this residue; backbone atoms frequently showed ΔQ of 0.06−0.10, as shown in the Supporting Information.

would describe it. As a consequence, the IPolQ protocol continues to model protein backbones with more polarity than the Cornell and Duan charge sets, maintaining a principal finding of our earlier study.

4.2. Preliminary Torsion Parameters for Amber ff14ipq. Torsion Fourier series terms for all amino acids were fitted by a linear least-squares approach as described in the Methods. The preliminary set of 28 000 structures and MP2/cc-pvTZ single-point energies is, to our knowledge, the most extensive potential energy surface ever employed for this type of molecular mechanics parameter development. In addition to the atom types present in Amber ff99SB and new atom types for C_β atoms introduced by the Simmerling group, we included new types for each of the Lennard−Jones modifications introduced to adjust hydration free energies of amino acid side-chain analogues in the initial IPolQ derivation. We also included new types for the glycine C_α atom and proline backbone nitrogen because of the unique chemical and bonding structures around these atoms.

When the electrostatics of each di-, tri-, and tetrapeptide are described by Q^{vac}, the fitted torsion Fourier series terms complete a force field and describe the molecular mechanics energy of the system in vacuo. The rmse of these energy estimates relative to the MP2/cc-pvTZ single-point energies for di- and tetrapeptides is given in Table 4. Also given in this table are the molecular mechanics (MM) energies obtained by reoptimizing a much smaller set of torsion Fourier series terms, the parameters found in Amber ff99SB,[7] in conjunction with our newly derived Q^{vac} and the MM energy estimates that would have been obtained if

Table 4. Accuracy of MM Energy Estimates with Different Torsion Parameters[a]

system	force field accuracy V-ff14[b]	V-ff99[c]	V-ff14nt[d]	term count V-ff14	V-ff99	torsion energy sum V-ff14	V-ff99
Arg	0.92	1.20	2.19	63	35	31.53	13.65
Ash	1.01	1.38	3.69	53	31	27.25	16.77
Asn	0.80	1.35	1.99	53	28	8.12	17.95
Asp	1.74	3.44	3.61	41	28	14.70	11.57
Cys	0.99	1.24	2.23	43	29	16.50	12.09
Cyx	1.22	1.48	1.94	38	28	19.43	10.93
Glh	0.80	0.99	2.16	60	33	29.16	14.44
Gln	0.67	0.86	2.00	60	30	20.40	18.33
Glu	1.27	1.80	2.04	48	30	16.97	12.32
Hid	0.79	0.99	1.93	52	34	25.44	11.75
Hie	0.78	1.03	1.87	52	34	23.26	11.72
Hip	1.50	1.63	2.63	51	33	24.28	11.78
Ile	0.64	1.14	2.28	63	33	24.46	13.25
Leu	0.77	0.99	2.01	48	33	18.87	13.75
Lys	1.20	1.57	2.73	56	34	17.28	14.43
Met	0.79	0.96	2.13	49	29	15.65	12.86
Phe	0.75	0.99	1.90	42	30	25.54	11.80
Ser	0.81	1.00	1.97	45	33	24.83	12.15
Thr	0.89	1.35	2.75	61	36	75.76	13.28
Trp	0.79	1.12	2.40	55	37	40.57	12.03
Tyr	0.79	0.90	1.95	45	32	24.48	13.09
Val	0.63	0.81	1.65	41	30	24.43	11.97
Ala3	1.14	1.23	1.99	30	28	61.62	25.03
Gly3	0.96	1.17	1.85	21	19	43.20	23.71

[a]In all cases, the charge set Q^{vac} fitted to reproduce the electrostatic potentials of blocked dipeptides in vacuo was used to estimate the molecular mechanics energy of each blocked dipeptide in vacuum. All energies are given in kcal/mol. [b]The V-ff14 force field: Q^{vac} has been substituted for the implicitly polarized charge set in the release version. [c]The ff99 force field, with Q^{vac} as derived for V-ff14 (identical to the force field in the first column, but with a smaller torsion parameter space). [d]V-ff14 with no torsion Fourier series terms.

no torsion terms were used. The overall contributions from the torsion terms are often small, only reducing the rmse of MM energy estimates by 1.3 to 1.5 kcal/mol relative to a model that has no such terms. The overall size of the torsion terms' contributions may understate their importance, given the number of published force field improvements based on changes in these terms.

The new ff14ipq force field has many more torsion Fourier series terms than the ff99SB force field: 427 to 67. Most of the new parameters were added by including the Simmerling group's unique C_β atom types. While the total number of torsion parameters increases nearly 7-fold from ff99SB to ff14ipq, Table 4 shows that the number of parameters expressed in any particular system doubles at most. New atom types C8, 3C, and 2C for C_β have added new parameters for χ_1 and χ_2, which distinguish the side-chain rotamer energetics for protonated His, Arg, and Lys, the amino acids Ile, Thr, and Val, and other amino acids. This partitioning is just one of many possible approaches, but it seems to have achieved a similar effect to an earlier extension of the ff99SB force field, which improved the rotamer propensities of residues Ile, Leu, Asp, and Asn by adding new atom types and fitting the newly minted torsion parameters to MP2/(aug)cc-pvTZ energy profiles.[31] The ff14ipq parameter space makes notable improvements over the ff99SB parameter space for Ile, Asp, and Asn; it appears that providing only a single

dx.doi.org/10.1021/ct500643c | J. Chem. Theory Comput. 2014, 10, 4515−4534

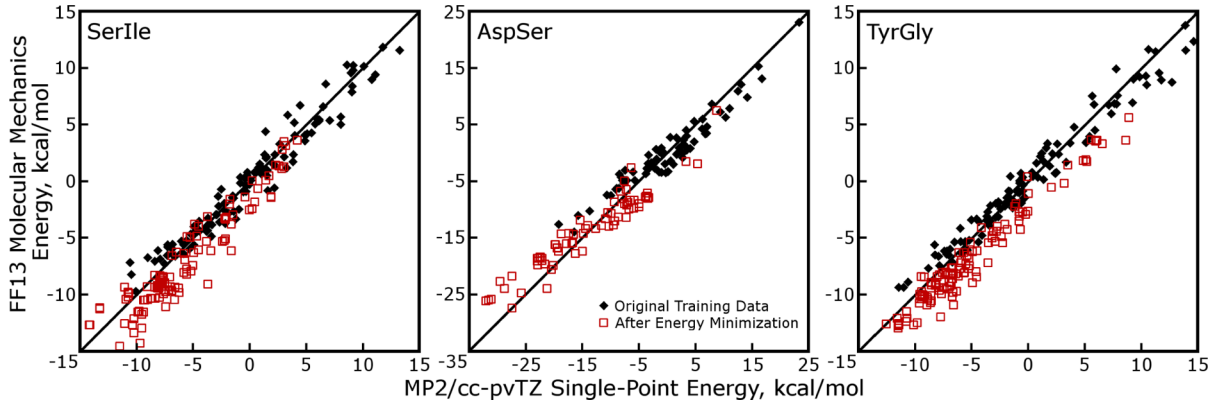

Figure 2. Energies of dipeptide and tripeptide conformations before and after energy minimization with the preliminary V-ff14 force field. All molecular mechanical energies are adjusted according to the adjustment constants found while fitting torsion parameters; quantum mechanical energies are normalized to a mean of zero. Hence, the energies of conformations found in the fitting data (black diamonds) lie directly on the trendlines, and the energies of system conformations optimized according to the preliminary V-ff14 (red, open squares) may not track it.

atom type for C_β forces the same set of torsion parameters to average between disparate potential energy surfaces and causes difficulty fitting the data. Including the new "TG" atom type for glycine C_α likewise decouples two distinct potential energy surfaces: the dihedral angles between a backbone α-hydrogen and the backbone polar hydrogen, or the backbone carbonyl oxygen, are sampled twice (roughly 120° apart) in glycine and only once in all other amino acids. These dihedrals contribute directly to the protein ϕ and ψ propensities, and the new TG atom type improves the fit for glycine tetrapeptide by roughly 0.5 kcal/mol rmse and the fit for all other dipeptides by up to 0.1 kcal/mol rmse (data not shown). We also tried adding distinct atom types to the glycine α-hydrogens, but this was not as effective as adding the TG atom type. In contrast, distinguishing the proline backbone nitrogen as TN has a negligible effect on the fit for any residues but the proline itself. Given the inherent difficulty in sampling dihedrals related to this atom type, we considered whether to include it at all, but it does decouple the proline ϕ profile from other amino acids and serves as a handle for future parameter development.

With the expanded parameter set, the torsion potentials contribute more and more to the total molecular mechanics energy; their total contributions in ff14ipq are frequently double their contributions in the refitted ff99 parameter set. However, as shown by the final column of Table 4, a model with no torsion potentials is only perhaps 1 to 2 kcal/mol less accurate. The larger torsion potential contributions in ff14ipq are being counterbalanced by larger energy adjustment constants C_{adj}, which never appear in the force field. The number of adjustable terms and the size of C_{adj} are only weakly correlated (Pearson coefficient 0.40), perhaps because many copies of the torsion terms can appear in the total MM energy and also because the amplitudes of the terms themselves are so variable. Despite this, the sizes of C_{adj} may be a useful indicator of whether a model is overfitted.

The increases in C_{adj} led us to track the sampling of each torsion parameter throughout the fitting set, as shown in the Supporting Information. With our exceptionally thorough data set, the torsion parameters related to rotatable bonds seem to be well-sampled, on the whole and in each system-dependent context in which they appear. Even with this degree of coverage, however, these fitted torsion parameters are not the final settings distributed as ff14ipq, whether in combination with the partial charge set Q^{vac} or Q^{IPol}; they are, rather, a first draft. The following section presents results obtained when the draft model was allowed to guide geometry optimizations and thereby expose the ways in which its parameters could conspire to accumulate errors over the course of a simulation.

