


Fast, High-Order Methods for Scattering by Inhomogeneous Media

Thesis by
E. McKay Hyde

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

2003
(Defended August 14, 2002)

© 2003

E. McKay Hyde

All Rights Reserved

For

Marie, Cooper and Miriam

Acknowledgements

The deeply enriching experience that I have enjoyed during my years at Caltech is largely due to the individuals that have assisted and influenced me and my work. I would be ungrateful if I did not acknowledge their invaluable contributions. First, I thank my advisor, Professor Oscar P. Bruno, for the countless hours spent advising my research and discussing my future career. He has always made me feel like a valued colleague. I also thank the remaining Applied and Computational Mathematics faculty and staff for their instruction, advice, and assistance as well as the members of my defense committee, Professors Oscar P. Bruno, Brent T. Fultz, Daniel I. Meiron, and Niles A. Pierce for their time and effort.

I extend my gratitude to the funding agencies that have supported my research. During my first year of research, I was supported through a grant from the Air Force Office of Scientific Research (AFOSR). I continue to attend the AFOSR's excellent annual workshop on electromagnetics. In the remaining years, I have been supported through a Computational Science Graduate Fellowship from the Department of Energy (DOE). This fellowship program has not only provided necessary funds, but has also greatly broadened my education and my associations by introducing me to DOE programs, laboratories, and scientists. In particular, as part of the fellowship program, I spent a summer in the Center for Applied Scientific Computing (CASC) at Lawrence Livermore National Laboratory. My experience at CASC was thoroughly rewarding due in large part to my supervisors, Dr. Milo R. Dorr and Dr. F. Xabier Garaizar. Finally, I am deeply grateful to the ARCS Foundation, whose fellowship has made supporting a family during graduate studies financially bearable. I will not forget their kindness and generosity.

My friends at Caltech have helped make life enjoyable even in those times when research was not. There are simply far too many people to mention here, but I am particularly grateful to Dave Amundsen, Rhett Brewer, Julian Chaubell, Christophe Geuzaine, Leonid Kunyansky, Randy Paffenroth, Matthew Pohlman, Chad Schmutzer, and Mayya Tokman

for many enjoyable hours of discussion on topics ranging from mathematics and computer science to religion and politics.

Most of all, I am grateful to my family. I am deeply grateful to my parents, Gerald and Carma Hyde. They gave me a love of learning and a belief in myself as well as faith in something greater. Their love and support have been a great strength to me. Finally, I express my deepest gratitude to my dear wife, Marie, and my wonderful children, Cooper and Miriam, who have sacrificed greatly to follow me as I pursue my dream. They are my motivation and my happiness. I dedicate this work to them.

Abstract

In this thesis, we introduce a new, fast, high-order method for scattering by inhomogeneous media in three dimensions. As in previously existing methods, the low ($\mathcal{O}(N \log N)$) complexity of our integral equation method is obtained through extensive use of the fast Fourier transform (FFT) in evaluating the required convolutions. Unlike previous FFT-based methods, however, this method yields high-order accuracy, even for scatterers containing geometric singularities such as discontinuities, corners, and cusps.

We begin our discussion with a thorough theoretical analysis of an efficient, high-order method recently introduced by Bruno and Sei (IEEE Trans. in Antenn. Propag., 2000), which motivated the present work. This two-dimensional method is based on a Fourier approximation of the integral equation in polar coordinates and a related, generally low-order, Fourier smoothing of the scatterer. The claim that use of this low-order approximation of the scatterer leads to a high-order accurate numerical method generated considerable controversy. Our proofs establish that this method indeed yields high-order accurate solutions. We also introduce several substantial improvements to the numerical implementation of this two-dimensional algorithm, which lead to increased numerical stability with decreased computational cost.

We then present our new, fast, high-order method in three dimensions. An immediate generalization of the polar coordinate approach in two dimensions to a spherical coordinate approach in three dimensions appears less advantageous than our chosen approach: Fourier approximation and integration in *Cartesian coordinates*. To obtain smooth and periodic functions (which are approximated to high-order via Fourier series), we 1) decompose the Green's function into a smooth part with infinite support and a singular part with compact support; and 2) replace, as in the two-dimensional approach, the (possibly discontinuous) scatterer with its truncated *Cartesian* Fourier series.

The accuracy of our three-dimensional method is approximately equal to that of the

two-dimensional method mentioned above and, interestingly, is actually much simpler than the two-dimensional approach. In addition to our theoretical discussion of these high-order methods, we present a parallel implementation of our three-dimensional Cartesian approach. The efficiency, high-order accuracy, and overall performance of both the polar and Cartesian methods are demonstrated through several computational examples.

Contents

Acknowledgements	iv
Abstract	vi
1 Introduction	1
1.1 Applications	1
1.2 Scattering Equations	3
1.3 Integral Equation Formulation	4
1.4 Previous Work	6
1.5 High-Order Convergence in FFT-Based Methods	8
1.6 Overview of Chapters	12
2 A Fast, High-Order Method in Two Dimensions: Theoretical Analysis	14
2.1 Approximate Integral Equation	14
2.2 Error Bounds	17
2.2.1 Error in Approximated Fourier Modes	18
2.2.2 Total Error in the Interior and Exterior Fields	22
2.3 Computation of the Angular Integral	25
3 A Fast, High-Order Method in Two Dimensions: Numerical Implementation	28
3.1 High-Order Numerical Integration	29
3.1.1 Angular Integration	29
3.1.2 Radial Integration	30
3.1.2.1 Resolution of Logarithmic Singularity	35
3.1.2.2 Evaluation of Integral Moments	37

3.2	Solution of the Linear System	41
3.2.1	Linear Solver	41
3.2.2	Preconditioner	42
4	A Fast, High-Order Method in Three Dimensions	48
4.1	High-Order Convolution with Smooth Part	51
4.2	High-Order Convolution with Singular Part	52
4.3	Fourier-Smoothed Scatterers	54
4.4	Computation of the Fourier Coefficients of the Green's Function	55
4.5	Implementation	57
5	Computational Examples	61
5.1	Two-Dimensional Computational Examples	62
5.1.1	Verification of Predicted Convergence Rates	63
5.1.2	Convergence in Radial Integration	67
5.1.3	Complex Scatterers and Preconditioning	73
5.2	Three-Dimensional Computational Examples	79
6	Conclusions	90
A	Technical Lemmas	93
A.1	Bound on Fourier Coefficient of Green's Function	93
A.2	Bound on Integral Operator	97
B	Accurate and Efficient Computation of Scaled Bessel Functions	100
C	High-Order Evaluation of Fourier Integrals	103
C.1	Second-Order Interpolating Polynomials	105
C.2	Fourth-Order Interpolating Polynomials	107
	Bibliography	109

List of Figures

1.1	Smooth and Periodic Function – $f(x) = e^{\cos^2 x}$	9
1.2	Example of Fourier Smoothing	11
5.1	Visualizations for Radially Layered Scatterer – Diameter = 4λ	65
5.2	Visualizations for Off-Center Disc – Centered at $(1\lambda, 0)$, Diameter = 1λ	66
5.3	Visualizations for Discontinuous Scatterer – Diameter = 10λ	68
5.4	Visualizations for $C^{0,\alpha}$ Scatterer – Diameter = 10λ	69
5.5	Visualizations for $C^{1,\alpha}$ Scatterer – Diameter = 10λ	70
5.6	Visualizations for Square Scatterer – Diagonal Length = 5λ	76
5.7	Visualizations for Star Scatterer – Diameter = 1λ	78
5.8	Visualizations for Bumpy Scatterer – Diameter = 5λ	80
5.9	Visualizations for Layered Sphere – $\kappa a = 4$	82
5.10	Far Field Intensity ($ u_\infty ^2$) for Layered Sphere – $\kappa a = 4$	83
5.11	Visualizations for Array of Smooth Scatterers – $5\lambda \times 5\lambda \times 5\lambda$	85
5.12	Far Field Intensity ($ u_\infty ^2$) for Array of Smooth Scatterers – $5\lambda \times 5\lambda \times 5\lambda$. .	86
5.13	Complex Scatterer – Close-Up Image of Scatterer	87
5.14	Visualizations for Complex Scatterer – Size $\approx 2.5\lambda \times 5\lambda \times 2.5\lambda$	88
5.15	Far Field Intensity ($ u_\infty ^2$) for Complex Scatterer – Size $\approx 2.5\lambda \times 5\lambda \times 2.5\lambda$.	89

List of Tables

1.1	Trapezoidal Rule Convergence for $f(x) = e^{\cos^2 x}$	9
1.2	High-Order Trapezoidal Rule Integration via Fourier Smoothing	10
5.1	Convergence Rate for Radially Layered Scatterer – Diameter = 4λ	64
5.2	Convergence for Off-Center Disc – Centered at $(1\lambda, 0)$, Diameter = 1λ	64
5.3	Convergence Rate for Discontinuous Scatterer – Diameter = 10λ	67
5.4	Convergence Rate for $C^{0,\alpha}$ Scatterer – Diameter = 10λ	71
5.5	Convergence Rate for $C^{1,\alpha}$ Scatterer – Diameter = 10λ	71
5.6	Convergence of Radial Integration for Disc Centered at Origin – 2λ Diameter	72
5.7	Convergence of Radial Integration for Disc Centered at $(1\lambda, 0)$ – 1λ Diameter	74
5.8	Convergence for Square Scatterer – Diagonal Length = 5λ	77
5.9	Convergence for Star Scatterer – Diameter = 1λ	77
5.10	Convergence for Bumpy Scatterer – Diameter = 5λ	79
5.11	Convergence for Layered Sphere – $\kappa a = 4$	83
5.12	Convergence for Array of Smooth Scatterers – $5\lambda \times 5\lambda \times 5\lambda$	84
5.13	Convergence for Complex Scatterer – Size $\approx 2.5\lambda \times 5\lambda \times 2.5\lambda$	89

Chapter 1

Introduction

Scattering theory remains an active and challenging field in science, engineering, and mathematics. In a broad sense, the term scattering refers to any situation in which a wave impinges on an obstacle and is thereby distorted, reflected, transmitted, or in some other way “scattered.” Thus, a scattered wave clearly contains information about the scattering obstacle itself. Hence, an understanding of the interactions between waves and obstacles should allow one to extract information about an obstacle from the waves scattered by it. This fact is one of the primary motivations behind the study of scattering phenomena.

Although most scattering problems take on similar mathematical forms, they find application in a wide range of fields including communications, materials science, plasma physics, biology and medicine, radar and remote sensing, etc. However, producing useful numerical solutions to such problems remains a challenge, requiring novel mathematical approaches and powerful computational tools. In the next two sections, we describe a number of interesting engineering and scientific applications where computational wave scattering plays an important role and we introduce the scattering equations to be used throughout this text.

1.1 Applications

As mentioned above, scattering applications are found in a wide variety of fields. Possibly the most familiar scattering applications are in radar and other techniques of remote sensing. These have a wide range of commercial, environmental, and military applications. For example, radar facilities track aircraft while remote sensing satellites, using radar and other technologies, collect atmospheric data, map the surface of the earth, and measure wind speeds at the oceans’ surface [22]. More sophisticated mathematical and computational

tools seek to extract even more information from the scattered waves than one can currently obtain. This could enable, for example, detection and identification of underground or underwater structures.

Of course, biological and medical applications are also of great practical importance. Ultrasound and x-ray imaging as well as the more recent technique of optical coherence tomography (OCT) are based on scattering phenomena. Like ultrasound imaging, OCT uses interference of incoherent waves (in this case infrared light) to determine biological microstructure [33, 43]. The combination of these techniques with new mathematical and computational methods may lead to much more powerful imaging technologies. Such efforts have paid dividends in the past. For example, Cormack and Hounsfield [34] developed computerized tomography (CT) by combining the physics of x-rays with the mathematics of tomography and efficient numerical methods. For this work, they shared the Nobel Prize in Medicine in 1979.

Materials science, particularly in the fields of microscopy and diffractometry, contains many important modern applications of scattering. Of course, transmission electron microscopy (TEM) and x-ray diffractometry have long been used to analyze material microstructure. More recently, techniques such as neutron diffraction and reflection high energy electron diffraction (RHEED) have been developed to further probe the structure of materials. In neutron diffraction, intense neutron beams interact either with nuclei or unpaired electrons in a material giving information on its structure [27, p. 156]. RHEED uses low-angle incident electron beams for analysis of the surface of a material [44, 50]. This provides, for example, real-time surface structure data during semiconductor film growth via molecular beam epitaxy (MBE).

Finally, although perhaps less familiar, the study of laser-plasma interactions is also closely related to scattering theory. When high-intensity laser light strikes a target, the material composing the target ablates and ionizes, producing a high-density plasma. This plasma then interacts with subsequent laser pulses producing a variety of waves and instabilities. Computational simulations can provide insight into the physics of these interactions [7, 21, 35].

These examples provide a glimpse of the wide range of fields in which scattering phenomena play an important role. Despite the great benefit that these techniques provide, much more potential benefit remains to be gained through the innovative use of mathematical

and computational tools.

1.2 Scattering Equations

Reflecting the diversity of fields in which scattering applications arise, several types of wave equations are used to describe scattering phenomena including the acoustic wave equation, Maxwell's equations and the Schrödinger equation. Under the assumption of time-harmonic scattering, many of these formulations reduce to the Helmholtz equation. Hence, although the Helmholtz equation does not capture all scattering phenomena, it does encompass many of the most important mathematical and computational issues arising in scattering theory. The precise formulation of the wave equation and the associated boundary conditions depend strongly on the type of scattering obstacle one considers. Many numerical approaches have been developed to treat various problem classes; these include the finite element and finite difference methods, Fourier-based methods and the fast multipole method.

In the following paragraphs, we describe a few of the most common problem classes. Throughout, we denote the total field by u , which is given by the sum of a given incident field u^i and a scattered field u^s , i.e.,

$$u = u^i + u^s.$$

The incident field satisfies

$$\Delta u^i + \kappa^2 u^i = 0, \tag{1.1}$$

in all of $\mathbb{R}^2$ or $\mathbb{R}^3$ where $\kappa = \frac{2\pi}{\lambda}$ is the wave number and λ is the wavelength of the incident wave. To guarantee that the scattered wave is outgoing, u^s satisfies the following Sommerfeld radiation conditions. In $\mathbb{R}^3$, u^s satisfies

$$\lim_{r \rightarrow \infty} r \left(\frac{\partial u^s}{\partial r} - i\kappa u^s \right) = 0, \tag{1.2}$$

where $r = |x|$ for $x \in \mathbb{R}^3$. In $\mathbb{R}^2$, u^s satisfies

$$\lim_{r \rightarrow \infty} \sqrt{r} \left(\frac{\partial u^s}{\partial r} - i\kappa u^s \right) = 0, \tag{1.3}$$

where $r = |x|$ for $x \in \mathbb{R}^2$.

The first class of problems we describe involves surface scattering by two-dimensional perfect electrical conductors in electromagnetics and two- or three-dimensional sound-soft obstacles in acoustics. Then, for a closed and bounded scatterer D , u satisfies the equations

$$\Delta u + \kappa^2 u = 0, \quad x \notin D,$$

$$u|_{\partial D} = 0. \tag{1.4}$$

Corresponding to the physical properties of the scattering object, other boundary conditions on ∂D may need to be used, including a vanishing normal derivative or an impedance boundary condition [17, pp. 2–7].

Another main class of scattering problems, often called volumetric scattering, involves scattering by penetrable, inhomogeneous media including, for example, dielectric material, biological tissue and ion-electron plasmas. In such cases we consider a bounded inhomogeneity with a variable refractive index n such that $n(x) = 1$ for x outside of some bounded set. (Note that if $n \neq 1$ but is constant outside the scatterer, we can scale κ such that $n = 1$ outside the scatterer. Note also that n need not represent the refractive index: for example, in acoustic scattering, n depends on the material density and in electron diffraction, n depends on a scattering potential.) Given $n(x)$, the total field u satisfies [17, p. 2]

$$\Delta u + \kappa^2 n^2(x) u = 0. \tag{1.5}$$

Thus, the presence of the inhomogeneous scatterer results in a Helmholtz equation with a variable coefficient.

1.3 Integral Equation Formulation

The subject of this thesis is the development and analysis of efficient, high-order computational methods for the solution of the volumetric scattering equation (1.5) in two and three dimensions. The algorithms available for the numerical solution of this equation fall into two broad classes. The first approach is the finite element or finite difference method. These methods have the advantage that, unlike other methods, they lead to sparse linear

systems. Their primary disadvantage, on the other hand, lies in the fact that in order to satisfy the Sommerfeld radiation condition (see (1.3) and (1.2)), one must use a large computational domain containing the scatterer and impose absorbing boundary conditions on the computational boundary [20, 36, 37, 46, 52]. Clearly, this procedure gives rise to large numbers of unknowns and correspondingly large linear systems.

A second class of algorithms is based on the use of integral equations (see (1.6)). These approaches have the advantage that the equation must only be discretized *on the scatterer itself*. Furthermore, the condition of radiation at infinity is *automatically* satisfied. On the other hand, integral equation methods have the disadvantage that they lead to dense linear systems. Therefore, a straightforward computation of the integral operator requires $\mathcal{O}(N^2)$ operations per iteration of an iterative linear solver. However, by reducing this complexity to $\mathcal{O}(N \log N)$ per iteration, integral equation methods become highly competitive with methods based on finite elements or finite differences. Thus, in this thesis, we develop computational methods based on a Lippmann-Schwinger integral equation of the form [17, p. 214]

$$u(x) = u^i(x) - \kappa^2 \int g(x-y)m(y)u(y)dy, \quad (1.6)$$

where $m = 1 - n^2$ and g is the Green's function for the Helmholtz equation. In two spatial dimensions, $g(x) = \frac{i}{4}H_0^1(\kappa|x|)$ whereas in three dimensions, $g(x) = \frac{e^{i\kappa|x|}}{4\pi|x|}$.

The fast multipole method (FMM) [15, 31, 47, 48] is perhaps the most widely known method for reducing the complexity of *surface* integral equations. Although the FMM has not been applied to the volumetric scattering problem (1.6), which we address in this thesis, its popularity requires that we mention it. We choose not to pursue an FMM-based approach, however, for the following reasons. As described several places in the literature, the FMM exhibits numerical instabilities at subwavelength spatial scales [18, 19, 30], thus limiting the attainable accuracy of the method. The authors of [30] suggested a low-frequency version of the FMM that is intended to be combined with the standard FMM to address this problem. However, although this article appeared four years ago, to our knowledge, no implementation of this approach has been developed. Finally, although the FMM exhibits $\mathcal{O}(N \log N)$ complexity, the large constant factor in this asymptotic bound requires a rather large value of N before any advantage over the straightforward $\mathcal{O}(N^2)$ approach is observed.

Another well-known method for reducing the complexity of integral equation methods to $\mathcal{O}(N \log N)$ operations per iteration is the k -space or conjugate-gradient FFT method (CG-FFT) [9, 51, 53]. In this method, the convolution with the Green’s function is computed via fast Fourier transforms (FFTs) and multiplication in Fourier space. Additionally, because of the small constant factor in the asymptotic complexity bound, efficiency gains over the straightforward $\mathcal{O}(N^2)$ approach are observed even for small values of N . However, although this method provides a reduced complexity, it is only first-order accurate. This low-order accuracy arises because the FFT provides a poor approximation to the Fourier transform when, as in this case, the function is not *smooth and periodic*.

Although our methods also use FFTs to achieve a reduced complexity, they yield, in addition, *high-order accuracy*. One reasonable measure of the effectiveness of a numerical scheme is the time required to obtain a solution for a given problem to within a given accuracy. In problems that require a high degree of accuracy, a high-order method with roughly the same computational complexity as available low-order methods will result in significant time and memory savings because many fewer points are required to obtain the given accuracy. This is especially true in three dimensions for which the number of unknowns (and hence the time and memory) scales as n^3 , where n is the number of points in each dimension. Furthermore, even when relatively low accuracy is required, high-order methods may exhibit advantages over low-order methods for two reasons. First, as we will show in our examples, a high-order method may obtain even a low level of accuracy with fewer points than a low-order method. Second, to estimate the accuracy of a computed solution, one typically compares the solution with a more accurate solution computed on a more refined computational grid; a high-order method yields such an increase in accuracy with only a slight increase in the number of unknowns, whereas a low-order method requires a much larger increase in problem size.

1.4 Previous Work

Despite these advantages, to our knowledge, only limited attempts have been made to develop high-order methods for this problem. Liu and Gedney [42] suggested a locally corrected Nyström scheme for scattering in two dimensions. (This approach is closely related to the high-order surface scattering method described in [14, 32].) This volumetric scattering

method provides high-order convergence rates that are not limited by the regularity of the scatterer. However, to achieve these rates, one must provide curvilinear cells for each level of discretization that accurately represent boundaries in the scatterer where discontinuities occur. Such a discretization seems rather difficult to obtain and would be even more difficult to obtain for general scatterers in three dimensions (although the method is only presented in two dimensions). In addition, this method, as presented, requires $\mathcal{O}(N^2)$ operations per iteration of the linear solver.

A fast, high-order method for *smooth* three-dimensional scatterers was proposed by Vainikko [49]. In this method, the integral equation is modified to produce a periodic solution by cutting off the Green's function (either smoothly or discontinuously) outside a cube that is at least twice as large as the scatterer. The solution to the modified integral equation is smooth and periodic on this larger cube and, furthermore, it agrees with the true solution on the support of the scatterer. Thus, for smooth scatterers, the solution is smooth and periodic and can, therefore, be approximated to high-order with a truncated Fourier series. As will be apparent, this method is somewhat related to our three-dimensional approach. The convergence rates of this approach, however, lag significantly behind those of our approach—producing only first-order convergence in the case of discontinuous scatterers. (This difference in convergence rates results primarily from our substitution of the scatterer by a truncated Fourier series, see Sections 1.5 and 4.3.) Vainikko introduces another, completely different, approach for piecewise smooth (discontinuous) scatterers that produces $\mathcal{O}(h^2(1 + \log h))$ convergence in both the near and far fields, where h is the discretization spacing in each direction. This approach requires that for each level of discretization, one must approximate the volume fraction of each cell that lies on each side of a discontinuity in the refractive index. This seems rather difficult to obtain, especially for complicated scatterers in three dimensions. In comparison, our two-dimensional approach requires only limited information about the geometry of the scatterer (namely, the Fourier coefficients $(mu)_\ell(r)$ along with the singularity points of these coefficients, see Chapter 3). Furthermore, our three-dimensional approach requires only the truncated Fourier series of the scatterer (see Chapter 4 for details).

1.5 High-Order Convergence in FFT-Based Methods

As mentioned previously, this thesis presents high-order accurate, FFT-based methods for scattering in two and three dimensions. The efficiency and high-order accuracy of these methods result from the following key facts. First, as is well known, for periodic integrands, the trapezoidal rule can be used to evaluate convolution integrals and Fourier coefficients and, in these cases, is algorithmically equivalent to the FFT. Second, the trapezoidal rule yields high-order convergence when integrating over the period of a *smooth and periodic* function and, similarly, a truncated Fourier series exhibits high-order convergence when approximating a *smooth and periodic* function. Finally, the Fourier smoothing of a discontinuous integrand, i.e., the replacement of a discontinuous integrand by its truncated Fourier series, allows high-order trapezoidal rule integration. We explain each of these concepts in more detail in the following paragraphs.

First, we consider the use of the trapezoidal rule in computing convolution integrals and Fourier coefficients. More precisely, in one dimension, using the trapezoidal rule to evaluate a convolution integral at N equally spaced points yields


$$\begin{aligned} \int_a^b f(x-y)g(y)dy &\approx h \sum_{k=0}^{N-1} f([a+jh] - [a+kh])g(a+kh) \\ &= h \sum_{k=0}^{N-1} f((j-k)h)g(a+kh), \end{aligned} \quad (1.7)$$

where $h = (b-a)/N$ and $k = 0, \dots, N-1$. Equation (1.7) is a discrete convolution, which can be evaluated in $\mathcal{O}(N \log N)$ operations by means of FFTs [45, pp. 531–537]. Similarly, using the trapezoidal rule to evaluate Fourier coefficients yields

$$\begin{aligned} f_\ell &= \frac{1}{b-a} \int_a^b f(x) e^{-2\pi i \ell (x-a)/(b-a)} dx \\ &\approx \frac{1}{N} \sum_{k=0}^{N-1} f(a+kh) e^{-2\pi i \ell k/N}. \end{aligned} \quad (1.8)$$

Clearly, (1.8) is a discrete Fourier transform, which we can evaluate in $\mathcal{O}(N \log N)$ operations by means of the FFT. Similarly, it is not difficult to see that FFTs can efficiently evaluate a truncated Fourier series on a set of equally spaced grid points.

We now consider the high-order accuracy obtained by means of trapezoidal rule integra-

Figure 1.1: Smooth and Periodic Function – $f(x) = e^{\cos^2 x}$

N	Abs. Error	Ratio
1	4.77(-2)	
2	1.19(-2)	4.03
4	2.95(-3)	4.02
8	7.36(-4)	4.01
8192	7.01(-10)	

(a) Convergence for $\int_0^{\pi/4} f(x)dx$

N	Abs. Error	Ratio
1	5.50(-1)	
2	6.03(-2)	9.12
4	3.10(-4)	1.95(2)
8	7.17(-10)	4.32(5)
16	2.10(-23)	3.42(13)

(b) Convergence for $\int_0^{2\pi} f(x)dx$ Table 1.1: Trapezoidal Rule Convergence for $f(x) = e^{\cos^2 x}$

tion or Fourier approximation of smooth and periodic functions. The high-order accuracy in these cases follows from the rapid decay of the Fourier coefficients of smooth and periodic functions (see Lemma 2.4). For example, consider the integration of the analytic function $f(x) = e^{\cos^2 x}$ (see Figure 1.1), over one quarter of its period $[0, \pi/4]$ and over its full period $[0, \pi]$ (see Table 1.1(a) and (b), respectively). One easily observes the second-order convergence when integrating over one quarter of its period and the *super-algebraic* convergence when integrating over the full period. This high-order accuracy results because the trapezoidal rule integrates the first N Fourier modes of the function *exactly* and thus, the convergence rate depends on the decay rate of the Fourier coefficients, which is exponential in this case.

The final key aspect of our approach is the Fourier smoothing of discontinuous scatterers. Some of the integrands that we encounter in our approach contain the product of a non-smooth (often discontinuous) scatterer and a significantly smoother periodic function. Direct trapezoidal rule integration of this product yields only first-order accuracy. On the

N	Abs. Error
4	0.264
8	6.42(-2)
16	4.71(-2)
32	1.20(-2)
64	1.07(-2)
128	5.13(-3)
256	2.62(-3)

(a) Convergence for $\int_{-1}^1 f(x)g(x)dx$

N	F	Abs. Error
4	2	6.93(-2)
8	4	4.11(-4)
16	8	4.87(-4)
32	16	3.86(-5)
64	32	4.96(-6)
128	64	7.25(-7)
256	128	6.68(-8)

(b) Convergence for $\int_{-1}^1 f^F(x)g(x)dx$

Table 1.2: High-Order Trapezoidal Rule Integration via Fourier Smoothing

other hand, although perhaps counterintuitive, substitution of the scatterer in the integrand by a truncated Fourier series leads to high-order accuracy in the integration. Because this approach has generated some controversy, we present a simple example of this fact in one dimension. In the case of a discontinuous scatterer, the solution $u \in C^1$ because of the regularizing properties of the integral operator [24, p. 78]. Hence, consider the integral of the product of a discontinuous function $f(x)$ and a C^1 , periodic function $g(x)$ over the period of g . We replace f by its truncated Fourier series with the same period as g

$$f^F(x) = \sum_{\ell=-F}^F f_{\ell} e^{2\pi i \ell / (b-a)x},$$

where the interval $[a, b]$ is the period of g (see Figure 1.2). (Note that the Fourier coefficients f_{ℓ} must be known either analytically or be computed very accurately.) Table 1.2 compares the accuracy obtained by means of the trapezoidal rule with and without the substitution of f by f^F . As expected, without the Fourier smoothing, one obtains only first-order convergence. With the Fourier smoothing, however, we observe approximately third-order convergence.

This is a rather surprising result at first glance since the truncated Fourier approximation of the discontinuous function $f(x)$ converges quite slowly. Of course, this first intuition is correct if one were to attempt to approximate the function $(fg)(x)$ itself in this manner. When approximating the *integral* of this product, however, this example shows that Fourier smoothing does indeed yields high-order accuracy. (Note that the convergence rate in this example is somewhat unsteady, perhaps due to the high-order Fourier modes of g that appear in the error. Despite this unsteady behavior, the convergence rate exceeds third-order in the sense of geometric mean. Somewhat similar convergence behavior is observed




(a) Discontinuous Function f (b) Fourier-Smoothed Function f^F (c) C^1 Function g

Figure 1.2: Example of Fourier Smoothing

in the computational results of the three-dimensional method as given in Section 5.2.)

