IMPROVED INTEGRAL EQUATION METHODS

FOR TRANSIENT WAVE SCATTERING

Byoung Hwa Lee.

A thesis submitted for the degree of Doctor of Philosophy

of the University of London and Diploma of Membership

of the Imperial College of Science, Technology and Medicine.

To my Mother and Father.

Abstract

As a consequence of the rapid development of computers, there is a strong interest in the solution of larger and larger scattering problems. In particular, radar cross section (RCS) problems which involve typically scattering from complex targets many wavelengths long are of considerable interest. Though there exist many numerical methods which are suitable for RCS problems, the boundary integral equation (BIE) method is one of the more efficient tools for their analysis, with costs scaling with the fifth power of the incident frequency. The algorithm modifications described here offer the prospect of large reductions in cost, and a possible reduction in cost scaling to the fourth power of the frequency. This has been achieved by the combination of a re-ordering of the normal BIE algorithm to a "project forward" form, and implementation of a modest further physical approximation which this re-ordering makes possible. This approach is demonstrated for a variety of geometries, including relatively difficult ones such as the NASA almond and cone-sphere and gap benchmarks.

The project forward algorithm has some advantages over the normal retrospective algorithm both in parallelisation and in hybridisation with time domain physical optics (POTD). The parallelisation and hybridisation of the project forward algorithm are developed in this thesis and their results are demonstrated. The results show that the project forward algorithm is very effective in such analysis.

Acknowledgements

There are many people whom I would like to thank for their help and support throughout the course of this research. Above all I am greatly indebted to Dr Simon Walker, for his guidance and encouragement throughout the research and the writing of this thesis.

I must also pay particular thanks to my colleagues Dr Michael Bluck and Dr Martin Pocock for their valuable advice in the areas of time domain and frequency domain electromagnetic work.

I would like to thank Steve Dodson, with whom I have had many valuable discussions.

Finally, thanks are due to the following people for their support and encouragement throughout my time here at Imperial College: Chi, Markku, Yoon Seok, Do Hyung, Young Kwan, Jin Seok, Tae Hyun and all other members of the Korean society.

Contents

List of Figures and Tables

Chapter 2.

Chapter 7.

Chapter 1 Introduction

1.1 Overview of the Thesis

This thesis is concerned with the application of the boundary integral equation (BIE) method to three dimensional transient electromagnetic (EM) wave scattering problems. The governing equations for these problems have been known for a long time. However, the solution for all but the simplest cases was impossible until the appearance of the first scientific mainframes in the early 1960's. With the approximately 10^6 times speed increase from the 10^3 floating-point operations per second (FLOP's)/sec of the UNIVAC-1 to the GigaFLOP speed of the present mainframes, the size and complexity of problems which can be solved have continuously increased.

As a consequence of such developments, there is nowadays a strong interest in the solution of larger and larger scattering problems. In particular, great attention has been paid to radar cross section (RCS) problems which involve typically scattering from complex targets many wavelengths long. For the solution of realistic problems, huge computational resources are required, and because of this it is often still impossible to solve these problems on present computers.

A variety of numerical methods have been investigated for the solution of these problems, and the BIE method has been proven to be one of the more efficient techniques, with its well known advantages of surface only discretisation, automatic satisfaction of the radiation boundary condition,

and so on. However, so as fully to exploit the computing power offered by available computers, existing algorithms must be re-examined with a focus on their efficiency. This is the motivation of this thesis.

Analysis of EM scattering by conventional BIE methods has computational costs which scale with the fifth power of the frequency. In this thesis, a novel modification to the normal (BIE) algorithm is presented. This modified algorithm is then combined with an approximation to the physics, involving in effect omitting integration over regions of low field magnitude. Whilst this does increase error slightly, it also provides large reductions in cost, and a reduction in cost scaling to the fourth power of the frequency. In the remainder of this thesis, the modified algorithm will be termed the 'project forward' algorithm to contrast it with the normal algorithm which is essentially retrospective.

Parallel computation offers a way of increasing computational power for large scattering problems. The parallelisation of the novel 'project forward' algorithm has been performed as a part of this work. Besides the 'project forward' algorithm being intrinsically faster in both serial and parallel form, parallelisation of the new algorithm has other advantages over the parallelisation of the normal BIE algorithm. It can be parallelised in a simpler way, and efficiently implemented on distributed memory parallel computer systems such as the CRAY T3D.

As briefly mentioned earlier, full field solutions for large scattering problems become prohibitively expensive for electrically large bodies. Fortunately, broadly 'optical' methods become accurate as larger bodies are

considered. However, often large bodies have significant features which are not electrically large, and here hybrid approaches are appropriate. In this thesis is presented a novel hybridisation of the 'project forward' algorithm with time domain physical optics (POTD). As with its parallelisation, the new algorithm can be hybridised in a simpler way than the normal BIE algorithm. Computations involving sample 'electrically large, small feature' problems will be presented later.

In the remainder of this chapter, the characteristics of the various numerical techniques suitable for wave scattering problems, and their computational costs, will be investigated. The algebraic derivation and numerical formulation of the normal BIE algorithm are described in chapter 2. Chapter 3 explains the physical basis of the potential cost saving, which leads to the need for the normal algorithm to be modified, and the applicability of the physical basis is investigated by showing some results obtained by simulating the new algorithm by trivial modifications to the conventional form. In chapter 4, the project forward algorithm is described in detail algebraically and numerically, and differences between it and the conventional algorithm are highlighted. The project forward algorithm is parallelised in chapter 5, demonstrating advantages over the parallelisation of the normal algorithm. Chapter 6 demonstrates the efficiency and accuracy of the modified algorithm for a variety of geometries on workstations and the CRAY T3D. In chapter 7, the project forward algorithm is hybridised with POTD, and some results are presented. Finally chapter 8 concludes this piece of work and suggests some associated areas for future study.

1.2 Large EM Wave Scattering Problems

The prediction of RCS response has been a great interest in many application areas, in particular military industries because it can be very important in the design of much military equipment. In spite of this importance, existing numerical methods, though extremely powerful tools, are not truly suitable for RCS prediction on many bodies of practical interest, because of their high computational costs.

Wave scattering problems can require very sophisticated modelling of a geometry; both if the geometry is intricate, and if the geometry is very smooth and 'stealthy'. Even if a geometry is simple, the field which also needs accurate modelling varies very rapidly if the body is many wavelengths long. In addition to these difficulties, typically RCS problems generally involve scattering from bodies many wavelengths long. Monostatic, or backscatter, cross section is the usual quantity of interest for most radar systems where the receiver and transmitter are collocated, and the number of 'look angles' required increases sharply as the bodies' electrical size increases.

There have been many efforts to apply a variety numerical methods to RCS problems. Loosely speaking, all works in this field are mainly concerned with obtaining accurate solutions as efficiently as possible (and indeed this present work is one of these). These will be reviewed in the next section.

1.3 Methods for EM Wave Scattering Problems

There exist many numerical methods for the prediction of the wave scattered from a body of arbitrary shape. They may be categorised broadly as optical (asymptotic or high frequency) methods, low frequency methods, and hybrid methods. Arguably, optical methods are not numerical methods because instead of solving the wave equations numerically, they embody physical approximations to the scattering behaviour. The main advantages of optical methods are that they can be reliably used for a body that is (very) large compared to the wavelength of the incident wave, and that their costs for obtaining solutions are much cheaper than those of other methods. However, their primary shortcoming is that they are only applicable for simple geometries, where length scales are significant multiples of a wavelength.

On the other hand, low frequency methods can be used for geometrically highly complex problems. They can be divided into differential equation methods which solve the differential equation directly by discretising the volume, and integral equation methods which first transform the differential equation to an integral equation before discretising the surface. The overwhelming disadvantage of both these classes of method is that existing computational resources limit the size of problem which can be solved.

Many problems combine both a large overall size, and important features which are not large compared to a wavelength; optical methods would be inaccurate, and low frequency methods too expensive. Hybrid methods are suitable for these problems, incorporating the optical methods and low

frequency methods in the region where each is most applicable. In the following sections, the relevant literature will be discussed in more detail.

1.3.1 Optical Methods

There are two basic optical methods; physical optics (PO) and geometrical optics (GO)[1]. The well known PO is the approximation of the surface current induced by the incident wave. The PO current $J_s = 2\hat{n} \times H_{inc}$ has long been used for evaluating the scattering properties of large objects. In the context of the magnetic field integral equation (MFIE), this solution results from completely neglecting mutual interaction effects, as represented by the integral term, on the illuminated portion of the scatterer, and further implies that in the shadow region the mutual interaction which occurs through the integral term completely cancels the incident field term. GO assumes that the incident and reflected rays are coplanar and the local angles of incidence and reflection are equal. Thus, in the absence of a specular point for a given scattering direction GO predicts that the field scattered in that direction is zero.

Some work has been done to extend the validity of PO and GO[2,3]. As a result, physical theory of diffraction (PTD) and geometrical theory of diffraction (GTD) have been developed by Ufimtsev and Keller respectively. Ufimtsev allows calculation of the induced surface currents which are due to various types of discontinuity. The extension of GO to GTD is based upon the incorporation of asymptotic evaluations of known, exact solutions of Maxwell's equations for a number of canonical problems, e.g., half-plane, wedge, cylinder, and sphere. The optical methods so far have been used

primarily for electrically large bodies, for which low frequency methods cannot be used because of their high computational costs.

1.3.2 Integral Equation Methods

Integral equation (IE) methods were the first to receive detailed development in the EM field, but mainly in the frequency domain. Time domain IE methods are probably least advanced, possibly because these are perceived to be the most difficult of all approaches. The main IE methods are the method of moment (MoM) and the boundary integral equation (BIE) methods. These are briefly reviewed below, and a more detailed investigation will be presented in the next chapter.

The MoM was introduced to electromagnetic problems by Harrington, and to acoustics by Banaugh and Goldsmith. Electromagnetics and acoustics are the main application areas of MoM. The frequency domain MoM may be the most widely used method[4,5,6], but recently time domain MoM has received increased attention[7,8].

The BIE approach starts from the same integral equation formulation as MoM, but it was developed by different researchers, and for different application areas. In particular, Brebbia, Banerjee and Butterfield contributed to the development of the current form of BIE[9,10]. The main application areas of BIE have been potential and stress analysis problems; BIE methods have had limited application to electromagnetic wave propagation problems. In the frequency domain, a good review is presented by Kress[11], and in the time domain a few works have been reported[12,13] recently.

Frequency domain IE methods have costs scaling with the sixth power of incident frequency, arising from the need for inversion of a dense matrix whose size scales with the square of frequency. The time domain IE cost scales with the fifth power, the reasons for which will be investigated in the next chapter. Although the scaling is more attractive than in the frequency domain, it is from a higher base. For modest sizes of scatterer (say below ~10 wavelengths long) the frequency domain approach is actually cheaper.

1.3.3 Differential Equation Methods

The two common differential equation methods are finite difference time domain (FDTD) and the finite element (FE) methods. In 1966, the first FDTD algorithm was introduced to the EM community by Yee[14]. He used a staggered grid mesh for the solutions of E and H fields for a perfectly conducting square illuminated by a pulse. After his work, many analyses have been reported[15]. FDTD is usually cast as an explicit time domain method. At each time step, the field at a point is calculated from nearby field values at the previous step. Matrix solution is thus not required, and this helps to enable the application of FDTD to electrically large problems.

Many EM problems are exterior, and this presents some difficulties in any differential equation method. It is necessary to terminate the computational solution domain at some finite distance from the body. The high cost of analyses provides a strong incentive to minimise this distance; this is counterbalanced by the difficulty of approximating sufficiently non-reflective boundary conditions[16] on this outer surface.

With discretisation at a constant number of nodes per wavelength, the costs of FDTD scale with frequency to the fourth power; three spatial dimensions, giving a cubic dependence per timestep, and with the number of timesteps required generally scaling with frequency. This is ostensibly the most benign scaling of the field solution methods. However, there is some evidence[17,18] that the fineness of discretisation itself must be increased as the frequency increases, to maintain a constant level of phase error. For the four dimensions involved, a modest dependence of discretisation on frequency can have significant effects, resulting in perhaps one or more additional power of frequency overall.

The large majority of FDTD work employs regular, 'sugar cube' meshes. When applied to curved geometries these can require 50 or more cells per wavelength to achieve acceptable accuracy[19]. There is some work on extending FDTD to curvilinear co-ordinate systems. The contour path (CP) method has been suggested by Jurgens *et al*[20], which can accurately model the curved surface with retaining the ability to model corners and edges. Recently conformal FDTD meshes have been implemented in two dimensional problems by Holland[21]. Possible savings over regular meshing by a factor of 256 are reported there when wavelength/10 cell sizes are employed instead of wavelength/50 .

The complexity of curvilinear FDTD anyway approaches that of finite elements, which have all the well known advantages of good surface modelling through curved elements. The use of FE for low frequency problems is well established[22,23], and there is now a growing interest in

applying FE to EM propagation problems[24]. In the time domain, like FDTD, FE is used as a 'marching-on-in-time' method. Like FDTD, the FE solution at a time step is generally an explicit combination of neighbouring nodal values. Timesteps of one eighth of the smallest element size have been used to overcome instability[25].

As in FDTD, FE requires a mesh truncation scheme when the domain of interest is exterior and infinite. This mesh truncation scheme is classified into two. One approach is an exact mesh truncation scheme which can reduce the computational domain. This involves in effect the use of a boundary integral formulation to model the infinite exterior domain surrounding the domain discretised by finite elements. However, it destroys the sparsity of the FEM matrix, and the BIE part is as usual dense. The other approach is an approximate truncation scheme. It needs an absorbing boundary condition (ABC) like FDTD and retains the sparsity of the FEM matrix, though it increases the computational domain to ensure acceptable level of accuracy. A recent review about this topic has been produced by Volakis *et al*[26].

1.3.4 Hybrid Methods

There have been many instances of hybridisation of optical and integral equation methods, especially MoM, since 1970's, and these works are reviewed by Medgyesi-Mitschang and Wang[27]. However, all this work has been done in frequency domain. There has been an increased interest in time domain methods of late. Very recently, as an exception, Walker and Vartiainen[28] combined the time domain BIE method with POTD. They

showed hybridisation to be as effective as in the frequency domain, whilst retaining all the other advantages of the time domain. As mentioned above, a related physical optics hybrid has been devised based on the 'project forward' algorithm developed here. This will be described more fully later.

Chapter 2 Time Domain Boundary Integral Equation Method for EM Wave Scattering Problems

As a precursor to describing the lower cost-scaling algorithm in chapter 4, the application of normal three dimensional BIE analysis of EM scattering from perfect electric conductors (PEC) is presented in this chapter. Section 2.1 gives a full derivation of the time domain magnetic field integral equation (MFIE), using the vector form of the Green's function. In section 2.2, some computational issues which arise in the governing integral equation are discussed, and some advantages of the present approach are explained. Numerical discretisation of a form of the MFIE which is suitable for computation is presented in section 2.3. Section 2.4 explains some features of pulsed excitation which can be dealt with only in the time domain. In section 2.5, the cost scaling of the BIE method in the time domain is discussed. Sections 2.6 and 2.7 respectively explain how to implement the methods and show some results. Finally, section 2.8 uses these results to discuss the present approaches.

2.1 Derivation of the Magnetic Field Integral Equation

The derivation of the MFIE is presented in various texts, but mainly in the frequency domain. Recently a detailed dyadic formulation in the time domain has been presented by Bluck[29]. There is an alternative approach based on the scalar Green function, 'vectorised' by multiplying it by an arbitrary vector[30]. The arbitrariness of this vector allows it to be eliminated

part way through the development. This vector formulation is presented here.

Whilst certainly the result, and most probably the derivation itself, are not novel, we are not aware of a published comprehensive derivation. Accordingly, it seemed desirable, and potentially helpful, to include the following fairly full treatment.

Derivation of the Non-Dimensionalised Vector Wave Equations:

The derivation starts from the differential form of the free space Maxwell's equations for a uniform perfect dielectric (source free) region,

$$\tilde{\nabla}' \times \tilde{E} + \frac{\partial \tilde{B}}{\partial \tilde{t}} = 0 \tag{2.1.1}$$

$$\tilde{\nabla}' \times \tilde{H} - \frac{\partial \tilde{D}}{\partial \tilde{t}} = 0 \tag{2.1.2}$$

$$\tilde{\nabla}' \cdot \tilde{B} = 0 \tag{2.1.3}$$

$$\tilde{\nabla}' \cdot \tilde{D} = 0 \tag{2.1.4}$$

together with the following associated constitutive equations,

$$\tilde{B} = \tilde{\mu}\tilde{H} \tag{2.1.5}$$

$$\tilde{D} = \tilde{\varepsilon}\tilde{E} \tag{2.1.6}$$

where the tilde ~ represents a dimensional quantity.

The first step involves taking the curl of equation (2.1.1) and (2.1.2), and then using Maxwell's equations again to give;

$$\tilde{\nabla}' \times \tilde{\nabla}' \times \tilde{E} + \tilde{\varepsilon}\tilde{\mu}\frac{\partial^2 \tilde{E}}{\partial \tilde{t}'^2} = 0 \qquad (2.1.7)$$

$$\tilde{\nabla}' \times \tilde{\nabla}' \times \tilde{H} + \tilde{\varepsilon}\tilde{\mu}\frac{\partial^2 \tilde{H}}{\partial \tilde{t}'^2} = 0 \qquad (2.1.8)$$

These are the familiar vector wave equations. By introducing a typical length scale, $\tilde{L}_0$ and a typical potential difference, $\tilde{V}_0$, the vector wave equations are now non-dimensionalised as follows:

$$\mathbf{r}' = \frac{\tilde{\mathbf{r}}}{\tilde{L}_0} \qquad (2.1.9)$$

$$t' = \frac{\tilde{t}}{\tilde{L}_0\sqrt{\tilde{\varepsilon}\tilde{\mu}}} \qquad (2.1.10)$$

$$\mathbf{E} = \tilde{\mathbf{E}}\frac{\tilde{L}_0}{\tilde{V}_0} \qquad (2.1.11)$$

$$\mathbf{H} = \tilde{\mathbf{H}}\frac{\tilde{L}_0}{\tilde{V}_0}\sqrt{\frac{\tilde{\mu}}{\tilde{\varepsilon}}} \qquad (2.1.12)$$

The non-dimensionalised vector wave equations may now be written in terms of the above non-dimensional quantities, showing their dependence on the position $\mathbf{r}'$ and time t'.

$$\nabla' \times \nabla' \times \mathbf{E}(\mathbf{r}',t') + \frac{\partial^2 \mathbf{E}(\mathbf{r}',t')}{\partial t'^2} = 0 \qquad (2.1.13)$$

$$\nabla' \times \nabla' \times H(r',t') + \frac{\partial^2 H(r',t')}{\partial t'^2} = 0 \qquad (2.1.14)$$

The electric field integral equation (EFIE) and the MFIE are the integral forms of equations (2.1.13) and (2.1.14) respectively. The following sections will develop these, with emphasis (where the derivations diverge) on the MFIE.

Preliminaries for the MFIE Derivation :

Figure 2.1.1 shows the domain Ω with locations r and r'. Other notation will be introduced as the derivation progresses.

Consider first the 'point source' inhomogeneous, infinite domain, vector wave equation

$$\nabla' \times \nabla' \times G(r,t:r',t') - \nabla'(\nabla' \cdot G) + \frac{\partial^2 G}{\partial t^2} = a\delta(t-t')\delta(r-r') \qquad (2.1.15)$$

Here G at position r' and time t' is 'caused' by (is the field due to) a unit impulsive vector source at position r and time t. The vector a is uniform and constant, but otherwise arbitrary.

The solution to (2.1.15) can be shown to be:

$$G(r,t:r',t') = a\frac{1}{4\pi}\frac{1}{|r-r'|}\delta(t-t'-|r-r'|) = aG \qquad (2.1.16)$$

Note that in (2.1.16) we have, on the extreme right, expressed G as the product of the vector a and the scalar G. This 'vector times scalar' form will be found convenient later. Note also the single dirac delta in (2.1.16), with

spatial and temporal arguments, as opposed to the two dirac deltas in (2.1.15) each with only a spatial or a temporal argument.

Now to transform the differential equations (2.1.13) and (2.1.14) into integral equations involving only the surface of the domain, the vector form of Greens second identity[31] will be used for two suitably smooth and continuous vectors $\mathbf{p}$ and $\mathbf{q}$;

$$\int_{\Omega} \mathbf{p} \cdot \nabla' \times \nabla' \times \mathbf{q} - \mathbf{q} \cdot \nabla' \times \nabla' \times \mathbf{p} \, dv' = \int_{\partial\Omega} \mathbf{n}' \cdot \left[(\nabla' \times \mathbf{q}) \times \mathbf{p} - (\nabla' \times \mathbf{p}) \times \mathbf{q} \right] ds' \qquad (2.1.17)$$

where v' and s' refer respectively to the volume and surface of the domain in the primed co-ordinates, $\mathbf{p}$ is taken to mean $\mathbf{p(r')}$ and $\mathbf{q}$ similarly, and $\mathbf{n'}$ is the unit outward pointing normal at $\mathbf{r'}$.

Firstly $\mathbf{q}$ is replaced with $\mathbf{G}$, and (2.1.17) is integrated from $t'=0$ to $t'=t$, the present time. Since $\mathbf{G}$ is not smooth enough to satisfy the requirements of Greens second identity, we must exclude a small (spherical) region around $\mathbf{r}$. The excluded region and its surface are denoted by the subscript e. Equation (2.2.17) thus can be written as;

$$\int_{t'=0}^{t'=t} \int_{\Omega-\Omega_e} \mathbf{p} \cdot \nabla' \times \nabla' \times \mathbf{G} - \mathbf{G} \cdot \nabla' \times \nabla' \times \mathbf{p} \, dv' dt'$$

$$= \int_{t'=0}^{t'=t} \int_{\partial\Omega+\partial\Omega_e} \mathbf{n}' \cdot \left[(\nabla' \times \mathbf{G}) \times \mathbf{p} - (\nabla' \times \mathbf{p}) \times \mathbf{G} \right] ds' dt' \qquad (2.1.18)$$

Recall that $\mathbf{G}$ embodies an arbitrary vector $\mathbf{a}$. Now we will make re-arrangements such that every term is dotted with $\mathbf{a}$, and $\mathbf{a}$ appears nowhere else, allowing $\mathbf{a}$ to be eliminated.

The surface integral of (2.1.18) is now considered. In the first term, $\nabla' \times \mathbf{G}$ can be written as;

$$\nabla' \times \mathbf{G} = \nabla' \times \mathbf{a}G = \nabla'G \times \mathbf{a} + G\nabla' \times \mathbf{a} = \nabla'G \times \mathbf{a} \qquad (2.1.19)$$

since the curl of the uniform $\mathbf{a}$ is naturally zero. Using a standard vector identiy, the first term of the surface integrand can thus be written as;

$$\mathbf{n}' \cdot [(\nabla' \times \mathbf{G}) \times \mathbf{p}] = \mathbf{n}' \cdot [(\nabla'G \times \mathbf{a}) \times \mathbf{p}] = \mathbf{a} \cdot (\mathbf{p} \times \mathbf{n}') \times \nabla'G \qquad (2.1.20)$$

where the a-dot is now extracted as required.

The second term of the surface intergal of (2.1.18), using again a standard vector identity, can be written as;

$$\mathbf{n}' \cdot (\nabla' \times \mathbf{p}) \times \mathbf{G} = \mathbf{G} \cdot \mathbf{n}' \times (\nabla' \times \mathbf{p}) = \mathbf{a} \cdot G\mathbf{n}' \times (\nabla' \times \mathbf{p}) \qquad (2.1.21)$$

and again a-dot has been extracted.

The entire surface integral I_s can thus be written as;

$$I_s = \mathbf{a} \cdot \int_{t'=0}^{t'=t} \int_{\partial\Omega + \partial\Omega_e} (\mathbf{p} \times \mathbf{n}') \times \nabla'G - G\mathbf{n}' \times (\nabla' \times \mathbf{p}) ds' dt' \qquad (2.1.22)$$

where we have also taken the $\mathbf{a}$ outside the integral. The volume integral of (2.1.18) will now be manipulated. First $\nabla' \times \nabla' \times \mathbf{G}$ in (2.1.18) is replaced with terms from (2.1.15).

$$\nabla' \times \nabla' \times \mathbf{G}(\mathbf{r}, t : \mathbf{r}', t') = \nabla'(\nabla' \cdot \mathbf{G}) - \frac{\partial^2 \mathbf{G}}{\partial t^2} = \nabla'(\nabla' \cdot \mathbf{G}) - \mathbf{a}\frac{\partial^2 G}{\partial t^2} \qquad (2.1.23)$$

Note that the right hand side of (2.1.15) is zero except inside our excluded region, and that the time derivatives of the constant vector **a** are zero. Insertion of this will generate a term $\mathbf{p} \cdot \nabla'(\nabla' \cdot \mathbf{G})$, which is now manipulated to extract a-dot.

Now using $\nabla \cdot (\phi \mathbf{A}) = \phi \nabla \cdot \mathbf{A} + \nabla \phi \cdot \mathbf{A}$, where $(\nabla' \cdot \mathbf{G}) = \phi$, $\mathbf{p} \cdot \nabla'(\nabla' \cdot \mathbf{G})$ can be written as;

$$\mathbf{p} \cdot \nabla'(\nabla' \cdot \mathbf{G}) = \nabla' \cdot ((\nabla' \cdot \mathbf{G})\mathbf{p}) - (\nabla' \cdot \mathbf{G})\nabla' \cdot \mathbf{p} \qquad (2.1.24)$$

The vector **G** is also replaced with aG in (2.1.24) :

$$\nabla' \cdot \mathbf{G} = \nabla' \cdot (\mathbf{a}G) = G\nabla' \cdot \mathbf{a} + \nabla'G \cdot \mathbf{a} = \nabla'G \cdot \mathbf{a} \qquad (2.1.25)$$

where the zero divergence of the constant vector **a** is used. Putting this into (2.1.24) gives:

$$\begin{aligned} \mathbf{p} \cdot \nabla'(\nabla' \cdot \mathbf{G}) &= \nabla \cdot ((\nabla'G \cdot \mathbf{a})\mathbf{p}) - (\nabla'G \cdot \mathbf{a})\nabla' \cdot \mathbf{p} \\ &= \nabla \cdot ((\nabla'G \cdot \mathbf{a})\mathbf{p}) - \mathbf{a} \cdot \nabla'G(\nabla' \cdot \mathbf{p}) \end{aligned} \qquad (2.1.26)$$

The second term on the right of (2.1.26) has the a-dot extracted as required.

