and TM polarizations have the same behavior for both cases, but the magnitude for the TM case is higher than the magnitude for the TE case.

In Fig. 9 the normalized bistatic and backscattering cross sections are plotted for two identical conducting spheroids A and B of axial ratio 5 and semi-major axis length $\lambda/4$, separated by a center-to-center distance of $\lambda/2$ along the x axis of A, with the rotation of B relative to A specified by the Euler angles 45°, 90°, 45°. The behavior of both the E-plane and the H-plane patterns is quite similar. However, the backscattering cross section pattern for TM polarization has more oscillations as compared to the pattern for TE polarization, with two minima around $\vartheta_i = 35°$, 145° and a maximum around $\vartheta_i = 90°$.

Numerical results for two identical dielectric spheroids A and B of axial ratio 5, semi-major axis length $\lambda/4$, and relative permitivity 3, in a parallel configuration, are given in Fig. 10. The two spheroids are separated by a center-to-center distance of $\lambda/2$ along the x axis. In this case the E-plane pattern has a sharp minimum around $\vartheta = 90°$, whereas the H-plane pattern presents some oscillation about the same minimum point. The backscattering cross section has a minimum around $\vartheta_i = 30°$ for both polarizations and exhibits a similar behavior, with a higher magnitude for the TM case than for the TE case at higher angles of incidence, as expected. Because of symmetry about $\vartheta_i = 90°$, the variation of the backscattering cross section is shown only for the range from $\vartheta_i = 0°$ to $\vartheta_i = 90°$.

Chapter 5

The Discrete Dipole Approximation for Light Scattering by Irregular Targets

Bruce T. Draine
Princeton University Observatory
Princeton, New Jersey 08544

I. INTRODUCTION

The discrete dipole approximation—also referred to as the coupled dipole approximation—is a flexible technique for studying scattering and absorption of electromagnetic radiation by targets with sizes comparable to the wavelength. Advances in numerical techniques coupled with the increasing speed and memory of scientific workstations now make it possible for calculations to be carried out for targets with dimensions as large as several times the wavelength of the incident

Light Scattering by Nonspherical Particles: Theory, Measurements, and Applications

radiation. Although the technique is not well suited for targets with very large complex refractive index m, it works well for materials with $|m-1| \lesssim 3$ and target dimension $D \lesssim 5\lambda$, where λ is the wavelength in the surrounding medium. The discrete dipole approximation (DDA) has been applied to compute scattering and absorption by targets of size comparable to the wavelength in a broad range of problems, including interstellar dust grains (Draine and Malhotra, 1993; Draine and Weingartner, 1996, 1997; Wolff *et al.*, 1998), ice crystals in the atmosphere of the Earth (Okamoto *et al.*, 1995; Lemke *et al.*, 1998) and other planets (West and Smith, 1991), interplanetary dust (Mann *et al.*, 1994; Kimura and Mann, 1998), cometary dust (Okamoto *et al.*, 1994; Xing and Hanner, 1997; Yanamandra-Fisher and Hanner, 1999), soot produced in flames (Ivezić and Mengüc, 1996; Ivezić *et al.*, 1997), surface features on semiconductor devices (Schmehl *et al.*, 1997; Nebeker *et al.*, 1998), and optical characteristics of human blood cells (Hoekstra *et al.*, 1998).

The discrete dipole approximation was introduced by Purcell and Pennypacker (1973) and has undergone a number of theoretical developments since then, including the introduction of radiative reaction corrections (Draine, 1988), application of fast Fourier transform techniques (Goodman *et al.*, 1991), and a prescription for dipole polarizabilities based on the lattice dispersion relation (Draine and Goodman, 1993). The discrete dipole approximation was reviewed recently by Draine and Flatau (1994).

II. WHAT IS THE DISCRETE DIPOLE APPROXIMATION?

There has been some confusion about exactly *what* is being approximated in the "discrete dipole approximation." The actual approximation can be simply stated:

> The discrete dipole approximation consists of approximating the actual target by an array of polarizable points (the "dipoles").

This is the *only* essential approximation—once the location and polarizability of the points are specified, calculation of the scattering and absorption of light by the array of polarizable points can be carried out to whatever accuracy is required (within the practical limits imposed by the computational hardware).