4.3. Energy Optimization with the Preliminary Amber ff14ipq. Energy optimization served as a preliminary test of our fitting program and also of the robustness of our fitting data. Energy minimization of the fitting set's structures using the Amber pmemd program and the preliminary V-ff14 force field confirmed that the MM energy computed for each initial structure matched that produced by the mdgx fitting module after solving eq 8. After energy minimization in vacuo, some initial structures converged to the same final configuration, but a number of local minima were still produced for each di-, tri-, and tetrapeptide system. We computed MP2/cc-pvTZ single-point energies for new structures in three tripeptide systems shown in Figure 2 and compared them to the MM energy estimates to test whether the new model's 427 parameters were prone to overfitting.

The results in Figure 2 show that V-ff14 is often able to guide small peptides into configurations that MP2/cc-pVTZ calculations agree have lower potential energy. However, after energy minimization, V-ff14 tends to estimate the energy of each configuration as being more favorable than the MP2/cc-pVTZ single-point energy. The degree to which V-ff14 departs from the 1:1 trendline with its benchmark may be small, as shown in Table 5. However, geometry optimizations in systems with His, Arg, and Lys led to MM energies that departed severely from the QM benchmark, despite the agreement maintained in the training set. By adding the structures freely optimized by V-ff14 back into the training set, however, the fitted V-ff14 parameters became more robust. The second generation V-ff14 nearly eliminated departures from the QM benchmark in Arg and His and also showed minor improvements in its ability to optimize the structures of most other residues.

Lysine presented more of a challenge: as shown in Figure 3, the second generation V-ff14 apparently contained a new artificial minimum that was subsequently found in all optimizations of the dipeptide. However, the third generation V-ff14 eliminated this trap as well. We examined the molecular mechanics energies of lysine conformations in the original training set, the second-generation training set, and the final training set with respect to each generation of torsion parameters. Torsion parameters describing two of the rotatable bonds generate the severe

dx.doi.org/10.1021/ct500643c | J. Chem. Theory Comput. 2014, 10, 4515−4534

Table 5. Overstatement of Energy Minimization Results by Successive Generations of the V-ff14 Force Field[a]

	generation		
amino acid	1	2	3
Ala(3)	−0.8744	0.1627	
Arg	−17.3268	−1.0679	
Asn	−0.9843	0.421	
Asp	0.359	0.4314	
Cys	−0.8042	−0.4301	
Gln	−0.7356	−0.0738	
Glu	−0.301	−0.431	
Gly	−3.2341	−0.9707	
Hid	−14.4731	−1.0202	
Hip	−68.7312	2.5557	
Ile	−1.288		
Leu	−1.1398		
Lys	−68.8004	−28.5527	0.9739
Met	−0.5633	0.1562	
Phe	0.0119		
Ser	−0.4869	0.6037	
Thr	0.224		
Trp	−1.8196	−0.3194	
Tyr	−0.9153	−0.4697	
Val	−1.8519	−0.0154	

[a]V-ff14 is able to guide unrestrained energy minimizations of structures in its own training set and reduce the internal potential energy by up to tens of kcal/mol. (This may be realistic, as most training set structures were restrained in one or more torsional degrees of freedom.) However, when re-evaluated at the MP2/cc-pVTZ level, the resulting structures were often not as optimal as molecular mechanics depicted. Negative numbers in the table indicate that V-ff14 strayed from its MP2 benchmark and estimated its optimizations to be too favorable. Gaps in the table indicate that a system was omitted from one generation, due to compute cluster downtime or sufficiently low error in the previous generation.

departures from the benchmark seen in the first generation of the force field. First, a wildcard parameter with 3-fold periodicity describing rotation of the lysine side chain amino terminus takes on an amplitude of −4.99. In the original training set, every conformation sampled the orientation of the terminus in a staggered conformation of the hydrogens at the crest of the cosine wave. The eclipsed conformation, favored by nearly 10 kcal/mol by this spurious parameter, was not sampled in the training set, and the result was six permutations of amino and aliphatic hydrogen interactions creating a nearly 60 kcal/mol fictitious energy release by adopting the unnatural, eclipsed conformation in every structure optimized by the first generation of the force field. A more serious problem occurred in the four-term Fourier series describing interactions between the atom type C8, coined by the Simmerling group for lysine and arginine Cβ atoms, and the backbone carbonyl carbons. Because of its unique appearance in lysine and arginine, this Fourier series strongly influences the backbone ϕ angles adopted by the residues. The original training set did not sample a 100° arc in ϕ for either residue, however, and the Fourier series became fitted to produce an unnaturally low energy in the first generation, approximately 14 kcal/mol more favorable than any conformation of the ϕ angle should have allowed. Only some of the lysine structures fell into this trap when optimized by steepest descent energy minimization under the first generation of the force field, and the result was the two striations, which can be seen in Figure 3. In the second-generation training set, sampling in these torsion parameters improved considerably by including the structures trapped in the first generation's spurious minima, as shown in Table 6. Rotation of the lysine amino terminus needed to be sampled completely in order to get a transferable model, but this was accomplished by the third generation.

The set of torsion parameters optimized in the third generation of training was chosen for pairing with Q^{IPol} and distributed as ff14ipq in AmberTools14 (see http://ambermd.org).

4.4. Dipeptide Potentials of Mean Force: The Impact of ΔQ and Torsion Parameter Refinement. Blocked dipeptides, the simplest systems exhibiting protein-like backbone and side-chain dynamics, were studied extensively to characterize the effects of slight alterations in the charge set, ΔQ, which differentiate ff14ipq, our force field for simulations of proteins in water, from V-ff14, a force field used as a parameter fitting apparatus that is otherwise appropriate only for simulations of proteins in vacuum. Furthermore, while we found that torsion parameter refinement could eliminate catastrophic traps in the

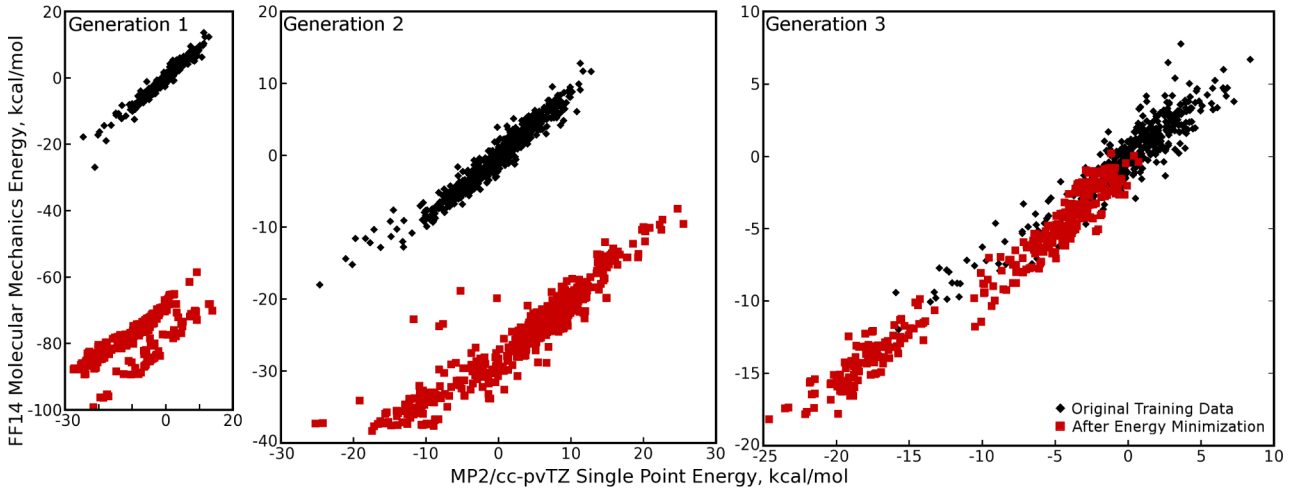

Figure 3. Energies of lysine dipeptide conformations estimated by three generations of the V-ff14 force field. Molecular and quantum mechanical energies are adjusted as described in Figure 2. Each generation's fitting set contained all of the initial fitting set data plus conformations created by all previous generations.

dx.doi.org/10.1021/ct500643c | J. Chem. Theory Comput. 2014, 10, 4515−4534

Table 6. Sampling and Amplitudes of Torsion Fourier Series Terms in Lysine and Arginine Residues[a]

Term	Generation	Sampling Histogram[b] −π 0 π	n=1[c]	n=2	n=3	n=4
C -N- CX-C8	1	X=.OXXU XXXXXX	-0.76	-2.22	-3.40	-5.11
C -N- CX-C8	2	X=.OXXOoXXXoXXXXXX	0.26	0.02	0.23	-0.37
C -N- CX-C8	3	X=.OXXXXXXXXXXXXXX	0.17	-0.23	0.18	-0.34
X -C8-NL-X	1	X .XX XX .X	-4.99			
X -C8-NL-X	2	X XX.XX XX XX XX.X	3.08			
X -C8-NL-X	3	XXXXXXXXXXXXXXXXXX	-0.28			

[a]The torsion fourier series terms describing dihedral interactions of atom types C−N−CX−C8 (backbone carbonyl carbon of any residue N-terminal to lysine or arginine, backbone nitrogen, Cα, and Cβ of lysine or arginine) and X−C8−NL−X (generic torsion affecting amino terminal hydrogens) evolve rapidly over three generations as sampling of the backbone ϕ angles and amino headgroup orientations becomes more complete. [b]Number of conformations in the data set displaying each angle. (0).: = e o U O 0 @ X (>10). [c]Periodicity of each Fourier series term.

potential energy surfaces of some amino acids, the approach made modest improvements in the behavior of every other amino acid as well. We computed multiple two-dimensional potentials of mean force (PMFs) for alanine, glycine, and serine dipeptides to assess backbone propensities with either charge set and with alanine dipeptide to examine the impact of torsion refinement.