For the first term on the right of (2.1.26) the volume integral of the divergence can be transformed to a surface integral, allowing the a-dot to be extracted as required:

$$\begin{aligned} I &= \int_{t'=0}^{t'=t} \int_{\Omega-\Omega_e} \nabla \cdot ((\nabla'G \cdot \mathbf{a})\mathbf{p}) dv' dt' = \int_{t'=0}^{t'=t} \int_{\partial\Omega+\partial\Omega_e} ((\nabla'G \cdot \mathbf{a})\mathbf{p}) \cdot \mathbf{n}' ds' dt' \\ &= \mathbf{a} \cdot \int_{t'=0}^{t'=t} \int_{\partial\Omega+\partial\Omega_e} \nabla'G(\mathbf{p} \cdot \mathbf{n}') ds' dt' \end{aligned} \qquad (2.1.27)$$

Thus the volume integral of $p \cdot \nabla'(\nabla' \cdot G)$ has been converted into:

$$\int_{t'=0}^{t'=t} \int_{\Omega-\Omega_\epsilon} p \cdot \nabla'(\nabla' \cdot G) dv' dt'$$

$$= -a \cdot \int_{t'=0}^{t'=t} \int_{\Omega-\Omega_\epsilon} \nabla' G(\nabla' \cdot p) dv' dt' + a \cdot \int_{t'=0}^{t'=t} \int_{\partial\Omega+\partial\Omega_\epsilon} \nabla' G(p \cdot n') ds' dt' \qquad (2.1.28)$$

The volume integral I_v of (2.1.18) now can be written with the arbitrary vector extracted as

$$I_v = -a \cdot \int_{t'=0}^{t'=t} \int_{\Omega-\Omega_\epsilon} \frac{\partial^2 G}{\partial t^2} p + G\nabla' \times \nabla' \times p + \nabla' G(\nabla' \cdot p) dv' dt' + a \cdot \int_{t'=0}^{t'=t} \int_{\partial\Omega+\partial\Omega_\epsilon} \nabla' G(p \cdot n') ds' dt'$$

$$\qquad (2.1.29)$$

The arbitrary vector a dots everything, and so can be dropped. Doing so, and putting the surface and volume integrals (2.1.29) and (2.1.22) back together, finally equation (2.1.18) can be written as;

$$\int_{t'=0}^{t'=t} \int_{\Omega-\Omega_\epsilon} \frac{\partial^2 G}{\partial t^2} p + G\nabla' \times \nabla' \times p - \nabla' G(\nabla' \cdot p) dv' dt' =$$

$$\int_{t'=0}^{t'=t} \int_{\partial\Omega+\partial\Omega_\epsilon} (n' \times p) \times \nabla' G + Gn' \times (\nabla' \times p) + \nabla' G(p \cdot n') ds' dt' \qquad (2.1.30)$$

Now in effect one of the two arbitrary vectors, q, in the Greens identity has been replaced with one, G, which is useful. Now p can be replaced with either E or H, depending on whether it is the MFIE or the EFIE which we seek to derive. In this thesis, only the MFIE will be considered.

Derivation of the MFIE :

Now vector p is replaced by H, and we immediately use the wave equation (2.1.13) to replace $\nabla' \times \nabla' \times H$ with the time derivative:

$$\int_{t'=0}^{t'=t} \int_{\Omega-\Omega_e} \frac{\partial^2 G}{\partial t^2} \mathbf{H} - G \frac{\partial^2 \mathbf{H}}{\partial t^2} - \nabla'G(\nabla' \cdot \mathbf{H}) dv' \, dt'$$

$$= \int_{t'=0}^{t'=t} \int_{\partial\Omega+\partial\Omega_e} G\mathbf{n}' \times (\nabla' \times \mathbf{H}) + (\mathbf{n}' \times \mathbf{H}) \times \nabla'G + \nabla'G(\mathbf{H} \cdot \mathbf{n}') ds' \, dt' \qquad (2.1.31)$$

The first term on the left hand side is easily integrable with respect time, since it is a perfect differential. Thus (2.1.31) becomes

$$\int_{\Omega-\Omega_e} \left[\mathbf{H} \frac{dG}{dt} - G \frac{d\mathbf{H}}{dt} \right]_{t'=0}^{t'=t} dv' - \int_{t'=0}^{t'=t} \int_{\Omega-\Omega_e} \nabla'G(\nabla' \cdot \mathbf{H}) dv' \, dt'$$

$$= \int_{t'=0}^{t'=t} \int_{\partial\Omega+\partial\Omega_e} G\mathbf{n}' \times (\nabla' \times \mathbf{H}) + (\mathbf{n}' \times \mathbf{H}) \times \nabla'G + \nabla'G(\mathbf{H} \cdot \mathbf{n}') ds' \, dt' \qquad (2.1.32)$$

At $t'=t$ the only place $\mathbf{G}$ is non-zero is at $\mathbf{r}'=\mathbf{r}$, by virtue of (2.1.16). Similarly, the only location where it would sift out the derivative of $\mathbf{H}$ is at $\mathbf{r}'=\mathbf{r}$ also. Since this location is excluded, the upper limit is zero.

The lower limit is an integral over the domain of (some function of) the initial conditions; $\mathbf{H}$ and its rate of change. It is via this that the contribution of the initial conditions to the eventual solution is included.

The form and value of the initial conditions is naturally problem dependent. For now we will simply write the volume integral of the initial conditions as some vector $\mathbf{W}(\mathbf{r})$. We will come back to this point later.

Equation (2.1.32) thus becomes

$$\mathbf{W} - \int_{t'=0}^{t'=t} \int_{\Omega-\Omega_e} \nabla'G(\nabla' \cdot \mathbf{H}) dv' \, dt'$$

$$= \int_{t'=0}^{t'=t} \int_{\partial\Omega+\partial\Omega_e} G\mathbf{n}' \times (\nabla' \times \mathbf{H}) + (\mathbf{n}' \times \mathbf{H}) \times \nabla'G + \nabla'G(\mathbf{H} \cdot \mathbf{n}') ds' \, dt' \qquad (2.1.33)$$

In equation (2.1.33) we have the term $\nabla'G$ to evaluate. Denoting $r-r'$ as $\mathbf{R}$, and $|\mathbf{R}|$ as R, $\nabla'G$ can be written as;

$$\nabla'G = \frac{1}{4\pi}\left(\frac{-\hat{\mathbf{R}}}{R^2}\delta(t-t'-R)+\frac{\hat{\mathbf{R}}}{R}\frac{\partial}{\partial R}\delta(t-t'-R)\right)$$

(2.1.34)

We now need to differentiate the delta function. As usual, this has no meaning till under an integral, but is of the form $\int\nabla\delta(x_0)f(x)dx = -\nabla f_{x=x_0}$; i.e. where the delta function sifts the value of the associated multiplying function, the derivative of the delta function sifts minus its derivative with respect to the variable of integration.

Integration over the Surface of the Excluded Region

Here we will perform only the surface integrations of (2.1.33), over only the surface of the excluded region. This will extract the 'free term'; the value of the field at the location r. This value will then be seen to be given by the sum of the remaining (volume and physical surface) integrals of (2.1.33).

Inserting the particular form of G and noting that for the excluded sphere $n'=\mathbf{R}/|\mathbf{R}|$, for the first term of the surface integral of (2.1.33) we have

$$\mathbf{I}_{s1} = \frac{1}{4\pi}\int_{t'=0}^{t'=t}\int_{\partial\Omega_e}\frac{1}{R}\delta(t-t'-R)n'\times(\nabla'\times\mathbf{H})ds'dt'$$

(2.1.35)

Integrating (2.1.35) through time (remembering that R is small, and shortly will be caused to tend to zero) sifts the present, t, value:

$$I_{s1} = \frac{1}{4\pi} \int_{\partial\Omega_e} \frac{1}{R} \mathbf{n}' \times (\nabla' \times \mathbf{H}) ds' \tag{2.1.36}$$

Writing $ds' = R^2 d\theta$, with θ the solid angle the elemental surface subtends at $\mathbf{r}$, this becomes

$$I_{s1} = \frac{1}{4\pi} \int_{4\pi} \frac{1}{R} \mathbf{n}' \times (\nabla' \times \mathbf{H}) R^2 d\theta \tag{2.1.37}$$

The limit of this as R shrinks to zero is zero. (However $\mathbf{n}'$ and $\nabla' \times \mathbf{H}$ vary with θ they remain finite, and are killed off by the net R term.)

For the second term of the surface integral of (2.1.33) we have

$$I_{s2} = \frac{1}{4\pi} \int_{t'=0}^{t'=t} \int_{\partial\Omega_e} (\hat{\mathbf{R}} \times \mathbf{H}) \times \left(\frac{-\hat{\mathbf{R}}}{R^2} \delta(t - t' - R) + \frac{\hat{\mathbf{R}}}{R} \frac{\partial}{\partial R} \delta(t - t' - R) \right) ds' dt' \tag{2.1.38}$$

where we have inserted $\nabla'G$ from (2.1.34) and have replaced $\mathbf{n}'$ with (unit) $\mathbf{R}$.

For the first term of (2.1.38) we have

$$I_{s2(i)} = \frac{1}{4\pi} \int_{t'=0}^{t'-t} \int_{\partial\Omega_e} (\hat{\mathbf{R}} \times \mathbf{H}) \times \left(\frac{-\hat{\mathbf{R}}}{R^2} \delta(t - t' - R) \right) ds' dt' \tag{2.1.39}$$

Integrating through time again sifts the present, t, value. Doing this, and simultaneously changing to an integration with respect to the solid angle, we obtain

$$I_{s2(i)} = \frac{1}{4\pi} \int_{4\pi} (\hat{\mathbf{R}} \times \mathbf{H}) \times \frac{-\hat{\mathbf{R}}}{R^2} R^2 d\theta = -\frac{1}{4\pi} \int_{4\pi} (\hat{\mathbf{R}} \times \mathbf{H}) \times \hat{\mathbf{R}} d\theta \tag{2.1.40}$$

This is the first non-zero integral we have found, and we will come back to this shortly.

Taking now the second term of (2.1.38), we have

$$I_{s2(ii)} = \frac{1}{4\pi} \int_{t'=0}^{t'=t} \int_{\partial\Omega_\epsilon} (\hat{\mathbf{R}} \times \mathbf{H}) \times \left(\frac{\hat{\mathbf{R}}}{R} \frac{\partial}{\partial t'} \delta(t - t' - R) \right) ds' dt' \tag{2.1.41}$$

where we have made use of the fact that $\dfrac{\partial}{\partial R} = \dfrac{\partial}{\partial t'}$ for the delta function here. When we integrate (2.1.41) through time this sifts minus the time derivative, here at time t, of the function it multiplies, giving

$$I_{s2(ii)} = \frac{1}{4\pi} \int_{\partial\Omega_\epsilon} -\left(\hat{\mathbf{R}} \times \frac{\partial \mathbf{H}}{\partial t'} \right) \times \left(\frac{\hat{\mathbf{R}}}{R} \right) ds' \tag{2.1.42}$$

Again writing $ds' = R^2 d\theta$, with θ the solid angle the elemental surface subtends at r, (2.1.42) becomes

$$I_{s2(ii)} = \frac{1}{4\pi} \int_{4\pi} -\left(\hat{\mathbf{R}} \times \frac{\partial \mathbf{H}}{\partial t'} \right) \times \left(\frac{\hat{\mathbf{R}}}{R} \right) R^2 d\theta \tag{2.1.43}$$

which again is zero as R shrinks to zero.

For the third term of the surface integral of (2.1.33) we have

$$I_{s3} = \frac{1}{4\pi} \int_{t'=0}^{t'=t} \int_{\partial\Omega_\epsilon} \left(\frac{-\hat{\mathbf{R}}}{R^2} \delta(t - t' - R) + \frac{\hat{\mathbf{R}}}{R} \frac{\partial}{\partial R} \delta(t - t' - R) \right) (\mathbf{H} \cdot \hat{\mathbf{R}}) ds' dt' \tag{2.1.44}$$

where we have inserted $\nabla'G$ from (2.1.34) and have replaced n' with (unit) R.

Integrating through time, using again the fact that $\dfrac{\partial}{\partial R} = \dfrac{\partial}{\partial t'}$ for the delta function here, and moving to the solid angle integration, gives

$$I_{s3} = \frac{1}{4\pi}\int_{4\pi}\left(\frac{-\hat{R}}{R^2}(H\cdot\hat{R}) - \frac{\hat{R}}{R}\left(\frac{\partial H}{\partial t'}\cdot\hat{R}\right)\right)R^2 d\theta = \frac{1}{4\pi}\int_{4\pi}-\hat{R}(H\cdot\hat{R})d\theta \qquad (2.1.45)$$

where we see the approach of R to zero eliminates the second term, but leaves the first.

The only non-zero terms we were left with were in (2.1.40) and (2.1.45). Gathering these (and changing the order to remove one minus sign) we have

$$I = \frac{1}{4\pi}\int_{4\pi}\hat{R}\times(\hat{R}\times H) - \hat{R}(H\cdot\hat{R})d\theta \qquad (2.1.46)$$

Using a standard vector identity, we have

$$I = \frac{1}{4\pi}\int_{4\pi}\hat{R}(\hat{R}\cdot H) - H(\hat{R}\cdot\hat{R}) - \hat{R}(H\cdot\hat{R})d\theta = \frac{1}{4\pi}\int_{4\pi}-H d\theta = -H(r,t) \qquad (2.1.47)$$

Integration over the Physical Volume

We now evaluate the volume integral term of (2.1.33), inserting $\nabla'G$ from (2.1.34):

$$I_v = \frac{-1}{4\pi}\int_{t'=0}^{t'=t}\int_{\Omega-\Omega_\epsilon}\left(\frac{-\hat{R}}{R^2}\delta(t-t'-R) + \frac{\hat{R}}{R}\frac{\partial}{\partial R}\delta(t-t'-R)\right)(\nabla'\cdot H)dv' dt' \qquad (2.1.48)$$

Changing the differentiation variable of the delta function, and performing the time integration (and thus sifting):

$$I_v = \frac{1}{4\pi}\int_\Omega \frac{\hat{R}}{R^2}(\nabla'\cdot H^*) + \frac{\hat{R}}{R}\left(\nabla'\cdot\frac{\partial H^*}{\partial t'}\right)dv' \tag{2.1.49}$$

where H^* is the field at the retarded time, and we have also ceased to exclude the spherical region; as its radius shrinks to zero, the $R^2dR\,d\theta$ of the integral keeps everything perfectly well behaved.

We have a volume integration over (almost) the whole domain of the divergence of the field (and its time derivative), at the retarded time as seen from **r**. There is no mathematical reason for this to be zero; the limiting process which eliminated the surface integrals obviously does not apply here, but we will later have recourse to the physical observation that for the equations we are trying to solve the divergence of **H** is indeed zero. The volume integral thus will disappear.

Integration over the Physical Surface

The first term of the surface integral of (2.1.33) generates, using (2.1.34):

$$I_{s1} = \frac{1}{4\pi}\int_{t'=0}^{t'=t}\int_{\partial\Omega}\left(\frac{1}{R}\delta(t-t'-R)\right)n'\times(\nabla'\times H)ds'\,dt' \tag{2.1.50}$$

and integration through time yields

$$I_{s1} = \frac{1}{4\pi}\int_{\partial\Omega}\frac{1}{R}n'\times(\nabla'\times H^*)ds' \tag{2.1.51}$$

where the delta function has sifted the (curl of the) field value at the retarded time.

The second term of the surface integral of (2.1.33) generates, using again (2.1.34):

$$I_{s2} = \frac{1}{4\pi} \int_{t'=0}^{t'=t} \int_{\partial\Omega} (n' \times H) \times \left(\frac{-\hat{R}}{R^2} \delta(t - t' - R) + \frac{\hat{R}}{R} \frac{\partial}{\partial R} \delta(t - t' - R) \right) ds' dt' \qquad (2.1.52)$$

Considering the first term of this and performing the time integration yields

$$I_{s2(i)} = \frac{1}{4\pi} \int_{\partial\Omega} (n' \times H^*) \times \frac{-\hat{R}}{R^2} ds' \qquad (2.1.53)$$

where we have sifted (n' cross) the field value at the retarded time.

Turning now to the second term of (2.1.52) we have

$$I_{s2(ii)} = \frac{1}{4\pi} \int_{t'=0}^{t'=t} \int_{\partial\Omega} (n' \times H) \times \left(\frac{\hat{R}}{R} \frac{\partial}{\partial R} \delta(t - t' - R) \right) ds' dt' \qquad (2.1.54)$$

As before, using $\dfrac{\partial}{\partial R} = \dfrac{\partial}{\partial t'}$ and performing the time integration we then sift once more, leaving

$$I_{s2(i)} = \frac{1}{4\pi} \int_{\partial\Omega} -\left(n' \times \frac{\partial H^*}{\partial t'} \right) \times \frac{\hat{R}}{R} ds' \qquad (2.1.55)$$

where we have sifted the time derivative of (n' cross) the field value at the retarded time.

The third term of the surface integral of (2.1.33) generates, using (2.1.34):

$$I_{s3} = \frac{1}{4\pi} \int_{t'=0}^{t'=t} \int_{\partial\Omega} \left(\frac{-\hat{R}}{R^2} \delta(t - t' - R) + \frac{\hat{R}}{R} \frac{\partial}{\partial R} \delta(t - t' - R) \right) (H \cdot n') ds' dt' \qquad (2.1.56)$$

The first term of (2.1.56) integrates and sifts to give

$$I_{s3(i)} = \frac{1}{4\pi} \int_{\partial\Omega} \frac{-\hat{R}}{R^2} (H^* \cdot n') ds'$$

(2.1.57)

Like the first term, the second term can be written as:

$$I_{s3(ii)} = \frac{1}{4\pi} \int_{\partial\Omega} -\frac{\hat{R}}{R} \left(\frac{\partial H^*}{\partial t'} \cdot n' \right) ds'$$

(2.1.58)

The contributions are (2.1.51), (2.1.53), (2.1.55), (2.1.57) and (2.1.58). Putting these together we have

$$I_s = \frac{1}{4\pi} \int_{\partial\Omega} \frac{1}{R} n' \times (\nabla' \times H^*) - (n' \times H^*) \times \frac{\hat{R}}{R^2} - \left(n' \times \frac{\partial H^*}{\partial t'} \right) \times \frac{\hat{R}}{R} - \frac{\hat{R}}{R^2} (H^* \cdot n') - \frac{\hat{R}}{R} \left(\frac{\partial H^*}{\partial t'} \cdot n' \right) ds'$$

(2.1.59)

Now equation (2.1.32) becomes, reordering a few terms for clarity:

$$H(r,t) = -W - \frac{1}{4\pi} \int_{\Omega} \frac{\hat{R}}{R^2} (\nabla' \cdot H^*) + \frac{\hat{R}}{R} \left(\nabla' \cdot \frac{\partial H^*}{\partial t'} \right) dv' +$$

$$\frac{1}{4\pi} \int_{\partial\Omega} \frac{1}{R} n' \times (\nabla' \times H^*) - (n' \times H^*) \times \frac{\hat{R}}{R^2} - \left(n' \times \frac{\partial H^*}{\partial t'} \right) \times \frac{\hat{R}}{R} - \frac{\hat{R}}{R^2} (H^* \cdot n') - \frac{\hat{R}}{R} \left(\frac{\partial H^*}{\partial t'} \cdot n' \right) ds'$$

(2.1.60)

We have here a formula for the field H at some location r in the interior of the domain, at some time t, in terms of:

(i) an integral over the domain of the initial conditions (the term W(r))

(ii) an integral over the surface of the domain of (some geometrical function of) the retarded (i.e. historical) values of the field and various spatial and temporal derivatives of the field.

(iii) an integral over the domain of the divergence of **H**, but as mentioned above that will be zero.

Derivation of the Boundary Integral Equation

Equation (strictly formula) (2.1.60) was for a point **r** restricted to lie somewhere in the interior of the domain. We now move the location **r** onto the boundary, and for generality we have selected a non-smooth portion of the boundary, where the exterior solid angle subtended is α. We will distort the boundary from this location, in much the same way as we excluded the small spherical region earlier. This is shown in figure 2.1.2. The physical boundary is indicated by the dotted line; outside this is a part-spherical distorted portion, centred on **r**. The physical boundary is again $\partial\Omega$, and the boundary of the distortion $\partial\Omega_\varepsilon$.

We are now at liberty to apply (2.1.60) at location **r**, integrating over the sum of $\partial\Omega$, and $\partial\Omega_\varepsilon$.

$$\mathbf{H}(\mathbf{r},t) = -\mathbf{W} - \frac{1}{4\pi}\int_\Omega \frac{\hat{\mathbf{R}}}{R^2}(\nabla'\cdot\mathbf{H}^*) + \frac{\hat{\mathbf{R}}}{R}\left(\nabla'\cdot\frac{\partial\mathbf{H}^*}{\partial t'}\right)dv' +$$

$$\frac{1}{4\pi}\int_{\partial\Omega+\partial\Omega_\varepsilon} \frac{1}{R}\mathbf{n}'\times(\nabla'\times\mathbf{H}^*) - (\mathbf{n}'\times\mathbf{H}^*)\times\frac{\hat{\mathbf{R}}}{R^2} - \left(\mathbf{n}'\times\frac{\partial\mathbf{H}^*}{\partial t'}\right)\times\frac{\hat{\mathbf{R}}}{R} - \frac{\hat{\mathbf{R}}}{R^2}(\mathbf{H}^*\cdot\mathbf{n}') - \frac{\hat{\mathbf{R}}}{R}\left(\frac{\partial\mathbf{H}^*}{\partial t'}\cdot\mathbf{n}'\right)ds'$$

$$(2.1.61)$$

We will concentrate on the integration over the distorted surface $\partial\Omega_\varepsilon$. Note that in what follows we have initially left in the asterisk to denote retarded values, but any term sifted out from a vanishingly small distance away will have an associated retarded time equal to the present time.

$$I_s = \frac{1}{4\pi} \int_{\partial\Omega_\epsilon} \frac{1}{R} \mathbf{n}' \times (\nabla' \times \mathbf{H}^*) - (\mathbf{n}' \times \mathbf{H}^*) \times \frac{\hat{R}}{R^2} - \left(\mathbf{n}' \times \frac{\partial \mathbf{H}^*}{\partial t'}\right) \times \frac{\hat{R}}{R} - \frac{\hat{R}}{R^2}(\mathbf{H}^* \cdot \mathbf{n}') - \frac{\hat{R}}{R}\left(\frac{\partial \mathbf{H}^*}{\partial t'} \cdot \mathbf{n}'\right) ds'$$

$$(2.1.62)$$

In the same way as for the excluded sphere, terms involving $1/R$ are eliminated once the $R^2\, d\theta$ substitution is made, and R is allowed to shrink to zero. Thus much as in (2.1.46) and (2.1.47), (2.1.62) becomes

$$I_s = \frac{-1}{4\pi} \int_\alpha \hat{R} \times (\hat{R} \times \mathbf{H}^*) - \hat{R}(\mathbf{H}^* \cdot \hat{R}) d\theta = \frac{\alpha}{4\pi} \mathbf{H}(\mathbf{r},t)$$

$$(2.1.63)$$

Equation (2.1.63) provides the value of the distorted-surface integral in (2.1.61). Using it, (and thus integrating now over only the physical boundary), we obtain

$$\mathbf{H}(\mathbf{r},t)\left(1 - \frac{\alpha}{4\pi}\right) = -\mathbf{W} - \frac{1}{4\pi} \int_\Omega \frac{\hat{R}}{R^2}(\nabla' \cdot \mathbf{H}^*) + \frac{\hat{R}}{R}\left(\nabla' \cdot \frac{\partial \mathbf{H}^*}{\partial t'}\right) dv' +$$

$$\frac{1}{4\pi} \int_{\partial\Omega} \frac{1}{R} \mathbf{n}' \times (\nabla' \times \mathbf{H}^*) - (\mathbf{n}' \times \mathbf{H}^*) \times \frac{\hat{R}}{R^2} - \left(\mathbf{n}' \times \frac{\partial \mathbf{H}^*}{\partial t'}\right) \times \frac{\hat{R}}{R} - \frac{\hat{R}}{R^2}(\mathbf{H}^* \cdot \mathbf{n}') - \frac{\hat{R}}{R}\left(\frac{\partial \mathbf{H}^*}{\partial t'} \cdot \mathbf{n}'\right) ds'$$

$$(2.1.64)$$

We see that the effect of moving to the boundary has been to modify the coefficient of the free term, from 1 to $\left(1 - \frac{\alpha}{4\pi}\right)$. Note that had the surface been smooth the modification would thus have been from 1 to 1/2.

We now in (2.1.64) have an expression for the field at a surface location involving fields at only other surface locations; in principle, a soluble system.

Application of the Boundary Conditions and Other Physics

We have as yet said nothing about boundary conditions and other aspects of the physics of our particular problem. We will be able to make various simplifications to our equation (2.1.64) (and corresponding ones could be made to the formula (2.1.60)).

We know that $\nabla' \cdot \mathbf{H}$ is zero. This eliminates the volume integral from (2.1.64) which then becomes:

$$\mathbf{H}(\mathbf{r},t)\left(1 - \frac{\alpha}{4\pi}\right) = -\mathbf{W} +$$

$$\frac{1}{4\pi}\int_{\partial\Omega}\frac{1}{R}\mathbf{n}'\times\left(\nabla'\times\mathbf{H}^*\right) - \left(\mathbf{n}'\times\mathbf{H}^*\right)\times\frac{\hat{\mathbf{R}}}{R^2} - \left(\mathbf{n}'\times\frac{\partial\mathbf{H}^*}{\partial t'}\right)\times\frac{\hat{\mathbf{R}}}{R} - \frac{\hat{\mathbf{R}}}{R^2}\left(\mathbf{H}^*\cdot\mathbf{n}'\right) - \frac{\hat{\mathbf{R}}}{R}\left(\frac{\partial\mathbf{H}^*}{\partial t'}\cdot\mathbf{n}'\right) ds'$$

$$(2.1.65)$$

It is often appropriate to model a scatterer as if it were the inhabitant of an infinite domain. In terms of our present derivation, we need a doubly connected domain, with a body wholly immersed in Ω, with its surface forming one of the boundaries of the domain. We must then allow the outer surface to recede to infinity.

We must now dispose of the integrations over this infinite boundary. It is perhaps most elegant to note that effects there at $t=0$, the earliest we need worry about, will not reach us till an infinite time later, by when our interest has declined considerably.

Physically our real world scattering problem will comprise some scatterer, and some transmitter; two bodies 'floating' in an infinite domain. The transmitter will have some impressed surface current (H field) distribution, and we will solve for the H field induced on the scatterer.

We could approach the modelling in this way, and then chose to describe the effect of integrations over the (remote) transmitter from points r on the scatterer in terms of an incident wave, H_{inc}. We would then impose zero initial conditions, such that W simply disappears.

Alternatively, we could ignore the transmitter, and say that the initial conditions throughout the domain were such as to result in some incident wave H_{inc} arriving at the scatterer; i.e. $H_{inc} = -W$.

If (2.1.61) is interpreted as integrations over the surface of only the scatterer, there is no practical difference between these approaches. Our boundary will subsequently be interpreted as the surface of the scatterer. We thus have

$$H(r,t)\left(1-\frac{\alpha}{4\pi}\right) = H_{inc} +$$

$$\frac{1}{4\pi}\int_{\partial\Omega}\frac{1}{R}n'\times\left(\nabla'\times H^*\right)-\left(n'\times H^*\right)\times\frac{\hat{R}}{R^2}-\left(n'\times\frac{\partial H^*}{\partial t'}\right)\times\frac{\hat{R}}{R}-\frac{\hat{R}}{R^2}(H^*\cdot n')-\frac{\hat{R}}{R}\left(\frac{\partial H^*}{\partial t'}\cdot n'\right)ds'$$

$$(2.1.66)$$

On a perfectly conducting surface there can be no tangential E field component: $n \times E = 0$. Similarly, there can be no normal H field component: $n \cdot H = 0$. We will apply these term by term to (2.1.66)

If $n \times E = 0$, $n \times \frac{\partial E}{\partial t} = 0$ also. By maxwell's equation, if this is so, $n \times \nabla' \times H = 0$ also. Therefore, the first term of integrand in (2.1.66) must vanish.

Note that it is here that we have decoupled the H and E fields. Although the H field has been chosen to be used here, in general it cannot be expected to solve using the H field alone, and it is here that the E field would have

featured. If we had not been able to state that there was no tangential component of E, we could not have deduced that the curl of H would be wholly normal, and so disappear when crossed with the normal. Keeping in the curl of H would have been equivalent to keeping the time derivative of E.

