The final publication is available at

https://doi.org/10.1016/j.jcp.2016.07.028

Copyright   Elsevier

Additional Information

Fast convolution with free-space Green's functions

Felipe Vico,[*] Leslie Greengard,[†] and Miguel Ferrando[*]

April 13, 2016

Abstract

We introduce a fast algorithm for computing volume potentials - that is, the convolution of a translation invariant, free-space Green's function with a compactly supported source distribution defined on a uniform grid. The algorithm relies on regularizing the Fourier transform of the Green's function by cutting off the interaction in physical space beyond the domain of interest. This permits the straightforward application of trapezoidal quadrature and the standard FFT, with superalgebraic convergence for smooth data. Moreover, the method can be interpreted as employing a Nystrom discretization of the corresponding integral operator, with matrix entries which can be obtained explicitly and rapidly. This is of use in the design of preconditioners or fast direct solvers for a variety of volume integral equations. The method proposed permits the computation of any derivative of the potential, at the cost of an additional FFT.

1 Introduction

Many problems in scientific computing require the solution of a constant coefficient elliptic partial differential equation subject to suitable boundary or radiation conditions. In many cases, the free-space Green's function for the corresponding equation is known but involves nonlocal (long-range) interactions. A typical example is the Helmholtz equation in $\mathbf{R}^d$

$$\Delta\phi + k^2\phi = f,$$

where ϕ can be thought of as an acoustic potential and f a known distribution of acoustic sources, which we assume to be supported in the bounded domain $D = [-\frac{1}{2}, \frac{1}{2}]^d$. This can be done without loss of generality by rescaling the Helmholtz parameter k. The solution which satisfies the Sommerfeld radiation condition is well-known to be

$$\phi(\boldsymbol{x}) = \int_D g_k(\boldsymbol{x} - \boldsymbol{y})\, f(\boldsymbol{y})\, d\boldsymbol{y}. \tag{1}$$

---

[*]Instituto de Telecomunicaciones y Aplicaciones Multimedia (ITEAM), Universidad Politècnica de València, 46022 València, Spain. *email*: `felipe.vico@gmail.com`, `mferrand@dcom.upv.es`.

[†]Courant Institute, New York University, New York, NY and Simons Foundation, New York, NY. *email*: `greengard@cims.nyu.edu`

1

where $g_k(\boldsymbol{r}) = \frac{1}{4i} H_0(kr)$ for $d = 2$ and $g_k(\boldsymbol{r}) = \frac{1}{4\pi} \frac{e^{ikr}}{r}$ for $d = 3$. Here, $r = \|\boldsymbol{r}\|_2$ and H_0 denotes the zeroth order Hankel function of the first kind.

Note that the interaction kernel is long-range, requiring fast algorithms to be practical, and singular at $r = 0$, requiring accurate quadrature techniques. In some applications, a third difficulty is that the source density f is highly inhomogeneous, requiring adaptive discretization. In such settings, intrinsically adaptive, hierarchical methods are required [1, 2, 3, 4]. When the density is smooth, however, and well-resolved by a uniform mesh, it is more convenient (and generally more efficient) to use Fourier methods. We restrict our attention to the latter case in the present paper.

There are two distinct ways in which Fourier methods can be applied to the computation of (1). The first is the direct discretization of the equation with a *locally-corrected* trapezoidal rule. In the two-dimensional setting, for example, one can discretize D with a uniform mesh of N^2 points and use the approximation

$$\phi(nh, mh) \approx \sum_{\substack{n',m' \in [-\frac{N}{2}, \frac{N}{2}] \\ |n-n'|, |m-m'| > k}} g_k((n - n')h, (m - m')h)\, f(n'h, m'h)\, h^2 \quad +$$

$$\sum_{\substack{n',m' \in [-\frac{N}{2}, \frac{N}{2}] \\ |n-n'|, |m-m'| \le k}} w_{n-n', m-m'} f(n'h, m'h)\ ,$$

where $h = \frac{1}{N}$. Several groups have shown that kth order accuracy can be achieved in this manner by a suitable choice of weights $w_{i,j}$ (see, for example, [5, 6, 7, 8, 9, 10]). The net sum takes the form of a discrete (aperiodic) convolution and, hence, can be computed using the FFT with zero-padding in $O(N^2 \log N)$ operations.

Alternatively, using the convolution theorem, one can write

$$\phi(\boldsymbol{x}) = \mathscr{F}^{-1}\left(\frac{F(s)}{|\boldsymbol{s}|^2 - k^2} \right) = \left(\frac{1}{2\pi} \right)^d \int_{\mathbf{R}^d} e^{i\boldsymbol{s}\cdot\boldsymbol{x}}\, \frac{F(\boldsymbol{s})}{|\boldsymbol{s}|^2 - k^2}\, d\boldsymbol{s}\,, \tag{2}$$

where

$$F(\boldsymbol{s}) = \mathscr{F}(f)(\boldsymbol{s}) = \int_D e^{-i\boldsymbol{s}\cdot\boldsymbol{x}} f(\boldsymbol{x})\, d\boldsymbol{x}\,. \tag{3}$$

$\mathscr{F}$ here denotes the Fourier transform. The fact that $f(\boldsymbol{x})$ is smooth permits us to compute the Fourier integral in (3) with "spectral" accuracy. It also ensures that the error in truncating the Fourier integral in the inverse transform (2) is rapidly decaying with $|\boldsymbol{s}|$. The principal difficulty in employing Fourier methods is the singularity $\frac{1}{|\boldsymbol{s}|^2 - k^2}$ in the integrand. In the case of the Poisson equation, this is simply $\frac{1}{|\boldsymbol{s}|^2}$.

It is possible to design high order rules for the inverse Fourier transform. In the case of the Poisson equation in three dimensions, for example, switching to spherical coordinates cancels the singularity entirely. Combining this with the nonuniform FFT yields more or

less optimal schemes in terms of CPU time (see [11] and the references therein). This approach becomes technically more complicated for the Helmholtz equation, where the singularity lives on the sphere $|\boldsymbol{s}| = k$.

It turns out that there is a simple method that works for all long-range Green's functions, independent of dimension, requires only the trapezoidal rule, achieves spectral accuracy, and is accelerated by the standard FFT. Moreover, the matrix entries corresponding to this high order method are easily computed - a useful feature for either preconditioning strategies or direct solvers when using volume integral methods to solve variable coefficient partial differential equations.

