Published in final edited form as:

Comput Phys Commun. 2014 March ; 185(3): 720–729. doi:10.1016/j.cpc.2013.10.028.

A biomolecular electrostatics solver using Python, GPUs and boundary elements that can handle solvent-filled cavities and Stern layers

Christopher D. Cooper^a, Jaydeep P. Bardhan^b, and L. A. Barba^{a,*}

^aMechanical Engineering, Boston University, Boston, MA, 02215 U.S.A

^bElectrical and Computer Engineering, Northeastern University, Boston, MA, 02115 U.S.A

Abstract

The continuum theory applied to biomolecular electrostatics leads to an implicit-solvent model governed by the Poisson-Boltzmann equation. Solvers relying on a boundary integral representation typically do not consider features like solvent-filled cavities or ion-exclusion (Stern) layers, due to the added difficulty of treating multiple boundary surfaces. This has hindered meaningful comparisons with volume-based methods, and the effects on accuracy of including these features has remained unknown. This work presents a solver called PyGBe that uses a boundary-element formulation and can handle multiple interacting surfaces. It was used to study the effects of solvent-filled cavities and Stern layers on the accuracy of calculating solvation energy and binding energy of proteins, using the well-known apbs finite-difference code for comparison. The results suggest that if required accuracy for an application allows errors larger than about 2% in solvation energy, then the simpler, single-surface model can be used. When calculating binding energies, the need for a multi-surface model is problem-dependent, becoming more critical when ligand and receptor are of comparable size. Comparing with the apbs solver, the boundary-element solver is faster when the accuracy requirements are higher. The cross-over point for the PyGBe code is in the order of 1–2% error, when running on one gpu card (nvidia Tesla C2075), compared with apbs running on six Intel Xeon cpu cores. PyGBe achieves algorithmic acceleration of the boundary element method using a treecode, and hardware acceleration using gpus via PyCuda from a user-visible code that is all Python. The code is open-source under MIT license.

Keywords

biomolecular electrostatics; implicit solvent; Poisson-Boltzmann; boundary element method; treecode; Python; CUDA

* New address: Mechanical and Aerospace Engineering, George Washington University, Washington, D.C., 20052, labarba@gwu.edu. cdcooper@bu.edu (Christopher D. Cooper), jbardhan@ece.neu.edu (Jaydeep P. Bardhan)

1. Introduction

Many vital biomolecular processes are dominated by electrostatic forces, such as solvation and binding. These processes occur in an aqueous environment, and it is crucial to include the effect of water (the solvent) in any model. In biology, salt ions in the solvent always play a role in the electrostatic potential field, and they must also be considered in the model. Thus, simulation of biomolecular systems is a challenging task, in particular if one aims for a detailed description of all molecules in the system, such as in molecular dynamics simulations. In some situations, the systems of interest are too large for a detailed description using molecular dynamics; in that case, we can look at mean-field potentials instead, using an implicit-solvent model [1, 2].

Implicit-solvent models rely on the solution of the Poisson-Boltzmann equation. Research through the years on the numerical solution of this equation has led to several codes that are open and widely used in the community, such as `APBS` [3], `Delphi` [4], `MEAD` [5, 6], `MIBPB` [7], `AFMPB` [8], `TABI` [9], among others. The most widely used are volumetric-based solvers, meaning that they create a discretization of the volume in which the equations apply. On the other hand, boundary-element methods (`BEM`), which discretize only surfaces, have become a popular alternative in biomolecular applications [10, 11, 12] and several `BEM` codes have been developed in recent years [8, 9, 13, 14].

In `BEM` solvers, it is more difficult to consider features like solvent-filled cavities and Stern layers around biomolecules, compared with volumetric-based solvers. In fact, it appears that apart from the work of Altman and co-workers [15], no other `BEM` implementation addresses solvent-filled cavities and Stern layers. Consequently, the literature does not explain the importance of including these features in a model. Moreover, without boundary-integral formulations that treat these features, comparisons between `BEM` and volumetric-based methods are not meaningful, because the latter always include them.

In this paper, we study the effect of including solvent-filled cavities and Stern layers in a Poisson-Boltzmann `bem` solver. This allows us to make a fair comparison with a volumetric-based method, in this case `APBS`, in order to determine in which cases a boundary integral approach may be preferable. To perform the study, we developed `PyGBe`, a treecode-accelerated boundary element solver [16]. `PyGBe` is written in Python and uses `PyCUDA` [17] to exploit gpu hardware for shorter runtimes. The code is free and open-source under an MIT license, and we welcome interactions with the community.

2. Methods

2.1. Implicit-solvent model

The implicit-solvent model describes the electrostatic field in a molecular system using continuum theory. We define a solvent region (water with salt) and a protein region, separated by the solvent-excluded surface (SES) [18]. The solvent region has a high dielectric constant (permittivity ~ 80) and is modeled by the linearized Poisson-Boltzmann equation to account for the presence of ions. The protein region has low permittivity (~ 2 –

4) and contains point charges at the locations of the atoms, and it is modeled by the Poisson equation.

Figure 1 depicts the two regions. The SES represents the closest a water molecule can get to the protein. It is defined by rolling a probe sphere of the size of a water molecule (1.4 Å) around the protein, described by the atoms with their corresponding van der Waals radius. This process is sketched in Figure 2.

Using the notation from Figure 1, the electrostatic potential can be described by the following system of partial differential equations:

$$\begin{aligned} \nabla^2 \phi_1(\mathbf{r}) &= -\sum_k \frac{q_k}{\epsilon_1} \delta(\mathbf{r}, \mathbf{r}_k) & \text{in protein region } \Omega_1 \\ \nabla^2 \phi_2(\mathbf{r}) &= \kappa^2 \phi_2(\mathbf{r}) & \text{in solvent region } \Omega_2 \\ \phi_1 &= \phi_2 & \text{on interface } \Gamma \\ \epsilon_1 \frac{\partial \phi_1}{\partial \mathbf{n}} &= \epsilon_2 \frac{\partial \phi_2}{\partial \mathbf{n}} \end{aligned} \quad (1)$$

where q_k is the charge associated to atom k and κ is the inverse of the Debye length. The boundary conditions in (1) enforce continuity of the potential and normal electric displacement through the interface Γ .

