Theor Chem Acc (2012) 131:1159
DOI 10.1007/s00214-012-1159-1

REGULAR ARTICLE

Polarization effects in protein–ligand calculations extend farther than the actual induction energy

Pär Söderhjelm

Received: 11 May 2011 / Accepted: 30 September 2011 / Published online: 2 March 2012

Abstract The various roles that polarizabilities play in the calculation of protein–ligand interaction energies with a polarizable force field are investigated, and the importance and distance dependence of some common approximations is determined for each of these roles separately, using quantum-mechanical calculations as the reference. It is found that the pure induction energy, if defined as the energetic gain from the charge redistribution upon interaction between the protein and ligand, is a rather short-ranged effect that becomes independent of the exact implementation at distances above $\sim$4 Å. On the other hand, the polarization between the protein residues in the assembly of a protein from separately computed fragments (as is routinely done in force field development) has a significant effect on the computed interaction energies, even for residues as far as 15 Å from the ligand. Finally, polarization improves the transferability of partial charges, but the simple polarization model used in, for example, the Amber force field explains only 14–19% of the conformational variation of the charges. In all cases, more advanced polarization models, especially involving anisotropic polarizabilities, seem to give significantly better descriptions of these effects. The study suggests that an accurate treatment of polarization can be important even in systems where the actual induction energy is small in magnitude.

Keywords Polarizable force field · Induction · Distributed polarizabilities · Amber · Intermolecular interactions · Ligand binding

1 Introduction

The reliability of a molecular simulation (e.g., a molecular dynamics or Monte Carlo simulation) is limited in at least three ways. First, the system that you simulate is typically much smaller than the real system. Second, the amount of sampling you can afford may not be enough to cover all relevant regions of phase space sufficiently well to give statistically converged results. Third, the potential-energy surface that you use in the simulation is only an approximation to the real (normally unknown) potential-energy surface.

For complex systems, such as macromolecules in water, the requirements in the two first aspects normally force you to use an empirical molecular mechanics (MM) force field as your potential-energy function. It has been noted that as computer resources are growing, the accuracy of the force field may become the limiting factor in many applications [1]. Force fields used in this area usually divide the potential energy into bonded and nonbonded terms. The bonded terms typically consist of bond, angle, and dihedral terms and describe the short-range part of the intramolecular interactions. The nonbonded terms include at least an electrostatic term and a Van der Waals term (containing dispersion and repulsion). It describes intermolecular interactions, as well as intramolecular interactions between

Published as part of the special collection of articles: From quantum mechanics to force fields: new methodologies for the classical simulation of complex systems.

Electronic supplementary material The online version of this article (doi:10.1007/s00214-012-1159-1) contains supplementary material, which is available to authorized users.

P. Söderhjelm (✉)
Department of Chemistry and Applied Biosciences, Computational Science, ETH Zürich, USI-Campus, via Giuseppe Buffi 13, 6900 Lugano, Switzerland
e-mail: par.soderhjelm@phys.chem.ethz.ch

parts of a big molecule that are not directly bound, but may be close in space.

Although it is widely recognized that electrostatics is a key to many interesting problems (e.g., protein–ligand interaction), most force fields still use a simple Coulomb interaction between atom-centered fixed partial charges. A promising route to increasing the accuracy is to explicitly include electronic polarization. The most common methods use either distributed point polarizabilities, fluctuating charges, or shell models [2–4]. Other ways to improve the electrostatics are to include higher-order multipoles (dipoles, quadrupoles, etc.) [5] or to use smeared charges that describe the charge distribution better than point charges [6].

Polarizable force fields date back to the 1960s [7] and were early applied to biological systems [8, 9], but only during the last decade, they have been incorporated into the most widely used simulation packages for macromolecules, either as polarized variants of established force fields, for example, Amber-02 and PFF [10, 11], or de novo developments such as Amoeba [12]. All of these are based on atomic isotropic dipole polarizabilities. The models differ mainly in the specific values of the polarizabilities and in the treatment of intramolecular polarization, and these choices are related.

Two early models have particularly influenced the development. Applequist showed that by allowing full coupling between all induced dipoles, it was possible to find atomic polarizabilities that reproduce the total polarizabilities of many molecules [13], but the anisotropy was often overestimated. Thole modified this model by introducing damping functions that corrected some of the error from using point dipoles at short range [14, 15]. In this way, the atomic polarizabilities could have larger and more realistic values, and the excessive anisotropy could be avoided. Roughly speaking, the Amber-02 model is based on the Applequist polarizabilities, albeit with modified coupling rules, whereas the Amoeba force field is based on the Thole model.

Other polarizable force fields, such as SIBFA, EFP, and NEMO, aim specifically at reproducing quantum-mechanical (QM) interaction energies and have a more complex functional form [5, 16, 17]. They normally derive their polarizabilities directly from QM calculations on the interacting monomers. Typically, they do not include coupling between the polarizabilities and thus have to use anisotropic atomic polarizabilities to reproduce the anisotropy of the molecular polarizabilities. Damping functions can nevertheless be used to account for the lack of Pauli effects in the intermolecular interactions [18–20].

Despite that the atomic polarizabilities obtained from QM calculations are not very similar for atoms of a certain element or atom type [21], most empirical force fields use only 8–15 atom types for polarizabilities, with 1–4 different polarizabilities for each element for the normal amino acids; some even use the same value for all non-hydrogen atoms [8]. It has been shown that improved accuracy is obtained using specific atomic polarizabilities [22]; thus, polarizabilities may well be treated in the same way as partial charges. Very recently, we showed that conformationally averaged QM-derived polarizabilities are a good choice for such atomic polarizabilities, as they give results that are fairly close to those obtained with conformation-specific polarizabilities [21].

It is difficult to directly assess the accuracy of a certain polarization model, mainly because of two reasons. First, the potential energy surface is extremely difficult to measure experimentally. One is normally limited to using quantities like solvation free energies or binding free energies and indirectly test the force field's ability to reproduce these quantities [23], which of course is made more difficult by the other simulation issues mentioned above. Even if studies have started to appear that compare the performance of nonpolarizable and polarizable force fields for biological problems [24], the precision of these studies is usually too poor to allow any detailed assessment of the polarization model itself. Alternatively, one can compare the results with QM interaction energies of high quality [25]. Very recently, an article pair has been published that systematically tries several variants of intramolecular polarization and tests the accuracy of the corresponding polarization models [26, 27]. As a reference, the authors use QM interaction energies at the MP2/aug-cc-pVTZ level. They find that all variants are better than additive force fields in reproducing the total interaction energy, but that a damping function [28] combined with exclusion of polarization through 1 or 2 bonds gives the best results.

Second, only the total potential energy is an observable; its decomposition is ambiguous. Thus, any improvements in the accuracy of the polarization term may be obscured by inaccuracies in the other terms and limits in the parametrization procedure. For example, in the above-mentioned assessment [27], the charges were adequately refitted for each polarization model, but it is still possible that the Van der Waals parameters, which are normally fitted for a particular treatment of electrostatics, unintentionally favor any of the models. One way to solve this issue is to use a specific energy decomposition scheme, for example, the restricted variational space method [29, 30], to match each contribution individually [16]. Another way to deal with the problem is to use the change in electrostatic potential as a reference, either at the monomer [31] level or the dimer level [32]. A third way is to only consider the total energy, but reassure that all the remaining terms are treated in a consistent way, for example, by

taking them from supermolecular calculations of smaller subsystems [33]. All these studies conclusively show that the inclusion of polarization can improve the accuracy of the total potential energy, but that a careful treatment is necessary for this to happen.

