Fast finite volume simulation of 3D electromagnetic problems with highly discontinuous coefficients

E. Haber* U. M. Ascher†

March 27, 2000

Abstract

We consider solving three-dimensional electromagnetic problems in parameter regimes where the quasi-static approximation applies and the permeability, permittivity and conductivity may vary significantly. The difficulties encountered include handling solution discontinuities across interfaces and accelerating convergence of traditional iterative methods for the solution of the linear systems of algebraic equations that arise when discretizing Maxwell's equations in the frequency domain.

In order to address the problem of slow convergence we reformulate Maxwell's equations in terms of potentials, applying a Helmholtz decomposition to either the electric field or the magnetic field. The null space of the **curl** operators can then be annihilated by adding a stabilizing term, using a gauge condition, thus obtaining a strongly elliptic differential operator. A staggered grid finite volume discretization is subsequently applied to the reformulated PDE system. This scheme works well for sources of various types, even in the presence of strong material discontinuities in both conductivity and permeability. The resulting discrete system is amenable to fast convergence of ILU-preconditioned Krylov methods.

*Departments of Computer Science and of Earth and Oceanography Sciences, University of British Columbia, Vancouver, BC, V6T 1Z4, Canada. (haber@cs.ubc.ca).

†Department of Computer Science, University of British Columbia, Vancouver, BC, V6T 1Z4, Canada. (ascher@cs.ubc.ca).

We test our method using several numerical examples and demonstrate its robust efficiency. We also compare it to the classical Yee method using similar iterative techniques for the resulting algebraic system, and show that our method is significantly faster, especially for electric sources.

Keywords: Maxwell's equations, solution discontinuities, Helmholtz decomposition, Coulomb gauge, finite volume, Krylov methods, mixed methods, preconditioning.

1 Introduction

The need for calculating fast, accurate solutions of three-dimensional electromagnetic equations arises in many important application areas including, among others, geophysical surveys and medical imaging [29, 32, 2]. Consequently, a lot of effort has recently been invested in finding appropriate numerical algorithms. However, while it is widely agreed that electromagnetic phenomena are generally governed by Maxwell's equations, the choice of numerical techniques to solve these equations depends on parameter ranges and various other restrictive assumptions, and as such is to a significant degree application-dependent [20, 32, 2].

The present article is motivated by remote sensing inverse problems, e.g. in geophysics, where one seeks to recover material properties – especially conductivity – in an isotropic but heterogeneous body, based on measurements of electric and magnetic fields on or near the earth's surface. The forward model, on which we concentrate here, consists of Maxwell's equations in the frequency domain over a frequency range which excludes high frequencies. Assuming a time-dependence $e^{-\imath\omega t}$, these equations are written as

$$\boldsymbol{\nabla} \times \boldsymbol{E} - \imath\omega\mu \boldsymbol{H} = \boldsymbol{0}, \tag{1a}$$

$$\boldsymbol{\nabla} \times \boldsymbol{H} - \widehat{\sigma} \boldsymbol{E} = \boldsymbol{J}^s, \tag{1b}$$

$$\nabla \cdot (\epsilon \boldsymbol{E}) - \rho = 0, \tag{1c}$$

$$\nabla \cdot (\mu \boldsymbol{H}) = 0, \tag{1d}$$

where μ is the magnetic permeability, σ is the conductivity, ϵ is the electrical permittivity,

$$\widehat{\sigma} = \sigma - \imath\omega\epsilon, \tag{2}$$

$\boldsymbol{J}^s$ is a known source current density, and ρ is the (unknown) volume density of free charges. In our work we assume that the physical properties $\mu > 0$, $\epsilon > 0$ and $\sigma \geq 0$ can vary with position, and $\mu\epsilon\omega^2 L^2 \ll 1$, where L is a typical length scale. The electric field $\boldsymbol{E}$ and the magnetic field $\boldsymbol{H}$ are the unknowns, with the charge density defined by (1c). Note that as long as $\omega \neq 0$, (1d) is redundant and can be viewed as an invariant of the system, obtained by taking the $\nabla\cdot$ of (1a). The system (1) is defined over a three-dimensional spatial domain Ω. In principle, the domain Ω is unbounded (i.e. $\Omega = \mathbb{R}^3$), but, in practice, a bounded subdomain of $\mathbb{R}^3$ is used for numerical approximations. In this paper we have used the boundary conditions

$$\boldsymbol{H} \times \boldsymbol{n}\Big|_{\partial\Omega} = 0 \tag{3}$$

although other boundary conditions are possible.

A number of difficulties arise when attempting to find numerical solutions for this three-dimensional PDE system. These difficulties include handling regions of (almost) vanishing conductivity, handling different resolutions in different parts of the spatial domain, handling the multiple scale lengths over which the physical properties can vary, and handling regions of highly varying conductivity, magnetic permeability, or electrical permittivity where jumps in solution properties across interfaces may occur.

On the other hand, the nature of the data (e.g. measurements of the electric and/or magnetic fields at the surface of the earth) is such that one cannot hope to recover to a very fine detail the structure of the conductivity σ or the permeability μ. We therefore envision, in accordance with the inverse problem of interest, a possibly nonuniform tensor product grid covering the domain Ω, where $\widehat{\sigma}$ and μ are assumed to be smooth or even constant inside each grid cell, but they may have significant jump discontinuities that can occur anywhere in Ω across cell interfaces. The source $\boldsymbol{J}^s$ is also assumed not to have jumps across interfaces. The relative geometric simplicity resulting from this modeling assumption is key in obtaining highly efficient solvers for the forward problem.

Denoting quantities on different sides of an interface by subscripts 1 and

2, it can be easily shown [33] that across an interface

$$\boldsymbol{n} \cdot (\widehat{\sigma}_1 \boldsymbol{E}_1 - \widehat{\sigma}_2 \boldsymbol{E}_2) = 0, \tag{4a}$$

$$\boldsymbol{n} \times (\boldsymbol{E}_1 - \boldsymbol{E}_2) = 0, \tag{4b}$$

$$\boldsymbol{n} \cdot (\mu_1 \boldsymbol{H}_1 - \mu_2 \boldsymbol{H}_2) = 0, \tag{4c}$$

$$\boldsymbol{n} \times (\boldsymbol{H}_1 - \boldsymbol{H}_2) = 0. \tag{4d}$$

These conditions imply that neither $\boldsymbol{E}$ nor $\boldsymbol{H}$ are continuous in the normal direction when $\widehat{\sigma}$, resp. μ, have a jump discontinuity across a cell face, and likewise, $\widehat{\sigma}\boldsymbol{E}$ and $\mu\boldsymbol{H}$ are not necessarily continuous in tangential directions. Care must therefore be exercised when numerical methods are employed which utilize these variables if they are to be defined where they are double-valued.

By far the most popular discretization for Maxwell's equations is Yee's method [37] (see discussions and extensions in [32, 26, 19]). This method employs a staggered grid, necessitating only short, centered first differences to discretize (1a) and (1b). In the more usual application of this method, the electric field components are envisioned on the cell's edges and the magnetic field components are on the cell's faces – see Fig. 1. It is further possible to

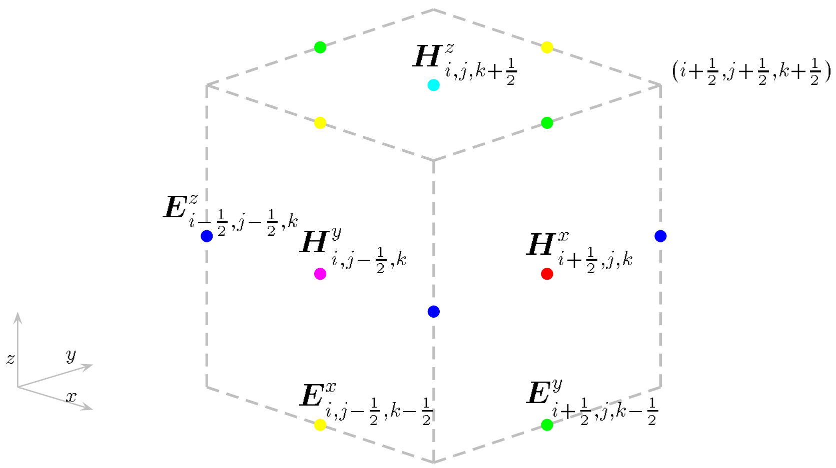

Figure 1: A staggered discretization of $\boldsymbol{E}$ and $\boldsymbol{H}$ in three dimensions: $\boldsymbol{E}$-components on edges, $\boldsymbol{H}$-components on faces.

eliminate the components of the magnetic field from the discrete equations, obtaining a staggered discretization for the second order PDE in $\boldsymbol{E}$,

$$\boldsymbol{\nabla} \times (\mu^{-1} \boldsymbol{\nabla} \times \boldsymbol{E}) - \imath\omega\widehat{\sigma}\boldsymbol{E} = \imath\omega \boldsymbol{J}^{\boldsymbol{s}}. \tag{5}$$

Related methods include the finite integration technique and certain mixed finite element methods [35, 5, 25]. Although these methods are often presented in the context of time-domain Maxwell's equations the issues arising when applying an implicit time-discretization (a suitable technique under our model assumptions) are often similar to the ones we are faced with here.