The ideas described above form the basis for our high-order accurate, FFT-based methods. Clearly, high-order accuracy in most of these examples required smooth and periodic functions. As initially posed, however, the scattering problem does not involve smooth and periodic functions; on the contrary, the Green's function is singular, the fields are not periodic and the scatterer is often discontinuous. Hence, our numerical methods center around a reformulation of the problem that allows the use of these high-order approaches.

1.6 Overview of Chapters

The main result of this thesis is a new, efficient, high-order method for volumetric scattering in three dimensions. This method achieves high-order convergence even for scatterers containing geometric singularities such as discontinuities, corners and cusps. Before introducing this method, we present in Chapter 2 a thorough theoretical analysis of an efficient, high-order method in two dimensions, first introduced in [13]. This method partially motivated our approach in three dimensions. High-order accuracy in this two-dimensional method is obtained by representing the total field and the Green's function as truncated Fourier series in polar coordinates, i.e., as truncated Fourier series in the angular variable at each radius. As will be shown, this representation implies a (generally low-order) Fourier smoothing of the scatterer. The claim that this low-order approximation of the scatterer nonetheless leads to a high-order accuracy numerical method generated considerable controversy. Hence, we prove that the method indeed yields high-order convergence (at least third-order in the far field) and relate the convergence rate to the regularity of the scatterer.

In Chapter 3, we present substantial practical improvements to the original numerical implementation of this two-dimensional approach. For example, we make use of a much more efficient and stable radial integration method based on Chebyshev polynomials. Furthermore, we present a new efficient preconditioner, which substantially reduces the number of required linear solver iterations for many scattering configurations. Finally, in Appendix B, we present an efficient and stable method for computing scaled high-order Bessel functions, which allow us to avoid underflow and overflow errors in large computations.

In Chapter 4, we present our new, efficient, high-order method in three dimensions. Instead of directly generalizing the two-dimensional polar coordinates approach to three-

dimensional spherical coordinates (with the associated requirement of a fast spherical harmonics transform), we choose to base this three-dimensional method on Fourier approximation and integration in *Cartesian coordinates*. High-order accuracy in this case is obtained through a smooth decomposition of the Green's function by means of a partition of unity into a smooth part with infinite support and a singular part with compact support as well as through Fourier smoothing of the scatterer as described above. Interestingly, this Cartesian approach in three dimensions is much simpler than the two-dimensional polar coordinates approach, while yielding approximately the same order of accuracy. Additionally, we describe our fully parallel implementation of this approach.

Chapter 5 contains several computational examples to illustrate the computational complexity, the high-order accuracy, and the overall performance of both the two- and three-dimensional methods. Finally, in Chapter 6, we present brief conclusions and describe possible future research directions.

Chapter 2

A Fast, High-Order Method in Two Dimensions: Theoretical Analysis

In this chapter, we present a theoretical analysis of the efficient, high-order method introduced in [13]. The high-order accuracy of this method derives from the concepts introduced in Section 1.5. However, this theoretical analysis is not simply an academic exercise, but rather puts to rest the considerable controversy generated by the claim of high-order accuracy. Additionally, as we will describe, this analysis reveals interesting and unexpected bounds on the convergence rates.

2.1 Approximate Integral Equation

As described in the introduction, the field scattered by an bounded inhomogeneity is given by the solution of the Lippmann-Schwinger integral equation (1.6). Although several existence and uniqueness results are known for this problem, the existing results do not address the full generality of the problem that we consider. Specifically, we assume that, given $m \in L^\infty$, (1.6) admits a unique solution $u \in L^\infty$. Certainly, this is the case for $m \in C^1$ as proven in [17, §8.1, 8.3]. Furthermore, it is well known that (1.6) admits a unique solution for constant inhomogeneities by consideration of the appropriate surface integral equation [16, pp. 100–101]; [39].

To obtain the numerical method, we construct an approximate integral equation. As mentioned in the introduction, we approximate the Green's function by a truncated Fourier series. We now prove that this approximate integral equation also admits a unique solution.

The related fact that the inverse operator for the approximate problem is uniformly bounded provides an essential part of our proof of high-order convergence.

After changing to polar coordinates, $x = ae^{i\phi}$ and $y = re^{i\theta}$, we obtain the Fourier representation of the Green's function by means of the addition theorem for the Hankel function [17, p. 67]

$$H_0^1(\kappa|ae^{i\phi} - re^{i\theta}|) = \sum_{\ell=-\infty}^{\infty} \mathcal{J}_\ell(a, r) e^{i\ell(\phi-\theta)},$$

where

$$\mathcal{J}_\ell(a, r) = H_\ell^1(\kappa \max(a, r)) J_\ell(\kappa \min(a, r)). \quad (2.1)$$

This allows us to expand the integral operator, Ku , (see (1.6)) in a Fourier series

$$(Ku)(a, \phi) = -\frac{i\kappa^2}{4} \int H_0^1(\kappa|x-y|) m(y) u(y) dy = \sum_{\ell=-\infty}^{\infty} (K_\ell u)(a) e^{i\ell\phi},$$

where

$$(K_\ell u)(a) = -i\frac{\kappa^2}{4} \int_{R_0}^{R_1} \mathcal{J}_\ell(a, r) \left[\int_0^{2\pi} m(r, \theta) u(r, \theta) e^{-i\ell\theta} d\theta \right] r dr. \quad (2.2)$$

Here $R_0 \leq a \leq R_1$ defines an annular region that contains the support of m .

To obtain the approximate integral equation, we truncate this Fourier series. Furthermore, we also replace the incident field by its truncated Fourier series. Although not necessary for our results, this simplifies the presentation somewhat. We thereby obtain

$$v(a, \phi) = u^{i,M}(a, \phi) + (K^M v)(a, \phi), \quad (2.3)$$

where

$$u^{i,M}(a, \phi) = \sum_{\ell=-M}^M u_\ell^i(a) e^{i\ell\phi} \quad (2.4)$$

$$(K^M v)(a, \phi) = \sum_{\ell=-M}^M (K_\ell v)(a) e^{i\ell\phi}. \quad (2.5)$$

(Note: we will use superscript M throughout this text to denote the truncated Fourier series of order M of a given function, e.g., (2.4) and (2.5).)

Decomposing (2.3) into Fourier modes, we observe that

$$v_\ell(a) = \begin{cases} u_\ell^i(a) + (K_\ell v)(a), & \text{for } |\ell| \leq M \\ 0, & \text{for } |\ell| > M. \end{cases}$$

Hence,

$$v(a, \phi) = v^M(a, \phi)$$

and solving (2.3) is equivalent to solving the following system of one-dimensional integral equations

$$v_\ell(a) - (K_\ell v^M)(a) = u_\ell^i(a), \ell = -M, \dots, M. \quad (2.6)$$

To prove existence and uniqueness for this approximate integral equation, we make use of the following technical lemma.

Lemma 2.1. *There exists a constant $C > 0$ depending only on R_0 , R_1 and κ such that*

$$\left\| \int_{R_0}^{R_1} |\mathcal{J}_\ell(a, r)| r dr \right\|_\infty \leq \frac{C}{\max\{1, \ell^2\}},$$

where $\mathcal{J}_\ell(a, r)$ is defined in (2.1).

This lemma is proved in Appendix A.1 and allows us to prove the following lemma.

Lemma 2.2. *For any $m \in L^\infty$,*

$$\|K - K^M\|_\infty \rightarrow 0, \text{ as } M \rightarrow \infty,$$

where the operator norm is the one induced by the L^∞ -norm.

Proof. Let $u \in L^\infty$. Then,

$$\int_0^{2\pi} \left| m(r, \theta) u(r, \theta) e^{-i\ell\theta} \right| d\theta \leq 2\pi \|m\|_\infty \|u\|_\infty.$$

Hence, for $M \geq 0$,

$$\begin{aligned}
\|(K - K^M)u\|_\infty &\leq \frac{\pi\kappa^2}{2} \|m\|_\infty \|u\|_\infty \sum_{|\ell| > M} \left\| \int_{R_0}^{R_1} |\mathcal{J}_\ell(a, r)| r dr \right\|_\infty \\
&= \mathcal{O} \left(\sum_{|\ell| > M} \frac{1}{\ell^2} \right) \|u\|_\infty \\
&= \mathcal{O}(M^{-1}) \|u\|_\infty.
\end{aligned}$$

Therefore,

$$\|K - K^M\|_\infty = \mathcal{O}(M^{-1}) \rightarrow 0,$$

as $M \rightarrow \infty$. □

This leads to the desired uniqueness and existence result.

Theorem 2.3. *Given $m \in L^\infty$, for M sufficiently large (2.3) admits a unique solution $v \in L^\infty$ for incident field u^i satisfying (1.1). Furthermore, for sufficiently large M , the operators $(I - K^M)^{-1}$ exist and are uniformly bounded.*

Proof. Since we have assumed that (1.6) admits a unique solution, $I - K$ has a bounded inverse. Since L^∞ is a Banach space, Lemma 2.2 and [38, Theorem 10.1, p. 142] imply that for all sufficiently large M the inverse operators $(I - K^M)^{-1}$ exist and are uniformly bounded. Therefore, (2.3) admits a unique solution $v \in L^\infty$, for M sufficiently large. □

2.2 Error Bounds

In summary, the approximate integral equation (2.3) is obtained by truncating the Fourier series of both the incident field u^i and the integral operator K at each radius to obtain an approximate solution v to (1.6), which itself is a truncated Fourier series. Roughly speaking, high-order accuracy is obtained because the integral operator Ku and the incident field u^i are smooth and periodic functions as a function of the angular variable. Therefore, their truncated Fourier series exhibit high-order convergence. In this section we derive bounds on the error in this approximation. Of course the full numerical implementation of the method introduces additional errors (e.g., quadrature errors), but here we consider only the accuracy of v , the solution of the approximate integral equation (2.3). All other errors are considered in our discussion of the numerical implementation of this method in Chapter 3.

2.2.1 Error in Approximated Fourier Modes

The error in the solution of the approximate integral equation (2.3) is given by

$$|u(x) - v^M(x)| \leq |u^M(x) - v^M(x)| + |u^T(x)|,$$

where u^T is the “tail” of the Fourier series of u ,

$$u^T(a, \phi) = u(a, \phi) - u^M(a, \phi) = \sum_{|\ell| > M} u_\ell(a) e^{i\ell\phi}.$$

In this section, we derive a bound on the first of these two terms. Note that

$$\begin{aligned} u^M &= u^{i,M} + K^M u, \\ v^M &= u^{i,M} + K^M v^M. \end{aligned}$$

Hence, taking the difference of these equations we obtain

$$\begin{aligned} u^M - v^M &= K^M(u - v^M) \\ &= K^M(u^M - v^M) + K^M u^T. \end{aligned}$$

Theorem 2.3 ensures that $I - K^M$ is invertible for M sufficiently large and that the operators $(I - K^M)^{-1}$ are uniformly bounded by some $B > 0$, i.e., $\|(I - K^M)^{-1}\|_\infty \leq B$ for all sufficiently large M . Then

$$\varepsilon_M = \|u^M - v^M\|_\infty \leq B \|K^M u^T\|_\infty.$$

Hence, we seek a bound on $K^M u^T$.

We first note that

$$\|K^M u^T\|_\infty \leq \sum_{\ell=-M}^M \|K_\ell u^T\|_\infty.$$

Furthermore, note that

$$\int_0^{2\pi} m(r, \theta) u^T(r, \theta) e^{-i\ell\theta} d\theta = \sum_{|j| > M} m_{\ell-j}(r) u_j(r).$$

Then we have

$$\|K_\ell u^T\|_\infty \leq \frac{C}{\max\{1, \ell^2\}} \sum_{|j|>M} \|m_{\ell-j}\|_\infty \|u_j\|_\infty. \quad (2.7)$$

To bound this expression we use bounds on the Fourier coefficients of m and u . Let the space $C_{pw}^{k,\alpha}$ denote the space of functions whose k^{th} derivative is piecewise Hölder continuous. (Note: throughout this text we assume that $0 < \alpha < 1$.) We say that a function f is piecewise continuous on Ω if and only if there is a finite number of open, disjoint subsets of Ω , denoted by $\Omega_1, \dots, \Omega_n$, such that $\bigcup_{i=1}^n \overline{\Omega_i} = \overline{\Omega}$ where $f|_{\Omega_i}$ can be extended as a continuous function to $\overline{\Omega_i}$ for each $i = 1, \dots, n$. Given these definitions, the following lemma follows by repeated integration by parts.

Lemma 2.4. *If $g \in C^k([0, 2\pi]) \cap C_{pw}^{k+2}([0, 2\pi])$, then the Fourier coefficients of g ,*

$$\begin{aligned} |c_\ell| &= \left| \frac{1}{2\pi} \int_0^{2\pi} g(\theta) e^{-i\ell\theta} d\theta \right| \\ &\leq \frac{C}{\max\{1, |\ell|^{k+2}\}}, \end{aligned}$$

for some constant C .

If $g \in L^\infty([0, 2\pi]) \cap C_{pw}^1([0, 2\pi])$, then

$$|c_\ell| \leq \frac{C}{\max\{1, |\ell|\}}.$$

To bound the discrete convolution in (2.7) we need a relationship between the regularity of m and the regularity of u . Variations on results in [8, p. 223], [24, pp. 97, 102] and [28, pp. 53, 56] give us the following.

Theorem 2.5. *Let D be a bounded, open set. If $m \in L^\infty(D)$, then $u \in C^1(D)$. Furthermore, if $m \in C^{k,\alpha}(D)$, $k \geq 0$, then $u \in C^{k+2,\alpha}(D)$.*

We emphasize that D is an arbitrary bounded, open set. Hence, the theorem relates the *local* regularity of u to the *local* regularity of m .

Using these results, we obtain the necessary bounds on (2.7) by means of the following technical lemma, whose proof is contained in Appendix A.2.

Lemma 2.6. Define the region $A = \{(a, \phi) : R_0 \leq a \leq R_1\}$. If $m \in C^{k,\alpha}(A) \cap C_{pw}^{k+2,\alpha}(A)$, $k \geq 0$, then there exists a constant $C > 0$ independent of ℓ and M such that for $\ell = 0, \dots, M$

$$\|K_\ell u^T\|_\infty \leq \frac{1}{M^{k+4}} \frac{4C}{\max\{1, \ell^2\}} \frac{1}{(M+1-\ell)^{k+1}},$$

Similarly, if $m \in L^\infty(A) \cap C_{pw}^{1,\alpha}(A)$, then

$$\|K_\ell u^T\|_\infty \leq \frac{2C}{\max\{1, \ell^2\}} \left\{ \frac{1}{3M^3} + \frac{1}{M^2} \frac{1}{(M+1-\ell)} \right\}.$$

This leads us to a bound on the approximated Fourier modes.

Theorem 2.7. If $m \in L^\infty(A) \cap C_{pw}^{1,\alpha}(A)$, then

$$\varepsilon_M = \|u^M - v^M\| \leq B \|K^M u^T\| = \mathcal{O}\left(\frac{1}{M^3}\right).$$

If $m \in C^{0,\alpha}(A) \cap C_{pw}^{2,\alpha}(A)$, then

$$\varepsilon_M = \mathcal{O}\left(\frac{1}{M^5}\right).$$

If $m \in C^{k,\alpha}(A) \cap C_{pw}^{k+2,\alpha}(A)$ for $k \geq 1$, then

$$\varepsilon_M = \mathcal{O}\left(\frac{1}{M^{k+6}}\right).$$

Proof. We use Lemma 2.6 to bound $\|K^M u^T\| \leq \sum_{\ell=-M}^M \|K_\ell u^T\|_\infty$. For $\ell = 0$ and $m \in C^{k,\alpha}(A) \cap C_{pw}^{k+2,\alpha}(A)$, we obtain

$$\|K_0 u^T\|_\infty = \mathcal{O}\left(\frac{1}{M^{2k+5}}\right).$$

For $m \in L^\infty(A) \cap C_{pw}^{1,\alpha}(A)$, in turn, we have

$$\|K_0 u^T\|_\infty = \mathcal{O}\left(\frac{1}{M^3}\right).$$

For the remaining part of the sum, it is sufficient to bound sums of the following form

$$\sum_{\ell=1}^M \frac{1}{\ell^2} \frac{1}{(M+1-\ell)^p},$$

for $p = 1, 2, \dots$. First consider $p \geq 2$.

$$\begin{aligned}
\sum_{\ell=1}^M \frac{1}{\ell^2} \frac{1}{(M+1-\ell)^p} &\leq \sum_{\ell=1}^M \frac{1}{\ell^2} \frac{1}{(M+1-\ell)^2} \\
&\leq 2 \sum_{\ell=1}^{\lfloor \frac{M}{2} \rfloor} \frac{1}{\ell^2} \frac{1}{(M+1-\ell)^2} \\
&= \mathcal{O}\left(\frac{1}{M^2}\right).
\end{aligned}$$

For $p = 1$, we write the summed quantity as partial fractions.

$$\begin{aligned}
\sum_{\ell=1}^M \frac{1}{\ell^2} \frac{1}{(M+1-\ell)} &= \frac{1}{(M+1)^2} \sum_{\ell=1}^M \frac{1}{\ell} + \frac{1}{M+1} \sum_{\ell=1}^M \frac{1}{\ell^2} + \frac{1}{(M+1)^2} \sum_{\ell=1}^M \frac{1}{(M+1-\ell)} \\
&= \mathcal{O}\left(\frac{1}{M}\right).
\end{aligned}$$

Combining these results, we arrive at the claims of the theorem. $\square$

Remark 2.8. Of course, there are many other conditions on m that could be proposed for which the corresponding convergence rates could be determined. For instance, one might remove the requirement of Hölder continuity. In every case, the convergence rates are directly determined by the rate of decay of the Fourier coefficients of m and u . We do not attempt to provide a comprehensive listing of all possible regularity conditions and their corresponding convergence rates.

Remark 2.9. We have taken great care in the proof to obtain tight bounds; the resulting convergence rates depend on k in a particularly interesting way. Proceeding with less care, one might have expected a simpler dependence on k as follows.

$$\begin{aligned}
\varepsilon_M &\leq B \|K^M\|_{\infty} \|u^T\|_{\infty} \\
&= \mathcal{O}\left(\sum_{\ell>M} \frac{1}{\ell^{k+4}}\right) = \mathcal{O}\left(\frac{1}{M^{k+3}}\right).
\end{aligned}$$

This bound predicts second-order convergence for $m \in L^{\infty}(A) \cap C_{pw}^{1,\alpha}(A)$, third-order convergence for $m \in C^{0,\alpha}(A) \cap C_{pw}^{2,\alpha}(A)$, and fourth-order convergence for $m \in C^{1,\alpha}(A) \cap C_{pw}^{3,\alpha}(A)$. However, as we have proven, these simple bounds are not tight. In fact, the method exhibits *third-order* convergence for $m \in L^{\infty}(A) \cap C_{pw}^{1,\alpha}(A)$, *fifth-order* convergence for

$m \in C^{0,\alpha}(A) \cap C_{pw}^{2,\alpha}(A)$ and *seventh-order* convergence for $m \in C^{1,\alpha}(A) \cap C_{pw}^{3,\alpha}(A)$. This rather interesting and unexpected convergence behavior can be observed in the far field convergence rates of the computational examples in Section 5.1.1.

2.2.2 Total Error in the Interior and Exterior Fields

To this point, we have only computed the convergence rate in the approximated modes, i.e., the first M modes. Given these convergence rates, we can now easily estimate the total error. We make a distinction here between two types of error: the *interior field error* (the error on the domain of integration $A = \{(a, \phi) : R_0 \leq a \leq R_1\}$) and the *exterior field error* (the error outside of A). The interior field error is simply the difference between the true solution $u(x)$ and the solution $v^M(x)$ of (2.3) on A . Therefore, for $x \in A$

$$\begin{aligned} |u(x) - v^M(x)| &\leq |u^M(x) - v^M(x)| + |u^T(x)| \\ &\leq \varepsilon_M + \tau_M(|x|), \end{aligned}$$

where $\tau_M(|x|)$ is a bound on $|u^T(x)|$.

The important observation here is that τ_M depends on the decay rate of the Fourier modes of u at the radius $|x|$. Hence, even if m is discontinuous on A , if m is smooth as a function of angle on the circle with radius $|x|$ centered at the origin, then the Fourier modes of u at radius $|x|$ decay very rapidly. This result implies that, at a given point x , the method converges at a rate that depends on a combination of the regularity of m at the radius $|x|$ and the regularity of m in all of A . We state this result more precisely in the following corollary to Theorem 2.7.

Corollary 2.10 (Interior field error). *Let $x \in A$. Let S denote the circle with radius $|x|$ centered at the origin. Let $N(S)$ be a neighborhood of S in $\mathbb{R}^2$. If $m \in C^{p,\alpha}(N(S)) \cap C_{pw}^{p+2,\alpha}(N(S))$, then the interior field error is given by*

$$|u(x) - v^M(x)| \leq \varepsilon_M + \tau_M(|x|),$$

where bounds on ε_M are given in Theorem 2.7 and

$$|u^T(x)| \leq \tau_M(|x|) = \mathcal{O}\left(\frac{1}{M^{p+3}}\right).$$

This result holds with $p = -1$ for $m \in L^\infty(N(S)) \cap C_{pw}^{1,\alpha}(N(S))$.

Proof. We need only prove the bound on $u^T(x)$. By Theorem 2.5, $u \in C^{p+2,\alpha}(N(S)) \cap C^{p+4,\alpha}(N(S))$. Hence, by Lemma 2.4,

$$\begin{aligned} |u^T(x)| &\leq \tau_M(|x|) = \sum_{\ell > M} \frac{C}{\ell^{p+4}} \\ &= \mathcal{O}\left(\frac{1}{M^{p+3}}\right). \end{aligned}$$

The proof for $m \in L^\infty(N(S)) \cap C_{pw}^{1,\alpha}(N(S))$ is similar. $\square$

Remark 2.11. Although the approximate Fourier modes v^M converge rapidly to u^M , the decay of u^T dominate the maximum interior field error. For example, in the case of $m \in L^\infty(A) \cap C_{pw}^{1,\alpha}(A)$, third-order convergence of v^M to u^M will be dominated by second-order decay of u^T in the maximum error. At the same time, if the interior field error is evaluated at a radius for which m is smooth, we will observe the more rapid convergence rate predicted by Theorem 2.7.

Before we can discuss convergence rates in the exterior field, we must describe how we extend our approximate solution v^M , which we have computed only on the interior of A , to the exterior field. Since the integration in (1.6) is performed only over the support of m , one can easily see that given the exact solution u on the boundary of A , the solution in the rest of $\mathbb{R}^2$ can be computed simply by an appropriate scaling of the Fourier modes of u . More precisely, define $A = \{(r, \theta) : R_0 \leq r \leq R_1\}$. Then,

$$u_\ell^s(r) = \begin{cases} \frac{J_\ell(\kappa r)}{J_\ell(\kappa R_0)} u_\ell^s(R_0), & \text{if } 0 \leq r < R_0 \\ \frac{H_\ell^1(\kappa r)}{H_\ell^1(\kappa R_1)} u_\ell^s(R_1), & \text{if } r > R_1. \end{cases} \quad (2.8)$$

(Note that this result can also be obtained directly from the differential equation by means of separation of variables.) Similarly, to extend our approximate solution v^M to the exterior of A , we simply scale its Fourier modes in the same way.

Corollary 2.12 (Exterior field error). *Given $x \notin A$, extend the approximate solution*

v^M to the exterior of A as prescribed above. Then, the exterior field error is given by

$$|u(x) - v^M(x)| = \mathcal{O}(\varepsilon_M),$$

where $\varepsilon_M = \|u^M - v^M\|_\infty$ has bounds given by Theorem 2.7.

Proof. Denote the scaling factors for the given radius $r = |x|$ by $\beta_\ell(r)$. Assume that $r > R_1$.

The proof for $0 \leq r < R_0$ is similar. We have

$$\begin{aligned} |u(x) - v^M(x)| &\leq \sum_{\ell=-M}^M |\beta_\ell(r)| \left| u_\ell^{s,M}(r) - v_\ell^{s,M}(r) \right| + |u^T(x)| \\ &\leq \varepsilon_M \sum_{\ell=-M}^M |\beta_\ell(r)| + |u^T(x)|. \end{aligned}$$

As before, let S denote the circle of radius r about the origin. Since $r = |x| > R_1$, there exists a neighborhood $N(S)$ of S such that $m|_{N(S)} = 0$. Therefore, $u \in C^\infty(N(S))$ and $|u^T(x)| \leq \frac{C}{M^p}$ for any integer $p > 0$. This implies that $|u^T(x)|$ is always dominated by ε_M .

To complete the proof, we use the asymptotic expressions for J_ℓ and Y_ℓ as found in [2, p. 365], i.e., for fixed z and as $\ell \rightarrow \infty$ through positive real values,

$$\begin{aligned} J_\ell(z) &\sim \frac{1}{\sqrt{2\pi\ell}} \left(\frac{ez}{2\ell} \right)^\ell \\ Y_\ell(z) &\sim -\sqrt{\frac{2}{\pi\ell}} \left(\frac{ez}{2\ell} \right)^{-\ell}. \end{aligned}$$

Therefore,

$$\begin{aligned} |\beta_\ell(r)|^2 &= \left| \frac{Y_\ell(\kappa r)}{Y_\ell(\kappa R_1)} \right|^2 \frac{1 + \left| \frac{J_\ell(\kappa r)}{Y_\ell(\kappa r)} \right|^2}{1 + \left| \frac{J_\ell(\kappa R_1)}{Y_\ell(\kappa R_1)} \right|^2} \\ &\sim \left(\frac{R_1}{r} \right)^{2\ell}. \end{aligned}$$

This implies that $|\beta_\ell(r)|$ is summable. Hence, $\sum_{\ell=-M}^M |\beta_\ell(r)|$ is bounded for all integers $M > 0$. We conclude that

$$|u(x) - v^M(x)| = \mathcal{O}(\varepsilon_M).$$

□

Remark 2.13. Note that $u \in C^\infty$ on the exterior of A and u may be much less regular on the interior of A (in general, $u \in C^1$ for a discontinuous scatterer). Hence, the decay of u^T on the exterior of A is superalgebraic, whereas u^T may decay as slowly as $\mathcal{O}(M^{-2})$ on the interior of A . This fact is responsible for the interesting result that the method converges more rapidly on the exterior of A than on the interior (where u^T may dominate ε_M).

This is particularly relevant in the evaluation of radar cross sections, an important measure in many applications. The evaluation of radar cross sections requires the computation of the *far field*. Although Corollary 2.12 does not directly address the error in the far field, we obtain an approximate far field by a scaling of the Fourier modes of v^M just as in the computation of the exterior field. As in [10, p. 6], we define the far field, u_∞ , by the asymptotic representation of the scattered field as $r \rightarrow \infty$, i.e.,

$$u^s(r, \phi) = e^{i(\kappa r - \frac{\pi}{4})} \sqrt{\frac{2}{\pi \kappa r}} [u_\infty(\phi) + \mathcal{O}(r^{-1})].$$

From (2.8) and the asymptotic expression for $H_\ell^1(z)$ for fixed ℓ as $z \rightarrow \infty$ [2, p. 364], we obtain the Fourier modes of u_∞ by a simple scaling of the Fourier modes of u^s .

$$(u_\infty)_\ell = \frac{u_\ell^s(R_1)}{i^\ell H_\ell^1(\kappa R_1)}.$$

If we define the approximate far field v_∞ by scaling the Fourier modes of $v^{s,M}$ in the same way, we can prove that

$$\|u_\infty - v_\infty\| = \mathcal{O}(\varepsilon_M).$$

The proof of this fact is nearly identical to the proof of Corollary 2.12.

The predicted convergence rates in both the interior field and the far field are verified through several computational examples in Section 5.1.1.

2.3 Computation of the Angular Integral

Thus we have shown that the method achieves high-order convergence even in the case of discontinuous scatterers. However, to this point, we have not discussed any methods for computing the required angular and radial integrals for each mode of the solution (2.6). Since this chapter primarily addresses the theoretical aspects of the method, we leave a

discussion of a particular efficient, high-order radial integrator to the next chapter. On the other hand, the required Fourier coefficients of $m(r, \theta)v^M(r, \theta)$ can be computed efficiently and *exactly* (except for round-off error). Furthermore, the approach taken in computing the angular integrals was the primary source of controversy surrounding the method. Therefore, we briefly discuss the angular integration here before moving on to a discussion of the numerical implementation in the next chapter.

