Available online at www.sciencedirect.com
ScienceDirect
Journal of Computational Physics 227 (2008) 7917–7928
JOURNAL OF COMPUTATIONAL PHYSICS
www.elsevier.com/locate/jcp

An iterative boundary potential method for the infinite domain Poisson problem with interior Dirichlet boundaries ☆

G.H. Miller *

University of California, Department of Applied Science, 1 Shields Avenue, Davis, CA 95616, United States
Lawrence Berkeley National Laboratory, Applied Numerical Algorithms Group, 1 Cyclotron Road, Berkeley, CA 94720, United States

Received 16 September 2007; received in revised form 15 February 2008; accepted 1 May 2008
Available online 16 May 2008

Abstract

An iterative method is developed for the solution of Poisson's problem on an infinite domain in the presence of interior boundaries held at fixed potential, in three dimensions. The method combines pre-existing fast multigrid-based Poisson solvers for data represented on Cartesian grids with the fast multipole method. Interior boundaries are represented with the embedded boundary formalism. The implementation is in parallel and uses adaptive mesh refinement. Examples are presented for a smooth interior boundary for which an analytical result is known, and for an irregular interior boundary problem. Second-order accuracy in L_1 with respect to the grid resolution is demonstrated for both problems.
Published by Elsevier Inc.

Keywords: Poisson problem; Fast multipole method; Multigrid; Embedded boundary method

1. Introduction

The Poisson problem is central to a wide variety of applications in computational physics, from electrostatics to projection methods for incompressible flow. For gridded data, or grid-mediated point data (*e.g.*, the particle-in-cell method) the easily implemented boundary conditions are Dirichlet, Neumann, or periodic. However, for many problems the most appropriate choice, on physical grounds, is the infinite domain condition. Solutions to the infinite domain problem have been estimated using the easily implemented boundary conditions in conjunction with very large computational domains, or with stretched grids, employed to remove the boundary from the region of interest. Of course such approaches are only approximate, and can be very demanding of resources especially in 3D. More rigorous boundary potential methods have been developed that determine the inhomogeneous Dirichlet conditions on a finite domain that are consistent with the desired

☆ This work was partially supported by a contract from Capschell Inc., Chicago, IL, USA, by the US DOE MICS Division under Contract Number DE-FG02-03ER25579, and under Subcontract B565606 with the Lawrence Livermore National Laboratory under the auspices of the Department of Energy contract No. W-7405-Eng-48.
* Tel.: +1 530 754 7059; fax: +1 530 752 2444.
E-mail address: grgmiller@ucdavis.edu

0021-9991/$ - see front matter Published by Elsevier Inc.
doi:10.1016/j.jcp.2008.05.005

infinite domain properties [22,14,12,23,1,16]. These methods exploit the free space Green's function to construct a boundary potential from a set of screening charges.

This work is concerned with an extension of boundary potential methods to infinite domain Poisson problems that contain also surfaces with fixed potential. One possible solution to this combined problem is the superposition of the solution of an external Dirichlet Laplace problem (*e.g.*, [19]) with the solution to an infinite domain Poisson problem constructed without interior boundaries. Such external Dirichlet Laplace problems involve quadrature of a codimension 1 Fredholm equation with singular kernel, integrated over the interior boundary. This results in a dense matrix equation for the charge density on the interior surface [15,19], and is similar to the 2D capacitance matrix method of Hockney and Eastwood [7]. Instead of pursuing such non-iterative approach, an iterative method based on existing fast solvers is developed.

Consider a three dimensional rectangular domain Ω_{dom} which contains space charges prescribed through the charge density ρ, and one or more closed regions Ω_{int} with prescribed surface potentials ϕ_{int}. The objective is the solution Φ to the Poisson problem

$$\Delta\Phi = \rho \tag{1a}$$

$$\Phi = \phi_{\text{int}} \quad \text{on } \partial\Omega_{\text{int}} \tag{1b}$$

$$\Phi(\mathbf{x}) \sim -\frac{Q}{4\pi|\mathbf{x}|} \quad \text{as } \mathbf{x} \to \infty, \tag{1c}$$

where Q is the sum of all charges in Ω_{dom}, consisting of space charges ρ, and also surface charges on $\partial\Omega_{\text{int}}$. After [22,14,12] we decompose Φ as sum of two fields, $\Phi = \phi + \Psi$, where ϕ is given by

$$\Delta\phi = \rho \tag{2a}$$

$$\phi = \phi_{\text{int}} \quad \text{on } \partial\Omega_{\text{int}} \tag{2b}$$

$$\phi = 0 \quad \text{on } \partial\Omega_{\text{dom}}, \tag{2c}$$

and $\phi = 0$ everywhere outside Ω_{dom}. It is possible to express ϕ as a free space Green's function convolution over the space charge density ρ and surface charges densities ϱ:

$$\phi(\mathbf{x}) = \int_{\Omega_{\text{dom}}} dV' G(\mathbf{x}|\mathbf{x}')\rho(\mathbf{x}') + \int\int_{\partial\Omega_{\text{int}}} dS' G(\mathbf{x}|\mathbf{x}')\varrho_{\text{int}}(\mathbf{x}') + \int\int_{\partial\Omega_{\text{dom}}} dS' G(\mathbf{x}|\mathbf{x}')\varrho_{\text{dom}}(\mathbf{x}'). \tag{3}$$