The popularity of Yee's method is due in part to its conservation properties and other ways in which the discrete system mimics the continuous system [19, 15, 5, 4]. However, iterative methods to solve the discrete system may converge slowly in low frequencies, due to the presence of the rich, nontrivial null space of the **curl** operator, and additional difficulties arise when highly discontinuous coefficients are present [29, 24, 16]. There are two major reasons for these difficulties. First, the conductivity can essentially vanish (for example, in the air, which forms part of Ω); from an analytic perspective, the specific subset of Maxwell's equations used typically forms an almost-singular system in regions of almost-vanishing $\widehat{\sigma}$. Even in regions where the conductivity is not close to vanishing, the resulting differential operator is strongly coupled and not strongly elliptic [6, 1]. Second, in cases of large jump discontinuities, care must be taken to handle $\boldsymbol{H}$ and $\widehat{\boldsymbol{J}} = \widehat{\sigma}\boldsymbol{E}$ carefully, since these are located as in Fig. 1 where they are potentially discontinuous.

In [1], we addressed the often slow convergence of iterative methods when used for the equations resulting from the discretization of (5) by applying a Helmholtz decomposition first, obtaining a potential formulation with a Coulomb gauge condition. This change of variables (used also in [3, 12, 22, 28], among many others) splits the electric field into components in the active and the null spaces of the **curl** operator. A further reformulation, reminiscent of the pressure-Poisson equation for the incompressible Navier-Stokes equations [14, 31], yields a system of strongly elliptic, weakly coupled PDEs, for which more standard preconditioned Krylov space methods are directly applicable.

In [15], we further addressed possible significant jumps in the conductivity while μ is assumed constant, by employing a finite volume discretization on a staggered grid, akin to Yee's method with the locations of $\boldsymbol{E}$- and $\boldsymbol{H}$-components exchanged, as in Fig. 2. The normal components of $\boldsymbol{E}$ are now double-valued, but this is taken care of in an elegant way by the Helmholtz decomposition of $\boldsymbol{E}$ and by introducing the (generalized) current

$$\widehat{\boldsymbol{J}} = \widehat{\sigma}\boldsymbol{E} = (\sigma - \imath\omega\epsilon)\boldsymbol{E} \tag{6}$$

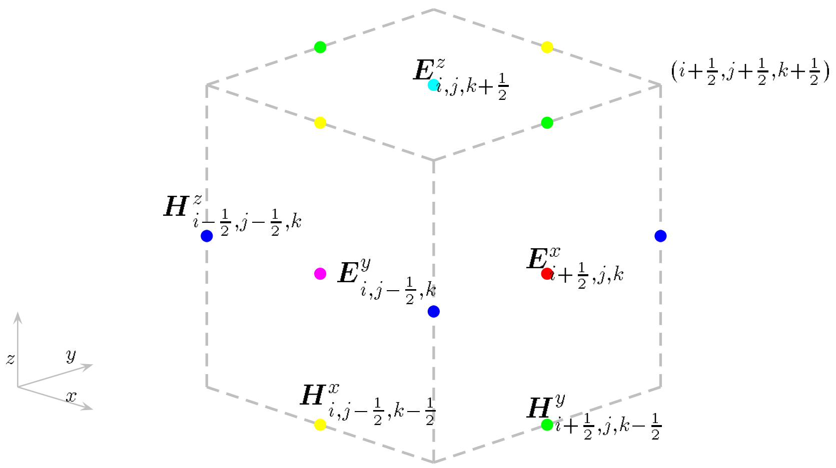

Figure 2: A staggered discretization of $\boldsymbol{E}$ and $\boldsymbol{H}$ in three dimensions: $\boldsymbol{E}$-components on faces, $\boldsymbol{H}$-components on edges.

into the equations. The **curl** operators in (5) are replaced by the vector Laplacian according to the vector identity

$$\nabla \times \nabla \times = -\nabla^2 I + \nabla\nabla\cdot \tag{7}$$

for sufficiently smooth vector functions (not $\boldsymbol{E}$).

In this paper we generalize our approach from [1, 15] to the case where the magnetic permeability μ may be highly discontinuous as well. This is a realistic case of interest in geophysical applications, although usually the jump in conductivity dominates the jump in permeability. Now the roles of $\boldsymbol{E}$ and $\boldsymbol{H}$ are essentially dual, and it is possible to apply a Helmholtz decomposition to either $\boldsymbol{E}$ or $\boldsymbol{H}$, keeping the other unknown vector function intact. We choose to decompose the electric field $\boldsymbol{E}$, referring to Fig. 2 for the locations of the $\boldsymbol{H}$-unknowns in the ensuing discretization. The major departure from our previous work is in the fact that the identity (7) does not directly extend for the operator $\nabla \times (\mu^{-1}\nabla\times)$ appearing in (5). We can, however, stabilize this operator by subtracting $\nabla(\mu^{-1}\nabla\cdot)$ (see, e.g., [20]), and this forms the basis for our proposed method. In cases of constant magnetic permeability or electric conductivity the formulation can be reduced to our previous formulation in [15] or a variant thereof.

Our approach to dealing with possible discontinuities can be viewed as using fluxes (which are continuous on cell faces) and vorticities (which are continuous on cell edges). The introduction of such unknowns is strongly connected to mixed finite elements which are used for highly discontinuous problems [7, 8, 16].

The paper is laid out as follows. In Section 2, we reformulate Maxwell's equations in a way which enables us to extend our methods. The resulting system is amenable to discretization using a finite-volume technique described in Section 3.

The extension and generalization of our approach from [1] through [15] to the present article is not without a price. This price is an added complication in the sparsity structure of the resulting discrete system and a corresponding added cost in the iterative solution of such systems. We briefly describe the application of Krylov space methods to solve the system of algebraic equations in Section 4. We use BICGSTAB (see, e.g., [30]) together with one of two preconditioners: an incomplete block LU-decomposition (which is a powerful preconditioner in the case of diagonally dominant linear systems) and SSOR. The system's diagonal dominance is a direct consequence of our analytic formulation.

We present the results of numerical experiments in Section 5 and compare results obtained using our method with those obtained using a more traditional Yee discretization. If the source is not divergence-free, as is the case for electric (but not magnetic) sources, then our method is better by more than two orders of magnitude. The method works well also for a case where the problem coefficients vary rapidly. We conclude with a short summary and further remarks.

2 Reformulation of Maxwell's Equations

Maxwell's equations (1a) and (1b) can be viewed as flux balance equations, i.e. each term describes the flux which arises from a different physical consideration, and the equations are driven by the conservation of fluxes. (In fact, this was how they were originally developed [23].) Therefore, in both (1a) and (1b) we have a flux term which should be balanced. In (1b) the *generalized current density* $\widehat{\boldsymbol{J}}$ defined in (6) is balanced with the source and the flux which arise from magnetic fields, and in (1a) the *magnetic flux*

$$\boldsymbol{B} = \mu \boldsymbol{H} \tag{8}$$

is balanced with the flux which arises from electric fields. In our context these fluxes are well-defined on cell faces but they may be multi-valued at cell edges.[1]

Furthermore, the leading differential operator in (5), say, has a nontrivial null space. Rather than devising iterative methods which directly take this into account (as, e.g., in [2, 16]), we transform the equations before discretization.