The required angular integrals are given by

$$I_\ell(r) = \int_0^{2\pi} m(r, \theta)v^M(r, \theta)e^{-i\ell\theta}d\theta, \quad (2.9)$$

where v^M solves the approximate integral equation (2.3). We can express this integral in terms of the Fourier coefficients of m and v , i.e.,

$$\begin{aligned} I_\ell(r) &= \int_0^{2\pi} \left(\sum_{j=-\infty}^{\infty} m_j(r)e^{ij\theta} \right) \left(\sum_{k=-M}^M v_k(r)e^{ik\theta} \right) e^{-i\ell\theta} \\ &= 2\pi \sum_{k=-M}^M m_{\ell-k}(r)v_k(r), \end{aligned} \quad (2.10)$$

where $\ell = -M, \dots, M$.

Hence, we obtain a finite discrete convolution of Fourier coefficients of m and v at each radius. Since $|\ell| \leq M$ and $|k| \leq M$, we have $|\ell - k| \leq 2M$. Thus, given the Fourier coefficients $m_\ell(r)$ for $|\ell| \leq 2M$, we can compute the required angular integrals *exactly*. Furthermore, as is well known, such discrete convolutions may be evaluated (with no approximation) with the help of FFTs [45, pp. 531–537] yielding a complexity of $\mathcal{O}(M \log M)$ at each radial point. (As we will demonstrate in our numerical examples, Chapter 5, these Fourier coefficients can be computed quite easily for a wide range of scatterers.)

This method of computing the angular integrals has an interesting implication. Since the computation involves only modes m_ℓ , $|\ell| \leq 2M$, replacing m with m^{2M} in the integral equation yields no additional error, i.e.,

$$I_\ell(r) = \int_0^{2\pi} m^{2M}(r, \theta)v^M(r, \theta)e^{-i\ell\theta}d\theta \quad (2.11)$$

Hence, in a sense, the truncation of the Fourier series of the integral operator *implies* an associated truncation of the Fourier series of the scatterer—as a result of the band-

limited nature of the solution v^M . Thus, surprisingly, the low-order approximation of a discontinuous scatterer at each radius by its truncated Fourier series yields *no more error* than our original, high-order truncation of the Fourier series of K . This illustrates the interesting *cancellation of errors* that underlies the power of this approach.

Note that this discrete convolution method of computing $I_\ell(r)$ for $\ell = -M, \dots, M$ is equivalent to trapezoidal rule integration of (2.11) with a sufficiently large number of integration points N_θ . More precisely, it is not difficult to see that the trapezoidal rule with N_θ points integrates Fourier modes $e^{ik\theta}$ for $|k| < N_\theta$ exactly. Since the largest mode in the integrand of (2.11) is $2M + M + M = 4M$, if we choose $N_\theta > 4M$, the trapezoidal rule computes (2.11) exactly (assuming exact arithmetic) and the use of FFTs yields a complexity of $\mathcal{O}(M \log M)$. Algorithmically, this is entirely equivalent to computing the discrete convolution (2.10) using FFTs.

Chapter 3

A Fast, High-Order Method in Two Dimensions: Numerical Implementation

In this chapter, we present several significant improvements to the original numerical implementation of the two-dimensional method introduced in [13], as described in the previous chapter. The numerical solution of the associated approximate integral equation (2.3) consists of two main parts: efficient, high-order numerical quadrature rules and an efficient linear solver.

The numerical evaluation of the integrals in (2.6) requires both angular and radial integration. In Section 2.3, we described a method for computing the required angular integrals (2.9), *exactly* by means of Fourier smoothing of the scatterer. Furthermore, this method requires only $\mathcal{O}(M \log M)$ operations for each radial point. (For smooth scatterers, however, the angular integrals can instead be computed efficiently and with high-order accuracy by direct application of the trapezoidal rule without replacing m with m^{2M} (see Section 3.1.1).)

On the other hand, to compute the radial integrals $(K_\ell v^M)(a)$ (2.2), we present an improved scheme based on Chebyshev polynomials approximation. More precisely, we approximate $I_\ell(r)$ in (2.9) with Chebyshev polynomials on each of several subintervals (see Section 3.1.2). When the integration domain contains the origin, we resolve the logarithmic singularity in the Hankel function by integrating by parts (see Section 3.1.2.1). Then by means of appropriate pre-computed integral moments (see Section 3.1.2.2), we obtain a high-order approximation of the required integrals. Computation of these integrals requires only $\mathcal{O}(N)$ operations.

This high-order discretization of the required integrals in (2.6) results in a linear system with unknowns $v_\ell(a_j)$ for $\ell = -M, \dots, M$ and $j = 0, \dots, N_r$, where a_n are the radial points in the discretization. We use the Generalized Minimal Residual (GMRES) iterative method to solve this linear system (see Section 3.2.1). The number of GMRES iterations required to achieve a given residual tolerance increases rapidly with increasing problem size as measured by interior wavelengths (proportional to $\kappa n(R_1 - R_0)$ for a constant refractive index). We describe a preconditioner (see Section 3.2.2) that reduces this required number of iterations for a variety of scattering configurations. We precondition the problem with an integral equation corresponding to a piecewise constant, radially layered approximating scatterer. By means of an equivalent partial differential equation, it can be shown that such an approximate integral equation admits an inverse, which can be computed in closed form. Application of this inverse requires only $\mathcal{O}(N)$ operations.

3.1 High-Order Numerical Integration

3.1.1 Angular Integration

We first consider the angular integration (2.9). As discussed in Section 2.3, $I_\ell(r)$ can be computed exactly even for discontinuous scatterers by replacing m by m^{2M} and choosing a sufficient number of trapezoidal rule integration points. Direct application of the trapezoidal rule in the case of a discontinuous inhomogeneity (without replacing m by m^{2M}) would yield only first-order convergence.

For smooth inhomogeneities m , however, we can obtain high-order convergence simply through direct application of the trapezoidal rule. This follows from the fact that, for smooth scatterers, the integrand $m(r, \theta)v^M(r, \theta)$ is a high-order approximation to the smooth and periodic function $m(r, \theta)u(r, \theta)$ and the trapezoidal rule yields high-order accuracy for smooth and periodic integrands. Hence, although $I_\ell(r)$ may always be computed exactly by replacing m by m^{2M} , for sufficiently smooth inhomogeneities, the direct application of the trapezoidal rule is somewhat simpler (since the Fourier coefficients $m_\ell(r)$ are not required) and produces nearly the same accuracy.

In this case, given the values $v_\ell(r)$, we first compute $v^M(r, \theta_j)$ for $\theta_j = 2\pi j/N_\theta$, where

N_θ is the number of trapezoidal rule integration points, i.e.,

$$v^M(r, \theta_j) = \sum_{\ell=-M}^M v_\ell(r, \theta_j) e^{2\pi i j \ell / N_\theta}$$

for $j = 0, \dots, N_\theta - 1$. This sum is computed in $\mathcal{O}(N_\theta \log N_\theta)$ operations for each radial point r by means of an FFT, where N_θ must be chosen such that $N_\theta > 2M$. Using the periodicity in θ of m and v^M , the trapezoidal rule gives

$$I_\ell(r) \approx I_\ell^{N_\theta}(r) = \sum_{j=0}^{N_\theta-1} m(r, \theta_j) v^M(r, \theta_j) e^{-2\pi i j \ell / N_\theta}, \quad (3.1)$$

which we compute for $\ell = -M, \dots, M$ by means of an FFT. Hence, given N_r radial points, this algorithm computes the integrals $I_\ell(r)$ with high-order accuracy in a total of $\mathcal{O}(N \log N)$ operations where $N = \mathcal{O}(N_r M)$.

It is interesting to note the relationship between this direct trapezoidal rule approach and the discrete convolution approach described in Section 2.3. When m is replaced by m^{2M} and N_θ is chosen such that $N_\theta > 4M$, these two approaches yield identical results. In other words, the trapezoidal rule approach described above is *identical* to the FFT-based method for computing a discrete convolution (when m^{2M} is used and $N_\theta > 4M$) [45, pp. 531–537]. Hence, these algorithms require only a single implementation; the approach that is actually used is determined by the choice of scatterer, either m or m^{2M} , and the related choice of N_θ .

3.1.2 Radial Integration

The radial integration requires the computation of two functions

$$\begin{aligned} F_\ell(a) &= \int_{R_0}^a Y_\ell(\kappa a) J_\ell(\kappa r) I_\ell(r) r dr \\ G_\ell(a) &= \int_a^{R_1} J_\ell(\kappa a) Y_\ell(\kappa r) I_\ell(r) r dr, \end{aligned}$$

for $R_0 \leq a \leq R_1$. In terms of these functions, $(K_\ell v^M)(a)$ is given by

$$(K_\ell v^M)(a) = -i \frac{\kappa^2}{4} \left[\int_{R_0}^a H_\ell^1(\kappa a) J_\ell(\kappa r) I_\ell(r) r dr + \int_a^{R_1} J_\ell(\kappa a) H_\ell^1(\kappa r) I_\ell(r) r dr \right] \quad (3.2)$$

$$= \frac{\kappa^2}{4} \left[F_\ell(a) + G_\ell(a) - i \frac{J_\ell(\kappa a)}{Y_\ell(\kappa R_1)} F_\ell(R_1) \right] \quad (3.3)$$

High-order accuracy in evaluating $F_\ell(a)$ and $G_\ell(a)$ requires that we treat discontinuities and/or singularities in $I_\ell(r)$ appropriately. The singularities in $I_\ell(r)$ are due to the singularities in the Fourier coefficients of the scatterer, $m_\ell(r)$. Thus, unless $m(r, \theta)$ is smooth within the annulus $R_0 \leq r \leq R_1$, $m_\ell(r)$ will exhibit singularities.

Following the discussion in [13], we assume that $m(r, \theta)$ is piecewise smooth, i.e., we assume that there are a finite number of disjoint sets $\mathcal{D}_i$ with piecewise smooth boundaries such that $m(r, \theta)$ is smooth on each $\mathcal{D}_i$. Singularities in $I_\ell(r)$ can occur at values of r that correspond either to non-smooth points in the boundary of a set $\mathcal{D}_i$ or to a point of tangency between the circle C_r of radius r and a set $\mathcal{D}_i$. Non-smooth points in the boundary of a set $\mathcal{D}_i$ lead to corner-type singularities. Points of tangency lead to singularities of type $(r - r_0)^\alpha$ for $0 \leq \alpha \leq 1$. For example, most of the examples of Section 5.1 have square-root type singularities, $\sqrt{r - r_0}$. Discontinuities in $I_\ell(r)$ may also arise, for example, when the circle C_r intersects a set $\mathcal{D}_i$ along a finite segment of its boundary.

Discontinuities and corner-type singularities are handled simply by subdividing the integration domain $[R_0, R_1]$ at the singularity points into two or more intervals. Singularities of the type $(r - r_0)^\alpha$ can be resolved by changing variables in the radial dimension. For example, as suggested in [13], we resolve square-root singularities occurring at the endpoints of the interval $[a, b]$ with the following change of variable

$$\cos(\phi) = \sqrt{\frac{r^2 - a^2}{b^2 - a^2}}. \quad (3.4)$$

Of course, in order to resolve such singularities, one must first deduce the type of singularity. However, this is not difficult to accomplish in many situations.

Thus, by breaking the integration domain $[R_0, R_1]$ into multiple intervals and by appropriately changing variables to resolve remaining $(r - r_0)^\alpha$ -type singularities, we are left with a smooth $I_\ell(r)$ on each interval of integration. Since the extension to this fully general setting is straightforward but also tedious, we describe only how the computation $F_\ell(a)$

and $G_\ell(a)$ to high-order accuracy for a *single* interval on which no change of variables is required. This special case illustrates all of the major aspects of the fully general method without unduly complicating the exposition.

We divide the integration interval $[R_0, R_1]$ into several subintervals of equal length. We then approximate $I_\ell(r)$ by a high-order polynomial expansion on each subinterval

$$I_\ell(r) \approx \sum_{n=0}^{N_c-1} c_{nk} p_n(r), \quad (3.5)$$

where $c_{nk} \in \mathbb{C}$, $k = 1, \dots, N_i$ and N_i is the number of subintervals. Also, $p_n(r)$ is a polynomial of order n . Thus, we thereby obtain a polynomial expansion of order $N_c - 1$.

There are several possible choices for the $p_n(r)$. In [13], Lagrange interpolating polynomials are used. However, for Lagrange interpolating polynomials, the computation of the coefficients c_{nk} is numerically unstable for large N_c . This instability can be mitigated somewhat by using a more stable algorithm, at the cost of computational complexity. In [13], the chosen routine requires $\mathcal{O}(N_c^3)$ operations.

Hence, we use Chebyshev polynomials instead. Chebyshev polynomials, T_n , provide excellent approximation (nearly equal to the minimax polynomial [3, pp. 225–236]) while also allowing stable evaluation of the c_{nk} in $\mathcal{O}(N_c^2)$ operations. For r in the k^{th} subinterval, we have

$$I_\ell(r) \approx \sum_{n=0}^{N_c-1} c_{nk\ell} T_n(x_k(r)), \quad (3.6)$$

where the functions $x_k(r)$ are simply linear maps of the k^{th} subinterval to the interval $[-1, 1]$, which is the standard interval of definition for the Chebyshev polynomials. More precisely, if the k^{th} subinterval is given by $[a, b]$, then, for $r \in [a, b]$, we have

$$x_k(r) = \frac{r-a}{b-a} - \frac{b-r}{b-a}.$$

We will also make use of the inverse of this map, which, for $y \in [-1, 1]$, is given by

$$x_k^{-1}(y) = \frac{1}{2} [(b-a)y + (b+a)].$$

The evaluation of the Chebyshev coefficients for a function defined on the standard

interval $[-1, 1]$ requires the function values at

$$y_j = \cos\left(\frac{\pi(j - 1/2)}{N_c}\right), \quad (3.7)$$

for $j = 1, \dots, N_c$. Hence, we require the value of $I_\ell(r)$ at the corresponding points in each of the subintervals, i.e., at the points

$$a_{jk} = x_k^{-1}(y_j) \quad (3.8)$$

for $j = 1, \dots, N_c$ and $k = 1, \dots, N_i$. Notice that the a_{jk} do not include the endpoints of each subinterval since the y_j do not include the endpoints of $[-1, 1]$. Since we will need these endpoints for parts of our calculation, we include the left and right endpoints of each subinterval as a_{0k} and a_{N_pk} , respectively, where $N_p = N_c + 1$. (Note that by this definition $a_{0k} = a_{N_pk-1}$ for $k > 1$.)

Hence, these a_{jk} are our radial discretization points giving a total of $N_r = \mathcal{O}(N_p N_i)$ radial points. Our other discretization parameter is the number of modes M in the approximate solution v^M . Hence, the total number of unknowns $N = \mathcal{O}(N_r M) = \mathcal{O}(N_p N_i M)$. The number of points per subinterval N_p determines the *order* Chebyshev approximation is held *fixed* while the number of subintervals N_i and the number of modes M is increased to obtain more accuracy. Given the values of $I_\ell(a_{jk})$, one can compute the Chebyshev coefficients $c_{nk\ell}$ in $\mathcal{O}(N_p^2)$ operations per subinterval and per mode giving a total of $\mathcal{O}(N_i N_p^2 M) = \mathcal{O}(N)$ operations for all subintervals. (One can reduce the N_p^2 complexity for each subinterval to $\mathcal{O}(N_p \log N_p)$ by use of FFTs. However, since N_p remains fixed, this does not change the overall complexity and since we typically use a relatively small value of N_p , e.g., $N_p = 9, 17$, we have found that this approach provides little benefit.)

Thus, given these Chebyshev approximations on each subinterval, we must only compute the following integral moments

$$P_{njk\ell} = \int_{a_{0k}}^{a_{jk}} Y_\ell(\kappa a_{jk}) J_\ell(\kappa r) T_n(x_k(r)) r dr \quad (3.9)$$

$$Q_{njk\ell} = \int_{a_{jk}}^{a_{N_pk}} J_\ell(\kappa a_{jk}) Y_\ell(\kappa r) T_n(x_k(r)) r dr. \quad (3.10)$$

Since all of the functions in the integrand are known analytically, these moments are only

computed once at the beginning of each run. The storage of these moments requires $\mathcal{O}(N_c N_p N_i M) = \mathcal{O}(N)$ memory. These integrals are problematic because of the rapid decay of J_ℓ for large ℓ near the origin as well as the logarithmic and polar singularities in Y_ℓ at the origin. In Section 3.1.2.1, we present a method for resolving the logarithmic singularity in Y_ℓ , which is necessary when the integration domain includes the origin, i.e., when $R_0 = 0$. We then discuss our method for computing the moments in Section 3.1.2.2.

Another practical obstacle when computing these moments concerns the rapid decay of J_ℓ and the rapid growth of Y_ℓ near the origin for large values of ℓ . Hence, for large values of ℓ , J_ℓ may underflow and Y_ℓ may overflow while their product remains machine-representable. We overcome this difficulty by computing scaled versions of J_ℓ and Y_ℓ , thus allowing accurate computation of $J_\ell(a)Y_\ell(b)$, $Y_\ell(a)/Y_\ell(b)$, etc. We describe our method for computing these scaled Bessel functions in Appendix B.

Once we have computed these moments, we can compute the values $F_{jkl} = F_\ell(a_{jk})$ and $G_{jkl} = G_\ell(a_{jk})$. We first compute the quantities A_{jkl} and B_{jkl} , which involve integration over a single subinterval.

$$\begin{aligned} A_{jkl} &= \int_{a_{0k}}^{a_{jk}} Y_\ell(\kappa a_{jk}) J_\ell(\kappa r) I_\ell(r) r dr \\ &\approx \sum_{n=0}^{N_c-1} c_{nkl} P_{njk\ell} \\ B_{jkl} &= \int_{a_{jk}}^{a_{Npk}} J_\ell(\kappa a_{jk}) Y_\ell(\kappa r) I_\ell(r) r dr \\ &\approx \sum_{n=0}^{N_c-1} c_{nkl} Q_{njk\ell}. \end{aligned}$$

Hence, given the Chebyshev coefficients, which require $\mathcal{O}(N)$ operations to compute, the computation of A_{jkl} and B_{jkl} requires $\mathcal{O}(N_c N_p N_i M) = \mathcal{O}(N)$ operations. We then sum and scale these values appropriately to compute F_{jkl} and G_{jkl} .

$$\begin{aligned} F_{jkl} &= \int_{a_{01}}^{a_{jk}} Y_\ell(\kappa a_{jk}) J_\ell(\kappa r) I_\ell(r) r dr \\ &\approx A_{jkl} + \alpha_{jkl} F_{Npk-1\ell}, \text{ for } k = 2, \dots, N_i \\ F_{j1\ell} &\approx A_{j1\ell}, \end{aligned}$$

where $\alpha_{jkl} = \frac{Y_\ell(\kappa a_{jk})}{Y_\ell(\kappa a_{N_p k-1})}$.

$$\begin{aligned} G_{jkl} &= \int_{a_{jk}}^{a_{N_p N_i}} J_\ell(\kappa a_{jk}) Y_\ell(\kappa r) I_\ell(r) r dr \\ &\approx B_{jkl} + \beta_{jkl} G_{0j+1\ell}, \text{ for } k = 1, \dots, N_i - 1 \\ G_{jN_i\ell} &\approx B_{jN_i\ell}, \end{aligned}$$

where $\beta_{jkl} = \frac{J_\ell(\kappa a_{jk})}{J_\ell(\kappa a_{0j+1})}$. Clearly, given A_{jkl} and B_{jkl} , the computation of F_{jkl} and G_{jkl} requires $\mathcal{O}(N_p N_i M) = \mathcal{O}(N)$ operations.

Therefore, the total complexity of the radial integration, given the integral moments P_{njkl} and Q_{njkl} , is $\mathcal{O}(N)$. Furthermore, assuming that the integral moments are computed with sufficient accuracy, the accuracy of the F_{jkl} and G_{jkl} is determined by the accuracy of the Chebyshev approximation on each subinterval. Since $I_\ell(r)$ is smooth (or can be made smooth by changing variables and breaking up the integration domain), the Chebyshev approximation is high-order accurate with the order of convergence dependent on the value of N_p , the number of points per subinterval. Thus, we obtain an efficient and high-order accurate method for computing the radial integrals. (See the numerical examples in Section 5.1.2.)

3.1.2.1 Resolution of Logarithmic Singularity

Besides the singularities in $I_\ell(r)$, which we resolve as discussed previously, the Bessel function $Y_\ell(\kappa r)$ exhibits logarithmic and polar singularities at $r = 0$. In this section, we show that the logarithmic singularity can be removed through integration by parts.

This method results as a slight modification of the approaches discussed above: we evaluate F_{jkl} as discussed previously, but we change the computation of G_{jkl} . Note that $Y_\ell(z)$ can be written as

$$Y_\ell(z) = \frac{2}{\pi} \log\left(\frac{z}{2}\right) J_\ell(z) + Y_\ell^{\text{polar}}(z),$$

where $Y_\ell^{\text{polar}}(z) = \mathcal{O}(z^{-\ell})$ as $z \rightarrow 0$. Assume that $R_0 = 0$, since otherwise the logarithmic

singularity is not present. Integrating the logarithmic term by parts gives

$$\begin{aligned}
\int_a^{R_1} \log\left(\frac{\kappa r}{2}\right) J_\ell(\kappa r) I_\ell(r) r dr &= \log\left(\frac{\kappa r}{2}\right) \int_0^r J_\ell(\kappa \rho) I_\ell(\rho) \rho d\rho \Big|_a^{R_1} \\
&\quad - \int_a^{R_1} dr \frac{1}{r} \int_0^r J_\ell(\kappa \rho) I_\ell(\rho) \rho d\rho \\
&= \log\left(\frac{\kappa R_1}{2}\right) S_\ell(R_1) - \log\left(\frac{\kappa a}{2}\right) S_\ell(a) - \int_a^{R_1} r^{-1} S_\ell(r) dr,
\end{aligned}$$

where $S_\ell(a) = \int_0^a J_\ell(\kappa r) I_\ell(r) dr = \frac{F_\ell(a)}{Y_\ell(\kappa a)}$.

$S_\ell(r)$ vanishes to first order at $r = 0$ and thus, $r^{-1}S_\ell(r)$ is smooth near the origin. Furthermore, since we have computed F_{jkl} , i.e., the values of $F_\ell(a)$ at Chebyshev points in each subinterval, the values $S_\ell(a_{jk})$ as well as the Chebyshev coefficients of $r^{-1}S_\ell(r)$ on each subinterval are easily computed. Therefore, in $\mathcal{O}(N)$ operations, we compute the Chebyshev coefficients d_{nkl} such that

$$r^{-1}S_\ell(r) \approx \sum_{n=0}^{N_c-1} d_{nkl} T_n(x_k(r)).$$

The Chebyshev coefficients d_{nkl} of the integrand are easily mapped into the Chebyshev coefficients D_{nkl} of its indefinite integral [45, pp. 189, 190] as

$$D_{nkl} = \frac{d_{n-1kl} - d_{n+1kl}}{2n}$$

for $n \geq 1$. The constant of integration D_{0kl} is arbitrary. Clearly, this mapping requires only $\mathcal{O}(N)$ operations. Furthermore, since the integral $\int_a^{R_1} r^{-1}S_\ell(r) dr$ is smooth, the Chebyshev approximation on each subinterval with coefficients D_{nkl} gives a high-order representation of this integral.

To complete the computation of $G_\ell(a_{jk})$, we define

$$G_{jkl}^{\log} = \int_{a_{jk}}^{R_1} J_\ell(\kappa a_{jk}) \frac{2}{\pi} \log\left(\frac{\kappa r}{2}\right) J_\ell(\kappa r) I_\ell(r) r dr$$

and

$$G_{jkl}^{\text{polar}} = \int_{a_{jk}}^{R_1} J_\ell(\kappa a_{jk}) Y_\ell^{\text{polar}}(\kappa r) I_\ell(r) r dr.$$

Therefore, $G_{jkl} = G_{jkl}^{\log} + G_{jkl}^{\text{polar}}$. We compute G_{jkl}^{polar} exactly as we computed G_{jkl} in the

previous section with $Q_{njk\ell}$ replaced by

$$Q_{njk\ell}^{polar} = \int_{a_{jk}}^{a_{N_p k}} J_\ell(\kappa a_{jk}) Y_\ell^{polar}(\kappa r) I_\ell(r) r dr. \quad (3.11)$$

We compute $G_{jk\ell}^{\log}$ by first defining, for $k = 1, \dots, N_i - 1$,

$$\begin{aligned} R_{jk\ell} &= \int_{a_{jk}}^{R_1} r^{-1} S_\ell(r) dr \\ &= R_{0k+1\ell} + \int_{a_{jk}}^{a_{N_p k}} r^{-1} S_\ell(r) r dr \\ &\approx R_{0k+1\ell} + \sum_{n=0}^{N_c-1} D_{nk\ell} [T_n(x_k(a_{N_p k})) - T_n(x_k(a_{jk}))], \end{aligned}$$

and, for $k = N_i$, $R_{jN_i\ell} = \sum_{n=0}^{N_c-1} D_{nk\ell} [T_n(x_k(a_{N_p N_i})) - T_n(x_k(a_{jk}))]$. Given the coefficients $D_{nk\ell}$, computation of $R_{jk\ell}$ for $j = 0, \dots, N_p$, $k = 1, \dots, N_i$ and $\ell = -M, \dots, M$ requires $\mathcal{O}(N_c N_p N_i M) = \mathcal{O}(N)$ operations to compute. Thus, we have

$$G_{jk\ell}^{\log} \approx \frac{2}{\pi} J_\ell(\kappa a_{jk}) \left[\log \left(\frac{\kappa R_1}{2} \right) S_{N_p N_i \ell} - \log \left(\frac{\kappa a_{jk}}{2} \right) S_{jk\ell} - R_{jk\ell} \right],$$

where $S_{jk\ell} = \frac{F_{jk\ell}}{Y_\ell(a_{jk})}$. Clearly, the computation of $G_{jk\ell}^{\log}$ requires only $\mathcal{O}(N)$ operations.

Therefore, given accurate values of the integral moments $P_{njk\ell}$ and $Q_{njk\ell}^{polar}$, this approach yields high-order accurate values of the required radial integrals in $\mathcal{O}(N)$ operations. We wish to emphasize that, as stated above, we need only resolve the logarithmic singularity in $Y_\ell(\kappa r)$ when the integration domain includes the origin, i.e., when $R_0 = 0$. It is straightforward to generalize this approach to the case in which we resolve singularities in $I_\ell(r)$ by breaking up the integration domain and/or changing variables.

3.1.2.2 Evaluation of Integral Moments

To complete the method, we need only find an efficient and accurate method for computing the integral moments $P_{njk\ell}$ and $Q_{njk\ell}$. As noted in the previous section, if the integration domain contains the origin, we can resolve the logarithmic singularity in Y_ℓ independently of the moment computation. In such a case, we compute $Q_{njk\ell}^{polar}$ (see equation (3.11)) instead of $Q_{njk\ell}$. However, even when the integration domain does not contain the origin, the rapid growth of the Y_ℓ and the rapid decay of the J_ℓ makes the development of accurate

quadrature schemes for computing the moments challenging. In particular, the Y_ℓ grow like $r^{-\ell}$ and the J_ℓ decay like r^ℓ as either $\ell \rightarrow \infty$ or as $r \rightarrow 0$. Since quadrature rules, for the most part, depend on accurate polynomial interpolation, accurate integration of these functions would seem to require a large number of integration points.

In this particular case, however, accurate values for these moments are actually not too difficult to obtain. The key insight is that a small *absolute* error, as opposed to a small *relative* error, is required in the computation of the moments. Since we desire a small relative accuracy in the *field* v^M , we require only a small absolute error in the Fourier coefficients of the field v_ℓ —small values of v_ℓ , regardless of their relative error, contribute very little to the value of the field v^M .

