Published in final edited form as:

J Chem Theory Comput. 2009 ; 5(12): 3260–3278.

A “Reverse-Schur” Approach to Optimization With Linear PDE Constraints: Application to Biomolecule Analysis and Design

Jaydeep P. Bardhan[†], Michael D. Altman[‡], B. Tidor^{¶,§}, and Jacob K. White[¶]

[†]Department of Molecular Biophysics and Physiology, Rush University Medical Center, Chicago IL

[‡]Merck Research Laboratories, Boston, MA

[¶]Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge MA

[§]Department of Biological Engineering

Abstract

We present a partial-differential-equation (PDE)-constrained approach for optimizing a molecule’s electrostatic interactions with a target molecule. The approach, which we call *reverse-Schur co-optimization*, can be more than two orders of magnitude faster than the traditional approach to electrostatic optimization. The efficiency of the co-optimization approach may enhance the value of electrostatic optimization for ligand-design efforts—in such projects, it is often desirable to screen many candidate ligands for their viability, and the optimization of electrostatic interactions can improve ligand binding affinity and specificity. The theoretical basis for electrostatic optimization derives from linear-response theory, most commonly continuum models, and simple assumptions about molecular binding processes. Although the theory has been used successfully to study a wide variety of molecular binding events, its implications have not yet been fully explored, in part due to the computational expense associated with the optimization. The co-optimization algorithm achieves improved performance by solving the optimization and electrostatic simulation problems simultaneously, and is applicable to both unconstrained and constrained optimization problems. Reverse-Schur co-optimization resembles other well-known techniques for solving optimization problems with PDE constraints. Model problems as well as realistic examples validate the reverse-Schur method, and demonstrate that our technique and alternative PDE-constrained methods scale very favorably compared to the standard approach. Regularization, which ordinarily requires an explicit representation of the objective function, can be included using an approximate Hessian calculated using the new BIBEE/P (boundary-integral-based electrostatics estimation by preconditioning) method.

1 INTRODUCTION

The problem of optimizing electrostatic interactions is a task of particular importance in molecular design. One asks whether a candidate designed molecule, or *ligand*, is optimal for binding the target molecule, which is called a *receptor*, and if not, what chemical modifications might be made to improve binding affinity or specificity. A variety of factors contribute to the binding free energy, including conformational entropy, although often the contributions are dominated by packing effects and electrostatics. Although the short-range

Correspondence to: B. Tidor; Jacob K. White.

Supporting Information Available

None available.

packing interactions can be conceptualized relatively easily, analysis of the electrostatic component is more complex. The electrostatic component of the binding free energy can be particularly non-intuitive due to the interactions' long range and the trade-off between favorable ligand–receptor interactions in the bound state and the unfavorable desolvation penalties paid on binding.¹ These non-intuitive features have led to the important but challenging goal of designing optimal electrostatic interactions as an approach to designing useful molecular binding partners.^{1,2} Questions in molecular biology regarding the evolution of biomolecules, whether to serve specific functions or to bind targets with high affinity and specificity, may also be interpreted as questions regarding optimization of a particular objective function.^{2–4}

Lee and Tidor presented the first work describing the possibility of optimizing electrostatic interactions between molecules,¹ showing that linear-response theory and simple assumptions about binding events—in particular, that the ligand binds rigidly and that no charge redistribution occurs on binding—give rise to a quadratic model for the electrostatic contribution to the binding free energy. Their primarily analytical study used a multipole-based representation of the ligand charge distribution and spherical geometries for the unbound ligand and the ligand–receptor complex. Chong *et al.* applied this theory to an idealized model of the protein barnase and found that small sets of biochemically reasonable charge distributions resembled the computed optimal charge distribution.⁵ Kangas and Tidor later proved that the electrostatic component of the binding free energy is a convex function under reasonable assumptions and extended the theory to address nonspherical geometries, alternative basis sets, and measures of binding specificity.^{6–8}

Following these developments, Lee and Tidor studied the interactions between two proteins, the extremely tight-binding partners barnase and barstar;^{3,9} their analysis suggested that the inhibitor barstar is electrostatically optimized to bind to the enzyme barnase. In another application of the optimization theory, Kangas and Tidor studied the enzyme *B. subtilis* chorismate mutase.² This investigation indicated a particularly promising modification to improve the binding affinity of a transition-state analog inhibitor—the replacement of a carboxylate group by a nitro group. Mandal and Hilvert synthesized the proposed inhibitor; in agreement with the computational analysis, the resulting ligand bound the enzyme more tightly, and was in fact the tightest-binding chorismate mutase inhibitor reported to date in the literature.¹⁰

Several groups have applied the optimization theory to study a number of other molecular systems. Sulea and Purisima have studied cation–protein binding, the optimization of protein–protein interfaces, and the use of the charge optimization framework as a means to identify “hot spots” for binding.^{11,12} Sims *et al.* studied two protein kinases, protein kinase A (PKA) and cyclin-dependent kinase 2 (CDK2), and several inhibitors.¹³ Green and Tidor have applied charge optimization theory to two systems.^{4,14} In one, they demonstrated that glutaminyl-tRNA synthetase is optimized for its substrates;⁴ more recently, they proposed optimization-theory-based mutations to 5-Helix, which inhibits HIV-1 membrane fusion by gp41.¹⁴ Armstrong *et al.* have studied several inhibitors of neuraminidase and the relation between charge optimization and lead progression.¹⁵ Gilson has explored the theory allowing the optimization of flexible ligands.¹⁶ Schreiber and collaborators have also focused on optimizing ligand–receptor electrostatics.^{17,18} Brock *et al.* have used a theory similar to Lee and Tidor's in their analysis of protein–protein complexes.¹⁹

The computational expense associated with optimization has limited the broad application of the electrostatic-optimization theory. Traditional approaches to optimization begin with an explicit representation of the second-derivative, or *Hessian*, matrix or a means by which to multiply a vector by the Hessian. For electrostatic optimization problems, a large cost

associated with calculating the Hessian matrix explicitly is the successive simulation of the bound and unbound systems with each of the point charges (more generally, the basis functions) used to describe the ligand charge distribution.¹ The cost to form the Hessian must be paid before optimization can be performed, and it scales essentially linearly with the number of basis functions for reasonably-sized problems. Lee and Tidor emphasized, in their original optimization paper, the importance of using sufficiently complete basis function sets to achieve convergence to the optimum affinity.¹ However, the fixed-location point-charge basis sets used in most electrostatic optimization work offer little insight into geometric sensitivity or to basis-set completeness. The availability of more efficient computational methods may therefore enable not only greater numbers of ligands to be optimized in design efforts, but also a more thorough exploration of the optimization theory itself and the extent to which biology may have employed electrostatic optimization to achieve desired binding affinities and specificities.

This paper presents a new, highly efficient approach, which we call *reverse-Schur co-optimization*, to solving the electrostatic optimization problem. The theory and implementation of efficient methods for optimization problems constrained by partial differential equations (PDEs) have become a progressively more important research topic over the past several years,^{20–23} and we show that electrostatic optimization is actually a special case of a PDE-constrained optimization problem. Most PDE-constrained optimization techniques following one of two approaches. *All-at-once* approaches incorporate the PDE state variables (for electrostatic optimization, the electrostatic potentials in the bound and unbound states) directly into the optimization problem. The state variables together with the decision variables (the point charge values) satisfy the PDE, which is included as an equality constraint.^{22–25} Such techniques are often termed *simultaneous analysis and design* (SAND) approaches.²⁰ The second general strategy, sometimes called a *black-box* approach, hides the PDE from the optimization algorithm.²³ The *nested analysis and design* (NAND) paradigm is a black-box method,^{26,27} as are techniques that directly invert the PDE constraint before initiating optimization. The calculation of an explicit Hessian is effectively a black-box approach, because the mathematical details of the PDE simulation are entirely hidden from the optimization procedure. In electrostatic optimization problems, the decision variables and the state variables are related by a linear matrix equation. As we show in this paper, this linearity allows optimal charge distributions to be found without calculating the Hessian explicitly and without using the discretized PDE as an equality constraint.

To demonstrate the co-optimization method's performance on problems of therapeutic relevance, we have applied the optimization methodology to two protein–small-molecule ligand complexes. The first is a complex between HIV-1 protease and the small-molecule inhibitor darunavir (TMC-114).²⁸ HIV-1 protease is an essential enzyme in the life cycle of HIV, and small-molecule inhibitors of the protease have been successful components of combinatorial strategies for treating HIV infection.^{29,30} The second protein–ligand system studied is a complex between the protein cyclin-dependent kinase 2 (CDK2) and a small-molecule inhibitor.³¹ The CDK family of proteins involve regulating cell growth, and inhibitors of these enzymes are cancer therapies.^{32,33} It may be possible to use charge optimization to identify regions of these small-molecule ligands that are suboptimal for binding their protein target, and chemical modification at these locations may lead to improved inhibitors.

The following section describes a linear-response continuum model for biomolecule electrostatics, two boundary-integral formulations of the PDE problem, boundary-element methods (BEM) for solving the integral equations numerically, the optimization problem based on the linear-response model, and methods for convex quadratic optimization. Section

3 presents the new *reverse-Schur co-optimization* method. In addition, we describe two more widely used approaches to PDE-constrained optimization problems, partially to highlight differences between these methods and the reverse-Schur approach, and partially to illustrate that the performance gains are not necessarily specific to the reverse-Schur method. Techniques for constrained co-optimization are also presented. Important details of the implementation—regularization methods and preconditioning—are described in Section 4. In Section 5 we present computational results that validate the method, demonstrate its computational efficiency, and show that realistic problems in biomolecule design can be studied using PDE-constrained optimization methods. Section 6 summarizes the paper and suggests future research directions.

2 THEORY

2.1 A Linear-Response Model for Estimating the Electrostatic Contribution to the Free Energy of Binding Between Biomolecules

Free energies of binding are commonly estimated using a thermodynamic cycle such as that shown in Figure 1.34 The lower set of images represents the ligand, receptor, and ligand–receptor complex in aqueous solvent, and the lower horizontal arrow represents the binding free energy ΔG_{bind}^0 to be estimated. The unbound ligand and receptor and the bound complex are assumed to be at infinite dilution. The upper cartoons represent the three species in a homogeneous low-dielectric environment with zero ionic strength throughout, and the horizontal arrow denoted by $\Delta G_{\text{bind}}^{0,\text{ref}}$ represents the free energy change on binding in the low-dielectric environment, the electrostatic component of which is simply the ligand–receptor Coulomb-interaction energy. The three steps illustrated by vertical arrows involve the transfer of a molecule or complex between the low-dielectric environment and the solvent. The difference in a molecule's free energy as it is transferred into solvent from the reference low-dielectric medium is called its solvation free energy,³⁴ and this free energy is frequently decomposed into non-polar and electrostatic terms so that

$$\Delta G_{\text{solv}}^0 = \Delta G_{\text{solv}}^{0,\text{np}} + \Delta G_{\text{solv}}^{0,\text{es}}. \quad (1)$$

In many models, the non-polar free energy $\Delta G_{\text{solv}}^{0,\text{np}}$ is proportional to the molecular surface area, although recently more sophisticated models have been developed and parameterized (see, for example,^{35,36}). The electrostatic component $\Delta G_{\text{solv}}^{0,\text{es}}$ is often estimated using a macroscopic continuum electrostatic model,^{34,37} shown in Figure 2.

The molecule–solvent boundary, denoted as Γ in Figure 2, is taken to be the Richards molecular surface³⁸ and separates the molecular interior, region I , from the solvent exterior, region II . The interior is modeled as a homogeneous dielectric with low dielectric constant ϵ_I and a charge distribution $\rho(r)$; in this work, we assume that the charge distribution consists of n_c discrete point charges, the i th of which is located at r_i and has charge q_i . For many biomolecules, n_c ranges from a few dozen to several thousand. The electrostatic potential in region I , $\phi_I(r)$, satisfies a Poisson equation:

$$\nabla^2 \phi_I(r) = - \sum_{i=1}^{n_c} \frac{q_i}{\epsilon_I} \delta(r - r_i). \quad (2)$$

The solvent region is modeled as a homogeneous dielectric with high dielectric constant ϵ_{II} ; in this region, the electrostatic potential $\phi_{II}(r)$ satisfies the Laplace equation

$$\nabla^2 \varphi_{II}(r) = 0, \quad (3)$$

for non-ionic solutions, or for dilute ionic solutions, the linearized Poisson–Boltzmann equation (LPBE):

$$\nabla^2 \varphi_{II}(r) = \kappa^2 \varphi_{II}(r), \quad (4)$$

where κ is the inverse Debye length. The continuity of the potential and normal displacement furnish boundary conditions for both regions.³⁹ In the remainder of this paper, we assume that $\kappa = 0$, noting that the PDE-constrained optimization techniques apply equally well when the LPBE is used to model the potential in the solvent.