If $n \cdot H = 0$ then $n \times H$ is sufficient to define H; H lies in the surface , and we could identify the surface current, J_s, to write in place of $n' \times H$ in the second and third terms of integrand in (2.1.66). The fourth and fifth terms of the integrand are eliminated by $n \cdot H = 0$.

The incident wave term does not obey any boundary condition. However, we could cross H with n in the left hand side free term without any loss of the information because just as for terms two and three of the integrand, if $n \cdot H = 0$ then $n \times H$ is sufficient to define H.

Now the surface H field is uniquely defined by the two components and the third component can be obtained just from the linear combination of the other two components. However, the decision for solving for two components or three components should be made in terms of the efficiency of the numerical calculation process though only two components are independent of each other.

We thus finally arrive at the form we solve in practice

$$H(r,t)\left(1-\frac{\alpha}{4\pi}\right) = H_{inc} + \frac{1}{4\pi}\int_{\partial\Omega} -\left(n'\times H^*\right)\times\frac{\hat{R}}{R^2} - \left(n'\times\frac{\partial H^*}{\partial t'}\right)\times\frac{\hat{R}}{R}ds' \qquad (2.1.67)$$

Equation (2.1.67) is the MFIE for a perfect electric conductor (PEC) scatterer, immersed in an infinite dielectric, subject to some incident wave.

Note that we choose computationally to define normals outwards from the body, rather than outwards from the domain, as is 'mathematically' conventional. If this changed direction of normal is understood, the two minus signs in the integrand of (2.1.67) disappear. Thus finally we obtain

$$H(r,t)\left(1-\frac{\alpha}{4\pi}\right) = H_{inc} + \frac{1}{4\pi}\int_{\partial\Omega}(n'\times H^*)\times\frac{\hat{R}}{R^2} + \left(n'\times\frac{\partial H^*}{\partial t'}\right)\times\frac{\hat{R}}{R}ds' \tag{2.1.68}$$

2.2 Computational Issues for the Governing Equation

There exists a variety of approaches to disretisation and solution of equation (2.1.68), but to predict scattering from electrically large and complex bodies, efficient solution methods are required, as well as sophisticated modelling techniques. In this section, various approaches in use will be reviewed, and the present approach will be briefly introduced.

2.2.1 Geometry Representation

As a means of better representing curved surfaces while avoiding the excessive computational cost imposed by small mesh sizes, much work has been done in many application areas. However, in general the sophistication of electromagnetic BIE treatments lags somewhat behind the numerical techniques used in BIE treatment of areas such as elastostatics.

In the 1960's, the wire-grid modelling approach was widely used for the prediction of far field quantities such as radiation and radar cross sections[32].

This approach has many limitations such as ill-conditioned moment matrices and difficulties in calculating near-field and surface quantities[33]. Most of these limitations were overcome by the surface patch approach. In particular, Rao, Wilton and Glisson[4] used a planar triangular element which enforced continuity of the surface current at edges in 1982. It has been used for a long time[5,6,8]. Using flat elements for a problem with sharply curved geometry is naturally less desirable. Recently special attention has been paid to higher order geometric approximation by using quadrilateral elements[12,34], and more recently similar works have been done by Chao *et al* and Bluck and Walker[13,35].

In this thesis quadratic quadrilateral elements will be employed in representing a geometry. This implies that scatterers which require sophisticated geometrical treatment can be modelled in a relatively easy way. The elements employed will be isoparametric; as will be discussed below, similar modelling will be adopted for the field variation as for the geometry.

2.2.2 Field Representation in Space and Time

In recent years, the cost of solving RCS problems has matched and in some cases even surpassed that of computational fluid dynamics. The difficulty arises from the need to model not only the geometry but also the field variation. Even if the geometry is simple, modelling of the rapid variation of field can be expensive. For wavelengths which are short compared with the length of the body, this may lead to impractical computation times.

Generally, low order (piecewise constant or linear) polynomial interpolation of the field variable is used in MoM techniques. However, for the accurate modelling of the field variation, this approach requires more nodes per wavelength than does higher order polynominal interpolation.

On the other hand, the BIE method in other application areas has long used techniques based on FE methods to obtain more sophisticated field representation, by the use of higher order polynomial interpolation. These have had limited application to EM wave scattering problems. In the frequency domain, Ingber and Ott[34] used linear field representation in 1991 and *Chao et al*[35] used quadratic polynomial interpolation in representing field in 1995. In the time domain, quadratic representation of field variation in space and time has been used by only a very few researchers[12,13].

In this thesis, quadratic modelling of the field variation in both space and time will be used. As mentioned earlier, the use of higher order polynomials means that fewer nodes are required to describe any given field variation, in both space and time.

2.2.3 Integration

All integral methods involve the integration of a kernel function over a given domain. In most cases, this integration can be performed by employing standard numerical integration. However, when the field point lies in the element over which integration is being performed, the kernel function becomes singular over the domain of integration. This singular kernel function may be separated into two classes: weakly singular, and

strongly singular. Some works associated with the integration of such kernel functions are discussed below.

Regular Integration

Historically, the numerical integration applied to integral equations in electromagnetics has generally been rather basic, reflecting the field and geometry approximation themselves. As an instance, the simple midpoint-rule has been widely used for planar elements. In this thesis, standard Gaussian quadrature will be used, where the order can be varied at will to provide any desired degree of accuracy of integration.

Weakly Singular Integration

The weakly singular integration may be evaluated via Lachat and Watson[36] partitioning, followed by standard Gaussian Quadrature. This is akin to a transformation to polar co-ordinates and is discussed in detail elsewhere[29]. This partitioning scheme will be used in this thesis and details will be explained in the next section.

Strongly Singular Integration

The strongly singular integration is a more difficult problem. For the many researchers who have employed simple flat triangular elements[7,8,37,38,39,40,41] these difficult integrands very conveniently vanish. However, as explained above, flat triangular elements are best used only for simple shapes, and very many must be used in describing curved geometry and rapid field variations.

Cauchy principal value (CPV) integrations can be calculated by direct computation or indirect computation. Indirect methods can sometimes be used to obtain the diagonal matrix coefficients from knowledge of the values of the more easily obtained off-diagonal matrix elements. Methods to evaluate them more directly have been developed by several researchers. Recently a method using a parallelogram element has been reported by Liao and Xu[42] for three dimensional elastic problems. Similarly, a tangential element was developed by Bluck *et al*[43] for the solution of time domain EM wave scattering problems. The CPV integration is changed into a weakly singular integration and a contour integration. The weakly singular parts are evaluated by means of the partitioning scheme mentioned above, while the strongly singular parts of all contour integrals cancel with corresponding terms on neighbouring elements. This method has many advantages such as cheap computational cost, accuracy on curved (and stealthy) bodies, and easy extension to other element topologies. Thus it is adopted here and will be described further in the next section.

2.2.4 Implicitness, Explicitness and Stability

The usual approach in the numerical solution of transient wave scattering problems is the so-called marching-on-in-time method. Normally this method is implemented by forcing the treatment to be explicit, leaving no equation ever actually needing to be 'solved'. This indeed is thought of as an important advantage of this method; no matrix solution is needed if the time step is smaller than a certain upper limit which is determined by the spatial discretisation of the object.

However, an important disadvantage of this procedure is the possible occurrence of instabilities. Various techniques for solving these instabilities have been reported. Rynne[37,38] showed that the convergence of an explicit marching-on-in-time method can depend on the properties of the time derivative approximation (of function f, at time step $k+1$), and that careful evaluation of this term was required. The use of more stable approximations, a retrospective averaging approach: $f^k = \left(f^{k+1} + 2f^k + f^{k-1}\right)/4$, he found to give a marked improvement in the stability of the results with the modest loss in accuracy. An averaging process approach was also suggested by Smith[39]. It is suggested that the instability is related to the existence of resonant frequencies at which the corresponding frequency domain integral equation has more than one solution, and as a stabilising method, averaging solution over three time steps as the computation proceeds was found to be reasonably successful in suppressing instability.

In 1991 Rao and Wilton[8] presented explicit time domain scattering computations, and observed instabilities. In the following year Vechinski and Rao[7] proposed methods to stabilise their work. They employed a temporal retrospective averaging, which despite some correspondence[44] between them seems for practical purposes identical to that of Rynne mentioned above.

In this thesis, the implicit approach will be used, which seems in practice to eliminate the stability difficulties normally associated with explicit approaches. Although this implicit approach is accompanied by the need for

repeated matrix equation solution, the use of an iterative solver allows the additional cost of matrix solution to be trivial. Details will be explained in the next section.

2.3 Numerical Discretisation of MFIE

Numerical solution of equation (2.1.68) requires discretisation of the surface $\partial\Omega$. In this section, the numerical discretisation of equation (2.1.68) for a smooth surface ($\alpha = 2\pi$) is described by using quadrilateral elements in space and quadratic elements in time during which the field is to be modelled. Only an outline description will be given here, as a more detailed presentation will be given when showing the modified 'project forward' algorithm later.

2.3.1 Geometry Representation

The surface is divided into elements, over which the geometry is approximated by a set of polynomial functions;

$$S_\alpha(\xi,\eta) \qquad \alpha = 1,...,9 \qquad\qquad (2.3.1.1)$$

where ξ,η are the parameterised spatial co-ordinates. Here nine noded quadrilaterals are used, but in principle a wide range may be employed.

The geometry in each element is thus described by

$$\mathbf{r}^m(\xi,\eta) = \sum_{\alpha=1}^{9} S_\alpha(\xi,\eta)\mathbf{r}_{j(m,\alpha)} \qquad\qquad (2.3.1.2)$$

where $j = j(m,\alpha)$ are the global node numbers of the local nodes on element m, and $r_{j(m,\alpha)}$ is the position vector of each of these nodes.

The radial distance vector from the ith spatial node to some ξ,η co-ordinate location on the mth spatial element is then;

$$\mathbf{R}_m(\mathbf{r}_i;\xi,\eta) = \mathbf{r}_i - \sum_{\alpha-1}^{9} S_\alpha(\xi,\eta)\mathbf{r}_{j(m,\alpha)} \qquad (2.3.1.3)$$

The above transformation onto a bi-unit square in ξ,η co-ordinates allows a integration to be carried between +1 and -1 in both intrinsic co-ordinates ξ and η. The Jacobian of the transformation, essentially a position dependant scaling between areas in Cartesian co-ordinates and areas in ξ,η co-ordinates, must be evaluated. The area of an element of surface of sides $d\xi$ and $d\eta$ is obtained by

$$ds' = \left|\mathbf{e}_\xi \times \mathbf{e}_\eta\right| d\xi d\eta \qquad (2.3.1.4)$$

where $\mathbf{e}_\xi$ and $\mathbf{e}_\eta$ are the tangent vectors

$$\mathbf{e}_\xi = \frac{\partial x_1}{\partial \xi}\mathbf{e}_{x_1} + \frac{\partial x_2}{\partial \xi}\mathbf{e}_{x_2} + \frac{\partial x_3}{\partial \xi}\mathbf{e}_{x_3} \qquad (2.3.1.5)$$

$$\mathbf{e}_\eta = \frac{\partial x_1}{\partial \eta}\mathbf{e}_{x_1} + \frac{\partial x_2}{\partial \eta}\mathbf{e}_{x_2} + \frac{\partial x_3}{\partial \eta}\mathbf{e}_{x_3}$$

Thus the Jacobian can be written as;

$$|J(\xi,\eta)| = \left|\mathbf{e}_\xi \times \mathbf{e}_\eta\right| \qquad (2.3.1.6)$$

In addition to the Jacobian, the unit normal vector over the surface is required. The unit normal vector at a point (ξ,η) can be obtained from

$$\mathbf{n}(\xi,\eta) = \frac{\mathbf{e}_\xi \times \mathbf{e}_\eta}{\left|\mathbf{e}_\xi \times \mathbf{e}_\eta\right|} \qquad (2.3.1.7)$$

With the Jacobian of transformation $|J(\xi,\eta)|$ and the unit normal vector $\mathbf{n}(\xi,\eta)$, equation (2.2.37) can be written

$$2\pi\mathbf{H}(\mathbf{r}_i,t) = 4\pi\mathbf{H}_{inc}(\mathbf{r}_i,t)$$

$$+\sum_{m=1}^{M}\iint_{\xi,\eta}\left\{\sum_{\alpha=1}^{9}S_\alpha(\xi,\eta)\left(\mathbf{n}'(\xi,\eta)\times\mathbf{H}\left(\xi,\eta;t^*(R_m,t)\right)\right)\times\frac{\mathbf{R}_m}{R_m^3}\right\}|J(\xi,\eta)|d\xi d\eta \qquad (2.3.1.8)$$

$$+\sum_{m=1}^{M}\iint_{\xi,\eta}\left\{\sum_{\alpha=1}^{9}S_\alpha(\xi,\eta)\left(\mathbf{n}'(\xi,\eta)\times\mathbf{H}\left(\xi,\eta;t^*(R_m,t)\right)\right)\times\frac{\mathbf{R}_m}{R_m^2}\right\}|J(\xi,\eta)|d\xi d\eta$$

where the dependence of t^* on R is explicitly indicated.

2.3.2 Field Representation in Space and Time

Using an isoparametric formulation, the spatial variation of the surface magnetic field over an element m may be written

$$\mathbf{H}(\xi,\eta;t)\big|_m = \sum_{\alpha=1}^{9}S_\alpha(\xi,\eta)\mathbf{H}_{j(m,\alpha)}(t) \qquad (2.3.2.1)$$

where again, $j = j(m,\alpha)$ are the global node numbers of the local nodes on element m.

The temporal variation of the magnetic field over a temporal element l is modelled by using quadratic elements of length $2\Delta t$, with associated basis functions $T_\beta(\tau)$, with τ the intrinsic time.

$$\mathbf{H}_{j(m,\alpha)}(\tau)\big|_l = \sum_{\beta-1}^{3} T_\beta(\tau)\mathbf{H}_{j(m,\alpha)}^{k(l,\beta)} \qquad\qquad (2.3.2.2)$$

where $k = k(l,\beta)$ are the global timestep numbers of the local temporal nodes on temporal element l, and $\mathbf{H}_{j(m,\alpha)}^{k(l,\beta)}$ are the field values at the spatial nodes, at the three time steps of the temporal element in which the time of interest falls.

Now the field at some intrinsic location (ξ,η) within spatial element m, at an intrinsic time τ within temporal element l, is approximated as;

$$\mathbf{H}(\xi,\eta,\tau)\big|_{m,l} = \sum_{\alpha=1}^{9}\sum_{\beta=1}^{3} S_\alpha(\xi,\eta)T_\beta(\tau)\mathbf{H}_{j(m,\alpha)}^{k(l,\beta)} \qquad\qquad (2.3.2.3)$$

Any particular location in (space, time) will lie in a single (spatial, temporal) element. Thus for a nine noded quadrilateral spatial element, the field at some location on it, at some time, is expressed as the weighted sum of the 27 (spatial and temporal) nodal values which surround it in the 'space-time' box.

Now the field is evaluated for a particular timestep $k+1$, with $t = (k+1)\Delta t$ For the lossless dielectric of interest, historical field values at retarded times are required. It is convenient to arrange that the present timestep, $k+1$, for which the field is being found, forms the final node of a temporal element. Then the number l^* of elements ago that the element in which the relevant retarded time falls is computed , via;

$$l^* = \text{int}\left\{\frac{R}{2\Delta t}\right\} \qquad (2.3.2.4)$$

and we identify the set of timestep numbers $k(l^*,\beta) = k - 2l^* - 2 + \beta$ associated with that element. The intrinsic time τ corresponding to the retarded time is then found via

$$\tau(R) = \frac{(2l^* + 1)\Delta t - R}{\Delta t} \qquad (2.3.2.5)$$

Then equation (2.3.1.8) may be written, using (2.3.2.1) to (2.3.2.5), as

$$2\pi \mathbf{H}_i^{k+1} = 4\pi \mathbf{H}_{inc,i}^{k+1} \qquad (2.3.2.6)$$

$$+ \sum_{m=1}^{M} \iint_{\xi,\eta} \left\{ \sum_{\alpha=1}^{9} S_\alpha(\xi,\eta) \sum_{\beta=1}^{3} \left(\frac{T_\beta(\tau(R_m))}{R_m^3} + \frac{\dot{T}_\beta(\tau(R_m))}{\Delta t R_m^2} \right) \left[\left(\mathbf{n}'(\xi,\eta) \times \mathbf{H}_{j(m,\alpha)}^{k(l^*,\beta)} \right) \times \mathbf{R}_m \right] \right\} |J(\xi,\eta)| d\xi d\eta$$

2.3.3 Integration

It is clear that the integrals contained in equation (2.3.2.6) must be evaluated numerically. The elements where the field point is not on the element being integrated, which will be termed the 'non-self element', can be treated by standard Gaussian Quadrature (regular integration). However, if the field point lies in the element being integrated, integrations which will be termed 'self element', the integrand becomes badly behaved, with the $1/R$ and $1/R^2$ terms giving rise to weakly and strongly singular integrands respectively. For each case, a brief explanation is given below.

Regular Integration

For a single one of the non-self elements, integration in the summation in equation (2.3.2.6) can be performed, using standard Gaussian quadrature of order $NG \times NG$:

$$\delta H_m = \iint_{\xi,\eta} \left\{ \sum_{\alpha=1}^{9} S_\alpha(\xi,\eta) \sum_{\beta=1}^{3} \left(\frac{T_\beta(\tau(R_m))}{R_m^3} + \frac{\dot{T}_\beta(\tau(R_m))}{\Delta t R_m^2} \right) \left[\left(\mathbf{n}'(\xi,\eta) \times \mathbf{H}_{j(m,\alpha)}^{k(l^*,\beta)} \right) \times \mathbf{R}_m \right] \right\} |J(\xi,\eta)| d\xi d\eta$$

$$= \sum_{p=1}^{NG} \sum_{q=1}^{NG} \sum_{\alpha=1}^{9} S_\alpha(\xi,\eta) \sum_{\beta=1}^{3} \left(\frac{T_\beta(\tau(R_m))}{R_m^3} + \frac{\dot{T}_\beta(\tau(R_m))}{\Delta t R_m^2} \right) \left[\left(\mathbf{n}'(\xi,\eta) \times \mathbf{H}_{j(m,\alpha)}^{k(l^*,\beta)} \right) \times \mathbf{R}_m \right] |J(\xi,\eta)| \omega_p \omega_q$$

(2.3.3.1)

Part of the kernel of equation (2.3.3.1) is conveniently expressed by writing $(\mathbf{n} \times \mathbf{H}) \times \mathbf{R}$ as a $[A']\mathbf{H}$, with the matrix $[A']$ given by:

$$[A'] = \begin{bmatrix} n_2'R_2 + n_3'R_3 & -n_1'R_2 & -n_1'R_3 \\ -n_2'R_1 & n_1'R_1 + n_3'R_3 & -n_2'R_3 \\ -n_3'R_1 & -n_3'R_2 & n_1'R_1 + n_2'R_2 \end{bmatrix}$$

(2.3.3.2)

where in Cartesian co-ordinates $\mathbf{n}' = (n_1', n_2', n_3')$ and $\mathbf{R}_m = (R_1, R_2, R_3)$.

Thus equation (2.3.3.1) becomes

$$\delta H_m = \sum_{p=1}^{NG} \sum_{q=1}^{NG} \sum_{\alpha=1}^{9} S_\alpha(\xi_p, \eta_q) \sum_{\beta=1}^{3} \left(\frac{T_\beta(\tau(R_m))}{R_m^3} + \frac{\dot{T}_\beta(\tau(R_m))}{\Delta t R_m^2} \right) [A'] \mathbf{H}_{j(m,\alpha)}^{k(l^*,\beta)} |J(\xi_p, \eta_q)| \omega_p \omega_q$$

(2.3.3.3)

where Gauss weights are represented as ω_p and ω_q, and Gauss locations as ξ_p and η_q.

Note that the retarded time associated with different Gaussian locations in a given spatial element may well fall in different temporal elements. Each such location would generate 3 sets of 9 coefficients. Since the temporal

elements would be adjacent, the first and last sets of nine would generally be common to the adjacent temporal elements, resulting in 9×3, 9×5, 9×7, ...and so on possible distinct coefficients resulting from integration over a single element.

Weakly Singular Integration

The self-element is partitioned into triangles, whose number NP depends on the location of the field node within the element. These triangles allow a polar co-ordinate system to be used which removes the singularity, and this in turn is transformed into an integral over bi-unit square in new intrinsic $\hat{\xi}, \hat{\eta}$ co-ordinates, to allow Gaussian quadrature to be used. The integration over a single one of the self- elements can then be written as ;

$$\delta H_m = \sum_{s=1}^{NP} \sum_{p=1}^{NG} \sum_{q=1}^{NG} \sum_{\alpha=1}^{9} S_\alpha(\xi_p, \eta_q) \sum_{\beta=1}^{3} \left(\frac{T_\beta(\tau(R_m))}{R_m^3} + \frac{\dot{T}_\beta(\tau(R_m))}{\Delta t R_m^2} \right) [A'] H_{j(m,\alpha)}^{k(l',\beta)} |J(\xi_p, \eta_q)| |J_s(\hat{\xi}_p, \hat{\eta}_q)| \omega_p \omega_q$$

$$(2.3.3.4)$$

where $|J_s(\hat{\xi}_p, \hat{\eta}_q)|$ is the Jacobian for the transformation of triangle into the bi-unit square. Fuller details are given elsewhere[29].

Strongly Singular Integration

The method presented, presented first by Bluck et al[43], here involves the subtraction of a related, similarly singular kernel from the true kernel. This technique weakens the singularity sufficiently for the above $O(1/R)$ singular integration method to be used for the residue. The singular kernel introduced is constructed in such a way that it has the crucial property of

cancelling exactly with this same kernel introduced on neighbouring elements.

Consider the following integral, which is typical of those arising in equation (2.3.2.6)

$$\delta \mathbf{H}_m = \int_m S_\alpha(\xi,\eta) \left[\frac{T_\beta(\tau)}{R_m^3} + \frac{\dot{T}_\beta(\tau)}{\Delta t R_m^2} \right] (\mathbf{n}' \times \mathbf{H}') \times \mathbf{R}_m ds' \qquad (2.3.3.5)$$

This integral is computed over an element m, in which a strongly singular integration arises.

Now a 'tangential element', $\tilde{m}$ is constructed as shown in figure 2.3.1, chosen to be a parallelogram tangent to the element m at the singular node. Using the element $\tilde{m}$, equation (2.3.3.5) can be written in the following form:

$$\delta \mathbf{H}_m = \int_m S_\alpha(\xi,\eta) \left[\frac{T_\beta(\tau)}{R_m^3} + \frac{\dot{T}_\beta(\tau)}{\Delta t R_m^2} \right] (\mathbf{n}' \times \mathbf{H}') \times \mathbf{R}'_m ds'$$

$$- \int_{\tilde{m}} (\tilde{\mathbf{n}}' \times \mathbf{H}') \times \frac{\tilde{\mathbf{R}}_m}{\tilde{R}_m^3} ds' + \int_{\tilde{m}} (\tilde{\mathbf{n}}' \times \mathbf{H}') \times \frac{\tilde{\mathbf{R}}_m}{\tilde{R}_m^3} ds' \qquad (2.3.3.6)$$

All quantities with the $\sim$ symbol are evaluated on the newly constructed tangential element $\tilde{m}$, and the second and third integrals on the right hand side of equation (2.3.3.6) are computed over it. By construction, and from figure 2.3.1, it is seen that as R approaches to zero, $\tilde{R}$, $\tilde{J}$, and $\tilde{\mathbf{n}}'$ get closer and closer to R, J, and $\mathbf{n}'$ respectively.

By changing to parameterized space and using Stokes' theorem, equation (2.3.3.6) can be written as;

$$\delta\mathbf{H}_m = \iint_{\xi\eta} S_\alpha\left[\frac{T_\beta(\tau)}{R_m^3} + \frac{\dot{T}_\beta(\tau)}{\Delta t R_m^2}\right](\mathbf{n}'\times\mathbf{H}')\times\mathbf{R}_m|\mathbf{J}(\xi,\eta)|d\xi d\eta - \iint_{\xi\eta}(\tilde{\mathbf{n}}'\times\mathbf{H}')\times\frac{\tilde{\mathbf{R}}_m}{\tilde{R}_m^3}|\tilde{\mathbf{J}}|d\xi d\eta$$

$$-\oint_{\partial(\tilde{m})}\frac{1}{\tilde{R}_m}d\mathbf{l}\times\mathbf{H}' \tag{2.3.3.7}$$

$|\mathbf{J}(\xi,\eta)|$ and $|\tilde{\mathbf{J}}|$ denote the Jacobian of the transformation from both the true element and the tangent element respectively onto the bi-unit square. As the notation implies, the latter is a constant by construction.

Now the integrand in the first term of equation (2.3.3.7) is at worst $O(1/R)$ singular and so the technique discussed above can be used. Finally using matrix notation, equation (2.3.3.7) can be expressed as;

$$\delta\mathbf{H}_m = \left\{\sum_{s=1}^{NP}\sum_{p=1}^{NG}\sum_{q=1}^{NG}\left\{\begin{array}{c}S_\alpha(\xi_p,\eta_q)\left(\dfrac{T_\beta(\tau)}{R_m^3} + \dfrac{\dot{T}_\beta(\tau)}{\Delta t R_m^2}\right)[A']|\mathbf{J}(\xi_p,\eta_q)||\mathbf{J}_s(\hat{\xi}_p,\hat{\eta}_q)| \\ -\dfrac{1}{\tilde{R}_m^3}|\tilde{\mathbf{J}}||\mathbf{J}_s(\hat{\xi}_p,\hat{\eta}_q)|[\tilde{A}']\end{array}\right\}\omega_p\omega_q - [P]\right\}\mathbf{H}_{j(m,\alpha)}^{k(l^*,\beta)}$$

$$\tag{2.3.3.8}$$

where

$$[\tilde{A}'] = \begin{bmatrix} \tilde{n}_2'\tilde{R}_2 + \tilde{n}_3'\tilde{R}_3 & -\tilde{n}_1'\tilde{R}_2 & -\tilde{n}_1'\tilde{R}_3 \\ -\tilde{n}_2'\tilde{R}_1 & \tilde{n}_1'\tilde{R}_1 + \tilde{n}_3'\tilde{R}_3 & -\tilde{n}_2'\tilde{R}_3 \\ -\tilde{n}_3'\tilde{R}_1 & -\tilde{n}_3'\tilde{R}_2 & \tilde{n}_1'\tilde{R}_1 + \tilde{n}_2'\tilde{R}_2 \end{bmatrix} \tag{2.3.3.9}$$

and

$$[P] = \begin{bmatrix} 0 & -\hat{\mathbf{k}}\cdot\mathbf{p} & -\hat{\mathbf{j}}\cdot\mathbf{p} \\ \hat{\mathbf{k}}\cdot\mathbf{p} & 0 & -\hat{\mathbf{i}}\cdot\mathbf{p} \\ -\hat{\mathbf{j}}\cdot\mathbf{p} & \hat{\mathbf{i}}\cdot\mathbf{p} & 0 \end{bmatrix} \tag{2.3.3.10}$$

$$P = \oint_{\partial(\tilde{m})} \frac{1}{R} d\mathbf{l}$$

(2.3.3.11)

Here equation (2.3.3.7) is the particular case of equation (2.3.3.4) when $\beta = 3$, $l^* = 1$ and $j(m, \alpha) = i$, the singular node. The contour integral of equation (2.3.3.11) can be treated analytically, as described in the paper by Bluck *et al*[43].

2.3.4 Implicitness, Explicitness and Stability

Evaluation of the summations in equation (2.3.3.1) provides an expression for the contribution to the field at a location in terms of a weighted sum of historical field values. The field on the left of equation (2.3.3.1) is (as yet) a three component vector, expressed as a weighted sum of historical three component vectors. The weights thus take the form of 3 by 3 (sub) matrices. Application of the PEC boundary conditions allows these to be reduced to 2 by 2 matrices. This elimination process is described in various texts[45,46].