Integral formulation—If we take the convolution of the PDEs in (1) with their corresponding free-space Green's functions, then evaluate the expressions (via integration by parts) on the interface Γ applying the continuity conditions, the system is transformed into a system of integral equations along Γ , as shown by Yoon and Lenho [10]:

$$\begin{aligned} \frac{\phi_1(\mathbf{r})}{2} - \oint_{\Gamma} \frac{\partial \phi_1}{\partial \mathbf{n}}(\mathbf{r}') \frac{1}{4\pi|\mathbf{r}-\mathbf{r}'|} d\Gamma' + \oint_{\Gamma} \frac{\partial}{\partial \mathbf{n}} \left[\frac{1}{4\pi|\mathbf{r}-\mathbf{r}'|} \right] \phi_1(\mathbf{r}') d\Gamma' &= \sum_{k=0}^{N_c} \frac{q_k}{\epsilon_1} \frac{1}{4\pi|\mathbf{r}-\mathbf{r}_k|}, \\ \frac{\phi_1(\mathbf{r})}{2} + \frac{\epsilon_2}{\epsilon_1} \oint_{\Gamma} \frac{\partial \phi_1}{\partial \mathbf{n}}(\mathbf{r}') \frac{e^{-\kappa|\mathbf{r}-\mathbf{r}'|}}{4\pi|\mathbf{r}-\mathbf{r}'|} d\Gamma' - \oint_{\Gamma} \frac{\partial}{\partial \mathbf{n}} \left[\frac{e^{-\kappa|\mathbf{r}-\mathbf{r}'|}}{4\pi|\mathbf{r}-\mathbf{r}'|} \right] \phi_1(\mathbf{r}') d\Gamma' &= 0. \end{aligned} \quad (2)$$

Here, $\frac{1}{4\pi|\mathbf{r}-\mathbf{r}'|}$ and $\frac{e^{-\kappa|\mathbf{r}-\mathbf{r}'|}}{4\pi|\mathbf{r}-\mathbf{r}'|}$ are the free-space Green's functions of the and linearized Poisson-Boltzmann equations, respectively. The position vectors $\mathbf{r}$ and $\mathbf{r}'$ are defined along the boundary. The singularity is treated via a limiting process in the classic way [19].

Solvation energy—The electrostatic contribution to the free energy of solvation—the solvation energy—is a common parameter of interest that can be computed from the electric potential. It is defined as the energy required to move the protein from a reference state, often a uniform dielectric, i.e., ϵ_1 in the exterior region, to the solvated state. The total electrostatic potential is the sum of the direct Coulomb potential and the *reaction potential*, which arises due to the presence of the two dielectric materials:

$$\phi_1(\mathbf{r}) = \phi_{Coul}(\mathbf{r}) + \phi_{reac}(\mathbf{r}). \quad (3)$$

Since ϕ_{Coul} is exactly the right hand side of the top equation in (2), we can compute ϕ_{reac} anywhere in Ω_1 using the following integral relation:

$$\phi_{reac}(\mathbf{r}) = \oint_{\Gamma} \frac{\partial \phi_1}{\partial \mathbf{n}}(\mathbf{r}') \frac{1}{4\pi|\mathbf{r}-\mathbf{r}'|} d\Gamma' - \oint_{\Gamma} \frac{\partial}{\partial \mathbf{n}} \left[\frac{1}{4\pi|\mathbf{r}-\mathbf{r}'|} \right] \phi_1(\mathbf{r}') d\Gamma'. \quad (4)$$

We obtain (4) using the same derivation described for Equation (2). However, in this case, $\mathbf{r}$ is not taken to the interface Γ , but is valid anywhere in the domain Ω_1 , and ϕ_{Coul} is subtracted out.

The energy due to any potential is:

$$\Delta G = \int_{\Omega} \rho \phi d\Omega, \quad (5)$$

where ρ is the charge distribution. In this case, solvation energy is generated by ϕ_{reac} and the integral can be computed analytically since ρ is a collection of point charges. Equation (5) becomes:

$$\Delta G_{solv} = \frac{1}{2} \sum_{k=1}^{N_c} q_k \phi_{reac}(\mathbf{r}_k). \quad (6)$$

Binding energy—We can calculate the electrostatic contribution to the binding free energy —referred to as binding energy—with the following relation:

$$\Delta \Delta G_{bind} = \Delta G_{complex} - \Delta G_{receptor} - \Delta G_{ligand}, \quad (7)$$

where G_i is the total electrostatic energy of species i :

$$\Delta G_i = \Delta G_{i,solvation} + \Delta G_{i,Coulomb}. \quad (8)$$

where $G_{solvation}$ is calculated as described above and $G_{Coulomb}$ corresponds to the Coulombic energy of the point charges.

We approximate the binding energy of a complex assuming that the geometry of the receptor and the ligand remain constant from the unbound to the bound states.

2.2. Boundary element method (BEM)

Single-surface problems—In BEM, we discretize the boundary into panels as sketched in Figure 3, and assume a distribution of the unknown quantities within those panels. We use flat triangular panels and take ϕ_1 and $\frac{\partial \phi_1}{\partial \mathbf{n}}$ to be constant on each one of them. Then, we use central collocation to obtain a linear system. The discretized form of (2) becomes

$$\begin{aligned} \frac{\phi_1(\mathbf{r}_i)}{2} - \sum_{j=1}^{N_p} \frac{\partial \phi_1}{\partial \mathbf{n}}(\mathbf{r}_j) \int_{\Gamma_j} \frac{1}{4\pi|\mathbf{r}_i-\mathbf{r}'|} d\Gamma' + \sum_{j=1}^{N_p} \phi_1(\mathbf{r}_j) \int_{\Gamma_j} \frac{\partial}{\partial \mathbf{n}} \left[\frac{1}{4\pi|\mathbf{r}_i-\mathbf{r}'|} \right] d\Gamma' &= \sum_{k=0}^{N_c} \frac{q_k}{\epsilon_1} \frac{1}{4\pi|\mathbf{r}_i-\mathbf{r}_k|} \\ \frac{\phi_1(\mathbf{r}_i)}{2} + \sum_{j=1}^{N_p} \frac{\partial \phi_1}{\partial \mathbf{n}}(\mathbf{r}_j) \frac{\epsilon_1}{\epsilon_2} \int_{\Gamma_j} \frac{e^{-\kappa|\mathbf{r}_i-\mathbf{r}'|}}{4\pi|\mathbf{r}_i-\mathbf{r}'|} d\Gamma' - \sum_{j=1}^{N_p} \phi_1(\mathbf{r}_j) \int_{\Gamma_j} \frac{\partial}{\partial \mathbf{n}} \left[\frac{e^{-\kappa|\mathbf{r}_i-\mathbf{r}'|}}{4\pi|\mathbf{r}_i-\mathbf{r}'|} \right] d\Gamma' &= 0, \end{aligned} \quad (9)$$

where $\mathbf{r}_i$ corresponds to the center of panel Γ_i , the collocation panel.

