CONF-9510199--1

SAND95-0264C

Time Efficient 3-D Electromagnetic Modeling on Massively Parallel Computers

David L. Alumbaugh and Gregory A. Newman
Sandia National Laboratories
P.O. Box 5800 MS 0750
Albuquerque, NM 87008

Abstract

A numerical modeling algorithm has been developed to simulate the electromagnetic response of a three dimensional earth to a dipole source for frequencies ranging from 100Hz to 100MHz. The numerical problem is formulated in terms of a frequency domain - modified vector Helmholtz equation for the scattered electric fields. The resulting differential equation is approximated using a staggered finite difference grid which results in a linear system of equations for which the matrix is sparse and complex symmetric. The system of equations is solved using a preconditioned quasi-minimum-residual method.

Dirichlet boundary conditions are employed at the edges of the mesh by setting the tangential electric fields equal to zero. At frequencies less than 1MHz, normal grid stretching is employed to mitigate unwanted reflections off the grid boundaries. For frequencies greater than this, absorbing boundary conditions must be employed by making the stretching parameters of the modified vector Helmholtz equation complex which introduces loss at the boundaries.

To allow for faster calculation of realistic models, the original serial version of the code has been modified to run on a massively parallel architecture. This modification involves three distinct tasks; 1) mapping the finite difference stencil to a processor stencil which allows for the necessary information to be exchanged between processors that contain adjacent nodes in the model, 2) determining the most efficient method to input the model which is accomplished by dividing the input into "global" and "local" data and then reading the two sets in differently, and 3) deciding how to output the data which is an inherently non-parallel process.

1 Introduction

Great strides have been made over the last decade in finite difference electromagnetic (EM) modeling of three dimensional (3-D) structures for geophysical purposes. Druskin and Knizherman (1988 and 1994), Smith (1992), Wang and Hohmann(1993), and Newman (1995) all employ some type of staggered finite difference grid (Yee,1966) to solve for the EM fields in both the time and/or frequency domain. However, even with these computationally efficient solutions, the complexity, and thus the realism of the models that can be simulated on traditional serial computers is limited by memory and flop rate of the processor.

With the rapid advancements in massively parallel computers the limitations posed by serial computers is disapearing. This is due to the fact that the rate at which the simulations can proceed is dramatically increased because thousands of processors can operate on the problem simultaneously. In this paper we examine the implementation of a frequency domain-finite difference (FD-FD) scheme on a parallel platform and demonstrate its usefulness over a wide frequency range for different types of geophysical scenarios.

2 Theoretical Devlopment

In order to simulate the EM response of a 3-D earth, we numerically solve the frequency domain version of the vector Helmholtz equation for the scattered electric fields using a finite difference approximation on a staggered grid(Yee ,1966). The FD solution we shall outline has been designed

to compute the 3-D EM response for a wide variety of earth properties for frequencies ranging from approximately 100 hz up to 100 Mhz. This scheme is similar to those outlined in Alumbaugh and Newman (1994) and Newman and Alumbaugh (1995a), but has been extended as described in Alumbaugh et al.(1995) to include both variable magnetic permeability as well as absorbing boundary conditions (ABCs). The ABCs are required to simulate the response for frequencies greater that 10 Mhz as without them erroneous results are produced. We have chosen to employ the "perfectly matched layer" (PML) absorbing boundary conditions originally developed by Berenger(1993) for 2-D time-domain calculations and later modified for 3-D by Katz et al. (1994) and Chew and Weedon (1994). This method uses a modified form of the Helmholtz equation in which the absorption is incorporated through the use of complex grid stretching.

Because the responses we are interested in simulating, for example airborne EM simulations, often employ a dipole source located far away from zones of anomalous electrical properties, we have chosen to work with the scattered field versions of the governing equations. Often this allows us to employ a coarser discretization about the source location than would be employed with a total field solution and helps to limit storage overhead. In addition, because the scattered field versions of Maxwell's equations numerically decouple in the air at low frequencies (< 10MHz) we have chosen to work with the scattered electric field version of the modified Helmholtz equation instead of the coupled Maxwell system, which has the form

$$\nabla_h \times \frac{\mu_p}{\mu} \nabla_e \times \mathbf{E^S} = -i\omega\mu_p(\sigma + i\omega\epsilon)\mathbf{E^S} - i\omega\mu_p[(\sigma - \sigma_p) + i\omega(\epsilon - \epsilon_p)]\mathbf{E^P} - i\omega\mu_p \nabla_h \times [(\frac{\mu - \mu_p}{\mu})\mathbf{H^P}] \quad (1)$$

where

$$\nabla_e = \mathbf{i}\frac{1}{e_x}\frac{\partial}{\partial x} + \mathbf{j}\frac{1}{e_y}\frac{\partial}{\partial y} + \mathbf{k}\frac{1}{e_z}\frac{\partial}{\partial z} \quad (2)$$

and

$$\nabla_h = \mathbf{i}\frac{1}{h_x}\frac{\partial}{\partial x} + \mathbf{j}\frac{1}{h_y}\frac{\partial}{\partial y} + \mathbf{k}\frac{1}{h_z}\frac{\partial}{\partial z}. \quad (3)$$

In these equations the electrical conductivity , magnetic permeability and dielectric permitivity are denoted by σ, μ, and ϵ, respectively, with the 'p' designating a whole-space background value, $\mathbf{E^S}$ and $\mathbf{E^P}$ are the scattered and primary electric field whose sum is equal to the total electric field, and e_i and h_i for i=x,y,z are coordinate stretching variables which stretch the x,y, and z coordinates. As shown in Chew and Weedon (1994), when e_i and h_i are complex then the medium is perfectly absorbing. Note that the terms at the end of equation (1) are "equivalent source" terms which are calculated wherever the properties of the medium are different from that of the assumed background. The boundary conditions employed are Dirichlet conditions, i.e., the tangential component of $\mathbf{E^S}$ is set to zero on the grid boundary.

