

Coupling finite and boundary element methods to solve the Poisson–Boltzmann equation for electrostatics in molecular solvation

Michał Bosy^{*†}, Matthew W. Scroggs^{‡§}, Timo Betcke^{†¶},
Erik Burman^{†||}, Christopher D. Cooper^{**†}

September 29, 2023

Abstract

The Poisson–Boltzmann equation is widely used to model electrostatics in molecular systems. Available software packages solve it using finite difference, finite element, and boundary element methods, where the latter is attractive due to the accurate representation of the molecular surface and partial charges, and exact enforcement of the boundary conditions at infinity. However, the boundary element method is limited to linear equations and piecewise constant variations of the material properties. In this work, we present a scheme that couples finite and boundary elements for the linearised Poisson–Boltzmann equation, where the finite element method is applied in a confined *solute* region and the boundary element method in the external *solvent* region. As a proof-of-concept exercise, we use the simplest methods available: Johnson–Nédélec coupling with mass matrix and diagonal preconditioning, implemented using the Bempp-cl and FEniCSx libraries via their Python interfaces. We showcase our implementation by computing the polar component of the solvation free energy of a set of molecules using a constant and a Gaussian-varying permittivity. As validation, we compare against well-established finite difference solvers for an extensive binding energy data set, and with the finite difference code APBS (to 0.5%) for Gaussian permittivities. We also show scaling results from protein G B1 (955 atoms) up to immunoglobulin G (20 148 atoms). For small problems, the coupled method was efficient, outperforming a purely boundary integral approach. For Gaussian-varying permittivities, which are beyond

^{*}School of Computer Science and Mathematics, Kingston University London, Penrhyn Road, Kingston upon Thames, KT1 2EE, UK

[†]ORCID: 0000-0003-2723-6913

[‡]Department of Mathematics, University College London, 25 Gordon Street, WC1H 0AY London, UK

[§]ORCID: 0000-0002-4658-2443

[¶]ORCID: 0000-0002-3323-2110

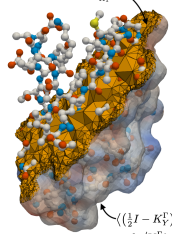
^{||}ORCID: 0000-0003-4287-7241

^{**}Department of Mechanical Engineering and Centro Científico Tecnológico de Valparaíso, Universidad Técnica Federico Santa María, Valparaíso 2390123, Chile

^{††}ORCID: 0000-0003-0282-8998

the applicability of boundary elements alone, we were able to run medium to large-sized problems on a single workstation. The development of better preconditioning techniques and the use of distributed memory parallelism for larger systems remains an area for future work. We hope this work will serve as inspiration for future developments that consider space-varying field parameters, and mixed linear-nonlinear schemes for molecular electrostatics with implicit solvent models.

Keywords: Finite element method, Boundary element method, Poisson–Boltzmann, Implicit solvent model, Electrostatics. ■

$$\begin{aligned}
& \langle \epsilon_1 \nabla \phi_1, \nabla v \rangle_{\Omega_1} - \langle \epsilon_1 \lambda_1, v \rangle_{\Gamma} \\
& = \left\langle \sum_{k=1}^{N_s} q_k \delta(\mathbf{x}_k, \mathbf{x}_k), v \right\rangle_{\Omega_1}
\end{aligned}$$


$$\begin{aligned}
& \left(\left(\frac{1}{2} I - K_{\Gamma}^{\epsilon} \right) \gamma \phi_1, \zeta \right)_{\Gamma} \\
& + \frac{1}{\epsilon_1} \left(V_{\Gamma}^{\epsilon} \lambda_2, \zeta \right)_{\Gamma} = 0
\end{aligned}$$

The boundary element method is a popular numerical algorithm to solve the Poisson–Boltzmann equation for molecular electrostatics. However, this technique is limited to linear equations and piecewise constant variations of the field parameters. Here, we overcome such limitations by coupling it with a finite element method, implemented using the Bempp-cl and FEniCSx libraries via their Python interfaces. This results in an efficient, accurate, flexible, and easy-to-use computational tool for model development.

INTRODUCTION

In biologically relevant settings, the structure and function of biomolecules are largely determined by the surrounding water, which usually contains salt. To describe these systems accurately, we need to account for the solvent correctly, which has given rise to a wide range of models.¹ Highly detailed models consider every water molecule and salt ion explicitly. For solutions with large numbers of molecules, however, these models can be very computationally expensive, so implicit-solvent models—approximated models that use continuum theory to represent the ionic solution—are often used instead.^{2,3} In the case of electrostatics, the implicit-solvent model is mathematically characterized by the Poisson–Boltzmann equation (PBE)^{4,5}, which is widely used to compute solvation free energies and mean-field potentials.

The implicit-solvent model for electrostatics describes the dissolved molecule as an infinite medium with a low-dielectric solute-shaped cavity, which contains a charge distribution from the partial charges—usually a sum of Dirac deltas at the atom’s locations. The outer solvent region is represented with a high dielectric constant and considers the presence of salt. These two regions are interfaced by the molecular surface where the continuity of the electrostatic potential and electric displacement are enforced. The molecular surface can be defined in various ways⁶.

The PBE has been solved numerically with finite difference^{7–10}, finite element^{11–13}, boundary element^{14–19}, and analytic^{20–24} methods. In particular, the boundary element method (BEM) has proven to be very efficient for high-accuracy calculations^{18,19}, mainly due to the precise description of the molecular surface and point charges. However, BEM is limited to constant material properties in each region and the linear version of the PBE. Even though these limitations are acceptable in a wide range of applications, there are cases when BEM falls short: for example, if a variable permittivity is required inside the solute^{25,26}, or the solute is highly charged such that the linear approximation breaks²⁷.

The present article describes a methodology to overcome some of those limitations, by coupling finite and boundary element methods. This approach brings the best of both worlds—the flexibility of FEM and the efficiency of BEM—all in an accurate description of the solute molecule. FEM-BEM coupling is a popular technique in the context of mixed

linear-nonlinear models,^{28,29} fracture mechanics,³⁰ fluid-structure interaction,³¹ acoustics,³² and electromagnetics.^{33–35} On the other hand, the PBE has been solved with hybrid numerical methods in the past: for example, by coupling finite differences with boundary elements³⁶ or finite elements^{37,38} to solve the nonlinear PBE in a specific region, and to implement modifications to the PBE model (i.e. the size-modified PBE). To the best of our knowledge, this is the first time finite and boundary elements have been combined in this application. We are proposing a coupling scheme for the linear Poisson-Boltzmann equation.

In this paper, we prototype this principle with the simplest implementation possible: a Johnson–Nédélec³⁹ coupling, where we solve using mass-matrix and diagonal preconditioning, and without distributed memory execution. This limits the size of problems we can access currently, however, it sets the basis for future developments that use more elaborate formulations and algorithms that are readily available in open-source software libraries. In this work, we use the boundary element library Bempp-cl⁴⁰ and the finite element library FEniCSx^{41,42}. These libraries are easy to use and their full functionalities may be accessed via their Python interfaces, making them ideal tools for easily implementing problems like this as well as for exploring computational efficiency and moving towards tackling large-scale problems, such as a full viral capsid.^{43,44} We hope this work will inspire research along these lines.