We decompose $\boldsymbol{E}$ into its components in the active space and in the null space of the **curl** operator:

$$\boldsymbol{E} = \boldsymbol{A} + \boldsymbol{\nabla}\phi, \tag{9a}$$
$$\nabla \cdot \boldsymbol{A} = 0. \tag{9b}$$

We could decompose $\boldsymbol{H}$ instead in a similar way, but we would not decompose both. Here we have chosen to concentrate on the decomposition of $\boldsymbol{E}$. Substituting (9) into Maxwell's equations we obtain

$$\boldsymbol{\nabla} \times \boldsymbol{A} - \imath\omega\mu\boldsymbol{H} = \boldsymbol{0}, \tag{10a}$$
$$\boldsymbol{\nabla} \times \boldsymbol{H} - \widehat{\sigma}(\boldsymbol{A} + \boldsymbol{\nabla}\phi) = \boldsymbol{J^s}, \tag{10b}$$
$$\nabla \cdot \boldsymbol{A} = 0. \tag{10c}$$

Furthermore, (5) becomes

$$\boldsymbol{\nabla} \times (\mu^{-1}\boldsymbol{\nabla} \times \boldsymbol{A}) - \imath\omega\widehat{\sigma}(\boldsymbol{A} + \boldsymbol{\nabla}\phi) = \imath\omega\boldsymbol{J^s}, \tag{11a}$$
$$\nabla \cdot \boldsymbol{A} = 0. \tag{11b}$$

Note that across an interface between distinct conducting media we have, in addition to (4),

$$\boldsymbol{n} \times (\boldsymbol{A}_1 - \boldsymbol{A}_2) = 0, \tag{12a}$$
$$\boldsymbol{n} \cdot (\boldsymbol{A}_1 - \boldsymbol{A}_2) = 0, \tag{12b}$$
$$\boldsymbol{n} \cdot (\epsilon_1\boldsymbol{\nabla}\phi_1 - \epsilon_2\boldsymbol{\nabla}\phi_2) = \rho_s, \tag{12c}$$
$$\boldsymbol{n} \cdot (\widehat{\boldsymbol{J}}_1 - \widehat{\boldsymbol{J}}_2) = 0, \tag{12d}$$

where ρ_s in (12c) is an electric surface charge density. These conditions and the differential equations (1) imply that while $\widehat{\boldsymbol{J}} \cdot \boldsymbol{n}$ is continuous, $\boldsymbol{E} \cdot \boldsymbol{n}$ is

[1]Under appropriate conditions, $\boldsymbol{B}, \widehat{\boldsymbol{J}} \in H(div; \Omega)$, whereas $\boldsymbol{E}, \boldsymbol{H} \in H(\mathbf{curl}; \Omega)$; see, e.g., [13].

not. Moreover, $\boldsymbol{\nabla}\phi$ inherits the discontinuity of $\boldsymbol{E}\cdot\boldsymbol{n}$, while $\boldsymbol{A}$ is continuous, and both $\nabla\cdot\ \boldsymbol{A}$ and $\boldsymbol{\nabla}\times\ \boldsymbol{A}$ are bounded (cf. [13]).

In [15] we had the relation $\boldsymbol{\nabla}\times\ \boldsymbol{\nabla}\times\ \boldsymbol{A} = -\nabla^2\boldsymbol{A}$ holding. However, when μ varies, the identity (7) does not extend directly, and we must deal with the null space of $\boldsymbol{\nabla}\times\ \boldsymbol{A}$ in a different way.

Let us define the Sobolev spaces

$$\boldsymbol{W}(\Omega) = \{\boldsymbol{v}\in[L^2(\Omega)]^3;\ \boldsymbol{\nabla}\times\ \boldsymbol{v}\in[L^2(\Omega)]^3, \nabla\cdot\ \boldsymbol{v}\in L^2(\Omega)\} \tag{13a}$$

equipped with the usual norm (see, e.g., [13])

$$\|\boldsymbol{v}\|_{\boldsymbol{W}(\Omega)} = \{\|\boldsymbol{v}\|^2 + \|\boldsymbol{\nabla}\times\ \boldsymbol{v}\|^2 + \|\nabla\cdot\ \boldsymbol{v}\|^2\}^{1/2} \tag{13b}$$

(the $L^2(\Omega)$- and $[L^2(\Omega)]^3$- norms are used on the right hand side of (13b)), and

$$\boldsymbol{W}_0(\Omega) = \{\boldsymbol{v}\in\boldsymbol{W}(\Omega);\ \boldsymbol{v}\cdot\boldsymbol{n}\Big|_{\partial\Omega} = 0\}. \tag{13c}$$

Green's formula yields, for any $\boldsymbol{u}\in\boldsymbol{W}(\Omega)$, $\boldsymbol{v}\in\boldsymbol{W}_0(\Omega)$,

$$\begin{aligned}&(\boldsymbol{\nabla}\times\ (\mu^{-1}\boldsymbol{\nabla}\times\ \boldsymbol{u}),\boldsymbol{v}) - (\boldsymbol{\nabla}(\mu^{-1}\nabla\cdot\ \boldsymbol{u}),\boldsymbol{v}) =\\ &\mu^{-1}[(\boldsymbol{\nabla}\times\ \boldsymbol{u},\boldsymbol{\nabla}\times\ \boldsymbol{v}) + (\nabla\cdot\ \boldsymbol{u},\nabla\cdot\ \boldsymbol{v})],\end{aligned}$$

where the usual notation for inner product in $L^2(\Omega)$ and $[L^2(\Omega)]^3$ is used. Thus, for any $\boldsymbol{u}\in\boldsymbol{W}_0(\Omega)$, if $\boldsymbol{u}\neq\boldsymbol{0}$ then

$$(\boldsymbol{\nabla}\times\ (\mu^{-1}\boldsymbol{\nabla}\times\ \boldsymbol{u}),\boldsymbol{u}) - (\boldsymbol{\nabla}(\mu^{-1}\nabla\cdot\ \boldsymbol{u}),\boldsymbol{u}) > 0.$$

We may therefore stabilize (11a) by subtracting a vanishing term:

$$\boldsymbol{\nabla}\times\ (\mu^{-1}\boldsymbol{\nabla}\times\ \boldsymbol{A}) - \boldsymbol{\nabla}(\mu^{-1}\nabla\cdot\ \boldsymbol{A}) - \imath\omega\widehat{\sigma}(\boldsymbol{A}+\boldsymbol{\nabla}\phi) = \imath\omega\boldsymbol{J^s}, \tag{14}$$

obtaining a strongly elliptic operator for $\boldsymbol{A}$, provided $\boldsymbol{A}\in\boldsymbol{W}_0(\Omega)$. The latter condition is guaranteed by the choice (17b) below.

A similar stabilization (or penalty) approach was studied with mixed success in [10, 20]. However, our experience and thus our recommendation are more positive because of the discretization we utilize. We elaborate further upon this point in the next section.

Using (10c), we can write (10b) as

$$\begin{aligned}\boldsymbol{\nabla}\times\boldsymbol{H} - \boldsymbol{\nabla}\psi - \widehat{\boldsymbol{J}} &= \boldsymbol{J^s} &(15a)\\ \psi - \mu^{-1}\nabla\cdot\ \boldsymbol{A} &= 0. &(15b)\end{aligned}$$

This may be advantageous in the case of discontinuous μ, similarly to the mixed formulation used for the simple *div-grad* system

$$\nabla \cdot (\sigma \boldsymbol{\nabla} \phi) = q$$

in [7, 27, 8, 11] and elsewhere.

