Iterative solution of multiple right hand side matrix equations for frequency domain monostatic radar cross section calculations

M D Pocock, S P Walker

Computational Mechanics, Imperial College of Science Technology and Medicine, London SW7 2BX
s.p.walker@ic.ac.uk

Abstract

Monostatic rcs characterisation using integral equation methods in the frequency domain requires the solution of very large matrix equations with multiple right hand sides. Although costly for a single right hand side, direct methods are attractive in that subsequent right hand sides are very cheap. Iterative methods are much cheaper for a single right hand side, but if the whole solution must be repeated for each, they become much more expensive. We investigate here the performance of a simple modification to the GCR algorithm, which allows solutions for an essentially unlimited number of right hand sides to be obtained for a modest multiple of the cost of the first. For the cases investigated, with up to 360 right hand sides on bodies up to 15 wavelengths long, with matrices up to 20,440 by 20,440 in size, this multiple was below ~10. Costs seem to rise with the number of right hand sides till the surface field is in some sense characterised, and thereafter subsequent illumination angles are essentially free. An investigation of cost scaling on a set of spheres, ranging from ~1 to 7 wavelengths in diameter, seems to indicate the cost of full monostatic characterisation to scale with about the fifth power of frequency.

1. Introduction

There are two main cost components in the solution of large scattering problems, such as rcs evaluation, via frequency domain integral equation methods. The first[1], scaling with frequency f to the fourth power, is the cost of formation of the dense matrix describing interactions between parts of the scatterer. The second is the cost of solving the resulting matrix equation.

For the direct solution schemes most usually employed, this solution cost scales with N^3, where N is the order of the matrix equation, or equivalently with f^6. For all but small problems matrix solution is thus the dominant cost. Because of this large cost, there has been increasing attention paid of late to the use of iterative methods for matrix solution[2-8]. It is generally concluded from such studies that iterative schemes require much less time to compute, often by an order of magnitude or more, than would be taken to solve the same system by direct methods. Cost reduction by a large multiple is obviously welcome; even more so would be if that multiple itself increased with problem size, corresponding to a reduction in cost scaling below the f^6 of the direct approach. This has indeed been suggested[7], but the evidence is as yet inconclusive[8].

For practical radar cross section analyses, it is necessary to determine the scattered field caused by fields incident from a (large) number of directions. Direct solution methods are then attractive. Once the f^6 cost of matrix factorisation has been paid, solutions for large numbers of incident waves (right hand side vectors in matrix equation terms) can each be found at a cost scaling with f^4. Even if the required number of right hand sides scales with the second power of frequency, the overall cost of a full monostatic evaluation still only(!) scales with frequency to the sixth power.

This interest in multiple illumination angles is a discouragement to the use of iterative methods. One approach using them is simply to start the solution afresh for each right hand side. Naturally, if more than a few right hand sides are required, any cost saving over the direct approach is lost. If the required number of right hand sides scales with the second power of frequency, the overall cost then scales with something approaching f^8.

We present here methods which seem able to provide iterative solutions for additional monostatic illumination angles (right hand sides) at essentially no cost, once sufficient right hand sides have been analysed to characterise, in some sense, the surface field.

There have been a few attempts to address this multiple right hand side issue for iterative solutions, which we will discuss prior to describing the present approach.

In general the quality of the initial guess has only a modest affect on the number of iterations required for convergence of an iterative method[2-4,7]. Consequently, the approach of using the solution from one incident wave as the starting point for the next offers little benefit.

For the sparse matrices resulting from a finite element discretisation, Smith and colleagues[6] effectively expanded several solutions in terms of a single set of search vectors, with some success.

In a very recent paper[9] Boyse and Seidl use the GMRES algorithm (and see references cited in Boyse for more details of the block GMRES and the multiple right hand side variant thereof). In essence the problem is first solved simultaneously for a modest number of right hand sides, distributed uniformly over the span of right hand sides of eventual interest, using a block GMRES approach. Whilst this costs more than for a single right hand side, it is not a large multiple of the cost. Intermediate values are then found, using the orthonormal basis for the Krylov sub-space which was found during the solution for the initial right hand sides. It was found necessary to ensure that the number of right hand sides solved for initially was carefully chosen. This number of right hand sides was noted as needing to be sufficient to represent in some sense the RCS distribution sought. An angular separation such as to allow rather more than two samples per wavelength was suggested. If this is not observed, intermediate solutions found subsequently tended to have significant errors. Too many initial right hand sides, however, caused slow convergence of the block GMRES solution. Overall, significant reduction in time relative to a direct solution was observed.

The GCR (generalised conjugate residual) algorithm[10] was used by Soudais[5] for solving the matrix equations resulting from analysis of scattering from a ~1 wavelength mixed dielectric-perfect conductor target. A finite element and symmetric Stratton-Chu treatment were combined, giving a matrix system which was symmetric, and sparse in the finite element regions. Multiple right hand sides were tackled by what seemed to be a very effective, and attractively simple, modification to the normal GCR algorithm. As with most iterative methods, the solution is changed at each iteration by moving a particular distance (the step length) in some direction (the search direction). The computationally expensive part of the algorithm is the finding of this search direction. In the modified algorithm, the search direction is found for only that right hand side currently exhibiting the largest residual. The solutions to all right hand sides are advanced in this direction, but for a different (and optimal) step length in that direction for each.