After collocating on all panels, we obtain the following $2N \times 2N$ linear system of equations:

$$\begin{bmatrix} \frac{1}{2}I + K_L & -V_L \\ \frac{1}{2}I - K_Y & \frac{\epsilon_1}{\epsilon_2}V_Y \end{bmatrix} \begin{bmatrix} \phi_1 \\ \frac{\partial \phi_1}{\partial \mathbf{n}} \end{bmatrix} = \begin{bmatrix} Q \\ 0 \end{bmatrix}, \quad (10)$$

where,

$$\begin{aligned} K_{L,ij} &= \int_{\Gamma_j} \frac{\partial}{\partial \mathbf{n}} \left(\frac{1}{|\mathbf{r}_{ij}|} \right) d\Gamma & V_{L,ij} &= \int_{\Gamma_j} \frac{1}{|\mathbf{r}_{ij}|} d\Gamma \\ K_{Y,ij} &= \int_{\Gamma_j} \frac{\partial}{\partial \mathbf{n}} \left(\frac{\exp(-\kappa|\mathbf{r}_{ij}|)}{|\mathbf{r}_{ij}|} \right) d\Gamma & V_{Y,ij} &= \int_{\Gamma_j} \frac{\exp(-\kappa|\mathbf{r}_{ij}|)}{|\mathbf{r}_{ij}|} d\Gamma \\ Q_i &= \sum_{k=0}^{N_c} \frac{q_k}{\epsilon_1} \frac{1}{|\mathbf{r}_{ik}|}. \end{aligned} \quad (11)$$

Here, V and K are the discretized versions of the single- and double-layer potentials, respectively, for the case of constant elements. The unknown vector in (10) corresponds to ϕ_1 and $\frac{\partial \phi}{\partial \mathbf{n}}$ evaluated at the center of each panel. Once we solve the linear system, we can use the discretized form of (4) to find ϕ_{reac} , the potential due to the solvent polarization, at the locations of q_k . We then compute the solvation energy from (6).

Multi-surface problems—Following Altman et al. [15], PyGBe has the capacity to solve problems that involve multiple surfaces, such as when there are ion-exclusion (Stern) layers and solvent-filled cavities (depicted in Figure 4).

An ion-exclusion or Stern layer is a 1 to 2 Å-thick layer around the solvent-excluded surface that has the solvent's permittivity and no ions, and the potential is described by the Laplace equation. The Stern layer compensates for the fact that the linearized Poisson-Boltzmann equation overestimates the concentration of ions near charged surfaces. This is an artifact of having a continuous representation of ions which does not consider steric effects, particularly important at charged interfaces.

In the folded state, proteins may have cavities inside that are big enough to contain solvent molecules. Inside these cavities, the permittivity is that of the solvent, and the potential obeys the Laplace equation. If, in addition, the cavity is large enough to hold ions, it can have its own interior Stern layer as well, and is governed by the linearized Poisson–Boltzmann equation.

Volumetric solvers include solvent-filled cavities and Stern layers automatically in the grid definition. But in boundary-integral formulations, each region adds a surface that generates two extra equations, making the linear system of (10) more complicated. In the design of PyGBe, we followed the method of Altman et al. [15] to automatically determine the correct boundary-integral formulation for problems with multiple surfaces. Readers are referred to that work for the details.

2.3. Fast summation via treecode

The treecode is an algorithm that reduces the complexity from $O(N^2)$ to $O(N \log N)$ for N -body calculations like evaluating the following sum in all points i :

$$V(\mathbf{x}_i) = \sum_{j=1}^N q_j \psi(\mathbf{x}_i, \mathbf{y}_j). \quad (12)$$

It is based on clustering the “sources” $\mathbf{y}_j$ in boxes of an oc-tree, and performing approximations when the “target” $\mathbf{x}_i$ is far away. The effect of the far-away sources is approximated using a series expansion about the cluster center. In this work, we use Taylor series for the expansions, writing for ψ :

$$\psi(\mathbf{x}_i, \mathbf{y}_j) = \sum_{\|\mathbf{k}\|=0}^P \frac{1}{\mathbf{k}!} D_{\mathbf{y}}^{\mathbf{k}}(\mathbf{x}_i, \mathbf{y}_j) \psi(\mathbf{y}_j - \mathbf{y}_c)^{\mathbf{k}}. \quad (13)$$

Here, $\mathbf{y}_c$ is the center of the box, P the order of the expansion, $D_{\mathbf{y}}^{\mathbf{k}} = D_{y_1}^{k_1} D_{y_2}^{k_2} D_{y_3}^{k_3}$ the derivative operator, and $\mathbf{k} = (k_1, k_2, k_3)$, $\mathbf{k}! = k_1! k_2! k_3!$, $\mathbf{y}^{\mathbf{k}} = y_1^{k_1} y_2^{k_2} y_3^{k_3}$.