METHODOLOGY

The implicit solvent model

The implicit solvent model^{2,3} can be described mathematically as a coupled system of partial differential equations, where the Poisson equation governs in the solvent region (Ω_1 in Figure 1), and the Poisson–Boltzmann equation governs in the solute region (Ω_2 in Figure 1). These regions are interfaced by the molecular surface (Γ), where the potential (ϕ) and electric displacement ($\epsilon\partial\phi/\partial\mathbf{n}$) are continuous across the surface. The linearised Poisson-Boltzmann




Figure 1: Representation of a molecule in an implicit solvent.

equation that results from approximating $\sinh(x)$ with x follows

$$\nabla^2 \phi_1(\mathbf{x}) = \frac{1}{\epsilon_1} \sum_{k=1}^{N_q} q_k \delta(\mathbf{x}, \mathbf{x}_k) \quad \mathbf{x} \in \Omega_1, \quad (1a)$$

$$(\nabla^2 - \kappa^2) \phi_2(\mathbf{x}) = 0 \quad \mathbf{x} \in \Omega_2, \quad (1b)$$

$$\phi_1(\mathbf{x}) = \phi_2(\mathbf{x}), \quad \mathbf{x} \in \Gamma, \quad (1c)$$

$$\epsilon_1 \frac{\partial \phi_1}{\partial \mathbf{n}}(\mathbf{x}) = \epsilon_2 \frac{\partial \phi_2}{\partial \mathbf{n}}(\mathbf{x}) \quad \mathbf{x} \in \Gamma. \quad (1d)$$

Here, ϵ_1 and ϵ_2 are the dielectric constants in the solute and solvent, respectively; κ is the inverse of the Debye length, related to the salt concentration; and q_k are the values of the partial charges, located at \mathbf{x}_k .

The electrostatic potential in Ω_1 can be further decomposed into singular and regular components as $\phi_1 = \phi_c + \phi_r$, where ϕ_c is the solution to

$$\begin{aligned} \nabla^2 \phi_c(\mathbf{x}) &= \frac{1}{\epsilon} \sum_{k=1}^{N_q} q_k \delta(\mathbf{x}, \mathbf{x}_k) & \mathbf{x} \in \Omega_1 \cup \Omega_2, \\ \phi_c(\mathbf{x}) &= 0 & \text{as } |\mathbf{x}| \rightarrow \infty. \end{aligned} \quad (2)$$

Physically, ϕ_c can be interpreted as the Coulomb-type potential from the point charges, whereas ϕ_r , also known as reaction potential, is originated by the polarization of the solvent and reorganization of the free ions. Usually, ϵ_1 is considered a constant value, yielding an analytic expression for ϕ_c . However, this is not the general case.

Regularized versions of the equations in (1)^{45,46} are widely used to numerically solve the Poisson–Boltzmann equation with finite element or finite difference methods, and have

recently been extended to super-Gaussian permittivities in the solute.⁴⁷ However, here we use the standard formulation in (1), as it offers more flexibility when dealing with, for example, variable permittivities, beyond super-Gaussian descriptions.

A common quantity of interest in implicit solvent models is the polar component of the solvation free energy, which is the change in Gibbs free energy as the molecule charges move from vacuum into a solute-shaped cavity inside the solvent. Considering the charge distribution ρ consists of point charges, this can be calculated as

$$\Delta G_{\text{solv}} = \frac{1}{2} \int_{\Omega_1} \rho(\mathbf{x}) \phi_r(\mathbf{x}) = \frac{1}{2} \sum_{k=1}^{N_q} q_k \phi_r(\mathbf{x}_k). \quad (3)$$

BEM-BEM coupling

Several possible boundary integral formulations of (1) exist.⁴⁸ The simplest form was presented by Yoon and Lenhoff in 1991⁴⁹ and is known as the *direct formulation*. However, the resulting system is usually ill-conditioned. Here, we use the better-conditioned formulation presented by Lu and co-workers,^{15,50,51} which reads

$$\begin{aligned} \frac{\phi_2}{2} \left(1 + \frac{\epsilon_1}{\epsilon_2}\right) - \left(K_Y^\Gamma - \frac{\epsilon_1}{\epsilon_2} K_L^\Gamma\right) \phi_2 + (V_Y^\Gamma - V_L^\Gamma) \lambda_2 &= \sum_{k=0}^{N_q} \frac{q_k}{4\pi\epsilon_2 |\mathbf{x}_\Gamma - \mathbf{x}_k|}, \\ \frac{\epsilon_1}{\epsilon_2} (W_Y^\Gamma - W_L^\Gamma) \phi_2 + \frac{\lambda_2}{2} \left(1 + \frac{\epsilon_1}{\epsilon_2}\right) + \left(\frac{\epsilon_1}{\epsilon_2} K_Y^{\Gamma'} - K_L^{\Gamma'}\right) \lambda_2 &= \sum_{k=0}^{N_q} \frac{\partial}{\partial \mathbf{n}_x} \left(\frac{q_k}{4\pi\epsilon_2 |\mathbf{x}_\Gamma - \mathbf{x}_k|} \right), \end{aligned} \quad (4)$$

where $\phi_2 = \phi_2(\mathbf{x}_\Gamma)$ is the potential on Γ as we approach from Ω_2 (*i.e.* the exterior Dirichlet trace), and $\lambda_2 = \frac{\partial}{\partial \mathbf{n}} \phi_2$ is the normal derivative Γ as we approach from Ω_2 (*i.e.* the exterior Neumann trace). Note that we can obtain ϕ_1 and λ_1 from ϕ_2 and λ_2 by applying the interface conditions (1c) and (1d). The operators V , K , K' , and W that appear in Equation (4) are the single-layer, double-layer, adjoint double-layer, and hypersingular operators, respectively, for the Laplace (subscript L) and Yukawa (subscript Y , also known as modified Helmholtz)

kernels. These are defined by

$$\begin{aligned}
V_i^\Gamma \varphi(\mathbf{x}) &= \oint_\Gamma g_i(\mathbf{x}, \mathbf{x}') \varphi(\mathbf{x}') d\mathbf{x}', \\
K_i^\Gamma \varphi(\mathbf{x}) &= \oint_\Gamma \frac{\partial g_i}{\partial \mathbf{n}'}(\mathbf{x}, \mathbf{x}') \varphi(\mathbf{x}') d\mathbf{x}', \\
K_i'^\Gamma \varphi(\mathbf{x}) &= \oint_\Gamma \frac{g_i}{\partial \mathbf{n}}(\mathbf{x}, \mathbf{x}') \varphi(\mathbf{x}') d\mathbf{x}', \\
W_i^\Gamma \varphi(\mathbf{x}) &= - \oint_\Gamma \frac{\partial^2 g_i}{\partial \mathbf{n}' \partial \mathbf{n}}(\mathbf{x}, \mathbf{x}') \varphi(\mathbf{x}') d\mathbf{x}',
\end{aligned} \tag{5}$$

where $i \in \{L, Y\}$, $\varphi(\mathbf{x})$ can be any distribution over Γ , and

$$\begin{aligned}
g_L(\mathbf{x}, \mathbf{x}') &= \frac{1}{4\pi|\mathbf{x} - \mathbf{x}'|} \\
g_Y(\mathbf{x}, \mathbf{x}') &= \frac{e^{-\kappa|\mathbf{x} - \mathbf{x}'|}}{4\pi|\mathbf{x} - \mathbf{x}'|}
\end{aligned} \tag{6}$$

are the corresponding free-space Green's functions.