As shown in Figure 4, the PMFs of simple, nonpolar amino acids are nearly identical for each charge set, particularly in the populated regions of the Ramachandran plot. By inspection, the polarization of charges slightly increases the alanine propensity toward (right-handed) α-helices, decreases its propensity toward β-sheets, and leaves the model's strongest tendency, toward poly proline II backbone conformations, unchanged. Larger differences in all dipeptides appear near (ϕ,ψ) = (0,0), as the backbone N−H and C═O dipoles become favorably aligned to create a more favorable free energy under the ff14ipq model, which strengthens both dipoles considerably. However, the (ϕ,ψ) = (0,0) arrangement remains strongly disfavored by steric clashes. The PMFs for alanine and glycine amino acids are also very close to those of the ff99SB force field.[7] In contrast, the serine dipeptide is not often studied. The lowest row of panels in Figure 4 suggests that solvent effects that make the side-chain charges more polar also facilitate transitions between poly proline II and left-handed α-helicies, but the relative weights of each major conformation are unchanged. While it is not certain what differences in protein folding or dynamics a ΔΔG of 0.25 kcal/mol in a low-energy region of (ϕ,ψ) space could lead to, these plots suggest that ΔQ changes protein folding transition states more than equilibria.

The importance of fitting torsion parameters in the context of an appropriate charge model is shown in Figure 5. In contrast to the changes introduced by fitting torsion parameters for V-ff14 to a gas-phase QM PES and then transferring them to work in ff14ipq, fitting torsion parameters in the context of IPolQ charges, retroactively forcing this charge model to mimic a gas-phase QM PES, makes more substantial changes to the alanine PMF. The difference between this PMF and the ff14ipq PMF is much more frustrated than the difference between the ff14ipq and V-ff14 PMFs. The strong gradients evident in the difference map drive the α-helical minimum approximately 15° to the south if the torsion parameters are fitted in the context of Q^{IPol} and gas-phase quantum data. β-sheet, poly proline II, and L-α-helical minima are also distorted. When working against an already polarized electrostatic model, the torsion parameters were also less effective at reproducing the gas-phase QM PES (data not shown). These contrasts probably originate in the opposing nature of internal solute and solute−solvent electrostatic interactions.

Refining the torsion parameters has its own effects on the alanine dipeptide PMF, as shown in Figure 6. Even though the quantum mechanics target is unchanged and the initial set of torsion parameters had no serious problems optimizing the geometries of alanine (tetra)peptide, the third-generation model produces noticeable differences in the PMF. Notably, what was a saddle point between α-helical and poly proline II conformations in the first-generation ff14ipq becomes a local minimum in the third. Transitions between poly proline II, α-helical, and L-α-helical conformations also tranverse lower barriers according to the third generation of torsion parameters, and in some of the high-energy regions of the PMF, the differences between the first and final generations are as great as those shown in Figure 5. While the magnitudes of the differences are comparable to those in Figure 5, however, the locations of all the major backbone conformational minima remain in the same places, and the most significant differences reside in high-energy regions of the PMF, which stands in contrast to the consequences of fitting torsion parameters with the wrong charge set. As before, the influence of these changes on protein folding is not discernible from the PMF alone, but the differences again appear mostly in highly strained configurations, suggesting that the refined torsion parameters depict more frequent transitions between major backbone conformations.

4.5. Polypeptides in Water: Simulations with the Rederived IPolQ Charge Set. Peformance of the release version of ff14ipq, specifically the third generation of torsion parameters paired with the condensed-phase appropriate Q^{IPol}, was evaluated on Ala(5), β-hairpins chignolin and protein G C-terminal fragment, the α-helical miniprotein K19, two variants of the miniprotein Trp Cage, and globular proteins GB3 and lysozyme. Sources of each protein structure, as well as sequences of the miniproteins, are given in Table 7. Each of the miniproteins simulated in this study was selected to evaluate ff14ipq's performance on a particular secondary structure element.

The Ala(5) system has recently become a standard diagnostic of a force field's ability to balance three major backbone configurations and reproduce NMR *J*-coupling results. Best and colleagues made a comprehensive assessment of modern force fields[38] with respect to NMR data from Graf and co-workers,[39] calculating mean χ^2 values for each model's reproduction of 11 order parameters. The Karplus relations used to calculate order parameters from the MD simulations are sensitive to their own coefficients, but Best and colleagues took three different sets of Karplus coefficients and posited that a χ^2 value of 2.25 or less under all three Karplus relations indicated a high-quality force field. The results for ff14ipq using the same Karplus coefficients and quadruplicate 375 ns simulations of the unblocked peptide with protonated C-terminus are shown in Table 8. The details of the simulated system were intended to match the acidic conditions of the NMR experiments; new charges were derived specifically for the protonated C-terminal alanine and are given in the Supporting Information. The mean χ^2 values obtained with the original Karplus coefficients used by Graf and two sets of DFT-based Karplus coefficients from Case and colleagues[40] were 1.3 ± 0.0, 2.6 ± 0.1, and 1.5 ± 0.0, respectively. (The error bars