On the other hand, there are many issues that are common to all polarizable force fields and therefore can be studied in a more general sense. One example is whether the transferability of the partial charges between various systems and conformations is improved by the inclusion of polarization. Another relevant question (for reasons of computational efficiency) is how far out from the active site we have to keep a good description of the polarization, that is, where we can shift to a simpler model. In studies of protein–ligand interaction and protein shift of absorption spectra, we tried to answer the latter question [34, 35]. To our surprise, there was a long-range (10–20 Å) sensitivity to changes in the treatment of polarization; for example, the effect of switching from anisotropic to isotropic polarizabilities in all residues separated by more than 10 Å from the ligand was 13 kJ/mol for a charged ligand [34].

There are two contributions to this energy difference. First, there is the actual induction energy, caused by charge redistribution in the protein due to the electric field from the ligand and vice versa. Second, there is a portion of the electrostatic interaction energy between the protein and the ligand, caused by polarization of each protein residue in the electric field from other residues. The aim of the current study is to study these effects separately and characterize the error convergence of each of them. For completeness, we also consider a third effect that polarizabilities may have in force fields: we investigate whether the intramolecular polarization improves the transferability of the electrostatic model among various conformations, as earlier studies have indicated [36–39]. For all tests, we use the same test system as in ref. [34], the avidin–biotin interaction. From these tests, a more complete picture of the role of polarization in protein–ligand systems will be obtained.

2 Methods

2.1 The distributed point polarizability model

When a molecule is subjected to an electric field $\boldsymbol{F}$, the induced dipole moment within the linear-response approximation is given by

$$\boldsymbol{\mu}_{\text{ind}} = \boldsymbol{\alpha} \cdot \boldsymbol{F} \quad (\text{e.g. } \mu_x = \alpha_{xx}F_x + \alpha_{xy}F_y + \alpha_{xz}F_z) \tag{1}$$

where $\boldsymbol{\alpha}$ is called the *polarizability tensor* of the molecule and is sometimes replaced by a scalar quantity (the *isotropic polarizability*), defined by $(\alpha_{xx} + \alpha_{yy} + \alpha_{zz})/3$.

To accurately describe the polarization occurring when two molecules interact, knowledge of the molecular polarizability tensors is normally not sufficient. First, the electric field is not homogeneous, that is, it varies in different parts of the molecule. Second, the induced dipole moment is seldom a useful quantity, because it does not specify the local charge redistribution. A solution to these problems is the *distributed point polarizability model*, which is the most common model in polarizable force fields, used in QM-mimicking methods such as SIBFA, EFP, and NEMO [5, 16, 17], as well as in simpler methods [8, 10, 12]. In this model, the response of each molecule is described as a set of induced dipoles, located at certain positions (most commonly the atomic nuclei). At the position of the polarizability $\boldsymbol{\alpha}_i$, the induced dipole $\boldsymbol{\mu}_i^{\text{ind}}$ is given by:

$$\boldsymbol{\mu}_i^{ind} = \boldsymbol{\alpha}_i \cdot \boldsymbol{F}_i = \boldsymbol{\alpha}_i \cdot \left[\boldsymbol{F}_i^{stat} - \sum_{j \neq i} g_{ij} \boldsymbol{\mu}_j^{ind} \nabla\nabla \left(\frac{1}{\boldsymbol{r}_{ij}} \right) \right] \tag{2}$$

where the electric field $\boldsymbol{F}_i$ in that position has been divided into contributions from the static charge distribution and from other induced dipoles, and a (possibly distance-dependent) scale factor g_{ij} has been introduced to allow for damping or exclusion of close-lying interactions. Equation 2 defines a linear system of equations, which can be solved through matrix inversion or by iteration.

The induction energy is given by

$$E^{\text{ind}} = -\frac{1}{2} \sum_{i=1}^{N} \boldsymbol{\mu}_i^{\text{ind}} \cdot \boldsymbol{F}_i^{\text{stat}} \tag{3}$$

where the factor 1/2 comes from the fact that the (positive) work of polarization cancels half of the interaction energy of the dipoles with the field. The energy contribution from a pair of induced dipoles is completely canceled by the work of polarization, so Eq. 3 contains only the static field; the coupling is included only through the values of $\boldsymbol{\mu}_i^{\text{ind}}$.

The static field $\boldsymbol{F}_i^{\text{stat}}$ is normally computed from the multipole expansion:

$$\boldsymbol{F}_i^{stat} = -\sum_{j \neq i} h_{ij} \left[q_j \nabla \left(\frac{1}{\boldsymbol{r}_{ij}} \right) + \boldsymbol{\mu}_j \cdot \nabla\nabla \left(\frac{1}{\boldsymbol{r}_{ij}} \right) + \boldsymbol{\Theta}_j \cdot \nabla\nabla\nabla \left(\frac{1}{\boldsymbol{r}_{ij}} \right) \cdots \right] \tag{4}$$

where h_{ij} is another scale factor possibly modifying the field from close-lying multipoles. In the simplest versions of polarizable force fields, only the first term (i.e. partial charges) is used.

The reasons for letting g_{ij} or h_{ij} deviate from unity vary in the two cases. For the coupling between polarizabilities in Eq. 2, the induced dipoles may become infinite at small distances (the "polarization catastrophe"), and the

polarization can become unphysical even before that happens. Although it is possible to reproduce the molecular polarizability by a set of fully coupled atomic polarizabilities, significantly smaller values must be used than with an uncoupled set [13]. A better solution is to damp the dipole interaction tensor for overlapping charge distributions [14, 28]. In the Amber polarizable force field (ff02) [10], no damping is used, but contributions from atoms that are directly bound (1–2 interactions) or separated by two bonds (1–3 interactions) are omitted, and the 1–4 interactions are scaled by 5/6. In contrast, for most methods using polarizabilities from QM calculations, the intramolecular coupling of polarizabilities is completely omitted, because it is already implicitly included through the QM calculation. The explicit and implicit coupling schemes have been compared [40]. For medium-sized molecules such as amino acids, the implicit scheme has been found adequate (except possibly for the aromatic amino acids) [32]. A mixed approach have been proposed for fragmentation approaches [33]; this scheme is illustrated in Fig. 1 and will be denoted the *LoProp exclusion rules*.

