

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

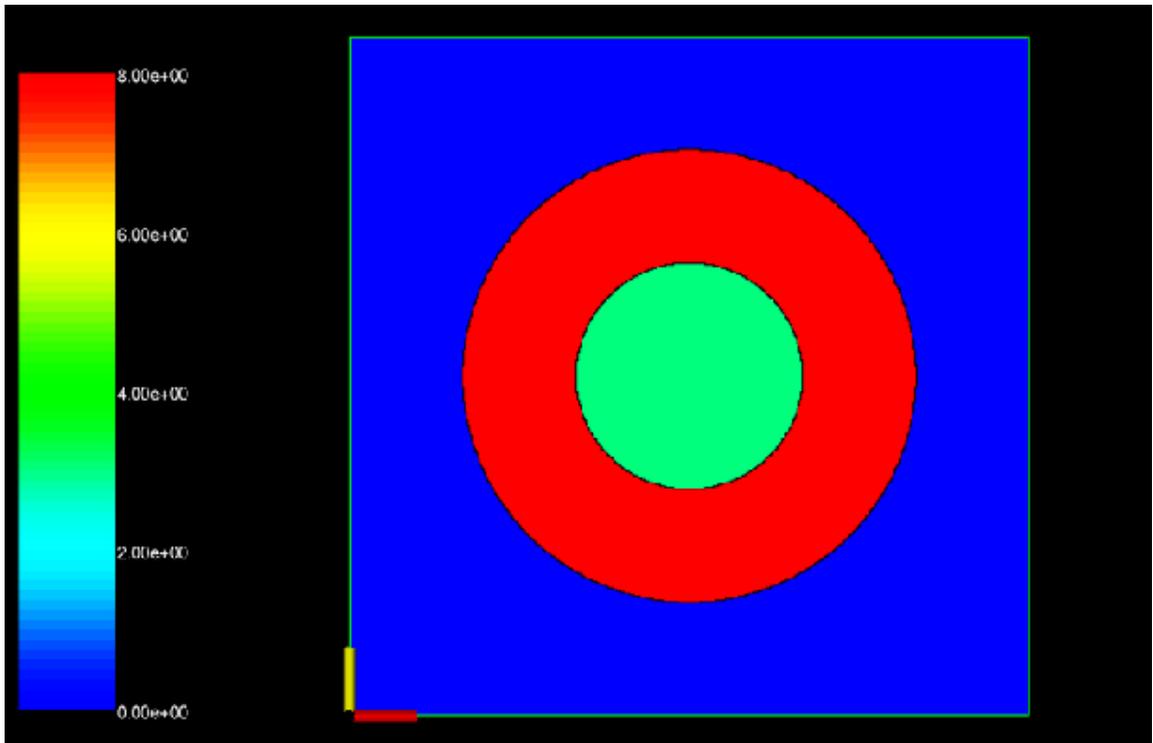
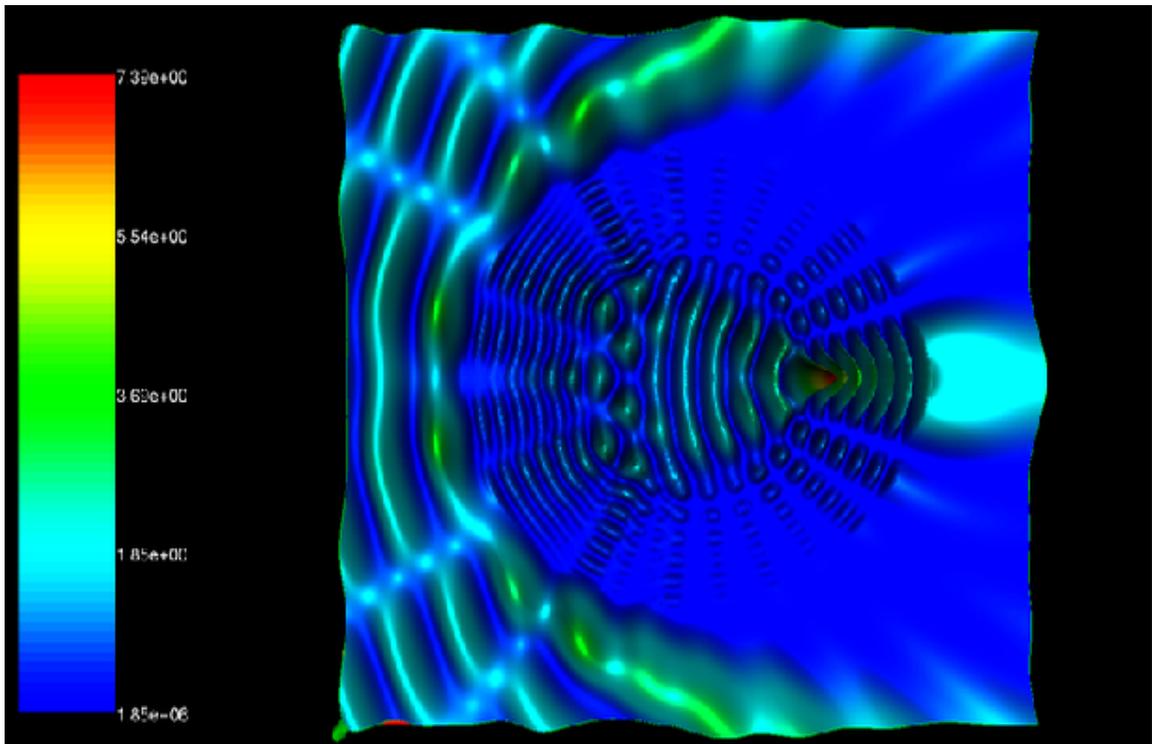
Table 5.2: Convergence for Off-Center Disc – Centered at $(1\lambda, 0)$, Diameter = 1λ

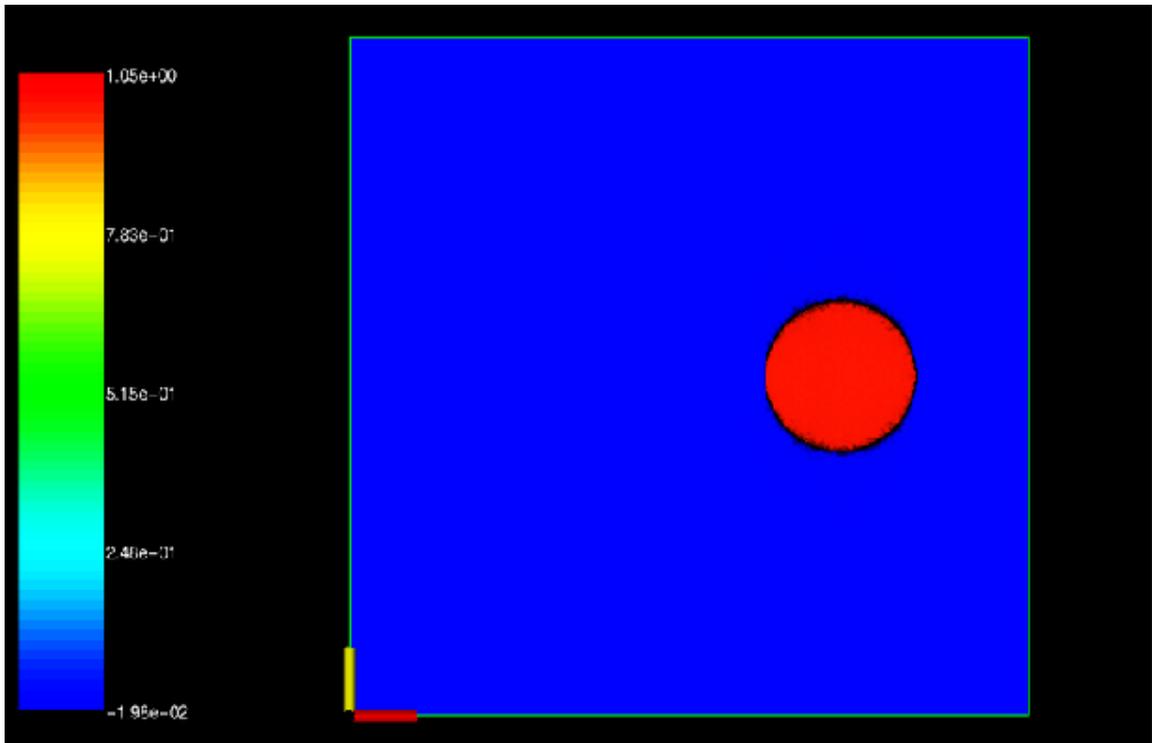
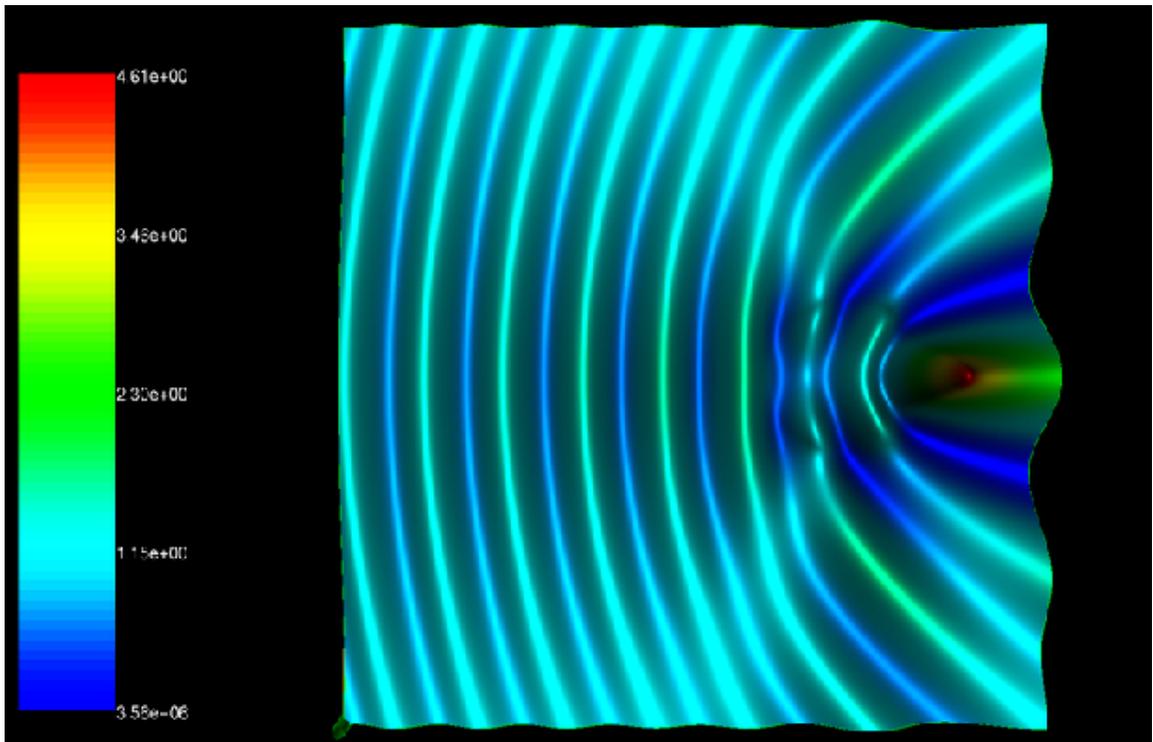
observe that the memory depends linearly on the number of unknowns and the the time required is consistent with an $\mathcal{O}(M \log M)$ complexity.