dx.doi.org/10.1021/ct500643c | J. Chem. Theory Comput. 2014, 10, 4515−4534

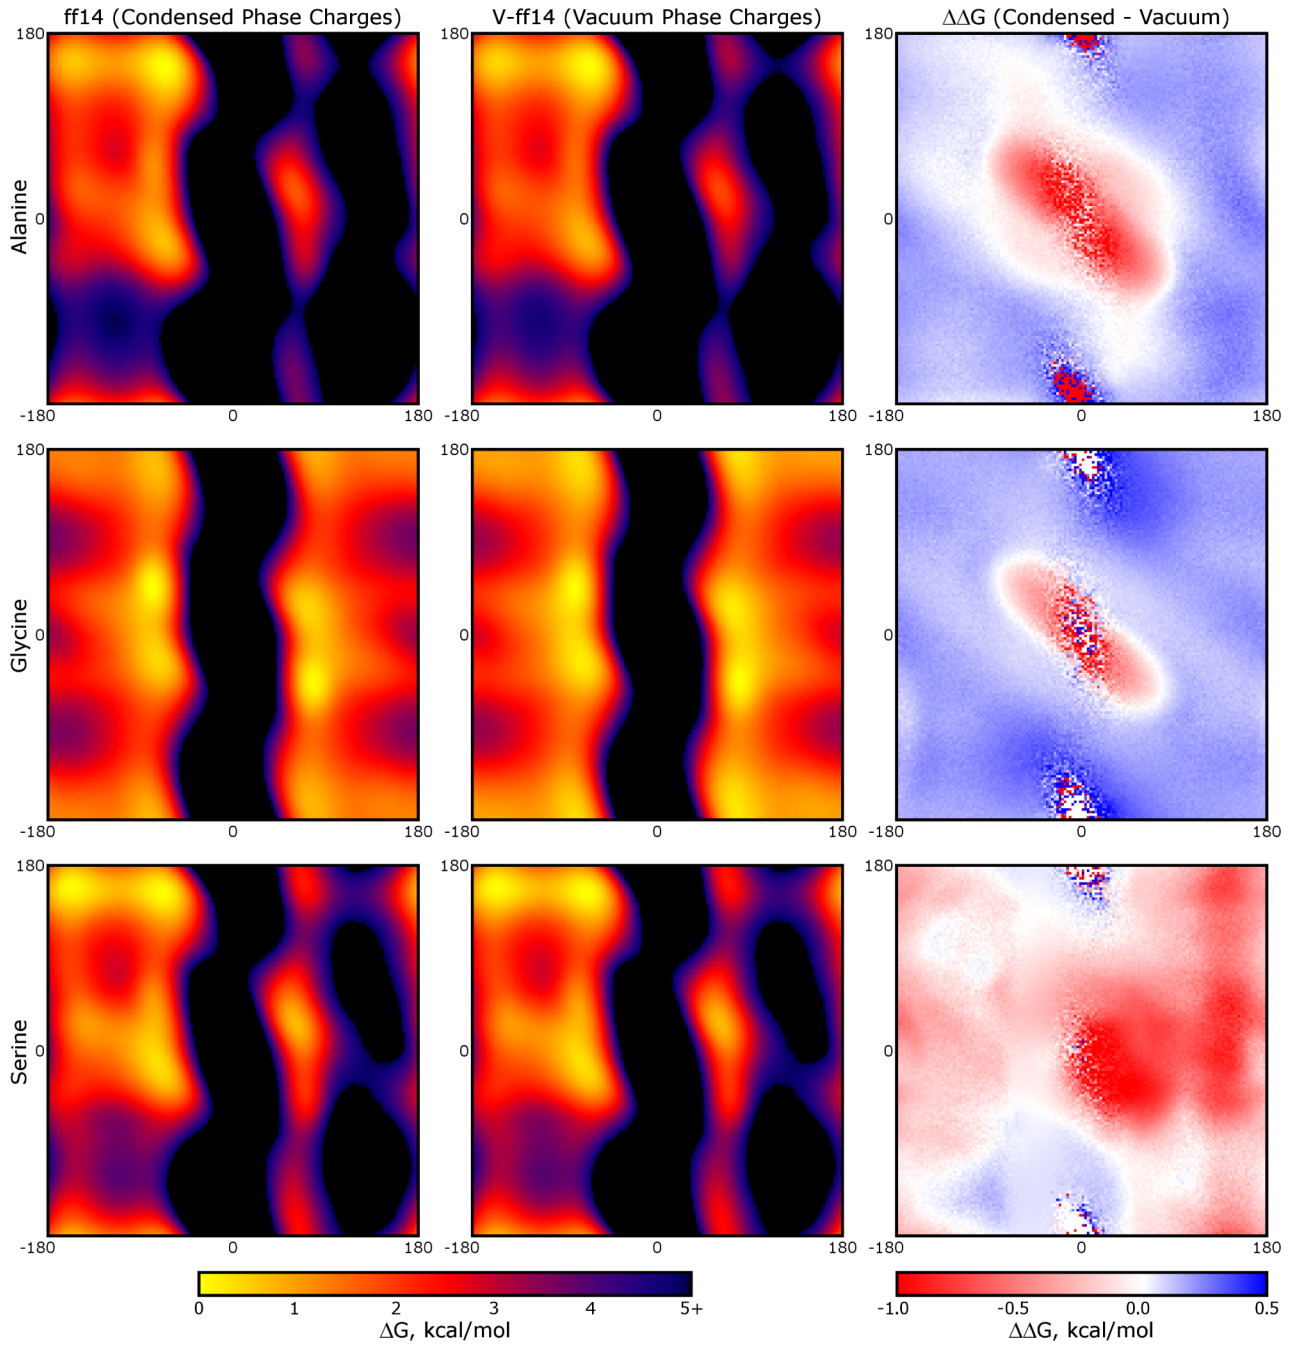

Figure 4. Potentials of mean force for blocked alanine, glycine, and serine dipeptides in water. The color scale for the preliminary versions of ff14ipq (leftmost panels) and V-ff14 (middle panels) measures ΔG, the energy difference between any point in (ϕ,ψ) space and the minimum free energy attainable in each model at 298 K. Differences between models are shown on the rightmost panels in a separate color scheme.

are standard deviations of the χ^2 from each of the four 375 ns simulations.) The major drivers of the scores are disagreement with the $^3J(C, C)$ coupling for the second alanine residue and the $^3J(H_N, C_\beta)$ coupling for the third residue. While ff14ipq does not meet Best's χ^2 standard by one of the DFT results, it does score very well by the other two, and the $^3J(H_N, C_\beta)$ coupling is known to be difficult for the Karplus relation itself.[40] By Best's definitions of each secondary structure element, ff14ipq models the central residue of Ala(5) in poly proline II, α-helical, and β-sheet conformations for 56, 18, and 14% of the pooled trajectories, respectively.

ff14ipq also stabilizes larger β-sheet structures, as indicated by the plots in Figure 7. The GB1 hairpin from Protein G[41−43] appears to be challenging for other fixed-charge force fields to stabilize (Emilio Gallicchio, personal communication) but maintains its secondary structure throughout the 250 ns simulations whether simulated with blocking groups or without. While it is reassuring that ff14ipq stabilizes a β-sheet structure, the hairpin should be only approximately 30% folded at 298 K. One element that appears to stabilize the system considerably is the ionic interaction between the termini: it remains intact for virtually the entire simulation with a mean and modal distance between the termini of 3.5 ± 0.7 Å. We investigated the stability of the blocked peptide with quadruplicate 600 ns runs and found no significant unfolding of the antiparallel β-sheet arrangement (data not shown; the 250 ns simulation reported in Figure 7 is

dx.doi.org/10.1021/ct500643c | J. Chem. Theory Comput. 2014, 10, 4515−4534

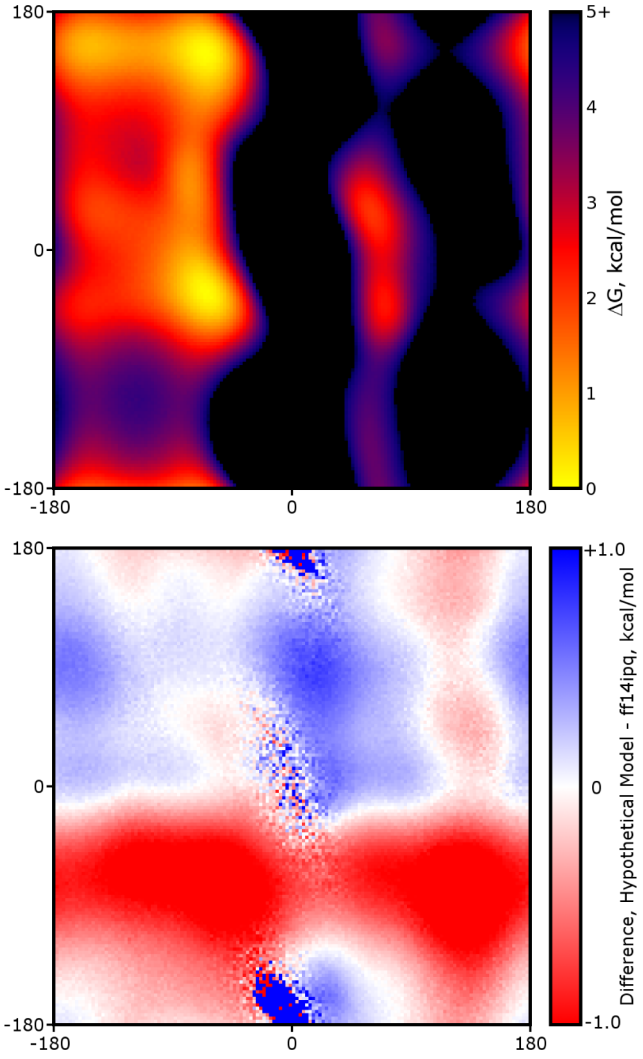

Figure 5. Difference plot of the alanine dipeptide PMF with torsion parameters derived for Q^{IPol} rather than Q^{vac}. The color scheme is similar to difference plots in Figure 4: here, solid red implies that a hypothetical (and incorrect) model fitting gas-phase quantum data in the context of charges appropriate to the solution-phase estimates a point in ϕ/ψ space more than 1 kcal/mol more favorably than a properly tuned model; solid blue would imply that the incorrect model disfavors the conformation.

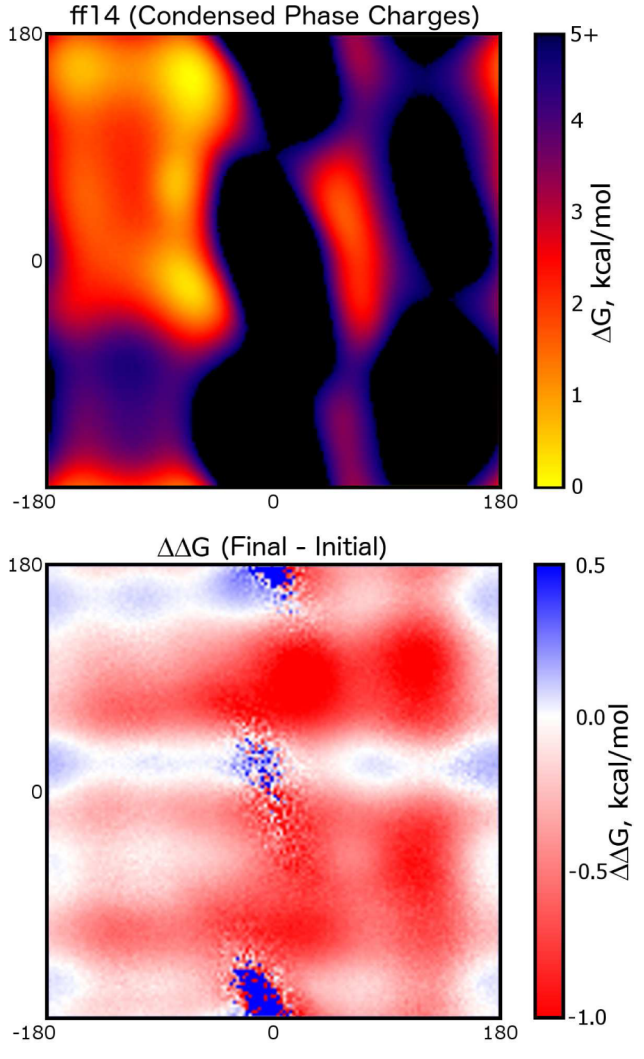

Figure 6. Potential of mean force for blocked alanine dipeptide in the release version of ff14ipq and comparison to the initial model. Both plots refer to combinations of condensed-phase charges with torsion parameters fitted to reproduce gas-phase quantum data. The color scale in the difference plot follows from Figure 4.

representative). Given that the melting curves for this peptide[43] and other β-hairpins are well-established, it may be possible to chip away at the stability of these structures in ff14ipq, and this is probably easier than to try and obtain the right structure from a model that cannot stabilize the peptide at all.

The chignolin system could be simulated economically on equally long time scales and appears to fluctuate between two major backbone states. Neither of them is very far from the structure obtained after equilibration; Figure 8 shows backbone rmsd relative to the first structure in the NMR ensemble, but if it is calculated relative to the first frame of the simulation after restrained equilibration, then the rmsd would fluctuate between 0.7 and 1.4 Å. In addition to calculating the standard, overall backbone positional root-mean-squared deviation (rmsd) relative the first NMR model, we computed the way in which individual residues deviate from the NMR ensemble as a whole. This was accomplished by making optimal alignments of the simulated backbones to each of the 18 chignolin NMR models found in the PDB and then computing the rmsd of each particular residue without further position alignment. The minimum residue rmsd to any of the models in the NMR ensemble was recorded for each frame of the trajectory to indicate how far from any of the plausible structures our simulations might have departed. The lower panel of Figure 8 shows histograms of these per-residue backbone position rmsds for each residue over the course of the simulation. This confirms that the backbone adopts two major conformations, with larger fluctuations in the zwitterionic, unblocked terminal residues. The arrangement of the termini follows the GB1 hairpin system: the zwitterionic termini begin the simulation in contact with one another. In chignolin, the termini fluctuate in concert most of the time, although for 7% of the simulation the two ionic groups separate by more than 4.5 Å. The β-hairpin fold is maintained throughout the simulation, but the state adopted in the middle of the simulation appears to be slightly outside the NMR structure ensemble. Further simulations may be able to fold the peptide, test the pathway against the known mechanism,[32] and suggest whether changes to the backbone−backbone interactions would tighten up the equilibrium structural results.

dx.doi.org/10.1021/ct500643c | J. Chem. Theory Comput. 2014, 10, 4515−4534

Table 7. Systems Simulated with ff14ipq[a]

system	PDB ID	sequence[b]	type
Ala(5)	none	AAAAA	backbone fragment
Trp cage	1L2Y	NLYIQWLKDGGPSSGRPPPS	miniprotein
Trp cage II	1RIJ	Ace-ALQELLGQWLKDGPSSGRPPPS-Nme	miniprotein
chignolin	1UAO	GYDPETGTWG	β-hairpin
GB1 hairpin		Ace-GEWTYDATKTFTVTE-Nme	β-hairpin
GB1 hairpin		GEWTYDATKTFTVTE	β-hairpin
K19 peptide		Ace-GGGKAAAAKAAAAKAAAAK-Nme	α-helix
GB3	1P7E	Unblocked; see PDB file	globular protein
lysozyme	4LZT	Unblocked; see PDB file	globular protein

[a]Simulations combined third-generation torsion parameters with Q^{IPol}. [b]Protein sequence; blocking groups are indicated by Ace- and -Nme.