Now all of the M equations (2.3.3.1) can be assembled, and non-zero matrix entries which are related to timesteps before the present one at which the field is being sought can be multiplied by the associated historical field value, forming a known contribution to the field which can be added to the incident wave term. Coefficients multiplying as yet unknown (i.e. nearby) field values are moved to the left hand side, and together with the free term form a sparse matrix equation for these new field values:

$$[a]H^{k+1} = [c]$$

(2.3.4.1)

Note that neighbouring new field values affect each other; a matrix equation is solved at each time step, and the approach used here is thus not an explicit but an implicit one. Implicitness has important consequences. In spite of the disadvantage of requiring solution of matrix equation, this is for all practical purposes not subject to instability, as the explicit approach is. Furthermore, and of crucial importance for computational costs, it allows local refinement to be used at edges, apertures, corners and so on, without requiring a corresponding reduction in timestep.

2.3.5 Matrix Solution

The rational choice of time step Δt, providing broadly equally good modelling of temporal and spatial variations of field, ensures that in general only a handful of distances between field point and Gauss points on boundary elements are shorter than $c\Delta t$. The consequence is that the $[a]$ matrix in equation (2.3.4.1) is sparse and this sparse matrix equation is solved at every time step.

Using iterative methods such as the conjugate gradient (CG) method, and with an excellent initial guess available in the form of the field at the previous time step, convergence is generally achieved in a handful of iterations. These iterative techniques involve a simple matrix vector multiplication, where the matrix $[a]$ has very small number of entries per row. The result of this makes the matrix solution cost a trivial fraction of the total cost, and a fraction that itself declines as larger problems are tackled.

2.4 Pulsed Excitation

The time domain Maxwell equations represent a more general form than the frequency domain Helmholz equation, which are usually restricted to solving scattering problems having time harmonic fields. A time-domain approach can obtain continuous wave (single frequency, harmonic) as well as single-pulse (broad-band frequency) transient response.

However in general, the RCS response is desired in the frequency domain, whereas the computational solution is in the time domain. The time domain results of electric and magnetic scattered waves are processed using a spectral technique such as Fourier deconvolution to obtain the response in the frequency domain. For a continuous wave representing a single frequency, the spectral analysis will provide the response at the given frequency, whereas for a pulse case containing many frequencies, the spectral analysis allows one to compute the RCS response for all frequencies contained in the incident pulse from a single time-domain transient calculation. An upper limit to the frequency which can in practice be extracted arises from the 'cut-off' implicit in the discretisation.

The pulse width and the frequency range which can be extracted from the pulse are related. Generally, as the pulse becomes narrow, the wider is the range of frequency that can be extracted. However, the pulse must be wide enough for the timestep to model its temporal variation, and the spatial discretisation its spatial variation. Therefore, choosing a proper pulse width may be explained in terms of the associated number of timesteps. Some

results will be demonstrated by relating them to the timestep size in section 2.7.

In this thesis, a Gauss pulse is used for the pulsed excitation, and thus the incident wave $\mathbf{H}_{inc}$ has the form:

$$\mathbf{H}_{inc}(\mathbf{r},t) = \mathbf{H}_0 \exp\left\{ -\frac{4\ln 2(t - R/c)^2}{g^2} \right\} \qquad (2.4.1)$$

where the pulse width parameter g is selected to provide the desired width at half maximum of this pulse.

2.5 Cost Scaling of BIEM for RCS problems

In the section below are obtained order of magnitude estimates of the computational costs of BIE techniques for transient wave scattering problems. Two main computational costs are of interest; matrix formation and matrix-vector multiplication for forming the right hand side vector. The dominant cost is dependent on the particular problem, principally the number and frequencies of the incident waves. For the MFIE, computational costs will be calculated below as a function of the number and frequencies of the incident waves.

A perfectly conducting sphere of P wavelengths diameter is considered as a geometry. The body is covered in nodal points at a spatial separation of n per wavelength and the temporal discretisation is at the same rate n per period of radiation. In the discussion below τ_F and τ_M are defined as the

costs of matrix formation and matrix-vector multiplication per coefficient respectively.

The approximate number of surface nodes and time steps per transit can then be calculated as;

$$\pi P^2 n^2 \qquad (2.5.1)$$

and

$$Pn \qquad (2.5.2)$$

respectively.

The computational costs can be investigated by estimating the number of non-zero matrix coefficients. It is assumed, due to the implicit nature of the method in this thesis, that the time step Δt is bounded by the maximum spatial nodal spacing. The maximum number of temporal elements required to 'cover' any spatial element is then three. The minimum will be obviously one, that is, one temporal element may completely 'cover' the spatial element. This is likely to be the case in particular in a problem with local mesh refinement.

If the maximum number of temporal elements per spatial element is three, then at most seven and at least three contributions may arise from the integral over this element. Considering 3 by 3 (sub) matrices, thus the number of non-zero matrix entries $N_{matrix-entries}$ can then be shown to be;

$$(3 \times 3) \times 3 \times (\pi P^2 n^2)^2 \leq N_{matrix-entries} \leq (3 \times 3) \times 7 \times (\pi P^2 n^2)^2 \qquad (2.5.3)$$

or equivalently,

$$27\pi^2 P^4 n^4 \leq N_{matrix-entries} \leq 63\pi^2 P^4 n^4 \qquad (2.5.4)$$

Therefore, the number of non-zero matrix entries can be written as;

$$N_{matrix-entries} = \alpha\pi^2 P^4 n^4 \qquad (27 \leq \alpha \leq 63) \qquad (2.5.5)$$

Now the matrix formation cost C_{form} can be directly calculated by multiplying the number of non-zero matrix entries by τ_F, the formation time of each, without any consideration of the incident waves, because the matrix is only dependent on geometry. Thus C_{form} will be shown to scale with the fourth power of the frequency of the incident wave;

$$C_{form} = \alpha\pi^2 P^4 n^4 \times \tau_f \qquad (2.5.6)$$

The number of incident waves must be considered for the calculation of the matrix-vector multiplication cost $C_{multiplication}$, and this multiplication is performed as many times as the total number of time steps, which is typically perhaps three times the number of time steps per transit. It should also noted that by now the 3 by 3 (sub) matrices have been reduced to 2 by 2 by the application of the PEC boundary conditions. Thus $C_{multiplication}$ can be written as below:

$$C_{multiplication} = N_{Waves} \times 3Pn \times \frac{4}{9}\alpha\pi^2 P^4 n^4 \times \tau_m = N_{Waves} \times \frac{4}{3}\alpha\pi^2 P^5 n^5 \times \tau_m \qquad (2.5.7)$$

where N_{Waves} is the number of incident waves.

In the above equation, N_{Waves} should increase with frequency due to the fact that as the frequency of incident wave increases, more illumination angles are required so as to characterise the response. Thus $C_{multiplication}$ will be by far the dominant cost, which scales with the fifth or sixth power of the frequency as body size increases, even though the cost ratio τ_F/τ_M is typically ~100 in time.

These two main costs are compared in figure 2.5.1, where N_{Waves} is approximated as $20P$ for the multiple illumination angles, and the values n and α are taken to be 10 and 54 respectively. Figure 2.5.1 shows that $C_{multiplication}$ is an increasing fraction of the total for a single illumination angle and as would be expected is by far the dominant cost for the multiple illumination angles.

2.6 Computational Implementation of BIEM

The methods described in section 2.3 have been implemented in Fortran 77, mainly on workstations including DEC alpha and Silicon Graphics. Creating solid models and meshes, and displaying results are via a direct interface to the widely used commercial CAD package; MSC 'Patran'.

2.6.1 Pre-Processing

There are two main activities in this stage. One is converting a neutral file from Patran to a mesh file which is easily recognisable and usable as an input file. This file has nodal co-ordinates and element information such as the node numbers each element has.

The other is performing tests on the quality of mesh. This process has many advantages because the test results are sometimes useful in analysing results obtained finally, and much information such as node normals and Cartesian co-ordinates of the intrinsic co-ordinates origin is pre-calculated and stored.

2.6.2 Surface Current Calculation

This stage is the most important and costly. The first step involved in this stage is the matrix formation. The integration of the kernel over all elements is performed for all (field) nodes. For each field point and element combination, there will be an array of linear relationships between the H field at a field node and the H field at a boundary node on an element at a small number of neighbouring time steps. All the contributions from integrating from a field point over all elements are inserted into a matrix row. This row is generally sparse, with entries clustered around a few time steps, and therefore is rearranged in a compact form so as to reduce storage requirements. It is then written to a matrix file in the specified location.

As the second step, the matrix is read from the disk, and the corresponding right hand side vector is calculated for each timestep. Finally the matrix equation is solved by the conjugate gradient squared (CGS) method at each location, for each time step.

2.6.3 Far Field Solution

Once the surface field has been calculated, the far field solution can be obtained using a spectral technique. Here the Fourier decomposition

method is used for obtaining RCS response to harmonic waves at a given frequency.

2.7 Some Examples of EM Scattering Problems

In this section the methods described in previous sections are demonstrated, using three different geometries illuminated with a Gaussian pulse. In section 2.7.1, numerical results are compared with analytical results for a sphere. The proper pulse width is investigated in terms of accuracy in section 2.7.2, and some results for a more difficult geometry than a sphere are shown using this pulse width in section 2.7.3. The surface current is calculated in all cases, and as a far field solution, a number of bistatic RCS evaluations are made.

2.7.1 Sphere

Figure 2.7.1.2 shows the bistatic RCS of the one wavelength long (674 node) sphere in figure 2.7.1.1, calculated by the Mie series, and by BIE time domain method. The bistatic RCS is extracted from the results of a pulsed illumination. Good agreement can be seen between the analytical and calculated results. Figure 2.7.1.3 shows the time variation of the surface field, H, for a single surface location, at (-0.15, -0.79, -0.59) when the wave propagates in the +y direction. The peak when the wave arrives, at time, ~2.8, is easily seen.

2.7.2 Choosing the Pulse Width

The question of the optimum combination of mesh nodal spacing, timestep length, and illumination pulse width, to achieve some desired mix of accuracy and computational cost, is complicated. It is also to a degree problem dependant, and in practice some essentially empirical, pragmatic approach must be used.

Here we will show some results on a dipole geometry; they are typical of a large number, from which the 'rules of thumb' actually employed have been derived.

Figure 2.7.2.2 shows results on a 10:1 cylindrical dipole, with hemispherical ends, shown in figure 2.7.2.1. It was illuminated from head-on with a Gaussian pulse with several different pulse widths to extract the bistatic RCS results at a certain frequency.

The bistatic RCS in figure 2.7.2.2 is extracted for a wavelength of ~1/4 of the dipole length. This corresponds to a maximum nodal separation of 1/10 of a wavelength. This is known from frequency domain work, where obviously timestep issues do not arise, to provide a good compromise between cost and accuracy. With quadratic modelling of both spatial and temporal variation, we use "the same" timestep; that is, $c\Delta t \sim \Delta x$, where c is the wave speed, and Δt and Δx the timestep and maximum nodal spacing respectively. The different lines on this figure correspond to illuminations using different pulse widths. Pulse widths are described in terms of the number of timesteps used to model the width at half maximum of the Gaussian pulse. A broad pulse provides 'less' of the higher frequency components of interest, but is easier to model; a narrow pulse, the reverse. It

is seen from the figure that good agreement is found between time and frequency domain results with the combination of a timestep size and maximum mesh spacing each about 1/3 of the pulse half-width. This combination, or there about, is generally used in this work.

2.7.3 NASA Almond

The almond is one of a set of benchmark RCS targets described by Woo[47]. It probably provides a good test of the approaches adopted in this thesis, as the backscatter RCS is low, and the main backscatter is probably due to waves creeping around the surface.

In figure 2.7.3.1, a 690 almond mesh is shown, and the 1.19 GHz (approximately 1 wavelength long) bistatic RCS for vertical(VV) polarisation is presented in figure 2.7.3.2. It shows good agreement at the backscatter between the measured (~38 dB sq m) and calculated (38.257 dB sq m) RCS.

2.8 Discussion

This chapter has described in detail the application of a conventional boundary integral equation method to the solution of transient EM wave scattering problems. The use of quadratic interpolation in representing a geometry and field has been shown, and the singular integration and implicitness associated with stability problem are considered.

The methods adopted in this thesis have been demonstrated for a few simple problems, and good results have been obtained. Satisfactory means

has been identified to help select the combinations of pulse width, mesh refinement and timestep to employ. In particular, it has been shown that these methods, incorporating smooth (non-facetted) modelling, allows accurate results to be obtained for difficult problems (for example, almond), even using relatively modest computational resources.

Figure 2.1.1   Domain and domain boundaries (I)

Figure 2.1.2   Domain and domain boundaries (II)
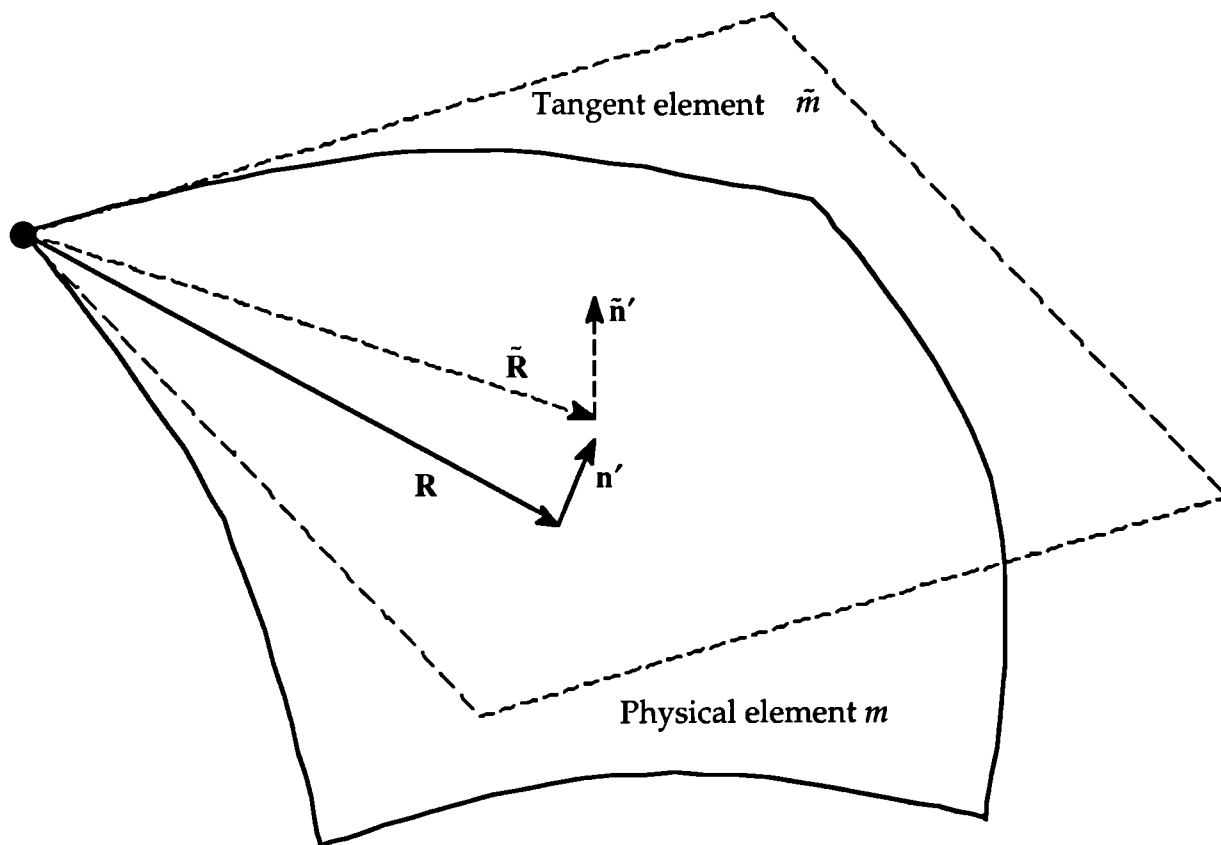
Figure 2.3.1   Tangential element constructed to quadratic element
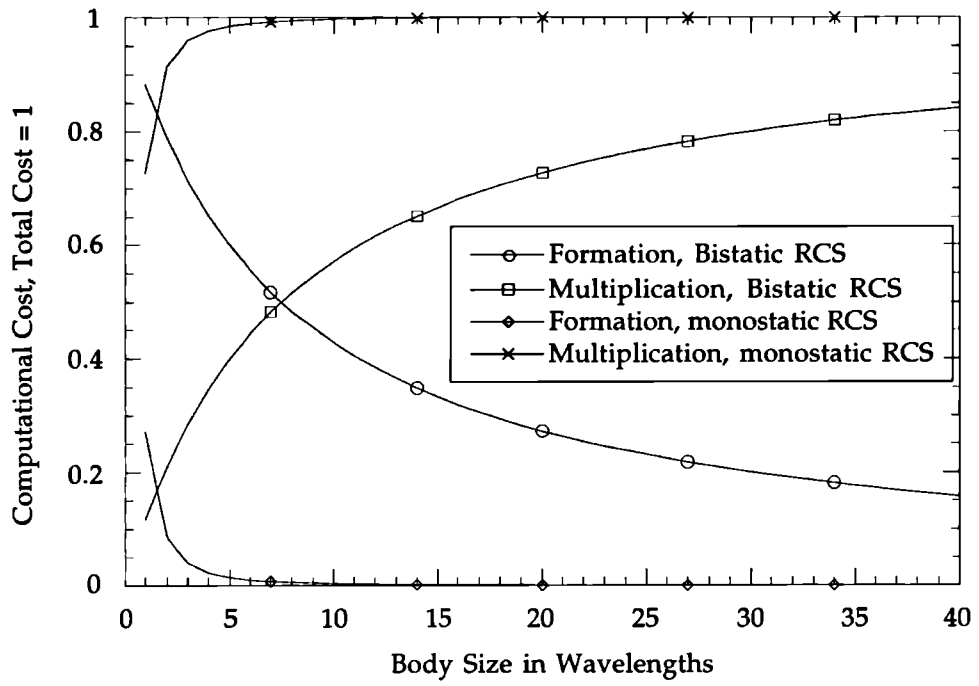
Figure 2.5.1   Computational costs versus body size in wavelengths

Figure 2.7.1.1   674 node sphere mesh
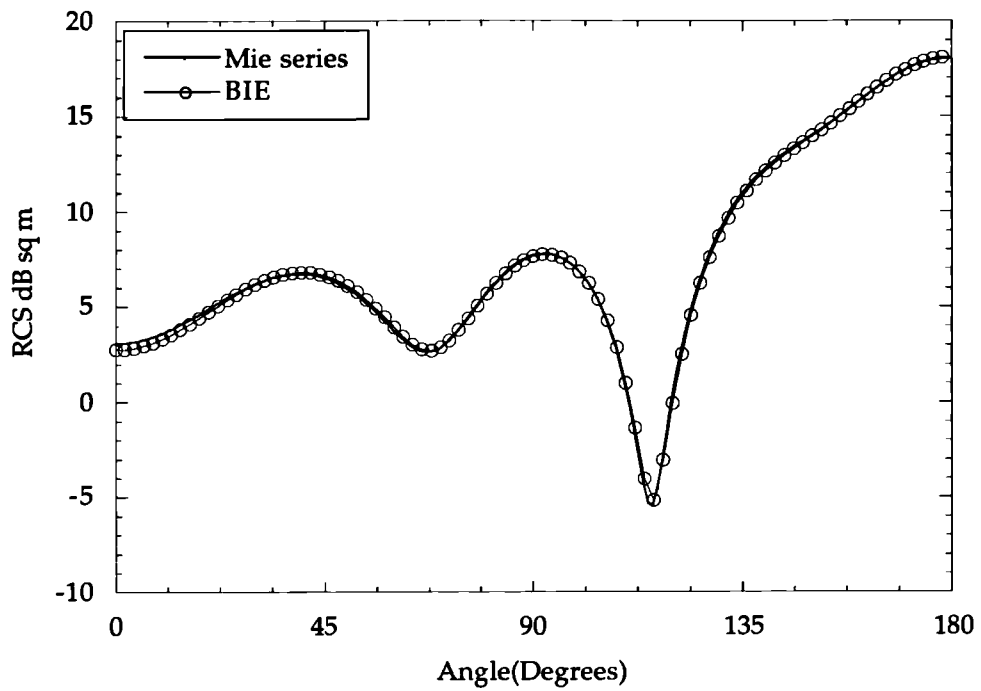
Figure 2.7.1.2   Bistatic RCS of sphere : 1.3 wavelengths long in diameter

Figure 2.7.1.3   Time variation of the surface field at a single location

Time Domain Boundary Integral Equation Method for EM Wave Scattering   77

Figure 2.7.2.1   674 node dipole mesh

Figure 2.7.2.2   Bistatic RCS of 10:1 dipole : 4 wavelengths long case

Figure 2.7.3.1   690 node almond mesh

Figure 2.7.3.2   Bistatic RCS of a NASA almond : 1 wavelength long

Chapter 3 Physical Basis of the Computational Cost Reduction

In this chapter, the physical observation which allows the project forward algorithm to achieve cost savings is described. It was presented earlier by Walker[48], along with an investigation into the likely loss in solution accuracy the approximation involved would cause. For completeness, however, it is briefly summarised here based closely on the reference cited, and demonstrations of the physical basis, without any modification of the algorithm to reduce costs, are given here.

This chapter is organised as follows. Section 3.1 briefly introduces the cost reduction methods. As a main section, section 3.2 describes these methods and shows some ways of implementing them. In section 3.3 are presented the results of investigations into the accuracy loss which the additional physical approximation causes, and into the accompanying cost saving which might be achieved if that accuracy loss is acceptable. Finally section 3.4 discusses possible ways of implementing the cost reduction methods.

3.1 Introduction

One approach to time domain modelling is to illuminate with a harmonic wave, following the calculation for long enough for transient effects to become negligible. The harmonic surface field can then be interrogated to obtain the response at that single frequency.

An alternative is to use a (for example Gaussian) pulse. This is in practice less straightforward; the interactions between optimum mesh fineness,

timestep size and maximum frequency which can reliably be extracted are rather complicated. It does have the major advantage that responses at other frequencies up to maximum can also be obtained at the minimal cost.

For present purposes, the pulsed illumination has a more crucial role to play. It is on the pulsed nature of the excitation, with the primary excitation being in length a small fraction of the body length, that the cost saving measures depend. These physical approximations, described in section 3.2, will provide an opportunity to reduce power dependence of computational costs with the frequency for this pulsed excitation form.

3.2 Physical Basis of Cost Saving

A pulse which is short compared to the body will tend to move over the body as a fairly narrow 'active' band, of a length in the direction of travel roughly equal to the pulse width plus a 'wake' of maybe several pulse widths. Fields on the surface will tend to be small outside this band, save perhaps for wholesale reflections off geometrical features. Such reflections will themselves, however, tend to take the form of an active band of a few pulse widths.

The main work at each time step in the regular time domain BIE approach is the integration over the surface to obtain the effect 'here and now' of historical (retarded) surface fields. There may be little benefit in integration over regions outside this active band, where the (historical) field was small. As suggested above, it seems likely that this active band may be of a width related to the pulse width; perhaps a modest multiple of it. In that case the

area of the surface over which it is necessary to integrate is reduced by a factor of order Θ, where Θ is the size of the body in wavelengths in the direction of propagation.

Now the methods to neglect all but this active band must identify some essentially arbitrary field value below which we treat the surface field as negligible, because naturally no part of any body is absolutely quiescent. This fact leads one to the idea of a threshold, below which the field is treated as zero. There are many ways of thresholding, but they should all aim to fulfil the requirement that at any time step, sufficient portions of the surface satisfy the thresholding criterion for appreciable saving to be possible without significant loss in overall accuracy. One easy and simple thresholding technique suggests itself; the threshold value, taken as constant throughout all time steps, is a certain fraction of the incident wave amplitude. This simple thresholding technique will be used in this thesis, and various alternative approaches will be discussed later.

The physical basis explained above may be viewed in some respects as lying in some sense between the conventional full field solution and the PO approximation. One of the short-comings of the PO approximation is its inability accurately to represent the current near the terminator (the illuminated region-shadow boundary). This can be a source of significant error because of the abrupt discontinuity in current which occurs. A typical way of circumventing this difficulty is to use a "pseudo" PO approximation which employs the integral equation to find the current in the shadow region. The above approximation can be an alternative approach. The

integral equations are employed both in illuminated and shadow regions, but more subtle approximation are made within the integral equation.

It needs to be ascertained that the pulse width is indeed related to the amount of cost saving. For a given body, as the pulses become shorter, pulsed illumination will tend to leave increasing fractions of the surface of the body quiescent. This means that possibly greater savings can be achieved by using narrower and narrower pulse widths. However, as mentioned in the previous chapter, the fineness of temporal and spatial discretisation needs to be sufficient to model the pulse accurately. The overall cost scalings still mean that it is not likely to be cost effective to mesh more finely than is needed to represent properly the highest frequency of interest.

3.3 Thresholding Results

In the section below, the results to illustrate the accuracy loss given by this neglect of all but 'small' fields are presented for various geometries. As mentioned earlier, simple thresholding technique is used; the threshold value is the quoted fraction of the incident pulse amplitude.

3.3.1 10:1 Dipole

The first results are for a 10:1 cylindrical dipole, with hemispherical ends. It was illuminated 20^0 degrees off-axis with a Gaussian pulse with the pulse width selected to allow extraction of results at relevant frequencies.

The bistatic radar cross section (RCS) is shown in Figure 3.3.1.1, extracted for a wavelength of 1/4.4 of the dipole length. The different lines on figure

3.3.1.1 correspond to different levels of threshold. There is a steady worsening of the result with increasing threshold, as would be expected. However, even the 5% and 10% threshold results are close to the normal BIE time domain result over most of the range, and for most practical purposes the 1% and 2% cases are probably indistinguishable from it.

Figure 3.3.1.2 shows the time variation of the surface field on-axis at the far end in the direction of travel of the pulse, for the same cases. The peak when the wave arrives, at timestep ~60, is obvious. The main wave then travels back, leaving the field at the end jittering close to zero. Most of this jitter is missed in the threshold cases, and indeed thresholding is such that the return wave seen at timestep 140 is supressed. If it were desired to capture this second bounce this is a perhaps a case where time dependent thresholding based on the present maximum field would be required.

3.3.2 Bent Dipole

The bistatic RCS of a hemispherically capped 10:1 aspect ratio dipole, with a 90^0 bend at its centre is calculated. The mesh used is shown in figure 3.3.2.1. This was illuminated by a half-sinusoid pulse, such that each arm of the dipole was 2 wavelengths long. It is a problem of modest size, and thus not one where dramatic savings are to be expected. For example, even if thresholding were to eliminate essentially all but the incident wave contribution, it would be expected that the active region in this case would be ~1/4 to 1/2 of the body surface during the first transit.

The calculation was performed with a range of thresholds, and the bistatic RCS for each case is indicated on figure 3.3.2.2. There is a steady increase in deviation from the 0% base case, although with even large 15% threshold, the main deviations involve only a deepening of troughs. Figure 3.3.2.3 shows the number (out of the total of 1094) nodes active versus time for three cases. This number of nodes active is a good indication of the potential cost saving. It is seen from figure 3.3.3.3 that as would be expected the number of nodes active declines and the body becomes quiescent more quickly as a larger threshold is used.