The scattered electric fields are assigned to each cell following the staggered grid scheme given in Figure 1. For node (i, j, k), the x, y and z components of electric field are sampled at $(i + 1/2, j, k)$, $(i, j+1/2, k)$ and $(i, j, k+1/2)$, respectively. For modeling the coupled Maxwell's equations this corresponds to assigning the electric fields to the edges of the cell and the magnetic field to its faces. In addition this formulation requires that the conductivity and dielectric permitivity be computed halfway along a given cell edge in Figure 1, and the magnetic pemeability to be computed in the center of the cell face. This is accomplished through the simple averaging schemes described in Alumbaugh et al. (1995).

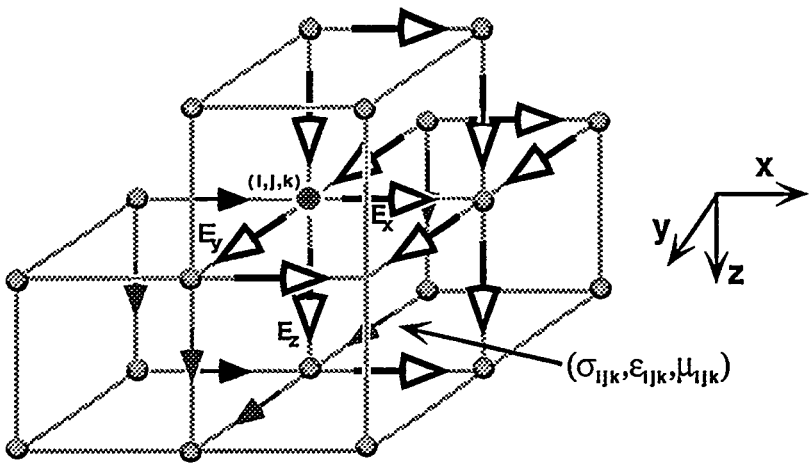

Figure 1: Finite difference stencil for solving the scattered electric field Helmholtz equation. The arrows represent the position of the electric fields, with the open arrows representing those unknowns needed to form the equation for Ex.

After numerically discretizing equation (1) to form the finite difference equations a linear system is assembled,

$$\mathbf{Af} = \mathbf{s} \tag{4}$$

where **A** is the stiffness matrix containing the numerical approximations to the derivatives as well as the electrical properties of the medium, **f** is the unknown vector for the scattered electric field and **s** is the equivalent source vector. Alumbaugh et al.(1995) show that **A** is complex symmetric, even when complex grid stretching is employed. The solution vector can be obtained using the quasi-minimum residual (QMR) (Freund ,1992) technique with preconditioning to iteratively determine the solution within a predetermined error leveli, which is defined here to be

$$er = \frac{\|\mathbf{Af} - \mathbf{s}\|^2}{\|\mathbf{s}\|^2}. \tag{5}$$

Tests with different types of incomplete decomposition and polynomial preconditioners has found that simple Jacobi scaling provides a simple, time efficient method of preconditioning.

After the scattered fields at the grid points have been determined, the fields at the receivers must be calculated. The electric field is simply calculated using bi-linear interpolation while the magnetic field is calculated by first taking a numerical approximation of Faraday's law for the scattered electric fields on the grid surrounding the receiver,

$$\nabla_e \times \mathbf{E^S} = -i\omega\mu\mathbf{H^S} + (\mu - \mu_p)\mathbf{H^P} \tag{6}$$

and then interpolating the result to the point of interest. In this expression $\mathbf{H^S}$ and $\mathbf{H^P}$ are the scattered and primary magnetic fields. Note, for each new source position and/or frequency a new system must be solved, although some time savings can be implemented by using the previous solution vector as an initial guess.

3 Properties of the PML Absorbing Boundary Condition

Although their calculations employ the coupled modified Maxwell's equations in the time domain, Chew and Weedon (1994) develop theory in the frequency domain to demonstrate how lossy, non-reflecting conditions are created along the mesh boundaries. The complex stretching parameters are assigned a value of the form $1+a-ib$. On the internal portion of the mesh, $a = b = 0$ such that the modified Helmholtz equation reduces to the normal form. Near the edges of the mesh $\mathbf{E^S}$ and $\mathbf{E^P}$ are allowed to vary over several cells, but only in the direction that is perpendicular to the boundary. For example along the +z boundary $e_x = e_y = h_x = h_y = 1$ and only e_z and h_z are allowed values of a and b that are not equal to zero.

Because we are solving an implicit rather than explicit system, we have found that in order to incorporate a given amount of loss, or attenuation, across a number of cells serving as the absorbing boundary, it is better to set a and b constant rather than gradually increasing their value toward the mesh boundaries as suggested by Berenger (1993); gradually increasing their value results in a greater number of iterations needed to achieve convergence. Simple MATLAB experiments have shown that this is due to the fact that the condition number of $\mathbf{A}$ increases as the ratio between the largest cell and smallest cell in the mesh increases. Thus gradually increasing the stretching parameters outward will produce a cell along the edge of the mesh which is effectively much larger than any of the cells employing constant stretching. Because the smallest cell size is the same in either case, the solution of the model that employs the gradual stretching will take longer to converge.