Having computed the electrostatic potential with Eq. (4), we obtain the reaction potential (ϕ_r) in Eq. (3) by subtracting the Coulombic component from ϕ_1 . This gives¹⁹

$$\phi_r(\mathbf{r}) = -K_L^\Gamma \phi_1(\mathbf{r}) + V_L^\Gamma \lambda_1(\mathbf{r}), \tag{7}$$

where $\mathbf{r} \in \Omega_1$. This is valid as long as ϵ_1 is a constant, so we use Eq. (7) for both BEM-BEM and FEM-BEM coupling simulations with constant ϵ_1 .

Novel numerical solution of the Poisson–Boltzmann equation

The boundary element method (BEM) is a standard tool for the numerical solution of the linear version of the Poisson–Boltzmann equation in molecular electrostatics.^{52,53} This was implemented in numerous codes, such as AFMPB,¹⁵ TABI,¹⁸ PyGBe,^{19,54} and more recently, with Bempp-cl⁴⁸. One strong advantage of BEM is that it reduces the dimension of the problem by using the boundary integral formulation. Unfortunately, this advantage comes at a cost: it requires a fundamental solution to be known in order for the method to be applied. There are many linear problems for which fundamental solutions are not known: this is the case for the Poisson-Boltzmann equation with a heterogeneous permittivity inside

the molecule, *i.e.*

$$\begin{aligned}
\nabla \cdot (\epsilon_1(\mathbf{x}) \nabla \phi_1(\mathbf{x})) &= \sum_{k=1}^{N_q} q_k \delta(\mathbf{x}, \mathbf{x}_k) & \mathbf{x} \in \Omega_1, \\
(\nabla^2 - \kappa^2) \phi_2(\mathbf{x}) &= 0 & \mathbf{x} \in \Omega_2, \\
\phi_1(\mathbf{x}) &= \phi_2(\mathbf{x}), & \mathbf{x} \in \Gamma, \\
\epsilon_1(\mathbf{x}) \frac{\partial \phi_1}{\partial \mathbf{n}}(\mathbf{x}) &= \epsilon_2 \frac{\partial \phi_2}{\partial \mathbf{n}}(\mathbf{x}) & \mathbf{x} \in \Gamma.
\end{aligned} \tag{8}$$

A finite element method (FEM) is more suitable for such a case.

The coupling of FEM and BEM is an approach that can be used to solve a wide range of multiphysics problems on unbounded domains^{55,56}, as it takes advantage of both methods. On the one hand, the BEM satisfies the infinite boundary conditions exactly when they decay to zero, and only approximates boundary conditions on surfaces; hence, it is commonly used for problems involving infinite or semi-infinite domains. On the other hand, the FEM is known for its robustness and universal applicability, even for problems of inhomogeneous or non-linear nature. Here, we use the Johnson–Nédélec formulation,³⁹ which is the simplest formulation of FEM-BEM coupling, and we detail it next.

We start with the variational formulation of the internal problem. Applying integration by parts to the first equation of (8) we have, for every $v \in H_0^1(\Omega_1)$,

$$\langle \epsilon_1 \nabla \phi_1, \nabla v \rangle_{\Omega_1} - \langle \epsilon_1 \partial_n \phi_1, v \rangle_{\Gamma} = \left\langle \sum_{k=1}^{N_q} q_k \delta(\mathbf{x}, \mathbf{x}_k), v \right\rangle_{\Omega_1}, \tag{9}$$

where v is a test function, and $\langle \varphi, v \rangle_{\Gamma} = \int_{\Gamma} \varphi(\mathbf{x}) v(\mathbf{x}) \, d\mathbf{x}$ and $\langle \varphi, v \rangle_{\Omega_1} = \int_{\Omega_1} \varphi(\mathbf{x}) v(\mathbf{x}) \, d\mathbf{x}$ are the inner products on the surface and in the domain, respectively. Slightly abusing notation, the product $\langle \cdot, \cdot \rangle_{\Omega_1}$ on the right-hand side denotes the duality pairing. For the external problem with BEM, we define the Dirichlet trace⁵⁷

$$\gamma : H^1(\Omega_2) \rightarrow H^{\frac{1}{2}}(\Gamma), \quad \gamma f(x) := \lim_{\Omega_2 \ni \mathbf{y} \rightarrow \mathbf{x} \in \Gamma} f(\mathbf{y}),$$

and we use the direct formulation from the second equation of Eq. (1) to obtain

$$\frac{\phi_1}{2} - K_Y^\Gamma \gamma \phi_1 + \frac{\epsilon_1}{\epsilon_2} V_Y^\Gamma \lambda_1 = 0.$$

V and K are the single-layer and double-layer operators as defined in Equation (5).

Then, the coupling problem can be written as: *Find* $\phi_1 \in H^1(\Omega_1)$ and $\lambda_1 \in H^{-\frac{1}{2}}(\Gamma)$ such that for all $v \in H^1(\Omega_1)$ and $\zeta \in H^{-\frac{1}{2}}(\Gamma)$,

$$\langle \epsilon_1 \nabla \phi_1, \nabla v \rangle_{\Omega_1} - \langle \epsilon_1 \lambda_1, v \rangle_{\Gamma} = \left\langle \sum_{k=1}^{N_q} q_k \delta(\mathbf{x}, \mathbf{x}_k), v \right\rangle_{\Omega_1}, \quad (10a)$$

$$\left\langle \left(\frac{1}{2}I - K_Y^\Gamma \right) \gamma \phi_1, \zeta \right\rangle_{\Gamma} + \frac{1}{\epsilon_2} \langle V_Y^\Gamma \epsilon_1 \lambda_1, \zeta \rangle_{\Gamma} = 0. \quad (10b)$$

Note that in the case where ϵ_1 is constant, then each ϵ_1 can be moved outside of the inner product it is inside.

When discretised, this can also be written in matrix form. Let $\vec{\phi}_1 := [\phi_1^1, \dots, \phi_1^j]^T$ be the vector of canonical basis functions of the finite element space $V_h^j \subset H^1(\Omega_1)$, and let $\vec{\lambda}_1 := [\lambda_1^1, \dots, \lambda_1^l]^T$ be the vector of canonical basis functions of $\Lambda_h^l \subset H^{-\frac{1}{2}}(\Gamma)$. We define the following matrices associated with the corresponding bilinear forms

$$\begin{aligned} A_{\alpha\beta} &= \left\langle \epsilon_1 \nabla \phi_1^\alpha, \nabla \phi_1^\beta \right\rangle_{\Omega_1}, & \widetilde{M}_{\alpha\beta} &= \left\langle \epsilon_1 \lambda_1^\beta, \gamma \phi_1^\alpha \right\rangle_{\Gamma}, \\ K_{\alpha\beta} &= \left\langle K_Y^\Gamma \gamma \phi_1^\alpha, \lambda_1^\beta \right\rangle_{\Gamma}, & V_{\alpha\beta} &= \left\langle V_Y^\Gamma \epsilon_1 \lambda_1^\alpha, \lambda_1^\beta \right\rangle_{\Gamma}, \\ M_{\alpha\beta} &= \left\langle \gamma \phi_1^\alpha, \lambda_1^\beta \right\rangle_{\Gamma}, \end{aligned}$$

and vector associated with the corresponding linear form

$$\vec{f}_\beta := \left\langle \sum_{k=1}^{N_q} q_k \delta(\mathbf{x}, \mathbf{x}_k), \phi_1^\beta \right\rangle_{\Omega_1}.$$