3.3.3 NASA Almond

The NASA almond[47] presents a more difficult test. This is one of a series of stealthy RCS benchmark targets, where things such as late time effects and travelling waves are significant contributors to the response. Here the 2.5 wavelength long version is analysed. The monostatic RCS is shown in figure 3.3.3.1 for thresholds of 1% and 5%. Also shown are unthresholded BIE time domain results, measured results (extracted from enlarged photocopies of the reference), and frequency domain results[49].

Both the normal BIE time domain and frequency domain results show reasonable agreement with the measurements, particularly bearing in mind both the experimental uncertainty and the additional errors introduced in making measurements from photocopies. The regular BIE time domain result is generally close to the frequency domain solution. More relevant for present purposes is the change with threshold in the results of the BIE computations; this is generally modest, even for the 5% case.

3.4 Discussion

In section 3.3, a relatively crude threshold was applied; a uniform, constant value expressed as a fraction of the incident pulse magnitude. After the first transit, and the incident pulse has left the body, fields tend to fall sharply. One approach to deal with this would be the specification of a time dependent threshold, expressing the threshold in terms of the present maximum field on the body.

For some geometries the field in the shadowed region may be small, but may play a large role in determining the scattered field. In such cases it may be good to apply position dependent thresholds, with a lower value in such generally low-field regions.

It is worth commenting here that the thresholding technique described above is not a 'near neighbour' approach, where remote regions are ignored. A remote portion of the surface could have been 'active' at the relevant retarded time, and would be included 'exactly', just as normal, while nearby portions of surface could have been quiescent, and integration over them would not be performed. A near neighbour approach could indeed be incorporated in addition, and this would perhaps be an interesting area for future work.

Figure 3.3.1.1   Bistatic RCS of 10:1 dipole : 4.4 wavelengths long case

Figure 3.3.1.2   Time dependent surface field on-axis, down stream end

Physical Basis of the Computational Cost Reduction   89

Figure 3.3.2.1   1094 bent dipole mesh

Figure 3.3.2.2   Bistatic RCS of bent dipole : 2 wavelengths long case

Figure 3.3.2.3   Number of nodes active versus timestep

Physical Basis of the Computational Cost Reduction   91

Figure 3.3.3.1   Monostatic VV RCS of NASA almond at 3 GHz

Chapter 4 Modification of the BIE Algorithm to Implement Cost Reduction

In the present chapter, a novel modification to the normal boundary integral equation (BIE) algorithm is presented. This algorithm is then combined with the approximation to the physics presented in the previous chapter, involving omitting integration over regions of the low field magnitude. Whilst this does error increases slightly, it also provides large reductions in cost, and a possible reduction in cost scaling to the fourth power.

Section 4.1 introduces some differences between the normal and modified algorithms. In section 4.2 the modified algorithm is derived algebraically and its numerical implementation presented. Section 4.3 describes its computational implementation, and a general discussion of the characteristics of the new algorithm is presented in section 4.4.

4.1 Introduction

The normal algorithm for the boundary integral equation (BIE) method in the time domain involves calculating the field at a node by integrating over all the rest of the surface, summing the effects of the historical field at the relevant retarded time (weighted by the appropriate integral of the Green function, embodied in the main system matrix coefficient). This is shown algebraically in equation (2.1.68) of Chapter 2. This can be described as a 'retrospective' approach, summing historical values relating to widely different times, once they are all known.

Here will be introduced an inversion of the normal algorithm, which will be termed 'project forward'. In the 'project forward' approach, once a field value is known, we can use it immediately, to begin to form *future* field values over the whole of the rest of the surface, simultaneously, for times up to one transit time in the future. For each other surface location, it is possible to determine the time in the future at which the field in question will have an effect, and the size of the effect. The field at each location and future time will then gradually be built up by contributions from various parts of the surface, at various times, until the current time itself becomes the future time under discussion. Computationally, this corresponds to a re-ordering of the loops implicit in equation (2.1.68). The situation is actually made more complicated by the implicit approach adopted, but is unchanged in principle.

At any time, we have now a set of future solutions for other times up to one transit time in the future. They will be partial solutions, with the temporally more remote ones less complete, and the nearer ones more complete. It is then no longer necessary to store the field history at each node; we have already made such use of it as we need. Its place is taken by this set of accumulating future surface fields.

The main work in the algorithm is now this projection forward of each surface field value as it is calculated. This is analogous to, and equal in effort to, the normal repeated surface integration of equation (2.1.68).

The above re-ordering of the algorithm is applicable to the 'normal' (i.e. unthresholded) approach, and displays the normal fifth power cost scaling.

However, once implemented the re-ordering does offer a relatively simple way to exploit the physical approximation, related to the small active area, discussed in the previous chapter. Once a field has been calculated, the normal action would be to project it forward, incurring the main cost of the time domain BIE approach. A lower cost, and possibly 'fourth power' approach is achieved if this projection forward is only done if the field was large enough to be deemed active when some arbitrary threshold is imposed.

4.2 Description of the Modified Algorithm

This description, and indeed the implementation of the reduced scaling version, are based on the curvilinear isoparametric treatment presented in chapter 2, but the modifications described are of general applicability. They could, for example, be applied similarly to a flat-facetted or wire grid discretisation.

As derived in chapter 2, for any point $\mathbf{r}$ on the smooth surface, the field is given by an integral over the rest of the surface ($\mathbf{r}'$ and s') of the history of the field there. For the field at some time t, we have

$$2\pi \mathbf{H}(\mathbf{r},t) = 4\pi \mathbf{H}_{inc}(\mathbf{r},t) + \int_{t'=0}^{t'=t} \int_{\partial\Omega} (\mathbf{n}' \times \mathbf{H}(\mathbf{r}',t')) \times \frac{\hat{\mathbf{R}}}{R^2} \delta(|t-t'|-|\mathbf{r}-\mathbf{r}'|) ds' dt'$$

$$+ \int_{t'=0}^{t'=t} \int_{\partial\Omega} \left(\mathbf{n}' \times \frac{\partial \mathbf{H}}{\partial t'}(\mathbf{r}',t') \right) \times \frac{\hat{\mathbf{R}}}{cR} \delta(|t-t'|-|\mathbf{r}-\mathbf{r}'|) ds' dt' \quad (4.2.1)$$

where we denote $\mathbf{r}' - \mathbf{r}$ as $\mathbf{R}$, and $|\mathbf{R}|$ as R.

The normal use of this equation involves finding present field values essentially by evaluating weighted summations of historical values. We will re-order this process. When we find the field at any timestep, we will immediately use the field value so found to begin to form future fields. Any field value will have some effect at every surface location at some time in the future; just how far into the future is a function of how far away the particular location is. We thus will have no need to store surface field histories, but instead will store accumulating partial sets of future surface fields.

We write (4.2.1) for some future time t_f, when the present time is t. We will shortly discretise time into timesteps of length Δt, and so will divide the range of time integration in (4.2.1) into the period up till Δt ago, the current timestep, and all future ones. Our concern at any one timestep will be only integration over the last period Δt, so denoting the past and future parts of the integral $H^P(r, t_f)$ and $H^F(r, t_f)$ respectively, we can write (4.2.1) as:

$$2\pi H(r, t_f) = 4\pi H_{inc}(r, t_f) + H^P(r, t_f) + H^F(r, t_f)$$

$$+ \int_{t'=t-\Delta t}^{t'=t} \int_{\partial\Omega} (n' \times H(r', t')) \times \frac{\hat{R}}{R^2} \delta(|t_f - t'| - |r - r'|) ds' dt' \quad (4.2.2)$$

$$+ \int_{t'=t-\Delta t}^{t'=t} \int_{\partial\Omega} \left(n' \times \frac{\partial H}{\partial t'}(r', t') \right) \times \frac{\hat{R}}{R} \delta(|t_f - t'| - |r - r'|) ds' dt'$$

We will now address the discretisation of (4.2.2).

The surface is divided into M elements m, over which the geometry is approximated using polynomial shape functions $S_\alpha(\xi, \eta)$ where ξ, η are the

parameterised spatial co-ordinates. We will here use 9 noded quadrilaterals, but in principle a wide range of element types may be employed. The geometry of each element is thus described by

$$\mathbf{r}^m(\xi,\eta) = \sum_{\alpha=1}^{9} S_\alpha(\xi,\eta)\mathbf{r}_{j(m,\alpha)}$$ (4.2.3)

where $j = j(m,\alpha)$ are the global node numbers of the local nodes on element m, and $\mathbf{r}_{j(m,\alpha)}$ is the position vector of each of these nodes.

The vector from the ith spatial node to some (ξ,η) co-ordinate location on the mth spatial element is then

$$\mathbf{R}(\mathbf{r}_i;\xi,\eta)\big|_m = \mathbf{r}_i - \sum_{\alpha=1}^{9} S_\alpha(\xi,\eta)\mathbf{r}_{j(m,\alpha)}$$ (4.2.4)

Using an isoparametric formulation, the spatial variation of the surface magnetic field over an element m may be written

$$\mathbf{H}(\xi,\eta;t)\big|_m = \sum_{\alpha=1}^{9} S_\alpha(\xi,\eta)\mathbf{H}_{j(m,\alpha)}(t)$$ (4.2.5)

We model the temporal variation using quadratic elements of length $2\Delta t$, with associated basis functions $T_\beta(\tau)$, with τ the intrinsic time. The temporal variation of the magnetic field over a temporal element is thus

$$\mathbf{H}_{j(m,\alpha)}(\tau)\big|_l = \sum_{\beta=1}^{3} T_\beta(\tau)\mathbf{H}_{j(m,\alpha)}^{k(l,\beta)}$$ (4.2.6)

where $k = k(l,\beta)$ are the global timestep numbers of the local temporal nodes on temporal element l, and $\mathbf{H}_{j(m,\alpha)}^{k(l,\beta)}$ are the field values at the spatial nodes, at

the three timesteps of the temporal element within which the time of interest falls.

We thus approximate the field at some intrinsic location ξ, η within spatial element m, at an intrinsic time τ within temporal element l, as

$$\left. H(\xi, \eta, \tau) \right|_{m,l} = \sum_{\alpha=1}^{9} \sum_{\beta=1}^{3} S_{\alpha}(\xi, \eta) T_{\beta}(\tau) H_{j(m,\alpha)}^{k(l,\beta)} \tag{4.2.7}$$

From these representations, all the various quantities needed (normals, tangents, various spatial derivatives, surface divergence, Jacobian J, and so on) may readily be obtained.

We return to consideration of (4.2.2). Because of the delta function, for any location r there will be only one future timestep k_f which will be influenced by the field at a particular location r' during the timestep from t_{k-1} to t_k over which we are integrating. This future timestep is given by

$$k_f = k + \text{int}\left[\frac{r - r'}{\Delta t}\right] \tag{4.2.8}$$

This field at k_f will be influenced by the field at an intrinsic time τ^* in the present time element, given by:

$$\tau^* = 1 - \left(\left[\frac{r - r'}{\Delta t}\right] - \text{int}\left[\frac{r - r'}{\Delta t}\right]\right) \tag{4.2.9}$$

However, note that as r' varies over an element, the future timestep k_f and the intrinsic time both in general vary over a single element m. Since the support of the shape function associated with any one node extends over

several elements, the field at a given boundary node j will affect the future field at a given field node i for a handful of future timesteps. Conversely, or equivalently, the field at any future timestep is affected by the field at a handful of earlier timesteps.

We see the timestep k and (intrinsic) time τ^* identify what would be the retarded time at $\mathbf{r'}$ as seen from the future time kf at $\mathbf{r}$.

Inserting our shape function representations into (4.2.2), making use of (4.2.8) and (4.2.9), we obtain

$$2\pi \mathbf{H}_i^{k_f} = 4\pi \mathbf{H}_{inc,i}^{k_f} + \mathbf{H}_i^{Pk_f} + \mathbf{H}_i^{Fk_f}$$

$$+ \sum_{m=1}^{M} \iint_{\xi\eta} \sum_{\alpha=1}^{9} S_\alpha(\xi,\eta) \sum_{\beta=1}^{3} \left(\frac{T_\beta(\tau^*)}{R^3} + \frac{\dot{T}_\beta(\tau^*)}{\Delta t R^2} \right) \left[\left(\mathbf{n'} \times \mathbf{H}_{j(m,\alpha)}^{k+\beta-3} \right) \times \mathbf{R} \right] J(\xi,\eta) |d\xi d\eta \qquad (4.2.10)$$

Writing $(\mathbf{n'} \times \mathbf{H}) \times \mathbf{R}$ as $[A']\mathbf{H}$, with the matrix $[A']$ given by:

$$[A'] = \begin{bmatrix} n_2'R_2 + n_3'R_3 & -n_1'R_2 & -n_1'R_3 \\ -n_2'R_1 & n_1'R_1 + n_3'R_3 & -n_2'R_3 \\ -n_3'R_1 & -n_3'R_2 & n_1'R_1 + n_2'R_2 \end{bmatrix} \qquad (4.2.11)$$

and integrating via Gaussian quadrature (with singular and hypersingular integrals treated essentially as in chapter 2), yields

$$2\pi \mathbf{H}_i^{k_f} = 4\pi \mathbf{H}_{inc,i}^{k_f} + \mathbf{H}_i^{Pk_f} + \mathbf{H}_i^{Fk_f}$$

$$\sum_{m=1}^{M} \sum_{p=1}^{NG} \sum_{q=1}^{NG} \sum_{\alpha=1}^{9} S_\alpha(\xi_p,\eta_q) \sum_{\beta=1}^{3} \left(\frac{T_\beta(\tau(R))}{R^3} + \frac{\dot{T}_\beta(\tau(R))}{\Delta t R^2} \right) [A'] \mathbf{H}_{j(m,\alpha)}^{k+\beta-3} J(\xi_p,\eta_q) w_p w_q \qquad (4.2.12)$$

Thus far we have implicitly allowed the field $\mathbf{H}$ to have three (say) Cartesian components, such that integrations of matrix $[A']$ would result in 3 by 3

matrices. As mentioned in chapter 2, application of the perfect conduct boundary conditions permits one of the three H components to be expressed as a linear combination of the other two, in turn reducing the results of the integrations to 2 by 2 (sub)matrices.

Denoting as α_{ij} the now 2 by 2 (sub)matrix which results from integration of the matrix $[A']$ for each i,j node pair, we can write (4.2.12) in the form

$$2\pi \mathbf{H}_i^{k_f} = 4\pi \mathbf{H}_{inc,i}^{k_f} + \mathbf{H}_i^{Pk_f} + \mathbf{H}_i^{Fk_f} + \sum_{j=1}^{N} \alpha_{ij}^{(k)} \mathbf{H}_j^k + \alpha_{ij}^{(k-1)} \mathbf{H}_j^{k-1} + \alpha_{ij}^{(k-2)} \mathbf{H}_j^{k-2} \qquad (4.2.13)$$

where we have indicated explicitly that there will be one such submatrix associated with each of the three timesteps k, k-1 and k-2, associated with the present temporal element.

This equation is evaluated at each timestep, for each node, for each of the future timesteps associated with each i,j node pair. For the particular case of k_f being the present timestep, we have naturally no more future contributions to add. The terms involving field values at timesteps k-1 and k-2 are multiplied out and added to the vector on the right hand side, to which is added also the incident wave vector, giving a right hand side vector c, say. Those terms α involving the present timestep are brought to the left hand side, and together with the free term 2π, form a matrix equation for the new field value at each node:

$$[a]\mathbf{H}^k = [c] \qquad (4.2.14)$$

This matrix equation is solved at each timestep to determine the new field values. The actual solution cost is low, as the matrix in (4.2.14) is very sparse, and is a tiny fraction of the total computational work. With the new field found, its future effects are accumulated via (4.2.13).

The dominant work is in these repeated 'projection forward' activities involved in (4.2.13). At each timestep, the effect of the field at every node, at every other node (at some time in the future) is calculated. This has costs which scale with nodes squared (i.e. frequency to the fourth power) at each timestep. With the number of timesteps generally scaling with frequency, this gives an overall fifth power cost scaling. This scaling, and indeed the actual computational work, is the same as in the conventional 'retrospective' approach to time domain integral equation calculations.

4.3 Computational Implementation of the Method

In the project forward algorithm, there are two main cost components in terms of storage; the main system matrix, the collection of submatrices α of (4.2.13), and the future right hand side vectors. For a given number of nodes N, which scales with the second power of frequency, $2N$ components of each right hand side vector need to be stored, where the number of future right hand side vectors is the number of time steps in a single transit. Thus the storage requirement for future right hand side vectors scales with the third power of the frequency.

Each component of the main matrix has 2×2 entries, and for each boundary point / field point pair there are from 3 up to about 7 time step 'layers'.

Therefore, the number of entries of matrix is up to about $28 \times N^2$, a quantity scaling with the fourth power of the frequency. This matrix storage is the dominant storage requirement, and as the frequency of the incident wave increases it rapidly becomes very large.

If the storage requirement for a problem is not greater than the RAM available on the computer being used, the matrix is best calculated in advance and kept in core. As shown in section 2.5, therefore, the cost for matrix formation is a trivial fraction of the total work. If the storage requirement for a problem is greater than that available, one approach is to form the matrix, write it to disk, and read it for each time step. On a single processor machine this incurs only a modest penalty, particularly if multiple illumination angles are required, allowing the cost of reading to be shared amongst the look angles. However, parallel computation gives a much greater increase in computational power in the processing stage than in the input / output stage. (Additionally, although it is not very costly, additional disk storage is needed.) Thus repeated reading from disk is not attractive.

The alternative approach is to form the matrix when it is needed. Normally, the entire matrix is formed only once at each time step because the matrix is only geometry dependent, even if multiple illumination angles are required. Thus it can be seen that the penalty associated with the inability to store matrix decreases as the frequency of incident wave increase, due to the fact that the number of angles required to characterise the response rises with frequency. If this penalty is accepted, the whole storage requirement is greatly relieved. The dominant storage becomes the future right hand side

vectors. Thus this approach allows in practice a vast reduction in the amount of storage required.

For the above approach, the matrix formation time becomes dominant when a single illumination angle is required, and remains so unless a very large number of illumination angles is required. The project forward algorithm offers the opportunity to save matrix formation time for this case. When the field is negligible, the formation of the matrix entry which will be multiplied by this field value can be omitted. This is naturally only applicable for a single illumination angle because the active region depends on the illumination angle. If multiple illumination angles are required it is likely that all the matrix will be used by one or other angle, so formation of the entire matrix prior to any 'projections forward' is probably simplest. However note that the right hand side vector formation, the projection forward activity itself, again becomes dominant for multiple illumination angles, thus recovering the cost saving.

The project forward algorithm has here been implemented in two different ways. If the matrix fits in core, it is formed once in advance, and used in the formation of right hand side vector $[c]$ at each time step. This we will term a 'not on the fly' calculation. If it does not fit in core, the matrix entries are calculated when they are required, termed here an 'on the fly' calculation. It is explained below how the modified algorithm is implemented in each case.

4.3.1 'Not on the fly' Calculation

As shown in chapter 2, there are three main stages to obtaining the RCS response at a given frequency in the time domain; pre-processing, surface field calculation, and far field (RCS) calculation. The only change which needs to be made is to this surface current calculation stage. The system matrix itself is formed exactly as usual. The project forward algorithm may be thought of at its simplest as a re-ordering of the process of right hand side formation in the normal algorithm. For each time step, the surface field is obtained by solving the matrix equation, and then at each location the magnitude of surface field is compared to threshold values. If the surface field is negligible, the multiplications associated with using this field value for forming future right hand side vectors are simply omitted.

4.3.2 'On the fly' Calculation

The 'on the fly' version of the project forward algorithm is necessarily implemented in a rather more complicated way. There are two main computational issues in this version; formation of the [a] matrix in equation (4.2.14) and finding the elements which should be integrated over at each time step. These are investigated below.

Generally the [a] matrix is very sparse, and manipulations and storage involving it is relatively cheap, as it scales in size with frequency squared. The storage cost involved in the [a] matrix are thus not important at all, and here it is formed at once in advance and kept in core. It is nonetheless worth noting that for modest sizes of scatterer (say below ~10 wavelengths long) which also are locally (very) refined, the storage for the [a] matrix could be important and sometimes exceeds that of the future right hand side vectors

which scales with frequency cubed, because of its higher base. If this is so, to overcome this difficulty, the [a] matrix entries may be calculated easily when they are required.

For each time step, the active elements, that is those elements which include active nodes where the surface field is not negligible, must be found. These are then integrated over, from every other field point, and the result multiplied by the surface field value to form the future right hand side vectors. Here, for the efficient finding of the active elements, firstly the list of active elements which each active node lies on is calculated. Normally this list has many duplicated elements, and thus secondly it is rearranged and entries made unique by using a sorting algorithm.

4.4. Discussion

As the pulse becomes shorter, pulsed illumination will tend to leave increasing fractions of the surface of a body quiescent. Exploitation of this has been identified as a means to reduce the cost and possibly cost scaling of time domain integral equation scattering analyses. A novel 'project forward algorithm' has been devised and implemented to achieve just such exploitation. The approach becomes easier to implement, and the cost savings closer to their asymptotic fourth power scaling, as the bodies become bigger, which will be shown in chapter 6.

In particular, in the case of 'on the fly' calculation, the coefficients of α constitute in effect a large matrix of size (several times) nodes by nodes, and scaling with frequency to the fourth power in storage requirements. Only a

fraction $1/f$ of this is needed at each timestep, and in practice we choose to re-evaluate this portion of the matrix as it is needed, rather than store it all. Storage requirements otherwise scale with frequency to the fourth power. This approach incurs a one-off operation count increase by a factor of perhaps 5, depending on the order of quadrature adopted. The benefit is that the remaining storage is only of the future surface fields; a much more manageable cost, scaling with frequency to the third power.

Chapter 5 Parallelisation of the Project Forward Algorithm

Large scattering computations by any technique are very expensive, and the computational costs of all approaches rise sharply with frequency. One way to begin to meet this computational requirement is via the use of massively parallel computers. In this chapter the parallelisation of the project forward algorithm is developed, for implementing on distributed memory multiple instruction multiple data (MIMD) computers.

In section 5.1 a few general topics associated with parallel computation are introduced. Section 5.2 reviews the (very few) prior works on time domain BIE parallelisation, and section 5.3 develops the parallelisation of the project forward algorithm. The differences between parallelisation of the normal 'retrospective', and the project forward algorithms are explained in this section.

The matrix solver, used at each timestep to solve the sparse matrix equation, is parallelised in section 5.4. In particular, the conjugate gradient squared (CGS) method is considered here.

Section 5.5 describes the particular computational implementation on a CRAY T3D, and finally the entire parallelisation activity is discussed in section 5.6 .

5.1 Parallel Computation

In this section are introduced some of the different kinds of parallel computers, and the approaches involved in the parallel computation.

5.1.1 Architecture of Parallel Computers

Loosely speaking, a parallel computer has the ability to perform several different processes at the same time. The architecture of parallel computers falls into two classes:

(a) Single-instruction multiple data systems (SIMD) such as ICL DAP and the Connection Machine, where each processor executes the same instruction simultaneously using different data, and

(b) Multiple-instruction multiple data systems (MIMD), where each processor executes possibly different instructions on different data.

There are two different types of MIMD computers: shared memory and distributed memory. In a shared memory machine such as CRAY X-MP, CRAY Y-MP and CRAY 2 series of computers, each processor can access a common memory, while in a distributed memory machine such as the MEIKO Computing Surface and the CRAY-T3D, each processor has its own memory which cannot be accessed directly by the other processors (although it is actually possible to do this to a degree on the T3D). With distributed memory machines, data and instructions must be passed between processors using communication links and the outputs from all processors must be combined to yield the final result.

5.1.2 Parallel Implementation

Currently, there are three major communication standards or libraries which support parallel computing: High Performance Fortran (HPF),

Parallel Virtual Machine (PVM) and Message Passing Interface (MPI). Which one is the best choice for parallel computing is still very much a matter of debate. They are briefly introduced below.

There have been several extensions of Fortran 77, and its sucessor Fortran 90, for parallel computing. High Performance Fortran (HPF), is one of the latest[50]. It is an extension of Fortran 90 to include data distribution features, data parallel execution features and extended intrinsic functions to be used with parallel machines. It is still in a very early stage of development.

PVM is a communication library and a programming environment for parallel programming on various machines including 'virtual' parallel machines. It is can be used with Fortran and C. It is general in nature, but slow.

Over the last ten years numerous message passing systems have been created all using very similar concepts with some variation. Among them, MPI arguably has taken and combined the best features of all present message passing systems including PVM, PICL, and Zipcode. The development of MPI has been so fast that MPI-2 was in preparation while as far as the author is aware, only one book on MPI-1 was published[51] very recently. In this thesis, the MPI-1 standard is used with Fortran 77 for parallel implementation.

5.1.3 Parallel Performance Metrics

There are many ways to measure the performance of a parallel algorithm. Among them, the conventional measure of speedup is the most widely used. This is defined as;

$$Speedup(N) = \frac{T(1)}{T(N)} \qquad (5.1.3.1)$$

where T(1) is the execution time using 1 processor, and T(N) is the execution time using N processors

The efficiency of algorithm is then defined as

$$Efficiency(N) = \frac{Speedup(N)}{N} = \frac{T(1)}{N \bullet T(N)} \qquad (5.1.3.2)$$

This measure of efficiency gives an indication of how well the application utilises the processors available, although for the same algorithm it can generally be improved by increasing the size of the problem solved.

5.2 Parallelisation of the Time Domain Integral Equation Method

Parallel computers have been applied very little to BIE techniques relative to say finite element or finite difference approaches, and the large majority of such work as there is addresses frequency domain or steady state problems[52,53,54,55]. There has been very little attention paid to parallelisation of transient BIE treatments. The only exceptions are the works of Walker and Leung[56,57] where acoustic and electromagnetic scattering problems are solved on distributed memory systems; a Meiko surface and the T3D. Their main concern is the parallelisation of the right hand side formation stage. For the wholly 'in-core' cases they addressed, this is by far the dominant cost,

and becomes an increasing fraction of total cost as larger problems are addressed. Various different approaches for this are investigated in their works, and these are summarised below.

The first approach is the allocation of a sequential set of field nodes to each processor. For each field node is then performed integration over all the surface patches. Both for this and for right hand side vector formation, there is almost exactly equal computation to be performed for each field node, allowing very good load balancing. At the end of this activity, each processor has generated a portion of the required right hand side, and these partial vectors are then communicated to a single processor for solving the matrix equation. This approach may be the most simple to implement, but it has significant drawbacks. The history of the field over the whole surface needs to be stored on each processor. Whilst it is not of itself the dominant storage cost, if the history is duplicated in this way it quickly becomes intolerably large. Thus this approach was not in practice used in their work, despite being so easy to implement.

The second approach is domain decomposition, where a major objective of the approach is to avoid storing (i.e. duplicating) the history of the all the field on each processor. Firstly, an allocation of elements to processors is made, dividing the domain (surface) between processors. Each processor then integrates over 'its' elements from every field node. Matrix formation via this approach exhibits essentially perfect parallelisation. However, its disadvantage is that nodes on element edges are dealt with separately by two

processors in the right hand side vector formation stage, leading to some duplication of effort.

In their third approach, a sequential set of boundary nodes is allocated to each processor. Each processor then integrates over its portion of the surface from every field node. This led to some inefficiencies in the matrix formation stage, because computations within a given element are repeated by all processors allocated a node lying in that element. However, it did optimise the dominant right hand side formation activity. Thus the quality of the parallelisation is not significantly affected by the inefficiency in matrix formation stage. Walker and Leung adopted this third approach.

5.3 Parallelisation of Project Forward Algorithm

As explained in an earlier chapter, when the storage requirement for a problem is greater than the storage available, forming the matrix at each time step is the more efficient approach in parallel computation than is the repeated reading of the matrix from a disk. Once this operation cost is accepted the whole storage requirement on each processor can be greatly relieved. This is probably a sensible course even for the conventional approach, and is certainly so for the reduced cost version, where only a small part of the matrix is used at each timestep. Consequently, this approach has been adopted here.