On the other hand, for the response to the static field, which we will refer to as *static polarization*, the exclusion rules are dictated by the parametrization procedure. Clearly, when the multipoles of a molecule are obtained from a single QM calculation, h_{ij} should be zero within the molecule, as all intramolecular effects are already included in the multipoles. A corresponding rule has been derived for fragmentation approaches [33], as illustrated in Fig. 1. However, the choice of h_{ij} is arbitrary, as long as the

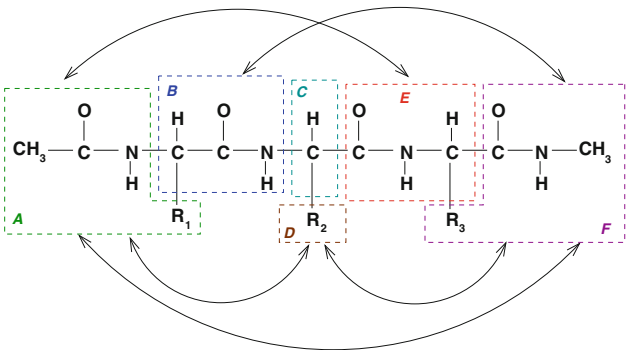

Fig. 1 Illustration of the MFCC procedure and the LoProp exclusion rules in the computation of properties (multipoles and polarizabilities) for a capped peptide. The atoms are divided into six groups. Individual QM calculations are performed on the subsystems *ABC*, *BCDE*, *CEF* and the concaps *BC* and *CE*, respectively. In all cases, the dangling bonds are capped with hydrogen atoms. The properties of the whole system are then computed as $TOT = ABC + BCDE + CEF - BC - CE$. As shown in ref. [33], intramolecular static polarization should only occur between atoms that have not been in the same QM calculation, that is, in this example only between the pairs of groups connected by *arrows*. The same rule is used for polarizability coupling

permanent multipoles are adjusted accordingly, as in, for example, the Amber ff02 model, in which $h_{ij} = g_{ij}$ [10]. For a rigid molecule, this choice does not affect the total interaction energies, whereas for a flexible molecule, the static polarization can be used to approximate how the charge distribution varies with changes in the molecular geometry [28, 36, 38, 39, 41].

The individual electrostatic and induction contributions are not comparable for models with different exclusion rules [33]. Therefore, we define the *pure induction energy* between two molecules as the energy caused by the charge redistribution of each molecule in response to the field from the *other* molecule, that is, when making the particular choice $h_{ij} = 0$ for all intramolecular interactions. The pure induction energy is always negative, in contrast to the direct difference of the energies in Eq. 3 which can have any sign [33].

Several methods to obtain distributed polarizabilities from QM calculations have been suggested [3]. A common approach is to partition the molecular polarizabilities, either in real space (e.g., the atoms-in-molecules approach) [42] or in terms of the basis set [43–46]. Alternatively, the polarizabilities can be fitted to molecular polarizabilities or to large sets of induction energies [22, 47–50]. For the most accurate descriptions in this study, we use LoProp multipoles (up to quadrupoles, $L = 2$) and anisotropic polarizabilities, located at the atomic nuclei and the midpoints of covalent bonds [46]. The multipoles are obtained by distributing the molecular charge density using an orthogonal basis of "nearly local" functions; the polarizabilities are obtained by distributing the molecular response to homogeneous electric fields using the same basis. The accuracy of the LoProp method has been tested before [31].

2.2 System

As an example of a protein–ligand interaction, we use avidin interacting with biotin. This system has previously been subject to several studies [51–54]. In particular, it has been used for testing the conformational dependence of charges [55] and polarizabilities [21] and for benchmarking the *Polarizable Multipole Interaction with Supermolecular Pairs* (PMISP) method [33, 34, 56].

For all calculations, we use a single snapshot of the avidin–biotin system, obtained from a molecular dynamics simulation using the polarizable Amber-02 force field [10]. The setup of the molecular dynamics simulation has been described before (the 02ohp simulation in ref. [54]). The geometry corresponds to one of the snapshots used to compute binding affinities with the PMISP method [56], but we removed the three extra biotin molecules that were considered part of the protein. The coordinate file used in all tests is provided as supplementary information.

The protein is cut into 494 fragments: the residues of the protein (all being standard amino acids). The cuts are done through the peptide bonds and each fragment is capped with $–COCH_3$ and $–NHCH_3$ groups at the N and C termini, respectively. The four cystine linkages are cut through the disulfide bonds and each fragment is capped with a $–SCH_3$ group.

The protein fragments are mainly taken to interact separately with the ligand, but in some calculations fragment pairs are used, some of which are covalently linked. To assemble the properties (distributed multipoles and polarizabilities) of such fragment pair, the molecular fractionation with conjugate caps (MFCC) approach [57] is used (see Fig. 1), in which a third fragment (*concap*) is constructed, consisting of the capping groups of the two fragments merged together, for example, forming a $CH_3CONHCH_3$ molecule for a peptide link. The properties of the concap fragment are subtracted from the sum of properties of the two capped fragments, and the same procedure is used for energies [33].

The distance between the ligand and a fragment is defined as the minimal distance between any ligand atom and any fragment atom. Quantities like the distance-dependent mean absolute error are computed by dividing the systems into bins of width 1 Å according to their distance; the error at a distance r then includes systems with distance between r and $r + 1$ Å.

2.3 Computational details

In the test of induction energy, the supermolecular energies are first computed at the MP2/cc-pVTZ level, whereas the properties are computed at the B3LYP/6-31G* level, as this has been found to be a good approximation [34]. For the further analysis of the error, Hartree–Fock (HF) theory with the 6-31G* basis set is used for both supermolecular energies and properties. In the test of protein assembly, the HF/6-31G* level is used throughout. The derivation of charges in the test of conformational dependence is done at the B3LYP/cc-pVTZ level to be consistent with both the procedure used in Amber ff02 [10] and with the supermolecular MP2/cc-pVTZ energies. All supermolecular energies are corrected for basis set superposition errors by the *counterpoise* procedure [58].

All QM calculations are done with the MOLCAS software [59] except the derivation of charges in the test of conformational dependence, which is done using the Gaussian software [60] followed by the RESP procedure [61], possibly employed in an iterative procedure where the AMBER 10 program [62] is used to calculate the induced dipoles and the corresponding contribution to the electrostatic potential is subtracted from the QM potential [10]. Except for this procedure, all classical interaction energies are computed by local programs that can handle all sorts of exclusion rules. The conversion from the AMBER topology file is exact; for a given system, the local software gives the same polarization energy as AMBER.

3 Results

We investigate the effect of polarizabilities on the induction energy, protein assembly, and conformational dependence, respectively. The main tests performed are schematically summarized in Fig. 2.

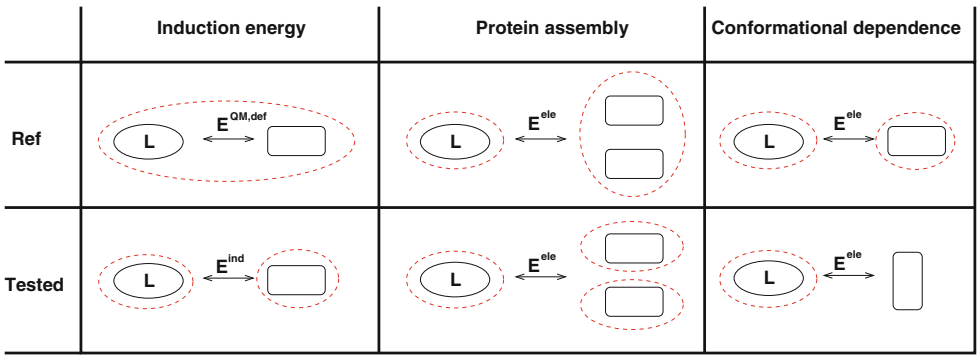

Fig. 2 Schematic overview of the three main tests performed in this article. Each column shows one test, labeled by the corresponding section heading and with the *top square* representing the reference quantity and the *bottom square* representing the tested quantity. The L-labeled *ellipse* is the ligand and the *rounded rectangle* any one of the (capped) protein residues. For the induction energy, the reference is the QM deformation energy, whereas the tested quantities are classical induction energies. For the protein assembly, the reference is the electrostatic interaction energy using multipoles computed from the residue pair (covalent or noncovalent neighbors) treated as a single molecule, whereas the tested quantity is the corresponding energy using multipoles assembled from separate residue calculations. For the conformational dependence, the reference is the electrostatic energy with the residue in its right conformation, whereas the tested quantities use charges not for the right conformation (e.g., averaged)

3.1 Induction energy

Long-range effects of approximations in the polarization model have been observed before [34, 35]. We now investigate how much of these effects are related to the pure induction energy, as opposed to the effects of the assembly of the protein from multiple fragments. The induction energy itself is inherently nonadditive. However, as the nonadditive effects have been analyzed before [33, 34], we focus instead on the interactions between the ligand and one fragment at a time.