Let us suppose, for the sake of simplicity, that we seek the restriction of the solution $\phi(\boldsymbol{x})$ to the unit box $D \subset \mathbf{R}^d$. Then, the maximum distance between any source and target point in D is $\sqrt{d}$. We define

$$g_k^L(\boldsymbol{r}) = \begin{cases} \dfrac{1}{4i} H_0(kr) \operatorname{rect}\left(\dfrac{\mathrm{r}}{2\mathrm{L}}\right) & \text{if } d = 2 \\[2ex] \dfrac{1}{4\pi} \dfrac{e^{ikr}}{r} \operatorname{rect}\left(\dfrac{\mathrm{r}}{2\mathrm{L}}\right) & \text{if } d = 3 \end{cases} \tag{4}$$

with rect(x) defined to be the characteristic function for the unit interval:

$$\operatorname{rect}(\mathrm{x}) = \begin{cases} 1 & \text{for } |x| < 1/2 \\ 0 & \text{for } |x| > 1/2. \end{cases}$$

If we set $L > \sqrt{d}$ in d dimensions, then the solution (1) is clearly indistinguishable from

$$\phi(\boldsymbol{x}) = \int_D g_k^L(\boldsymbol{x} - \boldsymbol{y}) \, f(\boldsymbol{y}) \, d\boldsymbol{y}. \tag{5}$$

Since g_k^L is compactly supported, the Paley-Wiener theorem implies that its Fourier transform G_k^L is entire (and C^∞). Moreover, as we shall see below, it is straightforward to compute. In the case of the Laplace operator in three dimensions, for example, $G_0^L = 2(\frac{\sin(Ls/2)}{s})^2$. Thus, the Poisson equation in three dimensions has the solution

$$\phi(\boldsymbol{x}) = \frac{2}{(2\pi)^3} \int_{\mathbf{R}^3} e^{i\boldsymbol{s}\cdot\boldsymbol{x}} \left(\frac{\sin(L|\boldsymbol{s}|/2)}{|\boldsymbol{s}|}\right)^2 F(\boldsymbol{s}) \, d\boldsymbol{s} \,. \tag{6}$$

Discretization by the trapezoidal rule on the domain $[-\frac{N}{2}, \frac{N}{2}]^d$ permits rapid evaluation using nothing more than the FFT. The achieved accuracy is controlled by the rate of decay of $F(\boldsymbol{s})$, with spectral accuracy achieved for sufficiently smooth $f(\boldsymbol{x})$ [12].

Remark 1. *The approach described here is both elementary and quite general, but seems to have been overlooked in the numerical analysis literature. An exception is the paper [13] by Vainikko, who used volume Helmholtz potentials for the iterative solution of the Lippmann-Schwinger equation.*

3

2 Computing the Fourier transform of truncated translation-invariant kernels

Suppose now that $f(\boldsymbol{x})$ is a radially symmetric function: $f(\boldsymbol{x}) = f(r)$, where $r = |\boldsymbol{x}|$. Then its Fourier transform $F(\boldsymbol{s})$ is also radially symmetric. For $d = 3$, it is easy to verify that

$$F(\boldsymbol{s}) = 4\pi \int_0^\infty \frac{\sin(sr)}{sr} f(r)\, r^2\, dr \tag{7}$$

where $s = |\boldsymbol{s}|$. For $d = 2$,

$$F(\boldsymbol{s}) = 2\pi \int_0^\infty J_0(sr)\, f(r)\, r\, dr. \tag{8}$$

For the Laplace equation in three dimensions, with Green's function

$$g_0^L(\boldsymbol{x}) = \frac{1}{4\pi|\boldsymbol{x}|} \mathrm{rect}\left(\frac{|\boldsymbol{x}|}{2\mathrm{L}}\right), \tag{9}$$

we have

$$G_0^L(\boldsymbol{s}) := \mathscr{F}\left(g_0^L(\boldsymbol{x})\right)(\boldsymbol{s}) = 4\pi \int_0^L \frac{\sin(sr)}{sr} \frac{1}{4\pi r} r^2\, dr = 2\left(\frac{\sin(Ls/2)}{s}\right)^2, \tag{10}$$

an analytic function expressible as a power series in s^2.

In $\mathbb{R}^2$, where the Green's function for the Laplace equation is

$$g_0(\boldsymbol{x}) = \frac{-1}{2\pi} \log|\boldsymbol{x}|, \tag{11}$$

we obtain the Fourier transform:

$$G_0^L(\boldsymbol{s}) := 2\pi \int_0^{+\infty} J_0(sr) g_0^L(r) r\, dr = \frac{1 - J_0(Ls)}{s^2} - \frac{L\log(L)J_1(Ls)}{s}. \tag{12}$$

We set $L = 1.8 > \sqrt{3}$ in three dimensions and $L = 1.5 > \sqrt{2}$ in two dimensions. Note that, in the inverse Fourier transform (6), the frequency content of the integrand in the variable of integration $\boldsymbol{s}$ is determined by the maximum magnitude of $\boldsymbol{x}$, the magnitude of L and $F(\boldsymbol{s})$ itself. It is straightforward to check that the integrand is sufficiently sampled with a mesh that is four times finer than in the original box: a factor of two from the fact that we are carrying out an aperiodic convolution so that the frequency content of $e^{i\boldsymbol{s}\cdot\boldsymbol{x}}F(\boldsymbol{s})$ is twice greater and a factor of two from the oscillatory behavior of the Fourier transform of the truncated kernel. Thus, if the unit box is discretized using N points in each dimension, we now require a grid of size $4N$ in each dimension. We will see in section 4 that, after a precomputation step, this can be reduced to a factor of $2N$.

3 Truncated kernels of mathematical physics

We now apply the technique described above to a collection of Green's functions that arise in mathematical physics. The resulting kernels in Fourier space are always C^∞, as noted above, by the Paley-Wiener theorem [14]. The method is easily extended to the calculation of any derivative using spectral differentiation.

Tables 1 and 2 summarize the results for various PDEs in three and two dimensions, respectively. We omit the derivations which are straightforward.

Table 1: Spectral representations of Green's Functions in 3D

Diff. Operator	Green's function	Truncated Spectral Representation
Δ	$g_0(r) = \dfrac{1}{4\pi r}$	$G_0^L(s) = 2\left(\dfrac{\sin(Ls/2)}{s}\right)^2$
$\Delta + k^2$	$g_k(r) = \dfrac{e^{ikr}}{4\pi r}$	$G_k^L(s) = \dfrac{-1 + e^{iLk}(\cos(Ls) - i\frac{k}{s}\sin(Ls))}{(k-s)(k+s)}$
Δ^2	$g_b(r) = \dfrac{r}{8\pi}$	$G^L(s) = \dfrac{(2 - L^2 s^2)\cos(Ls) + 2Ls\sin(Ls) - 2}{2s^4}$
$\Delta(\Delta + k^2)$	$g_{0k}(r) = \dfrac{e^{ikr}}{4\pi r} - \dfrac{1}{4\pi r}$	$G_{0k}^L(s) = G_k^L(s) - G_0^L(s)$
$(\Delta + \mathbf{h}\cdot\nabla)$	$g_{\mathbf{h}}(\boldsymbol{x}) = \dfrac{e^{i|\mathbf{h}||\boldsymbol{x}|}}{4\pi|\boldsymbol{x}|}e^{i\mathbf{h}\cdot\boldsymbol{x}}$	$G_{\mathbf{h}}^L(\mathbf{s}) = G_{|\mathbf{h}|}^L(|\mathbf{s} - \mathbf{h}|)$

For illustration, we plot the spectral representations of the free-space and truncated Laplace and Helmholtz Green's functions in Fig. 1.