Suppose we have an array of points $\mathbf{x}_j$, $j = 1, \ldots, N$, each with complex polarizability tensor $\boldsymbol{\alpha}_j$, and with a monochromatic incident wave $\mathbf{E}_{\mathrm{inc},j} e^{-i\omega t}$ at each location, where ω is the angular frequency and t is time. Each of the dipoles is subject to an electric field that is the sum of the incident wave plus the electric fields resulting from all of the other dipoles. The self-consistent solution for the dipole moments $\mathbf{P}_j e^{-i\omega t}$ satisfies a system of $3N$ linear equations, which can be

written as

$$\sum_{j=1}^{N} \mathbf{A}_{ij}\mathbf{P}_j = \mathbf{E}_{\mathrm{inc},i}, \qquad i = 1, \ldots, N, \tag{1}$$

where the elements $\mathbf{A}_{ij}$ are 3×3 matrices (see Draine and Flatau, 1994). The diagonal elements $\mathbf{A}_{ii} = \boldsymbol{\alpha}_i^{-1}$, where $\boldsymbol{\alpha}$ is the 3×3 complex polarizability tensor (see Section VI). The off-diagonal elements $\mathbf{A}_{i\neq j}$ depend only on $k = \omega/c$, where c is the speed of light, and the vector displacement $\mathbf{r}_{ij} \equiv \mathbf{x}_i - \mathbf{x}_j$. Equation (1) is a system of $3N$ complex linear equations; the computational challenge is to find the solution $\mathbf{P}_j$ satisfying this equation.

Once the solution $\mathbf{P}_j$ has been found, it is straightforward to calculate the complete scattering matrix for the target, as well as other quantities, such as the absorption and extinction cross sections, the intensity and polarization of scattered radiation, and the force and torque exerted on the target by the electromagnetic field (Draine and Weingartner, 1996).

There are some obvious issues surrounding the discrete dipole approximation:

1. How many dipoles are required for the dipole array to adequately approximate the target?
2. If a given lattice is to approximate a given continuum target, what choice of dipole polarizabilities will result in the most accurate approximation?
3. For a given target geometry, size, complex refractive index m, and incident wavelength λ, how accurate is the discrete dipole approximation?
4. What are the computational requirements of the discrete dipole approximation?

These are addressed next.

III. THE DDSCAT SCATTERING CODE

DDSCAT is a portable f77 code developed by B. T. Draine and P. J. Flatau to carry out calculations using the discrete dipole aproximation. The DDSCAT code is publicly available[1] (Draine and Flatau, 1994), and a comprehensive user guide is now available (Draine and Flatau, 1997). The calculations reported in the following discussion were carried out using the current "release," DDSCAT.5a8.

[1] http://astro.princeton.edu/~draine/, or anonymous ftp to astro.princeton.edu, directory draine/scat/ddscat.

IV. DIPOLE ARRAY GEOMETRY

Equation (1) applies for any array geometry. For general geometry, there are $(3N)^2$ distinct complex elements of the matrix $\mathbf{A}$; when considering $N \gtrsim 10^4$ dipoles, it is apparent that computing this matrix would be very central processing unit (CPU) intensive, and storing the elements for reuse would require large amounts of random access memory (RAM). For example, in the following discussion we show results computed for a target with $N \approx 200{,}000$ dipoles; storing the complete $\mathbf{A}$ matrix, with 8 bytes per complex number, would require 2.8 TB! There are great advantages to be gained if the dipoles are located on a lattice, because now many different pairs i, j have identical $\mathbf{r}_{ij}$, and hence identical $\mathbf{A}_{ij}$.

Because of this, DDSCAT requires the target array to reside on a cubic lattice. It is then possible to use fast Fourier transform (FFT) techniques to evaluate matrix–vector products $\mathbf{A}\cdot\mathbf{v}$ (Goodman *et al.*, 1991); because they allow evaluation of the product in $O[(3N)\ln(3N)]$ rather than $O[(3N)^2]$ operations, FFT techniques allow use of much larger values of N.

DDSCAT includes routines to generate dipole arrays for a variety of target geometries, including ellipsoids, rectangular prisms, hexagonal prisms, and tetrahedra; it can also accept a user-supplied list of occupied lattice sites. The target material can be anisotropic and the target can be inhomogeneous.

V. TARGET GENERATION

There is some arbitrariness in the construction of the array of point dipoles intended to represent a solid target of specified geometry. DDSCAT uses a straightforward procedure: In the "target frame," construct a target of volume V. Let the target centroid define the origin of coordinates. Choose a "trial" lattice spacing d and construct a lattice $(x, y, z) = (n_x, n_y, n_z)d + (o_x, o_y, o_z)d$, where the n_j are integers and the "offset" vector (o_x, o_y, o_z) allows the target centroid to be located at a lattice point or between lattice points, as appropriate.

Having chosen d and (o_x, o_y, o_z), the target array is now taken to consist of the lattice points located within the target volume; let N be the number of such points. With these N lattice points now determined, we make a small adjustment to the lattice spacing and set $d = (V/N)^{1/3}$ (when N is large, the d so obtained is nearly the same as the original "trial" d).[2]

In Fig. 1 we show an $N = 59{,}728$ dipole representation of a sphere. The array fits within a $48 \times 48 \times 48$ region on the lattice. In Fig. 2 we show an $N = 61{,}432$

[2]Because each dipole "represents" a volume d^3 of material, we require the array volume $Nd^3 = V$.