The most natural question is how well the supermolecular energy is modeled by a given force field. In our case, we start with a very good description (a LoProp model with multipoles up to quadrupoles and anisotropic polarizabilities) and introduce approximations, one at a time. A similar test with a wider range of interactions has recently been performed for other polarization models [27].

To avoid careful parametrization of the short-range terms, we define the error as

$$R^{\text{tot}} = E^{\text{sup}} - E^{\text{ele}} - E^{\text{ind}} - E^{\text{vdw}} \tag{5}$$

where E^{sup} is the BSSE-corrected supermolecular interaction energy, E^{ele} is the electrostatic energy, E^{ind} is the induction energy, and E^{vdw} is the Van der Waals term from the Amber force field [63]. It has been shown that, for distances larger than $\sim$6 Å, the repulsion and dispersion contributions to the QM energy are well approximated by the Van der Waals term [34].

We calculated this error for all 494 fragment–ligand dimers. Unfortunately, even using a high multipole level, the errors in the electrostatic energy dominate in the whole distance range. Therefore, it was impossible to accurately assess the induction energy. For example, the isotropic approximation has a negligible effect compared with the total error, as shown in Fig. 3. Note that the fragments interacting directly with the ligand show substantial errors because of the crude short-range potential; for example, 10 fragments give an error larger than 2 kJ/mol. To proceed, we need to compare the induction energy with the particular part of the supermolecular interaction energy that depends on the deformation of the charge density, often called the deformation energy. However, this quantity is only rigorously defined at the HF level. Therefore, we first verified that the distribution of the errors (for distances larger than 6 Å) was roughly independent of the quantum-chemical method, with mean values for MP2 and HF of −0.03 and 0.06 kJ/mol, respectively, and standard deviations of 0.11 and 0.15 kJ/mol, respectively (the Van der Waals energy was only subtracted in the MP2 case as it contains mainly dispersion at these distances). At the HF level, we then define the error in the induction energy as

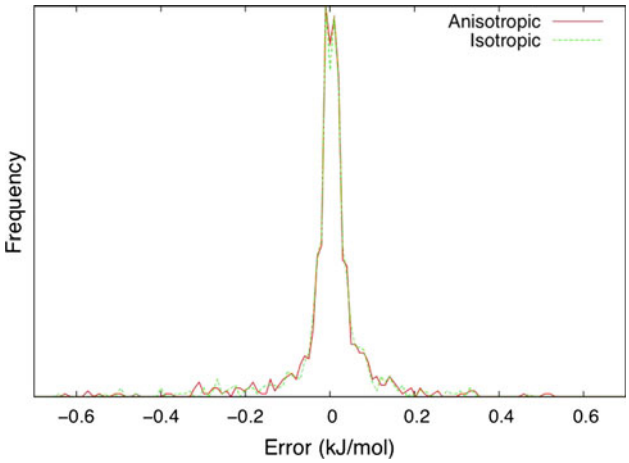

Fig. 3 Statistical distribution of the error in the total potential energy for each ligand–residue dimer computed with anisotropic and isotropic polarizabilities, respectively. The following outliers resulting from close interactions were removed from the figure: residue N12 (−10 kJ/mol error for anisotropic model), D13 (−1), Y33 (−4), T35 (−6), V37 (−1), T38 (−14), A39 (−19), T40 (−2), Q70 (3), F72 (3), S73 (−15), S75 (1), F79 (1), W97 (1), S101 (−1), N118 (−15), concap T38-A39 (−18), concap A39-T40 (−7), and concap F72-S73 (1 kJ/mol). The residue numbering refers to PDB structure 1AVD [64]

$$R^{\text{ind}} = E^{\text{sup}} - E^{\text{HL}} - E^{\text{ind}} \tag{6}$$

where E^{HL} is the Heitler–London energy, also known as the first-order energy E_1, that is, the energy obtained in a supermolecular HF calculation if the unperturbed monomer orbitals are used without any subsequent SCF iteration. Because E^{HL} describes the electrostatic and exchange-repulsion energies exactly, the remaining error R^{ind} is a good indicator of the accuracy of the induction energy. The first two terms on the right-hand side of Eq. 6 constitute the deformation energy, but it should be noted that in this study, unlike many intermolecular decomposition schemes, this quantity is corrected for BSSE.

The average error as a function of the distance from the ligand is shown in Fig. 4 for three multipole levels ($L = 0$, $L = 1$, and $L = 2$, i.e. up to charges, dipoles, and quadrupoles, respectively) and two polarizability levels (anisotropic and isotropic). Higher multipoles ($L = 3$) were also tested but showed no significant difference from $L = 2$. Detailed results for representative complexes at different distances, as well as structures of these complexes, are given in Table S1 and Figure S1 in the supplementary information.

The large errors at short distances (below 4 Å) are mainly due to the short-range effects of charge penetration, charge transfer, and other neglected effects in the point-polarizability model. Note that the errors in the final force field would possibly be smaller, because you would try to model these effects (and the corresponding electrostatic

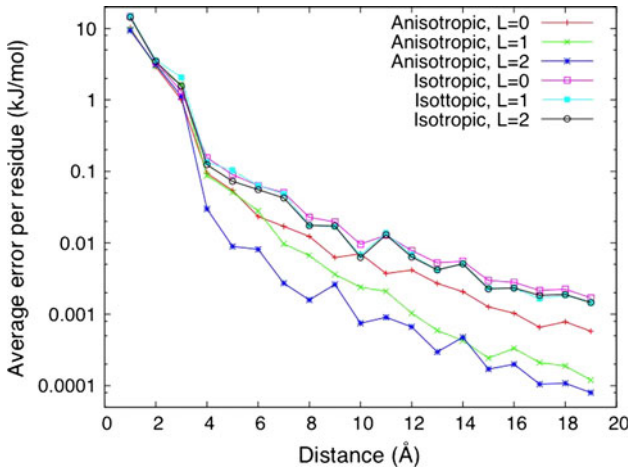

Fig. 4 Mean absolute error in the induction energy per residue as a function of the distance from the ligand