By rearranging terms and using (13) in (12), we get the approximation of $V(\mathbf{x}_i)$ as

$$V(\mathbf{x}_i) = \sum_{\|\mathbf{k}\|=0}^P \frac{1}{\mathbf{k}!} \underbrace{D_{\mathbf{y}}^{\mathbf{k}} \psi(\mathbf{x}_i, \mathbf{y}_c)}_{a^{\mathbf{k}}} \underbrace{\sum_j q_j (\mathbf{y}_j - \mathbf{y}_c)^{\mathbf{k}}}_{m_c^{\mathbf{k}}}, \quad (14)$$

where $a^{\mathbf{k}}$ are the coefficients of the Taylor expansion and $m_c^{\mathbf{k}}$ the multipoles. Equation (14) shows that the calculations involving the “sources” $\mathbf{y}_j$ can be decoupled from calculations involving the “targets” $\mathbf{x}_i$, which has the effect of reducing the complexity from $O(N^2)$ to $O(N \log N)$.

The octree is built stipulating that boxes at the lowest level—called twigs or leaves—should have less than n_{crit} sources. Once built, the tree is traversed for each target starting from the root box (corresponding to the whole domain) querying if the box is far enough to use the approximation from (14), or else, to look at the child boxes. This process is recursively repeated, ending at the lowest level of the tree, where the interactions are computed directly. The criterion to decide if a box is far enough from a target particle to use the series is called the multipole acceptance criterion (MAC), given by

$$\theta > \frac{r_b}{r_{tb}}, \quad (15)$$

where r_b is the size of the box and r_{tb} the target-box distance. We use the approximation in (14) when the criterion in (15) is met, with the value of θ usually chosen as either $\frac{1}{2}$ or $\frac{2}{3}$.

The multipole terms $m_c^{\mathbf{k}}$ in (14) are calculated first at the twigs of the tree, and then these values are shifted and added as necessary to obtain the multipole terms of all boxes in upper levels. The next step is to traverse the tree and compute the rest of Equation (14) where the mac is met. For this step, we need to calculate the coefficients of the Taylor expansion $a^{\mathbf{k}}$, using recursive relations for the Green's function of the Poisson equation [20] and linearized

Poisson-Boltzmann equation [21]. The final step is to directly compute interactions for boxes that are too close to targets in order to meet the MAC.

Treecode-accelerated BEM—A BEM solver using an iterative method for the linear system of equations requires multiple dense matrix-vector multiplications, which can be accelerated to $O(N \log N)$ with a treecode. Each element of the BEM matrix is an integral that can be approximated with Gauss quadrature, and we can formulate a matrix-vector product as a N -body problem like the one shown in (12), where quadrature points serve as sources, collocation points as targets and the function ψ is the free-space Green's function of the Laplace or linearized Poisson-Boltzmann equations.

Integrating accurately near the singularity of the free-space Green's function poses a major challenge in BEM. PyGBe computes singular integrals with a semi-analytical technique presented by Zhu et al. [22] and defines a near-singular region where it places more Gauss quadrature points. In practice, the near-singular region is always contained in an area where the MAC criterion (15) is not met, so special integration using more Gauss points and the semi-analytical techniques are implemented in the direct-interaction step of the treecode.

PyGBe uses a Krylov method for the solution of the linear system. The solver needs to perform one evaluation of the treecode per iteration, but uses the same source and target points. Because of this, we need to compute the tree structure and interaction list only once.

2.4. Richardson extrapolation

Richardson extrapolation is a fundamental technique used in numerical computing to obtain an estimated exact solution from a series of computations using consecutive resolutions from coarser to finer. To be correctly applied, all of the computations used need to be converging to the exact value at a constant rate, meaning, they are in the asymptotic range. Under these conditions, an estimate of the exact solution can be written as:

$$f_{exact} \approx f_1 + \frac{f_1 - f_2}{r^p - 1}, \quad (16)$$

where f_1 and f_2 are the fine-grid and coarse-grid solutions, respectively; r is the mesh refinement ratio (h_2/h_1 , where h can be spacing, area or volume) and p is the order of convergence.

From the calculations alone, we cannot ensure that the chosen grids are in the asymptotic range. Before using Equation (16), then, one should check that f_1 and f_2 are in that range. We can extract the *observed* order of convergence from three grid resolutions refined with constant ratio r , that is: $r = h_2/h_1 = h_3/h_2$. In that case:

$$p = \frac{\log \left(\frac{f_3 - f_2}{f_2 - f_1} \right)}{\log(r)}. \quad (17)$$

If the result from (17) matches the *expected* order of convergence of the method, it indicates that numerical computations for f_1, f_2 and f_3 are in the asymptotic range. For this reason, we

need in general three calculations in the asymptotic convergence range to obtain an extrapolated value with Richardson extrapolation.

2.5. Verification and Reproducibility

We previously verified the PyGBe code using analytical solutions and published a technical report [16]. The Barba research group has a consistent policy of making science codes freely available, in the interest of reproducibility. In line with this, we release the entire code that was used to produce the results of this study. PyGBe is made available under the MIT open-source license and we maintain a version-control repository.¹ We also release all the files needed for running the numerical experiments reported in this paper: input, configuration, and postprocessing. To support our open-science and computational reproducibility goals, in addition to open-source sharing of the code and the running scripts, several of the plots themselves are also available and usable under a CC-BY license, as indicated in the respective captions.

3. Results using PyGBe and comparison with APBS

Rationale for the experiments and setup

In the communities of biochemistry and biophysics, APBS is a trusted software for biomolecular electrostatics, actively developed since the late 1990s [3]. The APBS code uses a finite-difference solution to the Poisson-Boltzmann equation on a Cartesian volumetric mesh, with more recent versions of the code providing a finite-element solution [23]. Our objective was to test PyGBe using actual molecular geometries, comparing the results with those obtained using APBS. The comparisons address two questions: the accuracy that is obtained with these two different models and codes, and the conditions under which PyGBe can be competitive in terms of computational effort.

We set up the PyGBe and APBS runs using the parameters listed in Table 1, unless otherwise noted, with APBS running on 6 cores of an Intel Xeon X5650 cpu and PyGBe running on one NVIDIA Tesla C2075 GPU card. To generate meshes, we used the free MSMS software [24] and to determine charges and van der Waals radii, we used the free software pdb2pqr [25].

In Table 1, ϵ_1 and ϵ_2 represent the permittivity of the protein and solvent regions, respectively, κ is the inverse of the Debye length, K is the number of Gauss points per triangular element, P is the order of the Taylor expansion used for the treecode and $\text{GMRES}_{\text{tol}}$ is the tolerance of the Krylov iterative solver.