Using the above definitions, the discrete problem in (10) can be written in the following blocked matrix form:

$$\begin{bmatrix} A & -\widetilde{M}^T \\ \frac{1}{2}M - K & \frac{1}{\epsilon_2}V \end{bmatrix} \begin{bmatrix} \vec{\phi}_1 \\ \vec{\lambda}_1 \end{bmatrix} = \begin{bmatrix} \vec{f} \\ 0 \end{bmatrix}. \quad (11)$$

RESULTS AND DISCUSSION

This section presents the verification and performance results of the presented FEM-BEM coupling scheme for molecules modelled as cavities with constant and varying permittivity. With a constant permittivity inside the molecule, we tested convergence against an analytic

expression of the solvation energy of a sphere⁵⁸, and then compared a more realistic geometry (arginine) with a purely BEM implementation. Then, as a more challenging test, we performed binding energy calculations from a set of 153 structures,⁶ and compared them to well-established codes. We also considered a Gaussian-varying permittivity^{25,26} inside the molecular cavity of arginine, and used APBS⁷ to verify our results. The final tests show the scaling of the FEM-BEM coupling as the molecule size grows.

All runs were done on a Lenovo ThinkStation P620 with AMD Ryzen ThreadRipper PRO 3975WX (32-core and 3.5 GHz) and 128 GB RAM.

Software environment

For the finite element computations, we use the software package FEniCSx^{41,42} while for the boundary element computations, we use Bempp-cl⁴⁰ together with Exafmm-t⁵⁹. FEniCSx is the successor of the widely used FEniCS finite element library^{60,61}. It provides a convenient Python interface, describing problems using Unified Form Language (UFL)⁶², a convenient domain-specific language specifically designed for finite element discretisations of partial differential equations. During assembly, the UFL description is transformed into efficient low-level C++ code and just-in-time compiled^{63,64}. Bempp-cl is a Python package that uses low-level OpenCL kernels written in C99 to provide optimised assembly routines⁶⁵. The built-in dense assembly routines are highly efficient for moderate discretisation sizes up to a few ten thousand elements.

For very large grid sizes the user can enable fast multipole method (FMM) assembly which internally is handled in Bempp-cl through an interface to the Exafmm-t FMM library. For N surface elements this reduces the memory and computational complexity from $\mathcal{O}(N^2)$ in the dense assembly case to $\mathcal{O}(N)$ in the FMM case, making large boundary element problems tractable on standard workstations. In this work, we only use the FMM in the performance analysis for larger structures, as all other cases are small enough to run efficiently with a dense assembler.

To couple FEniCSx with Bempp-cl we load a volume mesh with FEniCSx. We then export the corresponding boundary mesh into Bempp-cl and assemble the boundary spaces

there. Bempp-cl provides numerical trace operators that can translate from the degree of freedom (DOF) representation in FEniCSx to the DOF representation in Bempp-cl. The corresponding translation work is handled opaquely and the user only needs to deal with high-level interfaces of FEniCSx operators, Bempp-cl operators, and trace operators. FMM assembly fits automatically into this framework and can be enabled or disabled as a simple configuration option. Once the discrete block matrices in Eq. (11) are built with Bempp-cl and FEniCSx, we assemble them into one matrix and solve the linear system with Scipy’s⁶⁶ GMRES, in this case, with a tolerance of 10^{-5} .

Docker images containing FEniCSx, Bempp-cl, and Exafmm-t are publicly available (<https://bempp.com/installation.html>), and all codes used to generate the results in this section are available as Jupyter Notebooks that can be reproducibly executed in an appropriate environment. The results in this section were obtained using version 0.6.0 of FEniCSx and version 0.3.0 of Bempp-cl. All codes to reproduce the results of this manuscript can be found in the GitHub repository https://github.com/MichalBosy/FEM_BEM_coupling/.

Results with constant permittivity

In implicit-solvent models, the molecule is usually considered as a region with constant permittivity: in our computations, we use $\epsilon_1 = 2$. In the solvent region, we used the permittivity of water ($\epsilon_2 = 80$) and an inverse of the Debye length of $\kappa = 0.125 \text{ \AA}^{-1}$. In this case, there is a known analytic solution for ϕ_c in Eq. (2), so it is enough to compute ϕ_r . We do this using Eq. (7) with both BEM-BEM and FEM-BEM coupling approaches. For FEM-BEM, the integral over Γ in Eq. (7) corresponds to the trace of the solution vector from Eq. (11).

Convergence of a spherical cavity

The Kirkwood sphere⁵⁸ is a standard benchmark test for the Poisson–Boltzmann equation in molecular electrostatics. In this case, we considered a spherical cavity of radius $R = 2 \text{ \AA}$, with three charges ($q_1 = 1$, $q_2 = 1$, and $q_3 = 0.75$) placed at $\mathbf{x}_1 = (1, 0, 0)$, $\mathbf{x}_2 = (0.7, 0.7, 0)$, and $\mathbf{x}_3 = (-0.5, -0.5, 0)$. Figure 2 shows the percentage error of the FEM-BEM approach,

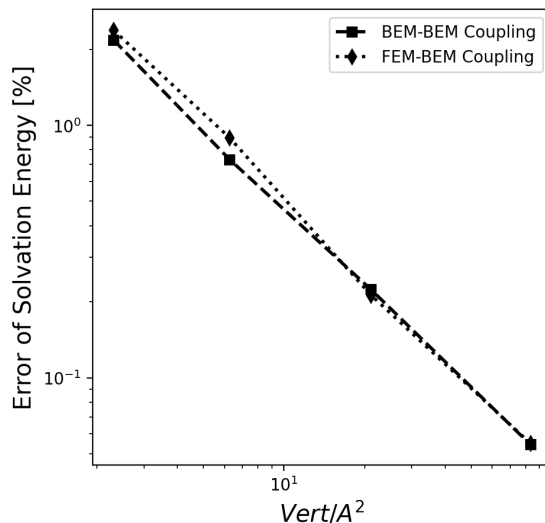


Figure 2: Error for the Kirkwood sphere.

and a reference BEM-BEM implementation, compared to the analytic solution ($\Delta G_{\text{solv}} = -336.0396 \text{ kcal/mol}$). In these runs, the FEM mesh was generated using GMSH⁶⁷ with 2, 6, 21 and 83 vertices per \AA^2 on the SES. The triangular mesh on the boundary surface of the FEM mesh was used for the BEM runs. The error in Figure 2 decays linearly with the number of vertices on the surface. This agrees with the expected convergence for P1 elements, indicating a correct implementation of the numerical scheme.

Performance with arginine

As a more realistic test, we assessed the performance of the FEM-BEM coupling technique against a BEM-BEM implementation for arginine. The structure of arginine was taken from the protein data bank, and parameterized with the Amber⁶⁸ force field. We generated surface meshes containing 4.1, 6.7, 8.6, 17, and 24.5 vertices per \AA^2 with Nanoshaper.⁶⁹ These densities correspond to a grid-scale parameter in Nanoshaper equal to 1.6, 2.0, 2.4, 3.4, and 4.0, respectively, where the grid scale is the reciprocal of the average characteristic length of the triangles. For our BEM-BEM solver, we used these meshes directly. For the FEM-BEM solver, we created volume meshes from these using pyGAMer,⁷⁰ which invoked TetGen⁷¹ with a quality parameter (radius-edge ratio) of 1.0.

