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

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

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

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

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]

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

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

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

where $h=(b-a)/N$ and $k=0,\ldots,N-1$. 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

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

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-

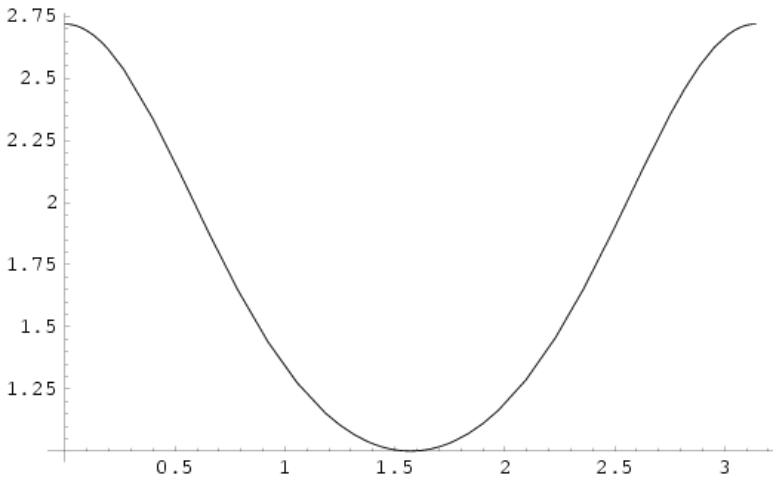

Figure 1.1: Smooth and Periodic Function – $f(x) = e^{\cos^2 x}$

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

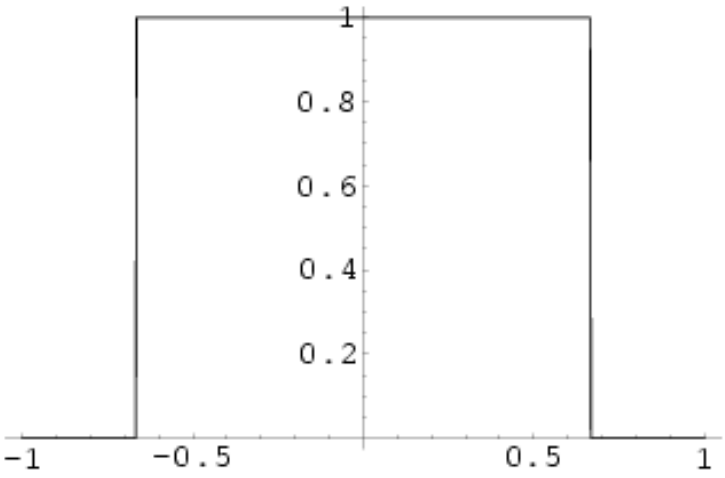

(a) Discontinuous Function f

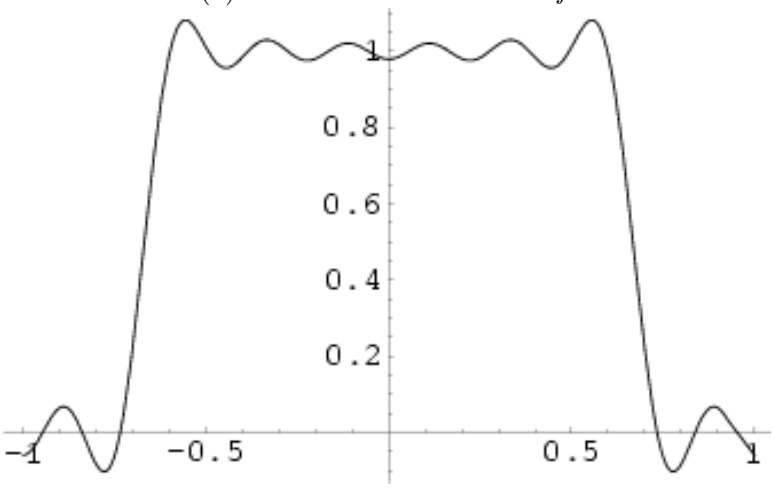

(b) Fourier-Smoothed Function f^F

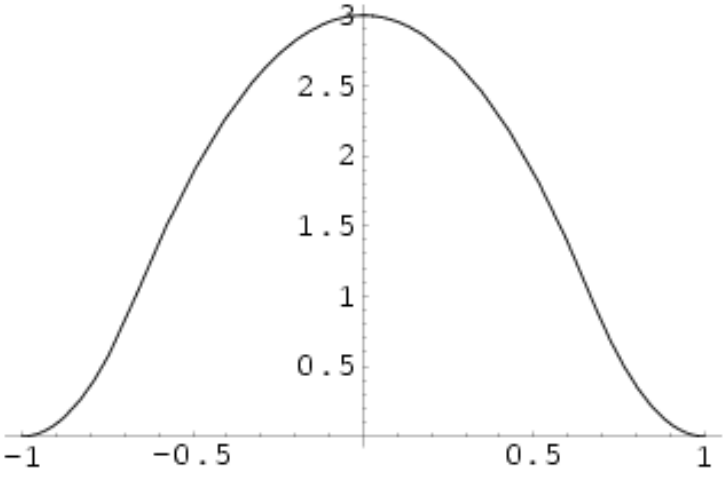

(c) C^1 Function g

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 gemetric 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.