Table 8. NMR *J* Couplings Calculated from Simulations of the Ala(5) System[a]

		simulation			
J coupling	residue	orig.[b]	DFT-1[c]	DFT-2	experiment
$^1J(N, C_\alpha)$	2	11.17	11.17	11.17	11.36
$^1J(N, C_\alpha)$	3	10.81	10.81	10.81	11.26
$^2J(N, C_\alpha)$	2	8.01	8.01	8.01	9.20
$^2J(N, C_\alpha)$	3	8.32	8.32	8.32	8.55
$^3J(C, C)$	2	0.86	0.77	0.85	0.19
$^3J(H_\alpha, C)$	2	1.66	1.43	1.59	1.85
$^3J(H_\alpha, C)$	3	1.91	1.67	1.84	1.86
$^3J(H_N, C)$	2	1.37	1.45	1.10	1.10
$^3J(H_N, C)$	3	1.33	1.38	1.09	1.15
$^3J(H_N, C_\beta)$	2	1.88	3.57	2.84	2.30
$^3J(H_N, C_\beta)$	3	1.89	3.57	2.84	2.24
$^3J(H_N, H_\alpha)$	2	5.68	5.21	5.79	5.59
$^3J(H_N, H_\alpha)$	3	5.73	5.30	5.84	5.74
$^3J(H_N, C_\alpha)$	2	0.58	0.58	0.58	0.67
$^3J(H_N, C_\alpha)$	3	0.61	0.61	0.61	0.68

[a]Calculated scalar *J* couplings pertain to averages over all four 375 ns trajectories. [b]Original Karplus coefficients used by Graf[39]. [c]DFT-based Karplus coefficients from Case[40].

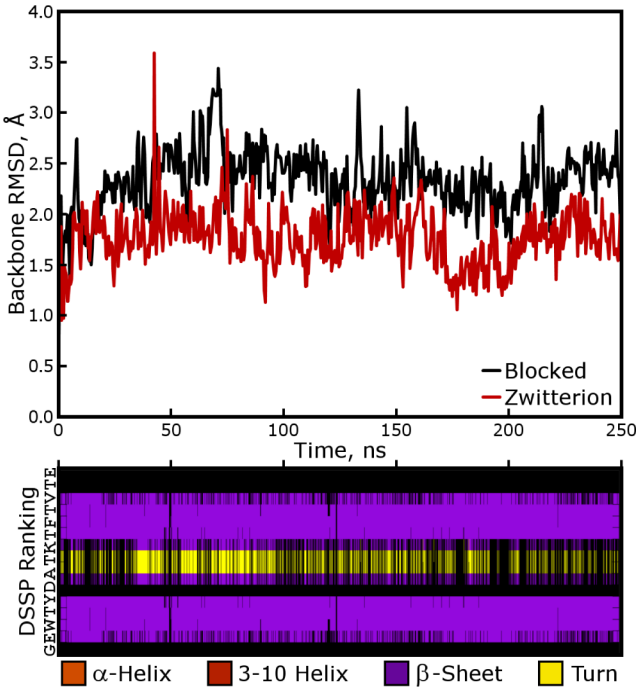

Figure 7. Backbone stability of the β-hairpin from Protein G over 250 ns. Blocked and unblocked forms of the peptide were simulated, and backbone rmsd is plotted for both variants. The DSSP chart below the rmsd plots refers to the blocked peptide system and indicates that the antiparallel β-sheet is maintained throughout the simulation. The DSSP chart for the unblocked case is essentially identical.

The K19 system[44] tested α-helical stability in ff14ipq. Over the course of 1 μs, residues 1−12 exhibit consistent helicity, while the C-terminus transiently explores alternative coil conformations and sometimes even forms a packed double-helical structure, as shown in Figures 9 and 10. The experimental data suggest that the protein's first 18 residues should consistently maintain an α-helix,[44] but the ff14ipq results are otherwise in qualitative agreement with experiment and previous simulations of this system. Of particular interest was the degree to which highly solvent-exposed lysine residues in the C-terminal region might make hydrogen bonds with the backbone. Polar atom Lennard−Jones radii in both the lysine amino headgroup and the backbone oxygen were altered when interacting with water for agreement with hydration free energies but allowed to maintain their original interaction with one another in the interest of keeping repulsion between nearby atoms from disrupting the internal potential energy surface. Interactions between lysine and the backbone oxygen may not be well tuned: much has been done to reduce the backbone oxygen's affinity for water, but its affinity for amino groups is only balanced, if at all, by the fact that amino groups' Lennard−Jones radii were reduced to provide greater affinity for water. We calculated the minimum distance between each lysine headgroup on K19 with any backbone oxygen atom throughout the simulation. As shown in Figure 11, lysines 4 and 9 rarely make contact with backbone carbonyl groups, but lysines 14 and 19 are much more likely to form hydrogen bonds to the peptide backbone. Defining an amino group to be hydrogen-bonded to a carbonyl if the nitrogen and oxygen atoms come within 3.2 Å of each other, lysines 4, 9, 14, and 19 are bonded to the backbone in 4, 7, 17, and 27% of the snapshots, respectively. While the interaction between some groups may be stronger than is realistic, their interactions remain transient. The nitrogen atom types in the lysine headgroup and the charged amino teminus of our β-hairpin systems are the same, and the oxygen atom types in carbonyl and carboxylate groups were altered in a similar manner with respect to their precursors in ff99. The artificial behaviors seen in each system are likely to have a common origin. Minor alterations may be warranted in either case.

As shown in Figure 12, both Trp Cage simulations showed very stable backbone configurations over the 500 ns simulations. The baseline rmsd of approximately 1.2 Å may indicate different preferences of the IPolQ charge set compared to the parameters used in NMR refinement, but larger departures from the

dx.doi.org/10.1021/ct500643c | J. Chem. Theory Comput. 2014, 10, 4515−4534

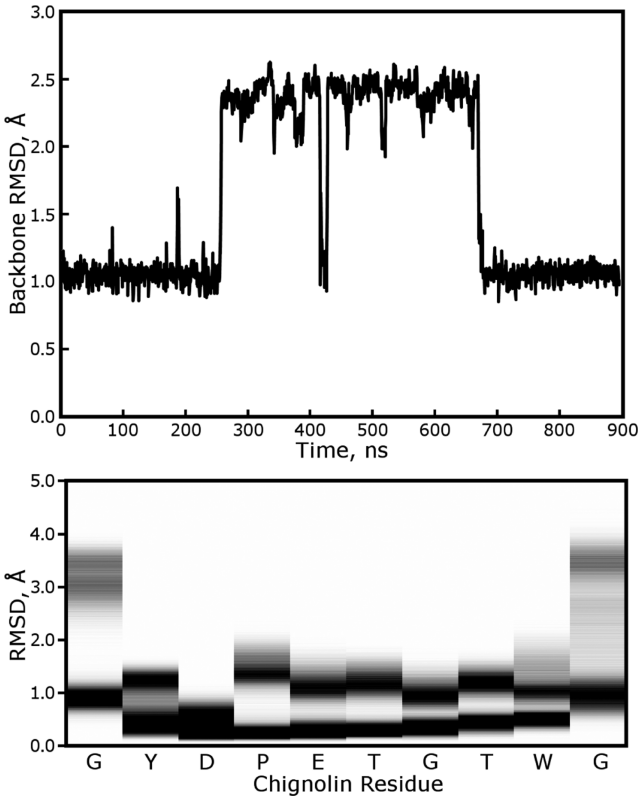

Figure 8. Stability of chignolin over 900 ns of dynamics. Backbone rmsd in the top plot is calculated relative to the first NMR structure; per-residue backbone rmsd reflects the deviation of each residue's backbone from the closest possible match out of the entire NMR ensemble.