Our goal, therefore, is to show that, given a maximum absolute error $\varepsilon > 0$ and an initial radial discretization, we can choose the number of Gaussian quadrature points N_g such that the absolute error in the moment integrals is less than ε for all ℓ and for all subsequent refinements of the radial discretization. Near the origin or for large ℓ , the asymptotic behavior of J_ℓ and Y_ℓ for $\ell \geq 1$ are given by [2, p. 360]

$$J_\ell(z) \sim \frac{1}{\ell!} \left(\frac{z}{2}\right)^\ell$$

and

$$Y_\ell(z) \sim -\frac{(\ell-1)!}{\pi} \left(\frac{z}{2}\right)^{-\ell}.$$

Hence, for $\ell > 2$, the moment integrals,

$$P_\ell(a, b) = Y_\ell(\kappa b) \int_a^b J_\ell(\kappa r) T_n(x_k(r)) r dr \quad (3.12)$$

and

$$Q_\ell(a, b) = J_\ell(\kappa a) \int_a^b Y_\ell(\kappa r) T_n(x_k(r)) r dr \quad (3.13)$$

have roughly the same properties as

$$-\frac{1}{\pi \ell b^\ell} \int_a^b r^{\ell+1} dr = -\frac{a^2}{\pi \ell (\ell+2)} \left[\left(\frac{b}{a}\right)^2 - \left(\frac{a}{b}\right)^\ell \right] \quad (3.14)$$

and

$$\frac{a^\ell}{\pi \ell} \int_a^b r^{-\ell+1} dr = -\frac{b^2}{\pi \ell (\ell-2)} \left[\left(\frac{a}{b}\right)^2 - \left(\frac{a}{b}\right)^\ell \right], \quad (3.15)$$

respectively. Thus, by considering the relatively simple quantities (3.14) and (3.15), we gain insight into the behavior of integration rules for evaluating (3.12) and (3.13). We will concentrate on integration rules for (3.13) since it is the more difficult of the two moment integrals because of the singularity in the integrand at $r = 0$. In addition, we restrict our attention to the case $a > 0$, since for $a = 0$, $Q_\ell(a, b) = 0$ for $\ell \geq 1$. (For $\ell = 0$, the only singularity in the integrand is the logarithmic singularity, which is resolved by integrating by parts as described in Section 3.1.2.1.)

Gauss-Legendre integration of a function f defined on the interval $[a, b]$ is given by [3, pp. 276–279]

$$\begin{aligned} \int_a^b f(t) dt &= \left(\frac{b-a}{2} \right) \int_{-1}^1 f \left(\frac{b+a}{2} + \frac{b-a}{2} x \right) dx \\ &\approx \left(\frac{b-a}{2} \right) \sum_{j=1}^{N_g} w_j f \left(\frac{b+a}{2} + \frac{b-a}{2} x_j \right), \end{aligned} \quad (3.16)$$

where the points $-1 < x_j < 1$ are zeroes of the degree N_g Legendre polynomial on $[-1, 1]$ and $w_j > 0$ are the corresponding weights. We first consider the decay of the sum (3.16) itself for (3.15). We have

$$\frac{a^\ell}{\ell} \int_a^b r^{-\ell+1} dr \approx \frac{b-a}{2\ell} \sum_{j=1}^{N_g} w_j r_j \left(\frac{a}{r_j} \right)^\ell \leq \frac{b(b-a)}{\ell}, \quad (3.17)$$

where $r_j = \frac{(b+a)}{2} + \frac{(b-a)}{2} x_j$. Here we have also used the fact that $w_j > 0$ and that the $\sum_{j=1}^{N_g} w_j = \int_{-1}^1 dx = 2$ since Gauss-Legendre quadrature integrates polynomials of degree less than $2N_g$ exactly. The sum (3.17) decays with increasing ℓ *independently of* N_g . Note that (3.15) *itself* also decays with ℓ as $\mathcal{O}(\ell^{-2})$. Thus, we conclude that, given an initial discretization, there is a positive integer L such that both (3.15) and (3.17) are smaller than $\varepsilon/2$ for all $\ell > L$, *independently of* N_g . Therefore, the absolute error for $\ell > L$ is smaller than ε . Crucially important in this result is the fact that the moments decay with ℓ . Thus, for sufficiently large ℓ , although the computed moments may have over 100% relative error, their absolute error remains smaller than ε .

We now consider the effect of discretization refinement. Refinement of the radial discretization is accomplished by increasing the number of subintervals N_i . (On each of these subintervals, there is a fixed number of Chebyshev points N_c .) Hence, for two adjacent

discretization points $a = R_0 + \alpha$ and $b = R_0 + \beta$, increasing N_i by some factor γ decreases both α and β by the factor γ . For example, doubling N_i halves both α and β . Similarly, $b - a$ also decreases by the factor γ when N_i increases by the factor γ . It follows from (3.17) that the absolute error for $\ell > L$ decays as we refine the radial discretization. Therefore, our choice of L is not only independent of N_g , but is also independent of subsequent refinements of the discretization.

Now we must consider the absolute error for $\ell \leq L$. The absolute error in the integral of a function f as computed with Gaussian quadrature is bounded by [3, pp. 276–279]

$$E_{N_g}(f) \leq \frac{\pi}{(2N_g)!} \left(\frac{b-a}{2} \right) \left(\frac{b-a}{4} \right)^{2N_g} \max_{a \leq t \leq b} |f^{(2N_g)}(t)|,$$

where the bound holds asymptotically as $N_g \rightarrow \infty$. One can then show that the absolute error $E_{N_g}^Q$ in computing (3.15) is bounded by

$$\begin{aligned} E_{N_g}^Q &= \frac{a^\ell}{\pi \ell} E_{N_g}(r^{-\ell+1}) \\ &\leq \frac{a(b-a)}{2} \left(\frac{b-a}{4a} \right)^{2N_g} \frac{(\ell-2+2N_g)!}{\ell(\ell-2)!(2N_g)!} \\ &\leq \frac{a(b-a)}{2} \left(\frac{b-a}{4a} \right)^{2N_g} \frac{(L-2+2N_g)!}{L(L-2)!(2N_g)!}, \end{aligned}$$

where we have used the fact that $\ell \leq L$. Hence, if we require that $(b-a)/4a < 1$, then $[(b-a)/4a]^{2N_g}$ exhibits exponential decay while $(L-2+2N_g)!/(2N_g)!$ exhibits polynomial growth as N_g increases. Therefore, for N_g sufficiently large $E_{N_g}^Q < \varepsilon$ for $\ell \leq L$. Clearly, $E_{N_g}^Q$ also decays with subsequent discretization refinements. Hence, our choice of N_g depends only on the value of L and the initial discretization.

We conclude that given an initial discretization satisfying $(b-a)/4a < 1$ for all points a, b , we can choose N_g sufficiently large such that the absolute error is smaller than ε for all ℓ and for all subsequent refinements of the discretization. Similar arguments obtain the same result for the absolute error in (3.14). Note that we have only proven this fact in the asymptotic regime of large ℓ or, equivalently, small a and b . However, outside of this asymptotic regime, the Bessel functions have much milder behavior, simply oscillating with wavenumber κ . These oscillations must be resolved by the radial discretization anyway to obtain even minimal accuracy in solving the integral equation. Hence, computing the inte-

gral moments in the oscillatory regime presents no significant difficulties. We do not present any theoretical estimates of the required values of N_g for various initial discretizations and error tolerances ε . However, the examples of Section 5.1 show that the value $N_g = 8$ suffices to achieve machine precision accuracy in our solutions.

3.2 Solution of the Linear System

3.2.1 Linear Solver

The previous sections describe our efficient, high-order method for computing the integral operator $K^M v^M$. The approximate solution v^M is then obtained by solving the linear system

$$v_\ell(a_{jk}) - (K_\ell v^M)(a_{jk}) = u_\ell^i(a_{jk}), \quad (3.18)$$

where $\ell = -M, \dots, M$, $j = 0, \dots, N_p$ and $k = 1, \dots, N_i$. If we denote the matrix associated with the left-hand side by A , the unknown vector $v_\ell(a_{jk})$ by x and the right-hand side by b , equation (3.18) becomes simply $Ax = b$.

In general, A is dense. Therefore, because of the large amount of memory that would be required, we do not actually construct A in solving the linear system. Instead we use an iterative method that requires only the value of the matrix-vector product $y = Ax$ for any given vector x . Thus, y is computed in $\mathcal{O}(N \log N)$ operations by means of the integration schemes described in this chapter.

We use the Generalized Minimal Residual (GMRES) method. This iterative method builds a Krylov subspace associated with the matrix A and an initial guess x_0 . The Krylov subspace at the k^{th} iteration is defined as

$$\mathcal{K}_k = \text{span}\{r_0, Ar_0, A^2 r_0, \dots, A^k r_0\},$$

where $r_0 = b - Ax_0$. Typically, GMRES builds an orthonormal basis for this subspace using a modified Gram-Schmidt procedure. The approximate solution to the linear system at the k^{th} iteration, $x_k = x_0 + y_k$, is the vector that minimizes (in the least-squares sense) the residual $r = b - A(x_0 + y)$ for $y \in \mathcal{K}_k$ [29, pp. 38–41].

We chose GMRES because of its generality (it applies to non-Hermitian matrices) and because of its “optimal” convergence properties (optimal in the sense that it produces the

residual with the smallest 2-norm from the Krylov subspace [29, p. 49]). On the other hand, GMRES stores the basis vectors for the Krylov subspace and, hence, requires $\mathcal{O}(kN)$ memory, where k is the number of iterations and N is the number of unknowns. Furthermore, the construction of the orthonormal basis requires $\mathcal{O}(k^2N)$ operations. To avoid the large memory requirement, one can restart GMRES after a specified number of iterations, which has the side-effect of slowing the convergence.

In this context, it is useful to consider how the number of required iterations depends on the problem size. Unlike many classical methods for finite-difference equations, in this integral equation method, the number of iterations required to obtain a given residual tolerance *does not* depend on the mesh size. On the other hand, the number of iterations does depend on the size of the scatterer (as measured in interior wavelengths). In other words, for a given scatterer, as we increase the frequency of the incident wave (increase κ) and/or increase the value of m , the number of iterations required increases significantly (see Section 3.2.2).

Of course, other methods exist for non-Hermitian matrices that do not have the memory and complexity issues of GMRES, such as the Quasi-Minimal Residual (QMR), Biconjugate-Gradient (BiCG), stabilized Biconjugate-Gradient (BiCGSTAB) and Conjugate-Gradient Squared (CGS) methods [29, pp. 92–94]. In general, these methods sacrifice convergence rates for memory and complexity, but may be preferable in some cases. (See Section 4.5 for more discussion on these solvers.)

3.2.2 Preconditioner

In the previous sections, we described the numerical implementation of the method as well as the linear solver. However, when the scatterer is large ($\kappa R \gg 1$), or when the contrast is large ($|m| \gg 1$), the linear solver may converge so slowly that it becomes infeasible to obtain the desired residual value. We try to remedy this problem by means of a preconditioning matrix P for the linear system. The expectation is that pre- or post-multiplication of A (as well as the right-hand side) with an appropriate P will yield a new matrix with an eigenvalue spectrum that allows the linear solver to converge more rapidly. Of course, if P were the exact inverse of A , we would have $PA = AP = I$ and the linear solver would converge in a single iteration. Hence, roughly speaking, an effective preconditioner P will approximate the inverse of A as closely as possible while still allowing efficient pre- or post-multiplication.

Since our numerical method has a complexity of $\mathcal{O}(N \log N)$ where $N = N_p N_i M$, we want the multiplication by the preconditioner to be at least as efficient.

In this section, we introduce a preconditioner based on an approximate scatterer: we approximate the true inhomogeneity m by a piecewise constant, radially layered inhomogeneity $\tilde{m}$. In a sense, this inhomogeneity forms a zeroth-order approximation to m in both geometry and value. As we will show, because of the relatively simple character of this scatterer, one can invert the associated integral equation in closed form. This inverse operator, which we use as our preconditioner P , can be computed in $\mathcal{O}(N)$ operations using radial integration methods identical to those described previously in Section 3.1.2.

As described above we define $\tilde{m}$ as

$$\tilde{m}(x) = \sum_{j=1}^q m_j \chi_{A_j}(x),$$

where m_j are constants. The sets $A_j, j = 1, 2, \dots, q$ are the annular regions

$$\begin{aligned} A_j &= \{x : a_{j-1} \leq |x| \leq a_j\}, \\ 0 &= a_0 < a_1 < \dots < a_q = R_1, \end{aligned}$$

where χ_{A_j} is the characteristic function of the set A_j . Here we have assumed that $R_0 = 0$ to simplify the discussion somewhat; the case for $R_0 > 0$ proceeds similarly. The preconditioner P is given by the inverse of the associated integral equation

$$v(x) + \kappa^2 \sum_{j=1}^q m_j \int_{A_j} \Phi(\kappa|x-y|) v(y) dy = w(x), \quad (3.19)$$

where $|x| \leq R_1$ and $\Phi(z) = i/4H_0^1(z)$. Given the inverse P , we then left-precondition our original equation as follows

$$P(I - K^M)v^M = Pu^{i,M}.$$

This linear system is solved using an iterative solver, as described previously. Since the right-hand side of (3.19) is given by $w = (I - K^M)v^M$, $w = w^M$.

To solve (3.19), we derive an equivalent differential equation. Define a new unknown $u = v - w$, where v solves (3.19) and w is the given right-hand side. Although v is only a

solution for $|x| \leq R_1$, we can define u in all of $\mathbb{R}^2$ by

$$u = -\kappa^2 \sum_{j=1}^q m_j \int_{A_j} \Phi(\kappa|x-y|)v(y)dy.$$

For $x \in A_j$, we have

$$(\Delta + \kappa^2 n_j^2)u(x) = \kappa^2 m_j w, \quad j = 1, 2, \dots, q$$

where $m_j = 1 - n_j^2$. For $|x| > R_1$ (and $|x| < R_0$ when $R_0 \neq 0$),

$$(\Delta + \kappa^2)u(x) = 0.$$

Furthermore, it is not difficult to show that u satisfies the Sommerfeld radiation condition as $|x| \rightarrow \infty$ [17, pp. 216–217]. This yields a differential equation in all of $\mathbb{R}^2$ that is equivalent to the preconditioning integral equation (3.19).

This equivalent differential equation can be solved in closed form as follows. For $x \in A_j$, we write the solution as a sum of a particular solution and a homogeneous solution, $u(x) = u_p(x) + u_h(x)$. Hence, for $x \in A_j$, a particular solution to the equation is

$$u_p(x) = -\kappa^2 m_j \int_{A_j} \Phi(\kappa n_j |x-y|)w(y)dy.$$

Since $w = w^M$, $u_p = u_p^M$ and in polar coordinates (a, ϕ)

$$(u_p)_\ell(a) = -\kappa^2 m_j \left[H_\ell^1(\kappa n_j a) \int_{a_{j-1}}^a J_\ell(\kappa n_j r) w_\ell(r) r dr + J_\ell(\kappa n_j a) \int_a^{a_j} H_\ell^1(\kappa n_j r) w_\ell(r) r dr \right], \quad (3.20)$$

for $a_{j-1} \leq a \leq a_j$. Clearly, these integrals can be computed with high-order accuracy and in $\mathcal{O}(N)$ operations using the *same methods* for radial integration that we discussed in Section 3.1.2. Furthermore, since the integration methods require the value of $w_\ell(r)$ at the Chebyshev points, which are exactly the values we have, application of these methods to the computation of $(u_p)_\ell$ requires no interpolation.

The homogeneous solution on the other hand is given by

$$u_h(a, \phi) = \begin{cases} \sum_{\ell=-\infty}^{\infty} \alpha_{\ell}^{(1)} J_{\ell}(\kappa n_j a) e^{i\ell\phi}, & \text{if } j = 1, \\ \sum_{\ell=-\infty}^{\infty} [\alpha_{\ell}^{(j)} J_{\ell}(\kappa n_j a) + \beta_{\ell}^{(j-1)} Y_{\ell}(\kappa n_j a)] e^{i\ell\phi}, & \text{if } j = 2, 3, \dots, q, \end{cases} \quad (3.21)$$

for $a_{j-1} \leq a \leq a_j$. (If $R_0 \neq 0$, the homogeneous solution for $j = 1$ is given by a linear combination of J_{ℓ} and Y_{ℓ} instead of J_{ℓ} alone. Furthermore, if $R_0 > 0$, we must consider the homogeneous solution in the additional region $a < R_0$, which takes the same form as for $j = 1$ in (3.21).) Finally, for $a > R_1$, we have

$$u(a, \phi) = \sum_{\ell=-\infty}^{\infty} \beta_{\ell}^{(q)} H_{\ell}^1(\kappa a) e^{i\ell\phi}.$$

Clearly, given the correct values of the coefficients $\alpha_{\ell}^{(j)}$ and $\beta_{\ell}^{(j)}$, we can compute $(u_h)_{\ell}(a)$ for $\ell = -M, \dots, M$ and all radial discretization points in $\mathcal{O}(N)$ operations. We thus obtain a closed-form solution to the integral equation (3.19).

To compute the $2q$ coefficients $\alpha_{\ell}^{(j)}$ and $\beta_{\ell}^{(j)}$, we require that $u \in C^1(\mathbb{R}^2)$ [28, pp. 53, 56]. Hence, by enforcing this condition at each $a_j, j = 1, 2, \dots, q$, we obtain $2q$ equations, for $2q$ unknowns. For $j = 1$,

$$\begin{aligned} \alpha_{\ell}^{(1)} J_{\ell}(\kappa n_1 a_1) + (u_p)_{\ell}(a_1^-) &= \alpha_{\ell}^{(2)} J_{\ell}(\kappa n_2 a_1) + \beta_{\ell}^{(1)} Y_{\ell}(\kappa n_2 a_1) + (u_p)_{\ell}(a_1^+) \\ \alpha_{\ell}^{(1)} n_1 J'_{\ell}(\kappa n_1 a_1) + \frac{1}{\kappa} (u_p)'_{\ell}(a_1^-) &= \alpha_{\ell}^{(2)} n_2 J'_{\ell}(\kappa n_2 a_1) + \beta_{\ell}^{(1)} n_2 Y'_{\ell}(\kappa n_2 a_1) + \frac{1}{\kappa} (u_p)'_{\ell}(a_1^+) \end{aligned}$$

For $j = 2, 3, \dots, q-1$, we have

$$\begin{aligned} \alpha_{\ell}^{(j)} J_{\ell}(\kappa n_j a_j) + \beta_{\ell}^{(j-1)} Y_{\ell}(\kappa n_j a_j) + (u_p)_{\ell}(a_j^-) &= \alpha_{\ell}^{(j+1)} J_{\ell}(\kappa n_{j+1} a_j) \\ &\quad + \beta_{\ell}^{(j)} Y_{\ell}(\kappa n_{j+1} a_j) + (u_p)_{\ell}(a_j^+) \\ \alpha_{\ell}^{(j)} n_j J'_{\ell}(\kappa n_j a_j) + \beta_{\ell}^{(j-1)} n_j Y'_{\ell}(\kappa n_j a_j) + \frac{1}{\kappa} (u_p)'_{\ell}(a_j^-) &= \alpha_{\ell}^{(j+1)} n_{j+1} J'_{\ell}(\kappa n_{j+1} a_j) \\ &\quad + \beta_{\ell}^{(j)} n_{j+1} Y'_{\ell}(\kappa n_{j+1} a_j) + \frac{1}{\kappa} (u_p)'_{\ell}(a_j^+) \end{aligned}$$

Finally, for $j = q$, we have

$$\begin{aligned} \alpha_{\ell}^{(q)} J_{\ell}(\kappa n_q R_1) + \beta_{\ell}^{(q-1)} Y_{\ell}(\kappa n_q R_1) + (u_p)_{\ell}(R_1^-) &= \beta_{\ell}^{(q)} H_{\ell}^1(\kappa R_1) + (u_p)_{\ell}(R_1^+) \\ \alpha_{\ell}^{(q)} n_q J'_{\ell}(\kappa n_q R_1) + \beta_{\ell}^{(q-1)} n_q Y'_{\ell}(\kappa n_q R_1) + \frac{1}{\kappa} (u_p)'_{\ell}(R_1^-) &= \beta_{\ell}^{(q)} (H_{\ell}^1)'(\kappa R_1) \end{aligned}$$

Here $(u_p)_\ell(a_j^+) = \lim_{a \rightarrow a_j^+} (u_p)_\ell(a)$ and $(u_p)_\ell(a_j^-) = \lim_{a \rightarrow a_j^-} (u_p)_\ell(a)$ with corresponding definitions for the derivatives.

For each ℓ , the matrix associated with this linear system is constant and banded with five diagonals. Hence, we compute the LU -decomposition of all of these matrices in $\mathcal{O}(qM)$ operations only once at the beginning of each run. In each iteration, after computing the $(u_p)_\ell(a_j^+)$ and $(u_p)_\ell(a_j^-)$ and their derivatives, we use the LU -decomposition to solve for the values of $\alpha_\ell^{(j)}$ and $\beta_\ell^{(j)}$ for all $j = 1, 2, \dots, q$. This again requires a total of $\mathcal{O}(qM)$ operations.

Finally, we compute $(u_p)_\ell(a_j^+)$ and $(u_p)_\ell(a_j^-)$ and their derivatives. For $a_{j-1} < a < a_j$,

$$(u_p)_\ell(a) = -\frac{i\pi}{2}\kappa^2 m_j \left[H_\ell^1(\kappa n_j a) \int_{a_{j-1}}^a J_\ell(\kappa n_j r) w_\ell(r) r dr + J_\ell(\kappa n_j a) \int_a^{a_j} H_\ell^1(\kappa n_j r) w_\ell(r) r dr \right].$$

$$(u_p)_\ell(a_{j-1}^+) = -\frac{i\pi}{2}\kappa^2 m_j J_\ell(\kappa n_j a_{j-1}) \int_{a_{j-1}}^{a_j} H_\ell^1(\kappa n_j r) w_\ell(r) r dr$$

and

$$(u_p)_\ell(a_j^-) = -\frac{i\pi}{2}\kappa^2 m_j H_\ell^1(\kappa n_j a_j) \int_{a_{j-1}}^{a_j} J_\ell(\kappa n_j r) w_\ell(r) r dr$$

Similarly,

$$(u_p)_\ell'(a_{j-1}^+) = -\frac{i\pi}{2}\kappa^3 n_j m_j J_\ell'(\kappa n_j a_{j-1}) \int_{a_{j-1}}^{a_j} H_\ell^1(\kappa n_j r) w_\ell(r) r dr$$

and

$$(u_p)_\ell'(a_j^-) = -\frac{i\pi}{2}\kappa^3 n_j m_j (H_\ell^1)'(\kappa n_j a_j) \int_{a_{j-1}}^{a_j} J_\ell(\kappa n_j r) w_\ell(r) r dr$$

These integrals are easily obtained from the values of $(u_p)_\ell(a)$ in $\mathcal{O}(qM)$ operations.

Thus, application of this preconditioner requires only $\mathcal{O}(N)$ operations per iteration. We expect to observe the greatest benefits in using this preconditioner when the original inhomogeneity can be approximated reasonable well by a piecewise constant radially layered scatterer. Of course, the size of the annuli A_j as well as the approximating values m_j can be tuned to improve the convergence rate. We demonstrate the performance of the

preconditioner through the computational examples in Section 5.1.3.

Chapter 4

A Fast, High-Order Method in Three Dimensions

In this chapter, we present a new, fast, high-order method in three dimensions. This method is motivated in part by the two-dimensional polar coordinates approach discussed in the previous two chapters. One could propose a direct generalization of the two-dimensional approach to a three-dimensional spherical coordinates method. However, to obtain the same $\mathcal{O}(N \log N)$ complexity, one would need to replace the FFTs with fast spherical harmonics transforms, robust versions of which are not readily available. Furthermore, as evidenced by the method for radial integration in the two-dimensional method, evaluation of special functions of high-order presents several mathematical as well as practical difficulties. Therefore, instead of extending the two-dimensional polar coordinates approach, we present a three-dimensional approach based on trapezoidal rule integration and Fourier series approximation in *Cartesian* coordinates. The fast, high-order method we obtain is actually much simpler than the two-dimensional approach while yielding approximately the same accuracy and efficiency.

In particular, while the polar coordinates approach in two dimensions required the identification and resolution of singularities in the Fourier coefficients $(mu)_\ell(r)$, this Cartesian approach in three dimensions requires no special consideration of scatterer geometry. This enables the construction of quite complicated scatterers containing discontinuities, corners, and cusps by simply summing the truncated Fourier series representations of each of the individual components of the scatterer (see Figures 5.13 and 5.14 as well as Table 5.13). There is no need to locate and resolve singularities.

As in the two-dimensional approach, the core of the numerical method is an efficient,

high-order scheme for computing the integral operator $(Ku)(x_j)$ (see (1.6)) at the given discretization points x_j . We thereby obtain the linear system

$$u(x_j) - (Ku)(x_j) = u^i(x_j), \quad (4.1)$$

whose solution $u(x_j)$ approximates the total field. An iterative solver then provides the solution at the discretization points $u(x_j)$. We discuss the advantages and disadvantages of various iterative solvers in Section 4.5.

Let m be contained within a box $\Omega_{[a,b]}$ with corners $a, b \in \mathbb{R}^3$, i.e., $\Omega_{[a,b]} = \{x : a_q \leq x_q \leq b_q, q = 1, 2, 3\}$. Then,

$$(Ku)(x) = -\kappa^2 \int_{\Omega_{[a,b]}} g(x-y)m(y)u(y)dy,$$

where u is the solution of (1.6) in $\mathbb{R}^3$.

Before specifying the details of our approach, we sketch its key aspects. As mentioned in the introduction, we decompose the Green's function into a smooth part with infinite support, $g_{smth}(x)$, and a singular part with compact support, $g_{cmp}(x)$. More precisely, we define $g_{smth}(x)$ and $g_{cmp}(x)$ by

$$g(x) = g(x)(1 - p(x)) + g(x)p(x) = g_{smth}(x) + g_{cmp}(x),$$

where $p(x) \in C^\infty$ is a partition of unity function such that $p(x) = 1$ near $x = 0$ and $p(x) = 0$ outside some neighborhood of $x = 0$. (Of course, there are many such partition of unity functions, and we do not specify a particular choice at this time.) It is then necessary to compute the two convolutions

$$(K_{smth}u)(x) = -\kappa^2 \int_{\Omega_{[a,b]}} g_{smth}(x-y)m(y)u(y)dy \quad (4.2)$$

$$(K_{cmp}u)(x) = -\kappa^2 \int_{\Omega_{[a,b]}} g_{cmp}(x-y)m(y)u(y)dy. \quad (4.3)$$

The following two sections describe the two different high-order methods we use to evaluate $K_{smth}u$ and $K_{cmp}u$, respectively. For both of these convolutions, the substitution of the scatterer by an appropriate Fourier-smoothed scatterer is absolutely key in obtaining high-order accuracy. As will be shown, the methods used to evaluate these convolutions

require the computation of an integral of the form

$$\int_{\Omega_{[a,b]}} m(y)w(y)dy,$$

where w is defined in $\mathbb{R}^3$. In each case, w is the product of the total field u and a known C^∞ function. Hence, the regularity of w is given by the regularity of u , which is related to the regularity of m . In particular, if $m \in L^\infty$, $u \in C^1$; and if $m \in C^{k,\alpha}$, $u \in C^{k+2,\alpha}$ (see Theorem 2.5).

Note that the integrands in each case are quite similar to the example in Figure 1.2, with one important exception: Although w is smooth (at least C^1), it is generally not periodic. Hence, direct substitution of the scatterer by its truncated Fourier series will not yield high-order convergence. However, we observe that since m vanishes outside of $\Omega_{[a,b]}$, we can extend the domain of integration without affecting the value of the integral. Similarly, any modification to w outside of the support of m , $\text{supp}(m)$, does not affect the integral.

With these observations in mind, we increase the computational domain to the box $\Omega_{[a-\delta,b+\delta]}$ for some $\delta \in \mathbb{R}^3$ such that the components $\delta_q > 0$. This gives us u and hence, w on $\Omega_{[a-\delta,b+\delta]}$. We then multiply $w(y)$ by a smooth cutoff function $p_m(y)$ such that $p_m \in C^\infty$, $p_m(y) = 1$ for $y \in \text{supp}(m)$ and $p_m(y) = 0$ for $y \notin \Omega_{[a-\delta,b+\delta]}$. (Of course, as with the partition of unity function, p , there are many such functions, p_m , and we do not specify a particular choice at this time.) This gives us

$$\int_{\Omega_{[a,b]}} m(y)w(y)dy = \int_{\Omega_{[a-\delta,b+\delta]}} m(y)p_m(y)w(y)dy.$$

Since $p_m(y)$ vanishes outside of $\Omega_{[a-\delta,b+\delta]}$, one can extend $p_m(y)w(y)$ as a smooth and periodic function. Following the example in the introduction, we can now substitute m by its truncated Fourier series to obtain high-order accuracy when integrating with the trapezoidal rule (see Figure 1.2 and Table 1.2), i.e., replace m by

$$m^F(x) = \sum_{\ell_1=-F_1}^{F_1} \sum_{\ell_2=-F_2}^{F_2} \sum_{\ell_3=-F_3}^{F_3} m_\ell e^{2\pi i c_\ell \cdot x},$$

where $\ell = (\ell_1, \ell_2, \ell_3)$ and the q^{th} component of c_ℓ , $(c_\ell)_q = \ell_q / [(b_q - a_q) + 2\delta_q]$ for $q = 1, 2, 3$

and $F = (F_1, F_2, F_3)$. This gives

$$\int_{\Omega_{[a,b]}} m(y)w(y)dy \approx \int_{\Omega_{[a-\delta, b+\delta]}} m^F(y)p_m(y)w(y)dy.$$

Thus, we obtain high-order accuracy in such integrals essentially by replacing m by $\tilde{m} = m^F p_m$ and by replacing $\Omega_{[a,b]}$ by $\Omega_{[\tilde{a}, \tilde{b}]}$, where $\tilde{a} = a - \delta$ and $\tilde{b} = b + \delta$.