Currently, we are investigating methods for choosing optimal stretching parameters for a given frequency and background wave number, defined as

$$k_p = \sqrt{-i\omega\mu_p(\sigma_p + i\omega\epsilon_p)} = \alpha - i\beta. \tag{7}$$

where α and β are both real. This analysis is based on the assumption that the loss that is incoporated through complex grid stretching is caused by 'pseudo' electrical parameters within cells of constant size. Through this assumption we can develop a pseudo-skin depth in each cell which is defined by

$$\delta^{ps} = \frac{1}{(\alpha b + \beta a)} \tag{8}$$

and a pseudo-wavelength defined as

$$\lambda^{ps} = \frac{1}{(\alpha a - \beta b)}. \tag{9}$$

To this point we have found that for frequencies greater than 1MHz, accurate results are achieved when a and b are chosen such that three pseudo-skin depth's of attenuation are provided for across the stretching region without out significatly changing the pseudo wavelength from that of the natural background wavelength. At frequencies below 10 khz the analysis seems to become more difficult as the manner in which the grid is stretched can significantly alter the convergence of the system. In general at these frequencies we have obtained good results using only real grid stretching; i.e. setting b=0 and varying only a.

4 Implementation on Massively Parallel Computers

In order to simulate larger, more realistic models than has previously been possible, the original serial version of the code has been modified to run on massively parallel MIMD (multiple instruction multiple data) machines which can have thousands of processors. These parallel machines are

employed by assigning a given number of processors in each direction of the model (nx in x, ny in y and nz in z) and then breaking up the model across the processor bank such that each individual processor is in charge of a 3-D subset. Because each processor needs only to make the necessary calculations for this subset, and because all of the processors are making their appropriate calculations simultaneously, the solution time is reduced by a factor which is approximately equal to the total number of processors employed (nx*ny*nz).

The first step in converting the serial version of the code to a parallel version is to divide the problem up among the processors such that it is optimally load balanced. This preprocessing step is necessary to ensure that large banks of processors are not standing idle for long periods of time while a single or small number of processors complete their calculations. As one would imagine, this type of scenario is an extremely inefficient use of resources. Rather the problem is broken up such that each processor has as close to an equal number of unknowns as possible for which to solve. The second issue that needs to be addressed is the manner in which the model is input. To accomplish this, we have decomposed the input data into two different sets: a global data set and a local data set. Global data are those variables that each processor needs to know such as the source and receiver positions, the frequencies, what type of solver is being employed, the location of the mesh nodes, ect. These form a fairly small data set which can easily be read in by a "lead" processor and then "broadcast" to all other processors. The second type of input is the local data, or local model parameters (conductivity, dielectric permitivity and magnetic permeability) that are assigned to each cell within the model. Because each processor needs only a small subset of this data and contains only a small amount of local memory, the local data is broken up into multiple files, one for each processor, which are then read in individually from a parallel disk system which allows several files to be read in simultaneously.

After the data have been accessed, each processor constructs its own portion of the stiffness matrix A and the source vector s, and then proceeds to solve for its portion of the solution vector. However, each iteration within the QMR solver requires one matrix- vector multiply and several vector dot products. These operations pose problems because in order to complete them, information must be exchanged both between all of the processors as well as small subsets of processors. The dot products are fairly easy to implement as they involve; 1) a local calculation in which each processor computes the dot-product of its portion of the vector and 2) a global calculation in which all the local calculations are "gathered" by the lead processor, summed, and the result broadcast across the machine.

The vector-matrix multiply is more difficult to implement because it requires that each processor communicates with those "neighboring" processors that are solving for the scattered electric fields in adjacent portions of the model. Determining these neighboring processors and the actual unknowns that need to be communicated is accomplished in the following manner. If we assume that each processor contains only a single node, then we can imagine it as a cubic shape enclosing node (i, j, k) as well as all other nodes in Figure 1. Careful examination then indicates that there are two types of communication that each processor needs to execute with its appropriate neighbors. The first type of communication will occur across the "faces" of the cube. For node (i, j, k) this implies communication with those nodes directly connected to it by the gray lines of the finite difference stencil, i.e., nodes $(i-1, j, k), (i+1, j, k), (i, j-1, k), (i, j+1, k), (i, j, k-1)$ and $(i, j, k+1)$. For these communications either two or three unknowns are exchanged per nodal position. The second type of communication occurs across certain "edges" of the cube, and involves those nodes which are not directly connected to (i, j, k) by the stencil lines, for example node $(i+1, j, k-1)$. This type of communication requires only one unknown per node being communicated each way. If we now expand the idea such that each processor cube contains a 3-D distribution of nodal points, then we can develop the processor communication stencil shown in Figure 2.