This set of coupled partial differential equations (PDEs) cannot be solved analytically except for relatively simple geometries. For general problems and realistic treatments of molecular geometries, numerical methods such as the finite-difference, finite-element, or boundary-element methods must be employed.^{39–61} In contrast to the finite-difference and finite-element methods, which discretize the differential form of the PDE, boundary-element methods discretize boundary-integral-equation formulations of the PDE problem.^{39,40,42,43,47,53–55,59,62–66} Boundary-integral formulations possess attractive theoretical and numerical properties such as reduced dimensionality, the possibility of exact treatment of the dielectric boundary, and exact treatment of point charge effects and boundary conditions at infinity.³⁹ Many integral-equation approaches have been described in the literature;^{39,40,42,47,55,62,67–69} in this paper, we present only the polarizable continuum model (PCM) formulation introduced by Miertus *et al.*^{40,51,70,71} (which was independently derived by Shaw and Zauhar, who also called it the apparent-surface-charge (ASC) formulation^{42,43,45}), and the formulation introduced by Yoon and Lenhoff.⁴⁷

Using the electrostatic model in Figure 2, the computational challenge associated with calculating $\Delta G_{solv}^{0,es}$, the electrostatic contribution to a solute's solvation free energy, is evaluating the *reaction-potentials* at each of the n_c charge locations that is induced by solvent polarization in response to the charges themselves. Because we have assumed linear response, the vector of reaction potentials at the charge locations, Φ_R , can be written as a weighted combination of the responses due to each of the individual point charges,

$$\varphi_R = S q, \quad (5)$$

where we have defined S to be the *reaction-potential*, or *solvation* matrix. The electrostatic free energy is then a quadratic function of q :

$$\Delta C_{solv}^{0,es} = \frac{1}{2} q^T S q. \quad (6)$$

We now derive expressions for the reaction-potential matrix S such that it may be written as

$$S = M_3 M_2^{-1} M_1, \quad (7)$$

where M_3 , M_2^{-1} , and M_1 are linear operators.

2.1.1 The Apparent-Surface-Charge Formulation—Numerous groups have derived the boundary-integral equation for the surface charge that develops at a dielectric boundary in response to a distribution of charge;^{40,42,72,73} it is known variously as the polarizable-

continuum model (PCM) and apparent-surface-charge (ASC) formulation,^{40,42,51,53} and has been widely used in biomolecular simulations.^{43,45,53,63,68,74} Rather than solving for the potential throughout space in the original mixed-dielectric PDE problem, one solves an equivalent problem with uniform dielectric constant ϵ_I everywhere, finding a distribution of charge $\sigma_p(r)$ on Γ such that $\sigma_p(r)$ reproduces the continuity conditions of the original mixed-dielectric problem. This surface charge satisfies the second-kind integral equation⁴²

$$\frac{\epsilon_I + \epsilon_{II}}{2\epsilon_I(\epsilon_I - \epsilon_{II})} \sigma_p(r) + \int_{\Gamma} \frac{\partial}{\partial n(r)} \frac{\sigma_p(r') dA'}{4\pi\epsilon_I \|r - r'\|} = - \frac{\partial}{\partial n(r)} \sum_{i=1}^{n_c} \frac{q_i}{4\pi\epsilon_I \|r - r_i\|}, \quad (8)$$

where $\oint$ denotes a principal-value integral^{75,76} and $n(r)$ denotes the outward normal direction into solvent. The surface charge distribution $\sigma_p(r)$ produces in the molecular interior (region I) a potential equal to that induced by the polarization of the solute. The *reaction potential* at a solute charge location r_i is the result of convolving the free-space Green's function with the surface charge distribution:

$$\varphi_R(r_i) = \int_{\Gamma} \frac{\sigma_p(r')}{4\pi\epsilon_I \|r_i - r'\|} dA'. \quad (9)$$

The set of reaction potentials at all the charge locations is therefore the image of the charge distribution under three linear operators:

$$\varphi_R = M_3 M_2^{-1} M_1 q. \quad (10)$$

The operator M_1 maps the solute charge distribution to the induced normal-displacement field at the dielectric boundary; that is, the application of M_1 to q generates the right-hand side (RHS) in Equation (8).

The operator M_2 generates the left-hand side in Equation (8) when applied to σ_p , and M_2^{-1} is used to denote the operator's inverse. That is, M_2^{-1} applied to the RHS in Equation (8) generates $\sigma_p(r)$. Finally, the integral operator M_3 maps the induced surface charge to the reaction potentials at the charge locations via Equation (9). Note that $M_3 M_2^{-1} M_1$ is an n_c -by- n_c matrix, even though M_1 , M_2 , and M_3 are operators.

Because the charge distribution is a set of discrete point charges, the difference in electrostatic free energy between the uniform ϵ_I domain and the mixed-dielectric problem is a finite-dimensional inner product:

$$\Delta C_{\text{solv}}^{0,\text{es}} = \frac{1}{2} \varphi_R^T q, \quad (11)$$

where φ_R denotes the vector of reaction potentials computed at the n_c charge locations. For problems in which $\epsilon_{II} > \epsilon_I$, $S = M_3 M_2^{-1} M_1$ is symmetric positive definite (SPD).

2.1.2 The Non-derivative Green's-Theorem Formulation—Using Green's theorem, Yoon and Lenhoff derived a pair of coupled integral equations capable of modeling solutes in dilute ionic solutions (that is, when the LPBE holds in the solvent).⁴⁷ The integral equations are

$$\frac{1}{2}\varphi(r) + \int_{\Gamma} \varphi(r') \frac{\partial G_I(r;r')}{\partial n} dA' - \int_{\Gamma} \frac{\partial \varphi}{\partial n}(r') G_I(r;r') dA' = \sum_{i=1}^{n_c} \frac{q_i}{\epsilon_I} G_I(r, r_i); \quad (12)$$

$$\frac{1}{2}\varphi(r) + \int_{\Gamma} \varphi(r') \frac{\partial G_{II}(r;r')}{\partial n} dA' - \frac{\epsilon_I}{\epsilon_{II}} \int_{\Gamma} \frac{\partial \varphi}{\partial n}(r') G_{II}(r;r') dA' = 0, \quad (13)$$

where $G_I(r;r')$ and $G_{II}(r;r')$ are the free-space Green's functions in the solute and solvent

regions and the unknown surface variables $\varphi(r)$ and $\frac{\partial \varphi}{\partial n}(r)$ are the surface potential and its normal derivative. These equations are derived by applying Green's theorem in regions *I* and *II*, finding the potential at arbitrary points in these regions by substituting the relevant Green's functions, and then letting the points approach the surface by taking appropriate

limits.^{47,64,66} After solving Equation (12) and Equation (13) for φ and $\frac{\partial \varphi}{\partial n}$, the reaction potential at the *i*th charge location induced by solvent polarization can be written as

$$\varphi_R(r_i) = \int_{\Gamma} \left[G_I(r_i;r') \frac{\partial \varphi}{\partial n}(r') - \varphi(r') \frac{\partial G_I}{\partial n}(r_i;r') \right] dA', \quad (14)$$

and again the electrostatic solvation free energy can be written as a product of three linear operators as in Equation (7).

2.2 Numerical Solution of the Integral Equations Using Boundary-Element Methods and Fast Algorithms

The boundary-element method (BEM) is a popular technique for solving boundary-integral equations numerically. To solve an integral equation such as Equation (8) using the BEM, one first introduces a set of basis functions defined on the surface. Representing the unknown surface variable as a weighted combination of the basis functions reduces the exact infinite-dimensional problem to an approximation problem with a finite number of unknowns, the weights used to scale the basis functions. A set of constraints on the weights are then written to force the approximate representation of the surface variable to satisfy the discretized integral equation as closely as possible in some metric (see, for example,⁷⁵). The resulting problem—that of finding the basis function weights that minimize some function of the residual—is a finite-dimensional matrix equation.

Usually, it is convenient to discretize the surface into a set of surface patches, or *boundary elements*, before defining the basis functions. In biomolecule electrostatic simulations, these elements are commonly planar triangles,^{47,63} although curved-element discretizations of molecule–solvent interfaces have been described by several groups.^{39,56,59,61,77,78} We present a boundary-element method for solving the ASC formulation. The Green's-theorem formulation (Equation (12) and Equation (13)) can be solved analogously, but requires two weights for each basis function: one for the potential and one for its normal derivative. Full details for solving the Green's-theorem formulation numerically can be found in.^{47,64,66}

First, the molecule–solvent interface is discretized using n_p boundary elements, and then a set of n_p piecewise-constant basis functions is defined such that

$$\chi_i(r) = \begin{cases} 1 & \text{if } r \text{ is on panel } i \\ 0 & \text{otherwise.} \end{cases} \quad (15)$$

The unknown surface charge density $\sigma_p(r)$ is then represented approximately as

$$\sigma_p(r) \approx \sum_{i=1}^{n_p} x_i \chi_i(r), \quad (16)$$

where the weights x_i are unknown. Using a Galerkin discretization⁷⁵ of the PCM/ASC formulation in which the inner integral is evaluated via one-point quadrature,^{68,79,80} one obtains the dense linear system $M_2 x = M_1 q$, with the entries of M_2 and M_1 given by

$$M_{2,ii} = \frac{\widehat{\varepsilon}}{2\varepsilon_i} \alpha_i + \int_{\text{panel } i} \frac{\partial}{\partial n(r)} \frac{\alpha_i dA'}{4\pi\varepsilon_i \|r - r_{ci}\|} \quad (17)$$

$$M_{2,ij} = \int_{\text{panel } i} \frac{\partial}{\partial n(r)} \frac{\alpha_j dA'}{4\pi\varepsilon_i \|r - r_{cj}\|} \quad (i \neq j) \quad (18)$$

$$M_{1,ij} = - \int_{\text{panel } i} \frac{\partial}{\partial n(r)} \frac{q_j dA'}{4\pi\varepsilon_i \|r - r_{cj}\|}, \quad (19)$$

where α_i denotes the area of panel i , $\widehat{\varepsilon} = \frac{\varepsilon_i + \varepsilon_{ii}}{\varepsilon_i - \varepsilon_{ii}}$, $n(r)$ denotes the outward normal at r , and r_{ci} denotes the centroid of panel i . The approach presented here differs slightly from the commonly used centroid-collocation method, which essentially approximates the outer Galerkin integral using one-point quadrature; the method described here offers superior accuracy.^{68,79} We note that the matrix entries of Equation (17), Equation (18), and Equation (19) are specific to the PCM/ASC formulation; if the Green's theorem formulation,^{47,66} or other boundary-integral formulations, are employed to define the solvation matrix, similar matrices are defined that play analogous roles.⁸⁰

Protein-sized systems often require more than 10^5 unknowns and boundary elements to accurately represent the molecule-solvent interfaces and surface variables. Because solving the n -dimensional dense matrix equation $M_2 x = M_1 q$ using LU factorization requires $O(n^3)$ time, and even storing M_2 requires prohibitively large $O(n^2)$ memory, more efficient methods have been developed whose time and memory requirements scale linearly or near-linearly in the number of unknowns.^{54,63,65,81,82} These fast-solver approaches combine Krylov-subspace iterative methods⁸³ such as GMRES⁸⁴ with fast, approximate algorithms to apply the discretized integral operator matrix to a vector. At the k th iteration of a Krylov-subspace algorithm, one finds an approximate solution x_k that lies in the k th Krylov subspace, which is formed by repeatedly applying A to b :

$$x_k \in \{b, Ab, A^2b, \dots, A^{k-1}b\}. \quad (20)$$

The fast multipole method^{81,85} and the precorrected-FFT algorithm⁸⁶ represent two algorithms that compute BEM matrix-vector products in linear or near-linear time. The results in this paper were computed using the FFTSVD algorithm, which offers several advantages for biomolecule electrostatic problems.^{65,66}

