CHAPTER 6
HIGHER ORDER APPROXIMATIONS TO THE SCATTERED FIELD

6.1 Introduction

Reconstruction algorithms based on the theory presented in Chapter 2 generate severe artifacts for objects larger than a few wavelengths or a change in refractive index greater than a few percent. These reconstruction algorithms are limited by the first Born or the first Rytov approximations and thus any improvement in the reconstructions will be accomplished by modeling the scattered fields more accurately. With this more accurate model it will again be necessary to invert the equation to arrive at an estimate of the object.

This section will describe iterative techniques that more accurately model the scattered fields. In addition to the theoretical derivations a number of simulations will be presented and the convergence of each series will be discussed.

Two approaches to more accurately compute the scattered field will be discussed here. Both approaches are iterative but they differ widely in their philosophy. The first approach, known as a fixed point method, assumes an initial guess that is a small perturbation from the correct solution. If the scattering integral is a "contracting" operator then this procedure converges to the correct answer. The second approach, based on a matrix formulation of the discrete wave equation, refines an initial guess by projecting it onto hyperplanes. Since each projection reduces the error this method always converges. Unfortunately the matrix formulation is more expensive (requires more calculations) than the fixed point method by a factor of several hundred.

6.2 The Singularity of the Green's Function

In each of the procedures to be discussed in this chapter it is necessary to evaluate an integral of the form

$$\int f(\vec{r}\,')g(\vec{r}-\vec{r}\,')d\vec{r}\,' \tag{6.1}$$

where $g(\vec{r}-\vec{r}\,')$ is the Green's function described in Chapter 2 or

$$g(\vec{r}-\vec{r}\,') = \frac{-N_0(R)+jJ_0(R)}{4} \tag{6.2}$$

and $R = |\vec{r}-\vec{r}\,'|$. This integral in equation (6.1) would be straightforward except for the fact that $N_0(R)$ tends to ∞ as R goes to zero.

The effect of this singularity is further complicated since for computer simulation it is necessary to sample the function at discrete positions. If straightforward sampling is performed on such a function any small change in the location of the sampling grid would cause a large change in the sampled value at the origin.

Fortunately it is possible to derive an approximate value for the sample at the origin. The actual value at the origin is not so important but when the Green's function is multiplied by the sampled kernel and integrated the result should be identical to a sampled version of the continuous integral.

Since all functions have been sampled they can be assumed to be reasonably smooth and the integral approximated

$$\sum_{r_i} f(\vec{r}_i)g(\vec{r}_j,\vec{r}_i). \tag{6.3}$$

When $\vec{r}_i$ is equal to $\vec{r}_j$ the area around the singularity is evaluated and the contribution of the $\vec{r}_j$ term of the summation should be equal to the original integral in a square region around the origin.

All functions are sampled with interval T, in both the x and the y directions, so the following equality is necessary

$$f(\vec{r}_j)g(\vec{r}_j,\vec{r}_i) = \int_{-\frac{T}{2}}^{\frac{T}{2}}\int_{-\frac{T}{2}}^{\frac{T}{2}} f(\vec{r}\,')g(\vec{r},\vec{r}\,')d\vec{r}\,'. \tag{6.4}$$

If $f(\vec{r}_i)$ is smooth enough then it can be considered a constant within the small region of the integral and brought out of the integral. The sample of the Green's function is now written

$$g(\vec{r}_j,\vec{r}_i) = \int_{-\frac{T}{2}}^{\frac{T}{2}} \int_{-\frac{T}{2}}^{\frac{T}{2}} g(\vec{r},\vec{r}\,')d\vec{r}\,'. \tag{6.5}$$

While it is possible to evaluate this integral using approximations to the Bessel functions, a better solution is to evaluate the integral numerically. The approximate expansions for the Bessel functions are only valid for small values of R and are no longer accurate as the sampling interval, T, approaches a quarter of a wavelength.

A better solution is to numerically integrate this function by dividing the original TxT region into a NxN grid of points and summing up each of the samples. The samples of the real part of $g(\vec{r}-\vec{r}\,')$ are shown in Figure 6.1 for a grid size of 4x4 and 16x16. As shown in this figure the point exactly at the origin can be ignored by offsetting the grid so that the singularity at the origin is never sampled. The effect of ignoring this undefined location can then be seen by examining the progression of average values as the number of samples in the T^2 region near the origin is increased. These values are shown in Table 6.1 for a sampling interval of $\frac{1}{4}$ wavelengths. Notice that this is a Cauchy sequence and converges to an average value for the Green's function over the TxT region. It is this value that will be used as the value of the Green's function at the origin in each of the computational procedures to be discussed in the remainder of this chapter.

6.3 Fixed Point Methods

The fixed point method is the most straightforward computational approach to solve an equation. In the remainder of this section it will be used to find a solution of both the Born integral and the Ricatti equation.

Consider a fixed point solution to the equation $x = f(x)$ defined over the region $x \in [a,b]$. Given an initial guess, x_0 a better estimate of the solution, ξ, can be found by [Sto80]

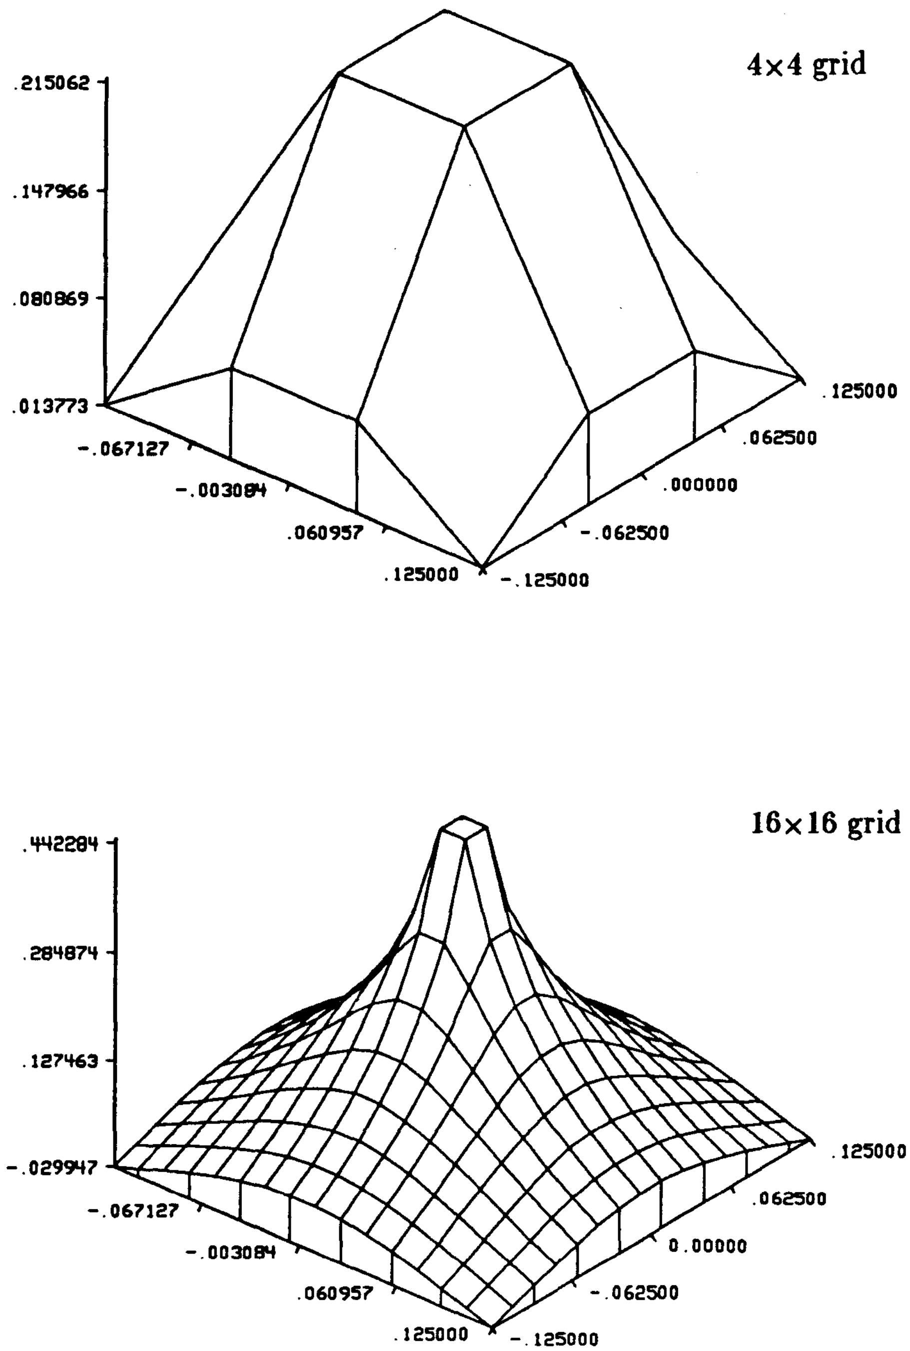

Figure 6.1 The real part of the function g($\vec{r}-\vec{r}'$) is shown here sampled on a 4x4 grid (top) and a 16x16 grid (bottom).

Table 6.1. Average value of the Green's function over a $\frac{1}{4}\lambda$ region sampled on an NxN grid.

N	G(0)
4	0.0925259 + j0.226659
8	0.0927197 + j0.225568
16	0.0927666 + j0.225296
32	0.0927782 + j0.225228
64	0.0927811 + j0.225211
128	0.0927818 + j0.225207
256	0.092782 + j0.225206
512	0.092782 + j0.225206
1024	0.092782 + j0.225206

$$x_1 = f(x_0) \tag{6.6}$$

or more generally

$$x_{i+1} = f(x_i). \tag{6.7}$$

This is illustrated in Figure 6.2.

While this method doesn't always converge, it is possible to show that it will converge to a unique solution if the absolute value of the function's first derivative is always less than one. Mathematically this condition can be written

$$|f'(x)| < 1 \qquad x \in [a,b]. \tag{6.8}$$

If this condition is true then it is also possible to show that f() is a contraction operator or

$$|f(x_1) - f(x_2)| \leq K|x_1 - x_2| \qquad x \in [a,b] \tag{6.9}$$

where $K < 1$.

That these two conditions are equivalent is easy to show by considering two cases. First, assume that f() is a contracting operator. From the definition of the derivative

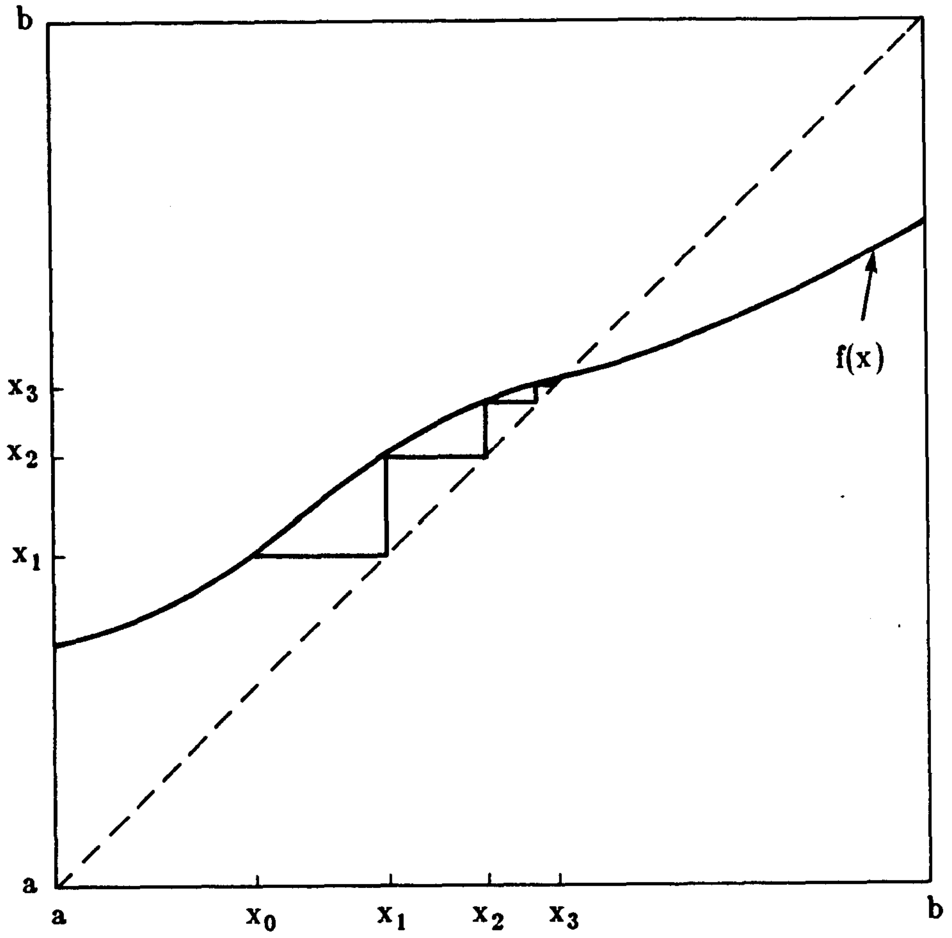

Figure 6.2 An initial estimate for the solution of the equation $f(x) = 0$ is refined using the iteration $x_i = f(x_{i-1})$.

$$|f'(x)| = \lim_{h\to 0}\left|\frac{f(x+h)-f(x)}{(x+h)-x}\right|. \tag{6.10}$$

But from equation (6.9) this ratio and the derivative are less than K.

To show the converse or that

$$|f'(x)|<1 \qquad x\in[a,b] \tag{6.11}$$

implies a contracting operator first assume the opposite case $|f(x_2)-f(x_1)| > |x_2-x_1|$. Without loss of generality assume x_2 is greater than x_1 and write the following integral

$$|f(x_2)-f(x_1)| = \int_{x_1}^{x_2} f'(x)dx. \tag{6.12}$$

There are two cases to consider. First assume that $f(x_2)-f(x_1)>0$ and write

$$0 \leq f(x_2)-f(x_1) = \int_{x_1}^{x_2} f'(x)dx. \tag{6.13}$$

But if $f'(x)<K$ then this integral can be bounded from above by

$$0 \leq |f(x_2)-f(x_1)| = \int_{x_1}^{x_2} f'(x)dx < (x_2-x_1)K. \tag{6.14}$$

This contradicts the original assumption.

For the opposite case, $f(x_2)-f(x_1) < 0$,

$$|f(x_2)-f(x_1)| = f(x_1)-f(x_2) = -\int_{x_1}^{x_2} f'(x)dx. \tag{6.15}$$

But if $f'(x)>-K$ then

$$0 > f(x_2)-f(x_1) = \int_{x_1}^{x_2} f'(x)dx > -(x_2-x_1)K. \tag{6.16}$$

For both cases, $f(x_1)>f(x_2)$ and $f(x_2)>f(x_1)$, the original assumption is contradicted therefore proving that if $|f'(x)|<K$ then the function is a contracting operator.

To study the convergence of the fixed point method to the correct solution, ξ, assume a value for x_i. Then the error after the i+1 iteration is given by

$$\text{Error} = |x_{i+1}-\xi| = |f(x_i)-\xi| . \tag{6.17}$$

But since ξ is a fixed point solution of the function f this expression for the error can be written

$$\text{Error} = |x_{i+1}-\xi| = |f(x_i)-f(\xi)| . \tag{6.18}$$

By the contraction property of Equation (6.9) this last term is bounded from above by

$$|x_{i+1}-\xi| = |f(x_i)-f(\xi)| < K|x_i-\xi| . \tag{6.19}$$

Thus the error after the i+1th iteration will always be less than the error after i iterations. In addition since K is less than one the sequence $|x_i-\xi|$ is bounded from above by the converging geometric series x_0K^i therefore the fixed point method converges to a correct answer.

That this solution is unique can be seen by assuming two solutions, ξ_1 and ξ_2, to the equation $x=f(x)$. But these two solutions violate the contraction property since

$$|\xi_1-\xi_2| = |f(\xi_1)-f(\xi_2)| > K|\xi_1-\xi_2| . \tag{6.20}$$

Thus if f() is a contracting operator then the fixed point solution will converge to the correct answer.

6.3.1 The Born Series

Recall from Chapter 2 the scattered field can be written

$$u_s(\vec{r}) = \int g(\vec{r}-\vec{r}\,')o(\vec{r}\,')u_0(\vec{r}\,')d\vec{r}\,' + \int g(\vec{r}-\vec{r}\,')o(\vec{r}\,')u_s(\vec{r}\,')d\vec{r}\,' . \tag{6.21}$$

The first Born approximation, as already described, simply represents the first iteration of the fixed point method with the initial guess $u_s^{(0)}=0$. Thus the first iteration is written

$$u_s(\vec{r}) \simeq u_B(\vec{r}) = \int g(\vec{r}-\vec{r}\,')o(\vec{r}\,')u_0(\vec{r}\,')d\vec{r}\,' . \tag{6.22}$$

If the kernel, $g(\vec{r}-\vec{r}\,')o(\vec{r}\,')$, is a contracting operator an even better estimate can be found by substituting $u_0(\vec{r})+u_B(\vec{r})$ for $u_0(\vec{r})$ in equation (6.22) to find

$$u_B^{(2)}(\vec{r}) = \int g(\vec{r}-\vec{r}\,')o(\vec{r}\,')[u_0(\vec{r}\,')+u_B(\vec{r}\,')]d\vec{r}\,' . \tag{6.23}$$

In general the i-th order Born field can be written

$$u_B^{(i+1)}(\vec{r}) = \int g(\vec{r}-\vec{r}\,')o(\vec{r}\,')[u_0(\vec{r}\,')+u_B^{(i)}(\vec{r}\,')]d\vec{r}\,' . \tag{6.24}$$

An alternate representation is possible if the total field is written

$$u(\vec{r}) = u_0(\vec{r}) + u_1(\vec{r}) + u_2(\vec{r}) + \cdots \tag{6.25}$$

where

$$u_{(i+1)}(\vec{r}) = \int u_i(\vec{r}\,')o(\vec{r}\,')g(\vec{r}-\vec{r}\,')d\vec{r}\,'. \tag{6.26}$$

By expanding equation (6.24) it is possible to see that the fixed point approximation for the scattered field, $u_B^{(i)}$, is

$$u_B^{(i)}(\vec{r}) = \sum_{j=0}^{i} u_j(\vec{r}) \tag{6.27}$$

and in the limit

$$u(\vec{r}) = u_0(\vec{r}) + u_1(\vec{r}) + u_2(\vec{r}) + u_3(\vec{r}) + \cdots. \tag{6.28}$$

This representation (6.28) has a more intuitive interpretation analogous to the Huygens principle. The Green's function gives the scattered field due to a point scatterer and thus the integral of equation (6.24) can be interpreted as calculating the first order scattered field due to the field u_i. For this reason the first order Born approximation represents the first order scattered field and u_i represents the i'th order scattered field.

A variation of the integral Born series presented here was first described and implemented by Azimi and Kak [Azi83]. In [Azi83] the scattered fields are calculated for an object that consists of multiple cylinders by considering the interactions of the scattered field with the other objects. Using the equations presented in Chapter 5 and [Wee64] and [Mor68] it is possible to calculate the exact scattered field from a cylindrical object illuminated with a plane wave. Unfortunately it is not possible to calculate the exact fields scattered by a pair of cylinders using this approach because the field from one object interferes with the other.

Instead Azimi and Kak propose a multiple scattering approach where the incident field is first scattered against each cylinder and then the scattered fields from each cylinder are propagated to the other cylinder(s) where they are scattered again. This is illustrated in Figure 6.3 where the incident field is denoted by u_0 and the field u_{ijk} denotes the field that has scattered off of object i, then object j and finally object k.

Since the field scattered by a cylinder is not a plane wave the key to this procedure is to calculate the scattered fields along lines between the cylinders. The field along this line can then be decomposed into plane waves [Goo68] and each plane wave scattered separately by the cylinder.

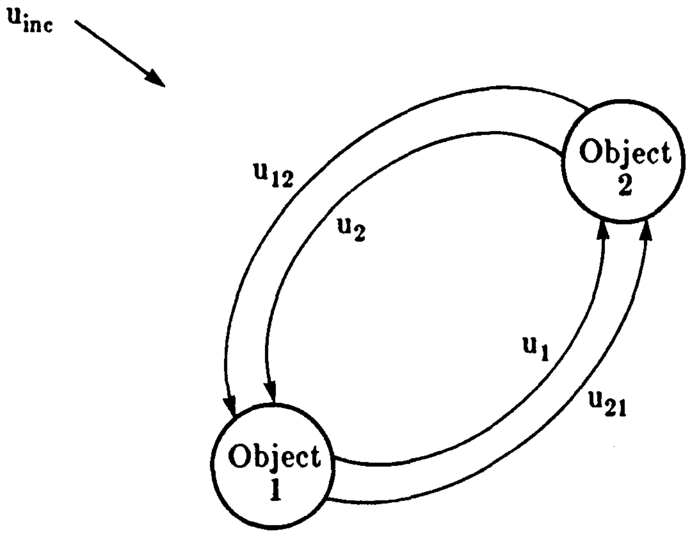

Figure 6.3 The incident field is scattered against each cylinder and the resulting fields are propagated to other cylinder for additional scattering.

While nothing is known about the convergence of this series its use in diffraction tomography is limited because it is only practical when the field scattered from each object can be computed exactly. Thus it has been used as a method to generate data for testing diffraction tomography algorithms.

Under both the Born and Rytov approximations an integral of the form

$$u_{i+1}(\vec{r}) = \int u_i(\vec{r}\,')o(\vec{r}\,')g(\vec{r}-\vec{r}\,')d\vec{r}\,' \tag{6.25}$$

is to be evaluated. A naive approach to this integral would be to evaluate it numerically but doing this requires on the order of N^4 operations. Performing the integration over a 128 x 128 grid, for example, would require over 270 million operations per iteration. This is clearly unreasonable.

The computational requirements can be greatly reduced by noting that the Green's function, $g(\vec{r}-\vec{r}\,')$, is only a function of the difference between the two points and that the integral can also be interpreted as a convolution. Thus representing the convolution in the frequency domain allows for an efficient implementation requiring only $6N^2\log N$ operations to do the integration. For a 128 x 128 grid this is only 700 thousand operations or a reduction in the computational complexity by almost 400. This approach is efficient using an array processor and for an object sampled over a 64 x 64 grid this integral can be computed in under 2 seconds using a Floating Point Systems AP120B array processor.

To evaluate the Born integral on a digital computer it is necessary to use a truncated and sampled version of the object function and the field. Truncating the two functions is error free because the object function is assumed to have finite support. In addition since only a finite number of receivers are present this further limits the number of data points that need to be calculated. On the other hand, since any function with finite support has infinite bandwidth it is not possible to sample the data without introducing errors. For smoothly varying data it is possible to approximate the data with discrete samples. In the work to follow the samples are taken over a rectangular grid using different sampling intervals, all less than a wavelength.

If the series u_i decays to zero then the total field is given by a summation of each scattered field, or

$$u_s(\vec{r}) = \sum_{i=1}^{\infty} u_i(\vec{r}) \tag{6.29}$$

Two different studies were performed to verify this approach to solving the wave equation. Most importantly it was necessary to verify that the total scattered field converged to the same answer as predicted by an exact solution to the wave equation. In addition it was necessary to determine the region of convergence of the Born series. These issues will be discussed later in this chapter.

As mentioned before, the integral equation in (6.25) is efficiently evaluated by implementing the convolution in frequency domain. Recapitulating this discussion, the frequency domain implementation can be summarized as follows.