In the sections which follow we will extend the application of this approach, to the dense and unsymmetric matrices which result from a normal integral equation discretisation, and to analysing scattering from multi-wavelength 'stealthy' targets. A major aspect of interest will be the interaction of the number of iterations required (a measure of the computational work), with the number of right hand sides analysed, and the body size in wavelengths. It is this latter which determines the 'jaggedness' with angle of the monostatic radar cross section, and consequently the number of illumination angles required to characterise it. If the number of iterations continues to rise in proportion to the number of illumination angles, till the point that the response is fully characterised, there may be no reduction in cost scaling; if otherwise, there could be.

In the next section we will summarise the integral equation formulation, and outline the normal GCR algorithm, and the modification to handle multiple right hand sides. Section three will present results from its application to a number of rcs problems.

2. Formulation

The formulation is well known, and only a brief description will be given here. The MFIE for scattering from a smooth, closed perfect conductor is

$$\tfrac{1}{2}\mathbf{H}(\mathbf{r}) = \mathbf{H}^{inc}(\mathbf{r}) + \frac{1}{4\pi}\int_{\partial\Omega}\left[\mathbf{n}'\times\mathbf{H}(\mathbf{r}')\right]\times\mathbf{R}\left(jk-\frac{1}{R}\right)\frac{e^{jkR}}{R^2}ds' \qquad (1)$$

where the integrations to obtain the field at surface location $\mathbf{r}$ are over the rest of the surface of the scatterer s' and $\mathbf{r}'$, with $\mathbf{n}$ the unit normal, and $\mathbf{R}$ the vector $\mathbf{r}-\mathbf{r}'$. The incident wave is $\mathbf{H}^{inc}$. We employ a curvilinear, isoparametric discretisation, with Gaussian quadrature. Fuller details are given elsewhere[11]; the result is a (complex) matrix equation of the form:

$$[\mathbf{A}]\mathbf{H} = \mathbf{H}^{inc} \qquad (2)$$

This is to be solved with tens or possibly hundreds of incident wave vectors $\mathbf{H}^{inc}$.

We use the unpreconditioned GCR algorithm, which we first outline in its normal 'single right hand side' form.

To solve

$$\mathbf{A}\mathbf{x} = \mathbf{b} \tag{3}$$

Initialise:

$$\begin{aligned} \mathbf{x}^{(0)} &= \mathbf{0} \\ \mathbf{r}^{(0)} &= -\mathbf{b} \\ \mathbf{p}^{(0)} &= -\mathbf{r}^{(0)} \end{aligned} \tag{4}$$

For each iteration k, for k = 0, ...:

Calculate the step length α:

$$\alpha^{(k)} = -\frac{\left\langle \mathbf{r}^{(k)}, \mathbf{A}\mathbf{p}^{(k)} \right\rangle}{\left\| \mathbf{A}\mathbf{p}^{(k)} \right\|^2} \tag{5}$$

Increment the solution vector

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \alpha^{(k)}\mathbf{p}^{(k)} \tag{6}$$

and residual vector:

$$\mathbf{r}^{(k+1)} = \mathbf{r}^{(k)} + \alpha^{(k)}\mathbf{A}\mathbf{p}^{(k)} \tag{7}$$

Check if converged sufficiently: stop if

$$\frac{\left\| \mathbf{r}^{(k+1)} \right\|}{\left\| \mathbf{r}^{(k)} \right\|} \le 10^{-m} \tag{8}$$

with m typically in the range 2 to 8.

If not, calculate search direction vector and matrix vector product for the next iteration:

$$\sigma_i^{(k+1)} = \frac{\left\langle \mathbf{A}\mathbf{r}^{(k+1)}, \mathbf{A}\mathbf{p}^{(i)} \right\rangle}{\left\| \mathbf{A}\mathbf{p}^{(i)} \right\|^2} \quad \text{for } i = 1,\ldots k \tag{9}$$

$$\mathbf{p}^{(k+1)} = -\mathbf{r}^{(k+1)} + \sum_{i=1}^{k} \sigma_i^{(k+1)}\mathbf{p}^{(i)} \tag{10}$$

$$\mathbf{A}\mathbf{p}^{(k+1)} = -\mathbf{A}\mathbf{r}^{(k+1)} + \sum_{i=1}^{k} \sigma_i^{(k+1)}\mathbf{A}\mathbf{p}^{(i)} \tag{11}$$

As well as the main system matrix, we see that the sets of vectors $\mathbf{p}^{(k)}$ and $\mathbf{A}\mathbf{p}^{(k)}$ must be stored.

The modifications required to solve for multiple right hand sides are very simple. In the initialisation stage a residual and initial search direction is obtained for each right hand side.

At equation (5) a step length is calculated for each right hand side, and the solution and residual vectors incremented. Between (7) and (8) the largest (normalised) residual is found, by searching through the residuals associated with all right hand sides.

Assuming it fails the test (8), a single search direction is calculated in (9) and (10), to suit that right hand side found to have the largest residual. This step, involving matrix vector multiplication, is the one where the principal computation lies. A different step length for each right hand side is then calculated in (5), which is used with the common search direction to increment each solution in (6). We see that at each iteration the parameters selected most suit that right hand side currently furthest from solution. As a result that right hand side improves rapidly, with some other right hand then side taking its place as the current worst. Thus are all right hand sides shepherded to solution more or less in step.