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

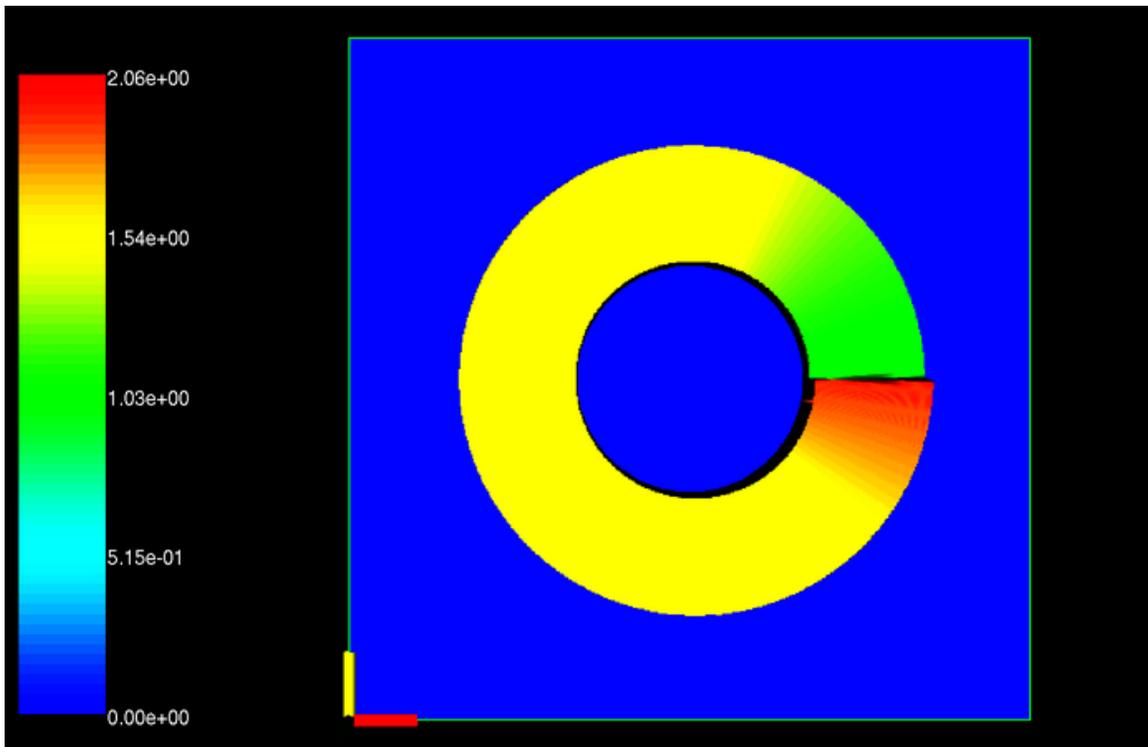
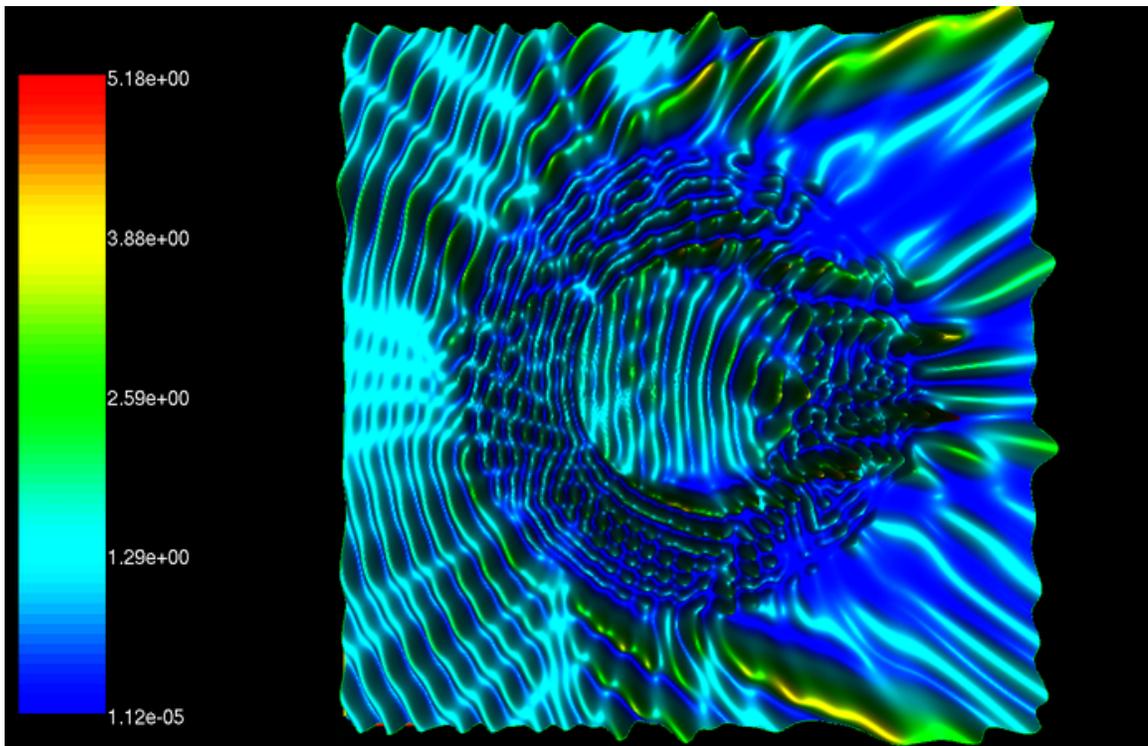
Table 5.3: Convergence Rate for Discontinuous Scatterer – Diameter = 10λ