4 An explicit construction of the discretized volume integral operators

The method described above requires a grid of dimension $(4N)^d$ points in order to compute an accurate volume integral without aliasing error. We show now that, after a precomputation step, only an FFT of dimension $(2N)^d$ is required. To see this, let us consider the three dimensional setting, with the data in the unit box denoted by $\rho_{ijk} = \rho(ih, jh, kh)$ where $h = 1/N$ and $i, j, k \in \{-N/2 + 1, .., N/2\}$. The solution must then take the form of

5

Table 2: Spectral representations of Green's Functions in 2D

Diff. Operator	Green's function	Truncated Spectral Representation
Δ	$g_0(r) = \dfrac{-1}{2\pi}\log(r)$	$G_0^L(s) = \dfrac{1 - J_0(Ls)}{s^2} - \dfrac{L\log(L)J_1(Ls)}{s}$
$\Delta + k^2$	$g_k(r) = \dfrac{i}{4}H_0^{(1)}(kr)$	$G_k^L(s) = \dfrac{1 + \dfrac{i\pi}{2}LsJ_1(Ls)H_0^{(1)}(Lk)}{s^2 - k^2}$ $-\dfrac{\dfrac{i\pi}{2}LkJ_0(Ls)H_1^{(1)}(Lk)}{s^2 - k^2}$
Δ^2	$g_b(r) = -\dfrac{r^2}{8\pi}\big(\log(r) - 1\big)$	$G^L(s) = \dfrac{J_0(Ls) - 1}{s^4} - \dfrac{L^3(\log(L) - 1)J_1(Ls)}{4s}$ $+\dfrac{(L\log L)J_1(Ls)}{s^3} - \dfrac{L^2(2\log L - 1)J_0(Ls)}{4s^2}$
$\Delta(\Delta + k^2)$	$g_{0k}(r) = g_k(r) + g_0(r)$	$G_{0k}^L(s) = G_k^L(s) - G_0^L(s)$
$(\Delta + \mathbf{h}\cdot\nabla)$	$g_{\mathbf{h}}(r) = \dfrac{i}{4}H_0^{(1)}(|\mathbf{h}||\boldsymbol{x}|)e^{i\mathbf{h}\cdot\boldsymbol{x}}$	$G_{\mathbf{h}}^L(\mathbf{s}) = G_{|\mathbf{h}|}^L(|\mathbf{s} - \mathbf{h}|)$

a discrete convolution operator:

$$\phi_{i'j'k'} = \sum_{i,j,k} T(i' - i, j' - j, k' - k)\rho_{ijk}. \tag{13}$$

Thus, all entries of T can be determined by simply applying the operator to the special right-hand side $\rho_{ijk} = \delta_{i0}\delta_{j0}\delta_{k0}$. Subsequent applications of T to a vector can then be carried out using standard aperiodic convolution, which only requires zero-padding by a factor of 2.

Remark 2. *A side effect of this precomputation is that we have generated a discrete matrix corresponding to a high order accurate Nyström discretization of the original volume integral operator. This is useful when implementing linear algebraic tools such as hierarchical direct solvers, incomplete LU preconditioners, etc.*

It is worth plotting the resulting entries of T and comparing them to a naive trapezoidal approximation (which blows up when $i = i'$, $j = j'$ and $k = k'$). As can be seen in Fig. 2, our high order rule takes the form of a mollified Green's function - with no significant oscillations in sign or other difficulties that plague many high order quadrature generation techniques.
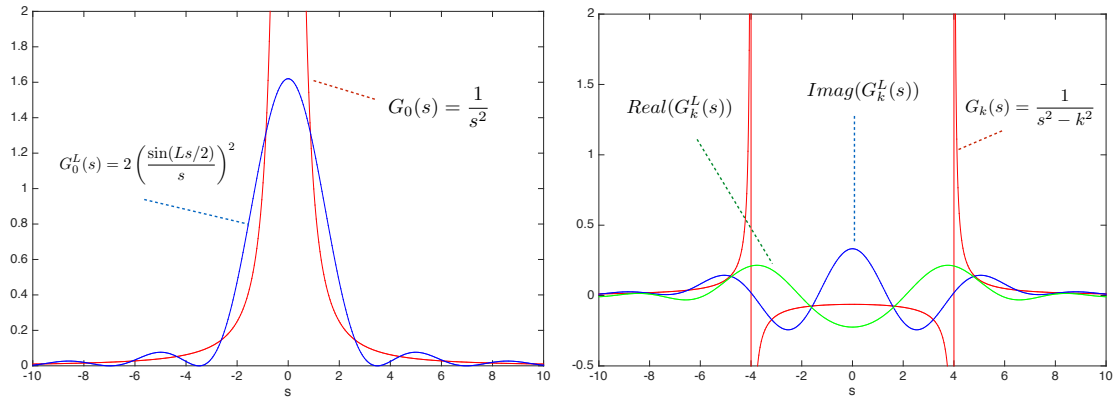
Figure 1: (l) Spectrum of the free-space Laplace kernel and the truncated Laplace kernel. (r) Spectrum of the free-space Helmholtz kernel for $k = 4$ and the truncated kernel with $L = 1.8$. Note that the truncated kernels are smooth but have introduced a slight oscillation.
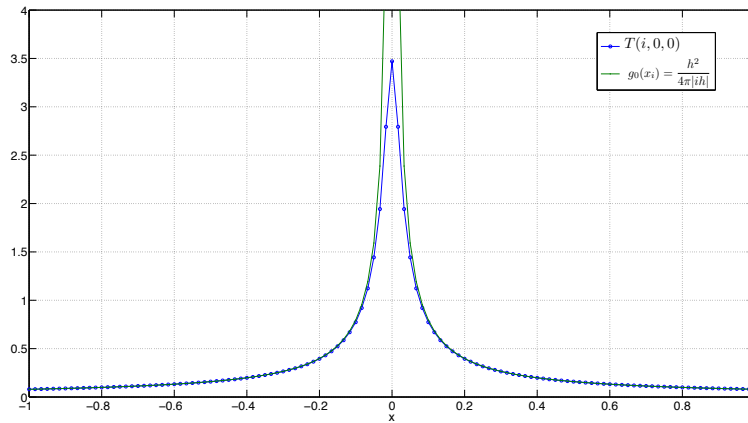

Figure 2: Comparing the naive trapezoidal rule on the original Green's function and the high order mollified Green's function along the line $y = z = 0$.

7

5 Numerical results

In this section, we illustrate the performance of the method described above. Our first examples simply involve convolution of the free-space Green's function with a Gaussian source, since the exact solution is available analytically. We also solve a variable dielectric Poisson-Boltzmann equation and a Lippmann-Schwinger type integral equation for variable medium scattering problems.