The solvation energy computed with the two schemes is presented in Figure 3, which,

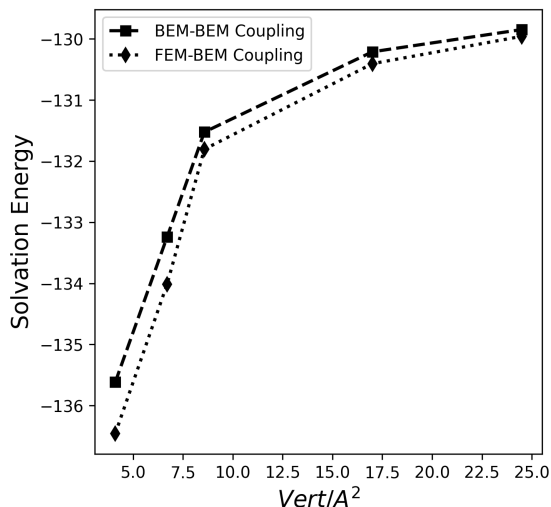


Figure 3: Solvation energy for arginine with a constant permittivity.

as expected, converges to a similar answer as the mesh is refined. Figure 4 compares the iteration count and time-to-solution. The left plot shows that using only BEM outperforms the coupled FEM-BEM approach in terms of the iteration count. However, if we look at the total time that solvers take to obtain the solution, we can see the advantage of using the FEM-BEM coupling. The higher computational cost is caused by the need to use a hypersingular operator in the BEM formulation, and the fact that we are not using any acceleration method (i.e. FMM). The timings for the FEM-BEM coupling scale are at a greater rate than the pure BEM counterpart, indicating that work on preconditioners and other acceleration methods will be required in order to make the FEM-BEM approach viable for large problems.

Results with variable permittivity

Note on modelling differences between finite difference and finite/boundary element methods

To facilitate verification, we conducted a comparative analysis between our FEM-BEM approach and established Poisson-Boltzmann solvers, including APBS,⁷ MIBPB,^{72,73} CPB,⁷⁴ PBSA,^{75,76} and Delphi.⁷⁷ These benchmarks primarily employ finite difference approxima-




Figure 4: Iteration count (left) and time-to-solution (right) for arginine with a constant permittivity.

tions, setting the stage for inherent modelling disparities with respect to our FEM-BEM method, and making this comparison a difficult task. Firstly, the boundary integral formulation in the external domain naturally enforces the decay of the electrostatic potential to zero at infinity, while finite difference-based solvers necessitate the imposition of effective boundary conditions along the mesh periphery. This boundary condition critically influences the electrostatic potential and may deviate from the true solution if the size of the mesh box is insufficient. Secondly, the definition of the molecular surface differs slightly across software, resulting in different domain geometries. Our FEM-BEM approach employs a triangulation of the solvent-excluded surface (based on NanoShaper⁶⁹), while finite difference meshes produce a staircase representation of this surface. This discrepancy is addressed by the matched interface technique in MIBPB,^{72,73} however, the resulting surface does not exactly match NanoShaper’s triangulation.

Binding energy calculations for a large data set

The binding energy of two molecules in a complex is

$$\Delta G_{bind} = \Delta G_{solv,complex} - \Delta G_{solv,1} - \Delta G_{solv,2} + \Delta G_{coul} \quad (12)$$

where $\Delta G_{solv,complex}$ is the solvation energy of the complex, $\Delta G_{solv,1}$ and $\Delta G_{solv,2}$ are the solvation energies of the two compounds, and ΔG_{coul} is the difference in Coulomb energy between bound and unbound states. This calculation presents a major challenge, as small errors in solvation energy may have a large impact result in in binding energy, making accuracy crucial.

Following the data set proposed by Fenley and co-workers,⁶ available in http://www.sb.fsu.edu/~mfenley/convergence/downloads/convergence_pqr_sets.tar.gz, and later used by Wei *et al.*⁷⁸ to assess the accuracy of different Poisson-Boltzmann solvers, we tested our method by computing binding energies. This data set consists of 51 complexes, totalling 153 structures, separated in three families: DNA-drug complexes, different versions of the barnase-barstar complex, and RNA-peptide complexes. Following the problem setting for these data sets, we use $\epsilon_1 = 1$, $\epsilon_2 = 80$ and an inverse of the Debye length of $\kappa = 0.10265 \text{ \AA}^{-1}$.

According to the arginine results in Figure 3, a mesh density of 24 vertices per \AA^2 gives a solvation energy that is close to the true solution for a realistic geometry. In that case, the difference in solvation energy between calculations using finite element grids generated from a surface mesh of 8.6 and 24.5 vertices per \AA^2 (a $3\times$ refinement) is less than 1%. Then, we used surface meshes with 24 vertices per \AA^2 and a quality parameter of 1.1 in TetGen to generate the meshes used to compute binding energies in Fenley’s database.

Several structures in the database contained voids inside the molecular region where a water molecule fits. In a finite difference description, these voids are naturally considered as pockets of solvent, with the permittivity of water ($\epsilon = \epsilon_2$). Regardless if this is an appropriate physical model, we also placed the permittivity of bulk water inside cavities for comparison purposes. In the FEM formalism, we had to detect the elements contained by cavities and use a piece-wise constant varying description of the permittivity inside the molecular region, which can be done without any modification in the formulation. A BEM-BEM formulation can also account for cavities as regions with water’s bulk permittivity,^{16,19} but involves certain adjustments in the stiffness matrix.

Figures 5, 6, and 7 show the binding energies of structures of DNA-drug complexes, barnase-barstar complexes, and RNA-peptide complexes, respectively, presented in Fenley’s data set.⁶ The data for CPB⁶ and MIBPB, Delphi, APBS and PBSA were taken from the




Figure 5: Binding energies for structures corresponding to DNA-drug complexes.

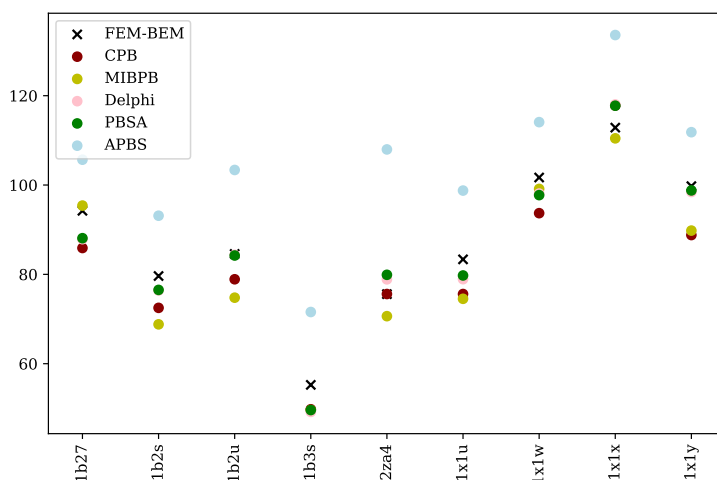


Figure 6: Binding energies for structures corresponding to the barnase-barstar complex.

literature,⁷⁸ for the finest mesh reported. From Figures 5, 6, and 7 it is evident that our FEM-BEM approach gives results that are within the range of the difference between other software for all 51 complexes. Given the challenging nature of binding energy calculations, and that our results are consistently comparable to other widely-trusted solvers that use a different numerical scheme for all 51 complexes, we can positively conclude on the correctness of our code.