The last issue to be addressed is the data output. Because for any given source we only need

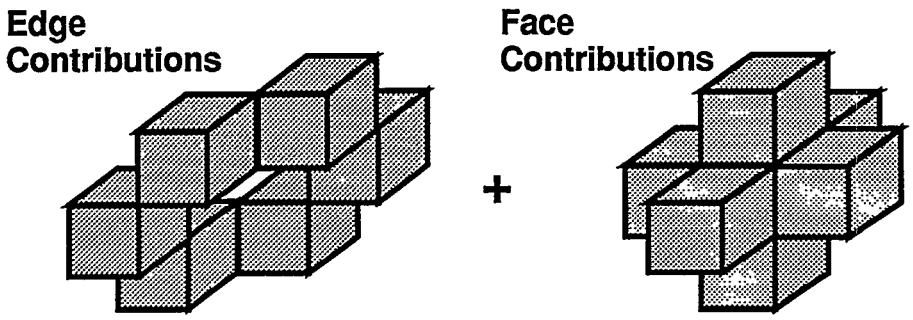

Figure 2: Processor stencil employed for message passing in order to correctly complete the matrix-vector multiply.

to know the results at only a limited number of receiver positions, all of which may lie on the same processor, the data output is inherently non-parallel and is accomplished in the following steps. 1) Each processor determines which processor holds the portion of the model that contains the receiver position. 2) This "receiver" processor then determines if it need any values from adjacent processors, completes the necessary point-to -point communication with those processors, and then does the necessary bilinear interpolation. 3) The results are then sent to the lead processor which outputs them to disk.

To this point the code has been implemented on two different MIMD machines available at Sandia National Laboratories, the 1840 processor Intel Paragon and 1024 processor NCUBE, and run time characteristics for the Paragon are given below. To provide for the required message passing on these two machines we have chosen to employ the Message Passing Interface (MPI, Skjellum et al.,1993) rather than using machine specific commands. This provides portability to the code as it will be able to run on any parallel machine and/or distributed network of machines on which this public domain library is available.

5 Demonstration of the Finite Difference Solution

To illustrate the versatilty and speed of the numerical solution when implemented on a parallel platform, we have simulated two different models which represent measurement configurations that might be employed in the field and employ a wide range of frequencies. The first simulation will be invlove frequncies in the low GPR range while the second will simulate a portion of helicopter EM survey. In the two cases the Krylov solver was assumed to have converged to an adequate error level when equation (5) was found to be less than or equal to 10^{-7} and 10^{-8}, respectively. These error levels are empirical and are based on extensive comparisons of the solution with other numerical solutions and scale model experiments (eg., Alumbaugh and Newman, 1994).

5.1 High Frequency Simulation for the 'VETEM' project

The 'VETEM" (Very Early Time ElectroMagnetic) project is an attempt to build an electromagnetic prospecting system that operates above traditional geophysical induction frequencies (100 kHz) yet below ground penetrating Radar frequencies (100 MHz) (Pellerin et al., 1995).To illustrate the ability of the code to simulate the electromagnetic response at these frequencies, variations of the model shown in Figure 3 has been employed. This example was designed to simulate a test site at the Colorado School of Mines where a prototype of the VETEM system known as the High Frequency Sounder (HFS) (Stewart et al., 1994) was first tested. The model is particularly difficult to simulate

because of two conflicting conditions that are imposed by the material properties; 1) the wavelength in the block at 28.5 MHz is approximately 1.6m which requires a maximum cell dimension of 0.16m to avoid grid dispersion (Chew, 1990, p 244) and 2) the skin depth in the first layer at that same frequency is 17.8m which requires the boundaries to be placed very far away to avoid reflections off the grid. The small cell size coupled with the large distance to the boundaries produces a very large mesh if no absorbing boundary conditions are employed.

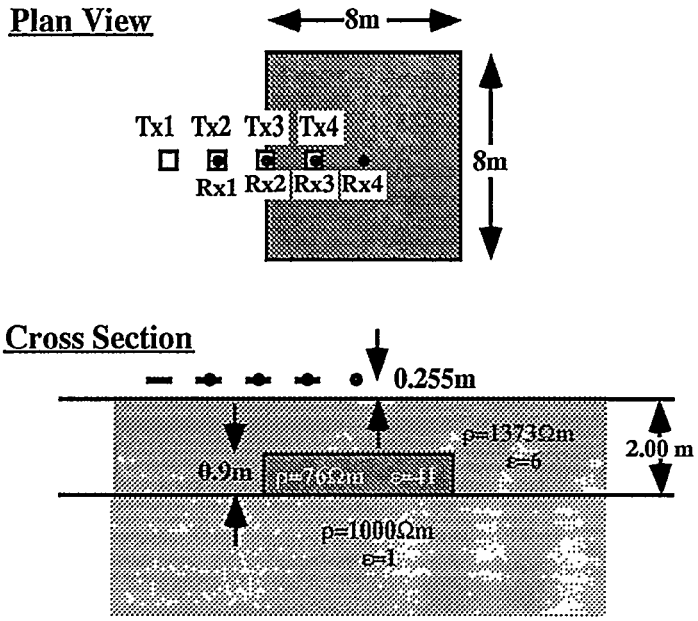

Figure 3: The Colorado School of Mines 3-D (CSM3D) model.

To simulate this example a 120 x 120 x 120 cell mesh was employed with a constant cell size of 0.15m in the x and y directions. This places the total distance across the mesh at 18m. In z, the maximum cell size was also 0.15m, with a minimum cell size of 0.13m to accommodate the layer thicknesses. Note, this mesh produces a total of $5x10^6$ unknowns for which to solve, which is much too large a problem for all but a supercomputer. The VMD source was placed at the center of the mesh in x and y, i.e., 9m from each boundary, and a background conductivity of $\sigma=10^{-16}$S/m was assumed. To incoporate loss and thus avoid reflections off the edges of the mesh, b was set equal to 0.6 over 20 cells along each edge of the mesh.