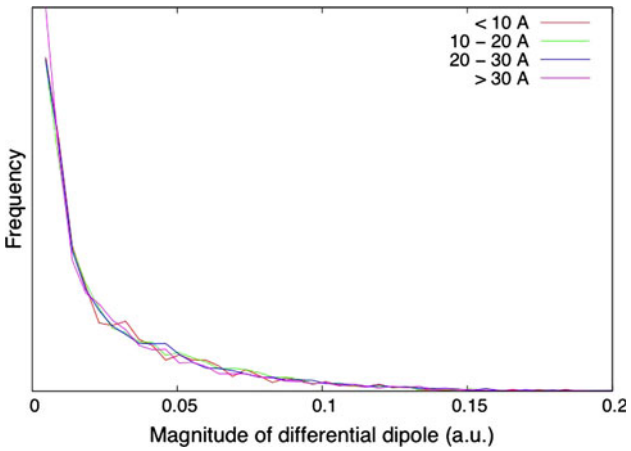

Fig. 5 Distribution of the magnitude of the difference between the induced dipole using anisotropic and isotropic polarizabilities. Each curve includes only the residues in a particular distance range from the ligand

effects) in an effective way through the repulsion term. At larger distances, on the other hand, the possibility to reduce the error by fitting disappears, and the error can be seen as a lower bound on the error of any force field, although it should be noted that the errors in the other terms are usually larger. In this far range, clear trends are obtained. As expected, anisotropic polarizabilities give lower errors than isotropic ones. With anisotropic polarizabilities, the accuracy also increases with increasing multipole level, but this effect is almost absent with isotropic polarizabilities, showing that once the isotropic approximation has been introduced, there is no need of having higher multipoles than point charges to compute the electrostatic field (although it may of course be important for the electrostatic term).

However, the main result from Fig. 4 is that, even for the worst model using point charges and isotropic polarizabilities, the errors quickly become negligible as the distance from the ligand increases. Thus, we can conclude that the pure induction is not responsible for the observed long-range effects of polarizabilities.

3.2 Protein assembly

A force field for a protein is usually assembled using parameters for smaller fragments, taken either directly from QM calculations or from a library. If the force field is polarizable, each fragment is polarized by the field from the surrounding fragments so that its charge distribution within the protein differs from that of the isolated fragment. This change, in turn, affects the electrostatic interaction with the ligand. Because this is an indirect effect (many-body effect), one would expect it to be smaller than the pure induction. However, electric fields in the protein are often strong due to the proximity of charged residues. Thus, the statically induced dipoles may become significant and interact strongly with the charged ligand, as these interactions has a formal r^{-2} dependence (compared to r^{-4} for the ion–induced dipole interaction).

Of course, this effect depends on the polarization model used. A comparison of the individual induced dipoles in the assembled avidin protein (without the ligand) modeled by anisotropic and isotropic polarizabilities, respectively, shows that the deviation is randomly distributed with average magnitude of 0.026 a.u. (corresponding to $\sim$80% of the average magnitude of the induced dipoles themselves) and independent of the distance from the ligand, as displayed in Fig. 5. Although the energetic effect of each of these individual differences is randomly distributed and rather small, in average 0.3 kJ/mol at 5 Å and 0.1 kJ/mol at 15 Å, the large number of contributions add up to a total energy contribution that is typically 3–6 kJ/mol for the residues outside of 15 Å and 4–11 kJ/mol outside 5 Å, depending on the particular geometry and QM method (results not shown). This is in agreement with Fig. 2 of ref. [34].

To investigate this effect more systematically, we use smaller subsystems consisting of only two fragments (each fragment being one capped amino acid) and the ligand (see Fig. 2). Two sets of fragment pairs were created. The first set consists of the 1,008 fragment pairs that have at least one atom–atom distance that is <2.5 Å, but are not covalently linked. The second set consists of the 494 fragment pairs that are directly covalently linked (eight of which share a cystine link, the rest a peptide bond). We test how well various polarization models describe the charge redistribution within each fragment pair upon association. As a measure of the charge redistribution for a given polarization model M, we calculate the change in

electrostatic interaction energy between the fragment pair (F_i, F_j) and the ligand:

$$\Delta E_{ij}^{M} = E_{\text{ele}}(L \leftrightarrow F_{ij}^{M}) - E_{\text{ele}}(L \leftrightarrow F_{i}^{M}) - E_{\text{ele}}(L \leftrightarrow F_{j}^{M}) \tag{7}$$

where $E_{\text{ele}}(L \leftrightarrow F_i^M)$ denotes the classical electrostatic interaction energy between the ligand and F_i. Whereas the ligand is always treated in the same way, with distributed multipoles up to quadrupoles, obtained from a LoProp calculation, we vary the representation M of the fragments. This variation includes the source of the polarizabilities (from a quantum-chemical LoProp calculation or from the Amber library), the form of the polarizabilities (anisotropic or isotropic; in the bonds or only in the atomic nuclei), the origin of the multipoles (LoProp multipoles up to quadrupoles or Amber charges), and the choice of exclusion rules for the polarizability coupling (LoProp or Amber). A *null model* without polarization is also included for comparison. All the tested models are listed in Table 1.

The fragment dimer F_{ij} is constructed by taking the two separate fragments described by M and letting them polarize each other. For the covalently bound pairs, the MFCC procedure is used, so that a concap term is added to Eq. 7 (see Sect. 2).

As a reference for these calculations, Eq. 7 is evaluated using multipoles for the fragment pair computed in a *single* QM calculation (again by the LoProp approach and up to quadrupoles). The mean unsigned error

$$R^{M} = \frac{1}{N}\sum_{i<j}^{N} |\Delta E_{ij}^{M} - \Delta E_{ij}^{\text{ref}}| \tag{8}$$

for each method over each set of dimers is shown in Table 2. To verify that the trends are not influenced by short-ranged effects, the average over only the distant pairs (minimal distance from the ligand larger than 5 Å) is also shown. The full distance dependence for some of the methods is shown in Fig. 6 for the nonbonded set.

The setup of these calculations ensures that we specifically test the error in the polarization part of the assembly. As expected, the most accurate treatment (*a*), using LoProp multipoles and anisotropic polarizabilities in both atoms and bonds, gives the lowest error, from ~0.04 kJ/mol for distant pairs up to 0.5 kJ/mol at short range. Replacing the anisotropic polarizabilities with their isotropic counterparts has a significant effect, increasing the error by 50–65%. In agreement with the results for the pure induction, the further approximation to use point charges instead of multipoles has a smaller effect, and these two effects are almost additive. The removal of the polarizabilities in the bonds (by dispersing them onto the atoms) has a negligible effect.

The change from LoProp (*ix0*) to Amber (*ambf*) polarization increases the error by 66–74 %. This change can be divided into two steps: the change of the values of the polarizabilities and the change of exclusion rules (the g_{ij} in Eq. 2). As shown in Table 2, both steps give significant (and almost additive) contributions to the error, but the change of values has the largest effect. The Amber polarizabilities were not devised to reproduce QM calculations, as has been pointed out before [21]. Therefore, the results for the Amber model are not alarming. Although there is much room for improvement—the error is three times larger than for the best model—the Amber model still gives significantly better results than the null model, which assembles the pairs without considering polarization at all. It should be noted that, although the mean absolute error per pair is <1 kJ/mol for all models, the many pairs may add up to a substantial total error, for example, 70 kJ/mol for the null model.