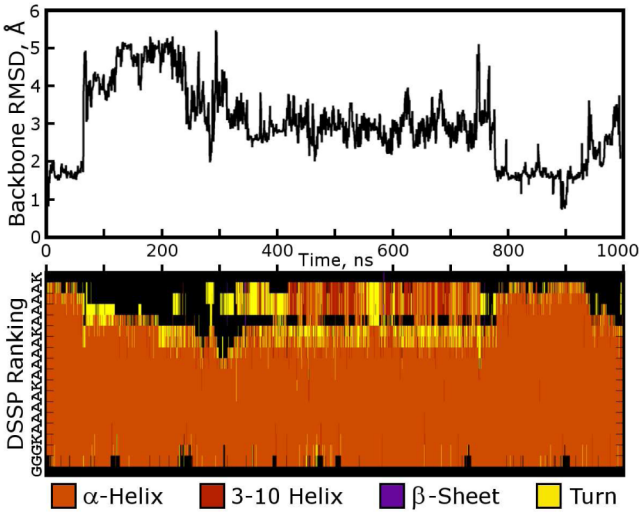

Figure 9. Backbone stability of K19 over a 500 ns simulation. The time axis applies to both the rmsd plot in the top panel and the DSSP plot in the lower panel. The K19 peptide is predominantly α-helical, with some instability at the C-terminus. Residues 16−19 begin to adopt a metastable 3−10 helical conformation near the middle of the simulation.

backbone configuration depicted in the NMR models were merely transient. Figure 12 shows histograms of per-residue backbone position rmsds as were calculated for chignolin. Most residues display low rmsds under this test, and most importantly the Pro-Pro-Pro sequence near the C-terminus remains stable. Some residues, in particular the Asp-Gly-Gly-Pro sequence in simulations of Neidigh's Trp Cage (PDB code 1L2Y[45]), show a

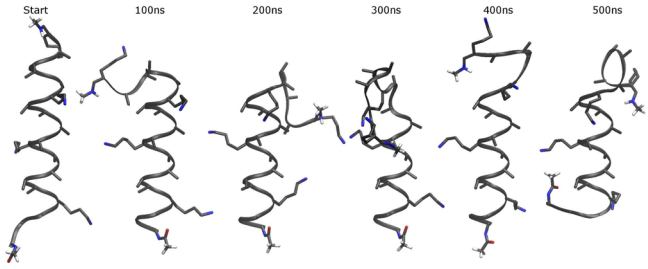

Figure 10. Conformations of K19 over 500 ns of dynamics. All conformations have been aligned relative to the stable backbone of residues 2−14. Lysine and alanine side chains are show in stick representation.

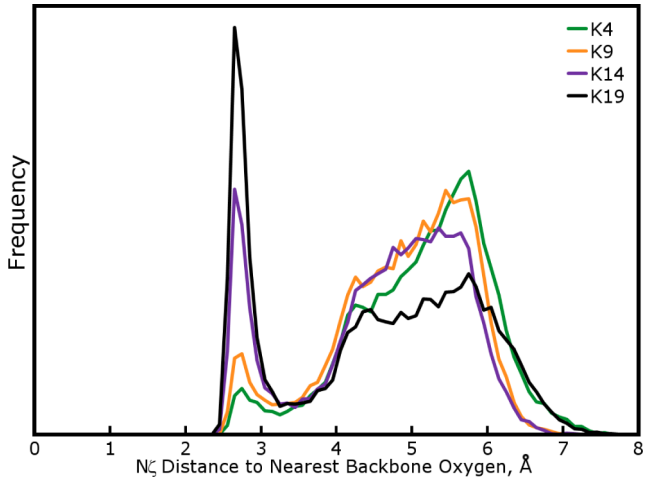

Figure 11. Radial distributions of lysine head groups and backbone oxygen atoms in the K19 system..

weakly bimodal distribution, suggesting that the simulations explore an alternate conformation not seen in the NMR ensemble. When simulating a hyperstable Trp Cage mutant (PDB code 1RIJ[46]), the Asp-Gly-Gly-Pro sequence again departs from the NMR ensemble more significantly than other regions of the protein. The conformations of these residues are good candidates for future analysis and refinement of ff14ipq.

Even though the IPolQ charge model is directed toward solvated molecules, ff14ipq is expected to have application to globular proteins. Most residues on proteins of even a few hundred residues have some degree of solvent exposure, and the majority of buried residues will be nonpolar and therefore little different when represented by IPolQ as opposed to other charge models.[20] Simulations of both GB3 and lysozyme showed that, overall, ff14ipq stabilizes the crystallographic backbone configurations of both proteins (Figure 13).

Because we began simulations of GB3 before realizing the importance of iterative torsion parameter refinement, the system provides an indication of the symptoms of underfitted parameters when simulating a complex biomolecule. The first generation of torsion parameters caused one of the lysine-containing loops (residues 9−16) to take on alternative conformations that drove the overal rmsd significantly higher. Application of the second-generation torsion parameters, which fixed an artificial minimum in the lysine backbone ϕ angle, greatly diminished these excursions, as shown in Figure 14. However, another very short loop of the protein, residues 39−41 connecting β-sheet to α-helical structures, makes excursions from the X-ray backbone structure in any generation. It may be

dx.doi.org/10.1021/ct500643c | J. Chem. Theory Comput. 2014, 10, 4515−4534

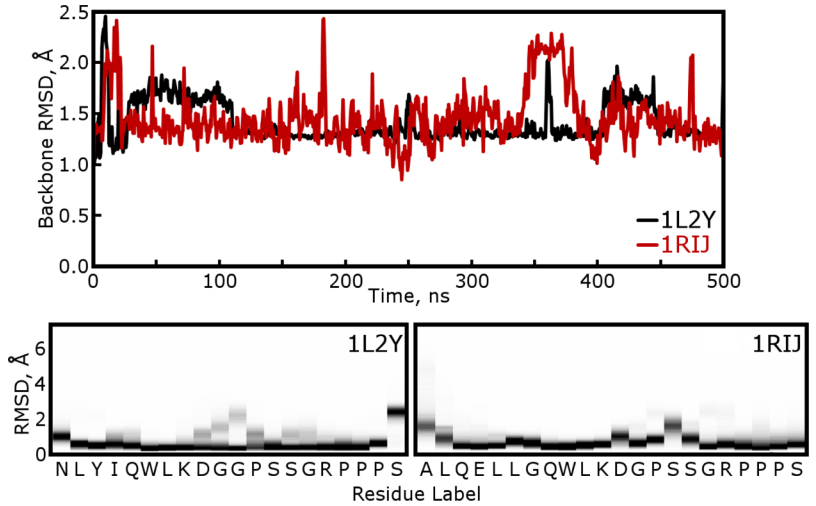

Figure 12. Backbone positional root-mean-squared deviations (rmsds) for Trp Cage miniprotein simulations. Simulations of each of two Trp Cage proteins are indicated by their respective PDB codes. The top panel shows overall backbone rmsd to the first published NMR model for each system. Lower panels show histograms of per-residue backbone rmsd to the closest possible match out of all published NMR models, darkened to indicate increasing occupancy at a particular deviation.

significant that this loop contains an aspartate residue: in the torsion fitting results (Table 4), aspartate presented one of the most difficult potential energy surfaces for our torsion parameters to capture. While ff14ipq seems to provide a better fit than would be possible with ff99, the greatly expanded parameter space that is almost certainly the basis for this improvement may have allowed some overfitting that we have not yet been able to eliminate with our iterative scheme. Alternatively, aspartate polarization in solution may not be well captured by our approximations. (A third possibility remains, whereby in solution this loop really does take on conformations not seen in the crystal lattice.) We intend that future releases of ff14ipq will build on all of the existing quantum data with further refinements to improve the description of larger protein systems.

The simulations performed to date with ff14ipq suggest that the force field is ready to guide much longer time scale experiments with proteins and solvated peptides. While some results indicate possible weaknesses in ff14ipq, it is straightforward to develop hypotheses for their origins in the parameter development. It is encouraging to get this level of performance from a new force field and to see the model improve over generations of new fitting data.

5. DISCUSSION

5.1. Charges Derived by the Extended IPolQ Protocol. Before completing our molecular mechanics model based on the IPolQ charge set, we needed a set of charges compatible with the gas-phase quantum calculations available to guide our torsion parameter optimization. This, in turn, required us to revisit the way in which we derived charges, decomposing the IPolQ charges into foundational values describing the electrostatics of solutes in vacuo and perturbations describing the manner in which atoms polarize in aqueous solution.

Because of the iterative nature of the IPolQ procedure, we must consider whether this latest change will require additional iterations of molecular simulations with fixed solutes followed by charge fitting. We do not believe that this is the case, for two reasons. First, although the polarities of certain bonds have changed somewhat, the individual charges of polar side-chain atoms and overall dipoles of the molecules did not change very much. Restraints on ΔQ seem to have the strongest effects on buried atoms or groups: increasing V_{GP} reveals indeterminacy in the model fitting but is not likely to affect hydration free energies. Indeed, we chose this approach to exploit the fact that there exist many partial charge models that reproduce a molecule's electrostatic potential with nearly the same accuracy: increasing V_{GP} chooses two models that are most similar in terms of the mean-squared value of ΔQ. Second, the objective of the IPolQ procedure is a solvent reaction field potential (SRFP) that is consistent with the way a solute polarizes; the result is a set of charges that is consistent with the SRFP. Restraints on ΔQ can make minor changes to the result, but they should not change the objective. It would therefore be appropriate to apply the original IPolQ procedure until convergence and then to make other stipulations about how the model's partial charges reproduce the fitting data.

One surprising result of this analysis is that the electrostatic potentials of solutes in vacuo are harder to fit than their potentials in the condensed phase when the wave function has been computed in the presence of a polarizing charge density. This occurs in spite of the fact that the charges for polar groups, including the backbone, become significantly larger in the condensed phase to represent stronger electrostatic fields. While an analysis of why this occurs is beyond the scope of this work, it is a worthy subject for future studies and may inform the design of polarizable charge models. In another study, Zeng and colleagues introduced the "dRESP" extension of the standard RESP protocol,[47] which is similar in spirit to our extended IPolQ scheme if one takes the vacuum charges to be dRESP's baseline charge model. dRESP employs variable restraint stiffnesses on polar and nonpolar atoms to accommodate different polarizabilities, something that our arbitrary constant V_{GP} does not support. Future IPolQ derivations may benefit from an environment-dependent V_{GP}.