In the first case we simulate two 1-D models, a two layer which assumes the block is absent, and a 3 layer model which assumes that the block extends to infinity in the x and y directions. This allows us to compare to a 1-D code devloped by Ki Ha Lee at Lawrence Berkeley Laboratory. Because HFS directly measures tilt angle and ellipticity of the magentic field (Smith and Ward, 1974) the results have been plotted in terms of these parameters rather than amplitude and phase of the different components. As can be seen in Figure 4a for the layered models, the 3-D code reproduces the 1-D calculations extremely well.

In Figure 4b the 3-D results for four different source positions are plotted with the results for the two 1-D models. Notice that the 3-D responses never reproduce the 1-D results even when the source-receivers are completely outside or within the block. This indicates that 3-D effects are measureable at greater distances than is immediately evident and that 1-D inversions probably would not accurately reproduce the structure of the subsurface.

To demonstrate some of the questions that must be answered when using the parallel machines,

the solution time as well as the flop rate has been plotted against the number of processors employed for the 3 layer model at 10.1 Mhz on the Intel Paragon. Figure 5 shows that a large decrease in run time occurs with an increasing number of processors from 200 up to 1000. This corresponds to solving for 24000 to 3000 unknowns per processor.and indicates that the processors are spending the majority of their time performing calculations rather than communicating. However, the relatively small decrease in run time with increasing number of processors over 1000 indicates the solution time is begining to be dominated by message passing if less than 3000 unknowns are being solved for on each processor. Thus we are left with a decision to make. If we wish to use the machine most efficiently, we would employ less than 1000 processors such that the internal computations are dominating the solution time. We could then run several jobs simultaneously such that the efficiency increases proportionally to the number of jobs. On the other hand if we desire as quick a turn around time as possible for a single computation, then we would want to operate near the right end of the curve.

5.2 Airborne Simulation

The second example simulates an helicopter EM survey flown to define the location of a buried paleo-channel through which conductive salt water is migrating, and is designed after a survey flow in Austrailia in the early 1990's (Doug Frazer, personnel communication). Figure 6 shows a plan view of the model at 5m depth below the earths surface as well as two cross sections through the model. The flight lines are 30m above the earths surface, are spaced at 200m intervals from top to bottom in Figure 6, and along each line the sampling interval is 100m. This yields a total of 187 source positions. A VMD source is operating at 0.9 kHz, 7.2 kHz and 56 kHz, with the receiver located 8m to the right of the source. The three frequencies coupled with the 187 positions yields a total of 561 forward solves.

To calculate this with the 3-D finite difference code, the earth and air were divided into a 208 x 184 x 49 cell grid which yields a total of 5.6×10^6 unknowns for which to solve. To avoid reflections off the mesh boundaries normal grid stretching (i.e. b=0) was employed to move them out to 400m from the nearest sampling point. The smallest cell size employed was 5m x 5m by 2.5m and was employed at the air-earth interface underneath each source array. The largest cell size employed was in the corners of the mesh and was 20m x 20m x 20m. A background conductivity of $s=10^{-16}$S/m was assumed to simulate the electrical properties of the air.

Notice in Figure 7 that for all three frequencies the channel is clearly defined, although its resisitivity is not as accurately defined at higher frequencies. This is due to the increased sensitivity at higher frequencies to the surface resistivie layer. To run this model on 1360 processors of the Intel Paragon took approximately two days. This reasonably quick turn around time for this complicated model illustrates the utility of these machines for solving realistic geologic problems. In addition, although this soluttion may of taken too much time to be employed for inversion of this large of a model, the near future development of machines with tens-of -thousands of processors should allow for it.

6 Conclusions and Discussion

In this paper we have presented a scheme to solve for the frequency domain electromagnetic response of a 3-D earth over a wide band of frequencies using massively parallel computers. The problems associated with providing absorbing boundary conditions and porting the serial version of the scheme to a parallel machine have been outlined, and two simulations have been included to demonstrate the versatilty of the code.

Although the demostrations shown here would have been nearly impossible prior to the parallel implementation due to the size of the models and/or the number of frequencies and sources involved, we believe that there is still much research to be done with regards to the implementation of this type of scheme. The most notable location for improvement is in the area of preconditioners. We are currently considering the use of multigrid preconditioners, and methods to separately treat the real and imaginary components of the matrix system. In addition, a scheme to accelerate the convergence for very low frequency simulations where channeling currents dominate needs to be developed in order to simulate natural field measurements as well as extend the frequency band down below 100 hz; Smith (1992) has found that a static correction can be incorporated to accommodate this. Finally, methods of dealing with the air-earth interface need to be more closely examined. We have found that this interface tremendously complicates the numerical problem, especially when electric dipole sources are employed on the surface.

Lastly, it must be mentioned that the development of this forward modeling code is not the final product. Rather this will be used as the forward solver within a 3-D non-linear electromagnetic inversion code being developed to image the electrical properties of the earth (Newman and Alumbaugh,1995b). This is the end product which has driven us to port the finite difference scheme to a parallel platform as without these large computers, full 3-D inversion would be impossible.