In Sections 4.1 and 4.2, we describe in more detail the evaluation of $K_{smth}u$ and $K_{cmp}u$, respectively. In Section 4.3, we discuss a few additional details associated with this substitution of m by m^F . Section 4.4 briefly describes the method we use to compute the Fourier coefficients of g_{cmp} . Finally, in Section 4.5, we describe our parallel implementation of the method and discuss the relative advantages of the various linear solvers.

4.1 High-Order Convolution with Smooth Part

We compute $K_{smth}u$ by means of the trapezoidal rule after substituting m by $\tilde{m}$ and $\Omega_{[a,b]}$ by $\Omega_{[\tilde{a}, \tilde{b}]}$ as described above. Given a number of discretization points $N \in \mathbb{N}^3$. Define the discretization spacing $h \in \mathbb{R}^3$ such that $h_q = (\tilde{b}_q - \tilde{a}_q)/N_q$ for $q = 1, 2, 3$. Define the associated evenly spaced discretization points $x_j, y_k \in \mathbb{R}^3$ such that $(x_j)_q = \tilde{a}_q + j_q h_q$ and $(y_k)_q = \tilde{a}_q + k_q h_q$, where j and k are three-dimensional indices such that their components satisfy $0 \leq j_q, k_q \leq N_q$ for $q = 1, 2, 3$.

Since $\tilde{m}(x)$ vanishes on the boundary of $\Omega_{[\tilde{a}, \tilde{b}]}$ (because $\tilde{m}$ is smooth and compactly supported in $\Omega_{[\tilde{a}, \tilde{b}]}$), the trapezoidal rule gives

$$(K_{smth}u)(x_j) \approx Prod(h) \sum_{k_1=0}^{N_1-1} \sum_{k_2=0}^{N_2-1} \sum_{k_3=0}^{N_3-1} g_{smth}(x_j - y_k) \tilde{m}(y_k) u(y_k),$$

where the notation, $Prod(h)$, which we use throughout this chapter, stands for $Prod(h) = h_1 h_2 h_3$. To further simplify the notation, we denote the triple sum above as $\sum_{k=0}^{N-1}$, which allows us to write

$$\begin{aligned} (K_{smth}u)(x_j) &\approx Prod(h) \sum_{k=0}^{N-1} g_{smth}(x_j - y_k) \tilde{m}(y_k) u(y_k) \\ &= Prod(h) \sum_{k=0}^{N-1} (g_{smth})_{j-k} \tilde{m}_k u_k. \end{aligned}$$

Here $(g_{smth})_k = g_{smth}((k_1 h_1, k_2 h_2, k_3 h_3))$, $\tilde{m}_k = \tilde{m}(y_k)$ and $u_k = u(y_k)$.

Hence, using the trapezoidal rule to evaluate this integral is algorithmically equivalent to computing a discrete convolution. We compute this convolution using FFTs in $\mathcal{O}(Prod(N) \log Prod(N))$ operations [45, pp. 531–536]. Thus, we obtain an efficient and high-order accurate method for computing $K_{smth}u$.

4.2 High-Order Convolution with Singular Part

To obtain high-order accuracy in the computation of (4.3), we approximate $K_{cmp}u$ by a truncated Fourier series. As described in the introduction, a truncated Fourier series provides high-order accuracy if the approximated function is smooth and periodic. Since g_{cmp} and m are both compactly supported, $(K_{cmp}u)(x)$ vanishes for points x sufficiently far from the inhomogeneity. More precisely, assume that the support of g_{cmp} is contained in a box $\Omega_{[-d,d]}$. Then, for $x \notin \Omega_{[a-d,b+d]}$, $(K_{cmp}u)(x) = 0$. Furthermore, since the convolution is a smoothing operation, $(K_{cmp}u)(x)$ is a smooth function, *even in the case of a discontinuous scatterer* (see Theorem 2.5). Therefore, $(K_{cmp}u)$ can be extended as a smooth and periodic function on $\mathbb{R}^3$.

Hence, we may represent $K_{cmp}u$ accurately by a truncated Fourier series if we choose a period for the expansion that contains $\Omega_{[a-d,b+d]}$, i.e.,

$$(K_{cmp}u)(x) \approx \sum_{\ell=-M}^M (K_{cmp}u)_\ell e^{2\pi i c_\ell \cdot (x-a)}, \quad (4.4)$$

where $(c_\ell)_q = \frac{\ell_q}{B_q - A_q}$ and $A_q \leq a_q - d_q \leq b_q + d_q \leq B_q$ for $q = 1, 2, 3$. Note that according to our convention, we have denoted the triple sum as $\sum_{\ell=-M}^M$. We have also shifted the Fourier basis functions by a . This simplifies the presentation somewhat.

We must now compute the Fourier coefficients $(K_{cmp}u)_\ell$. We have

$$\begin{aligned} (K_{cmp}u)_\ell &= -\frac{\kappa^2}{Prod(B-A)} \int_{\Omega_{[A,B]}} (K_{cmp}u)(x) e^{-2\pi i c_\ell \cdot (x-a)} dx \\ &= -\frac{\kappa^2}{Prod(B-A)} \int_{\Omega_{[A,B]}} m(y) u(y) e^{-2\pi i c_\ell \cdot (y-a)} dy \int_{\Omega_{[-d,d]}} g_{cmp}(z) e^{-2\pi i c_\ell \cdot z} dz \\ &= -\kappa^2 (g_{cmp})_\ell (mu)_\ell, \end{aligned}$$

where

$$(g_{cmp})_\ell = \int_{\Omega_{[-d,d]}} g_{cmp}(z) e^{-2\pi i c_\ell \cdot z} dz$$

and

$$(mu)_\ell = \frac{1}{Prod(B-A)} \int_{\Omega_{[A,B]}} m(y) u(y) e^{-2\pi i c_\ell \cdot (y-a)} dy.$$

Note that the coefficients $(g_{cmp})_\ell$ are computed with respect to a slightly different set of basis functions (they are not shifted by a) and that we integrate over $\Omega_{[-d,d]}$ instead of $\Omega_{[A,B]}$. Because $g_{cmp}(x)$ is known analytically, we need only compute its Fourier coefficients once at the beginning of each run. Furthermore, because of the singularity in g_{cmp} , we must take special care to compute these coefficients accurately. Our method for computing these coefficients is presented in Section 4.4.

To compute the Fourier coefficients $(mu)_\ell$, we again use the trapezoidal rule. As in the previous section, we obtain high-order accuracy by substituting m by $\tilde{m}$

$$\begin{aligned} (mu)_\ell &= \frac{1}{Prod(B-A)} \int_{\Omega_{[A,B]}} m(y) u(y) e^{-2\pi i c_\ell \cdot (y-a)} dy \\ &\approx \frac{1}{Prod(B-A)} \int_{\Omega_{[\tilde{a},\tilde{b}]}} \tilde{m}(y) u(y) e^{-2\pi i c_\ell \cdot (y-a)} dy \\ &\approx \frac{Prod(h)}{Prod(B-A)} \sum_{j=0}^{N-1} \tilde{m}_j u_j e^{-2\pi i c_\ell \cdot (j_1 h_1, j_2 h_2, j_3 h_3)}. \end{aligned}$$

Note that, as indicated in the second equation above, we need only integrate over the domain $\Omega_{[\tilde{a},\tilde{b}]}$ since this domain contains the support of $\tilde{m}$.

We can evaluate this expression with an FFT by choosing $A = (A_1, A_2, A_3)$ and $B = (B_1, B_2, B_3)$ such that $(B_q - A_q)/h_q$ is an integer for each $q = 1, 2, 3$. In other words, the larger domain $\Omega_{[A,B]}$ must be exactly an integer number of cells larger than the smaller domain $\Omega_{[\tilde{a},\tilde{b}]}$ in each dimension. Then, defining $\tilde{N} \in \mathbb{N}^3$ such that $\tilde{N}_q = (B_q - A_q)/h_q$ for $q = 1, 2, 3$, and $(c_\ell)_q = \ell_q/(B_q - A_q)$, we obtain

$$(mu)_\ell \approx \frac{1}{Prod(\tilde{N})} \sum_{j=0}^{\tilde{N}-1} m_j u_j e^{-2\pi i \ell \cdot (j_1/\tilde{N}_1, j_2/\tilde{N}_2, j_3/\tilde{N}_3)},$$

where $m_j u_j = 0$ if $j_q > N_q$ for any $q = 1, 2, 3$. Hence, this discrete Fourier transform can be evaluated by means of an FFT in $\mathcal{O}(Prod(\tilde{N}) \log Prod(\tilde{N}))$ operations for $|\ell_q| \leq \tilde{N}_q/2$.

Finally, given this high-order approximation of $(mu)_\ell$ and the pre-computed $(g_{cmp})_\ell$, we can sum the Fourier series (4.4) also by means of an FFT to obtain a high-order approximation to $K_{cmp}u$.

4.3 Fourier-Smoothed Scatterers

As we have shown, replacing m by the smooth and periodic function $\tilde{m}$ and enlarging the integration domain to $\Omega_{[\tilde{a}, \tilde{b}]}$ is key in obtaining our high-order method. Table 5.11 in Section 5.2 compares the convergence rates for a discontinuous inhomogeneity with and without this substitution. Although we do not present here a complete theoretical analysis of the method, we expect convergence rates similar to those proved for the two-dimensional method (see Theorem 2.7). For example, a discontinuous, piecewise C^1 scatterer is expected to yield second- and third-order convergence on the interior and exterior of the scatterer, respectively; a C^0 , piecewise C^2 scatterer is expected to yield third- and fifth-order convergence on the interior and exterior of the scatterer, respectively. Although gains in the asymptotic rate of convergence may always be expected when substituting m by $\tilde{m}$, real practical gains are most significant for scatterers with a low degree of regularity. (See Section 3.1.1 for a related discussion on the two-dimensional method.) For this reason, one need not typically perform this substitution for very smooth scatterers such as the C^∞ scatterer considered in Figure 5.11 in Section 5.2. In such cases, the trapezoidal rule alone provides high-order accuracy.

Of course when treating discontinuous scatterers, the user is relatively free to choose the smooth cutoff function p_m and the number of modes F in the truncated Fourier approximation of m . There are couple of competing issues that should figure into this decision. In particular, the smaller δ is, the smaller h becomes for a given N . However, this does not necessarily lead to higher accuracy because a small δ implies a relatively steep rise of p_m from 0 to 1. The more abrupt this rise, the smaller we must make h to achieve a given accuracy, thereby increasing N . The choice of F , on the other hand, depends on the choice of N . The value of N must be large enough to resolve the Fourier oscillations in the integrand. In particular, if F is chosen to be too large, the trapezoidal rule will yield very little accuracy in integrating the highly oscillatory modes in m^F . Hence, to obtain high-order convergence, we choose F to be a fixed fraction of N . In our experience, $F = N/2$ is a good

choice.

4.4 Computation of the Fourier Coefficients of the Green's Function

It remains to compute the coefficients $(g_{cmp})_\ell$, which are essentially the Fourier coefficients of g_{cmp} . As defined previously, $g_{cmp}(x) = g(x)p(x)$ where $p(x)$ has support in $\Omega_{[-d,d]}$ for some $d \in \mathbb{R}^3$ with $d_q > 0$, $q = 1, 2, 3$. Note that our choice of d depends on the same issues that affect our choice of δ for the cutoff function p_m as discussed in the previous section.

Of course, there are many possible choices of such partition of unity functions. Partition of unity functions in three dimensions can be constructed from a partition of unity function ϕ in one dimension as described in in [12],

$$\phi(t) = \begin{cases} 1, & \text{for } |t| \leq t_0 \\ \exp\left(\frac{2e^{-1/x}}{x-1}\right), & \text{for } t_0 < |t| < t_1, \text{ where } x = (|t| - t_0)/(t_1 - t_0) \\ 0, & \text{for } |t| \geq t_1. \end{cases} \quad (4.5)$$

For example, $p(x) = \phi(|x|)$ and $p(x) = \phi(x_1)\phi(x_2)\phi(x_3)$ are both partition of unity functions in three-dimensions that are centered at the origin. Of course, one may shift the center of the functions to any point in $\mathbb{R}^3$ without difficulty.

When computing the coefficients $(g_{cmp})_\ell$, we choose the spherically symmetric function $p(x) = \phi(|x|)$ to simplify the computation and choose $t_1 = R$ such that $R \leq d$. Changing to spherical coordinates in the integration gives

$$\begin{aligned} (g_{cmp})_\ell &= \int_{\Omega_{[-d,d]}} g_{cmp}(z) e^{-2\pi i c_\ell \cdot z} dz \\ &= \int_0^R \int_{S^1} \frac{e^{i\kappa\rho}}{4\pi\rho} p(\rho) e^{-2\pi i \rho c_\ell \cdot \hat{z}} \rho^2 d\rho d\sigma(\hat{z}) \\ &= \int_0^R g_{cmp}(\rho) j_0(2\pi|c_\ell|\rho) \rho d\rho \\ &= \frac{1}{2\pi|c_\ell|} \int_0^R p(\rho) e^{i\kappa\rho} \sin(2\pi|c_\ell|\rho) d\rho, \end{aligned}$$

where $\int_{S^1} d\sigma(\hat{z})$ denotes integration over the unit sphere and the second to last equality follows from [17, p. 32].

We wish to emphasize that the simplification hereby achieved is quite significant. The Jacobian associated with the change to spherical coordinates cancels the ρ^{-1} singularity in Green's function. Furthermore, since the integration on the unit sphere can be performed analytically, we are left only with a one-dimensional integral to be evaluated for various values of $|c_\ell|$.

We rewrite the required integral as follows

$$\begin{aligned}
(g_{cmp})_\ell &= \frac{1}{\alpha} \int_0^R p(\rho) e^{i\kappa\rho} \sin(\alpha\rho) d\rho \\
&= \frac{1}{2i\alpha} \left[\int_0^R p(\rho) e^{i(\kappa+\alpha)\rho} d\rho - \int_0^R p(\rho) e^{i(\kappa-\alpha)\rho} d\rho \right] \\
&= \frac{1}{2i\alpha} \{ \mathcal{H}[p](\kappa + \alpha) - \mathcal{H}[p](\kappa - \alpha) \},
\end{aligned} \tag{4.6}$$

where $\alpha = 2\pi|c_\ell|$ and

$$\mathcal{H}[p](\omega) = \int_0^R p(\rho) e^{i\omega\rho} d\rho. \tag{4.7}$$

Note that since $p(\rho)$ vanishes for $\rho > R$, $\mathcal{H}[p]$ is related to the Laplace transform, $\mathcal{L}[p]$, as follows

$$\mathcal{H}[p](\omega) = \mathcal{L}[p](-i\omega).$$

Observe that $\mathcal{H}[p](-\omega) = \overline{\mathcal{H}[p](\omega)}$ since $p(\rho)$ is real-valued. It is important to note that we can only use (4.6) to compute $(g_{cmp})_\ell$ when $|\ell| \neq 0$. For $|\ell| = 0$, on the other hand, it is not difficult to see that

$$(g_{cmp})_0 = \int_0^R \rho p(\rho) e^{i\kappa\rho} d\rho = \mathcal{H}[\rho p(\rho)](\kappa).$$

Therefore, to compute $(g_{cmp})_\ell$, we need an accurate and efficient method for computing $\mathcal{H}[g](\omega)$ for $g(\rho) = p(\rho)$ and $g(\rho) = \rho p(\rho)$. This problem is not trivial since the value of $\omega \leq \kappa + \alpha$ can be quite large, thus producing a highly oscillatory integrand. Furthermore, straightforward integration by means of the trapezoidal rule will give only first-order accuracy since $p(\rho)$ and $\rho p(\rho)$ cannot be extended as smooth and periodic functions. We are able to compute these integrals accurately and efficiently using a modification of the method suggested in [45, pp. 577–584] as described in Appendix C.

4.5 Implementation

Because of the large memory and CPU-time requirements of realistic problems in three dimensions, an efficient parallel implementation of the numerical method is quite useful. An advantage of the method is its relative simplicity: roughly, it requires only an efficient (parallel) FFT implementation and an effective (parallel) iterative solver for the linear system. We make use of the parallel FFT implementation *fftw* [25, 26] and the parallel iterative solvers in PETSc [4–6]. In addition, PETSc provides excellent vector scatter and gather routines as well as useful I/O routines. These packages make development of an efficient parallel implementation relatively simple.

The bulk of the method lies in computing the required convolutions (see Sections 4.1 and 4.2). As described previously, we approximate the convolution with the smooth part of the Green's function by means the following discrete convolution

$$(K_{smth}u)(x_j) \approx Prod(h) \sum_{k=0}^{N-1} (g_{smth})_{j-k} \tilde{m}_k u_k,$$

where the components $j_q = 0, \dots, N_q$, $h_q = (\tilde{b}_q - \tilde{a}_q)/N_q$ for $q = 1, 2, 3$. This discrete convolution requires the values of $(g_{smth})_{j-k}$ for $0 \leq j_q \leq N_q$ and $0 \leq k_q \leq N_q - 1$. Hence, $-N_q + 1 \leq j_q - k_q \leq N_q$. Therefore, computing this convolution by means of FFTs requires three-dimensional arrays with dimensions $2N$, where the array containing $m_k u_k$ is padded with an appropriate number of zeroes [45, pp. 531–537]. More precisely, we first compute

$$(\hat{g}_{smth})_\ell = \sum_{j=0}^{2N-1} (g_{smth})_j e^{-2\pi i \ell \cdot (j_1/2N_1, j_2/2N_2, j_3/2N_3)}$$

and

$$\widehat{m}u_\ell = \sum_{j=0}^{2N-1} \tilde{m}_j u_j e^{-2\pi i \ell \cdot (j_1/2N_1, j_2/2N_2, j_3/2N_3)},$$

where $(g_{smth})_j$ is defined by periodic extension for j outside the range $-N_q + 1 \leq j_q \leq N_q$ and where $m_j u_j = 0$ outside the range $0 \leq j_q \leq N_q - 1$. We then use these values to compute the discrete convolution

$$\sum_{k=0}^{N-1} (g_{smth})_{j-k} m_k u_k = \sum_{\ell=0}^{2N-1} (\hat{g}_{smth})_\ell \widehat{m}u_\ell e^{2\pi i \ell \cdot (j_1/2N_1, j_2/2N_2, j_3/2N_3)},$$

for $j_q = 0, \dots, N_q$.

(Note that this straightforward approach requires a factor of $2^3 = 8$ more memory to store these convolution arrays than to store the unknowns u_j , the smoothed inhomogeneity $\tilde{m}_j$ and the incident field u_j^i . If memory usage becomes the limiting factor, it is possible to break the $m_j u_j$ array into pieces and compute the convolution with each piece separately. This saves memory, but substantially increases CPU-time. Hence, we use the straightforward approach with $2N$ -sized arrays.)

On the other hand, the approximation of the convolution with the singular part of the Green's function requires computation of the following sum

$$(K_{cmp}u)(x_j) \approx \sum_{\ell=-M}^M (g_{cmp})_{\ell} (mu)_{\ell} e^{2\pi i \ell \cdot (j_1/\tilde{N}_1, j_2/\tilde{N}_2, j_3/\tilde{N}_3)},$$

where $M_q < \tilde{N}_q/2$, $j_q = 0, \dots, N_q - 1$ and

$$(mu)_{\ell} \approx \frac{1}{Prod(\tilde{N})} \sum_{j=0}^{\tilde{N}-1} m_j u_j e^{-2\pi i \ell \cdot (j_1/\tilde{N}_1, j_2/\tilde{N}_2, j_3/\tilde{N}_3)}.$$

Hence, these sums may also be computed using FFTs. However, in this case, we use FFTs of three-dimensional arrays of size $\tilde{N}$.

In theory, the FFTs used to compute each of these convolutions need not be related. In practice, however, we save both time and memory by choosing $\tilde{N} = 2N$. If $\tilde{N} \neq 2N$, we need to compute the FFT of $(g_{cmp})_{\ell} (mu)_{\ell}$ and the FFT of $(\hat{g}_{smth})_{\ell} \widehat{mu}_{\ell}$ separately. Furthermore, we need to store $(g_{cmp})_{\ell}$, $\ell_q = -M_q, \dots, M_q$ and $(\hat{g}_{smth})_{\ell}$, $\ell_q = 0, \dots, 2N_q - 1$ in two separate arrays with a total of $Prod(\tilde{N}) + 8Prod(N) > 9Prod(N)$ elements. On the other hand, if $\tilde{N} = 2N$, our approximation for $(mu)_{\ell}$ exactly equals $\widehat{mu}_{\ell}$. Therefore, we need only compute a single FFT of $\hat{g}_{\ell} \widehat{mu}_{\ell}$, where $\hat{g}_{\ell} = (g_{cmp})_{\ell} + (\hat{g}_{smth})_{\ell}$. This single array $\hat{g}_{\ell}$ has size $2N$ and hence a total of $8Prod(N)$ elements.

A further savings that becomes quite significant in a parallel implementation is the communication costs. To compute the convolutions, we must compute FFTs of sizes $\tilde{N}$ and $2N$ and then we copy a portion of the results into an array of size N . In our parallel code, each of these arrays will, in general, be distributed differently. Hence, copying the results to and from the array of size N involves communication between processors. Therefore, if $\tilde{N} \neq 2N$, one must first communicate the values of mu into the arrays of size $\tilde{N}$ and $2N$

for computation of the FFTs, and then communicate the results back from each of these arrays into the original array of size N . On the other hand, if $\tilde{N} = 2N$, roughly half the communication is required. Hence, because of these memory and communication savings, we choose $\tilde{N} = 2N$ in our computations.

Once we have implemented the method for computing the matrix-vector product, PETSc implements a wide variety of linear solvers. However, only a few of these linear solvers are appropriate for our linear system, which is non-symmetric and indefinite, namely GMRES, CGS, BiCGSTAB and two different transpose-free QMR methods. Of these, only GMRES and BiCGSTAB perform consistently well. In fact, for each of the other methods, we found an example in which it either rapidly diverged or stagnated. At the same time, the performance of these iterative solvers seems to be somewhat problem dependent: it is certainly possible that there exists a scattering configuration for which BiCGSTAB may perform less well than one of these other approaches. GMRES, on the other hand, performs consistently well, at the price of increased memory requirements and increased computational complexity.

As mentioned in Section 3.2.1, GMRES always requires fewer matrix-vector products to converge to a given residual error than BiCGSTAB. However, at each iteration GMRES stores a new Krylov subspace basis vector whereas BiCGSTAB does not. Hence, in problems which require many iterations, GMRES may rapidly exhaust the system memory. Of course, in such cases, one may restart GMRES after a given number of iterations, thereby limiting the memory used. However, we found that restarted GMRES loses much of its advantage over BiCGSTAB. Therefore, in problems requiring many iterations, BiCGSTAB has become our method of choice.

As in the two-dimensional method, the number of iterations required to achieve a given residual tolerance does not depend on mesh size. However, as the size of the scatterer is increased (as measured in interior wavelengths), the number of required iterations increases. Furthermore, in our numerical experiments, it appears that the number of required iterations increases more for BiCGSTAB than for GMRES. Thus, BiCGSTAB loses its edge over GMRES as the problem size increases.

The numerical examples of Section 5.2 illustrate the versatility, efficiency, and high-order accuracy of the three-dimensional method described in this chapter. Since we have not optimized for parallel performance, we do not present parallel speed-up results. However,

we do include results from parallel computations, from which one can obtain some idea of the parallel performance.

Chapter 5

Computational Examples

This chapter demonstrates through several computational examples the efficiency, high-order accuracy, and overall performance of the two- and three-dimensional methods discussed in this text. These examples include scatterers for which an exact solution is known, fully inhomogeneous scatterers (non-constant m) and scatterers containing geometric singularities such as discontinuities, corners and cusps. This thesis has focused on the high-order accuracy of these methods. Hence, in each example, we have attempted to show the convergence of the method over a fairly wide range of discretizations. This becomes increasingly difficult to accomplish as the size of the scatterer (as measured in interior wavelengths) increases because of the additional memory and linear solver iterations required. Therefore, these examples do not accurately portray the full range of our computational methods.

Despite this fact, however, the examples in this chapter are similar in size to the existing examples in the electromagnetics literature. The articles [41] and [51] are perhaps the most recent results that we have found. In [41], published in 1996, the authors proposed a three-dimensional method for scattering by inhomogeneous media. They computed the scattered field from a layered sphere with interior wavelength diameter $\approx 2.7\lambda$ as well as from a dielectric cube of less than one interior wavelength on a side. In [51], published in 2001, the authors proposed a modification to the CG-FFT method in two dimensions to make use of the non-uniform FFT. It appears that the largest scatterer they consider is a $13.3\lambda \times 13.3\lambda$ rectangular cylinder, where λ is the incident wavelength. However, in their example, the wave speed is faster inside the inhomogeneity than it is outside. Hence, it appears that the rectangular cylinder has approximate *interior* wavelength dimensions of $8\lambda \times 8\lambda$. The examples we present in this chapter are of comparable size while at the same time achieving very high accuracies for much more complicated scatterers.

5.1 Two-Dimensional Computational Examples

In the following sections, we illustrate the performance of the two-dimensional algorithm for a variety of scattering configurations. First, in Section 5.1.1, we seek to verify the convergence rates predicted by Theorem 2.7. We first illustrate the convergence of the method for two scatterers that admit analytical solutions. We then verify that the algorithm achieves the predicted convergence rates for three scatterers of differing regularity. In Section 5.1.2, we demonstrate the convergence of the radial integration scheme described in Section 3.1.2 for both a smooth and a discontinuous scatterer. Finally, in Section 5.1.3, we demonstrate the performance of the method with and without the preconditioner for three rather complex scatterers containing inhomogeneities, discontinuities, corners and cusps. To illustrate the performance of the preconditioner, we present the computational time and the number of GMRES iterations required to compute the solution with and without the use of the preconditioner for a few of the runs in each of these examples.

In each case, we compute the near and far fields produced under plane wave incidence, $u^i(x, y) = e^{i\kappa x}$. To compute the maximum error in the near field, we interpolate the solution computed by our method to an evenly spaced polar grid. On this grid, we compute the maximum absolute error as compared with either the analytical solution (when it is available) or the solution computed with a finer discretization. The maximum error in the far field is computed similarly by interpolating to an evenly-spaced angular grid. The results for each example are given in the accompanying figures and tables. The figures include visualizations of $q(x) = -m(x) = n^2(x) - 1$ and the computed near field intensity. The tables provide values for a subset of the following: the number of modes in the approximate solution M , the number of radial points N_r , the total number of unknowns $N = \mathcal{O}(MN_r)$, the approximate memory required, the number of GMRES iterations used, the wall-clock time required and the maximum absolute errors in the near and far field denoted by ϵ_u^{nf} and ϵ_u^{ff} , respectively. Additionally, in many of the tables the ratios of the errors at successive levels of discretization are listed to illustrate the convergence rates. At times the accuracy in the computed solution reaches the accuracy of the analytical solution, the accuracy of the radial integration or the tolerance of the linear solver. In such a case, we observe no improvement in the error of the solution as we refine the discretization and hence, to indicate a converged solution, we write “Conv.” in the ratio column.

5.1.1 Verification of Predicted Convergence Rates

Our main goal in this section is to verify the convergence rates proved in Theorem 2.7. We also seek to demonstrate the $\mathcal{O}(M \log M)$ complexity of the angular integration method. Theorem 2.7 does not actually address the convergence of a numerical method, but rather focuses on the convergence of the analytical solution of an approximate integral equation to the true solution of the scattering equation (1.6). Hence, in this section, we are primarily concerned with the convergence in the number of Fourier modes M , rather than the convergence in the number of radial points (which is discussed in Section 5.1.2). Therefore, we fix the number of radial points N_r at a sufficiently large value and we further hold the number of GMRES iterations fixed at a value that produces a sufficiently accurate solution of the linear system. This isolates the dependence of the times and errors on M and allows us to confirm the computational complexity and the predicted convergence rates. We note that all of these results were computed using a 700MHz Pentium III Xeon workstation with 3GB of RAM.