Here the surface charge densities ϱ_{int} and ϱ_{dom} are implicit functions given by (3) with the boundary conditions (2b) and (2c). Alternatively, Green's second theorem may be written

$$\begin{aligned}\phi(\mathbf{x}) = {} & \int_{\Omega_{\text{dom}}} dV' G(\mathbf{x}|\mathbf{x}')\Delta'\phi(\mathbf{x}') - \int\int_{\partial\Omega_{\text{int}}} d\mathbf{S}' \cdot \nabla'\phi(\mathbf{x}')G(\mathbf{x}|\mathbf{x}') + \int\int_{\partial\Omega_{\text{int}}} d\mathbf{S}' \cdot \nabla' G(\mathbf{x}|\mathbf{x}')\phi(\mathbf{x}') \\ & - \int\int_{\partial\Omega_{\text{dom}}} d\mathbf{S}' \cdot \nabla'\phi(\mathbf{x}')G(\mathbf{x}|\mathbf{x}') + \int\int_{\partial\Omega_{\text{dom}}} d\mathbf{S}' \cdot \nabla' G(\mathbf{x}|\mathbf{x}')\phi(\mathbf{x}')\end{aligned} \tag{4}$$

from which, using (2c) and comparing with (3), one may deduce

$$\varrho_{\text{dom}}(\mathbf{x}) = -\mathbf{n} \cdot \nabla\phi\mathbf{x} \quad \text{on } \partial\Omega_{\text{dom}}. \tag{5}$$

The correction field Ψ must solve

$$\Delta\Psi = -\int\int_{\partial\Omega_{\text{dom}}} dS'\delta(\mathbf{x} - \mathbf{x}')\varrho_{\text{dom}}(\mathbf{x}') \tag{6a}$$

$$\Psi = 0 \quad \text{on } \partial\Omega_{\text{int}} \tag{6b}$$

$$\Psi \sim -\frac{Q'}{4\pi|\mathbf{x}|} \quad \text{as } \mathbf{x} \to \infty, \tag{6c}$$

where Q' is the sum of boundary charges on $\partial\Omega_{\text{int}}$ and $\partial\Omega_{\text{dom}}$. The right hand side of (6a) is $-\varrho_{\text{dom}}$ expressed as a space charge density, required by the condition that $\Phi = \phi + \Psi$ have no charge density on the artificial boundary $\partial\Omega_{\text{dom}}$. As a convolution over the free space Green's function, Ψ may be written

$$\Psi(\mathbf{x}) = \int\int_{\partial\Omega_{\text{int}}} \mathrm{d}S' G(\mathbf{x}|\mathbf{x}')\tilde{\varrho}_{\text{int}}(\mathbf{x}') - \int\int_{\partial\Omega_{\text{dom}}} \mathrm{d}S' G(\mathbf{x}|\mathbf{x}')\varrho_{\text{dom}}(\mathbf{x}'), \tag{7}$$

where $\tilde{\varrho}_{\text{int}}$ is an implicit function determined by (7) with the boundary condition (6b). Comparing this result with Green's second theorem one determines

$$\tilde{\varrho}_{\text{int}} = -\mathbf{n}\cdot\nabla\Psi\mathbf{x} \quad \text{on } \partial\Omega_{\text{int}}. \tag{8}$$

Combining (5), (7), and (8) yields the integral equation

$$\Psi(\mathbf{x}) = -\int\int_{\partial\Omega_{\text{int}}} \mathrm{d}\mathbf{S}'\cdot\nabla\Psi(\mathbf{x}')G(\mathbf{x}|\mathbf{x}') + \int\int_{\partial\Omega_{\text{dom}}} \mathrm{d}\mathbf{S}'\cdot\nabla\phi(\mathbf{x}')G(\mathbf{x}|\mathbf{x}'). \tag{9}$$

Our approach will be to solve (2a) by finite volume methods, and to solve (9) through simple iteration, namely

$$\Psi^{(k+1)}(\mathbf{x}) = -\int\int_{\partial\Omega_{\text{int}}} \mathrm{d}\mathbf{S}'\cdot\nabla\Psi^{(k)}(\mathbf{x}')G(\mathbf{x}|\mathbf{x}') + \int\int_{\partial\Omega_{\text{dom}}} \mathrm{d}\mathbf{S}'\cdot\nabla\phi(\mathbf{x}')G(\mathbf{x}|\mathbf{x}') \tag{10}$$

subject to the condition $\Psi^{(k)} = 0$ on $\partial\Omega_{\text{int}}$. The algorithm is stated below in Section 2, and demonstrated through an analytical example in Section 3. In Section 4 the established numerical methods used to implement the algorithm are described briefly. Second-order accuracy in L_1 for C^∞ and C^0 interior boundaries is demonstrated in Section 5, and the convergence properties of the algorithm are discussed.

2. Algorithm

The iterative algorithm proceeds as follows:

I. First, solve the homogeneous Dirichlet domain boundary problem

$$\Delta\phi = \rho \tag{11a}$$

$$\phi = \phi_{\text{int}} \quad \text{on } \partial\Omega_{\text{int}} \tag{11b}$$

$$\phi = 0 \quad \text{on } \partial\Omega_{\text{dom}}. \tag{11c}$$

II. Second, compute the charge density on the domain boundary and compute an estimated boundary potential using the free space Green's function. The superscript in parenthesis will denote an iteration index.

$$\varrho_{\text{dom}} = -\mathbf{n}\cdot\nabla\phi \quad \text{on } \partial\Omega_{\text{dom}}. \tag{12a}$$

$$\phi^{(0)}_{\text{dom}}(\mathbf{x}) = -\int\int_{\partial\Omega_{\text{dom}}} \mathrm{d}S' G(\mathbf{x}|\mathbf{x}')\varrho_{\text{dom}}(\mathbf{x}'). \tag{12b}$$

Now proceed with iteration index $k = 1$.