Krylov-subspace iterative methods are commonly preconditioned so that instead of solving $Ax = b$ for x , one solves $P Ax = P b$, where P is a matrix that approximates A^{-1} such that the iterates x_k converge more rapidly towards the exact solution; that is, the use of a

preconditioner reduces the number of matrix–vector products required to find a suitably accurate approximation to the actual solution x . Preconditioning the ASC formulation is easily accomplished using a diagonal matrix in which $P_{ii}=M_{2,ii}^{-1}$. Methods for preconditioning the Green's-theorem formulation are presented in. 64·66

2.3 Biomolecule Electrostatic Optimization

Using the thermodynamic cycle in Figure Figure 1, the total electrostatic contribution to the binding free energy can be written as

$$\Delta G_{\text{bind}}^{0,\text{es}} = \left(-\Delta G_{\text{solv}}^{0,R,\text{es}} - \Delta G_{\text{solv}}^{0,L,\text{es}} \right) + \Delta G_{\text{bind}}^{0,\text{ref},\text{es}} + \Delta G_{\text{solv}}^{0,L:R,\text{es}}, \quad (21)$$

where the solvation free energies for the ligand, receptor, and complex are denoted by the subscripts L , R , and $L : R$, and the ligand–receptor Coulomb-interaction energy is written $\Delta G_{\text{bind}}^{0,\text{ref},\text{es}}$. Substituting the appropriate ligand, receptor, and ligand–receptor reaction potential matrices, one obtains

$$\Delta G_{\text{bind}}^{0,\text{es}} = -\frac{1}{2}q_R^T R_{\text{unbound}} q_R - \frac{1}{2}q_L^T L_{\text{unbound}} q_L + (Gq_R)^T q_L + \frac{1}{2}q_C^T C_{\text{bound}} q_C, \quad (22)$$

where q_L and q_R denote the n_{cL} - and n_{cR} -length vectors of ligand and receptor charge values, $q_C = (q_L, q_R)^T$, and L_{unbound} , R_{unbound} , and C_{bound} denote the appropriate solvation matrices; the electrostatic component of the low-dielectric binding free energy has been written $(Gq_R)^T q_L$, where the n_{cL} -by- n_{cR} Coulomb matrix G maps receptor-charge values to Coulomb potentials at the ligand-charge locations given the bound-state geometry.

The optimizable component of $\Delta G_{\text{bind}}^{0,\text{es}}$, which is the portion of $\Delta G_{\text{bind}}^{0,\text{es}}$ that is dependent on the ligand charges, is called the variational electrostatic binding free energy $\Delta G_{\text{bind}}^{0,\text{var}}$. The second term in Equation (22) does not contribute to $\Delta G_{\text{bind}}^{0,\text{var}}$, nor does the component of the final term that depends only on the receptor charges. Writing $\frac{1}{2}q_C^T C_{\text{bound}} q_C$ as

$$\Delta G_{\text{solv},L-R}^{0,\text{es}} = \frac{1}{2} \begin{bmatrix} q_L^T & q_R^T \end{bmatrix} \begin{bmatrix} L_{\text{bound}} & C_{\text{bound}}^{L,R} \\ C_{\text{bound}}^{R,L} & R_{\text{bound}} \end{bmatrix} \begin{bmatrix} q_L \\ q_R \end{bmatrix} \quad (23)$$

and exploiting the symmetry of C_{bound} allows the variational electrostatic binding free energy to be written as

$$\Delta G_{\text{bind}}^{0,\text{var}} = -\frac{1}{2}q_L^T L_{\text{unbound}} q_L + \frac{1}{2}q_L^T L_{\text{bound}} q_L + q_R^T G^T q_L + q_R^T C_{\text{bound}}^{L,R} q_L. \quad (24)$$

The final two terms in Equation (24) are linear in the ligand charge values, and the vector

$$c = Gq_R + C_{\text{bound}}^{L,R} q_R, \quad (25)$$

which represents the total field induced by the receptor charges at the ligand-charge locations in the bound state, may be used to further simplify Equation (24):

$$\Delta G_{\text{bind}}^{0,\text{var}} = \frac{1}{2}q_L^T (L_{\text{bound}} - L_{\text{unbound}}) q_L + c^T q_L. \quad (26)$$

Equation (26) is the objective function for optimizing the electrostatic component of the free energy of binding. Kangas and Tidor showed that the difference $L_{\text{bound}} - L_{\text{unbound}}$, which is the Hessian of the objective function, is positive definite if one assumes that the ligand binds rigidly, that the ligand charge distribution is unchanged on binding, and that the molecules have finite size.⁶ The variational electrostatic binding free energy $\Delta G_{\text{bind}}^{0,\text{var}}$ is therefore a convex function with respect to the ligand charge distribution, and there exists a unique minimal free energy.

Often, it is of interest to enforce sum-of-charge constraints over subsets of the charges and possibly over the entire set. Defining the matrix $H = L_{\text{bound}} - L_{\text{unbound}}$ and including the linear constraint $Aq = b$ gives rise to the constrained optimization problem

$$\begin{aligned} &\text{minimize} \quad \frac{1}{2}q^T Hq + c^T q \\ &\text{subject to} \quad Aq = b. \end{aligned} \quad (27)$$

In Equation (27) and for the remainder of the paper the vector q is used instead of q_L to represent the ligand charges. In addition to linear equality constraints, linear inequality constraints are sometimes imposed on the variables to ensure that the computed charges are physically reasonable. The resulting inequality-constrained problem

$$\begin{aligned} &\text{minimize} \quad \frac{1}{2}q^T Hq + c^T q \\ &\text{subject to} \quad Aq = b \\ &\quad \text{and} \quad m \leq q_i \leq M, \forall i \in \{1, \dots, n_c\}, \end{aligned} \quad (28)$$

where m_i and M_i represent the lower and upper bounds for q_i . Assuming without loss of generality that A has full rank, this problem can be transformed into the standard form for a convex quadratic problem,

$$\begin{aligned} &\text{minimize} \quad \frac{1}{2}x^T \hat{H}x + \hat{c}^T x \\ &\text{subject to} \quad \hat{A}x = \hat{b} \\ &\quad \text{and} \quad x \geq 0, \end{aligned} \quad (29)$$

using the substitutions

$$\begin{aligned} x &= \begin{bmatrix} t \\ r \end{bmatrix} \\ \hat{H} &= \begin{bmatrix} H & 0 \\ 0 & 0 \end{bmatrix} \\ \hat{c} &= \begin{bmatrix} c + Hm \\ 0 \end{bmatrix} \\ \hat{A} &= \begin{bmatrix} A & 0 \\ I & I \end{bmatrix} \\ \hat{b} &= \begin{bmatrix} b - Am \\ M - m \end{bmatrix}, \end{aligned} \quad (30)$$

where the slack variables t and r satisfy $m + t = q$ and $q + r = M$. This notation for inequality-constrained biomolecule optimization problems will be used throughout the rest of the paper.

Titratable chemical groups in the ligand warrant a brief discussion. One of the assumptions inherent in the electrostatic optimization theory is that the ligand geometry and charge distribution do not change on binding.¹ Thus, the charge optimization scheme presented

here, the number of charges n_c remains unchanged during optimization, as do their locations, and therefore charge optimization is performed for a particular titration state. In reality, of course, the ligand–receptor binding event may perturb the ligand geometry, its charge distribution, or state of protonation, or any combination of these. Exploring the dependence of the optimal charges (and the optimized binding free energy) on the titration state is the most obvious way to assess the impact of the assumptions underlying the theory, although in general such an undertaking is likely to be computationally expensive.

2.4 Solving Convex Quadratic Optimization Problems

This section presents methods for minimizing the quadratic function

$$f(x) = \frac{1}{2}x^T Hx + c^T x, \quad (31)$$

where x is a vector of length n_{primal} and the matrix H , also known as the *Hessian* matrix, is symmetric and positive definite (SPD). The global minimizer x^* can be found by setting $\nabla f(x) = 0$ and solving the resulting linear system

$$Hx^* = -c. \quad (32)$$

Optimization problems with constraints require more sophisticated approaches (see, for example, 87,88). The quadratic program

$$\begin{aligned} &\text{minimize} && \frac{1}{2}x^T Hx + c^T x \\ &\text{subject to} && Ax = b \end{aligned} \quad (33)$$

can be solved using Lagrange multipliers.⁸⁷ The optimal solution is a point x^* and a corresponding vector of multipliers λ^* that together satisfy the matrix equation

$$\begin{bmatrix} H & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} x^* \\ \lambda^* \end{bmatrix} = \begin{bmatrix} -c \\ b \end{bmatrix}. \quad (34)$$

Inequality-constrained problems require the introduction of a vector of slack variables s in addition to the Lagrange multipliers λ . Because the optimization problem in Equation (29) satisfies a constraint qualification,⁸⁷ an optimal solution (x^*, λ^*, s^*) can be calculated by finding a point that satisfies the Karush-Kuhn-Tucker (KKT) optimality conditions:

$$\begin{aligned} s^* &= Hx^* + c - A^T \lambda^* \\ Ay^* &= b \\ x_i^* s_i^* &= 0 \quad \forall i \in \{1, 2, \dots, n_{\text{primal}}\} \\ (x^*, s^*) &\geq 0. \end{aligned} \quad (35)$$

These conditions can be interpreted as the zeros of the nonlinear function

$$F(x, \lambda, s) = \begin{bmatrix} Hx + c - A^T \lambda - s \\ b - Ax \\ Xs \end{bmatrix}, \quad (36)$$

where X represents the diagonal matrix with $X_{ii} = x_i$. Primal-dual interior point methods find the roots of this equation using a modified Newton–Raphson iteration, with the Newton–Raphson updates biased to ensure convergence and scaled to ensure positivity of the

elements of x and s .⁸⁸ The k th update of a primal-dual iterative method satisfies the linear system of equations

$$\begin{bmatrix} H & -A^T & -I \\ A & 0 & 0 \\ S^k & 0 & X^k \end{bmatrix} \begin{bmatrix} \Delta x^{k+1} \\ \Delta \lambda^{k+1} \\ \Delta s^{k+1} \end{bmatrix} = \begin{bmatrix} -c + s^k - Hx^k + A^T \lambda^k \\ b - Ax^k \\ X^k S^k e \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \sigma \frac{x^{k,T} s^k}{n_{\text{primal}}} e \end{bmatrix}, \quad (37)$$

where e is the vector of ones, S^k is diagonal with $S_{ii}^k = s_i^k$, and the second term on the right-hand side (RHS) biases the update towards a point with equal pairwise products $x_i s_i$.⁸⁸ The parameter σ , which is between 0 and 1, determines the strength of the bias. It can be a fixed value over all iterations, or set dynamically based on the progress of the previous iterations.⁸⁹ Smaller values for σ allow faster convergence in most cases, but larger values offer superior robustness.⁸⁸

The reverse-Schur optimization method is specialized to PDE-constrained problems in which the relationships between the decision variables x , the PDE state variables y_I and the external state variables y_E are affine. That is, the three vectors satisfy a matrix equality

$$\begin{bmatrix} B & A & 0 \\ D & C & -I \end{bmatrix} \begin{bmatrix} x \\ y_I \\ y_E \end{bmatrix} = \begin{bmatrix} z_I \\ z_E \end{bmatrix} \quad (38)$$

for some vectors z_I and z_E ; in the electrostatic optimization, $z_I = 0$ and $z_E = 0$.