First for all x_j and y_k, the scattering potential, S, is calculated from the product of the "incident" field and the object,

$$S(x_j,y_k) = U_i(x_j,y_k)O(x_j,y_k). \tag{6.30}$$

Then by using two-dimensional FFT's, determine the following two transforms

$$S(u_j,w_k) = T^2\, \mathrm{FFT}\left\{(S(x_j,y_k))\right\} \tag{6.31}$$

$$G_i(u_j,w_k) = T^2\, \mathrm{FFT}\left\{(G(x_j,y_k))\right\}. \tag{6.32}$$

To find the Fourier transform of the scattered field form the following product in the frequency domain for all u_j and w_k

$$U_{i+1}(u_j,w_k) = S(u_j,w_k)\, G(u_j,w_k). \tag{6.33}$$

The i+1'th scattered field is then found by inverting the above expression to find

$$u_{i+1}(x_j,y_k) = \frac{1}{T^2}\, \mathrm{IFFT}\left\{U_{i+1}(u_j,w_k)\right\}. \tag{6.34}$$

To properly calculate the integral using the FFT approach it is necessary to remember that the discrete multiplication implements a circular convolution. A circular convolution can be turned into an aperiodic convolution by zero padding the data [Opp75]. For example a Floating Point Systems (FPS) AP120B Array Processor with 65,536 words of main data memory can only deal with arrays up to 128 x 128 elements. This memory limitation and the need to implement an aperiodic convolution limits the evaluation of the field to a 64 x 64 grid.

The exact scattered field from a dielectric cylinder with plane wave illumination is well known in the literature. This exact solution to the wave equation can be used to check the results of the Born iteration. Figure 6.4 shows a simulated experiment. In this experiment a plane wave is incident on a cylinder of radius 2λ and a refractive index of 1.1. The scattered field was calculated at 64 receiver positions along a line 7.75λ from the center of the object.

Figure 6.5 show the exact scattered field along the receiver line. In each of the plots that follow the real components of the field will be shown with a solid line, while the dotted lines represent the imaginary part of the field.

Figure 6.6 shows the result of iterating the Born integral. The first Born approximation (Figure 6.6a) gives a very poor estimate of the exact field since this estimate is based only on first order scattering.

In order to accurately calculate the scattered fields it is necessary to include the higher order scattered fields. This is shown in figure 6.6. Clearly in this case even the 30'th to 100'th order scattered fields are significant to the total field.

By comparing figures 6.6 and 6.5 it can be seen that in this case the Born series converges to the exact scattered fields. This simple example shows the correctness of the Born iteration code.

The Born integral defines an infinite series of partial scattered fields that are summed to find the total scattered field. An important measure of any infinite series is its region of convergence. The region of convergence defines a class of objects where the Born iteration converges to the exact scattered field. For an arbitrarily complex object the region of convergence is defined over an infinite dimensional space since an infinite number of parameters are needed to describe the object.

There are two trivial objects that can be analyzed analytically. First consider a pair of point scatterers separated by a distance of R located at $\vec{r}_1$ and $\vec{r}_2$. If each scatter has an area (or volume) of σ and the object function for each of these scatterers is equal to O then an approximation to the first order Born field at $\vec{r}_2$ due to the scatterer at $\vec{r}_1$ is

$$u_i^1(\vec{r}_2) \simeq \sigma G(R) O \tag{6.35}$$

Clearly there exists a value of O such that u_i^1 will have a magnitude greater than the incident field. This field, u_i^1, can then be scattered with the point scatterer at $\vec{r}_2$ and measure the field at $\vec{r}_1$ to find

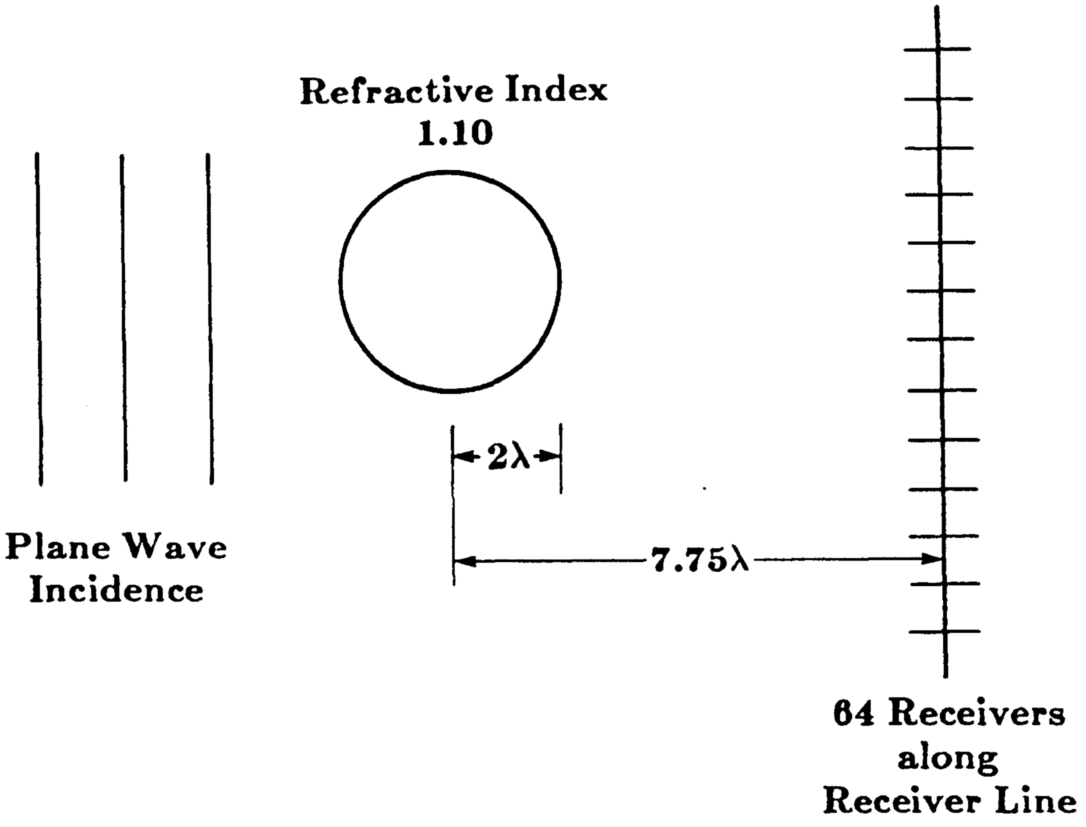

Figure 6.4 An experiment used to illustrate the higher order Born series.

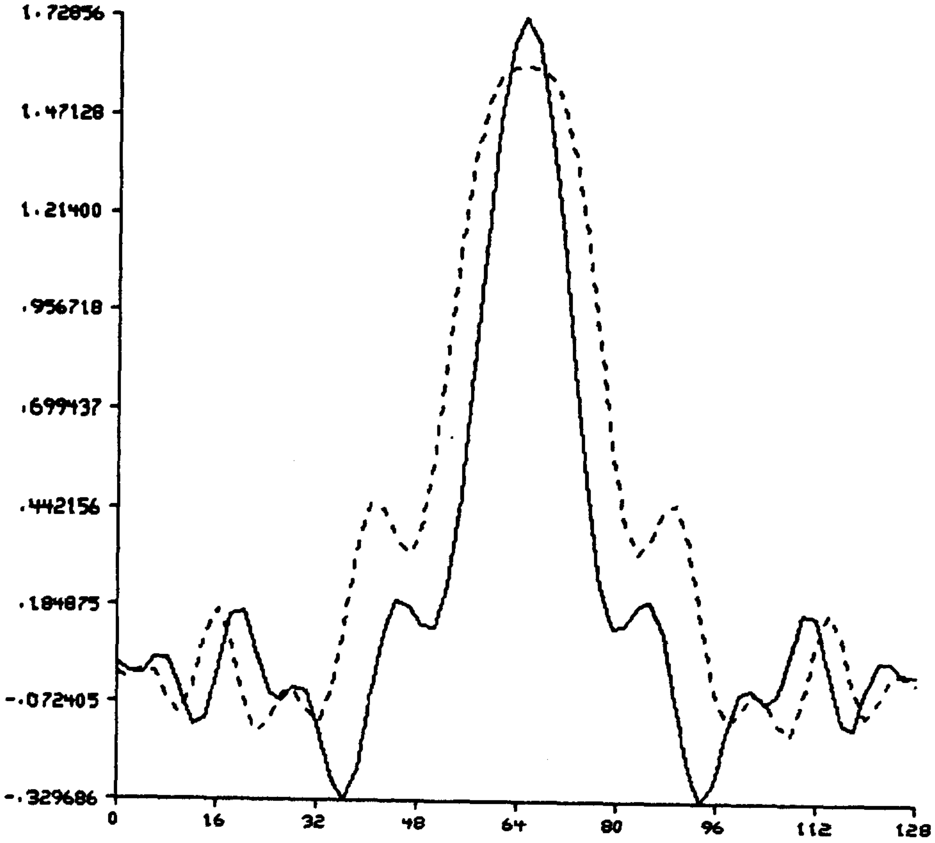

Figure 6.5 Exact scattered field for the object shown in Figure 6.4. The real component of the field is shown as a solid line while the imaginary component is dashed.

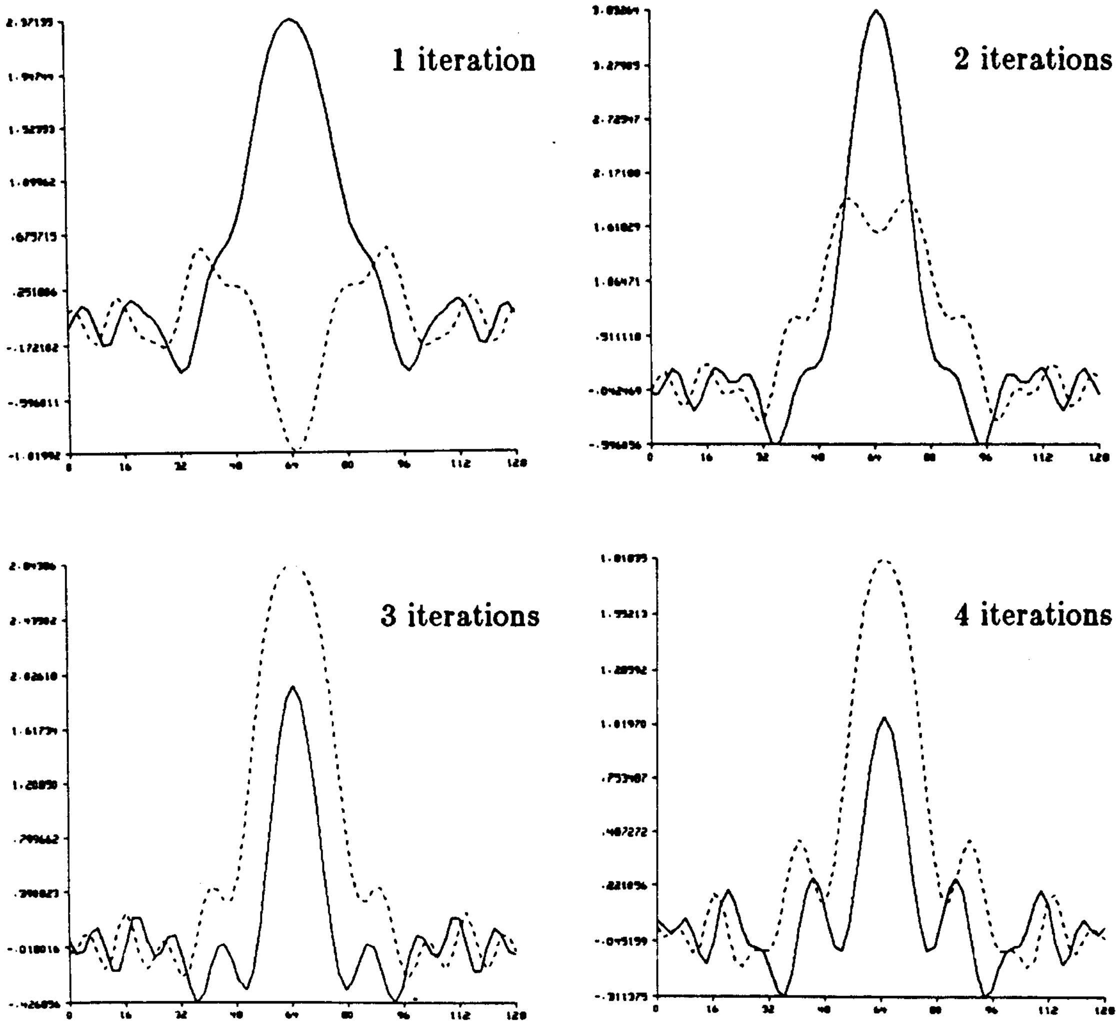

Figure 6.6 The scattered field as calculated from the Born series using terms numbered from 1 to 100.

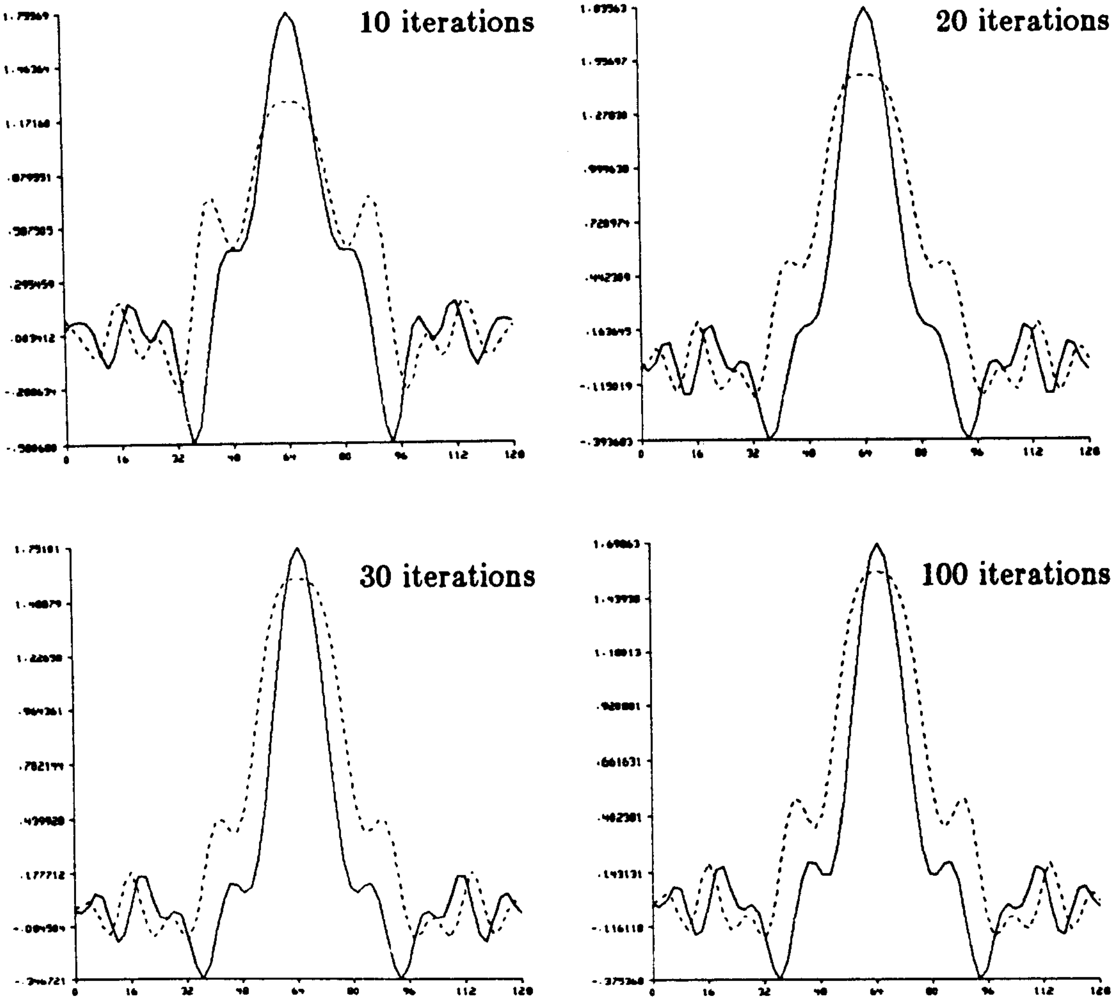

Figure 6.6 Continued.

$$u_2^{12}(\vec{r}_1) \simeq \sigma G(R)O[u_1^1(\vec{r}_2)] \tag{6.36}$$

If the term $\sigma G(R)O$ is greater than one then each succeeding term of the series u_i will diverge. This analysis has ignored the effect of the other scatterer and the field at $\vec{r}_1$ caused by the scatterer at $\vec{r}_1$ but this doesn't change the basic conclusion, that the Born series can converge or diverge depending on the object.

This effect can also be analyzed in the frequency domain. By taking the Fourier transform of equation (6.25) the convolution with the Green's function can be expressed as a multiplication in the spatial frequency domain. Equation (6.25) becomes [Sla83, Sla84]

$$U_{i+1}(\vec{K}) = G(\vec{K})\left\{\tilde{O}(\vec{K})*U_i(K)\right\}. \tag{6.37}$$

where '*' represents convolution in the frequency domain.

If the object function is assumed to be constant for all space then $O(\vec{K})$ becomes an impulse and equation (6.37) above becomes

$$U_{i+1}(\vec{K}) = OG(\vec{K})U_i(K). \tag{6.38}$$

This is a simple result and it is easy to see that the series, u_i, will diverge if there is any frequency where

$$G(\vec{K})O > 1. \tag{6.39}$$

In this simulation study the region of convergence for a single homogeneous cylindrical object was examined. Since any cylinder can be completely described by its size and refractive index the region of convergence is defined over a two dimensional space. More complicated objects could be studied but the results would not be as graphical.

It has already been shown that for either small cylinders or small changes in the refractive index the first Born approximation provides a good estimate to the exact scattered field. This is equivalent to saying that the higher order scattered fields decay quickly towards zero. If is easy to see that under these two conditions the Born integral series will quickly converge to the exact solution.

The total energy in the two dimensional field is used as a simple measure of convergence. Obviously if the total energy in the i'th scattered field is decaying towards zero as the Born integral is iterated then the series is converging. On the other hand if the total energy is increasing then the Born integral can not possibly converge.

For a given object radius the region of convergence is determined by conducting a binary search for the largest refractive index where the Born series converges. For each combination of object size and refractive index it is necessary to make a decision of convergence or divergence and adjust the refractive index accordingly.

The decision of convergence or divergence is made by studying the total energy in the partial scattered field as the iterations are performed. For this purpose a series is defined to be convergent if during each of four iterations the total energy in the partial scattered fields is monotonically decreasing. While if the energy is monotonically increasing then it is decided that the series diverges. As long as the last four terms are not monotonic than the iterations continue. The energy versus iteration number for the experiment of figure 6.4 is shown in figure 6.7.

Figure 6.8 displays the region of convergence for sampling intervals of 0.125λ, 0.25λ and 0.5λ. Each plot shows the maximum refractive index as a function of cylinder radius. For all experiments with a refractive index below the line the Born series converges, while for all experiments above the line the series diverges.

In each case the shape of the curve agrees with the original observation; the Born series converges for either small objects or small changes in the refractive index. The dependence of the region of convergence on the sampling interval is still under study. One possible explanation is that the numerical errors are larger for the larger sampling intervals.

While more complicated objects can be simulated it is more difficult to present the region of convergence in a simple fashion. In general the region of convergence will be described over a multidimensional space but can be reduced to two dimensions be keeping some of the parameters fixed. Two simple families of object that can be reduced to a two dimensional space will be described next. In each case the results will be compared to the results for a cylindrical object.

A simple extension of the previous work for a single cylinder is to consider two cylinders separated by a small distance. Figure 6.9 shows the region of convergence for two orientations of the cylinders with respect to the incident field. In both cases two identical cylinders are separated by a distance of $1\ \lambda$. Thus the region of convergence is reduced to two dimensions, the radius of the

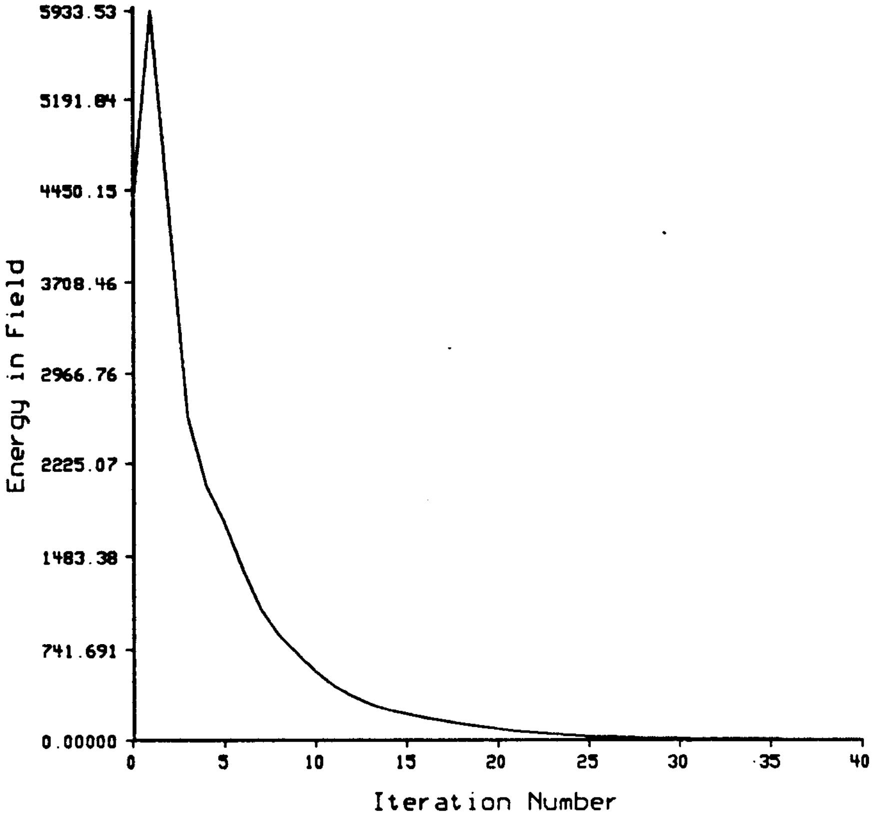

Figure 6.7 Total energy in the higher order scattered field versus iteration number.

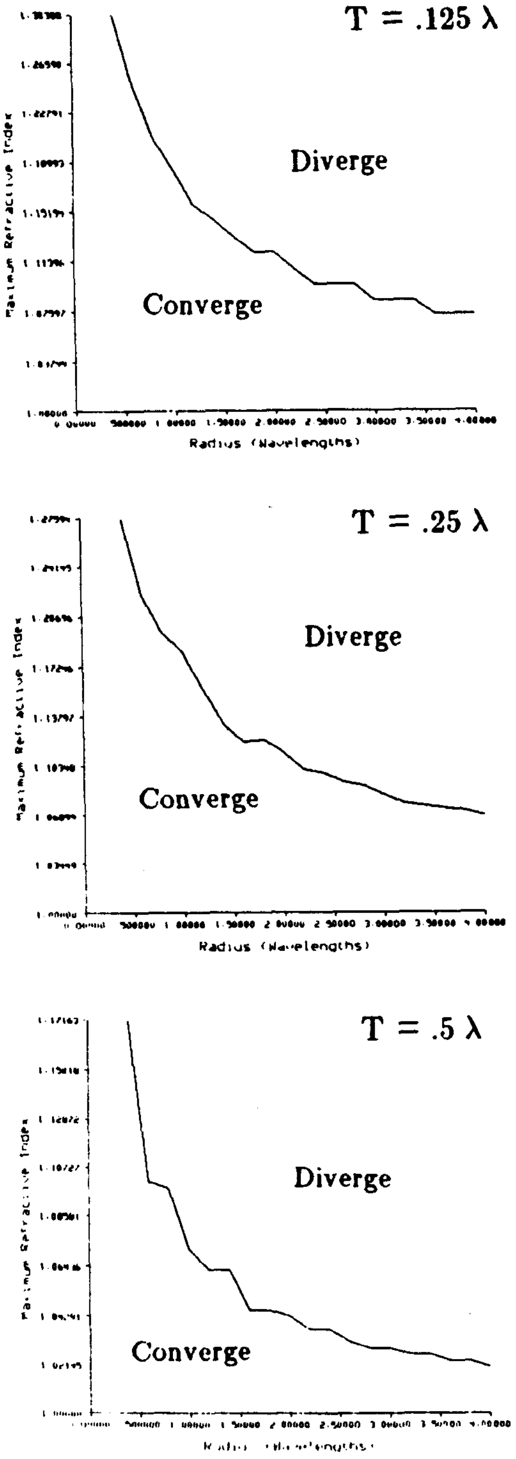

Figure 6.8 Region of convergence for the Born series using sampling intervals of .125λ, .25λ and .5λ.