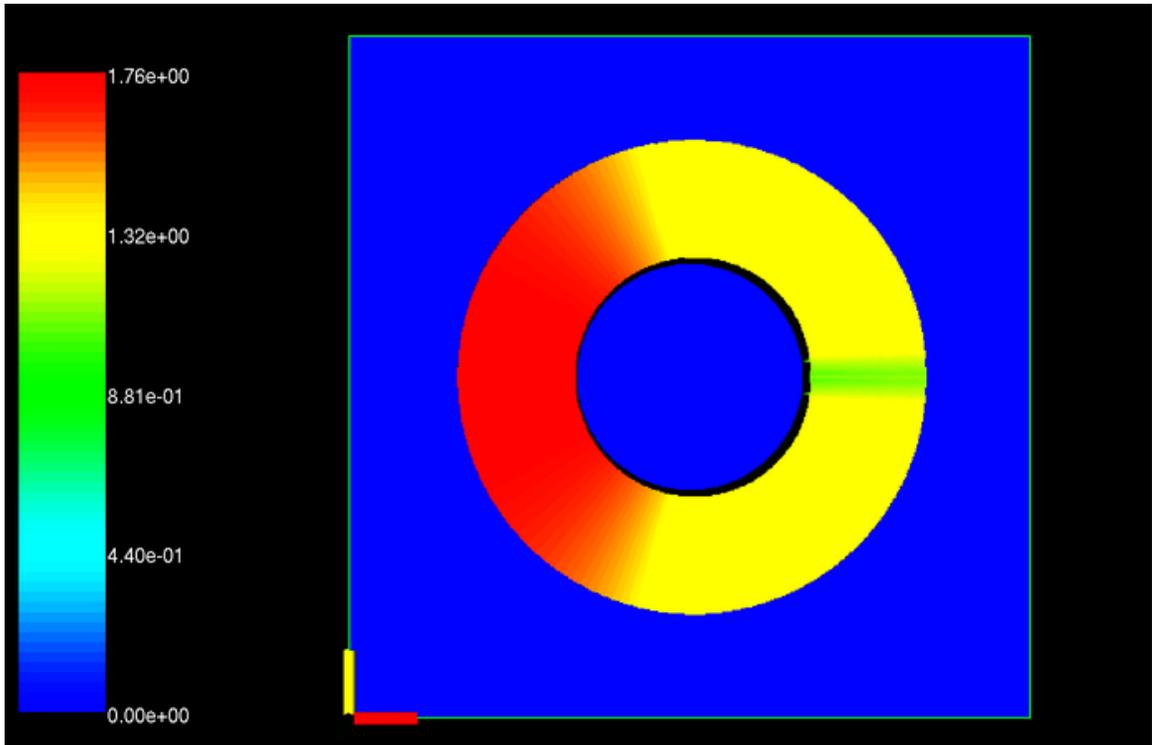
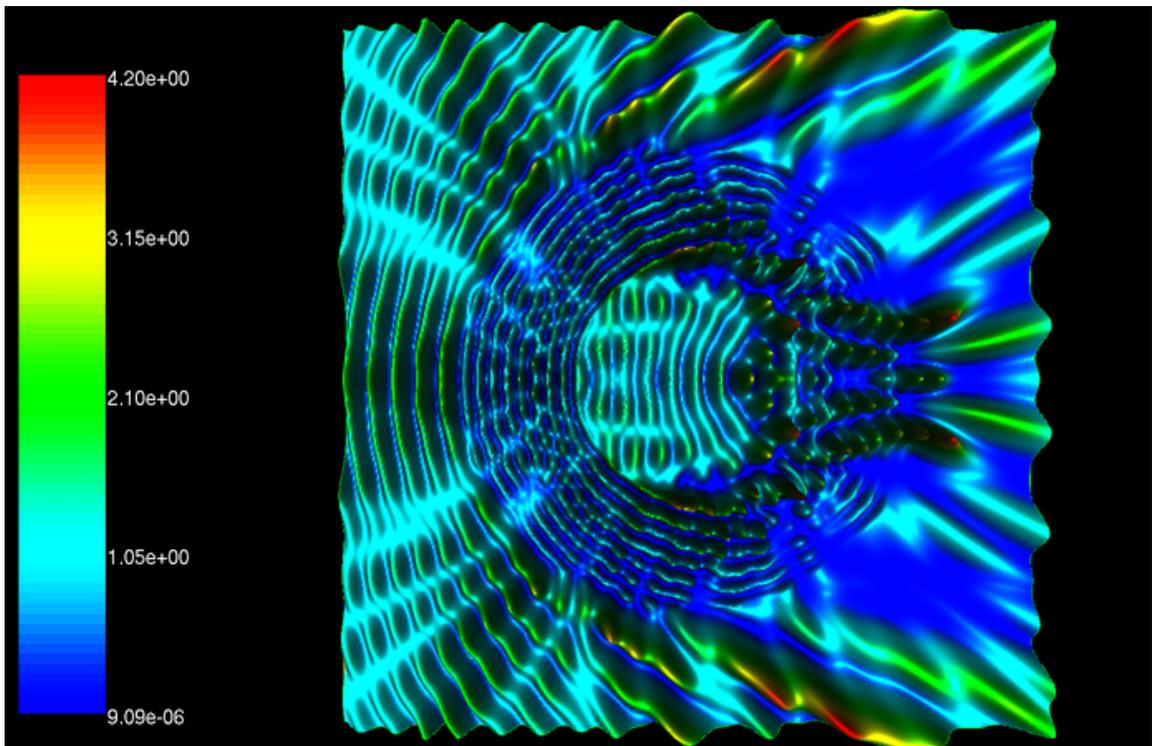
Note that for each integer k , this series becomes a sine or cosine series with real coefficients. If $k = -1$, m is discontinuous and piecewise smooth as a function of θ . Further, for any integer $k \geq 0$, $m \in C^{k,\alpha} \cap C_{pw}^\infty$ as a function of θ . The three examples that follow 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 .

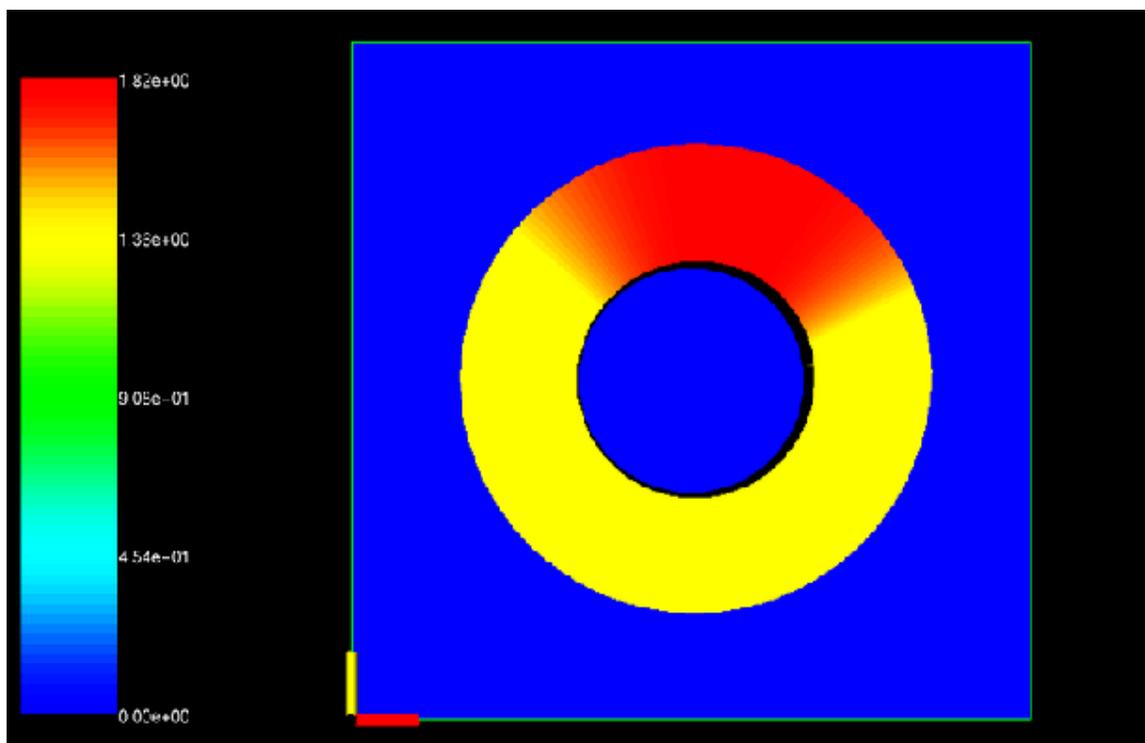
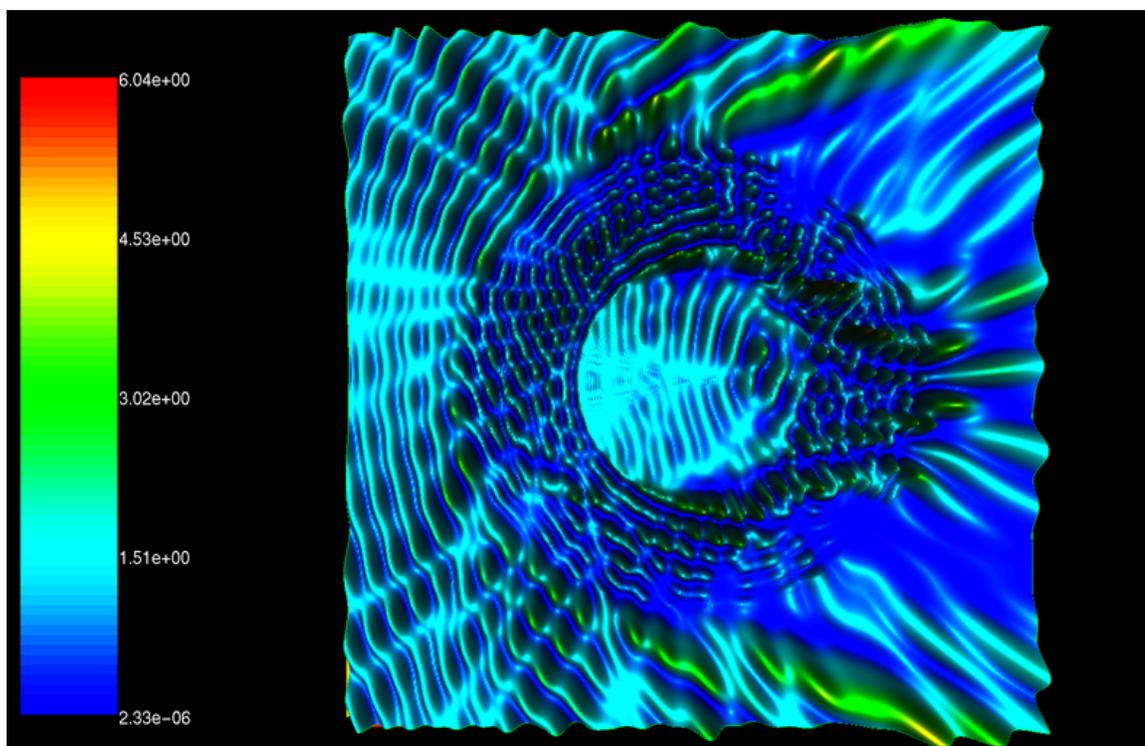
The results for $k = -1$ are found in Figure 5.3 and Table 5.3. The predicted second-order convergence in the near field is exceeded and the third-order convergence in the far field is readily observed. The results for $k = 0$ are found in Figure 5.4 and Table 5.4. In this case, the predicted third-order 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