5.1 Convolution with a Gaussian source

Suppose now that, in three dimensions, the source distribution is given by a simple Gaussian:

$$\rho(r) := \frac{1}{\sigma^3 (2\pi)^{3/2}} e^{-\frac{r^2}{2\sigma^2}} . \tag{14}$$

Then, the solution to the Poisson equation is given by

$$[g_0 * \rho](\boldsymbol{x}) = \frac{1}{4\pi r} \operatorname{erf}\left(\frac{r}{\sqrt{2}\sigma}\right). \tag{15}$$

For the Helmholtz equation, the solution is a little more complicated but also straightforward to compute:

$$[g_k * \rho](\boldsymbol{x}) = \frac{1}{4\pi r} e^{-\frac{\sigma^2 k^2}{2}} \left[Real\left(e^{-ikr} \operatorname{erf}\left(\frac{2\sigma^2 ik - 2r}{2\sqrt{2\sigma^2}}\right)\right) - i\sin(kr) \right]. \tag{16}$$

For the biharmonic equation, we have

$$[g_b * \rho](\boldsymbol{x}) = \frac{1}{8\pi} \left[\sigma\sqrt{\frac{2}{\pi}} + \operatorname{erf}\left(\frac{r}{\sigma\sqrt{2}}\right)\left(\frac{\sigma^2}{r} + r\right) \right]. \tag{17}$$

Similarly, in two dimensions, with

$$\rho(r) := \frac{1}{2\pi\sigma^2} e^{-\frac{r^2}{2\sigma^2}} , \tag{18}$$

we have the following solution for the Poisson equation:

$$[g_0 * \rho](\boldsymbol{x}) = \frac{-1}{4\pi} \left[\operatorname{Ei}\left(\frac{r^2}{2\sigma^2}\right) + \log(r^2) \right]. \tag{19}$$

For the Helmholtz equation, we have

$$[g_k * \rho](\boldsymbol{x}) = \frac{H_0(kr)}{4\sigma^2} \int_0^r J_0(ky) \, e^{-\frac{y^2}{2\sigma^2}} \, ydy + \frac{J_0(kr)}{4\sigma^2} \int_r^{+\infty} H_0(kr) \, e^{-\frac{y^2}{2\sigma^2}} \, ydy \tag{20}$$

and for the biharmonic equation, we have

$$[g_b * \rho](\boldsymbol{x}) = -\frac{\sigma^2}{8\pi}\left[\left(\frac{r^2}{2\sigma^2} + 1\right)\widetilde{\mathrm{Ei}}\left(\frac{r^2}{2\sigma^2}\right) - e^{-\frac{r^2}{2\sigma^2}}\right] + c_2 r^2 + c_1\,r, \tag{21}$$

where

$$\widetilde{\mathrm{Ei}}(x) := \mathrm{Ei}(x) + \log(x) + \gamma$$

$$c_1 := \frac{\sigma^2}{8\pi}\left(\gamma + \log\left(\frac{1}{2\sigma^2}\right)\right) \tag{22}$$

$$c_2 := \frac{1}{8\pi}\left(\frac{\gamma}{2} + \frac{1}{2}\log\left(\frac{1}{2\sigma^2}\right) + 1\right).$$

In Fig. 3, we plot the error in each of these solutions when computed using the truncated Green's function Fourier method. Spectral accuracy is evident in each case.
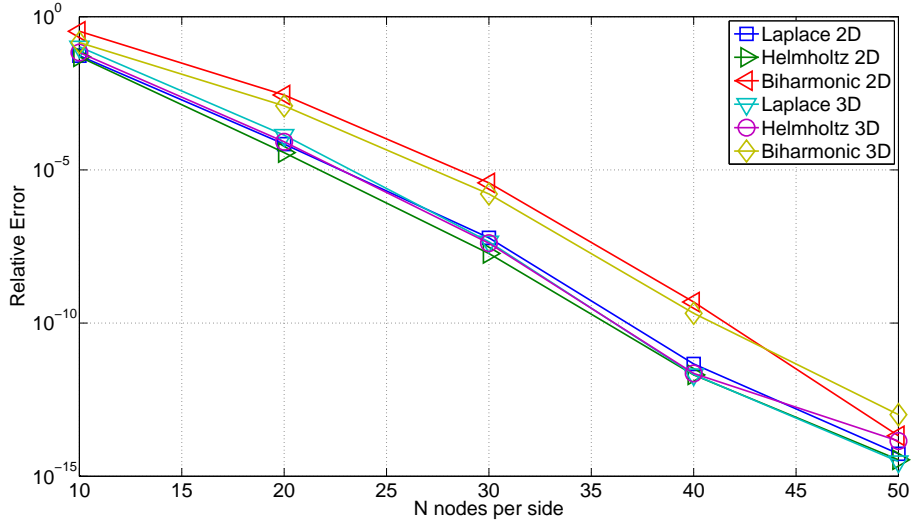

Figure 3: Convergence of the truncated Green's function Fourier method in solving the Poisson, Helmholtz and biharmonic equations in two and three dimensions for a single Gaussian source with $\sigma = 0.05$ (see eqs. (15-17), (19-21)). The Helmholtz parameter was set to $k = 2$.

5.2 Non-oscillatory elliptic equations with variable coefficients

A variety of problems in computational physics require the solution of the divergence-form elliptic partial differential equation

$$\nabla \cdot \epsilon(\boldsymbol{x})\nabla\phi - \lambda^2\phi = \rho(\boldsymbol{x}) \tag{23}$$

9

where ϵ is a known, smooth perturbation of a background constant $\epsilon + 0$, and where both $\rho(\boldsymbol{x})$ and $\epsilon - \epsilon_0$ have compact support.

In molecular electrostatics, this equation is referred to as the linearized Poisson-Boltzmann equation. While most models make use of a sharp dielectric interface (so that ϵ is piecewise constant), there is also interest in using approximations of the dielectric that vary smoothly [15, 16]. Following these references, a formula for ϵ is determined by first assuming we are given a macromolecule with M atoms, with a density

$$\alpha_i(\boldsymbol{x}) = exp[-r_i^2/(\mu^2 R_i^2)]$$

centered on the ith atom, where r_i denotes the distance of $\boldsymbol{x}$ from the atomic center, R_i is the van der Waals radius of the atom and μ is a user-specified variance. From this, a function

$$q(\boldsymbol{x}) = 1 - \prod_i [1 - \alpha_i(\boldsymbol{x})]$$

is constructed and, finally,

$$\epsilon(\boldsymbol{x}) = q(\boldsymbol{x})\,\epsilon_{in} + (1 - q(\boldsymbol{x})\,\epsilon_{out}.$$

In our example, we let $M = 1235$ with, $R_i = 0.022$ and $\mu^2 = 2$. Figure 4 shows the molecule represented as a union of spheres and the associated dielectric function $\epsilon(\boldsymbol{x})$ in darker gray.

Restricting our attention to the case $\lambda = 0$ for the sake of simplicity, we may represent the solution in the form

$$\phi(\boldsymbol{x}) = \int_{\mathbb{R}^3} \frac{1}{4\pi|\boldsymbol{x} - \boldsymbol{y}|}\sigma(\boldsymbol{y})dV_{\boldsymbol{y}}\,. \tag{24}$$

This leads directly to the following second kind Fredholm equation for the unknown density $\sigma(\boldsymbol{x})$:

$$-\epsilon(\boldsymbol{x})\sigma + \nabla\epsilon(\boldsymbol{x}) \cdot \nabla \int_{\mathbb{R}^3} \frac{1}{|\boldsymbol{x} - \boldsymbol{y}|}\sigma(\boldsymbol{y})dV_{\boldsymbol{y}} = \rho(\boldsymbol{x})\,. \tag{25}$$

We discretize σ on a uniform mesh with N^3 points and use the truncated Green's function Fourier method described above to convert (25) into a dense system of equations which we solve iteratively using GMRES. Each matrix vector product requires $O(N^3 \log N)$ operations using the FFT.