In parallel computation, in general the most important criteria are that the duplicated work or storage, and the communications between the processors, should be minimised. In terms of work and communications,

partitioning over field nodes seems to be the most attractive, but as described above the consequent duplication of storage renders it unsuitable for the conventional algorithm.

A significant consequence of the project forward algorithm developed here is that it does however allow this most attractive approach to be used. As noted earlier, the new algorithm replaces storage of the history of the field values with storage of the future right hand side vectors. The former are quantities associated with the 'boundary elements'; the latter with the field nodes. Thus when a set of field nodes are allocated to each processor, each processor need store only future right hand sides vectors for its 'own' set of nodes.

Our objective in the parallel implementation is to solve as large as possible a problem. As it is often the amount of storage which a processor has which limits the size of problem which can be solved, the efficient use of this storage is crucial to solve larger and larger problems.

The first main storage considered in this work is the [a] matrix (see equation (4.2.14)) which scales in size with frequency squared. (Recall, it is a sparse matrix, with as many rows as there are unknowns, with each row comprising a more or less fixed number of non-zero coefficients). For a large problem, the relevant part of the [a] matrix (that is, those rows associated with the nodes allocated to that processor) should be allocated to each processor. This requires that the matrix solution stage be parallelised, although viewed in isolation the matrix solution is not costly at all in operations. Each processor is given a portion of the [a] matrix, and performs

a subset of the matrix-vector multiplications involved in the matrix solution. Details associated with this will be explained in the next section.

The method adopted here requires the whole surface (all elements) to be integrated over 'from' a given set of field points on each processor. Thus all the mesh information required for element integration, such as the Cartesian co-ordinates of Gaussian locations, and the Jacobians at Gauss points, on every element, should be duplicated on every processor. This storage requirement, whilst intrinsically relatively modest and scaling with only the square of frequency, may eventually limit the size of problem due to this duplication over processors. If it is, one approach is wherever possible to calculate these quantities whenever they are needed. For large problems such those as which actually require parallel computation, this recalculation is a trivial fraction of the operations required for the whole calculation. This approach is adopted here.

5.4 Parallelisation of the matrix solver

As has been explained above, the actual computational time for repeated matrix solution is a small fraction of total, but the matrix solution stage should be parallelised to permit the allocation of the relevant part of the [a] matrix to each processor. Here the conjugate gradient squared (CGS) algorithm[58] has been parallelised.

The outline of the CGS algorithm and its parallelisation algorithm are shown in figures 5.4.1 and 5.4.2 respectively. Typically the matrix-vector mulitplication is the main calculation in iterative solvers, and indeed in

CGS each iteration involves two matrix-vector multiplications. Before performing those multiplications, the partial-length vector on each processor should thus be assembled into a full-length vector and sent to every processor because the matrix has full-length "rowise". The actual multiplication generates again a partial-length vector on each processor.

In the domain decompostion approach, the vector multiplied by the system matrix is partial, but is of full length. (That is, each coefficient is non-zero, but is of magnitude only a fraction of what it will be once agglomeration of all such partial vectors has been performed.) The consequence is that the domain decomposition approach requires much more data to be communicated than does the present approach, and is consequently less efficient in the solution stage too.

5.5 Computational Implementation of the Method

The method of the allocation of sequential sets of field nodes is applied to the project forward algorithm, and thus each processor performs the calculations involving its own set of field nodes. At every time step each processor forms its own part of the matrix and builds its own part of the future right hand side vectors. It is obvious that no communication is required in those stages.

The only stage requiring some communication is the matrix solution which is performed at every time step. For example, the residuals calculated on each processor must be added together after each iteration to compare with

the stopping criterion, and the solutions must be assembled for the output. Details are explained in section 5.4.

5.6 Discussion

The approach of the sequential sets of field nodes is very attractive because it does not require any communication between processors at the two main calculation stages (the matrix and right hand side vector formations), and duplicated calculations are performed only for the element-related mesh information, which is almost negligible for the large problems which require parallel computation.

The approach described in this chapter has been successfully implemented with the project forward algorithm. As expected, it exhibits good performance both for the matrix and right hand side formations; results demonstrating this will be shown in the next chapter.

$$\mathbf{x}_0 = \mathbf{x}_{previous\_time\_step}$$

$$\mathbf{r}_0 = \mathbf{b} - \mathbf{A}\mathbf{x}_0$$

$$\tilde{\mathbf{r}}_0 = \mathbf{r}_0$$

$$\mathbf{p}_0 = 0$$

$$\mathbf{q}_0 = 0$$

$$\rho_0 = 0$$

if $\|\mathbf{r}_0\| \le \varepsilon$ go to 20

do 10 $k = 1,2,...,MAXITER$

$$\rho_1 = \tilde{\mathbf{r}}_0^{\mathsf{T}} \cdot \mathbf{r}_{k-1}$$

$$\beta = \frac{\rho_1}{\rho_0}$$

$$\mathbf{u}_k = \mathbf{r}_{k-1} + \beta\mathbf{q}_{k-1}$$

$$\mathbf{p}_k = \mathbf{u}_{k-1} + \beta(\mathbf{q}_{k-1} + \beta\mathbf{p}_{k-1})$$

$$\mathbf{v}_k = \mathbf{A}\mathbf{p}_k$$

$$\alpha = \frac{\rho_1}{\tilde{\mathbf{r}}_0^{\mathsf{T}} \cdot \mathbf{v}_k}$$

$$\mathbf{q}_k = \mathbf{u}_k - \alpha\mathbf{v}_k$$

$$\mathbf{r}_k = \mathbf{r}_{k-1} - \alpha\mathbf{A}(\mathbf{u}_k + \mathbf{q}_k)$$

$$\mathbf{x}_k = \mathbf{x}_{k-1} - \alpha(\mathbf{u}_k + \mathbf{q}_k)$$

if $\|\mathbf{r}_k\| \le \varepsilon$ go to 20

10   end do

Convergence criterion was not satisfied

20   continue

$$\mathbf{x} = \mathbf{x}_k$$

Figure 5.4.1   Conjugate Gradient Squared (CGS) algorithm

$$\mathbf{x}_0 = \mathbf{x}_{previous\_time\_step}$$

$$\mathbf{r}_0 = \mathbf{b} - \mathbf{A}\mathbf{x}_0$$

$$\tilde{\mathbf{r}}_0 = \mathbf{r}_0$$

$$\mathbf{p}_0 = 0$$

$$\mathbf{q}_0 = 0$$

$$\rho_0 = 0$$

Summing of the partial residual, $\mathbf{r}_0$

if $\|\mathbf{r}_0\| \le \varepsilon$ go to 20

do 10 $k = 1, 2, \ldots, MAXITER$

$$\rho_1 = \tilde{\mathbf{r}}_0^{\mathbf{T}} \cdot \mathbf{r}_{k-1}$$

$$\beta = \frac{\rho_1}{\rho_0}$$

$$\mathbf{u}_k = \mathbf{r}_{k-1} + \beta \mathbf{q}_{k-1}$$

$$\mathbf{p}_k = \mathbf{u}_{k-1} + \beta(\mathbf{q}_{k-1} + \beta \mathbf{p}_{k-1})$$

Assembling $\mathbf{p}_k$ and broadcasting it to all processors.

$$\mathbf{v}_k = \mathbf{A}\mathbf{p}_k$$

$$\alpha = \frac{\rho_1}{\tilde{\mathbf{r}}_0^{\mathbf{T}} \cdot \mathbf{v}_k}$$

$$\mathbf{q}_k = \mathbf{u}_k - \alpha \mathbf{v}_k$$

Assembling $\mathbf{u}_k + \mathbf{q}_k$ and broadcasting it to all processors.

$$\mathbf{r}_k = \mathbf{r}_{k-1} - \alpha \mathbf{A}(\mathbf{u}_k + \mathbf{q}_k)$$

Summing of the partial residual, $\mathbf{r}_k$

$$\mathbf{x}_k = \mathbf{x}_{k-1} - \alpha(\mathbf{u}_k + \mathbf{q}_k)$$

if $\|\mathbf{r}_k\| \le \varepsilon$ go to 20

10   end do

Convergence criterion was not satisfied

20   continue

$\mathbf{x} = \mathbf{x}_k$

Assembling $\mathbf{x}$ and broadcasting it all processors

Figure 5.4.2   Parallelisation of CGS algorithm

Chapter 6 Results of the Modified Algorithm

In chapter 3, the accuracy loss occasioned by the neglect of all but 'small' fields has been investigated. In this chapter will be investigated the reductions in cost and possible reduction in cost scaling to the fourth power provided by that neglect. Many results will be shown for a variety of geometries by implementing on workstations and CRAYT3D.

6.1 Accuracy, Cost and Thresholds

The speed of computers has increased very rapidly even during the performing of this work, and varies widely across the computers on which these examples were run. As a consequence, making time comparisons is a difficult problem, particularly as the absolute cost of the runs performed varies by more than an order of magnitude.

In this section is presented a pragmatic measure of cost, expressed in terms of the number of calls of the principal subroutine of the code, and later costs and cost scalings will be presented in terms of this.

6.1.1 Spheres 1.3 and 4 wavelengths in diameter

As a first example, we consider two very different size problems; two spheres, of 1.3 and 4 wavelengths in diameter.

Meshes comprising 674 and 6146 nodes were used, with Gaussian pulse illumination. This is not a particularly attractive geometry for demonstrating

the cost reduction approach, but does allow us to compare with analytical solutions.

Table 6.1.1.1 shows the fraction of the CPU time expended on each subroutine in the calculation. The routine for the integration over an element, and the return of a set of weights (the coefficients α of chapter 4), IPEC, is evidently the dominant cost. Further, the associated routines IPECNSQ9, IPECSQ9 and ELIMAT are each called once per call of IPEC. Thus the number of calls of IPEC characterises at least 96.25% and 96.93% of the total CPU time for the job in the 674 and 6146 node cases respectively. Additionally, these calls are known to be a part of the computation which has the highest, characteristic cost scaling, with these percentages then increasing as larger bodies are analysed. For comparisons between cases we thus conclude that the number of calls of IPEC should be a robust, cross-platform, measure of the computational cost of a run, and it will be used in this thesis.

Figure 6.1.1.1 and 6.1.1.2 shows the bistatic RCS of 1.3 and 4 wavelength diameter spheres. With the spheres each of unit radius, the pulse widths were 0.483 and 0.168 respectively. Maximum nodal separation on the two meshes was 0.107 and 0.112, with corresponding timesteps of 120 and 200. In each case the frequency extracted corresponds to a nodal separation of ~1/10 of a wavelength, typical of the relationship between pulse width, mesh parameters and frequency we employ.

Good agreement between analytical and calculated data is found with a zero threshold, with a steady degradation of the results up to the maximum 5% threshold employed. Figures 6.1.1.3 and 6.1.1.4 show the number of IPEC calls versus different levels of threshold for the two cases. It is seen that though, as noted, the sphere is not an ideal geometry to demonstrate cost savings, and these problems are not really large enough to lead us to expect particularly dramatic results, the cost reduction method performs reasonable well. The 5% threshold provided a cost savings by factors of ~4 and ~6 for the small and larger sphere cases respectively with very slight loss of accuracy.

6.1.2 NASA Almond

We now show results on the NASA almond[47] of chapter 2, using various different levels of threshold.

The almond mesh employed has 690 nodes, and is of length 9.936 units, with a pulse width of 1.925 used. Maximum nodal separation is 0.055, corresponding ~1/18 of the wavelength to be extracted. This wavelength makes the almond one wavelength long, and corresponds to a frequency of 1.19GHz, one of the frequencies at which the RCS of the almond has been measured.

The bistatic RCS for vertical polarisation(VV) is shown in figure 6.1.2.1, for each level of threshold. The number of calls of IPEC for each case is shown in figure 6.1.2.2. There is seen to be some degradation of accuracy with threshold, but for say even the 10% case results are still very little different

from the unthresholded case, with for example the head-on RCS perturbed by about 0.02 dB sq m.

This is a relatively small problem, and large cost saving is not expected. However, over the range of thresholds employed, cost savings by up to a factor of 11 are gained.

We now consider a second, larger almond example. The mesh employed now has 3266 nodes, with a pulse width of 0.805. Maximum nodal separation is now 0.1, corresponding to 1/10 of the wavelength to be extracted. This wavelength makes the almond ~4.3 wavelengths long.

A single illumination from 20^0 degrees off-axis was performed. The 1% threshold in this case modifies the backscatter RCS trivially, from -38.9 to -38.7 dB sq m (although it should be said that there may well have been found larger differences at some angles had the full monostatic calculation been performed) whilst costs are reduced by a factor of ~9 relative to the conventional BIE time domain.

6.2 Cost Scaling

On small problems where the pulse width is similar to the body length, no cost reduction is expected. As bodies become larger, and specifically larger in the direction of travel of the incident wave, the factor by which costs are reduced is expected to increase roughly in proportion to the body size itself.

An investigation of this will now be made for the dipole of section 3.3. The dipole is of unit radius, is 22 units long in total, with hemispherical ends. A number of meshes are employed, ranging from 674 to 6034 nodes.

We will use head-on illumination. This is then a rather idealised geometry and illumination, and if the method does have any promise, it should be apparent in such a case.

We will be considering a range of frequencies, with our primary interest being in the computational cost. However, first we show results for one case, computed both with the basic code, and with the 1% threshold we will employ in the rest of this section. Figure 6.2.1 shows the bistatic RCS for the 6.53 wavelength long dipole (1218 node mesh, pulse width 1.5, average nodal separation 0.337) for these two cases. Whilst acceptable accuracy loss is naturally problem dependant, the degradation is small.

We now consider scattering from a series of such dipoles, of increasingly refined mesh, illuminated with shorter and shorter pulses. In all cases this 1% threshold is employed, along with corresponding calculations (where practicable) using the basic code. Our interest is to investigate the degree of cost saving, and in particular the relationship between cost and body size / frequency. As discussed earlier, here 'cost' will be expressed in terms of the number of calls of the routine IPEC.

Figure 6.2.2 shows the variation of computing cost with body size, along with results for the normal BIE time domain approach. Here the body size is characterised in terms of the wavelength of the highest frequency which

could be extracted from the run, using the usual criteria regarding mesh spacing, timestep and pulse width.

The reduction in cost is obvious. To provide some more practical indication of cost, the 6.53 wavelength case (1% threshold), showing a cost of 1.06×10^7 IPEC calls, actually took about 1 hour to run on a DEC alpha workstation. The number of IPEC calls for the conventional BIE time domain approach can be easily computed, and is shown for cases up to 14.7 wavelengths long on the figure 6.2.2. However, actual runs using the conventional approach (0% case) were only performed up to a ~10 wavelength long case. This took ~50 hours on a DEC alpha workstation. Using the lower cost approach, the ~10 wavelength case required only ~8 hours, a factor of ~6 smaller. Using the conventional approach the ~15 wavelength case would have required about $(15/10)^5$ times as long; i.e. about 400 hours. Using the reduced cost approach, the actual time required was 35 hours, a factor of ~11 smaller.

Plotting logarithmically allows cost scalings to be estimated, and this is done in figure 6.2.3. The 'basic code' scaling line on the figure indicates a cost scaling with frequency to the power, ~5.1. The other line shows the threshold results. Fitting to the observed results indicates a cost scaling with frequency to the power ~4.3.

It should be noted here that it is difficult to show a single definitive cost scaling with the frequency, because some physical phenomena, such as creeping waves (which depend on the shape of the target geometry) may affect the active band which is involved in the cost reduction method. Additionally, if in some regions the mesh fineness is determined by

geometrical considerations the simple, expected scaling behaviour will not be observed. For example, the coarser dipole meshes (up to 8.23 wavelengths long) used here are indeed locally refined at the cap ends. If however we observe the last three points on the line of thresholded results, which are relatively little affected by the local refinement, a cost scaling with frequency to the power 3.96 is observed.

6.3 Computational Work, and Active and Quiescent Regions

It is of interest to examine more closely the extent of quiescence and so on; even such a simple shape as this, the interactions and cancellations are actually rather subtle. Recall that at each timestep we project forward only 'significant' surface field values, to all the rest of the surface, to affect at each location a subset of future times. On large targets the net result of these projections forward to any given location will mostly be zero; they will cancel. Occasionally they will not all cancel, giving on those occasions a significant field, itself to be projected forward in turn.

Figure 6.3.1 shows the fraction of the surface (nodes) above the 1% threshold as the calculation progresses for the 6.5 and 12.7 wavelength long cases. The time average of this fraction is a good indication of the fractional cost saving. It is seen that as would be expected the fraction of the surface contributing, and its time average, declines as shorter pulses (or bigger bodies) are used.

We now will examine in rather more detail the process by which the field is built up (and indeed will observe that this lower cost approach still seems to

perform much redundant work, a point to which we will return later). We will consider for a particular location the variation with time of

(i) the number of elements from which significant fields were projected forward on to the node, and

(ii) the field value at the node, the net result of the projected forward fields.

We will consider a set of three nodes; one at the incident end on-axis, (a), one half way along, (b) and one on-axis at the far end on-axis, (c). These are indicated in figure 6.3.2 .

The three lines on figure 6.3.3 show results for each of the three locations (a), (b) and (c). In each case is shown the variation through time of the number of surface elements contributing to the field. Note that the times at which these elements were contributing will vary according to their position. The field at remote elements from long time ago, and nearby elements recently, will combine to produce the field at any particular location and time.

Shown on figure 6.3.4 is the field history itself at location (c); in effect a summation of the various contributions just described. It is seen that there is a large field when the incident pulse arrives; prior to that the field naturally zero, with causality naturally preventing any historical contributions. The duration of the large field is much shorter than the duration of the period when many elements contribute to the field. During time steps ~70 to 100 many elements contribute to the field there, but their contributions all cancel, and their net effect is zero. Much computational work is incurred in

calculating these zero values. Similar observations apply to locations, (a) and (b).

6.4 Perfectly Conducting Cone-Sphere with Gap

Another of the set of 'stealthy' RCS targets presented by Woo[47] is the cone-sphere with gap, shown in cross section in figure 6.4.1 . In this section we will investigate the effectiveness of the reduced cost methods in predicting its RCS. It is both a large and stealthy target, with the very low backscattered response a consequence of subtle cancellations, and thus provides a fairly severe test of the approach.

A single, head on, illumination was run with different levels of threshold on the Cray T3D. For the extraction of the frequency, 9 GHz (18 wavelengths long) results, a 7373 node mesh was used. This frequency corresponds to a nodal separation of ~1/6.6 of a wavelength, and is probably at the upper end of the frequency range for which the mesh is adequate. The mesh employed is shown in figure 6.4.2 .

The calculated and measured backscattered RCS for the VV polarisation is shown in table 6.4.1. It shows that the backscattered RCS is perturbed by about ~1.8 and ~0.5 dB sq m for 1% and 5% threshold cases respectively. Although the 5% threshold result is anomolously close to the measured RCS for backscatter, this is not necessarily the case over all scattering angles; in general we would expect better results for the 1% case.

We can investigate the cost saving that the threshold provides. Figure 6.4.3 shows the number of IPEC calls for the different levels of threshold.

Additionally, it is straightforward to calculate the number of such calls required with a zero threshold, and this too is indicated on the graph. We see, for example, that the 5% threshold, which degraded accuracy only slightly, nonetheless provided a cost saving by a factor of about 20.

6.5 Performance of the Parallel Algorithm

The strategy and algorithms associated with the parallel implementation have been discussed earlier. We here will discuss the actual performance of the algorithms on the T3D, addressing the 674 node dipole problem.

Figure 6.5.1 shows the CPU times to solve the matrix equation, and to form the matrix coefficients and to project forward, for those elements where it is necessary, for 32, 64 and 128 processors.

Lack of T3D time prevented a more comprehensive study of parallelisation, but is possible to infer that the parallelisation of forming coefficients and projecting forward is good. Time is about quartered by moving to 128 processors from 32. However, the parallelisation of solving the matrix is less good. This reflects that for this small problem, the matrix-vector multiplication which is actually parallelised is not a dominant cost in the solution stage because of the inevitable communications between processors. For example, it is seen that the solution of the matrix using 128 processors took essentially the same time as using 64 processors.

Calculation	1.3 wavelengths (% cycles)	4 wavelengths (% cycles)
Integration over an non-self element (IPECNSQ9)	85.19	88.93
Integration over an self element (IPECSQ9)	3.56	0.40
Reduction of 3 by 3 (sub) matrices to 2 by 2 (ELIMAT)	7.50	7.60
RHS formation	2.56	2.61

Table 6.1.1.1   Fraction of the CPU time expended on each calculation

Figure 6.1.1.1   Bistatic RCS of sphere : 1.3 wavelength long case

Figure 6.1.1.2   Bistatic RCS of sphere : 4 wavelength long case

Results of the Modified Algorithm   130
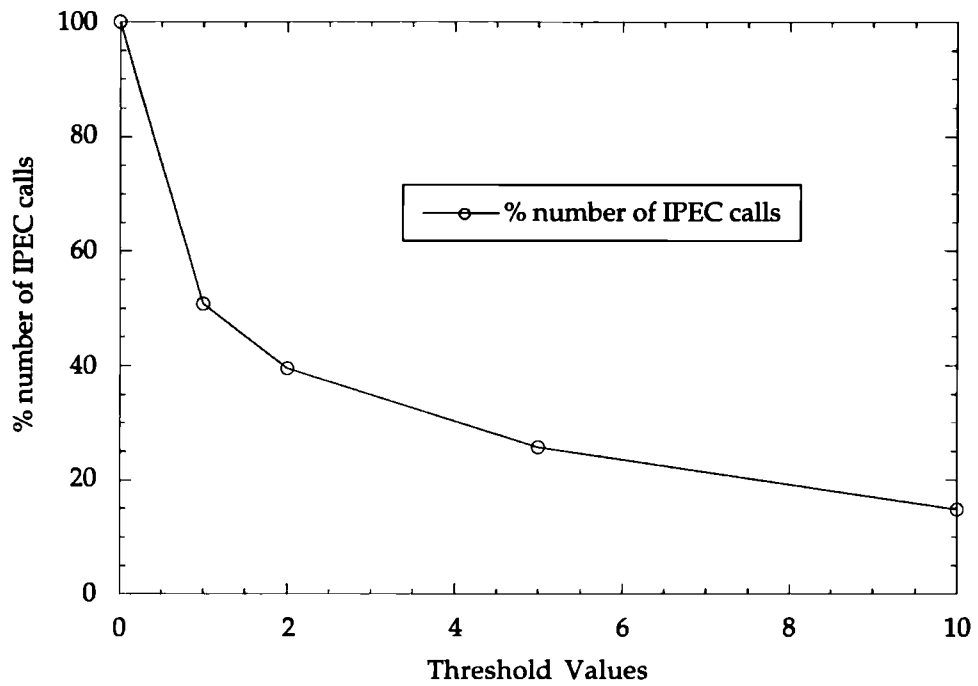
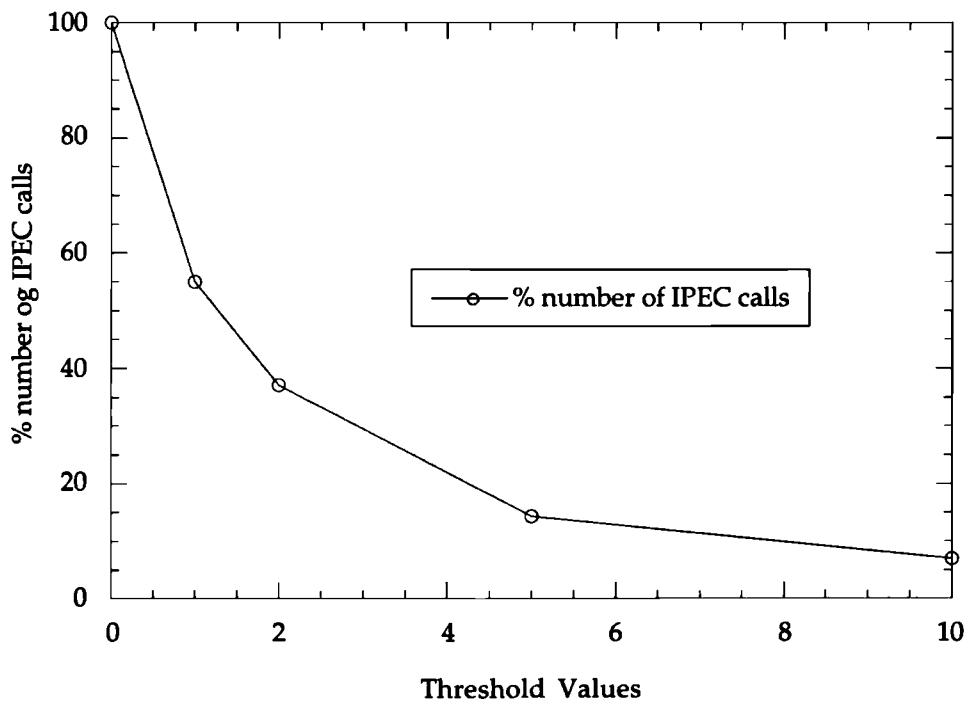Figure 6.1.1.3   % number of IPEC calls versus threshold values (1.3 wl case)

Figure 6.1.1.4   % number of IPEC calls versus threshold values (4 wl case)

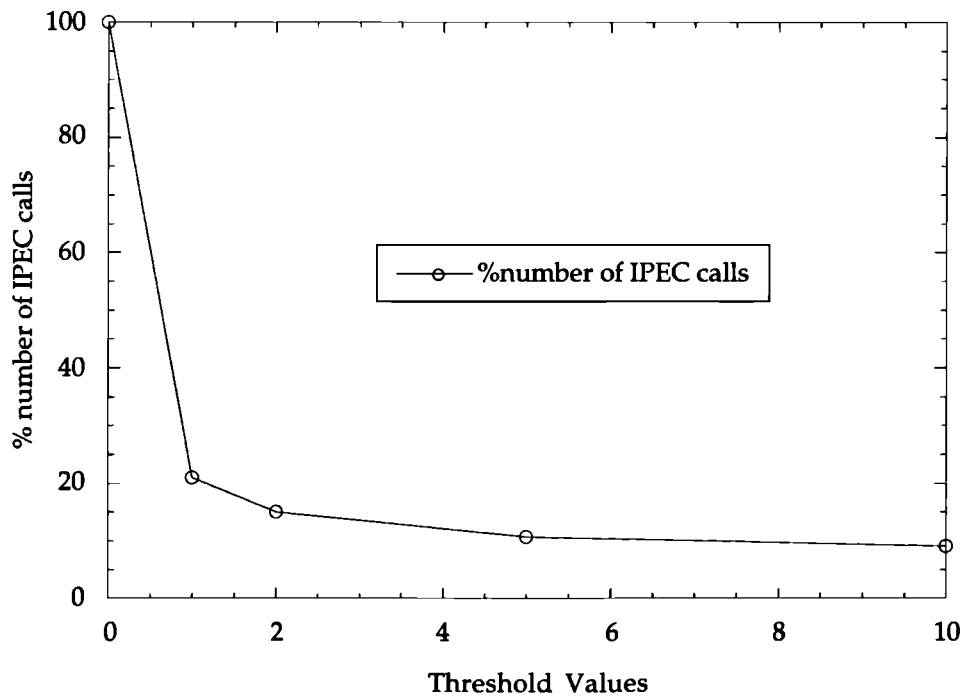Figure 6.1.2.1   Bistatic VV RCS of NASA almond at 1.19 GHz