3 THE REVERSE-SCHUR METHOD FOR ELECTROSTATIC OPTIMIZATION

Some problems in computational science can be solved more efficiently using a Schur complement, such that one solves not a block linear system such as

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} a \\ b \end{bmatrix}, \quad (39)$$

but rather a smaller (or better-conditioned) system like

$$(D - CA^{-1}B)y = b - CA^{-1}a. \quad (40)$$

The reverse-Schur co-optimization method uses the exactly opposite approach. An unconstrained quadratic program

$$\text{minimize } \frac{1}{2} x^T H x + c^T x, \quad (41)$$

in which the Hessian is of the form $H = M_3 M_2^{-1} M_1$, can be solved by setting the gradient equal to zero, leading to the linear system

$$M_3 M_2^{-1} M_1 x^* = -c. \quad (42)$$

Equation (42) resembles Equation (40) with $D = 0$ and $a = 0$, and therefore the Hessian may be said to have the structure of a Schur complement. In reverse-Schur co-optimization, one solves the larger “reverse-Schur-complement” system

$$\begin{bmatrix} 0 & M_3 \\ -M_1 & M_2 \end{bmatrix} \begin{bmatrix} x^* \\ y \end{bmatrix} = \begin{bmatrix} -c \\ 0 \end{bmatrix}. \quad (43)$$

As discussed in Section 2, the Hessian matrix $H = L_{\text{bound}} - L_{\text{unbound}}$ is a difference of two matrices, each of which has Schur structure. The reverse-Schur form of the unconstrained electrostatic optimization problem therefore has two reverse-Schur complements, and the optimizing ligand distribution q^* can be found by solving

$$\begin{bmatrix} 0 & M_{3,b} & -M_{3,u} \\ -M_{1,b} & M_{2,b} & 0 \\ -M_{1,u} & 0 & M_{2,u} \end{bmatrix} \begin{bmatrix} q^* \\ y_b \\ y_u \end{bmatrix} = \begin{bmatrix} -c \\ 0 \\ 0 \end{bmatrix}, \quad (44)$$

where the subscripts b and u denote the bound and unbound systems and the variables y_b and y_u are the surface variables for the corresponding BEM problems when the ligand charge distribution is the optimizing distribution q^* . In the ASC formulation, for instance, y_b represents the weights for the bound-state basis functions. The bound- and unbound-geometry state variables are therefore found simultaneously with the optimal decision variables q^* . The co-optimization approach in this respect resembles “all-at-once” methods; however, the corresponding state variables are not explicitly included in the optimization problem. Section 3.1.2 details that approach to Hessian-implicit optimization. It is important to note that the matrices associated with reverse-Schur co-optimization are not necessarily symmetric. Also, we note that the bound-state and unbound-state surface are different, with the bound state representing the larger ligand–receptor complex; consequently, the vector y_b is typically a longer vector than y_u .

As discussed in Section 2, electrostatic charge optimization problems usually have many fewer decision variables than there are degrees of freedom associated with the BEM problems. Reverse-Schur co-optimization systems such as Equation (44) are therefore only slightly larger than the corresponding BEM systems. Efficient preconditioning strategies, which will be described in Section 4, allow systems such as Equation (44) to be solved with approximately the same computational cost as would be required to solve a single bound-state and a single unbound-state electrostatic problem. We emphasize that the dense boundary-element matrices $M_{2,b}$ and $M_{2,u}$ are almost always too large to be calculated or stored, and so their LU factorizations cannot be computed. Forming H in Equation (41) explicitly requires n_c solves of the bound and unbound geometries. The reverse-Schur method is therefore more computationally efficient than the explicit-Hessian approach.

The next section presents alternative Hessian-free PDE-constrained methods, such as a nested-Krylov method, or NAND-like approach, and a traditional PDE-constrained formulation, following the SAND paradigm. Section 3.2 describes biomolecule co-optimization techniques for constrained problems.

3.1 Alternative PDE-Constrained Approaches

3.1.1 Nested-Krylov Approach—One alternative to the co-optimization approach would be to use nested Krylov methods to solve the linear systems associated with the explicit-Hessian approach. The unconstrained problem

$$(M_{3,b}M_{2,b}^{-1}M_{1,b} - M_{3,u}M_{2,u}^{-1}M_{1,u})q^* = -c \quad (45)$$

would then require two inner Krylov solves for every matrix–vector multiplication required for the outer Krylov method: one for the bound-state problem and one for the unbound-state

problem. This approach represents an implementation of a *nested analysis and design* (NAND) approach to PDE-constrained optimization.

3.1.2 Incorporating the PDE as Constraints—The electrostatic optimization problem can also be formulated in a traditional PDE-constrained approach in which the surface variables of the bound- and unbound-state boundary-element problems are included as optimization variables, with the boundary-element method equations added as equality constraints. The resulting problem,

$$\begin{aligned} &\text{minimize } \frac{1}{2} \begin{bmatrix} q \\ y_b \\ y_u \end{bmatrix}^T \begin{bmatrix} 0 & \frac{1}{2}M_{3,b} & -\frac{1}{2}M_{3,u} \\ \frac{1}{2}M_{3,b}^T & 0 & 0 \\ -\frac{1}{2}M_{3,u}^T & 0 & 0 \end{bmatrix} \begin{bmatrix} q \\ y_b \\ y_u \end{bmatrix} + \begin{bmatrix} c \\ 0 \\ 0 \end{bmatrix}^T \begin{bmatrix} q \\ y_b \\ y_u \end{bmatrix} \\ &\text{subject to } \begin{bmatrix} -M_{1,b} & M_{2,b} & 0 \\ -M_{1,u} & 0 & M_{2,u} \end{bmatrix} \begin{bmatrix} q \\ y_b \\ y_u \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \end{aligned} \quad (46)$$

can be solved by setting the gradient to zero and solving the resulting linear system

$$\begin{bmatrix} 0 & \frac{1}{4}M_{3,b} & -\frac{1}{4}M_{3,u} & -M_{1,b}^T & -M_{1,u}^T \\ \frac{1}{4}M_{3,b}^T & 0 & 0 & M_{2,b}^T & 0 \\ -\frac{1}{4}M_{3,u}^T & 0 & 0 & 0 & M_{2,u}^T \\ -M_{1,b} & M_{2,b} & 0 & 0 & 0 \\ -M_{1,u} & 0 & M_{2,u} & 0 & 0 \end{bmatrix} \begin{bmatrix} q^* \\ y_b^* \\ y_u^* \\ \lambda_b^* \\ \lambda_u^* \end{bmatrix} = \begin{bmatrix} -c \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}. \quad (47)$$

The matrix in Equation (47) is symmetric, which allows the system to be solved using specialized Krylov-subspace iterative methods. However, every matrix–vector product requires twice as much calculation as the matrix–vector products required for Krylov methods to solve Equation (44). It can also be difficult or impractical to apply the transposed matrices $M_{2,b}^T$ and $M_{2,u}^T$. Thus, both the reverse-Schur method and the SAND-type approach have strengths and weaknesses.

3.2 Constrained Co-Optimization

It is straightforward to use the reverse-Schur method to solve the constrained optimization problems presented in Section 2.4. After transforming Equation (34), the co-optimization system for solving the problem with linear-equality constraints is

$$\begin{bmatrix} 0 & A^T & M_{3,b} & -M_{3,u} \\ A & 0 & 0 & 0 \\ -M_{1,b} & 0 & M_{2,b} & 0 \\ -M_{1,u} & 0 & 0 & M_{2,u} \end{bmatrix} \begin{bmatrix} q^* \\ \lambda^* \\ y_b \\ y_u \end{bmatrix} = \begin{bmatrix} -c \\ b \\ 0 \\ 0 \end{bmatrix}. \quad (48)$$

Similarly, the co-optimization system associated with the k th iteration of a primal-dual interior-point method is transformed from Equation (37) to

$$\begin{bmatrix} 0 & -\widehat{A}^T & -I & \widehat{M}_{3,b} & -\widehat{M}_{3,u} \\ \widehat{A} & 0 & 0 & 0 & 0 \\ S^k & 0 & X^k & 0 & 0 \\ -\widehat{M}_{1,b} & 0 & 0 & M_{2,b} & 0 \\ -\widehat{M}_{1,u} & 0 & 0 & 0 & M_{2,u} \end{bmatrix} \begin{bmatrix} \Delta x^{k+1} \\ \Delta \lambda^{k+1} \\ \Delta s^{k+1} \\ \Delta y_b^{k+1} \\ \Delta y_u^{k+1} \end{bmatrix} = \begin{bmatrix} -\widehat{c} - s^k - \widehat{H}x^k + \widehat{A}^T \lambda^k \\ \widehat{b} - \widehat{A}x^k \\ X^k S^k e + \sigma \frac{y^{k,T} s^k}{n_c} e \\ \widehat{M}_{1,b} x^k - M_{2,b} y_b^k \\ \widehat{M}_{1,u} x^k - M_{2,u} y_u^k \end{bmatrix}, \quad (49)$$

where

$$\widehat{M}_{1,b} = \begin{bmatrix} M_{1,b} & 0 \end{bmatrix} \quad (50)$$

$$\widehat{M}_{3,b} = \begin{bmatrix} M_{3,b} \\ 0 \end{bmatrix} \quad (51)$$

with the zero submatrices of the appropriate size given the transformation of the inequality-constrained problem into standard form, and $\widehat{M}_{1,u}$ and $\widehat{M}_{3,u}$ are similarly defined. The term $\widehat{H}x^k$ on the right-hand side is computed as

$$\widehat{H}x^k = \widehat{H}x^{k-1} + \widehat{M}_{3,b} \Delta y_b^k - \widehat{M}_{3,u} \Delta y_u^k \quad (52)$$

and $\widehat{H}x^0$ must be found before the first iteration.

4 IMPLEMENTATION

In this section we present two implementation details that are important for the co-optimization method to achieve high efficiency and accuracy. First, regularization of the optimization presents a critical challenge. The need for regularization arises due to numerical error in simulation; although the exact Hessian is positive definite, typically a numerically computed H is not. Given H explicitly, the eigendecomposition or singular value decomposition can be used to penalize or eliminate the nonphysical part of the matrix. However, for the reverse-Schur (or any other implicit-Hessian) method to produce results comparable to those obtained by current methods, accurate regularization techniques must be found so that the appropriate directions can be penalized. Second, effective preconditioning schemes need to be developed because the co-optimization linear systems are solved using Krylov iterative methods.

A simple model geometry, shown in Figure 3, is used to demonstrate the performance of the presented regularization and preconditioning methods. The unbound ligand is a sphere of radius 8 Å; the ligand–receptor complex is modeled as the solvent-excluded surface produced by rolling a 1.4-Å probe sphere over the union of the ligand sphere and a 24-Å radius sphere representing the receptor, where the sphere centers are separated by 20 Å.³⁸ The internal dielectric constant is taken to be 4 and the solvent external dielectric constant to be 80. For simplicity in this paper, we assume that the Laplace equation holds in the solvent region (that is, that there are no mobile ions in solution). However, the implicit-Hessian methods can be used also for the case when the LPBE holds in the solvent region.

The receptor charge distribution consists of 2000 randomly placed charges located such that no charge is within 1.5 Å of another charge, the dielectric boundary, or the ligand volume; the charge values have been chosen from a uniform distribution on $[-1e, 1e]$. The ligand charge distribution is built from a set of charge locations randomly placed in the ligand sphere subject to the constraints that no charges are within 1.5 Å of one another or of the

ligand boundary. Several discretizations of each geometry have been generated. One surface discretization uses planar triangle boundary elements generated using MSMS,⁹⁰ with 71922 and 7924 elements used to approximate the bound and unbound surfaces; the other uses 2132 and 1810 curved boundary elements that exactly represent the two geometries.⁷⁸ Coarser discretizations have been used for some examples and are described where appropriate. Full computational details are deferred to Section 5.

4.1 Regularization

The explicit-Hessian approach to biomolecule electrostatic optimization allows straightforward regularization. The numerically calculated Hessian matrix is first symmetrized to remove numerical-error-based asymmetries. Because the Hessian is available explicitly, it is possible to calculate its eigendecomposition. The eigenspace corresponding to the smallest eigenvectors is then heavily penalized, but not removed explicitly, so that the existence of a feasible solution is assured even in the presence of other constraints. The co-optimization approach, in contrast, does not permit the use of direct eigendecomposition, because the Hessian matrix is not available. Co-optimization achieves its performance advantage by leaving the Hessian implicit, and therefore little information is available regarding its spectrum or corresponding eigenvectors. In order to solve the same optimization problems using the current protocols without sacrificing performance, the minimal eigenspace must be approximated inexpensively.

In,⁸⁹ linear-equality constrained co-optimization problems were preconditioned using an approximate Hessian $\hat{H}$ of the form

$$\hat{H} = M_{3,b} P_{2,b} M_{1,b} - M_{3,u} P_{2,u} M_{1,u} \quad (53)$$

where $P_{2,b}$ and $P_{2,u}$ represent the bound- and unbound-state BEM preconditioners. Using a diagonal approximation of the ASC integral equation as the preconditioner P_2 corresponds to the recent BIBEE/P electrostatic model.⁹¹ When performing co-optimization using the Green's theorem formulation, the operators M_1 , M_2 , M_3 , and P_2 differ from those employed in the ASC formulation (i.e., the entries of M_2 are no longer defined according to Equation (17)). However, the Green's theorem co-optimization linear systems are written the same way, and an approximate Hessian can still be defined according to Equation (53).