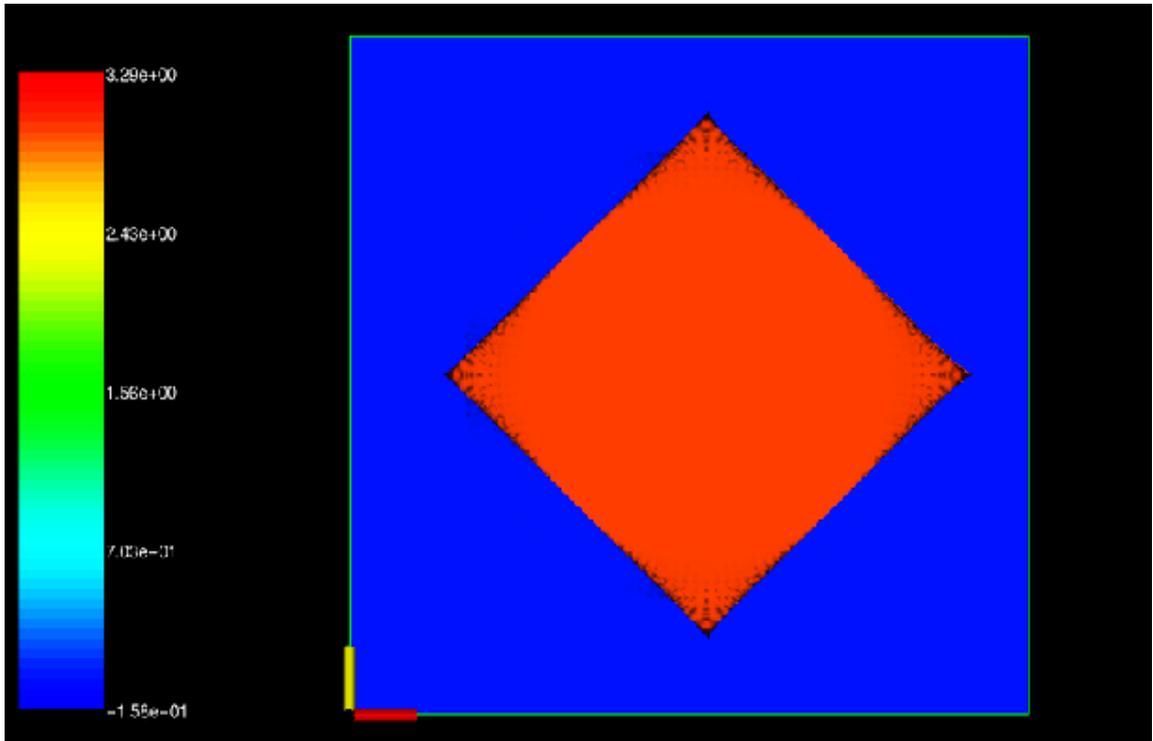
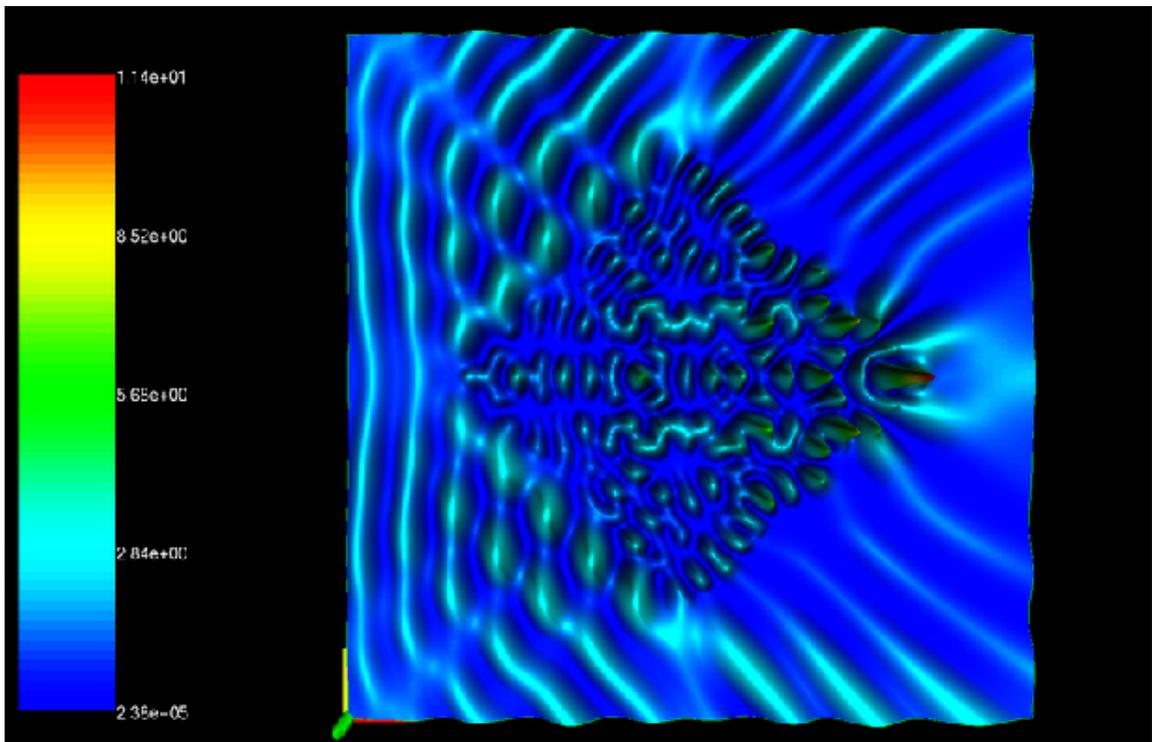
number of GMRES iterations, the memory and the time required with and without the use of our preconditioner. These results were obtained on a 1.7GHz Pentium IV Xeon with 2GB of RAM.

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 $[5\lambda/2\sqrt{2}, 5/2\lambda]$. 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 Gibbs 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