7 Acknowledgements

We express our thanks to Dr. Ki Ha Lee of Lawrence Berkeley Laboratory for the use of the 1-D layered earth code used in the model comparisons. This work was performed at Sandia National Laboratories, which is operated for the U.S. Department of Energy. Funding for this project was provided by 1) DOE's office of Basic Energy Sciences, Division of Engineering and Geoscience under contract DE-AC04-94AL85000, 2) the Sandia National Laboratory Director's Research funds and 3) the VETEM project which is supported by the DOE Office of Technology and Development.

References

Alumbaugh, D. L. and Newman G. A., 1994, Fast, frequency domain electromagnetic modeling using finite differences: 64th Ann. Intern. Mtg., Soc. Explor. Geophys., Los Angeles CA, Expanded Abstracts, 369-373.

Alumbaugh, D.L., and Newman, G.A., Prevost, L., and Shadid, J. N., 1995, Three dimesional wide band electromagnetic modeling on massively parallel computers; submitted to Radio Science for review.

Berenger, J., 1993, A perfectly matched layer for the absorption of electromagnetic waves: J. Comp. Phy., 114, 185-200.

Chew, W. C., 1990, Waves and Fields in Inhomogeneous Media, Van Nostrand Reinhold International Company Limited.

Chew, W.C., and Weedon, W.H., 1994, A 3D perfectly matched medium from modified Maxwell's equations with stretched coordinates; Microwave and Optical Tech. Let.,7, 599- 604.

Druskin, V. and Knizhnerman, L., 1988, A spectral semi-discrete method for the numerical solution of three-dimensional nonstationary problems of electric prospecting: Izvestiya, Earth Physics, 24, 641-648.

Druskin, V. and Knizhnerman, L., 1994, A spectral approach to solving three-dimensional diffusion Maxwell's equations in the time and frequency domains: Radio Science, 29, 937- 953.

Freund, R., 1992, Conjugate gradient type methods for linear systems with complex symmetric coefficient matrices: SIAM J. Sci. Statist. Comput., 13, 425-448.

Katz. D.S., Thiele, E.T., and Taflove, A.,1994, Validation and extension to three dimensions of the Berenger PML absorbing boundary condition for FD-TD meshes; IEEE Micro. and Guided Wave Let., 4,268-270.

Newman, G. A., 1995, Cross well electromagnetic inversion using integral and differential equations: Geophysics, 60,899-910.

Newman, G. A., and Alumbaugh, D. L., 1995a, Frequency-domain modeling of airborne electromagnetic responses using staggered finite differences: accepted for publication by Geophysical Prospecting.

Newman, G. A., and Alumbaugh, D. L., 1995b, 3D massively parallel electromagnetic inversion: International Symposium on Three-Dimensional Electromagnetics, Oct. 4-6, Schlumberger-Doll Resaerch, Ridgefield, Conneticut, USA.

Pellerin, L., Labson, V. F., and Pfeifer, M. C., 1995, VETEM - a very early time electromagnetic system; to be presented at SAGEEP '93 (Symposium on the Application of Geophysics to Environmental and Engineering Problems), Orlando, Florida.

Skjellum, A., Doss, N.E., and Bangalore, P.V., 1993, Writing libraries in MPI, in A. Skjellum and D. Reese, Ed., Proceedings of the Scalable Parallel Libraries Conference, IEEE Comput. Sci. Press, 166-173.

Smith, T. J., 1992, Conservative modeling of 3-D electromagnetic fields; International Association of Geomagnetism and Aeronomy, 11th Workshop on Electromagnetic Induction in the Earth, Wellington, New Zealand, Meeting Abstracts.

Stewart, D.C., Anderson, W.L., Grover, T.P., and Labson, V.F., 1994, Shallow subsurface mapping by electromagnetic sounding in the 300 khz to 30 Mhz range: model studies and prototype system assessment; Geophysics,59,1201-1210.

Smith, B. D., and Ward, S. H., 1974, On the computation of polarization ellipse parameters: Geophysics, 39, 867-869.

Wang T. and Hohmann, G. W., 1993, A finite difference time-domain solution for three- dimensional electromagnetic modeling: Geophysics, 58, 797-809.

Yee, K. S., 1966, Numerical solution of initial boundary problems involving Maxwell's equations in isotropic media: Inst. of Electr. and Electron. Eng., Trans. Ant. Prop, AP- 14, 302-309.