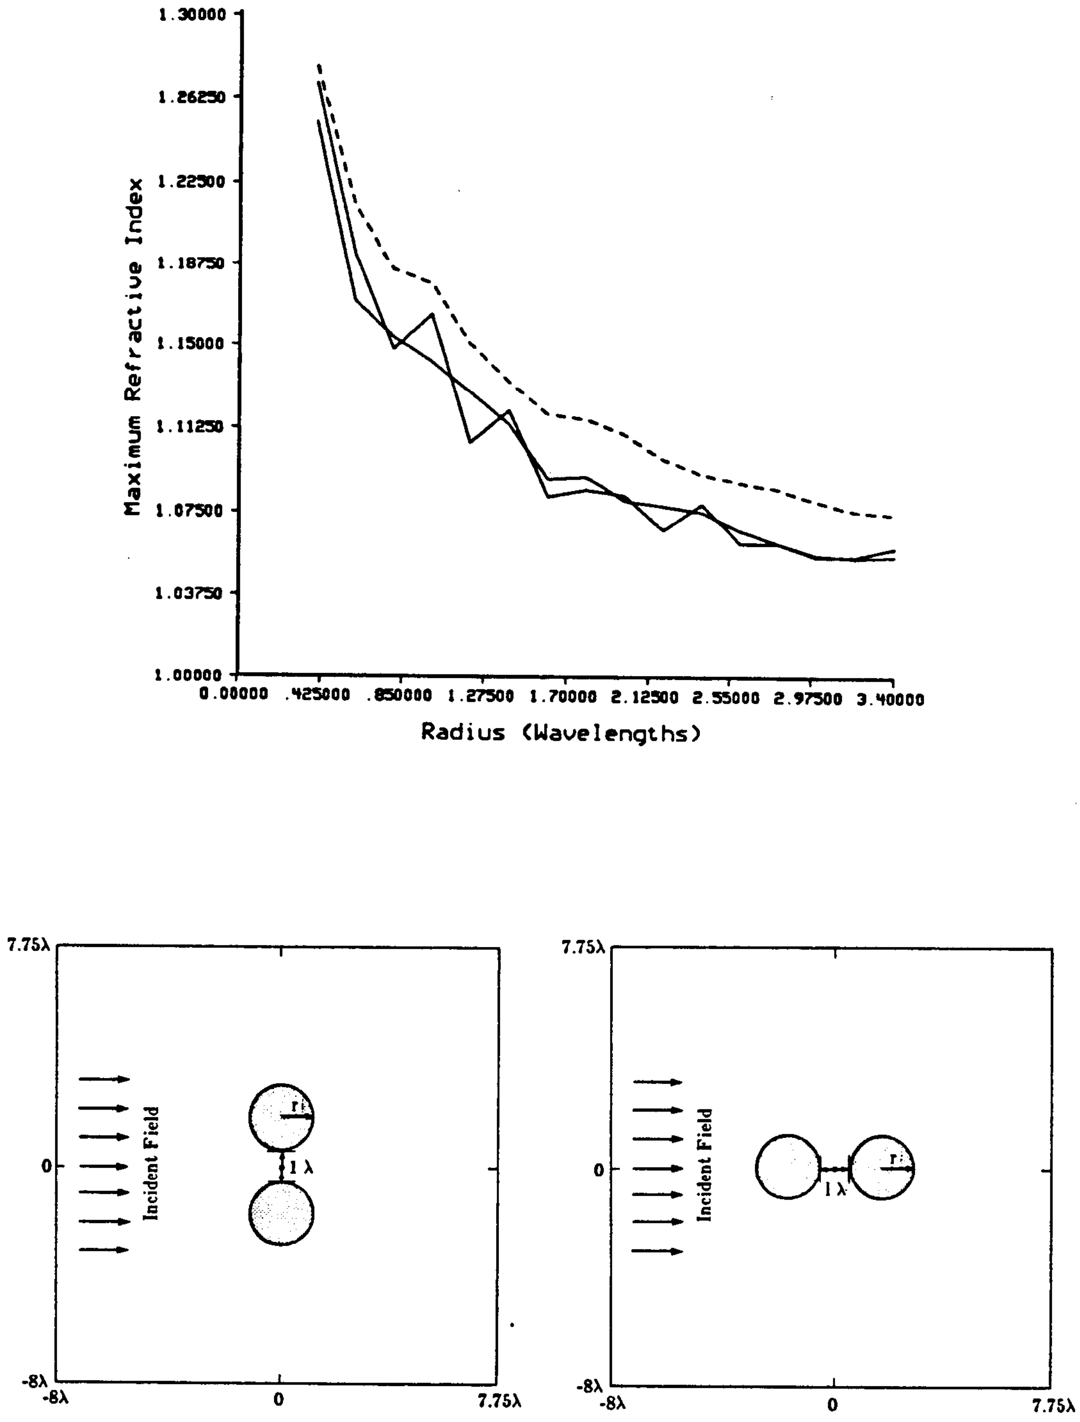

Figure 6.9 Born series region of convergence for two cylinders of radius r separated by 1λ. The solid lines show the convergence for two orientations. The dashed line shows the convergence for a single cylinder of radius r.

two cylinders and their refractive index. The solid lines in Figure 6.9 show the region of convergence for the two cylinders compared to that for a single cylinder (shown as a dashed line).

As would be expected the region of convergence for two cylinders is always smaller than that for a single cylinder. This is true since the convergence is plotted as a function of the radius of one of the cylinders and adding another cylinder increases the scattered field. As already described increasing the scattered field can only cause the Born series to converge more slowly.

Figures 6.10 and 6.11 show the region of convergence for a single ellipse with an eccentricity of 2 (length of major axis over minor axis is 2.) Since an ellipse can not be described by its radius the region of convergence is plotted as a function of the length of both the major and minor axis of the ellipse. The two figures differ only in the orientation of the ellipse with the incident field and in both cases the convergence for a single cylinder as a function of its radius is plotted as a dashed line.

The convergence of the Born series for an ellipse (solid lines) compared to a cylinder (dashed lines) is consistent with the idea that more scattering leads to the divergence of the Born series. The upper solid lines show the convergence plotted as a function of the ellipse's major axis and is above the dashed lines for all lengths. This is because an ellipse with major axis of length 2r has less area than a circle of radius r. Conversely the lowest line (solid) plots the convergence as a function of the ellipse's minor axis and is always below the dashed lines. (Note that the two solid lines represent the same ellipse. The only difference is that region of convergence is plotted with respect to different parameters.)

Figure 6.12 shows the same data as the previous two figures but now the two orientations are superimposed. From this figure it is easy to see that the convergence of the Born series is sensitive to the orientation of the ellipse. When the major axis of the ellipse is parallel to the direction of the incident field the region of convergence is reduced. This is consistent with the limitation described in Chapter 2 that the phase change in the field as it travels through the object is a good indication of the validity of the Born approximation. When the major axis is aligned with the incident field the phase change is larger than when it is perpendicular to the field and thus the Born series is more likely to diverge.

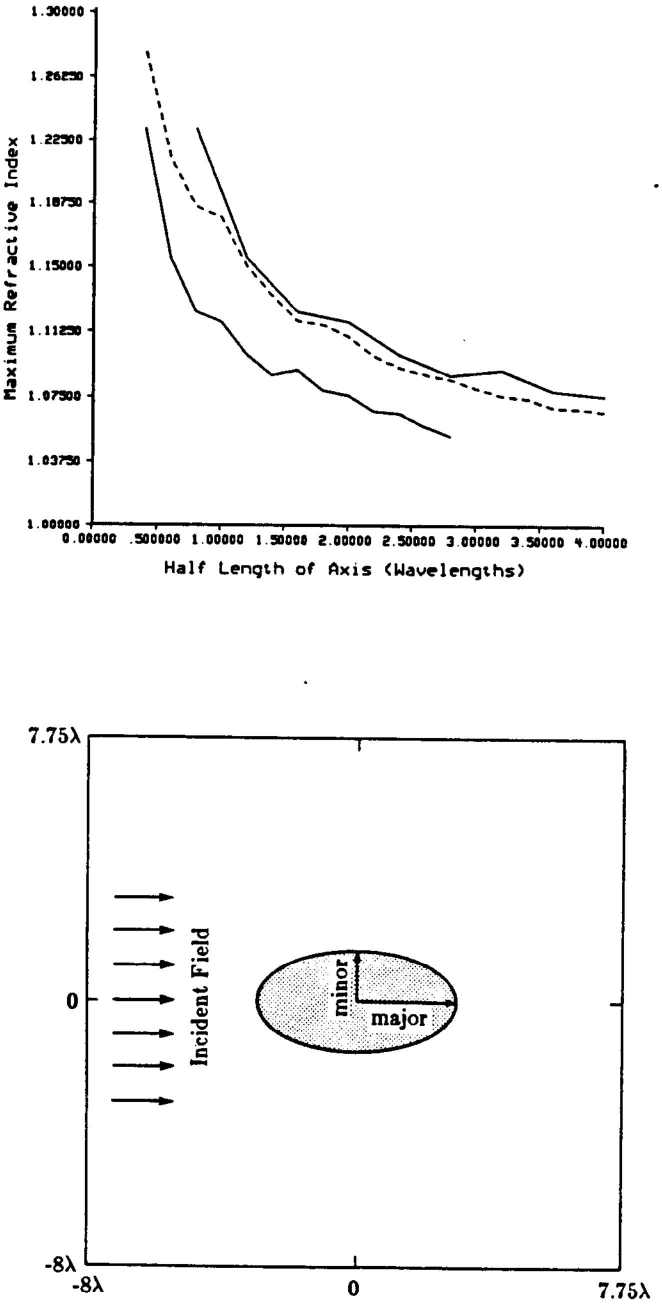

Figure 6.10 Born series region of convergence for an ellipse. The solid lines show the convergence for the ellipse as a function of its major axis (upper line) and minor axis (lower line). The dashed line shows the convergence for a single cylinder.

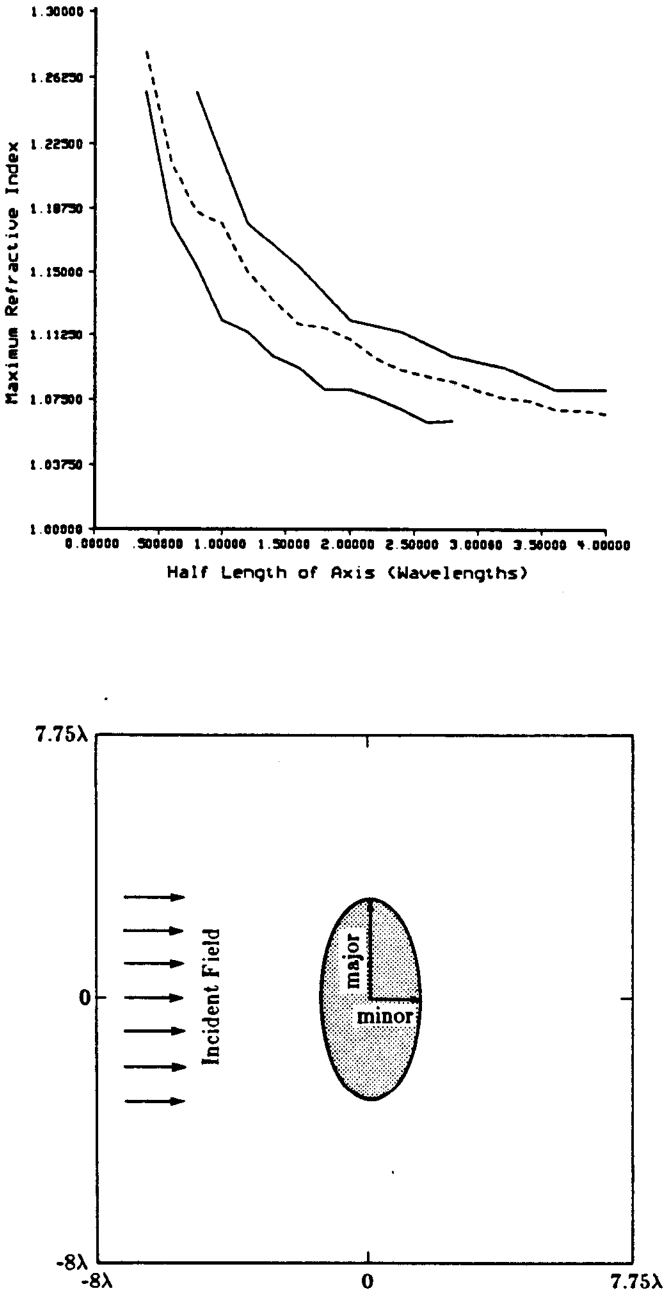

Figure 6.11 Born series region of convergence for the object shown in Figure 6.10 but rotated 90 degrees.

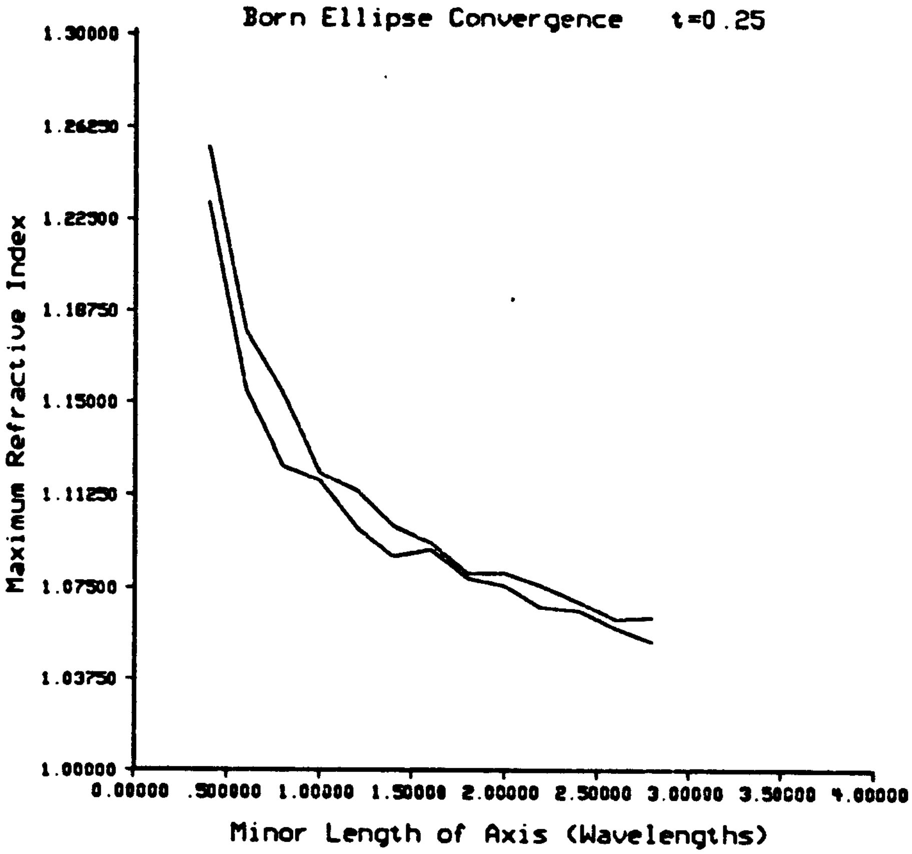

Figure 6.12 The region of convergence for two orientations of an ellipse are compared. The upper curve represents the convergence for the ellipse shown in Figure 6.11 while the lower curve represents the convergence shown in Figure 6.10.

6.3.2 Born Series with Attenuation

While to this point only fields in non-attenuating media have been discussed it is easy to also talk about the Born series when attenuation is present. Now the refractive index and the wavenumber are no longer just real valued and include an imaginary component to represent the attenuation.

The relationship between the real and imaginary components of the wavenumber are easily seen by examining one solution to the homogeneous wave equation

$$u_0(\vec{r}) = e^{jk_0x} \tag{6.40}$$

where $\vec{r} = (x,y)$. If k_0 is complex and equal to

$$k_0 = k_r + jk_i \tag{6.41}$$

then the real component will continue to represent the periodic component of the field. The imaginary component, k_i, contributes a multiplicative term, e^{-k_ix}, that causes attenuation of the plane wave with increasing distance.

There are two approaches that can be used to deal with attenuation. In the simpler approach the average wavenumber, k_0, is real and all of the attenuation is a perturbation from the average wavenumber. Thus the object function is complex and as will be shown shortly the region of convergence is reduced for large attenuations. In the second approach the attenuation is included in the average wavenumber thus reducing the magnitude of the object function. The only difference in the formulation is that the Green's function changes but now increasing attenuation leads to a larger region of convergence.

When the attenuation of the object is treated as just a perturbation of the object function from the real valued wavenumber then the effect is to reduce the region of convergence. This is shown in Figure 6.13 where it can be seen that the magnitude of the object function, not just the refractive index, determines the convergence of the Born.

The work first done with the Born approximation assumed that the average refractive index was real valued only. Since the real part of the wavenumber represents the speed of the wave and the imaginary part its attenuation, any attenuation in the object is included in the unknown perturbation. The magnitude of the perturbation determines the applicability of the Born and Rytov series therefore a more accurate estimate of the average refractive index will lead to smaller perturbations and better results with first order diffraction tomography algorithms.

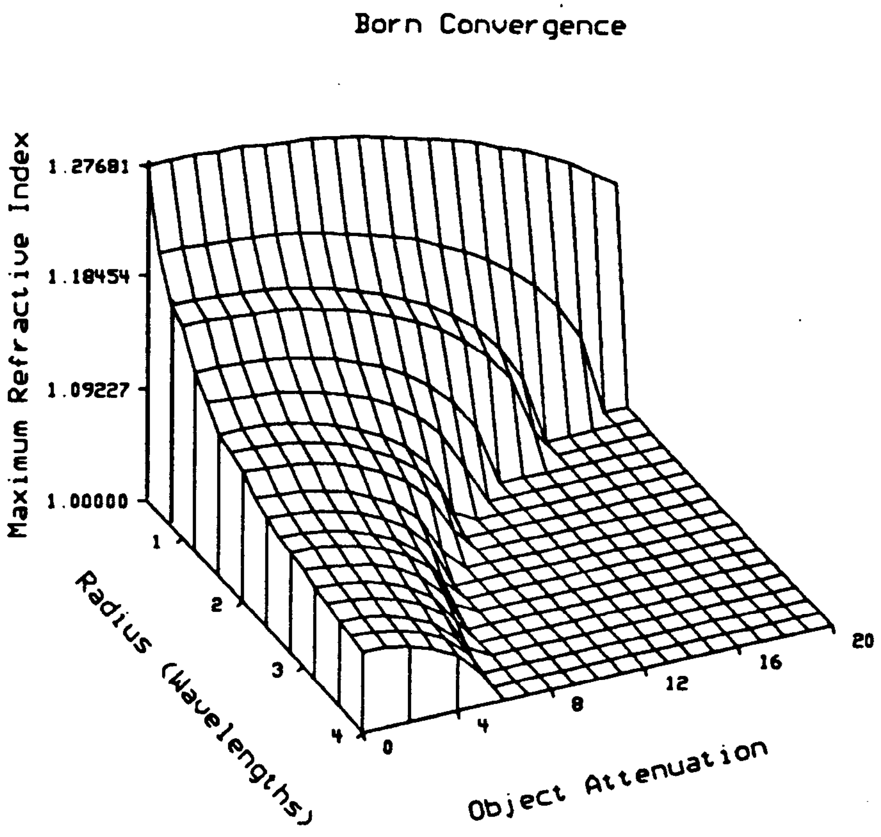

Figure 6.13 The region of convergence for the Born series as a function of cylinder radius and attenuation. The attenuation is plotted in units of nepers (1 neper represents an attenuation of the fields amplitude by 63% per wavelength.).

Consider an experiment where an object and the surrounding media are both highly attenuating. This might be typical of a microwave tomography experiment where the attenuation of microwaves is predominately due to the water molecule.

In this case a small perturbation model would be more accurate if the average value of the wavenumber, k_0, is assumed to include an imaginary component. Thus the real part continues to represent the spatial frequency of the wave while the imaginary part indicates the bulk attenuation of the wave as it travels through the media.

An important part of this derivation is to remember that a field can be described in two different manners; if you like, there are two sets of basis functions that can be used. A field that satisfies the Helmholtz equation is usually described in terms of plane waves. A plane wave is an exponential solution to the Helmholtz equation and for a plane field described by

$$u(\vec{r}) = e^{j\vec{k}\cdot\vec{r}} \tag{6.42}$$

where

$$\vec{k} = (k_x, k_y) \tag{6.43}$$

then valid plane waves satisfy

$$k_x^2 + k_y^2 = k_0^2. \tag{6.44}$$

k_0 in this equation represents the wavenumber of the media.

A problem with this approach is that both k_x and k_y can be complex. The complex valued wavevectors lead to evanescent waves which attenuate with distance. While for most applications the evanescent fields can be ignored (they tend to be small compared to the non attenuating components) the same assumption can not be made when calculating the field inside the object or when the object is in an attenuating media.

Consider an attenuating plane wave propagating in the x direction

$$u(\vec{r}) = e^{j\vec{k}\cdot\vec{r}} = e^{jkx} = e^{j(\alpha + j\beta)x} \qquad \beta \geq 0 \tag{6.45}$$

In this case α represents the phase term, while β represents the attenuation. For positive x this is an attenuating plane wave, but for negative x the wave grows exponentially. Thus it is normally necessary to specify that the wave is zero for $x < 0$.

Due to the non symmetry of attenuating plane waves and the efficiency of Fast Fourier Transform algorithm a much more natural set of basis functions is provided by the Fourier domain. In this approach the field is represented as a sum of Fourier components. If the field does not have any evanescent components the Fourier and the plane wave representations are identical.

The distinction becomes important when attenuating fields are considered. Mathematically the two approaches are equally valid but while an attenuating plane wave has a single component in the plane wave representation it has an infinite number of Fourier components. The difference is further illuminated if an attenuating plane wave is propagating through a homogeneous media. This field is represented as a single plane wave that satisfies the wave equation but its Fourier representation has an infinite number of components. While each of its Fourier components do not satisfy the wave equation they do represent a set of basis functions for describing linear operators. For this reason the Fourier approach is optimum for propagation problems [Goo68] and convolution integrals.

Again, like was done for the non attenuating Born, Figure 6.14 shows the components of the Born integral for a complex wavenumber. The major differences are that the incident field becomes an attenuating plane wave and the attenuating Green's function is more straightforward to calculate since it no longer has a singularity in the frequency domain.