The tests used three molecular systems: (i) a lysozyme molecule, (ii) a trypsin-BPTI complex; and (iii) a peptide-RNA complex, all considered with and without cavities and/or Stern layer. See Table 2 for the protein-data-bank entries. In addition, we performed convergence tests using Richardson extrapolation with both PyGBe and APBS, for which the test problem was a spherical molecule with an embedded charge.

¹<https://bitbucket.org/cdcooper/pygbe>

3.1. Convergence tests with both PyGBe and APBS

These results show the power of Richardson extrapolation. Using a spherical molecule of radius $4\text{\AA}$ with a centered charge and no ions in the solvent, we computed the electrostatic potential using APBS and PyGBe with a variety of meshes and obtained the solvation energy. The parameters for these runs are detailed in Table 3 and the results are shown in Figure 5. The data sets, figure files and plotting scripts for Figure 5 are available under a *CC-BY* license [27].

The plot in Figure 5a corresponds to the solvation energy obtained with different meshes using the two codes. The analytical solution for our spherical test is -41.2465 [kJ/mol] and is represented by a dotted line on this plot. The extrapolated values calculated using Equation (16) are -41.2471 [kJ/mol] with APBS and -41.2467 [kJ/mol] with PyGBe. These are both very close to the analytical solution and would require a very fine mesh to be obtained in a calculation with each code.

Figure 5b shows the error in solvation energy obtained with PyGBe and APBS , calculated using the corresponding extrapolated value as the base solution. The plot shows how error scales with the average area of the boundary elements for PyGBe and with mesh spacing for APBS . The observed order of convergence calculated with Equation (17) was 1.014 with respect to mesh spacing for APBS and 1.013 with respect to area for PyGBe. These values match the expected order of convergence for these methods, and thus we are confident that the computations are in the asymptotic convergence range.

3.2. Lysozyme

Calculating solvation energy for a lysozyme molecule (code 1HEL) with different mesh refinements, we obtained the results shown in Figure 6. This molecule is modeled with 1323 point charges, has 3 cavities, and we also considered a $2\text{\AA}$ -thick Stern layer. The structure for lysozyme was the only one not prepared using *pdb2pqr*; instead, we used the data from a previous publication [28]. The calculation parameters are detailed in Table 4, where: K_{fine} is the number of Gauss points in the near-singular region, N_k the number of Gauss points per side of the triangle for the semi-analytical integral technique used for the singular elements, thres the distance from the singularity defining the near-singular region—with L being a characteristic length of the integration panel—and θ the multipole-acceptance criterion of the treecode. The data sets, figure files and plotting scripts for Figure 6 are available under a *CC-BY* license [29].

Figure 6a shows the computed solvation energy for different mesh sizes. As the mesh density increases, the discretization elements become smaller and the solution tends towards an asymptote, represented by the dotted line, which we calculated with Richardson extrapolation. This value is -2070.47 [kJ/mol] for APBS and -2082.5 [kJ/mol] for PyGBe using the full model (with cavities and Stern layer).

The errors plotted in Figure 6b were calculated using the corresponding extrapolated values as the base solution. The fact that the error is scaling with area for PyGBe and mesh spacing for APBS proves that those calculations are in the regime with expected convergence, and that the error is dominated by the discretization. It is important to mention that mesh density for

Δ_{PBS} corresponds to number of cells per unit volume, whereas for PyGBe it is number of boundary elements per unit area, thus we cannot directly compare the discretizations, and we simply used these metrics for graphical purposes to have both cases in one plot.

We measured the time to solution and plotted the comparison between Δ_{PBS} and PyGBe for different errors in Figure 6c, using the errors from Figure 6b. For low-accuracy calculations, Δ_{PBS} is faster than PyGBe; however, Δ_{PBS} has worse scaling, and they cross-over near the 2%-error mark. This indicates that a boundary element approach is more suitable when there is a need for more accuracy.

The results shown in Figure 6d examine the importance of considering cavities and Stern layers. They correspond to a mesh-refinement study of solvation energy with Lysozyme modeled in four different ways: considering only the dielectric interface (“Single”), the dielectric interface plus cavities (“Cavity”), dielectric interface with Stern layer and no cavities (“Stern”) and including cavities and Stern layer (“Full”). For the most refined mesh, the difference between the “Single” result and the “Full” results is 35 [kJ/mol]. The Richardson extrapolated value using the single-surface model is -2047.16 [kJ/mol], which is also around 35 [kJ/mol] away from the extrapolated value for the “Full” simulation.

3.3. Trypsin-BPTI

Using the method explained in Section 2.1, we obtained the binding energy of the trypsin-BPTI complex (code 3BTK) using both PyGBe and Δ_{PBS} (BPTI stands for Bovine Pancreatic Trypsin Inhibitor). We first computed the solvation energy for each case and then added the Coulombic energy separately. The Coulombic energies for the trypsin-BPTI complex, receptor and ligand are shown in Table 5.

The PyGBe runs were done with two models: a multi-surface model, where the protein is modeled considering cavities and a $2\text{\AA}$ -thick Stern layer, and a single-surface model, where only the dielectric interface or SES is considered. The parameters used in these calculations are listed in Table 6. The data sets, figure files and plotting scripts for the results in this section are made available under a CC-BY license [30].

Trypsin-BPTI complex—Figure 7 shows the results of computing solvation energy for the trypsin-BPTI complex, which we later use for obtaining binding energy. We modeled the trypsin-BPTI complex using 4074 point charges with an Amber force field, and it contains 7 cavities.

The solvation-energy values obtained with different mesh densities are plotted in Figure 7a. The dotted lines represent the extrapolated values obtained using Richardson extrapolation, and shown in Table 7 for each case: Δ_{PBS} , multi-surface PyGBe and single-surface PyGBe. We calculated the errors shown in Figure 7b using the extrapolated values in Table 7 as the base solution. The errors seem to be scaling with mesh spacing for Δ_{PBS} and triangle area for PyGBe, at the expected rates.

Figure 7c shows a plot of time to solution with respect to estimated error. It displays the same behavior as the lysozyme molecule in Figure 6c, where for low-accuracy calculations a

volumetric approach is faster, but scales worse than the boundary integral approach, crossing over at a level of about 3% error. Of course, the single-surface model requires less elements and takes less time to solve compared to the multi-surface model.