To illustrate that an approximation to the ASC formulation can be used to regularize the optimization method but not the Green's-theorem method approximation, explicit Hessians and their approximations were calculated using both integral formulations and both planar and curved boundary elements. The four Hessian matrices and their approximations were decomposed using the singular value decomposition (SVD), and the right singular vectors of the approximate Hessians were projected onto those from the corresponding explicitly calculated Hessians. The Green's-theorem and ASC formulations produced very similar explicit actual Hessians ($\|H_A - H_G\|/\|H_A\| < 0.01$), regardless of whether planar or curved boundary elements were used.

To ensure comparable regularization between explicit-Hessian and implicit-Hessian optimization procedures, implicit-Hessian methods must penalize not only the same number of search directions as explicit-Hessian methods, but also the search directions themselves. To illustrate how the singular vectors of an approximate Hessian $\hat{H}$ are aligned with the singular vectors of the actual Hessian H , we calculate the matrix

$$V_H^T V_{\hat{H}}, \quad (54)$$

which represents the projection of the singular vectors of $\hat{H}$ onto those of H . Perfect alignment between the sets of vectors would produce a diagonal matrix X whose diagonal entries all have unit magnitude. Similarly, the degree to which the singular vectors of $\hat{H}$ are imperfectly aligned with those of H is reflected in the presence of non-zero entries off the diagonal. Figure 4(a) is a pseudo-color plot of the magnitudes of the entries of X , using approximate and actual Hessians computed from the Green's-theorem formulation and the planar-element discretization. The analogous plot, computed using the ASC formulation and planar elements, is shown in Figure 4(b). The ASC-based approximate singular vectors are clearly much better aligned with the corresponding singular vectors of the explicitly computed ASC Hessian. Similarly, plotted in Figures 5(a) and (b) are the results of curved-element simulations of the Green's-theorem and ASC formulations, respectively, with the approximate Hessian right singular vectors projected onto the explicit-Hessian right singular vectors. Because the ASC formulation generates superior approximate Hessians to the mixed formulation for both kinds of discretizations, we attribute the fidelity to the superior conditioning of purely second-kind integral operators (see, for instance, 56). In principle, it is possible to regularize the Green's-theorem co-optimization system using the penalty matrix derived from the ASC-based approximate Hessian. However, the results in Section 4.2 illustrate that the superior conditioning of the second-kind integral equation produces convergence of the co-optimization GMRES in many fewer iterations than are required for the Green's-theorem based co-optimization.

Figure 6(a) and (b) are plots of the singular values for the planar- and curved-element discretizations. Predictability of the relation between the approximate singular values $\hat{\sigma}_i$ and the actual values σ_i is important so that the appropriate number of search directions can be penalized. The singular values of $\hat{H}_A$ were much closer to those of H_A than the singular values of $\hat{H}_G$ were to H_G .

Based on the results in Figures 4, Figure 5, and Figure 6 we adopted the following scheme to regularize the co-optimization solutions. The ASC-based Hessian approximation

$\hat{H}_A$ was computed first, and the eigendecomposition of the symmetrized matrix $\frac{1}{2}(\hat{H}_A + \hat{H}_A^T)$ was taken. The first right singular vector $\hat{v}_1$ was multiplied by the Hessian H_A using BEM simulation of the bound- and unbound-states, and the Rayleigh quotient

$$\hat{\lambda}_1 = \hat{v}_1^T H_A \hat{v}_1 \quad (55)$$

was then used as an estimate for the maximum eigenvalue of H_A . The penalty matrix

$$W = \alpha V_{[i,I]} V_{[i,I]}^T \quad (56)$$

was then created, where the penalty parameter $\alpha = 100 \text{ kcal/mol/e}^2$, the eigenvalue tolerance $\gamma = 10^{-4}$, and the set of penalized directions $I = \{i | \hat{\sigma}_i < \gamma \hat{\lambda}_1\}$. The quadratic penalty term

$\frac{1}{2} q^T W q$ was then added to the objective function, and optimization could begin. The unconstrained co-optimization system with a penalty matrix is

$$\begin{bmatrix} W & A_{3,b} & -A_{3,u} \\ -A_{1,b} & A_{2,b} & 0 \\ -A_{1,u} & 0 & A_{2,u} \end{bmatrix} \begin{bmatrix} q^* \\ \sigma_{p,b} \\ \sigma_{p,u} \end{bmatrix} = \begin{bmatrix} -c \\ 0 \\ 0 \end{bmatrix}, \quad (57)$$

and the systems for constrained problems are similarly modified.

Even though $\hat{\lambda}_1$ usually approximated λ_1 to within a few percent, sometimes the number of directions penalized was slightly different from the number that would be penalized in an explicit-Hessian method using the same tolerance γ . As a result, it was desirable to have an inexpensive means of obtaining approximations to the optimal distributions for problems where different numbers of directions were penalized. For the unconstrained and linear-equality constrained problems, such estimates could be obtained using an approximation to the Sherman-Woodbury-Morrison formula

$$(H + UV^T)^{-1} = H^{-1} - H^{-1}U(I + V^TH^{-1}U)^{-1}V^TH^{-1}, \quad (58)$$

which specifies how the inverse of a matrix H changes when H is perturbed by the low-rank update UV^T . Update-approximation methods for inequality-constrained problems represent an area of current research.

4.2 Preconditioning

The Hessian-implicit linear systems in Equation (44), Equation (48), and Equation (49) share a similar structure, and therefore can be preconditioned by similar methods. For a problem with n_{primal} decision variables and Σn_{SV} total unknowns associated with the BEM simulations, we define the desolvation operators

$$\widehat{J}_1 = \begin{bmatrix} -M_{1,b} & 0 \\ -M_{1,u} & 0 \end{bmatrix} \quad (59)$$

$$\widehat{J}_2 = \begin{bmatrix} M_{2,b} & 0 \\ 0 & M_{2,u} \end{bmatrix} \quad (60)$$

$$\widehat{P}_2 = \begin{bmatrix} P_{2,b} & 0 \\ 0 & P_{2,u} \end{bmatrix} \quad (61)$$

$$\widehat{J}_3 = \begin{bmatrix} M_{3,b} & -M_{3,u} \\ 0 & 0 \end{bmatrix}, \quad (62)$$

where the zero blocks are sized such that $\widehat{M}_1 \in \mathbb{R}^{\Sigma n_{SV} \times n_{\text{primal}}}$ and $\widehat{M}_3 \in \mathbb{R}^{n_{\text{primal}} \times \Sigma n_{SV}}$.

We define the preconditioners by block-factorizing the corresponding linear systems using the BEM preconditioners rather than the inverses of the BEM matrices. The Hessian-implicit preconditioners would therefore be exact if the BEM preconditioners were actually the BEM matrix inverses. The resulting preconditioners can be written as the product $\tilde{P} = \tilde{P}_4 \tilde{P}_3 \tilde{P}_2 \tilde{P}_1$. Other preconditioners, which for example block-triangularize the Hessian-implicit linear systems, can also be used but are generally less effective than the block factorization. Figures Figure 7(a) and (b) are plots of the preconditioned relative GMRES residuals as a function of iteration count for the unconstrained problem using the the Green's theorem formulation and the apparent surface-charge formulation. Each solves the same 50-charge unconstrained problem using curved-element discretizations of the model problem in Figure Figure 3, using different preconditioners. It is clear that the ASC co-optimization converges in many fewer iterations than the corresponding Green's-theorem co-optimization, regardless of which approach to preconditioning is employed.

5 COMPUTATIONAL RESULTS

5.1 Efficiency of Co-Optimization and PDE-constrained Approaches

The standard “all-at-once” approach and nested-Krylov approach to the electrostatic optimization problem were implemented in MATLAB using the geometry in Figure Figure 3 and relatively coarse discretizations of 142 and 124 curved elements for the bound and unbound geometries. The methods’ implementations were verified by direct inspection of the optimal charges computed by dense factorization of the systems in Equation (45) and Equation (47). These techniques and the co-optimization solver were then used with preconditioned GMRES to solve a set of unconstrained problems of varying dimension. Computational expense was measured by counting the total number of applications of the BEM operator $\hat{M}_2$, because the computational cost of optimization is dominated by the application of the BEM operators. The all-at-once solver required two $\hat{M}_2$ matrix–vector products for every Krylov iteration.

The number of required applications for the nested-Krylov method was estimated using the fact that the nested-Krylov method, which relies on an implicit Hessian, and an explicit-Hessian Krylov method, require the same number of GMRES iterations to achieve convergence. Therefore we estimated the nested-Krylov computational expense using the explicitly-calculated Hessian, rather than a true nested-Krylov code. The GMRES solve was preconditioned using the ASC-based Hessian approximation $\hat{H}_A$. The all-at-once system was preconditioned using a block-factorization method similar to the co-optimization preconditioning schemes, so that the all-at-once preconditioner would be exact if the BEM preconditioner were exact. Every GMRES iteration for the all-at-once system requires two applications of the BEM operator $\hat{M}_2$, as shown in Equation (47). Figure Figure 8 is a plot of the computational cost of these methods for solving unconstrained problems as the number of optimization variables varies from 5 to 130. The PDE-constrained approaches scale very favorably compared to the explicit-Hessian approach, and exhibit essentially comparable reductions in cost. It should be noted that the performance differences between the PDE-constrained approaches may reflect the relatively unoptimized implementations of the methods, and no significant conclusions should be drawn regarding the merits of one PDE-constrained technique over another. We note also that constrained problems exhibit similar performance trends.^{89,92}

5.2 Realistic Protein–Ligand Systems

The three-dimensional structures of a complex between HIV-1 protease and darunavir (accession code 1T3R) and a complex between CDK2 and a small-molecule inhibitor (accession code 1OIT) were obtained from the Protein Data Bank (PDB). For both structures, protein side chains with missing density were rebuilt in their default geometry using the CHARMM computer program,⁹³ and in cases of multiple occupancy, the first entry listed was used. The final chi angles for as-paragine, glutamine, and histidine side chains were flipped by 180 degrees as necessary to improve the hydrogen bonding network. Hydrogen atoms were added to both structures using the HBUILD module⁹⁴ of CHARMM and the PARAM22 parameter set,⁹⁵ using a distance-dependent dielectric constant of 4. Ionizable residues were left in their standard states at pH 7. In the case of HIV-1 protease, the catalytic dyad was left doubly deprotonated. The receptor protonation states are assumed to be the same in both the bound and unbound states. For electrostatic simulations, atomic radii were taken from the PARSE parameter set.⁹⁶ Partial atomic charges for protein atoms were also taken from the PARSE parameter set; quantum-mechanically derived partial atomic charges for the small molecule inhibitors were calculated as follows. The geometry of each small molecule inhibitor was optimized using quantum mechanical calculations at the RHF/6–31G* level of theory as implemented in the program Gaussian 98.⁹⁷ After

geometry optimization, partial atomic charges were fit to the quantum mechanical electrostatic potential using the RESP methodology.⁹⁸⁻⁹⁹ Both the CDK2 and HIV-protease systems were optimized using the co-optimization method and an explicit Hessian; the same curved-element discretizations were used for both methods. As in Section 4, solute-solvent interfaces were defined as solvent-excluded (molecular) surfaces using a probe of radius 1.4 Å, and the solute and solvent dielectric constants were 4 and 80, respectively.

5.2.1 CDK2 and Inhibitor—The CDK2 inhibitor described by Anderson *et al.* has 40 atoms.³¹ Optimization of the partial atomic charges for a small-molecule inhibitor of CDK2 did not lead to significantly improved predicted electrostatic binding free energy; this inhibitor appears to already be very well optimized for its protein target. From finite-difference simulations, the total electrostatic contribution to binding for the quantum-mechanically-derived (wild-type) charge distribution, which has net zero charge, is 8.75 kcal/mol, and the optimal charge distribution leads to an electrostatic binding free energy of 5.54 kcal/mol.

Figure 9 is a plot of the inhibitor with RESP-derived nominal charge values (in red) and the co-optimization unconstrained optimal charge values (in blue) labeling each atom, computed using curved boundary elements. It can be seen that the inhibitor atoms that directly hydrogen bond to the protein, especially those in the aminopyrimidine core, have optimal partial atomic charges that closely match those determined through quantum mechanics.