We first compute the scattering by two obstacles for which an analytical solution exists: (1) a cylindrically symmetric scatterer centered at the origin with piecewise-constant refractive index, and (2) a disc centered at $(1\lambda, 0)$ with constant refractive index. Closed-form solutions for such piecewise-constant, cylindrically symmetric scatterers are found by separation of variables.

Perhaps the best indication of the difficulty of a scattering problem is given by its size in terms of interior wavelengths, since the numerical method must resolve these wavelengths sufficiently to provide any accuracy. For example, a scattering disc with refractive index n has a diameter of nd interior wavelengths, where d is the diameter in incident wavelengths.

The results for the first example are presented in Figure 5.1 and Table 5.1. Here the inner disc has a radius equal to 1λ and a refractive index $n = 2$; the outer annulus has an outer radius of 2λ and a refractive index $n = 3$. Thus, in terms of interior wavelengths, this scatterer has a diameter of 10λ . For this simple geometry, the preconditioner provides a nearly exact inverse. Hence, we require only two GMRES iterations to obtain a residual nearing machine precision. One may also observe that the method obtains an exponential convergence rate. This occurs despite the discontinuity in the refractive index because, at each radius, the refractive index is a C^∞ function of the angular variable. Finally, we

M	N_r	N	Memory	Iter.	Time	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
15	136	4K	3M	2	3.05s	8.50e-2		4.28e-2	
30	136	8K	9M	2	3.83s	1.13e-9	7.52e+7	5.46e-13	7.83e+10
60	136	17K	14M	2	5.46s	1.68e-12	6.73e+2	4.97e-13	Conv.

Table 5.1: Convergence Rate for Radially Layered Scatterer – Diameter = 4λ

M	N_r	N	Memory	Iter.	Time	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
15	544	17K	10M	14	7s	6.22e-2			
30	544	33K	18M	14	13s	5.95e-3	10.45	1.58e-3	18.80
60	544	66K	40M	14	25s	1.13e-3	5.27	1.83e-4	8.63
120	544	131K	78M	14	49s	2.83e-4	3.99	2.27e-5	8.06
240	544	262K	168M	14	99s	5.99e-5	4.72	2.84e-6	7.99
480	544	524K	300M	14	194s	6.65e-6	9.01	3.56e-7	7.98
960	544	1047K	633M	14	386s	1.99e-6	3.34	4.42e-8	8.05
1920	544	2093K	1332M	14	808s	2.75e-7	7.24	4.21e-9	10.50


Table 5.2: Convergence for Off-Center Disc – Centered at $(1\lambda, 0)$, Diameter = 1λ

observe that the memory depends linearly on the number of unknowns and the the time required is consistent with an $\mathcal{O}(M \log M)$ complexity.


The results for the second example are presented in Figure 5.2 and Table 5.2. Here the disc is centered at $(1\lambda, 0)$, has a diameter of 1λ and a refractive index $n = \sqrt{2}$. In terms of interior wavelengths, it has a diameter of $\sqrt{2}\lambda$. As opposed to the previous example, however, we do not observe an exponential rate of convergence despite the fact that the disc has a constant refractive index. Since the disc is not centered at the origin, the refractive index at each radius is actually a *discontinuous* function of the angular variable. Since the analytical solution in this case is known, the off-center disc provides direct verification of the predicted convergence rates for a *discontinuous* scatterer. The table shows excellent agreement with the predicted third-order convergence in the far field. The convergence in the near-field is less steady, but seems consistent with the predicated second-order convergence in the near-field (see Theorem 2.7). As in the previous example, we observe that the time and memory scale appropriately with M .

We now illustrate the convergence of the method for a series of three simple scatterers of increasing regularity. In each case, $m(x) = 1 - n^2(x)$ is given in the following form.

$$m(r, \theta) = \begin{cases} -\frac{3}{2} - \frac{1}{2\pi} \sum_{|\ell| \geq 1} \left(\frac{i}{\ell}\right)^{k+2} e^{i\ell\theta}, & \text{for } \frac{5}{2}\lambda \leq r \leq 5\lambda \\ 0, & \text{otherwise.} \end{cases}$$





(a) Scatterer ($q = -m = n^2 - 1$)



(b) Near Field Intensity ($|u|^2$)

Figure 5.1: Visualizations for Radially Layered Scatterer – Diameter = 4λ

(a) Scatterer ($q = -m = n^2 - 1$)(b) Near Field Intensity ($|u|^2$)Figure 5.2: Visualizations for Off-Center Disc – Centered at $(1\lambda, 0)$, Diameter = 1λ

M	N_r	N	Memory	Iter.	Time	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
60	68	8K	15M	71	27s	3.24e-2		2.07e-2	
120	68	17K	30M	71	52s	4.69e-3	6.91	1.95e-3	10.62
240	68	33K	62M	71	109s	6.23e-4	7.53	2.32e-4	8.41
480	68	66K	120M	71	228s	9.71e-5	6.42	2.87e-5	8.08
960	68	133K	238M	71	458s	1.04e-5	9.34	3.53e-6	8.13
1920	68	238K	483M	71	898s	1.45e-6	7.17	3.83e-7	9.22



Table 5.3: Convergence Rate for Discontinuous Scatterer – Diameter = 10λ



Note that for each integer k , this series becomes a sine or cosine series with real coefficients. If $k = -1$, m is discontinuous and piecewise smooth as a function of θ . Further, for any integer $k \geq 0$, $m \in C^{k,\alpha} \cap C_{pw}^\infty$ as a function of θ . The three examples that follow illustrate the convergence of the method for $k = -1, 0, 1$. Because these scatterers are fully inhomogeneous, their size in terms of interior wavelengths is not well defined. Note, however, that each scatterer has a reasonably large incident wavelength diameter of 10λ .



The results for $k = -1$ are found in Figure 5.3 and Table 5.3. The predicted second-order convergence in the near field is exceeded and the third-order convergence in the far field is readily observed. The results for $k = 0$ are found in Figure 5.4 and Table 5.4. In this case, the predicted third-order convergence in the near field and fifth-order convergence in the far field are both matched quite precisely. This example clearly illustrates the interesting jump in the far field convergence rate from third-order for a discontinuous scatterer to *fifth-order* for a $C^{0,\alpha}$ scatterer. Here the far field converges to the precision of the solution before the near field does. Hence, the last two runs result in no improvement in the far field accuracy. Finally, the results for $k = 1$ are found in Figure 5.5 and Table 5.5. In this case, the predicted fourth- and seventh-order convergence rates in the near and far fields, respectively, are obtained if not exceeded. However, because convergence is so rapid, it is difficult to observe a definite pattern, especially in the far field convergence. In each of these cases, we note that the memory and time scale appropriately with M . Finally, we mention that all of these results required less than 20 minutes and less than 700MB of memory.

5.1.2 Convergence in Radial Integration

In this section, we demonstrate the high-order convergence of the radial integration method described in Section 3.1.2. The degree of accuracy in the radial integration is determined by the number of subintervals N_i and the number N_c of Chebyshev points per subinterval

(a) Scatterer ($q = -m = n^2 - 1$)(b) Near Field Intensity ($|u|^2$)Figure 5.3: Visualizations for Discontinuous Scatterer – Diameter = 10λ

(a) Scatterer ($q = -m = n^2 - 1$)(b) Near Field Intensity ($|u|^2$)Figure 5.4: Visualizations for $C^{0,\alpha}$ Scatterer – Diameter = 10λ

(a) Scatterer ($q = -m = n^2 - 1$)(b) Near Field Intensity ($|u|^2$)Figure 5.5: Visualizations for $C^{1,\alpha}$ Scatterer – Diameter = 10λ

M	N_r	N	Memory	Iter.	Time	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
60	102	12K	18M	44	23s	9.33e-4		7.06e-6	
120	102	25K	34M	44	50s	8.91e-5	10.47	1.30e-7	54.31
240	102	50K	68M	44	105s	1.15e-5	7.75	3.86e-9	33.68
480	102	99K	138M	44	212s	1.46e-6	7.88	1.17e-10	32.99
960	102	198K	275M	44	565s	1.83e-7	7.97	1.73e-11	Conv.
1920	102	396K	548M	44	1136s	1.98e-8	9.24	1.85e-11	Conv.

Table 5.4: Convergence Rate for $C^{0,\alpha}$ Scatterer – Diameter = 10λ

M	N_r	N	Memory	Iter.	Time	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
60	102	12K	19M	54	36s	2.16e-5		7.33e-9	
120	102	25K	39M	54	72s	4.81e-7	44.91	1.06e-11	691.51
240	102	50K	75M	54	160s	1.05e-8	45.81	4.50e-12	Conv.
480	102	99K	150M	54	331s	4.76e-10	22.06	4.52e-12	Conv.
960	102	198K	305M	54	561s	1.36e-11	35.0	4.61e-12	Conv.
1920	102	396K	609M	54	1172s	1.94e-12	Conv.	4.72e-12	Conv.

Table 5.5: Convergence Rate for $C^{1,\alpha}$ Scatterer – Diameter = 10λ

used to approximate $I_\ell(r)$ as well as the number N_g of Gaussian points used to compute the moment integrals (3.9) and (3.10). As mentioned in Section 3.1.2.2, we have found $N_g = 8$ to be sufficient in all of the examples, i.e., an increase in N_g does not lead to any increase in the solution accuracy. In this section, we examine the dependence of the convergence rate on the choice of N_c and N_i .

We demonstrate the convergence in two examples: 1) a disc with constant refractive index centered at the origin, and 2) a disc with constant refractive index centered away from the origin. As mentioned in the previous section, given an incident plane wave, the value of the scattered field can be computed analytically for these simple scatterers. For each set of values N_c and N_i , we compute the maximum near and far field errors. We choose the number of modes M and the number of angular points N_θ sufficiently large so that these errors are dominated by the error in the radial integration. The value of N_c determines the order of the radial integration and is held fixed while N_i is increased to obtain the desired accuracy in the computed solution. In these examples, we choose relatively small values of N_c , for which the convergence rates are easily observed. In practice, we use significantly larger values of N_c .

The results for the disc centered at the origin are given in Table 5.6. This disc has a diameter of 2λ and a refractive index $n = \sqrt{2}$. We fix the number of modes $M = 30$

N_i	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
4	0.760		1.39	
8	0.216	3.51	0.373	3.72
16	5.92e-2	3.66	9.47e-2	3.94
32	1.41e-2	4.18	2.38e-2	3.98
64	3.67e-3	3.85	5.95e-3	4.00
128	9.34e-4	3.93	1.49e-3	4.00

(a) $N_c = 2$

N_i	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
4	9.77e-2		6.30e-2	
8	1.24e-2	7.88	3.53e-3	17.9
16	1.98e-3	6.28	2.31e-3	1.52
32	3.83e-4	5.15	6.74e-4	3.43
64	8.73e-5	4.39	1.74e-4	3.86
128	2.05e-5	4.27	4.40e-5	3.97

(b) $N_c = 3$

N_i	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
4	1.15e-2		4.31e-3	
8	8.25e-4	18.0	6.50e-4	6.64
16	5.41e-5	15.3	4.59e-5	14.2
32	3.49e-6	15.5	2.95e-6	15.5
64	2.11e-7	16.5	1.86e-7	15.9
128	1.37e-8	15.4	1.16e-8	16.0

(c) $N_c = 4$

N_i	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
4	1.19e-3		2.51e-4	
8	4.61e-5	25.8	2.50e-5	10.1
16	1.85e-6	24.9	2.34e-6	10.7
32	9.69e-8	19.1	1.60e-7	14.7
64	5.58e-9	17.4	1.02e-8	15.7
128	3.28e-10	17.0	6.41e-10	15.9

(d) $N_c = 5$ Table 5.6: Convergence of Radial Integration for Disc Centered at Origin – 2λ Diameter

and the number of angular points $N_\theta = 64$. Table 5.6 shows the convergence rates for various values of N_c as N_i is increased. It is not difficult to see from these results that the radial integrals converge as $N_i^{-N_c}$ if N_c is even and as $N_i^{-(N_c-1)}$ if N_c is odd. Hence, we always choose N_c to be even. Furthermore, the computation of the radial integrals requires $\mathcal{O}(N_c^2 N_i) = \mathcal{O}(N_c N_r)$ operations. At the same time, this $\mathcal{O}(N_c^2)$ complexity is not observed until N_c becomes quite large. We choose N_c to balance the order of the convergence with the efficiency of the computation. In most examples, we have found that $N_c = 16$ strikes this balance quite well.

The results for the disc centered at $(1\lambda, 0)$ are given in Table 5.7. In this case, the disc has a diameter of 1λ and a refractive index of $n = \sqrt{2}$. As in the previous example, we fix the number of modes and examine the convergence rates as N_i increases for various values of N_c . In this case, since the inhomogeneity m is a discontinuous function on the domain of integration, our computed solution converges as $\mathcal{O}(M^{-2})$ in the near field and $\mathcal{O}(M^{-3})$ in the far field. In the case of the disc centered at the origin, on the other hand, we observe super-algebraic convergence since the inhomogeneity is C^∞ on the domain of integration. Because the solution converges more slowly in this example, we must choose a much larger value of M (and N_θ) so that the dependence of the error on the radial integration can be observed. We choose $M = 1920$ and $N_\theta = 8192$.

As discussed in Section 3.1.2, since the disc centered away from the origin is a discontinuous scatterer, it produces singularities in $I_\ell(r)$, which must be resolved in order to obtain high-order convergence. It is not difficult to show that a disc with radius R , centered at $(d, 0)$ has Fourier coefficients

$$m_\ell(r) = \begin{cases} \frac{\bar{m}\alpha(r)}{\pi}, & \text{if } \ell = 0 \\ \frac{\bar{m}\sin(\ell\alpha(r))}{\pi\ell}, & \text{if } \ell \neq 0, \end{cases} \quad (5.1)$$

for $d-R \leq r \leq d+R$ and where $\bar{m} = 1-n^2 = 1-2 = -1$ and $\alpha(r) = \arccos\left[\frac{r^2-(d^2-R^2)}{2dR}\right]$. It is also not difficult to see that $m_\ell(r)$ and, hence, $I_\ell(r)$ (see (2.9)) has square-root singularities at $r = d - R$ and $r = d + R$. Therefore, we make use of the change of variables (3.4) to resolve these singularities. In Table 5.7, we give the near and far field errors obtained with and without this change of variables. We only present results for even values of N_c .

The convergence rates for this discontinuous scatterer are more erratic. It is clear that the change of variable increases the convergence rate rather significantly. However, it is difficult to ascertain a precise convergence rate from these numerical examples. Furthermore, although the convergence rate with the change of variable is high-order (at least in average), there seems to be little advantage in choosing a value of $N_c \geq 4$, i.e., there is only slight improvement in the results for $N_c = 8$ as compared with the results for $N_c = 4$. This is likely due to the rapid radial oscillations in the slowly decaying Fourier coefficients $m_\ell(r)$ (see (5.1)). The more rapid convergence rates shown in Table 5.6 are only obtained when these oscillations are sufficiently resolved by the radial discretization. Hence, when M is large, as in this example, we observe a more moderate, albeit still high-order, convergence rate for the radial integration.

5.1.3 Complex Scatterers and Preconditioning

In this section, we illustrate the power and flexibility of the two-dimensional method by considering three rather complex scatterers. The first two contain geometric singularities, corners and cusps, respectively. The last example contains smooth indentations and protrusions in a constant background, providing an example of a truly inhomogeneous, but smooth, medium. In each example, we present the maximum near and far field errors as we increase N_i , M and N_θ while fixing $N_c = 16$ and $N_g = 8$. Furthermore, we present the

N_i	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
1	2.26		4.94	
2	1.50	1.51	2.01	2.46
4	0.486	3.07	0.552	3.89
8	0.156	3.12	0.180	3.06
16	4.49e-2	3.46	5.58e-2	3.23
32	1.27e-2	3.54	1.71e-2	3.27
64	4.33e-3	2.93	5.50e-3	3.10
128	1.29e-3	3.35	1.94e-3	2.84
256	4.73e-4	2.73	6.78e-4	2.86

(a) $N_c = 2$ without change of variable

N_i	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
1	2.95		5.69	
2	1.94	1.52	2.58	2.21
4	0.675	2.88	1.02	2.53
8	0.200	3.37	0.262	3.88
16	5.11e-2	3.92	6.85e-2	3.82
32	1.41e-2	3.62	1.73e-2	3.96
64	3.55e-3	3.98	4.34e-3	3.99
128	8.41e-4	4.22	1.09e-3	4.00
256	1.89e-4	4.46	2.72e-4	4.00

(b) $N_c = 2$ with change of variable

N_i	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
1	1.84		1.57	
2	0.155	11.84	0.153	10.30
4	2.82e-2	5.50	2.18e-2	7.02
8	6.60e-3	4.27	6.52e-3	3.34
16	4.18e-3	1.58	2.07e-3	3.15
32	7.10e-4	5.89	6.90e-4	3.00
64	3.96e-4	1.79	2.37e-4	2.91
128	8.72e-5	4.54	8.28e-5	2.86

(c) $N_c = 4$ without change of variable

N_i	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
1	1.88		4.07	
2	0.795	2.36	0.550	7.39
4	3.84e-2	20.68	1.09e-2	50.56
8	2.97e-3	12.95	4.28e-4	25.45
16	4.28e-4	6.94	4.65e-5	9.20
32	2.48e-5	17.25	3.43e-6	13.55
64	7.50e-6	3.30	2.44e-7	14.05
128	4.11e-7	18.24	1.48e-8	16.43

(d) $N_c = 4$ with change of variable

N_i	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
1	2.73e-2		1.72e-2	
2	9.35e-3	2.92	5.32e-3	3.23
4	4.16e-3	2.25	1.63e-3	3.26
8	1.13e-3	3.70	5.60e-4	2.91
16	4.72e-4	2.38	1.96e-4	2.86
32	1.81e-4	2.61	6.94e-5	2.83
64	2.89e-5	6.26	2.46e-5	2.82

(e) $N_c = 8$ without change of variable

N_i	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
1	0.551		0.325	
2	9.65e-3	2.36	9.94e-4	327.3
4	1.37e-3	20.68	4.46e-5	22.27
8	1.76e-4	12.95	4.80e-6	9.31
16	5.23e-5	6.94	5.49e-7	8.73
32	4.27e-6	17.25	6.15e-8	8.94
64	3.45e-7	3.30	4.24e-9	14.50

(f) $N_c = 8$ with change of variableTable 5.7: Convergence of Radial Integration for Disc Centered at $(1\lambda, 0) - 1\lambda$ Diameter



number of GMRES iterations, the memory and the time required with and without the use of our preconditioner. These results were obtained on a 1.7GHz Pentium IV Xeon with 2GB of RAM.

The example of the square scatterer with constant refractive index in Figure 5.6 and Table 5.8 is another example of a discontinuous scatterer like the off-center disc. The length of the diagonal is 5λ and $n = 2$. Hence, in terms of interior wavelengths, the square has a diagonal length of 10λ . The implementation of this example requires no additional care to deal with the corners. We simply compute analytically the Fourier coefficients of m as a function of radius, $m_\ell(r)$.

In this example, we divide the radial integration domain into two intervals $[0, 5\lambda/2\sqrt{2}]$ and $[5\lambda/2\sqrt{2}, 5\lambda/2]$. On the first interval, m is constant. It is not too difficult to see that, on the second interval, the Fourier coefficients $m_\ell(r)$ have a square-root singularity at $r = 5\lambda/2\sqrt{2}$ from the right and, hence, we use the change of variable (3.4) on this interval. In Table 5.8, we list the number of subintervals in each of these intervals (as $N_i^{(1)}/N_i^{(2)}$) as well as the total number of radial points N_r . Finally, we precondition the resulting linear system with an approximate inhomogeneity $\tilde{m}$, which is equal to the true value of m on the first interval and equal to zero on the second interval. We have determined through experiment that this configuration produces the greatest decrease in the number of GMRES iterations required to obtain a given residual tolerance.

Note that the slight noise in the visualization of the refractive index in Figure 5.6(a) is caused by the Gibb's phenomenon in summing the truncated Fourier series. (This noise can be observed in the visualizations of the other discontinuous scatterers as well.) We emphasize that, as discussed in Section 2.3, this approximation of m by m^{2M} *does not* degrade the solution accuracy, but instead *improves* it by allowing efficient and *exact* computation of the Fourier coefficients of $m(r, \theta)v^M(r, \theta)$. As in previous examples, we observe second-order convergence in the near field and third-order convergence in the far field as predicted. Note that the time and memory scale linearly with the number of unknowns $N = (N_r + 1)(2M + 1)$.

The star-shaped scatterer in Figure 5.7 contains cusps and, hence, would likely pose difficulties for a boundary integral method. However, as with the square scatterer, no special care need be taken to treat this geometry. As before, we simply compute the Fourier coefficients for this scatterer analytically. This star-shaped scatterer is given by the region

(a) Scatterer ($q = -m = n^2 - 1$)(b) Near Field Intensity ($|u|^2$)Figure 5.6: Visualizations for Square Scatterer – Diagonal Length = 5λ

M	N_θ	N_i	N_r	Memory		Iter.		Time		ϵ_u^{nf}	ϵ_u^{ff}
60	256	2/2	68	13M	33M	52	218	8.56s	23.0s	6.15e-2	4.32e-2
120	512	2/3	85	31M	79M	56	226	21.4s	60.8s	1.22e-2	3.60e-3
240	1024	2/5	119	90M	228M	58	235	61.8s	181s	1.70e-3	3.21e-4
480	2048	2/10	204	311M	795M	62	238	218s	623s	4.25e-4	3.64e-5
960	4096	2/20	374	1183M		66		238s		6.72e-5	3.04e-6

Table 5.8: Convergence for Square Scatterer – Diagonal Length = 5λ

M	N_θ	N_i	N_r	Memory		Iter.		Time		ϵ_u^{nf}	ϵ_u^{ff}
30	128	1/2	51	6M	6M	32	73	3.16s	3.46s	0.792	0.581
60	256	1/4	85	14M	15M	35	77	7.76s	8.95s	0.112	9.34e-2
120	512	1/6	119	36M	42M	37	80	19.5s	23.7s	2.02e-2	1.28e-2
240	1024	1/10	187	113M	143M	40	82	61.3s	72.7s	3.87e-3	1.65e-3
480	2048	1/18	323	403M	543M	43	85	219s	261s	9.01e-4	2.06e-4
960	4096	1/34	595	1538M		45		882s		1.50e-4	2.13e-5


Table 5.9: Convergence for Star Scatterer – Diameter = 1λ

between the four unit discs centered at $(1, 1)$, $(1, -1)$, $(-1, 1)$ and $(-1, -1)$, respectively. It is scaled to have a diameter of 1λ and has a refractive index $n = 8$. Hence, in terms of interior wavelengths, this scatterer measures 8λ in diameter.


As with the square scatterer, we divide the radial integration domain into two intervals $[0, (\sqrt{2} - 1)/2\lambda]$ and $[(\sqrt{2} - 1)/2\lambda, 1/2\lambda]$. We again use the change of variable (3.4) to resolve the square-root singularity at the left endpoint of the second interval. In Table 5.9, we report the number of subintervals in each of these intervals as $N_i^{(1)}/N_i^{(2)}$. As with the square, we precondition with the value of m itself on the inner interval and with zero on the outer interval. In this example, the preconditioner yields only minor decreases in the time and memory. Generally, the preconditioner’s effectiveness is diminished for inhomogeneities that are poorly approximated by a piecewise constant, radially layered inhomogeneity (see Section 3.2.2).

The visualization of the near field intensity, Figure 5.7(b), clearly shows the relatively long wavelength of the field outside the scatterer as well as the relatively short wavelength of the field inside the scatterer. Table 5.9 demonstrates the second-order convergence in the near field and third-order convergence in the far field. We also observe the proper scaling of the memory and time with M , N_r and the number of GMRES iterations.

The final example is given in Figure 5.8. This scatterer is generated by adding (or subtracting) several smooth “bumps” with randomly generated heights, radii and centers



(a) Scatterer ($q = -m = n^2 - 1$)



(b) Near Field Intensity ($|u|^2$)

Figure 5.7: Visualizations for Star Scatterer – Diameter = 1λ

M	N_θ	N_i	N_r	Memory		Iter.		Time		ϵ_u^{nf}	ϵ_u^{ff}
30	128	3	51	5M	27M	12	439	1.38s	21.6s	1.38e-2	5.32e-3
60	256	7	119	14M	107M	15	445	5.69s	120s	2.22e-4	3.41e-5
120	512	12	204	47M	368M	20	457	22.5s	426s	5.24e-6	3.46e-7
240	1024	26	442	221M		23	468	131s		7.50e-8	2.83e-10


Table 5.10: Convergence for Bumpy Scatterer – Diameter = 5λ

to a constant refractive index in a disc. The constant background in this example has a diameter of 5λ and a refractive index $n = 2$ (interior wavelength diameter = 10λ). The bumps are simple smooth cylindrically symmetric functions as given by (4.5), which are then multiplied by a height (positive or negative). Since the background disc with constant refractive index is centered at the origin, we obtain a scatterer that is C^∞ at each radius as a function of the angular variable. Hence, we obtain rapid convergence in both the radial integration and the angular integration and thus, unlike the previous examples, we do not substitute m^{2M} for m , but rather compute $I_\ell(r)$ by means of direct trapezoidal rule integration with N_θ points (see Section 3.1.1).


Since this inhomogeneity is smooth in the integration domain, we need not resolve any singularities. Hence, we have only a single interval for the radial integration. We precondition by approximating the scatterer m by the background value $\tilde{m} = 1 - 2^2 = -3$. In this case, because of the accuracy of this preconditioning approximation as well as the relatively large size of the scatterer, Table 5.10 shows large decreases in the required time, memory and GMRES iterations when using the preconditioner.

5.2 Three-Dimensional Computational Examples

In this section, we demonstrate the $\mathcal{O}(N \log N)$ complexity as well as the high-order convergence rate of the three-dimensional method through three computational examples. First, we consider a piecewise-constant layered sphere scatterer, for which the analytical solution is known. Second, we consider a $5 \times 5 \times 5$ array of smooth potentials. Finally, we consider a rather complex scatterer containing inhomogeneous regions as well as geometric singularities (discontinuities, corners, and cusps). Because of the large number of unknowns required, some of the examples are computed in parallel. We list the wall-clock time required as well as the number of processors used as $T(s) \times P$. Thus, although not one of the primary goals of this paper, one can thereby obtain some idea of the parallel performance of the



(a) Scatterer ($q = -m = n^2 - 1$)



(b) Near Field Intensity ($|u|^2$)

Figure 5.8: Visualizations for Bumpy Scatterer – Diameter = 5λ

method. These examples were computed on 1–16 dual processor nodes (1–32 processors) of a Beowulf cluster. Each node consisted of two 1.7GHz Pentium IV Xeon processors and 1 GB of RAM. These nodes were connected via Myrinet.

To isolate the performance of the matrix-vector product, we hold the number of iterations of the linear solver fixed. Hence, given the $\mathcal{O}(N \log N)$ complexity, a doubling of the unknowns in each direction should yield roughly a factor of eight increase in the time required. This pattern is easily observed in each of the examples.



The convergence rate is illustrated by computing the maximum error in the near field solution (ϵ_u^{nf}) as well as in the far field solution (ϵ_u^{ff}). Given the near field solution u , the far field solution u_∞ is computed by means of the following integral [17, p. 223]

$$u_\infty(\hat{x}) = -\frac{\kappa^2}{4\pi} \int_{\Omega_{[a,b]}} e^{-i\kappa\hat{x}\cdot y} m(y) u(y) dy,$$

where $\hat{x}$ is a point on the unit sphere. We compute this integral with high-order accuracy by means of the trapezoidal rule (after replacing m with $\tilde{m}$ and $\Omega_{[a,b]}$ with $\Omega_{[\tilde{a},\tilde{b}]}$). To illustrate the convergence rates more clearly, we also display the ratios of errors between successive levels of discretization. When possible, we compare the computed solution with the analytical solution; in the other cases, we compare the computed solution with the solution at a finer discretization.