Our final step in the reformulation is to replace, as in [15], the gauge condition (10c) on $\boldsymbol{A}$ by an indirect one, obtained upon taking $\nabla\cdot$ of (15a) and simplifying using (15b) and (10c). This achieves only a weak coupling in the resulting PDE system. We note that replacing of the gauge condition (10c) is similar to the pressure Poison equation in CFD [31, 14]. The complete system, to be discretized in the next section, can now be written as

$$\begin{aligned}
\boldsymbol{\nabla} \times \boldsymbol{H} - \boldsymbol{\nabla} \psi - \widehat{\boldsymbol{J}} &= \boldsymbol{J^s}, && (16a)\\
\imath\omega\mu \boldsymbol{H} - \boldsymbol{\nabla} \times \boldsymbol{A} &= \mathbf{0}, && (16b)\\
\mu\psi - \nabla \cdot \boldsymbol{A} &= 0 && (16c)\\
-\nabla \cdot \widehat{\boldsymbol{J}} &= \nabla \cdot \boldsymbol{J^s}, && (16d)\\
\widehat{\boldsymbol{J}} - \widehat{\sigma}(\boldsymbol{A} + \boldsymbol{\nabla}\phi) &= \mathbf{0}. && (16e)
\end{aligned}$$

In order to complete the specification of this system we must add appropriate Boundary Conditions (BCs). First, we note that the original BC (3) can be written as

$$(\boldsymbol{\nabla} \times \boldsymbol{A}) \times \boldsymbol{n}\Big|_{\partial\Omega} = \mathbf{0}. \tag{17a}$$

An additional BC on the normal components of $\boldsymbol{A}$ is required for the Helmholtz decomposition (9) to be unique. Here we choose (corresponding to (13c))

$$\boldsymbol{A} \cdot \boldsymbol{n}\Big|_{\partial\Omega} = 0. \tag{17b}$$

This, together with (9), determines $\boldsymbol{A}$ for a given $\boldsymbol{E}$.

Moreover, since (16d) was obtained by taking the $\nabla\cdot$ of (16a) additional BCs are required on either $\partial\phi/\partial n$ or ϕ. For this we note that the original BC (3) together with the PDE (1b) imply also

$$\widehat{\boldsymbol{J}} \cdot \boldsymbol{n}\Big|_{\partial\Omega} = -\boldsymbol{J^s} \cdot \boldsymbol{n}\Big|_{\partial\Omega}. \tag{17c}$$

The latter relation (17c), together with (17b), implies $\frac{\partial \phi}{\partial n} = 0$ at the boundary.[2][3]

The above conditions determine ϕ up to a constant. We pin this constant down arbitrarily, e.g. by requiring

$$\int_\Omega \phi \, dV = 0. \tag{17d}$$

Finally, we note that (17c) together with (16a) and (3), imply in turn that $\frac{\partial \psi}{\partial n} = 0$ at the boundary. Since (16a) and (16d) imply that

$$\nabla^2 \psi = 0$$

we obtain $\psi \equiv 0$, and thus retrieve (10c), by pinning ψ down to 0 at one additional point.

The system (16) subject to the boundary conditions (17) and ψ pinned at one point is now well-posed.

3 Deriving a discretization

As in [15], we employ a finite volume technique on a staggered grid, where $\widehat{\boldsymbol{J}}$ and $\boldsymbol{A}$ are located at the cell's faces, $\boldsymbol{H}$ is at the cell's edges, and ϕ and ψ are located at the cell's center. Correspondingly, the discretization of (16d) and (16c) are centered at cell centers, those of (16e) and (16a) are centered at cell faces, and that of (16b) is centered at cell edges. The variables distribution over the grid cell is summarized in Table 1.

To approximate $\nabla \cdot \boldsymbol{u}$ for $\boldsymbol{u} = (u^x, u^y, u^z)^T$ over a grid cell $e_{i,j,k}$ we integrate at first over the cell using the Gauss divergence theorem

$$\frac{1}{|e_{i,j,k}|} \int_{e_{i,j,k}} \nabla \cdot \boldsymbol{u} \, dV = \frac{1}{|e_{i,j,k}|} \int_{\partial e_{i,j,k}} \boldsymbol{u} \cdot \boldsymbol{n} \, dS,$$

[2]In cases where the original BC is different from (3) we still use the boundary condition $\partial\phi/\partial n = 0$ as an asymptotic value for an infinite domain. Alternatively, note the possibility of applying the Helmholtz decomposition to $\boldsymbol{H}$, although generally there are good practical reasons to prefer the decomposition (9) of $\boldsymbol{E}$.

[3]In our geophysical applications $\boldsymbol{J^s} \cdot \boldsymbol{n}$ vanishes at the boundary.

${A^x}_{i+\frac{1}{2},j,k} \approx A^x(x_{i+\frac{1}{2}}, y_j, z_k)$	$\hat{J}^x_{i+\frac{1}{2},j,k} \approx \hat{J}^x(x_{i+\frac{1}{2}}, y_j, z_k)$
${A^y}_{i,j+\frac{1}{2},k} \approx A^y(x_i, y_{j+\frac{1}{2}}, z_k)$	$\hat{J}^y_{i,j+\frac{1}{2},k} \approx \hat{J}^y(x_i, y_{j+\frac{1}{2}}, z_k)$
${A^z}_{i,j,k+\frac{1}{2}} \approx A^z(x_i, y_j, z_{k+\frac{1}{2}})$	$\hat{J}^z_{i,j,k+\frac{1}{2}} \approx \hat{J}^z(x_i, y_j, z_{k+\frac{1}{2}})$
$H^x_{i,j+\frac{1}{2},k+\frac{1}{2}} \approx H^x(x_i, y_{j+\frac{1}{2}}, z_{k+\frac{1}{2}})$	
$H^y_{i+\frac{1}{2},j,k+\frac{1}{2}} \approx H^y(x_{i+\frac{1}{2}}, y_j, z_{k+\frac{1}{2}})$	
$H^z_{i+\frac{1}{2},j+\frac{1}{2},k} \approx H^z(x_{i+\frac{1}{2}}, y_{j+\frac{1}{2}}, z_k)$	
$\psi_{i,j,k} \approx \psi(x_i, y_j, z_k)$	
$\phi_{i,j,k} \approx \phi(x_i, y_j, z_k)$	

Table 1: Summary of the discrete grid functions. Each scalar field is approximated by the grid functions at points slightly staggered in each cell $e_{i,j,k}$ of the grid.

and then use midpoint quadrature on each face to evaluate each component of the surface integrals appearing on the right-hand side above. Thus, define

$$\nabla_{i,j,k} \cdot \boldsymbol{u} = \frac{u^x_{i+\frac{1}{2},j,k} - u^x_{i-\frac{1}{2},j,k}}{h^x_i} + \frac{u^y_{i,j+\frac{1}{2},k} - u^y_{i,j-\frac{1}{2},k}}{h^y_j} + \frac{u^z_{i,j,k+\frac{1}{2}} - u^z_{i,j,k-\frac{1}{2}}}{h^z_k} \tag{18}$$

and express the discretization of (16d) and (16c) on each grid cell as

$$\nabla_{i,j,k} \cdot \widehat{\boldsymbol{J}} + \nabla_{i,j,k} \cdot \boldsymbol{J^s} = 0. \tag{19a}$$

$$\psi_{i,j,k} = \mu^{-1}_{i,j,k} \nabla_{i,j,k} \cdot \boldsymbol{A} \tag{19b}$$

where $\mu_{i,j,k} = \mu(x_i, y_j, z_k)$. Note that we are not assuming a uniform grid: each cell may have different widths in each direction. The boundary conditions (17) are used at the end points of (19).

Next, consider the discretization at cell faces. Following [15], we define the *harmonic average* of $\widehat{\sigma}$ between neighboring cells in the x-direction by

$$\widehat{\sigma}_{i+\frac{1}{2},j,k} = h^x_{i+\frac{1}{2}} \left(\int_{x_i}^{x_{i+1}} \widehat{\sigma}^{-1}(x,y,z)dx \right)^{-1}, \tag{20a}$$

where $h^x_{i+\frac{1}{2}} = x_{i+1} - x_i = (h^x_{i+1} + h^x_i)/2$. If $\widehat{\sigma}$ is assumed constant over each cell, this integral evaluates to

$$\widehat{\sigma}_{i+\frac{1}{2},j,k} = h^x_{i+\frac{1}{2}} \left(\frac{h^x_i}{2\widehat{\sigma}_{i,j,k}} + \frac{h^x_{i+1}}{2\widehat{\sigma}_{i+1,j,k}} \right)^{-1}. \tag{20b}$$

Then, the resulting approximation for the x-component of (16e) is [15]

$$\hat{J}^x_{i+\frac{1}{2},j,k} = \widehat{\sigma}_{i+\frac{1}{2},j,k}\left({A^x}_{i+\frac{1}{2},j,k} + \frac{\phi_{i+1,j,k}-\phi_{i,j,k}}{h^x_{i+\frac{1}{2}}}\right). \tag{20c}$$