The wild-type molecule has zero net charge, and the co-optimization optimal solution from Figure 10 has a net charge of $-0.26 e$. For comparison, the explicit-Hessian boundary-element approach leads to an unconstrained optimum with $-0.33 e$ net charge, and the finite-difference optimal charges sum to $0.06 e$; the preponderance of the difference is localized in a small number of atoms whose optimal charges are of large magnitude (Figure 10). As expected, optimal charges computed using the explicit Hessian and the co-optimization methods were nearly identical; three unconstrained optimizations were performed, and the results are shown in Figure 10. Explicit-Hessian calculations were performed using both finite-difference methods¹⁰⁰ and boundary-element methods,⁶⁶ as was the boundary-element based co-optimization approach. The first instructive comparison, of the explicit-Hessian approaches, demonstrates that even the vastly different approaches to numerical simulation produce optimal charges that correspond very closely. The methods must give exactly the same results in the limit of infinitely fine discretizations, of course, but it is valuable to know that such good agreement can be obtained even for discretizations that can be easily solved on a personal workstation. The excellent agreement between the boundary-element methods demonstrates the correct implementation of the co-optimization approach and that numerical errors associated with the implicit representation of the dense boundary-integral operators do not materially affect the computed optimal solution.

The total charge on the inhibitor was also constrained to different net charge values. For these problems with different constraints, very little change was observed in the partial atomic charges for atoms directly interacting with the receptor, and the calculated optimal binding free energy (the objective function at the optimum) changed minimally. Figure 11 shows the box-constrained optimal charges when the total inhibitor charge was constrained to be -1 , 0 , or $+1$, computed using co-optimization. The co-optimization charges again correspond closely with calculations performed using either boundary-element or finite-difference methods with explicit Hessians (data not shown). The finite-difference calculations gave rise to optimized binding free energies of 6.88 kcal/mol for the $-1 e$ -constrained problem, 5.54 for the zero-charge problem, and 6.55 for the $+1 e$ -constrained problem.

The partial atomic charge values of the solvent-exposed sulfonamide group changed the most to accommodate these different net charges, largely because their solvent-exposed nature results in small desolvation penalties on binding. It should be noted that the $0.85e$ bound on the sulfur partial charge may be too stringent. These box constraints were introduced following earlier work³ and the observation that few biomolecular systems are modeled as having partial charges of larger magnitude, regardless of whether the partial charges are taken from molecular mechanics force fields such as CHARMM95 or derived from electronic structure calculations and RESP fitting. Our primary purpose here, however, is to demonstrate that the co-optimization method is fully capable of treating inequality constraints using a primal-dual interior point method. Also, the constraining of the total charge to either -1 , 0 , or $+1$ is not meant to test protonation states, given that the same number of charges and the same set of charge locations are used in each test. Optimization under these varying constraints suggests whether the optimized binding affinity or the optimal charge distribution is sensitive to the overall inhibitor charge, and for this problem neither appears to be the case.

The unbound- and bound-state geometries consisted of 3,821 and 138,770 curved boundary elements, respectively. Calculating the explicit Hessian required 461 applications of the bound-state integral operator $M_{2,b}$. In contrast, using reverse-Schur co-optimization, the unconstrained and linear-equality constrained problems each required at most $2 \hat{M}_2$ matrix–vector products, which is essentially the same cost as $2 M_{2,b}$ matrix–vector products owing to the small size of the unbound system. The more than 200-fold reduction in the number of matrix–vector products for such a small problem suggests that the model problems based on Figure 3 may actually be more computationally challenging than realistic problems, because a model problem of comparable size showed an acceleration of only a factor of ten using co-optimization over an explicit-Hessian method (see Figure 8).

This example illustrates one weakness of our chosen metric for performance improvement—the reduction in the number of $\hat{M}_2$ matrix–vector products. This metric neglects the startup cost that must be paid regardless of whether one intends to use explicit-Hessian or co-optimization methods and thus the overall compute time required to obtain an unconstrained optimum is not being reduced by a factor of 200. Our decision to use the number of $\hat{M}_2$ matrix–vector products as an improvement metric is based on the consideration that it is impossible to perform the optimization without paying the initialization cost. Thus, from a theoretical perspective it is appropriate to neglect this cost in comparing optimization methods.

Nevertheless, as a practical matter, it is important to understand the impact of co-optimization on overall computational cost. On a 2-GHz Intel MacBook pro, the planar-boundary-element simulations require approximately 100 seconds of setup time, and the bound-state simulations require an average 5 seconds (6 GMRES iterations are generally required for bound-state simulations). The overall unconstrained co-optimization cost is thus about 102 seconds (100 seconds for initialization and 2 GMRES iterations to solve the unconstrained co-optimization problem), whereas the overall optimization cost for the BEM explicit-Hessian approach is about 300 seconds (100 seconds for initialization and 5 seconds for calculating each of the 40 columns of the Hessian). Thus the total reduction in computational cost is almost a factor of three. For large problems such as optimizing protein–protein interactions, it can be expected that the total reduction will be even larger.

5.2.2 HIV-1 Protease and Inhibitor—The 75-atom inhibitor darunavir binds tightly to HIV-1 protease.²⁸ In Figure 12, the atoms are labeled with indices corresponding to the entries in Table 1, which lists the RESP-derived charge values, the unconstrained co-optimized partial atomic charge values, optimal charges under a zero total charge

equality constraint, and the optimal charges computed with box constraints such that no charge exceeded $0.85e$ in magnitude, with sum-of-charge constraints set to -1 , 0 , and 1 .

We emphasize that the equality constraints have not been introduced to evaluate protonation effects, but only to estimate the influence of total charge on the optimal solution (and the associated binding free energy). Our results illustrate that the faster co-optimization method generates results consistent with the traditional approach. The unbound- and bound-state geometry discretizations consisted of 5,892 and 133,067 curved boundary elements. Computation of the explicit Hessian required 576 applications of the operator $M_{2,b}$ and unconstrained co-optimization required 15 $M_{2,b}$ matrix–vector products. In Figure 13 are plotted the unconstrained optimal charges computed using the explicit-Hessian and the co-optimization methods and again, the answers agree extremely well.

Electrostatic optimization of darunavir in the HIV-1 protease active site led to a significant improvement in the predicted electrostatic binding free energy. The ligand is net neutral and the wild-type charge distribution gives an finite-difference-calculated electrostatic binding free energy of 27.54 kcal/mol, and the unconstrained optimal solution computed using co-optimization gives an electrostatic binding free energy of 5.48 kcal/mol. The resulting optimized binding free energies for the bound-constrained problems were: for the $-1 e$ problem, 10.40 kcal/mol; for the neutral problem, 6.89 kcal/mol; for the $+1 e$ problem, 5.73 kcal/mol.

The improvement on optimization can be attributed mainly to an accumulation of positive charge in the center of the ligand, near the negatively charged aspartyl dyad. These atoms are buried in both the bound and unbound states due to the molecular shape, and consequently can take larger charge values without incurring significant desolvation penalties. For the unconstrained problems, the net ligand charge for the co-optimization method was $0.64 e$, $0.64 e$ for the curved-boundary-element explicit-Hessian method, and $1.00 e$ using the finite-difference method.

Inhibitor atoms that make direct hydrogen-bonding interactions with the protease, such as the aniline nitrogen and hydrogen atoms (atom indices 1, 39, and 40), hydroxyl group (indices 18 and 57), and bis-tetrahydrofuran oxygen atoms (indices 26 and 28) have optimal charges very similar to their quantum-mechanically-derived values.

6 DISCUSSION

In this paper we have described an efficient technique, which we call reverse-Schur co-optimization, for calculating the molecular charge distribution that optimizes the electrostatic component of the free energy of binding to another molecule. The approach exhibits substantially better performance than traditional optimization approaches, which explicitly calculate the Hessian matrix before optimization. The co-optimization approach is a PDE-constrained optimization technique and exhibits comparable performance to alternative PDE-constrained optimization techniques that are well known in other areas of computational science and engineering. Although this paper has presented an approach based on a boundary-element method (BEM) for solving the electrostatics problem, no fundamental issues seem to preclude the use of other numerical methods in a co-optimization approach. The critical elements for efficient co-optimization appear to be the availability of good preconditioners and sufficiently accurate, but computationally inexpensive, Hessian approximations.

The reverse-Schur approach to PDE-constrained optimization is only possible because the PDE state variables and the decision variables are linearly related. This structure enables the PDE to be incorporated as a final algebraic manipulation before numerically solving the

linear systems associated with quadratic programming. The all-at-once and nested approaches PDE-constrained optimization techniques, in contrast, are much more flexible with respect to the relationships between the state and decision variables.

Regularization—the penalization of certain search directions associated with the smallest eigenvalues—is more complicated for PDE-constrained approaches than for methods that rely on explicit Hessians. However, the BIBEE/P approach to estimating electrostatic interactions⁹¹ has been demonstrated to generate a sufficiently accurate Hessian approximation whose eigendecomposition can be used as the basis for deriving a penalty function. The superior conditioning of purely second-kind integral-equation formulations^{56,75} relative to first-kind or mixed first-second-kind equations has an unexpected consequence in that the Yoon and Lenhoff formulation cannot be used to generate a preconditioner-based Hessian approximation.

A number of extensions to the co-optimization technique may make it still more efficient. For instance, primal-dual interior-point methods are less efficient than active-set methods for “warm start” problems, in which one begins optimizing from a neighborhood of the optimal solution. Coupling co-optimization to an active-set solver might therefore significantly reduce the cost required to solve problems that differ only by the inclusion of varying constraints. Furthermore, a co-optimization warm-start method may be faster than the present implementation because there exist ways (such as Gasteiger-Marsili charges¹⁰¹) to rapidly estimate a wild-type charge distribution that could be used as an initial guess for the optimal distribution. One might also save the Hessian-vector products as they are formed, in essence allowing the Hessian to be “built” such that after a sufficient number of optimizations have been performed, the solver is using a completely explicit Hessian.

The efficiency gains afforded by PDE-constrained methods in general, not just the reverse-Schur approach presented here, should allow significantly larger and more complex biological systems to be studied. Computational redesign of proteins, for instance, may produce optimization problems with dimension greater than one thousand. For these problems, co-optimization can offer a cost reduction of over two orders of magnitude;⁹² such an acceleration may allow the evaluation of many more candidate ligands or ligand poses. Also, geometry can now be varied to assess the best-case energetic cost of adding of a functional group as it changes a binding partner’s desolvation penalty, or the effects of molecular flexibility.¹⁶ For example, it may be computationally feasible to use co-optimization to study the influence of different protonation states on binding free energies and on the optimal charge distributions associated with each state. Finally, the original electrostatic optimization paper by Lee and Tidor noted an unexpected asymmetry between the receptor charge distribution and the calculated optimal ligand distribution.¹ Because co-optimization allows the use of substantially larger basis sets, it may be able to develop techniques that can identify optimal charge placement as well as value.

Acknowledgments

The authors thank J. Nocedal, J.-H. Lee, D. J. Willis, D. F. Green, M. Knepley, and S. Leyffer for useful discussions. J. P. Bardhan gratefully acknowledges support from a Department of Energy Computational Science Graduate Fellowship, as well as from a Wilkinson Fellowship in Scientific Computing funded by the Mathematical, Information, and Computational Sciences Division Subprogram of the Office of Advanced Scientific Computing Research, Office of Science, U. S. Dept. of Energy, under Contract DE-AC02-06CH11357. J. White acknowledges support from grants from the National Science Foundation and the Singapore–MIT Alliance. This work was supported in part by the National Institutes of Health (GM065418).