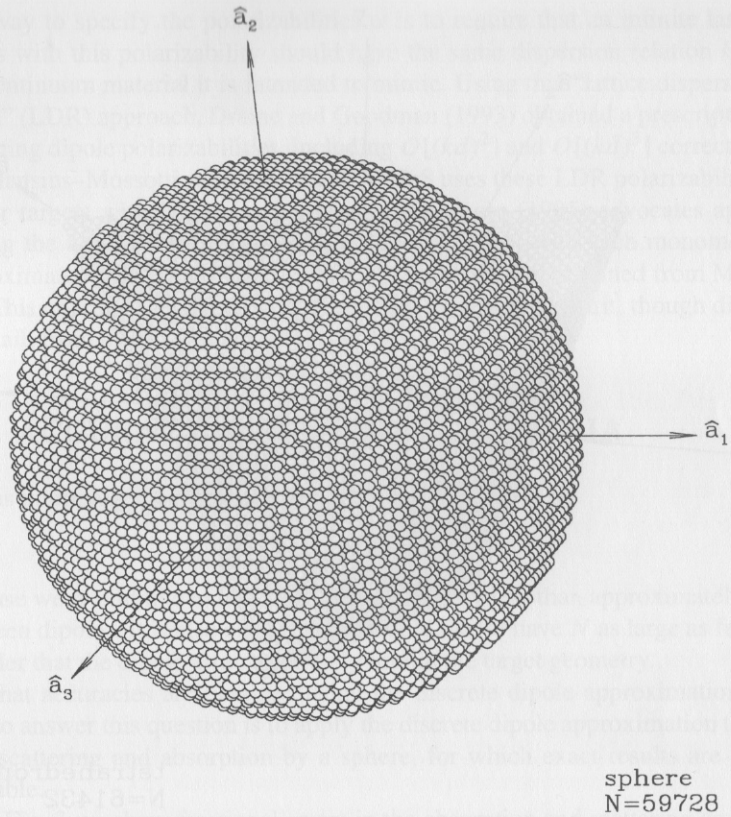

Figure 1 $N = 59{,}728$ dipole array representing a sphere.

dipole representation for a regular tetrahedron. In this case the dipole array fits within a $78 \times 82 \times 96$ region on the lattice.

It is convenient to characterize the target size by the "effective radius"

$$a_{\rm eff} \equiv \left(\frac{3V}{4\pi}\right)^{1/3} = \left(\frac{3N}{4\pi}\right)^{1/3} d; \qquad (2)$$

$a_{\rm eff}$ is simply the radius of a sphere of equal volume. DDSCAT reports dimensionless scattering, absorption, and extinction efficiency factors, $Q_{\rm sca}$, $Q_{\rm abs}$, and $Q_{\rm ext}$—these are simply the corresponding cross sections divided by $\pi a_{\rm eff}^2$.

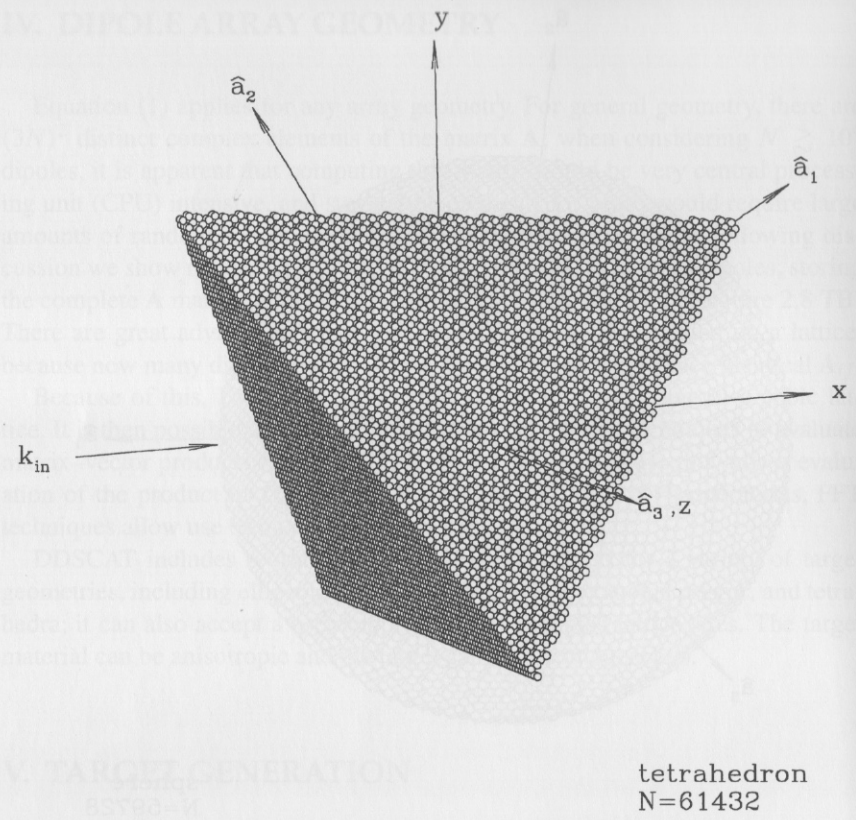

Figure 2 Array of $N = 61{,}432$ dipoles representing a regular tetrahedron.