Each right hand side is seen to increase by one the number of vector-vector manipulations required at each iteration (multiplications in equation (5)), and vector-vector additions in equations (6), (7) and (8)). Storage is increased by the need to store the evolving solution and residual for each right hand side. For problems of practical interest, these increases are a modest fraction of the requirements for a single right hand side. As the solutions are attained, it would be possible to continue to iterate only for those whose residual is not yet below the specified tolerance. Although this would save some cost, the saving would be small, and the refinement has not been implemented.

3. Demonstrations

In the sections which follow we present results of the analysis of targets illuminated by many different incident waves. Of particular interest will be the variation of the total number of iterations with the number of illumination angles, and, for practical purposes, the variation of the rcs over that range of illumination angles. Does the

number of iterations rise with illumination angles as long as 'new' information is being gained?

3.1 Equatorial scans

We consider first a monostatic equatorial scan of the NASA almond[12], illuminated at 7 GHz, making it ~6 wavelengths long. Discretisation employed an average nodal separation of about 1/9 of a wavelength, and throughout a termination residual of 10^{-4} was employed. This is a value recently suggested[13] as appropriate for this discretisation, for the curvilinear isoparametric modelling employed.

Figure 1 shows the variation in the number of iterations required with the number of illumination angles. The monostatic scan was computed as a single run, with respectively 1, 5, 10, ... up to 180 illumination angles (right hand sides) in turn. Cases were studied with these illumination angles distributed uniformly over 180^0, 90^0, 45^0 and 15^0, forming the lines shown on the figure.

We see that in each case the number of iterations required initially rises with the number of illumination angles examined. For the 180^0 case, this rise is essentially complete some time before ~30 illumination angles, by when about six times as many iterations are required as are required for a single illumination angle. By this point we are gaining ~30 solutions for a multiple of ~6 in the work required. From then on solutions for additional angles are obtained without further iterations; the maximum 180 considered are similarly gained with this same multiple of ~6 over the cost of one. (As discussed in section 2 above, there is still some modest angle-dependent cost, in the vector manipulations of (5) - (8). This is very small in practice, and will be neglected in subsequent discussion.) Similar comments can be made about cases with the illumination angles distributed over a smaller range. As the extent is reduced, in each case the 'plateau' in iterations occurs at a lower number of iterations, and after fewer illumination angles. For example, for a 15^0 degree range, the plateau in number of iterations is reached after about 10 illumination angles compared to the ~30 of the 180^0 case. This is consistent with the number of iterations being a function of the amount of information being sought; reducing the range reduces the amount.

There is an interesting parallel here with the observations of Boyes and Seidl[9]. Using their very different approach, they found costs to rise essentially linearly with number of illumination angles till the response was partially characterised, and that subsequent illumination angles were essentially free.

However, the response in that reference was considered primarily in terms of rcs, although the actual unknown for which we are solving a set of equations is the surface field. The monostatic rcs, the quantity actually sought, and the surface field at any particular surface location, are very different functions of incident illumination angle. The (monostatic) rcs generally is a much less smooth a function of incident illumination angle than is the surface field. (For simple geometries it can of course be much more smooth, as the case of the monostatic rcs of a multi-wavelength sphere attests.)

We find some empirical evidence that the plateau is reached when the density of illumination angles is such as to represent reasonably the variation of surface field with illumination angle. In figure 2 is shown an indication of the variation of surface field with illumination angle for the six wavelength almond, by selecting a node (at -4.1", 0, 0) and component (the real part of the y component) at random. The locations of 30 uniformly spaced illumination angles are marked on this figure, and it could be argued that results from only this many illumination angles provide a reasonably good representation of the surface field variation.

The corresponding rcs is shown in figure 3, with the solid line obtained by illumination from 180 angles, and the 30 locations again marked. (This present paper is not concerned with characterisation of an rcs code, but for completeness this figure also shows the measured rcs[12], albeit with the measured values obtained from an enlarged photocopy of the published measurements. As is seen, agreement of the computed values with this is good.) We see that much of the detailed structure of the rcs is not revealed by using only these 30 angles; lobes on the RCS plot are generally seen to be significantly narrower than on that of the surface field. However, with the 30 illumination angles corresponding to the start of the plateau, this detail is obtained essentially free. (We have not investigated the point, but this might be interesting to consider in the context of the 'monostatic - bistatic approximation'[14]. Here relatively widely separated illumination angles can be used to provide a good approximation to the finely sampled monostatic response. It may indeed be that similar criteria apply.)

Timings may be of interest. On a Dec Alpha 600 workstation, the matrix took about 16 minutes to form, and 25 minutes to solve for each illumination, giving a total time of 75 hours for a 180 illumination angle characterisation. Using the multiple right hand side approach, these same 180 solutions were obtained in about 3 hours.