References

1. Lee L-P, Tidor B. J. Chem. Phys. 1997; 106:8681–8690.

2. Kangas E, Tidor B. J. Phys. Chem. B. 2001; 105:880–888.
3. Lee L-P, Tidor B. Nat. Struct. Biol. 2001; 8:73–76. [PubMed: 11135675]
4. Green DF, Tidor B. J. Mol. Biol. 2004; 342:435–452. [PubMed: 15327945]
5. Chong LT, Dempster SE, Hendsch ZS, Lee L-P, Tidor B. Protein Sci. 1998; 7:206–210. [PubMed: 9514276]
6. Kangas E, Tidor B. J. Chem. Phys. 1998; 109:7522–7545.
7. Kangas E, Tidor B. Phys. Rev. E. 1999; 59:5958–5961.
8. Kangas E, Tidor B. J. Chem. Phys. 2000; 112:9120–9131.
9. Lee L-P, Tidor B. Protein Sci. 2001; 10:362–377. [PubMed: 11266622]
10. Mandal A, Hilvert D. J. Am. Chem. Soc. 2003; 125:5598–5599. [PubMed: 12733875]
11. Sulea T, Purisima EO. J. Phys. Chem. B. 2001; 105:889–899.
12. Sulea T, Purisima EO. Biophys. J. 2003; 84:2883–2896. [PubMed: 12719221]
13. Sims PA, Wong CF, McCammon JA. J. Comput. Chem. 2004; 25:1416–1429. [PubMed: 15185335]
14. Green DF, Tidor B. Proteins. 2005; 60:644–657. [PubMed: 16001428]
15. Armstrong KA, Tidor B, Cheng AC. J. Med. Chem. 2006; 49:2470–2477. [PubMed: 16610790]
16. Gilson MK. J. Chem. Theory Comput. 2006; 2:259–270.
17. Selzer T, Albeck S, Schreiber G. Nat. Struct. Biol. 2000; 7:537–541. [PubMed: 10876236]
18. Shaul Y, Schreiber G. Proteins. 2005; 60:341–352. [PubMed: 15887221]
19. Brock K, Talley K, Coley K, Kundrotas P, Alexov E. Biophys. J. 2007; 93:3340–3352. [PubMed: 17693468]
20. Haftka R. AIAA J. 1985; 23:1099–1103.
21. Orozco C, Ghattas O. AIAA J. 1992; 30:1877–1885.
22. Biegler L, Nocedal J, Schmid C. SIAM J. Optim. 1995; 5:314–347.
23. Dennis J, Heinkenschloss M, Vicente L. SIAM J. Control Optim. 1998; 36:1750–1794.
24. Biros G, Ghattas O. SIAM J. Sci. Comput. 2005; 27:687–713.
25. Biros G, Ghattas O. SIAM J. Sci. Comput. 2005; 27:714–739.
26. Brayton RK, Hachtel GD, Sangiovanni-Vincentelli AL. Proc. IEEE. 1981; 69:1334–1362.
27. Fisher, M.; Nocedal, J.; Tremolet, Y.; Wright, SJ. Data assimilation in weather forecasting: a case study in PDE-constrained optimization. Optimization and Engineering. 2008 [accessed August 1, 2009]. DOI: 10.1007/s11081-008-9051-5. <http://www.springerlink.com/content/e47441q543236t31/>
28. Surleraux DL, Tahri A, Verschueren WG, Pille GM, de Kock HA, Jonckers TH, A. Peeters A, Meyer SD, Azijn H, Pauwels R, de Bethune MP, King NM, Prabu-Jeyabalan M, Schiffer CA, Wigerinck PB. J. Med. Chem. 2005; 48:1813–1822. [PubMed: 15771427]
29. Huff JR. J. Med. Chem. 1991; 34:2305–2314. [PubMed: 1875332]
30. Flexner C. N. Engl. J. Med. 1998; 338:1281–1292. [PubMed: 9562584]
31. Anderson M, Beattie J, Breault G, Breed J, Byth K, Culshaw J, Ellston R, Green S, Minshull C, Norman R, Paupit R, Stanway J, Thomas A, Jewsbury P. Bioorg. Med. Chem. Lett. 2003; 13:3021. [PubMed: 12941325]
32. Vandenheuve S, Harlow E. Science. 1993; 262:2050–2054. [PubMed: 8266103]
33. Hartwell LH, Kastan MB. Science. 1994; 266:1821–1828. [PubMed: 7997877]
34. Honig B, Sharp K, Yang AS. J. Phys. Chem. 1993; 97:1101–1109.
35. Levy RM, Zhang LY, Gallicchio E, Felts AK. J. Am. Chem. Soc. 2003; 125:9523–9530. [PubMed: 12889983]
36. Wagoner JA, Baker NA. Proc. Natl. Acad. Sci. USA. 2006; 103:8331–8336. [PubMed: 16709675]
37. Sharp KA, Honig B. Annu. Rev. Biophys. Biophys. Chem. 1990; 19:301–332. [PubMed: 2194479]
38. Richards FM. Annu. Rev. Biophys. Bioeng. 1977; 6:151–176. [PubMed: 326146]
39. Juffer AH, Botta EFF, van Keulen BAM, van der Ploeg A, Berendsen HJC. J. Comput. Phys. 1991; 97:144–171.

40. Miertus S, Scrocco E, Tomasi J. *J. Chem. Phys.* 1981; 55:117–129.
41. Warwicker J, Watson HC. *J. Mol. Biol.* 1982; 157:671–679. [PubMed: 6288964]
42. Shaw PB. *Phys. Rev. A.* 1985; 32:2476–2487. [PubMed: 9896362]
43. Zauhar RJ, Morgan RS. *J. Mol. Biol.* 1985; 186:815–820. [PubMed: 4093987]
44. Klapper I, Hagstrom R, Fine R, Sharp K, Honig B. *Proteins.* 1986; 1:47–59. [PubMed: 3449851]
45. Zauhar RJ, Morgan RS. *J. Comput. Chem.* 1988; 9:171–187.
46. Gilson MK, Honig B. *Proteins.* 1988; 4:7–18. [PubMed: 3186692]
47. Yoon BJ, Lenhoff AM. *J. Comput. Chem.* 1990; 11:1080–1086.
48. Nicholls A, Honig B. *J. Comput. Chem.* 1991; 12:435–445.
49. Holst, MJ. Ph.D.thesis. Univ. of Ill. at Urbana-Champaign: 1993.
50. You TJ, Harvey SC. *J. Comput. Chem.* 1993; 14:484–501.
51. Cammi R, Tomasi J. *J. Comput. Chem.* 1995; 16:1449–1458.
52. Madura JD, Briggs JM, Wade RC, Davis ME, Luty BA, Ilin A, Antosiewicz J, Gilson MK, Bagheri B, Ridgway-Scott L, McCammon JA. *Comput. Phys. Comm.* 1995; 91:57–95.
53. Purisima EO, Nilar SH. *J. Comput. Chem.* 1995; 16:681–689.
54. Bharadwaj R, Windemuth A, Sridharan S, Honig B, Nicholls A. *J. Comput. Chem.* 1995; 16:898–913.
55. Cancès E, Mennucci B, Tomasi J. *J. Chem. Phys.* 1997; 107:3032–3041.
56. Liang J, Subramaniam S. *Biophys. J.* 1997; 73:1830–1841. [PubMed: 9336178]
57. Holst M, Baker N, Wang F. *J. Comput. Chem.* 2000; 21:1319–1342.
58. Rocchia W, Alexov E, Honig B. *J. Phys. Chem. B.* 2001; 105:6507–6514.
59. Bordner AJ, Huber GA. *J. Comput. Chem.* 2003; 24:353–367. [PubMed: 12548727]
60. Boschitsch AH, Fenley MO, Zhou H-X. *J. Phys. Chem. B.* 2002; 106:2741–2754.
61. Boschitsch AH, Fenley MO. *J. Comput. Chem.* 2004; 25:935–955. [PubMed: 15027106]
62. Zhou HX. *Biophys. J.* 1993; 65:955–963. [PubMed: 8218918]
63. Purisima EO. *J. Comput. Chem.* 1998; 19:1494–1504.
64. Kuo, SS.; Altman, MD.; Bardhan, JP.; Tidor, B.; White, JK. Fast Methods for Simulation of Biomolecule Electrostatics. *Proc. Int. Conf. Comput.-Aided Des. (ICCAD)*; San Jose, CA. New York, NY: ACM; 2002. p. 466-473.
65. Altman MD, Bardhan JP, Tidor B, White JK. *IEEE Trans. Comput.-Aided Des.* 2006; 25:274–284.
66. Altman MD, Bardhan JP, White JK, Tidor B. *J. Comput. Chem.* 2009; 30:132–153. [PubMed: 18567005]
67. Chipman DM. *J. Chem. Phys.* 2004; 120:5566–5575. [PubMed: 15267432]
68. Altman, MD.; Bardhan, JP.; White, JK.; Tidor, B. An Efficient and Accurate Surface Formulation for Biomolecule Electrostatics in Non-ionic Solution. *Proc. 2005 IEEE Eng. Med. Biol. Conf. (IEEE-EMBS) 2005*; Shanghai. Piscataway, NJ: IEEE; 2005. p. 7591-7595.
69. Grandison S, Penfold R, Vanden-Broeck J-M. *J. Comput. Phys.* 2007; 224:663–680.
70. Barone V, Cossi M, Tomasi J. *J. Chem. Phys.* 1997; 107:3210–3221.
71. Mennucci B, Cancès E, Tomasi J. *J. Phys. Chem. B.* 1997; 101:10506–10517.
72. Levitt DG. *Biophys. J.* 1978; 22:209–219. [PubMed: 656542]
73. Attard P. *J. Chem. Phys.* 2003; 119:1365–1372.
74. Boda D, Voliskó M, Eisenberg B, Nonner W, Henderson D, Gillespie D. *J. Chem. Phys.* 2006; 125 034901.
75. Atkinson, KE. *The Numerical Solution of Integral Equations of the Second Kind.* Cambridge University Press; 1997.
76. Hsiao GC, Wendland WL. *Encyclopedia of Computational Mechanics.* 2004
77. Zauhar RJ. *J. Comput.-Aided Mol. Des.* 1995; 9:149–159. [PubMed: 7608746]
78. Bardhan JP, Altman MD, White JK, Tidor B. *J. Chem. Phys.* 2007; 127 014701.
79. Tausch J, Wang J, White J. *IEEE Trans. Comput.-Aided Des.* 2001; 20:1398–1405.
80. Bardhan JP. *J. Chem. Phys.* 2009; 130 094102.

81. Nabors K, Korsmeyer FT, Leighton FT, White J. *SIAM J. Sci. Comput.* 1994; 15:713–735.
82. Lu BZ, Cheng XL, Huang J, McCammon JA. *Proc. Natl. Acad. Sci. USA.* 2006; 103:19314–19319. [PubMed: 17148613]
83. Golub, GH.; Loan, CFV. *Matrix Computations*. 3rd ed.. Baltimore, MD: The Johns Hopkins University Press; 1996.
84. Saad Y, Schultz M. *SIAM J. Sci. Stat. Comput.* 1986; 7:856–869.
85. Greengard L, Rokhlin V. *J. Comput. Phys.* 1987; 73:325–348.
86. Phillips JR, White JK. *IEEE Trans. Comput.-Aided Des.* 1997; 16:1059–1072.
87. Bertsekas, DP. *Nonlinear Programming*. 2nd ed.. Nashua, NH: Athena Scientific; 1999.
88. Wright, SJ. *Primal-Dual Interior Point Methods*. Philadelphia, PA: SIAM; 1997.
89. Bardhan, JP.; Lee, JH.; Altman, MD.; Benson, S.; Leyffer, S.; Tidor, B.; White, JK. Biomolecule Electrostatic Optimization with an Implicit Hessian. In: Laudon, M.; Romanowicz, B., editors. *Tech. Proc. 2004 Nanotech. Conf. Trade Show*, v. 1; Boston MA. 2004; Cambridge MA: NSTI; 2004. p. 164-167.
90. Sanner M, Olson AJ, Spehner JC. *Biopolymers.* 1996; 38:305–320. [PubMed: 8906967]
91. Bardhan JP. *J. Chem. Phys.* 2008; 129 144105.
92. Bardhan, JP.; Lee, JH.; Kuo, SS.; Altman, MD.; Tidor, B.; White, JK. Fast Methods for Biomolecule Charge Optimization. In: Laudon, M.; Romanowicz, B., editors. *Tech. Proc. 2003 Nanotech. Conf. Trade Show*, v. 2; San Francisco, CA. 2003; Cambridge, MA: NSTI; 2003. p. 508-511.
93. Brooks BR, Bruccoleri RE, Olafson BD, States DJ, Swaminathan S, Karplus M. *J. Comput. Chem.* 1983; 4:187–217.
94. Brünger AT, Karplus M. *Proteins.* 1988; 4:148–156. [PubMed: 3227015]
95. MacKerell AD Jr, Bashford D, Bellott M, Dunbrack RL Jr, Evanseck JD, Field MJ, Fischer S, Gao J, Guo H, Ha S, Joseph-McCarthy D, Kuchnir L, Kuczera K, Lau FTK, Mattos C, Michnick S, Ngo T, Nguyen DT, Prodhom B, Reiher WE III, Roux B, Schlenkrich M, Smith JC, Stote R, Straub J, Watanabe M, Wiorkiewicz-Kuczera J, Yin D, Karplus M. *J. Phys. Chem. B.* 1998; 102:3586–3616.
96. Sitkoff D, Sharp KA, Honig B. *J. Phys. Chem. B.* 1994; 98:1978–1988.
97. Frisch, MJ.; Trucks, GW.; Schlegel, HB.; Scuseria, GE.; Robb, MA.; Cheeseman, JR.; Zakrzewski, VG.; Montgomery, JA., Jr; Stratmann, RE.; Burant, JC.; Dapprich, S.; Millam, JM.; Daniels, AD.; Kudin, KN.; Strain, MC.; Farkas, O.; Tomasi, J.; Barone, V.; Cossi, M.; Cammi, R.; Mennucci, B.; Pomelli, C.; Adamo, C.; Clifford, S.; Ochterski, J.; Petersson, GA.; Ayala, PY.; Cui, Q.; Morokuma, K.; Salvador, P.; Dannenberg, JJ.; Malick, DK.; Rabuck, AD.; Raghavachari, K.; Foresman, JB.; Cioslowski, J.; Ortiz, JV.; Baboul, AG.; Stefanov, BB.; Liu, G.; Liashenko, A.; Piskorz, P.; Komáromi, I.; Gomperts, R.; Martin, RL.; Fox, DJ.; Keith, T.; Al-Laham, MA.; Peng, CY.; Nanayakkara, A.; Challacombe, M.; Gill, PMW.; Johnson, B.; Chen, W.; Wong, MW.; Andres, JL.; Gonzalez, C.; Head-Gordon, M.; Replogle, ES.; Pople, JA. *Gaussian 98*. Pittsburgh, PA: Gaussian, Inc.; 1998.
98. Green DF, Tidor B. *J. Phys. Chem. B.* 2003; 107:10261–10273.
99. Bayly CI, Cieplak P, Cornell WD, Kollman PA. *J. Phys. Chem.* 1993; 97:10269–10280.
100. Altman MD. Unpublished results.
101. Rizzo RC, Aynechi T, Case DA, Kuntz ID. *J. Chem. Theory Comput.* 2006; 2:128–139.