VI. DIPOLE POLARIZABILITIES

After constructing a target array with lattice spacing d, it remains to specify the dipole polarizabilities $\boldsymbol{\alpha}_j$. In the limit $d/\lambda \to 0$, the Clausius–Mossotti prescription would apply: $\alpha^{(\mathrm{CM})} = (3d^3/4\pi)(\varepsilon - 1)/(\varepsilon + 2)$, where $\varepsilon = m^2$ is the complex dielectric constant. For finite d/λ, however, this is not the best choice. It is important to include radiative-reaction corrections, which are of $O[(kd)^3]$ (Draine, 1988). However, there are also corrections of $O[(kd)^2]$, and there has been some controversy about what these should be taken to be (Goedecke and O'Brien, 1988; Iskander *et al.*, 1989a; Hage and Greenberg, 1990; Dungey and Bohren, 1991; Lumme and Rahola, 1994; Okamoto, 1995). Because it is intended that the dipole array mimic a continuum material of dielectric constant ε, a natural way to specify the polarizabilities α is to require that an infinite lattice of points with this polarizability should have the same dispersion relation $k(\omega)$ as the continuum material it is intended to mimic. Using this "lattice dispersion relation" (LDR) approach, Draine and Goodman (1993) obtained a prescription for assigning dipole polarizabilities, including $O[(kd)^2]$ and $O[(kd)^3]$ corrections to the Clausius–Mossotti estimate. DDSCAT.5a8 uses these LDR polarizabilities.

For targets with large size parameters, Okamoto (1995) advocates approximating the target as a cluster of spherical monomers, with each monomer then approximated by a point dipole, but with a polarizability obtained from Mie theory. This approach (Okamoto and Xu, 1998) is similar in spirit, though different in detail, from that of Dungey and Bohren (1991).

VII. ACCURACY AND VALIDITY CRITERIA

It is intuitively clear that one validity criterion should be

$$|m|kd \lesssim 1 \tag{3}$$

because we would like to have the "phase" vary by less than approximately 1 rad between dipoles. It is also clear that we would like to have N as large as feasible, in order that the dipole array accurately mimic the target geometry.

What accuracies are obtained using the discrete dipole approximation? One way to answer this question is to apply the discrete dipole approximation to compute scattering and absorption by a sphere, for which exact results are readily available.

In Fig. 3 we show fractional errors in the absorption and scattering cross sections C_{abs} and C_{sca} for a sphere with refractive index $m = 1.7 + 0.1i$. For $N = 17{,}904$, accuracies are better than 2% up to $x = 9.6$, and accuracies of better than 1% are attainable for size parameter $x = ka$ as large as 21.4 using $N \approx 2 \times 10^5$ dipoles, where a is the radius.

The accuracy does depend on the refractive index m; for smaller values of $|m - 1|$ the accuracy for a given N and $|m|kd$ is generally better. Draine and Flatau (1994) show accuracies for $m = 1.33 + 0.01i$, $1.7 + 0.1i$, $2 + i$, and $3 + 4i$, and comparison to exact results for touching spheres has been made by Flatau *et al.* (1993).

VIII. SOLUTION METHOD

When considering $N \gtrsim 10^4$, it is apparent that direct solution of Eq. (1) by standard techniques, requiring approximately $O[(3N)^3]$ operations, is utterly infeasible. However, iterative techniques exist that find excellent approximations