The incident field

$$u(\vec{r}) = e^{j(\alpha+j\beta)x} \qquad x>0 \tag{6.46}$$

is a complex (2 dimensional) function and its Fourier transform can be found by considering it as a multiplication of a complex sinusoid by a one sided exponential. The following one dimensional Fourier pairs are used:

$$e^{j\omega_0 t} \quad \leftrightarrow \quad 2\pi\delta(\omega-\omega_0) \tag{6.47}$$

and

$$e^{-\alpha t} \quad t\geq 0 \quad \leftrightarrow \quad \frac{1}{j\omega+\alpha}. \tag{6.48}$$

Multiplication in the space domain corresponds to convolution in the frequency domain so the Fourier transform of an attenuating sinusoid is written

$$\frac{2\pi}{j(\omega-\omega_0)+\alpha}. \tag{6.49}$$

In two dimensions the Fourier transform of the incident field is written

$$U_0(\vec{K}) = \frac{\delta(k_y)2\pi}{j(k_x-\alpha)+\alpha} \tag{6.50}$$

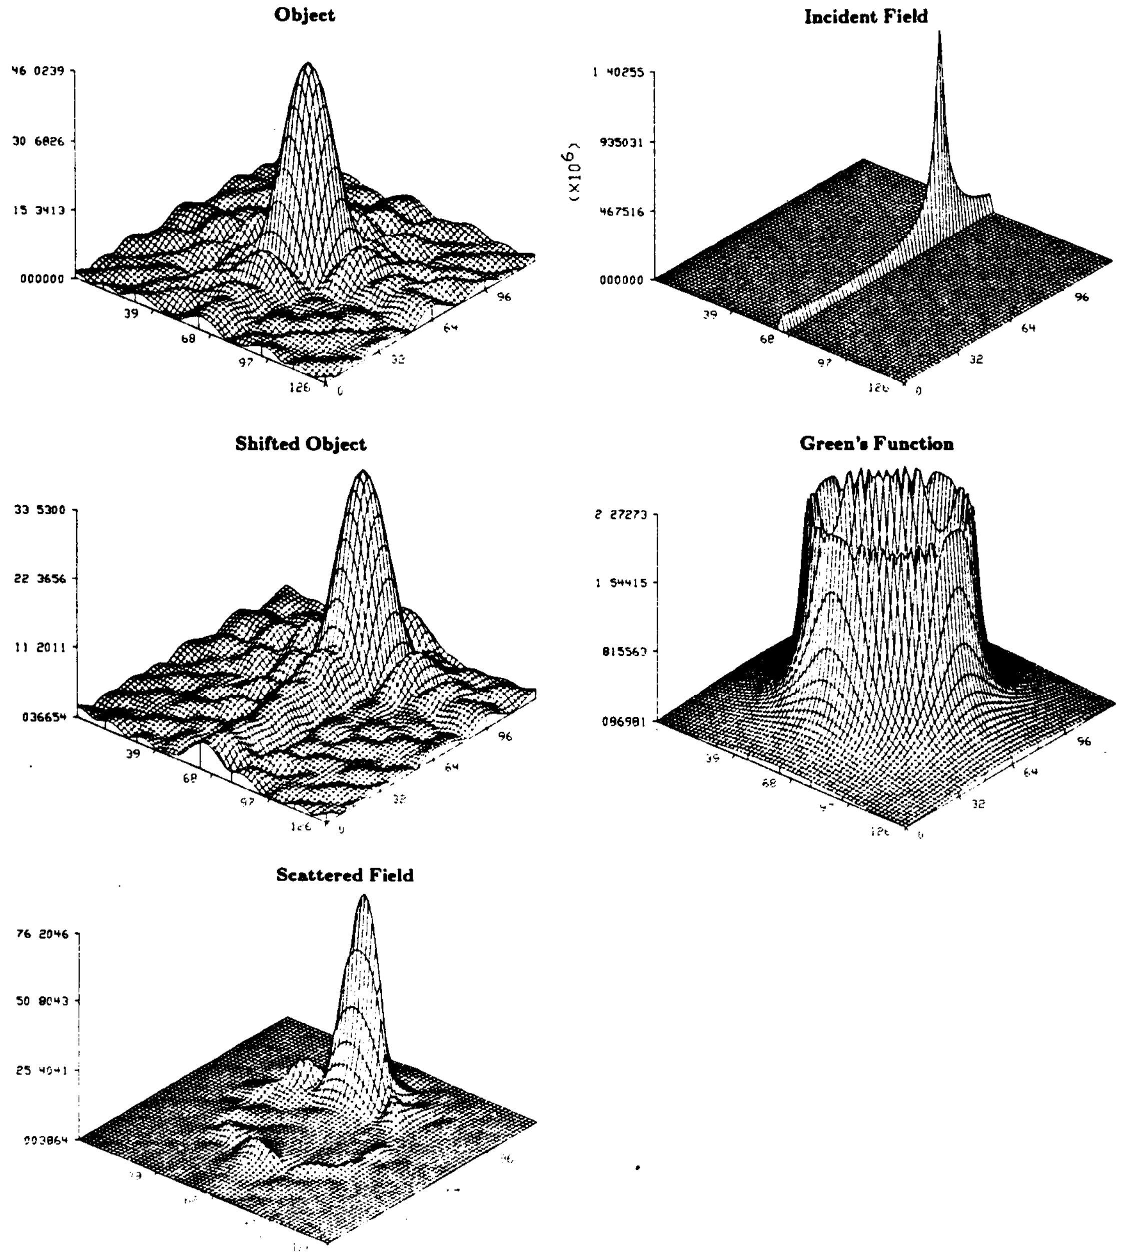

Object

Incident Field

Shifted Object

Green's Function

Scattered Field

Figure 6.14 The Fourier transform of several steps in the derivation of the scattered field in an attenuating media.

where the wavevector of the incident field is given by $\vec{\Lambda} = (\alpha,\beta)$. This is shown in Figure 6.14(Incident Field).

The Green's function for complex values of k_0 is even simpler than the non-attenuating case. By taking the Fourier transform of the wave equation with a delta forcing function it is easy to see that

$$G(\vec{\Lambda}) = \frac{1}{|\vec{\Lambda}|^2 - k_0^2}. \tag{6.51}$$

Since $|\vec{\Lambda}|^2$ is real and k_0^2 is complex there are no singularities in this function. The Fourier transform of the Green's function is still circularly symmetric and is shown in Figure 6.14(Green's Function).

The Fourier transform of the two dimensional scattered field can be written now as

$$U_s(\vec{K}) = G(\vec{K}) \, [\, U_0(\vec{K}) * O(\vec{K}) \,]. \tag{6.52}$$

The convolution of the incident field and the object is a shifted and a smeared version of the Fourier transform of the object. The convolution can be considered in two parts. The smearing caused by the width of the incident field in the spatial frequency domain can be ignored since it just redistributes some of the energy of the object function's Fourier transform. The remaining component, the shift in the frequency domain, is identical to the non attenuating Born case.

As derived above, the Green's function for an attenuating media does not have a ring of singularities and therefore samples a semi-circular region of the modified object function. This is shown in Figure 6.14(Scattered Field).

This procedure naturally leads to the Born series for objects with bulk attenuation. Figure 6.15 shows a composite graph of the region of convergence for a number of different attenuating media between 0 and 1 nepers per wavelength. It is interesting to note that the region of convergence gets larger as the attenuation increases. This is due to the reduction in multiple scattering because to the attenuating term in each wave. Thus for an average attenuation of 1 neper per wavelength (the amplitude of the field drops by 3dB per wavelength) the Born series converges for **all** objects with a refractive index less than 20%. In addition the convergence of the Born is less sensitive to the size of the object since the waves are attenuated before they travel the complete distance of the object.

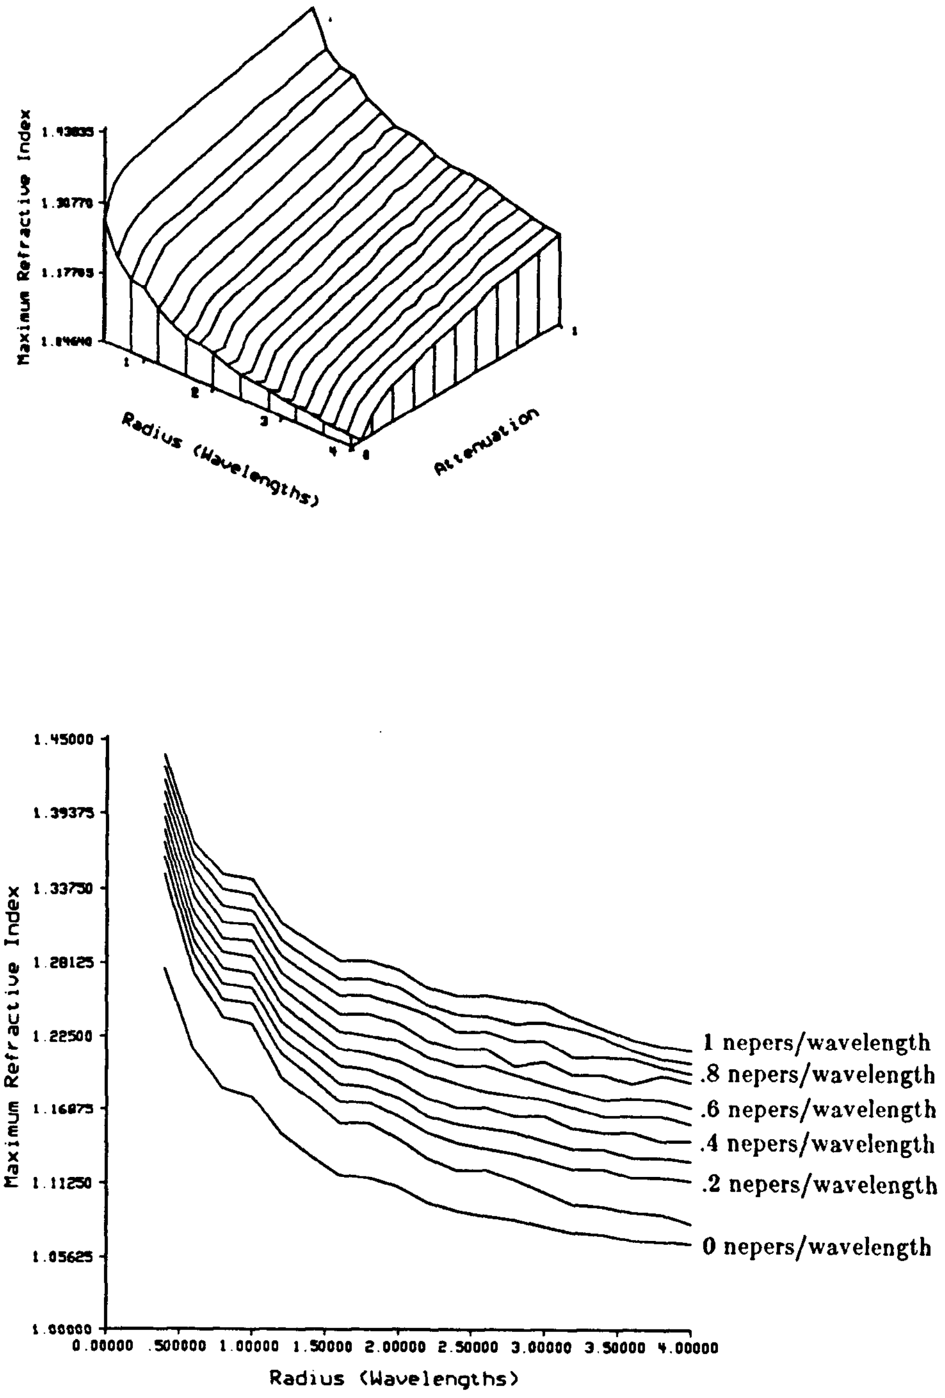

Figure 6.15 The region of convergence for the Born series as a function of cylinder radius and average attenuation of the media.

6.3.3 Rytov Series

The implementation of the Rytov series is much like that of the Born. Thus from Chapter 2 the Rytov integral for the scattering phase is written

$$\psi_1(\vec{r}) = \frac{1}{u_0(\vec{r})} \int_v g(\vec{r}-\vec{r}\,') \left[\nabla\psi_1(\vec{r}\,')\cdot\nabla\psi_1(\vec{r}\,') + o(\vec{r}\,')\right] u_o(\vec{r}\,')\, d\vec{r}\,'. \qquad (6.53)$$

A fixed point solution to this equation is possible if a guess for $\psi_1^{(0)}$ is used in the right side of this equation and a new value for the scattered phase, $\psi_1^{(1)}$, is computed. If the kernel of equation (6.53) is a contracting operator then $\psi_1^{(1)}$ will be a better estimate of the scattered phase. This iteration step can be carried out as often as desired until the change in the scattered phase is small.

The computer implementation of the Rytov series is more difficult than the Born due to the derivatives inside the integral. In the Born approximation it is possible to decompose the series and think of the entire iteration as modeling higher order scattering. On the other hand in the Rytov integral the $[\nabla\psi_1(\vec{r}\,')\cdot\nabla\psi_1(\vec{r}\,')]$ term is a non-linear function of the scattered phase and thus it is not possible to consider each term separately.

The computer implementation of the Rytov series is made especially difficult since the scattered potential does not have finite support as it does in the Born series. Recall, in the Born series the scattering potential is a function of the product of the scattered field and the object function and the convolution integral need be evaluated only where the object function is non-zero. Thus carrying out the Born integration over a finite region does not introduce any errors.

This is not the case with the Rytov integral since the scattering potential is now given by the expression

$$[\nabla\psi_1(\vec{r}\,')\cdot\nabla\psi_1(\vec{r}\,')]u_0(\vec{r}). \qquad (6.54)$$

In general, the scattered phase, ψ_1, is not equal to zero so limiting the integration to a finite region will always introduce errors.

Since the derivative is a linear operator it can be implemented as a convolution integral. Unfortunately, as will be shown shortly, the structure of the problem is such that the most accurate method, based on FFT's, is not workable.

When using an FFT to implement a convolution integral it is necessary to zero pad the original data so that the FFT represents an aperiodic convolution. While this technique works very well for most signals it has disastrous consequences when calculating the derivative of the field. Since the field never goes completely to zero there is always a sharp transition between the field and the start of the zero padding. This transition leads to a large value of the derivative at this point and eventually to large errors in the scattered field.

The standard procedure for dealing with this problem is to use a window to smooth out the transition. This solution is not viable here since the problem is to calculate the scattered phase at the outer edge of the grid, exactly where the effect of the window is greatest. Thus even though the kernel for the derivative operator is very compact the long tails lead to errors with an FFT based implementation.

A better solution is to approximate the derivative operator with a two point kernel and make appropriate adjustments at the edges of the grid. A third order polynomial is fit to the three points $(-t,y_{-1})$, $(0,y_0)$ and (t,y_1) with the function [Sto80]

$$f(x) = \frac{y_{-1}x(x-t)}{2t^2} + \frac{y_0(x-t)(x+t)}{-t^2} + \frac{y_1(x+t)x}{2t^2}. \tag{6.55}$$

The first derivative of this polynomial is found to be

$$f'(x)\Big|_{x=0} = \frac{y_1-y_{-1}}{2t}. \tag{6.56}$$

At the edge of the grid all three values of the field are not available and so a second order polynomial is fit to the two points and the derivative becomes

$$f'(x)\Big|_{x=0} = \frac{y_0-y_{-1}}{t} \tag{6.57}$$

or

$$f'(x)\Big|_{x=0} = \frac{y_1-y_0}{t}. \tag{6.58}$$

These operators are straightforward and allow the ∇^2 operator to be computed quickly in the time domain.

The convergence of the Rytov series is much like that of the Born. That the Rytov series converges to the correct answer is shown in Figures 6.16 and 6.17. Figure 6.16 shows the exact scattered field from a 2λ cylinder with a refractive index of 1.13. The field was measured at a receiver line 7.75λ from the center of the cylinder and sampled every $1/4\lambda$. Finally Figure 6.17 shows

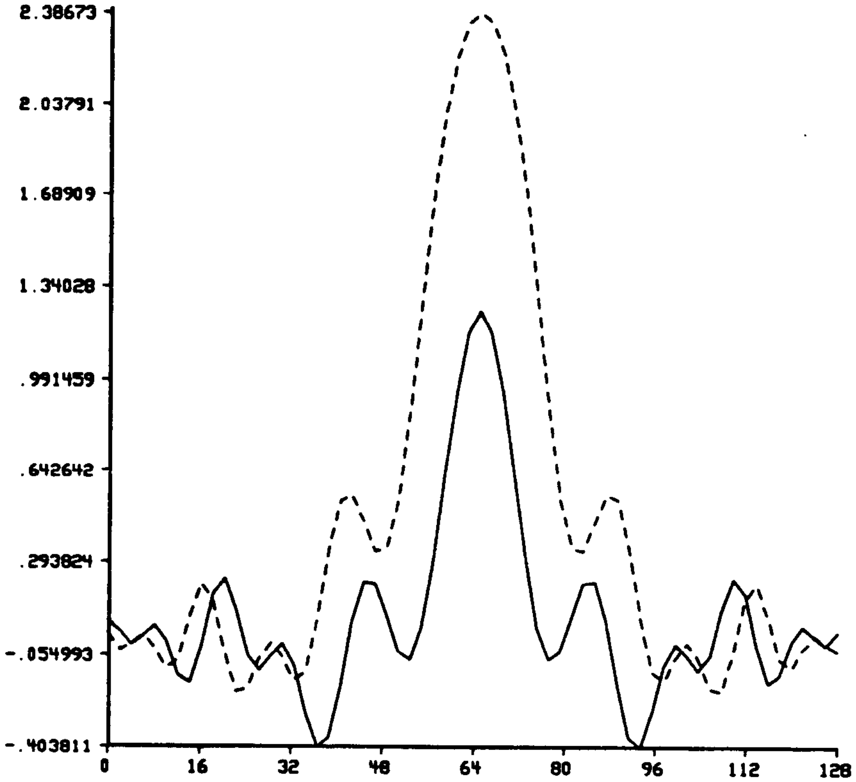

Figure 6.16 The exact scattered field from a 2λ cylinder with a refractive index of 1.13 is shown here. The real part of the field is shown as a solid line while the imaginary component is represented as a dashed line.

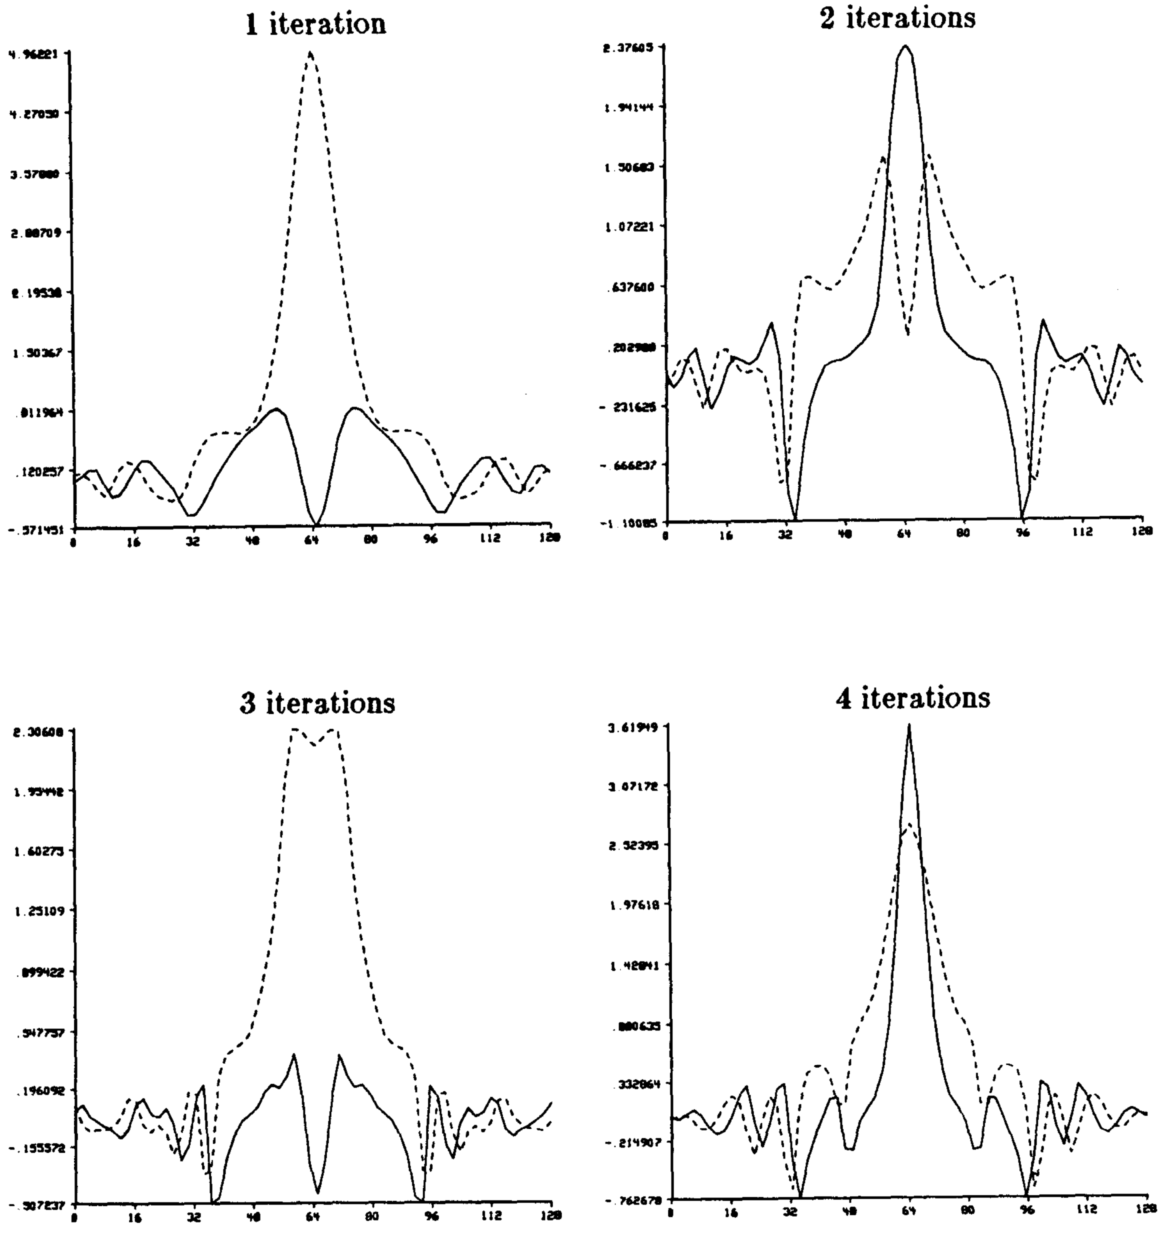

1 iteration

2 iterations

3 iterations

4 iterations

Figure 6.17 Twenty iterations of the Rytov series are shown demonstrating the convergence of the Rytov series to the field shown in Figure 6.16.