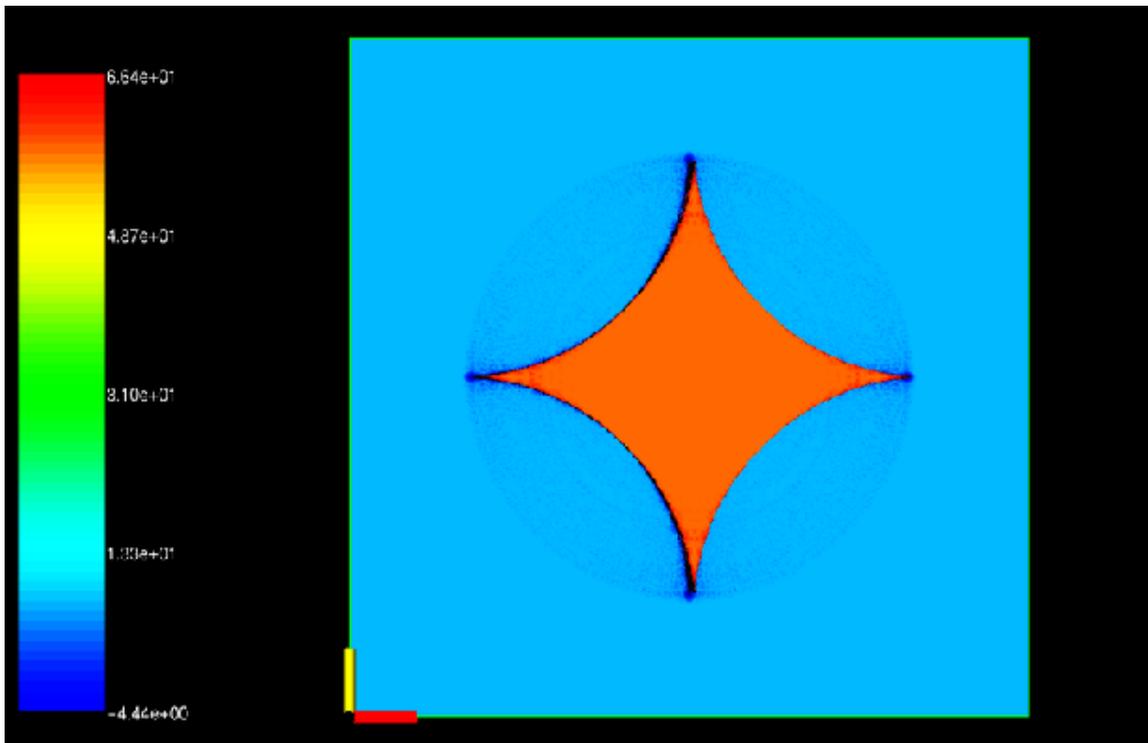
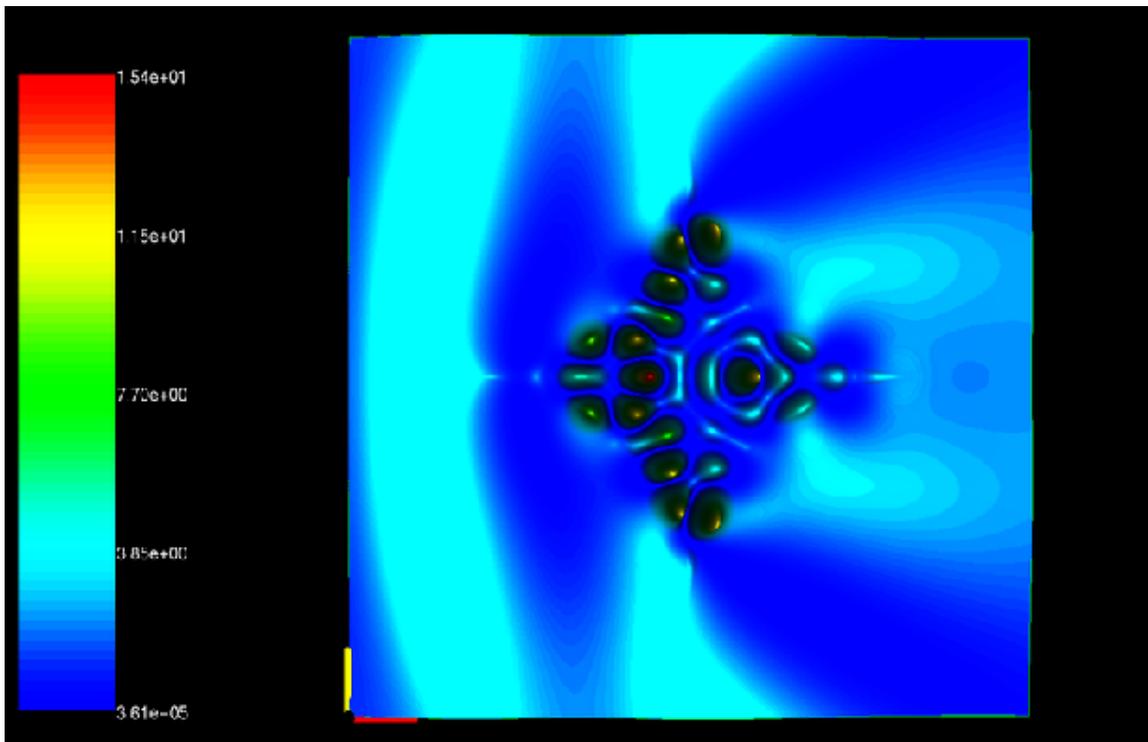
Table 5.9: Convergence for Star Scatterer – Diameter = 1λ

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

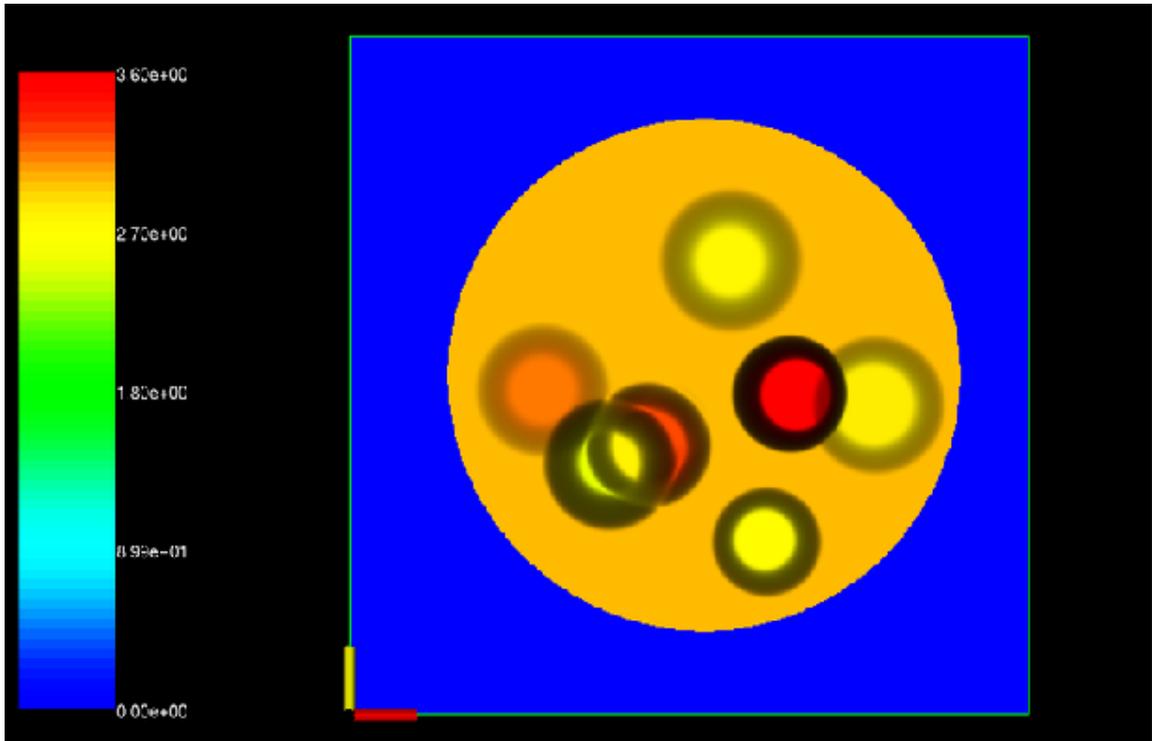
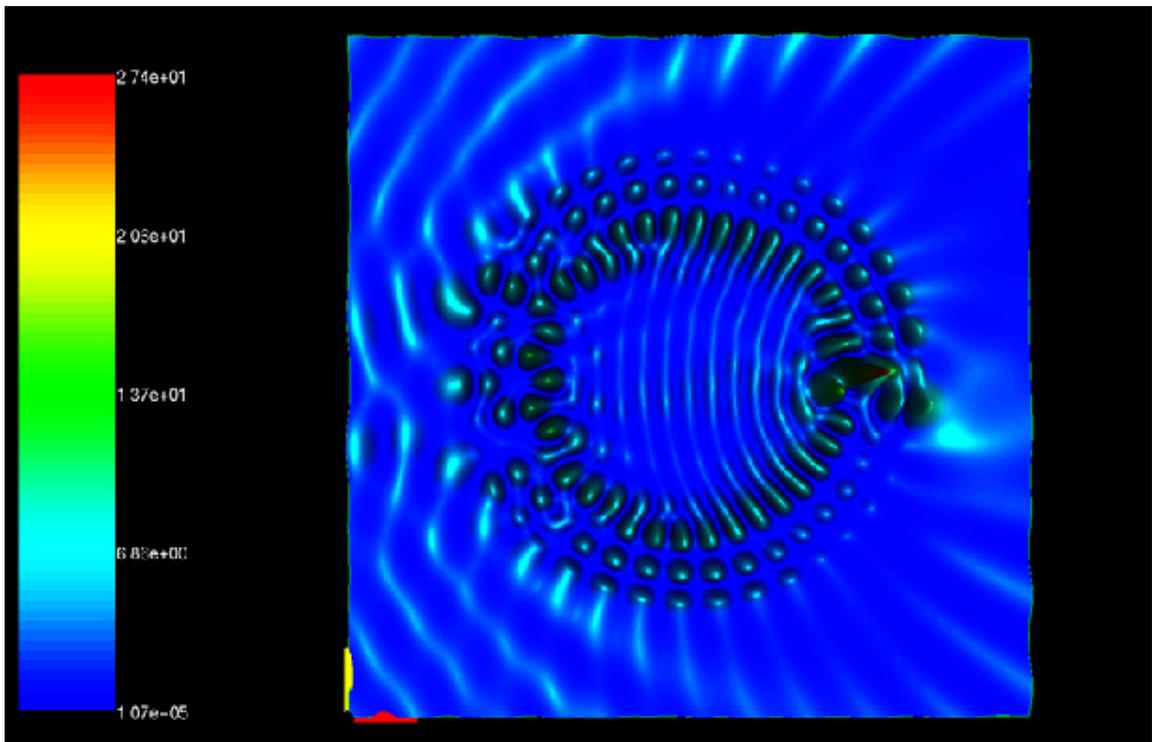
Table 5.10: Convergence for Bumpy Scatterer – Diameter = 5λ

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.