The results of the computations with the piecewise-constant layered sphere are presented in Figures 5.9 and 5.10 as well as in Table 5.11. As discussed in Section 4.3, m is replaced by $\tilde{m}$ to obtain high-order convergence in this example (see the columns labeled ϵ_u^s and $\epsilon_{u_\infty}^s$). Although the convergence rates jump around somewhat, the near field solution seems to converge as h^2 while the far field solution converges (at least in geometric average) as h^3 . For the purpose of comparison, we have also included the convergence rates that one observes when m is *not* replaced by $\tilde{m}$ (columns labeled ϵ_u^d and $\epsilon_{u_\infty}^d$). In this case, the convergence rates are significantly slower in both the near and the far fields, as expected. Furthermore, even at rather coarse discretizations that yield relatively low accuracy, the high-order method is clearly superior.

The results of the second example, the $5 \times 5 \times 5$ array of smooth inhomogeneous scatterers (potentials), are presented in Figures 5.11 and 5.12 as well as Table 5.12. Each of these scatterers is given by a spherically symmetric partition of unity function of the form (4.5).

(a) Scatterer ($q = n^2 - 1$)(b) Near Field Intensity ($|u|^2$)Figure 5.9: Visualizations for Layered Sphere – $\kappa a = 4$




Figure 5.10: Far Field Intensity ($|u_\infty|^2$) for Layered Sphere – $\kappa a = 4$

N	$T(s) \times P$	$\epsilon_u^{d,nf}$	$\epsilon_u^{s,nf}$	Ratio	$\epsilon_u^{d,ff}$	$\epsilon_u^{s,ff}$	Ratio
$10 \times 10 \times 10$	2.15×1	1.47	0.250		0.647	0.146	
$20 \times 20 \times 20$	15.6×1	0.614	2.25(-2)	11.1	0.314	4.56(-3)	31.95
$40 \times 40 \times 40$	125×1	0.167	5.70(-3)	3.96	8.42(-2)	9.55(-4)	4.78
$80 \times 80 \times 80$	1119×1	5.40(-2)	1.48(-3)	3.86	2.83(-2)	5.43(-5)	17.57
$160 \times 160 \times 160$	475×32	2.38(-2)	2.38(-4)	6.20	4.12(-3)	7.11(-6)	7.64

Table 5.11: Convergence for Layered Sphere – $\kappa a = 4$

N	$T(s) \times P$	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
$10 \times 10 \times 10$	5.65×32	3.70		43.0	
$20 \times 20 \times 20$	6.39×32	1.35	2.73	10.6	4.05
$40 \times 40 \times 40$	15.1×32	4.80(-2)	28.2	8.66(-2)	122
$80 \times 80 \times 80$	107×32	8.28(-3)	5.79	4.47(-2)	1.94
$160 \times 160 \times 160$	875×32	6.48(-5)	128	7.76(-5)	576

Table 5.12: Convergence for Array of Smooth Scatterers – $5\lambda \times 5\lambda \times 5\lambda$

The centers of the scatterers are separated by 1λ and each has a radius of $1/2\lambda$. Therefore, the support of the array is contained in a $5\lambda \times 5\lambda \times 5\lambda$ box. Since $m \in C^\infty$, we do not replace m by $\tilde{m}$, but instead integrate directly with the trapezoidal rule. As expected, we observe a very rapid convergence rate. Likely due to the heterogeneity in the scatterer, which must be sufficiently resolved, we observe a sharp dip in the convergence rate for $N = (40, 40, 40)$ followed by an immediate resumption of the rapid convergence.

The results of our final computational example are contained in Figures 5.13, 5.14 and 5.15 as well as Table 5.13. This scatterer is created by adding together a cube, two spheres, two layered spheres, and six smooth, inhomogeneous scatterers (similar to those in the previous example) to create a rather complex scatterer. More precisely, we first center at the origin a cube with $m = -1$ and sides of length 4. Next, two unit spheres with $m = +1$ and centered at $(0, -1, 0)$ and $(0, 1, 0)$, respectively, are added to the interior of this cube. This essentially cuts two spheres out of the cube. Then, two layered spheres with unit radii are placed tangent to the two faces of the cube that are orthogonal to the y -axis. Finally, two sets of three smooth, inhomogeneous scatterers are placed along the faces of the cube that are orthogonal to the z -axis. Figures 5.13 and 5.14(a) display this structure with two orthogonal slices of the scatterer.

This scatterer contains inhomogeneous regions, discontinuities, corners, and cusps. Despite these geometric singularities, we substitute m by

$$\tilde{m}(y) = p_{m_1}(y)m_1^F(y) + p_{m_2}(y)m_2^F(y) + \cdots,$$

where m_j^F is the truncated Fourier series for the j^{th} discontinuous component of the scatterer. (As in the previous example, we do not replace the C^∞ components of the scatterer by their Fourier-smoothed approximations.) One must only compute the Fourier coefficients of each of the individual discontinuous components of the scatterer, i.e., the cube, the spheres



(a) Scatterer ($q = n^2 - 1$)(b) Near Field Intensity ($|u|^2$)Figure 5.11: Visualizations for Array of Smooth Scatterers – $5\lambda \times 5\lambda \times 5\lambda$



Figure 5.12: Far Field Intensity ($|u_\infty|^2$) for Array of Smooth Scatterers – $5\lambda \times 5\lambda \times 5\lambda$






Figure 5.13: Complex Scatterer – Close-Up Image of Scatterer

and the layered sphere—the corners, cusps and inhomogeneities present *no additional difficulties*. We emphasize that no additional consideration of the scatterer geometry is required in our three-dimensional approach, whereas the two-dimensional approach required identification and resolution of singularities in the Fourier coefficients $(mu)_\ell(r)$.

Note that because of the asymmetry in the scatterer, we do not use the same number of discretization points in each direction. With $\kappa = 4$, the scatterer has an interior wavelength size (as measured along the y -axis) of approximately 7.2λ . We again observe high-order convergence in the near and far fields.

This illustrates one of the most powerful features of this method. We can treat quite complicated scatterers simply by adding the Fourier coefficients of the component pieces; no additional effort is needed to handle singularities that appear when these components are added. Furthermore, unlike other methods, the scatterer does not need to be “remeshed” as we refine the discretization. As demonstrated in this example, the high-order accuracy is maintained even for such complicated scatterers.

(a) Scatterer ($q = n^2 - 1$)(b) Near Field Intensity ($|u|^2$)Figure 5.14: Visualizations for Complex Scatterer – Size $\approx 2.5\lambda \times 5\lambda \times 2.5\lambda$




Figure 5.15: Far Field Intensity ($|u_\infty|^2$) for Complex Scatterer – Size $\approx 2.5\lambda \times 5\lambda \times 2.5\lambda$

N	$T(s) \times P$	ϵ_u^{nf}	Ratio	ϵ_u^{ff}	Ratio
$8 \times 16 \times 8$	4.39×1	3.98		13.6	
$16 \times 32 \times 16$	10.8×4	0.554	7.17	1.45	9.37
$32 \times 64 \times 32$	59.5×4	2.99(-2)	18.5	7.02(-2)	20.6
$64 \times 128 \times 64$	96.0×32	3.37(-3)	8.87	1.03(-3)	68.1
$128 \times 256 \times 128$	781×32	4.04(-4)	8.34	8.64(-5)	11.9

Table 5.13: Convergence for Complex Scatterer – Size $\approx 2.5\lambda \times 5\lambda \times 2.5\lambda$

Chapter 6

Conclusions

In this thesis, we have introduced a new, fast, high-order method for scattering by inhomogeneous media in three dimensions. This approach was motivated in part by the fast, high-order method in two dimensions introduced in [13]. In an attempt to allay the controversy generated by the claim of high-order accuracy, we proved, in Chapter 2, that this method indeed achieves high-order accuracy even in the case of discontinuous scatterers—yielding, in this case, second-order convergence in the near field and third-order convergence in the far field. We emphasized the interesting dependence of the convergence rate on the regularity of the scatterer, i.e., the convergence rates in the far field jump from third-order for discontinuous scatterers to fifth-order for $C^{0,\alpha}$ scatterers and to seventh-order for $C^{1,\alpha}$ scatterers (assuming they are also piecewise smooth).

In Chapter 3, we presented several improvements to the numerical implementation of this two-dimensional method. In particular, we achieved increased efficiency and stability through a new Chebyshev-based radial integration scheme and a new preconditioner. Also, the new method for computing scaled Bessel functions (see Appendix B) proved to be of great practical importance.

The high-order accuracy in the two-dimensional method is based on high-order trapezoidal rule integration and Fourier approximation of smooth and periodic functions. These ideas motivated the development of the three-dimensional method, as introduced in Chapter 4. By decomposing the Green’s function into a smooth part with infinite support and a singular part with compact support, we were able to make use of both high-order trapezoidal rule integration and high-order Fourier approximation in computing the required convolutions. The rather counterintuitive method of Fourier smoothing played a central role in achieving high-order accuracy for discontinuous scatterers.

In Chapter 5, we demonstrated the efficiency and high-order accuracy of these numerical methods through several computational examples. These examples served to verify the theoretical convergence rates of the two-dimensional method. We also demonstrated the high-order convergence of the radial integration scheme. We illustrated the effectiveness of the preconditioner. Finally, we presented results from parallel runs of the three-dimensional method. In particular, we sought to emphasize the power and versatility of the three-dimensional method in constructing complicated scatterers for which high-order accurate solutions can then be computed.

Many interesting problems still remain. Perhaps one of the most important research problems is the development of a preconditioner for the three-dimensional case. The dramatic growth of the number of required iterations with problem size remains one of the most fundamental obstacles to solving realistic problems of hundreds or even thousands of wavelengths in size.

A related issue concerns the scaling of the unknowns with the wavelength. In three dimensions, doubling the frequency requires a doubling of the grid points in each direction, yielding a factor of *eight* increase in unknowns. This doubling of discretization points is required to resolve the highly oscillatory fields. If, instead, one could factor out the dominant, highly oscillatory modes of the solution so that one would not need to explicitly resolve them, then the remaining smoothly varying function could be discretized with relatively few points. A similar approach, often called enveloping, is taken in paraxial approximations since the dominant propagation direction is known [7, 21, 23]. Also, previous work focusing on the application of multigrid to scattering problems [11, 40] made use of similar ideas.

Since FFT-based methods require, in general, equally spaced discretization points, it is not straightforward to implement adaptive discretization strategies. At the same time, there are certainly problems for which the availability of an adaptive method could save considerable time and memory. Hence, we want to consider adaptive approaches such as (smoothly) decomposing the scatterer into several pieces, each of which has a different discretization level. Interactions between these pieces could be computed through use of equivalent sources (see [12] for more details).

Finally, we are interested in extending these methods to other specific application fields such as materials science, particularly electron diffraction. In electron diffraction, the

Lippman-Schwinger integral equation (1.6) arises from the Schrödinger equation [27, p. 141]. In spite of similarities with the problems considered in this text, such materials science problems give rise to important differences. In particular, a crystal lattice, because of its size relative to the lattice spacings, is considered an infinite periodic structure. Extension of our methods to this problems is not straightforward, but the benefits of a fast, high-order accurate method would, we believe, prove quite useful in this field.

We believe research along these lines would lead to significantly improved capabilities in computational scattering. It is our hope that such methods in computational scattering will play an important role in advancing scientific understanding and engineering capabilities in a variety of fields.

Appendix A

Technical Lemmas

Here we present the proofs of two technical lemmas, Lemmas 2.1 and 2.6, that play an important role in the convergence rate results of Chapter 2. Neither of these results is very difficult to obtain, but since the technical nature of the proofs obscures the essence of the convergence rate results, we place them here in the appendix.

A.1 Bound on Fourier Coefficient of Green's Function

To prove that solutions to the approximate integral equation (2.3) exist and to prove the convergence rates, we need a bound on the decay rate of the Fourier coefficients of the Green's function, $\mathcal{J}_\ell(a, r)$ as defined in (2.1). This decay rate is given in Lemma 2.1. However, although the proof of this lemma is not difficult, it is somewhat technical. Hence, we first derive two simple bounds that aid in the proof.

Lemma A.1. *Given $\alpha > 0$ and a positive integer k , there exists a constant $C > 0$ such that*

$$\frac{\alpha^\ell}{\ell!} \leq \frac{C}{\max\{1, \ell^k\}},$$

for all integers $\ell \geq 0$.

Proof. If we let $\ell = 0$, then we require that $C \geq 1$ for the lemma to hold. Now consider $\ell \geq 1$. Equivalently we require that

$$\ell \log \alpha - \sum_{p=2}^{\ell} \log \ell + k \log \ell \leq \log C.$$

We bound the sum by the following integral

$$\sum_{p=2}^{\ell} \log p \geq \int_1^{\ell} \log x dx = \ell \log \ell - \ell + 1.$$

Thus, we require that

$$\ell \left[\log \alpha + 1 + k \frac{\log \ell}{\ell} \right] - \ell \log \ell \leq \log C + 1.$$

It is not difficult to see that the left-hand side of this equation is bounded from above for $\ell \geq 1$ since $\ell \log \ell$ eventually dominates. Hence, we can choose a constant $C \geq 1$ such that the result holds. $\square$

According to [2, p. 362], for all integers $\ell \geq 0$ and for any real, non-negative z ,

$$|J_{\ell}(z)| \leq \frac{1}{\ell!} \left(\frac{z}{2} \right)^{\ell}. \quad (\text{A.1})$$

The following lemma provides a similar bound for $|Y_{\ell}(z)|$.

Lemma A.2. *For all integers $\ell \geq 1$ and for any real, non-negative z ,*

$$|Y_{\ell}(z)| \leq \frac{(\ell-1)!}{\pi} \left(\frac{z}{2} \right)^{-\ell} e^{\left(\frac{z}{2}\right)^2} + \frac{2}{\pi \ell!} \left| \log \left(\frac{z}{2} \right) \right| \left(\frac{z}{2} \right)^{\ell} + \frac{2}{\pi} \left(\frac{z}{2} \right)^{\ell} e^{\left(\frac{z}{2}\right)^2}. \quad (\text{A.2})$$

For $\ell = 0$,

$$|Y_{\ell}(z)| \leq \frac{2}{\pi} \left| \log \left(\frac{z}{2} \right) \right| + \frac{2}{\pi} e^{\left(\frac{z}{2}\right)^2}.$$

Proof. By [10, p. 51], $Y_{\ell}(z)$ for any non-negative integer ℓ is given by

$$Y_{\ell}(z) = \frac{2}{\pi} J_{\ell}(z) \log \left(\frac{z}{2} \right) - \frac{1}{\pi} \sum_{k=0}^{\ell-1} \frac{(\ell-k-1)!}{k!} \left(\frac{z}{2} \right)^{2k-\ell} - \frac{1}{\pi} \sum_{k=0}^{\infty} \frac{\psi(\ell+k) + \psi(k)}{(-1)^k k! (k+\ell)!} \left(\frac{z}{2} \right)^{2k+\ell},$$

where $\psi(0) = -\gamma \approx -0.5772$ and $\psi(k) = -\gamma + \sum_{j=1}^k \frac{1}{j}$ for $k \geq 1$.

For $\ell \geq 1$,

$$\begin{aligned} \sum_{k=0}^{\ell-1} \frac{(\ell-k-1)!}{k!} \left(\frac{z}{2}\right)^{2k} &\leq (\ell-1)! \sum_{k=0}^{\infty} \frac{1}{k!} \left[\left(\frac{z}{2}\right)^2\right]^k \\ &\leq (\ell-1)! e^{\left(\frac{z}{2}\right)^2}. \end{aligned}$$

Also note that for $k \geq 1$, $|\psi(0)| \leq 1$ and

$$\begin{aligned} 0 \leq \psi(k) &\leq -\gamma + \sum_{j=1}^k 1 \\ &= -\gamma + k \\ &\leq k. \end{aligned}$$

Hence, $|\psi(k)| \leq \max\{1, k\}$.

Now observe that $|\psi(\ell+k) + \psi(k)| \leq 2 \max\{1, \ell+k\}$ and

$$\frac{|\psi(k)|}{k!} \leq 1,$$

for $k \geq 0$. Therefore, for $\ell \geq 0$

$$\begin{aligned} \sum_{k=0}^{\infty} \frac{|\psi(\ell+k) + \psi(k)|}{k!(k+\ell)!} \left(\frac{z}{2}\right)^{2k} &\leq 2 \sum_{k=0}^{\infty} \frac{1}{k!} \left(\frac{z}{2}\right)^{2k} \\ &\leq 2e^{\left(\frac{z}{2}\right)^2}. \end{aligned}$$

Hence, by these individual bounds and (A.1) we obtain the desired result. $\square$

With these two results in hand, we turn to the proof of the main lemma.

Proof of Lemma 2.1. First note that

$$\begin{aligned} \int_{R_0}^{R_1} |\mathcal{J}_\ell(a, r)| r dr &= |H_\ell^1(\kappa a)| \int_{R_0}^a |J_\ell(\kappa r)| r dr + |J_\ell(\kappa a)| \int_a^{R_1} |H_\ell^1(\kappa r)| r dr \\ &\leq |J_\ell(\kappa a)| \int_{R_0}^{R_1} |J_\ell(\kappa r)| r dr + |J_\ell(\kappa a)| \int_a^{R_1} |Y_\ell(\kappa r)| r dr \\ &\quad + |Y_\ell(\kappa a)| \int_{R_0}^a |J_\ell(\kappa r)| r dr \\ &\leq I_{J,J} + I_{J,Y} + I_{Y,J}, \end{aligned}$$

where

$$\begin{aligned} I_{J,J} &= |J_\ell(\kappa a)| \int_0^{R_1} |J_\ell(\kappa r)| r dr, \\ I_{J,Y} &= |J_\ell(\kappa a)| \int_a^{R_1} |Y_\ell(\kappa r)| r dr, \\ I_{Y,J} &= |Y_\ell(\kappa a)| \int_0^a |J_\ell(\kappa r)| r dr. \end{aligned}$$

Note that $|J_{-\ell}(z)| = |(-1)^\ell J_\ell(z)| = |J_\ell(z)|$ and similarly $|Y_{-\ell}(z)| = |Y_\ell(z)|$. Hence, it suffices to bound these integrals for $\ell \geq 0$.

We use (A.1) as well as Lemmas A.1 and A.2 to bound each of these integrals. For $\ell \geq 0$

$$\begin{aligned} I_{J,J} &\leq \frac{1}{(\ell!)^2} R_1^2 \left(\frac{\kappa R_1}{2} \right)^{2\ell} \\ &\leq \frac{C_{J,J}}{\max\{1, \ell^2\}}, \end{aligned}$$

where the last inequality follows from Lemma A.1 and $C_{J,J}$ depends only on κ and R_1 .

The bound for $I_{J,Y}$ consists of three parts from each of the three terms in (A.2). For $\ell > 2$,

$$\begin{aligned} \frac{\kappa^\ell}{2^\ell \ell!} a^\ell \int_a^{R_1} \frac{2^\ell (\ell-1)!}{\pi \kappa^\ell} r^{-\ell+1} e^{(\frac{\kappa r}{2})^2} dr &\leq \frac{R_1^2}{\pi \ell (\ell-2)} e^{(\frac{\kappa R_1}{2})^2} \left(\left(\frac{a}{R_1} \right)^2 - \left(\frac{a}{R_1} \right)^\ell \right) \\ &\leq \frac{C_{J,Y}^{(1)}}{\max\{1, \ell^2\}}. \end{aligned}$$

For $\ell = 0$ this term does not appear in (A.2) and for $\ell = 1, 2$, a similar argument yields the same bound. Continuing with the next term in (A.2)

$$\begin{aligned} \frac{(\kappa a)^\ell}{2^\ell \ell!} \int_a^{R_1} \frac{2\kappa^\ell}{\pi 2^\ell \ell!} \left| \log \left(\frac{\kappa r}{2} \right) \right| r^{\ell+1} dr &\leq \left(\frac{\kappa R_1}{2} \right)^{2\ell} \frac{2R_1}{\pi (\ell!)^2} \int_0^{R_1} \left| \log \left(\frac{\kappa r}{2} \right) \right| dr \\ &\leq \frac{C_{J,Y}^{(2)}}{\max\{1, \ell^2\}}, \end{aligned}$$

by Lemma A.1 and since $\int_0^{R_1} \left| \log \left(\frac{\kappa r}{2} \right) \right| dr$ is bounded. Finally, by similar arguments one can show that the third term is similarly bounded.

We also bound $I_{Y,J}$ by considering the three terms in (A.2). For $\ell \geq 1$,

$$\begin{aligned} \frac{2^\ell(\ell-1)!}{\pi\kappa^\ell} a^{-\ell} e^{(\frac{\kappa a}{2})^2} \int_0^a \frac{\kappa^\ell}{2^\ell \ell!} r^{\ell+1} dr &\leq \frac{1}{\pi\ell(\ell+2)} e^{(\frac{\kappa a}{2})^2} R_1^2 \\ &\leq \frac{C_{Y,J}^{(1)}}{\max\{1, \ell^2\}}. \end{aligned}$$

Continuing with the second term,

$$\begin{aligned} \frac{2}{\pi\ell!} \left| \log\left(\frac{\kappa a}{2}\right) \right| \left(\frac{\kappa a}{2}\right)^\ell \int_0^a \frac{1}{\ell!} \left(\frac{\kappa r}{2}\right)^\ell r dr &\leq \frac{1}{\pi(\ell!)^2} \left(\frac{\kappa R_1}{2}\right)^{2\ell} a^2 \left| \log\left(\frac{\kappa a}{2}\right) \right| \\ &\leq \frac{C_{Y,J}^{(2)}}{\max\{1, \ell^2\}}, \end{aligned}$$

since $a^2 \left| \log\left(\frac{\kappa a}{2}\right) \right|$ is bounded for $0 \leq a \leq R_1$. The bound for the last term also takes this form and can be obtained similarly. Hence,

$$I_{J,J} + I_{J,Y} + I_{Y,J} \leq \frac{C}{\max\{1, \ell^2\}},$$

for some constant $C > 0$ that depends only on κ and R_1 . $\square$

A.2 Bound on Integral Operator

In this section, we prove the bound on $K_\ell u^T$ as given in Lemma 2.6, which plays a primary role in our derivation of the convergence rates. As mentioned previously, some care is required to obtain tight bounds on the convergence rates. Hence, the proof of the lemma is somewhat technical.

Proof of Lemma 2.6. Define the annular region $A = \{(a, \phi) : R_0 \leq a \leq R_1\}$. Then (2.7), Lemma 2.4 and Theorem 2.5 imply that if $m \in C^{k,\alpha}(A) \cap C_{pw}^{k+2,\alpha}(A)$, then $u \in C^{k+2,\alpha}(A) \cap C_{pw}^{k+4,\alpha}(A)$ and there exists a constant $C > 0$ such that

$$\begin{aligned} \|K_\ell u^T\|_\infty &\leq \frac{C}{\max\{1, \ell^2\}} \sum_{|j|>M} \frac{1}{\max\{1, |\ell-j|^{k+2}\}} \frac{1}{\max\{1, |j|^{k+4}\}} \\ &= \frac{C}{\max\{1, \ell^2\}} \sum_{j>M} \frac{1}{j^{k+4}} \left(\frac{1}{(j-\ell)^{k+2}} + \frac{1}{(j+\ell)^{k+2}} \right) \\ &\leq \frac{2C}{\max\{1, \ell^2\}} \sum_{j>M} \frac{1}{j^{k+4}} \frac{1}{(j-|\ell|)^{k+2}}, \end{aligned} \tag{A.3}$$

for $\ell = -M, \dots, M$. This expression also holds for $m \in L^\infty(A) \cap C_{pw}^{1,\alpha}(A)$ with $k = -1$.

Clearly, we need only consider $\ell = 0, \dots, M$.

$$\begin{aligned} \sum_{j>M} \frac{1}{j^{k+4}} \frac{1}{(j-\ell)^{k+2}} &\leq \frac{1}{(M+1)^{k+4}} \frac{1}{(M+1-\ell)^{k+2}} + \int_{M+1}^{\infty} \frac{1}{x^{k+4}} \frac{1}{(x-\ell)^{k+2}} \\ &= \frac{1}{(M+1)^{k+4}} \frac{1}{(M+1-\ell)^{k+2}} \\ &\quad + \frac{1}{2k+5} \frac{1}{(M+1)^{2k+5}} F\left(k+2, 2k+5; 2k+6; \frac{\ell}{M+1}\right), \end{aligned}$$

where F is the hypergeometric function [2, p. 556]

$$\begin{aligned} F(a, b; c; z) &\equiv {}_2F_1(a, b; c; z) = \frac{\Gamma(c)}{\Gamma(b)\Gamma(c-b)} \int_0^1 t^{b-1} (1-t)^{c-b-1} (1-tz)^{-a} dt \\ &= \frac{\Gamma(c)}{\Gamma(a)\Gamma(b)} \sum_{n=0}^{\infty} \frac{\Gamma(a+n)\Gamma(b+n)}{\Gamma(c+n)} \frac{z^n}{n!}, \end{aligned}$$

for $|z| < 1$.

We need a few simple bounds on F in order to obtain the final result. It is easily verified that for positive integers a, b ,

$$F(a, b; b; z) = \sum_{n=0}^{\infty} \frac{(a+n-1)!}{(a-1)!n!} z^n = (1-z)^{-a}.$$

Using this result, if $a > 1$ and $b > 0$ are integers and $a \leq b+1$, then

$$\begin{aligned} F(a, b; b+1; z) &= \frac{b}{a-1} \sum_{n=0}^{\infty} \frac{a+n-1}{b+n} \frac{(a+n-2)!}{(a-2)!n!} z^n \\ &\leq \frac{b}{a-1} (1-z)^{-(a-1)}. \end{aligned}$$

Finally, for integers $a = 1$ and $b > 0$, we have

$$\begin{aligned} F(1, b; b+1; z) &= b \sum_{n=0}^{\infty} \frac{1}{b+n} z^n \\ &\leq 1 - b \log(1-z). \end{aligned}$$

Hence, for $k \geq 0$ and $\ell = 0, \dots, M$, we obtain

$$\begin{aligned} \sum_{j>M} \frac{1}{j^{k+4}} \frac{1}{(j-\ell)^{k+2}} &\leq \frac{1}{(M+1)^{k+4}} \frac{1}{(M+1-\ell)^{k+2}} + \frac{1}{k+1} \frac{1}{(M+1)^{k+4}} \frac{1}{(M+1-\ell)^{k+1}} \\ &\leq \frac{2}{M^{k+4}} \frac{1}{(M+1-\ell)^{k+1}}. \end{aligned} \quad (\text{A.4})$$

For $k = -1$ and $\ell = 0, \dots, M$,

$$\begin{aligned} \sum_{j>M} \frac{1}{j^3} \frac{1}{j-\ell} &\leq \frac{1}{(M+1)^3} \frac{1}{M+1-\ell} + \frac{1}{3(M+1)^3} + \frac{1}{(M+1)^3} \log \left(\frac{M+1}{M+1-\ell} \right) \\ &\leq \frac{1}{3M^3} + \frac{1}{M^2} \frac{1}{M+1-\ell}, \end{aligned} \quad (\text{A.5})$$

where we have used the fact that $\log x \leq x$. Combining (A.3) with (A.4) and (A.5) give the desired results for $m \in L^\infty(A) \cap C_{pw}^{1,\alpha}(A)$ and $m \in C^{k,\alpha}(A) \cap C_{pw}^{k+2,\alpha}(A)$, respectively. $\square$

Remark A.3. Finally, we note that the bounds for $m \in L^\infty(A) \cap C_{pw}^{1,\alpha}(A)$ and $m \in C^{0,\alpha}(A) \cap C_{pw}^{2,\alpha}(A)$ can be obtained more simply as follows.

$$\begin{aligned} \|K_\ell u^T\| &\leq \frac{2C}{\max\{1, \ell^2\}} \sum_{j>M} \frac{1}{j^{k+4}} \frac{1}{(j-|\ell|)^{k+2}} \\ &\leq \frac{2C}{\max\{1, \ell^2\}} \frac{1}{(M+1-|\ell|)^{k+2}} \sum_{j>M} \frac{1}{j^{k+4}} \\ &\leq \frac{\hat{C}}{\max\{1, \ell^2\}} \frac{1}{(M+1-|\ell|)^{k+2}} \frac{1}{M^{k+3}}, \end{aligned}$$

for some constant $\hat{C} > 0$. However, this simple bound does not capture the interesting convergence rate jumps in the results of Theorem 2.7 for $m \in C^{k,\alpha}(A) \cap C_{pw}^{k+2,\alpha}(A)$, $k \geq 1$.