Next, we discretize (16a) as in [37]. Writing the x-component of these equations,

$$(\partial_z H^y - \partial_y H^z) - \partial_x \psi - \hat{J}^x = s^x,$$

where we denote $\boldsymbol{J^s} = (s^x, s^y, s^z)^T$, a discretization centered at the center of the cell's x-face results:

$$\begin{aligned}
&\frac{H^y_{i+\frac{1}{2},j,k+\frac{1}{2}} - H^y_{i+\frac{1}{2},j,k-\frac{1}{2}}}{h^z_k} - \frac{H^z_{i+\frac{1}{2},j+\frac{1}{2},k} - H^z_{i+\frac{1}{2},j-\frac{1}{2},k}}{h^y_j} - \\
&\frac{\psi_{i+1,j,k}-\psi_{i,j,k}}{h^x_{i+\frac{1}{2}}} - \hat{J}^x_{i-\frac{1}{2},j,k} = s^x_{i+\frac{1}{2},j,k}.
\end{aligned} \tag{21}$$

Similar expressions to those in (20c) and (21) can be derived in the y- and z-directions. The boundary conditions (3) are used to close (21). Using (20c) we can eliminate $\widehat{\boldsymbol{J}}$ from (19a) and obtain a discrete equation in which the dominant terms all involve ϕ. The resulting stencil for ϕ has 7 points. We also apply the obvious quadrature for the single condition (17d).

Finally we discretize the edge-centered (16b). Consider, say, the x-component of (16b), written as

$$\partial_z A^y - \partial_y A^z = \imath\omega\mu H^x.$$

Integrating this equation over the surface of a rectangle with corners

$$(x_i, y_j, z_k),\ (x_i, y_{j+1}, z_k),\ (x_i, y_{j+1}, z_{k+1}),\ (x_i, y_j, z_{k+1})$$

the expression on the left hand side is integrated using the Gauss **curl** theorem, and μ on the right hand side is averaged in both directions to obtain a value on the edge. We do not divide through by μ before this integration because we wish to integrate the magnetic flux, which is potentially less smooth around an edge than the magnetic field. This yields

$$\mu_{i,j+\frac{1}{2},k+\frac{1}{2}} = [h^y_{j+\frac{1}{2}} h^z_{k+\frac{1}{2}}]^{-1}\left(\int_{y_j}^{y_{j+1}}\int_{z_k}^{z_{k+1}} \mu(x_i,y,z)\,dydz\right). \tag{22a}$$

If μ is assumed to be constant over each cell, this integral evaluates to

$$\mu_{i,j+\frac{1}{2},k+\frac{1}{2}} = \tag{22b}$$
$$[h^y_{j+\frac{1}{2}} h^z_{k+\frac{1}{2}}]^{-1} \left(\frac{\mu_{i,j,k} h^y_j h^z_k + \mu_{i,j+1,k} h^y_{j+1} h^z_k + \mu_{i,j+1,k+1} h^y_{j+1} h^z_{k+1} + \mu_{i,j,k+1} h^y_j h^z_{k+1}}{4} \right).$$

Then, the resulting approximation for the x-component of (16b) is

$$H^x_{i,j+\frac{1}{2},k+\frac{1}{2}} = (\imath\omega\mu_{i,j+\frac{1}{2},k+\frac{1}{2}})^{-1} \left(\frac{A^y_{i,j+\frac{1}{2},k+1} - A^y_{i,j+\frac{1}{2},k}}{h^z_{k+\frac{1}{2}}} - \frac{A^z_{i,j+1,k+\frac{1}{2}} - A^z_{i,j,k+\frac{1}{2}}}{h^y_{j+\frac{1}{2}}} \right) \tag{22c}$$

Using (19b) as well as (22c) and similar expressions derived in the y- and z-directions, we substitute for $\boldsymbol{H}$ and ψ in (21) and obtain a discrete system of equations for $\boldsymbol{A}$. The resulting stencil for $\boldsymbol{A}$ has 19 points and the same structure as for the discretization of the operator

$$\boldsymbol{\nabla} \times (\mu^{-1} \boldsymbol{\nabla} \times) - \boldsymbol{\nabla}(\mu^{-1} \nabla \cdot).$$

The difference between this discretization and a direct discretization of the latter is that μ at the interface is naturally defined as an arithmetic average and not a harmonic average.

The discretization described above can be viewed as a careful extension of Yee's method, suitable for discontinuous coefficients and amenable to fast iterative solution methods. It is centered, conservative and second order accurate. Note that throughout we have used a consistent, compact discretization of the operators $\nabla\cdot$, $\boldsymbol{\nabla}\times$ and $\boldsymbol{\nabla}$. We can denote the corresponding discrete operators by $\nabla_h\cdot$, $\boldsymbol{\nabla}_h\times$ and $\boldsymbol{\nabla}_h$ and immediately obtain the following identities (cf. [18, 17]),

$$(\boldsymbol{\nabla}_h \times)\boldsymbol{\nabla}_h = \boldsymbol{0}, \tag{23a}$$
$$(\nabla_h \cdot)(\boldsymbol{\nabla}_h \times) = 0, \tag{23b}$$
$$\boldsymbol{\nabla}_h(\nabla_h \cdot) - (\boldsymbol{\nabla}_h \times)(\boldsymbol{\nabla}_h \times) = (\nabla_h \cdot)\boldsymbol{\nabla}_h. \tag{23c}$$

These are of course analogues of vector calculus identities which hold for sufficiently differentiable vector functions. The boundary conditions (17) are discretized using these discrete operators as well.

Next, note that upon applying $\nabla_h\cdot$ to (21) and using (23b) and (19a) we obtain

$$\nabla_h \cdot \boldsymbol{\nabla}_h \psi = 0.$$

Moreover, from (21) and (17c), the discrete normal derivative of ψ vanishes at the boundaries as well. Setting ψ to 0 arbitrarily at one point,

$$\psi_{1,1,1} = 0$$

then determines that $\psi \equiv 0$ throughout the domain (as a grid function). We obtain another conservation property of our scheme, namely, a discrete divergence-free $\boldsymbol{A}$:

$$\nabla_h \cdot \boldsymbol{A} = 0. \tag{24}$$

Recall the stabilizing term added in (14). For the exact solution this term obviously vanishes. Now, (24) assures us that the corresponding discretized term vanishes as well (justifying use of the term 'stabilization', rather than 'penalty'). This is not the case for the nodal finite element method which was considered in [20, 10]. For an approximate solution that does not satisfy (24) the stabilization term may grow in size when μ varies over a few orders of magnitude, or else $\nabla_h \cdot \boldsymbol{A}$ grows undesirably in an attempt to keep $\mu^{-1}\nabla_h \cdot \boldsymbol{A}$ approximately constant across an interface [20].

Our particular way of averaging across discontinuities, namely, arithmetic averaging of μ at cell edges and harmonic averaging of $\widehat{\sigma}$ at cell faces, can be important. The averaging can be viewed as a careful approximation of the constitutive relationship for discontinuous coefficients. To show that, we look first at the relation

$$\hat{J}^x = \widehat{\sigma} E^x$$

across a face whose normal direction is x. This flux flows in series, and therefore an approximate $\widehat{\sigma}$ that represents the bulk property of the flow through the volume is given by the harmonic average (corresponding to an arithmetic average of the resistivities). Next, we look at the relation

$$B^x = \mu H^x$$

where μ is an edge variable and B^x is the flux through four cells which share that edge. Here the flow is in parallel, which implies that we need to approximate μ on the edge by an arithmetic average.

Note also that if we use the more common implementation of Yee's method (i.e. $\boldsymbol{H}$ on the cell's faces and $\boldsymbol{E}$ on the cell's edges) then the roles of $\widehat{\sigma}$ and μ interchange and we need to average μ harmonically and $\widehat{\sigma}$ arithmetically.