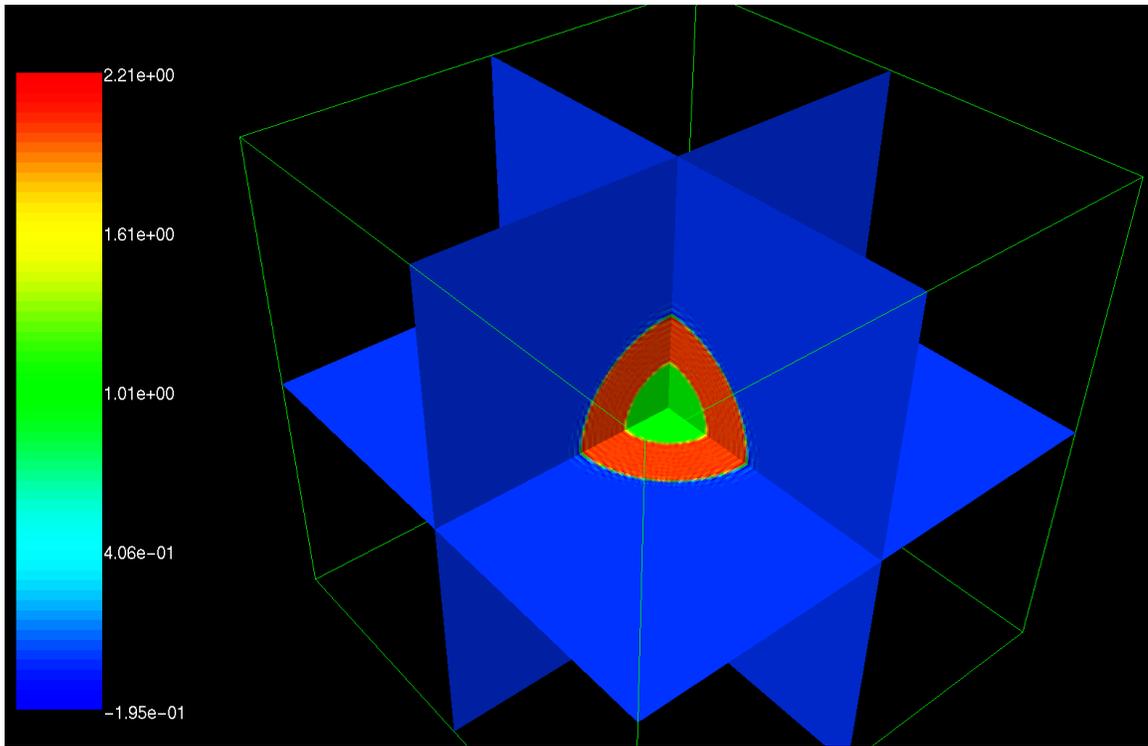
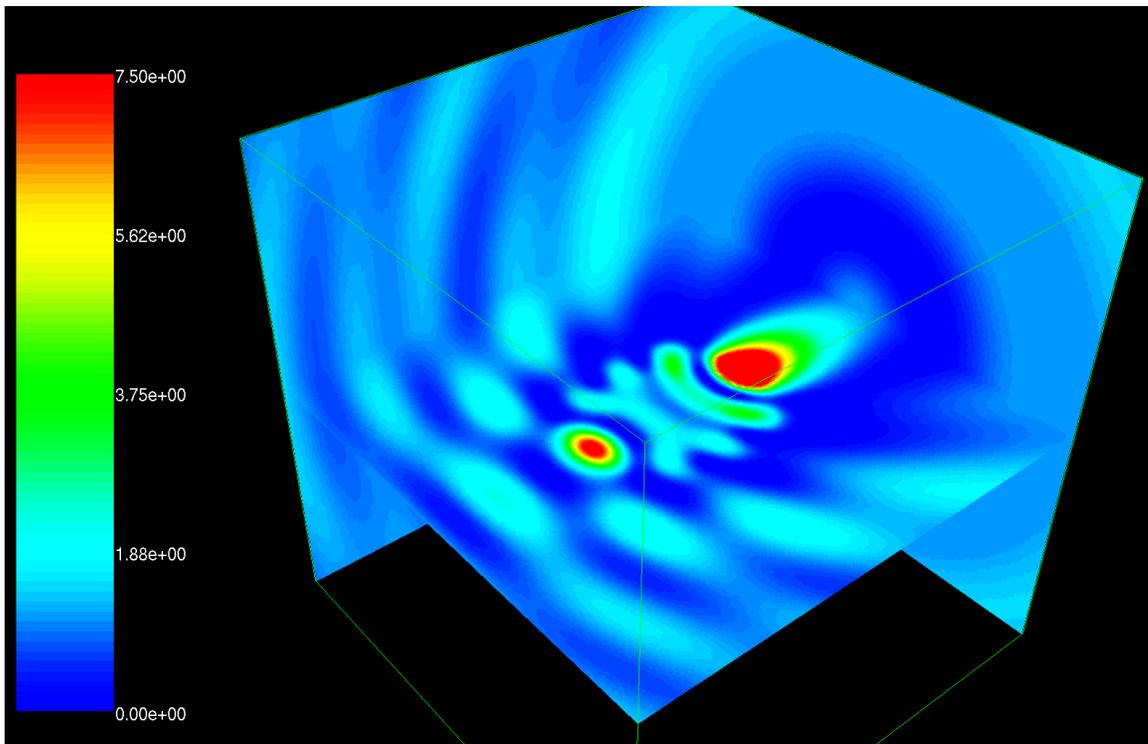
The convergence rate is illustrated by computing the maximum error in the near field solution (ϵ_u^{nf}) as well as in the far field solution (ϵ_u^{ff}). Given the near field solution u , the far field solution u_∞ is computed by means of the following integral [17, p. 223]

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

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 – $\kappa a = 4$

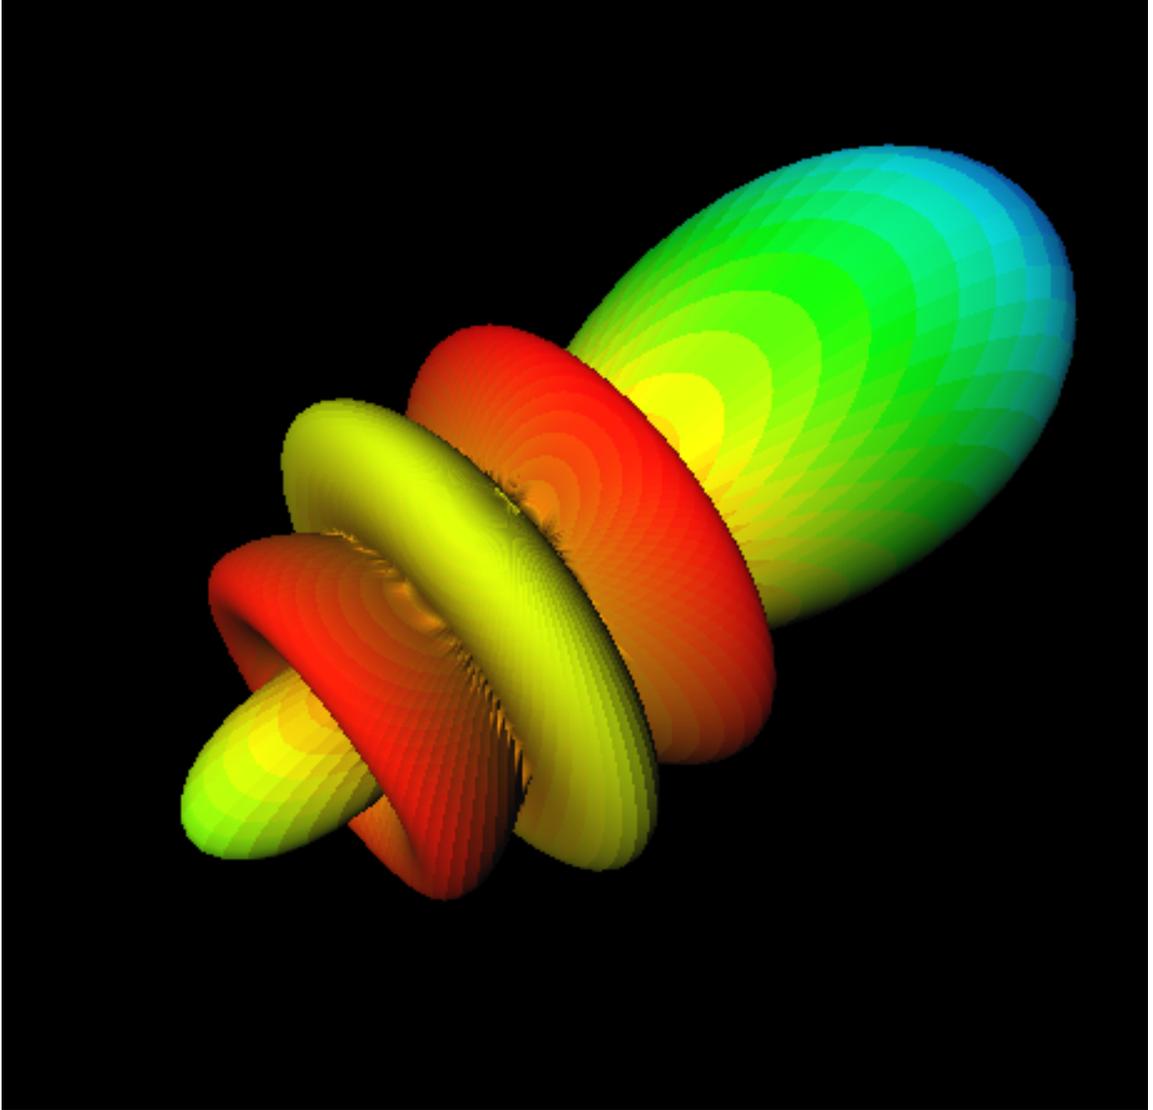


Figure 5.10: Far Field Intensity ($|u_\infty|^2$) for Layered Sphere – $\kappa a = 4$