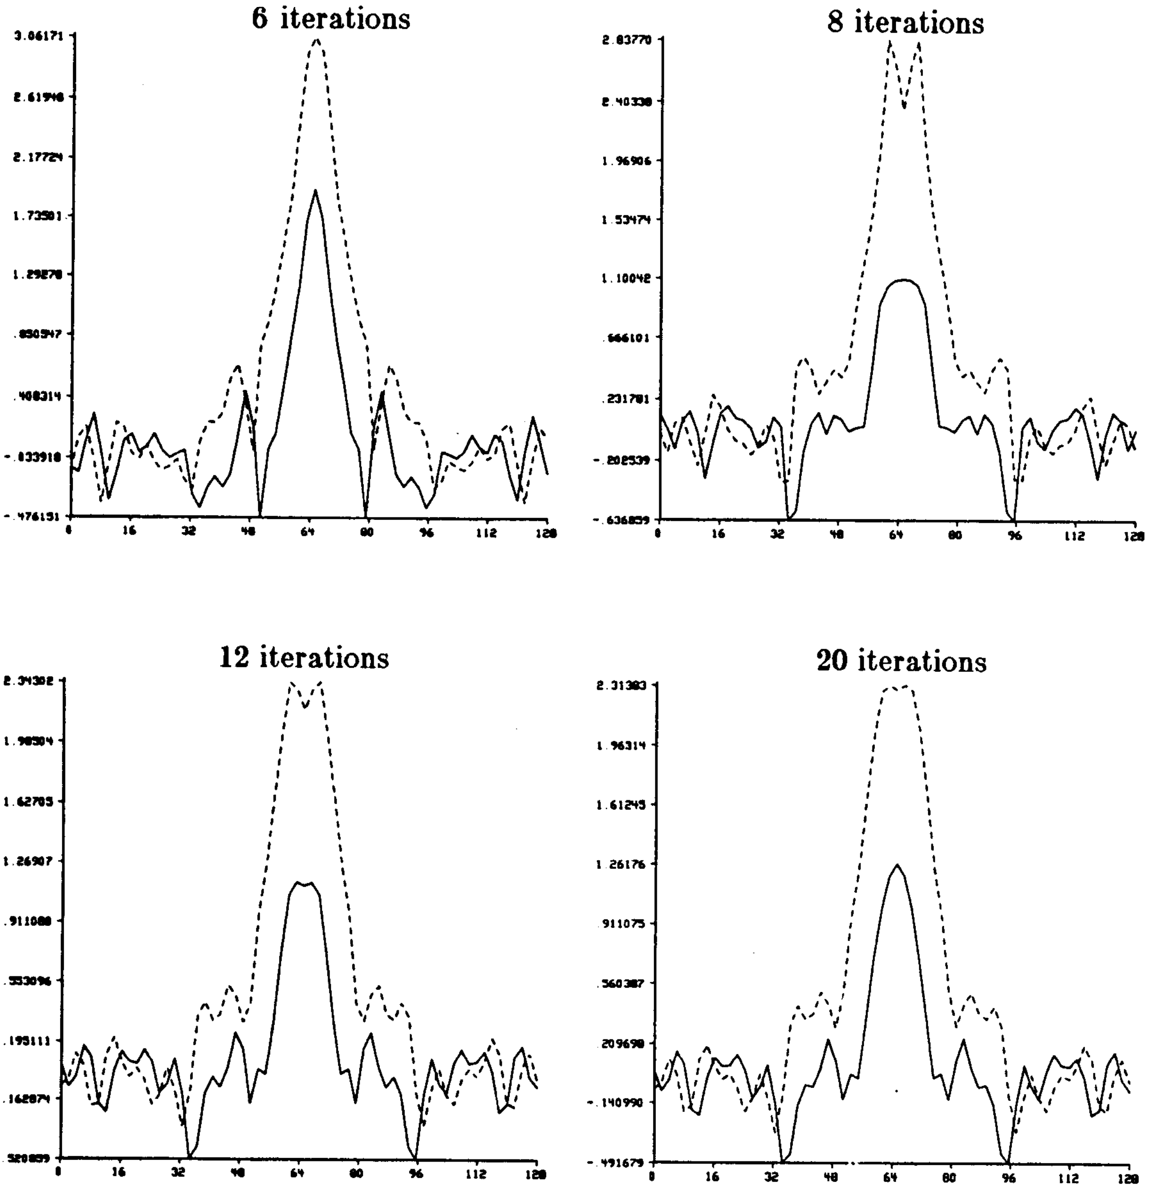

Figure 6.17 Continued.

that the Rytov series does converge to the exact solution for the scattered field.

The region of convergence of the Rytov series is compared to that of the Born in Figure 6.18. A number of works [Kel69, San70, Man70 and Wes84] have discussed this issue but it is not clear from the theoretical work which series is superior. These numerical implementations of the Born and Rytov series show the Rytov to converge for a large range of objects than the Born does. This is especially surprising considering the different domains of validity of the Born and Rytov approximations.

The difference between the Born and Rytov series is highlighted in Figures 6.19 and 6.20. These two figures show the convergence of the Rytov series for an object made up of two cylinders and an elliptical object. While in the Born approximation the orientation of the object changed the convergence of the series the same is not true for the Rytov series. The convergence results for the Rytov series are identical for either orientation (0 and 90 degrees).

The behavior of the Rytov series with an attenuating object is shown in Figure 6.21. Like the Born series the Rytov series is relatively insensitive to attenuation in the object until the attenuation becomes large enough. The attenuation at which the convergence of the Rytov series falls to zero is dependent on the radius of the cylinder.

If the object and the media have an average attenuation then the Rytov series will converge more easily. This is shown in Figure 6.22. Using an attenuating Green's function reduces the field at distances far from the object and thus makes it easier for the Rytov series to converge.

6.4 Matrix Formulation

An alternative to the fixed point methods like the Born and the Rytov was shown by Kaczmarz [Kac37], applied to the forward scattering problem by Richmond [Ric65] and extended to inverse scattering by Johnson [Joh83, Tra83 and Joh84]. The Kaczmarz algorithm has found widespread use in tomographic imaging based on ray tracing. Its use for this type of problem and a discussion of several possible optimizations and tricks is discussed in [Her73, Her76 and Her80]. While the Born and the Rytov series use discrete math to implement a continuous solution to the Helmholtz equation, a different approach is possible if the field and the object are first discretized. The Helmholtz equation now becomes a matrix equation and with appropriate manipulations can be put in the form

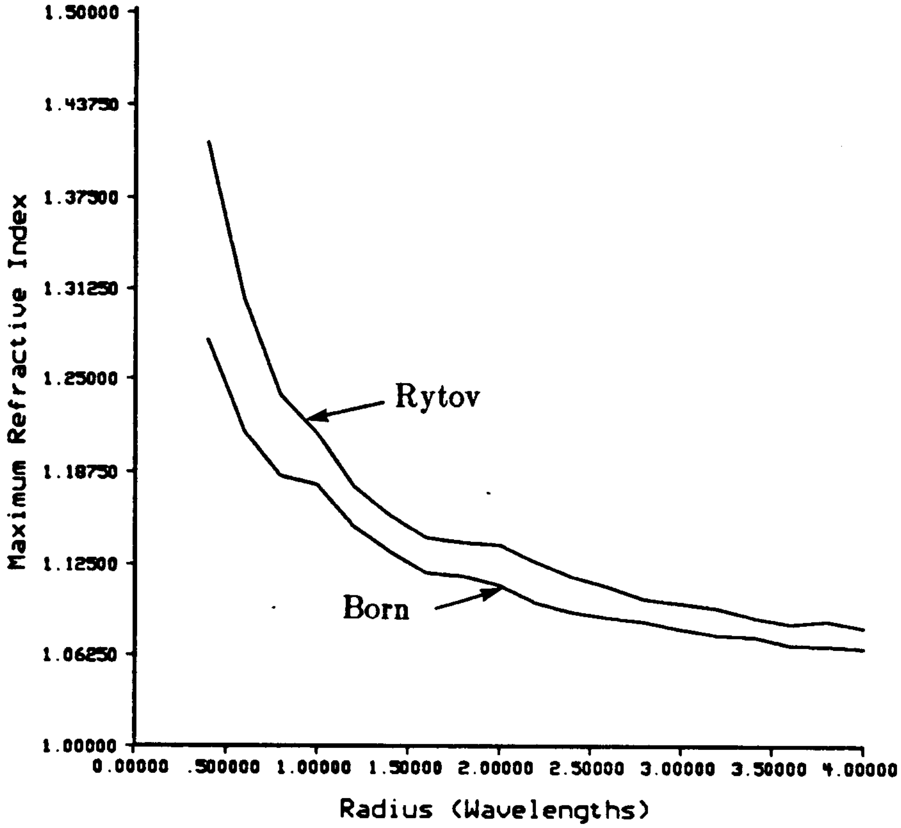

Figure 6.18 The convergence of the Born and the Rytov series are compared. The upper line represents the Rytov series.

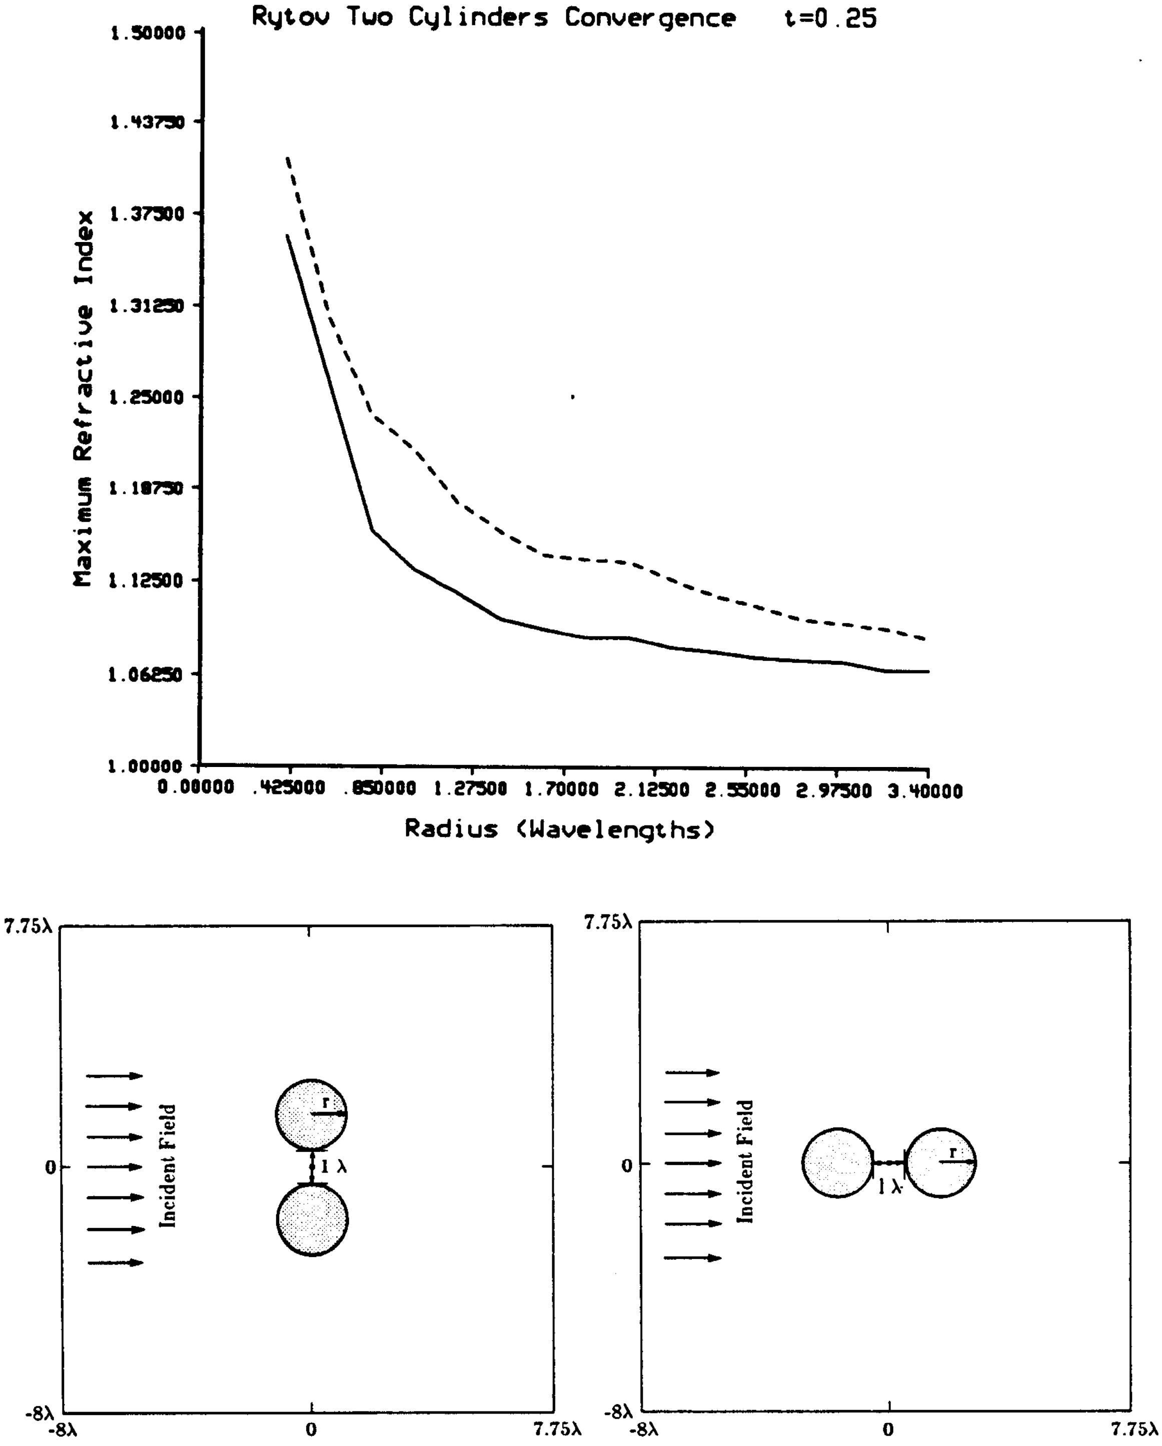

Figure 6.19 The region of convergence of the Rytov series for two cylinders. The experiment here is identical to that shown in Figure 6.9 but the results are independent of orientation.

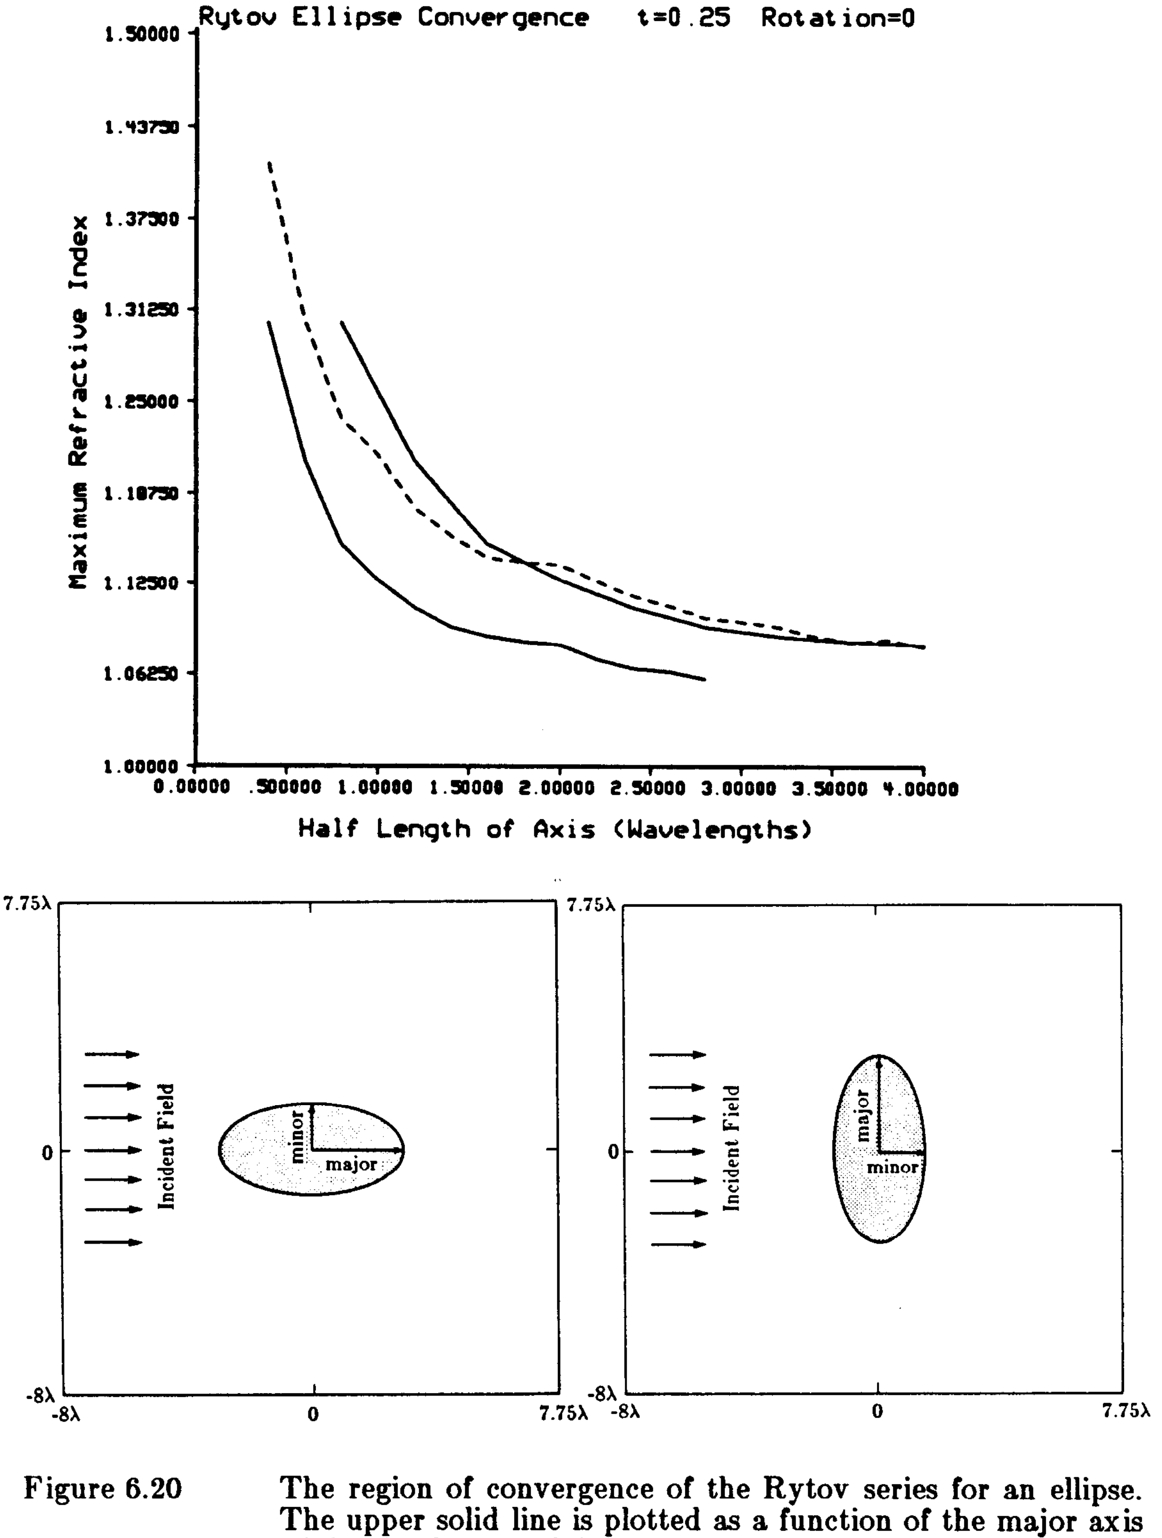

Figure 6.20 The region of convergence of the Rytov series for an ellipse. The upper solid line is plotted as a function of the major axis while the lower solid line is a function of the minor axis. The solid line represents the convergence for a single cylinder.

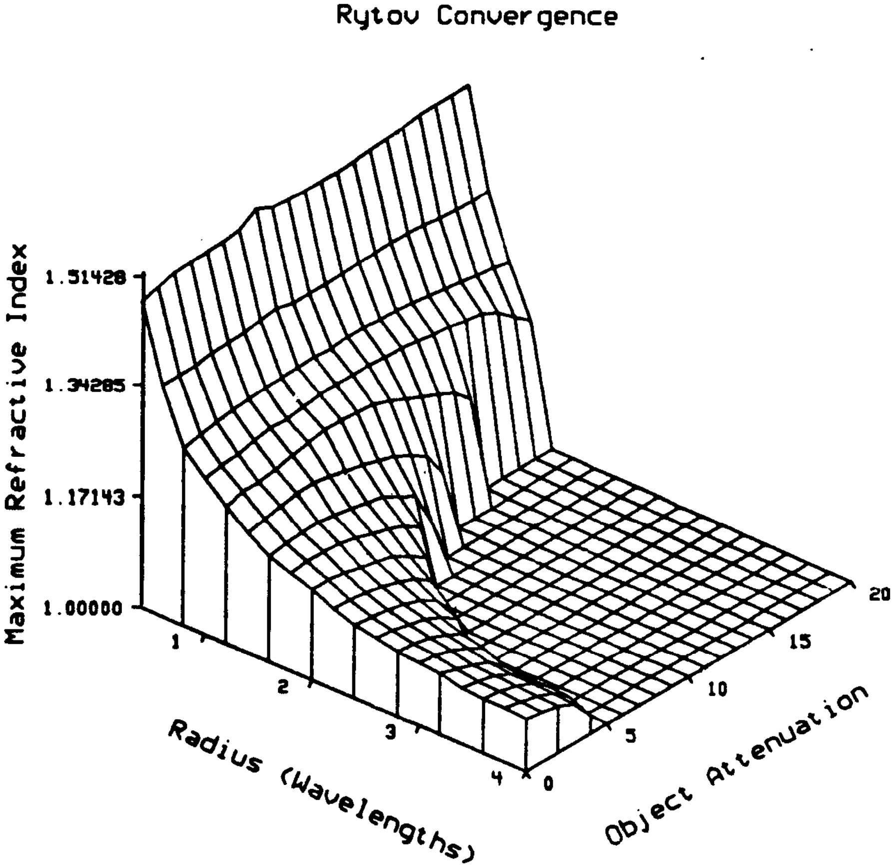

Figure 6.21 The convergence of the Rytov series is shown as a function of the cylinders radius and attenuation. The attenuation of the object is shown in nepers.

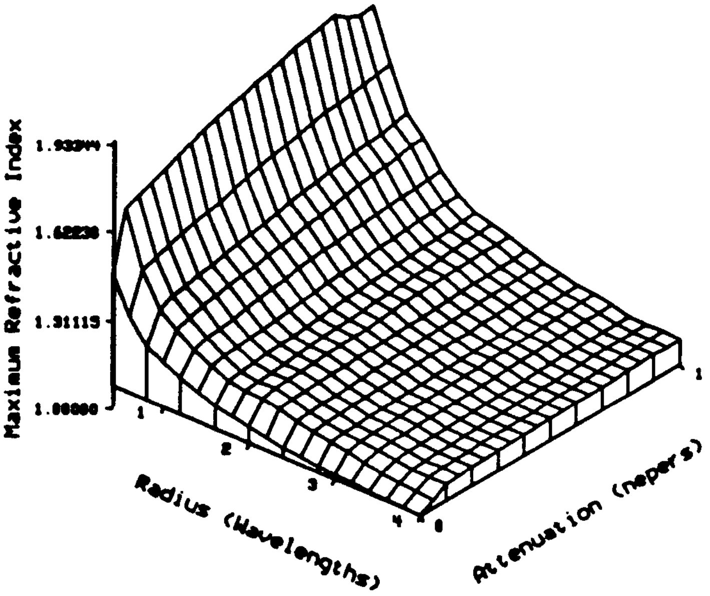

Figure 6.22 The convergence of the Rytov series is shown as a function of the cylinder radius and the average attenuation of the media. The attenuation is shown in nepers.

$$Ax = b \tag{6.59}$$

where x is the continuous field, A is a function of the object and the Green's function and b represents the effects of the incident field.

The Kaczmarz algorithm is used to solve this matrix equation because it operates on only one row of the matrix at a time. This property is important due to the large size of the matrix involved (as many as 4000 equations and unknowns.) The Kaczmarz algorithm belongs to a class of operators known as Row Action Methods and is described by Tanabe [Tan71] and by Censor [Cen81].

Most work with large matrices assumes that a significant fraction of the matrix is zero. Thus it is only necessary to store the location and the value of the non zero elements and then it is possible to use sparse matrix techniques to solve the equation (see, for example, Chapter 3 of [Hey85].) Since every element of the Green's matrix is non zero it is not possible to use these techniques and instead the Kaczmarz approach allows the calculation of a solution using modest amounts of computer memory.