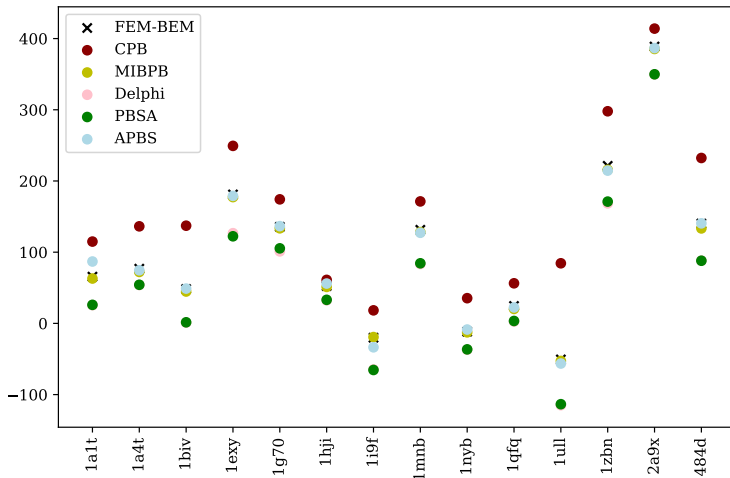


Figure 7: Binding energies for structures corresponding to RNA-peptide complexes.

Modeling the solute with a Gaussian-based variable permittivity

In contrast to a purely BEM approach, FEM-BEM coupling gives the flexibility to consider space-varying field parameters. A popular description of the molecule is to consider a permittivity that varies like a Gaussian around each atom,²⁵ which has shown enhanced accuracy in some applications, like pKa calculations.²⁶ In this setting, we define a density function ρ depending on the position \mathbf{r} as

$$\rho(r) := \prod_i \left(1 - \exp\left(-\frac{\|\mathbf{r} - \mathbf{x}_i\|}{\sigma^2 R_i^2}\right) \right), \quad (13)$$

where the product runs over all the atoms of the solute, R_i is the van der Waals radius of atom i , and we used $\sigma = 1$. Then, we can compute the permittivity as

$$\epsilon := (1 - \rho) \epsilon_1 + \rho \epsilon_2 \quad (14)$$

As ϵ is variable, Equation (2) does not have a known analytic solution, and the electrostatic potential in the vacuum state has to be computed numerically. For vacuum calculations, we considered the same distribution of ϵ inside the molecule as in the solvated case, but the solvent permittivity was set to $\epsilon_2 = 2$. Other implementations of Gaussian permittivities also modify the solute permittivity in vacuum calculations, according to a set cutoff.²⁶ We did not consider a cutoff in our calculations.

Mesh size (Å)	Grid points	ΔG_{solv} (kcal/mol)	Mesh density (vert/Å ²)	DOFs	ΔG_{solv} (kcal/mol)
$0.52 \times 0.52 \times 0.52$	$97 \times 97 \times 97$	-107.6186	4.1	3 491	-109.931
$0.39 \times 0.39 \times 0.39$	$129 \times 129 \times 129$	-107.8752	6.7	5 787	-110.237
$0.26 \times 0.26 \times 0.26$	$193 \times 193 \times 193$	-108.3378	8.6	8 844	-109.661
$0.195 \times 0.195 \times 0.195$	$257 \times 257 \times 257$	-108.5837	17.0	19 911	-109.369
$0.098 \times 0.098 \times 0.098$	$513 \times 513 \times 513$	-108.8844	24.5	32 302	-109.315

Table 1: Solvation energy of arginine with a Gaussian-like permittivity, computed using APBS (left) and the FEM-BEM approach (right). The mesh density for FEM-BEM corresponds to the vertex density of the surface mesh used to generate the volumetric mesh.

Convergence for arginine with APBS

We used Equation (14) to generate dielectric maps, which we ran on APBS⁷ for comparison. We chose APBS because it provides an easy interface to control dielectric maps to ensure their agreement with the maps imposed in our FEM-BEM coupled approach.

Table 1 shows a comparison of the solvation energy computed with APBS and our FEM-BEM coupling approach. We can see that they are both converging to equivalent values, where the solutions on the finest meshes agree to within 0.5% (0.5 kcal/mol). The coarsest meshes used in both cases are within the recommended densities for accurate solvation energy simulations with constant permittivity: a finite difference meshes with $h = 0.5$ Å or less is recommended for binding energy calculations⁷⁹ (the result of subtracting two solvation energies), and a similar study with BEM¹⁹ showed that a mesh with 2 vertices per Å² gives acceptable results when computing binding and solvation energies. We see a jump in the solvation energy for the FEM-BEM results going from 4.1 to 6.7 vertices per Å², but later the energy monotonically decreases, converging to a solution. This is an indication that mesh requirements for variable permittivities may be tighter than with a constant permittivity, even though the jump in dielectric constant across the molecular surface is usually smaller.

Performance analysis for larger structures

In the experiments presented so far, we have only tested the FEM-BEM coupled approach with small structures. In this section, we study its behaviour with larger structures to eval-

uate its applicability in more realistic settings, and in test cases where BEM-BEM coupling is not an alternative. Using the same Gaussian permittivity field inside the molecule, we computed the solvation free energy of protein G B1 (955 atoms, PDB code 1pgb), lysozyme (1960 atoms, PDB code 1lyz), the barnase-barstar complex (9464 atoms, PDB code 1x1u), and immunoglobulin G (20148 atoms, PDB code 1igt). All structures were parameterized with the Amber⁸⁰ force field using `pdb2pqr`.⁸¹ The surfaces were meshed with `Nanoshaper`⁶⁹ using a grid scale of 1.5, and then volume meshes were generated using `pyGAMer`⁷⁰ and `TetGen`⁷¹ with a radius-edge ratio of 1.0. The solvation free energy and timings for these runs are presented in Table 2. These results demonstrate that our FEM-BEM coupling implementation can reach medium-to-large-sized proteins on a simple workstation. For these larger structures, we enabled the FMM capabilities of `Bempp-cl`.

Table 2 shows that the iteration count increases with the problem size. This is expected for Johnson–Nédélec coupling, which is the simplest coupling strategy. To isolate the effect of the increased iterations in the analysis, we separate timings into the setup and solving time, where the setup time is independent of the number of iterations, and the solving time corresponds to the time spent in the GMRES solver. The time-per-iteration is computed by dividing the solving time by the number of iterations. The time per iteration scales closer to $\mathcal{O}(N)$, which is expected as we are using FMM for the BEM portion of the matrix. This is an indication that the high solving time is mainly due to the increase in iteration count, and having better-conditioned coupling methods, such as the so-called hybrid approach⁸², or more effective preconditioners for the blocked system would have a large impact on the time to solution.

Even though this scheme is capable of calculating the electrostatic potential in medium-to-large proteins, the largest test case in Table 2 (1igt) used up 30% of the available RAM. If we were aiming at larger structures, such as full viruses,^{43,44} we would require the use of optimised fast algorithms^{83,84} and parallelisation of the storage and solver, alongside the necessary improvements in the conditioning of the system.