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

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

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

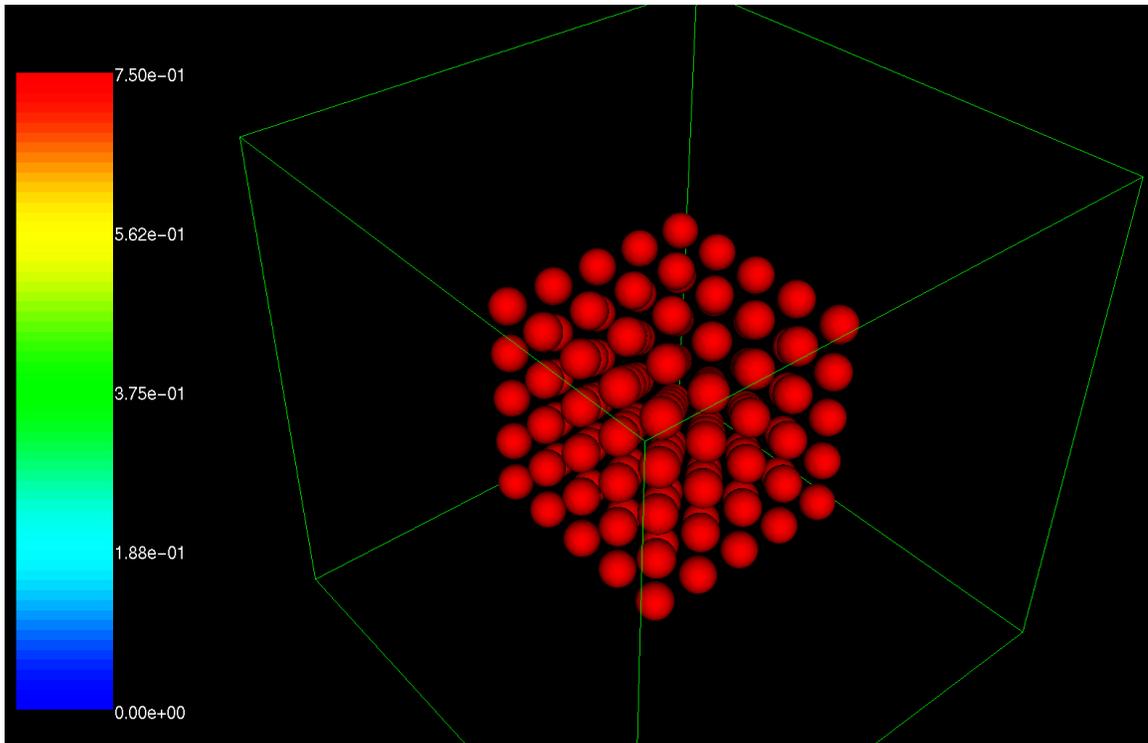
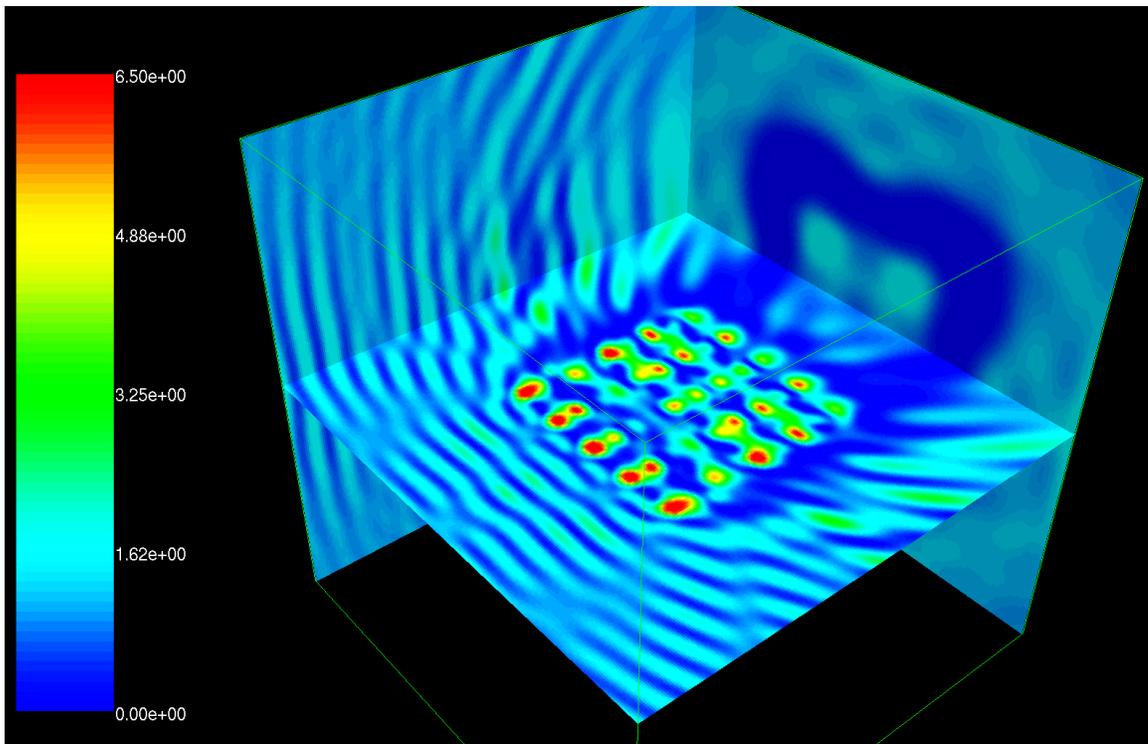
(a) Scatterer ($q = n^2 - 1$)(b) Near Field Intensity ($|u|^2$)Figure 5.11: Visualizations for Array of Smooth Scatterers – $5\lambda \times 5\lambda \times 5\lambda$