Table 1 Summary of the polarization variants used to test the protein assembly

Name	Polarizabilities		Multipoles	Exclusion rule
	Source	Anisotropic		
a	LoProp	Yes	LoProp	LoProp
i	LoProp	No	LoProp	LoProp
a0	LoProp	Yes	Amber	LoProp
i0	LoProp	No	Amber	LoProp
ix0	LoProp (no bonds)	No	Amber	LoProp
ix0f	LoProp (no bonds)	No	Amber	Amber
amb	Amber	No	Amber	LoProp
ambf	Amber	No	Amber	Amber
null	No polarization			

The polarizabilities can come from LoProp or Amber ff02 and can be anisotropic or isotropic. The multipoles used in the polarization can be either LoProp multipoles (up to quadrupoles) or Amber ff02 charges. The exclusion rule can be either LoProp (no intramolecular polarization) or Amber (only 1–2 and 1–3 interactions omitted). For comparison, a null model without polarization is also included

Table 2 Mean absolute error in ΔE (Eq. 8; thousandths of kJ/mol) for the two sets of fragment pairs (the nonbonded pairs and the covalent pairs) using various polarization methods for treating the charge redistribution within the pair

Averages are taken over all pairs (all) or those outside of 5 Å (far). Numbers within brackets are the corresponding errors including also the induction energy

Method	Nonbonded pairs		Covalent pairs	
	All (N = 1,008)	Far (N = 922)	All (N = 494)	Far (N = 452)
a	73 (86)	48 (48)	72 (79)	47 (47)
i	112 (213)	79 (82)	89 (164)	56 (58)
a0	117	75	76	50
i0	137	97	97	62
ix0	129	93	99	62
ix0f	162	117	110	65
amb	198 (343)	134 (137)	97 (243)	62 (67)
ambf	225 (374)	154 (158)	114 (288)	73 (78)
null	503 (1,325)	347 (384)	189 (1,049)	129 (183)

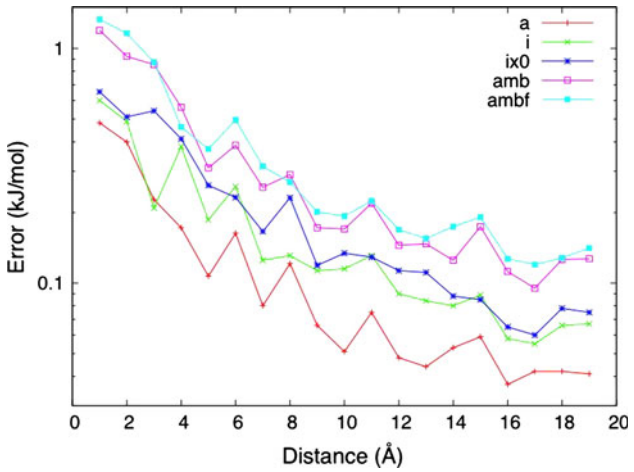

Fig. 6 Mean absolute error (Eq. 8) in the assembly of the nonbonded pairs as a function of the minimal distance between any of the pair fragments and the ligand

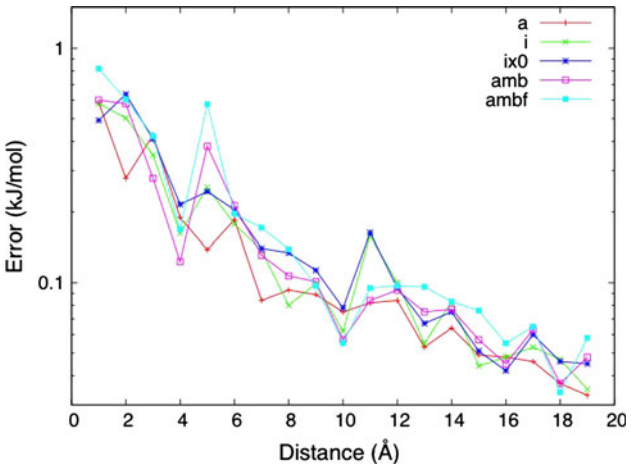

Fig. 7 Mean absolute error (Eq. 8) in the assembly of the covalent pairs as a function of the minimal distance between any of the pair fragments and the ligand

The fact that the LoProp exclusion rules give better results than the Amber exclusion rules, independently of the values of the polarizabilities, offers some physical insight. As described in Fig. 1, the coupling between polarizabilities within each fragment is ignored with the LoProp exclusion rules, because the coupling is implicitly included under the assumption that the electric field is uniform over the fragment. The Amber rules, on the other hand, try to model the coupling explicitly by only excluding coupling between atoms separated by one or two bonds. The results indicate that, for fragments that are as small as amino acids, it is better to ignore the coupling between polarizabilities.

The corresponding distance dependence for the covalently bound set is shown in Fig. 7. For this set, the results are less clear, but most of the trends remain. As expected, the treatment of intramolecular polarization is more difficult: all models give larger errors relative to the null model (e.g., 36% over the distant covalent pairs for the best model, compared to 14% for the distant nonbonded pairs; see Table 2). On the other hand, the polarization effects for the assembly of covalently linked residues are smaller in absolute value and thus less important than those in, for example, hydrogen bonds.

For some of the methods, the errors including the actual induction energy between the ligand and the fragment pairs are also shown in Table 2. As can be seen, the additional error for the far set introduced by the induction energy is negligible for all models (except the null model, which does not include the induction energy). This confirms the result from the previous section that the approximations affect the induction energy only at short range.

3.3 Conformational dependence

In the final test, we examine the possible advantage of using a polarizable model to enhance the transferability of partial charges between various conformations of the same amino acid. To this end, we derived RESP charges for all capped residues (in their particular conformation) of

Table 3 Mean absolute error per residue (kJ/mol) for various charge sets over all 494 residues (all) or the 453 residues with distance >6 Å (far), using either the *qm0* results or supermolecular (super) energies as the reference

Reference	qm0 (all)	qm0 (far)	Super (far)
qm0	–	–	0.10
qm1 + pol	0.05	0.03	0.10
cons0	0.31	0.18	0.22
cons1 + pol	0.25	0.15	0.18
ff02 + pol	1.28	1.00	1.03

avidin, using an iterative procedure [10] that makes them consistent with the Amber ff02 polarizabilities. For comparison, we also computed RESP charges at the same QM level but without polarization. Thus, for each capped residue, we have two descriptions that nearly reproduce the same QM electrostatic potential: one with only partial charges (*qm0*) and one with partial charges and isotropic polarizabilities (*qm1*). Two other charge sets (*cons0* and *cons1*) were constructed by averaging the *qm0* and *qm1* charges, respectively, over all occurrences of a given amino acid in the protein. Finally, we include the standard Amber-02 charges (*ff02*). There is no standard charge set for a nonpolarizable model at this exact QM level. In all cases involving polarization, we employ the Amber-02 polarizabilities [10]. For the ligand, we always use the same charges (*qm0*) and no polarization.

We use the electrostatic interaction energy between the ligand and each residue described by *qm0* as the reference and report the mean absolute errors over all residues, or over residues with a distance >6 Å, in Table 3. First, we verify that the *qm0* and *qm1* descriptions give the same interaction energies. As shown in Table 3, the mean absolute difference between the interaction energies with these models is only 0.05 kJ/mol. Thus, before averaging, the nonpolarizable (*qm0*) and polarizable (*qm1*) descriptions are roughly equivalent.