Trypsin—Figure 8 shows results of solvation energy for trypsin, which we later use for calculating binding energy. We model the trypsin using 3220 point charges with an Amber force field, and it contains 6 cavities.

Convergence of the solvation energy with mesh density is plotted in Figure 8a. The dotted lines represent the extrapolated values using Richardson extrapolation, listed in Table 7. We calculated the errors shown in Figure 8b using the extrapolated values in Table 7 as the base solution. The errors seem to be scaling with mesh spacing for $_{APBS}$ and with area for PyGBe, at the expected rates.

The plot of time to solution versus error in Figure 8c shows the same behavior as the lysozyme molecule in Figure 6c and the trypsin- $_{BPTI}$ complex in Figure 7c, where time to solution for the boundary element method has better scaling with error.

Bovine pancreatic trypsin inhibitor ($_{BPTI}$)—Figure 9 shows results of solvation energy calculations for bovine pancreatic trypsin inhibitor ($_{BPTI}$), which is modeled using 854 point charges with a Amber force field, and no cavities.

The plot of solvation energy for different values of mesh density in on Figure 9a, where the dotted lines represent the extrapolated value using Richardson extrapolation, shown in Table 7. We calculated the errors shown in Figure 9b using the extrapolated values in Table 7 as the base solution. The errors seem to be scaling with mesh spacing for $_{APBS}$ and area for PyGBe, in the expected rates. Figure 9c shows the plot of time to solution vs. error. As before, time to solution for the boundary element method has better scaling with error.

Trypsin- $_{BPTI}$ binding energy—Binding energies can be computed from the solvation energies as detailed in Equations (7) and (8). Figure 10 shows the computed binding energy for different values of mesh density. Using the extrapolated values of solvation energy for PyGBe with multi- and single-surface models, shown in Table 7, we calculated binding energies. These were 74.42 [kJ/mol] for the multi-surface model and 78.15 [kJ/mol] for the single-surface model, and are represented in Figure 10 by the dotted lines.

3.4. Peptide- $_{RNA}$ complex—We performed the same set of tests than in Section 3.3 but for a 22-residue long α -helical peptide of protein λ bound the “box B” $_{RNA}$ hairpin structure [26]. In this case, the complex has 998 atoms, the $_{RNA}$ 619 atoms and the peptide 379 atoms, and the structures are available online.² Neither molecule has cavities and PyGBe runs used the multi-surface model with a 2Å-thick Stern layer. Figure 11 shows the binding energy of this complex for different mesh densities using $_{APBS}$ and PyGBe with a multi-surface and single-surface model. The dotted lines correspond to binding energies calculated using the extrapolated values of solvation energy for the multi-surface model (93.7 [kJ/mol]) and

²<http://www.poissonboltzmann.org/apbs/examples/>

single-surface model (101.29 [kJ/mol]). The data sets, figure files and plotting scripts for the results in this section are made available under a CC-BY license [31].

4. Discussion

The results presented in the previous section show that the extrapolated values for PyGBe and APBS are converging to values that differ by $\sim 0.5\%$. Some difference is to be expected, because APBS is a volumetric, Cartesian-mesh solver that constrains the molecular-surface definition to a staircase representation, quite different from the surface-mesh representation used in PyGBe. Also, for volumetric approaches, point charges inside the molecule in general do not coincide with the mesh, and the required interpolation introduces smoothening of the point charges.

Figure 6d shows the impact of including cavities and Stern layers in solvation-energy calculations, comparing results for lysozyme —a molecule with 3 cavities— modeled with and without these features. The extrapolated values of the “Single” and “Full” models differ by 35[kJ/mol], which is a $\sim 2\%$ difference. This suggests that, for calculating solvation energy, if the required accuracy allows errors larger than 2%, then the simpler, single-surface model can be used, requiring less discretization elements and less runtimes.

In the lysozyme test, we found also that the solvation energy is more affected by the presence of cavities than a Stern layer (Figure 6d). And as expected, results with APBS match the multi-surface PyGBe model better, since APBS considers a Stern layer and cavities automatically.

Although 35 [kJ/mol] may be a small amount in the context of solvation energies, it may become important for derived quantities, for example, binding energies. Binding energy results from the difference between two solvation energies, which may be small. In the trypsin-BPTI complex, the binding energy difference between multi- and single-surface models was only 4 [kJ/mol], or $1.5RT$, with R being the gas constant and RT representing thermal fluctuations. This value is very close to thermal fluctuations, and corresponds to only $\sim 5\%$ of the total binding energy, indicating that for this case a single-surface model is good enough. On the other hand, the peptide-RNA complex of Section 3.4 shows a larger difference in binding energy when a Stern layer is considered (the components here have no cavities) of around 8 [kJ/mol] or $3RT$, which corresponds to $\sim 10\%$ of the total binding energy. Our results suggest that if the size difference between ligand and receptor is large, like in the trypsin-BPTI complex, the resulting binding energy is less sensitive to Stern layers or cavities. However, for more comparable sizes in the molecules, like the peptide-RNA complex, using all features becomes important. We think this happens because the receptor looks almost identical to the complex when it is much larger than the ligand, making the error introduced by using a single-surface model in the solvation energy calculation very similar for the receptor and complex, and being subtracted out in the binding energy calculation of Equation (7).

Our results in Section 3 also include time to solution. Timings for PyGBe do not include the mesh generation time with MSMS, this time is negligible compared to the total time to solution. The trend of the time to solution plots is very similar in all cases, where lower-

accuracy calculations are faster using a volumetric approach, but the scaling of the boundary integral technique with accuracy is better and they cross over at around 1% error. This indicates that the choice of solver should be made considering the required accuracy of the application.

5. Conclusion

In this paper, we studied the use of a boundary element method to solve biomolecular systems in the context of implicit-solvent models. We noted that BEM, compared to volumetric-based methods, has difficulties to model situations where more than one surface is required to include all features of the biological system. The importance of including those features has not been treated before in the literature. We found that the answer is problem-dependent, and we noticed that, for example, in binding-energy calculations involving a ligand and receptor of comparable sizes, cavities and Stern layers should be included. Our work also included a comparison of the time-to-solution between *APBS* and *PyGBe*. The results suggest that the choice of computational tool should be based on the desired accuracy, and when an error smaller than 1% is desired, a boundary element approach is more adequate.