4 Solution of the discrete system

After the elimination of $\boldsymbol{H}$, ψ and $\widehat{\boldsymbol{J}}$ from the discrete equations we obtain a large, sparse system for the grid functions corresponding to $\boldsymbol{A}$ and ϕ:

$$K\ u\ = \begin{pmatrix} H_1 - \imath\omega S & \imath\omega S D^T \\ -DS & H_2 \end{pmatrix} \begin{pmatrix} \boldsymbol{A} \\ \phi \end{pmatrix} = \begin{pmatrix} b_A \\ b_\phi \end{pmatrix} = b. \tag{25}$$

Here

$$H_1 = C^T M C + D^T M_c D$$

is the result of the discretization of the operator $\boldsymbol{\nabla}\times\mu^{-1}\boldsymbol{\nabla}\times\ -\boldsymbol{\nabla}\mu^{-1}\nabla\cdot$; C corresponds to the discretization of the operator $\boldsymbol{\nabla}\times$; D likewise corresponds to the discretization of the operator $\nabla\cdot$; the diagonal matrix S results from the discretization of the operator $\widehat{\sigma}(\cdot)$; M and M_c similarly arise from the operator $\mu^{-1}(\cdot)$ at cell edges and at cell centers respectively; and $H_2 = DSD^T$ represents the discretization of $\nabla\cdot(\widehat{\sigma}\boldsymbol{\nabla}(\cdot))$. In regions of constant μ the matrix H_1 simplifies into a discretization of the vector Laplacian. The blocks are weakly coupled through interfaces in μ. A typical sparsity structure of H_1 for variable μ, as contrasted with constant μ, is displayed in Fig. 3. The structure of the obtained system is similar to that in [15], although the main block diagonal is somewhat less pleasant.

Note that as long as the frequency is low enough that the diffusion number satisfies

$$d = \omega\mu\sigma h \ll 1$$

(where h is the maximum grid spacing) the matrix is dominated by the diagonal blocks. This allows us to develop a block preconditioner good for low frequencies ω based on the truncated ILU (ILU(t)) decomposition of the major blocks [30]. Thus, we approximate K by the block diagonal matrix

$$\hat{K} = \begin{pmatrix} H_1 & 0 \\ 0 & H_2 \end{pmatrix} \tag{26}$$

and then use ILU(t) to obtain a sparse approximate decomposition of the matrix $\hat{K}$. This decomposition is used as a preconditioner for the Krylov solver. Note that although K is complex the approximation $\hat{K}$ is real and therefore we need to apply the ILU decomposition only to two real matrices, which saves memory.

Note that the block approximation (26) makes sense for our discretization but not for the direct staggered-grid discretization of (1). Thus, the

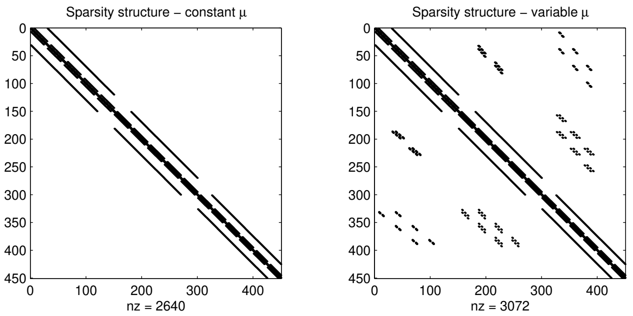

Figure 3: Sparsity structure of the matrix H_1 corresponding to variable μ (right) and to constant μ (left).

reformulation and subsequent discretization of Maxwell's equations allow us to easily obtain a good block preconditioner in a modular manner, while the discretization of (1) does not.

5 Numerical examples

Our goal in this section is to demonstrate the efficacy of our proposed method and to compare it to the more common discretization of the system (1), utilized e.g. in [24, 29], using standard Krylov-type methods and preconditioners for solving the resulting discrete systems. We vary the type of source used, the size of jumps in the coefficients σ and μ, the preconditioner and the grid size.

In the tables below, 'iterations' denotes the number of BICGSTAB iterations needed for achieving a relative accuracy of 10^{-7}; 'operations' denotes the number of giga-flops required; the SSOR parameter (when used) equals 1; and the ILU threshold (when used) equals 10^{-2}. The latter threshold is such that in our current MATLAB implementation iterations involving these two preconditioners cost roughly the same.

5.1 Example Set 1

We derive the following set of experiments. Let the air's permeability be $\mu_0 = 4\pi \cdot 10^{-7}$ H/m and its permittivity $\epsilon_0 = 8.85 \cdot 10^{-12}$ F/m. We assume a cube of constant conductivity σ_c and permeability μ_c embedded in an otherwise homogeneous earth with conductivity $\sigma_e = 10^{-3}$ S/m and permeability $\mu_e = \mu_0$; see Fig. 4. In a typical geophysical application, the conductivity may range over four orders of magnitude and more, whereas the permeability rarely changes over more than one order of magnitude. Therefore, we experiment with conductivity σ_c ranging from 10^{-2} S/m to 10^3 S/m and permeability μ_c ranging from μ_0 to $1000\mu_0$.

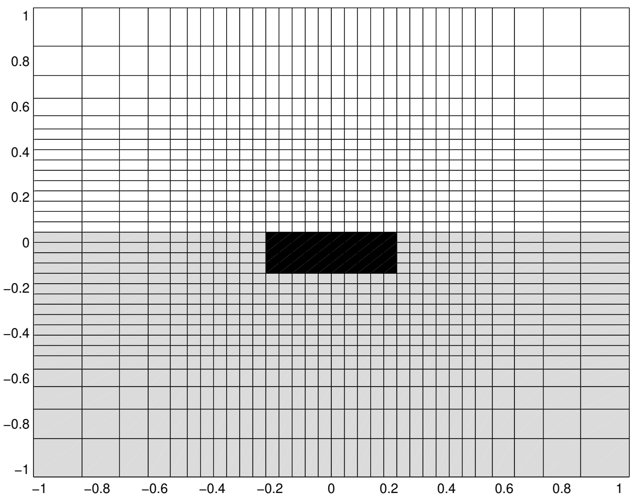

Figure 4: The setting of our first numerical experiment. A cube of conductivity σ_c and permeability μ_c is embedded inside a homogeneous earth with conductivity σ_e and permeability μ_e. Also, in the air $\widehat{\sigma} = -\imath\omega\epsilon_0$ and $\mu = \mu_0$.

We experiment with two different sources: (i) a magnetic source (a plane wave); and (ii) an electric dipole source in the x direction centered at $(0, 0, 0)$. The fact that the first source is magnetic implies that it is divergence-free. This source lies entirely in the active space of the **curl** operator. In contrast, the electric source is not divergence-free.

Both sources are assumed to oscillate with different frequencies ranging from 1 to 10^6 Hz. The solution is obtained, unless otherwise noted, on a nonuniform tensor grid (see Fig. 4) consisting of 32^3 cells. There are 95232 (complex) $\boldsymbol{E}$ unknowns corresponding to this grid and 128000 (complex) $\boldsymbol{A}, \phi$ unknowns. We then solve the system using the method described in the previous section.

5.1.1 Example 1a.

In order to be able to compare the resulting linear algebra solver with that corresponding to Yee's method we discretize the system

$$\nabla \times [\mu^{-1} \nabla \times (\widehat{\boldsymbol{J}}/\widehat{\sigma})] - \imath\omega \widehat{\boldsymbol{J}} = \imath\omega \boldsymbol{J}^s \tag{27}$$

using the staggered grid depicted in Fig. 2, i.e., where $\widehat{\boldsymbol{J}}$ is on cells' faces, which is similar to the discretization in [24]. This yields the discrete system

$$(C^T M C - \imath\omega S)e = \hat{b}$$

for the unknowns vector e corresponding to grid values of $\widehat{\boldsymbol{J}}/\widehat{\sigma}$, where the matrices C, M and S are defined in Section 4 and $\hat{b}$ depends on the source. In order to solve this system as well as ours we use BICGSTAB and an SSOR preconditioner. The comparison between the methods for the case $\sigma_c = 10^3\sigma_e$, $\mu_c = 100\mu_0$, and different frequencies is summarized in Table 2.

	Electric source				Magnetic source			
ω	iterations		operations		iterations		operations	
	$\boldsymbol{A}, \phi$	$\boldsymbol{E}$	$\boldsymbol{A}, \phi$	$\boldsymbol{E}$	$\boldsymbol{A}, \phi$	$\boldsymbol{E}$	$\boldsymbol{A}, \phi$	$\boldsymbol{E}$
10^0	46	9856	3.8	640	77	86	5.0	5.6
10^2	66	10323	4.3	670	82	91	5.3	5.9
10^4	77	10878	5.0	706	89	103	5.8	6.7
10^6	97	11212	6.3	728	99	113	6.4	7.3

Table 2: Iteration counts and computational effort for our method ($\boldsymbol{A}, \phi$) and the traditional implementation of Yee's method (applied to $\boldsymbol{E}$, or $\widehat{\boldsymbol{J}}$) for Example Set 1 using both an electric source and a magnetic source.