Our results are summarized in Table 3. $N_{tot} = N^3$ denotes the total number of unknowns, N_{iter} denotes the total number of GMRES iterations, E_2 denotes the relative error with respect to the reference solution in L_2 for $N = 250$, E_{inf} denotes the relative error in L_∞, and T_{solve} denotes the solution time in seconds required on a workstation with two Intel Xeon E5-2450 processors with 8 cores per processor and 64 GB of memory. $T_{precomp}$ denotes the time required for precomputation, as discussed in section 4, which requires a single FFT of dimension $(4N)^3$.
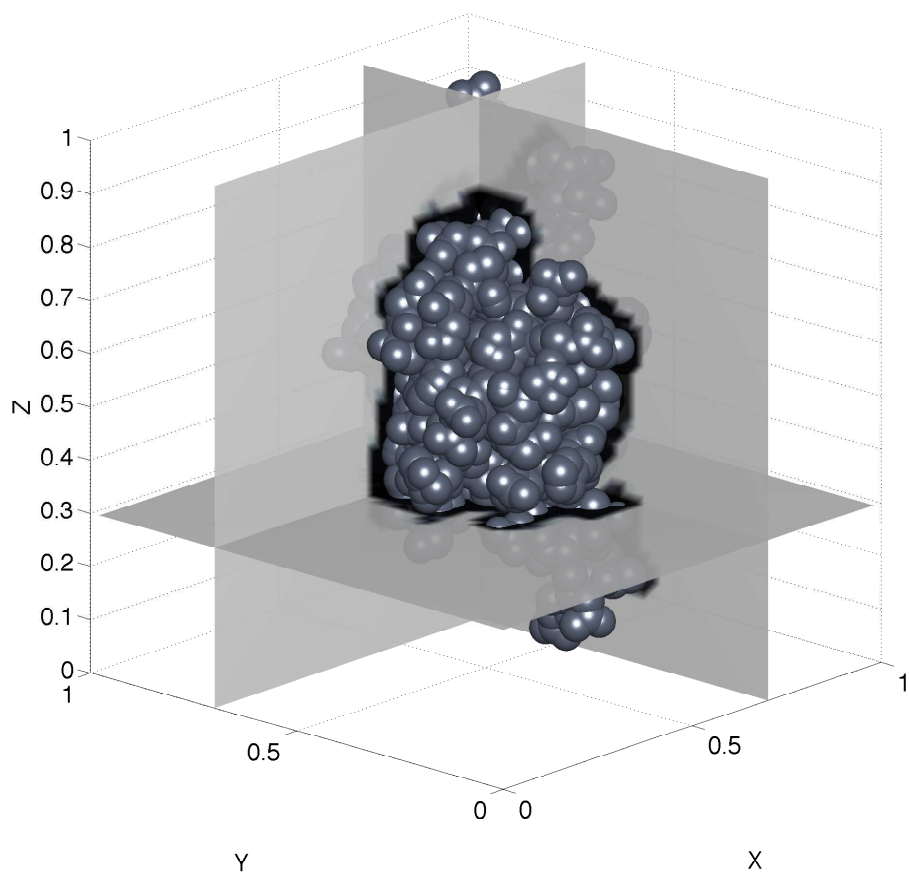
Figure 4: A protein molecule and the associated smooth dielectric function.

Table 3: Fast, iterative solution of the linearized Poisson-Boltzmann equation using the truncated Green's function Fourier method.

N_{tot}	N	E_2	E_{inf}	N_{iter}	T_{solve} (s)	$T_{precomp}$ (s)
1000000	100	2.08×10^{-8}	2.54×10^{-6}	12	14.3	4.4
3375000	150	2.74×10^{-10}	3.97×10^{-8}	16	76.9	13
8000000	200	7.92×10^{-12}	1.06×10^{-9}	16	205	28.9
15625000	250	$-$	$-$	16	421	75.8

5.3 Lippmann-Schwinger equation for wave scattering

In our last set of examples, we study the performance of the Lippmann-Schwinger integral equation for solving variable coefficient scattering problems in $\mathbb{R}^2$ and $\mathbb{R}^3$. The governing equation is a Helmholtz equation of the form

$$\Delta\phi^{\text{scat}} + k^2(1 + q(\boldsymbol{x}))\phi^{\text{scat}} = -k^2 q(\boldsymbol{x})\phi^{\text{inc}} \tag{26}$$

where ϕ^{scat} is assumed to satisfy the usual Sommerfeld radiation condition. We assume $q(\boldsymbol{x})$ has compact support. Using a volume integral representation for the solution in $\mathbb{R}^2$:

$$\phi^{\text{scat}}(\boldsymbol{x}) = \int_D H_0(k|\boldsymbol{x} - \boldsymbol{y}|)\,\sigma(\boldsymbol{y})\,dV_{\boldsymbol{y}}\,, \tag{27}$$

we obtain the second kind integral equation

$$-\sigma + k^2 q(\boldsymbol{x})\int_D H_0(k|\boldsymbol{x} - \boldsymbol{y}|)\,\sigma(\boldsymbol{y})\,dV_{\boldsymbol{y}} = -k^2 q(\boldsymbol{x})\phi^{\text{inc}}\,. \tag{28}$$

Similarly, in $\mathbb{R}^3$ we get:

$$-\sigma + k^2 q(\boldsymbol{x})\int_D \frac{e^{ik|\boldsymbol{x}-\boldsymbol{y}|}}{4\pi|\boldsymbol{x} - \boldsymbol{y}|}\,\sigma(\boldsymbol{y})\,dV_{\boldsymbol{y}} = -k^2 q(\boldsymbol{x})\phi^{\text{inc}}\,. \tag{29}$$

(This is the dual of the usual Lippmann-Schwinger equation.)

We consider four cases: a smoothly filtered flat dielectric disk in 2D, the 2D "Luneburg" lens, the 2D "Eaton" lens and a smoothed dielectric cube in 3D. The smoothly filtered disk (Fig. 5) has a contrast function given by

$$q(\boldsymbol{x}) = e^{-\frac{1}{2}\left(\frac{|\boldsymbol{x}|}{0.25}\right)^8}\,. \tag{30}$$

The Luneburg lens (Fig. 6) is designed to focus an incoming wave to a single point [17], with q given by

$$q(\boldsymbol{x}) = 1 - \left(\frac{|\boldsymbol{x}|}{0.45}\right)^2\,. \tag{31}$$

12

The Eaton lens (Fig. 7) is designed to bend light through an angle [18], with $q(\boldsymbol{x}) = n^2(\boldsymbol{x}) - 1$, where the refractive index n is given by the implicit equation

$$n^2(\boldsymbol{x}) = \frac{0.45}{n(\boldsymbol{x})|\boldsymbol{x}|} + \sqrt{\left(\frac{0.45}{n(\boldsymbol{x})|\boldsymbol{x}|}\right)^2 - 1} \tag{32}$$