As will be shown later the Kaczmarz algorithm always converges to a proper solution of the discrete equation Ax = b. If the discrete representation of the field, the object and the Green's function accurately model the true functions then the Kaczmarz solution will satisfy the Helmholtz equation.

An exact solution to the wave equation is given by the integral equation

$$u_s(\vec{r}) = \int o(\vec{r}\,') \, [u_0(\vec{r}\,') + u_s(\vec{r}\,')] \, g(\vec{r}-\vec{r}\,') d\vec{r}\,' . \tag{6.60}$$

By sampling faster than the Nyquist rate, each of the terms in the above equation can be discretized without errors as

$$u_s^{i,j} = u_s(iT, jT) \tag{6.61}$$

$$u_0^{i,j} = u_o(iT, jT) \tag{6.62}$$

$$o^{i,j} = o(iT, jT) \tag{6.63}$$

$$g^{i,j} = g(iT, jT). \tag{6.64}$$

The discrete version of the Helmholtz equation can now be written

$$u_s^{i,j} = \sum_{k=-\infty}^{\infty} \sum_{l=-\infty}^{\infty} o^{k,l} \, [\, u_0^{k,l} + u_s^{k,l} \,] \, g^{i-k,j-l} \tag{6.65}$$

Again, as in the implementation of the Born and Rytov series, assume that the object has finite support. Since the object multiplies each term in the summation the summation need be carried out only for those values of k and l where $o^{k,l}$ is non-zero. Thus without loss of generality the object will be assumed to exist only for $1 \leq i \leq N$ and $1 \leq j \leq N$ and all summations will be assumed to go from 1 to N.

The forward process will be described first. To do this rearrange the discrete version of the wave equation to find

$$\sum_k \sum_l o^{k,l} u_0^{k,l} g^{i-k,j-l} = u_s^{i,j} - \sum_k \sum_l o^{k,l} u_s^{k,l} g^{i-k,j-l}. \tag{6.66}$$

The $u_s^{i,j}$ term can then be moved into the summation to find

$$\sum_k \sum_l o^{k,l} u_0^{k,l} g^{i-k,j-l} = \sum_k \sum_l [\, \delta_{i,j}^{k,l} - o^{k,l} g^{i-k,j-l} \,] \, u_s^{k,l} \tag{6.67}$$

where the Kroenker delta, $\delta_{i,j}^{k,l}$, is defined as

$$\delta_{i,j}^{k,l} = \begin{cases} 1 & \text{for } i=k \text{ and } j=l \\ 0 & \text{elsewhere} \end{cases} \tag{6.68}$$

To put this in standard matrix form represent each field as a one-dimensional vector as follows

$$x_0 = \Big[x_0^{0,0}, x_0^{0,1}, \;\cdots, x_0^{0,N-1}, x_0^{1,0}, x_0^{1,1}, \;\cdots, x_0^{1,N-1}, \ldots, \\ x_0^{N-1,0}, \;\cdots, x_0^{N-1,N-1} \Big]^T \tag{6.69}$$

and

$$x = \Big[x^{0,0}, x^{0,1}, \;\cdots, x^{0,N-1}, x^{1,0}, x^{1,1}, \;\cdots, x^{1,N-1}, \ldots, \\ x^{N-1,0}, \;\cdots, x^{N-1,N-1} \Big]^T \tag{6.70}$$

where $[]^T$ represents vector transpose. An $N^2 x N^2$ matrix, A, can now be defined that represents the effects of all summations. Since the object and the Green's function are both known they can be combined into a single matrix and the discrete version of the wave equation can be written

$$Ax = b \tag{6.71}$$

where the terms of the A matrix are given by

$$A = \begin{bmatrix} 1-o^{0,0}g^{0,0} & -o^{0,1}g^{0,-1} & \cdots & -o^{N-1,N-1}g^{-N+1,-N+1} \\ -o^{0,0}g^{0,1} & 1-o^{0,1}g^{0,0} & \cdots & -o^{N-1,N-1}g^{-N+1,-N+2} \\ -o^{0,0}g^{0,2} & -o^{0,1}g^{0,1} & \cdots & -o^{N-1,N-1}g^{-N+1,-N+3} \\ -o^{0,0}g^{0,3} & -o^{0,1}g^{0,2} & \cdots & -o^{N-1,N-1}g^{-N+1,-N+4} \\ \cdots & \cdots & \cdots & \cdots \\ -o^{0,0}g^{N-1,N-2} & -o^{0,1}g^{N-1,N-3} & \cdots & -o^{N-1,N-1}g^{0,-1} \\ -o^{0,0}g^{N-1,N-1} & -o^{0,1}g^{N-1,N-2} & \cdots & 1-o^{N-1,N-1}g^{0,0} \end{bmatrix} \tag{6.72}$$

The field, x, is a one dimensional matrix with N^2 elements and the constant vector b is given by

$$b = Ax_0 \tag{6.73}$$

At first glance this equation (6.71) represents an especially simple form for the scattered field; that is until the size of the vectors are considered. For a small 64x64 reconstruction the b and x vectors have 4096 complex elements and the A matrix is a square 4096x4096 matrix with more than 16 million complex elements. Inverting a matrix of this size would require over 32 Megawords of memory or more than that which exists on all but a handful of processors today.

There are two tricks to solving this problem. The first of which is to realize that it is not necessary to find the inverse of the A matrix but only to find a vector field, x, which satisfies the discrete wave equation. Secondly by using a row action method such as that proposed by Kaczmarz and Johnson et al, it is no longer necessary to store the entire A matrix in memory. Thus it is possible to solve the system of equations storing only $4N^2$ equations at a time. For a 64x64 field this represents only 16000 elements so the storage requirements are reduced by a factor of 1000.

As described in [Ros82], the Kaczmarz method finds a vector x that satisfies the equation $Ax=b$ by considering each row of the matrix A to represent a separate equation. Thus the matrix

$$A = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \tag{6.74}$$

can be considered to be three equations of the form

$$\begin{aligned} b_1 &= a_{11}x_1 + a_{12}x_2 + a_{13}x_3 \\ b_2 &= a_{21}x_1 + a_{22}x_2 + a_{23}x_3 \\ b_3 &= a_{31}x_1 + a_{32}x_2 + a_{33}x_3 \end{aligned} \qquad (6.75)$$

In terms of an n-dimensional space each of these equations represents a single hyperplane and the intersection of all 'n' planes describes a single solution point x.

The Kaczmarz algorithm iteratively refines an initial guess by projecting the point onto each hyper-plane in sequence. The process of computing the projection onto the hyperplane also represents finding the point on the hyperplane closest to the original guess. As will be shown later this new x will always lie closer to the new solution vector than the original guess.

If the i'th row of the A-matrix is denoted as a_i and $< , >$ represents the dot product then a better solution to the equation $Ax^j=b$ is given by

$$x^{j+1} = x^j - \frac{<a_i,x^j> - b_i}{<a_i,a_i>} a_i. \qquad (6.76)$$

This operation is illustrated in Figure 6.23.

Simple geometrical arguments should convince the reader that this equation, (6.75) will always produce a better estimate of the solution vector x. For a two dimensional case this situation is illustrated in Figure 6.24. In this example a point, x_1, on the line CD will first be projected onto the line AB or a solution of the equation

$$a_{11}x_1 + a_{12}x_2 = b_1. \qquad (6.77)$$

For any point on the line CD, the point x_2 on the line AB is always a better estimate of the solution then the original point. Thus the Kaczmarz algorithm always converges. The distance between two points x_j and x_{j+1} will be defined in the normal Euclidean sense or

$$\text{Distance}^2 = <x^j-x^{j+1},x^j-x^{j+1}>. \qquad (6.78)$$

Since the solution vector, x, lies along the line AB an initial estimate x^j can always be improved by projecting the point onto the line. The point x^{j+1} is closer to the solution than all other points on the line CD so this always represents a better estimate. Since all points on the line CD project to the point x^{j+1} this procedure always reduces the error. Thus it is easy to see that the error is monotonicly decreasing and in addition will always converge to zero. (For now the cases where the system is either over determined or underdetermined are ignored. Both of these cases represent systems of

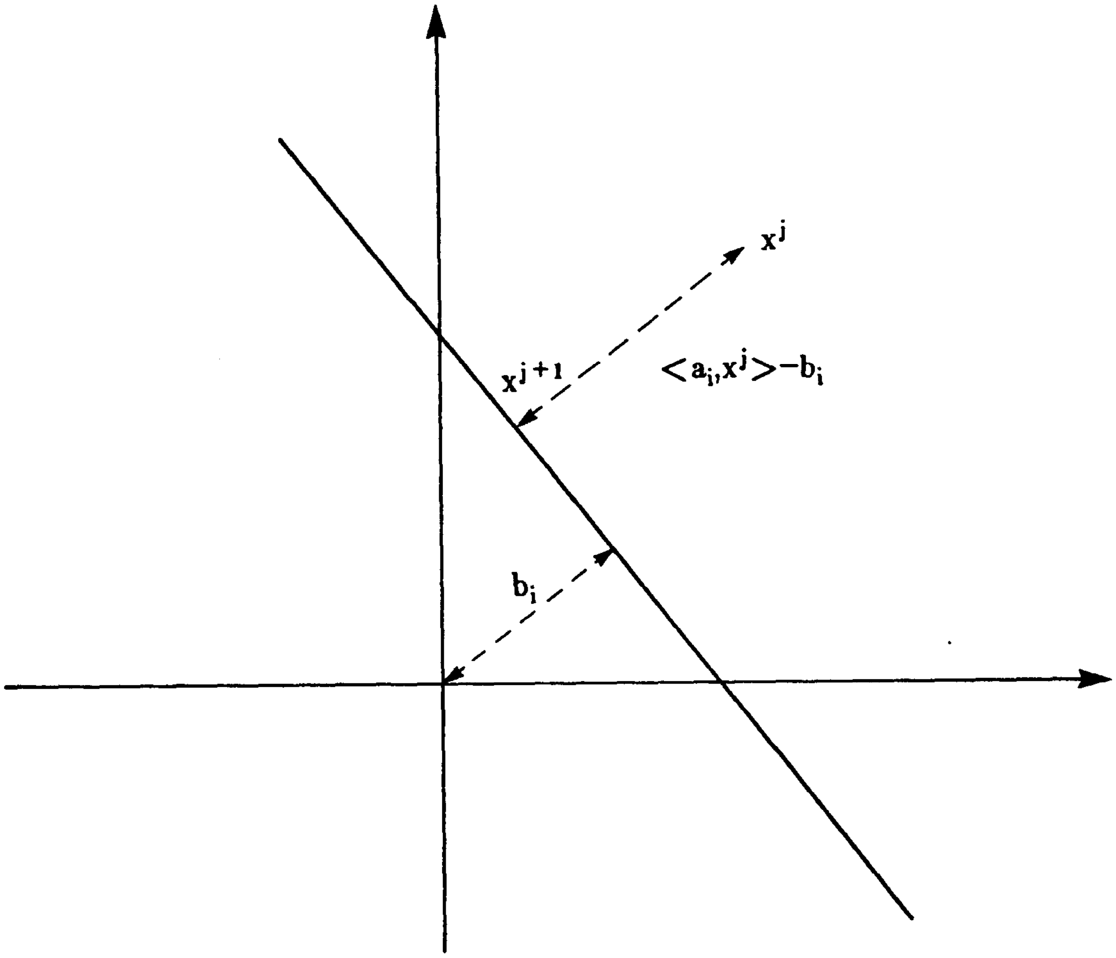

Figure 6.23 An initial estimate for the solution, x^j, is refined by finding the closest point on the line, x^{j+1}.

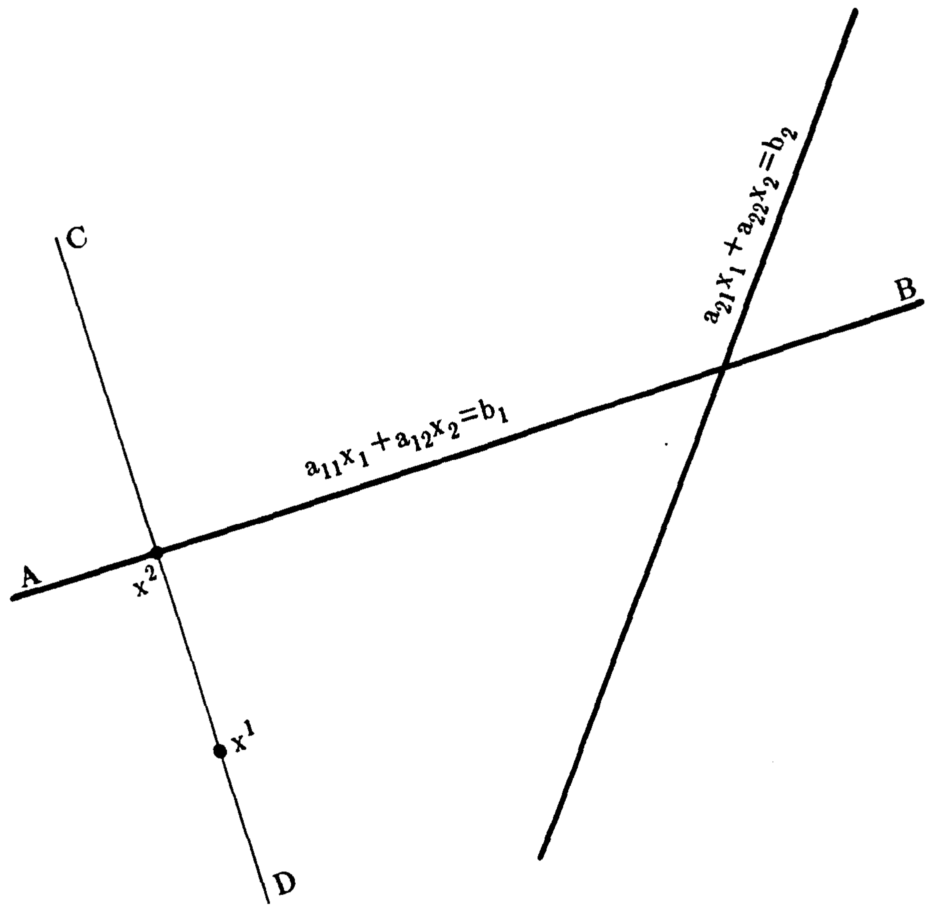

Figure 6.24 The point x^2 is a better estimate of the solution point then any other point on the line DC.

equations that don't necessarily have a unique solution.)

The speed of convergence is proportional to the independence of the rows of the A matrix. Figure 6.25 shows the convergence for two widely separated cases. In Figure 6.25a the two hyperplanes are perpendicular and the Kaczmarz algorithm converges to the solution in one iteration, no matter what the starting point is. (One iteration is defined as projecting the x vector onto each of the hyperplanes.) On the other hand in Figure 6.25b the hyperplanes are nearly parallel and it will take Kaczmarz algorithm many iterations to converge to the correct solution.

Since the Kaczmarz method only works with a single row of the A matrix at a time it is possible to make a space-time tradeoff. The form of the A matrix is simple enough that it is relatively inexpensive to recompute each row of the matrix as it is needed. While on many computer systems it is possible to precompute the A matrix and store it on disk, getting the data back off can be time consuming. For the computers accessible at Purdue (Floating Point Systems AP120B and Control Data Corporation Cyber 205) it is more cost effective to recompute the A matrix as needed. This approach has made the problem solvable for more than trivial sized matrices.

For an implementation of the forward scattering problem on a CDC Cyber 205 super computer calculating one row of the A matrix requires 98,000 floating point operations (64x64 grid). Since calculating the projection requires 57,000 floating point operations this implementation takes 2.7 times longer than the ideal (all values of the A matrix available immediately.) On the other hand retrieving the data off disk, either explicitly or using virtual memory, could take 100 times longer than the ideal situation*. Thus for this system of equations the tradeoff is easy to make. With a Cyber 205 one iteration of the Kaczmarz algorithm takes 233 million floating point operations and can be computed in 1 second of CPU time. This represents a real cost of 21¢ per iteration as charged at the standard rate by the Purdue University Computer Center.

To make the implementation of the Kaczmarz algorithm as fast as possible several quantities are precomputed and stored for quick access. The x and y coordinates of each grid point do not change during the course of the problem

---

* Theoretically it should be possible to organize the data on the disk in a sequential fashion and overlap execution time with IO time. With high enough bandwidth the total execution time should be nearly equal to the ideal situation. Unfortunately, most operating systems are not this cooperative.

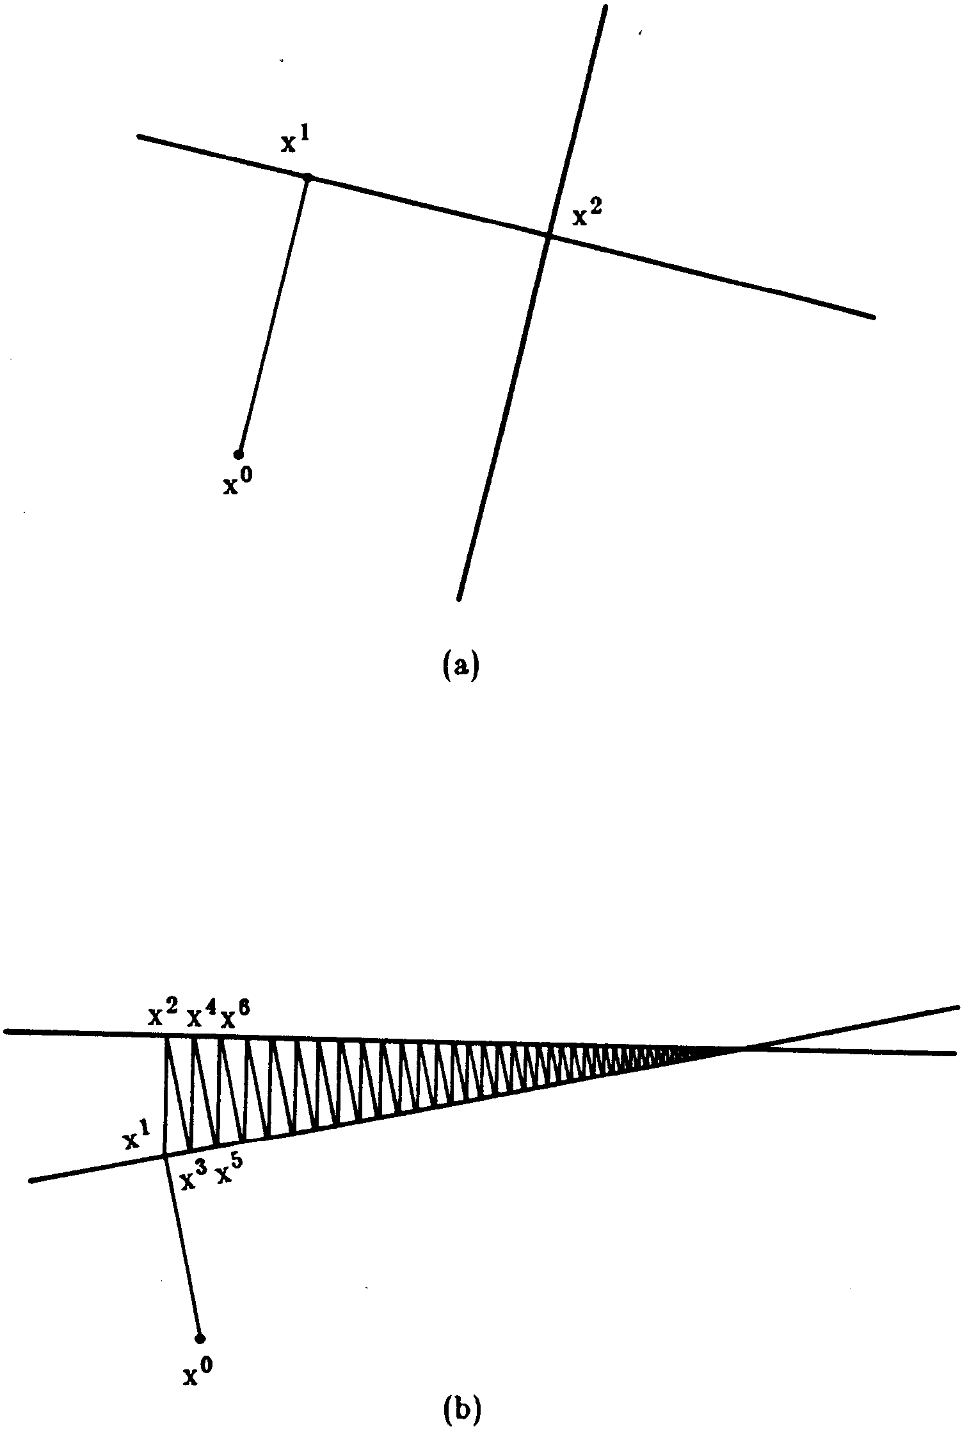

Figure 6.25 The orthogonality of the hyperplanes determines the rate of convergence. If the hyperplanes are perpendicular then the solution will be reached in only one iteration (a) while it will take much longer if the hyperplanes are nearly parallel (b).

and thus it is possible to store these in main memory and not recalculate them. In addition the Green's function is circularly symmetric and thus a function of only the radial distance. By precomputing the values of the Green's function along a single radial it is possible to use bilinear interpolation to quickly compute the Green's function for each grid point given the radial distance and the values along the radial.

The algorithm for one iteration of the Kaczmarz algorithm can be written as

```
For each equation (representing the scattered field at a single point)
    Compute the Green's Function
        •  Subtract a vector representing the x position of each grid point from the current equation.
        •  Subtract a vector representing the y position of each grid point from the current equation.
        •  Square each distance and add
        •  Find square root and multiply by scaling factor to find the argument of the Green's function.
        •  Use bilinear interpolation to find the Green's function at every point.
    Find the A matrix
        •  Multiply the Green's function and the object
        •  Subtract from the identity matrix $(\delta_{ij})$
    Project $x^j$ onto the hyperplane.
```

An additional complication in this approach is caused by the use of complex numbers. While the dot product operation is defined for complex vectors better results are obtained if each complex equation is considered to be two real valued equations. This simple change reduces the error by a factor of 100 or more but does increase the number of equations for a 64x64 image from 4096 to 8192. While the number of unknowns remains the same (4096 complex values or 8192 real values) there are now two projections that are needed for each row. Fortunately the number of operations remains the same.