5.2. Torsion Fourier Series Terms in ff14ipq. We had thought that, with a complete charge parameter scheme and no obvious reasons to make further changes to bonded parameters or Lennard–Jones terms, producing viable torsion parameters would be the simplest step in developing ff14ipq. Beyond the aforementioned issues with double-counting solvent contribu-

dx.doi.org/10.1021/ct500643c | J. Chem. Theory Comput. 2014, 10, 4515−4534

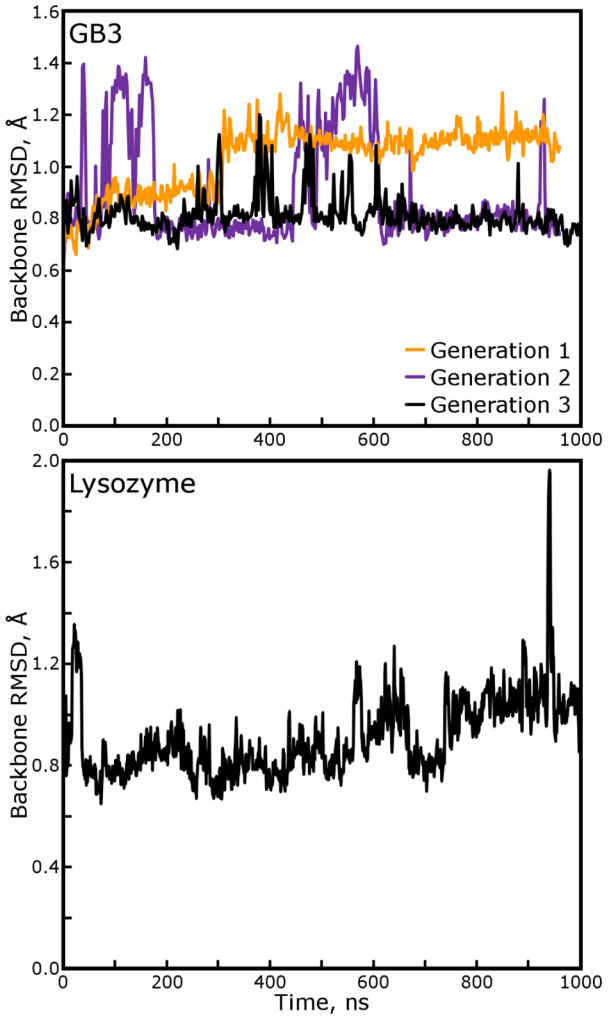

Figure 13. Backbone rmsd for globular proteins in water. GB3 and lysozyme (56 and 129 residues, respectively) were simulated with ff14ipq in baths of TIP4P-Ew water for 1 μs. Multiple simulations of GB3 are shown, each performed with a different generation of the ff14ipq torsion parameters. See Figure 3 for the accuracy of each generation in predicting energetics of lysine dipeptide.

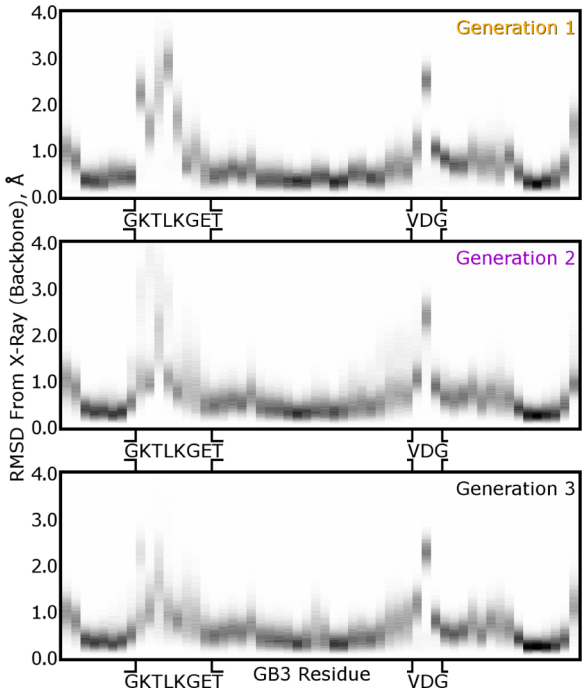

Figure 14. Per-residue backbone rmsd for GB3, with three generations of the ff14ipq model. Per-residue rmsd was calculated in the same manner as was done for Trp Cage, but the reference ensemble comprised only the one X-ray structure. Loops that make significant departures from the X-ray structure in solution-phase simulations are emphasized on the x axis.

tions, there were many hurdles in deriving a transferable torsion parameter set, and this aspect of the model development required as much effort as the charge derivation. We hope that our experience and the programs we designed to meet each challenge will benefit future ab initio force field development.

We found inclusion of new atom types and the new Fourier series terms they generated to be very helpful in fitting the vacuum-phase potential energy data. The new types introduced to accommodate the IPolQ charge set did not introduce many new parameters, as the atom types they replaced were already found in unique bonded arrangements. However, introducing atom types specific to individual residues introduced many new terms that could be fitted to the unique potential energy surface of the residue and relieve the burden that would have been placed on generic terms. In this manner, including a new atom type for C_α made improvements in the rmse for nearly all dipeptide systems, as the ϕ and ψ torsions of the uniquely achiral center of glycine could then be decoupled from those of all other, chiral amino acids. While we had wanted to include separate atom types for the C_α atoms of positively and negatively charged residues to match the unique charges derived for all backbone atoms in these residues, we were not confident in our ability to generate enough fitting data to sample the ϕ and ψ torsion space for these individual residues. In particular, sampling ϕ and ψ for consecutive arginine or lysine residues would require MP2/cc-pvTZ calculations on large tripeptide systems at very high computational cost. Expansion of the data set in this manner will likely be part of the next version of ff14ipq.

Our study also highlights the perils of adding new parameters without adequate sampling. While we believe it is correct to assign unique parameters to unique chemical environments, if the degrees of freedom governed by each parameter might take on values in simulations that are not present in the training set, then catastrophic excursions into artificial minima can occur. Data collected with the first-generation ff14ipq showed unnatural backbone ϕ angles adopted by lysine residues in the GB3 protein, which disappeared in subsequent generations trained with complete sampling of this rotatable bond. While such cases might be apparent to an especially careful investigator, we programmed the mdgx fitting routines to produce a great deal of information on the sampling of each parameter, and we did not connect higher backbone fluctuations in the GB3 protein loops to undersampling in particular parameters until significant improvements were seen in the second generation of the force field and a complete breakdown of the molecular mechanics energy was determined for each of the hundreds of lysine structures in each training set. Perhaps ore important is our finding that the descriptions of residues that do not exhibit such catastrophic artificial minima can still be improved by the simple approach of reintroducing the products of molecular simulations back into the model's training data. Whatever artificial minima were removed from the other residues by the third generation of ff14ipq must have been shallow and probably involved

dx.doi.org/10.1021/ct500643c | J. Chem. Theory Comput. 2014, 10, 4515−4534

combinations of several parameters, implying a very large conformational space to sample in order to obtain transferable parameters on the first attempt.

If allowed to drive simulations, then any molecular model should be expected to overpopulate regions of the conformational space in which it scores too favorably. However, because all of the torsion parameters discussed in this work are based on cosines, an unrefined model's propensity to overstate the favorable nature of some conformations may conversely create a tendency to overestimate barrier heights. If this is true, then it may not be a coincidence that the third generation of ff14ipq models transition states in alanine dipeptide with lower barriers than the first generation.

5.3. Consequences of Breaking the Lennard−Jones Combining Rule: The Road to ff15ipq. In our charge parameter development,[20] we were forced to alter the radii of certain polar atoms, mostly by making them larger, to adjust hydration free energies of amino acid side-chain analogues into agreement with experiment. In this work, we were forced to break the combining rule when integrating these atoms into a complete force field because the larger radii of nearby atoms would clash and make the internal potential energy surfaces of polar amino acids very hard to model. We have covered the possible implications of this decision in the K19 system, where lysine head groups appear to interact too strongly with the backbone. The special combining rules for these polar atoms would also affect the stability of backbone−backbone contacts, particularly β-sheets and α-helicies, and are likely the source of the overstabilization reported for the GB1 β-hairpins.

Other investigators have applied the IPolQ nonbonded parameters with some success. Goetz and colleagues reported strong performance for the IPolQ charge set when studying aggregation of ionic amino acids,[48] but in that system the correct result was for the amino acids to aggregate. While a prototype version of ff14ipq was able to form these aggregates with the correct radial distributions, their study did not test whether the contacts formed might, in fact, be too strong. More recently, Debiec, Gronenborn, and Chong reported that the IPolQ charge set produced the correct association constants for charged amino acid side chains when the full amino acids were considered: positively charged side chains showed affinity for both negatively charged side chains and the backbone carbonyl group, and the competition drove the equilibrium of side chains contacting one another in the right direction.[49] However, they did not rule out the possibility that the side chain to backbone contacts were themselves overstated, as our results on K19 lead us to believe. Debiec and co-workers also did not have available to them the special combining rules we adopted for ff14ipq. When these rules are applied, they found that ff14ipq overstabilized salt bridge formation between charged amino acid side-chain analogues by about 1 kcal/mol (personal communication).

We are collaborating with Debiec and others to find a pair-specific Lennard−Jones matrix that will not create excessive 1:4 strain but will properly balance the strength of polar interactions. We also intend to add new residues to the palette, including nucleic acids and phospholipids. The folding and aggregation of these residues is dominated by electrostatic interactions, which will demand the proper balance between charged groups interacting with one another and also with water. For proteins, we expect that the ff15ipq matrix of Lennard−Jones interactions will still contain special combining rules but that the departures from a standard rule will be smaller than they are in ff14ipq. Some of the Lennard−Jones interactions in IPolQ itself can be traced to limitations of the nuclear-centered charge model. For instance, the nuclear-centered charges cannot portray the hydroxyl group with as strong a polarity as it should have without making larger errors elsewhere in the electrostatic field.[20] The hydroxyl oxygen σ radius was reduced to correct alcohol hydration free energies, but this probably would not have been necessary with a virtual site scheme that could model the correct polarity. For lipids and nucleic acids, any new departures from a standard Lennard−Jones combining rule may hold clues to the inherent limitations of a nonpolarizable, nuclear-centered charge model.