III. Compute the induced charge on the interior boundary. This begins by solving the Laplace problem

$$\Delta\Psi^{(k)} = 0 \tag{13a}$$

$$\Psi^{(k)} = 0 \quad \text{on } \partial\Omega_{\text{int}} \tag{13b}$$

$$\Psi^{(k)} = \phi^{(k-1)}_{\text{dom}} \quad \text{on } \partial\Omega_{\text{dom}}, \tag{13c}$$

where $\Psi^{(k)}$ is taken to be zero everywhere within Ω_{int}. The induced surface charges are then

$$\tilde{\varrho}^{(k)}_{\text{int}} = -\mathbf{n}\cdot\nabla\Psi^{(k)} \text{on } \partial\Omega_{\text{int}}. \tag{14}$$

IV. These induced surface charges modify the boundary potential,

$$\phi^{(k)}_{\text{dom}}(\mathbf{x}) = \phi^{(0)}_{\text{dom}}(\mathbf{x}) + \int\int_{\partial\Omega_{\text{int}}} \mathrm{d}S' G(\mathbf{x}|\mathbf{x}')\tilde{\varrho}^{(k)}_{\text{int}}(\mathbf{x}'). \tag{15}$$

V. If the boundary potential has converged, the iteration ceases and $\Phi = \phi + \Psi^{(k)}$. Convergence is assessed by measuring the L_2 norm of $|\phi^{(k)}_{\text{dom}} - \phi^{(k-1)}_{\text{dom}}|$. Otherwise, $k := k + 1$ and recompute the induced charge (Step III).

The behavior and convergence properties of the algorithm can be shown with a small analytical example. In Section 3 it will be shown that the algorithm outlined above will fail to converge when the interior boundary and domain boundary are close (in a sense that will be made clear). One remedy for this convergence problem is to replace (15) with

$$\phi_{\mathrm{dom}}^{(k)}(\mathbf{x}) = \omega\phi_{\mathrm{dom}}^{(0)}(\mathbf{x}) + \omega\int\int_{\partial\Omega_{int}} \mathrm{d}S' G(\mathbf{x}|\mathbf{x}')\tilde{\varrho}_{\mathrm{int}}^{(k)}(\mathbf{x}') + (1-\omega)\phi_{\mathrm{dom}}^{(k-1)}(\mathbf{x}) \tag{16}$$

for $0 < \omega < 1$. (16) has the same stationary point as (15) and enables convergence for essentially any well-resolved geometry where Ω_{int} is contained within Ω_{dom}.

3. Analytical example

For an analytical exposition in 3D it is convenient to use spherical geometry. Let Ω_{dom} be a sphere of radius $a > 1$, and let Ω_{int} be a concentric sphere of radius 1, and omit space charges ρ. The potential can at all times be expanded in spherical harmonics. With three simple observations, the analytical example is easily carried out: (i) the spherical harmonic modes are fully decoupled from one another; (ii) the ℓm spherical harmonic component of the potential varies with r like $\alpha r^l + \beta r^{-(\ell+1)}$; and (iii) if the ℓm component of charge density on a spherical shell of radius r_1 is ϱ_ℓ^m, then the potential due to this charge density has amplitude $\phi_\ell^m = -\frac{r_1^{\ell+2}}{r^{\ell+1}}\frac{\varrho_\ell^m}{2\ell+1}$ evaluated on a spherical shell with radius $r_2 \geqslant r_1$.

Let ϕ_{int}^2 have amplitude f_ℓ^m in harmonic mode ℓm. The solution to algorithm step I is

$$\phi_\ell^m = f_\ell^m \frac{\left(\frac{r}{a}\right)^\ell - \left(\frac{a}{r}\right)^{\ell+1}}{s\left(\frac{1}{a}\right)^\ell - \left(\frac{a}{1}\right)^{\ell+1}}. \tag{17}$$

Associated with this solution is the domain surface charge density ϱ_{dom} and associated domain surface compensating potential, step II:

$$\varrho_{\mathrm{dom}_\ell}^m = -\frac{f_\ell^m}{a}\frac{(2\ell+1)}{\left(\frac{1}{a}\right)^\ell - \left(\frac{a}{1}\right)^{\ell+1}} \tag{18a}$$

$$\phi_{\mathrm{dom}_\ell}^{(0)\ m} = -f_\ell^m \frac{1}{\left(\frac{1}{a}\right)^\ell - \left(\frac{a}{1}\right)^{\ell+1}}. \tag{18b}$$

The general solution to the Laplace problem of step III is

$$\Psi_\ell^{(k)m} = \phi_{\mathrm{dom}\ell}^{(k-1)m} \frac{\left(\frac{r}{1}\right)^\ell - \left(\frac{1}{r}\right)^{\ell+1}}{\left(\frac{a}{1}\right)^\ell - \left(\frac{1}{a}\right)^{\ell+1}}, \tag{19a}$$

and the corresponding induced surface charge on the interior boundary is

$$\tilde{\varrho}_{\mathrm{int}\ell}^{(k)m} = \phi_{\mathrm{dom}\ell}^{(k-1)m} \frac{(2\ell+1)}{\left(\frac{a}{1}\right)^\ell - \left(\frac{1}{a}\right)^{\ell+1}} \tag{19b}$$

giving a new boundary potential, step IV:

$$\phi_{\mathrm{dom}\ell}^{(k)\ m} = \phi_{\mathrm{dom}\ell}^{(0)\ m} - \phi_{\mathrm{dom}\ell}^{(k-1)m}\frac{1}{a^{2\ell+1}-1}. \tag{20}$$

writing $\lambda_\ell = 1/(1-a^{2\ell+1})$ we have

$$\phi_{\mathrm{dom}\ell}^{(k)\ m} = \phi_{\mathrm{dom}\ell}^{(0)\ m}\frac{1-\lambda_\ell^{k+1}}{1-\lambda_\ell}, \tag{21}$$

which converges to

$$\phi_{\mathrm{dom}\ell}^{*\ m} = \phi_{\mathrm{dom}\ell}^{(0)\ m}\frac{a^{2\ell+1}-1}{a^{2\ell+1}} = f_\ell^m\frac{a^\ell}{a^{2\ell+1}-1} \tag{22}$$