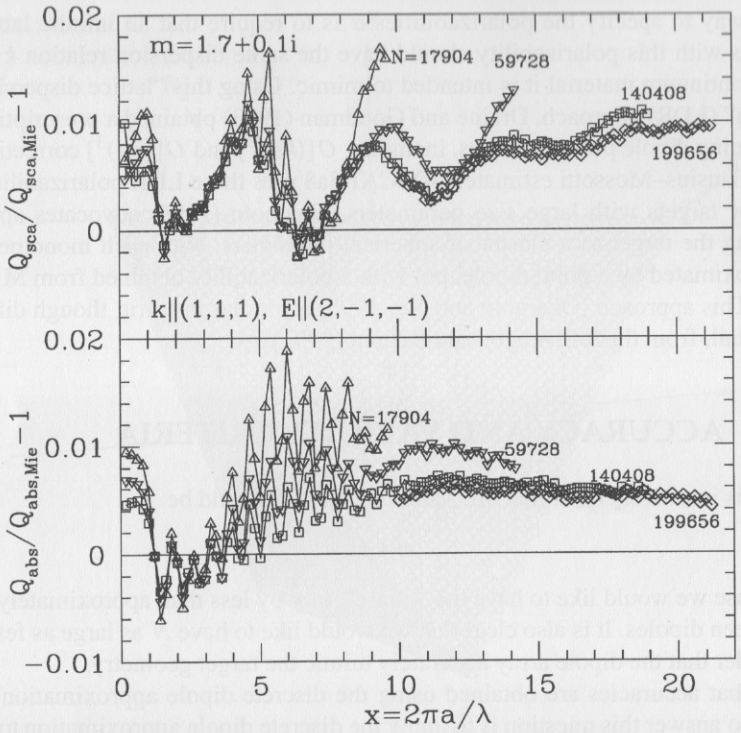

Figure 3 Fractional errors for scattering and absorption cross sections for an $m = 1.7+0.1i$ pseudosphere for different numbers N of dipoles. For each N, results are shown up to the value of x for which $|m|kd \approx 1$ [i.e., $x = (3N/4\pi)^{1/3}/|m|$]. In the "target frame," the incident radiation is propagating along the $(1,1,1)$ direction and polarized along the $(2,-1,-1)$ direction.

to the true solution $\mathbf{P}$ in a modest number (often only 10–100) iterations (Draine, 1988; Flatau, 1997). Each iteration involves computation of a matrix–vector product $\sum_{j=1}^{N} \mathbf{A}_{ij}\mathbf{v}_j$ or $\sum_{j=1}^{N} \mathbf{A}_{ij}^{\dagger}\mathbf{v}_j$, where $\mathbf{v}_j$ is a $3N$-dimensional vector and $\mathbf{A}^{\dagger}$ is the Hermitian conjugate of $\mathbf{A}$. The products are evaluated using FFT techniques, as described previously. Hoekstra *et al.* (1998) have demonstrated that the FFT calculations can be parallelized. A number of different iterative procedures of the complex conjugate gradient type are available (Flatau, 1997). Draine (1988) originally employed the method of Petravic and Kuo-Petravic (1979), which DDSCAT retains as an option. Lumme and Rahola (1994) and Nebeker *et al.* (1998) recommend the "quasi-minimum residual" method (Freund, 1992) as the most computationally efficient. Flatau (1997) compared a number of different methods and recommends the "stabilized bi-conjugate gradient"

method (van der Vorst, 1992) with preconditioning, available as an option within DDSCAT.

For an approximation $\mathbf{P}$, we define the fractional error to be

$$\text{err} \equiv \frac{|\mathbf{AP} - \mathbf{E}_{\text{inc}}|}{|\mathbf{E}_{\text{inc}}|}, \tag{4}$$

where $\mathbf{P}$ and $\mathbf{E}_{\text{inc}}$ are $3N$-dimensional vectors and $\mathbf{A}$ is the $3N \times 3N$ matrix from Eq. (1). We routinely iterate until err $< 10^{-5}$.

IX. COMPUTATIONAL REQUIREMENTS

Numerical techniques for computing scattering by irregular bodies are computationally intensive, and their utility may be limited by computational requirements.

A. Memory

Because of the use of FFT techniques, the memory requirements of DDSCAT are proportional to the number $N_{\text{FFT}} = N_x N_y N_z$ of sites in the "computational volume"—an $N_x \times N_y \times N_z$ region of the lattice containing the "occupied" lattice sites.[3] For a rectangular target, $N_{\text{FFT}} = N$, but for other targets $N_{\text{FFT}} > N$. For spherical targets, $N_{\text{FFT}} \approx (6/\pi)N$; for a tetrahedron, $N_{\text{FFT}} \approx 6N$.

Using 8 bytes per complex number, DDSCAT requires approximately $1.0 + 0.61(N_{\text{FFT}}/1000)$ MB. Thus a 32^3 computational volume requires only 21 MB, but a 64^3 volume would require 161 MB.

B. Central Processing Unit Time

Most of the computing time is spent iterating until the solution vector $\mathbf{P}$ satisfies Eq. (1) to the required accuracy. The time spent per iteration scales approximately as N_{FFT}. For a given scattering problem (target geometry, refractive index m, and $x = ka_{\text{eff}}$), the number of iterations required is essentially independent of N_{FFT}, so the overall CPU time per scattering problem scales approximately linearly with N_{FFT}.