The results for this paper using *PyGBe* were obtained with version e34c338 of the code, available from the version-controlled repository at <https://bitbucket.org/cdcooper/pygbe>. Input files, meshes and plotting scripts are also available [27, 29, 30, 31].

Acknowledgments

This work was supported by ONR via a grant from the Applied Computational Analysis Program, # N00014-11-1-0356. LAB also acknowledges support from NSF CAREER award OCI-1149784 and from NVIDIA, Inc. via the CUDA Fellows Program. The work of JPB has been supported in part by the National Institute of General Medical Sciences (NIGMS) of the National Institutes of Health (NIH) under award number R21GM102642. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

References

1. Roux B, Simonson T. Implicit solvent models. *Biophys. Chem.* 1999; 78:1–20. [PubMed: 17030302]
2. Bardhan JP. Biomolecular electrostatics — I want your solvation (model). *Comput. Sci. Disc.* 5:013001.
3. Baker NA, Sept D, Holst MJ, McCammon JA. Electrostatics of nanosystems: Application to microtubules and the ribosome. *P. Natl. Acad. Sci. USA.* 2001; 98:10037–10041.
4. Gilson MK, Sharp KA, Honig B. Calculating the electrostatic potential of molecules in solution: method and error assessment. *J. Comp. Chem.* 1987; 9(4):327–335.
5. Bashford D, Gerwert K. Electrostatic calculations of the pka values of ionizable groups in bacteriorhodopsin. *J. Mol. Biol.* 1992; 224(2):473–486. [PubMed: 1313886]
6. Bashford, D. An object-oriented programming suite for electrostatic effects in biological molecules. an experience report on the MEAD project.. In: Ishikawa, Y.; Oldehoeft, R.; Reynders, J.; Tholburn, M., editors. *Scientific Computing in Object-Oriented Parallel Environments*, Vol. 1343 of *Lecture Notes in Computational Science*. Springer; Berlin Heidelberg: 1997. p. 233–240.
7. Geng WH, Yu SN, Wei GW. Treatment of charge singularities in implicit solvent models. *J. Chem. Phys.* 2007; 127:114106. doi:10.1063/1.2768064. [PubMed: 17887827]

8. Lu B, Cheng X, Huang J, McCammon JA. Order N algorithm for computation of electrostatic interactions in biomolecular systems. *P. Natl. Acad. Sci. USA*. 2006; 103(51):19314–19319. doi: 10.1073/pnas.0605166103.
9. Geng WH, Krasny R. A treecode-accelerated boundary integral poisson-boltzmann solver for solvated biomolecules. *J. Comp. Phys.* 2013; 247:62–78.
10. Yoon BJ, Lenhoff AM. A boundary element method for molecular electrostatics with electrolyte effects. *J. Comput. Chem.* 1990; 11(9):1080–1086.
11. Juffer AH, Botta EFF, van Keulen BAM, van der Ploeg A, Berendsen HJC. The electric potential of a macromolecule in a solvent: A fundamental approach. *J. Comp. Phys.* 1991; 97(1):144–171.
12. Zhou HX. Boundary-element solution of macromolecular electrostatics—interaction energy between 2 proteins. *Biophys. J.* 1993; 65:955–963. [PubMed: 8218918]
13. Altman MD, Bardhan JP, White JK, Tidor B. An accurate surface formulation for biomolecule electrostatics in non-ionic solutions. *Proceedings of the 27th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC 2005)*. 2005
14. Bajaj C, Chen SC, Rand A. An efficient higher-order fast multipole boundary element solution for poisson-boltzmann-based molecular electrostatics. *SIAM Journal of Scientific Computing*. 33(2)
15. Altman MD, Bardhan JP, White JK, Tidor B. Accurate solution of multi-region continuum electrostatic problems using the linearized Poisson–Boltzmann equation and curved boundary elements. *J. Comput. Chem.* 2009; 30:132–153. [PubMed: 18567005]
16. Cooper CD, Barba LA. Validation of the PyGBe code for Poisson-Boltzmann equation with boundary element methods. Technical Report on figshare, CC-BY license. doi:10.6084/m9.figshare.154331 (January 25 2013). doi:10.6084/m9.figshare.154331.
17. Klöckner A, Pinto N, Lee Y, Catanzaro B, Ivanov P, Fasih A. Py-CUDA and PyOpenCL: A scripting-based approach to GPU run-time code generation. *Parallel Computing*. 2012; 38(3):157–174.
18. Connolly ML. Analytical molecular surface calculation. *J. Appl. Cryst.* 1983; 16:548–558.
19. Courant, R.; Hilbert, D. *Methods of Mathematical Physics*. Vol. 1. Interscience Publishers, Inc.; New York: 1953. free ebook on the Internet Archive <http://archive.org/details/MethodsOfMathematicalPhysicsVolume1>
20. Duan Z-H, Krasny R. An adaptive treecode for computing nonbonded potential energy in classical molecular systems. *J. Comp. Chem.* 2001; 22(2):184–195.
21. Li P, Johnston H, Krasny R. A Cartesian treecode for screened Coulomb interactions. *J. Comp. Phys.* 2009; 228:3858–3868.
22. Zhu Z, Huang J, Song B, White J. Improving the robustness of a surface integral formulation for wideband impedance extraction of 3D structures. *Proceedings of the 2001 IEEE/ACM Int. Conf. on Computer-Aided Design*. 2001:592–597.
23. Baker NA, Sept D, Holst MJ, McCammon JA. The adaptive multilevel finite element solution of the poisson-boltzmann equation on massively parallel computers. *IBM Journal of Research and Development*. 2001; 45:427–438.
24. Sanner, MF. Molecular surface computation home page. 1996. <http://mgltools.scripps.edu/packages/MSMS> (last checked Feb. 4, 2010)
25. Dolinsky TJ, Nielsen JE, McCammon JA, Baker NA. PDB2PQR: an automated pipeline for the setup of Poisson–Boltzmann electrostatics calculations. *Nucleic Acids Research*. 2004; 32:W665–W667. [PubMed: 15215472]
26. García-García C, Draper DE. Electrostatic interactions in peptide-rna complex. *J. Mol. Biol.* 2003; 331(1):75–88. [PubMed: 12875837]
27. Cooper, CD.; Bardhan, JP.; Barba, LA. Convergence of pygbe with kirkwood sphere, Data, figures and plottings script on figshare. CC-BY license. <http://dx.doi.org/10.6084/m9.figshare.799692> (September 2013). doi:10.6084/m9.figshare.799692
28. Yokota R, Bardhan JP, Knepley MG, Barba LA, Hamada T. Biomolecular electrostatics using a fast multipole BEM on up to 512 GPUs and a billion unknowns. *Comput. Phys. Comm.* 2011; 182(6):1271–1283. doi:10.1016/j.cpc.2011.02.013.