provided $|\lambda_\ell| < 1$ or $a^{2\ell+1} > 2$ for all ℓ. The 00 mode is the slowest to converge, requiring $a > 2$. The same conclusion regarding the 00 mode may be drawn through analysis of [11, problem 1.6]. The converged correction field Ψ_ℓ^m (19a), added to the initial calculation ϕ_ℓ^m (17), gives the exact analytic result Φ_ℓ^m for the model problem:

$$\Phi_\ell^m = f_\ell^m r^{-(\ell+1)}. \tag{23}$$

With the relaxed convergence scheme of (16), algorithm step IV would give

$$\phi_{\mathrm{dom}\ell}^{(k)\ m} = \omega\phi_{\mathrm{dom}\ell}^{(0)\ m} - \omega\phi_{\mathrm{dom}\ell}^{(k-1)m}\frac{1}{a^{2\ell+1}-1} + (1-\omega)\phi_{\mathrm{dom}\ell}^{(k-1)m} \tag{24}$$

in place of (20). Writing $\lambda'_\ell = 1 + \omega(\lambda_\ell - 1)$ we have

$$\phi_\ell^{(k)m} = \phi_\ell^{(0)m}\left[\omega\frac{1-\lambda'^{k+1}}{1-\lambda'} + \lambda'^{k+1}\right]. \tag{25}$$

Convergence requires $|\lambda'_\ell| < 1$ for all ℓ. This condition implies

$$0 < \omega < 2(1-a^{-1}), \tag{26}$$

and the converged result is again given by (22). Note that a viable solution to (26) is possible for any $a > 1$, whereas $a > 2$ was required in the absence of relaxation.

If a space charge density were included in this example, the potential fields ϕ and Φ would differ from the formulae given above by addition of the convolution $G * \rho$, where the particular Green's function used here maintains zero potential on the interior sphere. A particularly efficient solution is given by the method of images [11, Section 2.4].

4. Implementation

The numerical method uses a Cartesian grid finite volume approach with cell-centered data. Complex interior geometries are represented by overlaying the interior domain on the Cartesian grid. Cut cells are characterized by volume fraction, cell surface area fractions, and an interface normal vector. The implementation is built using the `Chombo` library of block-structured data management and operation tools. The treatment of complex geometry is based on ideas in [13], and has been developed for hyperbolic, parabolic, and elliptic problems in [4,17,20].

4.1. Poisson solvers

Algorithm steps I and III solve Dirichlet boundary condition Poisson problems using the second-order accurate approach of [20]. The approach uses a divergence form of the Laplace operator, with cell-edge fluxes at irregular domains being computed through specialized high-order stencils. Geometric multigrid is employed, with graph-based coarsening operators making possible coarse-grid relaxation at resolutions below which the interior geometry may be resolved. Additional detail on the stencil operations may be found in [18].

4.2. Boundary charge calculation

Domain boundary charges (algorithm Step II) are estimated using a one-sided finite difference approximation to $\mathbf{n}\cdot\nabla\phi$. When adaptive mesh refinement is used estimates of surface charge may be available at several grid resolutions. In that case the values computed at the highest resolution are used, and underlying coarse cell approximations are discarded.

Charges computed on the interior boundary (algorithm Step III) are estimated in each cell containing an interface. The derivative $\mathbf{n}\cdot\nabla\phi$ is obtained by solving a least squares problem for the first derivatives, second derivatives, and the second-order cross derivative terms of ϕ. The centering point of the Taylor series is the centroid of the intersection of the interior boundary with the computational cell. This least squares problem

uses data from all "regular" cells (cells that do not contain an interface and are not contained inside the interior boundary), and that are in a line of sight with the cell in which the gradient is desired. This last condition eliminates from the least squares fit those data that are "close", but which may be separated from the centering point by thin boundary walls. The least squares problem is solved by Householder QR reduction, and the result is stored as a stencil – a set of pairs of weights and cell coordinates. The relatively expensive stencil construction occurs once in the method, and each iterative evaluation of algorithm Step III involves a computationally inexpensive evaluation of the stencil. The interior boundary surface charge in a cell is proportional to the product of the derivative $\mathbf{n} \cdot \nabla\phi$ with the area of the interior boundary. As with domain boundaries, with adaptive mesh refinement the charge may be computed at multiple resolutions. The values used are obtained at the finest available level of resolution.

4.3. Free space Green's functions

Algorithm steps II and IV involve computing domain boundary potentials from surface charge densities. The free space Green's function convolution is solved using an implementation of the fast multipole method (FMM) [6,5,3], which follows closely the non-adaptive scheme presented in [5].[1] The rotation of multipole and local expansions make use of rotation matrices constructed by the very efficient recursive method of Ivanic and Ruedenberg [9,10].[2]

The FMM calculates local expansions which represent the potential in a neighborhood from charges well-separated from that neighborhood. Charges that are not well-separated are evaluated by direct calculation, $-1/4\pi|\mathbf{r}|$. When calculating the domain boundary potential due to domain boundary charges, it is necessary to account for the case where the charge and potential are co-located. This is accomplished using an integral formulation for the potential at a point due to a uniform charge density over a 2D square source (*e.g.*, [8, Eqs. (4)–(6)]. All domain potentials computed directly from domain charges make use of this strategy. If ever an interior boundary charge need to be transferred directly to an domain boundary potential, a point-to-point $-1/4\pi|\mathbf{r}|$ calculation is used, the points being respective centroids.

The FMM uses a hierarchy of Cartesian data layouts on a cubic domain. The finest level of FMM refinement corresponds with a coarsened version of the Poisson domain, suitably extended to be cubic and have length be a power of 2. The coarsening is not essential, but improves performance by balancing the cost of multipole $\rightarrow$ local operations against the number of direct $-1/4\pi|\mathbf{r}|$ operations required.

When adaptive mesh refinement is used, the boundary potential is calculated on all available levels of refinement that abut $\partial\Omega_{\text{dom}}$.