Using the averaged charges instead of those derived for exactly the right conformation gives an average error of 0.25–0.31 kJ/mol over all residues and 0.15–0.18 kJ/mol over the distant residues. Interestingly, the error is consistently lower for the polarizable model (by 14–19%), showing that the polarization accounts for some of the conformational dependence of the charges. However, earlier investigations have reported much larger improvements for more advanced polarization models [36, 38, 39], so the Amber model is not optimal. Clearly, the conformational dependence of charges (and higher multipoles) is an unsolved problem.

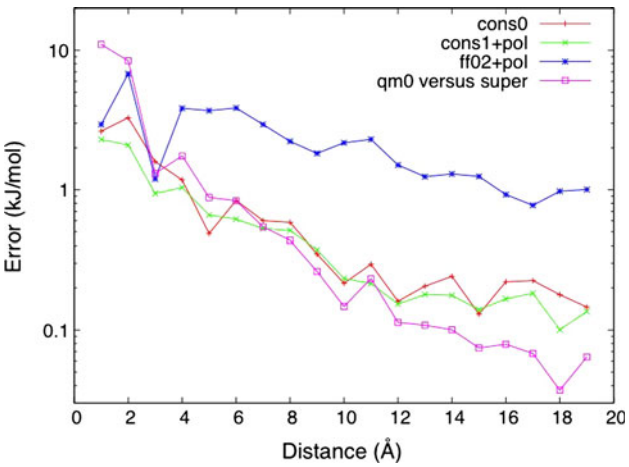

Fig. 8 Mean absolute error per residue (kJ/mol) introduced by the averaging of charges over conformations, with (*cons1*) and without (*cons0*) an additional polarization model, as well as with the standard *ff02* charges. For comparison, the *qm0 versus super* curve shows the deviation between the supermolecular MP2/cc-pVTZ interaction energy and the electrostatic interaction energy with the *qm0* charges

To see how the accuracy relates to the overall force field accuracy, we also compare these electrostatic interaction energies with the supermolecular MP2/cc-pVTZ interaction energies, with the Amber Van der Waals energy subtracted as in Sect. 3.1. The results are shown in Table 3 for the far set (the comparison is not relevant for shorter distances because several terms are replaced by the standard Van der Waals energy). As expected, the errors are larger, but not by much.

The errors with the standard charges from the *ff02* protein library are much larger: the average error is 1.3 kJ/mol over all residues and 1.0kJ/mol over the distant residues. This is somewhat unexpected, as one would expect the *ff02* charges to be similar to the *cons1* charges, as they are derived in similar ways. However, a detailed analysis shows that most of the error comes from the capping $-CH_3$ groups, which are significantly more negative (in both the reference calculations and the QM calculations) than the corresponding protein atoms that *ff02* tries to model.

The full distance dependence of the errors is shown in Fig. 8. For short distances, all the charge sets give large errors, although the use of polarization (*cons1*) seems to avoid the largest problems. At a distance of $\sim$4 Å, the errors with *cons0* and *cons1* charges quickly drop below 1 kJ/mol per residue, whereas the *ff02* error remains large. The comparison with supermolecular results shows that for distances larger than $\sim$7 Å, the conformational averaging is a larger source of inaccuracy than the point charge representation itself.

4 Conclusions

Based on the results of these tests, one may look at the computation of a protein–ligand interaction energy with a polarizable force field as a three-step process, with the polarization playing different roles in each step.

First, each protein fragment is internally polarized. Force fields often assume that different conformations of the same amino acid have the same partial charges. Polarization enables a physically motivated variation of the electrostatic properties of different conformations. We tested this effect by constructing artificial Amber-like polarizable and nonpolarizable force fields with conformationally averaged partial charges and comparing the resulting electrostatic interaction energies with those using charges for the right conformation. Indeed, with the Amber polarization model, the error caused by conformational averaging is 14–19% smaller with polarization than without.

Second, each protein fragment is polarized by the surrounding fragments, so that the electrostatic properties are altered compared with what they would be in isolated amino acids. We found that approximations in the polarization model have a large effect on the induced dipoles in the whole protein and thus on the electrostatic interaction energy with the ligand, even for fragments as far as 15 Å from the ligand (e.g., $\sim$0.1 kJ/mol per residue for the isotropic approximation). By investigating the assembly of residue pairs using a series of polarization models, it was found that polarization between residues interacting by, for example, hydrogen bonds is more important than between covalently linked residues, but on the other hand it is also easier to model. Upon switching from a very accurate model to the Amber model, the steps particularly found to increase the error are the removal of anisotropy and the change of the values of the polarizabilities. However, the introduction of intramolecular coupling of the polarizabilities also has a consistently negative impact on the accuracy. Encouragingly, all tested polarization models were found to give significantly better results than a nonpolarizable model.

Finally, the protein is polarized by the ligand and vice versa, giving the pure induction energy. By comparing the induction energy from various polarization models with its quantum-mechanically computed counterpart, we found that the approximation to make the polarizabilities isotropic has a larger effect than, for example, the reduction of the multipole level, but that approximations done in residues separated by more than $\sim$4 Å from the ligand have a negligible effect on the energy compared with the total error of the force field.

These results indicate that a careful treatment of polarization may be important even in cases where the actual induction energy is small. They also provide some guidance how to improve current polarization models. Apparently, anisotropic polarizabilities and rigorous exclusion rules are essential to achieve quantitative agreement with QM calculations.

Acknowledgments Financial support from the Swedish research council (623-2009-821) is greatly acknowledged.