In order to avoid the blowup in n at the origin, the refractive index is truncated at a maximum value of $n_{max} = \sqrt{3}$, corresponding to $q_{max} = 2$. Finally, the smoothed cube has a contrast function given by:

$$q(\boldsymbol{x}) = e^{-\frac{1}{2}\left(\left(\frac{x}{0.25}\right)^8 + \left(\frac{y}{0.25}\right)^8 + \left(\frac{z}{0.25}\right)^8\right)} . \tag{33}$$
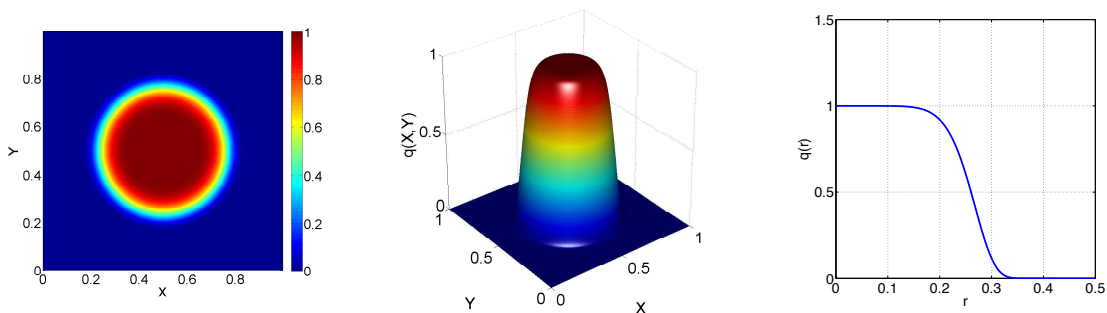

Figure 5: Contrast function $q(\boldsymbol{x})$ for the smoothly filtered disk plotted as colored contours (left), as a surface (center), and as a function of radius (right).

We solve each Lippmann-Schwinger equation iteratively, using Bi-CGStab with a tolerance of 10^{-12} for the iteration, since it has minimal storage requirements. This requires two matrix-vector products per iteration, each involving two applications of the FFT using the truncated Green's function Fourier method.

Except for the Eaton lens, the incoming wave is chosen to be a plane wave propagating to the right. The incoming wave for the Eaton lens is given by a Gaussian beam of the form

$$\phi^{\mathrm{inc}} = \overline{H_0(kR)}e^{-0.5k}$$

where

$$R = \sqrt{(x - x_0^c)^2 + (y - y_0^c)^2}, \ x_0^c = -0.01 - 0.5i, y_0^c = 0.77 .$$

Tables 4-7 show timings and errors for various frequencies and discretizations, while Figs. 8-10 show the computed solution. In these tables, size denotes the dimensions of
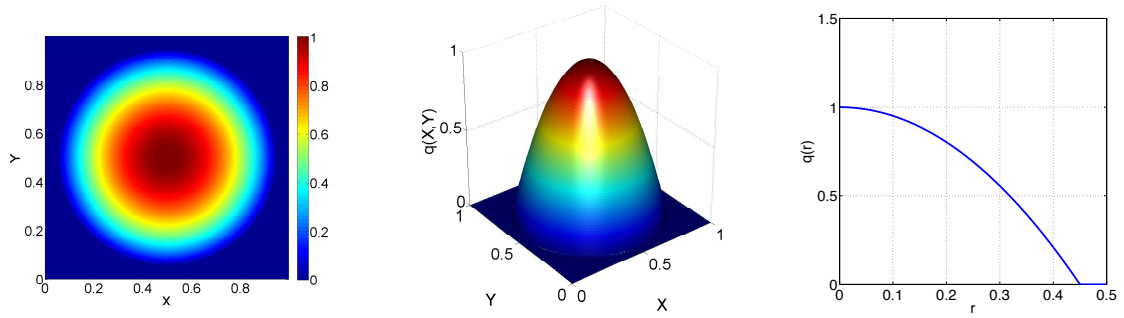
13

Figure 6: Contrast function $q(\boldsymbol{x})$ for the Luneburg lens plotted as colored contours (left), as a surface (center), and as a function of radius (right).
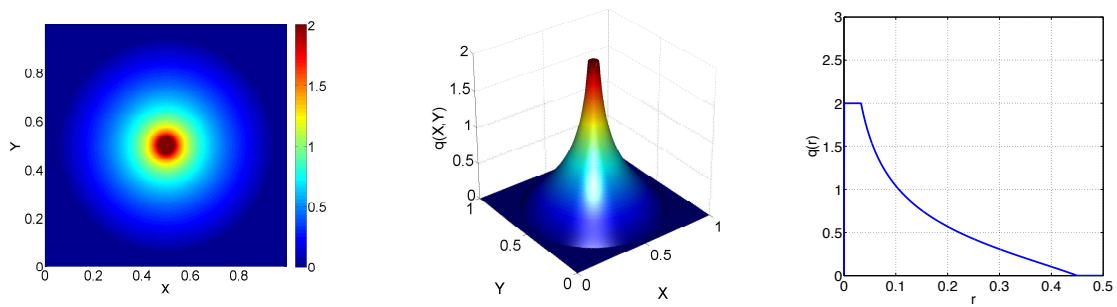

Figure 7: Contrast function $q(\boldsymbol{x})$ for the Eaton lens plotted as colored contours (left), as a surface (center), and as a function of radius (right).

14

the unit box in wavelengths, N_{tot} denotes the total number of points in the discretization, N denotes the number of points in a linear dimension, E_2 denotes the relative error with respect to the reference solution in L_2, E_{inf} denotes the relative error in L_∞, and N_{matvec} denotes the total number of matrix-vector products needed in the Bi-CGStab iteration. As above, T_{solve} denotes the solution time in seconds on a workstation with two Intel Xeon E5-2450 processors with 8 cores per processor and 64 GB of memory, and $T_{precomp}$ denotes the time required for precomputation, as discussed in section 4. A reference solution is computed using 6400×6400 points in the two-dimensional examples and using $300 \times 300 \times 300$ points in the three-dimensional example.

Note that spectral convergence rates are evident for smooth dielectric contrast functions. For the non-smooth Eaton and Luneburg lenses, the numerical convergence rate is closer to second order accuracy but with a small constant, so that high precision is achieved with a modest number of points per wavelength.
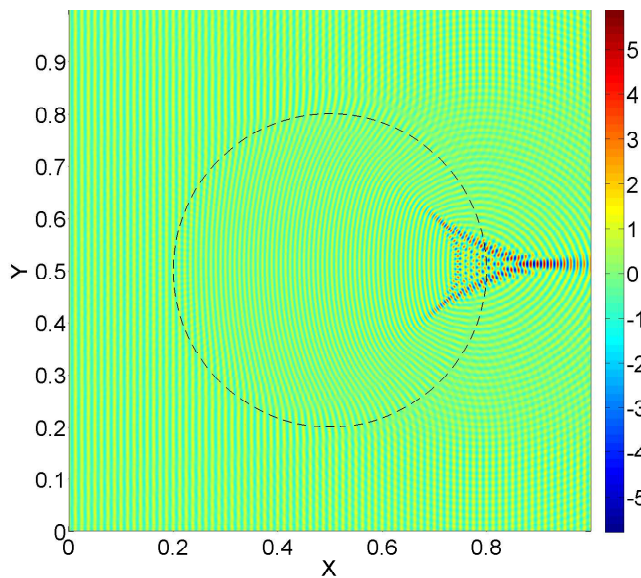

Figure 8: Scattering from a smoothly filtered disk with radius $R = 40\lambda_0$. in a unit square of size $80\lambda_0 \times 80\lambda_0$. We compute a reference solution with $N_{tot} = 40960000 = 6400^2$ points.