5. Results and conclusions

The first example verifies the method by comparing calculations against analytical solutions. The interior boundary is a sphere of radius 0.1, centered in a cubic domain of length 1, resembling the cartoon of Fig. 1. A collection of 200 point charges surround the sphere. Their coordinates are chosen randomly to lie inside the torus with major radius 0.2 and minor radius 0.05, bisected by the horizontal midplane of the sphere. The sphere has fixed potential with 00 and 22 spherical harmonic modes (using the real part with Greengard's choice of normalization), with magnitude 1 and 10, respectively. The spherical coordinates are rotated relative to the cube in order to break the symmetry of the problem, via $R_z(45^\circ) \circ R_y(15^\circ) \circ R_z(30^\circ)$. The analytic reference solution for the potential due to the sphere is treated in a straightforward manner. The Cartesian grid representation of the analytic contribution of the point charges is more delicate. In each computational cell, an analytical volume average of the $-1/4\pi|\mathbf{r}|$ potential is calculated for the given charges and for the image charges used to respect the interior boundary condition. For the numerical solution, the point charges are distributed on the grid using a cloud-in-cell distribution [2]. Second-order convergence in L_1 is found (Table 1). In

[1] The imaginary exponent in [5, Eq. (7.17)] and [3, Eq. (49)] should be negative. [6, Eq. (3.60)] is correctly rewritten as [5, Eq. (5.6)] and [3, Eq. (17)].

[2] Table 2 is correct in [9] not [10].

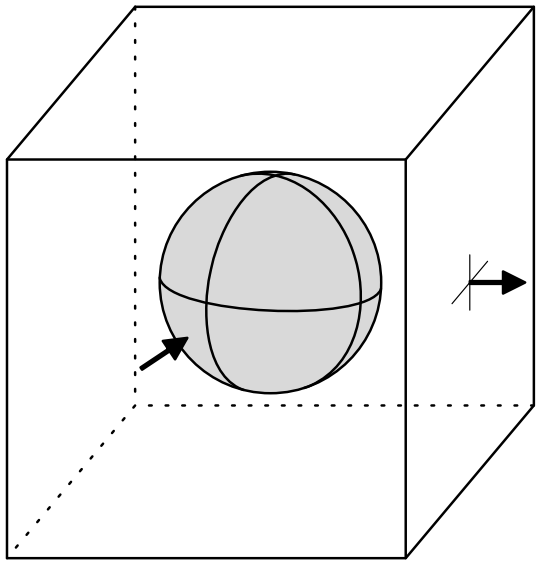

Fig. 1. Schematic view of proposed problem. The rectangular computational domain (Ω_{dom}) contains an interior object (Ω_{int}) whose surface potential ϕ_{int} is fixed. The volume $\Omega_{\text{dom}} \setminus \Omega_{\text{int}}$ may also contain a distribution of space charges ρ. Arrows signify boundary normal vectors **n**.

Table 1
Convergence study for sphere with 200 point charges

$1/h$	L_∞	Rate	L_1	Rate	L_2	Rate
16	0.151×10^2		0.175×10^1		0.209×10^1	
32	0.341×10^1	2.14	0.145×10^0	3.60	0.158×10^0	3.73
64	0.171×10^1	1.00	0.337×10^{-1}	2.10	0.381×10^{-1}	2.05
128	0.829×10^0	1.04	0.734×10^{-2}	2.20	0.100×10^{-1}	1.93

Errors relative to analytical solution.

L_∞ the convergence is $\mathcal{O}(h^1)$: an inherent problem of particle-in-cell type methods. With the point charges omitted, the corresponding Laplace problem is second-order accurate in all norms (Table 2).

The sphere-box problem has also been solved with adaptive mesh refinement, again omitting space charges. The sphere is offset relative to the center of the domain box (Fig. 2) to demonstrate that the method accommodates data with multiple resolutions on the domain boundary. In these computations the hierarchy of nested domains was held fixed. The grid spacing changes by a factor of two across levels. Approximately second-order convergence is again observed, Table 3.

The second example demonstrates the robustness of the method by showing second-order convergence for a case where the interior boundaries are not smooth. The interior surface is C^0 on a set of codimension 2, and C^∞ elsewhere. In some cells containing these C^0 regions a sufficiently large stencil of points is not always avail-

Table 2
Convergence study for sphere without point charges

$1/h$	L_∞	Rate	L_1	Rate	L_2	Rate
16	0.129×10^2		0.113×10^0		0.189×10^0	
32	0.335×10^1	1.95	0.132×10^{-1}	3.10	0.438×10^{-1}	2.11
64	0.126×10^1	1.41	0.389×10^{-2}	1.76	0.123×10^{-1}	1.83
128	0.237×10^0	2.41	0.799×10^{-3}	2.29	0.248×10^{-2}	2.31

Errors relative to analytical solution.

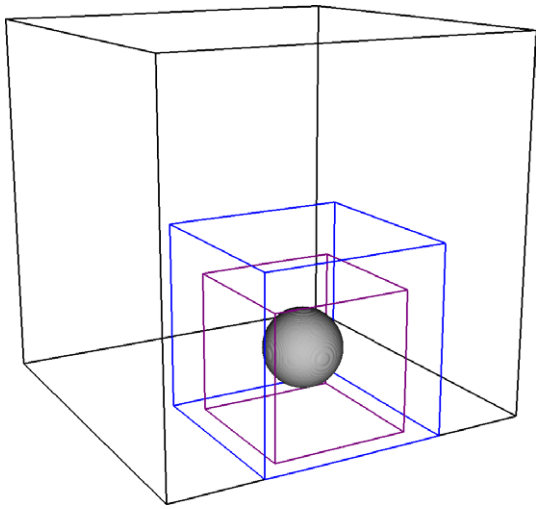

Fig. 2. AMR grid layout.