References

1. Ponder JW, Case DA (2003) Adv Protein Chem 66:27
2. Rick SW, Stuart SJ (2002) Potentials and algorithms for incorporating polarizability in computer simulations. In: Lipowitz KB, Boyd DB (eds) Reviews in computational chemistry, Wiley, New York, p 89
3. Cieplak P, Dupradeau FY, Duan Y, Wang JM (2009) J Phys Condens Matter 21(33):ARTN 333102
4. Warshel A, Kato M, Pisliakov AV (2007) J Chem Theory Comput 3:2034
5. Engkvist O, Åstrand P-O, Karlström G (2000) Chem Rev 100:4087
6. Piquemal J-P, Cisneros GA, Reinhardt P, Gresh N, Darden TA (2006) J Chem Phys 124:104101
7. Pullman B, Claverie P, Caillet J (1967) Proc Natl Acad Sci 57(6):1663
8. Warshel A, Levitt M (1976) J Mol Biol 103:227
9. Gresh N, Pullman B (1980) Biochim Biophys Acta 608(1):47–53
10. Cieplak P, Caldwell J, Kollman P (2001) J Comput Chem 22:1048
11. Maple JR, Cao Y, Damm W, Halgren TA, Kaminski GA, Zhang LY, Friesner RA (2005) J Chem Theory Comput 1:694
12. Ponder J, Wu C, Ren P, Pande VS, Chodera JD, Schnieders MJ, Haque I, Mobley DL, Lambrecht DS, DiStasio RA Jr (2010) J Phys Chem B 114(8):2549–2564
13. Applequist J, Carl JR, Fung K-K (1972) J Am Chem Soc 94:2952
14. Thole BT (1981) Chem Phys 59:341
15. Van Duijnen PT, Swart M (1998) J Phys Chem A 102(14):2399–2407
16. Gresh N, Cisneros GA, Darden TA, Piquemal J-P (2007) J Chem Theory Comput 3:1960
17. Gordon MS, Freitag MA, Bandyopadhyay P, Jensen JH, Kairys V, Stevens WJ (2001) J Phys Chem A 105:293
18. Gresh N (1995) J Comput Chem 16(7):856–882
19. Day PN, Jensen JH, Gordon MS, Webb SP, Stevens WJ, Krauss M, Garmer D, Basch H, Cohen D (1996) J Chem Phys 105:1968
20. Masia M, Probst M, Rey R (2005) J Chem Phys 123:164505
21. Söderhjelm P, Kongsted J, Ryde U (2011) J Chem Theory Comput 7:1404–1414
22. Elking D, Darden T, Woods RJ (2007) J Comput Chem 28(7):1261–1274
23. Shirts MR, Pitera JW, Swope WC, Pande VS (2003) J Chem Phys 119:5740
24. Jiao D, Golubkov PA, Darden TA, Ren P (2008) Proc Natl Acad Sci 105:6290–6295
25. Xantheas SS, Burnham CJ, Harrison RJ (2002) J Chem Phys 116:1493
26. Wang JM, Cieplak P, Li J, Hou TJ, Luo R, Duan Y (2011) J Phys Chem B 115(12):3091–3099
27. Wang JM, Cieplak P, Li J, Wang J, Cai Q, Hsieh MJ, Lei HX, Luo R, Duan Y (2011) J Phys Chem B 115(12):3100–3111
28. Ren P, Ponder JW (2002) J Comput Chem 23:1497
29. Stevens WJ, Fink WH (1987) Chem Phys Lett 139:15

30. Bagus PS, Hermann K, Bauschlicher CWJ (1984) J Chem Phys 80:4378
31. Söderhjelm P, Krogh JW, Karlström G, Ryde U, Lindh R (2007) J Comput Chem 28:1083
32. Söderhjelm P, Öhrn A, Ryde U, Karlström G (2008) J Chem Phys 128:014102
33. Söderhjelm P, Ryde U (2009) J Phys Chem A 113:617
34. Söderhjelm P, Aquilante F, Ryde U (2009) J Phys Chem B 113:11085
35. Söderhjelm P, Husberg C, Strambi A, Olivucci M, Ryde U (2009) J Chem Theory Comput 5:649
36. Engkvist O, Åstrand P-O, Karlström G (1996) J Phys Chem 100:6950
37. Tiraboschi G, Fournié-Zaluski MC, Roques BP, Gresh N (2001) J Comput Chem 22(10):1038–1047
38. Holt A, Karlström G (2008) J Comput Chem 29:1084
39. Holt A, Karlström G (2008) J Comput Chem 29:1905
40. Nakagawa S, Mark P, Ågren H (2007) J Chem Theory Comput 3:1947
41. Rasmussen TD, Ren P, Ponder JW, Jensen F (2007) Int J Quantum Chem 107:1390
42. Ángyán JG, Jansen G, Loos M, Hättig C, Heß BA (1994) Chem Phys Lett 219:267
43. Garmer DR, Stevens WJ (1989) J Phys Chem 93:8263
44. Stone AJ (1985) Mol Phys 56:1065
45. Le Sueur CR, Stone AJ (1993) Mol Phys 78:1267
46. Gagliardi L, Lindh R, Karlström G (2004) J Chem Phys 121:4494
47. Stout JM, Dykstra CE (1995) J Am Chem Soc 117(18):5127–5132
48. Celebi N, Ángyán JG, Dehez F, Millot C, Chipot C (2000) J Chem Phys 112:2709
49. Williams GJ, Stone AJ (2003) J Chem Phys 119:4620
50. Soteras I, Curutchet C, Bidon-Chanal A, Dehez F, Ángyán JG, Orozco M, Chipot C, Luque FJ (2007) J Chem Theory Comput 3(6):1901–1913
51. Miyamoto S, Kollman PA (1993) Proteins Struct Funct Genet 16:226
52. Wang J, Dixon R, Kollman PA (1999) Proteins Struct Funct Genet 34:69
53. Kuhn B, Kollman PA (2000) J Med Chem 43:3786
54. Weis A, Katebzadeh K, Söderhjelm P, Nilsson I, Ryde U (2006) J Med Chem 49:6596
55. Söderhjelm P, Ryde U (2009) J Comput Chem 30:750
56. Söderhjelm P, Kongsted J, Ryde U (2010) J Chem Theory Comput 6:1726
57. Zhang DW, Zhang JZH (2003) J Chem Phys 119:3599
58. Boys SF, Bernardi F (1970) Mol Phys 19:553
59. Molcas 7, University of Lund, Sweden (2007) See http://www.teokem.lu.se/molcas
60. Frisch MJ, Trucks GW, Schlegel HB, Scuseria GE, Robb MA, Cheeseman JR, Scalmani G, Barone V, Mennucci B, Petersson GA, Nakatsuji H, Caricato M, Li X, Hratchian HP, Izmaylov AF, Bloino J, Zheng G, Sonnenberg JL, Hada M, Ehara M, Toyota K, Fukuda R, Hasegawa J, Ishida M, Nakajima T, Honda Y, Kitao O, Nakai H, Vreven T, Montgomery JA Jr, Peralta JE, Ogliaro F, Bearpark M, Heyd JJ, Brothers E, Kudin KN, Staroverov VN, Kobayashi R, Normand J, Raghavachari K, Rendell A, Burant JC, Iyengar SS, Tomasi J, Cossi M, Rega N, Millam JM, Klene M, Knox JE, Cross JB, Bakken V, Adamo C, Jaramillo J, Gomperts R, Stratmann RE, Yazyev O, Austin AJ, Cammi R, Pomelli C, Ochterski JW, Martin RL, Morokuma K, Zakrzewski VG, Voth GA, Salvador P, Dannenberg JJ, Dapprich S, Daniels AD, Farkas Ö, Foresman JB, Ortiz JV, Cioslowski J, Fox DJ (2009) Gaussian 09 revision A.1. Gaussian Inc., Wallingford, CT
61. Bayly CI, Cieplak P, Cornell WD, Kollman PA (1993) J Phys Chem 97:10269
62. Case DA, Darden TA, Cheatham TE III, Simmerling CL, Wang J, Duke RE, Luo R, Crowley M, Walker RC, Zhang W, Merz KM, Wang B, Hayik S, Roitberg A, Seabra G, Kolossvary I, Wong KF, Paesani F, Vanicek J, Wu X, Brozell SR, Steinbrecher T, Gohlke H, Yang L, Tan C, Mongan J, Hornak V, Cui G, Mathews DH, Seetin MG, Sagui C, Babin V, Kollman PA (2008) Amber 10. University of California, San Francisco
63. Cornell W, Cieplak P, Bayly C, Gould I, Merz KM, Ferguson D, Spellmeyer D, Fox T, Caldwell J, Kollman P (1995) J Am Chem Soc 117:5179
64. Pugliese L, Coda A, Malcovati M, Bolognesi M (1993) J Mol Biol 231:698