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

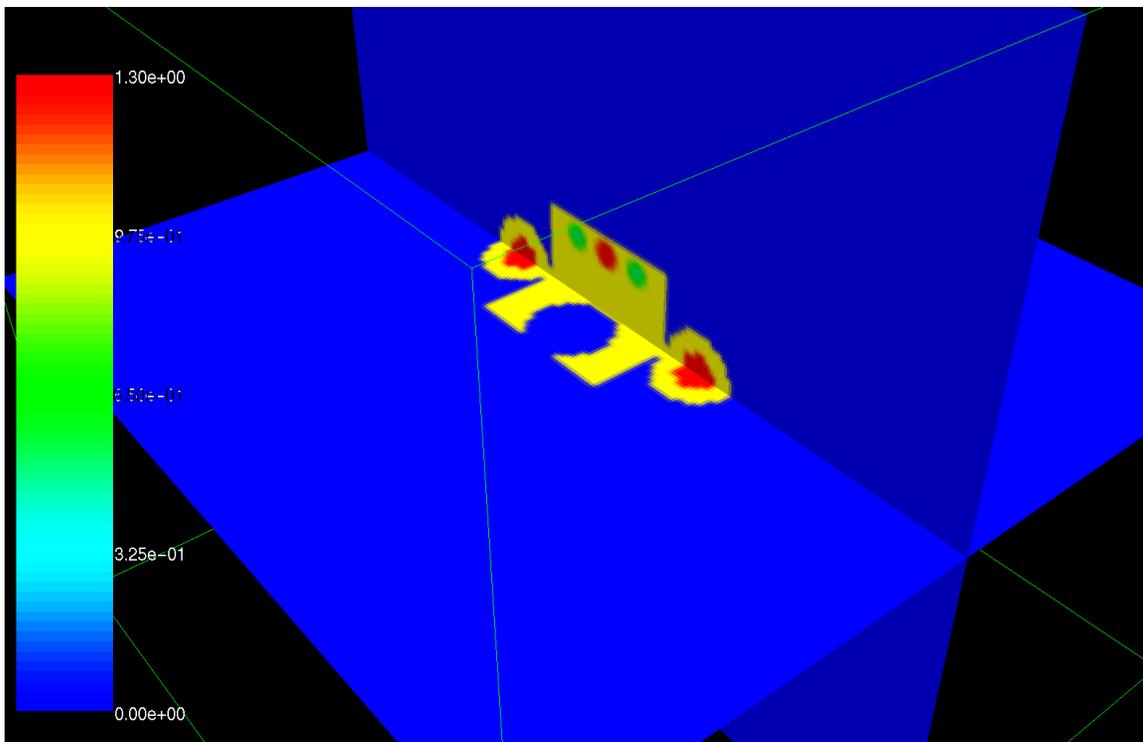
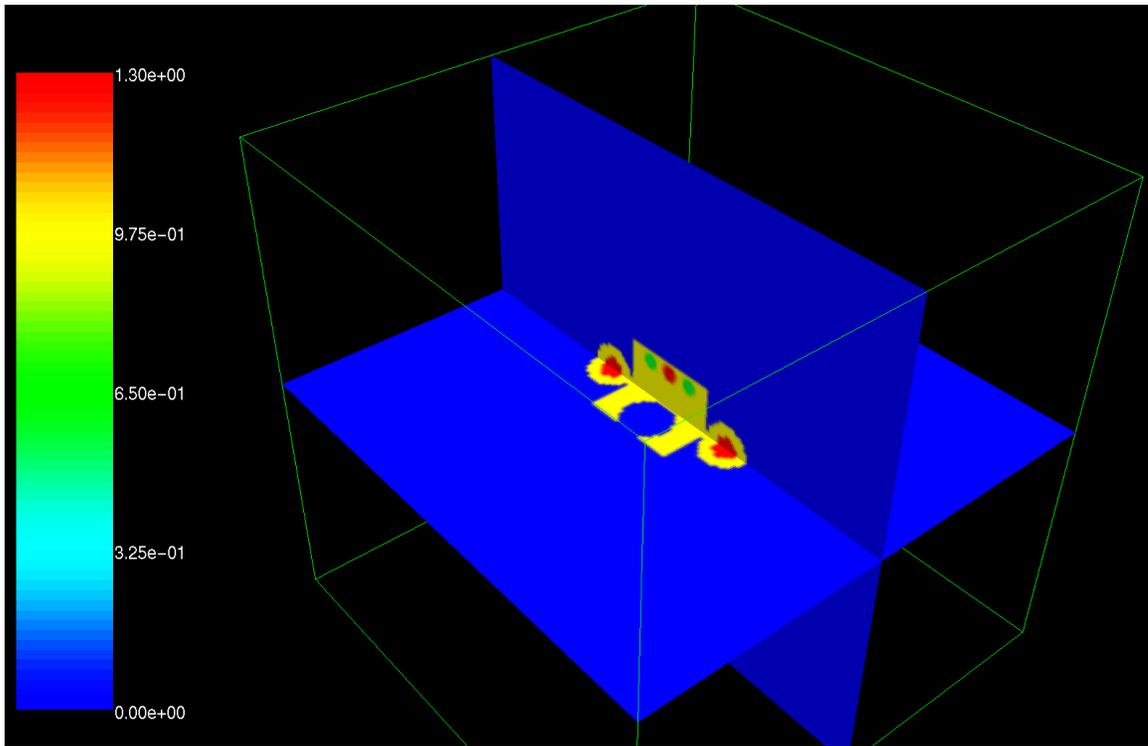
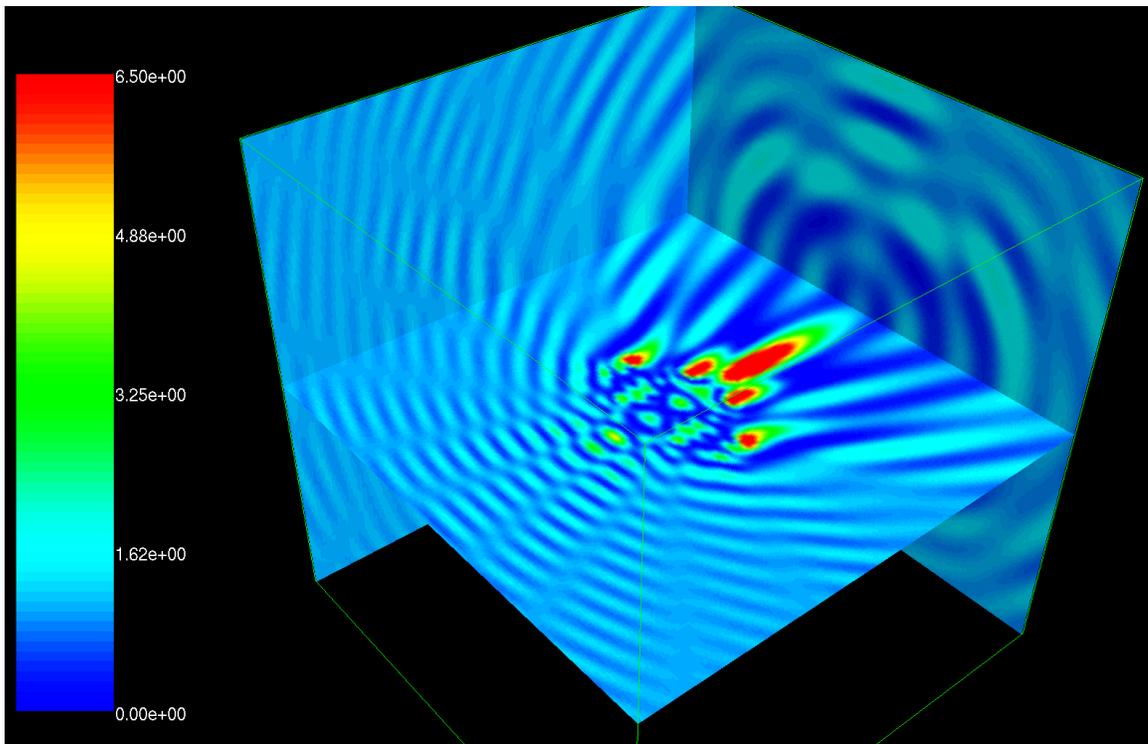


Figure 5.13: Complex Scatterer – Close-Up Image of Scatterer

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 $\approx 2.5\lambda \times 5\lambda \times 2.5\lambda$

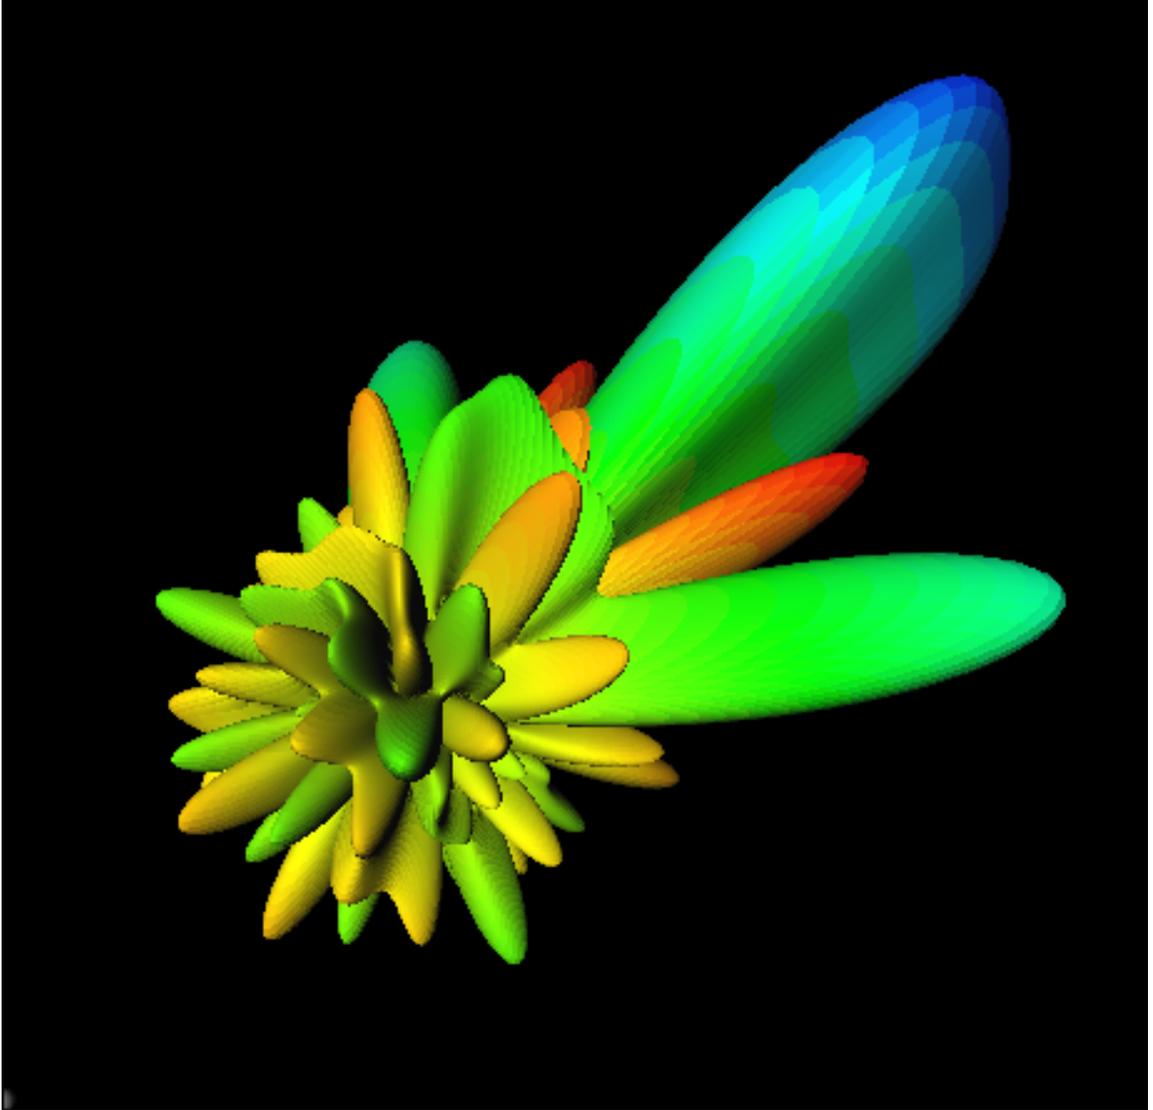


Figure 5.15: Far Field Intensity ($|u_\infty|^2$) for Complex Scatterer – Size $\approx 2.5\lambda \times 5\lambda \times 2.5\lambda$

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 $\approx 2.5\lambda \times 5\lambda \times 2.5\lambda$