For example, for an $m = 1.7+0.1i$ sphere and $x = 9$, solving for two incident polarizations using a 167-MHz Sun Ultrasparc required 1400 CPU-s for the $N = 17{,}904$ sphere and 14,800 CPU-s for $N = 140{,}408$.

[3] When Temperton's (1992) "generalized prime factor algorithm" is used, N_x, N_y, and N_z must each be of the form $2^p 3^q 5^r$, with p, q, r integers.

X. BENCHMARK CALCULATIONS: SCATTERING BY TETRAHEDRA

In addition to the discrete dipole approximation, there are other approaches that can be used to calculate scattering and absorption by irregular targets, including the "extended boundary condition method," often referred to as the "T-matrix method" (Barber and Yeh, 1975; Mishchenko, 1991a; Chapter 6); the "finite difference time domain method" (e.g., Yang and Liou, 1996a; Chapter 7); and the "volume-integral method" (e.g., Eremin and Ivakhnenko, 1998; Chapter 2; and references therein).

The DDA has been compared to some of these other methods for various target shapes, including spheroids, cylinders, and bispheres (Hovenier *et al.*, 1996) and cubes (Wriedt and Comberg, 1998).

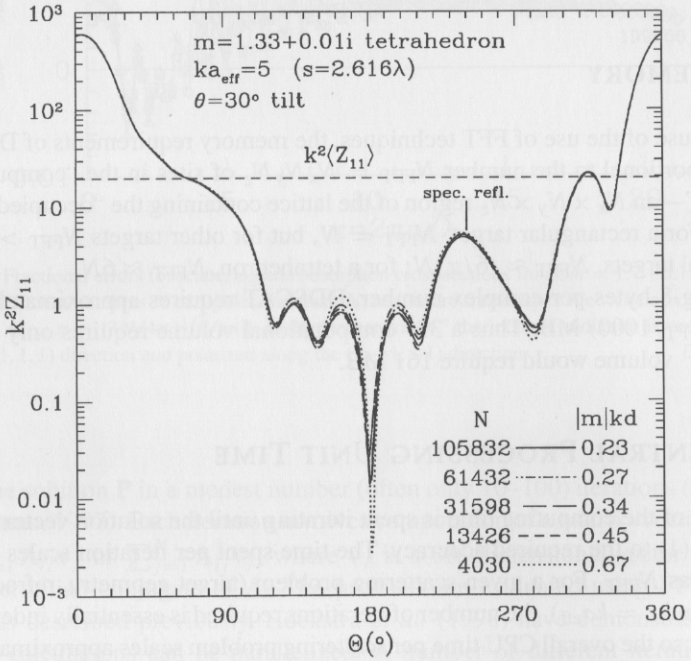

Figure 4 Scattering of unpolarized incident light by a tetrahedron with $m = 1.33 + 0.01i$ and $ka_{\rm eff} = 5$. The incident radiation is propagating in the x direction, and k^2Z_{11} is shown for scattering in the x–y plane. Axis $\hat{a}_1$ of the tetrahedron (see Fig. 2) is in the x–y plane at an angle $\theta = 30°$ from the x axis. Axis $\hat{a}_2$ of the tetrahedron is in the x–y plane. The broken line labeled $k^2\langle Z_{11}\rangle$ is k^2Z_{11} averaged over all scattering directions for this orientation. The peak at $\Theta = 240°$ is from "specular" reflection. Accurate results for Z_{11} are obtained for $|m|kd \approx 0.5$, which appears to be sufficient for good accuracy.

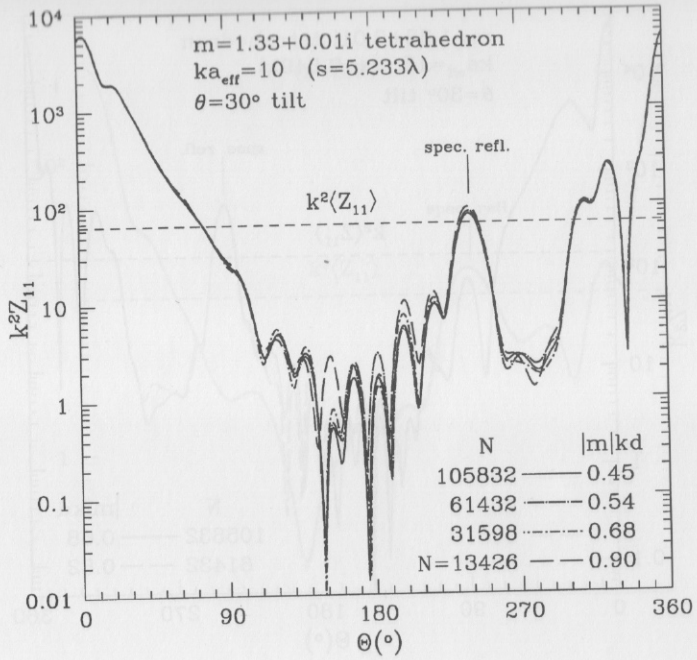

Figure 5 Same as Fig. 4, but for $ka_{\rm eff} = 10$. The specular reflection peak at $\Theta = 240°$ has become more pronounced. Once again, good accuracy is obtained for $|m|kd \approx 0.5$.