The implementation of this algorithm was tested by comparing the results of the exact solution for a small cylinder. Using the incident field as the first "guess" for the solution several iterations were calculated and the real part of the solution is shown in Figure 6.26. This compares favorably with the exact

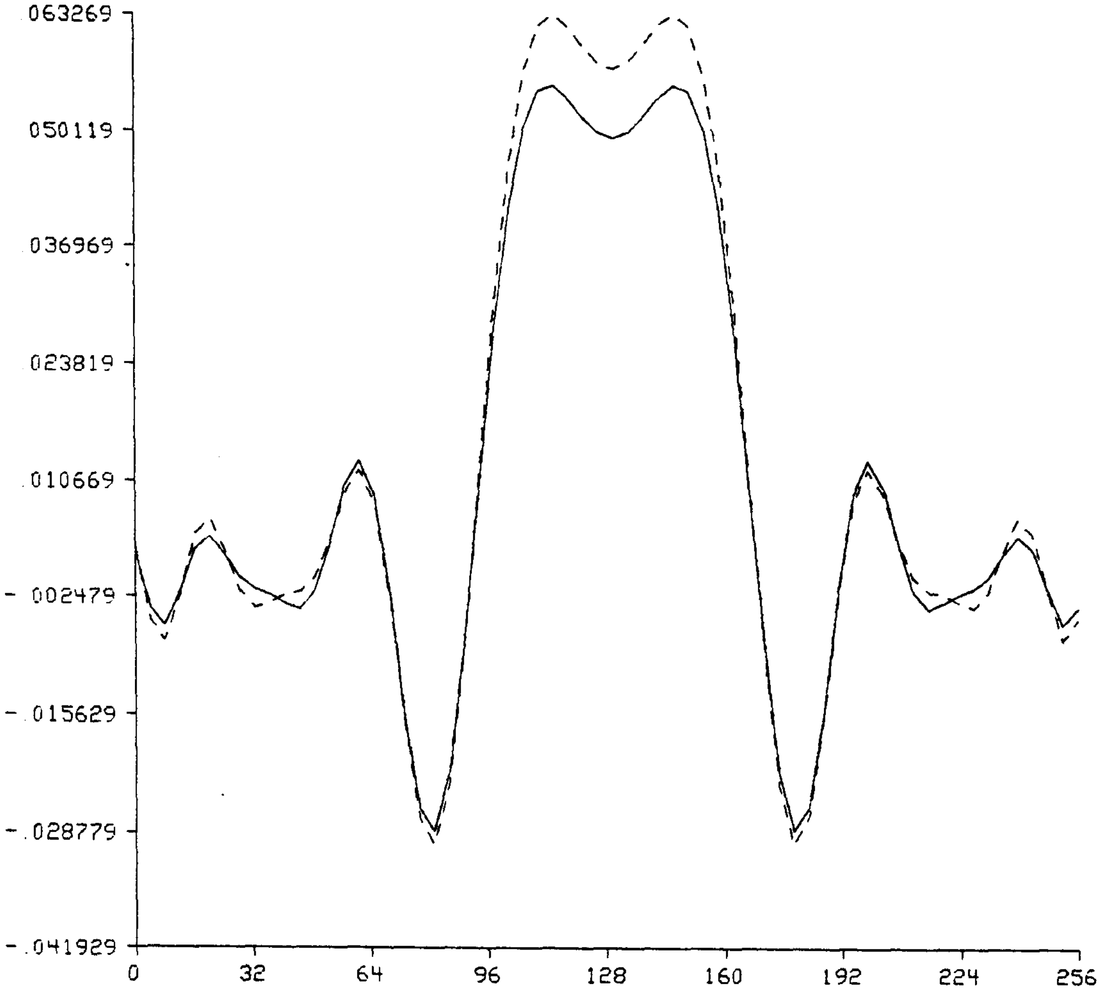

Figure 6.26 The real parts of the scattered field are compared as computed by the Kaczmarz approach (solid line) and the exact solution to the boundary conditions (dashed line).

solution shown as a dashed line. In this example and the work to follow one iteration is defined as projecting a vector, x^0, onto each of the N^2 hyperplanes.

The speed at which this algorithm converges to the correct answer is shown in Figure 6.27. Here the real part of the field for a section of space near the origin is shown before iterating (the incident field is the initial guess) and then after one, two and three iterations. This result is especially encouraging because the first iteration has changed so rapidly towards the correct answer.

Ramakrishnam [Ram79] proposed that faster convergence of a projection algorithm could be obtained by using pair wise orthogonalization. As already described, the Kaczmarz approach will converge in one step when all the equations are orthogonal and as the hyperplanes become more parallel it will take longer for the method to converge.

Certainly the best way to speed up convergence is to first orthogonalize the system of equations. Then it would be possible to solve the system of equations in a single iteration. Unfortunately the work required to orthogonalize the system is identical to that needed to find the inverse of the matrix. In addition it then would be necessary to compute and store all the elements of the matrix. Due to the large size of the matrix this is not practical.

Ramakrishnam proposed that pair wise orthogonalization be used to make each hyperplane perpendicular to the previous hyperplane. With this approach the equation for the field at each point is made orthogonal to the preceding one using the relation

$$\tilde{A}_i = A_i - \tilde{A}_{i-1}\left(\frac{<\tilde{A}_i, A_{i-1}>}{<\tilde{A}_{i-1}, \tilde{A}_{i-1}>}\right) \tag{6.79}$$

and

$$\tilde{b}_i = b_i - \tilde{b}_{i-1}\left(\frac{<\tilde{A}_i, A_{i-1}>}{<\tilde{A}_{i-1}, \tilde{A}_{i-1}>}\right) \tag{6.80}$$

Here the new orthogonalized system of equations are denoted with the matrix $\tilde{A}$ and the vector $\tilde{b}$, A_i is one row of the A matrix and b_i is the corresponding element of the b vector.

While this approach is not optimum it does have the advantage that at each step storage for only one extra row of the A matrix is needed. Ramakrishnam showed for a simple restoration problem that pair wise orthogonalization reduced the number of iteration needed to obtain a given

Real Part of Scattered Field
Radius = 3λ, Refractive Index = 1.01

Incident Field

1 Iteration

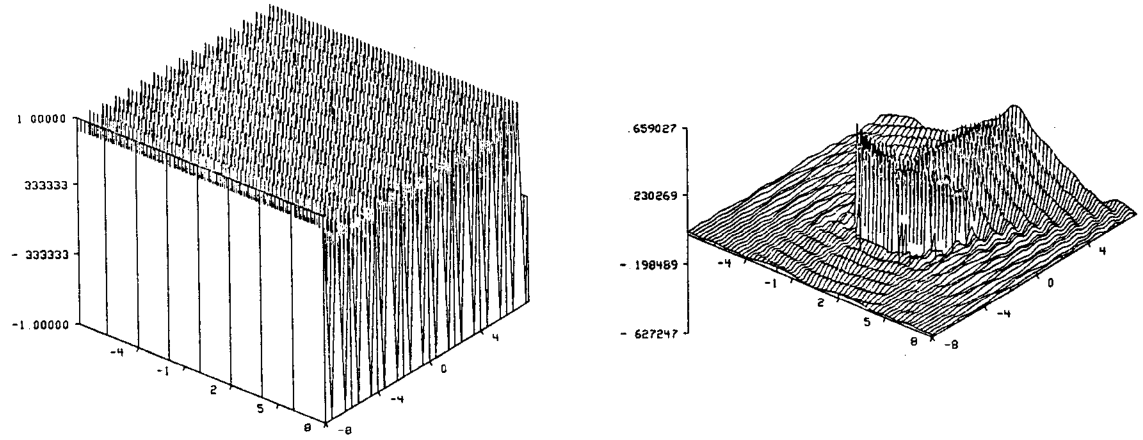

2 Iterations

3 Iterations

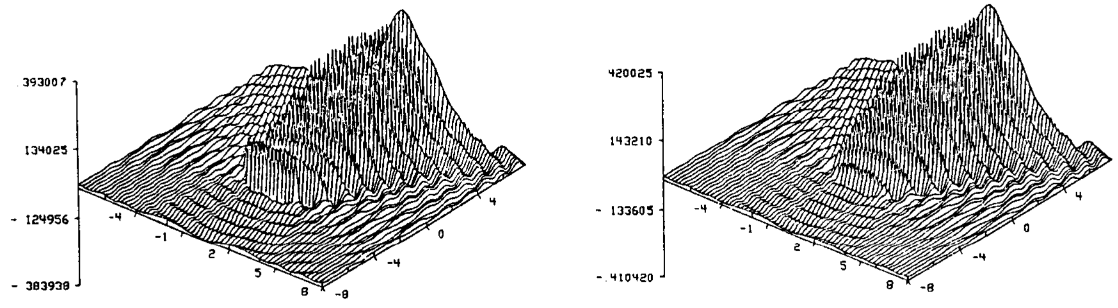

Figure 6.27 Three iterations of the Kaczmarz algorithm are shown to demonstrate the convergence of the approach to a single answer.

mean squared error by a factor of two. Since the orthogonalization is done with several dot products the reduction in iterations more than balances the extra work.

An alternative approach is rearrange the order of the equations to reduce their interdependency. This idea was first published by Hounsfield in the original patent for CT imaging [Hou72]. It seems reasonable that the hyperplanes describing the field at adjacent points (less than a quarter wavelength apart) would be nearly parallel. When the solution vector is first projected onto one plane then projecting onto a second parallel hyperplane will not significantly improve the answer. Thus convergence will be faster if the parallel hyperplane is saved till later in the iteration sequence.

This idea is illustrated in Figure 6.28 for a two dimensional case with four hyperplanes. In each case the order projections are considered is indicated as the hyperplane number at the end of the line. First in Figure 6.28a the two sets of parallel planes are considered separately while in Figure 6.28b the two sets are interleaved. It is easy to see that the first ordering will take twice as many iterations as the second.

To calculate the scattered field using the Kaczmarz algorithm a system of equations is set up that represents the field at each point as a function of the refractive index distribution. For ease of programming the x vector, the field at each point in the grid, is organized so that adjacent elements in the vector represent the field at adjacent points in the grid. Thus if the hyperplanes are considered in order there will be a high degree of correlation between the equations and convergence will be slow.

The degree of independence of two equations can be found by finding the angle between the two hyperplanes. If two equations contain nearly the same information then the angle between their respective hyperplanes will be zero while if the two equations are independent there will be an angle of 90 degrees. From standard vector theory the angle between two hyperplanes is defined as

$$\cos\theta = \frac{<A_i,A_j>}{\sqrt{(<A_i,A_i><A_j,A_j>)}} \tag{6.81}$$

where A_i and A_j represent the rows of the A matrix or the normal vector to the two hyperplanes.

The order equations are considered is a function of a parameter called δ_E that represents the change in equation number. For each iteration of a N x N grid the parameter i steps from 0 to N^2-1 and is mapped into an equation number, j, by the relation

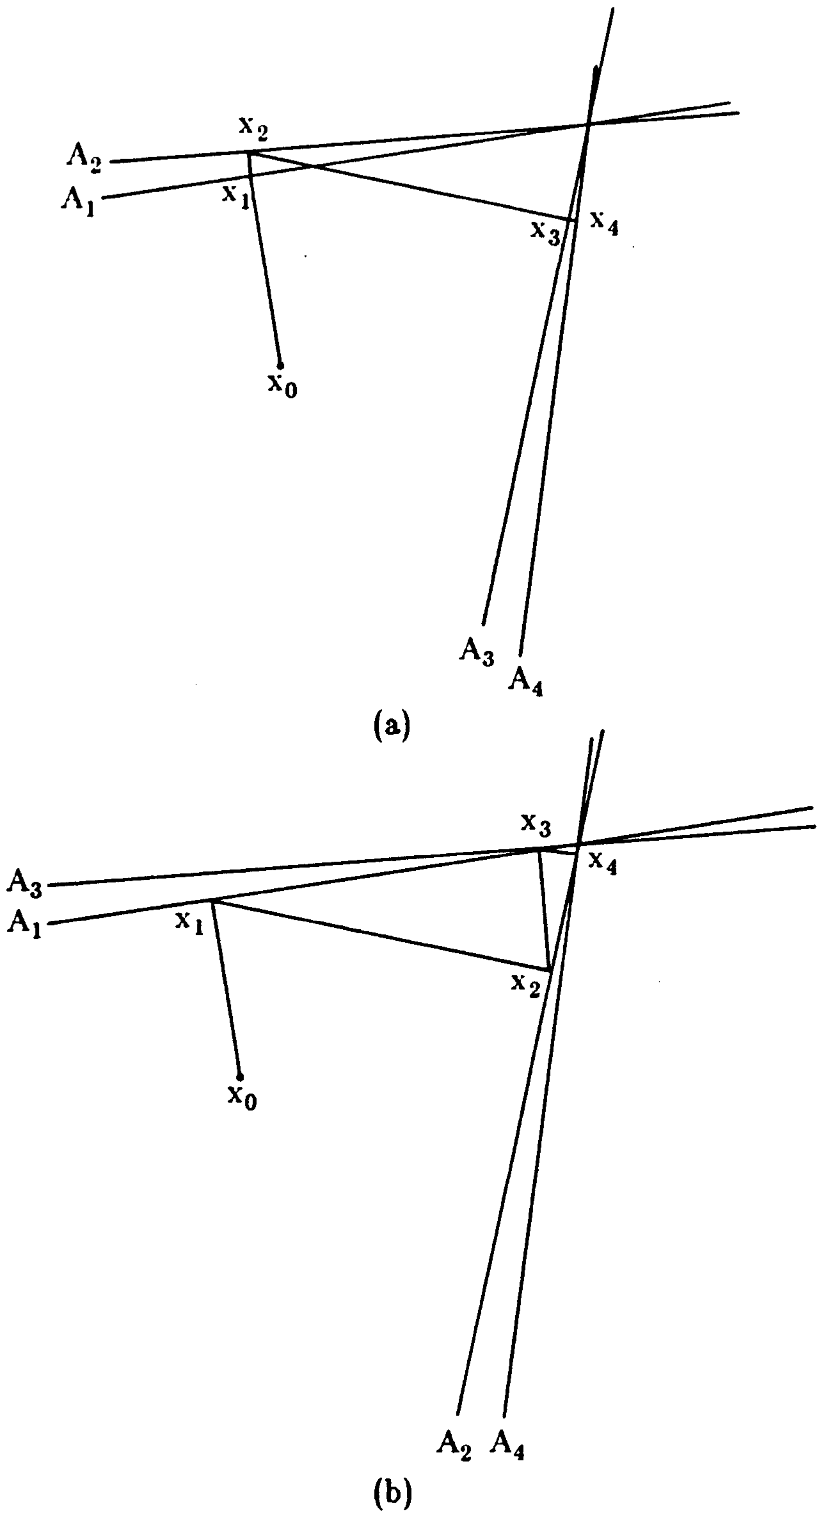

Figure 6.28 The order equations are considered can affect the rate of convergence. If the two sets of parallel lines are considered separately (a) then the convergence is twice as slow as it would be if they are interleaved (b).

$$j = (i*\delta_E) \bmod N^2. \tag{6.82}$$

With the proper choice of δ_E the equation number j will step through all N^2 equations.

The table below 6.2 shows the average and maximum cosine of the angle for a range of refractive indices between 1.01 and 1.5. A system of 1024 equations is used to define the scattered field over a 32 x 32 grid. Thus when the parameter δ_E is 1 adjacent equations are compared while the equations are as far apart as possible when δ_E is set to 513.

Table 6.2. The average and maximum cosine of the angle between hyperplanes is shown as a function of the refractive index and the number of equations skipped.

Refractive Index	δ_E	Average $\cos\theta$	Maximum $\cos\theta$
1.01	1	.000592	.010902
	513	.000125	.001338
1.1	1	.019414	.1682
	513	.001981	.013842
1.2	1	.064287	.377699
	513	.005448	.025288
1.5	1	.238268	.623253
	513	.022828	.06990

From the information in this table two points are apparent. First, as the refractive index is increased both the average and the maximum cosine of the angle increases. For small refractive indices the A matrix is dominated by the diagonal terms and thus each equation is nearly independent of the others. This also explains why the Kaczmarz algorithm converges much faster for small refractive indices.

Secondly, equations describing the field at widely spaced points have a large angle between their respective hyperplanes. As can be seen from the above table by comparing equations that are widely separated the average and the maximum cosine of the angle are reduced by a factor of 10. Thus even with a refractive index as great as 1.5 the minimum angle between hyperplanes is increased to 86 degrees by skipping a large number of equations.

The advantages gained by considering non adjacent equations is confirmed by considering the convergence of the algorithm. For the system of equations defined by $Ax = b$ the residue is given by

$$\text{Residue} = \langle A_i, x \rangle - b. \tag{6.83}$$

The residue measures the distance between the solution vector x and the hyperplane described by A_i and the total residue is defined as the sum of the squares of the residues from each row of A. This figure can then be used as a measure of quality of the solution.

The total residue when calculating the scattered field from a cylinder of radius 1λ and refractive index 1.1 is shown in Figure 6.29. The study was done for 16 iterations and compares the total residue when the equations are considered in order (δ_E is 1) and when the equations are widely separated (δ_E is 513). To reach any given total residue, iterating the solution by considering adjacent equations takes twice as long to converge as when the equations are widely separated.

Considering widely separated points gives the same benefits as pair wise orthogonalization (convergence is twice as fast) but without the extra work. In addition when the equations are widely separated the planes are nearly perpendicular and thus there is little to be gained by pair wise orthogonalization.

It is also possible to study the effect of projecting the field in a non sequential fashion by considering the field after one iteration. Figure 6.30 shows the exact (dashed line) and the Kaczmarz field (solid line) from a cylinder of radius 1λ. The cylinder has a refractive index of 1.01 and the Kaczmarz iteration is carried out over a 32 x 32 grid. In addition, to emphasize the difference made by changes in δ_E, the initial guess, x_0, for the field was the incident field.

While at first glance all of the Kaczmarz fields are very poor approximations to the exact field it is important to note that after one more iteration all of them have converged to the correct answer. Thus the only effect of altering the order of the equations is to change the rate of convergence.

The reason for the difference in one iteration of the Kaczmarz is best seen by comparing the field when δ_E is one and 1023. Recall that skipping 1023 equations in a system of 1024 equations projects adjacent equations, like the

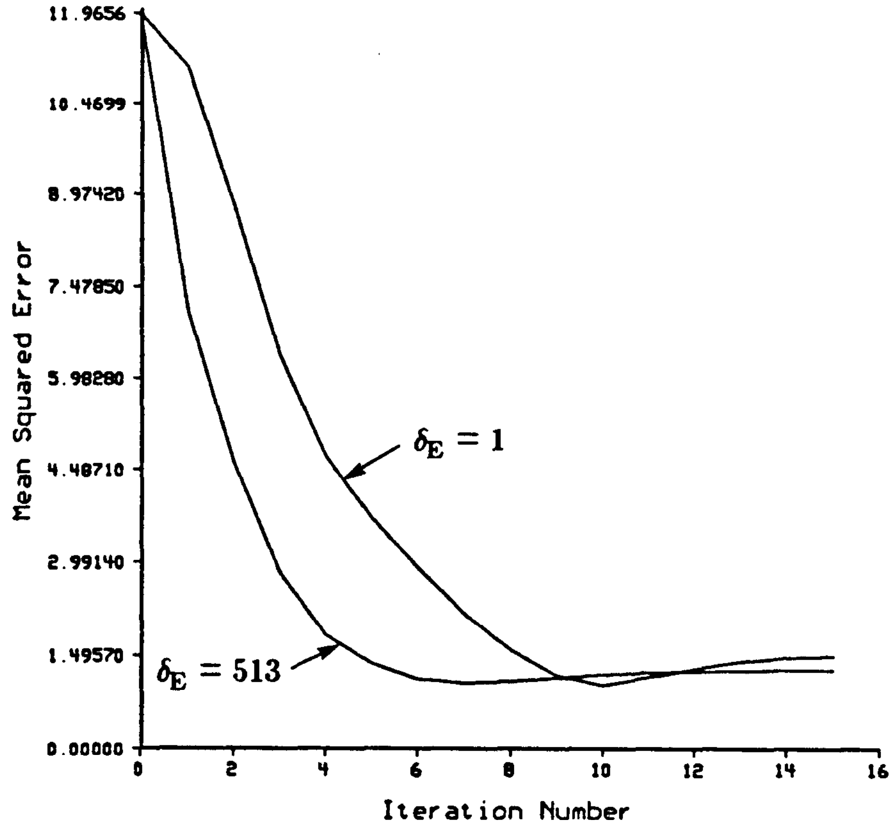

Figure 6.29 The residue remaining after the first 16 iterations of the Kaczmarz algorithm are shown as a function of the number of equations skipped (δ_E).

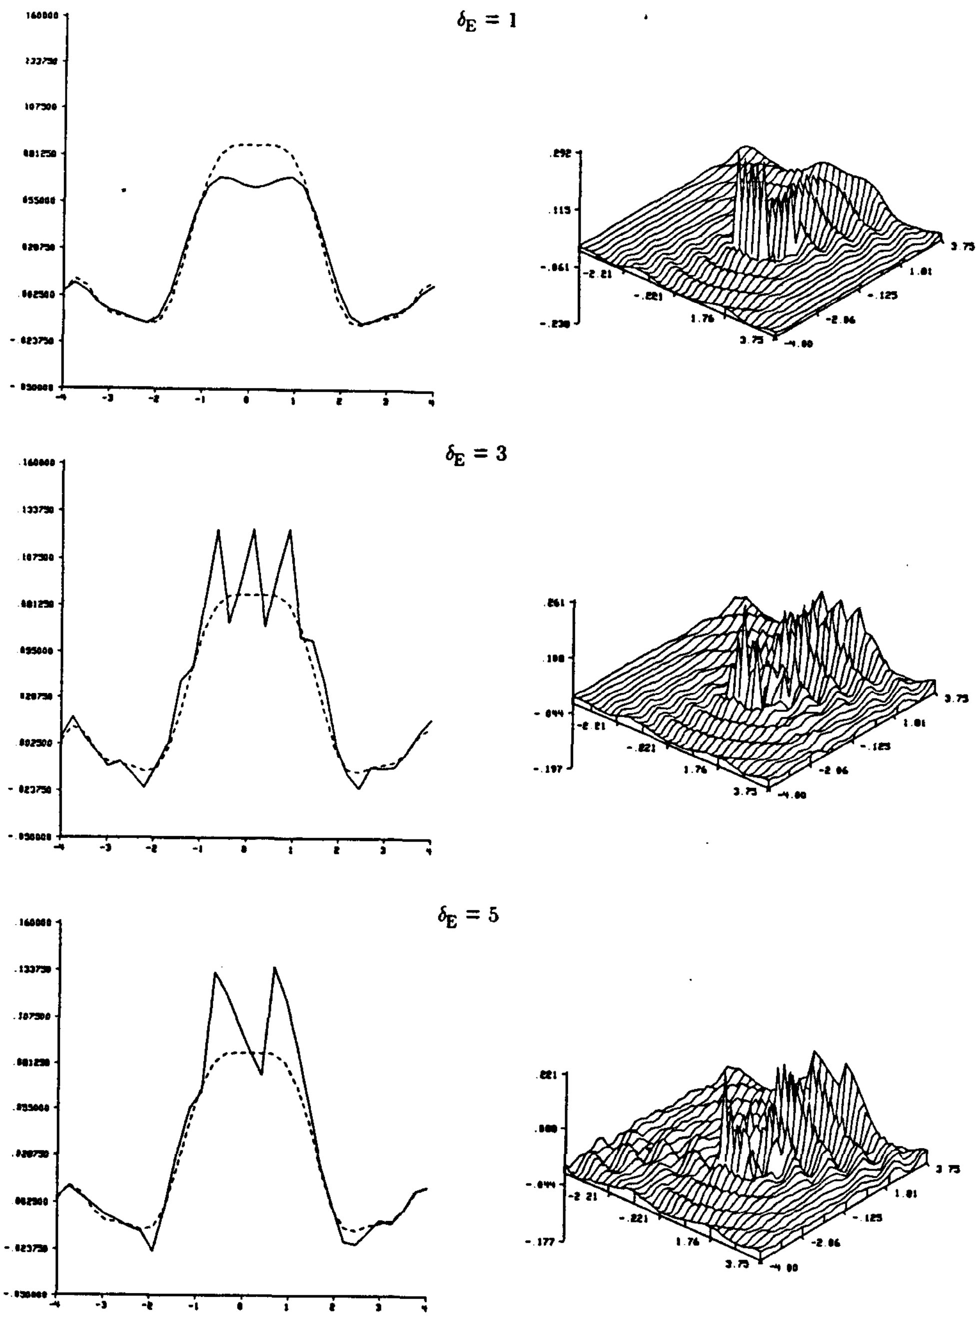

Figure 6.30 The field after one iteration is shown as a function of the number of equations skipped.