Table 4: Data for the smoothly filtered disk in two dimensions (see text for discussion). Timings are in seconds.

| Size (λ_0)	N_{tot}	N	E_2	E_{inf}	N_{matvec}	T_{solve}	$T_{precomp}$
1	400	20	1.4×10^{-4}	2.1×10^{-4}	17	0.382	0.222
1	2500	50	3.2×10^{-8}	3.2×10^{-8}	15	0.388	0.225
1	10000	100	8.7×10^{-13}	1.1×10^{-12}	15	0.632	0.5
1	40960000	6400	–	–	15	149	152
20	6400	80	4.2×10^{-5}	6.7×10^{-5}	332	1.38	0.26
20	10000	100	4.5×10^{-8}	8.8×10^{-8}	333	1.83	0.40
20	19600	140	4.1×10^{-11}	6.3×10^{-11}	335	2.39	0.303
20	40960000	6400	–	–	335	3170	143
80	62500	250	6.7×10^{-5}	1.0×10^{-4}	2938	58.5	0.503
80	72900	270	1.2×10^{-7}	2.2×10^{-7}	2990	67.1	0.518
80	102400	320	1.6×10^{-10}	2.8×10^{-10}	2906	83.1	0.61
80	40960000	6400	–	–	2948	29283	150
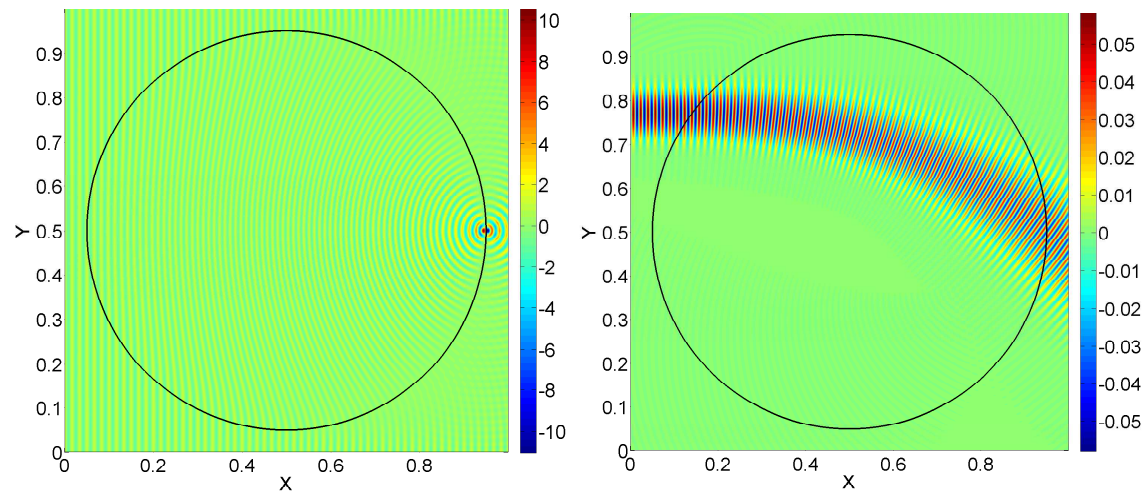
Figure 9: (left) Scattering by the two-dimensional Luneburg Lens, (right) Bending of an incoming beam by the two-dimensional Eaton Lens. In both cases, the lens radius is $R = 27\lambda_0$ in a unit square of size $60\lambda_0 \times 60\lambda_0$. The reference solution was computed with $N_{tot} = 40960000 = 6400^2$ points.

Table 5: Data for the two-dimensional Luneburg lens, where λ_0 denotes the *free-space* wavelength. A reference solution is computed using a 6400×6400 grid (except for $\lambda_0 = 60$), where a 3200×3200 grid is used. Timings are in seconds.

Size (λ_0)	N_{tot}	N	E_2	E_{inf}	N_{matvec}	T_{solve}	$T_{precomp}$
1	640000	800	1.26×10^{-7}	2.89×10^{-7}	17	2.41	2.09
1	2560000	1600	2.09×10^{-8}	5.09×10^{-8}	17	6.86	8.83
1	10240000	3200	2.93×10^{-9}	7.50×10^{-9}	17	40.3	34.9
20	640000	800	3.18×10^{-5}	4.01×10^{-5}	582	67.3	2.33
20	2560000	1600	5.84×10^{-6}	8.42×10^{-6}	581	243	9.02
20	10240000	3200	8.45×10^{-7}	1.39×10^{-6}	590	1190	34.3
40	640000	800	7.53×10^{-5}	8.59×10^{-5}	1415	163	2.37
40	2560000	1600	1.6×10^{-5}	1.88×10^{-5}	1393	740	9.24
40	10240000	3200	3.21×10^{-6}	3.89×10^{-6}	1321	2969	34.5
60	640000	800	1.26×10^{-4}	1.40×10^{-4}	3844	449	2.45
60	2250000	1500	3.13×10^{-5}	3.54×10^{-5}	3482	1611	8.27
60	10240000	3200	$-$	$-$	5220	12322	34.5

Table 6: Data for the two-dimensional Eaton lens, where λ_0 denotes the *free-space* wavelength. A reference solution is computed using a 6400×6400 grid and timings are in seconds.

| Size (λ_0)	N_{tot}	N	E_2	E_{inf}	N_{matvec}	T_{solve}	$T_{precomp}$
1	640000	800	1.01×10^{-7}	1.62×10^{-7}	15	1.82	1.85
1	2560000	1600	9.36×10^{-8}	1.75×10^{-7}	15	6.63	8.74
1	10240000	3200	3.53×10^{-9}	6.86×10^{-9}	15	35.5	34.4
20	640000	800	4.96×10^{-6}	1.54×10^{-5}	388	51.1	2.49
20	2560000	1600	8.50×10^{-7}	2.96×10^{-6}	390	203	9.10
20	10240000	3200	1.25×10^{-7}	4.77×10^{-7}	386	850	34.8
40	640000	800	1.4×10^{-5}	3.67×10^{-5}	958	105	2.26
40	2560000	1600	2.34×10^{-6}	8.77×10^{-6}	956	419	10.4
40	10240000	3200	3.98×10^{-7}	1.55×10^{-6}	968	2163	40
60	640000	800	2.81×10^{-5}	6.90×10^{-5}	2064	276	2.68
60	2560000	1600	4.61×10^{-6}	1.45×10^{-5}	2038	1065	9.06
60	10240000	3200	7.46×10^{-7}	2.87×10^{-6}	2024	4550	33
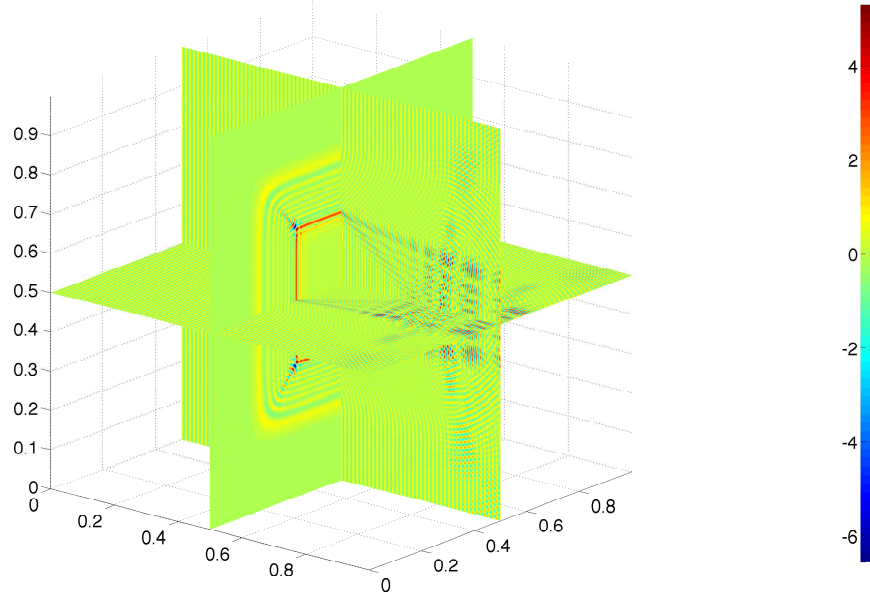
Figure 10: Three-dimensional scattering by a smoothed cube. We discretize a unit cell of dimension $80\lambda_0 \times 80\lambda_0 \times 80\lambda_0$ and compute a reference solution with $N_{tot} = 300^3$ points.