As a further example, figure 4 (inset) shows the monostatic rcs of a 11:1 cylindrical dipole, with hemispherical end caps, discretised with 3026 nodes, and illuminated such as to make it ~15 wavelengths long. This figure was drawn from results evaluated with 360 illumination angles 0.5^0 apart, and shows approximately 60 distinct peaks over the 180^0. It is not shown, but the variation of surface field with illumination angle is not surprisingly very much smoother than the rcs on this 15 wavelength body; indeed it is very smooth for many surface locations, with ~15 peaks being the maximum found. Figure 4 shows the variation of the number of iterations with the number of illumination angles. Again, we see a distinct plateau occurring, here after ~90 illumination angles. This is sufficient to represent reasonably the surface field variation, but not to characterise the equatorial rcs variation. Here the full equatorial monostatic result, employing 360 illumination angles, is obtained for ~4 times the cost of a single iterative solution. For this size of matrix, 6052 by 6052 (complex), we generally find a single iterative solution to cost rather more than an order of magnitude less than a single direct solution.

3.2 Near head-on rcs

Whilst it is conventional to analyse equatorial sweeps as above, practical interest may probably be concentrated on a relatively small solid angle centred around head-on, possibly biased towards illumination from slightly below. We have analysed this same 6 wavelength almond, with illumination in the range 0 to 16^0 vertically and 0 to 16^0 horizontally from head-on (where symmetry naturally makes only this one quadrant necessary). A uniform increment in each angular coordinate was used, with computations made with 16 by 16 (256) illumination angles, 8 by 8, 4 by 4, 3 by 3 and 1 (head on) illumination angle. The inset in figure 5 shows a representation of the rcs, plotted from the 256 illumination angle result. Figure 5 shows the variation in the number of iterations with the number of illumination angles, again exhibiting the characteristic 'plateau', here at about 7 times the number of iterations required for a single illumination angle.

3.3 Cost scaling

As noted earlier, there is no clear evidence regarding cost scaling for single illumination angle (single right hand side) solutions via iterative methods. The position is naturally made more complicated once multiple illumination angles are included. We can present here some empirical evidence, but only tentative observations and conclusions can be drawn.

As discussed elsewhere[8], fineness of discretisation (expressed in terms of degrees of freedom per incident wavelength) itself can influence the number of iterations required. We will here employ spherical scatterers of a range of sizes, as uniform discretisation is difficult to ensure on (say) the almond.

Whilst obviously it is machine dependant, some actual times might be helpful. All jobs were run on a 96Mb SGI Indy R5000. The matrix for the biggest mesh, 20440 (single precision complex) x 20440 in size, occupied 3.2 Gb. This was formed once, and read in for each iteration. Each such iteration took ~10 minutes, of which almost half was reading from disk. At the plateau of figure 6, mentioned below, some 680 iterations were required, corresponding to a time of 110 hours for the full characterisation.

We plot in figure 6 the number of iterations required versus illumination angles for a 180^0 scan of a series of spheres. These range in size from 1.6 wavelengths in diameter (1060 by 1060 matrix) to 7.12 wavelengths (20440 by 20440 matrix). They display behaviour qualitatively identical to that of the almond, with a flat plateau in number of iterations required being reached after only a modest number of illumination angles. Here, though, the monostatic rcs is of course characterised fully by a single illumination, whereas the computed surface field distribution is dependant on the illumination angle, lending further evidence to the observation that it is the latter which is the determinant of computational work.

In figure 7 we plot the variation of the number of iterations required at this plateau versus the diameter of the sphere in wavelengths. This number of iterations for full characterisation seems to rise roughly linearly with problem size. As the matrix size varies with the square of the body size or frequency, the cost of each iteration scales with

the fourth power of frequency, giving a total cost scaling in this particular case with about the fifth power of frequency.

4. Discussion and Conclusions

A very simple modification of the GCR approach has been shown to be effective for analysis of multiple right hand sides for the large, dense and unsymmetric matrices of multi-wavelength monostatic rcs calculations.

It seems possible to obtain solutions for a large number of different illumination angles for a modest multiple of the cost for a single right hand side. Typically, for the cases examined, this multiple is ~<10 for an essentially unlimited number of illumination angles.

As noted, the cost scaling of iterative solutions for single look angles is unclear. Starting from this point, we conclude from the present study:

- Costs for multiple look angles are independent of the number of look angles once more than some threshold number of look angles is considered.

- This threshold seems to be related to the number required to characterise the surface field, not the number required to characterise the rcs.

- How the number required to characterise the surface field varies with frequency is obviously geometry-dependant.

- On many geometries (eg the almond studied) the surface field needs far fewer angles to characterise it than does the rcs.

- The net result is that, for a body where this is true, monostatic rcs characterisation can be obtained for a small fraction of the cost of repeated *ab initio* iterative solution of the matrix equation.

References

1. Miller, E.K. A selective survey of computational electromagnetics. *IEEE Transactions on Antennas and Propagation* 36:pp1281-1305, (1988).

2. Peterson, A.F. and Mittra, R. Method of Conjugate Gradients for the Numerical Solution of Large-body Electromagnetic Scattering problems. *J.Opt.Soc.Am.* 2:pp971-977, (1985).

3. Peterson, A.F. and Mittra, R. Convergence of the Conjugate Gradient Method when Applied to Matrix Equations Representing Elctromagnetic Scattering Problems. *IEEE Transactions on Antennas and Propagation* 34:pp1447-1453, (1986).

4. Smith, C.F., Peterson, A.F. and Mittra, R. The Biconjugate Gradient Method for Electromagnetic Scattering. *IEEE Transactions on Antennas and Propagation* 38:pp938-940, (1990).