Table 3
Convergence for three-level adaptive mesh refinement, without point charges

$1/h$	L_∞	Rate	L_1	Rate	L_2	Rate
16/32/64	0.235×10^{1}		0.281×10^{-2}		0.109×10^{-1}	
32/64/128	0.169×10^{1}	0.47	0.667×10^{-3}	2.08	0.224×10^{-2}	2.29

Errors relative to analytical solution.

able for a second-order computation of the interior boundary charge. However, since the set of such points is only codimension 2, the overall solution displays second-order convergence in L_1. The interior boundaries are held at fixed potential $\phi_{\text{int}} = 0$, and a C^0 distribution of space charges ρ is prescribed (Figs. 3 and 4). Because of the proximity of the interior boundaries to the domain boundary, relaxation is used with $\omega = 0.7$. The con-

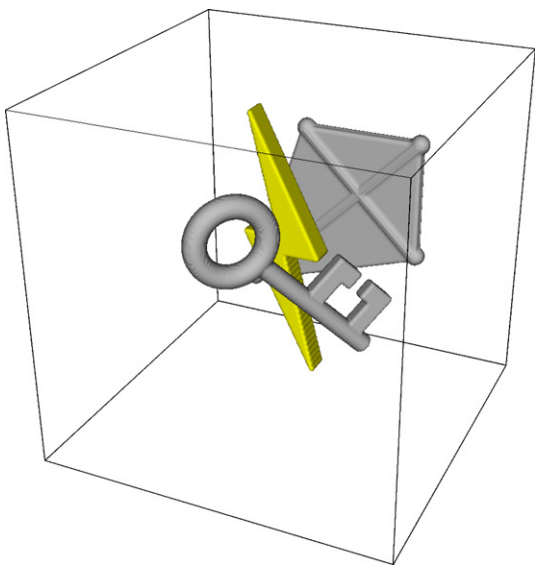

Fig. 3. Geometry of second test: Franklin's experiment with a bolt of charge.

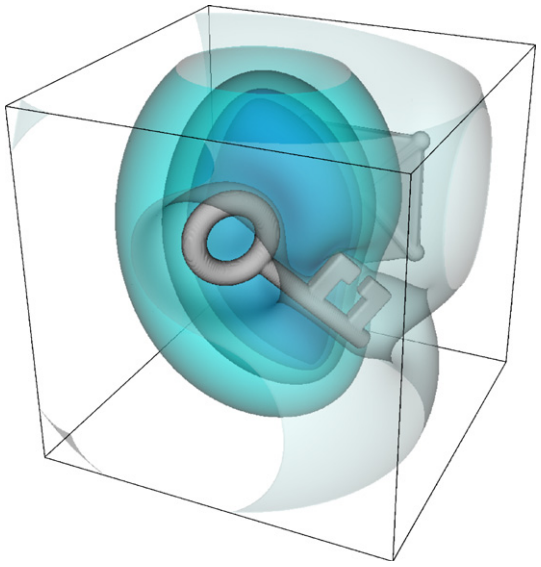

Fig. 4. Contours of Φ corresponding to geometry of figure 3.

Table 4
Convergence study for C^0 geometry

$1/h$	L_∞	Rate	L_1	Rate	L_2	Rate
16 vs. 32	0.231×10^{-1}		0.273×10^{-3}		0.365×10^{-3}	
32 vs. 64	0.745×10^{-3}	1.63	0.394×10^{-4}	2.78	0.407×10^{-4}	3.17
64 vs. 128	0.292×10^{-3}	1.35	0.171×10^{-4}	1.22	0.174×10^{-4}	1.22
128 vs. 256	0.289×10^{-3}	0.015	0.335×10^{-5}	2.36	0.342×10^{-5}	2.35

Errors estimated by Richardson extrapolation.

vergence results (Table 4) are second-order in L_1 and L_2. A rate of order h^0 is found for L_∞: this is the expected result for C^0 boundaries in L_∞ for Poisson's problem with the embedded boundary formalism [20, Section 2.2].

Fig. 5 displays the convergence properties for a sequence of test problems consisting of a sphere of diameter D centered in a cubic domain of length L, and $h = L/64$. Space charges are omitted. The sphere has the same Dirichlet boundary conditions as the first example problem, and $\omega = 1$ for all calculations. From the nested sphere analytical example we expect convergence for $D/L < \frac{1}{2}$ (approximately), and this expectation is borne out.

Fig. 6 examines the effect of varying ω when $D/L = 0.8$ – a problem that diverges with $\omega = 1$. Again, $h = L/64$. Qualitatively, relaxation values $\omega < 0.6$ have the effect of stabilizing the method, and an optimum relaxation parameter in the neighborhood of 0.4 is shown to exist.

An estimate of the optimal relaxation parameter for the general case may be deduced from a numerical measurement of the possibly divergent behavior when $\omega = 1$. We note first that the overall iterative scheme is characterized by a linear operator $\mathcal{A}$, and the convergence properties of the method depend on the eigenvalue spectrum of this operator. Let $\mathcal{L}$ denote the inverse operator of (13a), *i.e.*, $\Psi^{(k)} = \mathcal{L}\phi^{(k-1)}_{\text{dom}}$. Then $\mathcal{A}$ is given schematically as

$$\mathcal{A} = -G(\mathbf{x}_{\delta\Omega_{\text{dom}}}|\mathbf{x}_{\delta\Omega_{\text{int}}}) * \left.\frac{\partial}{\partial n}\right|_{\partial\Omega_{\text{int}}} \mathcal{L}, \tag{27}$$