Table 2 shows that our method converges in a moderate number of iterations for both sources, despite the presence of significant jumps in μ and

σ. On the other hand, the more traditional discretization performs poorly for the electric source and reasonably well for the magnetic source. Slow convergence of the direct staggered discretization of Maxwell's equations in the case of an electric source was also reported in [29], where $\boldsymbol{E}$ was defined on the grid's edges.

These results clearly show an advantage of our formulation over the original Yee formulation, even for a simple preconditioning, especially for electric sources and in low frequencies. In such circumstances, the discretized first term on the left hand side of (27) strongly dominates the other term, and the residual in the course of the BICGSTAB iteration has a nontrivial component in the null space of that operator; hence its reduction is very slow. The magnetic source, on the other hand, yields a special situation where, so long as the discrete divergence and the roundoff error are relatively small, the residual component in the null space of the leading term operator is also small, hence the number of iterations using the traditional method is not much larger than when using our method.

5.1.2 Example 1b.

Next, we test the effect of discontinuities on our method. We use the electric source and record the number of BICGSTAB iterations needed for our method to converge for various values of $\tilde{\mu} = \mu_c/\mu_e$ and $\tilde{\sigma} = \sigma_c/\sigma_e$, when using block-ILU(t) preconditioning as described in the previous section. The results are summarized in Table 3.

Note that large jump discontinuities in σ do not significantly affect the rate of convergence of the iterative linear system solver for our method, but large jump discontinuities in μ have a decisive effect. Results in a similar spirit were reported in [16] regarding the effect of discontinuities in μ on a specialized multigrid method for an edge-element discretization. However, even for large discontinuities in μ the number of iterations reported in Table 3 remains relatively small compared with similar experiments reported in [29, 9]. We attribute the increase in the number of iterations as the jump in μ increases in size to a corresponding degradation in the condition number of K in (25). This degradation, however, does not depend strongly on grid size, as we verify next.

		$\tilde{\sigma}$					
ω	$\tilde{\mu}$	10^1	10^2	10^3	10^4	10^5	10^6
0	10^1	35	36	36	31	27	28
	10^2	67	55	54	52	47	51
	10^3	143	143	140	142	140	148
1	10^1	41	40	35	34	35	32
	10^2	63	51	50	53	63	51
	10^3	165	158	169	160	161	165
10	10^1	42	39	39	39	35	36
	10^2	75	59	60	61	61	60
	10^3	201	166	182	178	185	187
10^2	10^1	46	43	44	40	47	39
	10^2	71	64	66	64	63	63
	10^3	201	190	199	184	220	177
10^3	10^1	44	44	39	38	38	41
	10^2	75	69	59	61	67	69
	10^3	219	224	202	240	216	204
10^4	10^1	47	44	41	44	44	44
	10^2	85	72	80	80	65	68
	10^3	238	246	201	259	293	252
10^5	10^1	57	50	44	44	47	48
	10^2	91	79	78	82	75	81
	10^3	321	232	245	244	261	242
10^6	10^1	62	56	50	52	57	53
	10^2	105	95	89	89	91	80
	10^3	365	290	301	270	286	287

Table 3: Iteration counts for different frequencies, conductivities and permeabilities in Example Set 1. The conductivity/permeability structure is a cube in a half-space.

5.1.3 Example 1c.

In the next experiment, we use the cube model with the electric source to evaluate the influence of the grid on the number of iterations. We fix $\omega = 10^2$ and test our method on a case with a modest coefficient jump $\tilde{\mu} = 10$ and on a case with a large jump $\tilde{\mu} = 10^3$. A set of uniform grids in the interval $[-1, 1]^3$ is considered. For each grid we record the resulting number of iterations using both the SSOR and the block ILU preconditioners. The results of this experiment are gathered in Table 4.

We observe that the number of iterations increases as the number of un-

	$\tilde{\mu} = \tilde{\sigma} = 10$		$\tilde{\mu} = \tilde{\sigma} = 1000$	
Grid size	ILU	SSOR	ILU	SSOR
8^3	6	20	58	268
16^3	10	32	108	453
32^3	23	52	213	789
64^3	31	76	396	1302

Table 4: Iteration counts for different grids, for two sets of problem coefficients and using two preconditioners.

knowns increases. The increase appears to be roughly proportional to the number of unknowns to the power 1/3. The growth in number of iterations as a function of grid size is also roughly similar for both preconditioners, although the block ILU requires fewer iterations (about 1/4 for $\tilde{\mu} = 1000$) for each grid. However, ILU requires more memory than SSOR, which may prohibit its use for very large problems. The increase rate is also similar, as expected, for both values of $\tilde{\mu}$. Thus, the increase in number of iterations as a function of $\tilde{\mu}$ essentially does not depend on the grid size. Practically, however, this increase is substantial and may be hard to cope with for (perhaps unrealistically) large values of $\tilde{\mu}$ using the present techniques.

5.2 Example 2

In our next experiment we consider a more complicated earth structure. We employ a random model, which is used by practitioners in order to simulate stochastic earth models [21]. Two distinct value sets σ_1, μ_1; σ_2, μ_2 and a probability P are assumed for the conductivities and permeabilities: for each cell, the probability of having values σ_1, μ_1 is P and the probability of having values σ_2, μ_2 is $1-P$. This can be a particularly difficult model to work with, as the conductivity and permeability may jump anywhere from one cell to the next, not necessarily just across a well-defined manifold. A cross-section of such a model is plotted in Fig. 5. We then carry out experiments as before for frequencies ranging from 0 to 10^6 Hz.

We use the random model with different conductivities, permeabilities and frequencies and the electric source, with $P = 0.5$ on the domain $[-1, 1]^3$ (in km). We employ a uniform grid of size 44^3 and use both the block ILU and the SSOR preconditioners. The results of this experiment in terms

Figure 5: The setting of Example Set 2.

of iteration counts are summarized in Table 5. The results show that our solution method is effective even for highly varying conductivity. As before, the method deteriorates when very large variations in μ are present.

We can also see from Table 5 that the block ILU preconditioner works very well for low frequencies, but it is not very effective for high frequencies. It is easy to check that in all cases where the block ILU preconditioner fails to achieve convergence (denoted 'nc', meaning failure to achieve a residual of 10^{-7} in 800 iterations) the maximum grid spacing h satisfies $\omega\mu\sigma\ h \gg 1$. In such a case the discretization of the leading terms of the differential operator no longer yields the dominant blocks in the matrix equations (25), and therefore our block ILU preconditioner fails. Thus, for high frequency and high conductivity we require more grid points in order for this preconditioner to be effective. This is also consistent with the physics, as the skin depth [36] decreases and the attenuated wave can be simulated with fidelity only on a finer grid.

6 Summary and further remarks

In this paper we have developed a fast finite volume algorithm for the solution of Maxwell's equations with discontinuous conductivity and permeability. The major components of our approach are as follows.

- Reformulation of Maxwell's equations: The Helmholtz decomposition is applied to $\boldsymbol{E}$; then a stabilizing term is added, resulting in a strongly elliptic system; the system is written in first order form to allow flexibility

		σ_2/σ_1					
		10^2		10^4		10^6	
ω	μ_2/μ_1	ILU	SSOR	ILU	SSOR	ILU	SSOR
0	10^1	12	35	13	32	13	48
	10^2	27	85	29	69	31	83
	10^3	143	215	142	272	148	229
1	10^1	14	34	13	35	15	47
	10^2	28	80	34	81	26	81
	10^3	158	270	160	306	165	269
10	10^1	18	47	17	39	16	44
	10^2	33	90	34	94	35	105
	10^3	201	329	182	303	185	354
10^2	10^1	17	47	19	45	19	59
	10^2	37	104	35	101	34	119
	10^3	190	393	184	380	177	365
10^3	10^1	21	44	22	45	25	60
	10^2	42	141	47	113	57	136
	10^3	224	434	240	453	204	434
10^4	10^1	24	46	23	51	46	65
	10^2	60	140	41	149	nc	151
	10^3	246	503	259	508	252	487
10^5	10^1	25	58	27	54	nc	76
	10^2	66	151	72	149	nc	157
	10^3	232	581	244	588	242	591
10^6	10^1	26	62	55	62	nc	77
	10^2	56	180	nc	170	nc	157
	10^3	290	641	270	643	241	655

Table 5: Iteration counts for different frequencies, conductivities and permeabilities for Example Set 2.

in the ensuing discretization; and finally, the divergence-free Coulomb gauge condition is eliminated using differentiation and substitution, which yields a weakly coupled PDE system enabling an efficient preconditioner for the large, sparse algebraic system which results from the discretization.