DISCLAIMER

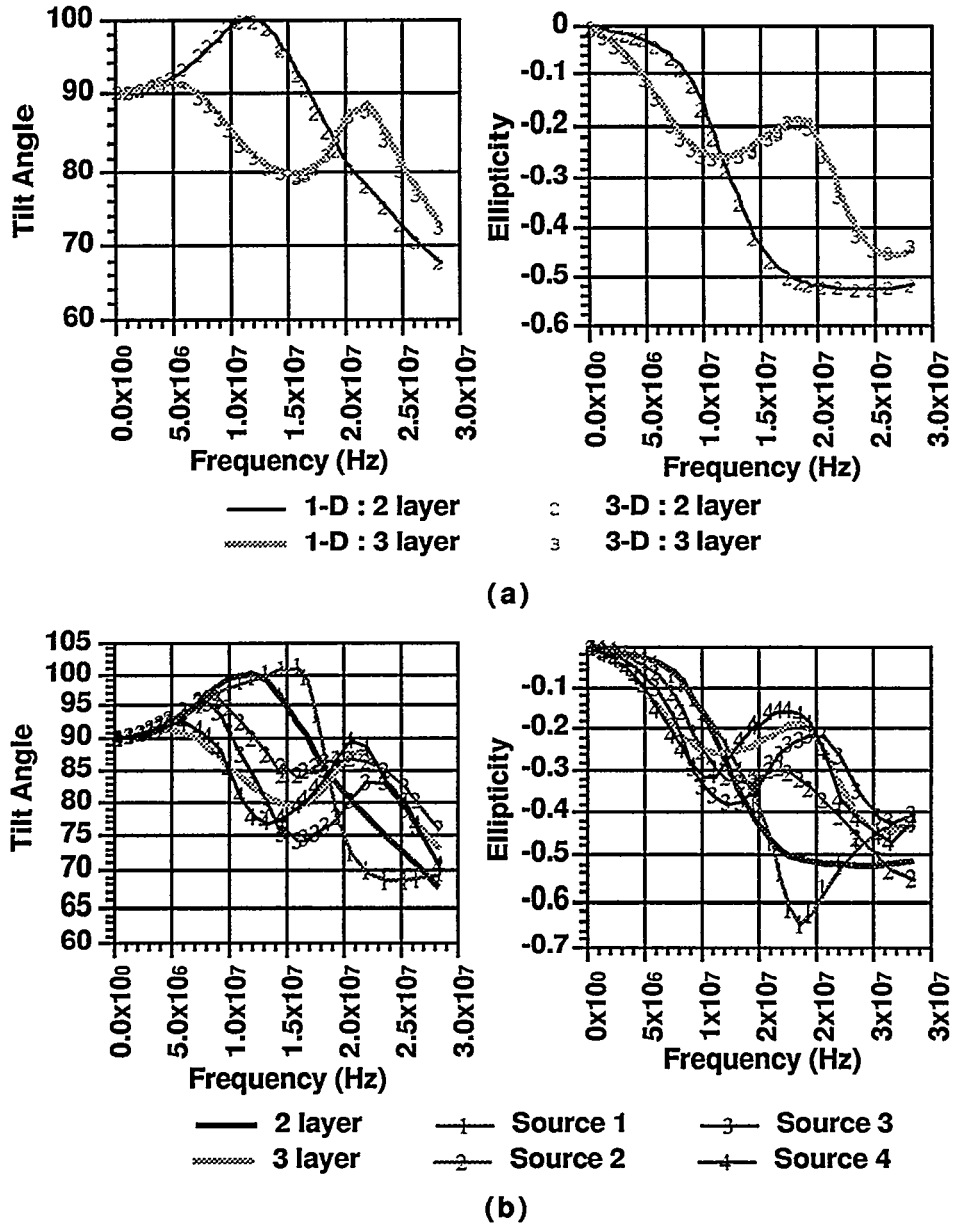

Figure 4: Results for the CSM3D model. a) 1-D comparisions for 2 and 3 layer models. b) 3-D comparison for different source positions.

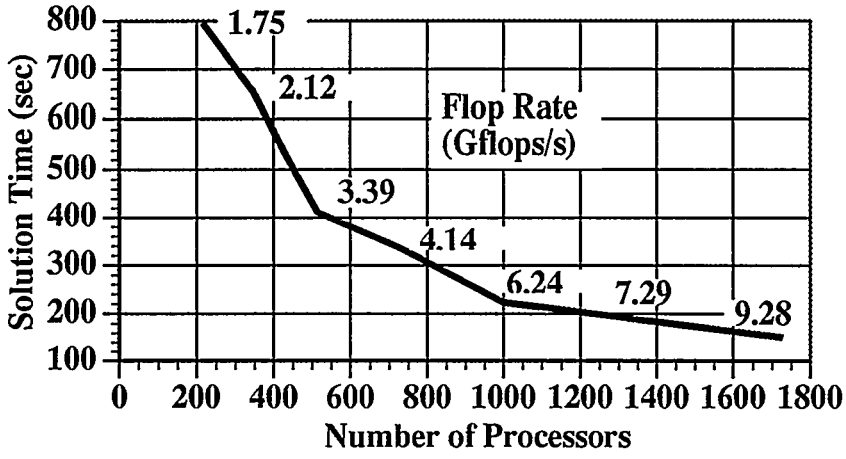

Figure 5: Run time versus number of processors employed for the 3 layer model at 10.1 MHz.