and

$$\mathrm{e}^{(k)} = \mathcal{A}\mathrm{e}^{(k-1)} \tag{28a}$$

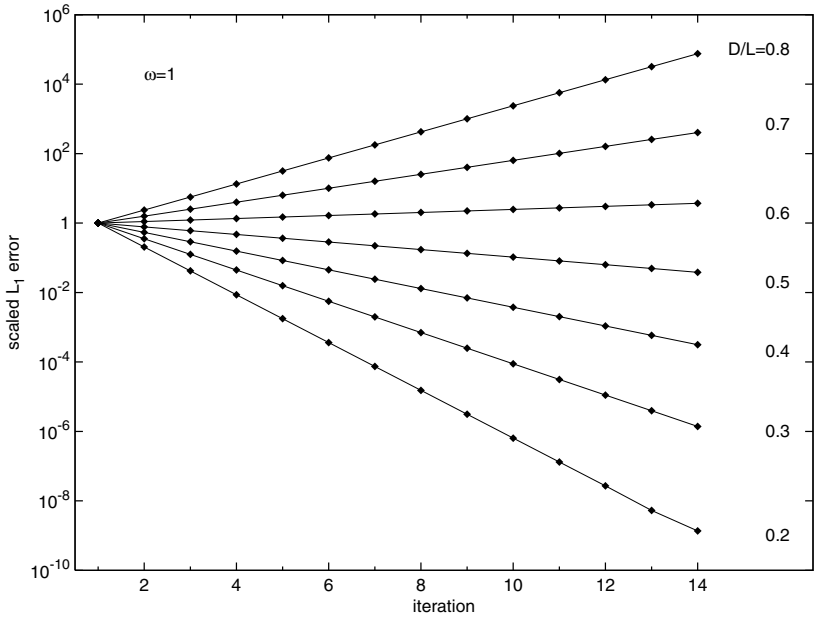

Fig. 5. Convergence behavior with $\omega = 1$, varying the ratio of sphere diameter D to cubic domain length L. The L_1 error is measured relative to the converged numerical result (convergence requires $\omega < 1$ for $D/L > 0.5$), and normalized to give unity at the first iteration.

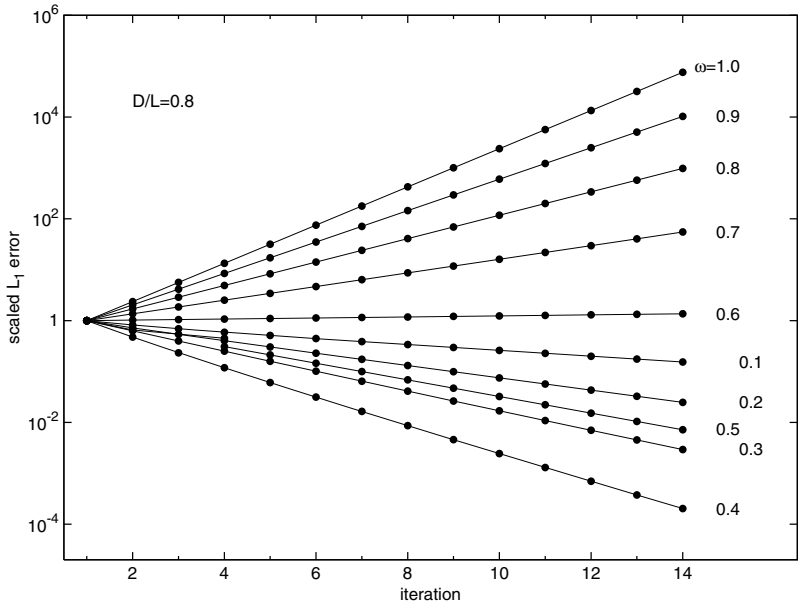

Fig. 6. Convergence behavior with sphere diameter fixed at 0.8 of the domain length. The relaxation parameter ω varies. The L_1 error is measured relative to the converged numerical result and normalized to give unity at the first iteration.

with

$$\mathrm{e}^{(k)} \equiv \phi^{(k)}_{\mathrm{dom}} - \phi^{*}_{\mathrm{dom}} \tag{28b}$$

the boundary potential error at iteration k when ϕ^{*}_{dom} is the converged result.

In consideration of the parity of operators $\mathcal{L}$ and G, operator $\mathcal{A}$ will have negative eigenvalues λ, and convergence will be governed by $\sup(|\lambda|)$. The proposed relaxation scheme is governed by an operator $\mathcal{A}_\omega$,

$$\mathcal{A}_\omega = (1-\omega)I + \omega\mathcal{A} \tag{29}$$

with eigenvalues $\lambda_\omega = 1 + \omega(\lambda - 1)$.

Let A denote the smallest eigenvalue of $\mathcal{A}$, and B denote its largest eigenvalue. An approximation to the optimal relaxation parameter ω can be had by seeking a mapping $(A,B)\mapsto(A_\omega, B_\omega)$, with $A_\omega = -B_\omega$ characterizing the optimal choice. That is,

$$\omega = \frac{2}{2 - A - B}. \tag{30}$$

Owing to the distinct linearity observed in Fig. 5, an estimate of A can be had by observing the convergence behavior of the method using $\omega = 1$ and a small number of iterations. With $D/L = 0.8$ the slope of the curve gives $A \approx -2.37$. The same slope is observed for the L_2 solution error, and we assume here that the solution error is proportional to the boundary potential error. There is no direct measure of B, but an upper bound of 0 can be inferred by analogy to the nested sphere example. Using this estimate one obtains $\omega \approx 0.42$ as an approximation to the optimal relaxation parameter. The observed optimum (Fig. 7) is approximately 0.435, and is consistent with B having been overestimated (the optimum $\omega \approx 0.435$ implies $B \approx -0.23$).

Although the spatial extent of the FMM grid hierarchy (Section 4.3) may be larger than the Poisson domain, ϕ and Ψ are not computed on the extended domain. In this sense the present method is very compact. In contrast, the methods of [12] and [16] solve fast (FFT-based) Poisson problems on extended domains as part of their boundary potential calculation. Note that the method presented here, applied to a problem with no interior domain, solves the same problem as [12,16], non-iteratively ($\Phi = \phi + \Psi^{(0)}$), on a compact domain.