5. Soudais, P. Iterative solution of a 3-D scattering problem from arbitrary shaped multidielectric and multiconducting bodies. *IEEE Transactions on Antennas and Propagation* 42:pp954-959, (1994).

6. Smith, C.F., Peterson, A.F. and Mittra, R. The Treatment of Multiple excitations by Iterative methods for problems of Electromagnetic Scattering. *IEEE A-P Soc, Proc Int Symp Dig Blackburg, VA* pp530-533, (1987).

7. Woodworth, M.B. and Yaghjian, A.D. Multiwavelength three dimensional scattering with dual surface integral equations. *Journal of the Optical Society of America A* 11:pp1399-1413, (1994).

8. Pocock, M.D. and Walker, S.P. The complex bi-conjugate gradient solver applied to large electromagnetic scattering problems; computational costs, and cost scalings. *IEEE Transactions on Antennas and Propagation* 45:pp140-146, (1997).

9. Boyse, W.E. and Seidl, A.A. Mono-Static RCS Computation with a Block GMRES Iterative Solver. *Applied Computational Electromagnetics Society Journal* 11:pp63-69, (1996).

10. Eisenstat, S.C., Elman, H.C. and Schultz, M.H. Variational iteration methods for non symmetric systems of linear equations. *SIAM Journal of Numerical Analysis* 20:pp345-357, (1983).

11. Pocock, M.D and Walker, S.P. Radar cross section prediction using boundary integral equation methods with isoparametric quadratic surface modelling and iterative solvers. *Electromagnetics* 16:pp651-669, (1996).

12. Woo, A.C., Wang, H.T.G. and Schuh, M.J. Benchmark radar targets for the validation of computational electromagnetics programs. *IEEE Antennas and Propagation Magazine* 35:pp84-89, (1993).

13. Walker, S.P. and Lee, B.H. Termination criteria in iterative solution of large scattering problems using integral equation methods. *Comm in Numerical Methods in Engng* 13:pp199-206, (1997).

14. Schuh, A. and Woo, A. The Monostatic - bistatic rcs approximation. *IEEE Antennas and Propagation Magazine* 36:pp76-80, (1994).

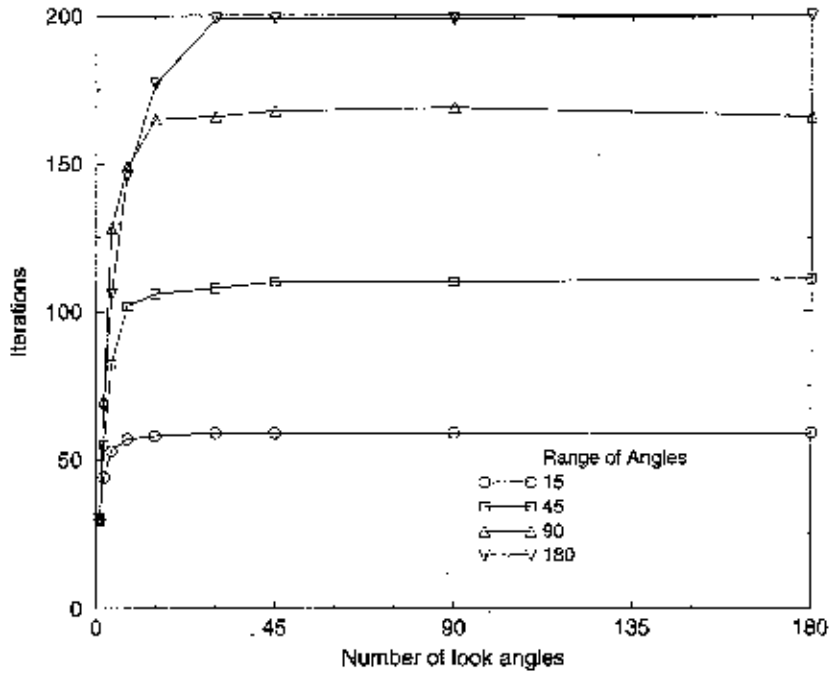

Figure 1

NASA almond, 7 GHz (6 wavelengths long); Iterations required versus number of illumination angles, with angular range spanned by the illumination angles, measured from head-on, as a parameter.