$\delta_E = 513$

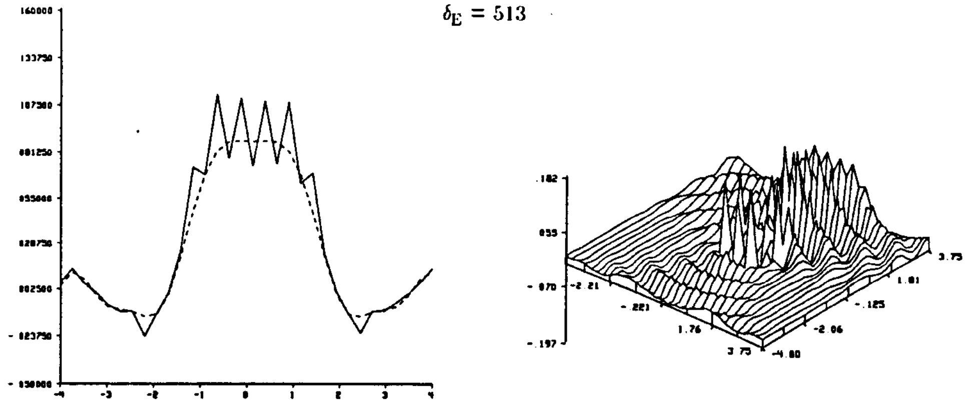

$\delta_E = 1023$

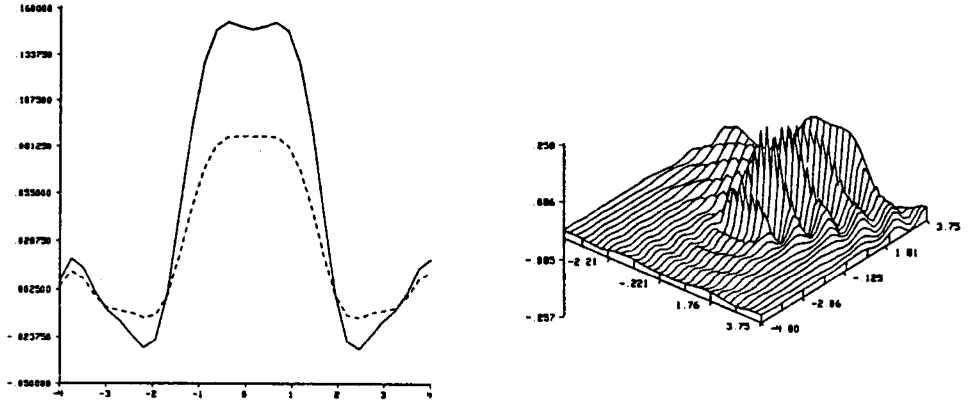

Exact Solution

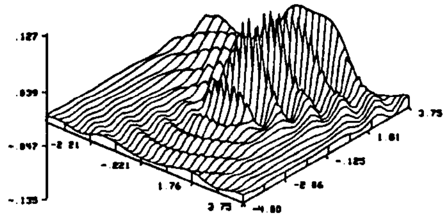

Figure 6.30 Continued.

original case with $\delta_E=1$, but instead the equations are treated in reverse order.

Since the diagonal elements of the A matrix dominate the matrix with a refractive index of 1.01 each projection only modifies the field at one point in the grid. Thus with $\delta_E = 1$ the field is calculated in the same direction as the incident field travels, from negative to positive y. This means that the field inside the cylinder has been calculated before the field is calculated at the receiver line. On the other hand when the equations are projected with $\delta_E = 1023$ the field at the receiver line is calculated before the field has been adjusted for the scattering inside the cylinder.

Unlike the fixed point algorithms (Born and Rytov series) the Kaczmarz algorithm should always converge to the correct scattered field. The speed the series converges is related to the average orthogonality of the defining equations and as shown in Table 6.2 the equations become more dependent as the object's refractive index increases. Thus the Kaczmarz algorithm converges more slowly for objects with large refractive index.

Two other factors determine the rate of convergence of the Kaczmarz series. As shown in Figure 6.29 the order the equations are considered changes the rate of convergence of the series. For this reason most of the simulations shown in this work were done using a δ_E equal to $N^2/2+1$ where the size of the grid is NxN. While it hasn't been investigated, even better performance could possibly be obtained by alternating different values for δ_E.

Finally quantization and sampling errors slow the rate of convergence. This is shown in Figure 6.31 for a sampling interval of .1λ and .25λ. In both cases the radius of the cylinder is equal to eight times the sampling interval and the calculations were performed over a 32x32 grid. The plots show the real (solid line) and imaginary (dashed line) components of the scattered field as calculated by the Kaczmarz algorithm and an exact algorithm based on the Bessel function expansions described in [Wee64 and Mor68].

Each of the plots in Figure 6.31 shows the scattered field after 32 iterations. While the exact limits of the Kaczmarz algorithm are difficult to define it is possible to say that the Kaczmarz algorithm has converged for refractive indices up to 1.4 with a sampling interval of .1λ and up to 1.2 with a sampling interval of .25λ. Calculating the scattered field from objects with larger refractive indices might be possible but will need more accurate implementations.

All the results shown here were sampled using rectangular basis functions. While according to the sampling theorem a rectangular set of basis functions

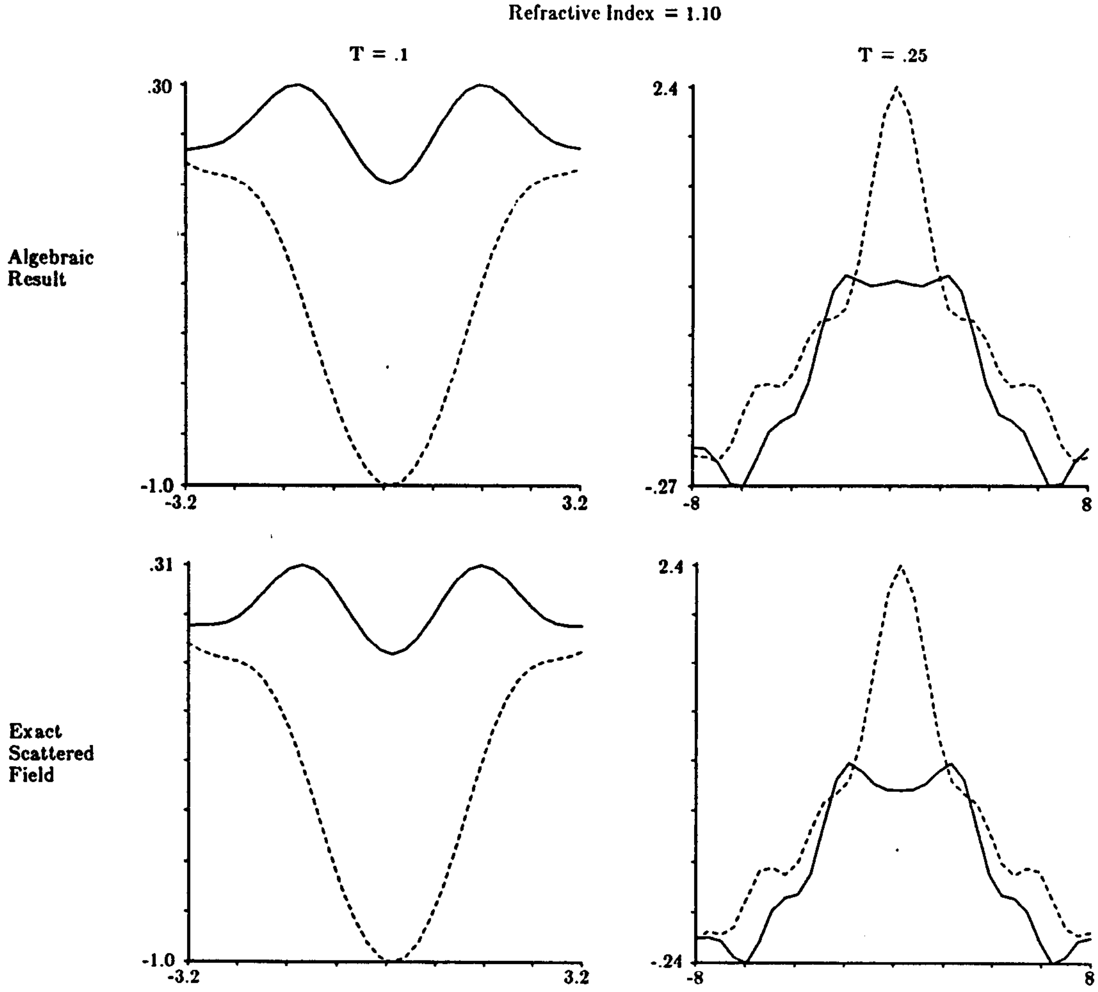

Figure 6.31 The exact scattered field and the result of 32 iterations of the Kaczmarz algorithm are compared here. In each case the real component of the field is shown as a solid line while the imaginary component is shown as a dashed line.

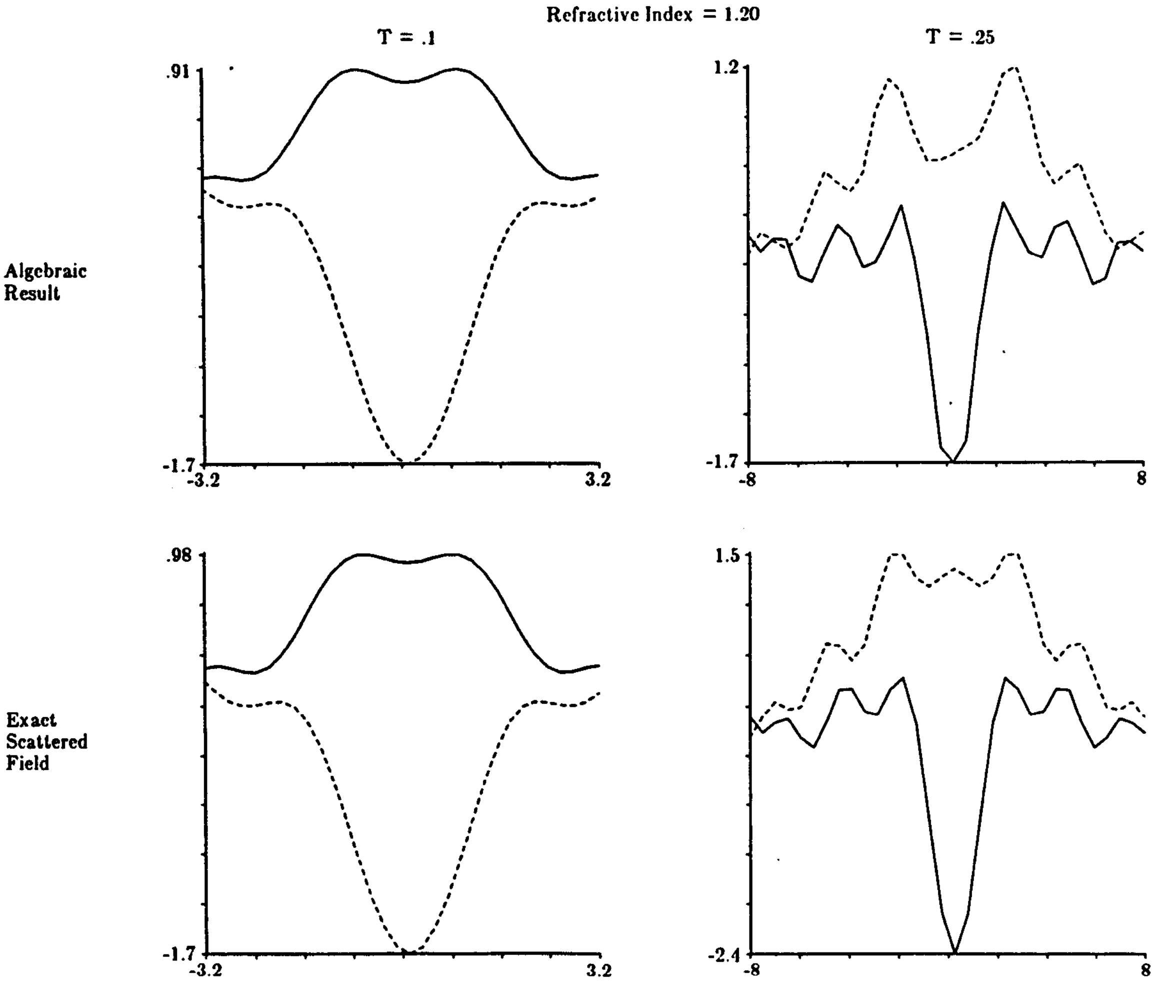

Figure 6.31 Continued.

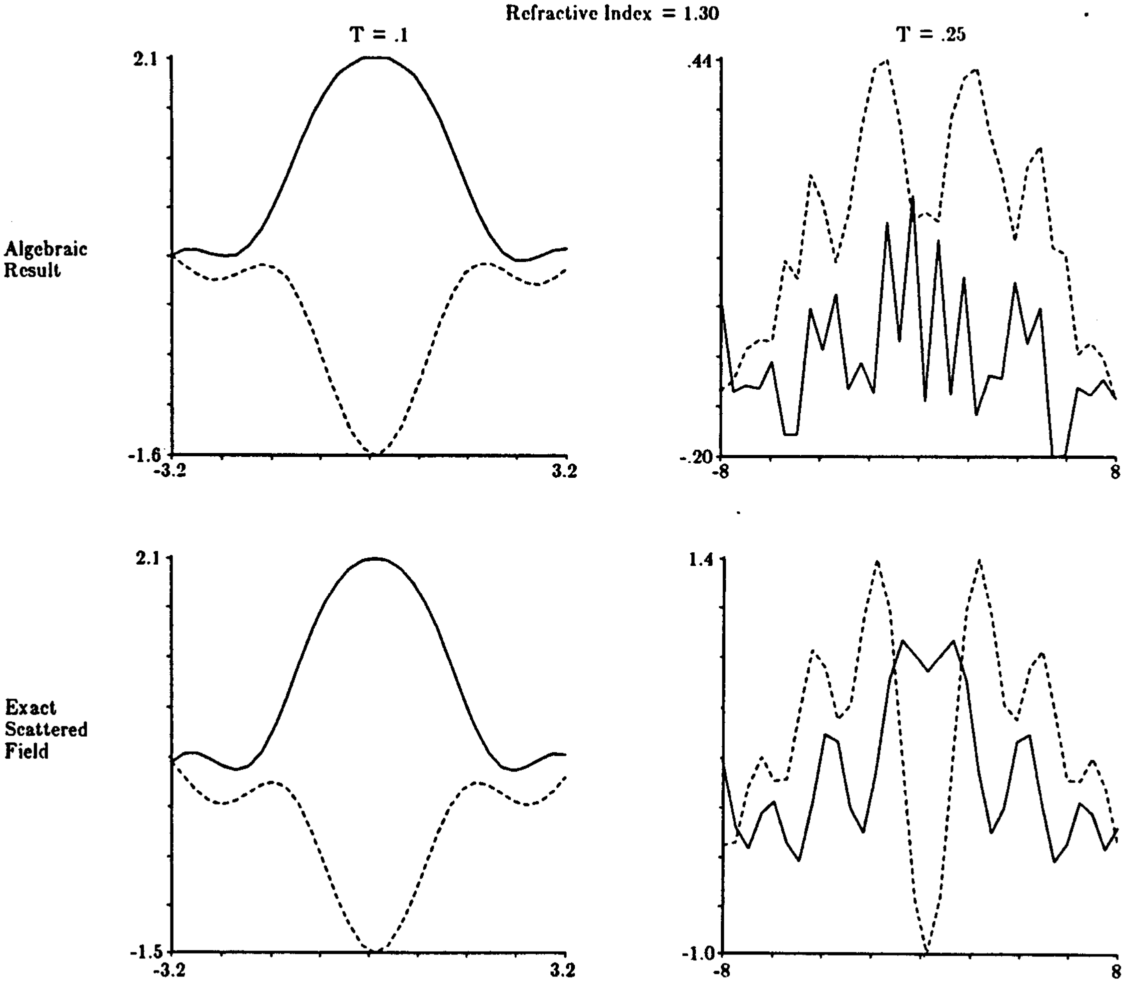

Figure 6.31 Continued.

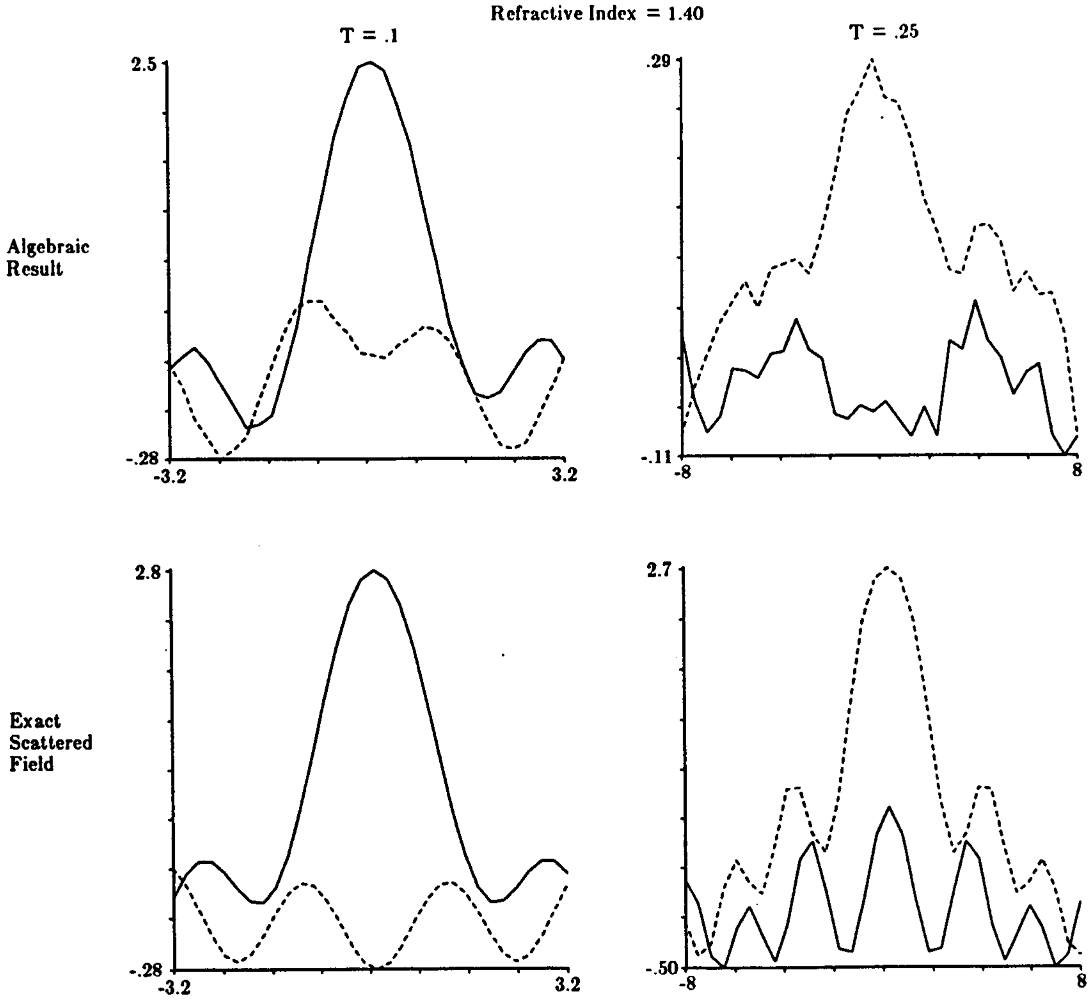

Figure 6.31 Continued.

can represent the original continuous function without errors more accurate integrations might be possible using sinc functions [Joh83, Ste81]. This idea has not been explored here.

References

[Azi83] M. Azimi and A. C. Kak, "Distortion in diffraction imaging caused by multiple scattering," *IEEE Transactions on Medical Imaging,* Vol. MI-2, December 1983, pp. 176-195.

[Cen81] Y. Censor, "Row-actions methods for huge and sparse systems and their applications," *SIAM Review,* Vol. 23, October 1981, pp. 444-466.

[Goo68] J. W. Goodman, *Introduction to Fourier Optics,* McGraw Hill Book Company, San Francisco, 1968.

[Her73] G. T. Herman, A. Lent, and S. Rowland, "ART: Mathematics and applications: A report on the mathematical foundations and on applicability to real data of the algebraic reconstruction techniques," *J. Theor. Biol.,* Vol. 43, 1973, pp. 1-32.

[Her76] G. T. Herman, A. V. Lakshminarayanan, A. Naparstek, E. L. Ritman, R. A. Robb, and E. H. Wood, "Rapid computerized tomography," *Med. Data Process.,* 1976, pp. 582-598.

[Hey85] G. T. Heydt, *Computer Analysis Methods for Power Systems,* MacMillan, New York, 1985.

[Hou72] G. N. Hounsfield, *A method of apparatus for examination of a body by radiation such as x-ray or gamma radiation,* Patent Specification 1283915, The Patent Office, 1972.

[Joh83] S. A. Johnson and M. L. Tracy, "Inverse scattering solutions by a sinc basis, multiple source, moment method -- Part I: Theory," *Ultrasonic Imaging,* Vol. 5, 1983, pp. 361-375.

[Kac37] S. Kaczmarz, "Angenaherte auflosung von systemen linearer gleichungen," *Bull. Acad. Polon. Sciences et Lettres A,* 1937, pp. 355-357.

[Kel69] J. B. Keller, "Accuracy and validity of the Born and Rytov approximations," *Journal of the Optical Society of America,* Vol. 59, 1969, pp. 1003-1004.

[Man70] K. Mano, "Interrelationship between terms of the Born and Rytov expansions," *Proceedings of the IEEE,* Vol. 58, July 1970, pp. 1168-1169.

[Mor68] P. M. Morse and K. U. Ingard, *Theoretical Acoustics,* McGraw Hill Book Company, New York, 1968.

[Opp75] A. V. Oppenheim and R. W. Schafer, *Digital Signal Processing,* Prentice-Hall, Englewood Cliffs, NJ, 1975.

[Ram79] R. S. Ramakrishnan, S. K. Mullick, R. K. S. Rathore, and R. Subramanian, "Orthogonalization, Bernstein polynomials, and image restoration," *Appl. Opt.,* Vol. 18, 1979, pp. 464-468.

[Ric65] J. Richmond, "Scattering by a dielectric cylinder of arbitrary cross section shape," *IEEE Transactions on Antennas and Propagation,* Vol. AP-13, 1965, pp. 334-341.

[Ros82] A. Rosenfeld and A. C. Kak, *Digital Picture Processing,* Academic Press, Second Edition, 1982.

[San70] M. I. Sancer and A. D. Varvatsis, "A comparison of the Born and Rytov methods," *Proceedings of the IEEE,* Vol. 58, January 1970, pp. 140-141.

[Sla83] Malcolm Slaney and A. C. Kak, "Diffraction tomography," *Proceedings of the SPIE,* Vol. 413, April 1983, pp. 2-19.

[Sla84] Malcolm Slaney, A. C. Kak, and L. E. Larsen, "Limitations of imaging with first order diffraction tomography," *IEEE Transactions on Microwave Theory and Techniques,* August 1984, pp. 860-873.

[Ste81] F. Stenger, "Numerical methods based on Whittaker cardinal, or sinc functions," *SIAM Review,* Vol. 23, April 1981, pp. 165-224.

[Sto80] Josef Stoer and Roland Bulirsch, *Introduction to Numerical Analysis,* Springer-Verlag, New York , 1980.

[Tan71] K. Tanabe, "Projection method for solving a singular system," *Numer. Math.,* Vol. 17, 1971, pp. 203-214.

[Tra83] M. L. Tracy and S. A. Johnson, "Inverse scattering solutions by a sinc basis, multiple source, moment method -- Part II: Numerical evaluations," *Ultrasonic Imaging,* Vol. 5, 1983, pp. 376-392.

[Wee64] W. L. Weeks, *Electromagnetic Theory for Engineering Applications*, John Wiley and Sons, Inc., New York, 1964.