■ ASSOCIATED CONTENT

Supporting Information

Program outputs from mdgx runs on charge and torsion parameter fitting; complete set of structures, Amber topologies, and corresponding MP2/cc-pVTZ energies used in torsion fitting for the third and final generation of ff14ipq. This material is available free of charge via the Internet at http://pubs.acs.org.

■ AUTHOR INFORMATION

Corresponding Authors

*(D.S.C.) E-mail: david.cerutti@schrodinger.com.

*(D.A.C.) E-mail: case@biomaps.rutgers.edu. Phone: (848) 445-5885. Fax: (732) 445-5958.

Funding

This work was supported by NIH grant GM45811.

Notes

The authors declare no competing financial interest.

■ ACKNOWLEDGMENTS

D.S. Cerutti is grateful to Professor Carlos Simmerling and Dr. James Maier for helpful discussions on torsion parameter optimization and force field testing. The authors thank Scott Legrand for timely implementation of pair-specific Lennard−Jones combining rules in the pmemd-CUDA dynamics code for graphics processors.

■ REFERENCES

(1) Case, D.; Cheatham, T.; Darden, T.; Gohlke, H.; Luo, R.; Merz, K. J.; Onufriev, A.; Simmerling, C.; Wang, B.; Woods, R. *J. Comput. Chem.* **2005**, *26*, 1668−1688.

(2) Kirschner, K.; Yongye, A.; Tschampel, S.; Gonzăles-Outeiriño, J.; Daniels, C.; Foley, B.; Woods, R. *J. Comput. Chem.* **2008**, *29*, 622−655.

(3) Buck, M.; Bouguet-Bonnet, S.; Pastor, R.; MacKerrel, A. *Biophys. J.* **2006**, *90*, L36−L38.

(4) Beauchamp, K.; Lin, Y.-S.; Das, R.; Pande, V. *J. Chem. Theory Comput.* **2012**, *8*, 1409−1414.

(5) Best, R.; Xiao Zhu, X.; Shim, J.; Lopes, P.; Mittal, J.; Feig, M.; MacKerell, A. D., Jr. *J. Chem. Theory Comput.* **2012**, *8*, 3257−3273.

(6) Klauda, J.; Venable, R.; Freites, J.; O'Connor, J.; Tobias, D.; Mondragon-Ramirez, C.; Vorobyov, I.; MacKerell, A.; Pastor, R. *J. Phys. Chem. B* **2010**, *114*, 7830.

(7) Hornak, V.; Abel, R.; Okur, A.; Strockbine, B.; Roitberg, A.; Simmerling, C. *Proteins* **2006**, *65*, 712−725.

(8) Wang, Z.; Zhang, W.; Wu, C.; Lei, H.; P, C.; Duan, Y. *J. Comput. Chem.* **2005**, *27*, 781−790.

(9) MacKerell, A., Jr. *J. Comput. Chem.* **2004**, *25*, 1584−1604.

(10) Cornell, W.; Cieplak, P.; Bayly, C.; Gould, I.; Merz, K. J.; Ferguson, D.; Spellmeyer, D.; Fox, T.; Caldwell, J.; Kollman, P. *J. Am. Chem. Soc.* **1995**, *117*, 5179−5197.

(11) Bayly, C.; Cieplak, P.; Cornell, W.; Kollman, P. *J. Phys. Chem.* **1993**, *97*, 10269−10280.

(12) Wang, J.; Cieplak, P.; Kollman, P. *J. Comput. Chem.* **2000**, *21*, 1049−1074.

(13) Dunning, T. J.; Peterson, K.; Wilson, A. *J. Chem. Phys.* **2001**, *114*, 9244−9253.
(14) Rassolov, V.; Ratner, M.; Pople, J.; Redfern, P.; Curtiss, L. *J. Comput. Chem.* **2001**, *22*, 976.
(15) Woon, D.; Dunning, T. J. *J. Chem. Phys.* **1993**, *98*, 1358−1371.
(16) Kendall, R.; Dunning, T. J.; Harrison, R. *J. Chem. Phys.* **1992**, *96*, 6796−6806.
(17) Dunning, T. J. *J. Chem. Phys.* **1989**, *90*, 1007−1023.
(18) Duan, Y.; Wu, C.; Chowdhury, S.; Lee, M.; Xiong, G.; Zhang, W.; Yang, R.; Cieplak, P.; Luo, R.; Lee, T. *J. Comput. Chem.* **2003**, *24*, 1999−2012.
(19) Zgarbová, M.; Luque, F.; Ŝponer, J.; Otyepka, M.; Jurečka, P. *J. Chem. Theory Comput.* **2012**, *8*, 3232−3242.
(20) Cerutti, D.; Swope, W.; Rice, J.; Case, D. *J. Phys. Chem. B* **2013**, *117*, 2328−2338.
(21) Karamertzanis, P.; Raiteri, P.; Galindo, A. *J. Chem. Theory Comput.* **2010**, *6*, 3153−3161.
(22) Horn, H.; Swope, W.; Pitera, J.; Madura, J.; Dick, T.; Hura, G.; Head-Gordon, T. *J. Chem. Phys.* **2004**, *120*, 9665−9678.
(23) Badyal, Y.; Saboungi, M.; Price, D.; Shastri, S.; Haeffner, D.; Soper, A. *J. Chem. Phys.* **2000**, *112*, 9206−9208.
(24) Nobes, R.; Rodwell, W.; Radom, L. *J. Comput. Chem.* **2004**, *3*, 561−564.
(25) Reif, M.; Winger, M.; Oostenbrink, C. *J. Chem. Theory Comput.* **2013**, *9*, 1247−1264.
(26) E; Simon, S.; Marquant, G.; Garcia, E.; Klimerak, G.; Delepine, J.; Cieplak, P.; Dupradeau, F.-Y. *Nucleic Acids Res.* **2011**, *39*, W511−W517.
(27) Dupradeau, F.-Y.; Pigache, A.; Zaffran, T.; Savineau, C.; Lelong, R.; Grivel, N.; Lelong, D.; Rosanski, W.; Cieplak, P. *Phys. Chem. Chem. Phys.* **2010**, *12*, 7821−7839.
(28) Frisch, M. et al. *Gaussian 09*, Revision D.01; Gaussian Inc.: Wallingford, CT, 2009.
(29) Neese, F. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2012**, *2*, 73−78.
(30) Onufriev, A.; Bashford, D.; Case, D. *Proteins* **2004**, *55*, 383−394.
(31) Lindorff-Larsen, K.; Piana, S.; Palmo, K.; Maragakis, P.; Klepeis, J.; Dror, R.; Shaw, D. *Proteins* **2010**, *78*, 1950−1958.
(32) Honda, S.; Yamasaki, K.; Sawada, Y.; Morii, H. *Structure* **2004**, *12*, 1507−1518.
(33) Ulmer, T.; Ramirez, B.; Delaglio, F.; Bax, A. *J. Am. Chem. Soc.* **2003**, *125*, 9179−9191.
(34) Walsh, M.; Schneidr, T.; Sieker, L.; Dauter, Z.; Lamzin, V.; Wilson, K. *Acta Crtystallogr., Sect. D* **1998**, *54*, 522−546.
(35) Ryckaert, J.; Ciccotti, G.; Berendsen, H.; Hirasawa, K. *J. Comput. Phys.* **1977**, *23*, 327−341.
(36) Miyamoto, S.; Kollman, P. *J. Comput. Chem.* **1992**, *13*, 952−962.
(37) Izaguirre, J.; Catarello, D.; Wozniak, J.; Skeel, R. *J. Chem. Phys.* **2001**, *114*, 2090−2098.
(38) Best, R.; Buchete, N.-V.; Hummer, G. *Biophys. J.* **2008**, *108*, L07.
(39) Graf, J.; Nguyen, P.; Stock, G.; Schwalbe, H. *J. Am. Chem. Soc.* **2007**, *129*, 1179−1189.
(40) Case, D.; Scheurer, C.; Brüschweiler, R. *J. Am. Chem. Soc.* **2000**, *122*, 10390−10397.
(41) Honda, S.; Kobayashi, N.; Munekata, E. *J. Mol. Biol.* **2000**, *295*, 269−278.
(42) Kobayashi, N.; Honda, S.; Yoshii, H.; Munekata, E. *Biochemistry* **2000**, *39*, 6564−6571.
(43) Muñoz, V.; Thompson, P.; Hofrichter, J.; Eaton, W. *Nature* **1997**, *390*, 196−199.
(44) Song, K.; Stewart, J.; Fesinmeyer, R.; Andersen, N.; Simmerling, C. *Biopolymers* **2008**, *89*, 747−760.
(45) Neidigh, J.; Fesinmeyer, R.; Andersen, N. *Nat. Struct. Biol.* **2002**, *9*, 425−430.
(46) Liu, Y.; Liu, Z.; Androphy, E.; Chen, J.; Baleja, J. *Biochemistry* **2004**, *43*, 7421−7431.
(47) Zeng, J.; Duan, L.; Zhang, J. *J. Comput. Chem.* **2013**, *34*, 847−853.
(48) Goetz, A.; Bucher, D.; Lindert, S.; McCammon, J. *J. Chem. Theory Comput.* **2014**, *10*, 1631−1637.
(49) Debiec, K.; Gronenborn, A.; Chong, L. *J. Phys. Chem. B* **2014**, *118*, 6561−6569.