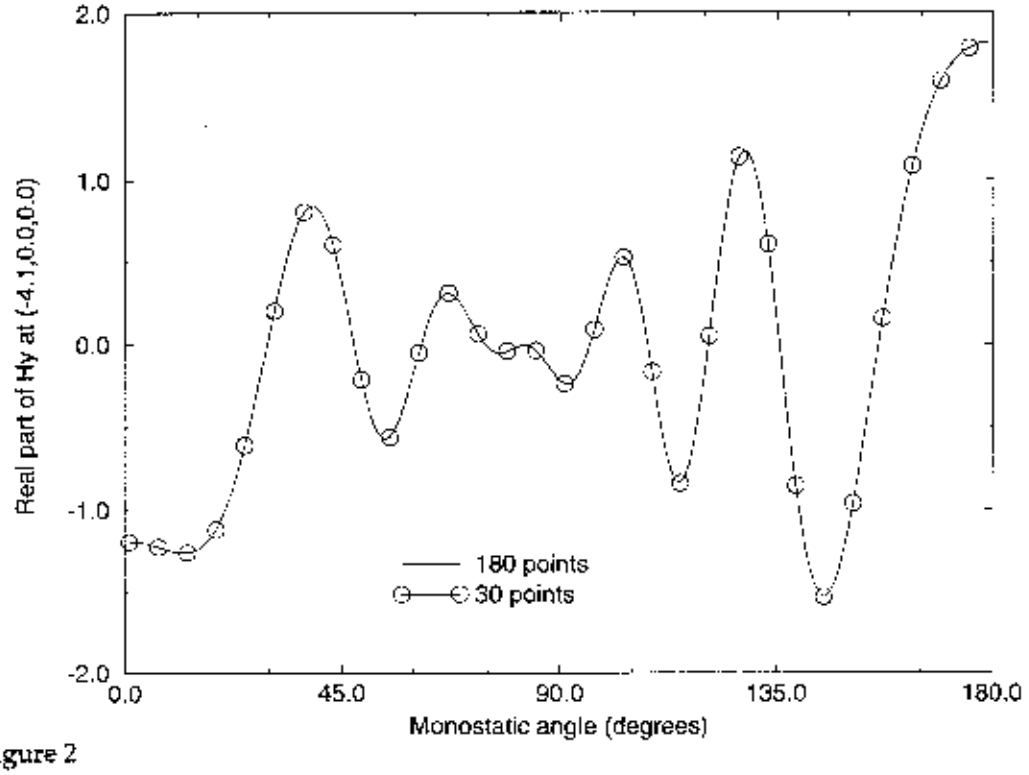

Figure 2

NASA almond (7 GHz, 6 wavelengths long): Real part of y-component of surface H field at surface location (-4.1", 0, 0) as a function of monostatic incident illumination angle. The illumination ranges from 0 to 180° in the equatorial (VV) plane (solid line), with results every 6 degrees (30 uniformly spaced illumination angles) additionally marked by circles.

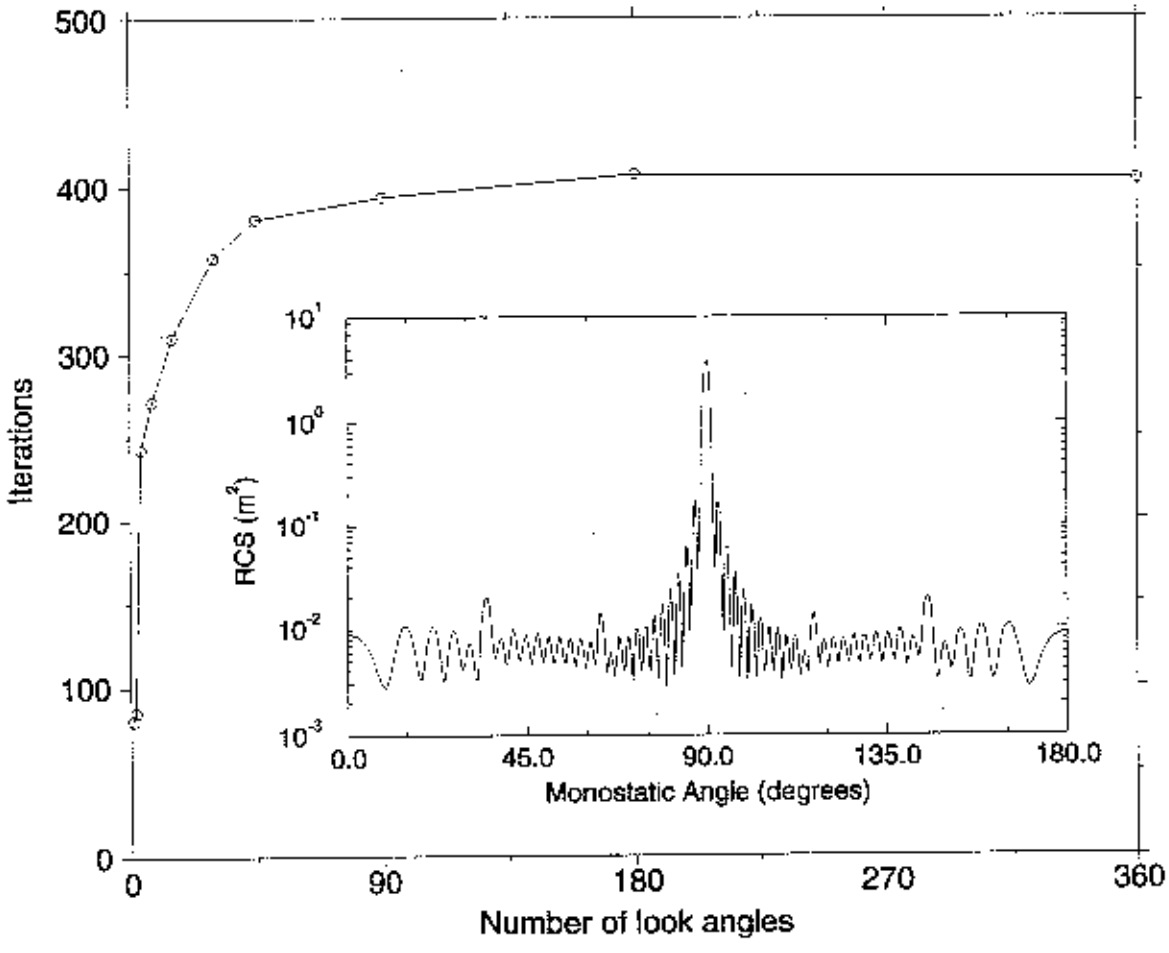

Figure 3

NASA almond (7 GHz, 6 wavelengths long): Monostatic rcs as a function of monostatic incident illumination angle. The illumination ranges from 0 to 180° in the equatorial (VV) plane (solid line), with results every 6 degrees (30 uniformly spaced illumination angles) additionally marked by solid circles. (Also shown are measured results, marked by open circles.)