Molecule	FEM DOFs	BEM DOFs	ΔG_{solv} (kcal/mol)	Iterations	Setup time (s)	Solving time (s)	Time per iter. (s)	Total time (s)
1pgb	29 434	10 058	-300.888	703	86	472	0.67	789
1lyz	56 114	18 606	-599.310	1 073	336	1 350	1.26	1 686
1x1u	263 181	81 258	-1 982.085	2 791	8 340	11 000	3.94	19 340
1igt	597 575	187 712	-3 294.0157	5 366	41 998	40 100	7.47	82 098

Table 2: Results for larger structures with a variable permittivity.

CONCLUSIONS

This paper presents the first implementation of a FEM-BEM coupling approach to solve the Poisson–Boltzmann equation for molecular electrostatics. This brings the best of both worlds: the accuracy and efficiency of BEM to exactly enforce the boundary conditions at infinity, and the flexibility of FEM to account for space variations of the material properties and nonlinearities. After presenting verification results of solvation energy for a sphere and arginine with a constant permittivity inside the solute, and binding energy for an extensive data set, we showcased our implementation with an advanced modelling technique that considers Gaussian-varying permittivities in a confined region, with the results validated against the widely-used APBS software. The final scaling results for larger molecules start from protein G B1 (955 atoms) and go up to immunoglobulin G (20 148 atoms), proving the applicability of this approach for realistic problems. Even though our implementation was able to reach medium-to-large systems, we recognise the need for further research towards better preconditioning of the linear system, optimising the coupling technique, acceleration algorithms, and parallel execution, especially as we look towards much larger solutes, like viruses. We hope this proof-of-concept work will serve as motivation for future model development that considers space-varying permittivities and Debye lengths, and mixed linear-nonlinear techniques, especially for highly-charged systems, like nucleic acids.

ACKNOWLEDGMENTS

MB acknowledges the support from Kingston University through the First Kingston University Grant.

MS acknowledges support from EPSRC grant EP/W007460/1

TB acknowledges support from EPSRC grants EP/W007460/1 and EP/W026260/1.

EB acknowledges support from EPSRC grants EP/V050400/1 and EP/W007460/1.

CDC acknowledges the support from CCTVal through ANID PIA/APOYO AFB220004 and Universidad Técnica Federico Santa María through Proyectos Internos USM 2023 PI-LIR-23-03.

We would like to thank Jørgen Dokken for his help with various implementation details related to FEniCSx.

DATA AVAILABILITY STATEMENT

All codes to reproduce the results of this manuscript can be found in the GitHub repository https://github.com/MichalBosy/FEM_BEM_coupling/, alongside links to Docker images that include the codes alongside installations of appropriate versions of Bempp-cl and FEniCSx. The numerical results of the binding energy calculations, alongside all the required input files to reproduce them, are available in the Zenodo repository [10.5281/zenodo.8393170](https://zenodo.org/record/105281/versions/8393170).

References

1. A. V. Onufriev and S. Izadi, Wiley Interdisciplinary Reviews: Computational Molecular Science **8**, e1347 (2018).
2. B. Roux and T. Simonson, Biophys. Chem. **78**, 1 (1999).
3. S. Decherchi, M. Masetti, I. Vyalov, and W. Rocchia, European journal of medicinal chemistry **91**, 27 (2015).
4. N. A. Baker, Meth. Enzymol. **383**, 94 (2004).
5. J. P. Bardhan, Comput. Sci. Disc. **5** (2012).
6. R. C. Harris, A. H. Boschitsch, and M. O. Fenley, Journal of chemical theory and computation **9**, 3677 (2013).
7. N. A. Baker, D. Sept, M. J. Holst, and J. A. McCammon, P. Natl. Acad. Sci. USA **98**, 10037 (2001).
8. M. K. Gilson, A. Rashin, R. Fine, and B. Honig, J. Mol. Biol. **184**, 503 (1985).
9. E. Jurrus, D. Engel, K. Star, K. Monson, J. Brandi, L. E. Felberg, D. H. Brookes, L. Wilson, J. Chen, K. Liles, et al., Protein Science **27**, 112 (2018).
10. C. Li, Z. Jia, A. Chakravorty, S. Pahari, Y. Peng, S. Basu, M. Koirala, S. K. Panday, M. Petukh, L. Li, et al., Journal of computational chemistry **40**, 2502 (2019).
11. M. Holst, J. A. Mccammon, Z. Yu, Y. Zhou, and Y. Zhu, Communications in computational physics **11**, 179 (2012).
12. S. D. Bond, J. H. Chaudhry, E. C. Cyr, and L. N. Olson, Journal of Computational Chemistry **31**, 1625 (2010).
13. S. Nakov, E. Sobakinskaya, T. Renger, and J. Kraus, Journal of Computational Chemistry **42**, 1832 (2021).

14. A. H. Boschitsch, M. O. Fenley, and H.-X. Zhou, *The Journal of Physical Chemistry B* **106**, 2741 (2002).
15. B. Lu, X. Cheng, J. Huang, and J. A. McCammon, *P. Natl. Acad. Sci. USA* **103**, 19314 (2006).
16. M. D. Altman, J. P. Bardhan, J. K. White, and B. Tidor, *J. Comput. Chem.* **30**, 132 (2009).
17. C. Bajaj, S.-C. Chen, and A. Rand, *SIAM Journal on Scientific Computing* **33**, 826 (2011).
18. W. H. Geng and R. Krasny, *J. Comput. Phys.* **247**, 62 (2013).
19. C. D. Cooper, J. P. Bardhan, and L. A. Barba, *Comput. Phys. Commun.* **185**, 720 (2014).
20. E.-H. Yap and T. Head-Gordon, *Journal of chemical theory and computation* **6**, 2214 (2010).
21. L. E. Felberg, D. H. Brookes, E.-H. Yap, E. Jurrus, N. A. Baker, and T. Head-Gordon, *Journal of computational chemistry* **38**, 1275 (2017).
22. C. Quan, B. Stamm, and Y. Maday, *SIAM Journal on Scientific Computing* **41**, B320 (2019).
23. S. V. Siryk and W. Rocchia, *The Journal of Physical Chemistry B* **126**, 10400 (2022).
24. A. Jha and B. Stamm, arXiv preprint arXiv:2309.06862 (2023).
25. J. A. Grant, B. T. Pickup, and A. Nicholls, *Journal of computational chemistry* **22**, 608 (2001).
26. L. Li, C. Li, Z. Zhang, and E. Alexov, *Journal of chemical theory and computation* **9**, 2126 (2013).
27. F. Fogolari, P. Zuccato, G. Esposito, and P. Viglino, *Biophysical journal* **76**, 1 (1999).