Table 7: Data for the smoothed cube in three dimensions with timings in seconds.

Size (λ_0)	N_{tot}	N	E_2	E_{inf}	N_{matvec}	T_{solve}	$T_{precomp}$
1	125000	50	4.08×10^{-8}	6.09×10^{-8}	15	1.17	1.07
1	343000	70	1.01×10^{-10}	1.25×10^{-10}	15	2.05	2.23
1	1000000	100	6.4×10^{-14}	7.91×10^{-14}	15	5.04	6.44
20	343000	70	7.23×10^{-4}	9.23×10^{-4}	449	44.2	2.11
20	1000000	100	3.84×10^{-8}	5.22×10^{-8}	441	124	6.15
20	3375000	150	8.57×10^{-13}	1.64×10^{-13}	434	411	20.1
40	3375000	150	4.6×10^{-7}	4.34×10^{-7}	895	891	19.3
40	8000000	200	6.76×10^{-12}	1.15×10^{-11}	907	1957	43.9
40	15625000	250	3.38×10^{-12}	5.76×10^{-12}	905	3428	99.1
60	3375000	150	4.76×10^{-1}	4.78×10^{-1}	8548	8534	19.8
60	8000000	200	5.67×10^{-6}	4.99×10^{-6}	1471	3538	48.5
60	15625000	250	3.38×10^{-9}	5.86×10^{-9}	1505	5688	97.8
80	15625000	250	4.86×10^{-5}	4.74×10^{-5}	2988	11332	97.5

6 Conclusions

We have presented a simple fast algorithm for computing volume potentials based on translation-invariant free-space Green's functions with compactly supported, smooth source distributions. By truncating the range of interaction in physical space, the Fourier transform of the kernel becomes an entire function, so that the trapezoidal rule yields superalgebraic convergence. Moreover, the transforms of the truncated kernels can be computed analytically. The principal advantages of our approach are that the standard FFT can be used for acceleration and that matrix entries for a high-order accurate Nyström discretization are available "on the fly". The latter is important in constructing hierarchical direct solvers or incomplete LU preconditioners.

We have illustrated the performance of the scheme on a variety of problems in two and three dimensions. For non-oscillatory problems, iterative methods are quite effective at solving variable coefficient partial differential equations when recast as volume integral equations. For scattering problems, it is well-known that the condition number grows with the size of the domain (measured in wavelengths). For problems up to approximately one hundred wavelengths in size, however, iterative schemes appear to be viable without preconditioning.

We will explore the use of these methods for full electromagnetic scattering problems in three dimensions in future work.

Acknowledgments

References

[1] F. Ethridge and L. Greengard. A New Fast-Multipole Accelerated Poisson Solver in Two Dimensions. *SIAM Journal on Scientific Computing*, 23(3):741–760, January 2001.

[2] H. Langston, L. Greengard, and D. Zorin. A free-space adaptive fmm-based pde solver in three dimensions. *Comm. Appl. Math. and Comp. Sci.*, 6:79–122, 2011.

[3] D. Malhotra and G. Biros. A parallel kernel independent fmm for particle and volume potentials. *Communications in Computational Physics*, 18:808–830, 2015.

[4] P. Mccorquodale, P. Colella, G. T. Balls, and S. B. Baden. A scalable parallel poisson solver in three dimensions with infinite-domain boundary conditions. In *In 7th International Workshop on High Performance Scientific and Engineering Computing*, pages 814–822, 2005.

[5] J. C. Aguilar and Y. Chen. High-Order Corrected Trapezoidal Quadrature Rules for Functions with a Logarithmic Singularity in 2-D. *Computers & Mathematics with Applications*, 44:1031–1039, 2002.

[6] J. C. Aguilar and Y. Chen. High-order corrected trapezoidal quadrature rules for the coulomb potential in three dimensions. *Computers and Mathematics with Applications*, 49(4):625–631, 2005.

[7] J. T. Beale and M.-C. Lai. A method for computing nearly singular integrals. *SIAM J. Numer. Anal*, 38:1902–25, 2001.

[8] R. Duan and V. Rokhlin. High-order quadratures for the solution of scattering problems in two dimensions. *Journal of Computational Physics*, 228(6):2152–2174, apr 2009.

[9] J. Goodman, T. Y. Hou, and J. Lowengrub. The convergence of the point vortex method for the 2-d euler equations. *Communications on Pure and Applied Mathematics*, 43:415–430, 1990.

[10] J. Lowengrub, M. Shelley, and B. Merriman. High-order and efficient methods for the vorticity formulation of the euler equations. *SIAM Journal on Scientific Computing*, 14:1107–1142, 1993.

[11] S. Jiang, L. Greengard, and W. Bao. Fast and accurate evaluation of nonlocal coulomb and dipole-dipole interactions via the nonuniform fft. *SIAM Journal on Scientific Computing*, 36:B777–B794, 2014.

[12] L. N. Trefethen. *Spectral methods in MATLAB*. SIAM, Philadelphia, 2000.

[13] G. Vainikko. Fast solvers of the Lippmann-Schwinger equation. *Direct and Inverse Problems of Mathematical Physics*, 5:423–440, 2000.

[14] L. Hörmander. *Linear Partial Differential Equations*. Springer, Berlin, 1976.

[15] L. Li, C. Li, Z. Zhang, and E. Alexov. On the dielectric constant of proteins: smooth dielectric function for macromolecular modeling and its implementation in delphi. *Journal of chemical theory and computation*, 9(4):2126–2136, 2013.

[16] J. A. Grant, B. T. Pickup, and A. A. Nicholls. A smooth permittivity function for Poisson-Boltzmann solvation methods. *Journal of Computational Chemistry*, 230(22):608–640, 2001.

[17] R. K. Luneburg and M. Herzberger. *Mathematical theory of optics.* Univ of California Press, 1964.

[18] A. J. Danner and U. Leonhardt. Lossless design of an Eaton lens and invisible sphere by transformation optics with no bandwidth limitation. In *Conference on Lasers and Electro-Optics*, page JThC4. Optical Society of America, 2009.