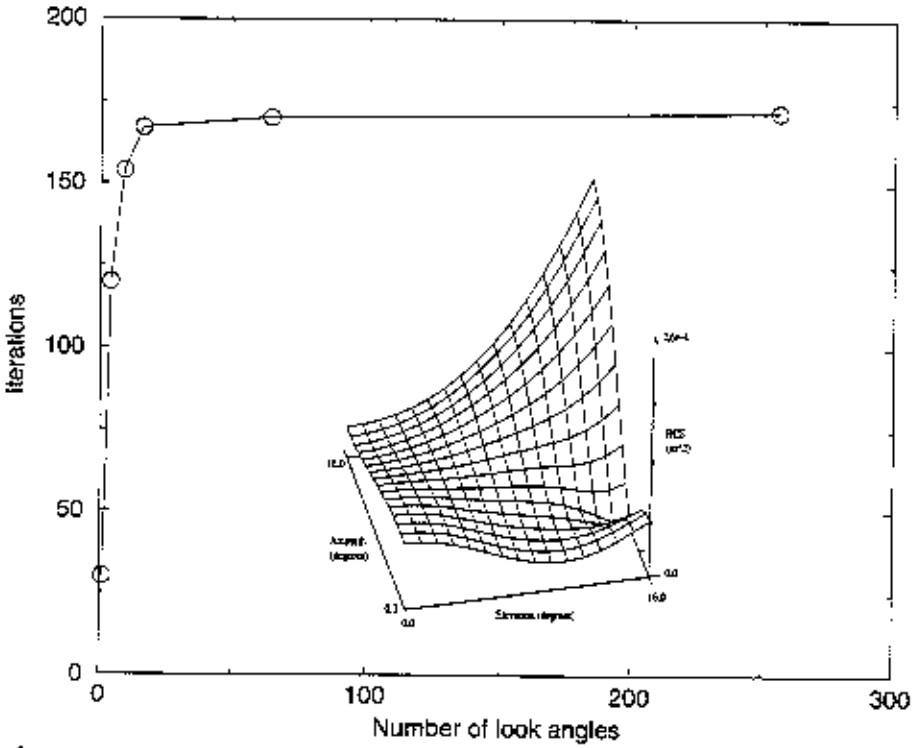

Figure 4

15 wavelength dipole: Iterations required versus number of illumination angles (distributed over 180° in all cases). Inset; Monostatic rcs.

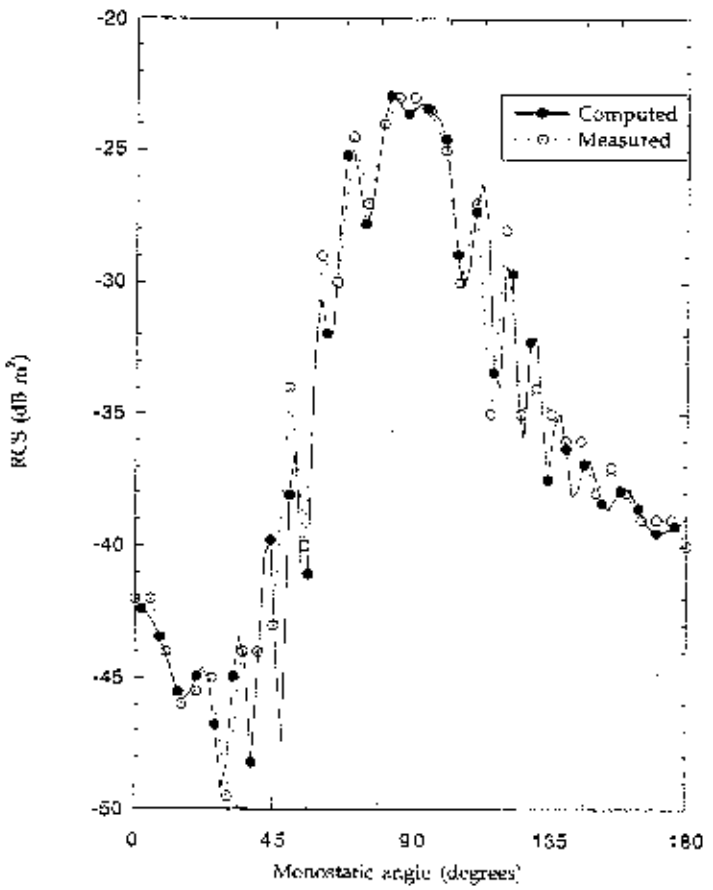

Figure 5

NASA almond (7 GHz, 6 wavelengths long): Iterations required versus number of illumination angles, with uniformly spaced illumination angles over the range 0 - 16° horizontally and vertically from head-on. Inset: Monostatic (VV) rcs.