28. C. Carstensen and E. Stephan, Numerical Methods for Partial Differential Equations **11**, 539 (1995).
29. M. Aurada, M. Feischl, T. Führer, M. Karkulik, J. M. Melenk, and D. Praetorius, Computational Mechanics **51**, 399 (2013).
30. B. Aour, O. Rahmani, and M. Nait-Abdelaziz, International Journal of Solids and Structures **44**, 2523 (2007).
31. O. v. Estorff and H. Antes, International Journal for Numerical Methods in Engineering **31**, 1151 (1991).
32. R. Hiptmair and P. Meury, SIAM journal on numerical analysis **44**, 2107 (2006).
33. F. Matsuoka and A. Kameari, IEEE Transactions on Magnetics **24**, 182 (1988).
34. R. Hiptmair and P. Meury, in *Modeling and computations in electromagnetics* (Springer, 2008), pp. 1–38.
35. F. Bruckner, C. Vogler, M. Feischl, D. Praetorius, B. Bergmair, T. Huber, M. Fuger, and D. Suess, Journal of magnetism and magnetic materials **324**, 1862 (2012).
36. A. H. Boschitsch and M. O. Fenley, Journal of computational chemistry **25**, 935 (2004).
37. D. Xie and J. Ying, Journal of Computational and Applied Mathematics **307**, 319 (2016).
38. J. Ying and D. Xie, Applied Mathematical Modelling **58**, 166 (2018).
39. C. Johnson, J. Nédélec, et al., Mathematics of computation **35**, 1063 (1980).
40. T. Betcke and M. W. Scroggs, Journal of Open Source Software **6**, 2879 (2021).
41. M. W. Scroggs, I. A. Baratta, C. N. Richardson, and G. N. Wells, Journal of Open Source Software **7**, 3982 (2022).
42. M. W. Scroggs, J. S. Dokken, C. N. Richardson, and G. N. Wells, ACM Transactions on Mathematical Software **48**, 18:1 (2022).

43. M. Martínez, C. D. Cooper, A. B. Poma, and H. V. Guzman, *Journal of Chemical Information and Modeling* **60**, 974 (2019).
44. T. Wang, C. D. Cooper, T. Betcke, and L. A. Barba, arXiv preprint arXiv:2103.01048 (2021).
45. B. Z. Lu, Y. C. Zhou, M. J. Holst, and J. A. McCammon, *Communications in Computational Physics* **3**, 973 (2008).
46. A. Lee, W. Geng, and S. Zhao, *Journal of Computational Physics* **426**, 109958 (2021).
47. S. Wang, Y. Shao, E. Alexov, and S. Zhao, *Journal of Computational Physics* **464**, 111340 (2022).
48. C. D. Cooper, E. van 't Wout, et al., *Journal of Computational Chemistry* (2022).
49. B. J. Yoon and A. M. Lenhoff, *J. Comput. Chem.* **11**, 1080 (1990).
50. B. Lu, X. Cheng, J. Huang, and J. A. McCammon, *J. Chem. Theory Comput.* **5**, 1692 (2009).
51. J. DeBuhr, B. Zhang, A. Tsueda, V. Tilstra-Smith, and T. Sterling, *Communications in Computational Physics* **20**, 1106 (2016).
52. R. Zauhar and R. Morgan, *Journal of molecular biology* **186**, 815 (1985).
53. P. B. Shaw, *Phys. Rev. A* **32**, 2476 (1985).
54. C. D. Cooper, N. C. Clementi, G. Forsyth, and L. A. Barba, *J. Open Source Softw.* **1**, 43 (2016).
55. R. Hiptmair, *SIAM Journal on Numerical Analysis* **40**, 41 (2002).
56. L. Banjai, C. Lubich, and F.-J. Sayas, *Numerische Mathematik* **129**, 611 (2015).
57. O. Steinbach, *Numerical approximation methods for elliptic boundary value problems* (Springer, New York, 2008), ISBN 978-0-387-31312-2, finite and boundary elements, Translated from the 2003 German original, URL <https://doi.org/10.1007/978-0-387-68805-3>.

58. J. G. Kirkwood, *J. Chem. Phys.* **2**, 351 (1934).
59. T. Wang, R. Yokota, and L. A. Barba, *Journal of Open Source Software* **6**, 3145 (2021), URL <https://doi.org/10.21105/joss.03145>.
60. M. S. Alnaes, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. E. Rognes, and G. N. Wells, *Archive of Numerical Software* **3** (2015), URL <https://doi.org/10.11588/ans.2015.100.20553>.
61. A. Logg, K. Mardal, G. N. Wells, et al., *Automated Solution of Differential Equations by the Finite Element Method* (Springer, 2012).
62. M. S. Alnaes, A. Logg, K. B. Igaard, M. E. Rognes, and G. N. Wells, *ACM Transactions on Mathematical Software* **40** (2014).
63. R. C. Kirby and A. Logg, *ACM Transactions on Mathematical Software* **32** (2006).
64. A. Logg, K. B. Igaard, M. E. Rognes, and G. N. Wells, in *Automated Solution of Differential Equations by the Finite Element Method*, edited by A. Logg, K. Mardal, and G. N. Wells (Springer, 2012), vol. 84 of *Lecture Notes in Computational Science and Engineering*, chap. 11.
65. T. Betcke and M. W. Scroggs, *IEEE Computing in Science & Engineering* **23**, 18 (2021).
66. P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, et al., *Nature Methods* **17**, 261 (2020).
67. C. Geuzaine and J.-F. Remacle, *International journal for numerical methods in engineering* **79**, 1309 (2009).
68. J. W. Ponder and D. A. Case, *Advances in protein chemistry* **66**, 27 (2003).
69. S. Decherchi and W. Rocchia, *PloS one* **8**, e59744 (2013).
70. C. T. Lee, J. G. Laughlin, J. B. Moody, R. E. Amaro, J. A. McCammon, M. Holst, and P. Rangamani, *Biophysical journal* **118**, 1003 (2020).

71. S. Hang, *ACM Trans. Math. Softw* **41**, 11 (2015).
72. W. Geng, S. Yu, and G. Wei, *The Journal of chemical physics* **127** (2007).
73. D. Chen, Z. Chen, C. Chen, W. Geng, and G.-W. Wei, *Journal of computational chemistry* **32**, 756 (2011).
74. A. H. Boschitsch and M. O. Fenley, *Journal of chemical theory and computation* **7**, 1524 (2011).
75. J. Wang, C. Tan, Y.-H. Tan, Q. Lu, and R. Luo, *Commun Comput Phys* **3**, 1010 (2008).
76. Q. Cai, M.-J. Hsieh, J. Wang, and R. Luo, *Journal of Chemical Theory and Computation* **6**, 203 (2010).
77. W. Rocchia, E. Alexov, and B. Honig, *The Journal of Physical Chemistry B* **105**, 6507 (2001).
78. D. D. Nguyen, B. Wang, and G.-W. Wei, *Journal of computational chemistry* **38**, 941 (2017).
79. J. Sørensen, M. O. Fenley, and R. E. Amaro, in *Computational Electrostatics for Biological Applications* (Springer, 2015), pp. 39–71.
80. J. M. J. Swanson, S. A. Adcock, and J. A. McCammon, *Journal of Chemical Theory and Computation* **1**, 484 (2005).
81. T. J. Dolinsky, J. E. Nielsen, J. A. McCammon, and N. A. Baker, *Nucleic Acids Research* **32**, W665 (2004).
82. T. Betcke, M. Bosy, and E. Burman, *Numerical Algorithms* **91**, 997 (2022).
83. T. Wang, R. Yokota, and L. A. Barba, *Journal of Open Source Software* **6**, 3145 (2021).
84. S. Kailasa, T. Wang, L. A. Barba, and T. Betcke, *arXiv preprint arXiv:2303.08394* (2023).