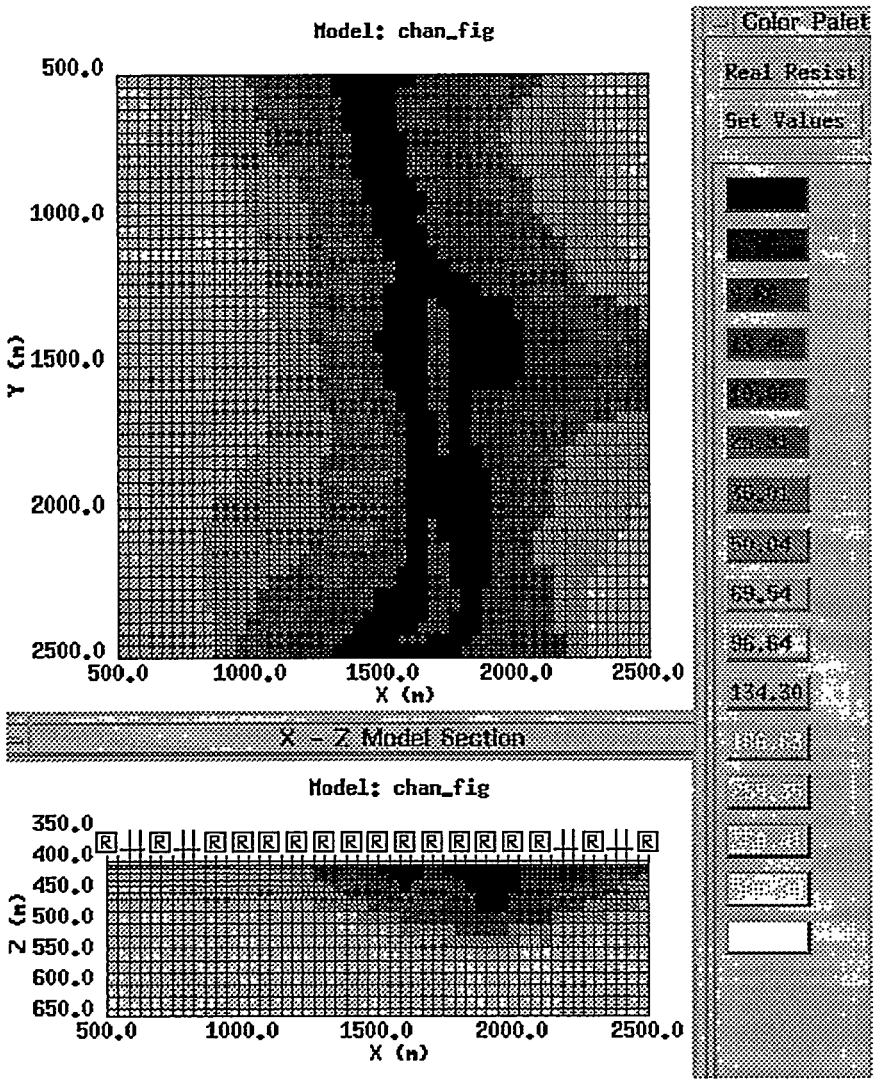

Figure 6: The subsurface channel model employed for the airborne simulation. The upper figure is a plan view at the top of the channel, and the bottom figure is a vertically exagerated cross-section at y=1500 m. Although it is difficult to see, a 5 m thick- 500 Ωm layer exists from the earth's surface down to the top of the channel. The gray scale varies logarithmically from 5 Ωm (black) to 697 Ωm (white).

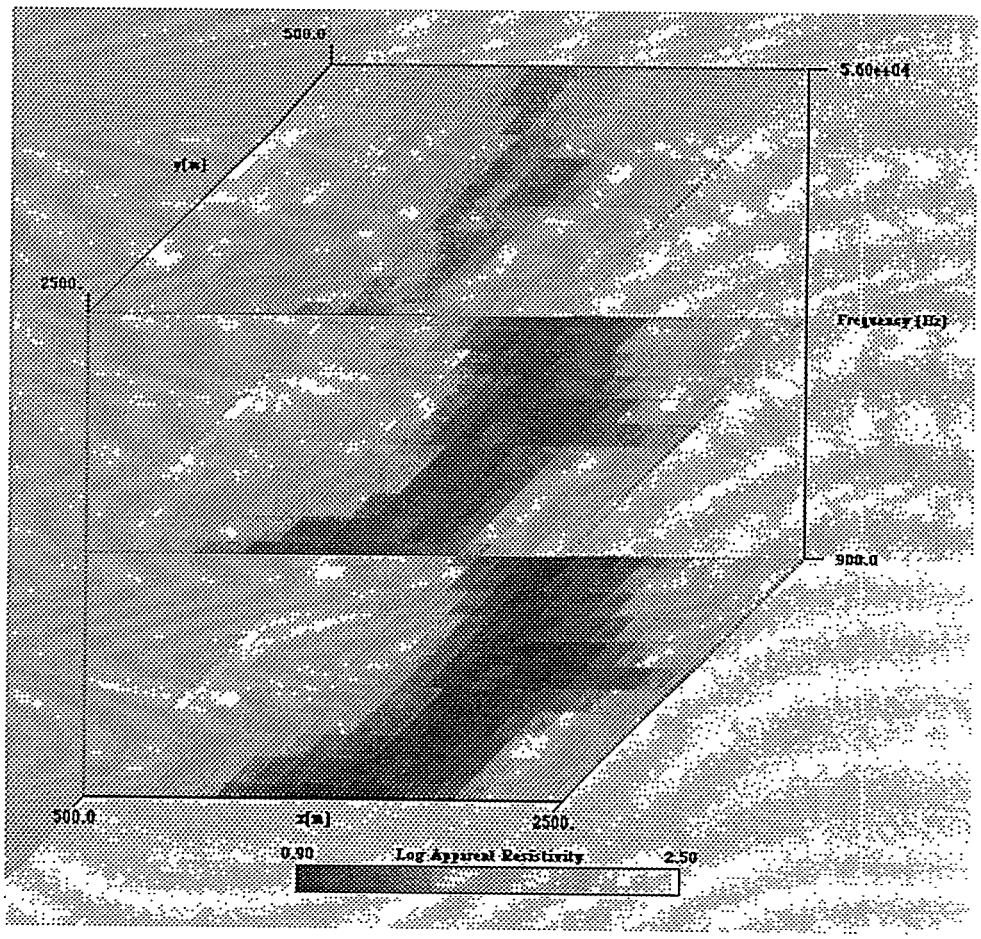

Figure 7: Calculated apparent resistivities for the subsurface channel model.