Figure 6.1.2.2   % number of IPEC calls versus threshold values

Figure 6.2.1 Bistatic RCS of 10:1 dipole : 6.53 wavelength long case
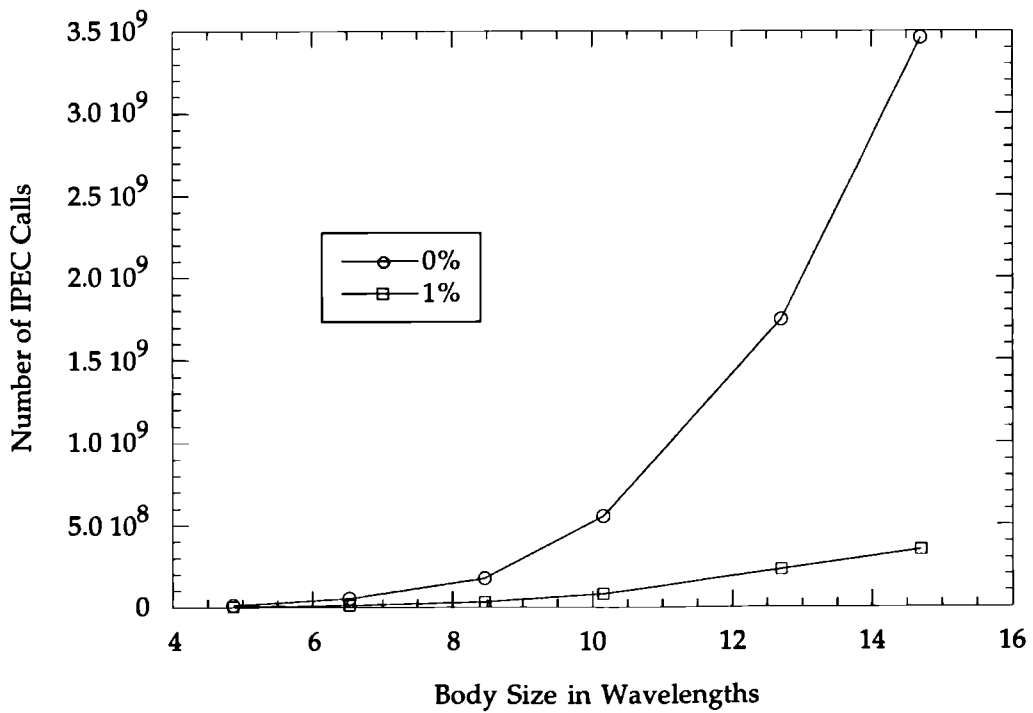

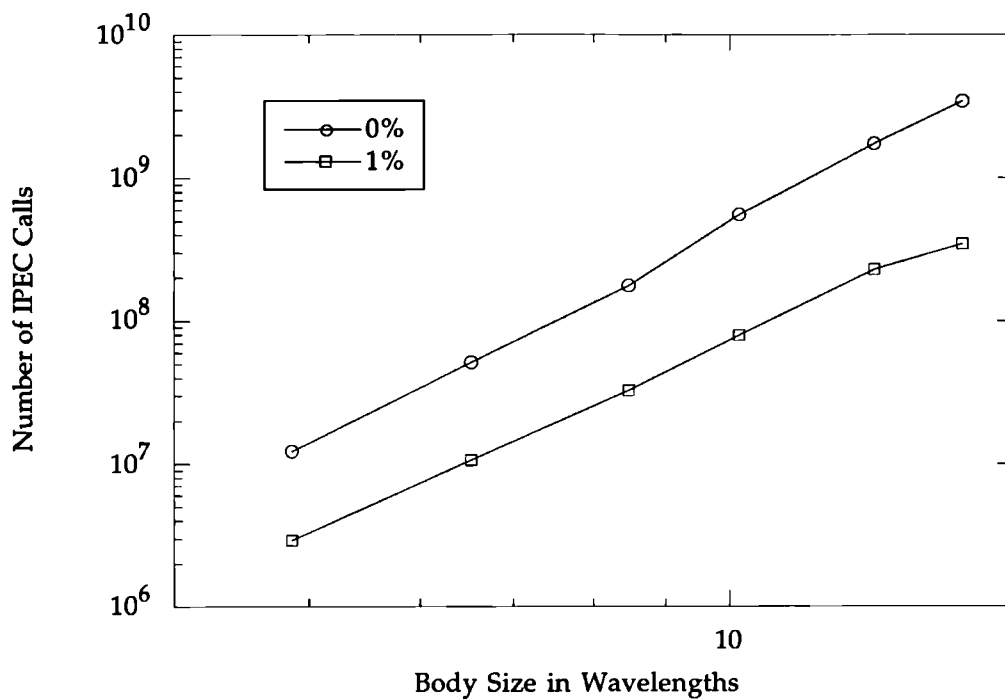**Figure 6.2.2** Number of IPEC calls versus body size in wavelengths

Figure 6.2.3　　　Number of IPEC calls versus body size in wavelengths

Figure 6.3.1　　　Fraction of dipole nodes with field above 1% threshold

Incident Wave

(b)

(a)

(c)

z

y

Figure 6.3.2   location (a), (b) and (c) on the dipole geometry


Figure 6.3.3   Number of elements contributing : 6.5 wavelength long case

Figure 6.3.4   Surface field values at (c) versus timestep

Figure 6.4.1   Cross section of the cone-sphere with gap

Figure 6.4.2   7373 node cone-sphere gap mesh

	RCS (dB sq m)
Measured	~ -6.3
1% threshold	-8.123
5% threshold	-5.866

Table 6.4.1   Backscattered RCS; cone-sphere with a gap at 9 GHz

Figure 6.4.3   % number of IPEC calls versus different levels of threshold

Figure 6.5.1   CPU time taken versus number of processors

Chapter 7 Hybridisation of the Project Forward Algorithm

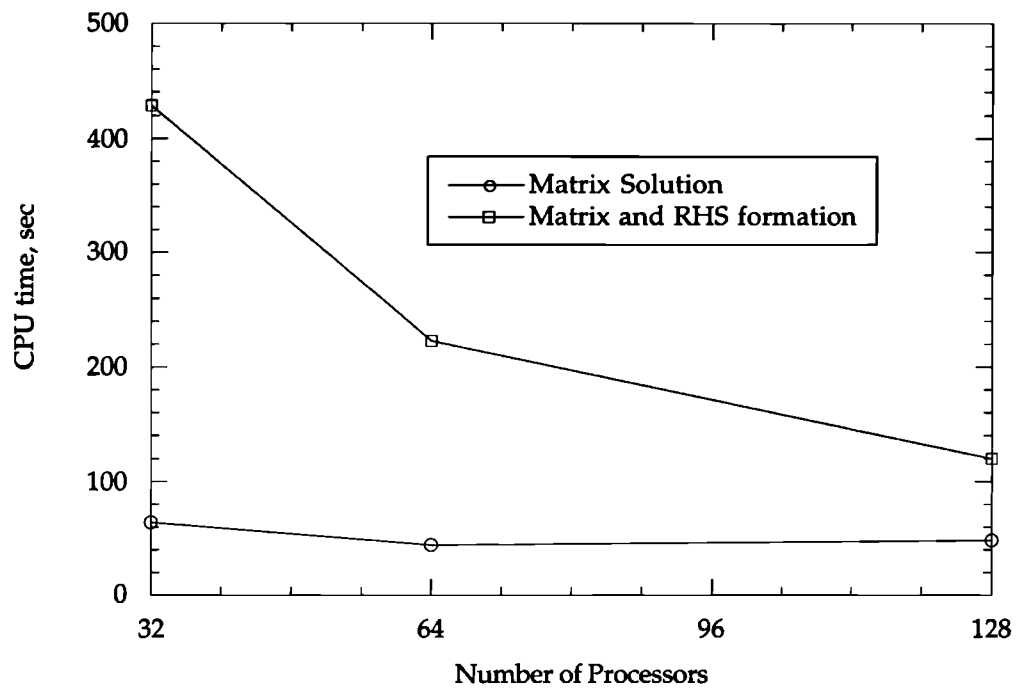The motivation for the entirety of the work described in this thesis is the high cost of full field solutions. Whilst this cost may be greatly reduced by the methods developed here, it remains high. A complementary approach to extending the size of problem which can be tackled is to hybridise two methods, using a cheap one in regions where its associated greater approximations are reasonably valid, and reserving the fuller, more expensive approach for regions where it is truly necessary. In essence, regions where the surface orientation changes over a length scale of at least several wavelengths can often be treated by a simple approach. In geometrically complex regions, where this is not so, a full field solution is employed. In the present context, this translates into using some (broadly) optical approach wherever feasible, and using the BIE approach where necessary.

There are many such problems where hybridisation of low frequency methods and optical methods is appropriate; where optical methods in isolation would be inaccurate, and field solutions too expensive. It is because of this that there is a considerable body of work addressing the hybridisation of various forms of optical treatment with full field solutions, allowing each to be used for that portion of the body for which it is most suited. All, bar one recent paper, has addressed the frequency domain, where issues and difficulties are markedly different from the time domain.

One recent work, by Walker and Vartinainen[28], has described hybridisation of the conventional 'retrospective' IETD approach with POTD. Besides its being an

intrinsically faster IETD algorithm (when thresholding is employed), there are features of the 'project forward' approach which also make it particularly suited to hybridisation. In this chapter a novel hybridisation of the project forward algorithm and time domain physical optics (POTD) is developed, exploiting simplicities this algorithm offers over that of the normal algorithm.

Section 7.1 introduces briefly the hybridisation of integral equation methods and physical optics. In section 7.2 is discussed the "normal" (that is, Walker & Vartinainen) time domain BIE hybridisation, and section 7.3 develops the hybridisation of the project forward algorithm . Some results are demonstrated in section 7.4, and finally the incorporation of the cost reduction algorithm into the BIE region is discussed in section 7.5.

It should be noted that the earlier referenced time domain hybridisation is itself a very recent piece of work, and the present hybridisation is thus even more recent. As a consequence, it is probably true to say that the implementation may not be as efficient as should be possible with more time available, nor has the investigation of its performance been as thorough. Nonetheless, it is believed that the principles of the new approach have been demonstrated, and opportunities for further research identified.

7.1 Introduction

In recent years, it has been shown that hybrid approaches incorporating high frequency (asymptotic) techniques and low frequency techniques (especially integral equation methods) have the potential to solve many problems which can not be solved by either one of these alone. Broadly speaking, these hybrid

approaches may be classified as either current- or field-based ones depending whether the mathematical *Ansatz* is a current or a field.

In this thesis is developed a hybridisation of integral equation (IE) techniques and physical optics (PO). There have been many reported hybridisations of IE techniques and PO in the frequency domain[59,60,61,62], but not in the time domain in spite of increasing interest in time domain methods of late.

Very recently Walker and Vartiainen[28] have presented a hybridisation of time domain BIE and POTD and investigated its cost savings. This employed the normal retrospective approach. It seems that one (initially unsuspected) feature of the project forward algorithm may be that it allows a simple and efficient implementation of hybridisation. The sections below will show how the new algorithm can be hybridised in a simpler way than the normal BIE algorithm, and will note the principal differences between the hybridisations of the two algorithms.

7.2 Hybridisation of the Time Domain Integral Equation Method

Consider a three dimensional body where it is identified that in some region, region (I), the local geometry requires a full field solution, and that in some region (II) POTD can be used. We will take the quantity θ as the fraction of nodes falling within (I).

Algebraic Form of the Hybridisation Approach

For nodes within region (I) partitioning the domain of integration into region (I) and (II) equation (2.1.68) for a smooth surface in chapter 2 can be written as;

$$2\pi H(r,t) = 4\pi H_{mc}(r,t) + \int_{\partial\Omega} \left(n' \times H(r',t^*)\right) \times \frac{\hat{R}}{R^2} + \left(n' \times \frac{\partial H(r',t^*)}{\partial t'}\right) \times \frac{\hat{R}}{R} ds'$$

$$+ \int_{\partial\Omega_{II}} \left(n' \times H(r',t^*)\right) \times \frac{\hat{R}}{R^2} + \left(n' \times \frac{\partial H(r',t^*)}{\partial t'}\right) \times \frac{\hat{R}}{R} ds' \qquad (7.2.1)$$

For nodes in (I) the surface H field is found from the usual integrations over the BIE region (I), plus the integrations over the optical region (II). Since in both regions the required historical H field is known (albeit via different methods, as discussed immediately below), the calculations in region I proceed exactly as normal.

For region (II) we do not employ (7.2.1) at all; the surface H field at any location and time is simply, and very cheaply, given by the physical optics approximation

$$H(r,t) = -2n \times \left(n \times H_{mc}(r,t)\right) \qquad (7.2.2)$$

Numerical Implementation of the Hybridisation of the Normal Algorithm

As shown in (7.2.2) the physical optics approximation gives the surface field values at nodes in region (II). The integration in (7.2.1) is of known quantities, and thus it can be easily evaluated. However, some computational issues arise in terms of efficiency. These are investigated below.

Normally historical fields only in region (I) are stored because storing historical PO solutions in region (II) is obviously wasteful evenif only the active band of region (II) for the pulsed illumination is considered. It is much cheaper to store the effect of these many POTD nodes on the (relatively) few IETD nodes, reducing the storage needs by the factor Θ. However, this requires that the time

in the future at which the field in question in the region (II) will affect some location in region I, and the size of the effect, should be determined in advance. As is clear in the work cited, incorporation of this into the conventional retrospective IETD approach is rather cumbersome. Conversely, it is apparent that it fits very naturally into the structure of the 'project forward' algorithm developed here.

7.3 Hybridisation of the Project Forward Algorithm

The hybridisation of the project forward algorithm with POTD can be more easily implemented than that of normal retrospective algorithm. For the region (II), again the POTD approximation is used for the surface H field solution, which is the same as in the hybridisation of the normal algorithm.

For the first integration in (7.2.1), over the IETD region I, the normal retrospective algorithm finds a geometrically weighted integral of earlier surface fields over all the rest of the body. On the other hand the approach in the project forward algorithm is to predict the size of effect which the fields in question in region (I) will have at times in the future.

This is also essentially exactly what is required to compute efficiently the influence of region II on region I. Since this is hardly special at all in the context of the project forward algorithm, hybridisation of it allows equation (7.2.1) to be solved by means of almost the same calculation procedure as used in the non-hybrid project forward approach.

7.4 Results

One difficulty with such a hybrid treatment is in finding suitable test cases to validate the implementations; if the test cases are small enough to be calculable by other means, they are too small for the hybrid approach to be expected to perform well. In this thesis, comparisons will be made with the full field solution for a small problem, and performance on some realistic problems will be demonstrated, but without any 'correct' solution to assess them against.

7.4.1 Correctness of Hybridisation

The correctness of the implementation of the hybridisation can be demonstrated by analysis of a problem small enough to be soluble by a full field solution (and thus a problem for which the hybrid is of little practical benefit). In such a case comparisons can be made with the full solution, and we would expect the analytical solution to be approached as the size of the optical region is reduced.

Figure 7.4.1.1 shows the bistatic RCS of a 2 wavelength (1583 node) sphere, calculated by the Mie series, and a set of hybrid analyses. In each case the central circular portion of the face of the sphere on which the wave impinges was analysed using POTD. The size of the POTD region is characterised by the angle its diameter subtends at the center. Good agreement can be seen between the wholly IETD result and the Mie series, with a degradation as more of the front face is analysed using POTD. Whilst use in such a way is of limited practical interest, it fulfils its primary purpose here, of indicating that the hybridisation itself is performing in a satisfactory manner.

7.4.2 Ball and plate

Now is considered a large problem, where much of the target is electrically large and electrically smooth, whilst some is neither. The geometry analysed is taken from Walker and Vartinainen[28]; a unit radius sphere placed centrally 1.75 diameters in front of a square plate of side 50 diameters. The sphere and a circular region on the plate directly behind the sphere are treated by BIE time domain approach, and the balance of the plate, the large majority of the total problem surface, by POTD. Here we will take as the diameter of the IETD circle 6.5 diameters of the sphere. Behaviour of the surface field H is investigated for the case of pulsed illumination from 6 degrees off-axis.

In figure 7.4.2.2 is shown the surface field at a location (a) in figure 7.4.2.1. The main feature of the response, and in particular the double peak of the incident wave and its reflection from the plate past (a), are well captured. The small later peak occurs later by about one transit time from the plate edge back to the centre, and presumably corresponds to a reflection from the plate edge.

7.4.3 Missile on Wing

As another large problem, a pointed cylinder located near and parallel to the surface of a large plate (akin to a missile suspended below a wing) is considered. The missile is placed centrally 1 diameter in front of a flat square plate of side 60 diameters of the missile. The missile and plate are meshed with 1794 nodes and 18225 nodes respectively. The missile itself, and a region extending about two diameters beyond the projection of the missile onto the plate, are treated by the BIE time domain approach. The balance of the plate is treated by POTD. The behaviour of the surface field H is investigated, using 15 degrees off-axis illumination.

Figures 7.4.3.1 and 7.4.3.2 show the mesh employed, and the overall arrangements. The time variation of the surface field magnitude is shown graphically in figure 7.4.3.3. It is seen from the figure that the surface field magnitude peaks sharply as the incident wave passes and exhibits no wake on that part of the wing which is a PO region, as expected from a PO approximation.

Figure 7.4.3.4 shows the interaction between the missile and wing. Two narrow active bands on the missile are clearly seen in this figure. One is caused by the incident wave, and the other presumably from the reflection from the wing.

In the results shown here, the cost reduction via thresholding, presented in chapter 3, was not applied. (As noted, it was not possible to devote very much time to the hybridisation itself, and it seemed prudent not to confuse the issue at so early a stage.) Nonetheless, the example shown here in 7.4.3 clearly offers the opportunity to apply it. The missile and adjacent plate has a relatively narrow band of active region, while the large majority of the rest of the IETD region is quiescent most of the time. The application of a threshold in the IETD region could thus result in considerable cost savings.

7.5 Discussion

The ability to hybridise the project forward approach with time domain physical optics has been demonstrated. This hybridisation offers clear gains in simplicity of implementation over the conventional retrospective approach. It similarly offers the ability to perform truly time domain scattering calculations on problems which would be way beyond the capacity of computers to perform

full field solutions. Additionally, this present approach permits the straightforward incorporation of thresholding, gaining the same cost reductions for the IETD region as have been demonstrated in the non-hybrid parts of this thesis.
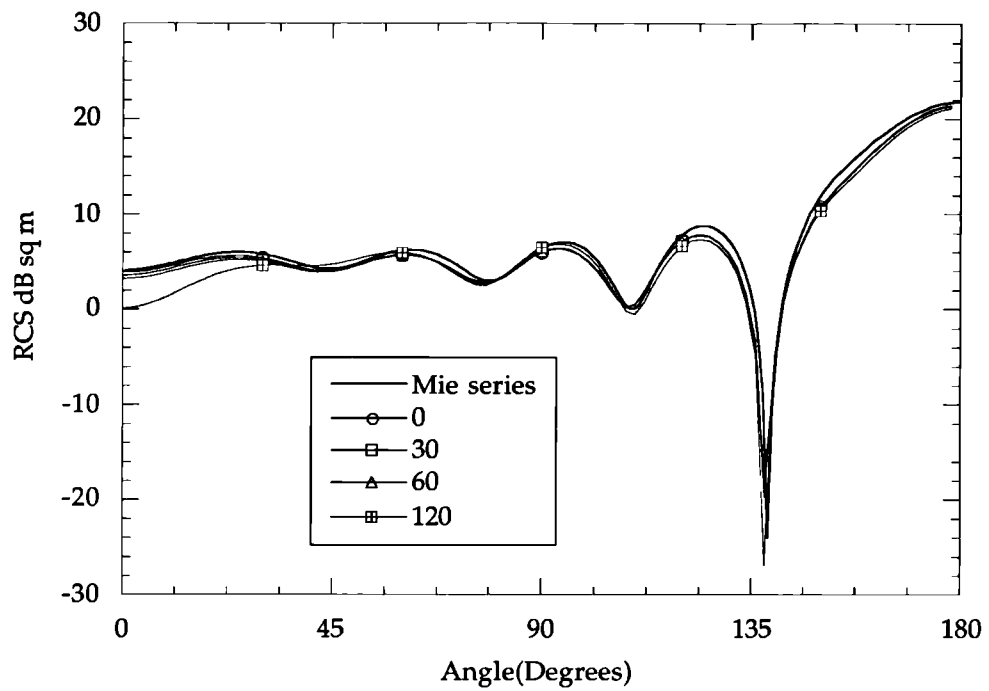
Figure 7.4.1.1   Bistatic RCS of sphere : 2 wavelength long case, for different angular extents of POTD region.
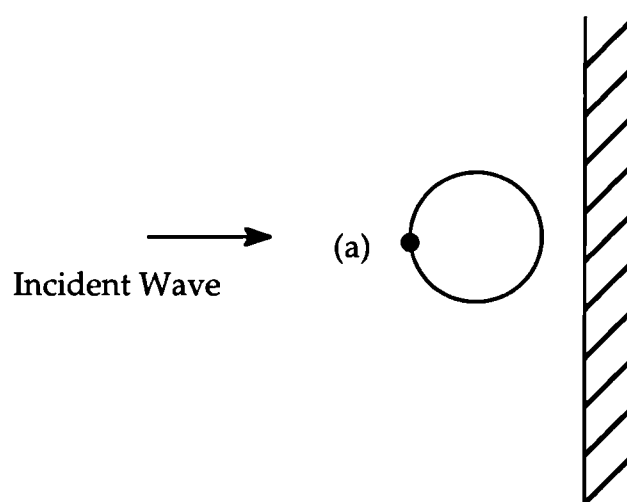

Figure 7.4.2.1   Target geometry and location (a)

Figure 7.4.2.2   Suraface field values versus time at (a)

Figure 7.4.3.1   20019 node "missile on wing" mesh

Figure 7.4.3.2 Cross section of the mesh indicating the incident wave

Figure 7.4.3.3 Time Variation of the surface field magnitude

FRINGE PLOT LC=2.15 RES=1.1(VEC-MAG) MSC/PATRAN R-1.4 EMsolver 24

FRINGE PLOT LC=2.21 RES=1.1(VEC-MAG) MSC/PATRAN R-1.4 EMsolver 24

FRINGE PLOT LC=2.12 RES=1.1(VEC-MAG) MSC/PATRAN R-1.4 EMsolver 24

FRINGE PLOT LC=2.18 RES=1.1(VEC-MAG) MSC/PATRAN R-1.4 EMsolver 24

Hybridisation of the project forward algorithm

154

Figure 7.4.3.4 The interaction of the missile and wing at a certain time

Chapter 8. Conclusions

The prediction of the behaviour of wave scattering is very important in many branches of engineering. Of significant interest in this thesis are RCS problems. For stealthy objects, RCS analyses combine requirements for good accuracy with the need to model bodies many wavelengths long. However these require huge computational resources, and present computers and numerical methods limit the size and complexity of problems which can be solved. Another broad area of application, and one where the true time domain calculation is particularly valuable, is in electromagnetic pulse (EMP) calculations. Here too, similar comments about computational costs apply.

Though there exist many numerical methods which are suitable for RCS and EMP problems, the time domain BIE method is one of the more efficient tools for their analysis. Its computational costs scale with the fifth power of the frequency. While this is broadly comparable to or better than the cost scaling of alternative techniques (such as finite difference time domain, or frequency domain integral equations), because of this cost scaling it is impossible to solve realistic problems on present computers.

The project forward algorithm presented here offers the prospect of large reductions in cost, and a reduction in cost scaling to the fourth power of the frequency. This has been achieved by the combination of a re-ordering of the algorithm, and implementation of a further physical approximation which this reordering makes possible. This approach is demonstrated for a variety

of geometries, including relatively difficult ones such as the NASA almond and the 'cone-sphere and gap'. Results in these cases have shown that the accuracy loss occasioned by the physical approximation is acceptable, and that large cost savings can be achieved, with larger cost savings being observed as the body size increases. On the examples studied cost savings by upto factors of 20 have been achieved.

The approach obviously has potential application in a variety of other fields. One is naturally acoustics, which is computationally essentially a subset of computational electromagnetics. Another area, where the same physical observation should hold, is elastodynamics.

In addition to allowing the easy implementation of the physical approximation, the project forward algorithm has some advantages over the normal retrospective algorithm in parallelisation. The partitioning over the field nodes seems to be the most attractive approach in terms of the simplicity and efficiency, but the consequent duplication of storage renders it unsuitable for the retrospective algorithm. The project forward algorithm does however allow this most attractive approach to be used, by its particular and crucial characteristic of the replacement of storage of the historical field values with storage of the future right hand side vectors. The efficiency of parallelisation of the project forward algorithm has been investigated by solving the same problem with different numbers of processors. Good parallelisation is observed in the matrix formation and the future right hand side formation stages, which are the dominant costs.

It has been shown that the project forward algorithm can be hybridised with POTD. As is anyway the underlying approach in the BIE region when using project forward, it is possible to determine the time in the future at which the field in question in the POTD region will have an effect, and the size of that effect. As a consequence of that, relatively minor changes are required for implementation of the hybridisation. Some large problems are analysed, showing reasonable results in terms of the physical behaviour of the EM wave. Further, the cost of the BIE time domain portion of the hybrid calculation can be reduced in just the same fashion as the non-hybrid case, using the physical approximations discussed above though this has not been done.

Topics for further Research

There seem two main areas for further research:

In the application of the physical basis for the cost reduction, only a relatively crude threshold was used. Some more sophisticated thresholding techniques could be implemented, such as the threshold in terms of the present maximum field, or even spatially or time dependent thresholds. A 'near neighbour' approach as explained in chapter 3 might be incorporated in addition. More study about those would be beneficial.

The hybridisation of the project forward algorithm with POTD was developed in the previous chapter. The application of physical basis for the cost reduction may be extended to the BIE part of the hybrid approach. This also would be a fruitful area for further work.

APPENDIX

For cost saving in the frequency domain, a study of the use of iterative methods for the solution of the large dense scattering matrices was made. In particular it involved development of a rational criterion for termination of the iteration. This work[63] will shortly appear in the journal of "Communications in numerical methods in engineering", and the text of the forthcoming paper forms this appendix.

Termination Criteria in Iterative Solution of Large Scattering problems Using Intergral Equation Methods

Abstract

Iterative methods are increasingly used for solution of the extremely large matrix equations generated by integral equation analysis of multi-wavelength frequency domain scattering. Whilst much cheaper than direct methods, the matrix solution remains the dominant cost, and is very costly.

The criterion adopted for termination of the iteration can have a marked effect on this cost. We show that for large scattering problems a robust and rational prior choice of termination criterion can be made, based only on discretisation. This allows confident use of a much larger termination residual than those commonly used, with consequent cost reduction.

1. Introduction

There is considerable interest in the solution of larger and larger scattering problems, in acoustic, elastodynamic and, perhaps in particular, electromagnetic applications. 'Large' in this context relates to the body size, expressed in wavelengths of the incident field. Frequency domain integral equation methods[1,2] (method of moments, boundary integral equations) are widely used, with their well known advantages of surface- only discretisation, automatic satisfaction of the radiation boundary condition, and so on.

The two components of computational work in such an analysis, matrix formation and solution, scale with size to the fourth and sixth powers respectively. There is thus considerable interest in reducing solution time, as increased interest in iterative methods[3-9] indicates.

Any iterative method is a gradual approach to the exact solution of the matrix equation; a solution which could in principle be found by direct methods. It is necessary to decide when to terminate the iteration; when the approach of the iterative solution to the exact solution of the matrix equation is sufficiently close. Since the solution cost is dominant, there is obviously a strong incentive to terminate as soon as possible. In particular, there is no benefit in spending effort in seeking, via prolonged iteration, an iterative solution which is an extremely close approximation to the exact solution of the matrix equation, when that 'exact' solution itself embodies by then much larger errors due to the discretisation.