Here I present additional "benchmark" calculations that can be used to compare different techniques. From the standpoint of the discrete dipole approximation, there are no "special" shapes—a target is just a list of occupied lattice sites.

A regular tetrahedron is a simple target shape with "edges" that can be used to test different approaches to computing light scattering. The orientation is as shown in Fig. 2, with the tetrahedron axis tilted $\theta = 30°$ away from the direction of propagation of the incident radiation.

Scattering of incident unpolarized light is measured by Z_{11}, one element of the 4×4 phase matrix (see Chapter 1). Figure 4 shows k^2Z_{11} for an $m = 1.33+0.01i$ tetrahedron with $ka_{\rm eff} = 5$, for scattering in the x–y plane. Z_{11} measures the scattered intensity for incident unpolarized light. Results are shown for different numbers N of dipoles. For $ka_{\rm eff} = 5$, the tetrahedron side $s = 2.616\lambda$.

Increasing the size of the target by a factor of 2 in linear extent, we obtain the results shown in Fig. 5. A further increase to $s = 7.849\lambda$ gives the scattering properties shown in Fig. 6. Figures 7 and 8 show the results for a different refractive index.

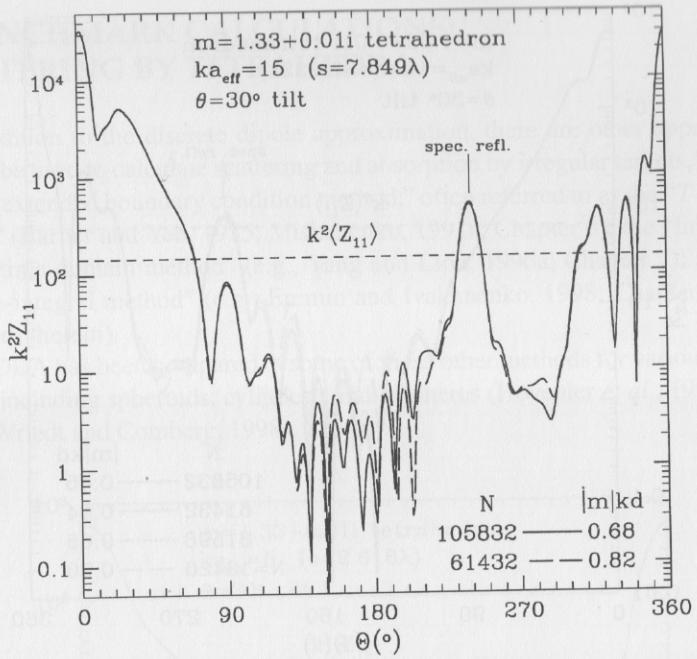

Figure 6 Same as Fig. 4, but for $x = ka_{\rm eff} = 15$. Notice that the specular reflection peak at $\Theta = 240°$ has become more pronounced. The results are well converged where $Z_{11} \gtrsim 0.1\langle Z_{11}\rangle$, but do not appear to be fully converged for directions where the scattering is weak.

Each plot shows the value of $k^2 Z_{11}$ averaged over all scattering directions for this orientation: $k^2\langle Z_{11}\rangle = k^2 C_{\rm sca}/4\pi$, where $C_{\rm sca}$ is the scattering cross section. We see that for $|m|kd \lesssim 1$ the fractional errors in Z_{11} are large only for scattering directions where the scattering is relatively weak to begin with.

In each of these plots it is interesting to note the peak in Z_{11} at a scattering angle $\Theta = 240°$, corresponding to "specular reflection" off the face of the tetrahedron upon which the radiation is incident. This specular reflection peak becomes narrower and stronger as the target size is increased, as expected from diffraction theory.

The CPU time required for these calculations is given in Table I. It will be of interest to compare these scattering results for tetrahedra with the results of other computational techniques.[4]

[4]For those wishing to repeat these benchmark calculations on their system, the ddscat.par files can be obtained by anonymous ftp from astro.princeton.edu, directory draine/scat/ddscat/benchmarks.