Figure 1.

A thermodynamic cycle for estimating binding free energies. The shaded regions on the lower set of cartoons represent aqueous solvent. The upper cartoons represent a uniform low dielectric (the same as that of the ligand and receptor) with zero ionic strength throughout.

Figure 2.

A mixed discrete-continuum model for estimating the electrostatic component of a solute's solvation free energy; ϵ_I and ϵ_{II} represent the dielectric constants of the solute and solvent regions, Γ is the boundary between the dielectric regions and is typically a molecular (solvent-excluded) surface, and q_1 and q_2 are representative discrete point charges in the solute.

Figure 3. Schematic of a model ligand–receptor complex for studying implementation details for the co-optimization method. The surface of the ligand–receptor complex is defined by rolling a probe sphere of radius $1.4 \text{ \AA}$ over the union of the two spheres. All distances are in Angstroms.

Figure 4.

Comparison of the explicit and approximate Hessians of the sample problem in Figure Figure 3 when discretized using planar boundary elements. The alignment between the singular vectors of exact and approximate Hessians is obtained by projecting the right singular vectors of an approximate Hessian $\hat{H}$ onto the right singular vectors of the explicit Hessian H and taking the magnitude of the resulting entries. Each row and column therefore has 2-norm of one. (a) Explicit and approximate Hessians obtained using the Green's theorem formulation. (b) Explicit and approximate Hessians obtained using the polarizable continuum model/apparent surface charge formulation.

Figure 5.

Comparison of the explicit and approximate Hessians of the sample problem in Figure Figure 3 when discretized using curved boundary elements. The alignment between the singular vectors of exact and approximate Hessians is obtained by projecting the right singular vectors of an approximate Hessian $\hat{H}$ onto the right singular vectors of the explicit Hessian H and taking the magnitude of the resulting entries. Each row and column therefore has 2-norm of one. (a) Explicit and approximate Hessians obtained using the Green's theorem formulation. (b) Explicit and approximate Hessians obtained using the polarizable continuum model/apparent surface charge formulation.

Figure 6.

The magnitudes of the singular values of the explicit and approximate Hessians computed using (a) planar boundary elements and (b) curved boundary elements. The singular values of the approximate Hessians calculated using the Green's-theorem formulation are less accurate than the singular values of the approximate Hessians computed using the apparent-surface-charge (ASC) formulation, regardless of whether planar or curved boundary elements are employed.

Figure 7. Preconditioning effects on GMRES convergence for a 50-charge unconstrained optimization problem using (a) the non-derivative Green's-theorem formulation of Yoon and Lenhoff⁴⁷ and (b) the Shaw apparent-surface-charge formulation.⁴²

Figure 8.

The cost to solve unconstrained optimization problems of varying dimension using the co-optimization method, the two alternative implicit-Hessian approaches presented in Section 3.1, and by calculating the Hessian explicitly.

Figure 9.

The Anderson *et al.* inhibitor of CDK2.31 The atoms are labeled with (red) partial atomic charge values derived from quantum mechanics via RESP98-99 fitting and (blue) optimized partial atomic charges computed with unconstrained minimization using curved-element BEM and co-optimization.

Figure 10.

The unconstrained optimal partial atomic charges computed using finite-difference and boundary-element explicit Hessians, as well as using the reverse-Schur co-optimization method. The boundary-element simulations employed curved boundary element discretizations.

Figure 11.

The Anderson *et al.* inhibitor of CDK2.31 Partial atomic charges have been optimized using box inequality constraints to enforce that charge values are less than $0.85e$ in magnitude, and sum-of-charge constraints have been imposed such that the total inhibitor charge is -1 (red label), 0 (blue label), or $+1$ (green label).

Figure 12.

The HIV-1 protease inhibitor darunavir28 with atom indices for reference to Table Table 1. Hydrogen atom indices are indicated in parentheses adjacent to the atoms to which they are bonded.

Figure 13.

Unconstrained optimal partial atomic charges for the HIV-1 protease inhibitor darunavir computed using finite-difference and boundary-element explicit Hessians, as well as using the reverse-Schur co-optimization method. The boundary-element simulations employed curved boundary elements.

Table 1

Partial atomic charges for darunavir. All charge values are given as multiples of the electron charge magnitude e and have been rounded to the nearest 0.01 e . Box-constrained optimization enforced that each charge had maximum magnitude 0.85 e .

Atom Index	RESP-fit	Unconstrained Optimal	Equality-Constrained $\sum q_i = 0$	Box-Constrained	
				$\sum q_i = -1$	$\sum q_i = +1$
1	-0.87	-0.52	-0.45	-0.33	-0.44
2	0.40	0.16	0.09	-0.05	0.16
3	-0.29	0.14	0.22	0.35	0.10
4	-0.18	-0.38	-0.46	-0.61	-0.37
5	0.02	0.37	0.43	0.62	0.47
6	-0.18	-0.09	-0.12	-0.24	-0.20
7	-0.29	-0.14	-0.12	-0.03	-0.11
8	0.86	0.09	0.07	-0.33	-0.35
9	-0.52	0.00	0.00	0.00	0.01
10	-0.50	0.07	0.06	0.08	0.10
11	-0.36	-2.35	-2.25	-0.85	-0.85
12	-0.06	0.45	0.19	-0.75	-0.48
13	0.15	0.55	0.80	0.85	0.81
14	-0.14	-0.19	-0.22	-0.02	-0.11
15	-0.31	-0.22	-0.31	-0.16	-0.20
16	-0.08	1.62	1.66	0.85	0.85
17	0.09	-0.17	-0.22	-0.21	-0.19
18	-0.67	0.39	0.37	0.31	0.33
19	0.05	-1.08	-0.87	0.15	0.10
20	-0.38	-0.41	-0.46	-0.85	-0.79
21	0.68	0.76	0.76	0.82	0.85
22	-0.56	-0.38	-0.39	-0.42	-0.39
23	-0.39	-0.13	-0.13	-0.03	-0.05
24	0.12	-0.71	-0.67	-0.82	-0.85

Atom Index	RESP-fit	Unconstrained Optimal	Equality-Constrained $\Sigma q_i = 0$	Box-Constrained		
				$\Sigma q_i = -1$	$\Sigma q_i = 0$	$\Sigma q_i = +1$
25	0.05	0.88	0.86	0.79	0.83	0.85
26	-0.49	-0.61	-0.58	-0.52	-0.56	-0.59
27	0.44	0.21	0.17	-0.03	0.06	0.16
28	-0.51	-0.34	-0.32	-0.26	-0.30	-0.33
29	0.07	-0.35	-0.48	-0.54	-0.37	-0.19
30	-0.12	1.70	1.89	0.85	0.85	0.85
31	-0.01	-0.14	-0.18	0.43	0.35	0.25
32	-0.08	1.36	1.26	0.84	0.85	0.85
33	-0.16	0.24	0.26	0.23	0.17	0.12
34	-0.19	-0.09	-0.13	-0.15	-0.07	-0.01
35	-0.18	0.11	0.16	0.20	0.10	0.01
36	-0.19	-0.11	-0.10	-0.04	-0.04	-0.05
37	-0.05	0.12	0.10	-0.03	-0.02	-0.00
38	-0.06	-0.63	-0.61	-0.43	-0.39	-0.35
39	0.39	0.22	0.19	0.16	0.19	0.22
40	0.39	0.10	0.03	-0.08	0.03	0.14
41	0.18	0.14	0.13	0.11	0.13	0.14
42	0.19	-0.13	-0.22	-0.37	-0.22	-0.07
43	0.21	0.03	0.03	0.04	0.04	0.04
44	0.21	0.20	0.20	0.21	0.21	0.21
45	0.07	0.14	0.17	0.25	0.26	0.25
46	0.11	0.06	0.09	0.17	0.18	0.18
47	0.02	0.07	0.03	0.05	0.05	0.11
48	0.03	0.04	0.03	-0.01	0.02	0.04
49	0.00	-0.07	-0.06	-0.09	-0.07	-0.07
50	0.06	0.04	0.03	-0.15	-0.11	-0.10
51	0.09	0.03	0.01	-0.04	0.02	0.06

Atom Index	RESP-fit	Unconstrained Optimal	Equality-Constrained $\Sigma q_i = 0$	Box-Constrained		
				$\Sigma q_i = -1$	$\Sigma q_i = 0$	$\Sigma q_i = +1$
52	0.06	-0.12	-0.19	-0.33	-0.18	-0.06
53	0.08	0.05	0.06	0.03	0.03	0.01
54	0.14	-0.04	-0.06	0.04	0.04	0.05
55	0.11	-0.28	-0.28	-0.15	-0.15	-0.16
56	0.14	0.34	0.34	0.28	0.29	0.30
57	0.41	0.25	0.25	0.28	0.29	0.29
58	0.18	0.13	0.08	-0.22	-0.22	-0.18
59	0.07	-0.21	-0.20	-0.16	-0.16	-0.15
60	0.07	-0.12	-0.11	-0.09	-0.08	-0.06
61	0.09	-0.06	-0.05	-0.03	-0.03	-0.04
62	0.16	0.01	0.00	0.01	0.02	0.03
63	0.17	-0.00	-0.01	-0.04	-0.02	-0.00
64	0.17	-0.08	-0.08	-0.09	-0.09	-0.08
65	0.17	-0.10	-0.22	-0.40	-0.20	-0.01
66	0.10	0.14	0.12	0.14	0.15	0.16
67	0.06	-0.05	-0.06	-0.05	-0.04	-0.03
68	0.11	-0.14	-0.14	-0.14	-0.14	-0.13
69	0.04	0.03	0.04	-0.04	-0.03	-0.02
70	0.06	0.22	0.18	0.12	0.19	0.25
71	0.07	-0.43	-0.43	-0.09	-0.17	-0.24
72	0.05	-0.44	-0.70	-0.69	-0.37	-0.05
73	0.11	0.08	0.11	0.19	0.13	0.07
74	0.09	0.13	0.16	0.22	0.18	0.13
75	0.23	0.24	0.24	0.27	0.28	0.29