This general issue, of the interplay between the error due to truncation of an iterative solution, and the discretisation error, is an area in which considerable work has been done in a number of branches of numerical mathematics. The interested reader is referred to chapter 4 of Briggs[10], and the references cited therein. In this paper we address the issue of what constitutes 'sufficiently close' in the context of large harmonic scattering problems; how far from the exact solution of the matrix equation can one safely terminate the iteration, on the grounds that the matrix equation solution is itself only an approximation to the physical problem.

However, we naturally do not know the analytical solution to the physical problem, and so do not know how far from it is the exact solution of the matrix equation representing the discretised problem. However, for the practically important case of multi-wavelength scattering problems, it is almost invariably the modelling of the field variation which is the determinant of the error due to discretisation.

It will be shown that, via its expression in the same terms as the residual in the iterative method by transforming it into the right hand side space, a robust and reliable prediction can be made of the difference between the analytical solution and the exact solution of the discretised problem. This difference is primarily a function of the fineness of discretisation; the principal determinant of solution cost, and under the direct control of the analyst.

The closeness expressed in these terms of the exact to the analytical solution provides a rational target for the closeness of the exact and iterative solution which need be sought. This allows a more rational, albeit necessarily still

empirical, approach to the choice of the value of the iteration residual at which to terminate. In particular, it allows values lower than those commonly reported to be used with confidence, with significant savings in time.

2. The method

We consider either acoustic or electromagnetic scattering of an incident ϕ_{inc} wave from the exterior of a hard (perfectly conducting) body, transformed into a second kind integral equation. The surface of the body is then discretised, and a matrix equation formed[1,2].

We then seek the solution to

$$A\phi_{DIS} - \phi_{inc} = 0 \tag{1}$$

where A is the matrix resulting from discretisation of the integral equation and ϕ_{DIS} the (exact) solution of the matrix equation. We will define additional vectors: ϕ_{AN} the values of the analytical (ie correct physical) solution at the nodes, and $\phi^{(k)}_{Iter}$ the k'th attempt to find the solution via an iterative method.

We define a discretisation error vector:

$$e_{DIS} = \phi_{AN} - \phi_{DIS} \tag{2}$$

The iterative method calculates a series of iterates via

$$\phi^{(k)}_{Iter} = \phi^{(k-1)}_{Iter} + \alpha^{(k-1)}d^{(k-1)} \tag{3}$$

where the difference between methods lies essentially in the choice of step lengths α and step directions d. Associated residuals are

$$A\phi_{Iter}^{(k)} - \phi_{inc} = r^{(k)} \tag{4}$$

The iteration is terminated when

$$\frac{\left\|r^{(k)}\right\|}{\left\|\phi_{inc}\right\|} \leq R_{Iter} \tag{5}$$

The selection of a small value for R_{Iter} ensures that the iteration is not terminated till the iterative solution $\phi^{(k)}_{Iter}$ is sufficiently close the exact solution ϕ_{DIS} of the matrix equation.

There has been relatively little attention paid to the actual value adopted for the quantity R_{Iter} , and figures ranging from 10^{-3} to 10^{-7} have been used in the iterative solutions of integral equations referenced above, with essentially no discussion of the rationale behind the choice.

3. Discretisation error as a quasi-residual

The error due to discretisation can be re-expressed in a form similar to the residual. Since the discretised solution, not the analytical solution, is the solution to the matrix equation (1), we can identify a discretisation residual vector

$$A\phi_{AN} - \phi_{inc} = r_{DIS} \tag{6}$$

and thus define the 'discretisation-residual':

$$\frac{\left\|r_{DIS}\right\|}{\left\|\phi_{inc}\right\|} = R_{DIS} \tag{7}$$

This of course cannot normally be evaluated, as the analytical solution is unknown.

However effective (3) is in generating small residuals in (4), it is still only finding the discretised solution, ϕ_{DIS}, ever more precisely. The discretised solution remains just that; an approximation to the analytical one. A measure of how far it is from the analytical one, in terms comparable to the iteration residual, is provided by (7).

Iteration to a vanishingly small value of R_{Iter} still would correspond to a 'true' residual R_{DIS} as given by (7) above, and is thus wasteful. Iteration till $R_{Iter} \sim R_{DIS}$ may be worthwhile; beyond this the return for the extra computational work is likely to be small.

4. Determination of the discretisation residual

It remains to estimate the discretisation residual; the analytical solution, used in (6) and (7), is obviously not available in any case of real interest.

The error arises from imperfect discretisation; imperfect representation of the geometry, imperfect integrations to form the coefficients of the system matrix, and imperfect representation of the field variation over the surface. For the multi-wavelength bodies of interest, the field variation is much greater than the geometry variation, and is the dominant contributor to error. It is to be expected then that a suitable parameter against which to correlate error will be the number of nodes per wavelength used in the discretisation.

It is this observation, that the discretisation error is likely to depend primarily on a quantity available in advance, which permits its use as a benchmark against which to compare the iteration error.

This dominant dependence of error for large problems on field discretisation is readily demonstrated. A series of sphere meshes, comprising quadratic quadrilateral elements, with 266, 530, 770, 1130 and 1658 nodes N, were each excited with incident scalar waves of a variety of frequencies. Figure 1a shows error versus nodal separation. It is seen that the nodal separation is indeed a good criterion for accuracy, and that as expected it becomes increasingly so as bigger bodies are modelled; for small bodies the influence of representation of the geometry is clear. The actual separation selected obviously will depend on the needs of the work. For large scattering problems such as radar cross section evaluation typically separations of about $1/5$ to $1/10$ wavelength are employed.

In figure 1b is plotted the variation of the 'discretisation residual' R_{DIS}. Nodal separation correlates well with this, especially for the cases of most relevance (large problems where representation of the geometry does not intrude). If for example a nodal separation of ~$1/6$ of a wavelength is employed, the figure indicates there is nothing to be gained by using an iteration residual below ~ 6×10^{-3}, and that in practice a value of ~ 10^{-2} can probably be employed with little loss in accuracy.

5. Demonstrations of the method

Here we apply the above methods for selecting the iteration residual, and investigate their effectiveness in reducing the number of iterations. We employ

the complex bi-conjugate gradient method[11], although others produce similar results. Since the residual does not decline monotonically, use of the moving average of the last few values introduces a further slight conservatism and produces more consistent termination.

We will first investigate the behaviour on two spheres, discretised with very different fineness, and consequently producing very different discretisation errors. Three other more practically relevant geometries are then investigated, with 'analytical' solutions approximated by separate analyses with very fine meshes and double precision direct solutions.

5.1 Large and small spherical scatterers

Figure 2a shows error and residual during solution of a 4610 node sphere at 2 and 8 wavelengths diameter, with nodal separations of 0.05 and 0.2 .

Figure 1b would suggest termination with residuals of $\sim 0.5 \times 10^{-3}$ and $\sim 10^{-2}$; a factor of 20 apart. This would cause termination essentially as soon as the final discretisation error is reached, with modest increase in error. Relative to a residual of say 10^{-6}, savings for both would be by factors of ~ 2.

This simple approach has thus allowed us to select, *a priori*, residuals such that we stop very close to the eventual errors in each case, despite the errors themselves and the selected residuals differing by a factor of 20 between the cases.

5.2 Prismatic scatterer

We consider now a rectangular prismatic scatterer, of sides 1, 2 and 10, subject to an incident wave 20^0 off-axis, of wavelength 1.9. The course of the iteration with a nodal separation of 0.14 wavelengths is shown in figure 2b. From figure 1b termination at a residual of ~4 x 10^{-3} is predicted to be acceptable, and would reduce the computational work relative to say 10^{-6} by a factor of about 2.

5.3 Slender dipole

Figure 2c shows analysis of scattering from a slender dipole (hemispherical end caps, aspect ratio 10:1, ~8 wavelengths long, nodal separation 0.14, 30^0 off-axis illumination) .

A residual of ~4 x 10^{-3} would again be suggested. This does not increase the error arising from the discretisation itself, but the consequent termination after ~12 iterations reduces by ~4 the work required to reach 10^{-6}.

5.4 Almond

Figure 2d shows analysis of scattering from a 866 node 'almond', a widely analysed benchmark[12] (here, 4.3 wavelengths long, nodal separation 0.13, head-on illumination). Figure 1(a) suggests a residual of ~2 x 10^{-3}, and figure 2d shows this leaves unchanged the error arising from the discretisation itself. This is achieved after ~7 iterations, reducing by a factor of about two the work required to reach 10^{-6}.

6. Conclusions

For large scatterers, many wavelengths in size, the need to model the field variation is the determinant of the discretisation needed. In such cases the error is predominantly a function of nodal separation in wavelengths. It has been shown that it is possible to express this error as a 'discretisation residual', the value of which can be characterised reliably in terms of this separation, and which is thus available, in advance, for practical cases in which the solution is naturally not known. This discretisation residual can then be used as a criterion for terminating iterative solution of the matrix equation. Its successful application in cases with markedly different errors and residuals indicates the robustness of the empirical criterion developed.

Use of this residual value gives an 'iteration error' generally small compared to, and at worst comparable to, the discretisation error. For typical discretisations employed its use results in a reduction in computational work by a factor of two or more compared to the residuals of $\sim 10^{-6}$ commonly employed.

Acknowledgements

References

1. Harrington, R.F. *Field Computation by Moment Methods*, New York:MacMillan, 1968.

2. Umashankar, K.R. Numerical Analysis of Electromagnetic Wave Scattering and Interaction Based on Frequency Domain Integral Equation and Method of Moment Techniques. *Wave Motion* 10:493-525, 1988.

3. Mansur, W.J., Araujo, F.C. and Malaghini, J.E.B. Solution of BEM Systems of Equations via Iterative Techniques. *International Journal for Numerical Methods in Engineering* 33:1823-1841, 1992.

4. Barra, L.P.S., Coutinho, A.L.G.A., Mansur, W.J. and Telles, J.C.F. Iterative solution of BEM equations by GMRES algorithm. *Computers and Structures* 44 no. 6:1249-1253, 1992.

5. Kleinman, R.E. and van den Berg, P.M. Iterative Methods for Solving Integral Equations. *Radio Science* 26:175-181, 1991.

6. Sarkar, T.K., Siarkiewicz, K.R. and Stratton, R.F. Survey of Numerical Methods for Solution of Large Systems of Linear Equations for Electromagnetic Field Problems. *IEEE Transactions on Antennas and Propagation* 29 No 6:847-856, 1981.

7. 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* 1996.(In Press)

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* (in press):1996.

9. Kane, J.H., Keyes, D.E. and Guru Prasad, K. Iterative Solution Techniques in Boundary Element Analysis. *International Journal for Numerical Methods in Engineering* 31:1511-1536, 1991.

10. Briggs, W.L. *A Multigrid Tutorial*, Philadelphia:SIAM, 1987.

11. David, T. and Blyth, G. Parallel algorithms for panel methods. *International Journal for Numerical Methods in Fluids* 14:95-108, 1992.

12. Woo, A.C., Wang, H.T.G. and Schuh, M.J. Benchmark radar targets for the validation of computational electromagnetics programs. *IEEE Ant.Prop.Mag* 35:84-89, 1993.
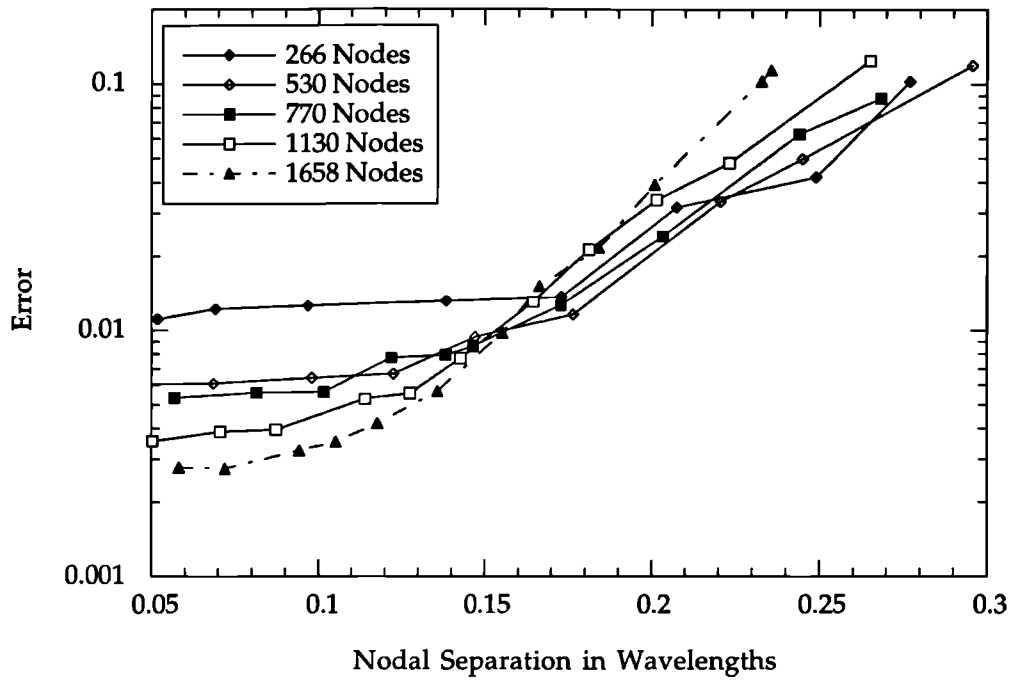
Figure captions

Figure 1

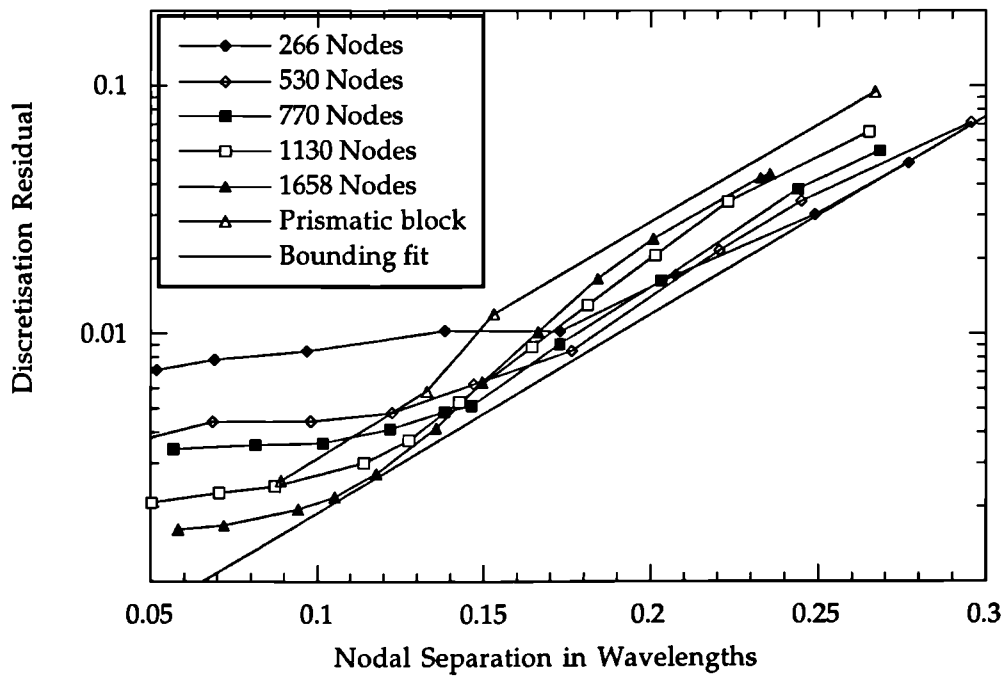Error (a) and discretisation residual (b) for spheres discretised with different numbers of nodes, illuminated at a range of frequencies.

Figure 2

Error and residual versus iteration number, for: (a) spheres, 4610 nodes, 0.05 wavelength nodal separation, 2.0 wavelengths in size (diameter) and 4610/0.2/8.0, (b) prismatic bar, 488/0.14/5.3, (c) dipole 10:1 aspect ratio, 722/0.14/8.0, (d) almond 866/0.13/4.3.

(a)

(b)

(a)

(b)

(c)

(d)

Termination Criteria in Iterative Solutions of Large Scattering Problems
Using Integral Equation Methods

References

1. Bowman, J.J., Senior, T.B.A. and Uslenghi, P.L.E. *Electromagnetic and Acoustic Scattering by Simple Shapes*, 1969.

2. Ufimtsev, P. *The Method of Fringes Waves in the Physical Theory of Diffraction*, 1962.

3. Keller, J.B. Geometrical theory of diffraction. *J.Opt.Soc.Am.* 52:116-130, 1962.

4. Rao, S.M., Wilton, D.R. and Glisson, A.W. Electromagnetic Scattering by Surfaces of Arbitrary Shape. *IEEE Transactions on Antennas and Propagation* 30:409-418, 1982.

5. Kirby, P., Lizius, J.G., Morgan, J.G. and Wood, J.C. CLASP - a new, genral-purpose Method of Moments Scattering Code. *IEE Proc.Int.Conf.Comput.Electrom.* 1:137-140, 1991.

6. Caorsi, S., Moreno, D. and Sidoti, F. Theoretical and Numerical Treatment of Surface Integrals Involving the Free-Space Green's Function. *IEEE Transactions on Antennas and Propagation* 41:1296-1301, 1993.

7. Vechinski, D.A. and Rao, S.M. A stable procedure to calculate the transient scattering by conducting surfaces of arbitrary shape. *IEEE Transactions on Antennas and Propagation* 40:661-665, 1992.

8. Rao, S.M. and Wilton, D.R. Transient scattering by conducting surfaces of arbitrary shape. *IEEE Transactions on Antennas and Propagation* 39:56-61, 1991.

9. Brebbia, C.A. Walker S. *Boundary Element Techniques in Engineering* 1980.

10. Banerjee, P.K. and Butterfield, R. *Boundary Element Methods in Engineering Science*, London:McGraw Hill, 1981.

11. Kress, R. Numerical Solution of Boundary Integral Equations in Time-harmonic electromagnetic scattering. *Electromagnetics* 10:1-20, 1990.

12. Schlemmer, E., Rucker, W.M. and Richter, K.R. Boundary Element Computations of 3D Transient Scattering from Lossy Dielectric Objects. *IEEE Transactions on Magnetics* 29:1524-1527, 1993.

13. Bluck, M.J. and Walker, S.P. Time Domain BIE Analysis of Large Three Dimensional Electromagnetic Scattering Problems. *IEEE Transactions on Antennas and Propagation* (submitted):1995.

14. YEE, K.S. Numerical Solution of Initial Boundary Value Problems Involving Maxwell's Equations in Isotropic Media. *IEEE Tr.Ant.& Prop.* 14,n.3:302-307, 1966.

15. Taflove, A. and Umashankar, K.R. Review of FD-TD Numerical Modeling of Electromagnetic Wave Scattering and Radar Cross Section. *IEEE Proc.* 77:682-699, 1989.

16. Mur, G. Absorbing Boundary Conditions for the Finite-difference approximation of the Time-domain Electromagnetic Field Equations. *IEEE Trans.Electromagn.Compat.* 23:377-382, 1981.

17. Petropoulos, P.G. A Simple Guideline for Choosing the Points per Wavelength in FD-TD Simulations. *IEEE Ant.Prop.Mag* 36:67-68, 1994.

18. Petropoulos, P.G. Phase Error Control for FD-TD Methods of Second and Fourth Order Accuracy. *IEEE Transactions on Antennas and Propagation* 42:859-862, 1994.

19. Tirkas, P.A., Balanis, C.A., Purchine, M.P. and Barber, G.C. Finite Difference Time-Domain Method for Electromagnetic Radition, Interference and Interaction with Complex Structures. *IEEE Trans.Electromagn.Compat.* 35:192-203, 1993.

20. Jurgens T.G., Taflove A., Umashankar K. and Moore T. G. Finite Difference Time-Domain Modeling of Curved Surfaces. *IEEE Transactions on Antennas and Propagation* 40:357-366, 1992.

21. Holland, R. Pitfalls of Staircase Meshing. *IEEE Trans.Electromagn.Compat.* 35:434-439, 1993.

22. Lavers, J.D. Electromagnetic Field Computation in Power Engineering. *IEEE Transactions on Magnetics* 29:2347-2353, 1993.

23. Silvester, P.P. and Ferrari, R.L. *Finite Elements for Electrical Engineers*, Cambridge:Cambridge University Press, 1990. Ed. 2

24. Wesson, P. and Parrot, A.K. Performance of a Parallel Unstructured Electromagnetic Scattering Code. *UCINA Workshop on Computational Electromagnetics* 1994.

25. Mur, G. and de Hoop, A.T. A Finite Element Method for Computing Three-dimensioanl Electromagnetic Fields in inhomogeneous Media. *IEEE Transactions on Magnetics* 21:2188-2191, 1985.

26. Volakis, J.L., Chatterjee, A. and Kempel, L.C. Rewiev of finite-element method for three-dimensional electromagnetic scattering. *J.Opt.Soc.Am.* 11:1422-1433, 1994.

27. Medgyesi-Mitschang, L.N. and Wang, D.-S. Hybrid methods for analysis of complex scatterers. *Proceedings of the IEEE* 77 no. 5:770-779, 1989.

28. Walker, S.P. and Vartiainen, M.J. Hybridisation of curvilinear time domain integral equation and time domain optical methods for electromagnetic scattering analysis. *IEEE Transactions on Antennas and Propagation* (submitted):1996.

29. Bluck, M.J. *Integral Equation Methods for Transient Wave Propagation*, PhD Thesis, University of London, 1993.

30. Poggio, A.J. and Miller, E.K. Integral Equation Methods of Three-Dimensional Scattering Problems. In: *Computer Techniques for Electromagnetics*, edited by Mittra, R. Oxford: Pergamon Press, 1973, p. 159-265.

31. Morse, P.M. and Feshbach, H. *Methods of Theoretical Physics*, New York:McGraw-Hill, 1953.

32. Richmond, J.H. A wire grid model for scattering by conducting bodies. *IEEE Trans. Antennas Propagatation* 14:782-786, 1966.

33. Lee, K.S.H., Martin, L. and Castillo, J.P. Limitations of wire-grid modelling of a closed surface. *IEEE Trans. Electromag. Compat.ability* 18:123-129, 1976.

34. Ingber, M.S. and Ott, R.H. An Application of the Boundary Element Method to the Magnetic Field Integral Equation. *IEEE Transactions on Antennas and Propagation* 39:606-611, 1991.

35. Chao, J.C., Liu, Y.J., Rizzo, F.J., Martin, P.A. and Udpa, L. Regularized Integral Equations and Curvilinear Boundary Elements for Wave Scattering in Three Dimensions. *IEEE Transactions on Antennas and Propagation* 43:1416-1421, 1995.

36. Lachat, J.C. and Watson, J.O. Effective Numerical Treatment of Boundary Integral Equations: A Formulation for Three-Dimensional Elastostatics. *International Journal for Numerical Methods in Engineering* 10:991-1005, 1976.

37. Rynne, B.P. Stability and convergence of time marching methods in scattering problems. *IMA Journal of Applied Mathematics* 35:297-310, 1985.

38. Rynne, B.P. Instabilities in time marching methods for scattering problems. *Electromagnetics* 6:129-144, 1986.

39. Smith, P.D. Instabilities in time marching methods for scattering; cause and rectification. *Electromagnetics* 10:439-451, 1990.

40. Rynne, B.P. Time domain scattering from dielectric bodies. *Electromagnetics* 14:181-193, 1994.

41. Vechinski, D.A., Rao, S.M. and Sarkar, T.K. Transient scattering from three-dimensional arbitrary shaped dielectric bodies. *J.Opt.Soc.Am.* 11:1458-1470, 1994.

42. Liao, H. and Xu, Z. A Method for Direct Evaluation of Singular Integral in Direct Boundary Element Method. *International Journal for Numerical Methods in Engineering* 35:1473-1485, 1992.

43. Bluck, M.J., Pocock, M.D. and Walker, S.P. An accurate method for the calculation of singular integrals arising in time-domain integral equation analysis of electromagnetic scattering. *IEEE Transactions on Antennas and Propagation* (submitted):1996.

44. Rynne, B.P. Comments on "A Stable Procedure to Calculate the Transient Scattering by Conducting Surfaces of Arbitrary Shape". *IEEE Transactions on Antennas and Propagation* 41:517-519, 1993.

45. Leung, C.Y. *Computational Methods for Integral Equation Analysis*, London:PhD Thesis, University of London, 1995.

46. Pocock, M.D *Integral equation methods for harmonic wave propagation*, PhD Thesis, University of London, 1995.

47. Woo, A.C., Wang, H.T.G. and Schuh, M.J. Benchmark radar targets for the validation of computational electromagnetics programs. *IEEE Ant.Prop.Mag* 35:84-89, 1993.

48. Walker, S.P. Scattering Analysis via Time Domain Integral Equations; Methods to reduce the Cost Scaling of Costs with Frequency. *IEEE Ant.Prop.Mag*(submitted):1996.

49. 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* 1996.(In Press)

50. Richardson, H. *High performance Fortran History, Overview and Current Status.* Technical Report, Edinburgh Parallel Computing Centre, University of Edinburgh, September 1995.

51. Gropp, W., Lusk, E. and Skjellum, A. *Using MPI: Portable Programming with the Message-Passing Interface,* the MIT Press, 1994.

52. Cwik, T., Partee, J. and Patterson, J. Method of Moment Solutions to Scattering Problems in a Parallel Processing Environment. *IEEE Transactions on Magnetics* 27:3837-3840, 1991.

53. Cwik, T. Parallel decomposition methods for the solution of electromagnetic scattering problems. *Electromagnetics* 12:343-357, 1992.

54. Drake, J. and Gray, L.J. *Parallel Implementation of the Boundary Element Method on the IPSC2 hypercube for electroplating applications,* Tennessee:Oak Ridge National Laboratory ORNL-5615, 1988.

55. Daoudi, E.M. and Lobry, J. Implementation of a Boundary Element Method on Distributed Memory Computers. *Parallel Computing* 18:1317-1324, 1992.

56. Walker, S.P. and Leung, C.Y. Parallel computation of integral equation methods for three dimensional transient wave propagation. *Communications in Numerical Methods in Engineering* 11:515-524, 1995.

57. Walker, S.P. and Leung, C.Y. Parallel computation of time domain integral equation analyses of electromagnetic scattering and rcs. *IEEE Transactions on Antennas and Propagation* 1996.(In Press)

58. Sonneveld, P. CGS, A Fast Lanczos-Type Solver for Nonsymmetric Linear Systems. *Siam Journal of Scientific and Statistical Computing* 10:36-52, 1989.

59. Jakobus, U. and Landstorfer, F.M. Current-Based Hybrid Moment Method Analysis of Electromagnetic Radiation and Scattering Problems. *Applied Computational Electromagnetics Society Journal* 10:38-39, 1995.

60. Jakobus, U. and Landstorfer, F.M. Improved PO-MM Hybrid Formulation for Scattering from Three-Dimensional Perfectly Conducting Bodies of Arbitrary Shape. *IEEE Transactions on Antennas and Propagation* 43:162-169, 1995.

61. Jakobus, U. and Landstorfer, F.M. Improvement of the PO-MoM Hybrid Method by Accounting for Effects of Perfectly Conducting Wedges. *IEEE Transactions on Antennas and Propagation* 43:1123-1129, 1995.

62. Medgyesi-Mitschang, L.N. and Wang, D.-S. Hybrid Solutions for Scattering from Perfectly Conducting Bodies Revolution. *IEEE Transactions on Antennas and Propagation* 31:570-583, 1983.

63. Walker, S.P. and Lee, B.H. Termination criteria in iterative solution of large scattering problems using integral equation methods. *Communications in Numerical Methods in Engineering* (accepted):1996.