The method presented here may be adapted to solve the infinite domain Poisson problem with Neumann boundary conditions on the interior domain. Three changes are required: (i) ϕ of step I would have an inhomogeneous Neumann condition on the interior surface; (ii) $\Psi^{(k)}$ of step III would have a homogeneous Neumann condition on the interior boundary; and (iii) the induced charge distribution on the interior boundary would be a double layer. Across a double layer the jump in $\partial\Psi/\partial\mathbf{n}$ is zero, thus the potential inside the interior domain is constant. By application of the divergence theorem, with there being no charges in Ω_{int}, it follows that the constant must be zero. Finally, the jump in potential across a double layer gives σ [21, L.15 Section 2]. With the interior value of Ψ being zero,

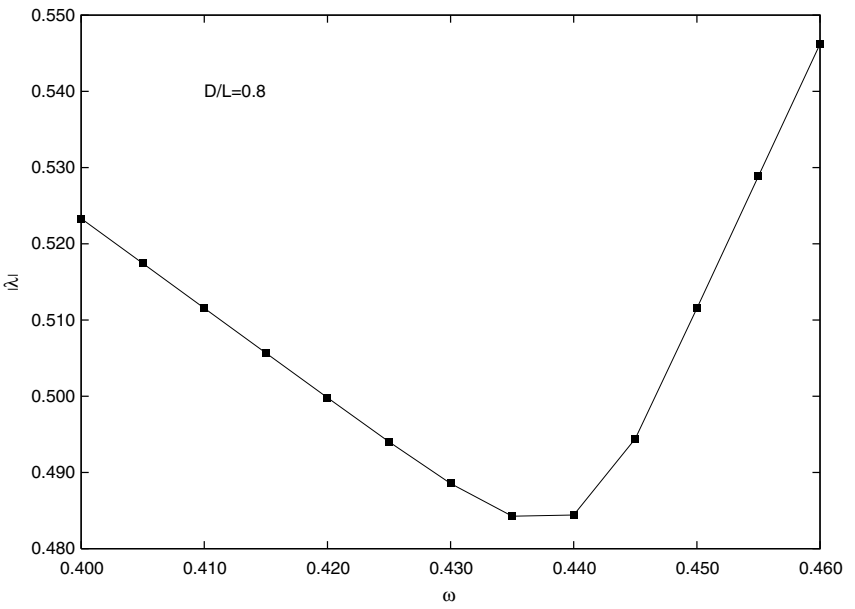

Fig. 7. Convergence behavior in the neighborhood of the predicted optimum $\omega = 0.42$.

$$\sigma(\mathbf{x}_0) = \lim_{x \to \mathbf{x}_0} \frac{-\Psi(\mathbf{x})}{4\pi}, \mathbf{x}_0 \quad \text{on } \delta\Omega_{\text{int}}, \tag{31}$$

where the limit is approached from the outside the interior domain. The Green's function used to compute the change in boundary potential $\phi^{(k)}_{\text{dom}}$ due to the double layer is the negative of the normal derivative of the free space Green's function. The FMM method is easily adapted to this case [6].

Acknowledgements

GHM thanks P. Colella, D. T. Graves, and D. Trebotich for many insightful discussions. Two anonymous reviewers provided very helpful suggestions that significantly improved the exposition of the method.

References

[1] G.T. Balls, P. Colella, J. Comput. Phys. 180 (2002) 25–53.
[2] C.K. Birdsall, D. Fuss, J. Comput. Phys. 3 (1969) 494–511.
[3] H. Cheng, L. Greengard, V. Rokhlin, J. Comput. Phys. 155 (1999) 468–498.
[4] P. Colella, D.T. Graves, B.J. Keen, D. Modiano, J. Comput. Phys. 211 (2006) 347–366.
[5] L. Greengard, V. Rokhlin, Acta Numer. 6 (1997) 229–269.
[6] L.F. Greengard, The Rapid Evaluation of Potential Fields in Particle Systems, MIT Press, Cambridge, MA, 1988.
[7] R.W. Hockney, J.W. Eastwood, Computer Simulation Using Particles, McGraw-Hill, New York, 1981.
[8] G. Hummer, J. Electrostat. 36 (1996) 285–291.
[9] J. Ivanic, K. Ruedenberg, J. Phys. Chem. 100 (1996) 6342–6347.
[10] J. Ivanic, K. Ruedenberg, J. Phys. Chem. A 102 (1998) 9099–9100.
[11] J.D. Jackson, Classical Electrodynamics, second ed., Wiley, New York, 1975.
[12] R.A. James, J. Comput. Phys. 25 (1977) 71–93.
[13] H. Johansen, P. Colella, J. Comput. Phys. 147 (1998) 60–85.
[14] K. Lackner, Comput. Phys. Commun. 12 (1976) 33–44.
[15] A. Mayo, SIAM J. Numer. Anal. 21 (1984) 285–299.
[16] P. McCorquodale, P. Colella, G.T. Balls, S.B. Baden, Commun. Appl. Math. Comput. Sci. 2 (2007) 57–81.
[17] P. McCorquodale, P. Colella, H. Johansen, J. Comput. Phys. 173 (2001) 620–635.
[18] A.J.T. Nonaka, A higher-order upwind method for viscoelastic flow, PhD thesis, University of California, Davis, 2007.
[19] V. Rokhlin, J. Comput. Phys. 60 (1985) 187–207.
[20] P. Schwartz, M. Barad, P. Colella, T. Ligocki, J. Comput. Phys. 211 (2006) 531–550.
[21] S.L. Sobolev, Partial Differential Equations of Mathematical Physics, Pergamon, Oxford, 1964.
[22] K. von Hagenow, K. Lackner, Computation of axisymmetric MHD equilibria, in: Proceedings of the Seventh Conference on the Numerical Simulation of Plasmas, pp. 140–143, 1975.
[23] Z.J. Wang, J. Comput. Phys. 153 (1999) 666–670.