29. Cooper, CD.; Bardhan, JP.; Barba, LA. Convergence and time to solution of pygbe with lysozyme molecule, Data, figures and plottings script on figshare. CC-BY license. <http://dx.doi.org/10.6084/m9.figshare.799702> (September 2013). doi:10.6084/m9.figshare.799702
30. Cooper CD, Bardhan JP, Barba LA. Binding energy of trypsinbpti complex with pygbe and apbs, Data, figures and plottings script on figshare. CC-BY license. <http://dx.doi.org/10.6084/m9.figshare.799703> (September 2013). doi:799703.
31. Cooper, CD.; Bardhan, JP.; Barba, LA. Binding energy of peptide-rna complex with pygbe and apbs, Data, figures and plottings script on figshare. CC-BY license. <http://dx.doi.org/10.6084/m9.figshare.799704> (September 2013). doi:799704

Figure 1.

Regions in the implicit-solvent model: Ω_1 is the protein region, with permittivity ϵ_1 and potential ϕ_1 ; Ω_2 is the solvent region, with permittivity ϵ_2 , inverse Debye length κ , and potential ϕ_2 ; Γ is the solvent-excluded surface and the dots represent the locations of atoms.

Figure 2.

Generation of the solvent-excluded surface. The dots are the locations of the atoms and the dotted circles represent the atomic spheres with appropriate radii. The solid ball represents the spherical probe used to compute the SES.

Figure 3. Sketch of the discretized interface Γ : Γ_i is the panel containing the collocation point and Γ_j the panels over which we integrate.

Figure 4.
Sketch of a molecule modeled with more than two regions: an ion-exclusion or Stern layer and one cavity are present.

(a) Solvation energy convergence with mesh refinement.

(b) Error convergence with mesh refinement.

Figure 5.

Convergence results for PyGBe and $_{APBS}$ using a spherical molecule of radius $4\text{\AA}$ with a centered charge of $1e^-$. The errors were calculated with respect to the corresponding extrapolated values (obtained from Richardson extrapolation). Data sets, figure files and plotting scripts available under CC-BY [27].

(a) Solvation energy convergence with mesh refinement.

(b) Error convergence with mesh refinement.

(c) Runtime vs. estimated error.

(d) Solvation energy comparison of various PyGBe models.

Figure 6.

Results for the lysozyme test, using PyGBe and _{APBS}. Data sets, figure files and plotting scripts available under CC-BY [29].

(a) Solvation energy convergence with mesh refinement.

(b) Error convergence with mesh refinement.

(c) Runtime vs. estimated error.

Figure 7.

APBS and PyGBe results for trypsin-BPTI complex. Data sets, figure files and plotting scripts available under CC-BY [30].

(a) Solvation energy convergence with mesh refinement.

(b) Error convergence with mesh refinement.

(c) Runtime vs. estimated error.

Figure 8.

APBS and PyGBe results for trypsin as receptor of trypsin-BPTI complex. Data sets, figure files and plotting scripts available under CC-BY [30].

(a) Solvation energy convergence with mesh refinement.

(b) Error convergence with mesh refinement.

(c) Runtime vs. estimated error.

Figure 9.

APBS and PyGBe results for BPTI as ligand of trypsin-BPTI complex. Data sets, figure files and plotting scripts available under CC-BY [30].

Figure 10.

APBS and PyGBe results for binding energies of the trypsin-BPTI complex. The multi-surface model of PyGBe considers cavities and Stern layer, and the single-surface model considers neither. Data sets, figure files and plotting scripts available under CC-BY [30].

Figure 11.

APBS and PyGBe results for binding energies of a 22-residue long α helix peptide with “box B” hairpin structure of RNA. PyGBe runs used a multi-surface model (considers Stern layer), and a single-surface model. Data sets, figure file and plotting script available under CC-BY [31].

Table 1

Parameters for setting up the runs with PyGBe and APBS.

ε_1	ε_2	κ	K	P	$\text{GMRES}_{\text{tol}}$
4	80	0.125	1	2	10^{-4}

Table 2

The three molecular systems used in the tests, and the corresponding entry in the Protein Data Bank.

Case	Name	PDB code
(i)	lysozyme	1HEL
(ii)	trypsin-BPTI	3BTK
(iii)	peptide-RNA	[26]

Table 3

Parameters for the convergence tests using a spherical molecule.

ε_1	ε_2	κ	K	P	$\text{GMRES}_{\text{tol}}$
4	80	0	3	12	10^{-6}

Table 4

Calculation parameters for the lysozyme case.

K_{fine}	N_{k}	thres	θ
19	5	1.25L	0.6

Table 5

Coulombic energies for the trypsin-BPTI complex.

Coulombic Energy [kJ/mol]		
Complex	Trypsin	BPTI
-79763.68	-62956.98	-17046.18

Table 6

Calculation parameters for trypsin-BPTI tests.

K_{fine}	N_{k}	thres	θ
19	5	$2L$	0.5

Table 7

Extrapolated solvation energies.

Code	Solvation energy [kJ/mol]		
	Complex	Trypsin	BPTI
APBS	-3384.08	-2037.88	-1218.86
PyGBe multi	-3421.0	-2039.92	-1216.02
PyGBe single	-3405.0	-2023.74	-1219.93