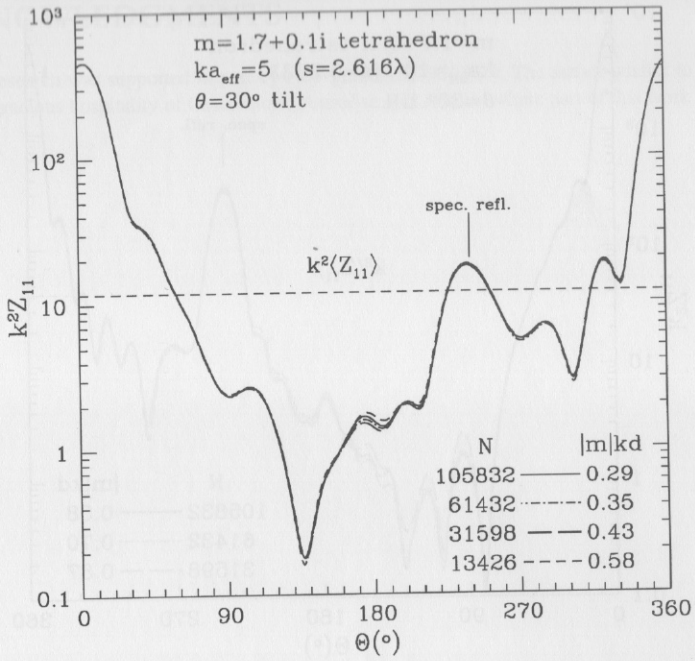

Figure 7 Same as Fig. 4, but for refractive index $m = 1.7 + 0.1i$. The results appear to be well converged for $|m|kd \lesssim 0.5$.

Table I

Timings on Sun Ultrasparc 170 (167 MHz)

$m = 1.33 + 0.01i$ tetrahedron: CPU time (s)			
N	$ka_{\rm eff} = 5$	$ka_{\rm eff} = 10$	$ka_{\rm eff} = 15$
4,030	78	—	—
13,426	311	565	—
31,598	776	1,390	2,280
61,432	1,580	2,850	4,650
105,832	3,600	6,320	11,300

$m = 1.7 + 0.1i$ tetrahedron: CPU time (s)		
N	$ka_{\rm eff} = 5$	$ka_{\rm eff} = 10$
13,426	767	—
31,598	1,920	4,180
61,432	4,080	8,420
105,832	9,590	19,800

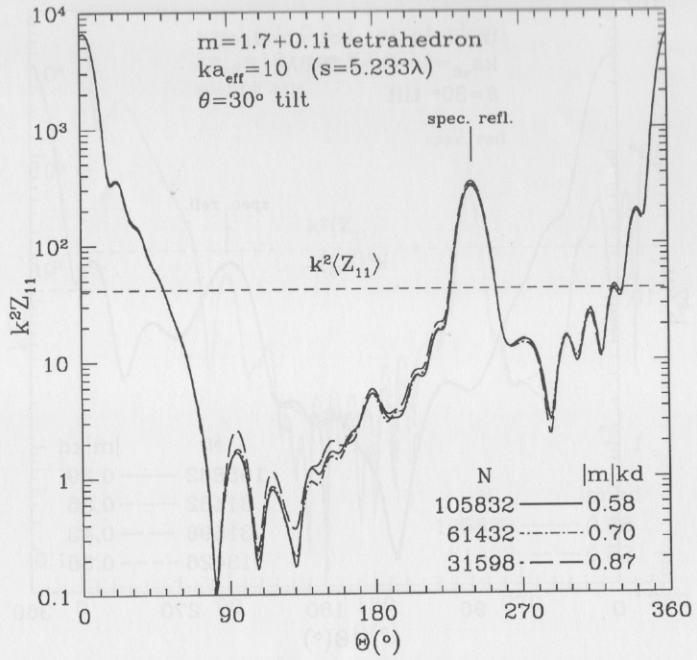

Figure 8 Same as Fig. 7, but for $x = ka_{\rm eff} = 10$. Notice the strong specular reflection peak at $\Theta = 240°$. The good agreement between the results for $N = 61{,}432$ and 105,832 indicates that $N = 105{,}832$ gives a good approximation to the exact results.

XI. SUMMARY

The discrete dipole approximation is an effective technique for computing scattering and absorption by irregular targets, provided the target dimension $D \lesssim 5\lambda$ and $|m-1| \lesssim 3$, where m is the complex refractive index. The portable f77 code DDSCAT can be used to apply the discrete dipole approximation to a broad range of scattering problems. It is quite easy to apply DDSCAT to study new target geometries.

A set of benchmark calculations has been proposed, using tetrahedral targets with $m = 1.33 + 0.01i$ and $m = 1.7 + 0.1i$. It is hoped that other methods for computing scattering from irregular targets will be applied to these test problems, as this will make possible direct comparison of the accuracy and computational demands of the different numerical approaches.

ACKNOWLEDGMENTS

This research was supported in part by NSF grant AST-9619429. The author wishes to acknowledge the gracious hospitality of Osservatorio Astrofisico di Arcetri, where part of this work was completed.