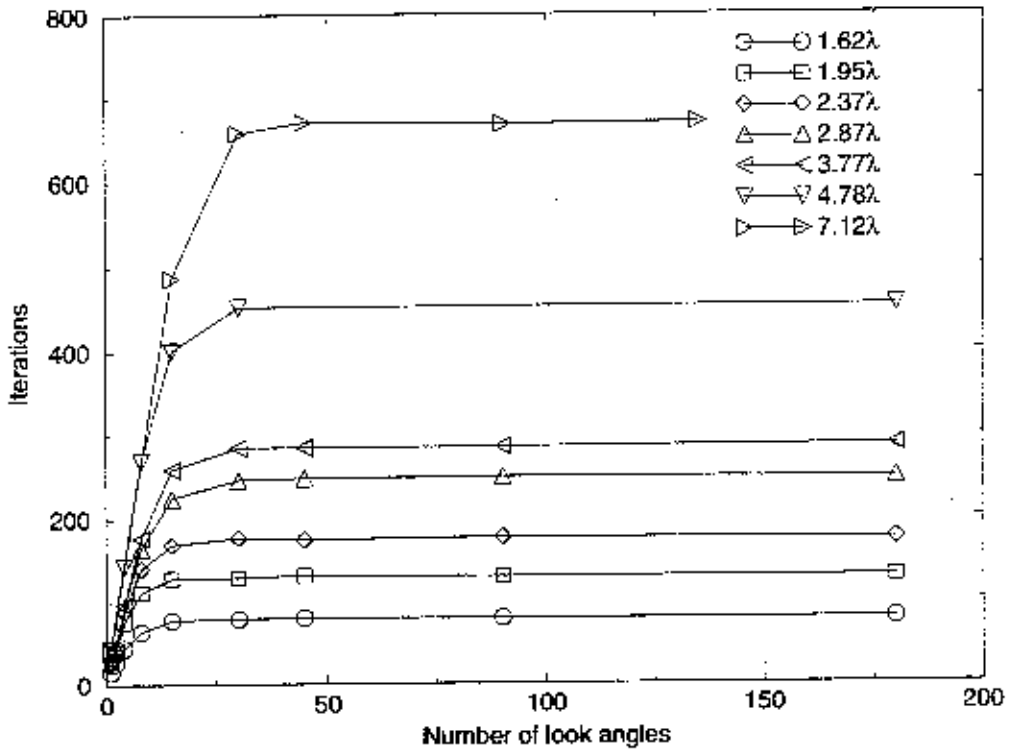

Figure 6

Iterations required versus number of illumination angles for a 180° scan of a series of spheres of diameters (in wavelengths) indicated.

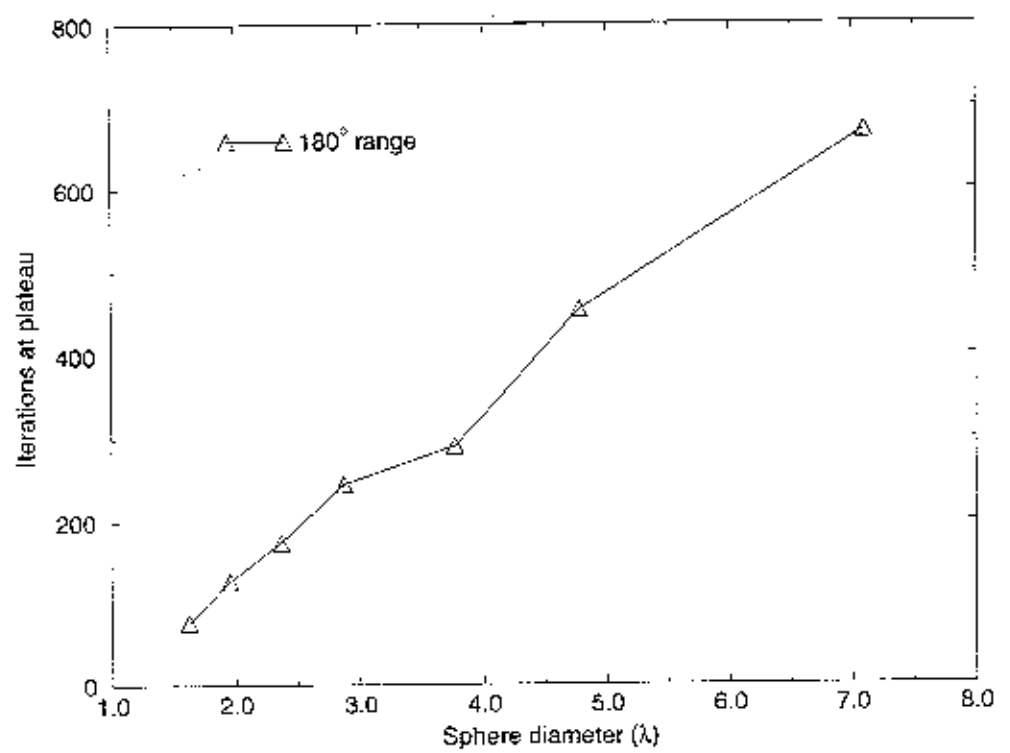

Figure 7

The 'plateau' number of iterations for full 180 degree characterisation, versus sphere diameter in wavelengths.