- Discretization using staggered grids, respecting continuity conditions and carefully averaging material properties across discontinuities. For this discretization, the stabilizing term vanishes at the exact solution of the discrete equations, which is important for cases with large contrasts in μ.

- Solution of the resulting linear system using fast preconditioned Krylov methods.

The resulting algorithm was tested on a variety of problems. We have shown dramatic improvement over the more standard algorithm when the source is electric. Good performance was obtained even when the coefficients μ and σ were allowed to vary rapidly on the grid scale – a case which should prove challenging for multigrid methods.

The project that has motivated us is electromagnetic inverse problems in geophysical prospecting [34]. Solving the forward problem, i.e. Maxwell's equations as in the present article, is a major bottleneck for the data inversion methods – indeed it may have to be carried out dozens, if not hundreds, of times for each data inversion. Thus, extremely fast solvers of the problem discussed in our paper are needed. Based on the algorithm described here an implementation has been carried out which solves realistic instances of this forward problem in less than two minutes on a single-processor PC, enabling derivation of realistic algorithms at low cost for the inverse problem.

Acknowledgment: We wish to thank Drs. Doug Oldenburg and Dave Moulton for fruitful discussions and an anonymous referee for valuable comments on the first two sections of our exposition.

References

[1] D. Aruliah, U. Ascher, E. Haber, and D. Oldenburg. A method for the forward modelling of 3D electromagnetic quasi-static problems. Technical report, Institute of Applied Mathematics, University of British Columbia, 1998.

[2] R. Beck, P. Deuflhard, R. Hiptmair, R. Hoppe, and B. Wohlmuth. Adaptive multilevel methods for edge element discretizations of maxwell's equations. *Surveys Math. Indust.*, 1999. to appear.

[3] O. Bíró and K. Preis. On the use of the magnetic vector potential in the finite element analysis of three-dimensional eddy currents. *IEEE Transactions on Magnetics*, 25:3145–3159, 1989.

[4] A. Bossavit. Whitney forms: a class of finite elements for three-dimensional computations in electromagnetism. *IEEE Proc. A*, 135:493–500, 1988.

[5] A. Bossavit and L. Kettunen. Yee-like schemes on staggered cellular grids: a synthesis between fit and fem approaches. *COMPUMAG*, 1999. Short contribution.

[6] A. Brandt and N. Dinar. Multigrid solution to elliptic flow problems. In *Numerical Methods for PDE's*, pages 53–147. Academic Press, New York, 1979.

[7] F. Brezzi and M. Fortin. *Mixed and hybrid finite element methods.* Springer, 1991.

[8] G. Carey and T. Oden. *Finite Elements, a second course.* Prentice-Hall, 1983.

[9] M. Clemens and T. Weiland. Numerical algorithms for the FDITD and FDFD simulation of slowly varying electromagnetic feilds. *University of Technology, Darmstadt*, 1980.

[10] N. A. Demerdash and R. Wang. Theoretical and numerical difficulties in 3D vector potential methods in finite element magnetostatic computations. *IEEE Trans. on Magnetics*, MAG-26:1656–1658, 1990.

[11] J. E. Dendy. Black box multigrid. *Journal of Computational Physics*, 48:366–386, 1982.

[12] M.E. Everett and A. Schultz. Geomagnetic induction in a heterogenous sphere: Azimuthally symmetric test computations and the response of an undulating 660-km discontinuity. *J. Geophys. Res.*, 101:2765–2783, 1996.

[13] V. Girault and P.A. Raviart. *Finite Element Methods for Navier-Stokes Equations.* Springer, Berlin-Heidelberg, 1986.

[14] P. M. Gresho and R. L. Sani. On pressure boundary conditions for the incompressible Navier-Stokes equations. *Internat. J. Numer. Methods Fluids*, 7:1111–1145, 1987.

[15] E. Haber, U. Ascher, D. Aruliah, and D. Oldenburg. Fast modelling of 3D electromagnetic using potentials. Technical report, Institute of Applied Mathematics, University of British Columbia, 1999.

[16] R. Hiptmair. Multigrid method for maxwell's equations. *SIAM J. Numer. Anal*, 36:204–225, 1998.

[17] J. M. Hyman and M. Shashkov. Natural discretizations for the divergence gradient and curl on logically rectangular grids. *International Journal of Computers and Mathematics with Applications*, 33:81–104, 1997.

[18] J. M. Hyman and M. Shashkov. The orthogonal decomposition theorems for mimetic finite difference methods. *SIAM J. Numer. Anal.*, 36:788–818, 1999.

[19] J.M. Hyman and M. Shashkov. Mimetic discretizations for Maxwell's equations and equations of magnetic diffusion. Technical report, Los Alamos Nat. Lab., 1998. LA-UR-98-1032.

[20] J. Jin. *The Finite Element Method in Electromagnetics.* John Wiley and Sons, 1993.

[21] R.J. Knight and P. Tercier. Electrical conduction in sandstones: Examples of two dimensional and three dimensional percolation. *Transactions, American Geophysical Union*, 75, 1992.

[22] D.J. LaBrecque. A scalar-vector potential solution for 3D EM finite-difference modeling. In M. Oristaglio and B. Spies, editors, *International Symposium on Three-Dimensional Electromagnetics*, pages 143–149. Schlumberger-Doll Research, 1995.

[23] M. S. Longer. *Theoretical concepts in physics.* Cambridge university press, 1991.

[24] R.L. Mackie, T.R. Madden, and P.E. Wannamaker. Three dimensional magnetotelluric modeling using difference equations - theory and comparison to integral equation solutions. *Geophysics*, 58:215–226, 1993. vol 2.

[25] C. Mattiussi. An analysis of finite volume, finite element, and finite difference methods using some concepts from algebraic topology. *Journal of Computational Physics*, 133:289–309, 1997.

[26] P. Monk and E. Suli. A convergence analysis of Yee's scheme on nonuniform grids. *SIAM J. Numer. Anal*, 31:393–412, 1994.

[27] J. D. Moulton. *Numerical implementation of nodal methods*. PhD thesis, Institute of Applied Mathematics, University of British Columbia, 1996.

[28] T. Nakata, N. Takahashi, and K. Fujiwara. Solutions of 3D eddy current problems by finite elements. *in: Finite element electromagnetics and design*, pages 36–72, 1995. Elsevier.

[29] G.A. Newman and D.L. Alumbaugh. Frequency-doman modelling of airborne electromagnetic responses using staggered finite differences. *Geophys. Prospecting*, 43:1021–1042, 1995.

[30] Y. Saad. *Iterartive Methods for Sparse Linear Systems*. PWS Publishing Company, 1996.

[31] D. Sidilkover and U. Ascher. A multigrid solver for the steady state navier-stokes equations using the pressure-poisson formulation. *Comp. Appl. Math. (SBMAC)*, 14:21–35, 1995.

[32] A. Taflove. *Computational Electrodynamics: the Finite-Difference Time-Domain Method*. Artech House Publishers, 1995.

[33] Y. Tai. *Diadic Green Functions in Electromagnetics*. Elsevier, 1991.

[34] S.H. Ward and G.W. Hohmann. Electromagnetic theory for geophysical applications. *Electromagnetic Methods in Applied Geophysics*, 1:131–311, 1988. Soc. Expl. Geophys.

[35] T. Weiland. Time domain electromagnetic field computation with finite difference methods. *Int. J. Numer. Modelling*, 9:295–319, 1996.

[36] K. P. Whittall and D. W. Oldenburg. *Inversion of Magnetotelluric Data for a One Dimensional Conductivity*, volume 5. SEG monograph, 1992.

[37] K.S. Yee. Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media. *IEEE Trans. on antennas and propagation*, 14:302–307, 1966.