Appendix B

Accurate and Efficient Computation of Scaled Bessel Functions

As explained in Section 3.1.2, the rapid decay of the $J_\ell(z)$ and the rapid growth of the $Y_\ell(z)$ as ℓ increases produces factors that underflow and overflow, respectively, but whose product is machine-representable. We overcome these and other related issues by computing scaled versions of the Bessel functions. The leading order asymptotic behavior of $J_\ell(z)$ and $Y_\ell(z)$ near the origin are given respectively by

$$J_\ell(z) \sim \frac{1}{\ell!} \left(\frac{z}{2}\right)^\ell$$

and

$$Y_\ell(z) \sim -\frac{(\ell-1)!}{\pi} \left(\frac{z}{2}\right)^{-\ell}$$

for $\ell > 0$. Thus, we scale the Bessel functions by their asymptotic representations, i.e., for $\ell > 0$

$$\tilde{J}_\ell(z) = \ell! \left(\frac{z}{2}\right)^{-\ell} J_\ell(z), \tag{B.1}$$

$$\tilde{Y}_\ell(z) = -\frac{\pi}{(\ell-1)!} \left(\frac{z}{2}\right)^\ell Y_\ell(z). \tag{B.2}$$

We use these scaled Bessel functions to compute products and quotients of J_ℓ , Y_ℓ and H_ℓ^1 in many combinations. For example, we can compute the product $J_\ell(z_1)Y_\ell(z_2)$ for $\ell > 0$ as

$$J_\ell(z_1)Y_\ell(z_2) = -\frac{1}{\pi\ell} \tilde{J}_\ell(z_1)\tilde{Y}_\ell(z_2). \tag{B.3}$$

Typically, the unscaled Bessel functions are computed by means of their recurrence relations

$$J_{\ell+1}(z) = \frac{2\ell}{z}J_{\ell}(z) - J_{\ell-1}(z), \quad (\text{B.4})$$

$$Y_{\ell+1}(z) = \frac{2\ell}{z}Y_{\ell}(z) - Y_{\ell-1}(z). \quad (\text{B.5})$$

Note that these recurrence relations are identical, signifying the fact that the underlying difference equation has linearly independent solutions $J_{\ell}(z)$ and $Y_{\ell}(z)$ for each z . Hence, in theory, given either $J_0(z)$ and $J_1(z)$ or $Y_0(z)$ and $Y_1(z)$, these recurrence relations allow us to obtain $J_{\ell}(z)$ or respectively $Y_{\ell}(z)$ for all ℓ . In practice, this works well for $Y_{\ell}(z)$. However, this procedure works for $J_{\ell}(z)$ only if $z > \ell$. Unfortunately, when computing $J_{\ell}(z)$ for $z < \ell$ (in which we are primarily interested), since $Y_{\ell}(z)$ is an exponentially growing solution, the recurrence relation is numerically unstable for increasing ℓ , i.e., the round-off error in $J_0(z)$ and $J_1(z)$ is rapidly amplified by the recurrence. On the other hand, this instability for increasing ℓ also implies that the recurrence relation is stable when computing $J_{\ell}(z)$ for *decreasing* ℓ , i.e., the round-off error in the starting values is quickly damped by the recurrence. Thus, we can begin the downward recurrence with two arbitrary values and the recurrence will rapidly converge to $\alpha J_{\ell}(z)$ where α is an unknown normalization constant. Finally, to obtain the correct values of $J_{\ell}(z)$, we compute α by means of the relationship

$$1 = J_0(z) + 2 \sum_{n=1}^{\infty} J_{2n}(z). \quad (\text{B.6})$$

This sum is not difficult to approximate accurately since $J_{\ell}(z)$ decays exponentially for $z < \ell$ (for more details, see [45, pp. 173–175]).

Hence, when computing $Y_{\ell}(z)$ for arbitrary z and ℓ or when computing $J_{\ell}(z)$ for $z > \ell$, we first compute $Y_0(z)$ and $Y_1(z)$ or, respectively, $J_0(z)$ and $J_1(z)$ (perhaps by means of an asymptotic expansion). We then use the recurrence relation with increasing ℓ to compute $Y_{\ell}(z)$ or $J_{\ell}(z)$. For $z < \ell$, on the other hand, we iterate several times through the downward recurrence with two arbitrary starting values to converge onto the correct sequence of $\alpha J_{\ell}(z)$. Then, in the process of the downward recurrence, we collect the sum (B.6) and normalize the sequence by the result.

To compute the scaled Bessel functions, we derive a new set of recurrence relations

related to the recurrence relations (B.4) and (B.5). From the definitions (B.1) and (B.2), and using (B.4) and (B.5), we obtain

$$\tilde{J}_{\ell+1}(z) = \ell(\ell+1) \left(\frac{2}{z}\right)^2 \left[\tilde{J}_{\ell}(z) - \tilde{J}_{\ell-1}(z) \right], \quad (\text{B.7})$$

$$\tilde{Y}_{\ell+1}(z) = \tilde{Y}_{\ell}(z) - \frac{1}{\ell(\ell+1)} \left(\frac{z}{2}\right)^2 \tilde{Y}_{\ell-1}. \quad (\text{B.8})$$

In this case, as a result of the scaling, neither $\tilde{J}_{\ell}(z)$ or $\tilde{Y}_{\ell}(z)$ grows (or decays) exponentially. It is not difficult to show by means of a few numerical experiments that the recurrence for $\tilde{J}_{\ell}$ is unstable and the recurrence for $\tilde{Y}_{\ell}$ is stable for increasing ℓ . Hence, after computing $Y_0(z)$ and $Y_1(z)$ as done previously, we then scale these values and use the recurrence relation (B.8) to compute $\tilde{Y}_{\ell}(z)$. To compute $\tilde{J}_{\ell}(z)$, we use a downward recurrence and the normalization sum (B.6) as before.

The implementation of this algorithm involves only relatively simple modifications to any existing algorithm for computing Bessel functions $J_{\ell}(z)$ or $Y_{\ell}(z)$. In our application, we modified the Fortran77 routines *rybesl* and *rybesl*, which one can easily obtain from the Netlib repository [1].

Appendix C

High-Order Evaluation of Fourier Integrals

Given a smooth, compactly supported, real-valued function $g(t)$ for $t \in \mathbb{R}$, we seek to compute the integral

$$I(\omega) = \int_a^b g(t) e^{i\omega t} dt$$

for various values of $\omega \in [\omega_{min}, \omega_{max}]$. Clearly, since $g(t)$ is real-valued, $I(-\omega) = \overline{I(\omega)}$ and, therefore, we may restrict our attention to $\omega_{min} \geq 0$.

We present a modified version of the method suggested in [45, pp. 577–584]. Through appropriate combinations of Lagrange interpolating polynomials of order q , we obtain a high-order approximation of $g(t)$. In particular, there exist piecewise smooth interpolating polynomials $\psi(s)$ of order q where $-q \leq s \leq q$ such that $\psi(0) = 1$ and $\psi(s) = 0$ for integer values $s = -q, \dots, q$. To further simplify the approach, we consider only even functions $\psi(s)$. (We describe specific choices of $\psi(s)$ for $q = 2, 4$ in the following sections.) Thus, we can construct a high-order approximation of $g(t)$ as

$$g(t) \approx \sum_{k=-(q-1)}^{N+(q-1)} g_k \psi\left(\frac{t - t_k}{\delta}\right),$$

where $\delta = (b - a)/N$, $t_k = a + k\delta$ and $g_k = g(t_k)$. Note that this approximation requires knowledge of g outside of the interval $[a, b]$. This presents no difficulties in our application since the integrands $p(\rho)$ and $\rho p(\rho)$ are given by analytic expressions.

After some simplification, the integral becomes

$$I(\omega) \approx \delta e^{i\omega a} \left[W(\theta) S(\theta) + \nu(\theta) + e^{i\omega(b-a)} \overline{\mu(\theta)} \right],$$

where $\theta = \omega\delta$,

$$S(\theta) = \sum_{k=0}^N g_k e^{i\theta k},$$

$$W(\theta) = \int_{-p}^p \psi(s) \cos(\theta s) ds,$$

$$\nu(\theta) = g_0 \gamma_0(\theta) + \sum_{k=1}^{q-1} \left[g_k \gamma_k(\theta) - g_{-k} \overline{\gamma_k(\theta)} \right],$$

$$\mu(\theta) = g_N \gamma_0(\theta) + \sum_{k=1}^{q-1} \left[g_{N-k} \gamma_k(\theta) - g_{N+k} \overline{\gamma_k(\theta)} \right],$$

and

$$\gamma_k(\theta) = e^{i\theta k} \int_k^q \psi(s) e^{i\theta s} ds.$$

Thus, the computation involves a simple sum of $N + 1$ terms, $S(\theta)$, the quantity $W(\theta)$ and a relatively small number of endpoint corrections, $\nu(\theta)$ and $\mu(\theta)$. Furthermore, since $\psi(s)$ is known analytically, $W(\theta)$ and $\gamma_k(\theta)$ can be computed exactly for each choice of ψ .

The only approximation in the method to this point is the high-order interpolation of $g(t)$. Thus, we require only enough points to accurately approximate $g(t)$ instead of the highly oscillatory function $g(t)e^{i\omega t}$. Furthermore, since we only approximate $g(t)$, the accuracy of the approximation is *independent of ω* . Hence, given any $\epsilon > 0$, we can choose N sufficiently large so that the error in $I(\omega)$ is less than ϵ , *uniformly in ω* .

Note that to decrease the error one may either increase the number of interpolation points N or increase q (thereby increasing the order of the interpolation). As can be easily demonstrated, the order of the method depends on q in much the same way as with Newton-Cotes integration methods. More precisely, for q odd, the error decays like $\mathcal{O}(N^{-(q+1)})$ and for q even, the error decays like $\mathcal{O}(N^{-(q+2)})$. Hence, we most generally choose $q = 2$ (fourth-order convergence) or $q = 4$ (sixth-order convergence). The values of $W(\theta)$ and $\gamma_k(\theta)$ for $q = 2$ and $q = 4$ are found in Sections C.1 and C.2, respectively.

In general, we may need to evaluate $I(\omega)$ for many different values of ω . (In our application, $\omega = \kappa \pm 2\pi|c_\ell|$ with $(c_\ell)_q = \ell_q/(B_q - A_q)$ and where $|\ell_q| \leq \tilde{N}_q/2$.) This is not difficult to obtain for $W(\theta)$, $\nu(\theta)$ and $\mu(\theta)$ since they involve only a few of the g_k . However, straightforward evaluation of the sum $S(\theta)$ has quadratic complexity. To reduce the complexity, we use an FFT to compute $S(\theta)$ at $\theta_n = 2\pi n/N_F$ for $n = 0, \dots, N_F - 1$, where

$N_F > N$. More precisely,

$$S(\theta) = \sum_{k=0}^N g_k e^{i\theta_n k} = \sum_{k=0}^{N_F-1} g_k e^{2\pi i k n / N_F},$$

where we set $g_k = 0$ for $k > N$. Since $S(\theta)$ is periodic in θ with period 2π , we thereby obtain the value of the $S(\theta)$ at $\theta = \theta_n + 2\pi r$, $r \in \mathbb{Z}$.

Thus, given an arbitrary $\theta = \omega\delta$, we interpolate to find $S(\theta)$, which together with $W(\theta)$, $\nu(\theta)$ and $\mu(\theta)$ give us $I(\omega)$. The number of interpolation points N_p determines the order of the interpolation. While a large value of N_p yields high-order accuracy, it is well known that choosing N_p too large can lead to numerical instabilities. Hence, we generally choose $N_p \leq 10$. Furthermore, although increasing the value of N_F also increases the accuracy of the interpolated value $S(\theta)$, the actual value of N_F is less important than the ratio N_F/N , called the oversampling rate β . This is the number of points per wavelength with which the most oscillatory mode in $S(\theta)$ is sampled. We have found that for a partition of unity $p(\rho)$ with $t_0 = 1/2$ and $t_1 = 1$ (see (4.5)) the values $q = 4$, $N = 1024$, $\beta = 128$, $N_p = 10$ as well as $q = 2$, $N = 8192$, $\beta = 128$, $N_p = 10$ give us nearly full double precision accuracy. The choice between these two possibilities depends on the problem size. When $q = 4$, the FFT is faster since $N_F = \beta N$ is smaller. When $q = 2$, the interpolation is faster since the endpoint corrections, ν and μ , are simpler. Hence, in smaller problems, we prefer the $q = 4$ values, and in larger problems, we prefer the $q = 2$ values.

C.1 Second-Order Interpolating Polynomials

For the case of $q = 2$, $\psi(s)$ is given by a sum of second-order Lagrange interpolating polynomials. More precisely, define

$$\begin{aligned} \psi_1(s) &= \begin{cases} \frac{(s+2)}{2} \frac{(s+1)}{1}, & \text{if } -2 \leq s \leq 0 \\ 0, & \text{otherwise,} \end{cases} \\ \psi_2(s) &= \begin{cases} \frac{(s+1)}{1} \frac{(s-1)}{-1}, & \text{if } -1 \leq s \leq 1 \\ 0, & \text{otherwise,} \end{cases} \\ \psi_3(s) &= \begin{cases} \frac{(s-1)}{-1} \frac{(s-2)}{-2}, & \text{if } 0 \leq s \leq 2 \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

Notice that ψ_1 and ψ_3 form the usual piecewise second-order Lagrange interpolation scheme when the point $s = 0$ lies on the boundary of two subintervals. On the other hand, ψ_2 is the usual Lagrange interpolating polynomial when the point $s = 0$ lies at the center of the subinterval. Addition and normalization leads to

$$\psi(s) = \frac{1}{2} [\psi_1(s) + \psi_2(s) + \psi_3(s)] .$$

The functions $W(\theta)$ and $\gamma_k(\theta)$ can now be obtained:

$$\begin{aligned} W(\theta) &= \frac{4 \sin^3 (\theta/2) [2 \cos (\theta/2) + \theta \sin (\theta/2)]}{\theta^3}, \\ \gamma_0(\theta) &= -\frac{2i + (3 + 4i\theta)\theta - 4(\theta + i)e^{i\theta} + (\theta + 2i)e^{2i\theta}}{4\theta^3}, \\ \gamma_1(\theta) &= -\frac{e^{i\theta} [-2i + \theta + (2 + it)e^{i\theta}]}{4\theta^3}. \end{aligned}$$

It is important to note that for $\theta \ll 1$ the numerical evaluation of these functions can produce a significant amount of cancellation error. To avoid this problem, for sufficiently small θ , we approximate $W(\theta)$ and $\gamma_k(\theta)$ with a power series. Through experiment, we determine the value of θ at which to switch (for approximately double precision accuracy) from one method to the other. For example, for the function $W(\theta)$ above, we switch to the power series method for $\theta < 10^{-4}$; and for $\gamma_1(\theta)$, we switch for $\theta < 8/10$.

C.2 Fourth-Order Interpolating Polynomials

For $q = 4$, we similarly construct $\psi(s)$ as a sum of fourth-order Lagrange interpolating polynomials. Define

$$\begin{aligned}\psi_1(s) &= \begin{cases} \frac{(s+4)}{4} \frac{(s+3)}{3} \frac{(s+2)}{2} \frac{(s+1)}{1}, & \text{if } -4 \leq s \leq 0 \\ 0, & \text{otherwise,} \end{cases} \\ \psi_2(s) &= \begin{cases} \frac{(s+3)}{3} \frac{(s+2)}{2} \frac{(s+1)}{1} \frac{(s-1)}{-1}, & \text{if } -3 \leq s \leq 1 \\ 0, & \text{otherwise,} \end{cases} \\ \psi_3(s) &= \begin{cases} \frac{(s+2)}{2} \frac{(s+1)}{1} \frac{(s-1)}{-1} \frac{(s-2)}{-2}, & \text{if } -2 \leq s \leq 2 \\ 0, & \text{otherwise,} \end{cases} \\ \psi_4(s) &= \begin{cases} \frac{(s+1)}{1} \frac{(s-1)}{-1} \frac{(s-2)}{-2} \frac{(s-3)}{-3}, & \text{if } -1 \leq s \leq 3 \\ 0, & \text{otherwise,} \end{cases} \\ \psi_5(s) &= \begin{cases} \frac{(s-1)}{-1} \frac{(s-2)}{-2} \frac{(s-3)}{-3} \frac{(s-4)}{-4}, & \text{if } 0 \leq s \leq 4 \\ 0, & \text{otherwise.} \end{cases}\end{aligned}$$

Then, as with $q = 2$, $\psi(s)$ is given by the normalized sum of these piecewise smooth polynomials

$$\psi(s) = \frac{1}{4} \sum_{j=0}^5 \psi_j(s).$$

In this case for $W(\theta)$ and $\gamma_k(\theta)$, we obtain

$$\begin{aligned}W(\theta) &= \frac{4 \sin^5\left(\frac{\theta}{2}\right)}{3\theta^5} \left\{ 2\theta [12 - \theta^2 + 3(6 - \theta^2) \cos \theta] \sin\left(\frac{\theta}{2}\right) \right. \\ &\quad \left. + (12 + \theta^2) \cos\left(\frac{\theta}{2}\right) + (12 - 11\theta^2) \cos\left(\frac{3\theta}{2}\right) \right\},\end{aligned}$$

$$\begin{aligned}
\gamma_0 &= \frac{1}{48\theta^5} \left[(12i + 30\theta - 35i\theta^2 + 25\theta^3 - 48i\theta^4) + (-48i - 108\theta + 104i\theta^2 + 48\theta^3) e^{i\theta} \right. \\
&\quad + (72i + 144\theta - 114i\theta^2 - 36\theta^3) e^{2i\theta} + (-48i - 84\theta + 56i\theta^2 + 16\theta^3) e^{3i\theta} \\
&\quad \left. + (12i + 18\theta - 11i\theta^2 - 3\theta^3) e^{4i\theta} \right], \\
\gamma_1 &= \frac{1}{48\theta^5} \left[(-36i - 66\theta + 33i\theta^2 - 29\theta^3) e^{i\theta} + (72i + 144\theta - 114i\theta^2 - 36\theta^3) e^{2i\theta} \right. \\
&\quad + (-48i - 84\theta + 56i\theta^2 + 16\theta^3) e^{3i\theta} + (12i + 18\theta - 11i\theta^2 - 3\theta^3) e^{4i\theta} \left. \right], \\
\gamma_2 &= \frac{1}{48\theta^5} \left[(36i + 42\theta + 3i\theta^2 + 7\theta^3) e^{2i\theta} + (-48i - 84\theta + 56i\theta^2 + 16\theta^3) e^{3i\theta} \right. \\
&\quad \left. + (12i + 18\theta - 11i\theta^2 - 3\theta^3) e^{4i\theta} \right], \\
\gamma_3 &= \frac{1}{48\theta^5} \left[(-12i - 6\theta - i\theta^2 - \theta^3) e^{3i\theta} + (12i + 18\theta - 11i\theta^2 - 3\theta^3) e^{4i\theta} \right].
\end{aligned}$$

Bibliography

- [1] Netlib repository. <http://www.netlib.org>. 102
- [2] M. Abramowitz and I. A. Stegun, editors. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, 1965. 24, 25, 38, 94, 98
- [3] K. E. Atkinson. *An Introduction to Numerical Analysis*. Wiley, second edition, 1989. 32, 39, 40
- [4] S. Balay, K. Buschelman, W. D. Gropp, D. Kaushik, L. C. McInnes, and B. F. Smith. PETSc home page. <http://www.mcs.anl.gov/petsc>, 2001. 57
- [5] S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith. Efficient management of parallelism in object oriented numerical software libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhauser Press, 1997. 57
- [6] S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith. PETSc users manual. Technical Report ANL-95/11 - Revision 2.1.0, Argonne National Laboratory, 2001. 57
- [7] R. L. Berger, B. F. Lasinski, T. B. Kaiser, E. A. Williams, A. B. Langdon, and B. I. Cohen. Theory and three-dimensional simulation of light filamentation in laser-produced plasma. *Phys. Fluids B*, 5(7):2243–2258, July 1993. 2, 91
- [8] L. Bers, F. John, and M. Schechter. *Partial Differential Equations*. Interscience, 1964. 19
- [9] N. N. Bojarski. The k -space formulation of the scattering problem in the time domain. *J. Opt. Soc. Amer.*, 72:570–584, 1982. 6

- [10] J. J. Bowman, T. B. A. Senior, and P. L. E. Uslenghi, editors. *Electromagnetic and Acoustic Scattering by Simple Shapes*. North-Holland Publishing Company, 1969. 25, 94
- [11] A. Brandt and I. Livshits. Wave-ray multigrid method for standing wave equations. *Electronic Transactions on Numerical Analysis*, 6:162–181, December 1997. 91
- [12] O. P. Bruno and L. A. Kunyansky. A fast, high-order algorithm for the solution of surface scattering problems: Basic implementation, tests and applications. *J. Comput. Phys.*, 169:80–110, 2001. 55, 91
- [13] O. P. Bruno and A. Sei. A fast high-order solver for EM scattering from complex penetrable bodies: TE case. *IEEE Trans. in Antenn. Propag.*, 48(12):1862–1864, December 2000. 12, 14, 28, 31, 32, 90
- [14] L. F. Canino, J. J. Ottusch, M. A. Stalzer, J. L. Visher, and S. M. Wandzura. Numerical solution of the Helmholtz equation in 2D and 3D using a high-order Nyström discretization. *J. Comput. Phys.*, 146:627–663, 1998. 6
- [15] R. Coifman, V. Rokhlin, and S. Wandzura. The fast multipole method for the wave equation: A pedestrian prescription. *IEEE Antennas and Propagation*, 35(7), 1993. 5
- [16] D. Colton and R. Kress. *Integral Equation Methods in Scattering Theory*. Wiley-Interscience, 1983. 14
- [17] D. Colton and R. Kress. *Inverse Acoustic and Electromagnetic Scattering Theory*. Springer-Verlag, second edition, 1998. 4, 5, 14, 15, 44, 55, 81
- [18] E. Darve. The fast multipole method: Numerical implementation. *J. Comput. Phys.*, 160:195–240, 2000. 5
- [19] B. Dembart and E. Yip. The accuracy of fast multipole methods for Maxwell’s equations. *IEEE Comput. Sci. Eng.*, 4:48, 1998. 5
- [20] A. Ditkowski, K. Dridi, and J. S. Hesthaven. Convergent Cartesian grid methods for Maxwell’s equations in complex geometries. *J. Comput. Phys.*, 170(1):39–80, 2001. 5

- [21] M. R. Dorr, F. X. Garaizar, and J. F. A. Hittinger. Simulation of laser plasma filamentation using adaptive mesh refinement. Submitted to the Journal of Comput. Phys., June 2001. 2, 91
- [22] C. Elachi. *Introduction to the Physics and Techniques of Remote Sensing*. John Wiley & Sons, 1987. 1
- [23] M. D. Feit and Jr. J. A. Fleck. Beam nonparaxiality, filament formation, and beam breakup in the self-focusing of optical beams. *Journal of the Optical Society of America B*, 5(3), 1988. 91
- [24] G. B. Folland. *Introduction to Partial Differential Equations*. Princeton University Press, 1976. 10, 19
- [25] M. Frigo and S. G. Johnson. FFTW home page. <http://www.fftw.org>. 57
- [26] M. Frigo and S. G. Johnson. FFTW: An adaptive software architecture for the FFT. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, volume 3, pages 1381–1384, 1998. 57
- [27] B. Fultz and J. M. Howe. *Transmission Electron Microscopy and Diffractometry of Materials*. Springer, 2001. 2, 92
- [28] D. Gilbarg and N. S. Trudinger. *Elliptic Partial Differential Equations of Second Order*. Springer-Verlag, 1977. 19, 45
- [29] A. Greenbaum. *Iterative Methods for Solving Linear Systems*. SIAM, 1997. 41, 42
- [30] L. Greengard, J. F. Huang, V. Rokhlin, and S. Wandzura. Accelerating fast multipole methods for the Helmholtz equation at low frequencies. *IEEE Comput. Sci. Eng.*, 5: 32, 1998. 5
- [31] L. Greengard and V. Rokhlin. A fast algorithm for particle simulations. *J. Comput. Phys.*, 135:280, 1997. 5
- [32] L. R. Hamilton, J. J. Ottusch, M. A. Stalzer, R. S. Turley, J. L. Visher, and S. M. Wandzura. Numerical solution of 2-D scattering problems using high-order methods. *IEEE Antenn. Propag.*, 47:683–691, 1999. 6

- [33] J. W. Hettinger, M. d. l. P. Mattozzi, W. R. Myers, M. E. Williams, A. Reeves, R. L. Parsons, R. C. Haskell, D. C. Petersen, R. Wang, and J. I. Medford. Optical coherence microscopy. A technology for rapid, in vivo, non-destructive visualization of plants and plant cells. *Plant Physiology*, 123:3–15, May 2000. 2
- [34] G. N. Hounsfield. Computerized transverse axial scanning (tomography). 1. Description of system. *British Journal of Radiology*, 46:1016–1022, 1973. 2
- [35] E. M. Hyde. The validity of a paraxial approximation in the simulation of laser-plasma interactions. Technical Report UCRL-JC-139805, Lawrence Livermore National Laboratory, August 2000. *Presented at Student Sym. 2000, Technical Presentations & Career Fair, August 10, 2000, Sandia NL, Albuquerque, NM.* 2
- [36] A. Kirsch and P. Monk. Convergence analysis of a coupled finite element and spectral method in acoustic scattering. *IMA Journal of Numerical Analysis*, 1990. 5
- [37] A. Kirsch and P. Monk. An analysis of the coupling of finite-element and Nyström methods in acoustic scattering. *IMA Journal of Numerical Analysis*, 14:523–544, 1994. 5
- [38] R. Kress. *Linear Integral Equations*. Springer-Verlag, 1989. 17
- [39] R. Kress and G. F. Roach. Transmission problems for the Helmholtz equation. *J. Math. Phys.*, 19:1433–1437, 1979. 14
- [40] B. Lee, T. A. Manteuffel, S. F. McCormick, and J. Ruge. First-order system least-squares for the helmholtz equation. *SIAM Journal on Scientific Computing*, 21(5):1927–1949, May 2000. 91
- [41] J. H. Lin and W. C. Chew. BiCG-FFT T-Matrix method for solving for the scattering solution from inhomogeneous bodies. *IEEE Transactions on Microwave Theory and Techniques*, 44(7):1150–1155, July 1996. 61
- [42] G. Liu and S. Gedney. High-order Nyström solution of the volume EFIE for TM-wave scattering. *Microwave and Optical Technology Letters*, 25(1):8–11, April 2000. 6
- [43] B. R. Masters. Early development of optical low-coherence reflectometry and some

recent biomedical applications. *Journal of Biomedical Optics*, 4(2):236–247, April 1999.

2

- [44] L. M. Peng, K. Gjønnnes, and J. Gjønnnes. Bloch wave treatment of symmetry and multiple beam cases in reflection high energy electron diffraction and reflection electron microscopy. *Microscopy Research and Technique*, 20:360–370, 1992. 2
- [45] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in Fortran 77: The Art of Scientific Computing*, volume 1. Cambridge University Press, 2 edition, 1992. 8, 26, 30, 36, 52, 56, 57, 101, 103
- [46] W. Rachowicz and L. Demkowicz. An hp-adaptive finite element method for electromagnetic Part 1: Data structure and constrained approximation. *Computer Methods in Applied Mechanics and Engineering*, 187:307–335, 2000. 5
- [47] V. Rokhlin. Rapid solution of integral equations of classical potential theory. *J. Comput. Phys.*, 60:187–207, 1985. 5
- [48] V. Rokhlin. Rapid solution of integral equations of scattering theory in two dimensions. *J. Comput. Phys.*, 86:414, 1990. 5
- [49] G. M. Vainikko. Fast solvers of the Lippmann-Schwinger equation. In R. P. Gilbert, J. Kajiwara, and Y. S. Xu, editors, *Direct and Inverse Problems of Mathematical Physics*, pages 423–440. Kluwer, 2000. 7
- [50] K. Watanabe, S. Hara, Y. Horio, and I. Hashimoto. A two-dimensional Bloch-wave method for dynamical RHEED calculations. *Acta Cryst.*, A(54):452–459, 1998. 2
- [51] X. M. Xu and Q. H. Liu. Fast spectral-domain method for acoustic scattering problems. *IEEE Trans. Ultrasonics, Ferroelectrics, and Frequency Control*, 48(2):522–529, March 2001. 6, 61
- [52] B. Yang, D. Gottlieb, and J. S. Hesthaven. Spectral simulations of electromagnetic wave scattering. *J. Comput. Phys.*, 134(2):216–230, 1997. 5
- [53] P. Zwamborn and P. Van den Berg. Three dimensional weak form of the conjugate gradient FFT method for solving scattering problems. *IEEE Trans. Microwave Theory Tech.*, 40:1757–1766, Sept. 1992. 6